

Shortest Paths in Near-Linear Time

(FOLS 22 paper of Bernstein, Nanorzakai and Wulff-Nielsen).

• Non negative weights: Dijkstra $O(m+n \log n)$
+ FibHeaps.

- Integer Weights
- Single-Source Shortest Paths
- Negative weights (but no neg. cycles, etc.)

• Negative weights: Bellman-Ford-Moore $O(mn)$.

• Better for integer weights of small magnitude (see our old notes for Andrew Goldberg's algo in $O(mn \log C)$)

↑ max neg weight in absolute value

• This paper: $O(m \text{ poly} \log(n))$!!!

randomized algo using ideas we've seen (~~but~~)

(but using them in interesting ways).

↑ extending; even ↑ but simple

Ingredients

- Dijkstra's Algorithm (~~but~~ but the k -fold version)
- Bellman-Ford-Moore (but run for smaller # of iterations)
- Low-diameter decompositions (but on directed graphs)
- Prices / Feasible potentials
- SSSP on DAGs in Linear time

Other refs: see Youtube talks by Danupon N. Aaron B.

← my lecture follows this talk fairly closely


Recall: prices / potentials φ

give new weights $w_\varphi(u, v) = \varphi(u) + w(u, v) - \varphi(v)$.

(i) w_φ has same shortest paths as w , just different lengths.

(ii) if $w_\varphi \geq 0$ then can use Dijkstra.

(iii) ~~find~~ one feasible φ obtained by computing SPs in graph \hat{G} and setting $\varphi(v) = \hat{d}(s', v)$




b/c. $\varphi(v) = \hat{d}(s', v) \leq \hat{d}(s', u) + w_{uv} = \varphi(u) + w_{uv}$

$\Rightarrow \varphi$ is feasible potential.

But may take too much time, since need SSSP on \hat{G} (chicken & egg).

So use a more gradual approach.



~~Now since weights are integers, $w_e \ge -c/2$ means $w_e \ge 0$ and we're all set in $O(\log c)$ rounds~~

See later for what to do here

ultimate goal: $w_e \ge -c$ $\xrightarrow{\text{transform to find } \varphi}$ $w_e^\varphi \ge -c/2$

current goal: $w_e \ge -2$ $\xrightarrow{\text{find } \varphi}$ $w_e^\varphi \ge -1$.

(will extend to the ultimate goal easily)

feasible prices $\downarrow \uparrow$ SSSP

"scaling"-based approach

I Bellman-Ford-More (K)
 $d(s) = 0 \quad d(u) = \infty \quad \forall u \neq s$
 for $i = 1$ to K

for $e \in \text{edges of } G, e = (u, v)$
 $d(v) \leftarrow \min \{ d(v), d(u) + w_{uv} \}$

Thm 1: if all shortest paths have at most K edges then BFM(K) correctly returns the shortest path labels (i.e. $d(u) = \text{dist}(s, u) \quad \forall u$).

But, of course, shortest paths may be longer than $\text{poly}(n)$!

II Dijkstra's Algorithm (K)

$d(s) = 0, \quad d(u) = \infty \quad \forall u \neq s$

for $i = 1$ to $K+1$

all vertices are unmarked.

pick v with smallest label $d(v)$, v is unmarked

mark v , "relax" all outgoing edges $d(u) \leftarrow \min \{ d(u), d(v) + w_{vu} \}$
 $\forall u \text{ st } (v \rightarrow u) \in E$

Thm 2: if all shortest paths have at most K negative weight edges then

(HW 0) Dijkstra(K) correctly returns SSSP labels, runs in time $\tilde{O}(mK)$

Thm 3: Let $\eta(v)$ be # of neg. weight edges on shortest path from s to v .

A variant of Dijkstra(K) correctly computes SSSP, runs in time

$$\tilde{O}\left(\sum_{v \in V} \eta(v)\right)$$

↑ subsumes Thm 2. (pf: maybe later)

III

FixDag(G):

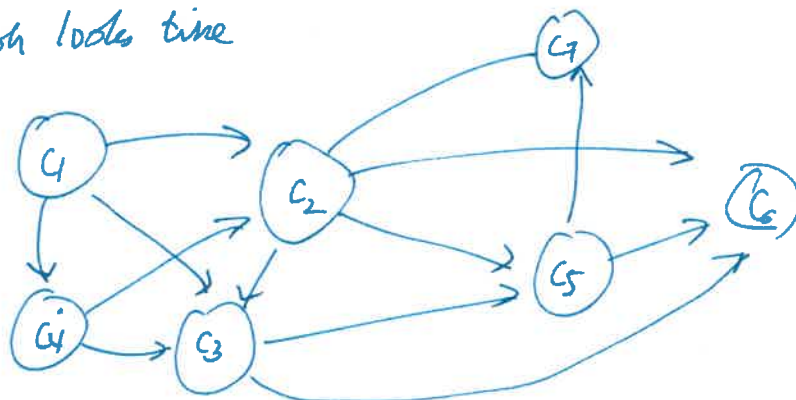
4

Recall the Strongly connected component decomposition of a Digraph G.

SCCs (C is an SCC if $\forall u, v \in C, \exists$ path from u to v
 & from v to u)


this gives an equivalence relation on V

So the graph looks like



The edges outside the SCCs form a DAG (directed acyclic graph)


Can put these in topological order so that edges only go from Left \rightarrow Right



can find this in linear time.

Now can make all reduced costs of these DAG-edges ≥ 0 .

How: set



$$\text{then } w_p(u, v) = (-iL) + w_{uv} - (-jL) = \underbrace{(j-i)L + w_{uv}}$$

for $u \in C_i$
 and $v \in C_j, j > i$

make L large enough
 to make this ≥ 0 .

Thm 4: Can "fix" all DAG edges of SCC in time $O(m)$.

Low diameter decompositions

Recall the definition from the last lecture. But of course we now have digraphs, and negative weights. Can't handle negative wts, but can get analog for digraphs with $w_e \geq 0$.

Thm: Given H with $l_e \geq 0$ (directed), and parameter $D \geq 0$.

Then \exists randomized algo that deletes edge set E_{del} s.t.

- \forall strongly connected components C in $H \setminus E_{del}$, $\text{diameter}_C \leq D$

$$\boxed{\max_{u, v \in C} \text{dist}_C(u, v) \leq D.}$$

- $\forall e, \Pr[e \in E_{del}] \leq \frac{l_e}{D} \cdot O(\log^2 n)$

\uparrow was $O(\log n)$ for undir graphs

Pf: (omit at first reading)

OK, we have most components. Let's see the plan.

(6)

Given G , ~~with~~ with $w(e) \geq -2$, want φ st. $w_\varphi(e) \geq -1$.

Plan [1st level]: define G^{+1} as graph with weights $w^{+1}(e) := w(e)+1$, G^{+2} has $w^{+2}(e) := w(e)+2 \forall e$.

• Find φ st. G_φ^{+1} has non negative lengths $\Rightarrow G_\varphi$ has $w_\varphi(e) \geq -1$.
 $w_\varphi^{+1}(e) \geq 0$.

How?

1. G^{+2} has $w^{+2}(e) \geq 0 \Rightarrow$ perform LDD with some param D .
to delete Edges.

2. Now on G^{+1}/Edges

(i) Find φ_1 to fix the SCC edges. (use that SCC have low diam).

(ii) Find φ_2 to fix the DAG edges (use DAG fixing).

3. now: on G^{+1}

• Find φ_3 to fix Edges (use that SPs have few negative edges due to LDD; then use Dijkstra⁺⁺).

\Rightarrow get φ st G_φ^{+1} has $w_\varphi^{+1}(e) \geq 0 \Rightarrow G_\varphi$ has $w_\varphi(e) \geq -1$.

This gives $O(m\sqrt{m})$ when $D = \sqrt{m}$.

Then (a) can recurse in step 2(i) instead.

(b) and choose better choice of $D \Rightarrow$ get $O(\text{poly}(\log m))$.

Then scale on distance to get $w(e) \geq -\Delta \Rightarrow w_\Delta(e) \geq -\Delta/2$.

Butter?

Details: (1) is easy form definition.


Step 2. For ~~any~~ graph G^+ , define $\hat{S} + G^+$
want to find SPs from \hat{S} to get prices.



(i) drop $E_{del} \cup E_{dag}$.

Find SPs for each SCC separately.

~~Want~~ Want to use BF (D rounds).



Alert! that was that ~~all~~ ^{all} SPs have few edges.
But we only have bounds on the weighted diameter !!

Main Lemma: sps C^+ has no negative cycles, and diameter $(C^+) \leq D$.

then $\hat{S} + C^+$ has shortest paths with $\leq D+1$ edges.

Pf.: sps. shortest path ~~has~~ has $\geq D+2$ edges.
from \hat{S} to v



\Rightarrow P (tail of path) has $\geq D+1$ edges.

note: $w^+(P) \leq 0$ b/c of construction of $\hat{S} + C^+$

so $w(P) \leq -(D+1)$ b/c $w^+ = w + 1$.

but $d(v, u) \leq D$ b/c diameter

\Rightarrow neg. cycle, contradiction 😊.


$\Rightarrow \hat{S} + C^+$ has few edges in each SP. ($\leq D + 1$)

\Rightarrow can just use BF ($D + 1$ rounds.)

time = $O(mD)$

#2(ii) DAG edges put back in, and fixed in $O(m)$ time.

#3: How many ~~edges~~ ~~neg~~ negative edges in E_{del}^+ need to be fixed? (in expectation?)



Note: $\begin{cases} e \in E_{del}^+ \text{ and } e \text{ negative} \Rightarrow w^{+2}(e) \in [0, 1]. \\ \text{b/c } w(e) \geq -2 \text{ and } w^{+1}(e) < 0 \end{cases}$

$$\Rightarrow Pr[e \in E_{del}] = \frac{w(e)}{D} \cdot O(\log^2 n) \leq \frac{O(\log^2 n)}{D}$$

$$\Rightarrow E[\text{# of } E_{del} \cap \text{negative in } w^{+1}] \leq \tilde{O}\left(\frac{m}{D}\right)$$

$$\Rightarrow E[\text{runtime of Dijkstra to fix all these}] \leq \tilde{O}(m) \times \tilde{O}\left(\frac{m}{D}\right)$$

\uparrow repeated Dijkstra!

total time: $O(mD + m + m^2/D) \Rightarrow$ set $D = \sqrt{m}$ to get $O(m^{3/2})$.

Let's do better:

SSSP (G, Δ)

// invariant: $(G$ has each SCC having diam $\leq \Delta$) and $w(e) \geq -2$

(1) LDD with $D = \Delta/2 - 1$, delete E_{del}

(2) (i) recurse on $(G \setminus E_{del}, \Delta/2)$

// invariant maintained,

get back $\forall e \in E \setminus E_{del} \Rightarrow w^{+1}(e) \geq 0$

(3) Fix E_{del}^+ .

But now: Each vertex v has $\eta(v) \leq \sum_{e \in sp} \frac{1}{(\Delta/2)} \cdot O(\log^2 n) \leq O(\log^2 n)$

$$\Rightarrow E \sum \eta(v) \leq \tilde{O}(m)$$

overall $\tilde{O}(m \log \Delta)$

BTW: DAG fix in in the base case $m=1$ if $E_{del}=0$.

Final Piece: Find φ st. $w(e) \geq -D \Rightarrow w_{\varphi}(e) \geq -D/2$