

1 Approximation Algorithms via LP rounding

In last few lectures, we have seen a method of designing approximation algorithms via linear programming, namely, primal-dual scheme. In this lecture, we will look another method of doing the same – rounding of linear programming relaxations. The first step, as usual, is to exactly express our combinatorial problem as an integer program. Then relax the integer program to a linear program. We can find an optimal solution for this linear program by any generic LP solver, for example, the ellipsoid algorithm. However, it is possible that this optimal solution is not integral, and thus, does not correspond to a solution of the combinatorial problem. In that case, we need to **round** this solution x^* to an integer solution $x_{rounded}$. Recall that by design, any integer solution of the LP corresponds to a valid solution of the combinatorial problem. However, the rounded solution might not be an *optimal* integral solution.

To give an approximation guarantee, one needs to bound the gap between the value of the rounded solution $f(x_{rounded})$ and the integer optimal value $f(x_{int}^*)$. It might be hard to bound this gap, as we do not have any clue about the integer optimal value. Instead, we might hope to bound the gap between $f(x^*)$ and $f(x_{rounded})$, since we know that $x_{rounded}$ is obtained from x^* via a rounding procedure. Observe that this gap is at least as large as the gap we desire to bound. Formally, in case of a maximization problem we can say the following for any $\alpha \leq 1$.

$$f(x_{rounded}) \geq \alpha f(x^*) \implies f(x_{rounded}) \geq \alpha f(x_{int}^*)$$

Figure 1 shows the relative order between the three values in the case of a maximization problem.




Figure 1: LP optimal, LP integral optimal and the value of the rounded solution. The red bar shows the gap we want to upper bound and the blue bar shows the gap we are actually going to upper bound.

Rounding Schemes. There is no one fixed way to round a fractional solution to an integral solution. The exact rounding procedure depends on the combinatorial problem at hand. It is often useful to have a randomized rounding procedure, and usually, each coordinate is rounded independently. We list some common rounding schemes.

- For some fixed threshold λ , assign $x_i = 1$ if $x_i^* > \lambda$ and assign $x_i = 0$ otherwise.
- Assign $x_i = 1$ with probability $g(x_i^*)$ and $x_i = 0$ with probability $1 - g(x_i^*)$ for some function $g: [0, 1] \rightarrow [0, 1]$.

In the next part of the lecture, we will talk about the maximum satisfiability problem and an approximation algorithm for it using randomized rounding. Before that let us look at some basics of probability and expectation.

2 Probability and Expectation

Any random experiment has an associated *sample space* Ω , that is the set of possible outcomes of the experiment. For example,

1. A coin toss has two possible outcomes, head or tail, so $\Omega = \{H, T\}$.
2. For two coin tosses, $\Omega = \{HH, HT, TH, TT\}$.
3. If you randomly permute the set of three numbers $\{1, 2, 3\}$, then

$$\Omega = \{(1, 2, 3), (1, 3, 2), (2, 3, 1), (2, 1, 3), (3, 1, 2), (3, 2, 1)\}.$$

Probability is a function $\Pr: \Omega \rightarrow [0, 1]$ from the sample space to real numbers between 0 and 1 with the property that

$$\sum_{x \in \Omega} \Pr(x) = 1.$$

For any $x \in \Omega$, $\Pr(x)$ denotes the chances that the outcome of the random experiment will be x . An *event* is a subset of the sample space, while the elements of the sample space can be called elementary events. The probability of an event $S \subseteq \Omega$ will be $\sum_{x \in S} \Pr(x)$. Let us consider some examples.

- In the second experiment, the event that we have exactly one head is given by $S = \{HT, TH\}$ and has probability $1/4 + 1/4 = 1/2$ (assuming all four outcomes are equi-probable).
- In the third experiment, the event that 2 comes before 3 is given by $S = \{(1, 2, 3), (2, 3, 1), (2, 1, 3)\}$. Assuming that all permutations have an equal probability of $1/6$, the event ‘2 comes before 3’ has probability $1/6 + 1/6 + 1/6 = 1/2$.
- If you toss n unbiased coins, the event that there will be exactly k heads will have probability $\binom{n}{k} 2^{-n}$.

Independent Events. Two events S_1, S_2 are said to be *independent* if occurrence of one does not affect the chances of other one occurring. Let us formalize this a bit. By $\Pr(S_1 | S_2)$ we would mean the probability of S_1 given that S_2 occurred. To compute this, one can imagine S_2 as the sample space and scale the probabilities uniformly so as to ensure that $\sum_{x \in S_2} \Pr(x) = 1$. This way, we can derive the following.

$$\Pr(S_1 | S_2) = \frac{\Pr(S_1 \cap S_2)}{\Pr(S_2)}. \quad (1)$$

S_1 and S_2 are said to be independent if

$$\Pr(S_1) = \Pr(S_1 | S_2).$$

The above equation is equivalent to $\Pr(S_2) = \Pr(S_2 | S_1)$. Consider an example. Suppose we pick a random permutation of four elements $\{1, 2, 3, 4\}$ where each permutation has probability $1/24$. Consider three events.

- S_1 : 1 comes before 2.
- S_2 : 1 comes before 3.
- S_3 : 3 comes before 4.

Which pair among S_1, S_2, S_3 is independent and which pair is not? First let us compute the probabilities of these events. If you go over all 24 permutations, you will find that exactly 12 of them have 1 before 2. Same is true for the other two events. Hence,

$$\Pr(S_1) = \Pr(S_2) = \Pr(S_3) = 12/24 = 1/2.$$

Now, let us find the conditional probabilities $\Pr(S_1 | S_3)$ and $\Pr(S_2 | S_3)$. For these, we need to find $\Pr(S_1 \cap S_3)$ and $\Pr(S_2 \cap S_3)$. You can go over all permutations and verify that the permutations where ‘1 comes before 2’ and ‘3 comes before 4’ are

$$S_1 \cap S_3 = \{(1, 2, 3, 4), (1, 3, 2, 4), (1, 3, 4, 2), (3, 1, 2, 4), (3, 1, 4, 2), (3, 4, 1, 2)\}.$$

Also verify that the permutations where ‘1 comes before 3’ and ‘3 comes before 4’ are

$$S_2 \cap S_3 = \{(1, 2, 3, 4), (1, 3, 2, 4), (1, 3, 4, 2), (2, 1, 3, 4)\}.$$

Thus,

$$\Pr(S_1 \cap S_3) = 6/24 = 1/4 \text{ and } \Pr(S_2 \cap S_3) = 4/24 = 1/6.$$

By (1), we can write

$$\Pr(S_1 | S_3) = (1/4)/(1/2) = 1/2 \text{ and } \Pr(S_2 | S_3) = (1/6)/(1/2) = 1/3.$$

Thus, we conclude: S_1 and S_3 are independent as $\Pr(S_1) = \Pr(S_1 | S_3)$ and S_2 and S_3 are not independent as $\Pr(S_2) \neq \Pr(S_2 | S_3)$.

Random Variable. A random variable is a variable associated with a random experiment. Formally, a random variable $v: \Omega \rightarrow \mathbb{R}$ is a function from the sample space to real numbers (for example). For any $\alpha \in \mathbb{R}$, the probability that v takes value α is given by

$$\Pr(v = \alpha) = \sum_{x: v(x) = \alpha} \Pr(x).$$

Similarly,

$$\Pr(v \geq \alpha) = \sum_{x: v(x) \geq \alpha} \Pr(x).$$

For example, if you toss n coins and define v to be the number of heads obtained then

$$\Pr(v = k) = \binom{n}{k} / 2^n.$$

Expectation. Expectation of a random variable is the weighted average of the values it takes with probabilities being weights.

$$\mathbb{E}[v] = \sum_{x \in \Omega} v(x) \Pr(x) = \sum_{\alpha} \alpha \Pr(v = \alpha).$$

Intuitively, if you repeat a random experiment a large number of times, the average value of the random variable will be its expectation.

Linearity of Expectation. For two random variables v_1, v_2 , the expectation of the random variable $v_1 + v_2$ is same as the sum of the expectations of v_1 and v_2 , i.e.,

$$\mathbb{E}[v_1 + v_2] = \mathbb{E}[v_1] + \mathbb{E}[v_2].$$

More generally for k random variables,

$$\mathbb{E} \left[\sum_{i=1}^k v_i \right] = \sum_{i=1}^k \mathbb{E}[v_i].$$

We leave the proof of linearity of expectation as an exercise.

Linearity of expectation proves to be quite useful in simplifying the computation of expectation. Consider an example. Suppose we pick a uniformly random permutation σ of $(1, 2, \dots, n)$. Let v be the number of consecutive pairs in the permutation. That is,

$$v = |\{i : \sigma(i+1) - \sigma(i) = 1\}|.$$

For example, for $n = 5$, $v(1, 3, 2, 4, 5) = 1$ and $v(1, 4, 5, 2, 3) = 2$. We want to compute the expectation of v . By definition,

$$\mathbb{E}[v] = \sum_{\sigma} (1/n!) v(\sigma),$$

where the sum is over all permutations of $(1, 2, \dots, n)$. This sum looks quite cumbersome to compute as it is not obvious how many permutations have exactly k consecutive pairs for some given k . Let us see how linearity of expectation makes this computation easy. For each $1 \leq j \leq n-1$, define a random variable v_j to be the number of times the particular pair $(j, j+1)$ appears consecutively. Clearly, v_j can take value 0 or 1. Observe that v is related to v_j 's as follows

$$v = \sum_{j=1}^{n-1} v_j.$$

By linearity of expectation

$$\mathbb{E}[v] = \sum_{j=1}^{n-1} \mathbb{E}[v_j]. \quad (2)$$

Now, we just need to compute $\mathbb{E}[v_j]$ for each j . By definition,

$$\mathbb{E}[v_j] = 1 \cdot \Pr(v_j = 1) + 0 \cdot \Pr(v_j = 0).$$

What is the probability that v_j is 1? This is the probability that j and $j+1$ appear together in the order $(j, j+1)$. The number of permutations where $(j, j+1)$ appear consecutively is $(n-1)!$. To see this, imagine $(j, j+1)$ together as one element, then you will have all permutations of $n-1$ elements. Thus, we get

$$\mathbb{E}[v_j] = \Pr(v_j = 1) = \frac{(n-1)!}{n!} = \frac{1}{n}.$$

Putting this in (2),

$$\mathbb{E}[v] = (n-1)/n.$$

Exercise. If you toss a coin n times, what is the expected number of two consecutive heads?