

Online Set Cover: Introduction and LP Duality

Thomas Kesselheim

Last Update: October 18, 2025

Today, we will learn about a fundamental technique in the design of online algorithms. As our motivating example, we consider the set cover problem in its weighted variant. In the offline version, you are given a universe of m elements $U = \{1, \dots, m\}$ and a family of n subsets of U called $\mathcal{S} \subseteq 2^U$. For each $S \in \mathcal{S}$, there is a cost c_S . Your task is to find a *cover* $\mathcal{C} \subseteq \mathcal{S}$ of minimum cost $\sum_{S \in \mathcal{C}} c_S$. A set \mathcal{C} is a cover if for each $e \in U$ there is an $S \in \mathcal{C}$ such that $e \in S$. Alternatively, you could say $\bigcup_{S \in \mathcal{C}} S = U$.

We assume that each element of U is included in at least one $S \in \mathcal{S}$. So in other words \mathcal{S} is a feasible cover. Otherwise, there might not be a feasible solution.

Note that the problem is NP-hard in the offline case, so this already limits our expectations. We will consider the online version, in which the universe U arrives online, one element at a time. Whenever an element is revealed, we get to know which sets $S \in \mathcal{S}$ it is contained in and have to make sure that it is covered, potentially by adding a set from \mathcal{S} to \mathcal{C} . We may never remove sets from \mathcal{C} . Our goal is to eventually select sets so as to minimize $\sum_{S \in \mathcal{C}} c_S$.

Example 3.1. *A special case is the ski-rental problem. As a simplification, we can assume that every day is a skiing day but we do not know the overall number of days. We could capture this by setting $\mathcal{S} = \{\{1\}, \dots, \{m\}, \{1, \dots, m\}\}$, so each element can be covered individually (which means renting the skis for this particular day) or all can be covered simultaneously (which means buying the skis). The costs are set to $c_{\{e\}} = 1$ for all $e \in U$ and $c_U = B$.*

1 LP Relaxation

We can state the set cover problem as an integer program as follows

$$\begin{aligned}
 & \text{minimize } \sum_{S \in \mathcal{S}} c_S x_S && \text{(minimize the overall cost)} \\
 & \text{subject to } \sum_{S: e \in S} x_S \geq 1 && \text{for all } e \in U \quad \text{(cover every element at least once)} \\
 & x_S \in \{0, 1\} && \text{for all } S \in \mathcal{S} \quad \text{(every set is either in the set cover or not)}
 \end{aligned}$$

We can relax the problem by exchanging the constraints $x_S \in \{0, 1\}$ by $0 \leq x_S \leq 1$. (These are the only constraints requiring integrality of the solution.) We get the following LP relaxation¹

$$\begin{aligned}
 & \text{minimize } \sum_{S \in \mathcal{S}} c_S x_S \\
 & \text{subject to } \sum_{S: e \in S} x_S \geq 1 && \text{for all } e \in U \\
 & x_S \geq 0 && \text{for all } S \in \mathcal{S}
 \end{aligned}$$

Also the online problem has a fractional relaxation: We know the variables and the objective function in advance. We get to know one constraint at a time, have to maintain a feasible solution and are not allowed to reduce the values of the variables. So the difficulty is that we do not know what constraints will come later when we choose which variables to increase.

¹We could also include that $x_S \leq 1$ for all S but this will not change the optimal solution as values greater than 1 do not make sense.

Example 3.2. *Let us come back to the ski rental problem. Its LP relaxation is (for readability we rename the variables)*

$$\begin{aligned} & \text{minimize } Bx_{\text{buy}} + \sum_t x_{\text{rent},t} \\ & \text{subject to } x_{\text{buy}} + x_{\text{rent},t} \geq 1 && \text{for all } t \\ & x_{\text{buy}}, x_{\text{rent},t} \geq 0 && \text{for all } t \end{aligned}$$

A fractional solution might tell us to buy half the skis and to rent the other half. While this makes no sense at first sight, we can interpret such values later as probabilities.

Indeed, we will first consider the fractional problem and devise an algorithm for it and then derive a feasible integral solution from the fractional one.

2 LP Duality

We will use LP duality for our algorithm. It is not necessary to know LP duality in its generality. What we need to know is that the dual LP gives us a *lower bound* on all feasible solutions. Let us start with a simple example.

Example 3.3. *Consider $U = \{1, 2, 3\}$, $\mathcal{S} = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$, $c_S = 1$ for all $S \in \mathcal{S}$. The optimal set cover solution has cost 2 because we need to take two sets. However, setting $x_{\{1,2\}} = x_{\{1,3\}} = x_{\{2,3\}} = \frac{1}{2}$ for all $S \in \mathcal{S}$ is a feasible solution to the LP relaxation of cost $\frac{3}{2}$. There is no cheaper solution than this: From the three constraints, we get that for any feasible x*

$$x_{\{1,2\}} + x_{\{1,3\}} + x_{\{2,3\}} = \frac{1}{2}(x_{\{1,2\}} + x_{\{1,3\}}) + \frac{1}{2}(x_{\{1,2\}} + x_{\{2,3\}}) + \frac{1}{2}(x_{\{1,3\}} + x_{\{2,3\}}) \geq \frac{3}{2}.$$

The point of LP duality is to find coefficients just as $\frac{1}{2}, \frac{1}{2}, \frac{1}{2}$ above to derive such a lower bound on the LP value. Interestingly, possible choices for such coefficients can again be found as the solution of a linear program, which is then called the *dual LP*. The dual of the Set Cover LP relaxation is

$$\begin{aligned} & \text{maximize } \sum_{e \in U} y_e \\ & \text{subject to } \sum_{e \in S} y_e \leq c_S && \text{for all } S \in \mathcal{S} \\ & y_e \geq 0 && \text{for all } e \in U \end{aligned}$$

As always in LP duality, we get a dual variable for each primal constraint and a dual constraint for each primal variable. In this case, this means that the dual LP has one variable for each element and one constraint for each set. A usual interpretation, which is also helpful here, is that dual variables correspond to “how expensive” it is to cover a certain element.

Lemma 3.4 (Weak Duality). *Let x and y be feasible solutions to the primal and dual program respectively. Then $\sum_{S \in \mathcal{S}} c_S x_S \geq \sum_{e \in U} y_e$.*

Proof. We have $\sum_{e \in U} y_e \leq \sum_{e \in U} (\sum_{S: e \in S} x_S) y_e = \sum_{S \in \mathcal{S}} x_S \sum_{e \in S} y_e \leq \sum_{S \in \mathcal{S}} x_S c_S$. \square

Example 3.5. *The dual of the ski rental LP is*

$$\begin{aligned} & \text{maximize} && \sum_t y_t \\ & \text{subject to} && \sum_t y_t \leq B \\ & && y_t \leq 1 && \text{for all } t \\ & && y_t \geq 0 && \text{for all } t \end{aligned}$$

So, these constraints tell us that the optimal solution does not spend more than B in total and not more than 1 per day.

3 Example Algorithm: Ski Rental

Let us understand the notions that we have just defined in the light of ski rental. Consider the following algorithm. It does (more or less) the same as our initial deterministic algorithm.

- If $x_{\text{buy}} < 1$
 - Increase x_{buy} by $\frac{1}{B}$, $x_{\text{rent},t}$ by 1.
 - Set $y_t = 1$.
- Else: Set $y_t = 0$.
- If $x_{\text{buy}} = 1$, buy the skis (unless already done so), otherwise rent them.

This algorithm already follows the approach that we will use for general Set Cover. The idea is to maintain (fractional) vectors x and y and to update them in every step. Entries in x can only be increased; in y only the variable that arrived in this step can be set.

The x vector is not a feasible solution to our actual problem because it is fractional. Therefore, we still have to derive an integral decision, which in this case is straightforward in the last step. Note that the cost of actual buying and renting is clearly bounded by $Bx_{\text{buy}} + \sum_t x_{\text{rent},t}$. So, it is enough if we only talk about the latter quantity.

Theorem 3.6. *Let x^* be any feasible (possibly fractional) solution to the ski rental LP, let x be the one computed by the algorithm. Then $Bx_{\text{buy}} + \sum_t x_{\text{rent},t} \leq 2(Bx_{\text{buy}}^* + \sum_t x_{\text{rent},t}^*)$.*

Proof. Let $x^{(t)}$ be the state of x after the t -th step, $x^{(0)} = 0$. Observe that $B(x_{\text{buy}}^{(t)} - x_{\text{buy}}^{(t-1)}) + \sum_{t'} (x_{\text{rent},t'}^{(t)} - x_{\text{rent},t'}^{(t-1)}) \leq 2y_t$ for all t . So, by a telescoping sum

$$Bx_{\text{buy}} + \sum_t x_{\text{rent},t} \leq 2 \sum_t y_t .$$

Furthermore, observe that y is dual feasible: No individual entry is bigger than 1, the sum of all entries is no bigger than B . Therefore, by Lemma 3.4, we have

$$Bx_{\text{buy}}^* + \sum_t x_{\text{rent},t}^* \geq \sum_t y_t .$$

Combining the two inequalities, the theorem follows. □

4 LP Duality for Online Algorithms

Like the example algorithm for ski rental, our algorithm for fractional online Set Cover will use LP duality. We will not only maintain (feasible) primal solutions $x^{(t)}$ but also construct a (possibly infeasible) dual solution y . Both start from $x^{(0)} = 0$ and $y = 0$. In step t , we will increase the primal variables from $x^{(t-1)}$ to $x^{(t)}$ and set the dual variable y_t . Furthermore, the primal increase will be bounded by the dual increase as in the following lemma.

Lemma 3.7. *If*

(a) *in every step t the primal increase is bounded by β times the dual increase, that is*

$$P^{(t)} - P^{(t-1)} \leq \beta y_t, \text{ where } P^{(t)} = \sum_{S \in \mathcal{S}} c_S x_S^{(t)}$$

(b) $\frac{1}{\gamma} y$ *is dual feasible,*

then the algorithm is $\beta\gamma$ -competitive.

Proof. First, observe that by a telescoping-sum argument, we have $P^{(t)} = \sum_{t'=1}^t (P^{(t')} - P^{(t'-1)}) \leq \beta \sum_{t'=1}^t y_{t'}$.

Let x^* be an optimal offline solution. Then, by weak duality, we know $\sum_{S \in \mathcal{S}} c_S x_S^* \geq \sum_{e \in U} \tilde{y}_e$ for any dual feasible \tilde{y} , in particular $\tilde{y} = \frac{1}{\gamma} y^{(t)}$. So, $\sum_{S \in \mathcal{S}} c_S x_S^* \geq \frac{1}{\gamma} \sum_{e \in U} y_e$.

Combined with $P^{(t)} \leq \beta \sum_{t'=1}^t y_{t'}$, we get $\sum_{S \in \mathcal{S}} c_S x_S^{(t)} \leq \beta \cdot \sum_{t'=1}^t y_{t'} \leq \beta\gamma \sum_{S \in \mathcal{S}} c_S x_S^*$. This means exactly that the online solution $x^{(t)}$ is within an $\beta\gamma$ factor of the offline solution x^* . \square

We have already seen one example to apply this lemma: Our primal-dual algorithm for Ski Rental fulfills the conditions with $\beta = 2$ and $\gamma = 1$. Next time, we will see a smarter way to increase the variables so that β is smaller.

References

- N. Buchbinder, J. Naor: The Design of Competitive Online Algorithms via a Primal-Dual Approach. Foundations and Trends in Theoretical Computer Science 3(2-3): 93-263 (2009)