



CSE332: Data Abstractions

Lecture 17: Shortest Paths

Tyler Robison

Summer 2010



Single source shortest paths (from some specific point a)

- ▶ Done: BFS to find the minimum path length from v to u in $O(|E|)$
- ▶ Actually, can find the minimum path length from v to *every node*
 - ▶ Still $O(|E|)$
 - ▶ 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
- ▶ Aside: We can find the shortest path from every vertex to every other vertex in $O(|V|^3)$

Not as easy




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

We will assume there are no negative weights


- Problem is ill-defined if there are negative-cost *cycles*
- What is the shortest path here?
- Even without negative cycles, can still get wrong answer if negative weights are involved

Dijkstra's Algorithm (for shortest paths)

- ▶ Named after its creator Edsger Dijkstra (1930-2002)
 - ▶ Truly one of the “founders” of computer science; this is just one of his many contributions
 - ▶ Quotation: “computer science is no more about computers than astronomy is about telescopes”
- ▶ The idea: reminiscent of BFS, but adapted to handle weights
 - ▶ A priority queue will prove useful for efficiency (later)
 - ▶ Will grow the set of nodes whose shortest distance has been computed
 - ▶ Nodes not in the set will have a “best distance so far”




Dijkstra's Algorithm: Idea



▶ Conceptually:

- ▶ Grow our 'cloud' of known vertices by 1 each step
- ▶ Pick a vertex outside the cloud, that's closest to our starting point
- ▶ Guaranteed that we have the shortest path to everything within the cloud (more later)

Dijkstra's Algorithm: Idea




- ▶ Initially, start node has cost 0 and all other nodes have cost ∞
- ▶ Mark each vertex as 'unknown' (also referred to as 'unvisited', 'unexplored')
- ▶ At each step:
 - ▶ Pick closest unknown vertex v (will be start node for first step)
 - ▶ Add it to the "cloud" of known vertices
 - ▶ Update distances for nodes with edges from v
- ▶ That's it! (Have to prove it produces correct answers)

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 ,

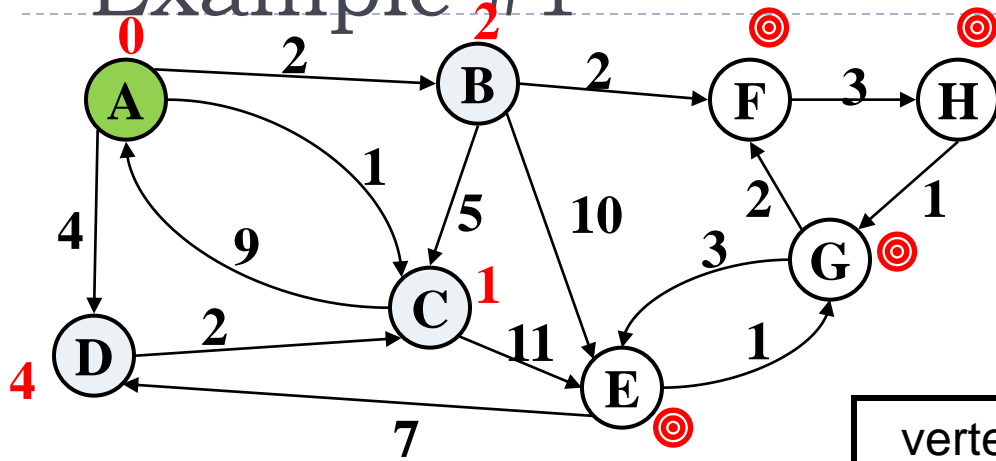
```
c1 = v.cost + w // cost of best path through v to u
c2 = u.cost // cost of best path to u previously known
if (c1 < c2) { // if the path through v is better
    u.cost = c1
    u.path = v // for computing actual paths
}
```

Example #1




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

Example #1




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

Example #1




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

Example #1




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

Example #1



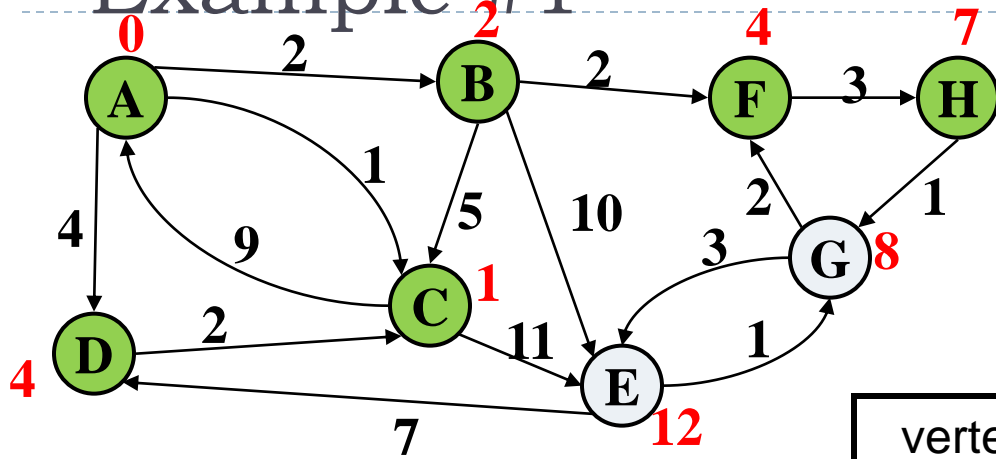
vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E		≤ 12	C
F		≤ 4	B
G		??	
H		??	

Example #1




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

Example #1




vertex	known?	cost	path
A	Y	0	
B	Y	2	A
C	Y	1	A
D	Y	4	A
E		≤ 12	C
F	Y	4	B
G		≤ 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		≤ 11	G
F	Y	4	B
G	Y	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	Y	11	G
F	Y	4	B
G	Y	8	H
H	Y	7	F

Important features

- ▶ Once a vertex is marked 'known', the cost of the shortest path to that node is known
 - ▶ As is the path itself
- ▶ While a vertex is still not known, another shorter path to it might still be found

Interpreting the results

- ▶ Now that we're done, how do we get the path from, say, A to 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 G?

- ▶ A to 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




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

Example #2




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

Example #2



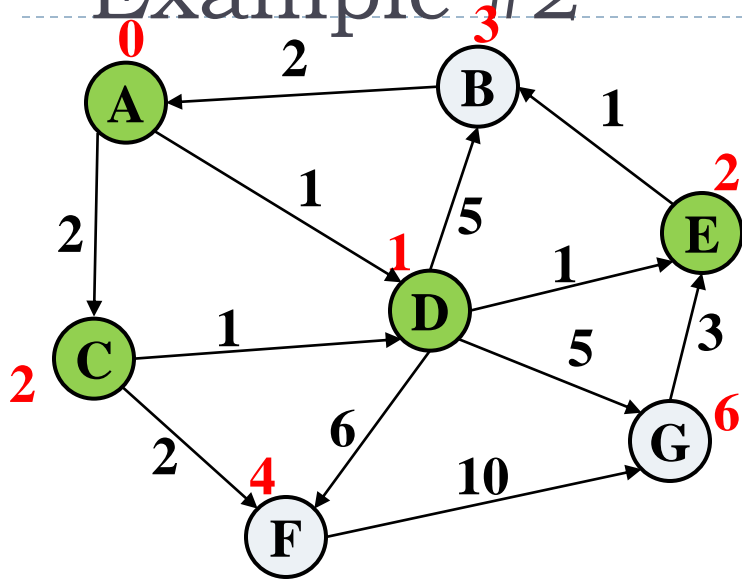
vertex	known?	cost	path
A	Y	0	
B		≤ 6	D
C		≤ 2	A
D	Y	1	A
E		≤ 2	D
F		≤ 7	D
G		≤ 6	D

Example #2




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

Example #2




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

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		≤ 4	C
G		≤ 6	D

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		≤ 6	D

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

True or false:


- ▶ If we were to count out all the edges 'found' by Dijkstra's, we would have $|V|-1$ edges

Example #3



How will the best-cost-so-far from X to Y proceed?

Example #3



How will the best-cost-so-far from X to Y proceed?
90, 81, 72, 63, 54

Where are we?

- ▶ Have described 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, irrevocably does the best thing it can at that step
 - ▶ Because of the way the algorithm is structured, the 'apparent best' actually is the best
- ▶ What should we do after learning an algorithm?
 - ▶ Prove it is correct
 - ▶ Not obvious!
 - ▶ We will sketch the key ideas
 - ▶ Analyze its efficiency

Correctness: Intuition

Rough intuition:

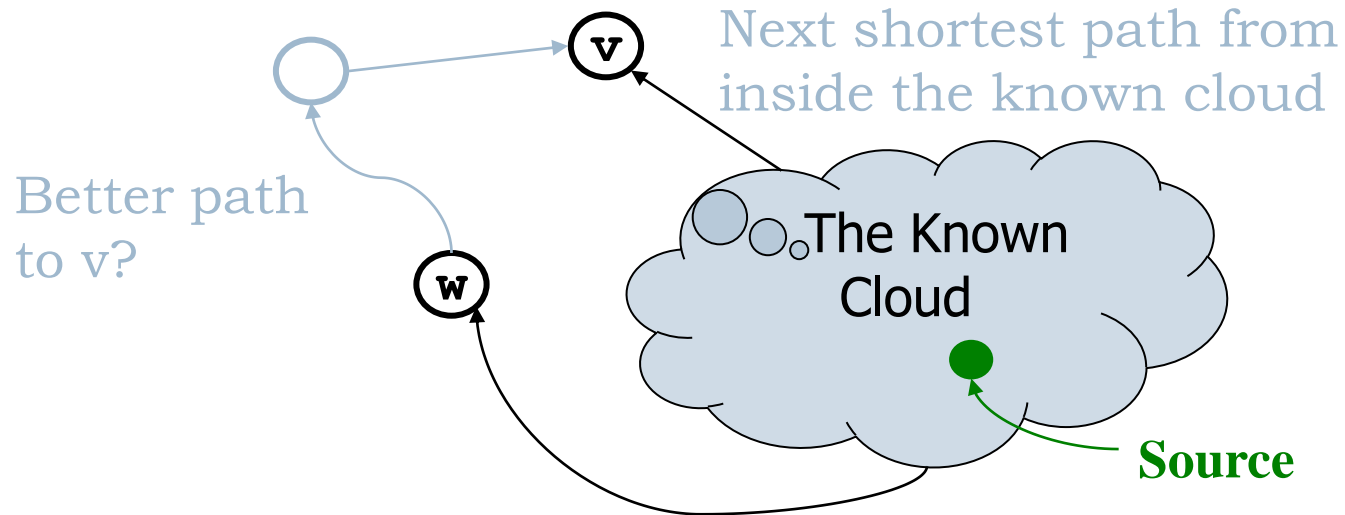
All the “known” vertices have the correct shortest path

- ▶ 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”

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 Idea)



Suppose **v** is the next node to be marked known (“added to the cloud”)

- The *best-known path* to **v** must have only nodes “in the cloud”
 - Since we’ve selected it, and we only know about paths through the cloud to a node right outside
- Assume the *actual shortest path* to **v** is different than the best-known
 - 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 ‘actual shortest 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.

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  
        }  
  }  
}
```

$O(|V|)$

$O(|V|^2)$

$O(|E|)$

$O(|V|^2)$

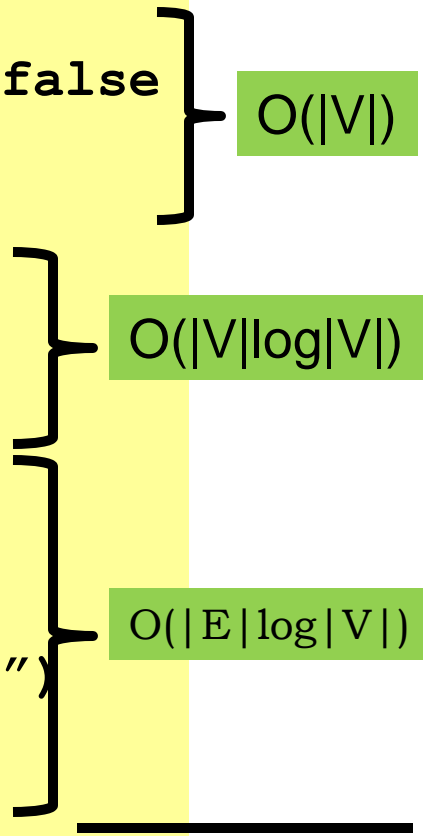
Improving asymptotic running time

- ▶ 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?
 - ▶ A priority queue holding all unknown nodes, sorted by cost
 - ▶ But must support **decreaseKey** operation
 - ▶ Must maintain a reference from each node to its position in the priority queue

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
 - ▶ On the other hand, for “normal graphs”, we might call **decreaseKey** rarely (or not percolate far), making $|E|\log|V|$ more like $|E|$