

Lecture 2,3 (Sept 13 and 15, 2011): Set Cover, Introduction to LP

Lecturer: Mohammad R. Salavatipour

Scribe: based on older notes

2.1 Set Cover Problem

Now we turn our attention to the Set Cover problem, which is (perhaps) the most central problem in the study of approximation algorithms. There are different algorithms for this problem. In this course we will see at least 4 different approximation algorithms for this using different methods.

Set Cover:

- Input:
 - A set of n elements $U = \{e_1, \dots, e_n\}$, called the Universe.
 - A set $S = \{S_1, \dots, S_m\}$ of m subsets of U such that each $e \in U$ is in some $S_i \in S$
 - A cost function $c : S \rightarrow \mathbb{Q}^+$
- Goal: Find a minimum cost subset S' of S such that each $e \in U$ is in some $S_i \in S'$.

Note that vertex cover is a special case of set cover where U is the set of all edges and each vertex v is a subset in S which contains all edges incident to v . In this case, each element is in exactly two subsets in S . We present a greedy approximation algorithm for Set Cover. This is probably the most natural greedy algorithm for this problem. The idea is, at each iteration pick a set where the ratio of the cost of the set divided by the number of new elements it covers is minimized. This general idea of “covering” elements iteratively by finding good partial solutions has been used in many other problems. The analysis of set-cover (we present here) can typically be extended to those other covering algorithms that behave similarly.

Definition 1 Given a subset C of U , define the cost effectiveness of set $S_i \in S$ as $\frac{c(S_i)}{|S_i - C|}$. If $S_i \subseteq C$ then say the cost effectiveness is $+\infty$.

Algorithm SC1

```

C ← ∅
S' ← ∅
while C ≠ U do
  select S_i ∈ S with minimum cost effectiveness α =  $\frac{c(S_i)}{|S_i - C|}$  with respect to C
  for each e ∈ S_i, define price(e) as α
  S' ← S' ∪ {S_i}
  C ← C ∪ S_i
return S'
```

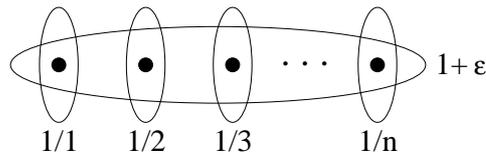


Figure 2.1: A tight example for SC1.

Obviously, all elements are eventually covered by S' since the algorithm terminates only when $C = U$. Note that the final cost of set S' is $\sum_{e \in U} \text{price}(e)$ since, for each $S_i \in S'$, the cost of S_i is distributed among all elements in S_i that were covered for the first time when S_i was picked.

Lemma 1 *Algorithm SC1 is an $\ln n$ -approximation algorithm; more precisely it has ratio at most H_n , where H_n is the n 'th harmonic number.*

Proof. Let $T_{OPT} \subseteq S$ be a set cover with minimum cost OPT . Order the elements of U by the time they were covered by algorithm SC1 (breaking ties arbitrarily) as e_1, e_2, \dots, e_n .

Consider the time just before e_k is covered. The remaining at least $n - k + 1$ elements can be covered at a price of no more than OPT by adding the currently unselected sets of T_{OPT} to S' . In other words, each element can be covered at a price of no more than $\frac{OPT}{n-k+1}$ on average.

We claim that there must be a set with cost effectiveness at most $\frac{OPT}{n-k+1}$. If this were not true, then the cost of covering the remaining uncovered elements would be strictly greater than $(n - k + 1) \cdot \frac{OPT}{n-k+1} = OPT$ which contradicts the fact that the remaining elements can be covered at a cost of at most OPT by selecting T_{OPT} . Thus, $\text{price}(e_k) \leq \frac{OPT}{n-k+1}$ which yields

$$\sum_{k=1}^n \text{price}(e) \leq \sum_{k=1}^n \frac{OPT}{n-k+1} = OPT \cdot \sum_{k=1}^n \frac{1}{k} = OPT \cdot H_n$$

where H_n is the n 'th harmonic number. By comparison with $\int \frac{dx}{x}$ we see that $\ln n \leq H_n \leq \ln n + 1$. Therefore SC1 is an $O(\log n)$ approximation algorithm. ■

Through similar analysis, we can show that SC1 is an $O(\log k)$ -approximation where $k = \max |S_i|$. Note that this proves the ratio of $O(\log \Delta)$ for the greedy vertex cover algorithm where Δ is the size of maximum degree of nodes. The analysis of SC1 is also tight. For any $\epsilon > 0$ being a small constant, consider the following instance of set cover (illustrated in Figure 2.1):

- $U = \{e_1, \dots, e_n\}$
- $S = \{S_0, S_1, \dots, S_n\}$
- $c(S_0) = 1 + \epsilon$ and $c(S_i) = \frac{1}{i}$ for all $1 \leq i \leq n$.

The optimum solution is S_0 with a cost of $1 + \epsilon$ while SC1 returns the solution $\{S_1, \dots, S_n\}$ with a cost of $H_n = OPT \cdot \frac{H_n}{1+\epsilon}$. Since this holds for any small constant $\epsilon > 0$, the analysis is tight (even up to the constant).

Interestingly, the above algorithm is essentially the best possible for set cover.

Theorem 1 (Lund and Yannakakis (92), Feige (96), Raz and Safra (97), Sudan (97))

- Unless $P = NP$, there is no $c \ln n$ approximation algorithm for set cover for some constant $0 < c < 1$.
- Unless $NP \subseteq \text{DTIME}(n^{O(\log \log n)})$, there is no $(1 - \epsilon) \ln n$ approximation for set cover for all $\epsilon > 0$.

2.2 Linear Programming

Definition 2 Linear Programming (or LP) is the problem of optimizing a linear function of variables subject to some linear constraints.

An LP with n variables x_1, \dots, x_n takes the form

$$\begin{aligned} \text{minimize : } & \sum_{i=1}^n c_i x_i \\ \text{subject to : } & \sum_{i=1}^n a_{ji} x_i \geq b_j, \quad \forall 1 \leq j \leq m \\ & x_i \geq 0, \quad \forall 1 \leq i \leq n \end{aligned}$$

for some constants $c_i, a_{ij}, b_j, 1 \leq i \leq n, 1 \leq j \leq m$. Here, the first expression is called the *objective function*. In matrix form, an LP looks like

$$\begin{aligned} \text{minimize : } & x \cdot c \\ \text{subject to : } & x \cdot A \geq b, \quad \forall 1 \leq j \leq m \\ & x \geq 0, \quad \forall 1 \leq i \leq n \end{aligned}$$

where x is an n -dimensional vector of variables. In such a case, A is called the *constraint matrix*. An assignment of values to the variables in x is said to be *feasible* if all the constraints are satisfied.

Consider the following LP with two variables and four constraints.

$$\begin{aligned} \text{minimize : } & x_2 \\ \text{subject to : } & 3x_1 - x_2 \geq 0 \\ & x_1 + x_2 \geq 6 \\ & -x_1 + 2x_2 \geq 0 \\ & x_1 \geq 2 \\ & x_1, x_2 \geq 0 \end{aligned}$$

In this example, the assignment $x_1 = 2$ and $x_2 = 4$ is feasible and has objective function value 4. See figure 2.2 for a visualization of the LP as well as how the following terminology applies to the example.

Each LP falls into exactly one of the following three categories:

- There is a feasible $x \in \mathbb{R}^n$ such that for all feasible $x' \in \mathbb{R}^n$, the objective function value under assignment x' is at least the objective function value under x . We say the LP has a *finite optimum*.
- Every $x \in \mathbb{R}^n$ is not feasible. In this case, we say the LP is *infeasible*.
- For every $z \in \mathbb{R}$, there exists a feasible $x \in \mathbb{R}^n$ such that the objective function value under assignment x exceeds the objective function value under assignment z . Here, the LP is said to be *unbounded*.

Each constraint defines a half-space of \mathbb{R}^n . The feasible region is the intersection of all half-spaces defined by the constraints of the LP. Consider an LP with a non-empty feasible region. If this feasible region is bounded then the LP has a finite optimum which also implies that the feasible region is infinite if the LP is unbounded. Finally, if the LP is infeasible, then the feasible region is empty.

Every solution to an LP is a vector in \mathbb{R}^n . Every constraint corresponds to a half-space. The set of feasible solutions to an LP is *convex*. That is, for all feasible $x, y \in \mathbb{R}^n$ we have $\alpha x + (1 - \alpha)y$ being feasible for each

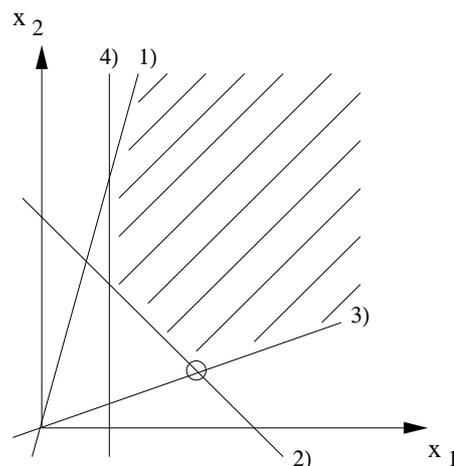


Figure 2.2: A visualization of the example LP. The four lines correspond to the four constraints and are numbered in the order they appear in the example. The feasible region (shaded above) is infinite while the LP has a finite optimum. The optimum value $x_2 = 2$ occurs at the basic feasible solution (circled above) defined by the second and third constraints.

$0 \leq \alpha \leq 1$. Intuitively, all points along the straight line between two feasible points are also feasible. A feasible region for an LP is called a *polyhedron*. If the LP is bounded, then the feasible region is called a *polytope*. A feasible point $x \in \mathbb{R}^n$ is called a *basic feasible solution* or a *vertex* of the polytope if some n constraints are satisfied with equality under assignment x . Geometrically, a basic feasible solution is found where the boundaries of n distinct half-spaces intersect. By definition, a vertex solution is not a convex combination of two or more distinct feasible solution. It can be shown that if an LP has a finite optimum, then some optimum occurs at a basic feasible solution.

A linear program can be solved in polynomial time using, for example, the ellipsoid method or an interior point method. An *oracle* for an LP is a method that decides if a proposed point $x \in \mathbb{R}^n$ is feasible and, if not, produces a violated constraint. The ellipsoid method, while not nearly as practical as interior point methods, has the advantage of solving an LP in polynomial time if there is a polynomial time oracle for the problem even if the number of constraints is exponential.

2.3 Set Cover by deterministic rounding

Let's formulate the Set Cover problem as an integer program and find the LP relaxation. Summarizing the definition of Set Cover we have:

- Input
 - $U = \{e_1, \dots, e_n\}$ is the *universe* of elements.
 - $S = \{S_1, \dots, S_m\}$ is a collection of subsets of U with costs $c(S_1), \dots, c(S_m)$.
- Goal: to find the minimum cost set $S' \subseteq S$ such that the union of all elements of S' give U .

$$\begin{aligned} \text{minimize: } & \sum_{i=1}^n c(S_i) \cdot x_i \\ \text{s.t. } & \sum_{e \in S_j} x_j \geq 1, \quad \forall e_i \in U \quad (1) \\ & x_i \in \{0, 1\} \quad (2) \end{aligned}$$

If we relax constraint (2) to $0 \leq x_i \leq 1$ then we obtain the LP relaxation to the IP. Note that in any IP/LP, the solution to the LP is a lower bound for the IP. Note that, we can remove the condition $x_i \leq 1$ as in any optimal solution there cannot be a value larger than 1. The reason is that in any feasible solution we can round all the $x_i > 1$ to 1 and still get a feasible solution with a lower cost.

A common technique in design of approximation algorithms is to first formulate the problem as an IP/LP, then solve the LP relaxation, and then round this fractional solution to an integer one while keeping a bound on the total value of the objective function. There are different ways to round a fractional solution to an integer one. Here we discuss a deterministic rounding for Set Cover.

We define the frequency of an element $e \in U$ to be the number of the sets of S that contain e and let f be the largest frequency among all elements. For instance, for the case of vertex cover (where the elements are edges and the sets are the vertices) the frequency of each element is 2, so $f = 2$. Below we describe an f -approximation rounding algorithm for set cover.

SC2: Deterministic Rounding for Set Cover

- Let x^* be an optimal solution to the LP relaxation of set-cover.
- For each set i define $\tilde{x}_i = 1$ if $x_i^* \geq 1/f$ and 0 otherwise.
- return the set of sets with $\tilde{x}_i = 1$.

First we claim that this algorithm returns a feasible solution. The reason is that for each element e , constraint (1) implies that among all the at most f sets containing e , at least one set has x^* value at least $1/f$. Therefore, for at least one of those sets \tilde{x} value is 1. To see that this is an f -approximation it is sufficient to observe that for each set S_i picked in our solution, the fractional solution is paying at least $\frac{1}{f} \cdot c(S_i)$ since $x_i^* \geq 1/f$.

The above algorithm when applied to vertex cover implies a 2-approximation even for the case of weighted vertex cover in which every node $v \in V$ has a weight (or cost) c_v and the goal is to find a feasible solution with minimum total cost (the cardinality vertex cover is the special case where $c_v = 1$ for all nodes).

2.4 Integrality Gap

For an instance I of a minimization problem, let $OPT_f(I)$ be the cost of the optimal LP solution. Then, integrality gap of this instance is given by $OPT(I)/OPT_f(I)$. The integrality gap of the problem Π is $\max_I OPT(I)/OPT_f(I)$, i.e. the largest ratio between the fractional and integral solution over all possible instances of Π . Note that in any LP rounding algorithm, we use the LP solution as a lower bound for the optimal solution. Therefore, any such algorithm cannot have a performance ratio better than the integrality gap. In general, it is difficult to design algorithms with ratio better than the integrality gap. Therefore, large integrality gaps typically are indications of difficulty of a problem. Consider for example the Vertex Cover problem, and let I be the complete graph K_n . Clearly by assigning a value of $\frac{1}{2}$ to each vertex we get a fractional solution of value $\frac{n}{2}$ whereas any integer solution must contain at least $n - 1$ nodes (or else one edge is not covered). Thus the integrality gap of V.C. is at least $2 - \frac{2}{n}$.

Exercise. Prove that the integrality gap for Set Cover problem is $\Omega(\log n)$.

2.4.1 Weighted Vertex Cover.

In this section we show that the standard LP for the weighted version of V.C has some other nice features. Recall the following LP relaxation of the vertex cover problem:

$$\begin{aligned} \text{minimize : } & \sum_{i=1}^n c_v \cdot x_v \\ \text{subject to : } & x_u + x_v \geq 1, \quad \forall uv \in E \\ & x_v \geq 0 \end{aligned}$$

Definition 3 (Half-integer) A variable is half-integer, if it is an integer factor of $\{0, \frac{1}{2}, 1\}$.

Our main goal is to prove the following lemma:

Lemma 2 Any basic solution to the above LP for the Vertex Cover is half-integer

It's easy to see that by rounding each $x_i \geq \frac{1}{2}$ to 1 we get a 2-approximation algorithm for V.C.

Proof. In order to prove this lemma, we take advantage of the fact that all non-basic feasible solutions can be written as a convex combination of two other feasible solutions. However, basic solutions cannot be written in such a way. Let x be any basic solution to the LP. Assume x is not half-integer. Define:

$$\begin{aligned} V^> &= \left\{ v \mid \frac{1}{2} < x_v < 1 \right\} \\ V^< &= \left\{ v \mid 0 < x_v < \frac{1}{2} \right\} \end{aligned}$$

Assuming that x is not half-integer, we must have $V^> \cup V^< \neq \emptyset$. We would like to show that there exist feasible solutions y_v and z_v such that x can be written as a convex combination of them, which is a contradiction. Define y_v and z_v as follows:

$$y_v = \begin{cases} x_v + \epsilon & x \in V^> \\ x_v - \epsilon & x \in V^< \\ x_v & \text{Otherwise} \end{cases} \quad \text{and} \quad z_v = \begin{cases} x_v - \epsilon & x \in V^> \\ x_v + \epsilon & x \in V^< \\ x_v & \text{Otherwise} \end{cases}$$

Observation 1. If $V^> \cup V^< \neq \emptyset$, then $y \neq x$ and $z \neq x$.

Observation 2. $x = \frac{1}{2}(y + z)$; i.e. x is written as a combination of y and z .

We can easily make sure that $0 \leq y_v, z_v \leq 1$ by choosing ϵ sufficiently small. In order to prove y and z are both feasible solutions, we consider two cases:

1. For every edge uv with $x_u + x_v > 1$, we can easily choose ϵ so small that we have $(x_u - \epsilon) + (x_v - \epsilon) \geq 1$. In this case, both $y_u + y_v \geq 1$ and $z_u + z_v \geq 1$ since y_u and z_u are both $\geq x_u - \epsilon$ (similarly for y_v and z_v)

2. Now, suppose $x_u + x_v = 1$. In this case one of x_u and x_v belongs to $V^>$ and the other belongs to $V^<$. Assume $x_u \in V^>$ and $x_v \in V^<$. Therefore $\underbrace{(x_u + \epsilon - \epsilon)}_{y_u} + \underbrace{(x_v - \epsilon + \epsilon)}_{y_v} = 1 \Rightarrow y_u + y_v = 1$ and

$$\underbrace{(x_u - \epsilon + \epsilon)}_{z_u} + \underbrace{(x_v + \epsilon - \epsilon)}_{z_v} = 1 \Rightarrow z_u + z_v = 1$$

Therefore, in z, y are both feasible and x (which is a basic solution) is a convex combination of two feasible solution, a contradiction. ■