



Universiteit
Leiden
The Netherlands

Chaotic Dynamics in N-body systems

Boekholt, T.C.N.

Citation

Boekholt, T. C. N. (2015, November 10). *Chaotic Dynamics in N-body systems*. Retrieved from <https://hdl.handle.net/1887/36077>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/36077>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/36077> holds various files of this Leiden University dissertation

Author: Boekholt, Tjarda

Title: Chaotic dynamics in N-body systems

Issue Date: 2015-11-10

On the Reliability of N-body Simulations

Based on: *On the Reliability of N-body Simulations* by T. C. N. Boekholt and S. F. Portegies Zwart in *Computational Astrophysics and Cosmology*, Volume 2, article id. #2, 21 pp. (2015), Ch.4-7

As mentioned in Chapter 2, the general consensus in the N-body community is that statistical results of an ensemble of collisional N-body simulations are accurate, even though individual simulations are not. In order to test this assumption, we developed the new N-body code Brutus that solves the N-body problem to a pre-defined precision.

Using this new, brute force N-body code, we test the reliability of N-body simulations by a controlled numerical experiment. In this experiment we perform a series of resonant 3-body simulations, where the term resonant implies a phase or multiple phases during the interaction where the stars are more or less equidistant (Hut & Bahcall, 1983). These phases are intermingled by ejections, where a binary and single star are clearly separated. We perform the simulations with conventional double-precision, and with arbitrary-precision to reach the converged solution. In Sec. 4.1 we explain the experiment in more detail, and in Sec. 4.2 we compare the solutions individually to investigate the distribution of the errors. We also compare the global statistical distributions using two-sample Kolmogorov–Smirnov tests (??).

In summary, we find that on average at least half of the conventional simulations diverge from the converged solution, such that the two solutions are microscopically incomparable. For the solutions which have not diverged significantly, we observe that if the integrator has a bias in energy and angular momentum, this propagates to a bias in the statistical properties of the binaries. In the case when the conventional solution has diverged onto an entirely different trajectory in phase-space, we find that the errors are centred around zero and symmetric; the error due to divergence is unbiased, as long as the time-

step parameter, $\eta \leq 2^{-5}$ and when simulations which violate energy conservation by more than 10% are excluded. For resonant 3-body interactions, we conclude that the statistical results of an ensemble of conventional solutions are indeed accurate.

4.1 PRECISION OF STATISTICAL RESULTS: EXPERIMENTAL SETUP

In the previous section we demonstrated that it is possible to obtain a converged solution for a particular initial condition. We have also shown that a solution obtained by Hermite diverges from the converged solution, even up to the point that the microscopic solution given by Hermite is beyond recognition. We now perform a statistical study, to examine the hypothesis that double-precision N-body simulations produce statistically indistinguishable results, from those obtained from an ensemble of converged solutions with the same set of initial conditions. Because it is computationally expensive to reach convergence, we start investigating the hypothesis above by exploring the accuracy of 3-body statistics.

The $N = 3$ experiment is inspired by the Pythagorean problem, where after a complex 3-body interaction, a binary and an escaper are formed. As a variation to this, we define four different sets of initial conditions as follows:

1. Plummer distribution equal mass
2. Plummer distribution with masses 1:2:4
3. Plummer distribution equal mass with zero velocities
4. Plummer distribution with masses 1:2:4 and zero velocities.

The positions and velocities of the three stars are selected randomly from a virialised Plummer distribution (??). For the cold collapse systems, we set the velocities to zero. Then we rescale the positions and velocities to virialise the systems if the initial velocities are non-zero, or we set the total energy equal to $E = -0.25$ if the system starts out cold. We adopt standard Hénon units (??) throughout.

In the case of the cold initial conditions, the systems start democratically, i.e. the minimal distance between each pair of particles is greater than N^{-1} . We reject initial conditions in which this criterion is not satisfied. This is to prevent initial realisations where two stars which are very near, fall to each other radially causing very long

wall-clock times for the integration. When starting with a democratic configuration, there will also be an initial close triple encounter (?), which is hard to integrate accurately and is therefore a good test. A total of 10000 random realisations are generated for each set of initial conditions and can be found in the accompanying data files.

We stop the simulations when the system is dissolved into a permanent binary and an escaper. The criteria used to detect an escaper are the following:

1. escaper has a positive energy, $E > 0$,
2. is a certain distance away from the center of mass, $r > 2 r_{\text{virial}}$,
3. is moving away from the center of mass, $r \cdot v > 0$,

The energy of the escaper is calculated in the barycentric frame of the three particles and r_{virial} is the virial radius of the system, which is of the order unity in Hénon units.

There may be situations in which a star is ejected without actually escaping from the binary. After a long excursion the star turns around and once again engages the binary in a 3-body resonance (Hut & Bahcall, 1983). Because these systems need to be integrated for a longer time, they also require higher precision to reach convergence, which takes a long time to integrate (see also ?). To deal with this issue, we perform the simulations iteratively by increasing the final integration time t_{end} . Starting with $t_{\text{end}} = 50$ Hénon time units, we evolve every system and detect those that are dissolved. Then we increase t_{end} to 100, 150, 200 etc., but only for those systems which have not yet dissolved. A complete ensemble of solutions is obtained up to $t_{\text{end}} \sim 500$, or equivalently ~ 180 crossing times where the crossing time has a value of $2\sqrt{2}$ in Hénon units (??). Systems which take a longer time to integrate are not taken into account in this research. The fraction of long-lived systems is however a statistic we measure. We gathered the final, converged configurations in the accompanying data files.

Each initial realisation is run with the Hermite code, using standard double-precision, and with Brutus, using arbitrary-precision until a converged solution is obtained. At the end of each simulation, we investigate the nature of the binary and the escaper. In addition to the BS tolerance, word-length, CPU time and dissolution time, we record the mass, speed and escape direction of the escaping single star, and the semimajor axis, binding energy and eccentricity of the binary. In this way, we obtain statistics for $N = 3$ generated by a conventional N-body solver and by Brutus.

4.2 RESULTS

Before we perform a detailed comparison between results obtained by Hermite and Brutus, we first compare the Brutus results with analytical distributions from the literature in order to relate to previous studies. We compare Hermite and Brutus on a global level by performing two-sample Kolmogorov–Smirnov tests (??) to see whether global distributions are statistically indistinguishable. We also compare the distribution of lifetimes of triples to see whether precision influences the stability and we measure the typical CPU time and BS tolerance needed to obtain a converged solution. After this, we compare Hermite and Brutus per individual system, with the aim of investigating the nature of the differences of every individual outcome. Finally, we define categories which classify a conventional simulation as a preservation or exchange, depending on whether the identity of the escaping star is consistent between Hermite and Brutus.

4.2.1 Brutus versus Analytical Distributions

In Fig. 4.1, the distributions obtained by converged solutions are given for the following quantities: velocity and kinetic energy of the escaper in the barycentric reference frame, and semimajor axis, binding energy and eccentricity of the binary. We start by looking at the eccentricity distributions (bottom panel in Fig. 4.1). These distributions can be estimated analytically by assuming that the probability of a certain configuration is proportional to the associated volume in phase space (??) or by considering an equilibrium distribution of binary stars in a cluster (Heggie, 1975). The resulting thermal distribution in the three-dimensional case is given by

$$f(e) = 2e, \quad (4.1)$$

and in the two-dimensional case by

$$f(e) = \frac{e}{\sqrt{1-e^2}}, \quad (4.2)$$

The 3-body cold collapse problem is essentially a two-dimensional problem. We compare the empirical and theoretical distributions by means of the K–S test (see also next section). It turns out that the distributions in eccentricity are statistically distinguishable. By inspection by eye we observe that in the virialised case, there are slight deviations at high eccentricities. In the case of the equal-mass, cold

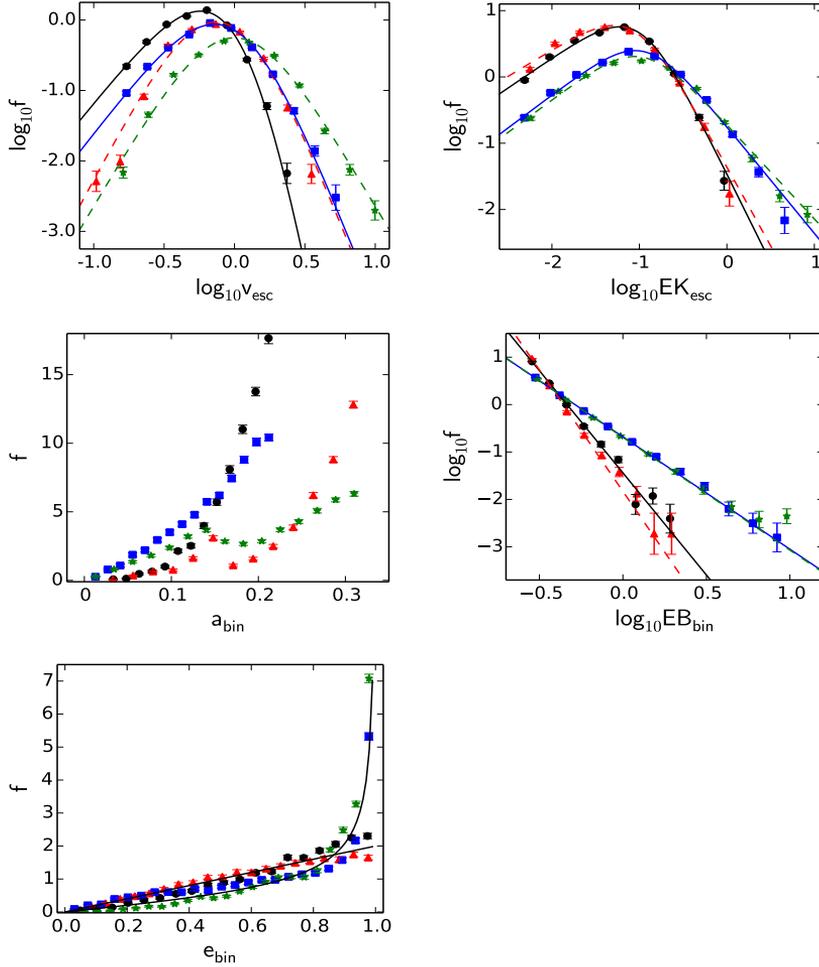


Figure 4.1: Comparison of Brutus results and analytical distributions. Distributions are given for the escaper speed (top left) and kinetic energy (top right), binary semimajor axis (middle left), binding energy (middle right) and binary eccentricity (bottom). The results from the Brutus simulations are represented by the data points, for each of the four sets of initial conditions: Plummer equal mass (bullets), Plummer with different masses (triangles), cold Plummer equal mass (squares) and cold Plummer with different masses (stars). Note that we use standard Hénon units (??). Analytical models from the literature are fitted to the empirical distributions represented by the curves. For the eccentricities we plot the thermal distributions.

Velocity	α	β		
Plummer equal mass	2.5 ± 0.09	6.7 ± 1.02		
Plummer mass ratio	3.8	0.16	4.4	0.43
Cold Plummer equal mass	2.6	0.19	3.8	0.28
Cold Plummer mass ratio	3.4	0.45	3.4	0.19
Kinetic energy				
Plummer equal mass	0.9	0.02	1.8	0.04
Plummer mass ratio	0.8	0.02	1.6	0.04
Cold Plummer equal mass	0.99	0.02	1.3	0.03
Cold Plummer mass ratio	0.98	0.03	1.2	0.02
Binding energy				
Plummer equal mass	4.31	0.13		
Plummer mass ratio	5.12	0.32		
Cold Plummer equal mass	2.37	0.11		
Cold Plummer mass ratio	2.38	0.12		

Table 4.1: Fitted power law indices for the velocity and kinetic energy distributions of the escaping stars and for the binding energy distribution of the binary stars. Note that we use equal intervals in logarithmic space.

systems, there are more low eccentricity binaries compared to the theoretical prediction. They coincide at an eccentricity of about 0.7, after which they deviate again. For the cold systems with unequal masses, this behaviour is the other way around. The analytical predictions are able to capture the empirical distributions only in a qualitative manner.

The velocity distribution of the single escaping star can be estimated analytically in a similar way as was done for the eccentricities. The resulting distribution is predicted to be a double power law given by (??):

$$f(v) \propto \frac{v^\alpha}{(1 + \gamma v^2)^\beta}. \quad (4.3)$$

We fit this model to the data (see Fig. 4.1, first panel) and obtain values for α and β which are given in Table 4.1. The power law indices vary with mass ratio and total angular momentum. To remove the dependence on mass ratio, we plot the kinetic energy of the escaper (see Fig. 4.1, top right panel). Again, we fit a double power law of a similar form as Eq. 4.3, and the power law indices are given in Table 4.1. Both the escaper velocity and kinetic energy are consistent with a double power law distribution.

The binary semimajor axis and binding energy are related quantities. We fit the binding energy distribution (see Fig. 4.1, middle right panel) to a power law (Heggie, 1975; ?; ?):

$$f(E_B) \propto E_B^{-\alpha}. \quad (4.4)$$

The fitted power law indices are given in Table 4.1. The empirical distributions are consistent with a power law, although somewhat steeper than predicted (Heggie, 1975; ?; ?). The slopes do tend to vary somewhat as a function of angular momentum (??).

The empirical distributions obtained by Brutus are in qualitative agreement with the analytical estimates present in the literature (Heggie, 1975; ?; ?). Slight variations are present due to the dependence on total angular momentum, a limited statistical sampling and assumptions made in the derivation of the analytical distributions. Nevertheless, a similar qualitative agreement has been obtained between the analytical distributions discussed above and empirical distributions from an ensemble of conventional numerical solutions, e.g. not converged (? , chapters 7–8 and references therein). The question remains to what extent conventional and converged solutions agree quantitatively.

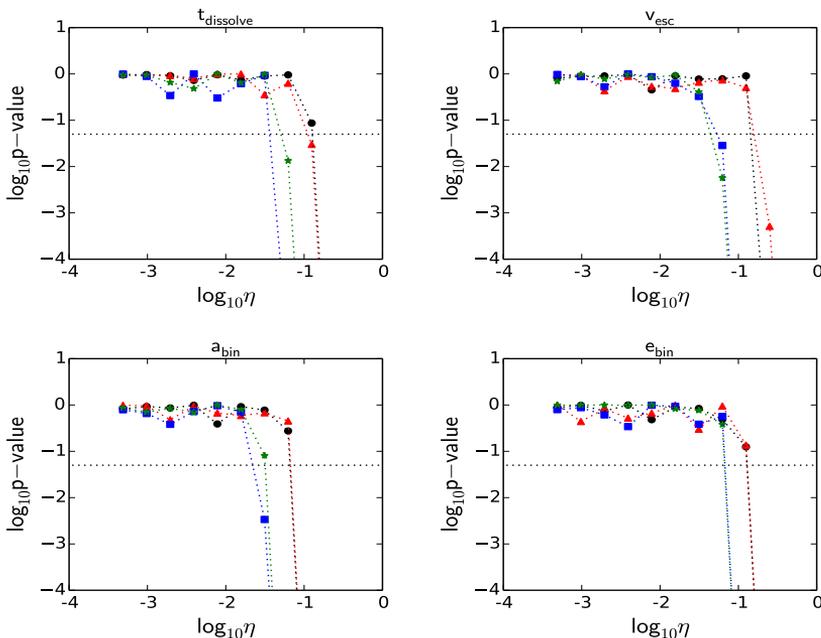


Figure 4.2: Two-sample K–S tests on distributions obtained by Hermite and Brutus. We compare distributions of dissolution time (top left), escaper speed (top right), binary semimajor axis (bottom left) and binary eccentricity (bottom right). Two-sample K–S tests are performed and the p -value is plotted versus Hermite time-step parameter η . The dashed line represents the 5% significance level. For $\eta < 2^{-5}$, the distributions are not significantly different. (The different curves represent the different data sets similar as in Fig. 4.1)

4.2.2 Brutus versus Hermite: Global Comparison

A quantitative way to compare global distributions is by performing two-sample Kolmogorov–Smirnov tests (K–S tests) (??). The K–S test gives the likelihood that two samples are drawn from the same distribution, quantified by the value called p . When the p -value is below five percent, the distributions are considered to be significantly different.

In Fig. 4.2 we plot the p -value obtained by comparing the Brutus distribution with the Hermite distribution versus time-step parameter η used for Hermite. In the panel showing the data for the binary semimajor axis, the distributions of the cold systems become significantly different for $\eta > 2^{-6}$. The distributions from the initially virialised systems start to differ for $\eta > 2^{-4}$. The cold systems are harder to model accurately, because of the close encounters that occur shortly after

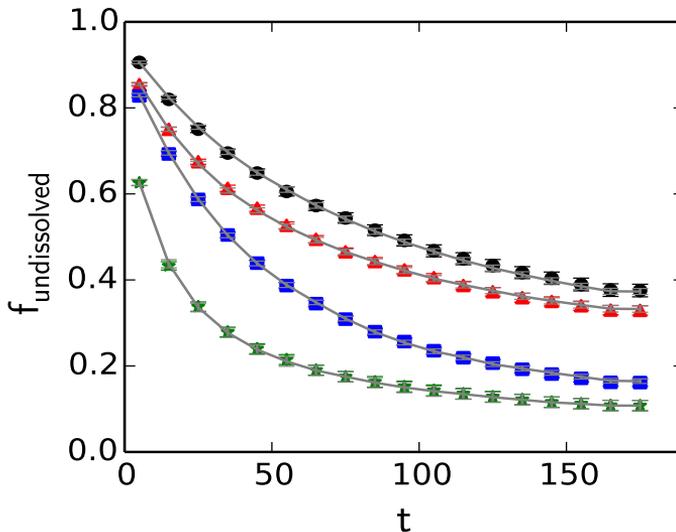


Figure 4.3: Lifetime of triple systems. We plot the fraction of triple systems that have not dissolved yet into a permanent binary and escaping single star configuration, as a function of simulation time (in units of crossing time). The curves through the data points represent the interpolated Hermite results with a time-step parameter $\eta = 2^{-5}$. (The different curves represent the different data sets similar as in Fig. 4.1)

the start. The reason the distributions start to become significantly different at large time-steps is because at these large time-steps most simulations violate energy conservation by $|\Delta E/E| > 0.1$. When this occurs, solutions might reach regions in $6N$ -dimensional phase-space, which theoretically are forbidden. The distribution then becomes biased by these outlier solutions.

4.2.3 Lifetime of Triple Systems

In Fig. 4.3, we present the fraction of triple systems which are undissolved, i.e. still interacting, as a function of time. The results by Brutus are represented by the data points: equal-mass Plummer (black bullets), Plummer with different masses (red triangles), equal-mass cold Plummer (blue squares) and cold Plummer with different masses (green stars). The results by Hermite for a time-step parameter $\eta = 2^{-5}$ are represented by the curves appearing to go through the data points.

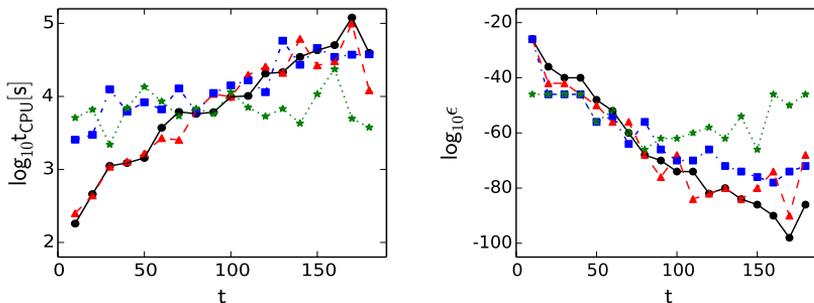


Figure 4.4: CPU time and precision as a function of time for Brutus. On the left, we plot the CPU time of the simulation which took the longest, as a function of dissolution time. On the right, we plot the Bulirsch–Stoer tolerance of the simulation which needed the highest precision, as a function of dissolution time. The different curves represent the four sets of initial conditions as in the previous plots.

The initially cold systems dissolve faster than the initially virialised systems. This is somewhat expected due to the close triple encounter resulting from the initial cold collapse: the rate of energy exchange can be very high for these encounters (?). After ~ 180 crossing times, about 40% of the systems which started with an equal-mass Plummer initial configuration, are undissolved, compared to about 10% for the cold Plummer with different masses. Systems which include stars with different masses dissolve faster than their equal mass counterparts. Energy equipartition tends to cause the lightest particle to quickly reach the escape velocity.

In Fig. 4.3, the grey curves through the data points represent the interpolated Hermite results. Even though Hermite and Brutus use different algorithms and precisions to solve the equations of motion, we find that the lifetime of an unstable triple is statistically indistinguishable between converged Brutus and non-converged Hermite solutions (but see also Sec. 4.3.3).

In Fig. 4.4, we plot the maximum CPU time and minimum BS tolerance, both as a function of dissolution time. This is shown for the Brutus simulations, for the four different initial conditions. The longer it takes for a system to dissolve, the longer the CPU time and the higher the precision needed to reach a converged solution. To reach ~ 180 crossing times, there are systems which require a BS tolerance of the order 10^{-100} , with the final converged run taking of the order a few days. The average CPU time as a function of time is about an order of magnitude smaller than the maximum CPU time. The average BS

tolerance ranges from $\sim 10^{-20}$ to 10^{-30} . For systems which dissolve within 100 crossing times, Brutus is on average about a factor 120 slower than Hermite.

We were able to obtain a complete ensemble of systems dissolving within ~ 180 crossing times. Simulations which take longer than this are not taken into account in this experiment. The fraction of long-lived systems as obtained by Hermite and Brutus are consistent. For our purpose of comparing results from conventional integrators with the converged solution, integrating up to ~ 180 crossing times is sufficient, in the sense that there is enough time for conventional solutions to diverge from the true solution (see Sec. 4.2.4). Including the long-lived triple systems may however influence the statistical distributions and biases on the long term.

4.2.4 Brutus versus Hermite: Individual Comparison

For the individual comparison, we take a certain initial realisation and compare the solutions of Hermite and Brutus. In Fig. 4.5 we show scatter plots of the Hermite solution (with time-step parameter $\eta = 2^{-5}$) versus the converged Brutus solution for the equal-mass Plummer data set.

Data points on the diagonal represent accurate solutions, whereas the scatter around it represents inaccurate Hermite solutions. The diagonal is present in each panel and extends throughout the range of possible outcomes. The width of the diagonal is very narrow. When the normalized phase-space distance between the Hermite and Brutus solution $\delta < 10^{-1}$, then the coordinates are accurate enough to produce derived quantities accurate to at least one decimal place and Hermite and Brutus will give similar results. Once $\delta > 10^{-1}$, the solution has diverged to a different trajectory in phase-space leading to a different outcome. This outcome could in principle be any of the possible outcomes as can be derived from the amount of scatter in the Hermite solutions at a fixed Brutus solution.

In the scatter plot of the dissolution time, we observe that for small times ($t < 10$), Hermite and Brutus agree on the solution in the sense that the data points lie on the diagonal. Systems which dissolve after a short time don't have sufficient time to accumulate enough error to diverge to another trajectory in phase-space. Once however this level of divergence is reached, the scatter immediately covers the entire, available outcome space. This randomisation is also observed in the other panels.

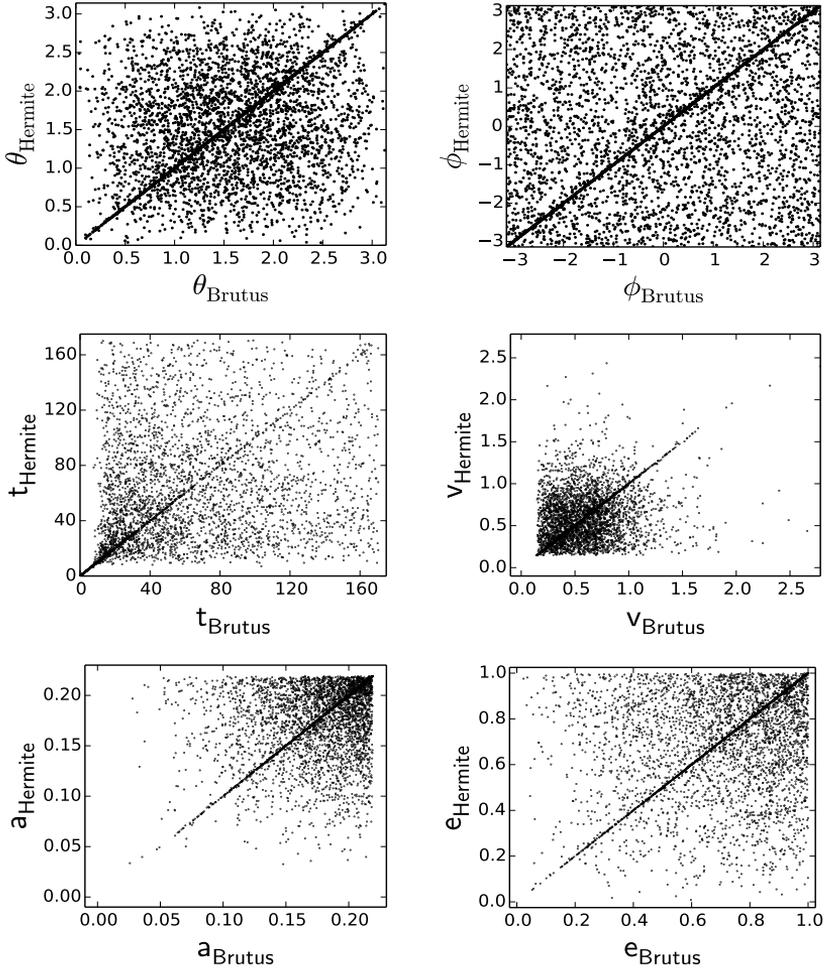


Figure 4.5: Direct comparison of Brutus and Hermite results per individual simulation. The results are shown only for the $N = 3$ equal mass Plummer data set and for a Hermite time-step parameter $\eta = 2^{-5}$. Each dot in a panel represents a different initial realisation. The value on the ordinate is the value obtained using Hermite and the value on the abscissa the value obtained by Brutus. We compare the direction of the escaper: polar angle (top left) and azimuthal angle (top right), (with respect to the plane of the binary and pericentre direction), dissolution time (middle left), escaper velocity (middle right), binary semimajor axis (bottom left) and binary eccentricity (bottom right). The diagonal represents accurate Hermite solutions. The scatter around it represents solutions where Hermite and Brutus have diverged.

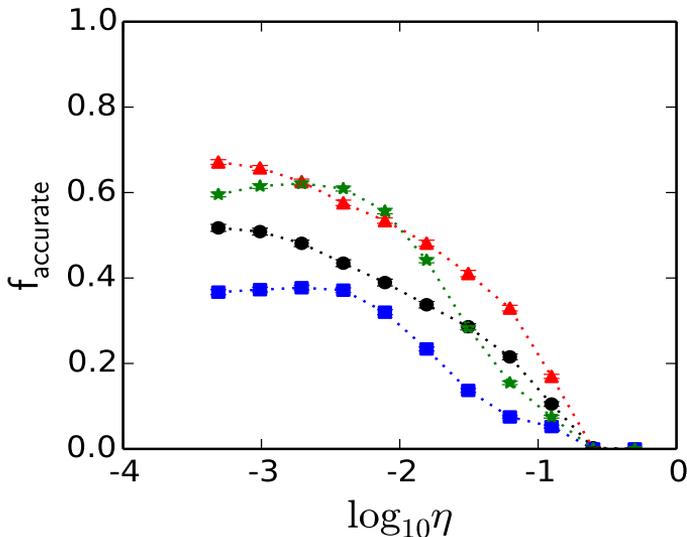


Figure 4.6: The fraction of accurate Hermite simulations as a function of Hermite time-step parameter η . The different curves represent the different data sets: equal mass Plummer (bullets), Plummer with different masses (triangles), equal mass cold Plummer (squares) and cold Plummer with different masses (stars). As η decreases, the accurate fraction increases. However, for $\eta < 2^{-7}$, the fraction starts to saturate, more so for the cold data sets. At this point the effect of round-off error becomes important.

The Fraction of Accurate Solutions

In Fig. 4.6 we estimate the fraction of data points on the diagonal as a function of the Hermite time-step parameter, η . We only include the data points for which the normalized phase-space distance $\delta < 10^{-1}$. For the largest time-step parameters used ($\eta > 10^{-1}$) the fraction on the diagonal, or the accurate fraction, varies from zero to about 0.2. By reducing the time-step parameter, the accurate fraction increases until it saturates at about 0.4 to 0.7 depending on the initial conditions. Even though by reducing η , the discretisation error decreases, the number of integration steps increases, which then increases the round-off error. For the data sets with zero angular momentum, the maximum accurate fraction is obtained for $\eta \sim 2^{-9}$. For the initially virialised systems this seems to occur between $\eta \sim 10^{-3} - 10^{-4}$, although the actual saturation point is not visible yet. This dependence on angular momentum is due to the initial cold collapse and subsequent close encounters, which increases the round-off error.

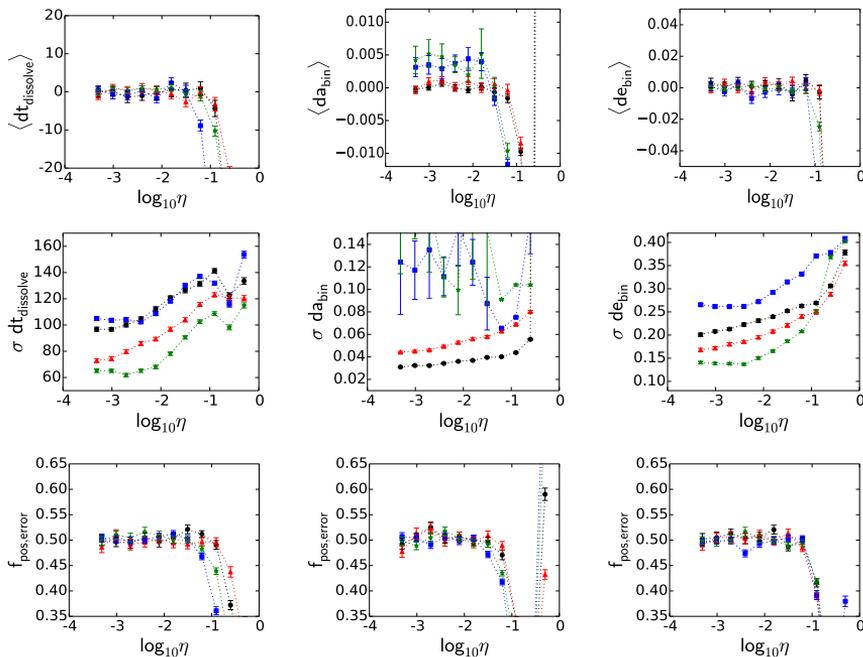


Figure 4.7: Statistics on the error distribution of Hermite results. We present the average error (top row), the standard deviation of the error distribution (middle row) and the fraction of errors which are positive (bottom row). The errors are given for the dissolution time (left column), binary semimajor axis (middle column) and eccentricity (right column). The different curves represent the different data sets similar as in Fig. 4.6.

The Error Distribution

In Fig. 4.7 we present statistics on the distribution of the errors, i.e. $S_{Hermite} - S_{Brutus}$, with S a statistic. For the dissolution time and the eccentricity, the average error converges to zero for $\eta < 10^{-1}$. For larger time-steps, simulations which grossly violate energy conservation ($|\Delta E/E| > 0.1$) cause biases in the average error. For the binary semimajor axis however, the data representing the cold collapse simulations also seem to be systematically biased for small time-steps, in the sense that Hermite makes fewer tight binaries.

The width of the error distributions converge to a non-zero value. This can be understood because with decreasing time-step, round-off errors will become more important so that the standard deviation of the errors will never reach zero. For the dissolution time, the width of the error distribution for the smallest time-step parameter adopted, varies from 60 to 100 crossing times. For the eccentricities the width

is on average ~ 0.2 . For the semimajor axis the width approaches ~ 0.05 (in Hénon units). In the case of the semimajor axis, the data representing the cold collapse simulations behave differently, because the width is much larger than the width for the data representing the initially virialised systems.

If we regard the results given by Brutus and Hermite as random variables drawn from the same distribution, then we can write the variance in a certain statistic, in this example the eccentricity, as:

$$\langle (e_H - e_B)^2 \rangle = \langle e_H^2 \rangle + \langle e_B^2 \rangle - 2\langle e_H \rangle \langle e_B \rangle. \quad (4.5)$$

Here e stands for eccentricity and the subscripts for Brutus and Hermite. For a thermal eccentricity distribution (Eq. 4.1), we obtain a standard deviation of $1/3$. However, this only applies to inaccurate Hermite results, which had enough time to diverge through outcome space. If we multiply the theoretical standard deviation calculated above by the inaccurate fraction, we obtain a range in the standard deviation from 0.17 to 0.27, as η ranges from the most precise value to $\eta = 10^{-1}$.

Symmetry of the Error Distribution

To measure the symmetry of the error distribution, we count the fraction of positive errors (Fig. 4.7, bottom panels). Again for an $\eta < 10^{-1}$, this fraction converges to 0.5. A more detailed comparison is given in Fig. 4.8, where we compare distribution functions of positive and negative errors. In Sec. 2.2.3, we mentioned that in our experiment we define the Brutus solution to be converged when at least 3 decimal places of every coordinate have converged. To investigate the symmetry up to higher precision, we repeated a subset of 1000 simulations. We did this only for the initial conditions with equal-mass stars picked randomly from a virialised Plummer distribution and this time we obtain solutions converged up to the first 15 decimal places.

We observe that the majority of errors are larger than $\sim 10^{-3}$ and within the statistical error, the positive and negative errors have a similar distribution. For the smallest errors however, we observe an asymmetry in the sense that there are more negative, small errors. The magnitude of the error where this excess occurs is determined by the precision of the integration. For the smallest η , the excess is below double-precision and thus not observable anymore (see Sec. 4.3.2 for more explanation).

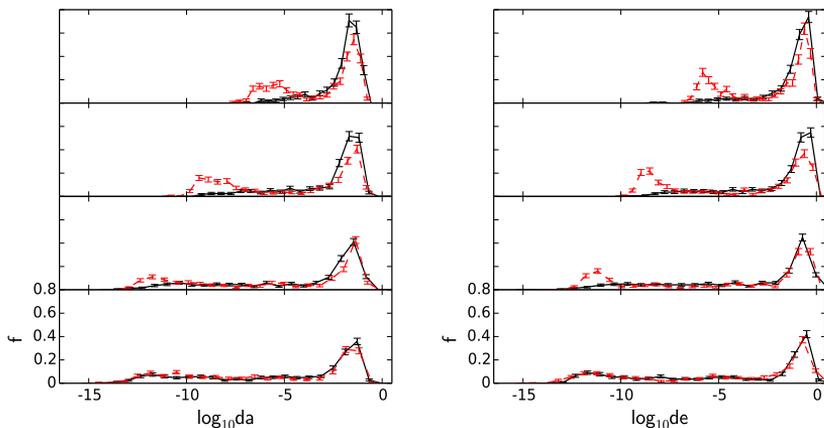


Figure 4.8: Symmetry of the error distributions. We show distributions of the errors in semimajor axis (left column) and eccentricity (right column) of the binaries formed in the equal-mass Plummer data set. This is shown separately for the positive errors (solid line) and negative errors (dashed line), to investigate the symmetry of the error distribution. From the panels at the top to the bottom, the time-step parameter for Hermite varies as 2^{-5} , 2^{-7} , 2^{-9} and 2^{-11} . An asymmetry can be observed at the smallest errors.

4.2.5 Escaper Identity

In this section we compare the solutions obtained with Hermite and Brutus individually, by looking at which star eventually becomes the escaper and which form the binary. We define preservation if the Hermite and the Brutus solution both have the same star as the escaper. We define it as exchange if the escaping star is different. A further distinction can be made in the preservation category, if the Hermite simulation is also accurate. We can typify each Hermite simulation as follows:

- **Accurate:** The coordinates are accurate, up to at least two digits.
- **Preservation:** The coordinates are inaccurate, but same star escapes.
- **Exchange:** Different star escapes.

In Fig. 4.9 we present the fraction of each category as a function of time. As expected, systems which dissolve quickly, hardly have time to develop errors and are categorized as accurate simulations. In time

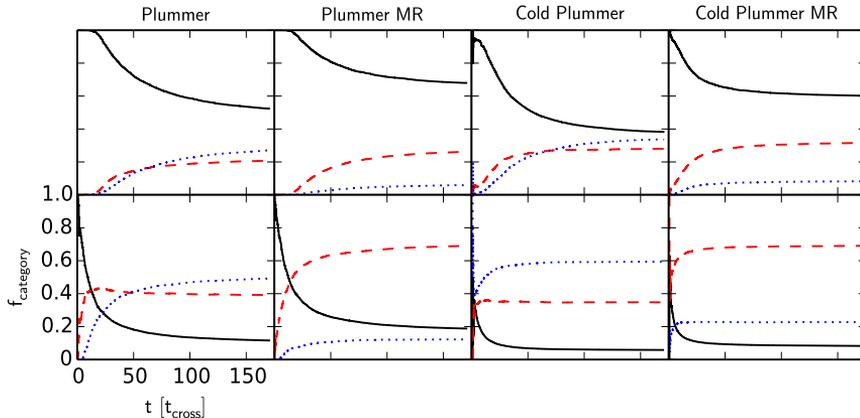


Figure 4.9: The evolution of the relative fraction of categories. The different curves represent the different categories: accurate (solid curves), preservation (dashed curves) and exchange (dotted curves). These three categories are defined in the text. From left to right, the data are from the Plummer, Plummer with different masses, cold Plummer and cold Plummer with different masses data sets. In the top panels we show the results for a Hermite time-step parameter $\eta = 2^{-11}$ and in the bottom for $\eta = 2^{-3}$.

however, because errors grow exponentially, the solutions become inaccurate. The fractions of preservation and exchange start to grow. For a small time-step parameter ($\eta = 2^{-11}$, top row in Fig. 4.9), this growth starts after ~ 20 crossing times for the initially virialised systems. For the initially cold systems, the inaccurate fractions already start to grow after a single crossing time.

The cold collapse with equal-mass stars is the hardest problem to integrate as the accurate fraction is of comparable magnitude as the preservation and exchange fractions. The accurate fraction generally remains dominant, with a final fraction varying from about 0.4 for the equal-mass cold Plummer to about 0.7 for the Plummer with different masses. For the lesser precision ($\eta = 2^{-3}$, bottom row in the figure), the accurate fractions decrease to below 0.2.

In the panels in Fig. 4.9, which include the data for the systems with different masses, preservation is more common than exchange. This can be understood, because due to energy equipartition, the lightest particle will be more likely to escape and therefore the identity is more often correct than in the equal mass case. For the equal mass case, the fraction of preservation and exchange is comparable, except in the case of the equal-mass cold Plummer with the low precision ($\eta = 2^{-3}$, the bottom row). If we regard the identity of the escaping star to be

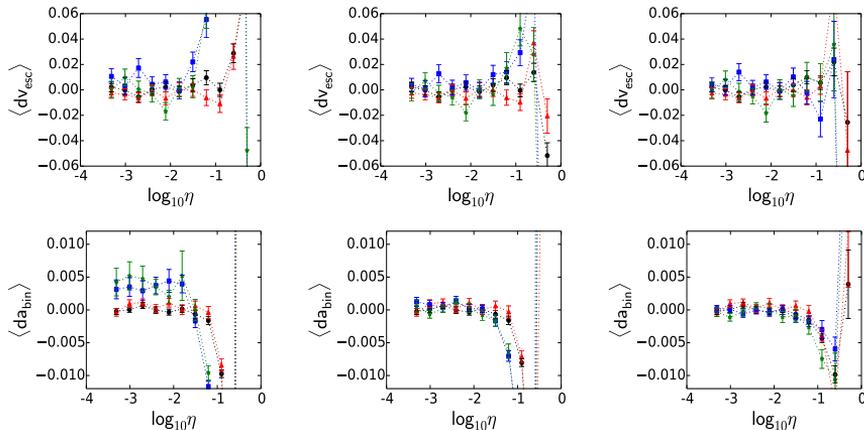


Figure 4.10: The effect of cuts in final relative energy conservation. We plot the average error in the velocity of the escaping star (top row) and the error in the binary semimajor axis (bottom row) as a function of Hermite time-step parameter η (with the same relation between the curves and the data sets as in Fig. 4.6). The three columns differ in the maximum allowed level of relative energy conservation. In the left column we show the results for the total ensemble of solutions, in the middle column for a maximum level of unity and in the right column for 10^{-1} . The bias in the left column for the binary semimajor axis is caused by solutions which grossly violate energy conservation. Note that this only happens for the cold collapse simulations. When these outliers are taken out of the ensemble, the bias vanishes.

completely random once the solution has become inaccurate, we would expect the fraction of exchange to be twice the fraction of preservation. This is roughly what we observe in the equal mass cold collapse case with low precision. Because of the low precision and the initial close encounter, solutions will diverge very quickly. In the panel with the higher precision this trend is not observed because the solutions are less randomised. The preservation category includes solutions which slightly differ from the converged solution only in the escape angle of the escaper. Also the long-lived triples are not taken into account here, which will alter these fractions.

4.3 DISCUSSION

4.3.1 Energy Conservation

In every ensemble of Hermite solutions there are some that grossly violate conservation of energy $|\Delta E/E| > 0.1$. This deformation of the energy hyper-surface in phase-space can allow solutions to reach

parts of phase-space which are theoretically forbidden. This affects the global statistical distributions. In Fig. 4.10, we replot the average error in the binary semimajor axis as a function of the time-step parameter. We produce similar diagrams as presented in Fig. 4.7, but this time we introduce a maximum allowed error in the energy. If we filter out simulations with a relative energy conservation $|\Delta E/E| > 1$, or $|\Delta E/E| > 0.1$, we observe that the bias in the average error of the semimajor axis of the binaries vanishes. We conclude that this bias is caused by a few simulations which grossly violate energy conservation. A similar bias in the velocity of the escaping star is less pronounced.

Time-reversible, symplectic integrators should in principle conserve energy to a better level than non-symplectic integrators, since there is no drift present in the energy error. Therefore, by using a symplectic integrator, the number of simulations with large energy error could be reduced. Using a Leapfrog integrator with constant time-steps, we tested this assumption and we find that for resonant 3-body interactions, it is challenging to obtain accurate solutions. The main reason is that, contrary to regular systems like, for example, the solar system, resonant 3-body interactions often include very close encounters, which need a very small time-step size to be resolved accurately. This is especially the case for the initially cold systems. Adopting such a small time-step size for the whole simulation, will increase the wall-clock time to that of Brutus or beyond.

4.3.2 Asymmetry at Small Errors

In Sec. 4.2.4, we discussed an asymmetry at small errors. In Fig. 4.11, we present similar diagrams as in Fig. 4.8 for the positive and negative errors. This time we add the errors in the total energy and angular momentum of the system and the error in the velocity of the escaper.

We also vary the integration method because different methods produce different (biased) error distributions in energy and angular momentum. We use a standard Leapfrog integrator, a standard Hermite integrator and a Hermite integrator which uses the $P(EC)^n$ method (we adopted $n=3$) (?). This last method adds an iterative procedure to the algorithm to improve the predictions and corrections, which improves the time-symmetry. For each method we implement a shared, adaptive time-step criterion as in Eq. 2.1, with a time-step parameter $\eta = 2^{-7}$. As a consequence they will not be time-symmetric nor symplectic.

We first look at the error distributions in the total energy and angular momentum. We observe that none of them are symmetric, in the

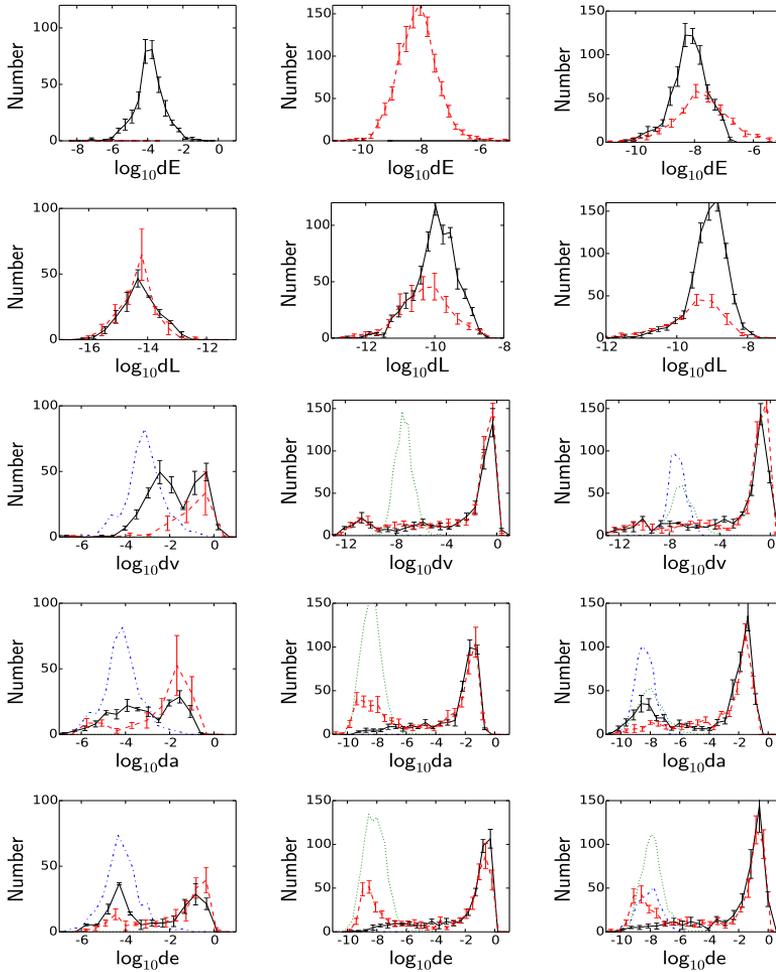


Figure 4.11: Explanation of the asymmetry at small errors. We show distributions of the positive (solid curves) and negative (dashed curves) errors in the total energy (top row), total angular momentum (second row), escaper velocity (third row), binary semimajor axis (fourth row) and eccentricity (bottom row). This is shown for different algorithms: Leapfrog (left column), standard Hermite (middle column) and Hermite with $P(EC)^n$ method (right column, $n = 3$). Each method implements a shared, adaptive time-step criterion according to Eq. 2.1, with a time-step parameter $\eta = 2^{-7}$. Each of these three integrators has a different asymmetry in the conservation of energy and angular momentum. By propagating these asymmetric errors as a small perturbation to the converged solution, we can estimate the resulting asymmetry in the derived quantities. These estimated error distributions are also given separately for the positive (dot-dash, light curves) and negative (dotted, light curves) errors. We observe that the estimated error distributions are located at the asymmetry in the empirical error distributions. The asymmetry at small errors is caused by a bias in the integrator.

sense that the positive and negative errors have identical distributions, except for the angular momentum in the Leapfrog simulations. The Leapfrog solutions tend to gain energy, whereas the standard Hermite loses energy. The Hermite with the P(EC)ⁿ method produces both positive and negative errors in the energy, but not in a symmetric manner.

To investigate whether the bias in energy and angular momentum conservation propagates to a bias in the binary and escaper properties, we estimate what the errors should be if we regard the error in the energy and angular momentum as a small perturbation to the converged solution. For the error in the velocity of the escaper, using the derivative of the kinetic energy with respect to velocity, we obtain the following expression:

$$\delta v = \frac{1}{mv} \delta E. \quad (4.6)$$

Here m is the mass of a star, v the velocity as obtained by Brutus, δE the energy error and δv the error in the velocity due to this energy error. For the binary semimajor axis we obtain:

$$\delta a = \frac{2}{m^2} a^2 \delta E. \quad (4.7)$$

Here a is the semimajor axis from the Brutus solution. For the eccentricity we obtain:

$$\delta e = \frac{1}{\sqrt{1 + \frac{2\epsilon l^2}{\mu^2}}} \left(\frac{l^2}{\mu^2} \delta \epsilon + \frac{2\epsilon l}{\mu^2} \delta l \right). \quad (4.8)$$

Here μ is the total mass of the binary, ϵ and l the specific energy and specific angular momentum of the binary as obtained by Brutus. The error in the eccentricity δe has contributions from errors in the energy $\delta \epsilon$ and angular momentum δl .

If we compare the resulting error distributions to the actual error distributions, we find that the approximated error distribution is positioned at the asymmetry in the empirical error distribution. This is most clearly seen for the semimajor axis and eccentricity (see Fig. 4.11).

The reason why the approximated error distribution overestimates the excess, is because not all errors are solely due to an error in the energy and angular momentum. In time, the numerical solution diverges from the true solution and this error due to divergence will become

more dominant. With this in mind, we can approximate the error in a statistic as follows:

$$\delta S = \delta S_{\text{conservation}} + \delta S_{\text{divergence}}. \quad (4.9)$$

Here S is a statistic that is related to energy and/or angular momentum, $\delta S_{\text{conservation}}$ is the error due to a small perturbation in the energy and/or angular momentum and $\delta S_{\text{divergence}}$ is the error due to divergence of the solution. When the solution has not diverged appreciably yet, the first type of error will dominate and possible biases can be observed. When the second type of error dominates, we observe that the symmetry is restored to within the statistical error.

Upon inspection of the velocity data, we observe no asymmetry in the Hermite results. When we measure which fraction of the energy error is reserved for the binary and which fraction for the escaper, we find that in most cases the error propagates to the binary. For the Leapfrog however, the asymmetry is still present.

4.3.3 Preservation of the Macroscopic Properties

Valtonen et al. (?) state that the final statistical distributions forget the specific initial conditions and only depend on globally conserved quantities. This assumption makes predictions which are verified by our experiment. The results show that for a time-step parameter $\eta < 2^{-5}$, the distributions are statistically indistinguishable, even though at least half of the solutions diverged from the converged solution. If however, energy conservation is grossly violated, biases are introduced in the statistics. In our experiment, a maximum level of relative energy conservation of $|\Delta E/E| = 0.1$ was sufficient to remove the biases. This is a much milder constraint than the $|\Delta E/E| \sim 10^{-6}$ usually adopted in collisional simulations. Whether 0.1 is also sufficient for systems with more stars, should be verified experimentally. Heggie (?) for example, finds that the energy of escaping stars in higher-N systems, depends sensitively on integration accuracy. The maximum required level of energy conservation should be such that it is below the energy taken away from the cluster by the escaping stars.

The chaoticity of the 3-body problem is illustrated by the scatter diagrams in Fig. 4.5. For a certain value of a statistic obtained by Brutus, any other value in the allowed outcome space is reachable for the Hermite integrator. For example, if the converged solution gives an eccentricity for the binary of 0.6, a diverged solution can produce any eccentricity between 0 and 1. Once the solution has diverged from the true solution, it will start a random walk through or near the allowed

phase-space until the 3-body system has dissolved. We observed that this randomisation happens in such a way that the available outcome space is still completely sampled and that it preserves global statistical distributions.

In Sec. 4.2.3, we discussed that the lifetime of an unstable triple does not depend on the integrator used nor on the accuracy of that integrator. This last point should be interpreted in the sense that when more effort is put into performing simulations with higher precision, that this does not change the global statistics, even though individual solutions will change with precision (see for example the Hermite results in Fig. 2.1). If instead we continue to decrease the precision, there will be a point where biases start to appear. Urminky (?) analysed the 3-body Sitnikov problem and showed that the precision of the integration influences the average lifetime of triple systems, contrary to our results. The integration times in our experiment however, are much shorter. Obtaining a converged solution for a resonant 3-body system for longer than 200 crossing times, is still computationally challenging. Therefore any statistical difference on the long term will not be visible in our experiment.

4.4 CONCLUSION

Brutus is an N-body code that uses the Bulirsch–Stoer method to control discretisation errors, and arbitrary-precision arithmetic to control round-off errors. By using the method of convergence, where we systematically vary the Bulirsch–Stoer tolerance parameter and the word-length, we can obtain a solution for a particular N-body problem, for which the first p digits in the mantissa are independent of the time-step size and word-length. We call this solution converged to p decimal places.

Obtaining the converged solution is computationally very expensive, mainly because of the exponential divergence of the solution. In some cases, Bulirsch–Stoer tolerances of 10^{-100} are needed to reach convergence. We estimate that the time for simulating a star cluster up to core collapse, until convergence, scales approximately exponentially with the number of stars. Simulations with 256 stars however, may be performed within a year of computing time.

The motivation to obtain expensive, converged solutions is to test the assumption that the statistics of an ensemble of approximate solutions, are indistinguishable from the statistics of an ensemble of true solutions. To put this assumption to the test, we have investigated the statistics on the breakup of 3-body systems. In our experiment, a

bound triple system will eventually dissolve into a binary and an escaping star. Solutions to every initial realisation were obtained using the standard Hermite integrator and using Brutus.

For systems with a long lifetime it is challenging to obtain the converged solution. Due to repeated ejections and resonances, many accurate digits will be lost and so a very small Bulirsch–Stoer tolerance is required. Therefore, we have set an integration limit at ~ 180 crossing times. For equal-mass, virialised systems, $\sim 40\%$ of the random initial realisations were not dissolved by this time. For the initially cold systems with different masses this was $\sim 10\%$. Hermite and Brutus are consistent on the average lifetime of an unstable triple system. However, possible differences on the long term are not visible in this experiment.

When we compare the results on an individual basis, we find that on average about half of the Hermite solutions give accurate results, i.e. at most a 1% relative difference compared to Brutus. For the inaccurate results, the error distribution becomes unbiased and symmetric for a time-step parameter $\eta \leq 2^{-5}$ and implementing a maximum level of relative energy conservation of $|\Delta E/E| < 0.1$.

Once the conventional solution has diverged from the converged solution, it will start a random walk through or near the allowed region in phase space. such that any allowed outcome of a statistic is reachable. This randomisation process completely samples the available outcome space of a statistic and it also preserves the global statistical distributions.

Kolmogorov–Smirnov tests were performed to compare the global distributions produced by Hermite and Brutus. No significant differences were detected when using the criteria mentioned above for the time-step parameter η and relative energy conservation. This research for the 3-body problem supports the assumption that results from conventional N-body simulations are valid in a statistical sense. We observed however that a bias is introduced for the smallest errors, if the algorithm used to solve the equations of motion, is biased in the conservation of energy and angular momentum. In this research however, this bias did not have an appreciable effect. It is important to see whether this remains true for statistics of higher-N systems or systems with a dominant mass. An example of a higher-N system where precision might play a role is a young star cluster (without gas) going through the process of cold collapse (?). At the moment of deepest collapse, a fraction of stars will obtain large accelerations, so that a small error in the acceleration can cause large errors in the position

and velocity. The rate of divergence can increase up to about 5 digits per Hénon time unit for 128 particles and it increases with N .

