Preface

Since this book is aimed at a broad audience within the physical sciences, we expect most of our readers not to be experts in either astrophysics or mathematics. For those readers, the title of this book may seem puzzling at least. Why should they be interested in the *gravitational* attraction between bodies? What is so special about a *million*indexmillion-body problem, rather than a billion or a trillion bodies? What kind of *bodies* do we have in mind? And finally, what is the *problem* with this whole topic?

In physics, many complex systems can be modeled as an aggregate of a large number of relatively simple entities with relatively simple interactions between them. It is one of the most fascinating aspects of physics that an enormous richness can be found in the collective phenomena that emerge out of the interplay of the much simpler building blocks. Smoke rings and turbulence in air, for example, are complex manifestations of a system of air molecules with relatively simple interactions, strongly repulsive at small scales and weakly attractive at larger scales. From the spectrum of avalanches in sand piles to the instabilities in plasmas of more than a million degrees in labs to study nuclear fusion, we deal with one or a few constituents with simple prescribed forces. What is special about *gravitational* interactions is the fact that gravity is the only force that is mutually attractive. Unlike a handful of protons and electrons, where like charges repel and opposite charges attract, a handful of stars shows attraction between every pair of stars. As a result, a star cluster holds itself together: there is no need for a container (as with a gas or plasma in a lab) or a table (as with a sand pile). And for astronomers, an extra reason to study star clusters is simply: because they are there.

Now why are we interested in a *million* stars? Usually in physics we analyze a system of interacting components in two limits: one in which the number of components is one, or a few; and one in which the number of components tends to infinity. And at first sight, a million and a billion stars both seem to be 'close enough' to infinity to allow a common treatment. However, there is a large difference in behavior between the two types of systems. In a typical rich star cluster, with a million stars, each star feels enough of the granularity of the gravitational field of the other stars that the consequent perturbations lead to a total loss of memory of the initial conditions of its orbit, within the life time of the Universe. In contrast, in a typical galaxy, with a number of stars between a billion and a trillion, an individual star for all intents and purposes feels only the smooth average background field. This is the reason that galaxies still retain much prettier shapes than the shapeless (indeed globular) clusters for which orbital memories have been wiped out. In other words, in this part of astronomy, when we count 'one, two, a few, many', a few means a million.

The term *body* is just an archaic term for a material object, be it a molecule or a stone or a star. In many cases in astronomy, when we study the motion of a group of stars we can completely neglect the finite size of the stars, and we may as well treat them as if they were simple mass points, with a mass but without an extension. When stars come closer, we may want to approximate them as finite-size bodies, with a simple description of their mass distribution. Only when stars physically collide, as they sometimes do in the dense inner regions of star clusters, do we have to remind ourselves how complex these celestial bodies really are*, in order to say something about the transformations that stars undergo when they merge.

Finally, why do we call all this a *problem*? In mathematics, more than in physics, we talk about specific topics as problems, as long as they have not (yet) been solved. For example, the famous four-colour problem, the question of how many different colours we need to introduce to colour a map such that neighboring countries can always be given different colours, remained a real problem until the seventies. Partly with the help of computers, this problem was solved (the answer was: four colours suffice), and then the four-body problem turned into the four-colour theorem. In the days of Newton, at the beginning of mathematical physics, the same term 'problem' was used to describe the challenge to find the motions of two, three, or more bodies under the influence of their

* Addressing an astronomer who had just remarked that "after all, a star is a pretty simple thing", R.O. Redman pointed out, that, "at a distance of 10 parsecs *you'd* look pretty simple!"

mutual gravitational attractions: hence the two-body problem, the three-body problem, etc. Unlike the four-colour problem, though, we don't expect to ever 'solve' the gravitational N-body problem, for arbitrary N – and even for $N = 2$, where we *do* have analytic solutions, we still follow tradition in calling this the two-body *problem*. So apart from this quaint piece of history, we could have called this book 'the gravitational million-body system'.

Why would someone want to study the gravitational million-body problem? There are at least four quite different motivations, centering on the fields of astronomy, theoretical physics, computational physics, and mathematics.

Let us first take the point of view of *astronomy*. Throughout the last hundred years or more, astronomers have been studying an important class of objects called *globular star clusters*, which are roughly spherical collections of stars, each much smaller than a galaxy but much bigger than the solar system. The number of stars they contain varies from one object to another, but a million is the right order of magnitude. Astronomers study globular clusters for many reasons, of which we can touch on only a few. It is thought that they were among the first recognisable structures that were born in galaxies like ours, and their age is a vital constraint on that of the universe. All the stars in each cluster were born with different masses at roughly the same time, and their present stellar population gives a snapshot of the results of about 10 billion years of stellar evolution. They are exceptionally rich in some of the more exotic kinds of star which astronomers now study: binary X-ray sources, millisecond pulsars, etc. Finally, the stars inside a cluster are sufficiently densely packed that they can and do collide with each other, and so clusters provide us with a sort of laboratory where we may hope to understand the more dramatic effects of the dense stellar environment in the nuclei of galaxies. Thus the gravitational million-body problem has an important place in modern astronomy.

Now let us consider the point of view of *theoretical physics*. Whichever way a physicist approaches them, gravitational problems pose particular difficulties. Classical thermodynamics is poorly developed for such problems, because of the long-range nature of the gravitational force. Kinetic theory requires ad hoc approximations, for the same reasons. The methods of plasma physics fail because gravitational forces, being attractive, are unshielded. Gravitational systems are intractable with many traditional methods because the natural equilibria are not spatially uniform. The gravitational many-body problem also has a wider significance in physics, because of its historical roots. The science of mechanics became firmly rooted, in large part, because of the success of Newton's programme for the study of planetary motion. The further development of this theory led to the development of Lagrangian methods, and then, in the hands of Poincaré, to the foundation of the study of chaos. Much of the development of physics since the time of Newton has rested on foundations set in place with the aid of the gravitational many-body problem, which became the model for how a successful physical theory should look. Finally, the most difficult models to predict and to interpret their behavior are the ones where we are dealing with neither a very large nor a very small number of particles. When there are huge numbers of particles, as in a gas, there are successful approximate methods, like fluid mechanics, which greatly simplify the problems. When there are very few particles, especially one or two, the problem is either completely soluble or can be understood approximately. When there are an intermediate number (a million or less in our case), elements of both extremes are in play simultaneously, thereby thwarting either type of approximation.

This brings us naturally to *computational physics*. Just as the theoretical physicist does, the computational physicist immediately grasps the fact that the gravitational million-body problem poses formidable problems whichever way it is approached. It is not even clear whether the solution can be obtained reliably at all. Even if we lay such fundamental issues aside, running an N-body simulation with $N \gtrsim 10^4$ is exceptionally time-consuming for even the fastest generally available computers. This in turn has led to the development of special computers (the GRAPE family), whose sole task is to solve this problem very quickly. Around 1995 these were the fastest computers in the world, a place they recaptured in 2001. This approach has become a model of the way in which future work in computational science in other areas may be carried out economically and quickly, given the right flair and ingenuity. Providing the hardware is only half the problem though; developing the software for running N-body simulations is equally challenging. In our particular case of globular cluster simulations, we are confronted with length scales spanning the range from kilometers to parsecs (a factor of more than 10^{13}), and with time scales spanning the range from milliseconds to the life time of the universe (a factor of more than 10^{20}). As a result, stellar dynamicists have developed special integration methods that are not encountered in any text books on differential equations.

And finally we turn to *mathematics*. When Newton laid the foundations of classical physics in the *Principia*, among the mathematical tools he deployed was the infinitesimal calculus, which he had invented for the purpose. This illustrates the rich potential for the invention of new mathematics which results from intense scrutiny of gravitational problems, and other issues in dynamics. Newton set a precedent which has repeated itself several times since. The work of Poincaré was developed in the context of celestial mechanics, and the famous theorem of Kolmogorov, Arnold and Moser, foreshadowed as it was by the work of C.L. Siegel in his book "Vorlesungen über Himmelsmechanik"*, owes much to this discipline, and the emphasis which it helped to place on Hamiltonian problems. And yet when one turns to the *million*-body problem, as opposed to the *few*-body problems of the solar system, the flow of information has been from mathematics to stellar dynamics, rather than the other way around. A number of techniques introduced by mathematicians for the solution of quite abstract problems have turned out to be just what was needed to improve the numerical solution of our problem. This flow of ideas has been sufficiently influential that the astronomer's understanding of the problem would not have advanced so far without it. One of our purposes in writing this book, perhaps a far-fetched one, is the hope that it may help reverse the flow, and stimulate the birth of new mathematics. At the very least, we hope that mathematicians will enjoy learning how their work has been put to use. As one of us has enjoyed a place in the tolerant community of mathematicians for many years, it is a way of giving something back.

We believe that the gravitational million-body problem has a seminal role to play in all four areas: astronomy, physics, mathematics, computational science. But while the above remarks have sometimes emphasised the difficulties it poses, we will have failed in one of our aims if, after reading the book, the reader from any of these disciplines is not impressed by how much of the problem can be understood, and how much it has to offer. In other words, one of our goals is to convey some of the beauty and simplicity underlying classical dynamics, as illustrated through the gravitational many-body problem. This is a book that is meant to be read by a variety of scientists who share a curiosity for the roots of physics, the recent fruits that have sprouted directly from those old roots, and its interconnections with neighbouring sciences.

That having been said, it is mostly within the astronomical community that this subject has been developed. And it has been another firm intention in writing the book that it may serve the role of a *graduate textbook* on the theory of stellar dynamics of dense stellar systems. Wherever possible, therefore, our statements and results are developed from first principles. We include some exercises and problems, and hints for their solution.

In writing the book we were aware that the reader may or may not be an astronomer, and may or may not be interested in learning the details of

* Lectures on Celestial Mechanics

the theory. We realised also that the variety of topics which make up the gravitational million-body problem are best treated at a variety of levels, and that the interconnectedness of these topics forces interconnections between the chapters. Some of the topics we address might even strike some readers as being faintly whimsical. Recognising, therefore, that not all chapters will be equally accessible or interesting to all readers, we start each part of the book (each of which consists of several connected chapters) with an outline of its contents. We hope that readers will thus find their own optimal route between the contents page and the index.

For much the same variety of reasons we have relegated some material to boxes. Often these contain details of some derivation or discussion, and a reader in a hurry could avoid them. Sometimes, however, they contain background which might help a reader over some difficulty. Sometimes they collect some useful results which might act as a reference resource. Sometimes they explore interesting or amusing by-ways that would otherwise interrupt the text. Perhaps, therefore, they should not be passed over too readily.

We hope the index will both help and intrigue the reader, but the quirky and personal selection of names there is not meant as a comprehensive answer to the question of who did what. (Last time we looked, we could not even find our own names there.) Even in the extensive list of references to published work we have not attempted completeness. While this risks antagonizing the reader for the omission of his or her own most cherished paper, for which we apologise, it could be just as hazardous if we tried to attribute every advance to its originator. Our main reason for giving references is to give the reader an entry into the research literature. Nowadays it is quite easy to move both forwards and backwards from a suitable entry point, using such invaluable bibliographical resources as the Astronomy Abstract Service of the NASA Astrophysics Data System.

And now we have a confession to make. Studying problems in theoretical physics, the solved as well as the unsolved, is interesting and enjoyable all by itself. Finding new and unexpected insights, and extending the world's body of knowledge about these problems, is an added pleasure. Sharing these interests with others is still more rewarding. But we have other reasons for our interest in the gravitational N-body problem, and among these is its status as one of the *oldest* unsolved problems in the exact sciences, and one with an exceptionally distinguished pedigree. Some of the great names of the subject have been mentioned in this preface, but we would add some other much admired names of the past, including those of Jeans, Chandrasekhar and Spitzer. With equal pleasure we call to mind the many contemporaries with whom we have worked and continue to work on these problems. The community is world-wide, but it is not a large one, it is rather close-knit, and works together and openly in the best tradition of the "community of scholars". We hope that the common enjoyment of this joint enterprise surfaces from time to time throughout this book.

Acknowledgements

We have a number of these colleagues to thank with particular regard to the help they have given with parts of the book. Large sections have been read and commented on by Haldan Cohn, Herwig Dejonghe, Mirek Giersz, Jeremy Goodman, Jarrod Hurley, Jun Makino, Simon Portegies Zwart, Rainer Spurzem and Peter Teuben, and we thank them warmly for their comments. If there are problems with the book, it is probably because we did not *always* take their advice. Simon Portegies Zwart and Chris Eilbeck also gave considerable help with appendix A, and we also thank Maurizio Salaris and Max Ruffert for their input. For permission to reproduce illustrations we would like to thank O. Gnedin, J. Goodman, L. Hernquist, K. Johnston, S. Kulkarni, C. Lada, J. Makino, S. McMillan, J. McVean, E. Milone, F. Rasio, the American Astronomical Society, Blackwell Publishing (publishers of Monthly Notices of the Royal Astronomical Society), and other individuals and institutions named beside individual figures. We are also most grateful to staff at Cambridge University Press, especially A. Black (who has since moved on), J. Garget and S. Mitton, for their advice, encouragement and patience. DCH is deeply grateful to the Institute for Advanced Study, Princeton, for its generous hospitality over several visits during which much of the planning and drafting of the book was carried out.

Edinburgh and Princeton
2002