

Greedy Algorithms II

[Chapter 24 Cormen 2001]

- Today's class will discuss Dijkstra's algorithm which is a greedy algorithm

- But first we need to talk about single-source shortest paths

Single-source shortest paths:

- In a shortest-paths problem we are given a weighted directed graph $G = (V, E)$ with weight function $w: E \rightarrow \mathbb{R}$ mapping edges to real-valued weights

- The weight $w(p)$ of path $p = \langle v_0, v_1, \dots, v_n \rangle$ is the sum of the weights of its constituent edges:

$$w(p) = \sum_{i=1}^k w(v_{i-1}, v_i)$$

- We define the shortest-path weight $\delta(u, v)$ from u to v by:

$$\delta(u, v) = \begin{cases} \min\{w(p) : u \xrightarrow{p} v\}, & \text{if there is a path from } \underline{u} \text{ to } \underline{v} \\ \infty & \text{, otherwise} \end{cases}$$

- A shortest path from vertex u to vertex v is then defined as any path p with weight $w(p) = d(u, v)$

- Single-source shortest-paths problem: given a graph $G = (V, E)$, we want to find a shortest path from a given source vertex $s \in V$ to each vertex $v \in V$

- Shortest-paths algorithms typically rely on the property that a shortest path between two vertices contains other shortest paths within it

- Optimal substructure property
 ↳ We saw this on the class "Greedy Algorithms I"
 ↳ Remember MSTs?

- The following Lemma states the optimal-substructure property of shortest paths more precisely:

Lemma 24.1 - Subpaths of shortest paths are shortest paths

Given a weighted, directed graph $G = (V, E)$ with weight function $w: E \rightarrow \mathbb{R}$, let $p = \langle v_0, v_1, \dots, v_k \rangle$ be a shortest path from vertex v_0 to vertex v_k and, for any i and j such that $0 \leq i \leq j \leq k$, let $p_{ij} = \langle v_i, v_{i+1}, \dots, v_j \rangle$ be the subpath of p from vertex v_i to v_j . Then p_{ij} is a shortest

- Proof:

- If we decompose path p into:

$$v_0 \xrightarrow{p_{0i}} v_i \xrightarrow{p_{ij}} v_j \xrightarrow{p_{jk}} v_k$$

Then we have that:

$$w(p) = w(p_{0i}) + w(p_{ij}) + w(p_{jk})$$

- Now, assume that there is a path p'_{ij} from v_i to v_j with weight $w(p'_{ij}) < w(p_{ij})$. Then:

$$v_0 \xrightarrow{p_{0i}} v_i \xrightarrow{p'_{ij}} v_j \xrightarrow{p_{jk}} v_k$$

is a path from v_0 to v_k whose weight $w(p_{0i}) + w(p'_{ij}) + w(p_{jk})$ is less than $w(p)$, which contradicts the assumption that p is a shortest path from v_0 to v_k .

- Issues regarding cycles (1/2):

- Negative-weight edges: If the graph contains a negative-weight cycle reachable from s , ~~there is~~ no path from s to a vertex on the cycle can be a shortest path. We can always find a path with lower weight by following the proposed "shortest" path and then traversing the negative weight cycle. We would thus obtain another "shortest" path.

- Issues regarding cycles (2/2):

• "Normal" cycles: A shortest path cannot contain a normal cycle, since removing the cycle from the path produces a path with the same source and destination and a lower weight.

∴ Shortest paths cannot contain: $\left. \begin{matrix} \text{negative} \\ \text{positive} \end{matrix} \right\}$ cycles

• That leaves only 0-weight cycles. We can remove a 0-weight cycle from any path to produce another path whose weight is the same.

∴ We can assume, without loss of generality, that when we are finding shortest paths we have no cycles.

- We often wish to compute not only shortest-path weights, but the vertices on shortest paths as well. Given a graph $G = (V, E)$ we maintain for each vertex $v \in V$ a predecessor $v.\pi$ that is either another vertex or \underline{NIL} .

- Relaxation:

- The algorithms in Chapter 24 of the book use the technique of relaxation.

- For each vertex $v \in V$ we maintain an attribute $v.d$ which is an upper bound on the weight

- We call $v.d$ a shortest-path estimate. We initialize the shortest-path estimates and predecessors by the following $O(V)$ -time procedure:

```

Initialize Single Source (G, s)
for each vertex v ∈ G.V
    v.d = ∞
    v.π = Nil
s.d = 0

```

- The process of relaxing an edge (u, v) consists of testing whether we can improve the shortest path to v found so far by going through u and, if so, updating v.d and v.π

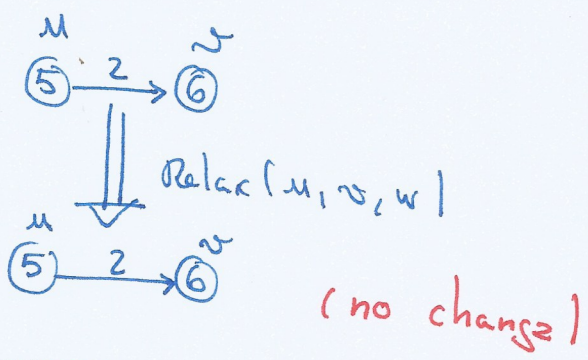

- A relaxation step may decrease the value of the shortest-path estimate v.d and update v's predecessor attribute v.π. The following code performs a relaxation step on edge (u, v) in $O(1)$ time:

```

Relax (u, v, w)
    If v.d > u.d + w(u, v)
        v.d = u.d + w(u, v)
        v.π = u

```


- Example:



Notes:

- The shortest path estimate of each vertex appears within the vertex

- Relaxation is the only means by which shortest path estimates and predecessors change.

- Properties of shortest paths and relaxation

- Proofs can be found on section 24.5 of Cormen 2004

- Triangle inequality (Lemma 24.10)

For any edge $(u, v) \in E$, we have $\delta(s, v) \leq \delta(s, u) + w(u, v)$
shortest-path weight

- Upper-bound property (Lemma 24.11)

We always have $n.d \geq \delta(s, v)$ for all vertices $v \in V$, and once $n.d$ achieves the value $\delta(s, v)$ it never changes

- No-path property (Lemma 24.12)

If there is no path from s to v , then we always have $n.d = \delta(s, v) = \infty$

Convergence Property (Lemma 24.14)

If $s \xrightarrow{\text{path to } u} u \xrightarrow{\text{edge } (u, v)} v$ is a shortest path in G for some $u, v \in V$ and if $u.d = \delta(s, u)$ at any time prior to relaxing edge (u, v) then $v.d = \delta(s, v)$ at all times afterward

Path-Relaxation Property (Lemma 24.15)

○ If $p = \langle v_0, v_1, \dots, v_k \rangle$ is a shortest path from $s = v_0$ to v_k and we relax the edges of p in the order $(v_0, v_1), (v_1, v_2), \dots, (v_{k-1}, v_k)$ then $v_k.d = \delta(s, v_k)$

This property holds regardless of any other relaxation steps that occur, even if they are intermixed with relaxations of the edges of p .

Predecessor-subgraph property (Lemma 24.17)

Once $v.d = \delta(s, v)$ for all $v \in V$, the predecessor subgraph is a shortest-paths tree rooted at s

24.3 Dijkstra's algorithm

- Bellman-Ford algorithm run in $O(VE)$
- Dijkstra's algorithm solves the single-source shortest path problem on a weighted directed graph $G = (V, E)$ for the case in which all edge weights are nonnegative.
- I.e. assume $\forall (u, v) \in E \quad w(u, v) \geq 0$
- With a good implementation the running time is better than Bellman-Ford
- Dijkstra's algorithm maintains a set S of vertices whose final shortest-path weights from the source have already been determined
- The algorithm repeatedly selects the vertex $u \in V - S$ with the minimum shortest-path estimate, adds u to S and relaxes all edges leaving u
- The following implementation uses a min-priority queue of vertices, keyed by their d values.

Question:

Do this remind you of anything?

Dijkstra (G, w, s)

Initialize Single Source (G, s)

$S = \emptyset$

$Q = G.V$

while $Q \neq \emptyset$

$u = \text{ExtractMin}(Q)$

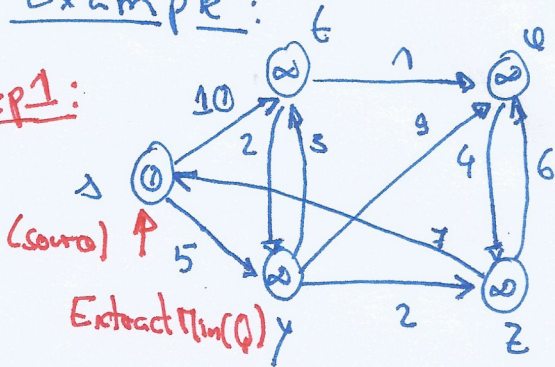
$S = S \cup \{u\}$

for each vertex $v \in G.Adj[u]$

Relax (u, v, w) // Can we improve the shortest path estimate? * /

- Example:


Step 1:



Caption:

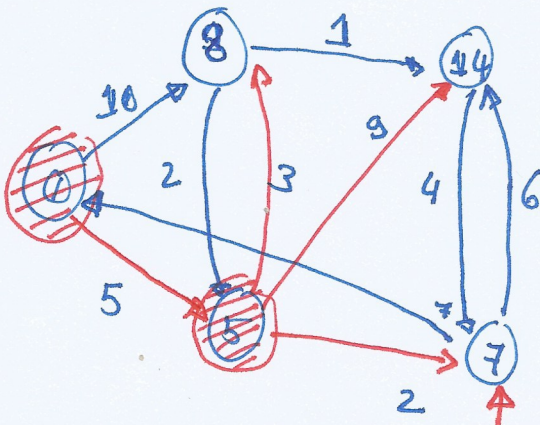
- Shortest path estimates appear within the vertices
- ~~Shaded~~ edges marked with - indicate predecessor values
- Vertices marked as ⊙ are in the set S
- Vertices not marked are in the min priority queue.

Step 2:



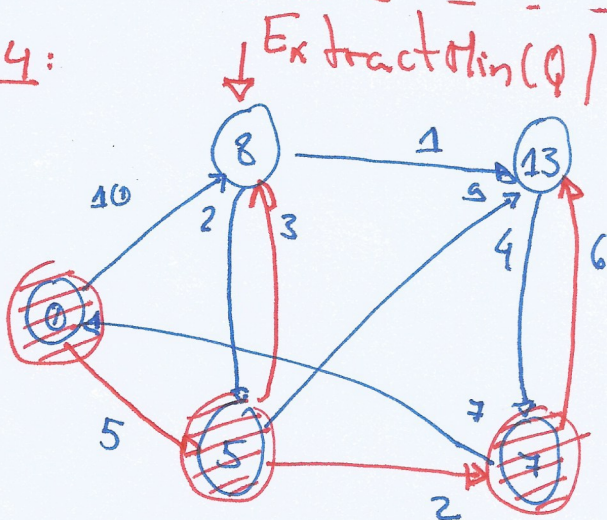
Extract Min(Q)

Step 3:

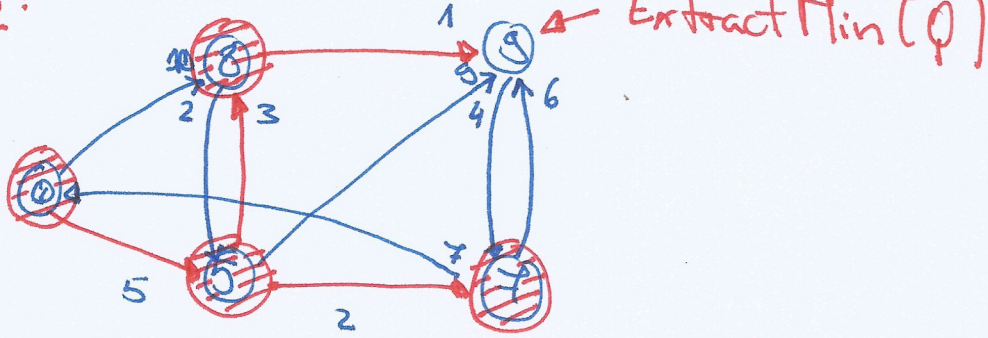


Extract Min(Q)


Step 4:



Step 5:



Step 6:



Nothing to update

Notes:

- The algorithm never inserts vertices into Q after the initialization
- Each vertex is extracted from Q and added to S exactly once, so the loop iterates $|V|$ times
- Because Dijkstra's algorithm always chooses the closest vertex in $V-S$ to add to set S we say it uses a greedy strategy.

Question:

But how can we know that Dijkstra's algorithm does indeed compute shortest-paths?

Idea: Show that each time a vertex u is added to set S we have $u.d = \delta(S, u)$

Theorem 24.6 - Correctness of Dijkstra's algorithm

- At the start of each iteration of the loop:

$$u.d = \delta(S, u) \quad \forall u \in S$$

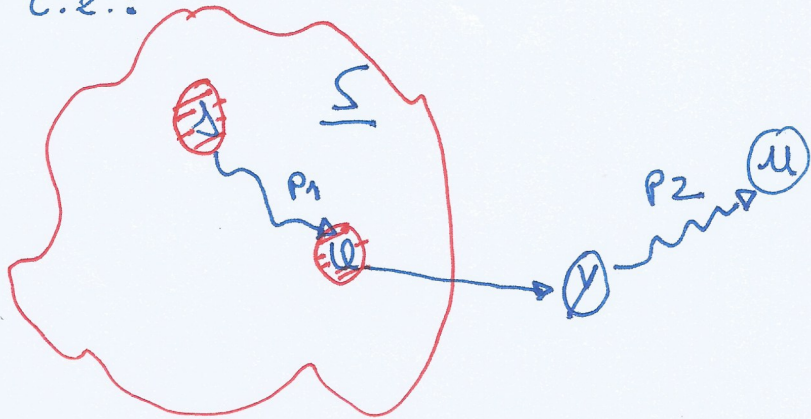
We need to show this condition always holds

- First iteration (initialization): there are no vertices in S so the condition is true

- Maintenance:

- We wish to show that $u.d = \delta(S, u)$ for the vertex added to S
- For the purpose of contradiction, let u be the first vertex for which $u.d \neq \delta(S, u)$ when u is added to S
- We must have $u \neq s$ because s is the first vertex added to S and $s.d = \delta(S, s) = 0$.

- Because $u \neq s$ we also have that $S \neq \emptyset$ just before \underline{u} is added to \underline{S}
- There must be some path from \underline{s} to \underline{u} for otherwise $u.d = \delta(s, u) = \infty$ by the no-path property, which would violate our assumption that $u.d \neq \delta(s, u)$
- Because there is at least one path, there is a shortest path p from \underline{s} to \underline{u}
- Prior to adding \underline{u} to \underline{S} path p connects a vertex in \underline{S} namely \underline{s} , to a vertex in $\underline{V-S}$, namely \underline{u} .
- Let us consider the first vertex y along p such that $y \in \underline{V-S}$ and let $u \in \underline{S}$ be y 's predecessor along p
- Then we can decompose path p into $s \xrightarrow{p_1} u \rightarrow y \xrightarrow{p_2}$
i.e.:



- We claim that $y.d = \delta(s, y)$ when \underline{u} is added to \underline{S} .
This happens because we chose \underline{u} as the first vertex for which $u.d \neq \delta(s, u)$ when it is added to \underline{S} .
- We had $u.d = \delta(s, u)$ when \underline{u} was added to \underline{S} .
- Edge (u, y) was relaxed at that time and the claim follows from the convergence property (page 7 of these notes).

• We can now obtain a contradiction to prove that
 $u.d$ must be $\delta(s, u)$, i.e.: $u.d = \delta(s, u)$

- Because \underline{y} appears before \underline{u} on a shortest path from \underline{s} to \underline{u} and all edge weights are non-negative we have $\delta(s, y) \leq \delta(s, u)$ and thus

$$y.d = \delta(s, y) \leq \delta(s, u)$$

$$\leq u.d \quad (\text{by the upper-bound prop.})$$

- But because both vertices \underline{u} and \underline{y} were in $\underline{V-S}$ when \underline{u} was chosen we have $u.d \leq y.d$. ~~I.e. we~~

~~have not yet updated our estimate for \underline{u}~~ This happens because the ExtractMin(\mathcal{Q}) operation will return vertex \underline{u} .

• Thus

$$\begin{cases} y \cdot d \leq u \cdot d \\ u \cdot d \leq y \cdot d \end{cases} \Rightarrow y \cdot d = u \cdot d \Leftrightarrow y \cdot d = \delta(\underline{s}, y) = \delta(\underline{s}, u) = u \cdot d$$

• Consequently: $u \cdot d = \delta(\underline{s}, u)$ which contradicts our choice of \underline{u} .

• $\therefore u \cdot d = \delta(\underline{s}, u)$ when \underline{u} is added to \underline{s} and that this equality is maintained all times thereafter.

Termination:

• At termination $\emptyset = \emptyset$ but initially $\emptyset = V - S$

$$\begin{cases} \emptyset = \emptyset \\ \emptyset = V - S \end{cases} \Rightarrow \emptyset = V - S \Leftrightarrow V = S$$

• Thus, when combined with our maintenance result we know that $\delta(\underline{s}, u) \forall u \in V$

▣ end of proof
(finally ☺)

Question: How fast is Dijkstra's algorithm?

• Algorithm maintains min-priority queue Q by the operations:

- Insert
 - Extract Min
 - Decrease key (From the relax operation)
- } These are called once per vertex

• Each vertex u is added to S exactly once and each edge in the adjacency list $adj[u]$ is also examined exactly once

• Total number of edges is $|E|$

↳ Therefore Decrease key is called $|E|$ times

• Overall complexity depends on how Q is implemented:

Q	Insert	Extract Min	Decrease key	Total
Array	$O(1)$	$O(V)$	$O(1)$	$O(V^2 + E)$
Heap	$O(\lg V)$	$O(\lg V)$	$O(\lg V)$	$O((V+E)\lg V)$
Fib Heap	$O(1)$	$O(\lg V)$	$O(1)$	$O(V \lg V + E)$

Total cost all vertices the number for (oo)