

# CS 6363.005.19S Lecture 15–March 12, 2019

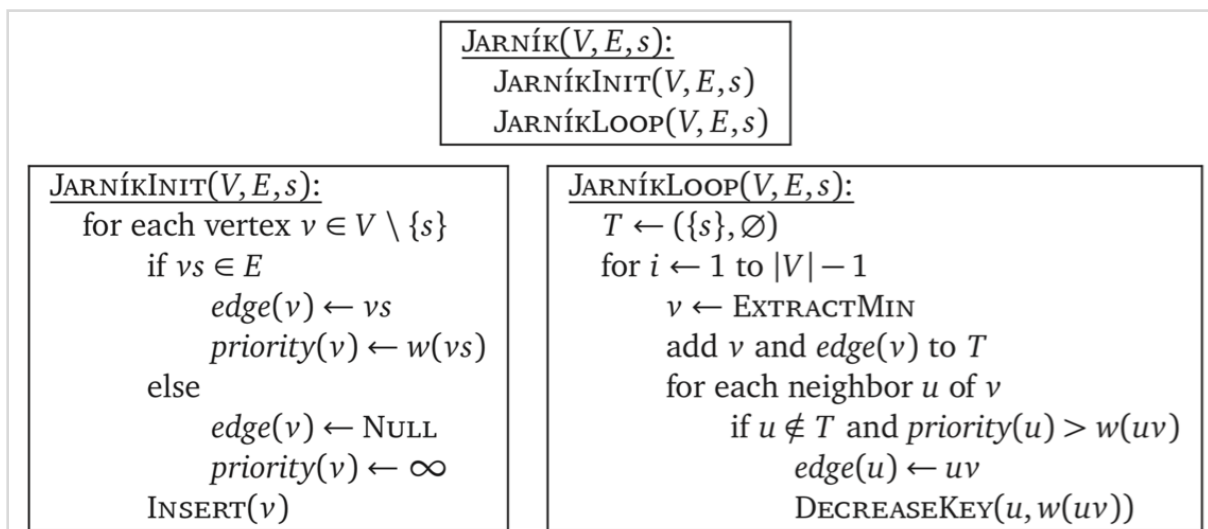
Main topics are `#single_source_shortest_paths`.

## Prelude

- Homework 3 is due **Thursday** March 14.
- Jiashuai is holding a guest lecture on matroids on Thursday.
- I'm holding extra office hours this afternoon from 2 to 3 since I'll be out town that day.

## Prim-Jarník Algorithm

- Last time, I put this algorithm up on the board for computing minimum spanning trees. I want to spend a couple minutes analyzing it before we move on to shortest paths.
- **[time permitting, write your own version that sets everything to infinity except  $\text{priority}(s) = 0$  and  $\text{edge}(s) = \text{emptyset}$ . then it's more clear that we're using an adjacency list]**



- There are  $O(E)$  DecreaseKey operations,  $O(V)$  Inserts, and  $O(V)$  ExtractMins.
- Using a normal binary heap for the priority queue, operations takes  $O(\log V)$  time each, so the algorithm takes  $O(E \log V)$  time total.
- But, we can get fancy and use something called a Fibonacci heap instead. These handle DecreaseKeys in  $O(E)$  time on average across all the operations, so the whole thing takes  $O(E + V \log V)$  time instead.
- But this is probably slower in practice unless you have very large dense graphs.

## Single Source Shortest Paths

- For the next problem, let's say you're given a *directed* graph  $G = (V, E, w)$  where  $w : E \rightarrow \mathbb{R}$  is another weight function. The shortest path between two vertices  $s$  and  $t$  is the  $s,t$ -path  $P$

minimizing  $w(P) = \sum_{\{u \rightarrow v \text{ in } P\}} w(u \rightarrow v)$ . The minimum value is the *distance* from  $s$  to  $t$ .

- Most algorithms for shortest paths actually end up solving the more general *single source shortest paths* (SSSP) problem: find the shortest path from  $s$  to every vertex in  $G$ .
- A subpath of a shortest path is itself a shortest path, and we can always pick our shortest paths consistently so they form a spanning tree, rooted at  $s$ . We'll focus on computing the shortest path tree from  $s$  and the distances from  $s$  to every other vertex.
- Please please please don't confuse minimum spanning trees with shortest path trees. They're both optimal trees but minimum spanning trees are for undirected graphs while shortest path trees are best described for directed graphs and are themselves directed away from their root. If edge weights are distinct, there is exactly one minimum spanning tree, but there is a different shortest path tree for every choice of source vertex  $s$ .
- If you want to do shortest paths in an undirected graph, replace every edge  $uv$  with a pair of edges  $u \rightarrow v$  and  $v \rightarrow u$  of the same weight. Again, the shortest path trees may all be different from the minimum spanning tree.
- Like with minimum spanning trees, there's nothing saying we can't have negative weights. If you think of positive weights as a cost for following certain edges, negative weight would represent some benefit. Maybe you're trying to plan a trip and there's a few particularly pretty roads you'd like to drive down.
- However, we run into trouble if there's a directed cycle with negative total weight. What we're *really* going to compute are shortest walks from  $s$  to every vertex. If there are negative weight cycles, a "shortest" walk would go around and around and around and infinite number of times before reaching its destination; i.e. it's not well-defined.
- If there are no negative weight cycles, though, then we can compute shortest walks and they'll be paths (no reason to repeat a vertex if coming back to it has non-negative cost).
- The trick of turning an undirected graph into a directed one only works if there are non-negative weights, because otherwise you'd create tiny negative weight cycles. You can do shortest paths in undirected graphs with negative weights, but the algorithms are well beyond the scope of this course.

## The Only SSSP Algorithm

- Like minimum spanning trees, there's really only one shortest path tree algorithm independently discovered by Lester Ford, George Dantzig, and George Minty around the same time.
- The idea is that we'll keep an educated guess on the distance and shortest path to each vertex.
- $\text{dist}(v)$  is the length of a tentative shortest  $s$  to  $v$  path, or infinity if we haven't found one yet.
- $\text{pred}(v)$  is the predecessor of  $v$  in the tentative shortest  $s$  to  $v$  path, or Null if we haven't found one yet.

- At the beginning of the algorithm, we know  $\text{dist}(s) = 0$  and  $\text{pred}(s) = \text{Null}$ . For every other vertex  $v \neq s$ , we initially set  $\text{dist}(v) = \text{infinity}$  and  $\text{pred}(v) = \text{Null}$ , because we haven't found *any* paths to those vertices yet!

<p><u>INITSSSP(s):</u>  <math>\text{dist}(s) \leftarrow 0</math>  <math>\text{pred}(s) \leftarrow \text{NULL}</math>  for all vertices <math>v \neq s</math>  <math>\text{dist}(v) \leftarrow \infty</math>  <math>\text{pred}(v) \leftarrow \text{NULL}</math></p>
---

- Call an edge  $u \rightarrow v$  *tense* if  $\text{dist}(u) + w(u \rightarrow v) < \text{dist}(v)$ .
- If an edge is tense, then the distances are certainly not shortest path distances.
- We want to *relax* tense edges to represent our newly found shorter path.

<p><u>RELAX(<math>u \rightarrow v</math>):</u>  <math>\text{dist}(v) \leftarrow \text{dist}(u) + w(u \rightarrow v)</math>  <math>\text{pred}(v) \leftarrow u</math></p>
--

- The only SSSP algorithm repeatedly finds some tense edge and relaxes it.

<p><u>FORDSSSP(s):</u>  INITSSSP(s)  while there is at least one tense edge  RELAX any tense edge</p>
---

- Now, I could prove correctness for this algorithm, but I can't say too much about running time, because that depends heavily on the order in which we relax edges.
- In particular, there are different relaxation orders depending on the time of graph and weights you use.
- If all weights are 1 (your graph is unweighted), use a BFS in  $O(V + E)$  time.
- If the graph is a DAG, use the dynamic programming algorithm I showed you but with a min, again in  $O(V + E)$  time.
- There are two more common algorithms. To avoid repeating myself too much, I'll prove correctness and analyze running more-or-less separately for the two special cases.
- One observation we'll need for both algorithms, though, is that if  $\text{dist}(v) \neq \text{infinity}$ , then it is the length of *some* walk from  $s$  to  $v$ .
  - We can use induction on the number of relaxations.
  - If the last change to  $\text{dist}(v)$  was setting  $\text{dist}(v) \leftarrow \text{dist}(u) + w(u \rightarrow v)$  then  $\text{dist}(u)$  at that moment was the length of some  $s$  to  $u$  walk.
  - We just added  $u \rightarrow v$  to that walk to make an  $s$  to  $v$  walk of length  $\text{dist}(u) + w(u \rightarrow v) = \text{dist}(v)$ .
- In particular,  $\text{dist}(v)$  is always *at least* the shortest path distance from  $s$  to  $v$ .
- We can use a similar induction proof to show the predecessors give shortest paths when no edges are tense, but I'll only focus on distances today.

## No (or few) Negative Edges: Dijkstra's Algorithm

- The first algorithm was independently discovered by many many many different researchers, but by now everybody has decided to call it Dijkstra's algorithm.
- It's similar to Prim-Jarník. We'll use a priority queue on vertices with the key of vertex  $v$  being  $\text{dist}(v)$ . We repeatedly extract the minimum distance vertex and relax all outgoing edges. When  $\text{dist}(w)$  changes for some vertex  $w$ , we'll insert it in the priority queue or decrease its key.

```
DIJKSTRA( $s$ ):  
  INITSSSP( $s$ )  
  INSERT( $s, 0$ )  
  while the priority queue is not empty  
     $u \leftarrow \text{EXTRACTMIN}()$   
    for all edges  $u \rightarrow v$   
      if  $u \rightarrow v$  is tense  
        RELAX( $u \rightarrow v$ )  
      if  $v$  is in the priority queue  
        DECREASEKEY( $v, \text{dist}(v)$ )  
      else  
        INSERT( $v, \text{dist}(v)$ )
```

- Now, this is an instance of Ford's general strategy, so it computes shortest paths as long as there are no negative cycles in the graph.
- But something special happens if there are no negative edges at all.
- Intuitively, you could imagine a wavefront spreading out from  $s$ , passing over vertices in increasing order of their distance and never returning to a vertex that's already been passed over.
- Let's analyze this algorithm and prove correctness assuming no negative weight edges. Let  $u_i$  be the vertex returned by the  $i$ th ExtractMin call (so  $u_1 = s$ ) and  $d_i$  be  $\text{dist}(u_i)$  just after the Extraction (so  $d_1 = 0$ ). We cannot assume yet these vertices are distinct. For all  $i < j$  we know  $u_i = u_j$  for some  $i < j$ .
- Lemma: For all  $i < j$ , we have  $d_i \leq d_j$ . (Vertices are extracted in non-decreasing order of distance.)
  - Fix some  $i$ . By induction, it suffices to prove  $d_{i+1} \geq d_i$ , meaning the distances are non-decreasing.
  - If  $u_i \rightarrow u_{i+1}$  is relaxed during the  $i$ th iteration, then afterward  $d_{i+1} = \text{dist}(u_{i+1}) = \text{dist}(u_i) + w(u_i \rightarrow u_{i+1}) \geq \text{dist}(u_i) = d_i$ .
  - Otherwise,  $u_{i+1}$  is already in the priority queue when we extract  $u_i$ . But we extracted  $u_i$  so  $d_i = \text{dist}(u_i) \leq \text{dist}(u_{i+1}) = d_{i+1}$ .
- Lemma: Each vertex is extracted from the priority queue at most once.
  - Suppose  $v = u_i = u_k$  for  $i < k$ .

- $\text{dist}(v)$  never increases and  $v$  is reinserted in the queue only when  $\text{dist}(v)$  decreases, so  $d_k < d_i$ . But that contradicts the previous lemma.
- Lemma: When Dijkstra ends,  $\text{dist}(v)$  is the length of the shortest path from  $s$  to  $v$  for every vertex  $v$ .
  - For any vertex  $v$ , consider some shortest path  $v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_{\text{ell}}$  where  $v_0 = s$  and  $v_{\text{ell}} = v$ . Let  $L_j$  be the length of the subpath  $v_0 \rightarrow \dots \rightarrow v_j$ . We'll prove by induction on  $j$  that  $\text{dist}(v_j) \leq L_j$ .
  - $\text{dist}(v_0) = \text{dist}(s) = 0 = L_0$ .
  - Consider  $j > 0$ . By induction, we set  $\text{dist}(v_{j-1})$  and at some point Extract  $v_{j-1}$  from the queue. At that moment either  $\text{dist}(v_j) \leq \text{dist}(v_{j-1}) + w(v_{j-1}, v_j)$  already or we set  $\text{dist}(v_j) = \text{dist}(v_{j-1}) + w(v_{j-1}, v_j)$  by the end of the iteration. Either way
    - $\text{dist}(v_j) \leq \text{dist}(v_{j-1}) + w(v_{j-1}, v_j) \leq L_{j-1} + w(v_{j-1}, v_j) = L_j$ .
  - In particular,  $\text{dist}(v) \leq L_{\text{ell}} =$  the length of the whole path.
  - Again,  $\text{dist}(v)$  is at least the shortest path distance and therefore equal to it.
- Just like Prim-Jarník, we have  $V$  Insert and ExtractMin operations and  $E$  DecreaseKey operations. With a min-heap that all takes  $O(E \log V)$  time. With a Fibonacci heap, it's only  $O(E + V \log V)$  time.
- Again, this algorithm, as I wrote it, works just fine if you have negative edge lengths but no negative cycles. In fact, it's likely to be faster than the next algorithm I present if you only have a few negative length edges.
- You could also write a version that never puts a vertex back in the priority queue as CLRS does, but then it's incorrect if there's some negative length edges.

## If All Else Fails: Bellman-Ford

- OK, so what if you have some negative weights and you don't have a DAG and you want to prove a good performance guarantee?
- Again, this algorithm was proposed by many people, but everybody calls it Bellman-Ford now.
- We just relax all tense edges and then recurse.

```

BELLMANFORD(s)
  INITSSSP(s)
  while there is at least one tense edge
    for every edge  $u \rightarrow v$ 
      if  $u \rightarrow v$  is tense
        RELAX( $u \rightarrow v$ )

```

- This algorithm is somehow a bit easier to analyze.
- Let  $\text{dist}_{\leq i}(v)$  denote the length of the shortest *walk* in  $G$  from  $s$  to  $v$  with *at most*  $i$  edges. So  $\text{dist}_{\leq 0}(s) = 0$  and  $\text{dist}_{\leq 0}(v) = \text{infinity}$  for all  $v \neq s$ .

- Lemma: For every vertex  $v$  and non-negative integer  $i$ , after  $i$  iterations we have  $\text{dist}(v) \leq \text{dist}_{\leq i}(v)$ .
- Proof:
  - If  $i = 0$ , the lemma is trivially true.
  - Let  $W$  be a shortest walk from  $s$  to  $v$  with at most  $i$  edges. By definition,  $W$  has length  $\text{dist}_{\leq i}(v)$ .
  - If  $W$  has no edges, it goes from  $s$  to  $s$ , meaning  $v = s$  and  $\text{dist}_{\leq i}(v) = 0$ .  $\text{dist}(s) \leftarrow 0$  in  $\text{InitSSSP}$  and  $\text{dist}(s)$  never increases, so  $\text{dist}(s) \leq 0$ .
  - Otherwise, let  $u \rightarrow v$  be the last edge of  $W$ . After  $i - 1$  iterations,  $\text{dist}(u) \leq \text{dist}_{\leq i-1}(u)$ .
  - In the  $i$ th iteration, we consider edge  $u \rightarrow v$ . Either  $\text{dist}(v) \leq \text{dist}(u) + w(u \rightarrow v)$  already or we set  $\text{dist}(v) \leftarrow \text{dist}(u) + w(u \rightarrow v)$ . Either way,  $\text{dist}(v) \leq \text{dist}_{\leq i-1}(u) + w(u \rightarrow v) = \text{dist}_{\leq i}(v)$ . Again,  $\text{dist}(v)$  does not increase after that, although it may decrease further by the time the loop ends.
- This lemma is true even if there are negative length cycles!
- Again,  $\text{dist}(v)$  is always at least the shortest path distance.
- If there are no negative cycles, the shortest walk from  $s$  to any  $v$  has at most  $V - 1$  edges, so  $\text{dist}(v)$  must be the true shortest path distance by the end of  $V - 1$  iterations.
- Each iteration takes  $O(E)$  time, so the algorithm takes  $O(VE)$  time if there are no negative length cycles.
- That said, maybe there are negative length cycles and your distances are not shortest path distances. You can do yet another proof by induction to show there will always be a tense edge if there are any dist values that are too high.
- So there will be a tense edge after those  $V - 1$  iterations, and we can modify the algorithm slightly to detect negative cycles.

```

BELLMANFORD(s)
  INITSSSP(s)
  repeat V - 1 times
    for every edge  $u \rightarrow v$ 
      if  $u \rightarrow v$  is tense
        RELAX( $u \rightarrow v$ )
  for every edge  $u \rightarrow v$ 
    if  $u \rightarrow v$  is tense
      return "Negative cycle!"

```

- This version runs in  $O(VE)$  time even if there are negative cycles.