

Lecture 1 (Sept 5, 2013): Introduction, Vertex Cover

*Lecturer: Mohammad R. Salavatipour**Scribe: Mohammad R. Salavatipour*

1.1 Introduction

In this course we will be studying how to design approximation algorithms for NP-hard optimization problems through a series of problems. Some of our objectives are: to learn some techniques for design and analysis of such algorithms; have a better understanding of problems that are NP-hard and do a (partial) classification based on the level of difficulty of these problems. This is done by study of hardness of approximation for these problems.

Recall that P is the class of problems solvable in polynomial time and NP (informally) are those (decision) problems whose solutions can be verified by a polynomial time algorithm. We will be studying NP-hard optimization problems. For these problems, we would like to:

1. find the optimal solution,
2. find the solution fast, often in polynomial time
3. find solutions for any instance

Unfortunately, with the assumption of $P \neq NP$, we cannot have all the above three at the same time. Therefore, we need to relax at least one of them. If we:

1. relax (3), then we are into study of special cases of the problem.
2. relax (2), we will be in the field of integer programming and the techniques there such as branch-and-bound, etc.
3. relax (1), we are into study of heuristics and approximation algorithms.

We are going to focus on the (3) in this course, and in particular, on the field of approximation algorithms. We are interested in:

- finding solutions that are within a guaranteed factor of the optimal solutions, and
- We want to find this solution fast (which often means polynomial time).

Why is it important to study approximation algorithms?

- We need to solve optimization problems (they appear everywhere!). Saying that because they are NP-hard we don't know of any efficient way of solving them is not enough. If you cannot find the optimal solution, try your best!
- We can prove how good the solution is (this is sometimes critical in applications).

- We can have an understanding of how hard the problem is, and this in turn can help to design better algorithms.
- Tools and algorithmic ideas can often give practical and good heuristics and have applications to other areas.
- Solving very challenging problems using beautiful (sometimes neat and sometimes sophisticated) ideas is fun!

Throughout the course, we often ignore optimizing the running time as long as it is polynomial. We also ignore implementation and engineering aspects or how close the model we study is to the application it came from. Instead we will be focusing on designing algorithms with best possible performance ratio.

Let's start with an example. Consider the Vertex-Cover problem.

Vertex-Cover:

- Input
 - G : an undirected graph $G = (V, E)$
 - c : a cost function on vertices, $c : V \rightarrow \mathbb{Q}^+$
- Goal: find a minimum cost vertex cover, i.e., a set $V' \subseteq V$ such that every edge has at least one endpoint incident at V' .

The special case, in which all vertices are of unit cost, is called the *cardinality vertex cover problem*. Let's consider the cardinality VC for now. Perhaps the most natural greedy algorithm for this problem is the following algorithm:

Algorithm VC1

- $S \leftarrow \emptyset$
- **while** $E \neq \emptyset$ **do**
 - let v be a vertex of maximum degree in G
 - $S \leftarrow S \cup \{v\}$
 - remove v and all its edges from G
- return S

We will see that the approximation ratio for Algorithm VC1 is $O(\log \Delta)$, where Δ is the maximum degree in graph G .

1.2 NP Optimization problems

An NP optimization problem, Π , is a minimization (maximization) problem which consists of the following items:

Valid instances: each valid instance I is recognizable in polynomial time. (note: 'Polynomial time' means polynomial time in terms of the size of the input.) The set of all valid instances is denoted by D_Π . The size of an instance $I \in D_\Pi$, denoted by $|I|$, is the number of bits required to represent I in binary.

Feasible solutions: Each $I \in D_{\Pi}$ has a set $S_{\Pi}(I)$ of feasible solutions and for each solution $s \in S_{\Pi}(I)$, $|s|$ is polynomial (in $|I|$).

Objective function: A polynomial time computable function $f(s, I)$ that assigns a non-negative rational value to each feasible solution s for I .

We often have to find a solution s such that this objective value is minimized (maximized). This solution is called optimal solution for I , denoted by $OPT(I)$.

Examples:

- Vertex Cover: In this case we have:

valid instances : set of graphs with weighted vertices.

feasible solutions : all the vertex covers of the given graph.

objective functions : minimizing the total weight of a vertex cover.

- Minimum Spanning Tree (MST) problem: Given a connected graph $G(V, E)$, with each edge $(u, v) \in E$ assigned a weight $w(u, v)$, find an acyclic subset $T \subseteq E$ that connects all the vertices and its total weight is minimized. Since T is acyclic and connects all of the vertices, it is a tree.

valid instances : a graph with weighted edges.

feasible solutions : all the spanning trees of the given weighted graph.

objective functions : minimizing the total weight of a spanning tree.

1.3 Approximation algorithms

An α -factor approximation algorithm (or simply an α -approximation) is a polynomial time algorithm whose solution is always within α factor of optimal solution.

Definition 1 For a minimization problem Π , algorithm A has approximation factor α if it runs in polynomial time and for any instance $I \in D_{\Pi}$ it produces a solution $s \in S_{\Pi}(I)$ such that $f(s, I) \leq \alpha(|I|) \cdot OPT(I)$. α can be a constant or a function of the size of the instance.

We use $A(I)$ to denote the value of the solution returned by algorithm A for instance I ; therefore, from Definition 1: $A(I) \leq \alpha(|I|) \cdot OPT(I)$.

Similarly, for a maximization problem, we have

- $OPT(I) \leq \alpha(|I|) \cdot A(I)$, if $\alpha > 1$;
- or $OPT(I) \leq \frac{A(I)}{\alpha(|I|)}$, if $\alpha < 1$.

Definition 2 Algorithm A has asymptotic α -approximation ratio if the ratio $\lim_{|I| \rightarrow \infty} \frac{A(I)}{OPT(I)} \leq \alpha$

Intuitively, for large enough instances, the approximation ratio of the algorithm is almost α (although for small instances it might be larger). Just as there are randomized algorithms that compute exact solutions, there are randomized algorithms that compute approximate solutions.

Definition 3 An algorithm A is a randomized factor α -approximation for problem Π if for any instance $I \in D_\Pi$, A produces a feasible solution s such that $\Pr[f(s, I) \leq \alpha(|I|) \cdot \text{OPT}(I)] \geq \frac{1}{2}$.

This probability can be increased to $(1 - \frac{1}{2^x})$ by repeating the same algorithm A for x times.

Definition 4 A (uniform) PTAS (Polynomial Time Approximation Scheme) for an optimization problem is an approximation algorithm A that takes as input not only an instance of the problem, but also a value $\epsilon > 0$, and runs in time polynomial in n and returns a solution s that satisfies: $f(s, I) \leq (1 + \epsilon) \cdot \text{OPT}(I)$, i.e. it is a $(1 + \epsilon)$ -approximation algorithm.

In case of non-uniform PTAS we have a class of algorithms $\{A_\epsilon\}$, one for every possible value of ϵ . A running time $O(n^{\frac{1}{\epsilon}})$ is still polynomial for any fixed ϵ , but computations with ϵ values very close to 1 may turn out to be practically not feasible. This leads to the definition of FPTAS, a more restricted version of PTAS and much faster than PTAS.

Definition 5 FPTAS (Fully Polynomial-Time Approximation Scheme)

For a minimization problem Π , a FPTAS algorithm A takes an instance I and error bound $\epsilon > 0$, and returns a solution s such that $f(s, I) \leq (1 + \epsilon) \cdot \text{OPT}(I)$. A runs in polynomial time in terms of both $|I|$ and ϵ (i.e. ϵ can not appear in exponent).

Example: $O(\frac{n^2}{\epsilon^2})$ vs. $O(n^{\frac{1}{\epsilon}})$

Some problems like Knapsack, Euclidean TSP, and some scheduling problems belong to class PTAS. Problems like MAX-SAT, Vertex-Cover, are much harder and don't admit a PTAS. These problems belong to a class called APX-hard, which is the class of problems that have a constant approximation but don't have a PTAS.

1.4 A 2-Approximation for Cardinality Vertex Cover

Recall the definition of the cardinality vertex cover problem.

- Input: An undirected graph $G = (V, E)$.
- Goal: Find a minimum cardinality set of vertices $S \subseteq V$ such that every edge in E has at least one endpoint in S .

Definition 6 A matching in a graph $G = (V, E)$ is a subset of edges $M \subseteq E$ such that no two edges in M share a common endpoint. A matching M is called maximal if it is not strictly contained in any other matching (i.e. no edges in $E \setminus M$ can be added to M).

Our second vertex cover algorithm simply finds any maximal matching M . The cover is defined as all endpoints of edges in M .

Algorithm VC2

```

S ← ∅
while E ≠ ∅ do

```

```

    let  $uv$  be any edge in  $E$ 
     $S \leftarrow S \cup \{u, v\}$ 
    delete  $u, v$ , and all their incident edges from  $G$ 
return  $S$ 

```

Lemma 1 *The set S returned by VC2 is a vertex cover.*

Proof. Consider any edge e deleted in an iteration of the loop. If e was selected as the edge in the first line of the loop, then both of its endpoints were added to S . Otherwise, e must share an endpoint with the edge selected in the first line of the loop so one of its endpoints is added to S . Since all edges are eventually deleted in some iteration, the final set S is a vertex cover. ■

Lemma 2 *Let M be any matching and S be any vertex cover. Then, $|M| \leq |S|$.*

Proof. Each $e \in M$ must have at least one of its endpoints covered by S . Since M is a matching then no two edges in e share an endpoint. Therefore, each $e \in M$ is covered by a vertex that covers no other $e' \in M$ so $|M| \leq |S|$. ■

Lemma 3 *VC2 is a 2-approximation.*

Proof. Let M be the set of all edges selected in the first line of the loop. Then M is a matching since any edge sharing an endpoint with some $e \in M$ was deleted in the third line of the loop. Since S consists exactly of the endpoints of edges in M then $|S| = 2|M|$. Say OPT is the size of the smallest vertex cover. Then, by lemma 2, $|S| = 2|M| \leq 2 \cdot OPT$. ■

The analysis of the approximation ratio of VC2 is tight. Consider the complete bipartite graph $K_{n,n}$ with n vertices on each side of the bipartition. All $2n$ vertices will be selected by VC2 whereas it is sufficient to only select all n vertices from one of the bipartitions.

The best known algorithm for vertex cover has ratio $2 - \Theta(\log^{-1/2} |V|)$ (Karakostas, 04). Currently, the best lower bound that assumes $P \neq NP$ shows that vertex cover cannot be approximated within a factor of $10\sqrt{5} - 21 > 1.3606$ (Dinur and Safra, 02). A tighter lower-bound can be obtained under the *unique games conjecture*; a much stronger assumption than $P \neq NP$. This lower-bound states that vertex cover cannot be approximated within $2 - \epsilon$ for any constant $\epsilon > 0$ (Khot and Regev, 03).

1.5 Steiner Tree

Steiner tree problem is a very well studied problem that generalizes the Minimum Spanning Tree problem. This problem (and its generalizations that we see in future) has a wide range of applications in fields like VLSI design, computer networks, and computational biology.

- **Steiner tree problem:** Given an undirected graph $G = (V, E)$ with cost function $c : E \rightarrow Q^+$, and a subset $T \subseteq V$ called terminals, the goal is to find a minimum cost tree spanning all vertices in T . The vertices in $S = V - T$ are called Steiner nodes.

One of the special cases of Steiner tree is Spanning tree in which $T = V$. As we all know, minimum Spanning tree problem can be solved in polynomial time. But the Minimum Steiner problem is NP-hard.

Before we study this problem we discuss a notion that we will use later on in the course very frequently. We say a cost function c satisfies metric property (or is a metric) if:

- $c(u, u) = 0$ for all points u ,
- $c(u, v) = c(v, u)$ for all pairs of points u, v ,
- $c(u, v) \leq c(u, w) + c(w, v)$ for all u, v, w (triangle inequality)

By this definition, when we say that $G = (V, E)$ has metric cost function it means that $c(u, v)$ is defined for every pair of nodes (so G is a complete graph) and in particular it satisfies the triangle inequality condition.

We will consider the restriction of the problem to the metric case. We show that this restricted version of the problem is in fact as hard as the general version:

Theorem 1 *There is an approximation factor preserving reduction from the Steiner tree problem to the metric Steiner tree problem.*

Proof. Given $G(V, E)$ as an instance for the general case, construct the complete graph G' on V by assigning $c_{G'}(uv)$ (cost of uv in G') to be the cost of the shortest uv -path in G . The set of terminals in G' is the same as in G .

Trivially this is a metric instance (follows directly from the definition and property of shortest paths).

Claim 1 *The cost of optimal solution to G' is less than or equal to the cost of optimal solution of G .*

Proof. This is because for any edge uv in G , $c_{G'}(uv) \leq c_G(uv)$. ■

Claim 2 *The cost of optimal solution to G is less than or equal to the cost of optimal solution to G' .*

Proof. Take any optimal Steiner tree \mathcal{T}' for G' and replace each edge vw in \mathcal{T}' by the shortest path between v and w in G to obtain a subgraph of G . Remove the extra edges to obtain that create cycles to obtain a tree \mathcal{T} . Clearly, the cost does not increase during this process. Therefore: $cost(\mathcal{T}) \leq cost(\mathcal{T}')$. ■

The proof of theorem follows from the above two claims. ■

By Theorem 1, it is enough to consider only metric instances of the Steiner tree problem. Graph G' built above is called the metric completion of G . Now we present a simple 2-approximation for Steiner tree.

Algorithm: STEINER-TREE(G, T)

1. Find a minimum spanning tree in subgraph of G induced by T , $G[T]$.
2. Return this as the solution.

Lemma 4 STEINER-TREE(G, T) is a factor $2 - \frac{2}{|T|}$ approximation algorithm for Steiner tree problem.

Proof. It is simple to see that the algorithm finds a feasible solution in polynomial time. To prove its factor, let T_{opt} be an optimal Steiner tree. Double every edge of T_{opt} and find an Eulerian walk on this graph. Traverse Eulerian walk and shortcut over Steiner Points and visited terminal vertices to reach a simple cycle C on terminals. Then, delete the heaviest edge from the cycle to reach simple path P on terminals.

Let OPT be the cost of T_{opt} and OPT_{mst} be the cost of minimum spanning tree in $G[T]$. Since G satisfies triangular inequality, $c(C)$ (*i.e.* the cost of C) is not more than the cost of initial Eulerian walk and that is $2OPT$. Because the weight of heaviest edge in this cycle is at least $\frac{c(C)}{|C|} = \frac{c(C)}{|T|}$ and P is a spanning tree in $G[T]$, we have:

$$OPT_{mst} \leq c(P) \leq \left(1 - \frac{1}{|T|}\right) c(C) \leq 2 \left(1 - \frac{1}{|T|}\right) OPT,$$

which completes the proof. ■

The best known approximation algorithm for Steiner tree uses iterative rounding (an LP solution) and has ratio $\ln(4) + \epsilon < 1.39$.