


Shortest paths

$G = (V, E)$

Edge weights / lengths / distances.

Visited set = $\{s, u\}$

\downarrow
 $\{s, u, v\}$



$d'[y]=4, d'[v]=2, d'[x]=3$

Start from s and pick nearest neighbours. Mark s as "visited".

Explore all the neigh

$d'[x] = \min \{ d[u] + 2, d[v] + 4 \}$

Question: Given $G=(V,E), s,t \in V(G)$, what is the shortest distance between s and t ? What is that path.

Dijkstra's algorithm (G, L, s)

start node.

list of weights/lengths of edges $\in \mathbb{R}_{\geq 0}$

Output: Return a list of shortest distances of every node from s .

$S \leftarrow \{s\}; d(s)=0$ // S is the set of all "visited" nodes.

For every $v \neq s \in V: d(v) = \infty$. // d maintains shortest distance from s for every node

While $S \neq V$:

Select $v \notin S$ w/ at least one edge from S s.t

$d'(v) = \min_{(u,v) \in E} \{ d(u) + l_{u,v} \}$ is minimized. // Greedy

$$d[v] = d'(v)$$

Add v to S' .

Correctness:

Lemma: Consider the set S at an arbitrary point of the algorithm's execution. For all $u \in S$, $d[u]$ is the shortest $S \rightsquigarrow u$ distance.

Proof: Induction on $|S|$.

Base case: $|S|=1$. $S = \{s\}$. $d[s] = 0$.

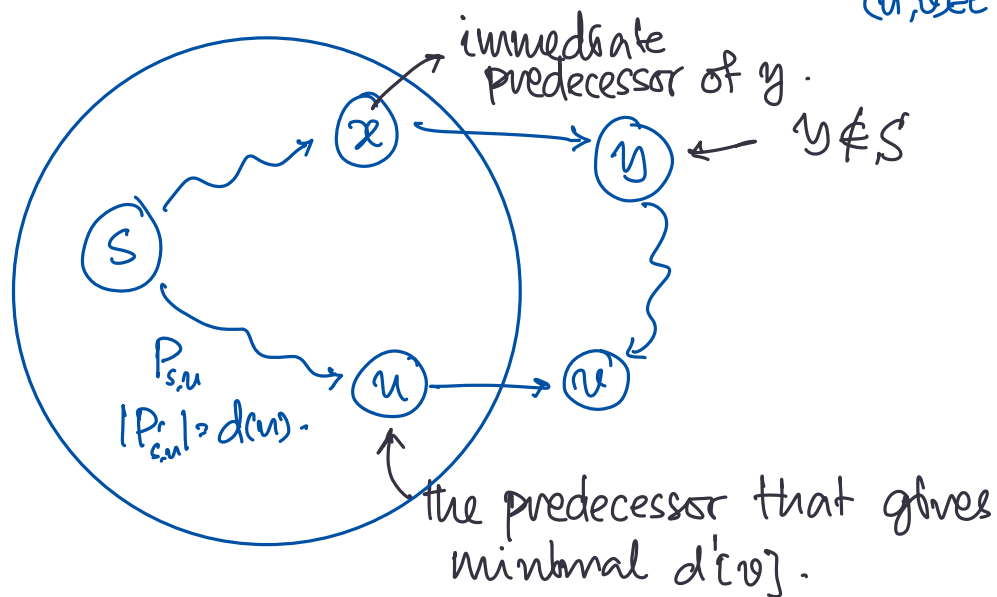
Induction hypothesis: The statement is true for $|S|=k$ for some $k \geq 1$. $k \leq n-1$.

→ Algorithm picks v s.t. $d'[v]$ is the minimum over $\{d'[v'] : v' \in V \setminus S\}$.

$$d'[v] = \min \{d'[v'] : v' \notin S\}.$$

→ Algo set $d[v]$ to $d'[v]$.

$$d'[v] = \min_{(u, v) \in E} \{d[u] + l_{u,v}\}.$$



Algo computed d' values for all vertices not in S .

↳ Algo picked v over $y \Rightarrow d'[v] \leq d'[y]$.

For the sake of contradiction, let the shortest path from s to v be via $s \sim x \sim y \sim v$. $y \neq v$.

→ Since v was picked no path from s to y was shorter than s to v .

$$\begin{array}{l} \overbrace{s \sim x \rightarrow y \sim v} \\ \text{dist: } \underbrace{l(s,x) + l_{x,y} + l(y,v)} \\ \geq \underbrace{d[x] + l_{x,y} + l(y,v)} \\ \geq \underbrace{d'[y] + l(y,v)} \\ \geq \underbrace{d'[v] + l(y,v)} \end{array} \quad \left. \begin{array}{l} s \sim u \rightarrow v \\ \underline{l(s,u)} + l_{u,v} \\ = \underline{d[u]} + l_{u,v} \\ = \underline{d[v]} \end{array} \right\} \geq$$

distance for any other path (say through $y \neq s$) will at least be $\underline{d'[y] + l(y,v)}$.
and if v was picked before y , then $d'[y] \geq d'[v] = d[v]$

$\Rightarrow d[v] = d'[v]$ is in fact the shortest s to v distance.

Implementation and running time:

- $O(n)$ initialization
- Priority queue: maintain $d'[u]$ for every $u \notin S$.
 - ExtractMin operation for every iteration

ChangeKey ← m many operations overall.

$O(n) + n$ ExtractMin + m ChangeKey + $O(n)$ overhead.