EXACT COHERENT STRUCTURES AND DYNAMICAL CONNECTIONS IN A QUASI 2D KOLMOGOROV LIKE FLOW

A Dissertation Presented to The Academic Faculty

By

Ravi Kumar Pallantla

In Partial Fulfillment of the Requirements for the Degree Doctor of Philosophy in the School of Physics

Georgia Institute of Technology

May 2018

Copyright © Ravi Kumar Pallantla 2018

EXACT COHERENT STRUCTURES AND DYNAMICAL CONNECTIONS IN A QUASI 2D KOLMOGOROV LIKE FLOW

Approved by:

Dr. Roman O. Grigoriev , Advisor School of Physics *Georgia Institute of Technology*

Dr. Michael F. Schatz School of Physics *Georgia Institute of Technology*

Dr. Predrag Cvitanović School of Physics *Georgia Institute of Technology* Dr. Flavio H. Fenton School of Physics *Georgia Institute of Technology*

Dr. Yingjie Liu School of Mathematics *Georgia Institute of Technology*

Dr. Rafael de la Llave School of Mathematics *Georgia Institute of Technology*

Date Approved: December 5, 2017

Dedicated to my mom, Satyavathi. Love you amma.

ACKNOWLEDGEMENTS

I thank my advisor, Dr. Roman Grigoriev, without whom this thesis would not have been possible. His expert guidance has trained me to be a better researcher. The immense freedom he gave has helped me to work on several problems, try various solutions and paved the way for learning a lot of concepts. His systematic approach to understand science using the fundamentals has helped me in effectively solving the problems. His enthusiasm for science has always been very motivating. Roman was always available and approachable, which was required especially during the times I encountered roadblocks in the research. He being an excellent teacher, has clearly explained and answered my questions. Particularly, the 'Pattern Formation' course offered by him was very helpful and formed the basis of developing several codes used in the simulations carried out in this research work. In addition to the academic guidance, he also supported me during various difficult phases in the past years. I can never forget his help in the last six months of my graduate student life, especially his feedback in writing this thesis. He is the best advisor one could ask for, and I thank him for all his help. I also thank my co-advisor, Dr. Michael Schatz, whose expertise in experiments has made possible several of our projects. His constant encouragement for lucid explanation of concepts and studying them from the basics has helped me in gaining clarity in understanding the problems. I am very thankful for the financial support he provided me in my last semester as a PhD student. I thank my thesis committee members for their invaluable time.

The seeds of my enthusiasm for physics have been sown by the excellent teachings of my high school teachers, K.V. Srinivasa Rao and Samir Biswas. This enthusiasm was boosted by the lectures of Dr. Arul Lakshminarayan, Dr. Suresh Govindarajan, Dr. Pattabiraman, Dr. Lakshmi Bala and Dr. Balakrishnan during my undergraduate studies at IIT Madras. This was further motivated by Dr. Haridass and Dr. Srikanth during my internship as an undergraduate at PoornaPrajna Institute. All these teachers have encouraged me to pursue research, and have also helped me in their best ways. I thank all of them for their encouragement.

I thank the School of Physics at GeorgiaTech for offering me an admission into their graduate program. My graduate studies would not have been possible without the financial support provided by the School in my first two years as a PhD student, and by the National Science Foundation and the Defense Advanced Research Projects Agency in the later years. The School of Physics has a very active and vibrant atmosphere for conducting research. The School's Center for Nonlinear Sciences has provided me a platform for several collaborations and discussions with some of the best minds in the research area of nonlinear sciences. I had very exciting discussions with its students – Chris Marcotte, Tongran Qin, Radford Mitchell, Xiong Ding, Mohammad Farazmand, Adam Kamor, Simon Berman, Doga Murat, Burak Budanur, Matthew Gudorf, Chris Crowley, Balachandra Suri, Jeff Tithof, Logan Kageorge, Daniel Borrero, Grace Chambers, Patrick Reinbold, Bret, Joshua Bernett – to name a few. The collaborations with Jeff Tithof, Logan Kageorge, Balachandra Suri, Patrick Reinbold have been very successful and I am looking forward for more collaborations with them in the future. I thank Balachandra Suri and Raford Mitchell, who developed some of the basic codes required for performing simulations in this research work.

The Atlanta life had been a fun journey having some of my best moments with the company of Ananda Barua, Sabyasachi Deyati, Sudanth Gudladona, Rohit Doi, Naomi Chopra, Atul Hodge, Ruhollah Heydari, Mansi Agrawal, Tanushree Ganguly, Tanvi Olivia Singh, Chris Crowley and John Indergrad. The Thursday night trips to Athens with Sudanth, Rohit, Nercy, parties in the physics estate with Chris, John, David Caroll, David Jackson, Johnathan, Samee, Raj, Bharat, Palit, Sabyasachi, trips outside Atlanta with Sudanth, Mansi, Sabyasachi were very fun moments. I also had very thought provoking discussions outside academics with Bala, Chris, Sabyasachi, Mansi, Ananda, Palit, Tanvi, Pradeep, Sarvani, Trilochan, Prasanth. My interactions in the past six years with people

from diverse backgrounds exposed me to a diversity of cultures and thoughts. It included several people – Ananda Barua, Sabyasachi Deyati, Aritra Benerjee, Debashis Benerjee, Arindham Khan, Ayan Chakraborthi, Mohan Rajendran, Piyush Sao, Atul Hodge, Aseem Grover, Sahil Kapoor, Faisal Ahmed, Ruhollah Heydari, Nimit Nigania, Rohit Doi, Anusha Garapaty, Naomi Chopra, Nercy, Kshitiz Gupta, Mason Bogue, Owen Vail, Mansi Agrawal, Sudanth Gudladona, Naveen Bommakanti, Krishna Chaitanya, Tanushree Ganguly, Tanvi Singh, Chris Crowley, John Indergrad, Balachandra Suri, Palit Subhobrata, Manu Mohan, Pradeep Prathapa, Sathya Sarvani, Trilochan Rambhatla, Prasanth Alapati, and the list goes on.

Last but not the least, I thank the almighty and my family for their blessings, love, well wishes and support throughout.

TABLE OF CONTENTS

Acknov	vledgme	ents	iv
List of]	Figures	· · · · · · · · · · · · · · · · · · ·	ĸi
Summa	ry		ii
Chapte	r 1: Int	roduction and background	1
1.1	Introdu	uction	1
1.2	Backg	round	4
1.3	Motiva	ation for the present work	0
1.4	Thesis	outline	4
Chapte	r 2: Mo	del system	5
2.1	Experi	mental setup	6
2.2	Numer	rical model	8
	2.2.1	Incompressible model	8
	2.2.2	Weakly compressible model	9
	2.2.3	Nondimensionalization	25
	2.2.4	The models and experiment	6
	2.2.5	Symmetries of the system	28

	2.2.6	State space	28
2.3	Summ	ary	30
Chapte	r 3: Tra	ansition from laminar to turbulent flow	31
3.1	Pre-tu	rbulent flow	32
	3.1.1	Straight flow	32
	3.1.2	Primary instability	35
	3.1.3	Secondary instability	41
3.2	Two-d	imensional turbulence	44
	3.2.1	Transition to turbulence	44
	3.2.2	Recurrences	49
	3.2.3	Intermittency	50
	3.2.4	Periodic window	52
3.3	Summ	ary	54
Charte	- 1 . No	retar based wethods for commuting FCS and dynamical company	
Cnapte	r 4: INe tio	ns	57
4.1	Comp	uting equilibria and periodic orbits	57
	4.1.1	Newton's method	59
	4.1.2	GMRES	62
	4.1.3	Hook step	64
4.2	Comp	uting dynamical connections	66
	4.2.1	Shooting method	66
	4.2.2	Symmetry constraint	72

4.3	Summary	6
Chapte	r 5: Adjoint-based methods for computing ECS and dynamical connec- tions	8
5.1	Adjoint evolution equation	8
	5.1.1 Illustrative example	0
	5.1.2 Accelerating convergence	3
5.2	Computing equilibria	4
	5.2.1 Choice of the weight and acceleration matrices	6
	5.2.2 Numerical solution of the adjoint equation	7
	5.2.3 Results	9
5.3	Computing dynamical connections	0
	5.3.1 Spectral representation	1
	5.3.2 Tuning the spectral representation	5
	5.3.3 Compressible and incompressible model	6
	5.3.4 Results	7
	5.3.5 Optimal choice of the scaling parameter	3
	5.3.6 Discussion	4
5.4	Computing periodic orbits	7
5.5	Summary	9
Chapte	r 6: Conclusions	1
6.1	Main scientific contributions	1
6.2	Open questions	2

Appendix A: Spatial discretization and numerical integration					
A.1	Gauge condition				
A.2	Spatial discretization				
A.3	Numerical integration				
A.4	Jacobian computation				
Append	Appendix B: Parametric continuation				
Appendix C: Spectral formulation for periodic orbits					
Referen	i ces				

LIST OF FIGURES

1.1	Cross sectional view of the coherent structures and exact coherent structures in a pipe flow. The top panel shows the coherent structures observed in experiment at (A) $Re = 2000$ (C) $Re = 2500$ (E) $Re = 5300$. The bottom panel shows the ECS computed in numerical simulations at (B) $Re = 1250$ (D) $Re = 1360$ (F) $Re = 2900$. Reproduced from Ref. [24]	5
1.2	Plane Couette flow in a minimal flow unit. Shown are the snapshots of the flow along a periodic orbit P97 at intervals $\Delta t = 15$ marked by open magenta dots in Figure 1.3, starting at the point labeled P97. Reproduced from Ref. [44].	7
1.3	A state-space portrait of turbulent plane Couette flow. This 3D projection shows a turbulent trajectory u_{turb} (solid and dotted black lines) shadowing the P97 periodic orbit (bold magenta line) and the unstable manifolds (blue, red, and green lines) of symmetry-related equilibria (solid blue, red, and green dots; the black dot at the origin is the laminar flow state). Dynamical connections are shown as bold red and blue lines connecting different equilibria (filled dots). The Figure is reproduced from Ref. [44].	8
2.1	Schematic of the experimental setup with the (a) top view and (b) side view.	17
2.2	The x-component of normalized forcing profile, $\overline{f}_0(x, y)$, used in the simulations, as a function of the nondimensional coordinates. The y-component of the forcing is identically zero.	25
3.1	The dependence of Re (red) and \tilde{Re} (blue) on the current I . Panel (b) shows a zoomed-in version of the circled region in panel (a)	33
3.2	(a) x-component and (b) y-component of velocity field of the straight flow in experiment at $Re = 8.06$. The driving current used in the experiment is 7 mA	34

3.3	(a) Vorticity field of the straight flow in experiment at $Re = 9.25$ ($I = 8$ mA) and (b) in simulation at $Re = 9.26$ ($I = 7.65$ mA). Comparison of the vorticity fields on the lines (c) $x = 0$ and (d) $y = 0$ between experiment (red) and simulation (blue).	36
3.4	(a) Vorticity field of the modulated flow in experiment at $Re = 14.01$ and (b) in simulation at $Re = 14.03$. The driving currents used are, respectively, $I = 14$ mA and $I = 13.89$ mA	37
3.5	(a) Leading stability exponents of the 3 equilibria near the imperfect pitch- fork bifurcation. The branch containing E_1 and E_2 is plotted in magenta while E_3 , E_4 are plotted respectively in blue and black. E_3 undergoes pitch- fork bifurcation (shown as cyan sphere) by breaking the rotational symme- try \mathcal{R} . (b) Projections of these 3 equilibria along with E_5 and E_6 (in red), plotted as the driving current varies. The saddle node bifurcation of E_3 and E_4 is shown as brown sphere	38
3.6	Vorticity fields (from simulations) of the 3 equilibria present in the imper- fect pitchfork bifurcation - (a) \mathbf{E}_2 (b) \mathbf{E}_3 and (c) \mathbf{E}_4 at $I = 10.56$ mA. This current, far away from the bifurcation, is chosen as the flow fields look indistinguishable to the naked eye close to the bifurcation	39
3.7	Vorticity fields which correspond to eigenvectors (a) e_1 and (b) e_2	40
3.8	Order parameter \hat{V} (cf. Eq. (3.3)) as a function of Re for the modulated flow. The experimental data is plotted in red and the simulation data in blue.	41
3.9	(a) Stability exponents of the steady modulated state at $I_2 = 19.4$ mA $(Re_2 = 17.38)$. The two complex conjugate exponents crossing the imaginary axis indicate a Hopf bifurcation. (b) Floquet multipliers of the periodic orbit PO ₁ at $I_4 = 20.3$ mA $(Re_4 = 17.96)$. Complex conjugate Floquet multipliers crossing the unit circle (red) indicate a Hopf bifurcation. In both panels 25 leading eigenvalues are shown.	42
3.10	Power spectrum intensity as a function of frequency for various Re in (a) experiment and (b) the compressible model.	43
3.11	Order parameter (peak power) at various Re and a square root fit for the data are plotted for experiment (red), the compressible model (blue), and the incompressible model (black).	44

3.12	The power spectrum $P(f)$ at different I in the compressible model demon- strates the hysteresis near $I_4 = 20.3$ mA. The top panel shows the power spectrum when I is quasistatically increased beyond I_4 , and the bottom panel shows the power spectrum when I is decreased. The magnitude of I (in mA) is shown at the top of each plot. All the plot have the same range, as shown for the plot in the top left corner	46
3.13	Poincare sections at various driving currents at (a) $I = 20.2$ mA, (b) $I = 20.3$ mA, (c) $I = 20.31$ mA, and (d) $I = 20.35$ mA showing the transition from quasiperiodicity to turbulence. To illustrate the breakup of the 2-torus, different colors have been used in panels (b) and (c) to label the eight disconnected components of the resulting attractor (e.g., <i>n</i> th crossing is color-coded based on the values of $n \mod 8$).	48
3.14	Sample recurrence plots at various driving currents. (a) At $I = 15$ mA the flow approaches an equilibrium, which is a global attractor. (b) At $I = 20$ mA the flow approaches a periodic orbit, which is also a global attractor	50
3.15	Recurrence plots showing intermittency just above Re_4 in (a) simulation at $Re = 18.2$ and (b) experiment at $Re = 18.1$.	51
3.16	Recurrence plots in the simulations of compressible model for $I > I_4$. (a) $I = 20.4$ mA and (b) $I = 20.6$ mA demonstrate that the laminar periods in the intermittent dynamics become shorter as I increases farther away from I_4 . (c) $I = 21.1$ mA shows chaotic dynamics. (d) $I = 21.4$ mA shows intermittency near the left edge of the periodic window. (e) $I = 21.5$ mA shows dynamics inside the periodic window. (f) $I = 21.9$ mA and (g) $I = 22.6$ mA demonstrate that the laminar periods in the intermittent dynamics become shorter as I increases farther away from I_8 . (h) $I = 25$ mA shows chaotic dynamics.	53
3.17	Recurrence plots in the stable island, where the state asymptotically settles in to a periodic orbit. The recurrence plots are computed in (a) simulation at $Re = 19.03$ and (b) experiment at $Re = 20.1$	54
3.18	The dominant 25 Floquet multipliers of the periodic orbit, PO ₂ , plotted in the complex plane for at $I_7 = 21.6$ mA ($Re_7 = 19.03$). At this driving, the (black) Floquet multipliers cross the unit circle (in red) resulting in a period-doubling bifurcation.	55

3.19	A sketch showing periodic solutions PO_2 , PO_3 , and PO_4 and their stability inside the periodic window. Stable (unstable) branches are shown using solid (dashed) lines and the circles indicate solutions computed using New- ton iterations. The horizontal axis shows the driving current I in mA and the vertical axis corresponds to a direction in the state space along the lead- ing Floquet multiplier of PO_2 at I_7 . The blue portions of the branches are based on continuation, while the red portions are speculated based on PO_3 computed at one value of I (red open circle)	55
4.1	Sample plot of state space speed $S(t)$. Local minima below the threshold $S_0 = 0.1$ are denoted with black circles. The values of $S(t)$ are $O(1)$ because the system has been nondimensionalized such that both the relevant length and time scales are unity.	58
4.2	Sample recurrence plot at $I = 25$ mA, where the dynamics is turbulent. A local minimum in the recurrence below a threshold of $d_0 = 0.1$ is circled in black.	59
4.3	Heteroclinic connection (green) between unstable equilibria E_3 (blue) and E_4 (black) computed at a set of different driving currents to illustrate the continuation procedure. The equilibria at intermediate values of the current are shown as dashed lines and bifurcations as spheres (brown for the saddle-node, cyan for the pitchforks, and red for the Hopf bifurcation. The same projection is used as in Figure 3.5b.	71
4.4	Bifurcation diagram showing equilibrium states E_1 through E_8 (lines). Pitch- fork bifurcations are labeled with cyan spheres and Hopf bifurcations are labeled with red spheres. Saddle-node bifurcations are not labeled; they correspond to points where equilibria E_3 and E_4 or E_7 and E_8 meet. The same projection is used as in Figure 3.5b.	72
4.5	Vorticity fields of (a) \mathbf{E}_7 , (b) \mathbf{E}_8 and the eigenvectors (c) \mathbf{e}_1 , (d) \mathbf{e}_2 of \mathbf{E}_7 at $I = 11.9$ mA, where the heteroclinc connection is computed	74
4.6	(a) Projections of the trajectories that start on the unstable manifold of the origin equilibrium \mathbf{E}_7 (blue sphere) onto the subspace spanned by the orthogonal vectors \mathbf{d}_1 , \mathbf{d}_2 , \mathbf{d}_3 which are linear combinations of \mathbf{e}_1 , \mathbf{e}_2 , and $\mathbf{E}_8 - \mathbf{E}_7$. Without the symmetry constraint, all trajectories (black curves) terminate at the global attractor \mathbf{E}_2 (magenta sphere). With the symmetry constraint, the trajectory (green curve) from $\mathbf{E}_7 + \epsilon \mathbf{e}_2$ terminates at the destination equilibrium \mathbf{E}_8 (red sphere). (b) Zoomed view of (a) near \mathbf{E}_7 showing the focusing effect of the dominant unstable direction.	75

4.7	The distance from the destination equilibrium \mathbf{E}_8 versus time for the symmetry constrained trajectory with the initial condition $\mathbf{E}_7 + \epsilon \mathbf{e}_2$.	76
5.1	Comparison of the solutions to the adjoint evolution equations with and without acceleration. (a) The residual $E(\tau)$. (b) The "trajectories" $\mathbf{X}(\tau)$, the level sets of E (cyan curves), and the solution \mathbf{X}^* (magenta circle). Blue lines correspond to the steepest descent, black lines – to the accelerated descent, and red – to the momentum method. The same choice of Q and W has been used in the accelerated descent and the momentum method	81
5.2	The histogram showing the success rate of finding equilibria using the adjoint method, as a function of I .	90
5.3	Equilibria found using hybrid adjoint-Newton solver at a driving current $I = 24$ mA. Shown are the vorticity fields of the three equilibria (a) E_9 , (b) E_{10} and (c) E_{11} .	91
5.4	Convergence of (a) Newton/GMRES method and (b) the adjoint method for the heteroclinic connection between \mathbf{E}_3 and \mathbf{E}_4 at $I = 9.6$ mA. The Krylov subspace dimension (10, 100, 200, or 300) is shown in the figure legend	98
5.5	(a) Spectral coefficients for the connection between \mathbf{E}_3 and \mathbf{E}_4 for various choices of m_0 and ω . (b) A 2-D projection of the connection onto the plane spanned by the two unstable eigenvectors \mathbf{e}_1 and \mathbf{e}_2 of the origin \mathbf{E}_3 , with the continuous green line representing the result of shooting and the symbols showing the collocation points for the spectral representation. Both (a) and (b) use the same color-coding: $m_0 = 10$ and $\omega = 0.1$ (blue), $m_0 = 41$ and $\omega = 0.01$ (red), $m_0 = 79$ and $\omega = 0.01$ (black).	100
5.6	Zoomed view of the connection (a) near the origin E_3 and (b) near the des- tination E_4 at $I = 11.31$ mA. The connection computed using the spectral method is shown in blue and the connection computed using the shooting method is shown in green.	101
5.7	Stability spectra of the origin equilibrium (a) \mathbf{E}_3 and the destination equilibrium (b) \mathbf{E}_4 at $I = 11.31$ mA.	102
5.8	Projections of the heteroclinic connection between the origin E_7 (blue sphere) and destination E_8 (red sphere). The connection computed using the sym- metry constraint is shown as the green curve. The collocation points on the connection computed using the spectral/adjoint method are shown as black asterisks. The same projection as in Figure 4.6 is used here.	103

- 5.9 The effect of the scaling parameter ω on the accuracy of the spectral representation for the connection between \mathbf{E}_3 and \mathbf{E}_4 at I = 9.6 mA. (a) The evolution of the residual. (b) The decay of spectral coefficients. The values of ω used are 0.01 (blue), 0.05 (red), and 0.1 (black); in all cases $m_0 = 41$. 105
- A.1 Schematic of the staggered grid with $n_x = 3$, $n_y = 4$. The locations of the fields u,v,p are respectively shown using blue circle, cross and dot, and the boundaries of the domain are represented by the red lines. The figure also shows the ghost points that are defined outside the physical domain. 119

SUMMARY

Turbulence in fluid flows is ubiquitous. It is present in rivers, oceans, and the atmosphere. The flow of air past the wings of an airplane or around a car is turbulent, the flow of water around a boat or down a water pipe is also turbulent. Scientists and engineers have been trying to understand turbulence for centuries, yet it remains a rather mysterious phenomenon. This thesis explores a modern approach for describing turbulence using unstable nonchaotic (e.g., equilibrium or temporally-periodic) solutions of the Navier-Stokes equation, called exact coherent structures (ECS). These solutions are closely related to classical coherent structures – characteristic spatiotemporal patterns that appear fleetingly and disappear only to reappear at a different place and time – which are found in a variety of wall-bounded turbulent flows, both in experiment and in simulations. It has been conjectured that a hierarchy of ECS forms the skeleton of fluid turbulence: each ECS guides the dynamics of the flow in its vicinity, with apparent randomness arising from the flow moving from the neighborhood of one ECS to the neighborhood of another.

Some of the most unexpected discoveries in developing and testing this conjecture have been made in a system – Kolmogorov-like flow in a thin fluid layer of electrolyte driven by Lorentz force suspended above a thin lubricating layer of a fluid dielectric – which is essentially two-dimensional. Due to its effective two-dimensionality, this flow offers an unprecedented level of access both experimentally and numerically, allowing the kinds of analysis that would be either prohibitively expensive or simply too difficult in most threedimensional flows. The price one pays for this simplification is the requirement of using an approximate two-dimensional model, which relies on some assumptions. We have developed, and implemented numerically, an improved "weakly-compressible" model of the flow which retains the simplicity of the "incompressible" two-dimensional model introduced previously, but also accounts for the thickness variation of the two fluid layers. This compressible model has been shown to offer a more accurate description of the transition to turbulence in this system, compared with its incompressible predecessor.

Most of the previous work has focused on identifying collections of ECS in various canonical wall-bounded flows and using periodic orbit theory to connect temporal averages with state averages. However, relatively little attention has been devoted to using the hierarchy of ECS to describe and predict the global dynamics of turbulent flows. In particular, there are almost no studies that explain how and why the flow moves from the neighborhood of one ECS to the neighborhood of another. One possibility is that the turbulent flow follows another type of unstable solutions to Navier-Stokes – heteroclinic connections that lie at the intersection of the unstable manifold of an origin ECS and the stable manifold of the destination ECS. To investigate this possibility, these connections have to be computed, however no reliable methods for computing them for such complicated systems as a fluid flow have been developed previously. Existing algorithms have been developed mainly for low-dimensional systems and become computationally prohibitive for fully-resolved discretizations of Navier-Stokes in realistic geometries. We have developed and tested several robust and efficient numerical algorithms for computing both ECS and dynamical connections in high-dimensional dynamical systems. These algorithms should facilitate the next step in the development of the geometric, deterministic description of fluid turbulence.

CHAPTER 1 INTRODUCTION AND BACKGROUND

1.1 Introduction

The flow of Newtonian fluids such as water and air, for which the viscous stress is proportional to the velocity gradient, is governed by the Navier-Stokes equation (NSE), which was introduced in early 19th century and represents momentum conservation

$$\rho\left(\partial_t \mathbf{v} + \left(\mathbf{v} \cdot \nabla\right) \mathbf{v}\right) = -\nabla p + \mu \nabla^2 \mathbf{v} + \mathbf{f}_b,\tag{1.1}$$

where ρ is the density of the fluid under consideration, v is its velocity, μ is the dynamic viscosity, p is the pressure, and \mathbf{f}_b is the density of the body force acting on the fluid. For flows with velocity much less than the speed of sound, mass conservation leads to the incompressibility condition

$$\nabla \cdot \mathbf{v} = 0. \tag{1.2}$$

The NSE is deterministic: in principle, given an initial flow state, it unambiguously predicts the flow state at any later time. Nevertheless, in practice this determinism only holds when the flow is sufficiently slow (and hence laminar). For fast flows, due to its nonlinearity, the NSE possesses chaotic solutions (which describe turbulent flows) characterized by exponentially fast separation of close initial conditions, such that in practice the deterministic predictions can only be made over a finite period of time. Also as a result of nonlinearity, solutions cannot be expressed in closed form in any geometry, which substantially complicated their mathematical analysis.

The first systematic investigation to understand fluid turbulence was performed circa

1883 by Osborne Reynolds, who injected dye into the flow of fluid through a pipe to visualize the flow [1]. The experiments showed that the complexity of the flow was determined solely by the dimensionless parameter

$$Re = \frac{v_s \ell}{\nu} \tag{1.3}$$

where $\nu = \mu/\rho$ is the kinematic viscosity of the fluid, v_s is the characteristic velocity scale of the flow, and ℓ is the characteristic length scale of the problem, such as the pipe diameter. Originally introduced by Stokes in 1851, this parameter now carries Reynolds' name and can be physically interpreted as the ratio of inertial and viscous forces in the flow, At small Re, viscous forces dominate and the flow is laminar, i.e., smooth and simple. As the Reynolds number increases, inertial forces become dominant and the flow becomes turbulent, i.e., complicated and unpredictable. Though Reynolds was able to explain the transition from laminar to turbulent flow qualitatively, he failed to produce any kind of a quantitative description. In fact, in his subsequent paper [2], he concluded that turbulent flow is too complicated for detailed understanding, and by borrowing ideas from the kinetic theory of gases he laid the foundation for a statistical description of turbulence by factoring the flow into the mean and fluctuating components. Since then, most of the research in fluid turbulence has focused on its statistical properties [3, 4, 5, 6]. However, given the deterministic nature of the NSE, a deterministic description should be possible. Not only can it complement – and justify – the statistical description, the deterministic description is crucial for describing the physical mechanisms responsible for initiating and sustaining turbulent flow.

The pioneering effort in constructing a deterministic description of turbulence was made in 1944 by Landau [7]. He proposed that the laminar flow, stable at low Re, becomes unstable and gives rise (through a Hopf bifurcation) to a time-periodic flow as Re is increased further, this time-periodic flow undergoes another Hopf

bifurcation, giving rise to a a flow with two temporal frequencies. The process continues with increasing Re, leading to a quasi-periodic flow with a large number of incommensurate frequencies. Landau argued that the number of different frequencies tends to infinity as Re increases, resulting in an irregular and apparently random flow behavior, characteristic of turbulence [7]. In 1948, Eberhard Hopf constructed a theoretical model that reproduced this scenario [8] and proposed using the corresponding solutions to compute the statistical properties of turbulence. Though never observed in experiments, this scenario marks a paradigm shift in our approach to the turbulence problem. In 1971, Ruelle and Takens showed that the scenario based on an infinite sequence of Hopf bifurcations is not generic and introduced a refined version [9]. By viewing NSE as a dynamical system, they showed that the transition to turbulence can involve a finite number of bifurcations. This was later confirmed in experiments [10, 11], suggesting a possible link between dynamical systems theory and turbulence.

To make matters more complicated, in addition to being nonlinear, the NSE also involves a large number of degrees of freedom. This can be seen by spatially discretizing it, which formally yields an infinite set of coupled ordinary differential equations (ODEs). The first successful attempt to describe temporally complex fluid flows using dynamical systems theory is due to Edward Lorenz, who in 1963 derived a drastically simplified model of convection rolls in the atmosphere by truncating a spectral discretization of the NSE obtained by Saltzman [12]. His simplified model involving three coupled nonlinear ODEs is presently known as the Lorenz system [13]. Using numerical simulations, Lorenz discovered that his simple model generates solutions that are chaotic, i.e., highly sensitive to initial conditions, similar to fluid turbulence thus strengthening the connection between turbulence and the deterministic chaos. The seminal work of Lorenz in many ways defines the modern view of turbulence through the lens of dynamical systems theory. The last decade has seen a lot of progress in this direction in part due to advancements in computing power, development of sophisticated numerical algorithms, and carefully designed experiments [14]. Although still in its early phase, a deterministic description of turbulence is poised to revolutionize our understanding of fluid turbulence.

1.2 Background

Despite their apparent complexity and unpredictability, turbulent flows commonly exhibit characteristic patterns known as "coherent structures". For example, in wall-bounded shear flows (e.g., the boundary layer in flow past an airfoil), coherent structures are typically observed in the form of streamwise vortices [15] and streaks of streamwise velocity [16]. An example of coherent structures in a pipe flow is shown in Figure 1.1. Coherent structures in the form of vortices are frequently observed in the ocean and atmosphere [17]. Coherent structures typically play a central role in turbulent transport and mixing [18]. Coherent structures are transient; for example, in wall turbulence, vortices and streaks exhibit a bursting process [19, 20, 21]. Nevertheless, coherent structures are robust features of turbulent flows in that they frequently and repeatedly reappear for a wide range of Reynolds numbers [22, 21]. Empirical criteria have typically been used to characterizate coherent structures in turbulence; in wall-bounded flows, for example, turbulent streaks are identified by well-specified protocols for conditional sampling [23].

Recent theoretical work assisted by large scale numerical simulations has shown that observed coherent structures can be connected directly to *unstable* exact solutions of the NSE [25, 26, 27, 28, 29, 30, 31, 32, 21, 33, 34, 35] as shown in Figure 1.1 for a pipe flow. These exact solutions therefore became known as "exact coherent structures" (ECS). In wall-bounded shear flows, ECS prominently feature, like classic coherent structures, both streamwise vortices and streaks. ECS also exhibit bursting behaviors and can account for much of the turbulence energy production attributed to classic coherent structures [26, 27, 21]. Most theoretical studies of ECS have focused on flows in straight channels and pipes at moderate Reynolds numbers ($\sim 10^3$); nevertheless, evidence suggests that this approach extends to much larger Reynolds numbers ($\sim 10^4$ or more) [31].



Figure 1.1: Cross sectional view of the coherent structures and exact coherent structures in a pipe flow. The top panel shows the coherent structures observed in experiment at (A) Re = 2000 (C) Re = 2500 (E) Re = 5300. The bottom panel shows the ECS computed in numerical simulations at (B) Re = 1250 (D) Re = 1360 (F) Re = 2900. Reproduced from Ref. [24].

These recent results suggest that ECS could provide the building blocks for novel models of turbulent flow. The foundation stone for this line of thought can be traced back to Poincaré [36] who realized that unstable periodic orbits provide a skeletal structure that organizes chaotic dynamics of celestial objects. This idea was later developed in the context of quantum chaos by Gutzwiller [37] and subsequently applied to high-dimensional chaos generated by nonlinear PDEs such as the Kuramoto-Sivashinski equation [38, 39] and the Ginzburg-Landau equation [40]. Although it took the fluid dynamics community some time to take notice, the same idea eventually started producing new insights into the dynamical mechanisms underlying turbulence [41, 26]. The last decade has seen an explosion of research in weakly turbulent flows, mostly in Europe and Japan (see [14] for a recent review), giving us hope that the "the greatest unsolved problem of classical physics" might finally surrender, paving the way for our understanding, and simplified deterministic description, of spatiotemporally complex dynamics in numerous other physical and biological systems.

Although the mathematical framework is still under development, its structure has

started to crystallize. PDEs, such as the NSE (1.1), formally describe infinite-dimensional systems, but the drastic reduction (due to dissipation) of the effective number of degrees of freedom suggests that it should also be possible to reduce the complexity of the dynamical description. As a system is driven further out of equilibrium (e.g., Reynolds number Re of a fluid flow is increased), the degree of nonlinearity is increased, formerly stable spatially and temporally simple solutions become unstable, and new (mostly unstable) solutions appear through sequences of bifurcations. Bifurcation cascades generate an infinity of unstable solutions which lead to the emergence of chaotic sets, leading to chaotic dynamics (and turbulence, in particular). This (standard in low-dimensional dynamical systems) scenario has recently been verified in high dimensions through numerical simulations for several canonical fluid flows (plane Poiseuille flow [42] and pipe flow [43]).

In principle, the dynamics of turbulence takes place in an infinite-dimensional state space, the space formed by the set of all possible flow states. Each point in the state space corresponds to a snapshot of a particular solution of the governing PDEs (e.g., a flow state in the physical space for fluids [45]). In particular, each unstable solution in the state space corresponds to an ECS in the physical space, as Figure 1.2 illustrates. For practical purposes (e.g., real-time prediction of turbulent flow), a low-dimensional description of the dynamics must be constructed. Previous efforts have largely relied on global lowdimensional models constructed by projecting the NSE onto a collection of modes (i.e., disturbances around the laminar flow profile) extracted using either proper orthogonal decomposition (POD) [46, 47] and, more recently, balanced truncation methods [48, 49], dynamic mode decomposition [50, 51] and Koopman operator technique [52]. These efforts have achieved some measure of success: they provide a *qualitative* description of the self-sustaining processes leading to regeneration of turbulent coherent structures. However, due to fundamental limitations stemming from their inherent linearity, these methods are largely unsuitable for *quantitative* description of turbulent dynamics, where both nonlinearity and fine-scale structure must be properly represented, especially at higher Reynolds



Figure 1.2: Plane Couette flow in a minimal flow unit. Shown are the snapshots of the flow along a periodic orbit P97 at intervals $\Delta t = 15$ marked by open magenta dots in Figure 1.3, starting at the point labeled P97. Reproduced from Ref. [44].



Figure 1.3: A state-space portrait of turbulent plane Couette flow. This 3D projection shows a turbulent trajectory u_{turb} (solid and dotted black lines) shadowing the P97 periodic orbit (bold magenta line) and the unstable manifolds (blue, red, and green lines) of symmetry-related equilibria (solid blue, red, and green dots; the black dot at the origin is the laminar flow state). Dynamical connections are shown as bold red and blue lines connecting different equilibria (filled dots). The Figure is reproduced from Ref. [44].

numbers.

The discovery of ECS suggests a radically different modeling strategy to overcome the quantitative failures of previous low-dimensional models by describing turbulence as *a walk within a network of unstable ECS*. The turbulent dynamics or the spatiotemporally chaotic dynamics take place on a chaotic set (attractor or repeller) which is embedded in a relatively low-dimensional inertial manifold (of order ten or so for weakly turbulent fluid flows in small three-dimensional computational domains [53]). Inertial manifold is, in turn, embedded within an infinite-dimensional state space. As Figure 1.3 illustrates for a weakly turbulent plane Couette flow (PCF) [54], unstable *nonchaotic* solutions (often referred to as invariant solutions in the literature) inhabit the same region of state space as chaotic solutions and are dense within the chaotic set: a sufficiently long chaotic trajectory approaches any of these unstable solutions arbitrarily closely. These unstable solutions (equilibria, periodic and quasi-periodic solutions, etc.) are all saddles, i.e., they have both attracting (stable) and repelling (unstable) manifolds. A *statistical* description of the dynamics can

be obtained by invoking the ergodic hypothesis, assuming it is valid, to relate temporal averages of any physical observable to an appropriately weighted average over the saddles embedded in the chaotic set. The weights are related to the stability of saddles, with the least unstable ones providing the largest contribution to the average, according to the periodic orbit theory (POT) [55, 56], hence there is a clear hierarchy of saddles.

The *dynamical* description based on unstable solutions is considerably less developed, but numerical simulations suggest the following general picture. Once a turbulent trajectory (dotted line in Figure 1.3) enters the neighborhood of a particular saddle (e.g., timeperiodic orbit shown as the magenta curve) it follows (shadows) that saddle for a period of time (the solid black line on the right side of the figure), before leaving that neighborhood and moving towards another saddle. This process then repeats, with the turbulent trajectory wandering from neighborhood to neighborhood. If one associates the saddles with the nodes of a network or graph, then the sequence of the nodes visited by the turbulent flow provides a *coarse* description of the dynamics. If the initial conditions are well specified, this sequence is deterministic. In particular, the turbulent trajectory is expected to follow dynamic (or heteroclinic) connections between the saddles (solid red and blue lines). The corresponding segments are shown as solid black lines on the left of the figure. If the initial conditions *are not* well specified, then one can associate a probability with each of the dynamical connections originating at a given node and then the sequence becomes a random walk on a Markov chain, from which one can readily obtain the statistical description. It is worth emphasizing that for these descriptions, the ECS and their connections need only be determined once as these invariant objects do not change in time.

The *geometrical* interpretation associated with the state space representation of the dynamics has similarly proved to be extremely useful. In particular, the shape of ECS in the state space and the evolution of the corresponding flows structures in the physical space were found to have a deep physical significance. In particular the dynamical mechanisms that underlie turbulence have previously been understood, with the help of coherent structures, only qualitatively. Identification of ECS allowed a quantitative description, the selfsustaining process identified by Walleffe [57] being the most famous example, as well as the realization that their is an entire hierarchy of such mechanisms. Some saddles (termed edge states) were found to define the local shape of the chaotic set, with their stable manifolds define the boundary between the basins of attraction of the laminar flow and the chaotic attractor supporting turbulent flow [58, 59]. In contrast, the *topological* information associated with the network of dynamical connections that links distant parts of the chaotic set, on the other hand, has been largely unexplored.

The framework based on unstable solutions also allows construction of a *fine* description of the dynamics. Although global embedding of the entire chaotic set may require a vector space of very high dimensionality (determined by the discretization of the PDEs required to fully resolve the dynamics), locally the inertial manifold is low-dimensional, and in the neighborhood of any saddle/connection it can be parametrized using a relatively small number of degrees of freedom (O(10)). This should be contrasted with more conventional model reduction techniques such as POD or balanced truncation, which may require many thousands of modes to represent quantitatively the dynamics near even the simplest time-periodic solutions. The local degrees of freedom can be obtained conveniently by linearizing the infinite-dimensional governing equations around each saddle/connection, keeping the slow (i.e., unstable, marginal, and weakly stable) modes and discarding an infinity of strongly stable modes. The linearization around ECS determines the probability for the turbulent evolution to follow various dynamical connections, providing the last missing link between the dynamical and the statistical descriptions.

1.3 Motivation for the present work

Although substantial progress has been made in developing the foundation of the deterministic, dynamical description of fluid turbulence in recent years, many of its ingredients remain largely unexplored, untested, and not understood. In particular, we still don't know what types of ECS (equilibria, periodic orbits, quasi-periodic orbits) play the most important role in the dynamics. For instance, the dynamically dominant solutions appear to be relative periodic orbits in pipe flow [60] and in two-dimensional Kolmogorov-like flow *without* the Rayleigh friction term [61], while in the same flow *with* the Rayleigh friction term equilibria appear to be dynamically dominant [62]. The role of quasi-periodic solutions is not currently understood, primarily due to the practical difficulty in computing solutions featuring several incommensurate frequencies. The role of dynamical connections is equally poorly understood, again due to the difficulty of computing this type of solutions. We have very limited data regarding the statistics of visits to the neighborhoods of different solutions by the turbulent flow and, in the cases when such data is available it appears to contradict the predictions of POT [61].

The dynamical description of turbulence requires the computation of both the ECS and the connections between them. However, the combination of non-linearity and highdimensionality of the NSE make the problem of computing these solutions extremely challenging. Several different approaches for computing absolute and relative equilibria and time-periodic solutions have been proposed and tested. The shooting method, which relies on time-integration of the NSE combined with Newton-hook-step iterations [63], is the most common approach. It solves a system of coupled nonlinear equations that defines a point in the state space the either corresponds to an equilibrium or lies on a closed orbit. A variational approach, where the entire orbit is discretized using finite differences [64] or spectral representation [65], offers a more robust alternative for highly unstable solutions for which the shooting method generally breaks down. Both of these approaches rely on the Newton's method to solve the resulting system of equations, which is computationally expensive and prone to stagnation at local minima. To overcome these limitations, an adjoint-based method has been developed for computing absolute and relative equilibria [66].

In comparison, no systematic methods for computing dynamical (homoclinic or hetero-

clinic) connections for the NSE currently exist. Most existing methods are shooting-based [67, 68, 69] and were only tested for low-dimensional systems. Whether they are applied for computing time-periodic orbits or dynamical connections, shooting-based methods, especially ones that employ Newton iterations, require a good initial guess. Such good initial guesses are generally unavailable for dynamical connections of the NSE and can only be generated under very restrictive conditions. For instance, in all of the examples where dynamical connections have been computed for turbulent fluid flows [27, 31, 54], the origin ECS possessed a very low-dimensional unstable manifold (one- or two-dimensional), which made a brute-force search for good initial guesses tractable. Such brute-force approach becomes intractable for typical ECS that have more than a couple of unstable directions. Variational methods, originally developed for computing time-periodic solutions, offer a viable alternative: a curve in the state space connecting the origin and destination ECS is discretized and evolved in pseudo-time in such a manner that the tangent along this curve is everywhere aligned with the direction of the vector flow. Existing implementation of the variational method [70] relies on Newton iterations, which makes it impractical for computing the connection in fully-resolved simulations of fluid turbulence with physical boundary conditions due to the immense computational power and storage requirements. A lack of memory efficient, robust numerical algorithm for computing the dynamical connections motivated some of the work presented here. In this thesis, this problem is addressed by developing an adjoint-based numerical solver for computing the dynamical connections between unstable equilibria of the NSE, which is both memory-efficient and robust. In addition, it is shown that adjoint-based approach can also be used for computing unstable equilibria and time-periodic orbits, though computation of periodic orbits has not been pursued in this thesis.

Another open problem addressed by this thesis is the lack of a sufficiently accurate two-dimensional model that can be directly compared with an experimentally realizable (nearly) two-dimensional flow, such as the Lorentz-force-driven flow in a thin layer of electrolyte supported by a thin lubricated layer of fluid dielectric above a horizontal rigid boundary [71]. Although both the fluid layers in such an experimental setup have a finite thickness, the flow in the electrolyte layer can be made very nearly two-dimensional due to the strong confinement in the vertical direction [72]. Two-dimensional flows have been widely used to investigate fluid turbulence [73, 74, 75, 62], since they reproduce much of the phenomenology of the three-dimensional turbulence, including subcritical transition and the energy and enstrophy cascades, but are much easier to study experimentally and numerically.

The flow in the experimental system considered in this thesis is very similar to the Kolmogorov flow – a theoretical model of transition to turbulence that has been extensively studied theoretically, numerically, and experimentally [76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89]. The Kolmogorov flow is assumed to be strictly two-dimensional and is driven by a steady body force that varies sinusoidally in space,

$$\mathbf{f}(x,y) = f_0 \sin(ky)\hat{x},\tag{1.4}$$

where f_0 is the strength of the forcing, which defines the characteristic flow velocity v_s , and $2w = 2\pi/k$ is the forcing wavelength. Since theoretical and numerical studies typically ignore the effect of the lateral boundaries and instead assume periodic boundary conditions in both directions, w – one half of the forcing period – defines the characteristic length scale ℓ of the flow. In experiment, the forcing is nearly sinusoidal and the flow is nearly two-dimensional, hence we refer to it as Kolmogorov-like. A two-dimensional model for this Kolmogorov-like flow has previously been developed [71, 72] by ignoring the thickness variation of fluid layers. Comparisons of this model with the experiment show reasonable agreement at low Reynolds numbers, when the flow is steady. However, as the *Re* increases, the vertical component of the velocity becomes non-negligible and the thickness of fluid layers becomes time-dependent. Taking these effects into account allows extending

the validity of the model to higher Re.

1.4 Thesis outline

The rest of the thesis is organized as follows. We start by presenting a derivation of the improved, weakly compressible, two-dimensional model of the Kolmogorov-like flow in Chapter 2. The model is validated by comparing its predictions with experimental observations in Chapter 3. This chapter also discusses the sequence of bifurcations leading from laminar flow to turbulence as the Reynolds number increases. Chapter 4 describes the traditional Newton-based shooting methods for computing ECS and dynamical connections. Chapter 5 introduces the new, robust, memory efficient, adjoint-based solver for computing equilibria and the dynamical connections between them. Finally, the main results of this work and some remaining open questions are discussed in Chapter 6.

CHAPTER 2 MODEL SYSTEM

The deterministic, dynamical description of turbulence requires validation against both numerical simulations and experiments. As discussed in the introduction, turbulent flows in two spatial dimensions are much easier to study, compared with their 3D counterparts. This motivated the studies of 2D turbulence in soap films [90] and thin layers of liquid metals [91] and liquid electrolytes [92]. In these experiments, the fluid is strongly confined in one direction, making the flow effectively two-dimensional. Soap-film experiments are poorly suited for quantitative comparisons, since the thickness (and hence the inertia) of the film cannot be made sufficiently uniform. On the other hand, experiments which use thin layers of fluids supported by a rigid horizontal boundary feature flows that are not formally two-dimensional: the no-slip boundary condition at the bottom of the container imposes a nonuniform vertical flow profile. The effect of the bottom wall has been traditionally modeled by adding an empiric Rayleigh friction term to the two-dimensional Navier-Stokes equation. A formal first-principles derivation of both the vertical flow profile and the 2D evolution equations for the horizontal flow profile featuring the Rayleigh friction term have been made in Ref. [71] by assuming that the vertical component of the velocity is negligible. Predictions of this model were found to be in good agreement with experiment for steady flows in pre-turbulent regime [72]. However, at higher Re, when the flow becomes time-dependent, substantial discrepancy between the experiment and the model predictions is found. To extend the validity of the 2D model to higher Re, where the flow becomes turbulent, a new 2D model has been derived in this chapter by accounting for the vertical component of the velocity. This improved model has been compared with experiments and found to reproduce well the details of the observed subcritical transition to turbulence.

2.1 Experimental setup

The experimental setup employing a shallow layer of fluid electrolyte driven by Lorentz force was originally introduced by Bondarenko [92]. The no-slip boundary condition at the bottom generates a vertical gradient of the horizontal velocity, which gives rise to a nonvanishing vertical component through Ekman pumping [93, 94]. To address this problem, a setup using two immiscible fluid layers was proposed by Rivera *et al.* [95], where a lighter layer of electrolyte is suspended on a heavier layer of dielectric fluid. The flow is driven by an electromagnetic force in the top electrolyte layer, while the bottom dielectric layer acts as a liquid lubricant and helps generating a nearly two-dimensional flow in the top layer. Immiscibility of the two layers suppresses the flow along the vertical direction and also enhances the two-dimensionality of the flow.

The Lorentz force

$$\mathbf{f} = \mathbf{J} \times \mathbf{B} = J(\hat{x}B_z - \hat{z}B_x). \tag{2.1}$$

driving the flow is generated by the interaction of a DC electric current with density $\mathbf{J} = \hat{y}J$ passed through the electrolyte and the magnetic field **B** produced by an array of magnets with alternating polarity placed below the fluid. The experimental setup considered here consists of 14 magnets with magnetization in the vertical direction (*z*-axis), as shown in Figure 2.1a, in order to generate a nearly sinusoidal forcing profile (in the *y*, or transverse, direction), like in Kolmogorov flow. A thin glass plate of thickness 0.079 ± 0.005 cm is placed above the magnets to compensate for the surface irregularities at the bottom. The shallow layers of dielectric (thickness $h_d = 3$ mm, viscosity = 1.30 mPa·s, density = 1769 kg/m³ at 23.0°C) and electrolyte (thickness $h_c = 3$ mm, viscosity = 5.85 mPa·s, density = 1192 kg/m³ at 23.0°C) are placed on the top of the glass plate. The entire setup is placed in a rectangular aluminum container of dimensions $L_x \times L_y = 22.86$ cm \times 17.78 cm, with a buffer region of $d_x = 1.27$ cm and $d_y = 2.54$ cm between the magnet array and the



Figure 2.1: Schematic of the experimental setup with the (a) top view and (b) side view.

container walls, as shown in Figure 2.1a. The fluid flow, driven by the Lorentz force in the x-direction, is constrained by the walls normal to x-axis (henceforth referred to as the end walls). Two electrodes of length L_x each, mounted to the walls parallel to the x-axis, are used to drive the electric current through the electrolyte (these are henceforth referred to as the side walls). Since the dielectric is nonconducting, J = 0 there. In the electrolyte layer the current density is assumed uniform; a net electric current I corresponds to a current density $J = I/(h_c L_x)$.

Each of the magnets in the magnet array is 15.24 cm long and 1.27 cm wide. Both the electrolyte and the dielectric layer are thin (compared with the forcing length scale which is determined by the width of one magnet w), essentially making the fluid flow in both layers horizontal. The speed of the flow, however, depends not only on the x and y coordinates, but also on the z coordinate, so we will refer to this flow as a quasi-2D flow. This makes the flow different from strictly 2D flows such as that considered by, e.g., Chandler and Kerswell [61].

2.2 Numerical model

2.2.1 Incompressible model

Let the origin of the coordinate system be at the geometric center of the bottom of the dielectric layer. The velocity field in the fluids is inherently three-dimensional, in the sense that it generally depends on all three spatial coordinates and time t, $\mathbf{v} = \mathbf{v}(x, y, z, t)$. It is governed by the Navier-Stokes equation (1.1) with $\mathbf{f}_b = \rho \mathbf{g} + \mathbf{f}$, where $\rho \mathbf{g} = -\rho g \hat{z}$ is the gravitational force and \mathbf{f} is the Lorentz force.

For shallow layers of fluids driven by a weak, in-plane forcing, the velocity component along the vertical direction is small compared to the horizontal ones, so lubrication approximation applies [96] and the horizontal velocity can be written in the form

$$\mathbf{v}_{\parallel}(x, y, z, t) = q(z)\mathbf{u}(x, y, t), \tag{2.2}$$

where \mathbf{v}_{\parallel} is the horizontal component of \mathbf{v} , and q(z) describes the dependence of the horizontal velocity on z. The no-slip boundary condition at the bottom of the dielectric layer is imposed by setting q(0) = 0, and the stress-free boundary condition at the surface of the electrolyte layer (z = h) requires q'(h) = 0, where $h = h_c + h_d$. Furthermore, a normalization condition

$$q(h_c + h_d) = 1, (2.3)$$

is imposed to make the factorization unique, so $\mathbf{u}(x, y, t)$ can be interpreted as the velocity of the free surface of the electrolyte layer.

Assuming that the vertical component of the velocity is negligible and integrating (1.1) along the *z*-direction, a 2D evolution equation for u has been derived in [71]

$$\partial_t \mathbf{u} + \beta \left(\mathbf{u} \cdot \nabla_{\parallel} \right) \mathbf{u} = -\nabla_{\parallel} \bar{p} + \nu \nabla_{\parallel}^2 \mathbf{u} - \alpha \mathbf{u} + \bar{\mathbf{f}}, \qquad (2.4)$$
where $\nabla_{\parallel} = \hat{x}\partial_x + \hat{y}\partial_y$, \bar{p} is the 2D analogue of pressure, \bar{f} is the depth-averaged Lorentz force, and the constant coefficients β , ν , and α will be defined later. This equation is formally equivalent to the 2D Navier-Stokes equation, where ν is the effective kinematic viscosity, with two modifications: (i) the coefficient β of the advection term is not equal to unity and (ii) there is a new term $-\alpha \mathbf{u}$, known as the Rayleigh friction, which describes the effect of the bottom boundary. The 2D velocity \mathbf{u} is divergence-free

$$\nabla_{\parallel} \cdot \mathbf{u} = 0 \tag{2.5}$$

under the assumption that the vertical component of the velocity, $v_z = 0$. Hence we call this an "incompressible model." This model has been compared with experiment and its predictions were found to be in good agreement with experimental observations for steady flows at $Re \leq 17$, where the Reynolds number (1.3) has been defined using an appropriate characteristic velocity v_s (to be discussed later in more detail), effective kinematic viscosity ν , and the width w of the bar magnets as the characteristic length scale ℓ [72].

2.2.2 Weakly compressible model

In order to extend the validity of the model to higher Re when the flow becomes unsteady, the vertical component of the velocity is nonnegligible, and the variation in the thickness of the fluid layers has to be taken into account, a new model is derived here. When $v_z \neq 0$, (2.2) generalizes to

$$\mathbf{v}(x,y,z,t) = q(z)\mathbf{u}(x,y,t) + v_z(x,y,z,t)\hat{z}$$
(2.6)

Substitution of (2.6) into (1.1) gives

$$\rho q \partial_t \mathbf{u} + \rho q^2 \left(\mathbf{u} \cdot \nabla_{\parallel} \right) \mathbf{u} + \rho v_z q' \mathbf{u} = -\nabla_{\parallel} p + \mu q \nabla_{\parallel}^2 \mathbf{u} + \mu q'' \mathbf{u} + J B_z \hat{x},$$
(2.7)

$$\rho \partial_t v_z + \rho q \left(\mathbf{u} \cdot \nabla_{\parallel} \right) v_z + \rho v_z \partial_z v_z = \mu \partial_z^2 v_z - \rho g - J B_x - \partial_z p, \tag{2.8}$$

Integrating (2.7) over the z coordinate from the bottom of the fluid layer (z = 0) to the free surface (z = h), we obtain

$$\partial_t \mathbf{u} + \beta \left(\mathbf{u} \cdot \nabla_{\parallel} \right) \mathbf{u} = \mathbf{a} + \nu \nabla_{\parallel}^2 \mathbf{u} - (\alpha + \theta) \mathbf{u} + \bar{\mathbf{f}},$$
(2.9)

where

$$\beta = \varrho^{-1} \int_0^h \rho q^2 dz,$$
$$\nu = \varrho^{-1} \int_0^h \mu q dz,$$
$$\alpha = \varrho^{-1} \mu q'(0)$$

are the parameters common to both models and

$$\varrho = \int_0^h \rho q dz. \tag{2.10}$$

The forcing term F represents the horizontal component of the depth-averaged Lorentz force

$$\bar{\mathbf{f}} = \varrho^{-1} J \hat{x} \int_0^h B_z dz \tag{2.11}$$

and

$$\mathbf{a} = -\varrho^{-1} \int_0^h \nabla_{\parallel} p dz. \tag{2.12}$$

Equation (2.9) has the same form as (2.4), with one exception: the Rayleigh friction coefficient has an additional contribution

$$\theta = \varrho^{-1} \int_0^h \rho v_z \partial_z q dz \tag{2.13}$$

which is a function of the velocity. It describes advective transport of momentum as opposed to the diffusive (viscous) transport represented by the coefficient α . The vertical component of velocity v_z can be computed by substituting the ansatz (2.6) into the incompressibility condition (1.2) and integrating the result along the vertical direction from the

bottom surface to arbitrary height z. This gives

$$v_z = -\chi(z)\nabla_{\parallel} \cdot \mathbf{u}, \qquad (2.14)$$

where

$$\chi(z) = \int_0^z q(z')dz'$$
 (2.15)

Substituting (2.14) into (2.13) gives

$$\theta = \gamma \nabla_{\parallel} \cdot \mathbf{u}, \tag{2.16}$$

where

$$\gamma = -\varrho^{-1} \int_{0}^{h} \rho \chi \partial_{z} q dz$$

$$= -\varrho^{-1} \left[\int_{0}^{h} \rho \partial_{z}(\chi q) dz - \int_{0}^{h} \rho q \partial_{z} \chi dz \right]$$

$$= -\varrho^{-1} \left[\rho_{d} \int_{0}^{h_{d}} \partial_{z}(\chi q) dz + \rho_{c} \int_{h_{d}}^{h} \partial_{z}(\chi q) dz - \int_{0}^{h} \rho q^{2} dz \right]$$

$$= \beta - \varrho^{-1} \left[(\rho_{d} - \rho_{c}) \chi_{d} q(h_{d}) + \rho_{c}(\chi_{c} + \chi_{d}) \right].$$
(2.17)

Here we have defined the constants

$$\chi_{d} = \int_{0}^{h_{d}} q(z')dz',$$

$$\chi_{c} = \int_{h_{d}}^{h} q(z')dz'.$$
(2.18)

Equation (2.14) shows that the vertical component of the velocity is noticeably smaller than the horizontal components,

$$v_z = -\chi(z)\nabla_{\parallel} \cdot \mathbf{u} \le \mathcal{O}(\epsilon)v_s, \tag{2.19}$$

where $\epsilon = h_s/w$ and $h_s = \max(h_c, h_l)$ is the characteristic thickness of the liquid layers. In the present experimental setup $\epsilon \approx 0.24$ but, for the range of Reynolds numbers considered here ($Re \leq 25$), the typical magnitude of v_z is even smaller (around 1% of v_s). The terms on the left-hand-side of (2.8) are of order $\epsilon \rho v_s^2/w$ or less and become negligible compared to ρg when $Re \ll \sqrt{gw^4 h_s}/(h_s \nu) \approx 2827$. Similarly, the term $\mu \partial_z^2 v_z$ on the right-handside is of order $\epsilon \rho \nu v_s/h_s^2$ and is negligible compared with ρg when $Re \ll gh_s w^2/\nu^2 \approx 10^5$. Finally, the magnitude of the term JB_x on the right-hand-side is less than the magnitude of the term JB_z which appears on the right-hand-side of (2.7). For inertia-dominated flows the Lorentz force is balanced by the fluid inertia, i.e., JB_z is of order $\rho v_s^2/w$ and hence the term JB_x in (2.8) is negligible compared with ρg for $Re \ll \sqrt{gw^3}/\nu \approx 1375$. Dropping these terms yields a dramatically simplified equation

$$\partial_z p = -\rho g. \tag{2.20}$$

The pressure can be obtained directly by integrating this equation from the top of the electrolyte layer z = h, where p is equal to the constant atmospheric pressure p_a , to a given depth z. The local thickness of the two layers can be obtained using the relation (2.14), which gives for each of the two layers

$$\partial_t h_c = v_z(h) - v_z(h_d) = -\chi_c \nabla_{\parallel} \cdot \mathbf{u}, \qquad (2.21)$$

$$\partial_t h_d = v_z(h_d) = -\chi_d \nabla_{\parallel} \cdot \mathbf{u}, \qquad (2.22)$$

If we define the mean thickness of the two layers as \bar{h}_c and \bar{h}_d , then $h_c = \bar{h}_c + \chi_c \phi$ and $h_d = \bar{h}_d + \chi_d \phi$, where ϕ is the non-dimensional variation in thickness which satisfies the differential equation

$$\partial_t \phi = -\nabla_{\parallel} \cdot \mathbf{u}. \tag{2.23}$$

Generally $\partial_t \phi \neq 0$, so the horizontal flow u becomes weakly compressible.

Integrating the equation (2.20), we find for the electrolyte layer:

$$p = \rho_c g(\bar{h}_c + \chi_c \phi + \bar{h}_d + \chi_d \phi - z) - \sigma_c(\chi_c + \chi_d) \nabla_{\parallel}^2 \phi$$
(2.24)

and for the dielectric layer:

$$p = \rho_c g(\bar{h}_c + \chi_c \phi) + \rho_d g(\bar{h}_d + \chi_d \phi - z) - \sigma_c (\chi_c + \chi_d) \nabla_{\parallel}^2 \phi - \sigma_d \chi_d \nabla_{\parallel}^2 \phi, \qquad (2.25)$$

where σ_c and σ_d are the surface tensions at the electrolyte-air and dielectric-electrolyte interface. Substituting (2.24) and (2.25) into (2.12), we find

$$\mathbf{a} = -\kappa \nabla_{\parallel} \phi + \kappa' \nabla_{\parallel} \nabla_{\parallel}^2 \phi, \qquad (2.26)$$

where

$$\kappa' = \varrho^{-1} [(\chi_c + \chi_d) h \sigma_c + \chi_d h_d \sigma_d],$$

$$\kappa = \varrho^{-1} g [\chi_c \rho_c h + \chi_d (\rho_c h_c + \rho_d h_d)].$$
(2.27)

The thicknesses h_c and h_d depend on x, y, and t. This means that all the parameters (i.e., β , ν , α , γ , κ , κ') also vary with space and time. However, the thickness variation of the two layers in the experiment is very small (a few μ m), compared to their mean values $\bar{h}_c \approx \bar{h}_d \approx 3$ mm, so we can use these constant mean values to compute all the parameters. In addition, the second term in equation (2.26) is small compared to the first one

$$\frac{\kappa' \nabla_{\parallel} \nabla_{\parallel}^2 \phi}{\kappa \nabla_{\parallel} \phi} = O(Bo^{-1}) \ll 1,$$
(2.28)

since the Bond number

$$Bo = \frac{\rho_s g w^2}{\sigma_s} \approx 40, \tag{2.29}$$

defined using $\rho_s = \max(\rho_c, \rho_d)$ and $\sigma_s = \max(\sigma_c, \sigma_d)$ is large. With these simplifications we finally arrive at the compressible 2D model of the flow

$$\partial_t \mathbf{u} + \beta \left(\mathbf{u} \cdot \nabla_{\parallel} \right) \mathbf{u} + \gamma (\nabla_{\parallel} \cdot \mathbf{u}) \mathbf{u} = -\nabla_{\parallel} p + \nu \nabla_{\parallel}^2 \mathbf{u} - \alpha \mathbf{u} + \bar{\mathbf{f}}, \qquad (2.30)$$

$$\partial_t p = -\kappa \nabla_{\parallel} \cdot \mathbf{u}. \tag{2.31}$$

where we defined the effective (2D) pressure $p = \kappa \phi$. This pair of coupled equations have to be solved subject to no-slip boundary conditions u at the lateral walls $x = \pm L_x/2$ and $y = \pm L_y/2$. The effective pressure p satisfies Neumann boundary conditions as the contact angle between the electrolyte surface and the container walls is nearly 0°.

The magnetic field B(x, y, z) required to evaluate the forcing term $\overline{\mathbf{f}}$ is tedious to measure in the experiment. Therefore, we will follow Ref. [72] modeling each magnet as a collection of magnetic dipoles pointing in either $+\hat{z}$ or $-\hat{z}$ direction and summing their magnetic fields to compute the net field B(x, y, z), which yields a good approximation of the experimental measurements [72]. Substituting the resulting magnetic field and current density into (2.11) we find

$$\bar{\mathbf{f}} = \frac{I}{\varrho h_c L_x} \int_{h_d}^h B(x, y, z) dz \, \hat{x} = I a_0 \bar{f}_0(x, y) \, \hat{x}, \tag{2.32}$$

where a_0 is a normalization constant chosen such that the maximum value of the nondimensional forcing profile $\bar{f}_0(x, y)$ is unity. This profile is shown in Figure 2.2. The magnet array used in the experiment corresponds to $a_0 = 1.1 \times 10^{-4}$ m/(amp s²).

In order to determine the parameters ρ , α , β , γ , ν , and κ , the vertical profile q(z) of the horizontal velocity has to be computed first. This profile is weakly dependent on the hor-



Figure 2.2: The x-component of normalized forcing profile, $\bar{f}_0(x, y)$, used in the simulations, as a function of the nondimensional coordinates. The y-component of the forcing is identically zero.

izontal velocity profile and can be calculated with reasonable accuracy by considering the flow in a laterally unbounded domain driven by a sinusoidal forcing profile $\bar{f}_0 = \sin(ky)\hat{x}$, where $k = \pi/w$ [71]. For sufficiently weak driving, the resulting flow is steady and has the same form, $\mathbf{u} = u_0 \sin(ky)\hat{x}$, and the vertical profile q(z) can be computed from the full 3D Navier-Stokes equation (1.1). The explicit expression for q(z) is not particularly illuminating and is available in Ref. [71]. The corresponding parameters are $\alpha = 0.64 \text{ s}^{-1}$, $\beta = 0.83$, $\gamma = -0.16$, $\nu = 3.26 \times 10^{-6} \text{ m}^2/\text{s}$, and $\kappa = 0.0549 \text{ m}^2/\text{s}^2$.

2.2.3 Nondimensionalization

The dynamical system (\mathbf{u}, p) governed by the evolution equations (2.30) and (2.31) have been nondimensionalized for numerical simulations. Choosing a length scale w, velocity scale v_s , time scale $\tau = w/v_s$ and nondimensionalizing the spatial coordinates, time, velocity field \mathbf{u} , and pressure p yields the following evolution equations

$$\partial_{t}\mathbf{u} + \beta \left(\mathbf{u} \cdot \nabla_{\parallel}\right)\mathbf{u} + \gamma (\nabla_{\parallel} \cdot \mathbf{u})\mathbf{u} = - \nabla_{\parallel}p + Re^{-1} \left(\nabla_{\parallel}^{2}\mathbf{u} - \tilde{\alpha}\mathbf{u}\right) + \frac{Ia_{0}w}{v_{s}^{2}}\bar{f}_{0}\hat{x}, \qquad (2.33)$$

$$\partial_t p = -\tilde{\kappa} \nabla_{\parallel} \cdot \mathbf{u}. \tag{2.34}$$

where $\tilde{\kappa} = \kappa/v_s^2$ and $\tilde{\alpha} = \alpha w^2/\nu = 3.14$ describes the relative magnitude of the Rayleigh friction term $\alpha \mathbf{u}$ and viscous term $\nabla_{\parallel}^2 \mathbf{u}$. The free velocity scale

$$v_s = \sqrt{Ia_0w} \tag{2.35}$$

can be chosen such that the coefficient of the forcing profile \bar{f}_0 in (2.33) is eliminated (it corresponds to the balance between the external forcing and the inertia of the fluid mentioned previously). With these scales, the nondimensional evolution equations become

$$\partial_t \mathbf{u} + \beta \left(\mathbf{u} \cdot \nabla_{\parallel} \right) \mathbf{u} + \gamma (\nabla_{\parallel} \cdot \mathbf{u}) \mathbf{u} = -\nabla_{\parallel} p + \tilde{R} e^{-1} \left(\nabla_{\parallel}^2 \mathbf{u} - \tilde{\alpha} \mathbf{u} \right) + \bar{f}_0 \hat{x}, \qquad (2.36)$$

$$\partial_t p = -\tilde{\kappa} \nabla_{\parallel} \cdot \mathbf{u}, \tag{2.37}$$

where \tilde{Re} denotes the Reynolds number computed using the velocity scale (2.35).

The driving current I ranges from 7mA to 25mA in this work, which corresponds to \tilde{Re} varying between 12.1 and 22.8 and $\tilde{\kappa}$ varying between 5.7×10^3 and 1.6×10^3 , which means that this system of PDEs is stiff and care is required to solve it numerically. In the remainder of the thesis, all the dimensional quantities are specified using their units. Otherwise, all the quantities are assumed to be nondimensional.

2.2.4 The models and experiment

The compressible model (2.36)-(2.37) is more accurate than its incompressible counterpart (2.4)-(2.5). However the difference between the solutions is fairly subtle. Since $\tilde{\kappa}$ is large,

the pressure evolves at a much faster rate than the velocity fields. If the initial velocity field is not close to being divergence-free, the pressure rapidly evolves to suppress the divergence, quickly making the flow almost incompressible. In particular, the equilibria (or) fixed points of the compressible model defined by $(\partial_t \mathbf{u}, \partial_t p) = (\mathbf{0}, 0)$ lie in divergencefree subspace, and so are also the equilibria of the incompressible model.

Quick relaxation of the flow towards a divergence-free subspace is an attractive feature of the compressible model. Consider, for instance, initialization of the simulations using experimentally accessible data. The system (2.36)-(2.37) requires an initial condition for velocity field u and an initial condition for the pressure p. While the former can be easily obtained using particle image velocimetry [72]. However, it is extremely difficult to measure the thickness of the two layers with sub-micron precision necessary to reconstruct the pressure field p. On the other hand, quick relaxation of the pressure means that inaccuracy in its initialization plays an almost negligible role on the evolution of the flow field, allowing one to set, e.g., p = 0 as an initial condition.

An alternative procedure can be applied to initialize both the compressible and the incompressible model. The measured flow field u_{exp} can be decomposed into divergence-free and curl-free components using the Helmoltz decomposition

$$\mathbf{u}_{exp} = \mathbf{u}_{exp}^* - \nabla_{\parallel} p, \qquad (2.38)$$

where $\nabla_{\parallel} \cdot \mathbf{u}_{exp}^* = 0$. Taking the gradient of both sides yields a Poisson equation for p

$$\nabla_{\parallel}^2 p = -\nabla_{\parallel} \cdot \mathbf{u}_{exp} \tag{2.39}$$

whose solution gives the initial condition for the pressure, while $\mathbf{u}_{exp}^* = \mathbf{u}_{exp} + \nabla_{\parallel} p$ gives a divergence-free initial condition for the horizontal velocity. In constructing the magnetic field profile (2.32) using dipole summation we assume that the magnet array is symmetric with respect to the origin, with polarity being an odd function of y. This means that if the frame of reference is rotated by π about the vertical z axis, the forcing function merely changes sign

$$\mathcal{R}\bar{f}_0(x,y) = \bar{f}_0(-x,-y) = -\bar{f}_0(x,y), \qquad (2.40)$$

where \mathcal{R} is the corresponding rotation (or inversion) operator.

The same operation also changes the sign of all the odd spatial derivatives. Writing the evolution equations (2.36) and (2.37) in the rotated frame of reference and using the symmetry of the forcing function, we find that the pair $(\mathbf{u}_{\mathcal{R}}, p_{\mathcal{R}})$, where

$$\mathbf{u}_{\mathcal{R}} = -\mathcal{R}\mathbf{u} \tag{2.41}$$

$$p_{\mathcal{R}} = \mathcal{R}p \tag{2.42}$$

is also a solution of the evolution equations [72].

In general, (\mathbf{u}, p) and its symmetric copy, $(\mathbf{u}_{\mathcal{R}}, p_{\mathcal{R}})$, need not be the same, but if an initial condition is chosen to lie in the symmetric subspace $(\mathbf{u}_{\mathcal{R}}, p_{\mathcal{R}}) = (\mathbf{u}, p)$, the solution remains symmetric under inversion at all times.

2.2.6 State space

As described in the previous chapter, the evolution of the flow can be visualized as motion of a point in the state space. So, a flowfield in the physical space corresponds to a point in the state space and *vice versa*. By discretizing the physical domain, the state can be written as a column with each entry representing the value of the pressure field or component of velocity field at a particular spatial location, i.e., the state space vector X can be written as

$$\mathbf{X} = \begin{bmatrix} u_{11}, \cdots, u_{N_x N_y}, v_{11}, \cdots, v_{N_x N_y}, p_{11}, \cdots, p_{N_x N_y} \end{bmatrix}^T$$
(2.43)

where p_{ij} , u_{ij} , v_{ij} respectively are the pressure, and x and y components of \mathbf{u} at the grid location (i, j), where the flow domain is discretized using a grid with $N_x \times N_y$ points. So, the state space vector has a dimension of $N = 3N_xN_y$ and hence the state space is a subset of \mathbb{R}^N . Note that the choice for the order of the elements in defining the state space is a mere convenience, and need not necessarily be the way described above. Using equation (2.43), the evolution equations (2.36) and (2.37) can be written in compact form as

$$\dot{\mathbf{X}} = \partial_t \mathbf{X} = \mathbf{V}(\mathbf{X}), \tag{2.44}$$

where the vector function V(X) represents the spatial discretization discussed in more detail in Appendix A.

Writing (2.43) in compact form as

$$\mathbf{X} = \begin{bmatrix} u(x,y) \\ v(x,y) \\ p(x,y) \end{bmatrix},$$
(2.45)

the action of the symmetry operator on the flow state can be written as

$$\mathcal{R}\mathbf{X} = \begin{bmatrix} -u(-x, -y) \\ -v(-x, -y) \\ p(-x, -y) \end{bmatrix}.$$
(2.46)

Since $\mathcal{R}^2 = \mathbb{1}$, together the identity operator $\mathbb{1}$ and \mathcal{R} form a group $\mathcal{G} = \{\mathbb{1}, \mathcal{R}\}$.

The system also has an approximate discrete symmetry. Indeed, if the magnet array

were infinitely extended in the y direction, the evolution equations would not change under a mirror reflection in the x direction followed by the shift by w in the y direction:

-

$$\mathcal{P}\mathbf{X} = \begin{bmatrix} -u(-x, y+w) \\ v(-x, y+w) \\ p(-x, y+w) \end{bmatrix}.$$
(2.47)

The finite extent of the magnet array and the side walls of the container weakly break this symmetry, however the presence of this weakly broken symmetry has a major effect on the structure of solutions, as shown in Ref. [72] and discussed in more detail in Chapter 3. Unlike \mathcal{R} , the square of \mathcal{P} is not an identity operator, rather it represents another approximate discrete symmetry $\mathcal{T} = \mathcal{P}^2$ – the translation by 2w in the y direction:

$$\mathcal{T}\mathbf{X} = \begin{bmatrix} u(x, y + 2w) \\ v(x, y + 2w) \\ p(x, y + 2w) \end{bmatrix}.$$
(2.48)

Together the symmetries \mathcal{T}^n , $\mathcal{T}^n \mathcal{R}$, and $\mathcal{T}^n \mathcal{P}$, where *n* is any (small) integer make up an approximate symmetry group \mathcal{G}' of the flow over a finite-size magnet array.

2.3 Summary

In this Chapter, a 2D model has been developed for the flow inside a thin fluid layer driven by a nearly sinusoidal in-plane forcing profile. The model has been derived from first principles by depth-averaging the 3D Navier-Stokes equation. By taking the thickness variation of the fluid layers in to account, the model is expected to show improved agreement with the experiment at higher *Re*, where the flow becomes turbulent. This model is validated by comparing its predictions with the experimental observations over a range of Reynolds numbers in the next Chapter.

CHAPTER 3

TRANSITION FROM LAMINAR TO TURBULENT FLOW

Although the main focus of this Chapter is on the quality of predictions made by the 2D model of the flow in the turbulent regime, direct tests of the model against experimental observations are complicated by the chaotic nature of turbulent flow. For instance, even the tiniest error in defining the initial conditions in the model will result in its predictions diverging exponentially fast from the experimental observation. The comparison therefore has to be performed at lower Re, when the flow in not chaotic. Alternatively, in the turbulent regime at higher Re one should compare metrics that are not affected by the chaotic nature of the dynamics. Therefore, we will mainly focus on the comparison of various bifurcations at lower Re and temporal recurrences at higher Re.

The 2D compressible model of the flow is as complicated mathematically as the original Navier-Stokes equation coupled with the incompressibility condition, and so can only be studied numerically at *Re* of interest. A numerical representation of the model was obtained by finite-difference spatial discretization of the physical domain on a staggered grid and using velocity-pressure formulation (see Appendix A for details), following the previous analysis carried out using the incompressible 2D model [72]. Unlike that previous study, which focused on the bifurcations of steady flows, the present analysis explores the entire sequence of bifurcations leading to turbulence and even considers the periodic window embedded in the turbulent regime.

The 2D models of the Kolmogorov-like flow offer a dramatic simplification compared to the full description of the problem based on a pair of 3D Navier-Stokes equations describing each of the fluid layers, both of which involve free surfaces. However, there is a price: both 2D models of the Kolmogorov-like flow involve a set of parameters, none of which can be measured directly, but rather have to be computed with the help of the ansatz (2.6), the accuracy of which can only be assessed indirectly. Unlike the rest of the parameters, the Reynolds number can be obtained using the velocity field u at the top of the electrolyte layer, which can be directly measured in experiment or computed in the numerics. Specifically, we will define the characteristic velocity that enters the definition (1.3) of the Reynolds number as

$$v_s = \frac{1}{T_m} \int_0^{T_m} \sqrt{\langle \mathbf{u} \cdot \mathbf{u} \rangle} \, dt, \qquad (3.1)$$

where T_m is a sufficiently long observation time and $\langle \cdot \rangle$ denotes the spatial average in the central square region of size $4w \times 4w$, with w being the magnet width. The central region has been used as opposed to the full domain because the experimental data has larger errors near the edges of the domain.

Although the corresponding Reynolds number Re enables unambiguous comparison between the experiment and simulations, this definition has a notable drawback. Unlike the characteristic velocity (2.35), the one defined using (3.1) depends on the flow and may not be unique, if several different attractors coexist at a given current. This means that Re is not a one-to-one function of the current I, which may result in an ambiguity. Since there does not appear to be a satisfactory way to address this issue, we will refer to both the Reynolds number Re and the corresponding current I in comparison of the simulations and experiment in this Chapter. A comparison of the Reynolds numbers defined using the two definitions of v_s is presented in Figure 3.1.

3.1 **Pre-turbulent flow**

3.1.1 Straight flow

The spatial pattern of a steady flow, which corresponds to an equilibrium in the state space, can be understood qualitatively using the dominant balance in the governing equation (2.36). In particular, for weak driving (low Re) the dominant terms are $\bar{f}_0\hat{x}$ (Lorentz



Figure 3.1: The dependence of Re (red) and \tilde{Re} (blue) on the current *I*. Panel (b) shows a zoomed-in version of the circled region in panel (a).

force) and $Re^{-1}(\nabla_{\parallel}^{2}\mathbf{u} - \tilde{\alpha}\mathbf{u})$ (dissipation due to viscosity). Balancing these terms gives a flow which resembles the forcing profile. In the case of Kolmogorov flow (for which $\bar{f}_{0} = \sin(ky)$ with $k = \pi/w$), the forcing has only the *x* component. Consequently the flow is in the *x* direction and its *y* component is negligible on a domain with periodic boundary conditions [61, 71]

$$\mathbf{u} = \frac{Re}{k^2\nu + \tilde{\alpha}}\sin(ky)\hat{x}.$$
(3.2)

For the Kolmogorov-like flow in a bounded domain considered here, the flow in the middle of the domain closely resembles the analytical solution (3.2), but slows down considerably in the buffer region between the edge of the magnet array and the lateral walls. This flow state is the analogue of laminar flow, and will be referred to as "straight flow", denoted as E_1 . The "edge effects" due to both the no-slip boundary conditions at the lateral walls and the decrease in the strength of the magnetic field outside of the magnet array are responsible for the emergence of a nonzero *y* component of the flow which is mostly constrained to the buffer region. The two velocity components of the straight flow in experiment are shown



Figure 3.2: (a) x-component and (b) y-component of velocity field of the straight flow in experiment at Re = 8.06. The driving current used in the experiment is 7 mA.

in Figure 3.2.

At low values of the driving current I, the steady straight flow state is stable and is a global attractor. Hence it can be computed by performing a long time-integration of the system starting from an arbitrary initial condition. To speed up convergence, we instead computed most stable (and all unstable) solutions using numerical continuation in I (described in Appendix B) which relies on a Newton-based solver (described in Chapter 4). The solution of the model is compared with experiment in Figure 3.3, which shows the vorticity fields $\omega = (\nabla_{\parallel} \times \mathbf{u}) \cdot \hat{z}$. Vorticity is a scalar field and offers a more convenient representation of the flow compared with the vector field \mathbf{u} . As the Figure illustrates, the model reproduces all features of the experimental flow field, such as the slight tilt of the shear band due to the global recirculation, the very accurately, justifying the accuracy of the model. The sections x = 0 and y = 0 of the vorticity field shown in Figure 3.3 provide a more quantitative comparison of the experimental measurements and numerical results. Again we find good agreement (the slight discrepancy in Figure 3.3 for $-5 \leq x \leq 5$ is

mainly due to the insufficient spatial resolution in the experiment).

3.1.2 Primary instability

As the driving current is increased, the straight flow in the experiment undergoes a qualitative change at $I_1 = 9.6$ mA ($Re_1 = 11.1$), with the straight shear bands developing spatial modulation in the transverse (y) direction that eventually gives rise to a pattern of distinct stationary vortices (cf. Figure 3.4). We will denote this steady modulated flow state as E_2 below. As Ref. [72] demonstrated, the transition from E_1 to E_2 in the presence of lateral boundaries is not a result of an instability. Indeed, both states correspond to the same solution branch, and the leading stability exponent associated with this branch becomes close to marginal near I_1 , but never vanishes, as Figure 3.5a shows. Instead, the transition from E_1 to E_2 is associated with an *imperfect* pitchfork bifurcation, as described in Ref. [72] for the incompressible model. Since the equilibria of both the compressible and the incompressible models are the same, the exact same bifurcation occurs in both the models at $I_1 \approx 8.98$ mA ($Re_1 = 10.81$).

Two new solution branches – \mathbf{E}_3 and \mathbf{E}_4 – are created in a saddle-node bifurcation at I_1 as Figure 3.5b illustrates. They both describe equilibria with the corresponding flow fields shown in Figure 3.6. The saddle \mathbf{E}_3 represents the "extension" of the straight flow \mathbf{E}_1 on the subcritical side $I < I_1$ to the supercritical side $I > I_1$, and has been computed using the parametric continuation in I as described in appendix B. The node \mathbf{E}_4 represents the symmetry-related analogue of the modulated flow \mathbf{E}_2 (see below) . Ignoring the imperfect nature of the pitchfork bifurcation, at I slightly greater than I_1 , we can write $\mathbf{E}_2 \approx \mathbf{E}_3 + p\mathbf{e}_1$ and $\mathbf{E}_4 \approx \mathbf{E}_3 - p\mathbf{e}_1$, where \mathbf{e}_1 is the marginal direction of \mathbf{E}_3 associated with this bifurcation and p is a function of $I - I_1$. Hence, once \mathbf{E}_3 is computed for a particular $I > I_1$, \mathbf{E}_4 can be found using Newton's method with the initial condition $\mathbf{E}_4 \approx 2\mathbf{E}_3 - \mathbf{E}_2$.

As discussed in Ref. [72], the fact that the pitchfork bifurcation is imperfect is due to the weak symmetry breaking by the lateral boundaries and a finite size of the magnet array.



Figure 3.3: (a) Vorticity field of the straight flow in experiment at Re = 9.25 (I = 8 mA) and (b) in simulation at Re = 9.26 (I = 7.65 mA). Comparison of the vorticity fields on the lines (c) x = 0 and (d) y = 0 between experiment (red) and simulation (blue).



Figure 3.4: (a) Vorticity field of the modulated flow in experiment at Re = 14.01 and (b) in simulation at Re = 14.03. The driving currents used are, respectively, I = 14 mA and I = 13.89 mA.

Recall that, unlike the rotation symmetry \mathcal{R} , the discrete symmetry \mathcal{P} which corresponds to a mirror reflection in the *x* direction followed by the shift in the *y* direction by *w* is not exact. A *perfect* pitchfork bifurcation represents breaking of an *exact* symmetry. An example is shown in Figure 3.5b, where the flow state \mathbf{E}_3 which is symmetric under \mathcal{R} undergoes a pitchfork bifurcation at I = 9.04 mA, producing two new solution branches \mathbf{E}_5 and \mathbf{E}_6 (shown in red) that are nearly symmetric under \mathcal{P} and are mapped into each other by \mathcal{R} (these branches were hypothesized, but never computed in Ref. [72]). The straight flow solution \mathbf{E}_1 which is symmetric under \mathcal{R} and *nearly* symmetric under \mathcal{P} , undergoes an *imperfect* pitchfork bifurcation which *completely* breaks \mathcal{P} , but preserves \mathcal{R} . As a result, all three solution branches \mathbf{E}_2 , \mathbf{E}_3 , and \mathbf{E}_4 remain symmetric under \mathcal{R} , while \mathbf{E}_2 and \mathbf{E}_4 are approximately mapped into each other by \mathcal{P} .

Figure 3.5b has been generated by projecting the solutions described above from the high-dimensional state space onto the two dominant eigenvalues associated with the pitch-



Figure 3.5: (a) Leading stability exponents of the 3 equilibria near the imperfect pitchfork bifurcation. The branch containing E_1 and E_2 is plotted in magenta while E_3 , E_4 are plotted respectively in blue and black. E_3 undergoes pitchfork bifurcation (shown as cyan sphere) by breaking the rotational symmetry \mathcal{R} . (b) Projections of these 3 equilibria along with E_5 and E_6 (in red), plotted as the driving current varies. The saddle node bifurcation of E_3 and E_4 is shown as brown sphere.



Figure 3.6: Vorticity fields (from simulations) of the 3 equilibria present in the imperfect pitchfork bifurcation - (a) \mathbf{E}_2 (b) \mathbf{E}_3 and (c) \mathbf{E}_4 at I = 10.56 mA. This current, far away from the bifurcation, is chosen as the flow fields look indistinguishable to the naked eye close to the bifurcation.

fork bifurcation at I = 9.04 mA. Their spatial structure is shown in Figure 3.7. The eigenvector e_2 defines the marginal direction at the bifurcation (it represents the deviation of E_5 and E_6 from E_3 , while the eigenvector e_1 defines the unstable direction. The latter is nearly indistinguishable from the eigenvector which defines the marginal direction at the saddle-node bifurcation and corresponds to the deviation of E_2 and E_4 from E_3 . The symmetry of the eigenvectors reflects the nature of the two bifurcations and the respective symmetry breaking: e_1 is symmetric under \mathcal{R} and approximately antisymmetric under \mathcal{P} , while e_2 is antisymmetric under \mathcal{R} and approximately symmetric under \mathcal{P} . These results verify and refine the description of the respective bifurcations proposed in Ref. [72].

The four new equilibria, \mathbf{E}_3 , \mathbf{E}_4 , \mathbf{E}_5 , \mathbf{E}_6 , being either unstable or weakly stable, are not observed in experiment (or simulations starting from arbitrary initial conditions). Instead, as *I* increases, the flow follows the continuous solution branch, which smoothly transitions from the straight flow \mathbf{E}_1 to the modulated flow \mathbf{E}_2 at $Re \approx Re_1$. This transition can be quantified using the magnitude of the transverse component of the velocity, which plays the role of an order parameter. Defined as the normalized spatial mean square of the *y* component of velocity in the central square region of the flow domain with dimensions



Figure 3.7: Vorticity fields which correspond to eigenvectors (a) e_1 and (b) e_2 .

 $4w \times 4w$,

$$\hat{V} = \frac{\langle v^2 \rangle}{\langle u^2 \rangle + \langle v^2 \rangle},\tag{3.3}$$

the order parameter increases monotonically with Re (and I), as shown in Figure 3.8. The Figure also shows that the numerical model accurately captures the critical Re at which the bifurcation occurs in the experiment, with the critical Reynolds number in simulations being within 2.6% of that in the experiment.

Qualitatively, the compressible model reproduces the experimental flow structure extremely well, as Figure 3.4 illustrates. However, comparison of the order parameters shown in Figure 3.8 indicates that there is a quantitative discrepancy. As discussed in Ref. [72], this discrepancy is due to the choice of the model parameters which are computed using the vertical profile q(z) of the velocity which corresponds to the straight flow E_1 . As discussed in Chapter 2, the vertical profile q(z) (and hence all model parameters) is weakly dependent on the horizontal flow profile u which, in turn, depends on *Re*. The agreement between the model predictions and experimental observations can be improved by allowing the model parameters to vary with *Re*. One promising possibility that is currently being explored is



Figure 3.8: Order parameter \hat{V} (cf. Eq. (3.3)) as a function of Re for the modulated flow. The experimental data is plotted in red and the simulation data in blue.

to extract the values of all the parameters from time series of experimental flow fields using a sparse regression method [97].

3.1.3 Secondary instability

The modulated flow remains stable over the range $I_1 < I < I_2$ ($Re_1 < Re < Re_2$), where $I_2 = 19.6$ mA and $Re_2 = 17.6 \pm 0.1$ in experiment. As the current is increased beyond I_2 , the modulated flow looses stability and undergoes a Hopf bifurcation, giving way to a stable time-periodic flow PO₁ with the temporal period $T_1 = 120 \pm 1$ seconds (and frequency $f_1 \approx 0.008$ Hz). Figure 3.10 shows the normalized intensity of the spatially averaged temporal power spectrum

$$P(f) = \frac{\left\langle \left| \sum_{j=0}^{N_0 - 1} e^{-i2\pi f j dT} \mathbf{u} \right|^2 \right\rangle}{N_0^2 \left\langle \mathbf{u} \cdot \mathbf{u} \right\rangle}$$
(3.4)

of the flow as a function of frequency f at different $Re > Re_2$. Here the power spectrum intensity is computed using $N_0 = 1800$ consecutive states uniformly sampled with a time



Figure 3.9: (a) Stability exponents of the steady modulated state at $I_2 = 19.4$ mA ($Re_2 = 17.38$). The two complex conjugate exponents crossing the imaginary axis indicate a Hopf bifurcation. (b) Floquet multipliers of the periodic orbit PO₁ at $I_4 = 20.3$ mA ($Re_4 = 17.96$). Complex conjugate Floquet multipliers crossing the unit circle (red) indicate a Hopf bifurcation. In both panels 25 leading eigenvalues are shown.

step dT = 1 sec. Both the integral of the power spectrum and its maximum

$$P_{\max} = \max_{f} P(f), \tag{3.5}$$

which characterize the amplitude of oscillation, can be used to define an order parameter for this bifurcation. We choose to use the latter.

A similar bifurcation is found in both the models, as illustrated using the stability exponents of the modulated state at I_2 in the compressible model. The period at the onset of time-periodic oscillations is found to be 120.5 seconds in both the models but the critical Re is different. In compressible model, $I_2 = 19.4$ mA or $Re_2 = 17.38$, which is within 2% of the experimental value. This should be compared with the incompressible model, which predicts $Re_2 = 16.97$ (a 3.8% discrepancy with experiment), illustrating the improvement that can be obtained by taking the thickness variation into account. A comparison of the order parameter (3.5) as a function of Re for the experiment and the two models is shown in Figure 3.11, and again we find the compressible model to be substantially more accu-



Figure 3.10: Power spectrum intensity as a function of frequency for various Re in (a) experiment and (b) the compressible model.



Figure 3.11: Order parameter (peak power) at various Re and a square root fit for the data are plotted for experiment (red), the compressible model (blue), and the incompressible model (black).

rate than the incompressible one. The residual discrepancy in the order parameter and/or critical Re can be further reduced by a more precise choice of the model parameters, but this is outside the scope of this thesis. In both models and experiment the amplitude of the oscillation vanishes at I_2 and increases monotonically with I, while the period remains fairly constant. The scaling of the order parameter $P_{\text{max}} \propto \sqrt{I - I_2}$ is consistent with the supercritical Hopf bifurcation. It has to be noted here that the orbit, PO₁, is stable only in a narrow range of Reynolds numbers (and a similarly narrow range of currents I) and so, the experiment should be carefully designed to enable fine adjustments of the current, to investigate this time-periodic state.

3.2 Two-dimensional turbulence

3.2.1 Transition to turbulence

An interesting feature that is observed in the power spectrum (cf. Figure 3.10a) of the experiment is the appearance (at $Re_3 \approx 17.77$) of a second peak with frequency $f_2 \ll f_1$, which grows in magnitude with increased current (or Reynolds number). Interestingly, the power spectrum produced by the model does not show this second frequency for most Re in

the range where the oscillatory state is stable (cf. Figure 3.10b). This puzzling discrepancy was resolved by investigating the mechanism which destabilizes the oscillatory state. The spectrum of Floquet multipliers

$$\Theta_i = e^{\Lambda_i T} \tag{3.6}$$

associated with this solution at $Re_4 = 17.96$ or $I_4 = 20.3$ mA is shown in Figure 3.9b. It shows that the periodic orbit PO₁ becomes unstable by undergoing another Hopf bifurcation, which results in the creation of a new, quasiperiodic solution QP₁. Furthermore, the Floquet spectrum shows that the phase of the marginal complex conjugate pair associated with this second Hopf bifurcation at Re_4 is $\phi \approx \pi/5$ (cf. Figure 3.9b), which corresponds to a period $(2\pi/\phi)T_1 \approx 10T_1 \approx 1200$ seconds or a frequency of roughly 0.0008 Hz. If the Hopf bifurcation at $Re = Re_4$ were subcritical, this could well explain the quasiperiodic dynamics for $Re < Re_4$.

Figure 3.12 shows the power spectra of the solutions of the compressible model obtained by quasistatically increasing the driving current above I_4 and then decreasing it below I_4 . The top row, which corresponds to increasing I shows a single peak corresponding to the stable time-periodic solution for $Re < Re_4$. Above Re_4 we find a broad spectrum characteristic of chaotic dynamics and, indeed, the corresponding flow state becomes weakly turbulent. If subsequently I is decreased, the turbulent flow persists below I_4 , but eventually gives way to a stable quasiperiodic state, as the power spectra in the bottom row of Figure 3.12 illustrate. The associated frequencies of the quasiperiodic state are 0.008 Hz and 0.0008 Hz, consistent with the frequency calculated from the Floquet spectrum of PO₁. It is this quasiperiodic state that is found in the experiment for $Re_3 < Re < Re_4$, as Figure 3.10a illustrates, where $Re_4 = 17.88$ (in the compressible model $Re_3 = 17.84$ and $Re_4 = 17.96$). The experimental data set has a duration of $\Delta T = 1800$ seconds, which limits the frequency resolution to $\Delta f = 2/\Delta T \approx 0.001Hz$, clearly insufficient to accu-



Figure 3.12: The power spectrum P(f) at different I in the compressible model demonstrates the hysteresis near $I_4 = 20.3$ mA. The top panel shows the power spectrum when I is quasistatically increased beyond I_4 , and the bottom panel shows the power spectrum when I is decreased. The magnitude of I (in mA) is shown at the top of each plot. All the plot have the same range, as shown for the plot in the top left corner.

rately determine the second frequency $f_2 \approx 0.0008$ Hz associated with the quasiperiodic state. Instead, we see a peak at a frequency $O(\Delta f)$ which may or may not correspond to f_2 .

The presence of the hysteresis below Re_4 implies that the second Hopf bifurcation is subcritical. The compressible model predicts that in the range of Re where the hysteresis is found, two different attractors are present: the first one is the stable oscillatory state discussed in the previous section and the second one represents a more complicated dynamical regime (quasiperiodic at lower Re, chaotic at higher Re). The basins of attraction of the two stable states split the state space into two separate parts. The boundary between them (or at least some part of it) should correspond to the stable manifold of the unstable quasiperiodic state QP₁ which is created at Re_4 and appears to exists over the entire range $Re_3 < Re < Re_4$. Which of the two stable states is selected depends on the choice of initial conditions, i.e., on the protocol used in the simulations or experiment.

The Poincare section shown in Figure 3.13a shows a stable quasiperiodic solution QP_2 ,

which exists over a portion of the range $Re_3 < Re < Re_4$. It corresponds to a 2-torus in the state space and a loop in the Poincare section. The Poincare section was computed by calculating the intersections of the trajectory with a hyperplane normal to PO₁ at some point along the orbit. Then, these intersections were projected onto a plane spanned by the two unstable complex conjugate Floquet vectors of PO₁. The resulting sections are plotted in Figure 3.13, where the horizontal and vertical axes correspond to the real and imaginary components of the projection.

The quasiperiodic solution QP_2 is created at Re_3 , with both amplitudes being finite, which indicates a saddle-node bifurcation of two quasi-periodic solutions, one stable and one unstable. In all likelihood, the unstable one is the quasiperiodic solution QP_1 created in the secondary Hopf bifurcation of PO₁. The Poincare section shown in Figure 3.13b shows that the stable 2-torus QP_2 becomes unstable and undergoes a bifurcation (or sequence of bifurcations) at some $Re > Re_3$, giving rise to an attracting set with eight components, which are disconnected in the Poincare section. This suggests resonant interaction between the two frequencies of the 2-torus. In the full state space these components are, of course, connected, with each component representing an attractor of the period-8 Poincare map. As Re increases these components grow, overlap, and eventually get transformed into the chaotic set underlying turbulent dynamics at $Re > Re_4$, as Figure 3.13c and Figure 3.13d show.

We have restricted our analysis of the hysteresis to numerical simulations. Even within the numerical model, a detailed study of all the bifurcations and the structure of the resulting chaotic sets is quite expensive and so has not been performed. Experimental studies have their own challenges. The range ($\Delta Re = Re_4 - Re_3 = 0.12$) over which the hysteresis is found is very narrow and adjusting the current with precision necessary to place the experimental system into this range is not easy. Furthermore, the frequency associated with the second Hopf bifurcation at Re_4 corresponds to a rather long period of 1200 seconds (or 20 minutes). Reliably extracting this frequency in experiment would require processing



Figure 3.13: Poincare sections at various driving currents at (a) I = 20.2 mA, (b) I = 20.3 mA, (c) I = 20.31 mA, and (d) I = 20.35 mA showing the transition from quasiperiodicity to turbulence. To illustrate the breakup of the 2-torus, different colors have been used in panels (b) and (c) to label the eight disconnected components of the resulting attractor (e.g., *n*th crossing is color-coded based on the values of $n \mod 8$).

many hours worth of PIV data. Processing the latter is in itself a challenge. Due to these complications hysteresis has not been thoroughly investigated in the experiment.

Although the flow undergoes a number of supercritical bifurcations before transitioning to turbulence, this transition has all the hallmarks of a subcritical transition to turbulence found in numerous canonical three-dimensional flows (e.g., channel flow [98], plane Couette flow [99], pipe flow [100], etc.).

3.2.2 Recurrences

Before discussing the dynamics in the turbulent regime, we introduce recurrence analysis [101], which is a convenient tool that can be used to analyze the temporal aspect of the dynamics in both numerical simulations and experiment. It can be used to quickly identify whether the flow approaches an equilibrium or a time-periodic state, even when such states are unstable. For instance, if the turbulent flow passes near an unstable solution with temporal period T, the flow state will not exactly recur time T later, but rather one will observe a near-recurrence. In order to identify such near-recurrences, we will compute a recurrence diagram obtained by comparing a state at time t with a state at time $t - \tau$. The closeness of the recurrence can be quantified by using the normalized distance

$$d(t,\tau) = \frac{\min_{g \in \mathcal{G}} \left(\|g\mathbf{X}(t) - \mathbf{X}(t-\tau)\| \right)}{\|\mathbf{X}(t)\|}$$
(3.7)

where $\|\cdot\|$ represents the 2-norm and the minimum over the symmetry group \mathcal{G} is taken to account for the rotation symmetry. In this thesis, all the norms used are 2-norms, unless otherwise specified.

Examples of the recurrence diagram for an asymptotically steady and a time-periodic state are shown in Figure 3.14, where the horizontal and vertical axis corresponds to the physical time and time delay, respectively. For example, Figure 3.14a shows that $d(t, \tau) \rightarrow 0$ for all τ as $t \rightarrow \infty$, which implies that the state approaches a stable equi-



Figure 3.14: Sample recurrence plots at various driving currents. (a) At I = 15 mA the flow approaches an equilibrium, which is a global attractor. (b) At I = 20 mA the flow approaches a periodic orbit, which is also a global attractor.

librium. Figure 3.14b, on the other hand, shows that $d(t, \tau)$ approaches a function which is periodic in τ (with the same period T) and which vanishes at $\tau = nT$ with *n*-integer, indicating that the flow approaches a stable time-periodic orbit with period T.

3.2.3 Intermittency

At values of Re just above Re_4 we find a spatiotemporally intermittent flow which alternates between two very different dynamical regimes which we can broadly refer to as laminar and turbulent. This type of intermittency, which is triggered by a subcritical Hopf bifurcation is called "type-2" intermittency, following Pomeau *et al.* [102]. In this particular case, laminar flow corresponds to the oscillatory state discussed previously. This is well illustrated by Figure 3.15 which shows the recurrence plots obtained using numerical solutions and experimental flow fields. The recurrence plots are qualitatively similar, with the depth of the minima being the most noticeable difference attributable to the significantly higher level of noise in the experimental recordings compared with the essentially noiseless numerics. In both numerics and experiment we find extended intervals of time where dynamics are nearly periodic, as indicated by the the minima of the distance function (3.7)



Figure 3.15: Recurrence plots showing intermittency just above Re_4 in (a) simulation at Re = 18.2 and (b) experiment at Re = 18.1.

at $\tau = nT_0$, where T_0 is the period of PO₁ and *n* is an integer. During those intervals, the flow is essentially indistinguishable from that described by PO₁, even though this solution is (weakly) unstable at this *Re*.

Since PO₁ is unstable, the flow will eventually be repelled from this state. Indeed the recurrence diagrams show that laminar intervals are broken up by short turbulent bursts where the dynamics are clearly aperiodic (cf. Figure 3.16a). After some excursion away from PO₁, the trajectories again get reinjected into the vicinity of this state and the process repeats, producing intermittency. The reinjection mechanism is not understood at present. As the driving current is increased, the orbit PO₁ becomes progressively more unstable, which results in laminar periods becoming shorter and, consequently, turbulent bursts occurring more frequently (cf. Figure 3.16b). Eventually, the laminar intervals completely disappear and one finds persistent turbulence, with any trace of periodicity completely disappearing from the temporal dynamics (cf. Figure 3.16c). For *Re* of order a few tens, this two-dimensional turbulence is weak in the sense that there are no dynamics at multiple length scales. Instead dissipation occurs on the same length scale *w* at which the energy

is injected into the flow. Such flows can be equally accurately described as being chaotic, rather than turbulent.

3.2.4 Periodic window

As often happens in other nonlinear systems, a window of periodic dynamics is found within the range of Re where chaotic dynamics are found (cf. Figure 3.16e). This window corresponds to the driving currents in the interval $I_5 < I < I_8$, where $I_5 \approx 21.5$ mA and $I_8 \approx 21.7$ mA, and is associated with the emergence of a stable periodic orbit PO₂. In the simulations, this orbit has a period $T_2 \approx 460$ s, which is essentially independent of the driving current. A periodic window is also observed in the experiment at comparable I, although in that case the period of the orbit is approximately 600 s. As Figure 3.17 illustrates, in both experiment and simulations, following an initial transient, the flow approaches a time-periodic state.

The orbit PO₂ has been computed numerically at I = 21.55 mA using a minimum in the recurrence diagram as an initial condition and then continued both up and down in I to investigate its relevance to the chaotic dynamics on either side of the window. This orbit undergoes a subcritical bifurcation at I_5 and becomes unstable below I_5 , which is manifested in intermittent flow observed near the left edge of the periodic window (cf. Figure 3.16d). Unexpectedly, PO₂ was found to undergo a period-doubling bifurcation at $I_7 = 21.6$ mA, i.e., *inside* the window, with the leading Floquet multiplier crossing the unit circle at -1, as shown in Figure 3.18.

To find the periodic orbit born out of this period-doubling bifurcation, Newton's iterations with an initial condition along the marginal direction of PO₂ have converged to a stable periodic orbit PO₄ with a period of approximately $2T_2$. Continuation of PO₄ in *I* showed that it exists and is stable both below and above I_7 , as illustrated in Figure 3.19, suggesting that the bifurcation of PO₂ at I_7 is subcritical. We also found an unstable periodic orbit PO₃ of period close to $2T_2$ below I_7 , which suggests the presence of a new



Figure 3.16: Recurrence plots in the simulations of compressible model for $I > I_4$. (a) I = 20.4 mA and (b) I = 20.6 mA demonstrate that the laminar periods in the intermittent dynamics become shorter as I increases farther away from I_4 . (c) I = 21.1 mA shows chaotic dynamics. (d) I = 21.4 mA shows intermittency near the left edge of the periodic window. (e) I = 21.5 mA shows dynamics inside the periodic window. (f) I = 21.9 mA and (g) I = 22.6 mA demonstrate that the laminar periods in the intermittent dynamics become shorter as I increases farther away from I_8 . (h) I = 25 mA shows chaotic dynamics.



Figure 3.17: Recurrence plots in the stable island, where the state asymptotically settles in to a periodic orbit. The recurrence plots are computed in (a) simulation at Re = 19.03 and (b) experiment at Re = 20.1

solution branch that bifurcates off of PO₂ at I_7 and disappears in a saddle-node bifurcation with PO₄ at some driving current I_6 , where $I_5 < I_6 < I_7$, as shown in Figure 3.19.

The flow remains time-periodic as the driving current is increased up to I_8 , which suggests that PO₄ remains stable over that range of *I*. As the current is increased above I_8 , the flow becomes intermittent, as shown in Figure 3.16f, suggesting that PO₄ undergoes a subcritical bifurcation at I_8 . The nature of this bifurcation (and the existence of PO₄ above I_8) has not been investigated. Further increase in the driving current leads to the gradual decrease in the fraction of time the flow spends in the time-periodic regime (cf. Figure 3.16g), and eventually dynamics again become fully chaotic (cf. Figure 3.16h).

3.3 Summary

The weakly compressible model was found to be in qualitative and quantitative agreement with experiment in predicting the sequence of bifurcations leading up to turbulence. Furthermore, it qualitatively described the hysteretic subcritical transition between turbulence


Figure 3.18: The dominant 25 Floquet multipliers of the periodic orbit, PO₂, plotted in the complex plane for at $I_7 = 21.6$ mA ($Re_7 = 19.03$). At this driving, the (black) Floquet multipliers cross the unit circle (in red) resulting in a period-doubling bifurcation.



Figure 3.19: A sketch showing periodic solutions PO_2 , PO_3 , and PO_4 and their stability inside the periodic window. Stable (unstable) branches are shown using solid (dashed) lines and the circles indicate solutions computed using Newton iterations. The horizontal axis shows the driving current *I* in mA and the vertical axis corresponds to a direction in the state space along the leading Floquet multiplier of PO_2 at I_7 . The blue portions of the branches are based on continuation, while the red portions are speculated based on PO_3 computed at one value of *I* (red open circle).

and time-periodicity that was observed in experiment. The model also correctly predicted and explained the appearance of various spatiotemporally intermittent flows and a periodic window embedded in the chaotic regime, giving us confidence in using it to investigate quasi-two-dimensional turbulent flows. The remaining qualitative disagreement can likely be resolved by tuning the model parameters, which is outside the scope of this work.

CHAPTER 4

NEWTON-BASED METHODS FOR COMPUTING ECS AND DYNAMICAL CONNECTIONS

This chapter describes Newton-based numerical solvers that have been used to compute Exact Coherent Structures (ECS) and dynamical connections between them. Due to the presence of lateral boundaries, the types of ECS that are possible in this system are equilibria, time-periodic orbits (and quasiperiodic orbits which are not considered in detail). Neither relative periodic orbits nor traveling waves (relative equilibria) are present because the lateral boundaries (primarily end walls) break the continuous translational symmetry in the longitudinal (x) direction. Due to the presence of a discrete symmetry – rotation about the center of the flow domain by 180 degrees – it may also be possible to find pre-periodic orbits $\mathbf{X}(t + T) = \mathcal{R}\mathbf{X}(t)$ (pairs of such pre-periodic orbits make up the periodic orbits $\mathbf{X}(t + 2T) = \mathbf{X}(t)$). Section 4.1 describes a matrix-free Newton-Krylov/GMRES method for computing ECS and Section 4.2 discusses computing dynamical connections between equilibria. Computation of dynamical connections for which the origin and/or destination is a periodic orbit is not pursued in this thesis.

4.1 Computing equilibria and periodic orbits

The standard numerical approach to computing unstable solutions of the dynamical system (2.44) relies on solving an appropriately defined (nonlinear) equation $\mathbf{F}(\mathbf{X}) = 0$ using an iterative Newton's method [63]. Both isolated stable and unstable solutions and ECS embedded in the chaotic set can be computed using an essentially identical procedure, with the main difference being how the initial guess for the iteration is chosen. In particular, for isolated equilibria such as \mathbf{E}_3 and \mathbf{E}_4 , initial conditions could be generated using coarse/fine parameter continuation and/or combinations of known stable/unstable solutions or their



Figure 4.1: Sample plot of state space speed S(t). Local minima below the threshold $S_0 = 0.1$ are denoted with black circles. The values of S(t) are O(1) because the system has been nondimensionalized such that both the relevant length and time scales are unity.

eigenvectors and their symmetry properties near bifurcations. Such an approach was used to compute for instance various branches of solutions in Figure 3.5b.

The initial guesses for ECS can be obtained by performing a long time-integration of the evolution equation, discarding the initial transient, and then looking for signatures of a turbulent trajectory $\mathbf{X}(t)$ passing near an ECS of a particular type. For instance, initial guesses for equilibria were obtained by finding the local minima of the state space speed

$$S(t) = \left\|\partial_t \mathbf{X}(t)\right\| \approx \left\|\frac{\mathbf{X}(t+\delta t) - \mathbf{X}(t)}{\delta t}\right\|,\tag{4.1}$$

where δt is the time step of the numerical integrator. Sufficiently deep minima, such as the ones circled in Figure 4.1 indicate close passes of the turbulent trajectory near an unstable fixed point. The state $\mathbf{X}(t_0)$ at the time instants t_0 which correspond to a minimum of S(t) serves as good initial guess that can be refined to find the nearby fixed point using Newton iterations.

For time-periodic (or pre-periodic) ECS, initial guesses can be obtained by identifying near-recurrences as discussed in Section 3.2.2. Again, a local, sufficiently deep minimum in a recurrence plot, such as the one circled in Figure 4.2, at some (t_0, τ_0) implies that over the time interval $(t_0 - \tau_0, t_0)$ the turbulent trajectory shadowed an unstable periodic or pre-periodic orbit of period $T \approx \tau_0$. The snapshot of the turbulent trajectory $\mathbf{X}(t_0 - \tau_0)$ can then be taken as an initial guess that will be refined to find the shadowed (pre-)periodic



Figure 4.2: Sample recurrence plot at I = 25 mA, where the dynamics is turbulent. A local minimum in the recurrence below a threshold of $d_0 = 0.1$ is circled in black.

orbit using Newton iterations.

4.1.1 Newton's method

The equilibria of the dynamical system (2.44) correspond to zeros of its right-hand-side

$$\mathbf{F}(\mathbf{X}) = \mathbf{V}(\mathbf{X}) = 0. \tag{4.2}$$

This equation can be solved iteratively, generating a sequence X_n , which starts at the initial guess X_0 and converges to X. For small deviation $\delta X_n = X - X_n$ from the solution X, Eq. (4.2) can be linearized

$$\mathbf{F}(\mathbf{X}_n + \delta \mathbf{X}_n) = \mathbf{F}(\mathbf{X}_n) + J(\mathbf{X}_n)\delta \mathbf{X}_n = 0,$$
(4.3)

where $J = \partial \mathbf{F} / \partial \mathbf{X}$ is the corresponding Jacobian matrix. Solving this for $\delta \mathbf{X}_n$, yields the following iterative algorithm:

$$\mathbf{X}_{n+1} = \mathbf{X}_n - J_{\mathbf{X}}(\mathbf{X}_n)^{-1} \mathbf{F}(\mathbf{X}_n).$$
(4.4)

For the model considered here, the Jacobian matrix J is sparse, since finite differencing has been used for spatial discretization, so it can be explicitly computed and stored. For the problem size considered here, the linear equation involving this Jacobian matrix can be solved exactly using LU factorization without large memory or time requirements. Hence, a direct solve has been used to solve the linear equation (4.4). In cases when the resulting Newton step did not decrease the norm of the residual F, a parabolic line search [103] has been used to adjust the step magnitude. Newton iteration was then repeated until the residual norm decreases to a desired tolerance, an absolute tolerance of 10^{-12} used here for equilibria. The dominant stability exponents { Λ_i } of the converged solution and the corresponding eigenmodes are given by the smallest eigenvalues and eigenvectors of J, which were computed using the function **eigs** in Matlab. The algorithm optimized for sparse matrices is described in Refs. [104, 105].

Periodic (and pre-periodic) orbits instead correspond to solutions of the equation

$$\mathbf{F}(\mathbf{X},T) = g\mathbf{\Phi}_T(\mathbf{X}) - \mathbf{X} = 0, \tag{4.5}$$

where g = 1 or \mathcal{R} and $\Phi_T(\mathbf{X}_0)$ denotes the solution $\mathbf{X}(T)$ of (2.44) with the initial condition $\mathbf{X}(0)$. The time period T is *a priori* unknown and has to be solved for along with \mathbf{X} . The initial guess $T_0 = \tau_0$ is also obtained from the recurrence plot. Starting with the initial guess (\mathbf{X}_0, T_0) , solutions of equation (4.5) can again be found iteratively. Linearizing Eq. (4.5) for small deviations $\delta \mathbf{X}_n = \mathbf{X} - \mathbf{X}_n$ and $\delta T_n = T - T_n$ from the solution

$$\mathbf{F}(\mathbf{X}_n + \delta \mathbf{X}_n, T_n + \delta T_n) = \mathbf{F}(\mathbf{X}_n) + (gJ_{\mathbf{X}}(\mathbf{X}_n) - 1)\delta \mathbf{X}_n + gJ_T(\mathbf{X}_n)\delta T_n = 0, \quad (4.6)$$

where $J_{\mathbf{X}} = \partial \Phi_T / \partial \mathbf{X}$ and $J_T = \partial \Phi_T / \partial T$. Equation (4.6) yields an under-determined linear system for $\delta \mathbf{X}_n$ and δT (there are N equations for N + 1 unknowns). To find a unique solution, we will require that the Newton step be normal to the direction of the state space flow

$$\delta \mathbf{X}_n \cdot \mathbf{V}(\mathbf{X}_n) = 0. \tag{4.7}$$

Imposing this constraint yields an augmented linear system

$$J(\mathbf{X}_n) \begin{pmatrix} \delta \mathbf{X}_n \\ \delta T_n \end{pmatrix} = - \begin{pmatrix} \mathbf{F}(\mathbf{X}_n) \\ 0 \end{pmatrix}, \qquad (4.8)$$

where

$$J = \begin{pmatrix} gJ_{\mathbf{X}} - \mathbb{1} & gJ_T \\ \mathbf{V}^T & \mathbf{0} \end{pmatrix}$$
(4.9)

and the superscript T denotes matrix transpose. Once the system (4.8) is solved, the next iterate can be computed as

$$\mathbf{X}_{n+1} = \mathbf{X}_n + \delta \mathbf{X}_n,$$

$$T_{n+1} = T_n + \delta T_n.$$
 (4.10)

The dimension of the state space is large $N = O(10^5)$ and the $(N + 1) \times (N + 1)$ matrix J is dense. It cannot be stored in memory, and constructing it explicitly would be prohibitively expensive. Direct inversion or iterative methods like Jacobi, Guass-Siedel, etc. cannot be used to solve systems of equations of such high dimensions either. Hence, we use a Krylov space method to iteratively solve the linear system (4.8). The method computes the solution by projecting the linear system onto a Krylov subspace, which is spanned by the vectors $\{\mathbf{F}_1, \dots, \mathbf{F}_m\}$ generated iteratively as $\mathbf{F}_{m+1} = J\mathbf{F}_m$, starting with an appropriately chosen initial vector \mathbf{F}_1 .

It should be noted that, in order to construct this sequence, the matrix J does not have to be constructed in explicit form, rather a prescription is needed which defines its action on an arbitrary vector \mathbf{F} . In our case, $J_T(\mathbf{X}) = \mathbf{V}(\Phi_T(\mathbf{X}))$ and $\mathbf{V}(\mathbf{X})$ can be computed directly, while the action of the $N \times N$ Jacobian $J_{\mathbf{X}}$ can be approximated using

$$J_{\mathbf{X}}(\mathbf{X})\mathbf{F} \approx \frac{\mathbf{\Phi}_{T}(\mathbf{X} + \epsilon \mathbf{F}) - \mathbf{\Phi}_{T}(\mathbf{X})}{\epsilon}$$
(4.11)

for a suitably small ϵ . Since the Jacobian is not computed explicitly, this method is referred to as matrix-free. Different variations of Krylov space methods have been developed for various types of problems. For example, the conjugate gradients method [106] is very efficient for solving systems with symmetric J. For non-symmetric J a generalized minimum residual (GMRES) method [107, 108] is commonly employed. This is the method that has been used to solve the system (4.8).

4.1.2 GMRES

GMRES approximates the exact solution of the system $J\mathbf{u} = \mathbf{b}$ (where **b** is assumed to be normalized, $\|\mathbf{b}\| = 1$), by a vector $\mathbf{u}_m \in K_m$, which minimizes the norm of the residual $\mathbf{F}_m = J\mathbf{u}_m - \mathbf{b}$. Here K_m is the *m*-dimensional Krylov subspace spanned by the set of vectors $\{\mathbf{b}, J\mathbf{b}, \dots, J^{m-1}\mathbf{b}\}$. A more convenient orthonormal basis $\{\hat{\mathbf{q}}_1, \dots, \hat{\mathbf{q}}_m\}$ for K_m can be constructed using the stabilized Gram-Schmidt method. The algorithm is known as Arnoldi iteration and generates an associated $(m + 1) \times m$ upper Hessenberg matrix H_m whose eigenvalue spectrum approximates that of J.

$$\begin{split} \hat{\mathbf{q}}_{1} &= \mathbf{b} \\ & \text{for } l = 1 \text{ to } m \text{ do} \\ & \left| \begin{array}{c} \mathbf{q}_{l+1} &= J\hat{\mathbf{q}}_{l} \ // \ \text{Generate new vector} \\ & \text{for } j = 1 \text{ to } l \text{ do} \\ & \left| \begin{array}{c} H_{j,l} &= \mathbf{q}_{l+1} \cdot \hat{\mathbf{q}}_{j} \ // \ \text{Projection onto previous vector} \\ & \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} - H_{j,l} \hat{\mathbf{q}}_{j} \ // \ \text{Orthogonalization} \\ & \text{end} \\ & H_{l+1,l} &= \left\| \mathbf{q}_{l} \right\| \ // \ \text{Magnitude of new orthogonal vector} \\ & \hat{\mathbf{q}}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{end} \\ & \text{end} \\ & \text{end} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{Normalization} \\ & \text{for } \mathbf{q}_{l+1} &= \mathbf{q}_{l+1} / H_{l+1,l} \ / \ \text{for } \mathbf{q}_{l+1} \\ & \text{for } \mathbf{q}_{l+1} \\$$

Algorithm 1: Arnoldi iteration. The hat indicates a unit vector.

Since \mathbf{u}_m , $\mathbf{b} \in K_m$, these vectors can be expanded in the basis $\{\hat{\mathbf{q}}_1, \cdots, \hat{\mathbf{q}}_{m+1}\}$ as

$$\mathbf{u}_m = Q_m \mathbf{z}_m,\tag{4.12}$$

$$\mathbf{b} = Q_{m+1}\mathbf{v}_1 \tag{4.13}$$

where $Q_m = [\hat{\mathbf{q}}_1, \cdots, \hat{\mathbf{q}}_m]$, $\mathbf{z}_m \in \mathbb{R}^m$ and $\mathbf{v}_1 = (1, 0, \cdots, 0)^T \in \mathbb{R}^{m+1}$. Furthermore, using the relations $JQ_m = Q_{m+1}H_m$ and $Q_{m+1}^TQ_{m+1} = \mathbb{1}$ it can be shown that

$$\|\mathbf{F}_{m}\| = \|J\mathbf{u}_{m} - \mathbf{b}\| = \|Q_{m+1}H_{m}\mathbf{z}_{m} - Q_{m+1}\mathbf{v}_{1}\| = \|H_{m}\mathbf{z}_{m} - \mathbf{v}_{1}\|.$$
 (4.14)

This result can be easily generalized to the case of $\beta = \|\mathbf{b}\| \neq 1$ by rescaling

$$\|\mathbf{F}_m\| = \|H_m \mathbf{z}_m - \beta \mathbf{v}_1\|. \tag{4.15}$$

Therefore, minimizing the residual in the full N-dimensional state space is equivalent to solving a linear least squares problem (4.15) in an m + 1 dimensions whose solution is

$$\mathbf{z}_m = \beta H_m^* \mathbf{v}_1, \tag{4.16}$$

where H_m^* is the pseudoinverse of H_m . Successive approximations \mathbf{u}_m with increasing m can be generated until $\|\mathbf{F}_m\|$ becomes sufficiently small, resulting in two nested loops: GMRES iteration (over m) within each Newton iteration (over n).

In practice, the linear system typically does not have to be solved exactly, hence this method is called inexact Newton method. A variety of criteria can be used to stop the GMRES iteration, as the ultimate objective is to solve the non-linear equation $\mathbf{F}(\mathbf{X}) = 0$. The stopping criterion used in this thesis is based on the eigenvalues of the Hessenberg matrix and the following procedure is followed. A minimum of 10 GMRES iterations is always performed. The eigenvalues $h_{m,i}$ of H_m sorted from largest to smallest in magnitude

are then computed and the products $h_m^k = |h_{m,1}| |h_{m,2}| \cdots |h_{m,k}|$ are evaluated. GMRES iterations are then continued until $m \ge 2k_0$, where k_0 is the smallest k for which $h_m^k < 1$. The motivation behind this constraint is to achieve a balance between unstable and weakly stable modes for unstable solutions. If the solution is stable, then $k_0 = 1$, and the GMRES iterations are instead continued until the dominant eigenvalues of H_m converge (to the eigenvalues of J), i.e.,

$$|h_{m,m-1} - h_{m-1,m-1}| \le \varepsilon |h_{m-1,1}| \tag{4.17}$$

with $\varepsilon = 10^{-3}$ used here. If none of the above criteria are met, GMRES iterations are continued until a maximum $m_{\text{max}} = 40$.

The next step after computing the correction $\delta \mathbf{X}_n$ (and δT_n for (pre-)periodic orbits) using GMRES is to compare the norm $\|\mathbf{F}(\mathbf{X}_{n+1})\|$ of the new non-linear residual with the norm $\|\mathbf{F}(\mathbf{X}_n)\|$ of the old one. If the residual decreases, the algorithm proceeds to the next Newton iteration. However, occasionally the non-linear residual does not decrease. This means that \mathbf{X}_n is far from the solution \mathbf{X} , the linear approximation underlying the Newton method is not valid, and so computed correction is inaccurate. In such a case, a "hook step" is performed [61] as explained below.

4.1.3 Hook step

In the hook step algorithm, a correction is also computed in the Krylov subspace K_m , however the linear residual \mathbf{F}_m is minimized subject to the constraint that \mathbf{z}_m lies on a hypersphere of radius Δ , known as the trust region radius. This is equivalent to a constrained minimization problem defined by the cost function

$$\mathcal{L} = \|H_m \mathbf{z}_m - \beta \mathbf{v}_1\|^2 + \mu (\|\mathbf{z}_m\|^2 - \Delta^2),$$
(4.18)

where μ is the Lagrange multiplier. It can be simplified by rotating the basis using the singular value decomposition of the Hessenberg matrix

$$H_m = M_{m+1} D_m N_m^T, (4.19)$$

where M_{m+1} and N_m are square orthogonal matrices and D_m is an $(m+1) \times m$ diagonal matrix with the last row of all zeros. Then, the residual given by Eq. (4.15) will be

$$\|\mathbf{F}_{m}\| = \|H_{m}\mathbf{z}_{m} - \beta\mathbf{v}_{1}\|$$

$$= \|M_{m+1}D_{m}N_{m}^{T}\mathbf{z}_{m} - \beta\mathbf{v}_{1}\|$$

$$= \|M_{m+1}(D_{m}N_{m}^{T}\mathbf{z}_{m} - M_{m+1}^{-1}\beta\mathbf{v}_{1})\|$$

$$= \|D_{m}\mathbf{y}_{m} - \mathbf{p}_{m}\|$$
(4.20)

where

$$\mathbf{p}_{m} = \beta M_{m+1}^{-1} \mathbf{v}_{1} = \beta M_{m+1}^{T} \mathbf{v}_{1},$$
$$\mathbf{y}_{m} = N_{m}^{T} \mathbf{z}_{m}.$$
(4.21)

Since N_m is an orthogonal matrix, $\|\mathbf{y}_m\| = \|\mathbf{z}_m\|$, so (4.18) becomes

$$\mathcal{L} = \|D_m \mathbf{y}_m - \mathbf{p}_m\|^2 + \mu \left(\|\mathbf{y}_m\|^2 - \Delta^2\right).$$
(4.22)

The solution can be computed by setting the partial derivatives of \mathcal{L} with respect to the components $y_{m,i}$ of \mathbf{y}_m to zero, which yields

$$y_{m,i} = \frac{p_{m,i}(D_m)_{ii}}{(D_m)_{ii}^2 + \mu}, \quad 1 \le i \le m,$$
(4.23)

where $(D_m)_{ii}$ is the *i*th diagonal element of D_m . The unknown constant μ is then computed by requiring that $\|\mathbf{y}_m\| = \Delta$.

In practice, one does not know Δ and starts with a trust region radius computed in the previous Newton iteration. If the hook step succeeds (i.e., the new nonlinear residual is smaller than the old one), then the trust region radius Δ is increased and the hook step is repeated as long as the non-linear residual keeps decreasing. This way, a larger hook step is taken, reducing the number of expensive Newton iterations. If the hook step fails, then Δ is reduced and the hook step is repeated with the new, reduced Δ . The state is updated using the correction solved by the GMRES-hook step, and the Newton iterations are continued until the residual norm $\|\mathbf{F}\|$ decreases to a desired tolerance, an absolute tolerance of 5×10^{-6} used here for periodic orbits.

As pointed out previously, the eigenvalues of the Hessenberg matrix H generated by GMRES approximate the leading eigenvalues of the Jacobian J when the Krylov subspace has a sufficiently large dimensionality. Similarly, the eigenspectrum of the Hessenberg matrix H_X generated in the arnoldi iteration using the action of J_X approximates its dominant eigenspetrum. These eigenvalues, evaluated at the solution $\mathbf{X}(T) = \mathbf{X}(0)$, are simply the Floquet multipliers (3.6) of the corresponding orbit, which determine its stability.

It is worth mentioning that the Newton-Krylov/GMRES method described here, which is designed primarily to find periodic orbits, can also be (and often is) used to compute the equilibria of the dynamical system (2.44). This requires a judicious choice of T that is held fixed between Newton iterations.

4.2 Computing dynamical connections

4.2.1 Shooting method

Dynamical (or heteroclinic or homoclinic) connections between unstable equilibria can be computed using a shooting method that has a lot in common with the method for computing periodic orbits described in the previous section. Let \mathbf{E}_{-} be the origin equilibrium where the connection starts and \mathbf{E}_+ be the destination equilibrium where the connection terminates. The connection lies within the unstable manifold of \mathbf{E}_- and so can, in principle, be computed by time-integration of (2.44) starting from a suitable initial condition $\mathbf{X}(0)$ which lies on that manifold [109]. The manifold is tangent to a hyperplane spanned by all the unstable eigenvectors \mathbf{e}_k of \mathbf{E}_- , so we can choose

$$\mathbf{X}(0) = \mathbf{E}_{-} + \sum_{k=1}^{m_u} \alpha_k \mathbf{e}_k, \qquad (4.24)$$

where m_u is the number of unstable eigenvalues of E_- . The initial condition X should be close to E_- , i.e.,

$$\|\mathbf{X}(0) - \mathbf{E}_{-}\| = \epsilon_{-},\tag{4.25}$$

with ϵ_{-} sufficiently small to make sure the hyperplane is an accurate approximation to the manifold. On the other hand, $\mathbf{V}(\mathbf{X}) \to 0$ for $\mathbf{X} \to \mathbf{X}_{-}$, so ϵ_{-} should be large enough to make the total integration time reasonable. Eq. (4.25) defines a hypersphere of radius ϵ_{-} around \mathbf{X}_{-} , so the initial condition $\mathbf{X}(0)$ lies at the intersection of the hypersphere and the hyperplane.

Furthermore, the connection should lie within the stable manifold of \mathbf{E}_+ which is tangent to the hyperplane spanned by all the stable eigenvectors of \mathbf{E}_+ and so $\mathbf{X}(T)$ should lie on this hyperplane. While it is possible to express this condition in the form similar to (4.24), the number of stable eigenvalues for a system such as the one considered here is very large, $m_s = \mathcal{O}(N)$. An alternative to impose this constraint is to require

$$P_{+u}\left(\mathbf{X}(T) - \mathbf{E}_{+}\right) = 0 \tag{4.26}$$

where P_{+u} is the projection operator onto the low-dimensional unstable eigenspace of \mathbf{E}_+ and $\mathbf{X}(T) = \mathbf{\Phi}_T(\mathbf{X}(0))$ [67]. However, the projection requires the knowledge of the corresponding adjoint (or left) eigenvectors, which are not generally straightforward to compute. On the other hand, $\mathbf{V}(\mathbf{X}) \to 0$ for $\mathbf{X} \to \mathbf{E}_+$, so we should require that the numerically computed trajectory that originates at $\mathbf{X}(0)$ terminate *anywhere* inside a hypershere of finite radius ϵ_+ centered at \mathbf{E}_+ ,

$$\|\mathbf{F}(\mathbf{X}(0), T)\| = \|\mathbf{X}(T) - \mathbf{E}_{+}\| = \epsilon_{+}$$
(4.27)

at some large, but again finite, time T. The constraint (4.27) does not guarantee that $\mathbf{X}(T)$ lies exactly in the stable manifold of \mathbf{E}_+ . However, by choosing a sufficiently small ϵ_+ , we ensure that the trajectory comes arbitrarily close to \mathbf{X}_+ , which is only possible only if $\mathbf{X}(T)$ lies very close to the stable manifold of the destination. Constraints (4.24), (4.25), and (4.27) therefore give us a very good approximation for the connection, provided a small number of "coordinates" α_k is chosen correctly.

This is essentially a shooting-type problem, where the initial condition parametrized by $\{\alpha_k\}$ is varied to minimize the $\|\mathbf{F}\|$ for a sufficiently large T, with the tolerance determined by (4.27). Provided that a good initial guess is available, the solution can be obtained iteratively by refining it using Newton's method. Minimization of the norm $\|\mathbf{F}\|$ is, of course, equivalent to solving the nonlinear equation

$$\mathbf{F}(\mathbf{X},T) = 0,\tag{4.28}$$

so we can more or less follow the procedure outlined in the previous section. Let the guess for the initial condition $\mathbf{X}(0)$ at the *n*th iteration be $\mathbf{X}_n = \mathbf{X} - \delta \mathbf{X}_n$, where

$$\delta \mathbf{X}_n = \sum_{k=1}^{m_u} \delta \alpha_k^{(n)} \mathbf{e}_k.$$
(4.29)

Linearizing (4.28) about X and $T = T_n + \delta T_n$ and ignoring the higher order terms, we get

$$\mathbf{F}(\mathbf{X}_n + \delta \mathbf{X}_n, T_n + \delta T_n) = \mathbf{F}(\mathbf{X}_n, T_n) + \sum_{k=1}^{m_u} \delta \alpha_k^{(n)} \mathbf{w}_k + \delta T_n \mathbf{w}_T = 0$$
(4.30)

where

$$\mathbf{w}_{k} = \frac{\partial \mathbf{F}}{\partial \alpha_{k}^{(n)}} \approx \frac{\mathbf{F}(\mathbf{X}_{n} + \epsilon \mathbf{e}_{k}, T) - \mathbf{F}(\mathbf{X}_{n}, T)}{\epsilon}$$
$$= \frac{\Phi_{T}(\mathbf{X}_{n} + \epsilon \mathbf{e}_{k}) - \Phi_{T}(\mathbf{X}_{n})}{\epsilon},$$
$$\mathbf{w}_{T} = \frac{\partial \mathbf{F}}{\partial T} \approx \frac{\mathbf{F}(\mathbf{X}_{n}, T + \epsilon) - \mathbf{F}(\mathbf{X}_{n}, T)}{\epsilon}$$
$$= \frac{\Phi_{T+\epsilon}(\mathbf{X}_{n}) - \Phi_{T}(\mathbf{X}_{n})}{\epsilon}$$
(4.31)

with sufficiently small ϵ . The solution to (4.30) is obtained by minimizing

$$E^{(n)} = \left\| \mathbf{F}(\mathbf{X}_n, T_n) + \sum_{k=1}^{m_u} \delta \alpha_k^{(n)} \mathbf{w}_k + \delta T_n \mathbf{w}_T \right\|^2.$$
(4.32)

Setting the partial derivatives $\partial E^{(n)}/\partial \delta \alpha_k^{(n)}$ to zero yields m_u linear equations

$$\sum_{k=1}^{m_u} \delta \alpha_k^{(n)} (\mathbf{w}_j \cdot \mathbf{w}_k) + \delta T_n (\mathbf{w}_j \cdot \mathbf{w}_T) = -\mathbf{w}_j \cdot \mathbf{F}(\mathbf{X}_n, T), \qquad 1 \le j \le m_u.$$
(4.33)

Furthermore, the constraint (4.25) requires

$$\sum_{j=1}^{m_u} C_j \delta \alpha_j^{(n)} = 0, \tag{4.34}$$

where

$$C_j = \sum_{i=1}^{m_u} (\mathbf{e}_i \cdot \mathbf{e}_j) \alpha_i^{(n)}.$$
(4.35)

Solving the linear system of equations (4.33) and (4.34) uniquely determines the solution $\{\delta \alpha_k^{(n)}\}\$ and δT_n .

The temporal length of the connection is infinite, so we cannot use this result for δT_n , which is based on linearization, to compute the new guess T_{n+1} . Instead, once the new guess \mathbf{X}_{n+1} has been computed using $\{\delta \alpha_k^{(n)}\}$, the trajectory is computed over a time interval $[0, T_{max}]$, with $T_{max} \gg T_n$. Then T_n is set as the value of T which minimizes the left-hand-side of (4.27). It is possible that the residual norm $||\mathbf{F}||$ does not decrease after the Newton update, which means that the calculated $\delta \mathbf{X}_n$ is far outside the region where the linearization can be considered accurate. In that case, a parabolic line search (see [103]) is performed by taking a fraction of the computed $\delta \mathbf{X}_n$, and updating \mathbf{X}_{n+1} and T_{n+1} accordingly.

Since Newton search is based on linearization, while the heteroclinic connection is an unstable solution which exponentially amplifies errors in the choice of the initial condition, an extremely good initial guess is required to find the heteroclinic connection using shooting. In general, even moderately good guesses are not readily available. Moreover, it is in general unknown which pairs of equilibria possess sufficiently short connections. Both of these problems can be addressed, however, in the following systematic way. Let us take an equilibrium and continue it, in any parameter(s), until it undergoes a bifurcation. At the bifurcation it merges with another unstable (or stable) solution. At an infinitesimal distance away from the bifurcation, the two solutions are connected, with the connection lying in (or even coinciding with) their center manifold. In many instances (e.g., for saddle-node, transcritical, and pitchfork bifurcations of equilibria) the center manifold between the two solutions is one-dimensional, locally straight, and serves as an excellent initial guess for the connection.

After computing the connection near the bifurcation, the origin and destination equilibria and the connection between them can be continued back to the original values of parameter(s). To illustrate this method, we have computed the heteroclinic connection between the equilibria \mathbf{E}_3 and \mathbf{E}_4 born in a saddle-node bifurcation at $I = I_1$ (the connection is created at the same bifurcation). The calculations used the following parameters: $\epsilon_- = \epsilon_+ = 5 \times 10^{-4} ||\mathbf{X}_+||$ and $T_{max} = 2000$ seconds. The connection computed at various values of the current I, is plotted in Figure 4.3. The destination equilibrium \mathbf{E}_4 has no unstable directions below and one unstable direction above the pitchfork bifurcation



Figure 4.3: Heteroclinic connection (green) between unstable equilibria \mathbf{E}_3 (blue) and \mathbf{E}_4 (black) computed at a set of different driving currents to illustrate the continuation procedure. The equilibria at intermediate values of the current are shown as dashed lines and bifurcations as spheres (brown for the saddle-node, cyan for the pitchforks, and red for the Hopf bifurcation. The same projection is used as in Figure 3.5b.



Figure 4.4: Bifurcation diagram showing equilibrium states E_1 through E_8 (lines). Pitchfork bifurcations are labeled with cyan spheres and Hopf bifurcations are labeled with red spheres. Saddle-node bifurcations are not labeled; they correspond to points where equilibria E_3 and E_4 or E_7 and E_8 meet. The same projection is used as in Figure 3.5b.

(shown as cyan sphere), while the number of unstable directions for the origin E_3 increases from one to two (at the pitchfork bifurcation) to four (at the Hopf bifurcation, shown as red sphere, which takes place at I = 11.38 mA). While in general dynamical connections between ECS may be created or destroyed when either the origin or the destination undergoes a bifurcation, the connection between E_3 and E_4 persists despite both equilibria undergoing several bifurcations.

4.2.2 Symmetry constraint

The equilibrium \mathbf{E}_4 undergoes a Hopf bifurcation at I = 11.74 mA. Just above it, two equilibria \mathbf{E}_7 and \mathbf{E}_8 are born out of a saddle node bifurcation at I = 11.76 mA. A projection of the bifurcation diagram is shown in Figure 4.4. We have tried to compute the connection between the saddle \mathbf{E}_7 and node \mathbf{E}_8 at a driving current of I = 11.9 mA, which is just above the location of saddle-node bifurcation, using the shooting method. For this current, the destination \mathbf{E}_8 has one unstable direction, while the origin \mathbf{E}_7 has two unstable directions with the corresponding stability exponents $\Lambda_1 = 0.037$ and $\Lambda_2 = 0.004$ differing by an order of magnitude. This large disparity caused the Newton method to fail due to the extreme amplification of the perturbations in the eigendirection \mathbf{e}_1 for the connection that is aligned approximately in the eigendirection \mathbf{e}_2 (since the latter corresponds to the center manifold for this bifurcation).

Since the unstable manifold of the origin is just two-dimensional, all trajectories lying in this manifold can be parameterized using initial conditions lying on the circle of radius ϵ_{-} in the plane spanned by e_1 and e_1 . Hence the trajectory connecting E_7 to E_8 can, in principle, be found by applying bisection method to this one-dimensional set of initial conditions. However, even bisection failed to determine the initial condition with sufficient accuracy to generate a trajectory that would approach the destination equilibrium numerical accuracy was exceeded first. Instead, all initial conditions produced trajectories that terminated at the stable equilibrium E_2 (which corresponds to the modulated flow) as illustrated in Figure 4.6b.

Although computation of the connection using shooting in the full symmetry-unrestricted state space fails, the basic approach can be salvaged by noticing that both \mathbf{E}_7 and \mathbf{E}_8 lie in the rotationally-symmetric subspace (cf. Figure 4.5). So does the connection, at least close to the bifurcation, where it coincides with the center manifold, whose direction is given by the eigenvector \mathbf{e}_2 . Indeed, near the saddle-node bifurcation, \mathbf{e}_2 is proportional to the difference $\mathbf{E}_7 - \mathbf{E}_8$ and so also lies in the symmetry subspace. Since there is no mechanism that would break its symmetry, the connection remains symmetric away from the bifurcation as well. On the other hand, \mathbf{e}_1 is not rotationally-symmetric, as can be seen from Figure 4.5. Hence, an initial condition of the form $\mathbf{E}_7 + \epsilon_- \mathbf{e}_2$, when evolved forward, should lie in the symmetric subspace.

The numerical noise at every time step weakly breaks rotational symmetry of the state, which leads to a gradual drift of the trajectory away from the symmetric subspace. These tiny perturbations due to the numerical noise are exponentially amplified in the domi-



Figure 4.5: Vorticity fields of (a) \mathbf{E}_7 , (b) \mathbf{E}_8 and the eigenvectors (c) \mathbf{e}_1 , (d) \mathbf{e}_2 of \mathbf{E}_7 at I = 11.9 mA, where the heteroclinc connection is computed.



Figure 4.6: (a) Projections of the trajectories that start on the unstable manifold of the origin equilibrium E_7 (blue sphere) onto the subspace spanned by the orthogonal vectors d_1 , d_2 , d_3 which are linear combinations of e_1 , e_2 , and E_8-E_7 . Without the symmetry constraint, all trajectories (black curves) terminate at the global attractor E_2 (magenta sphere). With the symmetry constraint, the trajectory (green curve) from $E_7+\epsilon_-e_2$ terminates at the destination equilibrium E_8 (red sphere). (b) Zoomed view of (a) near E_7 showing the focusing effect of the dominant unstable direction.



Figure 4.7: The distance from the destination equilibrium \mathbf{E}_8 versus time for the symmetryconstrained trajectory with the initial condition $\mathbf{E}_7 + \epsilon_- \mathbf{e}_2$.

nant unstable direction, completely destroying the symmetry. In order to eliminate the symmetry-breaking effect of numerical noise, we projected the solution back into the symmetry subspace at every time step. With this correction, the shooting method generates the instability and produces a trajectory that lies in the symmetric subspace and terminates at the destination equilibrium as shown in Figure 4.6. To confirm this we computed the distance from the trajectory to the destination equilibrium \mathbf{E}_8 , which is shown in Figure 4.7. Note that since the intersection of the symmetry subspace with the unstable manifold of \mathbf{E}_7 is one-dimensional, no Newton iteration is necessary to find the connection. The choice of the initial condition along \mathbf{e}_2 guarantees convergence to \mathbf{E}_8 , since \mathbf{E}_8 is stable within the symmetry subspace (although there is an unstable direction that lies outside of the subspace).

4.3 Summary

In this chapter, Newton-based methods for computing ECS and their dynamical connections have been described. Close to the solution, Newton steps are second-order accurate, i.e., the residual at the n^{th} iteration satisfies $||F_n|| \leq C||F_{n-1}||^2$ for some constant C. This leads to a super-exponential convergence over multiple iterations [110]. However, such fast convergence is only observed when the residual is already very small. Such good initial guesses are usually unavailable or, as in the case of parameter continuation, costly to obtain. If/when a very good initial guess is not available, the convergence can be much slower. In fact, Newton-based methods can completely fail to converge. One reason is the nonlinearity of the problem. Another, more serious one is the exponential sensitivity of the numerical solutions obtained by shooting to small perturbations in the initial conditions, which is exacerbated for connecting orbits that have a formally infinite duration. The algorithm completely breaks down when the initial condition has to be specified with an accuracy exceeding the accuracy of computer arithmetic. One possible solution to handle such highly unstable solutions is to use multi-shooting [111], where the trajectory is broken up into multiple shorter segments, such that on each segment the sensitivity to initial condition is dramatically decreased. However, multi-shooting increases the dimensionality of the problem, requiring a substantial increase in both memory and processing power.

The symmetry-based method for computing connections partially overcomes the problems of the Newton-based methods. While it works extremely well and is very efficient, it is limited to pairs of solutions that lie in the same symmetry subspace. Furthermore, it requires that the destination have no unstable directions in the symmetry subspace. An alternative approach that is both very general and overcomes the limitations of the methods described above is introduced in the next chapter.

CHAPTER 5

ADJOINT-BASED METHODS FOR COMPUTING ECS AND DYNAMICAL CONNECTIONS

In this chapter we describe an alternative method for computing unstable solutions of the Navier-Stokes equation. It uses the adjoint-based formulation that was introduced into this class of problems by Farazmand [66], who illustrated its power by computing several new solutions of the two-dimensional Kolmogorov flow described by (2.4) with $\beta = 1$ and $\tilde{\alpha} = 0$ and (2.5). In its original form, this method was limited only to equilibria and traveling waves (relative equilibria) on physical domains with periodic boundary conditions. We have generalized this method in several ways. In particular, the generalized method allows computing solutions on domains with physical no-slip boundary conditions. Furthermore, it can be used for computing periodic orbits and heteroclinic connections between equilibria. The method is described in this chapter, starting with its application to a system of two linear equations which provides a simple geometrical illustration of the main concepts.

5.1 Adjoint evolution equation

Let $\mathbf{F}(\mathbf{X})$ be a function from \mathbb{R}^N to \mathbb{R}^N . Finding a solution \mathbf{X}^* of the vector equation

$$\mathbf{F}(\mathbf{X}) = 0 \tag{5.1}$$

is equivalent to finding a root of the non-negative scalar function $E = \mathbf{F}^T W \mathbf{F}$ for any symmetric positive definite weight matrix W (for F-complex, we should use Hermitian conjugate rather than transpose). Assuming **X** is a function of a fictitious time τ , we can write

$$\partial_{\tau} E = 2\mathbf{F}^T W \partial_{\tau} \mathbf{F} = 2\mathbf{F}^T W J \partial_{\tau} \mathbf{X} = 2(J^T W \mathbf{F})^T \partial_{\tau} \mathbf{X}, \tag{5.2}$$

where $J = \partial \mathbf{F} / \partial \mathbf{X}$ is the Jacobian of \mathbf{F} . If we were to require that \mathbf{X} evolve in fictitious time according to

$$\partial_{\tau} \mathbf{X} = -Q J^T W \mathbf{F}(\mathbf{X}), \tag{5.3}$$

where $Q = D^T D$ is a symmetric positive-definite acceleration matrix, then $E(\mathbf{X})$ effectively becomes a Lyapunov function for the dynamical system (5.3). Indeed, it is straightforward to check that E decreases in time

$$\partial_{\tau} E = -2 \| D J^T W \mathbf{F} \|^2 \le 0, \tag{5.4}$$

except when

$$J^T W \mathbf{F} = 0. \tag{5.5}$$

The last equation is only satisfied either when $\mathbf{F} = 0$ (i.e., a root of (5.1) has been found) or when \mathbf{F} lies in the kernel of $J^T W$. The latter possibility is a potential drawback of the method that might prevent global convergence, as discussed below.

Since (5.3) involves the transpose of the Jacobian, which defines the adjoint eigenvalue problem for J, we will refer to this method as "adjoint-based". The matrices Q and W play the role of preconditioners and can be used to improve the convergence speed, as will be shown below. In the special case Q = W = 1 we have $E = ||\mathbf{F}||^2$ and (5.3) reduces to

$$\partial_{\tau} \mathbf{X} = -J^T \mathbf{F} = -\frac{1}{2} \nabla E.$$
(5.6)

The evolution of X in fictitious time described by (5.6) clearly corresponds to the steepest descent direction for E.

5.1.1 Illustrative example

To gain some geometrical intuition about the role of the matrices Q and W, the adjointbased method is applied for finding the root of a linear function $\mathbf{F}(\mathbf{X}) = A\mathbf{X}$ in \mathbb{R}^2 , where $\mathbf{X} = [x, y]$ and

$$A = \begin{bmatrix} a & 0 \\ 0 & 1 \end{bmatrix}, \tag{5.7}$$

with a = 1/4 and $\mathbf{X}(0) = (1, 0.5)^T$. In this problem, $J = J^T = A$. For W = 1 we obtain the steepest descent equation (5.6), which can be solved analytically:

$$\partial_{\tau} \begin{bmatrix} x \\ y \end{bmatrix} = -A^2 \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} -a^2 x \\ -y \end{bmatrix}, \qquad (5.8)$$

such that

$$x(\tau) = x(0)e^{-a^{2}\tau},$$

$$y(\tau) = y(0)e^{-\tau}.$$
(5.9)

The corresponding residual is also found to decay exponentially fast

$$E(\tau) = a^2 x^2(\tau) + y^2(\tau) = a^2 x^2(0) e^{-2a^2\tau} + y^2(0) e^{-2\tau},$$
(5.10)

as shown in Figure 5.1a.

The trajectory $\mathbf{X}(\tau)$, along with the level sets of E, is shown in Figure 5.1b. The trajectory is everywhere normal to the level set of E, so since the level sets of E are stretched along the x-direction relative to the y-direction, the steepest descent trajectory



Figure 5.1: Comparison of the solutions to the adjoint evolution equations with and without acceleration. (a) The residual $E(\tau)$. (b) The "trajectories" $\mathbf{X}(\tau)$, the level sets of E (cyan curves), and the solution \mathbf{X}^* (magenta circle). Blue lines correspond to the steepest descent, black lines – to the accelerated descent, and red – to the momentum method. The same choice of Q and W has been used in the accelerated descent and the momentum method.

will approach the solution in the y-direction much faster than that in the x-direction, as Figure 5.1b illustrates. The x-direction, along which the residual decreases slowly (at a rate $a^2 = 0.0625 \ll 1$ in our example), corresponds to a slowly decaying mode. It is such slowly decaying modes that control the convergence speed [112].

Convergence speed can be increased by requiring that $\partial_{\tau} \mathbf{X}$ be rotated to increase its component(s) along the slowly decaying mode(s), but such that it remains a descent direction. This can be achieved by choosing the positive-definite symmetric matrices Q and/or W accordingly, such that (5.6) is replaced with (5.3). For instance, choosing

$$W = \begin{bmatrix} w & 0 \\ 0 & 1 \end{bmatrix}, \qquad Q = \begin{bmatrix} q & 0 \\ 0 & 1 \end{bmatrix}, \tag{5.11}$$

with w > 0 and q > 0 yields

$$\partial_{\tau} \begin{bmatrix} x \\ y \end{bmatrix} = -QAWA \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} -a^2 qwx \\ -y \end{bmatrix}, \qquad (5.12)$$

such that

$$x(\tau) = x(0)e^{-a^2qw\tau},$$

$$y(\tau) = y(0)e^{-\tau}$$
(5.13)

and

$$E(\tau) = a^2 w x^2(0) e^{-2a^2 q w \tau} + y^2(0) e^{-2\tau}.$$
(5.14)

This does not change the convergence rate in the y-direction, but speeds up convergence in the x-direction, increasing the convergence rate from a^2 to a^2qw . The results for q = 16and w = 1 are compared with steepest descent in the Figure 5.1a. It is clear that this choice notably accelerates convergence.

5.1.2 Accelerating convergence

When the Jacobian is not symmetric, close to the solution we will find

$$\partial_{\tau} \delta \mathbf{X} = -Q J^T W J \delta \mathbf{X}, \tag{5.15}$$

where $\delta \mathbf{X} = \mathbf{X}(\tau) - \mathbf{X}^*$ is the deviation from the root, so it is the smallest eigenvalues of the product QJ^TWJ that control the convergence speed, with the corresponding eigenvalues defining the slowly decaying modes. In general, the matrices Q and/or W should be chosen to remove these small eigenvalues. For instance, choosing $Q = (J^TJ)^{-1}$ and W = 1 yields $QJ^TWJ = 1$, such that

$$\partial_{\tau} \delta \mathbf{X} = -\delta \mathbf{X},\tag{5.16}$$

with all degrees of freedom converging at exactly the same rate (unity). For a discretization of non-linear PDEs, evaluating the Jacobian, let alone its inverse, is an intractable problem. However, it is still possible to optimize the choice of Q and W using linearization of the PDEs. This is discussed in sections 5.2, 5.3, and 5.4, where solutions of 2D Navier-Stokes equation are computed.

Far from the solution, where (5.15) does not apply, the decay of slow modes can be accelerated by using a classical momentum acceleration method [113] or its variation proposed by Nesterov [114]. In both cases, once a correction has been computed, a fraction of the correction at the previous step is added to the correction predicted by the linear evolution equation (5.3), which is effectively equivalent to adding a "momentum" to the trajectory $\mathbf{X}(\tau)$. For instance, if (5.3) is solved numerically using forward Euler with a constant time step $\delta \tau$, the modified correction during the $k^{\rm th}$ iteration will be given by

$$\delta \mathbf{X}_k = -QJ^T W \mathbf{F}|_{\mathbf{X}_k} \delta \tau + p_k \delta \mathbf{X}_{k-1}.$$
(5.17)

Note that, to leading order in $\delta \tau$, (5.17) corresponds to the discretization of the secondorder differential equation

$$m\partial_{\tau}^{2}\mathbf{X} + \partial_{\tau}\mathbf{X} + QJ^{T}W\mathbf{F} = 0, \qquad (5.18)$$

where $p_k = m/\delta \tau$, so p_k does indeed correspond to inertia or momentum.

In Nesterov's algorithm the inertia is variable, $p_k = (\mu_k - 1)/(\mu_k + 1)$ with

$$\mu_k = \frac{1 + \sqrt{1 + 4\mu_{k-1}^2}}{2} \tag{5.19}$$

and $\mu_1 = 1$. The corresponding $\delta \mathbf{x}$ may not be a descent direction, so a monotonic decrease of the residual is not guaranteed. Whenever the residual E increases during the iteration, the momentum is reset, $\mu_k = 1$ such that $p_k = 0$, after which the iteration continued. Figure 5.1a illustrates that adding momentum to the trajectory further accelerates convergence of the adjoint evolution towards the solution.

5.2 Computing equilibria

In this section, the adjoint based method for finding equilibria of compressible 2D Navier-Stokes equations (which are also the equilibria of incompressible flow) is described. The compressible model (2.33)-(2.37) can be rewritten in the form

$$\dot{u} = F^{u} \equiv -\beta(u\partial_{x} + v\partial_{y})u - \gamma(\partial_{x}u + \partial_{y}v)u + \tilde{R}e^{-1} \left(\nabla^{2}u - \tilde{\alpha}u\right) + \bar{f}_{0} - \partial_{x}p,$$

$$\dot{v} = F^{v} \equiv -\beta(u\partial_{x} + v\partial_{y})v - \gamma(\partial_{x}u + \partial_{y}v)u + \tilde{R}e^{-1} \left(\nabla^{2}v - \tilde{\alpha}v\right) - \partial_{y}p,$$

$$\dot{p} = F^{p} \equiv -\tilde{\kappa}(\partial_{x}u + \partial_{y}v),$$
(5.20)

where $\mathbf{u} = u\hat{x} + v\hat{y}$ and we dropped the subscript \parallel of the gradient operator and the bar on p, since this cannot cause any confusion in 2D.

With the help of the relation (2.45) between the state space and the physical space variables, we can see that the equilibria of (5.20) in the physical space correspond to zeros of the function

$$\mathbf{F}(\mathbf{X}) = \begin{bmatrix} F^u \\ F^v \\ F^p \end{bmatrix}$$
(5.21)

or global minima of $E = \mathbf{F}^T W \mathbf{F}$ in the state space. As discussed above, these minima can be found by solving the evolution equation (5.3) in fictitious time, where

$$J = \frac{\partial \mathbf{F}}{\partial \mathbf{X}} = \begin{bmatrix} \frac{\partial F^u}{\partial u} & \frac{\partial F^u}{\partial v} & -\partial_x \\ \frac{\partial F^v}{\partial u} & \frac{\partial F^v}{\partial v} & -\partial_y \\ -\tilde{\kappa}\partial_x & -\tilde{\kappa}\partial_y & 0 \end{bmatrix}.$$
 (5.22)

For a representation of the differential equations (5.20), where the fields are spatially discretizated using finite differences, all the operators in (5.22) are replaced with the respective finite-dimensional sparse matrices that can be used to compute the matrix vector product $-J^T \mathbf{F}$, which enters the evolution equation (5.3) for \mathbf{X} .

5.2.1 Choice of the weight and acceleration matrices

As discussed previously, to achieve reasonably fast convergence, the operator QJ^TWJ should not have any small eigenvalues. For Q = W = 1, small eigenvalues do arise due to the differences in the time scales associated with the evolution of the velocity and pressure caused by the large value of the constant $\tilde{\kappa}$ in the compressible model. Indeed, discretizing the steepest descent evolution equation (5.6) in fictitious time using forward Euler yields $\delta \mathbf{X} = -J^T \mathbf{F} \delta \tau$. According to (5.2) and (5.22), the corresponding change in the residual E is given by

$$\frac{\delta E}{\delta \tau} = -2 \left\| J^T \mathbf{F} \right\|^2 = -2 \left\| \left(\frac{\partial F^u}{\partial u} \right)^T F^u + \left(\frac{\partial F^v}{\partial u} \right)^T F^v - \tilde{\kappa} (\partial_x)^T F^p \right\|^2 -2 \left\| \left(\frac{\partial F^u}{\partial v} \right)^T F^u + \left(\frac{\partial F^v}{\partial v} \right)^T F^v - \tilde{\kappa} (\partial_y)^T F^p \right\|^2 -2 \| (\partial_x)^T F^u + (\partial_y)^T F^v \|^2,$$
(5.23)

Since $\tilde{\kappa} = \mathcal{O}(10^3)$, variation of E is highly sensitive to the value of F^p , with the associated eigenvalues being $\mathcal{O}(\tilde{\kappa})$ larger than those associated with F^u or F^v . Our previous analysis suggests that the corresponding degrees of freedom (i.e., pressure) should be rescaled by a factor of $\tilde{\kappa}^{-2}$ in (5.3), which can be achieved, for instance, by choosing

$$Q = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \tilde{\kappa}^{-2} \end{bmatrix}.$$
 (5.24)

Farazmand only considered the case Q = 1 and $W \neq 1$ in his study [66]. To explore the alternative, in the following I will only consider the case W = 1 and $Q \neq 1$.

A disparity in the magnitude of the eigenvalues of QJ^TWJ defined using (5.24) can also arise due to the presence of the term ∇^2 in (5.20). The degrees of freedom associated with high wavenumbers in u will have eigenvectors much larger than those associated with low wavenumbers. As a result, either the low frequency modes converge extremely slowly or, if a larger step $\delta \tau$ is used, the high frequency modes become numerically unstable. This can be addressed by using the optimized acceleration matrix

$$Q = \begin{bmatrix} (\tilde{\alpha}\mathbb{1} - \nabla_D^2)^{-1} & 0 & 0\\ 0 & (\tilde{\alpha}\mathbb{1} - \nabla_D^2)^{-1} & 0\\ 0 & 0 & \tilde{\kappa}^{-2}(\nabla_N^2)^{-1} \end{bmatrix},$$
(5.25)

where ∇_D^2 and ∇_N^2 denote the discretization of the Laplacian with Dirichlet and Neumann boundary conditions, respectively.

Two comments should be made here. The operator $(\nabla_N^2)^{-1}$ is singular, since ∇_N^2 has a zero eigenvalue associated with the mode p = const. This mode does not affect the dynamics as pressure appears in the governing equations (5.20) only as ∇p . Hence, this mode of p is set to zero by imposing an explicit constraint in the solution procedure. Furthermore, the choice of the acceleration operator $Q_u = Q_v = (\tilde{\alpha}\mathbb{1} - \nabla_D^2)^{-1}$ is not optimal, based on the analysis presented above. It is, however, cheaper to implement numerically, compared with the more optimal choice $Q_u = Q_v = (\tilde{\alpha}\mathbb{1} - \nabla_D^2)^{-2}$.

5.2.2 Numerical solution of the adjoint equation

The adjoint evolution equation (5.3) has to be integrated in τ until the residual becomes sufficiently small. Integration using an implicit scheme, which is stable for large $\delta\tau$, is very expensive (and often practically impossible) owing to the large dimensionality of the system. Integration using an explicit scheme is also expensive, owing to the stiffness of the system, which requires very small "time" steps. However, we are interested only in the asymptotic state of the system and not the entire trajectory generated by (5.3). So, in order to find the asymptotic state of the system, we do not perform an accurate "time" integration but use an iterative method to evolve the state **X** as follows. Since our governing equations (5.20) are quadratic in X, the residual $E(\mathbf{X} + \mathbf{W}\delta\tau)$ can be written as a 4th degree polynomial in $\delta\tau$ for any X and W. Hence, once the descent direction $\mathbf{W}_k = -QJ^T W \mathbf{F}(\mathbf{X}_k)$ has been computed, $E(\mathbf{X}_k + \mathbf{W}_k \delta\tau_k)$ is fit to a 4th degree polynomial along the descent direction and the optimal "time" step $\delta\tau_k$, which minimizes E along W, is calculated. This gives us the intermediate step, $\delta \mathbf{X}_k^* = \mathbf{W}_k \delta\tau_k$. A fraction of the previous step is added to $\delta \mathbf{X}_k^*$ following Nesterov's algorithm, yielding the step at the k^{th} iterate

$$\delta \mathbf{X}_{k} = \mathbf{W}_{k} \delta \tau_{0} + \frac{\mu_{k} - 1}{\mu_{k} + 1} \delta \mathbf{X}_{k-1}, \qquad (5.26)$$

where μ_k is given by equation (5.19). The next iterate $\mathbf{X}_{k+1} = \mathbf{X}_k + \delta \mathbf{X}_k$ is then evaluated and the process proceeds until convergence.

As discussed in Section 5.1.1, momentum should be reset, $\mu_k = 1$, if the residual does not decrease. However, it is possible that the residual does not decrease even with $\mu_k = 1$, which means that the iterative algorithm reached a *local* rather than *global* minimum of *E* and got stuck there. In such cases the search was simply restarted with a new initial guess.

Equilibria of the compressible model lie in the divergence-free subspace, just like the equilibria of the incompressible model, for which $\nabla \cdot \mathbf{u} = 0$. Hence, in computing the equilibria of the incompressible model, Farazmand [66] used initial guesses that are divergence-free and used a pressure projection to map the velocity field into the divergence-free subspace after every iteration during the adjoint evolution. Since the divergence-free condition is imposed using a projection, rather than as a constraint, the resulting procedure is not guaranteed to descrease the residual. Furthermore, in Farazmand's case computing the pressure that leads to a divergence-free flow was cheap due to the assumption of periodic boundary conditions. The physical no-slip boundary conditions considered here make projection to the divergence-free condition is unnecessary. In the formulation we have derived for the compressible model, according to

(5.3), (5.22), and (5.25),

$$\partial_{\tau} p = -\tilde{\kappa}^{-2} (\nabla_N^2)^{-1} (\partial_x^T F^u + \partial_y^T F^v)$$

$$= \tilde{\kappa}^{-2} (\nabla_N^2)^{-1} (\partial_x F^u + \partial_y F^v).$$
(5.27)

Hence, the divergence-free condition is not enforced, instead the pressure is chosen according to (5.27) at all times. The incompressibility condition is restored once the solution has converged to an equilibrium, where $\partial_{\tau} p = 0$ and hence $\nabla \cdot \mathbf{u} = 0$ according to (5.20).

5.2.3 Results

We used a hybrid adjoint-Newton method introduced by Farazmand [66] to compute the equilibria solutions of the compressible model (5.20). In this method, adjoint evolution equation (5.3) along with the momentum is used to reduce the residual norm $||\mathbf{F}||$ to a certain tolerance, which we chose to be 10^{-3} . The final state produced by the adjoint evolution is then used as an initial condition for the Newton-Krylov solver described in chapter 4. Close to an equilibrium Newton-based methods converge faster (i.e., super-exponentially) than the adjoint-based method, which "only" converge exponentially, according to (5.15). To assess the sensitivity of the hybrid solver to the choice of the initial guess, we performed several tests at various values of the current I when the asymptotic state is turbulent. In both cases, the simulation started with fluid at rest, $\mathbf{u} = 0$ with a small amount of white noise, and the system was evolved forward for 1000 s, so the flow would settle onto the chaotic attractor. In the first test, a state after every 500 seconds was chosen as an initial conditions, if the minimum was below a certain threshold. In both cases 100 initial conditions have been generated. The success rate is summarized in Figure 5.2.

In particular, at I = 24 mA, of the 100 initial conditions in the first test, only three converged to equilibria E_9 , E_{10} and E_{11} shown in Figure 5.3. In the second test, four of



Figure 5.2: The histogram showing the success rate of finding equilibria using the adjoint method, as a function of I.

the 100 initial conditions converged, two to E_9 and the other two to E_{11} . The rest of the 193 initial conditions got stuck at local minima during the adjoint evolution. This is in stark contrast with the results obtained by Farazmand [66], who found that arbitrary initial conditions have converged with 100% success rate to the equilibria of the Kolmogorov flow with periodic boundary conditions. The success rate of the adjoint-based solver apparently depends on numerous factors: model parameters, the size of the physical domain, and, as will be discussed below, on the type of the solutions sought. When it fails, the adjoint-based solver always converged to a local minimum. Developing efficient methods that can avoid local minima should enable construction of a globally convergent solver. This is, however, beyond the scope of this thesis.

5.3 Computing dynamical connections

Although it is simple, the shooting method for computing dynamical connections between unstable equilibria discussed in Section 4.2 has a number of serious drawbacks. First of all, choosing the radii ϵ_{\pm} of the two hyperspheres that the connection originates and terminates


Figure 5.3: Equilibria found using hybrid adjoint-Newton solver at a driving current I = 24 mA. Shown are the vorticity fields of the three equilibria (a) \mathbf{E}_9 , (b) \mathbf{E}_{10} and (c) \mathbf{E}_{11} .

on is far from trivial. They need to be small enough for the linearization to be remain accurate, but large enough for limit the temporal length of the unstable trajectory and the associated exponential amplification of perturbations. The typical lack of a good initial guess required for convergence of Newton iterations is an even greater concern. Below we present an adjoint-based approach for computing dynamical connections that is more robust method and does not have these drawbacks.

5.3.1 Spectral representation

The problems of the shooting method all stem from (i) the infinite temporal length and (ii) the temporal instability of the connection. Both of these problems can be sidestepped by avoiding time-integration altogether, following the approach previously introduced by Liu *et al.* [115] in the context of low-dimensional dynamical systems. They represented the connection between equilibria using a spectral representation on the inifinite time-inteval $t \in \mathbb{R}$. The spectral basis is formed by orthogonal rational functions

$$R_k(t;\omega,t^*) = \cos[k \cot^{-1}(\omega(t-t^*))], \qquad k = 0, 1, 2, \cdots.$$
(5.28)

with some fixed parameters ω and t^* , which are discussed later. The orthogonal rational functions correspond to remapped Chebyshev polynomials, and their properties are discussed in great detail by Boyd [116]. Truncating the spectral expansion to the first $m_0 + 2$ basis functions, for some positive integer m_0

$$\mathbf{X}(t) = \sum_{k=0}^{m_0+1} \mathbf{c}_k R_k(t;\omega,t^*)$$
(5.29)

allows us to represent a (smooth) connection in terms of m_0+2 vector coefficients $\mathbf{c}_k \in \mathbb{R}^N$. Differentiating (5.29) with respect to time gives

$$\dot{\mathbf{X}}(t) = \sum_{k=0}^{m_0+1} \mathbf{c}_k \dot{R}_k(t;\omega,t^*).$$
(5.30)

Let us define the collocation points

$$t_{j} = t^{*} + \frac{1}{\omega} \cot\left(\frac{j\pi}{m_{0}+1}\right), \qquad j = 1, \cdots, m_{0},$$

$$t_{0} = +\infty,$$

$$t_{m_{0}+1} = -\infty.$$
(5.31)

At these points

$$\dot{R}_k(t_j;\omega,t^*) = \omega k \sin^2\left(\frac{j\pi}{m_0+1}\right) \sin\left(\frac{kj\pi}{m_0+1}\right)$$
(5.32)

and

$$\mathbf{X}_{j} \equiv \mathbf{X}(t_{j}) = \sum_{k=0}^{m_{0}+1} \mathbf{c}_{k} \cos\left(\frac{kj\pi}{m_{0}+1}\right)$$
(5.33)

The system (5.33) can be inverted to yield

$$\mathbf{c}_{k} = \frac{2}{(m_{0}+1)d_{k}} \sum_{m=0}^{m_{0}+1} \frac{1}{d_{m}} \cos\left(\frac{mk\pi}{m_{0}+1}\right) \mathbf{X}_{m}$$
(5.34)

where $d_m = 2$ for m = 0 or $m_0 + 1$ and $d_m = 1$ otherwise.

Since the connection is a solution of (2.44), we should have

$$\mathbf{V}_{j} \equiv \mathbf{V} \left(\mathbf{X}_{j} \right) = \sum_{k=0}^{m_{0}+1} \mathbf{c}_{k} \dot{R}_{k}(t_{j}; \omega, t^{*})$$
$$= \omega \sin^{2} \left(\frac{j\pi}{m_{0}+1} \right) \sum_{k=0}^{m_{0}+1} k \mathbf{c}_{k} \sin \left(\frac{k j \pi}{m_{0}+1} \right)$$
$$= \omega \sum_{m=0}^{m_{0}+1} \hat{A}_{jm} \mathbf{X}_{m}$$
(5.35)

for $j = 1, \cdots, m_0$, where

$$\hat{A}_{jm} = \frac{2}{m_0 + 1} \sin^2 \left(\frac{j\pi}{m_0 + 1}\right) \sum_{k=0}^{m_0 + 1} \frac{k}{d_k d_m} \cos\left(\frac{mk\pi}{m_0 + 1}\right) \sin\left(\frac{kj\pi}{m_0 + 1}\right)$$
(5.36)

This system of m_0 equations has to be supplemented with the "boundary conditions"

$$\mathbf{X}_0 = \mathbf{X}_-,$$
$$\mathbf{X}_{m_0+1} = \mathbf{X}_+,$$
(5.37)

where \mathbf{X}_{-} and \mathbf{X}_{+} are the origin and the destination equilibria. Defining

$$\hat{A} = [\mathbf{a}^-, A, \mathbf{a}^+], \tag{5.38}$$

were A is an $m_0 \times m_0$ matrix and \mathbf{a}^- , \mathbf{a}^+ are m_0 -dimensional vectors which depend only

on m_0 , the system (5.35) can be rewritten as

$$\mathbf{V}_{j} = \omega \left(\sum_{m=1}^{m_{0}} A_{jm} \mathbf{X}_{m} + a_{j}^{-} \mathbf{X}_{-} + a_{j}^{+} \mathbf{X}_{+} \right), \qquad j = 1, \cdots, m_{0}.$$
(5.39)

Finally, if we define

$$\tilde{\mathbf{X}} = \begin{bmatrix} \mathbf{X}_1 \\ \vdots \\ \mathbf{X}_{m_0} \end{bmatrix}, \qquad \tilde{\mathbf{V}} = \begin{bmatrix} \mathbf{V}_1 \\ \vdots \\ \mathbf{V}_{m_0} \end{bmatrix}, \qquad \tilde{\mathbf{X}}_{\pm} = \begin{bmatrix} a_1^{\pm} \mathbf{X}_{\pm} \\ \vdots \\ a_{m_0}^{\pm} \mathbf{X}_{\pm} \end{bmatrix}$$
(5.40)

and

$$\tilde{A} = \begin{bmatrix} A_{11}\mathbb{1} & \cdots & A_{1m_0}\mathbb{1} \\ \vdots & \ddots & \vdots \\ A_{m_01}\mathbb{1} & \cdots & A_{m_0m_0}\mathbb{1} \end{bmatrix},$$
(5.41)

the system (5.39) can be rewritten as

$$\tilde{\mathbf{F}}(\tilde{\mathbf{X}}) = \tilde{\mathbf{V}} - \omega \tilde{A}\tilde{\mathbf{X}} - \omega \tilde{\mathbf{X}}_{-} - \omega \tilde{\mathbf{X}}_{+} = 0.$$
(5.42)

This system of equations can be solved using either a Newton-Krylov or an adjointbased method discussed in previous sections, where we have redefined $E = \tilde{\mathbf{F}}^T \tilde{W} \tilde{\mathbf{F}}$ with a block-diagonal weight matrix

$$\tilde{W} = \begin{bmatrix} W & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & W \end{bmatrix},$$
(5.43)

so that (5.3) is replaced with

$$\partial_{\tau} \tilde{\mathbf{X}} = -\tilde{Q} \tilde{J}^T \tilde{W} \tilde{\mathbf{F}},\tag{5.44}$$

where $\tilde{J} = \partial \tilde{\mathbf{F}} / \partial \tilde{\mathbf{X}}$ and

$$\tilde{Q} = \begin{bmatrix} Q & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & Q \end{bmatrix},$$
(5.45)

is a symmetric positive-definite acceleration matrix.

5.3.2 Tuning the spectral representation

The fixed parameters ω and t^* perform a linear transform of the infinite time interval defined by t to $\omega(t-t^*)$ and map the infinite interval back onto itself. The location of the collocation points, however, changes if either the temporal origin t^* or the temporal scale ω is changed. Along with m_0 , these parameters can be used to fine-tune the spectral representation. It is easy to see that m_0 controls the number of the collocation points, ω – (the inverse of) the temporal separation between the collocation points, and t^* – the symmetry of their distribution between the beginning and the end of the connection. The number m_0 of basis functions has to be chosen sufficiently large for the spectral approximation to converge, and the minimal value clearly depends on the problem (e.g., on how smooth the connection is). We have not explored the effect of t^* (all calculations were done using $t^* = 0$).

The scaling factor ω controls the density of the collocation points near the ends of the connection versus its central portion. A poor choice results in numerical instabilities (e.g., spurious oscillations) [115], so one has to be careful in choosing ω . A detailed analysis was performed by Boyd [117], who showed that, for entire functions, the optimal choice of ω depends on the type of the problem, and also showed that the results are insensitive

to choice of ω around the optimal value. It is unclear if Boyd's results are applicable to the class of problems considered here, due to his focus on entire functions. The work of Liu *et al.* [115] also leaves the choice of scaling as an open question. Our own analysis is presented in Section 5.3.5).

5.3.3 Compressible and incompressible model

The formulation described in section 5.3.1 can be directly applied to the compressible model (5.20) as the time derivative exists for all of the three fields - u, v, p, but it needs slight modification for the incompressible model where the first-order ODE for the pressure is replaced with the divergence-free condition $\nabla \cdot \mathbf{u} = 0$. Writing $\tilde{\mathbf{F}}$ defined by equation (5.42) as

$$\tilde{\mathbf{F}} = \begin{bmatrix} \mathbf{F}_1 \\ \vdots \\ \mathbf{F}_{m_0} \end{bmatrix}, \qquad \mathbf{F}_j = \begin{bmatrix} F_j^u \\ F_j^v \\ F_j^p \end{bmatrix}, \qquad (5.46)$$

it can be seen that, for the compressible flow,

$$F_j^p = -\tilde{\kappa}\nabla \cdot \mathbf{u}(t_j) - \sum_{m=1}^{m_0} \omega A_{jm} p(t_m) - \omega a_j^- p_- - \omega a_j^+ p_+, \qquad (5.47)$$

where the last three terms represent $\partial_t p$. For the incompressible model, the expressions F^u and F^v are the same as those for the compressible one (with $\gamma = 0$), but the pressure component (5.47) would change to

$$F_j^p = -\nabla \cdot \mathbf{u}(t_j),\tag{5.48}$$

while the general form of the nonlinear problem (5.42) would remain the same.

Both $\tilde{\mathbf{X}}$ and F are Nm_0 -dimensional vectors, where the dimension of the state space $N = \mathcal{O}(10^6)$ for the spatial discretizations of the model (5.20) considered here. Hence the number of unknowns is comparable to that (Nm_{max}) in the shooting-based methods

considered in the previous Chapter. As discussed there, solving such a large system using an exact Newton method is not practically feasible, although inexact Newton-Krylov methods could again be used. Alternatively, an adjoint-based method presented above can be employed.

5.3.4 Results

In order to test the spectral method, we recomputed the connection between the unstable equilibria \mathbf{E}_3 and \mathbf{E}_4 for I = 9.6 mA. An intentionally poor initial guess was chosen: We used a set of $m_0 = 10$ points uniformly distributed along the straight line (in the state space) joining the origin (\mathbf{E}_3) and the destination (\mathbf{E}_4) and set the scaling parameter to $\omega = 0.1$. Solving the resulting system (5.42) using Newton/GMRES method was not successful. For a Krylov subspace of dimension m = 10 the residual does not decrease noticeably and quickly stagnates, as Figure 5.4a shows. Increasing the dimension of Krylov subspace tenfold to m = 100 decreases the residual somewhat more, but convergence rate slows down dramatically after a few tens of iterations, as Figure 5.4a illustrates. Choosing an even larger Krylov subspace (with dimension as high as m = 300) does not change anything qualitatively, with the residual (nearly) stagnating at high value of the residual.

Given the difficulties experienced by the Newton's method, we have tried to solve the system (5.42) using the adjoint-based method with momentum. The solution $\tilde{\mathbf{X}}$ was evolved using equation (5.44) with the acceleration operators Q and W chosen to be the same as those used for finding the equilibria in Section 5.2. Unlike Newton/GMRES, the adjoint-based method quickly reduces the residual by several orders of magnitude, as shown in Figure 5.4b. Overall, the adjoint-based method decreased the residual by almost five orders of magnitude, compared with the initial condition, to $\|\tilde{\mathbf{F}}\| = 10^{-4}$.

Eventually the residual stagnates, which is expected, because a small number of basis functions does not allow a sufficiently accurate representation of the connection. This can be easily seen from the "spectrum" $\|\mathbf{c}_k\|$ of the converged solution shown in Figure 5.5a.



Figure 5.4: Convergence of (a) Newton/GMRES method and (b) the adjoint method for the heteroclinic connection between E_3 and E_4 at I = 9.6 mA. The Krylov subspace dimension (10, 100, 200, or 300) is shown in the figure legend.

The spectral representation of the connection for $m_0 = 10$ is clearly under-resolved. We repeated the calculation with $m_0 = 41$ and $m_0 = 79$. As Figure 5.5a shows, both solutions are well-resolved, with the last spectral coefficients associated with the highest "wavenumber" being between six and seven orders of magnitude smaller than the first spectral coefficient associated with the lowest "wavenumber". The nearly exponential decay of spectral coefficients clearly demonstrates the advantages of the spectral formulation. This is in stark contrast with the finite-difference scheme used to represent the time derivatives by Dong and Lan [70], which required $O(10^3)$ collocation points, compared with a few tens of collocation points for the spectral formulation.

The sparse set of collocation points does not allow reconstructing the connection, which is a continuous trajectory, at all times. However, there is a one-to-one correspondence (5.33) between the collocation points X_j and the spectral coefficients c_k . Once the solution \tilde{X} has been obtained, the spectral coefficients can be calculated and the entire trajectory representing the dynamical connection can be computed at any time t using the spectral expansion (5.29). The results are plotted in Figure 5.5b, where a two-dimensional projection of the connection computed using the adjoint-based method is compared with the result of the shooting-based calculation discussed in Section 4.2. Once again, we see that $m_0 = 10$ produces an under-resolved result, while $m_0 = 41$ and $m_0 = 79$ produce a connection which is indistinguishable from that computed using shooting.

Figure 5.5b also illustrates and compares the distribution of the collocation points describing the three choices of parameters. For $m_0 = 10$ and $\omega = 0.1$ the density of collocation points is clearly too low. The distribution is also strongly skewed towards the destination equilibrium E_4 . A more balanced distribution could, in principle, be obtained by using a nonzero value of the parameter t^* . Choosing $m_0 = 41$ or $m_0 = 79$ and $\omega = 0.01$ produces the distribution of collocation points that is highly nonuniform, with a much higher density of points near the ends of the connection. This is crucial for achieving a small value of the residual, since the connection has fine structure near the ends as shown in Figure 5.6.



Figure 5.5: (a) Spectral coefficients for the connection between \mathbf{E}_3 and \mathbf{E}_4 for various choices of m_0 and ω . (b) A 2-D projection of the connection onto the plane spanned by the two unstable eigenvectors \mathbf{e}_1 and \mathbf{e}_2 of the origin \mathbf{E}_3 , with the continuous green line representing the result of shooting and the symbols showing the collocation points for the spectral representation. Both (a) and (b) use the same color-coding: $m_0 = 10$ and $\omega = 0.1$ (blue), $m_0 = 41$ and $\omega = 0.01$ (red), $m_0 = 79$ and $\omega = 0.01$ (black).



Figure 5.6: Zoomed view of the connection (a) near the origin E_3 and (b) near the destination E_4 at I = 11.31 mA. The connection computed using the spectral method is shown in blue and the connection computed using the shooting method is shown in green.



Figure 5.7: Stability spectra of the origin equilibrium (a) \mathbf{E}_3 and the destination equilibrium (b) \mathbf{E}_4 at I = 11.31 mA.

As can be seen in the figure, this fine structure is only partially resolved by the shooting method.

The fine structure of the connection near the origin and destination equilibrium can be understood using their stability spectra shown in Figure 5.7. In particular, the destination equilibrium E_4 undergoes Hopf bifurcation at a driving current (of 11.78 mA) slightly above the value at which the connection was computed (11.31 mA). Correspondingly, the stability spectrum features a pair of weakly stable complex conjugate eigenvalues, which explains the spiraling behavior of the connection in the neighborhood of E_4 observed in Figure 5.6b.

On the other hand, the origin E_3 has two real unstable eigenvalues, so we should expect the connection to align in the direction of one of the two unstable eigenvectors, which is consistent with the shape of the connection computed using the shooting method. The spiraling structure of the connection computed using the spectral representation (cf. Figure 5.6a) is in all likelihood a result of over-resolving this region. Indeed, near the origin both \dot{X} and V(X) are very small, and so is their difference, even if these two vectors point in completely opposite directions at the corresponding collocation points. This artifact illustrates that care should be taken in choosing the parameters ω and t^* which control the



Figure 5.8: Projections of the heteroclinic connection between the origin E_7 (blue sphere) and destination E_8 (red sphere). The connection computed using the symmetry constraint is shown as the green curve. The collocation points on the connection computed using the spectral/adjoint method are shown as black asterisks. The same projection as in Figure 4.6 is used here.

clustering of collocation points.

As another test of the spectral/adjoint method, we used it to compute the connections between equilibria E_7 and E_8 for which the shooting method fails (without the symmetry constraint) due to the disparity in the magnitude of the unstable eigenvalues of the origin E_7 . The initial guess was again chosen as a set of uniformly distributed collocation points along the line joining the origin and the destination. The adjoint evolution quickly converged from this initial condition to the connection, with the result shown in Figure 5.8, which shows a low dimensional projection of the connections computed using various methods. As one can easily see, the spectral/adjoint method finds the same solution as the one obtained by the symmetry-constrained shooting method.

5.3.5 Optimal choice of the scaling parameter

When m_0 is chosen sufficiently large, the accuracy of the adjoint-based method can be optimized by choosing the scaling parameter ω [115, 117]. Recall that the inverse of ω determines the temporal separation between the collocation points, so that choosing too large a value of ω will lead to the collocation points clustering in the central portion of the connection, while too small a small value will cause the collocation points to cluster near the ends of the connection. A poor choice of ω will lead to either the ends or the central portion of the connection to become under-resolved, which will cause the residual to stagnate. We have investigated the influence of ω on the computation of the connections between equilibria \mathbf{E}_3 and \mathbf{E}_4 at I = 9.6 mA for $m_0 = 41$. The results are shown in Figure 5.9, which illustrates that choosing $\omega = 0.01$ results in the residual decreasing below the threshold (here $\|\tilde{\mathbf{F}}\| = 5 \times 10^{-5}$), while choosing a value much lower or higher than the optimal causes stagnation of the residual before reaching the threshold.

Given that $\|\tilde{\mathbf{F}}\|$ possesses a minimum for some optimal value of ω , this optimal value can be determined along with the solution $\tilde{\mathbf{X}}$ by allowing ω to vary with τ . This yields a new minimization problem with respect to $\tilde{\mathbf{X}}$ and ω , which can again be solved using the adjoint-based method. The procedure is described in Section 5.4 below, where computation of periodic orbits is discussed, in which case ω represents the circular frequency of the orbit.

5.3.6 Discussion

Let us conclude this section with a discussion of the general properties of the spectral/adjoint method in comparison with alternative approaches to computing dynamical connections.

- As we have shown, the spectral formulation coupled with the adjoint method offers a general and efficient approach for computing heteroclinic connections between unstable equilibria. Unlike Newton's method, this method does not need a good initial guess to start with.
- 2. A similar approach can be used for computing heteroclinic connections between other types of unstable solutions. When the origin and/or the destination is a periodic orbit, not only its temporal duration, but also the arc length of the connection in the state space becomes infinitely long. This can be handled by clipping the connection at a Poincare section once it enters the linear neighborhood of the origin/destination periodic orbit, similar to Section 4.2.1.



Figure 5.9: The effect of the scaling parameter ω on the accuracy of the spectral representation for the connection between \mathbf{E}_3 and \mathbf{E}_4 at I = 9.6 mA. (a) The evolution of the residual. (b) The decay of spectral coefficients. The values of ω used are 0.01 (blue), 0.05 (red), and 0.1 (black); in all cases $m_0 = 41$.

- 3. Since it does not involve temporal integration, this method can handle even highly unstable connections. This has been demonstrated in this thesis by computing the connection between equilibria E_7 and E_8 for which the shooting method, which relies on temporal integration, fails. While, in this particular example, the shooting method can be "salvaged" by applying a symmetry constraint, in general this is not possible.
- 4. There are two different ways to parallelize the calculations. One relies on temporal decomposition, where each compute node performs all the calculations for one collocation point. Indeed, the state space velocity V_j = V(X(t_j)) at each collocation point can be evaluated independently. The calculation of the product Q̃J^TW̃F in (5.44) can also be trivially parallelized. However the computation of F and matrix vector product with J^T requires the flow states X_j at all collocation points. Assembling this information is the bottleneck of the temporal decomposition.
- 5. An alternative parallelization strategy relies on domain decomposition, where each compute node performs all the calculations for all the collocation points in a sub-region of the physical domain. Since finite differences have been used for spatial discretization, computing the relevant components of the state space velocity vector \tilde{V} only requires exchanging the information along the boundary of the corresponding sub-region, which should be much faster.
- 6. The calculation of \tilde{J}^T is straightforward for arbitrary boundary conditions when finite differences are used for spatial discretization of the model. For other types of discretization, computing \tilde{J}^T , especially for non-periodic boundary conditions, may be rather nontrivial. This is however outside of the scope of this thesis.
- 7. Spectral representation allows parameterization of the entire connection using relatively few spectral coefficients, making this approach more memory-efficient than single- or multi-shooting methods coupled with Newton/GMRES iterations. By dis-

tributing these collocation points across multiple computing nodes, connections can be computed even for memory demanding, high-dimensional systems such as 3D fluid flows on large computational domains.

- 8. Unlike the case of equilibria, the computation of heteroclinic connections never got stuck at a local minimum. The likely reason for this is that the problem of computing the connection is substantially more constrained (due to the requirement that the connection begin at the origin equilibrium and terminate at the destination equilibrium). As a result, at least for relatively short connections, the chances of local minima of *E* emerging are much slimmer.
- 9. Although we focused entirely on heteroclinic connections in this section, the same approach clearly applies to computing homoclinic connections, where the origin and the destination equilibria are the same. The only substantial difference concerns the choice of the initial guess, which is far from obvious for a homoclinic connection. Homoclinic connections are far less common, and in all likelihood a search for one will only be performed if there is concrete numerical evidence, e.g., near-recurrence that can be used to initialize the method.

5.4 Computing periodic orbits

The spectral formulation described in Section 5.3.1 can be trivially generalized to periodic orbits. Due to temporal periodicity of the corresponding solutions, the basis of Fourier modes is more appropriate. Let $\mathbf{X}_j = \mathbf{X}(t_j)$ with $t_j = jT/m_0$ be the locations of the collocation points, where $j = 1, \dots, m_0$ and T is the (unknown) period of the orbit. As shown in Appendix C, the periodic orbit corresponds to a root of the system

$$\tilde{\mathbf{F}}(\tilde{\mathbf{X}};\omega) = \tilde{\mathbf{V}} - \omega \tilde{A} \tilde{\mathbf{X}} = 0,$$
(5.49)

where $\tilde{\mathbf{V}}$ and $\tilde{\mathbf{X}}$ are again given by (5.40), the constant matrix \tilde{A} depends only on m_0 (cf. Appendix C), and $\omega = 2\pi/T$ is the associated frequency, the inverse of which again determines temporal separation between the collocation points. In practice, rather than solving for the root of (5.49) we compute the minimum of $E = \tilde{\mathbf{F}}^T \tilde{W} \tilde{\mathbf{F}}$. Similar to the problem of finding optimal ω in the case of heteroclinic connections, to compute the periodic orbit, in addition to solving for $\tilde{\mathbf{X}}$, one also needs to solve for ω . This can be accomplished in several different ways described below:

1. Writing the adjoint evolution equation for $\tilde{\mathbf{F}}(\tilde{\mathbf{X}};\omega)$ defined by (5.49) results in the the following equations for $\tilde{\mathbf{X}}$ and ω :

$$\partial_{\tau} \tilde{\mathbf{X}} = -\tilde{Q} \tilde{J}^T \tilde{W} \tilde{\mathbf{F}}$$
(5.50)

$$\partial_{\tau}\omega = -R\tilde{\mathbf{Y}}^{T}\tilde{W}\tilde{\mathbf{F}},\tag{5.51}$$

where $\tilde{\mathbf{Y}} = \tilde{A}\tilde{\mathbf{X}}$ and R is an arbitrary positive constant. The two unknowns, $\tilde{\mathbf{X}}$ and ω , can be evolved in pseudo-time in parallel using the above equations.

Since F̃ is linear in ω, E is a quadratic function of ω for a fixed X̃. The minimum of E is achieved for

$$\omega = \frac{\tilde{\mathbf{Y}}^T \tilde{W} \tilde{\mathbf{V}} + \tilde{\mathbf{V}}^T \tilde{W} \tilde{\mathbf{Y}}}{2 \tilde{\mathbf{Y}}^T \tilde{W} \tilde{\mathbf{Y}}}.$$
(5.52)

Hence, $\tilde{\mathbf{X}}$ can be evolved using equation (5.50) with ω chosen according to (5.52) at every pseudo-time step.

3. Once the root of (5.49) has been found, we should have $\tilde{F}_k = 0$ and therefore

$$\omega = \frac{V_k}{\tilde{Y}_k} \tag{5.53}$$

for all k. So ω can be eliminated from (5.49) by redefining $\tilde{\mathbf{F}}$ as

$$\tilde{\mathbf{F}}(\tilde{\mathbf{X}}) = \tilde{\mathbf{V}} - \frac{\tilde{V}_k}{\tilde{Y}_k} \tilde{\mathbf{Y}}$$
(5.54)

where the index k is chosen such that $|\tilde{Y}_k|$ is bounded away from zero. The resulting system can be solved using the adjoint evolution equation (5.50).

Using one of the formulations described above, a periodic orbit can be found by solving (5.49) using the corresponding adjoint evolution equation(s). Convergence can be further sped up by switching to Newton/GMRES iterations once the residual decreases sufficiently. Preliminary investigation shows that all three of the above options allow periodic orbits to be computed for low-dimensional systems (we tested them on the Lorenz system). However, we have not explored any of these options in the context of fluid flows thoroughly enough to meaningfully compare them.

5.5 Summary

An adjoint-based approach allows the problem of computing a variety of different types of solutions (equilibria, periodic orbits, heteroclinic and homoclinic connections) to be formulated in a systematic and extremely flexible way. It leads to an evolution equation which involves the adjoint of the Jacobian, whose action on an arbitrary vector, unlike that of the inverse of the Jacobian, can be explicitly and relatively easily computed. As a result, adjoint-based methods are both faster, and require less memory, than Newton-Krylov methods.

The adjoint-based method for computing equilibria and periodic orbits has not been fully tested or optimized. For instance, we have not systematically explored the choice of the weight matrix W and the acceleration matrix Q. This choice could be exploited to both speed up convergence and avoid local minima in the case of equilibria. Similarly, neither the choice of the preconditioner matrices nor the choice of the evolution equation for the frequency ω has been explored for periodic orbits.

The main focus of this Chapter is on computing heteroclinic connections, for which no established systematic methods had existed previously. Combining the adjoint method with spectral representation yields a robust and efficient way to compute connections between equilibria, which, unlike Newton/shooting method, does not require a good initial guess. Furthermore, since there is no time-integration involved, the spectral/adjoint method can be used to compute even highly unstable connections.

CHAPTER 6 CONCLUSIONS

This thesis makes a number of contributions towards developing a deterministic, geometrical description of turbulent fluid flows. The most significant ones are summarized below.

6.1 Main scientific contributions

First of all, a new 2D model has been derived, and implemented numerically, to describe the Kolmogorov-like flow in the experiment at moderate Re. The model offers a dramatic simplification of the problem, compared with the computationally much more expensive full 3D model of the flow, making possible calculations that would otherwise be essentially intractable. Although approximate, the model improves on the one introduced previously [71] by taking into account the small vertical component of the velocity and the associated deflection of the free surfaces of the two layers (dielectric and electrolyte), which makes the horizontal flow field weakly compressible. This compressible model is in better agreement with the experiment than the incompressible one, especially at higher Re. In particular, numerical simulations based on the compressible model faithfully reproduce such experimental details as a hysteretic transition between periodic and turbulent flow featuring quasiperiodic dynamics and a periodic window inside the turbulent regime.

Another significant contribution is the development of several new approaches for computing dynamical connections between ECS. In particular, we have developed a procedure that identifies at least some of the ECS a particular one is connected to (either as an origin or as a destination) by performing a parameter continuation until a bifurcation where the two ECS merge. This procedure provides a good initial guess for the connection to be identified (it coincides with the center manifold) and the connection is computed by refining this initial guess using Newton iterations. Finally, both ECS and the connection between them can be continued back to the original value(s) of the paramater(s). This approach can fail, however, when the connection is aligned along the near-marginal direction at the origin, due to strong amplification of numerical noise by other unstable directions. We have shown that, when both ECS lie in the same symmetric subspace, such problems can be circumvented by restricting the connection to lie in that subspace.

To address these difficulties, we have developed a more general and robust approach which completely avoids time-integration and the associated stability issues, by combining a spectral representation of the connection in terms of mapped Chebyshev polynomials [115] with an adjoint-based solver [66]. In this approach, a continuous curve is evolved in pseudo-time to gradually align its tangent at a set of collocation points with the direction of the vector field, until it converges to the connection which is everywhere parallel to the vector field, which is similar in spirit to the variational approach [64] for computing periodic orbits. The spectral discretization affords a substantial reduction in the number of degrees of freedom, compared with the finite-difference discretization proposed by Dong et al [70] and allows a continuous connection to be easily reconstructed without resorting to time-integration. Furthermore, we showed that the adjoint-based formulation also allows determining the optimal parameters, such as the scaling factor ω in a systematic manner, solving a problem which was not properly addressed previously [117, 115]. We demonstrated that this approach successfully finds every known unstable dynamical connection without requiring a good initial guess to begin with. Furthermore, it can be easily generalized to find time-peroidic solutions.

6.2 Open questions

While the compressible two-dimensional model of the flow shows good qualitative agreement with the experiment in the pre-turbulent regime, it fails to show perfect quantitative agreement, especially at higher Re. The quantitative discrepancies are not intrinsic to the model, but result from some assumptions made in deriving a reduced two-dimension description of what is nominally a three-dimensional flow. In particular, the dimensional model features several parameters (α , β , γ , ν , and κ) obtained by depth-averaging the velocity ansatz (2.6) with the vertical profile function q(z) which weakly depends on the horizontal velocity profile $\mathbf{u}(x, y, t)$. The horizontal flow profile changes rather substantially as Re increases, so it is quite reasonable to expect that q(z), and hence all of the above parameters, should also change compared with the values evaluated for the straight-flow solution \mathbf{E}_1 at low Re. The values of these parameters matching a particular experimental setup can be obtained, for instance, using a variation of the sparse identification method proposed by Rudy *et al.* [97]. That method assumes that *all* state variable are observable. However, in the compressible model, the pressure p representing with the distortion of the free surfaces is not directly observable, so not all of the library functions representing various terms in the evolution equations (2.30)-(2.31) can be reconstructed from the experimental flow field $(\mathbf{u}(x, y, t))$ measurements. This can be addressed by using the gradient of (2.31) to eliminate the term $\partial_t \nabla p$ from the time derivative of (2.30), yielding an equation that only involves \mathbf{u} and incorporates all of the unknown parameters:

$$\partial_t^2 \mathbf{u} = -\beta \partial_t [(\mathbf{u} \cdot \nabla_{\parallel})\mathbf{u}] - \gamma \partial_t [(\nabla \cdot \mathbf{u})\mathbf{u}] + \kappa \nabla (\nabla \cdot \mathbf{u}) + \nu \partial_t \nabla_{\parallel}^2 \mathbf{u} - \alpha \partial_t \mathbf{u}.$$
(6.1)

With a proper choice of the parameters, the validity of the model can likely be extended to much higher Re.

The compressible model, with the proper choice of parameters, can be used to describe the subcritical transition to turbulence in the experimental Kolmogorov-like flow. Our investigation has shown that the transition involves a stable quasi-periodic state with two frequencies, which corresponds to a 2-torus in the state space. The two frequencies are roughly in 1:8 ratio and numerical simulations show that the Poincare section of this 2-torus breaks up into eight disconnected components. While some sort of resonant interaction between the two frequencies (like in KAM theory [118]) appears likely, it is unclear what the exact mechanism is that leads to this breakup. Understanding it would shed new light on the different mechanisms that could generate chaotic saddles underlying turbulent dynamics of fluid flows.

We have demonstrated that dynamical connections can be computed using the spectral representation even when they have a rather complex structure near the origin and/or the destination ECS. However, if the connection has complex structure at multiple locations along the connection, this basis may not be optimal. One possible workaround involves dividing the connection into several segments, with a separate spectral expansion for each segment [119]. Such an approach has several advantages, e.g., it can be exploited to construct a numerical algorithm which can be parallelized more efficiently. Since the computation on each of the segments of the connection only requires the information from the neighboring segments, the communication overhead is reduced considerably compared with the global basis considered in this thesis. The main challenge here is incorporating the missing the boundary conditions at the ends of the neighboring segments.

The global spectral representation yields a formulation that can be solved relatively efficiently using an adjoint-based method. However, computing the adjoint of the Jacobian is not always easy, possible, or convenient. Newton-Krylov methods, which do not require computing the adjoint, were found to quickly stagnate for global spectral representation. Segmenting the connection into multiple pieces appears to help in the sense that Newton iterations can be used to solve the resulting system, at least in low-dimensional cases [119]. It is worth exploring whether this strategy also allows Newton-Krylov methods to succeed for high-dimensional systems resulting from discretization of fully resolved models of fluid flows.

As discussed previously, the main problem of the shooting methods is their inherent instability associated with time-integration. However, this problem can also be circumvented by using multiple shooting [120]. The idea of this approach is to introduce multiple Poincare sections along the connection, which limits the growth of the perturbations when a trajectory is time-integrated between one section and the next. Just like the segmented spectral approach, multi-shooting formulation (i) yields a system of equations which can be solved using Newton-Krylov methods and (ii) the numerical implementation can be relatively efficiently parallelized [121]. Since these alternative methods have never been directly compared, it is presently unclear which of them is the best in terms of the speed and memory requirements and how robust they are with respect to the choice of the initial guess.

The problem of computing connections between ECS one (or both) of which is a timeperiodic (or quasiperiodic) state has not been considered in this thesis. The main difficulty in this case is that the connection becomes infinitely long not just temporally, but also in terms of its arclength. Hence, the connection has to be broken up into three pieces: the infinitely long "ends" which can be computed analytically using linearization about the origin/destination and the finite-length "middle" which is computed numerically. Such an approach has been used, for instance, to compute connections between an equilibrium and a time-periodic state [122] and between a pair of time-periodic states [123] using the singleshooting method. The same approach can be used if the central portion of the connection is computed using multi-shooting or spectral representation.

One substantial advantage of the adjoint-based solver is its extreme generality. As we have demonstrated, it can be used to compute a variety of different types of solutions, from equilibria to periodic solutions, to homo/heteroclinic connections. There are some hurdles though that still must be overcome, for this method to become truly useful. As discussed previously, the adjoint integration can get stuck at a local minimum, i.e., $\partial_{\tau} E = 0$ when $J^T W \mathbf{F} = 0$ even though $\mathbf{F} \neq 0$. This means that $W \mathbf{F}$ lies in the kernel of J^T and, rather than choosing W = 1 and restarting the search from a new initial condition, as we have done, one can simply change the weight matrix W such that $W \mathbf{F}$ is *not* in the kernel of J^T . The main challenge is choosing the operator $W(\tau)$ dynamically, which achieves this objective while retaining its positive-definite property.

Another challenge is choosing the acceleration operator Q to eliminate the slow modes which result in slow convergence. A simple choice of acceleration operator that aims to speed up convergence has been proposed and tested in this thesis. While constant and inexpensive to apply, this operator is *not* optimal, as the simple example considered in Section 5.1.2 illustrates. A more efficient approach would involve constructing a non-constant acceleration operator $Q(\mathbf{X}(\tau))$ using the linearization of the model about the current, evolving state $\mathbf{X}(\tau)$, which remains positive-definite, just like in the case of the weight W. Appendices

APPENDIX A

SPATIAL DISCRETIZATION AND NUMERICAL INTEGRATION

A.1 Gauge condition

In the evolution equations (2.36) and (2.37), the pressure p appears only under the sign of the gradient. It means that adding an arbitrary constant to p does not change the evolution equations, which results in the Jacobian of the evolution equation being singular. This will cause a problem in numerically evolving the equations, for instance, if implicit scheme has been used to update the state. In order to eliminate this singularity, we impose the gauge condition

$$p(x_0, y_0) = 0, (A.1)$$

where (x_0, y_0) is an arbitrary spatial location. With this gauge condition, the evolution equations are transformed to

$$\partial_{t}\mathbf{u} + \beta \left(\mathbf{u} \cdot \nabla_{\parallel}\right)\mathbf{u} + \gamma (\nabla_{\parallel} \cdot \mathbf{u})\mathbf{u} = \tilde{R}e^{-1} \left(\nabla_{\parallel}^{2}\mathbf{u} - \tilde{\alpha}\mathbf{u}\right) + \bar{f}_{0}\hat{x} - \nabla_{\parallel}p, \quad (A.2)$$

$$\partial_t p = -\tilde{\kappa} \left(\nabla_{\parallel} \cdot \mathbf{u} - \nabla_{\parallel} \cdot \mathbf{u} \Big|_{(x_0, y_0)} \right)$$
(A.3)

A.2 Spatial discretization

The spatial discretization scheme, the grid parameters, and the method of computing spatial derivatives described here has been originally developed and thoroughly tested by Balachandra Suri [62] for the simulations of 2D incompressible Navier-Stokes on a rectangular domain given by equations (2.4). We have slightly modified it to fit to the weakly compressible 2D Navier-Stokes in this thesis.

For the simulation of the transformed evolution equations given by (A.2) and (A.3), the fields (\mathbf{u},p) have been discretized using finite differences on a rectangular "staggered"



Figure A.1: Schematic of the staggered grid with $n_x = 3$, $n_y = 4$. The locations of the fields u,v,p are respectively shown using blue circle, cross and dot, and the boundaries of the domain are represented by the red lines. The figure also shows the ghost points that are defined outside the physical domain.

grid or marker and cell (MAC) grid, proposed by Harlow *et al* [124]. In staggered grid configuration, the domain is divided in to $n_x \times n_y$ identical, rectangular cells with each cell having dimensions of $\delta x \times \delta y$. *p* field is defined at the cell centers whereas the *x* and *y* components of u, respectively *u*, *v*, are defined on the edges of the cell. A schematic of the distribution of the fields on the grid is shown in Figure A.1. We have used 20 points for unit non-dimensional length along *x* and *y* directions, i.e., $\delta x = \delta y = 1/20$. This would give $n_x = 280$ and $n_y = 360$ for the dimensions of the experimental setup considered in this thesis.

Shifting the origin of the coordinate system to the bottom left corner of the domain and using (i, j) to denote the location $(i\delta x, j\delta y)$, it can be seen from the Figure A.1 that the u, v, p are defined respectively at (i, j + 1/2), (i + 1/2, j) and (i + 1/2, j + 1/2) for integers i and j. Since these fields are defined at different locations in space, one needs to do interpolation to get the contribution from one field to the evolution of another field. For example, the term $\left(
abla _{\parallel} \cdot \mathbf{u}
ight) u$ required for evolution of u is computed as

$$\left(\nabla_{\parallel}\cdot\mathbf{u}\right)u\big|_{(i,j+\frac{1}{2})} = \frac{1}{2}\left[\left(\nabla_{\parallel}\cdot\mathbf{u}\right)\big|_{(i+\frac{1}{2},j+\frac{1}{2})} + \left(\nabla_{\parallel}\cdot\mathbf{u}\right)\big|_{(i-\frac{1}{2},j+\frac{1}{2})}\right]u\left(i,j+\frac{1}{2}\right) \quad (A.4)$$

where $\nabla_{\parallel} \cdot \mathbf{u}$ is computed using the nearest neighbors as

$$\nabla_{\parallel} \cdot \mathbf{u} \Big|_{(i+\frac{1}{2},j+\frac{1}{2})} = \frac{u(i,j+\frac{1}{2}) - u(i-1,j+\frac{1}{2})}{\delta x} + \frac{v(i+\frac{1}{2},j) - v(i+\frac{1}{2},j-1)}{\delta y}.$$
(A.5)

The term required for temporal evolution of p is computed using the nearest neighbors and the gauge condition as

$$\partial_t p|_{(i+\frac{1}{2},j+\frac{1}{2})} = -\tilde{\kappa} \left(\nabla_{\parallel} \cdot \mathbf{u} \Big|_{(i+\frac{1}{2},j+\frac{1}{2})} - \nabla_{\parallel} \cdot \mathbf{u} \Big|_{(i_0+\frac{1}{2},j_0+\frac{1}{2})} \right), \tag{A.6}$$

where $(i_0 + \frac{1}{2}, j_0 + \frac{1}{2})$ is the location that fixes the gauge for p. The laplacian term is computed using the standard 5-point stencil as

$$\nabla_{\parallel}^{2} u \Big|_{(i,j+\frac{1}{2})} = \frac{u(i-1,j+\frac{1}{2}) - 2u(i,j+\frac{1}{2}) + u(i+1,j-\frac{1}{2})}{\delta x^{2}} + \frac{u(i,j-\frac{1}{2}) - 2u(i,j+\frac{1}{2}) + u(i,j+\frac{3}{2})}{\delta y^{2}}.$$
(A.7)

To compute the laplacian of u at the locations adjacent to left and right boundaries, we need its values on these boundaries. These are set to zero owing to the physical no-slip boundary conditions of the velocity, i.e.,

$$u\left(0, j + \frac{1}{2}\right) = u\left(n_x, j + \frac{1}{2}\right) = 0.$$
 (A.8)

To compute laplacian of u at the locations adjacent to top and bottom boundaries, ghost points are defined outside these boundaries and the boundary conditions for the ghost points are chosen such that no-slip boundary conditions are satisfied at the bottom and top boundaries, i.e., u at the ghost points satisfies

$$u(i,0) = \frac{u(i,-\frac{1}{2})+u(i,+\frac{1}{2})}{2} = 0 \quad \Rightarrow u\left(i,-\frac{1}{2}\right) = -u\left(i,+\frac{1}{2}\right)$$
$$u(i,n_y) = \frac{u(i,n_y+\frac{1}{2})+u(i,n_y-\frac{1}{2})}{2} = 0 \quad \Rightarrow u\left(i,n_y+\frac{1}{2}\right) = -u\left(i,n_y-\frac{1}{2}\right).$$
(A.9)

Similarly, Neumann boundary conditions for p are implemented using the ghost points defined outside the boundary as

$$p\left(-\frac{1}{2}, j + \frac{1}{2}\right) = p\left(\frac{1}{2}, j + \frac{1}{2}\right),$$

$$p\left(n_x + \frac{1}{2}, j + \frac{1}{2}\right) = p\left(n_x - \frac{1}{2}, j + \frac{1}{2}\right),$$

$$p\left(i + \frac{1}{2}, -\frac{1}{2}\right) = p\left(i + \frac{1}{2}, \frac{1}{2}\right),$$

$$p\left(i + \frac{1}{2}, n_y + \frac{1}{2}\right) = p\left(i + \frac{1}{2}, n_y - \frac{1}{2}\right).$$
(A.10)

The non-linear term $(u\partial_x + v\partial y)u$ can be computed using central difference or upwind scheme. Central difference scheme has less numerical dissipation but is unstable and upwind scheme is stable but has high numerical dissipation. So, the non-linear term is computed using a blending of central and upwind schemes, to include the merits of both the schemes, and the blending parameter b, which is the fraction of contribution from upwind scheme, is chosen as

$$b = \min(1.2\,C, 1) \tag{A.11}$$

where the CFL number C is given by

$$C = \max\left(\|u\|_{\infty}\frac{\delta t}{\delta x}, \|v\|_{\infty}\frac{\delta t}{\delta y}\right)$$
(A.12)

and $\|\cdot\|_{\infty}$ denotes the infinity or max-norm. For the numerical stability, we require the CFL number to be less than 1. Since u, v are $\mathcal{O}(1)$ and $\delta x = \delta y = 1/20$, this condition would translate to the non-dimensional time step δt being less than 1/20. For the range of driving currents considered here, this would mean the dimensional time step should be less than 0.1. As a safety factor, all the simulations in this thesis are carried out with a dimensional time step of 1/32, which corresponds to a CFL number less than 0.6.

The procedure thus described to compute the spatial derivatives required for evolution of u can be similarly extended to compute the spatial derivatives required for evolution of v. The evolution equations are written only inside the physical domain and the values of the fields u, v and p only inside the domain are stored thus making their sizes respectively $n_y \times (n_x - 1), (n_y - 1) \times n_x$ and $(n_y - 1) \times (n_x - 1)$. The ghost points are calculated using the interior of the fields as and when required. Once the spatial derivatives have been computed using this discretization, the system is integrated using IMEX scheme as described in the following section.

A.3 Numerical integration

The evolution equations have been integrated forward in time using Implicit-Explicit scheme (IMEX), in which the stiff linear terms are evolved implicitly and the non-linear terms are evolved explicitly using 2nd order Adams-Bashforth method, as described below.

The transformed evolution equations (A.2) and (A.3) can be written as

$$\partial_t \mathbf{u} = \mathcal{L}(\mathbf{u}) + \mathcal{N}(\mathbf{u}) + \bar{f}_0 \hat{x} - \tilde{\kappa} \nabla_{\parallel} p$$
 (A.13)

$$\partial_t p = -\left(\nabla_{\parallel} \cdot \mathbf{u} - \nabla_{\parallel} \cdot \mathbf{u} \Big|_{(i_0 + \frac{1}{2}, j_0 + \frac{1}{2})}\right).$$
(A.14)

where the linear and non-linear terms, respectively $\mathcal{L}(\mathbf{u})$ and $\mathcal{N}(\mathbf{u})$, are given by

$$\mathcal{L}(\mathbf{u}) = \tilde{Re}^{-1} \left(\nabla_{\parallel}^{2} \mathbf{u} - \tilde{\alpha} \mathbf{u} \right) = \tilde{Re}^{-1} \left(\nabla_{\parallel}^{2} - \tilde{\alpha} \right) \mathbf{u}$$
(A.15)

$$\mathcal{N}(\mathbf{u}) = -\beta \left(\mathbf{u} \cdot \nabla_{\parallel} \right) \mathbf{u} - \gamma (\nabla_{\parallel} \cdot \mathbf{u}) \mathbf{u}.$$
(A.16)

Using implicit evolution for the linear terms, 2nd order Adams Bashforth for the nonlinear terms and forward difference for the temporal derivative, equations (A.13) and (A.14) become

$$\frac{\delta \mathbf{u}^{n}}{\delta t} = \mathcal{L}(\mathbf{u}^{n}) + \mathcal{L}(\delta \mathbf{u}^{n}) + 1.5\mathcal{N}(\mathbf{u}^{n}) - 0.5\mathcal{N}(\mathbf{u}^{n-1}) + \bar{f}_{0}\hat{x} - \nabla_{\parallel}p^{n} - \nabla_{\parallel}\delta p^{n}$$
(A.17)

$$\frac{\delta p^{n}}{\delta t} = -\tilde{\kappa} \left(\nabla_{\parallel} \cdot \mathbf{u}^{n} + \nabla_{\parallel} \cdot \delta \mathbf{u}^{n} - \nabla_{\parallel} \cdot \mathbf{u}^{n} \Big|_{(i_{0} + \frac{1}{2}, j_{0} + \frac{1}{2})} \right).$$
(A.18)

where (\mathbf{u}^n, p^n) are the fields at the n^{th} time step and $(\delta \mathbf{u}^n, \delta p^n)$ are the changes in the fields at this time step. Equations (A.17) and (A.18) can be combined and written as a single matrix equation for the unknown $(\delta \mathbf{u}^n, \delta p^n)$ as

$$\begin{bmatrix} \mathbb{1} - \delta t \mathcal{L} & \delta t \nabla_{\parallel} \\ \tilde{\kappa} \delta t \nabla_{\parallel} \cdot & \mathbb{1} \end{bmatrix} \begin{bmatrix} \delta \mathbf{u}^n \\ \delta p^n \end{bmatrix} = \begin{bmatrix} \mathcal{L}(\mathbf{u}^n) + 1.5\mathcal{N}(\mathbf{u}^n) - 0.5\mathcal{N}(\mathbf{u}^{n-1}) + \bar{f}_0 \hat{x} - \nabla_{\parallel} p^n \\ -\tilde{\kappa} \left(\nabla_{\parallel} \cdot \mathbf{u}^n - \nabla_{\parallel} \cdot \mathbf{u}^n \big|_{(i_0 + \frac{1}{2}, j_0 + \frac{1}{2})} \right) \end{bmatrix} \delta t,$$
(A.19)

where $\mathbb{1}$ is the identity matrix, and the operators \mathcal{L} , ∇_{\parallel} , ∇_{\parallel} are mere finite dimensional sparse matrices as finite differencing has been used for the grid discretization. The boundary conditions are embedded in the operators \mathcal{L} , ∇_{\parallel} , ∇_{\parallel} , which act on the fields in the interior of the domain. After solving the above matrix equation, the updated fields at the next time step are computed as

$$\begin{bmatrix} \mathbf{u}^{n+1} \\ p^{n+1} \end{bmatrix} = \begin{bmatrix} \mathbf{u}^n \\ p^n \end{bmatrix} + \begin{bmatrix} \delta \mathbf{u}^n \\ \delta p^n \end{bmatrix}.$$
 (A.20)

The coefficient matrix in equation (A.19) is a sparse and constant matrix. So, it is computed just once before the time integration starts and its LU factorization performed at the beginning. Using this LU factorization, equation (A.19) is solved at every time step during the integration.

A.4 Jacobian computation

The Jacobian matrix of the state and the operators \mathcal{L} , ∇_{\parallel} , ∇_{\parallel} in equation (A.19) can be explicitly computed for the grid discretization described here, and they will be a sparse matrices. For illustration purposes, we explain how the Jacobian of Laplacian of u can be computed. Let U represent the interior 2D field u when it is written as a single column vector using column major ordering. Then, an element $u(i, j + \frac{1}{2})$ gets mapped to U(k), where $k = n_y(i-1) + j$. If $(i, j + \frac{1}{2})$ represents an interior location that is not adjacent to any of the boundary walls, the Laplacian of u at this location is given by

$$\begin{split} \nabla_{\parallel}^{2} u \big|_{(i,j+\frac{1}{2})} &= \frac{u(i-1,j+\frac{1}{2}) - 2u(i,j+\frac{1}{2}) + u(i+1,j-\frac{1}{2})}{\delta x^{2}} \\ &+ \frac{u(i,j-\frac{1}{2}) - 2u(i,j+\frac{1}{2}) + u(i,j+\frac{3}{2})}{\delta y^{2}} \\ &= \frac{U(k-n_{y}) - 2U(k) + U(k+n_{y})}{\delta x^{2}} + \frac{U(k-1) - 2U(k) + U(k+1)}{\delta y^{2}} \\ &= \nabla_{\parallel}^{2} U(k). \end{split}$$
(A.21)

So, $\nabla^2_{\parallel}U(k)$ depends only on 5 elements of U and hence the k^{th} row of its Jacobian, (say) J_0 , has only 5 non-zero elements thus making it a very sparse matrix. It can be written in compact form as

$$J_{0}(k,l) = \begin{cases} \frac{1}{\delta x^{2}}, & l = k \pm n_{y}, \\ \frac{1}{\delta y^{2}}, & l = k \pm 1, \\ -\frac{2}{\delta x^{2}} - \frac{2}{\delta y^{2}}, & l = k, \\ 0, & \text{otherwise.} \end{cases}$$
(A.22)

For the interior points adjacent to the boundary walls, boundary conditions have to be included while computing the Jacobian. For example, the Laplacian of u at the interior point $(1, \frac{1}{2})$ is

$$\nabla_{\parallel}^{2} u \big|_{(1,\frac{1}{2})} = \frac{u(0,\frac{1}{2}) - 2u(1,\frac{1}{2}) + u(2,\frac{1}{2})}{\delta x^{2}} + \frac{u(1,-\frac{1}{2}) - 2u(1,\frac{1}{2}) + u(1,\frac{3}{2})}{\delta y^{2}}.$$
(A.23)

Using the boundary conditions given by equations (A.8), (A.9) and (A.9), the above equation becomes

$$\begin{aligned} \nabla_{\parallel}^{2} u \Big|_{(1,\frac{1}{2})} &= \frac{-2u(1,\frac{1}{2}) + u(2,\frac{1}{2})}{\delta x^{2}} + \frac{-3u(1,\frac{1}{2}) + u(1,\frac{3}{2})}{\delta y^{2}} \\ &= \frac{-2U(1) + U(n_{y}+1)}{\delta x^{2}} + \frac{-3U(1) + U(2)}{\delta y^{2}} \\ &= \nabla_{\parallel}^{2} U(1), \end{aligned}$$
(A.24)

which implies that the corresponding row of the Jacobian will be

$$J_{0}(1,l) = \begin{cases} \frac{1}{\delta x^{2}}, & l = 1 + n_{y}, \\ \frac{1}{\delta y^{2}}, & l = 2, \\ -\frac{2}{\delta x^{2}} - \frac{3}{\delta y^{2}}, & l = 1, \\ 0, & \text{otherwise.} \end{cases}$$
(A.25)

Likewise, the row of the Jacobian corresponding to all the points adjacent to the boundaries can be computed.

Similarly, differentiating the finite difference formulas of the evolution equations with the state vector and using appropriate boundary conditions, one gets explicit formulas to compute the Jacobian of the evolution equations as a sparse matrix. Computation of this Jacobian and the operators defined in equation (A.19) is performed in PETSc [125, 126], a toolkit very efficient in handling sparse matrices, and is transformed to MATLAB, in which majority of the computations are performed.
APPENDIX B PARAMETRIC CONTINUATION

Suppose that solutions X_0, X_1, \dots, X_k on a branch have been computed respectively, at the monotonic sequence of parameter values $\lambda_0, \lambda_1, \dots, \lambda_k$, then a polynomial extrapolation using the computed solutions can be used to form a good initial guess to compute the solution at a parameter value λ_{k+1} . However, if the degree of the polynomial is large, then extrapolation can lead to errors. Hence, we have restricted attention to quadratic extrapolation as explained below.

Using a quadratic approximation for the solution $\mathbf{X}(\lambda)$ over the interval $[\lambda_{k-2}, \lambda_{k+1}]$, it can be written as

$$\mathbf{X}(\lambda) = \mathbf{a}_0 + \mathbf{a}_1 \lambda + \mathbf{a}_2 \lambda^2, \tag{B.1}$$

where \mathbf{a}_0 , \mathbf{a}_1 , \mathbf{a}_2 are the coefficients of the quadratic polynomial that need to be determined. These coefficients can be determined from the previously computed solutions: the quadratic fit (B.1) requires

$$\begin{aligned} \mathbf{X}_{k-2} &= \mathbf{a}_0 + \mathbf{a}_1 \lambda_{k-2} + \mathbf{a}_2 \lambda_{k-2}^2, \\ \mathbf{X}_{k-1} &= \mathbf{a}_0 + \mathbf{a}_1 \lambda_{k-1} + \mathbf{a}_2 \lambda_{k-1}^2, \\ \mathbf{X}_k &= \mathbf{a}_0 + \mathbf{a}_1 \lambda_k + \mathbf{a}_2 \lambda_k^2. \end{aligned}$$
(B.2)

The above set of equations can be written in the form of a matrix equation as

$$\begin{bmatrix} \mathbf{X}_{k-2} & \mathbf{X}_{k-1} & \mathbf{X}_{k} \end{bmatrix} = \begin{bmatrix} \mathbf{a}_{0} & \mathbf{a}_{1} & \mathbf{a}_{2} \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ \lambda_{k-2} & \lambda_{k-1} & \lambda_{k} \\ \lambda_{k-2}^{2} & \lambda_{k-1}^{2} & \lambda_{k}^{2} \end{bmatrix}.$$
 (B.3)

In the above equation, the matrix that is dependent on the parameter values is a 3×3 matrix because quadratic approximation has been used for the solution. Had an n^{th} degree polynomial been used, it will be a square matrix of size n + 1, which is small. Inverting this matrix gives

$$\begin{bmatrix} \mathbf{a}_0 & \mathbf{a}_1 & \mathbf{a}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{X}_{k-2} & \mathbf{X}_{k-1} & \mathbf{X}_k \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ \lambda_{k-2} & \lambda_{k-1} & \lambda_k \\ \lambda_{k-2}^2 & \lambda_{k-1}^2 & \lambda_k^2 \end{bmatrix}^{-1} . \quad (B.4)$$

Substituting these coefficients into (B.1), initial guess at parameter value λ_{k+1} can be computed. Starting the solver from this initial guess, solution \mathbf{X}_{k+1} can be found. If \mathbf{X} represents a periodic orbit, then there is a time translation symmetry associated with it. So, after the solution \mathbf{X}_k is computed, it has to be shifted along the temporal direction, so that it is in phase with \mathbf{X}_{k-1} , otherwise the polynomial extrapolation would result in very bad guesses.

The procedure described above smoothly continues the solution along the branch as long as the branch does not undergo bifurcation in which the stability exponent of the solution crosses zero. When such bifurcation happens, the Jacobian matrix becomes singular causing problems in the solver, if the latter is based on Newton's method. To overcome this problem near such bifurcation, we choose the increment in parameter value to be such that the new solution (to be computed) is sufficiently far away from the bifurcation so that the singularity of the Jacobian does not cause numerical errors, and not too far away that the polynomial extrapolation results in a bad guess. This increment in parameter can be estimated from the rate of change of floquet exponent with the parameter, $\partial_{\lambda} \Lambda|_{\lambda \approx \lambda_b}$, which can be approximately computed using the stability exponents of the earlier computed solutions near the bifurcation.

Once the increment in parameter is obtained, initial guess is constructed from the polynomial extrapolation, as explained previously, and the solution is computed using this initial guess. After the post bifurcation solution (say X_b) has been computed, its near marginal

direction defined by the eigenvector associated with the stability exponent that is close to zero, assuming there is only one, is found. Using this near-marginal direction, we search for other branches that (possibly) are born out of this bifurcation by initiating the solver using a guess which is a state on this near marginal direction and at a small distance from the computed solution. If the solver converges back to the original solution, X_b , then the distance of the initial guess from X_b is increased and the solver is restarted. This search is performed for four different cases – along both positive and negative directions of the marginal direction and on either side of the bifurcation, as it is *a priori* unknown whether the bifurcation is subcritical or supercritical.

APPENDIX C

SPECTRAL FORMULATION FOR PERIODIC ORBITS

A periodic orbit of time period T can be expanded using the Fourier basis,

$$\mathbf{X}(t) = \sum_{k=-\infty}^{\infty} \mathbf{c}_k e^{-ik\omega t},$$
(C.1)

where $\omega = 2\pi/T$ is the angular frequency of the orbit, and c_k are the Fourier coefficients. Truncating the above expansion to $m_0 + 1$ (for an even m_0) modes, the above equation becomes

$$\mathbf{X}(t) = \sum_{k=-m_0/2}^{m_0/2} \mathbf{c}_k e^{-ik\omega t}.$$
 (C.2)

Differentiating (C.2) with respect to time gives

$$\dot{\mathbf{X}}(t) = \omega \sum_{k=-m_0/2}^{m_0/2} (-ik) \mathbf{c}_k e^{-ik\omega t}.$$
(C.3)

Evaluating (C.2) and (C.3) at the collocation points $t_j = jT/(m_0+1)$ for $j = 1, \dots, m_0+1$ yields

$$\mathbf{X}_{j} = \mathbf{X}(t_{j}) = \sum_{k=-m_{0}/2}^{m_{0}/2} \mathbf{c}_{k} e^{-i\frac{2\pi k_{j}}{m_{0}+1}}$$
(C.4)

and

$$\dot{\mathbf{X}}_{j} = \omega \sum_{k=-m_{0}/2}^{m_{0}/2} (-ik) \mathbf{c}_{k} e^{-i\frac{2\pi kj}{m_{0}+1}}.$$
(C.5)

The Fourier coefficients c_k can be computed using the Discrete Fourier Transform

(DFT) as

$$\mathbf{c}_{k} = \frac{1}{m_{0}+1} \sum_{m=1}^{m_{0}+1} \mathbf{X}_{m} e^{i\frac{2\pi km}{m_{0}+1}}.$$
 (C.6)

Substituting (C.6) into (C.5) gives

$$\dot{\mathbf{X}}_{j} = \omega \sum_{m=1}^{m_{0}+1} A_{jm} \mathbf{X}_{m}, \tag{C.7}$$

where

$$A_{jm} = -\frac{i}{m_0 + 1} \sum_{k=-m_0/2}^{m_0/2} k e^{i\frac{2\pi k(m-j)}{m_0 + 1}}$$
(C.8)

Finally, with the help of (5.40) and (5.41) the condition that the collocation points lie on a periodic orbit can be written as

$$\tilde{\mathbf{F}}(\tilde{\mathbf{X}};\omega) = \tilde{\mathbf{V}} - \omega \tilde{A}\tilde{\mathbf{X}} = 0.$$
(C.9)

REFERENCES

- [1] O. Reynolds, "An experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channels," *Proc. Roy. Soc. London*, vol. 35, pp. 84–99, 1883.
- [2] O. Reynolds, "On the dynamical theory of incompressible viscous fluids and the determination of the criterion," *Poc. Roy. Soc. London*, vol. 56, pp. 40–45, 1894.
- [3] G. I. Taylor, "Statistical theory of turbulence," *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 151, no. 873, pp. 421–444, 1935.
- [4] T. von Kármán, "On the statistical theory of turbulence.," *Proceedings of the National Academy of Sciences of the United States of America*. 1937, vol. 23, no. 2, pp. 98–105,
- [5] T. von Kármán and L. Howarth, "On the statistical theory of isotropic turbulence," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 164, no. 917, pp. 192–215, 1938.
- [6] W. Heisenberg, "On the theory of statistical and isotropic turbulence," *Proc. Roy. Soc. A*, vol. 195, pp. 402–406, 1948.
- [7] L. D. Landau, "On the problem of turbulence," *Akad. Nauk. Doklady*, vol. 44, p. 339, 1944.
- [8] E. Hopf, "A mathematical example displaying features of turbulence," *Commun. Pur. Appl. Math.*, vol. 1, no. 4, pp. 303–322, 1948.
- [9] D. Ruelle and F. Takens, "On the nature of turbulence," *Commun. Math. Phys*, vol. 20, no. 3, pp. 167–192, 1971.
- [10] J. P. Gollub and H. L. Swinney, "Onset of turbulence in a rotating fluid," *Phys. Rev. Lett.*, vol. 35, no. 14, p. 927, 1975.
- [11] G. Ahlers, "Low-temperature studies of the rayieigh-bénard instability and turbulence," *Phys. Rev. Lett.*, vol. 33, pp. 1185–1188, 20 1974.
- [12] B. Saltzman, "Finite amplitude free convection as an initial value problem I," *Journal of the Atmospheric Sciences*, vol. 19, no. 4, pp. 329–341, 1962.

- [13] E. N. Lorenz, "Deterministic nonperiodic flow," J. Atmos. Sci., vol. 20, no. 2, pp. 130–141, 1963.
- [14] G. Kawahara, M. Uhlmann, and L. van Veen, "The significance of simple invariant solutions in turbulent flows," *Annu. Rev. Fluid Mech.*, vol. 44, pp. 203–225, 2012.
- [15] P. Bakewell and J. L. Lumley, "Viscous sublayer and adjacent wall region in turbulent pipe flow," *Phys. Fluids*, vol. 10, p. 1880, 1967.
- [16] H. Kim, S. Kline, and W. Reynolds, "The production of turbulence near a smooth wall in a turbulent boundary layer," *J. Fluid Mech.*, vol. 50, pp. 133–160, 1971.
- [17] D. Dritschel and B. Legras, "Modeling oceanic and atmospheric vortices," *Phys. Today*, vol. 46, p. 44, 1993.
- [18] J. L. Lumley and P. Blossey, "Control of turbulence," Ann. Rev. Fluid Mech., vol. 30, p. 311, 1998.
- [19] J. D. Swearingen and R. F. Blackwelder, "The growth and breakdown of streamwise vortices in the presence of a wall," J. Fluid Mech., vol. 182, p. 255, 1987.
- [20] J. M. Hamilton, J. Kim, and F. Waleffe, "Regeneration mechanisms of near-wall turbulence structures," J. Fluid Mech., vol. 287, pp. 317–348, 1995.
- [21] J. Jiménez and R. D. Moser, "What are we learning from simulating wall turbulence?" *Phil. Trans. Royal Soc. A*, vol. 367, p. 715, 2007.
- [22] J. C. Klewicki, M. M. Metzger, E. Kelner, and E. M. Thurlow, "Viscous sublayer flow visualizations at $R_{\theta} \sim 1500000$," *Phys. Fluids*, vol. 7, p. 857, 1995.
- [23] R. F. Blackwelder and R. Kaplan, "On the wall structure of the turbulent boundary layer," *J. Fluid Mech.*, vol. 76, p. 89, 1976.
- [24] B. Hof, C. W. H. van Doorne, J. Westerweel, F. T. M. Nieuwstadt, H. Faisst, B. Eckhardt, H. Wedin, R. R. Kerswell, and F. Waleffe, "Experimental observation of nonlinear traveling waves in turbulent pipe flow," *Science*, vol. 305, no. 5690, pp. 1594–1598, 2004.
- [25] M. Nagata, "Three-dimensional finite-amplitude solutions in plane Couette flow: Bifurcation from infinity," J. Fluid Mech., vol. 217, pp. 519–527, 1990.
- [26] F. Waleffe, "On a Self-Sustaining Process in shear flows," *Phys. Fluids*, vol. 9, pp. 883–900, 1997.

- [27] G. Kawahara and S. Kida, "Periodic motion embedded in plane Couette turbulence: Regeneration cycle and burst," *J. Fluid Mech.*, vol. 449, pp. 291–300, 2001.
- [28] F. Waleffe, "Exact coherent structures in channel flow," J. Fluid Mech., vol. 435, pp. 93–102, 2001.
- [29] F. Waleffe, "Homotopy of exact coherent structures in plane shear flows," *Phys. Fluids*, vol. 15, pp. 1517–1543, 2003.
- [30] H. Wedin and R. R. Kerswell, "Exact coherent structures in pipe flow: Traveling wave solutions," *J. Fluid Mech.*, vol. 508, pp. 333–371, 2004.
- [31] J. Wang, J. F. Gibson, and F. Waleffe, "Lower branch coherent states in shear flows: Transition and control," *Phys. Rev. Lett.*, vol. 98, no. 20, p. 204 501, 2007.
- [32] D. Viswanath, "Recurrent motions within plane Couette turbulence," *J. Fluid Mech.*, vol. 580, pp. 339–358, 2007.
- [33] J. F. Gibson, J. Halcrow, and P. Cvitanović, "Visualizing the geometry of state space in plane Couette flow," *J. Fluid Mech.*, vol. 611, pp. 107–130, 2008.
- [34] P. Cvitanović and J. F. Gibson, "Geometry of turbulence in wall-bounded shear flows: Periodic orbits," *Phys. Scr. T*, vol. 142, p. 014 007, 2010.
- [35] T. M. Schneider, J. F. Gibson, and J. Burke, "Snakes and ladders: Localized solutions of plane Couette flow," *Phys. Rev. Lett.*, vol. 104, p. 104 501, 2010.
- [36] H. Poincaré, *Les Méthodes Nouvelles de la Méchanique Céleste*. Gauthier-Villars, 1899.
- [37] M. C. Gutzwiller, "Periodic orbits and classical quantization conditions," *J. Math. Phys.*, vol. 12, no. 3, pp. 343–358, 1971.
- [38] F. Christiansen, P. Cvitanović, and V. Putkaradze, "Hopf's last hope: Spatiotemporal chaos in terms of unstable recurrent patterns," *Nonlinearity*, vol. 10, pp. 55–70, 1997.
- [39] Y. Lan and P. Cvitanović, "Unstable recurrent patterns in Kuramoto-Sivashinsky dynamics," *Phys. Rev. E*, vol. 78, p. 026 208, 2008.
- [40] V. López, P. Boyland, M. T. Heath, and R. D. Moser, "Relative periodic solutions of the complex Ginzburg-Landau equation," *SIAM J. Appl. Dyn. Syst.*, vol. 4, pp. 1042–1075, 2006.

- [41] F. Waleffe, "Hydrodynamic stability and turbulence: Beyond transients to a Self-Sustaining Process," *Stud. Applied Math.*, vol. 95, pp. 319–343, 1995.
- [42] S. Zammert and B. Eckhardt, "Crisis bifurcations in plane Poiseuille flow," *Physical Review E*, vol. 91, no. 4, p. 041 003, 2015.
- [43] P. Ritter, F. Mellibovsky, and M. Avila, "Dynamical-systems approach to localised turbulence in pipe flow," arXiv:1512.02382, 2015.
- [44] P. Cvitanović and J. F. Gibson, "Geometry of state space in plane Couette flow," in Advances in Turbulence XII, A. V. Johansson and B. Eckhardt, Eds., ser. Proc. 12th EUROMECH Eur. Turb. Conf. Berlin: Springer, 2009, pp. 41–43.
- [45] E. Hopf, "A mathematical example displaying features of turbulence," *Comm. Pure Appl. Math.*, vol. 1, pp. 303–322, 1948.
- [46] N. Aubry, P. Holmes, J. L. Lumley, and E. Stone, "The dynamics of coherent structures in the wall region of turbulent boundary layer," *J. Fluid Mech.*, vol. 192, pp. 115–173, 1988.
- [47] P. Holmes, J. L. Lumley, and G. Berkooz, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge: Cambridge Univ. Press, 1996.
- [48] K. Lee, L. Cortelezzi, J. Kim, and J. L. Speyer, "Application of reduced order controller to turbulent flows for drag reduction," *Phys. Fluids*, vol. 13, p. 1321, 2001.
- [49] C. W. Rowley, "Model reduction for fluids, using balanced proper orthogonal decomposition," *Int. J. Bif. Chaos*, vol. 15, p. 997, 2005.
- [50] P. J. Shmid, "Dynamic mode decomposition of numerical and experimental data," *Journal of Fluid Mechanics*, vol. 656, 528, 2010.
- [51] M. S. Hemati, M. O. Williams, and C. W. Rowley, "Dynamic mode decomposition for large and streaming datasets," *Physics of Fluids*, vol. 26, no. 11, p. 111701, 2014.
- [52] I. Mezić, "Analysis of fluid flows via spectral properties of the koopman operator," *Annual Review of Fluid Mechanics*, vol. 45, no. 1, pp. 357–378, 2013.
- [53] J. F. Gibson, J. Halcrow, and P. Cvitanović, "Equilibrium and traveling-wave solutions of plane Couette flow," *J. Fluid Mech.*, vol. 638, pp. 1–24, 2009.
- [54] J. Halcrow, J. F. Gibson, P. Cvitanović, and D. Viswanath, "Heteroclinic connections in plane Couette flow," J. Fluid Mech., vol. 621, pp. 365–376, 2009.

- [55] P. Cvitanović, "Periodic orbits as the skeleton of classical and quantum chaos," *Physica D*, vol. 51, p. 138, 1991.
- [56] P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay, *Chaos: Classical and Quantum*. Copenhagen: Niels Bohr Inst., 2011, ChaosBook.org.
- [57] F. Waleffe, "On a self-sustaining process in shear flows," *Physics of Fluids*, vol. 9, no. 4, pp. 883–900, 1997.
- [58] J. D. Skufca, J. A. Yorke, and B. Eckhardt, "Edge of chaos in a parallel shear flow," *Phys. Rev. Lett.*, vol. 96, no. 17, p. 174 101, 2006.
- [59] T. M. Schneider, B. Eckhardt, and J. Yorke, "Turbulence, transition, and the edge of chaos in pipe flow," *Phys. Rev. Lett.*, vol. 99, p. 034 502, 2007.
- [60] A. P. Willis, K. Y. Short, N. B. Budanur, M. Farazmand, and P. Cvitanović, "Relative periodic orbits form the backbone of turbulent pipe flow," *arXiv preprint arXiv:1705.03720*, 2017.
- [61] G. J. Chandler and R. R. Kerswell, "Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow," J. Fluid Mech., vol. 722, pp. 554– 595, Apr. 2013.
- [62] B. Suri, J. Tithof, R. O. Grigoriev, and M. F. Schatz, "Forecasting fluid flows using the geometry of turbulence," *Phys. Rev. Lett.*, vol. 118, p. 114 501, 11 2017.
- [63] D. Viswanath, "The Lindstedt-Poincaré technique as an algorithm for finding periodic orbits," *SIAM Review*, vol. 43, pp. 478–496, 2002.
- [64] Y. Lan and P. Cvitanović, "Variational method for finding periodic orbits in a general flow," *Phys. Rev. E*, vol. 69, p. 016 217, 2004.
- [65] Y. Lan, C. Chandre, "Newton's descent method for the determination of invariant tori," *Phys. Rev. E*, vol. 74, p. 046 206, 4 2006.
- [66] M. Farazmand, "An adjoint-based approach for finding invariant solutions of navierstokes equations," *Journal of Fluid Mechanics*, vol. 795, 278312, 2016.
- [67] W. J. Beyn, "The numerical computation of connecting orbits in dynamical systems," *IMA J. Numer. Anal.*, vol. 10, no. 3, pp. 379–405, 1990.
- [68] W. J. Beyn, "On well-posed problems for connecting orbits in dynamical systems," *Contemporary Mathematics*, vol. 172, no. 3, pp. 131–168, 1994.

- [69] L. Dieci and J. Rebaza, "Point-to-periodic and periodic-to-periodic connections," *BIT Numerical Mathematics*, vol. 44, p. 41, 2004.
- [70] C. Dong and Y. Lan, "A variational approach to connecting orbits in nonlinear dynamical systems," *Physics Letters A*, vol. 378, no. 9, pp. 705–712, 2014.
- [71] B. Suri, J. Tithof, R. Mitchell, R. O. Grigoriev, and M. F. Schatz, "Velocity profile in a two-layer Kolmogorov-like flow," *Phys. Fluids*, vol. 26, no. 5, 053601, p. 053 601, 2014.
- [72] J. Tithof, B. Suri, R. K. Pallantla, R. O. Grigoriev, and M. F. Schatz, "Bifurcations in a quasi-two-dimensional kolmogorov-like flow," *Journal of Fluid Mechanics*, vol. 828, pp. 837–866, 2017.
- [73] H. Kellay and W. Goldburg, "Two-dimensional turbulence: A review of some recent experiments," *Rep. Prog. Phys.*, vol. 65, p. 845, 2002.
- [74] P. Tabeling, "Two-dimensional turbulence: A physicist approach," *Phys. Rep.*, vol. 362, p. 1, 2002.
- [75] B. Suri, "Quasi-two-dimensional Kolmogorov flow: Bifurcations and exact coherent structures," PhD thesis, Georgia Institute of Technology, 2017.
- [76] W. I. Klazkin, "Nonlinear stability theory of periodic flows," J. Appl. Math. Mech. (Prikl. Math. Mekh.), vol. 36, p. 263, 1972.
- [77] J. S. A. Green, "Two-dimensional turbulence near the viscous limit," *J. Fluid Mech.*, vol. 62, p. 273, 1974.
- [78] G. I. Sivashinsky, "Weak turbulence in periodic flows," *Physica D*, vol. 17, p. 243, 1985.
- [79] Z. S. She, "Metastability and vortex pairing in the Kolmogorov flow," *Phys. Lett.*, vol. 124, p. 161, 1987.
- [80] Z. S. She, in *Progress in Astronautics and Aeronautics*, H. Branover, M. Mond, and Y. Unger, Eds., vol. 112, Washington, DC: AIAA, 1987, pp. 374–396.
- [81] D. Armbruster, E. J. K. R. Heiland, and B. Nicolaenko, "Phase-space analysis of bursting behavior in Kolmogorov flow," *Physica D*, vol. 58, pp. 392–401, 1992.
- [82] S. C. Kim and H. Okamoto, "Vortices of large scale appearing in the 2d stationary navierstokes equations at large reynolds numbers," *Japan J. Indust. Appl. Math*, vol. 27, 4771, 2010.

- [83] N. F. Bondarenko, M. Z. Gak, and F. V. Dolzhansky, "Laboratory and theoretical models of a plane periodic flow," *Izv. Akad. Nauk (Fiz. Atmos. Okeana)*, vol. 15, p. 1017, 1979.
- [84] A. M. Batchaev and V. A. Dowzhenko, "Experimental modeling of stability loss in periodic zonal flows," *Dokl. Akad. Nauk*, vol. 273, p. 582, 1983.
- [85] V. Dolzhansky, "On the influence of external friction on stability of plane parallel flows of homogeneous incompressible fluid," *Izv. Akad. Nauk (Fiz. Atmos. Okeana)*, vol. 23, p. 348, 1987.
- [86] J. Sommeria, "Experimental study of the two-dimensional inverse energy cascade in a square box," *J. Fluid Mech.*, vol. 170, p. 139, 1986.
- [87] Y. B. Kolesnikov, "Investigation offlat shear flow instability in a magnetic field," *Magn. Gidrodin.*, vol. 1, p. 60, 1985.
- [88] A. M. Batchaev, "Experimental study of supercritical regimes of Kolmogorov flow on a cylindrical surface," *Izv. Akad. Nauk (Fiz. Atmos. Okeana)*, vol. 24, p. 844, 1988.
- [89] A. M. Obuchov, "Laboratory modeling of Kolmogorov flow," Russ. Math. Surv. (Usp. Mat. Nauk), vol. 38, p. 101, 1983.
- [90] Y. Couder, "Two-dimensional grid turbulence in a thin liquid film," J. Phys. Lett., vol. 45, no. 8, pp. 353–360, 1984.
- [91] J. Sommeria and R. Moreau, "Why, how, and when, MHD turbulence becomes two-dimensional," *J. Fluid Mech.*, vol. 118, pp. 507–518, Apr. 1982.
- [92] N. F. Bondarenko, M. Z. Gak, and F. V. Dolzhanskiy, "Laboratory and theoretical models of plane periodic flows," *Izv. Akad. Nauk SSSR, Fiz. Atmos. Okeana*, vol. 15, no. 10, pp. 711–716, 1979.
- [93] R. A. D. Akkermans, A. R. Cieslik, L. P. J. Kamp, R. R. Trieling, H. J. H. Clercx, and G. J. F. van Heijst, "The three-dimensional structure of an electromagnetically generated dipolar vortex in a shallow fluid layer," *Phys. Fluids*, vol. 20, p. 116601, 2008.
- [94] D. H. Kelley and N. T. Ouellette, "Onset of three-dimensionality in electromagnetically driven thin-layer flows," *Phys. Fluids*, vol. 23, no. 4, 045103, p. 045103, 2011.
- [95] M. K. Rivera and R. E. Ecke, "Pair dispersion and doubling time statistics in twodimensional turbulence," *Phys. Rev. Lett.*, vol. 95, p. 194 503, 2005.

- [96] V. A. Dovzhenko, A. M. Obukhov, and V. M. Ponomarev, "Generation of vortices in an axisymmetric shear flow," *Fluid Dyn.*, vol. 16, no. 4, pp. 510–518, 1981.
- [97] S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Data-driven discovery of partial differential equations," *Science Advances*, vol. 3, e1602614, 2017.
- [98] W. C. Reynolds and M. C. Potter, "Finite-amplitude instability of parallel shear flows," *J. Fluid Mech.*, vol. 27, p. 465, 1967.
- [99] S. J. Chapman, "Subcritical transition in channel flows," *J. Fluid Mech.*, vol. 451, pp. 35–97, 2002.
- [100] B. Eckhardt, T. M. Schneider, B. Hof, and J. Westerweel, "Turbulence transition in pipe flow," Ann. Rev. Fluid Mech., vol. 29, pp. 447–468, 2007.
- [101] J.-P. Eckmann, S. O. Kamphorst, and D. Ruelle, "Recurrence plots of dynamical systems," *EPL (Europhysics Letters)*, vol. 4, no. 9, p. 973, 1987.
- [102] Y. Pomeau and P. Manneville, "Intermittent transition to turbulence in dissipative dynamical systems," *Commun. Math. Phys.*, vol. 74, no. 2, pp. 189–197, 1980.
- [103] C. Kelley, Solving Nonlinear Equations with Newton's Method. SIAM, 2003.
- [104] R. B. Lehoucq and D. C. Sorensen, "Deflation techniques for an implicitly restarted arnoldi iteration," *SIAM Journal on Matrix Analysis and Applications*, vol. 17, no. 4, pp. 789–821, 1996.
- [105] D. C. Sorensen, "Implicit application of polynomial filters in a k-step Arnoldi method," *SIAM Journal on Matrix Analysis and Applications*, vol. 13, no. 1, pp. 357– 385, 1992.
- [106] M. R. Hestenes and E. Stiefel, *Methods of conjugate gradients for solving linear systems*, 1. NBS, 1952, vol. 49.
- [107] Y. Saad, *Iterative Methods for Sparse Linear Systems*, Second. Society for Industrial and Applied Mathematics, 2003.
- [108] Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 3, pp. 856–869, 1986.
- [109] J. Halcrow, J. F. Gibson, P. Cvitanović, and D. Viswanath, "Heteroclinic connections in plane Couette flow," J. Fluid Mech., vol. 621, pp. 365–376, 2009.

- [110] P. Deuflhard, *Newton methods for nonlinear problems: affine invariance and adaptive algorithms.* Springer Science & Business Media, 2011, vol. 35.
- [111] P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay, *Chaos: Classical and quantum*, chapter "*Fixed points and how to get them*", Copenhagen, 2016.
- [112] T. Lakoba and J. Yang, "A mode elimination technique to improve convergence of iteration methods for finding solitary waves," *Journal of Computational Physics*, vol. 226, no. 2, pp. 1693 –1709, 2007.
- [113] N. Qian, "On the momentum term in gradient descent learning algorithms," *Neural Networks*, vol. 12, 145151, 1999.
- [114] Y. Nesterov, "A method of solving a convex programming problem with convergence rate $O(1/k^2)$."
- [115] Y. Liu, L. Liu, and T. Tang, "The numerical computation of connecting orbits in dynamical systems: A rational spectral approach," *Journal of Computational Physics*, vol. 111, no. 2, pp. 373 –380, 1994.
- [116] J. P. Boyd, "Spectral methods using rational basis functions on an infinite interval," *Journal of Computational Physics*, vol. 69, no. 1, pp. 112–142, 1987.
- [117] J. P. Boyd, "The optimization of convergence for chebyshev polynomial methods in an unbounded domain," *Journal of Computational Physics*, vol. 45, no. 1, pp. 43 -79, 1982.
- [118] R. De la Llave *et al.*, "A tutorial on kam theory," in *Proceedings of Symposia in Pure Mathematics*, Providence, RI; American Mathematical Society; 1998, vol. 69, 2001, pp. 175–296.
- [119] J. James and M. Murray, "Chebyshev-taylor parameterization of stable/unstable manifolds for periodic orbits: Implementation and applications," *arXiv preprint arXiv:1706.03345*, 2017.
- [120] J. Sánchez and M. Net, "On the multiple shooting continuation of periodic orbits by newton-krylov methods," *International Journal of Bifurcation and Chaos*, vol. 20, no. 01, pp. 43–61, 2010.
- [121] L. Van Veen, G. Kawahara, and M. Atsushi, "On matrix-free computation of 2d unstable manifolds," *SIAM Journal on Scientific Computing*, vol. 33, no. 1, pp. 25– 44, 2011.

- [122] E. Doedel, B. Kooi, G. van Voorn, and Y. Kuznetsov, "Continuation of connecting orbits in 3D-ODEs (I): Point-to-cycle connections," *International Journal of Bifurcation and Chaos*, vol. 18, no. 07, pp. 1889–1903, 2008.
- [123] E. Doedel, B. Kooi, G. van Voorn, and Y. Kuznetsov, "Continuation of connecting orbits in 3D-ODEs (II): Cycle-to-cycle connections," *International Journal of Bifurcation and Chaos*, vol. 19, no. 01, pp. 159–169, 2009.
- [124] F. H. Harlow and J. E. Welch, "Numerical calculation of time dependent viscous incompressible flow of fluid with free surface," *Phys. Fluids*, vol. 8, no. 12, pp. 2182– 2189, 1965.
- [125] S. Balay, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang, "PETSc users manual," Argonne National Laboratory, Tech. Rep. ANL-95/11 - Revision 2.1.5, 2004.
- [126] Portable, Extensible Toolkit for Scientific Computation, http://www.mcs.anl.gov/petsc.