

23.1 Introduction

We spent the last two lectures proving that for certain problems, we can't expect to find the optimal solution in polynomial time. What do we do in the face of NP-completeness? There are a few options, including just giving up on proving theorems and designing algorithms that we hope give “good-enough” solutions. Let's not give up quite yet, though. Instead, let's try to design *approximation algorithms*: algorithms which run in polynomial time, and give solutions that are *provably* not too far from the optimal solution. We might not be able to find the optimal solution on polynomial time, but maybe we can find a solution that costs at most twice as much.

Definition 23.1.1 *Let \mathcal{A} be some (minimization) problem, and let I be an instance of that problem. Let $OPT(I)$ be the value of the optimal solution on that instance. Let ALG be a polynomial-time algorithm for \mathcal{A} , and let $ALG(I)$ denote the value of the solution returned by ALG on instance I . Then we say that ALG is an α -approximation if*

$$\frac{ALG(I)}{OPT(I)} \leq \alpha$$

for all instances I of \mathcal{A} .

One interesting thing to note is that, as we saw last week, the theory of NP-completeness tells us that all NP-complete problems are in some sense equally hard – if we could solve one of them in polynomial time, then we could solve all of them. However, it is *not* true that they are all equally hard to approximate. Some can be approximated extremely well in polynomial time, and some cannot be approximated at all. So sometimes it is useful to think of approximability as a more “fine-grained” notion of hardness: while all NP-complete problems are hard, problems which can be approximated within 2 are easier than problems which can only be approximated within $\Theta(\log n)$, which are easier than problems which can only be approximated within $\Theta(\sqrt{n})$, etc.

23.2 Vertex Cover

For our first example, let's go back to the vertex cover problem.

Definition 23.2.1 *Let $G = (V, E)$ be a graph. Then $M \subseteq V$ is a vertex cover if for every edge $\{u, v\} \in E$, at least one of u, v is in M .*

In other words, a vertex cover is a set of vertices with the property that every edge has at least one endpoint in the set. In the vertex cover problem, we are given a graph $G = (V, E)$ and are asked to find the *smallest* vertex cover. This problem is NP-complete by an easy reduction from the Independent Set problem (which we didn't quite have time for last lecture, but is in the lecture notes). So we cannot expect to actually solve vertex cover. What can we do instead?

1. Idea 1: Pick an arbitrary vertex with at least one uncovered edge incident on it, add it to the cover, and repeat. Unfortunately this is arbitrarily far from optimal: see the star graph.
2. Idea 2: Instead of picking arbitrarily, let's try to pick smartly. In particular, an obvious thing to try is the greedy algorithm: pick the vertex with the largest number of uncovered edges incident to it, and add it to the cover. Repeat until all edges are covered.

While this is a better idea, it's still not very good. We'll give a quick overview of the counterexample, but fleshing out the details is a good exercise to do at home. Consider a set U of t nodes. For every $i \in \{2, 3, \dots, t\}$, divide U into $\lfloor t/i \rfloor$ disjoint groups of size exactly i (if i does not divide t , then there will be less than i nodes which are not in any group). Let the groups for value i be $G_1^i, G_2^i, \dots, G_{\lfloor t/i \rfloor}^i$. For every $i \in \{2, \dots, t\}$ and every $j \in [\lfloor t/i \rfloor]$, we create a vertex v_j^i and add edges between v_j^i and every node in G_j^i . We refer to the nodes $\{v_j^i\}_{j \in [\lfloor t/i \rfloor]}$ as *layer i nodes*.

It is easy to see that in this graph, every node in U has degree at most $t-1$ since it is adjacent to at most one node from each layer. Every node in layer i has degree exactly i . So at the beginning of the algorithm, the maximum degree node is the one node v_1^t in layer t . After we pick this node, the nodes in U now have degree at most $t-2$, so the algorithm would next pick the node in layer $t-1$. It is easy to see by induction that this will continue: the algorithm will always choose the nodes in the largest remaining layer rather than the nodes in U .

Notice that $OPT(G) \leq t$, since U itself is a vertex cover of size t . On the other hand, we just argued that

$$ALG(G) = \sum_{i=2}^t \left\lfloor \frac{t}{i} \right\rfloor \geq \sum_{i=2}^t \left(\frac{1}{2} \cdot \frac{t}{i} \right) = \frac{t}{2} \sum_{i=2}^t \frac{1}{i} = \Omega(t \log t).$$

Since $ALG(G)/OPT(G) \geq \log t$ is not a constant, the greedy algorithm is not an α -approximation for any constant α .

OK, so now let's find some algorithms which work better. Here's an algorithm which sounds stupid but is actually pretty good:

1. Pick an arbitrary edge which is not yet covered. Add *both* endpoints to the cover, and repeat.

I claim that this algorithm is a 2-approximation. To see this, suppose that our algorithm took k iterations, and let L be the set of edges that it selected as uncovered and included both of the endpoints (so $|L| = k$). The algorithm returns a vertex cover of size $2k$. On the other hand, note that these edges form a matching, i.e. they all have distinct endpoints. This means that *any* vertex cover needs to have size at least k , just in order to cover the edges in L , and hence $OPT \geq k$. Thus $ALG = 2k \leq 2 \times OPT$, so it is a 2-approximation.

Now let's see a more involved way to get a 2-approximation. At first this will seem unnecessarily complicated, but we'll see later why it is useful.

2. Let $G = (V, E)$ be an instance of vertex cover, and consider the following linear program:

$$\begin{aligned} \min \quad & \sum_{u \in V} x_u \\ \text{subject to} \quad & x_u + x_v \geq 1 \quad \forall \{u, v\} \in E \\ & 0 \leq x_u \leq 1 \quad \forall u \in V \end{aligned}$$

Our algorithm solves this LP, and then for every vertex v we include v in our cover if $x_v \geq 1/2$. We first claim that our algorithm does indeed give a vertex cover. To see this, consider an arbitrary edge $\{u, v\} \in E$. Then the LP has a constraint $x_u + x_v \geq 1$, so clearly either x_u or x_v (or both) is at least $1/2$. Thus at least one endpoint of the edge will be contained in the cover. Since this holds true of all edges, we get a valid vertex cover.

So now we want to prove the the vertex cover we return is a 2-approximation. Let \vec{x}^* be the optimal solution to the above LP on G , let $LP(G) = \sum_{u \in V} x_u^*$ denote the value of the above LP on G , let $OPT(G)$ denote the size of the minimum vertex cover, and let $ALG(G)$ denote the size of the vertex that we return.

First, note that $LP(G) \leq OPT(G)$: this is true because for each vertex cover, there is an LP solution with value exactly equal to the size of the cover (just set $x_u = 1$ if u is in the cover and $x_u = 0$ otherwise). Since the LP is trying to minimize an objective, it does *at least* as well as the smallest vertex cover. On the other hand, we claim that the algorithm does not do too much worse than the LP. To see this let $\hat{x}_u = 1$ if $x_u^* \geq 1/2$ and let $\hat{x}_u = 0$ if $x_u^* < 1/2$. Note that $\hat{x}_u \leq 2x_u^*$. Hence

$$ALG(G) = \sum_{u \in V} \hat{x}_u \leq 2 \sum_{u \in V} x_u^* = 2 \times LP(G).$$

Thus $ALG(G) \leq 2 \times LP(G) \leq 2 \times OPT(G)$, so the algorithm is a 2-approximation.

The second algorithm is known as an *LP rounding* algorithm: we first solve an LP which if we could enforce integrality would correspond exactly to the problem we want to solve. Since we can't enforce integrality we instead get back fractional value. Then we somehow want to round these fractional values to integers (usually to 0 or 1). This algorithm is a version of *threshold rounding*, since all we did was set a threshold and round up to 1 any value above the threshold and round down to 0 any value below the threshold. But there are many other types of LP rounding. The most famous is probably *randomized rounding*, where we interpret each LP variable as a probability and set it to 1 with probability equal to its value.

Since we already had a super simple 2-approximation, why did we bother with the LP rounding algorithm? One really nice feature of LP rounding algorithms is that they tend to be extremely flexible to slight changes in the problem. For example, consider the *weighted* vertex cover problem, in which each vertex v also has a positive weight w_v and we try to find the minimum *weight* vertex cover. It is not at all clear how to adapt our first algorithm to this setting. But with the LP algorithm it's trivial – we just change the objective function to $\min \sum_{v \in V} w_v x_v$. Then the rest of the analysis essentially goes through without change! Rounding with a threshold of $1/2$ still gives a 2-approximation!

Reductions and Approximation. One more note about Vertex Cover: remember that we proved that Vertex Cover was NP-hard by reducing from Independent Set. However, the type of reduction we did is not *approximation-preserving*: it's true that if we could *solve* Vertex Cover then we could solve Independent Set (we proved this), but it is not true that an α -approximation for Vertex Cover implies an α -approximation for Independent Set. In fact, Independent Set is much harder: unless $P=NP$, it is not possible even to *approximate* Independent Set better than $n^{1-\varepsilon}$ for any constant $\varepsilon > 0$.

23.3 Max-E3SAT

Recall that in the 3SAT problem we are given a 3-CNF boolean formula (i.e. a formula which was the AND of a collection of clauses, and each clause was the OR of at most 3 literals) and asked to determine if there exists a satisfying assignment. A natural optimization version of this problem is Max-3SAT, in which we are again given a collection of clauses, each of which is the OR of at most 3 literals, but where our goal is to find an assignment which maximizes the number of satisfied clauses. So if there is a fully satisfying assignment then we hope to find it, but if not then we try to find an assignment which satisfies as many clauses as possible. We'll actually spend most of our time discussing Max-E3SAT (or maximum *exact* 3SAT) in which every clause has *exactly* three literals.

First, one technicality: since this is a maximization problem, we want to redefine an α -approximation to mean that $ALG(I)/OPT(I) \geq \alpha$ for all instances I (rather than $\leq \alpha$).

Here's an easy randomized algorithm: simply use a random assignment! It turns out that this is a pretty good approximation.

Theorem 23.3.1 *A random assignment is (in expectation) a 7/8-approximation for Max-E3SAT.*

Proof: Consider a clause. Since there are exactly 3 literals in it, there are exactly $2^3 = 8$ possible assignments to it, of which 7 are satisfying. So the probability that a random assignment satisfies it is exactly 7/8. Now linearity of expectations lets us apply this to all of the clauses: the expected number of clauses satisfied is exactly $(7/8)m$, where m is the total number of clauses. Since $OPT(I) \leq m$, this implies that a random assignment is in expectation a 7/8-approximation.

■

What if we want a *deterministic* algorithm? This turns out to be reasonably simple based on a derandomization of the randomized algorithm using what is known as the “method of conditional expectations”.

Consider the first variable x_1 . If we set it to 0, then some clauses become satisfied (those with \bar{x}_1 as a literal) and some others can no longer be satisfied by x_1 (those with x_1 as a literal). So we can calculate the number of expected clauses that would be satisfied if we set x_1 to 0 and set all of the other variables randomly. Similarly, we can calculate the number of expected clauses that would be satisfied if we set x_1 to 1 and set all of the other variables randomly. Then we'll just set x_1 according to whichever of these expectations is larger. And now that x_1 is set we can do the same thing with x_2 (where now x_1 is fixed), x_3 , etc., until in the end we have fixed an assignment. Since we always pick whichever of the two expectations is larger, it never goes down, and thus stays at

least $7/8$. So in the end we have satisfied at least $7/8$ of the clauses.

Amazingly, this is essentially the best possible approximation!

Theorem 23.3.2 ([1]) *Assuming $P \neq NP$, for all constant $\epsilon > 0$ there is no polytime $(\frac{7}{8} + \epsilon)$ -approximation for Max-E3SAT.*

This turns out to be a consequence of an amazing theorem known as *Håstad's 3-bit PCP Theorem* [1], which is unfortunately way outside the scope of this class.

23.4 Set Cover (likely not covered in lecture)

A fundamental problem which we have not really talked about yet is *Set Cover*. The input to Set Cover is a set X of n items, and m subsets S_1, \dots, S_m of X . The goal is to find the minimum number of these subsets necessary to cover all of the items. In other words, we must find a set $A \subseteq [m]$ of indices so that $\cup_{i \in A} S_i = X$, and our goal is to minimize $|A|$. It's known that Set Cover is NP-complete.

23.4.1 Greedy Algorithm

It turns out that the simple greedy algorithm does pretty well. This algorithm picks the set which covers the most items, throws out all of the items covered, and repeats. In other words, we just always pick the set which covers the most uncovered items.

Theorem 23.4.1 *The greedy algorithm is a $(\ln n)$ -approximation algorithm for Set Cover.*

Proof: Suppose that the optimal set cover has size k . Then there must be some set which covers at least a $1/k$ fraction of the items. Since we always pick the best remaining set, this means that after the first iteration there are at most $(1 - \frac{1}{k})n$ items still uncovered. But now the logic continues to hold! No matter what items have been covered already, the original optimal solution certainly covers all of the remaining items (since it covers all items). So there must be some set from that solution which covers at least a $1/k$ fraction of the remaining items.

Thus after i iterations, the number of uncovered items is at most $(1 - \frac{1}{k})^i n$. So after $k \ln n$ iterations, the number of uncovered items is at most $(1 - \frac{1}{k})^{k \ln n} n < e^{-\ln n} n = 1$. Thus we are finished after $k \ln n$ iterations and so we chose at most $k \ln n$ sets, which immediately implies that it's the greedy algorithm is a $(\ln n)$ -approximation ■

Note that this analysis is quite similar to our analysis of Edmonds-Karp #1: we prove that in each iteration, we can reduce the "amount to go" by some fraction x , and then use the inequality $(1 - x)^t \leq e^{-xt}$ to bound the "amount to go" after t iterations.

In fact, we know that this is the best possible approximation for Set Cover: it is NP-hard to give an $o(\log n)$ -approximation ratio, and under a stronger complexity assumption (that problems in NP cannot be solved in $n^{O(\log \log n)}$ time) the lower bound can be improved to $(1 - \epsilon) \ln n$ for arbitrarily small constant $\epsilon > 0$ (this is due to Uriel Feige).

23.4.2 LP Rounding 1

There are also algorithms for approximating Set Cover based on LP rounding. The obvious LP relaxation for Set Cover is the following:

$$\begin{aligned} \min \quad & \sum_{i \in [m]} x_i \\ \text{subject to} \quad & \sum_{i \in [m]: e \in S_i} x_i \geq 1 \quad \forall e \in X \\ & 0 \leq x_i \leq 1 \quad \forall i \in [m] \end{aligned}$$

As with vertex cover, let x^* denote the optimal solution to the LP, let $LP^* = \sum_{i \in [m]} x_i$ denote the cost of this optimal solution, let A^* denote the optimal solution, and let $OPT = |A^*|$ denote the cost (number of subsets) of the optimal solution. Then as before, $LP^* \leq OPT$ since if we set $x_i = 1$ for all $i \in A^*$ and $x_i = 0$ for all $i \notin A^*$, we get a solution to the LP of cost OPT . So if we want to prove that some algorithm is an α -approximation, it suffices to prove that it gives a solution of cost at most $\alpha \cdot LP^*$.

Consider the following threshold rounding algorithm. For every element $e \in X$, let $f(e) = \{i : e \in S_i\}$ denote the *frequency* of e . Let $f = \max_{e \in X} f(e)$. Our algorithm first solves the LP (in polynomial time), and then we let $A = \{i : x_i \geq 1/f\}$.

Theorem 23.4.2 *Threshold rounding with threshold $1/f$ is an f -approximation to Set Cover.*

Proof: This is very similar to the Vertex Cover analysis. First, note that the algorithm clearly runs in polynomial time, since it solves an LP in polynomial time and thus does a simple threshold rounding (taking at most polynomial time). Consider some element $e \in X$. Since $f(e) \leq f$, and $\sum_{i \in [m]: e \in S_i} x_i \geq 1$ (by the LP constraint), we know that $x_i \geq 1/f(e) \geq 1/f$ for at least one element in $\{i : e \in S_i\}$. Thus our solution A is indeed a set cover.

Now we just need to analyze its cost. Let $\hat{x}_i = 1$ iff $i \in A$, i.e., $\hat{x}_i = 1$ iff $x_i \geq 1/f$. Thus $|A| = \sum_{i \in [m]} \hat{x}_i \leq \sum_{i \in [m]} f x_i = f \sum_{i \in [m]} x_i = f \cdot LP^*$, and thus the algorithm is an f -approximation. ■

23.4.3 LP Rounding 2

The previous LP rounding is all well and good, but the f -factor it gives is incomparable to the $\log n$ of the greedy algorithm (it could be larger or smaller). Can we use the same LP to also get an $O(\log n)$ -approximation. It turns out that the answer is yes, and moreover that this is relatively easy if we use *randomization*.

Consider the following algorithm based on *randomized rounding*. Instead of doing threshold rounding, we instead include every index $i \in [m]$ in A with probability $\min(2x_i^* \ln n, 1)$.

Theorem 23.4.3 *A is a set cover with probability at least $1 - 1/n$, and $\mathbf{E}[|A|] \leq 2 \ln n \cdot LP^*$.*

Proof: Consider some element e . The probability that it is not covered by A is the probability that none of the sets containing it are added to A . If some set S_i containing e has $x_i \geq 1/(2 \ln n)$

then this set is added to A with probability 1, so e is covered. Otherwise, the probability that e is not covered is

$$\prod_{i \in [m]: e \in S_i} (1 - 2x_i^* \ln n) \leq \prod_{i \in [m]: e \in S_i} e^{-2x_i^* \ln n} = e^{-2 \ln n \sum_{i \in [m]: e \in S_i} x_i^*} \leq e^{-2 \ln n} = 1/n^2.$$

Now taking a union bound over all n elements implies that the probability that there exists an uncovered element is at most $n(1/n^2) = 1/n$, and thus A is a valid set cover with probability at least $1 - 1/n$.

It is even easier to see the cost bound. Let X_i be an indicator random variable for the event that $i \in A$. Then

$$\mathbf{E}[|A|] = \mathbf{E}\left[\sum_{i \in [m]} X_i\right] = \sum_{i \in [m]} \mathbf{E}[X_i] \leq \sum_{i \in [m]} 2x_i^* \ln n = 2 \ln n \cdot LP^*.$$

■

In other words, randomized rounding is an $O(\log n)$ -approximation (it's a good exercise to do at home to figure out why it's OK to give a high probability bound on being a feasible solution with an expectation bound on the cost).

References

- [1] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, jul 2001.