

More Coping with NP- hardness

CSE 417 Winter 24
Lecture 24

Announcements

Don't be scared when you open up HW8

It's longer, and we're letting you submit more problems, but the expectation when I'm making grades is the same as the other homeworks.

You can do more if you want to review and/or catch-up, but you don't have to.

The last resubmission cycle is due Monday Mar 11 (i.e., right after HW8)

That's the Monday of finals week.

Part of the reason for having so many problems this week is to give you room to do an "extra" one if you're really close to a grade-guarantee and want some insurance on getting a particular score.

Where Are We?

Any NP-hard problem has the property that if you find a polynomial time algorithm for it, you find a polynomial time algorithm for all problems in NP.

There are thousands of NP-hard problems! We don't think any of them have polynomial time algorithms (but we haven't proven it).

To show your problem B is NP-hard, reduce a known NP-hard problem A to your problem B .

Direction matters!

Examples

There are literally thousands of NP-complete problems.
And some of them look weirdly similar to problems we do know efficient algorithms for.

In P

Short Path

Given a directed graph, report if there is a path from s to t of length at most k .

NP-Complete

Long Path

Given a directed graph, report if there is a path from s to t of length at least k .

Examples

In P

Light Spanning Tree

Given a weighted graph, find a spanning tree (a set of edges that connect all vertices) of weight at most k .

NP-Complete

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 .

The electric company just needs a greedy algorithm to lay its wires.
Amazon doesn't know a way to optimally route its delivery trucks.

Examples

In P

2-Coloring

Given an undirected graph, can the vertices be labeled red and blue with no edge having the same colors on both endpoints?

NP-Complete

3-Coloring

Given an undirected graph, can the vertices be labeled red, blue, and green with no edge having the same colors on both endpoints?

Just changing a number by one takes us from one of the first problems we solved (and one of the fastest algorithms we've seen) to something we don't know how to solve efficiently at all.

Dealing with NP-hardness


Thousands of times someone has wanted to find an efficient algorithm for a problem...

...only to realize that the problem was NP-hard.

Takeaway 1:

Even though we haven't proven $P \neq NP$ (i.e. we haven't proven any of these problems **don't** have an efficient algorithm), this is good evidence that we shouldn't be trying to solve *NP*-hard problems.

It's probably not just a matter of finding the "right representation"/"right angle on the problem" we've tried a few thousand of them.



Why Should you care?

Why should you care about P vs. NP

Most computer scientists are convinced that $P \neq NP$.

A survey of experts (PhDs in CS) found 98% of them thought $P \neq NP$.

And the median guess was that we're at least 50 years from getting the answer.

Why should you care about this problem?

It's your chance for:

\$1,000,000. The Clay Mathematics Institute will give \$1,000,000 to whoever solves P vs. NP (or any of the 5 remaining mathematical conjectures they listed)

To get a Turing Award

Why should you care about P vs. NP

Most computer scientists are convinced that $P \neq NP$.

A survey of experts (PhDs in CS) found 98% of them thought $P \neq NP$.

And the median guess was that we're at least 50 years from getting the answer.

Why should you care about this problem?

It's your chance for:

\$1,000,000. The Clay Mathematics Institute will give \$1,000,000 to whoever solves P vs. NP (or any of the 5 remaining mathematical conjectures they listed)

To get a ~~Turing Award~~ the Turing Award renamed after you.

Why Should You Care if $P=NP$?

Suppose $P=NP$.

Specifically that we found a genuinely in-practice efficient algorithm for an NP-complete problem. What would you do?

\$1,000,000 from the Clay Math Institute obviously, but what's next?

Why Should You Care if $P=NP$?

We found a genuinely in-practice efficient algorithm for an NP-complete problem. What would you do?

Another \$5,000,000 from the Clay Math Institute

Put mathematicians out of work?

Decrypt (essentially) all current internet communication. See the real world question.

A world where $P=NP$ is a very very different place from the world we live in now.

Why Should You Care if $P \neq NP$?

We already expect $P \neq NP$. Why should you care when we finally prove it?

$P \neq NP$ says something fundamental about the universe.

For some questions there is not a clever way to find the right answer
Even though you'll know it when you see it.

There is actually a way to obscure information, so it cannot be found quickly no matter how clever you are.

Why Should You Care if $P \neq NP$?


To prove $P \neq NP$ we need to better understand the differences between problems.

Why do some problems allow easy solutions and others don't?

What is the structure of these problems?

We don't care about P vs NP just because it has a huge effect about what the world looks like.

We will learn a lot about computation along the way.



Dealing With NP-hardness

Dealing with NP-hardness

Thousands of times someone has wanted to find an efficient algorithm for a problem...

...only to realize that the problem was NP-hard.

Takeaway 2:

Sooner or later it will happen to one of you.

What do you do if you think your problem is NP-complete?

Dealing with NP-completeness

You just started your new job at Amazon. Your boss asks you to look into the following problem

You have a graph, each vertex is where a specific truck has to do a delivery. Starting from the warehouse, how do you make all the deliveries and return to the warehouse using the minimum amount of gas.

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 .

Step 1: Make sure your problem is really *NP*-hard

Understand **exactly** what your inputs and outputs are.

2-coloring and 3-coloring are very different.

Finding a vertex cover of a general graph is *NP*-hard. Finding a vertex cover of a bipartite graph can be done in multiple very efficient ways.

Understand **exactly** what you're being asked to solve.

Realizing you're trying to solve a problem on a tree instead of a general graph almost always makes DP possible.

Are your constraints linear (can you use an LP)?

Are your constraints simple (are you solving 2SAT instead of 3SAT)?

Step 2: It still looks hard

Now that you know exactly what you're trying to solve, and you still can't solve it...

Next try to prove hardness (i.e. do a reduction).

Usually there's a similar problem you can convert from!

It's easier to do a reduction from 3-coloring to 5-coloring than from Hamiltonian Path to 5-coloring.

Both reductions **exist** but there's no need to flex here, look up a list of NP-complete problems and see what's similar looking.

Step 3: ???

So you go to your boss and say

“Sorry, problem’s NP-hard. I proved it.”

And your boss says:

“that’s a cool proof and all, but really. We need to tell the drivers where to go tomorrow...and we need to use less gas.”

Step 3: Band-aids

Can you write your problem as a *SAT* instance?

Ok, you definitely can if it's in *NP*, that's what *NP*-hardness means...can you write it as a reasonably-sized SAT instance (n^3 instead of $1000000n^{100}$)?

There are SAT libraries that **often** run pretty fast. In the worst-case they're still exponential, but you don't always hit the worst case!

Can you write your problem as an integer program?

Run an integer programming library and see what happens!

Can you write your problem as a graph problem?

Many are very well studied for "simple" graphs (e.g. "planar" graphs, ones that can be drawn on a piece of paper without edges crossing).

Step 4 – Permanent Solutions

But what if you need a guarantee?

Your code is running every night, and you need an answer by 6 AM or the delivery trucks don't go out.

Step 4 – Permanent Solutions


Two good options:

Exponential algorithms that aren't-as-slow-as-others

Give you an exact answer; won't take polynomial time but will be guaranteed take you less time than brute force.

Approximation algorithms

Don't give you the best answer, but guarantees a reasonable amount of time, and a guaranteed-pretty-good-answer.



Not-Terrible Exponential Time

Polynomial vs. Exponential

If you have an algorithm that takes exactly $f(n)$ microseconds, how large of an n can you handle in the given time?

	1 second	1 minute	1 hour	1 day	1 month	1 year	1 century
$\lg n$	2^{10^6}	$2^{6 \cdot 10^7}$	$2^{36 \cdot 10^8}$	$2^{864 \cdot 10^8}$	$2^{25920 \cdot 10^8}$	$2^{315360 \cdot 10^8}$	$2^{31556736 \cdot 10^8}$
\sqrt{n}	10^{12}	$36 \cdot 10^{14}$	$1296 \cdot 10^{16}$	$746496 \cdot 10^{16}$	$6718464 \cdot 10^{18}$	$994519296 \cdot 10^{18}$	$995827586973696 \cdot 10^{16}$
n	10^6	$6 \cdot 10^7$	$36 \cdot 10^8$	$864 \cdot 10^8$	$2592 \cdot 10^9$	$31536 \cdot 10^9$	$31556736 \cdot 10^8$
$n \lg n$	62746	2801417	133378058	2755147513	71870856404	797633893349	68654697441062
n^2	1000	7745	60000	293938	1609968	5615692	56175382
n^3	100	391	1532	4420	13736	31593	146677
2^n	19	25	31	36	41	44	51
$n!$	9	11	12	13	15	16	17

Vertex Cover

Input: Graph G , integer k

Output: Is there a set of at most k vertices such that every edge has at least one endpoint in the set?

The problem is NP-complete.

In the worst-case, we need exponential time.

For every subset S of vertices

 Check if every edge has at least one endpoint
in the set

Time? $O(2^n(n + m))$

Vertex Cover

We can do better (sometimes)!

Don't check every subset, just the biggest allowed subsets. How does the running time change as k changes?

When k is a constant (say, $k \leq 3$)

How many subsets are there of size 3? n^3 , running time: $O(n^3(n + m))$

That's not too terrible!

Vertex Cover

When k is a constant (say, $k \leq 3$)

Running time $O(n^3(n + m))$

k is a little bigger (say, $k = \log_2 n$)

Running time $O(n^{\log n}(n + m))$ not polynomial anymore

Worst value of k ($k = n/2$)

Running time $O\left(2^{\frac{n}{2}}(n + m)\right)$ VERY SLOW

k very very big ($k = n - 3$)

Running time $O(n^3(n + m))$ (not many very large vertex sets)

We can do better

When k is big, not much we can do. What about when it's small?

Our running time depends on k anyway, let's focus in on making our algorithm better when k is small.

Key idea: pick an edge (u, v)

There is a vertex cover of size k if and only if

There is a vertex cover of size $k - 1$ in $G - u$ or $G - v$.

i.e. at least one of u, v in the minimum vertex cover.

Key Idea – Let's Prove it!

If there is a vertex cover of size k , then there is a vertex cover of size $k - 1$ in $G - u$ or $G - v$.

Every vertex cover has to cover (u, v) . So at least one of u or v is included. Delete that vertex (one arbitrarily if both are in the vertex cover) and all edges that touch it. Every other edge was covered by another vertex (since we deleted all the edges touching the deleted vertex). What remains is a vertex cover of size $k - 1$ on $G - u$ or $G - v$.

Key Idea – Let's Prove it!

If there is a vertex cover of size $k - 1$ in $G - u$ or $G - v$ then there is a vertex cover of size k in G .

Assume that the vertex cover of size $k - 1$ is in $G - u$ (the argument is the same if it's in $G - v$ instead). Take the vertex cover of $G - u$ and add in u . Every edge of $G - u$ is covered by the vertex cover. The only other edges in G touch u , so u covers them.

Algorithm

```
VertexCover(graph G, int k)
    if(G has no edges) //we've covered them all!
        return true
    if(k < 0)
        return false
    H1 = copy of G
    H2 = copy of G
    pick any edge (u,v)
    H1 = H1.remove(u) //removes u and all edges with u as
an endpoint
    H2 = H2.remove(v) //removes v and all edges with v as
an endpoint
    return VertexCover(H1, k-1) || VertexCover(H2, k-1)
```

Running Time

$$\text{Recurrence: } T(k) = \begin{cases} 2T(k-1) + O(n+m) & \text{if } k \geq 1 \\ O(1) & \text{if } k < 0 \end{cases}$$

Running time? Unroll or use recursion tree

Running Time

$$\text{Recurrence: } T(k) = \begin{cases} 2T(k-1) + O(n+m) & \text{if } k \geq 1 \\ O(1) & \text{if } k < 0 \end{cases}$$

Running time? Unroll or use recursion tree

$$O\left((n+m) \cdot 2^k\right)$$

Vertex Cover

When k is a constant (say, $k \leq 3$)

Running time $O(2^3(n + m)) = O(n + m)$

k is a little bigger (say, $k = \log_2 n$)

Running time $O(2^{\log_2 n}(n + m)) = O(n(n + m))$ still polynomial!

$k = n/2$

Running time $O(2^{n/2}(n + m))$ very slow

k very very big ($k = n - 3$)

Running time $O(2^{n-3}(n + m))$ very very slow

Comparison

Sample values of k	Brute Force	Recurse by edge
3	$O(n^3(n + m))$	$O(n + m)$
$\log n$	$O(n^{\log n}(n + m))$	$O(n(n + m))$
$n/2$	$O\left(2^{\frac{n}{2}}(n + m)\right)$	$O\left(2^{\frac{n}{2}}(n + m)\right)$
$n - 3$	$O(n^3(n + m))$	$O(2^{n-3}(n + m))$

Takeaway

If your vertex cover is small you can get a pretty efficient algorithm.
For k at most $O(\log n)$ it even becomes polynomial.

A “simple case” you can carve off.

More Generally


Measuring the complexity in terms of something other than the size of the input is called “parameterized complexity”

Common parameters:

The answer (like Ford-Fulkerson! And vertex cover)

How “complicated” the input is (e.g. for graphs, do you have a tree, something very close to a tree, or nothing like a tree).

Another example: SAT – an instance with few variables and many constraints is very different from an instance with many variables and few constraints.



Approximation Algorithms

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 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?

[Pollev.com/robbie](https://pollev.com/robbie)

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)

$$x_u + x_v \geq 1.$$

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.
2. 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 x).

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 $\epsilon \approx 10^{-36}$

Developed by three researchers at UW **last 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.