

Lecture 7: Approximation Algorithms

Notes by Ola Svensson¹


We first finish the formal analysis of the Hungarian algorithm, and then talk about settings in which relaxing integrality has a cost, i.e. about approximation algorithms.

1 The Hungarian Algorithm

We start with some intuition before we give the formal description of the algorithm.

1.1 Intuition

Consider the following instance of the min-cost perfect matching problem:




The thin edges have cost 1, whereas the thick edge has cost 2. The Hungarian algorithm will use Lemma ?? in the following way: we will maintain a dual solution u, v that is *feasible at all times*. Then, for a fixed dual solution, the lemma tells us that our perfect matching is only allowed to contain edges that are *tight*, i.e., edges $e = (a, b)$ for which $u_a + v_b = c(e)$. This reduces our problem to finding any perfect matching in the subgraph consisting only of tight edges, i.e., in the graph (V, E') where $E' = \{e = (a, b) \in E : u_a + v_b = c(e)\}$. *Intuitively, we have thus reduced our weighted problem to an unweighted one!*

Let us return to our example. We initialize our procedure with the trivial dual solution

$$v_b = 0, \quad u_a = \min_{b \in B} c_{ab}.$$

(We could have also started from $u = v = 0$.) So in the example, $v_4 = v_5 = v_6 = 0$ and $u_1 = u_2 = u_3 = 1$. The set E' of tight edges is thus:



We then try to find a perfect matching in this graph using e.g. the augmenting path algorithm we saw in Lecture 3. However, the considered graph has *no* perfect matching! This is because we started with a poor lower bound (dual solution). So we will use the fact that there is no perfect matching to improve our dual solution. We make use of *Hall's condition*, which you will prove in the exercise session:

¹**Disclaimer:** These notes were written as notes for the lecturer. They have not been peer-reviewed and may contain inconsistent notation, typos, and omit citations of relevant works.


Theorem 1 (Hall's Theorem) *An n -by- n bipartite graph $G = (A \cup B, E')$ has a perfect matching if and only if $|S| \leq |N(S)|$ for all $S \subseteq A$.*

Here $N(S) = \{b \in B : \text{there is an } a \in S \text{ such that } (a, b) \in E'\}$ denotes the neighborhood of the vertices in S .

In the above example, we have $S = \{1, 2\}$ and $N(S) = \{4\}$, which violates Hall's condition (and thus there is no perfect matching). We now use the set S to improve our dual lower bound.

We gradually increase u_a for every $a \in S$ and at the same time decrease v_b for $b \in N(S)$ at the same rate. Let us see what happens to all edges in E' . Notice that the tight edges between S and $N(S)$ will remain tight. Similarly, the tight edges between $A \setminus S$ and $B \setminus N(S)$ will remain tight. Any tight edges between $A \setminus S$ and $N(S)$ will stop being tight. Finally, by definition, there are no edges from S to $B \setminus N(S)$ initially. We continue to gradually change the dual solution until some such edge becomes tight.

In the above example, this will result in updating $u_1 = u_2 = 2$ and $v_4 = -1$. We have thus increased two variables by one unit and decreased one variable by one unit. In total, the dual lower bound was thus increased by one. The set E' of tight edges with respect to the new dual solution is now



Our (augmenting-path) algorithm will now find a perfect matching in this graph, which is optimal by Lemma ?? (since it only uses tight edges). In summary, our algorithm always maintains a dual feasible solution. We then solve the *unweighted* perfect matching problem on the edges allowed by Lemma ??. In the case of failure, we update the dual lower bound to a strictly better lower bound and repeat.

1.2 Formal description

Idea: Maintain a feasible dual solution (u, v) . Try to construct a feasible primal solution that satisfies complementarity slackness (Lemma ??) and is thus optimal.

Algorithm:

- **Initialization:**

$$v_b = 0, \quad u_a = \min_{b \in B} c_{ab}.$$

- **Iterative step:**

- consider $G' = (A \cup B, E')$ where $E' = \{e = (a, b) \in E : u_a + v_b = c(e)\}$ (E' is the set of all tight edges);
- find a maximum-cardinality matching in G' :
 - * if it is a perfect matching, then we are done (this is a primal feasible solution and it satisfies complementarity slackness, because we consider only the edges in E' and all of them satisfy slackness by construction),
 - * otherwise the algorithm finds a set $S \subseteq A$ s.t. $|S| > |N(S)|$ (which is guaranteed to exist by Hall's theorem)

- we can choose a small $\varepsilon > 0$ and improve the dual solution:

$$u'_a = \begin{cases} u_a + \varepsilon & \text{if } a \in S, \\ u_a & \text{if } a \notin S, \end{cases}$$

$$v'_b = \begin{cases} v_b - \varepsilon & \text{if } b \in N(S), \\ v_b & \text{if } b \notin N(S), \end{cases}$$

which remains dual feasible because:

- * edges in $S \times N(S)$ are unchanged ($+\varepsilon - \varepsilon$),
- * edges in $(A \setminus S) \times (B \setminus N(S))$ are unchanged,
- * edges in $(A \setminus S) \times N(S)$ are decreased by ε ,
- * edges in $S \times (B \setminus N(S))$ are increased by ε but they were not tight (they were not in E'),
- the dual value increases by $(|S| - |N(S)|)\varepsilon$;
- to make as much progress as possible (and to get a new tight edge), we should choose ε as large as possible for which the dual remains feasible, that is

$$\varepsilon = \min_{e=(a,b) \in S \times (B \setminus N(S))} c(e) - u_a - v_b > 0.$$

The above algorithm can be implemented to run in time $O(n^3)$. (This is quite non-trivial to see. It is somewhat easier to implement in $O(n^4)$.)

1.3 An alternative proof that the bipartite perfect matching polytope is integral

We have shown that for any cost function c we can obtain a minimum-cost perfect matching and it will be integral. By carefully choosing the cost function, one can make any extreme point of the polytope to be the unique optimum solution to the minimum-cost perfect matching problem. This shows that all extreme points are integral, i.e., it is an integral polytope. See Figure 1 for an example.




Figure 1: Example of a perfect matching polytope and a cost function vector c for the highlighted vertex.

2 When Relaxing Integrality is Not for Free: Approximation Algorithms

Thousands (or millions :)) of optimization problems have been proved to be NP-hard, which means that finding optimal solutions for them is very likely to be intractable. More specifically, unless $P = NP$, there is no algorithm for an NP-hard optimization problem that satisfies all of the following three desiderata:

1. The algorithm is efficient (runs in polynomial time).
2. The algorithm is reliable (works for any input instance).
3. The algorithm finds an optimal solution (optimality).

Therefore, when confronted with such problems, we need to relax one of the above conditions when dealing with NP-hard optimization problems. If we relax reliability, then we get heuristics that are designed to work well on instances that commonly appear in practice. If we relax the third condition, we obtain *approximation algorithms*. The notion of approximation algorithms allows us to obtain a more fine-grained picture of the difficulty of NP-hard optimization problems: some problems turn out to have very good approximations, whereas others resist such attempts.

The formal definition of approximation algorithms is as follows.

Definition 2 *An α -approximation algorithm for a given optimization problem is an algorithm that runs in polynomial time and outputs a solution S such that:*

- $\frac{\text{cost}(S)}{\text{cost}(\text{Optimal solution})} \leq \alpha$ if the problem is a minimization problem,
- $\frac{\text{profit}(S)}{\text{profit}(\text{Optimal solution})} \geq \alpha$ if the problem is a maximization problem.

It is clear that we will have $\alpha \geq 1$ for minimization problems and $\alpha \leq 1$ for maximization problems. Moreover, if $\alpha = 1$, then we have an efficient exact algorithm (and the problem is in P). The goal when designing approximation algorithms is to achieve a guarantee α that is as close to 1 as possible. We now introduce a general framework for using linear programming in the design of approximation algorithms followed by an application: vertex cover.


3 Using Linear Programming to Design Approximation Algorithms

Consider a minimization problem. When considering the definition of approximation algorithms, it seems very hard, even for a specific instance of the problem, to analyze the approximation guarantee $\frac{\text{cost}(S)}{\text{cost}(\text{Optimal solution})}$ of the algorithm on the instance, as we are comparing ourselves with an optimal solution (that is most likely very hard to compute and does not possess any nice structure). The solution to this is to compare ourselves with a *lower bound* on the optimum. Indeed, we then have that

$$\frac{\text{cost}(S)}{\text{cost}(\text{Optimal solution})} \leq \frac{\text{cost}(S)}{\text{lower bound on opt}} \leq \alpha.$$

From the above, it is clear that we need a good lower bound to be able to claim a good guarantee α . To obtain such a lower bound (and to design the approximation algorithm), linear programming is super handy. A popular framework is as follows:

1. Give an exact formulation of the problem as Integer LP – usually with binary variables ($x_i \in \{0, 1\}$).
2. Relax to LP – $x_i \in [0, 1]$



3. Solve LP to get a optimal solution x^* to the LP which is a lower (upper) bound on the optimal solution to Integer LP and thus the original problem. Then somehow round x^* to an integral solution "without losing too much" (which will determine the guarantee α).

We now use the above framework for the vertex cover problem, and then introduce the concept of integrality gap.

3.1 Vertex Cover

Here, we apply the framework to get an approximation algorithm for the *Vertex Cover* problem. First recall the definition:

Definition 3 (Vertex Cover (VC) Problem) *Given a graph $G = (V, E)$ and a weight function on the vertices $w : V \rightarrow \mathbb{R}_+$, output a set $C \subseteq V$ of minimum weight such that for all $\{u, v\} \in E$, $u \in C$ or $v \in C$.*

First, we define an ILP solving VC: For all $v \in V$, we introduce a variable x_v , which is 1 if $v \in C$, and 0 otherwise. The objective function is

$$\min \sum_{v \in V} w(v)x_v$$

and for each $\{u, v\} \in E$ we introduce the constraint $x_u + x_v \geq 1$. This ensures that for each edge, at least one endpoint is in C . Additionally, we require that for each $v \in V$ we have $x_v \in \{0, 1\}$. Note that at this point our Integer LP formulation is exactly equivalent to the original problem.

Second, we relax the ILP to an LP by allowing that $x_v \in [0, 1]$. Actually, it's sufficient to require that $x_v \geq 0$, because having an x_v greater than 1 will not satisfy any additional constraint, but only increase the value of the objective function, so this will not happen in an optimal solution.

Third, we have to do the rounding: Suppose we've solved the LP and got an optimal solution x^* . We will return $C = \{v \in V : x_v^* \geq \frac{1}{2}\}$ as our solution to VC.

Claim 4 C is a feasible solution.

Proof Consider any edge $\{u, v\}$ and its constraint. Since $x_u^* + x_v^* \geq 1$, at least one of x_u^*, x_v^* is $\geq \frac{1}{2}$ and thus in C . ■

Claim 5 The weight of C is at most twice the value of the optimal solution of VC.

Proof We have

$$\sum_{v \in C} w(v) = \sum_{v \in V : x_v^* \geq \frac{1}{2}} w(v) \leq \sum_{v \in V : x_v^* \geq \frac{1}{2}} 2x_v^* w(v) \leq \sum_{v \in V} 2x_v^* w(v) = 2 \sum_{v \in V} x_v^* w(v) = 2LP_{OPT} \leq 2VC_{OPT}$$

where the first inequality holds because $2x_v^* \geq 1$. ■

So, we have designed a 2-approximation algorithm for Vertex Cover. This is a simple algorithm using linear programming. To appreciate the power of the framework, I challenge you to find a 2-approximation algorithm for Vertex Cover (with node-weights) without the use of LPs. (It is possible but quite difficult.)

3.2 Integrality Gap

The notion of the *integrality gap* allows us to bound the power of our linear programming relaxation. Let $\mathcal{I}$ be the set of all instances of a given problem. In the case of a minimization problem, the integrality gap g is defined as

$$g = \max_{I \in \mathcal{I}} \frac{OPT(I)}{OPT_{LP}(I)}.$$

As an example, suppose $g = 2$, and that LP found that $OPT_{LP} = 70$. Then, since our problem instance might be the one which maximizes the expression for g , all we can guarantee is that $OPT(I) \leq 2 \cdot OPT_{LP}(I) = 140$, so it's not possible to find an approximation algorithm (using only this linear programming relaxation) that approximates better than within a factor of $g = 2$.

3.2.1 Integrality Gap for Vertex Cover

Claim 6 *The integrality gap for vertex cover is at least $2 - \frac{2}{n}$.*

Proof Consider the complete graph on n vertices. We have $OPT = n - 1$, because if there are 2 vertices that we don't choose, the edge between them is not covered. However, $LP_{OPT} \leq \frac{n}{2}$, because assigning $\frac{1}{2}$ to each vertex is a feasible solution of cost $\frac{n}{2}$, so the optimum can only be smaller. Now,

$$g \geq \frac{n-1}{\frac{n}{2}} = 2 - \frac{2}{n}.$$

■

We remark that our 2-approximation algorithm for vertex cover implies that the integrality gap is *at most* 2.

4 Set Cover via Randomized Rounding

Let us now apply the framework to the *Set Cover* problem. It can be seen as a generalization of the vertex cover problem and its definition is as follows:

Definition 7 (Set Cover Problem) *Given a universe $\mathcal{U} = \{e_1, e_2, \dots, e_n\}$, and a family of subsets $\mathcal{T} = \{S_1, S_2, \dots, S_m\}$ and a cost function $c : \mathcal{T} \rightarrow \mathbb{R}_+$, find a collection C of subsets of minimum cost that cover all elements.*

As for vertex cover, we start by giving an exact Integer LP formulation. For each $i \in \{1, \dots, m\}$, define x_i , which is 1 if $S_i \in C$, and 0 otherwise. The objective function is

$$\min \sum_{i=1}^m x_i \cdot c(S_i)$$

and for each element $e \in \mathcal{U}$, we add the constraint $\sum_{S_i : e \in S_i} x_i \geq 1$. This ensures that each element is covered by at least one set in C . And for each x_i , we require that $x_i \in \{0, 1\}$ in the ILP. The LP relaxation is then obtained by replacing the boolean constraints $x_i \in \{0, 1\}$ by $x_i \in [0, 1]$.

Now suppose that each element belongs to at most f sets. Then, as in your exercise on vertex cover on k -uniform hypergraphs, we can do the following rounding: $C = \{S_i : x_i^* \geq \frac{1}{f}\}$. In each constraint, there's at least one x_i^* which is at least $\frac{1}{f}$, so each constraint is satisfied. Using the same reasoning as in the analysis of the vertex cover rounding, we can show that this approximation is within a factor of f .

4.1 A better approximation for Set Cover

If we introduce randomness and allow our algorithm to output non-feasible solutions with some small probability, we can get much better results (in expectation).

We use the same LP as in the previous section, and will run the following algorithm:

1. Solve the LP to get an optimal solution x^* .
2. Choose some positive integer constant d (we will see later how d affects the guarantees we get). Start with an empty result set C , and repeat step 3 $d \cdot \ln(n)$ times.
3. For $i = 1, \dots, m$, add set S_i to the solution C with probability x_i^* , choosing independently for each set.

Now let us analyze what guarantees we can get:

Claim 8 *The expected cost of all sets added in one execution of Step 3 is*

$$\sum_{i=1}^m x_i^* c(S_i) = LP_{OPT}$$

Proof

$$\mathbb{E}[\text{rounded cost}] = \sum_{i=1}^m c(S_i) \Pr[S_i \text{ is added}] = \sum_{i=1}^m c(S_i) x_i^* = LP_{OPT}$$

■

From this, we can immediately derive

Corollary 9 *The expected cost of C after $d \cdot \ln(n)$ executions of Step 3 is at most*

$$d \cdot \ln(n) \cdot \sum_{i=1}^m c(S_i) x_i^* \leq d \cdot \ln(n) \cdot LP_{OPT} \leq d \cdot \ln(n) \cdot OPT$$

Note that we have $LP_{OPT} \leq OPT$ because LP is a relaxation of the original problem, so its optimum can only be better.

That sounds good, but we should also worry about feasibility:

Claim 10 *The probability that a constraint remains unsatisfied after a single execution of Step 3 is at most $\frac{1}{e}$.*

Proof Suppose our constraint contains k variables, and let us write it as $x_1 + x_2 + \dots + x_k \geq 1$. Then,

$$\begin{aligned} \Pr[\text{constraint unsat.}] &= \Pr[S_1 \text{ not taken}] \dots \Pr[S_k \text{ not taken}] \\ &= (1 - x_1^*) \dots (1 - x_k^*) \\ &\leq e^{-x_1^*} \dots e^{-x_k^*} \end{aligned} \tag{1}$$

$$\begin{aligned} &= e^{-\sum_{i=1}^k x_i^*} \\ &\leq e^{-1} \end{aligned} \tag{2}$$

where (1) follows from the inequality $1 - x \leq e^{-x}$ and (2) from the fact that $\sum_i x_i^* \geq 1$. ■

Claim 11 *The output C is a feasible solution with probability at least $1 - \frac{1}{n^{d-1}}$.*

Proof Using claim 10, we find that the probability that a given constraint is unsatisfied after $d \cdot \ln(n)$ executions of step 3 is at most

$$\left(\frac{1}{e}\right)^{d \cdot \ln(n)} = \frac{1}{n^d}$$

and by union-bound, the probability that there exists any unsatisfied constraint is at most

$$n \cdot \frac{1}{n^d} = \frac{1}{n^{d-1}}$$

■

Now we have an expected value for the cost, and also a bound on the probability that an infeasible solution is output, but we still might have a bad correlation between the two: It could be that all feasible outputs have a very high cost, and all infeasible outputs have a very low cost.

The following claim deals with that worry.

Claim 12 *The algorithm outputs a feasible solution of cost at most $4d \ln(n) OPT$ with probability greater than $\frac{1}{2}$.*

Proof Let μ be the expected cost, which is $d \ln(n) \cdot OPT$ by corollary 9. We can upper-bound the bad event that the actual cost is very high: By Markov's inequality, we have $\Pr[\text{cost} > 4\mu] \leq \frac{1}{4}$. The other bad event that we have to upper bound is that the output is infeasible, and by claim 11, we know that this happens with probability at most $\frac{1}{n^{d-1}} \leq \frac{1}{n}$. Now in the worst case, these two bad events are completely disjoint, so the probability that no bad event happens is at least $1 - \frac{1}{4} - \frac{1}{n}$, and if we suppose that n is greater than 4, this probability is indeed greater than $\frac{1}{2}$. ■

We have thus designed a randomized $O(\log n)$ -approximation algorithm for the set cover problem.

We remark that the used framework has the following general advantage (compared to worst-case guarantees): we can often get better per-instance guarantee than the general approximation factor: Suppose we have an instance where $LP_{OPT} = 100$, and our algorithm found a solution of cost 110. Since we know that $LP_{OPT} \leq OPT$, we can say that our solution on this instance is at most 10% away from the optimal solution for this instance.

References

- [1] Mateusz Golebiewski, Maciej Duleba: *Scribes of Lecture 5 in Topics in TCS 2015*.
<http://theory.epfl.ch/courses/topicstcs/Lecture52015.pdf>