

DAT038/TDA417

Data structures and algorithms

Written examination

Thursday, 2021-01-14, 14:00–18:00, in Canvas and Zoom

- Examiner(s)** Peter Ljunglöf, Nick Smallbone, Christian Sattler.
- Instructions** Read all instructions on the front page of the Canvas exam room.
- Allowed aids** Course literature, other books, the internet.
You are not allowed to discuss the problems with anyone!
- Individualised** In most questions you will be asked to work on an individual problem instance. You will get your individualised problem from a specific Canvas page.
- Submitting** Submit your answers as answers to Canvas quizzes.
There is one Canvas quiz per question.
- Pen & paper** Some questions are easier to answer using pen and paper, which is fine!
In that case, take a photo of your answer and upload to your quiz answer.
Make sure the photo is readable!
- Assorted notes** You may answer in English or Swedish.
Excessively complicated answers might be rejected.
We need to be able to read and understand your answer!
- Exam review** If you want to discuss the grading, please contact Peter via email.

There are 6 basic, and 3 advanced questions, and two points per question. So, the highest possible mark is 18 in total (12 from the basic and 6 from the advanced). Here is what you need for each grade:

- To **pass** the exam, you must get 8 out of 12 on the basic questions.
- To get a **four**, you must get 9 out of 12 on the basic questions, plus 2 out of 6 on the advanced.
- To get a **five**, you must get 10 out of 12 on the basic questions, plus 4 out of 6 on the advanced.

Grade	Total points	Basic points	Advanced
3	≥ 8	≥ 8	—
4	≥ 11	≥ 9	≥ 2
5	≥ 14	≥ 10	≥ 4

Question 1: Complexity

Get a program and a part B question from here (“Q1: complexity”):

<https://chalmers.instructure.com/courses/13943/pages/individualised-problem-instances>

Part A

- What is the worst-case time complexity of the above program?
- Justify that the program has this complexity.

Note: Your complexity should be as simple as possible, and as low as possible.

Part B

Get a part B question from the link above.

- What is your answer to the question?
- Justify your answer.

Question 2: Sorting

In this question, you will show how the *partitioning* algorithm in quicksort works.

Start by getting an individualised array from here (“Q2: Sorting”):

<https://chalmers.instructure.com/courses/13943/pages/individualised-problem-instances>

Part A

Take the array you got from the link above, and partition it, using the *median-of-three* strategy to choose the pivot. (Do not make any recursive calls to quicksort.) Your answer should show:

- The sequence of swaps that the partitioning algorithm performs.
- What the array looks like after partitioning.
- Which parts of the array are the two partitions.

You can indicate the partitions using square brackets like this: [11 22 33] 44 [55 77 66].

If you use a different partitioning algorithm from the one in the course, write down what algorithm you used (either its name or a description).

Part B

Now go back to the array you were given at the start of the question (*not* your answer to Part A).

Swap two array elements, so that afterwards the array has the following property: if we partition it using the median-of-three strategy, one of the partitions will only contain one element.

Note: you are only allowed to swap two elements, not more.

Write down which elements you swapped, and briefly explain why partitioning using the median-of-three strategy will produce a partition containing only one element.

Hint: Think about which element must be chosen as the pivot for this to happen.

Question 3: Search trees

In this question you will show how insertion works in AVL trees.

Start by getting an individualised AVL tree from here (“Q3: Search trees”):

<https://chalmers.instructure.com/courses/13943/pages/individualised-problem-instances>

Part A

Using the AVL insertion algorithm, insert an element into the AVL tree, so that:

- During insertion, **the tree becomes unbalanced**, and the AVL insertion algorithm rebalances it
- The node that gets rebalanced is either a **left-left or right-right case**

You choose which element to insert, **but it must create either a left-left or right-right unbalanced tree**. You should then use the AVL insertion algorithm to rebalance the tree.

Your answer should show:

- What element you inserted
- The unbalanced tree after the element is added (in other words, the tree that would have been created if you just used ordinary BST insertion)
- Which node the AVL insertion algorithm needs to rebalance, and which kind of imbalance it has (left-left or right-right)
- What rotations the algorithm performed to rebalance the tree
- The final balanced tree

If you have trouble scanning your answers, you can instead write your tree as a nested list (but pictures are fine too!). The individualised AVL tree you got comes with an example of this.

Part B

Repeat part A, but this time **create either a left-right or right-left unbalanced tree**. Start from the individualised AVL tree you got at the beginning, *not* the tree you created in Part A.

Just the same as Part A, write down:

- What element you inserted
- The unbalanced tree after the element is added (in other words, the tree that would have been created if you just used ordinary BST insertion)
- Which node the AVL insertion algorithm needs to rebalance, and which kind of imbalance it has (left-right or right-left)
- What rotations the algorithm performed to rebalance the tree
- The final balanced tree

Question 4: Priority queues

In this question, you will implement an algorithm that takes as input two priority queues, and returns a *sorted list* containing all the elements of both priority queues, in ascending order. For example, if the algorithm is called on the inputs $\{1,2,3,5\}$ and $\{1,2,4\}$, it should return a list containing $\{1,1,2,2,3,4,5\}$.

Start by getting an unfinished implementation of the algorithm from here (“Q4: Priority queues”): <https://chalmers.instructure.com/courses/13943/pages/individualised-problem-instances>

The code defines a method `mergePQs` which takes two `PriorityQueues` as input, and returns its result in a dynamic array. It’s unfinished – the main loop of the method currently just reads “???”. Read the implementation and make sure you understand what it does so far.

To keep the code simple, the method `mergePQs` only works on priority queues containing integers.

Part A

Complete the code of `mergePQs`, by filling in the part labelled `???`. After `mergePQs` has finished running, it should return the combined contents of both input priority queues, *in ascending order*. You are allowed to modify the priority queues, for example by adding or removing elements.

Hint: remove one item in each iteration of the while-loop. Remember to handle the case when one of the priority queues is empty.

In your solution, you are *only allowed to call the following methods* from the [course API](#):

- `PriorityQueue` interface: `add`, `getMin`, `removeMin`, `isEmpty`, `size`
- `List` interface: `addFirst`, `addLast`

When you call these methods, *you do not need to write out their implementation*. In addition, you can use ordinary integer operations such as `+` and `<` on `Int`.

Give your answer in *pseudocode* or a suitable programming language. An example of pseudocode can be seen in question 1.

Part B

What is the worst-case complexity of `mergePQs`, in big-O notation? Assume that both priority queues have a size of n , and are implemented using binary heaps. Briefly justify your answer.

Note: Your complexity should be as simple as possible, and as low as possible.

Question 5: Hash tables

Get an individualised hash table and hash values from here (“Q5: hash tables”):

<https://chalmers.instructure.com/courses/13943/pages/individualised-problem-instances>

Part A

Assume that the open addressing hash table you got has never been resized, and that no elements have been deleted from it. Also assume that we use the hash value modulo the array size to get the array index, and standard linear probing for conflict resolution (i.e., probe by increasing by one).

- In which order can the elements have been added to the hash table, under the conditions above? It’s enough that you give one possible answer.
- Explain what hash collisions there were, and how they were resolved.


Part B

Now, add the same elements in the order from part A, to a *separate chaining hash table of size 5*. Once again, use modulo the array size to get the array index. Assume that you use linked lists as the container.

- How will the hash table look like when all elements have been added?
- Explain how elements are added to the linked list (in the front or back or somewhere else), and what conflicts you got when adding the elements to the hash table.

Question 6: Graphs

Assume the following undirected graph:









This graph is actually weighted, so get your individualised edge weights from here (“Q6: graphs”): <https://chalmers.instructure.com/courses/13943/pages/individualised-problem-instances>

Part A

Classify each of the following subgraphs A–F of the graph above as:


- **MST**: a minimum spanning tree
- **SPT from S**: not an MST, but a shortest path tree starting from node S
- **ST**: neither an MST nor a SPT from S, but a spanning tree
- **None**: not a spanning tree at all



For each of the graphs, briefly justify why it is classified as MST, SPT, ST or None.

Part B

Assume that you are running Prim's or Kruskal's algorithm on the graph, and after some iterations you have reached the following state (where the bold edges are the ones already added to the final MST):



Which edge will be added next to the MST by each of the algorithms?

- by Prim's algorithm? (Assume that the search started in node S)
- by Kruskal's algorithm?

Briefly justify your answers.

Advanced question 7: Complexity

Get an individualised two-variable orders of growth A, B, C, D, E from here ("Q7: Complexity II"):
<https://chalmers.instructure.com/courses/13943/pages/individualised-problem-instances>

Part A

Simplify the given orders of growth. Use O-notation. Justify your answer.

Part B

Draw, or list the edges of, a directed graph that describes the relationships between the given orders of growth: there should be a path from $O(X)$ to $O(Y)$ exactly if X is $O(Y)$. Justify why these relations hold and no other.

Advanced question 8: Cyclic order binary search

Let A be an array of $n \geq 1$ integers. Assume that A is in *strict cyclic order*: $A[i] < A[(i + 1) \bmod n]$ for all indices i except one.

(Note that this implies that all elements of A are distinct.)

Part A

Find an algorithm that computes the index of the minimum of A and runs in time $O(\log n)$.

```
int findMinIndex(int[] A)
```

Part B

Find an algorithm that checks if A contains an element v and runs in time $O(\log n)$:

```
boolean contains(int[] A, int v)
```

Give your algorithms in pseudocode or a suitable programming language.

Justify why your algorithms work, and why they have these complexities.

Advanced question 9: Bidirectional search

Modify the UCS/Dijkstra algorithm so that it can do *bidirectional search*:

- “Bidirectional search is a graph search algorithm that finds a shortest path from an initial vertex to a goal vertex in a directed graph. It runs two simultaneous searches: one forward from the initial state, and one backward from the goal, stopping when the two meet. The reason for this approach is that in many cases it is faster”
(https://en.wikipedia.org/wiki/Bidirectional_search)

Here’s a formulation of the UCS algorithm, similar to the one from the lecture slides:

```
List<Node> searchUCS(Node start, Node goal):
    pqueue = new PriorityQueue(comparing (e) -> e.costToHere)
    pqueue.add(new PQEntry(start, 0, null))
    Set<Node> visited = new Set()

    while pqueue is not empty:
        PQEntry entry = pqueue.removeMin()

        if entry.node == goal:
            return extractPath(entry)

        if entry.node not in visited:
            visited.add(entry.node)
            for Edge edge in graph.outgoingEdges(entry.node):
                costToNext = entry.costToHere + edge.weight
                pqueue.add(new PQEntry(edge.to, costToNext, entry))

    return null // Failure
```

It uses an auxiliary class of priority queue entries, and a method that extracts a path from the start node to the current:

```
class PQEntry:
    Node node
    double costToHere
    PQEntry backPointer

List<Node> extractPath(PQEntry entry):
    ... // you don't have to know how this method works
```

The graph API is the following:

```
class Edge<Node>:
    Node from
    Node to
    double weight

class Graph<Node>:
    Collection<Edge<Node>> outgoingEdges(Node from)
    Graph<Node> reversedCopy()
```

This graph API is very similar to the one from lab 4 or the lecture slides, with one important addition: the `reversedCopy()` method. This will create a new copy of the graph with the directions of all edges reversed (so that from becomes to and to becomes from). The existing graph will not be changed by calling this method – it creates a new graph instead.

Hints:

- Don't try to use concurrency or parallelisation. Instead alternate between the forward and the backward search – remove from the first, then from the second, then from the first, then from the second, etc.
- Think about how you can make use of the “active” nodes – the nodes that are currently in the priority queue, but have not yet been visited.
- You can use any class or method from [the course API](#) or the Java standard API (`java.lang.*` or `java.util.*`).

Give your algorithm in pseudocode or a suitable programming language.

Explain how your algorithm works.

Side note: This simple version of bidirectional search does not guarantee that it will return the optimal solution, as explained in footnote 1 of the Wikipedia article. There are some additional things you have to do to get an optimal algorithm, but your solution will return an almost-optimal solution in most cases.

DAT038/TDA417

Data structures and algorithms

Written examination

Solution suggestions

Thursday, 2021-01-14, 14:00–18:00, in Canvas and Zoom

Examiner(s) Peter Ljunglöf, Nick Smallbone, Christian Sattler.

Allowed aids Course literature, other books, the internet.

You are not allowed to discuss the problems with anyone!

Individualised In most questions you will be asked to work on an individual problem instance. You will get your individualised problem from a specific Canvas page.

The solutions in this document are for one specific problem instance!

Assorted notes You may answer in English or Swedish.

Excessively complicated answers might be rejected.

We need to be able to read and understand your answer!

Exam review If you want to discuss the grading, please contact Peter via email.

There are 6 basic, and 3 advanced questions, and two points per question. So, the highest possible mark is 18 in total (12 from the basic and 6 from the advanced). Here is what you need for each grade:

- To **pass** the exam, you must get 8 out of 12 on the basic questions.
- To get a **four**, you must get 9 out of 12 on the basic questions, plus 2 out of 6 on the advanced.
- To get a **five**, you must get 10 out of 12 on the basic questions, plus 4 out of 6 on the advanced.

Grade	Total points	Basic points	Advanced
3	≥ 8	≥ 8	—
4	≥ 11	≥ 9	≥ 2
5	≥ 14	≥ 10	≥ 4

Question 1

This is how your individual program could have looked like:

```
boolean f(int[] xs):
    ys = new dynamic array of int
    for i in 0 .. xs.length-1:
        for j in i .. xs.length-1:
            ys.add(xs[i] + xs[j])

    boolean r = false
    for m in ys:
        for n in ys:
            if m + n == 0:
                r = true
    return r
```

And this is a possible individual question:

When you run the program on your computer in 2021 on an array of size 1000, it takes one day. Your computer has gotten faster by a factor of 4 every 9 years. What is the largest array size you could have processed in one day in 1967?

Part A solution

- The first group of for-loops has complexity $O(n) \cdot O(n) \cdot O(1) = O(n^2)$: each loop has $O(n)$ iterations and the body has amortized complexity $O(1)$.
- After it has executed, **ys** has $O(n) \cdot O(n) = O(n^2)$ elements.
- With this information, the second group of for-loops has complexity $O(n^2) \cdot O(n^2) \cdot O(1) = O(n^4)$ (same reasoning as before).
- Altogether, the complexity is $O(n^2) + O(n^4) = O(n^4)$.

Part B solution

- We have $(2021 - 1967) / 9 = 54 / 9 = 6$, so my computer is $4^6 = 2^{12}$ times faster in 2021 than it was in 1967. So one day in 2021 corresponds to 2^{12} times as many computation steps than one day in 1967.
- Let $T(n)$ denote the run time in terms of number of computation steps on an array of size n . From Part A, we have $T(n) \approx C n^4$ for some constant C .
- We want to find n such that $T(1000) = 2^{12} T(n)$. This means $1000^4 \approx 2^{12} n^4$, i.e., $1000 \approx 2^3 n = 8 n$.
- So the answer is $n \approx 125$.

Question 2

Here is a possible individual input array:

23 36 82 91 77 45 19 63 50

Part A solution

- We use the partitioning algorithm from the lectures here, but there are several possible algorithms that all are correct.
- The median of (23, 77, 50) = 50 so we start by swapping 23 and 50.
- The following table shows the steps in the partitioning. The numbers in **red bold** are the numbers that are going to be swapped at each stage.


50	36	82	91	77	45	19	63	23
50	36	23	91	77	45	19	63	82
50	36	23	19	77	45	91	63	82
50	36	23	19	45	77	91	63	82
[45	36	23	19]	50	[77	91	63	82]

Part B solution

- Swap 19 with 50, or 19 with 77
- Then the three median elements will be (19, 23, 50) or (19, 23, 77), and the median of three will be 23. This is the second-smallest element of the array, and then there will be one partition with only one element, 19.


Question 3

Here is one possible individual tree:




Part A solution

1) To create a left-left tree, we can insert a node as a **child of 50**. E.g. inserting 45 (but e.g. 55 also works):




2) Now, **node 80 is a left-left tree**. We solve it by **rotating 80 right**, to get:




Part B solution

1) To create a left-right tree, we can insert a node as a **child of 70**. E.g. inserting 75 (but e.g. 65 also works):



2) Now, **node 80 is a left-right tree**. We solve it by **first rotating 60 left**, then **rotating 80 right**, to get:



Question 4

Here's a possible individual skeleton code:

```
List<Int> mergePQs(PriorityQueue<Int> pq1, PriorityQueue<Int> pq2):
    List<Int> result = new DynamicArrayList<Int>()
    while not (pq1.isEmpty() and pq2.isEmpty()):
        ???
    return result
```

Part A solution

Here is one possible solution:

```
List<Int> mergePQs(PriorityQueue<Int> pq1, PriorityQueue<Int> pq2):
    List<Int> result = new DynamicArrayList<Int>()
    while not (pq1.isEmpty() and pq2.isEmpty()):
        if pq1.isEmpty():
            // pq2 must be non-empty
            result.add(pq2.removeMin())
        else if pq2.isEmpty():
            // pq1 must be non-empty
            result.add(pq1.removeMin())
        else:
            // neither pq is empty
            if pq1.getMin() <= pq2.getMin():
                result.add(pq1.removeMin())
            else:
                result.add(pq2.removeMin())
    return result
```

But any solution is allowed as long as it is correct and doesn't use extra data structures/library functions.

Part B solution

- The while-loop runs $O(n)$ times taking $O(\log n)$ time per iteration.
 - `PriorityQueue.removeMin` is $O(\log n)$
 - `DynamicArrayList.add` is $O(1)$ (amortised)
- So the total complexity is $O(n \log n)$.

Question 5

Here is one possible individual instance:

[J, M, -, Q, K, N, -, L, P]

And with individual hash values:

hash(J) = 9
hash(K) = 13
hash(L) = 16
hash(M) = 8
hash(N) = 3
hash(P) = 7
hash(Q) = 12

Part A solution

- This gives the following table indices, after mod 9, and conflict resolution:

Indices 3-5:	Q = 3
	K = 4
	N = 3 --> 4 --> 5
Indices 7-8, 0-1:	J = 0
	L = 7
	P = 7 --> 8
	M = 8 --> 0 --> 1

- Which in turn gives the following constraints on the insertion order:

$L < P < M, Q < N, J < M, K < N$

- So one possible insertion order is: L P J M Q K N

Part B solution

- The hash values will give the following table indices (after mod 5):

$L = 1, P = 2, J = 4, M = 3, Q = 2, K = 3, N = 3$

- So the final hash table will look like this (if we add elements in the insertion order, and always add elements to the front of the linked list):

0: -
1: L
2: Q --> P
3: N --> K --> M
4: J


Question 6

Here is one possible instance of the graph:

Edge weights:

QT:11, QX:13, PT:6, PX:5, PS:7, PZ:3, SX:1, SZ:4, RZ:10, RU:12, RY:8, UY:9

Or equivalently, inserted into the graph:



Part A solution

- A is a spanning tree
- B is not a spanning tree, because it is not connected
- C is a shortest path tree (SPT) from node S
- D is a spanning tree
- E is a minimum spanning tree (MST)

Part B solution

Prim's algorithm:

- will add **edge RZ** with weight 10
- (the minimum-weight edge that has exactly one node in the current MST)

Kruskal's algorithm:

- will add **edge RY** with weight 8
- (the minimum-weight edge that will not create a cycle when added)

Question 7

Here is one possible set of order-of-growth functions:

- $A(n, k) = O((n + 18 \sqrt{n}) (1 + 18 \sqrt{k}))$.
- $B(n, k)$ is the complexity of the below algorithm for cs of size n and j equal to 0:

```
boolean has_sum(int[] cs, int j, int y):  
    if j == cs.length: return y == 0  
    return has_sum(cs, j + 1, y) or has_sum(cs, j + 1, y - cs[j])
```
- $C(n, k)$ is the complexity of binary search on a sorted array of size n k .
- $D(n, k) = O((n + n k + k)^{12})$.
- $E(n, k)$ is the complexity of the below algorithm for set of size n and ys of size k :

```
void track_large(RedBlackSet<Int> set, int[] ys):  
    for y in ys:  
        set.add(y)  
        set.remove(set.min())
```

Part A solution

- **$A(n, k) = O(n \sqrt{k})$**
because only the asymptotically biggest summand matters and constant factors don't matter (and $O(-)$ preserves products).
- **$B(n, k) = O(2^n)$**
because the program considers all possible subsets of xs to sum, of which there are 2^n . Everything outside the recursive structure is constant. k does not play a role.
- **$C(n, k) = O(\log(n k)) = O(\log(n) + \log(k))$**
because binary search is logarithmic in the array size.
- **$D(n, k) = O((n k)^{12}) = O(n^{12} k^{12})$**
because only the asymptotically biggest summand matters (and $O(-)$ preserves products).
- **$E(n, k) = O(\log(n) k)$**
because k times three operations logarithmic in n are performed.

Part B solution

A graph with the edges $C \rightarrow A \rightarrow D$ and $C \rightarrow E \rightarrow D$, and node B is isolated. We have:

- $C \leq A$, because $O(\log(n)), O(\log(k)) \leq O(n \sqrt{k})$.
- $C \leq E$, because $O(\log(n)), O(\log(k)) \leq O(\log(n) k)$.
- $A \leq D$, because $O(n) \leq O(n^{12})$ and $O(\sqrt{k}) \leq O(k^{12})$.
- $E \leq D$, because $O(\log(n)) \leq O(n^{12})$ and $O(k) \leq O(k^{12})$.
- A is not related to E because, for $n = 1$, A is strictly smaller than E , and for $k = 1$, it is strictly larger than E .
- B is not related to any of A, C, D, E because for $n = 1$ it is strictly smaller than the rest and for $k = 1$ it is strictly larger than the rest.

Question 8

Here is an example of an array in strict cyclic order: [53, 62, 81, 94, 15, 22, 43]

Part A solution

Calling `findMinIndex` on the example array above should return 4 because this is the index of the smallest element (15). There are several possible solutions, but all are variants of binary search. The main idea is to select the interval where the first element is larger than the last – because the smallest element must be in that interval:

```
int findMinIndex(int[] A):
    // First check if the array is ordered.
    if A[0] ≤ A[A.length - 1]:
        return 0
    // Invariant: left < right, A[left] > A[right]
    int left = 0
    int right = A.length - 1
    // Goal: right = left + 1. Then right is the index of the minimum.
    while right - left > 1:
        mid = (left + right) / 2 // Rounding doesn't matter.
        if A[left] > A[mid]:      // A[left] > A[mid] < A[right]
            right = mid
        else:                    // A[left] < A[mid] > A[right]
            left = mid
    return right
```

Note that we have to handle the special case when the array is already ordered, i.e., when the first element is the smallest.

Part B solution

First we find the index of the smallest element using part A. Then we can compare the searched element v with the first array element $A[0]$ to decide which part of the array we have to search in. This part is then guaranteed to be sorted, so we can use standard binary search to look for v . Like this:

```
boolean contains(int[] A, int v):
    orderStart = findMinIndex(A)
    if v ≥ A[0]:
        return binarySearch(A, 0, orderStart - 1)
    else:
        return binarySearch(A, orderStart, A.length - 1)
```

There are many other possibilities, but all involve binary search in some way. E.g.:

- It's possible to tweak the part A solution directly instead of calling it.
- Binary search in the *virtual list* of same size as A that sends index i to $A[(orderStart + i) \bmod A.length]$

Question 9

One solution idea is to swap the graphs after each iteration. Like this:

```
List<Node> searchBi(Node start, Node goal):
    pqueue = new PriorityQueue(comparing (e) -> e.costToHere)
    pqueue.add(new PQEntry(start, 0, null))
    Map<Node,PQEntry> visited = new Map()

    graphR = graph.reverseCopy()
    pqueueR = new PriorityQueue(comparing (e) -> e.costToHere)
    pqueueR.add(new PQEntry(goal, 0, null))
    Map<Node,PQEntry> visitedR = new Map()

    while pqueue and pqueueR are not empty:
        PQEntry entry = pqueue.removeMin()

        if entry.node in visitedR:
            // now we've found a solution
            entryR = visitedR.get(entry.node)
            return extractPath(entry) + reverse(extractPath(entryR))

        if entry.node not in visited:
            visited.set(entry.node, entry)
            for Edge edge in graph.outgoingEdges(entry.node):
                costToNext = entry.costToHere + edge.weight
                pqueue.add(new PQEntry(edge.to, costToNext, entry))

        swap pqueue and pqueueR
        swap visited and visitedR

    // if we end here we have failed
    return null
```

Note that we change the goal test, from checking if the current node is the goal, to checking if the current node has been visited by the reverse search.

To extract the solution path we just concatenate the two solution paths – we just have to remember to reverse one of them. To be able to extract the path for the reverse search, we have to find the PQEntry that led to the node, and therefore we use a map for the visited nodes, instead of a set as in the original one-directional search.

One minor problem that this solution doesn't handle is that the extracted path might be reversed (from the goal to the start). This can be handled in many ways – one is to keep a boolean flip-flop variable, saying if the current pqueue is the original or the reversed.