

Lecture 15: Single-Source Shortest Paths

Michael Dinitz

October 15, 2024

601.433/633 Introduction to Algorithms

Introduction

Setup:

- ▶ Directed graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$
- ▶ Length $\ell(\mathbf{x}, \mathbf{y})$ on each edge $(\mathbf{x}, \mathbf{y}) \in \mathbf{E}$ (equivalent: $\ell : \mathbf{E} \rightarrow \mathbb{R}$)
- ▶ Length of path \mathbf{P} is $\ell(\mathbf{P}) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathbf{P}} \ell(\mathbf{x}, \mathbf{y})$
- ▶ $\mathbf{d}(\mathbf{x}, \mathbf{y}) = \min_{\mathbf{x} \rightarrow \mathbf{y} \text{ paths } \mathbf{P}} \ell(\mathbf{P})$

Introduction

Setup:

- ▶ Directed graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$
- ▶ Length $\ell(\mathbf{x}, \mathbf{y})$ on each edge $(\mathbf{x}, \mathbf{y}) \in \mathbf{E}$ (equivalent: $\ell : \mathbf{E} \rightarrow \mathbb{R}$)
- ▶ Length of path \mathbf{P} is $\ell(\mathbf{P}) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathbf{P}} \ell(\mathbf{x}, \mathbf{y})$
- ▶ $\mathbf{d}(\mathbf{x}, \mathbf{y}) = \min_{\mathbf{x} \rightarrow \mathbf{y} \text{ paths } \mathbf{P}} \ell(\mathbf{P})$

Today: source $\mathbf{v} \in \mathbf{V}$, want to compute shortest path from \mathbf{v} to every $\mathbf{u} \in \mathbf{V}$

- ▶ $\mathbf{d}(\mathbf{u}) = \mathbf{d}(\mathbf{v}, \mathbf{u})$ for all $\mathbf{u} \in \mathbf{V}$
- ▶ Representation: “shortest path tree” out of \mathbf{v} .
- ▶ Often only care about distances – can reconstruct tree from distances.

Bellman-Ford

Dynamic Programming Approach

Subproblems:

- ▶ $OPT(u, i)$: shortest path from v to u that uses at most i hops (edges)
- ▶ If no such path, set to “infinitely long” fake path.
- ▶ For simplicity, create loop (edge to and from the same node) at every node, length 0

Dynamic Programming Approach

Subproblems:

- ▶ $OPT(u, i)$: shortest path from v to u that uses at most i hops (edges)
- ▶ If no such path, set to “infinitely long” fake path.
- ▶ For simplicity, create loop (edge to and from the same node) at every node, length 0

Theorem (Optimal Substructure)

$$\ell(OPT(u, k)) = \begin{cases} 0 & \text{if } u = v, k = 0 \\ \infty & \text{if } u \neq v, k = 0 \\ & \text{otherwise} \end{cases}$$

Dynamic Programming Approach

Subproblems:

- ▶ $OPT(u, i)$: shortest path from v to u that uses at most i hops (edges)
- ▶ If no such path, set to “infinitely long” fake path.
- ▶ For simplicity, create loop (edge to and from the same node) at every node, length 0

Theorem (Optimal Substructure)

$$\ell(OPT(u, k)) = \begin{cases} 0 & \text{if } u = v, k = 0 \\ \infty & \text{if } u \neq v, k = 0 \\ \min_{w:(w,u) \in E} (\ell(OPT(w, k-1)) + \ell(w, u)) & \text{otherwise} \end{cases}$$

Proof of Optimal Substructure

Induction on k .

$k = 0$: \checkmark . So let $k \geq 1$.

Proof of Optimal Substructure

Induction on k .

$k = 0$: \checkmark . So let $k \geq 1$.

\leq : Let $x = \arg \min_{w:(w,u) \in E} (\ell(\mathit{OPT}(w, k-1)) + \ell(w, u))$

$\implies \mathit{OPT}(x, k-1) \circ (x, u)$ is a $v \rightarrow u$ path with at most k edges, length $\ell(\mathit{OPT}(x, k-1)) + \ell(x, u)$

$\implies \ell(\mathit{OPT}(u, k)) \leq \min_{w:(w,u) \in E} (\ell(\mathit{OPT}(w, k-1)) + \ell(w, u))$

Proof of Optimal Substructure

Induction on k .

$k = 0$: \checkmark . So let $k \geq 1$.

\leq : Let $x = \arg \min_{w:(w,u) \in E} (\ell(\mathit{OPT}(w, k-1)) + \ell(w, u))$

$\implies \mathit{OPT}(x, k-1) \circ (x, u)$ is a $v \rightarrow u$ path with at most k edges, length $\ell(\mathit{OPT}(x, k-1)) + \ell(x, u)$

$\implies \ell(\mathit{OPT}(u, k)) \leq \min_{w:(w,u) \in E} (\ell(\mathit{OPT}(w, k-1)) + \ell(w, u))$

\geq : Let z be node before u in $\mathit{OPT}(u, k)$, and let P' be the first $k-1$ edges of $\mathit{OPT}(u, k)$.
Then

$$\begin{aligned} \ell(\mathit{OPT}(u, k)) &= \ell(P') + \ell(z, u) \geq \ell(\mathit{OPT}(z, k-1)) + \ell(z, u) \\ &\geq \min_{w:(w,u) \in E} (\ell(\mathit{OPT}(w, k-1)) + \ell(w, u)) \end{aligned}$$

Bellman-Ford Algorithm

Obvious dynamic program!

$$M[u, 0] = \infty \text{ for all } u \in V, u \neq v$$

$$M[v, 0] = 0$$

for($k = 1$ to $n - 1$) {

 for($u \in V$) {

$$M[u, k] = \min_{w:(w,u) \in E} (M[w, k - 1] + \ell(w, u))$$

 }

}

Bellman-Ford Algorithm

Obvious dynamic program!

$$M[u, 0] = \infty \text{ for all } u \in V, u \neq v$$

$$M[v, 0] = 0$$

for($k = 1$ to $n - 1$) {

 for($u \in V$) {

$$M[u, k] = \min_{w:(w,u) \in E} (M[w, k - 1] + \ell(w, u))$$

 }

}

Running Time:

Bellman-Ford Algorithm

Obvious dynamic program!

$$M[u, 0] = \infty \text{ for all } u \in V, u \neq v$$

$$M[v, 0] = 0$$

for($k = 1$ to $n - 1$) {

 for($u \in V$) {

$$M[u, k] = \min_{w:(w,u) \in E} (M[w, k - 1] + \ell(w, u))$$

 }

}

Running Time:

- ▶ Obvious: $O(n^3)$

Bellman-Ford Algorithm

Obvious dynamic program!

$$M[u, 0] = \infty \text{ for all } u \in V, u \neq v$$

$$M[v, 0] = 0$$

for($k = 1$ to $n - 1$) {

 for($u \in V$) {

$$M[u, k] = \min_{w:(w,u) \in E} (M[w, k - 1] + \ell(w, u))$$

 }

}

Running Time:

- ▶ Obvious: $O(n^3)$
- ▶ Smarter: $O(mn)$

Bellman-Ford: Correctness

Theorem

After algorithm completes, $M[u, k] = \ell(OPT(u, k))$ for all $k \leq n - 1$ and $u \in V$.

Bellman-Ford: Correctness

Theorem

After algorithm completes, $M[u, k] = \ell(\text{OPT}(u, k))$ for all $k \leq n - 1$ and $u \in V$.

Proof.

Induction on k . Obviously true for $k = 0$.

Bellman-Ford: Correctness

Theorem

After algorithm completes, $M[u, k] = \ell(\mathit{OPT}(u, k))$ for all $k \leq n - 1$ and $u \in V$.

Proof.

Induction on k . Obviously true for $k = 0$.

$$\begin{aligned} M[u, k] &= \min_{w:(w,u) \in E} (M[w, k-1] + \ell(w, u)) && \text{(algorithm)} \\ &= \min_{w:(w,u) \in E} (\ell(\mathit{OPT}(w, k-1)) + \ell(w, u)) && \text{(induction)} \\ &= \ell(\mathit{OPT}(u, k)) && \text{(optimal substructure)} \end{aligned}$$



Negative Weights and Cycle

Suppose weights are negative. Does the problem make sense?

Negative Weights and Cycle

Suppose weights are negative. Does the problem make sense?

- ▶ Negative-weight cycle: not really!

Negative Weights and Cycle

Suppose weights are negative. Does the problem make sense?

- ▶ Negative-weight cycle: not really! Go around cycle forever, make distances arbitrarily negative

Negative Weights and Cycle

Suppose weights are negative. Does the problem make sense?

- ▶ Negative-weight cycle: not really! Go around cycle forever, make distances arbitrarily negative
- ▶ No negative-weight cycle: everything we did before is fine!

Negative Weights and Cycle

Suppose weights are negative. Does the problem make sense?

- ▶ Negative-weight cycle: not really! Go around cycle forever, make distances arbitrarily negative
- ▶ No negative-weight cycle: everything we did before is fine!

Detecting negative-weight cycle:

Negative Weights and Cycle

Suppose weights are negative. Does the problem make sense?

- ▶ Negative-weight cycle: not really! Go around cycle forever, make distances arbitrarily negative
- ▶ No negative-weight cycle: everything we did before is fine!

Detecting negative-weight cycle: One more round of Bellman-Ford!

Negative Weights and Cycle

Suppose weights are negative. Does the problem make sense?

- ▶ Negative-weight cycle: not really! Go around cycle forever, make distances arbitrarily negative
- ▶ No negative-weight cycle: everything we did before is fine!

Detecting negative-weight cycle: One more round of Bellman-Ford!

Fun fact: best-known algorithm with negative (real) edge weights until this year!

Jeremy Fineman. *Single-Source Shortest Paths with Negative Real Weights in $\tilde{O}(mn^{8/9})$ Time.*
STOC '24

Relaxations

Common primitive in shortest path algorithms

- ▶ Reinterpret Bellman-Ford via relaxations
- ▶ Use relaxations for Dijkstra's algorithm

Relaxations

Common primitive in shortest path algorithms

- ▶ Reinterpret Bellman-Ford via relaxations
- ▶ Use relaxations for Dijkstra's algorithm

$\hat{d}(u)$: upper bound on $d(u)$

- ▶ Initially: $\hat{d}(v) = 0$, $\hat{d}(u) = \infty$ for all $u \neq v$

Relaxations

Common primitive in shortest path algorithms

- ▶ Reinterpret Bellman-Ford via relaxations
- ▶ Use relaxations for Dijkstra's algorithm

$\hat{d}(u)$: upper bound on $d(u)$

- ▶ Initially: $\hat{d}(v) = 0$, $\hat{d}(u) = \infty$ for all $u \neq v$

Intuition for $\text{relax}(x, y)$: can we improve $\hat{d}(y)$ by going through x ?

Relaxations

Common primitive in shortest path algorithms

- ▶ Reinterpret Bellman-Ford via relaxations
- ▶ Use relaxations for Dijkstra's algorithm

$\hat{d}(u)$: upper bound on $d(u)$

- ▶ Initially: $\hat{d}(v) = 0$, $\hat{d}(u) = \infty$ for all $u \neq v$

Intuition for $\text{relax}(x, y)$: can we improve $\hat{d}(y)$ by going through x ?

```
relax(x, y) {  
    if( $\hat{d}(y) > \hat{d}(x) + \ell(x, y)$ ) {  
         $\hat{d}(y) = \hat{d}(x) + \ell(x, y)$   
        y.parent = x  
    }  
}
```

Bellman-Ford as Relaxations

```
for( $i = 1$  to  $n$ ) {  
  foreach( $u \in V$ ) {  
    foreach(edge ( $x, u$ )) {  
      relax( $x, u$ )  
    }  
  }  
}
```

Bellman-Ford as Relaxations

```
for( $i = 1$  to  $n$ ) {  
  foreach( $u \in V$ ) {  
    foreach(edge ( $x, u$ )) {  
      relax( $x, u$ )  
    }  
  }  
}
```

Not precisely the same: freezing/parallelism

Dijkstra's Algorithm

High Level

Intuition: “greedy starting at v ”

- ▶ BFS but with edge lengths: use priority queue (heap) instead of queue!

Pros: faster than Bellman-Ford (super fast with appropriate data structures)

Cons: Doesn't work with negative edge weights.

Dijkstra's Algorithm


$$\mathcal{T} = \emptyset$$

$$\hat{d}(v) = 0$$


$$\hat{d}(u) = \infty \text{ for all } u \neq v$$

```
while(not all nodes in  $\mathcal{T}$ ) {  
  let  $u$  be node not in  $\mathcal{T}$  with minimum  $\hat{d}(u)$   
  Add  $u$  to  $\mathcal{T}$   
  foreach edge  $(u, x)$  with  $x \notin \mathcal{T}$  {  
    relax(u,x)  
  }  
}
```


Dijkstra Example






(a)



(b)



(c)



Dijkstra Correctness

Theorem

Throughout the algorithm:

1. \mathcal{T} is a shortest-path tree from \mathbf{v} to the nodes in \mathcal{T} , and
2. $\hat{\mathbf{d}}(\mathbf{u}) = \mathbf{d}(\mathbf{u})$ for every $\mathbf{u} \in \mathcal{T}$.

Dijkstra Correctness

Theorem

Throughout the algorithm:

1. \mathcal{T} is a shortest-path tree from \mathbf{v} to the nodes in \mathcal{T} , and
2. $\hat{\mathbf{d}}(\mathbf{u}) = \mathbf{d}(\mathbf{u})$ for every $\mathbf{u} \in \mathcal{T}$.

Proof. Induction on $|\mathcal{T}|$ (iterations of algorithm)

Dijkstra Correctness

Theorem

Throughout the algorithm:

1. \mathcal{T} is a shortest-path tree from \mathbf{v} to the nodes in \mathcal{T} , and
2. $\hat{\mathbf{d}}(\mathbf{u}) = \mathbf{d}(\mathbf{u})$ for every $\mathbf{u} \in \mathcal{T}$.

Proof. Induction on $|\mathcal{T}|$ (iterations of algorithm)

Base Case: After first iteration (when $|\mathcal{T}| = 1$), added \mathbf{v} to \mathcal{T} with $\hat{\mathbf{d}}(\mathbf{v}) = \mathbf{d}(\mathbf{v}) = 0 \checkmark$

Correctness: Inductive Step (Sketch)


Consider iteration when u added to T , let $w = u.\text{parent}$

$$\implies \hat{d}(u) = \hat{d}(w) + \ell(w, u) = d(w) + \ell(w, u) \text{ (induction)}$$

Correctness: Inductive Step (Sketch)

Consider iteration when u added to T , let $w = u.\text{parent}$

$$\implies \hat{d}(u) = \hat{d}(w) + \ell(w, u) = d(w) + \ell(w, u) \text{ (induction)}$$




- ▶ Red path P actual shortest path, black path found by Dijkstra
- ▶ w' predecessor of u on P . Can't be in T .
 - ▶ If it was, would have $\hat{d}(w') = d(w')$ by induction, would have relaxed (w', u) , so would have $w' = u.\text{parent}$
- ▶ x first node of P outside T , previous node y

Correctness: Inductive Step (Sketch)

Consider iteration when u added to T , let $w = u.\text{parent}$

$$\implies \hat{d}(u) = \hat{d}(w) + \ell(w, u) = d(w) + \ell(w, u) \text{ (induction)}$$




- ▶ Red path P actual shortest path, black path found by Dijkstra
- ▶ w' predecessor of u on P . Can't be in T .
 - ▶ If it was, would have $\hat{d}(w') = d(w')$ by induction, would have relaxed (w', u) , so would have $w' = u.\text{parent}$
- ▶ x first node of P outside T , previous node y

$$\hat{d}(x) \leq \hat{d}(y) + \ell(y, x) = d(y) + \ell(y, x) < \ell(P) = d(u) \leq \hat{d}(u)$$

Correctness: Inductive Step (Sketch)

Consider iteration when u added to T , let $w = u.\text{parent}$

$$\implies \hat{d}(u) = \hat{d}(w) + \ell(w, u) = d(w) + \ell(w, u) \text{ (induction)}$$



- ▶ Red path P actual shortest path, black path found by Dijkstra
- ▶ w' predecessor of u on P . Can't be in T .
 - ▶ If it was, would have $\hat{d}(w') = d(w')$ by induction, would have relaxed (w', u) , so would have $w' = u.\text{parent}$
- ▶ x first node of P outside T , previous node y

$$\hat{d}(x) \leq \hat{d}(y) + \ell(y, x) = d(y) + \ell(y, x) < \ell(P) = d(u) \leq \hat{d}(u)$$

Contradiction! Algorithm would have chosen x next, not u .

Running Time

Algorithm needs to:

- ▶ Select node with minimum \hat{d} value n times
- ▶ Decrease a \hat{d} value at most once per relaxation $\implies \leq m$ times.

Running Time

Algorithm needs to:

- ▶ Select node with minimum \hat{d} value n times
- ▶ Decrease a \hat{d} value at most once per relaxation $\implies \leq m$ times.

Nothing fancy, keep $\hat{d}(u)$ in adjacency list: selecting min \hat{d} value takes $O(n)$ time
 $\implies O(n^2 + m) = O(n^2)$ total.

Running Time

Algorithm needs to:

- ▶ Select node with minimum \hat{d} value n times
- ▶ Decrease a \hat{d} value at most once per relaxation $\implies \leq m$ times.

Nothing fancy, keep $\hat{d}(u)$ in adjacency list: selecting min \hat{d} value takes $O(n)$ time
 $\implies O(n^2 + m) = O(n^2)$ total.

Keep \hat{d} values in a heap!

- ▶ Insert n times
- ▶ Extract-Min n times
- ▶ Decrease-Key m times

Running Time

Algorithm needs to:

- ▶ Select node with minimum \hat{d} value n times
- ▶ Decrease a \hat{d} value at most once per relaxation $\implies \leq m$ times.

Nothing fancy, keep $\hat{d}(u)$ in adjacency list: selecting min \hat{d} value takes $O(n)$ time
 $\implies O(n^2 + m) = O(n^2)$ total.

Keep \hat{d} values in a heap!

- ▶ Insert n times
- ▶ Extract-Min n times
- ▶ Decrease-Key m times

Binary heap: $O(\log n)$ per operation (amortized)
 $\implies O((m + n) \log n)$ running time.

Running Time

Algorithm needs to:

- ▶ Select node with minimum \hat{d} value n times
- ▶ Decrease a \hat{d} value at most once per relaxation $\implies \leq m$ times.

Nothing fancy, keep $\hat{d}(u)$ in adjacency list: selecting min \hat{d} value takes $O(n)$ time
 $\implies O(n^2 + m) = O(n^2)$ total.

Keep \hat{d} values in a heap!

- ▶ Insert n times
- ▶ Extract-Min n times
- ▶ Decrease-Key m times

Binary heap: $O(\log n)$ per operation (amortized)
 $\implies O((m + n) \log n)$ running time.

Fibonacci Heap:

- ▶ Insert, Decrease-Key $O(1)$ amortized
 - ▶ Extract-Min $O(\log n)$ amortized
- $\implies O(m + n \log n)$ running time