Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **457,** 2107–2121 (2016) doi:10.1093/mnras/stw106

A review of action estimation methods for galactic dynamics

Jason L. Sanders[1]★ and James Binney[2]

[1]*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK*
[2]*Rudolf Peierls Centre for Theoretical Physics, Oxford University, 1 Keble Road, Oxford, OX1 3NP, UK*

Accepted 2016 January 8. Received 2015 January 7; in original form 2015 November 20

ABSTRACT

We review the available methods for estimating actions, angles and frequencies of orbits in both axisymmetric and triaxial potentials. The methods are separated into two classes. Unless an orbit has been trapped by a resonance, convergent, or iterative, methods are able to recover the actions to arbitrarily high accuracy given sufficient computing time. Faster non-convergent methods rely on the potential being sufficiently close to a separable potential, and the accuracy of the action estimate cannot be improved through further computation. We critically compare the accuracy of the methods and the required computation time for a range of orbits in an axisymmetric multicomponent Galactic potential. We introduce a new method for estimating actions that builds on the adiabatic approximation of Schönrich & Binney and discuss the accuracy required for the actions, angles and frequencies using suitable distribution functions for the thin and thick discs, the stellar halo and a star stream. We conclude that for studies of the disc and smooth halo component of the Milky Way, the most suitable compromise between speed and accuracy is the Stäckel Fudge, whilst when studying streams the non-convergent methods do not offer sufficient accuracy and the most suitable method is computing the actions from an orbit integration via a generating function. All the software used in this study can be downloaded from https://github.com/jls713/tact.

Key words: methods: numerical – galaxies: kinematics and dynamics.

1 INTRODUCTION

Galaxies are complex dynamical systems in which stars and dark-matter particles move on orbits in the gravitational field that is generated by all the other stars and dark-matter particles. Early experiments with typical smooth, axisymmetric galactic potentials demonstrated that these orbits possess three integrals of motion (Ollongren 1962). By Jeans' theorem, the distribution function (DF) of an equilibrium galaxy model can be assumed to be a function of these integrals of motion. Any function of the integrals of motion is also an integral but action integrals stand out from the rest because only action integrals J_i can be complemented by canonically conjugate variables θ_i to form a complete system $(\boldsymbol{\theta}, \boldsymbol{J})$ of canonical coordinates. The actions $\boldsymbol{J}$ label orbits and each angle variable θ_i increases linearly in time at the rate Ω_i. For a discussion of the merits of these variables, we refer readers to Binney & Tremaine (2008).

It has been demonstrated that DFs that are analytic functions of the actions can successfully model near-equilibrium structures in our Galaxy such as the disc (Binney 2010, 2012b; Bovy & Rix 2013; Piffl et al. 2014; Sanders & Binney 2015b), the stellar halo (Williams & Evans 2015; Das & Binney, in preparation) and the dark halo (Binney & Piffl 2015). Moreover, angle-action variables have proved useful for modelling structures that are not phase-mixed, such as tidal streams (Helmi & White 1999; Sanders & Binney 2013a; Bovy 2014; Sanders 2014) and substructures in the velocity space of solar-neighbourhood stars (Sellwood 2010; McMillan 2011).

Although orbits in typical potentials appear to admit three action integrals, the actions can be computed analytically only in a very few cases. Hence, practical use can be made of angle-action coordinates only to the extent that numerical methods enable us to compute angle-action coordinates from normal $(\boldsymbol{x}, \boldsymbol{v})$ coordinates. In recent years, and with the arrival of the *Gaia* data on the horizon, much effort has been invested in developing such methods. Many of these methods rely on, or are inspired by, the rare cases in which we can compute the actions analytically. This paper summarizes and collates these methods, and critically compares the various approaches. We seek to guide and advise readers on the best methods for approaching different types of data. In addition, we have made available code for all the approaches detailed in this paper at https://github.com/jls713/tact.

The paper is organized as follows. In Section 2, we recall the cases in which the actions can be computed exactly. We then describe numerical methods for generic potentials under two headings:

★ E-mail: jls@ast.cam.ac.uk

'convergent' methods in Section 3 and 'non-convergent' methods in Section 4. One of these methods, the spheroidal adiabatic approximation (SAA), has not been previously published. In Section 5, we critically compare the methods under two headings: accuracy and computational cost. In Section 6, we compare the accuracy required in various applications to the accuracy achieved by the methods that deliver best value. The use of the presented methods in non-integrable Hamiltonians and the impact of resonant trapping are discussed. We also briefly discuss how changes of parameters affect the accuracy of one of the most powerful methods, and we outline the structure of the code we have made available. Section 7 sums up. An appendix gives functional forms of DFs that provide good models of the thin and thick discs, the stellar halo and a stellar stream.

2 SEPARABLE POTENTIALS

Angle-action coordinates are intimately linked to separable potentials. A separable potential is one for which the Hamilton–Jacobi equation can be solved by separation of variables. The separation process introduces two constants of motion I_2 and I_3 in addition to the energy E and yields an expression for each momentum $p_i(q_i; E, I_2, I_3)$ as a function of one coordinate q_i and the integrals of motion. Given these expressions, the actions $J_i = (2\pi)^{-1} \oint \mathrm{d}q_i\, p_i$ can be computed as functions $J_i(E, I_2, I_3)$ of the 'classical' integrals of motion.

2.1 Spherical potentials

The Hamilton–Jacobi equation for any spherical potential, $\Phi(r)$, is always separable. The Hamiltonian for a particle in a spherical potential is

$$H = \frac{1}{2}p_r^2 + \frac{L^2}{2r^2} + \Phi(r), \tag{1}$$

where $L = |\boldsymbol{L}|$ is the length of the angular momentum vector. The orbit is confined to the plane normal to $\boldsymbol{L}$, and in this plane the angular motion separates from the radial motion. The angular motion is quantified by the angular action L. The Hamiltonian (1) is one-dimensional, so the radial action is

$$J_r = \frac{1}{\pi}\int_{r_\mathrm{p}}^{r_\mathrm{a}} \mathrm{d}r \sqrt{2E - 2\Phi - \frac{L^2}{r^2}}, \tag{2}$$

where r_p is the pericentric radius and r_a the apocentric radius.

In a spherical potential, the third action is arbitrary in the sense that it defines the orientation of the orbital plane with respect to some coordinate system. It is convenient to take it to be $J_\phi = L_z$, i.e. the component of angular momentum along the z-axis. We have the freedom to use as new actions linear combinations of old actions and we use this freedom to use the set $\boldsymbol{J} = (J_r, L_z, L - |L_z|)$, which proves to be the analogue of the sets we will work with in the axisymmetric but non-spherical case.

2.1.1 Analytic cases

The integral (2) can rarely be evaluated analytically. An exception of great importance is the isochrone potential (Henon 1959)

$$\Phi(r) = -\frac{GM}{b + \sqrt{r^2 + b^2}}, \tag{3}$$

where M is the mass and b is a scale radius. In this potential, the actions and angles are analytic functions of $(\boldsymbol{x}, \boldsymbol{v})$. In the limit $b \to 0$, the isochrone tends to the Kepler potential, whilst in the limit $b \to \infty$ the isochrone tends to the spherical harmonic-oscillator potential. More generally, we have analytic expressions for angle-action coordinates as functions of Cartesian phase-space coordinates for any triaxial harmonic-oscillator potential

$$\Phi(\boldsymbol{x}) = \frac{1}{2}\sum_{i=1}^{3} \omega_i^2 {x_i}^2, \tag{4}$$

where ω_i are the oscillator's frequencies.

2.2 Stäckel potentials

The most general class of separable potentials is that of triaxial Stäckel potentials. Below we will discuss the subclass of axisymmetric Stäckel potentials, and spherical potentials may be considered to lie within this subclass. Triaxial Stäckel potentials are intrinsically linked to confocal ellipsoidal coordinates. Here we briefly detail the relevant properties of these coordinates.

2.2.1 Confocal ellipsoidal coordinates

The confocal ellipsoidal coordinates (λ, μ, ν) of the point with Cartesian coordinates (x, y, z) are defined to be the three roots $\nu \le \mu \le \lambda$ of the cubic in τ

$$\frac{x^2}{(\tau - a^2)} + \frac{y^2}{(\tau - b^2)} + \frac{z^2}{(\tau - c^2)} = 1. \tag{5}$$

Here a, b and c are constants that define the coordinate system. We adopt the convention that the x-axis is the potential's major axis, y is the intermediate axis and z is the minor axis. For this to be the case, we require $c^2 \le \nu \le b^2 \le \mu \le a^2 \le \lambda$. Surfaces of constant λ are ellipsoids, surfaces of constant μ are hyperboloids of one sheet (flared tubes of elliptical cross-section that surround the x-axis), and surfaces of constant ν are hyperboloids of two sheets that have their extremal points on the z-axis. In the plane $z = 0$, lines of constant λ are ellipses with foci at $y = \pm\Delta_1 \equiv \pm\sqrt{a^2 - b^2}$, whilst in the plane $x = 0$ lines of constant μ are ellipses with foci at $z = \pm\Delta_2 \equiv \pm\sqrt{b^2 - c^2}$. Adding the same constant to a^2, b^2 and c^2 simply adds that constant to each of λ, μ and ν, leaving invariant Δ_1 and Δ_2 and the shapes of the coordinate curves. One may exploit this degeneracy to set one of a, b or c to zero.

The generating function, S, of the canonical transformation between Cartesian (x, y, z, p_x, p_y, p_z) and ellipsoidal coordinates $(\lambda, \mu, \nu, p_\lambda, p_\mu, p_\nu)$ is

$$S(p_x, p_y, p_z, \lambda, \mu, \nu) = p_x x(\lambda, \mu, \nu) + p_y y(\lambda, \mu, \nu) + p_z z(\lambda, \mu, \nu). \tag{6}$$

The momentum conjugate to any ellipsoidal coordinate τ is $p_\tau = \partial S/\partial\tau$.

When we express the Hamiltonian as a function of ellipsoidal coordinates, we find

$$\begin{aligned} H &= \frac{1}{2}\left(p_x^2 + p_y^2 + p_z^2\right) + \Phi(x, y, z) \\ &= \frac{1}{2}\left(\frac{p_\lambda^2}{P_\lambda^2} + \frac{p_\mu^2}{P_\mu^2} + \frac{p_\nu^2}{P_\nu^2}\right) + \Phi(\lambda, \mu, \nu), \end{aligned} \tag{7}$$

where

$$P_\lambda^2 = \frac{(\lambda - \mu)(\lambda - \nu)}{4(\lambda - a^2)(\lambda - b^2)(\lambda - c^2)}, \tag{8}$$

and P_μ and P_ν are given by cyclic permutations of (λ, μ, ν).

2.2.2 *Triaxial Stäckel potentials*

The most general triaxial Stäckel potential is

$$\Phi_{\rm S}(\lambda,\mu,\nu) = \frac{F(\lambda)}{(\lambda-\mu)(\nu-\lambda)} + \frac{F(\mu)}{(\mu-\nu)(\lambda-\mu)} + \frac{F(\nu)}{(\nu-\lambda)(\mu-\nu)}. \quad (9)$$

$\Phi_{\rm S}$ is composed of three functions of one variable. Here we denote the three functions with the same letter, F, as their domains are distinct. Moreover, for $\Phi_{\rm S}$ to be finite at $\lambda=\mu=a^2$ and $\mu=\nu=b^2$, $F(\tau)$ must be continuous at $\tau=a^2$ and b^2. Solution of the Hamilton–Jacobi equation (de Zeeuw 1985) yields equations for the momenta of the form

$$2(\tau-a^2)(\tau-b^2)(\tau-c^2)p_\tau^2 = \tau^2 E - \tau A + B + F(\tau), \quad (10)$$

where A and B are separation constants. Given a phase-space point $(\boldsymbol{x}_0,\boldsymbol{v}_0)$, we find $\tau_0(\boldsymbol{x}_0,\boldsymbol{v}_0)$ and $p_{\tau 0}(\boldsymbol{x}_0,\boldsymbol{v}_0)$ from the coordinate transformations, and can then find the integrals A and B by solving equation (10) simultaneously for two coordinates, say $\tau=\mu$ and λ (see de Zeeuw 1985, for more details). These integrals are related to the classical integrals I_2 and I_3 in a simple way:

$$I_2 = \frac{a^4E - a^2A + B}{c^2-a^2},$$
$$I_3 = \frac{c^4E - c^2A + B}{a^2-c^2}. \quad (11)$$

As p_τ is a function of only τ, the actions are then given by the 1D integrals

$$J_\tau = \frac{2}{\pi}\int_{\tau_-}^{\tau_+} {\rm d}\tau\, |p_\tau(\tau)|, \quad (12)$$

where (τ_-,τ_+) are the roots of $p_\tau(\tau)=0$, which we locate with Brent's method. We compute the integral (and all similar integrals) using a 10-point Gauss–Legendre quadrature. de Zeeuw (1985) chose to define the actions in a triaxial Stäckel potential in this way such that the actions are continuous across the transitions from one orbit family to the next. However, for the loop orbits, the oscillation in one of the coordinates is covered twice. In the axisymmetric case, this means that the radial action J_λ is half the value given by the right-hand side of equation (12).

2.2.3 *Axisymmetric limits*

In the limit $a^2\to b^2$, the confocal ellipsoidal coordinate system reduces to a prolate spheroidal coordinate system. In this case, $\Delta_1\to 0$ and the μ coordinate is replaced by the familiar polar angle ϕ. The coordinate system is now described by a single pair of foci separated by $\Delta=\Delta_2=\sqrt{a^2-c^2}$. The most general potentials that are separable in these coordinate are the oblate Stäckel potentials

$$\Phi_{\rm S,obl} = -\frac{F(\lambda)-F(\nu)}{\lambda-\nu}. \quad (13)$$

The potential is defined by two functions of a single variable $F(\tau)$ that have distinct domains. In these potentials $I_2=\frac{1}{2}J_\phi^2$. Note that taking the limit $b^2\to c^2$ reduces the confocal ellipsoidal coordinate system to an oblate spheroidal coordinate system and the corresponding separable potentials are the prolate Stäckel potentials. There is convincing evidence that the potential of the Milky Way is oblate, so we will not discuss prolate potentials.

In the axisymmetric case, equation (10) yields $p_\tau(\tau)$ in the form

$$2(\tau-a^2)(\tau-c^2)p_\tau^2 = (\tau-c^2)E - \frac{(\tau-c^2)}{(\tau-a^2)}I_2 - I_3 + F(\tau). \quad (14)$$

From equation (13) we see that the derivative

$$\frac{\partial^2}{\partial\lambda\partial\nu}\left[(\lambda-\nu)\Phi_{\rm S,obl}\right] = 0.$$

This expression can be rewritten in terms of R and z derivatives as[1]

$$\Delta^2 = a^2 - c^2 = z^2 - R^2 + \left[3z\frac{\partial\Phi}{\partial R} - 3R\frac{\partial\Phi}{\partial z} + Rz\left(\frac{\partial^2\Phi}{\partial R^2} - \frac{\partial^2\Phi}{\partial z^2}\right)\right] \Big/ \frac{\partial^2\Phi}{\partial R\partial z}, \quad (15)$$

where we have dropped the subscript 'S,obl' for brevity. We will use this expression to find an approximate Δ for general axisymmetric potentials in Section 4.

3 CONVERGENT METHODS

In this and the following section, we present the available methods for estimating actions. We present each method briefly since full details can be found in cited papers. Whenever numerical integration of the equations of motion is required, we use an adaptive embedded Runge–Kutta Prince-Dormand (8,9) scheme provided in the GNU Scientific Library (Galassi et al. 2009) with a relative accuracy of 10^{-8}.

We split methods for estimating actions into two categories: *convergent* and *non-convergent*. In the absence of resonant trapping, convergent methods are such that the errors in the computed actions can be reduced to arbitrarily small values given sufficient computation. Non-convergent methods, by contrast, rely on an approximation that cannot be refined, so the error in the computed actions cannot be arbitrarily diminished. In this section, we discuss convergent methods.

These methods all involve numerical construction of the generating function that maps the analytic angle-action variables $(\boldsymbol{\theta}',\boldsymbol{J}')$ of a 'toy' potential such as the isochrone or harmonic oscillator, into the angle-action variables $(\boldsymbol{\theta},\boldsymbol{J})$ of the true potential.

3.1 Iterative torus construction (ItTC)

McGill & Binney (1990) introduced the technique of 'torus mapping' to obtain actions for a general axisymmetric potential. They used the isochrone potential as the toy potential, and showed that the required generating function can be written as

$$S(\boldsymbol{\theta}',\boldsymbol{J}) = \boldsymbol{\theta}'\cdot\boldsymbol{J} + 2\sum_{\boldsymbol{n}>\boldsymbol{0}} S_{\boldsymbol{n}}(\boldsymbol{J})\sin\boldsymbol{n}\cdot\boldsymbol{\theta}', \quad (16)$$

where the vector $\boldsymbol{n}$ has integer components: on account of the periodicity of the angle coordinates, the generating function is a Fourier series with coefficients $S_{\boldsymbol{n}}$. On account of the time-reversal symmetry of the Hamiltonian and assumed mirror symmetry of the potential around the plane $z=0$, we require only a sine series and only even

[1] This equation corresponds to equation 8 of Sanders (2012) except in that paper a minus sign is missing on the right-hand side of the equation.

values of n_z occur. From the oddness of $\sin(x)$, the sum over $\boldsymbol{n}$ can be restricted to only half of $\boldsymbol{n}$ space.

The relation between the toy and true actions is

$$\boldsymbol{J}' = \frac{\partial S}{\partial \boldsymbol{\theta}'} = \boldsymbol{J} + 2\sum_{\boldsymbol{n}>\boldsymbol{0}} \boldsymbol{n} S_{\boldsymbol{n}}(\boldsymbol{J}) \cos \boldsymbol{n}\cdot\boldsymbol{\theta}', \tag{17}$$

while the angles are related by

$$\boldsymbol{\theta} = \frac{\partial S}{\partial \boldsymbol{J}} = \boldsymbol{\theta}' + 2\sum_{\boldsymbol{n}>\boldsymbol{0}} \frac{\partial S_{\boldsymbol{n}}}{\partial \boldsymbol{J}}(\boldsymbol{J}) \sin \boldsymbol{n}\cdot\boldsymbol{\theta}'. \tag{18}$$

Given a true action $\boldsymbol{J}$ and a trial set of coefficients $S_{\boldsymbol{n}}$, for each point on a regular grid in toy angle $\boldsymbol{\theta}'$ the corresponding toy actions $\boldsymbol{J}'$ are found from equation (17). The coordinates $(\boldsymbol{x}, \boldsymbol{v})$ are then recovered from the toy potential's analytic formulae, and the true Hamiltonian is evaluated at this phase-space point. The Marquardt–Levenberg algorithm is used to adjust the $S_{\boldsymbol{n}}$ iteratively until the spread in energy at the sampled points is minimized. Binney & Kumar (1993) showed how the procedure could be extended to recover $\partial S_{\boldsymbol{n}}/\partial \boldsymbol{J}$ so the true angles can be recovered from equation (18). Kaasalainen & Binney (1994b) showed that torus mapping can also be used for potentials that admit more than one orbit family and introduced a new algorithm for the recovery of $\partial S_{\boldsymbol{n}}/\partial \boldsymbol{J}$.

This algorithm involves integrating short sections of the orbit starting from several points on the torus. For each point along each orbit section, the toy angles $\boldsymbol{\theta}'$ are known and the true angles $\boldsymbol{\theta}$ must satisfy $\boldsymbol{\theta} = \boldsymbol{\theta}(0) + \boldsymbol{\Omega} t$. Equating these values to those given by equation (18), we obtain a series of linear equations for $\boldsymbol{\theta}(0)$, $\boldsymbol{\Omega}$ and $\partial S_{\boldsymbol{n}}/\partial \boldsymbol{J}$:

$$\boldsymbol{\theta}(0) + \boldsymbol{\Omega} t = \boldsymbol{\theta}' + 2\sum_{\boldsymbol{n}>\boldsymbol{0}} \frac{\partial S_{\boldsymbol{n}}}{\partial \boldsymbol{J}}(\boldsymbol{J}) \sin \boldsymbol{n}\cdot\boldsymbol{\theta}'. \tag{19}$$

More equations are obtained than one has unknowns and they are solved in a least-squares sense by standard methods.

With the $S_{\boldsymbol{n}}$ and $\partial S_{\boldsymbol{n}}/\partial \boldsymbol{J}$ found, the corresponding $(\boldsymbol{x}, \boldsymbol{v})$ for any point $\boldsymbol{\theta}$ on the torus $\boldsymbol{J}$ can be found. This is ideal for many applications but for others one needs machinery to convert $(\boldsymbol{x}, \boldsymbol{v})$ to $(\boldsymbol{\theta}, \boldsymbol{J})$ (e.g. McMillan & Binney 2013). McMillan & Binney (2008) obtained $(\boldsymbol{\theta}, \boldsymbol{J})$ for given points $(\boldsymbol{x}, \boldsymbol{v})$ by repeatedly constructing tori until one was obtained that passed through the given phase-space point. Their procedure was slow but Sanders & Binney (2015a) proposed a much faster iterative procedure that combines torus mapping with one of the non-convergent methods described below. We call this method *iterative torus construction* (ItTC).

One first estimates the angle-actions $(\boldsymbol{\theta}_{\rm S}, \boldsymbol{J}_{\rm S})$ from $(\boldsymbol{x}, \boldsymbol{v})$ using the axisymmetric Stäckel Fudge of Section 4.3.1. A torus with actions $\boldsymbol{J}_{\rm S}$ is constructed and the point $(\boldsymbol{x}_{\rm S}, \boldsymbol{v}_{\rm S})$ on this torus that is nearest to the given point $(\boldsymbol{x}, \boldsymbol{v})$ is found by minimization of the distance

$$\eta = |\boldsymbol{\Omega}|^2(\boldsymbol{x}-\boldsymbol{x}_{\rm S})^2 + (\boldsymbol{v}-\boldsymbol{v}_{\rm S})^2,$$

with respect to the angles $\boldsymbol{\theta}$ using $\boldsymbol{\theta}_{\rm S}$ as an initial guess. $\boldsymbol{\Omega}$ is the frequency of the constructed torus. One now uses the Stäckel Fudge to estimate the angle-actions $(\boldsymbol{\theta}_{\rm P}, \boldsymbol{J}_{\rm P})$ of the point $(\boldsymbol{x}_{\rm S}, \boldsymbol{v}_{\rm S})$. One now has in $\Delta\boldsymbol{J} = \boldsymbol{J}_{\rm P} - \boldsymbol{J}_{\rm S}$ an estimate of the error in the actions produced by the Stäckel Fudge. Under the assumption that this error varies slowly in phase space, one's new estimate of the action is then $\boldsymbol{J}_{\rm S} - \Delta\boldsymbol{J} = 2\boldsymbol{J}_{\rm S} - \boldsymbol{J}_{\rm P}$.

The procedure can be repeated using this improved action estimate to construct another torus. Sanders & Binney (2015a) found that a single torus construction produced actions accurate to 0.01 per cent and further torus constructions only reduced this by a factor of 2. We decide that the algorithm has converged when the distance η falls below some threshold η^*. For our standard setup, we use $\eta^* = (0.1\,{\rm km\,s^{-1}})^2$, a maximum of five torus constructions and a relative error for each torus construction of $\Delta J/J = 10^{-3}$.

3.2 Generating function from orbit integration (O2GF)

Sanders & Binney (2014) proposed a method for constructing the generating function (equation 16) that is based on orbit integration rather than Marquardt–Levenberg minimization of the variance in the Hamiltonian over a trial torus. For brevity, we will refer to this as the O2GF method. One starts by computing $N_{\rm samp}$ phase-space points along an orbit integrated for a time $N_T T$, where T is the period of a circular orbit with the same energy. At each phase-space point, the actions and angles in a toy potential are computed. These toy angle-actions are used to set up the linear equations (17) and (19). Finally, the entire set of equations is solved for the unknowns: the true actions $\boldsymbol{J}$, the Fourier coefficients, $S_{\boldsymbol{n}}$, and their derivatives $\partial S_{\boldsymbol{n}}/\partial \boldsymbol{J}$ plus $\boldsymbol{\theta}(0)$ and $\boldsymbol{\Omega}$. The vectors included in the sums are limited to $\boldsymbol{n} < N_{\rm max}$.

In Sanders & Binney (2014), the parameters of the toy potential were chosen by minimizing the deviation of the toy Hamiltonian around the orbit. This procedure was computationally costly and sub-optimal in that it tended to produce very sharp variations of the toy Hamiltonian along the orbit that required a large number of Fourier coefficients to remove. Here we find the minimum and maximum radius of the orbit sample and require the radial force of the toy and true potentials to agree there. Since we only consider two toy parameters (the scale mass and radius of the isochrone), this procedure is sufficient.

Several criteria for a sufficient sampling of toy angle space were discussed by Sanders & Binney (2014). For each vector $\boldsymbol{n}$, we compute the minimum and maximum values of $\boldsymbol{n}\cdot\boldsymbol{\theta}'$ (where $\boldsymbol{\theta}'$ are unrolled continuous angles, i.e. not 2π periodic). If the difference between the minimum and maximum value of $\boldsymbol{n}\cdot\boldsymbol{\theta}'$ is less than 2π, we repeat the algorithm with $N_T \to 2N_T$. Secondly, if for any mode range$(\boldsymbol{n}\cdot\boldsymbol{\theta}')/N_T > \pi$, the density of the sampling is too low to constrain the mode $\boldsymbol{n}$ and we repeat the algorithm with $N_{\rm samp} \to 2N_{\rm samp}$. For our standard setup, we use $N_T = 8$, $N_{\rm samp} = 300$ and $N_{\rm max} = 8$.

Using this approach, actions can be found in a triaxial potential that supports several orbital families (Sanders & Binney 2014). A different toy potential is used for each class of orbit, so the class (box, long-axis loop, etc.) to which an orbit belongs must be established before the equations are set up. This can be done by determining whether the sign of the angular momentum components L_x and L_z changes over the orbit. If both components change sign, the orbit is a box orbit and a triaxial harmonic-oscillator potential is appropriate. If L_z does not change sign, the orbit is a short-axis loop and an isochrone potential with its symmetry axis along z is suitable. If L_x does not change sign, the orbit is a long-axis loop and an isochrone with symmetry axis along x is suitable.

3.3 Averaged generating function (AvGF)

As noted by Bovy (2014) and Fox (2014), averaging equation (17) over the toy angles yields an expression for the true actions:

$$\boldsymbol{J} = \int \frac{{\rm d}^3\boldsymbol{\theta}'}{(2\pi)^3}\left(\boldsymbol{J}' - 2\sum_{\boldsymbol{n}>\boldsymbol{0}} \boldsymbol{n} S_{\boldsymbol{n}}(\boldsymbol{J}) \cos \boldsymbol{n}\cdot\boldsymbol{\theta}'\right) = \int \frac{{\rm d}^3\boldsymbol{\theta}'}{(2\pi)^3}\,\boldsymbol{J}'. \tag{20}$$

The task now is to estimate the value of the integral on the right given values of $\boldsymbol{J}'$ at irregularly distributed points $\boldsymbol{\theta}'$. This can be done to reasonable accuracy provided the points $\boldsymbol{\theta}'$ provide good coverage of the basic cube of angle space. This 'AvGF' method avoids solving elaborate matrix equations, so it is simple to code. For the standard setup, we use the same parameters as for the O2GF method.

4 NON-CONVERGENT METHODS

The non-convergent methods below are all based on the simplicity with which the actions can be computed in the separable potentials of Section 2.1.1. The basic idea is to proceed as if the potential were of Stäckel form. As we require the actions for an individual orbit, it is not necessary for a Stäckel potential to provide a good global fit to the potential; it suffices for a Stäckel potential to provide a good fit over the region explored by the orbit. For orbits that stay close to the plane, we will see that the assumption of cylindrical separability works well, whilst for more vertically and radially extended orbits a better approximation is found to be separability in prolate spheroidal coordinates. For even more radially and vertically extended orbits, the assumption of separability over the orbit region breaks down, and convergent methods are required for decent accuracy.

The ideas presented here are connected to work on estimating a third integral I_3 for general axisymmetric potentials. Any third integral can be used as the argument of a DF although, as argued in the introduction and elsewhere (e.g. Binney & McMillan 2016), actions are the integrals of choice from several perspectives. Two approaches to the construction of I_3 have been tried: explicit fitting of a Stäckel potential to the real potential or using the real potential in formulae derived for Stäckel potentials.

Dejonghe & de Zeeuw (1988) presented a general formalism for fitting a Stäckel potential to a general axisymmetric potential either globally or locally, and Batsleer & Dejonghe (1994) and Famaey & Dejonghe (2003) fitted a Stäckel potential to the available constraints for our Galaxy. However, it seems the inner regions of the Galaxy cannot be accurately represented with a Stäckel potential. De Bruyne, Leeuwin & Dejonghe (2000) produced a series of locally fitted Stäckel potentials and found that I_3 varied by $\sim$10 per cent around an orbit.

Kent & de Zeeuw (1991) pioneered using the real potential in formulae derived for Stäckel potentials by considering the target potential along lines of constant prolate spheroidal coordinate. These authors proposed several methods for estimating I_3 based on this idea. For general disc orbits, they found that the most accurate technique was the 'least-squares method' that minimized the variation in the expression for I_3 in a Stäckel potential using an approximation for $F(\lambda) \approx -(\lambda - c^2)\Phi(R(\lambda, c^2), 0)$. More recently, Bienaymé, Robin & Famaey (2015) used what appears to be an identical approximation for $F(\lambda)$ to that used by Kent & de Zeeuw (1991) to demonstrate that a third integral for disc orbits in the Besançon Galaxy model is conserved to $\sim$1 per cent.

4.1 Cylindrical adiabatic approximation (CAA)

The cylindrical adiabatic approximation (CAA) was introduced by Binney (2010) to model the distribution of stars in the Galactic disc. He argued that since the vertical frequency Ω_z of a disc star is significantly larger than its radial frequency Ω_r, the potential that governs vertical oscillations may be considered to vary slowly as the star oscillates radially, with the consequence that J_z is adiabatically invariant. Binney & McMillan (2011) showed that a better approximation to orbits can be obtained by replacing J_ϕ by $|J_\phi| + J_z$ in the formula for the effective radial potential. Schönrich & Binney (2012) observed that conservation of energy requires that variation along the orbit in the energy E_z of vertical motion is balanced by variation in the energy of radial motion. This principle leads to a slightly different modification of the effective radial potential to that proposed by Binney & McMillan (2011). Here we present the CAA in this refined form.

Vertical motion at radius R is controlled by the potential $\Psi_z(z) = \Phi(R, z) - \Phi(R, 0)$, so the vertical energy is

$$E_z = \frac{1}{2}v_z^2 + \Psi_z(z) \tag{21}$$

and the vertical action is

$$J_z = \frac{2}{\pi}\int_0^{z_{\max}} \mathrm{d}z\, v_z, \tag{22}$$

where $z_{\max}$ is the height above the plane where v_z vanishes. On a grid in (R, E_z) we use these formulae to tabulate J_z, and by interpolation in this table can recover $E_z(J_z, R)$. The assumption that J_z is constant along the orbit then yields the variation of E_z with R. Conservation of radial plus vertical energy implies that the radial motion is governed by the effective potential

$$\Psi_R(R) = \Phi(R, 0) + \frac{J_\phi^2}{2R^2} + E_z(J_z, R) - E_z(J_z, R_c), \tag{23}$$

where R_c is the guiding-centre radius. Using this effective potential, we estimate the radial action as

$$J_r = \frac{1}{\pi}\int_{R_p}^{R_a} \mathrm{d}R\, v_R, \tag{24}$$

where R_p and R_a are the radii where the radial velocity, v_R, vanishes. For our standard setup, we use a linearly spaced grid of 100 points in R_0 and a linearly spaced grid of 100 points in E_z^2.

4.2 Spheroidal adiabatic approximation (SAA)

The CAA assumes that stars oscillate parallel to the z-axis with the consequence that orbits are bounded by cylinders R = constant. Orbits that move significant distances from the Galactic plane are much more nearly bounded by ellipses than straight vertical lines in the (R, z) plane (e.g. Binney & Tremaine 2008, fig. 3.27). Consequently, we now make the assumption that adiabatically invariant oscillations occur along the spheroidal surfaces of prolate spheroidal coordinates. We call this method the spheroidal adiabatic approximation (SAA).

For a general axisymmetric potential, $\Phi(R, z)$, the Hamiltonian in prolate spheroidal coordinates is (cf. equation 7),

$$H = \frac{1}{2}\left(\frac{p_\lambda^2}{P_\lambda^2} + \frac{p_\nu^2}{P_\nu^2} + \frac{J_\phi^2}{R^2(\lambda, \nu)}\right) + \Phi(\lambda, \nu), \tag{25}$$

where

$$\begin{aligned} P_\lambda^2 &= \frac{\lambda - \nu}{(\lambda - a^2)(\lambda - c^2)} \\ P_\nu^2 &= \frac{\nu - \lambda}{(\nu - a^2)(\nu - c^2)}. \end{aligned} \tag{26}$$

Now we assume that the 'vertical' motion follows an ellipse of constant λ such that the ν coordinate is determined by the potential

$$\Psi_\nu(\nu) = \frac{J_\phi^2}{2R^2(\lambda, \nu)} - \frac{J_\phi^2}{2R^2(\lambda, c^2)} + \Phi(\lambda, \nu) - \Phi(\lambda, c^2). \tag{27}$$

The energy of the ν oscillations is

$$E_\nu = \frac{p_\nu^2}{2P_\nu^2} + \Psi_\nu(\nu) \tag{28}$$

and the vertical action is

$$J_z = \frac{2}{\pi}\int_{c^2}^{\nu_+} \mathrm{d}\nu\, p_\nu = \frac{2}{\pi}\int_{c^2}^{\nu_+} \mathrm{d}\nu \sqrt{2P_\nu^2(\lambda,\nu)}\sqrt{E_\nu - \Psi_\nu(\nu)}, \tag{29}$$

where ν_+ is the root of $E_\nu = \Psi_\nu(\lambda, \nu)$. Now given a 6D phase-space point $(\boldsymbol{x}, \boldsymbol{v})$, we find the best prolate spheroidal coordinate system using equation (15), evaluate E_ν at this point and then obtain J_z by evaluating the integral of equation (29) along the curve of constant λ that passes through the given phase-space point.

J_r is determined from

$$J_r = \frac{1}{\pi}\int_{\lambda_-}^{\lambda_+} \mathrm{d}\lambda\, p_\lambda = \frac{1}{\pi}\int_{\lambda_-}^{\lambda_+} \mathrm{d}\lambda \sqrt{2P_\lambda^2(\lambda,c^2)}\sqrt{E_{\rm tot} - \Psi_\lambda(\lambda)}, \tag{30}$$

where λ_+ and λ_- are the roots of $E_{\rm tot} = \Psi_\lambda(\lambda)$ and $\Psi_\lambda(\lambda)$ is an effective radial potential. The latter is defined by requiring that the sum of the radial and vertical energies is conserved as E_ν varies along the orbit. On a grid in (R, J_ϕ, E_ν), where R is the radius at which the curve of constant λ cuts the plane, we tabulate J_z from equation (29). Then by interpolation we can recover E_ν from given values of (R, J_ϕ, J_z). Clearly, J_ϕ and J_z are constant along an orbit, so for an individual orbit E_ν is a function of $\lambda(R)$ alone. Hence, the radial effective potential

$$\Psi_\lambda(\lambda) = \Phi(\lambda, c^2) + \frac{J_\phi^2}{2R^2(\lambda,c^2)} + E_\nu(R = \sqrt{\lambda - a^2}, J_\phi, J_z) \tag{31}$$

is well defined. It is straightforward to show that with this definition of Ψ_λ, the orbit's energy is

$$E_{\rm tot} = \frac{p_\lambda^2}{2P_\lambda^2} + \Psi_\lambda(\lambda). \tag{32}$$

The grid on which J_z is evaluated is defined in terms of R rather than λ so that we can use different choices of Δ for each orbit.

An illustration of how this method differs from the CAA is shown in Fig. 1. We show an example thick disc orbit integrated in a realistic Galactic potential in the meridional plane. The blue line shows the line along which we integrate to determine J_z with the CAA, while the red line shows the line of constant λ along which we integrate with the SAA. The points at which they intersect give the $(R, \pm|z|)$ coordinates of the input phase-space point. The SAA clearly better captures the shape of the boundaries of the orbit.

As we will see in Section 5, the SAA is significantly more accurate than the CAA. However, on account of E_z now being a function of three variables (R, J_ϕ, J_z), the SAA takes slightly longer to use than the CAA. Fortunately, the tabulation of E_ν requires a very small number of J_ϕ values (here we use 10) over the required J_ϕ range. Note that the CAA is contained within the SAA in that the surfaces of constant λ and ν tend to those of constant R and z in the limit of very large interfocal distance. In this limit, the vertical action becomes independent of J_ϕ and the SAA tends to the CAA.

For our standard setup, we use a linearly spaced grid of 100 points in λ, a linearly spaced grid of 100 points in E_ν^2 and a 10-point linearly spaced grid in J_ϕ.

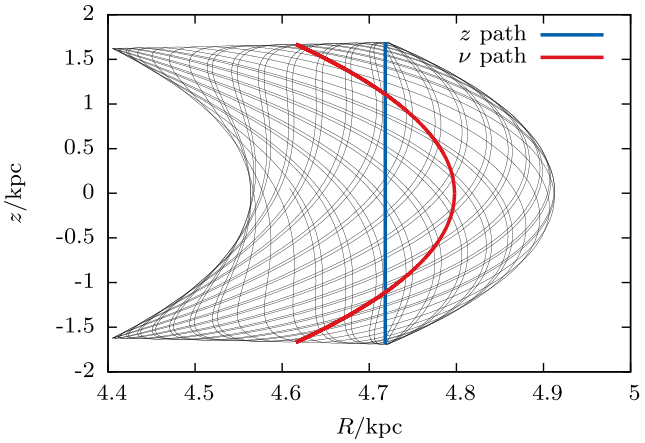

Figure 1. Illustration of the difference between the CAA and SAA. An example orbit is shown by the black line. The blue line shows the line along which the CAA vertical action integration is performed. The red line shows the line of constant λ along which the SAA vertical action integration is performed. The points at which they intersect give the $(R, \pm|z|)$ coordinates of the input phase-space point.

4.3 Stäckel Fudge

The essence of the adiabatic approximations is the definition of one-dimensional effective potentials that enable one to obtain J_z and J_r from quadratures. Binney (2012a) presented an alternative way of defining one-dimensional potentials, also using prolate spheroidal coordinates. Sanders & Binney (2015a) extended this idea to the triaxial case.

4.3.1 Axisymmetric case

For a general oblate axisymmetric potential, Φ, we define

$$\chi_\lambda(\lambda,\nu) \equiv -(\lambda - \nu)\Phi,$$
$$\chi_\nu(\lambda,\nu) \equiv (\lambda - \nu)\Phi. \tag{33}$$

If Φ were a Stäckel potential, these quantities would be given by

$$\chi_\lambda(\lambda,\nu) = F(\lambda) - F(\nu),$$
$$\chi_\nu(\lambda,\nu) = F(\nu) - F(\lambda). \tag{34}$$

Therefore, for a general potential, we can write

$$F(\tau) \approx \chi_\tau(\lambda,\nu) + D_\tau, \tag{35}$$

where D_τ are constants provided we evaluate χ_λ at constant ν and vice versa. We can write the equation for $p_\tau(\tau)$ (equation 14) as

$$2(\tau - a^2)(\tau - c^2)p_\tau^2 = (\tau - c^2)E - \left(\frac{\tau - c^2}{\tau - a^2}\right)\frac{J_\phi^2}{2} - B_\tau + \chi_\tau(\lambda,\nu), \tag{36}$$

where we have defined the integrals of motion $B_\tau = I_3 - D_\tau$. Given an initial phase-space point, we use equation (15) to find a suitable coordinate system, calculate λ, ν, p_λ and p_ν, and use equation (36) to find the integrals B_τ. The momentum p_τ from equation (36) is then integrated over an oscillation in τ to find each action as in equation (12). We note that for the λ integral, we keep ν fixed at the input value, and vice versa.

Binney (2014) used an alternative to equation (15) as a method of estimating Δ. At each energy and J_ϕ, a shell orbit $J_r = 0$ is computed and an ellipse is fitted to the orbit in the meridional

plane. The location of the focus of the ellipse gives an estimate of Δ. When Δ is determined from equation (15), we shall refer to the method as the Stäckel Fudge v1, and when Δ is determined from shell orbits as just described, we will call the method Stäckel Fudge v2. We test both methods below.

4.3.2 *Triaxial case*

The above procedure is simply generalized to triaxial potentials. Given a general triaxial potential, we define the quantities

$$\begin{aligned}\chi_\lambda(\lambda,\mu,\nu) &\equiv (\lambda-\mu)(\nu-\lambda)\Phi(\lambda,\mu,\nu),\\ \chi_\mu(\lambda,\mu,\nu) &\equiv (\mu-\nu)(\lambda-\mu)\Phi(\lambda,\mu,\nu),\\ \chi_\nu(\lambda,\mu,\nu) &\equiv (\nu-\lambda)(\mu-\nu)\Phi(\lambda,\mu,\nu).\end{aligned} \tag{37}$$

If Φ were a Stäckel potential, these quantities would be given by, for instance,

$$\chi_\lambda(\lambda,\mu,\nu) = F(\lambda) - \lambda\frac{F(\mu)-F(\nu)}{\mu-\nu} + \frac{\nu F(\mu)-\mu F(\nu)}{\mu-\nu}. \tag{38}$$

Therefore, for a general potential, we can write

$$F(\tau) \approx \chi_\tau(\lambda,\mu,\nu) + C_\tau\tau + D_\tau, \tag{39}$$

where C_τ and D_τ are constants provided we always evaluate χ_τ with two of the ellipsoidal coordinates fixed. For instance, we always evaluate χ_λ at fixed μ and ν.

When we substitute these expressions into equation (10), we find

$$2(\tau-a^2)(\tau-b^2)(\tau-c^2)p_\tau^2 = \tau^2 E - \tau A_\tau + B_\tau + \chi_\tau(\lambda,\mu,\nu). \tag{40}$$

For each coordinate τ, there are two new integrals of motion given by $A_\tau = a - C_\tau$ and $B_\tau = b + D_\tau$. Using a single 6D coordinate and a choice of coordinate system gives us a single constraint on a combination of A_τ and B_τ. Due to the separability of Stäckel potentials, the derivative of the Hamiltonian with respect to the ellipsoidal coordinates will be zero for a true Stäckel potential. Setting it equal to zero for a general potential gives a further constraint on A_τ and B_τ allowing us to solve for these integrals given only a single $(\boldsymbol{x},\boldsymbol{v})$ coordinate. Then we have an approximate expression for $p_\tau(\tau)$ that can be integrated to estimate the actions.

4.4 Stäckel fitting

Whereas the Stäckel Fudge uses the real potential as if it were a Stäckel potential, Sanders (2012) explicitly fitted a Stäckel potential to the real potential. This method uses the procedure from Dejonghe & de Zeeuw (1988) to find the locally best-fitting Stäckel potential. One minimizes the difference between the auxiliary function $(\lambda-\nu)\Phi$ formed from the true potential and the Stäckel auxiliary function $F(\nu)-F(\lambda)$. The best-fitting F can be computed by an integral over a region of the prolate spheroidal coordinate system that is defined on an orbit-by-orbit basis. We first integrate the orbit for several time-steps to form an average estimate of Δ^2 via equation (15). We further integrate the orbit to determine the orbit boundaries in λ and ν. With the boundaries found, we compute the best-fitting F on a grid (we use 40 grid points) and hence the best-fitting Stäckel potential in which the actions can be computed as detailed in Section 2.2.

The fitting procedure of Dejonghe & de Zeeuw (1988) uses weight functions which allow some flexibility in the fitting of the potential. We use weight functions $\Lambda(\lambda) \propto (\lambda-\Delta^2)^{-4}$ and $N(\nu) \propto (\nu-c^2)^{0.5}$.

4.5 Angle and frequency estimation

The angle coordinates are related to the generating function, S, by

$$\boldsymbol{\theta} = \frac{\partial S}{\partial \boldsymbol{J}}. \tag{41}$$

In an axisymmetric Stäckel potential, this generating function has the form

$$S(\lambda,\phi,\nu,J_r,J_\phi,J_z) = \int^\lambda \mathrm{d}\lambda\, p_\lambda + \int^\phi \mathrm{d}\phi\, p_\phi + \int^\nu \mathrm{d}\nu\, p_\nu. \tag{42}$$

The angle coordinates can be computed by writing

$$\theta_i = \sum_k \frac{\partial S}{\partial I_k}\frac{\partial I_k}{\partial J_i}. \tag{43}$$

The first term is computed by differentiating the momenta p_i with respect to the classical integrals I_k and integrating along the coordinate path. The second term is computed by inverting the matrix $\partial J_i/\partial I_k$ which again is found by differentiating the momenta in the action integrals and integrating over the full orbital path (see for example the appendix of Sanders 2012). This approach can be extended to compute the angles using the presented approximation schemes. In the adiabatic approximations, the three integrals are taken to be (E, J_ϕ, J_z). When computing the derivatives of the generating function and J_r with respect to J_z, we require the derivatives $\partial J_z/\partial E_z|_R$ and $\partial J_z/\partial E_\nu|_{\lambda,J_\phi}$ which we choose to compute from the grids (note these quantities can be computed by integration but as they are required for each integrand call this would be computationally expensive). In the Stäckel Fudge, the classical integrals are (E, J_ϕ, B_τ) and finally in the Stäckel-fitting procedure the classical integrals are (E, J_ϕ, I_3).

The frequencies are given by $\partial H/\partial \boldsymbol{J}$, so these are just single components of the inverse of the matrix $\partial J_i/\partial I_k$.

There is a significant difference between how the actions and the frequencies are evaluated. Each action is an integral of the momenta, while a frequency is the integral of the inverse of the momentum. Therefore, the error in the action is dominated by the regions of high momentum in the middle of the range of integration while the error in the frequency is dominated by errors in the integrand near the end points of the range of integration, where the integrand diverges. For this reason, the frequencies are much more sensitive to the location of the orbit's turning points while the actions are more sensitive to how closely the potential is modelled within the body of the orbit.

5 METHOD COMPARISONS

In this section, we critically compare the methods presented above. We focus on two quantities for each method: the accuracy of the action computation (i.e. how constant are the actions around an orbit) and how much computing time is required. The latter quantity is slightly subjective as it depends on computational details. However, we will see that the time differences between the methods are orders of magnitude, so we do not expect programming upgrades to disturb the rank-ordering of the methods.

Since there are more competing methods for the axisymmetric than for the triaxial case, we limit our comparison to the axisymmetric versions of the algorithms – in the triaxial case the choice currently is between the O2GF method of Section 3.2 if one requires a few accurate actions and the Stäckel Fudge method of Section 4.3.2 if one requires many actions with a lower level of accuracy. The axisymmetric methods we compare are as follows:

(i) cylindrical adiabatic approximation (CAA), Section 4.1;
(ii) spheroidal adiabatic approximation (SAA), Section 4.2;
(iii) Stäckel Fudge with Δ chosen using equation (15) (Fudge v1), Section 4.3;
(iv) Stäckel Fudge with Δ chosen from shell orbits (Fudge v2), Section 4.3;
(v) locally fitting Stäckel potentials (Fit), Section 4.4;
(vi) iterative torus construction (ItTC), Section 3.1;
(vii) computing the generating function from an orbit (O2GF), Section 3.2;
(viii) averaging toy actions over toy angles (AvGF), Section 3.3.

We begin by inspecting numerically estimated actions at a series of times along single orbits (Fig. 2) and we go on to investigate the action, angle and frequency variances for a broader range of orbits (Section 5.2).

5.1 Four representative orbits

We analyse results for four representative Galactic orbits: a typical *thin* disc orbit, a typical *thick* disc orbit, a typical *halo* orbit and a typical orbit for the progenitor of a tidal *stream*. The initial conditions of each orbit are as follows:

(i) thin,

$\boldsymbol{x} = (8.29, 0.1, 0.1)\,\mathrm{kpc}$, $\boldsymbol{v} = (30.22, 211.1, 19.22)\,\mathrm{km\,s^{-1}}$,

(ii) thick,

$\boldsymbol{x} = (8.29, 0.1, 0.1)\,\mathrm{kpc}$, $\boldsymbol{v} = (50.22, 187.1, 54.22)\,\mathrm{km\,s^{-1}}$,

(iii) halo,

$\boldsymbol{x} = (8.29, 0.1, 0.1)\,\mathrm{kpc}$, $\boldsymbol{v} = (100.22, 109.1, 101.22)\,\mathrm{km\,s^{-1}}$,

(iv) stream,

$\boldsymbol{x} = (26., 0.1, 0.1)\,\mathrm{kpc}$, $\boldsymbol{v} = (0.1, 141.8, 83.1)\,\mathrm{km\,s^{-1}}$.

The thin, thick and halo orbits were chosen arbitrarily but the stream orbit is selected as a likely orbit for the progenitor of the GD-1 stream (Koposov, Rix & Hogg 2010; Sanders & Binney 2013b). Each orbit was integrated for 10 orbital periods of a circular orbit of the same energy and 1000 time samples were recorded. The orbits were integrated in the multicomponent Galactic potential of Piffl et al. (2014). This potential was fitted to the kinematics of Radial Velocity Experiment (RAVE) stars, the vertical density profile of the solar cylinder measured by the Sloan Digital Sky

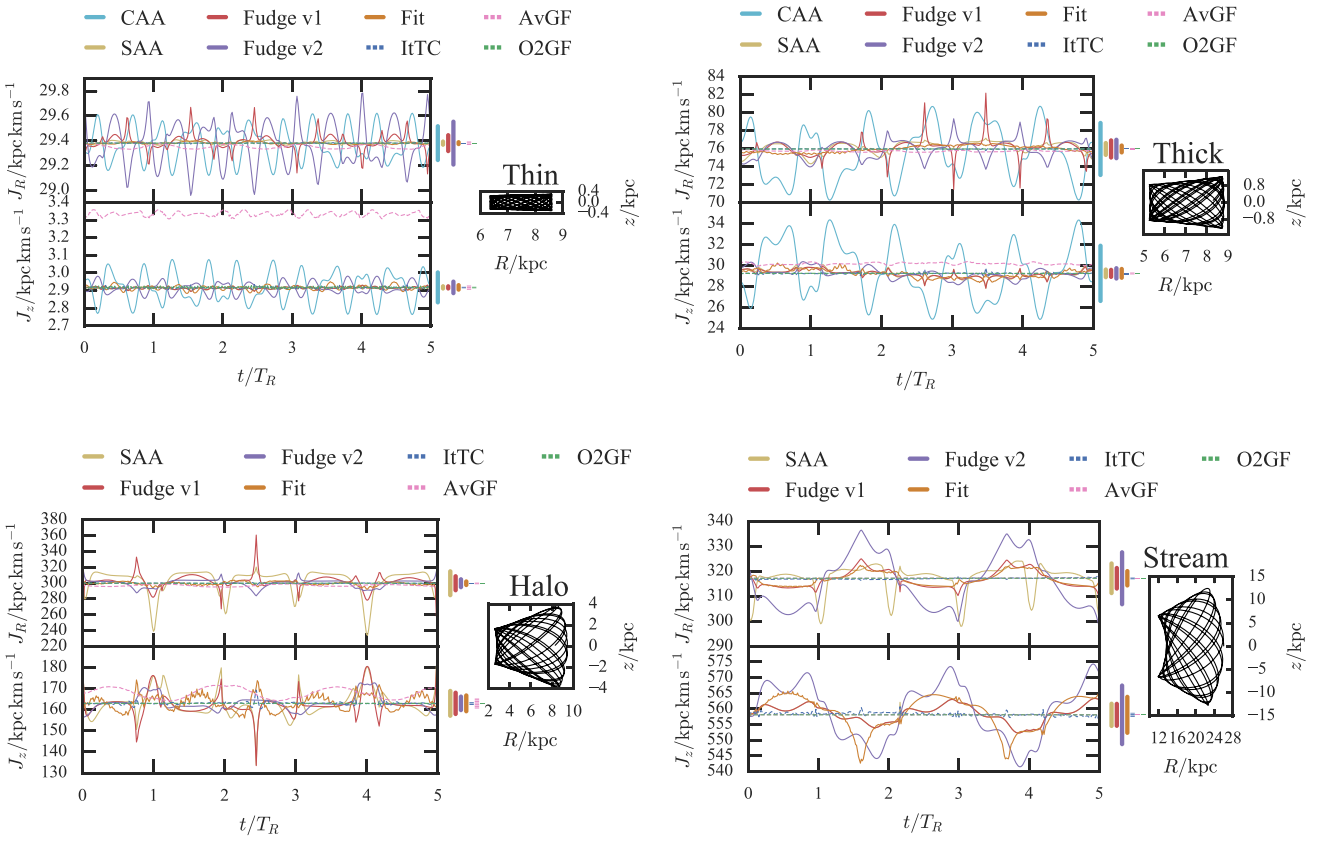

Figure 2. Comparison of action estimation methods for single orbits. In each set of panels, the orbit's trace in the meridional plane (R, z) is shown on the right. The set of panels at top left is for a typical *thin* disc orbit, the set at top right is for a typical *thick* disc orbit, that at bottom left is for a typical *halo* orbit and the bottom-right set is for a typical orbit of the progenitor of a tidal *stream*. Each set of panels shows on the left the radial action, J_r (top panel), and the vertical action, J_z (lower panel), as a function of time in units of the radial period. Values from *non-convergent* methods are shown with solid lines whilst values from *convergent* methods are shown with dashed lines. Cyan solid line is for the CAA, mustard-yellow solid line is for the SAA, red solid line is for the Stäckel Fudge with variable Δ estimation (v1), the purple solid is for the Stäckel Fudge with Δ estimated from shell orbits (v2), the solid orange line is for the Stäckel-fitting method, the blue dashed line is for the ItTC method, the green dashed line is for the O2GF method and the pink dashed line is for the AvGF method. To the right of each sub-panel, the vertical lines show the $\pm$ standard deviation spread in the action estimates from each method. All orbits are computed in the multicomponent potential from Piffl et al. (2014).

Table 1. Errors in the action for four example orbits: in the first line we give the radial and vertical actions for the four orbits in units of kpc km s^{-1}. Below the double horizontal separators, we give the rms deviations of the radial and vertical action estimates relative to the mean of the O2GF estimates for the eight methods. The methods above the horizontal separator are the non-convergent methods whilst those below are the convergent methods.

		Thin		Thick		Halo		Stream	
		J_R	J_z	J_R	J_z	J_R	J_z	J_R	J_z
	Actions	29.38	2.92	75.96	29.25	299.79	163.02	317.20	558.09
Method	CAA	0.1	0.08	3	3	40	50	100	200
	SAA	0.01	0.007	0.7	0.4	10	6	6	4
	Fudge v1	0.07	0.007	1	0.3	9	5	4	3
	Fudge v2	0.2	0.03	1	0.5	5	3	10	9
	Fit	0.007	0.01	0.4	0.4	2	4	3	6
	ItTC	0.002	0.008	0.03	0.06	0.3	0.8	0.4	0.6
	AvGF	0.03	0.4	0.2	2	3	6	0.2	0.3
	O2GF	0.0003	0.005	0.002	0.007	0.03	0.03	0.01	0.01

Survey (Jurić et al. 2008) and the constraints on the circular speed as a function of Galactocentric radius assembled by McMillan (2011). The generating mass distribution consists of three exponential discs (thin, thick and gas), a central bulge and an NFW dark halo. The actions were computed using the eight different methods for each time sample. These estimates are plotted in Fig. 2 and their standard deviations are given in Table 1.

The ItTC and O2GF methods produce root-mean-squared (rms) deviations that are small enough for all purposes for all the inspected orbits, so we will not comment on them further in this section.

5.1.1 *Thin*

From the plot of the orbit in the meridional plane, we can clearly see that the orbit is nearly separable in cylindrical polar coordinates. However, the Stäckel Fudge and SAA methods give better results than the CAA. Of the non-convergent methods, Stäckel fitting gives the most accurate results.

For this very low eccentricity orbit, locally choosing Δ around the orbit (Fudge v1) gives better results than using shell orbits to choose the parameter Δ of the ellipsoidal coordinate system (Fudge v2), but this ordering probably reflects the grid used to compute Δ for the closed orbits.

The AvGF method gives slightly biased estimates of both actions. One reason for this is that the toy actions are always positive, so for small actions, averaging produces a positively biased action. This might be improved by using a better toy potential but we do not explore this option here.

5.1.2 *Thick*

The left and right boundaries of the thick disc orbit show significant curvature, so the assumption of separability in cylindrical polar coordinates is not a good one. For this reason, the CAA is inferior to the other methods by a factor of a few. All four of the methods based on Stäckel potentials give comparable accuracy for J_z but explicit Stäckel fitting gives a result that is twice as accurate as the other methods for J_r. As for the thin disc orbit, the AvGF method gives slightly biased estimates of both J_r and J_z, with the error in $J_z \sim 1$ kpc km s^{-1}.

5.1.3 *Halo*

For the halo and stream orbits, we have not plotted results from the CAA because the latter is not competitive at these higher eccentricities. The hierarchy for the Stäckel-type methods is SAA, Fudge v1, Fudge v2 and Stäckel fitting with the Stäckel-fitting method producing a value of J_r that is a factor of 5 more accurate than that from the SAA. All three methods produce equally accurate results for J_z. The AvGF method again produces a small systematic shift in J_z.

Since on this orbit the z oscillations do not occur faster than the R oscillations, there is no reason why J_z should be adiabatically conserved as the SAA assumes. Hence, it is remarkable that the SAA performs quite well on this orbit.

5.1.4 *Stream*

The stream-like orbit ranks the Stäckel-type methods in the same order as the halo orbit. The Stäckel Fudge v1 and the SAA produce very similar results for both J_r and J_z. Stäckel fitting produces more accurate results for J_r but a less accurate value for J_z, and the Stäckel Fudge v2 yields slightly larger errors in both J_r and J_z. Evidently, estimating Δ from shell orbits is not optimal for this more radially extended orbit. There are no obvious systematic shifts in the actions from the AvGF method.

5.2 Multiple orbits

We now analyse results for a larger sample of orbits. All orbits have $J_\phi = R_0 V_c(R_0)$, where R_0 is the solar radius and V_c is the circular speed. At a series of steadily increasing energies, we set $v_z = 0.8 v_R$ and launch the orbit from the solar position. We integrate orbits for 10 periods of the circular orbit with the same energy and compute the actions, angles and frequencies at 40 points along the orbit. The most extreme orbit considered has pericentre $\sim$5 kpc, apocentre $\sim$28 kpc and maximum Galactic height $\sim$15 kpc.

In Fig. 3, we show as a function of $(J_r + J_z)/|J_\phi|$, which is a measure of eccentricity, the rms of the differences between the actions from a particular method and the mean of the values obtained from the O2GF method. The use of the rms does not give the full story as inspection of Fig. 2 shows that the action fluctuations can be quite peaky, particularly around the turning points of the orbit. However, it serves as a useful summary statistic. Since the green dashed curve of the O2GF method runs along the bottom of both the upper panel for rms in J_r and the lower panel for rms in J_z, the O2GF method emerges as the most accurate method regardless of orbital eccentricity. It must, however, be acknowledged that our playing field is not entirely level in that rms fluctuations are computed relative to the mean of the O2GF results, so for O2GF to show no rms it merely needs to extract the same actions from two, strongly

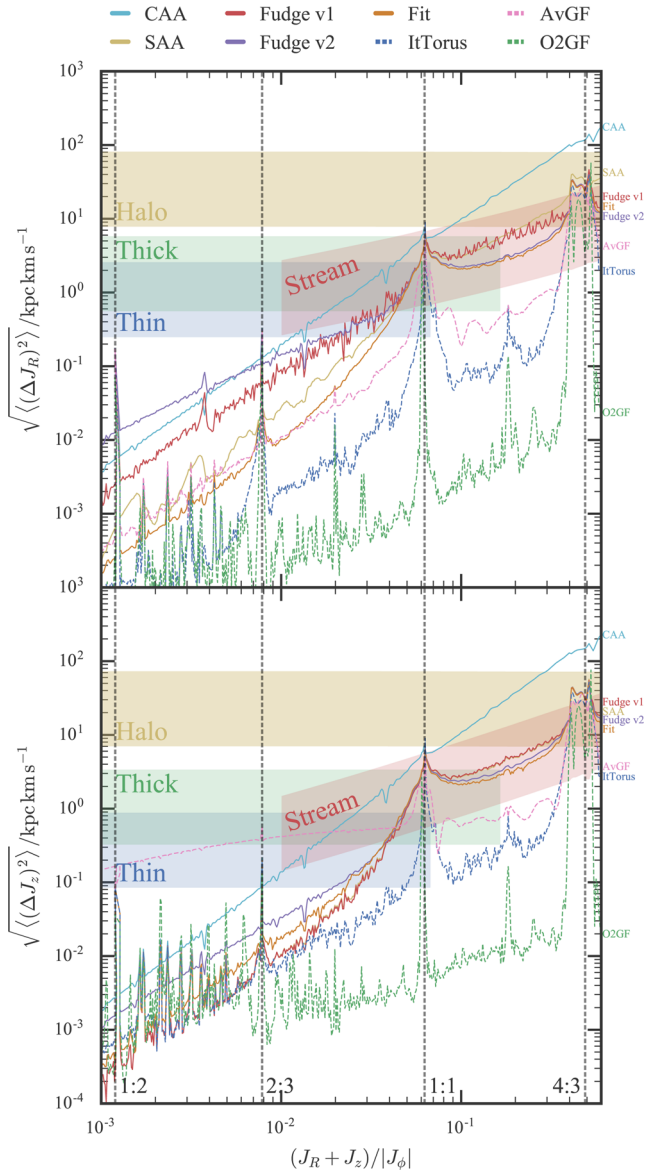

Figure 3. Comparison of the accuracy of the action estimation methods for multiple orbits: each panel shows the rms deviations of the estimates from the mean computed using the O2GF method. The upper panel shows the radial action J_r and the lower panel shows the vertical action J_z. The *non-convergent* methods are shown with solid lines whilst the *convergent* methods are shown with dashed lines. The cyan solid line is for the CAA, the mustard-yellow solid line is for the SAA, the red solid line is for the Stäckel Fudge with variable Δ estimation, the purple solid line is for the Stäckel Fudge with Δ estimation from shell orbits, the solid orange line is for the Stäckel-fitting method, the blue dashed line is for the ItTC method, the green dashed line is for the O2GF method and the pink dashed line is for the AvGF method. The coloured bands show the region in which the relative errors in the action-based DFs for the thin (blue), thick (green) and halo (yellow) components are between 1 and 10 per cent. The upper and lower boundaries of the red band correspond to the action-space widths of a stream shed from a progenitor that has velocity dispersions of 5 and 0.5 km s^{-1}. The chosen horizontal limits of these bands are $\log f/f_{\max} = -5$ for the thin and thick components and $(J_r + J_z)/|J_\phi| = 0.01$ for the stream component. The vertical grey dashed lines correspond to resonances $x : y$ when $\Omega_R/\Omega_z = x/y$.

overlapping segments of orbit. In particular, its rms would be zero even if it consistently returned erroneous actions.

The next most accurate method is the ItTC method. As regards J_z, it is followed by the Stäckel Fudge v1 and SAA methods at small eccentricities and by the AvGF method at larger eccentricities. However, consideration of J_r yields a different ordering with both the Stäckel-fitting and AvGF methods proving more accurate than the Stäckel Fudge v1 at all eccentricities. The SAA and Stäckel Fudge methods calculate the vertical action in very similar ways so at all eccentricities Stäckel Fudge v1 method has near identical errors to the SAA method. However, inspection of Fig. 2 shows that the individual action estimates around an orbit are not identical for the two methods. At high eccentricities, the Stäckel-fitting method produces the best results of the non-convergent methods followed by Fudge v2 and then the SAA and Fudge v1. At the largest eccentricities, the CAA produces errors that are an order of magnitude larger than the Stäckel-based non-convergent methods.

In Fig. 3, vertical lines mark eccentricities at which orbits may become trapped by resonances. At these eccentricities, there are sharp peaks in the errors for all methods. The peaks are particularly sharp for the O2GF method. For eccentricities smaller than that at which the 2:3 resonance occurs, the errors in J_z are similar for the O2GF, ItTC and Stäckel Fudge methods. Even at these low eccentricities, the errors in J_r from the Stäckel Fudge v1 are significantly larger than those from the two convergent methods.

In Fig. 4, we show the rms deviations of the frequencies from the mean of the frequencies returned by the O2GF method, and in Fig. 5 we show the rms deviation of the angles from those computed with the O2GF method. Again the O2GF method shows clear superiority at all eccentricities, followed by the ItTC method. Of the Stäckel-based non-convergent methods, the broad picture is that Stäckel fitting performs best followed by the Stäckel Fudge, the SAA and finally the CAA. An exception to this is the azimuthal frequency Ω_ϕ where the SAA performs worst and the CAA is competitive with Stäckel fitting. Additionally, the SAA angle estimation is very competitive at low eccentricities. The vertical frequency estimates for both the CAA and SAA are peaky due to the interpolation grid that is used. Again we see a striking deterioration in the performance of the O2GF method at resonances. The deleterious impact of resonances is particularly marked in the case of the angle variables.

We have also computed rms deviations for two series of orbits generated in the same way as those used for Figs 3–5 but starting from the circular orbits at 5 kpc and at 13 kpc rather than at R_0. The results are qualitatively the same as those given above. We further explored the impact on these tests of replacing the realistic Galactic potential by a flattened logarithmic potential ($q = 0.9$). The errors with the actions are much smoother in this case and resonances are not evident. The magnitudes of the errors in the actions are very similar to those in the more realistic potential with all the non-convergent Stäckel methods producing very similar magnitude errors that are a factor of 10 smaller errors than the CAA, and the AvGF and ItTC methods producing errors another factor of 10 smaller for the most eccentric orbits. The O2GF method produces significantly smaller errors in the logarithmic potential than in the multicomponent potential ($\sim$ a factor of 10) as the smoothness of the logarithmic potential means fewer terms are required in the generating function.

5.3 Computational time

In Fig. 6, we show for each method the average time for computing a single action as a function of $(J_r + J_z)/|J_\phi|$. The actions

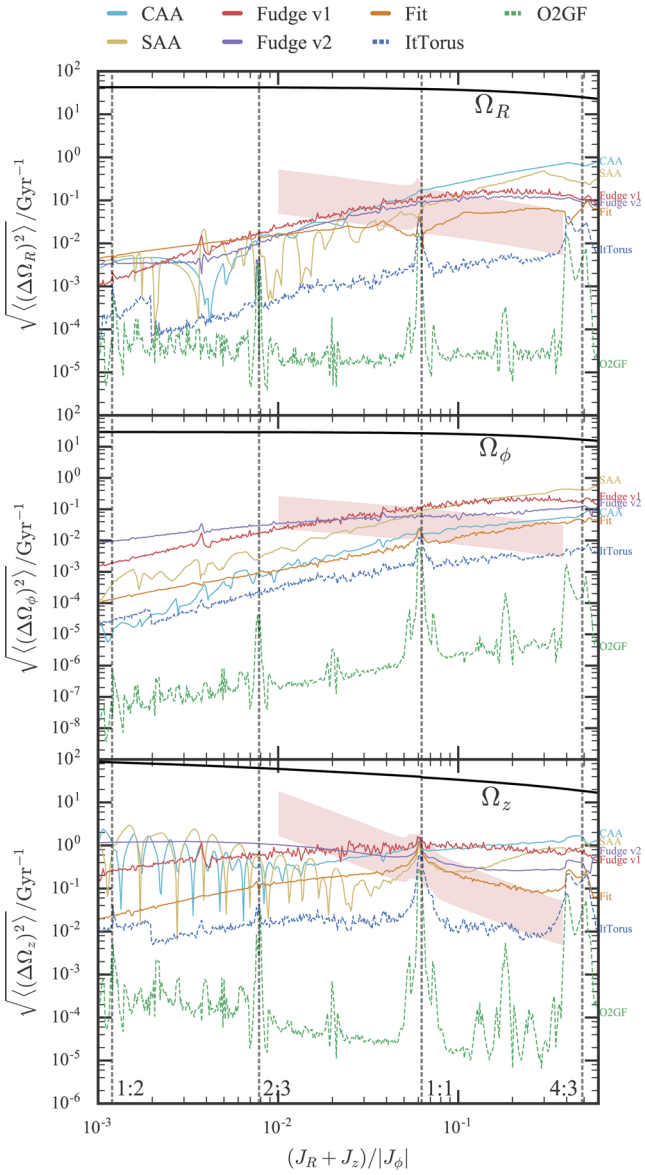

Figure 4. Comparison of the accuracy of frequencies estimated using different methods: the top panel shows the radial frequency, middle panel the azimuthal and bottom panel the vertical frequency. Each panel shows the rms deviations of the estimates from the mean computed using the O2GF method. The colour-coding and range of orbits explored are as in Fig. 3. The solid black lines show the frequencies of the orbits computed using the O2GF method. The upper and lower boundaries of the red band correspond to the frequency widths of streams with velocity dispersions of 5 and 0.5 km s^{-1}. The vertical grey dashed lines correspond to resonances $x : y$ when $\Omega_R/\Omega_z = x/y$.

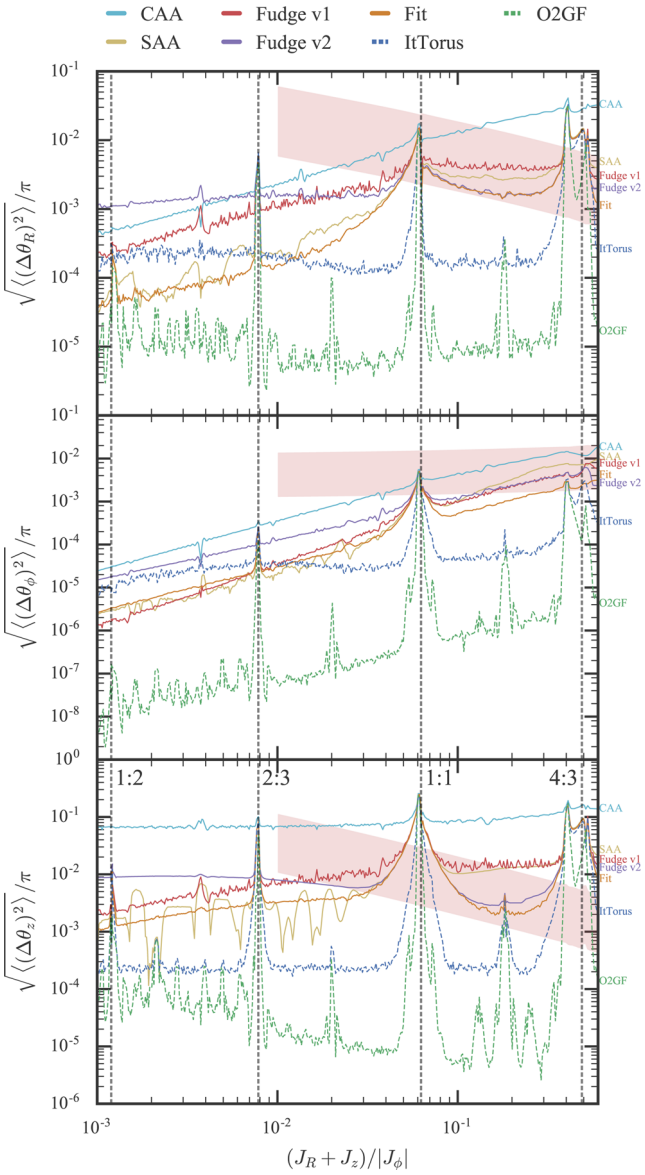

Figure 5. Comparison of the accuracy of angle estimation using different methods for multiple orbits: the top panel shows the radial angle, middle panel the azimuthal and bottom the vertical. The colour-coding and the range of orbits explored are as in Fig. 3. Each panel shows the rms deviations of the estimates from the angles computed using the O2GF method. The O2GF line shows the rms deviations from the angles computed using $\theta_i(0) + \langle\Omega_i\rangle t$. The upper and lower boundaries of the red band correspond to the angle widths of streams with scale radii of 100 and 10 pc. The vertical grey dashed lines correspond to resonances $x : y$ when $\Omega_R/\Omega_z = x/y$.

were computed on a single 4GHz Intel i7 processor. For an orbit of given eccentricity, these times vary by more than three orders of magnitude, with the Stäckel Fudges being cheapest ($\sim 5 \times 10^{-6}$ s) and the ItTC method most expensive. In the bargain basement, we have the Stäckel Fudges and the adiabatic approximations CAA and SAA. Given that the CAA and SAA are two to three times more expensive (because they involve interpolations in two- and three-dimensional tables), and are never significantly more accurate (Fig. 3 and Table 1) than the Stäckel Fudges, it is clear that anyone shopping for a cheap and cheerful method of action determination should choose a Stäckel Fudge. Versions 1 and 2 of the Stäckel Fudge have identical cost but version 1 is the more accurate except for highly eccentric orbits, when its disadvantage is marginal. Consequently, for general low-precision work Stäckel Fudge v1 is the clear recommendation.

Fig. 6 shows the mid-price methods to be the AvGF, Stäckel Fit and O2GF. Of these, the O2GF method clearly stands out as the best buy because it is orders of magnitude more accurate than its peers. Its only drawback is that it is a complex method to code. The dashed green curve plotting its cost is jagged because for nearly resonant orbits a longer or more densely sampled orbit integration is required.

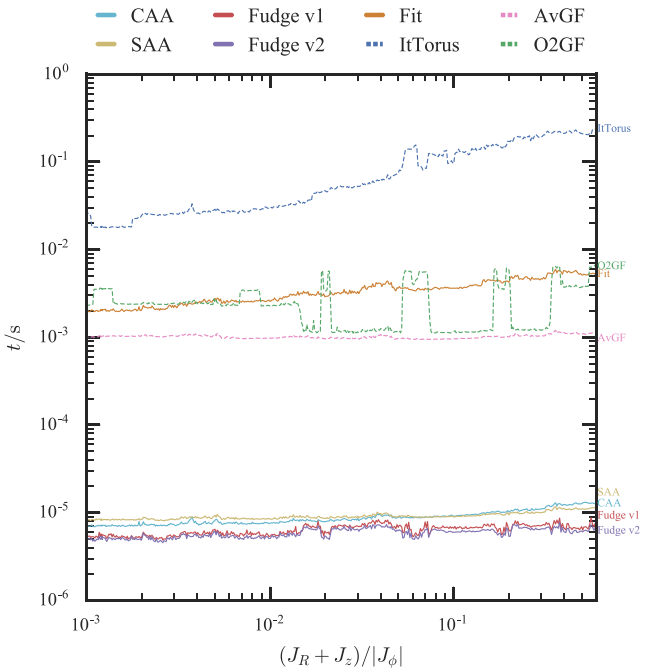

Figure 6. Comparison of the speed of the methods. The colour-coding and range of orbits explored are as in Fig. 3. The time quoted is the time per estimated action averaged from 50 estimates per orbit.

The ItTC method is roughly an order of magnitude more expensive than any other method. Its strength is robustness in the presence of resonant trapping. Each torus construction involves inverting two matrices (one for the S_n and another for their derivatives) and we have allowed a maximum of five torus constructions. In some applications, a significant speed-up should be possible by constructing tori by interpolating on a grid of pre-computed tori as described by Binney & McMillan (2016).

6 DISCUSSION

We have seen that the costs and performance of action finders vary considerably, with the Stäckel Fudge v1 delivering lots of medium-quality actions in the time required for the O2GF method to deliver a single high-quality result. Consequently, it is important to understand what accuracy one really requires for a particular application.

An important application is the computation of observable quantities such as stellar densities and kinematics from a model with a specified DF $f(\boldsymbol{J})$. In this case, the key criterion is the accuracy with which either $f(\boldsymbol{J})$ or its moments can be evaluated. Binney & McMillan (2011) compared the recovery of the velocity distributions and density profiles of model stellar discs when the actions were obtained from either an early version of the CAA or torus mapping. They found that within $\sim$1 kpc of the Galactic plane, the differences were very small, but further from the plane the adiabatic approximation produced results that are in error by $\sim$10 per cent. Similarly, Bovy & Rix (2013) showed a comparison between an early version of the CAA and the Stäckel Fudge. They found that the difference in the vertical velocity dispersions obtained from the two methods was $\sim$1 per cent in a model with a very small vertical gradient in velocity dispersion.

Binney (2014) showed that densities and velocity dispersions computed for flattened isochrone spheres using the Stäckel Fudge v2 satisfy the Jeans equations to good precision. Similarly, Sanders & Binney (2014) demonstrated that actions estimated with the triaxial Stäckel Fudge are sufficiently accurate for the resulting density and velocity dispersions of a triaxial Navarro–Frenk–White dark-matter halo to satisfy the Jeans' equations. These studies demonstrate that although the error in the actions for individual orbits can be large for some methods, the resulting moments of the DF are well recovered as they are presumably dominated by those low-action orbits that have good action estimates.

In Appendix A, we give details of four DFs that describe components of our Galaxy. With these DFs we are able to compare the accuracy of the action calculations required to evaluate the DFs to a certain accuracy. In each panel of Fig. 3, we show four coloured bands. The three labelled *thin*, *thick* and *halo* give the accuracy required to evaluate the corresponding DF accurate to between 1 and 10 per cent. In Fig. 3, the band labelled *stream* shows the action-space spreads for a stream created from a progenitor with a velocity dispersion of between 0.5 and 5 km s^{-1}. The yellow bands in Figs 4 and 5 show the width in frequencies and angles of particles very recently stripped from such a progenitor. Since the spread in the angles grows linearly in time as the stream grows, the given width is the smallest one would ever wish to resolve. Wherever the curves showing a method's accuracy lie below all these coloured bands, that method provides sufficient accuracy for the job in hand. We see that the Stäckel Fudge v1 comes too high for safety only in a small range of values of $(J_r + J_z)/|J_\phi|$, so even this cheap and cheerful method is adequate for all but a few thick disc orbits.

The spread in frequencies and angles we have computed is averaged over many orbital periods. In reality, when modelling a stream, the stream particles are confined to a much narrower range of angles. The frequency estimation for the collection of particles may be systematically offset from the truth due to the phase of the stream but we anticipate the interparticle frequency separations will be well resolved. Therefore, we can perhaps get away with slightly larger errors in the frequency estimation than those suggested by the *stream* band in Fig. 4.

6.1 Trapping by resonances

We have seen that results from different methods tend to diverge for orbits on which Ω_r: Ω_z approaches a simple ratio such as 1:2 or 2:3. This finding arises because these orbits have become trapped by a resonance. The pre-requisite to a clear discussion of resonant trapping is the existence of a Hamiltonian H_0 that is similar to the actual Hamiltonian H but in which resonant trapping does not occur. That is, for an orbit $\boldsymbol{J}^{(0)}$ in H_0, $\Omega_r/\Omega_z = m/n$ for small integers m and n, but on neighbouring orbits Ω_r/Ω_z is an irrational number close to m/n. Given the existence of H_0, one can then understand resonant trapping by considering the impact on orbits in H_0 of the small perturbation $\Delta \equiv H - H_0$ (e.g. Binney 2013).

Perturbation theory allows one to estimate the extent of the region in action space within which the perturbation Δ will cause orbits to become trapped by the m: n resonance. It may happen that this region overlaps with the region of entrapment by another resonance, associated with different integers. In this case, Δ causes orbits in H_0 that lie in the overlap region to become chaotic in the sense that they cease to be quasi-periodic (Chirikov 1979; Binney 2013).

Resonantly trapped quasi-periodic orbits have actions, but we cannot take two of these actions to be J_r and J_z; instead, we must use a linear combination of J_r and J_z complemented by a new action that quantifies the extent of their libration about the underlying resonant orbit (see section 3.7.2 of Binney & Tremaine 2008). Chaotic orbits do not admit a full complement of action integrals.

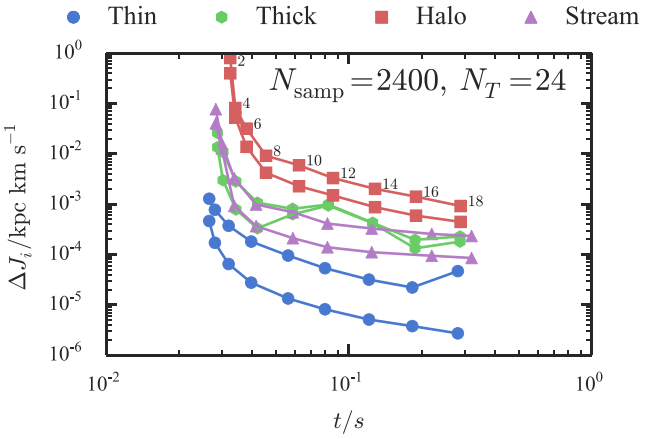

Figure 7. Error in the actions as a function of computation time for the O2GF method using different numbers of Fourier coefficients: $N_{\rm max}$ is the maximum magnitude of the Fourier vectors $\boldsymbol{n}$ considered. We show the radial (bottom for all) and vertical action errors for thin- (blue circles), thick- (green hexagons), halo- (red squares) and stream-like (purple triangles) orbits. The halo points are labelled by $N_{\rm max}$. For all cases, we fix $N_T = 24$ and $N_{\rm samp} = 2400$.

The applications we have in mind for actions assume that negligibly few orbits are resonantly trapped, so all orbits can be characterized by J_r, J_ϕ and J_z. Consequently, what we ask of a method of action determination is that when it is confronted by a trapped orbit, it returns the actions of one of the orbits in H_0 that is trapped into the given orbit. We say 'one of' because a given orbit in H_0 can be trapped into many different orbits in H, and a given orbit in H can be reached from many orbits in H_0. Hence, resonant trapping obliges us to allow a certain fuzziness in the meaning of actions. If the DF returns similar values for all actions in the range of a given region of entrapment, this fuzziness is likely to be unimportant. Thus, the criterion for resonant trapping to be unimportant is closely related to the criterion we have explored for errors from action finders being unimportant.

The nearby integrable Hamiltonian H_0 can be constructed by torus mapping (Kaasalainen & Binney 1994a; Binney & McMillan 2016), and we hope shortly to present practical details of how this is best achieved.

6.2 Convergence of the O2GF method

At the start of the previous section, we remarked that the two convergent methods produced results that by some distance out-rank the results from the non-convergent methods. However, the results we gave were obtained using just one set of parameters in each convergent method. Here we explore how the error in the actions changes as a function of the parameters used in the O2GF method.

There are three parameters to choose for the O2GF method: the integration time N_T, the number of samples $N_{\rm samp}$ and the number of Fourier coefficients (controlled by the parameter $N_{\rm max}$). The error in the action proves to be most strongly a function of the number of Fourier coefficients. For the frequencies and angles, the error is also affected by the total integration time.

In Fig. 7, we show the error in the actions for the four representative orbits as a function of $N_{\rm max}$ and computation time using $N_T = 24$ and $N_{\rm samp} = 2400$. As expected, the errors generally decrease and the computing time increases as $N_{\rm max}$ increases. However, increasing $N_{\rm max}$ sometimes *increases* the errors. For small $N_{\rm max}$, the computing time is dominated by the orbit integration, with the consequence that the time required to compute the actions of the thin disc orbit, which is easy to integrate, is significantly smaller than that required to compute the actions of the halo orbit. For large $N_{\rm max}$, by contrast, the computing time is dominated by equation solving, and similar times are required to compute the actions of all four orbits of Fig. 2.

6.3 TACT code

'The Action Computation Tool' or TACT is a publicly available code implementing all the algorithms given in this paper. It is available at https://github.com/jls713/tact. We also provide programs to produce the plots given in this paper, and there is documentation describing each of the algorithms. The code is written in C++ but code is provided to compile the routines into a PYTHON library. The code can work in tandem with the code TM (Torus Mapper), which is described by Binney & McMillan (2016).

Each algorithm discussed in the paper is implemented as an `Action_Finder` class. These classes take as arguments instances of a potential class (`Potential_JS`), and all have methods `actions` and `angles` that return the actions, and angles and frequencies respectively given some Cartesian 6D phase-space point (x, y, z, v_x, v_y, v_z). It should be simple for users to specify their own potential by implementing an inherited class of `Potential_JS` and defining a method `Phi` that returns the potential and a method `Forces` that returns the force. We also provide a simple `Orbit` class that may be used with custom potentials.

7 CONCLUSIONS

We have reviewed the currently available algorithms for estimating actions, angles and frequencies in general potentials. We began by presenting a discussion of the separable potentials in which these quantities can be expressed as simple quadratures. The most general such class of potentials is the triaxial Stäckel potentials of which axisymmetric Stäckel potentials and spherical potentials are limiting cases. The only case in which the actions can be computed analytically is the isochrone potential. The existence of actions in more general systems has been demonstrated from numerical experiments but there is no way to simply compute the actions in a general potential. The results from the special separable cases have inspired and aided work in creating numerical algorithms for action estimation in general potentials, and we have collected together and summarized these algorithms.

We focused on axisymmetric potentials and for tests used a realistic model of the potential of our Galaxy. The presented methods fall into two groups: non-convergent methods and convergent methods. The convergent methods should produce ever more accurate action estimates given increased computing resource. The non-convergent methods all centre on defining one-dimensional effective potentials that enable actions to be obtained from a one-dimensional integral. Directly or indirectly they all derive from Stäckel potentials, and their accuracy deteriorates with the distance of the actual potential from the nearest Stäckel potential. The convergent and non-convergent methods are entirely complementary. For some applications, one requires a few accurate action computations (e.g. stream modelling) whilst for other applications one requires many less accurate actions (e.g. disc or halo modelling).

We began by comparing the actions estimated along four representative orbits in a multicomponent axisymmetric Galactic potential. We went on to compare the accuracy of the actions, frequencies and angles for a large number of orbits in the same potential, and we finally discussed the computational speed of each method. The

time required to compute an action varies by three to four orders of magnitude between the cheapest method, the Stäckel Fudge, and the most expensive method, ItTC. We gave DFs for components of the Milky Way and compared the accuracy of the action estimates yielded from the various methods with the relative error in the corresponding evaluation of the DF. The accuracy of Stäckel Fudge v1 is sufficient for modelling the smooth phase-mixed components of the Galaxy (the disc and halo) with the possible exception of some thick disc orbits. The convergent methods O2GF and ItTC provide significantly more accuracy than one normally requires, except possibly at powerful resonances and when modelling cold structures such as tidal streams.

The accuracy of all methods deteriorates in the vicinity of major resonances such as the resonance $\Omega_r/\Omega_z = 1$. Orbits can become trapped at such resonances, and the actions we seek to evaluate here are not the true actions of a trapped orbit. Instead one independent action should be a linear combination of the actions sought here, and the other independent action quantifies the amplitude of the trapped orbit's libration about an underlying closed orbit. The non-convergent methods typically lack the ability to identify a resonance and the returned actions are at best the actions of a similar orbit in a nearby integrable Hamiltonian (in which trapping does not occur). Currently, there is no general-purpose tool that will evaluate the true actions of a trapped orbit, although Kaasalainen & Binney (1994b) showed that the actions and angles of trapped orbits can be evaluated by torus mapping.

Action, angle and frequency variables have proved important tools for modelling dynamical systems, and with the arrival of the *Gaia* data on the horizon, the need for robust, accurate and rapid routines for estimating the actions is great. We have demonstrated that there are powerful algorithms that are appropriate for a wide range of modelling requirements and that promise to be indispensable for understanding and dissecting the *Gaia* data.

ACKNOWLEDGEMENTS

JLS acknowledges the financial support of the Science and Technologies Facilities Council (STFC). This project was made possible through the use of following open-source software: the GNU Scientific Library (GSL; Galassi et al. 2009), the `cuba` multidimensional integration library (Hahn 2005), and the PYTHON packages NUMPY (Oliphant 2007) and MATPLOTLIB (Hunter 2007). We thank Angus Williams for reading a draft of this paper and constructing the backronym 'The Action Computation Tool' for the TACT code.

REFERENCES

Batsleer P., Dejonghe H., 1994, A&A, 287, 43
Bienaymé O., Robin A. C., Famaey B., 2015, A&A, 581, A123
Binney J., 2010, MNRAS, 401, 2318
Binney J., 2012a, MNRAS, 426, 1324
Binney J., 2012b, MNRAS, 426, 1328
Binney J., 2013, New Astron. Rev., 57, 29
Binney J., 2014, MNRAS, 440, 787
Binney J., Kumar S., 1993, MNRAS, 261, 584
Binney J., McMillan P., 2011, MNRAS, 413, 1889
Binney J., McMillan P. J., 2016, MNRAS, in press
Binney J., Piffl T., 2015, MNRAS, 454, 3653
Binney J., Tremaine S., 2008, Galactic Dynamics, 2nd edn. Princeton Univ. Press, Princeton, NJ
Bovy J., 2014, ApJ, 795, 95
Bovy J., Rix H.-W., 2013, ApJ, 779, 115
Chirikov B. V., 1979, Phys. Rep., 52, 263
De Bruyne V., Leeuwin F., Dejonghe H., 2000, MNRAS, 311, 297
de Zeeuw T., 1985, MNRAS, 216, 273
Dejonghe H., de Zeeuw T., 1988, ApJ, 329, 720
Eyre A., Binney J., 2011, MNRAS, 413, 1852
Famaey B., Dejonghe H., 2003, MNRAS, 340, 752
Fox M. F. J., 2014, Masters Thesis, University of Oxford
Galassi M., Davies J., Theiler J., Gough B., Jungman G., 2009, GNU Scientific Library – Reference Manual, 3rd edn. Network Theory Ltd.
Hahn T., 2005, Comput. Phys. Commun., 168, 78
Helmi A., White S. D. M., 1999, MNRAS, 307, 495
Henon M., 1959, Ann. Astrophys., 22, 126
Hunter J. D., 2007, Comput. Sci. Eng., 9, 90
Jurić M. et al., 2008, ApJ, 673, 864
Kaasalainen M., Binney J., 1994a, Phys. Rev. Lett., 73, 2377
Kaasalainen M., Binney J., 1994b, MNRAS, 268, 1033
Kent S. M., de Zeeuw T., 1991, AJ, 102, 1994
Koposov S. E., Rix H.-W., Hogg D. W., 2010, ApJ, 712, 260
McGill C., Binney J., 1990, MNRAS, 244, 634
McMillan P. J., 2011, MNRAS, 414, 2446
McMillan P. J., Binney J. J., 2008, MNRAS, 390, 429
McMillan P. J., Binney J. J., 2013, MNRAS, 433, 1411
Oliphant T. E., 2007, Comput. Sci. Eng., 9, 10
Ollongren A., 1962, Bull. Astron. Inst. Neth., 16, 241
Piffl T. et al., 2014, MNRAS, 445, 3133
Sanders J., 2012, MNRAS, 426, 128
Sanders J. L., 2014, MNRAS, 443, 423
Sanders J. L., Binney J., 2013a, MNRAS, 433, 1813
Sanders J. L., Binney J., 2013b, MNRAS, 433, 1826
Sanders J. L., Binney J., 2014, MNRAS, 441, 3284
Sanders J. L., Binney J., 2015a, MNRAS, 447, 2479
Sanders J. L., Binney J., 2015b, MNRAS, 449, 3479
Schönrich R., Binney J., 2012, MNRAS, 419, 1546
Sellwood J. A., 2010, MNRAS, 409, 145
Williams A. A., Evans N. W., 2015, MNRAS, 454, 698

APPENDIX A: RELATIVE ERRORS IN DFs FOR GALACTIC COMPONENTS

A1 Thin and thick discs

For both the thin and thick discs, we use the quasi-isothermal DF presented in Binney (2010). This DF is nearly-separable in the actions such that the radial and vertical dependence is given by

$$f_{\rm disc}(J_r, J_z) \propto \exp\left(-\frac{\kappa J_r}{\sigma_R^2} - \frac{\nu J_z}{\sigma_z^2}\right), \tag{A1}$$

where κ and ν are the epicyclic frequencies and σ_R and σ_z are velocity dispersion parameters. We adopt the parameter values from Piffl et al. (2014) that were found to produce a good fit to the RAVE data. From a DF of this form, the relative error in the DF is given by

$$\left(\frac{\Delta f_{\rm disc}}{f_{\rm disc}}\right)^2 = \sqrt{\left(\frac{\kappa \Delta J_r}{\sigma_R^2}\right)^2 + \left(\frac{\nu \Delta J_z}{\sigma_z^2}\right)^2}. \tag{A2}$$

Note that due to the form of the DF, the relative error in the DF is independent of the actions.

A2 Stellar halo

For the stellar halo, we use the DF from Williams & Evans (2015). This DF was fitted to the blue horizontal branch stars from the Sloan Extension for Galactic Understanding and Exploration survey and has the form

$$f_{\rm halo} \propto \mathcal{L}^{-q}(\mathcal{L}^2 + J_b^2)^{-(p-q)/2}, \tag{A3}$$

where

$$\mathcal{L} = \mathcal{F} D_0 J_r + |J_\phi| + J_z, \tag{A4}$$

$p = 0.83$ and $q = 9.16$ govern the inner and outer slopes of the density profile and $J_b = 3600\,\mathrm{kpc\,km\,s^{-1}}$ governs the break radius of the density profile. $D_0 = 1.52$ produces the isotropic model and $\mathcal{F} = 0.59$ controls the anisotropy of the model.

As with the disc DFs we use this expression to find the relative error in the DF as a function of the actions. Here we note that, unlike with the discs, the relative error *does* depend on the actions. However, we see from Fig. 3 that the dependence is relatively weak.

A3 Streams

Streams are formed from material tidally stripped from a progenitor and tend to form cold thin structures that are well approximated by a Gaussian in action space. If the progenitor has a velocity dispersion σ_v, the widths of the resulting action distributions are well approximated by (Eyre & Binney 2011)

$$\begin{aligned}
\sigma_{J_r} &= \frac{\sigma_\mathrm{v}(R_\mathrm{a} - R_\mathrm{p})}{\pi},\\
\sigma_{J_\phi} &= \frac{\sigma_\mathrm{v} R_\mathrm{p}}{\pi},\\
\sigma_{J_z} &= \frac{2\sigma_\mathrm{v} z_\mathrm{max}}{\pi},
\end{aligned} \tag{A5}$$

where R_p is the pericentric radius of the progenitor orbit, R_a the apocentric radius and z_max the maximum height above the Galactic plane.

Bovy (2014) and Sanders (2014) demonstrate that a simple, but realistic stream model can be constructed in angle–frequency space. Therefore, it is of interest to know the error required to resolve the angle–frequency structure of the stream. The frequency spread of the stream can be related to the action spread via the Hessian matrix as

$$\frac{\partial^2 H}{\partial J_i \partial J_j} \equiv \frac{\partial \Omega_i}{\partial J_j}, \tag{A6}$$

which is computed by finite differencing the frequencies for a series of tori constructed around the true torus. We use this matrix to convert the action spreads in equation (A5) into frequency spreads $\boldsymbol{\sigma}_{\Omega i}$.

McMillan (in preparation) shows that the initial angle spreads of particles released into a tidal stream from a progenitor are given by

$$\begin{aligned}
\sigma_{\theta_R} &= \frac{\pi r_\mathrm{s}}{(R_\mathrm{a} - R_\mathrm{p})},\\
\sigma_{\theta_\phi} &= \frac{\pi r_\mathrm{s}}{R_\mathrm{p}},\\
\sigma_{\theta_z} &= \frac{\pi r_\mathrm{s}}{2 z_\mathrm{max}},
\end{aligned} \tag{A7}$$

where r_s is the scale radius of the stream progenitor. However, the angle spread in the stream increases linearly in time at a rate governed by the frequency separation from the progenitor, so for most streams we do not need to resolve the angle spreads in such fine detail.

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