Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

Mon. Not. R. Astron. Soc. **403,** 525–544 (2010) doi:10.1111/j.1365-2966.2009.16192.x

The orbital evolution induced by baryonic condensation in triaxial haloes

Monica Valluri,[1⋆] Victor P. Debattista,[2⋆†] Thomas Quinn[3] and Ben Moore[4]

[1]*Department of Astronomy, University of Michigan, Ann Arbor, MI 48109, USA*
[2]*Jeremiah Horrocks Institute, University of Central Lancashire, Preston PR1 2HE*
[3]*Astronomy Department, University of Washington, Box 351580, Seattle, WA 98195-1580, USA*
[4]*Department of Theoretical Physics, University of Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland*

Accepted 2009 December 8. Received 2009 December 4; in original form 2009 June 23

ABSTRACT

Using spectral methods, we analyse the orbital structure of prolate/triaxial dark matter (DM) haloes in *N*-body simulations in an effort to understand the physical processes that drive the evolution of shapes of DM haloes and elliptical galaxies in which central masses are grown. A longstanding issue is whether the change in the shapes of DM haloes is the result of chaotic scattering of the major family of box orbits that serves as the backbone of a triaxial system, or whether they change shape adiabatically in response to the evolving galactic potential. We use the characteristic orbital frequencies to classify orbits into major orbital families, to quantify orbital shapes and to identify resonant orbits and chaotic orbits. The use of a frequency-based method for distinguishing between regular and chaotic *N*-body orbits overcomes the limitations of Lyapunov exponents which are sensitive to numerical discreteness effects. We show that regardless of the distribution of the baryonic component, the shape of a DM halo changes primarily due to changes in the shapes of individual orbits within a given family. Orbits with small pericentric radii are more likely to change both their orbital type and shape than orbits with large pericentric radii. Whether the evolution is regular (and reversible) or chaotic (and irreversible), it depends primarily on the radial distribution of the baryonic component. The growth of an extended baryonic component of any shape results in a regular and reversible change in orbital populations and shapes, features that are not expected for chaotic evolution. In contrast, the growth of a massive and compact central component results in chaotic scattering of a significant fraction of both box and long-axis tube orbits, even those with pericentre distances much larger than the size of the central component. Frequency maps show that the growth of a disc causes a significant fraction of halo particles to become trapped by major global orbital resonances. We find that despite the fact that shape of a DM halo is always quite oblate following the growth of a central baryonic component, a significant fraction of its orbit population has the characteristics of its triaxial or prolate progenitor.

Key words: methods: *N*-body simulations – galaxies: evolution – galaxies: haloes – galaxies: kinematics and dynamics – galaxies: structure – dark matter.

1 INTRODUCTION

The condensation of baryons to the centres of dark matter (DM) haloes is known to make them more spherical or axisymmetric (Dubinski 1994; Kazantzidis et al. 2004a; Debattista et al. 2008, hereafter D08). D08 found that the halo shape changes by $\Delta(b/a) \gtrsim 0.2$ out to at least half the virial radius. This shape change reconciles the strongly prolate-triaxial shapes found in collisionless *N*-body simulations of the hierarchal growth of haloes (Bardeen et al. 1986; Barnes & Efstathiou 1987; Frenk et al. 1988; Dubinski & Carlberg 1991; Jing & Suto 2002; Bailin & Steinmetz 2005; Allgood et al. 2006) with observations, which generally find much rounder haloes (Schweizer, Whitmore & Rubin 1983; Sackett & Sparke 1990; Franx & de Zeeuw 1992; Huizinga & van Albada 1992; Buote & Canizares 1994; Franx, van Gorkom & de Zeeuw 1994; Kuijken & Tremaine 1994; Bartelmann, Steinmetz & Weiss 1995; Kochanek 1995; Olling 1995, 1996; Schoenmakers, Franx & de Zeeuw 1997; Koopmans, de Bruyn & Jackson 1998; Olling & Merrifield 2000; Andersen et al. 2001; Buote et al. 2002; Barnes & Sellwood 2003; Debattista 2003; Iodice et al. 2003; Oguri, Lee & Suto 2003; Diehl & Statler 2007; Banerjee & Jog 2008).

What is the physical mechanism driving shape change? Options suggested in the literature include two possibilities. The first is that the presence of a central mass concentration scatters box orbits that

⋆E-mail: mvalluri@umich.edu(MV) ; vpdebattista@uclan.ac.uk (VPD)
†RCUK Fellow.

serve as the backbone of a triaxial potential, rendering them chaotic (Gerhard & Binney 1985; Merritt & Valluri 1996). Chaotic orbits in a stationary potential do not conserve any integrals of motion other than the energy E and consequently are free to uniformly fill their allowed equipotential surface. Since the potential is, in general, rounder than the density distribution, chaotic diffusion results in evolution to a more oblate or even a spherical shape (Merritt & Quinlan 1998; Kalapotharakos 2008). The second possibility is that the change of the central potential occurs because the growth of the baryonic component causes orbits of collisionless particles in the halo to respond by changing their shapes in a regular (and therefore reversible) manner (Holley-Bockelmann et al. 2002).

Time dependence in a potential is also believed to result in chaotic mixing (Terzić & Kandrup 2004; Kandrup & Novotny 2004) and has been invoked as the mechanism that drives violent relaxation. However, a more recent analysis of mixing during a major merger showed that the rate and degree of mixing in energy and angular momentum are not consistent with chaotic mixing, but rather the particles retain strong memory of their initial energies and angular momenta even in strongly time-dependent potentials (Valluri et al. 2007).

One of the principal features of chaotic evolution is irreversibility. This irreversibility arises from two properties of chaotic orbits. First, chaotic orbits are exponentially sensitive to small changes in initial conditions even in a collisionless system. Secondly, chaotic systems display the property of *chaotic mixing* (Lichtenberg & Lieberman 1992). Using this principle of irreversibility, D08 argued that if chaotic evolution is the primary driver of shape change, then if, subsequently, the central mass concentration is artificially 'evaporated', the system would not be able to revert to its original triaxial distribution. D08 showed that growing baryonic components inside prolate/triaxial haloes led to a large change in the shape of the halo. Despite these large changes, by artificially evaporating the baryons, they showed that the underlying halo phase-space distribution is not grossly altered unless the baryonic component is too massive or centrally concentrated, or transfers significant angular momentum to the halo. This led them to argue that the chaotic evolution alone cannot explain the shape change since such a process is irreversible. They speculated that at most only slowly diffusive chaos occurred in their simulations. Using test particle orbit integrations, they also showed that box orbits largely become deformed, possibly changing into tube orbits, during disc growth, but do not become strongly chaotic.

D08 employed irreversibility as a convenient proxy for the presence of chaos. In this paper, we undertake an orbital analysis of some of the models studied by D08 to better understand the mechanism that drives shape change. Our goal is to understand whether chaotic orbits are an important driver of shape change and if so under what conditions they are important. We also wish to understand how the orbital populations in haloes change when a centrally concentrated baryonic component grows inside a triaxial DM halo. Finally, we would like to understand under what circumstances orbits change their classification.

This paper is organized as follows. In Section 2, we describe the simulations used in this paper and briefly describe three models from D08 as well as two additional simulations. In Section 3, we describe the principal technique: numerical analysis of fundamental frequencies (NAFF) that we use to obtain frequency spectra and fundamental frequencies and describe how these frequencies are used to characterize orbits. In Section 4, we describe the results of our analysis of five different simulations. In Section 5, we summarize our results and discuss their implications.

2 NUMERICAL SIMULATIONS

We formed prolate/triaxial haloes via mergers of systems, as described in Moore et al. (2004). The initially spherical NFW (Navarro, Frenk & White 1996) haloes were generated from a distribution function using the method described in Kazantzidis, Magorrian & Moore (2004b) with each halo composed of two mass species arranged on shells. The outer shell has more massive particles than the inner one, similar to the method described by Zemp et al. (2008), which allows for higher mass resolution at small radii. Our model halo A was generated by the head-on merger of two prolate haloes, themselves the product of a binary merger of spherical systems. The first merger placed the concentration $c = 10$ haloes 800 kpc apart approaching each other at 50 km s^{-1}, while the second merger starts with the remnant at rest, 400 kpc from an identical copy. The resulting halo is highly prolate with a mild triaxiality. Halo model B was produced by the merger of two spherical haloes starting at rest, 800 kpc apart and is prolate, with $\langle b/a\rangle = \langle c/a\rangle \simeq 0.58$. Halo A has $\langle b/a\rangle \simeq 0.45$ and $\langle c/a\rangle \simeq 0.35$ while halo B has $\langle b/a\rangle = \langle c/a\rangle \simeq 0.58$ (see fig. 3 of D08 for more details). Both haloes A and B consist of 4×10^6 particles. The outer particles are $\sim$18 times more massive in halo A and $\sim$5 times more massive in halo B. A large part of the segregation by particle mass persists after the mergers and the small radius regions are dominated by low-mass particles (cf. Dehnen 2005). We used a softening parameter $\epsilon = 0.1$ kpc for all halo particles. The radius, r_{200}, at which the halo density is 200 times the mean density of the Universe and the total mass within this radius, M_{200}, are given in Table 1.

Once we produced the prolate/triaxial haloes, we inserted a baryonic component, either a disc of particles that remains rigid throughout the experiments or softened point particles. The parameters that describe the distribution of the baryonic components are given in

Table 1. The simulations in this paper.

Run number (from D08)	Run description	Halo	r_{200} (kpc)	M_{200} (10^{12} M$_\odot$)	M_b (10^{11} M$_\odot$)	f_b	R_b (kpc)	t_g (Gyr)	t_e (Gyr)
SA1	*Triax + Disc*	A	215	4.5	1.75	0.039	3.0	5	2.5
P$_l$A3	*Triax + Bulg*	A	215	4.5	1.75	0.039	1.0	5	2.5
P$_f$B2	*Prolt + Ellip*	B	106	0.65	0.7	0.108	3.0	10	4
P$_l$A4	*Triax + hardpt*	A	215	4.5	1.75	0.039	0.1	5	2.5
P$_l$B3	*Prolt + hardpt*	B	106	0.65	0.35	0.054	0.1	5	5

Note. M_b is the mass in baryons and f_b is the baryonic mass fraction. For the particle simulations (P$_f$B2, P$_l$B3, P$_l$A3, P$_l$A4), R_b refers to the softening of the spherical baryonic distribution particle(s). For simulation SA1, R_b refers to the scalelength of the baryonic disc.

Table 1. In four of the models (P_lA3, P_lA4, P_fB2 and P_lB3), the baryonic component is simply a softened point mass with softening scalelength given by R_b. In model SA1, the density distribution of the disc was exponential with scalelength of the baryonic component R_b and Gaussian scaleheight $z_b/R_b = 0.05$. The disc was placed with its symmetry axis along the triaxial halo's short axis in model SA1 (additional orientations of the disc relative to the principal axes were also simulated but their discussion is deferred to a future paper). Initially, the disc has negligible mass, but it grows adiabatically and linearly with time to a mass M_b during a time t_g. After this time, we slowly evaporated it during a time t_e. We stress that this evaporation is a numerical convenience for testing the effect of chaos on the system, and should not be mistaken for a physical evolution. The disc is composed of 300 K equal-mass particles each with a softening $\epsilon = 100$ pc. From $t = 0$ to $t_g + t_e$, the halo particles are free to move and achieve equilibrium with the baryons as their mass changes, but all disc particles are frozen in place. The masses of models with single softened particles are also grown in the same way; these are models P_fB2, P_lB3 from D08, and P_lA3 and P_lA4 which are new to this paper. D08's naming convention for these experiments used 'P' subscripted by '*f*' for particles frozen in place and by '*l*' for live particles free to move.

Three different baryonic components are grown in the triaxial halo A: in model SA1, the baryons are in the form of a disc grown perpendicular to the short axis and the model is referred to as *Triax + Disc*; in model P_lA3, the baryonic component is a softened central point mass resembling a bulge and the model is referred to as *Triax + Bulg*; finally, in model P_lA4, the baryonic component is a hard central point mass with a softening of 0.1 kpc and the model is referred to as *Triax + hardpt*. Two different baryonic components are grown in the prolate halo B: in model P_fB2, the baryonic component loosely resembles an elliptical galaxy so this run is referred to as *Prolt + Ellip*; in model P_lB3, the baryonic component is a hard central point mass with a softening of 0.1 kpc and is referred to as *Prolt + hardpt*. For model P_fB2, D08 showed that there is no significant difference in the evolution if the central particle is live instead of frozen, all other things being equal. P_lA3 was constructed specifically for this paper in order to have a triaxial halo model with a moderately soft spherical baryonic distribution which can be contrasted with the prolate halo model P_fB2, while P_lA4 is a triaxial halo model which can be contrasted with the prolate halo model P_lB3.

For each model, there were five phases in evolution. The initial triaxial or prolate halo without the baryonic component is referred to as *phase a*. There is then a phase (of duration t_g) during which the halo's shape is evolving as the baryonic component is grown adiabatically. We do not study orbits in this phase since the potential is evolving with time. When the baryonic component has finished growing to full strength and the halo has settled to a new equilibrium, the model is referred to as being in *phase b*. The baryonic component is 'adiabatically evaporated' over a time-scale of duration t_e listed in Table 1. Again we do not study orbits during this period when the potential is evolving with time. After the baryonic component has been adiabatically evaporated completely and the halo has returned to an equilibrium configuration, the halo is referred to as being in *phase c*. We only study the orbits in the halo during the three phases when the halo is in equilibrium and is stationary (not evolving with time).

The growth of the baryonic component induces several changes in the distributions of the DM halo particles: first is an increase in central density relative to the original NFW halo due to increase in the depth of the central potential (an effect commonly referred to as 'baryonic compression'); secondly, the haloes become more oblate especially within $0.3r_{vir}$. The details of the changes in the density and velocity distributions of DM particles differ slightly depending on the nature of the baryonic component (D08).

All the simulations in this paper, which are listed in Table 1, were evolved with PKDGRAV an efficient, multistepping, parallel tree code (Stadel 2001). We used cell opening angle for the tree code of $\theta = 0.7$ throughout.[1] Additional details of the simulations can be found in D08.

2.1 Computing orbits

In each of the haloes studied, we selected a subsample of between 1000–6000 particles and followed their orbits in each of the three stationary phases of the evolution described in the previous section. The particles were randomly chosen in the haloes at $t = 0$ such that they were inside a fixed outer radius (either 100 or 200 kpc). Since the particles were selected at random from the distribution function, they have the same overall distribution as the entire distribution function within the outer radius selected. We integrated the motion of each a test particle while holding all the other particles fixed in place. We used a fixed time-step of 0.1 Myr and integrated for 50 Gyr, storing the phase-space coordinates of each test particle every 1 Myr. We used such long integration times to ensure we are able to obtain accurate measurements of frequencies (as described in the next section). We carried out this operation for the same subset of particles at *phases a*, *b* and *c*. In model SA1 (*Triax + Disc*), we integrated the orbits of 6000 particles which in *phase a* were within $r = 200$ kpc. In model P_lA3 ($Triax + Bulg$) and model P_lA4 (*Triax + hardpt*), we considered a subsample of 5000 particles starting within $r = 100$ kpc. In models P_fB2 and P_lB3, we considered orbits of 1000 particles within $r = 200$ kpc. We integrated their orbits as above but we used a smaller time-step $\delta t = 10^4$ yr in the case of P_lA4 and P_lB3, which had harder central point masses. The orbit code computes forces in a frozen potential using an integration scheme that uses forces calculated from the PKDGRAV tree; we used the orbit integration parameters identical to those used for the evolution of the self-consistent models.

3 FREQUENCY ANALYSIS

In a three-dimensional (3D) galactic potential that is close to integrable, all orbits are quasi-periodic. If an orbit is quasi-periodic (or regular), then any of its coordinates can be described explicitly as a series,

$$x(t) = \sum_{k=1}^{\infty} A_k e^{i\omega_k t}, \quad (1)$$

where the ω_ks are the oscillation frequencies and the A_ks are the corresponding amplitudes. In a 3D potential, each ω_k can be written as an integer linear combination of three fundamental frequencies $\omega_1, \omega_2, \omega_3$ (one for each degree of freedom). If each component of the motion of a particle in the system [e.g. $x(t)$] is followed for several ($\sim$100) dynamical times, a Fourier transform of the trajectory yields a spectrum with discrete peaks. The locations of the peaks in the spectrum correspond to the frequencies ω_k and their amplitudes A_k can be used to compute the linearly independent

[1] Opening angle θ is used in tree codes to determine how long-range forces from particles acting at a point are accumulated (Barnes & Hut 1986).

fundamental frequencies (Boozer 1982; Kuo-Petravic et al. 1983; Binney & Tremaine 2008).

Binney & Spergel (1982, 1984) applied this method to galactic potentials and obtained the frequency spectra using a least-squares technique to measure the frequencies ω_k. Laskar (1990, 1993) developed a significantly improved numerical technique (NAFF) to decompose a complex time series of the phase-space trajectory of an orbit of the form $x(t) + i\,v_x(t)$, (where v_x is the velocity along the x coordinate). Valluri & Merritt (1998) developed their own implementation of this algorithm that uses integer programming to obtain the fundamental frequencies from the frequency spectrum. In this paper, we use this latter implementation of the NAFF method.

We refer readers to the above papers and to section 3.7 of Binney & Tremaine (2008) for a detailed discussion of the main idea behind the recovery of fundamental frequencies. For completeness, we provide a brief summary here. The NAFF algorithm for frequency analysis allows one to quickly and accurately compute the fundamental frequencies that characterize the quasi-periodic motion of regular orbits. The entire phase space at a given energy can then be represented by a frequency map which is a plot of ratios of the fundamental frequencies of motion. A frequency map is one of the easiest ways to identify families of orbits that correspond to resonances between the three degrees of freedom.

The structure of phase space in 3D galactic potentials is quite complex and we summarize some of its properties here to enable the reader to more fully appreciate the results of the analysis that follows. When an integrable potential is perturbed, its phase-space structure is altered, resulting in the appearance of resonances (Lichtenberg & Lieberman 1992). Resonances are regions of phase space where the three fundamental frequencies are not linearly independent of each other, but two or more of them are related to each other via integer linear relations. As the perturbation in the potential increases, the potential deviates further and further from integrability, and a larger and larger fraction of the phase space becomes associated with resonances.

In a 3D potential, orbits that satisfy one resonance condition such as $l\omega_x + m\omega_y + n\omega_z = 0$ are referred to as 'thin orbit' resonances since they cover the surface of a two-dimensional surface in phase space (Merritt & Valluri 1999). If two independent resonance conditions between the fundamental frequencies exist, then the orbit is a closed periodic orbit. Orbits that have frequencies close to the resonant orbit frequencies are said to be resonantly trapped. Such orbits tend to have properties similar to that of the parent resonance, but get 'thicker' as their frequencies move away from the resonance. At the boundary of the region of phase space occupied by a resonant family is a region called the 'separatrix'. The separatrix is the boundary separating orbits with different orbital characteristics. In this case, it is the region between orbits that have frequencies that are similar to the resonant orbits and orbits that are not resonant. Chaotic orbits often occur in a 'stochastic layer' close to resonances and at the intersections of resonances. In fact, one of the primary factors leading to an increase in the fraction of chaotic orbits is the overlap of resonances (Chirikov 1979). Chaotic orbits that are close to a resonant family are referred to as 'resonantly trapped' or 'sticky orbits' (Habib, Kandrup & Mahon 1997) and are often only weakly chaotic. Orbits that are 'sticky' behave like the resonant parent orbit for extremely long times and therefore do not diffuse freely over their energy surface or undergo significant chaotic mixing.

The frequency analysis method allows one to map the phase-space structure of a distribution function and to easily identify the most important resonances by plotting ratios of pairs of frequencies (e.g. ω_x/ω_z versus ω_y/ω_z) for many thousands of orbits in the potential. In such a frequency map, resonances appear as straight lines. Stable resonances appear as filled lines about which many points cluster, and unstable resonances appear as 'blank' or depopulated lines. The strength of the resonances can be determined by the number of orbits that are associated with them.

3.1 Overcoming microchaos in *N*-body simulations

In N-body systems like those considered in this paper, the galactic mass distribution is realized as a discrete set of point masses. The discretization of the potential is known to result in exponential deviation of nearby orbits, even in systems where all orbits are expected to be regular (Miller 1964; Kandrup & Smith 1991; Goodman, Heggie & Hut 1993; Valluri & Merritt 2000; Hemsendorf & Merritt 2002). However, as the number of particles in a simulation is increased, and when point masses are softened, the majority of orbits begin to appear regular despite the fact that their non-zero Lyapunov exponent implies that they are chaotic. Hemsendorf & Merritt (2002) showed that this Lyapunov exponent saturates at a finite value beyond a few hundred particles and corresponds to an e-folding time-scale of 1/20 of a system crossing time (for systems with $N \sim 10^5$ particles). Despite having large Lyapunov exponents (i.e. short e-folding times), these orbits behave and look much like regular orbits (Kandrup & Sideris 2001; Kandrup & Siopis 2003). This property of N-body orbits to have non-zero Lyapunov exponents has been referred to as 'microchaos' (Kandrup & Sideris 2003) or the 'Miller Instability' (Hemsendorf & Merritt 2002; Valluri et al. 2007) and suggests that Lyapunov exponents, while useful in continuous potentials, are not a good measure of chaotic behaviour resulting from the global potential when applied to N-body systems. This is a strong motivation for our use of a frequency-based method which, as we demonstrate, is extremely effective at distinguishing between regular and chaotic orbits and is apparently largely unaffected by microchaos.

We now discuss how frequency analysis can be used to distinguish between regular and chaotic orbits. In realistic galactic potentials, most chaotic orbits are expected to be weakly chaotic and lie close to regular orbits mimicking their behaviour for long times. The rate at which weakly chaotic orbits change their orbital frequencies can be used as a measure of chaos. Laskar (1993) showed that the change in the fundamental frequencies over two consecutive time intervals can be used as a measure of the stochasticity of an orbit. This method has been used to study the phase-space structure in galactic potentials (Papaphilippou & Laskar 1996, 1998; Valluri & Merritt 1998). Examples of frequency spectra for each component of motion, and their resolution into three fundamental frequencies are given by Papaphilippou & Laskar (1998) for different types of orbits. For each time series, the spectrum is analysed and the three fundamental frequencies are obtained. In Cartesian coordinates, the frequencies would be $\omega_x, \omega_y, \omega_z$.

For each orbit, we therefore divide the integration time of 50 Gyr into two consecutive segments and use NAFF to compute the fundamental frequencies $\omega_x, \omega_y, \omega_z$ (note that all frequencies in this paper are in units of Gyr^{-1}, therefore units are not explicitly specified everywhere). We compute the three fundamental frequencies $\omega_x(t_1), \omega_y(t_1), \omega_z(t_1)$ and $\omega_x(t_2), \omega_y(t_2), \omega_z(t_2)$ in each of the two intervals t_1 and t_2, respectively. We compute the 'frequency drift' for each frequency component as:

$$\log(\Delta f_x) = \log \left| \frac{\omega_x(t_1) - \omega_x(t_2)}{\omega_x(t_1)} \right|, \tag{2}$$

$$\log(\Delta f_y) = \log\left|\frac{\omega_y(t_1) - \omega_y(t_2)}{\omega_y(t_1)}\right|, \quad (3)$$

$$\log(\Delta f_z) = \log\left|\frac{\omega_z(t_1) - \omega_z(t_2)}{\omega_z(t_1)}\right|. \quad (4)$$

We define the frequency drift parameter $\log(\Delta f)$ (logarithm to base 10) to be the value associated with the largest of the three frequencies f_x, f_y, f_z. The larger the value of the frequency drift parameter, the more chaotic the orbit.

Identifying truly chaotic behaviour, however, also requires that we properly account for numerical noise. In previous studies, orbits were integrated with high numerical precision for at least 100 orbital periods, resulting in highly accurate frequency determination. For instance, Valluri & Merritt (1998) found that orbital frequencies in a triaxial potential could be recovered with an accuracy of 10^{-10} for regular orbits and 10^{-4}–10^{-6} for stochastic orbits using integration times of at least 50 orbital periods per orbit.

In order to use frequency analysis to characterize orbits as regular or chaotic in N-body systems, it is necessary to assess the numerical accuracy of orbital frequencies obtained by the NAFF code. To quantify the magnitude of frequency drift that arises purely from discretization effects (the microchaos discussed above), we select a system that is spherically symmetric and in dynamical equilibrium. All orbits in a smooth spherically symmetric potential are rosettes confined to a single plane (Binney & Tremaine 2008) and are regular. Hence, any drift in orbital frequencies can be attributed entirely to discretization errors (including minute deviations of the N-body potential from perfect sphericity). As a test of our application of the NAFF code to N-body potentials, we analyse orbits in spherical NFW haloes of two different concentrations ($c = 10$ and 20). The haloes are represented by 10^6 particles and have mass $\sim 2 \times 10^{12}\,\mathrm{M}_\odot$. Particles in both cases come in two species with softening of 0.1 and 0.5 kpc. We carried out the frequency analysis of 1000 randomly selected orbits which were integrated for 50 Gyr in the frozen N-body realizations of each of the NFW haloes.

Fig. 1 shows the distribution of values of $\log(\Delta f)$ for both haloes. In both cases, the distribution has a mean value of $\log(\Delta f) = -2.29$,

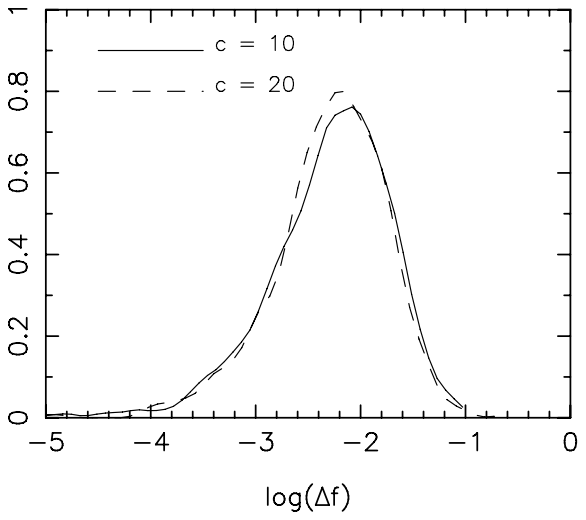

Figure 1. Distributions of frequency drift parameter $\log(\Delta f)$ for 1000 orbits in two different spherical NFW haloes. Despite the difference in concentration $c = 10$ and 20, the distributions are almost identical having a mean of $\log(\Delta f) = -2.29$ and a standard deviation $\sigma \simeq 0.56$, with a significant skewness towards small values of $\log(\Delta f)$.

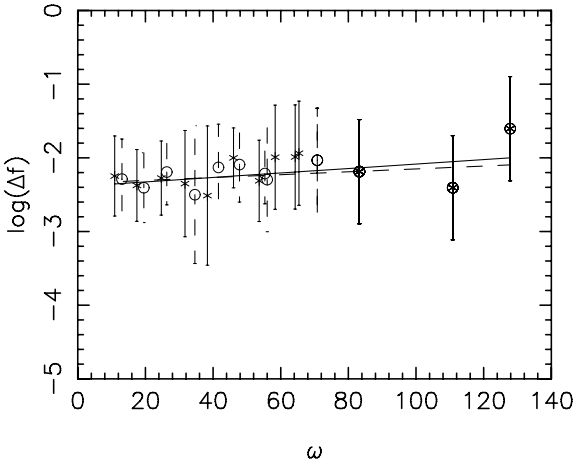

Figure 2. $\log(\Delta f)$ versus ω (in units of Gyr^{-1}) for orbits with $n_\mathrm{p} > 20$ in the two spherical NFW haloes. Stars are for the halo with $c = 10$ and the open circles are for the halo with $c = 20$. The 1000 particles are binned in ω so that each bin contains the same number of particles. The vertical error bars represent the standard deviation in each bin. The straight lines are fits to the data, the slopes of both lines are consistent with zero.

with standard deviations of 0.58 (for the $c = 10$ halo) and 0.54 (for the $c = 20$ halo). Both distributions are significantly skewed towards small values of $\log(\Delta f)$ (skewness $= -0.85$) and are more peaky than Gaussian (kurtosis $= 1.95$). Despite the fact that the two NFW haloes have different concentrations, the distributions of $\log(\Delta f)$ are almost identical, indicating that our chaotic measure is largely independent of the central concentration.

To define a threshold value of $\log(\Delta f)$ at which orbits are classified as chaotic, we note that 99.5 per cent of the orbits have values of $\log(\Delta f) < -1.0$. Since all orbits in a stationary spherical halo are expected to be regular, we attribute all larger values of $\log(\Delta f)$ to numerical noise arising from the discretization of the potential. Henceforth, we classify an orbit in our N-body simulations to be regular if it has $\log(\Delta f) < -1.0$.

To accurately measure the frequency of an orbit it is necessary to sample a significant part of its phase-space structure (i.e. the surface of a 2-torus in a spherical potential or the surface of the 3-torus in a triaxial potential). Valluri & Merritt (1998) showed that the accuracy of the frequency analysis decreases significantly when orbits were integrated for less than 20 oscillation periods. Inaccurate frequency determination could result in misclassifying orbits as chaotic (since inaccurate frequency measurement can also lead to larger frequency drifts). We test the dependence of $\log(\Delta f)$ on the number of orbital periods n_p by plotting the frequency drift parameter against the largest orbital frequency (for orbits with $n_\mathrm{p} > 20$) in both NFW haloes in Fig. 2.[2] We use ω instead of n_p since $n_\mathrm{p} \propto \omega$, but is harder to compute accurately. Particles are binned in equal intervals in ω and the error bars represent the standard deviation in each bin. The straight lines are best fits to the data points. The slopes of the correlation for the $c = 10$ halo (solid line) and for the $c = 20$ halo (dot–dashed line) are both consistent with zero, indicating that $\log(\Delta f)$ is largely independent of ω (and

[2] We use the fractional change in the largest of the three fundamental frequencies measured over two contiguous time intervals (frequency drift) as a measure of chaos (Laskar 1990). For situations where a large fraction of orbits is resonant, it may be more appropriate to use the smallest of the three frequencies or the component with the largest amplitude.

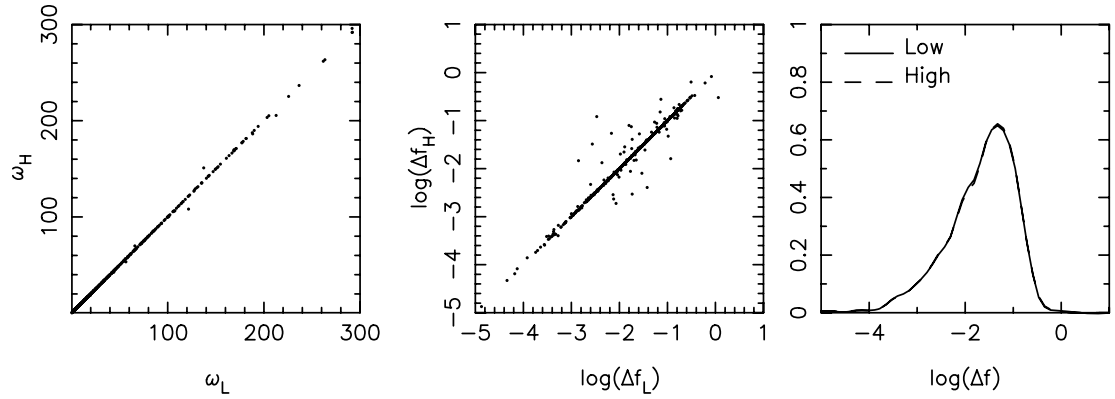

Figure 3. Left-hand panel: comparison of frequencies computed from the low and high time-sampling runs $\omega_{\rm H}$ versus $\omega_{\rm L}$, respectively; middle panel: comparison of diffusion parameter $\log(\Delta f)$ measured in low and high sampling runs; right-hand panel: histograms of diffusion parameter $\log(\Delta f)$.

hence of $n_{\rm p}$). Henceforth, we only use orbits which execute more than 20 orbital periods in the 50 Gyr over which they are integrated. The excluded orbits lie predominantly at large radii and are not significantly influenced by the changes in the inner halo that are investigated here. This rejection criterion affects about 25 per cent of the orbits in the triaxial DM haloes that we consider later.

The effect of central concentration on the accuracy of frequency estimation is of particular concern during *phase b*, when the potential is deepened due to the growth of a baryonic component. In this phase, frequencies of those orbits which are strongly influenced by the deepened potential are increased. Consequently, some orbits execute many more orbital periods during *phase b* than they do in *phase a* or *c*. However, we have fixed the orbital sampling time period (*not integration time-step*) to 1 Myr in all phases. In principle, coarse time sampling should not be a concern since the long integration time can still ensure a proper coverage of the phase-space torus. To ensure that the sampling frequency per orbital period does not significantly alter the frequency estimation, we resimulated one model (SA1) in *phase b* and stored the orbits five times more often (i.e. at time intervals of 0.2 Myr). We compared the frequencies of orbits computed for the low ($\omega_{\rm L}$) and high ($\omega_{\rm H}$) time-sampling runs. Fig. 3 (left-hand panel) shows that there is a strong correspondence between frequencies obtained with the two different samplings. We also found (Fig. 3, middle panel) that the frequency drift parameter $\log(\Delta f)$ obtained from the two runs are highly correlated although there is some increase in scatter for orbits with values of $\log(\Delta f) > -2$. Since the scattered points lie roughly uniformly above and below the 1:1 correlation line, there is no evidence that the higher sampling rate gives more accurate frequencies. The right-hand panel shows that the overall distribution of $\log(\Delta f)$ is identical for the two runs. We find that 95 per cent of the particles showed a frequency difference <0.1 per cent between the two different sampling rates. From these tests we conclude that our choice of sampling rate in the *phase b* is unlikely to significantly affect the frequency measurements of the majority of orbits. We therefore adopt the lower orbit sampling frequency for all the analysis that follows.

3.2 Orbit classification

Carpintero & Aguilar (1998, hereafter CA98) showed that once a frequency spectrum of an orbit is decomposed into its fundamental frequencies, the relationships between the values of the frequencies $(\omega_x, \omega_y, \omega_z)$ can be used to classify the orbits in a triaxial potential into the major orbit families as boxes, long (x) axis tubes and short (z) axis tubes. (CA98 point out that it is difficult to distinguish between the inner long-axis tubes (L-tubes) and outer L-tubes from their frequencies alone. Therefore, we do not attempt to distinguish between these two families with our automatic classification scheme.) In addition to classifying orbits into these three broad categories, they showed that if one or more of the fundamental frequencies is an integer linear combination of the other frequencies, the orbit can be shown to be resonant (either a periodic orbit or an open resonance). We followed the scheme outlined by CA98 to develop our own algorithm to classify orbits as boxes, L-tubes and S-tubes (abbreviated as S-tubes) and to also identify orbits that are associated with low-order resonances. We do not describe the classification scheme here since it is essentially identical to that described by CA98, the main difference lies in that we use NAFF to obtain the fundamental frequencies of orbits in the N-body model, whereas they used a method based on that of Binney & Spergel (1984). We tested our automated classification by visually classifying 60 orbits that were randomly selected from the different models. We then ran our automated orbit classifier on this sample, and compared our visual classification with that resulting from the automated classifier. The two methods agreed for 58/60 orbits (a 96 per cent accuracy rate assuming that the visual classification is perfectly accurate). Hereafter, we assume that our automated classification is accurate 96 per cent of the time and therefore any orbit fractions quoted have an error of ± 4 per cent.

3.3 Quantifying orbital shapes

In any self-consistent potential, the distribution of shapes of the majority of the orbits match the overall shape of the density profile. The elongation along the major axis is provided either by box orbits or by inner L-tubes. The ratios of the fundamental frequencies of orbits can be used to characterize their overall shape. Consider a triaxial potential in which the semimajor axis (along the x-axis) has a length a_x, the semi-intermediate axis has length a_y and the semiminor axis has a length a_z. The fact that $a_x > a_y > a_z$ implies that the oscillation frequencies along these axes are $|\omega_x| < |\omega_y| < |\omega_z|$ for any (non-resonant) orbit with the same overall shape as the density distribution (we consider only the absolute values of the frequencies since their signs only signify the sense of oscillation). We can use this property to define an average 'orbit shape parameter' ($\chi_{\rm s}$) for any orbit. For an orbit whose overall shape matches the

shape of the potential

$$|\omega_z| > |\omega_y| > |\omega_x|$$
$$\Rightarrow \frac{|\omega_y|}{|\omega_z|} > \frac{|\omega_x|}{|\omega_z|}$$
$$\chi_s \equiv \frac{|\omega_y|}{|\omega_z|} - \frac{|\omega_x|}{|\omega_z|} > 0. \quad (5)$$

The orbit shape parameter χ_s is positive for orbits with elongation along the figure. The larger the value of χ_s, the greater the degree of elongation along the major axis. Very close to the centre of the potential, it is possible for orbits to have greater extent along the y-axis than along the x-axis, as is sometimes the case with outer L-tubes. For such orbits, χ_s is slightly negative. An orbit for which all fundamental frequencies are almost equal would enclose a volume that is almost spherical. For such an orbit, $\chi_s \sim 0$ (which we refer to as 'round'). Note that orbits which are close to axisymmetric about the short (z) axis (i.e. the S-tubes) also have $\chi_s \sim 0$ because $\omega_x \sim \omega_y$ regardless of the value of ω_z. Our definition of shape parameter does not permit us to distinguish between truly round orbits for which $\omega_x \sim \omega_y \sim \omega_z$ and S-tubes, but both contribute to a more oblate axisymmetric potential.

4 RESULTS

For every model in Table 1, the three fundamental frequencies $\omega_x, \omega_y, \omega_z$ of each of the orbits in a selected subsample were computed separately in each of the three phases. For each orbit, the largest of the three fundamental frequencies is assumed to represent the dominant frequency of motion. The absolute value of this quantity is referred to as the largest fundamental frequency: we use $\omega_a, \omega_b, \omega_c$ to refer to the largest fundamental frequencies of an orbit in each of the three phases *a, b, c*. In addition to computing the fundamental frequencies over the entire 50 Gyr interval, we split the interval into two equal halves and computed the frequencies in each to compute the frequency drift parameter $\log(\Delta f)$ defined in Section 3. All orbits with $\log(\Delta f) < -1.0$ are identified as regular and the rest are identified as chaotic. For each orbit we also compute the total energy (E), the absolute value of the total specific angular momentum ($|j_{\rm tot}|$), the number of orbital periods ($n_{\rm p}$), and the pericentre and apocentre distance from the centre of the potential ($r_{\rm peri}, r_{\rm apo}$).

In this section, we consider results of five simulations, SA1 (*Triax + Disc*), P$_l$A3 (*Triax + Bulg*), P$_l$A4 (*Triax + hardpt*), P$_f$B2 (*Prolt + Ellip*) and P$_l$B3 (*Prolt + hardpt*). We shall show that haloes A and B have very different initial and final orbital properties despite the fact that their shapes in the presence of baryonic components are very similar (D08).

4.1 Distributions of orbital frequencies

The frequency distribution of randomly selected orbits in a triaxial halo can be used to characterize the orbital structure of phase space. It is useful to begin by discussing our expectations for how orbital frequencies change in response to growth of a central baryonic component. The potential is significantly deeper in *phase b* compared to *phase a*, consequently the most tightly bound orbits in the initial potential increase their frequencies. In contrast, orbits which largely lie outside the central mass concentration do not experience much deformation or much change in their frequencies. The higher the initial frequency the greater will be the frequency increase. Hence, we expect a faster-than-linear increase in frequency in *phase b* relative to the frequency in *phase a*.

When the baryonic component is evaporated, the halo expands once more and the haloes regain their triaxiality in models SA1, P$_l$A3 and P$_f$B2 but are irreversibly deformed in runs P$_l$A4 and P$_l$B3. One way to investigate the cause of the difference in these behaviours is to look for correlations between largest fundamental frequencies of each orbit in each of the three phases. When the growth of the baryonic component causes an adiabatic change in orbits, one expects that their frequencies ω_b will change in a regular (i.e. monotonic) way so that the particles which are deepest in the potential experience the greatest frequency increase. In Figs 4 and 5, we plot correlations between the frequencies ω_a, ω_b and ω_c.

In Figs 4, we show results for the three models whose baryonic scalelength is greater than 1 kpc: the left-hand panels show that ω_b increases faster-than-linearly with ω_a, as expected with fairly small scatter. The right-hand panels show that ω_c is quite tightly correlated with ω_a in all three models with the tightest correlation for simulation SA1 (the dashed line shows the 1:1 correlation between the two frequencies). The deviation from the dashed line and the scatter is only slightly larger in simulations P$_l$A3 and P$_f$B2. The strong correlation between ω_c and ω_a in these models supports the argument by D08 that the growth of the baryonic component resulted in regular rather than chaotic evolution. In all three models, only a small fraction of points deviate from the dashed line for the highest frequencies.

In contrast, the two models with a hard central point mass, P$_l$B3 and P$_l$A4 (Fig. 5), show a non-monotonic change in ω_b in response to the growth of the central point mass as well as a higher degree of scattering in frequency space. In particular, we note that orbits with small values of ω_a (i.e. those which are most weakly bound and have large apocentres) have the largest values of ω_b, which is in striking contrast to the situation in Fig. 4. We also see that ω_b sometimes decreased instead of increasing – again evidence for a scattering in frequency space rather than an adiabatic change. There is also greater scatter in the right-hand panels pointing to a less complete recovery in the frequencies ω_c after the baryonic component is evaporated. Thus, we see that when the central point mass is hard and compact there is significant orbit scattering.

It is clear from a comparison of Figs 4 and 5 that a baryonic component with a scalelength of $R_b \sim 1$ kpc or larger generally causes a regular adiabatic change in the potential while a hard point mass ($R_b \sim 0.1$ kpc) can produce significant chaotic scattering.

In both Figs 4 and 5, we see in the right-hand panels that $\omega_c < \omega_a$ especially at large values of ω_a (i.e. all points in the figures lie systematically below the line, indicating a decrease in the frequency ω_c). Thus, particles must have gained some energy implying that there has been a slight expansion in the DM distribution following the evaporation of the baryonic component.

What fraction of orbits experience a large fractional change in frequency of an orbit from *phase a* to *phase b*, and from *phase a* to *phase c*? To investigate this we define

$$\Delta\omega_{ab} = |(\omega_a - \omega_b)/\omega_a|, \quad (6)$$

$$\Delta\omega_{ac} = |(\omega_a - \omega_c)/\omega_a|. \quad (7)$$

The first quantity is a measure of the change in frequency distribution of orbits induced by the presence of the baryonic component, while the latter quantity measures the irreversibility of the evolution following the 'evaporation of the baryonic component'. In Fig. 6, we plot kernel density histograms of the distribution of the

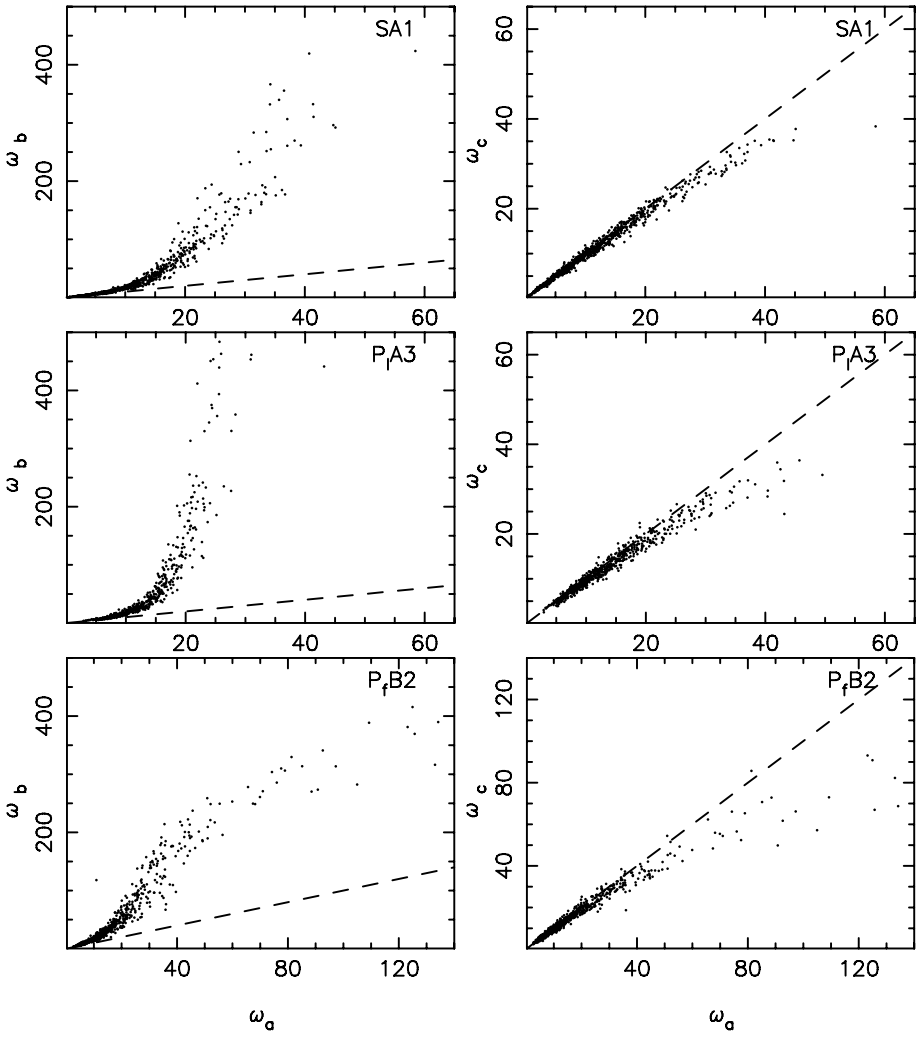

Figure 4. For the three models with extended baryonic components the left-hand panels show ω_b versus ω_a and right-hand panels show ω_c versus ω_a (frequencies in Gyr^{-1}). Dashed lines in each panel show the 1:1 correlation between each pair of frequencies. From top to bottom the models contain a baryonic disc (SA1), a spherical bulge (P$_l$A3) and a spherical elliptical (P$_f$B2).

frequency change $\Delta\omega_{ab}$ (left-hand panel) and $\Delta\omega_{ac}$ (right-hand panel) for orbits in all five models as indicated by the line legends. Each curve is normalized so that the area under it is unity.

The distribution of $\Delta\omega_{ab}$ is much wider for models P$_l$A3, P$_l$A4 and P$_l$B3 than for the other two models. In these three models, the scalelength of the baryonic component is $\leq$1 kpc and results in a broad distribution of $\Delta\omega_{ab}$, indicating that orbits over a wide range of frequencies experience significant frequency change. For model P$_l$A4 (dotted line), the histogram of values of $\Delta\omega_{ab}$ appears almost flat on the scale of this figure because it is spread out over a much larger range of abscissa values indicating that many more particles are significantly scattered in *phase b*. The full distribution for P$_l$A4 (see inset panel) is similar in form to P$_l$A3 and P$_l$B3.

The right-hand panels show that only a small number of orbits in models SA1, P$_l$A3 and P$_f$B2 experience an irreversible frequency change $\Delta\omega_{ac} > 20$ per cent, with the majority of particles experiencing less than 10 per cent. In contrast, in models P$_l$A4 and P$_l$B3, the distribution of $\Delta\omega_{ac}$ is much broader: a significant fraction of particles have experienced a large (20–50 per cent) permanent change in their frequencies, reflecting the fact that the models with a hard-compact point mass are the only ones which experience irreversible chaotic scattering.

Are there specific orbital characteristics that contribute to a large permanent frequency change, $\Delta\omega_{ac}$, between the two triaxial phases? We address this by determining how this quantity relates to other orbital properties. In Fig. 7, we plot $\Delta\omega_{ac}$ versus ω_a (left-hand panels) and versus ω_b (right-hand panels) for four of our models. In the top two panels (SA1 and P$_f$B2 – models with an extended baryonic component), there is no evidence of a dependence of frequency change on ω_b and only a slight increase in $\Delta\omega_{ac}$ at the highest values of ω_a (results for P$_l$A3 are not shown but are very similar to those for SA1).

On the other hand, the lower two panels (P$_l$A4 and P$_l$B3 – models with a compact hard baryonic component) show that there is a strong correlation between $\Delta\omega_{ac}$ and orbital frequency ω_a indicating that the orbits with the highest frequencies (ω_a) experience the largest frequency change, $\Delta\omega_{ac}$. This is evidence that scattering by the hard central point mass is greatest for particles that are most tightly bound and therefore closest to the central potential, confirming previous expectations (Gerhard & Binney 1985; Merritt & Valluri 1996). The absence of an appreciable correlation with ω_b is the consequence of scattering of orbits in frequency.

In Fig. 8, we plot $\Delta\omega_{ac}$ versus $r_{\rm peri}$ (left-hand panels) and versus $|j_{\rm tot}|$ (the total specific angular momentum of an orbit averaged

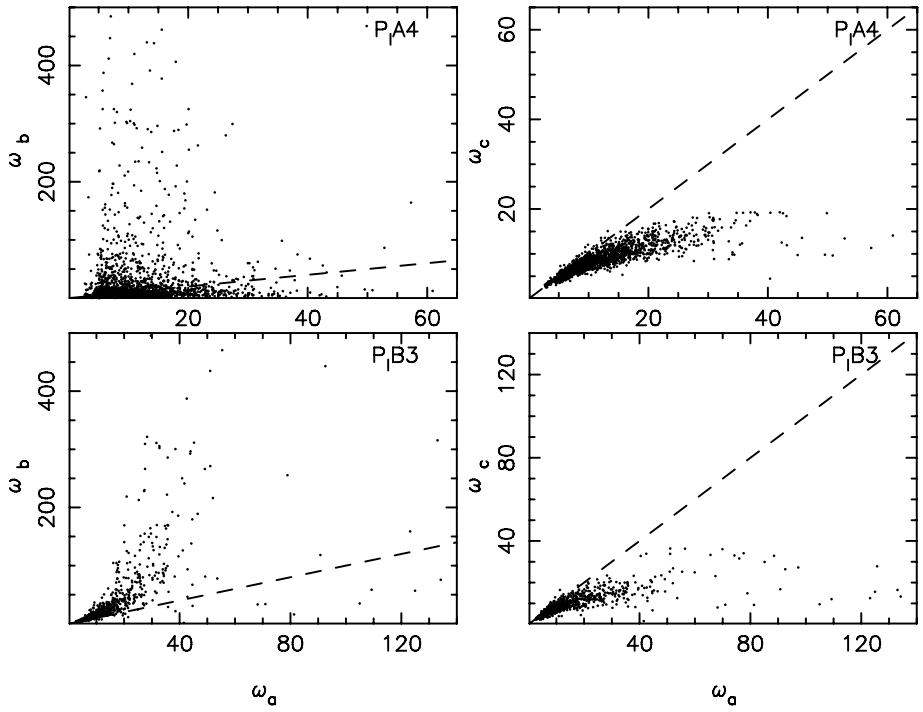

Figure 5. Same as in Fig. 4, but for the two models with hard point masses of 0.1 kpc softening. Top panels show the effect in the triaxial halo (dominated by box orbits) and bottom panel shows the effect on the prolate halo (dominated by L-tube orbits).

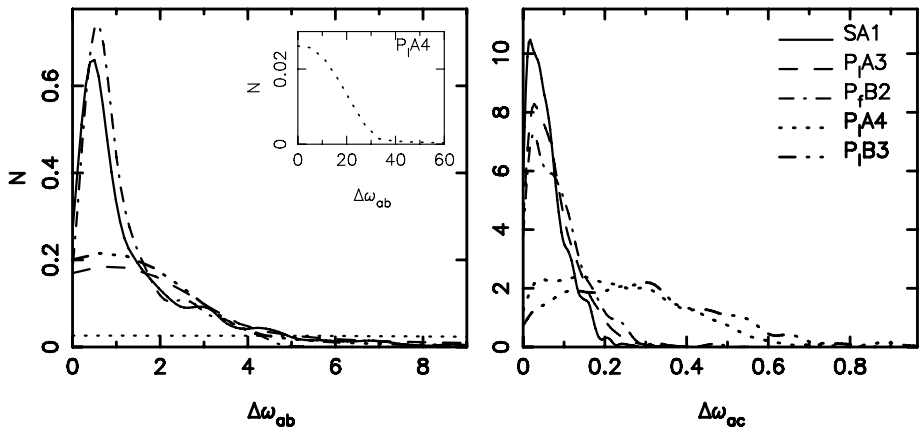

Figure 6. Kernel density distributions of $\Delta\omega_{ab}$ (left-hand panel) and of $\Delta\omega_{ac}$ (right-hand panel) for particles in the haloes SA1, P$_f$B2, P$_l$A3, P$_l$A4 and P$_l$B3 and as indicated by line legends. Each curve normalized to unit integral. The inset shows the full histogram for P$_l$A4 plotted on a different scale.

over its entire orbit in *phase a*) (right-hand panels) for orbits in the two models with compact central point mass (P$_l$A4 and P$_l$B3). (We do not show plots for models SA1, P$_l$A3 and P$_f$B2, because they show no correlation between $\Delta\omega_{ac}$ and either $|j_{\rm tot}|$ or $r_{\rm peri}$.) The left-hand panels of Fig. 8 shows that orbits which pass closest to the central point mass experience the most significant scattering. The absence of a correlation with $|j_{\rm tot}|$ however indicates that scattering *is independent of the angular momentum* of the orbit. In the next section, we show that halo A (the initial triaxial halo for model P$_l$A4) is dominated by box orbits while halo B is initially prolate, and is dominated by L-tubes which circulate about the long axis. Contrary to the prevailing view that centrophilic box orbits are more strongly scattered by a central point mass than centrophobic tube orbits, these figures provide striking evidence that the chaotic scattering is equally strong for the centrophobic L-tubes that dominate model P$_l$B3 as it is for the centrophilic box orbits that dominate model P$_l$A4. We return to a fuller discussion of the cause of this scattering in Section 5.

4.2 Changes in orbital classification

As we discussed in Section 3.2, relationships between the fundamental frequencies of a regular orbit can be used to classify it as a box orbit, a L-tube or a S-tube orbit. Quantifying the orbital composition of the two different haloes A and B and how their compositions change in response to the growth of a baryonic component yields further insight into the factors that lead to halo shape change. Orbits were first classified as regular or chaotic based on

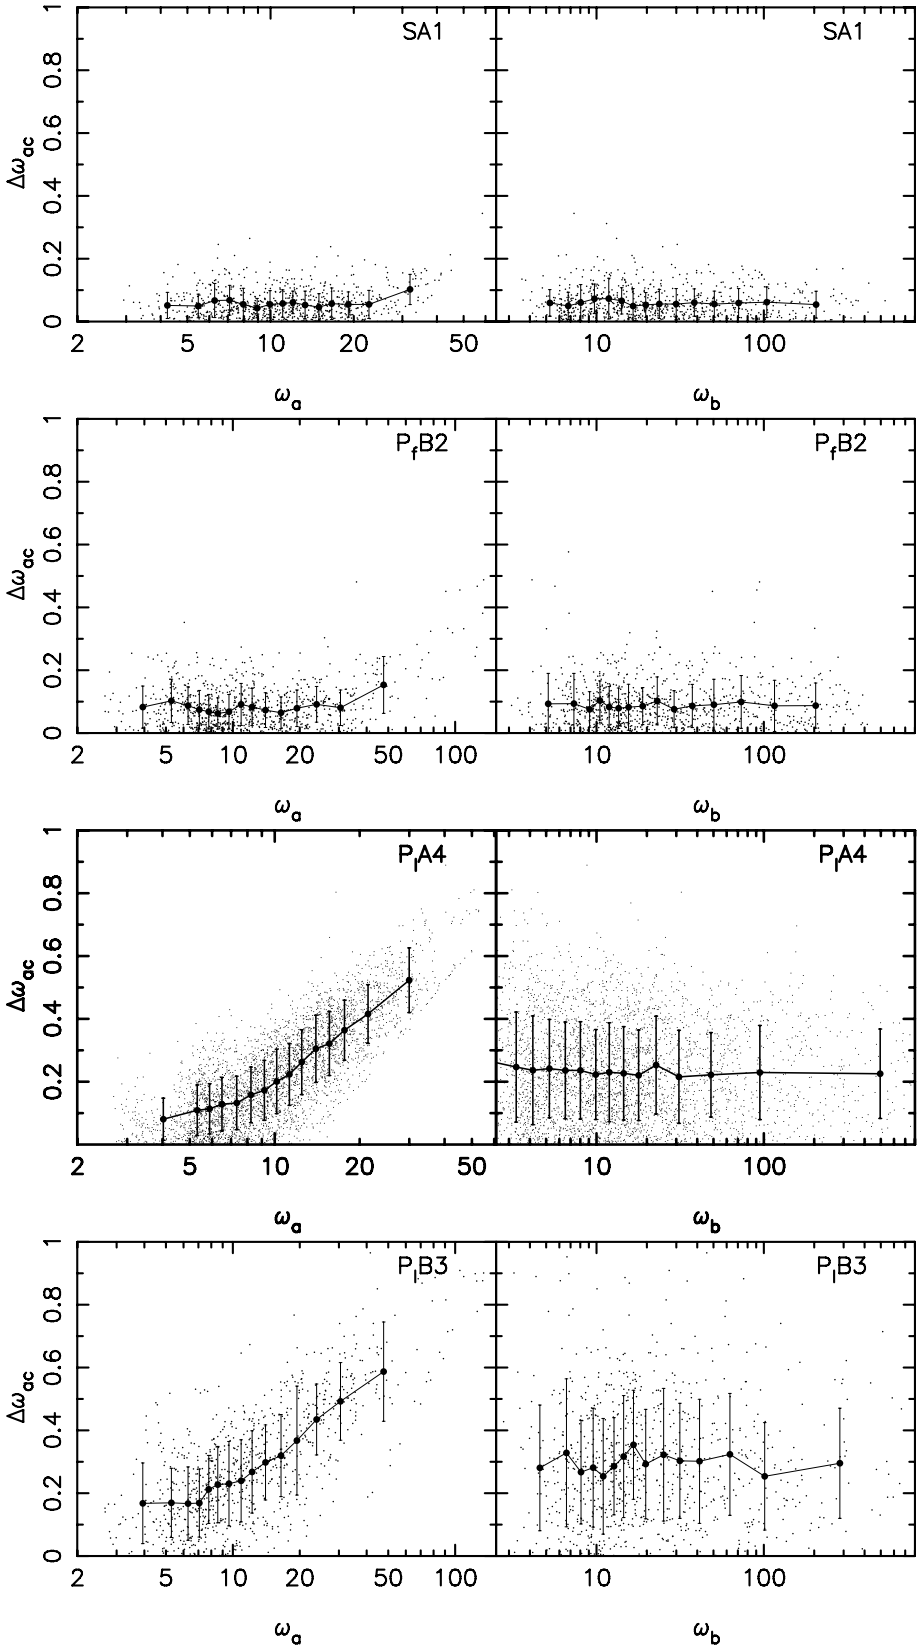

Figure 7. For four models, small dots show $\Delta\omega_{ac}$ versus ω_a (left) and $\Delta\omega_{ac}$ versus ω_b (right) for all particles analysed. Large solid dots with error bars show mean and standard deviation of particles in 15 bins in frequency. Top two panels are for models with an extended baryonic component (SA1, P$_f$B2), lower two panels show models (P$_l$A4, P$_l$B3) with a compact baryonic component.

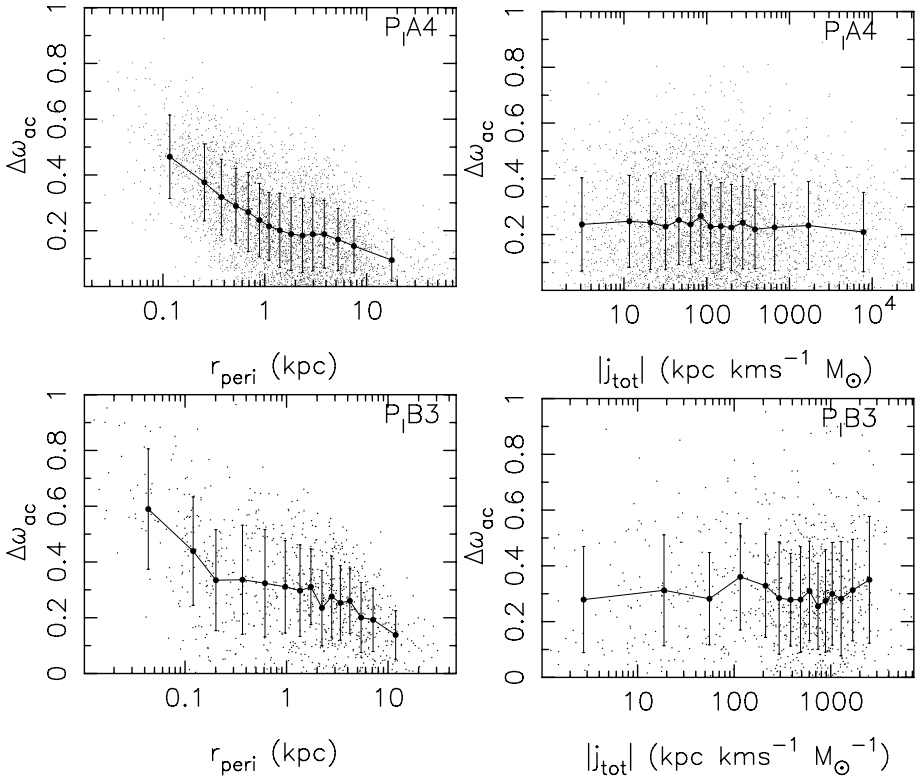

Figure 8. For model P$_l$A4 (top panels) and P$_l$B3 (bottom panels) left: $\Delta\omega_{ac}$ versus $r_{\rm peri}$; right: $\Delta\omega_{ac}$ versus $|j_{\rm tot}|$.

Table 2. Orbit composition of the models.

Type	Run SA1			Run P$_l$A3			Run P$_f$B2			Run P$_l$A4			Run P$_l$B3		
Phase	*a*	*b*	*c*	*a*	*b*	*c*	*a*	*b*	*c*	*a*	*b*	*c*	*a*	*b*	*c*
Boxes	0.86	0.43	0.83	0.84	0.16	0.76	0.15	0.09	0.29	0.84	0.17	0.80	0.15	0.03	0.21
L-tubes	0.11	0.09	0.12	0.12	0.43	0.15	0.78	0.75	0.54	0.12	0.35	0.11	0.78	0.78	0.59
S-tubes	0.02	0.27	0.03	0.02	0.33	0.06	0.07	0.09	0.16	0.02	0.26	0.04	0.07	0.11	0.14
Chaotic	0.01	0.21	0.02	0.02	0.08	0.03	0.00	0.07	0.01	0.02	0.21	0.05	0.00	0.08	0.06

Note. The numbers represent the fraction of orbits in each family.

their drift parameter $\log(\Delta f)$ as described in Section 3.1. Regular orbits were then classified into each of three orbital families using the classification scheme outlined in Section 3.2. The results of this orbit classification for each model, in each of the three phases, are given in Table 2.

The most striking difference between the initial triaxial models is that halo A (*phase a* of models SA1, P$_l$A3 and P$_l$A4) is dominated by box orbits (84–86 per cent) while halo B (*phase a* of models P$_f$B2 and P$_l$B3) is dominated by L-tubes (78 per cent). (The small differences between models SA1, P$_l$A3 and P$_l$A4 in *phase a* is purely a consequence of the selection of different subsets of orbits from halo A.) None of the initial models has a significant fraction of S-tubes or of chaotic orbits.

The very different orbit compositions of haloes A and B in *phase a* results in rather different evolutions of their orbital populations in response to the growth of a central baryonic component. Although the growth of the disc results in a significant decrease in the box orbit fraction (from 86 to 43 per cent) with boxes being converted to either S-tubes or becoming chaotic in *phase b*, model SA1 is highly reversible suggesting an adiabatic change in the potential. In model P$_l$A3 and P$_l$A4, the more compact spherical baryonic components decrease the fraction of box orbits even more dramatically (from 84 per cent down to 16–17 per cent), pointing to the vulnerability of box orbits to perturbation by a central component. Despite the similar changes in the orbital populations of the two models, P$_l$A3 is much more reversible than model P$_l$A4, indicating that both the shape of the central potential and its compactness play a role in converting box orbits to other families and that the change in orbit type is not evidence for chaotic scattering. It is striking that the more compact point mass in model P$_l$A4 results in significantly more chaotic orbits (21 per cent) compared to 8 per cent in P$_l$A3.

While halo A is initially dominated by box orbits, halo B is initially dominated by L-tubes, which dominate the orbit population in halo B in all three phases. The growth of the baryonic component in *phase b* causes the box orbit fraction to decrease (especially in model P$_l$B3) while the fraction of chaotic orbits increases slightly. The more extended point mass in P$_f$B2 causes a larger fraction of L-tubes to transform to orbits of another type than does the harder point mass in P$_l$B3, despite the fact that there is much greater scattering in the latter model. A comparison between the model P$_l$A4 and P$_l$B3 show that their orbit populations in the presence of a baryonic component differ significantly due to the different original orbit populations, while their degree of irreversibility is identical (e.g. Fig. 6) since the point mass in the two models is identical.

A significant fraction (21 per cent) of the orbits in *phase b* of model SA1 and P_lA4 are classified as chaotic [orbits with drift rate $\log(\Delta f) > -1.0$], in comparison with 9, 7 and 8 per cent in models P_lA3, P_fB2 and P_lB3, respectively. While the presence of such a large fraction of chaotic orbits in *phase b* of model P_lA4 may be anticipated from previous work, the high fraction of chaotic orbits in SA1 (*Triax* + *Disc*) is puzzling. To address concerns about classification error that could arise from errors in the accuracy of our frequency computation, we showed, in Fig. 3, that changing the frequency at which orbits were sampled by a factor of 5 did not result in any change in the overall distribution of $\log(\Delta f)$, and hence should not affect our classification of orbits as regular or chaotic. Another puzzling fact is that, although model SA1 in *phase b* has such a significant fraction of chaotic orbits, the orbit fractions essentially revert almost exactly to their original ratios once the disc is evaporated in *phase c*. Hence, the large fraction of chaotic orbits in *phase b* do not appear to cause much chaotic mixing. We will return to a more complete investigation of this issue in Section 4.4.

In Fig. 9, we investigate how orbits of different types (boxes, L-tubes, S-tubes, chaotic) are distributed with $r_{\rm peri}$, and how this distribution changes from *phase a* (black curves) to *phase b* (red curves). In *phase a*, the initially triaxial halo A models (black curves) are dominated by box orbits. The fraction of box orbits is significantly decreased in *phase b*. In particular, box orbits with large $r_{\rm peri}$ are transformed equally to S-tubes and chaotic orbits, while some box orbits at small $r_{\rm peri}$ are converted to L-tubes. In contrast, halo B models are dominated by L-tubes in both phases. Rather striking is how little the fraction of L-tubes in the halo B models changes, despite the fact that the haloes are significantly more oblate axisymmetric in *phase b* than in *phase a*. We saw in Fig. 6 that a significant fraction of orbits experience strong scattering that manifests as a change in their orbital frequencies, and in Fig. 8 we noted that the orbits with the smallest pericentre radii experience the largest change in frequency. In both models P_lA4 and P_lB3, the compact central point mass significantly reduced the box orbits. However, model P_lB3 only has a small fraction (15 per cent) of box orbits and it seems unlikely that the chaotic scattering of this small fraction of orbits off the central point mass is entirely responsible for driving the evolution of halo shape. Also P_fB2 which has a much more extended baryonic component shows a change in orbit population which closely parallels P_lB3 and we saw that P_fB2 is quite reversible and shows little evidence for chaotic scattering. It is clear (from Fig. 9) that in the prolate models (halo B) the majority of the orbits are L-tubes with large pericentre radii ($\langle r_{\rm peri}\rangle \sim 3$ kpc) and these remain L-tubes in *phase b*. How then do these prolate models evolve to more spherical models while retaining their dominant orbit populations? To address this question, we will now investigate the distribution of orbital shapes in each model at each phase of the evolution.

4.3 Changes in orbital shape

A parameter, χ_s, to quantify the shape of an orbit was defined in equation (3) of Section 3.3. Recall that this quantity is positive

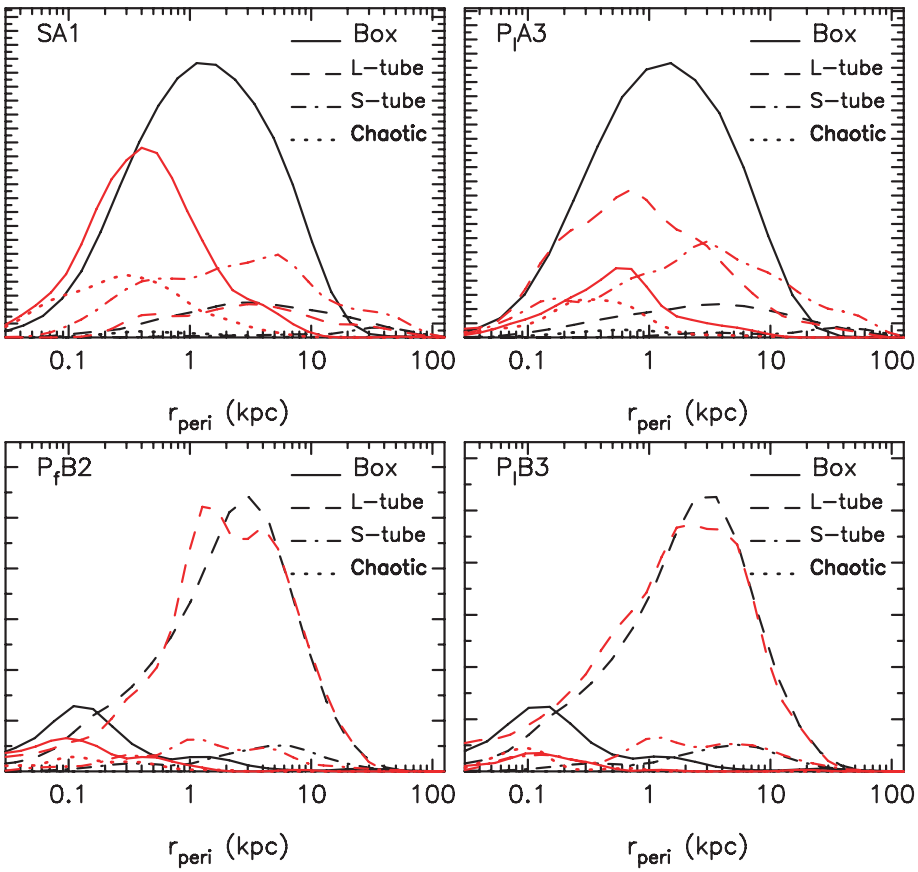

Figure 9. Distributions of $r_{\rm peri}$ for different orbit types. Distributions of each of the four different orbital types as indicated by the line legends. Distribution in *phase a* is given by black curves and distribution in *phase b* is shown by red curves. The integral under each curve is proportional to the number of orbits of that orbital type.

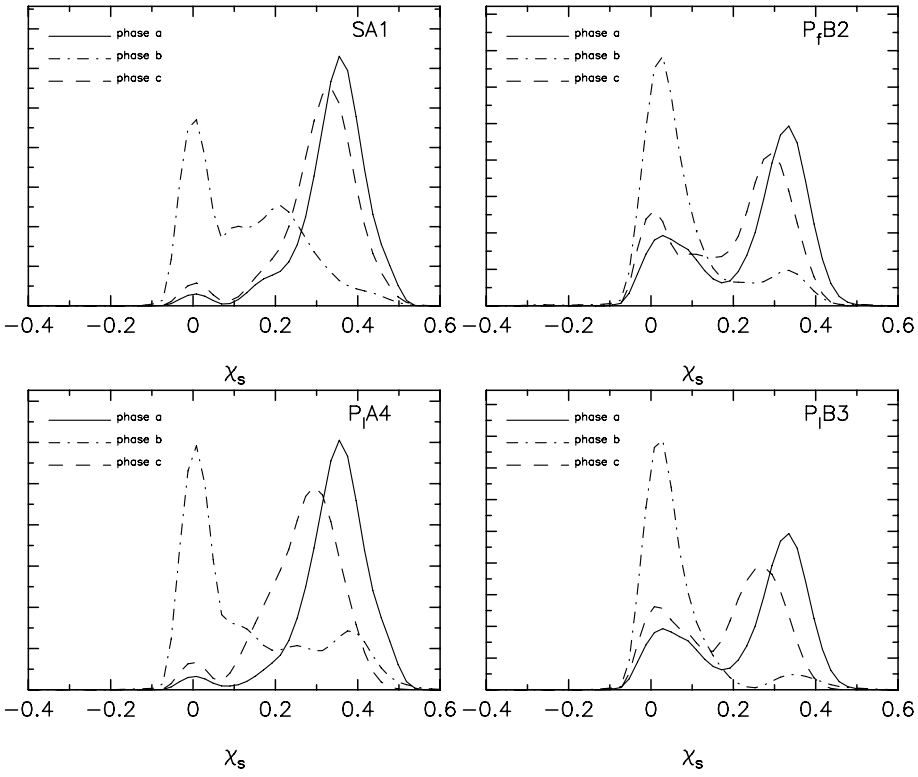

Figure 10. Kernel density histograms of the distribution of orbital shape parameter χ_s for each of the four models: SA1 (top left), P$_l$A4 (bottom left), P$_f$B2 (top right) and P$_l$B3 (bottom right) (P$_l$A3 is not shown since it is very similar to P$_l$A4). Distributions of χ_s in *phase a* are shown by solid curves, in *phase b* by dot–dashed curves and in *phase c* by dashed curves. In all models, a large fraction of orbits in *phase b* are 'round' ($\chi_s \simeq 0$).

when the orbit is elongated along the major axis of the triaxial figure, negative when elongated along the intermediate axis, and almost zero when the orbit is 'round' ($\omega_x \sim \omega_y \sim \omega_z$) or roughly axisymmetric about the minor axis ($\omega_x \sim \omega_y$). In Fig. 10, we show the shape distributions for the orbits in four of our five models. For each model, we show kernel density histograms for models in *phase a* (solid curves), *phase b* (dot–dashed curve) and *phase c* (dashed curves). In each plot, the curves are normalized such that the integral under each curve is unity. We define orbits to be elongated if $\chi_s \gtrsim 0.25$, and to be 'round' if $|\chi_s| \leq 0.1$.

Before the growth of the baryonic component (*phase a*: solid curves), the halo A models (left-hand panels SA1, P$_l$A4) have a distribution of orbital shapes that has a large peak at $\chi_s \sim 0.35$, arising from elongated orbits and a very small peak at $\chi_s \sim 0$ due to round orbits (model P$_l$A3 is not shown but is similar to P$_l$A4.) In halo B models, (right-hand panels P$_f$B2, P$_l$B3) on the other hand, the distribution of shapes is double peaked with about one third of all orbits contributing to the peak at $\chi_s \sim 0$. This implies one third of its orbits in the initially prolate halo B are 'round'. In both halo A and B, however, the larger of the two peaks has a value of $\chi_s \sim 0.35$ corresponding to quite elongated orbits. Despite the quite different underlying orbital distributions (halo A models dominated by box orbits while the halo B models are dominated by L-tubes). This illustrates that despite having different orbital compositions, a significant fraction of their orbits are similarly elongated.

The dot–dashed curves in all the panels show the distribution of orbital shapes in *phase b*. In all four models, there is a dramatic increase in the peak at $|\chi_s| \sim 0$, pointing to a large increase in the fraction of round (or S-tubes) at the expense of the elongated (L-tube or box) orbits. In the halo B models, the elongated orbits are significantly diminished indicating that the elongated L-tubes in *phase a* are easily deformed to 'round' orbits in *phase b* (most likely squat inner L-tubes). However, in model SA1, there is a large fraction of orbits with intermediate values of elongation $0.1 \leq \chi_s \leq 0.4$.

In *phase c* (dashed curves), all models show the dominant peak shifting back to quite high elongation values of $\chi_s \sim 0.3$ (although this is slightly lower than $\chi_s \sim 0.35$ in *phase a*). The downward shift in the peak is most evident in model P$_l$B3 (*Prolt + hardpt*), which as we saw before, exhibits the greatest irreversibility in shape. The scattering of a large fraction of the orbits by the hard central potential in model P$_l$B3 seen in Figs 7 and 8 is the major factor limiting reversibility of the potential. The smallest shift is for model SA1 (*Triax + Disc*), which exhibited the greatest reversibility.

We can also investigate how the shapes of orbits vary with pericentric radius. We expect that orbits closer to the central potential should become rounder ($\chi_s \to 0$) than orbits further out. We see that this expectation is borne out in Fig. 11 where we plot orbital shape parameter χ_s versus $r_{\rm peri}$ in both *phase a* (left-hand plots) and *phase b* (right-hand plots). In each plot, the dots show values for individual orbits. The solid curves show the mean of the distribution of points in each of 15 bins in $r_{\rm peri}$. Curves are only plotted if there are more than 30 particles in a particular orbital family (P$_l$A4 is not shown since it is similar to P$_l$A3). For models SA1 (*Triax + Disc*) and P$_l$A3 (*Triax + Bulg*), the figure confirms that elongated orbits in the initial halo A were box orbits (black dots and curves) and L-tubes (red dots and curves). The S-tubes (blue dots and curves) are primarily responsible for the 'round' population at $\chi_s \sim 0$. In *phase b* (right-hand panels) of both SA1 and P$_l$A3, there is a clear tendency for the elongated orbits (boxes, L-tubes and chaotic) to become rounder at small pericentre distances, but they continue to

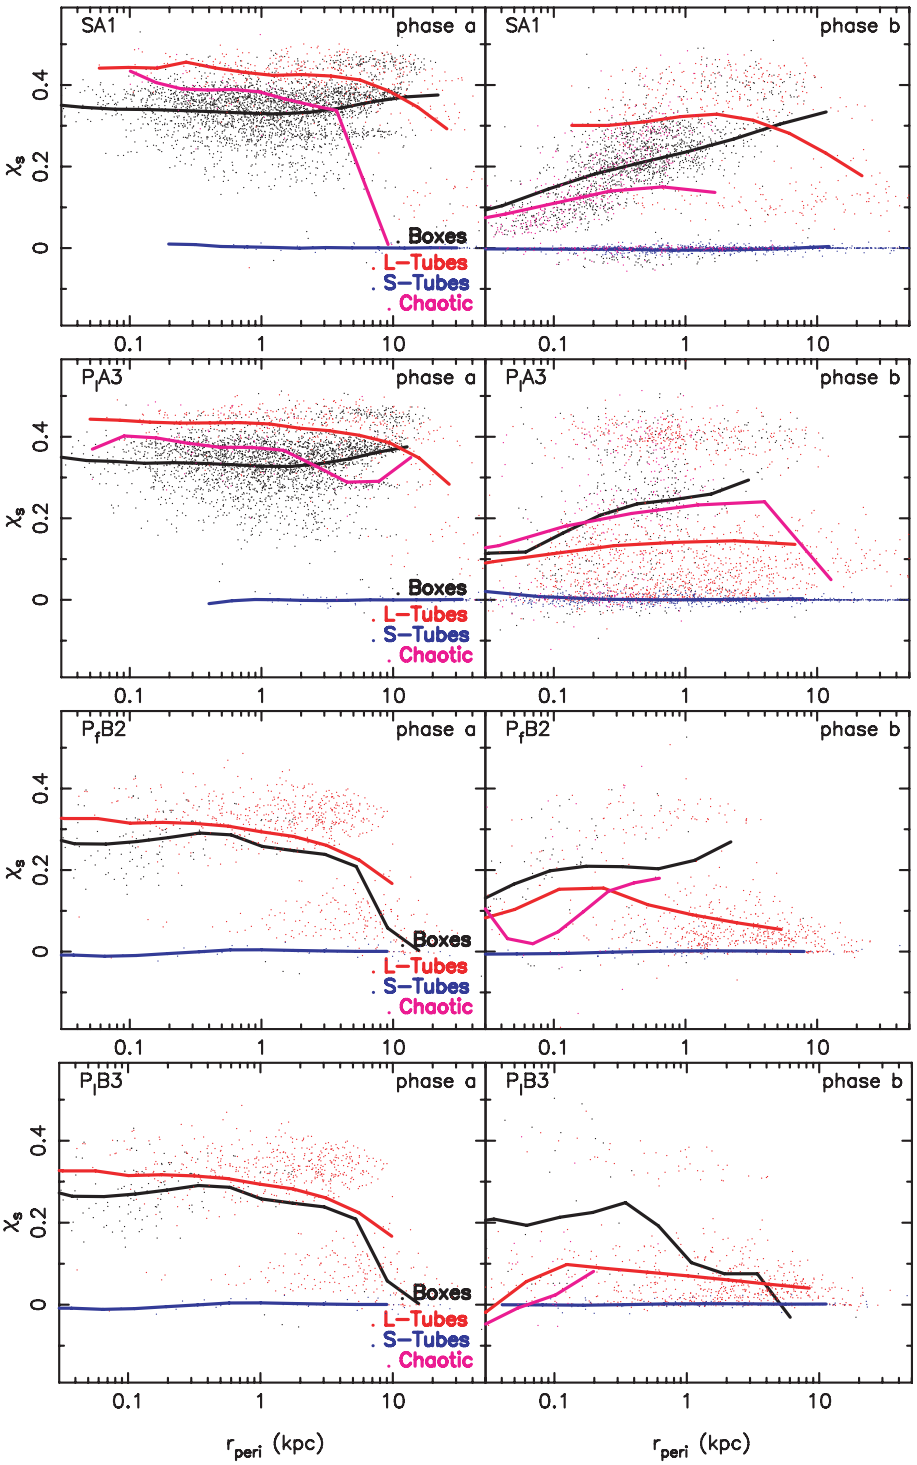

Figure 11. For models SA1, P_lA3, P_fB2 and P_lB3, the orbital shape parameter χ_s for each orbit is plotted against its pericentric radius $r_{\rm peri}$ as a small dot. The orbits of each of the four major orbital families are colour coded as in the figure legends. Left-hand panels are for *phase a* and right-hand panels are for *phase b*. The solid curves show the mean value of χ_s for all particles of that particular family, in 15 bins in $r_{\rm peri}$. Curves are not plotted if there are fewer than 30 orbits in a given orbital family. We used a kernel regression algorithm to smooth the curves.

be somewhat elongated at intermediate to large radii. Chaotic orbits in *phase b* of model SA1 appear to span the full range of pericentric radii and are not confined to small radii. (Note that the density of dots of a given colour is indicative of the number of orbits of a given type but the relative fractions are better judged from Fig. 9 and Table 2.)

For *phase a* in the models P_fB2 (*Prolt* + *Ellip*) and P_lB3 (*Prolt* + *hardpt*) (left-hand panels of each plot), boxes and L-tubes are elongated ($\chi_s \geq 0.25$), except at large $r_{\rm peri} \geq 8$ kpc where they become rounder. We see a trend for the average orbital shape (as indicated by the curves) in *phase b* to become rounder at small pericentre radii.

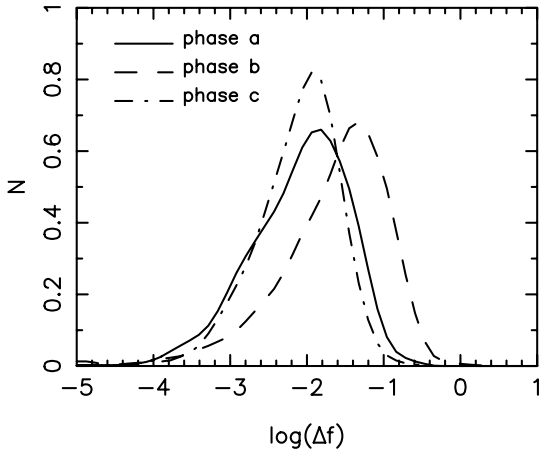

Figure 12. Histograms of frequency drift parameter $\log(\Delta f)$ for the three phases of model SA1 as indicated by the line legends.

Note that, in all the plots, the curves only show the average shape of orbits of a given type at any radius. The points show that in the case of the L-tubes in particular, the red dots tend to be distributed in two 'clouds': one with large elongations $\chi_s > 0.3$ and one with small elongation $\chi_s \sim 0.1$.

Thus, in all four models, it is clear that orbits that are elongated along the major axis of the triaxial potential in *phase a* become preferentially rounder at small pericentre radii in *phase b*. It is this change in orbital shape that plays the most significant role in causing the overall change in the shape of the density in the baryonic phase.[3]

4.4 Frequency maps and chaotic orbits

We saw, in Table 2, that *phase b* of model SA1 (*Triax + Disc*) and of model P_lA4 (*Triax + hardpt*) have a significant fraction (21 per cent) of chaotic orbits (i.e. orbits with $\log(\Delta f) > -1$). While P_lA4 shows significant lack of reversibility, which we can attribute to the presence of this high fraction of chaotic orbits, model SA1 does not show evidence for irreversibility.

Fig. 12 shows kernel density histograms of the chaotic drift parameter $\log(\Delta f)$ for orbits in each of the three phases in model SA1. It is obvious that, in *phases a* and *c*, there is only a small fraction of chaotic orbits [i.e. orbits with $\log(\Delta f) > -1$], whereas a much more significant fraction of orbits lie to the right of this value in *phase b*. Even the peak of the distribution in *phase b* is quite significantly shifted to higher drift values.

In this section, we investigate the surprising evidence that the chaotic orbits in *phase b* of model SA1 do not appear to mix. One possible reason for the lack of diffusion of the chaotic orbits is that the time-scale for evolution is not long enough. Indeed, D08 report that evolving run SA1 with the disc at full mass for an additional 5 Gyr after the growth of the disc is complete leads to a larger irreversible evolution (see their Fig. 3a). None the less, even in that case the irreversible evolution was only marginally larger than when the disc was evaporated right after it grew to full mass. Moreover, the growth time was 5 Gyr which means that the halo was exposed to a massive disc for a cosmologically long time.

A second possible reason for the lack of chaotic diffusion is that most of the chaotic orbits in this phase of the simulation are 'sticky'. The properties of 'sticky chaotic orbits' were discussed in Section 3. In a series of experiments designed to measure the rate of chaotic mixing, Merritt & Valluri (1996) showed that while ensembles of strongly chaotic orbits diffused and filled an equipotential surface on time-scales between 30–100 dynamical times, similar ensembles of 'sticky' or resonantly trapped orbits diffused much less quickly and only filled a small fraction of the allowed surface after very long times.

Laskar (1990) showed that frequency maps are a powerful way to identify resonances in dynamical systems. Frequency maps are obtained by plotting ratios of the three fundamental frequencies for each individual orbit. If a large and representative orbit population is selected, they can provide a map of the phase-space structure of the potential including all the resonances. Resonances appear as straight lines on the frequency map since their fundamental frequencies satisfy a condition like $l\omega_x + m\omega_y + n\omega_z = 0$. This method of mapping the phase space has the advantage that, since it only depends on the ratios of the frequencies and not on the frequencies themselves, it can be used to map phase space for large ensembles of particles without requiring them to be iso-energetic. This is a significant advantage over mapping schemes like Poincaré surfaces-of-section, when applied to an *N*-body simulation where particles, by design, are initialized to be smoothly distributed in energy. Thus, one can use the method to identify global resonances spanning a large range of orbital energies in *N*-body simulations.

In Fig. 13, we present frequency maps for four of the five models in *phase a* (left-hand panels) and *phase b* (right-hand panels) (P_fB2 is not shown since the frequency maps for this model are indistinguishable from those for P_lB3). For each orbit, the ratios of the fundamental frequencies ω_y/ω_z and ω_x/ω_z are plotted against each other. Particles are colour coded by their energy in *phase a*. The energy range in *phase a* was divided into three broad energy bins, with equal numbers of particles per bin. The most tightly bound particles are coloured blue, the least bound particles are coloured red and the intermediate energy range is coloured green.

Resonance lines are seen in the clustering of particles in all the maps. The most striking of the frequency maps is that for *phase b* of model SA1 (*Triax + Disc*). This map has significantly more prominent resonance lines, around which many points cluster, than any of the other maps. Three strong resonances and several weak resonances are clearly seen as prominent straight lines. The horizontal line at $\omega_y/\omega_z = 1$ corresponds to the family of orbits associated with the 1:1 closed (planar) orbit that circulates around the *x*-axis, namely the family of 'thin shell' L-tubes. The diagonal line running from the bottom left corner to the top right corner with a slope of unity ($\omega_y/\omega_z = \omega_x/\omega_z$) corresponds to the family of orbits that circulates about the *z*-axis: the family engendered by the 'thin shell' S-tubes. Since this latter family shares the symmetry axis of the disc, it is significantly strengthened in *phase b* by the growth of the disc. In addition to having many more orbits associated with it, this resonance extends over a much wider range in energy as evidenced by the colour segregation along the resonance line (blue points to the bottom left and red points at the top right). This segregation is the result of an increase in the gradient of the potential along the *z*-axis due to the growth of the disc, which results in an increase in ω_z. The more tightly bound a particle, the greater the increase in ω_z, and the greater the decrease in both its ordinate and abscissa. The most bound particles (blue points) therefore move away from their original positions towards the bottom left-hand corner of the plot. The least bound particles (red points) are farthest from the centre of the potential and these points experience the least displacement – although these points also shift slightly towards the resonance lines.

[3] Due to our chosen definition of shape parameter, S-tubes generally have $\chi_s \sim 0$ regardless of radius, because $\omega_x \sim \omega_y$.

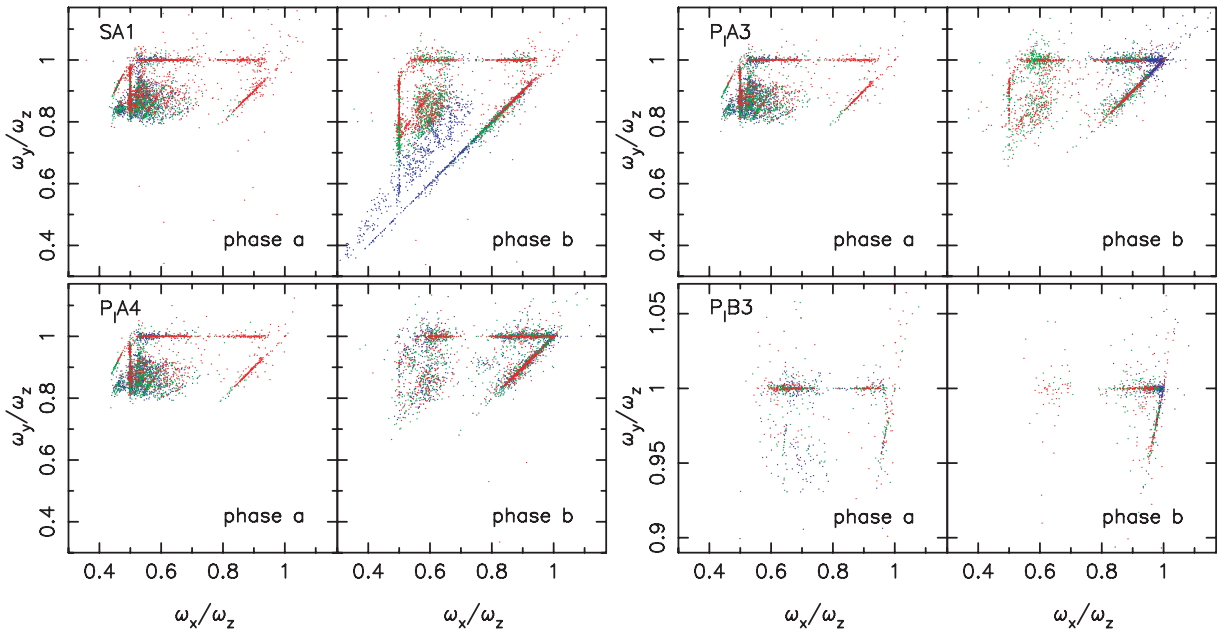

Figure 13. Frequency maps of particles in *phases a* and *b* for four models. For each particle, the ratio of the fundamental frequencies ω_y/ω_z is plotted versus ω_x/ω_z is plotted by a single dot. The dots are colour coded by the energy of the particle in *phase a*. The most tightly bound particles are coloured blue, and the least bound particles are coloured red. Model SA1 has 6000 particles, model P$_l$A4 and P$_l$A3 have 5000 particles, while P$_l$B3 has 1000 particles.

A third prominent resonance is the vertical line at $\omega_x/\omega_z = 0.5$ that corresponds to orbits associated with the family of banana (1:2 resonant) orbits. This banana (boxlet) resonance is also enhanced by the growth of the disc since this family of orbits, while not axisymmetric, is characterized by large excursions along the x-axis and smaller excursions in the z direction. Several shorter resonance lines are seen but are too sparsely populated in this plot to properly identify.

The frequency map for model P$_l$A3 in *phase b* shows that a spherical baryonic component produces a rather different phase-space structure than that produced by the disc. In particular, it is striking that the most tightly bound (blue) points are now clustered at the intersection of the horizontal and diagonal resonances namely around the closed period orbits 1:1:1. This may be understood as the consequence of the growth of the spherical baryonic point mass around which all orbits are rosettes and since no direction is preferred all orbits are 'round'. The 1:2 banana resonance is also less prominent in this model (largely because the deep central potential destabilizes this boxlet family).

The frequency map for model P$_l$A4 *phase b* shows the greatest degree of scattering, as evidenced by the thickest resonance lines. We attribute this to the large number of chaotic orbits in this model. Apart from the broad clustering of points around the diagonal (S-tube) and horizontal (L-tube) resonances there are no strong resonance lines seen in this map. Unlike the map for P$_l$A3 which shows a clustering of tightly bound (blue) points at the 1:1:1 periodic orbit resonances, the blue points are widely scattered in the frequency map of P$_l$A4.

The frequency maps for model P$_l$B3 shows that most of the orbits in this model are associated with the (1:1) L-tube family (horizontal line). A smaller number of orbits is associated with the 1:1 S-tube resonance (diagonal line). We saw previously that the growth of the baryonic components in this prolate halo caused little change in the orbit families. This is confirmed by the fact that the frequency maps in both phases are remarkably similar except for an increase in the clustering of points at the intersection of the horizontal (L-tube) and the diagonal (S-tube) resonance, which occurs for the same reason as in P$_l$A3. Since halo B is initially a highly prolate model, it has (as we saw previously) only a small fraction of box (and boxlet) orbits and in particular no banana orbits.

It is quite striking that in *phase a* the frequency maps show significantly less segregation by energy, and only a few resonances. This is because the initial triaxial models were generated out of mergers of spherical NFW haloes which were initially constructed so that orbits were smoothly distributed in phase space. The increase in the number of resonances following the growth of a baryonic component is one of the anticipated consequences of resonant trapping that occurs during the adiabatic change in a potential (e.g. Tremaine & Yu 2000; Binney & Tremaine 2008).

To test the conjecture that the majority of chaotic orbits in model SA1 (*phase b*) are resonantly trapped, we compute the number of chaotic orbits that lie close to a major resonance line. We define 'closeness' to the resonance by identifying those orbits whose frequency ratios lie $\pm\alpha$ of the resonant frequency ratio. For example, we consider an orbit to be close to the (1:1) L-tube resonance (horizontal line in map), if $|\omega_y/\omega_z - 1| \leq \alpha$. We find that the fraction of chaotic orbits in *phase b*, that lie close to one of the three major resonances identified above, is 51 per cent when $\alpha = 0.01$ and 62 per cent when $\alpha = 0.03$. Weaker resonances lines (which are hard to recognize due to the sparseness of the data points) may also trap some of the chaotic orbits. This supports our conjecture that the main reason that model SA1 does not evolve in *phase b*, despite the presence of a significant fraction of chaotic orbits, is that the majority of the chaotic orbits are trapped around resonances and therefore behave like regular orbits for very long times.

In Fig. 14, we plot four examples of chaotic orbits in *phase b* of SA1, which illustrate how resonantly trapped or 'sticky' chaotic orbits look. Each panel of four subplots shows a single orbit plotted in two Cartesian projections (side-by-side). The top pair of subplots show the orbit over the first 10 Gyr long time segment, while the

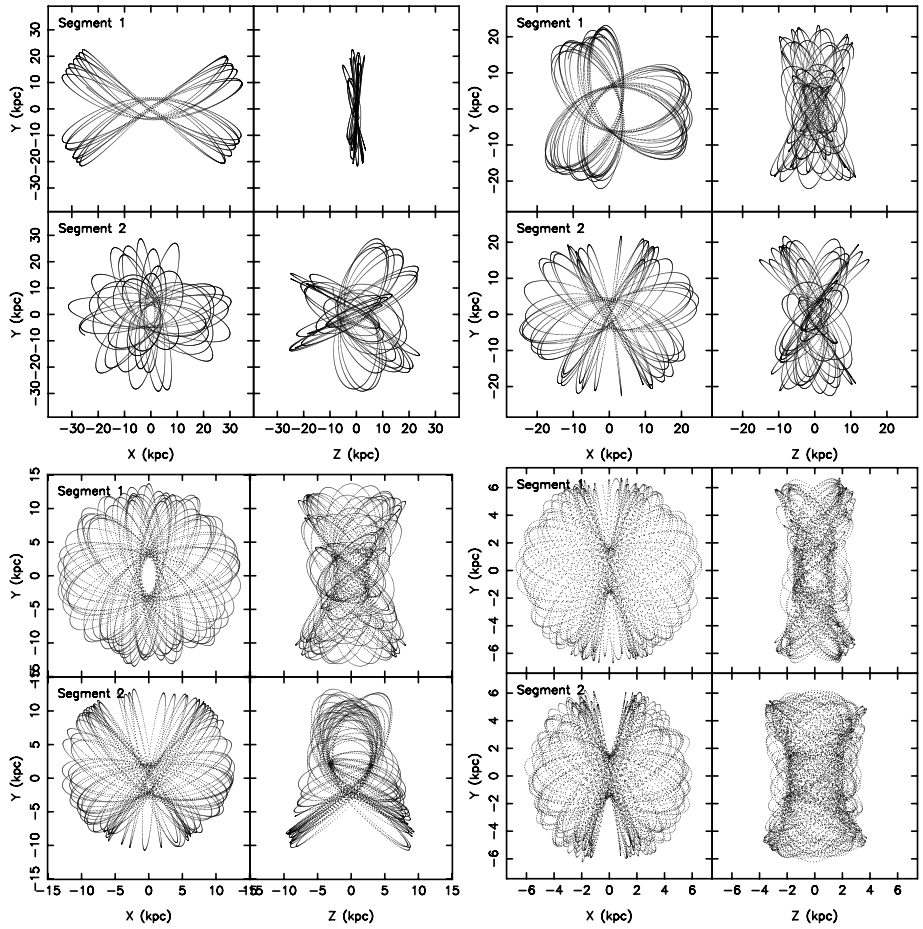

Figure 14. Several chaotic orbits in *phase b* of model SA1. Orbits are plotted in two Cartesian projections over two different time segments of 10 Gyr (top two plots in each panel show the first time segment, while bottom panels show the second time segment). See text for details.

bottom pair shows the same orbit over a second 10 Gyr time segment. The two time segments were separated by 10 Gyr. For illustration, we selected orbits with a range of drift parameters. The orbit in the top left-hand panel is an example of an orbit that conforms to our notion of a chaotic orbit that explores more phase space as time progresses, and has a large drift parameter of $\log(\Delta f) = -0.48$. The top right-hand panel shows a S-tube orbit that suddenly migrates to a box orbit [this orbit has a $\log(\Delta f) = -0.54$] and was probably in the separatrix region between the S-tube and box families. The bottom left-hand panel shows an orbit that is originally a S-tube that becomes trapped around a resonant boxlet ('fish') family [with $\log(\Delta f) = -0.66$], while the bottom right-hand panel shows a weakly chaotic box orbit [with $\log(\Delta f) = -0.94$]. Of the 21 per cent of orbits in *phase b* that are chaotic [$\log(\Delta f) \geq -1$], only $\sim$5 per cent have $\log(\Delta f) \geq -0.5$. This fraction is small enough that one does not expect it to result in significant chaotic mixing.

5 SUMMARY AND DISCUSSION

Since it was first proposed, the idea that a central black hole would scatter centrophilic box orbits in triaxial galaxies resulting in more axisymmetric potentials (Gerhard & Binney 1985) has frequently been used to explain the shape change in a variety of systems from the destruction of bars by central black holes (Norman, Sellwood & Hasan 1996) to the formation of more oblate galaxy clusters in simulations with gas (Kazantzidis et al. 2004a).

Experiments on chaotic mixing indicated that the time-scales for such mixing is about 30–100 dynamical times (Merritt & Valluri 1996), which is much longer than the time-scales for evolution of DM haloes in simulations with gas (Kazantzidis et al. 2004a). In addition, recent detailed studies of *N*-body simulations with controlled experiments have shown that the role of chaotic mixing may be less dramatic than conjectured by these previous studies. A study of relaxation of collisionless systems following the merger of two spherical galaxies showed that despite the fact that a large fraction of the orbits in a system undergoing violent relaxation are chaotic, the time-scales for chaotic diffusion and mixing are too long for this process to play a significant role (Valluri et al. 2007). In fact, even after violent relaxation, orbits retain strong memories of their initial energies and angular momenta.

D08 argued that since chaos introduces an irreversible mixing, numerical experiments in which evolution is driven by chaotic orbits should not be reversible. These authors studied the macroscopic shapes of triaxial DM haloes in response to the growth of a baryonic component. Unless the baryons were too centrally concentrated, or transported angular momentum to the halo, the evolution they saw was reversible, from which they concluded that much of the shape change arises from deformations in the shape of individual orbits rather than significant chaotic scattering. In this paper, we investigated this issue in significantly greater detail by applying the NAFF technique that allows us to quantify the degree to which chaotic diffusion drives evolution and to identify the primary physical

processes that cause halo shape change. The frequency-based method is able to distinguish between regular and chaotic orbits, making it more useful than Lyapunov exponents which are known to be sensitive to discretization effects in N-body systems (Hemsendorf & Merritt 2002). We use the method to quantify the drift in frequencies of large representative samples of orbits, thereby quantifying the degree of chaos in the systems we study. It also allowed us to map the phase-space structure of the initial and final haloes and to quantify the relationship between the change in the shapes of individual orbits and the shape of the halo as a whole.

Applying various analysis methods to orbits in five systems we demonstrated that the conclusion reached by D08 that chaos is not an important driver of shape evolution when the baryonic component is extended is indeed valid. As did D08, we also found that significant chaotic scattering does occur when the baryonic component is in the form of a hard central point mass (of scalelength $\sim$0.1 kpc). It is interesting that regardless of the original orbital composition of the triaxial or prolate halo, and regardless of the shape of radial scalelength of the baryonic component, haloes become more oblate following the growth of a baryonic component. Thus, two quite different processes (chaotic scattering and adiabatic deformation) result in similar final halo shapes even in haloes with very different orbital compositions.

We explored two different initial haloes, one in which box orbits were the dominant elongated population (halo A) and the other in which L-tubes dominated the initial halo (halo B). Despite the different orbit compositions both models exhibit similar overall evolution with regard to the shapes of orbits. In the halo A models, the box orbits were much more likely to change to either L-tubes or S-tubes, whereas in the halo B models, the dominant family of L-tubes largely retained their orbital classification while deforming their shapes.

Below we list the main results of this paper:

(i) Correlations between the orbital frequencies in the three different phases ω_a, ω_b and ω_c are a useful way to search for regular versus chaotic evolution of orbits. The orbital frequencies in the three phases are found to be strongly correlated with each other when the baryonic component is extended (Fig. 4), but show significant scattering when the baryonic component is a compact point mass (of scalelength about 0.1 kpc) (Fig. 5). In the more extended distributions, only a small fraction of the orbits experience significant change in their original orbital frequencies when the baryons are evaporated, while both the magnitude of scattering in orbital frequency as well as the fraction of orbits experiencing scattering, increases as the baryonic component becomes more compact (Fig. 6).

(ii) In the three models with relatively extended baryonic components, the change in orbital frequency between *phases c* and a ($\Delta\omega_{ac}$) is not correlated with orbital frequency (Fig. 7), pericentre distance or orbital angular momentum. When the baryonic component is a hard point mass, however, the frequency change is greater for orbits that are deeper in the potential and therefore have both a higher initial orbital frequency (Fig. 7) and smaller pericentric radius (Fig. 8). Scattering in frequency affects both the centrophilic box orbits as well as centrophobic L-tubes.

(iii) The growth of a baryonic component in halo A (either disc or softened point mass) causes box orbits with large pericentre radii to be converted to S-tubes, L-tubes or become chaotic (Fig. 9 top panel). While this change is almost completely reversible in the case of the disc or a diffuse point mass, it is less so when the baryonic component is a hard point mass. In halo B, which is dominated by L-tubes, the growth of the baryonic component causes almost no change in the orbital composition of the halo, indicating that the L-tubes are not destroyed but deformed (Fig. 9 bottom panel). Even though P_lB3 (*Prolt* + *hardpt*) is a model with significant orbit scattering by the hard central point mass, the process appears to mainly convert elongated inner L-tube orbits to somewhat rounder outer L-tubes. In model P_lA4 (*Triax* + *hardpt*), the box orbits are scattered on to S-tubes or chaotic orbits. The significant amount of scattering seen for even centrophobic L-tube orbits shows that the evolution is not due to direct scattering by a central point mass as sometimes assumed. Two alternative possibilities are more likely to account for the significant scattering in frequency. First the change in the symmetry and depth of the central potential is a perturbation to the potential that gives rise to an increase in the region of phase space occupied by resonances (Kandrup 1998 – private communication). As the resonances overlap, there is an increase in the degree of chaotic behaviour (Chirikov 1979). The second option is that the point mass attains equipartition with the background mass distribution, resulting in Brownian motion (Merritt 2005). The Brownian motion can cause the centre of the point mass to wander within a region of radius $\sim$0.1$-$1 kpc which can result in a significant change in the maximum gravitational force experienced by an orbit from one pericentre encounter to the next. This change in the maximum gravitational force manifests as scattering of the orbit which is equally effective for both box orbits and L-tubes. Indeed, a small wandering of the central massive point is seen in the N-body simulations of P_lA4; this motion is not included in our orbit calculations since all particles are frozen in place when calculating orbits. While it is beyond the scope of this paper to explore this issue further, we caution that the motion of the point mass in our N-body simulations is likely to overestimate the magnitude of the Brownian motion, since this depends on the mass resolution of the background particles. This suggests that in a real galaxy the evolution of the shape is much more likely to be driven by smooth adiabatic deformation of orbits than chaotic scattering.

(iv) In triaxial haloes, the orbital shapes sharply peaked distribution with the most elongated orbits ($\chi_s > 0.25$) are either boxes or L-tubes. In the prolate haloes, the second peak at $\chi_s \sim 0$ contains a third of the orbits and is composed of squat outer L-tubes and some box orbits. The growth of a baryonic component of any kind causes orbits of all types to become 'rounder', especially at small pericentre radii. This change in orbital shape distribution with radius is the primary cause of the change of halo shapes in response to the growth of a baryonic component. This is consistent with the findings of D08 who also found that the orbits in the models became quite round.

(v) The growth of a disc causes a large fraction of halo orbits to become resonantly trapped around major resonances. The three most important resonances are those associated with the 1:1 tube (thin shell) orbit that circulates about the short axis in the x–y plane, the 1:1 tube (thin shell) orbit that circulates about the long axis in the y–z plane and the 1:2 banana resonance in the x–z plane. We saw from the frequency maps that the resonant trapping of the halo particles depends both on the form of the baryonic component grown in the halo as well as on the initial orbital population of the halo.

Thus, we conclude that the evolution of galaxy and halo shapes following the growth of a central component occurs primarily due to regular adiabatic deformation of orbital shapes in response to the changing central potential. Chaotic scattering of orbits may be important particularly for orbits with small pericentre radii but

only when the central point mass is extremely compact. Contrary to previous expectations, chaotic scattering is only slightly more effective for centrophilic box orbits than it is for centrophobic L-tubes. Boxes can be scattered on to both L- and S-tube orbits and a significant fraction becomes chaotic. When the compact central point mass scatters L-tubes as it does in a prolate halo, they are scattered on to other L-tube orbits rather than on to S-tube orbits. The strong chaotic scattering that we see on centrophobic L-tube orbits has not been previously anticipated.

An important implication of our analysis is that while the shapes of haloes (and by extension elliptical galaxies) become more oblate (especially at small radii), following the growth of a baryonic component, the majority of their orbits are not S-tubes as might be predicted from their shapes. Instead, our analysis shows that orbits prefer to maintain their orbital characteristics, and the majority of the orbits are those which would be generally found in triaxial galaxies. This is particularly important for studies of the internal dynamics of elliptical galaxies since the fact that their shapes appear nearly axisymmetric need not imply that their orbital structure is as simple as the structure of oblate elliptical galaxies. Modifying the shapes to slightly triaxial could result in significant changes in their orbit populations and consequently could affect both the inferred dynamical structure as well as the estimates of the masses components such as the supermassive black holes in these galaxies (van den Bosch & de Zeeuw 2009).

Finally, our finding that the growth of a stellar disc can result in a large fraction of halo orbits becoming trapped in resonances could have important implications for observational studies of the Milky Way's stellar halo. The computation expense of the orbit calculations forced us to restrict the size of the frequency map for model SA1 to 6000 particles. This is only a tiny fraction of the particles in the original simulation. Despite the smallness of the sample, the frequency maps (Fig. 13) show a rich resonant structure which implies that the particles (either stars or DM) in the stellar and DM haloes of our Galaxy, particularly those close to the plane of the disc, are likely to be associated with resonances, rather than being smoothly distributed in phase space (this is in addition to structures arising due to tidal destruction of dwarf satellites). Although significantly greater resolution is required to resolve such resonances than is currently available, this could have significant implications for detection of structures in current and upcoming surveys of the Milky Way such as SDSS-III (Segue) and *Gaia* (Perryman et al. 2001; Wilkinson et al. 2005) and in the on-going direct detection experiments which search for DM candidates.

ACKNOWLEDGMENTS

MV is supported by the University of Michigan and NSF grant AST-0908346. VPD thanks the University of Zürich for hospitality during part of this project. Support for one of these visits by Short Visit Grant #2442 within the framework of the ESF Research Networking Programme entitled 'Computational Astrophysics and Cosmology' is gratefully acknowledged. Support for a visit by MV to the University of Central Lancashire at an early stage of this project was made possible by a Livesey Grant held by VPD. All simulations in this paper were carried out at the Arctic Region Supercomputing Center. We thank the referee Fred Adams for his comments which helped improve this paper.

REFERENCES

Allgood B., Flores R. A., Primack J. R., Kravtsov A. V., Wechsler R. H., Faltenbacher A., Bullock J. S., 2006, MNRAS, 367, 1781
Andersen D. R., Bershady M. A., Sparke L. S., Gallagher J. S., Wilcots E. M., 2001, ApJ, 551, L131
Bailin J., Steinmetz M., 2005, ApJ, 627, 647
Banerjee A., Jog C. J., 2008, ApJ, 685, 254
Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, ApJ, 304, 15
Barnes J., Efstathiou G., 1987, ApJ, 319, 575
Barnes J., Hut P., 1986, Nat, 324, 446
Barnes E. I., Sellwood J. A., 2003, AJ, 125, 1164
Bartelmann M., Steinmetz M., Weiss A., 1995, A&A, 297, 1
Binney J., Spergel D., 1982, ApJ, 252, 308
Binney J., Spergel D., 1984, MNRAS, 206, 159
Binney J., Tremaine S., 2008, Galactic Dynamics, 2nd edn. Princeton Univ. Press, Princeton, NJ
Boozer A. H., 1982, Phys. Fluids, 25, 520
Buote D. A., Canizares C. R., 1994, ApJ, 427, 86
Buote D. A., Jeltema T. E., Canizares C. R., Garmire G. P., 2002, ApJ, 577, 183
Carpintero D. D., Aguilar L. A., 1998, MNRAS, 298, 1 (CA98)
Chirikov B. V., 1979, Phys. Rep., 52, 263
Debattista V. P., 2003, MNRAS, 342, 1194
Debattista V. P., Moore B., Quinn T., Kazantzidis S., Maas R., Mayer L., Read J., Stadel J., 2008, ApJ, 681, 1076 (D08)
Dehnen W., 2005, MNRAS, 360, 892
Diehl S., Statler T. S., 2007, ApJ, 668, 150
Dubinski J., 1994, ApJ, 431, 617
Dubinski J., Carlberg R. G., 1991, ApJ, 378, 496
Franx M., de Zeeuw T., 1992, ApJ, 392, L47
Franx M., van Gorkom J. H., de Zeeuw T., 1994, ApJ, 436, 642
Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, ApJ, 327, 507
Gerhard O. E., Binney J., 1985, MNRAS, 216, 467
Goodman J., Heggie D. C., Hut P., 1993, ApJ, 415, 715
Habib S., Kandrup H. E., Mahon M. E., 1997, ApJ, 480, 155
Hemsendorf M., Merritt D., 2002, ApJ, 580, 606
Holley-Bockelmann K., Mihos J. C., Sigurdsson S., Hernquist L., Norman C., 2002, ApJ, 567, 817
Huizinga J. E., van Albada T. S., 1992, MNRAS, 254, 677
Iodice E., Arnaboldi M., Bournaud F., Combes F., Sparke L. S., van Driel W., Capaccioli M., 2003, ApJ, 585, 730
Jing Y. P., Suto Y., 2002, ApJ, 574, 538
Kalapotharakos C., 2008, MNRAS, 389, 1709
Kandrup H. E., Novotny S. J., 2004, Celest. Mech. Dyn. Astron., 88, 1
Kandrup H. E., Sideris I. V., 2001, Phys. Rev. E, 64, 056209
Kandrup H. E., Sideris I. V., 2003, ApJ, 585, 244
Kandrup H. E., Siopis C., 2003, MNRAS, 345, 727
Kandrup H. E., Smith H. J., 1991, ApJ, 374, 255
Kazantzidis S., Kravtsov A. V., Zentner A. R., Allgood B., Nagai D., Moore B., 2004a, ApJ, 611, L73
Kazantzidis S., Magorrian J., Moore B., 2004b, ApJ, 601, 37
Kochanek C. S., 1995, ApJ, 445, 559
Koopmans L. V. E., de Bruyn A. G., Jackson N., 1998, MNRAS, 295, 534
Kuijken K., Tremaine S., 1994, ApJ, 421, 178
Kuo-Petravic G., Boozer A. H., Rome J. A., Fowler R. H., 1983, J. Comput. Phys., 51, 261
Laskar J., 1990, Icarus, 88, 266
Laskar J., 1993, Celest. Mech. Dyn. Astron., 56, 191
Lichtenberg A. J., Lieberman M. A., 1992, Regular and Chaotic Motion. Applied Mathematical Sciences. Springer, New York
Merritt D., 2005, ApJ, 628, 673
Merritt D., Quinlan G. D., 1998, ApJ, 498, 625
Merritt D., Valluri M., 1996, ApJ, 471, 82
Merritt D., Valluri M., 1999, AJ, 118, 1177
Miller R. H., 1964, ApJ, 140, 250
Moore B., Kazantzidis S., Diemand J., Stadel J., 2004, MNRAS, 354, 522
Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563
Norman C. A., Sellwood J. A., Hasan H., 1996, ApJ, 462, 114
Oguri M., Lee J., Suto Y., 2003, ApJ, 599, 7
Olling R. P., 1995, AJ, 110, 591
Olling R. P., 1996, AJ, 112, 481

Olling R. P., Merrifield M. R., 2000, MNRAS, 311, 361
Papaphilippou Y., Laskar J., 1996, A&A, 307, 427
Papaphilippou Y., Laskar J., 1998, A&A, 329, 451
Perryman M. A. C. et al., 2001, A&A, 369, 339
Sackett P. D., Sparke L. S., 1990, ApJ, 361, 408
Schoenmakers R. H. M., Franx M., de Zeeuw P. T., 1997, MNRAS, 292, 349
Schweizer F., Whitmore B. C., Rubin V. C., 1983, AJ, 88, 909
Stadel J. G., 2001, PhD thesis, University of Washington
Terzić B., Kandrup H. E., 2004, MNRAS, 347, 957
Tremaine S., Yu Q., 2000, MNRAS, 319, 1
Valluri M., Merritt D., 1998, ApJ, 506, 686
Valluri M., Merritt D., 2000, in Gurzadyan V. G., Ruffini R., eds, Adv. Ser. Astrophys. Cosmol. Vol. 10, The Chaotic Universe, Proc. the Second ICRA Network Workshop. World Scientific, Singapore, p. 229
Valluri M., Vass I. M., Kazantzidis S., Kravtsov A. V., Bohn C. L., 2007, ApJ, 658, 731
van den Bosch R. C. E., de Zeeuw P. T., 2009, MNRAS in press
Wilkinson M. I. et al., 2005, MNRAS, 359, 1306
Zemp M., Moore B., Stadel J., Carollo C. M., Madau P., 2008, MNRAS, 386, 1543

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