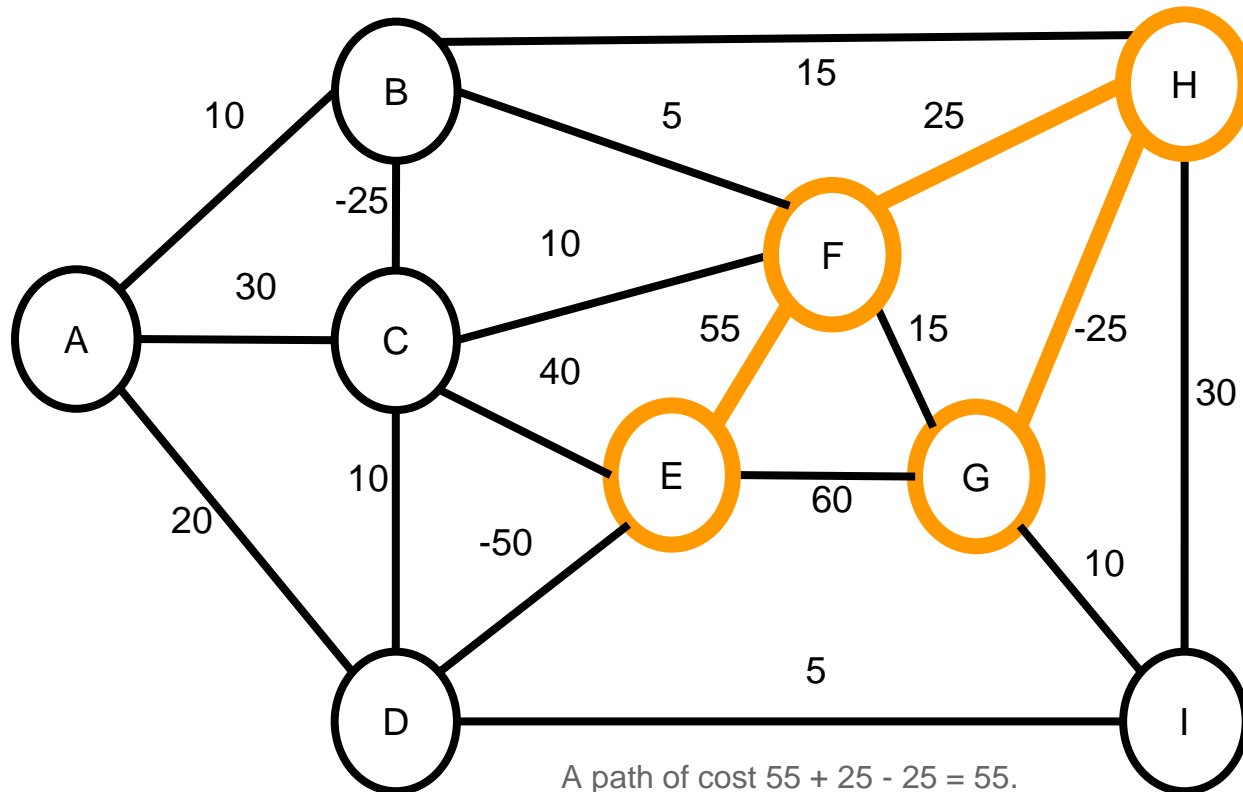


Graph Applications: Shortest paths


Lecture 29 by *Marina Barsky*

Paths with costs

In a weighted graph, the **cost of a path** is the sum of the weights (costs) of the edges along the path.



A path from x to y is a **minimum-cost path** if it has the smallest cost among all paths from x to y .




Single-source Minimum-Cost Paths

without negative edge weights

Algorithm by Dijkstra

Minimum Cost Paths: will simple greedy work?



The straightforward greedy approach: from each node on the path, take the edge with the smallest cost


Is 6 the cost of the minimum-cost path from s to t ?

Simply taking the smallest-weight edge out of the current node does not work!

Storing the Minimum Cost

We store the minimum cost from the start node in a `min_cost` array.

- The minimum cost from the start node to x is `min_cost[x]`.
- The start node s has `min_cost[s] = 0`.



There is a path of cost 20 from s to u .

There is a path of cost 10 from s to v .


Question: Will we ever need to change a `min_cost` value during the algorithm?

Or will the value be set once and then never change?

Storing the Minimum Cost

We store the minimum cost from the start node in a `min_cost` array.

- The minimum cost from the start node to x is `min_cost[x]`.
- The start node s has `min_cost[s] = 0`.



There is a path of cost 20 from s to u .


There is a path of cost 10 from s to v .

Question: Will we ever need to change a `min_cost` value during the algorithm?

Or will the value be set once and then never change? **The value can change**

Property of Minimum-Cost Paths

Suppose that a minimum cost path from x to z goes through the nodes y_1, y_2, \dots, y_k . Notice that the subpath from x to y_i is also a minimum cost path from x to y_i for all i .



A minimum cost path from x to z . One of its subpaths is a path from x to the intermediate node y_3 .

If there is another path from x to y_3 that has lesser cost than this subpath, then the original path from x to z was not minimum cost (could have been improved).

Therefore, if we want to build minimum cost paths, then we never want to extend a path that is not itself minimum-cost.

- Don't add node to the solution until we know that we have a minimum cost path to it.
- We need to keep track of the current minimum cost path to each node.
- We will compute the shortest path from 'source' node x to all other nodes y .

Dijkstra Algorithm: intuition

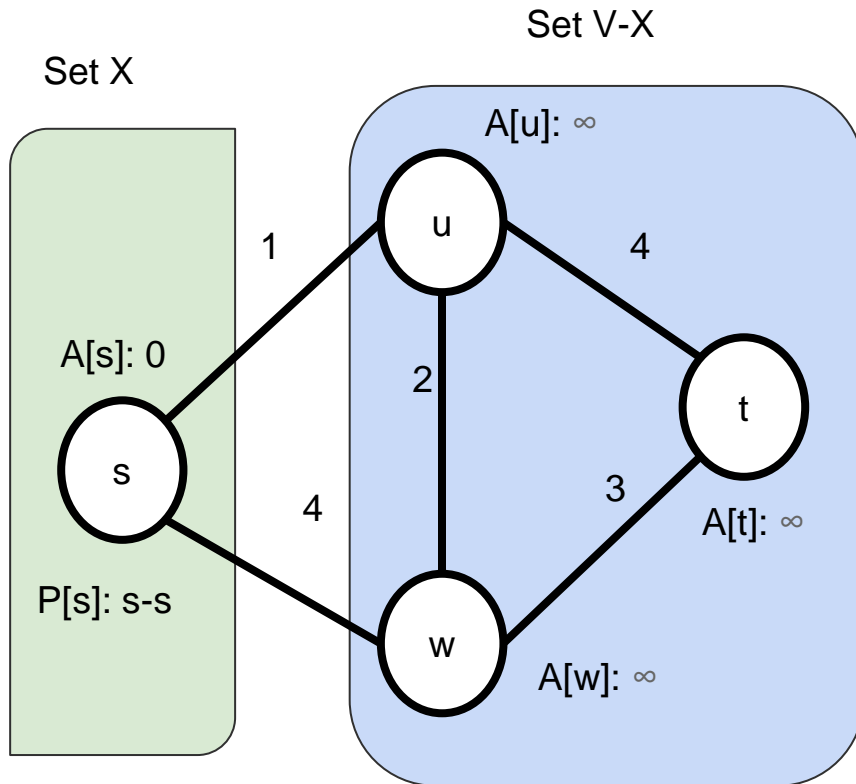
We maintain 2 sets of nodes:

Set X for nodes for which we already know the final cost of the min-cost paths from S (**Processed**).

Set $V-X$ of remaining nodes for which the min-cost path is yet to be found (**Unprocessed**).

- We perform n iterations of the main loop:
 - At each iteration we choose one node from $V-X$ and add it to X with its corresponding cost (and path if required).
 - The node is chosen according to the minimum *Dijkstra Greedy Score* (**DGS**).
 - We store the current greedy score for vertex v in the *min_cost* array
 - We grow set X until all n vertices from V are added to X .

Dijkstra algorithm: short illustration




A is a *min_cost* array containing DGS for each of n nodes

P is an array containing min-cost paths from s to each of n nodes

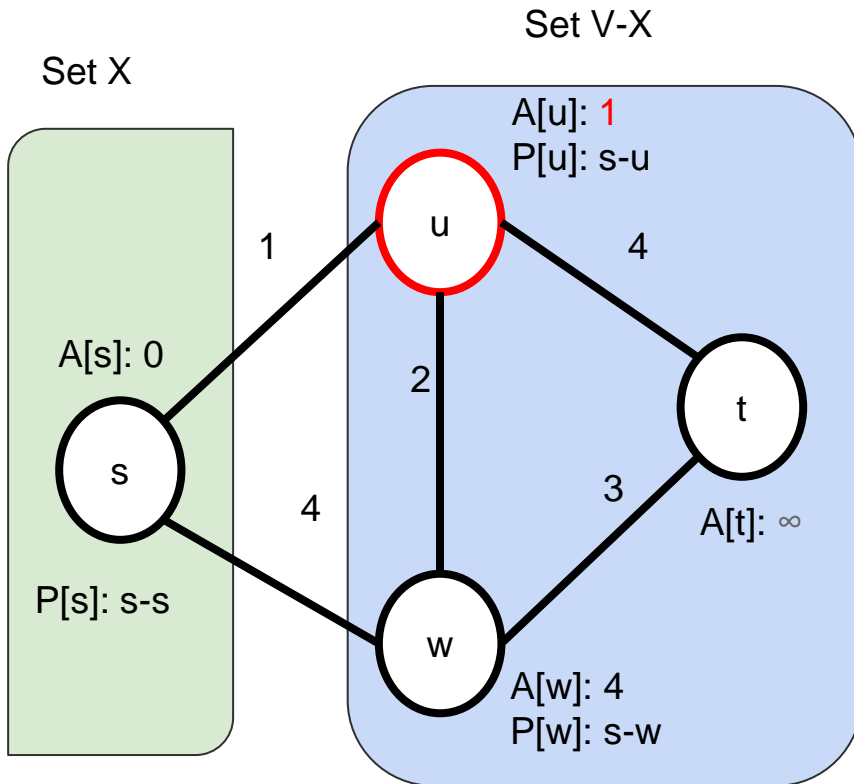
- Originally, only the source vertex S is in X : the cost of the path $S-S$ is 0.
- For paths from S to other nodes the cost is unknown, we mark them as ∞ .
- At each step, there will be edges inside X , inside $V-X$, and the edges between the 2 sets.
- **We are interested only in edges that “cross the border” - they will allow us to improve the DGS for each remaining node**

Dijkstra algorithm: short illustration




- The goal is to add more nodes to X.
- The only 2 edges extending already known min-cost path are (s, u) and (s, w) .
- For both u and w , we update their DGS to the sum of $A[s] + \text{cost}(s, u)$ and $A[s] + \text{cost}(s, w)$ respectively.
- This will be a new Dijkstra Greedy Score for these nodes.

Dijkstra algorithm: illustration



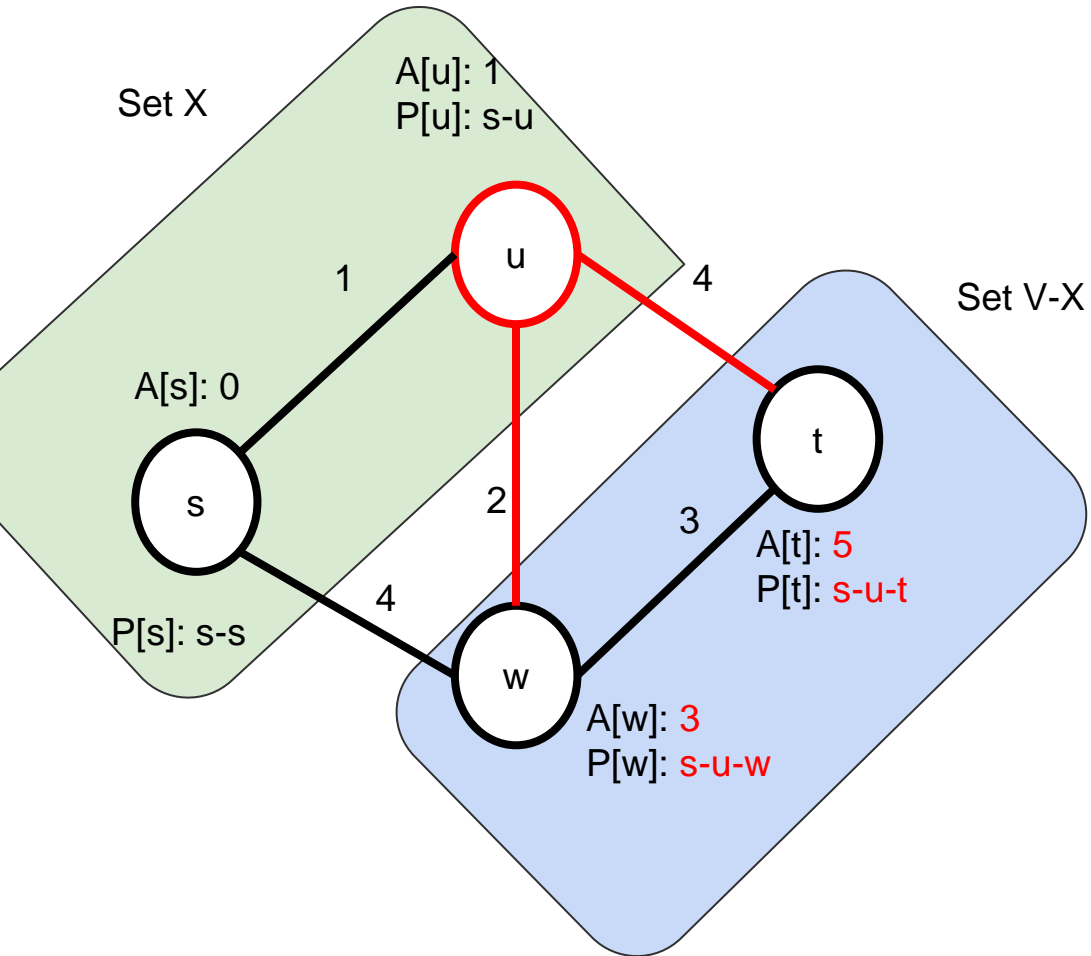
- Next, we select the vertex with the minimum DGS - vertex u - and add it to X

Dijkstra algorithm: illustration




- Now we have a new node u in X , and we know that $s \rightarrow u$ is the next smallest min-cost path from s
- There are 2 new edges out of u which cross the border between X and $V-X$
- They may help improve the DGS of remaining nodes

Dijkstra algorithm: illustration



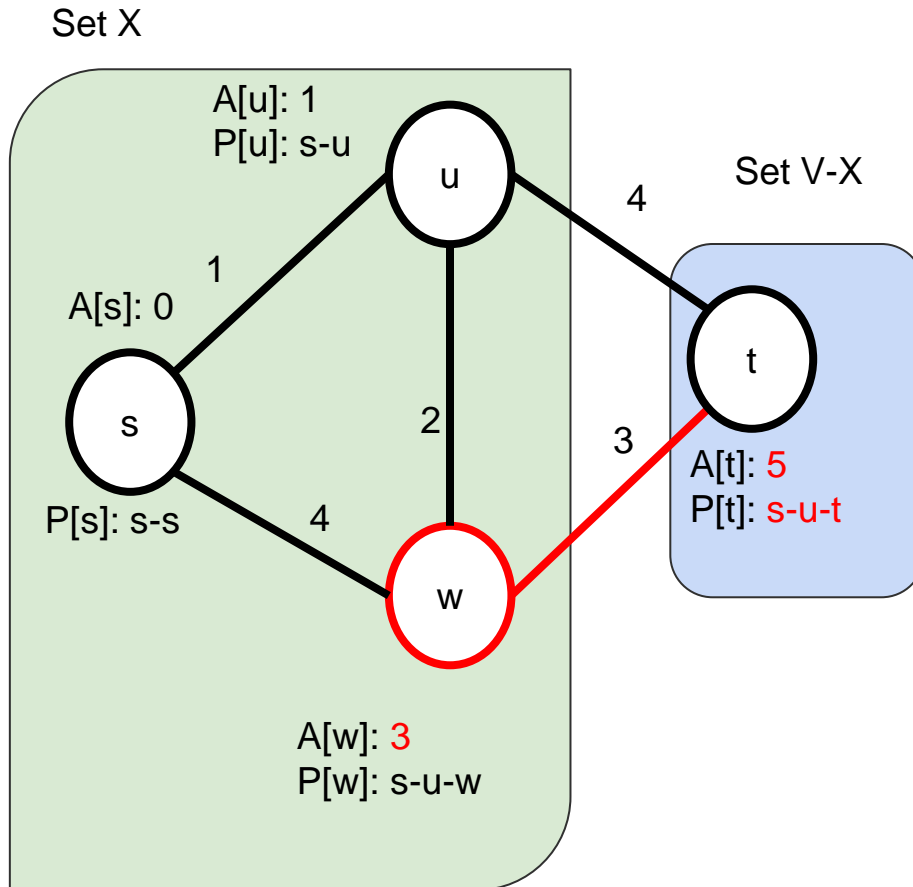
- We check if we can update the DGS using $A[u] + \text{cost}(u,t)$ and $A[u] + \text{cost}(u,w)$

Dijkstra algorithm: illustration



- Next we select the node with the smallest DGS and add it to X

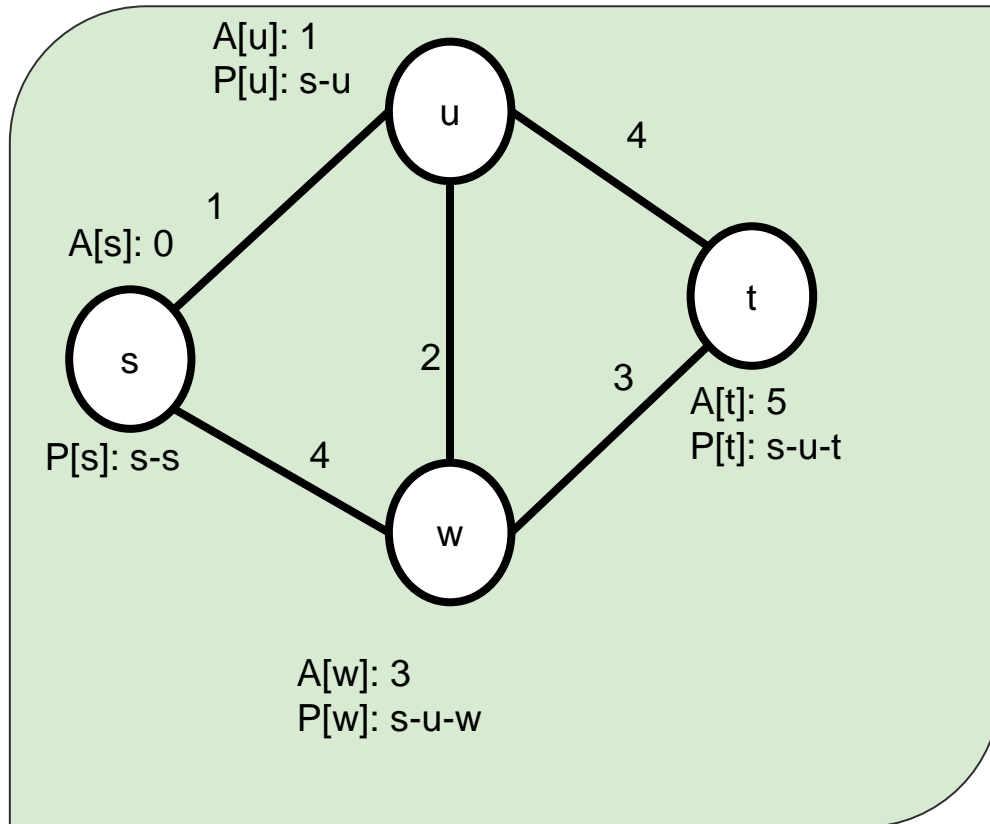
Dijkstra algorithm: illustration



- The only new edge that can update DGS for t is (w,t).
- We check if the new score going through w is better, it is not

Dijkstra algorithm: illustration


Set X



- The last vertex is added to X
- At this point all min-cost paths from s to each other vertex have been computed

Dijkstra algorithm: the paths

Set X



- Do we really need to store the paths themselves?
- No, instead of storing the min path for each node, we could just record the parent node when we update DGS, and we will be able to recover the shortest path from any node to s

Dijkstra Algorithm: correctness

Intuitively: the algorithm is correct because we transfer the node v into set X by extending the shortest paths from the nodes for which we already know that the paths from s are optimal.

Proof by induction (sketch):

- Base case: $A[s] = 0$
- Inductive hypothesis:
for all $v \in X$, $A[v]$ is the cost of the shortest path $s \rightsquigarrow v$
- In each iteration:
We pick the vertex $w \notin X$ with the lowest DGS among all vertices $\notin X$.
The path from s to w extends some shortest path $s \rightsquigarrow v$ for some $v \in X$.
We updated the $DGS(w)$ with the lowest possible cost of extending any such path
- Then any alternative path from s to w which we did not explore yet must go through some vertex z in $V-X$. But for any z , $DGS(z) \geq DGS(w)$, so any such path will have the cost at least $A[w]$ (not shorter).
- Hence, if we assume that each path from s to $v \in X$ was a shortest path, the extension of one of such paths will be a shortest path too.

A full formal correctness proof of Dijkstra's algorithm can be found [here](#)

Pseudocode

Algorithm *Dijkstra*(G , array of edge weights w , start)

```
unprocessed: = empty set
min_cost:= empty dictionary
for each  $u$  in vertices of  $G$ 
    min_cost[ $u$ ]: =  $\infty$ 
    unprocessed.add( $u$ )

min_cost[start]: = 0
processed: = empty set
processed.add(start)

while unprocessed is not empty
     $v$ : = remove  $v$  with min_cost from unprocessed
    processed.add( $v$ )
    for each edge ( $v, u$ )
        if  $u$  in unprocessed:
            min_cost[ $u$ ]: = min(min_cost[ $u$ ], min_cost[ $v$ ] +  $w_{v,u}$ )
```

Naive Dijkstra Algorithm

Running time of Dijkstra's Algorithm

Algorithm *Dijkstra*(G, w, start)

```
unprocessed: = empty set  
min_cost := empty dictionary
```

```
for each u in vertices of G  
    min_cost[u] :=  $\infty$   
    unprocessed.add(u)
```

```
min_cost[start] := 0  
processed: = empty set
```

```
while unprocessed is not empty
```

```
    v := remove v with min_cost from unprocessed  
    processed.add(v)
```

```
    for each edge (v,u)
```

```
        if u in unprocessed:
```

```
            min_cost[u] := min(min_cost[u], min_cost[v] +  $w_{v,u}$ )
```

Loop is
executed
 $O(n)$ times

Search for
min in set of
size $O(n)$

Each node may have degree $O(n)$ - but total
amortized $O(m)$ edges to process

The running time: $n \cdot n + m = O(n^2)$

Recap: Min-Priority Queue

A *min-priority queue* is an ADT for fast retrieval of min element.

Implementations: binary heap, balanced BST, *Fibonacci heap* (retrieval in time $O(1)$ but large constants).

For Dijkstra: Priority Queue ADT should be enhanced with the *update** operation.

	Priority queue
enqueue	$O(\log n)$ -time
dequeue	$O(\log n)$ -time
update	$O(\log n)$ -time

*The *update* operation decreases the associated value of a given item.

In other words, it increases its priority.

We can keep pointers to each queue node to locate it quickly.

However if the priority of the heap node changed, we need then repair the heap – $O(\log n)$

Dijkstra's Algorithm with Priority Queue

The `min_cost` priority queue (`min_pq`) stores tuples `(node, DGS)` prioritized by DGS.

```
Algorithm Dijkstra Improved(G, start)
```

```
min_pq := empty priority queue
```

```
for each u in vertices of G
```

```
    min_pq.enqueue((u, ∞))
```

```
processed = empty set
```

```
min_pq.update((start, 0))
```

```
while min_pq is not empty
```

```
    (cost_v, v) := min_pq.dequeue()
```

```
    processed.add(v, cost_v)
```

```
    for each edge (v, u):
```

```
        if u in min_pq:           # we have pointer to each node in the extended priority queue
```

```
            cost_u := min_pq.get(u).cost
```

```
            if cost_v +  $w_{v,u}$  < cost_u:
```

```
                min_pq.update(u, cost_v +  $w_{v,u}$ )
```

Dijkstra's with Priority Queue: running time

Algorithm *Dijkstra Improved*(G, start)

```
min_pq:= empty priority queue
for each u in vertices of G
    min_pq.enqueue((u, ∞))
```

```
processed:= empty set
min_pq.update((start, 0))
```

```
while min_pq is not empty
    (cost_v, v) := min_pq.dequeue()
    processed.add(v, cost_v)
```

Loop is
executed
 $O(n)$ times

Each dequeue in time $\log n$

```
for each edge (v, u):
```

```
    if u in min_pq:
```

```
        cost_u:= min_pq.get(u).cost
```

```
        if cost_v +  $w_{v,u}$  < cost_u:
```

```
            min_pq.update(u, cost_v +  $w_{v,u}$ )
```

In sum $O(m)$ edges to process

Quickly finds u in
min_pq, and
updates only if
new DGS is
better: rebalance
in time $O(\log n)$

Running time $O(n \log n) + O(m \log n) = O(m \log n)$

Dijkstra Algorithm: running time

Running time with Priority Queue: $O(m \log n)$

- If $m = O(n)$ [sparse graphs], then running time $O(n \log n)$
- If $m = O(n^2)$ [dense graphs], then running time is $O(n^2 \log n)$


Dijkstra Algorithm: non-negative weights

- The algorithm combines ideas from both greedy and iterative improvement techniques
- It iteratively improves DGS of each node until no more improvement is possible and at this point the node is transferred into a processed set X
- However if edges are allowed to have a negative cost, then some of them could potentially improve the DGS of already processed nodes (including the source node s !)
- Then we would never have the processed set to start with
- Therefore this algorithm is not applicable for graphs with negative edge weights


Group Activity: step-by-step run of
Dijkstra's algorithm

Dijkstra's Algorithm: full example

Find all minimum cost paths from the source node C.



Dijkstra's Algorithm: full example



We start by assigning Dijkstra Greedy Score (DGS) to each node as ∞

The only known min-cost path is C-C of length 0.

We know that it cannot be improved so we add it to the Processed nodes (green)


Known
shortest paths
from C

To v_i	Shortest path
C	C-C: 0

Remaining nodes
with their Dijkstra
Greedy Score

	DGS
A	∞
B	∞
D	∞
E	∞
F	∞
G	∞

Dijkstra's Algorithm: full example



Update DGS for all nodes adjacent to C.
 Improve their DGS using edges that cross Processed and Unprocessed sets.


Known shortest paths from C

To v_i	Shortest path
C	C-C: 0

Remaining nodes with their Dijkstra Greedy Score

	DGS
A	5
B	20
D	10
E	15
F	25
G	∞

Dijkstra's Algorithm: full example



Select the node with min DGS and add it to known min-cost paths.


Known shortest paths from C

To v_i	Shortest path
C	C-C: 0

Remaining nodes with their Dijkstra Greedy Score

	DGS
A_{c-a}	5
B_{c-b}	20
D_{c-d}	10
E_{c-e}	15
F_{c-f}	25
G	∞

Dijkstra's Algorithm: full example



Update DGS for every node v adjacent to A: cost of path(C-A) + cost of edge(A, v)
 Select the node with min DGS and add it to Processed


Known shortest paths from C

To v_i	Shortest path
C	C-C: 0
A	C-A: 5

Remaining nodes with their Dijkstra Greedy Score

	DGS
B_{C-a-b}	20 15
D_{C-d}	10
E_{C-e}	15
F_{C-f}	25
G	∞

Dijkstra's Algorithm: full example



Select the node with min DGS and add it to known min-cost paths.


Known shortest paths from C

To v_i	Shortest path
C	C-C: 0
A	C-A: 5
D	C-D: 10

Remaining nodes with their Dijkstra Greedy Score

	DGS
B_{C-a-b}	15
E _{C-e}	15
F _{C-f}	25
G	∞

Dijkstra's Algorithm: full example



Update DGS for all unprocessed nodes v adjacent to B: $\text{cost of min path}(C-B) + \text{cost}(B,v)$
 Select the node with min DGS and add it to known min-cost paths


Known shortest paths from C

To v_i	Shortest path
C	C-C: 0
A	C-A: 5
D	C-D: 10
B	C-A-B: 15

Remaining nodes with their Dijkstra Greedy Score

	DGS
E_{c-e}	15
F_{c-f}	25
G	∞

Dijkstra's Algorithm: full example



Update DGS for all unprocessed nodes v adjacent to E

Select the node with min DGS and add it to known min-cost paths


Known shortest paths from C

Remaining nodes with their Dijkstra Greedy Score

T o v _i	Shortest path
C	C-C: 0
A	C-A: 5
D	C-D: 10
B	C-A-B: 15
E	C-E: 15

	DGS
F_{c-e-f}	25 20
G_{c-e-g}	40

Dijkstra's Algorithm: full example



Update DGS for all unprocessed nodes v adjacent to F :
 $\text{len}(C-F) + \text{len}(F,v)$.
 This is the last node - mark it as processed.


Known
shortest paths
from C

To v_i	Shortest path
C	C-C: 0
A	C-A: 5
D	C-A-D: 10
B	C-A-B: 15
E	C-E: 15
F	C-E-F: 20

Remaining nodes
with their Dijkstra
Greedy Score

	DGS
$G_{c-e-f-g}$	40 30

Dijkstra's Algorithm: full example




All min-cost paths from C to any other node have been computed.

All shortest paths
from C

To v_i	Shortest path
C	C-C: 0
A	C-A: 5
D	C-D: 10
B	C-A-B: 15
E	C-E: 15
F	C-E-F: 20
G	C-E-F-G: 30

Traceback



Of course, instead of storing the actual min path for each node, we could just store the cost of the path and the link to the parent node when we update DGS, and we will be able to find the shortest path from any node to C