

# CSE 373: Data Structures and Algorithms

## Lecture 16: Dijkstra's Algorithm (Graphs)

Instructor: Lilian de Greef  
Quarter: Summer 2017

# Today

- Announcements
- Graph Traversals Continued
  - Remarks on DFS & BFS
  - Shortest paths for weighted graphs:  
**Dijkstra's Algorithm!**

# Announcements:

Homework 4 is out!


- Due next Friday (August 4<sup>th</sup>) at 5:00pm
- May choose to pair-program if you like!
  - Same cautions as last time apply: choose partners and when to start working wisely!
- Can almost entirely complete using material by end of this lecture
- Will discuss some software-design concepts next week to help you prevent some (potentially non-obvious) bugs

# Another midterm correction... ( 🤪 & 🤨 )

## 1. True or False: (6 points)

Circle whether the statement is either true or false.

- f. (true / **false**): In an AVL tree, the longest and shortest paths (i.e. number of edges) from the root to a leaf do not differ by more than one.



Bring your midterm to \*any\* office hours to get your point back.

I will have the final exam *quadruple-checked* to avoid these situations!  
(I am so sorry)

# Graphs: Traversals Continued

And introducing Dijkstra's Algorithm for shortest paths!

# Graph Traversals: Recap & Running Time


- Traversals: General Idea

- Starting from one vertex, repeatedly explore adjacent vertices
- Mark each vertex we visit, so we don't process each more than once (cycles!)

Marking a vertex adds it to the "set of visited vertices"

- Important Graph Traversal Algorithms:

	Depth First Search (DFS)	Breadth First Search (BFS)
Explore...	as far as possible before backtracking	all neighbors first before next level of neighbors
Choose next vertex using...	recursion or a <u>stack</u>	a <u>queue</u>



push/pop/enqueue/dequeue

- Assuming "choose next vertex" is  $O(1)$ , entire traversal is
  - Use graph represented with adjacency list

~~$O(V+E)$~~   
 $O(V)$

# Comparison (useful for Design Decisions!)

- Which one finds **shortest** paths? *BFS!*
  - i.e. which is better for “what is the shortest path from **x** to **y**” when there’s more than one possible path?
- Which one can use less space in finding a path? *(memory) DFS*
- A third approach:
  - *Iterative deepening (IDFS)*:
    - Try DFS but disallow recursion more than  $K$  levels deep
    - If that fails, increment  $K$  and start the entire search over
  - Like BFS, finds shortest paths. Like DFS, less space.

# Graph Traversal Uses

In addition to finding paths, we can use graph traversals to answer:

- What are all the vertices *reachable* from a starting vertex?
- Is an undirected graph connected?
- Is a directed graph strongly connected?

*continue until mark  
all vertices*

- But what if we want to actually output the path?
- How to do it:
  - Instead of just “marking” a node, store the previous node along the path
  - When you reach the goal, follow `path` fields back to where you started (and then reverse the answer)
  - If just wanted path *length*, could put the integer distance at each node instead once



# Single source shortest paths

- Done: BFS to find the minimum path length from  $\mathbf{v}$  to  $\mathbf{u}$  in  $O(|E|+|V|)$
- Actually, can find the minimum path length from  $\mathbf{v}$  to *every node*
  - Still  $O(|E|+|V|)$
  - No faster way for a “distinguished” destination in the worst-case
- Now: Weighted graphs

Given a weighted graph and node  $\mathbf{v}$ ,  
find the minimum-cost path from  $\mathbf{v}$  to every node

- As before, asymptotically no harder than for one destination


# A Few Applications of Shortest Weighted Path

- Driving directions — traffic (distance)
- Cheap flight itineraries — price
- Network routing — traffic (bandwidth)
- Critical paths in project management — importance

# Not as easy as BFS

Why BFS won't work: Shortest path may not have the fewest edges

- Annoying when this happens with costs of flights




We will assume there are no negative weights

- *Problem* is *ill-defined* if there are negative-cost *cycles*
- Today's *algorithm* is *wrong* if *edges* can be negative
  - There are other, slower (but not terrible) algorithms

# Algorithm: General Idea


**Goal:** From one starting vertex, what are the shortest paths to each of the other vertices (for a weighted graph)?



**Idea:** Similar to BFS

- Repeatedly increase a “set of vertices with known shortest distances”
- Any vertex not in this set will have a “best distance so far”
- Each vertex has a “cost” to represent these shortest/best distances
- Update costs (i.e. “best distances so far”) as we add vertices to set

# Shortest Path Example #1




Known Set (in order added):

A, C, B, D, F, H, G, E

next vertex is the one with the lowest cost so far

vertex	known?	cost	path
A	Y	0	
B	Y	<del>∞</del> 2	A
C	Y	<del>∞</del> 1	A
D	Y	<del>∞</del> 4	A
E	Y	<del>∞</del> <del>11</del> 11	<del>G</del>
F	Y	<del>∞</del> 4	B
G	Y	<del>∞</del> 8	H
H	Y	<del>∞</del> 7	E


# Example #1



Order Added to Known Set:

vertex	known?	cost	path
A		0	
B		??	
C		??	
D		??	
E		??	
F		??	
G		??	
H		??	

# Example #1




vertex	known?	cost	path
A	Y	0	
B		$\leq 2$	A
C		$\leq 1$	A
D		$\leq 4$	A
E		??	
F		??	
G		??	
H		??	

Order Added to Known Set:

A

# Example #1




Order Added to Known Set:

A, C

vertex	known?	cost	path
A	Y	0	
B		$\leq 2$	A
C	Y	1	A
D		$\leq 4$	A
E		$\leq 12$	C
F		??	
G		??	
H		??	



# Example #1




Order Added to Known Set:

A, C, B

vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D		$\leq 4$	A
E		$\leq 12$	C
F		$\leq 4$	B
G		??	
H		??	

# Example #1




Order Added to Known Set:

A, C, B, D

vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E		$\leq 12$	C
F		$\leq 4$	B
G		??	
H		??	

# Example #1




Order Added to Known Set:

A, C, B, D, F

vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E		$\leq 12$	C
F	Y	4	B
G		??	
H		$\leq 7$	F

# Example #1




Order Added to Known Set:

A, C, B, D, F, H

vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E		$\leq 12$	C
F	Y	4	B
G		$\leq 8$	H
H	Y	7	F

# Example #1




vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E		$\leq 11$	G
F	Y	4	B
G	Y	8	H
H	Y	7	F

Order Added to Known Set:

A, C, B, D, F, H, G

# Example #1



vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E	Y	11	G
F	Y	4	B
G	Y	8	H
H	Y	7	F


Order Added to Known Set:

A, C, B, D, F, H, G, E

# This is called... Dijkstra's Algorithm

Named after its inventor Edsger Dijkstra (1930-2002)

Truly one of the “founders” of computer science;  
this is just one of his many contributions



*“Computer science is no more about computers  
than astronomy is about telescopes.”*

- Edsger Dijkstra

# Dijkstra's Algorithm (Pseudocode)

**Dijkstra's Algorithm** – the following algorithm for finding single-source shortest paths in a weighted graph (directed or undirected) with no negative-weight edges:

1. For each node  $v$ , set  $v.cost = \infty$  and  $v.known = false$
2. Set  $source.cost = 0$
3. While there are unknown nodes in the graph
  - a) Select the unknown node  $v$  with lowest cost
  - b) Mark  $v$  as known
  - c) For each edge  $(v, u)$  with weight  $w$ ,

```
c1 = v.cost + w // cost of best path through v to u
c2 = u.cost      // cost of best path to u previously known
if(c1 < c2){     // if the path through v is better
    u.cost = c1
    u.path = v   // for computing actual paths
}
```



# Dijkstra's Algorithm: Features

- When a vertex is marked known, the cost of the shortest path to that node is known
  - The path is also known by following back-pointers
- While a vertex is still not known, another shorter path to it *\*might\** still be found

Note: The “Order Added to Known Set” is not important

- A detail about how the algorithm works (client doesn't care)
- Not used by the algorithm (implementation doesn't care)
- It is sorted by path-cost, resolving ties in some way
  - Helps give intuition of why the algorithm works

# Dijkstra's Algorithm: Commentary

Dijkstra's Algorithm is one example of...

for Dijkstra's  
always will

- A **greedy algorithm**:

- Make a locally optimal choice at each stage to (hopefully) find a global optimum
- i.e. Settle on the best looking option at each repeated step
- **Note**: for some problems, greedy algorithms cannot find best answer!


vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E	Y	11	G
F	Y	4	B
G	Y	8	H
H	Y	7	F

- **Dynamic programming**:

- Solve a complex problem by breaking it down into a collection of simpler subproblems, solve each of those subproblems just once, and store their solutions.
- i.e. Save partial solutions, and use it to solve further parts to avoid repeating work




# Dijkstra's Algorithm: Practice Time!



An order of adding vertices to the known set:


- A) ~~A, D, C, E, F, B, G~~
- B) A, D, C, E, B, F, G
- C) ~~A, D, E, C, B, G, F~~
- D) ~~A, D, E, C, B, F, G~~

correct.



vertex	known?	cost	path
A	Y	0	
B	Y	<del>∞</del> ≤ 6 ≤ 3 3	D
C	Y	<del>∞</del> ≤ 2 2	A
D	Y	<del>∞</del> ≤ 1 1	A
E	Y	<del>∞</del> ≤ 2 2	D
F	Y	<del>∞</del> ≤ 7 ≤ 4 4	<del>D</del> C
G	Y	<del>∞</del> ≤ 6 6	D

# Dijkstra's Algorithm: Practice Time!



An order of adding vertices to the known set:

A) A, D, C, E, F, B, G


B) A, D, C, E, B, F, G ← correct

C) A, D, E, C, B, G, F

D) A, D, E, C, B, F, G ← also correct!

vertex	known?	cost	path
A	Y	0	
B		$\leq 6$	D
C		$\leq 2$	A
D	Y	1	A
E		$\leq 2$	D
F		$\leq 7$	<del>D</del> C
G		$\leq 6$	D


## Example #2



Order Added to Known Set:

vertex	known?	cost	path
A		0	
B		??	
C		??	
D		??	
E		??	
F		??	
G		??	

## Example #2




Order Added to Known Set:

A

vertex	known?	cost	path
A	Y	0	
B		??	
C		$\leq 2$	A
D		$\leq 1$	A
E		??	
F		??	
G		??	

## Example #2




Order Added to Known Set:

A, D

vertex	known?	cost	path
A	Y	0	
B		$\leq 6$	D
C		$\leq 2$	A
D	Y	1	A
E		$\leq 2$	D
F		$\leq 7$	D
G		$\leq 6$	D

## Example #2




Order Added to Known Set:

A, D, C

vertex	known?	cost	path
A	Y	0	
B		$\leq 6$	D
C	Y	2	A
D	Y	1	A
E		$\leq 2$	D
F		$\leq 4$	C
G		$\leq 6$	D



## Example #2




Order Added to Known Set:

A, D, C, E

vertex	known?	cost	path
A	Y	0	
B		$\leq 3$	E
C	Y	2	A
D	Y	1	A
E	Y	2	D
F		$\leq 4$	C
G		$\leq 6$	D

## Example #2




Order Added to Known Set:

A, D, C, E, B

vertex	known?	cost	path
A	Y	0	
B	Y	3	E
C	Y	2	A
D	Y	1	A
E	Y	2	D
F		$\leq 4$	C
G		$\leq 6$	D

## Example #2




Order Added to Known Set:

A, D, C, E, B, F

vertex	known?	cost	path
A	Y	0	
B	Y	3	E
C	Y	2	A
D	Y	1	A
E	Y	2	D
F	Y	4	C
G		$\leq 6$	D

## Example #2




Order Added to Known Set:

A, D, C, E, B, F, G



vertex	known?	cost	path
A	Y	0	
B	Y	3	E
C	Y	2	A
D	Y	1	A
E	Y	2	D
F	Y	4	C
G	Y	6	D

## Example #3



- How will the “best-cost-so-far” for Y proceed? 90, 80, 70, ...
- Is this expensive? No  
→ because each *edge* is processed only once

# Where are We?

- Had a problem: Compute shortest paths in a weighted graph with no negative weights
- Learned an algorithm: Dijkstra's algorithm
- What should we do after learning an algorithm?
  - Prove it is correct 
    - Not obvious!
    - We will sketch the key ideas
  - Analyze its efficiency 
    - Will do better by using a data structure we learned earlier!


# Correctness: Intuition

Rough intuition:

All the “known” vertices have the correct shortest path

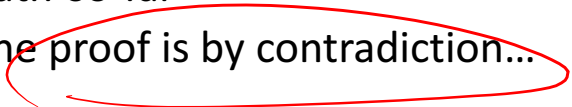
- True initially: shortest path to start node has cost 0
- If it stays true every time we mark a node “known”, then by induction this holds and eventually everything is “known”

*induction!*




Key fact we need: When we mark a vertex “known” we won’t discover a shorter path later!

- This holds only because Dijkstra’s algorithm picks the node with the next shortest path-so-far
- The proof is by contradiction...



# Correctness: The Cloud (Rough Sketch)



- Suppose  $v$  is the next node to be marked known (next to add to “the cloud of known vertices”)
- The **best-known path** to  $v$  must have only nodes “in the cloud”
  - Else we would have picked a node closer to the cloud than  $v$
- Suppose the **actual shortest path** to  $v$  is different
  - It won’t use only cloud nodes, or we would know about it
  - So it must use non-cloud nodes. Let  $w$  be the *first* non-cloud node on this path.
  - The part of the path up to  $w$  is **already known** and must be shorter than the best-known path to  $v$ .
  - So  $v$  would not have been picked. Contradiction!