CSE 332 Data Structures & Parallelism

Shortest Paths

Melissa Winstanley Spring 2024

Shortest Path Applications

• Network routing

. . .

- Driving directions
- Cheap flight tickets
- Critical paths in project management (see textbook)

BFS is great if all we care about is path length

- But what if we care about path **cost** (ie a weighted graph)?

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*
- Today's algorithm is *wrong* if edges can be negative

Dijkstra's Algorithm

- Named after its inventor Edsger Dijkstra (1930-2002)
 - Truly one of the "founders" of computer science; 1972 Turing Award; this is just one of his many contributions
 - Sample 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



- Initially, start node has cost 0 and all other nodes have cost
- 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! (Have to prove it produces correct answers)

The Algorithm

```
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.pred = b
        }
```



vertex	known?	cost	pred
A			
В			
С			
D			
E			
F			
G			
Н			



A, C, B, D

vertex	known?	cost	pred
Α	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E		≤ 12	С
F		≤ 4	В
G		ø	
Н		ø	



A, C, B, D, **F**

vertex	known?	cost	pred
A	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E		≤ 12	С
F	Т	4	В
G		×	
Н		∞ ≤ 4+3 = 7	F



A, C, B, D, F, **H**

vertex	known?	cost	pred
Α	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E		≤ 12	С
F	Т	4	В
G		∞ ≤ 7+1 = 8	Н
Н	Т	7	F



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

vertex	known?	cost	pred
А	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E		12 ≤ 8+3=11	G
F	Т	4	В
G	Т	8	Н
Н	Т	7	F



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

vertex	known?	cost	pred
Α	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E	Т	11	G
F	Т	4	В
G	Т	8	Н
Н	Т	7	F



Path from A to E:

vertex	known?	cost	pred
A	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E	Т	11	G
F	Т	4	В
G	Т	8	Н
Н	Т	7	F



Path from A to E:

vertex	known?	cost	pred
A	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E	Т	11	G
F	Т	4	В
G	Т	8	Н
Н	Т	7	F

Stopping short

- How would this have worked differently if we were only interested in:
 - The path from A to G?
 - The path from A to D?



How would this have worked differently if we were only interested in:

- The path from A to G?
- The path from A to D?

vertex	known?	cost	pred
А	Т	0	
В	Т	2	А
С	Т	1	А
D	Т	4	А
E	Т	11	G
F	Т	4	В
G	Т	8	Н
Н	Т	7	F
	vertex A B C D E E F G H	vertexknown?ATBTCTDTFTGTHT	vertex known? cost A T 0 B T 2 C T 1 D T 4 E T 11 F T 4 G T 8 H T 7



- 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 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

Greedy failure example

Making change

Use smallest # of coins to make \$\frac{1}{5}\$ cents Options: 25, 10, 5, 1 10,5

Use smallest # of coins to make 15 cents Options: 25, 13, 10, 5, 1 7/1 \3, \, \

Where are we?

What should we do after learning an algorithm?

- Prove it is correct
 - Not obvious!
 - \circ $\,$ We will sketch the key ideas
- Analyze its efficiency
 - Will do better by using a data structure we learned earlier!



Suppose \mathbf{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 the cloud
- Assume 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 \mathbf{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!

The Algorithm - asymptotic running time

```
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.pred = b
        }
```

The Algorithm - asymptotic running time



The Algorithm - asymptotic running time



Improving asymptotic running time

● So far: O(|V²+ |E|)

- 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?

Efficiency, second approach

```
dijkstra(Graph G, Node start) {
  for each node: x.cost=infinity, x.known=false
  start.cost = 0
\rightarrow build-heap with all nodes O(\vee)
  while(heap is not empty) {
 \rightarrow b = deleteMin() \log(\sqrt{)}
    b.known = true
    for each edge (b,a) in G
      if(!a.known)
         if(b.cost + weight((b,a)) < a.cost)
       \rightarrowdecreaseKey(a, "new cost - old cost") \09\vee
           a.pred = b
         }
```

Efficiency, second approach

```
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.pred = b
       V + c_1 + V + V * (logV) + c_2 + d*(c_3 + logV)) => V + VlogV + V + E(logV)
```

Dense vs sparse again

- First approach:
- Second approach:



So which is better?

- Sparse: O(|V|log|V|+|E|log|V|) (if |E| > |V|, then O(|E|log|V|))
- Dense: O(|V|²+ |E|) , or: 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 <u>decreaseKey</u> rarely (or not percolate far), making |E|log|V| more like |E|



vertex	known?	cost	pred
A		0	
В			
С			
D			
E			
F			
G			



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

vertex	known?	cost	pred
А	Т	0	
В	Т	3	E
С	Т	2	А
D	Т	1	А
E	Т	2	D
F	Т	4	С
G	Т	6	D

But what about negative weight edges?

- We said that Dijkstra's algorithm doesn't work for **negative-weight edges**
- But what if we want to support that?

- Problem: negative-weight edges ruin our correctness proof
 - The shortest path might involve nodes outside "the cloud"
- Solution: just do our edge calculations for all edges, but |V|-1 times
 - That way we MUST consider all paths that contain all nodes

• Enter Bellman-Ford

Bellman-Ford algorithm bellmanFord(Graph G, Node start) { for each node: x.cost=infinity, x.known=false start.cost = 0 \checkmark for (i = 0; i < |V| - 1) ->for each edge (b,a) in G (if(b.cost + weight((b,a)) < a.cost){</pre> a.cost = b.cost + weight((b,a)) a.pred = b // Relax one more time to find a cycle for each edge (b,a) in G if(b.cost + weight((b,a)) < a.cost) // we found a cycle!