

**Part A**

This problem is very similar in flavor to the segmented least squares problem. We observe that the last line ends with word  $\omega_n$  and has to start with some word  $\omega_j$ ; breaking off words  $\omega_j, \dots, \omega_n$  we are left with a recursive sub-problem on  $\omega_1, \dots, \omega_{j-1}$ .

Thus, we define  $OPT[i]$  to be the value of the optimal solution on the set of words  $W_i = \{\omega_1, \dots, \omega_i\}$ , for any  $i \leq j$ , let  $S_{i,j}$  denote the slack of a line containing the words  $\omega_i, \dots, \omega_j$ ; as a notational device, we define  $S_{i,j} = \infty$  if these words exceed total length  $L$ . For each fixed  $i$ , we can compute all  $S_{i,j}$  in  $O(n)$  time by considering values of  $j$  in increasing order; thus we can compute all  $S_{i,j}$  in  $O(n^2)$  time.

As noted above, the optimal solution must begin the last line somewhere (at word  $\omega_j$ ), and solve the sub-problem on the earlier lines optimally. We thus have the recurrence

$$OPT[n] = \min_{1 \leq j \leq n} S_{j,n}^2 + OPT[j - 1]$$

and the line of words  $\omega_j, \dots, \omega_n$  is used in an optimum solution if and only if the minimum is obtained using index  $j$ .

Finally, we just need a loop to build up all these values:

```

Compute all values  $S_{i,j}$  as described above
Set  $OPT[0]=0$ 
For  $k=1, \dots, n$ 
     $OPT[k] = \min_{1 \leq j \leq k} S_{j,k}^2 + OPT[j-1]$ 
Endfor
Return  $OPT[n]$ 

```

As noted above, it takes  $O(n^2)$  time to compute all values  $S_{i,j}$ . Each iteration of the loop takes time  $O(n)$ , and there are  $O(n)$  iterations. Thus the total running time is  $O(n^2)$ .

By tracing back through the array  $OPT$ , we can recover the optimal sequence of line breaks that achieve the value  $OPT[n]$  in  $O(n)$  additional time.

**Part B 1a**

Suppose by way of contradiction that  $T$  and  $T'$  are two distinct MCST of  $G$ . Since  $T$  and  $T'$  have the same number of edges, but are not equal, there is some edge  $e'$  in  $T'$  but not in  $T$ . If we add  $e'$  to  $T$ , we get a cycle  $C$ . Let  $e$  be the most expensive edge on this cycle. Then by the *Cycle property*,  $e$  does not belong to any MCST, contradicting the fact that it is in at least one of  $T$  or  $T'$ .

*Cycle property:* Assume that all edge costs in  $G$  are distinct. Let  $C$  be any cycle in  $G$ , let  $e$  be the most expensive edge in  $C$ . Then  $e$  does not belong to any MCST of  $G$ .

The cycle property follows directly from the exchange lemma.

### Part B 1b

Label the edges arbitrarily as  $e_1, \dots, e_m$  with the property that  $e_{m-n+1}, \dots, e_m$  belong to  $T$ . Let  $\delta$  be the minimum difference between any two non-equal edge weights; subtract  $\delta i/n^3$  from the weight of edge  $i$ . Note that all edge weights are now distinct, and the sorted order of the new weights is the same as some valid ordering of the original weights. Over all spanning trees of  $G$ ,  $T$  is the one whose total weight has been reduced by the most; thus, it is now the unique minimum spanning tree of  $G$  and will be returned by Kruskal's algorithm on this valid ordering.

### Part B 2

There are many correct ways to solve this. We will list here the main versions.

**Version 1.** Sort by “stream rate”. Rate of stream  $i$  is  $r_i = b_i/t_i$ . Send the streams in increasing order of rates. Assume for simplicity of notation that the streams are given in this order.

Check if the total rate  $\sum_{i=1}^n b_i \leq r \sum_{i=1}^n t_i$  holds. We claim the following

(1) *If this test fails than no ordering produce a feasible schedule. If the test succeeds, then we claim the above order gives a feasible schedule.*

*Proof.* If the test fails then no order can produce a feasible schedule, as the total time is  $\sum_{i=1}^n t_i$ , and we need to send  $\sum_{i=1}^n b_i$ , no matter what way we order it.

We claim that if the above sorted order sends too much for any initial time period  $[0, t]$  than the above test will also fail. To see this consider a time  $t$ , and let  $i$  be the stream sent during the last time period. If rate of stream  $i$  is at most  $r$  (i.e.,  $r_i \leq r$ ) then all streams sent so far have a rate at most  $r$ , so the total send is at most  $rt$ , contradicting the assumption that we sent too much in  $t$  time. So we must have  $r_i > r$ . However, streams are ordered in increasing rate, so in any time step after  $t$  we will also send at least  $r$  bits, and hence, the total rate at the end of all streams will also violate the “average rate at most  $r$ ” rule.

The running time is  $O(n)$ . The problem only asked to decide if an ordering exists, and that is done by testing that if the one inequality  $\sum_{i=1}^n b_i \leq r \sum_{i=1}^n t_i$  holds. If we also want to output the ordering we need to spend  $O(n \log n)$  time sorting rates. It is also okay to test if the above ordering is feasible after each job (taking  $O(n)$  time).

**Version 2.** Prove that sorting by rates is an optimal schedule by a “greedy stays ahead” argument. The key here is to state in what way the schedule is optimal, and in what way the greedy algorithm stays ahead.

(2) *In the above greedy schedule, after any time  $t$  the amount of data transmitted in the first  $t$  time steps is as low as possible.*

*Proof.* Each of the  $t_i$  seconds of the schedule transmitting the stream  $i$  will have transmission

rate of  $r_i$ . For any  $t$ , the first  $t$  seconds are the  $t$  lowest transmission rates, so the first  $t$  seconds send the lowest total number of bits. If this schedule violates the bit-rate requirements after  $t$  seconds, then all other schedules will also violate the requirement as they send at least as many bits.

An alternate way of phrasing this argument is to talk about the slack: allowed rate of  $rt$  minus the number of bits sent till time  $t$ . Using this notion, the *Greedy schedule is ahead as for any  $t$ , the first  $t$  seconds have the highest total slack possible for any schedule*. This is true for the same reason as used above: the schedule sends the  $t$  lowest transmission rates in the first  $t$  seconds. The schedule violates the bit-rate requirements after  $t$  seconds, if it has negative slack after  $t$  seconds, and then all other schedules will also violate the requirement as they have at most as much slack as this schedule.

**Version 3.** Prove that sorting by rates is an optimal schedule by an “exchange argument”.

Assume that there is a feasible schedule  $O$ , and let the algorithm’s schedule be  $A$ . We say that two jobs are inverted if they occur in different order in  $O$  and in  $A$ . As in earlier exchange arguments in the text, we know that if the two orders are different, then there are two adjacent jobs  $i$  and  $j$  in  $O$  (say  $i$  immediately follows  $j$ ) that are inverted. We need to argue that if  $O$  is feasible schedule, and  $i$  and  $j$  are inverted, then the schedule  $O$  obtained by swapping  $i$  and  $j$  is also feasible. Let  $T$  be the time  $j$  starts in  $O$ , and assume that the schedule  $O$  sends  $B$  bits in the first  $T$  seconds. The only times when the total amount sent so far is affected by the swap are the times in the range  $[T, T + t_j + t_i]$ . Let  $T + t$  be such a time. Assume that in the schedule  $O$  we send  $b_o$  bits during these  $t$  seconds. The schedule  $O$  is feasible, and so  $B + b_o \leq r(T + t)$ . By the ordering used of our algorithm (and the fact that the jobs  $i$  and  $j$  were inverted), the number of bits sent by the same  $t$  seconds in the swapped schedule  $O'$  is  $b' \leq b_o$ . Therefore, the new schedule satisfies  $B + b' \leq B + b_o \leq r(T + t)$ . Swapping a pair of adjacent inverted jobs decreases the number of inversions and keeps the schedule feasible. So we can repeatedly swap adjacent inverted jobs until the schedule  $O$  gets converted to the schedule of the greedy algorithm. This proves that the greedy algorithm produces a feasible schedule.

**Version 4.** In fact, we do not need to sort at all to produce a feasible schedule. For stream  $i$  compute its slack  $s_i = rt_i - b_i$ . We claim that if a feasible schedule exists, then any order that starts with all streams that have positive slack is feasible. To see why, observe that an ordering is feasible if the sum of slacks is non-negative for any initial segment of the order. Starting with all streams of positive slack creates the highest possible total slack before we start adding the jobs with negative slacks. This observation allows us to create the feasible ordering in  $O(n)$  time without sorting, by simply computing the sign of a slack for each stream.

Note that this argument also shows that sorting streams in decreasing in order of slacks works too, as this also orders streams with positive slack before those with negative slack.

Finally, note that one ordering that does not always work is to order streams in increasing order of bits  $b_i$ .