



Graphs and graph algorithms

A **graph** is formed of a (finite) set of vertices/nodes and a set of edges between them. We distinguish four types of graphs:



Unweighted

Weighted

Undirected



Directed



If the graph is undirected, an edge between nodes u and w can be thought of as having two edges $u \rightarrow w$ and $w \rightarrow u$.

Examples

Undirected unweighted graph:

- vertices = registered people on Facebook
- edges = friendships between people (it is mutual!)

Directed unweighted graph:

- vertices = registered people on Twitter
- edges = who is following who

Undirected weighted graph:

- vertices = train stops/stations
- edges = rail lines connecting train stops together with their length

Graph represented as an adjacency matrix

Assume that graph's vertices are numbered $V = \{0, 1, 2, \dots, n - 1\}$.

Adjacency matrix G is a two-dimensional array/matrix $n \times n$ described as follows.

Unweighted graphs:

- $G[v][w] = 1$ if there is an edge going from v to w
- $G[v][w] = 0$ if there is no such edge


Weighted graphs:

- $G[v][w] = \text{weight of the edge going from } v \text{ to } w$
- $G[v][w] = \infty$ if there is no such edge
- $G[v][v] = 0$

Remark: The graph is undirected if $G[v][w] = G[w][v]$ for all vertices v and w .

Example: Adjacency matrix


Unweighted undirected:



$$G = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$$

For example: $G[2][0] = 0$ and $G[2][1] = 1$.

Weighted directed:



$$G = \begin{pmatrix} 0 & \infty & \infty \\ 4 & 0 & 1 \\ \infty & 3 & 0 \end{pmatrix}$$

For example: $G[2][0] = \text{infy}$ and $G[2][1] = 3$.

Graph represented as adjacency lists

To represent a graph on vertices $V = \{0, 1, 2, \dots, n - 1\}$ by *adjacency lists* we have an array \mathbf{N} of n -many linked lists (one list for every vertex).

Unweighted:


- $\mathbf{N}[v]$ is the list of *neighbours* of v .
(w is a neighbour of v if there is an edge $v \rightarrow w$)

Weighted:

- $\mathbf{N}[v]$ is the list of *neighbours* of v together with the weight of the edge that connects them with v .


Example: adjacency lists

Unweighted undirected:



$N[v]$	neighbours
A	B
B	A, C
C	B

Weighted directed:




$N[v]$	neighbours & weights
A	
B	(A, 4), (C, 1)
C	(B, 3)

We said that representing a graph by adjacency lists means that we will have an array N of n -many linked list (where n is the number of vertices). Then, for example, $N[2]$ stores the address of the head of the linked list of all neighbours of the 2nd vertex. If we name our vertices by letters A , B , C , for example, we need to find a way to assign indexes of the array N to the letters A , B , C . One way to do this is to use hash tables.

However, in the example given here, we don't care how this is done. We assume that we have lists of neighbours stored in $N[A]$, $N[B]$, $N[C]$.

In the weighted case, $N[B]$ also stores the weights of the edges:



But instead of drawing this and we just say that $N[B]$ stores the list $(A, 4)$, $(C, 1)$.

Comparison of those two methods

Set n = the number of vertices, m = the total number of edges.

	Adjacency matrix	Adjacency lists
Checking if there is an edge $v \rightarrow w$:	Reading $G[v][w]$ (which is in $\mathcal{O}(1)$).	Checking if w is in the list $N[v]$.
Allocated space:	n arrays of size n $= \mathcal{O}(n \times n)$ space.	n linked lists storing m edges in total $= \mathcal{O}(n + m)$ space.
Traversing v 's neighbours:	Traversing all $G[v][0]$, $G[v][1]$, ..., $G[v][n-1]$. $= \mathcal{O}(n)$ time.	Traversing only the linked list $N[v]$.

In the third case (with adjacency lists) we only traverse the actual neighbours of v . This is better whenever the graph is **sparse** (= not dense), that is, if there are relatively few edges.

A graph is **sparse** if m is approximately equal to n .

An example of a sparse graph would be the graph of Facebook users with edges representing friendships. Facebook has hundreds of millions of users but each user has only a few hundreds of friends. In other words, every vertex of the graph has only a few hundreds of neighbours.

From the table we see that checking whether an edge exists is much faster for adjacency matrix. On the other hand, if our graph is sparse, then the allocated space of adjacency lists is much smaller than adjacency matrix and also traversing neighbours is faster for adjacency lists than adjacency matrix.

Paths and shortest paths

A **path** from v to z is a sequence of edges

$$v \rightarrow w_1 \rightarrow w_2 \rightarrow \dots \rightarrow z$$


connecting v with z .

The **shortest path** is the path such that the sum of weights of its edges is the minimal such. (In unweighted graphs, set weights to 1.)

Example

1. $A \rightarrow B \rightarrow E \rightarrow G$
2. $A \rightarrow C \rightarrow D \rightarrow F \rightarrow G$
3. ...

The shortest: $A \rightarrow B \rightarrow D \rightarrow F \rightarrow E \rightarrow G$



Dijkstra's algorithm to find the shortest path from v to z

For each vertex w of the graph, we keep track of the following:

- i. $d[w]$ = the shortest distance from v to w so far (initially: ∞)
- ii. $p[w]$ = the predecessor on the path from v (initially: w)
- iii. $f[w]$ = is computation of $d[w]$ finished? (initially: `false`)

The algorithm (idea):

1. Set $d[v] = 0$.
2. Set $w =$ the yet unfinished vertex with the smallest $d[w]$.
3. Set $f[w] = \text{true}$ (mark w as *finished*).
4. Update $d[u]$ and $p[u]$ of neighbours of w :
For every neighbour u of w such that $d[w] + \text{weight}(w,u) < d[u]$,
set $d[u] = d[w] + \text{weight}(w,u)$ and $p[u] = w$.
5. If still $f[z] == \text{false}$, go to step 3.

(Where $\text{weight}(w,u)$ is the weight of the edge $w \rightarrow u$.)

The input of the algorithm is a graph (represented as an adjacency matrix or adjacency lists) and two vertices v and z . The aim is to find the shortest path from v to z .


As the algorithm runs it changes the values $d[w]$, $p[w]$ and $f[w]$. Initially $d[w] = \text{infinity}$, $p[w] = w$ and $f[w] = \text{false}$ for every vertex w .

The arrays d and f obey the following *invariants*:

- $d[w]$ is the length of the shortest path from v to w when using only the finished vertices (i.e. those w such that $f[w] == \text{true}$).
- If w is finished then $d[w]$ is the actual length of the shortest path from v to w .

After the algorithm finishes, we compute the found shortest path by using the array p . Lastly, $\text{weight}(w,u)$ is the weight of the edge $w \rightarrow u$ which we obtain the adjacency matrix/lists of the graph.


Example: Execution of Dijkstra's algorithm



Example: Execution of Dijkstra's algorithm

finished


A	B	C	D	E	F
0,A	∞ ,B	∞ ,C	∞ ,D	∞ ,E	∞ ,F



Example: Execution of Dijkstra's algorithm

finished


A	B	C	D	E	F	
0,A	∞ ,B	∞ ,C	∞ ,D	∞ ,E	∞ ,F	
0,A, \checkmark	3,A	1,A	7,A	∞ ,E	∞ ,F	A



Example: Execution of Dijkstra's algorithm


finished

A	B	C	D	E	F	
0,A	∞ ,B	∞ ,C	∞ ,D	∞ ,E	∞ ,F	
0,A, \checkmark	3,A	1,A	7,A	∞ ,E	∞ ,F	A
0,A, \checkmark	2,C	1,A, \checkmark	6,C	2,C	∞ ,F	C




Example: Execution of Dijkstra's algorithm

A	B	C	D	E	F	finished
0, A	∞ , B	∞ , C	∞ , D	∞ , E	∞ , F	
0, A, \checkmark	3, A	1, A	7, A	∞ , E	∞ , F	A
0, A, \checkmark	2, C	1, A, \checkmark	6, C	2, C	∞ , F	C
0, A, \checkmark	2, C, \checkmark	1, A, \checkmark	6, C	2, C	∞ , F	B




Example: Execution of Dijkstra's algorithm




A	B	C	D	E	F	finished
0,A	∞ ,B	∞ ,C	∞ ,D	∞ ,E	∞ ,F	
0,A,✓	3,A	1,A	7,A	∞ ,E	∞ ,F	A
0,A,✓	2,C	1,A,✓	6,C	2,C	∞ ,F	C
0,A,✓	2,C,✓	1,A,✓	6,C	2,C	∞ ,F	B
0,A,✓	2,C,✓	1,A,✓	5,E	2,C,✓	8,E	E

Example: Execution of Dijkstra's algorithm



A	B	C	D	E	F	finished
0,A	∞ ,B	∞ ,C	∞ ,D	∞ ,E	∞ ,F	
0,A,✓	3,A	1,A	7,A	∞ ,E	∞ ,F	A
0,A,✓	2,C	1,A,✓	6,C	2,C	∞ ,F	C
0,A,✓	2,C,✓	1,A,✓	6,C	2,C	∞ ,F	B
0,A,✓	2,C,✓	1,A,✓	5,E	2,C,✓	8,E	E
0,A,✓	2,C,✓	1,A,✓	5,E,✓	2,C,✓	7,F	D


Example: Execution of Dijkstra's algorithm



A	B	C	D	E	F	finished
0,A	∞ ,B	∞ ,C	∞ ,D	∞ ,E	∞ ,F	
0,A,✓	3,A	1,A	7,A	∞ ,E	∞ ,F	A
0,A,✓	2,C	1,A,✓	6,C	2,C	∞ ,F	C
0,A,✓	2,C,✓	1,A,✓	6,C	2,C	∞ ,F	B
0,A,✓	2,C,✓	1,A,✓	5,E	2,C,✓	8,E	E
0,A,✓	2,C,✓	1,A,✓	5,E,✓	2,C,✓	7,F	D
0,A,✓	2,C,✓	1,A,✓	5,E,✓	2,C,✓	7,F,✓	F

Example: Execution of Dijkstra's algorithm

finished



A	B	C	D	E	F	finished
0,A	∞ ,B	∞ ,C	∞ ,D	∞ ,E	∞ ,F	
0,A,✓	3,A	1,A	7,A	∞ ,E	∞ ,F	A
0,A,✓	2,C	1,A,✓	6,C	2,C	∞ ,F	C
0,A,✓	2,C,✓	1,A,✓	6,C	2,C	∞ ,F	B
0,A,✓	2,C,✓	1,A,✓	5,E	2,C,✓	8,E	E
0,A,✓	2,C,✓	1,A,✓	5,E,✓	2,C,✓	7,F	D
0,A,✓	2,C,✓	1,A,✓	5,E,✓	2,C,✓	7,F,✓	F

The shortest path from A to F is obtained (in the reversed order) by reading out $p[w]$'s, starting from F:

$A \rightarrow C \rightarrow E \rightarrow D \rightarrow F$.

Every iteration of the algorithm corresponds to one row in the table and each such row shows the content of the three arrays `d[-]`, `p[-]` and `f[-]`. (Check marks denote finished vertices.)

In the graph, the two circles adjacent to a vertex mark the current state of `d[w]` and `p[w]`. They turn blue whenever is the vertex marked as finished.

For a detailed explanation of Dijkstra's run see the solution file for exercises of week 11:

[https://canvas.bham.ac.uk/courses/27506/files/5177678?
module_item_id=895128](https://canvas.bham.ac.uk/courses/27506/files/5177678?module_item_id=895128)

Dijkstra's time complexity (adjacency matrix)

n = the number of vertices, m = the total number of edges.

We do the following *up to* n -times:

2. Find w which is unfinished and with the smallest $d[w]$.
3. Mark w as finished.
4. Update every neighbour of w .

Representing the graph by an *adjacency matrix*, means that it takes $\mathcal{O}(n)$ to do step 4.

We can also do step 2. in $\mathcal{O}(n)$ by going through all vertices.

\implies The time complexity is $\mathcal{O}(n^2)$.

Dijkstra's time complexity (adjacency lists)

We do the following *up to* n -times:

2. Find w which is unfinished and with the smallest $d[w]$.
3. Mark w as finished.
4. Update every neighbour of w .

With *adjacency lists*, executions of step 4. will (in total) update

neighbours of the 1st selected w ,
neighbours of the 2nd selected w ,
neighbours of the 3rd selected w ,

...

Over all iterations combined we update m -many times! $\Rightarrow \mathcal{O}(m)$

Speeding up steps 2. and 4.:

Use min-priority queue! The priority of u is $d[u]$.

- We call `deleteMin` once per iteration, i.e. up to n -times.
- Whenever $d[u]$ changes, we `update` the priority of u .

\Rightarrow The total time complexity

$$\mathcal{O}(n \times \text{"cost of deleteMin"} + m \times \text{"cost of update"})$$

What omitted in the analysis is the time complexity of initialising the heap. This is usually done by `heapify` and its time complexity was always $\mathcal{O}(n)$ for all heaps we had. Alternatively, we can do `insert` n -times which will result in the time complexity $\mathcal{O}(n \log n)$ or $\mathcal{O}(n)$ depending on the heap that we are using. Either way, the initialisation will not play any role in the total time complexity.

Dijkstra's time complexity – comparison

Adjacency matrices	Adjacency lists	
	Binary Heaps	Fibonacci Heaps
$\mathcal{O}(n^2)$	$\mathcal{O}((n + m) \log n)$	$\mathcal{O}((n \log n) + m)$

Min-priority queues:

- Binary heaps: both `update` and `deleteMin` are in $\mathcal{O}(\log n)$.
- Fibonacci heaps: `update` is in $\mathcal{O}(1)$ and `deleteMin` is in $\mathcal{O}(\log n)$ (both amortized).

Remark: Dijkstra's algorithm works only if all weights are ≥ 0 .

Remark: If the graph is *dense*, that is if the number of edges is approximately n^2 , then using adjacency lists together with binary heaps has the time complexity $\mathcal{O}((n + n^2) \log n) = \mathcal{O}(n^2 \log n)$ which is slower than just using adjacency matrices. This problem disappears when using Fibonacci heaps where, for dense graphs, we the time complexity becomes $\mathcal{O}(n \log n + n^2) = \mathcal{O}(n^2)$.

On the other hand, if the graph is not dense, using adjacency lists with binary heaps or Fibonacci heaps is faster than using adjacency matrices.

Dijkstra's algorithm (pseudocode with adjacency matrix)

```
1  dijkstra_with_matrix(int [][] G, int v, int z) {
2      n = G.length;
3      d = new int[n]; p = new int[n]; f = new bool[n];
4
5      for (int w=0; w<n; w++) {
6          d[w] = inf; p[w] = w; f[w] = false;
7      }
8      d[v] = 0;
9
10     while (true) {
11         w = min_unfinished(d, f);
12         if (w == -1) break;
13
14         for (int u=0; u<n; u++) update(w, u, d, p);
15         f[w] = true;
16     }
17     // Compute output in a required form:
18     return compute_result(v, z, G, d, p);
19 }
```

```
1 int min_unfinished(int [] d, bool [] f) {
2     int min = +infty;
3     int idx = -1;
4
5     for (int i=0; i<d.length; i++) {
6         if (f[i] == false && d[i] < min) {
7             idx = i;
8             min = d[i]
9         }
10    }
11
12    return idx;
13 }
```

```
1 void update(w, u, G, d, p) {
2     if (d[w] + G[w][u] < d[u]) {
3         d[u] = d[w] + G[w][u];
4         p[u] = w;
5     }
6 }
```

Dijkstra's algorithm (pseudocode with adjacency lists)

```
1  dijkstra_with_lists(List<Edge>[] N, int v, int z) {
2      n = G.length;
3      d = new int[n];    p = new int[n];
4      Q = new MinPriorityQueue();
5
6      for (int w=0; w<n; w++) {
7          d[w] = inf;    p[w] = w;
8          Q.add(w, d[w]);
9      }
10     d[v] = 0;
11     Q.update(v, 0);
12
13     while (Q.notEmpty()) {
14         w = Q.deleteMin();
15
16         for (Edge e : N[w]) { // iterate over edges to neighbours
17             u = e.target;
18             if (d[w] + e.weight < d[u]) { // should we update?
19                 d[u] = d[w] + e.weight;
20                 p[u] = w;
21                 Q.update(u, d[u]);
22             }
23         }
24     }
25
26     return compute_result(v, z, G, d, p);
27 }
```

```
1  class Edge {
2      // target node
3      int target;
4
5      int weight;
6  }
```


The initialisation happens on lines 6–9.

Lines 10–11 make sure that the first selected w will be v .

We use the class `Edge` to store neighbours together with the weight of the edge that connects them. For example, if the vertex `A` has neighbours `B`, `C` and `D` with the edge $A \rightarrow B$ of weight 3, $A \rightarrow C$ of weight 1, and $A \rightarrow D$ of weight 8, then we will have that the linked list `N[v]` stores `Edge(B, 3)`, `Edge(C, 1)` and `Edge(D, 8)`.

Minimal spanning tree


Such that there is a path between any two vertices.

Assumption: Consider only *undirected* and *connected* graphs!


A **spanning tree** is a minimal possible selection of edges which connects all vertices. (That is, a spanning tree does not contain any cycles.)

Minimum spanning tree is a spanning tree such that the sum of weights of its edges is the minimal such.


Example




Spanning tree.



Not a spanning tree.
A is not connected.



Not a spanning tree.
Any of the edges CV,
CD, BD could be re-
moved.







Minimum
spanning tree!

Example: Execution of Jarník-Prim algorithm

Idea: Iteratively extend the tree with an edge which has the smallest weight and which connects a yet unconnected node.

Example

Start from
any node,
e.g. B.



The minimal spanning tree consists of edges:
AB, BC, AE, ED.

Jarník-Prim algorithm for finding the minimal spanning tree

For each vertex w of the graph, we keep track of the following:

- i. $d[w]$ = the current distance from the *tree* (initially: ∞)
- ii. $p[w]$ = the vertex which connects to the tree (initially: w)
- iii. $f[w]$ = has w been added to the tree? (initially: *false*)

The algorithm (idea):

1. Set $d[0] = 0$. (vertex 0 could be replaced by any other vertex)
2. Set $w =$ the yet unfinished node with the smallest $d[w]$.
3. Set $f[w] = \text{true}$ (mark w as *finished*).
4. Update $d[u]$ and $p[u]$ of neighbours of w :
For every neighbour u of w such that $\text{weight}(w,u) < d[u]$, set
 $d[u] = \text{weight}(w,u)$ and $p[u] = w$.
5. If there are still some nodes unfinished, go to step 3.

(Where $\text{weight}(w,u)$ is the weight of the edge $w \rightarrow u$.)

Jarník-Prim's algorithm works similarly to Dijkstra's algorithm. The differences are marked by red. Although the principle is similar, the interpretation of the execution and the result is different. The main idea of Jarník-Prim is that we are building a “spanning tree” iteratively, in steps.

The following invariants hold for Jarník-Prim:

1. The vertices marked as finished are connected/added to the tree.
2. For those vertices which are not connected yet, $d[w]$ denotes the smallest weight of an edge that connects w to the tree.
3. $p[w]$ denotes the vertex of the tree such that the edge between w and $p[w]$ is the edge with weight $d[w]$.

In step 1. we pick an arbitrary vertex (stored on 0th position) and mark its distance as 0. As a consequence, we start building tree from this vertex.

After the algorithm finishes, i.e. all vertices are marked finished, we can read out the spanning tree from the array $p[-]$. To obtain the minimum spanning tree, for every vertex w (except for $w == 0$), add the edge $w - p[w]$.

Jarník-Prim's time complexity

The time complexity is the same as for Dijkstra's algorithm!

Adjacency matrices	Adjacency lists	
	Binary Heaps	Fibonacci Heaps
$\mathcal{O}(n^2)$	$\mathcal{O}((n + m) \log n)$	$\mathcal{O}((n \log n) + m)$

Remark: Unlike Dijkstra's algorithm, Jarník-Prim's algorithm would also work for graphs with edges that have negative weights.