

HW 5 (due Monday, at noon, March 4, 2013)

CS 473: Fundamental Algorithms, Spring 2013

Version: 1.1

Make sure that you write the solutions for the problems on separate sheets of paper. Write your name and netid on each sheet.

Collaboration Policy: The homework can be worked in groups of up to 3 students each.

1. (40 PTS.) Simultaneous Climbs.

One day, Kris (whom you all know and love from past climbing competitions) got tired of climbing in a gym and decided to take a very large group of climber friends (after, all he is a popular guy) outside to climb. The climbing area where they went, had a huge wide boulder, not very tall, with various marked hand and foot holds. Kris took a look and quickly figured out an "allowed" set of moves that his group of friends would do so that they get from one hold to another. He also figured out the difficulty of each individual move and assigned a grade (weight) to it. The higher the weight, the harder the move. Let $G = (V, E)$ be the (undirected) graph with a vertex for each hold and an edge between two holds (u, v) if v can be reached from u (and vice versa) by one of the moves that Kris decided. For an edge $(u, v) \in E$, we have a weight $w(uv)$ associated with (u, v) which represents the difficulty that he assigned to that particular move (it is always positive, negative weights would mean that climbers can defy gravity).

A *k-climb* is a sequence where a climber performs k moves in sequence. In graph G it is represented by a simple path with exactly k edges in it. Two k -climbs are disjoint if they do not share any vertex. A collection M of k -climbs is a *k-climb packing* if all pairs of climbs of M are disjoint (for $k = 1$ the set M is a matching in the graph). The total weight of a k -climb packing is the total weight of the edges used by the climbers.

Kris and his friends decided to play a game (they are all very good climbers), where as many climbers as possible are simultaneously on the wall and each climber needs to perform a set of k moves in sequence. In other words, they are interested in the problem of computing the maximum weight k -climb packing in G . In general, this problem seems hard.

Describe an efficient algorithm (i.e., provide pseudo-code, etc), as fast as possible, for computing the maximum weight k -climb packing when G is a rooted tree (fortunately for the tree case this is much easier) or a forest (collection of rooted trees).

Your algorithm should be recursive and use memoization to achieve efficiency. (You can not assume G is a binary tree - a node might have arbitrary number of children.) What is the running time of your algorithm as function of $n = |V(G)|$ and k ?

For an example, the figure shows a tree with a possible 3-climb packing.

2. (40 PTS.) Help Fellow Climber Out and ARC Scheduling.

(A) (20 PTS.) Since Kris is a professional climber, the moves he had in mind for his friends to climb, unfortunately, did not result to the graph G from the previous question being a tree (or a collection of trees). In order for his friends to have fun on the boulder, you need to help Kris (who hasn't brushed up his algorithms in a long time) reduce his set

of moves (edges in G) so that G ends up being a tree or a forest but also that the total difficulty of moves he ignores is as small as possible (he doesn't want to be too easy on his friends). Formally, for the graph $G = (V, E)$ that appears in the previous question, describe an algorithm that removes the smallest weight subset of edges such that the remaining graph is a tree or a forest.

- (B) (20 PTS.) Consider a variant of interval scheduling except now the intervals are arcs on a circle. The goal is to find the maximum number of arcs that do not overlap. More formally, let C be the circle on the plane centered at the origin with unit radius. Let A_1, \dots, A_n be a collection of arcs on the circle where an arc A_i is specified by two angles $\alpha_i \in [0, 2\pi]$ and $\beta_i \in [0, 2\pi]$: the arc starts at the point on the circle C with angle α_i and goes counter-clockwise till the point on C at angle β_i (the end points are part of the arc). Two arcs overlap if they share a point on the circle. Describe an algorithm to find the maximum number of non-overlapping arcs in the given set of arcs.

3. (20 PTS.) Process these words.

In a word processor the goal of “pretty-printing” is to take text with a ragged right margin, like this

```
I guess it takes two
to progress
from an apprentice to a
legitimate surfer. Two digits
in front of the
size in feet of the
wave one needs to take, two double overhead waves that holds
one under after
the wipeout, and two equally sized pieces that
one's previously intact board comes up
on the surface as.
```

and turn it into text whose right margin is as “even” as possible, like this

```
I guess it takes two to progress
from an apprentice to a legitimate
surfer. Two digits in front of
the size in feet of the wave one
needs to take, two double overhead
waves that holds one under after
the wipeout, and two equally sized
pieces that one's previously intact
board comes up on the surface as.
```

To make this precise enough for us to start thinking about how to write a pretty-printer for text, we need to figure out what it means for the right margins to be “even”. So suppose our text consists of a sequence of *words*, $W = \{w_1, w_2, \dots, w_n\}$, where w_i consists of c_i characters. We have a maximum line length of L . We will assume we have a fixed-width font and ignore issues of punctuation or hyphenation.

A *formatting* of W consists of an ordered partition of the words in W into *lines*. In the words assigned to a single line, there should be a space after each word except the last; and so if w_j, w_{j+1}, \dots, w_k are assigned to one line, then we should have

$$\left[\sum_{i=j}^{k-1} (c_i + 1) \right] + c_k \leq L.$$

We will call an assignment of words to a line *valid* if it satisfies this inequality. The difference between the left-hand side and the right-hand side will be called the *slack* of the line—that is, the number of spaces left at the right margin.

Given a partition of a set of words W , the *penalty* of the formatting is the sum of the *squares* of the slacks of all lines (including the last line). Give an efficient algorithm to find a partition of a set of words W into valid lines, so that the penalty of the formatting becomes minimized.