



UNIVERSIDAD NACIONAL AUTÓNOMA DE MÉXICO
PROGRAMA DE MAESTRÍA Y DOCTORADO EN CIENCIAS MATEMÁTICAS Y
DE LA ESPECIALIZACIÓN EN ESTADÍSTICA APLICADA

Numerics on the Dissipative Spin-Orbit Model in Celestial Mechanics

TESINA
QUE PARA OPTAR POR EL GRADO DE:
MAESTRO (A) EN CIENCIAS

PRESENTA:
Nicolas Gonzalez Boileau

DIRECTOR DE LA TESINA:
Renato Carlos Calleja Castillo
IIMAS

CD. MX., Diciembre, 2017



Universidad Nacional
Autónoma de México



UNAM – Dirección General de Bibliotecas
Tesis Digitales
Restricciones de uso

DERECHOS RESERVADOS ©
PROHIBIDA SU REPRODUCCIÓN TOTAL O PARCIAL

Todo el material contenido en esta tesis esta protegido por la Ley Federal del Derecho de Autor (LFDA) de los Estados Unidos Mexicanos (México).

El uso de imágenes, fragmentos de videos, y demás material que sea objeto de protección de los derechos de autor, será exclusivamente para fines educativos e informativos y deberá citar la fuente donde la obtuvo mencionando el autor o autores. Cualquier uso distinto como el lucro, reproducción, edición o modificación, será perseguido y sancionado por el respectivo titular de los Derechos de Autor.

Contents

1	Introduction	2
2	Mathematical Model	2
3	Theoretical Setup	4
4	Numerical Scheme	8
4.1	Newton Step	8
4.2	Small Divisors	10
4.3	Numerical Setting	11
4.4	FFT	11
4.5	Continuation	12
4.6	Grid Extension	14
4.7	Filtering	15
5	Numerical Results	16
5.1	Test Case	16
5.2	Filtering	17
5.3	Small Divisors	18
5.4	Quasi-Periodic Solutions	19
5.5	Error	23
5.6	Satellite Angle	24
6	Conclusion	25

1 Introduction

In this work, we will consider a problem arising in the field of Celestial Mechanics named the Spin-Orbit problem. More concretely, we will study the motion of a satellite orbiting around a big-mass body described by a non-linear second order ODE. One of the peculiarities in our model is the inclusion of dissipation in the system. It is clear, intuitively, that periodic motion occurs as the satellite revolves around the big mass body and also spins on itself. However, when considering dissipation, the dynamics will be described by stable, attracting quasi-periodic motion. This is the main goal of our work. Studying numerically how the dynamics are affected when dissipation is taken into account.

In order to compute the “attracting” dynamics, we will need to reparametrize the problem on the 2-torus and transform the original ODE into a PDE. Numerically, we will take advantage of various tools such as the FFT and a variant of the Newton scheme sometimes referred as Newton-Nash-Moser, giving us a very fast and efficient computational algorithm. We first discuss the model from the physics standpoint, describing all variables and assumptions in its derivation. Secondly, we state the theoretical results on which this numerical study is based and which serve as a guideline for the numerical scheme. Naturally, the next section describes the numerical scheme in full details along with additional features in the computations. Lastly, we present and discuss the results and state some concluding remarks.

2 Mathematical Model




Figure 1: Visual representation of Spin-Orbit model

This work consists of studying the dissipative spin-orbit problem. The Hamiltonian (conservative) spin-orbit problem has already been studied extensively in [6], for example, and will be used as a foundation for this more general case. It is a natural and obvious question to ask how are the dynamics affected when transitioning into the dissipative setting. We aim to answer this question numerically.

To this end, we will consider a non-rigid body namely, a satellite, revolving around a big-mass body on a Keplerian ellipse. The satellite will be subject to the gravitational attraction of the big-mass body located at one of the ellipse’s foci. Being a non-rigid body, the satellite will obviously deform slightly when closer to the big-mass body, where the forces acting on it are

greater. In fact, one of the most important contribution of dissipation comes from the satellite's non-rigidity in the form of tidal torque. For our model, this tidal torque will depend linearly on the relative angular velocity.

We assume some symmetry and regularity in the satellite's trajectory. Such assumptions simplify the analysis. As we mentioned above, it orbits on a Keplerian ellipse. It therefore follows that the satellite's motion lies on a fixed plane in 3D space. The symmetry in our problem is precisely defined by this ellipse's major axis (see Figure 1).

The satellite's position and motion will be determined by the angle, $x(t)$ (see Figure 1). This angle is defined by the satellite's major physical axis or in other words, its own orientation, and the ellipse's major axis corresponding to the "equatorial" line defining the problem's symmetry. Clearly, both the satellite's orientation and position are periodic variables with respect to time as it spins around the big-mass body and on itself. Putting all these physical facts and assumptions together, we can describe the satellite's rotations by the following second-order ODE,

$$\ddot{x}(t) + \eta(\dot{x}(t) - \nu) + \varepsilon f_x(x(t), t) = 0, \quad (1)$$

where,

- $t \in [0, 2\pi)$, is the periodic time variable.
- $x(t) \in [0, 2\pi)$, is the periodic angle between the satellite's orientation and the major axis of the orbit.
- $(x, t) \in \mathbb{T}^2 = \mathbb{R}^2/2\pi\mathbb{Z}^2$, the tuple angle-time lives on the 2-torus given their periodic nature.
- $\dot{x}(t) \in \mathbb{R}$, the angular velocity lies on the whole real line.
- $\ddot{x}(t) \in \mathbb{R}$, the angular acceleration lies on the whole real line.
- $\eta > 0$, is the dissipation parameter, $\eta = 0$ would be equivalent to the Hamiltonian (conservative) case.
- ν , is an external parameter related to the orbit's eccentricity.
- $0 < A \leq B < C$, are the satellite's principal moments of inertia.
- $\varepsilon \geq 0$, is the magnitude of the perturbation. This parameter is actually determined by the previous quantities,

$$\varepsilon = \frac{3B - A}{2C}$$

although we will strictly deal with ε directly.

- ω , is the satellite's Diophantine frequency of motion. In mathematical terms, a Diophantine number $\omega \in \mathcal{D}_{\kappa, \tau}$ must satisfy the following condition,

$$|\omega n_1 + n_2| \geq \frac{\kappa}{|n_1|^\tau}, \forall (n_1, n_2) \in \mathbb{Z}^2, n_1 \neq 0, \text{ for some } \kappa, \tau > 0.$$

In other words, we must impose some irrationality on ω .

In order to obtain an expression for the potential, f , we need to declare the following various quantities beforehand,

- $e \in (0, 1)$, is the eccentricity of the Keplerian orbit on which the satellite is orbiting.
- $K \geq 0$, is a physical constant depending on the internal non-rigid structure of the satellite.

- $\Omega_e > 0$ is a known function of the eccentricity, $\Omega_e = (1 + 3e^2 + \frac{3}{8}e^4)(1 - e^2)^{-\frac{9}{2}}$.
- $u_e(t)$, is the 2π -periodic solution to Kepler's equation: $t = u_e - e \sin(u_e)$.
- $\rho_e(t)$, is the orbital radius (the distance between the ellipse's focus and the satellite) given by $\rho_e(t) = 1 - e \cos(u_e(t))$ (see Figure 1).
- $f_e(t)$, is the true anomaly (angle between the major axis and the satellite's orbital radius) given by $f_e(t) = 2 \arctan \left(\sqrt{\frac{1+e}{1-e}} \tan \left(\frac{u_e(t)}{2} \right) \right)$ (see Figure 1).
- $f(x, t) \in \mathbb{R}$, is the potential given as a real analytic function. It is defined as,

$$f(x, t) = \frac{-1}{2\rho_e(t)^3} \cos(2x(t) - 2f_e(t)).$$

observation: both $\rho_e(t)$ and $f_e(t)$ are naturally 2π -periodic functions.

3 Theoretical Setup

Our numerical study is based on various theoretical results proved in [1]. We shall solely state these results to put some context when describing the numerical scheme. Before going into any further details, we begin by considering the simplified Hamiltonian problem. When removing the dissipation parameter, i.e. $\eta = 0$, the equation (1) becomes,

$$\ddot{x}(t) + \varepsilon f_x(x(t), t) = 0,$$

or written in system form,

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} y \\ -\varepsilon f_x(x, t) \end{bmatrix},$$

giving us the following Hamiltonian equations,

$$\dot{y}(t) = -\partial_x H_\varepsilon, \quad \dot{x}(t) = \partial_y H_\varepsilon,$$

for the nearly integrable Hamiltonian,

$$H_\varepsilon(y, x, t) = h(y) - \varepsilon f_x(x(t), t) = \frac{y^2}{2} - \varepsilon f_x(x(t), t). \quad (2)$$

When setting $\varepsilon = 0$, we can trivially integrate this Hamiltonian system giving us,

$$\begin{aligned} y(t) &= y_0, \\ x(t) &= x_0 + y_0 t. \end{aligned}$$

Physically, such a solution means the satellite's angle evolves linearly with constant angular velocity throughout time. We therefore see that in phase space, $\mathbb{T}^2 \times \mathbb{R}$, the plane, $\mathbb{T}^2 \times \{y_0\}$ is invariant. In fact, this plane is actually an invariant torus since both $x(t)$ and t are periodic variables and therefore the extremities of the plane are connected. We also have that y_0 coincides with the frequency of motion, $\omega = y_0$. Considering now the perturbed Hamiltonian system, i.e. $\varepsilon \neq 0$, we look for other more interesting invariant surfaces with the same frequency of motion, ω . We observe that the unperturbed term of (2) is non-degenerate since

$$h''(y) = 2 \neq 0 \quad \forall y.$$

Standard KAM theory has already shown that, under the non-degeneracy condition and the Diophantine property above, if the ε parameter is small enough, there exists invariant tori on which the motion is quasi-periodic with frequency ω of the form,

$$x(t) = \omega t + u(\omega t, t). \quad (3)$$

Remark. *The linear part of these quasi-periodic solutions, ωt , correspond to the unperturbed Hamiltonian case since $y_0 = \omega$.*

The reader is referred to [6] for further details. By frequency of motion ω , we mean that as the solution revolves around the center of the torus, it will simultaneously complete ω revolutions around the tube of the torus or vice versa. We have therefore quickly shown that there exists many quasi-periodic solutions, when the frequency of motion is Diophantine, for the conservative spin-orbit problem. This gives us a strong lead for our numerical study since we want to know what happens to these quasi-periodic solutions when the dissipation is turned on, i.e. $\eta \neq 0$.

Basically, we will “prove” numerically that such solutions will actually persist with dissipation and moreover, will smoothly bifurcate into (local) quasi-periodic attractors as ε goes away from zero. The fact that we will deal with attractors is key since the numerical scheme we will later develop is iterative. We should therefore be able to converge to these attractors numerically when our initial approximation is close enough. Before getting into more technical details, we explicitly define the reparametrization we will use for our problem (1) on the 2-torus,

$$\theta = (\theta_1, \theta_2) = (\omega t, t) \in \mathbb{T}_\omega^2 := \mathbb{R}^2 / (2\pi\omega\mathbb{Z} \times 2\pi\mathbb{Z}).$$

Remark. *This 2-torus is not to be confused with the 2-torus for (x, t) even though it is just a rescaling. From this point on, we will mostly deal with the rescaled torus, \mathbb{T}_ω^2 , we have just defined.*

Our focus will now shift to computing the function $u : \mathbb{T}_\omega^2 \rightarrow \mathbb{R}$ since it is the quasi-periodic component of the solutions we are interested in. Of course, retrieving $x(t)$ afterwards is a trivial matter. We therefore need to rewrite (1) in terms of the variable $u(\theta)$. Firstly, the time derivative for $x(t)$ translates to,

$$\begin{aligned} \dot{x}(t) &= \frac{d}{dt}(\omega t + u(\theta_1, \theta_2)), \\ &= \omega + \omega \partial_{\theta_1} u + \partial_{\theta_2} u. \end{aligned}$$

The time derivative for x therefore corresponds to a directional derivative on \mathbb{T}^2 for u . We define this linear operator by,

$$\begin{aligned} \partial_\omega &= \omega \partial_{\theta_1} + \partial_{\theta_2}, \\ \implies \dot{x}(t) &= \omega + \partial_\omega u(\theta). \end{aligned}$$

The second time derivative $\ddot{x}(t)$ now corresponds to $\ddot{x}(t) = \partial_\omega^2 u(\theta)$. Plugging in both values in the original equation (1) gives us,

$$\begin{aligned} \partial_\omega^2 u + \eta(\omega + \partial_\omega u - \nu) + \varepsilon f_x(\theta_1 + u(\theta_1, \theta_2), \theta_2) &= 0, \\ \partial_\omega^2 u + \eta \partial_\omega u + \varepsilon f_x(\theta_1 + u(\theta_1, \theta_2), \theta_2) + \eta(\omega - \nu) &= 0. \end{aligned} \tag{4}$$

Interestingly, we have converted an ODE to a PDE problem, yet this will facilitate matters. We now state the theorem, presented in [1], confirming the existence of quasi-periodic solutions in the dissipative case.

Theorem 1. *Fix $0 < \kappa < 1 \leq \tau$ and $\eta_0 > 0$. Then $\exists 0 < \varepsilon_0 < 1$ such that for any $\varepsilon \in [0, \varepsilon_0]$, any $\eta \in [-\eta_0, \eta_0]$ and any $\omega \in \mathcal{D}_{\kappa, \tau}$, \exists a unique function $u = u_\varepsilon(\theta; \eta, \omega)$, real-analytic in $\theta \in \mathbb{T}_\omega^2$ and ε , C^∞ in η , with the following property,*

$$\langle u \rangle := \frac{1}{4\pi^2} \int_{\mathbb{T}_\omega^2} u(\theta) d\theta = 0$$

such that $x(t)$ defined as in (3) solves (1) with

$$\nu = \omega(1 + \langle (u_{\theta_1})^2 \rangle). \tag{5}$$

□

Remark. *The zero-average property of the function $u(\theta)$ is only required for uniqueness.*

This result is in fact more general than what is needed for our purposes. The reason being that the parameter η can take on negative values, which physically does not make much sense, at least not for the spin-orbit problem as energy would be added to the system. Nonetheless, we are guaranteed that, theoretically, quasi-periodic solutions of (1) exist and must satisfy (4), given the appropriate values of η and ε respectively. As we had previously claimed, invariant tori from the Hamiltonian case have persisted with dissipation.

Focusing specifically on the dissipative spin-orbit problem, we actually have an explicit expression for the parameter $\eta = K\Omega_e$. Rewriting (4) in terms of the eccentricity variables we obtain,

$$\begin{aligned} \partial_\omega^2 u + K\Omega_e \partial_\omega u + \varepsilon f_x(\theta_1 + u(\theta_1, \theta_2), \theta_2) + K\Omega_e(\omega - \nu) &= 0, \\ \partial_\omega^2 u + K\Omega_e \partial_\omega u + \varepsilon f_x(\theta_1 + u(\theta_1, \theta_2), \theta_2) + K\Omega_e \omega - K\Omega_e \nu &= 0. \end{aligned} \quad (6)$$

The dissipation now depends on the parameter K . We restate Theorem 1 but in the specific context of the spin-orbit problem. For further details, the reader is once again referred to [1],

Theorem 2. *Fix $\kappa, r \in (0, 1)$ and $\tau \geq 1$. $\exists \varepsilon^* > 0$ such that for any $\varepsilon \in [0, \varepsilon^*]$, any $K = \frac{\eta}{\Omega_e} \in [0, 1]$ and any $\omega \in \mathcal{D}_{\kappa, \tau} \cap [1 + r, \frac{1}{r}]$, \exists unique functions $\nu = \nu_\varepsilon(K, \omega)$ and $u = u_\varepsilon(\theta; K, \omega)$, real-analytic in $\theta \in \mathbb{T}_\omega^2$ and ε , C^∞ in K , with the following property, $\langle u \rangle = 0$, which solve (6). \square*

Remark. *i) The condition for ν in (5) is implicitly satisfied by the function ν_ε .*

ii) the notation $u_\varepsilon(\theta; K, \omega)$ means there exists a family of solutions u parametrized by the parameter K , ω as well as ε . The subscript ε emphasizes the “bifurcation” parameter.

We re-emphasize that Theorem 2 is simply a restatement of Theorem 1 in terms of the dissipative spin-orbit problem. Obviously, the dissipation parameter, K , is now restricted to non-negative values but must also be smaller or equal to 1. Additionally, the range for the Diophantine number, ω , has now changed, restricted to values greater than 1. This theorem is another theoretical confirmation that quasi-periodic solutions, and therefore invariant tori, exist for (1) for the right choices of parameters K, ω . Lastly, unlike in the previous theorem, we no longer know the upper bound for ε explicitly. This will also be of great interest for our work. We will numerically try to approach the value ε^* , giving us an approximate idea of its magnitude and when the breakdown of the invariant tori occurs.

Up to now, we have seen two theorems discussing the existence of quasi-periodic solutions. The next result will augment that conclusion and treat the (local) attractors. From this point on, even though the equation specific to the spin-orbit is (6), we will strictly deal with equation (4), since there are less parameters and therefore simplifies notation a little bit. This is basically the only real difference between the two equations. For convenience, we define some additional operators.

- $D_\eta : z \in C^1(\mathbb{T}_\omega^2) \rightarrow D_\eta z := \partial_\omega z + \eta z$, observation: $\langle D_\eta z \rangle = \eta \langle z \rangle$, (it follows trivially by considering the Fourier coefficients of z)
- $\Delta_\eta : z \in C^2(\mathbb{T}_\omega^2) \rightarrow \Delta_\eta z := D_\eta \partial_\omega z = \partial_\omega D_\eta z$
- $\mathcal{F}_\eta : (z, \zeta) \in C^2(\mathbb{T}_\omega^2) \times \mathbb{R} \rightarrow \mathcal{F}_\eta(z, \zeta) := \Delta_\eta z + \varepsilon f_x(\theta_1 + z, \theta_2) + \zeta = \Delta_\eta z + g_x(\theta_1 + z, \theta_2) + \zeta$

The PDE in (4) is equivalent to applying the operator \mathcal{F}_η . The actual problem we will therefore try to solve numerically is the following, i.e. look for quasi-periodic solutions satisfying

$$\begin{cases} \mathcal{F}_\eta(u, \gamma) = 0, & \gamma = \eta(\omega - \nu), \\ \langle u \rangle = 0, & 1 + \partial_{\theta_1} u \neq 0 \quad \forall \theta \in \mathbb{T}_\omega^2. \end{cases} \quad (7)$$

Remark. *i) The non-vanishing property $1 + \partial_{\theta_1} u \neq 0$ will appear in the numerical scheme. Theoretically, it ensures the map below is a diffeomorphism.*

$$M : \mathbb{T}_\omega^2 \rightarrow \mathbb{T}^2 := M(\theta) = (\theta_1 + u(\theta), \theta_2) = (x(t), t)$$

ii) The condition (5) in Theorem 1 implies that,

$$\eta\nu = \eta\omega(1 + \langle (u_{\theta_1})^2 \rangle) \implies \eta\omega \langle (u_{\theta_1})^2 \rangle + \gamma = 0.$$

The attractor is therefore described by the tuple (u, γ) and it also nullifies the operator \mathcal{F}_η . We shall compute it through some fixed-point iteration scheme. Before going on to explain how we will proceed numerically, we state the theorem not only guaranteeing us existence of quasi-periodic solutions, but quasi-periodic attractors for the spin-orbit problem reinterpreted in (7). We require some norm and functions spaces definitions first.

- We will use the following norm on the space of all continuous functions on the scaled 2-torus, $z : \mathbb{T}_\omega^2 \rightarrow \mathbb{R} \in C^1$

$$\|z\|_\xi := \sum_{n \in \mathbb{Z}^2} |\hat{z}_n| e^{|n|\xi}, \quad |n| = |n_1| + |n_2|,$$

where $\xi \in \mathbb{R}$, \hat{z}_n are the Fourier coefficients of the function z and e is Napier's constant.

- \mathcal{H}^ξ is the Banach space of all continuous functions on the scaled 2-torus with finite norm, i.e. $z : \mathbb{T}_\omega^2 \rightarrow \mathbb{R} \in C^1$ s.t. $\|z\|_\xi < \infty$.
- \mathcal{H}_0^ξ is the Banach space of all continuous functions on the scaled 2-torus with finite norm and zero average, i.e. $z : \mathbb{T}_\omega^2 \rightarrow \mathbb{R} \in C^1$ s.t. $\|z\|_\xi < \infty$ and $\hat{z}_0 = 0$.

Theorem 3. *Let $0 < \xi_* < \xi < \bar{\xi} \leq 1$, let $0 < \kappa < 1 \leq \tau$, let $I_0 = [-\eta_0, \eta_0]$ for some $\eta_0 > 0$, let $\omega \in \mathcal{D}_{\kappa, \tau}$ and $(x, t) \in \mathbb{T}^2 \rightarrow g(x, t) \in \mathcal{H}^{\bar{\xi}}$, let $M > 0$ be such that, $\|g_{xxx}\|_{\bar{\xi}} \leq M$, let also $0 < \lambda < \bar{\xi} - \xi$, $0 < \alpha < 1$ and $0 < \sigma < 1$. Then there exists $k > 1$ such that the following holds: Assume $\exists v \in C^\infty(\mathbb{T}_\omega^2 \times I_0)$ and $\beta \in C^\infty(I_0)$ (the initial approximation to the quasi-periodic attractor) satisfying the following hypotheses:*

H1) for each $\eta \in I_0$, $v(\theta; \eta) \in \mathcal{H}_0^\xi$,

H2) $\|v_{\theta_1}\|_\xi \leq \sigma\lambda$ for any $\eta \in I_0$,

H3) $V := 1 + v_{\theta_1}$, $W := V^2$, $\rho := \eta \frac{\langle W^{-1} D_\eta^{-1} v_{\theta_1} \rangle}{\langle W^{-1} \rangle}$ and assume that for any $\eta \in I_0$, $|\rho| \leq \sigma\alpha$,

H4) assume that for any $\eta \in I_0$, $k\|\mathcal{F}_\eta(v, \beta)\|_\xi \leq 1$.

Under the above hypotheses, \exists unique $u(\theta; \eta) \in C^\infty(\mathbb{T}_\omega^2 \times I_0)$ and $\gamma(\eta) \in C^\infty(I_0)$ such that (7) holds. \square

Once again, all technical details and proofs of the theorems can be found in [1]. Basically, if our initial approximate solution (v, β) is close enough, i.e. satisfies all the conditions and inequalities of Theorem 3, then we know there exists a quasi-periodic local attractor satisfying (7). The proof of this theorem is constructive meaning we converge to the actual attractor from the initial approximation or in other words, the initial guess is “attracted” to the quasi-periodic solution. Essentially, our numerical scheme mimics the proof presented in [1] and provides a numerical confirmation of Theorem 3. This is the content of the following section.

4 Numerical Scheme

4.1 Newton Step

Computing the attractor corresponds to solving the functional equation in (7). Our goal is to get a Newton scheme which will converge quadratically. To this end, we use Taylor expansion for the operator $\mathcal{F}_\eta(v, \beta)$ in both variables. Trivially, the operator \mathcal{F}_η applied to (u, γ) will be equal to zero since it is the quasi-periodic attractor. For any other approximation (v, β) , of course, we only have $\mathcal{F}_\eta(v, \beta) \approx 0$. The error, E , can then be naturally defined as by how much the operator \mathcal{F}_η differs from zero, i.e. its magnitude.

$$E = \|\mathcal{F}_\eta(v, \beta)\| = \|\Delta_\eta v + \varepsilon f_x(\theta_1 + v, \theta_2) + \beta\|.$$

We want to minimize the operator's magnitude as much as possible at every iteration. We therefore look for updates $(\tilde{v}, \tilde{\beta})$ to correct (v, β) and this new better approximation would ideally give us $\mathcal{F}_\eta(v + \tilde{v}, \beta + \tilde{\beta}) = 0$. For this purpose, we expand \mathcal{F}_η about (v, β) . From the definition of the operator \mathcal{F}_η , we want to have,

$$\begin{aligned} 0 &= \mathcal{F}_\eta(v + \tilde{v}, \beta + \tilde{\beta}) \\ &= \Delta_\eta v + \Delta_\eta \tilde{v} + \varepsilon f_x(\theta_1 + v + \tilde{v}, \theta_2) + \beta + \tilde{\beta}. \end{aligned}$$

More specifically, we expand the term f_x about $\theta_1 + v$,

$$\begin{aligned} &= \mathcal{F}_\eta(v, \beta) + \tilde{\beta} + \Delta_\eta \tilde{v} + \varepsilon \tilde{v} f_{xx}(\theta_1 + v, \theta_2) + \mathcal{O}(\|\tilde{v}\|^2), \\ \implies \mathcal{F}_\eta(v, \beta) + \tilde{\beta} + \Delta_\eta \tilde{v} + \varepsilon \tilde{v} f_{xx}(\theta_1 + v, \theta_2) &= \mathcal{O}(\|\tilde{v}\|^2). \end{aligned}$$

Now multiplying the equation by $V = 1 + \partial_{\theta_1} v$, gives us

$$V\mathcal{F}_\eta(v, \beta) + V\tilde{\beta} + V\Delta_\eta \tilde{v} + \varepsilon \tilde{v} f_{xx}(\theta_1 + v, \theta_2)V = \mathcal{O}(\|\tilde{v}\|^2). \quad (8)$$

The reason for multiplying by V will become clear after the following. We differentiate the operator \mathcal{F}_η with respect to θ_1 and observe that exactly the same terms as in (8) appear,

$$\begin{aligned} \partial_{\theta_1} \mathcal{F}_\eta(v, \beta) &= \partial_{\theta_1} (\Delta_\eta v + \varepsilon f_x(\theta_1 + v, \theta_2) + \beta) \\ &= \Delta_\eta (\partial_{\theta_1} v) + \varepsilon f_{xx}(\theta_1 + v, \theta_2)(1 + \partial_{\theta_1} v) \\ &= \Delta_\eta (1 + \partial_{\theta_1} v) + \varepsilon f_{xx}(\theta_1 + v, \theta_2)(1 + \partial_{\theta_1} v) \\ &= \Delta_\eta V + \varepsilon f_{xx}(\theta_1 + v, \theta_2)V, \\ \implies \tilde{v} \partial_{\theta_1} \mathcal{F}_\eta &= \tilde{v} \Delta_\eta V + \varepsilon \tilde{v} f_{xx}(\theta_1 + v, \theta_2)V. \end{aligned}$$

where the last line is simply multiplied by \tilde{v} . Terms cancel out when subtracting the last equation from equation (8), therefore resulting in,

$$V\Delta_\eta \tilde{v} - \tilde{v} \Delta_\eta V = -(V\mathcal{F}_\eta(v, \beta) + V\tilde{\beta}) + \mathcal{O}(\|\tilde{v}\|^2) - \tilde{v} \partial_{\theta_1} \mathcal{F}_\eta. \quad (9)$$

The magnitude of $\partial_{\theta_1} \mathcal{F}_\eta$ is roughly equal to the magnitude of \mathcal{F}_η . To see this, we need to recall the definition of the operator \mathcal{F}_η in (7). We know it is, in part, determined by a scalar and the potential function, f_x , which by definition are both analytic. The other term is the second directional derivative of our approximation, v , of the quasi-periodic solution which we know to be smooth by Theorem 3. Since all terms of \mathcal{F}_η are analytic or smooth, its magnitude is not greatly altered or affected by differential operators, in this case ∂_{θ_1} . With this reasoning in mind, we can conclude that,

$$\begin{aligned} \mathcal{O}(\|\partial_{\theta_1} \mathcal{F}_\eta(v, \beta)\|) &\approx \mathcal{O}(\|\mathcal{F}_\eta(v, \beta)\|), \\ \implies \mathcal{O}(\|\partial_{\theta_1} \mathcal{F}_\eta(v, \beta)\|) &\approx \mathcal{O}(E). \end{aligned}$$

Additionally, we will also assume,

$$\mathcal{O}(E) \approx \mathcal{O}(\|\tilde{v}\|).$$

This assumption is not so far-fetched since, once again, all terms of \mathcal{F}_η are analytic or smooth and therefore the size of the update of the approximate solution should not be too far off from the size of the error. Putting these facts together in (9) gives us a numerical iterative scheme with quadratic convergence, which by convention is referred to as a Newton scheme,

$$V\Delta_\eta\tilde{v} - \tilde{v}\Delta_\eta V = -(V\mathcal{F}_\eta(v, \beta) + V\tilde{\beta}) + \mathcal{O}(E^2).$$

From our initial tuple (v, β) , we have to solve for the update $(\tilde{v}, \tilde{\beta})$ from the equation,

$$V\Delta_\eta\tilde{v} - \tilde{v}\Delta_\eta V = -(V\mathcal{F}_\eta(v, \beta) + V\tilde{\beta}). \quad (10)$$

We then update our approximate solution $(v, \beta) = (v + \tilde{v}, \beta + \tilde{\beta})$ and repeat the process until convergence is reached. We now derive how to solve equation (10) for both unknowns $(\tilde{v}, \tilde{\beta})$. The left hand side of (10) can be simplified even further by noting that,

$$V\Delta_\eta\tilde{v} - \tilde{v}\Delta_\eta V = D_\eta \left(V^2 D_0 \left(\frac{\tilde{v}}{V} \right) \right). \quad (11)$$

It is obvious that in order for the expression on the right-hand side of (11) to be well-defined, $V \neq 0 \forall \theta \in \mathbb{T}^2$. Recalling the definition of $V = 1 + \partial_{\theta_1} v$, this is precisely the non-vanishing property assumption in (7) which, as we mentioned before, is necessary for this scheme to converge and even exist. Sending the D_η operator and V^2 to the right hand side reduces to,

$$\begin{aligned} D_\eta \left(V^2 D_0 \left(\frac{\tilde{v}}{V} \right) \right) &= -(V\mathcal{F}_\eta(v, \beta) + V\tilde{\beta}), \\ D_0 \left(\frac{\tilde{v}}{V} \right) &= -\frac{(D_\eta^{-1}(V\mathcal{F}_\eta) + \tilde{\beta}D_\eta^{-1}V)}{V^2}. \end{aligned} \quad (12)$$

The average property for the D_η operator, $\langle D_\eta w \rangle = \eta \langle w \rangle$, now comes into play. In particular we know that,

$$\langle D_0(w) \rangle = 0, \quad \forall w \in C(\mathbb{T}_\omega^2).$$

Since the operator D_0 is present in equation (12), we apply the average-over- \mathbb{T}_ω^2 operator, $\langle \cdot \rangle$, making the left hand side of (12) equal to zero,

$$\begin{aligned} \left\langle D_0 \left(\frac{\tilde{v}}{V} \right) \right\rangle &= -\left\langle \frac{D_\eta^{-1}(V\mathcal{F}_\eta)}{V^2} \right\rangle - \tilde{\beta} \left\langle \frac{D_\eta^{-1}V}{V^2} \right\rangle \\ 0 &= -\left\langle \frac{D_\eta^{-1}(V\mathcal{F}_\eta)}{V^2} \right\rangle - \tilde{\beta} \left\langle \frac{D_\eta^{-1}V}{V^2} \right\rangle, \\ \implies \tilde{\beta} &= -\left\langle \frac{(D_\eta^{-1}(V\mathcal{F}_\eta))}{V^2} \right\rangle \Big/ \left\langle \frac{D_\eta^{-1}V}{V^2} \right\rangle. \end{aligned}$$

Since this expression is independent of \tilde{v} and depends exclusively on the previous approximation (v, β) , we can compute the update for β straightforwardly. Once we have calculated $\tilde{\beta}$, we obtain \tilde{v} by solving (12),

$$\tilde{v} = VD_0^{-1} \left(\frac{D_\eta^{-1}(-V\mathcal{F}_\eta) + \tilde{\beta}D_\eta^{-1}(-V)}{V^2} \right).$$

Remark. The correction for the scalar parameter γ , will implicitly be applied to the parameter ν . Recalling its definition from (7), $\gamma = \eta(\omega - \nu)$, both η and ω are fixed parameters for the dissipation and frequency of motion respectively. The only parameter that can therefore be “tweaked” is ν .

We resume the whole process in algorithm 1, highlighting the main steps for one iteration of the derived Newton scheme.

Algorithm 1. *Constant Case Spin-Orbit Newton Step*

- 1) Fix the tolerance, tol , to check for convergence in either the error, E , or the update values, $(\tilde{v}, \tilde{\beta})$. Fix some norm for the analytic-periodic functions on the torus. Fix maximum number of iterations to perform for convergence, $MaxIter$.
- 2) Define or compute an initial approximation (v, β) .
- 3) Compute the error, $E = \|\mathcal{F}_\eta(v, \beta)\|$.
- 4) If the error, E , is smaller than the allowed tolerance, then return the computed solution; *if* $E < tol$ *then* $(u, \gamma) = (v, \beta)$.
- 5) Compute $V = 1 + \partial_{\theta_1} v$.
- 6) Compute $\phi = D_\eta^{-1}(-V)$.
- 7) Compute $\psi = D_\eta^{-1}(-V\mathcal{F}_\eta(v, \beta))$.
- 8) Compute the update for β , $\tilde{\beta} = \frac{-\langle \psi \rangle}{\langle \phi \rangle}$.
- 9) Compute the update for v , $\tilde{v} = VD_0^{-1}\left(\frac{\psi + \tilde{\beta}\phi}{-V^2}\right)$ (small-divisors problem).
- 10) Update (v, β) to $(v + \tilde{v}, \beta + \tilde{\beta})$.
- 11) If the norm of the computed updates, is smaller than the allowed tolerance, then return the computed solution; *if* $\|\tilde{v}\| < tol$ *and* $\|\tilde{\beta}\| < tol$ *then* $(u, \gamma) = (v, \beta)$.
- 12) If the maximum number of iterations is not yet reached then repeat, otherwise return the computed solution; *if* $\#Iter < MaxIter$ *then goto* 3) *else* $(u, \gamma) = (v, \beta)$.

Remark. Once a value of ε and η have been fixed, an initial guess close to the actual solution needs to be set or computed for the scheme to converge. As we have mentioned above, we are looking for **local** attractors, meaning we cannot start from any random value. In doing so, our scheme may not converge. In the continuation subsection we will give greater details on the starting point.

4.2 Small Divisors

Another critical aspect for this numerical scheme to work is the fact that the frequency of motion ω must be Diophantine. We discuss why such a condition is necessary for numerical convergence and actually responsible for making the operator D_0^{-1} well-defined. To see this, we take the Fourier transform of the variable $z(\theta)$, $\theta \in \mathbb{T}_\omega^2$,

$$z = \sum_{n \in \mathbb{Z}^2} \hat{z}_n \mathbf{e}^{in \cdot \theta} \quad \text{where } n = (n_1, n_2) \text{ and } \theta = (\theta_1, \theta_2) \in \mathbb{T}_\omega^2.$$

We quickly recall the definition of the differential operator $D_0 = \partial_\omega = \omega \partial_{\theta_1} + \partial_{\theta_2}$. It is a well established fact that differentiation in real space corresponds to algebraic multiplication in Fourier space. Therefore applying the inverse operator D_0^{-1} in Fourier space will result in algebraic division, leading us to the following expression,

$$D_0^{-1}z = \sum_{n \in \mathbb{Z}^2} \frac{\hat{z}_n \mathbf{e}^{in \cdot \theta}}{i(\omega n_1 + n_2)} \quad \text{where } n = (n_1, n_2) \text{ and } \theta = (\theta_1, \theta_2) \in \mathbb{T}_\omega^2.$$

We see that the denominator contains exactly the expression present in the definition of a Diophantine number. Precisely from the Diophantine condition, we know that this value will always be different from zero. However, depending on the combination of Fourier wavenumbers n_1 and n_2 , the denominator can come arbitrarily close to zero. Numerically, this is known as the “small-divisors” problem. As we will see later when we present the numerical results, the small-divisors actually play a significant role in dictating the overall error of the Newton scheme.

4.3 Numerical Setting

All the computations will be carried out on the scaled torus, \mathbb{T}_ω^2 . We therefore require a mesh on the space $\mathbb{T}_\omega^2 = [0, 2\pi\omega] \times [0, 2\pi]$. Given the natural periodicity of the spin-orbit problem, we will recur to the FFT to compute all differential and inverse differential operators. These calculations will therefore be of order $\mathcal{O}(N_1 N_2 \log(N_1 N_2))$ operations and $\mathcal{O}(N_1 N_2)$ in storage where $N_1 \times N_2$ is the number of points on the mesh. The remaining addition, multiplication, division and average operations will also clearly be $\mathcal{O}(N_1 N_2)$ operations. In conclusion, the algorithm will be of order $\mathcal{O}(N_1 N_2 \log(N_1 N_2))$ in computations and $\mathcal{O}(N_1 N_2)$ in storage.

The code is written in the C++ language. The obvious reason is that C++ is faster than the majority of other high-level programming languages. To some extent, it can be considered low-level, explaining why it performs better. Another advantage of C++ is the possibility of recurring to OOP (object-oriented programming). This feature allows us to make the code more readable and simpler to reuse for other purposes. This duality makes it the perfect environment for our needs. We will carry out the numerics in `long double` precision giving us an improvement of about two significant digits compared to `double`. Amongst the standard built-in datatypes in the C++ compiler GCC, `long double` is the most precise.

4.4 FFT

To compute the FFT (Fast Fourier Transform) in 2D, we use the well known library called FFTW which stands for the Fastest Fourier Transform in the West. It is actually considered by many as the fastest open-source software implementation of the FFT. One peculiarity of FFTW comes from one of the basic properties of the Fourier Transform. When the input is strictly real, i.e. the imaginary component is equal to zero for all points, the Fourier transform’s complex coefficients are conjugate symmetric,

$$\hat{u}[n_1, n_2] = \hat{u}[-n_1, -n_2]^*.$$

Taking advantage of this property, FFTW optimizes its memory allocation as depicted in Figure 2 where roughly only half of the Fourier transform’s last dimension’s coefficients are stored. The negative part of the Fourier spectrum for the second dimension can be retrieved by simply conjugating the stored Fourier coefficients. More concretely, supposing we have a mesh in real space of $N_1 \times N_2$ points, the size of the mesh in Fourier space equals $N_1 \times (\frac{N_2}{2} + 1)$ points, where the division is actually integer division (rounding down). As is common practice for numerical FFTs software, the wavenumbers are distributed in a particular ordering. For the second dimension, we have all non-negative wavenumbers in ascending order up to the value $\frac{N_2}{2}$. For the first dimension, however, the non-negative wavenumbers come first followed by the negative ones, always in ascending order. Lastly, FFTW executes the FFT through its datatype called a `plan`. For our specific spin-orbit problem, we will recur to both plans, `fftw_plan_dft_r2c_2d` which stands for real-to-complex and `fftw_plan_dft_c2r_2d` which stands for complex-to-real. The former executes the actual transform while the latter computes the inverse transform.

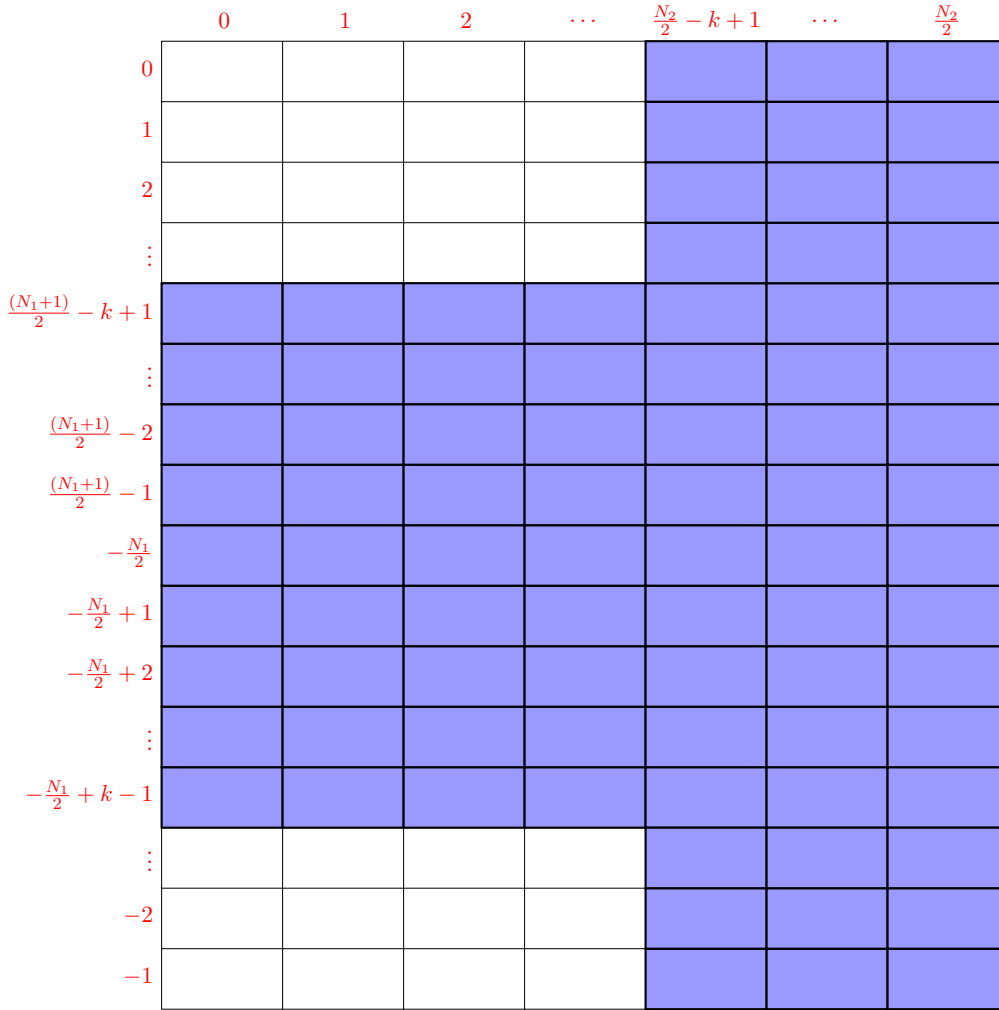


Figure 2: Visual representation of the 2D Fourier transform in FFTW for a grid size of $N_1 \times N_2$ in real space. The wavenumbers distribution in FFTW is shown by the red indices. The blue cells correspond to the filtering of width k applied to the Fourier coefficients. All divisions in the wavenumbers indices are integer divisions.

4.5 Continuation

Up to now, all that we have discussed applies for a fixed value of ε . From an initial approximation for the quasi-periodic attractor, (u, γ) , we iterate a few Newton steps to converge to, hopefully, machine precision. Recalling the statement in Theorem 2, this procedure is mathematically valid for any $\varepsilon \in [0, \varepsilon^*]$. Once past this critical value, the invariant torus breaks down and no quasi-periodic solutions exist to converge to numerically. “Continuing” our solution as ε grows away from zero allows us to approximate vaguely the range $[0, \varepsilon^*]$ for the invariant tori.

Numerically, estimating the asymptotic value of ε is rather subjective. As we will see in the results section, the error constantly increases along with the ε value and therefore setting the threshold error limit for the torus breakdown depends on one’s judgement. In our case, we are also affected by the limited computational resources at our disposal. Lastly, it is vital our initial guess be as close as possible to the actual solution in the continuation. Not only does this speed up convergence but in the worst case scenarios where the breakdown nears, convergence may not be attained if starting too far off from the (local) attractor.

Ultimately we are interested in studying how the quasi-periodic solutions change from $\varepsilon = 0$ to $\varepsilon = \varepsilon^*$. The analysis we did at the beginning of the theoretical section showed us that for

$\varepsilon = 0$,

$$\begin{aligned} y(t) &= \dot{x}(t) = \omega, \\ x(t) &= x_0 + \omega t, \end{aligned}$$

is an invariant plane in phase space $\mathbb{T}^2 \times \mathbb{R}$. We can actually say even more by looking at (1) with $\varepsilon = 0$ giving us the second-order linear ODE

$$\ddot{x}(t) + \eta \dot{x}(t) - \eta \nu = 0,$$

which has the general solution,

$$\begin{aligned} x(t) &= x_0 + \nu t + \frac{1 - e^{-\eta t}}{\eta} (v_0 - \nu), \\ \implies \dot{x}(t) &= \nu + e^{-\eta t} (v_0 - \nu), \end{aligned}$$

where x_0 is the initial angle and v_0 is the initial angular velocity. In dynamical terms,

$$\begin{aligned} x(t) &= \omega t = \nu t, \\ \dot{x}(t) &= \omega = \nu, \end{aligned}$$

is not only an invariant plane in phase space but actually a global attractor. Setting the velocity constant to ν nullifies the extra terms and in any case, the exponentials decay with time and therefore also vanish.

Remark. *WLOG we set $x_0 = 0$ for simplicity.*

This global attractor for $\varepsilon = 0$ will in fact bifurcate into a quasi-periodic attractor when $\varepsilon > 0$. Recalling the type of solutions (3) we are looking for, we see that initially for $\varepsilon = 0$,

$$\begin{aligned} x(t) &= \omega t + u(\omega t, t), \\ \implies u(\omega t, t) &= 0. \end{aligned}$$

Likewise, our γ parameter was defined as,

$$\begin{aligned} \gamma &= \eta(\omega - \nu), \\ \implies \gamma &= 0 \text{ for } \varepsilon = 0. \end{aligned}$$

The parametrizing tuple $(0_{\mathbb{T}_2^2}, 0)$ for the quasi-periodic attractor will therefore be the starting point for the continuation. ($0_{\mathbb{T}_2^2}$ refers to the 0-function on the 2-torus.)

Once the algorithm converges to (v, β) for some value $\varepsilon > 0$, we jump to a new greater value $\varepsilon_1 = \varepsilon + \delta$ and try to compute the parametrization of the quasi-periodic attractor for this greater ε_1 . Using (v, β) as our new initial approximations, we once again apply a maximum number of Newton steps, hoping to converge to the quasi-periodic attractor (v_1, β_1) . We repeat the procedure jumping to ε_2 , now setting (v_1, β_1) as our initial guess and converging to (v_2, β_2) , etc. We keep increasing the ε parameter as long as the numerics allow us to. The key point here is to use the previous tuple we converged to as the initial approximation for the next iteration.

There are however a potential few bumps in the road. First of all, as ε increases and comes “closer and closer” to ε^* , discontinuities start forming in the solution. The filtering discussed in the following section goes some way in dealing with this problem but is eventually rendered powerless past some ε value. As mentioned earlier, we also have to deal with the “small-divisors” problem constantly affecting the error in our algorithm.

We will need to slightly modify our algorithm to keep moving ε away from 0. As a first attempt, we simply rescale our jump δ by some fraction. It is possible that the jump is simply too large, meaning that our initial guess might be too far away from the solution to converge. The tradeoff is that we then jump with a smaller δ it takes more time and computations to

arrive at our numerical breakdown limit. Reducing δ consecutively many times is avoided as much as possible. Contrastingly, if the convergence is fast with very few Newton steps, we can increase the jump by some scalar greater than 1. On paper, this might seem to balance out but the reality is that closer to the critical value, divergence will be significantly more frequent than convergence, meaning our δ will overall scale down and eventually approach machine precision.

4.6 Grid Extension

We need a more robust approach to deal with such technical difficulties. Grid extension allows us to have a more precise representation of the solution. We are then better equipped to deal with the solution’s growing discontinuities giving us greater probability of reaching convergence. It also minimizes the need to reduce the jump, δ , completing the continuation process in less steps. Extending in real space would be extremely complicated and costly computationally with all the interpolation required. Once again, we take advantage of the FFT and extend in Fourier space.

The basic idea is fairly simple. We simply augment our solution’s Fourier transform with Fourier coefficients of value zero for higher order wavenumbers. This has no effect in real space when we apply the inverse on the extended Fourier grid since we have only added null values. We must be somewhat careful though when dealing with the specific order storage of FFTW. Figure 3 pictorially demonstrates how to align the coefficients from the original to the extended Fourier coefficients grid. In Figure 3, the orange cells correspond to the non-negative row coefficients which preserve their relative position in both matrices. The negative coefficients, corresponding to the green cells, however must drop a few rows down to also keep their relative position. The remaining white cells represent the new higher frequency nodes which we have set to zero. We also have to rescale the Fourier coefficients with the new dimensions in real space,

$$\rho = \frac{newN_1 \times newN_2}{oldN_1 \times oldN_2}$$

Once the grid in real space has been redimensioned accordingly, we obtain the corresponding extended solution in real space by taking the inverse transform on the lower matrix in Figure 3. After extending, every operation (differentiation, averaging, coefficient-wise multiplication, etc) takes more computing time but the jump δ will not need to be scaled down so frequently as before. Every iteration will now take more time, but we shall arrive at our “final” ε in less jumps.

We emphasize that extension only takes place once convergence has been reached and therefore all Fourier coefficients are considered. After every successful convergence, we check whether we have reached a critical stage where grid extension is required. After executing it or not, we increase ε and keep going with the continuation normally.

In order to determine when to extend, we look at the higher frequency nodes of the spectrum. We look at a band of width h to see if the higher frequency Fourier coefficients now have considerable magnitude that could be eventually troublesome. In a sense, we try and predict numerical instabilities before they happen. Once such a coefficient has exceeded the threshold, we compute the extension shown in Figure 3. h must be related to the grid size since the relative meaning of higher frequency wavenumbers changes as the grid dimensions do as well. The pattern of nodes we check is shown in Figure 2 except the width of the band is greater than the filtering band, $h > k$.

There is however a price to pay when extending. When the “small-divisors” problem was discussed, we concluded that the denominator value $\omega n_1 + n_2$ could be troublesome. The Diophantine property ensures its non-nullity, but may be very close to zero. Extension results in having more Fourier coefficients and more wavenumber values, therefore the range for n_1 and n_2 grows. The probability of landing on a tuple (n_1, n_2) resulting in a near null-denominator therefore rises every time we extend as more combinations of (n_1, n_2) are dealt with. This will

have an effect on the overall convergence and error of the scheme. In order to try to “smooth out” all these numerical issues, we recur to filtering.

4.7 Filtering

Being limited to finite precision, numerical instabilities will arise from various source as already mentioned. We must take this numerical noise, that will gradually start building up, into account and deal with it to avoid instability and/or divergence when looking for a new solution and executing Newton steps.

We take advantage of the fact that we are already using the FFT for the differential operators. We know that the numerical noise, if any, will appear gradually in the high frequency wavenumbers of the Fourier transform before possibly affecting the whole spectrum. Using a low-pass filter turns out to be a simple yet effective solution. We nullify a band of width k , as shown in Figure 2, of the extremities of the Fourier spectrum everytime the solution is updated in the Newton steps. In other words, we “smooth out” the solution in Fourier space as its core information is kept within the lower frequency modes. Everytime the solution is updated, it is immediately filtered.

The procedure is represented in Figure 2 by the blue grid cells. The width of the filtering band, k has been determined heuristically. It must not be too large otherwise, vital information about the solution might be deleted in the spectrum. Additionally, of course, the larger the width, the greater the number of nodes to delete. If this band is unnecessarily large, it affects the overall performance of the code as the filtering is executed quite frequently. As we have already stated, the filter band must be strictly smaller than the extension checking band, $h > k$. If that were not to be the case, we would never extend our grid since the Fourier coefficients checked for extension would have been nullified prior by the filtering. The continuation range would then be drastically reduced.

We want to stress that the methods discussed above are based on heuristics and observations over several trial and error runs. The numerical results below will show that filtering and grid extension are a necessity. Yet, the manner in which to do them (grid extension condition, width of the bands, etc) is very specific to the values of parameters chosen for the spin-orbit problem, i.e. the dissipation and frequency of motion. The computational resources available is also a factor. The bottom line is to try several different combinations of the above mentioned to find a somewhat optimal combination of the various numerical tools presented.

We resume the various ideas described above in the following guideline

Methodology 1.

- 1) Fix the maximum number of Newton steps for convergence, $maxIter$, as well as maximum number of Newton steps for “fast” convergence, $maxIterF$
- 2) Update and filter (k) the approximate solution at every Newton step
- 3) If # of Newton steps $< maxIter$ AND # of Newton Steps $< maxIterF$ increase jump by $(1 + \alpha Inc)$ for $\alpha Inc \in (0, 1)$. Else if convergence was not reached, we go to step 5)
- 4) If extension condition is TRUE (h), refine the solution’s mesh and go to step 6)
- 5) decrease jump by $(1 - \alpha Dec)$ for $\alpha Dec \in (0, 1)$
- 6) increase ε by jump and go to step 2)

Remark. The clear advantage in our methodology is that we start with a very coarse grid for small values of ε . This obviously speeds up the computations as a finer mesh is only necessary once ε approaches ε^* .

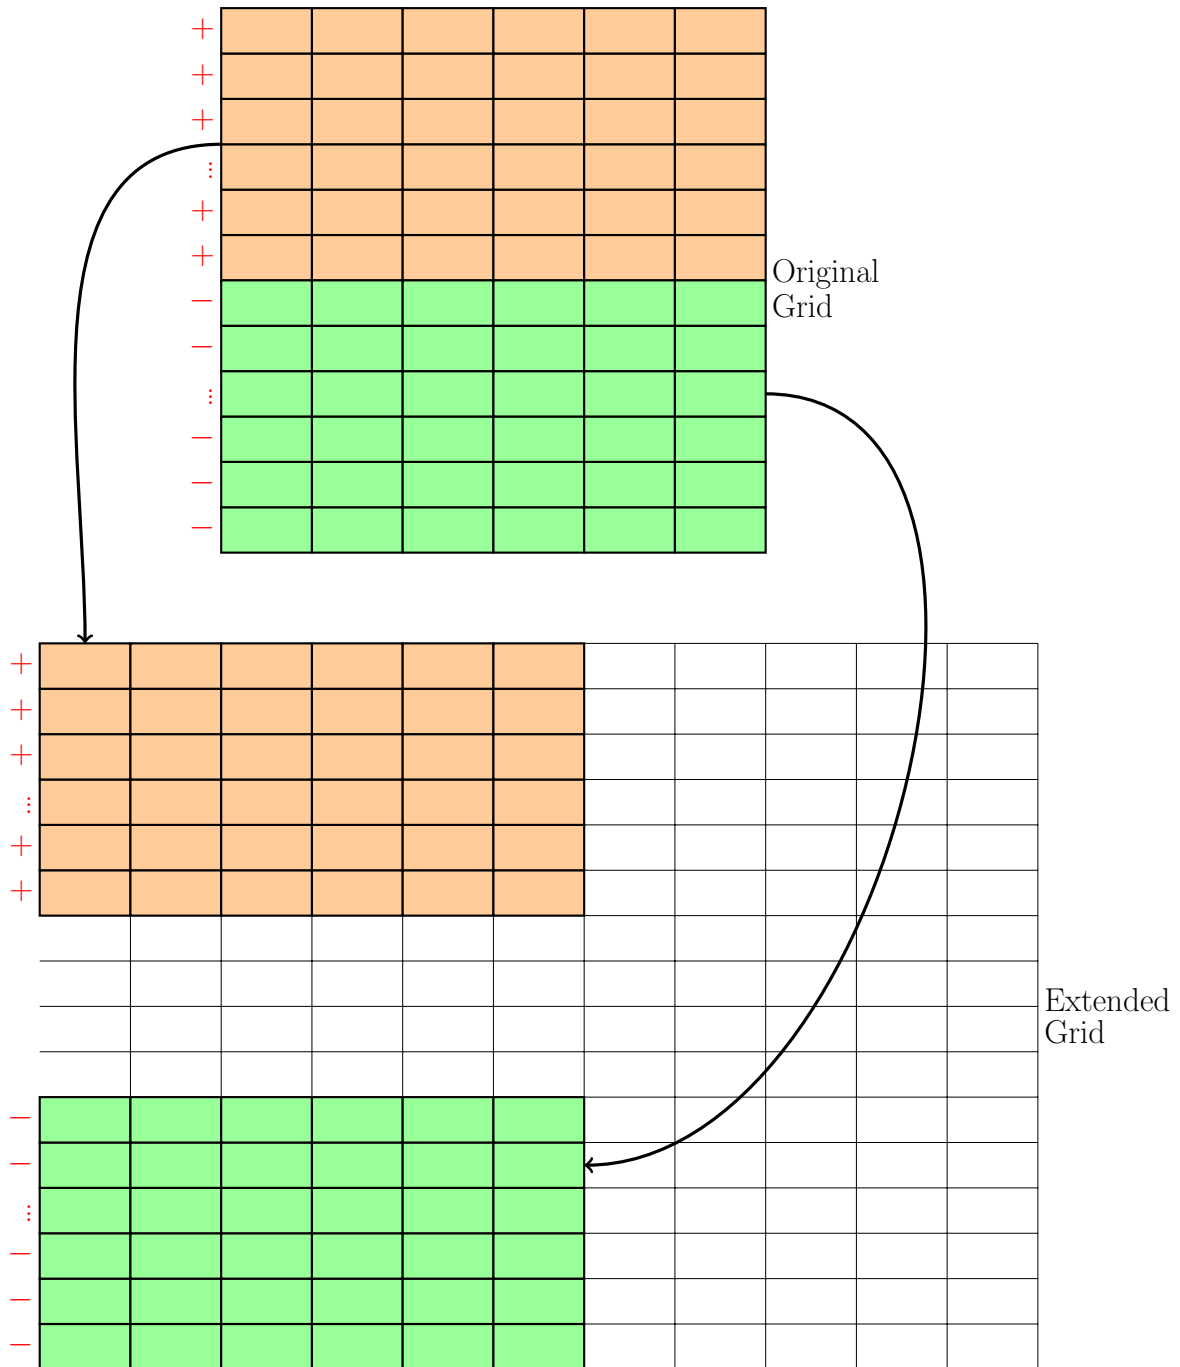


Figure 3: Visual representation of grid extension in FFTW

5 Numerical Results

5.1 Test Case

We enumerate the values used for our test case,

- computing specifications: 4GB of RAM, AMD Phenom(tm)II X4 945 (quadcore) processor
- $\omega = \frac{\sqrt{5}+1}{2}$, the golden ratio
- eccentricity, $e = 0.0001$

- $\eta = K\Omega_e = 0.25$
- maximum number of Newton steps for convergence = 10
- filter band width, $k = \frac{N_1}{10}$
- extension band width, $h = \frac{N_1}{5}$
- initial extension Fourier coefficients threshold = 10^{-11}
- initial tolerance, $tol = 10^{-16}$
- initial grid size: $N_1 = 66$, $N_2 = 108$ and maximum grid size: $N_1 = 752$, $N_2 = 1218$ (because of limited computational resources)
- $\|\cdot\|_1$ norm will be used to compute the error

5.2 Filtering

We show how the low-pass filter described in the previous section contributes to the convergence of the algorithm. Table 1 presented below illustrates the differences when running the algorithm directly or augmenting it with filtering. The test case is well behaved since the perturbation is quite small, $\varepsilon = 0.01$. The grid is coarse, its dimensions equal to the initial sizes. It does still highlight the benefits with the low-pass filter.

First of all and most importantly, the difference in minimum error order reached between both columns is quite staggering. For the filtering case in the third column, we reach the minimal error order possible under `long double` computations in C++ which is E^{-18} . However, the no-filter scheme can only go down to order E^{-15} meaning three orders of convergence less than with filtering. This is quite significant given the small perturbation considered. If we were to let ε grow, this difference would probably drastically increase.

Secondly, we clearly see how the no-filter scheme is not numerically stable. In the second column, from the fourth iteration where the minimum error order of E^{-15} is attained, the error starts growing quite impressively allowing us to conclude that the original basic algorithm is unstable. Consequently, we might encounter ourselves in a situation where without filtering the error comes really close to the desired tolerance but right before reaching it, ends up diverging extensively instead. Obviously, this is avoided in the filtering scenario and is another advantage with using the low-pass filter, making the algorithm numerically stable. As can be seen in the third column, by the fifth iteration we reach the minimum error possible. We let the algorithm run for a few more iterations yet the error does not budge, not even altering its digits in the same order of magnitude. To put it briefly, augmenting the numerical scheme with filtering results in numerical stability.

# Iterations	No Filter	Low-Pass Filter
1	0.0100028	0.0100028
2	7.78458e-05	7.78458e-05
3	1.6281e-09	1.6281e-09
4	3.50154e-15	4.3005e-18
5	9.20027e-13	4.3005e-18
6	2.46908e-10	4.3005e-18
7	6.67368e-08	4.3005e-18
8	1.80748e-05	4.3005e-18
9	0.0048976	4.3005e-18
10	1.32826	4.3005e-18

Table 1: Error comparison with/without filtering for the grid size: $N_1 = 66$ and $N_2 = 108$, $\varepsilon = 0.01$.

5.3 Small Divisors

We recall the issue arising from the small-divisors when applying the inverse operator D_0^{-1} . Since we require ω be Diophantine, we are guaranteed that

$$\omega n_1 + n_2 \neq 0$$

where as explained in the previous section, n_1 and n_2 are the Fourier wavenumbers. However, numerically such quantity can get really close to zero thus hindering the algorithm's convergence. Since we have now defined ω to be the golden ratio, we know that

$$\begin{aligned} \omega n_1 + n_2 &\approx 0 \\ \implies \frac{-n_2}{n_1} &\approx \omega \end{aligned}$$

It is also a well-known fact that the golden ratio is the limiting fraction of successive Fibonacci numbers,

$$\omega = \lim_{k \rightarrow \infty} \frac{F_{k+1}}{F_k}$$

Therefore the worst case scenario is when,

$$\frac{n_2}{n_1} = \frac{F_{k+1}}{-F_k}, \text{ for some } k \in \mathbb{N}$$

The reason we apply the minus sign to the denominator of the Fibonacci fraction is related to how FFTW manages the Fourier wavenumbers in 2D. As illustrated in the previous section, we know that the wavenumbers for the second dimension are nonnegative and only half of them are actually stored. The problem will therefore occur when the wavenumber for the first dimension lands on the negative Fibonacci numbers. Of course, we must extend our grid at some point to keep reaching a satisfactory convergence but we try to limit the problematic values originating from the small-divisors by only extending up to the next Fibonacci coefficients exclusively, i.e.

$$\begin{aligned} 0 &\leq n_2 \leq F_{k+1} - 1 \\ -F_k + 1 &\leq n_1 \leq F_k - 2 \end{aligned}$$

For example, our initial coarser grid has size $N_1 = 66$ and $N_2 = 108$. This is due to the fact that $F_9 = 34$ and $F_{10} = 55$. As we want to avoid the tenth Fibonacci number in the second dimension in Fourier space, we restrict the range of n_2 to $0 \leq n_2 \leq 54$. Since we only store half of the Fourier coefficients, this implies we actually have $54 \times 2 = 108$ grid points in real space for the second dimension. Similarly, for the first dimension, we want to avoid the ninth Fibonacci number for the first wavenumber in Fourier space. This means we can have up to 33 negative coefficients which implies a total of 66 coefficients when considering the nonnegative coefficients as well, i.e., $-33 \leq n_1 \leq 32$. As we predicted, the small-divisors greatest magnitude for each grid size occurs when the wavenumbers land on successive Fibonacci numbers.

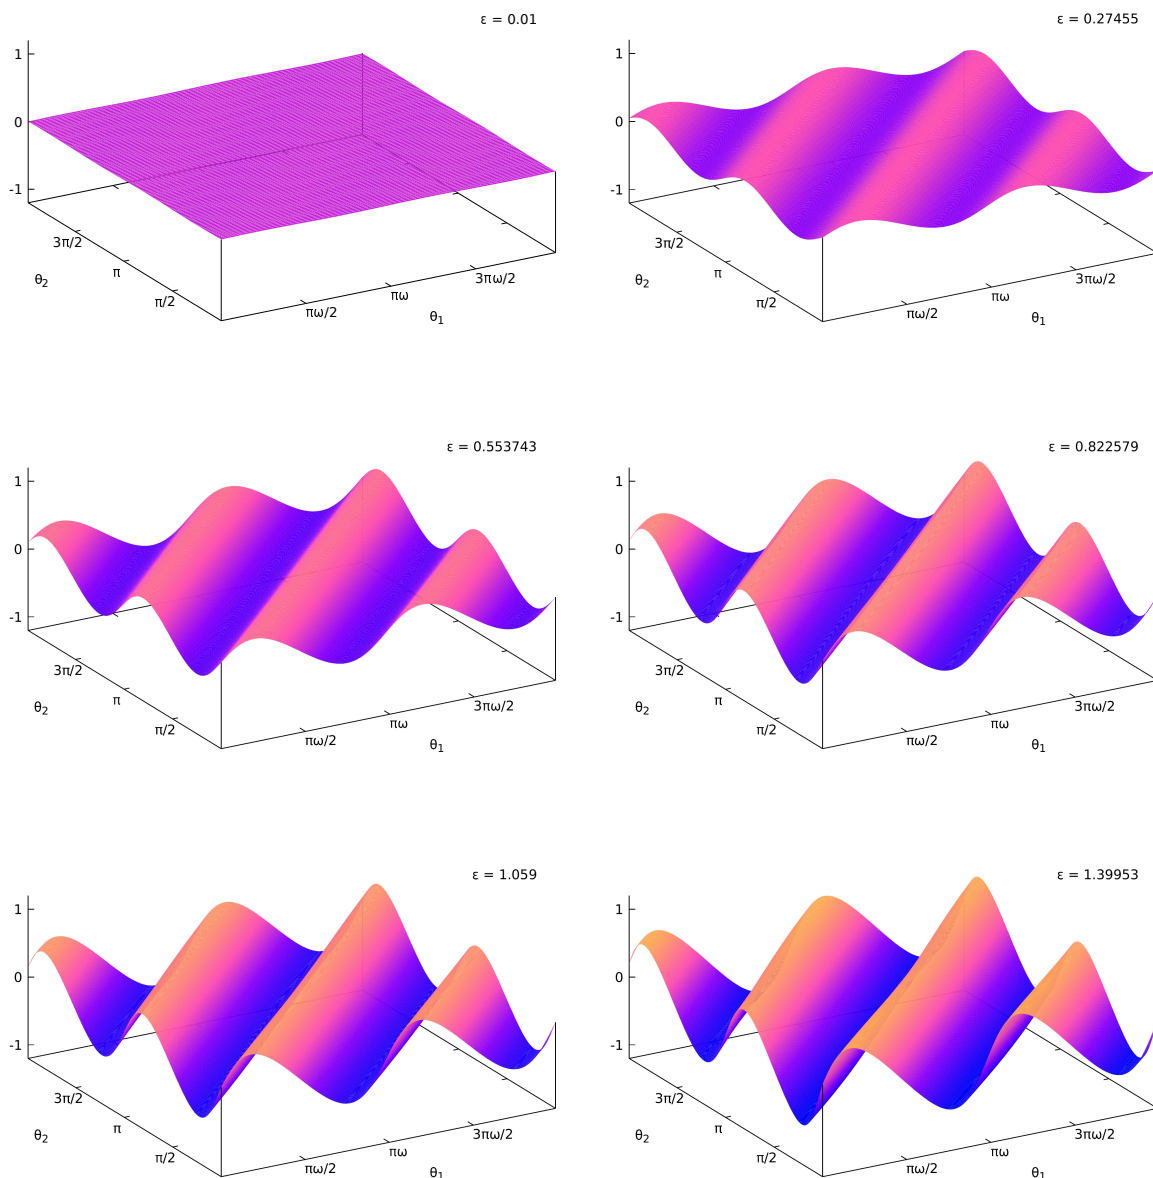
N_1	N_2	n_1	n_2	Small-divisor
66	108	-21	34	46.97871376374779026
108	176	-34	55	76.01315561749641806
176	286	-55	89	122.99186938124417127
286	464	-89	144	199.00502499874043902
464	752	-144	233	321.99689437998435773
752	1218	-233	377	521.00191937872542203

Table 2: Small-divisors values for the different grid sizes.

5.4 Quasi-Periodic Solutions

Here we present the actual quasi-periodic local attractor and more precisely, its evolution as ε grows. We can clearly see discontinuities forming in the solution as ε enlarges. Initially, we have the global attractor represented by the plane (first plot of Figure 4) as we had discussed in the theoretical section. In fact, this represents the regular 2-torus since both θ_1, θ_2 are periodic variables. We could retrieve the doughnut shape by connecting the extremities in both dimensions. The plane then deforms, compressing on itself as a zigzag, where crests and troughs start to appear. We observe that their amplitude grows showing the attractor loses smoothness with ε .

Near the breakdown of the torus, the amplitude almost reached 1.2 in both negative and positive z-axis directions. The peaks' orientation is related to ω as we will see in the contour plots below. One characteristic of the attractor which remains unchanged throughout the simulation is the number of crests and troughs, four for each. In the last row of plots of Figure 4, we are extremely close to the breakdown of the attractor. We see the peaks no longer form a straight line along their extremas and the surface is no longer smooth. Little kinks start to form all over the surface.







Figure 4: Quasi-Periodic Solution Continuation as ε grows.

The previous graphs somewhat fail to clearly show the increase in amplitude for the attractor as ε moves away from zero. From the perspective in Figure 5, we see that the amplitude reaches a maximum of about 1.2 in absolute value. The slope between each peak and crest becomes steeper, the surface's curvature also increases. In the second plot of Figure 5, the crests are described by a simple U-shape. The same applies for the peaks but with a \cap -shape. This is no longer possible for the fourth plot with little kinks and deformations appearing over the surface and the extremas no longer aligned. To some extent, the attractor twists on itself. For example, we see that the major peak in the fourth plot now resembles a saddle point. Basically, the surface becomes very noisy near the breakdown.






Figure 5: Quasi-Periodic Solution Amplitude Plots as ε grows.

Here we show one last perspective of the attractor, in this case its contour plots for different values of ε . As we have mentioned in the previous figures, the peaks retain the same slope throughout the continuation and in fact, this slope is equal to $\frac{1}{\omega}$. This is actually not surprising as we defined ω to be the frequency of motion. Another way of seeing this is recalling that $\theta_1 = \omega t$ and $\theta_2 = t$. As we specified before, the orientation of the attractor initially respects the frequency of motion. Its crests and troughs form aligning themselves with $\frac{1}{\omega}$. At some point, however, when approaching the quasi-periodic attractor's breakdown, this behaviour is lost and the crests and troughs lose this "linearity". The solution distorts itself.




Figure 6: Quasi-Periodic Solution Contour Plots as ε grows.

In the plot of Figure 7, we present the parameter γ vs. ε . In the numerical scheme, we defined the parameter gamma as, $\gamma = \eta(\omega - \nu)$. We stress again that the value for η, ω are fixed at the beginning of the numerical simulation meaning the correction in the parameter γ is directly applied to ν . The magnitude of ν is therefore always greater than ω and strictly increases throughout the continuation since γ is always negative and monotone decreasing.




Figure 7: γ parameter over the continuation

Another indication of the loss of analyticity in the quasi-periodic attractors is shown in Figure 8. The Sobolev norm of u is computed, more specifically $\|u\|_{H_5}$, along with the parameter ε . Unsurprisingly, its value exponentially blows up very close to the limiting case corresponding to approximately $\varepsilon \approx 2.7$ i.e. the breakdown of the attractor. The plot might suggest the Sobolev norm is close to being constant prior to the breakdown. This could not be further from truth. The norm is constantly increasing with ε yet nowhere near as drastically as when the attractor breakdown occurs.




Figure 8: Sobolev Norm over the continuation

Figure 9 presents the Fourier spectrum of the quasi-periodic solution for different values of ε to highlight the impact of the extension procedure. The obvious observation is that the magic happens on the negative coefficients of the first dimension, kx . The white line delimits the band where Fourier coefficients are checked for the extension. The first plot on the top left shows the spectrum of the solution for $\varepsilon = 1.53155$ where the extension is about to take place. We can see two strips of Fourier coefficients crossing the white line but the lower strip is responsible for the extension. The top right graph shows the spectrum after the extension has taken place at the next step of the continuation namely, $\varepsilon = 1.57133$. First of all, the graph is way smoother since we now have many additional Fourier coefficients. Secondly, we see that the upper strip still crosses the white line and yet we have just extended the spectrum, indicating that the upper strip is never responsible for the extension. The bottom plot shows the spectrum near the breakdown, $\varepsilon = 2.56804$. The interesting aspect here is that the two strips have now merged.




Figure 9: Fourier Spectrum of the Quasi-Periodic Attractor over the continuation.

The last aspect of the continuation that is of our interest is ε^* . Recalling the discussion in section 4.5, this parameter is the theoretical limit on ε , corresponding to the torus breakdown. Given our specific numerical setting (precision, grid extension method, etc.) we were able to reach as far as $\varepsilon^* = 2.66951$. We emphasize that this is only a numerical approximation of the theoretical value, only giving us a general idea on its magnitude. From that point onwards, our algorithm could no longer converge for any reasonable error tolerance.

5.5 Error

Unsurprisingly, the bottom plot of Figure 10 shows how the error blows up when the breakdown is (numerically) reached. Prior to the exponential blowup, we can observe that the error follows the trend of a step function (see Figure 10) yet always monotonically increasing. These increments or spikes can be better distinguished in the top graph of Figure 10 where we limit the ε range to before the exponential blowup. The spikes actually correspond to grid extensions and demonstrate how decisive the small-divisors are in the overall convergence. For example, referring to Table 2, the spike at around $\varepsilon \approx 1.5$ is due to the third grid extension taking place at $\varepsilon \approx 1.57$ and the small-divisor maximum value jumped from approximately 122 to 199. Another clear asymptotic increase in the top plot occurs when $\varepsilon \approx 2.3$. Again, around this ε value, we extended the grid for the fourth time bringing the maximum small-divisor to about 321.

It is clear from the plots that the small divisors play a major role in dictating the error magnitude. As is to be expected, the jumps in the error grow along with ε . The quasi-periodic solution plots show us that we start off with a perfectly smooth torus (plane) and discontinuities start to form as ε increases. It therefore follows that the variation in the error is less for small values of ε . The discontinuities in the solution and the small divisors only really start to notably affect the error once $\varepsilon > 0.5$. Another observation from the error plots of Figure 10 is the scale of the magnitude of the error. The range we have at the breakdown is in the order of magnitude

of E^{-12} yet we initially start off with an error of order magnitude E^{-18} . This difference shows how much the error tolerance must be relaxed during the continuation.




Figure 10: (top) Error over the continuation prior to blow-up. (bottom) Error over the whole continuation.

5.6 Satellite Angle

Most of the results presented deal with the quasi-periodic component of our original spin-orbit problem (1). It is important to mention how these quasi-periodic attractors relate to the actual variable at stake, the satellite's angle. Our original problem deals with the periodic variable $x(t)$ but we describe it with quasi-periodic motion. The physical intuition behind these results is we will never land on the same exact spot twice because of the Diophantine frequency, at least in finite time. This might seem paradoxical since the angle, $x(t)$, was defined as a periodic variable as the satellite revolves around the big mass body. What we actually mean is the combination of the satellite's own orientation and the angle with the ellipse's major axis will never be repeated at any point as it orbits around the planet.

Another aspect of our solution that we must mention is the eccentricity. The way we set up our problem, we fix the eccentricity, $e = 0.0001$, over the whole continuation and tweak the parameter $\nu = \nu_\epsilon$. The satellite therefore orbits around the big-mass body in almost circular fashion as the size of the perturbation, ϵ , increases. In fact, we have witnessed that convergence becomes more difficult for greater values of the eccentricity. However, this is a simplification of the more realistic model presented in [1] where the eccentricity is actually tweaked over the

continuation, $e = e_\varepsilon$, and the parameter $\nu = \nu_e$ depends on the eccentricity and not ε directly. The reason we modify the model from [1] is purely numerical. The continuation algorithm we derive in this work immediately diverges for $\varepsilon > 0$ when we consider the eccentricity, e_ε , as our continuation parameter and not a fixed value.

6 Conclusion

We were able to compute the quasi-periodic attractor u numerically through an efficient numerical Newton scheme. We used the FFT to develop a fast numerical implementation of the derived algorithm. We therefore confirmed numerically the various stated results in [1]. There are of course many obvious extensions arising from this work.

First of all, we could look at different ways to develop a criterion and implementation for the extension. In our case, we purely looked at the small divisors in real space to decide how large to grow the grid. Ultimately, since the extension takes place in Fourier space and the spectrums in Figure (9) display a visible pattern, extending our grid based on some criterion of the spectrum instead could be a viable option.

Secondly, a natural and rather important characteristic in making our model more realistic would be to consider the parameters η, ν not as constants anymore but really as functions of time, t . Of course, the algorithm itself would have to be modified since this assumption was key in our derivation of the numerical scheme. In fact, the spin-orbit model derived in [2] considers the parameters ν, η as functions of time and eccentricity. The reduction to the constant parameters case comes from averaging the functions over one period.

Lastly, it would be interesting to do these calculations with extended precision and more computational power to get more accurate predictions on the breakdown of the attractor.

References

- [1] Celletti A. and Chierchia L., *Quasi-Periodic Attractors in Celestial Mechanics*, Archive for Rational Mechanics and Analysis, Vol.191 (2009), p.311-345.
- [2] Alessandra Celletti, *Stability and Chaos in Celestial Mechanics*, Springer-Praxis, 2010.
- [3] Calleja R., Celletti A., and de la Llave R., *KAM Theory for Conformally Symplectic Systems: Efficient algorithms and their validation*, J. Differential Equations, Vol.255 (2013), no.5, p.978-1049.
- [4] Calleja R. and Celletti A., *Breakdown of Invariant Attractors for the Dissipative Standard Map*, Chaos, Vol.20 (2010), 013121.
- [5] Calleja R. and Figueras J.-Ll., *Collision of Invariant Bundles of Quasi-Periodic Attractors in the Dissipative Standard Map*, Chaos, Vol.22 (2012), 033114.
- [6] Arnold, V.I., *Encyclopedia of Mathematical Sciences, Dynamical Systems III, Vol. 3* Springer, Heidelberg, 1988
- [7] Matteo Frigo, http://www.fftw.org/fftw3_doc/, FFTW 3.3.6-pl1 User Manual, 15-01-2017.