



Pranay Pathole

@PPathole



Bae: Come over

Dijkstra: But there are so many routes to take and I don't know which one's the fastest

Bae: My parent aren't home

Dijkstra:

## Dijkstra's algorithm


Graph search algorithm

*Not to be confused with Dykstra's projection algorithm.*

**Dijkstra's algorithm** is an algorithm for finding the **shortest paths** between **nodes** in a **graph**, which may represent, for example, road networks. It was conceived by computer scientist **Edsger W. Dijkstra** in 1956 and published three years later.<sup>[1][2]</sup>

The algorithm exists in many variants; Dijkstra's original variant found the shortest path between two nodes,<sup>[3]</sup> but a more common variant fixes a single node as the "source" node and finds shortest paths from the source to all other nodes in the graph, producing a **shortest-path tree**.

Dijkstra's algorithm



@CodeDoesMeme

# CS 201: Data Structures Shortest Paths

Aaron Bauer  
Winter 2021

# Single source shortest paths

- Done: BFS to find the minimum path length from  $v$  to  $u$  in  $O(|E|+|V|)$
- Actually, can find the minimum path length from  $v$  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  $v$ ,  
find the minimum-cost path from  $v$  to every node

- As before, asymptotically no harder than for one destination
- Unlike before, BFS will not work

# *Applications*

- Driving directions
- Cheap flight itineraries
- Network routing
- Critical paths in project management

# Not as easy



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*
- BFS is *wrong* if *edges* can be negative
  - There are other, slower (but not terrible) algorithms

# *Dijkstra's algorithm*

- Algorithm named after its inventor Edsger Dijkstra (1930-2002)
  - A good quotation: “computer science is no more about computers than astronomy is about telescopes”
- The idea: reminiscent of BFS, but adapted to handle weights
  - Grow the set of nodes whose shortest distance has been computed
  - Nodes not in the set will have a “best distance so far”
  - A priority queue will turn out to be useful for efficiency

# Dijkstra's Algorithm: Idea



- Initially, start node has cost 0 and all other nodes have cost  $\infty$
- At each step:
  - Pick closest unknown vertex  $v$
  - Add it to the “cloud” of known vertices
  - Update distances for nodes with edges from  $v$
- That's it!

# The Algorithm

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
}
```

# *Important 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




# Example #1



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

Order Added to Known Set:

# 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




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		??	

Order Added to Known Set:

A, C

# Example #1




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		??	

Order Added to Known Set:

A, C, B

# Example #1




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		??	

Order Added to Known Set:

A, C, B, D

# Example #1




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

Order Added to Known Set:

A, C, B, D, F

# Example #1




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

Order Added to Known Set:

A, C, B, D, F, H

# Example #1




vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E		$\leq 11$	<b>G</b>
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

# 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

# Interpreting the Results

- Now that we're done, how do we get the path from, say, A to E?
  - Follow that path column in reverse E to G to H to F to B to A
  - So the path is A, B, F, H, G, E




Order Added to Known Set:

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

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

# Stopping Short

- How would this have worked differently if we were only interested in:
  - The path from A to F?
  - We can stop as soon as we add F to the known set




Order Added to Known Set:

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

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


# Example #2



Order Added to Known Set:

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

# Example #2




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

Order Added to Known Set:

A

# 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




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

Order Added to Known Set:

A, D, C



# Example #2




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

Order Added to Known Set:

A, D, C, E

# Example #2




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

Order Added to Known Set:

A, D, C, E, B

## Example #2




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

Order Added to Known Set:

A, D, C, E, B, F

## Example #2




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

Order Added to Known Set:

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


# Example #3



How will the best-cost-so-far for Y proceed?

Is this expensive?


## Example #3



How will the best-cost-so-far for Y proceed? 90, 81, 72, 63, 54, ...

Is this expensive?

## Example #3



How will the best-cost-so-far for Y proceed? 90, 81, 72, 63, 54, ...

Is this expensive? No, each *edge* is processed only once

# *A Greedy Algorithm*

- Dijkstra's algorithm
  - For single-source shortest paths in a weighted graph (directed or undirected) with no negative-weight edges
- An example of a *greedy algorithm*:
  - At each step, always does what seems best at that step
    - A locally optimal step, not necessarily globally optimal
  - Once a vertex is known, it is not revisited
    - Turns out to be globally optimal



# *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?
  - Analyze its efficiency
    - Will do better by using a data structure we learned earlier!

# *Efficiency, first approach*

Use pseudocode to determine asymptotic run-time

- Notice each edge is processed only once

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  while(not all nodes are known) {  
    b = find unknown node with smallest cost  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          a.cost = b.cost + weight((b,a))  
          a.path = b  
        }  
  }  
}
```



# Efficiency, first approach

Use pseudocode to determine asymptotic run-time

- Notice each edge is processed only once

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  while(not all nodes are known) {  
    b = find unknown node with smallest cost  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          a.cost = b.cost + weight((b,a))  
          a.path = b  
        }  
  }  
}
```

$O(|V|)$

# Efficiency, first approach

Use pseudocode to determine asymptotic run-time

- Notice each edge is processed only once

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  while(not all nodes are known) {  
    b = find unknown node with smallest cost  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          a.cost = b.cost + weight((b,a))  
          a.path = b  
        }  
  }  
}
```

$O(|V|)$

$O(|V|^2)$

# Efficiency, first approach

Use pseudocode to determine asymptotic run-time

- Notice each edge is processed only once

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  while(not all nodes are known) {  
    b = find unknown node with smallest cost  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          a.cost = b.cost + weight((b,a))  
          a.path = b  
        }  
  }  
}
```

$O(|V|)$

$O(|V|^2)$

$O(|E|)$

# Efficiency, first approach

Use pseudocode to determine asymptotic run-time

- Notice each edge is processed only once

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  while(not all nodes are known) {  
    b = find unknown node with smallest cost  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          a.cost = b.cost + weight((b,a))  
          a.path = b  
        }  
  }  
}
```

$O(|V|)$

$O(|V|^2)$

$O(|E|)$

$O(|V|^2)$

# *Improving asymptotic running time*

- So far:  $O(|V|^2)$
- We had a similar “problem” with topological sort being  $O(|V|^2)$  due to each iteration looking for the node to process next
  - We solved it with a queue of zero-degree nodes
  - But here we need the lowest-cost node and costs can change as we process edges
- Solution?

# *Improving (?) asymptotic running time*

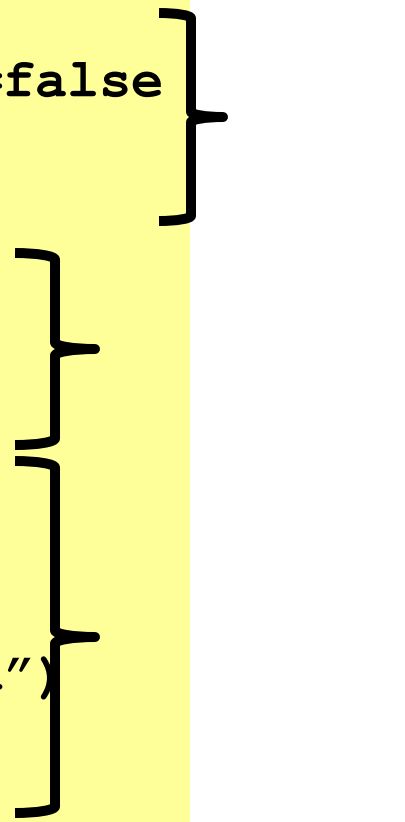
- So far:  $O(|V|^2)$
- We had a similar “problem” with topological sort being  $O(|V|^2)$  due to each iteration looking for the node to process next
  - We solved it with a queue of zero-degree nodes
  - But here we need the lowest-cost node and costs can change as we process edges
- Solution?
  - A priority queue holding all unknown nodes, sorted by cost
  - But must support **decreaseKey** operation
    - Must maintain a reference from each node to its current position in the priority queue
    - Conceptually simple, but can be a pain to code up



# Efficiency, second approach

Use pseudocode to determine asymptotic run-time

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  build-heap with all nodes  
  while(heap is not empty) {  
    b = deleteMin()  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          decreaseKey(a, "new cost - old cost")  
          a.path = b  
        }  
  }  
}
```



# Efficiency, second approach

Use pseudocode to determine asymptotic run-time

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  build-heap with all nodes  
  while(heap is not empty) {  
    b = deleteMin()  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          decreaseKey(a, "new cost - old cost")  
          a.path = b  
        }  
  }  
}
```

$O(|V|)$

# Efficiency, second approach

Use pseudocode to determine asymptotic run-time

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  build-heap with all nodes  
  while(heap is not empty) {  
    b = deleteMin()  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          decreaseKey(a, "new cost - old cost")  
          a.path = b  
        }  
  }  
}
```

$O(|V|)$

$O(|V|\log|V|)$

# Efficiency, second approach

Use pseudocode to determine asymptotic run-time

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  build-heap with all nodes  
  while(heap is not empty) {  
    b = deleteMin()  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          decreaseKey(a, "new cost - old cost")  
          a.path = b  
        }  
  }  
}
```

$O(|V|)$

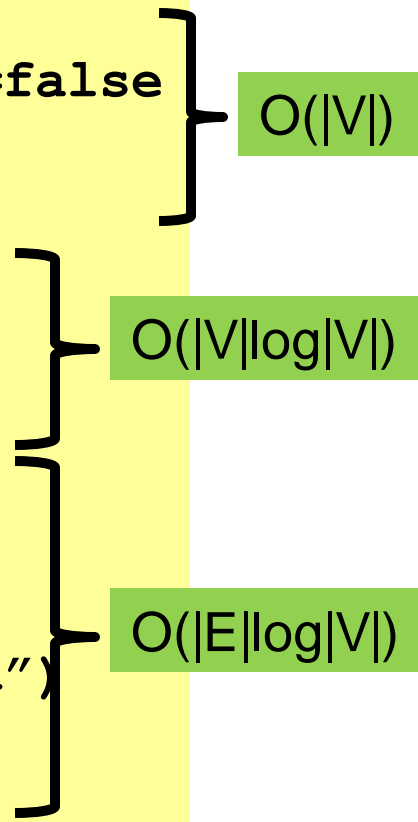
$O(|V|\log|V|)$

$O(|E|\log|V|)$

# Efficiency, second approach

Use pseudocode to determine asymptotic run-time

```
dijkstra(Graph G, Node start) {  
  for each node: x.cost=infinity, x.known=false  
  start.cost = 0  
  build-heap with all nodes  
  while(heap is not empty) {  
    b = deleteMin()  
    b.known = true  
    for each edge (b,a) in G  
      if(!a.known)  
        if(b.cost + weight((b,a)) < a.cost) {  
          decreaseKey(a, "new cost - old cost")  
          a.path = b  
        }  
  }  
}
```



$O(|V|)$

$O(|V|\log|V|)$

$O(|E|\log|V|)$

$O(|V|\log|V| + |E|\log|V|)$

# *Dense vs. sparse again*

- First approach:  $O(|V|^2)$
- Second approach:  $O(|V|\log|V|+|E|\log|V|)$
- So which is better?
  - Sparse:  $O(|V|\log|V|+|E|\log|V|)$  (if  $|E| > |V|$ , then  $O(|E|\log|V|)$ )
  - Dense:  $O(|V|^2)$
- But, remember these are worst-case and asymptotic
  - Priority queue might have slightly worse constant factors than a linear search
  - On the other hand, for “normal graphs”, we might call **decreaseKey** rarely (or not percolate far), making  $|E|\log|V|$  more like  $|E|$