

# Dijkstra's Algorithm:

Bellman-Ford tightens every edge  $n-1$  times. Inefficient!

Dijkstra tightens each edge once.  
Works harder to find the right edge to tighten.

In each round, Dijkstra "visits" a vertex, tightening all edges out of the vertex.

Chooses the unvisited vertex closest to  $S$ .  
[But we don't know all distances yet! So it picks the vertex of minimum  $c(\cdot)$ .]

Dijkstra  $(G, s)$ :

$$c(u) = \infty \quad \forall u$$

$$c(s) = 0$$

$$S = \{s\}$$

While  $S \neq V$ :

find  $u \in (V - S)$  minimizing  $c(u)$ .

$$S \leftarrow S + \{u\}$$

For each edge  $(u, v)$  from  $u$  in  $E$ :

$$c(v) \leftarrow \min(c(v), c(u) + \text{cost}(u \rightarrow v))$$

tighten  $(u, v)$

"visit  $u$ "

## Correctness

For simplicity, suppose  $c^*(u)$  all unique  
[Full proof on Piazza]

Can order  $u \in V$  by distance from  $s$ :  
 $s = u_1, u_2, \dots, u_n$   
 $c(u_1) < c(u_2) < c(u_3) < \dots < c(u_n)$

Lemma:

The  $k^{\text{th}}$  node visited =  $u_k$   
and  $c(u_k) = c^*(u_k)$  when it is visited.

PF Trivial for  $k=1$ .

If true for all  $k' < k$ ,  
then consider the state just before  
choosing the  $k^{\text{th}}$  node to visit.

Claim:  $c(u_k) = c^*(u_k)$ .

Let  $u' = \text{pred}(u_k)$  = previous node to  $u_k$   
in shortest  $s \rightarrow u_k$  path.

$$c^*(u') = c^*(u_k) - \text{cost}(u' \rightarrow u_k) \\ \leq c^*(u_k)$$

because **nonnegative edges**.

Uniqueness assumption  $\Rightarrow c^*(u') < c^*(u_k)$

$\Rightarrow u'$  before  $u_k$  in order

inductive hypo

$\Rightarrow$  already visited  $u'$ , and

$c(u') = c^*(u')$  when it was visited

$\Rightarrow$  when visited  $u'$  we set

$$c(u_k) \leftarrow \min(c(u_k), \underbrace{c^*(u') + \text{cost}(u' \rightarrow u_k)}_{c^*(u_k)}) \\ = c^*(u_k).$$

So we have  $c(u_k) = c^*(u_k)$

When deciding on the  $k^{\text{th}}$  node to visit,  
we've already visited  $u_i$   $\forall i < k$ ,  
and all other  $i$  have

$$c(u_i) \geq c^*(u_i) > c^*(u_k) = c(u_k).$$

Hence Dijkstra will choose  $u_k$  in the  $k^{\text{th}}$  round,  
with  $c(u_k) = c^*(u_k)$ .  $\square$

Since we visit every node, and  $c(u) = c^*(u)$   
when it is visited, Dijkstra eventually  
gets each  $c(u) = c^*(u)$ , proving correctness.

## Running time

Time =  $O(\text{time to tighten } m \text{ edges}$   
 $+ \text{time to find the } n \text{ vertices}$   
 $\text{to visit})$

Simplest approach:

Look through all  $V$  to decide  
node to visit,

$\Rightarrow O(1)$  tighten,  $O(n)$  time to find each  $u$ .

$\Rightarrow O(m+n^2) = O(n^2)$  running time.

Better than Bellman-Ford!

Better approach: store  $V-S$  in a binary heap keyed by  $c()$

node to visit = delete-min on heap  
=  $O(\log n)$  time,

but now  $\text{tighten}()$  changes a  $c()$

$\Rightarrow$  need to bubble up that node in heap  
"decrease-key operation"  
=  $O(\log n)$  time.

$\Rightarrow$  total time =  $O(m \log n + n \log n)$   
=  $O(m \log n)$ .

Fanciest approach: use a **Fibonacci heap**

delete-min:  $O(\log n)$

decrease-key:  $O(1)$

(amortized)

$\Rightarrow O(m + n \log n)$  time.

[in practice, use a binary heap]