

Homework 2

Instructor: Dieter van Melkebeek

TA: Nicollas Mocolin Schroevski

This homework covers the divide and conquer paradigm. **Problem 3 must be submitted for grading by 2:29pm on 9/27.** Please refer to the homework guidelines on Canvas for detailed instructions.

Warm-up problems

- (a) Design an algorithm to compute the binary representation of 10^n in $O(n^{\log_2 3})$ time.
(b) Design an algorithm that converts a given n -digit decimal number to binary in $O(n^{\log_2 3})$ time.
- Given an array $A[1 \dots n]$ of integers and a positive integer k , we want to find a rearrangement B of A such that the subarrays $B_1 \doteq B[1 \dots k]$, $B_2 \doteq B[k+1 \dots 2k]$, \dots , satisfy the following property: For every $i < j$, every element of B_i is less than or equal to every element of B_j . Design an $O(n \log(n/k))$ algorithm for this problem. You can assume that all elements in the array are distinct.

Regular problems

- [Graded] You are given a sequence of n real numbers a_1, a_2, \dots, a_n and a corresponding sequence of weights w_1, w_2, \dots, w_n . The weights are nonnegative reals that add up to 1, i.e., $\sum_{i=1}^n w_i = 1$. The weighted median of the sequence is the number a_k such that

$$\sum_{a_i < a_k} w_i < \frac{1}{2} \quad \text{and} \quad \sum_{a_i \leq a_k} w_i \geq \frac{1}{2}.$$

For example, the weighted median of the following instance is 2.5:

i	1	2	3	4	5	6	7
a_i	.40	-.5	4	0	2.5	6	-.2
w_i	.25	.1	.05	.18	.15	.2	.07

Design an $O(n)$ algorithm that finds the weighted median using elementary operations. An addition or a multiplication of two real numbers counts as one elementary operation.

- You are given an $n \times n$ grid, and a procedure $V(i, j)$ that assigns an integer value to each position (i, j) in the grid, where i and j are integers such that $1 \leq i, j \leq n$. Your goal is to find a local minimum (or sink) in the grid (i.e., integers i^*, j^* with $1 \leq i^*, j^* \leq n$ such that for all neighbors (i, j) of (i^*, j^*) in the grid, $V(i^*, j^*) \leq V(i, j)$). The neighbors of (i, j) are $(i-1, j)$, $(i+1, j)$, $(i, j+1)$, $(i, j-1)$; the elements along the diagonals do not count as neighbors.

Design an algorithm that makes $O(n)$ calls to V . Note that the grid has n^2 nodes.

- You are given a topographical map that provides the maximum altitude along the direct road between any two neighboring cities, and two cities s and t . All roads can be traversed in both directions. Your goal is to find a route from s to t that minimized the maximum altitude.

Design an algorithm that runs in time $O(n+m)$ on instances with n cities and m routes when the input is given in adjacency-list format (for each city, a list of all routes that go from that city to a neighboring city).

Challenge problem

The following is one of the nicest introductory algorithm problems I know. Give it a try!

- You are given n coins, at least one of which is bad. All the good coins weigh the same, and all the bad coins weigh the same. The bad coins are lighter than the good coins.

Design an algorithm that makes $O((\log n)^2)$ weighings on a balance to find the exact number of bad coins. Each weighing tells you whether the total weight of the coins on the left side of the balance is smaller than, equal to, or larger than the total weight of the coins on the right side.

Programming problem

- SPOJ problem [Aggressive Cows](#) (problem code AGGRCOW).

Homework 2 Solutions to Warm-up Problems

Instructor: Dieter van Melkebeek

TA: Nicollas Mocolin Schroevski

Problem 1

- Design an algorithm to compute the binary representation of 10^n in $O(n^{\log_2 3})$ time.
- Design an algorithm that converts a given n -digit decimal number to binary in $O(n^{\log_2 3})$ time.

Part (a)

The idea is to use fast powering (H1 01, prob. 1) with the nontrivial divide and conquer integer multiplication algorithm from class (denoted INTEGER-MULTIPLICATION below). Pseudocode is in algorithm 1.

Algorithm 1

Input: $n \in \mathbb{Z}$, $n \geq 1$

Output: 10^n in binary

```

1: procedure POWER-OF-TEN( $n$ )
2:   if  $n = 1$  then
3:     return 1010
4:   else
5:      $c \leftarrow$  POWER-OF-TEN( $\lfloor n/2 \rfloor$ )
6:      $s \leftarrow$  INTEGER-MULTIPLICATION( $c, c$ )
7:     if  $n$  is even then
8:       return  $s$ 
9:     else ( $n$  is odd)
10:      return  $(s \cdot 2^3) + (s \cdot 2)$  where multiplication by  $2^k$  is implemented by shifting

```

Correctness If $n = 1$, then the algorithm returns 1010 correctly. For $n > 1$, the basic observation is the following: if c is the binary representation of $10^{\lfloor n/2 \rfloor}$, then correctness of INTEGER-MULTIPLICATION implies that s is the binary representation of $10^{2\lfloor n/2 \rfloor}$. When n is even, $10^{2\lfloor n/2 \rfloor} = 10^n$, so the algorithm is correct to return s . When n is odd, the expression $(s \cdot 2^3) + (s \cdot 2)$ computes $s \cdot 10$. Since $s \cdot 10 = 10^{2\lfloor n/2 \rfloor + 1} = 10^n$, the algorithm is correct in this case as well. This can be formalized by induction on n , where the inductive hypothesis guarantees that the recursive call computing the value of c is correct.

Running Time We use the recursion tree method. The shape of the tree is a line, with input size n at the root and decreasing by half with each step down the tree. The depth of the tree is at most $\log n$. The work at a node with input size ℓ consists of a call to INTEGER-MULTIPLICATION with $O(\ell)$ -bit inputs, plus sums and shifts of numbers of $O(\ell)$ bits. This work is bounded by $c\ell^{\log_2 3}$

for some constant c . At depth d , the input size ℓ is bounded by $n/2^d$. Summing this over all d , the total work of the algorithm is

$$\sum_{d=0}^{\log n} c \left(\frac{n}{2^d} \right)^{\log_2 3} = cn^{\log_2 3} \sum_{d=0}^{\log n} \left(\frac{1}{3} \right)^d \leq cn^{\log_2 3} \sum_{d=0}^{\infty} \left(\frac{1}{3} \right)^d = cn^{\log_2 3} \cdot \frac{3}{2},$$

where the final equality uses the formula for summing a geometric series. The overall running time is therefore $O(n^{\log_2 3})$.

Part (b)

The idea is to divide the digits of the input into two halves, recursively compute the binary representations of the number represented by each half, and use the INTEGER-MULTIPLICATION routine and the solution from part (a) to combine the halves into the whole.

In more detail, given an n -digit decimal number X , we divide its digits into two parts of size about $n/2$: $X = L \cdot 10^{\lfloor n/2 \rfloor} + R$, where $L, R < 10^{\lfloor n/2 \rfloor}$ are integers. Next we convert L and R to binary, and use the subroutine from part (a) to compute the binary representation of $10^{\lfloor n/2 \rfloor}$. Finally, we use INTEGER-MULTIPLICATION to compute $L \cdot 10^{\lfloor n/2 \rfloor}$ in binary, and add to this the binary representation of R .

Algorithm 2

Input: X , an n -digit integer given in decimal

Output: the binary representation of X

```

1: procedure DECIMAL-TO-BINARY( $X$ )
2:   if  $n = 1$  then
3:     return binary representation of the digit of  $X$  [lookup table omitted for brevity]
4:   else
5:     Let  $R$  be the lowest  $\lfloor n/2 \rfloor$  digits of  $X$ , and  $L$  the remaining digits.
6:      $L' \leftarrow$  DECIMAL-TO-BINARY( $L$ )
7:      $R' \leftarrow$  DECIMAL-TO-BINARY( $R$ )
8:      $P \leftarrow$  POWER-OF-TEN( $\lfloor n/2 \rfloor$ )
9:     return INTEGER-MULTIPLICATION( $L', P$ ) +  $R'$ 

```

Correctness Again, when $n = 1$ the algorithm is clearly correct. For $n > 1$, the main observation is that when R are the lowest $\lfloor n/2 \rfloor$ digits of X and L the remaining digits, then $X = L \cdot 10^{\lfloor n/2 \rfloor} + R$. In that case, if L' and R' are the binary representations of L and R , respectively, and P is the binary representation of $10^{\lfloor n/2 \rfloor}$ (by correctness of part (a)), then the algorithm returns the correct value. As before, this can be formalized by induction on n , where the inductive hypothesis guarantees that the recursive calls made by the algorithm are correct.

Running Time We use the recursion tree method. Each non-leaf node in the recursion tree has two children, and the input size n shrinks by half with each level of recursion, so the tree has the same shape as MERGESORT. At a node with input size ℓ , the non-recursive work done is essentially the calls to POWER-OF-TEN and INTEGER-MULTIPLICATION on inputs of size $O(\ell)$.

The subroutines each run in $O(\ell^{b \log_2 3})$, so for some constant c the work at a node with input size ℓ is $c\ell^{b \log_2 3}$.

At depth d , the input size ℓ is at most $\frac{n}{2^d}$. As there are 2^d nodes at depth d , the total work at level d is

$$2^d \cdot c \left(\frac{n}{2^d}\right)^{b \log_2 3} = cn^{\log_2 3} \left(\frac{2}{3}\right)^d.$$

As in part (a), the sum over all d behaves as a geometric series with base less than 1. So the total work done is $O(n^{b \log_2 3})$.

Problem 2

Given an array $A[1 \dots n]$ of integers and a positive integer k , we want to find a rearrangement B of A such that the subarrays $B_1 \doteq B[1 \dots k]$, $B_2 \doteq B[k+1 \dots 2k]$, \dots , satisfy the following property: For every $i < j$, every element of B_i is less than or equal to every element of B_j . Design an $O(n \log(n/k))$ algorithm for this problem. You can assume that all elements in the array are distinct.

We first consider the case where $n = 2k$. Here the problem becomes to rearrange A into B so that, when divided into two halves, B_1 and B_2 , each element of B_1 is less than or equal to every element of B_2 . The running time requirement becomes $O(n \log(n/k)) = O(n)$.

One way to rephrase the problem is that we want to rearrange A into B so that B_1 consists exclusively of elements at most the median, and B_2 consists of elements at least the median. Using linear-time selection, (FAST-SELECT in the notes), we can find the median of A . Using the splitting subroutine (SPLIT in the notes), we can use the median to divide A into the requisite parts B_1 and B_2 , with $B = B_1 \cdot B_2$. Both algorithms are linear-time, so this addresses the $n = 2k$ case.

Plugging the above idea into a typical divide-and-conquer scheme, we can handle the case where n/k is a power of 2. We FAST-SELECT the median of A , use it to SPLIT A into its lower and upper halves, and then recurse on the left and right halves. This idea works also for general k , but we need to take care to select an element of A on the boundary of the middle group instead of just the median of A . Complete pseudocode is below.

Algorithm 3

Input: $A[1 \dots n], k$. A is an array of integers, and k is a positive integer.
Output: A permuted copy B of A so that the subarrays $B_1 \doteq B[1 \dots k]$, $B_2 \doteq B[k+1 \dots 2k]$, \dots satisfy that for every $i < j$, every element of B_i is less than or equal to every element of B_j .

```

1: procedure PARTIAL-SORT( $A$ )
2:   if  $n \leq k$  then
3:     return  $A$ 
4:   else
5:      $b \leftarrow \lceil n/k \rceil$  (the number of subarrays  $A_i$  of size  $k$ )
6:      $r \leftarrow k \lfloor b/2 \rfloor$  (number of elements in  $A_1, \dots, A_{\lfloor b/2 \rfloor}$ )
7:      $p \leftarrow$  FAST-SELECT( $A, r$ )
8:     ( $L, R$ )  $\leftarrow$  SPLIT( $A, p$ )
9:     Append  $p$  to the right of  $L$ 
10:     $L' \leftarrow$  PARTIAL-SORT( $L, k$ )
11:     $R' \leftarrow$  PARTIAL-SORT( $R, k$ )
12:    return  $L' \cdot R'$ 

```

Correctness When $n \leq k$, then there is no constraint on the ordering of A , so the algorithm returning $B = A$ is a correct result. For $n > k$, let $b = \lceil n/k \rceil$ be the number of subarrays B_1, B_2, \dots each of size k , and let $r = k \lfloor b/2 \rfloor$ be as in the algorithm. By the specifications of FAST-SELECT and SPLIT, after line 9, the smallest r elements of A are in array L , and the largest $n - r$ elements are in array R . After the two recursive calls, both L' and R' obey the output condition. Considering

the output $B = L' \cdot R'$, we now check that, for every $i < j$, each element of B_i is less than or equal to each element of B_j . There are three cases:

- When i and j are both at most $\lfloor b/2 \rfloor$, this follows from L' obeying the output condition.
- When i and j are both larger than $\lfloor b/2 \rfloor$, this similarly follows from R' obeying the output condition.
- When $i \leq \lfloor b/2 \rfloor$ and $j > \lfloor b/2 \rfloor$, each element of B_i is among the smallest r elements of A , and each element of B_j is among the largest $n - r$ elements of A , so correctness holds in this case as well.

And this completes the proof. As with the other problems, this argument can be made formal by induction on n .

Running Time We use the recursion tree method. Each non-leaf node has two children, and the input size n is reduced by half with each level of recursion. The leaves correspond to when $n = k$, so the depth is the number of times n must be halved (up to rounding) before becoming equal to k . This is $O(\log(n/k))$. For any node in the recursion tree, the work local to that node consists of calls to FAST-SELECT and SPLIT, which are linear time. So for some constant c , this work is bounded by $c\ell$. Nodes at depth d have input size $\ell \leq n/2^d$. Summing the work across all 2^d nodes at level d , we get a bound of cn . Summing these across all levels, we get a total running time of $cn \log(n/k) = O(n \log(n/k))$.

CS 577: Introduction to Algorithms

Fall 2022

Homework 2 Solutions to Regular Problems

Instructor: Dieter van Melkebeek

TA: Nicollas Mocolin Sdroieviski

Problem 3

You are given a sequence of n real numbers a_1, a_2, \dots, a_n , and a corresponding sequence of weights w_1, w_2, \dots, w_n . The weights are nonnegative reals that add up to 1, i.e., $\sum_{i=1}^n w_i = 1$. The weighted median of the sequence is the number a_k such that

$$\sum_{a_i < a_k} w_i < \frac{1}{2} \quad \text{and} \quad \sum_{a_i \leq a_k} w_i \geq \frac{1}{2}.$$

Design an $O(n)$ algorithm that finds the weighted median using elementary operations. An addition or a multiplication of two real numbers counts as one elementary operation.

First, note that if we were allowed $O(n \log n)$ time for our procedure, a very simple solution would exist—sort the elements $\{a_i\}_{i \in [n]}$ in nondecreasing order, then compute the total weight of the first element in the sorted list, then the first two elements, the first three, and so on until you find a prefix of the elements with total weight at least $1/2$. The last element of this prefix would be the weighted median. Unfortunately, we have only $O(n)$ time, so a more sophisticated approach is required.

In particular, we will use divide-and-conquer. When applying divide-and-conquer to arrays, a common approach is to use linear-time selection to find the median m of the array, and then to split the array into three parts: a part L consisting of all elements less than the median, a part M consisting of all elements equal to the median, and a part R consisting of all elements greater than the median. Trying this method here yields an interesting property.

Note that we can compute the total weight $w(L)$ of L in time linear in the size of L , and similarly for the analogously defined $w(M)$. After computing each of these, we have three cases:

- $w(L) \geq 1/2$. In this case, we know that the weighted median must be in L .
- $w(L) < 1/2$ and $w(L) + w(M) \geq 1/2$. In this case, we know that the weighted median must be exactly m .
- $w(L) + w(M) < 1/2$. In this case, since we know that all the weights in the array must sum to 1, we know that the weighted median must be in R .

This looks promising, because in every case we have either reduced by half the size of the array where we need to search for the weighted median, or solved the problem outright. However, there is a small issue: while in the first case, we can find the weighted median by recursively calling our procedure on L (even though the weights may not sum to 1), in the third case, we need to use a different threshold for the total weight. This is because the total weight of all elements in L and M should count towards the $1/2$ threshold, so we are really looking for the smallest element m' of R such that the total weight of all elements $\leq m'$ in R is $\geq 1/2 - w(M) - w(L)$.

This suggests that to solve the problem of weighted median, we will actually need to solve the slightly more general problem of weighted selection, specified and implemented in Algorithm 1. Some notation in this procedure must be explained.

- The method MEDIAN takes an input array and returns the median of that array in linear time.
- The method SUM takes an input array and returns the sum of its elements, also in linear time.
- The notation $[a_i \in A : a_i < m]$ should be parsed as “the subsequence of A consisting of all elements of A less than m , stored as an array”. Similarly, $[u_i \in W : a_i < m]$ is “the subsequence of W consisting of all elements of W corresponding to elements of A that are smaller than m , stored as an array”. Note that subsequences preserve the original ordering of the elements and allow duplicates. Each of these subsequences can be computed from A and W in linear time.

Additionally, the input A should be interpreted as the original sequence of real numbers, W as the weights of those numbers, and t as the target threshold (which is $1/2$ for the original call if we are looking for the weighted median).

Algorithm 1

Input: an array $A[1..n]$ of real numbers, an array $W[1..n]$ of nonnegative real numbers, and a real number t such that there exists k for which $\sum_{i:A[i]<A[k]} W[i] < t$ and $\sum_{i:A[i]\leq A[k]} W[i] \geq t$

Output: $A[k]$ for a k that satisfies the precondition

```

1: procedure WEIGHTEDSELECTION( $A, W, t$ )
2:    $m \leftarrow \text{MEDIAN}(A)$ 
3:    $L \leftarrow [a_i \in A : a_i < m], M \leftarrow [a_i \in A : a_i = m], R \leftarrow [a_i \in A : a_i > m]$ 
4:    $W_L \leftarrow [w_i \in W : a_i < m], W_M \leftarrow [w_i \in W : a_i = m], W_R \leftarrow [w_i \in W : a_i > m]$ 
5:    $w(L) \leftarrow \text{SUM}(W_L), w(M) \leftarrow \text{SUM}(W_M)$ 
6:   if  $w(L) < t$  and  $w(L) + w(M) \geq t$  then
7:     return  $m$ 
8:   else
9:     if  $w(L) \geq t$  then
10:      return WEIGHTEDSELECTION( $L, W_L, t$ )
11:    else
12:      return WEIGHTEDSELECTION( $R, W_R, t - w(M) - w(L)$ )

```

Correctness Correctness of WEIGHTEDSELECTION follows along the lines of the intuition given at the beginning of this solution. There are three cases:

- If $w(L) < t$ and $w(L) + w(M) \geq t$, then M must contain an entry with an index in A that satisfies the precondition. Since every element of M is equal to m , m is the correct value to return.
- If $w(L) \geq t$, then for any index k satisfying the precondition of WEIGHTEDSELECTION(A, W, t), all the elements of A with value at most $A[k]$ are included in L . In particular, there exists ℓ such that $L[\ell] = A[k]$, $\sum_{i:A[i]<A[k]} W[i] = \sum_{j:L[j]<L[\ell]} W_L[j]$, and $\sum_{i:A[i]\leq A[k]} W[i] =$

2

$\sum_{j:L[j]\leq L[\ell]} W_L[j]$. It follows that the arguments of the recursive call WEIGHTEDSELECTION(L, W_L, t) in line 10 satisfy the preconditions, and the call returns $L[\ell] = A[k]$, which is correct.

- Otherwise, we know that $w(L) + w(M) < t$. In that case any index k satisfying the precondition of WEIGHTEDSELECTION(A, W, t) has to correspond to an element of R , and all elements with value less than $A[k]$ are included in L and M . Thus, there exists ℓ such that $R[\ell] = A[k]$, $\sum_{i:A[i]<A[k]} W[i] = w(L) + w(M) + \sum_{j:R[j]<R[\ell]} W_R[j]$, and $\sum_{i:A[i]\leq A[k]} W[i] = w(L) + w(M) + \sum_{j:R[j]\leq R[\ell]} W_R[j]$. Since

$$w(L) + w(M) + \sum_{j:R[j]<R[\ell]} W_R[j] < t \Leftrightarrow \sum_{j:R[j]<R[\ell]} W_R[j] < t - w(L) - w(M)$$

and

$$w(L) + w(M) + \sum_{j:R[j]\leq R[\ell]} W_R[j] \geq t \Leftrightarrow \sum_{j:R[j]\leq R[\ell]} W_R[j] \geq t - w(L) - w(M),$$

it follows that the arguments of the recursive call WEIGHTEDSELECTION(L, W_L, t) in line 12 satisfy the preconditions, and that the call returns $R[\ell] = A[k]$, which is correct.

Running time We use the recursion tree method. The shape of a tree is a line, where the input size of a node reduces by half at each step. The local work done at a node of size ℓ is at most $c\ell$ for some constant c . This is because, as noted in the explanation of the algorithm’s notation, each of lines 2-5 only takes linear time to execute. If the recursion tree has depth d and the base case takes e operations to compute, then the total time taken is at most

$$e + \sum_{i=0}^{d-1} \frac{cn}{2^i} \leq e + cn \sum_{i=0}^{\infty} \frac{1}{2^i} = e + 2cn,$$

which is $O(n)$.

3

Problem 4


You are given an $n \times n$ grid, and a procedure $V(i, j)$ that assigns an integer value to each position (i, j) in the grid, where i and j are integers such that $1 \leq i, j \leq n$. Your goal is to find a local minimum (or sink) in the grid (i.e., integers i^*, j^* with $1 \leq i^*, j^* \leq n$ such that for all neighbors (i, j) of (i^*, j^*) in the grid, $V(i^*, j^*) \leq V(i, j)$). The neighbors of (i, j) are $(i - 1, j)$, $(i + 1, j)$, $(i, j + 1)$, $(i, j - 1)$; the elements along the diagonals do not count as neighbors.

Design an algorithm that makes $O(n)$ calls to V . Note that the grid has n^2 nodes.

We first observe that any $n \times n$ grid contains a local minimum, since any path constructed by repeatedly stepping toward a neighbor with a strictly smaller label must terminate at a local minimum. While this is indeed an algorithm, it isn’t efficient enough as it could potentially take $\Theta(n^2)$ probes.

We use divide-and-conquer to find a local-minimum in $O(n)$ probes. We think of dividing the input grid into four sub-grids, as in Figure 1. Generically speaking, a local minimum of any one of the sub-grids is a local minimum of the overall grid, so the idea is to do some $O(n)$ work to decide which of the four sub-grids to recurse into, and do that. Since n shrinks by half with each layer of recursion, the overall running time will behave like a geometric series, and be $O(n)$ overall.

Figure 1: Dividing a grid into four sub-grids.



This intuition about minima of sub-grids is not quite right. Our first hint of this is that, if every local minimum of a sub-grid were a local minimum of the overall grid, then we could pick any of the four sub-grids to recurse into. Thus we would be able to find a local minimum with zero probes! Clearly this is wrong, so something is missing.

And indeed, it is possible for an element to be a local minimum in a sub-grid and not be a local minimum in the overall grid. For example, there may be a smaller neighbor in a different sub-grid. That being said, it is true that if every neighbor of a sub-grid’s local minimum lies within the same sub-grid, then it is a local minimum of the overall grid. This suggests a more refined strategy: prior to recursing into a sub-grid, we try to ensure that the local minimum found is not on the boundary.

To manage this, first observe that there are only $O(n)$ many positions that have a neighbor in a different sub-grid. We will refer to these positions as boundary positions; they are highlighted as the blue region in Figure 1. We can afford to examine them all. We can use that information to decide which of the four sub-grids to recurse into.

4

An interesting boundary position is the one that has the smallest label among all the boundary positions. Call this position m . Its label is clearly no larger than the label of any of its neighbors that are also boundary positions. So the only way m is not a local minimum of the grid is if it has a smaller neighbor in the same sub-grid. We can imagine following a path starting from m and repeatedly stepping to an adjacent position with strictly smaller label, until reaching a local minimum. Because m is the minimum of all boundary positions, this path can never return to the boundary of the subgrid. In particular, there has to be a local minimum in the sub-grid containing m that is not on the boundary of that sub-grid. Thus suggests we should recurse into the sub-grid containing m .

This almost works! There is one issue though. The specification of our algorithm is that it finds some local minimum. When we recurse onto the sub-grid containing m , there is no guarantee the local minimum found is the one arrived at by repeatedly stepping to a neighbor with strictly smaller label as above. Indeed, the recursive call could very well return a boundary position. We need to strengthen our specification so that it not only returns a local minimum, but one that we can argue is not on the boundary.

To do this, we will require that the local minimum we find in a sub-grid has label no larger than the label of m . The local minimum arrived at by repeatedly stepping to a neighbor with strictly smaller label satisfies this property, so certainly it exists in the sub-grid containing m . Let’s carefully write down a specification for our algorithm:

Input: A grid G , a labeling procedure V , and a position t in the grid.

Output: A local minimum of the grid that has label no larger than the label of t .

At the top layer of the recursion there is no need for the additional argument; we can pass an arbitrary position. Since we adjusted the specification, we need to ensure our algorithm matches it as well. In particular, we now have the additional input of some position t in the grid, and need to make sure the local minimum we find has label smaller than the label of that position. If the label of m is at most the label of t , we are fine. Otherwise, the label of t is smaller than the label of m , and thus also smaller than every label of a boundary position, so we can instead recurse into the sub-grid containing t . This gives the following algorithm:

Correctness: We prove correctness by induction on n , the size of the grid. The n -th assertion is that FINDGRIDLOCALMINIMUM matches its specification for all grids of size at most n .

Base Case: For the case in which the grid has size $n = 1$, there is only one position, namely t . As it has no neighbors, it is a local minimum, so the algorithm is correct.

Inductive step: When $n > 1$, the algorithm falls into its recursive case. Let m be as in the algorithm and B represent the set of boundary positions. In both the case $V(t) \geq V(m)$ and the case $V(t) < V(m)$, we return the result of a recursive call FINDGRIDLOCALMINIMUM(G', u), where G' is a subgrid of G and u is a node in G' such that $V(u) \leq V(b)$ for every $b \in B$ and $V(u) \leq V(t)$. By the inductive hypothesis, the call returns a local minimum m' of the subgrid G' with $V(m') \leq V(u)$. Because m' is a local minimum in G' , every neighbor of m' in G' has no smaller label. Every other neighbor of m' is a position $b \in B$. Since $V(m') \leq V(u) \leq V(b)$, it follows that every neighbor of m' has no smaller label, and so m' is a local minimum in G . Finally, $V(m') \leq V(u) \leq V(t)$, so the condition that $V(m') \leq V(t)$ is satisfied as well.

5

Algorithm 2 FINDGRIDLOCALMINIMUM

Input: An $n \times n$ grid G ; a position t in G ; access to a procedure V assigning an integer to each position in G .

Output: A local minimum of G with label at most $V(t)$

```
1: procedure FINDGRIDLOCALMINIMUM( $G, t$ )
2:   if  $n = 1$  then
3:     return  $t$ 
4:   else
5:     Divide  $G$  as in Figure 1
6:     Let  $B$  represent the boundary positions
7:     Let  $m$  be a position s.t.  $V(m) = \min_{b \in B} V(b)$ 
8:     if  $V(t) \geq V(m)$  then
9:       Let  $G'$  be the quadrant containing  $m$ 
10:      return FINDGRIDLOCALMINIMUM( $G', m$ )
11:    else
12:      Let  $G'$  be the quadrant containing  $t$ 
13:      return FINDGRIDLOCALMINIMUM( $G', t$ )
```

Analysis: We use the recursion tree method. The shape of the recursion tree is a line, with input size n at the root and decreasing by half with each layer of recursion. The work done at a node with input size ℓ is some constant amount of book-keeping plus the cost of finding the minimum boundary position. As there are at most 4ℓ boundary positions, and finding a minimum takes linear time, this cost is $c\ell$ for some constant c . Nodes at depth d of the recursion tree have input size $\ell \leq n/2^d$, summing the work per node over the whole tree gives the bound $cn + c\frac{n}{2} + \dots \leq 2cn = O(n)$ on the total work done.

Problem 5

You are given a topographical map that provides the maximum altitude along the direct road between any two neighboring cities, and two cities s and t . All roads can be traversed in both directions. Your goal is to find a route from s to t that minimizes the maximum altitude.

Design an algorithm that runs in time $O(n + m)$ on instances with n cities and m routes when the input is given in adjacency-list format (for each city, a list of all routes that go from that city to a neighboring city).

Let G be the input graph on n vertices and m edges, with each vertex representing a city. Two vertices are connected by an edge if there is a direct road between them on our map, and the weight of each edge is the maximum altitude that occurs on that road. In this view, we are looking for a path from s to t such that the largest weight on that path is as small as possible.

To achieve this, we try to construct a path from s to t using the lightest edges possible. As a first step, we find the smallest threshold τ such that we can get from s to t using only edges with weight at most τ . Once we find τ , we can use it to find a path in which all edges have weight at most τ : First, we discard all edges with weight larger than τ . Then we perform a breadth-first search on the remaining graph to find a path. Deleting the edges and performing a breadth-first search can be done in $O(n + m)$ time. So from here on out, we focus on finding the smallest threshold τ such that we can get from s to t using only edges with weight at most τ .

We start by doing a breadth-first search on G to ensure that there is, in fact, a path from s to t . This takes time $O(n + m)$. If there is no path, then we are done. Otherwise, note that we can focus just on the connected component of G that contains s and t ; the rest of the graph is irrelevant to us. From now on, we will assume G is connected. In this case, $m \geq n - 1$, and we can express the running time of BFS and related algorithms in terms of m only. (This helps to simplify the analysis.) We also know that there exists a path from s to t . The question then becomes, how small can we set a threshold τ so that there exists a path from s to t among edges with weight at most τ ?

For this we can use divide and conquer, specifically binary search. For a given threshold τ , we can check whether there is a path from s to t that uses only edges of weight at most τ . This amounts to running a breadth-first search from s that ignores edges of weight more than τ . If the path exists, we need not consider larger values of τ ; if no such path exists, we need not consider smaller values. The question remains how to pick the values of τ to test.

It is clear that the value of τ we ultimately want is equal to the weight of some edge in the graph. Given that, we want to search for τ among the list of edge weights. Following the pattern of binary search, we maintain a list of viable thresholds (initially all edge weights), and use the median of that list as τ in the above test. Based on the result of the test, we either recursively search among thresholds less than the median or larger than the median; either way removes half the candidates.

This gives a complete algorithm. Initially there are m viable thresholds, so the binary search requires $\log m$ iterations to complete. Each iteration takes time $O(m)$ to run, as the test has to run a BFS on the whole graph, including the edges with nonviable weights. This gives a running time of $O(m \log m)$. This is too slow! We need to make it faster.

To do better, we need to keep the BFS from considering edges with nonviable weights. If we can do this, the work done by the BFS will shrink with the list of viable edge weights, i.e., it will halve from one step to the next. In total, the overall cost of the binary search behaves like a geometric

series, and the running time becomes $O(m)$.

Edges with weights too large to be viable are easy to handle. We just delete them from the graph. Then the BFS does not have to waste time seeing and ignoring them.

As for edges with weights too small to be viable, we want the BFS to treat sets of vertices connected by such edges in one fell swoop. A way to do this is to merge all vertices that are connected by such edges. In a bit more detail, we first form a subgraph G_1 of G with the same vertices and with only the edges of G with weights too small to be viable. From G_1 , we form another graph, G_2 , by merging the vertices in each connected component of G_1 . G_2 has one vertex per connected component of G_1 , and an edge for each edge in G that is still viable as the minimum maximum altitude. Then we recurse on G_2 . It only has edges corresponding to viable edges of G , so it is indeed half the size of G . We can compute G_2 from G in linear time: forming G_1 takes linear time; the connected components of G_1 are found in linear time with a breadth-first search; and figuring out what endpoints each edge should have in G_2 takes only a single pass through the viable edges of G .

This gives us our final procedure: we find the median edge-weight in linear time, and decide whether we can get from s to t using only the lighter edges. If we can, we delete all the heavier edges and recurse on a single subproblem half the size of the original. If we cannot, we merge vertices when they are connected by light edges, after which only heavy edges remain, and we again have a single subproblem of half the size. Pseudocode is given in Algorithm 3 on page 9. It is slightly more complex than this summary, as it deals with the possibility of repeated edge weights.

Correctness Correctness essentially follows from the above discussion; we give a more formal argument here. Correctness of MINWEIGHTPATH follows directly from correctness of MINWEIGHTPATHCONNECTED, so we will just establish the latter.

To prove correctness of MINWEIGHTPATHCONNECTED, we do induction on m . The m -th assertion is that the algorithm is correct for all graphs G with at most m edges. In the base case, there is one edge. Since s and t are distinct, but connected, they are connected by the unique remaining edge. Therefore, the algorithm correctly returns the weight of that edge.

For the inductive step, $m > 1$. Let τ be as in the algorithm. Every path connecting s and t by edges with weight less than τ is a path in G_1 , and every path connecting s and t by edges with weight at most τ is a path in \hat{G}_1 . We now break into three cases:

Case 1. There exists a path connecting s and t for which every edge has weight less than τ . In this case, the path continues to exist in G_1 , and the algorithm goes to line 18, where it constructs and recurses on G_1' . By construction, G_1' is connected. Moreover, every path from s to t with maximum edge weight less than τ remains in G_1' . Finally, at least one edge is removed from G to form G_1 (hence also G_1'), so we can apply the inductive hypothesis to conclude the recursive call is correct. This implies correctness in Case 1.

Case 2. There exists a path connecting s and t with maximum weight exactly τ , but there are no paths with smaller maximum weight. In this case, s and t are not connected in G_1 , but are connected in \hat{G}_1 , and the algorithm correctly returns τ on line 21.

Case 3. All paths connecting s and t have maximum-weight larger than τ . In this case, s and t are disconnected in both G_1 and \hat{G}_1 , so the algorithm goes to line 23, where it constructs G_2 and recurses to find a max-weight path from c_s to c_t . Since s and t are not connected in \hat{G}_1 , c_s and c_t are distinct vertices of G_2 . Every path in G gives rise to a similar path in

Algorithm 3

Input: Weighted graph G with m edges, distinct vertices s, t

Output: Minimum τ such that there is a path from s to t using only edges with weight at most τ , or indicate that no such path exists.

```
1: procedure MINWEIGHTPATH( $G, s, t$ )
2:   Compute connected components of  $G$ 
3:   if  $s = t$  then
4:     return no need to walk anywhere (max altitude is undefined)
5:   else if  $s$  and  $t$  are not connected then
6:     return impossible to walk from  $s$  to  $t$ 
7:   else
8:      $G' \leftarrow$  connected component containing  $s$  and  $t$ 
9:     return MINWEIGHTPATHCONNECTED( $G', s, t$ )
Input: Connected, weighted graph  $G$  with  $m$  edges, distinct vertices  $s, t$ 
Output: Minimum  $\tau$  such that there is a path from  $s$  to  $t$  using only edges with weight at most  $\tau$ 
10: procedure MINWEIGHTPATHCONNECTED( $G, s, t$ )
11:   if  $m = 1$  then
12:     return the weight of the edge in  $G$ 
13:   else
14:      $W \leftarrow$  list of all weights of edges in  $G$ 
15:      $\tau \leftarrow$  FAST-SELECT( $W, \lfloor m/2 \rfloor$ )
16:      $\hat{G}_1 \leftarrow$  copy of  $G$  with all edges of weight  $\geq \tau$  removed
17:      $G_1 \leftarrow$  copy of  $G$  with all edges of weight  $> \tau$  removed
18:     Compute connected components of  $G_1$  and  $\hat{G}_1$ 
19:     if  $s$  and  $t$  are connected in  $G_1$  then
20:        $G_1' \leftarrow$  connected component of  $G_1$  containing  $s$  and  $t$ 
21:       return MINWEIGHTPATHCONNECTED( $G_1', s, t$ )
22:     else if  $s$  and  $t$  are connected in  $\hat{G}_1$  then
23:       return  $\tau$ 
24:     else
25:        $G_2 \leftarrow$  new graph with one vertex per connected component of  $\hat{G}_1$ 
26:       for each edge  $\{u, v\}$  in  $G$  with weight  $u > \tau$  do
27:         Add an edge to  $G_2$  between the components of  $u$  and  $v$  with weight  $w$ 
28:        $c_s, c_t \leftarrow$  components of  $s, t$ , respectively
29:       return MINWEIGHTPATHCONNECTED( $G_2, c_s, c_t$ )
```

G_2 by following the edges with weight larger than τ . In particular, G_2 is connected, and the smallest maximum weight of a path from s to t in G_2 is the same as the smallest maximum weight of a path in G . Moreover, \hat{G}_1 contains at least one edge, so G_2 has fewer edges than G . Therefore the induction hypothesis applies to the recursive call on line 27, and this implies the algorithm is correct.

As this handles all the cases, the induction step is complete, and correctness for all inputs follows by induction.

Running time analysis We start with MINWEIGHTPATHCONNECTED. We use the recursion tree method. The recursion tree is shaped as a line. We measure the input size by the number of edges in G . At the root, this is m . From one node to the next, the number of edges shrinks by at least half. The work local to a node consists of a linear-time selection routine, a few breadth-first searches, and some elementary linear-time operations. Thus for some constant c , each node of the recursion tree working on ℓ edges does at most $c\ell$ work. Adding up all this work, we get

$$cm + c\frac{m}{2} + \dots + c\frac{m}{2^{\log m}} \leq cm \cdot \left(1 + \frac{1}{2} + \dots\right) = 2cm = O(m)$$

As for MINWEIGHTPATH, computing connected components takes $O(n+m)$ time, and the call to MINWEIGHTPATHCONNECTED takes $O(m)$ time. Combined, this is $O(n+m)$ time.