

1 Introduction

NP-hard optimization problems exhibit a rich set of possibilities, all the way from allowing approximability to any required degree, to essentially not allowing approximability at all. Despite this diversity, underlying the process of design of approximation algorithms are some common principles. We will explore these in the current chapter.

An optimization problem is polynomial time solvable only if it has the algorithmically relevant combinatorial structure that can be used as “footholds” to efficiently home in on an optimal solution. The process of designing an exact polynomial time algorithm is a two-pronged attack: unraveling this structure in the problem and finding algorithmic techniques that can exploit this structure.

Although **NP**-hard optimization problems do not offer footholds for finding optimal solutions efficiently, they may still offer footholds for finding near-optimal solutions efficiently. So, at a high level, the process of design of approximation algorithms is not very different from that of design of exact algorithms. It still involves unraveling the relevant structure and finding algorithmic techniques to exploit it. Typically, the structure turns out to be more elaborate, and often the algorithmic techniques result from generalizing and extending some of the powerful algorithmic tools developed in the study of exact algorithms.

On the other hand, looking at the process of designing approximation algorithms a little more closely, one can see that it has its own general principles. We illustrate some of these principles in Section 1.1, using the following simple setting.

Problem 1.1 (Vertex cover) Given an undirected graph $G = (V, E)$, and a cost function on vertices $c : V \rightarrow \mathbf{Q}^+$, 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, will be called the *cardinality vertex cover problem*.

Since the design of an approximation algorithm involves delicately attacking **NP**-hardness and salvaging from it an efficient approximate solution, it will be useful for the reader to review some key concepts from complexity theory. Appendix A and some exercises in Section 1.3 have been provided for this purpose.

It is important to give precise definitions of an **NP**-optimization problem and an approximation algorithm for it (e.g., see Exercises 1.9 and 1.10). Since these definitions are quite technical, we have moved them to Appendix A. We provide essentials below to quickly get started.

An **NP**-optimization problem Π is either a minimization or a maximization problem. Each valid instance I of Π comes with a nonempty set of feasible solutions, each of which is assigned a nonnegative rational number called its objective function value. There exist polynomial time algorithms for determining validity, feasibility, and the objective function value. A feasible solution that achieves the optimal objective function value is called an optimal solution. $\text{OPT}_\Pi(I)$ will denote the objective function value of an optimal solution to instance I . We will shorten this to OPT when there is no ambiguity. For the problems studied in this book, computing $\text{OPT}_\Pi(I)$ is **NP**-hard.

For example, valid instances of the vertex cover problem consist of an undirected graph $G = (V, E)$ and a cost function on vertices. A feasible solution is a set $S \subseteq V$ that is a cover for G . Its objective function value is the sum of costs of all vertices in S . A feasible solution of minimum cost is an optimal solution.

An approximation algorithm, \mathcal{A} , for Π produces, in polynomial time, a feasible solution whose objective function value is “close” to the optimal; by “close” we mean within a guaranteed factor of the optimal. In the next section, we will present a factor 2 approximation algorithm for the cardinality vertex cover problem, i.e., an algorithm that finds a cover of cost $\leq 2 \cdot \text{OPT}$ in time polynomial in $|V|$.

1.1 Lower bounding OPT

When designing an approximation algorithm for an **NP**-hard **NP**-optimization problem, one is immediately faced with the following dilemma. In order to establish the approximation guarantee, the cost of the solution produced by the algorithm needs to be compared with the cost of an optimal solution. However, for such problems, not only is it **NP**-hard to find an optimal solution, but it is also **NP**-hard to compute the cost of an optimal solution (see Appendix A). In fact, in Section A.5 we show that computing the cost of an optimal solution (or even solving its decision version) is precisely the difficult core of such problems. So, how do we establish the approximation guarantee? Interestingly enough, the answer to this question provides a key step in the design of approximation algorithms.

Let us demonstrate this in the context of the cardinality vertex cover problem. We will get around the difficulty mentioned above by coming up with a “good” polynomial time computable *lower bound* on the size of the optimal cover.

1.1.1 An approximation algorithm for cardinality vertex cover

We provide some definitions first. Given a graph $H = (U, F)$, a subset of the edges $M \subseteq F$ is said to be a *matching* if no two edges of M share an endpoint. A matching of maximum cardinality in H is called a *maximum matching*, and a matching that is maximal under inclusion is called a *maximal matching*. A maximal matching can clearly be computed in polynomial time by simply greedily picking edges and removing endpoints of picked edges. More sophisticated means lead to polynomial time algorithms for finding a maximum matching as well.

Let us observe that the size of a maximal matching in G provides a lower bound. This is so because *any* vertex cover has to pick at least one endpoint of each matched edge. This lower bounding scheme immediately suggests the following simple algorithm:

Algorithm 1.2 (Cardinality vertex cover)

Find a maximal matching in G and output the set of matched vertices.

Theorem 1.3 *Algorithm 1.2 is a factor 2 approximation algorithm for the cardinality vertex cover problem.*

Proof: No edge can be left uncovered by the set of vertices picked – otherwise such an edge could have been added to the matching, contradicting its maximality. Let M be the matching picked. As argued above, $|M| \leq \text{OPT}$. The approximation factor follows from the observation that the cover picked by the algorithm has cardinality $2|M|$, which is at most $2 \cdot \text{OPT}$. \square

Observe that the approximation algorithm for vertex cover was very much related to, and followed naturally from, the lower bounding scheme. This is in fact typical in the design of approximation algorithms. In Part II of this book, we show how linear programming provides a unified way of obtaining lower bounds for several fundamental problems. The algorithm itself is designed around the LP that provides the lower bound.

1.1.2 Can the approximation guarantee be improved?

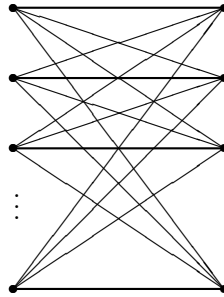
The following questions arise in the context of improving the approximation guarantee for cardinality vertex cover:

1. Can the approximation guarantee of Algorithm 1.2 be improved by a better analysis?

2. Can an approximation algorithm with a better guarantee be designed using the lower bounding scheme of Algorithm 1.2, i.e., size of a maximal matching in G ?
3. Is there some other lower bounding method that can lead to an improved approximation guarantee for vertex cover?

Example 1.4 shows that the answer to the first question is “no”, i.e., the analysis presented above for Algorithm 1.2 is tight. It gives an infinite family of instances in which the solution produced by Algorithm 1.2 is twice the optimal. An infinite family of instances of this kind, showing that the analysis of an approximation algorithm is tight, will be referred to as a *tight example*. The importance of finding tight examples for an approximation algorithm one has designed cannot be overemphasized. They give critical insight into the functioning of the algorithm and have often led to ideas for obtaining algorithms with improved guarantees. (The reader is advised to run algorithms on the tight examples presented in this book.)

Example 1.4 Consider the infinite family of instances given by the complete bipartite graphs $K_{n,n}$.



When run on $K_{n,n}$, Algorithm 1.2 will pick all $2n$ vertices, whereas picking one side of the bipartition gives a cover of size n . \square

Let us assume that we will establish the approximation factor for an algorithm by simply comparing the cost of the solution it finds with the lower bound. Indeed, almost all known approximation algorithms operate in this manner. Under this assumption, the answer to the second question is also “no”. This is established in Example 1.5, which gives an infinite family of instances on which the lower bound, of size of a maximal matching, is in fact half the size of an optimal vertex cover. In the case of linear-programming-based approximation algorithms, the analogous question will be answered by determining a fundamental quantity associated with the linear programming relaxation – its integrality gap (see Chapter 12).

The third question, of improving the approximation guarantee for vertex cover, is currently a central open problem in the field of approximation algorithms (see Section 30.1).

Example 1.5 The lower bound, of size of a maximal matching, is half the size of an optimal vertex cover for the following infinite family of instances. Consider the complete graph K_n , where n is odd. The size of any maximal matching is $(n - 1)/2$, whereas the size of an optimal cover is $n - 1$. \square

1.2 Well-characterized problems and min–max relations

Consider decision versions of the cardinality vertex cover and maximum matching problems.

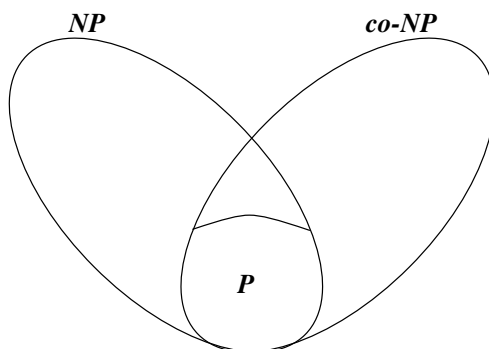
- Is the size of the minimum vertex cover in G at most k ?
- Is the size of the maximum matching in G at least l ?

Both these decision problems are in **NP** and therefore have Yes certificates (see Appendix A for definitions). Do these problems also have No certificates? We have already observed that the size of a maximum matching is a lower bound on the size of a minimum vertex cover. If G is bipartite, then in fact equality holds; this is the classic König-Egerváry theorem.

Theorem 1.6 *In any bipartite graph,*

$$\max_{\text{matching } M} |M| = \min_{\text{vertex cover } U} |U|.$$

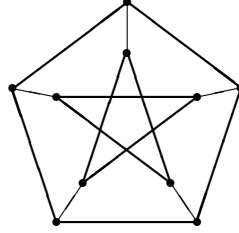
Therefore, if the answer to the first decision problem is “no”, there must be a matching of size $k + 1$ in G that suffices as a certificate. Similarly, a vertex cover of size $l - 1$ must exist in G if the answer to the second decision problem is “no”. Hence, when restricted to bipartite graphs, both vertex cover and maximum matching problems have No certificates and are in **co-NP**. In fact, both problems are also in **P** under this restriction. It is easy to see that any problem in **P** trivially has Yes as well as No certificates (the empty string suffices). This is equivalent to the statement that $\mathbf{P} \subseteq \mathbf{NP} \cap \mathbf{co-NP}$. It is widely believed that the containment is strict; the conjectured status of these classes is depicted below.



Problems that have Yes and No certificates, i.e., are in $\mathbf{NP} \cap \text{co-}\mathbf{NP}$, are said to be *well-characterized*. The importance of this notion can be gauged from the fact that the quest for a polynomial time algorithm for matching started with the observation that it is well-characterized.

Min-max relations of the kind given above provide proof that a problem is well-characterized. Such relations are some of the most powerful and beautiful results in combinatorics, and some of the most fundamental polynomial time algorithms (exact) have been designed around such relations. Most of these min-max relations are actually special cases of the LP-duality theorem (see Section 12.2). As pointed out above, LP-duality theory plays a vital role in the design of approximation algorithms as well.

What if G is not restricted to be bipartite? In this case, a maximum matching may be strictly smaller than a minimum vertex cover. For instance, if G is simply an odd length cycle on $2p + 1$ vertices, then the size of a maximum matching is p , whereas the size of a minimum vertex cover is $p + 1$. This may happen even for graphs having a perfect matching, for instance, the Petersen graph:



This graph has a perfect matching of cardinality 5; however, the minimum vertex cover has cardinality 6. One can show that there is no vertex cover of size 5 by observing that any vertex cover must pick at least $p + 1$ vertices from an odd cycle of length $2p + 1$, just to cover all the edges of the cycle, and the Petersen graph has two disjoint cycles of length 5.

Under the widely believed assumption that $\mathbf{NP} \neq \text{co-}\mathbf{NP}$, \mathbf{NP} -hard problems do not have No certificates. Thus, the minimum vertex cover problem in general graphs, which is \mathbf{NP} -hard, does not have a No certificate, assuming $\mathbf{NP} \neq \text{co-}\mathbf{NP}$. The maximum matching problem in general graphs is in \mathbf{P} . However, the No certificate for this problem is not a vertex cover, but a more general structure: an odd set cover.

An *odd set cover* C in a graph $G = (V, E)$ is a collection of disjoint odd cardinality subsets of V , S_1, \dots, S_k , and a collection v_1, \dots, v_l of vertices such that each edge of G is either incident at one of the vertices v_i or has both endpoints in one of the sets S_i . The *weight* of this cover C is defined to be $w(C) = l + \sum_{i=1}^k (|S_i| - 1)/2$. The following min-max relation holds.

Theorem 1.7 *In any graph,*
$$\max_{\text{matching } M} |M| = \min_{\text{odd set cover } C} w(C).$$

As shown above, in general graphs a maximum matching can be smaller than a minimum vertex cover. Can it be arbitrarily smaller? The answer is “no”. A corollary of Theorem 1.3 is that in any graph, the size of a maximum matching is at least half the size of a minimum vertex cover. More precisely, Theorem 1.3 gives, as a corollary, the following approximate min–max relation. Approximation algorithms frequently yield such approximate min–max relations, which are of independent interest.

Corollary 1.8 *In any graph,*

$$\max_{\text{matching } M} |M| \leq \min_{\text{vertex cover } U} |U| \leq 2 \cdot \left(\max_{\text{matching } M} |M| \right).$$

Although the vertex cover problem does not have No certificate under the assumption $\mathbf{NP} \neq \mathbf{co-NP}$, surely there ought to be a way of certifying that (G, k) is a “no” instance for small enough values of k . Algorithm 1.2 (more precisely, the lower bounding scheme behind this approximation algorithm) provides such a method. Let $\mathcal{A}(G)$ denote the size of vertex cover output by Algorithm 1.2. Then, $\text{OPT}(G) \leq \mathcal{A}(G) \leq 2 \cdot \text{OPT}(G)$. If $k < \mathcal{A}(G)/2$ then $k < \text{OPT}(G)$, and therefore (G, k) must be a “no” instance. Furthermore, if $k < \text{OPT}(G)/2$ then $k < \mathcal{A}(G)/2$. Hence, Algorithm 1.2 provides a No certificate for all instances (G, k) such that $k < \text{OPT}(G)/2$.

A No certificate for instances (I, B) of a minimization problem Π satisfying $B < \text{OPT}(I)/\alpha$ is called a *factor α approximate No certificate*. As in the case of normal Yes and No certificates, we do not require that this certificate be polynomial time computable. An α factor approximation algorithm \mathcal{A} for Π provides such a certificate. Since \mathcal{A} has polynomial running time, this certificate is polynomial time computable. In Chapter 27 we will show an intriguing result – that the shortest vector problem has a factor n approximate No certificate; however, a polynomial time algorithm for constructing such a certificate is not known.

1.3 Exercises

1.1 Give a factor $1/2$ algorithm for the following.

Problem 1.9 (Acyclic subgraph) Given a directed graph $G = (V, E)$, pick a maximum cardinality set of edges from E so that the resulting subgraph is acyclic.

Hint: Arbitrarily number the vertices and pick the bigger of the two sets, the forward-going edges and the backward-going edges. What scheme are you using for upper bounding OPT?

1.2 Design a factor 2 approximation algorithm for the problem of finding a minimum cardinality maximal matching in an undirected graph.

Hint: Use the fact that any maximal matching is at least half the maximum matching.

1.3 (R. Bar-Yehuda) Consider the following factor 2 approximation algorithm for the cardinality vertex cover problem. Find a depth first search tree in the given graph, G , and output the set, say S , of all the nonleaf vertices of this tree. Show that S is indeed a vertex cover for G and $|S| \leq 2 \cdot \text{OPT}$.

Hint: Show that G has a matching of size $\lceil |S|/2 \rceil$.

1.4 Perhaps the first strategy one tries when designing an algorithm for an optimization problem is the greedy strategy. For the vertex cover problem, this would involve iteratively picking a maximum degree vertex and removing it, together with edges incident at it, until there are no edges left. Show that this algorithm achieves an approximation guarantee of $O(\log n)$. Give a tight example for this algorithm.

Hint: The analysis is similar to that in Theorem 2.4.

1.5 A maximal matching can be found via a greedy algorithm: pick an edge, remove its two endpoints, and iterate until there are no edges left. Does this make Algorithm 1.2 a greedy algorithm?

1.6 Give a lower bounding scheme for the arbitrary cost version of the vertex cover problem.

Hint: Not easy if you don't use LP-duality.

1.7 Let $A = \{a_1, \dots, a_n\}$ be a finite set, and let " \leq " be a relation on A that is reflexive, antisymmetric, and transitive. Such a relation is called a *partial ordering* of A . Two elements $a_i, a_j \in A$ are said to be *comparable* if $a_i \leq a_j$ or $a_j \leq a_i$. Two elements that are not comparable are said to be *incomparable*. A subset $S \subseteq A$ is a *chain* if its elements are pairwise comparable. If the elements of S are pairwise incomparable, then it is an *antichain*. A *chain (antichain) cover* is a collection of chains (antichains) that are pairwise disjoint and cover A . The size of such a cover is the number of chains (antichains) in it. Prove the following min-max result. The size of a longest chain equals the size of a smallest antichain cover.

Hint: Let the size of the longest chain be m . For $a \in A$, let $\phi(a)$ denote the size of the longest chain in which a is the smallest element. Now, consider the partition of A , $A_i = \{a \in A \mid \phi(a) = i\}$, for $1 \leq i \leq m$.

1.8 (Dilworth's theorem, see [202]) Prove that in any finite partial order, the size of a largest antichain equals the size of a smallest chain cover.

Hint: Derive from the König-Egerváry Theorem. Given a partial order on n -element set A , consider the bipartite graph $G = (U, V, E)$ with $|U| = |V| = n$ and $(u_i, v_j) \in E$ iff $a_i \leq a_j$.

The next ten exercises are based on Appendix A.

1.9 Is the following an **NP**-optimization problem? Given an undirected graph $G = (V, E)$, a cost function on vertices $c : V \rightarrow \mathbf{Q}^+$, and a positive integer k , find a minimum cost vertex cover for G containing at most k vertices.

Hint: Can valid instances be recognized in polynomial time (such an instance must have at least one feasible solution)?

1.10 Let \mathcal{A} be an algorithm for a minimization **NP**-optimization problem Π such that the expected cost of the solution produced by \mathcal{A} is $\leq \alpha \text{OPT}$, for a constant $\alpha > 1$. What is the best approximation guarantee you can establish for Π using algorithm \mathcal{A} ?

Hint: A guarantee of $2\alpha - 1$ follows easily. For guarantees arbitrarily close to α , run the algorithm polynomially many times and pick the best solution. Apply Chernoff's bound.

1.11 Show that if SAT has been proven **NP**-hard, and SAT has been reduced, via a polynomial time reduction, to the decision version of vertex cover, then the latter is also **NP**-hard.

Hint: Show that the composition of two polynomial time reductions is also a polynomial time reduction.

1.12 Show that if the vertex cover problem is in **co-NP**, then **NP** = **co-NP**.

1.13 (Pratt [230]) Let L be the language consisting of all prime numbers. Show that $L \in \mathbf{NP}$.

Hint: Consider the multiplicative group $\text{mod } n$, $Z_n^* = \{a \in \mathbf{Z}^+ \mid 1 \leq a < n \text{ and } (a, n) = 1\}$. Clearly, $|Z_n^*| \leq n - 1$. Use the fact that $|Z_n^*| = n - 1$ iff n is prime, and that Z_n^* is cyclic if n is prime. The Yes certificate consists of a primitive root of Z_n^* , the prime factorization of $n - 1$, and, recursively, similar information about each prime factor of $n - 1$.

1.14 Give proofs of self-reducibility for the optimization problems discussed later in this book, in particular, maximum matching, MAX-SAT (Problem 16.1), clique (Problem 29.15), shortest superstring (Problem 2.9), and Minimum makespan scheduling (Problem 10.1).

Hint: For clique, consider two possibilities, that v is or isn't in the optimal clique. Correspondingly, either restrict G to v and its neighbors, or remove v from G . For shortest superstring, remove two strings and replace them by a legal overlap (may even be a simple concatenation). If the length of the optimal superstring remains unchanged, work with this smaller instance. Generalize the scheduling problem a bit – assume that you are also given the number of time units already scheduled on each machine as part of the instance.

1.15 Give a suitable definition of self-reducibility for problems in **NP**, i.e., decision problems and not optimization problems, which enables you to obtain a polynomial time algorithm for finding a feasible solution given an oracle for the decision version, and moreover, yields a self-reducibility tree for instances.

Hint: Impose an arbitrary order among the atoms of a solution, e.g., for SAT, this was achieved by arbitrarily ordering the n variables.

1.16 Let Π_1 and Π_2 be two minimization problems such that there is an approximation factor preserving reduction from Π_1 to Π_2 . Show that if there is an α factor approximation algorithm for Π_2 then there is also an α factor approximation algorithm for Π_1 .

Hint: First prove that if the reduction transforms instance I_1 of Π_1 to instance I_2 of Π_2 then $\text{OPT}_{\Pi_1}(I_1) = \text{OPT}_{\Pi_2}(I_2)$.

1.17 Show that

$$L \in \mathbf{ZPP} \text{ iff } L \in (\mathbf{RP} \cap \text{co-}\mathbf{RP}).$$

1.18 Show that if $\mathbf{NP} \subseteq \text{co-}\mathbf{RP}$ then $\mathbf{NP} \subseteq \mathbf{ZPP}$.

Hint: If SAT instance ϕ is satisfiable, a satisfying truth assignment for ϕ can be found, with high probability, using self-reducibility and the **co-RP** machine for SAT. If ϕ is not satisfiable, a “no” answer from the **co-RP** machine for SAT confirms this; the machine will output such an answer with high probability.

1.4 Notes

The notion of well-characterized problems was given by Edmonds [75] and was precisely formulated by Cook [53]. In the same paper, Cook initiated the theory of **NP**-completeness. Independently, this discovery was also made by Levin [193]. It gained its true importance with the work of Karp [171], showing **NP**-completeness of a diverse collection of fundamental computational problems.

Interestingly enough, approximation algorithms were designed even before the discovery of the theory of **NP**-completeness, by Vizing [263] for the minimum edge coloring problem, by Graham [119] for the minimum makespan problem (Problem 10.1), and by Erdős [79] for the MAX-CUT problem (Problem 2.14). However, the real significance of designing such algorithms emerged only after belief in the $\mathbf{P} \neq \mathbf{NP}$ conjecture grew. The notion of an approximation algorithm was formally introduced by Garey, Graham, and Ullman [97] and Johnson [157]. The first use of linear programming in approximation

algorithms was due to Lovász [199], for analyzing the greedy set cover algorithm (see Chapter 13). An early work exploring the use of randomization in the design of algorithms was due to Rabin [232] – this notion is useful in the design of approximation algorithms as well. Theorem 1.7 is due to Edmonds [75] and Algorithm 1.2 is due independently to Gavril and Yannakakis (see [225]).

For basic books on algorithms, see Cormen, Leiserson, Rivest, and Stein [56], Papadimitriou and Steiglitz [225], and Tarjan [254]. For a good treatment of min–max relations, see Lovász and Plummer [202]. For books on approximation algorithms, see Hochbaum [133] and Ausiello, Crescenzi, Gambosi, Kann, Marchetti, and Protasi [18]. Books on linear programming, complexity theory, and randomized algorithms are listed in Sections 12.5, A.6, and B.4, respectively.



<http://www.springer.com/978-3-540-65367-7>

Approximation Algorithms

Vazirani, V.V.

2003, XIX, 380 p., Hardcover

ISBN: 978-3-540-65367-7