Chapter 2

Complexity

In Chapter 1 we learned some nifty algorithmic techniques. However, sometimes we are dealing with a problem where none of these techniques seem to work! Should we give up? And what do we do after giving up?

2.1 P vs. NP

Definition 2.1 (Computational Problem). *A **computational problem** is defined as a (possibly infinite) set of inputs X and a (possibly empty) set of outputs Y for each input $x \in X$.*

Definition 2.2 (Optimization Problem). *In an **optimization problem**, given an input $x \in X$ and a quality measure function $q(x, y)$, the task is to find an output $y \in Y$, such that $q(x, y)$ is optimal.*

Remarks:

- We have seen several optimization problems in the previous chapter: What is the maximal value that can be packed into a knapsack? What is the optimal value of flow in a given network?
- We can also rewrite any linear program (Definition 1.18) as an optimization problem: we define the input to the problem as $x = (A, b, c)$, the set of admissible outputs as $Y = \{y|Ay \leq b\}$, and the function to be optimized as $f(x, y) = c^T y$.

Definition 2.3 (Decision Problem). *In a **decision problem**, given an input (x, d) and a quality measure function $q(x, y)$, where $x \in X$ and d is a decision value, the task is to check whether there exists a solution $y \in Y$ with $q(x, y) \geq d$. The output of the decision problem is a boolean value.*

Remarks:

- Note that we can reformulate any optimization problem as a decision problem. The corresponding decision problem asks a Yes/No question about quality of a solution. For example: Can we pack items of total value 100 into a knapsack? Is there a flow of value at least 5 in a given network?

Lemma 2.4. *By solving the optimization problem, we find a solution for the corresponding decision problem.*

Proof. Assume we have a decision problem, where we want to verify whether a problem has a solution of at least quality q. We can then run an algorithm for the corresponding optimization problem and check whether the solution is better or worse than q. □

Remarks:

- The opposite direction of Lemma 2.4 is usually also possible: if we have an algorithm for a decision problem, we can run it for many different decision values in order to determine the optimal value. This way, we can solve the corresponding optimization problem, but it is not clear how efficiently. The decision problem is therefore intuitively simpler than the optimization problem.
- In the first part of this chapter, we will focus on decision problems.
- In the previous chapter, we have discussed a measure for discussing the running time of algorithms – the time complexity (Definition 1.5). Algorithms for Knapsack seemed to have exponential time complexity, while network flow and matching were polynomial.
- When studying the complexity of computers, we are first and foremost interested in whether a problem can be solved in polynomial time. All problems that can be solved in polynomial time are considered "easy", and grouped together in a class.

Definition 2.5 (Complexity Class P). *The **complexity class** P contains all decision problems that can be solved in polynomial time, i.e. where we know an algorithm whose running time is in $\mathcal{O}(\mathsf{poly}(n))$, where n is the size of the input.*

Remarks:

- There are many problems that belong to class P: sorting, matrix multiplication, max value in an array, etc.
- In order to show that a problem is in P, we usually give an explicit algorithm and analyze the running time directly. There is also an indirect way of showing that a problem is in P: if we can show that some problem known to be in P can be used to solve our problem. Then our problem is "at most as difficult" as a problem in P and is therefore also in P.

Definition 2.6 (Polynomial Reduction). *Let $A = (x_A, q_A)$ and $B = (x_B, q_B)$ be two decision problems, and Y_A, Y_B denote the two admissible sets for the respective problems. We say that there exists a polynomial reduction from A to B, short $A \leq B$, if there exist a function $r : (X_A, Y_A) \to (X_B, Y_B)$ such that $f(x, y)$ has quality q_A if and only if $f(r(x, y))$ has quality q_B. In addition, there is an algorithm that can compute r in $\mathcal{O}(\mathsf{poly}(n))$, where n is the size of input A.*

Lemma 2.7. *Sorting $\leq$ Max*

Proof. Assume that the maximum value of an array with n values can be computed in polynomial time. We can compute the maximum value iteratively n times, and thus sort the array in polynomial time. □

Remarks:

- We can use polynomial reductions between problems in order to determine classes of problems that have roughly the same difficulty.

Problem 2.8 (Clique)**.** *The input is a graph $G = (V, E)$ with n nodes, and an integer $k < n$. In the* ***clique (decision) problem*** *we want to know whether there exists a subset of k nodes in G such that there is an edge between any two nodes.*

Problem 2.9 (Independent Set)**.** *The input is a graph $G = (V, E)$ with n nodes, and an integer $k < n$. In the* ***independent set problem*** *we want to know whether there exists a subset of k nodes in G such that there is* ***no*** *edge between any two nodes in the subset.*

Theorem 2.10. *Clique $\leq$ Independent Set and Independent Set $\leq$ Clique.*

Proof. There is a simple function that maps one problem to the other: Let $K = V \times V$ denote all possible edges in a graph G. Given the graph $G = (V, E)$ in Clique, remove all edges E of this graph and add all edges in $K - E$. Then, all cliques in G will become independent sets in the new graph, and vice versa. The transformation can be computed in polynomial time. □

Problem 2.11 (Vertex Cover)**.** *The input is a graph $G = (V, E)$ with n nodes, and an integer $k < n$. In the* ***vertex cover problem*** *we want to know whether there exists a subset S of k nodes in G such that every edge in E has at least one of its two adjacent nodes in S.*

Theorem 2.12. *Vertex Cover $\leq$ Independent Set and Independent Set $\leq$ Vertex Cover.*

Proof. Given a any independent set S, all other nodes $V - S$ will be a vertex cover. The nodes in S are not connected to each other by definition, so the nodes $V - S$ must cover all edges. Likewise, if the nodes $V - S$ cover all edges, then the nodes S cannot have any edges between them, so S is an independent set.

If S is a maximal independent set, the vertex cover $V - S$ must be minimal: If the vertex cover was not minimal, we could find a bigger independent set, so S was not maximal. And vice versa. □

Lemma 2.13. *The relation $\leq$ is transitive, i.e., if $A \leq B$ and $B \leq C$ then $A \leq C$.*

Proof. Let A, B and C be computational problems, such that $A \leq B$ and $B \leq C$. Then there exists an algorithm that can reduce problem A to C in $\mathcal{O}(\mathsf{poly}(\mathsf{poly}(n))) = \mathcal{O}(\mathsf{poly}(n))$ time. □

Remarks:

- Therefore, there are polynomial reductions between Clique and Vertex Cover as well. If you know how to solve one of these problems in polynomial time, you can solve the other two in polynomial time as well! Unfortunately, nobody knows whether these problems (and many others!) are in P.
- Using polynomial reductions, we can define a hierarchy of different problems.
- Even though we could not find an efficient (polynomial) algorithm for Knapsack as of now, we can try to see whether a simpler decision problem can be solved efficiently: Assume a friend (magically) suggests a solution to a seemingly difficult problem. Can you at least efficiently verify whether the proposed solution is admissible and gives the desired result?

Definition 2.14 (NP)**.** *The class of non-deterministic polynomial problems* (NP) *contains decision problems, for which, given an input x and a possible solution y, the problem whether $q(x, y)$ has a certain quality can be verified in time $\mathcal{O}(\mathsf{poly}(n))$, where n is the size of the input x.*

Remarks:

- The solution y in Definition 2.14 is also called certificate or witness.
- The term "non-deterministic" means that an algorithm is allowed to use parallelism and randomness. A non-deterministic algorithm can in parallel evaluate all possible solutions.
- The class NP contains a large set of computational problems. In fact, it is difficult to imagine a problem not in NP.

Lemma 2.15. *Knapsack is in* NP.

Proof. For Knapsack, a solution is a set of items. In order for Knapsack to be in NP, we would be given a set of items and a total value, for example 100. Verifying whether a given set S of items has a total value of 100 is just a sum. We also need to make sure that S is admissible. This can be also done in polynomial time. So Knapsack is in NP. □

Lemma 2.16. *Clique is in* NP.

Proof. Given a graph and a proposed set of k nodes for a clique, we need to check whether the k proposed nodes are connected. This can be done by going through the list of all edges and counting the edges between the k nodes. If we counted $k(k-1)/2$ edges, the proposed node set forms a clique. □

Lemma 2.17. $\mathsf{P} \subseteq \mathsf{NP}$.

Proof. Every problem in P is automatically in NP. For problems in P, we can assume that the solution is first computed in Definition 2.14. □

Lemma 2.18. *Let B be a problem in* NP *and A be a problem with $A \leq B$. Then, A is also in* NP.

Proof. Given an input x_A and a solution y_A to problem A, there exists a function r that transforms the pair (x_A, y_A) to a pair (x_B, y_B) of problem B. Since B is in NP, we can verify in polynomial time whether the pair (x_B, y_B) is a valid solution for B and therefore also whether (x_A, y_A) is valid for A. A is in NP, because r and the verification can be both computed in polynomial time in the input size of A. □

Remarks:

- Observe that all three problems (Clique, Independent Set, Vertex Cover) are in the class NP, and they seem to be more difficult than the problems in P. Are there even more difficult problems? Is Knapsack more difficult? What is the most difficult problem?!

2.2 NP-hard

Definition 2.19 (NP-hard)**.** *A decision problem H is called* NP*-hard, if there exists a polynomial reduction from every problem in* NP *to H.*

Remarks:

- In other words, an NP-hard problem is at least as difficult as *all* the problems in NP.
- NP-hard problems do not necessarily have to be in NP.
- Are there also NP-hard problems in NP?

Definition 2.20 (NP- complete)**.** *A decision problem is called* NP*- complete if it is* NP*-hard and it is contained in* NP.

Lemma 2.21. *Let B be an* NP*-complete problem and C be a problem for which exists a polynomial reduction $B \leq C$. Then, C is* NP*-hard.*

Proof. Since B is NP-complete, there exists a polynomial reduction from every problem in NP to B. By using Lemma 2.13, we can show there also exists a polynomial reduction from every problem in NP to C. Therefore, C is NP-hard. □

Remarks:

- Can we identify such an NP- complete problem? If yes, we could use polynomial reductions to show that there are many more problems contained in the NP- complete class. We will postpone this question to Section 2.3.
- Figure 2.22 visualizes some of the complexity classes and how they are related to each other. The figure just represents our current knowledge of how these classes relate to each other. There are still many open questions in complexity theory that could change the landscape once an answer is found. Also, there are many more classes.

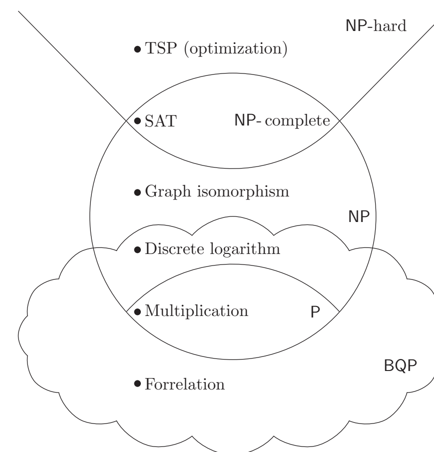

Figure 2.22: A complexity "pet" zoo: Diagram with different complexity classes and sample problems inside each class. The boolean satisfiability problem (SAT) will be introduced in the next section; the optimization version of the traveling salesperson problem (TSP) will be introduced in Section 2.5; the discrete logarithm and the graph isomorphism problems will be discussed in Chapter 3.

- It is generally believed that there are some problems which are in NP but not in P. Or simply:

Conjecture 2.23. P $\neq$ NP.

Remarks:

- The P versus NP question is one of the most important open scientific problems. There were many proof attempts, however, so far, without success.
- Another open question is whether quantum computers will allow to solve all difficult problems efficiently.

Definition 2.24 (BQP)**.** *A decision problem is in the class of* ***bounded-error quantum polynomial time*** (BQP) *problems, if it can be solved on a quantum computer in polynomial time and if for any input the probability to compute the wrong answer is at most* $1/3$.

Remarks:

- It is known that some problems from NP are also in BQP.
- The relation between the class of NP- complete problems and problems from BQP is unknown.
- It has been shown that there is a problem in BQP that is not in NP – the forrelation problem. In this problem, the question is whether a boolean function correlates to a Fourier transform of another boolean function. This result implies that only a quantum computer could solve (or even just verify) Forrelation efficiently.

2.3 Boolean Formulas

Definition 2.25 (Boolean Formula)**.** *Let $\boldsymbol{x} = (x_1, \ldots, x_n)$ be a vector of boolean variables (True/False, respectively 1/0). A* ***boolean formula*** *f consists of boolean variables and logical operations AND, OR and NOT. In particular, f maps a given assignment of variables to True or False (equivalently, 1 or 0). We say that a boolean formula is* ***satisfiable*** *if there exists an assignment of variables $\boldsymbol{x}$, such that $f(\boldsymbol{x}) = 1$. We will call x_i and its negation $\neg x_i$* ***literals****. Literals that are all connected by a logical AND (OR) operation we will call an* ***AND-clause*** *(OR-clause).*

Definition 2.26 (Conjunctive Normal Form CNF)**.** *A boolean formula f is in* ***conjunctive normal form (CNF)****, if it consists of OR-clauses that are connected by an AND operation. A CNF formula is satisfiable, if there is an assignment of variables such that all OR-clauses of the formula are satisfied.*

Remarks:

- Any boolean formula can be written as CNF.

Problem 2.27 (SAT)**.** *In the boolean satisfiability problem (SAT), the task is to determine whether a given CNF formula is satisfiable.*

Problem 2.28 (3-SAT)**.** *3-SAT is a special case of SAT, where the given formula is a CNF that has exactly three literals in each OR-clause.*

Theorem 2.29. *SAT is in* NP.

Proof. Assume we are given an arbitrary boolean formula and an assignment of its variables that acts as a solution. We need to verify in polynomial time whether the given assignment is True. Since the input (formula and assignment) has the size of the boolean formula, the verification can be done efficiently. □

Remarks:

- Intuitively, SAT seems difficult to solve. With n variables, there are 2^n possible assignments. If only one assignment satisfies the boolean formula, we are trying to find a needle in a haystack.

- The satisfiability problem with DNF (disjunctive normal form) formulas, on the other hand, is in P. In a DNF formula, the operations AND and OR are swapped. Therefore, only one AND-clause has to be satisfied in order to satisfy the whole formula. To do so, it is sufficient to verify that there is an AND-clause that contains no variable in its negated and non-negated form.
- Note that CNF $\leq$ DNF is not true. While every CNF can be expressed as DNF, the DNF may be exponentially bigger than the CNF, so the reduction function r in Definition 2.6 is not polynomial.

Theorem 2.30 (Cook/Levin Theorem)**.** *SAT is* NP*- complete.*

Proof. The proof of this theorem is beyond the scope of this script and hence omitted here. The theorem is usually proven by using an alternative definition of NP- completeness via non-deterministic Turing machines. Using this definition one can show that it is possible to encode a Turing Machine as a SAT formula in polynomial time. □

Lemma 2.31. *3-SAT is* NP*- complete.*

Proof. Observe that 3-SAT $\leq$ SAT holds since 3-SAT is a special case of SAT, and therefore, according to Lemma 2.18, 3-SAT is in NP. The other direction also holds, that is SAT $\leq$ 3-SAT: In order to show this, OR-clauses that contain more than three literals need to be split into smaller clauses, while clauses with less literals need to be filled up. Here, we will only consider some small examples. A clause with two literals $(x_1 \vee x_2)$ can simply be replaced by $(x_1 \vee x_2 \vee y) \wedge (x_1 \vee x_2 \vee \neg y)$, where y is a new variable. A clause of four literals $(x_1 \vee x_2 \vee x_3 \vee x_4)$ we can rewrite as $(x_1 \vee x_2 \vee z) \wedge (x_3 \vee x_4 \vee \neg z)$, where z is a new variable. □

Remarks:

- Not all instances of SAT are hard. In 2-SAT, every clause contains exactly two literals. In contrast to 3-SAT, a solution of 2-SAT can be found in polynomial time.
- In order to show that some of the previously introduced problems (e.g. Clique or Vertex Cover) are also NP- complete, we will search for polynomial-time reductions from SAT (or 3-SAT) to these problems.

Theorem 2.32. *Clique is* NP*-hard.*

Proof. We will prove that Clique is NP-hard by using Lemma 2.21 and showing the reduction SAT $\leq$ Clique. We therefore will model a boolean formula as a graph. Given a boolean formula as a CNF, we construct a graph G as follows: Each OR-clause of the form $(\neg x_1 \vee x_2 \vee x_3)$ in the boolean formula forms a cluster of 3 nodes. Each literal from the clause corresponds to a node in the cluster. In the sample clause, we label the corresponding three nodes $\neg x_1$, x_2 and x_3. The nodes inside the same cluster will not be connected. We will also not connect nodes that represent the same variable in a negated and non-negated state, i.e. x and $\neg x$ will not be connected. Each other pair of nodes in G we connect by an edge. Note that is construction is polynomial in the size of the boolean formula.

Next, we show that a given boolean formula of k OR-clauses is satisfiable if and only if the corresponding graph G has a clique of size k. Let us first assume that a given boolean formula is satisfiable, that is, there is an assignment of $\boldsymbol{x}$, such that in each clause there is a literal that evaluates to True. Select one such variable from each of the k clauses. Note that in the graph G, all these variables will be connected. This is because the corresponding nodes are all in different clusters and because nodes with labels x and $\neg x$ cannot both be True in the boolean assignment. Therefore, these k nodes form a clique in G.

Now assume that there exists a clique of size k in G. Note that this clique cannot contain nodes from the same cluster, as such nodes would not be connected. By definition of G, the clique will also not contain nodes with labels x and $\neg x$. Therefore, the clique of size k will have exactly one node from each cluster, and each variable will be represented in either its negated or its non-negated form. By setting the values of the node labels to True, we will receive an assignment of variables that satisfies the corresponding boolean formula. □

Remarks:

- Note that the opposite direction, Clique $\leq$ SAT, does not directly follow from above proof. Not all graphs G can be reached with the construction. In order to show Clique $\leq$ SAT, one would have to show that every graph G can be transformed into a boolean formula in polynomial time.
- In Lemma 2.16, we showed that Clique is in NP. Therefore, Clique is also NP- complete.
- Also Independent Set and Vertex Cover are NP- complete, since there are reductions in both directions between Clique and these problems, which we have shown in Theorem 2.10 and Theorem 2.12.
- In other words, all these problems are the most difficult problems in NP. If you could solve one of them in polynomial time, you could solve all problems in NP!
- What about Knapsack? We need yet another problem.

Problem 2.33 (Subset Sum)**.** *Given a number $s \in \mathbb{N}$ and a set of n natural numbers $x_1, \ldots, x_n$, does there exist a subset of these numbers that has a sum of s.*

Theorem 2.34. *Subset Sum is NP- complete.*

Proof. It is easy to verify whether a given subset of elements has the desired sum s. Therefore, Subset Sum is in NP. In order to show that Subset Sum is NP-hard, the reduction 3-SAT $\leq$ Subset Sum is used. This reduction is non-trivial and will be omitted in this script. □

Theorem 2.35. *The Knapsack (decision) problem is NP-complete.*

Proof. Subset Sum is a special case of Knapsack: let the natural numbers be items with values $x_1, \ldots, x_n$ and weights $x_1, \ldots, x_n$, and let the knapsack have capacity s. If items of value s can be packed into the knapsack, then the corresponding items form a subset of sum s. Therefore, Subset Sum $\leq$ Knapsack.

Since Subset Sum is NP- complete, Knapsack has to be NP-hard. In Lemma 2.15, we already showed that Knapsack is in NP. □

Remarks:

- In fact, thousands of interesting problems in many different areas are NP- complete: Achromatic Number, Battleship, Cut, Dominating Set, Equivalence Deletion, . . . , Super Mario Bros, Tetris, . . . , Zoe.
- Figure 2.36 below shows the reductions of NP- complete problems that we discuss in this chapter.

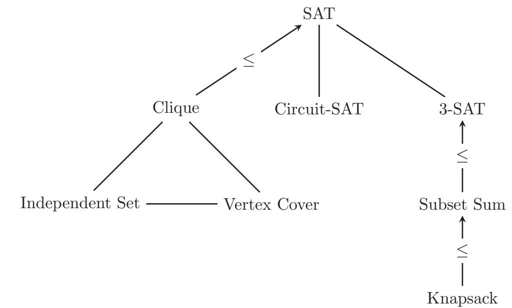

Figure 2.36: Reductions between NP- complete problems

2.4 Boolean Circuits

Definition 2.37 (Boolean Circuit)**.** *Let $n, m \in \mathbb{N}$. A **boolean circuit** is a directed acyclic graph (DAG) with n boolean **input nodes** and m **output nodes**. The input nodes have no incoming edges, while the output nodes have no outgoing edges. The nodes of the graph that are not input nodes represent logical operations (AND, OR, NOT) and are called **gate nodes**. We will label the corresponding nodes $\wedge, \vee$ and $\neg$. Each NOT gate has in-degree 1. The in- and out- degrees of the other gates do not have to be bounded. We further define the **size** of the circuit to be its total number of gates, and the **depth** of the circuit to be the length of the longest (directed) path from an input to an output node.*

A boolean circuit C represents a mapping $C : \{0,1\}^n \mapsto \{0,1\}^m$. The value of an input node is the value of the corresponding input bit of C. The value of each other gate is determined recursively by applying the logic operation that this gate represents to the values received through incoming edges. The output of the boolean circuit are the values of all output nodes.

Remarks:

- Boolean circuits provide a nice tool for understanding the complexity of computation. They are limited in computation power, yet may help to answer the $\mathsf{P} \neq \mathsf{NP}$ conjecture.
- In the literature, it is often assumed that AND and OR gates have a bounded number of input values, or that the circuit has a bounded depth.
- In the following, we will consider boolean circuits that have only one output value (True or False).

Problem 2.38 (Circuit-SAT). *Let $C : \{0,1\}^n \mapsto \{0,1\}$ be a boolean circuit with only one output value. In **Circuit-SAT**, the task is to determine whether a given circuit C has an input vector $\boldsymbol{z} \in \{0,1\}^n$, such that $C(\boldsymbol{z}) = 1$.*

Theorem 2.39. *Circuit-SAT is in* NP.

Proof. We will start by showing that Circuit-SAT is in NP by applying Lemma 2.18 and showing that Circuit-SAT $\leq$ SAT. We therefore need to construct a boolean formula that represents the gates of the given circuit. First, we introduce a new variable g_i representing gate i for each gate of the circuit. We next differentiate between the three possible gates: If g_j is a NOT gate of gate g_i, then we add the clauses $(g_i \wedge \neg g_j) \vee (\neg g_i \wedge g_j)$ to the formula. If g_k is an AND gate with inputs from two gates, g_i and g_j, we add the following clauses to the formula: $(g_k \wedge g_i \wedge g_j) \vee (\neg g_k \wedge g_i \wedge \neg g_j) \vee (\neg g_k \wedge \neg g_i \wedge g_j) \vee (\neg g_k \wedge \neg g_i \wedge \neg g_j)$. In the last case, if g_k is an OR gate with inputs from g_i and g_j, we add $(g_k \wedge g_i \wedge g_j) \vee (g_k \wedge g_i \wedge \neg g_j) \vee (g_k \wedge \neg g_i \wedge g_j) \vee (\neg g_k \wedge \neg g_i \wedge \neg g_j)$ to the formula. Observe that an AND or an OR gate with k inputs can be replaced by $k-1$ gates with two inputs by increasing the depth of the circuit.

Therefore, a boolean formula is satisfiable if and only if the corresponding circuit evaluates to 1. The construction was polynomial in the input size, and the resulting boolean formula can be transformed into a CNF, hence Circuit-SAT is in NP. □

Theorem 2.40. *Circuit-SAT is* NP*-complete.*

Proof. In the previous theorem we have shown that Circuit-SAT is in NP. In order to show that Circuit-SAT is NP-hard, we will reduce SAT to Circuit-SAT, i.e. SAT $\leq$ Circuit-SAT. Let each variable of the boolean formula be an input to the circuit. For each negated variable, we can add a NOT gate. Then, for each clause in the formula, we add an AND gate that has all variables from the clause as input to the gate. And, finally, we add one OR gate that has values from all AND gates as its input. This gives a polynomial construction of a boolean formula as a circuit thus showing that Circuit-SAT is NP-complete. □

Problem 2.41 (Minimum Circuit Size Problem or MCSP). *Given a boolean function $f : \{0,1\}^n \mapsto \{0,1\}$, the task of the **Minimum Circuit Size Problem (MCSP)** is to determine whether f can be represented by a boolean circuit of size s, i.e. a circuit with s gates.*

Remarks:

- There are several approaches of how boolean formulas can be simplified. Often the laws from boolean algebra are applied (Distributive law, Idempotent law, Identity law, Complement law, DeMorgan's law, Karnaugh maps). These minimization rules however do not guarantee that the resulting boolean formula uses the minimum number of gates.
- How difficult is MCSP?

Theorem 2.42. *MCSP is in* NP.

Proof. We will first show that for every n, there exists a boolean function $f : \{0,1\}^n \mapsto \{0,1\}$ which cannot be computed by a circuit of size $2^n/(4n)$. Note that there exist 2^{2^n} possible functions f. Each circuit of size s with in-degree 2 can be represented by at most $3s \log s$ bits. Assume for contradiction that every boolean function can be represented by a circuit of size $s = 2^n/(4n)$. There exist at most $2^{3s \log s}$ such circuits. Observe that

$$2^{3s \log s} = 2^{3 \cdot 2^n/(4n) \log(3 \cdot 2^n/(4n))} \leq 2^{3n \cdot 2^n/(4n)} < 2^{2^n}$$

holds, meaning that not all functions can be computed by a circuit of size $s = 2^n/(4n)$.

The first part showed that the input to MCSP is at least of size $2^n/(4n)$. In order for MCSP to be in NP, we would have to efficiently check that the circuit has the desired size s and that it computes a given boolean function. Since the truth table of the given boolean function has size $(n+1)2^n$, it is possible to compute the truth table for the given circuit in a time that is polynomial in the input size. □

Remarks:

- Is MCSP also NP-hard? As of now, there is no answer to this question. In fact, there is evidence that this result might be difficult to establish.
- Some complexity classes are defined with respect to boolean circuits.

Definition 2.43 (Class AC^0). *The class AC^0 contains all decision problems that can be decided by a family of circuits $\{C_n\}$, where C_n has a constant depth and the size of the circuit is polynomial in n.*

Remarks:

- Note that we need to consider a circuit family, because the input size of each circuit is fixed, while the input to a problem can vary.
- Note that we do not allow the boolean circuits to save output bits or use recursive operations.
- There exist many other classes based on the boolean circuits: The class NC^0 is defined in the same way as AC^0, but the number of incoming edges to the AND and the OR gate is restricted to only two edges. Both classes can be extended to versions that allow a non-constant depth of the circuit: AC^i and NC^i have depth $\mathcal{O}(\log^i n)$.

2.5 Solving Hard Problems

We have seen that many important computational problems are known to be NP-hard. So what do we do if we encounter such a hard problem?

Remarks:

- We will now consider optimization problems instead of decision problems. This often makes sense in practice, where you usually want to find the best solution to some problem, instead of just deciding if the problem has a solution with a specific cost.
- So how can we approach a hard problem in practice? So far, we have mostly considered exact algorithms.

Definition 2.44 (Exact Algorithm)**.** *An **exact algorithm** always returns the optimal solution to the problem.*

Remarks:

- We have seen that for NP- complete problems we do not know if there is an exact algorithm with polynomial running time.
- Even when the worst-case running time of an algorithm is exponential, one can often use tricks to considerably reduce this running time in practice, e.g. by using look-ahead or pruning techniques as in Definition 1.9.
- What if we also accept weaker-than-optimal solutions, but instead, we require our algorithm to always have a reasonable running time?

Definition 2.45 (Heuristic)**.** *A **heuristic** is an algorithm that is guaranteed to have a polynomial running time, at the cost of returning a suboptimal solution.*

Remarks:

- One example for a heuristic is our greedy algorithm for Knapsack, Algorithm 1.8.
- Good heuristics are usually based on insights into the structure or the behavior of the problem. Some heuristics are known to usually provide good solutions in practice.
- However, while heuristics are often good on some inputs, they may return very weak solutions on other inputs.
- In general, it is better to have algorithms which provide guarantees that they always return a solution which is reasonably close to the optimum. We want to understand algorithms with such guarantees.

2.6 Vertex Cover Approximation

Let us first revisit the optimization version of Vertex Cover: we want to find a vertex cover $S \subseteq V$ in an input graph $G = (V, E)$, with $|S|$ as small as possible. Can we come up with an algorithm where we can prove that the size of the returned vertex cover is always “reasonably close” to the optimum?

Remarks:

- A natural approach is the greedy method of Algorithm 2.46.

```
1 def VertexCover_Greedy_Naive(G):
2     S = ∅
3     while E ≠ ∅:
4         select an arbitrary edge (u,v) ∈ E
5         S = S ∪ {u}
6         remove all edges from E that are adjacent to u
7     return S
```

Algorithm 2.46: Naive greedy algorithm for Vertex Cover.

Lemma 2.47. *Algorithm 2.46 returns a vertex cover.*

Proof. An edge from E is only removed when one of its adjacent nodes is included in the set S. □

Remarks:

- However, the size of this vertex cover can be very far from the optimal size.

Theorem 2.48. *The size $|S|$ of the vertex cover returned by Algorithm 2.46 can be an $n-1$ factor larger than the optimum.*

Proof. Consider a “star” graph with n nodes and $n-1$ edges, where $n-1$ distinct leaf nodes u_1, u_2, ..., u_{n-1} are connected by an edge to the same center node v. Whenever Algorithm 2.46 selects an edge (u_i, v) in this graph, it can happen that it always chooses the leaf node u_i. In this case, the algorithm only removes a single edge (u_i, v) in each step, and the algorithm lasts for $n-1$ iterations of the loop. In the end, the algorithm returns the vertex cover $S = \{u_1, ..., u_{n-1}\}$, which consists of $n-1$ nodes.

On the other hand, the set $\{v\}$ is an optimal vertex cover of cost $c^* = 1$ in this graph, so we have $\frac{|S|}{c^*} = n - 1$. □

Remarks:

- It is a natural idea to try to improve Algorithm 2.46 with a clever tie-breaking rule: for example, to always select the higher-degree adjacent node of the chosen edge, or the highest-degree node altogether among the available nodes. This solves the star. However, there are more complicated counterexamples which show that even in this case, the solution we obtain can be a $\log n$ factor worse than the optimum.
- However, there is a slightly different greedy approach that provides much better guarantees.

```
1 def VertexCover_Greedy(G):
2     S = ∅
3     while E ≠ ∅:
4         select an arbitrary edge (u,v) ∈ E
5         S = S ∪ {u,v}
6         remove all edges from E that are adjacent to u or v
7     return S
```

Algorithm 2.49: Greedy algorithm for Vertex Cover.

Theorem 2.50. *Algorithm 2.49 always returns a vertex cover with* $|S| \leq 2 \cdot c^*$.

Proof. Algorithm 2.49 returns a correct vertex cover: an edge is only removed from E if at least one of its adjacent nodes is inserted into S.

Assume that the algorithm runs for c iterations of the main loop, i.e. it selects c different edges (u_i, v_i), and inserts both adjacent nodes u_i and v_i of these edges into S for $i \in \{1, ..., c\}$. Note that none of these edges (u_i, v_i) share a node, because all edges adjacent to either u_i or v_i are removed when (u_i, v_i) is selected (in graph theory, they refer to such a set of edges as a *matching*). This means that in the optimal vertex cover S^*, each vertex can only be adjacent to at most one of the edges (u_i, v_i); thus in order to cover all the edges (u_i, v_i), we already need at least c nodes. As a result, we get that $c^* = |S^*| \geq c$.

On the other hand, our algorithm returns a vertex cover of size $|S| = 2 \cdot c$, so we have $|S| = 2 \cdot c \leq 2 \cdot c^*$. □

Remarks:

- In a graph of $\frac{n}{2}$ independent edges, the solution returned by Algorithm 2.49 is indeed 2 times worse than the optimum, so our analysis is tight.
- This algorithm was somewhat counter-intuitive; one might expect Algorithm 2.46 to be more efficient. However, as we have seen, Algorithm 2.49 is always at most a factor 2 away from the optimum c^*, while Algorithm 2.46 can be an $(n-1)$-factor away.
- Algorithms that are proven to always be within an α factor from the optimum are called *approximation algorithms*. Theorem 2.50 shows that Algorithm 2.49 is a 2-approximation algorithm for Vertex Cover.

- For our formal definition of this notion, we assume a minimization problem.

Definition 2.51 (Approximation Algorithm)**.** *We say that an algorithm $\mathcal{A}$ is an α-**approximation algorithm** if for every possible input of the problem, it returns a solution with a cost of at most $\alpha \cdot c^*$, where c^* is the cost of the optimal solution (the one with minimal cost).*

Remarks:

- The definition is similar for maximization problems. In this case, we denote the value of the optimal solution (the one with highest value) by v^*. For an approximation algorithm, we require that the returned solution always has a value of at least v^*/α.
- This section only considers approximations algorithms that run in polynomial time. That is exactly the point of these algorithms: to find a reasonably good solution without having an unreasonably high running time.

Definition 2.52. *We say that a problem is α-approximable if there exists a polynomial-time algorithm $\mathcal{A}$ that is an α-approximation algorithm for the problem.*

Remarks:

- In general, α can be any constant value with $\alpha > 1$, or it can also be a function of n (the size of the input), e.g. $\alpha = \log n$ or $\alpha = n^{3/4}$.
- In many cases, we can already get some approximability results from the most trivial algorithms. Independent Set naturally has $v^* \leq n$, since we can find a solution of value 1 by simply outputting a set that consists of an arbitrary single node. This shows that Independent Set is (at least) n-approximable.
- This allows us to classify hard problems, based on how well their optimum solution can be approximated. E.g. the complexity class of problems that are α-approximable for some constant α is called APX. Since Theorem 2.50 has $\alpha = 2$, Vertex Cover is in APX.
- For Vertex Cover, there is currently no $2-\varepsilon$ approximation algorithm known for any $\varepsilon > 0$, so the algorithm above is indeed the best we have. Also, it is proven that no approximation better than 1.3606 is possible at all unless $\mathsf{P} = \mathsf{NP}$. This means that unless $\mathsf{P} = \mathsf{NP}$, we are unable to approximate the best solution of this problem arbitrarily well (in polynomial time).

2.7 Bin Packing Approximation

Let us now look at some other problems that are α-approximable to some constant α. We continue with the bin packing problem, which is similar to Knapsack.

Problem 2.53 (Bin Packing)**.** *We have a set of n items of size $x_1, ..., x_n$, and an unlimited number of available bins, each having a capacity of B. A set of items fits into a bin if its total size is at most B. Our goal is to put all of the items into bins, using the smallest possible number of bins.* → notebook

Remarks:

- You can easily imagine immediate applications of this, e.g. packing files on disks.
- In case of this problem, c^* denotes the number of bins that are used in the optimal solution.
- In Bin Packing, a simple greedy heuristic already allows us to obtain a 2-approximation.

```
1 def FirstFit(items, B):
2     for each item in items:
3       place item in the first bin where it still fits
4       if item does not fit into any bin:
5         open a new bin, and insert item into the new bin
```

Algorithm 2.54: First Fit algorithm for Bin Packing.

Lemma 2.55. *In Algorithm 2.54, at most 1 bin is filled at most half.*

Proof. Assume that there are at least 2 bins b_1 and b_2 that are filled at most half. This means that b_1 still has free space of at least $\frac{B}{2}$, and b_2 only contains items of size $x_i \leq \frac{B}{2}$. However, in this case, the items sorted into b_2 would also fit into b_1. This contradicts the First Fit algorithm, which only opens a new bin when the next item does not fit into any of the previous bins. □

Theorem 2.56. *Algorithm 2.54 is a 2-approximation.*

Proof. Assume that the First Fit algorithm uses m bins. Due to Lemma 2.55, at least $m-1$ of these bins are filled to a capacity of more than $\frac{B}{2}$. This implies

$$\sum_{i=1}^{n} x_i > (m-1) \cdot \frac{B}{2}.$$

On the other hand, we know that

$$c^* \geq \frac{\sum_{i=1}^{n} x_i}{B},$$

since all the items have to be sorted into a bin, and every bin can only take at most B. The two inequalities imply

$$B \cdot c^* > (m-1) \cdot \frac{B}{2},$$

and thus

$$2 \cdot c^* > m - 1.$$

Since c^* and m are integers, this implies $2 \cdot c^* \geq m$, proving our claim. □

Remarks:

- One can even improve on this algorithm by first sorting the elements in decreasing order, and then applying the same First Fit rule. With only a slightly more detailed analysis, one can show that this improved algorithm achieves an approximation ratio of 1.5.
- Can we get a better approximation ratio than 1.5? Unfortunately not, unless we have $\mathsf{P} = \mathsf{NP}$. One can show that getting a better than 1.5 approximation is already an NP-hard problem.
- To prove this, we present a reduction to Partition.

Problem 2.57 (Partition)**.** *In the Partition Problem, given a set of n items of positive size $x_1, ..., x_n$, we want to decide if we can partition them into two groups such that the total sum is equal in the two groups.*

Remarks:

- Partition is a special case of Subset Sum where $s = \frac{1}{2} \cdot \sum_{i=1}^{n} x_i$.
- It is known that this special case is still NP- complete.

Theorem 2.58. *Subset Sum $\leq$ Partition.*

Proof. The reduction is not too complex, but a bit boring, so we do not discuss it here. □

Theorem 2.59. *It is NP-hard to solve Bin Packing with an α-approximation ratio for any constant $\alpha < \frac{3}{2}$.*

Proof. Given an input of Partition with integers x_1, ..., x_n, we convert it into a Bin Packing problem: we consider items of size x_1, ..., x_n, and we define $B = \frac{1}{2} \cdot \sum_{i=1}^{n} x_i$.

If the original Partition problem is solvable, then this Bin Packing problem can also be solved with 2 bins. In this case, an α-approximation algorithm with $\alpha < \frac{3}{2}$ must always return a solution with $m \leq c^* \cdot \alpha < 2 \cdot \frac{3}{2} = 3$ bins; since m is an integer, this means a solution with $m = 2$ bins.

On the other hand, if the items cannot be partitioned into two sets of size B, then the algorithm can only return a solution with $m \geq 3$ bins.

Hence an α-approximation algorithm for Bin Packing would also solve Partition in polynomial time: we can just solve the converted Bin Packing problem, and if $m = 2$, then output 'Yes', whereas if $m \geq 3$, then output 'No'. □

Remarks:

- Thus, Bin Packing is a problem that can be approximated to some constant $\alpha = 1.5$, but not arbitrarily well, i.e. not for any $\alpha > 1$.

2.8 TSP Approximation

Let us now consider another natural graph question.

Problem 2.60 (Hamiltonian Cycle)**.** *The input is a graph $G = (V, E)$. In the **Hamiltonian cycle problem** we want to know whether there exists a cycle in G which contains each node of the graph exactly once.*

Remarks:

- Hamiltonian cycle was among the first problems that were shown to be NP- complete.
- In practice, cities often have different distances between each other.

Problem 2.61 (Traveling Salesperson or TSP)**.** *The input is a clique graph $G = (V, E)$ (there is an edge between any two nodes of G), with positive edge weights, i.e. a function $d : E \to \mathbb{R}^+$. The goal of **TSP** is to find a Hamiltonian cycle in G where the total weight of the edges contained in the cycle is minimal.* → notebook

Theorem 2.62. *In general graphs, TSP is NP-hard to approximate to any constant factor α.*

Proof. We provide a reduction to Hamiltonian Cycle, which is NP-complete.

Assume that we have an α-approximation algorithm for TSP. Given an input graph $G = (V, E)$ for Hamiltonian Cycle, we turn this into an instance of TSP on V. For each pair of nodes u and v, we define the weight of edge (u, v) in the TSP in the following way:

- if $(u, v) \in E$, then we assign $d(u, v) = 1$ in the TSP,
- if $(u, v) \notin E$, then we assign $d(u, v) = \alpha \cdot n + 1$ in the TSP.

If the initial graph had a Hamiltonian cycle, then the resulting TSP has a cycle of total weight n as the optimal solution. In this case, our approximation algorithm is guaranteed to return a solution of cost at most $\alpha \cdot n$. On the other hand, if there was no Hamiltonian cycle in the original graph, then the optimal TSP cycle must contain at least one edge of weight $\alpha \cdot n + 1$, and thus $c^* \geq \alpha \cdot n + 1$.

Hence an α-approximation returns a solution of cost at most $\alpha \cdot n$ if and only if the original graph had a Hamiltonian cycle; thus running such an approximation algorithm allows us to solve the Hamiltonian cycle problem. This shows that finding a polynomial-time α-approximation of TSP for *any* constant α is NP-hard. □

Remarks:

- This means that in case of general TSP, the best we can hope for is an approximation ratio that depends on n.
- On the other hand, there is a special case of the problem that does allow a constant-factor approximation.

Problem 2.63 (Metric TSP)**.** *In Metric TSP, the distances between any three nodes v_1, v_2, v_3 must satisfy the triangle-inequality: $d(v_1, v_2) \leq d(v_1, v_3) + d(v_3, v_2)$.*

Remarks:

- Note that the triangle inequality is a very realistic assumption in many real-world applications.
- This is another possible approach to solving hard problems in practice: maybe we can show that our actual problems are restricted to only a special case of the original problem, and that this special case is computationally more tractable.
- A Metric TSP already has a 2-approximation algorithm.

```
1 def Tree_Based_TSP(G)
2     find a Minimum Spanning Tree T in G
3     form a sequence of nodes P0 by traversing all nodes of T
4     simplify P0 to P by only keeping the first occurrence of
      ↪  each node
5     return P
```

Algorithm 2.64: Tree-based approximation for TSP.

Remarks:

- The Minimum Spanning Tree (MST) is, intuitively speaking, the subset of edges with smallest total weight which already connects the whole graph, i.e. there is a path between any two nodes through these edges. An MST in a graph can easily be found in polynomial time with some simple algorithms.

Theorem 2.65. *Algorithm 2.64 is a 2-approximation.*

Proof. See also Figure 2.66. Let t^* denote the total cost of the MST. Note that if we delete a single edge from any Hamiltonian cycle, we get a spanning tree, so t^* must be smaller than the cost of any Hamiltonian cycle. This implies $t^* \leq c^*$.

Now consider the cost of the Hamiltonian cycle returned by Algorithm 2.64. Since the traversal visits each edge of the spanning tree exactly twice, the total cost of P_0 is $2 \cdot t^*$. By only keeping the first occurrence of each node in P_0, we create "shortcuts": whenever we delete a node from P_0, we only shorten our tour due to the triangle inequality. As such, for the total costs we have

$$cost(P) \leq cost(P_0) = 2 \cdot t^* \leq 2 \cdot c^*.$$

□

Remarks:

- Being slightly more clever even allows a 1.5-approximation.
- For TSP, it is often reasonable to assume the even more special case of Euclidean TSP, where the nodes of the graph corresponds to specific points in a plane, for example. That is, each node v is associated with two coordinates v_x and v_y, and the weight of the edge between u and v is the Euclidean distance of u and v, i.e. $\sqrt{(u_x - v_x)^2 + (u_y - v_y)^2}$.

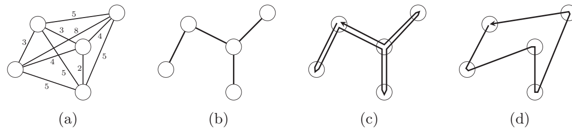

Figure 2.66: A weighted complete graph as an input to TSP (a), an MST in this graph (b), the tour P_0 obtained from this MST, starting from an arbitrary node (c), and the simplified tour P after removing repeated occurrences (d).

2.9 FPTAS: Knapsack

We have now seen many problems that are α-approximable for a specific constant α. Sometimes we can even get arbitrarily close to the optimum in polynomial time: there exists a $(1+\varepsilon)$-approximation for any $\varepsilon > 0$.

Definition 2.67 (FPTAS)**.** *We say an algorithm $\mathcal{A}$ is a fully polynomial-time approximation scheme (FPTAS) if for any $\varepsilon > 0$*

- *$\mathcal{A}$ is a $(1+\varepsilon)$-approximation algorithm, and*
- *the running time of $\mathcal{A}$ is polynomial in both n and $\frac{1}{\varepsilon}$.*

Remarks:

- While an FPTAS is not as good as a polynomial-time exact algorithm, it is almost as good: for any desired error ε, we can efficiently find a solution that is only ε away from the optimum.
- There is also a slightly weaker notion of PTAS (polynomial-time approximation scheme), where we only require the running time to be polynomial in n, but not in $\frac{1}{\varepsilon}$. That is, the running time is still polynomial in n for any fixed $\varepsilon > 0$, but it might increase very quickly in $\frac{1}{\varepsilon}$; for example, it includes a factor of $n^{2^{1/\varepsilon}}$. These algorithms are not as useful in practice: if we want to get very close to the optimum by selecting a small ε, the running time still becomes unreasonably large.
- As an example for a nicely approximable problem, we revisit the Knapsack problem, and present an FPTAS for it.
- Since we now study the problem from more of a mathematical than a programming-based perspective, we introduce a shorter notation for the inputs of the problem.

Definition 2.68 (Knapsack notation)**.** *We will use C to denote the capacity of our Knapsack, and we denote the value and weight of the i^{th} item as v_i and w_i, respectively. We denote the number of items by n as before. Finally, let us use $v_{\max}$ to denote the value of the highest value item, i.e. $v_{\max} := \max v_i$.*

Remarks:

- In case of this problem, we use v^* to denote the maximal value we can fit into the knapsack.
- Note that if any item has $w_i > C$, then it can never fit into the knapsack, so we might as well remove such items. We will assume that $v_{\max}$ denotes the maximal value of an item that still fits into the knapsack.
- Recall that we have discussed a dynamic programming solution for Knapsack, where each cell $V[i][c]$ of our DP table stored the maximum value that can be achieved with capacity c using only the first i items. The starting point of our FPTAS algorithm will be a slightly different variant of this DP method: in our new table, $W[i][v]$ will store the weight of the lowest-weight subset of items $1, ..., i$ that has a total value of v.
- One can show that this table W can also be computed with a similar dynamic programming method to Algorithm 1.12. As before, the table has n rows. The number of columns is now at most $n \cdot v_{\max}$, which is an upper bound on the value of any subset; we have n items, and each of them has value at most $v_{\max}$. From this alternative DP table, we can find the optimum v^* by taking the maximal v value in the table where $W[n][v] \leq C$.

→ notebook

- The main difference from the original DP algorithm is that now each column expresses the *value* of a specific subset of items, instead of the *weight* of a specific subset of items as before.
- We can now use this alternative DP algorithm to develop an FPTAS for Knapsack. The main idea is to slightly round up the values of the items, which results in some inaccuracy for the algorithm, but it also reduces the number of columns in our DP table.

```
1 def Knapsack_FPTAS(items, C):
2     k = ε · v_max / n
3     consider the modified Knapsack problem where item i has
    ↪  weight w_i and value v̂_i = ⌈v_i / k⌉ · k
4     run the alternative DP algorithm on this problem, with
    ↪  columns only corresponding to multiples of k
```

Algorithm 2.69: FPTAS algorithm for Knapsack.

Lemma 2.70. *Algorithm 2.69 has a running time of $\mathcal{O}(n^3 \cdot \frac{1}{\varepsilon})$.*

Proof. Given the parameter ε, the algorithm introduces a scaling parameter $k = \varepsilon \cdot v_{\max} / n$, and defines the *rounded-up value* of item i as $\hat{v}_i = \lceil \frac{v_i}{k} \rceil \cdot k$. This means that whichever subset of nodes we select, the total value of the subset is a multiple of k.

On the other hand, any subset of the items has a total value of $n \cdot v_{\max}$ at most. If each column represents a multiple of k, then the table has at most $n \cdot v_{\max} / k = n^2 \cdot \frac{1}{\varepsilon}$ columns. The table still has n rows, so the total running time of the DP algorithm with these rounded values is $O(n^3 \cdot \frac{1}{\varepsilon})$. □

Theorem 2.71. *Algorithm 2.69 is an FPTAS.*

Proof. The polynomial running time has already been established in Lemma 2.70. Let us use S to denote the set of items chosen by Algorithm 2.69. Since we round the values to multiples of k, the value of each item has been overestimated by k at most, so $\hat{v}_i - k \leq v_i$. Hence for the total value of items in S we have

$$\sum_{i \in S} v_i \geq \sum_{i \in S} (\hat{v}_i - k) = \sum_{i \in S} \hat{v}_i - |S| \cdot k = \sum_{i \in S} \hat{v}_i - |S| \cdot \frac{\varepsilon \cdot v_{\max}}{n} \geq \sum_{i \in S} \hat{v}_i - \varepsilon \cdot v_{\max}.$$

To simplify this last expression even more, we make two more observations. First, since the value of each item is only rounded upwards in this modified problem, and our algorithm finds the optimal solution, we have $\sum_{i \in S} \hat{v}_i \geq v^*$. Also, choosing the item with maximal value $v_{\max}$ is always a valid solution, so we also have $v_{\max} \leq v^*$. Hence for the total value in S we have

$$\sum_{i \in S} \hat{v}_i - \varepsilon \cdot v_{\max} \geq v^* - \varepsilon \cdot v^* = (1 - \varepsilon) \cdot v^*.$$

Note that this shows $v^* / v(S) \leq \frac{1}{1-\varepsilon}$, while our original FPTAS definition requires a slightly different relation: that $v^* / v(S) \leq (1 + \varepsilon)$. However, since we now have an $\alpha = \frac{1}{1-\hat{\varepsilon}}$-approximation algorithm for any $\hat{\varepsilon} > 0$, we can obtain $\alpha = (1 + \varepsilon)$ by choosing $\hat{\varepsilon} = \frac{\varepsilon}{1+\varepsilon}$. □

2.10 Non-Approximability: Independent Set

The most difficult problems are those that cannot be approximated to any constant α, but only to a factor $\alpha = f(n)$ depending on n, e.g. $\alpha = \log n$ or $\alpha = \sqrt{n}$. This means that as the size of the problem becomes larger, the difference between the optimum and our solution also grows larger.

Remarks:

- Remember the Independent Set problem? It turns out that this is one of the most difficult problems from this perspective: approximating it to any factor that is slightly smaller than n would already imply that $\mathsf{P} = \mathsf{NP}$.

Theorem 2.72. *For any positive constant $\delta > 0$, it is NP-hard to approximate the Independent Set problem to a ratio of $n^{1-\delta}$.*

Proof. The proof of this theorem is more involved than our previous claims, so we do not discuss it here. The idea of the proof is built on the notion of so-called probabilistically checkable proof systems. □

Remarks:

- This means that for any approximation algorithm $\mathcal{A}$, there are some graphs where $\mathcal{A}$ returns very weak solutions.
- This $n^{1-\delta}$ factor is indeed huge. For example, it might be that the largest independent set contains $n^{0.99}$ nodes, but $\mathcal{A}$ returns an independent set of only 1 or 2 nodes. As such, an approximation algorithm like this is not useful in practice.
- The same result also holds for the Clique problem.
- Recall from Theorem 2.12 that as a decision problem, Independent Set $\leq$ Vertex Cover; however, we have seen that Vertex Cover is 2-approximable. This shows that even if the decision version of two problems are equivalent, the approximability of the optimization version may differ significantly.
- So while all NP-complete decision problems have the same difficulty in theory, their optimization variants can behave very differently in practice.

Chapter Notes

The first formal definition of computation was provided by Alan Turing in 1936 [20]. In this paper, Turing presents an automatic machine (now known as Turing machine) that can compute certain classes of numbers. We will discuss the Turing machine in more detail in the last chapter. The birth of computational complexity as a field is attributed to Hartmanis and Stearns for their paper “On the computational complexity of algorithms” [14]. They proposed to measure time with respect to the input size and showed that some problems can only be solved if more time is given.

The first problem that was shown to be in P was the maximum matching problem, see Edmonds [9]. In another paper, Edmonds introduced a notion that is equivalent to the class NP [10]. In 1971, Cook [6] showed that SAT was NP-complete. Two years later, independently of Cook’s result, Levin [17] showed that six problems were NP-complete, the so-called universal search problems. Both authors have formulated the famous P versus NP problem in computer science. The name NP was however given to the class a year later by Karp [16], who proved that 21 (combinatorial) problems were NP-complete. He thereby first used the technique of polynomial reductions in these proofs. Since then, thousands of problems in different areas of science have been shown to be NP-hard. Also many new complexity classes have been proposed in the literature. Different “zookeepers” around the world keep track of the newly introduced complexity classes in their complexity zoos [1]. The question $\mathsf{P} \neq \mathsf{NP}$ still remains open. In 2000, the P versus NP problem was announced as one of the seven Millennium Prize Problems by the Clay Mathematics Institute [4]. Solving P versus NP (or BQP versus NP) will be a major milestone in science, in case of equality ($\mathsf{BQP} = \mathsf{NP}$) there will be amazing practical consequences.

Circuit theory is almost 200 years old, being studied by electrical engineers and physicists. In the 1940s, the invention of semiconductor devices and later

the transistor further boosted the quest for understanding circuits. The computational model of boolean circuits was introduced by Claude Shannon [19] in 1949, who showed that most boolean functions require circuits of exponential size (see Theorem 2.42). It was a natural step to restrict the circuits in size and depth and consider problems that are still computable on restricted circuits. Furst, Saxe and Sipser [11] for example showed that Parity is not in AC^0, but in NC^1. The first connection between circuits and Turing machines has been made by Savage [18]. It is generally believed that proving bounds using boolean circuits is easier than by using Turing machines.

While all known exact algorithms for NP-hard problems have a superpolynomial runtime in the worst case, there are numerous possible tricks and optimizations that can make these algorithms viable on special cases of the problem, even for very large inputs. For example, there are yearly SAT-competitions where SAT-solver algorithms try to decide the satisfiability of SAT formulas from real-life applications, with many of these formulas containing millions of variables and tens of millions of clauses [2].

Heuristic solutions to hard problems have also been extensively studied. In particular, there is a wide range of so-called metaheuristics, which are general approaches and techniques for developing a heuristic solution to a problem. This includes some simpler approaches like local search or gradient descent, and also some more sophisticated ones like simulated annealing, tabu search or genetic algorithms [13]. Note that many techniques in machine learning are also developed for this purpose: to provide heuristic solutions to hard problems.

The 2-approximation algorithm for Vertex Cover has long been known, discovered by both Gavril and Yannakakis independently [12]. The FPTAS algorithm for Knapsack has also been around for a long time, and it is one of the most popular FPTAS examples [21].

Bin Packing has also been studied for multiple decades [12]. A recent analysis of the First Fit heuristic is available by Dósa and Sgall, proving the even stronger result that First Fit is a 1.7-approximation [8]. The variant which first sorts the items in decreasing order is known to even provide a $\frac{11}{9}$-approximation up to an additive constant, and this bound is known to be tight [7].

For Travelling Salesperson, a slightly improved version of our approximation algorithm is due to Christofides, which improves the approximation ratio to only 1.5 [5]. In the special case of an Euclidean TSP, there even exists a (rather complicated) PTAS algorithm, which received a Gödel prize [3].

The inapproximability result for Independent Set was the final result of a line of lower bounds in the early 2000s; John Hastad also received a Gödel prize for this result [15, 22].

This chapter was written in collaboration with Darya Melnyk and Pál András Papp.

Bibliography

[1] Complexity zoo. `https://complexityzoo.uwaterloo.ca/Complexity_Zoo`.

[2] The international SAT competition. `http://www.satcompetition.org`.

[3] Sanjeev Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, September 1998.

[4] James A Carlson, Arthur Jaffe, and Andrew Wiles. *The millennium prize problems*. American Mathematical Society Providence, RI, 2006.

[5] Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.

[6] Stephen A Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158, 1971.

[7] György Dósa. The tight bound of first fit decreasing bin-packing algorithm is ffd(i) ¡ 11/9 opt(i)+6/9. In *Combinatorics, Algorithms, Probabilistic and Experimental Methodologies*, pages 1–11, Berlin, Heidelberg, 2007.

[8] György Dósa and Jiri Sgall. First Fit bin packing: A tight analysis. In *30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013)*, volume 20 of *LIPIcs*, pages 538–549, Dagstuhl, Germany, 2013.

[9] Jack Edmonds. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National Bureau of Standards B*, 69(125-130):55–56, 1965.

[10] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17:449–467, 1965.

[11] Merrick Furst, James B Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical systems theory*, 17(1):13–27, 1984.

[12] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.

[13] Fred W Glover and Gary A Kochenberger. *Handbook of metaheuristics*, volume 57. Springer Science & Business Media, 2006.

[14] Juris Hartmanis and Richard E Stearns. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306, 1965.

[15] Johan Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 627–636. IEEE, 1996.

[16] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.

[17] Leonid Anatolevich Levin. Universal sequential search problems. *Problemy peredachi informatsii*, 9(3):115–116, 1973.

[18] John E Savage. Computational work and time on finite machines. *Journal of the ACM (JACM)*, 19(4):660–674, 1972.

[19] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.

[20] Alan Mathison Turing. On computable numbers, with an application to the entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.

[21] Vijay V Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013.

[22] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, page 681–690, New York, NY, USA, 2006.