

Algorithms

Shortest Paths



Recap


- HW Posted
- Reading for Mon. + Wed. posted.
- Stay tuned for Friday

Next problem: Shortest paths

Goal: Find shortest path from s to v .

We'll think directed, but really could do undirected w/no negative edges:

Motivation:



- maps
- routing


Usually, to solve this, need to solve a more general problem:

Find shortest paths from s to every other vertex.

Called the single-source shortest path tree.

Some notes:

- Why a tree?



If $s \rightarrow a \rightarrow b \rightarrow c \rightarrow d \rightarrow v$ and $s \rightarrow a \rightarrow x \rightarrow y \rightarrow d \rightarrow u$ are shortest paths, then $s \rightarrow a \rightarrow b \rightarrow c \rightarrow d \rightarrow u$ is also a shortest path.

If 2 shortest paths cross twice, subpaths must be tied.

- Negative edges?




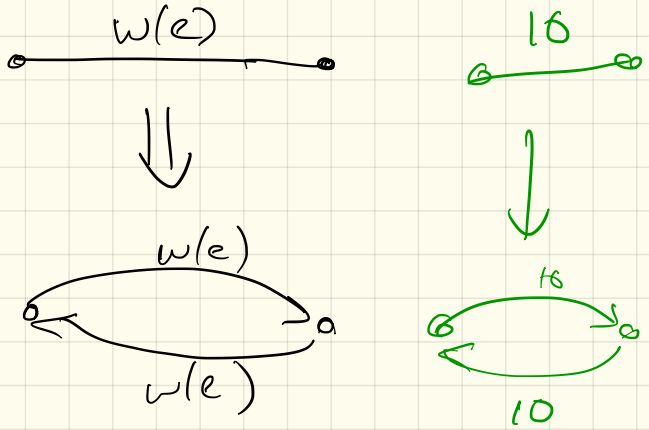
Figure 8.3. There is no shortest walk from s to t.

No repeat edges




Also: If undirected, can simulate w/ directed.

ie



Unless!! negatives



B/c gets weird:





Figure 8.4. An undirected graph where shortest paths from s are unique but do not define a tree.


How to solve ?? **Flows!**

Important to realize:
MST \neq SSSP

MST



SSSP



Why?

MST is optimizing some global structure

SSSP is local.

Computing a SSSP:

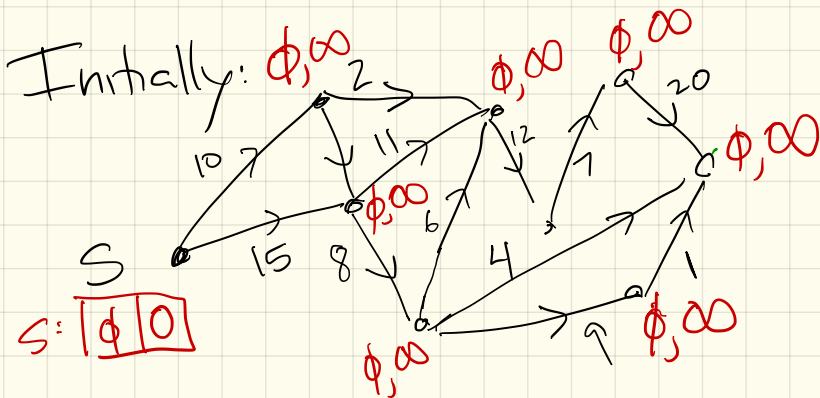
(Ford 1956 + Dantzig 1957)

Each vertex will store 2 values.

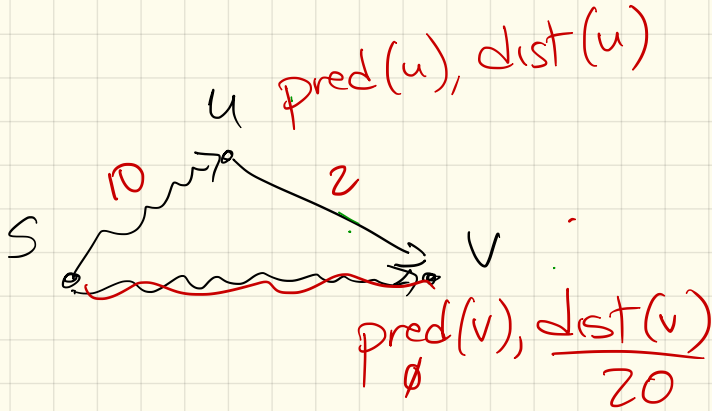
(Think of these as tentative shortest paths)

- $\text{dist}(v)$ is length of tentative shortest $S \rightsquigarrow v$ path
(or ∞ if don't have an option yet)

- $\text{pred}(v)$ is the predecessor of v on that tentative path $S \rightsquigarrow v$
(or NULL if none)



We say an edge \vec{uv} is tense
if $\text{dist}(u) + w(u \rightarrow v) \leq \text{dist}(v)$



If $u \rightarrow v$ is tense:
path via u is better

so: $\text{pred}(v) = u$
 $\text{dist}(v) = \text{dist}(u) + w(u \rightarrow v)$

So, relax:

RELAX($u \rightarrow v$):

$\text{dist}(v) \leftarrow \text{dist}(u) + w(u \rightarrow v)$

$\text{pred}(v) \leftarrow u$

Algorithm:

Repeatedly find tense edges & relax them.

When none remain, the $\text{pred}(v)$ edges form the SSSP tree.

INITSSSP(s):

$\text{dist}(s) \leftarrow 0$

$\text{pred}(s) \leftarrow \text{NULL}$

for all vertices $v \neq s$

$\text{dist}(v) \leftarrow \infty$

$\text{pred}(v) \leftarrow \text{NULL}$

GENERICSSSP(s):

INITSSSP(s)

put s in the bag

while the bag is not empty

take u from the bag

for all edges $u \rightarrow v$

if $u \rightarrow v$ is tense

RELAX($u \rightarrow v$)

put v in the bag

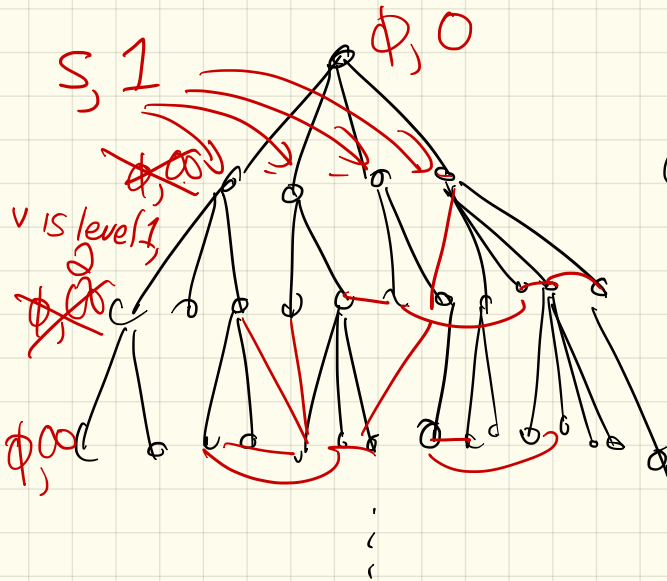
To do: which "bag"?
(+ proof)

"Easy" (??) warm-up:

What if unweighted?

→ use a queue

How does "tense" work?
(Hint: think BFS!)



Other edges?
edges?
(ie can anything be tense?)

What the heck is his forken??






Figure 8.6. A complete run of breadth-first search in a directed graph. Vertices are pulled from the queue in the order $s * b d * c a g * f e * h *$, where $*$ is the end-of-phase token. Bold vertices are in the queue at the end of each phase. Bold edges describe the evolving shortest path tree.

s's nbrs
other children
nbrs of s's children




2nd version

What if directed & acyclic?



Well, know something:
topological order!



So, use it!

DAGSSSP(s):


INITSSSP(s)

for all vertices v in topological order

for all edges $u \rightarrow v$

if $u \rightarrow v$ is tense

RELAX($u \rightarrow v$)



Dijkstra (59)

(actually Lexzorek et al '57,
Dantzig '58)

Make the bag a priority
queue:

Keep "explored" part of
the graph, S .

Initially, $S = \{s\} + \text{dist}(s) = 0$
(all others NULL $\rightarrow \infty$)


While $S \neq V$:

Select node $v \notin S$ with
one edge from S to v
with:

$$\min_{e=(u,v), u \in S} \text{dist}(u) + w(u \rightarrow v)$$

Add v to S , set $\text{dist}(v) + \text{pred}(v)$

Picture \rightarrow



Four phases of Dijkstra's algorithm run on a graph with no negative edges. At each phase, the shaded vertices are in the heap, and the bold vertex has just been scanned. The bold edges describe the evolving shortest path tree.

Correctness

Thm: Consider the set S at any point in the algorithm.

For each $u \in S$, the distance $\text{dist}(u)$ is the shortest path distance (so $\text{pred}(u)$ traces a shortest path).

pf: induction on $|S|$:

Base case:

IH: Spgs claim holds when
 $|S| = k-1$.

IS: Consider $|S| = k$:
algorithm is adding
some v to S \cup

Back to implementation +
run time:

For each $v \in S$, could check
each edge + compute
 $D[v] + w(e)$

runtime?

Better: a heap!

When v is added to S :

- look at v 's edges and either insert w with key $\text{dist}(v) + w(v \rightarrow w)$ or update w 's key, if $\text{dist}(v) + w(v \rightarrow w)$ beats current one

Runtime:

- at most m ChangeKey operators in heap \mathcal{D}
- at most n inserts/removes