

TI2306 Algorithm Design

Digital test – part 2 of 2

January 29, 2019, 13.30-15.30

- Usage of the book, notes or a calculator during this test is not allowed.
- This digital test contains 3 questions (worth a total of 10 points) contributing 10/20 to your grade for the digital test. The other digital test contributes the other 10/20 to your grade. Note that the final mark for this course also consists for $\frac{2}{3}$ of the unrounded score for the written exam (if both ≥ 5.0).
- The storyline throughout the questions can be ignored: each question can be answered independently of the others and in any order.
- The spec tests in WebLab **do not represent your grade**.
- There is an API specification of the Java Platform (Standard Edition 8) available at <https://weblab.tudelft.nl/java8/api>.
- When an implementation is asked, please provide the **most efficient** implementation in terms of asymptotic time complexity. Providing a suboptimal implementation can lead to a subtraction of points.
- The material for this tests consists of module 3 (Dynamic Programming) and 4 (Network Flow) of the course and the corresponding book chapters and lectures.
- Total number of pages (without this preamble): 1.

1. (3 points) A company behind an operating system supplies updates (e.g., with security-related bug fixes) via the Internet. At a certain moment in time, a number of updates are in development, and for each update we know the exact future date when it is ready. In order of these dates, the updates are denoted by their index in the sequence $1, 2, \dots, n$.

Shipping an update comes with some constant (and known) costs c . To minimise these shipping costs, the board of directors of the company are asking you to explore the possibility of *bundling* updates. A *bundle* is a series of consecutive updates $i, i+1, \dots, j$ that is shipped in one go.

The total costs of a bundle consists of shipping costs and risk-related costs: shipping one complete bundle is equal to the cost of a single update (c). However, postponing security updates until the complete bundle of updates is ready means users are at risk longer. These (estimated) costs of a bundle $i, i+1, \dots, j$ are given to us, denoted by $\text{costs}[i][j]$ for all $1 \leq i \leq j \leq n$.¹ The board would like to know how to bundle updates, so that the sum of the costs of all bundles is minimised.

Implement an efficient *iterative* dynamic programming algorithm to compute the minimal costs of bundling. You are given the following input: the number of updates n , the shipping costs c , and the table costs with dimensions $(n+1) \times (n+1)$.

2. (3 points) When you present your solution to the board members, one of them explains that the content of updates come from code changes that either resolve bugs or include new features. We call these code changes "jobs". Each job $i \in \{1, 2, \dots, n\}$ has some (estimated) costs $\text{costs}[i]$, representing the number of hours required for a developer, and some (estimated) benefits $\text{benefits}[i]$ to end-users. There is a dynamic program in place that finds the maximal total benefit of any list of selected jobs given a (monthly) cost budget of C : this program fills an array M with at position $M[i][c]$ for $i \in \{1, 2, \dots, n\}$ and $0 \leq c \leq C$ the benefit of the optimal selection of jobs from $\{1, \dots, i\}$ such that the total costs of the included jobs is not more than the budget c , i.e., this array contains the results of the following recursive function:

$$M(i, c) = \begin{cases} 0 & \text{if } i < 1 \\ M(i-1, c) & \text{if } \text{costs}[i] > c \\ \max\{M(i-1, c), \text{benefits}[i] + M(i-1, c - \text{costs}[i])\} & \text{otherwise} \end{cases}$$

Implement an algorithm that determines which jobs together have the maximum total benefit *using the table M* . As input you get: the number of jobs n , the total cost budget C , the array benefits of dimension $n+1$, the array costs of dimension $n+1$, and the array M filled as described in the text above, with dimensions $(n+1) \times (C+1)$. The algorithm should return a list of the indices of the included jobs *in increasing order*. This algorithm should run in linear time.

3. (4 points) When you bring your solution to the attention of the developers of the operating system, it appears that the board of directors (who gave you the assignment in the first place) completely misunderstood how such a selection of code changes ("jobs") is made:
 1. Jobs are sometimes dependent on each other: there are a number of pairs of jobs (i, j) for which it holds that if i is included, then also j must be included in the new release.
 2. There is no exact limit on the budget: when it is clear that the benefits outweigh the costs (e.g., with high-risk security bugs), developers from other departments jump in to help, or implementing all the jobs just takes longer (e.g., than a month).

You are asked to implement an algorithm to find a subset of jobs that maximises the net profit (total benefit minus total costs) under these (new) conditions.

You are again given a list of n jobs, with for each job $1 \leq i \leq n$ the $\text{costs}[i]$ and the $\text{benefits}[i]$ in arrays as for question 2. Also now you are given a list of pairs of code changes (i, j) where i is dependent on j . These are represented as a 2D-array dependencies of dimension $(n+1) \times (n+1)$ where $\text{dependencies}[i][j] = 1$ iff i depends on j and 0 otherwise.

Create a flow network and use the given implementation of Ford-Fulkerson to find the subset of jobs with maximal net profit (to include in the next release). Output only the (net) value of this release, i.e., the benefit minus the costs of the included jobs.

¹Where $\text{costs}[i][j] = 0$ for all $i \geq j$.