

CS 341: ALGORITHMS

Lecture 14: graph algorithms V – single source shortest path
Readings: see website

Trevor Brown
<https://student.cs.uwaterloo.ca/~cs341>
trevor.brown@uwaterloo.ca

1

DIJKSTRA'S ALGORITHM

Single-source shortest path
in a graph with non-negative edge weights

2

PROBLEM: SINGLE SOURCE SHORTEST PATHS (SSSP)

Let's study directed G .
Can also be defined for undirected G .

- Input: graph $G = (V, E)$ and a **non-negative** weight function $w(e)$ defined for every edge e
- Problem: for every node $v \neq s$, output a path $s \rightsquigarrow v$ with the **smallest total weight** (among all paths $s \rightsquigarrow v$)
- I.e., each path P should minimize $w(P) = \sum_{e \in P} w(e)$

"Shortest" means minimum weight

Suppose this is s

Shortest path to d

Shortest path to i

Shortest path to c

Shortest path to e

And so on... one path for each node.

3

APPLICATION: DRIVING DISTANCE TO MANY POSSIBLE DESTINATIONS

- Single source: from where you are
- Shortest path: to **all** destinations
 - Display a subset of destinations
 - Include the optimal distances computed using SSSP algorithm
 - Other heuristics... traffic? Lights?
 - Weights** can combine many factors

4

[video clip]

Game AI: path finding with waypoints

Divide game world into **linear paths**, then send game characters in **straight lines** between waypoints

If some linear paths are much faster/slower, use **weighted SSSP**

Otherwise use BFS to find shortest sequence of waypoints (with **fewest** waypoints)

5

DIJKSTRA'S ALGORITHM

ILLUSTRATIVE EXAMPLE

Start node s is here

Showing $dist$ -values

$d_c, d_b, d_e, d_f, d_g, d_h, d_i, d_j, d_k, d_l$ is optimal

Done!

This $dist$ is optimal!

Can we use this optimal $dist$ to improve the $dist$ of neighbours

We call this **relaxing** the neighbours

key insight: after relaxing all, the smallest $dist$ (that we didn't already know was optimal) is now optimal

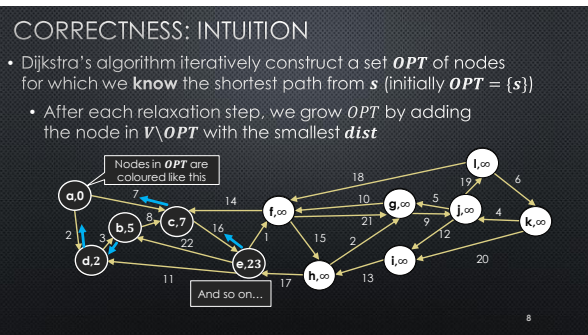
6

```

1 Dijkstra(adj[1..n], s)
2 pred[1..n] = [null, null, ..., null]
3 dist[1..n] = [infy, infy, ..., infy]
4 pq = new priority queue
5
6 dist[s] = 0
7 for u = 1..n
8   pq.enqueue(u, dist[u])
9
10 while pq is not empty
11   u = pq.dequeueMin()
12   for v in adj[u]
13     if dist[u] + w(u,v) < dist[v]
14       dist[v] = dist[u] + w(u,v)
15       pred[v] = u
16       pq.changePriority(v, dist[v])
17
18 return pred, dist
    
```

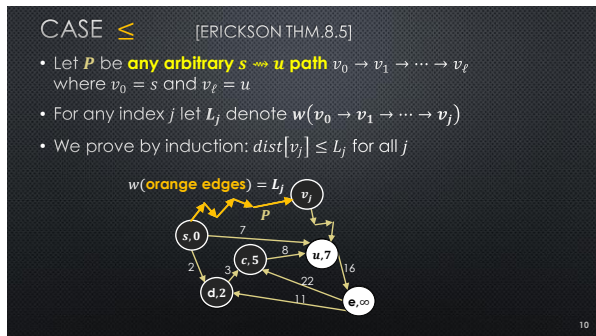
Annotations:

- Line 4: Maintain nodes in priority order, ordered by smallest distance
- Line 7-9: Enqueue all nodes with distance ∞ except for s with distance 0
- Line 10-11: Eventually dequeue all nodes (no more enqueues)
- Line 12: Each dequeued node u has optimal $dist$
- Line 13-16: Relax neighbour v



PROOF

- Theorem:** At the end of the algorithm, for all u , $dist[u]$ is exactly the total weight of the shortest $s \rightsquigarrow u$ path
- We prove this in two parts
 - $dist[u] \leq$ the total weight of the shortest $s \rightsquigarrow u$ path (case \leq)
 - $dist[u] \geq$ the total weight of the shortest $s \rightsquigarrow u$ path (case \geq)



- Prove by induction: $\forall j : dist[v_j] \leq L_j$
- Base case: $dist[v_0] = dist[s] = 0 = L_0$
- Ind. step: **suppose** $\forall_{j>0} : dist[v_{j-1}] \leq L_{j-1}$
 - When dequeueMin() returns v_{j-1} : we check if $dist[v_{j-1}] + w(v_{j-1}, v_j) < dist[v_j]$
 - If so, we set $dist[v_j] = dist[v_{j-1}] + w(v_{j-1}, v_j)$
 - If not, $dist[v_j] \leq dist[v_{j-1}] + w(v_{j-1}, v_j)$
 - In both cases, $dist[v_j] \leq dist[v_{j-1}] + w(v_{j-1}, v_j)$
 - By I.H. $dist[v_{j-1}] \leq L_{j-1}$ so $dist[v_j] \leq L_{j-1} + w(v_{j-1}, v_j)$
 - And $L_{j-1} + w(v_{j-1}, v_j) = L_j$ by definition
- So $dist[v_j] \leq L_j$**

So $dist[u] \leq$ the weight of EVERY $s \rightsquigarrow u$ path. Including the shortest $s \rightsquigarrow u$ path.

CASE \geq

- Let P' be the path $s \rightarrow \dots \rightarrow pred[pred[u]] \rightarrow pred[u] \rightarrow u$
 - i.e., the reverse of following pred pointers from u back to s
- We show $dist[u]$ is as long as P' path (and hence as long as the shortest path)
- Denote the nodes in P' by v_0, v_1, \dots, v_ℓ where $v_0 = s$ and $v_\ell = u$
- Let $L_j = w(v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_j)$
- Prove by induction:** $\forall_{j>0} : dist[v_j] = L_j$
- Base case: $dist[v_0] = dist[s] = 0 = L_0$

CASE \geq

- $P^i = v_0 \rightarrow \dots \rightarrow v_i \rightarrow s \rightarrow \dots \rightarrow pred[pred[u]] \rightarrow pred[u] \rightarrow u$
- $L_j = w(v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_j)$
- Inductive step:** suppose $\forall_{j>0} : dist[v_{j-1}] = L_{j-1}$
- When we set $pred[v_j] = v_{j-1}$, we set $dist[v_j] = dist[v_{j-1}] + w(v_{j-1}, v_j)$

Recall: if $dist[u] + w(u, v) < dist[v]$
 $dist[v] = dist[u] + w(u, v)$
 $pred[v] = u$

- By I.H., $dist[v_j] = L_{j-1} + w(v_{j-1}, v_j)$
- By definition $L_j = L_{j-1} + w(v_{j-1}, v_j)$
- So $dist[v_j] = L_j$

So $dist[u]$ = length of a particular path P^i in the graph
 And length of P^i is \geq length of shortest path
 So $dist[u] \geq$ length of shortest path $s \rightarrow u$
 So $dist[u]$ is both \leq and \geq to the length of the shortest $s \rightarrow u$ path!
 That means it's **equal** to the length of the shortest path!

RUNTIME

```

1 Dijkstra(adj[1..n], s)
2 pred[1..n] = [null, null, ..., null]
3 dist[1..n] = [infy, infy, ..., infy]
4 pq = new priority queue
5
6 dist[s] = 0
7 for u = 1..n
8   pq.enqueue(u, dist[u])
9
10 while pq is not empty
11   u = pq.dequeueMin()
12   for v in adj[u]
13     if dist[u] + w(u,v) < dist[v]
14       dist[v] = dist[u] + w(u,v)
15       pred[v] = u
16       pq.changePriority(v, dist[v])
17
18 return pred, dist
    
```

- Each node enqueued and dequeueMin'd once
- $O(n \log n)$
- For each dequeueMin, do $O(\log n)$ per neighbour
- $O(\log n)$ for **each edge**
- $O(m \log n)$ w/adjacency lists
- Total time $O((n + m) \log n)$**

Space complexity?

OUTPUTTING ACTUAL SHORTEST PATH(S)?

- To compute the actual shortest **path** $s \rightsquigarrow t$
- Inspect $pred[t]$
- If it is NULL, there is no such path
- Otherwise, follow $pred$ pointers back to s , and return the **reverse** of that path

AN ALTERNATIVE IMPLEMENTATION

```

1 Dijkstra(adj[1..n], s)
2 pred[1..n] = [null, null, ..., null]
3 dist[1..n] = [infy, infy, ..., infy]
4 OPT = [false, false, ..., false]
5
6 dist[s] = 0
7 OPT[s] = true
8 numOpt = 1
9
10 while numOpt < n
11   choose u such that OPT[u] == false
12   and dist[u] is minimized
13   OPT[u] = true
14   numOpt = numOpt + 1
15   for v = adj[u]
16     if dist[u] + w(u,v) < dist[v]
17       dist[v] = dist[u] + w(u,v)
18       pred[v] = u
19
20 return pred, dist
    
```

- Instead of using a **priority queue**
- Find the minimum dist[] node to add to OPT via **linear search**
- Runtime?**
 - $O(n^2)$
- Better or worse than $O((n + m) \log n)$?**

WEBSITE DEMONSTRATING DIJKSTRA'S ALG

- <https://www.cs.usfca.edu/~galles/visualization/Dijkstra.html>

BELLMAN-FORD

Single-source shortest path in a graph with possibly **negative** edge weights but **no negative cycles**

Shortest Paths and Negative Weight Cycles

Subsequent algorithms we will be studying will solve shortest path problems as long as there are no cycles having negative weight.

If there is a negative weight cycle, then there is no shortest path (why?).

There is still a shortest simple path, but there are apparently no known efficient algorithms to find the shortest simple paths in graphs containing negative weight cycles.

If there are no negative weight cycles, we can assume WLOG that shortest paths are simple paths (any path can be replaced by a simple path having the same weight).

Negative weight edges in an undirected graph are not allowed, as they would give rise to a negative weight cycle (consisting of two edges) in the associated directed graph.

BELLMAN-FORD

The *Bellman-Ford algorithm* solves the single source shortest path problem in any directed graph without negative weight cycles.


The algorithm is very simple to describe:

Repeat $n - 1$ times: *relax every edge* in the graph (where *relax* is the updating step in Dijkstra's algorithm).

$O(n)$	1	BellmanFord($n, E[1..n], s$)	
	2	$pred[1..n]$ = new array filled with null	
	3	$D[1..n]$ = new array filled with infinity	
	4	$D[s] = 0$	
$O(n)$	5	for $i = 1..n$	$O(1)$
$O(n)$ outer iterations	6	for (u,v,w) in E	
$O(m)$ inner iterations per outer iteration	7	if $D[u] + w < D[v]$	
	8	$D[v] = D[u] + w$	
	9	$pred[v] = u$	
Could be $O(n^2)$	10	return $(D, pred)$	
Total $O(nm)$			

BEST CASE EXECUTION

If technically suffices to do one iteration of the outer loop




Edges happen to be processed left to right by the inner loop

```

1 BellmanFord( $n, E[1..n], s$ )
2    $pred[1..n]$  = new array filled with null
3    $D[1..n]$  = new array filled with infinity
4    $D[s] = 0$ 
5   for  $i = 1..n$ 
6     for  $(u,v,w)$  in  $E$ 
7       if  $D[u] + w < D[v]$ 
8          $D[v] = D[u] + w$ 
9          $pred[v] = u$ 
10  return  $(D, pred)$ 
    
```

WORST CASE EXECUTION

Need n iterations of outer loop



Edges happen to be processed right to left by the inner loop

```

1 BellmanFord( $n, E[1..n], s$ )
2    $pred[1..n]$  = new array filled with null
3    $D[1..n]$  = new array filled with infinity
4    $D[s] = 0$ 
5   for  $i = 1..n$ 
6     for  $(u,v,w)$  in  $E$ 
7       if  $D[u] + w < D[v]$ 
8          $D[v] = D[u] + w$ 
9          $pred[v] = u$ 
10  return  $(D, pred)$ 
    
```

Since the longest possible path without a cycle can be $n - 1$ edges, the edges must be scanned $n - 1$ times to ensure the shortest path has been found for all nodes.

Dijkstra's is similar, but consistently achieves good ordering using its priority queue

WHY BELLMAN-FORD WORKS

- Not going to prove this (by induction), but the crucial lemma is:
 - After i iterations of the outer *for*-loop,
 - if $D[u] \neq \infty$, it is equal to the weight of some path $s \rightsquigarrow u$; and
 - if there is a path $P = (s \rightsquigarrow u)$ with at most i edges, then $D[u] \leq w(P)$
- So, after $n - 1$ iterations, if \exists path P with at most $n - 1$ edges, then $D[u] \leq w(P)$. (Note: any more edges would create a cycle.)
- So, if u is reachable from s , then $D[u]$ is the length of the shortest simple path (no cycles) from s to u

Of course every simple path has at most $n - 1$ edges

So what if we do another iteration, and some $D[u]$ improves?

There is a negative cycle!

A MORE DETAILED IMPLEMENTATION

```

1 BellmanFordCheck( $n, E[1..n], s$ )
2    $pred[1..n]$  = new array filled with null
3    $D[1..n]$  = new array filled with infinity
4    $D[s] = 0$ 
5   for  $i = 1..n$ 
6      $changed = false$ 
7     for  $(u,v,w)$  in  $E$ 
8       if  $D[u] + w < D[v]$ 
9          $D[v] = D[u] + w$ 
10         $pred[v] = u$ 
11         $changed = true$ 
12    if not  $changed$ 
13      exit loop
14    if  $i == n$  // assert:  $changed == true$ 
15      return NEGATIVE_CYCLE
16  return  $(D, pred)$ 
    
```

- With early stopping
- and checking for negative cycles

BONUS SLIDE

- Why can't you just modify a graph with negative weights by: finding the minimum edge weight W_{min} , and adding that to each edge, so you no longer have negative edges and can run Dijkstra's algorithm?
- **Exercise:** can you find a graph for which this will cause Dijkstra's algorithm to return the **wrong answer**?
- **Solution:**
 - Consider a graph with 5 nodes: s, a, b, c, t
 - And edges s->a with weight -10, b->t with weight 10
s->b weight -1, b->c weight -1, c->t weight -1
 - What happens if you modify this graph as proposed, then run Dijkstra's to find the shortest path from s to t?

25