# Problem Set 3

Please write your solutions in the LATEX and Python templates provided. Aim for concise solutions; convoluted and obtuse descriptions might receive low marks, even when they are correct.

**Problem 3-1.** [5 points] **Hash Practice**

(a) [2 points] Insert integer keys `A = [47, 61, 36, 52, 56, 33, 92]` in order into a hash table of size 7 using the hash function $h(k) = (10k + 4) \bmod 7$. Each slot of the hash table stores a linked list of the keys hashing to that slot, with later insertions being appended to the end of the list. Draw a picture of the hash table after all keys have been inserted.

(b) [3 points] Suppose the hash function were instead $h(k) = ((10k + 4) \bmod c) \bmod 7$ for some positive integer $c$. Find the smallest value of $c$ such that no collisions occur when inserting the keys from `A`.

**Problem 3-2.** [15 points] **Dorm Hashing**

MIT wants to assign $2n$ new students to $n$ rooms, numbered $0$ to $n - 1$, in Pseudorandom Hall. Each MIT student will have an ID: a postive integer less than $u$, with $u \gg 2n$. No two students can have the same ID, but new students are allowed to choose their own IDs after the start of term.

MIT wants to find students quickly given their IDs, so will assign students to rooms by hashing their IDs to a room number. So as not to appear biased, MIT will publish a family $\mathcal{H}$ of hash functions online before the start of term (before new students choose their IDs), and then after students choose IDs, MIT will choose a rooming hash function uniformly at random from $\mathcal{H}$.

New MIT freshmen Rony Stark and Tiri Williams want to be roommates. For each hash family below, show that either:

- Rony and Tiri can choose IDs $k_1$ and $k_2$ so as to guarantee that they'll be roommates, or
- prove that no such choice is possible and compute the highest probability they could possibly achieve of being roommates.

(a) [5 points] $\mathcal{H} = \{h_{ab}(k) = (ak + b) \bmod n \mid a, b \in \{0, \ldots, n - 1\}$ and $a \neq 0\}$

(b) [5 points] $\mathcal{H} = \{h_a(k) = \left(\lfloor \frac{kn}{u} \rfloor + a\right) \bmod n \mid a \in \{0, \ldots, u - 1\}\}$

(c) [5 points] $\mathcal{H} = \{h_{ab}(k) = ((ak + b) \bmod p) \bmod n \mid a, b \in \{0, \ldots, p - 1\}$ and $a \neq 0\}$
    for fixed prime $p > u$ (this is the universal hash family from Lecture 4)

**Problem 3-3.** [20 points] **The Cold is Not Bothersome Anyway**

Ice cores are long cylindrical plugs drilled out of deep glaciers, which are accumulations of snow piled on top of each other and compressed into ice. Scientists can divide an ice core into distinct slices, each representing one year of deposits. For each of the following scenerios, describe an **efficient**[1] algorithm to sort $n$ slices collected from multiple ice cores. Justify your answers.

(a) [5 points] Every ice core is given a unique ***core identifier*** for bookkeeping, which is a string of exactly $16\lceil \log_4(\sqrt{n}) \rceil$ ASCII characters.[2] Sort the slices by core identifier.

(b) [5 points] The deepest ice cores in the database are up to 800,000 years old. Sort the slices by their ***age***: the integer number of years since the slice was formed.

(c) [5 points] Variation in the amount of snowfall each year will cause a glacier to accumulate at different rates over time. Sort the slices by ***thickness***, a rational number of centimeters of the form $m/n^3$ between 0 and 4, where $m$ is an integer.

(d) [5 points] Elna of Northendelle has discovered that water has ***memory***, but is unable to quantify the memory of a given slice. Luckily, given two slices, she can distinguish which has more memory in $O(1)$ time using her "two-finger algorithm" (touching the slices with her two index fingers). Sort the slices by memory.

**Problem 3-4.** [20 points] **Pushing Paper**

Farryl Dilbin is a forklift operator at Munder Difflin paper company's central warehouse. She needs to ship exactly $r$ reams of paper to a customer. In the warehouse are $n$ boxes of paper, each one foot in width, lined up side-by-side covering an $n$-foot wall. Each box contains a known positive integer number of reams, where **no two boxes contain the same number of reams**. Let $B = (b_0, \ldots b_{n-1})$ be the number of reams per box, where box $i$ located $i$ feet from the left end of the wall contains $b_i$ reams of paper, where $b_i \neq b_j$ for all $i \neq j$. To minimize her effort, Pharryl wants to know whether there is a ***close*** pair $(b_i, b_j)$ of boxes, meaning that $|i - j| < n/10$, that will ***fulfill*** order $r$, meaning that $b_i + b_j = r$.

(a) [10 points] Given $B$ and $r$, describe an expected $O(n)$-time algorithm to determine whether $B$ contains a close pair that fulfills order $r$.

(b) [10 points] Now suppose that $r < n^2$. Describe a worst-case $O(n)$-time algorithm to determine whether $B$ contains a close pair that fulfills order $r$.

---

[1]By "efficient", we mean that asymptotically faster correct algorithms will receive more points than slower ones.

[2]You may assume a string of $k$ ASCII characters is a pointer to a contiguous sequence of $k$ bytes in memory, where each byte stores an integer from 0 to 127 inclusive representing an ASCII character.

https://en.wikipedia.org/wiki/ASCII

**Problem 3-5.** [40 points] **Anagram Archaeology**

String $A$ is an *anagram* of another string $B$ if $A$ is a permutation of the letters in $B$; for example, (`indicatory`, `dictionary`) and (`brush`, `shrub`) are pairs of words that are anagrams of each other. In this problem, all strings will be ASCII strings containing only the lowercase English letters `a` to `z`.

Given two strings $A$ and $B$, the *anagram substring count* of $B$ in $A$ is the number of contiguous substrings of $A$ that are anagrams of $B$. For example, if $A = $ `'esleastealaslatet'` and $B = $ `'tesla'`, then, of the 13 contiguous substrings in $A$ of length $|B| = 5$, exactly 3 of them are anagrams of $B$, namely (`'least'`, `'steal'`, `'slate'`), so the anagram substring count of $B$ in $A$ is 3.

(a) [12 points] Given string $A$ and a positive integer $k$, describe a data structure that can be built in $O(|A|)$ time, which will then support a single operation: given a different string $B$ with $|B| = k$, return the anagram substring count of $B$ in $A$ in $O(k)$ time.

(b) [3 points] Given string $T$ and an array of $n$ length-$k$ strings $\mathcal{S} = (S_0, \ldots, S_{n-1})$ satisfying $0 < k < |T|$, describe an $O(|T| + nk)$-time algorithm to return an array $A = (a_0, \ldots, a_{n-1})$ for which $a_i$ is the anagram substring count of $S_i$ in $T$ for all $i \in \{0, \ldots, n-1\}$.

(c) [25 points] Write a Python function `count_anagram_substrings(T, S)` that implements your algorithm from part (b). Note the built-in Python function `ord(c)` returns the ASCII integer corresponding to ASCII character `c` in $O(1)$ time. You can download a code template containing some test cases from the website.

```
1   def count_anagram_substrings(T, S):
2       '''
3       Input:   T | String
4                S | Tuple of strings S_i of equal length k < |T|
5       Output:  A | Tuple of integers a_i:
6                  | the anagram substring count of S_i in T
7       '''
8       A = []
9       ##################
10      # YOUR CODE HERE #
11      ##################
12      return tuple(A)
```