

Online Set Cover: Randomized Rounding and Extensions

Thomas Kesselheim

Last Update: October 28, 2025

So far, we have seen an algorithmic approach to solve the fractional version of Online Set Cover. That is, we are presented the LP

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

one constraint at a time. Whenever we see a constraint, we have to increase the variables to ensure it is fulfilled.

So far, the variables could take fractional values. By design, these values never exceed 1. In this lecture, we will see how the fractional solutions can be turned into integral ones. We will do this in a randomized way by interpreting fractional values as probabilities.

1 Ski Rental

Our first result is again for the *ski-rental problem*. Recall that we derived a primal-dual algorithm for the fractional problem, which is given by the following LP.

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

Our algorithm is deterministic and $\frac{e}{e-1}$ -competitive. Our goal now is to turn this algorithm into a randomized algorithm for the integral problem. To this end, we will interpret the fractions as probabilities. For example, for $B = 3$, our algorithm would set $x_{\text{buy}} = \frac{4}{19}$ on the first day, $x_{\text{buy}} = \frac{10}{19}$ on the second day, and $x_{\text{buy}} = 1$ on the third day. This translates to a randomized algorithm that buys skis with probability $\frac{4}{19}$ on the first day, $\frac{10}{19}$ within the first two days, and 1 within the first three days.

The easiest way to formalize this is as follows. We draw θ uniformly at random from $[0, 1]$ beforehand, buy the skis as soon as $x_{\text{buy}}^{(t)} \geq \theta$, and rent them before. Interestingly, this approach works for any deterministic algorithm for the fractional problem.

Theorem 5.1. *Given an α -competitive algorithm for the fractional relaxation of Ski Rental, the above algorithm is a randomized α -competitive algorithm for Ski Rental.*

Proof. Let Z_{buy} be the random variable indicating the cost of buying skis within the first m steps. We do so with probability $x_{\text{buy}}^{(m)}$. The expected cost from buying in the first m steps is $\mathbf{E}[Z_{\text{buy}}] = B \cdot x_{\text{buy}}^{(m)}$.

Let $Z_{\text{rent},t}$ be the random variable indicating the cost of renting skis in the t -th step. We rent skis with probability $1 - x_{\text{buy}}^{(t)}$. Note that $x_{\text{buy}}^{(t)} + x_{\text{rent},t}^{(t)} \geq 1$ because this is the new constraint

that has to be fulfilled. In other words, we rent skis with probability at most $x_{\text{rent},t}^{(t)}$. So,

$$\mathbf{E} [Z_{\text{rent},t}] \leq x_{\text{rent},t}^{(t)} \leq x_{\text{rent},t}^{(m)}$$

Overall, our expected cost is

$$\mathbf{E} \left[Z_{\text{buy}} + \sum_{t=1}^m Z_{\text{rent},t} \right] \leq B \cdot x_{\text{buy}}^{(m)} + \sum_{t=1}^m x_{\text{rent},t}^{(m)} \leq \alpha \left(B \cdot x_{\text{buy}}^* + \sum_{t=1}^m x_{\text{rent},t}^* \right) \leq \alpha \cdot \text{cost}(\text{OPT}(\sigma)) ,$$

where x^* is an optimal fractional offline solution. □

2 Rounding Fractional Solutions for Set Cover

Our algorithm for general Set Cover is a little more complicated but follows the same principle.

A priori, we choose for each $S \in \mathcal{S}$ a threshold θ_S uniformly at random from $[0, 1]$.

Upon arrival of an element e , update $(x_S^{(t)})_{S \in \mathcal{S}}$ according to the fractional algorithm.

- (a) Pick all sets S for which $x_S^{(t)} \geq \frac{1}{2 \ln t} \theta_S$.
- (b) If e is still uncovered, choose one set S from the probability distribution defined by $(\frac{1}{W} x_S^{(t)})_{S: e \in S}$ and pick it, where $W = \sum_{S: e \in S} x_S^{(t)}$. Note that these probabilities sum up to 1 and $W = \sum_{S: e \in S} x_S^{(t)} \geq 1$.

Note that our algorithm might choose more sets than it has to in part (a). For example, multiple sets may be picked in a step and even sets that do not contain the new element. However, this variant of the algorithm is easier to describe and analyze.

Theorem 5.2. *Using an α -competitive algorithm for the fractional problem, the algorithm for the integral problem is $O(\alpha \cdot \log m)$ -competitive.*

Let's define a bit of notation. We let $Y_{S,t} = 1$ if $x_S^{(t)} \geq \frac{1}{2 \ln t} \theta_S$ and 0 otherwise. Furthermore, we let $Z_{S,t} = 1$ if set S is chosen due to part (b) of step t .

Observe that if $Y_{S,t} = 1$, then set S is picked in one of the steps $1, \dots, t$ because the threshold is exceeded in step t the latest.

Lemma 5.3. *The probability that part (b) is executed in the t -th step is at most $\frac{1}{t^2}$.*

Proof. Note that for part (b) to be executed in the t -th step it is necessary that $Y_{S,t} = 0$ for all S with $e \in S$ because otherwise the respective set would have been picked and the element is covered.

Let's consider any fixed such set S . By design, if $(2 \ln t)x_S^{(t)} \leq 1$,

$$\Pr [Y_{S,t} = 1] = \Pr \left[x_S^{(t)} \geq \frac{1}{2 \ln t} \theta_S \right] = \Pr \left[\theta_S \leq (2 \ln t)x_S^{(t)} \right] = (2 \ln t)x_S^{(t)} .$$

As $1 - q \leq \exp(-q)$ for $q \in [0, 1]$, this implies

$$\Pr [Y_{S,t} = 0] = 1 - (2 \ln t)x_S^{(t)} \leq \exp(-(2 \ln t)x_S^{(t)}) .$$

If $(2 \ln t)x_S^{(t)} > 1$, then $\Pr [Y_{S,t} = 0] = 0$, so this bound holds as well.

Note that these choices are independent, so

$$\Pr \left[\bigwedge_{S: e \in S} Y_{S,t} = 0 \right] = \prod_{S: e \in S} \Pr [Y_{S,t} = 0] \leq \prod_{S: e \in S} \exp(-(2 \ln t)x_S^{(t)}) = \exp \left(-(2 \ln t) \sum_{S: e \in S} x_S^{(t)} \right) \leq \frac{1}{t^2} .$$

□

Lemma 5.4. *The expected cost due to set S within the first m rounds is at most $((2 \ln m) + \sum_{t=1}^m \frac{1}{t^2}) c_S x_S^{(m)}$.*

Proof. Observe that if $Y_{S,t} = 1$, then also $Y_{S,t'} = 1$ for $t' > t$ because $x_S^{(t')} \geq x_S^{(t)}$ and $\frac{1}{2 \ln t'} \leq \frac{1}{2 \ln t}$. Therefore if set S is picked in the first m rounds due to part (a), then $Y_{S,m} = 1$.

Recall our notation that $Z_{S,t} = 1$ if set S is chosen in part (b) statement of step t . We have $\mathbf{E}[Z_{S,t}] = \mathbf{Pr}[Z_{S,t} = 1] \leq \frac{1}{t^2} \cdot \frac{1}{W} \cdot x_S^{(t)} \leq \frac{1}{t^2} \cdot x_S^{(m)}$ because when arriving in part (b)—which happens with probability no more than $\frac{1}{t^2}$ —the set is picked with probability no more than $x_S^{(t)} \leq x_S^{(m)}$.

We have that set S is picked only if $Y_{S,m} + \sum_{t=1}^m Z_{S,t} \geq 1$ and

$$\begin{aligned} \mathbf{Pr} \left[Y_{S,m} + \sum_{t=1}^m Z_{S,t} \geq 1 \right] &\leq \mathbf{E} \left[Y_{S,m} + \sum_{t=1}^m Z_{S,t} \right] = \mathbf{E}[Y_{S,m}] + \sum_{t=1}^m \mathbf{E}[Z_{S,t}] \\ &\leq (2 \ln m) x_S^{(m)} + \sum_{t=1}^m \frac{1}{t^2} x_S^{(m)}. \end{aligned}$$

Therefore, expected cost due to set S within the first m rounds is at most

$$c_S \mathbf{Pr} \left[Y_{S,m} + \sum_{t=1}^m Z_{S,t} \geq 1 \right] \leq \left((2 \ln m) + \sum_{t=1}^m \frac{1}{t^2} \right) c_S x_S^{(m)}. \quad \square$$

Proof of Theorem 5.2. We consider the outcome after m rounds. Let x^* be an optimal offline fractional solution to the LP relaxation. As the fractional algorithm is α -competitive, we have $\sum_{S \in \mathcal{S}} c_S x_S^{(m)} \leq \alpha \cdot \sum_{S \in \mathcal{S}} c_S x_S^*$. Furthermore, by Lemma 5.4, $\mathbf{E}[\text{cost}(\text{ALG}(\sigma))] \leq \sum_{S \in \mathcal{S}} \left((2 \ln m) + \sum_{t=1}^m \frac{1}{t^2} \right) c_S x_S^{(m)}$.

Note that $\sum_{t=1}^{\infty} \frac{1}{t^2} = \frac{\pi^2}{6}$. So overall, $\mathbf{E}[\text{cost}(\text{ALG}(\sigma))] \leq \alpha \cdot \left((2 \ln m) + \frac{\pi^2}{6} \right) \sum_{S \in \mathcal{S}} c_S x_S^* = \alpha \cdot \left((2 \ln m) + \frac{\pi^2}{6} \right) \text{cost}(\text{OPT}(\sigma)). \quad \square$

3 Extensions

The primal-dual framework for Online Set Cover can be extended in various ways, some of which we will highlight here.

3.1 Arbitrary Covering Constraints

It is pretty straightforward to extend the approach to arbitrary covering constraints. That is, we have a covering LP of the form

$$\begin{aligned} &\text{minimize} && \sum_j c_j x_j \\ &\text{subject to} && \sum_j a_{i,j} x_j \geq 1 && \text{for all } i \\ &&& x_j \geq 0 && \text{for all } j \end{aligned}$$

where all c_j and $a_{i,j}$ are non-negative reals.

3.2 Packing LPs

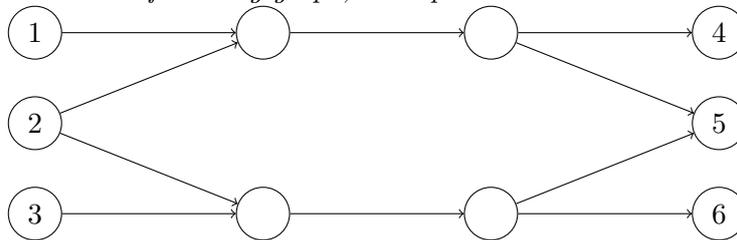
We can also consider packing problem, which are given by the following kind of linear program.

$$\begin{aligned}
 & \text{maximize} && \sum_i y_i \\
 & \text{subject to} && \sum_i a_{i,j} y_i \leq c_j && \text{for all } j \\
 & && y_i \geq 0 && \text{for all } i
 \end{aligned}$$

Now, we know all constraints in advance. In each step a variable appears and we have to decide on its value. Note that this is exactly what we already did with the dual in the covering problem. As it turns out, slight adaptations of our Set Cover algorithm give us an $O(\log n)$ -competitive algorithm for the fractional problem with n packing constraints if $a_{i,j} \in \{0, 1\}$. If each variable appears in at most f constraints, this can be improved to $O(\log f)$ because this corresponds to the frequency of the set system of the Set Cover instance.

One particular application is the following. We are given a directed graph $G = (V, E)$, in which each edge $e \in E$ has a capacity c_e . Now, we are presented pairs of a source node $s_k \in V$ and a sink node $t_k \in V$ one after the other. For each of these pairs (s_k, t_k) , we have to find a path connecting s_k and t_k or we can reject it. The number of paths using edge e may not exceed c_e . The goal is to connect as many source-sink pairs as possible.

Example 5.5. Consider the following graph, all capacities are 1.



In advance, we do not know which pairs of vertices should be connected. If the first request is to connect vertices 2 and 5, we may either choose the top or the bottom path or reject the request. Suppose we choose the top path. If then the second request is to connect vertices 1 and 4, we have to reject the request because there is no more capacity left. So, we served only one request whereas the offline optimum can serve both.

The fractional relaxation of the problem looks as follows. We have variables $f_{k,P}$ denoting the flow routed along path P for every path connecting s_k and t_k .

$$\begin{aligned}
 & \text{maximize} && \sum_k \sum_P f_{k,P} \\
 & \text{subject to} && \sum_{k,P:e \in P} f_{k,P} \leq c_e && \text{for all } e \in E \quad (\text{capacity constraints of edges}) \\
 & && \sum_P f_{k,P} \leq 1 && \text{for all } k \quad (\text{only one unit of flow between any } s_k \text{ and } t_k) \\
 & && f_{k,P} \geq 0 && \text{for all } k, P
 \end{aligned}$$

To design an algorithm for the fractional problem, when we see a pair (s_k, t_k) we add all variables $f_{k,P}$ for paths connecting s_k to t_k to the LP one after the other and run an algorithm for packing LPs.

Note that in this case the number of constraints n corresponds to the number of edges plus the number of rounds. However, this LP is sparse in the following sense: Let d be the maximum length of a path. Then each variable appears in at most $d + 1$ constraints. Therefore, there is an $O(\log d)$ -competitive algorithm for the fractional problem.

For more details and many other applications, see the survey by Buchbinder and Naor.

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)