

Approximation Algorithms

CSE 417 Winter 21
Lecture 26

Where Are We?

What do you do if you want to solve a problem that is *NP*-hard?

You shouldn't expect to design an algorithm that runs in polynomial time and always gives you the right answer.

Last time: Algorithms that run in (not-as-bad-as-it-could-be) exponential time and give you the right answer.

Today: Algorithms that run in polynomial time, but don't always give you the right answer.

Optimization Problems

Putting away decision problems, we're now interested in optimization problems.

Problems where we're looking for the "biggest" or "smallest" or "maximum" or "minimum" or some other "best"

Vertex Cover (Optimization Version)

Given a graph G find the **smallest** set of vertices such that every edge has at least one endpoints in the set.

Much more like the problems we're used to!

What does NP-hardness say?

NP-hardness says:

We can't tell (given G and k) if there is a vertex cover of size k .

And therefore, we can't find the minimum one (write the reduction! It's good practice. Hint: binary search over possible values of k).

It doesn't say (without thinking more at least) that we couldn't design an algorithm that gives you an ~~independent set~~ ^{vertex cover} that's only a tiny bit worse than the optimal one. Only 1% worse, for example.

How do we measure worse-ness?

Approximation Ratio

For a minimization problem (find the shortest/smallest/least/etc.)

If $OPT(G)$ is the value of the best solution for G , and $ALG(G)$ is the value that your algorithm finds, then ALG is an α approximation algorithm if for every G ,

$$\alpha \cdot OPT(G) \geq ALG(G)$$

i.e. you're within an α factor of the real best.

Finding an approximation for Vertex Cover

Take the idea from the clever exponential time algorithm.

But instead of checking which of u, v a good idea to add, just add them both!

```
While (G still has edges)
    Choose any edge (u,v)
    Add u to VC, and v to VC
    Delete u v and any edges touching them
EndWhile
```

Does it work?

Do we find a vertex cover?

Is it close to the smallest one?

But first, let's notice – we're back to polynomial time algorithms!

If we're going to take exponential time, we can get the exact answer. We want something fast if we're going to settle for a worse answer.

Do we find a vertex cover?

When we delete an edge, it is covered (because we added both u and v). And we only stop the algorithm when every edge has been deleted. So every edge is covered (i.e. we really have a vertex cover).

How big is it?

Let OPT be a minimum vertex cover.

Key idea: when we add u and v to our vertex cover (in the same step), at least one of u or v is in OPT .

Why? (u, v) was an edge! OPT covers (u, v) so at least one is in OPT .

So how big is our vertex cover? At most twice as big!

This is a 2-approximation for vertex cover!

Another Approximation Algorithm

Let's look at another approximation algorithm for vertex cover.

Remember the linear program for vertex cover?

Vertex Cover LP

Minimize $\sum w(u) \cdot x_u$

Subject to:

$x_u + x_v \geq 1$ for all $(u, v) \in E$

$0 \leq x_u \leq 1$ for all u .

Don't worry about the weights for today.

We got an exact solution for bipartite graphs....

What do we do


Increase x for the purple vertices, and decrease x for the gold vertices.

(at the same time at the same rate)

Every edge (in our example) has a purple and gold endpoint, so every constraint is still satisfied.

The objective (in our example) increases and decreases at the same rate.

So we still have an optimal (minimum) vertex cover



In General...

2-color the graph (call the vertices "purple" or "gold")

Increase all the purple vertices by some value δ

And decrease all the gold vertices by the same value δ

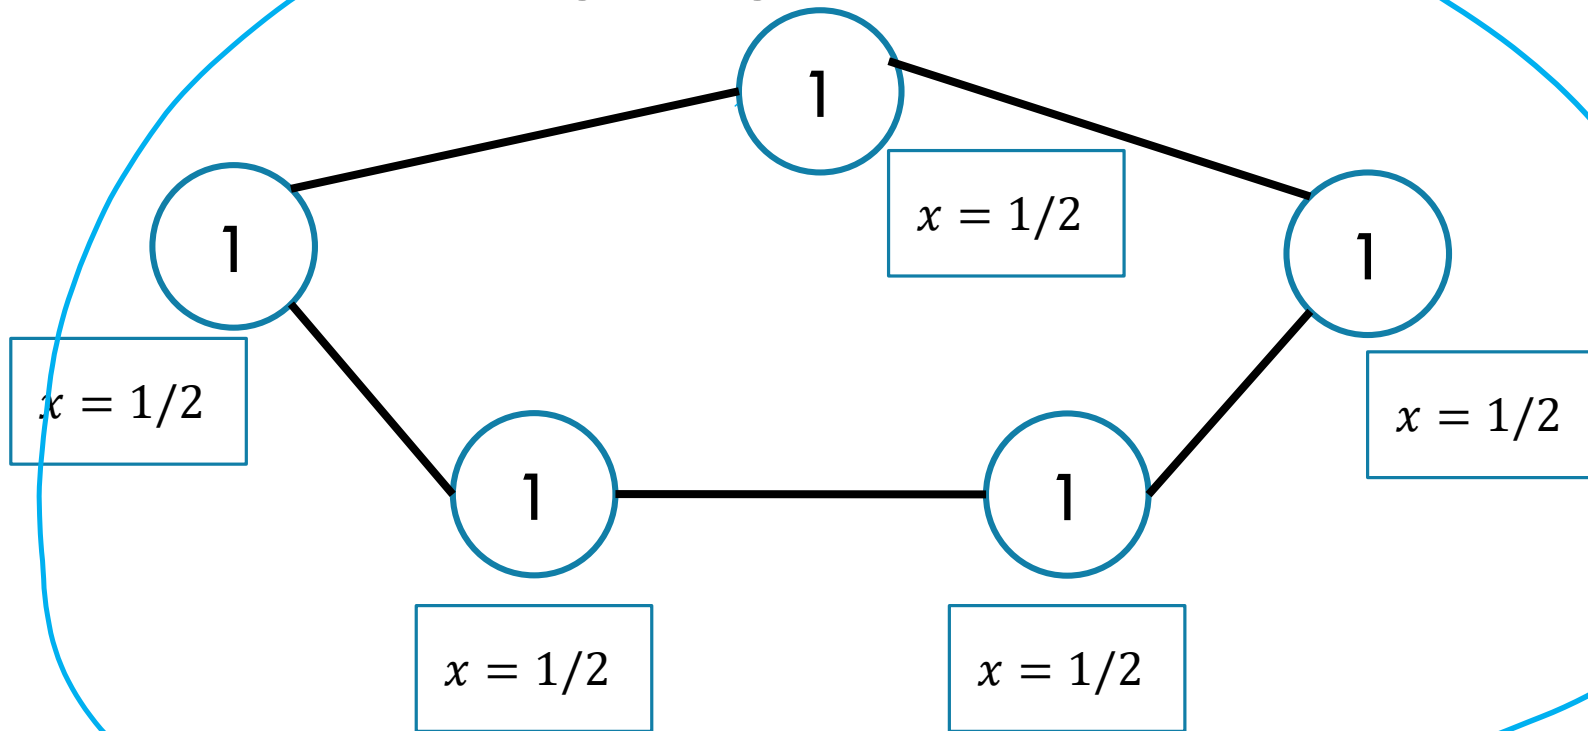
Choose δ so that we set at least one variable to 0 or 1 (but don't move any variables outside the $[0,1]$ range allowed).

Those vertices that just got set to 0 or 1 can be deleted. Start over with the remaining graph.

Non-Bipartite

We needed the graph to be bipartite to be able to 2-color it.

What if our original graph isn't bipartite?



The LP finds a fractional vertex cover of weight 2.5

There's no "real"/integral VC of weight 2.5. – lightest is weight 3.

There's a "gap" between integral and fractional solutions.

So, what if the graph isn't bipartite?

Big idea:

Just round!

If $x_u \geq \frac{1}{2}$, round up to 1.

If $x_u < \frac{1}{2}$, round down to 0

Two questions – is it a vertex cover? How far are we from the true minimum?

Fill out the poll everywhere for
Activity Credit!

Go to pollev.com/cse417 and login
with your UW identity

Minimize $\sum w(u) \cdot x_u$

Subject to:

$x_u + x_v \geq 1$ for all $(u, v) \in E$

$0 \leq x_u \leq 1$ for all u .

Is it a vertex cover?

Every edge was covered in the fractional matching

i.e. for every edge (u, v)

$$\underline{x_u + x_v \geq 1.}$$

At least one of those is getting rounded up!

So every edge is covered.

(And we've rounded to integers, so we have a "real" vertex cover.)

How good of an approximation is it?

Well, we might have doubled the value of the LP when we rounded. But we definitely didn't do any more than that.

$$2 \cdot LP \geq ALG$$

And the value of the LP is definitely not bigger than the true size of the vertex cover (because otherwise the LP would have found that).

$$OPT \geq LP$$

Combining:

$$2 \cdot OPT \geq ALG$$

So we're safe in calling this a 2-approximation.

Comparing to the LP value

We did a weird thing on that last slide.

We were supposed to compare the value of our vertex cover to the best vertex cover.

But instead we compared it to the value of the LP...which we know isn't always the value of the vertex cover!

That wasn't laziness, it's a very common technique. We know very little about the true value of the vertex cover (if we knew what it looked like VERY VERY precisely, why couldn't we just write an algorithm to find it? We actually won't know much). So we start with what the algorithm gave us (that we do understand).

Side Note

Could we do better?


Not just with the LP.

If you take a graph with n vertices and every possible edge, the LP's minimum is $n/2$, the true minimum vertex cover is size $n - 1$.

The ratio is $2 - 1/n$. So if we don't at least double the value **sometimes** we won't get a vertex cover at all.

Getting a 1.99999999 approximation is an open problem!

Another Algorithm



Lets try to approximate Travelling Salesperson.

Traveling Salesperson

Given a weighted graph, find a tour (a walk that visits every vertex and returns to its start) of weight at most k .

Some assumptions:

1. The graph is undirected.

2. The graph is complete (every edge is there) – the edges might represent series of roads rather than individual streets. Weight is how much gas you need to travel.


3. The weights satisfy the “triangle inequality” (it’s faster to go from x to y directly than it is to go from x to y through z).

TSP starting point

What would be a good baseline?

Something we **can** calculate that would at least connect things up for us.

A Minimum Spanning Tree!



From MST to Tour

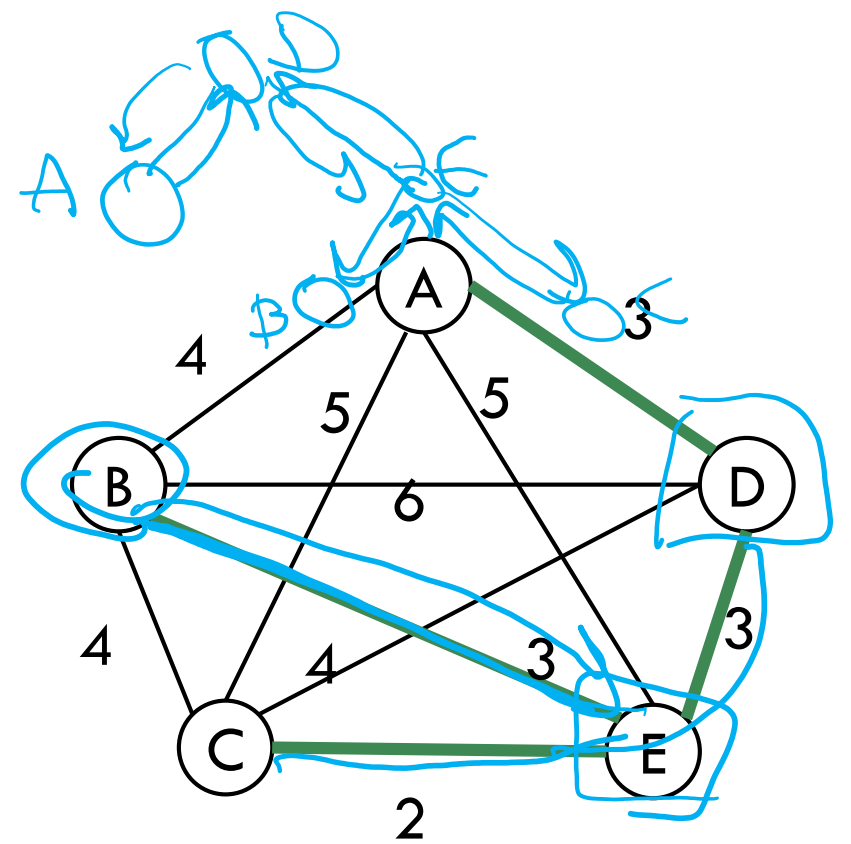
How do we get from start to every vertex and back?

Make the starting point the root, do a traversal (DFS) of the graph!



Why not BFS? Because the "next vertex" isn't always right next to you! (not a problem in this example, but very bad if you have a very tall tree)

How much gas do we use in DFS? We use each edge twice



If D is the starting point:
Go from D to A , back to D
To E Down to B back to E to C
Back to E back to D .


Doing a Little Better

Using each edge twice is potentially a little wasteful. Can we do better?

The biggest problem is vertices of odd degree. The last time we enter that vertex, the only way out is an already used edge.

And that's definitely not taking us somewhere new!

So lets add some possible ways out.



What would help?



A matching would help! (i.e. a set of edges that don't share endpoints)

Specifically a minimum weight matching.

You can find one of those efficiently (just trust me)

Add those edges in (if they're already in the MST, make an extra copy)

So we now have the MST AND the minimum weight matching on the odd edges.



Did It Help?



So...now every vertex has even degree...but there's not a nice order anymore.

We'll have to find one.

Start from the starting point, and just follow any unused edge!

Because every vertex has even degree, (except for the starting vertex) when you go in, you can come out! So you can't "get stuck"

What if you get back to the start and end up with unused edges? Find a visited vertex one is adjacent to and "splice in" the cycles.



D,A,B,E,D is found first. E,C,E found next. After splicing: D,A,B,E,C,E,D. is the final tour

Is it a good approximation algorithm?

We will visit every vertex at least once!

Every vertex had degree at least one (because we started with an MST!)

So by the end of the process, we had degree at least two on every vertex.

And we go back and use all the edges we selected. So we visit every vertex, and we start and end at the same place.

Is it a good approximation algorithm?

What does our algorithm produce?

At most $\frac{3}{2} OPT$ (at most 1.5 times the weight of the optimal tour)

Why? We use every edge once, that's one *MST* plus the weight of the matching.

How much is the *MST*? Less than *OPT*. (*OPT* has a spanning tree inside it!)

How much is the matching? Less than $\frac{1}{2} OPT$. (*OPT* is less than a tour on the odd vertices, and a tour on the odd vertices is made up of two matchings)

Approximating TSP

We found a $\frac{3}{2}$ -approximation for TSP!

The algorithm is called "Christofides Algorithm"

It's almost 50 years old.

The best approximation is $\frac{3}{2} - \epsilon$ where $\epsilon > 10^{-36}$

Developed by three researchers at UW **this year**.

<https://arxiv.org/pdf/2007.01409.pdf>

Summary

Coping with NP-hardness.

1. Understand your problem really well (make sure you're not solving an easy special case).
2. Prove the problem really is NP-hard.
3. Try a band-aid (SAT library, Integer programming library, etc.)
4. Try to find a good-enough exponential time algorithm or an approximation algorithm.