


Complexity & Algorithms, Spring 2026


Shortest Paths



Single Source Shortest paths (SSSP)

Some notes:

Why a tree?



If $s \rightarrow a \rightarrow b \rightarrow c \rightarrow d \rightarrow v$ and $s \rightarrow a \rightarrow x \rightarrow y \rightarrow d \rightarrow u$ are shortest paths, then $s \rightarrow a \rightarrow b \rightarrow c \rightarrow d \rightarrow u$ is also a shortest path.

What about negative cycles or edges?




Figure 8.3. There is no shortest walk from s to t .

Also: If undirected, can simulate with a directed graph:



Unless you have negative edges.
(It gets weird.)





Figure 8.4. An undirected graph where shortest paths from s are unique but do not define a tree.

Important to realize:

SSSP \neq Minimum Spanning Tree



Why?

based on source

vs global min object

Computing a SSSP.

(Ford 1956 + Dantzig 1957)

Each vertex will store 2 values.


(Think of these as tentative shortest paths.)

- $\text{dist}(v)$ is length of tentative shortest path $s \rightsquigarrow v$

(or ∞ if don't have an option yet)

- $\text{pred}(v)$ is the predecessor of v on that tentative path $s \rightsquigarrow v$ (or NULL if none)

Initially:

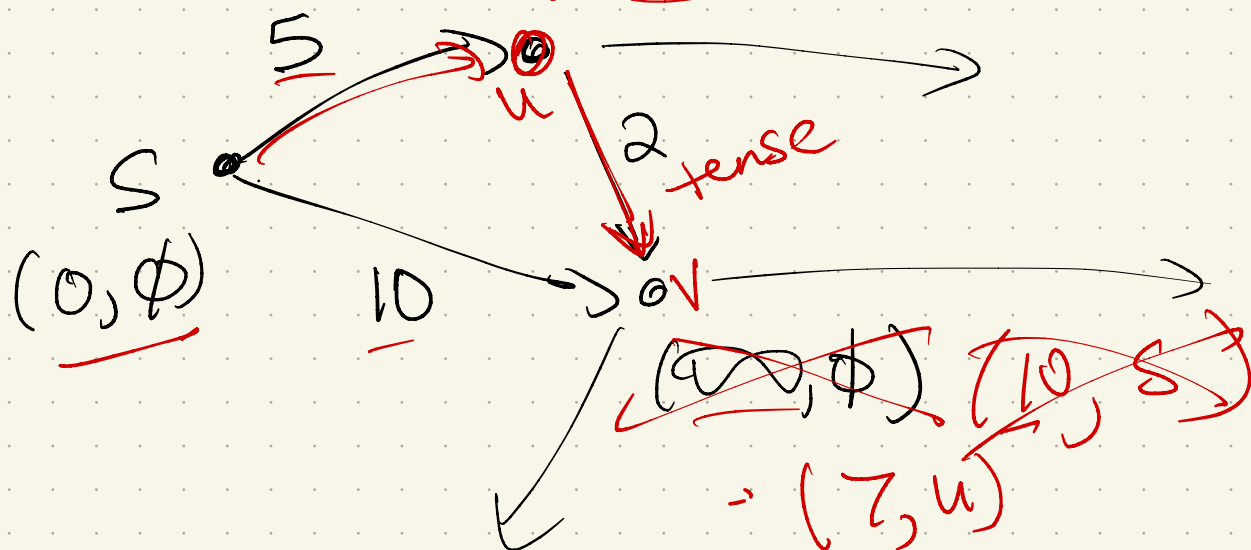


We say an edge \vec{uv} is tense if

$$\underline{\text{dist}(u)} + \underline{w(u \rightarrow v)} < \underline{\text{dist}(v)}$$

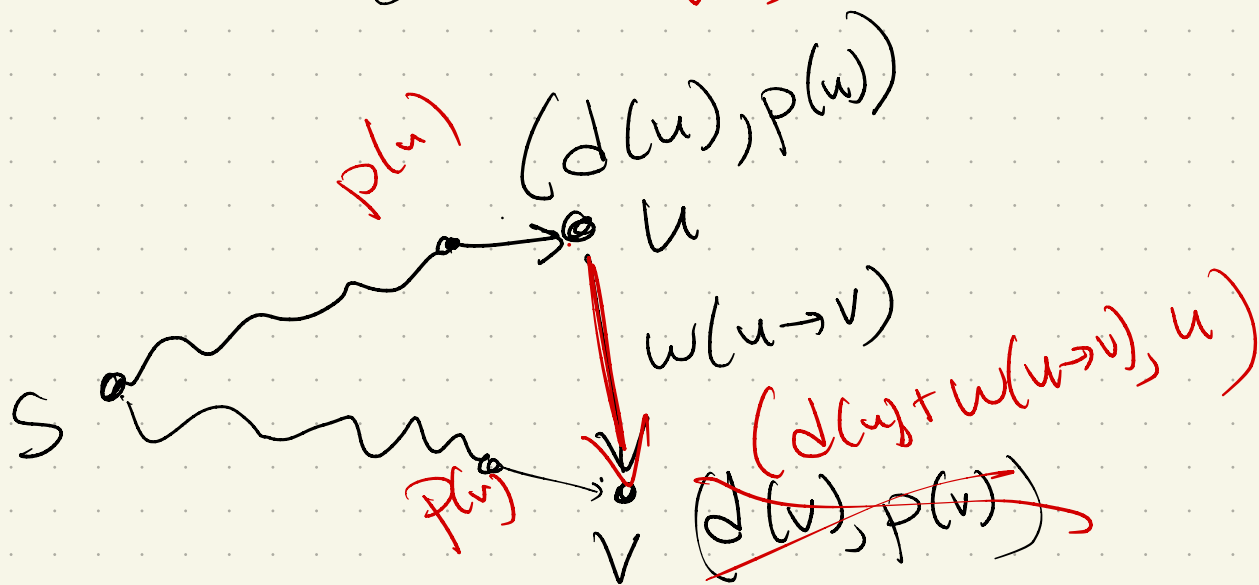
~~(∞, ϕ)~~ $(\underline{5}, \underline{8})$

Initially:



Here:

In general:



Key idea of all algorithms:

Find tense edges & relax them:

RELAX($u \rightarrow v$):

$dist(v) \leftarrow dist(u) + w(u \rightarrow v)$

$pred(v) \leftarrow u$

Then:

INITSSSP(s):

$dist(s) \leftarrow 0$

$pred(s) \leftarrow \text{NULL}$

for all vertices $v \neq s$

$dist(v) \leftarrow \infty$

$pred(v) \leftarrow \text{NULL}$

GENERICSSSP(s):

INITSSSP(s)

put s in the bag

while the bag is not empty

take u from the bag

for all edges $u \rightarrow v$

if $u \rightarrow v$ is tense

RELAX($u \rightarrow v$)

put v in the bag

Claim: At any point in time, $\text{dist}(v)$ is either ∞ or the length of some $s \rightsquigarrow v$ walk (which uses $\leq i$ edges, $i = \text{loop iteration}$)

Proof: Induction on while loop iterations.

Base case: loop iteration 1

at beginning, s has $\text{dist} = 0$ +

all others = ∞

at end, s has $\text{dist} = 0$ still,

+ all neighbors u now have

$\text{dist}(u) = w(s \rightarrow u)$, which is a length 1 walk. (Others are ∞)


Ind hyp:

In iteration $k-1$, the claim is

true: all are $= \infty$, or some $s \rightsquigarrow v$ walk

Ind Step:

In iteration k : At beginning, we take out some vertex u .



By IH, $\text{dist}(u)$ is the weight of some $s \rightsquigarrow u$ walk.

At end, all nbrs v of u are either unchanged (or so by IH are still either ∞ or length of $s \rightsquigarrow v$ walk)

or $u \rightarrow v$ was tense, \rightarrow

$$\text{now } \text{dist}(v) = \text{dist}(u) + w(u \rightarrow v).$$

Since $\text{dist}(u)$ is a $s \rightsquigarrow u$ walk,

then $\text{dist}(v)$ is weight of the

walk $(s \rightsquigarrow u) + (u \rightarrow v)$, which

is a walk with one more edge

at end.

(All other vertices are unchanged,
so by IH are still ∞ or a $s \rightsquigarrow v$
walk.) \square

Dijkstra's algorithm:

Algorithm 2: Dijkstra's Algorithm

Input: Digraph $G = (V, E)$ with edge-weights $w_e \geq 0$ and source vertex $s \in G$

Output: The shortest-path distances from s to each vertex

2.1 add s to heap with key 0

2.2 **for** $v \in V \setminus \{s\}$ **do**

2.3 | add v to heap with key ∞

2.4 **while** heap not empty **do**


2.5 | $u \leftarrow \text{deletemin}$

2.6 | **for** v a neighbor of u **do**

2.7 | | key(v) $\leftarrow \min\{\text{key}(v), \text{key}(u) + w_{uv}\}$ // relax uv

min heap:

*Extract Min
replace old value if tense*




Four phases of Dijkstra's algorithm run on a graph with no negative edges. At each phase, the shaded vertices are in the heap, and the bold vertex has just been scanned. The bold edges describe the evolving shortest path tree.

Logic:

at each iteration, grab
closest vertex.

All others can only have
greater distance



Runtime:

Algorithm 2: Dijkstra's Algorithm

Input: Digraph $G = (V, E)$ with edge-weights $w_e \geq 0$ and source vertex $s \in G$

Output: The shortest-path distances from s to each vertex

```
2.1 add  $s$  to heap with key 0
2.2 for  $v \in V \setminus \{s\}$  do
2.3   | add  $v$  to heap with key  $\infty$ 
2.4 while heap not empty do
2.5   |  $u \leftarrow \text{deletemin}$ 
2.6   | for  $v$  a neighbor of  $u$  do
2.7     |  $\text{key}(v) \leftarrow \min\{\text{key}(v), \text{key}(u) + w_{uv}\}$  // relax  $uv$ 
```

$V \log V$ time to do anything in heap

→ How many times can loop repeat?

V times, since we never re-add vertex

Each time: alter key + deleteMin

\swarrow \downarrow
 $O(\log_2 V)$ binary heap $O(\log_2 V)$

$V \log V \rightarrow$ initialize

repeat V times:

delete $\leftarrow \log V$

$d(x) \rightarrow$ for each incident edge
alter key $= \log V$

$$\sum_{x \in V} d(x) \cdot \log V = E \log V$$

$O(E \log V)$

Other data structures:

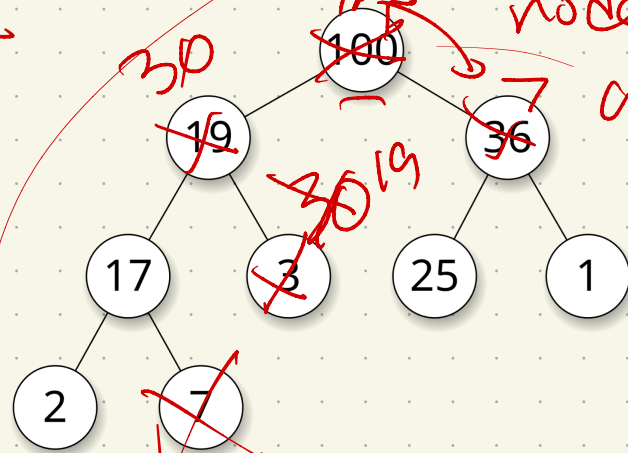
Heaps are classic.

→ Priority queue

You probably saw binary heaps:

(This is a max binary heap →)


Tree representation



For each node, children are less than node's value

$\log_2 n$

Array representation



$i \rightarrow 2i$
 $i \rightarrow 2i+1$

Next time?

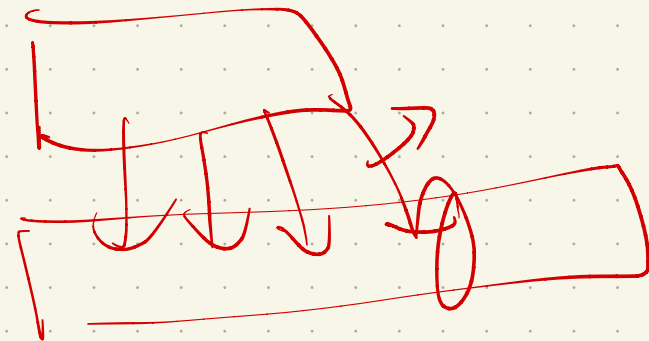
More data structures!

◦ Fibonacci Heaps

◦ Van Emde Boas → probably
a bit too long...

→ Binomial heap ★

Amortized analysis



Bellman-Ford:


Relax edges for a while.

Stop when every edge has been relaxed at least once

any path has length $\leq V-1$

If any one is still tense:

you've relaxed ≥ 2 times!



Runtime: $O(V \cdot E)$

```
BELLMANFORD(s)
INITSSSP(s)
repeat V - 1 times
  for every edge u → v
    ( if u → v is tense
      RELAX(u → v) )
for every edge u → v
  if u → v is tense
    return "Negative cycle!"
```


How to prove correctness?

Notation:

Let $\text{dist}_{\leq i}^o(v) :=$
 length of shortest s to v path using
 $\leq i$ edges

\hookrightarrow any path uses $\leq |V|$

Ex:



	<u>s^o</u>	<u>v_1^o</u>	<u>v_2^o</u>	<u>v_3^o</u>
≤ 1	0	10	1	∞
≤ 2	0	4	1	6
≤ 3	0	4	1	6

Claim: $\forall v \neq i$, after i iterations of B-F,
 $\text{dist}(v)$ $\leq \text{dist}_{\leq i}^0(v)$

Induction on i :


\uparrow # iterations

BC: $i=0$: $s=0$
all other v , $\text{dist}_{\leq 0}^0(v) = \infty$

IH: After $i-1$ iterations, all tentative guesses are $\leq \text{dist}_{\leq i-1}^0(v)$.

IS: Now consider $\text{dist}_{\leq i}^0(v)$:

built from a path \rightarrow



We know in round $i-1$, $\text{dist}(x) \leq d_{\leq i-1}(x)$
 $\forall x \in V$.

Consider $u \rightarrow v$ in next round:

It was tense:

got shorter

or not:

same iff holds

Then:

Any simple path uses $\leq V-1$ edges, so $d(x) \leq d_{\leq V-1}(x)$.

The result:

Algorithm 3: The Bellman-Ford Algorithm

Input: A digraph $G = (V, E)$ with edge weights $w_e \in \mathbb{R}$, and source vertex $s \in V$

Output: The shortest-path distances from s to each vertex, or report that a negative-weight cycle exists

```
3.1  $dist(s) = 0$  // the source has distance 0
3.2 for  $v \in V$  do
3.3 |  $dist(v) \leftarrow \infty$ 
3.4 for  $|V|$  iterations do
3.5 | for edge  $e = (u, v) \in E$  do
3.6 | |  $dist(v) \leftarrow \min\{dist(v), dist(u) + weight(e)\}$ 
3.7 If any distances changed in the last ( $n^{th}$ ) iteration, output "G has a negative weight cycle".
```

Runtime:

$O(V \cdot E)$

Shortest paths from one source:

$O(E \log V)$ if no negative edges

$O(V \cdot E)$ if negative

if no weights: BFS $O(V+E)$

~~II...x..I~~ ~~all~~ Computing MSSPs \leftarrow all pairs shortest paths

First attempt: [for each vertex v
Run SSSP(v)

Runtime : $O(V \times (\text{SSSP alg}))$

• no negative edge weights, Dijkstra was
 $O(E \log V) \Rightarrow O(VE \log V)$

• Bellman-Ford was $O(VE)$
 $\Rightarrow O(V^2 E) \leq O(V^4)$

Can we do better? Yes!

Side question:

Since negative edges are bad, why can't we just re-weight?

Idea: Increase all edge weights by some amount, so all positive. $+C$

Doesn't work:




Figure 9.1. Increasing all the edge weights by 2 changes the shortest path from s to t.

Why? # of edges is also additive: $+ (\text{path length}) \cdot C$


Another idea (that works):

Suppose each v has a price attached,

$\pi(v)$.

← anything

(Still have edge weights.)




Define a new function w' :

$$w'(u \rightarrow v) = \pi(u) + w(u \rightarrow v) - \pi(v)$$

What happens to paths? → goal is consistent sale

Claim: under w' , shortest paths are the same as under w .

Why? Consider 2 $S \rightarrow V$ paths:



$$\begin{aligned}
 & (\pi(S) + w(S \rightarrow v_1) - \pi(v_1)) + (\pi(v_1) + w(v_1 \rightarrow v_2) - \pi(v_2)) \\
 & + (\pi(v_2) + w(v_2 \rightarrow v_3) - \pi(v_3)) + \dots
 \end{aligned}$$

→ blue path:

$$\pi(S) + \text{length under } w - \pi(V) \Rightarrow \pi(S) + \text{length under } w - \pi(V)$$

Johnson's algorithm for MSSP:
use this new kind of weight
function!

- Run Bellman-Ford once from some $s \in V$
↳ no negative cycles or give up
- Reweight (if it works):

Let $\pi(v) =$ dist from s to v
compute via BF = $d(s, v)$

$$\begin{aligned} \& w'(u \rightarrow v) &= \pi(u) + w(u \rightarrow v) - \pi(v) \\ &= \text{dist}(s, u) + w(u \rightarrow v) - \text{dist}(s, v) \end{aligned}$$


Claim: Under w' , all weights are ≥ 0 .
Why?

use Dijkstra

No tense edges: s

$$\text{so } d(s, u) + w(u \rightarrow v) \geq d(s, v)$$

not $<$



$$\text{Then } \pi(u) = d(s, u)$$

$$\pi(v) = d(s, v)$$

$$\& w'(u \rightarrow v) = d(s, u) + w(u \rightarrow v) - d(s, v) \geq 0$$

So - Algorithm:

Runtime:

JOHNSONAPSP(V, E, w) :

⟨⟨Add an artificial source⟩⟩

add a new vertex s

for every vertex v

add a new edge $s \rightarrow v$

$w(s \rightarrow v) \leftarrow 0$

⟨⟨Compute vertex prices⟩⟩

$dist[s, \cdot] \leftarrow \text{BELLMANFORD}(V, E, w, s)$

if BELLMANFORD found a negative cycle

fail gracefully

⟨⟨Reweight the edges⟩⟩

for every edge $u \rightarrow v \in E$

$w'(u \rightarrow v) \leftarrow dist[s, u] + w(u \rightarrow v) - dist[s, v]$

⟨⟨Compute reweighted shortest path distances⟩⟩

for every vertex u

$dist'[u, \cdot] \leftarrow \text{DIJKSTRA}(V, E, w', u)$

⟨⟨Compute original shortest path distances⟩⟩

for every vertex u

for every vertex v

$dist[u, v] \leftarrow dist'[u, v] - dist[s, u] + dist[s, v]$

Figure 9.2. Johnson's all-pairs shortest paths algorithm

not
B-F
 $\sqrt{VE} \log V \rightarrow$

wird
II

$O(V^3)$

Compare to naive: $\underline{O(V^2 E)} = O(V^4)$

Can we get better?

Floyd-Warshall:

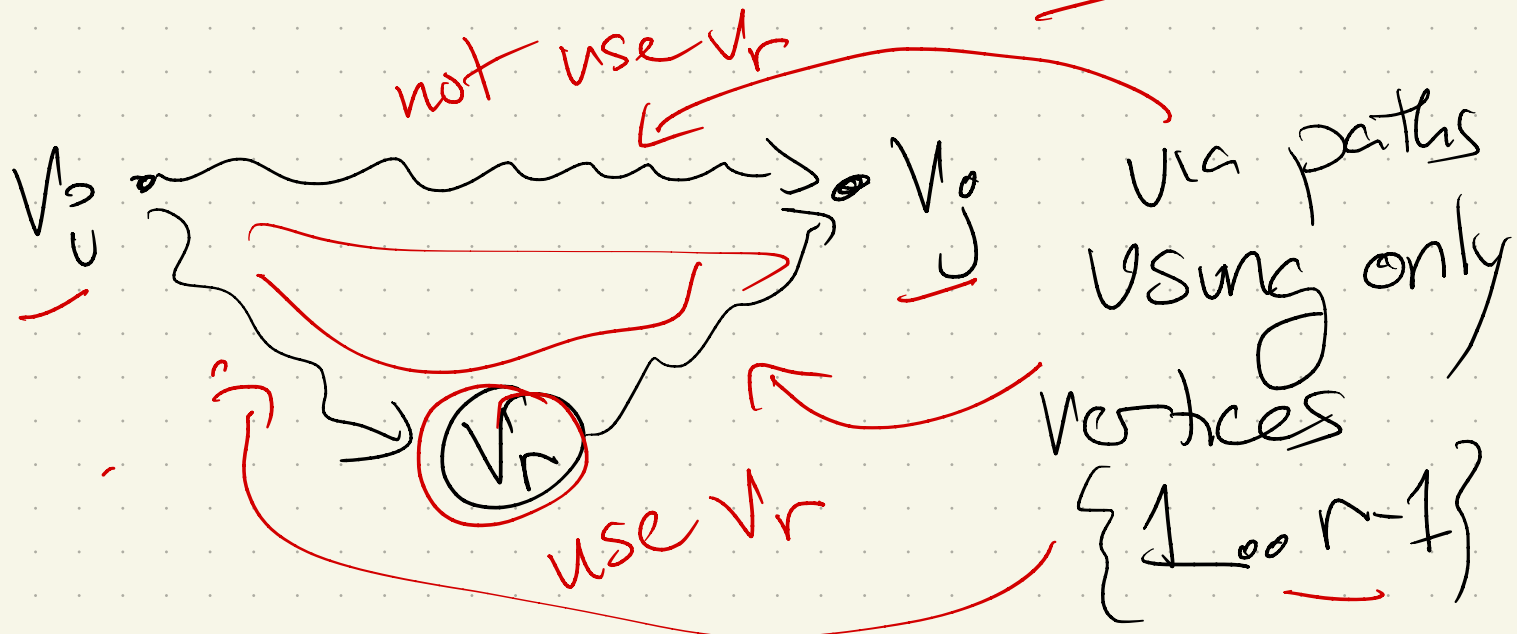
order vertices

$1, \dots, v \rightarrow v_1, \dots, v_r$

Let $d(x, y, r) =$

best length path via
vertices labeled $1, \dots, r$

Then
recursion




Recursion:

$$\text{dist}(u, v, r) = \begin{cases} w(u \rightarrow v) & \text{if } r = 0 \\ \min \left\{ \begin{array}{l} \text{dist}(u, v, r-1) \\ \text{dist}(u, r, r-1) + \text{dist}(r, v, r-1) \end{array} \right\} & \text{otherwise} \end{cases}$$

use vertex r

don't use vertex r



code \rightsquigarrow

KLEENEAPSP(V, E, w):

```
for all vertices  $u$ 
  for all vertices  $v$ 
     $dist[u, v, 0] \leftarrow w(u \rightarrow v)$ 

for  $r \leftarrow 1$  to  $V$ 
  for all vertices  $u$ 
    for all vertices  $v$ 
      if  $dist[u, v, r-1] < dist[u, r, r-1] + dist[r, v, r-1]$ 
         $dist[u, v, r] \leftarrow dist[u, v, r-1]$ 
      else
         $dist[u, v, r] \leftarrow dist[u, r, r-1] + dist[r, v, r-1]$ 
```

$O(V^3)$

Runtime:

$\leftarrow V$ matrices


$O(V^3)$

save space


FLOYDWARSHALL(V, E, w):

```
for all vertices  $u$ 
  for all vertices  $v$ 
     $dist[u, v] \leftarrow w(u \rightarrow v)$ 

for all vertices  $r$ 
  for all vertices  $u$ 
    for all vertices  $v$ 
      if  $dist[u, v] > dist[u, r] + dist[r, v]$ 
         $dist[u, v] \leftarrow dist[u, r] + dist[r, v]$ 
```



use r if it's better



vertices
1-r-1

