CAMBRIDGE

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

1

Introduction

1.1 Physics and computational physics

Solving a physical problem often amounts to solving an ordinary or partial differential equation. This is the case in classical mechanics, electrodynamics, quantum mechanics, fluid dynamics and so on. In statistical physics we must calculate sums or integrals over large numbers of degrees of freedom. Whatever type of problem we attack, it is very rare that analytical solutions are possible. In most cases we therefore resort to numerical calculations to obtain useful results. Computer performance has increased dramatically over the last few decades (see also Chapter 16) and we can solve complicated equations and evaluate large integrals in a reasonable amount of time.

Often we can apply numerical routines (found in software libraries for example) directly to the physical equations and obtain a solution. We shall see, however, that although computers have become very powerful, they are still unable to provide a solution to most problems without approximations to the physical equations. In this book, we shall focus on these approximations: that is, we shall concentrate on the development of computational methods (and also on their implementation into computer programs). In this introductory chapter we give a bird's-eye perspective of different fields of physics and the computational methods used to solve problems in these areas. We give examples of direct application of numerical methods but we also give brief and heuristic descriptions of the additional theoretical analysis and approximations necessary to obtain workable methods for more complicated problems which are described in more detail in the remainder of this book. The order adopted in the following sections differs somewhat from the order in which the material is treated in this book.

1.2 Classical mechanics and statistical mechanics

The motion of a point particle in one dimension subject to a force F depending on the particle's position x, and perhaps on the velocity $\dot{x}$ and on time t, is determined

© Cambridge University Press www.cambridge.org

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

by Newton's equation of motion:

$$m\ddot{x}(t) = F[x(t), \dot{x}(t), t]. \tag{1.1}$$

The (double) dot denotes a (double) derivative with respect to time. A solution can be found for each set of initial conditions $x(t_0)$ and $\dot{x}(t_0)$ given at some time t_0. Analytical solutions exist for constant force, for the harmonic oscillator ($F = \kappa x^2/2$), and for a number of other cases. In Appendix A7.1 a simple numerical method for solving this equation is described and this can be applied straightforwardly to arbitrary forces and initial conditions.

Interesting and sometimes surprising physical phenomena can now be studied. As an example, consider the Duffing oscillator [1], with a force given by

$$F[x, \dot{x}, t] = -\gamma\dot{x} + 2ax - 4bx^3 + F_0\cos(\omega t). \tag{1.2}$$

The first term on the right hand side represents a velocity-dependent friction; the second and third terms are the force a particle feels when it moves in a double potential well $bx^4 - ax^2$, and the last term is an external periodic force. An experimental realisation is a pendulum consisting of an iron ball suspended by a thin string, with two magnets below it. The pendulum and the magnets are placed on a table which is moved back and forth with frequency ω. The string and the air provide the frictional force, the two magnets together with gravity form some kind of double potential well, and, in the reference frame in which the pendulum is at rest, the periodic motion of the table is felt as a periodic force. It turns out that the Duffing oscillator exhibits *chaotic behaviour* for particular values of the parameters γ, a, b, F_0 and ω. This means that the motion itself looks irregular and that a very small change in the initial conditions will grow and result in a completely different motion. Figure 1.1 shows the behaviour of the Duffing oscillator for two nearly equal initial conditions, showing the sensitivity to these conditions. Over the past few decades, chaotic systems have been studied extensively. A system that often behaves chaotically is the weather: the difficulty in predicting the evolution of chaotic systems causes weather forecasts to be increasingly unreliable as they look further into the future, and occasionally to be dramatically wrong.

Another interesting problem is that of several particles, moving in three dimensions and subject to each other's gravitational interaction. Our Solar System is an example. For the simplest nontrivial case of three particles (for two particles, Newton has given the analytical solution), analytical solutions exist for particular configurations, but the general problem can only be solved numerically. This problem is called the *three-body problem* (N-body problem in general). The motion of satellites orbiting in space is calculated numerically using programs for the N-body problem, and the evolution of galaxies is calculated with similar programs using a large number of test particles (representing the stars). Millions of particles can

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

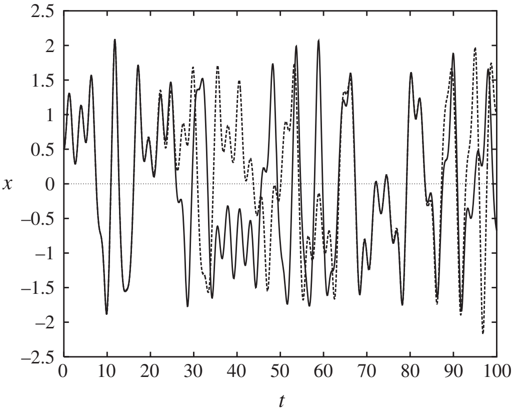

Figure 1.1. Solution of the Duffing oscillator. Parameters are $m = 1$, $a = 1/4$, $b = 1/2$, $F_0 = 2.0$, $\omega = 2.4$, $\gamma = 0.1$. Two solutions are shown: one with initial position $x_0 = 0.5$, the other with $x_0 = 0.5001$ ($\dot{x}_0 = 0$ in both cases). For these nearly equal initial conditions, the solutions soon become uncorrelated, showing the difficulty in predicting the time evolution of a chaotic system.

be treated using a combination of high-end computers and clever computational methods which will be considered in Chapter 8. Electrostatic forces are related to gravitational forces, as both the gravitational and the electrostatic (Coulomb) potential have a $1/r$ form. The difference between the two is that electrostatic forces can be repulsive or attractive, whereas gravitational forces are always attractive.

Neutral atoms interact via a different potential: they attract each other weakly through induced polarisation, unless they come too close – then the Pauli principle causes the electron clouds to repel each other. The problem of many interacting atoms and molecules is a very important subfield of computational physics: it is called *molecular dynamics*. In molecular dynamics, the equations of motion for the particles are solved straightforwardly using numerical algorithms similar to those with which a Duffing oscillator is analysed, the main difference being the larger number of degrees of freedom in molecular dynamics. The aim of molecular dynamics simulations is to predict the behaviour of gases, liquids and solids (and systems in other phases, like liquid crystals). An important result is the equation of state: this is the relation between temperature, number of particles, pressure and volume. Also, the microscopic structure as exhibited by the pair correlation function, which is experimentally accessible via neutron scattering, is an interesting property which can be determined in simulations. There are, however, many problems and pitfalls associated with computer simulations: the systems that can be simulated are always much smaller than realistic systems, and simulating a system at a predefined temperature or chemical potential is nontrivial. All these aspects will be considered in Chapter 8.

CAMBRIDGE

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

1.3 Stochastic simulations

In the previous section we have explained how numerical algorithms for solving Newton's equations of motion can be used to simulate liquids. The particles are moved around according to their mechanical trajectories which are governed by the forces they exert on each other. Another way of moving them around is to displace them in a random fashion. Of course this must be done in a controlled way, and not every move should be allowed, but we shall see in Chapter 10 that it is possible to obtain information in this way similar to that obtained from molecular dynamics. This is an example of a *Monte Carlo* method – procedures in which random numbers play an essential role. The Monte Carlo method is not suitable for studying dynamical physical quantities such as transport coefficients, as it uses artificial dynamics to simulate many-particle systems.

Random number generators can also be used in *direct simulations*: some process of which we do not know the details is replaced by a random generator. If you simulate a card game, for example, the cards are distributed among the players by using random numbers. An example of a direct simulation in physics is diffusion limited aggregation (DLA), which describes the growth of dendritic clusters (see Figure 1.2). Consider a square lattice in two dimensions. The sites of the lattice are either occupied or unoccupied. Initially, only one site in the centre is occupied. We release a random walker from the boundary of the lattice. The walker moves over the lattice in a stepwise fashion. At each step, the walker moves from a site to

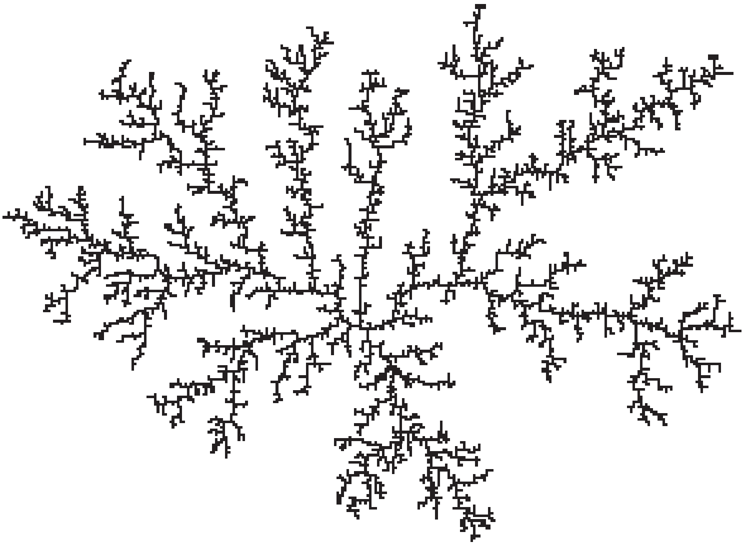

Figure 1.2. Dendritic cluster grown in a DLA simulation. The cluster consists of 9400 sites and it was grown on a 175 × 175 lattice.

© Cambridge University Press

CAMBRIDGE

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

one of its neighbour sites, which is chosen at random (there are four neighbours for each site in the interior of the lattice; the boundary sites have three neighbours, or two if they lie on a corner). If the walker arrives at a site neighbouring the occupied central site, it sticks there, so that a two-site cluster is formed. Then a new walker is released from the boundary. This walker also performs a random walk on the lattice until it arrives at a site neighbouring the cluster of two occupied sites, to form a three-site cluster, and so on. After a long time, a dendritic cluster is formed (see Figure 1.2), which shows a strong resemblance to actual dendrites formed in crystal growth, or by growing bacterial colonies [2], frost patterns on the window and numerous other physical phenomena.

This shows again that interesting physics can be studied by straightforward application of simple algorithms. In Chapter 10 we shall concentrate on the Monte Carlo method for studying many-particle systems at a predefined temperature, volume and particle number. This technique is less direct than DLA, and, just as in molecular dynamics, studying the system for different predefined parameters, such as chemical potential, and evaluating free energies are nontrivial aspects which need further theoretical consideration. The Monte Carlo method also enables us to analyse lattice spin models, which are important for studying magnetism and field theory (see below). These models cannot always be analysed using molecular dynamics methods, and Monte Carlo is often the only tool we have at our disposal in that case. There also exist alternative, more powerful techniques for simulating dendrite formation, but these are not treated in this book.

1.4 Electrodynamics and hydrodynamics

The equations of electrodynamics and hydrodynamics are partial differential equations. There exist numerical methods for solving these equations, but the problem is intrinsically demanding because the fields are continuous and an infinite number of variables is involved. The standard approach is to apply some sort of *discretisation* and consider the solution for the electric potential or for the flow field only on the points of the discrete grid, thus reducing the infinite number of variables to a finite number. Another method of solution consists of writing the field as a linear combination of smooth functions, such as plane waves, and solving for the best values of the expansion coefficients.

There exist several methods for solving partial differential equations: finite difference methods (FDM), finite element methods (FEM), Fourier transform methods and multigrid methods. These methods are also very often used in engineering problems, and are essentially the domain of numerical analysis. The finite element method is very versatile and therefore receives our particular attention in Chapter 13. The other methods can be found in Appendix A7.2.

CAMBRIDGE

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

1.5 Quantum mechanics

In quantum mechanics we regularly need to solve the Schrödinger equation for one or more particles. There is usually an external potential felt by the particles, and in addition there might be interactions between the particles. For a single particle moving in one dimension, the stationary form of the Schrödinger equation reduces to an ordinary differential equation, and techniques similar to those used in solving Newton's equations can be used. The main difference is that the stationary Schrödinger equation is an eigenvalue equation, and in the case of a discrete spectrum, the energy eigenvalue must be varied until the wave function is physically acceptable, which means that it matches some boundary conditions and is normalisable. Examples of this direct approach are discussed in Appendix A, in particular Problem A.4.

In two and more dimensions, or if we have more than one particle, or if we want to solve the time-dependent Schrödinger equation, we must solve a partial differential equation. Sometimes, the particular geometry of the problem and the boundary conditions allow us to reduce the complexity of the problem and transform it into ordinary differential equations. This will be done in Chapter 2, where we shall study particles scattering off a spherically symmetric potential.

Among the most important quantum problems in physics is the behaviour of electrons moving in the field generated by nuclei, which occurs in atoms, molecules and solids. This problem is treated quite extensively in this book, but the methods we develop for it are also applied in nuclear physics. Solving the Schrödinger equation for one electron moving in the potential generated by the atomic static nuclei is already a difficult problem, as it involves solving a partial differential equation. Moreover, the potential is strong close to the nuclei and weak elsewhere, so the typical length scale of the wave function varies strongly through space. Therefore, discretisation methods must use grids which are finer close to the nuclei, rendering such methods difficult. The method of choice is, in fact, to expand the wave function as a linear combination of fixed basis functions that vary strongly close to the nuclei and are smooth elsewhere, and find the optimal values for the expansion coefficients. This is an example of the *variational method*, which will be discussed in Chapter 3. This application of the variational method leads to a matrix eigenvalue problem which can be solved very efficiently on a computer.

An extra complication arises when there are many (say N) electrons, interacting via the Coulomb potential, so that we must solve a partial differential equation in $3N$ dimensions. In addition to this we must realise that electrons are fermions and the many-electron wave function must therefore be antisymmetric with respect to exchange of any pair of electrons. Because of the large number of dimensions, solving the Schrödinger equation is not feasible using any of the

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

standard numerical methods for solving partial differential equations, so we must make approximations. One approach is the Hartree–Fock (HF) method,developed in the early days of quantum mechanics, which takes into account the antisymmetry of the many-electron wave function. This leads to an independent particle picture, in which each electron moves in the potential generated by the nuclei plus an average potential generated by the other electrons. The latter depends on the electronic wave functions, and hence the problem must be solved *self-consistently* – in Chapter 4 we shall see how this is done. The HF method leads to wave functions that are fully antisymmetric, but contributions arising from the Coulomb interaction between the particles are taken into account in an approximate way, analogous to the way correlations are treated in the mean field approach in statistical mechanics.

Another approach to the quantum many-electron problem is given by *density functional theory* (DFT), which will be discussed in Chapter 5. This theory, which is in principle exact, can in practice only be used in conjunction with approximate schemes to be discussed in Chapter 5, the most important of which is the *local density approximation* (LDA). This also leads to an independent-particle Schrödinger equation, but in this case, the correlation effects resulting from the antisymmetry of the wave function are not incorporated exactly, leading to a small, unphysical interaction of an electron with itself (*self-interaction*). However, in contrast to Hartree–Fock,the approach does account (in an approximate way) for the dynamic correlation effects due to the electrons moving out of each other's way as a result of the Coulomb repulsion between them.

All these approaches lead in the end to a matrix eigenvalue problem, whose size depends on the number of electrons present in the system. The resulting solutions enable us to calculate total energies and excitation spectra which can be compared with experimental results.

1.6 Relations between quantum mechanics and classical statistical physics

In the previous two sections we have seen that problems in classical statistical mechanics can be studied with Monte Carlo techniques, using random numbers, and that the solution of quantum mechanical problems reduces to solving matrix eigenvalue problems. It turns out that quantum mechanics and classical statistical mechanics are related in their mathematical structure. Consider for example the partition function for a classical mechanics system at temperature T, with degrees of freedom denoted by the variable X and described by an energy function (that is, a classical Hamiltonian) $\mathcal{H}$:

$$Z_{\mathrm{Cl}} = \sum_X \mathrm{e}^{-\mathcal{H}(X)/(k_\mathrm{B}T)}, \tag{1.3}$$

© Cambridge University Press www.cambridge.org

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

and that of a quantum system with quantum Hamiltonian H:

$$Z_{\mathrm{QM}} = \mathrm{Tr}(\mathrm{e}^{-H/(k_{\mathrm{B}}T)}); \tag{1.4}$$

'Tr' denotes the trace of the operator following it. We will show in Chapter 12 that in the path-integral formalism, the second expression can be transformed into the same form as the first one. Also, there is a strong similarity between the exponent occurring in the quantum partition function and the quantum time-evolution operator $U(t) = \exp(-\mathrm{i}tH/\hbar)$, so solving the time evolution of a quantum system is equivalent to evaluating a classical or quantum partition function, the difference being an imaginary factor $\mathrm{i}t/\hbar$ replacing the real factor $1/(k_{\mathrm{B}}T)$, and taking the trace in the case of the quantum partition function rather than a sum over states in the classical analogue.

These mathematical analogies suggest that numerical methods for either classical statistical mechanics or quantum mechanics are applicable in both fields. Indeed, in Chapter 11, we shall see that it is possible to analyse classical statistical spin problems on lattices by diagonalising large matrices. In Chapter 12, on the other hand, we shall use Monte Carlo methods for solving quantum problems. These methods enable us to treat the quantum many-particle problem without systematic approximations, because, as will be shown in Chapter 12, Monte Carlo techniques are very efficient for calculating integrals in many dimensions. This, as we have seen above, was precisely the problem arising in the solution of interacting many-particle systems.

1.7 Quantum molecular dynamics

Systems of many interacting atoms or molecules can be studied classically by solving Newton's equations of motion, as is done in molecular dynamics. Pair potentials are often used to describe the atomic interactions, and these can be found from quantum mechanical calculations, using Hartree–Fock,density functional theory or quantum Monte Carlo methods. In a dense system, the pair potential is inadequate as the interactions between two particles in the system are influenced by other particles. In order to incorporate these effects in a simulation, it would be necessary to calculate the forces from full electronic structure calculations for all configurations occurring in the simulation. Car and Parrinello have devised a clever way to calculate these forces as the calculation proceeds, by combining density functional theory with molecular dynamics methods.

In the Car–Parrinello approach, electron correlations are not treated exactly because of the reliance on LDA (see Section 1.5), but it will be clear that it is an important improvement on fully classical simulations where the interatomic interactions are described by a simple form, such as pair potentials. It is possible

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

to include some damping mechanism in the equations of motion and then let the nuclei relax to their ground state positions, so that equilibrium configurations of molecules and solids can be determined (neglecting quantum fluctuations).

1.8 Quantum field theory

Quantum field theory provides a quantum description for fields: strings in one dimension, sheets in two dimensions, etc. Quantum field theory is also believed to describe elementary particles and their interactions. The best known example is quantum electrodynamics (QED) which gives a very accurate description of the interaction between charged spin-1/2 fermions (electrons) and electromagnetic fields. The results of QED are obtained using perturbation theory which works very well for this case, because the perturbative parameter remains small for all but the smallest length scales (at large length scales this is the fine structure constant).

In quantum chromodynamics (QCD), the theory which supposedly describes quarks bound together in a proton or neutron, the coupling constant grows large for large scales, and perturbation theory breaks down. One way to obtain useful results for this theory is to discretise space-time, and simulate the theory on this space-time lattice on a computer. This can be done using Monte Carlo or molecular dynamics techniques. The application of these techniques is far from easy as the QCD field theory is intrinsically complicated. A problem which needs to be addressed is efficiency, notably overcoming *critical slowing down*, which decreases the efficiency of simple Monte Carlo and molecular dynamics techniques for the cases which are of physical interest. The fact that quarks are fermions leads to additional complications.

QCD simulations relate quark masses to masses and interaction constants of hadrons (mesons, protons, neutrons).

1.9 About this book

In this book, the emphasis is on methods which do not merely involve straightforward application of numerical methods, and which are specific to problems studied in physics. In most cases, the theory is treated in some detail in order to exhibit clearly what the approximations are and why the methods work. However, some of this theoretical material can be skipped at first reading (this is the material in the sections marked with an asterisk *). Details on implementation are given for most of the methods described.

We start off with a chapter on quantum mechanical scattering theory. This is a rather straightforward application of numerical techniques, and is used as an

Cambridge University Press
978-0-521-83346-2 - Computational Physics, Second Edition
Jos Thijssen
Excerpt
More information

illustration of solving a rather simple (not completely trivial) physical problem on a computer. The results of a sample program are compared with experiment. In Chapters 3 to 5 we discuss computational methods for the electronic structure: variational calculus, Hartree–Fock and density functional theory. We apply these methods to some simple systems: the hydrogen and the helium atoms, and the hydrogen molecule. We calculate the energies of these systems. Chapter 6 deals with solving the independent-particle Schrödinger equation in solids.

In Chapters 7 to 12 we describe molecular dynamics and Monte Carlo techniques for classical and quantum many-particle systems. Chapter 7 contains an overview of classical statistical mechanics, with emphasis on ensembles and on critical phenomena, which are also important for field theory, as discussed in Chapter 15. The molecular dynamics and Monte Carlo techniques are treated in Chapters 8 and 10. The standard example of a molecular liquid, argon, is analysed, but simulations for liquid nitrogen and for lattice spin systems (Ising model) are also discussed. Chapter 9 deals with the quantum molecular dynamics technique.

The relations between classical and statistical mechanics are exploited in Chapter 11 where the transfer matrix method for lattice spin systems is described. The next chapter deals with the application of Monte Carlo methods to quantum mechanics, and we revisit the helium atom which is now treated without Hartree–Fock or DFT approximations.

In Chapter 15 we consider numerical methods for field theory. Techniques for analysing the simplest interesting field theory, the scalar ϕ^4 theory, are studied, and methods for studying more complicated field theories (QED and QCD) are discussed. Because of the relation between statistical and quantum mechanics, some of the techniques discussed in this chapter are also relevant for classical statistical mechanics.

Finally, in Chapter 16 modern computer architectures are briefly considered and an example of a parallel algorithm for molecular dynamics is given.

The algorithms presented, and the programs to be written in the exercises, can be coded in different languages: C, C++, Java, Fortran 77, Fortran 90 etc. Also, an integrated scientific computer environment such as MatLab may be used. They all have their pluses and minuses: Fortran 77 allows for dirty programming, but is quite efficient, and the same holds for C; Fortran 90 is efficient and neat. MatLab is easy to use, but not as efficient as using a high-level programming language. Perhaps the most structured way of programming is by using the objected-oriented programming paradigm, as implemented in the langauges C++ and Java. For large and complex projects, these languages are unbeatable. However, for smaller jobs MatLab or Fortran 90 is usually sufficient. It is my experience that students relatively new to programming get their programs to run correctly most quickly when using Fortran 90 or MatLab.

© Cambridge University Press www.cambridge.org