Gargantuan chaotic gravitational three-body systems and their irreversibility to the Planck length

T. C. N. Boekholt [●],[1,2]★ S. F. Portegies Zwart[3]★ and M. Valtonen[4,5]★

[1] *CFisUC, Department of Physics, University of Coimbra, 3004-516 Coimbra, Portugal*
[2] *Instituto de Telecomunicações, Campus Universitário de Santiago, 3810-193 Aveiro, Portugal*
[3] *Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*
[4] *Finnish Centre for Astronomy with ESO, University of Turku, FI-20014 Turku, Finland*
[5] *Department of Physics and Astronomy, University of Turku, FI-20014 Turku, Finland*

ABSTRACT

Chaos is present in most stellar dynamical systems and manifests itself through the exponential growth of small perturbations. Exponential divergence drives time irreversibility and increases the entropy in the system. A numerical consequence is that integrations of the *N*-body problem unavoidably magnify truncation and rounding errors to macroscopic scales. Hitherto, a quantitative relation between chaos in stellar dynamical systems and the level of irreversibility remained undetermined. In this work, we study chaotic three-body systems in free fall initially using the accurate and precise *N*-body code `Brutus`, which goes beyond standard double-precision arithmetic. We demonstrate that the fraction of irreversible solutions decreases as a power law with numerical accuracy. This can be derived from the distribution of amplification factors of small initial perturbations. Applying this result to systems consisting of three massive black holes with zero total angular momentum, we conclude that up to 5 per cent of such triples would require an accuracy of smaller than the Planck length in order to produce a time-reversible solution, thus rendering them fundamentally unpredictable.

Key words: methods: numerical – stars: black holes – stars: kinematics and dynamics.

1 INTRODUCTION

Chaos is an inherent property of most dynamical systems in the universe, ranging from small bodies in the Solar system (e.g. Wisdom, Peale & Mignard 1984; Boekholt et al. 2016; Correia 2018), small stellar systems (e.g. Hut & Bahcall 1983; Portegies Zwart & Boekholt 2014; Boekholt & Portegies Zwart 2015; Portegies Zwart & Boekholt 2018; Leigh & Wegsman 2018; Stone & Leigh 2019), star clusters (e.g. Miller 1964; Goodman, Heggie & Hut 1993), and galaxies (e.g. Valluri & Merritt 2000). The main signature of chaos is the exponential sensitivity to small changes in the initial conditions, which is quantified by the e-folding time scale within some finite-time interval, i.e. the finite-time Lyapunov time-scale (Heggie 1991).

The exponential sensitivity in *N*-body systems has both physical and numerical consequences. From a physical point of view, the rate of growth of perturbations determines the stability of a system. Such studies are well known for the Solar system, whose Lyapunov time-scale is about 5 Myr (Laskar 1989). Due to observational uncertainties in the orbital elements of the planets, we can only

predict the future evolution of the Solar system for a few million years, warranting a statistical study of its stability over Gyr time-scales (Laskar 1989; Sussman & Wisdom 1992; Ito & Tanikawa 2002; Hayes 2007). Hence, in contrast to regular and stable systems, high precision in the initial conditions is crucial for accurate modelling of chaotic systems.

In few-body stellar dynamical systems, it was first shown by Miller (1964) that two nearby trajectories in phase space tend to diverge exponentially. 'The divergence of the two trajectories from each other is a manifestation of the macroscopic irreversibility' and 'the rate of divergence yields information on the rate of entropy production' (Miller 1964). This rate is linear with time, because the rate of divergence is exponential, and the entropy proportional to the logarithm of the increasing phase space volume. The presence of chaos and macroscopic irreversibility can be related to the arrow of time, in the sense that it points in the direction of increasing entropy. Thus, the arrow of time points in the direction of diverging trajectories rather than converging ones. This leads to the idea that in a world consisting of only three bodies, there would already be a definite direction for the arrow of time (Lehto et al. 2008).

From a numerical point of view, errors in *N*-body simulations also act as small perturbations to the system, and their subsequent exponential magnification causes the solution to eventually diverge on to a completely different trajectory after only a few Lyapunov

★ E-mail: tjardaboekholt@gmail.com (TCNB); spz@strw.leidenuniv.nl (SFPZ); mvaltonen2001@yahoo.com (MV)

time-scales. The calculated system is not causally related to its initial condition anymore, in the same way a physical system is (Miller 1964). This raises suspicions on the reliability of *N*-body simulations. The common assumption is that approximate results from *N*-body simulations are valid in a statistical sense (Goodman et al. 1993). Empirically this has been shown to be the case for certain specific *N*-body systems (e.g. Portegies Zwart & Boekholt 2014; Boekholt & Portegies Zwart 2015; Portegies Zwart & Boekholt 2018), but a sound theoretical basis is still missing. Our trust in *N*-body simulations can potentially be made more robust if it can be shown that the 'numerically diverged trajectory' still has some physical connection to the initial condition space under consideration, but to a slightly different initial realization than the one used to start the simulation. In other words, we are still calculating physical trajectories, but to randomize initial conditions (Dejonghe & Hut 1986). This process can be demonstrated if it can be shown that approximate solutions have 'shadow orbits' (Quinlan & Tremaine 1992; Urminsky 2010). Such orbits remain close to the approximate trajectory for a time much longer than the Lyapunov time, but which have a physical connection to the initial condition space of the *N*-body problem under consideration.

Alternatively, one can apply brute-force computing power to try and reduce the magnitude of numerical errors. A robust way to test the accuracy of a specific *N*-body simulation is by performing a reversibility test. Since Newton's equations of motion are time reversible, a forward integration followed by a backward integration of the same time should recover the initial realization of the system (albeit with a sign difference in the velocities). The outcome of a reversibility test is thus exactly known. In practice, reversibility in simulations of chaotic systems is very difficult to achieve due to (1) exponential growth of perturbations due to chaos and (2) irreversible numerical errors. Time reversibility can be forced by using integer arithmetic, but this does not guarantee the solution is also accurate.

Recently, Portegies Zwart & Boekholt (2018) obtained a reversible solution to the Pythagorean problem (Burrau 1913; Szebehely & Peters 1967; Aarseth et al. 1994). This is a classic example of a three-body system in free fall initially exhibiting a prolonged chaotic triple interaction, and an eventual break-up into a permanent and unbound binary-single pair. They applied the Brutus *N*-body code and the method of convergence in which the accuracy and precision of the integration is systematically increased until convergence of the solution to the first few decimal places (Boekholt & Portegies Zwart 2015). Although Brutus is not formally time reversible, they manage to retrieve the initial condition to the first 10 decimal places in each coordinate of each body in the final snapshot. Whereas the forward integration is subject to exponential divergence, the backward integration is subject to exponential convergence to the initial size of the perturbation over nine orders of magnitude. This behaviour was called *definitive reversibility* by Portegies Zwart & Boekholt (2018).

We extend the initial condition space from the Pythagorean problem to the homology map of Agekyan & Anosova (1967) and Agekyan & Anosova (1968) [see also Anosova & Nebukin (1991), Anosova (1991), Anosova, Orlov & Aarseth (1994), Tanikawa, Umehara & Abe (1995), Martynova & Orlov (2014), and Orlov, Titov & Shombina (2016)]. For a definition and visualization of this map, see fig. 1 of Lehto et al. (2008). The Agekyan–Anosova map consists of every equal-mass triple system configuration with zero initial velocities (after potentially rescaling or rotating the system). The initial conditions thus specify three-body systems in free fall initially with varying initial mutual separations. Such trajectories may closely approach a triple collision, which are notoriously

challenging to solve, even using regularization techniques. After such close triple approaches, the triple can break up, or alternatively continue its evolution in a prolonged, chaotic, and resonant (Hut & Bahcall 1983) interaction. Reversibility tests for the 'homology map' have been performed by Lehto et al. (2008). They find that about half of the systems are reversible and the other half remains irreversible, regardless how much computer time you spend on the problem. They conclude that half of the three-body systems are so chaotic that they cannot be solved numerically (Valtonen & Karttunen 2006). This is also corroborated by results from Dejonghe & Hut (1986), who measure amplification factors of small initial perturbations of up to 10^{150}. However, the fraction of irreversible solutions can potentially be reduced if one goes beyond standard double-precision, i.e. using arbitrary-precision arithmetic.

2 RESULTS

We perform reversibility tests for the Agekyan–Anosova map (Agekyan & Anosova 1967, 1968; Lehto et al. 2008), using the arbitrary-precision *N*-body code Brutus (Boekholt & Portegies Zwart 2015). We control the accuracy by varying the Bulirsch–Stoer tolerance (Bulirsch & Stoer 1964), ϵ, and fix the arbitrary-precision word-length to 1024 bits (about 300 decimal places). More detail on the methods is given in Appendix A.

The main idea of our experiment is the following. Each triple system has a certain escape time, which is the time it takes for the triple to break-up into a permanent and unbound binary-single configuration. Given a numerical accuracy, ϵ, there is also a tracking time, which is the time that the numerical solution is still close to the physical trajectory that is connected to the initial condition. If the tracking time is shorter than the escape time, then the numerical solution has diverged from the physical solution, and as a consequence, it has become time irreversible. Only the systems with the smallest amplifications factors will pass the reversibility test. However, by systematically increasing the numerical accuracy (decreasing ϵ), we aim to increase the tracking time of each system. An increasing fraction of systems will obtain a tracking time exceeding its escape time, thus gradually decreasing the fraction of irreversible solutions.

In Fig. 1, we present our low-resolution Agekyan–Anosova map, where we plot the lifetime of the triple system as a function of initial condition. The triple lifetime is defined as the duration of the triple interaction until a permanent and unbound binary-single configuration is reached. When comparing the least accurate ($\epsilon = 10^{-6}$) and the most accurate ($\epsilon = 10^{-70}$) maps, we observe that there are 'microscopic' differences. However, in a 'macroscopic' sense, the maps look similar. This is confirmed by performing a Kolmogorov–Smirnoff test, which gives a *p*-value of 0.72. This implies that we cannot reject the hypothesis that the two distributions are statistically indistinguishable. Therefore, it seems that for the Agekyan–Anosova map, approximate computations are nevertheless reliable in a statistical sense (Goodman et al. 1993). This is another example of the concept of '*nagh-Hoch*' (Portegies Zwart & Boekholt 2018), which refers to the 'similar appearance' of statistical distributions, which are obtained with different numerical precisions.[1]

In Fig. 2, we plot the fraction of irreversible solutions as a function

---

[1]The term 'Nagh Hoch' was first defined by Portegies Zwart & Boekholt (2018) and comes from the Klingon dictionary meaning 'similar appearance' or 'set in stone'.
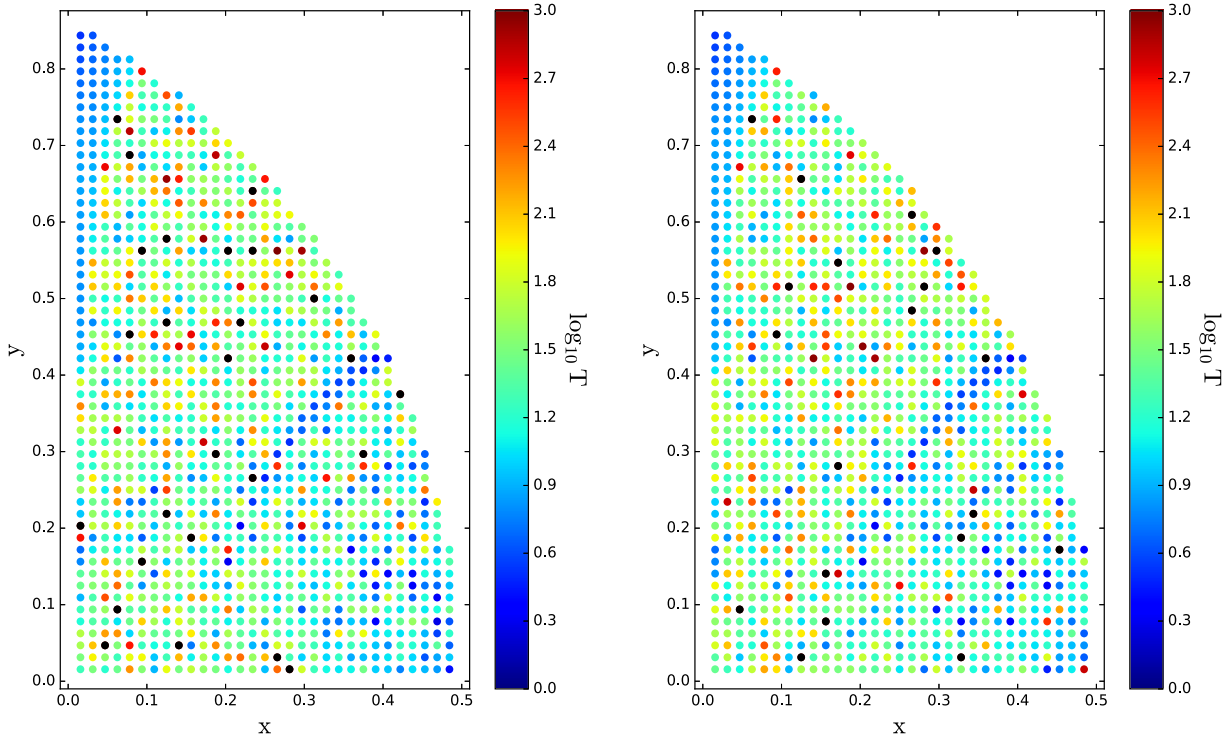
Figure 1. Duration of the triple interaction as a function of initial condition in the Agekyan–Anosova map (a higher resolution map can be found in fig. 1 of Lehto et al. 2008). We show the 'lifetime map' for a numerical accuracy of $\epsilon = 10^{-6}$ (left) and $\epsilon = 10^{-70}$ (right). Black dots represent systems that live for longer than a 1000 time units. Even though the same initial condition in the two maps can give very different lifetimes, the two maps are statistically indistinguishable according to a Kolmogorov–Smirnoff test.
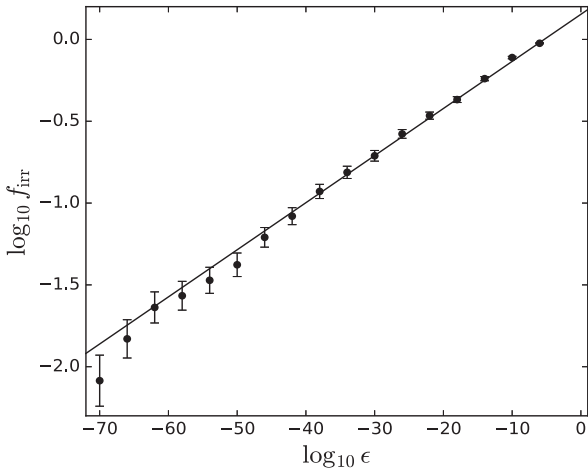

Figure 2. Fraction of irreversible solutions, f_{irr}, as a function of numerical accuracy, ϵ. The power-law fit gives $\log_{10} f_{irr} = \alpha \log_{10} \epsilon + \beta$, with $\alpha = 0.029 \pm 0.001$ and $\beta = 0.15 \pm 0.04$.
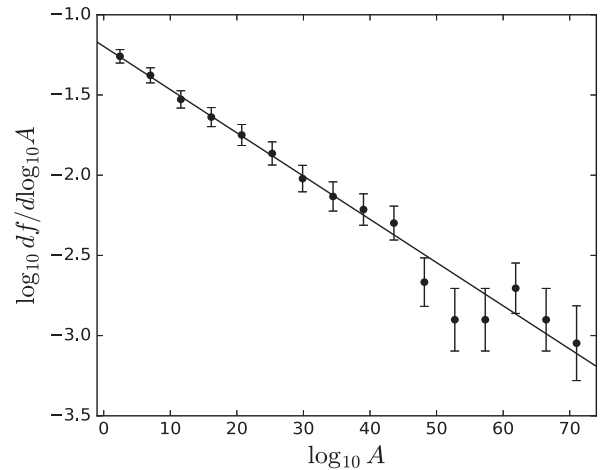
Figure 3. Distribution of amplification factors. The jackknife estimated errorbars increase towards large values of A due to the decrease in number of systems. The power-law fit gives $\log_{10} df/d\log_{10} A = \gamma \log_{10} A + \delta$, with $\gamma = -0.0270 \pm 0.0008$ and $\delta = -1.20 \pm 0.01$.

of numerical accuracy, i.e. the Bulirsch–Stoer tolerance, ϵ. For an accuracy close to double-precision, the fraction of irreversible solutions is about half, consistent with the results of Lehto et al. (2008). By increasing the numerical accuracy beyond machine-precision, we demonstrate that we are able to further decrease the fraction of irreversible solutions. The data are accurately fitted by a power law, given by

$$\log_{10} f_{irr} = \alpha \log_{10} \epsilon + \beta, \tag{1}$$

with $\alpha = 0.029 \pm 0.001$ and $\beta = 0.15 \pm 0.04$.

In Fig. 3, we plot the distribution function of the amplification factors of small initial perturbations. This quantity is defined as the Euclidean norm of the distance in position space between the forward and backward integration as a function of time.[2] In

[2]This is similar to the phase space distance (equation 2 of Miller 1964), but only considering the position coordinates.
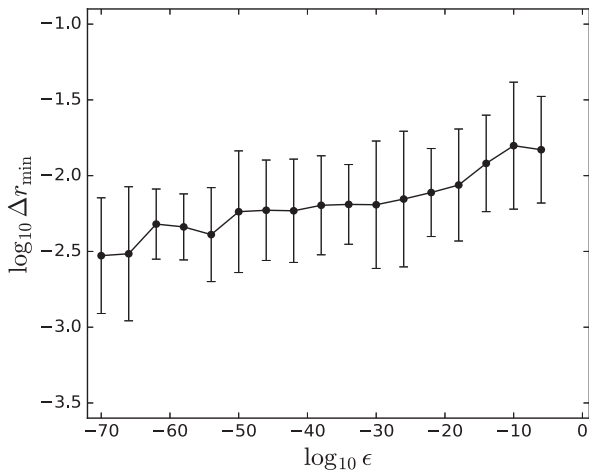
Figure 4. Minimum separation, $\Delta r_{\min}$, between two bodies during a simulation, which required a Bulirsch–Stoer tolerance, ϵ, in order to pass the reversibility test. The data points and errorbars correspond to the average and standard deviation.

a perfectly time reversible integration, the norm should remain zero, but during a numerical integration it will grow exponentially. The final distance divided by the initial distance (after a single integration step) is the amplification factor, A. We find that the distribution is accurately fitted by a power law, given by

$$\log_{10} \frac{\mathrm{d}f}{\mathrm{d}\log_{10} A} = \gamma \log_{10} A + \delta, \qquad (2)$$

with $\gamma = -0.0270 \pm 0.0008$ and $\delta = -1.20 \pm 0.01$. If this relation could be extrapolated, this would imply that for a very high sampling of the Agekyan–Anosova map, there should be some systems with amplification factors exceeding 10^{100}, which would take a long time to calculate up to convergence. The average wall-clock time of our simulations was about 7 h, with the longest run taking 1 month to complete.

We observe a very large difference in the ranges of the axes in Figs 2 and 3. Whereas the fractions vary over two orders of magnitude, the accuracy and amplification factors vary over 70 orders of magnitude in Figs 2 and 3, respectively. The slope is very small and only resolvable due to the very high accuracy and precision of the Brutus code. The results are inconsistent with a flat curve, which intuitively makes sense since with higher numerical accuracy we expect to reduce the fraction of irreversible solutions.

Finally, we show in Fig. 4 that there is only a mild dependence of the required numerical accuracy, ϵ, on the closest encounter between any two bodies during the simulation. This is because a close two-body encounter with a small perturbation from the third body is well approximated by a Keplerian orbit. It is the close encounter plus a large third-body perturbation that may lead to loss of numerical accuracy. The amplification factor is determined by both the lifetime of the triple system, and the magnitude of the finite-time Lyapunov exponent. It remains an open question what determines the rate of exponential growth and transitions in the growth (see e.g. fig. 1 of Portegies Zwart & Boekholt 2018). The exponential sensitivity, or 'the butterfly effect', can be explained in terms of separatrix crossings (Mardling 2008). A trajectory in phase space approaches a separatrix, which divides regions of libration and circulation. It might succeed in crossing the separatrix to a new region in phase space, with potentially different chaotic properties.

However, a nearest neighbour trajectory might fail to cross the separatrix and remain in its current region of phase space, resulting in an exponential magnification of its separation in phase space from the other trajectory. A detailed study of the relation between orbital geometries and the rate of exponential growth of small perturbations will be presented elsewhere.

The errorbars in Figs 2 and 3 result from a finite sampling of the Agekyan–Anosova map. The exact fraction of irreversible solutions as a function of accuracy is obtained when sampling the map with infinite resolution. However, this is impractical, and instead we generate a sample of 1212 initial conditions by overlaying a uniform grid on the map with a grid-spacing of 0.015625. The uncertainty due to the finite sample is estimated by the jackknife resampling method. The errorbars in Fig. 4 are much larger, simply because the variation in the minimal distance between two bodies varies a lot among different simulations with the same accuracy. This shows that the exponential growth is not driven by close two-body encounters, but rather by a prolonged phase of strong three-body encounters.

3 DISCUSSION

The absolute values of the coefficients α and γ are statistically indistinguishable. This suggests that a relation exists between the distribution of amplification factors and the fraction of irreversible solutions. Given a certain value of the Bulirsch–Stoer tolerance, ϵ, we can only resolve amplification factors of order ϵ^{-1}, i.e. with $\epsilon = 10^{-10}$ amplification factors of order $A = 10^{10}$ can be resolved. Hence, the fraction of irreversible solutions equals the fraction of systems with an $A > \epsilon^{-1}$, i.e. $f_{irr}(\epsilon) = F(A > \epsilon^{-1})$. Thus, the fraction of irreversible solutions for a given numerical accuracy, ϵ, is determined by the distribution of amplification factors of the systems in the initial condition map. In Fig. 3, we show that the distribution of amplification factors is also a power law. In Appendix B, we demonstrate that the two power laws in Figs 2 and 3 are related.

In this study, we limit our initial conditions to three-body systems in free fall, i.e. with zero angular momentum. The power law indices measured in Section 2 reflect these initial conditions and the way the homology map was sampled. It is likely that for a different set of initial conditions, such as a Plummer distribution with non-zero angular momentum, the power-law indices might be different. Nevertheless, large amplification factors can still be expected for some non-zero angular momentum systems, since the amplification factor is determined not only by the finite-time Lyapunov exponent (through close triple encounters) but also by the duration of the triple interaction, which is considerably longer for systems with a larger angular momentum (e.g. fig. 7 of Boekholt & Portegies Zwart 2015). An example for a binary-single scattering experiment is given in fig. 4 of Dejonghe & Hut (1986) who measure an amplification factor of about 10^{40} over a time interval of about 5400 N-body time units.

In the limit of infinite accuracy ($\epsilon \to 0$), we retrieve the microscopic time-reversibility of Newton's equations of motion. In the presence of perturbations of size ϵ, whether numerical or physical, a fraction of systems becomes irreversible. As a concrete application of our result, we consider three black holes, each of a million solar masses, and initially separated from each other by roughly one parsec. Such a configuration is not uncommon among supermassive black holes in the concordance model of cosmology and hierarchical galaxy formation (Amaro-Seoane et al. 2010). Here, we will focus specifically on the subset of triples that approach

zero angular momentum, consistent with the systems studied in this work. From Fig. 4, we estimate that the closest approach between any two black holes is on average between $10^{-2.5}$ and 10^{-2} parsec, during which the Newtonian approximation still holds. A parsec equals 10^{51} Planck lengths. Hence, from equation (1) and Fig. 2, we estimate that up to 5 per cent of triples with zero angular momentum are irreversible up to the Planck length, thus rendering them fundamentally unpredictable.

ACKNOWLEDGEMENTS

REFERENCES

Aarseth S. J., Anosova J. P., Orlov V. V., Szebehely V. G., 1994, Celest. Mech. Dyn. Astron., 58, 1

Agekyan T. A., Anosova Z. P., 1967, AZh, 44, 1261

Agekyan T. A., Anosova Z. P., 1968, SvA, 11, 1006

Amaro-Seoane P., Sesana A., Hoffman L., Benacquista M., Eichhorn C., Makino J., Spurzem R., 2010, MNRAS, 402, 2308

Anosova Z. P., 1991, Celest. Mech. Dyn. Astron., 51, 1

Anosova Z. P., Nebukin A. V., 1991, A&A, 252, 410

Anosova J. P., Orlov V. V., Aarseth S. J., 1994, Celest. Mech. Dyn. Astron., 60, 365

Boekholt T., Portegies Zwart S., 2015, ComAC, 2, 2

Boekholt T. C. N., Pelupessy F. I., Heggie D. C., Portegies Zwart S. F., 2016, MNRAS, 461, 3576

Bulirsch R., Stoer J., 1964, Numerische Mathematik. Springer-Verlag, Berlin, p. 413

Burrau C., 1913, Astron. Nachr., 195, 113

Correia A. C. M., 2018, Icarus, 305, 250

Dejonghe H., Hut P., 1986, in Hut P., McMillan S. L. W., eds, The Use of Supercomputers in Stellar Dynamics Vol. 267 of Lecture Notes in Physics, Round-Off Sensitivity in the N-Body Problem. Springer-Verlag, Berlin, p. 212

Goodman J., Heggie D. C., Hut P., 1993, ApJ, 415, 715

Hayes W. B., 2007, NatPh, 3, 689

Heggie D. C., 1991, in Roeser S., Bastian U., eds, Predictability, Stability, and Chaos in N-Body Dynamical Systems Chaos in the N-body Problem of Stellar Dynamics. Plenum Press, New York, p. 47

Heggie D. C., Mathieu R. D., 1986, in Hut P., McMillan S. L. W., eds, The Use of Supercomputers in Stellar Dynamics Vol. 267 of Lecture Notes in Physics, Standardised Units and Time Scales. Springer-Verlag, Berlin, p. 233

Hut P., Bahcall J. N., 1983, ApJ, 268, 319

Ito T., Tanikawa K., 2002, MNRAS, 336, 483

Laskar J., 1989, Nature, 338, 237

Lehto H. J., Kotiranta S., Valtonen M. J., Heinämäki P., Mikkola S., Chernin A. D., 2008, MNRAS, 388, 965

Leigh N. W. C., Wegsman S., 2018, MNRAS, 476, 336

Mardling R. A., 2008, Resonance, Chaos and Stability: The Three-Body Problem in Astrophysics. Cambridge University Press, Cambridge, p. 59

Martynova A. I., Orlov V. V., 2014, ARep, 58, 756

Miller R. H., 1964, ApJ, 140, 250

Orlov V. V., Titov V. A., Shombina L. A., 2016, ARep, 60, 1083

Portegies Zwart S., Boekholt T., 2014, ApJ, 785, L3

Portegies Zwart S. F., Boekholt T. C. N., 2018, CNSNS, 61, 160

Quinlan G. D., Tremaine S., 1992, MNRAS, 259, 505

Stone N. C., Leigh N. W. C., 2019, Nature, 576, 7787

Sussman G. J., Wisdom J., 1992, Science, 257, 56

Szebehely V., Peters C. F., 1967, AJ, 72, 876

Tanikawa K., Umehara H., Abe H., 1995, Celest. Mech. Dyn. Astron., 62, 335

Urminsky D. J., 2010, MNRAS, 407, 804

Valluri M., Merritt D., 2000, in Gurzadyan V. G., Ruffini R., eds, The Chaotic Universe Orbital Instability and Relaxation in Stellar Systems. World Scientific Publishing, Singapore, p. 229

Valtonen M., Karttunen H., 2006, The Three-Body Problem. Cambridge University Press, Cambridge

Wisdom J., Peale S. J., Mignard F., 1984, Icarus, 58, 137

APPENDIX A: NUMERICAL METHODS

We adopt the Agekyan–Anosova map (Agekyan & Anosova 1967, 1968) and sample it uniformly with a resolution of 0.015 625. This results in an ensemble of 1212 initial realizations. For a high-resolution version of the map, see the 'warrior shield' by Lehto et al. (2008).

We use the arbitrary-precision N-body code Brutus (Boekholt & Portegies Zwart 2015) and vary its two main parameters, ϵ, the Bulirsch–Stoer tolerance, and, L_w, the word-length in bits. In order to reduce the grid of parameters to vary, we fix $L_w = 1024$ bits, which corresponds to about 300 digits, which is sufficient for the calculations in this work.

We evolve each initial realization to the point of a permanent binary-single configuration, or a maximum time of 10 000 time units (Heggie & Mathieu 1986). The single body is defined to be permanently escaped if (1) it is separated from the binary centre of mass by more than 10 distance units, (2) it is moving away from the binary, and (3) its energy is positive. Note that for near parabolic escapes these criteria are only fulfilled at very large separations, potentially exceeding our maximum time limit.

Once the system has fulfilled the escape criteria at time $t = T$, we reverse the velocity of each body, and continue the integration up to $t = 2T$. Then, we compare the initial snapshot to the final snapshot by calculating the Euclidean norm of the distance between the solutions in position space. This is similar to the phase space distance defined by Miller (1964), but only using the position coordinates. On the one hand, this is done to avoid confusion between adding quantities with different units, but also because from experience, we noticed that if the Euclidean norm in position space is small, then this is also the case in momentum space and vice versa. Hence, if the initial positions are retrieved after performing the reversibility experiment, then this must also be the case for the momenta. Otherwise chaos would have caused the time-reversed solution to exponentially diverge from the forward-integrated solution. If the Euclidean norm of the distance in position space, Δ, is sufficiently small, then the simulation has passed the reversibility test. If not, then we redo the simulation with a higher accuracy (smaller ϵ), until for some accuracy the reversibility test is successful. This way, we iteratively increase the fraction of reversible solutions. Our criterion for deeming a solution reversible is $\log_{10} \Delta < -3$, i.e. each position coordinate of each body in the initial and final snapshot differs only in the third decimal place or beyond. The phenomenon that, after iteratively increasing the accuracy and precision, the first n decimal places of the solution have converged, and the solution has started

to become time-reversible up to *n* decimal places, is defined as *definitive reversibility* (Portegies Zwart & Boekholt 2018).

Once we have the ensemble of definitive reversible solutions for the Agekyan–Anosova map, we measure the lifetime, T, and amplification factor, A, for each system. The lifetime is measured by considering the final snapshot of the forward integration, consisting of the permanent binary+single, and then retracing their steps to the moment when the binary+single were closest to the moment of the final ejection. Especially for near parabolic escapes, this can cut off a significant fraction of the simulation as the escape criteria are only fulfilled when the binary and single are at very large separations. The amplification factor of a small initial perturbation is calculated by measuring the Euclidean norm of the distance in position space between the forward and backward integration as a function of time. The backward integration diverges exponentially from the forward integration at a rate given by the finite-time Lyapunov exponent (Heggie 1991) of the system. Note that the perturbation is smallest at the end of the forward integration. The amplification factor is then defined as the ratio of the initial and final Euclidean norm of the distance in position space, $A = \Delta_T/\Delta_0$.

APPENDIX B: SUPPLEMENTARY TEXT

In Fig. 1, we present our low-resolution Agekyan–Anosova map, where we plot the lifetime of the triple system as a function of initial condition. We observe the blue 'rivers' of systems that dissolve rapidly, and the surrounding chaotic landscape where nearest neighbours can have very different lifetimes. When comparing the least accurate ($\epsilon = 10^{-6}$) and the most accurate ($\epsilon = 10^{-70}$) maps, we observe that there are 'microscopic' differences. For example the black dots, which represent very long lived systems, are differently populated in the two maps, However, in a 'macroscopic' sense, the maps look similar. To make the comparison more quantitative, we take the distributions of triple lifetime and binding energy of the final, permanent binary, and perform a two-sample Kolmogorov–Smirnoff test. We find that the least and most accurate data sets are statistically indistinguishable according to the Kolmogorov–Smirnoff test, giving *p*-values of 0.72 (lifetimes) and 0.85 (binding energies). These results demonstrate that, for the Agekyan–Anosova map, approximate computations are nevertheless reliable in a statistical sense (Goodman et al. 1993), verifying the ergodic-like property of '*nagh-Hoch*' (Portegies Zwart & Boekholt 2018).

Correlation between amplification factor and fraction of irreversible solutions

In Fig. 2, we plot the fraction of irreversible solutions as a function of numerical accuracy, i.e. the Bulirsch–Stoer tolerance, ϵ. We observe that initially, at low accuracy, the fraction of irreversible solutions is close to unity. As we increase the accuracy to about standard double-precision, the reversible and irreversible fractions have become roughly equal. This result is consistent with that of Lehto et al.

(2008). By increasing the numerical accuracy beyond machine-precision, we demonstrate that we are able to further decrease the fraction of irreversible solutions. In Fig. 2, we show that the fraction of irreversible solutions is accurately fitted by a power law, given by

$$\log_{10} f_{\mathrm{irr}} = \alpha \log_{10} \epsilon + \beta, \tag{B1}$$

with $\alpha = 0.029 \pm 0.001$ and $\beta = 0.15 \pm 0.04$. By the time we reached a Bulirsch–Stoer tolerance of $\epsilon = 10^{-70}$, the fraction of irreversible solutions had dropped to about 1 per cent, which is when we decided to end the iteration for practical reasons.

In Fig. 3, we plot the distribution of amplification factors. This distribution is also accurately fitted by a power law, given by

$$\log_{10} \frac{\mathrm{d}f}{\mathrm{d}\log_{10} A} = \gamma \log_{10} A + \delta, \tag{B2}$$

with $\gamma = -0.0270 \pm 0.0008$ and $\delta = -1.20 \pm 0.01$. If this relation could be extrapolated, this would imply that for a very high sampling of the Agekyan–Anosova map, there should be some systems with amplification factors exceeding 10^{100}, which would take a long time to calculate up to convergence. The average wall-clock time of our simulations was about 7 h, with the longest run taking 1 month to complete.

The coefficients α and γ in equations (B1) and (B2) are equal to within 3σ. This suggests that there is a relation between the distribution of amplification factors and the fraction of irreversible solutions for some specified numerical accuracy. Given a Bulirsch–Stoer parameter, ϵ, we are only able to resolve amplification factors of order ϵ^{-1}, i.e. with an $\epsilon = 10^{-10}$ we can resolve amplification factors of order $A = 10^{10}$. Hence, the fraction of irreversible solutions should approximately be equal to the fraction of systems with an $A > \epsilon^{-1}$, i.e. $f_{\mathrm{irr}}(\epsilon) = F(A > \epsilon^{-1})$. Hence, using equation (B2), we can derive the following:

$$f_{\mathrm{irr}}\left(\log_{10}\epsilon\right) = F\left(\log_{10} A > -\log_{10}\epsilon\right), \tag{B3}$$

$$f_{\mathrm{irr}}\left(\log_{10}\epsilon\right) = \int_{-\log_{10}\epsilon}^{\infty} \frac{\mathrm{d}f}{\mathrm{d}\log_{10} A}\mathrm{d}\log_{10} A, \tag{B4}$$

$$f_{\mathrm{irr}}\left(\log_{10}\epsilon\right) = \int_{-\log_{10}\epsilon}^{\infty} 10^{\gamma \log_{10} A+\delta}\mathrm{d}\log_{10} A, \tag{B5}$$

$$f_{\mathrm{irr}}\left(\log_{10}\epsilon\right) \sim \left[10^{\gamma \log_{10} A}\right]_{-\log_{10}\epsilon}^{\infty}, \tag{B6}$$

$$f_{\mathrm{irr}}\left(\log_{10}\epsilon\right) \sim 10^{-\gamma \log_{10}\epsilon}. \tag{B7}$$

Finally, taking the logarithm, we can write

$$\log f_{\mathrm{irr}} \sim -\gamma \log_{10}\epsilon. \tag{B8}$$

Comparing this expression to equation (B1), we conclude that $-\gamma = \alpha$.

This paper has been typeset from a TEX/LATEX file prepared by the author.