

Algorithms (2020)

Minimum
Spanning Trees



Recap

- HW on greedy - due Wed.

- Office hours today:

3pm (not 1pm)

(email if you ever need
a non-office hour time)


Minimum Spanning Trees

Goal: Given a weighted graph G , find a spanning tree of G , T , that minimizes:

$$w(T) = \sum_{e \in T} w(e)$$

edge in T

edge: $w(e) = \text{value in } T$



Motivation:

roads
networks
etc.

First:

Does it have to be a tree?

Yes! Why?

Spps not - then

\exists a cycle

goal: connect every pair of vertices


remove any edge - still connected

Second:

These are unique!

obviously ~~not~~ not

Ex:



← unweighted!

tree? Any spanning tree works — use DFS or BFS

Things will be cleaner if we have unique trees. So:

Lemma: Assuming all edge weights are distinct, then MST is unique.

Pf: By contradiction:

Suppose T & T' are both MSTs, with $T \neq T'$

↳ at least one edge different

• $T \cup T'$ contains a

cycle C


↳ every $u \leftrightarrow v$ connected in T, T' has another edge e, e' that

• That cycle must have 2 edges of equal weight path cycle

⇒ Contradiction!

MST would lose one of those edges.

could be more than 1 edge!



Now, what if weights aren't unique?

Just need a way to consistently break ties.

↑ if $w(e) = w(e')$, one will always win.

SHORTEREDGE(i, j, k, l)

if $w(i, j) < w(k, l)$

then return (i, j)

if $w(i, j) > w(k, l)$

then return (k, l)

if $\min(i, j) < \min(k, l)$

then return (i, j)

if $\min(i, j) > \min(k, l)$

then return (k, l)

if $\max(i, j) < \max(k, l)$

then return (i, j)


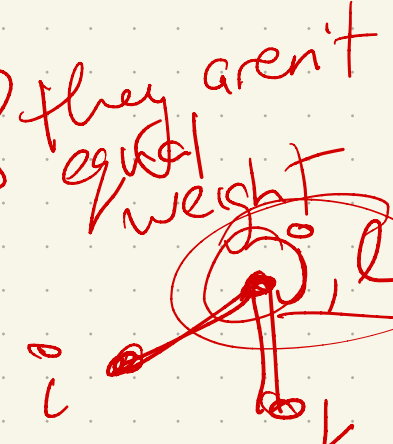
if $\max(i, j) > \max(k, l)$

then return (k, l)

if $\max(i, j) > \max(k, l)$

return (k, l)

break ties.



assuming vertices are stored in some array

$V: [0, \dots, n]$
adj list.

So, takeaway:

Can assume unique MST.

Next: an algorithm.

The magic truth of MSTs:

You can be SUPER greedy.

Almost any natural idea
will work!

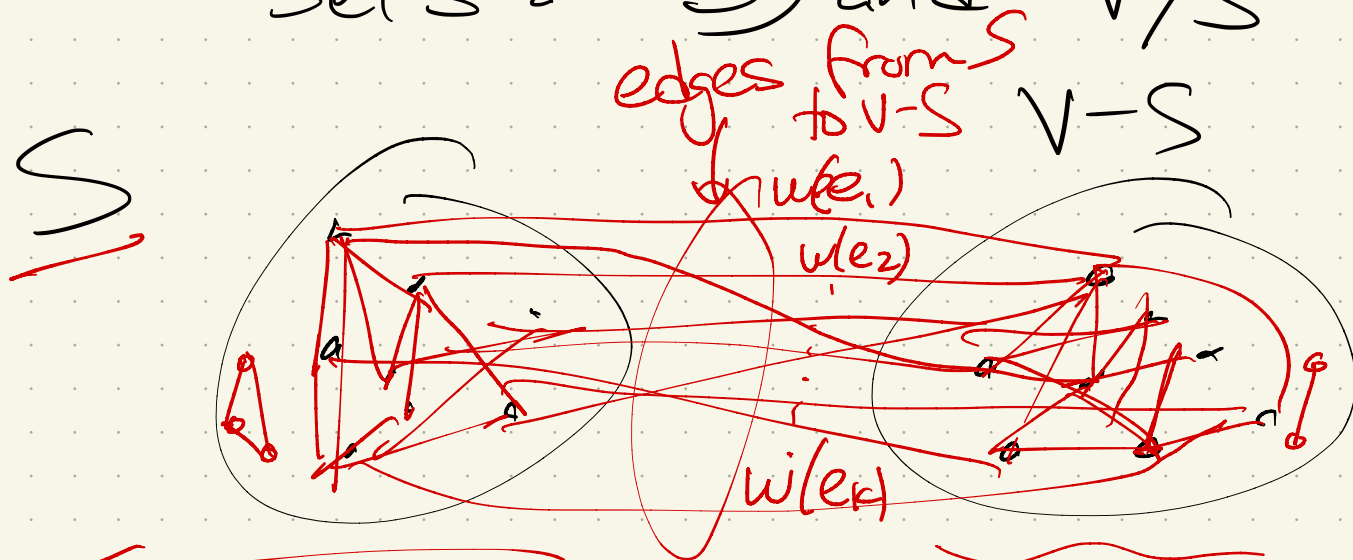
This is highly unusual, &
there's a reason for it:

→ these are a (rare) example
of something called a
matroid.

(Way beyond this class...)


Key property:

Consider breaking G into two sets: S and V/S




The MST will always contain the lowest edge connecting the two sides.

No matter what!



Proof: consider minimum edge e



Suppose MST does not contain e .

Then MST must have some other S to $V-S$ edge

Take $\underline{T} \cup \underline{e}^{uv}$: a tree + an edge has at least one cycle. why? since $e = uv$ & $u \rightarrow v$ have a path in T . Take cycle. & remove a different S to $V-S$ edge. $w(e') > w(e)$, since e was new tree, $< T$. smallest.

Generic Algorithm: → apply that lemma w/ $S \cup V-S$

Build a forest: an acyclic subgraph.


↑ less than a tree, no cycles.

Dfn: An edge is useless

if it connects 2 endpts in of F same component

also edges that are neither.

An edge is safe if it is minimum edge from some component of F to another.



So idea:

Add safe edges
until you get a tree


If everything isn't connected,
must have some safe edge.

not a tree

Why?

Pick a component of F .
Make it S + rest $V-S$,
+ apply that lemma.

Add it + recurse.



We'll see 3 ways:

① Find all safe edges.
Add them & recurse.

Boruvka


② Keep a single connected component

At each iteration, add
1 safe edge.


③ Sort edges & loop through them.

If edge is safe,
add it.

all 3 use this $S, V-S$
lemma



First one: (1926-ish)



So we need to:

loop: While more than 1 component: ↗ not a tree

⇒ Track components

— Find all safe edges

— Add them

More formally:

BORUVKA(V, E):

$F = (V, \emptyset)$

$\text{count} \leftarrow \text{COUNTANDLABEL}(F)$


while $\text{count} > 1$

$\text{ADDALLSAFEEDGES}(E, F, \text{count})$

$\text{count} \leftarrow \text{COUNTANDLABEL}(F)$

return F

Graph



Min edge

Safe:

uv	wv	uv	
1	2	3	4

each comp gets 1 safe edge

ADDALLSAFEEDGES(E, F, count):

for $i \leftarrow 1$ to count

$\text{safe}[i] \leftarrow \text{NULL}$

for each edge $uv \in E$

 if $\text{comp}(u) \neq \text{comp}(v)$

 if $\text{safe}[\text{comp}(u)] = \text{NULL}$ or $w(uv) < w(\text{safe}[\text{comp}(u)])$

$\text{safe}[\text{comp}(u)] \leftarrow uv$

 if $\text{safe}[\text{comp}(v)] = \text{NULL}$ or $w(uv) < w(\text{safe}[\text{comp}(v)])$

$\text{safe}[\text{comp}(v)] \leftarrow uv$

for $i \leftarrow 1$ to count

 add $\text{safe}[i]$ to F

Uses WFS-variant from Ch 5:

COUNTANDLABEL(G):

$\text{count} \leftarrow 0$

for all vertices v

 unmark v

for all vertices v

 if v is unmarked

$\text{count} \leftarrow \text{count} + 1$

$\text{LABELONE}(v, \text{count})$

return count

Label one component

LABELONE(v, count):

while the bag is not empty

 take v from the bag

 if v is unmarked

 mark v

$\text{comp}(v) \leftarrow \text{count}$

 for each edge vw

 put w into the bag

Correctness:

- MST must have any safe edge
- We keep computing safe edges & adding
- Stop when # connected components = 1

⇒ Have the MST!


Run time:

A bit trickier!

Depends on how many
safe edges we get.

Claim: There are at
least $\frac{\text{\#components}}{2}$
safe edges each time.

Why?



So: runtime:

ADDALLSAFEEDGES(E, F, count):

for $i \leftarrow 1$ to count

$\text{safe}[i] \leftarrow \text{NULL}$

for each edge $uv \in E$

 if $\text{comp}(u) \neq \text{comp}(v)$

 if $\text{safe}[\text{comp}(u)] = \text{NULL}$ or $w(uv) < w(\text{safe}[\text{comp}(u)])$

$\text{safe}[\text{comp}(u)] \leftarrow uv$

 if $\text{safe}[\text{comp}(v)] = \text{NULL}$ or $w(uv) < w(\text{safe}[\text{comp}(v)])$

$\text{safe}[\text{comp}(v)] \leftarrow uv$

for $i \leftarrow 1$ to count

 add $\text{safe}[i]$ to F

↑ Looks at each vertex & edge
in worst case:

BORŮVKA(V, E):

$F = (V, \emptyset)$

$\text{count} \leftarrow \text{COUNTANDLABEL}(F)$

 while $\text{count} > 1$

 ADDALLSAFEEDGES(E, F, count)

$\text{count} \leftarrow \text{COUNTANDLABEL}(F)$

 return F

BFS/DFS
on tree:

How many
iterations?