CSC373— Algorithm Design, Analysis, and Complexity — Spring 2018

Tutorial Exercise 3: Divide and Conquer

The first two questions gives you practice with divide and conquer. The third question re-addresses the minimum-sized blocking set problem from Assignment 1, but this time with a dynamic programming approach.

Work on these questions, preferably in your groups, before the next tutorial. Bring any questions you may have to the tutorial, where you will get a chance to discuss these with TAs and other student groups. (Raise your hands to get the TA's attention!)

1. Divide and . . .: Max Step. Given an array of $n \geq 2$ integers, say $[x(1), \ldots, x(n)]$, we want to find the largest step d, which is defined to be the max of $x(j)-x(i)$ over all $j > i$. For example, for $x = [22, 5, 8, 10, -3, 1]$

$$d = x(4) - x(2) = 10 - 5 = 5.$$

1a) Divide: Finish the pseudo-code LS below for computing d. Your algorithm must make essential use of a divide and conquer strategy which splits the current problem roughly in half. We realize that there are simpler ways to compute d; the point here is for you to practice divide and, well, see what you get (if not actually rule the world).

```
// For the function defined further below, the largest step d is given by,
[d] = LS(x, 1, n)
// Assume a function call can return more values than are actually assigned to variables on the left.


[ d, ____________________ ] = LS(x, a, b)
    // Input: Array x(1..n) of integers, and indicies 1 ≤ a < b ≤ n.
    //        We require n ≥ 2.
    // Output: Largest step d = x(j) - x(i), for any i, j with a ≤ i < j ≤ b
    // Feel free to add additional returned values or arrays, as needed. To do that
    // use "return [r1, r2, r3, ...]" to return more than one value in your pseudo-code.
```

1b) Runtime. Given an array x of length n, let $T(n)$ be the runtime of $LS(x, 1, n)$ in part (a). Write an equation expressing $T(n)$ in terms of $T(k)$ for some integer(s) $k < n$. Use the Master Theorem to determine the order $T(n)$. (Hint: You should be able to achieve $O(n)$, which is the same order as a simpler approach.)

2. Divide and Conquer: Power Mod. For positive integers y and b, we define the operation y mod b as: $z = y \bmod b$ if and only if $z = y - jb$ where j is the maximum integer such that $jb \leq y$. So, in particular, $0 \leq y \bmod b < b$. We are interested in computing the **power-mod**, which is defined as

$$\text{powerMod}(y, n, b) \equiv y^n \mod b, \tag{1}$$

where we are given integers $y > 0$, $n > 0$, and $b > 0$.

Suppose we can use a somewhat limited built-in mod function that runs in constant time, but it does not apply to really large input integers y. In particular, assume that for some positive constant Y, y mod b can only be computed with the built-in function when $|y| \leq Y$. In this problem we will also assume that $b^2 < Y$.

The difficulty in implementing powerMod is that it might be the case that y^n is much larger than Y and we therefore cannot directly use the built-in mod function. Instead we will need to make use of the following property of the (mathematical) mod operation:

Property 1. For any two integers $u, v > 0$, we have $(uv \bmod b) = ([(u \bmod b)(v \bmod b)] \bmod b)$.

2a. For more practice with proofs, prove Property 1. (Note that here we are using the mathematical mod function, not the restricted built-in function.)

2b. Divide: We wish to compute $a = y^n \bmod b$ for any integer $n > 0$, and y and b are as above. Use a divide and conquer approach to compute a. Your algorithm must run in $\Theta(\log(n))$ time. Take careful note of the restriction on the built-in mod function described above. Clearly explain your algorithm and show how you derived its runtime estimate. (Hint: On you first try you may obtain a $\Theta(n)$ algorithm. Use the Master Theorem to identify what you need to speed up to get $\Theta(\log(n))$.)

3. Dynamic Programming for Schedule Blocking. The following problem was on Assignment 1, but this time we approach it with dynamic programming.

The input information for this problem is the same as the interval scheduling problem considered in class. That is, suppose you are given a set of jobs, $\{J(k)\}_{k=1}^{K}$ where each job $J(k)$ must run in the (real-valued) time interval $(s_k, f_k) \subset \mathbb{R}$. Here we assume that these interval endpoints satisfy $0 \leq s_k < f_k$ where s_k and f_k are non-negative integers for all k.

We will refer to these jobs simply by their indicies $1, 2, \ldots K$. We will abuse this notation slightly and simply refer to $J(k)$ as "job k".

A subset of jobs $C \subseteq \{1, 2, \ldots K\}$ is said to be **compatible** iff for each pair $k, j \in C$, with $k \neq j$, we have $(s_k, f_k) \cap (s_j, f_j) = \emptyset$. Otherwise C is said to be incompatible.

Here we consider the first version of the blocking set considered on Asgn#1. that is,

> **B must be compatible.** A subset $B \subseteq \{1, 2, \ldots K\}$ is said to be a **blocking subset** iff for each $k \notin B$ and $1 \leq k \leq K$, it must be the case that job k is not compatible with at least one job in B. Moreover, we require that B itself is compatible.

Problem. Given input intervals of the above form $\{J(k)\}_{k=1}^{K}$, find a minimimum-sized blocking subset B.

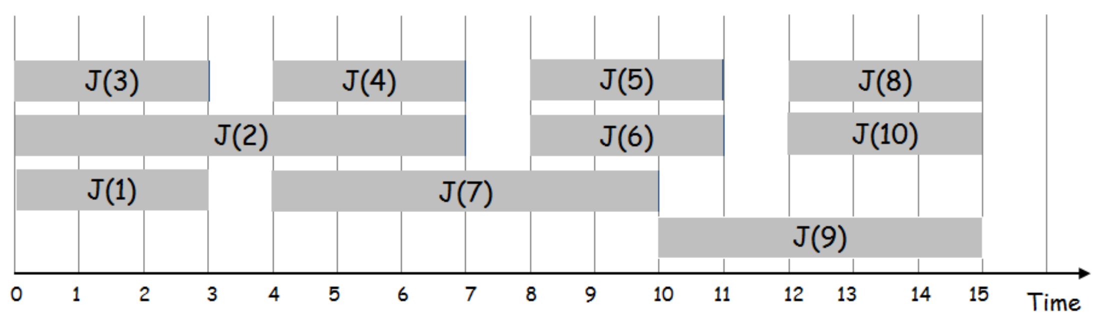

For the example above we have $K = 10$ input intervals. Note that, since we consider the jobs to be executed in open time intervals, the two jobs $J(7)$ and $J(9)$ are compatible even though one starts and the other ends at time 10. Observe that $C = \{3, 7, 10\}$ forms a blocking subset since there is no additional job that is compatible with jobs already in C. Can we find a blocking subset that has size less than $|C| = 3$? Indeed, $B = \{2, 9\}$ is such a blocking subset of size two, and this turns out to be of the minimum possible size for this example. So a possible solution for this case would be $B = \{2, 9\}$.

3a) A new instructor, Professor Field, suggests you try dynamic programming on this problem. To set this up, she suggests that you first sort the intervals by finish time. Suppose this has been done, and we have renumbered the sorted intervals, so now $f_k \leq f_j$ for all $k < j$. For each job $k \in \{1, 2, \ldots, K\}$, define the **set of immediate-precursor jobs** as

$$P(k) = \{j \mid f_j \leq s_k \text{ and there is no intervening compatible job } i > j \text{ with } f_j \leq s_i < f_i \leq s_k\}. \quad (2)$$

It is also useful to define a termination set of jobs which do not appear in any $P(k)$, say

$$T = \{j \mid 1 \leq j \leq K \text{ and, there no compatible job } i > j \text{ with } f_j \leq s_i\}. \tag{3}$$

Prof. Field suggests that it is possible to compute all the $P(k)'s$ and T in $O(dK)$ time, where d is the "depth" of the given input set. (That is, d is the maximum number of jobs in the input that need to run at any given real-valued time, t. For the example above, $d = 3$.) Clearly describe such an algorithm.

3b) Prof. Field then suggests that a dynamic programming solution may solve this problem if, working forward through the jobs in finish time, you simply keep track of one integer score, say $N(k) = |B(k)|$, per job. Moreover, here $B(k)$ is a minimum sized blocking set for the jobs $i \in \{1, 2, \ldots, k\}$ **where this blocking set $B(k)$ must contain job k**. That is, it must be the case that $k \in B(k)$. (Hint: The sets $P(k)$ defined in part (3a) are useful for marching forward in k and T may be useful as well.) Describe such an algorithm for computing these $N(k)$ for all $k \in \{1, 2, \ldots, K\}$. Give enough detail so that another group could implement it. What is the running time?

3c) Given $N(k)$ for all k, describe an algorithm for computing a minimum-sized blocking set for the original problem. What is the running time of this last step?

3d) During the forward pass through the computation of $N(k)$, at what values of k can you state that there must be an optimum solution for the whole problem that contains a particular job j, for some j that either equals k or conflicts with job k? That is, when can you ever decide that a specific paritial solution that you have after iteration k "is promising"? (For example, if we did not have job $J(7)$ in the above example then, after you process the fourth job, you know job $J(2)$ must be in any optimum solution.) To determine such points you may consider the conflict set for job k, and some of the jobs with even later start times (if any). (**Update:** It appears that in order to determine these points of certainty, it is useful to look at properties of rolling list in part (3a) near the time the k^{th} interval is pushed on the rolling list.)