

Algorithms - fall 2020

Shortest
paths (part 2)



Recap:

- HW: due in 1 week
- Next HW - possibly will be 2nd oral HW.
(Stay tuned)

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

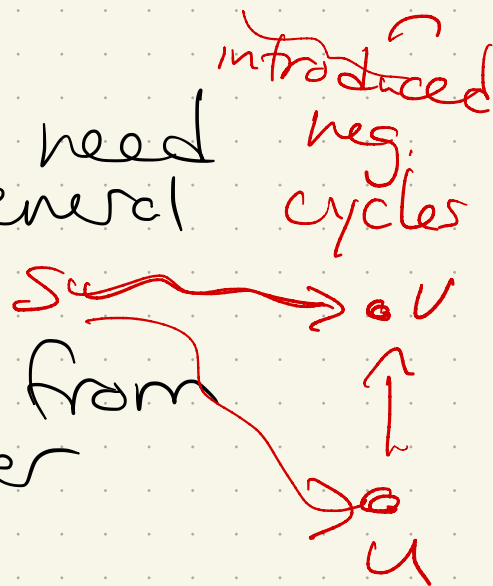
abstract
but
weighted



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.



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

"dist" is that best guess so far



If $u \rightarrow v$ is tense:

update w/ better path!

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

+ update previous node

So, relax:

RELAX($u \rightarrow v$):

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

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

take uv edge (plus u 's distance) as v 's new "guess"

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)



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 $\text{NULL} \rightarrow \infty$)


While $S \neq V$:

grow S
Select node $v \notin S$ with
one edge from S to v
with:
greedy! $\min \text{dist}(u) + w(u \rightarrow v)$
 $e = (u, v), u \in S$ "densest" edge

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.

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

$\text{RELAX}(u \rightarrow v)$

if v is in the priority queue

$\text{DECREASEKEY}(v, \text{dist}(v))$

else

INSERT($v, \text{dist}(v)$)

Figure 8.11. Dijkstra's algorithm.

PG: vertices, + best guesses (as priority)

insert remove Min get Min changekey

relax - found better path

sets s 's value to $(0, \emptyset)$ + all others to (∞, \emptyset)

at end, (v, dist) pairs, encode SP-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: if $|S| = 1$:

then $S = \{s\}$
(one vertex)

s has distance to itself!


so (S, ϕ) is correct



IH: Spps claim holds when $|S|=k$.

IS: Consider $|S|=k+1$:

algorithm is adding
some v to S , which
(by IH) we know has
only shortest paths



Claim: (b/c alg relaxed)
 \min is correct.

If no negative edges, no
other path can beat one in S .

Back to implementation +
run time:


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

runtime?

$\#$ reps

$O(VE)$

time to
search
all edge
lists



only
want
smallest

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 E ChangeKey operations in heap
- at most E inserts/removes in a heap

Total:

$$O(E \log V)$$

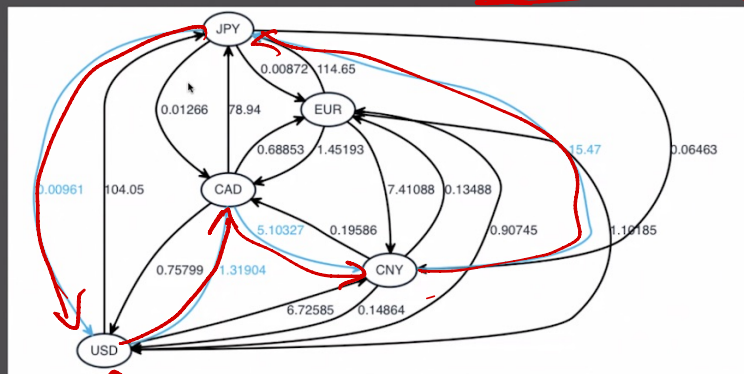
(no negative edges)

Negative edges!
Can happen!

Example!

2 weeks ago
in colloquium

Foreign Exchange Arbitrage



↑ dollars


Use Bellman-Ford on $-\log$ to find negative cycles

in finance
(not so much
in roads...)

What happens with negative edges? (in Dijkstra)

Well, for one thing, the induction breaks!

How bad is it?



→ If negative cycles:

completely fails → runs forever!

→ If negative edges

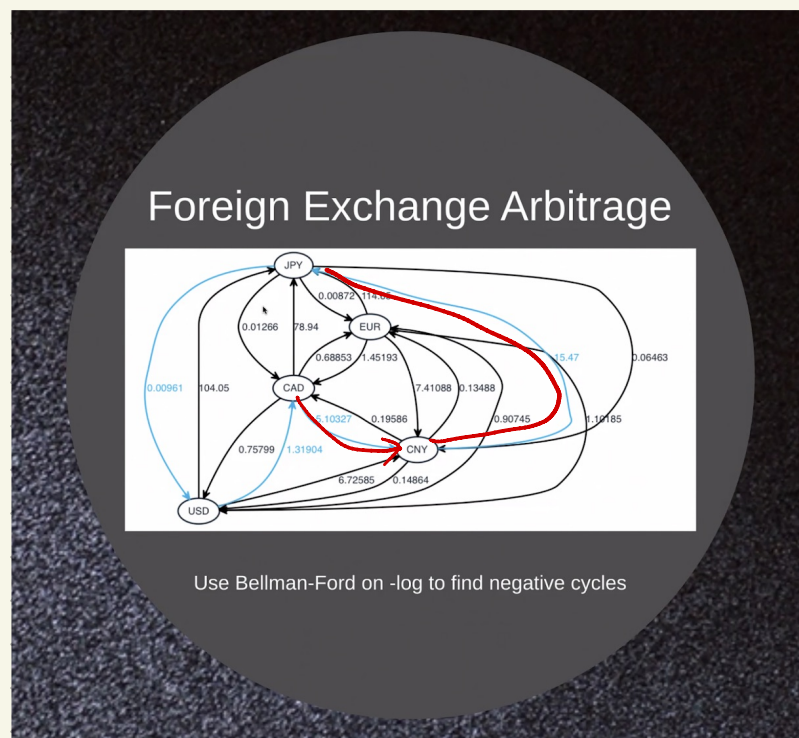
but no negative cycles:

exponential

but will stop eventually

What do we do with negative cycles?

Note: sometimes the entire goal is to find these!



Here, cycle = profit!




Well, can't be "greedy" & take the minimum each time.

Might repeat, so need to

- compute SP-tree
- or find a negative cycle

So, back to BFS style.

∞, ϕ



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($u \rightarrow v$)

while loop


all tense initially,
so relax them all.

Repeat - how long?? (problem)
& neg cycle, repeat forever

Result: how long can a path be? (shortest infinite loop)

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

$O(V)$
not $\infty \rightarrow V$ times
 $a \rightarrow a \rightarrow a \rightarrow a \rightarrow a$
relat all!
 $O(E)$
this is on cycle!
 $\leq V-1$



tense anywhere!



Runtime: $V-1$ rounds guarantees every vertex & edge have been "hit" \geq once $O(VE)$

Notation:

Let $\boxed{\text{dist}_{\leq i}(v)} :=$

length of the shortest S -to- v path using at most i edges

Ex:



S :

v_1 :

v_2 :

v_3 :

$$d_{\leq 0}(S) = 0$$

$$d_{\leq 0}(v_1) = \infty$$

$$d_{\leq 0}(v_2) = \infty$$

$$d_{\leq 0}(v_3) = \infty$$

$$d_{\leq 1}(S) = 1$$

$$d_{\leq 1}(v_1) = 10$$

$$d_{\leq 1}(v_2) = 1$$

$$d_{\leq 1}(v_3) = \infty$$

\vdots
all 0

$$d_{\leq 2}(v_1) = 4$$

$$d_{\leq 2}(v_2) = 1$$

$$d_{\leq 2}(v_3) = 6$$

$$d_{\leq 3}(v_1) = 4$$

$$d_{\leq 3}(v_2) = 1$$

$$d_{\leq 3}(v_3) = 6$$

=

So: $\text{dist}_{\leq 0}(s) = 0$ & for all $v \neq s$,
 $\text{dist}_{\leq 0}(v) = \infty$ BC - Init SSSP

Claim: $\forall v \neq i$, after i iterations of B-F,
 $\text{dist}(v) \leq \text{dist}_{\leq i}(v)$


Why?

Induction on i :

BC: $\text{dist}_{\leq 0}(s)$ & $\text{dist}_{\leq 0}(v)$ are good

IH: $\text{dist}_{\leq i-1}(v)$ is $\geq \text{dist}(v)$

IS: $\text{dist}_{\leq i}(v)$: Take shortest path of length i to v :



Either $u \rightarrow v$ is false:

$$\text{dist}(v) > \text{dist}_{\leq i-1}(u) + w(u \rightarrow v)$$

→ so relax:

better v guess

$$\text{dist}(v) \neq \text{dist}_{\leq i}(u) + w(u \rightarrow v)$$

OR: not

$$\text{dist}(v) < \underbrace{\text{dist}_{\leq i-1}(u) + w(u \rightarrow v)}$$

↑

worse,

so keep old
 $i-1$ path



Take away

Since any path has
length $\leq V-1$, don't need
to repeat more than that!

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!"

look
at every
edge



Runtime:

$O(VE)$

Why is B-F in practice slower?

Dijkstra $O(E \log V)$

B-F : $O(EV)$

cool data structure:
1 edge
looks at all edges

really 3 possibilities:

- G has positive weights
use Dijkstra

- G has negative edges


how many / how big?
→ use B-F (Dijkstra works but is slow)

- G has negative cycles.

→ use B-F

The rest: an (in practice) speed-up

Think of a BFS tree:



MOORE(s):

INITSSSP(s)

PUSH(s)

PUSH(★) *«start the first phase»*

while the queue contains at least one vertex

$u \leftarrow \text{PULL}()$

 if $u = \star$

PUSH(★) *«start the next phase»*

 else

 for all edges $u \rightarrow v$

 if $u \rightarrow v$ is tense

 RELAX($u \rightarrow v$)

 if v is not already in the queue

 PUSH(v)

queue:
S, I

more next time

Final version: Bellman's!

$$\text{dist}_{\leq i}(v) = \begin{cases} 0 & \text{if } i = 0 \text{ and } v = s \\ \infty & \text{if } i = 0 \text{ and } v \neq s \\ \min \left\{ \begin{array}{l} \text{dist}_{\leq i-1}(v) \\ \min_{u \rightarrow v} (\text{dist}_{\leq i-1}(u) + w(u \rightarrow v)) \end{array} \right\} & \text{otherwise} \end{cases}$$

Why??

Using i again as ~~#~~ of edges
in the path!

Since all paths are $\leq V-1$,

$\text{dist}_{V-1}(v)$ is $\text{dist}(v)$

(assuming no negative cycles)

Nicer:

BELLMANFORDDP(s)

```
dist[0,s] ← 0
for every vertex v ≠ s
  dist[0,v] ← ∞
for i ← 1 to V - 1
  for every vertex v
    dist[i,v] ← dist[i-1,v]
    for every edge u → v
      if dist[i,v] > dist[i-1,u] + w(u → v)
        dist[i,v] ← dist[i-1,u] + w(u → v)
```

Later observation:

Really don't need the i .

Just update those "tentative" distances, & trust it'll halt.

BELLMANFORDFINAL(s)

```
dist[s] ← 0
for every vertex v ≠ s
  dist[v] ← ∞
for i ← 1 to V - 1
  for every edge u → v
    if dist[v] > dist[u] + w(u → v)
      dist[v] ← dist[u] + w(u → v)
```

Same runtime as prior BF
(just a bit more confusing!)