

Computational Complexity for Physicists*

S. Mertens[†]

Inst. f. Theor. Physik, Otto-von-Guericke Univ., PF 4120, D-39016 Magdeburg, Germany

Abstract: These lecture notes are an informal introduction to the theory of computational complexity and its links to quantum computing and statistical mechanics.

1 Introduction

1.1 Motivation

Compared to the traditionally close relationship between physics and mathematics, an exchange of ideas and methods between physics and *computer science* barely takes places. The few interactions that go beyond Fortran programming and the quest for faster computers were often successful and provided surprising insights in both fields. This holds particularly for the mutual exchange between statistical mechanics and the theory of computational complexity.

The branch of theoretical computer science known as computational complexity is concerned with classifying problems according to the computational resources (CPU time, memory) required to solve them. This has led to precisely defined notions of *tractable* and *intractable* problems. It appears as if this notions can be transferred to the problem of *analytical solubility* of models from statistical physics, explaining to some extent why for example the Ising model is exactly soluble in two dimensions, but not in three (sec. 3.2).

The success of the theory of computational complexity is based on a pessimistic attitude: a problem's tractability is defined according to the worst possible instance. Quite often this worst case scenario differs considerably from the typical case, averaged over a reasonable ensemble of instances. A common observation is that hard problems are typically easy to solve. To get real hard instances, the parameters of the underlying ensemble must be carefully tuned to certain critical values. Varying the parameters across the critical region leads to abrupt changes in the typical complexity—very similar to the abrupt changes associated with phase

^{*}published in *Computing in Science & Engineering* 4 31–47 (May/June 2002)

[†]stephan.mertens@physik.uni-magdeburg.de

transitions in physical systems. Phase transitions in physical *and* computational systems are best studied within the framework of statistical mechanics (sec. 3.4).

Apart from phase transitions, statistical mechanics offer means for the general probabilistic analysis of computational problems. The usual way is to formulate an optimization problem as a spin glass and analyze the low temperature properties of the latter. This ‘physical’ approach often yield results which go beyond the results obtained by traditional methods (sec. 3.3).

Another exciting link between physics and computer science is provided by *quantum computing*. There is some theoretical evidence that computers using quantum systems as computational devices are more powerfull than computers based on classical devices. The hope is that problems which are intractable on a classical computer become tractable when put on a quantum computer. Results like Shor’s celebrated quantum algorithm for factorization nurture this hope, but a real breakthrough is still missing. Obviously some knowledge of computational complexity helps to understand the promises and limitations of quantum computers (sec. 3.1).

These notes are directed at physicists with no or little knowledge of computational complexity. Of course this is not the first informal introduction into the field, see e.g. [1, 2] or the corresponding chapters in textbooks on algorithms [3, 4]. For a deeper understanding of the field you are referred to the classical textbooks of Garey and Johnson [5] and Papadimitriou [6].

1.2 The measure of complexity

1.2.1 Algorithms

The computational complexity of a problem is a measure of the computational resources, typically time, required to solve the problem. What can be measured (or computed) is the time that a particular *algorithm* uses to solve the problem. This time in turn depends on the implementation of the algorithm as well as on the computer the program is running on. The theory of computational complexity provides us with a notion of complexity that is largely independent of implementational details and the computer at hand. Its precise definition requires a considerable formalism, however. This is not surprising, since it is related to a highly non trivial question that touches the fundamentals of mathematics: *What do we mean by saying a problem is solvable?* Thinking about this question leads to Gödel’s incompleteness theorem, Turing machines and the Church Turing Thesis on computable functions. See [7] for an entertaining introduction into these topics written by a physicist.

Here we will adopt a more informal, pragmatic point of view. A problem is solvable if it can be solved by a computer program written in your favourite programming language. The running time or *time complexity* of your program must then be defined with some care to serve as a meaningful measure for the complexity of the problem.

1.2.2 Time complexity

In general the running time depends on the size of the problem and on the specific input data, the instance. Sorting 1000 numbers takes longer than sorting 10 numbers. Some sorting algorithms run faster if the input data is partially sorted already. To minimize the dependency

on the specific instance we consider the *worst case time complexity* $T(n)$,

$$T(n) = \max_{|x|=n} t(x) \tag{1}$$

where $t(x)$ is the running time of the algorithm for input data x (in arbitrary units) and the maximum is taken over all problem instances of size n . The worst case time is an upper bound for the observable running time.

A measure of time complexity should be based on a unit of time that is independent of the clock rate of a specific CPU. Such a unit is provided by the time it takes to perform an elementary operation like the addition of two integer numbers. Measuring the time in this unit means counting the number of elementary operations executed by your algorithm. This number in turn depends strongly on the implementation details of the algorithm – smart programmers and optimizing compilers will try to reduce it. Therefore we will not consider the precise number $T(n)$ of elementary operations but only the asymptotic behavior of $T(n)$ for large values of n as denoted by the Landau symbols \mathcal{O} and Θ :

- We say $T(n)$ is of order at most $g(n)$ and write $T(n) = \mathcal{O}(g(n))$ if there exist positive constants c and n_0 such that $T(n) \leq cg(n)$ for all $n \geq n_0$.
- We say $T(n)$ is of order $g(n)$ and write $T(n) = \Theta(g(n))$ if there exist positive constants c_1, c_2 and n_0 such that $c_1g(n) \leq T(n) \leq c_2g(n)$ for all $n \geq n_0$.

Multiplying two $n \times n$ matrixes requires n^3 multiplications according to the textbook formula. Does this mean that the problem of multiplying two $n \times n$ matrices has complexity $\Theta(n^3)$? No. The textbook formula is a particular algorithm, and the time complexity of an algorithm is only an *upper bound* for the inherent complexity of a problem. In fact faster matrix multiplication algorithms with complexity $\mathcal{O}(n^\alpha)$ and $\alpha < 3$ have been found during the last decades, the current record being $\alpha = 2.376$ [8]. Since the product matrix has n^2 entries, α can not be smaller than 2, and it is an open question whether this lower bound can be achieved by an algorithm. A problem where the upper bound from algorithmic complexity meets an inherent lower bound is sorting n items. Under the general assumption that comparisons between pairs of items are the only source of information about the items, it can be shown that $\Theta(n \log n)$ is a lower bound for the number of comparisons to sort n items in the worst case [3, chap. 9]. This bound is met by algorithms like heapsort or mergesort.

1.2.3 Problem size

Our measure of time complexity still depends on the somewhat ambiguous notion of problem size. In the matrix multiplication example we tacitly took the number n of rows of one input matrix as the “natural” measure of size. Using the number of elements $m = n^2$ instead will “speed up” the $\mathcal{O}(n^3)$ algorithm to $\mathcal{O}(m^{3/2})$ without changing a single line of program code. You see that an unambiguous definition of problem size is required to compare algorithms.

In computational complexity, all problems which can be solved by a polynomial algorithm, i.e. an algorithm with time complexity $\Theta(n^k)$ for some k , are lumped together and called *tractable*. Problems which can only be solved by algorithms with non-polynomial running time like $\Theta(2^n)$ or $\Theta(n!)$ are also lumped together and called *intractable*. There are

practical as well as theoretical reasons for this rather coarse classification. One of the theoretical advantages is that it does not distinguish between the $\mathcal{O}(n^3)$ - and $\mathcal{O}(n^{3/2})$ -algorithm form above, hence we can afford some sloppiness and stick with our ambiguous “natural” measure of problem size.

1.3 Tractable and intractable problems

1.3.1 Polynomial vs. exponential growth

From a practical point of view, an exponential algorithm means a hard limit for the accessible problem size. Suppose that with your current equipment and you can solve a problem of size n just within the schedule. If your algorithm has complexity $\Theta(2^n)$, a problem of size $n + 1$ will need twice the time, bringing you definitely out of schedule. The increase in time caused by an $\Theta(n)$ or $\Theta(n^2)$ algorithm on the other hand is far less dramatic and can easily be compensated by upgrading your hardware. You might object that a $\Theta(n^{100})$ algorithm outperforms a $\Theta(2^n)$ algorithm only for problem sizes that will never occur in your application. A polynomial algorithm for a problem usually goes hand in hand with a mathematical *insight* into the problem which enables you to find a polynomial algorithm with small degree, typically $\Theta(n^k)$, $k = 1, 2$ or 3 . Polynomial algorithms with $k > 10$ are rare and arise in rather esoteric problems.

1.3.2 Tractable trees

As an example consider the following problem from network design. You have a business with several offices and you want to lease phone lines to connect them up with each other. The phone company charges different amounts of money to connect different pairs of cities, and your task is to select a set of lines that connects all your offices with a minimum total cost.

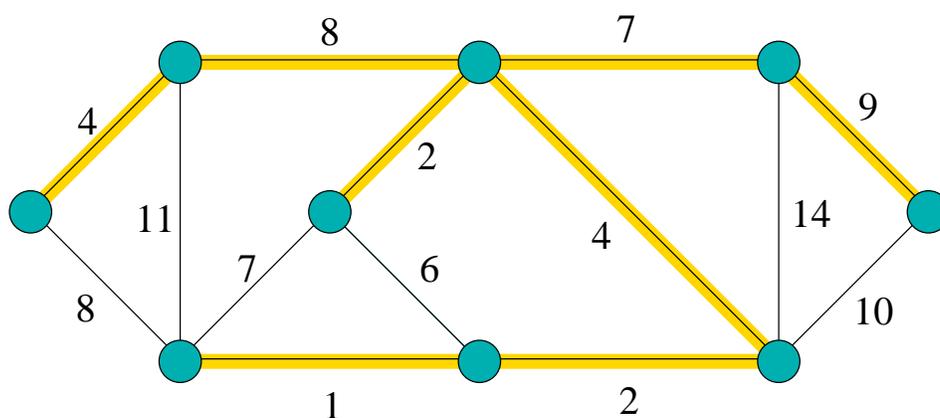


Figure 1: A weighted graph and its minimum spanning tree (colored edges).

In mathematical terms the cities and the lines between them form the vertices V and edges E of a weighted graph $G = (V, E)$, the weight of an edge being the leasing costs of the corresponding phone line. Your task is to find a subgraph that connects all vertices in the

graph, i.e. a spanning subgraph, and whose edges have minimum total weight. Your subgraph should not contain cycles, since you can always remove an edge from a cycle keeping all nodes connected and reducing the cost. A graph without cycles is a tree, so what you are looking for is a *minimum spanning tree* in a weighted graph (Fig. 1).

MINIMUM SPANNING TREE (MST): Given a weighted graph $G = (V, E)$ with non-negative weights. Find a spanning Tree $T \subseteq G$ with minimum total weight.

How to find a minimum spanning tree? A naive approach is to generate all possible trees with n vertices and keep the one with minimal weight. The enumeration of all trees can be done using Prüfer codes [9], but Cayley's formula tells us that there are n^{n-2} different trees with n vertices. Already for $n = 100$ there are more trees than atoms in the observable universe! Hence exhaustive enumeration is prohibitive for all but the smallest trees. The mathematical insight that turns MST into a tractable problem is this:

Lemma: Let $U \subset V$ be any subset of the vertices of $G = (V, E)$, and let e be the edge with the smallest weight of all edges connecting U and $V - U$. Then e is part of the minimum spanning tree.

Proof: By contradiction. Suppose T is a minimum spanning tree not containing e . Let $e = (u, v)$ with $u \in U$ and $v \in V - U$. Then because T is a spanning tree it contains a unique path from u to v , which together with e forms a cycle in G . This path has to include another edge f connecting U and $V - U$. Now $T + e - f$ is another spanning tree with less total weight than T . So T was not a minimum spanning tree.

The lemma allows to grow a minimum spanning tree edge by edge, using Prim's algorithm for example:

```

PRIM( $G$ )
Input: weighted graph  $G(V, E)$ 
Output: minimum spanning tree  $T \subseteq G$ 
begin
  Let  $T$  be a single vertex  $v$  from  $G$ 
  while  $T$  has less than  $n$  vertices
    find the minimum edge connecting  $T$  to  $G - T$ 
    add it to  $T$ 
  end
end

```

The precise time complexity of Prim's algorithm depends on the data structure used to organize the edges, but in any case $\mathcal{O}(n^2 \log n)$ is an upper bound. (see [10] for faster algorithms). Equipped with such a polynomial algorithm you can find minimum spanning trees with thousands of nodes within seconds on a personal computer. Compare this to exhaustive search! According to our definition, MST is a tractable problem.

1.3.3 Intractable itineraries

Encouraged by the efficient algorithm for MST we will now investigate a similar problem. Your task is to plan an itinerary for a travelling salesman who must visit n cities. You are given a map with all cities and the distances between them. In what order should the salesman visit the cities to minimize the total distance he has to travel? You number the cities arbitrarily and write down the distance matrix (d_{ij}) , where d_{ij} denotes the distance between city number i and city number j . A tour is given by a *cyclic permutation* $\pi : [1 \dots n] \mapsto [1 \dots n]$, where $\pi(i)$ denotes the successor of city i , and your problem can be defined as:

TRAVELING SALESMAN PROBLEM (TSP): Given a $n \times n$ distance matrix with elements $d_{ij} \geq 0$. Find a cyclic permutation $\pi : [1 \dots n] \mapsto [1 \dots n]$ that minimizes

$$c_n(\pi) = \sum_{i=1}^n d_{i\pi(i)} \quad (2)$$

The TSP probably is the most famous optimization problem, and there exists a vast literature specially devoted to it, see [11–14] and references therein. It is not very difficult to find good solutions, even to large problems, but how can we find the *best* solution for a given instance? There are $(n - 1)!$ cyclic permutations, calculating the length of a single tour can be done in time $\mathcal{O}(n)$, hence exhaustive search has complexity $\mathcal{O}(n!)$. Again this approach is limited to very small instances. Is there a mathematical insight that provides us with a shortcut to the optimum solution, like for MST? Nobody knows! Despite the efforts of many brilliant people, no polynomial algorithm for the TSP has been found. There are some smart and efficient (i.e. polynomial) algorithms that find good solutions but do not guarantee to yield the optimum [12]. According to our definition, the TSP is intractable.

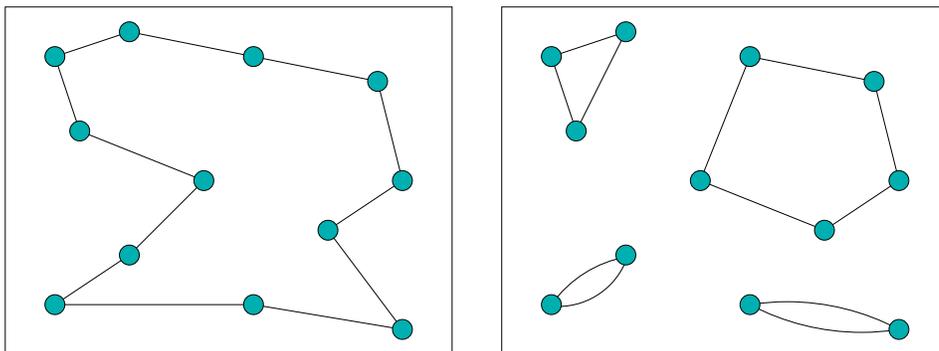


Figure 2: Same instance, different problems: A valid configuration of the TRAVELING SALESMAN PROBLEM (left) and the ASSIGNMENT problem (right). Whereas ASSIGNMENT can be solved in polynomial time, the TSP is intractable.

Why is the TSP intractable? Again, nobody knows. There is no *proof* that excludes the existence of a polynomial algorithm for TSP, so maybe someday someone will come up with

a polynomial algorithm and the corresponding mathematical insight. This is very unlikely, however, as we will see soon.

The intractability of the TSP astonishes all the more considering the tractability of a very similar, almost identical problem:

ASSIGNMENT: Given a $n \times n$ cost matrix with elements $d_{ij} \geq 0$. Find a permutation $\pi : [1 \dots n] \mapsto [1 \dots n]$ that minimizes

$$c_n(\pi) = \sum_{i=1}^n d_{i\pi(i)} \quad (3)$$

The only difference between TSP and ASSIGNMENT is that the latter allows all permutations on n items, not only the cyclic ones. If the d_{ij} denote distances between cities, ASSIGNMENT corresponds to total tour length minimization for a variable number of salesmen, each travelling his own *subtour* (Fig. 2).

The classical application of ASSIGNMENT is the assignment of n tasks to n workers, subject to the constraint that each worker is assigned exactly one task. Let d_{ij} denote the cost of having task j performed by worker i , and $\pi(i)$ denote the task assigned to worker i , ASSIGNMENT is the problem of minimizing the total cost.

There are $n!$ possible assignments of n tasks to n workers, hence exhaustive enumeration again is prohibitive. In contrast to the TSP, however, ASSIGNMENT can be solved in polynomial time, for example using the $\mathcal{O}(n^3)$ hungarian algorithm [4]. Compared to MST, the algorithm and the underlying mathematical insight are a bit more involved and will not be discussed here.

2 Complexity classes

2.1 Decision problems

So far we have discussed optimization problems: solving MST, TSP or ASSIGNMENT implies that we compare an exponential number of feasible solutions with each other and pick the optimum. Exhaustive search does this explicitly, polynomial shortcuts implicitly. Maybe we learn more about the barrier that separates tractable from intractable problems if we investigate simpler problems, problems where the solutions are recognizable without explicit or implicit comparison to all feasible solutions. So let us consider *decision problems*, problems whose solution is either ‘yes’ or ‘no’.

Any optimization problem can be turned into a decision problem by adding a bound B to the instance. Examples:

MST (DECISION): Given a weighted graph $G = (V, E)$ with non-negative weights and a number $B \geq 0$. Does G contain a spanning tree T with total weight $\leq B$?

TSP (DECISION): Given a $n \times n$ distance matrix with elements $d_{ij} \geq 0$ and a number $B \geq 0$. Is there a tour π with length $\leq B$?

In a decision problem, the feasible solutions are not evaluated relative to each other but to an ‘absolut’ criterion: a tour in the TSP either has length $\leq B$ or not.

MST(D) can be solved in polynomial time: simply solve the optimization variant MST and compare the result to the parameter B . For the TSP(D) this approach does not help. In fact we will see in section 2.3.1 that there exists a polynomial algorithm for TSP(D) if and only if there exists a polynomial algorithm for TSP. It seems as if we cannot learn more about the gap between tractable and intractable problems by considering decision variants of optimization problems. So lets look at other decision problems, not derived from optimization problems.

2.1.1 Eulerian circuits

Our first ‘genuine’ decision problem dates back into the 18th-century, where in the city of Königsberg (now Kaliningrad) seven bridges crossed the river Pregel and its two arms (Fig. 3). A popular puzzle of the time asked if it was possible to walk through the city crossing each of the bridges exactly once and returning home.

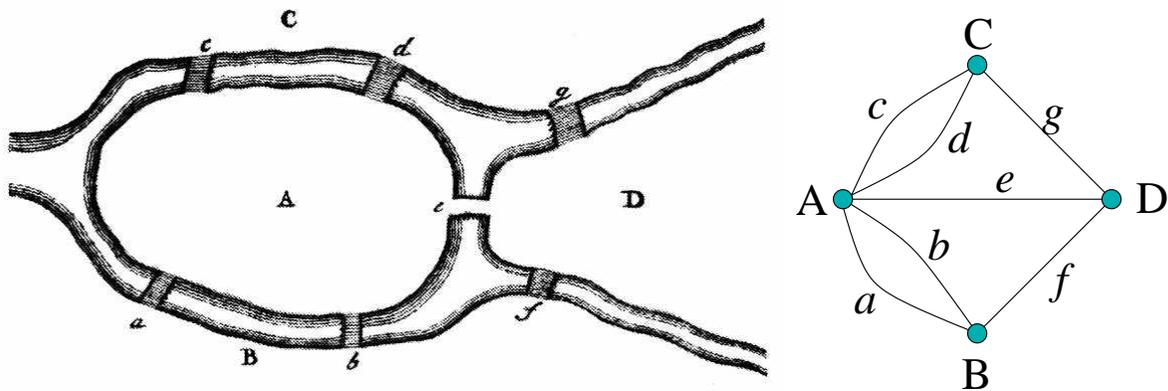


Figure 3: The seven bridges of Königsberg, as drawn in Euler’s paper from 1736 [15] (left) and represented as a graph (right). In the graph, the riverbanks and islands are condensed to points (vertices), and each of the bridges is drawn as a line (edge).

It was Leonhard Euler who solved this puzzle in 1736 [15]. First of all Euler recognizes that for the solution of the problem the only thing that matters is the pattern of interconnections of the banks and islands—a graph $G = (V, E)$ in modern terms. The graph corresponding to the puzzle of the Königsberg bridges has 4 vertices for the two banks and the two islands and 7 edges for the bridges (Fig. 3). Euler’s paper on the Königsberg bridges can be regarded as the birth of graph theory.

To generalize the Königsberg bridges problem we need some terminology from graph theory [9]. A *walk* in a graph $G = (V, E)$ is an alternating sequence of vertices $v \in V$ and edges $(v, v') \in E$,

$$v_1, (v_1, v_2), v_2, (v_2, v_3), \dots, (v_{l-1}, v_l), v_l.$$

Note that the sequence begins and ends with a vertex, and each edge is incident with the vertices immediately preceding and succeeding it. A walk is termed *closed* if $v_l = v_1$, *open* otherwise. A walk is called a *trail* if all its edges are distinct, and a closed trail is called a *circuit*. What the strollers in Königsberg tried to find was a circuit that contains all edges. To the honour of Leonhard Euler such a circuit is called *Eulerian circuit*. Equipped with all this terminology we are ready to define the generalization of the Königsberg bridges problem:

EULERIAN CIRCUIT: Given a graph $G = (V, E)$. Does G contain an Eulerian circuit?

Obviously this is a decision problem: the answer is either ‘yes’ or ‘no’, and any circuit can be checked to be Eulerian or not without resorting to all possible circuits.

Once again exhaustive search would solve this problem, but the intractability of this approach was already noticed by Euler. More than 200 years before the advent of computers he wrote “The particular problem of the seven bridges of Königsberg could be solved by carefully tabulating all possible paths, thereby ascertaining by inspection which of them, if any, met the requirement. This method of solution, however, is too tedious and too difficult because of the large number of possible combinations, and in other problems where many more bridges are involved it could not be used at all.” (cited from [1]).

Euler solved the Königsberg bridges problem not by listing all possible trails but by mathematical insight. He noticed that in a circuit you must leave each vertex via an edge different from the edge that has taken you there. In other words, the degree of the vertex (that is the number of edges adjacent to the vertex) must be even. This is obviously a necessary condition, but Euler proved that it is also sufficient:

Theorem: A connected graph $G = (V, E)$ contains an Eulerian circuit if and only if the degree of every vertex $v \in V$ is even.

Euler’s theorem allows us to devise a polynomial algorithm for EULERIAN CIRCUIT: Check the degree of every vertex in the graph. If one vertex has an odd degree, return ‘no’. If all vertices have been checked having even degree, return ‘yes’. The running time of this algorithm depends on the encoding of the graph. If $G = (V, E)$ is encoded as a $|V| \times |V|$ adjacency matrix with entries a_{ij} = number of edges connecting v_i and v_j , the the running time is $\mathcal{O}(|V|^2)$. Thanks to Euler, EULERIAN CIRCUIT is a tractable problem. The burghers of Königsberg on the other hand had to learn from Euler, that they would never find a walk through their hometown crossing each of the seven bridges exactly once.

2.1.2 Hamiltonian cycles

Another decision problem is associated with the mathematician and Astronomer Royal of Ireland, Sir William Rowan Hamilton. In the year 1859, Hamilton put on the market a new puzzle called the Icosian game (Fig. 4).

The generalization of the Icosian game calls for some definitions from graph theory: A closed walk in a graph is called a *cycle* if all its vertices (except the first and the last) are distinct. A Hamiltonian cycle is a cycle that contains all vertices of a graph. The generalization of the Icosian game then reads

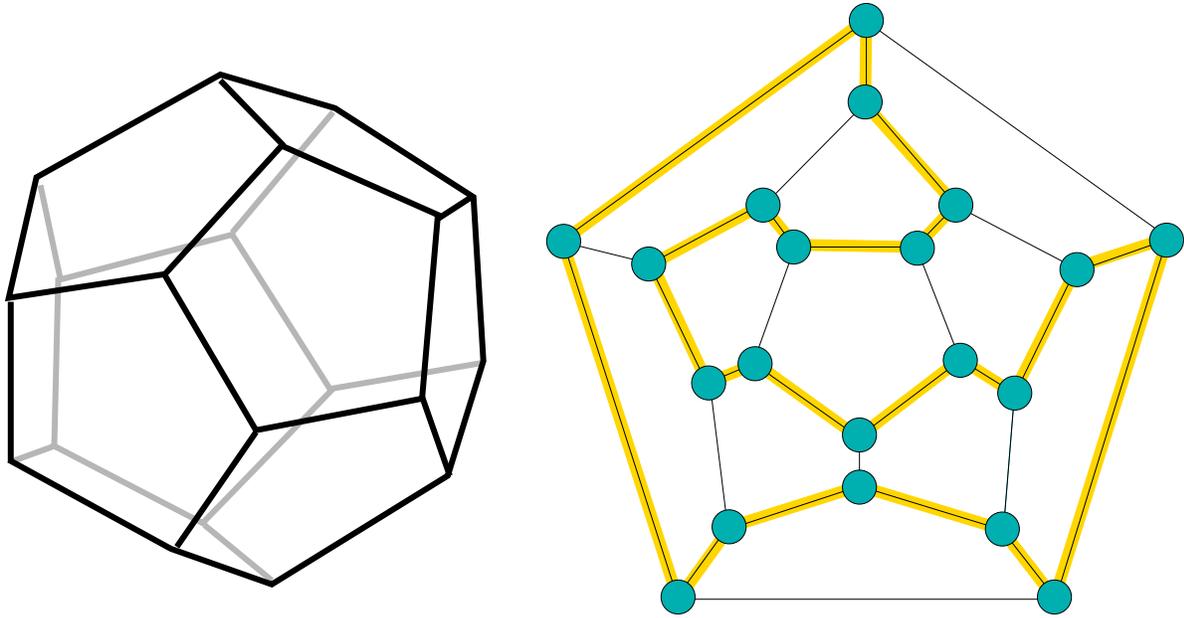


Figure 4: Sir Hamilton's Icosian game: Find a route along the edges of of the dodecahedron (left), passing each corner exactly once and returning to the starting corner. A solution is indicated (shaded edges) in the planar graph that is isomorphic to the dodecahedron (right).

HAMILTONIAN CYCLE: Given a graph $G = (V, E)$. Does G contain a Hamiltonian cycle?

There is a certain similarity between EULERIAN CIRCUIT and HAMILTONIAN CYCLE. In the former we must pass each edge once; in the latter, each vertex once. Despite this resemblance the two problems represent entirely different degrees of difficulty. The available mathematical insights into HAMILTONIAN CYCLE provide us neither with a polynomial algorithm nor with a proof that such an algorithm is impossible. HAMILTONIAN CYCLE is intractable, and nobody knows why.

2.1.3 Coloring

Imagine we wish to arrange talks in a congress in such a way that no participant will be forced to miss a talk he would like to hear. Assuming a good supply of lecture rooms enabling us to hold as many parallel talks as we like, can we finish the programme within k time slots? This problem can be formulated in terms of graphs: Let G be a graph whose vertices are the talks and in which two talks are adjacent (joined by an edge) if and only if there is a participant wishing to attend both. Your task is to assign one of the k time slots to each vertex in such a way that adjacent vertices have different time slots. The common formulation of this problem uses colors instead of time slots:

k -COLORING: Given a graph $G = (V, E)$. Is there a coloring of the vertices of G using at most k different colors such that no two adjacent vertices have the same color?

When $k = 2$ this problem is tractable — the construction of a polynomial algorithm is left as easy exercise. For $k = 3$ things change considerably: 3-COLORING is intractable. Note that for larger k the problem gets easier again: a *planar* graph is always colorable with 4 colors! This is the famous 4-color Theorem. 3-COLORING remains intractable even when restricted to planar graphs.

2.1.4 Satisfiability

We close this section with a decision problem that is not from graph theory but from Boolean logic. A Boolean variable x can take on the value 0 (false) or 1 (true). Boolean variables can be combined in *clauses* using the Boolean operators

- NOT $\bar{\cdot}$ (negation): the clause \bar{x} is true ($\bar{x} = 1$) if and only if x is false ($x = 0$).
- AND \wedge (conjunction): the clause $x_1 \wedge x_2$ is true ($x_1 \wedge x_2 = 1$) if and only if both variables are true: $x_1 = 1$ and $x_2 = 1$
- OR \vee (disjunction): the clause $x_1 \vee x_2$ is true ($x_1 \vee x_2 = 1$) if and only if at least one of the variables is true: $x_1 = 1$ or $x_2 = 1$.

A variable x or its negation \bar{x} is called a *literal*. Different clauses can be combined to yield complex Boolean formulas, e.g.

$$F_1(x_1, x_2, x_3) = (x_1 \vee \bar{x}_2 \vee x_3) \wedge (x_2 \vee \bar{x}_3) \wedge (\bar{x}_1 \vee x_2) \wedge (\bar{x}_1 \vee \bar{x}_3). \quad (4)$$

A Boolean formula evaluates to either 1 or 0, depending on the assignment of the Boolean variables. In the example above $F_1 = 1$ for $x_1 = 1, x_2 = 1, x_3 = 0$ and $F_1 = 0$ for $x_1 = x_2 = x_3 = 1$. A formula F is called *satisfiable*, if there is at least one assignment of the variables such that the formula is true. F_1 is satisfiable,

$$F_2(x_1, x_2) = (\bar{x}_1 \vee x_2) \wedge \bar{x}_2 \wedge x_1 \quad (5)$$

is not satisfiable.

Every Boolean formula can be written in *conjunctive normal form* (CNF) i.e. as a set of clauses C_k combined exclusively with the AND-operator

$$F = C_1 \wedge C_2 \wedge \cdots \wedge C_m \quad (6)$$

where the literals in each clause are combined exclusively with the OR-operator. The examples F_1 and F_2 are both written in CNF. Each clause can be considered as a constraint on the variables, and satisfying a formula means satisfying a set of (possibly conflicting) constraints simultaneously. Hence

SATISFIABILITY (SAT): Given disjunctive clauses C_1, C_2, \dots, C_m of literals, where a literal is a variable or negated variable from the set $\{x_1, x_2, \dots, x_n\}$. Is there an assignment of variables that satisfies all clauses simultaneously?

can be considered as prototype of a constraint satisfaction problem [16]. Fixing the number of literals in each clause leads to

k -SAT: Given disjunctive clauses C_1, C_2, \dots, C_m of k literals each, where a literal is a variable or negated variable from the set $\{x_1, x_2, \dots, x_n\}$. Is there an assignment of variables that satisfies all clauses simultaneously?

Polynomial algorithms are known for 1-SAT and 2-SAT [17]. No polynomial algorithm is known for general SAT and k -SAT if $k > 2$.

2.2 The classes \mathcal{P} and \mathcal{NP}

2.2.1 Tractable problems

Now we have seen enough examples to introduce two important complexity classes for decision problems.

The class of tractable decision problems is easy to define: it consists of those problems, for which a polynomial algorithm is known. The corresponding class is named \mathcal{P} for “polynomial”:

Definition: A decision problem P is element of the class \mathcal{P} if and only if it can be solved by a polynomial time algorithm.

EULERIAN CIRCUIT, 2-COLORING, MST(D) etc. are in \mathcal{P} .

2.2.2 Nondeterministic algorithms

The definition of the second complexity class involves the concept of a *nondeterministic algorithm*. A nondeterministic algorithm is like an ordinary algorithm, except that it may use one additional, very powerful instruction [18]:

goto both label 1, label 2

This instruction splits the computation into two parallel processes, one continuing from each of the instructions indicated by “label 1” and “label 2”. By encountering more and more such instructions, the computation will branch like a tree into a number of parallel computations that potentially can grow as an exponential function of the time elapsed (see Fig. 5). A nondeterministic algorithm can perform an exponential number of computations in polynomial time! In the world of conventional computers, nondeterministic algorithms are a *theoretical concept* only, but in quantum computing this may change (section 3.1). We need the concept of nondeterminism to define the class \mathcal{NP} of “nondeterministic polynomial” problems:

Definition: A decision problem P is in the class \mathcal{NP} , if and only if it can be solved in *polynomial* time by a *nondeterministic algorithm*.

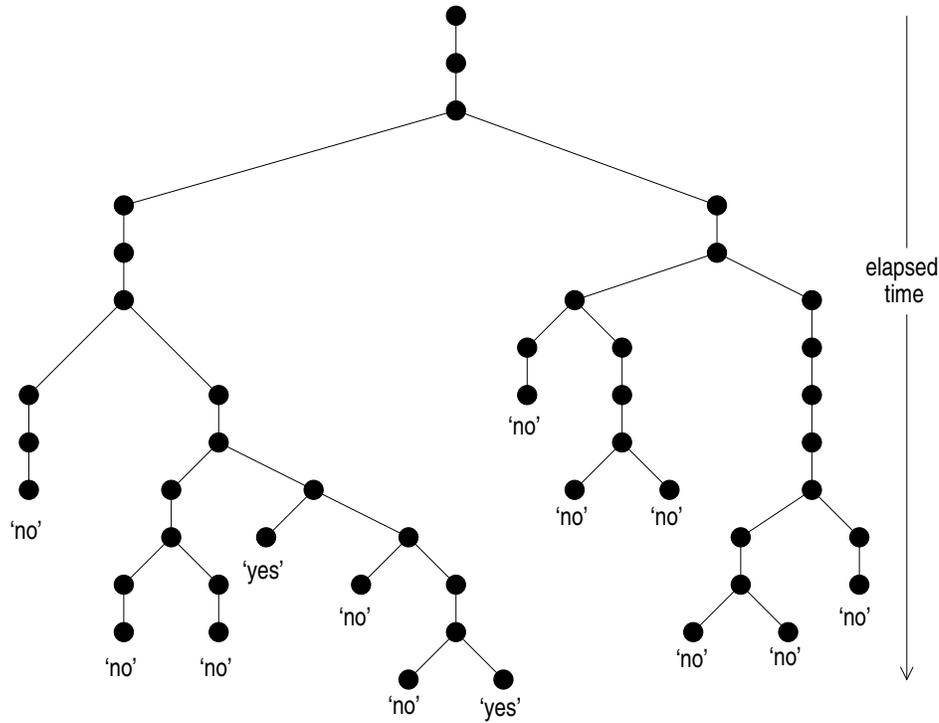


Figure 5: Example of the execution history of a nondeterministic algorithm.

Solubility by a nondeterministic algorithm means this: All branches of the computation will stop, returning either ‘yes’ or ‘no’. We say that the overall algorithm returns ‘yes’, if *any* of its branches returns ‘yes’. The answer is ‘no’, if none of the branches reports ‘yes’. We say that a nondeterministic algorithm solves a decision problem in polynomial time, if the number of steps used by the first of the branches to report ‘yes’ is bounded by a polynomial in the size of the problem.

We require polynomial solubility only for the ‘yes’-instances of decision problem. This asymmetry between ‘yes’ and ‘no’ reflects the asymmetry between the ‘there is’- and ‘for all’-quantifiers in decision problems: a graph G is a ‘yes’-instance of HAMILTONIAN CYCLE, if *there is* at least one Hamiltonian cycle in G . For a ‘no’-instance, *all* cycles in G have to be non Hamiltonian.

Note that the conventional (deterministic) algorithms are special cases of a nondeterministic algorithms (those nondeterministic algorithms that do not use the **goto both** instruction). It follows immediately that $\mathcal{P} \subseteq \mathcal{NP}$.

All decision problems we have discussed in the preceding section are members of \mathcal{NP} . Here’s a nondeterministic polynomial algorithm for SAT:

```

SATISFIABILITY( $F$ )
Input: Boolean formula  $F(x_1, \dots, x_n)$ 
Output: 'yes' if  $F$  is satisfiable, 'no' otherwise
begin
  for  $i = 1$  to  $n$ 
    goto both label 1, label 2
    label 1:  $x_i = \text{true}$ ; continue
    label 2:  $x_i = \text{false}$ ; continue
  end
  if  $F(x_1, \dots, x_n) = \text{true}$  then return 'yes'
    else return 'no'
end

```

The **for**-loop branches at each iteration: in one branch $x_i = \text{true}$, in the other branch $x_i = \text{false}$ (the **continue** instruction starts the next iteration of the loop). After executing the **for**-loop we have 2^n branches of computation, one for each of the possible assignments of n Boolean variables.

The power of nondeterministic algorithms is that they allow the exhaustive enumeration of an exponentially large number of candidate solutions in polynomial time. If the evaluation of each candidate solution (calculating $F(x_1, \dots, x_n)$ in the above example) in turn can be done in polynomial time, the total nondeterministic algorithm is polynomial. For a problem from the class \mathcal{NP} , the sole source of intractability is the exponential size of the search space.

2.2.3 Succinct certificates

There is a second, equivalent definition of \mathcal{NP} , based on the notion of a *succinct certificate*. A certificate is a proof. If you claim that a graph G has a Hamiltonian cycle, you can proof your claim by providing a Hamiltonian cycle. Certificates for EULERIAN CIRCUIT and k -COLORING are an Eulerian circuit and a valid coloring. A certificate is succinct, if its size is bounded by a polynomial in the size of the problem. The second definition then reads

Definition: A decision problem P is element of the class \mathcal{NP} if and only if for every 'yes'-instance of P there exists a *succinct certificate* that can be verified in polynomial time.

The equivalence of both definitions can easily be shown [18]. The idea is that a succinct certificate can be used to deterministically select the branch in a nondeterministic algorithm that leads to a 'yes'-output.

The definition based on nondeterministic algorithms reveals the key feature of the class \mathcal{NP} more clearly, but the second definition is more usefull for proving that a decision problem is in \mathcal{NP} . As an example consider

COMPOSITENESS: Given a positive integer N . Are there integer numbers $p > 1$ and $q > 1$ such that $N = pq$?

A certificate of a ‘yes’ instance N of COMPOSITENESS is a factorization $N = pq$. It is succinct, because the number of bits in p and q is less or equal the number of bits in N , and it can be verified in quadratic time by multiplication. Hence $\text{COMPOSITENESS} \in \mathcal{NP}$.

Most decision problems ask for the existence of an object with a given property, like a cycle which is Hamiltonian or a factorization with integer factors. In these cases, the desired object may serve as a succinct certificate. For some problems this does not work, however, like for

PRIMALITY: Given a positive integer N . Is N prime?

PRIMALITY is the negation or complement of COMPOSITENESS: the ‘yes’-instances of the former are the ‘no’-instances of the latter and vice versa. A succinct certificate for PRIMALITY is by no means obvious. In fact, for many decision problems in \mathcal{NP} no succinct certificate is known for the complement, i.e. it is not known whether the complement is also in \mathcal{NP} . For PRIMALITY however, one can construct a succinct certificate based on Fermat’s Theorem [19]. Hence $\text{PRIMALITY} \in \mathcal{NP}$.

2.2.4 A first map of complexity

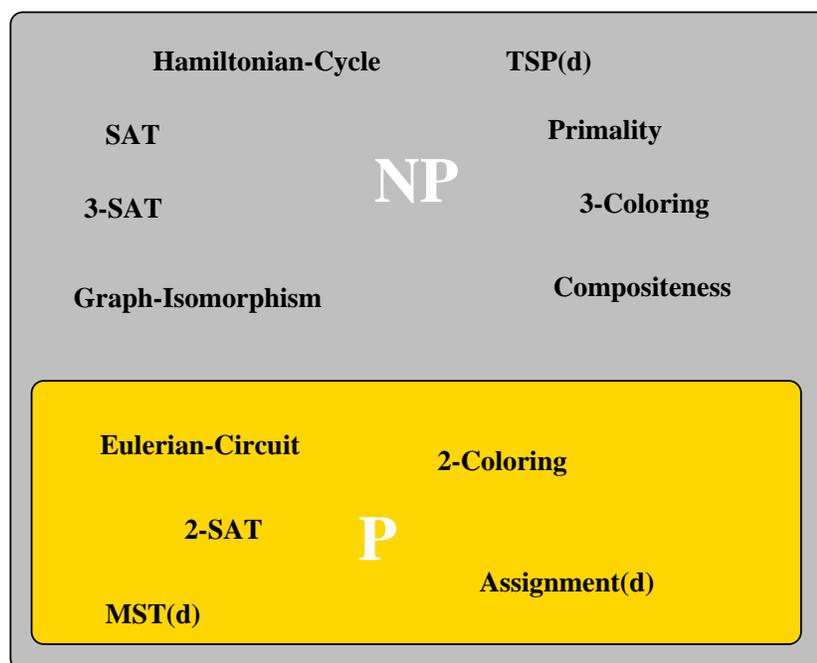


Figure 6: A first map of complexity. All problems indicated are defined within the text. Problems with a (D) are decision variants of optimization problems.

Fig. 6 summarizes what we have achieved so far. The class \mathcal{NP} consists of all decision problems whose sole source of difficulty is the size of the search space which grows exponentially with the size of the problem. These problems are intractable unless some mathematical

insight provides us with a polynomial shortcut to avoid exhaustive search. Such an insight promotes a problem into the class \mathcal{P} of polynomially soluble problems.

The class \mathcal{NP} not only contains a large number of problems with important applications, but additionally represents a real challenge: all problems in \mathcal{NP} still have a chance to be in \mathcal{P} . A proof of non-existence of a polynomial algorithm for a single problem from \mathcal{NP} would establish that $\mathcal{P} \neq \mathcal{NP}$. As long as such a proof is missing,

$$\mathcal{P} \stackrel{?}{=} \mathcal{NP} \tag{7}$$

represents the most famous open conjecture in theoretical computer science.

2.3 \mathcal{NP} -completeness

2.3.1 Polynomial reductions

The computational complexities of two problems P_1 and P_2 can be related to each other by constructing an algorithm for P_1 that uses an algorithm for P_2 as a “subroutine”. Consider the following algorithm that relates HAMILTONIAN CYCLE to TSP(D):

```

HAMILTONIAN CYCLE( $G$ )
Input: Graph  $G = (V, E)$ 
Output: 'yes' if  $G$  contains a Hamiltonian cycle, 'no' otherwise
(1)  begin
(2)     $n := |V|$ 
(3)    for  $i = 1$  to  $n$ 
(4)      for  $j = 1$  to  $n$ 
(5)        if  $(v_i, v_j) \in E$  then  $d_{ij} := 1$ 
(6)          else  $d_{ij} := 2$ 
(7)    if TSP-decision  $(d, B := n) = \text{'yes'}$  then return 'yes'
(8)    else return 'no'
(9)  end

```

This algorithm solves HAMILTONIAN CYCLE by solving an appropriate instance of TSP(D). In the **for**-loops (lines 3-5) a distance matrix d is set up with entries $d_{ij} = 1$ if there is an edge (v_i, v_j) in G and $d_{ij} = 2$ otherwise. Now a Hamiltonian cycle in G is valid tour in the corresponding TSP with all intercity distances having length 1, i.e. with total length n . Conversely, suppose that the TSP has a tour of length n . Since the intercity distances are either 1 or 2 and a tour sums up n such distances, a total length n implies that each pair of successively visited cities must have distance 1, i.e. the tour follows existing edges in G and corresponds to a Hamiltonian cycle. Hence the call to a subroutine that solves TSP(D) (line 7) yields a solution to HAMILTONIAN CYCLE.

How does this algorithm relate the computational complexity of HAMILTONIAN CYCLE to that of TSP(D)? Note that this is a polynomial algorithm if the call to the TSP(D)-solver

is considered as elementary operation. If someone comes up with a polynomial algorithm for TSP(D), we will instantly have a polynomial algorithm for HAMILTONIAN CYCLE! We say that HAMILTONIAN CYCLE is *polynomially reducible* to TSP(D) and write

$$\text{HAMILTONIAN CYCLE} \leq \text{TSP(D)}. \quad (8)$$

In many books, polynomial reducibility is denoted by ‘ \propto ’ instead of ‘ \leq ’. We follow [20] and use ‘ \leq ’ because this notation stresses an important consequence of polynomial reducibility: the existence of a polynomial reduction from P_1 to P_2 excludes the possibility that P_2 can be solved in polynomial time, but P_1 cannot. Hence $P_1 \leq P_2$ can be read as *P_1 cannot be harder than P_2* . Here is the (informal) definition:

Definition: We say a problem P_1 is *polynomially reducible* to a problem P_2 and write $P_1 \leq P_2$ if there exists a polynomial algorithm for P_1 provided there is a polynomial algorithm for P_2 .

2.3.2 \mathcal{NP} -complete problems

Here are some other polynomial reductions that can be verified similar to Eq. 8:

$$\begin{aligned} \text{SAT} &\leq \text{3-SAT} \\ \text{3-SAT} &\leq \text{3-COLORING} \\ \text{3-COLORING} &\leq \text{HAMILTONIAN CYCLE} \end{aligned} \quad (9)$$

See [5, 21] for the corresponding reduction algorithms. Polynomial reducibility is transitive: $P_1 \leq P_2$ and $P_2 \leq P_3$ imply $P_1 \leq P_3$. From transitivity and Eqs. 8 and 9 it follows that each of SAT, 3-SAT, 3-COLORING and HAMILTONIAN CYCLE reduces to TSP(D), i.e. a polynomial algorithm for TSP(D) implies a polynomial algorithm for all these problems! This is amazing, but only the beginning. The true scope of polynomial reducibility was revealed by Stephen Cook in 1971 [22] who proved the following, remarkable theorem:

Theorem: (Cook, 1971) All problems in \mathcal{NP} are polynomially reducible to SAT,

$$\forall P \in \mathcal{NP} : P \leq \text{SAT} \quad (10)$$

This theorem means that

1. no problem in \mathcal{NP} is harder than SAT or SAT is among the hardest problems in \mathcal{NP} .
2. a polynomial algorithm for SAT would imply a polynomial algorithm for *every* problem in \mathcal{NP} , i.e. would imply $\mathcal{P} = \mathcal{NP}$.

It seems as if SAT is very special, but according to transitivity and Eqs. 9 and 8 it can be replaced by 3-SAT, 3-COLORING, HAMILTONIAN CYCLE or TSP(D). These problems form a new complexity class:

Definition: A problem P is called \mathcal{NP} -complete if $P \in \mathcal{NP}$ and $Q \leq P$ for all $Q \in \mathcal{NP}$

The class of \mathcal{NP} -complete problems collects the hardest problems in \mathcal{NP} . If any of them has an efficient algorithm, then *every* problem in \mathcal{NP} can be solved efficiently, i.e. $\mathcal{P} = \mathcal{NP}$. This is extremely unlikely, however, considered the futile efforts of many brilliant people to find polynomial algorithms for problems like HAMILTONIAN CYCLE or TSP(D).

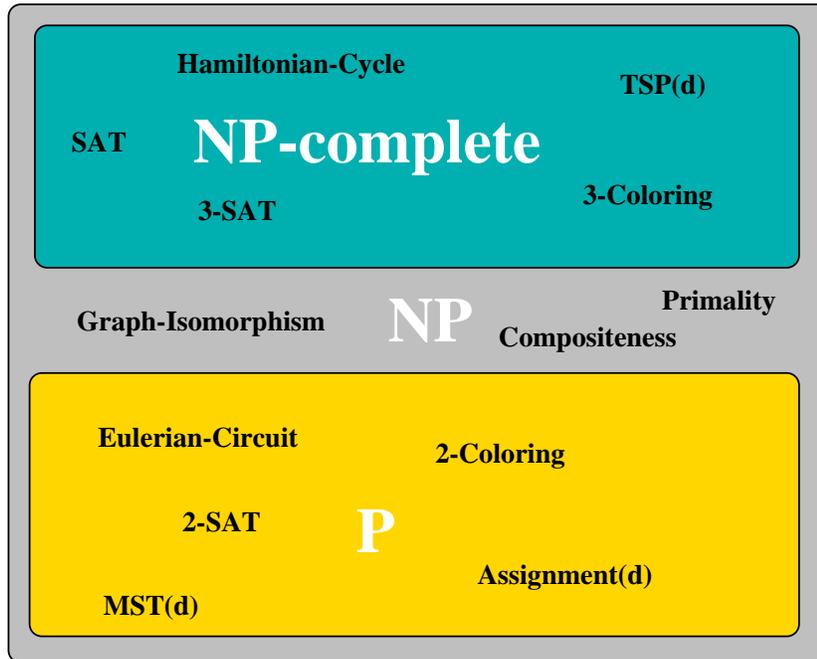


Figure 7: The map of complexity revisited.

2.3.3 The map of \mathcal{NP}

Since Cook's Theorem many problems have been shown to be \mathcal{NP} -complete. A comprehensive, up-to-date list of hundreds of \mathcal{NP} -complete problems can be found in the web [23]. Our map of \mathcal{NP} needs some redesign (Fig. 7). It turns out that all the intractable problems we have discussed so far are \mathcal{NP} -complete – except COMPOSITENESS and PRIMALITY. For both problems neither a polynomial algorithm is known nor a polynomial reduction that classifies them \mathcal{NP} -complete. Another \mathcal{NP} problem which resists classification in either \mathcal{P} or \mathcal{NP} is this:

GRAPH ISOMORPHISM: Given two graphs $G = (V, E)$ and $G'(V, E)$ on the same set of nodes. Are G and G' isomorphic, i.e. is there a permutation π of V such that $G' = \pi(G)$, where by $\pi(G)$ we denote the graph $(V, \{[\pi(u), \pi(v)] : [u, v] \in E\})$?

There are more problems in \mathcal{NP} that resists a classification in \mathcal{P} or \mathcal{NP} , but none of these problems has been *proven* not to belong to \mathcal{P} or \mathcal{NP} . What has been proven is

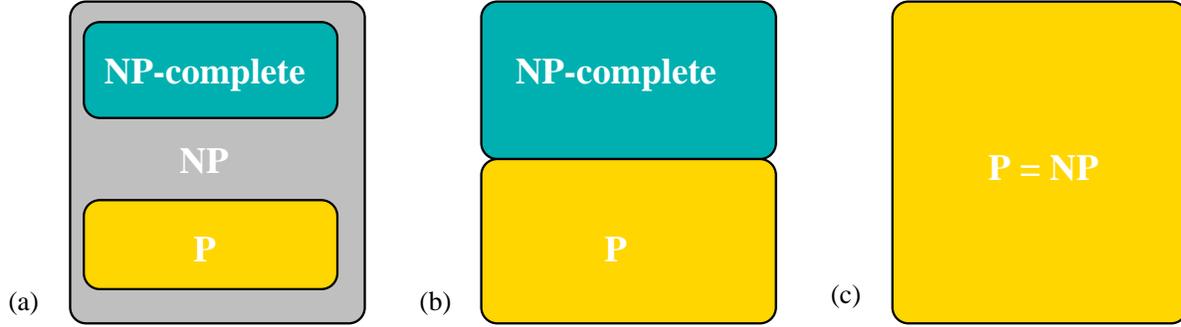


Figure 8: Three tentative maps of \mathcal{NP} . (b) can be ruled out. It is very likely (but not sure) that (a) is the correct map.

Theorem: If $\mathcal{P} \neq \mathcal{NP}$, then there exist \mathcal{NP} problems which are neither in \mathcal{P} nor are they \mathcal{NP} -complete.

This Theorem [24] rules out one of three tentative maps of \mathcal{NP} (Fig. 8).

2.4 Beyond \mathcal{NP}

2.4.1 Optimization problems

How does the classification of decision problems relates to optimization problems? The general instance of an optimization problem is a pair (F, c) , where F is the set of feasible solutions and c is a cost function $c : F \rightarrow \mathbb{R}$. We will consider only *combinatorial optimization* where the set F is countable. A combinatorial optimization problem P comes in three different flavors:

1. The optimization problem $P(O)$: Find the *feasible solution* $f^* \in F$ that minimizes the cost function.
2. The evaluation problem $P(E)$: Find the *cost* $c^* = c(f^*)$ of the minimum solution.
3. The decision problem $P(D)$: Given a bound $B \in \mathbb{R}$, is there a feasible solution $f \in F$ such that $c(f) \leq B$?

Under the assumption that the cost function c can be evaluated in polynomial time, it is straightforward to write down polynomial reductions that establish

$$P(D) \leq P(E) \leq P(O). \quad (11)$$

If the decision variant of an optimization problem is \mathcal{NP} -complete, there is no efficient algorithm for the optimization problem at all—unless $\mathcal{P} = \mathcal{NP}$. An optimization problem like TSP, whose decision variant is \mathcal{NP} -complete is denoted \mathcal{NP} -hard.

Does a polynomial algorithm for a decision problem imply a polynomial algorithm for the optimization or evaluation variant? For that we need to prove the reversal of Eq. 11,

$$P(O) \leq P(E) \leq P(D). \quad (12)$$

$P(E) \leq P(D)$ can be shown to hold if the cost of the optimum solution is an integer with logarithm bounded by a polynomial in the size of the input. The corresponding polynomial reduction evaluates the optimal cost c^* by asking the question “Is $c^* \leq B$?” for a sequence of values B that approaches c^* , similar to the bisection method to find the zeroes of a function.

There is no general method to prove $P(O) \leq P(E)$, but a strategy that often works can be demonstrated for the TSP: Let c^* be the known solution of TSP(E). Replace an arbitrary entry d_{ij} of the distance matrix with a value $c > c^*$ and solve TSP(E) with this modified distance matrix. If the length of the optimum tour is not affected by this modification, the link ij does not belong to the optimal tour. Repeating this procedure for different links one can reconstruct the optimum tour with a polynomial number of calls to a TSP(E)-solver, hence $\text{TSP}(O) \leq \text{TSP}(E)$.

2.4.2 Counting problems

So far we have studied two related styles of problems: Decision problems ask whether a desired solution exists, optimization problems require that a solution be produced. A third important and fundamentally different kind of problem asks, how many solutions exist. The *counting* variant of SAT reads

#SAT: Given a Boolean expression, compute the number of different truth assignments that satisfy it

Similarly, #HAMILTONIAN CYCLE asks for the number of different Hamiltonian cycles in a given graph, #TSP for the number of different tours with length $\leq B$ and so on.

Definition: A counting problem $\#P$ is a pair (F, d) , where F is the set of all feasible solutions and d is a decision function $d : F \mapsto \{0, 1\}$. The output of $\#P$ is the number of $f \in F$ with $d(f) = 1$. The class $\#\mathcal{P}$ (pronounced “number P”) consists of all counting problems associated with a decision function d that can be evaluated in polynomial time.

Like the class \mathcal{NP} , $\#\mathcal{P}$ collects all problems whose sole source of intractability is the number of feasible solutions. A polynomial algorithm for a counting problem $\#P$ implies a polynomial algorithm for the associated decision problem P : $P \leq \#P$. Hence it is very unlikely that #SAT can be solved efficiently. In fact one can define polynomial reducibility for counting problems and prove that all problems in $\#\mathcal{P}$ reduce polynomially to #SAT [6]:

Theorem: #SAT is $\#\mathcal{P}$ -complete.

As you might have guessed, #HAMILTONIAN CYCLE and #TSP are $\#\mathcal{P}$ -complete, too. Despite the similarity between \mathcal{NP} and $\#\mathcal{P}$, counting problems are inherently harder than decision problems. This is documented by those $\#\mathcal{P}$ -complete problems, for which the corresponding decision problem can be solved in polynomial time, the classical example being the problem of calculating the permanent of a matrix [25].

3 Computational complexity and physics

3.1 Quantum parallelism

In a seminal paper [26], Richard Feynman pointed out that a system of n quantum particles is exponentially hard to simulate on a classical computer. The idea of quantum computing is to reverse this situation and simulate a classically hard (i.e. exponential!) problem in polynomial time on a computer made of quantum devices.

A quantum computer processes qubits, quantum two-state systems $|0\rangle, |1\rangle$. A key feature of a quantum computer is that its registers can hold and process linear superpositions of all 2^n product states of n qubits like

$$\frac{1}{\sqrt{2^n}} \sum_{i_1, i_2, \dots, i_n=0}^1 |i_1 i_2 \dots i_n\rangle. \quad (13)$$

Using this feature it is not very difficult to construct a quantum computer capable of computing any function $f(x_1, \dots, x_n)$ of n Boolean variables simultaneously for all 2^n possible input values — in theory at least. This *quantum parallelism* resembles very much a nondeterministic algorithm with its **goto both** instruction and its exponentially branching execution tree. Is quantum parallelism the key to exponential computing power? The problem is how to extract the exponential information out of a quantum computer. When we defined nondeterministic solubility we did not care about how to “spot” the single ‘yes’-answer among the 2^n ‘no’-answers. This works fine for a theoretical concept, but for a practical computer reading the output matters a lot.

In order to gain advantage of exponential parallelism, it needs to be combined with another quantum feature known as interference. The goal is to arrange the computation such that constructive interference amplifies the result we are after and destructive interference cancels the rest. Due to the importance of interference phenomena it is not surprising that calculating the Fourier transform was the first problem that undergoes an exponential speedup: from $\mathcal{O}(n \log n)$ on a classical to $\mathcal{O}(\log^2 n)$ on a quantum computer. This speedup was the seed for the most important quantum algorithm known today: Shor’s algorithm to factor an integer in polynomial time [27].

Although Shor’s algorithm has some consequences for public key cryptography, it does not shatter the world of \mathcal{NP} : remember that COMPOSITENESS is in \mathcal{NP} , but not \mathcal{NP} -complete. Hence the holy grail of quantum computing, a polynomial time quantum algorithm for an \mathcal{NP} -complete problem, is yet to be discovered!

See [28] or www.qubit.org to learn more.

3.2 Analytical solubility of Ising models

Some problems in statistical physics have been exactly solved, but the majority of problems has withstand the efforts of generations of mathematicians and physicists. Why are some problems analytically solvable whereas others, often similar, are not? Relating this question to the algorithmic complexity of evaluating the partition function gives us no final answer

but helps to clarify the borderline that separates the analytically tractable from the intractable problems.

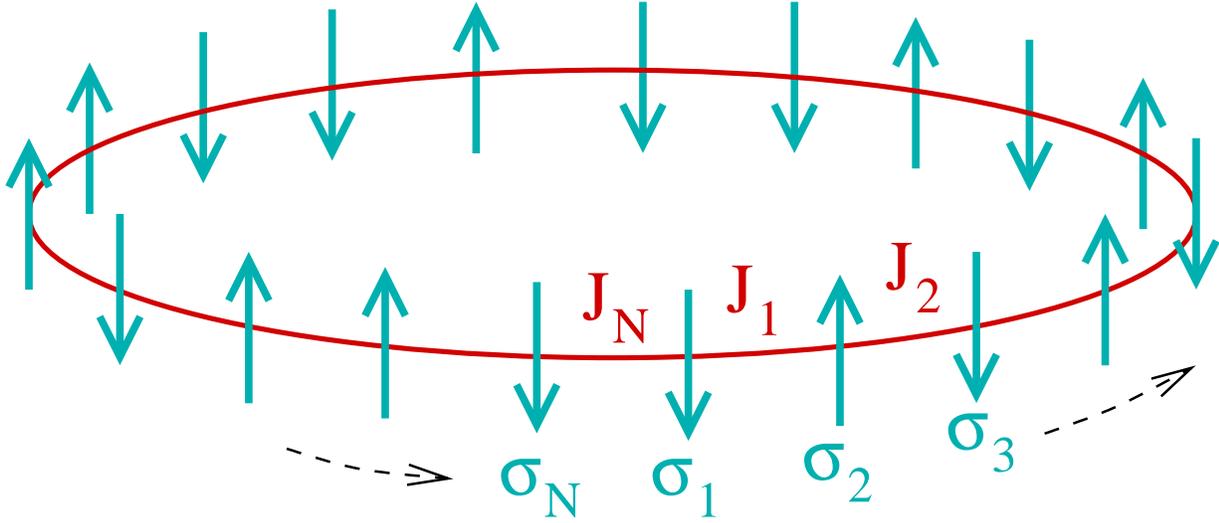


Figure 9: Onedimensional Ising spin glass with periodic boundary conditions. The partition sum of this system can be calculated in polynomial time.

As an example consider the Ising spin glass on a general graph G [29]. Let $\sigma = (\sigma_1, \dots, \sigma_N)$ be an assignment of Ising spins $\sigma_i = \pm 1$ (“up” or “down”). The energy of a configuration σ is

$$E(\sigma) = - \sum_{\langle i,j \rangle} J_{i,j} \sigma_i \sigma_j - H \sum_i \sigma_i, \quad (14)$$

where H is the external magnetic field, $J_{i,j}$ are the coupling constants and the first summation is over all edges in G . The fundamental problem in statistical mechanics is to determine the partition function

$$Z_N(G) = \sum_{\sigma} e^{-\beta E(\sigma)}. \quad (15)$$

Evaluating the sum directly requires $\mathcal{O}(2^N)$ operations. The notion “analytic solution” has no precise definition, but as a minimum requirement we want this number to be reduced from being exponential to being polynomial in N . As an example consider the well known transfer matrix solution of the one dimensional Ising glass with periodic boundary conditions and coupling J_k between spins σ_k and σ_{k+1} (Fig. 9),

$$Z_N(\text{ring}) = \text{Tr} \prod_{k=1}^N \begin{pmatrix} e^{\beta(J_k+H)} & e^{-\beta J_k} \\ e^{-\beta J_k} & e^{\beta(J_k-H)} \end{pmatrix}, \quad (16)$$

which can be evaluated in $\mathcal{O}(N)$ elementary operations. Since any exact solution must include all numbers J_k , this is the best we can expect. In the homogeneous case $J_k \equiv J$, where we

can calculate the product of transfer matrices,

$$Z_N = \lambda_+^N + \lambda_-^N \text{ with } \lambda_{\pm} = e^{\beta J} \left[\cosh(\beta H) \pm \sqrt{\cosh^2(\beta H) - 2e^{-2\beta H} \sinh(2\beta H)} \right] \quad (17)$$

the evaluation complexity drops to $\mathcal{O}(\log N)$ using fast exponentiation.

Writing the partition sum as

$$Z(G) = \sum_{E_k} n(E_k) e^{-\beta E_k} \quad (18)$$

where the sum is over all possible energy values, it becomes obvious that calculating Z is closely related to the $\#\mathcal{P}$ problem of determining $n(E_k)$. For general finite graphs G , this problem has proven to be $\#\mathcal{P}$ -complete [30, 31], so there is no hope to find an analytical solution (even in the weak sense above). This situation hardly improves if we restrict the graphs to the more “physical” crystal lattices: computing the partition function for a finite sublattice of a *non-planar crystal lattice* is $\#\mathcal{P}$ -complete [32]. This includes every crystal lattice in $d > 2$, the $d=2$ model with next-nearest neighbor interactions, two coupled $d = 2$ models etc.. It also includes all models with $d \geq 2$ and external field H , since these can be transformed into zero-field models on an augmented graph \hat{G} which is non-planar unless the underlying lattice graph G is one-dimensional. The construction of \hat{G} from G is easy, just adjoin an additional vertex (spin) σ_0 to G and let the additional edges $\sigma_0\sigma_i$ have constant interaction energy $J_{0,i} = H$. The partition function of the augmented system reads

$$Z(\hat{G}) = \sum_{\sigma} e^{-\beta[-\sum J_{i,j}\sigma_i\sigma_j]} \left(e^{-\beta H \sum \sigma_i} + e^{\beta H \sum \sigma_i} \right), \quad (19)$$

where the additional factor comes from the new spin $\sigma_0 = \pm 1$. From this expression it is easy to see that $Z(\hat{G})$ equals two times the partition function $Z(G)$ in field H .

But where are the soluble Ising models? It has been proven that the partition function of Ising systems on *planar crystal lattices* can be evaluated in polynomial time [33]. This includes the celebrated Onsager solution of the square ferromagnet [34] as well as the one-dimensional example from above. It turns out that the Ising model’s partition sum can be calculated in polynomial time for *all graphs of fixed genus g* [35, 36]. A graph has genus g if it can be drawn on an orientable surface of genus g (a sphere with g “handles” attached to it) without crossing the edges. Planar graphs have genus 0, toroidal graphs have genus 1 and so on. For the crystal lattices in $d > 2$, the genus increases monotonically with the number of spins, i.e. it is not fixed [35].

The mechanism for proving tractability or intractability is the same in statistical mechanics as in computational complexity: polynomial reduction. Barahona [33] for example applies a reduction of the \mathcal{NP} -hard problem MAX CUT to the Ising spin glass in $3d$ to prove the \mathcal{NP} -hardness of the latter. A reduction of the planar Ising model to MINIMUM WEIGHT MATCHING on the other hand proves the tractability of the former, since MINIMUM WEIGHT MATCHING can be solved in polynomial time.

Note that in our classification of spin systems the nature of the couplings is not significant. A frustrated, glassy system with random couplings $J_{i,j}$ of both signs is in the same class as

the homogeneous ferromagnet with $J_{i,j} \equiv J > 0$ as long as the underlying graph G is the same. In our onedimensional example we did not discriminate these cases: they are both polynomial. This situation changes of course if we consider the ground states rather than the complete partition function. Here the nature of the couplings matters a lot: finding the ground states in pure ferromagnetic systems is trivial on all lattices, whereas it is \mathcal{NP} -hard for glassy systems with positive and negative couplings on non-planar lattices [33].

A lot of other problems arising in statistical physics can be classified according to the computational complexity to evaluate their partition function, see [37] and references therein. It turns out that all the problems known to have an exact solution [38] can be evaluated in polynomial time. Problems that are $\#\mathcal{P}$ -complete, however, are very unlikely to be exactly solvable. Anyone looking for an exact solution of such a problem should keep in mind, that he or she is simultaneously attacking TSP, HAMILTONIAN CYCLE, SAT and all the other NP-hard problems. In statistical mechanics the focus is on results for the thermodynamic limit $N \rightarrow \infty$ rather than for finite systems, however. It is not clear how much of the “hardness” survives in this limit.

3.3 Probabilistic analysis of combinatorial problems

Problems from combinatorial optimization can formally considered as models in statistical mechanics. The cost function is renamed Hamiltonian, random instances are samples of quenched disorder, and the ground states of the formal model correspond to the solutions of the optimization problems. In this way methods developed in the framework of statistical mechanics of disordered systems become powerful tools for the probabilistic analysis of combinatorial problems [39].

Statistical mechanics methods have been applied for example to the TSP [40,41], GRAPH PARTITIONING [42] and k -SAT [43, 44]. A particular nice example of this approach is the analysis of ASSIGNMENT (also called BIPARTITE MATCHING): Given an $N \times N$ matrix with non-negative entries $a_{i,j} \geq 0$. Find

$$E_N^* = \min_{\sigma} \sum_{i=1}^N a_{i,\sigma(i)} \quad (20)$$

where the minimum is taken over all permutations σ of $(1, \dots, N)$.

A probabilistic analysis aims at calculating average properties for an ensemble of random instances, the canonical ensemble being random numbers $a_{i,j}$ drawn independently from a common probability density $\rho(a)$. Using the replica method from statistical physics, Marc Mézard and Giorgio Parisi [45] found (among other things)

$$\lim_{N \rightarrow \infty} \langle E_N^* \rangle = \frac{\pi^2}{6}, \quad (21)$$

where $\langle \cdot \rangle$ denotes averaging over the $a_{i,j}$. A rigorous proof of eq. 21 has been presented recently [46] and represents one of the rare cases where a replica result has been confirmed by rigorous methods.

Some years after the replica-solution, Parisi recognized that for exponentially distributed matrix elements ($\rho(a) = e^{-a}$) the average optimum for $N = 1$ and $N = 2$ is

$$\langle E_1^* \rangle = 1 \quad \langle E_2^* \rangle = 1 + \frac{1}{2^2}. \quad (22)$$

From this and the fact that the replica result for $N \rightarrow \infty$ can be written as

$$\frac{\pi^2}{6} = \sum_{k=1}^{\infty} \frac{1}{k^2} \quad (23)$$

he conjectured [47] that the average optimum for finite systems is

$$\langle E_N^* \rangle = \sum_{k=1}^N \frac{1}{k^2}. \quad (24)$$

Parisi's conjecture is supported by numerical simulations, but no formal proof has been found despite some efforts [48, 49].

Note that Eqs. 22 and 24 only hold for $\rho(a) = e^{-a}$, whereas eq. 21 is valid for all densities with $\rho(a \rightarrow 0) = 1$. For the uniform density on $[0, 1]$ the first terms are

$$\langle E_1^* \rangle = \frac{1}{2} \quad \langle E_2^* \rangle = \frac{23}{30}. \quad (25)$$

If you can you guess the expression for general, finite N in this case, please send me an email.

Sometimes a statistical mechanics analysis not only yields exact analytical results but also reveals features that are important to design and understand algorithms. A recent example is the analysis of the Davis-Putnam-Algorithm for SAT [50, 51]. Another example is given by the number partitioning problem NPP [52], an \mathcal{NP} -hard optimization problem, where it has been shown, that for this particular problem no heuristic approach can be better than stupid random search [53–55].

3.4 Phase transitions in computational complexity

The theory of computational complexity is based entirely on worst-case analysis. It may happen that an algorithm requires exponential time on pathological instances only. A famous example is the Simplex Method for LINEAR PROGRAMMING. Despite its exponential worst-case complexity it is used in many applications to solve really large problems. Apparently the instances that trigger exponential running time do not appear under practical conditions.

Now LINEAR PROGRAMMING is in \mathcal{P} thanks to the Ellipsoid algorithm [4], but similar scenarios are observed for \mathcal{NP} -hard problems. Take 3-SAT as an example. Generating random instances with N variables and M clauses and feeding them to a smart but exhaustive algorithm one observes polynomial running time *unless* the ratio M/N is carefully adjusted. If M/N is too low, the problem is *underconstrained*, i.e. has many satisfying solutions, and a clever algorithm will find one of these quickly. If M/N is too large, the problem is *overconstrained*, i.e. has a large number of contradictory constraints which again a clever algorithm

will discover quickly [56]. The real hard instances are generated for ratios $\alpha = M/N$ close to a critical value α_c [57].

The transition from underconstrained to overconstrained formulas in 3-SAT bears the hallmarks of a phase transition in physical systems. The control parameter is α , the order parameter is the probability of the formula being satisfiable. Similar phase transitions have been discovered and analyzed with methods from statistical mechanics in other computational problems. See the special issue of *Theoretical Computer Science* on “Phase Transitions in Combinatorial Problems” for a recent overview of this interdisciplinary field [58].

References

- [1] Harry R. Lewis and Christos H. Papadimitriou. The efficiency of algorithms. *Scientific American*, pages 96–109, 1 1978.
- [2] Alexander K. Hartmann. Introduction to complexity theory. www.theorie.physik.uni-goettingen.de/~hartmann/seminar/komplex.ps.gz, 1999.
- [3] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. The MIT Electrical Engineering and Computer Science Series. The MIT Press, Cambridge, Massachusetts, 1990.
- [4] C.H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization*. Prentice-Hall, Englewood Cliffs, New Jersey, 1982.
- [5] Michael R. Garey and David S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W.H. Freeman, New York, 1977.
- [6] Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1994.
- [7] Roger Penrose. *The Emperor’s new Mind*. Oxford University Press, 1989.
- [8] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Computation*, 9(3):251–280, 1990.
- [9] Béla Bollobás. *Modern Graph Theory*, volume 184 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1998.
- [10] H.N. Gabow, Z. Galil, T.H. Spencer, and R.E. Tarjan. Efficient algorithms for finding minimum spanning trees in undirected and directed graphs. *Combinatorica*, 6:109–122, 1986.
- [11] E.L. Lawler, J.K. Lenstra, A.H.G. Rinnoy Kan, and D.B. Shmoys, editors. *The Traveling Salesman Problem*. Wiley-Interscience series in discrete mathematics and optimization. John Wiley & Sons, New York, 1985.

- [12] Gerhard Reinelt. *The Travelling Salesman. Computational Solutions for TSP Applications*, volume 840 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin Heidelberg New York, 1994.
- [13] Martin Grötschel and Manfred W. Padberg. Die optimierte Odyssee. *Spektrum der Wissenschaft*, 4:76–85, 1999. See also <http://www.spektrum.de/odyssee.html>.
- [14] David Applegate, Robert Bixby, Vasek Chvatal, and William Cook. The tsp homepage. <http://www.keck.caam.rice.edu/tsp/>.
- [15] L. Euler. Solutio problematis ad geometriam situs pertinentis. *Comm. Acad. Sci. Imper. Petropol.*, 8:128–140, 1736.
- [16] Vipin Kumar. Algorithms for Constraint Satisfaction Problems: A Survey. *AI Magazine*, 13(1):32–44, 1992. available from <ftp://ftp.cs.umn.edu/dept/users/kumar/csp-aimagazine.ps>.
- [17] Bengt Aspvall, Michael F. Plass, and Robert Endre Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121–123, 1979.
- [18] D.S. Johnson and C.H. Papadimitriou. Computational complexity. In Lawler et al. [11], pages 37–85.
- [19] V.R. Pratt. Every prime has a succinct certificate. *SIAM J. Comput.*, 4:214–220, 1975.
- [20] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Prottasi. *Complexity and Approximation*. Springer-Verlag, Berlin Heidelberg New York, 1999.
- [21] R.M. Karp. Complexity of computer computations. In R.E. Miller and J.W. Thatcher, editors, *Reducibility Among Combinatorial Problems*, pages 85–103, New York, 1972. Plenum Press.
- [22] Stephen Cook. The complexity of theorem proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, pages 151–158, 1971.
- [23] Pierluigi Crescenzi and Viggo Kann. A compendium of NP optimization problems. <http://www.nada.kth.se/~viggo/wwwcompendium>.
- [24] R.E. Ladner. On the structure of polynomial time reducibility. *Journal of the ACM*, 22:155–171, 1975.
- [25] L.G. Valiant. The complexity of computing the permanent. *Theor. Comp. Sci.*, 8:189–201, 1979.
- [26] R. Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21(6/7):467–488, 1982.

- [27] P.W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comp.*, 26(5):1484–1509, 1997.
- [28] Dorit Aharonov. Quantum computation. In Dietrich Stauffer, editor, *Annual Reviews of Computational Physics VI*. World Scientific, 1998. arXiv:quant-ph/9812037.
- [29] Brian Hayes. Computing sciene: The world in a spin. *American Scientist*, 88(5):384–388, September-October 2000. <http://www.amsci.org/amsci/issues/-comsci00/compsci2000-09.html>.
- [30] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. In *Proc. 17th ICALP*. EATCS, 1990.
- [31] F. Jaeger, D.L. Vertigan, and D.J.A. Welsh. On the computational complexity of the Jones and Tutte polynomials. *Math. Proc. Camb. Phil. Soc.*, 108:35–53, 1990.
- [32] Sorin Istrail. Statistical mechanics, three-dimensionality and NP-completeness. In *Proc. 31st ACM Annual Symposium on the Theory of Computing (STOC 2000)*, Portland, Oregon, May 2000. ACM Press.
- [33] F. Barahona. On the computational complexity of Ising spin glass models. *J. Phys. A*, 15:3241–3253, 1982.
- [34] L. Onsager. Crystal statistics I. A two-dimensional model with an order-disorder transition. *Phys. Rev.*, 65, 1944.
- [35] Tullio Regge and Riccardo Zecchina. Combinatorial and topological approach to the 3D Ising model. *J. Phys. A*, 33:741–761, 2000.
- [36] Anna Galluccio, Martin Loeb1, and Jan Vondrák. New algorithm for the Ising problem: Partition function for finite lattice graphs. *Phys. Rev. Lett.*, 84(26):5924–5927, 2000.
- [37] D.J.A. Welsh. The computational complexity of some classical problems from statistical physics. In G.R. Grimmett and D.J.A. Welsh, editors, *Disorder in Physical Systems*, pages 307–321. Clarendon Press, Oxford, 1990.
- [38] R.J. Baxter. *Exactly Solved Models in Statistical Mechanics*. Academic Press, San Diego, 1982.
- [39] M. Mézard, G. Parisi, and M.A. Virasoro. *Spin glass theory and beyond*. World Scientific, Singapore, 1987.
- [40] M. Mézard and G. Parisi. Mean-field equations for the matching and the travelling salesman problems. *Europhys. Lett.*, 2(12):913–918, December 1986.
- [41] W. Krauth and M. Mézard. The cavity method and the travelling-salesman problem. *Europhys. Lett.*, 8(3):213–218, 1989.

- [42] Y. Fu and P. W. Anderson. Application of statistical mechanics to NP-complete problems in combinatorial optimization. *J. Phys. A*, 19:1605–1620, 1986.
- [43] Scott Kirkpatrick and Bart Selman. Critical behavior in the satisfiability of random boolean expressions. *Science*, 264:1297–1301, May 1994.
- [44] Remi Monasson and Riccardo Zecchina. Statistical mechanics of the random k -satisfiability model. *Phys. Rev. E*, 56(2):1357–1370, 8 1997.
- [45] M. Mézard and G. Parisi. Replicas and optimization. *J. Physique Lett.*, 46:L771–L778, 1985.
- [46] David J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Rand. Struct. Alg.*, 18:381–418, 2001.
- [47] Giorgio Parisi. A conjecture on random bipartite matching. arXiv:cond-mat/9801176, 1998.
- [48] V.S. Dotsenko. Exact solution of the random bipartite matching model. *J. Phys. A*, 33:2015–2030, 2000.
- [49] D. Coppersmith and G. Sorkin. Constructive bounds and exact expectations for the random assignment problem. *Rand. Struct. Alg.*, 15:113–144, 1999.
- [50] Simona Cocco and Remi Monasson. Trajectories in phase diagrams, growth processes and computational complexity: how search algorithms solve the 3-satisfiability problem. *Phys. Rev. Lett.*, 86:1654–1657, 2001.
- [51] Simona Cocco and Remi Monasson. Statistical physics analysis of the computational complexity of solving random satisfiability problems using backtrack algorithms. *European Physics Journal B*, 22:505–531, 2001. arXiv:cond-mat/0012191.
- [52] Brian Hayes. Computing science: The easiest hard problem. *American Scientist*, 90(2):113–117, March-April 2002. <http://www.amsci.org/amsci/issues/comsci02/-compsci2002-03.html>.
- [53] Stephan Mertens. Random costs in combinatorial optimization. *Phys. Rev. Lett.*, 84(7):1347–1350, February 2000.
- [54] Stephan Mertens. A physicist’s approach to number partitioning. *Theor. Comp. Sci.*, 265(1-2):79–108, 2001.
- [55] Christian Borgs, Jennifer Chayes, and Boris Pittel. Phase transition and finite-size scaling for the integer partitioning problem. *Rand. Struct. Alg.*, 19(3–4):247–288, 2001.
- [56] Brian Hayes. Computing science: Can’t get no satisfaction. *American Scientist*, 85(2):108–112, March-April 1997. <http://www.amsci.org/amsci/issues/Comsci97/-compsci9703.html>.

- [57] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400:133–137, July 1999.
- [58] Olivier Dubois, Remi Monasson, Bart Selman, and Riccardo Zecchina, editors. *Phase Transitions in Combinatorial Problems*, volume 265 of *Theor. Comp. Sci.*, 2001.