

1 Vertex Cover and Relax-and-Round

In the Vertex Cover problem, we are given a graph $G = (V, E)$ and the goal is to select a minimum cardinality set of vertices $S \subseteq V$ such that every edge is adjacent to at least one vertex in S . We will use this problem to introduce the Relax-and-Round framework.

Question: what is a natural algorithm for this problem? The first thing you might try is the *greedy* algorithm: take the vertex of largest degree, put it in the cover and delete it and its edges from the graph. Iterate until the graph has no edges. It turns out there are examples where this is a $\Omega(\log n)$ approximation.

1.1 A 2-approximation

Despite this $\Omega(\log n)$ bound for greedy, it turns out there is an easy 2-approximation: pick any uncovered edge and put both vertices in the cover. Iterate until no edges remain.

Fact 1.1. *The algorithm which iteratively picks both endpoints of an uncovered edge and puts them in the cover is a 2-approximation.*

Proof. Consider the uncovered edges e_1, \dots, e_k picked by the algorithm. They must form a matching, since if any two edges e_i and e_j for $j > i$ shared an endpoint, then e_j would have been covered when the two endpoints of e_i were put in the cover. Furthermore, the solution has size $2k$.

The optimal vertex cover by definition must contain at least one vertex adjacent to every edge. However every vertex is adjacent to at most one edge among e_1, \dots, e_k . Therefore the optimal matching has size at least k . Since the algorithm outputs a matching of size $2k$, this is a 2-approximation. \square

1.2 Relax and Round

Now consider the variant of Vertex Cover where the vertices have costs and the goal is to minimize the cost of the cover. The same 2-approximation does not work, and it is not obvious how to fix it. We will use the relax and round framework to make this problem easy for us. The framework has three ingredients. A big advantage of this framework is that it is highly re-usable, and similar ideas work for large groups of problems.

Step 1: Model the problem as an Integer Linear Program (ILP). For vertex cover, this is the following, where we have a variable x_v for each vertex v indicating whether or not it is in the cover and the cost of vertex v is c_v :

$$\begin{aligned} \min \quad & c(x) \\ \text{s.t.} \quad & x_u + x_v \geq 1 \quad \forall e = \{u, v\} \in E \\ & x_v \in \{0, 1\} \quad \forall v \in V \end{aligned} \tag{1}$$

Here $c(x) = \sum_{v \in V} c_v x_v$. This is an ILP because it consists of a linear objective function, linear constraints, and the variables are required to be integers.

This is clearly an equivalent problem to Vertex Cover, so it's still NP-Hard, and it doesn't seem like we've gained anything. That's why we do this next step:

Step 2: Relax the ILP into a Linear Program (LP). This just means we remove the requirement that the variables are integers.

$$\begin{aligned} \min \quad & c(x) \\ \text{s.t.} \quad & x_u + x_v \geq 1 \quad \forall e = \{u, v\} \in E \\ & 0 \leq x_v \leq 1 \quad \forall v \in V \end{aligned} \tag{2}$$

The result is an LP. It turns out *all* LPs can be solved in polynomial time by the ellipsoid method so long as the set of constraints has a separation oracle, as was proved by Khachiyan in 1979.

Definition 1.2 (Separation Oracle). *A separation oracle for a linear program is a polynomial time procedure that given a point $x \in \mathbb{R}^n$ either certifies that x satisfies all constraints or finds a constraint that x violates.*

In other words, the set of constraints in our linear program should either (i) have polynomial size or (ii) given a possible solution x , we should be able to determine in polynomial time if all constraints are met and, if not, which constraint is not met. In the case of vertex cover, we have only polynomially many constraints (just $|E| + 2n$) so we can solve the LP in polynomial time.

So, great: now we have some solution x to this LP. We also know that $c(x) \leq c(OPT)$. This is because OPT is a feasible solution to the LP and we found the cheapest LP solution. But, x has fractional coordinates. What do we do to get an actual cover? In comes the final part:

Step 3: Round the fractional solution to an integer one. For Vertex Cover, if we want a 2-approximation, this step turns out to be very easy. Think a moment and see if you can figure it out. Remember you can use that $c(x) \leq c(OPT)$.

The idea is just to let the cover consist of all vertices v with $x_v \geq \frac{1}{2}$. This is called **threshold rounding** and is arguably the simplest form of rounding.

Fact 1.3. *Applying a threshold rounding at $\frac{1}{2}$ is a 2-approximation for vertex cover.*

Proof. Let S be the cover returned by the threshold rounding. Consider any constraint $x_u + x_v \geq 1$ for some $e = (u, v)$. Either x_u or x_v must be at least $\frac{1}{2}$, as otherwise $x_u + x_v < 1$. Therefore, either u or v is in S , demonstrating that S is a feasible vertex cover.

To analyze the cost of S , consider:

$$c(S) = \sum_{v: x_v \geq \frac{1}{2}} c_v \leq \sum_{v \in V} 2x_v c_v = 2c(x) \leq 2c(OPT),$$

as desired. □

We can also study the "strength" of this relaxation by measuring how far a fractional solution can be from an integer one in the worst case.

Integrality Gap

The **integrality gap** of an ILP (and its associated LP) is the worst-case value of $\frac{c(OPT_{ILP})}{c(OPT_{LP})}$ over all instances.

Sometimes you will also see the integrality gap of an LP defined as the worst-case value of $\frac{c(OPT)}{c(OPT_{LP})}$ over all instances (with no reference to an ILP). But this is equivalent so long as the ILP is equivalent to the original problem.

Upper bounding the integrality gap can be accomplished in several ways, but designing an approximation algorithm using the LP as a bound is clearly one such way:

Fact 1.4. *Given a feasible solution x to an LP, if a rounding algorithm A always produces a feasible solution to the corresponding ILP of cost at most $\alpha \cdot c(x)$, then the integrality gap of the ILP is at most α .*

So, we know that the integrality gap of the polytope given by (1) is at most 2. Lower bounding the integrality gap is usually done using an example.




Figure 1: On the left is an optimal vertex cover where the four red nodes have $x_v = 1$ and the white node is set to 0. On the right is an optimal fractional vertex cover where all nodes have $x_v = 1/2$.

It turns out that there is also a lower bound of 2: a complete graph has an optimal integral solution of value $n - 1$, while it has an optimal fractional solution of value $\frac{n}{2}$, as demonstrated in Fig. 1. Therefore, the value of the optimal solution to (2) may be only half as large as value of the integer optimal solution (which would be faithfully returned by (1)).

1.3 Summary for Vertex Cover

So, the situation is quite good for vertex cover. We have an approximation algorithm with ratio 2, and we know the LP we used has an integrality gap of 2, so if we want to use $c(x)$ as a lower bound we cannot do better than 2. It turns out we cannot get an approximation algorithm with ratio better than 2 under the Unique Games Conjecture (UGC) [KR08]. It is NP-Hard to approximate better than $\sqrt{2} \approx 1.414$ [KMS23].

For most problems, we are farther from knowing the truth. Often, we do not even know the integrality gap for the natural ILP.

1.4 Strengthening the LP

Even though the ILP is equivalent to the original problem, we can still add constraints that may strengthen the LP to reduce the integrality gap.

For example, for the vertex cover problem, we can write the constraint that $x_u + x_v + x_w \geq 2$ for all triangles $e = \{u, v\}, f = \{v, w\}, g = \{w, u\}$ in the graph. This is a valid constraint, since at least two vertices in any triangle need to be in the cover. Now the integrality gap of some specific instances goes down. For example, it used to be that a triangle with all edges set to $1/2$ was a feasible LP solution of value $3/2$, whereas OPT was 2. Now on this instance, the LP solves it optimally.

Notice that the ILP *didn't need* this extra constraint, because it was implied by the integrality of the variables! We can derive it in the ILP as follows:

$$x_u + x_v \geq 1, x_v + x_w \geq 1, x_u + x_w \geq 1$$

$$2x_u + 2x_v + 2x_w \geq 3$$

$$x_u + x_v + x_w \geq \frac{3}{2}$$

But this means that $x_u + x_v + x_w \geq 2$ since all variables are integers, so the RHS must be an integer.

In some cases, adding constraints like this decreases the integrality gap of the LP. In this specific case, it unfortunately does not. In fact, it was recently shown that no polynomial sized LP has an integrality gap below 2 for Vertex Cover [Baz+19] (however, note that this doesn't rule out exponential sized LPs with separation oracles).

References

- [Baz+19] Abbas Bazzi, Samuel Fiorini, Sebastian Pokutta, and Ola Svensson. “No Small Linear Program Approximates Vertex Cover Within a Factor $2 - \epsilon$ ”. In: *Math. Oper. Res.* 44.1 (2019), 147–172. ISSN: 0364-765X. DOI: [10.1287/moor.2017.0918](https://doi.org/10.1287/moor.2017.0918) (cit. on p. 4).
- [KMS23] Subhash Khot, Dor Minzer, and Muli Safra. “Pseudorandom sets in Grassmann graph have near-perfect expansion”. In: *Annals of Mathematics* 198.1 (2023), pp. 1–92. DOI: [10.4007/annals.2023.198.1.1](https://doi.org/10.4007/annals.2023.198.1.1) (cit. on p. 3).
- [KR08] Subhash Khot and Oded Regev. “Vertex cover might be hard to approximate to within $2 - \epsilon$ ”. In: *J. Comput. Syst. Sci.* 74.3 (2008), 335–349. ISSN: 0022-0000. DOI: [10.1016/j.jcss.2007.06.019](https://doi.org/10.1016/j.jcss.2007.06.019) (cit. on p. 3).