

B16 Algorithms and Data Structures 1 - Example Sheet

Andrea Vedaldi

Academic Year 2023-24 (version 2.0)

Contents

Introduction	2
1 Linear-time sorting	2
1.1 Radix sort	2
2 Arrays	2
2.1 Complexity of <code>ArrayInsert</code>	2
2.2 Writing <code>ArrayDelete</code>	2
3 Stacks	3
3.1 Enhanced stack class	3
3.2 Reverse Polish calculator	4
4 Queues	5
4.1 Enhanced queue class	5
5 Lists	7
5.1 Deleting elements from a list	7
6 Binary trees	9
6.1 Enhanced <code>BinaryTree</code> class	9
6.2 Minimum and maximum element in a binary search tree	11
6.3 Traversing the nodes of a BST by non-decreasing value	12
7 Heaps	12
7.1 Building a heap	12
7.2 Building a heap vs sorting	12
7.3 Updating the priority of a queued element	12
8 Hashing	12
8.1 Worst case complexity for multiple chaining	12
8.2 Other usages of hash functions	13
8.3 Division method for large keys	13
8.4 Permutation invariance of the division method	13
9 Graphs	13
9.1 Shortest paths using the adjacency list representation	13
9.2 Decoding shortest paths	15
10 Further examples (optional)	17
10.1 Inserting movable objects in an array	17
10.2 Stack with move semantics	22

Introduction

This is the example sheet for B16 Part 3 Algorithms and Data Structures 1. The sheet involves completing and running several C++ programs, for which you should start from the code made available here.

1 Linear-time sorting

1.1 Radix sort

Consider non-negative integers of k binary digits each. We can represent each integer as a binary sequence $\langle c_{k-1}c_{k-2}\cdots c_0 \rangle$, where $c_i \in \{0, 1\}$ where c_{k-1} is the most significant digit.

Informally describe an algorithm that sort a sequence A of n such integers one digit per time, with an overall cost of $O(kn)$.

Optionally, write a C++ implementation `radix_sort.hpp` of this algorithm with an interface similar to `merge_sort.hpp` and test it using the following C++ driver program. Do this only if you have time as it is slightly tricky to get right.

File `radix_sort_driver.cpp`:

```

1  #include "radix_sort.hpp"
2  #include "utils.hpp"
3
4  int main(int argc, char **argv)
5  {
6      auto A = std::vector<int>{1, 19, 2, 9, 12, 18, 4, 8, 5, 6,
7                               17, 10, 11, 14, 16, 15, 7, 3, 13, 20};
8      print(A, "Before sorting: ");
9      radix_sort(begin(A), end(A));
10     print(A, "After sorting: ");
11 }

```

Before sorting: [1, 19, 2, 9, 12, 18, 4, 8, 5, 6, 17, 10, 11, 14, 16, 15, 7, 3, 13, 20]

After sorting: [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20]

2 Arrays

2.1 Complexity of ArrayInsert

What is the Big-O complexity of `ArrayInsert`?

2.2 Writing ArrayDelete

Design an algorithm `ArrayDelete` to delete the element at position i in an array A . What is the worst-case complexity of this algorithm?

Implement the algorithm as a C++ template function `void array_delete(std::vector<T>&, size_t)` by completing the file `array_delete.hpp`:

```

1  #ifndef __array_delete__
2  #define __array_delete__
3
4  #include <cstdint>

```

```

5  #include <utility>
6  #include <vector>
7
8  template <typename T>
9  void array_delete(std::vector<T> &A, std::size_t index)
10 {
11     // WRITE YOUR CODE HERE
12 }
13
14 #endif // __array_delete__

```

Use the following test driver `array_delete_driver.cpp` to test your code:

```

1  #include "array.hpp"
2  #include "array_delete.hpp"
3  #include "utils.hpp"
4
5  int main(int argc, char **argv)
6  {
7      auto A = std::vector<float>{0, 1, 2, 3, 4};
8      print(A, "Initially A = ");
9      while (!A.empty()) {
10         array_delete(A, 0);
11         print(A, "After deleting the element at position 0: A = ");
12     }
13     return 0;
14 }

```

Initially A = [0, 1, 2, 3, 4]

After deleting the element at position 0: A = [1, 2, 3, 4]

After deleting the element at position 0: A = [2, 3, 4]

After deleting the element at position 0: A = [3, 4]

After deleting the element at position 0: A = [4]

After deleting the element at position 0: A = []

3 Stacks

3.1 Enhanced stack class

Starting from the class `Stack` given in the notes, create a new stack class `StackEnhanced` with a member function `clear` to remove all elements from the stack. Also write an overload of `operator<<` so that elements can be pushed onto the stack via a notation like `stack << value1 << value2`. Write this into the `stack_enhanced.hpp` file, by completing as needed:

```

1  #ifndef __stack_enhanced__
2  #define __stack_enhanced__
3
4  #include "stack.hpp"
5
6  template <typename T> class StackEnhanced : public Stack<T>
7  {
8      public:
9         // Inherit the Stack<T> constructors as they are
10        // https://en.cppreference.com/w/cpp/language/using_declaration
11        using Stack<T>::Stack;

```

```

12
13     void clear()
14     {
15         // WRITE YOUR CODE HERE
16     }
17 };
18
19 template <typename T>
20 StackEnhanced<T> &operator<<(StackEnhanced<T> &stack, const T &value)
21 {
22     // WRITE YOUR CODE HERE
23 }
24
25 #endif // __stack_enhanced__

```

Use the following test driver to test your code.

File `stack_enhanced_driver.cpp`:

```

1  #include "stack.hpp"
2  #include "stack_enhanced.hpp" // Put your code in this file
3  #include <iostream>
4
5  int main(int argc, char **argv)
6  {
7      auto stack = StackEnhanced<int>(100);
8
9      stack << 1 << 2 << 3;
10     stack.clear();
11     stack << 4 << 5 << 6;
12
13     // Dump the stack content
14     std::cout << "Stack content:";
15     while (!stack.empty()) {
16         std::cout << ' ' << stack.top();
17         stack.pop();
18     }
19     std::cout << '\n';
20
21     return 0;
22 }

```

Stack content: 6 5 4

3.2 Reverse Polish calculator

The Reverse Polish Notation (RPN) for arithmetic expression lists first the operands and then the operation. For example, the expression $2 \times (2 + 3)$ is written `2 2 3 + ×`. In order to evaluate the expression, one reads it from left to right and immediately replaces any operation and its operands with the corresponding result. For instance, the expression above evaluates as:

$$\begin{array}{l}
 2 \ 2 \ 3 \ + \ \times \\
 2 \ 5 \ \times \\
 10
 \end{array}$$

A stack is the ideal data structure to store the intermediate results as an RPN expression is evaluated, because each operation always involves the last n values computed (i.e., the ones at the top of the stack).

Write functions `plus`, `minus`, `multiplies`, `divides` and `negate` to implement a basic RPN calculator. Each function takes as input a stack containing the state of the calculation and immediately replaces the top one or two elements with the operation result. Use the following test code to check that your code works.

File `stack_rpn_driver.cpp`:

```
1  #include <iostream>
2
3  #include "stack.hpp"
4  #include "stack_rpn.hpp" // Put your code in this file
5
6  int main(int argc, char **argv)
7  {
8      // Basic interface
9      auto stack = Stack<int>(100);
10     stack.push(2);
11     stack.push(2);
12     stack.push(3);
13     plus(stack);
14     multiplies(stack);
15     std::cout << "2 2 2 + * = " << stack.top() << '\n';
16
17     // Advanced interface (optional)
18     stack << 2 << 2 << 3 << plus << multiplies;
19     std::cout << "2 2 2 + * = " << stack.top() << '\n';
20
21     return 0;
22 }
```

```
2 2 2 + * = 10
2 2 2 + * = 10
```

Optionally, also implement the “advanced interface” used in the driver above, where operand and operations can be entered by using the `<<` symbol (hint: overload `operator<<`).

4 Queues

4.1 Enhanced queue class

Starting from the class `Queue` given in the notes, create a new stack class `Dequeue` (double-ended queue) with:

1. A member function `clear` to remove all elements from the queue.
2. A member function `enqueue_front` to enqueue elements at the front of the queue.
3. A member function `dequeue_back` to remove elements from the front of the queue.
4. A member function `back` to access the first element of the queue.

You can start from the following `queue_enhanced.hpp` file:

```
1  #ifndef __queue_enhanced__
2  #define __queue_enhanced__
3
4  #include <cassert>
5  #include <vector>
6
```

```

7  #include "queue.hpp"
8
9  template <typename T> class Dequeue : public Queue<T>
10 {
11 public:
12     // Inherit the constructors of Queue<T>
13     using Queue<T>::Queue;
14
15     // Access the element at the back of the queue
16     T &back()
17     {
18         // WRITE YOUR CODE HERE
19     }
20
21     // Const-access the element at the back of the queue
22     const T &back() const
23     {
24         // WRITE YOUR CODE HERE
25     }
26
27     // Add a new element to the front of the queue by copying
28     void enqueue_front(const T &value)
29     {
30         // WRITE YOUR CODE HERE
31     }
32
33     // Remove the element at the back of the queue
34     void dequeue_back()
35     {
36         // WRITE YOUR CODE HERE
37     }
38
39     // Remove all elements from the queue
40     void clear() {
41         // WRITE YOUR CODE HERE
42     }
43
44 protected:
45     // Return the index of the element at the back of the queue
46     size_t _tail() const
47     {
48         // WRITE YOUR CODE HERE
49     }
50 };
51
52 #endif // __queue_enhanced__

```

Use the following test driver to test your code.

File `queue_enhanced_driver.cpp`:

```

1  #include "queue_enhanced.hpp"
2  #include "utils.hpp"
3
4  #include <iostream>

```

```

5
6 int main(int argc, char **argv)
7 {
8     // Create a queue with space for a few elements
9     auto queue = Dequeue<float>(5);
10
11     // Keep pushing and popping elements from the dequeue for a while
12     for (int repetition = 0; repetition < 3; ++repetition) {
13         std::cout << "Enqueued front";
14         for (int i = 0; i < 3; ++i) {
15             queue.enqueue_front(i);
16             std::cout << ' ' << i;
17         }
18         std::cout << "\nDequeued front";
19         for (int i = 0; i < 3; ++i) {
20             std::cout << ' ' << queue.front();
21             queue.dequeue();
22         }
23         std::cout << "\nEnqueued back";
24         for (int i = 0; i < 3; ++i) {
25             queue.enqueue(i);
26             std::cout << ' ' << i;
27         }
28         std::cout << "\nDequeued back";
29         for (int i = 0; i < 3; ++i) {
30             std::cout << ' ' << queue.back();
31             queue.dequeue_back();
32         }
33         std::cout << '\n';
34     }
35
36     return 0;
37 }

```

```

Enqueued front 0 1 2
Dequeued front 2 1 0
Enqueued back 0 1 2
Dequeued back 2 1 0
Enqueued front 0 1 2
Dequeued front 2 1 0
Enqueued back 0 1 2
Dequeued back 2 1 0
Enqueued front 0 1 2
Dequeued front 2 1 0
Enqueued back 0 1 2
Dequeued back 2 1 0

```

5 Lists

5.1 Deleting elements from a list

Starting from the module `list.hpp` given in the notes, write a `list_delete_after` function adding it to a new `list_enhanced.hpp` module, completing the code below:

```

1
2 #ifndef __list_enhanced__
3 #define __list_enhanced__
4
5 #include "list.hpp"
6 #include <iostream>
7
8 template <typename T> void list_delete_after(Node<T> *node)
9 {
10     // WRITE YOUR CODE HERE
11 }
12
13 #endif // __list_enhanced__

```

Then test it using the following driver.

File `list_enhanced_driver.cpp`:

```

1 #include "list_enhanced.hpp"
2 #include "utils.hpp"
3
4 int main(int argc, char **argv)
5 {
6     auto list = Node<float>{};
7
8     // Insert some numbers to the front of the list.
9     auto last = &list;
10    for (int i = 0; i < 10; ++i) {
11        last = list_insert_after(last, static_cast<float>(i));
12        print(list_to_vector(list));
13    }
14
15    // Remove the elements from the beginning of the list.
16    for (int i = 0; i < 10; ++i) {
17        list_delete_after(&list);
18        print(list_to_vector(list));
19    }
20 }

```

```

[0]
[0, 1]
[0, 1, 2]
[0, 1, 2, 3]
[0, 1, 2, 3, 4]
[0, 1, 2, 3, 4, 5]
[0, 1, 2, 3, 4, 5, 6]
[0, 1, 2, 3, 4, 5, 6, 7]
[0, 1, 2, 3, 4, 5, 6, 7, 8]
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
[1, 2, 3, 4, 5, 6, 7, 8, 9]
[2, 3, 4, 5, 6, 7, 8, 9]
[3, 4, 5, 6, 7, 8, 9]
[4, 5, 6, 7, 8, 9]
[5, 6, 7, 8, 9]
[6, 7, 8, 9]

```


[7, 8, 9]
[8, 9]
[9]
[]

What is the worst-case complexity of this algorithm?

6 Binary trees

6.1 Enhanced BinaryTree class

Similar to the BinaryTree class given in the notes, write an enhanced binary tree class with a member function `parent()` returning the parent of a node. You can start from the file `binary_tree_enhanced.hpp` file:

```
1  #ifndef __binary_tree_enhanced__
2  #define __binary_tree_enhanced__
3
4  #include <cassert>
5  #include <vector>
6  #include <memory>
7
8  // A class representing a binary tree
9  template <typename V> struct BinaryTreeEnhanced {
10     V _value;
11     std::unique_ptr<BinaryTreeEnhanced<V>> _left;
12     std::unique_ptr<BinaryTreeEnhanced<V>> _right;
13
14     // WRITE YOUR CODE HERE
15
16     friend V &value(BinaryTreeEnhanced *t) { return t->_value; }
17     friend const V &value(const BinaryTreeEnhanced *t)
18     {
19         return t->_left_value;
20     }
21     friend BinaryTreeEnhanced *left(const BinaryTreeEnhanced *t)
22     {
23         return t->_left.get();
24     }
25     friend BinaryTreeEnhanced *right(const BinaryTreeEnhanced *t)
26     {
27         return t->_right.get();
28     }
29     friend BinaryTreeEnhanced *parent(const BinaryTreeEnhanced *t)
30     {
31         // WRITE YOUR CODE HERE
32     }
33 };
34
35 // A helper function to build an enhanced binary tree
36 template <typename V>
37 std::unique_ptr<BinaryTreeEnhanced<V>>
38 make_binary_tree_enhanced(const V &value,
39                          std::unique_ptr<BinaryTreeEnhanced<V>> l,
40                          std::unique_ptr<BinaryTreeEnhanced<V>> r)
```

```

41 {
42     // WRITE YOUR CODE HERE
43 }
44
45 #endif // __binary_tree_enhanced__

```

Use the following test driver to test it.

File `binary_tree_enhanced_driver.cpp`:

```

1  #include "binary_tree_enhanced.hpp"
2  #include "binary_tree_print.hpp"
3  #include "binary_tree_traversal.hpp"
4
5  int main(int argc, char **argv)
6  {
7      auto bt = make_binary_tree_enhanced(
8          1.0f,
9          make_binary_tree_enhanced(
10             2.0f, make_binary_tree_enhanced(4.0f, {}, {}),
11             make_binary_tree_enhanced(
12                 5.0f, {}, make_binary_tree_enhanced(8.0f, {}, {}))),
13             make_binary_tree_enhanced(
14                 3.0f, make_binary_tree_enhanced(6.0f, {}, {}),
15                 make_binary_tree_enhanced(7.0f, {}, {})));
16
17     std::cout << "Tree:\n";
18     print_binary_tree(bt.get());
19
20     auto action = [](const auto &tree) {
21         auto p = parent(tree);
22         if (p) {
23             std::cout << "The parent of " << value(tree) << " is "
24                 << value(p) << '\n';
25         } else {
26             std::cout << "Node " << value(tree) << " has no parent\n";
27         }
28     };
29
30     df_traversal(bt.get(), action);
31
32     return 0;
33 }

```

Tree:

```

1 -----v
2 -v      3 -v
4 5 -v    6 7
      8

```

```

The parent of 4 is 2
The parent of 2 is 1
The parent of 5 is 2
The parent of 8 is 5
Node 1 has no parent
The parent of 6 is 3

```

The parent of 3 is 1
The parent of 7 is 3

6.2 Minimum and maximum element in a binary search tree

Write a module `binary_search_tree_enhanced.hpp` defining functions `bst_min` and `bst_max` computing the minimum and maximum element in a BST so that the following test driver works.

You can start from the interface file:

```
1  #ifndef __binary_saerch_tree_enhanced__
2  #define __binary_saerch_tree_enhanced__
3
4  #include "binary_search_tree.hpp"
5
6  template <typename T> T bst_min(const T &tree)
7  {
8      // WRITE YOUR CODE HERE
9  }
10
11 template <typename T> T bst_max(const T &tree)
12 {
13     // WRITE YOUR CODE HERE
14 }
15
16 #endif // __binary_saerch_tree_enhanced__
```

Test it using the following driver.

File `binary_search_tree_enhanced_driver.cpp`:

```
1  #include <iostream>
2
3  #include "binary_search_tree_enhanced.hpp"
4  #include "binary_tree_print.hpp"
5
6  int main(int argc, char **argv)
7  {
8      std::unique_ptr<BinaryTree<int>> bt;
9
10     for (int x : {12, 5, 18, 2, 9, 15, 19, 13, 17}) {
11         bt = bst_insert(std::move(bt), x);
12     }
13
14     std::cout << "Tree:";
15     print_binary_tree(bt.get());
16     std::cout << "\n";
17
18     std::cout << "The smallest element is "
19               << value(bst_min(bt.get())) << '\n';
20
21     std::cout << "The largest element is " << value(bst_max(bt.get()))
22               << '\n';
23
24     return 0;
25 }
```

```

Tree:12 ---v
5 -v 18 -----v
2 9 15 -v 19
    13 17

```

The smallest element is 2
The largest element is 19

6.3 Traversing the nodes of a BST by non-decreasing value

Design an algorithm to traverse the nodes of a BST by non-decreasing value. You can omit the C++ implementation.

7 Heaps

7.1 Building a heap

In the notes, we have discussed the `BuildHeap` algorithm, for building a heap given an initial vector A of values:

`BuildHeap(A)`:

1. For $i = \lfloor |A|/2 \rfloor - 1, \dots, 0$:
 1. Interpret the subarray $(A_i, \dots, A_{|A|-1})$ as a complete binary tree S .
 2. Call `SiftDown(S)`.

This algorithm builds the heap starting from the leaves of the binary tree and working its way towards the root. Alternatively, we can build a heap starting from the root and moving towards the leaves. The idea is to progressively include elements A_0, A_1 etc in the heap, which amounts to adding one more leaf to the corresponding binary tree. Each time this is done, one calls `SiftUp` on the last element added to restore the heap property. In pseudo-code:

`BuildHeapAlt(A)`:

1. For $i = 1, \dots, |A| - 1$:
 1. Interpret the subarray (A_0, \dots, A_i) as a complete binary tree T and let S be the subtree rooted at A_i .
 2. Call `SiftUp(S)`.

Discuss the relation between `BuildHeapAlt` and `PriorityEnqueue` and compare the complexity of `BuildHeap` and `BuildHeapAlt`.

7.2 Building a heap vs sorting

Discuss whether sorting an array A can be an alternative implementation of `BuildHeap`.

7.3 Updating the priority of a queued element

Discuss how you could use `SiftUp` and `SiftDown` to update the priority of an element in a heap, used as a priority queue.

8 Hashing

8.1 Worst case complexity for multiple chaining

Consider n distinct keys k_1, \dots, k_n drawn from a finite set \mathcal{K} , called the *universe*. We wish to store the keys in a hash table that uses m chains. Assume that the universe is large, meaning that $|\mathcal{K}| > m \cdot n$, and

show that it is possible to choose the n keys so that they all hash to the same chain. Discuss what is the implication on the worst case complexity of retrieving the keys from the resulting hash table.

8.2 Other usages of hash functions

Consider a file such as a photo or a video. Explain how a hash function can be used to create a compact “signature” of the file to verify its integrity. Specifically, explain how knowledge of the signature can be used to tell with high probability whether the file has been modified without requiring knowledge of the original file. Explain how you would choose the size m of the hash space for this application.

8.3 Division method for large keys

Consider the problem of designing a hash function $h(k)$. The division method consists in computing the remainder $h(k) = k \bmod m$ of dividing the key k , regarded as a large natural number, by the number m of chains. This is convenient because $h(k) \in [0, \dots, m - 1]$ can be used directly to index a chain.

Suppose that the key is a string $\langle c_{q-1}c_{q-2} \dots c_0 \rangle$ where $c_i \in [0, 255]$ is a character or byte. Except for very short strings ($q \leq 8$), the corresponding number

$$\sum_{i=0}^{q-1} c_{q-1-i} \cdot 256^i$$

does *not* fit in any of the standard atomic data types such as `char`, `short`, `int`, etc. The size of these data types is bounded by the size of the underlying CPU registers, which are thus usually up to 64 bit long. The built-in C++ operator `%` can be used to compute the remainder for atomic data types, this cannot be used directly for numbers wider than 64 bits.

Write a function `uint32_t hash(const std::string& str, const uint32_t m)` that takes as input an arbitrary long string `str` and a divisor `m` and computes

$$h(c) = \left(\sum_{i=0}^{q-1} c_{q-1-i} \cdot 256^i \right) \bmod m$$

where c denotes the string `str` and q is the number of characters in the string.

Hint. Make use of the properties of modular arithmetic:

- $a + b \bmod m = ((a \bmod m) + (b \bmod m)) \bmod m$
- $ab \bmod m = ((a \bmod m) \cdot (b \bmod m)) \bmod m$

8.4 Permutation invariance of the division method

Show that choosing $m = 255$ in the previous question results in hashes that are invariant to permutation of the characters in the string (meaning for example that $h(\text{ciao}) = h(\text{oaic})$). Discuss whether this is a desirable property or not.

9 Graphs

9.1 Shortest paths using the adjacency list representation

Starting from the code given in the notes, write a version of the Bellman-Ford and Dijkstra algorithms that use the adjacency list representation of a graph instead of the adjacency matrix representation, as given in the notes. In particular, implement the following functions.

File `shortest_paths_sparse.hpp`:

```

1  #ifndef __shortest_paths_sparse__
2  #define __shortest_paths_sparse__
3
4  #include "graph.hpp"
5
6  std::vector<hop_t> bellman_ford(const SparseGraph &graph, const int source,
7                               bool &has_negative_cycle);
8
9  std::vector<hop_t> dijkstra(const SparseGraph &graph, const int source);
10
11 #endif // __shortest_paths_sparse__

```

You can start from the following `shortest_paths_sparse.cpp` template:

```

1  #include "shortest_paths_sparse.hpp"
2  #include <priority_queue.hpp>
3
4  #include <cmath>
5
6  std::vector<hop_t> bellman_ford(const SparseGraph &graph, const int source,
7                               bool &has_negative_cycle)
8  {
9      const int V = static_cast<int>(graph.size());
10     auto DP = std::vector<hop_t>(V, {inf, -1});
11
12     // WRITE YOUR CODE HERE
13
14     return DP;
15 }
16
17 struct triplet_t {
18     float d;
19     int r;
20     int v;
21 };
22
23 std::vector<hop_t> dijkstra(const SparseGraph &graph, const int source)
24 {
25     assert(source >= 0);
26     assert(source < (signed)graph.size());
27
28     auto DP = std::vector<hop_t>(graph.size(), {inf, -1});
29
30     // WRITE YOUR CODE HERE
31
32     return DP;
33 }

```

Use the following driver to test them.

File `shortest_paths_sparse_driver.cpp`:

```

1  #include "shortest_paths_sparse.hpp"
2  #include <utils.hpp>
3
4  int main(int argc, const char *argv[])

```

```

5 {
6     auto graph = sparse_test_graph;
7     print_graph(graph);
8
9     {
10        int source = 2;
11        bool has_negative_cycle;
12        std::cout << "Bellman-Ford from source " << source
13                << std::endl;
14        auto DP = bellman_ford(graph, source, has_negative_cycle);
15        print(DP);
16        std::cout << std::endl;
17    }
18
19    {
20        int source = 2;
21        std::cout << "Dijkstra from source " << source << std::endl;
22        auto DP = dijkstra(graph, source);
23        print(DP);
24        std::cout << std::endl;
25    }
26
27    return 0;
28 }

```

Are these implementations better for sparse graph?

9.2 Decoding shortest paths

Our shortest paths algorithms encode the shortest paths as a pair of vectors (or matrices) D and P , both coded in a single vector DP of type `std::vector<hop_t>`, where each `hop_t` element contains a pair (D_v, P_v) .

Write a module `shortest_path_decode.hpp` defining a function:

```
std::vector<int> decode(const std::vector<hop_t> &DP, int v);
```

You can start from the following `shortest_paths_decode.hpp` template:

```

1  #ifndef __shortest_paths_decode__
2  #define __shortest_paths_decode__
3
4  #include "graph.hpp"
5  #include <algorithm>
6  #include <vector>
7
8  inline std::vector<int> decode(const std::vector<hop_t> &DP, int v)
9  {
10     // WRITE YOUR CODE HERE
11 }
12
13 #endif // __shortest_paths_decode__

```

that takes DP and a destination vertex v index and returns the path from the source to the destination as a vector of vertex indices. Use the following test driver to test it.

File `shortest_paths_fw_decode_driver.cpp`:

```

1
2 #include <algorithm>
3 #include <iostream>
4 #include <vector>
5
6 #include "shortest_paths_decode.hpp"
7 #include "shortest_paths_fw.hpp"
8 #include <utils.hpp>
9
10 int main(int argc, const char *argv[])
11 {
12     Graph graph = test_graph;
13     print_graph(graph);
14
15     std::cout << "Floyd-Warshall ASPSP" << std::endl;
16     auto DP = floyd_warshall(graph);
17     for (const auto &row : DP) {
18         print(row);
19     }
20     std::cout << std::endl;
21
22     for (int u = 0; u < (signed)graph.size(); ++u) {
23         for (int v = 0; v < (signed)graph.size(); ++v) {
24             auto path = decode(DP[u], v);
25             if (path.size()) {
26                 std::cout << "Shortest path " << u << " --> " << v
27                     << " (weight " << DP[u][v].weight << "): ";
28                 print(path);
29             }
30         }
31     };
32
33     return 0;
34 }

```

```

digraph G {
    0 -> 1 [label= 4];
    0 -> 7 [label= 8];
    1 -> 7 [label= 11];
    2 -> 5 [label= 4];
    2 -> 8 [label= 2];
    3 -> 4 [label= 9];
    3 -> 5 [label= 14];
    4 -> 5 [label= 10];
    5 -> 6 [label= 2];
    6 -> 3 [label= 3];
    6 -> 7 [label= 1];
    6 -> 8 [label= 6];
    7 -> 8 [label= 7];
}

```

Floyd-Warshall ASPSP

```

[(0,-1), (4,0), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (8,0), (15,7)]
[(inf,-1), (0,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (11,1), (18,7)]

```



```

[(inf,-1), (inf,-1), (0,-1), (9,6), (18,3), (4,2), (6,5), (7,6), (2,2)]
[(inf,-1), (inf,-1), (inf,-1), (0,-1), (9,3), (14,3), (16,5), (17,6), (22,6)]
[(inf,-1), (inf,-1), (inf,-1), (15,6), (0,-1), (10,4), (12,5), (13,6), (18,6)]
[(inf,-1), (inf,-1), (inf,-1), (5,6), (14,3), (0,-1), (2,5), (3,6), (8,6)]
[(inf,-1), (inf,-1), (inf,-1), (3,6), (12,3), (17,3), (0,-1), (1,6), (6,6)]
[(inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (0,-1), (7,7)]
[(inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (inf,-1), (0,-1)]

```

```

Shortest path 0 --> 1 (weight 4): [0, 1]
Shortest path 0 --> 7 (weight 8): [0, 7]
Shortest path 0 --> 8 (weight 15): [0, 7, 8]
Shortest path 1 --> 7 (weight 11): [1, 7]
Shortest path 1 --> 8 (weight 18): [1, 7, 8]
Shortest path 2 --> 3 (weight 9): [2, 5, 6, 3]
Shortest path 2 --> 4 (weight 18): [2, 5, 6, 3, 4]
Shortest path 2 --> 5 (weight 4): [2, 5]
Shortest path 2 --> 6 (weight 6): [2, 5, 6]
Shortest path 2 --> 7 (weight 7): [2, 5, 6, 7]
Shortest path 2 --> 8 (weight 2): [2, 8]
Shortest path 3 --> 4 (weight 9): [3, 4]
Shortest path 3 --> 5 (weight 14): [3, 5]
Shortest path 3 --> 6 (weight 16): [3, 5, 6]
Shortest path 3 --> 7 (weight 17): [3, 5, 6, 7]
Shortest path 3 --> 8 (weight 22): [3, 5, 6, 8]
Shortest path 4 --> 3 (weight 15): [4, 5, 6, 3]
Shortest path 4 --> 5 (weight 10): [4, 5]
Shortest path 4 --> 6 (weight 12): [4, 5, 6]
Shortest path 4 --> 7 (weight 13): [4, 5, 6, 7]
Shortest path 4 --> 8 (weight 18): [4, 5, 6, 8]
Shortest path 5 --> 3 (weight 5): [5, 6, 3]
Shortest path 5 --> 4 (weight 14): [5, 6, 3, 4]
Shortest path 5 --> 6 (weight 2): [5, 6]
Shortest path 5 --> 7 (weight 3): [5, 6, 7]
Shortest path 5 --> 8 (weight 8): [5, 6, 8]
Shortest path 6 --> 3 (weight 3): [6, 3]
Shortest path 6 --> 4 (weight 12): [6, 3, 4]
Shortest path 6 --> 5 (weight 17): [6, 3, 5]
Shortest path 6 --> 7 (weight 1): [6, 7]
Shortest path 6 --> 8 (weight 6): [6, 8]
Shortest path 7 --> 8 (weight 7): [7, 8]

```

10 Further examples (optional)

All these examples are **optional**. Do them if you want to develop a better understanding of certain advanced features of the C++ language.

10.1 Inserting movable objects in an array

The implementation of `array_insert` given in the notes is quite inefficient as it requires copying elements in the array. While the complexity of the algorithm cannot be improved without switching to a different data structure, we can at least make the implementation faster by moving instead of copying the array elements.

Write a function

```
template <typename T>
void array_insert_movable(std::vector<T>& A, size_t index, T&& x);
```

that uses move operations instead of copy operations whenever possible.

In order to test this function, you can use the `Movable` type provided in `probes.hpp`. This class prints messages whenever different constructors are invoked, which allows you to track copy and move operation explicitly, as follows.

File `probes.hpp`:

```
1  #ifndef __probes_hpp__
2  #define __probes_hpp__
3
4  #include <iostream>
5
6  // A test class with copy and move constructor and assignment
7  struct Movable {
8      float x;
9
10     Movable() : x{0}
11     {
12         std::cout << "Movable: Default-constructed" << '\n';
13     }
14
15     Movable(float x) : x{x}
16     {
17         std::cout << "Movable: Constructed from \"" << x << "\"\n";
18     }
19
20     Movable(const Movable& m) : x{m.x}
21     {
22         std::cout << "Movable: Copy-constructed" << '\n';
23     }
24
25     Movable(Movable&& m) noexcept : x{m.x}
26     {
27         m.x = 0;
28         std::cout << "Movable: Move-constructed\n";
29     }
30
31     Movable& operator=(const Movable& m)
32     {
33         x = m.x;
34         std::cout << "Movable: Copy-assigned\n";
35         return *this;
36     }
37
38     Movable& operator=(Movable&& m)
39     {
40         x = m.x;
41         m.x = 0;
42         std::cout << "Movable: Move-assigned\n";
43         return *this;
44     }
```

```

45
46     ~Movable() { std::cout << "Movable: Destructed\n"; }
47 };
48
49 // A test class with copy constructor and assignment only
50 struct Copyable {
51     float x;
52
53     Copyable() : x{0}
54     {
55         std::cout << "Copyable: Default-constructed\n";
56     }
57
58     Copyable(float x) : x{x}
59     {
60         std::cout << "Copyable: Constructed from \"" << x << "\"\n";
61     }
62
63     Copyable(const Copyable& m) : x{m.x}
64     {
65         std::cout << "Copyable: Copy-constructed\n";
66     }
67
68     Copyable& operator=(const Copyable& m)
69     {
70         x = m.x;
71         std::cout << "Copyable: Copy-assigned\n";
72         return *this;
73     }
74
75     ~Copyable() { std::cout << "Copyable: Destructed\n"; }
76 };
77
78 inline std::ostream& operator<<(std::ostream& os, const Movable& m)
79 {
80     return os << m.x;
81 }
82
83 inline std::ostream& operator<<(std::ostream& os, const Copyable& c)
84 {
85     return os << c.x;
86 }
87
88 #endif // __probes_hpp__

```

Test `array_insert_movable` with the following code:

```

1 #include "array.hpp"
2 #include "array_insert_movable.hpp"
3 #include "utils.hpp"
4 #include "probes.hpp"
5
6 std::vector<Movable> make_test_array()
7 {

```

```

8     std::cout << "Constructing an array of Movable objects\n";
9     auto A = std::vector<Movable>{};
10    A.reserve(10);
11    A.push_back(Movable{1});
12    A.push_back(Movable{2});
13    A.push_back(Movable{3});
14    print(A, "Constructed array: ");
15    std::cout << "\n";
16    return A;
17 }
18
19 int main(int argc, char** argv)
20 {
21     {
22         auto A = make_test_array();
23         auto x = Movable{-1};
24         std::cout << "Inserting object by copy:" << '\n';
25         array_insert(A, 0, x);
26         print(A, "Array content: ");
27     }
28
29     std::cout << "\n";
30
31     {
32         auto A = make_test_array();
33         auto x = Movable{-1};
34         std::cout << "Inserting object by copying and moving:" << '\n';
35         array_insert_movable(A, 0, x);
36         print(A, "Array content: ");
37     }
38
39     std::cout << "\n";
40
41     {
42         auto A = make_test_array();
43         auto x = Movable{-1};
44         std::cout << "Inserting object by moving:" << '\n';
45         array_insert_movable(A, 0, std::move(x));
46         print(A, "Array content: ");
47     }
48
49     return 0;
50 }

```

```

Constructing an array of Movable objects
Movable: Constructed from "1"
Movable: Move-constructed
Movable: Destructed
Movable: Constructed from "2"
Movable: Move-constructed
Movable: Destructed
Movable: Constructed from "3"
Movable: Move-constructed
Movable: Destructed

```

Constructed array: [1, 2, 3]

Movable: Constructed from "-1"

Insertinig object by copy:

Movable: Copy-constructed

Movable: Copy-assigned

Movable: Copy-assigned

Movable: Copy-assigned

Array content: [-1, 1, 2, 3]

Movable: Destructed

Movable: Destructed

Movable: Destructed

Movable: Destructed

Movable: Destructed

Constructing an array of Movable objects

Movable: Constructed from "1"

Movable: Move-constructed

Movable: Destructed

Movable: Constructed from "2"

Movable: Move-constructed

Movable: Destructed

Movable: Constructed from "3"

Movable: Move-constructed

Movable: Destructed

Constructed array: [1, 2, 3]

Movable: Constructed from "-1"

Insertinig object by copying and moving:

Movable: Move-constructed

Movable: Move-assigned

Movable: Move-assigned

Movable: Copy-assigned

Array content: [-1, 1, 2, 3]

Movable: Destructed

Movable: Destructed

Movable: Destructed

Movable: Destructed

Movable: Destructed

Constructing an array of Movable objects

Movable: Constructed from "1"

Movable: Move-constructed

Movable: Destructed

Movable: Constructed from "2"

Movable: Move-constructed

Movable: Destructed

Movable: Constructed from "3"

Movable: Move-constructed

Movable: Destructed

Constructed array: [1, 2, 3]

Movable: Constructed from "-1"

Insertinig object by moving:

```

Movable: Move-constructed
Movable: Move-assigned
Movable: Move-assigned
Movable: Move-assigned
Array content: [-1, 1, 2, 3]
Movable: Destructed
Movable: Destructed
Movable: Destructed
Movable: Destructed
Movable: Destructed

```

The code test three cases: using `array_insert` from the notes, using `array_insert_movable` without actually moving the new element `x` (which causes it to be copied), and finally using `array_insert_movable` to move `x` into the array.

If you implemented `array_insert_movable`, from the printout above you should notice that only the move constructors/assignments of `Movable` are used in the last case.

10.2 Stack with move semantics

In the notes, the `Stack` class has two implementations of the `push` method: one version takes as input a lvalue reference to the value to be pushed, which internally creates a copy of the element. The other version takes a rvalue reference, which supports moving rather than copying. Besides being potentially faster, moving is required if you use automatic memory management using `unique_ptr`: by moving, you transfer ownership of objects to and from the stack.

This is demonstrated in the following driver.

File `stack_move_driver.cpp`:

```

1  #include <iostream>
2  #include <memory>
3
4  #include <stack.hpp>
5
6  template <typename T> class StackWithMove : public Stack<T>
7  {
8      using Stack<T>::Stack;
9
10     public:
11         // Move a value to the top of the stack (optional)
12         void push(T &&x)
13         {
14             assert(this->_head < this->_storage.size());
15             this->_storage[this->_head++] = std::move(x);
16         }
17 };
18
19 int main(int argc, char** argv)
20 {
21     auto stack = StackWithMove<std::unique_ptr<float>>(10);
22
23     // Create a new `float` object
24     auto data = std::make_unique<float>(1);
25     std::cout << "`data` points to " << data.get() << std::endl;
26
27     // Push the oobject on the stack by transferring ownership to it

```

```

28     stack.push(std::move(data));
29     std::cout << "`data` points to " << data.get() << std::endl;
30
31     // Retrieve the object and its ownership
32     auto popped = std::move(stack.top());
33     stack.pop();
34     std::cout << "`popped` points to " << popped.get() << std::endl;
35
36     // The object will be automatically deleted
37     // by exiting the function as the variable `retrieved`, which owns
38     // it, falls out of scope (RAII idiom)
39     return 0;
40 }

```

```

`data` points to 0x7f87d8405880
`data` points to 0x0
`popped` points to 0x7f87d8405880

```

Explain the output of `stack_move_driver`.

10.3 List iterators

In the examples so far we have “manually” iterated over a linked list by using pointers. Iterators are a concept used by the C++ standard library to abstract operations involving such pointers.

By implementing iterators, we can iterate over the list in a much nicer way. The following test driver shows how this can be done using iterators directly (`cbegin(list)`) and using a range-based for loop (`x : list`). It also shows the power of template by passing the list to our own `print` helper function defined in `utils.hpp`: the latter works because, internally, it is based on a range-based for loop that our list container now supports.

```

1  #include <iostream>
2
3  #include "list.hpp"
4  #include "list_iterator.hpp"
5  #include "utils.hpp"
6
7  int main(int argc, char** argv)
8  {
9      auto list = Node<float>{};
10
11     // Insert some numbers in the list
12     for (int i = 0; i < 10; ++i) {
13         list_insert_after(&list, static_cast<float>(i));
14     }
15
16     // Traverse the list using a const iterator
17     for (auto iter = cbegin(list); iter != cend(list); ++iter) {
18         std::cout << *iter << ' ';
19     }
20     std::cout << std::endl;
21
22     // Traverse the list using a range for loop
23     for (const auto& x : list) {
24         std::cout << x << ' ';
25     }

```

```
26     std::cout << std::endl;
27
28     // Traverse the list using our own `print` helper function from
29     // `utils.hpp`
30     print(list, "List content: ");
31
32     return 0;
33 }
```

9 8 7 6 5 4 3 2 1 0

9 8 7 6 5 4 3 2 1 0

List content: [9, 8, 7, 6, 5, 4, 3, 2, 1, 0]

Write a STL-compatible iterator for the `Node<T>` list class.