1
Introduction

1.1 Historical background

The origins of molecular dynamics – MD – are rooted in the atomism of antiquity. The ingredients, while of more recent vintage, are not exactly new. The theoretical underpinnings amount to little more than Newton's laws of motion. The significance of the solution to the many-body problem was appreciated by Laplace [del51]: 'Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective situation of the beings who compose it – an intelligence sufficiently vast to submit these data to analysis – it would embrace in the same formula the movements of the greatest bodies of the universe and those of the lightest atom; for it, nothing would be uncertain and the future, as the past, would be present to its eyes'. And the concept of the computer, without which there would be no MD, dates back at least as far as Babbage, even though the more spectacular hardware developments continue to this day. Thus MD is a methodology whose appearance was a foregone conclusion, and indeed not many years passed after digital computers first appeared before the first cautious MD steps were taken [ald57, gib60, rah64].

The N-body problem originated in the dynamics of the solar system, and the general problem turns out to be insoluble for three or more bodies. Once the atomic nature of matter became firmly established, quantum mechanics took charge of the microscopic world, and the situation became even more complicated because even the constituent particles seemed endowed with a rather ill-defined existence. But a great deal of the behavior of matter in its various states can still be understood in classical (meaning nonquantum) terms, and so it is that the classical N-body problem is also central to understanding matter at the microscopic level. And it is the task of the numerical solution of this problem that MD addresses.

For systems in thermal equilibrium, theory, in the form of statistical mechanics, has met with a considerable measure of success, particularly from the conceptual

point of view. Statistical mechanics provides a formal description – based on the partition function – of a system in equilibrium; however, with a few notable exceptions, there are no quantitative answers unless severe approximations are introduced, and even then it is necessary to assume large (essentially infinite) systems. Once out of equilibrium, theory has very little to say. Simulations of various kinds, including MD, help fill the gaps on the equilibrium side, but in the more general case it is only by means of simulation – principally MD – that progress is possible.

From the outset, the role of computers in scientific research has been a central one, both in experiment and in theory. For the theoretician, the computer has provided a new paradigm of understanding. Rather than attempting to obtain simplified closed-form expressions that describe behavior by resorting to (often uncontrolled) approximation, the computer is now able to examine the original system directly. While there are no analytic formulae to summarize the results neatly, all aspects of the behavior are open for inspection.

1.2 Computer simulation

Science requires both observation and comprehension. Without observation there are no facts to be comprehended; without comprehension science is mere documentation. The basis for comprehension is theory, and the language of theoretical science is mathematics. Theory is constructed on a foundation of hypothesis; the fewer the hypotheses needed to explain existing observations and predict new phenomena, the more 'elegant' the theory – Occam's razor.

The question arises as to how simulation is related to physical theory. University education abounds with elegant theoretical manipulation and is a repository for highly idealized problems that are amenable to closed-form solution. Despite the almost 'unreasonable applicability' of mathematics in science [wig60], the fact is that there is usually a chasm between the statement of a theory and the ability to extract quantitative information useful in interpreting experiment. In the real world, exact solutions are the notable exception. Theory therefore relies heavily on approximation, both analytical and numerical, but this is often uncontrolled and so reliability may be difficult to establish. Thus it might be said that simulation rests on the basic theoretical foundations, but tries to avoid much of the approximation normally associated with theory, replacing it by a more elaborate calculational effort. Where theory and simulation differ is in regard to cost. Theory requires few resources beyond the cerebral and is therefore 'cheap'; simulation needs the hardware and, despite plummeting prices, a computer system for tackling problems at the forefront of any field can still prove costly.

Simulation also draws from experiment. Experimental practice rests on a long (occasionally blemished) tradition; computer simulation, because of its novelty,

is still somewhat more haphazard, but methodologies are gradually evolving. The output of any simulation should be treated by the same statistical methods used in the analysis of experiments. In addition to estimating the reliability of the results (on the assumption that the measurements have been made correctly) there is also the issue of adequate sampling. This is particularly important when attempting to observe 'rare' events: quantitative studies of such events require that the entire occurrence be reproduced as many times as necessary to assure adequate sampling – if computer resources cannot accommodate this requirement it is presumptuous to expect reliable results.

What distinguishes computer simulation in general from other forms of computation, if such a distinction can be made, is the manner in which the computer is used: instead of merely performing a calculation, the computer becomes the virtual laboratory in which a system is studied – a numerical experiment. The analogy can be carried even further; the results emerging from a simulation may be entirely unexpected, in that they may not be at all apparent from the original formulation of the model. A wide variety of modeling techniques have been developed over the years, and those relevant for work at the molecular level include, in addition to MD, classical Monte Carlo [all87, lan00], quantum based techniques involving path-integral [ber86c, gil90] and Monte Carlo methods [sch92], and MD combined with electron density-function theory [rem90, tuc94], as well as discrete approaches such as cellular automata and the lattice–Boltzmann method [doo91].

Although the goal of science is understanding, it is not always obvious what constitutes 'understanding'. In the simulational context, understanding is achieved once a plausible model is able to reproduce and predict experimental observation. Subsequent study may lead to improvements in the model, or to its replacement, in order to explain further experiments, but this is no different from the way in which science is practiced in the broader context. Clearly, there is no inherent virtue in an excessively complex model if there is no way of establishing that all its features are essential for the desired results (Occam again). The practical consequence of this policy is that, despite any temptation to do otherwise, features should be added gradually. This helps with quality control in the notoriously treacherous world of computer programming; since the outcome of a simulation often cannot be predicted with enough confidence to allow full validation of the computation, the incremental approach becomes a practical necessity.

Simulation plays an important role in education. It takes little imagination to see how interactive computer demonstrations of natural phenomena can enrich any scientific presentation. Whether as an adjunct to experiment, a means of enhancing theoretical discussion, or a tool for creating hypothetical worlds, simulation is without peer. Especially in a conceptually difficult field such as physics, simulation can be used to help overcome some of the more counterintuitive concepts

encountered even at a relatively elementary level. As to the role of MD, it can bring to life the entire invisible universe of the atom, an experience no less rewarding for the experienced scientist than for the utter tyro. But, as with education in general, simulation must be kept honest, because seeing is believing, and animated displays can be very convincing irrespective of their veracity.

1.3 Molecular dynamics

Foundations

The theoretical basis for MD embodies many of the important results produced by the great names of analytical mechanics – Euler, Hamilton, Lagrange, Newton. Their contributions are now to be found in introductory mechanics texts (such as [gol80]). Some of these results contain fundamental observations about the apparent workings of nature; others are elegant reformulations that spawn further theoretical development. The simplest form of MD, that of structureless particles, involves little more than Newton's second law. Rigid molecules require the use of the Euler equations, perhaps expressed in terms of Hamilton's quaternions. Molecules with internal degrees of freedom, but that are also subject to structural constraints, might involve the Lagrange method for incorporating geometric constraints into the dynamical equations. Normal equilibrium MD corresponds to the microcanonical ensemble of statistical mechanics, but in certain cases properties at constant temperature (and sometimes pressure) are required; there are ways of modifying the equations of motion to produce such systems, but of course the individual trajectories no longer represent the solution of Newton's equations.

The equations of motion can only be solved numerically. Because of the nature of the interatomic interaction, exemplified by the Lennard-Jones potential with a strongly repulsive core, atomic trajectories are unstable in the sense that an infinitesimal perturbation will grow at an exponential rate†, and it is fruitless to seek more than moderate accuracy in the trajectories, even over limited periods of time. Thus a comparatively low-order numerical integration method often suffices; whether or not this is adequate emerges from the results, but the reproducibility of MD measurements speaks for itself. Where softer interactions are involved, such as harmonic springs or torsional interactions, either or both of which are often used for modeling molecules with internal degrees of freedom, a higher-order integrator, as well as a smaller timestep than before, may be more appropriate to accommodate the fast internal motion. The numerical treatment of constraints introduces an additional consideration, namely that the constraints themselves must be preserved to much higher accuracy than is provided by the integration method, and methods

† This is discussed in §3.8.

exist that address this problem. All these issues, and more, are covered in later chapters.

While MD is utterly dependent on the now ubiquitous computer, an invention of the twentieth century, it pays little heed to the two greatest developments that occurred in physics in the very same century – relativity and quantum mechanics. Special relativity proscribes information transfer at speeds greater than that of light; MD simulation assumes forces whose nature implies an infinite speed of propagation. Quantum mechanics has at its base the uncertainty principle; MD requires – and provides – complete information about position and momentum at all times. In practice, the phenomena studied by MD simulation are those where relativistic effects are not observed and quantum effects can, if necessary, be incorporated as semiclassical corrections – quantum theory shows how this should be done [mai81]. But, strictly speaking, MD deals with a world that, while intuitively appealing to late nineteenth-century science, not to mention antiquity, has little concern for anything that is 'nonclassical'. This fact has in no way diminished the power and effectiveness of the method.

Relation to statistical mechanics

Statistical mechanics (for example [mcq76]) deals with ensemble averages. For the canonical ensemble, in which the temperature T and number of particles N_m are fixed, the equilibrium average of some quantity G is expressed in terms of phase-space integrals involving the potential energy $U(\boldsymbol{r}_1, \ldots \boldsymbol{r}_{N_m})$,

$$\langle G \rangle = \frac{\int G(\boldsymbol{r}_1, \ldots \boldsymbol{r}_{N_m}) e^{-\beta U(\boldsymbol{r}_1, \ldots \boldsymbol{r}_{N_m})} \, d\boldsymbol{r}_1 \cdots \boldsymbol{r}_{N_m}}{\int e^{-\beta U(\boldsymbol{r}_1, \ldots \boldsymbol{r}_{N_m})} \, d\boldsymbol{r}_1 \cdots \boldsymbol{r}_{N_m}} \tag{1.3.1}$$

where $\{\boldsymbol{r}_i \,|\, i = 1, \ldots N_m\}$ are the coordinates, $\beta = 1/k_B T$, and k_B is the Boltzmann constant. This average corresponds to a series of measurements over an ensemble of independent systems.

The ergodic hypothesis relates the ensemble average to measurements carried out for a single equilibrium system during the course of its natural evolution – both kinds of measurement should produce the same result. Molecular dynamics simulation follows the dynamics of a single system and produces averages of the form

$$\langle G \rangle = \frac{1}{M} \sum_{\mu=1}^{M} G_\mu(\boldsymbol{r}_1, \ldots \boldsymbol{r}_{N_m}) \tag{1.3.2}$$

over a series of M measurements made as the system evolves. Assuming that the

sampling is sufficiently thorough to capture the typical behavior, the two kinds of averaging will be identical. The observation that MD corresponds to the microcanonical (constant energy) ensemble, rather than to the canonical (constant temperature) ensemble, will be addressed when it appears likely to cause problems.

Relation to other classical simulation methods

The basic Monte Carlo method [lan00] begins by replacing the phase-space integrals in (1.3.1) by sums over states

$$\langle G \rangle = \frac{\sum_s G(s) e^{-\beta U(s)}}{\sum_s e^{-\beta U(s)}} \tag{1.3.3}$$

Then, by a judicious weighting of the states included in the sum, which for the general case results in

$$\langle G \rangle = \frac{\sum_s W(s)^{-1} G(s) e^{-\beta U(s)}}{\sum_s W(s)^{-1} e^{-\beta U(s)}} \tag{1.3.4}$$

where $W(s)$ is the probability with which states are chosen, (1.3.4) can be reduced to a simple average over the S states examined, namely,

$$\langle G \rangle = \frac{1}{S} \sum_{s=1}^{S} G(s) \tag{1.3.5}$$

Clearly, we require

$$W(s) = e^{-\beta U(s)} \tag{1.3.6}$$

for this to be true, and much of the art of Monte Carlo is to ensure that states are actually produced with this probability; the approach is called importance sampling. The Monte Carlo method considers only configuration space, having eliminated the momentum part of phase space. Since there are no dynamics, it can only be used to study systems in equilibrium, although if dynamical processes are represented in terms of collision cross sections it becomes possible to study the consequences of the process, even if not the detailed dynamics [bir94].

Molecular dynamics operates in the continuum, in contrast to lattice-based methods [doo91], such as cellular automata, which are spatially discrete. While the latter are very effective from a computational point of view, they suffer from certain design problems such as the lack of a range of particle velocities, or unwanted effects due to lattice symmetry, and are also not easily extended. The MD approach is computationally demanding, but since it attempts to mimic nature it has few

inherent limitations. One further continuum-dynamical method, known as Brownian dynamics [erm80], is based on the Langevin equation; the forces are no longer computed explicitly but are replaced by stochastic quantities that reflect the fluctuating local environment experienced by the molecules.

Applications and achievements

Given the modeling capability of MD and the variety of techniques that have emerged, what kinds of problem can be studied? Certain applications can be eliminated owing to the classical nature of MD. There are also hardware imposed limitations on the amount of computation that can be performed over a given period of time – be it an hour or a month – thus restricting the number of molecules of a given complexity that can be handled, as well as storage limitations having similar consequences (to some extent, the passage of time helps alleviate hardware restrictions).

The phenomena that can be explored must occur on length and time scales that are encompassed by the computation. Some classes of phenomena may require repeated runs based on different sets of initial conditions to sample adequately the kinds of behavior that can develop, adding to the computational demands. Small system size enhances the fluctuations and sets a limit on the measurement accuracy; finite-size effects – even the shape of the simulation region – can also influence certain results. Rare events present additional problems of observation and measurement.

Liquids represent the state of matter most frequently studied by MD methods. This is due to historical reasons, since both solids and gases have well-developed theoretical foundations, but there is no general theory of liquids. For solids, theory begins by assuming that the atomic constituents undergo small oscillations about fixed lattice positions; for gases, independent atoms are assumed and interactions are introduced as weak perturbations. In the case of liquids, however, the interactions are as important as in the solid state, but there is no underlying ordered structure to begin with.

The following list includes a somewhat random and far from complete assortment of ways in which MD simulation is used:

- Fundamental studies: equilibration, tests of molecular chaos, kinetic theory, diffusion, transport properties, size dependence, tests of models and potential functions.
- Phase transitions: first- and second-order, phase coexistence, order parameters, critical phenomena.
- Collective behavior: decay of space and time correlation functions, coupling of translational and rotational motion, vibration, spectroscopic measurements, orientational order, dielectric properties.

- Complex fluids: structure and dynamics of glasses, molecular liquids, pure water and aqueous solutions, liquid crystals, ionic liquids, fluid interfaces, films and monolayers.
- Polymers: chains, rings and branched molecules, equilibrium conformation, relaxation and transport processes.
- Solids: defect formation and migration, fracture, grain boundaries, structural transformations, radiation damage, elastic and plastic mechanical properties, friction, shock waves, molecular crystals, epitaxial growth.
- Biomolecules: structure and dynamics of proteins, protein folding, micelles, membranes, docking of molecules.
- Fluid dynamics: laminar flow, boundary layers, rheology of non-Newtonian fluids, unstable flow.

And there is much more.

The elements involved in an MD study, the way the problem is formulated, and the relation to the real world can be used to classify MD problems into various categories. Examples of this classification include whether the interactions are short- or long-ranged; whether the system is thermally and mechanically isolated or open to outside influence; whether, if in equilibrium, normal dynamical laws are used or the equations of motion are modified to produce a particular statistical mechanical ensemble; whether the constituent particles are simple structureless atoms or more complex molecules and, if the latter, whether the molecules are rigid or flexible; whether simple interactions are represented by continuous potential functions or by step potentials; whether interactions involve just pairs of particles or multiparticle contributions as well; and so on and so on.

Despite the successes, many challenges remain. Multiple phases introduce the issue of interfaces that often have a thickness comparable to the typical simulated region size. Inhomogeneities such as density or temperature gradients can be difficult to maintain in small systems, given the magnitude of the inherent fluctuations. Slow relaxation processes, such as those typical of the glassy state, diffusion that is hindered by structure as in polymer melts, and the very gradual appearance of spontaneously forming spatial organization, are all examples of problems involving temporal scales many orders of magnitude larger than those associated with the underlying molecular motion.

1.4 Organization

Case studies

Case studies are used throughout. The typical case study begins with a review of the theoretical background used for formulating the computational approach. The

computation is then described, either by means of a complete listing of the functions that make up the program, or as a series of additions and modifications to an earlier program. Essential but often neglected details such as the initial conditions, organization of the input and output, accuracy, convergence and efficiency are also addressed.

Results obtained from running each program are shown. These sometimes reproduce published results, although no particular effort is made to achieve a similar level of accuracy since our goal is one of demonstration, not of compiling a collection of definitive measurements. Suggested extensions and assorted other projects are included as exercises for the reader.

We begin with the simplest possible example, to demonstrate that MD actually works. Later chapters extend the basic model in a variety of directions, improve the computational methods, deal with various kinds of measurement and introduce new models for more complex problems. The programs themselves are constructed incrementally, with most case studies building on programs introduced earlier. In order to avoid a combinatorial explosion, the directions explored in each chapter tend to be relatively independent, but in more ambitious MD applications it is quite likely that combinations of the various techniques will be needed. Some care is necessary here, because what appears obvious and trivial for simple atoms may, for example, require particular attention for molecules subject to constraints – each case must be treated individually.

Itinerary

Chapter 2 introduces the MD approach using the simplest possible example, and demonstrates how the system behaves in practice; general issues of programming style and organization that are used throughout the book are also introduced here. In Chapter 3 we discuss the methodology for simulating monatomic systems, the algorithms used, and the considerations involved in efficient and accurate computation. Chapter 4 focuses on measuring the thermodynamic and structural properties of systems in equilibrium; some of these properties correspond to what can be measured in the laboratory, while others provide a microscopic perspective unique to simulation. The dynamical properties of equilibrium systems are the subject of Chapter 5, including transport coefficients and the correlation functions that are associated with space- and time-dependent processes.

More complex systems and environments form the subject of subsequent chapters. Modifications of the dynamics to allow systems to be studied under conditions of constant temperature and pressure, as opposed to the constant energy and volume implicit in the basic MD approach, are covered in Chapter 6. In Chapter 7 we discuss further methods for measuring transport properties, both by modeling the

relevant process directly and by using a modified form of the dynamics designed for systems not in thermal equilibrium. The dynamics of rigid molecules forms the subject of Chapter 8; methods for handling the general problem are described and a model for water is treated in some detail. Flexible molecules are discussed in Chapter 9 and a model for surfactants examined. Molecules possessing internal degrees of freedom, but also subject to geometric constraints that provide a certain amount of rigidity, are analyzed in Chapter 10, together with a model used for simulating alkane chains. An alternative route to dealing with molecules having internal degrees of freedom, based on treating the internal coordinates directly, is described in Chapter 11. Approaches used for three-body and many-body interactions are introduced in Chapter 12. Specialized methods for treating long-range forces involving Ewald sums and multipole expansions are discussed in Chapter 13.

Chapter 14 describes an alternative approach to MD based on step potentials, rather than on the continuous potentials of earlier chapters; this calls for entirely different computational techniques. In Chapter 15 we focus on the study of time-dependent behavior and demonstrate the ability of MD to reproduce phenomena normally associated with macroscopic hydrodynamics. The methods developed for MD can also be applied to studying the dynamics of granular materials; a short introduction to this subject appears in Chapter 16. The special considerations that are involved in implementing MD computations on parallel and vector supercomputers form the subject of Chapter 17. Chapter 18 deals with a range of software topics not covered by the case studies. And, finally, some closing thoughts on where MD may be headed appear in Chapter 19. A concise alphabetical summary of the variables used in the software and a list of the programs that are available for use with the book appear in the Appendix.

1.5 Further reading

A great deal of information about MD methodology and applications is scattered throughout the scientific literature, and references to material relevant to the subjects covered here will appear in the appropriate places. Three volumes of conference proceedings include pedagogical expositions of various aspects of MD simulation [cic86a, cat90, all93b] and a monograph on liquid simulation covers both MD and Monte Carlo techniques [all87]. Another book devoted in part to MD is [hoo91]. Three evenly spaced reviews of the role of simulation in statistical mechanics are [bee66, woo76, abr86]. Two extensive literature surveys on liquid simulation [lev84, lev92] and a collection of reprints [cic87] are also available.