


- 1 Problem 9 in Jeff's note on counting inversions. This is also a solved problem in Kleinberg-Tardos book. This is the simpler version of the solved problem at the end of this home work.
- 2 We saw a linear time selection algorithm in class which is based on splitting the array into arrays of 5 elements each. Suppose we split the array into arrays of 7 elements each. Derive a recurrence for the running time.
- 3 Suppose we are given n points $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ in the plane. We say that a point (x_i, y_i) in the input is dominated if there is another point (x_j, y_j) such that $x_j > x_i$ and $y_j > y_i$. Describe an $O(n \log n)$ time algorithm to find all the *undominated* points in the given set of n points.
- 4 Solve some recurrences in Jeff's notes.
- 5 Suppose we have a stack of n pancakes of different sizes. We want to sort the pancakes so that the smaller pancakes are on top of the larger pancakes. The only operation we can perform is a *flip* - insert a spatula under the top k pancakes, for some k between 1 and n , and flip them all over.



- 5.A. Describe an algorithm to sort an arbitrary stack of n pancakes and give a bound on the number of flips that the algorithm makes. Assume that the pancake information is given to you in the form of an n element array A . $A[i]$ is a number between 1 and n and $A[i] = j$ means that the j 'th smallest pancake is in position i from the bottom; in other words $A[1]$ is the size of the bottom most pancake (relative to the others) and $A[n]$ is the size of the top pancake. Assume you have the operation $\text{Flip}(k)$ which will flip the top k pancakes. Note that you are only interested in minimizing the number of flips.
- 5.B. Suppose one side of each pancake is burned. Describe an algorithm that sorts the pancakes with the additional condition that the burned side of each pancake is on the bottom. Again, give a bound on the number of flips. In addition to A , assume that you have an array B that gives information on which side of the pancakes are burned; $B[i] = 0$ means that the bottom side of the pancake at the i 'th position is burned and $B[i] = 1$ means the top side is burned. For simplicity, assume that whenever $\text{Flip}(k)$ is done on A , the array B is automatically updated to reflect the information on the current pancakes in A .

No proof of correctness necessary.

- 6 Suppose you are given k sorted arrays A_1, A_2, \dots, A_k each of which has n numbers. Assume that all numbers in the arrays are distinct. You would like to merge them into single sorted array A of

kn elements. Recall that you can merge two sorted arrays of sizes n_1 and n_2 into a sorted array in $O(n_1 + n_2)$ time.


- 6.A. Use a divide and conquer strategy to merge the sorted arrays in $O(n \log k)$ time. To prove the correctness of the algorithm you can assume a routine to merge two sorted arrays.
- 6.B. In MergeSort we split the array of size N into two arrays each of size $N/2$, recursively sort them and merge the two sorted arrays. Suppose we instead split the array of size N into k arrays of size N/k each and use the merging algorithm in the preceding step to combine them into a sorted array. Describe the algorithm formally and analyze its running time via a recurrence.

7 7.A. Describe an algorithm to determine in $O(n)$ time whether an arbitrary array $A[1..n]$ contains more than $n/4$ copies of any value.

7.B. Describe and analyze an algorithm to determine, given an arbitrary array $A[1..n]$ and an integer k , whether A contains more than k copies of any value. Express the running time of your algorithm as a function of both n and k .

Do not use hashing, or radix sort, or any other method that depends on the precise input values.

8 Suppose we are given two sets of n points, one set $\{p_1, p_2, \dots, p_n\}$ on the line $y = 0$ and the other set $\{q_1, q_2, \dots, q_n\}$ on the line $y = 1$. Consider the n line segments connecting each point p_i to the corresponding point q_i . Describe and analyze a divide-and-conquer algorithm to determine how many pairs of these line segments intersect, in $O(n \log n)$ time. See the example below.



Seven segments with endpoints on parallel lines, with 11 intersecting pairs.

Your input consists of two arrays $P[1..n]$ and $Q[1..n]$ of x -coordinates; you may assume that all $2n$ of these numbers are distinct. No proof of correctness is necessary, but you should justify the running time.

Solution: We begin by sorting the array $P[1..n]$ and permuting the array $Q[1..n]$ to maintain correspondence between endpoints, in $O(n \log n)$ time. Then for any indices $i < j$, segments i and j intersect if and only if $Q[i] > Q[j]$. Thus, our goal is to compute the number of pairs of indices $i < j$ such that $Q[i] > Q[j]$. Such a pair is called an *inversion*. We count the number of inversions in Q using the following extension of mergesort; as a side effect, this algorithm also sorts Q . If $n < 100$, we use brute force in $O(1)$ time. Otherwise:

- Recursively count inversions in (and sort) $Q[1.. \lfloor n/2 \rfloor]$.
- Recursively count inversions in (and sort) $Q[\lfloor n/2 \rfloor + 1.. n]$.
- Count inversions $Q[i] > Q[j]$ where $i \leq \lfloor n/2 \rfloor$ and $j > \lfloor n/2 \rfloor$ as follows:
 - Color the elements in the Left half $Q[1.. \lfloor n/2 \rfloor]$ **blue**.
 - Color the elements in the Right half $Q[\lfloor n/2 \rfloor + 1.. n]$ **red**.
 - Merge $Q[1.. \lfloor n/2 \rfloor]$ and $Q[\lfloor n/2 \rfloor + 1.. n]$, maintaining their colors.
 - For each **blue** element $Q[i]$, count the number of smaller **red** elements $Q[j]$.

The last substep can be performed in $O(n)$ time using a simple for-loop:

```

COUNTREDBLUE( $A[1..n]$ ):
   $count \leftarrow 0$ 
   $total \leftarrow 0$ 
  for  $i \leftarrow 1$  to  $n$ 
    if  $A[i]$  is red
       $count \leftarrow count + 1$ 
    else
       $total \leftarrow total + count$ 
  return  $total$ 

```

In fact, we can execute the third merge-and-count step directly by modifying the MERGE algorithm, without any need for “colors”. Here changes to the standard MERGE algorithm are indicated in red.

```

MERGEANDCOUNT( $A[1..n], m$ ):
   $i \leftarrow 1$ ;  $j \leftarrow m + 1$ ;  $count \leftarrow 0$ ;  $total \leftarrow 0$ 
  for  $k \leftarrow 1$  to  $n$ 
    if  $j > n$ 
       $B[k] \leftarrow A[i]$ ;  $i \leftarrow i + 1$ ;  $total \leftarrow total + count$ 
    else if  $i > m$ 
       $B[k] \leftarrow A[j]$ ;  $j \leftarrow j + 1$ ;  $count \leftarrow count + 1$ 
    else if  $A[i] < A[j]$ 
       $B[k] \leftarrow A[i]$ ;  $i \leftarrow i + 1$ ;  $total \leftarrow total + count$ 
    else
       $B[k] \leftarrow A[j]$ ;  $j \leftarrow j + 1$ ;  $count \leftarrow count + 1$ 
  for  $k \leftarrow 1$  to  $n$ 
     $A[k] \leftarrow B[k]$ 
  return  $total$ 

```

We can further optimize this algorithm by observing that $count$ is always equal to $j - m - 1$. (Proof: Initially, $j = m + 1$ and $count = 0$, and we always increment j and $count$ together.)

```

MERGEANDCOUNT2( $A[1..n], m$ ):
   $i \leftarrow 1$ ;  $j \leftarrow m + 1$ ;  $total \leftarrow 0$ 

```

Rubric: 10 points = 2 for base case + 3 for divide (split and recurse) + 3 for conquer (merge and count) + 2 for time analysis. Max 3 points for a correct $O(n^2)$ -time algorithm. This is neither the only way to correctly describe this algorithm nor the only correct $O(n \log n)$ -time algorithm. No proof of correctness is required.

9 Consider the following restricted variant of the Tower of Hanoi puzzle. The pegs are numbered 0, 1, and 2, and your task is to move a stack of n disks from peg 1 to peg 2. However, you are forbidden to move any disk *directly* between peg 1 and peg 2; *every* move must involve peg 0.

Describe an algorithm to solve this version of the puzzle in as few moves as possible. *Exactly* how many moves does your algorithm make?

10 Consider the following cruel and unusual sorting algorithm.

```

CRUEL(A[1..n]):
  if n > 1
    CRUEL(A[1..n/2])
    CRUEL(A[n/2 + 1..n])
    UNUSUAL(A[1..n])
  
```

```

UNUSUAL(A[1..n]):
  if n = 2
    if A[1] > A[2]           the only comparison!
      swap A[1] ↔ A[2]
  else
    for i ← 1 to n/4         swap 2nd and 3rd quarters
      swap A[i + n/4] ↔ A[i + n/2]
    UNUSUAL(A[1..n/2])      // recurse on left half
    UNUSUAL(A[n/2 + 1..n])  // recurse on right half
    UNUSUAL(A[n/4 + 1..3n/4]) // recurse on middle half
  
```

Notice that the comparisons performed by the algorithm do not depend at all on the values in the input array; such a sorting algorithm is called *oblivious*. Assume for this problem that the input size n is always a power of 2.

- 10.A. Prove by induction that CRUEL correctly sorts any input array. (**Hint:** Consider an array that contains $n/4$ 1s, $n/4$ 2s, $n/4$ 3s, and $n/4$ 4s. Why is this special case enough? What does UNUSUAL actually do?)
- 10.B. Prove that CRUEL would *not* correctly sort if we removed the for-loop from UNUSUAL.
- 10.C. Prove that CRUEL would *not* correctly sort if we swapped the last two lines of UNUSUAL.
- 10.D. What is the running time of UNUSUAL? Justify your answer.
- 10.E. What is the running time of CRUEL? Justify your answer.

11 You are a visitor at a political convention (or perhaps a faculty meeting) with n delegates. Each delegate is a member of exactly one political party. It is impossible to tell which political party any delegate belongs to. In particular, you will be summarily ejected from the convention if you ask. However, you *can* determine whether any pair of delegates belong to the *same* party or not simply by introducing them to each other. Members of the same party always greet each other with smiles and friendly handshakes; members of different parties always greet each other with angry stares and insults.

- 11.A.** Suppose more than half of the delegates belong to the same political party. Describe and analyze an efficient algorithm that identifies every member of this majority party.
- 11.B.** Now suppose precisely p political parties are present and one party has a plurality: more delegates belong to that party than to any other party. Please present a procedure to pick out the people from the plurality party as parsimoniously as possible.¹ Do *not* assume that $p = O(1)$.

12 (100 PTS.) That's so Hanoi-ing. (Fall 2020)

Consider the following variants of the Towers of Hanoi. For each of variant, describe an algorithm to solve it in as few moves as possible. **Prove** that your algorithm is correct. Initially, all the n disks are on peg 1, and you need to move the disks to peg 2. In all the following variants, you are not allowed to put a bigger disk on top of a smaller disk.

- 12.A.** (30 PTS.) **Hanoi 0:** Suppose you are forbidden to move any disk directly between peg 1 and peg 2, and every move must involve (the third peg) 0. Exactly (i.e., not asymptotically) how many moves does your algorithm make as a function of n ?
- 12.B.** (30 PTS.) **Hanoi 2:** Suppose you are only allowed to move disks from peg 0 to peg 1, from peg 1 to peg 2, or from peg 2 to peg 0.
Provide an upper bound, as tight as possible, on the number of moves that your algorithm uses.
(One can derive the exact upper bound by solving the recurrence, but this is too tedious and not required here.)
- 12.C.** (40 PTS.) **Hanoi Bye Bye:** Finally consider the disappearing Tower of Hanoi puzzle where the largest remaining disk will disappear if there is nothing on top of it. The goal here is to get all the disks to disappear and be left with three empty pegs (in as few moves as possible).
Provide an upper bound, as tight as possible, on the number of moves your algorithm uses.

13 (100 PTS.) Divide and Merger (Fall 2020)

Suppose you are given k sorted arrays A_1, A_2, \dots, A_k (potentially of different sizes). Let $n_i > 0$ be the size of the i th array A_i , for $i = 1, \dots, k$, with $\sum_{i=1}^k n_i = n$. Assume that all the numbers in all the arrays are distinct. You would like to merge them into a single sorted array A of n elements.

- 13.A.** (30 PTS.) Use a divide and conquer strategy to derive an algorithm that sorts the given sorted arrays in $O(n \log k)$ time, into one big happy sorted array.

¹Describe and analyze an efficient algorithm that identifies every member of the plurality party.

- 13.B.** (30 PTS.) In **MergeSort** we split the array of size N into two arrays each of size $N/2$, recursively sort them and merge the two sorted arrays. Suppose we instead split the array of size N into k arrays of size N/k each and use the merging algorithm in the preceding step to combine them into a sorted array. Describe the algorithm formally and analyze its running time via a recurrence. You do not need to prove the correctness of the recursive algorithm.
- 13.C.** (40 PTS.) Describe an algorithm (not necessarily divide and conquer) for the settings of (??) that works in $O(n + \sum_{i=1}^k n_i \log \frac{n}{n_i})$ time. Prove the correctness of your algorithm, Note that this is potentially an improved algorithm if the n_i s are non-uniform. For example, if $n_i = n/2^i$, for $i = 1, \dots, k$, then the overall running time is linear. One can verify (but you do not need to do it – it is not immediate) that $O(\sum_{i=1}^k n_i \log \frac{n}{n_i}) = O(n \log k)$. This implies that this algorithm is a strict improvement over (??).


14 (100 PTS.) Fowl business. (Fall 2020)

You were given a kettle of n birds, which look all the same to you. To decide if two birds are of the same species, you perform the following experiment – you put the two of them in a cage together. If they are friendly to each other, then they are of the same species. Otherwise, you separate them quickly before survival of the fittest kicks in.

- 14.A.** (60 PTS.) Suppose that strictly more than half of the birds belong to the same species. Describe and analyze an efficient algorithm that identifies every bird among the n birds that belong to this dominant species.
- 14.B.** (40 PTS.) Now suppose that there are exactly p species present in your kettle of n birds. and one species has a plurality: more birds belong to that species than to any other species. Present a procedure to pick out the birds from the plurality species as efficiently as possible (i.e., minimize the number of experiments you have to do as a function of n and p). Do *not* assume that $p = O(1)$.

15 (100 PTS.) Incomparable, that's what you are. (Fall 2022)

Two points $p = (x, y)$ and $p' = (x', y')$ are **incomparable** if $x < x'$ and $y' < y$, or alternatively, $x' < x$ and $y < y'$. Thus, in the following example, the only incomparable pair is $\{b, c\}$:



- 15.A.** (50 PTS.) You are given as input two sequences of points in the plane: $S = s_1, s_2, \dots, s_n$ and $T = t_1, t_2, \dots, t_n$, such that $x(s_i) < \alpha < x(t_j)$, for all i, j , where α is some real number. Furthermore, assume that the points of S are sorted in increasing y value, and so are the points of T .
- Let $I(S, T)$ denote the number of incomparable pairs $\{s, t\}$ with $s \in S$ and $t \in T$. Describe an algorithm, as fast as possible, that computes and outputs $I(S, T)$.
- 15.B.** (50 PTS.) Given a set P of n points in the plane, let $I(P)$ denote the number of incomparable pairs of points of P . Describe a divide-and-conquer algorithm, using the algorithm in the first part, that computes and outputs $I(P)$. (The input P is not sorted in any way.) Your algorithm needs to be as fast as possible.

You can assume that all input points have unique x and y coordinates.

(Hint: Merge sort.)

16 (100 PTS.) Sort in bulk. (Fall 2022)

You are given n distinct items in an array A . You are also given a procedure **BulkSort**(B) that sorts any array B with at most k such items. Note, that you can not compare two items directly (but you can sort them, by calling **BulkSort** on an array containing only these two items).

- 16.A.** (50 PTS.) Describe a divide and conquer algorithm that sorts A using calls to **BulkSort**. Your algorithm should perform (asymptotically) as few calls to **BulkSort** as possible. What is the worst case number of calls to **BulkSort** that your algorithm performs as a function of n and k ? (We ignore the cost of other operations - such as copying or moving things around. We also assume that given two items, or copies of two items, we can decide if they are the same item.)

(Getting the optimal algorithm for this part is difficult – an answer that is optimal up to a polylog factor is acceptable here.)

- 16.B.** (50 PTS.) Describe a recursive algorithm that performs a minimum number of calls to **BulkSort** and outputs the median item in A . Specifically, given $A[1 \dots n]$ and a parameter t , it outputs the item of rank t in the sorted order of the items of A . What is the number of calls your algorithm performs to **BulkSort** as a function of n and k ?

(Stating a recurrence on the number of calls to **BulkSort**, together with a solution of this recurrence is enough – a formal proof of the correctness of the solution to the recurrence is not required since it is very painful.)

For sanity check purposes, think about how many calls your algorithm performs to **BulkSort** (for both parts) if $k = n/2$ or $k = \sqrt{n}$.