Mon. Not. R. Astron. Soc. **298,** 1–21 (1998)

Orbit classification in arbitrary 2D and 3D potentials

Daniel D. Carpintero[1,2] and Luis A. Aguilar[1]

[1] *Observatorio Astronómico Nacional, IAUNAM Ensenada, Apdo. postal 877, 22800 Ensenada, México*
[2] *Facultad de Ciencias Astronómicas y Geofísicas, Universidad Nacional de La Plata, Paseo del bosque S/N, 1900 La Plata, Argentina*

Accepted 1997 November 3. Received 1997 September 10; in original form 1996 October 9

ABSTRACT

A method of classifying generic orbits in arbitrary 2D and 3D potentials is presented. It is based on the concept of spectral dynamics introduced by Binney & Spergel that uses the Fourier transform of the time series of each coordinate. The method is tested using a number of potentials previously studied in the literature and is shown to distinguish correctly between regular and irregular orbits, to identify the various families of regular orbits (boxes, loops, tubes, boxlets, etc.), and to recognize the second-rank resonances that bifurcate from them. The method returns the position of the potential centre and, for 2D potentials, the orientation of the principal axes as well, should this be unknown. A further advantage of the method is that it has been encoded in a FORTRAN program that does not require user intervention, except for 'fine tuning' of search parameters that define the numerical limits of the code. The automatic character makes the program suitable for classifying large numbers of orbits.

Key words: celestial mechanics, stellar dynamics – galaxies: kinematics and dynamics.

1 INTRODUCTION

Stellar orbits constitute the basic set of building blocks for the dynamics of galaxies. This is so because galaxies are, to a large extent, collisionless systems and thus orbits are a well-defined concept. Although the most important function to a modeller is the phase-space distribution function, it is not the explicit dependence on phase-space coordinates, but an implicit one given by an underlying orbital structure, that is sought after. The fundamental problem of determining whether a self-consistent dynamical model can be built with a given potential, ultimately depends on the orbital structure supported by the potential. In recent times, a new approach to building 3D dynamical models has been developed that expressly makes use of a suitable exploration of the orbits supported by the potential (e.g. Schwarzschild 1979; Richstone 1980, 1984; Statler 1987; Levison & Richstone 1987; see de Zeeuw 1994, section 3.4, for a review).

Although the Jeans theorem (Jeans 1915) apparently allows one to bypass detailed orbital knowledge to build self-consistent dynamical models, questions as to irregularity, chaos, and non-analytical integrals of motion, limit its use to integrable systems, or to those cases in which irregular orbits are not important (Binney 1982b). Given the proliferation of irregular orbits on more realistic dynamical systems it seems that, in the end, we need to get back to a proper orbital assessment, even if it is only to ensure that the system can be modelled without the difficulties introduced by the irregular orbits.

The orbital structure may have a bearing on other important issues: irregular orbits in a model in which the curvature of isopotentials changes rapidly with radius may limit its flattening, as conjectured by Binney (1982a). Even regular orbits may limit the flattening in triaxial systems (Miralda-Escudé & Schwarzschild 1989; Lees & Schwarzschild 1992; Pfenniger & de Zeeuw 1989; Schwarzschild 1993). The time-scale on which irregular orbits cover their allowed region in phase space may introduce another relaxation time-scale, much shorter than the two-body relaxation time-scale (Schwarzschild 1993; Merritt & Fridman 1996).

Integrable potentials (e.g. Stäckel 1890; Kuzmin 1956; Lynden-Bell 1962; de Zeeuw 1985) provide us with the basic templates for the most important regular orbits on generic potentials, and an insight into the way in which other orbits may arise. Their study, however, can carry us some distance only, as the problem of finding the factors in a potential that give rise to a particular orbital structure, or whether it is integrable or not, is still one of the fundamental unsolved problems of dynamics. Thus the study of the orbital structure of potentials of interest in stellar dynamics has remained an active enterprise (e.g. Binney 1982a and Miralda-Escudé & Schwarzschild 1989 for 2D potentials; Schwarzschild 1979, 1993, Richstone 1982, Merritt & de Zeeuw 1983, Gerhard & Binney 1985, Lees & Schwarzschild 1992 and Merritt & Fridman 1996 for non-rotating 3D potentials; Heisler, Merritt & Schwarzschild 1982, Mulder & Hooimeyer 1984 and Pfenniger 1984 for rotating 3D potentials; Wilkinson & James 1982, Pfenniger & Friedli 1991, 1993, Hasan, Pfenniger & Norman 1993 for potentials extracted from N-body systems). In these and related studies, it is important to have a tool that can quickly, and efficiently, classify the orbits obtained. Such a tool is the goal of the present work.

In Section 2 we summarize some notions about orbits and methods to classify them, and a consistent orbit nomenclature is presented. Section 3 describes the procedure by which lines are

extracted from the Fourier spectra, the first step of our classification scheme. In Section 4, all the previous information is used to develop a spectral classification method for 2D orbits. Section 5 shows how to obtain information about the potential centre and principal axes. Section 6 discusses some numerical points of our classifier. Section 7 presents the results obtained for orbits in 2D potentials. In Section 8 we extend the spectral classification scheme to 3D orbits. In Section 9 we present results obtained in 3D potentials. Section 10 presents our conclusions.

2 PRELIMINARY NOTIONS

Here we review some notions about orbits, their structure in phase space, methods of classifying them, orbital resonances and their role in parenting orbit families. The basic characteristics of each orbit that we will use in the spectral classification are identified, and an orbital nomenclature that synthesizes the most important orbit properties is introduced.

2.1 Integrals of motion and types of orbits

Orbits are shaped, to a large extent, by their isolating integrals (Binney & Tremaine 1987, hereafter BT87, section 3.1.1). Each such integral lowers by one dimension the region open to an orbit. If the number M of isolating integrals equals the number N of degrees of freedom, the orbital manifold is diffeomorphic to an N-dimensional torus, i.e. we can find a one-to-one map between the manifold and the torus that covers both completely (Arnold 1989, chapter 10; Lichtenberg & Lieberman 1992, section 1.3). Such orbits are called *regular* and their motion is quasi-periodic. For them we can define a special set of canonical coordinates, the so-called action–angle variables, in which the motion is constant in the actions and there is uniform rotation in the angle variables.

A regular orbit with additional isolating integrals ($M > N$) is further constrained. This occurs when there are resonances between the rotation of two or more angle variables, in which case the orbit is no longer dense on the torus. Since the rotation frequencies in the angle coordinates are fixed on a given orbital manifold, all orbits sharing the same manifold will be identical, except for a phase difference (even when $M = N$, there is an infinite number of orbits sharing the same orbital manifold). We thus speak of non-resonant ($M = N$), and resonant ($M > N$) orbital tori; although both are dense in the integrable region of phase space, only the former form a set of non-zero measure.

If there is full resonance (i.e., $M = 2N - 1$), the orbit closes on itself after a finite number of turns around the torus and it becomes periodic. Such orbits are very important, because they can generate their own family of orbits: if a periodic orbit is stable (BT87, p. 175), neighbouring orbits will move on concentric tori nested around the stable periodic orbit and form an orbital family. A method for identifying all stable periodic orbits thus gives us the regular orbital families supported by a given potential. Although numerical methods have been devised (e.g. Contopoulos & Magenant 1985; Pfenniger & Friedli 1993), these require a long, detailed examination of a large set of orbits.

Irregular orbits, on the other hand, do not have such a torus-like structure. They have, in general, complicated shapes and may be chaotic, in the sense of having exponential divergence of neighbouring orbits. Irregular orbits make difficult the construction of self-consistent models; the ones dense on the energy manifold, for instance, are limited by the equipotential surfaces, which, except for a spherical configuration, are different in shape from the isodensity surfaces of the corresponding mass distribution.

2.2 Orbit classification methods

The dimensionality of the region in which an orbit moves is the feature exploited by a popular method used to classify orbits in 2D potentials. The surface of section (SoS) (see e.g. BT87, section 3.3) is usually taken as a 2D cut of phase space. Intersections of an orbit with this section lie in a region with one dimension less than the original orbital manifold. Thus, in 2D potentials, irregular orbits, open regular orbits, and closed regular orbits define a 'sea', a line, and a finite set of points, respectively. Unfortunately, distinguishing among these cases involves a visual inspection of the section, and this results in a subjective and time consuming procedure. Besides, this method is not yet generalized to 3D potentials: a 2D section of the corresponding 5D energy manifold does not suffice to disentangle orbits moving on three or fewer dimensions (regular) from those moving on higher dimensions (irregular).

Another, increasingly popular, way to ascertain whether or not an orbit is regular is by the computation of its Lyapunov characteristic exponents (see e.g. Lichtenberg & Lieberman 1992), which give the exponential rate at which nearby trajectories diverge from the original one. It can be shown that a regular orbit has vanishing Lyapunov exponents; unfortunately, it is not clear whether an irregular orbit will necessarily have at least a non-zero real exponent. A positive real exponent signals the onset of a more disordered behaviour called chaos. (Although it is generally assumed that irregular and chaotic orbits are the same in Hamiltonian systems, this has not been proven in general.) Another drawback is that their numerical computation is difficult and time consuming. Additionally, unlike SoSs and the method to be presented here, the Lyapunov exponents do not give any further information regarding regular orbits, for instance, whether or not they are closed (Merritt & Valluri 1996).

There is another technique first introduced in stellar dynamics by Binney & Spergel (1982, 1984), and more recently extended in a different form by Laskar (1993), that relies on one of the fundamental properties of regular orbits: the fact that they move winding on a torus-like manifold and are thus quasi-periodic. Then, the Fourier spectra of the time series of the coordinates of a regular orbit should consist of discrete lines the frequencies of which can be expressed as integer linear combinations of the frequencies associated to the N angle variables. We will call these latter frequencies the *base frequencies* (BFs) to stress this property. If the orbit is further constrained, this will manifest itself in a reduced number of incommensurate BFs. If the orbit is closed, only one BF will exist. Irregular orbits, not being quasi-periodic, will produce Fourier spectra with lines the frequencies of which cannot be reduced to integer combinations of less than N frequencies. If the orbit is chaotic, the spectrum will be continuous (Tabor 1989, section 4.5.b).

So, if we compute the Fourier transform of the coordinates of an orbit, identify its peaks, extract the corresponding frequencies, and look for the BFs (if any), then the orbit could be classified. It turns out that a closer inspection of the Fourier spectrum can further disclose the orbit family to which the orbit belongs as well as the resonance and resonance rank of its parent, so we need first to describe these orbit families and their resonant parents and examine the concept of resonance rank.

2.3 Orbit families and rank of resonant parents

We start with 2D potentials. Apart from the trivial case of a particle at rest at the centre of a potential, the simplest orbits are the axial

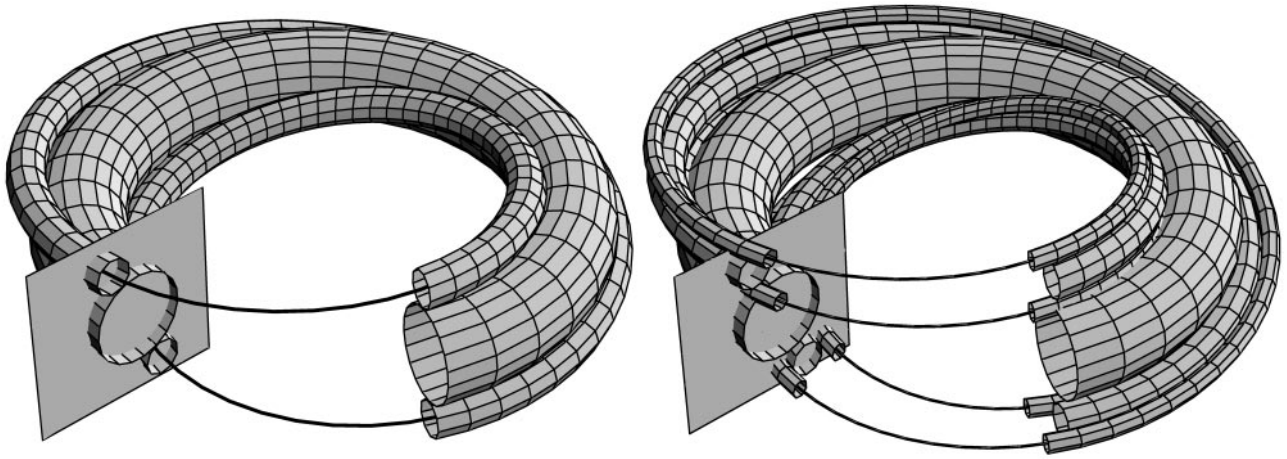

Figure 1. Left: a stable first-rank 2:1 resonance orbit, and surrounding torus of one of its daughter orbits, around an open zeroth-rank torus. Right: second-rank resonant orbit and daughter orbit torus. All tori have been cut open to show the closed resonant parent in each case. A plane corresponding to an SoS has been introduced in both graphs.

orbits that move along each axis; these orbits are simple loops in phase space and appear as single points at the origin in the SoS orthogonal to them. When they are stable, the circulation of the daughter orbits around the parent orbit introduces an additional frequency which, in general, is not commensurable with the unique BF of its parent; these orbits are dense in the nested tori. Additionally, these daughter orbits do not have a fixed sense of rotation around the centre and are dense within a box-like region in configuration space. Their intersections with an SoS lie on loops that encircle the single fixed point that corresponds to the axial orbit. We will call these *π-box orbits* to emphasize their frequency incommensurability and the fact that they do not circulate around the origin.

As we move through the nested tori, the BFs will change in general, and for an infinite but countable number of tori, the frequencies will be commensurable (only one frequency will be a BF) and the corresponding orbits will close again. These first-rank resonance orbits are analogous to closed Lissajous figures in configuration space and produce a finite set of points around the axial orbit in the SoS. There is, however, a fundamental difference between the 1:1 resonance and the others: while the former corresponds to a simple loop in configuration space and thus has a definite sense of rotation – except for the cases of 0 or π coordinate phase differences – the latter do not. To emphasize this important fact we will call the 1:1 resonance the *closed 1:1 loop*, and the rest, the *closed m:n boxes*. As we move in energy, some of these resonant orbits may become stable and generate, in turn, families of daughter orbits that move on nested tori around them and inherit the parent property of rotating, or not, around the centre (Fig. 1). Similar to the case of π-boxes, these daughter orbits will be open in general, as the newly introduced oscillating frequency will be, in general, incommensurable with the parent unique BF. We will call these orbits *open m:n boxes* and *open 1:1 loops*, depending on the parent identity.

When any of the daughter orbits spawned by a first-rank resonance becomes resonant itself, we obtain a second-rank resonant orbit. The orbit parenting process repeats itself on top of the orbits that become stable as we move further in energy. Their corresponding new daughter orbits appear in an SoS as islands that enclose the discrete set of points produced by the closed second-rank parent (Fig. 1). The second-rank resonance will be denoted by an extra integer k, which represents the number of turns which it has to give around the parent before closing on itself ($k = 2$ for the example in the figure). We will thus speak of *closed* and *open km:kn boxes* and *k:k loops*. The common factor k will signal the higher-rank resonance.

Table 1. 2D orbit classification according to resonance rank.

Resonance Rank		No circulation around the centre	Circulation around the centre
0^{th}	parent	axial	
	daughter	π-box (box)	
1^{st}	parent	closed *m:n* box (closed boxlet)	closed 1:1 loop (closed loop)
	daughter	open *m:n* box (open boxlet)	open 1:1 loop (open loop)
2^{nd}	parent	closed *km:kn* box (higher order resonance boxlet)	closed *k:k* loop (closed looplet)
	daughter	open *km:kn* box (?)	open *k:k* loop (open looplet)

In three dimensions, the 1:1 loops orthogonal to the long and short potential axes give rise to corresponding tubes, when perturbed orthogonal to the loop plane (the intermediate axis loops are unstable and do not generate tubes, BT87 p. 154). The planar closed *m:n* boxes, when stable, can also spawn families of 3D orbits. All of these orbits may have an additional resonance in the direction orthogonal to their planar parent, *l:m:n* resonances, which could give rise to higher rank orbits as well.

Table 1 summarizes our nomenclature for 2D orbits and includes, in parenthesis, the corresponding customary name. Although it is common to refer to *m:n* resonances as boxlets, and a rich and whimsical nomenclature exists for some of the resonances (banana 2:1, fish 2:3, pretzel 3:4, etc.), in our scheme we prefer to have all orbits that circulate around the origin be called 'loops' and all that

do not 'boxes'. Parent and daughter orbits share the name of the parent resonance, the distinction being the 'closed' or 'open' qualification. A common factor in a resonance specification naturally indicates a second-rank resonance. (Although there is no provision for third and higher-ranked resonances, this notation could be extended to them.) These are the basic properties that define the morphology of an orbit and which are of utmost importance when trying to build a self-consistent model. As we will see, the procedure presented in this work allows us to recognize all these orbital properties. The extension of this nomenclature to 3D orbits is presented in Section 8.

3 EXTRACTING WAVES FROM THE FOURIER SPECTRUM

We now describe in some detail the way in which line frequencies are extracted from the computed Fourier spectrum of the orbit, the first step in the orbit classification.

3.1 One wave

To classify, we must compute, from its Fourier spectrum, the sinusoidal components that build up the orbit. Here we develop a method for extracting them, which differs from that used by Binney & Spergel (1982). We first concentrate on the issue of a unique wave.

Let us suppose that at times $t_k = k\delta$, $k = 0, \ldots, N-1$ (N even), the values $z_k \equiv z(t_k)$ of a complex function are recorded. The discrete Fourier transform of the set $\{z_k\}$ is

$$Z_j = \frac{1}{N}\sum_{k=0}^{N-1} z_k \exp\left(-\frac{\mathrm{i}2\pi jk}{N}\right), \tag{1}$$

where $j = -N/2+1, \ldots, N/2$. The Fourier spectrum then consists of N waves of amplitudes $|Z_j|$, phases $\gamma_j \equiv \arg(Z_j)$, and frequencies $\omega_j = 2\pi f_j = 2\pi j/N\delta$. Let us suppose, to begin with, that $z(t)$ is a plane wave,

$$z(t) = A\mathrm{e}^{\mathrm{i}(\omega_s t+\phi)}. \tag{2}$$

Our goal is to recover the amplitude A, initial phase ϕ, and frequency ω_s from the Fourier spectrum. For convenience, we put $\omega_s = 2\pi s/N\delta$, where s is a (real) number. If s equals any of the j (i.e., if $\omega_s = \omega_j$ for some j), then the resulting spectrum will be simply $Z_j = A\exp(\mathrm{i}\phi)$ if $j = s$, and $Z_j = 0$ otherwise. This shows that the choice of a normalization $1/N$, a negative sign on the exponent, and indices $-N/2+1 \le j \le N/2$ for the Fourier transform, yields the correct results. Were any of the above chosen other way, further transformations would have been required to get the correct answers.

When s is not an integer, every frequency ω_j has a non-zero amplitude. To see this, we now obtain an expression for the amplitudes. Replacing equation (2) into equation (1) yields

$$Z_j = \frac{A\mathrm{e}^{\mathrm{i}\phi}}{2N\{1-\cos[2\pi(s-j)/N]\}}(\rho+\mathrm{i}\sigma), \tag{3}$$

where we have defined

$$\rho \equiv 1-\cos[2\pi(s-j)/N]-\cos[2\pi(s-j)] + \cos[2\pi(s-j)(N-1)/N], \tag{4a}$$

$$\sigma \equiv \sin[2\pi(s-j)/N]-\sin[2\pi(s-j)] + \sin[2\pi(s-j)(N-1)/N]. \tag{4b}$$

Now, from equation (3) we can compute

$$|Z_j| = \frac{A}{N}\left\{\frac{1-\cos[2\pi(s-j)]}{1-\cos[2\pi(s-j)/N]}\right\}^{1/2}. \tag{5}$$

This equation includes $s = j$ as a limiting case. Note that the phase ϕ has disappeared. To solve for the remaining unknowns A and s, since we have N equations – one for each value of j – we may take any two, namely, $j = m_1$ and $j = m_2$. Now we can eliminate A by dividing them:

$$\frac{1-\cos[2\pi(s-m_1)]}{1-\cos[2\pi(s-m_2)]} = \frac{1-\cos\left[\frac{2\pi(s-m_1)}{N}\right]}{1-\cos\left[\frac{2\pi(s-m_2)}{N}\right]}\frac{|Z_{m_1}|^2}{|Z_{m_2}|^2}. \tag{6}$$

m_1 and m_2, however, are integers, so the first member is equal to 1. Using the identity $1-\cos\alpha = 2\sin^2(\alpha/2)$, and defining $\kappa = \pi(m_2-m_1)/N$ and $\alpha = \pi(s-m_1)/N$ to simplify the notation, the foregoing equation becomes

$$|\cos\kappa - \cot\alpha\sin\kappa| = \frac{|Z_{m_1}|}{|Z_{m_2}|}. \tag{7}$$

Depending on whether $\cos\kappa - \cot\alpha\sin\kappa \gtrless 0$, we have

$$\tan\alpha = \frac{\sin\kappa}{\cos\kappa \mp |Z_{m_1}|/|Z_{m_2}|}. \tag{8}$$

Now, since $|m_2-m_1| < N/2$, then $|\kappa| < \pi/2$, and the inequalities $\cos\kappa - \cot\alpha\sin\kappa \gtrless 0$ can be written in the form $\cot\alpha\tan\kappa \lessgtr 1$. However, since $|s-m_1| < N/2$ also, α is an angle of the first or fourth quadrants. In the former case, the upper sign implies $m_1, m_2 < s$, and the lower sign $m_1 < s < m_2$. In the latter case, the upper sign implies $s < m_1, m_2$, and the lower sign $m_2 < s < m_1$. So, if we always choose m_1 and m_2 such that s is in between, we can safely take the lower sign. Now we are able to compute α (and therefore s) from equation (8) with the plus sign.

Our next goal is to compute the amplitude. Since we now know s, from equation (5) we have, taking any j,

$$A = N|Z_j|\left[\frac{1-\cos\frac{2\pi(s-j)}{N}}{1-\cos 2\pi(s-j)}\right]^{1/2}. \tag{9}$$

Now the phase ϕ. Equation (3) can be written

$$|Z_j|(\cos\gamma_j + \mathrm{i}\sin\gamma_j) = \frac{A(\cos\phi+\mathrm{i}\sin\phi)(\rho+\mathrm{i}\sigma)}{2N\{1-\cos[2\pi(s-j)/N]\}}. \tag{10}$$

From the foregoing equation we obtain

$$\tan\phi = \frac{\{1-\cos[2\pi(s-j)/N]\}(\rho\sin\gamma_j - \sigma\cos\gamma_j)}{\{1-\cos[2\pi(s-j)/N]\}(\rho\cos\gamma_j + \sigma\sin\gamma_j)}. \tag{11}$$

We keep the factors $\{1-\cos[2\pi(s-j)/N]\}$ so that we do not lose any signs in finding the quadrant of ϕ.

There only remains the choice of m_1, m_2, and the j of equations (8) and (11); so far, the only restriction is that s must be between m_1 and m_2. We note that, from equation (5), the amplitudes of the (discrete) spectrum peak at the value of j nearest to s. Now let us suppose that $z(t)$ contains another plane wave with frequency s'. If we choose m_1 close to s', the information of m_1 regarding s will be greatly shadowed by the presence of this second wave. Therefore, it is convenient to choose m_1 to be the value of j nearest to s (i.e., the frequency at the peak), and, for the same reason, to take m_2 such that $m_2 = m_1 \pm 1$ and s lies in between, such that m_2 is the adjacent frequency with the greater amplitude. It is also convenient to choose $j = m_1$ to save numerical work, for the computations involving m_1 can then be reused.

Thus, the computation proceeds as follows. First, we look for the greatest amplitude in the Fourier spectrum; its frequency j will be m_1, and that of its neighbour with the greater amplitude will be m_2. Then equations (8), (9), and (11) are used, in turn, to obtain s (or ω_s), A, and ϕ. Once this triplet has been computed, we say, following Binney & Spergel (1982), that a 'line' has been extracted.

3.2 More than one wave

If $z(t)$ has more than one plane wave, first we compute A, ϕ, and s for the wave corresponding to the greatest peak in the spectrum. Next, we subtract its contribution (equation 3):

$$Z_j \leftarrow Z_j - \frac{A\mathrm{e}^{\mathrm{i}\phi}}{2N\{1-\cos[2\pi(s-j)/N]\}}(\rho+\mathrm{i}\sigma). \tag{12}$$

This eliminates the peak from all the frequencies. Now we look for the second greatest peak, extract the corresponding line, and so on. Unfortunately, this is not an exact procedure as with a single wave, because the first line will be contaminated with information regarding the others; the closer two lines are, the more inaccurately their parameters will be computed. This will affect the subtraction, and so errors are carried into the next lines. To reduce this problem, we extract the lines twice. For the second extraction, we start from the naked spectrum resulting from the complete first extraction. Then we add the first extracted line which, being alone, will be much less contaminated than before. We re-extract this line, and subtract it with the new computed values. Then we add the second line, and so on. As we will see below, this procedure proves to be very successful, even in cases in which two nearby lines have comparable amplitudes, thus strongly influencing each other.

A subtlety remains to be considered, since, in the case of an orbit, $z(t)$ is a real function. Let us suppose that

$$z(t) = A\cos(\omega t+\phi) = \frac{A}{2}\left[\mathrm{e}^{\mathrm{i}(\omega t+\phi)} + \mathrm{e}^{-\mathrm{i}(\omega t+\phi)}\right]. \tag{13}$$

We are not interested in the values $A/2$, $\pm\phi$, and $\pm\omega$ of the plane waves, but in the values A, ϕ, and ω of the original cosine. We may merely take the positive portion of the spectrum, extract the line, and double the amplitude. However, since every peak found in the positive portion of the spectrum has a negative twin, the tails of the latter may make a non-negligible contribution to the former, particularly if the peak is at a frequency near zero. So instead of simply neglecting the negative portion, we subtract both peaks simultaneously.

We measured the accuracy of the method with the incompleteness parameter ι (Binney & Spergel 1982):

$$\iota^2 = \frac{\sum_{n=1}^{10}\left\{\omega_{\max}^2\left[(\delta x_n)^2+(\delta y_n)^2\right]+(\delta\dot{x}_n)^2+(\delta\dot{y}_n)^2\right\}}{\sum_{n=1}^{10}\left[\omega_{\max}^2\left(x_n^2+y_n^2\right)+\dot{x}_n^2+\dot{y}_n^2\right]}, \tag{14}$$

where $x_n \equiv x(t_n)$, etc; $\delta x_n = x_n - \bar{x}_n$, etc; $[x(t), y(t)]$ is the orbit, $[\bar{x}(t), \bar{y}(t)]$ is the orbit reconstructed from the lines; $t_1, \ldots, t_{10}$ are ten reference times, and $\omega_{\max}$ is the frequency of the line of maximum amplitude. To compare our method of extraction with that used by Binney & Spergel (1982), we choose an orbit in the potential $\Phi = \ln r^2$, with initial conditions $(x_0, y_0, \dot{x}_0, \dot{y}_0) = (0.6, 0, 0, 1.115)$, which yields an energy per unit mass $\mathcal{E} = -0.4$ and an eccentricity $e \simeq 0.22$. We extracted all the lines with amplitudes greater than 5×10^{-5} times the greatest (Fig. 2). Taking ten equidistant times along a period of integration of 100 units ($\simeq$45 orbital periods), we obtained $\iota = 0.003$. As a comparison, Binney & Spergel obtained $\iota = 0.025$ for this same orbit.

Also, the parameter ι allows us to measure the benefits of the double-extraction algorithm mentioned above. We generated synthetic orbits with two sines in x and two in y, all with the same amplitude, and with phases 0° and 90°, respectively. On each coordinate one wave had a fixed frequency, and the other had a frequency which was varied between experiments. We computed ι

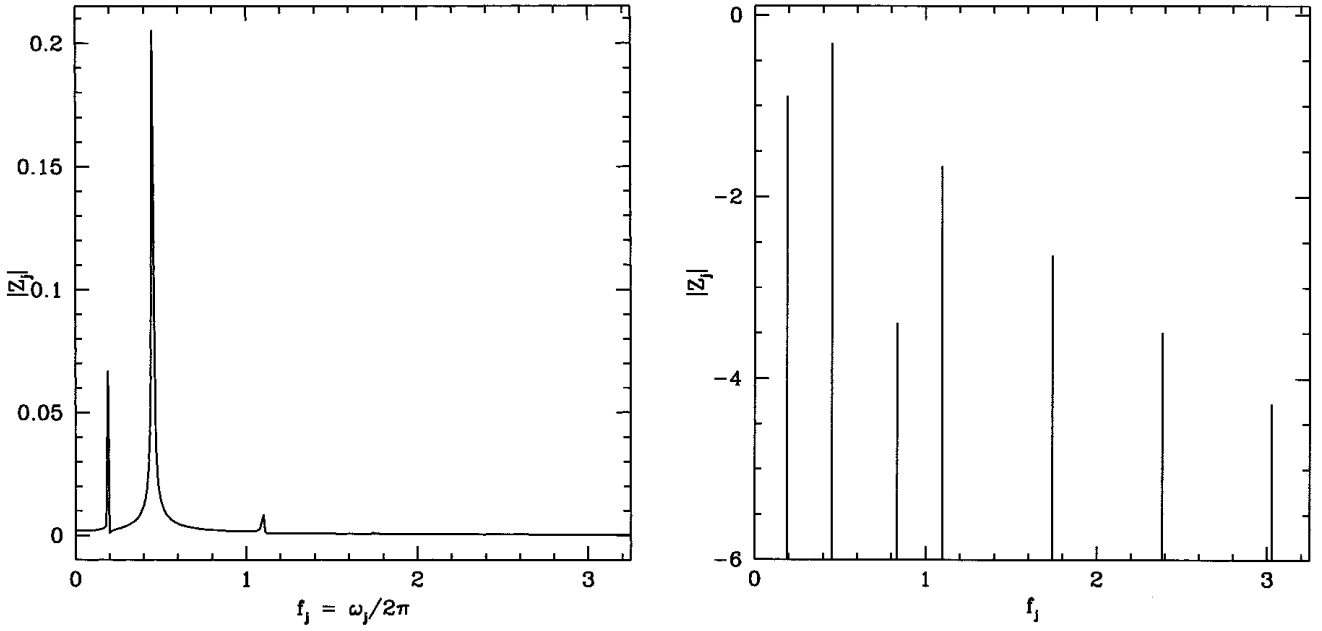

Figure 2. Left: Fourier spectrum of the x coordinate of the orbit mentioned in the text. The points have been joined with solid lines in order to improve visibility. Right: amplitudes of the extracted lines of the orbit, on a logarithmic scale, as a function of frequency.

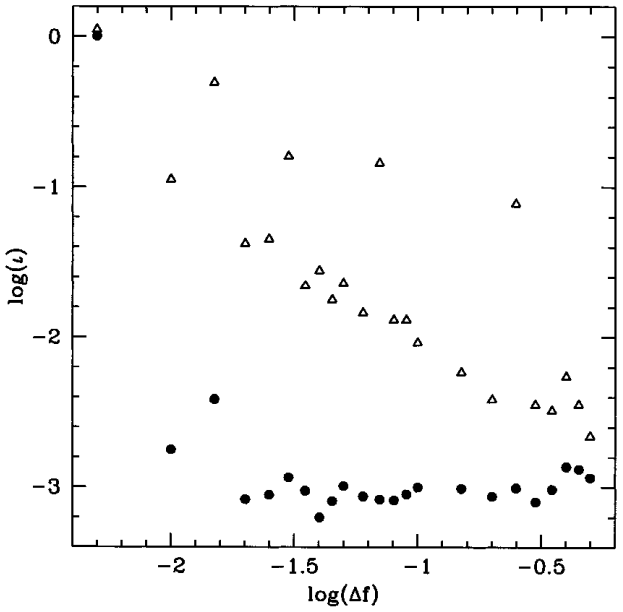

Figure 3. Incompleteness parameter ι as a function of the separation in frequency of two equal-amplitude lines. Open triangles show the results when a single extraction of the lines is performed; filled circles are the results when the lines are extracted twice.

as a function of the separation of the frequencies. Fig. 3 shows how a second extraction greatly improves the computation of the lines.

4 SPECTRAL CLASSIFICATION OF 2D ORBITS

Now we describe the spectral characteristics which allow us to classify 2D orbits. We denote as F_x and F_y the frequencies corresponding to the greatest amplitudes of the x and y line spectra, respectively, calling them *dominant frequencies*. Note that a dominant frequency is not neccesarily a base frequency. We assume a coordinate system with its x axis aligned with the major axis of the potential, and with its origin coinciding with the centre of the potential.

In all the examples in this section, we have used one of the following potentials, where v_0^2, R_c, q, and R_e are constants, and $R = (x^2 + y^2)^{1/2}$.

(i) Logarithmic potential:

$$\Phi_L(x, y) = \frac{v_0^2}{2} \ln\left[R_c^2 + x^2 + \frac{y^2}{q^2}\right]. \quad (15)$$

(ii) Binney potential (Binney 1982a):

$$\Phi_B(x, y) = \frac{v_0^2}{2} \ln\left[R_c^2 + x^2 + \frac{y^2}{q^2} - \frac{R(x^2 - y^2)}{R_e}\right]. \quad (16)$$

(iii) Hénon–Heiles potential (Hénon & Heiles 1964):

$$\Phi_H(x, y) = \frac{1}{2}\left(R^2 + 2x^2 y - \frac{2}{3}y^3\right). \quad (17)$$

All our orbits were integrated using a Runge–Kutta–Fehlberg integrator, with Cash–Karp coefficients (Press et al. 1994), with double precision. We get an energy conservation better than one part in 10^6 in all cases.

4.1 π-box orbits

Since a π-box orbit does not put any constraint on the ratio F_y/F_x, this will in general be an irrational number, and we are already left with two independent frequencies, i.e., neither of the two is a multiple of the other, nor they are multiples of a common third frequency. Since the y axis is aligned with the minor axis of the potential, one expects, in general, that $F_y > F_x$. Sometimes it occurs that $F_y < F_x$, but in these cases there is always another y line symmetric to F_y with respect to F_x, and with a similar amplitude; this is explained by a simple model of a librating oscillator (Binney & Spergel 1982). When we come upon this, we do not hesitate to call the latter frequency the dominant frequency.

The remaining lines must be integer linear combinations of F_x and F_y, in order to ensure regularity. Thus, we already have a criterion to classify an orbit as a π-box orbit: it must have an *irrational ratio of dominant frequencies, and integer linear dependence of the rest of the spectra with respect to the former*. Fig. 4 shows a typical π-box orbit and some of its extracted lines.

4.2 First- and second-rank box orbits

The basic spectral difference between m:n box orbits and π-box orbits is that F_x and F_y are commensurable in the former case, corresponding therefore to the same BF. Moreover, the resonance is simply the quotient F_y/F_x. We will call *unit frequency* the greatest common divisor of F_x and F_y, i.e., the frequency in terms of which they are integers. This unit frequency may or may not be occupied by a line. If the rest of the lines are multiples of the BF, the m:n box orbit is closed, for we have only one BF. However, if one line is independent (i.e., its ratio with the BF is irrational), and the rest of the lines are linearly dependent of both the BF and the new line, then we have an open m:n box orbit. This new frequency can in general be associated with the libration of the orbit around its parent. Figs 5 and 6 show examples of closed and open m:n box orbits, and their extracted lines.

As we saw, a resonant orbit may parent orbits with the same resonance but higher rank. This happens, for example, when a line, in what otherwise would have been an open m:n box orbit, matches a submultiple of the unit frequency: now this line can be regarded as a new unit frequency, thus multiplying m and n by the same factor, and generating a closed second-rank m:n box, i.e., a closed km:kn box. If there is an additional BF, the orbit becomes open. Fig. 7 shows an example of a km:kn box orbit and some of its lines.

4.3 First- and second-rank loop orbits

From a spectral point of view, loop orbits and m:n box orbits differ only in the order of the resonance: in the former, the ratio $F_y/F_x = m{:}n = 1{:}1$ (F_x is, in this case, the frequency of rotation around the centre of the potential, and we talk of a 1:1 *loop orbit*), whereas in the latter, this ratio is any other rational. As before, if the orbit is closed, the other lines will be multiples of a single BF; but if it is open, the rest of the lines must be integer linear combinations of two BFs. Fig. 8 shows an example of an open 1:1 loop orbit and its extracted lines.

A closed loop orbit can close itself after one turn (a 1:1 resonance), but it can also close itself after k turns; we speak in this case of a k:k loop orbit. Any one of these may give birth to a family of open librating loop orbits with the same k:k resonance. As before, we can compute this higher-rank resonance with the aid of

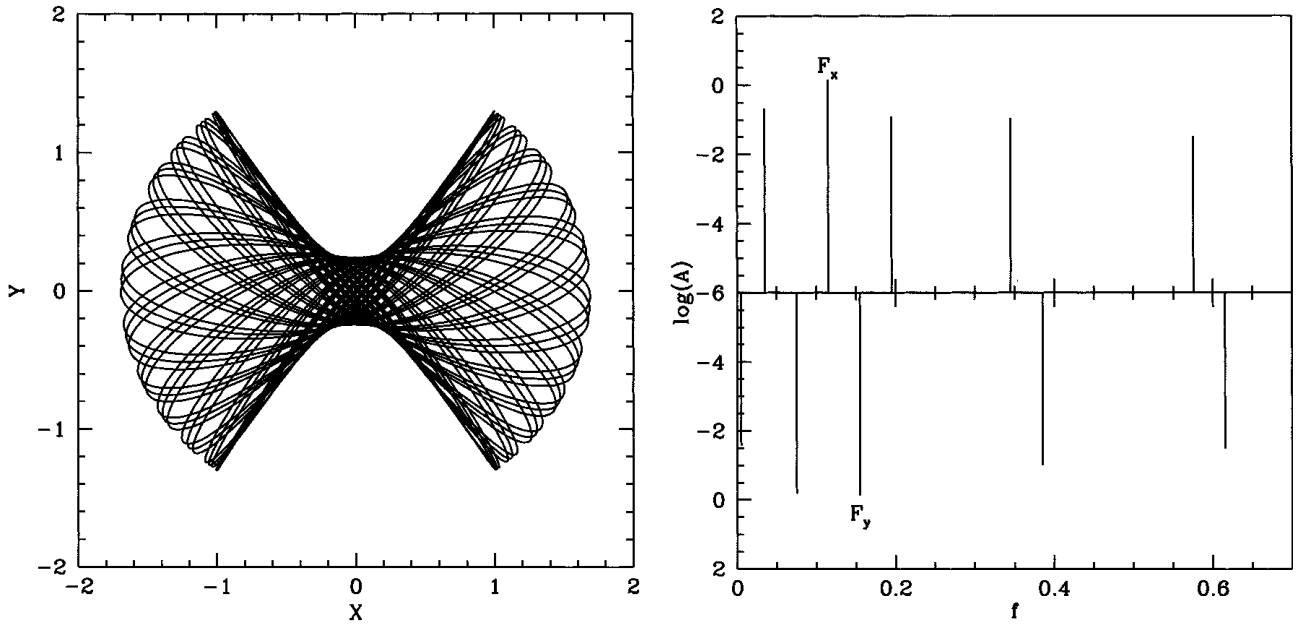

Figure 4. Left: a π-box orbit, obtained by launching a particle in the potential Φ_L, with $v_0^2 = 1$, $R_c = 0.14$, $q = 0.9$, and initial conditions $x_0 = 1$, $y_0 = 1.3$, $\dot{x}_0 = 0$, and $\dot{y}_0 = 0$. Right: amplitudes of the extracted lines versus frequency for the x (top) and y (bottom) coordinates. In all of the examples presented, only those lines with amplitudes greater than 2 per cent of the strongest line, in each coordinate, are shown; this number is further constrained to five lines at most (i.e., only the lines used to classify the orbit are shown). The actual spectra bear many more lines. The dominant frequencies are also marked.

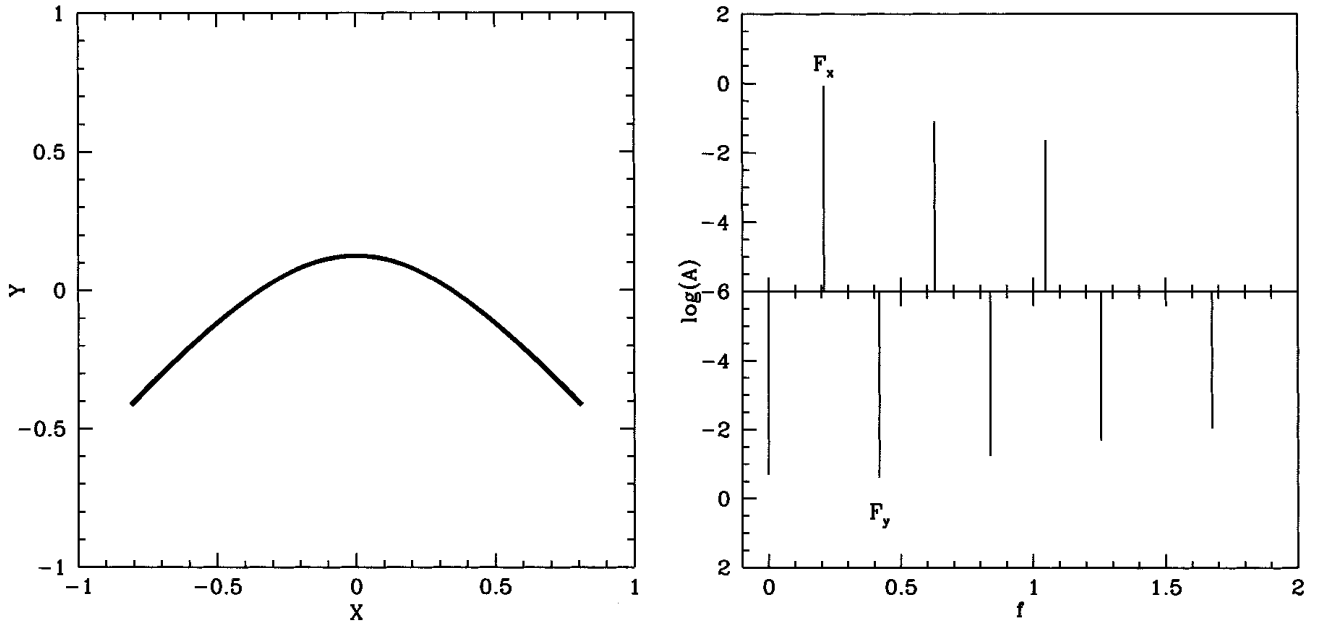

Figure 5. As Fig. 4, but for a closed 2:1 box orbit, obtained by launching a particle in the potential Φ_L, with $v_0^2 = 1$, $R_c = 0$, $q = 0.7$, and with initial conditions $x_0 = 0$, $y_0 = 0.13$, $\dot{x}_0 = 1.83496$, and $\dot{y}_0 = 0$.

the unit frequency. Figs 9 and 10 show examples of a closed 3:3 and an open 2:2 loop orbits, respectively, and their extracted lines.

4.4 Irregular orbits

Since irregular orbits do not have line-like spectra that can be expressed as integer linear combinations of BFs (in the case of a truly chaotic orbit, even a continuum will appear), the lines extracted by our algorithm do not describe the orbit, i.e., its incompleteness parameter ι is large. Sometimes, irregular orbits are close to regular orbits in phase space, and this is reflected in their spectra; but they always have some extra lines that are not integer combinations of the BFs, which gives them away. Fig. 11 shows an example of an irregular orbit and its spectra.

A related problem is that of irregular orbits highly confined between regular regions, sometimes called 'sticky' or 'semi-stochastic' orbits

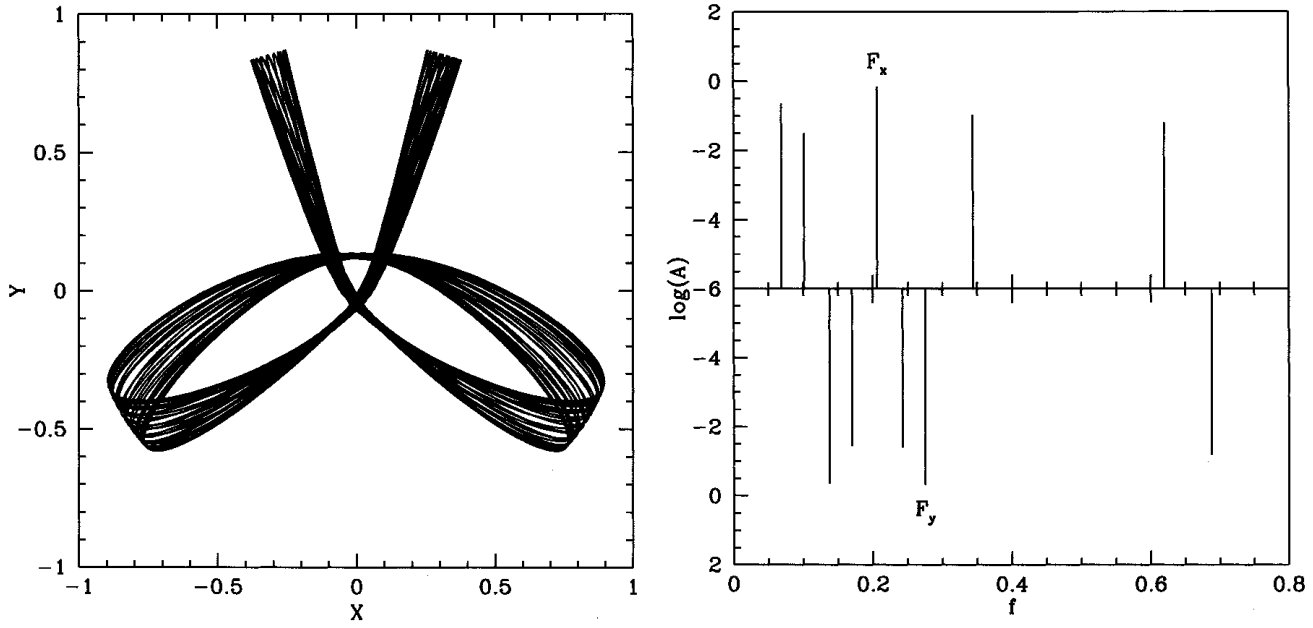

Figure 6. Same as Fig. 4, but for an open 4:3 box orbit, obtained by launching a particle in the potential Φ_L, with $v_0^2 = 1$, $R_c = 0$, $q = 0.9$, and with initial conditions $x_0 = 0$, $y_0 = 0.13$, $\dot{x}_0 = 1.95121$, and $\dot{y}_0 = 0.25$.

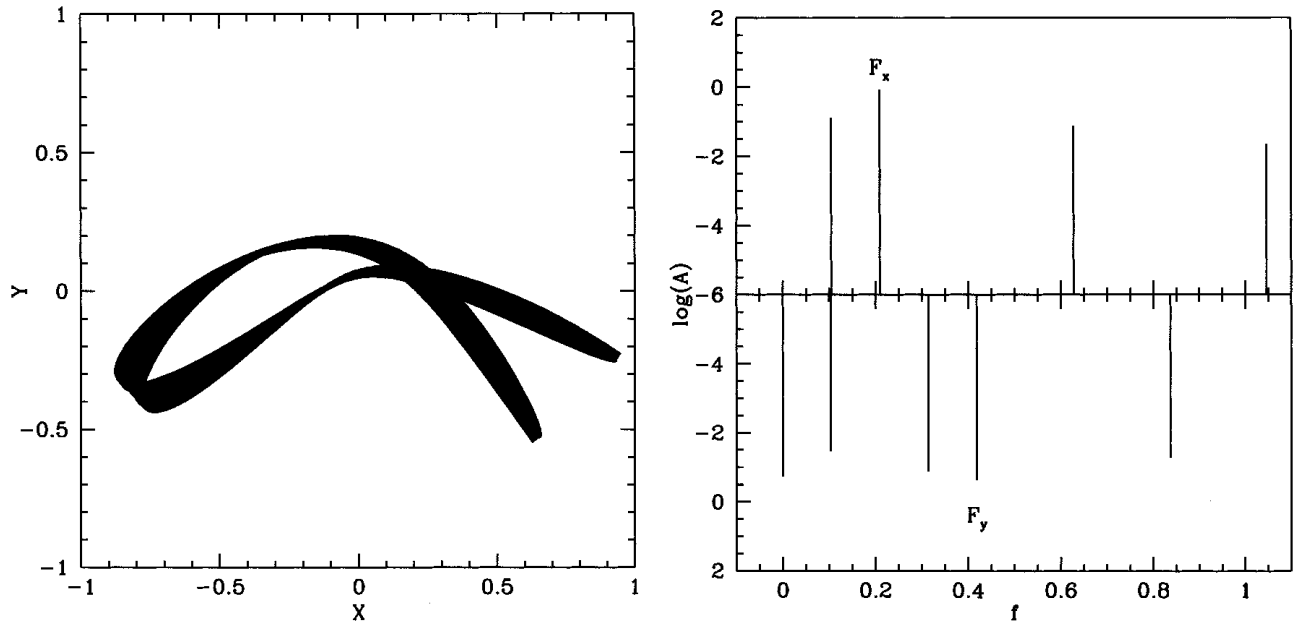

Figure 7. Same as Fig. 4, but for an open 4:2 (second-rank) box orbit, obtained by launching a particle in the potential Φ_L, with $v_0^2 = 1, R_c = 0, q = 0.7$, and with initial conditions $x_0 = 0$, $y_0 = 0.064$, $\dot{x}_0 = 2.15750$, and $\dot{y}_0 = 0.36$.

(Goodman & Schwarzschild 1981). Such orbits must be given sufficient time to reveal their irregularity without ambiguity, i.e. they need longer integrations to be properly classified. An example will illustrate this. Let us launch a particle in the potential Φ_L with $(v_0, R_c, q) = (1, 0, 0.7)$, and initial conditions $(x_0, y_0, \dot{x}_0, \dot{y}_0) = (0, 0.2663, 1.39029, 0)$, so that $\mathcal{E} = 0$. This orbit is classified as a 2:1 box, but only if we integrate it for fewer than about 400 orbital periods. For longer integrations, our algorithm tells us that it is an irregular orbit. Fig. 12 shows what happens. Fig. 12(a) shows the $(y, \dot{y}, x = 0, \dot{x} > 0)$ SoS of this orbit, sampled during some 400 orbital periods. We can see how closely it resembles a regular orbit. Fig. 12(b) shows the same SoS for an integration five times longer. We see that the phase space visited has grown, and now we clearly have an irregular orbit. This is just what the algorithm showed. Unfortunately, there is no way to tell in these cases whether or not a further integration period is needed to establish a good classification. On the other hand, the orbit *has behaved* as a regular orbit during the time in which it was classified as such. This orbit was also integrated with a different implementation of the Runge–Kutta–Fehlberg integrator (Fehlberg 1968), attaining an energy conservation of one part in 10^9.

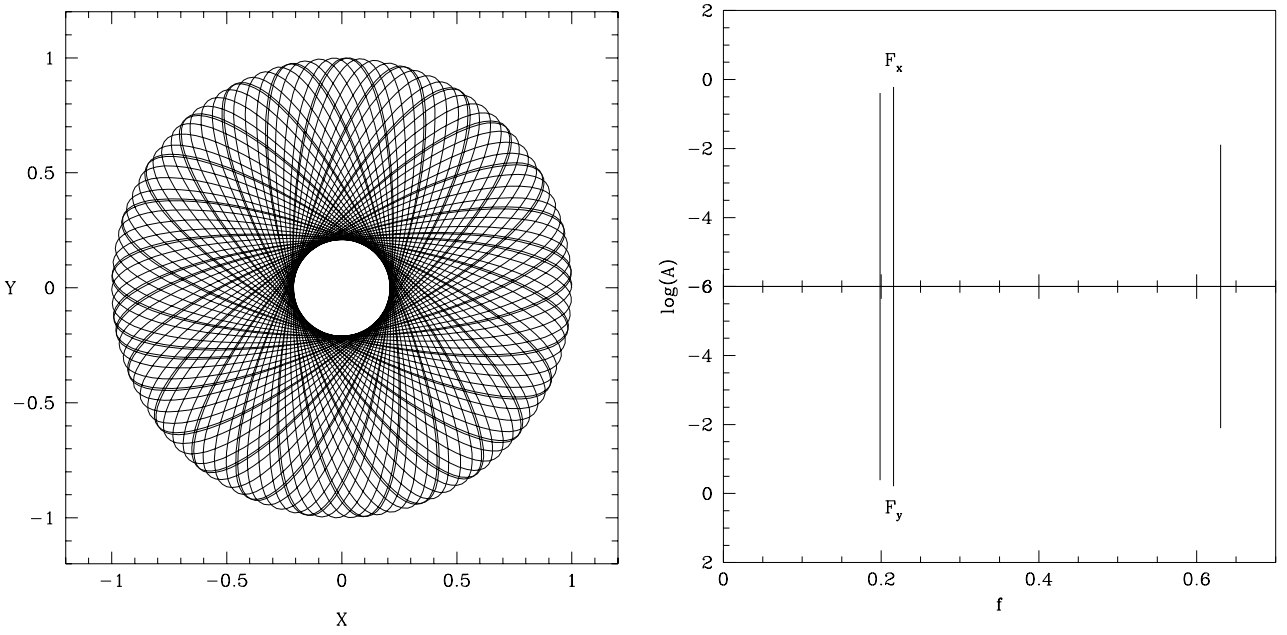

Figure 8. Same as Fig. 4, but for an open 1:1 loop orbit, obtained by launching a particle in the potential Φ_L, with $v_0^2 = 3$, $R_c = 1$, $q = 1$, and with initial conditions $x_0 = 1$, $y_0 = 0$, $\dot{x}_0 = 0$, and $\dot{y}_0 = 0.3$.

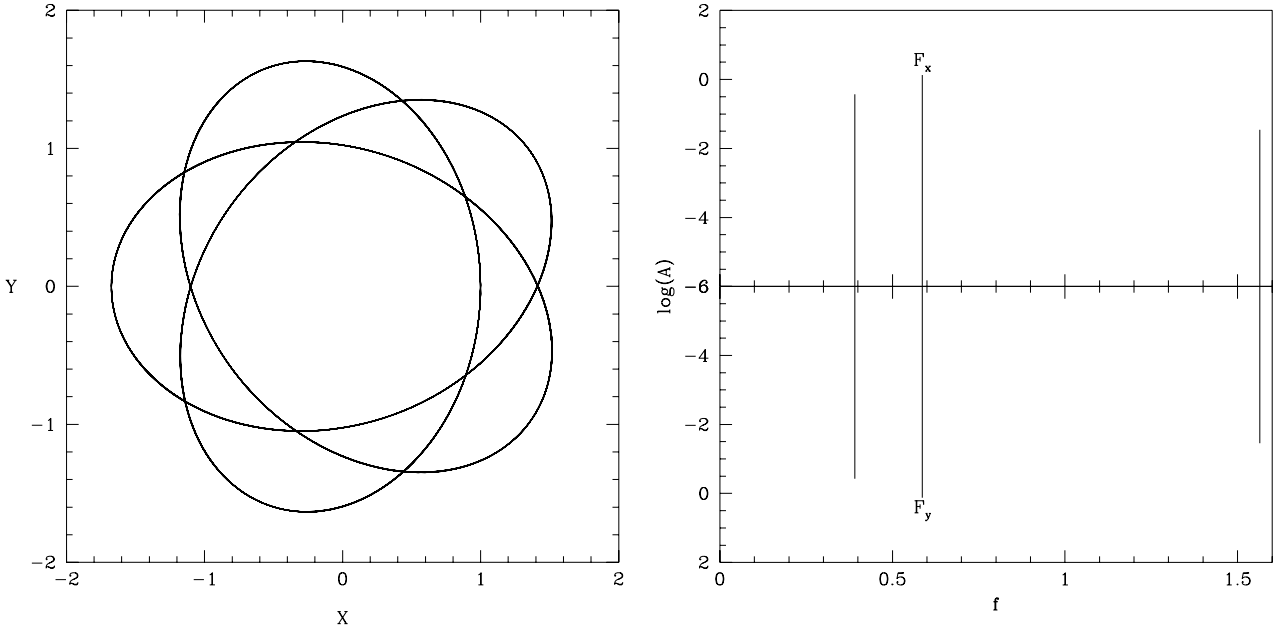

Figure 9. Same as Fig. 4, but for a closed 3:3 loop orbit, obtained by launching a particle in the potential Φ_L, with $v_0^2 = 40$, $R_c = 1$, $q = 1$, and with initial conditions $x_0 = 1$, $y_0 = 0$, $\dot{x}_0 = 0$, and $\dot{y}_0 = 6.28318$.

In the end we must distinguish between a mathematical definition of irregular orbit and a pragmatic one: although the above orbit is strictly speaking irregular, what matters to an astronomer is its behaviour during the time-scales of interest. The same orbit may play the role of a regular or irregular orbit, depending on the ratio of its diffusion time-scale in phase space to the dynamical time-scale of the modelled system (Goodman & Schwarzschild 1981; Schwarzschild 1993; Merritt & Fridman 1996). Our classifier gives the correct pragmatic answer.

Fig. 13 summarizes our 2D classification in terms of BFs and resonances between dominant frequencies.

5 WHEN THE ORIENTATION AND CENTRE OF THE POTENTIAL ARE UNKNOWN

The results of the previous section are based upon the assumptions that (i) the x axis of coordinates is on the long axis of the potential, and (ii) the origin of coordinates lies on the centre of the potential. This is fine if we know the potential a priori. However, we may want to classify an orbit which, for instance, has been obtained from an N-body simulation for which the orientation and centre of the system are unknown, or at least are not known with adequate accuracy. We will show here that the algorithm is able to classify

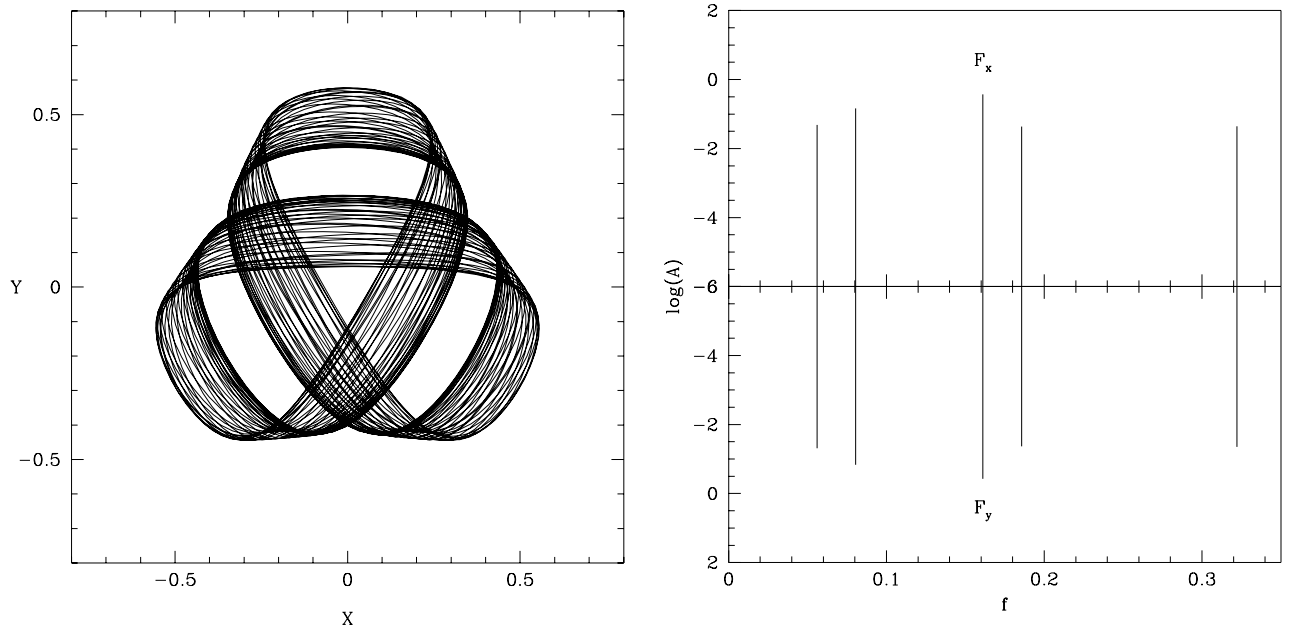

Figure 10. Same as Fig. 4, but for an open 2:2 loop orbit, obtained by launching a particle in the potential $\Phi_{\rm H}$ with initial conditions $x_0 = -0.48$, $y_0 = 0$, $\dot{x}_0 = 0$, and $\dot{y}_0 = 0.30768$.

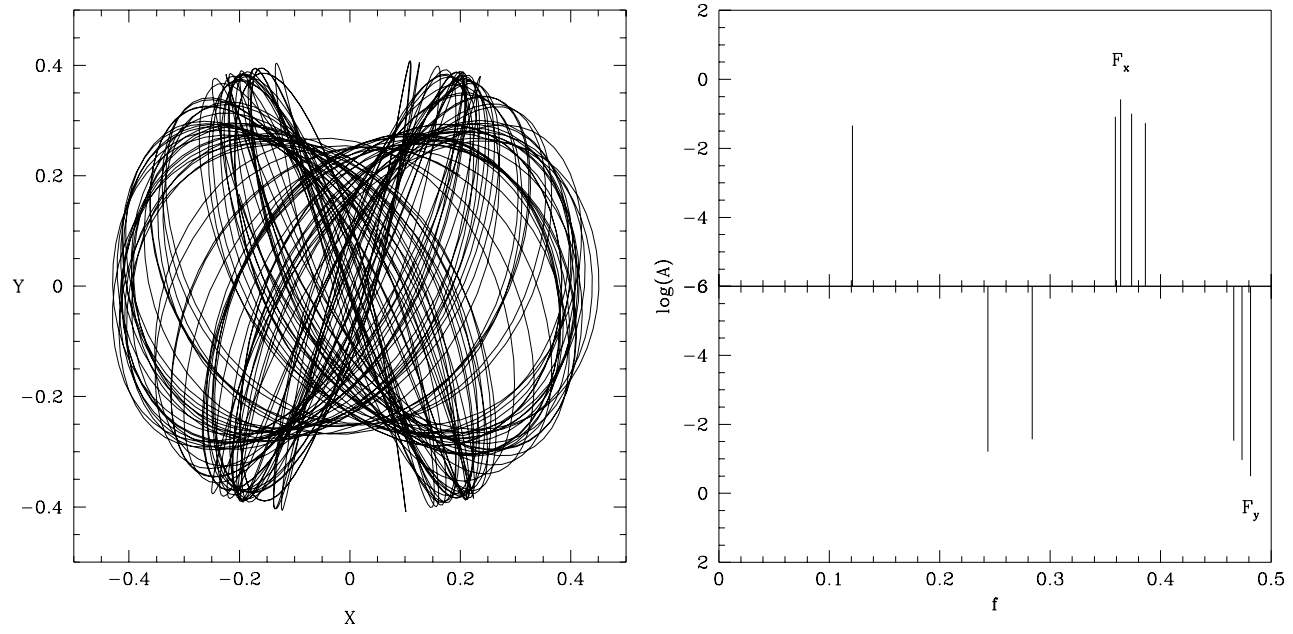

Figure 11. Same as Fig. 4, but for an irregular orbit, obtained by launching a particle in the potential $\Phi_{\rm B}$, with $v_0^2 = 1$, $R_{\rm c} = 0.14$, $q = 0.9$, $R_{\rm e} = 1.5$, and with initial conditions $x_0 = 0.2$, $y_0 = 0.2$, $\dot{x}_0 = -0.8$, and $\dot{y}_0 = 0.55401$.

orbits without knowledge of the orientation or of the centre of the potential, and that it can even provide us with this information.

5.1 Orientation

It is clear that loop orbits are immune to any rotation of the coordinate system; for other reasons, irregular orbits also are immune. Box orbits, on the other hand, are not. To see this, let us take the long axis oscillator, with amplitude A, with the coordinate system aligned with the potential. Then, a single line will appear on the x spectrum, with amplitude A, and nothing will appear on the y spectrum. Now let the coordinate system be rotated anticlockwise by an angle β. The amplitude of the x-line will shrink to $A\cos\beta$, and a line will appear on the y spectrum at the same frequency, with amplitude $A\sin\beta$. The phases of both lines will be the same. Had we rotated by a clockwise angle β, the amplitudes and frequencies would be the same as above, but the phases would differ by an angle π. Since, in this context, β and $\beta + \pi$ are equivalent angles, we have covered any possible angle with the example above.

To recover the original alignment, we seek to learn whether the frequencies are equal, and then compute $\tan\beta$ from the amplitudes. Since this computation always yields an angle β belonging to the

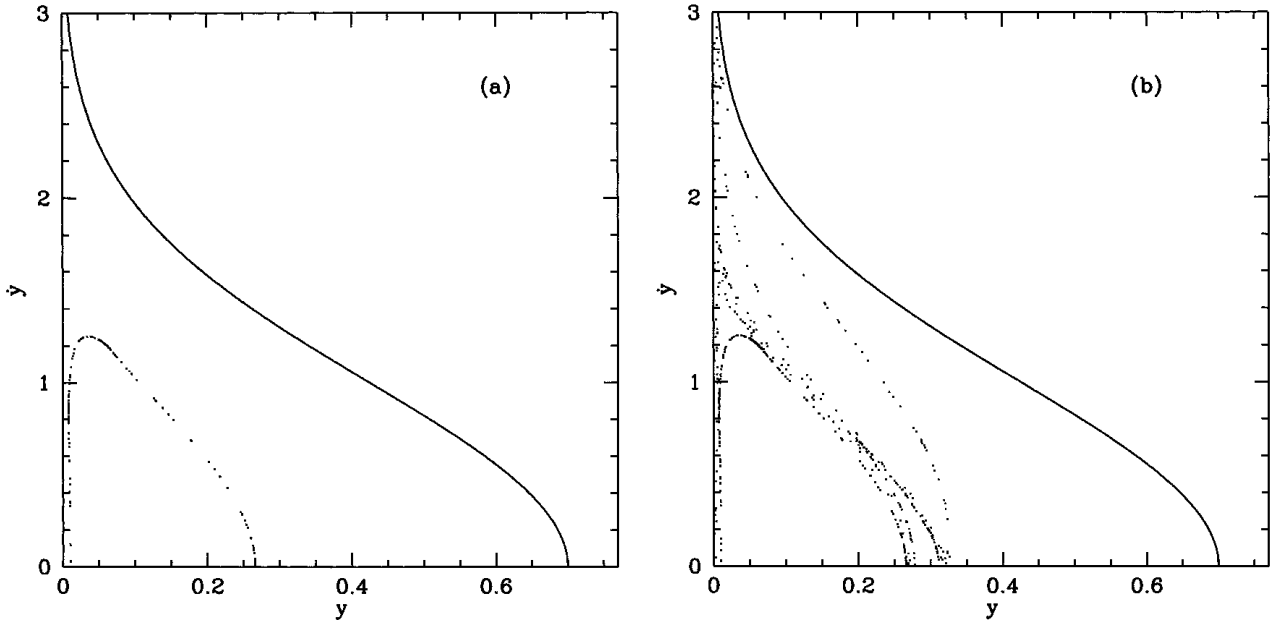

Figure 12. Left: SoS of an orbit (see text), integrated during 400 orbital periods. Right: the same, but for some 2000 orbital periods.

first quadrant, we fix the sign through the phases. This scheme is not in conflict with the recognition of a loop orbit (we recall that the presence of equal frequencies in both x and y spectra is the criterion for loop orbits) because once the coordinate system is rotated, then (i) if the orbit is a loop, nothing changes and we still have equal frequencies in both spectra; or (ii) if the orbit is a box, one of the lines vanishes.

If the spectra have more than one line, they are irrelevant to the above discussion, because, in order to make the rotation, it suffices to take the line with the maximum amplitude (whether it appears on the x or on the y spectrum) to make the rotation.

Before leaving the subject, we must work out a subtle problem. We have seen that the long axis of the potential, i.e., the prospective x axis of coordinates, bears the maximum amplitude of oscillation of a box orbit, and the above rotation is based on this assumption. However, it turns out that there are orbits in which the maximum amplitude lies on the short axis of the potential. If this is the case, we are left with a coordinate system wrongly rotated 90°. To avoid this, since the ratio F_y/F_x is always greater than one, we simply rotate the coordinate axes by 90° whenever $F_y/F_x < 1$.

5.2 Translation

A displacement of the coordinate system reflects on the spectra of an orbit only in its zero-frequency slot; the rest is left unchanged. Then the most simple way to centre the orbit is to nullify the amplitude of the zero frequency. This centring is most useful in dealing with orbits in the meridional plane of an axisymmetric potential.

One last comment about loop orbits: we might have tried to classify loop orbits based on their constant sense of motion, instead of through their spectra. If this were the case, a translation does change matters, because the constant sense of rotation is true only if measured from within the 'hole' of the orbit. We can still make a translation eliminating the frequency zero, but if the orbit has a very tiny hole, then the translation might fail to put the origin exactly inside the hole, and the sense of rotation will come out wrong. Also we would run into trouble if trying to classify a loop orbit in the meridional plane of an axisymmetric potential, for in this case one cannot say beforehand where the hole is (unless one visually inspects the orbit). It is safer to deal with the spectra.

			Number of base frequencies		
			1	2	3 or more
Number of resonances	0	0^{th} rank	axial	π–box	irregular
	1	1^{st} rank	closed m:n box closed 1:1 loop	open m:n box open 1:1 loop	
		2^{nd} rank	closed km:kn box closed k:k loop	open km:kn box open k:k loop	

Figure 13. Summary of 2D spectral orbit classification.

6 NUMERICAL ALGORITHM AND LIMITS OF THE CLASSIFICATION

Compromises that involve decisions are unavoidable when we translate the above considerations into a finite precision, numerical algorithm. Here we discuss the most important.

One decision to make is how many lines should be extracted in order to properly classify an orbit. We want to establish a limit, both to save time and to avoid tiny noisy lines. One way to do this is by limiting the number of lines, N_l; another one is to put a lower limit on the amplitude of the line, for example with respect to the strongest. We used both criteria and found that $N_l = 5$ per coordinate, and a minimum amplitude of 0.02 times that of the greatest suffices to achieve a good classification.

We also must decide whether or not two numerically computed frequencies a and b are equal. We have used the following

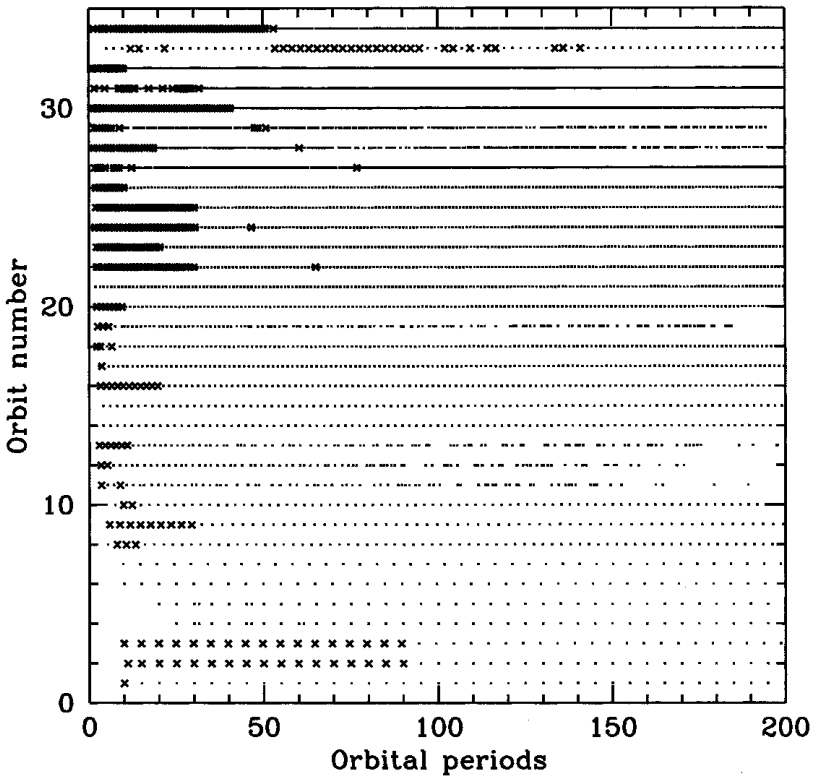

Figure 14. Orbit classification as a function of the time-span of orbit integration, for several assorted orbits. A dot indicates a successful classification, a cross marks a discrepant one.

criterion,

$$a = b \Leftrightarrow \frac{|a-b|}{|a|+1} < \varepsilon\,\Delta, \tag{18}$$

where Δ is the difference in frequency between two adjacent Fourier slots, and ε is a parameter to fix. We found that $\varepsilon \simeq 0.25$ is a good compromise.

A further problem arises from the time-span of the orbit integration. It is clear that an orbit that was integrated for too short a time may yield an incorrect Fourier spectrum. As Fig. 14 shows, we found that, with $N = 4096$ points that sample uniformly the orbit in time, the classification is reliable when the orbit has been integrated over 100 orbital periods. Fortunately, since the orbital period is, generally speaking, the inverse of F_x, the program itself can estimate the number of orbital periods over which an orbit was integrated, and can warn the user if neccesary.

As a further security, the classification is performed at least twice for each orbit: one with the original time-span, and once again with a portion of the orbit removed. If the classifications differ (which sometimes happens), a third pass is made in order to resolve the question. We found that this solves almost all uncertain cases. The rare cases left undecided are signalled by the classifier so that a more detailed examination may be performed by other means.

The most challenging numerical problem consists of determining whether two (finite precision) numbers have a low-order quotient. (Note that this is not the problem of reconstructing a numerator and a denominator from their rational quotient.) If a and b are those numbers, and m and n are (small) integers, we found that comparing a/b with m/n with the aid of equation (18) suffices in most cases. If the ratio is found to be rational, then the resonance is automatically

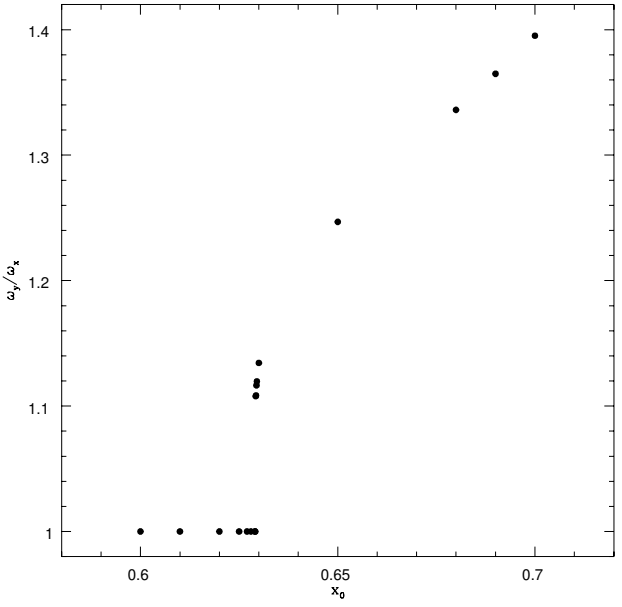

Figure 15. Ratio of dominant frequencies versus initial x coordinate in the potential Φ_L described in the text.

m:n, and the unit frequency arises at once. However, this procedure is not guaranteed to succeed, because we do not have an infinite frequency resolution. This means that resonances that do not generate orbit families will no longer be a set of measure zero, but will appear as ribbons of finite width. Decreasing ε in equation (18) will result in a narrowing of the ribbons, but the price paid will be that many real rational ratios, because of

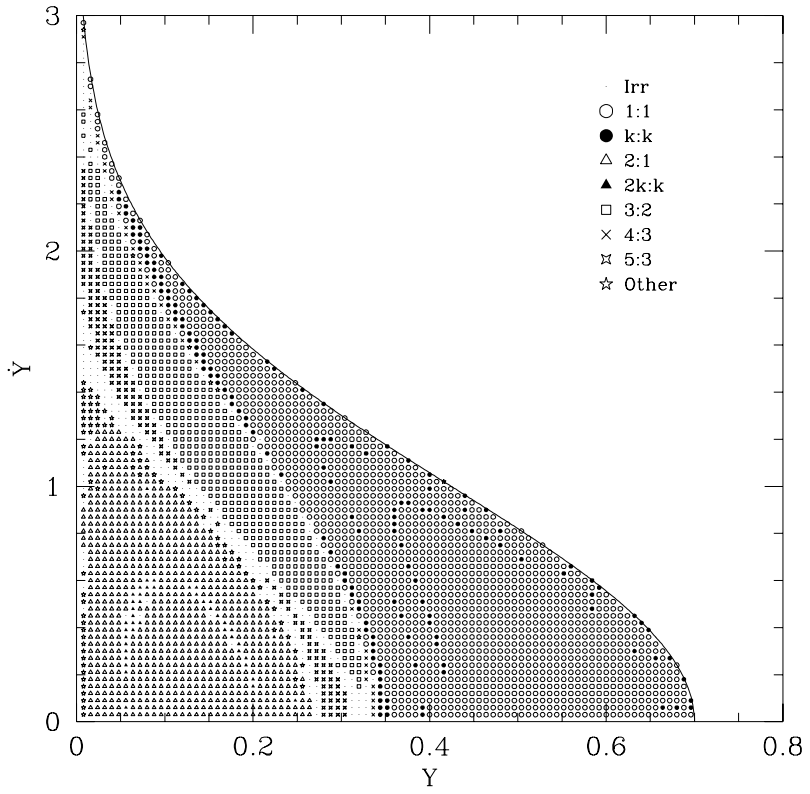

Figure 16. Orbital structure of the logarithmic potential $\Phi_{\rm L}$, with $v_0 = 1$, $R_{\rm c} = 0$ and $q = 0.7$, at $\mathcal{E} = 0$.

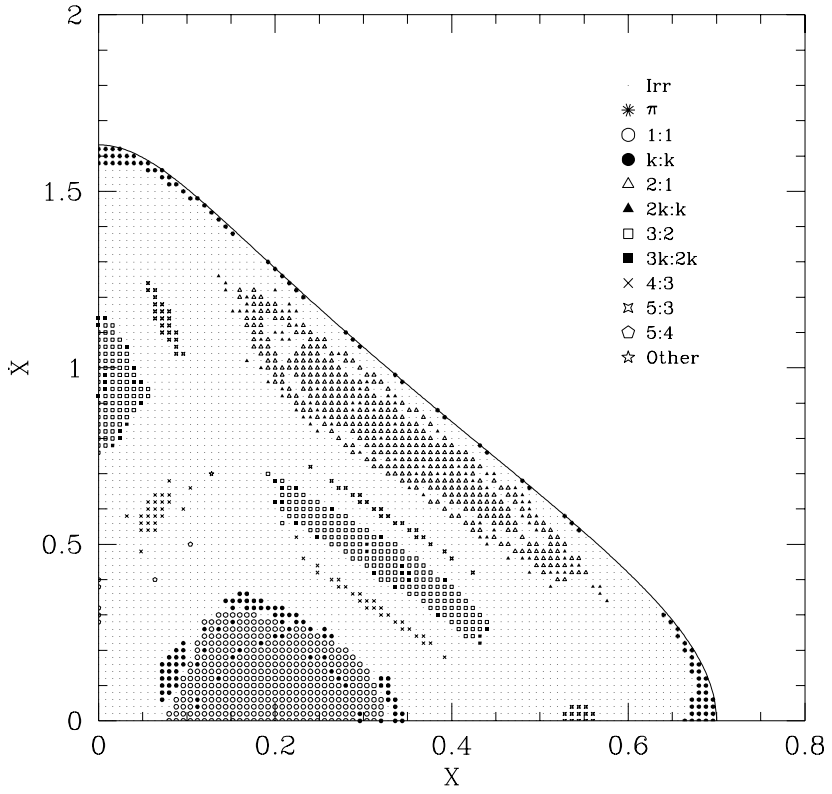

Figure 17. Orbital content of the Binney potential with $q = 0.9$, $R_{\rm c} = 0.14$, $v_0^2 = 1$, and $R_{\rm e} = 1.5$, at $\mathcal{E} = -0.6348$.

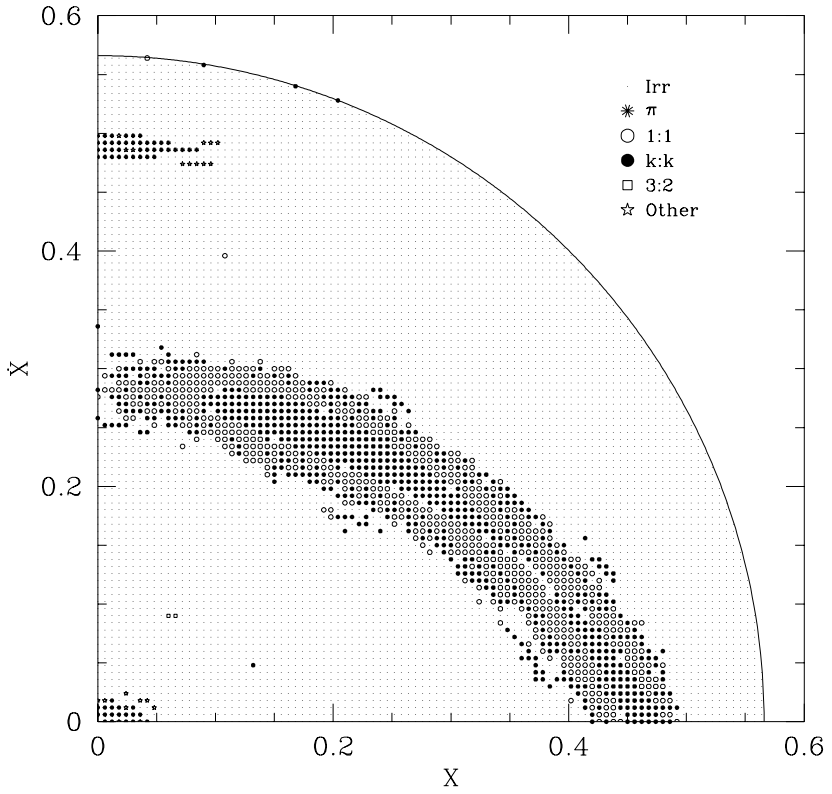

Figure 18. Orbital content of the Hénon–Heiles potential at $\mathcal{E} = 0.1602$.

inaccuracies in the determination of their frequencies, will be considered irrational, thus resulting in spurious irregular orbits. In the end, one must recognize that any numerical tool will have limits in resolution.

However, the classifier is still a powerful tool. We integrated a series of orbits in the potential Φ_L with $v_0 = 1$, $q = 0.9$, and $R_c = 0.14$, at an energy $\mathcal{E} = -0.337$ (BT87, fig. 3.8), with initial conditions $0.6 \leq x_0 \leq 0.7, \dot{x}_0 = y_0 = 0$, and $\dot{y}_0$ chosen to match the energy. Thus, this series crosses from the loop to the box region of the potential. Fig. 15 shows the ratio of dominant frequencies along this series. The classifier sharply recognizes the border between families ($\omega_y/\omega_x = 1$ for loops, > 1 for π-boxes). The first two orbits, from left to right, with ratios > 1, were classified as irregulars. The rest were classified as loops (left) or π-boxes (right).

7 RESULTS FOR 2D POTENTIALS

To test our classifier, we analysed the orbital structure of several potentials previously studied, so that we could compare the results.

Fig. 16 shows a grid of orbits that we have classified on the SoS for the logarithmic potential Φ_L, with $q = 0.7, R_c = 0, v_0^2 = 1$, that corresponds to an energy per unit mass $\mathcal{E} = 0$ (in this and other examples in this section, every symbol in the corresponding figure represents the initial point of an orbit that was integrated and classified separately). This is to be compared with fig. 1 of Miralda-Escudé & Schwarzschild (1989). As indicated by these authors, there is a complete lack of π-box orbits, and irregular orbits fill the phase space between the different classes of regular orbits. All resonances found by these authors have been found by our classifier, and, at the higher resolution of our grid, further fine detail is beginning to appear. Resonances as high as 8:5 appear close to $y \sim 0$ and $\dot{y} < 1.4$ (labelled as 'other' in the figure), narrow strips of k:k loops at the border and within the region occupied by the open 1:1 loops orbits, and also narrow strips of $2k$:k boxes within the region of 2:1 boxes. The 'irregular sea' in this SoS appears to be all connected as can be seen in Fig. 12(b), which shows an irregular orbit that loiters around almost all of the irregular bands between regular zones in this SoS. The few irregulars close to the $2k$:k orbits, when inspected, reveal themselves as 4:2 boxes, so we are here at the limit of frequency resolution of the classifier.

Fig. 17 shows a similar SoS grid, but for the Binney potential Φ_B, with $q = 0.9$, $R_c = 0.14$, $v_0^2 = 1$, and $R_e = 1.5$, at $\mathcal{E} = -0.6348$. As can be seen, there are islands of regular orbits in a sea of irregularity. Again, second-rank resonance orbits begin to show up. This figure can be compared with figs 3–27 of BT87, where some orbits were sketched. Our classifier combined with a finer grid gives us a higher resolution picture of the orbital structure of this particular SoS.

Fig. 18 shows an SoS of the Hénon–Heiles potential Φ_H, at $\mathcal{E} = 0.1602$. As is already known, at this relatively high energy most orbits are irregular. Our classifier shows also that there exists a large regular region populated mainly by 1:1 and k:k loops. There are also a few m:n box orbits with high resonances; however, these isolated orbits are most likely a consequence of the finite frequency resolution of our classifier.

Table 2 reports the percentage of orbits of each type found on the orbital grids studied in this section.

8 SPECTRAL CLASSIFICATION OF 3D ORBITS

The Fourier analysis and subsequent extraction of lines proceeds as

Table 2. Percentages of orbits in the studied potentials.

Orbit type	Φ_L	Φ_B	Φ_H
Irregular	8	68	82
π-box	0	2	1
1:1	41	8	8
k:k	4	2	9
2:1	22	9	0
2k:k	1	2	0
3:2	15	5	<1
3k:2k	0	1	0
4:3	1	1	0
5:3	4	1	0
5:4	0	<1	0
Other m:n	3	<1	<1
Total of orbits	3556	4237	7084

in 2D. The only change is in the classification itself. We assume that the x, y, and z axes correspond to the large, intermediate, and short axes of the potential, respectively. First, we determine the dominant frequencies in each coordinate, denoting them F_x, F_y, and F_z.

Now we take each couple of coordinates (x, y), (x, z), (y, z) in turn. For each case, we analyse whether the second-dominant line must bear the title of dominant, as we explained in Section 4.1. Then, again for each pair, we compute whether or not they have a rational ratio, i.e., whether a resonance is present. If this is the case, we compute the integer numerator and denominator, as described in Section 6. In each case, we search for the possibility of an improper quotient related to a higher-rank resonance, and compute numerators and denominators accordingly. Once the three pairs of coordinates have been surveyed, we are left with zero, one, two or three resonances encountered.

If we find three resonances, i.e., one resonance in each pair of coordinates, then all three dominant frequencies are multiples of a single unit frequency. We then build the relationship between the three coordinates in the form of a proportion among integers X:Y:Z, taking into account any present higher-rank resonance.

When the program finds only two resonances (which is, of course, not possible), what happens is that one of the quotients is likely to be formed by two large integers beyond the range searched by the classifier. We then reconstruct the lost resonance from the other two, ensuring that no spurious improper resonances leak in.

If we find only one resonance, say Y/X, then it means that F_z has an irrational relation with respect to F_x and F_y. In this case, we will denote the irrational ratios by putting $Z = \pi$, thus yielding the triplet X:Y:π. Of course, the same holds if Z/X or Z/Y were the case.

If no resonances are encountered, all three dominant frequencies are irrationally related. This corresponds to the 3D analogues of the 2D π-boxes. Extending our foregoing notation, we will call them 3D π-boxes.

We can now establish the number of BFs among the dominant frequencies: if there are three resonances, there is one BF; if there is one resonance, there are two BFs; and if no resonances were found at all, there are three BFs. With these figures, we go on searching for additional BFs, i.e., frequencies present in the orbit but that are not integer linear combinations of the dominant-born BFs. The procedure goes much as in the 2D case, but now we make room for up to four possible independent frequencies.

The last step is to classify the orbit from the data collected above. First, we consider the total number of BFs: if we find a total of more than three BFs, the orbit is irregular and the BFs are not really BFs. If fewer than four BFs are found, we have regular orbits that are *closed* (one BF) or *open* (three BFs). If there are two BFs, the orbit moves on a 2D manifold in configuration space. Following the usual nomenclature, we will call these orbits *thin*.

		Number of base frequencies			
		1	2	3	4 or more
Number of resonances	0	axial	2–D π–box	3–D π–box	irregular
	1	closed 0:m:n box[a] closed 0:1:1 loop	thin π:m:n box thin π:1:1 tube	open π:m:n box open π:1:1 tube	
	3	closed l:m:n box closed l:1:1 tube	thin l:m:n box thin l:1:1 tube	open l:m:n box open l:1:1 tube	

[a] The order of the indices does not change the classification name, just the spatial orientation.

Figure 19. Summary of 3D spectral orbit classification. [a] The order of the indices does not change the classification name, just the spatial orientation.

Secondly, we consider the number of resonances. When there are no resonances, then the orbit is an axial orbit (one BF), a 2D π-box (two BFs) or a 3D π-box (three BFs). If there is one resonance, we have a 2D closed box (one BF), a thin box (two BFs) or an open box (three BFs); except when the resonance is 1:1, in which case we have closed loops, thin tubes and open tubes, respectively. Finally, the orbits whose dominant frequencies are all in resonance, and with three or two BFs, are boxes or tubes the parents of which are the corresponding 3D closed resonant orbits with one BF.

The particular case of a 1:1 resonance between two coordinates gives rise to the familiar z-tubes, if x and y are resonant, or x-tubes, if y and z are resonant. (In rare cases, as in a fully harmonic 3D potential, we may also have y-tubes or two or three k:k relations). Although we were not able to find any distinguishing feature between so-called outer x-tubes and inner x-tubes (de Zeeuw 1985) spectra, we can fill this gap with a simple routine, computing from the coordinates of the orbit whether it is concave (outer x-tube) or convex (inner x-tube) in the x direction. However, it turned out that this routine could not handle all the cases properly, as sometimes the curvature was subtle. A better routine remains to be created.

Fig. 19 summarizes the classification. The columns arrange orbits according to whether they are closed, thin or open. The rows correspond to family sequences in which the parent is 1-, 2- or 3D. We have added a 0 in the notation to indicate an absent coordinate.

Our notation, however, has a difficulty which did not appear before. In 2D, a common factor was used to signal a second-rank resonance; in 3D, such a factor may be spurious. Let us suppose, for example, that we have a 9:4:2 resonance. Have we the resonances 9:4, 9:2 and the second-rank 4:2, or is the third pair a 2:1 resonance converted to 4:2 by the first two pairs? Clearly, we must extend our notation to break the ambiguity. We use a prime to signal a pair of numbers with a common factor which do not represent a real second-rank resonance. In the example above, a 4:2 resonance would yield the standard 9:4:2, whereas a 2:1 resonance would be written as 9:4′:2′. When there are two pairs of numbers having a common factor, as in 9:6:4, it may be that there are no second-rank resonances, in which case we put 9′:6′:4′ to indicate the resonances 3:2, 9:4 and 3:2. If there is only one second-rank resonance, say 6:4, we write 9′:6′:4, marking individual 3:2, 9:4 and 6:4 resonances. Finally, 9:6:4 signals that both 9:6 and 6:4 are second-rank resonances.

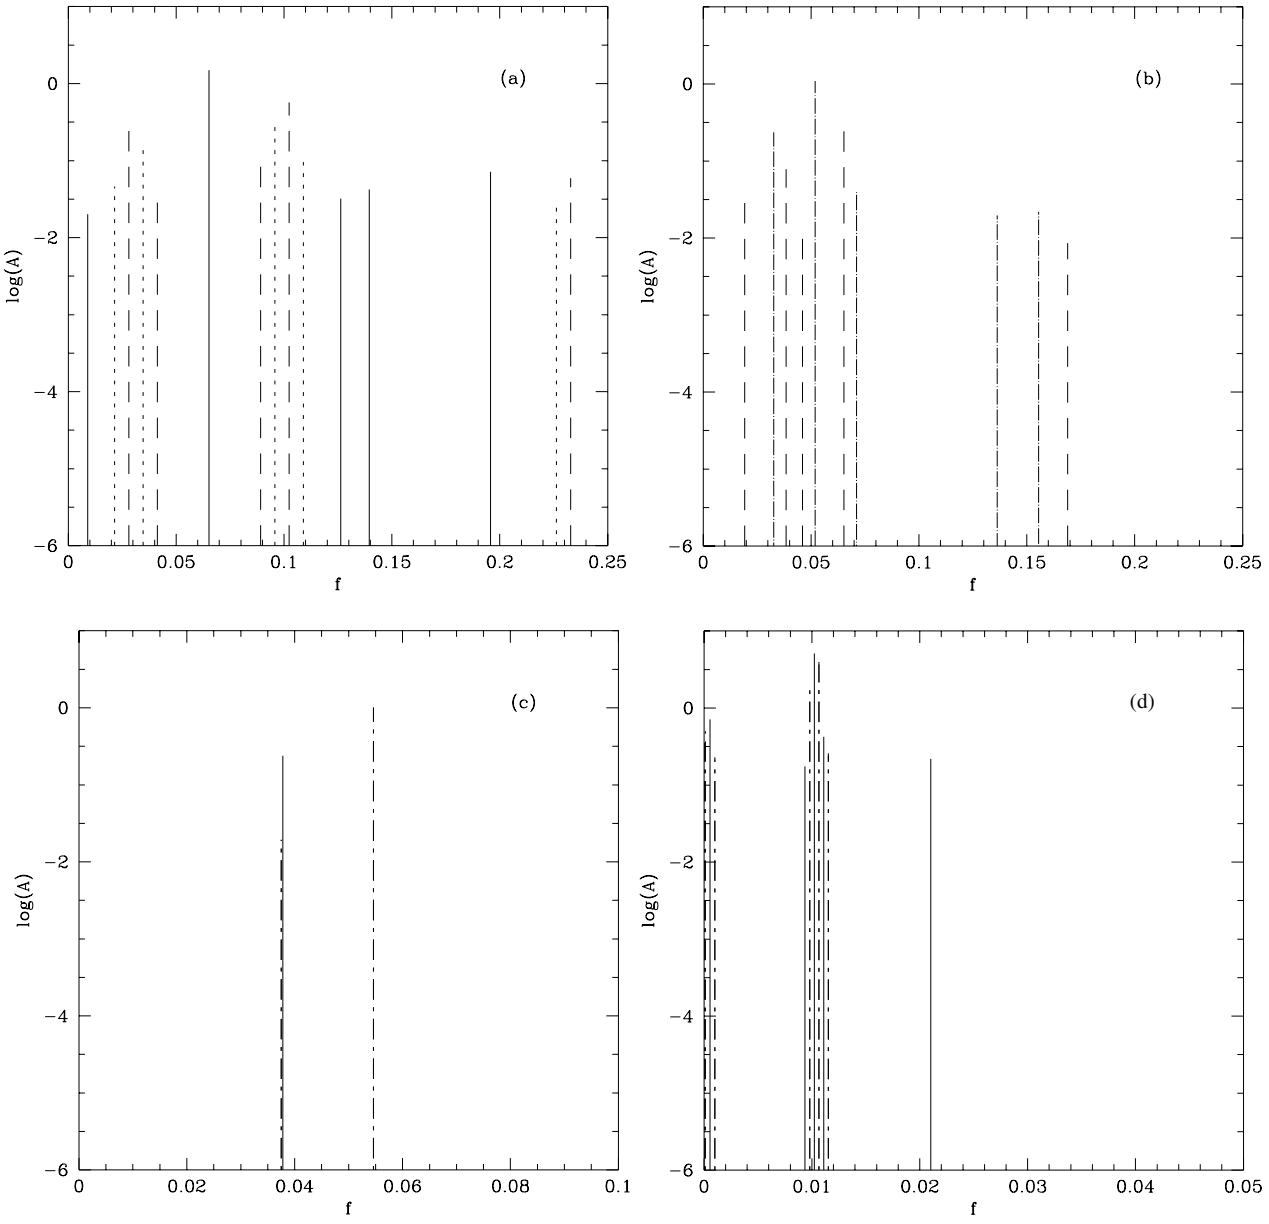

Figure 20. Spectra of orbits in the Stäckel potentials described in the text. (a) Spectra of the orbit obtained by launching a particle from $\boldsymbol{x}_0 \equiv (x_0, y_0, z_0) = (-1.3, -0.5, 0.3)$, and with initial velocity $\boldsymbol{v}_0 \equiv (v_{x0}, v_{y0}, v_{z0}) = (0.25, 0.1, 0.2)$. x-lines are solid, y-lines dotted, and z-lines dashed. The resulting classification is a π-box. (b) The same, but with $(\boldsymbol{x}_0, \boldsymbol{v}_0) = (1, 1, 0.1, -0.2901, 0.3, 0.1)$. The dashed–dotted lines belong to both x and y coordinates. It is an open 1:1:π z-tube. (c) The same, but with $(\boldsymbol{x}_0, \boldsymbol{v}_0) = (0.01, -1, 0, 0.05678, 0.01, 0.35)$. Dashed–dotted lines indicate both y and z lines. It is a thin 9:13′:13′ inner x-tube. (d) The same, but with $(\boldsymbol{x}_0, \boldsymbol{v}_0) = (1, -3.1, -5.4, 0.3, -0.1, 0.1)$. Dashed–dotted lines indicate both y and z lines. It is an open π:1:1 outer x-tube.

As in the 2D case, we perform a second classification by removing a portion of the orbit; if the classification differs, a third pass is made with another orbital portion removed.

The position of the centre of the 3D potential is obtained by means of a simple extrapolation of the method described in 2D. The orientation, on the other hand, is far from this simple. We have not yet developed the corresponding algorithm.

9 RESULTS FOR 3D POTENTIALS

As for the 2D case, we study some potentials the orbital structure of which is either simple, or has been determined independently, to test the validity of our classifier.

We first tested the classifier with orbits in a 3D harmonic oscillator, obtaining good classifications in all cases.

A non-trivial and fully regular and well-studied potential is the 'perfect ellipsoid' (de Zeeuw 1985; BT87), a member of the family of Stäckel potentials (Stäckel 1890). Although it has its simpler form in ellipsoidal coordinates, which is not very convenient for numerical integration, we have used this potential because it has the major four families of 3D orbits, and by computing the three integrals of motion, one can know a priori which type of orbit one is working with, and then confront this result with the output of our classifier.

Fig. 20 shows some examples. The orbit of Fig. 20(a) was integrated in the perfect ellipsoid potential with parameters

$\alpha = -1, \beta = -0.390\,625$, and $\gamma = -0.25$. It is a 3D π-box; lines of each coordinate are marked with different line types. Note how each pair of coordinates makes a 2D π-box spectrum on its own. Fig. 20(b) shows the spectra of an open 1:1:π tube (a z-tube), integrated in the same potential with parameters $\alpha = -2$, $\beta = -1$, and $\gamma = -0.5$; in this case, the lines belonging to the x and y coordinates have the same frequencies, and so the amplitude shown is only the greatest among x and y. Fig. 20(c) shows an example of a thin inner x-tube, integrated in the prolate limit of the perfect ellipsoid potential (de Zeeuw 1985) with parameters $\alpha = -4$ and $\beta = -1$; here the frequencies of the y- and z-lines are the same. Fig. 20(d) shows the spectra of a π:1:1 outer x-tube, integrated in the same potential of Fig. 20(a).

Schwarzschild (1993) has investigated the possibility of self-consistent models in a singular, triaxial logarithmic potential, expanded in spherical harmonics. We have used this reference because the orbital classification for 3600 orbits on six variants of this potential in two so-called 'start spaces' is given, and the precise resonance is listed for 64 stable, closed orbits, together with their initial conditions (tables 3 and 4 of Schwarzschild 1993). These orbits span a rich variety of resonances as high as 31:42:51. The classification accomplished in this work represents a *tour de force*, achieved by considering the symmetry and behaviour of the mean angular momenta and positions along each axis, as well as the corresponding extrema. The way in which the precise resonance order has been found is through a visual inspection of the plots of the individual time series for each coordinate (Schwarzschild, private communication), a daunting task indeed, when considering the high order of some resonances.

We have integrated all stable closed orbits listed in this work for 80 periods. We were pleased to see that for all but 4 of the 64 orbits, we obtained the same classification as Schwarzschild, including the resonance, the closeness, and all the second-rank resonances, which Schwarzschild calls 'higher period multiples'. Here follows an account of the discrepant cases.

(i) In table 3 of the above reference, we found that orbit number 3 of the $C = 0.7$, $T = 0.5$ model, and orbit number 3 of the $C = 0.3$, $T = 0.5$ model, were thin 1:π:2 boxes, i.e. there were two BFs, and therefore they were not closed. The classification given by Schwarzschild is 1:S:2. We note that a letter 'S' appears in his tables 3 and 4 whenever there is no motion on a coordinate, except on these two cases where, according to us, it should indicate a non-resonant coordinate; interestingly enough, an S is used to indicate precisely this on his table 1.

(ii) In his table 4, orbit number 7 of the $C = 0.7$, $T = 0.5$ model was found to be a π:42:51 box, instead of a 31:42:51 box. That is, our program found two irrational quotients instead of the relations 31:42 and 31:51, and a 14:17 resonance which turned out to be improper, becoming the 42:51 quotient. We can tune the algorithm in order to detect even these high quotients (we did it and found the right classification); this practice, however, may render the algorithm unable to find irrational quotients (see Section 6). We prefer to stick with a 'safe' algorithm, losing high resonances, which are not very important.

(iii) In table 4 of Schwarzschild (1993), orbit number 3 of the $C = 0.3$, $T = 0.98$ model was found to be a 2:3:0 box (a planar fish), instead of a reported 2:3:5. This is clearly a mere misprint in the reference, which must then read 2:3:S.

As a further step, we tried to reproduce the orbital content of the x-z start space for the six models used by Schwarzschild (fig. 4 in that reference). Fig. 21 shows the outcome. In this figure, the x loop stands for planar 0:1:1 orbits, which we are unable to classify as inner or outer long tubes. We have explicitly separated saucers (1:1:2) and fans (1:2:2) from the rest of the bananas and tubes, in order to compare our results with the work of Schwarzschild. The greatest difference turns out to be the lack of saucers in the third and fourth models, which are conspicuous in those models in Schwarzschild (1993). However, this may be due to a slight inconsistency in Schwarzschild's definition of saucers, since they are defined as 1:1:2 resonances in his table 1, but are searched for simply as 'significant L_z, asymmetric in z', conditions that orbits other than saucers can fulfil. Also, comparison between our Fig. 21 and fig. 4 of Schwarzschild shows that a finer resolution can easily be attained with our automated classification. However, there appear several suspicious orbits: inner tubes immersed in regions of outer tubes, isolated irregulars in otherwise regular zones, or box orbits embedded in a sea of tubes. We closely examined a number of these orbits. The inner tubes studied turned out to be outer tubes: their concavities were small, and so our routine was fooled. On the other hand, suspicious irregulars and boxes had very close lines in different coordinates, which should be at the same frequency if those orbits were to follow the family of their neighbours; in fact, by augmenting ε in equation (18), we recovered the correct classifications. This demonstrates how the value of ε is a compromise between a finer resolution in frequency space, and a compensating factor for the numerical inaccuracies in extracting the lines.

10 SUMMARY AND COMPARISON WITH PREVIOUS SPECTRAL METHODS

We have developed a method to automatically classify an orbit, given its coordinates as a function of time. It is based on the analysis of the Fourier spectrum of the orbit, and takes advantage of the quasi-periodic property of regular orbits. It can distinguish not only between regular and irregular orbits, but also between loop, box, and other resonant orbits up to orders comparable to those studied previously. It can also identify higher-rank resonances. An additional feature is the ability to classify orbits, even when the coordinate system is not centred, nor (for the 2D case), aligned with the potential. For orbits to be classified reliably, they should be integrated for about 100 orbital periods.

The main limitation of the present method is its resolution in frequency space, an unavoidable limitation of any finite precision computation. Its main advantage is the automatic character of the classification, which permits the reliable classification of a large number of orbits, leaving only a few doubtful cases to be examined further in detail. This advantage should greatly facilitate detailed orbital structure investigations of generic potentials. We have also introduced a consistent orbit nomenclature based on the spectral features used for the classification. This nomenclature has the virtue of including most types of orbits previously studied, in a notation that is compact and physically meaningful. Although this notation encompasses up to second-rank resonances, it can be extended in the future, should it be necessary to study higher-rank resonances.

Binney & Spergel (1982) first introduced the spectral method in galactic dynamics. In that first work they studied box and loop orbits in the 2D logarithmic potential. They recognized the line-like spectra as the basic distinguishing feature of regular orbits in frequency space, and interpreted the two BFs for these orbits in terms of a simple physical model. In a follow-up work (Binney & Spergel 1984), they showed the way in which the actions can be computed from the Fourier representation of the orbits. They came back to the logarithmic potential and were able to show several

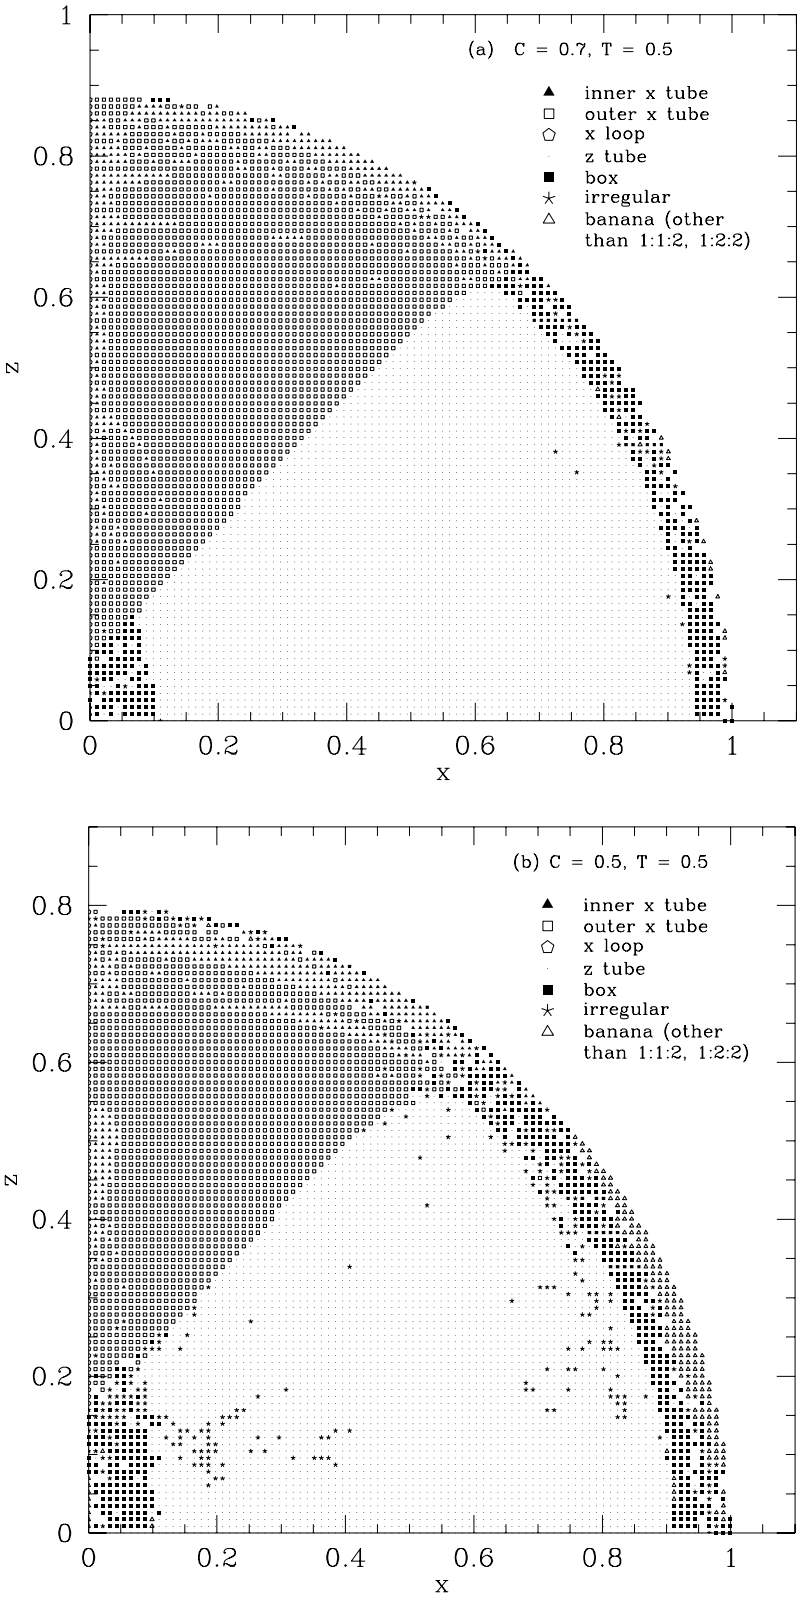

Figure 21. Orbital content of the x–z start space in the (a) first, (b) second, (c) third, (d) fourth, (e) fifth and (f) sixth scale-free logarithmic models used in Schwarzschild (1993).

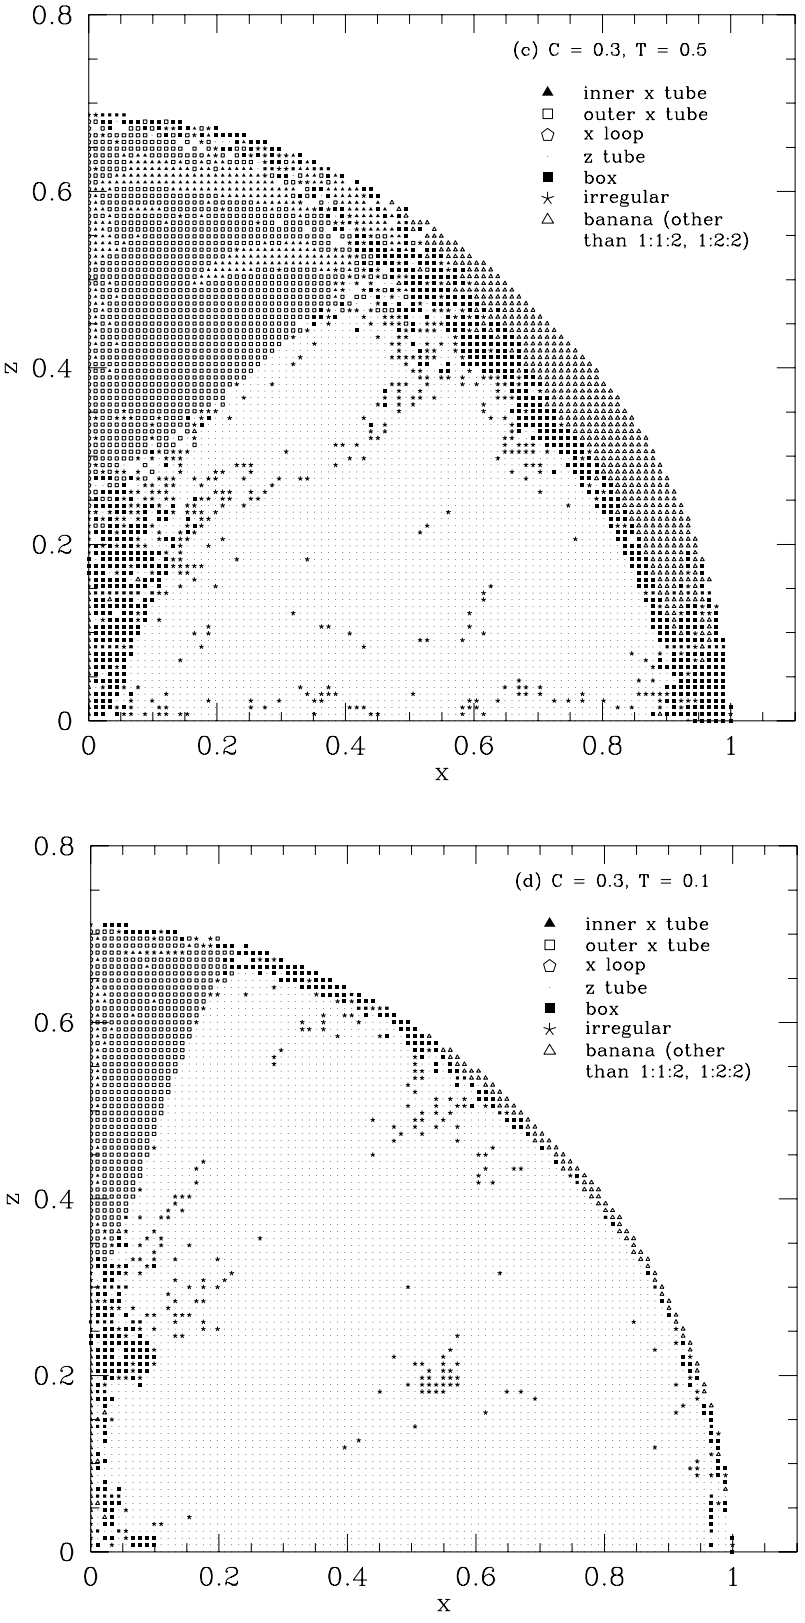

Figure 21 – *continued*

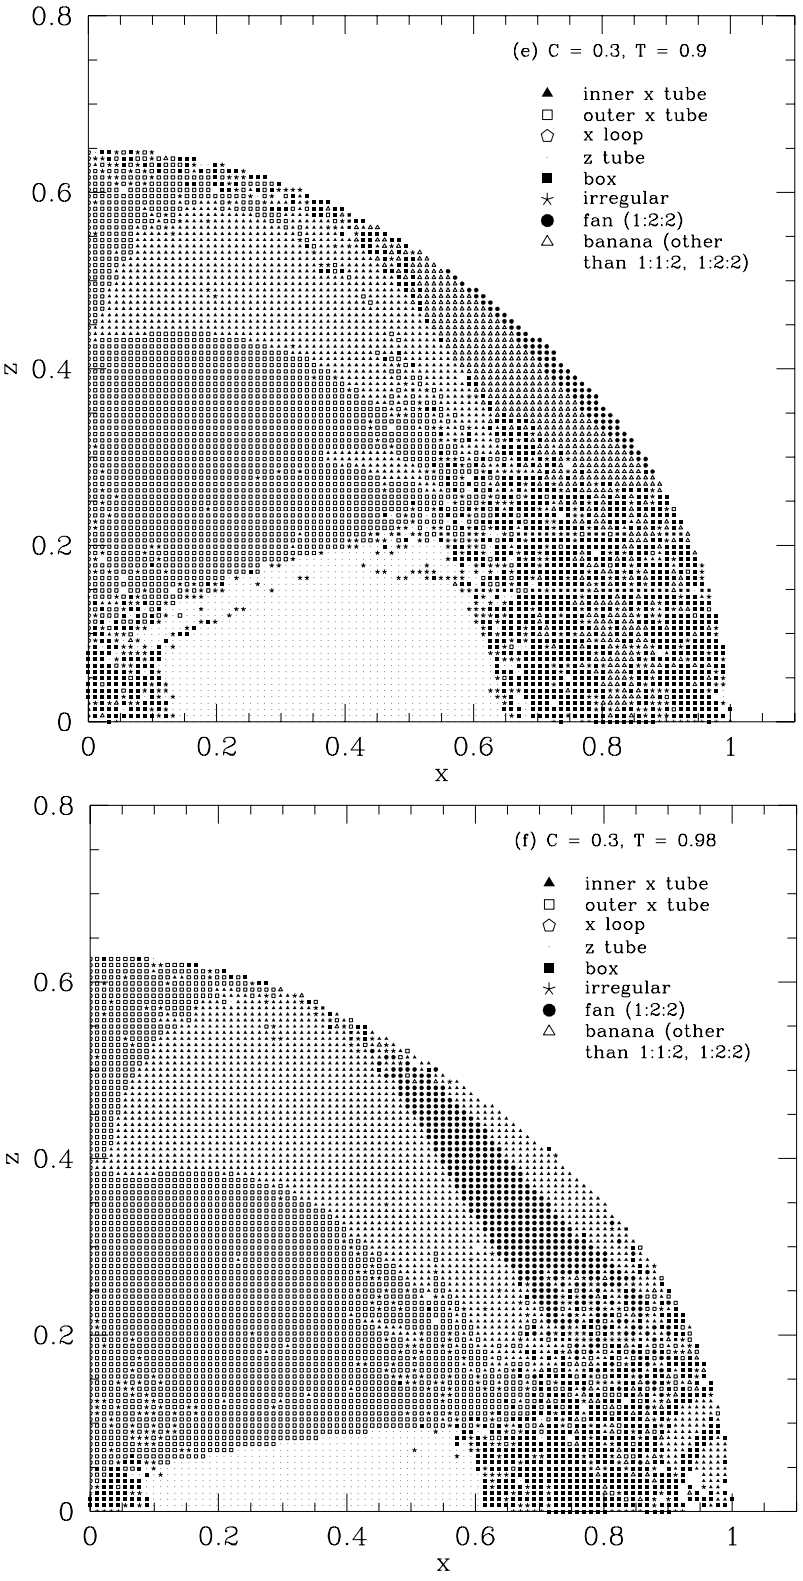

Figure 21 – *continued*

important results, e.g. the way in which regular orbit families fit in action space.

More recently, Laskar (1993) has reintroduced the spectral method within the context of stability studies of orbits in conservative dynamical systems. Laskar obtains the BFs using phase-space coordinates in an iterative numerical algorithm that results in a very accurate determination of frequencies and phases. Papaphilippou & Laskar (1996) have used this method to study the same logarithmic potential studied by Binney & Spergel. The actions of box and loop orbits are then approximated by means of the spectral representation. The main thrust of this work is the study of the orbital structure by means of a 1D map from one phase-space variable to frequency space. Resonances as high as 9:16 are identified by means of this technique.

Our approach has been different, as our main goal has been to get an automatic orbit classifier. Our procedure for extracting lines is better than the one originally used by Binney & Spergel, but does not seem to be as good as the one claimed by Laskar. Our precision, however, seems to be adequate to obtain the desired classification. The identification of the particular behaviour of line frequencies in Fourier space, which we use to accomplish the classification, was not contemplated in those earlier references.

In computing the BFs for an orbit, we are obtaining the basic frequencies that occur in action–angle theory. This being of importance itself, further work remains to be done; in particular, the computation of the action integrals for 3D orbits extending the method introduced by Binney & Spergel (1984) is of paramount importance, because they provide a natural coordinate system on which to study the phase-space distribution of regular dynamical systems. Another interesting line of work would be to investigate whether chaotic orbits can be identified by means of the continuum they produce in Fourier space. We hope to pursue these lines of investigation in the future.

The program developed to accomplish the orbit classification is available for general use upon request.

ACKNOWLEDGMENTS

We thank R. Montgomery and F. Fasso for illuminating discussions on the intrincancies of non-linear dynamics, and M. Aragón for providing us with a useful algorithm. We are also indebted to T. de Zeeuw, J. Binney, D. Merritt, and M. Schwarzschild, who read the manuscript and made valuable comments. M. Schwarzschild, in particular, made a very thorough examination that resulted in new insights into the problem. L. Aguilar acknowledges the hospitality of the Leiden Observatory, where part of this work was written. His stay there was supported in part by a Bezoerersbeurs from NWO (the Netherlands Foundation for Scientific Research), and in part by DGAPA (the National University of Mexico). D. Carpintero acknowledges Commision 38 of the I.A.U. for providing funds to travel to Mexico. This work was supported by CONACYT grant 3739-E.

REFERENCES

Arnold V. I., 1989, Mathematical Methods of Classical Mechanics. Springer-Verlag, New York
Binney J., 1982a, MNRAS, 201, 1
Binney J., 1982b, MNRAS, 201, 15
Binney J., Spergel D., 1982, ApJ, 252, 308
Binney J., Spergel D., 1984, MNRAS, 206, 159
Binney J., Tremaine S., 1987, Galactic Dynamics. Princeton Univ. Press, Princeton NJ (BT87)
Contopoulos G., Magenant P., 1985, Cel. Mech., 37, 387
de Zeeuw P. T., 1985, MNRAS, 216, 273
de Zeeuw P. T., 1994, in Muñoz-Tuñón C., Sánchez F., eds, The Formation and Evolution of Galaxies. Cambridge Univ. Press, Cambridge
Fehlberg E., 1968, NASA Technical Report TR R-287
Gerhard O. E., Binney J., 1985, MNRAS, 216, 467
Goodman J., Schwarzschild M., 1981, ApJ, 245, 1087
Heisler J., Merritt D., Schwarzschild M., 1982, ApJ, 258, 490
Hasan H., Pfenniger D., Norman C., 1993, ApJ, 409, 91
Hénon M., Heiles C., 1964, AJ, 69, 73
Jeans J. H., 1915, MNRAS, 76, 71
Kuzmin G. G., 1956, AZh, 33, 27
Laskar J., 1993, Physica D, 67, 257
Lees J. F., Schwarzschild M., 1992, ApJ, 384, 491
Levison H., Richstone D., 1987, ApJ, 314, 476
Lichtenberg A. J., Lieberman M. A., 1992, Regular and Chaotic Dynamics. Springer-Verlag, New York
Lynden-Bell D., 1962, MNRAS, 124, 95
Merritt D., Valluri M., 1996, ApJ, 471, 82
Merritt D., de Zeeuw T., 1983, ApJ, 267, L19
Merritt D., Fridman T., 1996, ApJ, 460, 136
Miralda Escudé J., Schwarzschild M., 1989, ApJ, 339, 752
Mulder W. A., Hooimeyer J. R. A., 1984, A&A, 134, 158
Papaphilippou Y., Laskar J., 1996, A&A, 307, 427
Pfenniger D., 1984, A&A, 134, 373
Pfenniger D., de Zeeuw P. T., 1989, in Merritt D., ed., Dynamics of Dense Stellar Systems. Cambridge Univ. Press, Cambridge, p. 81
Pfenniger D., Friedli D., 1991, A&A, 252, 75
Pfenniger D., Friedli D., 1993, A&A, 270, 561
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1994, Numerical Recipes in FORTRAN: The Art of Scientific Computing. Cambridge Univ. Press, Cambridge
Richstone D., 1980, ApJ, 238, 103
Richstone D., 1982, ApJ, 252, 496
Richstone D., 1984, ApJ, 281, 100
Schwarzschild M., 1979, ApJ, 232, 236
Schwarzschild M., 1993, ApJ, 409, 563
Stäckel P., 1890, Math. Ann., 35, 91
Statler T. S., 1987, ApJ, 321, 113
Tabor M., 1989, Chaos and integrability in Nonlinear Dynamics. John Wiley and Sons, Chichester
Wilkinson A., James R. A., 1982, MNRAS, 199, 171

This paper has been typeset from a $\TeX$/$\LaTeX$ file prepared by the author.