

ADVANCES IN THE DESIGN OF HIGH-PERFORMANCE FLOW CONTROL

KEVIN K. CHEN

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
MECHANICAL AND AEROSPACE ENGINEERING
ADVISERS: CLARENCE W. ROWLEY AND HOWARD A. STONE

JUNE 2014

© Copyright by Kevin K. Chen, 2015.
All Rights Reserved.

Abstract

This thesis tackles challenges in feedback control design for fluid flows, from multiple angles and approaches. It covers three major facets—stability theory, control, and reduced-order modeling—and it investigates three major challenges of these facets: nonlinearity, high dimensionality, and non-normality.

The dissertation begins with a discussion of global stability via linearized Navier–Stokes eigen-decompositions, including numerical algorithms for this analysis. This section then investigates the global stability of a pipe flow through a T-shaped bifurcation at mid-hundred Reynolds numbers, which exhibits vortex breakdown. The recirculation and sensitivity regions closely coincide, which we explain using an inviscid short-wavelength perturbation theory. We also discuss the stability and receptivity properties of this flow.

The second part discusses feedback control design for fluid flows, including optimal actuator and sensor placement. It presents an algorithm that computes the gradient of a control measure with respect to such placements, allowing an efficient gradient-based optimization. The implementation on the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire models of fluid flow reveals that common methods for placement, such as global mode analysis, are suboptimal. We discuss heuristics, including sensitivity, that may predict optimal placements.

The third part covers reduced-order flow modeling. It examines previously unknown properties of dynamic mode decomposition (DMD)—a data-based modeling technique—including the uniqueness of the numerical algorithm and the boundary conditions of DMD-based models. We also propose an “optimized” DMD that produces less spurious decompositions, and gives the user control over the number of output modes. We show examples from the two-dimensional laminar flow over a cylinder. This part also investigates the stability and performance of high dimensional (e.g., fluid) systems in closed-loop with reduced-order controllers, since such control design is typically necessary for computational tractability. Theorems based on the normalized coprime factorization and ν -gap metric provide sufficient conditions for stability and performance. These conditions can also determine model reduction orders for which the stability or performance of reduced-order control is guaranteed. We demonstrate this on the control of the linearized Ginzburg–Landau system.

Acknowledgments

The last five years of my life have been the greatest, the most fulfilling, and even the most fun up to this point. I owe this to certain individuals who very much deserve my thanks.

First, I would like to thank my advisers, Clancy Rowley and Howard Stone. The magnitude of academic and motivational support that Clancy and Howard have given me is immeasurable. They have challenged and pushed my mind to its limits; yet, they have always done so professionally and cordially, wishing nothing less than the best for me. I am proud to say that Clancy and Howard are role models for me to follow when I one day become a professor myself.

I also thank Phil Holmes, who has been on my Ph.D. committee from the start. My discussions with him never fail to be invigorating, and he has often given me clear directions for my research. Next, I thank my thesis readers Marcus Hultmark and Mihailo Jovanović (University of Minnesota), as well as my thesis defense examiners Clancy, Howard, and Phil. Their time and effort in reviewing my dissertation were critical to its success.

A number of esteemed faculty at other universities deserve mention as well. Through campus visits and conferences, my conversations with Peter Schmid (Imperial College), Mihailo Jovanović, Bernd Noack (Université de Poitiers), Jean-Marc Chomaz (École Polytechnique), Paolo Luchini (Università di Salerno), Igor Mezić (University of California, Santa Barbara), Shervin Bagheri (KTH Royal Institute of Technology), François Gallaire (École Polytechnique Fédérale de Lausanne), and many others have shaped my thesis. These individuals have also given me research directions for the future, and I hope to collaborate with them one day.

Continuing on, I am indebted to the more senior members of the Rowley lab, who were my first line of help when I needed it. The hours upon hours of questions covering fluid mechanics, control theory, dynamical systems, and computation have propelled me forward when I otherwise would have felt lost. Stefan Radl and Daniele Vigolo, two former postdocs under Howard Stone, were also instrumental at the start of my research presented in Chapter 7.

The research presented in this dissertation was primarily supported by the Department of Defense National Defense Science and Engineering Graduate Fellowship and the National Science Foundation Graduate Research Fellowship Program. I also thank Princeton University and the Department of Mechanical and Aerospace Engineering for providing significant stipends and awards. I performed a portion of my dissertation research on clusters provided by the TIGRESS High Performance Computing Center at Princeton University. I am extremely grateful for the university's Computational Science and Engineering Support team, who worked expertly and tirelessly to ensure that my computational research flowed smoothly.

I dedicate my thesis to two individuals who have completely changed my life. First and foremost, I dedicate this work to my wife Jennifer, who is the most tender and faithful person I have ever known, and who gently reminds me that I work too much. I will never forget my six-hour death-defying mission after Hurricane Sandy—careening around fallen trees, downed power lines, and flooded streets—to extract her from a wrecked New York City four days before our wedding. Even if I could live my life a hundred times over, I would not find a better wife. Secondly, I dedicate my research to my father Raymond, who proudly recalls the moment at age 11 when I decided to devote my life to the pursuit of aerospace engineering. As an electrical engineer, he has cheered me on by means of casual and technical conversations, embarrassingly bad engineering jokes, lighthearted competitions between our publication and patent records, and of course, great emotional support.

This dissertation carries T-3278 in the records of the Department of Mechanical and Aerospace Engineering.

Preface

This dissertation contains two parts. The first part provides a concise review of stability, control theory, and dynamical systems analyses of fluid flow. It discusses multiple aspects and ideas that are critical to the successful design of feedback flow control systems. These ideas include local and global theories of flow stability, optimal and robust control, model reduction and control design for high-dimensional systems, and dynamical systems approaches to flow modeling. I present this material toward an audience that does not necessarily have expertise in these subjects. Although the material in this dissertation is interdisciplinary by nature, I write specifically with a fluid mechanics audience in mind. Therefore, some of the control theory discussion begins from a more elementary standpoint.

This first part provides fundamental ideas needed for the second part, which contains published and submitted papers. I have reformatted the articles for this dissertation, but their content is essentially unchanged from the published versions. I outline my individual contributions to these papers in Chapter 6. Finally, the bibliography for the entire thesis, including the attached papers, is printed in the final pages of this dissertation.

Contents

Abstract	iii
Acknowledgments	iv
Preface	v
List of Tables	x
List of Figures	xi
 I Flow stability, control, and modeling	 1
1 Introduction	2
1.1 Overview	2
1.2 Brief survey of flow stability, control, and modeling	3
1.3 Challenges of feedback flow control	5
1.4 Contributions of this dissertation	7
1.5 Organization of Part I	8
 2 Global stability theory of fluid flows	 9
2.1 Overview	9
2.2 A prelude: local stability theory	10
2.3 Flow operators	12
2.4 Eigenmode analysis: stability, receptivity, and sensitivity	15
2.5 Computational methods	18
2.6 An example: application to an impacting pipe T-junction flow	23
 3 Feedback control theory of fluid flows	 27
3.1 Overview	27
3.2 \mathcal{H}_2 optimal control	28
3.3 Optimal actuator and sensor placement	31
3.4 Stability margins and the ν -gap metric	36
3.5 \mathcal{H}_∞ robust control	41
 4 Reduced-order modeling of fluid flows	 43
4.1 Overview	43
4.2 Standard reduced-order modeling techniques	44
4.3 Techniques for very high dimensional systems	46
4.4 Dynamic mode decomposition	48
4.5 Stability and performance with reduced-order controllers	53

5	Conclusion	57
5.1	Summary	57
5.2	Future outlook	59
II	Papers	61
6	Overview	62
6.1	Author contributions	63
7	Vortex breakdown in a T-junction: recirculation and sensitivity	64
7.1	Introduction	64
7.2	Global stability theory	66
7.3	Computational methods	67
7.4	Results	68
7.5	Conclusions	71
8	\mathcal{H}_2 optimal actuator and sensor placement in the Ginzburg–Landau system	72
8.1	Introduction	72
8.2	Complex Ginzburg–Landau equation	74
8.2.1	Continuous representation	74
8.2.2	Discrete representation	76
8.3	\mathcal{H}_2 optimal control and placement	77
8.3.1	State space setup	77
8.3.2	Linear quadratic Gaussian framework	78
8.3.3	Perturbation technique for optimal placement	80
8.4	Results	81
8.4.1	Brute force sampling	81
8.4.2	Conjugate gradient minimisation	84
8.4.3	Further remarks	85
8.5	Comparison with previous approaches	87
8.6	Conclusion	89
9	Fluid flow control applications of \mathcal{H}_2 optimal actuator and sensor placement	91
9.1	Introduction	91
9.2	Theory	92
9.2.1	\mathcal{H}_2 synthesis	92
9.2.2	The gradient of the \mathcal{H}_2 norm	93
9.2.3	Remarks	94
9.3	Examples	94
9.3.1	Overview	94
9.3.2	The linearized Ginzburg–Landau system	95
9.3.3	The Orr–Sommerfeld/Squire system	97
9.4	Comparison with previous approaches	99
9.4.1	Eigenmode analysis	99
9.4.2	Gramian analysis	100
9.4.3	Sensitivity analysis	100
9.5	Conclusion	101

10	Heuristics for effective actuator and sensor placement in feedback flow control	102
10.1	Introduction	102
10.2	\mathcal{H}_2 optimal control and placement	103
10.3	Fluid flow models	104
10.3.1	Overview	104
10.3.2	Linearized Ginzburg–Landau equation	105
10.3.3	Orr–Sommerfeld/Squire equations	106
10.4	Observations and heuristics	107
10.4.1	Fundamental requirements	108
10.4.2	Eigenmode-based analyses	110
10.4.3	Transient growth	111
10.4.4	Remarks	114
10.5	Conclusion	114
11	Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses	116
11.1	Introduction	116
11.2	The Koopman operator and dynamic mode decomposition	117
11.3	Uniqueness of dynamic mode decomposition	120
11.4	Boundary condition handling	122
11.4.1	Overview	122
11.4.2	Mean subtraction	123
11.4.3	Equivalence with temporal DFT and harmonic average	124
11.4.4	Implications of mean subtraction	125
11.4.5	Subtracting an equilibrium	125
11.5	Optimized dynamic mode decomposition	126
11.6	Application to a cylinder flow	129
11.6.1	Overview	129
11.6.2	Numerical method	130
11.6.3	Dynamic mode decomposition of the cylinder flow	131
11.6.4	Comparison with optimized DMD	136
11.6.5	Comparison with the temporal DFT	137
11.6.6	Comparison with equilibrium subtraction	139
11.6.7	Projection error	139
11.7	Conclusion	141
12	Normalized coprime robust stability and performance guarantees for reduced-order controllers	143
12.1	Introduction	143
12.2	The normalized coprime stability margin and the ν -gap metric	145
12.3	Guarantees on the stability and performance of $[\mathbf{P}, \mathbf{K}_r]$	147
12.4	Example: linearized Ginzburg–Landau control	150
12.5	Conclusion	154
	Bibliography	156

List of Tables

3.1	Orr–Sommerfeld/Squire optimal two-actuator, two-sensor placement	35
8.1	Ginzburg–Landau parameters	76
8.2	Optimal placement and cost function variation with state cost weighting	86
11.1	Strouhal number and growth rate comparison	137

List of Figures

1.1	Passive, open-loop, and closed-loop control of the cylinder wake	3
1.2	Cylinder flow impulse response	7
2.1	Conceptual picture of stability	9
2.2	Green's function examples of local stability	11
2.3	Weakly parallel stability regions of the linearized Ginzburg–Landau equation	12
2.4	Stability, receptivity, and sensitivity of a cylinder flow at $Re = 50$	16
2.5	One-dimensional Newton's method example	18
2.6	The T-junction geometry and flow	23
2.7	Streamlines and recirculation regions in the T-junction	24
2.8	Sensitivity regions for the T-junction's first and second instabilities	24
2.9	Stability and receptivity regions of the T-junction flow	25
3.1	Full-state feedback and observer	29
3.2	The LQG and the general control setup	30
3.3	Linearized Ginzburg–Landau MIMO optimal placement	34
3.4	Orr–Sommerfeld/Squire global and adjoint modes	35
3.5	Linearized Ginzburg–Landau global modes, adjoint modes, and overlap	36
3.6	\mathcal{H}_∞ norm and the Nyquist plot	38
3.7	Open-loop and closed-loop representations of perturbed systems	38
3.8	The standard feedback interconnection	39
3.9	The normalized right coprime factorization	39
3.10	The normalized right coprime factorization uncertainty	39
3.11	\mathcal{H}_∞ loop-shaping design	42
4.1	Two-dimensional POD example	47
4.2	Ritz value movement from equilibrium to limit cycle	51
4.3	Conceptual picture of DMD vs. optimized DMD	52
4.4	Approaches for designing reduced-order controllers	54
4.5	Robust stability and performance vs. weighted plant reduction order	56
7.1	The T-junction geometry and flow	65
7.2	Sensitivity function L_2 norm vs. Re , with bifurcations	68
7.3	Streamlines and recirculation regions in the T-junction	69
7.4	Sensitivity regions for the first and second instabilities	69
7.5	Sensitivity and separation at alternate junction corner radii of curvature	70
8.1	General control problem formulation	79
8.2	SISO optimal placement and example, disturbances everywhere	82
8.3	SISO optimal placement and example, disturbances upstream	83

8.4	MIMO optimal placement	84
8.5	5-input, 5-output LQG gain	85
8.6	Global and adjoint modes	87
8.7	Global and adjoint mode overlap	89
9.1	General control problem formulation	92
9.2	Linearized Ginzburg–Landau SISO optimal placement	96
9.3	Linearized Ginzburg–Landau MIMO optimal placement	96
9.4	Orr–Sommerfeld/Squire SISO optimal placement	98
9.5	Orr–Sommerfeld/Squire global and adjoint modes	99
9.6	Linearized Ginzburg–Landau global modes, adjoint modes, and sensitivity	100
10.1	The closed-loop system	104
10.2	Linearized Ginzburg–Landau simulations with and without control	106
10.3	Linearized Ginzburg–Landau MIMO optimal placement	106
10.4	Orr–Sommerfeld/Squire simulations with and without control	107
10.5	Orr–Sommerfeld/Squire optimal placement	108
10.6	Linearized Ginzburg–Landau SISO optimal placement	108
10.7	Linearized Ginzburg–Landau global modes, adjoint modes, and sensitivity	109
10.8	Orr–Sommerfeld/Squire SISO optimal placement	109
10.9	Orr–Sommerfeld/Squire global and adjoint modes	109
10.10	Orr–Sommerfeld/Squire optimal growth	112
10.11	Linearized Ginzburg–Landau optimal growth	112
10.12	Orr–Sommerfeld/Squire impulse responses	113
10.13	Orr–Sommerfeld/Squire impulse response output traces	113
11.1	Cylinder flow at $Re = 60$	129
11.2	Lift coefficient and Strouhal number for cylinder flow	130
11.3	Ritz values	132
11.4	DMD modes near equilibrium	133
11.5	DMD modes from transient regime	133
11.6	DMD mean mode near limit cycle	134
11.7	DMD modes near limit cycle	134
11.8	Ritz value movement from equilibrium to limit cycle	135
11.9	DMD modes from entire transient regime	135
11.10	DMD and optimized DMD comparison	136
11.11	Temporal DFT Ritz values	138
11.12	Ritz values from partial data sets	138
11.13	Ritz values from equilibrium-subtracted data	139
11.14	Projection error using different DMD mode sets	140
12.1	Approaches for designing reduced-order controllers	144
12.2	The standard feedback interconnection	146
12.3	Bode plots of the unweighted and weighted Ginzburg–Landau dynamics	152
12.4	Robust stability vs. controller reduction order	153
12.5	Robust performance vs. controller reduction order	153
12.6	Robust stability vs. weighted plant reduction order	153
12.7	Robust performance vs. weighted plant reduction order	154

Part I

Flow stability, control, and modeling

Chapter 1

Introduction

1.1 Overview

Although researchers have independently studied fluid mechanics and control theory for a considerable time, serious efforts to apply the latter in the former have only been prominent more recently. This synergy has yielded new physical insights into the behavior of fluid flows, as well as new control and dynamical systems approaches that are also applicable to other physical systems.

Techniques for flow control typically fall in three categories: passively placing structures in a flow, actively running an actuator in the flow in an open-loop sense, and closed-loop feedback mechanisms where mathematical controllers dictate actuator behavior based on sensor information. (See Figure 1.1.) Historically, many seminal flow control results focused largely on fluid physics, and utilized passive control (e.g., Strykowski and Sreenivasan, 1990) or very basic open-loop or closed-loop active control (e.g., Roussopoulos, 1993). Despite the successes of such studies, modern feedback control theory may facilitate significant improvements in flow control, such as increased robustness and wider operating ranges. For this reason, this dissertation focuses on topics related to feedback flow control.

Although there now exists a plethora of academic studies in feedback flow control, real-world applications currently remain limited. The potential that feedback flow control carries, however, should not be underestimated. With advanced research, we may be able to achieve drag reduction, acoustic noise reduction, and increased aerodynamic performance on vehicles of all types and scales. The reduction of skin-friction drag by using feedback control to smooth out turbulence, for example, could drastically cut fuel consumption. These benefits could translate into reduced pollution, billions of dollars in fuel savings, and the preservation of finite natural resources (Kim and Bewley, 2007). Furthermore, feedback flow control could substantially increase the performance of micro air vehicles (MAVs) and other unmanned aircraft, and even wind turbines for clean energy generation.

Why, then, has the penetration of feedback flow control in the aerospace industry been so limited? The union of fluid mechanics and control has been a decidedly difficult one to research and develop, because of significant challenges that the governing equations of fluid mechanics pose. Among the greatest of these are nonlinearity, high dimensionality, and non-normality. Individually, each of these issues is notably difficult to manage. The combination of these—as well as numerous other issues—enormously magnifies the challenge.

In this dissertation, we discuss these three issues in detail, and we present research that addresses these issues from varied approaches. These approaches include the effect of non-normality on stability analysis—an important precursor for control—as well as on control itself. We also discuss model reduction approaches for the computational tractability of high-dimensional control, as well

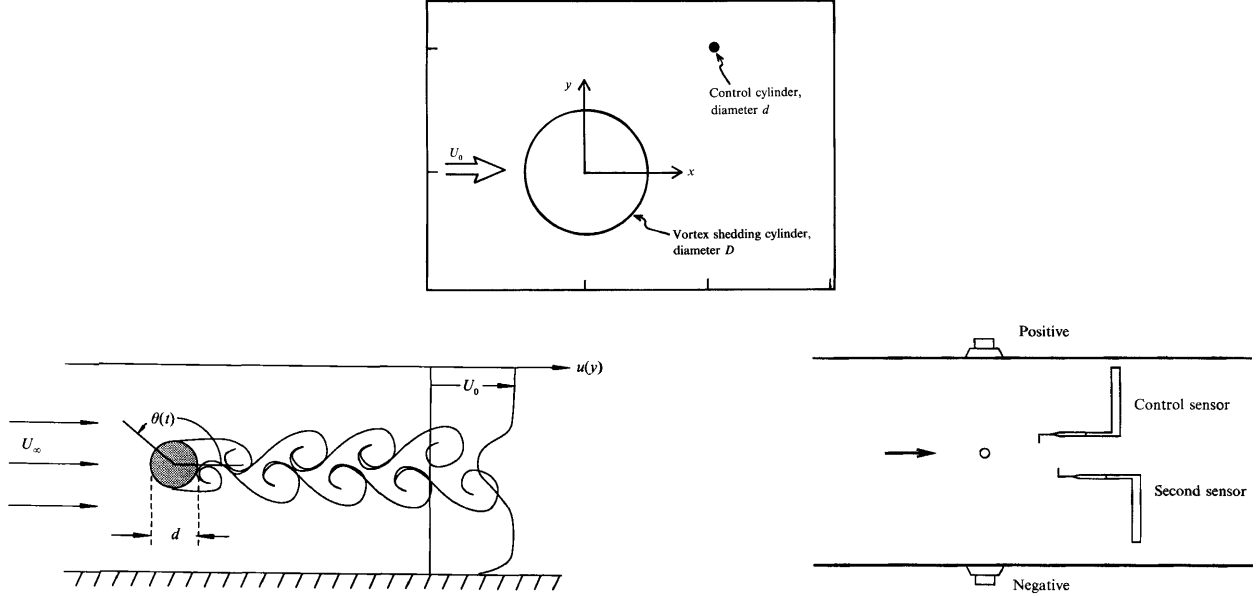


Figure 1.1: Three different control techniques, from least to most complex, for the canonical problem of reducing cylinder wake oscillations. Top: passive control by placing a small control cylinder to the upper right (Strykowski and Sreenivasan, 1990). Bottom left: open-loop active control by prescribing the cylinder angular displacement (Tokumaru and Dimotakis, 1991). Bottom right: closed-loop (i.e., feedback) control by determining the loudspeaker signals based on sensor information (Roussopoulos, 1993).

as control stability and performance with reduced-order control. Finally, we tackle nonlinearity in the discussion of dynamic mode decomposition (DMD) as a modeling technique for fluid flows.

In Section 1.2, we give a brief overview of recent research in flow stability, control, and modeling. Section 1.3 summarizes the difficulties of nonlinearity, high dimensionality, and non-normality; it discusses both previous successes and outstanding issues in these areas. Next, Section 1.4 describes the contributions that the research in this dissertation has made to these broad fields. Finally, Section 1.5 outlines the organization of the remaining chapters in Part I.

1.2 Brief survey of flow stability, control, and modeling

An overview of research and applications of flow stability, control, and modeling would be far too expansive to give here. Instead, throughout the following discussion, we highlight review articles and examples that lay out the general background of these topics.

Much of the early work in flow stability relied on a parallel or local flow approximation, in which we assume that streamwise variations in the flow occur over large length scales, or do not occur at all. Therefore, we assume that the fluid at any streamwise point behaves as if the flow were everywhere governed by the flow conditions at that point. Although this assumption is quite severe and restrictive, it is significantly computationally cheaper, since it generally reduces the number of spatial dimensions required for computation. This analysis has been able to predict the stability characteristics of many flows; see Huerre and Monkewitz (1990) and Chomaz (2005). It has even been successful in more complex flows, such as viscosity-stratified flows (Govindarajan and Sahu, 2014).

More recently, the focus of flow stability research has shifted somewhat to a global analysis, which makes no assumptions about streamwise variations in the flow. This approach, which Chomaz

(2005) and Theofilis (2011) review, is far more computationally expensive, but it gives a truer representation of the dynamics. In general, the only admissible assumption is that of linearization about a fixed point or trajectory. The linear analysis is based largely on direct and adjoint eigen-decompositions of the linearized flow operators, which yield information about long-term growth, receptivity to initial conditions and external forcing (Luchini and Bottaro, 2014), and sensitivity to spatially localized dynamical perturbations. Some recent work, which Schmid (2007) reviews, has considered alternate analyses such as pseudospectra for highly non-normal flows. In these flows, traditional eigenmode analyses are unable to explain behavior such as large transient growths and eigenvalue sensitivity.

Stability analysis is, in some sense, a precursor for flow control, even if research in the former does necessarily not come chronologically before research in the latter. Much of feedback control theory is based on the stability of systems in open-loop, or under closed-loop control (Skogestad and Postlethwaite, 2005). Although early research in feedback flow control has successfully applied classical control theory to fluid flows with limited models of flow dynamics (e.g., Roussopoulos, 1993), more recent work has focused on control design based on relatively faithful linear models of fluid flow. The review by Kim and Bewley (2007) discusses control theory concepts that researchers have effectively employed in linear fluid flow models.

Many reviews have also discussed successful implementations of control in specific flow applications. Lumley and Blossey (1998) and Kim (2003) discuss important factors relevant to the control of turbulence, a problem that remains open to this day. Rowley and Williams (2006) provide a treatise of cavity flow oscillations, as well as passive, open-loop, and closed-loop feedback control techniques to reduce these oscillations. Choi et al. (2008) survey various aspects of bluff body flow control, including vortex, separation, and wake effects. Choi et al. (2014) focus specifically on the aerodynamics of large vehicles such as trucks, buses, and trains, and explore passive techniques for drag reduction. Researchers have made progress on the feedback control of combustion oscillations (Dowling and Morgans, 2005), and research on spike-type stall control is improving the engineering of axial compressors (Tan et al., 2010). Cattafesta and Sheplak (2011) discuss the broad range of flow actuators that are applicable in real-life scenarios.

As previously suggested, the design of modern feedback flow control systems requires faithful models of the flow dynamics. The conception of these models itself is rather nontrivial, and has received notable attention in the flow control community. Traditional modeling techniques drawn from the general control theory community (e.g., balanced truncation and optimal Hankel norm reduction; see Skogestad and Postlethwaite 2005) are not immediately applicable because of the prohibitively large size of fluid dynamics operators in computations.

Perhaps the most venerable of the fluid modeling techniques is proper orthogonal decomposition (POD; Sirovich, 1987; Holmes et al., 1996), also known as principal component analysis and Karhunen–Loève analysis. This method decomposes experimental or computational data snapshots into orthogonal modes, from which we may derive low-dimensional fluid models via the computation of mode coefficients. A more recent method known as balanced POD (Rowley, 2005) approximates balanced truncation (see Skogestad and Postlethwaite, 2005), a modeling technique known to be generally effective for linear systems. This method, along with a related technique known as the eigensystem realization algorithm (ERA; Juang and Pappa, 1985), is based on the computation of input–output impulse responses. Lastly, Schmid (2010) introduced DMD as a numerical algorithm that “curve-fits” snapshots of experimental or computational data to linear trajectories, even if nonlinear dynamics generate the data. This method is a numerical approximation of a Koopman operator method (Rowley et al., 2009; Mezić, 2013) rather than linearization, and may soon yield accurate reduced-order fluid models.

The study of flow stability, control, and modeling is far from solved. In particular, a number of issues have consistently posed formidable challenges in this field. The following section outlines the nature of these difficulties.

1.3 Challenges of feedback flow control

As aforementioned, three of the key challenges in successful feedback flow control are nonlinearity, high dimensionality, and non-normality. Put together, the whole of these difficulties is perhaps greater than the sum of their parts. For feedback flow control systems to be successful in wide operating conditions and applications, they must typically address all three.

The source of nonlinearity is well-known to those familiar with the governing laws of fluid mechanics. The most general expression of the momentum conservation law for a velocity field \mathbf{u} , density field ρ , pressure field p , deviatoric stress tensor $\boldsymbol{\tau}$, and body force field $\boldsymbol{\gamma}$ is

$$\frac{\partial(\rho\mathbf{u})}{\partial t} + \nabla \cdot (\rho\mathbf{u}\mathbf{u}) = -\nabla p + \nabla \cdot \boldsymbol{\tau} + \boldsymbol{\gamma}, \quad (1.1)$$

where the second term on the left-hand side is clearly nonlinear. This poses a great challenge for control, because linear control theory is a far more mature field of study than nonlinear control theory. In fact, it is difficult or impossible to apply control techniques generally and universally across nonlinear systems.

By far the most common method for handling the nonlinearity in feedback flow control is to linearize the fluid dynamics about a base flow, and to use the linearized dynamics in the control design (see Kim and Bewley, 2007). This approach tends to be successful, and there exist too many examples to cite here. It understandably only works up to a limited extent, however, since the linearized equations are only accurate when perturbations away from the base flow are very small. An alternative to simply ignoring the nonlinear terms is to cast them as forcing terms in otherwise linear equations (e.g., Landahl, 1967, 1975). A number of authors (e.g., Trefethen et al., 1993; Farrell and Ioannou, 1993; Bamieh and Dahleh, 2001; Jovanović and Bamieh, 2005) have made key connections between the behaviors of nonlinear shear flows and their forced linearized counterparts. Others (e.g., Sharma et al., 2011) have successfully implemented linear control by retaining nonlinear terms as exogenous forcing. Nevertheless, there does not—and might not ever—exist a universally applicable method for implementing linear control in nonlinear fluid equations. As far as reduced-order modeling is concerned, some techniques such as POD-based Galerkin modeling and DMD are able to represent the fluid equations’ nonlinearities with reasonable success. Nevertheless, the successful control and modeling of the nonlinear fluid equations remains an open research problem.

The high dimensionality of feedback flow control arrives from the numerical spatial discretization of the flow equations into a finite number of grid points. Such discretization is not strictly necessary; as an example, there exist methods for the boundary control of simple partial differential equations (see Aamo et al., 2005, 2007). For sophisticated equations such as Navier–Stokes equations, however, domain discretization is the most universally applicable way of applying model-based linear control theory. In this approach, we recast the partial differential equations governing the flow (e.g., (1.1)) into ordinary differential equations of a large number of states. The states contain the relevant flow variables on the discrete grid points, and we cast the partial derivatives as operations on a vector of these states. From this approach, it should be clear how the high-dimensionality arises. Over the last few decades, numerical meshes for fluid simulations have increased from 14×14 two-dimensional grids (totaling 196 cells; Winters, 1987) to the extreme of $8,192 \times 8,192 \times 8,192$

three-dimensional grids (totaling $5.5 \cdot 10^{11}$ cells; Yeung and Sreenivasan, 2013). The dimensionality of discretized fluid flow equations, then, is typically this number of cells or grid points multiplied by the number of scalar flow variable fields needed to represent the dynamics.

Such a dimensionality is problematic for model-based control design. Many optimal and robust control techniques, for instance, rely on the solution of quadratic matrix equations known as continuous algebraic Riccati equations (CAREs; Doyle et al., 1989). Simple scaling experiments show that the execution time of a CARE solution in the MATLAB software scales by approximately $n^{2.5}$, where n is the state size. A notebook computer with an Intel Core i7-4960HQ processor could solve a 400-state double-precision CARE in 0.6 seconds. Yet, a simple extrapolation shows that the solution of a 10^5 -state CARE would require 1 week, and the solution of a 10^{10} -state CARE would require 4.4 times the current age of the universe! Even the storage of such matrix systems is prohibitive. The 10^5 -state system would involve matrices 75 GB in size—large, but manageable with some difficulty, by today’s standards. On the other hand, the 10^{10} -state system would involve matrices $8 \cdot 10^{20}$ bytes in size. No computer in the present date can store anywhere near this much data. Therefore, it is typically necessary to derive reduced-order approximations of the fluid dynamics to maintain the computational tractability of control design.

The essence of ongoing research in reduced-order modeling and control is the design of algorithms that are computationally cheap, yet yield effective control. Although countless papers have successfully used POD, this decomposition is not without fault. A key outstanding issue is that the POD modes that are dominant in the data are not necessarily the ones that are dynamically important for reduced-order modeling. Therefore, the selection of POD modes to retain in modeling is often arbitrary, and it is generally unclear how to arrive at a “best” set of modes. Although balanced POD remedies this issue by sorting its modes in order of controllability and observability, it requires the data from direct and adjoint impulse responses of the fluid system (Rowley, 2005). This computation is expensive, and experiments cannot provide adjoint data. The ERA approach is able to replicate balanced POD without the need for adjoint impulse responses (Ma et al., 2011), but it may be difficult to obtain direct impulse response data from experiments. Furthermore, balanced POD and ERA are restricted to linear fluid models. Finally, we note that DMD shows promise in potentially yielding physically meaningful nonlinear reduced-order models (e.g., Schmid, 2010; Rowley et al., 2009; Schmid, 2011; Schmid et al., 2011; Chen et al., 2012; Jovanović et al., 2014), but to the best of our knowledge, no one has accomplished such a task to date. In addition, like in POD, the selection of dynamically meaningful DMD modes may not always be straightforward. Furthermore, DMD exhibits a particular sensitivity to output noise because of its mathematical construction. To this date, model reduction techniques that are consistently effective, generally applicable, and easy to compute do not yet exist.

We also comment that a smaller body of research has been dedicated to the design of controllers directly from high-dimensional fluid models. These methods (see Kailath, 1973; Pralits and Luchini, 2010; Bewley et al., 2011; Semeraro et al., 2013) are generally computationally expensive, but may be able to yield higher-performing controllers than control design based on reduced-order fluid models.

Finally, we conclude this section with a few remarks on non-normality. A linear operator \mathcal{L} is said to be non-normal if, given the adjoint \mathcal{L}^* , $\mathcal{L}\mathcal{L}^* \neq \mathcal{L}^*\mathcal{L}$. This condition has many important implications. Perhaps the most relevant of these to flow control is that the eigenvectors of \mathcal{L} may be nearly parallel, which can give rise to large transient growths—even when \mathcal{L} is stable in the infinite-time sense. Figure 1.2 shows an example of a large transient growth in the stable wake of a cylinder. Most shear flows at respectably large Reynolds numbers are highly non-normal, where transient growths can trigger significant nonlinearities, leading ultimately to the failure of linear control. Furthermore, the eigenvalues and eigenmodes of \mathcal{L} do not individually predict




Figure 1.2: The response to an impulse in the cylinder’s displacement, at $Re = 35$; from Le Gal and Croquette (2000). Although the flow is stable at this Reynolds number, non-normality causes a significant transient growth.

the behavior of transient growths. (See Schmid and Henningson (2000); Schmid (2007) for more detailed discussions.)

Following the early investigation of non-normality by Butler and Farrell (1992), we now possess tools such as pseudospectra (Trefethen et al., 1993; Trefethen and Embree, 2005), sensitivity analyses (Giannetti and Luchini, 2007), impulse responses, optimal growth (see Schmid, 2007), and other analyses that better describe transient growths in non-normal flows. In addition, there have been observations that non-normality decreases the stabilizability of flows (Lauga and Bewley, 2003), and that certain feedback control techniques may be able to reduce non-normality (Lauga and Bewley, 2004). Nevertheless, the understanding of effective control techniques for non-normal flows remains an open research question.

In this dissertation, we present ideas that address nonlinearity, high dimensionality, and non-normality in specific ways. It is our hope that these studies will open doors for further research in these topics.

1.4 Contributions of this dissertation

This dissertation includes four separate studies that analyze specific topics in flow stability, control, and modeling.

Section 2.6 and Chapter 7 present a linear global stability analysis of the flow through a T-shaped pipe bifurcation. The laminar flow through this pipe element is very sophisticated, and the mathematical stability analysis is intrinsically connected to the physical mechanisms underlying the flow, including the breakdown of vortices in the junction. The chief contribution of this study is the three-dimensional demonstration of recirculation and dynamical sensitivity regions coinciding. Although prior studies have shown this phenomenon, the T-junction flow is—to the best of our knowledge—by far the most complicated flow for which anyone has observed this behavior.

Section 3.3 and Chapters 8–10 investigate the optimal placement of actuators and sensors in fluid flow models. This is one of comparatively few studies in the literature that have looked beyond controller design, and have specifically targeted actuator and sensor placement. This study improves a placement optimization technique that researchers in flexible structure control have

previously developed. It also demonstrates that certain placement techniques in the flow control literature are actually suboptimal and conceptually flawed. Finally, it explores heuristics based on tools from stability and control theory that may be able to predict optimal placements.

Section 4.4 and Chapter 11 explore previously unknown properties of DMD. First, this research reveals necessary and sufficient conditions for the numerical DMD algorithm to produce a unique decomposition. Next, it analyzes methods for computing DMD modes with homogeneous boundary conditions, since such modes are often necessary for reduced-order model design. This analysis reveals that the DMD of mean-subtracted data is analytically identical to the discrete Fourier transform. On the other hand, subtracting an equilibrium point prior to the DMD computation roughly preserves the DMD modes and eigenvalues, while giving the modes the needed homogeneous boundary conditions. Finally, this study presents an “optimized” variant of the DMD algorithm, which yields less spurious results than the original DMD algorithm.

Section 4.5 and Chapter 12 consider the stability and performance of high-dimensional plants (e.g., fluid dynamical systems) under feedback control by a reduced-order controller. This type of design is very common in feedback flow control, because the large dimensionality of fluid representations renders direct modern control design intractable. Using the normalized coprime stability margin and the ν -gap metric, this research reveals theorems that provide sufficient conditions for the stability or performance of such closed-loop systems. These conditions allow researchers to compute, *a priori*, the model reduction orders for which stability or performance is guaranteed.

1.5 Organization of Part I

Part I of this dissertation presents the fundamental theories of global stability, control, and reduced-order modeling, and it also summarizes the key results from the papers in Part II.

Chapter 2 briefly covers fundamental concepts of local stability theory, as a precursor to the global stability theory. Next, it reviews stability, receptivity, and sensitivity as important ideas that direct and adjoint eigendecompositions of the linearized Navier–Stokes operators reveal. This chapter also reviews computational methods for performing global linear stability analyses of high-dimensional systems, and it concludes by applying these methods and ideas to the pipe flow through a T-shaped bifurcation. Chapter 3 outlines general theories of linear optimal and robust control. It then analyzes the optimal placement of actuators and sensors in fluid systems. This chapter also discusses stability margins and the ν -gap metric, which play important roles in robustness analyses and robust control design. Chapter 4 summarizes reduced-order models for general linear systems, as well as for high-dimensional (e.g., fluid) systems. It then discusses recent advances in DMD theory. This chapter concludes by using the stability margin and ν -gap theory of Section 3.4 to examine the closed-loop stability and performance of high-dimensional systems in closed-loop with reduced-order controllers. Finally, Chapter 5 provides brief concluding remarks.

Chapter 2

Global stability theory of fluid flows

2.1 Overview

The rigorous investigation of fluid flow stability dates back to the famed pipe flow experiments of Reynolds (1883). Since then, studies in flow stability have typically started from a simple question: given some prescribed flow, will a small perturbation experience decay, drawing the flow back to its original state? Alternatively, will the perturbation grow, ultimately leading the flow away from this state?

A common conceptual visualization of this idea is that of a ball on a potentially uneven surface, under the effects of gravity and friction. A steady-state solution is *Lyapunov stable* if any infinitesimally small perturbation in the state results in a trajectory that is at most infinitesimally small. As shown in Figure 2.1(a, b), a ball in a bowl or on a flat surface is therefore Lyapunov stable. A Lyapunov stable steady-state solution is further said to be *asymptotically stable* if any infinitesimal perturbation in the state specifically results in a trajectory that approaches the steady-state solution. Therefore, the ball in a bowl (Figure 2.1(a)) is asymptotically stable. Finally, a steady-state solution that is not Lyapunov stable, such as the ball on top of an inverted bowl in Figure 2.1(c), is unstable.

Beyond the issue of perturbation growth and decay, the study of flow stability has spawned many more questions. What factors or flow features can trigger instabilities? What are the effects of initial conditions or external forcing on the flow behavior? How do transient growths appear, and what is their effect on the flow stability? What are the effects of nonlinearity, and how does nonlinearity alter linear stability results? Where is the flow most sensitive to perturbations in the underlying dynamics? The reviews by Huerre and Monkewitz (1990), Chomaz (2005), Schmid (2007), and Theofilis (2011) discuss these and many other related topics.

Historically, a major focus in the study of flow stability has been the use of a parallel or local flow approximation (Huerre and Monkewitz, 1990; Chomaz, 2005), where we assume that streamwise flow variations have large length scales, or that the flow is streamwise constant. This approximation




Figure 2.1: Conceptual picture of stability, using a ball under the effects of gravity and friction. (a) Lyapunov and asymptotically stable. (b) Lyapunov stable. (c) Unstable.

vastly simplifies the computation and analysis of the flow stability, since it is generally possible to represent instabilities using a reduced number of spatial dimensions. This research in this dissertation does not use local stability theory, but Section 2.2 presents an overview for the purpose of historical perspective.

Instead, the focus of this chapter is on the linear *global* theory of flow stability. In this theory, we make no assumptions other than the linearization of the underlying flow equations about a fixed point. The relaxation of the parallel flow assumption allows us to investigate flows where instability mechanisms may have a nontrivial streamwise dependence. The global theory stems primarily from the direct and adjoint eigendecompositions of the linearized flow operators. This yields information about multiple facets of flow, including the stability of the nonlinear and linearized dynamics, the long-term behavior of the linearized dynamics, the receptivity of the dynamics to external forcing and the choice of initial conditions, and the sensitivity of the flow behavior to localized perturbations in the linearized dynamics.

The eigendecomposition of linearized flow operators is typically a very computationally intensive task. Since computers at this time do not have sufficient power to perform these computations directly, it is necessary to use certain iterative techniques that converge toward the most important global stability solutions. A treatise of the various available computational methods would be too consuming to provide here. Instead, this chapter focuses on a particular set of methods: the Newton–GMRES iteration with the Armijo rule for seeking steady-state solutions, followed by the Arnoldi iteration for seeking leading eigenvalues and eigenmodes.

In this chapter, we apply these algorithms to investigate the linear global stability of a pipe flow through a T-shaped bifurcation at $Re \approx 560$. This investigation locates the regions of the geometry relevant to flow’s long-term behavior, as well as the flow’s receptivity to initial conditions and external forcing. This study also probes vortex breakdown regions in the flow, providing a complex three-dimensional example of a recent observation that regions of recirculation and sensitivity often coincide (Giannetti et al., 2010).

In this chapter, Section 2.2 briefly reviews local stability theory, and Section 2.3 reviews the flow operators that linear global stability theory utilizes. Section 2.4 then relaxes the parallel flow assumption and reviews this linear global stability theory, whose computational methods Section 2.5 discusses. Finally, Section 2.6 presents original research that applies the global stability theory to the flow through a pipe T-junction.

2.2 A prelude: local stability theory

As aforementioned, the key assumption in local stability theory is that the base flow varies slowly, or not at all, in the streamwise direction. This greatly simplifies the stability analysis and often still produces correct predictions of physical behavior.

This theory, which Huerre and Monkewitz (1990) and Chomaz (2005) review, typically analyzes the behavior of modal perturbations superimposed on top of a base flow. It decomposes flow perturbations into a linear combination of waves $A\phi(y, k, \omega) \exp(i(kx - \omega t))$, where A is an amplitude, ϕ is an eigenfunction, x is the streamwise direction, y is the transverse direction, $k \in \mathbb{C}$ is a spatial wavenumber, and $\omega \in \mathbb{C}$ is a temporal frequency. (In three spatial dimensions, it is common to assume a wave-like form in the spanwise dimension as well.) The insertion of these waves into the governing flow equations yields a dispersion relation $D(k, \omega; Re) = 0$. Two special cases of the dispersion relation are worthy of mention. First, if we select some $k \in \mathbb{R}$, then the perturbations are streamwise constant, and one may solve the dispersion relation for the unknown ω . By the Hartman–Grobman theorem, if $\text{Im}(\omega) < 0$, then the wavenumber- k mode is asymptotically stable;




Figure 2.2: Green's function examples of a (a) stable, (b) convectively unstable, and (c) absolutely unstable partial differential equation of one spatial dimension. From Bagheri et al. (2009).

if $\text{Im}(\omega) > 0$, then that mode is unstable. In the degenerate case where $\text{Im}(\omega) = 0$, the mode is neutrally stable in the linearized dynamics, but further tools are needed to determine its stability in the nonlinear dynamics. Second, if we select $\omega \in \mathbb{R}$, then the perturbations oscillate in time, and one may solve the dispersion relation for the unknown k . If $\text{Im}(k) < 0$, then the perturbation grows in the streamwise direction, and if $\text{Im}(k) > 0$, then it decays. The degenerate case $\text{Im}(k) = 0$ is as before.

An important feature of the parallel flow assumption is the categorization of flows as *stable*, *convectively unstable*, and *absolutely unstable* (see Huerre and Monkewitz, 1990). This determination stems from the nature of the system's response to an impulse forcing. Suppose a set of linearized dynamics are given for the perturbation variable $u(x, t)$ by the equation $\mathcal{L}u(x, t) = 0$. Now, consider the Green's function $G(x, t)$ —that is, the impulse response—satisfying $\mathcal{L}G(x, t) = \delta(x - \xi)\delta(t)$ for the Dirac delta function δ and impulse location $x = \xi$. The flow is considered

- *stable* if $\lim_{t \rightarrow \infty} G(x, t) = 0$ for $x - \xi = ct$, for all $c \in \mathbb{R}$ (see Figure 2.2(a));
- *convectively unstable* if $\lim_{t \rightarrow \infty} G(\xi, t) = 0$, but $\lim_{t \rightarrow \infty} |G(x, t)| = \infty$ for $x - \xi = ct$, with some $c \in \mathbb{R} \setminus \{0\}$ (see Figure 2.2(b)); and
- *absolutely unstable* if $\lim_{t \rightarrow \infty} |G(\xi, t)| = \infty$ (see Figure 2.2(c)).

In practical terms, a flow is stable if all perturbations decay with time, convectively unstable if there exists an unbounded growth that convects away from the perturbation source, and absolutely unstable if the unbounded growth does not convect away from the source.

By using a “weakly nonparallel” assumption, the local theory of flow stability allows the investigation of slowly varying flows. This formulation assumes that the dynamics at any given point behave as if the flow were everywhere like at that point. A key example is the investigation of the Ginzburg–Landau equation in Chomaz et al. (1988) and Cossu and Chomaz (1997). Given a streamwise spatial variable $x \in \mathbb{R}$, an amplification and decay function $\mu(x) : \mathbb{R} \rightarrow \mathbb{R}$, an advection speed $\nu = U + 2ic_u$ (with $U, c_u \in \mathbb{R}$), and a diffusion $\gamma = 1 + ic_d$ (with $c_d \in \mathbb{R}$), the linearized equation is

$$\frac{dq}{dt} = \mathcal{L}q, \quad (2.1a)$$

where

$$\mathcal{L} = \mu(x) - \nu \frac{\partial}{\partial x} + \gamma \frac{\partial^2}{\partial x^2}, \quad (2.1b)$$

subject to some choice of boundary conditions. Following Bagheri et al. (2009) and Chen and Rowley (2011), we choose $U = 2$, $c_u = 0.2$, and $c_d = -1$, and we require $\lim_{x \rightarrow \pm\infty} q(x, t) = 0$.




Figure 2.3: The stable (S), convectively unstable (CU), and absolutely unstable (AU) regions of the linearized Ginzburg–Landau example. The green line (at $\mu = -c_u^2$) indicates the threshold for convective instability, and the red line (at $\mu = \mu_t - c_u^2$) indicates the threshold for absolute instability.

In the parallel analysis, we consider a constant $\mu \in \mathbb{R}$. Following Cossu and Chomaz (1997) and Bagheri et al. (2009), the entire flow is stable if $\mu + c_u^2 \leq 0$. If $0 < \mu + c_u^2 < \mu_t$, where $\mu_t = (U + 2c_d c_u)/(4|\gamma|^2)$, then the entire flow is convectively unstable. Finally, if $\mu + c_u^2 \geq \mu_t$, then the entire flow is absolutely unstable. In the weakly nonparallel analysis, we now suppose that $\mu(x)$ has a small dependence on x ; in this example, we choose $\mu(x) = \mu_0 - c_u^2 + \mu_2 x^2/2$, with $\mu_0 = 0.41$ and $\mu_2 = -0.01$; thus, $\mu_t = 0.32$. The weakly nonparallel assumption allows us to consider the regions where $\mu(x) + c_u^2 \leq 0$, $0 < \mu(x) + c_u^2 < \mu_t$, and $\mu(x) + c_u^2 \geq \mu_t$, and respectively identify those regions as stable, convectively unstable, and absolutely unstable. (See Figure 2.3.) If the streamwise dependence of the flow is indeed small, then the local behavior of impulse responses in a region will roughly match the stability type identified with that region.

The weakly parallel analysis has an important consequence related to global stability: Chomaz et al. (1991) showed that the existence of an absolute instability region is necessary for global instability. In the nonparallel example above, the linearized Ginzburg–Landau operator \mathcal{L} is unstable when $\mu_0 > \mu_c$, where $\mu_c = \mu_t + |\sqrt{-2\mu_2\gamma}| \cos((\arg \gamma)/2)/2 > \mu_t$. For the values assigned above, $\mu_c = 0.40$, and absolute instability locally occurs at $\mu_0 > \mu_t = 0.32$. The choice of $\mu_0 = 0.41$ causes \mathcal{L} to have exactly one unstable complex eigenvalue.

The local stability analysis of fluid flows has been a powerful tool for many decades, and it has yielded successful predictions in wakes, jets, and even stratified flows. Nevertheless, we cannot escape the fact that the crux of the analysis is a crude and typically incorrect assumption on streamwise variations. Furthermore, this analysis is wholly incapable of determining the stability of flows with nontrivial streamwise dependence. With recent advances in computational power, the flow stability community has gradually been shifting its attention toward the *global* stability of fluid flows. The following sections focus on the linear formulation of such an analysis.

2.3 Flow operators

The linear global theory of flow stability is rooted in eigendecompositions of linearized flow operators. The direct and adjoint eigendecompositions of such operators yield information about the operator stability, the receptivity to initial conditions and external forcing, and the sensitivity to spatially localized perturbations. Adjoint-based analyses, in particular, have received notable attention in the fluid mechanics community; see Luchini and Bottaro (2014). This section discusses

the direct and adjoint linearizations of the Navier–Stokes equations, and it reviews the consequences of their eigendecompositions.

In the framework of incompressible, constant-density, constant-viscosity flows of Newtonian fluids, we define the nonlinear Navier–Stokes equation for a nondimensional velocity field $\mathbf{u}(\mathbf{x}, t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ (with $n = 2$ or 3 , typically), nondimensional pressure field $p(\mathbf{x}, t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$, and Reynolds number $Re > 0$ by

$$\frac{\partial \mathbf{u}}{\partial t} = \mathcal{N}\mathbf{u}, \quad (2.2a)$$

where

$$\mathcal{N} : \mathbb{R}^n \rightarrow \mathbb{R}^n \quad (2.2b)$$

$$\mathbf{u} \mapsto -\mathbf{u} \cdot \nabla \mathbf{u} - \nabla p + \frac{1}{Re} \nabla^2 \mathbf{u}, \quad (2.2c)$$

subject to the mass balance

$$\nabla \cdot \mathbf{u} = 0. \quad (2.2d)$$

The boundary conditions on \mathbf{u} and p depend on the choice of the geometry. If there exists $\{\mathbf{u}_0, p_0\}$ such that $\mathcal{N}\mathbf{u}_0 = \mathbf{0}$, then $\{\mathbf{u}_0, p_0\}$ is said to be an equilibrium, steady-state solution, or fixed point of \mathcal{N} .

We can then introduce velocity and pressure perturbations $\mathbf{u}'(\mathbf{x}, t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{C}^n$ and $p'(\mathbf{x}, t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{C}$ such that

$$\mathbf{u} = \mathbf{u}_0 + \mathbf{u}' \quad (2.3a)$$

$$p = p_0 + p', \quad (2.3b)$$

and $\|\mathbf{u}'\| \ll 1$ and $\|p'\| \ll 1$ everywhere in the domain. The linearization of \mathcal{N} about $\{\mathbf{u}_0, p_0\}$ is then

$$\frac{\partial \mathbf{u}'}{\partial t} = \mathcal{L}\mathbf{u}', \quad (2.4a)$$

where

$$\mathcal{L} : \mathbb{C}^n \rightarrow \mathbb{C}^n \quad (2.4b)$$

$$\mathbf{u}' \mapsto -\mathbf{u}' \cdot \nabla \mathbf{u}_0 - \mathbf{u}_0 \cdot \nabla \mathbf{u}' - \nabla p' + \frac{1}{Re} \nabla^2 \mathbf{u}' \quad (2.4c)$$

subject to the mass balance

$$\nabla \cdot \mathbf{u}' = 0. \quad (2.4d)$$

The boundary conditions on \mathbf{u}' and p' are typically homogeneous variants of the conditions on \mathbf{u} and p .

To define the adjoint linearized Navier–Stokes operator \mathcal{L}^* , we first define the velocity inner product over a control volume Ω by

$$\langle \cdot, \cdot \rangle : \mathbb{C}^n \times \mathbb{C}^n \rightarrow \mathbb{C} \quad (2.5a)$$

$$\mathbf{u}_1, \mathbf{u}_2 \mapsto \int_{\Omega} \bar{\mathbf{u}}_2 \cdot \mathbf{u}_1 dV, \quad (2.5b)$$

where $\overline{(\cdot)}$ denotes the complex conjugate, and the integrand contains the standard pointwise dot product. The adjoint operator \mathcal{L}^* must therefore satisfy $\langle \mathcal{L}\mathbf{u}', \hat{\mathbf{u}}' \rangle = \langle \mathbf{u}', \mathcal{L}^*\hat{\mathbf{u}}' \rangle$. Integration by

parts reveals that the adjoint operator for the adjoint velocity and pressure perturbations $\hat{\mathbf{u}}'(\mathbf{x}, t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{C}^n$ and $\hat{p}'(\mathbf{x}, t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{C}$ is

$$\mathcal{L}^* : \mathbb{C}^n \rightarrow \mathbb{C}^n \quad (2.6a)$$

$$\hat{\mathbf{u}}' \mapsto -(\nabla \mathbf{u}_0) \cdot \hat{\mathbf{u}}' + \mathbf{u}_0 \cdot \nabla \hat{\mathbf{u}}' - \nabla \hat{p}' + \frac{1}{Re} \nabla^2 \hat{\mathbf{u}}', \quad (2.6b)$$

where, in indicial notation, $((\nabla \mathbf{u}_0) \cdot \hat{\mathbf{u}}')_i = (\partial u_{0,j} / \partial x_i) \hat{u}'_j$. This operator is subject to the mass balance

$$\nabla \cdot \hat{\mathbf{u}}' = 0. \quad (2.6c)$$

Also, denoting the boundary of the volume Ω by $\partial\Omega$ and the normal vector on $\partial\Omega$ by \mathbf{n} , \mathcal{L}^* is subject to the boundary conditions

$$\oint_{\partial\Omega} \left(-(\mathbf{u}' \cdot \hat{\mathbf{u}}') \mathbf{u}_0 - \hat{\mathbf{u}}' p' + \mathbf{u}' \hat{p}' + \frac{1}{Re} ((\nabla \mathbf{u}') \cdot \hat{\mathbf{u}}' - (\nabla \hat{\mathbf{u}}') \cdot \mathbf{u}') \right) \cdot \mathbf{n} dS, \quad (2.7)$$

as a result of the integration by parts.

We briefly remark, firstly, that the derivation of the adjoint linearized Navier–Stokes operator may be clearer if we consider the expanded operators

$$\mathcal{L}_e : \mathbb{C}^{n+1} \rightarrow \mathbb{C}^{n+1} \quad (2.8a)$$

$$\begin{bmatrix} \mathbf{u}' \\ p' \end{bmatrix} \mapsto \begin{bmatrix} \mathcal{L} \mathbf{u}' \\ \nabla \cdot \mathbf{u}' \end{bmatrix} \quad (2.8b)$$

and

$$\mathcal{L}_e^* : \mathbb{C}^{n+1} \rightarrow \mathbb{C}^{n+1} \quad (2.8c)$$

$$\begin{bmatrix} \hat{\mathbf{u}}' \\ \hat{p}' \end{bmatrix} \mapsto \begin{bmatrix} \mathcal{L}^* \hat{\mathbf{u}}' \\ \nabla \cdot \hat{\mathbf{u}}' \end{bmatrix}, \quad (2.8d)$$

where the expanded inner product

$$\langle \cdot, \cdot \rangle_e : \mathbb{C}^{n+1} \times \mathbb{C}^{n+1} \rightarrow \mathbb{C} \quad (2.9a)$$

$$\begin{bmatrix} \mathbf{u}_1 \\ p_1 \end{bmatrix}, \begin{bmatrix} \mathbf{u}_2 \\ p_2 \end{bmatrix} \mapsto \int_{\Omega} (\bar{\mathbf{u}}_2 \cdot \mathbf{u}_1 + \bar{p}_2 p_1) dV \quad (2.9b)$$

sets the required adjoint relation

$$\left\langle \mathcal{L}_e \begin{bmatrix} \mathbf{u}' \\ p' \end{bmatrix}, \begin{bmatrix} \hat{\mathbf{u}}' \\ \hat{p}' \end{bmatrix} \right\rangle_e = \left\langle \begin{bmatrix} \mathbf{u}' \\ p' \end{bmatrix}, \mathcal{L}_e^* \begin{bmatrix} \hat{\mathbf{u}}' \\ \hat{p}' \end{bmatrix} \right\rangle_e. \quad (2.10)$$

Secondly, the above derivation is similar to the one found in Giannetti and Luchini (2007), who use a variant of these steps.

As is the case with most linear operators, the eigendecompositions of these operators yield particularly useful information about the underlying dynamics. The next section reviews this eigenmode-based analysis.

2.4 Eigenmode analysis: stability, receptivity, and sensitivity

In some sense, the eigendecomposition of a linearized flow operator is no different from that of other operators—it identifies modes that remain invariant (except in complex amplitude) under the action of the operator, along with associated growth rates and frequencies. The utility of linear global flow stability analysis, however, is that the eigendecomposition has the power to reveal structures that are intrinsically tied to the underlying flow physics.

The eigendecomposition of the linearized Navier–Stokes operator \mathcal{L} given in (2.4) reveals information about the flow *stability*. If we compute the eigenvalues $\lambda_j \in \mathbb{C}$ and eigenmodes $\phi_j(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{C}^n$ such that

$$\mathcal{L}\phi_j = \lambda_j\phi_j, \quad (2.11)$$

then a basic result comes from the Hartman–Grobman theorem, which Section 2.2 mentioned for local stability. In the case of global stability, the theorem states that if $\text{Re}(\lambda_j) > 0$ for any j , then \mathcal{L} and the nonlinear operator \mathcal{N} (2.2) are both unstable; if $\text{Re}(\lambda_j) < 0$ for all j , then \mathcal{L} and \mathcal{N} are both stable. In the degenerate case where $\text{Re}(\lambda_j) = 0$ for some j , then a nonlinear analysis is required to determine the stability of \mathcal{N} . Furthermore, the eigendecomposition (2.11) yields information about the long-term behavior of \mathcal{L} . When $t \rightarrow \infty$, the linearized fluid flow will be dominated by the eigenmode ϕ_j whose corresponding eigenvalue λ_j has the largest real part. As an example, Figure 2.4(a) shows the stability region—where ϕ_j is large—of a two-dimensional cylinder flow’s first instability at $Re = 50$.

Given adjoint eigenmodes $\psi_j(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{C}^n$, the eigendecomposition of the adjoint linearized Navier–Stokes operator is

$$\mathcal{L}^*\psi_j = \bar{\lambda}_j\psi_j. \quad (2.12)$$

The direct and adjoint eigenmodes satisfy the biorthogonality condition $\langle \phi_j, \psi_k \rangle = 0$ for $j \neq k$. Furthermore, the eigenvalues λ_j are the same in (2.11) and (2.12). This adjoint eigendecomposition reveals the dynamics’ *receptivity to initial conditions* in the following sense. If \mathcal{L} has a full set of eigenmodes, then the solution to the linearized Navier–Stokes equations (2.4) from some initial condition $\mathbf{u}'(0)$ (where the dependence on \mathbf{x} is implied) can be written as a sum of eigenmodes and coefficients

$$\mathbf{u}'(t) = \sum_j a_j(t)\phi_j, \quad (2.13a)$$

where

$$a_j(t) = \frac{\langle \mathbf{u}'(0), \psi_j \rangle}{\langle \phi_j, \psi_j \rangle} e^{\lambda_j t}. \quad (2.13b)$$

Therefore, a nontrivial overlap between the initial condition and the j th adjoint eigenmode indicates that the j th direct eigenmode plays a nontrivial role in the linearized solution $\mathbf{u}'(t)$. Figure 2.4(b) shows the receptivity region of the cylinder flow’s first instability, where ψ_j is large.

The adjoint eigendecomposition also reveals the dynamics’ *receptivity to external forcing*. If we consider a forced linearized Navier–Stokes equation

$$\frac{\partial \mathbf{u}'}{\partial t} = \mathcal{L}\mathbf{u}' + \mathbf{f}(t) \quad (2.14a)$$

with $\mathbf{f}(t)$ representing an external forcing, then the solution is once again (2.13a), where

$$\frac{da_j}{dt} = \lambda_j a_j(t) + \frac{\langle \mathbf{f}(t), \psi_j \rangle}{\langle \phi_j, \psi_j \rangle} \quad (2.14b)$$




Figure 2.4: The (a) stability, (b) receptivity, and (c) sensitivity regions of the two-dimensional flow over a cylinder, with the free stream flowing from left to right at $Re = 50$. From Giannetti and Luchini (2007).

with some choice of initial conditions $a_j(0)$. Similarly to before, the overlap between the forcing and the j th adjoint eigenmode indicates the contribution of j th direct eigenmode to the solution $\mathbf{u}'(t)$. The external forcing of the linearized Navier–Stokes equations is a well-studied topic with far-reaching consequences in transition and turbulence (Trefethen et al., 1993; Farrell and Ioannou, 1993; Bamieh and Dahleh, 2001; Jovanović and Bamieh, 2005).

Finally, the idea of the *sensitivity to localized perturbations* is derived from the combination of the direct and adjoint eigendecompositions. In particular, it is well-known that the direct and adjoint eigendecompositions do not individually provide intuitive explanations of dynamical trajectories from non-normal operators (Trefethen et al., 1993). Following the weakly nonparallel analysis of Section 2.2, regions of absolute instability in the flow can behave as “wavemakers” that shed waves both downstream and upstream (Chomaz, 2005). In the framework of global stability, the ideas of wavemakers and sensitivity are closely related. Furthermore, the sensitivity of linear operators to perturbations can be very large when the underlying dynamics are strongly non-normal (Chomaz, 2005).

We can formulate the sensitivity problem using a slight simplification of Giannetti and Luchini (2007). Suppose we perturb the linearized Navier–Stokes operator (2.4) from \mathcal{L} to $\mathcal{L} + d\mathcal{L}$. The perturbation of the linearized Navier–Stokes eigendecomposition (2.11) is then

$$d\mathcal{L}\phi_j + \mathcal{L}d\phi_j = d\lambda_j\phi_j + \lambda_jd\phi_j. \quad (2.15)$$

The inner product of the above with ψ_j yields

$$\langle d\mathcal{L}\phi_j, \psi_j \rangle + \langle \mathcal{L}d\phi_j, \psi_j \rangle = \langle d\lambda_j\phi_j, \psi_j \rangle + \langle \lambda_jd\phi_j, \psi_j \rangle. \quad (2.16a)$$

Note, however, that $\langle \mathcal{L}d\phi_j, \psi_j \rangle = \langle d\phi_j, \mathcal{L}^*\psi_j \rangle = \langle d\phi_j, \bar{\lambda}_j\psi_j \rangle = \langle \lambda_jd\phi_j, \psi_j \rangle$; therefore, (2.16a) simplifies to

$$\langle d\mathcal{L}\phi_j, \psi_j \rangle = \langle d\lambda_j\phi_j, \psi_j \rangle. \quad (2.16b)$$

If we now assume that $d\mathcal{L}$ is a spatially localized perturbation of strength ds at the location $\mathbf{x} = \boldsymbol{\xi}$, that is,

$$d\mathcal{L} = \delta(\mathbf{x} - \boldsymbol{\xi}) ds \quad (2.17)$$

with δ the Dirac delta function, then (2.16b) simplifies to

$$\frac{d\lambda_j}{ds} = \frac{\bar{\psi}_j(\boldsymbol{\xi}) \cdot \phi_j(\boldsymbol{\xi})}{\langle \phi_j, \psi_j \rangle}. \quad (2.18a)$$

Therefore, we may define the sensitivity function

$$\zeta(x) = \frac{\bar{\psi}_j(\mathbf{x}) \cdot \phi_j(\mathbf{x})}{\langle \phi_j, \psi_j \rangle} \quad (2.18b)$$

dictating the change in the j th eigenvalue as a result of a perturbation in \mathcal{L} localized at \mathbf{x} . Figure 2.4(c) shows the sensitivity region of the cylinder flow’s first instability, where ζ is large.

Since the numerical representations of the nonlinear flow operator \mathcal{N} , linearized flow operator \mathcal{L} , and adjoint linearized flow operator \mathcal{L}^* are typically very high-dimensional, the computation of steady states and eigendecompositions is a nontrivial task. The next section reviews a particular set of computational methods that iterate toward these solutions.

Algorithm 2.1: One-dimensional Newton's method

Data: Initial condition $u_1 \in \mathbb{R}$, function $f : \mathbb{R} \rightarrow \mathbb{R}$, small tolerance $\epsilon > 0$

Result: u_0 such that $|f(u_0)| \leq \epsilon$

```
1 for  $j = 1, 2, \dots$  do
2   if  $|f(u_j)| \leq \epsilon$  then
3     return  $u_j$ 
4   end
5    $u_{j+1} \leftarrow u_j - f(u_j)/f'(u_j)$ 
6 end
```




Figure 2.5: An example of the one-dimensional Newton's method with $f(u) = -\log(-x)$ and $u_1 = -0.05$, where the sequence u_1, u_2, \dots converges to the true root $u_0 = -1$.

2.5 Computational methods

The computation of the linear global stability analysis consists of two parts: the steady-state solution and the eigendecomposition. Although there exist many iterative algorithms for these calculations, we focus on one particular set that tends to be robust, fast, and accurate. Here, we discuss the Newton iteration with the Armijo rule, using the generalized minimal residual (GMRES) method for linear systems solving and a time-stepping approach. Later, we discuss the Arnoldi algorithm for estimating the leading eigenvalues and eigenmodes of the linearized flow operators.

Newton's method is a powerful tool that allows us to iterate towards solutions \mathbf{u}_0 satisfying $\mathcal{N}\mathbf{u}_0 = \mathbf{0}$ (see (2.2)). The one-dimensional Newton's method is a classic example of an iterative root-solving method, dating back to the 17th century. This method, which most texts in single variable calculus cover, seeks a solution $u_0 \in \mathbb{R}$ such that $f(u_0) = 0$ for a given nonlinear function $f : \mathbb{R} \rightarrow \mathbb{R}$. It iterates toward the solution by using the derivative at a given iterate to trace a tangent from the iterate to zero, where the location of the tangent's root provides the next iterate. Algorithm 2.1 presents the one-dimensional Newton's method, and Figure 2.5 shows an example. Excepting certain ill conditions, Newton's method converges quadratically (Kelley, 1995).

The representation of velocity fields \mathbf{u} and nonlinear Navier–Stokes operators \mathcal{N} , however, typically requires a very large number of dimensions. If the scalar u_{jk} is the value of the j th velocity component in the k th computational cell, then a common vector representation of a three-dimensional velocity field \mathbf{u} over q cells is the vector

$$\mathbf{u} = \begin{bmatrix} u_{11} & u_{21} & u_{31} & u_{12} & u_{22} & u_{32} & \cdots & u_{1q} & u_{2q} & u_{3q} \end{bmatrix}^T. \quad (2.19)$$

Therefore, in the context of solving for Navier–Stokes equilibria in three spatial dimensions, we must typically solve $\mathcal{N}\mathbf{u}_0 = \mathbf{0}$, where numerically, $\mathbf{u}_0 \in \mathbb{R}^{3q}$ and $\mathcal{N} : \mathbb{R}^{3q} \rightarrow \mathbb{R}^{3q}$. Presently, q is commonly in the thousands to billions. Fortunately, the multi-dimensional Newton’s method is little more than a generalization of the one-dimensional case. Instead of using derivatives to compute tangent lines, it uses gradients. Algorithm 2.2 presents the multi-dimensional Newton’s method, which Kelley (1995, 2003) present in greater detail. When searching for nonlinear Navier–Stokes equilibria, we set $\mathbf{f} = \mathcal{N}$.

When implementing the multi-dimensional Newton’s method, however, at least two major difficulties remain. The first is that line 5 of Algorithm 2.2 is computationally intractable when the dimension n of \mathbf{u} is very high. A fairly robust remedy for this problem is the use of the GMRES algorithm (Saad and Schultz, 1986; Trefethen and Bau, 1997) to solve line 5. The GMRES algorithm estimates the solution \mathbf{x} to the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ by identifying, after k iterations, the vector $\mathbf{x}_k \in \text{span}\{\mathbf{b}, \mathbf{A}\mathbf{b}, \dots, \mathbf{A}^{k-1}\mathbf{b}\}$ that minimizes $\|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|$. Here, $\|\cdot\|$ is the standard Euclidean 2-norm, and $\text{span}\{\mathbf{b}, \mathbf{A}\mathbf{b}, \dots, \mathbf{A}^{k-1}\mathbf{b}\}$ is known as a Krylov subspace. The utility of the GMRES algorithm is hidden in the above formulation: the algorithm never requires an explicit representation of \mathbf{A} ; instead, it only requires the ability to compute $\mathbf{A}\mathbf{b}$ from some input vector \mathbf{b} . Therefore, we may compute line 5 of Algorithm 2.2 by setting $\mathbf{A} = \partial\mathbf{f}/\partial\mathbf{u}|_{\mathbf{u}_j}$, $\mathbf{x} = \mathbf{h}$, and $\mathbf{b} = \mathbf{f}(\mathbf{u}_j)$ in the GMRES algorithm.

Before presenting the GMRES algorithm, we first outline the Arnoldi iteration (see Trefethen and Bau, 1997), which approximates the high-dimensional matrix decomposition $\mathbf{A} = \mathbf{Q}\mathbf{H}\mathbf{Q}^H$, with \mathbf{Q} unitary, \mathbf{H} upper Hessenberg (that is, zero beneath the first subdiagonal), and $(\cdot)^H$ denoting the complex transpose. Algorithm 2.3 presents this method, which GMRES requires. A careful observation of this algorithm reveals the mechanics of the Arnoldi iteration. In this procedure, the outer iteration advances the Krylov subspace by providing one more application of the linear operator \mathbf{A} , whereas the inner iteration projects the previously computed orthogonal directions out of the result.

With the Arnoldi iteration in place, we can define the matrices

$$\mathbf{Q}_k = \begin{bmatrix} \mathbf{q}_1 & \cdots & \mathbf{q}_k \end{bmatrix} \quad (2.20a)$$

$$\mathbf{H}_k = \begin{bmatrix} h_{1,1} & h_{1,2} & \cdots & h_{1,k} \\ h_{2,1} & h_{2,2} & \cdots & h_{2,k} \\ & \ddots & \ddots & \vdots \\ & & h_{k,k-1} & h_{k,k} \\ 0 & & & h_{k+1,k} \end{bmatrix}, \quad (2.20b)$$

which are “in-progress” versions of the unitary \mathbf{Q} and upper Hessenberg \mathbf{H} that the Arnoldi iteration builds, such that $\mathbf{A}\mathbf{Q}_k = \mathbf{Q}_{k+1}\mathbf{H}_k$. Also, let $\mathbf{e}_1 = \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix}$ be the first unit vector in \mathbb{R}^n . The GMRES algorithm for solving $\mathbf{A}\mathbf{x} = \mathbf{b}$ is then given by Algorithm 2.4. Note also that the Givens rotation solves line 4 of this algorithm more efficiently than standard least-square solvers, because of the upper Hessenberg matrix \mathbf{H}_k ’s special structure (Trefethen and Bau, 1997). We also comment that preconditioning can accelerate the convergence of the Arnoldi and GMRES algorithms, but we do not apply such methods here.

Even with the GMRES algorithm, the computation of the Jacobian–vector product $\partial\mathbf{f}/\partial\mathbf{u}|_{\mathbf{u}_j} \cdot \mathbf{q}_k$ (see line 3 of Algorithm 2.3) is usually not directly possible because of the large state dimension.

Algorithm 2.2: Multi-dimensional Newton's method

Data: Initial condition $\mathbf{u}_1 \in \mathbb{R}^n$, function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, small tolerance $\epsilon > 0$

Result: \mathbf{u}_0 such that $\|f(\mathbf{u}_0)\| \leq \epsilon$

```
1 for  $j = 1, 2, \dots$  do
2   if  $\|f(\mathbf{u}_j)\| \leq \epsilon$  then
3     return  $\mathbf{u}_j$ 
4   end
5   Solve  $\frac{\partial f}{\partial \mathbf{u}} \Big|_{\mathbf{u}_j} \cdot \mathbf{h} = f(\mathbf{u}_j)$  for  $\mathbf{h}$ 
6    $\mathbf{u}_{j+1} \leftarrow \mathbf{u}_j - \mathbf{h}$ 
7 end
```

Algorithm 2.3: The Arnoldi iteration

Data: $\mathbf{A} \in \mathbb{R}^{n \times n}$, random vector $\mathbf{b} \in \mathbb{R}^n$

Result: Approximate upper Hessenberg decomposition $\mathbf{A} \approx \mathbf{Q}\mathbf{H}\mathbf{Q}^H$ such that $\mathbf{H} = \mathbf{Q}^H \mathbf{A} \mathbf{Q}$

```
1  $\mathbf{q}_1 \leftarrow \frac{\mathbf{b}}{\|\mathbf{b}\|}$ 
2 for  $k = 1, \dots, p$  with  $p < n$  do
3    $\mathbf{v} \leftarrow \mathbf{A}\mathbf{q}_k$ 
4   for  $m = 1, \dots, k$  do
5      $h_{m,k} \leftarrow \langle \mathbf{v}, \mathbf{q}_m \rangle$ 
6      $\mathbf{v} \leftarrow \mathbf{v} - h_{m,k} \mathbf{q}_m$ 
7   end
8    $h_{k+1,k} \leftarrow \|\mathbf{v}\|$ 
9    $\mathbf{q}_{k+1} \leftarrow \frac{\mathbf{v}}{h_{k+1,k}}$ 
10 end
11  $\mathbf{Q} \leftarrow [\mathbf{q}_1 \ \dots \ \mathbf{q}_p]$ 
12  $\mathbf{H} \leftarrow \begin{bmatrix} h_{1,1} & h_{1,2} & \dots & h_{1,p} \\ h_{2,1} & h_{2,2} & \dots & h_{2,p} \\ & \ddots & \ddots & \vdots \\ 0 & & h_{p,p-1} & h_{p,p} \end{bmatrix}$ 
```

Algorithm 2.4: The GMRES algorithm

Data: $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{b} \in \mathbb{R}^n$, small tolerance $\epsilon > 0$

Result: $\mathbf{x}_k \in \text{span}\{\mathbf{b}, \mathbf{A}\mathbf{b}, \dots, \mathbf{A}^{k-1}\mathbf{b}\}$ minimizing $\|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|$, with $\|\mathbf{A}\mathbf{x}_k - \mathbf{b}\| \leq \epsilon \|\mathbf{b}\|$

```
1  $\mathbf{q}_1 \leftarrow \frac{\mathbf{b}}{\|\mathbf{b}\|}$ 
2 for  $k = 1, \dots, p$  with  $p < n$  do
3   Perform lines 3–9 of Algorithm 2.3
4   Find  $\mathbf{y} \in \mathbb{R}^k$  that minimizes  $r_k = \|\mathbf{H}_k \mathbf{y} - \|\mathbf{b}\| \mathbf{e}_1\|$ 
5   if  $r_k \leq \epsilon \|\mathbf{b}\|$  then
6     return  $\mathbf{Q}_k \mathbf{y}$ 
7   end
8 end
```

Algorithm 2.5: Multi-dimensional Newton–Armijo method

Data: Initial condition $\mathbf{u}_1 \in \mathbb{R}^n$, function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, small tolerance $\epsilon > 0$, small constant $\alpha > 0$, minimum step multiplier $\sigma_0 \in (0, 1)$, maximum step multiplier $\sigma_1 \in (0, 1)$

Result: \mathbf{u}_0 such that $\|\mathbf{f}(\mathbf{u}_0)\| \leq \epsilon$

```
1 for  $j = 1, 2, \dots$  do
2   if  $\|\mathbf{f}(\mathbf{u}_j)\| \leq \epsilon$  then
3     return  $\mathbf{u}_j$ 
4   end
5   Solve  $\left. \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \right|_{\mathbf{u}_j} \cdot \mathbf{h} = \mathbf{f}(\mathbf{u}_j)$  for  $\mathbf{h}$ 
6    $\lambda_1 \leftarrow 1$ 
7   for  $k = 1, 2, \dots$  do
8     if  $\|\mathbf{f}(\mathbf{u}_j - \lambda_k \mathbf{h})\| < (1 - \alpha \lambda_k) \|\mathbf{f}(\mathbf{u}_j)\|$  then
9       break
10    else
11      if  $k = 1$  then
12         $\lambda_2 \leftarrow \sigma_1 \lambda_1$ 
13      else
14        Estimate  $\lambda_{k+1} \in [\sigma_0 \lambda_k, \sigma_1 \lambda_k]$  minimizing  $\|\mathbf{f}(\mathbf{u}_j - \lambda_{k+1} \mathbf{h})\|$ 
15      end
16    end
17  end
18   $\mathbf{u}_{j+1} \leftarrow \mathbf{u}_j - \lambda_k \mathbf{h}$ 
19 end
```

Instead, it typically suffices to compute the finite difference approximation

$$\left. \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \right|_{\mathbf{u}_j} \cdot \mathbf{q}_k = \frac{1}{\epsilon} (\mathbf{f}(\mathbf{u}_j + \epsilon \mathbf{q}_k) - \mathbf{f}(\mathbf{u}_j)) + \mathcal{O}(\epsilon), \quad (2.21)$$

with ϵ small, via a time-stepping approach in the Arnoldi algorithm.

The GMRES algorithm tackles the first difficulty in implementing the multi-dimensional Newton’s method, which is the linear systems solution. A second difficulty is that this method has no guarantees on numerical convergence. A poor initial condition, a lack of sufficient smoothness near the solution, and errors in the function and derivative evaluations can individually cause the iteration to diverge. A way to improve the convergence of the Newton iteration is the implementation of the Armijo rule (Armijo, 1966; Kelley, 1995, 2003). In this variant of Newton’s method, we use the solution \mathbf{h} from line 5 in Algorithm 2.2 to determine the Newton step direction. In addition, we allow the possibility of scaling the Newton step size by $\lambda_k \in (0, 1]$ so that the step provides a sufficient decrease in $\|\mathbf{f}(\mathbf{u}_j)\|$. The calculation of an appropriate λ_k is iterative.

Algorithm 2.5 outlines the Newton–Armijo method, based on Kelley (1995, 2003). In this procedures, typical values (as specified in Kelley (2003)) are $\alpha = 10^{-4}$, $\sigma_0 = 0.1$, and $\sigma_1 = 0.5$. The $(1 - \alpha \lambda_k)$ term in line 8 helps prevent oscillations that occasionally occur in Newton iteration (Kelley, 1995). The constants σ_0 and σ_1 set lower and upper bounds on the allowable step size reduction at each Armijo iteration. Furthermore, the estimation in line 14 is possible using an

Algorithm 2.6: Arnoldi iteration for eigendecomposition

Data: $\mathbf{A} \in \mathbb{R}^{n \times n}$, random vector $\mathbf{b} \in \mathbb{R}^n$

Result: Estimates of $\{\lambda_j\}_{j=1}^p$ and $\{\phi_j\}_{j=1}^p$ such that $\mathbf{A}\phi_j = \lambda_j\phi_j$

- 1 Perform Algorithm 2.3
 - 2 Compute the eigendecomposition $\mathbf{H}\mathbf{v}_j = \lambda_j\mathbf{v}_j$
 - 3 $\phi_j \leftarrow \mathbf{Q}\mathbf{v}_j$ for $j = 1, \dots, p$
-

efficient three-point parabolic model based on the previous, current, and zero values of λ_k , along with the corresponding values of $\|\mathbf{f}(\mathbf{u}_j - \lambda_k \mathbf{h})\|$.

Once the Newton's method has identified the steady state solution \mathbf{u}_0 satisfying $\mathcal{N}\mathbf{u}_0 \approx \mathbf{0}$, we may employ the Arnoldi iteration once more to estimate the eigendecompositions of the direct and adjoint linearized operators, \mathcal{L} and \mathcal{L}^* . Because the Arnoldi iteration relies on the formation of a Krylov subspace, with repeated applications of the linear operator, it is generally preferable to eigendecompose the discrete-time versions of these operators. These results then transform easily back to the framework of continuous time. Therefore, define the discrete-time operators

$$\mathcal{L}_d : \mathbb{C}^n \rightarrow \mathbb{C}^n \quad (2.22a)$$

$$\mathbf{u}'(t) \mapsto \mathbf{u}'(t + \Delta t) \quad (2.22b)$$

$$\mathcal{L}_d^* : \mathbb{C}^n \rightarrow \mathbb{C}^n \quad (2.22c)$$

$$\hat{\mathbf{u}}'(t) \mapsto \hat{\mathbf{u}}'(t + \Delta t). \quad (2.22d)$$

(Note that the numerical computation performs these discrete-time operators simply by advancing the perturbations a prescribed number of time steps.) Algorithm 2.6 estimates the leading eigenvalues and eigenmodes of the spatially discretized \mathcal{L}_d and \mathcal{L}_d^* . Once we have approximated the discrete-time eigendecompositions

$$\mathcal{L}_d \phi_j = \mu_j \phi_j \quad (2.23a)$$

$$\mathcal{L}_d^* \psi_j = \bar{\mu}_j \psi_j, \quad (2.23b)$$

the usual relation $\lambda_j = (\log \mu_j)/\Delta t$ recovers the continuous-time eigenvalues. The convergence of the Arnoldi iteration toward the leading eigenvalues and eigenmodes of \mathcal{L}_d and \mathcal{L}_d^* is a topic beyond the scope of this discussion, but Lecture 34 of Trefethen and Bau (1997) describes this in detail.

Altogether, the application of Newton's method, the GMRES linear solver, the finite-difference approximation of the Jacobian-vector product, the Armijo line rule, and the Arnoldi eigendecomposition involves multiple layers of iteration and a large amount of approximation. Even so, a careful selection of initial conditions and numerical tolerances usually leads to successful calculations of steady-state solutions and eigendecompositions. Further steps, such as preconditioning the GMRES iteration, can further improve performance, but we do not explore this here. The next chapter presents an application of these numerical techniques to a fluid system with three velocity components and up to $7 \cdot 10^6$ grid cells, totaling a vector dimension of over $2 \cdot 10^7$.

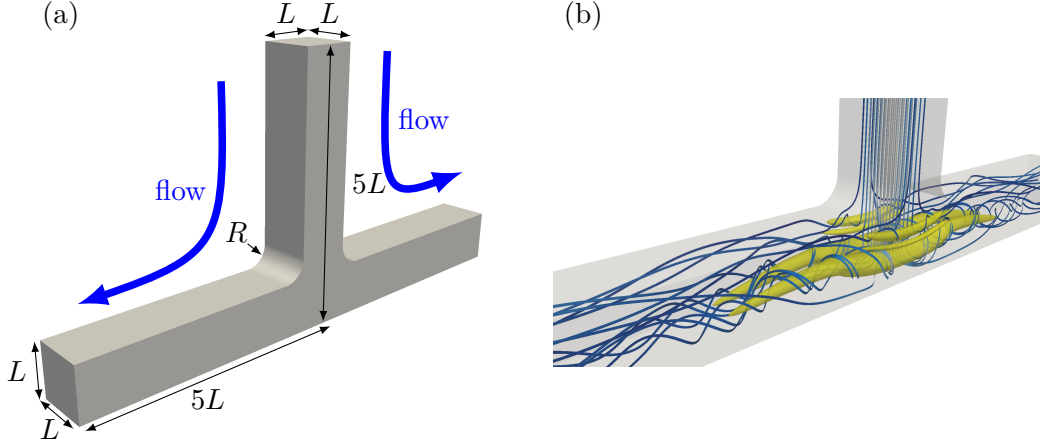


Figure 2.6: (a) The T-junction geometry. (b) A level set of the Q -criterion (yellow) with overlaid velocity streamlines (white: high speed; blue: medium speed; black: low speed) at $Re = 560$, with $R = 0.4L$. Repeated from Figure 7.1.

2.6 An example: application to an impacting pipe T-junction flow

This section demonstrates the application of the linear global stability analysis on a complex laminar three-dimensional flow. The discussion below is, in part, a concise summary of Chapter 7, which discusses vortex breakdown and sensitivity in greater detail.

The example in this section is the flow through a T-shaped pipe bifurcation, at a Reynolds number at or near 560. Impacting T-junction flows are rather ubiquitous in both natural and man-made systems, with examples including industrial pipe networks, microfluidic channels, and the pulmonary and basilar arteries in the human body. Specifically, the geometry under investigation is a pipe with a square cross section, with two 90° bends diverting the flow to two outlets. Figure 2.6(a) shows the geometry and the dimensions used in the computational model, where typically, $R = 0.4L$ to match accompanying experimental apparatuses. Given the pipe dimensions, as well as the average inlet flow speed U and the fluid kinematic viscosity ν , the Reynolds number is $Re = UL/\nu$. In this example, the laminar flow first loses stability at $Re = 556$, where it encounters a supercritical Hopf bifurcation. Figure 2.6(b) shows velocity streamlines for the steady-state solution at the nearby $Re = 560$. To aid in the visualization of the vortical behavior, this figure also shows a level set of the Q -criterion $Q = (\|\mathbf{\Omega}\|_F^2 - \|\mathbf{S}\|_F^2)/2$, where $\mathbf{\Omega} = (\nabla\mathbf{u} + (\nabla\mathbf{u})^T)/2$ is the velocity rotation tensor and $\mathbf{S} = (\nabla\mathbf{u} - (\nabla\mathbf{u})^T)/2$ is the velocity rate-of-strain tensor, and $\|\cdot\|_F$ indicates the Frobenius norm.

A major feature of interest in the impacting T-junction flow is the appearance of vortex breakdown and the subsequent formation of recirculation regions at $Re \geq 320$, when the steady-state solution is still stable. Figure 2.7(a) depicts the vortex breakdown in the flow at $Re = 560$. The downward momentum from the inlet flow produces a pair of outlet-oriented vortices in each outflow pipe. The four resulting vortices contain features common to breakdown phenomena: a very high swirl angle, approaching 90° near the breakdown points and near the center plane between the outlets; an adverse pressure gradient in the vortex core; and a streamline convergence toward the vortex core just upstream of the breakdown point (Hall, 1972). In addition, the vortex breakdown creates at least six stagnation points in the interior of the flow. Four of these are at the vortex breakdown points, with one in each quadrant of the outlet–depth plane; two additional stagnation points are on the center plane between the outlets, in the vortex core.

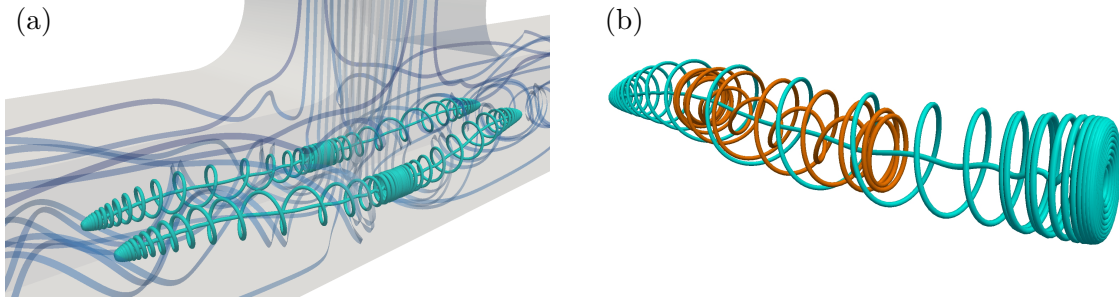


Figure 2.7: The vortex breakdown at $Re = 560$. (a) The four recirculation regions (cyan), with streamlines as in Figure 2.6(b). (b) An outer (cyan) and inner (orange) streamline from a single recirculation region. The cyan streamline's left end is the vortex breakdown point, and its right side is flush with the T-junction's center plane. Repeated from Figure 7.3.




Figure 2.8: (a) A level set of the first instability's sensitivity magnitude $|\zeta_1|$ in magenta, with streamlines of the recirculation regions in cyan. Here, $Re = 560$. (b) As (a), but with the second instability's sensitivity magnitude $|\zeta_2|$. Repeated from Figure 7.4.

The stability regions (i.e., direct eigenmodes) and receptivity regions (i.e., adjoint eigenmodes) of the impacting T-junction flow are consistent across different eigenmodes of the linearized Navier–Stokes operator. Figure 2.9(a) visualizes these regions for the first instability (i.e., the center mode) at $Re = 560$, near the supercritical Hopf bifurcation, and Figure 2.9(b) depicts the regions for the second least stable mode. The stability regions of the flow are most significant in the outflow pipes. In particular, the neutrally stable mode in Figure 2.9(a) is a modal representation of the corkscrew oscillations that the large vortices in Figure 2.6(b) undergo slightly above $Re = 556$ because of the supercritical Hopf bifurcation. The second least stable mode in Figure 2.9(b) resides further downstream, but all stability regions are effectively limited to these areas only.

Interestingly, the regions of receptivity to initial conditions and external disturbances have very little variation among the modes of the adjoint linearized Navier–Stokes operator. These regions are significant only in the junction and slightly above in the inlet; specifically, they are nontrivial only near the front and back walls of the junction. Since these regions are particularly compact, this flow only has a narrow area where the choice of initial conditions and the introduction of external disturbances would have a very large effect.

In the T-junction flow, the regions of sensitivity to spatially localized perturbations—where the direct and adjoint eigenmodes overlap—is of particular interest. Figure 2.8 shows the sensitivity regions, along with recirculation regions that Figure 2.7 previously showed. It is immediately evident that these two regions coincide in a nontrivial way. To close approximation, if not exactly, each recirculation region's outer shell is a two-dimensional manifold that is stable at the vortex




Figure 2.9: (a) The neutrally stable direct and adjoint eigenmodes of the linearized Navier-Stokes operator at $Re = 560$, shown as level sets of the direct modes' real (orange) and imaginary (red) velocity magnitudes, as well as the adjoint modes' real (green) and imaginary (purple) velocity magnitudes. Streamlines are overlaid, as in Figure 2.6(b). (b) As (a), but for the second least stable mode.

breakdown point and unstable at the fixed point on the center plane. The vortex core, on the other hand, is a backward-facing streamline (i.e., one-dimensional manifold) that is unstable at the vortex breakdown and stable at the center plane. Figure 2.7(b) depicts a single recirculation region from a different angle, as well as an interior streamline that may densely fill a two-dimensional torus if extended infinitely. The discussion of whether manifolds of different fixed points coincide or intersect is beyond the scope of this study, and would minimally require a very precise mesh refinement (Sotiropoulos et al., 2001). Nevertheless, the recirculation regions are at least approximately orbits where fluid particles roughly return to their starting point after one cycle.

The appendix of Giannetti et al. (2010) discusses the connection between recirculation and sensitivity. According to the inviscid short-wavelength approximation of Bayly (1988) and Lifschitz and Hameiri (1991), infinitesimally small flow perturbations advect with the underlying velocity field in the absence of viscosity. Our analysis does not consider infinitesimal perturbations, nor is the fluid inviscid. Nonetheless, this approximation presents the possibility that perturbations may positively feed back upon a complete recirculation, giving rise to a self-amplification mechanism. This, in turn, can cause a large sensitivity to localized perturbations in the recirculation regions, as Figure 2.8 shows. Previous research has demonstrated this connection between recirculation and sensitivity in bluff body wakes (Giannetti and Luchini, 2007; Giannetti et al., 2010) and cavity flows (Luchini et al., 2013). To the best of our knowledge, however, the T-junction flow is by far the most complex flow to date for which anyone has observed this connection.

This section presented original research in the stability, receptivity, and sensitivity regions of the flow through a square pipe with a T-shaped bifurcation. Chapter 7 extends the physical analysis of this flow. Among other topics, it discusses the effects of the radius of curvature R at the junction corners, as well as the separation regions near these corners.

Chapter 3

Feedback control theory of fluid flows

3.1 Overview

The use of engineering strategies for the automatic regulation of a system's behavior dates back many centuries. In the 18th century centrifugal (i.e., “flyball”) governor, for instance, a rapidly spinning engine shaft would cause two masses to shift outward, reducing the throttle to the engine. This, in turn, would help maintain a particular engine speed. The centrifugal governor utilizes a control strategy similar to what we now know as proportional control.

The formal study of control theory has progressed significantly, especially in the last several decades. This dissertation focuses specifically on linear time invariant (LTI) control theory, where we assume that plants (i.e., systems to be controlled) and controllers have linear dynamics that do not change with time. Although the LTI assumption is restrictive, this branch of control theory is very mature. More specifically, LTI control is split into “classical” and “modern” theories. The key difference between the two is the representation and analytical approaches for dynamical systems, with classical theory relying on frequency-domain representations, and modern theory relying on time-domain representations. These system representations are related by a simple Laplace transform.

In the typical *state space* representation of a finite-dimensional LTI input–output dynamical system, we assume n state variables, q actuator inputs, and p sensor outputs. Using the state variables $\mathbf{x}(t) : \mathbb{R} \rightarrow \mathbb{C}^n$, actuation signals $\mathbf{u}(t) : \mathbb{R} \rightarrow \mathbb{C}^q$, and sensor signals $\mathbf{y}(t) : \mathbb{R} \rightarrow \mathbb{C}^p$, along with the dynamical operator $\mathbf{A} \in \mathbb{C}^{n \times n}$, actuation-to-dynamics matrix $\mathbf{B} \in \mathbb{C}^{n \times q}$, dynamics-to-sensing matrix $\mathbf{C} \in \mathbb{C}^{p \times n}$, and direct-feedthrough matrix $\mathbf{D} \in \mathbb{C}^{p \times q}$, the state space is

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \quad (3.1a)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{D}\mathbf{u}(t). \quad (3.1b)$$

For the purpose of notation, we may call this system \mathbf{G} . The Laplace transform of \mathbf{G} (3.1), using an argument $s \in \mathbb{C}$ and transformed vectors $\hat{\mathbf{x}}(s)$, $\hat{\mathbf{u}}(s)$, and $\hat{\mathbf{y}}(s)$ in place of $\mathbf{x}(t)$, $\mathbf{u}(t)$, and $\mathbf{y}(t)$, is

$$s\hat{\mathbf{x}}(s) = \mathbf{A}\hat{\mathbf{x}}(s) + \mathbf{B}\hat{\mathbf{u}}(s) \quad (3.2a)$$

$$\hat{\mathbf{y}}(s) = \mathbf{C}\hat{\mathbf{x}}(s) + \mathbf{D}\hat{\mathbf{u}}(s), \quad (3.2b)$$

that is,

$$\hat{\mathbf{y}}(s) = \mathbf{G}(s)\hat{\mathbf{u}}(s) \quad (3.3a)$$

with the *transfer function* $\mathbf{G}(s)$ given by

$$\mathbf{G}(s) = \mathbf{C}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B} + \mathbf{D}. \quad (3.3b)$$

In the special case of real frequency analyses, it is then common to substitute $s = i\omega$, with $\omega \in \mathbb{R}$. In this case, the $\mathbf{G}(i\omega)$ dictates the steady-state gain from the input to the output if the system is stable (i.e., in state space form, all the eigenvalues of the \mathbf{A} matrix have a negative real part). Since the time-domain (3.1) and frequency-domain (3.3b) representations are equivalent, the label “ \mathbf{G} ” may refer to either.

In the feedback control of fluid flows, many early studies (e.g., Roussopoulos, 1993) only employed classical control methods, in part because these controllers are easier to design without high-accuracy fluid flow models. More recently, advanced research in feedback flow control has heralded a shift to modern methods, because of their generally superior ability to create high-performing and robust control. Kim and Bewley (2007) review the fundamentals of state space control methods for fluid flows. The time-domain and frequency-domain approaches, however, are not entirely separable. For instance, it is often natural to analyze the robustness of modern state space controllers from a transfer function perspective, identifying frequencies most likely to cause a system to lose stability or performance.

A review of modern control methods for feedback flow control would be far too large to provide here. Instead, this chapter focuses on particular concepts with far-reaching consequences for fluid flows. In particular, this chapter addresses optimal control, from the perspectives of both control design and control system architecture. Afterwards, it discusses robust control methods and analyses. This robustness discussion is particularly crucial for the real-life implementation of feedback flow control, since closed-loop control systems must be able to maintain stability and performance, even in the presence of modeling errors and other adverse external factors.

Section 3.2 reviews basics of optimal control, with a focus on \mathcal{H}_2 theory. Section 3.3 extends the \mathcal{H}_2 optimal theory by presenting original research on the simultaneous optimization over all LTI controllers, actuator placements, and sensor placements. Section 3.4 discusses robustness from the point of view of the normalized coprime stability margin and the ν -gap metric, which together form a powerful tool for general robustness analysis. Section 3.5 extends this discussion by briefly reviewing the robustifying \mathcal{H}_∞ loop-shaping controller.

3.2 \mathcal{H}_2 optimal control

The key idea in LTI optimal control is the design of the controller \mathbf{K} that—out of all possible LTI controllers—achieves the best control performance, for some choice of the control performance. This section first discusses the linear quadratic regulator (LQR) as the optimal full-state feedback controller and the Kalman filter as the optimal state estimator. Next, this section reviews the linear quadratic Gaussian (LQG)—which combines the LQR and the Kalman filter—as a special case of \mathcal{H}_2 optimal control. It also briefly comments on general \mathcal{H}_2 optimal control. Since optimal control is a standard topic in many texts (e.g., Skogestad and Postlethwaite, 2005, Section 9.3), we only provide a brief overview. Doyle et al. (1989) also provide a concise treatise on this topic.

The *linear quadratic regulator* is the full-state feedback matrix that, out of all possible matrices, minimizes a cost function based on the size of the state and actuator signals. Consider the input–state dynamics

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \quad (3.4)$$




Figure 3.1: (a): Full-state feedback system. (b): Observer.

(repeated from (3.1a)). Given a cost matrix $\mathbf{Q} \in \mathbb{C}^{n \times n}$ with $\mathbf{Q} \geq 0$ (i.e., positive semidefinite) and an input cost matrix $\mathbf{R} \in \mathbb{C}^{q \times q}$ with $\mathbf{R} > 0$ (i.e., positive definite), the LQR is the constant matrix \mathbf{F} such that full-state feedback $\mathbf{u}(t) = -\mathbf{F}\mathbf{x}(t)$ minimizes the cost

$$J = \int_0^\infty (\mathbf{x}^H(t)\mathbf{Q}\mathbf{x}(t) + \mathbf{u}^H(t)\mathbf{R}\mathbf{u}(t)) dt. \quad (3.5)$$

Therefore, the matrices \mathbf{Q} and \mathbf{R} determine how the cost function penalizes various components of the state and actuator inputs. The design of the LQR matrix is conceptually simple. First, solve the continuous algebraic Riccati equation (CARE)

$$\mathbf{A}^H\mathbf{X} + \mathbf{X}\mathbf{A} - \mathbf{X}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^H\mathbf{X} + \mathbf{Q} = \mathbf{0} \quad (3.6)$$

for the matrix $\mathbf{X} \in \mathbb{C}^{n \times n}$ satisfying $\mathbf{X} > 0$. Most computational algorithms solve this quadratic equation using a Schur decomposition of a Hamiltonian representation of (3.6). The LQR matrix is then

$$\mathbf{F} = \mathbf{R}^{-1}\mathbf{B}^H\mathbf{X}. \quad (3.7)$$

In the closed-loop feedback system (Figure 3.1(a)), the dynamics become $\dot{\mathbf{x}}(t) = (\mathbf{A} - \mathbf{B}\mathbf{F})\mathbf{x}(t)$.

Although the LQR controller has guarantees on the stability and robustness of the closed-loop system, its direct implementation is usually impossible. In most control systems (and almost certainly, in real-life fluid systems), it is not feasible to feed the controller the full knowledge of the entire state $\mathbf{x}(t)$. Therefore, it is typically necessary in optimal control to design and implement a state observer that estimates the state $\mathbf{x}(t)$ based on knowledge of the plant, actuator inputs $\mathbf{u}(t)$, and sensor outputs $\mathbf{y}(t)$. The typical observer formulation assumes an expansion of the plant (3.1) by the inclusion of state disturbances $\mathbf{d}(t)$ and sensor noise $\mathbf{n}(t)$, along with respective scaling matrices $\mathbf{W}^{1/2}$ and $\mathbf{V}^{1/2}$. That is, we define the plant

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}u(t) + \mathbf{W}^{1/2}\mathbf{d}(t) \quad (3.8a)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{V}^{1/2}\mathbf{n}(t), \quad (3.8b)$$

where we further assume the typical case $\mathbf{D} = \mathbf{0}$. In this formulation, $\mathbf{d}(t)$ and $\mathbf{n}(t)$ are white noise signals satisfying $E(\mathbf{d}\mathbf{d}^H) = \mathbf{I}$ and $E(\mathbf{n}\mathbf{n}^H) = \mathbf{I}$, where E denotes the expected value. The observer is then a dynamical system for the estimate $\hat{\mathbf{x}}(t)$ of $\mathbf{x}(t)$ given commonly by

$$\dot{\hat{\mathbf{x}}}(t) = \mathbf{A}\hat{\mathbf{x}}(t) + \mathbf{B}u(t) + \mathbf{L}(\mathbf{y}(t) - \mathbf{C}\hat{\mathbf{x}}(t)) \quad (3.9)$$

(see Figure 3.1(b)). Thus, it mimics (3.4), but also includes an error correction determined by the difference between the estimated sensor output $\mathbf{C}\hat{\mathbf{x}}(t)$ and the actual sensor output $\mathbf{y}(t)$.

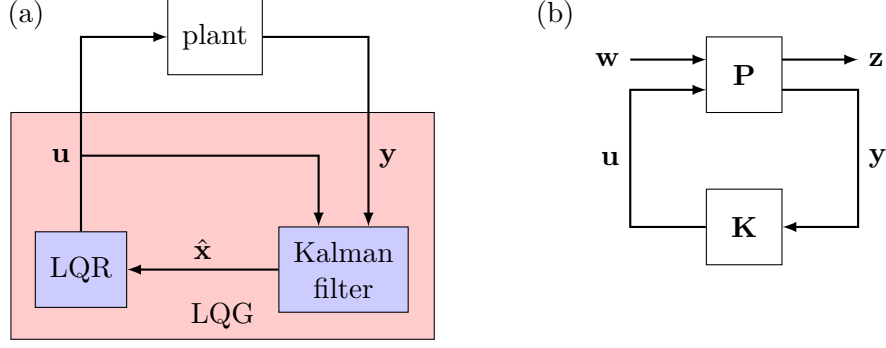


Figure 3.2: (a) The linear quadratic Gaussian controller (red) in closed-loop with a plant. (b) The lower linear fractional transform \mathbf{G} , with a plant \mathbf{P} and controller \mathbf{K} in closed-loop; repeated from Figure 8.1.

The *Kalman filter* (or linear quadratic estimator, or LQE) is specifically the design of \mathbf{L} that minimizes the mean square error $E((\hat{\mathbf{x}}(t) - \mathbf{x}(t))(\hat{\mathbf{x}}(t) - \mathbf{x}(t))^H)$. The design of the Kalman filter is the dual of the LQR design (that is, it is the LQR of the adjoint plant). Solve the CARE

$$\mathbf{A}\mathbf{Y} + \mathbf{Y}\mathbf{A}^H - \mathbf{Y}\mathbf{C}^H\mathbf{V}^{-1}\mathbf{C}\mathbf{Y} + \mathbf{W} = \mathbf{0} \quad (3.10)$$

for the matrix $\mathbf{Y} \in \mathbb{C}^{n \times n}$ with $\mathbf{Y} > 0$. The Kalman filter matrix, then, is

$$\mathbf{L} = \mathbf{Y}\mathbf{C}^H\mathbf{V}^{-1}. \quad (3.11)$$

When a feedback system consists of a LQR gain and the Kalman filter, the controller is called the *linear quadratic Gaussian*. Figure 3.2(a) shows this configuration. A key feature of the LQG controller is that it can stabilize any plant that is detectable and stabilizable (that is, the actuators can control and the sensors can observe all unstable eigenmodes of \mathbf{A}). The controller may not be robust, however; in certain cases—including in non-normal flows (see Lauga and Bewley, 2003)—LQG-controlled systems may be very close to unstable. Section 3.5 discusses control strategies that specifically attempt to improve the robustness of closed-loop systems.

The LQG controller is a special case of a more general class of optimal control known as \mathcal{H}_2 optimal control. This theory builds on the general control setup, also known as the lower linear fractional transform, shown in Figure 3.2(b). As previously, the plant outputs sensor signals \mathbf{y} to the controller, which computes actuator signals \mathbf{u} from \mathbf{y} and sends the signal back to the plant. In general, we also assume that exogenous inputs \mathbf{w} affect the plant. This encompasses all factors that are outside the direct control of the system designer, and could include state disturbances and sensor noise. In addition, we assume that we can measure some exogenous outputs \mathbf{z} from the plant, which we want the controller to minimize optimally. Thus, in the LQG controller, we set

$$\mathbf{w}(t) = \begin{bmatrix} \mathbf{d}(t) \\ \mathbf{n}(t) \end{bmatrix}, \quad \mathbf{z}(t) = \begin{bmatrix} \mathbf{Q}^{1/2}\mathbf{x}(t) \\ \mathbf{R}^{1/2}\mathbf{u}(t) \end{bmatrix}. \quad (3.12)$$

In the general \mathcal{H}_2 theory, we consider the Frobenius norm $\|\cdot\|_F$ and the transfer function representation $\mathbf{G}(s)$ of a system, and define the \mathcal{H}_2 norm as

$$\|\mathbf{G}\|_2 = \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \|\mathbf{G}(i\omega)\|_F^2 d\omega}. \quad (3.13a)$$

Therefore, the \mathcal{H}_2 norm is a measure of the average control performance over all frequencies, input directions, and output directions. Equivalently, if $\mathbf{w}(t)$ is a white noise input with unit intensity, then

$$\|\mathbf{G}\|_2 = E \left(\lim_{T \rightarrow \infty} \sqrt{\frac{1}{2T} \int_{-T}^T \|\mathbf{z}(t)\|_2^2 dt} \right). \quad (3.13b)$$

Hence, the \mathcal{H}_2 norm can also be thought of as the root-mean-square of the gain provided by \mathbf{G} for a white noise input.

The \mathcal{H}_2 *optimal controller* is the controller \mathbf{K} that, out of all LTI controllers, minimizes the \mathcal{H}_2 norm $\|\mathbf{G}\|_2$ of the closed-loop system \mathbf{G} from \mathbf{w} to \mathbf{z} , as Figure 3.2(b) shows. If the expanded plant \mathbf{P} in Figure 3.2(b) is

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{z} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B}_1 & \mathbf{B}_2 \\ \mathbf{C}_1 & \mathbf{0} & \mathbf{D}_{12} \\ \mathbf{C}_2 & \mathbf{D}_{21} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{w} \\ \mathbf{u} \end{bmatrix}, \quad (3.14)$$

and necessary conditions are met (see Doyle et al., 1989; Skogestad and Postlethwaite, 2005, Section 9.3.1), then the solutions of the CAREs

$$\mathbf{A}^H \mathbf{X} + \mathbf{X} \mathbf{A} - \mathbf{X} \mathbf{B}_2 (\mathbf{D}_{12}^H \mathbf{D}_{12})^{-1} \mathbf{B}_2^H \mathbf{X} + \mathbf{C}_1^H \mathbf{C}_1 = \mathbf{0} \quad (3.15a)$$

$$\mathbf{A} \mathbf{Y} + \mathbf{Y} \mathbf{A}^H - \mathbf{Y} \mathbf{C}_2^H (\mathbf{D}_{21} \mathbf{D}_{21}^H)^{-1} \mathbf{C}_2 \mathbf{Y} + \mathbf{B}_1 \mathbf{B}_1^H = \mathbf{0} \quad (3.15b)$$

for $\mathbf{X} > 0$ and $\mathbf{Y} > 0$ yield the matrices

$$\mathbf{F} = (\mathbf{D}_{12}^H \mathbf{D}_{12})^{-1} \mathbf{B}_2^H \mathbf{X} \quad (3.16a)$$

$$\mathbf{L} = \mathbf{Y} \mathbf{C}_2^H (\mathbf{D}_{21} \mathbf{D}_{21}^H)^{-1}. \quad (3.16b)$$

The \mathcal{H}_2 optimal controller \mathbf{K} is then given by

$$\begin{bmatrix} \dot{\hat{\mathbf{x}}} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{A} - \mathbf{B}_2 \mathbf{F} - \mathbf{L} \mathbf{C}_2 & \mathbf{L} \\ -\mathbf{F} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}} \\ \mathbf{y} \end{bmatrix}. \quad (3.17)$$

The \mathcal{H}_2 norm $\|\mathbf{G}\|_2$ of the closed-loop system \mathbf{G} (Figure 3.2(b)) is given by

$$\|\mathbf{G}\|_2^2 = \text{tr}(\mathbf{C}_1^H \mathbf{C}_1 \mathbf{Y}) + \text{tr}(\mathbf{C}_2 \mathbf{Y} \mathbf{X} \mathbf{L}) \quad (3.18a)$$

$$= \text{tr}(\mathbf{X} \mathbf{B}_1 \mathbf{B}_1^H) + \text{tr}(\mathbf{F} \mathbf{Y} \mathbf{X} \mathbf{B}_2). \quad (3.18b)$$

This section briefly reviewed the fundamentals of the LQG and \mathcal{H}_2 optimal controllers. This theory optimizes the choice of the controller \mathbf{K} given the plant dynamics, which include the effects of exogenous inputs and the calculation of cost functions. The next section extends this theory by including the placement of actuators and sensors in the space of optimization parameters.

3.3 Optimal actuator and sensor placement

Although the idea of optimal control has traditionally referred to the optimization of the feedback controller \mathbf{K} given the plant \mathbf{P} and a choice of weights, this is not the only parameter that a control designer can optimize. In every study or application of localized feedback control in partial

differential equation systems, the designer must decide—by some means or another—where he or she should physically place the actuators and sensors.

A survey of the literature in feedback flow control would reveal that researchers have utilized seemingly every reasonable placement possible: on an upstream, midstream, or downstream boundary (e.g., Roussopoulos, 1993; Aamo et al., 2007); in a near or far wake (e.g., Roussopoulos, 1993; Cohen et al., 2004); or anywhere else where an actuator or sensor would intuitively seem to be effective. Such a survey would also show, perhaps surprisingly, that few studies have rigorously analyzed which actuator and sensor locations would be most effective, and what the implications of any particular placement would be (e.g., Lauga and Bewley, 2003; Belson et al., 2013). A very large number of studies simply employ a guess-and-check or *ad hoc* procedure (e.g., Roussopoulos, 1993; Gillies, 2001), and fewer others employ relatively untested heuristics, such as placement by global direct and adjoint eigenmodes (e.g., Åkervik et al., 2007; Bagheri et al., 2009).

The original research in this section—which Chapters 8–9 and Chen and Rowley (2011) describe in greater detail—shows that the placement of actuators and sensors in a fluid flow can be just as important as the controller design itself. This section extends Section 3.2 by considering a closed-loop system \mathbf{G} , as in Figure 3.2(b), and demonstrating a perturbation technique that computes the gradient of the \mathcal{H}_2 norm $\|\mathbf{G}\|_2$ with respect to actuator and sensor locations. This allows a control designer to implement an efficient gradient-based local optimization procedure for iteration toward the optimal placement. This section then demonstrates the procedure on the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire models of fluid flow.

The perturbation technique for computing $\nabla\|\mathbf{G}\|_2$ is an improvement in computational efficiency on a method that Hiramoto et al. (2000) propose. In the general framework of \mathcal{H}_2 optimal control, we suppose that the actuation matrix \mathbf{B}_2 is a known function of the actuator positions $\{x_j^a\}_{j=1}^q$. Therefore, an infinitesimal perturbation δx_j^a in the j th actuator position induces a calculable perturbation $\delta\mathbf{B}_j$ in \mathbf{B}_2 , which by (3.15a) induces a perturbation $\delta\mathbf{X}_j$ in \mathbf{X} by the Lyapunov equation

$$(\mathbf{A} - \mathbf{B}_2\mathbf{F})^H\delta\mathbf{X}_j + \delta\mathbf{X}_j(\mathbf{A} - \mathbf{B}_2\mathbf{F}) = \mathbf{X}\delta\mathbf{B}_j\mathbf{F} + \mathbf{F}^H\delta\mathbf{B}_j^H\mathbf{X}. \quad (3.19)$$

After solving for $\delta\mathbf{X}_j$, the perturbation of (3.18a) yields

$$\frac{\partial\|\mathbf{G}\|_2^2}{\partial x_j^a} = \frac{\text{tr}(\mathbf{C}_2\mathbf{Y}\delta\mathbf{X}_j\mathbf{L})}{\delta x_j^a}. \quad (3.20)$$

In this formulation, we must solve (3.19) for each actuator in the system. Colburn et al. (2011) improved this computation by employing an adjoint so that we need only solve one Lyapunov equation. Define the linear operators

$$\mathbf{S}(\delta\mathbf{X}_j) = (\mathbf{A} - \mathbf{B}_2\mathbf{F})^H\delta\mathbf{X}_j + \delta\mathbf{X}_j(\mathbf{A} - \mathbf{B}_2\mathbf{F}) \quad (3.21a)$$

$$\mathbf{T}(\delta\mathbf{B}_j) = \mathbf{X}\delta\mathbf{B}_j\mathbf{F} + \mathbf{F}^H\delta\mathbf{B}_j^H\mathbf{X}, \quad (3.21b)$$

so that (3.19) simplifies to

$$\mathbf{S}(\delta\mathbf{X}_j) = \mathbf{T}(\delta\mathbf{B}_j). \quad (3.22)$$

Given the trace inner product $\langle\mathbf{Z}_1, \mathbf{Z}_2\rangle_{\text{tr}} = \text{tr}(\mathbf{Z}_2^H\mathbf{Z}_1)$, the adjoint of \mathbf{S} is

$$\mathbf{S}^*(\Phi) = (\mathbf{A} - \mathbf{B}_2\mathbf{F})\Phi + \Phi(\mathbf{A} - \mathbf{B}_2\mathbf{F})^H. \quad (3.23)$$

Thus, if we solve the Lyapunov equation

$$\mathbf{S}^*(\Phi) = \mathbf{L}\mathbf{C}_2\mathbf{Y} \quad (3.24)$$

exactly once for Φ , then it follows from (3.20) that

$$\frac{\partial \|\mathbf{G}\|_2^2}{\partial x_j^a} = \frac{\langle \mathbf{S}^*(\Phi), \delta \mathbf{X}_j \rangle_{\text{tr}}}{\delta x_j^a} \quad (3.25a)$$

$$= \frac{\langle \Phi, \mathbf{S}(\delta \mathbf{X}_j) \rangle_{\text{tr}}}{\delta x_j^a} \quad (3.25b)$$

$$= \frac{\langle \Phi, \mathbf{T}(\delta \mathbf{B}_j) \rangle_{\text{tr}}}{\delta x_j^a}. \quad (3.25c)$$

Note that for different actuator indices j , we retain the matrix Φ . The computation of the derivative with respect to a different actuator only involves matrix multiplication, addition, conjugate transpose, and trace operations.

The procedure for computing the gradient of $\|\mathbf{G}\|_2^2$ with respect to sensor positions is exactly analogous. Given a perturbation δx_k^s of the k th sensor position, compute the corresponding perturbation $\delta \mathbf{C}_k$ in \mathbf{C}_2 . Solve

$$(\mathbf{A} - \mathbf{L}\mathbf{C}_2)^H \Psi + \Psi (\mathbf{A} - \mathbf{L}\mathbf{C}_2) = \mathbf{X}\mathbf{B}_2\mathbf{F} \quad (3.26)$$

for Ψ . The derivative of $\|\mathbf{G}\|_2$ with respect to the k th sensor position is

$$\frac{\partial \|\mathbf{G}\|_2^2}{\partial x_k^s} = \frac{\langle \Psi, \mathbf{Y} \delta \mathbf{C}_k^H \mathbf{L}^H + \mathbf{L} \delta \mathbf{C}_k \mathbf{Y} \rangle_{\text{tr}}}{\delta x_k^s}. \quad (3.27)$$

Even without the above procedure for computing $\nabla \|\mathbf{G}\|_2$, it is always possible to approximate such a result with finite differencing. What have we gained by this gradient formulation, then? There are two chief advantages. First, the gradient algorithm is analytically exact. An approximation of $\nabla \|\mathbf{G}\|_2$ by finite differencing, on the other hand, will necessarily suffer from truncation error if the step ϵ in the actuator or sensor position is too large, and round-off error if ϵ is too small. Second, the algorithm is computationally cheaper than finite differencing. Using the adjoint-based procedure above, for any number of actuators and sensors, the computation of $\|\mathbf{G}\|_2$ requires the solution of two CAREs, and the computation of $\nabla \|\mathbf{G}\|_2^2$ requires the solution of two Lyapunov equations. If we were to estimate $\nabla \|\mathbf{G}\|_2^2$ for a system with n_a actuators and n_s sensors by finite differencing, we would have to solve $n_a + n_s$ CAREs. In the MATLAB software on a particular computer, the CARE and Lyapunov solvers have the same polynomial scaling with matrix size, with the former about 3.9 times slower than the latter. Therefore, the analytic gradient method is approximately $2.0(n_a + n_s)$ times faster than finite differencing.

The examples below compute the optimal actuator and sensor placements for two linear models of fluid flow, each given by a linear operator \mathcal{L} in some domain. In the extended plant (3.8a), \mathbf{A} is the discretization of \mathcal{L} , and \mathbf{B} and \mathbf{C} are skinny Gaussian functions modeling point actuation and sensing. Also, the state cost and state disturbance matrices \mathbf{Q} and \mathbf{W} are weighting matrices so that the LQR cost function (3.5) and the state disturbances in (3.8a) weight every spatial point in the domain x equally. The input cost and sensor noise matrices \mathbf{R} and \mathbf{V} are identity matrices scaled down to be very small compared to \mathbf{Q} and \mathbf{W} . In these examples, the chief concern is the minimization of the state perturbation, and we do not pay particular attention to sensor noise.




Figure 3.3: (a) Optimal actuator and sensor placements in the Ginzburg–Landau system. Each row shows the solution for a different number of actuators and sensors. The amplification region is shaded gray. Values of $\|\mathbf{G}\|_2$ are shown on the right. (b) As (a), but with flow disturbances introduced only at $x_d = -11.0$. The top row has two collocated actuators. Repeated from Figure 9.3.

The first demonstration of this procedure is on the linearized complex Ginzburg–Landau equation, which models perturbations in the streamwise direction for a shear flow. For the state variable $q(x) : \mathbb{R} \rightarrow \mathbb{C}$, this equation is $\dot{q} = \mathcal{L}q$, where

$$\mathcal{L} = \mu(x) - \nu \frac{\partial}{\partial x} + \gamma \frac{\partial^2}{\partial x^2} \quad (3.28)$$

and $q \rightarrow 0$ as $x \rightarrow \pm\infty$. As in Section 2.2, $\mu(x) = 0.37 - 5 \cdot 10^{-3}x^2$, $\nu = 2 + 0.4i$, and $\gamma = 1 - i$. The spatially discretized version of this equation uses Hermite pseudospectral derivatives and collocation points on 100 nodes. Figure 3.3(a) shows the optimal placements for various numbers of actuators and sensors in the system, where the gray region indicates the amplification region $\mu(x) > 0$. Figure 3.3(b) shows the optimal placement in the special case where the state disturbance matrix $\mathbf{W}^{1/2}$ only introduces a single-channel disturbance at the upstream point $x = -11$.

The second demonstration is on the combined Orr–Sommerfeld/Squire equations. These equations model the transverse variation of transverse velocity mode amplitudes $v(y, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$ and transverse vorticity mode amplitudes $\omega(y, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$, where the modes are sinusoidal in the streamwise and spanwise directions. Let y be the transverse direction, $\alpha \in \mathbb{R}$ be the streamwise wavenumber, $\beta \in \mathbb{R}$ be the spanwise wavenumber, $k = \sqrt{\alpha^2 + \beta^2}$, $u_0(y)$ be the velocity of the base flow in the streamwise direction, and Re be the Reynolds number. The combined Orr–Sommerfeld/Squire operator is

$$\mathcal{L} : \begin{bmatrix} \partial^2/\partial y^2 - k^2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix} \mapsto \begin{bmatrix} \mathcal{L}_{OS} & 0 \\ -i\beta u'_0 & \mathcal{L}_S \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix}, \quad (3.29a)$$

where

$$\mathcal{L}_{OS} = i\alpha \left(u_0'' - u_0 \left(\frac{\partial^2}{\partial y^2} - k^2 \right) \right) + Re^{-1} \left(\frac{\partial^4}{\partial y^4} - 2k^2 \frac{\partial^2}{\partial y^2} + k^4 \right) \quad (3.29b)$$

$$\mathcal{L}_S = -i\alpha u_0 + Re^{-1} \left(\frac{\partial^2}{\partial y^2} - k^2 \right). \quad (3.29c)$$

Table 3.1: Orr–Sommerfeld/Squire optimal two-actuator, two-sensor placement

Re	α	β	Orr–Sommerfeld		Squire	
			actuator	sensor	actuator	sensor
10^4	1	0	± 0.709	∓ 0.017	0	0
10^4	0	1	± 0.443	∓ 0.486	∓ 0.485	± 0.481
$5 \cdot 10^3$	1	1	± 0.628	± 0.293	∓ 0.728	∓ 0.399




Figure 3.4: (a) Leading direct (center peak, blue) and adjoint (side peaks, red) modes of the Orr–Sommerfeld system, at $Re = 10^4$, $\alpha = 1$, and $\beta = 0$, for which the Orr–Sommerfeld and Squire systems decouple. Real and imaginary parts are shown respectively as dashed and dotted lines, and the magnitude is shown as a solid line. (b) As (a), but for the Squire system. The two real parts overlap, as do the two magnitudes. Repeated from Figure 9.5.

At the endpoints of the domain, $v = v' = \omega = 0$. The base flow in this example is the plane Poiseuille flow through a channel bounded by $y = \pm 1$, with a velocity profile $u_0(y) = y^2 - 1$. To compute the discretized equations, we use Chebyshev pseudospectral differentiation on 125 nodes. Table 3.1 shows the optimal placement for different values of the Reynolds number, streamwise wavenumber, and spanwise wavenumber. In every case, the plant has one actuator for each of the vertical velocity and the vertical vorticity, and one sensor for each as well.

A common heuristic for placing actuators and sensors in a fluid flow is to put sensors where global direct eigenmodes are large, and actuators where global adjoint eigenmodes are large (e.g., Åkervik et al., 2007; Bagheri et al., 2009). This placement makes intuitive sense in light of the stability analysis in Section 2.4, and it ensures that the input–output system is stabilizable and detectable. The eigenmode approach, however, only predicts the optimal placement when the flow operator \mathcal{L} is close to normal. This is the case, for instance, in the Orr–Sommerfeld/Squire system when $\alpha = 1$ and $\beta = 0$; see Figure 3.4. In the other Orr–Sommerfeld/Squire experiments and in the Ginzburg–Landau equation (see Figures 3.3 and 3.5(a)), the eigenmode analysis poorly predicts the optimal placement. This is likely because of two factors. First, the eigenmode analysis fails to predict non-normal transient growth that the control ought to suppress. Secondly, non-normality produces a large separation between direct and adjoint eigenmodes. A large separation between actuators and sensors can introduce a nontrivial time delay for strongly convective systems.

Others have also suggested a controllability or observability analysis for actuator or sensor placement (e.g., Ma et al., 2011). We do not comment on this in detail here, besides to say that this analysis decouples actuator placement from sensor placement. This may once again lead to a large physical separation between actuators and sensors, ultimately giving rise to undesirable time delays. Chapters 8–9 and Chen and Rowley (2011) comment on this further.




Figure 3.5: (a) Leading direct (right, blue) and adjoint (left, red) mode of the Ginzburg–Landau system. Real and imaginary parts are shown respectively as dashed and dotted lines, and the magnitude is shown as a solid line. The amplification region is shaded gray. (b) The corresponding sensitivity magnitude $|\zeta(x)|$ (see (2.18b)). Repeated from Figure 9.6.

In the Ginzburg–Landau example, the sensitivity region of the flow, shown in Figure 3.5(b), is a better predictor of the optimal placements shown in Figure 3.3(a). This is in good agreement with Strykowski and Sreenivasan (1990), Roussopoulos (1993), Gillies (2001), Lauga and Bewley (2003, 2004), and Giannetti and Luchini (2007), who reported effective control by applying control in the sensitivity, wavemaker, or absolutely unstable regions of the flow. This approach is not effective for Orr–Sommerfeld/Squire control, however. In this system, it is not optimal to collocate actuators and sensors, so the localized feedback analysis breaks down.

This section briefly reviewed recent original research in optimal actuator and sensor placement. It presented an algorithm for computing the gradient of the closed-loop system’s \mathcal{H}_2 norm with respect to actuator and sensor positions, allowing for efficient optimization of the positions. It also demonstrated the procedure on the control of the Ginzburg–Landau and Orr–Sommerfeld/Squire models of fluid flow. Chapters 8–9 discuss the algorithm, basic results, and analyses in greater detail, and Chapter 10 describes characteristics of effective placement.

3.4 Stability margins and the ν -gap metric

Section 3.2 presented a branch of optimal control theory whose impact on engineering has been immeasurable. In academic studies and practical development, the number of successful optimal control applications is practically uncountable. Nonetheless, optimal control is far from a final and definitive control strategy. Real-life control systems will always contain modeling inaccuracies, as well as other disturbances and perturbations that are unaccounted for. This is especially the case in fluid systems, where LTI control ignores nonlinearities, and the use of reduced-order models can introduce further errors. Practically, a control strategy can only be successful if the feedback system can maintain stability and performance in the presence of these perturbations.

This section focuses on the normalized coprime stability margin and the ν -gap metric, which form a very powerful and generalizable framework for analyzing the robustness of feedback control systems. The essence of this theory boils down to two closely related questions. First, when a plant \mathbf{P} and controller \mathbf{K} are in a stable feedback system, how large of a perturbation can \mathbf{P} or \mathbf{K} withstand before the feedback system loses stability? Second, how far apart are a system \mathbf{P} (e.g., the fluid dynamics) and a perturbed system \mathbf{P}_p (e.g., the reduced-order fluid model) from the perspective of closed-loop stability? This section emphasizes the key results of Vinnicombe (1993) and Vinnicombe (2001), which provide more complete analyses.

To motivate this discussion, we pose the second question again, but this time more broadly: how different—in general terms—are a system \mathbf{P} and its perturbed version \mathbf{P}_p ? A natural and frequently used answer is to consider the \mathcal{H}_∞ norm of $\mathbf{P}_p - \mathbf{P}$. Given the maximum singular value function $\bar{\sigma}$, the \mathcal{H}_∞ norm of a transfer function $\mathbf{G}(s)$ is

$$\|\mathbf{G}\|_\infty = \begin{cases} \max_{\omega \in \mathbb{R}} \bar{\sigma}(\mathbf{G}(i\omega)), & \mathbf{G} \text{ stable} \\ \infty, & \text{otherwise;} \end{cases} \quad (3.30)$$

Therefore, the \mathcal{H}_∞ norm gives the worst-case gain of \mathbf{G} over all frequencies, input directions, and output directions. (This is contrasted with the \mathcal{H}_2 norm, which gives an average gain; see (3.13)). Figure 3.6 shows a visual interpretation of $\|G_p - G\|_\infty$ for a scalar system on a Nyquist plot, displaying $G(i\omega)$ and $G_p(i\omega)$ on the complex plane for $-\infty < \omega < \infty$. For a stable system $G_p - G$, the norm $\|G_p - G\|_\infty$ is the maximum distance on the complex plane, for any $\omega \in \mathbb{R}$, between $G(i\omega)$ and $G_p(i\omega)$.

The \mathcal{H}_∞ norm of $\mathbf{G}_p - \mathbf{G}$ is a simple and natural open-loop metric in the space of LTI systems. Nevertheless, a simple example can demonstrate its inadequacy as a closed-loop metric, in the presence of feedback control. Consider the three single-input, single-output (SISO) differential equations with corresponding scalar transfer functions below:

$$P_1 : \begin{cases} \frac{dx}{dt} = -0.2x + 20u \\ y = x \end{cases} \implies P_1(s) = \frac{100}{5s + 1} \quad (3.31a)$$

$$P_2 : \begin{cases} \frac{dx}{dt} = 0.2x + 20u \\ y = x \end{cases} \implies P_2(s) = \frac{100}{5s - 1} \quad (3.31b)$$

$$P_3 : \begin{cases} \frac{d}{dt} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & -5 \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} + \begin{bmatrix} 0 \\ 100 \end{bmatrix} u \\ y = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} \end{cases} \implies P_3(s) = \frac{100}{s^2 + 5s + 1}. \quad (3.31c)$$

Figure 3.7(a) shows their open-loop Nyquist plots. It is clear in this case that P_3 is more similar to P_1 in open-loop than P_2 ; in fact, $\|P_2 - P_1\|_\infty = \infty$, since P_2 is unstable. Yet, let us now consider the simple negative feedback system with unity gain in Figure 3.7(b), where u inserts a unit step signal. Figure 3.7(c) shows the resulting step responses. In closed-loop, we find that P_1 and P_2 behave quite similarly, but P_3 is very different from the two. As this example makes clear, the \mathcal{H}_∞ norm of the difference between two systems may not reveal very much about how different the systems are under feedback control.

To formalize the discussion of closed-loop stability, we first define $[\mathbf{P}, \mathbf{K}]$ to be stable if the eight transfer functions from $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$, and \mathbf{v}_4 to \mathbf{u} and \mathbf{y} in Figure 3.8 are stable. These eight transfer function relations are

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{P} \\ \mathbf{I} \end{bmatrix} (\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \begin{bmatrix} -\mathbf{K} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{v}_2 \\ \mathbf{v}_1 \end{bmatrix} + \begin{bmatrix} \mathbf{I} \\ \mathbf{K} \end{bmatrix} (\mathbf{I} - \mathbf{P}\mathbf{K})^{-1} \begin{bmatrix} \mathbf{I} & -\mathbf{P} \end{bmatrix} \begin{bmatrix} \mathbf{v}_4 \\ \mathbf{v}_3 \end{bmatrix} \quad (3.32)$$

(see Definition 12.1).

The development of a better closed-loop measure begins with the *normalized coprime factorization*. The normalized right coprime factorization is a particular construction of the system




Figure 3.6: Nyquist plot of an example system $G(i\omega)$ (blue) and a perturbed version $G_p(i\omega)$ (red). The \mathcal{H}_∞ norm of $G_p - G$ is indicated above.

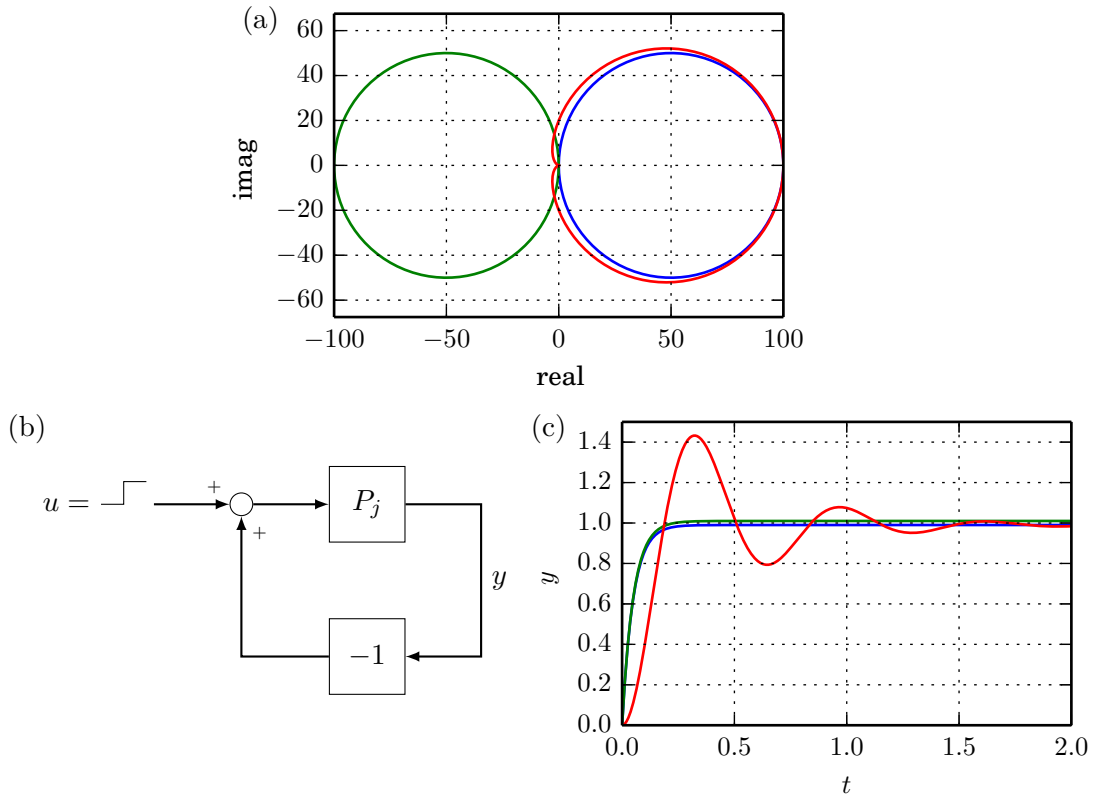


Figure 3.7: (a) The Nyquist plots of P_1 (blue), P_2 (green), and P_3 (red). (b) The closed-loop step response block diagram. (c) Their step responses in the time domain.




Figure 3.8: The standard feedback interconnection (repeated from Figure 12.2).




Figure 3.9: The normalized right coprime factorization $\mathbf{P} = \mathbf{N}\mathbf{M}^{-1}$.

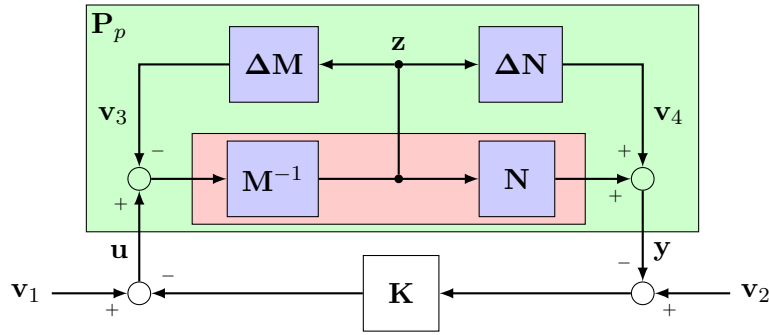


Figure 3.10: The normalized right coprime factorization uncertainty $\mathbf{P}_p = (\mathbf{N} + \Delta\mathbf{N})(\mathbf{M} + \Delta\mathbf{M})^{-1}$ (green) in closed-loop with a controller \mathbf{K} . (C.f. Figures 3.8 and 3.9.)

decomposition

$$\mathbf{P}(s) = \mathbf{N}(s)\mathbf{M}^{-1}(s) \quad (3.33)$$

(see Figure 3.9). In particular, we require that \mathbf{N} and \mathbf{M} are both stable and bounded on $i\mathbb{R}$ (even if \mathbf{P} is not), and $\mathbf{N}^H(s)\mathbf{N}(s) + \mathbf{M}^H(s)\mathbf{M}(s) = \mathbf{I}$. (For more details, see Definition 12.2, as well as Vidyasagar (1984), Meyer and Franklin (1987), Vidyasagar (1988), and Skogestad and Postlethwaite (2005, Section 4.1.5).) There also exists a normalized left coprime factorization, for which $\mathbf{P} = \mathbf{M}^{-1}\mathbf{N}$ and $\mathbf{N}\mathbf{N}^H + \mathbf{M}\mathbf{M}^H = \mathbf{I}$. In either formulation, all the right half-plane poles of $\mathbf{P}(s)$ (i.e., unstable eigenvalues of the \mathbf{A} matrix in state space form) are zeros of \mathbf{M} (i.e., values of s for which $\mathbf{M}(s)$ is rank-deficient), and all the right half-plane zeros of $\mathbf{P}(s)$ are contained in \mathbf{N} .

The utility of this decomposition is that it readily allows the representation of a perturbed plant \mathbf{P}_p by

$$\mathbf{P}_p(s) = (\mathbf{N}(s) + \Delta\mathbf{N}(s))(\mathbf{M}(s) + \Delta\mathbf{M}(s))^{-1} \quad (3.34)$$

in the case of the right factorization, where $\Delta\mathbf{N}$ and $\Delta\mathbf{M}$ model the perturbation; see Figure 3.10. Specifically, $\Delta\mathbf{M}$ models right half-plane pole uncertainties in \mathbf{P} , and $\Delta\mathbf{N}$ models right half-plane zero and time delay uncertainties in \mathbf{P} . It is for this reason that the normalized coprime factorization is a highly versatile method for modeling system uncertainties.

The key result that ties the normalized coprime uncertainty together with closed-loop uncertainty is that if $[\mathbf{P}, \mathbf{K}]$ is stable and

$$\left\| \begin{bmatrix} \Delta \mathbf{N} \\ \Delta \mathbf{M} \end{bmatrix} \right\|_{\infty} < \left\| \begin{bmatrix} \mathbf{P} \\ \mathbf{I} \end{bmatrix} (\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \begin{bmatrix} -\mathbf{K} & \mathbf{I} \end{bmatrix} \right\|_{\infty}^{-1} \quad (3.35)$$

(along with other conditions involving the representation of right half-plane poles in the system), then $[\mathbf{P}_p, \mathbf{K}]$ must be stable as well. This inequality—which is derived by analyzing the relation among \mathbf{z} , \mathbf{v}_3 , and \mathbf{v}_4 in Figure 3.10—provides strong hints as to how large a system perturbation can be before it destroys closed-loop stability, as well as how we may be able to construct a closed-loop measure of distance between systems. The inequality motivates the definition of the *normalized coprime stability margin* as

$$b_{\mathbf{P}, \mathbf{K}} = \begin{cases} \left\| \begin{bmatrix} \mathbf{P} \\ \mathbf{I} \end{bmatrix} (\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \begin{bmatrix} -\mathbf{K} & \mathbf{I} \end{bmatrix} \right\|_{\infty}^{-1}, & [\mathbf{P}, \mathbf{K}] \text{ stable} \\ 0, & \text{otherwise} \end{cases} \quad (3.36)$$

(see Definition 12.3). By construction, $b_{\mathbf{P}, \mathbf{K}} \in [0, 1]$ and $b_{\mathbf{P}, \mathbf{K}} = b_{\mathbf{K}, \mathbf{P}}$. A greater $b_{\mathbf{P}, \mathbf{K}}$ indicates a better ability of $[\mathbf{P}, \mathbf{K}]$ to reject external disturbances and to maintain robustness in the face of system uncertainties. Given the normalized coprime uncertainty representation of \mathbf{P}_p , the inequality (3.35) also motivates the ν -gap metric,

$$\delta_{\nu}(\mathbf{P}, \mathbf{P}_p) = \inf_{\Delta \mathbf{N}, \Delta \mathbf{M}} \left\| \begin{bmatrix} \Delta \mathbf{N} \\ \Delta \mathbf{M} \end{bmatrix} \right\|_{\infty}. \quad (3.37)$$

(See Definition 12.4 for a more rigorous definition.) The ν -gap metric, which satisfies $\delta_{\nu}(\mathbf{P}, \mathbf{P}_p) \in [0, 1]$ and $\delta_{\nu}(\mathbf{P}, \mathbf{P}_p) = \delta_{\nu}(\mathbf{P}_p, \mathbf{P})$ by construction, is both an open- and a closed-loop measure of the difference between two systems. In the open-loop sense, we may remove the controller from Figure 3.10 and only consider the size of $\Delta \mathbf{N}$ and $\Delta \mathbf{M}$ perturbing the plant. In the closed-loop sense, however, we have the result from (3.35) that $[\mathbf{P}_p, \mathbf{K}]$ must be stable if $\delta_{\nu}(\mathbf{P}, \mathbf{P}_p) < b_{\mathbf{P}, \mathbf{K}}$.

An even stronger statement relating the normalized coprime stability margin and the ν -gap metric, is possible. Notably, the most important relation governing the perturbation of systems is that

$$\sin^{-1} b_{\mathbf{P}_p, \mathbf{K}} \geq \sin^{-1} b_{\mathbf{P}, \mathbf{K}} - \sin^{-1} \delta_{\nu}(\mathbf{P}, \mathbf{P}_p); \quad (3.38a)$$

equivalently, since the inputs to the normalized coprime stability margin and the ν -gap metric commute,

$$\sin^{-1} b_{\mathbf{P}, \mathbf{K}_p} \geq \sin^{-1} b_{\mathbf{P}, \mathbf{K}} - \sin^{-1} \delta_{\nu}(\mathbf{K}, \mathbf{K}_p). \quad (3.38b)$$

These inequalities are tight, in that it is possible to construct plants and controllers for which these relations are equalities. We remark that mathematically, a simpler but weaker set of statements is that

$$b_{\mathbf{P}_p, \mathbf{K}} \geq b_{\mathbf{P}, \mathbf{K}} - \delta_{\nu}(\mathbf{P}, \mathbf{P}_p), \quad (3.39a)$$

$$b_{\mathbf{P}, \mathbf{K}_p} \geq b_{\mathbf{P}, \mathbf{K}} - \delta_{\nu}(\mathbf{K}, \mathbf{K}_p); \quad (3.39b)$$

(3.38) directly implies (3.39).

Returning to the previous example (3.31), we can compute specifically that $\|P_2 - P_1\|_{\infty} = \infty$ (with $\max_{\omega \in \mathbb{R}} \bar{\sigma}(P_2(i\omega) - P_1(i\omega)) = 200$) and $\|P_3 - P_1\|_{\infty} = 3.93$; yet, $\delta_{\nu}(P_1, P_2) = 0.02$ and

$\delta_\nu(P_1, P_3) = 0.702$. The ν -gap metrics are in agreement with Figure 3.7(c), which showed that in the closed-loop sense, P_2 is much closer to P_1 than P_3 is.

This section laid out the basic theory of stability margins and the ν -gap metric, which Section 4.5 and Chapter 12 implement to analyze closed-loop stability with the use of reduced-order modeling. Since reduced-order models of fluid systems are effectively perturbations of full-order systems, (3.38) can provide sufficient conditions for the stability of real fluid dynamical systems in closed-loop with reduced-order controllers. The next section briefly discusses \mathcal{H}_∞ loop-shaping, a robust control design that is based on this section's theory, and which Section 4.5 and Chapter 12 implement as well.

3.5 \mathcal{H}_∞ robust control

The discussion of the normalized coprime stability margin (3.36) and the ν -gap metric (3.37) in Section 3.4 further motivates a particular control design. Since the normalized coprime stability margin $b_{\mathbf{P}, \mathbf{K}}$ is a general measure of robustness in closed-loop, can we design a controller \mathbf{K} that maximizes $b_{\mathbf{P}, \mathbf{K}}$ for a given \mathbf{P} ? The answer is yes, and this control design is known as \mathcal{H}_∞ robust stabilization, which constitutes part of the \mathcal{H}_∞ loop-shaping technique. Since nonlinearity, non-normality, and the need for model reduction can all limit the accuracy of system models in feedback flow control design, \mathcal{H}_∞ loop-shaping is a technique that may be particularly useful for fluid flows. This brief section reviews this design, which McFarlane and Glover (1992), Vinnicombe (2001), and Section 9.4 of Skogestad and Postlethwaite (2005) present in greater detail.

First, for an n -state stable system \mathbf{G} with the standard state space representation (3.1), we briefly introduce the controllability and observability Gramians as the matrices $\mathbf{W}_c, \mathbf{W}_o \in \mathbb{C}^{n \times n}$, $\mathbf{W}_c, \mathbf{W}_o \geq 0$ given by

$$\mathbf{W}_c = \int_0^\infty e^{\mathbf{A}t} \mathbf{B} \mathbf{B}^H e^{\mathbf{A}^H t} dt \quad (3.40a)$$

$$\mathbf{W}_o = \int_0^\infty e^{\mathbf{A}^H t} \mathbf{C}^H \mathbf{C} e^{\mathbf{A}t} dt. \quad (3.40b)$$

These Gramians obey the Lyapunov equations

$$\mathbf{A} \mathbf{W}_c + \mathbf{W}_c \mathbf{A}^H + \mathbf{B} \mathbf{B}^H = \mathbf{0} \quad (3.41a)$$

$$\mathbf{A}^H \mathbf{W}_o + \mathbf{W}_o \mathbf{A} + \mathbf{C}^H \mathbf{C} = \mathbf{0}, \quad (3.41b)$$

which aid substantially in their computation. The Hankel norm of a stable \mathbf{G} is defined as the maximum ratio, over all input signals, between an output signal norm and the given input signal norm. With $\text{eig}(\cdot)$ indicating the eigenvalues of a matrix, this is given by $\|\mathbf{G}\|_H = \sqrt{\max \text{eig}(\mathbf{W}_c \mathbf{W}_o)} = \sqrt{\max \text{eig}(\mathbf{W}_o \mathbf{W}_c)}$. Returning now to \mathcal{H}_∞ robust stabilization, given a plant \mathbf{P} with normalized right coprime factorization $\mathbf{N} \mathbf{M}^{-1}$, the greatest $b_{\mathbf{P}, \mathbf{K}}$ —out of all choices of \mathbf{K} —is

$$\sup_{\mathbf{K}} b_{\mathbf{P}, \mathbf{K}} = b_{\text{opt}}(\mathbf{P}) = \sqrt{1 - \left\| \begin{bmatrix} \mathbf{M} \\ \mathbf{N} \end{bmatrix} \right\|_H^2}. \quad (3.42)$$

The computation of controllers optimizing $b_{\mathbf{P}, \mathbf{K}}$ is a standard procedure which we do not describe here. Robust control toolboxes (e.g., in MATLAB) often contain routines for this computation, and Glover and McFarlane (1989) describe suboptimal state space formulations.

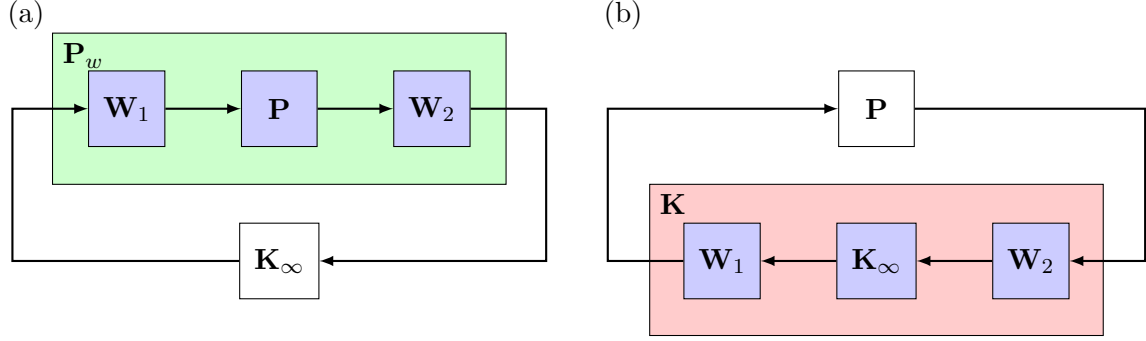


Figure 3.11: The use of weighting functions in \mathcal{H}_∞ loop-shaping design, emphasizing (a) the weighted plant, and (b) the final controller.

An important addition in \mathcal{H}_∞ loop-shaping, however, is the inclusion of weighting functions in the design procedure. In this approach, the control designer first picks a precompensator \mathbf{W}_1 and a postcompensator \mathbf{W}_2 —often by classical design techniques—so that the weighted plant $\mathbf{P}_w = \mathbf{W}_2 \mathbf{P} \mathbf{W}_1$ has an open-loop frequency response providing favorable performance; see Figure 3.11(a). The \mathcal{H}_∞ robust stabilization on \mathbf{P}_w then robustifies the loop shape by computing the controller \mathbf{K}_∞ maximizing $b_{\mathbf{P}_w, \mathbf{K}_\infty}$. To obtain \mathcal{H}_∞ loop-shaping controller \mathbf{K} , simply move the weights \mathbf{W}_1 and \mathbf{W}_2 into the robust stabilization controller, yielding $\mathbf{K} = \mathbf{W}_1 \mathbf{K}_\infty \mathbf{W}_2$; see Figure 3.11(b). The selection of weights \mathbf{W}_1 and \mathbf{W}_2 is a topic in many texts, including Vinnicombe (2001) and Skogestad and Postlethwaite (2005, Section 9.4.2). The example in Sections 4.5 and 12.4 uses the LQG as a weight; in this case, the \mathcal{H}_∞ loop-shaping can be thought of as a method of robustifying an optimal control.

This section briefly reviewed the fundamentals of \mathcal{H}_∞ robust stabilization and loop-shaping. Although \mathcal{H}_∞ robust stabilization is not an especially common controller choice for fluid flows, the maximization of robustness provides powerful safeguards for maintaining stability and performance.

Chapter 4

Reduced-order modeling of fluid flows

4.1 Overview

In LTI control theory, the idea of reduced-order modeling often applies specifically to *controller reduction*. Many modern control designs, such as \mathcal{H}_2 optimal control (see Section 3.2), \mathcal{H}_∞ optimal control (a variant employing \mathcal{H}_∞ norms instead of \mathcal{H}_2 norms), and \mathcal{H}_∞ robust control (see Section 3.5) create controllers that have at least as many states as the plant to be controlled. This kind of complexity is actually often unnecessary, and potentially burdensome in practical implementations. In many applications, it is sufficient to implement a controller approximation that has one or multiple orders of magnitude fewer states than the original controller does. With a good choice of a model reduction, a full-order controller and a reduced-order approximation can be nearly indistinguishable from a frequency response point of view. The reduced-order modeling of LTI systems is a mature field, and many rigorous results—including lower and upper bounds on errors between full- and reduced-order models—are available.

The mechanics of fluids, however, present particular challenges to the study and application of model reduction. To begin, direct modern control design from high-dimensional fluid models tends to be overwhelmingly computationally intractable. Because of modern control design algorithms' typical $\mathcal{O}(n^2)$ – $\mathcal{O}(n^3)$ computational complexity, there is an enormous difference between the 100–250-state control of one-dimensional fluid models in Section 3.3, and a hypothetical control of the 20 million-state three-dimensional T-junction flow in Section 2.6. Therefore, the application of reduced-order modeling in fluid flows typically revolves around the construction of low-order fluid dynamical approximations, from which the construction of approximately equal-order controllers would be tractable. Yet, even though model reduction techniques are designed precisely for high-dimensional systems, fluid models are typically *so* high-dimensional that modern computers cannot even compute these reductions at all!

Therefore, research in fluid flows has created a unique branch in the general study of reduced-order modeling, focusing specifically on systems so large that explicit representations of flow operators may not even be available. To date, there does not exist a single modeling technique that is vastly superior to all others, nor does there exist a single unifying approach for designing modeling techniques. Nevertheless, these methods are typically data-driven, meaning that they require little more than “snapshots” of flow fields from experiments or simulations.

In this chapter, Section 4.2 reviews standard reduced-order modeling techniques from LTI control theory, and Section 4.3 reviews some of the main reduced-order modeling techniques for high-dimensional fluid dynamics. Section 4.4 discusses one such technique, DMD, in greater detail. This section includes original research in mathematical properties of DMD, as well as a variant

of the original DMD algorithm. Finally, Section 4.5 presents original results on the stability and performance of feedback control systems that utilize reduced-order controllers.

4.2 Standard reduced-order modeling techniques

This section discusses a few standard techniques and results for the model reduction of LTI systems, including modal truncation, balanced truncation, optimal Hankel norm reduction, and normalized coprime factor model reduction. The primary purpose of this review is to place the latter sections of this chapter in the context of more firmly established theory. Section 4.5 and Chapter 12, however, implement balanced truncation and its error upper bound, which this section describes. Many standard texts in control theory, including Skogestad and Postlethwaite (2005, Chapter 11), provide a more thorough discourse on this topic.

Before discussing any particular model reduction methods, however, we first describe a concept known as Hankel singular values, which are prevalent through model reduction theory. Section 3.5 described the Hankel norm of a stable state-space system \mathbf{G} as $\sqrt{\max \text{eig}(\mathbf{W}_c \mathbf{W}_o)}$ or $\sqrt{\max \text{eig}(\mathbf{W}_o \mathbf{W}_c)}$, where \mathbf{W}_c and \mathbf{W}_o are respectively the controllability and observability Gramians of \mathbf{G} . More generally, the j th Hankel singular value $\sigma_j(\mathbf{G})$ is the square root of the j th eigenvalue of $\mathbf{W}_c \mathbf{W}_o$ or $\mathbf{W}_o \mathbf{W}_c$, with the ordering $\sigma_1(\mathbf{G}) \geq \sigma_2(\mathbf{G}) \geq \dots$. Conceptually, the Hankel singular values of a system are the singular values of the system's Hankel operator, which maps input signals for $t < 0$ to output signals for $t > 0$.

A very important and universal result for stable systems \mathbf{G} is that if \mathbf{G} has n states and \mathbf{G}_r is a different system with $r < n$ states, then the \mathcal{H}_∞ norm of the difference between the two obeys

$$\|\mathbf{G} - \mathbf{G}_r\|_\infty \geq \sigma_{r+1}(\mathbf{G}). \quad (4.1)$$

This is a model reduction error lower bound that applies to all model reduction techniques. The fact that model reduction error has a lower bound, computed only from the original system, has very important implications. It means that no matter how faithful of a modeling technique anyone may be able to develop, reduced-order models can only be as good as the original system and the reduction order permit them to be.

Modal truncation is one of the simplest methods for the model reduction of stable systems. We first produce a state space representation $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D})$ of \mathbf{G} such that \mathbf{A} is diagonal, if such a diagonalization is possible. Given the representation

$$\mathbf{A} = \begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{b}_1^H \\ \vdots \\ \mathbf{b}_n^H \end{bmatrix}, \quad \mathbf{C} = [\mathbf{c}_1 \quad \cdots \quad \mathbf{c}_n] \quad (4.2)$$

of \mathbf{G} with $\text{Re}(\lambda_1) \geq \cdots \geq \text{Re}(\lambda_n)$, the truncation of order $r < n$ yields the reduced-order system \mathbf{G}_r given by

$$\mathbf{A}_r = \begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_r \end{bmatrix}, \quad \mathbf{B}_r = \begin{bmatrix} \mathbf{b}_1^H \\ \vdots \\ \mathbf{b}_r^H \end{bmatrix}, \quad \mathbf{C}_r = [\mathbf{c}_1 \quad \cdots \quad \mathbf{c}_r]. \quad (4.3)$$

The matrix \mathbf{D} remains unchanged. This process, which eliminates the dynamics, actuation, and sensing of the most stable eigenvalues of \mathbf{A} , has the error upper bound

$$\|\mathbf{G} - \mathbf{G}_r\|_\infty \leq \sum_{j=r+1}^n \frac{\bar{\sigma}(\mathbf{c}_j \mathbf{b}_j^H)}{|\operatorname{Re}(\lambda_j)|}. \quad (4.4)$$

Although the truncation of the most stable modes is an intuitive approach to model reduction, it is usually far from the best procedure in terms of minimizing the reduction error.

A generally preferable approach, called *balanced truncation* (Moore, 1981), removes the states of stable systems that correspond to the lowest Hankel singular values. To do this, the state space is first expressed as the *balanced realization* $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D})$ such that the controllability and observability Gramians of the realization are

$$\mathbf{W}_c = \mathbf{W}_o = \begin{bmatrix} \sigma_1(\mathbf{G}) & & 0 \\ & \ddots & \\ 0 & & \sigma_n(\mathbf{G}) \end{bmatrix}. \quad (4.5)$$

As in the case of modal truncation, the balanced truncation retains the upper left $r \times r$ corner of \mathbf{A} , the first r columns of \mathbf{B} , the first r rows of \mathbf{C} , and the unchanged \mathbf{D} . Therefore, the balanced truncation of order r conceptually retains the r most equally controllable and observable modes of \mathbf{G} . Safonov and Chiang (1989) provide a numerical algorithm for computing balanced truncation.

In broad terms, we may intuitively expect this method to produce more accurate reduced-order models than modal truncation, because it specifically considers the action of the dynamics, actuation, and sensing together. Modal truncation, on the other hand, only considers the system dynamics and may retain modes that are difficult to actuate or sense. Balanced truncation has the analytical error upper bound

$$\|\mathbf{G} - \mathbf{G}_r\|_\infty \leq 2 \sum_{j=r+1}^n \sigma_j(\mathbf{G}) \quad (4.6)$$

(see Glover, 1984). Since the Hankel singular values of systems tend to fall quickly, (4.1) and (4.6) provide fairly tight *a priori* bounds on balanced truncation error, using only \mathbf{G} .

We remark briefly that the aforementioned truncation methods also have *residualization* variants, in which we force the extraneous states to have a zero time derivative, instead of setting their values to zero. Residualization forces the reduced-order model to have the same steady-state gain as the original dynamics, and is generally superior at maintaining low-frequency accuracy. Truncation, on the other hand, is generally superior at maintaining high-frequency accuracy.

A more advanced model reduction method worthy of mention is *optimal Hankel norm reduction* (Glover, 1984). Given a stable system \mathbf{G} and a desired order r , this technique finds the \mathbf{G}_r that obtains the minimum $\|\mathbf{G} - \mathbf{G}_r\|_H = \sigma_{r+1}(\mathbf{G})$. The optimal \mathbf{G}_r , however, is not unique. There exists a choice of \mathbf{G}_r , computed using an extension of balanced truncation, that has the error upper bound

$$\|\mathbf{G} - \mathbf{G}_r\|_\infty \leq \sum_{j=r+1}^n \sigma_j(\mathbf{G}). \quad (4.7)$$

Another common algorithm (Safonov et al., 1990) does not use balanced truncation and is better conditioned, but has the worse upper bound (4.6).

In the case of unstable systems, there are at least two common methods for model reduction. The first is to decompose the unstable \mathbf{G} into $\mathbf{G} = \mathbf{G}_s + \mathbf{G}_u$, where \mathbf{G}_s is stable and \mathbf{G}_u is anti-stable

with p unstable poles. After applying any model reduction method of order $r-p$ on \mathbf{G}_s to obtain the approximation \mathbf{G}_{sr} , the final reduced-order model of order r is simply $\mathbf{G}_r = \mathbf{G}_{sr} + \mathbf{G}_u$. Section 4.5 and Chapter 12 use this approach to reduce an unstable Ginzburg–Landau equation. A second common method is *normalized coprime factor model reduction* (McFarlane et al., 1990), in which we first compute the normalized right or left coprime factorization $\mathbf{G} = \mathbf{N}_r \mathbf{M}_r^{-1}$ or $\mathbf{G} = \mathbf{M}_l^{-1} \mathbf{N}_l$ (see Section 3.4). By computing the model reduction of either

$$\mathbf{G}_r = \begin{bmatrix} \mathbf{N}_r \\ \mathbf{M}_r \end{bmatrix} \quad \text{or} \quad \mathbf{G}_l = \begin{bmatrix} \mathbf{N}_l & \mathbf{M}_l \end{bmatrix}, \quad (4.8)$$

the reduced left and right factors can recombine to form the reduced system \mathbf{G}_r .

This section reviewed well-established results in the model reduction of LTI systems. Section 4.5 and Chapter 12 use some of these results to establish sufficient conditions for the stability and performance of systems in feedback with reduced-order controllers. The next section also extends this discussion by considering the case when the state size of systems is so large that this section’s methods are computationally intractable.

4.3 Techniques for very high dimensional systems

Historically, the fluid mechanics community has been chiefly responsible for developing very-high-dimensional reduced-order modeling techniques. Extremely high state sizes are necessary in the conversion of partial differential equations to ordinary differential equations such as (3.1) via spatial discretization. Of the many applications of feedback control, fluid mechanics is a unique example that consistently requires state sizes in the thousands, millions, and billions, and quite likely beyond in the future.

Unfortunately, very-high-dimensional model reduction is a far less mature field of study than the model reduction that Section 4.2 presented, and there generally exist fewer unifying themes among the existing techniques. This section briefly describes proper orthogonal decomposition (POD), balanced POD, and the eigensystem realization algorithm (ERA), which are reasonably well-established techniques. This dissertation does not use these techniques in detail, however, so the purpose of this section is only to present a context for the remaining chapters.

Proper orthogonal decomposition (see Sirovich, 1987; Holmes et al., 1996, Chapter 3) is perhaps the most venerable of the fluid model reduction techniques. This method, also known as principal component analysis and Karhunen–Loève analysis, is akin to a multi-dimensional linear regression. In this process, the model designer first collects a sequence of flow data from experimental or computational data, typically consisting of scalar or vector fields of flow variables from a domain of interest. We normally view this data on a vector space centered at the data mean. This serves two purposes: first, it removes any bias that the data mean may otherwise introduce into the modes, and second, it ensures that the data, and therefore the POD modes, will have homogeneous boundary conditions. Barring more sophisticated boundary conditions, the latter purpose gives the control designer the flexibility to superimpose any linear combination of POD modes on top of the data mean, with guaranteed satisfaction of the boundary conditions. With the mean subtracted from the data, POD finds orthogonal modes (i.e., vector directions) such that the orthogonal projection of the data onto the span of the first j modes is maximized in a least-squares sense, for any j . Figure 4.1 shows a two-dimensional example, in which the first POD captures the dominant direction of the data, and the second POD mode—which is orthogonal to the first—captures the remaining direction.




Figure 4.1: Example of POD on 100 data points in \mathbb{R}^2 . Black dots show the data points, and the blue line segment and square represent the data mean. The first POD mode is shown in red, and the second in green.

In mathematical notation, we assume a data set $\{\mathbf{x}_j\}_{j=1}^m$, with $\mathbf{x}_j \in \mathbb{C}^n$ for $j = 1, \dots, m$. Given the data mean $\bar{\mathbf{x}} = (1/m) \sum_{j=1}^m \mathbf{x}_j$, we first construct the mean-subtracted data $\mathbf{x}'_j = \mathbf{x}_j - \bar{\mathbf{x}}$. With an inner product $\langle \cdot, \cdot \rangle$, an induced norm $\|\cdot\|$, and an orthogonal projection \mathbf{P}_p onto the span of the first p POD modes, these modes $\{\mathbf{v}_j\}_{j=1}^m$ satisfy $\langle \mathbf{v}_j, \mathbf{v}_k \rangle = 0$ for $j \neq k$, and maximize $\sum_{j=1}^m \|\mathbf{P}_p \mathbf{x}'_j\|^2$ for $p = 1, \dots, m$.

When $\langle \mathbf{v}_j, \mathbf{v}_k \rangle = \mathbf{v}_k^H \mathbf{v}_j$, the numerical algorithm for computing POD is very simple. We assemble the data matrix $\mathbf{K} = [\mathbf{x}'_1 \ \cdots \ \mathbf{x}'_m]$ and compute the economy-sized singular value decomposition (SVD) $\mathbf{K} = \mathbf{U} \mathbf{\Sigma} \mathbf{W}^H$. If the diagonal of $\mathbf{\Sigma}$ is sorted by descending singular values, then the j th column of \mathbf{U} is the j th POD mode. We point out two remarks. First, if the inner product is $\langle \mathbf{x}_j, \mathbf{x}_k \rangle = \mathbf{x}_k^H \mathbf{H} \mathbf{x}_j$ for some $\mathbf{H} \in \mathbb{C}^{n \times n}$ —as is typically the case for meshes with variable cell volumes—then the orthogonality condition requires an SVD of $\mathbf{H}^{1/2} \mathbf{K}$, and the POD modes are the columns of $\mathbf{H}^{-1/2} \mathbf{U}$. Second, the *method of snapshots* is a more computationally efficient variant when $n \gg m$. Here, we compute the eigendecomposition or SVD of the small $m \times m$ matrix $\mathbf{K}^H \mathbf{K} = \mathbf{W} \mathbf{\Sigma}^2 \mathbf{W}^H$, and recover $\mathbf{U} = \mathbf{K} \mathbf{W} \mathbf{\Sigma}^{-1}$.

After computing the POD modes, the typical approach for constructing a reduced-order fluid model is a Galerkin projection (see Holmes et al., 1996, Chapter 4). In this process, we restrict the domain of the fluid dynamics to the data mean plus a linear combination of the POD modes. That is, given $p < m$ POD modes, we approximate the data $\mathbf{x}(t)$ by

$$\hat{\mathbf{x}}(t) = \bar{\mathbf{x}} + \sum_{j=1}^p a_j(t) \mathbf{v}_j. \quad (4.9)$$

By inserting this approximation into the flow dynamics, the Galerkin projection yields a (typically nonlinear) ordinary differential equation for the coefficients $\{a_j(t)\}_{j=1}^p$. Since typically $p \ll n$, this reduced-order system is far easier to iterate in time than the original dynamics.

Despite the fact that POD has been a favorite model reduction tool for many decades, it is not without fault. In particular, it suffers from two significant problems. First, although the POD procedure creates modes that maximize the “energy” $\sum_{j=1}^m \|\mathbf{P}_p \mathbf{x}'_j\|^2$ of the projected data, these are not necessarily the modes with the greatest dynamical significance. As a result, the “best” choice of POD modes to retain in Galerkin models is often difficult to determine (see, e.g., Ilak and Rowley, 2008). For a desired number $j < m$ of modes, the j modes corresponding to the greatest singular values may not yield the most accurate reduced-order model, and the selection of the j

best modes typically requires guess-and-check. Second, there is no guarantee on the stability of POD-based reduced-order models. Often, certain combinations of POD modes yield stable models, while other combinations lead to unstable models (see, e.g., Aubry et al., 1993). In these cases, there is typically little or no intuition as to which specific combinations yield unstable models.

Two attractive alternatives to POD for linear (or nearly linear) systems are balanced POD (Rowley, 2005) and ERA (Juang and Pappa, 1985), which are high-dimensional approximations of balanced truncation (see Section 4.2). Since the details of the numerical algorithms do not provide intuitive insight quite as readily, we only describe the procedures in general terms. Balanced POD—like balanced truncation, and unlike POD—assumes an input–output model (e.g., (3.1)), and it specifically considers the actions of actuators and sensors on the system dynamics. The algorithm requires an impulse response of the linear dynamics for each actuator, and an impulse response of the adjoint dynamics for each sensor. (If the number of sensors is large, then a method known as *output projection* can approximate this process with a smaller number of impulse responses.) The algorithm computes the SVD of a matrix of inner products between direct and adjoint impulse responses, and uses the decomposition to construct modes that transform the high-dimensional dynamics into a reduced-order approximation of the balanced realization. Tu and Rowley (2012) improve the numerical algorithm by using predictive techniques to reduce the amount of impulse response data needed for accurate modeling.

We briefly remark that the ERA technique is analytically equivalent to balanced POD (Ma et al., 2011), but has the advantage that it does not require adjoint impulse responses. Therefore, unlike balanced POD, it is theoretically possible to use ERA on experimental impulse response data to approximate balanced truncation. Furthermore, Ma et al. (2011) show that ERA can be less computationally expensive than balanced POD. A particular disadvantage of ERA, however, is that ERA does not reveal adjoint modes, which may otherwise be useful for observability analyses. Therefore, the choice of balanced POD and ERA depends on the system designer’s requirements. In general, however, balanced POD and ERA outperform POD when the fluid dynamics are approximately linear. Furthermore, balanced POD and ERA provide estimates of the full system’s Hankel singular values (see Section 4.2). As a result, these models obey the upper bounds on balanced truncation error (4.1, 4.6), not including any additional errors that the high-dimensional approximations may have added.

This section provided basic descriptions of POD, balanced POD, and ERA as established methods for fluid model reduction. These techniques are, in a historical sense, precursors to the novel dynamic mode decomposition, which the next section presents.

4.4 Dynamic mode decomposition

Section 4.3 discussed model reduction techniques for high-dimensional (e.g., fluid) systems that have been established in the literature. As previously mentioned, these methods typically do not apply universally—stable and faithful POD models can be difficult to construct, and balanced POD and ERA assume linear models. These restrictions can be quite limiting in practice. Dynamic mode decomposition (DMD), however, is a novel numerical technique that may be more generally applicable than other model reduction techniques. This section reviews the DMD theory that Schmid and Sesterhenn (2008) and Schmid (2010) introduced, and Rowley et al. (2009) expanded. It also presents original research that Chapter 11 and Chen et al. (2012) describe in greater detail.

Dynamic mode decomposition has at least two major interpretations. In the context of Schmid (2010), it is a numerical technique that approximates a time series of data by a linear system comprised of eigenvalues and eigenmodes. This approximation is akin to a curve fit, and is different

Algorithm 4.1: Companion matrix algorithm for DMD

Data: Data series $\{\mathbf{x}_k\}_{k=0}^m$, $\mathbf{x}_k \in \mathbb{C}^n$

Result: Modes $\{\mathbf{v}_j\}_{j=1}^m$, $\mathbf{v}_j \in \mathbb{C}^n$ and eigenvalues $\{\lambda_j\}_{j=1}^m$, $\lambda_j \in \mathbb{C}$ satisfying (4.10)

- 1 $\mathbf{K} \leftarrow [\mathbf{x}_0 \ \cdots \ \mathbf{x}_{m-1}]$
 - 2 $\mathbf{c} \leftarrow \mathbf{K}^+ \mathbf{x}_m$, with $(\cdot)^+$ the Moore–Penrose pseudoinverse
 - 3 $\mathbf{C} \leftarrow \begin{bmatrix} 0 & 0 & \cdots & 0 & c_0 \\ 1 & 0 & \cdots & 0 & c_1 \\ 0 & 1 & & 0 & c_2 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & & 1 & c_{m-1} \end{bmatrix}$, with $\mathbf{c} = \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ \vdots \\ c_{m-1} \end{bmatrix}$
 - 4 Solve the eigenvalues $\{\lambda_j\}_{j=1}^m$ of \mathbf{C}
 - 5 Construct $\mathbf{T} \in \mathbb{C}^{m \times m}$ such that $T_{jk} = \lambda_j^{k-1}$
 - 6 $\mathbf{V} = [\mathbf{v}_1 \ \cdots \ \mathbf{v}_m] \leftarrow \mathbf{K} \mathbf{T}^{-1}$
-

from linearization about a fixed point. An important factor in this interpretation is that the fitting is possible even if nonlinear dynamics generated the data. In the second interpretation of DMD, which Rowley et al. (2009) present, the algorithm numerically approximates the decomposition of an observable function into Koopman eigenfunctions and modes. In this dissertation, we focus on the first interpretation.

Given a data series $\{\mathbf{x}_k\}_{k=0}^m$ evenly spaced in time, with $\mathbf{x}_k \in \mathbb{C}^n$, the DMD of the data is the modes (or *Ritz vectors*) $\{\mathbf{v}_j\}_{j=1}^m$, $\mathbf{v}_j \in \mathbb{C}^n$ and eigenvalues (or *Ritz values*) $\{\lambda_j\}_{j=1}^m$, $\lambda_j \in \mathbb{C}$ such that

$$\mathbf{x}_k = \sum_{j=1}^m \lambda_j^k \mathbf{v}_j, \quad k = 0, \dots, m-1 \quad (4.10a)$$

$$\mathbf{x}_m = \sum_{j=1}^m \lambda_j^m \mathbf{v}_j + \mathbf{r}, \quad \mathbf{r} \perp \text{span}(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}), \quad (4.10b)$$

(repeated from (11.6)). This representation “curve-fits” the potentially nonlinearly generated data to the linear model by enforcing an exact representation for all data points but the last (4.10a), and allowing an orthogonal residual on the last data point (4.10b).

There are many algorithms for computing of DMD. Algorithm 4.1, which uses a companion and a Vandermonde matrix, is perhaps the most intuitive in light of the structure of (4.10), but can be ill-conditioned in practice. Algorithm 4.2 uses an SVD and is better conditioned. Finally, Algorithm 4.3 is a less computationally expensive version of Algorithm 4.2 that uses a method of snapshots approach. We provide two remarks. First, as in the case of POD (see Section 4.3), uneven cell sizes in discretized partial differential equations often lead to the inner product $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle = \mathbf{x}_2^H \mathbf{H} \mathbf{x}_1$ representing a volume integral, where $\mathbf{H} \in \mathbb{C}^{n \times n}$ is a weighting matrix. In this case, the orthogonality condition (4.10b) requires an application of the DMD algorithms on $\mathbf{H}^{1/2} \mathbf{K}$, in which case the modes are the columns of $\mathbf{H}^{-1/2} \mathbf{V}$. Second, the algorithms prove the existence of DMD formulation by construction. Chapter 11 and Chen et al. (2012) further prove that the choice of $\{\lambda_j\}_{j=1}^m$ and $\{\mathbf{v}_j\}_{j=1}^m$ in (4.10) is unique up to a reordering in j , if and only if $\{\mathbf{x}_k\}_{k=0}^{m-1}$ are linearly independent and $\{\lambda_j\}_{j=1}^m$ are distinct.

Algorithm 4.2: SVD algorithm for DMD

Data: Data series $\{\mathbf{x}_k\}_{k=0}^m$, $\mathbf{x}_k \in \mathbb{C}^n$

Result: Modes $\{\mathbf{v}_j\}_{j=1}^m$, $\mathbf{v}_j \in \mathbb{C}^n$ and eigenvalues $\{\lambda_j\}_{j=1}^m$, $\lambda_j \in \mathbb{C}$ satisfying (4.10)

- 1 $\mathbf{K} \leftarrow [\mathbf{x}_0 \ \cdots \ \mathbf{x}_{m-1}]$
 - 2 Solve the SVD $\mathbf{K} = \mathbf{U}\mathbf{\Sigma}\mathbf{W}^H$
 - 3 Truncate the parts of \mathbf{U} , $\mathbf{\Sigma}$, and \mathbf{W} corresponding to trivial singular values
 - 4 $\mathbf{K}_* \leftarrow [\mathbf{x}_1 \ \cdots \ \mathbf{x}_m]$
 - 5 Solve the eigendecomposition $\mathbf{U}^H \mathbf{K}_* \mathbf{W} \mathbf{\Sigma}^{-1} = \mathbf{Y} \mathbf{\Lambda} \mathbf{Y}^{-1}$
 - 6 $[\lambda_1 \ \cdots \ \lambda_m] \leftarrow \text{diag}(\mathbf{\Lambda})$
 - 7 $\mathbf{V} = [\mathbf{v}_1 \ \cdots \ \mathbf{v}_m] \leftarrow \mathbf{U} \mathbf{Y}$
-

Algorithm 4.3: SVD method of snapshots algorithm for DMD

Data: Data series $\{\mathbf{x}_k\}_{k=0}^m$, $\mathbf{x}_k \in \mathbb{C}^n$

Result: Modes $\{\mathbf{v}_j\}_{j=1}^m$, $\mathbf{v}_j \in \mathbb{C}^n$ and eigenvalues $\{\lambda_j\}_{j=1}^m$, $\lambda_j \in \mathbb{C}$ satisfying (4.10)

- 1 $\mathbf{K} \leftarrow [\mathbf{x}_0 \ \cdots \ \mathbf{x}_{m-1}]$
 - 2 Solve the eigendecomposition or SVD $\mathbf{K}^H \mathbf{K} = \mathbf{W} \mathbf{\Sigma}^2 \mathbf{W}^H$
 - 3 Truncate the parts of \mathbf{W} and $\mathbf{\Sigma}$ corresponding to trivial eigenvalues or singular values
 - 4 $\mathbf{U} \leftarrow \mathbf{K} \mathbf{W} \mathbf{\Sigma}^{-1}$
 - 5 Complete Algorithm 4.2, starting at line 4
-

There exists at least some evidence that—unlike POD and balanced POD—DMD may be able to represent dynamics in linear regimes, in periodic orbits, between these two, and perhaps in nonlinear dynamical systems in general. Rowley et al. (2009) demonstrates that in linear systems, the Koopman modes (which DMD numerically approximates) are exactly the system’s eigenmodes, and in periodic orbits, the Koopman modes are exactly the discrete Fourier modes. To demonstrate DMD’s ability to represent dynamics more generally, we may consider the fluid flow over a two-dimensional cylinder, as the flow gradually evolves from a linear regime to a limit cycle regime. Figure 4.2 shows that in this example, the DMD of a window of data also gradually morphs from a linear mode to a Fourier mode. This suggests that DMD may be broadly applicable.

With this versatile modal decomposition, then, is it possible to construct a Galerkin model as Section 4.3 did for POD? Since DMD modes have dynamical significance, and the significance is not restricted to specific dynamical regimes, we may expect such models to yield better Galerkin models than POD. The answer to this question is yes, but there is an important subtlety. In Section 4.3, we ensure that the Galerkin model (4.9) satisfies boundary conditions correctly by constructing a linear superposition of modes with homogeneous boundary conditions, on top of a data mean. Chapter 11 and Chen et al. (2012), however, show that the DMD of a mean-subtracted data set is analytically identical to the discrete Fourier transform (DFT). Given a data set $\{\mathbf{x}_k\}_{k=0}^m$, the DMD of the set

$$\mathbf{x}'_k = \mathbf{x}_k - \frac{1}{m+1} \sum_{j=0}^m \mathbf{x}_j \quad (4.11)$$

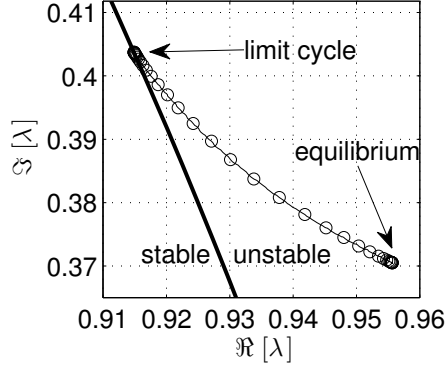


Figure 4.2: The movement of the primary unstable Ritz value near the equilibrium of a flow over a cylinder (see Figures 11.3(a) and 11.4(c–e)) to the primary oscillatory Ritz value near the limit cycle (see Figures 11.3(e) and 11.7(a–c)), as a window of roughly one period traverses the data. The unit circle is shown as a thick solid curve, and the movement of the Ritz value is shown as a thin solid curve with circles. Repeated from Figure 11.8.

for $k = 0, \dots, m$ always returns the eigenvalues

$$\lambda_j = \exp\left(\frac{2\pi i j}{m+1}\right), \quad j = 1, \dots, m \quad (4.12)$$

(repeated from (11.31)), and thus, the DMD reconstruction (4.10) becomes the DFT

$$\mathbf{x}'_k = \sum_{j=1}^m \exp\left(\frac{2\pi i j k}{m+1}\right) \mathbf{v}_j, \quad k = 0, \dots, m \quad (4.13)$$

(repeated from (11.32)). This is generally undesirable—the DFT assumes that all modes have neither growth nor decay, and it assumes that frequencies are evenly spaced on the unit circle by (4.12). In some sense, this defeats the original intention of DMD as a means of computing growth rates and frequencies for nonlinearly generated data. A seemingly better alternative for ensuring that modes have homogeneous boundary conditions is to subtract a stable or unstable fixed point from the data. This is akin to the affine transformation used in linearization, but does not otherwise alter the data. As examples in Chapter 11 will show, the equilibrium subtraction appears to preserve the DMD eigenvalues and the qualitative features of the DMD modes.

The DMD equations in (4.10) also inspires a question about the formulation itself. Why should we place a residual on the reconstruction of the last data point (4.10b), and not any other (4.10a)? Conceptually, it is not as if the last data point has any unique importance attached to it. An alternate setup that does not place a bias on the last data point is the optimized DMD of Chapter 11 and Chen et al. (2012). In this decomposition, the model designer can pick the number of desired output modes. The decomposition then produces a set of modes and eigenvalues such that the reconstruction of every data point may contain a residual, but the overall residual is minimized. Figure 4.3 shows a conceptual comparison between the reconstructions provided by DMD and optimized DMD.

In quantitative terms, suppose that there exist p data vectors, evenly sampled, given by $\{\mathbf{x}_k\}_{k=0}^{p-1}$, with $\mathbf{x}_k \in \mathbb{C}^n$. If the desired number of modes is $m < p$, then the optimized DMD consists of the




Figure 4.3: Conceptual picture of DMD (blue) vs. optimized DMD (red). DMD requires an exact representation at every point except the last, which typically has a residual. Optimized DMD allows a residual at every point, but minimizes the overall residual.

Algorithm 4.4: Iterative algorithm for optimized DMD

Data: Data series $\{\mathbf{x}_k\}_{k=0}^{p-1}$, $\mathbf{x}_k \in \mathbb{C}^n$
Result: Modes $\{\mathbf{v}_j\}_{j=1}^m$, $\mathbf{v}_j \in \mathbb{C}^n$ and eigenvalues $\{\lambda_j\}_{j=1}^m$, $\lambda_j \in \mathbb{C}$ satisfying (4.14)

- 1 Guess an initial condition $\{\lambda_j\}_{j=1}^m$, e.g., by DMD
- 2 $\mathbf{K} \leftarrow [\mathbf{x}_0 \ \cdots \ \mathbf{x}_{p-1}]$
- 3 **repeat**
- 4 Construct $\mathbf{T} \in \mathbb{C}^{m \times p}$ such that $T_{jk} = \lambda_j^{k-1}$
- 5 $J \leftarrow \text{tr}(\mathbf{T}^+ \mathbf{T} \mathbf{K}^H \mathbf{K})$, with $(\cdot)^+$ the Moore–Penrose pseudoinverse
- 6 Update $\{\lambda_j\}_{j=1}^m$ to maximize J , e.g., by Powell’s Method or global optimization
- 7 **until** J maximized
- 8 $\mathbf{V} = [\mathbf{v}_1 \ \cdots \ \mathbf{v}_m] \leftarrow \mathbf{K} \mathbf{T}^+$

modes $\{\mathbf{v}_j\}_{j=1}^m$, $\mathbf{v}_j \in \mathbb{C}^n$ and eigenvalues $\{\lambda_j\}_{j=1}^m$, $\lambda_j \in \mathbb{C}$ such that

$$\mathbf{x}_k = \sum_{j=1}^m \lambda_j^k \mathbf{v}_j + \mathbf{r}_k, \quad k = 0, \dots, p-1, \quad (4.14a)$$

and

$$\Gamma = \sum_{k=0}^{p-1} \|\mathbf{r}_k\|_2^2 \quad (4.14b)$$

is minimized (repeated from (11.35)). At present, there does not exist an analytical algorithm for computing the optimized DMD. Algorithm 4.4 presents a technique that uses an unconstrained optimization approach. Sparsity-promoting DMD—as introduced by Jovanović et al. (2014)—solves a problem similar to optimized DMD using convex optimization, and as such, is more readily applicable.

Optimized DMD has a number of significant advantages over the original DMD formulation. First, it allows the user to choose the number of output modes and eigenvalues. On the other hand, DMD must necessarily output a number of modes one less than the number of data vectors. For model reduction, this gives rise to a problem where—as with POD—the selection of modes to

retain may be rather arbitrary. Second, because of the first point, optimized DMD results are often less spurious than DMD results. Schmid (2010) and Bagheri (2013) report that DMD spectra often contain parabolic arrangements of eigenvalues where a single mode should exist. Such a spread of modes may arise from convective and diffusive processes, as well as noise and nonlinearity. The advantage of optimized DMD is that the restriction of the decomposition to a small number of modes allows the decomposition to “hone in” on the chief modes of interest, without cluttering the spectra with extraneous modes. Third, optimized DMD is more robust to sensor noise than DMD. The conceptual comparison between the DMD algorithms and polynomial fitting makes the reason clear. Optimized DMD is a global trajectory analysis, whereas DMD is chiefly concerned with the relation between data pairs $\{\mathbf{x}_k, \mathbf{x}_{k+1}\}$, for $k = 0, \dots, m-1$. Therefore, in the case of DMD, sensor noise can easily obfuscate the dynamics driving \mathbf{x}_k to \mathbf{x}_{k+1} . Furthermore, the perfect fit required by (4.10a)—like the order m polynomial fit over $m+1$ data points—can be meaningless for noisy data. On the other hand, if the sensor noise is unbiased, then optimized DMD can filter it out, as a low-order polynomial fit would.

Optimized DMD, however, also has a few disadvantages compared to DMD. The most serious concern is that there does not exist a direct algorithm for its computation. The iterative algorithm (Algorithm 4.4) can be slow, and if the user does not choose the iteration parameters well, it can converge to a globally suboptimal solution. Another disadvantage is that optimized DMD is more sensitive to process disturbances than DMD. The reason for this is, in some sense, the same reason that DMD is more sensitive to sensor noise. In the case of DMD, the comparison between data pairs $\{\mathbf{x}_k, \mathbf{x}_{k+1}\}$ for $k = 0, \dots, m-1$ can filter out zero-mean process disturbances. Optimized DMD, however, may fail to yield meaningful results, because process noise can alter the global trajectory of the data set.

This section reviewed the fundamentals of DMD, as well as specific contributions to the theory relating to uniqueness, boundary conditions and base flow subtraction, and algorithmic improvements. Since the numerical examples for these concepts are involved, we do not include them here. Instead, Chapter 11 and Chen et al. (2012) analyze these concepts for the fluid flow over a two-dimensional cylinder at a Reynolds number of 60. To the best of our knowledge, very few—possibly none—have successfully constructed and analyzed reduced-order fluid models based on DMD modes. Nevertheless, DMD shows great promise because of its ability to apply generally to nonlinear dynamics.

4.5 Stability and performance with reduced-order controllers

So far, this chapter has discussed model reduction techniques from standard LTI theory, as well as research in high-dimensional systems. As aforementioned, the typical application of these methods in feedback flow control is the construction of low-order approximations for fluid dynamics, from which control design is tractable. Presumably, a controller ought to stabilize the plant it is designed to control; yet, can a controller designed from a *reduced-order* plant successfully control the *full-order* plant? Most studies in feedback flow control simply assume so (or they do not discuss failed attempts at control design!), but a better understanding of this issue can streamline feedback flow control design. This section uses an analytical approach to answer this question.

Previously, Section 3.4 discussed the normalized coprime stability margin as a quantity that describes the performance and robustness of plants in closed-loop with controllers. That section showed that the ν -gap metric is a closed-loop measure of the difference between two systems, in that it provides tight bounds on the normalized coprime stability margin between a perturbed plant and a controller (3.38). This result has far-reaching effects, particularly in that if the closed-loop

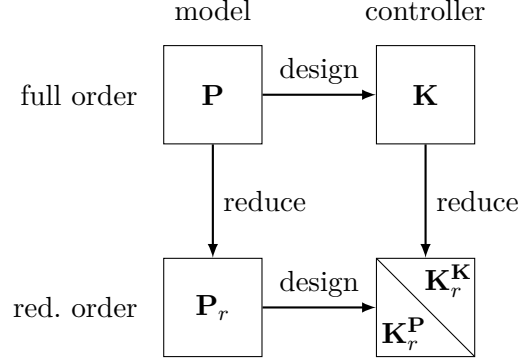


Figure 4.4: Two approaches for arriving at different reduced-order controllers \mathbf{K}_r from a full-order model \mathbf{P} , as used in this study (c.f. Obinata and Anderson, 2001, Figure 3.1.1). Repeated from Figure 12.1.

system $[\mathbf{P}, \mathbf{K}]$ is stable, then the ν -gap metric $\delta_\nu(\mathbf{P}, \mathbf{P}_p)$ between \mathbf{P} and a perturbed plant \mathbf{P}_p can potentially predict, *a priori*, if $[\mathbf{P}_p, \mathbf{K}]$ will be stable as well.

Although the ν -gap theory makes no particular assumptions about the relation between a system \mathbf{G} and its perturbation \mathbf{G}_p , a natural interpretation in the context of this chapter is that one of the systems is a reduced-order approximation of the other. In this context, the ν -gap metric $\delta_\nu(\mathbf{G}, \mathbf{G}_p)$ is a closed-loop measure of how “different” a system and its reduced-order model are. In turn, this metric can potentially predict if a controller designed from a reduced-order plant can stabilize the full plant, or alternatively, if a reduced-order approximation of a full-order controller can stabilize a plant (see Figure 4.4). Further results are available if the model reduction technique has known upper bounds on the approximation error, as Sections 4.2–4.3 describe. In this case, the ν -gap theory and the error upper bounds can reveal model reduction orders that are sufficient for the stability or performance of a full-order plant in closed-loop with a reduced-order controller.

This section summarizes the key results of Chapter 12 and Chen and Rowley (2013). The discussion picks up directly from Sections 3.4–3.5, and it omits the proofs of its results. In addition, it only examines the “reduce-then-design” approach (see Figure 4.4), which is the common feedback flow control method in which a designer constructs a controller $\mathbf{K}_r^{\mathbf{P}}$ from a reduced-order fluid model \mathbf{P}_r approximating the full fluid dynamics \mathbf{P} . The “design-then-reduce” results are exactly analogous.

An important result on the ν -gap metric is that

$$\delta_\nu(\mathbf{G}, \mathbf{G}_p) = \left\| (\mathbf{I} + \mathbf{G}_p \mathbf{G}_p^H)^{-\frac{1}{2}} (\mathbf{G} - \mathbf{G}_p) (\mathbf{I} + \mathbf{G}^H \mathbf{G})^{-\frac{1}{2}} \right\|_\infty \quad (4.15)$$

(repeated from (12.4)), if certain conditions on boundedness and right half-plane poles are satisfied. The first and third terms in the \mathcal{H}_∞ norm are conceptually like matrix versions of a scalar expression $1/\sqrt{1+|g|^2} < 1$; some linear algebra can therefore reveal that

$$\delta_\nu(\mathbf{G}, \mathbf{G}_r) \leq \|\mathbf{G} - \mathbf{G}_r\|_\infty \quad (4.16)$$

(repeated from (12.6)). In light of the model reduction error bounds in Sections 3.4–3.5, this relation is the key to the guaranteed stability or performance of feedback systems with reduced-order controllers. If the model reduction technique has an upper bound on $\|\mathbf{P} - \mathbf{P}_r\|_\infty$ (e.g., as balanced truncation and optimal Hankel norm reduction do, and balanced POD and ERA approximately do), then (4.16) provides a generally loose upper bound on $\delta_\nu(\mathbf{P}, \mathbf{P}_r)$ as well. This,

in turn, can provide a lower bound on $b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}}$ via (3.38a), if we consider \mathbf{P} to be a perturbation of the model \mathbf{P}_r from which we compute the controller $\mathbf{K}_r^{\mathbf{P}}$.

In more technical terms, if \mathbf{P}_r is an approximation of \mathbf{P} that satisfies $\|\mathbf{P} - \mathbf{P}_r\|_{\infty} \leq \beta$ for some $\beta \in [0, 1]$, and $\delta_{\nu}(\mathbf{P}, \mathbf{P}_r) < 1$, then the condition

$$\sin^{-1} b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}} > \sin^{-1} \beta + \sin^{-1} b_d \quad (4.17)$$

is sufficient for satisfying the performance criterion $b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}} > b_d$ for the desired $b_d \in [0, 1]$ (repeated from Theorem 12.4). In practice, β is the model reduction error upper bound; for instance, in the case of balanced truncation, $\beta = 2 \sum_{j=r+1}^n \sigma_j(\mathbf{G})$. In addition, we also remark that because \mathbf{P}_r and $\mathbf{K}_r^{\mathbf{P}}$ are typically small systems (with an order usually less than 100 in feedback flow control), the inequality (4.17) is much easier to compute than $b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}}$ itself. It is in this sense that this theorem can save significant computational effort, since it predicts the minimum performance of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$ without ever needing to test $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$ itself.

A corollary to this result is immediately evident if we take the desired performance level to be $b_d = 0$. In this case, the criterion $b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}} > 0$ becomes one of stability, as per the definition of the normalized coprime stability margin (3.36). This result, repeated from Corollary 12.2, states that if \mathbf{P}_r is an approximation of \mathbf{P} that satisfies $\|\mathbf{P} - \mathbf{P}_r\|_{\infty} \leq \beta$ for some $\beta \in [0, 1]$, and $\delta_{\nu}(\mathbf{P}, \mathbf{P}_r) < 1$, then the condition

$$b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}} > \beta \quad (4.18)$$

is sufficient for the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$. There is a stronger statement on robust stability, however, which follows from Obinata and Anderson (2001, Chapter 3.2). This theorem, repeated from Theorem 12.5, states that if $[\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}]$ is stable, \mathbf{P}_r is an approximation of \mathbf{P} that satisfies $\|\mathbf{P} - \mathbf{P}_r\|_{\infty} \leq \beta$ for some $\beta \in [0, 1]$, and \mathbf{P} and \mathbf{P}_r have the same number of right half-plane poles, then

$$\left\| (\mathbf{I} - \mathbf{K}_r^{\mathbf{P}} \mathbf{P}_r)^{-1} \mathbf{K}_r^{\mathbf{P}} \right\|_{\infty}^{-1} > \beta \quad (4.19)$$

is sufficient for the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$. This is a stronger theorem, because an inspection of the normalized coprime stability margin in (3.36) will reveal that (4.18) directly implies (4.19). We remark briefly that \mathbf{P} and \mathbf{P}_r will typically have the same number of right half-plane poles, since it is common to reduce unstable plants \mathbf{P} by decomposing them into additive stable and unstable parts, and reducing only the stable part before recombining (see Section 4.2).

To demonstrate these results, we use the same linearized complex Ginzburg–Landau equation and parameters as in Section 3.3. In summary, $\dot{q} = \mathcal{L}q$, where $\mathcal{L} = \mu(x) - \nu \partial / \partial x + \gamma \partial^2 / \partial x^2$, $\mu(x) = 0.37 - 5 \cdot 10^{-3} x^2$, $\nu = 2 + 0.4i$, and $\gamma = 1 - i$, and the boundary conditions are that $q \rightarrow 0$ as $x \rightarrow \pm\infty$. A single actuator and sensor are respectively at $x_a = -1$ and $x_s = 1$, which is the optimal placement for LQG control. The plant P is the 100-state discretization of the resulting input–output system. We implement an \mathcal{H}_{∞} loop-shaping controller $K_{wr} = w K_r^{\mathbf{P}}$ (see Section 3.5), where the weight w is given by a negative-feedback LQG controller nearly identical to that of Section 3.3, and $K_r^{\mathbf{P}}$ is the \mathcal{H}_{∞} robust stabilization controller for $P_w = Pw$. Therefore, the controller is a robustified \mathcal{H}_2 optimal controller designed from P_w .

This example uses balanced truncation for model reduction. The analytical results above provide two important tests. First, given a desired model reduction order r , the ν -gap metric between P_w and the order- r approximation P_{wr} can provide a lower bound on $b_{P_{wr}, K_r^{\mathbf{P}}}$ as per (3.38a). Second, the balanced truncation error upper bound is calculable from the Hankel singular values of P_w . Therefore, a comparison between $\beta = 2 \sum_{k=r}^{n-1} \sigma_k(P_w)$ and the robustness of $[P_{wr}, K_r^{\mathbf{P}}]$ can predict r for which $[P_w, K_r^{\mathbf{P}}]$ (that is, $[P, K_{wr}]$) is guaranteed to be stable or to have good performance.

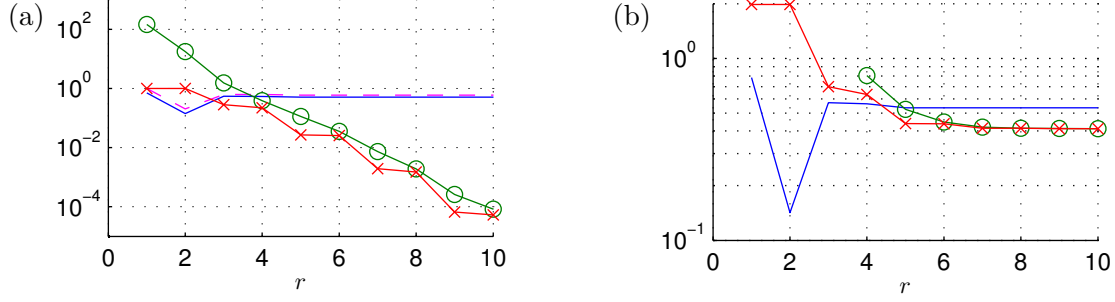


Figure 4.5: (a) Robust stability quantities as the balanced truncation order of the weighted plant P_w varies. Dashed line: $\|(I - K_r^P P_{wr})^{-1} K_r^P\|_\infty^{-1}$, where P_{wr} is the order r balanced truncation of P_w ; solid line: b_{P_{wr}, K_r^P} ; \circ : $\beta = 2 \sum_{k=r}^{n-1} \sigma_k(P_w)$; \times : $\delta_\nu(P_w, P_{wr})$. (b) Robust performance quantities as the balanced truncation order of the weighted plant P_w varies; here, $b_d = 0.4$. Solid line: $\sin^{-1} b_{P_{wr}, K_r^P}$; \circ : $\sin^{-1} \beta + \sin^{-1} b_d$; \times : $\sin^{-1} \delta_\nu(P_w, P_{wr}) + \sin^{-1} b_d$. Repeated from Figures 12.6–12.7.

Figure 4.5 shows these results. In the case of robust stability, (4.18) and (4.19) both predict that $[P, K_{wr}]$ is stable when $r \geq 4$, but the direct application of the ν -gap metric in (3.38a) yields the tighter bound $r \geq 3$. If we actually tested the stability of $[P, K_{wr}]$ for each r , we would find that the closed-loop system is in fact stable if and only if $r \geq 3$. If we required not only stability, but also a minimal performance level of $b_d = 0.4$, then Figure 4.5(b) predicts robust performance for $r \geq 5$, using both (4.17) and (3.38a). This bound is tight, in that $[P, K_{wr}]$ does in fact meet the desired performance level if and only if $r \geq 5$.

This section showed how the normalized coprime stability margin and the ν -gap metric relate to error upper bounds on model reduction. The model reduction error upper bound is generally not the best indicator of robust stability or performance, since it is essentially an open-loop measure. In fact, many model reduction techniques try to fit approximations to original systems over a larger frequency range than is necessary for closed-loop stability or performance. Nonetheless, the model reduction error leads to theorems that guarantee the stability or performance of closed-loop systems containing a full-order plant and a reduced-order controller, when the controller design involves model reduction. These theorems can then provide model reduction orders sufficient for stability or performance, without the need to test the closed-loop system itself. Chapter 12 and Chen and Rowley (2013) describe this theory in greater detail.

Chapter 5

Conclusion

5.1 Summary

In the study of feedback control, fluid mechanics is far from “just another application” of existing theory. Problems in fluid mechanics pose significant and unique challenges that are typically absent in other control applications. This dissertation outlined nonlinearity, high dimensionality, and non-normality as three of the most imposing challenges. These challenges pose difficulties in the broad elements of feedback control—stability theory, control design, and modeling—in different ways.

In the study of fluid stability theory, powerful results are available using a purely linear analysis. The eigenvalues of the linearized flow equations reveal the growth rates of direct eigenmodes, indicating the stability of the nonlinear flow systems in all but a few singular cases. The adjoint eigenmodes, on the other hand, reveal the receptivity of the linearized dynamics to external disturbances and the choice of initial conditions. Individually, these eigenmodes tend not to be particularly revealing for non-normal flows. In these flows, eigenmodes are nearly parallel, so they do not readily predict transient growths that occur even in globally stable flows. On the other hand, the theory of sensitivity to spatially localized dynamical perturbations can often yield relevant results in non-normal flows. Regions of sensitivity—which are related to regions of absolute instability and to wavemakers—often indicate the parts of the physical domain where perturbation growth can be large, and the dynamics respond in nontrivial ways to modifications. This dissertation did not discuss nonlinear stability theory, but Chomaz (2005) discusses nonlinear extensions of both local and global linear stability.

Computational methods for stability analysis are well-established at this point, even for high-dimensional systems. It is generally possible to solve for steady states of high-dimensional dynamical systems using Newton’s method, where an inner iteration such as GMRES or the biconjugate gradient stabilized method solves the linear Jacobian–vector equation that reveals the Newton step. It is further possible to stabilize Newton’s method using the Armijo line rule, and to speed up the convergence of linear solvers using preconditioning. After Newton’s method has revealed the steady-state solutions of a dynamical system, the Arnoldi method iterates toward the most unstable eigenvalues and eigenmodes of the linearization.

In this dissertation, we apply linear stability theory to the flow through a T-shaped pipe bifurcation, at Reynolds numbers in the mid-hundreds. The iterative techniques reveal the stability, receptivity, and sensitivity regions of this highly non-normal flow. In particular, the sensitivity regions coincide closely with recirculating regions undergoing vortex breakdown. This coincidence likely occurs because the recirculation causes certain perturbations in the flow to feed back positively, leading to self-amplification.

These stability results have important consequences for fluid flow control, at the very least because successful control must stabilize unstable modes and transient growths. This dissertation does not discuss nonlinear control, which tends to be a somewhat fragmented theory. In LTI control, however, the theory of \mathcal{H}_2 optimality yields powerful controllers that many have widely used for a few decades. The LQG—a particular \mathcal{H}_2 optimal controller—optimally minimizes a cost function based on the size of state variables and actuator inputs, in the presence of system disturbances and sensor noise.

Although optimal controller design is nearly ubiquitous in LTI control theory, comparatively few have previously studied the optimization of actuator and sensor placements in partial differential equation systems. This dissertation presents an algorithm that computes the gradient of a closed-loop \mathcal{H}_2 norm—e.g., from state disturbances and sensor noise to the state and input size—with respect to actuator and sensor locations. This gradient allows a control designer to implement a gradient-based local minimization that simultaneously optimizes over all LTI controllers, all actuator placements, and all sensor placements. In addition, the computation of the gradient is faster than the computation of the \mathcal{H}_2 norm itself. The optimization of actuator and sensor placements in the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire models of fluid flow reveal that eigenmode analyses can successfully predict optimal placements when the system dynamics are close to normal. When the flow contains significant non-normality, however, eigenmode and Gramian analyses fails to predict the optimal placement. The sensitivity region may be a more viable heuristic for optimal placement in certain systems. Ongoing research is attempting to establish firmer conclusions.

Next, this dissertation discussed stability margins and closed-loop system metrics that arise from the normalized coprime factorization. The normalized coprime stability margin also motivates \mathcal{H}_∞ robust stabilization and loop-shaping, which are easy-to-use methods of designing optimally robust controllers. This robustness theory, based heavily on the ν -gap metric, is particularly powerful in its ability to predict closed-loop stability and performance in the presence of plant or controller perturbations. These perturbations may include the effects of reduced-order modeling, which is typically required in high-dimensional control. In this interpretation, the comparison among the normalized coprime stability margin, the ν -gap metric, and upper bounds on model reduction error is a particularly powerful one. This dissertation presents theorems that use these quantities to determine sufficient conditions for the *a priori* stability or performance of full-order plants in feedback with reduced-order controllers. Not only does this save time and effort by eliminating the need to analyze such feedback loops in detail, but it also predicts the model reduction orders that are sufficient for the given control objectives.

The study of reduced-order modeling for specifically high-dimensional systems, however, tends not to be straightforward. Reduced-order models based on proper orthogonal decomposition are not easy to design well, and balanced POD and ERA assume that systems are linear. Furthermore, there does not exist a unifying theory of high-dimensional model reduction. The advent of dynamic mode decomposition takes a significant step in this direction by “curve fitting” time series of data to linear trajectories comprised of eigenvalues and eigenmodes, even if the data arise from nonlinear dynamics. The Koopman operator decomposition—which DMD approximates—accurately models both linear and periodic trajectories, and recent evidence suggests that DMD is general in its applicability to nonlinear systems. This thesis proves that the DMD of a given data set is unique under certain conditions. It also analyzes methods for constructing DMD modes with homogeneous boundary conditions, since Galerkin modeling typically requires such modes. This investigation reveals that subtracting the data mean from a time series actually reduces DMD to the DFT, which is generally undesirable. On the other hand, the subtraction of a steady-state solution appears to

preserve the DMD. Finally, the “optimized” variant of DMD better tailors the decomposition to a desired number of modes.

5.2 Future outlook

The system-based feedback control of fluid flows is, in many ways, a young field. On one hand, many decades have passed since the earliest applications of feedback control on fluid flows. Yet, the research community is still making large strides toward the understanding of high-dimensional and possibly nonlinear stability, control, and modeling of typically non-normal flows. The constant increase in computational power plays a critical role in this research progress, but it is not the sole driving factor, either. (If it were, we researchers might be better off taking an extended vacation while waiting for computers to become more powerful!) Pseudospectra (Trefethen et al., 1993), global sensitivity analyses (Giannetti and Luchini, 2007), and faithful modeling techniques for nonlinear systems (Schmid, 2010) are just a few examples of recent advances—not directly stemming from increased computational power—that may make high-performance feedback flow control a closer reality.

In the field of feedback flow control, many consider the laminarization of turbulent flow to be the “trophy” unsolved problem. Kim (2003) reviews significant progress toward this goal. More recently, Sharma et al. (2011) propose a control method that provably reduces the magnitude of flow perturbations monotonically. Yet, many of these computational studies assume actuation and sensing over a large domain (such as an entire surface or volume), which may not be physically realistic. The laminarization of turbulent flow in an experimental setting, then, is in some sense the ultimate goal. The feedback flow control community likely still has major strides to take before it could accomplish such a task.

Robustness, for instance, is a critically important property that feedback flow control designers must seek to increase. So long as computers are not powerful enough to design controllers directly from fluid models, there will exist approximations in light of which the feedback system must remain stable. The extreme non-normality of high Reynolds number flows can also amplify tiny perturbations by many magnitudes, potentially leading to nonlinear effects that—again—the feedback system must tolerate. Therefore, flow control research targeting non-normality, nonlinear terms in partial differential equations, and robustness must continue.

The placement of actuators and sensors in a fluid flow is also far from solved. At this point, the literature only contains very rough heuristics and simplified studies for optimal placement. Oftentimes, actuator and sensor placement can be just as important—if not more so—than controller design in fluid flows (see, for instance, Belson et al., 2013). A large amount of research in placement for optimal robustness and for high-dimensional fluid systems must take place before there could be any hope of controlling high Reynolds number flows in an experimental setting via feedback.

The design of reduced-order models for fluid flows is also tremendously important for successful feedback flow control. The model reduction techniques that are common in flow control design have significant shortcomings such as general applicability, convergence, model accuracy, and the ability to model nonlinearity. Therefore, there is still a need to develop modeling methods that can yield effective controllers for fluid flows. This can be a particularly difficult task, since there does not currently exist an encompassing or general theory for high-dimensional reduced-order model design.

The examples in this section do not at all comprise a comprehensive list of major unsolved problems in feedback flow control. This field is a highly complex one where every theoretical or applied study presents just another facet of the larger problem. Nevertheless, feedback flow control

has heralded complete transformations of many engineering applications, and the day may yet come when applications in fluid mechanics experience such a shift.

Part II

Papers

Chapter 6

Overview

Part II of this dissertation contains articles that are either in the literature, or likely will be in the near future. Only minor modifications (e.g., related to formatting) are present between the published articles and the chapters here. The papers are organized into chapters as follows.

- Chapter 7 presents a global stability analysis of a pipe flow through a T-shaped bifurcation at mid-hundred Reynolds numbers, which is a commonly-occurring flow in natural and man-made systems. It focuses on the relation between vortex breakdown and dynamical sensitivity to spatially localized perturbations.
- Chapter 8 discusses the placement of actuators and sensors in feedback flow control systems. It presents an algorithm that computes the gradient of an \mathcal{H}_2 optimal control performance with respect to actuator and sensor locations, allowing for an efficient gradient-based optimization of the placement. The paper applies the algorithm to the control of the linearized Ginzburg–Landau system, and it evaluates the efficacy of placement techniques used in the literature.
- Chapter 9 extends Chapter 8 by discussing an adjoint-based improvement to the original gradient algorithm, as Colburn et al. (2011) proposed. It also applies the algorithm to the control of the Orr–Sommerfeld and Squire models of shear flow perturbations.
- Chapter 10 extends Chapters 8 and 9 by discussing heuristics and qualities that characterize optimal and other effective actuator and sensor placements in feedback flow control.
- Chapter 11 analyzes previously unknown properties of dynamic mode decomposition (DMD), a novel method for reduced-order flow representation. This paper proves the uniqueness of the decomposition and addresses boundary conditions in DMD-based reduced-order models. It also proposes an “optimized” algorithm that produces less spurious decompositions and gives the user control over the number of output modes.
- Chapter 12 investigates the stability and performance of high-dimensional (e.g., fluid) systems in closed-loop with reduced-order controllers. Theorems based on the normalized coprime factorization and the ν -gap metric provide sufficient conditions for stability and performance. These conditions, which are applicable with typically known upper bounds on model reduction error, can determine adequate model reduction orders for reduced-order control.

6.1 Author contributions

In the following chapters, I developed most of the in-house software and most directly executed the theory development, analyses, and writing. Clarence W. Rowley advised me on nearly all aspects of the research, and revised and edited these papers. I outline specific contributions from co-authors below; most unlisted contributions are from myself.

- In Chapter 7, Howard A. Stone initially proposed the investigation of the pipe T-junction, and I proposed the stability and bifurcation analysis of the flow. Rowley advised me on the computational, stability, and bifurcation analyses; Stone advised me on the flow physics, and in particular, the vortex breakdown phenomenon. Stone also planned the layout of the paper, and Rowley wrote an early version of the Newton’s method software that I further developed.
- In Chapter 8, I initially proposed the investigation of optimal actuator and sensor placement. Rowley proposed the use of closed-loop \mathcal{H}_2 norms as a control performance measure, as well as the computation of \mathcal{H}_2 norm gradients via perturbations. Rowley also proposed the investigation of Gramians and wavemakers as heuristics for actuator and sensor placement, and he proposed the general experimentation of Ginzburg–Landau control.
- In Chapter 9, I proposed the use of the Orr–Sommerfeld/Squire equations as a transverse-direction test bed.
- In Chapter 10, Rowley proposed the use of optimal growth theory for characterizing effective actuator and sensor placements.
- In Chapter 11, Rowley and Jonathan H. Tu proposed original concepts that would evolve into the optimized DMD algorithm. Rowley also developed the portions of the optimized DMD algorithm involving least-squares calculations, and I developed the rest of the algorithm. Tu performed a portion of the mean subtraction analyses. Lastly, Rowley proposed the investigation of DMD as a means of bridging linear modes and limit cycle modes.
- In Chapter 12, Rowley proposed the use of the ν -gap metric as a means of measuring distance from the perspective of closed-loop systems. Rowley also proposed the use of this metric to investigate the stability of closed-loop systems with reduced-order controllers.

Chapter 7

Vortex breakdown in a T-junction: recirculation and sensitivity

Kevin K. Chen, Clarence W. Rowley, and Howard A. Stone

Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

To be submitted to *Physics of Fluids*.

Pipe bifurcations are a ubiquitous configuration in fluid flows. In this study, we explore the flow through T-shaped bifurcations with square cross-sections, at Re varying between approximately 320 and 650. When $Re \geq 320$, vortex breakdown occurs in four locations in the junction. The associated recirculation regions are particularly sensitive to spatially localized perturbations of the linearized Navier–Stokes operator. We explain this phenomenon using a physical interpretation of an inviscid Lagrangian short-wavelength approximation.

7.1 Introduction

The flow through pipe bifurcations is a common occurrence in both natural and man-made systems, including blood vessels, industrial pipe networks, and microfluidic channels. T-shaped pipe bifurcations, such as the one shown in Figure 7.1(a), are one of the most common and universal flow configurations across both natural and man-made systems. In the former category, familiar examples include the pulmonary and basilar arteries in the human body. The basilar artery is a particularly intriguing case, since aneurysms may occur at artery junctions when the artery walls are weak (Ropper and Samuels, 2009). Among man-made systems, examples include industrial pipe networks, microfluidic channels, and even surgeries such as the Norwood procedure for hypoplastic left heart syndrome patients (Norwood et al., 1980).

A distinguishing feature of the geometry is that the T-junction must curve flow streamlines. Historically, many studies have focused on numerical solutions of a simpler curved-streamline flow: the fully-developed Dean flow through a slowly bending pipe (see Winters (1987) and the references within). Although the Dean flow contains a rich set of solutions and bifurcations, it only roughly captures the physics of the T-junction flow. For instance, it correctly predicts the existence of multi-vortex solutions in the T-junction. Nonetheless, this essentially two-dimensional analysis cannot predict complex three-dimensional behavior such as vortex breakdown, which we document here.

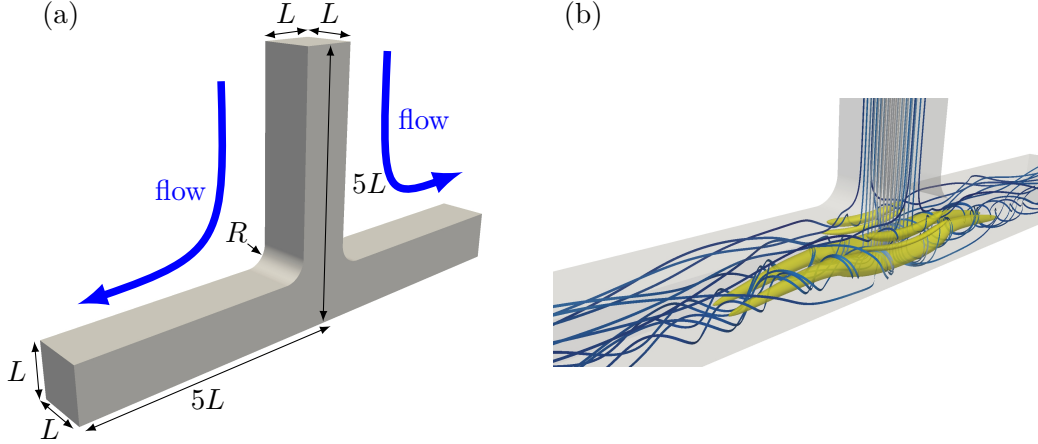


Figure 7.1: (a) The T-junction geometry. (b) Velocity streamlines (white: high speed; blue: medium speed; black: low speed) at $Re = 560$, with $R = 0.4L$. Given the Frobenius norm $\|\cdot\|_F$, we also define $Q = (\|\nabla \mathbf{u} - (\nabla \mathbf{u})^T\|_F^2 - \|\nabla \mathbf{u} + (\nabla \mathbf{u})^T\|_F^2)/8$ and depict a single level set of Q (yellow; see Chakraborty et al. (2005) for a discussion of vortex identification).

Despite the universality of the impacting T-junction flow and the closely related L-bend (Kockmann et al., 2005; Haller et al., 2009), a large gap in the physical understanding of these flows still persists. Many previous studies (Senn and Poulikakos, 2004; Kockmann et al., 2005; Haller et al., 2009; Doorly and Sherwin, 2009) offer detailed images and basic physical insight, particularly as they relate to microchannel flows. Yet, few provide quantitative characterizations tied to the physical behavior in these flows.

A recent insight in the study of impacting T-junction flows, however, is the discovery of recirculation regions arising from vortex breakdown (Vigolo et al., 2014). The flow traps bubbles, which are suspended in the liquid, in vortical structures near stagnation regions in the “T” when the Reynolds number is above approximately 350. The trapping occurs because of the vortex breakdown, as well as large pressure gradients and drag forces in the junction.

This bubble trapping phenomenon merits a further investigation of the single-phase T-junction flow. Of particular interest is a recent observation that simple closed particle trajectories may give rise to an increased sensitivity in the linearized dynamics to spatially localized perturbations (Giannetti et al., 2010). According to an inviscid short-wavelength theory (Bayly, 1988; Lifschitz and Hameiri, 1991), perturbations advect with the velocity field in the limit that the perturbation’s wavelength and the fluid viscosity approach zero. This local theory, based on the Wentzel–Kramers–Brillouin–Jeffreys approximation, has been used to study instabilities in two-dimensional flows (see the references contained in the Appendix of Giannetti et al. (2010)). The key observation in Giannetti et al. (2010) is that local perturbations may self-amplify after they have traversed a closed orbit, depending on the phase of the returned perturbation. In turn, the amplification gives rise to a notable increase in dynamical sensitivity to localized perturbations near closed orbits.

In this letter, we investigate the connection between recirculation and dynamical sensitivity in the T-junction flow, at Reynolds numbers where the steady-state solution loses stability, but the flow remains laminar. We employ a global linear analysis (Giannetti and Luchini, 2007) and show that the recirculation and sensitivity regions essentially coincide. Also, we briefly comment on the lack of appreciable effects from the junction corners’ radius of curvature and the associated flow separation on the dynamical sensitivity to localized perturbations.

7.2 Global stability theory

To define the sensitivity, we use a linear analysis framework. Given the average inlet fluid speed U , the pipe width L (see Figure 7.1(a)), and the fluid kinematic viscosity ν , we assume that there exists some steady-state velocity field \mathbf{u}_0 for the given Reynolds number $Re = UL/\nu$. That is, given a velocity field \mathbf{u} , a pressure field p , the nonlinear Navier–Stokes equation

$$\dot{\mathbf{u}} = \mathcal{N}\mathbf{u} = -\mathbf{u} \cdot \nabla \mathbf{u} - \nabla p + Re^{-1} \nabla^2 \mathbf{u}, \quad (7.1)$$

and the continuity equation $\nabla \cdot \mathbf{u} = 0$, we require $\mathcal{N}\mathbf{u}_0 = \mathbf{0}$. We then linearize the operator \mathcal{N} about \mathbf{u}_0 to derive the linearized Navier–Stokes equation for the velocity and pressure perturbations \mathbf{u}' and p' . This equation is

$$\dot{\mathbf{u}}' = \mathcal{L}\mathbf{u}' = -\mathbf{u}' \cdot \nabla \mathbf{u}_0 - \mathbf{u}_0 \cdot \nabla \mathbf{u}' - \nabla p' + Re^{-1} \nabla^2 \mathbf{u}', \quad (7.2)$$

subject to $\nabla \cdot \mathbf{u}' = 0$; these together define \mathcal{L} . Given the complex conjugate operator $\overline{(\cdot)}$ and some control volume Ω , we define the inner product

$$\langle \mathbf{u}_1, \mathbf{u}_2 \rangle = \int_{\Omega} \bar{\mathbf{u}}_2 \cdot \mathbf{u}_1 dV, \quad (7.3)$$

from which we can define the adjoint \mathcal{L}^* of the operator \mathcal{L} . Denoting the adjoint velocity and pressure perturbations by $\hat{\mathbf{u}}'$ and \hat{p}' , subject to $\nabla \cdot \hat{\mathbf{u}}' = 0$, the adjoint equation is

$$\dot{\hat{\mathbf{u}}} = \mathcal{L}^* \hat{\mathbf{u}}' = -(\nabla \mathbf{u}_0) \cdot \hat{\mathbf{u}}' + \mathbf{u}_0 \cdot \nabla \hat{\mathbf{u}}' - \nabla \hat{p}' + Re^{-1} \nabla^2 \hat{\mathbf{u}}'. \quad (7.4)$$

We define the *sensitivity* of the linearized dynamics to spatially localized feedback using a global mode analysis (Giannetti and Luchini, 2007). If the direct and adjoint eigendecompositions are $\mathcal{L}\phi_j = \lambda_j \phi_j$ and $\mathcal{L}^*\psi_j = \bar{\lambda}_j \psi_j$, respectively, then an infinitesimal perturbation of the former yields

$$d\mathcal{L}\phi_j + \mathcal{L}d\phi_j = d\lambda_j \phi_j + \lambda_j d\phi_j. \quad (7.5)$$

Following the definition in (7.3), the inner product of (7.5) with ψ_j leads to a cancellation of the second terms on the left- and right-hand sides, i.e.,

$$\langle d\mathcal{L}\phi_j, \psi_j \rangle = \langle d\lambda_j \phi_j, \psi_j \rangle. \quad (7.6)$$

Let us further restrict $d\mathcal{L}$ to be a spatially localized perturbation. Given the Dirac delta function δ , some perturbation location $\boldsymbol{\xi}$, and some perturbation size ds , we have $d\mathcal{L} = \delta(\mathbf{x} - \boldsymbol{\xi}) ds$. This perturbation could be, for instance, the application of feedback control with a collocated actuator and sensor (Giannetti and Luchini, 2007). From (7.6), it follows that

$$\bar{\psi}_j(\boldsymbol{\xi}) \cdot \phi_j(\boldsymbol{\xi}) ds = d\lambda_j \langle \phi_j, \psi_j \rangle. \quad (7.7)$$

Therefore, we may define a *sensitivity function*

$$\zeta_j(\mathbf{x}) = \frac{d\lambda_j}{ds} = \frac{\bar{\psi}_j(\mathbf{x}) \cdot \phi_j(\mathbf{x})}{\langle \phi_j, \psi_j \rangle}, \quad (7.8)$$

which dictates the change in the j th eigenvalue as a result of a perturbation in \mathcal{L} localized at \mathbf{x} with strength ds . The physical locations where $|\zeta_j(\mathbf{x})|$ is large are the regions where the dynamics are especially sensitive to these spatially localized perturbations.

7.3 Computational methods

To approximate the steady-state solution \mathbf{u}_0 for a given Re , we employ a Newton–Armijo iteration (Armijo, 1966; Kelley, 1995) on the operator \mathcal{N} (7.1) over the vector space of spatially-discretized velocity fields \mathbf{u} . We compute the Newton solver’s initial condition by running either a time-resolved flow solver for a long time, or a zeroth- or first-order extrapolation from solutions at other Reynolds numbers. In the solver itself, we use GMRES (Saad and Schultz, 1986; Kelley, 1995) to solve the Jacobian–vector equation $\partial\mathcal{N}/\partial\mathbf{u}|_{\mathbf{u}_k} \cdot \mathbf{h} = -\mathbf{u}_k$ for the Newton step \mathbf{h} from the k th iterate \mathbf{u}_k . In this approach, however, the large dimension prevents us from computing the Jacobian–vector product $\partial\mathcal{N}/\partial\mathbf{u}|_{\mathbf{u}_k} \cdot \mathbf{h}$ directly. Instead, we use a “time-stepping” approach (Mamun and Tuckerman, 1995), employing the finite difference $\partial\mathcal{N}/\partial\mathbf{u}|_{\mathbf{u}_k} \cdot \mathbf{h} = (\mathcal{N}(\mathbf{u}_k + \epsilon\mathbf{h}) - \mathcal{N}\mathbf{u}_k)/\epsilon + \mathcal{O}(\epsilon)$, with $\epsilon = 10^{-3}$. We terminate the Newton–Armijo search when \mathbf{u}_k satisfies $\|\mathcal{N}\mathbf{u}_k\|_2/V < 10^{-6}$, where $\|\mathbf{u}\|_2 = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}$ (see (7.3)) and V is the volume of the domain Ω .

Once we have obtained a satisfactory approximation for \mathbf{u}_0 , we employ an Arnoldi iteration (Trefethen and Bau, 1997) with discrete-time variants of \mathcal{L} (7.2) and \mathcal{L}^* (7.4) to approximate their leading eigenvalues and eigenmodes. In our computation, 100 Arnoldi iterations are more than sufficient for good convergence of the relevant modes. Also, we remark that when we compute any of the Navier–Stokes operators on a velocity field, we first solve the divergence of the corresponding equation (7.1, 7.2, 7.4) for the pressure field.

We implement the Navier–Stokes operators (7.1, 7.2, 7.4) using software based on the OpenFOAM suite’s icoFoam solver. This solver, which uses a finite-volume implicit pressure correction (Ferziger and Perić, 2002) on an unstructured mesh, is more primitive than other available solvers, but is sufficient for the Newton–Armijo, GMRES, and Arnoldi iterations. In our software, we interface these iterators with the Navier–Stokes operators. With 64 Intel Xeon E5-2670 processor cores on an FDR Infiniband network, the computation of a steady-state solution and its direct and adjoint eigendecompositions collectively requires about one day.

The T-junction geometry we employ has a square cross-section to match accompanying experiments (Vigolo et al., 2013, 2014), but our results may extend to other cross-sections as well. The cross-section has dimensions $L \times L$, and the inlet and two outlets are $5L$ long; see Figure 7.1(a). A shortening of the inlet to $3L$ and an extension of the outlets to $10L$ individually show no visible differences in the flow behavior and the eigenvalues of \mathcal{L} and \mathcal{L}^* . In our investigation, we focus on a junction corner radius of curvature $R = 0.4L$ (see Figure 7.1(a)) to match our experimental work, but we also examine the flow at $R = 0$ (i.e., a square corner) and $R = L$. The geometries have 5,459,520 finite-volume cells at $R = 0$, 6,096,840 cells at $R = 0.4L$, and 7,435,776 cells at $R = L$. The mesh is finer near the junction and the walls, where gradients are large, and coarser near the inlet and outlets, where gradients in the inflow and outflow directions are small. Finer meshes do produce small changes in velocity and pressure profiles, as well in the eigenvalues of \mathcal{L} and \mathcal{L}^* . The qualitative nature of our solutions, however, does not change.

In the nonlinear Navier–Stokes operator \mathcal{N} , we impose the fully-developed velocity profile for a square channel (White, 2005) at the inlet, as well as $\mathbf{n} \cdot \nabla p = 0$, with \mathbf{n} the normal vector. Although this pressure condition is analytically incorrect, the solver corrects the pressure gradient within a few downstream cells. We further impose $\mathbf{u} = \mathbf{0}$ and $\mathbf{n} \cdot \nabla p = 0$ at the walls, and $\mathbf{n} \cdot \nabla \mathbf{u} = \mathbf{0}$ and $p = 0$ at the outlets. The linearized operator \mathcal{L} has the same conditions on \mathbf{u}' and p' , except that




Figure 7.2: The sensitivity function’s L_2 norm $\|\zeta\|_2 = (\int_{\Omega} |\zeta|^2 dV)^{1/2}$ (see (7.8)), as the center (i.e., neutrally stable) mode at the first supercritical Hopf bifurcation numerically continues to lower and higher Re . Blue dashed line: $R = 0$; green solid line: $R = 0.4L$; red dotted line: $R = L$. Supercritical Hopf bifurcations are shown as squares.

$\mathbf{u}' = \mathbf{0}$ at the inlet. Finally, the adjoint linearized operator \mathcal{L}^* has the same boundary conditions as \mathcal{L} on $\hat{\mathbf{u}}'$ and \hat{p}' , except that $(\mathbf{u}_0 \cdot \mathbf{n})\hat{\mathbf{u}}' + Re^{-1}\mathbf{n} \cdot \nabla \hat{\mathbf{u}}' = \mathbf{0}$ at the outlets; this condition arises from the integration by parts in the derivation of \mathcal{L}^* .

7.4 Results

At $Re = \mathcal{O}(10^2)$, the flow is smooth in the inlet, and features a large counter-rotating vortex pair in the junction, which extends toward the outlets. Figure 7.1(b) depicts these flow features with streamlines and vortex visualization. Additional counter-rotating vortex pairs are also present in the junction, such as the secondary and tertiary pair visible in Figure 7.1(b).

We primarily consider the first instability of the T-junction flow. The computed steady-state solutions remain stable until the first supercritical Hopf bifurcation occurs at $Re = 587$ for $R = 0$, $Re = 556$ for $R = 0.4L$, and $Re = 552$ for $R = L$; see Figure 7.2. In this laminar flow regime, an increase in Re generally increases the sensitivity of the linearized dynamics to spatially localized perturbations (Giannetti and Luchini, 2007). The increase in Reynolds number triggers an increase in non-normality, thereby enlarging the ϵ -pseudospectrum of \mathcal{L} where ϵ -scale perturbations in \mathcal{L} may shift its eigenvalues by considerably more than ϵ (Chomaz, 2005). We will see that the regions of large sensitivity are intrinsically connected to recirculation regions in the flow.

At $Re \geq 320$, when the steady-state solution of interest is still stable, four of these recirculation regions appear in the junction, with one in each outlet–depth quadrant. This behavior, first reported for a bubbly T-junction flow (Vigolo et al., 2014) and shown in Figure 7.3 near the first Hopf bifurcation, arises as a result of vortex breakdown abruptly ending the tightly swirling vortex. The recirculation gives rise to stagnation points in the interior of the flow. Four of these stagnation points are found at the vortex breakdown points. On the center plane between the two outlets, additional stagnation points also appear in the vortex cores. Streamlines running against the primary direction of the outlet flow bring fluid from the breakdown point back to the center plane.

To a rough approximation, the recirculation regions resemble a Hill’s spherical vortex (Hill, 1894) with very high swirl, stretched in the outlet directions and further spatially distorted. As with the endpoints of the Hill’s spherical vortex, the vortex breakdown point is a saddle-focus fixed point with a two-dimensional stable manifold comprising a “shell” of the recirculation region, and a one-dimensional unstable manifold carrying fluid either toward the outlet, or back toward the center plane. The stagnation point on the center plane has the opposite behavior: a two-dimensional unstable manifold comprising the shell, and a one-dimensional stable manifold carrying fluid from

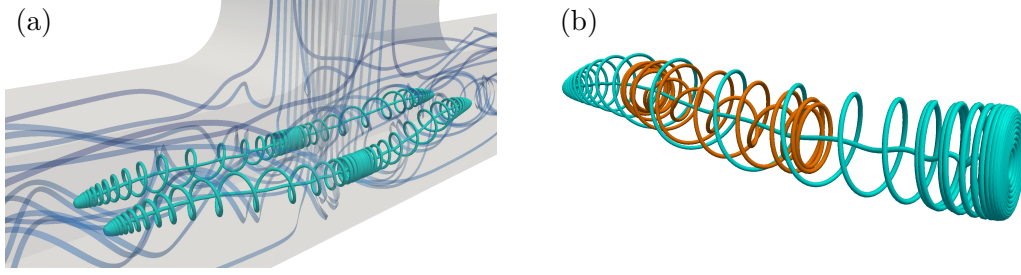


Figure 7.3: The vortex breakdown at $Re = 560$ and $R = 0.4L$. (a) The four recirculation regions (cyan), with streamlines as in Figure 7.1(b). (b) An outer (cyan) and inner (orange) streamline from a single recirculation region. The cyan streamline's left end is the vortex breakdown point, and its right side is flush with the T-junction's center plane.




Figure 7.4: A level set of the sensitivity magnitude $|\zeta|$ (magenta), with recirculation streamlines (cyan). Here, $Re = 560$ and $R = 0.4L$. (a) The first instability. (b) The second instability.

the breakdown points. We do not investigate whether the breakdown points' manifolds coincide or transversely intersect with those of the center plane's stagnation points. Numerically, they appear to coincide, but a thorough Lagrangian analysis would require a very careful mesh refinement (Sotiropoulos et al., 2001) that we do not attempt here. Inside the recirculation regions, the flow streamlines resemble the interior streamlines of the Hill's spherical vortex with swirl—they swirl downstream over an outer surface and swirl upstream over an inner surface. Figure 7.3(b) provides a detailed visualization of a recirculation region, and Sotiropoulos et al. (2001) provide an in-depth treatise of Lagrangian dynamics in a vortex breakdown scenario.

The qualitative nature of the T-junction flow's recirculation streamlines has significant implications in the flow dynamics' sensitivity, since there is—at least approximately—a dense family of closed streamline orbits. In our research, we do not attempt to apply the inviscid short-wavelength theory in its mathematically rigorous sense. At present, researchers have observed the connection between closed orbits and dynamical sensitivity in simple closed orbits, such as in stationary (Giannetti and Luchini, 2007) and periodic (Giannetti et al., 2010) flows behind a cylinder. The application of the technical theory has been restricted to such simple orbits, including driven cavity flows (Luchini et al., 2013). In contrast, the orbits in Figure 7.3 are likely too complex for a rigorous application of the theory. We argue, however, that the underlying idea is still physically applicable. Figure 7.4 demonstrates that in the T-junction flow, the recirculation and sensitivity regions do in fact coincide closely. This figure shows the sensitivity regions of the two least stable modes at the first supercritical Hopf bifurcation (i.e., $Re = 560$) with $R = 0.4$. An examination of the next several modes' sensitivity regions shows very little qualitative change. The connection between recirculation and sensitivity is in close agreement with the Appendix of Giannetti et al. (2010), but



Figure 7.5: (a) As Figure 7.1(b) and 7.4, but with junction radii $R = 0$ and $Re = 590$ (left), and $R = L$ and $Re = 550$ (right). (b) As (a, left), with additional streamlines (cyan) depicting the flow separation at the left square corner.

the recirculation in the T-junction flow is fully three-dimensional and far more complicated. To the best of our knowledge, the T-junction flow is the most complex flow for which anyone has reported a close relation between recirculation and sensitivity.

A comparison of Figures 7.1(b) and 7.4(a) reveals that the dynamical sensitivity is also nontrivial in the secondary and tertiary counter-rotating vortex pairs that appear above the large primary pair. Although these vortex pairs neither undergo vortex breakdown nor constitute recirculation regions, the swirl in these pairs is very large compared to the axial velocity. We posit that this swirl may contribute to the locally increased sensitivity in a similar way to the recirculation regions. Also, we remark that the sensitivity of the second instability (Figure 7.4(b)), as well as the third and fourth instabilities (not shown), is more pronounced away from the center plane between the two outlets. The sensitivity in these modes is also nontrivial in lobes immediately outside the recirculation regions.

At this point, it is clear that the recirculation regions and vortices in the junction and outlet pipes are key aspects of the impacting T-junction flow. Nevertheless, what is the role of the separated flow at the junction corners? Figure 7.5(a) shows the vortical flow and the first instability's sensitivity at the first supercritical Hopf bifurcation, for the square corner $R = 0$ and the large radius $R = L$. In both of these cases, the qualitative nature of the sensitivity regions is nearly identical to that of the medium radius $R = 0.4L$ (Figure 7.4).

Additional insight may be possible if we visualize the velocity streamlines comprising the separation regions at the square corners for $R = 0$. In Figure 7.5(b), a particular region of flow starts near the walls of the inlet, navigates around the square corner of the junction, and encounters an adverse pressure gradient in the outlet pipes. This pressure gradient causes these streamlines to turn backward, forming a vertical vortex pair that flushes fluid downward between the two large outlet-oriented vortices in the junction. The critical observation in this case is that the separation region is completely disjoint from the sensitivity region. Similar flow features also occur for $R = 0.4L$ and $R = L$ (not shown), where the separation region is much smaller and bears some resemblance to the flow in a diverging nozzle.

Since the separation and sensitivity regions are distinct, the flow separation at the corner is “inert” in the sensitivity analysis. In other words, a linearized dynamical perturbation in the separation region will not have a large effect on the stability of the steady-state flow. Physically, this lack of an appreciable effect may be because the separation region is structurally stable, in the sense that dynamical perturbations there may not significantly influence the flow structures. In comparison, the recirculation regions in the junction may be less structurally stable, enabling small perturbations in the dynamics there to have large effects on global stability.

7.5 Conclusions

In this research, we employed a global linear sensitivity analysis and confirmed that recirculation regions in the T-junction flow exhibit a high sensitivity to localized dynamical perturbations. According to the inviscid short-wavelength theory, the recirculation regions—which resemble a dense family of closed orbits—approximately advect short-wave perturbations that can positively feed back, causing a self-amplification mechanism. Secondary and tertiary vortex pairs also exhibit high sensitivity, possibly because of high swirl. The junction corners’ radius of curvature does not exhibit a significant effect on the flow sensitivity, however, and the flow separation at the junction corners is physically disjoint from the sensitivity regions. This feature suggests that the flow separation is structurally stable. In a future paper, we plan to discuss further aspects of the global stability analysis, including the regions of stability and receptivity given by the linearized Navier–Stokes operators’ direct and adjoint eigenmodes. These stability analyses will yield new insight into the behavior of this very common, though complex and, at times, counterintuitive flow.

K. K. C. was funded by the National Science Foundation through the Graduate Research Fellowship Program. We thank Stefan Radl and Daniele Vigolo for their assistance, and François Gallaire and Paolo Luchini for helpful discussions and guidance regarding vortex breakdown and sensitivity. We also thank Philip Holmes, John O. Dabiri, Matthew O. Williams, and Jonathan H. Tu for their insight on the behavior of the recirculation streamlines. We performed the computation on clusters at Princeton University’s TIGRESS High Performance Computing Center.

Chapter 8

\mathcal{H}_2 optimal actuator and sensor placement in the Ginzburg–Landau system

Kevin K. Chen and Clarence W. Rowley

Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

Appears in *Journal of Fluid Mechanics* 681:241–260, August 2011.

Digital Object Identifier [doi:10.1017/jfm.2011.195](https://doi.org/10.1017/jfm.2011.195)

The linearised complex Ginzburg–Landau equation is a model for the evolution of small fluid perturbations, such as in a bluff body wake. By implementing actuators and sensors and designing an \mathcal{H}_2 optimal controller, we control a supercritical, infinite-domain formulation of this system. We seek the optimal actuator and sensor placement that minimises the \mathcal{H}_2 norm of the controlled system, from flow disturbances and sensor noise to a cost on the perturbation and input magnitudes. We formulate the gradient of the \mathcal{H}_2 squared norm with respect to actuator and sensor placements, and iterate toward the optimal placement. When stochastic flow disturbances are present everywhere in the spatial domain, it is optimal to place the actuator just upstream of the origin and the sensor just downstream. With pairs of actuators and sensors, it is optimal to place each actuator slightly upstream of each corresponding sensor, and scatter the pairs throughout the spatial domain. When disturbances are only introduced upstream, the optimal placement shifts upstream as well. Global mode and Gramian analyses fail to predict the optimal placement; they produce \mathcal{H}_2 norms about five times higher than at the true optimum. The wavemaker region is a better guess for the optimal placement.

8.1 Introduction

Stability theory has been an active field of fluid mechanics research because of its prevalence in physical phenomena. A common characteristic of stability theory is the simplification of the full Navier–Stokes equations to an amplitude equation of reduced dimension. This is often done out of practicality, since reduced-dimensional models are more tractable and easier to analyse. The Ginzburg–Landau equation is one such model that is often used to study fluid instabilities in spatially developing flows (Huerre and Monkewitz, 1990); it has been applied to bluff body wakes,

jets, and other configurations (see Chomaz et al. 1988; Cossu and Chomaz 1997; Lauga and Bewley 2004; Bagheri et al. 2009).

The Ginzburg–Landau equation describes the temporal and spatial evolution of velocity perturbations in a given flow configuration. It is a complex partial differential equation first-order in time and second-order in one spatial dimension, the streamwise direction. As such, it serves as a convenient starting point for establishing a flow control model without the complexities of a full two- or three-dimensional Navier–Stokes system. If flow field perturbations are sufficiently small, then the equation’s cubic term can be truncated to form a linear partial differential equation.

The articles by Bagheri et al. (2009) and Lauga and Bewley (2004) employed a flow control framework for the linearised Ginzburg–Landau equation. They modelled the flow dynamics, actuator effect, and probe sensing as a standard linear state space. This representation carries the benefit that linear time-invariant (LTI) control theory is a highly developed field of study; results from optimal control, robust control, and model reduction are directly applicable. These papers analysed the transient and long-time stability of the Ginzburg–Landau system using a single-input, single-output (SISO) feedback loop—that is, with one actuator manipulating the flow at one point in space, acting on information from one sensor located at another point in space. Bagheri et al. emphasised model reduction of the system and linear quadratic Gaussian (LQG) control (which is a specific instance of \mathcal{H}_2 optimal control) in the presence of an upstream disturbance. Lauga and Bewley, on the other hand, focused on \mathcal{H}_∞ -based robust control.

The present study expands on the \mathcal{H}_2 optimal control framework by investigating where actuators and sensors should be placed for maximum reduction of the flow perturbation and actuator effort, given an exogenous flow disturbance. The current fluid mechanics literature has few analytically rigorous studies on optimal actuator and sensor placement. Bagheri et al. (2009) and Åkervik et al. (2007) stated that sensor locations should overlap with unstable global modes, and actuator locations with corresponding adjoint modes; they placed their sensor and actuator at the mode magnitude maxima. These detectability and stabilisability conditions are necessary for the construction of the \mathcal{H}_2 optimal control solution (Doyle et al., 1989). As will become evident in this paper, however, actuator and sensor placement by modal analysis yields suboptimal control. Although such a placement may be sufficient for preventing large-scale self-sustaining global oscillations, it may be insufficient for applications where transient growth needs to be reduced.

In an early study, Strykowski and Sreenivasan (1990) placed a small secondary cylinder in the wake of a larger primary cylinder, and plotted the placements where the secondary cylinder was passively able to suppress vortex shedding for given Reynolds numbers. Giannetti and Luchini (2007) showed that Strykowski and Sreenivasan’s placements corresponded with the wavemaker region of the Navier–Stokes operator linearised about the unstable equilibrium wake. Gillies (2001) similarly placed an actuator in the absolutely unstable part of the domain, and experimentally showed that his controllers were successful only when the sensors were also in the absolutely unstable domain. These results were consistent with those of Roussopoulos (1993), who experimented with sensor placement in a wind tunnel and a water channel. Lauga and Bewley (2004) placed an actuator and a sensor in the wavemaker region of the Ginzburg–Landau equation. In an earlier study, Lauga and Bewley (2003) employed a linear quadratic regulator with full-state feedback; they tracked the movement of the optimal actuator placement as the supercriticality parameter increased.

Much of the analytically rigorous literature on optimal actuator and sensor placement is related to the vibration control of flexible structures. For instance, Moheimani and Ryall (1999) placed collocated actuator and sensor pairs where they were able to control all important vibration modes. Kondoh, Yatomi, and Inoue (1990) employed Powell’s method of multidimensional function minimisation to locate \mathcal{H}_2 optimal actuator and sensor positions, but reported that the

iterative technique became difficult as the number of actuators and sensors increased. Hiramoto, Doki, and Obinata (2000) formulated the gradient of the controlled system’s \mathcal{H}_2 norm with respect to actuator and sensor locations, allowing the use of gradient-based minimisation techniques. The framework developed, however, assumed collocation of the actuator and sensor.

Vibration control of flexible structures benefits from the fact that the dynamical evolution operator is normal (Hiramoto et al., 2000), and eigenfunctions are orthogonal (Halim and Moheimani, 2003). On the other hand, the linearised Ginzburg–Landau and Navier–Stokes operators are non-normal, and their global modes do not form an orthogonal basis (Cossu and Chomaz, 1997; Chomaz, 2005; Bagheri et al., 2009). Therefore, methods for optimal placement in the flexible structures literature generally do not apply to fluid systems. Ideally, techniques should analytically or numerically compute optimal placement without iteration. In the absence of such techniques, the Powell’s method of Kondoh et al. (1990) and the gradient-based minimisation of Hiramoto et al. (2000) serve as good starting points for the Ginzburg–Landau system.

In this paper, we seek the actuator and sensor placement that minimises the perturbation magnitude and actuator effort of the Ginzburg–Landau system in an \mathcal{H}_2 sense. We investigate cases in which a white noise flow disturbance is present everywhere in the domain, or only in the upstream region. The underlying idea is that while \mathcal{H}_2 optimal control theory can provide the “best” controller given an actuator and sensor placement, the “best” placement should ideally be sought in conjunction. We develop an improved formulation of Hiramoto et al.’s \mathcal{H}_2 norm gradient. We employ the conjugate gradient method of multidimensional function minimisation for SISO as well as multiple-input, multiple-output (MIMO) control of the Ginzburg–Landau equation. For the SISO cases, we compare the conjugate gradient results to brute force sampling results, where we evaluate the \mathcal{H}_2 optimal controller and the corresponding system norm for a large test matrix of actuator and sensor positions. Finally, we discuss the failure of global mode and Gramian analyses to yield the optimal placement, and we suggest that the wavemaker region is a better indicator of the optimal placement.

8.2 Complex Ginzburg–Landau equation

8.2.1 Continuous representation

A complete discussion of the Ginzburg–Landau equation is too expansive to be given here. For a more comprehensive review, refer to the review article by Bagheri et al. (2009).

For a velocity or streamfunction perturbation amplitude given by the real part of $q(x, t)$, the complex Ginzburg–Landau equation is

$$\frac{\partial q}{\partial t} + \nu \frac{\partial q}{\partial x} = \mu(x) q + \gamma \frac{\partial^2 q}{\partial x^2} - a |q|^2 q. \quad (8.1)$$

We choose the spatial domain of q to be fully infinite, requiring $q(\pm\infty, t) = 0$ (Cossu and Chomaz, 1997; Chomaz, 2005). This allows perturbations to grow and decay throughout the entire streamwise dimension. With this choice of a domain, we typically assume that the physical feature of interest is at $x = 0$. The complex advection speed is $\nu = U + 2ic_u$, the complex diffusion parameter is $\gamma = 1 + ic_d$, and the amplification factor is $\mu(x) = \mu_0 - c_u^2 + \mu_2 x^2/2$ (Bagheri et al., 2009). The mean advection velocity is given by $U > 0$, and c_u , μ_0 , and a are also positive quantities; μ_2 and c_d are negative (Cossu and Chomaz, 1997; Chomaz, 2005). Numerical values of these parameters are determined empirically (Williamson, 1996; Roussopoulos and Monkewitz, 1996).

The nonlinear term $-a|q|^2 q$ causes a Hopf bifurcation expected in many geometrical configurations (e.g. Sreenivasan, Strykowski, and Olinger, 1987). In a linear analysis, the flow is assumed to be near the equilibrium state $q = 0$; thus, $-a|q|^2 q \approx 0$. The resulting linear equation is

$$\dot{q} = \mathcal{L}q, \quad (8.2a)$$

where

$$\mathcal{L} \triangleq -\nu \frac{\partial}{\partial x} + \mu(x) + \gamma \frac{\partial^2}{\partial x^2} \quad (8.2b)$$

defines the linear operator \mathcal{L} . From this point forward, we will always truncate the cubic term and refer only to the linear system.

We now comment on the stability properties of the linearised Ginzburg–Landau system, drawing from results found in Bagheri et al. (2009) and Cossu and Chomaz (1997). In the parabolic construction $\mu(x) = \mu_0 - c_u^2 + \mu_2 x^2/2$, the region of amplification is the part of the domain where $\mu(x) > 0$; this corresponds to $-\sqrt{-2(\mu_0 - c_u^2)/\mu_2} < x < \sqrt{-2(\mu_0 - c_u^2)/\mu_2}$. The quantity $\mu_0 - c_u^2$ is the greatest amplification factor present in the physical domain, occurring at $x = 0$; μ_0 is derived from the Reynolds number. If $\mu_0 - c_u^2 < 0$, then the flow is stable everywhere, and all perturbations decay with time. Otherwise, at least some part of the domain exhibits amplification (convective or absolute instability), and c_u is the most unstable spatial wavenumber.

In the analysis of parallel flows (where μ is constant), we construct a transitional value $\mu_t = U_{max}^2 / (4|\gamma|^2)$, with $U_{max} = U + 2c_d c_u$ the group velocity. When $0 < \mu + c_u^2 < \mu_t$,¹ the domain is convectively unstable. The response $G(x, t)$ to an impulse $\delta(x)\delta(t)$ blows up along $x = kt$ for at least one value of k as $t \rightarrow \infty$, but $\lim_{t \rightarrow \infty} G(0, t) = 0$. When $\mu_t < \mu + c_u^2$,¹ the domain is absolutely unstable, and $\lim_{t \rightarrow \infty} G(0, t) = \infty$ (Schmid and Henningson, 2000).

Returning to the non-parallel case, where μ is parabolic in space, we also define a critical value $\mu_c = \mu_t + |h| \cos((\text{Arg } \gamma)/2)/2$, where $h = \sqrt{-2\mu_2\gamma}$. Global stability is lost if and only if $\mu_0 > \mu_c$; this is a direct result of the eigendecomposition of \mathcal{L} . Eigenvalues, eigenmodes, and adjoint modes are given respectively by

$$\lambda_n = \mu_0 - c_u^2 - \frac{\nu^2}{4\gamma} - \left(n + \frac{1}{2}\right)h \quad (8.3a)$$

$$\phi_n(x) = \exp\left(\frac{1}{2}\left(\frac{\nu x}{\gamma} - \chi^2 x^2\right)\right) H_n(\chi x) \quad (8.3b)$$

$$\psi_n(x) = \exp\left(-\frac{\bar{\nu}x}{\bar{\gamma}}\right) \bar{\phi}_n(x), \quad (8.3c)$$

where $n = 0, 1, \dots$, $\chi = (-\mu_2/(2\gamma))^{1/4}$, H_n is the n^{th} Hermite polynomial, and $\overline{(\cdot)}$ indicates a complex conjugate. If $\text{Re}[\lambda_0] > 0$, then the system must be globally unstable. Otherwise, it is globally stable and perturbations eventually decay with time, but transient perturbation growth may still exist.

Finally, we note briefly that μ_2 denotes how non-parallel the flow is; it is partly responsible for determining the size of the amplification region. In addition, c_d is a dispersion parameter, but we will not discuss it further. See Bagheri et al. (2009) for additional details.

The Ginzburg–Landau parameters chosen for this study are drawn from Bagheri et al. (2009) and are shown in Table 8.1. This study primarily investigates a supercritical, globally unstable

¹In the published article, $\mu + c_u^2$ appears incorrectly as μ .

Table 8.1: Ginzburg–Landau parameters. Subcritical parameters are shown in parentheses.

variable	description	value
U	mean advection velocity	2.0
c_u	most unstable wavenumber	0.2
c_d	dispersion parameter	-1.0
μ_0	overall amplification	0.41 (0.38)
μ_2	degree of non-parallelism	-0.01
μ_t	transitional μ	0.32
μ_c	critical μ	0.40
a	nonlinearity	not used
x_1, x_N	extent of discretised domain	-56.06, 56.06
$\pm \sqrt{-2(\mu_0 - c_u^2)/\mu_2}$	extent of region of amplification	± 8.60 (± 8.25)
x_d	upstream disturbance location	-11.0
σ	Gaussian width	0.4

system, with $\mu_0 = 1.03\mu_c$. When we consider subcritical, globally stable systems, however, we use $\mu_0 = 0.96\mu_c$. We always assume supercriticality unless stated otherwise.

8.2.2 Discrete representation

To model the Ginzburg–Landau equation as an LTI state space, we spatially discretise the physical domain. We create a domain discretisation and a pseudospectral formulation of $\partial/\partial x$ and $\partial^2/\partial x^2$ using weighted Hermite polynomials, as in Bagheri et al. (2009). We draw the discretisation and associated differentiation operators from Weideman and Reddy (2000). The physical domain is represented using $N = 100$ grid points, and these points $\{x_1, x_2, \dots, x_N\}$ are the roots of $H_N(\text{Re}[\chi]x)$. With this choice of N , the discretised domain stretches from $x_1 = -56.06$ to $x_N = 56.06$. For the globally unstable system studied, the region of amplification is $-8.60 < x < 8.60$. We choose $N = 100$ because it achieves sufficient accuracy while maintaining problem tractability. To check for convergence of results in N , we test systems with one actuator and one sensor, and with two actuators and two sensors. In both cases, we introduce flow disturbances across the entire spatial domain (see Section 8.3.1 for details). Increasing N from 100 to 200 changes the computed optimal actuator and sensor placement by less than 1.0% of the amplification region’s length.

Using this discretisation scheme, we construct the state vector $\mathbf{q} = [q_1 \ \cdots \ q_N]^T$, where $q_j = q(x_j)$ for $j = 1, 2, \dots, N$. Let \mathbf{D} be the Hermite pseudospectral derivative matrix, and $\mathbf{\Delta}$ be the second derivative matrix. Also, let $\boldsymbol{\mu}$ be $\mu(x)$ discretised in the same way as \mathbf{q} . Then the discrete representation of the linearised Ginzburg–Landau equation is

$$\dot{\mathbf{q}} = \mathbf{A}\mathbf{q}, \quad (8.4a)$$

where

$$\mathbf{A} \triangleq -\nu\mathbf{D} + \text{diag}(\boldsymbol{\mu}) + \gamma\mathbf{\Delta} \quad (8.4b)$$

is the discrete linearised Ginzburg–Landau operator.

8.3 \mathcal{H}_2 optimal control and placement

In this section, we apply techniques developed in the control theory community, and we use these methods to develop a gradient-based formulation for optimising the actuator and sensor placement. The review article by Kim and Bewley (2007) provides a quick introduction to optimal control theory and its application to fluid mechanics. A complete reference can be found in Skogestad and Postlethwaite (2005). The seminal paper by Doyle et al. (1989) establishes the \mathcal{H}_2 optimal control theory used in this section. Finally, Schmid (2007) reviews important concepts from linear dynamics, such as transfer functions and stochastic dynamics, that we employ.

8.3.1 State space setup

Let us return momentarily to the continuous-space linearised Ginzburg–Landau equation. Instead of implementing a particular type of flow actuator or sensing probe, we apply generic actuation and sensing drawn from the standard state space setup. This way, we can easily take advantage of available LTI tools. We also insert a white noise disturbance d , which may be present everywhere in the spatial domain, or centred in the upstream domain. Assume a small sensor white noise n . For a single actuator located at $x = x_a$ and a single sensor at $x = x_s$, the state space is

$$\frac{\partial q}{\partial t}(x, t) = \mathcal{L}q(x, t) + \exp\left(-\frac{(x - x_a)^2}{2\sigma^2}\right)u(t) + d(x, t) \quad (8.5a)$$

$$y(t) = \left\langle q(x, t), \exp\left(-\frac{(x - x_s)^2}{2\sigma^2}\right) \right\rangle + n(t), \quad (8.5b)$$

where $\langle f_1, f_2 \rangle \triangleq \int_{-\infty}^{\infty} \bar{f}_2 f_1 dx$ is an inner product. The state q , input signal u , output signal y , disturbance d , and noise n are all complex, and we choose $\sigma = 0.4$.

In discretised form, with \mathbf{A} given by (8.4b) and all other variables explained below, we recover the standard state space

$$\dot{\mathbf{q}} = \mathbf{A}\mathbf{q} + \mathbf{B}_2\mathbf{u} + \mathbf{W}^{\frac{1}{2}}\mathbf{d} \quad (8.6a)$$

$$\mathbf{y} = \mathbf{C}_2\mathbf{q} + \mathbf{V}^{\frac{1}{2}}\mathbf{n}. \quad (8.6b)$$

In SISO control, \mathbf{u} and \mathbf{y} are complex scalars, and \mathbf{B}_2 is an N -by-1 vector discretising the distribution $\exp\left(-\frac{(x - x_a)^2}{(2\sigma^2)}\right)$. For discretised sensing, note that a weighting matrix must be used to carry out the integration implied by (8.5b). Let

$$\mathbf{M} = \frac{1}{2}\text{diag}\left(\begin{bmatrix} x_2 - x_1 & x_3 - x_1 & \cdots & x_{i+1} - x_{i-1} & \cdots & x_m - x_{m-2} & x_m - x_{m-1} \end{bmatrix}\right) \quad (8.7)$$

be the weighting matrix corresponding to the trapezoidal integration operator, and let \mathbf{s} be the n -by-1 discretisation of $\exp\left(-\frac{(x - x_s)^2}{(2\sigma^2)}\right)$. Then the inner product in (8.5b) transforms to $\langle \mathbf{q}, \mathbf{s} \rangle$, where $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle = \mathbf{v}_2^* \mathbf{M} \mathbf{v}_1$ is the discrete inner product and $(\cdot)^*$ indicates a conjugate transpose. This yields $\mathbf{C}_2 = \mathbf{s}^T \mathbf{M}$. In the case of MIMO control, suppose that there exist m_a actuators and m_s sensors. Let \mathbf{b}^j be the discretisation of the j th actuator, and \mathbf{s}^k the k th sensor. Then $\mathbf{B}_2 = [\mathbf{b}^1 \ \cdots \ \mathbf{b}^{m_a}]$, and $\mathbf{C}_2 = [\mathbf{s}^1 \ \cdots \ \mathbf{s}^{m_s}]^T \mathbf{M}$.

The disturbance and sensor noise are complex white noise stochastic processes. In the case of a disturbance pervading the entire domain, \mathbf{d} is an N -by-1 complex white noise signal with covariance $E(\mathbf{d}\mathbf{d}^*) = \mathbf{I}$ (with E the expected value), and we choose $\mathbf{W}^{\frac{1}{2}} = \mathbf{I}$. If the disturbance

is upstream, then \mathbf{d} is a complex white noise scalar with unit covariance, and $\mathbf{W}^{\frac{1}{2}}$ is the N -by-1 vector discretisation of $\exp\left(-(x - x_d)^2 / (2\sigma^2)\right)$; x_d is fixed at -11.0 . We do not wish to place any particular focus on sensor noise in this study, but the noise must be present for the LQG formulation to be well-posed. Therefore, \mathbf{n} is an m_s -by-1 complex white noise signal with covariance $E(\mathbf{n}\mathbf{n}^*) = \mathbf{I}$, and we choose $\mathbf{V}^{\frac{1}{2}} = 2 \cdot 10^{-4}\mathbf{I}$ so that the sensor noise is present but minimal. It will be apparent later that the LQG controller remains well-behaved for this small choice of $\mathbf{V}^{\frac{1}{2}}$.

8.3.2 Linear quadratic Gaussian framework

Having set up the governing dynamics and input/output relations in (8.5–8.6b), we now design a controller that determines the control signal $u(t)$ from the sensor output $y(t)$. In particular, we seek the relation from y to u that minimises the perturbation magnitude $\int_{-\infty}^{\infty} |q(x, t)|^2 dx$. This minimisation is the ultimate goal of Ginzburg–Landau control. Realistically, however, we must also keep the input size $|u(t)|^2$ within bounds; a very large input size may be non-physical. Therefore, we establish a cost function

$$J = \beta^2 \int_{-\infty}^{\infty} |q(x, t)|^2 dx + |u(t)|^2 \quad (8.8)$$

to minimise using the right choice of a controller. The positive scalar β controls the balance between maintaining a small perturbation size and maintaining a small input size. In this study, we choose $\beta = 7$. We discuss different choices of β in Section 8.4.3.

In discrete space, the cost function takes the form $J = \mathbf{q}^* \mathbf{Q} \mathbf{q} + \mathbf{u}^* \mathbf{R} \mathbf{u}$, where \mathbf{Q} and \mathbf{R} are user-specified penalties on the state and inputs sizes. To be consistent with (8.8), we use $\mathbf{Q} = \beta^2 \mathbf{M}$ and $\mathbf{R} = \mathbf{I}$. In vector form, the cost on the perturbation and input sizes are given by $\mathbf{j}_1 = \mathbf{Q}^{\frac{1}{2}} \mathbf{q}$ and $\mathbf{j}_2 = \mathbf{R}^{\frac{1}{2}} \mathbf{u}$, where the matrix square root $\mathbf{Z}^{\frac{1}{2}}$ satisfies $(\mathbf{Z}^{\frac{1}{2}})^* \mathbf{Z}^{\frac{1}{2}} = \mathbf{Z}$. Since the disturbance \mathbf{d} and sensor noise \mathbf{n} are exogenous inputs, minimising J is equivalent to minimising the gain from \mathbf{d} and \mathbf{n} to the cost vectors \mathbf{j}_1 and \mathbf{j}_2 .

To establish the relation from \mathbf{d} and \mathbf{n} to \mathbf{j}_1 and \mathbf{j}_2 , we first incorporate the cost vectors into the state space in (8.6b) to yield

$$\mathbf{P} : \begin{bmatrix} \dot{\mathbf{q}} \\ \mathbf{z} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B}_1 & \mathbf{B}_2 \\ \mathbf{C}_1 & \mathbf{0} & \mathbf{D}_{12} \\ \mathbf{C}_2 & \mathbf{D}_{21} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \mathbf{w} \\ \mathbf{u} \end{bmatrix}, \quad (8.9a)$$

where

$$\mathbf{z} = \begin{bmatrix} \mathbf{j}_1 \\ \mathbf{j}_2 \end{bmatrix} \quad \mathbf{w} = \begin{bmatrix} \mathbf{d} \\ \mathbf{n} \end{bmatrix} \quad (8.9b)$$

and

$$\mathbf{B}_1 = \begin{bmatrix} \mathbf{W}^{\frac{1}{2}} & \mathbf{0} \end{bmatrix}, \quad \mathbf{C}_1 = \begin{bmatrix} \mathbf{Q}^{\frac{1}{2}} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{D}_{12} = \begin{bmatrix} \mathbf{0} \\ \mathbf{R}^{\frac{1}{2}} \end{bmatrix}, \quad \mathbf{D}_{21} = \begin{bmatrix} \mathbf{0} & \mathbf{V}^{\frac{1}{2}} \end{bmatrix}. \quad (8.9c)$$

Equation (8.9a) is called the “plant” or simply \mathbf{P} . Figure 8.1 is a pictorial representation of the control problem. The controller \mathbf{K} determines the input signal \mathbf{u} based on the sensor output \mathbf{y} . The external input to the system is the flow disturbance and sensor noise in \mathbf{w} , and the system output is the cost \mathbf{z} to be minimised.

In the construction of the general LTI controller, we choose matrices \mathbf{K}_A , \mathbf{K}_B , \mathbf{K}_C , and \mathbf{K}_D such that

$$\mathbf{K} : \begin{bmatrix} \dot{\hat{\mathbf{q}}} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{K}_A & \mathbf{K}_B \\ \mathbf{K}_C & \mathbf{K}_D \end{bmatrix} \begin{bmatrix} \hat{\mathbf{q}} \\ \mathbf{y} \end{bmatrix} \quad (8.10)$$




Figure 8.1: General control problem formulation \mathbf{G} .

determines the actuation signal \mathbf{u} based on the output signal \mathbf{y} . The variable $\hat{\mathbf{q}}$ is the state of the controller, and is typically an estimate of the plant state \mathbf{q} . The LQG controller is the LTI controller that minimises the \mathcal{H}_2 norm of the transfer function from \mathbf{d} and \mathbf{n} to \mathbf{j}_1 and \mathbf{j}_2 . Suppose that $\mathbf{G}(s)$ is that matrix transfer function. In other words, for $s \in \mathbb{C}$ and disturbance and sensor noise inputs of the form e^{st} ,

$$\begin{bmatrix} \mathbf{j}_1 \\ \mathbf{j}_2 \end{bmatrix} = \mathbf{G}(s) \begin{bmatrix} \mathbf{d} \\ \mathbf{n} \end{bmatrix} \quad (8.11)$$

dictates the gain and phase from \mathbf{d} and \mathbf{n} to \mathbf{j}_1 and \mathbf{j}_2 . Then the \mathcal{H}_2 norm is defined as

$$\gamma_2 \triangleq \|\mathbf{G}\|_2 \triangleq \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{tr}(\mathbf{G}^*(i\omega) \mathbf{G}(i\omega)) d\omega}, \quad (8.12)$$

and is an average (instead of a worst-case) measure of the gain provided by \mathbf{G} . Thus, given exogenous inputs, minimising the perturbation magnitude and actuator effort is equivalent to minimising $\|\mathbf{G}\|_2$. We implement the LQG controller so that $\|\mathbf{G}\|_2$ is automatically minimised.

To construct the LQG controller, solve the continuous algebraic Riccati equations

$$\mathbf{A}^* \mathbf{X} + \mathbf{X} \mathbf{A} - \mathbf{X} \mathbf{B}_2 \mathbf{R}^{-1} \mathbf{B}_2^* \mathbf{X} + \mathbf{C}_1^* \mathbf{C}_1 = \mathbf{0} \quad (8.13a)$$

$$\mathbf{A} \mathbf{Y} + \mathbf{Y} \mathbf{A}^* - \mathbf{Y} \mathbf{C}_2^* \mathbf{V}^{-1} \mathbf{C}_2 \mathbf{Y} + \mathbf{B}_1 \mathbf{B}_1^* = \mathbf{0} \quad (8.13b)$$

for \mathbf{X} and \mathbf{Y} , and set $\mathbf{F} = \mathbf{R}^{-1} \mathbf{B}_2^* \mathbf{X}$ and $\mathbf{L} = \mathbf{Y} \mathbf{C}_2^* \mathbf{V}^{-1}$. The LQG controller is then given by

$$\mathbf{K} : \begin{bmatrix} \dot{\hat{\mathbf{q}}} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{A} - \mathbf{B}_2 \mathbf{F} - \mathbf{L} \mathbf{C}_2 & \mathbf{L} \\ -\mathbf{F} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{q}} \\ \mathbf{y} \end{bmatrix}. \quad (8.14)$$

The plant (8.9a) and LQG controller (8.14) can be combined to form a single system, known as the linear fractional transformation (LFT) of \mathbf{G} and \mathbf{K} . One representation of the LFT is

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\hat{\mathbf{q}}} \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & -\mathbf{B}_2 \mathbf{F} & \mathbf{B}_1 \\ \mathbf{L} \mathbf{C}_2 & \mathbf{A} - \mathbf{B}_2 \mathbf{F} - \mathbf{L} \mathbf{C}_2 & \mathbf{L} \mathbf{D}_{21} \\ \mathbf{C}_1 & -\mathbf{D}_{12} \mathbf{F} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \hat{\mathbf{q}} \\ \mathbf{w} \end{bmatrix}. \quad (8.15)$$

(Often, the LFT is expressed using an error $\mathbf{e} = \hat{\mathbf{q}} - \mathbf{q}$ instead of directly using the state estimate $\hat{\mathbf{q}}$.) Denote the upper-left two-by-two matrix block by \mathbf{Z}_A , the upper-right two-by-one block by \mathbf{Z}_B , and the lower-left one-by-two block by \mathbf{Z}_C . If we compute Gramians \mathbf{W}_c and \mathbf{W}_o by the Lyapunov

equations

$$\mathbf{Z}_A \mathbf{W}_c + \mathbf{W}_c \mathbf{Z}_A^* + \mathbf{Z}_B \mathbf{Z}_B^* = \mathbf{0} \quad (8.16a)$$

$$\mathbf{Z}_A^* \mathbf{W}_o + \mathbf{W}_o \mathbf{Z}_A + \mathbf{Z}_C^* \mathbf{Z}_C = \mathbf{0}, \quad (8.16b)$$

then the LQG-controlled system's \mathcal{H}_2 squared norm is

$$\Gamma_2 \triangleq \|\mathbf{G}\|_2^2 = \text{tr}(\mathbf{Z}_C \mathbf{W}_c \mathbf{Z}_C^*) \quad (8.17a)$$

$$= \text{tr}(\mathbf{Z}_B^* \mathbf{W}_o \mathbf{Z}_B). \quad (8.17b)$$

The above formulations, however, require solutions of $2N$ -by- $2N$ matrix equations. A simpler N -by- N formulation is given by

$$\Gamma_2 = \text{tr}(\mathbf{C}_1 \mathbf{Y} \mathbf{C}_1^*) + \text{tr}(\mathbf{V}^{-1} \mathbf{C}_2 \mathbf{Y} \mathbf{X} \mathbf{Y} \mathbf{C}_2^*) \quad (8.18a)$$

$$= \text{tr}(\mathbf{B}_1^* \mathbf{X} \mathbf{B}_1) + \text{tr}(\mathbf{R}^{-1} \mathbf{B}_2^* \mathbf{X} \mathbf{Y} \mathbf{X} \mathbf{B}_2). \quad (8.18b)$$

8.3.3 Perturbation technique for optimal placement

Suppose that a choice of \mathbf{B}_2 and \mathbf{C}_2 is \mathcal{H}_2 optimal. That is, the choice of \mathbf{B}_2 and \mathbf{C}_2 provides the best-performing \mathcal{H}_2 optimal controller, given matrix constraints described in Section 8.3.1. It is necessary, but not sufficient, for an appropriately constrained perturbation $\delta \mathbf{B}_2 \neq \mathbf{0}$ to produce $\delta \Gamma_2 = 0$; similarly, $\delta \mathbf{C}_2 \neq \mathbf{0}$ must produce $\delta \Gamma_2 = 0$. The constraint on $\delta \mathbf{B}_2$ is that it must equal the change in \mathbf{B}_2 when an actuator position moves from x_a to $x_a + \delta x_a$, for a small δx_a . If there is multiple actuation, then $\delta \mathbf{B}_2$ can be a perturbation along any one of the actuator positions. The constraint on $\delta \mathbf{C}_2$ is completely analogous, except that it is necessary to pay careful attention to the discretised integration operator \mathbf{M} .

If the actuator location is perturbed from \mathbf{B}_2 to $\mathbf{B}_2 + \delta \mathbf{B}_2$, then we can define $\delta \mathbf{X}$ such that (8.13a) produces a perturbed solution $\mathbf{X} + \delta \mathbf{X}$. The exact representation is

$$\mathbf{A}^* (\mathbf{X} + \delta \mathbf{X}) + (\mathbf{X} + \delta \mathbf{X}) \mathbf{A} - (\mathbf{X} + \delta \mathbf{X}) (\mathbf{B}_2 + \delta \mathbf{B}_2) \mathbf{R}^{-1} (\mathbf{B}_2 + \delta \mathbf{B}_2)^* (\mathbf{X} + \delta \mathbf{X}) + \mathbf{C}_1^* \mathbf{C}_1 = \mathbf{0}. \quad (8.19)$$

The $O(\delta^0)$ component returns (8.13a), as expected. The $O(\delta^1)$ component is

$$(\mathbf{A} - \mathbf{B}_2 \mathbf{R}^{-1} \mathbf{B}_2^* \mathbf{X})^* \delta \mathbf{X} + \delta \mathbf{X} (\mathbf{A} - \mathbf{B}_2 \mathbf{R}^{-1} \mathbf{B}_2^* \mathbf{X}) - \mathbf{X} (\delta \mathbf{B}_2 \mathbf{R}^{-1} \mathbf{B}_2^* + \mathbf{B}_2 \mathbf{R}^{-1} \delta \mathbf{B}_2^*) \mathbf{X} = \mathbf{0}, \quad (8.20)$$

which is a Lyapunov equation for $\delta \mathbf{X}$. The perturbation in \mathbf{X} then produces the cost perturbation $\delta \Gamma_2 = \text{tr}(\mathbf{V}^{-1} \mathbf{C}_2 \mathbf{Y} \delta \mathbf{X} \mathbf{Y} \mathbf{C}_2^*)$ from (8.18a). At the optimal actuator location,

$$\delta \Gamma_2 = \text{tr}(\mathbf{V}^{-1} \mathbf{C}_2 \mathbf{Y} \delta \mathbf{X} \mathbf{Y} \mathbf{C}_2^*) = 0. \quad (8.21)$$

In summary, given an appropriately constrained \mathbf{B}_2 and $\delta \mathbf{B}_2$, the \mathcal{H}_2 optimal \mathbf{B}_2 is one that produces $\text{tr}(\mathbf{V}^{-1} \mathbf{C}_2 \mathbf{Y} \delta \mathbf{X} \mathbf{Y} \mathbf{C}_2^*) = 0$, where $\delta \mathbf{X}$ is the solution to (8.20) and \mathbf{X} is the solution to (8.13a). For this choice of \mathbf{B}_2 , the actuator position and the controller are simultaneously optimised to minimise the magnitude of flow perturbations and actuator input. Additional steps are required, however, to ensure that the \mathbf{B}_2 satisfying $\delta \Gamma_2 = 0$ actually minimises Γ_2 , instead of yielding a maximum or saddle point.

The procedure for finding the optimal sensor location is similar to the previously described steps. For an appropriate choice of \mathbf{C}_2 and $\delta\mathbf{C}_2$, the optimal \mathbf{C}_2 is the one that first yields \mathbf{Y} by (8.13b), then yields $\delta\mathbf{Y}$ by the Lyapunov equation

$$\left(\mathbf{A} - \mathbf{Y}\mathbf{C}_2^*\mathbf{V}^{-1}\mathbf{C}_2\right)\delta\mathbf{Y} + \delta\mathbf{Y}\left(\mathbf{A} - \mathbf{Y}\mathbf{C}_2^*\mathbf{V}^{-1}\mathbf{C}_2\right)^* - \mathbf{Y}\left(\delta\mathbf{C}_2^*\mathbf{V}^{-1}\mathbf{C}_2 + \mathbf{C}_2^*\mathbf{V}^{-1}\delta\mathbf{C}_2\right)\mathbf{Y} = \mathbf{0}, \quad (8.22)$$

and finally satisfies

$$\delta\Gamma_2 = \text{tr}\left(\mathbf{R}^{-1}\mathbf{B}_2^*\mathbf{X}\delta\mathbf{Y}\mathbf{X}\mathbf{B}_2\right) = 0. \quad (8.23)$$

Since the Riccati equations in (8.13) must typically be solved numerically, we do not propose an analytical procedure for finding actuator and sensor placements yielding $\delta\Gamma_2 = 0$. Instead, we now create a framework where a gradient-based function minimisation can be used to iterate toward the optimal placements. For a set of actuator positions $\{x_a^1, \dots, x_a^{m_a}\}$ and sensor positions $\{x_s^1, \dots, x_s^{m_s}\}$, Γ_2 is obtained by solving (8.13) and (8.18). For a perturbation of the j th actuator position x_a^j by an amount δx_a^j , solve (8.20) for $\delta\mathbf{X}^j$ and set

$$\frac{\delta\Gamma_2}{\delta x_a^j} = \frac{\text{tr}\left(\mathbf{V}^{-1}\mathbf{C}_2\mathbf{Y}\delta\mathbf{X}^j\mathbf{Y}\mathbf{C}_2^*\right)}{\delta x_a^j}. \quad (8.24a)$$

For a perturbation of the k th sensor position x_s^k by an amount δx_s^k , solve (8.22) for $\delta\mathbf{Y}^k$ and set

$$\frac{\delta\Gamma_2}{\delta x_s^k} = \frac{\text{tr}\left(\mathbf{R}^{-1}\mathbf{B}_2^*\mathbf{X}\delta\mathbf{Y}^k\mathbf{X}\mathbf{B}_2\right)}{\delta x_s^k}. \quad (8.24b)$$

Note that the Riccati solutions \mathbf{X} and \mathbf{Y} are computed while solving for Γ_2 , and can be reused while solving the Lyapunov equations needed for $\nabla\Gamma_2$. Thus, we solve two Riccati equations for every evaluation of Γ_2 , and $m_a + m_s$ Lyapunov equations for every evaluation of the $(m_a + m_s)$ -by-1 vector $\nabla\Gamma_2$.

This technique for computing $\nabla\Gamma_2$ is similar to the method proposed by Hiramoto et al. (2000), but is believed to be more computationally efficient. Because Hiramoto et al. computed Γ_2 using the LFT of the controlled system, each component of $\nabla\Gamma_2$ required the solution to a $2N$ -by- $2N$ Lyapunov equation; see (8.15–8.17). On the other hand, the method we propose reduces this size to N -by- N . Using the Lyapunov equation solvers selected for this study, our proposed method should provide a speedup by a factor of approximately six.

We employ a Polak-Ribière conjugate gradient method, using a modified Brent’s method of line minimisation that uses derivative information. Brent’s method also requires a bracketing function that provides bounds on the line minimum. These methods are robust in the sense that they are unlikely to lose numerical stability, and they are generally efficient as well. Implementation details can be found in Press et al. (2007).

8.4 Results

8.4.1 Brute force sampling

To understand the effect of SISO actuator and sensor positions more fully, we first perform a brute force sampling of the LQG-controlled system’s \mathcal{H}_2 norm. Instead of implementing a minimisation algorithm, we map out $\gamma_2(x_a, x_s)$ over a test matrix spanning actuator and sensor positions in $\{-15.00, -14.75, \dots, 15.00\}$. The results for a white noise disturbance present across the entire spatial domain are plotted in Figure 8.2(a). In Figure 8.2(b), we run a simulation of the LQG-

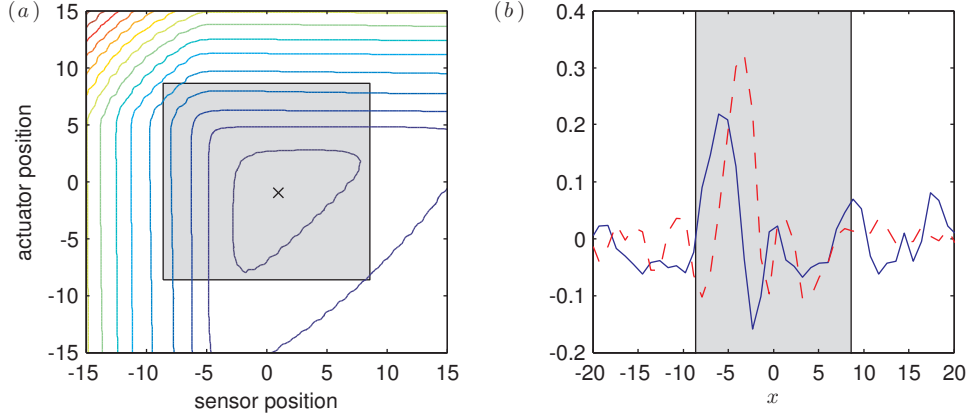


Figure 8.2: (a): Contours of $\log_{10} \gamma_2$ and the optimal placement (\times), with disturbances introduced everywhere. The innermost contour is $\log_{10} \gamma_2 = 2$, and each subsequent contour increments by 0.5. (b): A sample snapshot of the controlled perturbation q , shown as real (solid/blue) and imaginary (dashed/red) parts. The actuator is at $x_a = -1.00$ and the sensor is at $x_s = 1.00$. In both plots, the region of amplification is shaded grey.

controlled system and show a sample snapshot. Without control, the state would blow up because the system would be globally unstable.

Some qualitative features are immediately obvious from Figure 8.2(a). First, the optimal placement is approximately $x_a = -1.00$ and $x_s = 1.00$, at which $\gamma_2 = 46.1$; we will refine these numbers later through the use of iterative minimisation. Second, a penalty exists for placing the actuator too far downstream. This is an intuitive result—since the flow in question is convection-driven, a downstream actuator is able to influence less of the physical domain. Third, a similar penalty exists for placing the sensor too far upstream. Recall that the amplification factor $\mu(x)$ is an inverted parabola with roots at ± 8.60 . If the sensor is too far upstream, then the detected disturbances naturally dampen before they reach the region of amplification. Thus, the disturbances amplified in this region would not primarily be the ones observed by the sensor. Finally, excessive time delay (positioning the sensor too far downstream of the actuator) has a detrimental effect on perturbation control. This is also a result of the convection-driven nature of the system. If the sensor is too far downstream of the actuator, then the information it receives will have already passed the actuator for some time. Thus, the actuator would be acting on outdated information, and the feedback would effectively result in a time lag.

Data for a flow disturbance introduced around $x_d = -11.00$ are shown in Figure 8.3(a), with a sample snapshot shown in Figure 8.3(b). The qualitative features of this $\gamma_2(x_a, x_s)$ map are similar to those of Figure 8.2(a), but we highlight a few differences. First, the optimal placement is now approximately $x_a = -8.50$ and $x_s = -5.50$, which is much further upstream than in the previous case. In the sample snapshot shown in Figure 8.3(b), we note that the controller is extremely effective in reducing the perturbation magnitude within the region of amplification.

Intuitively, we might expect an optimal placement $x_s \approx x_d$, so that the sensor placement coincides with the disturbance source. Indeed, such a placement allows the state estimate $\hat{\mathbf{q}}$ to match the true state \mathbf{q} very closely in the disturbance region. This placement, however, is far from optimal. As a result of the strong amplification downstream of the disturbance source, even the smallest estimator error $\mathbf{e} = \hat{\mathbf{q}} - \mathbf{q}$ in the disturbance region can amplify significantly. By placing the sensor downstream of the disturbance source, the estimator is better able to reduce the error \mathbf{e} in the region of amplification, where accurate state estimation is critically important.

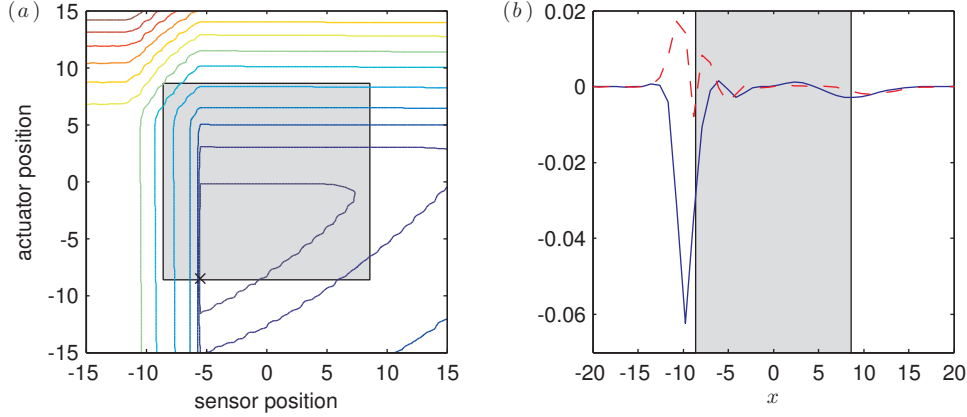


Figure 8.3: (a): Contours of $\log_{10} \gamma_2$ and the optimal placement (\times), for the case of an upstream disturbance introduced at $x_d = -11.00$. The innermost contour is $\log_{10} \gamma_2 = 1.0$, and each subsequent contour increments by 0.5. (b): A sample snapshot of the controlled perturbation q , shown as real (solid/blue) and imaginary (dashed/red) parts. The actuator is at $x_a = -8.50$ and the sensor is at $x_s = -5.50$. In both plots, the region of amplification is shaded grey.

The second difference in the case of upstream disturbances is that the optimal placement is not robust. Although the optimal placement is unique, Figure 8.3(a) shows that a small upstream shift in x_s drastically increases γ_2 . The topography of the $\gamma_2(x_a, x_s)$ function contains a cliff near $x_s = -5.50$ where the \mathcal{H}_2 norm increases suddenly. Holding $x_a = -8.50$ but moving $x_s = -5.50$ to -5.75 , the \mathcal{H}_2 norm γ_2 increases from 3.87 to 138. This occurs because the shift in the sensor location causes a small zero of the **u**-to-**y** transfer function to move from the left-half-plane to the right-half-plane. When this occurs, the small right-half-plane zero introduces significant bandwidth limitations for disturbance rejection. A sensitivity function analysis indicates considerably inferior disturbance rejection at frequencies $\omega < 10$.

This non-robustness appears to indicate the importance of placing the sensor firmly within the region of high amplification, lest estimator errors grow beyond acceptable limits. For instance, in an actual two- or three-dimensional flow setup without *a priori* knowledge of the optimal sensor placement, an experimenter could guess a sensor placement where amplification is expected. (Recall that the dynamics of fluid systems are often strongly non-normal (Chomaz, 2005), so regions of transient amplification may not correspond well with unstable eigenmodes of the linearised dynamics.)

Besides these two major differences, the general form of the γ_2 solution is similar to the case where disturbances are present everywhere. Level sets of γ_2 have a near-triangular form, with penalties on placing the actuator too far downstream, placing the sensor too far upstream, and creating excessive separation between sensors and actuators. The general similarities between figures 8.2(a) and 8.3(a) indicate that it is possible to find an actuator and sensor placement that is effective for both disturbance types. For example, if $x_a = -3.50$ and $x_s = 0.00$, then $\gamma_2 = 53.6$ when disturbances are present everywhere and $\gamma_2 = 5.90$ when disturbances are concentrated upstream. These are acceptably close to the respective optimal values of 46.1 and 3.86. If other disturbance types retain the same qualitative nature of $\gamma_2(x_a, x_s)$, then a “universally” optimal placement could feasibly be sought.

We also perform brute force sampling for the two-actuator, two-sensor configuration with disturbances present everywhere. The placement is optimal at about $x_a = -3.75, 2.75$ and $x_s = -2.75, 3.75$, for which $\gamma_2 = 34.0$. (Results are not plotted here.) As in the SISO case, these num-

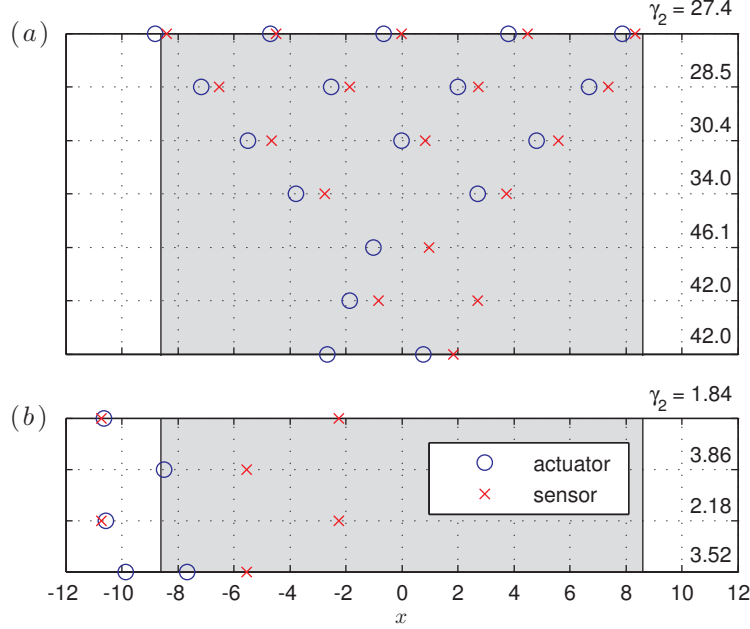


Figure 8.4: Optimal placement for different numbers of actuators and sensors. The \mathcal{H}_2 norm γ_2 is tabulated for each configuration, and the region of amplification is shaded grey. (a): Disturbances present everywhere. (b): Disturbances introduced at $x_d = -11.0$. The first row is the optimal placement for two actuators and two sensors; the actuators are collocated.

bers are not precise, and are later refined using an iterative procedure. One common feature in MIMO results is that if one actuator or sensor is placed near the optimal position, then the effect of the other actuator or sensor placement on γ_2 is generally not large.

8.4.2 Conjugate gradient minimisation

The optimal placement for different numbers of actuators and sensors, as found using conjugate gradient minimisation, is shown in Figure 8.4(a) for disturbances present everywhere. The iterative method more accurately puts the optimal SISO placement at $x_a = -1.03$, $x_s = 0.98$ and the optimal two-actuator, two-sensor placement at $x_a = -3.78, 2.71$, $x_s = -2.75, 3.74$. An approximate symmetry around $x = 0$ between actuator and sensor placement is apparent. Additionally, a simple pattern exists for configurations with an equal number of actuators and sensors. The optimal placement favors near-collocation of actuator and sensor pairs, with each actuator placed slightly upstream of each sensor. This allows the sensors to measure the effect of the actuators without a large time lag. As the number of actuators and sensors increases, the actuator-sensor pairs spread out across the spatial domain, and the spacing within pairs becomes tighter. Figure 8.4(a) shows that the placement of the pairs forms an orderly geometric pattern as the number of pairs increases.

A transfer function analysis of the controller gain \mathbf{K} from \mathbf{y} to \mathbf{u} shows that the communication between sensors and actuators primarily occurs within pairs, especially at higher frequencies. (See (8.14) for the formulation of \mathbf{K} .) This is illustrated in Figure 8.5, which supposes the optimally-placed five-input, five-output system shown in the top row of Figure 8.4(a). These three plots show the LQG gain from each sensor to each actuator, given that each sensor receives a signal $\exp(i\omega t)$. At lower frequencies (Figure 8.5a), the gain from downstream sensors to upstream actuators is still nontrivial. This feedback aids the performance of the upstream actuators. At higher frequencies,




Figure 8.5: The LQG gain in the optimally-placed five-input, five-output system, shown in decibels. The block in row j and column k shows the gain from a signal $\exp(i\omega t)$ in sensor k to actuator j , assuming that the sensors and actuators are ordered from upstream to downstream. (a): At $\omega = 10^{-1}$; (b): at $\omega = 10^1$; (c): at $\omega = 10^3$.

the primarily diagonal structure of the gain matrices indicates that communication outside the pairs is weak. The diffusive nature of the Ginzburg–Landau operator dampens high frequency oscillations, so such oscillations do not propagate effectively across large distances. Therefore, high-frequency control is most effective within actuator-sensor pairs.

The optimal placements in the case of upstream disturbances are illustrated in Figure 8.4(b). For the limited number of configurations tested, the optimal actuator and sensor placements are upstream of the origin. In SISO control, the optimal placement is $x_a = -8.48$ and $x_s = -5.55$. When employing one actuator and two sensors, the cost γ_2 is very insensitive to the precise placement of the downstream sensor. Moving that sensor’s location anywhere from $x = -5$ to 15 changes γ_2 by not more than 0.04%. Nevertheless, the combination of the upstream sensor near the disturbance source and the downstream sensor far in the wake allows for very accurate estimation in the entire domain of interest. The optimal two-actuator, two-sensor placement collocates the actuators and roughly reproduces the one-actuator, two-sensor placement. The fact that both actuators are optimally placed near the disturbance source suggests that no other actuator location can be as effective. Finally, we remark that the optimal two-actuator, one-sensor configuration is like the optimal SISO configuration, but it is possible to actuate near the disturbance source without causing serious time delays.

For each actuator-sensor configuration, Figure 8.4 also tabulates the \mathcal{H}_2 norm γ_2 from the disturbances to the perturbation and input costs. As expected, γ_2 decreases as the number of actuators and sensors increases. This simply indicates that the control objective is easier to attain when more actuators and sensors are available. In Figure 8.4(a), however, the decrease in γ_2 tapers off when increasing the number of actuator-sensor pairs from one to five. Furthermore, in Figure 8.4(b), the improvement to the base SISO case by adding one actuator is not as great as the improvement by adding one sensor. In choosing the number of actuators and sensors to use, an experimenter should consider a trade-off between how much the perturbation and input need to be minimised, and how many actuators and sensors can be afforded. Other factors, such as the way disturbances are introduced, also affect this decision. We may observe, for instance, that the multiplicative reduction in γ_2 from SISO to two-input, two-output control is more significant in Figure 8.4(b) than in Figure 8.4(a).

8.4.3 Further remarks

When performing an iterative technique, it is important to note the computational effort required. The linearised Ginzburg–Landau system presented here is accurately represented using only 100 states; however, iterating on higher-dimensional systems may be intractable if a large number of function calls is required. A complete discussion on the computational effort required to minimise Γ_2

Table 8.2: Optimal placement, perturbation size, and input size for different β , with disturbances present everywhere.

β	SISO				two-actuator, two-sensor				γ_q/β	γ_u
	x_a	x_s	γ_q/β	γ_u	x_a	x_s	γ_q/β	γ_u		
1	-1.39	0.04	7.14	4.15	-3.25	3.17	-2.77	2.86	5.23	3.24
3	-1.02	0.92	6.59	6.33	-3.72	2.57	-2.79	3.65	4.86	4.69
7	-1.03	0.98	6.48	8.22	-3.78	2.71	-2.75	3.74	4.77	6.20
15	-1.04	1.01	6.44	11.1	-3.79	2.70	-2.73	3.77	4.74	8.30
50	-1.04	1.03	6.42	19.0	-3.80	2.70	-2.71	3.79	4.72	14.3

is beyond the scope of this paper. The effort would heavily depend on the choice of a minimisation algorithm, its numerical tolerances, and the initial conditions chosen.

As previously mentioned, each evaluation of Γ_2 requires solving two N -by- N Riccati equations (8.13), and each evaluation of $\nabla\Gamma_2$ requires solving $m_a + m_s$ N -by- N Lyapunov equations (8.20, 8.22). The bracketing function additionally requires a number of N -by- N Riccati solutions that is generally *a priori* unknown. With the algorithms chosen, scaling experiments show that for $40 \leq N \leq 200$, the CPU times of one Riccati solution and one Lyapunov solution both scale by about $N^{2.5}$. On a 2.66 GHz dual-core Core i7-620M processor, each Riccati solution requires about 150 milliseconds for $N = 100$, and each Lyapunov equation requires about 59 milliseconds. The total number of equations to be solved depends heavily on the $\gamma_2(\{x_a\}, \{x_s\})$ topography and the choice of tolerances and initial conditions. It suffices to state that the CPU time for convergence increases significantly as the number of actuators and sensors increases.

For the results presented in this study, the $\gamma_2(\{x_a\}, \{x_s\})$ function topography is relatively simple. Global minima can be found by guessing multiple initial conditions and checking for common convergence. Performing a coarse brute force search also aids in picking good initial conditions. Beyond these simple configurations, however, convergence to a global minimum can be difficult. This is especially true when disturbances are only introduced upstream, or when unequal numbers of actuators and sensors are used. Good initial conditions are particularly important in these cases. Bad initial conditions may cause the iterator to converge very slowly, quit prematurely, or fail to converge to the global minimum. A global minimisation technique such as simulated annealing would then be necessary to ensure that a local minimum found is indeed the absolute minimum.

The choice of cost matrices also affects the computed optimal placement. As aforementioned, the vector costs used are $\mathbf{j}_1 = \mathbf{Q}^{\frac{1}{2}}\mathbf{q}$ and $\mathbf{j}_2 = \mathbf{R}^{\frac{1}{2}}\mathbf{u}$, where $\mathbf{Q}^{\frac{1}{2}} = \beta\mathbf{M}^{\frac{1}{2}}$ and $\mathbf{R}^{\frac{1}{2}} = \mathbf{I}$. The positive scalar β controls the balance between minimising \mathbf{q} and minimising \mathbf{u} . To quantify their sizes individually, let γ_q be the \mathcal{H}_2 norm of the transfer function from \mathbf{w} to \mathbf{j}_1 , and let γ_u be the \mathcal{H}_2 norm of the transfer function from \mathbf{w} to \mathbf{j}_2 . These are calculated using the LFT formulation (8.15) by respectively setting $\mathbf{Z}_C = \begin{bmatrix} \mathbf{C}_1 & \mathbf{0} \end{bmatrix}$ and $\mathbf{Z}_C = \begin{bmatrix} \mathbf{0} & -\mathbf{D}_{12}\mathbf{F} \end{bmatrix}$, and computing $\sqrt{\Gamma_2}$ by (8.16–8.17). It can be shown that $\gamma_q^2 + \gamma_u^2 = \gamma_2^2$. A measure of the perturbation size is then given by γ_q/β , and the input size by γ_u . For different choices of β , optimal positions and corresponding values of γ_q/β and γ_u are listed in Table 8.2. (Note that the optimal controller and positions depend not on \mathbf{Q} and \mathbf{R} individually, but rather only β .)

We immediately see that within a fairly large range of β , the optimal actuator and sensor placement is robust to the exact choice of β . This is consistent with findings by Lauga and Bewley

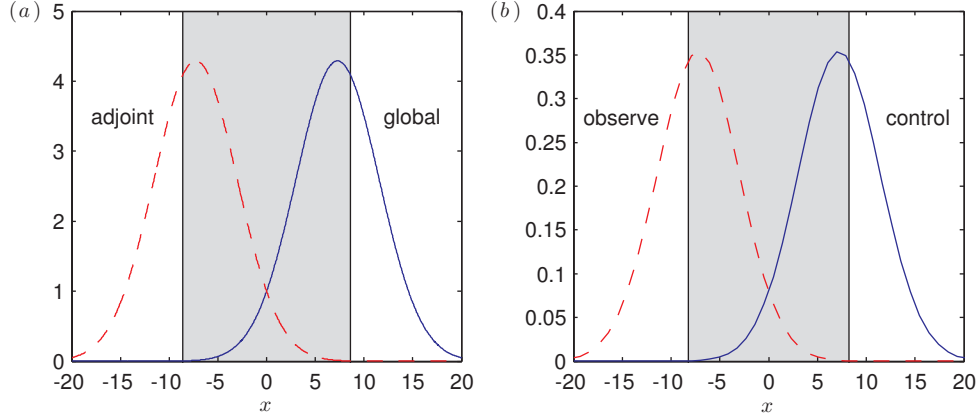


Figure 8.6: (a): Leading global and adjoint eigenmode magnitudes, for $\mu_0 = 1.03\mu_c$. (b): Leading modes of the observability and controllability Gramians (assuming full sensing and actuation), for $\mu_0 = 0.96\mu_c$. The region of amplification is shaded grey.

(2003). As expected, a larger β is more effective at reducing the perturbation magnitude, but at a greater input expense. In practice, one would choose β based on this trade-off. We select $\beta = 7$ because it maintains a small perturbation magnitude using reasonably-sized inputs. For a discussion on the effect of β on the closed-loop eigenvalues of the Ginzburg–Landau system, consult Lauga and Bewley (2004).

We also remark briefly that the controller designed from the linearised model is effective in the full, nonlinear Ginzburg–Landau model. The cubic term in (8.1) always drives the state q towards zero, aiding the controller in doing the same. Brief numerical experiments show, however, that the optimal placement in the nonlinear model does not necessarily correspond with that of the linear model. Further study is necessary to establish optimal placement in that case.

8.5 Comparison with previous approaches

We might naively believe that the optimal actuator and sensor placement could be determined directly from a global mode analysis. In this analysis, we would place sensors where the most unstable eigenmodes are large, so that growing disturbances would register a large output. Similarly, we would place actuators where the most unstable adjoint eigenmodes are large. This is supported by the fact that the LQG controller’s construction requires plant detectability and stabilisability (Doyle et al., 1989). Bagheri et al. (2009) and Åkervik et al. (2007) used this approach for actuator and sensor placement. While this method is plausible on intuitive grounds, it does not produce optimal results. With global and adjoint modes given by equation (8.3), we find for $\mu_0 = 1.03\mu_c$ that there is one unstable mode and one unstable adjoint mode (Bagheri et al., 2009). These two mode magnitudes are bell curves, with peaks respectively at $x = 7.28$ and $x = -7.28$ (see Figure 8.6(a)). If flow disturbances were introduced everywhere, and a sensor and actuator were placed at these two points, then (8.18) would yield $\gamma_2 = 259$. This is far from the optimal $\gamma_2 = 46.1$ achieved by placing the actuator at $x_a = -1.03$ and the sensor at $x_s = 0.98$.

The global mode analysis fails to reveal the optimal placement because the linearised Ginzburg–Landau operator is strongly non-normal. In general, the non-normality prevents global mode analysis from accurately predicting transient growth (Bagheri et al., 2009). In this system, as in Navier–Stokes systems, the disturbances that cause large transient growth are known not to

correspond well with the most unstable eigenmodes (Trefethen et al., 1993). In principle, one may be able to circumvent this issue in a stable system by placing an actuator where an optimal harmonic forcing causes the greatest steady-state solution. The optimal harmonic forcing in the adjoint system could also yield a sensor placement. Bagheri et al. (2009) performed a related analysis in which the initial condition leading to the largest transient growth was sought. The magnitude of this optimal initial condition was a bell curve centred near the upstream boundary of the amplification region, and therefore would not yield the optimal actuator location. Consult Trefethen et al. (1993) and Schmid and Henningson (2000) for more details.

A similar analysis could be done using controllability and observability Gramians of the uncontrolled system with full sensing and actuation. This system is governed by $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ representing

$$\dot{\mathbf{q}} = \mathbf{A}\mathbf{q} + \mathbf{B}\mathbf{d} \quad (8.25a)$$

$$\mathbf{y} = \mathbf{C}\mathbf{q}, \quad (8.25b)$$

where $\mathbf{B} = \mathbf{I}$ and $\mathbf{C} = \mathbf{M}$. The adjoint system is governed by $(\mathbf{A}^+, \mathbf{C}^+, \mathbf{B}^+)$. Because of the uneven grid spacing, the adjoint is not the same as the complex conjugate (Bagheri et al., 2009); rather,

$$\mathbf{A}^+ = \mathbf{M}^{-1}\mathbf{A}^*\mathbf{M} \quad (8.26a)$$

$$\mathbf{B}^+ = \mathbf{B}^*\mathbf{M} \quad (8.26b)$$

$$\mathbf{C}^+ = \mathbf{M}^{-1}\mathbf{C}^*. \quad (8.26c)$$

The controllability Gramian \mathbf{W}_c and the observability Gramian \mathbf{W}_o are the solutions to the Lyapunov equations

$$\mathbf{A}\mathbf{W}_c + \mathbf{W}_c\mathbf{A}^+ + \mathbf{B}\mathbf{B}^+ = \mathbf{0} \quad (8.27a)$$

$$\mathbf{A}^+\mathbf{W}_o + \mathbf{W}_o\mathbf{A} + \mathbf{C}^+\mathbf{C} = \mathbf{0}. \quad (8.27b)$$

In this analysis, we expect the leading eigenvector of \mathbf{W}_c to be large where the state is sensitive to the disturbance \mathbf{d} ; this would appear to be a good place to put a sensor. Similarly, we expect the leading eigenvector of \mathbf{W}_o to be large where the uncontrolled dynamics are most sensitive to perturbations; this would appear to be a good place to put an actuator. These eigenvectors are shown in Figure 8.6(b) with $\mu_0 = 0.96\mu_c$, since Gramians are only well-defined for globally stable systems.

Once again, we find that this analysis does not correctly predict the optimal SISO placement. For this subcritical Ginzburg–Landau system, the optimal placement is $x_a = -0.98$ and $x_s = 0.95$, yielding $\gamma_2 = 42.8$. The Gramian analysis, on the other hand, predicts $x_a = -7.0$ and $x_s = 7.0$, yielding $\gamma_2 = 188$. This analysis fails because Gramians, by construction, decouple actuation from sensing. This is apparent from (8.27). When actuation and sensing are decoupled, the convective nature of information propagation is completely ignored. Placing the sensor so far downstream of the actuator effectively introduces a large time lag in the feedback control.

A method proposed by Giannetti and Luchini (2007) yields a better estimate of the optimal actuator and sensor placement when disturbances are present everywhere. Giannetti and Luchini define the “wavemaker” region as the part of the domain where

$$\zeta(x) = \frac{\phi_0(x)\bar{\psi}_0(x)}{\langle \phi_0, \psi_0 \rangle} \quad (8.28)$$




Figure 8.7: The global and adjoint mode overlap $|\zeta(x)|$ for the unstable mode. The region of amplification is shaded grey.

has a large magnitude, given a global mode ϕ_0 and adjoint mode ψ_0 ; see (8.3). They employed a perturbation technique to show that in this wavemaker region, a change in \mathbf{A} creates the largest possible change in the corresponding eigenvalue λ_0 . As a validation of their theory, Giannetti and Luchini showed good agreement between the wavemaker region of the unstable cylinder wake and the effective placement of passive actuation in Strykowski and Sreenivasan (1990). Although we do not seek to change \mathbf{A} , the wavemaker region nevertheless indicates a region of high dynamical sensitivity (Chomaz, 2005). In such a region, an actuator has a powerful influence on the flow, and localised feedback is effective. Lauga and Bewley (2004) used the wavemaker to determine their actuator and sensor placement, and reported effective \mathcal{H}_∞ control.

As shown in Figure 8.7, the unstable global mode has the largest overlap with its corresponding adjoint mode at $x = 0$ when $\mu_0 = 1.03\mu_c$. Therefore, we propose that placing actuators and sensors near the origin is a sensible initial condition for iterative function minimisation. Although this is a crude approximation for optimal placement, it remains more accurate than global mode and Gramian analyses.

8.6 Conclusion

The linearised Ginzburg–Landau equation is a simple model describing the temporal and spatial evolution of small flow perturbations. In this paper, we control the linearised Ginzburg–Landau equation by establishing a standard LTI state space with some number of actuators and sensors. We perform a perturbation analysis on the continuous algebraic Riccati equations that determine the \mathcal{H}_2 optimal controller. From this, we derive a formulation for the gradient of the controlled system’s \mathcal{H}_2 squared norm. This formulation, which is similar to but more efficient than the formulation of Hiramoto et al. (2000), is applicable to any LQG-controlled system where actuation and sensing matrices are to be optimised. We use conjugate gradient function minimisation to iterate toward the linearised Ginzburg–Landau system’s optimal actuator and sensor placement. This is but one example of how the perturbation theory can be used.

The conjugate gradient results agree with brute force mappings of the \mathcal{H}_2 norm over a large domain. When one actuator and sensor are available, and flow disturbances are introduced everywhere, the optimal placement is $x_a = -1.03$ and $x_s = 0.98$ for the LQG cost matrices chosen. The brute force mappings show that there exist penalties for placing the actuator too far downstream,

the sensor too far upstream, or the actuator too far upstream of the sensor. When an equal number of actuators and sensors is employed, it is optimal to place actuators and sensors in distinct pairs. Each actuator is slightly upstream of the corresponding sensor, and pairs are distributed throughout the region of amplification. We show that the optimal placement is relatively robust to the precise scaling of cost matrices.

When the disturbance is centred at $x_d = -11.0$, the optimal placement is $x_a = -8.48$ and $x_s = -5.55$. Optimal MIMO placements are upstream of the origin, but no clear patterns emerge. Regardless of the way disturbances are introduced, the control system’s performance increases by adding more actuators and sensors, but the increase may not be significant. When optimising the control performance in an experimental setup, for instance, a greater improvement may be achieved by tuning the existing actuator and sensor placements than by simply adding more actuators and sensors.

The process of finding the optimal placement in the linearised Ginzburg–Landau system is completely tractable on a modern computer. Convergence to the global optimum can be difficult, however, when an unequal number of actuators and sensors is used, or when the flow disturbances are only introduced upstream.

Although global mode and Gramian analyses are standard techniques in LTI systems, neither is able to reproduce the optimal actuator and sensor placement. Both roughly suggest $x_a = -7$ and $x_s = 7$, yielding an \mathcal{H}_2 norm four to six times larger than the true minimum when disturbances are present everywhere. The global mode analysis fails because the underlying dynamics are strongly non-normal, and the Gramian analysis fails because it inherently decouples actuation from sensing. On the other hand, we propose that wavemaker analysis is a better way to seek the optimal placement. Although it does not immediately yield the optimal placement, it provides a sensible guess of $x_a = x_s = 0.0$. This is a good initial condition for iterative minimisation.

The methods and results presented may be extended to two- or three-dimensional flow control problems. In such cases, the state size would likely be so large that Riccati and Lyapunov solutions would be computationally intractable. Therefore, each evaluation of the \mathcal{H}_2 norm and its gradient would likely require the construction of a reduced-order model, for instance using balanced proper orthogonal decomposition or the eigensystem realisation algorithm (Ma et al., 2011). Alternatively, iterative methods that bypass the model reduction process may be preferred (Bewley et al., 2011). Potentially faster minimisation algorithms, such as quasi-Newton iteration, may also be substituted (Press et al., 2007). We posit that with such techniques, optimisation of actuator and sensor placement in Navier–Stokes systems should remain computationally tractable using high-performance computing.

This work was supported by the National Science Foundation, grant CMMI-0932928, and K. K. C. received support from the United States Department of Defense (DOD) through the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program. We thank Prof. Tim Colonius, Prof. Paolo Luchini, Shervin Bagheri, and Miloš Ilak for guidance and advising.

Chapter 9

Fluid flow control applications of \mathcal{H}_2 optimal actuator and sensor placement

Kevin K. Chen and Clarence W. Rowley

Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

To appear in the proceedings of the American Control Conference, Portland, OR, June 2014.

In the control of distributed parameter systems, the problem of effective actuator and sensor placement is not well understood, and the use of suboptimal placement methodology is common. We discuss a technique we recently proposed that iterates toward the optimal actuator and sensor placement for \mathcal{H}_2 synthesis, minimizing the closed-loop \mathcal{H}_2 norm from exogenous inputs to exogenous outputs. The iteration is based on an efficient computation of the \mathcal{H}_2 norm's gradient with respect to actuator and sensor placements. We demonstrate the technique on the Orr–Sommerfeld/Squire model of fluid flow, and review previous results from the linearized Ginzburg–Landau model. Finally, we use these results to analyze the validity of previous placement approaches.

9.1 Introduction

Historically, few studies have focused on actuator and sensor placement within distributed parameter systems, which can have as large an impact as controller design. Some work has focused specifically on optimal actuator and sensor placement for flexible structure control; see, for instance, Hiramoto et al. (2000). In vibration dynamics, however, the governing linear operator is typically normal (i.e., it commutes with its adjoint). The techniques in that subset of literature tend not to carry well into non-normal systems, of which fluid flows are a key example. Recent work, such as Morris (2011) and Darivandi et al. (2013), discusses optimal linear quadratic actuator placement in a more general setting, but these works also have limitations; for instance, they treat actuator and sensor placement independently of each other. The latter article is based on a method (Geromel, 1989) that converts the generally non-convex placement problem into a convex one, but only if the actuator location is modeled in a particular discrete way.

In flow control, an early study (Strykowski and Sreenivasan, 1990) placed a small cylinder in the wake of a larger cylinder, and plotted the placements where the passive control stabilized wake




Figure 9.1: Lower linear fractional transform \mathbf{G} .

oscillations. Later, Giannetti and Luchini (2007) discovered that the placements in Strykowski and Sreenivasan (1990) correspond to the so-called “wavemaker” region where the direct and adjoint eigenmodes of the fluid dynamical operator overlap. Similar studies (Roussopoulos, 1993; Lauga and Bewley, 2003) also obtained effective control with actuators and sensors in this general region.

Nonetheless, it is easy to overlook the rigorous methodology and justification of actuator and sensor placement. Two studies employing linear quadratic Gaussian control (Åkervik et al., 2007; Bagheri et al., 2009) each placed an actuator and sensor where the leading adjoint and direct eigenmodes are respectively large. Although stabilizability and detectability are crucial requirements for any control design, eigenmodes generally fail to predict optimal placements when the dynamical operator is non-normal (Chen and Rowley, 2011).

This study discusses a technique for seeking the \mathcal{H}_2 optimal actuator and sensor placement in a systematic way. The \mathcal{H}_2 optimal placement is defined as the actuator and sensor locations in space that—given the \mathcal{H}_2 synthesis—minimizes the \mathcal{H}_2 norm of the closed-loop dynamics from exogenous inputs to exogenous outputs. Therefore, it is the placement that minimizes the average closed-loop response over all linear time-invariant controllers, actuator positions, and sensor positions.

Our computational technique, which Chen and Rowley (2011) introduced and Colburn et al. (2011) improved, computes the gradient of the closed-loop \mathcal{H}_2 norm with respect to actuator and sensor positions. This technique is a more efficient variation of Hiramoto et al. (2000), which focuses on flexible structure control. It is also similar to Lou and Christofides (2003), which uses a more primitive control design. The gradient allows efficient local optimizations of actuator and sensor positions. We review the optimal placement results for the linearized Ginzburg–Landau model of fluid flow from Chen and Rowley (2011), and present new results for the Orr–Sommerfeld/Squire equations.

This paper is organized as follows. Section 9.2 presents the procedure for computing the \mathcal{H}_2 norm’s gradient. Section 9.3 demonstrates the technique on the two fluid flow models, and Section 9.4 compares the results to approaches that others have previously used to justify actuator and sensor placement. Finally, Section 9.5 summarizes the results of the paper.

9.2 Theory

9.2.1 \mathcal{H}_2 synthesis

Since \mathcal{H}_2 synthesis is a standard topic in many texts, we only provide a very concise review. Following the standard matrix notation, we assume that the plant \mathbf{P} in the lower linear fractional transform (Figure 9.1) is

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \mathbf{z} \\ \mathbf{y} \end{bmatrix} = \left[\begin{array}{c|cc} \mathbf{A} & \mathbf{B}_1 & \mathbf{B}_2 \\ \hline \mathbf{C}_1 & \mathbf{0} & \mathbf{D}_{12} \\ \mathbf{C}_2 & \mathbf{D}_{21} & \mathbf{0} \end{array} \right] \begin{bmatrix} \mathbf{q} \\ \mathbf{w} \\ \mathbf{u} \end{bmatrix}, \quad (9.1)$$

and that the matrices in this system obey the standard conditions for \mathcal{H}_2 synthesis, i.e., as specified in Doyle et al. (1989). To avoid rescaling \mathbf{w} and \mathbf{z} , however, we do not require that \mathbf{D}_{12} and \mathbf{D}_{21} be restricted to concatenations of the zero and identity matrices.

To solve the \mathcal{H}_2 synthesis, first solve for the positive-definite matrices \mathbf{X} and \mathbf{Y} in the continuous algebraic Riccati equations (CAREs)

$$\mathbf{A}^H \mathbf{X} + \mathbf{X} \mathbf{A} - \mathbf{X} \mathbf{B}_2 (\mathbf{D}_{12}^H \mathbf{D}_{12})^{-1} \mathbf{B}_2^H \mathbf{X} + \mathbf{C}_1^H \mathbf{C}_1 = \mathbf{0} \quad (9.2a)$$

$$\mathbf{A} \mathbf{Y} + \mathbf{Y} \mathbf{A}^H - \mathbf{Y} \mathbf{C}_2^H (\mathbf{D}_{21} \mathbf{D}_{21}^H)^{-1} \mathbf{C}_2 \mathbf{Y} + \mathbf{B}_1 \mathbf{B}_1^H = \mathbf{0}, \quad (9.2b)$$

where $(\cdot)^H$ denotes the conjugate transpose. With the matrices $\mathbf{F} = (\mathbf{D}_{12}^H \mathbf{D}_{12})^{-1} \mathbf{B}_2^H \mathbf{X}$ and $\mathbf{L} = \mathbf{Y} \mathbf{C}_2^H (\mathbf{D}_{21} \mathbf{D}_{21}^H)^{-1}$, the \mathcal{H}_2 synthesis controller \mathbf{K} is

$$\begin{bmatrix} \dot{\mathbf{z}} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{A} - \mathbf{B}_2 \mathbf{F} - \mathbf{L} \mathbf{C}_2 & \mathbf{L} \\ -\mathbf{F} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{z} \\ \mathbf{y} \end{bmatrix}. \quad (9.3)$$

The squared norm of the closed-loop system in Figure 9.1 is then

$$\|\mathbf{G}\|_2^2 = \text{tr}(\mathbf{C}_1^H \mathbf{C}_1 \mathbf{Y}) + \text{tr}(\mathbf{C}_2 \mathbf{Y} \mathbf{X} \mathbf{L}) \quad (9.4a)$$

$$= \text{tr}(\mathbf{B}_1 \mathbf{B}_1^H \mathbf{X}) + \text{tr}(\mathbf{F} \mathbf{Y} \mathbf{X} \mathbf{B}_2). \quad (9.4b)$$

The objective of this research is to find the actuator and sensor placements that minimize this quantity.

9.2.2 The gradient of the \mathcal{H}_2 norm

This section generalizes and expands our previous work (Chen and Rowley, 2011). We assume that the actuation matrix \mathbf{B}_2 is a differentiable function of r continuous scalar parameters $\{\xi_j\}_{j=1}^r$. Likewise, we assume the same for the sensing matrix \mathbf{C}_2 and p parameters $\{\eta_k\}_{k=1}^p$. The goal of this theory is to compute the derivative of $\|\mathbf{G}\|_2$ with respect to each ξ_j and η_k , to aid in the minimization of (9.4). In the applications that follow, these parameters are the actuator and sensor locations, but this assumption is not necessary in the general theory.

Using $\mathbf{d}(\cdot)_j$ as a shorthand notation for $\partial(\cdot)/\partial \xi_j$, the differentiation of (9.2a) by ξ_j yields the Lyapunov equation

$$(\mathbf{A} - \mathbf{B}_2 \mathbf{F})^H \delta \mathbf{X} + \delta \mathbf{X} (\mathbf{A} - \mathbf{B}_2 \mathbf{F}) = \mathbf{X} \delta \mathbf{B} \mathbf{F} + \mathbf{F}^H \delta \mathbf{B}^H \mathbf{X}. \quad (9.5)$$

This linear equation for the unknown $\delta \mathbf{X}$ is easier to solve than a CARE (9.2). The differentiation of (9.4a) then yields

$$\frac{\partial \|\mathbf{G}\|_2^2}{\partial \xi_j} = \text{tr}(\mathbf{C}_2 \mathbf{Y} \delta \mathbf{X} \mathbf{L}). \quad (9.6)$$

The computation of $\partial \|\mathbf{G}\|_2^2 / \partial \eta_k$ is exactly analogous.

This procedure requires the solution of $r + p$ Lyapunov equations per gradient. A related study (Colburn et al., 2011) illustrated an improvement that reduces this number to two, regardless of r and p . Define the linear operators

$$\mathbf{S}(\delta \mathbf{X}) = (\mathbf{A} - \mathbf{B}_2 \mathbf{F})^H \delta \mathbf{X} + \delta \mathbf{X} (\mathbf{A} - \mathbf{B}_2 \mathbf{F}) \quad (9.7a)$$

$$\mathbf{T}(\delta \mathbf{B}) = \mathbf{X} \delta \mathbf{B} \mathbf{F} + \mathbf{F}^H \delta \mathbf{B}^H \mathbf{X} \quad (9.7b)$$

so that the Lyapunov equation (9.5) is $\mathbf{S}(\delta\mathbf{X}) = \mathbf{T}(\delta\mathbf{B})$. Define the trace inner product $\langle \mathbf{Z}_1, \mathbf{Z}_2 \rangle_{\text{tr}} = \text{tr}(\mathbf{Z}_2^H \mathbf{Z}_1)$, so that the adjoint of \mathbf{S} is

$$\mathbf{S}^*(\phi) = (\mathbf{A} - \mathbf{B}_2 \mathbf{F})\phi + \phi(\mathbf{A} - \mathbf{B}_2 \mathbf{F})^H. \quad (9.8)$$

Thus, if we solve the Lyapunov equation

$$\mathbf{S}^*(\phi) = \mathbf{L} \mathbf{C}_2 \mathbf{Y} \quad (9.9)$$

exactly once for ϕ , then it follows from (9.6) that

$$\frac{\partial \|\mathbf{G}\|_2^2}{\partial \xi_j} = \langle \mathbf{S}^*(\phi), \delta\mathbf{X} \rangle_{\text{tr}} \quad (9.10a)$$

$$= \langle \phi, \mathbf{S}(\delta\mathbf{X}) \rangle_{\text{tr}} \quad (9.10b)$$

$$= \langle \phi, \mathbf{T}(\delta\mathbf{B}) \rangle_{\text{tr}}. \quad (9.10c)$$

Again, the procedure for $\partial \|\mathbf{G}\|_2^2 / \partial \eta_k$ follows the same way.

9.2.3 Remarks

The above gradient formulation provides a tangible improvement over naive finite differencing in the ξ_j - η_k space. It is about $2(r+p)$ times faster (independent of matrix size), and is exact up to roundoff error.

We also remark that in the application of the \mathcal{H}_2 norm gradient to actuator and sensor placements, (9.1) is typically a spatial discretization of a distributed parameter (i.e., partial differential equation, or PDE) system. In this context, Jones and Kerrigan (2010) discuss a theory that determines if the spatial discretization is acceptable for control stability or performance.

Finally, the $\|\mathbf{G}\|_2^2(\{\xi_j\}, \{\eta_k\})$ surface tends to be non-convex in optimal placement applications. In our examples, however, a coarse search of the $\|\mathbf{G}\|_2^2$ space is sufficient for determining a good initial condition for local minimization.

9.3 Examples

9.3.1 Overview

In this section, we use the theory of Section 9.2 to seek the optimal actuator and sensor placements in the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire models of fluid flow. In the discretized state space (9.1), \mathbf{A} is the spatial discretization of the PDE operator \mathcal{L} defining the flow model $\dot{q} = \mathcal{L}q$, and \mathbf{q} is the value of the dependent variable q at discrete nodes. To obtain accurate discrete-space models, we employ pseudospectral differentiation with software based on Weideman and Reddy (2000).

Let us denote the r actuator positions by $\{x_j^a\}_{j=1}^r$ and the p sensor positions by $\{x_k^s\}_{k=1}^p$. In the discretized state space models (9.1), the actuation matrix is $\mathbf{B}_2 = [\mathbf{b}_1 \ \cdots \ \mathbf{b}_r]$, where \mathbf{b}_j represents $\exp(-(x - x_j^a)^2 / (2\sigma^2)) / (\sqrt{2\pi}\sigma)$ at the polynomial nodes, and σ is a given Gaussian width. Since the nodes are unevenly spaced, we also introduce the weighting matrix \mathbf{H} so that $\mathbf{f}_2^H \mathbf{H} \mathbf{f}_1$ is the trapezoidal approximation of the inner product

$$\langle f_1, f_2 \rangle = \int_{\Omega} \bar{f}_2(x) f_1(x) dx \quad (9.11)$$

over the domain Ω , where $\overline{(\cdot)}$ denotes the complex conjugate. We set $\mathbf{C}_2 = [\mathbf{c}_1 \ \cdots \ \mathbf{c}_p]^H \mathbf{H}$, where \mathbf{c}_k represents $\exp(-(x - x_k^s)^2/(2\sigma^2))/(\sqrt{2\pi}\sigma)$ at the polynomial nodes. We pick a small σ to approximate point actuation and sensing.

We assume a standard linear quadratic Gaussian problem. Thus, with a state cost \mathbf{Q} , input cost \mathbf{R} , state disturbance covariance \mathbf{W} , and sensor noise covariance \mathbf{V} ,

$$\mathbf{B}_1 = \begin{bmatrix} \mathbf{W}^{1/2} & \mathbf{0} \end{bmatrix}, \quad \mathbf{C}_1 = \begin{bmatrix} \mathbf{Q}^{1/2} \\ \mathbf{0} \end{bmatrix}, \quad (9.12a)$$

$$\mathbf{D}_{12} = \begin{bmatrix} \mathbf{0} \\ \mathbf{R}^{1/2} \end{bmatrix}, \quad \mathbf{D}_{21} = \begin{bmatrix} \mathbf{0} & \mathbf{V}^{1/2} \end{bmatrix}. \quad (9.12b)$$

Here, we focus primarily on the general effects of external disturbances on the flow behavior. Thus, we choose \mathbf{Q} and \mathbf{W} to be effectively constant in space, and \mathbf{R} and \mathbf{V} to be respectively small compared to \mathbf{Q} and \mathbf{W} . As long as we weight the state cost and state disturbance more heavily than we weight the input cost and sensor noise, the optimal placements in fluid flow models tend to be robust to the actual weights we choose; see Chen and Rowley (2011).

9.3.2 The linearized Ginzburg–Landau system

This section is based on our earlier work (Chen and Rowley, 2011), which contains a more complete discussion.

The linearized Ginzburg–Landau equation models the behavior of small fluid perturbations near a Hopf bifurcation. Letting $x \in \mathbb{R}$ be the primary direction of the fluid flow, and choosing some $\mu(x) : \mathbb{R} \rightarrow \mathbb{R}$, $\nu \in \mathbb{C}$, and $\gamma \in \mathbb{C}$, the linearized Ginzburg–Landau operator is

$$\mathcal{L} = \mu(x) - \nu \frac{\partial}{\partial x} + \gamma \frac{\partial^2}{\partial x^2}. \quad (9.13)$$

The state variable is $q(x, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$, and $\text{Re}(q)$ represents a velocity perturbation in the flow. The boundary conditions are that $q \rightarrow 0$ as $x \rightarrow \pm\infty$. For a comprehensive review of the equation and its control, see Bagheri et al. (2009). We choose $\mu(x) = 0.37 - 5 \cdot 10^{-3}x^2$, $\nu = 2 + 0.4i$ (with i the imaginary unit), and $\gamma = 1 - i$, as in Bagheri et al. (2009) and Chen and Rowley (2011). This system is highly convective in the positive x direction, and the amplification region—where $\mu(x) > 0$ —is $[-8.60, 8.60]$. The system is also strongly non-normal, which will be important in Section 9.4. For this model, $\sigma = 0.4$.

To model \mathcal{L} in discrete space, we compute nodes and differentiation matrices with Hermite polynomials. Letting \mathbf{D}_n denote the n th spatial Hermite pseudospectral derivative on the Hermite nodes x_1, \dots, x_m , the discrete operator is $\mathbf{A} = \text{diag}(\mu) - \nu \mathbf{D}_1 + \gamma \mathbf{D}_2$. Also, we choose $\mathbf{Q}^{1/2} = 7\mathbf{H}^{1/2}$, $\mathbf{R}^{1/2} = \mathbf{I}$, $\mathbf{W}^{1/2} = \mathbf{I}$, and $\mathbf{V}^{1/2} = 2 \cdot 10^{-4}\mathbf{I}$. With $m = 100$ grid points, the leading eigenvalues of \mathbf{A} match the analytic eigenvalues of \mathcal{L} presented in Bagheri et al. (2009), and the optimal placement results are well-converged.

Figure 9.2(a) depicts values of $\|\mathbf{G}\|_2$ in a “brute force” search for the optimal placement by dense sampling, with one actuator and one sensor. The optimal placement is an actuator at $x_a = -1.03$ and a sensor at $x_s = 0.98$, yielding $\|\mathbf{G}\|_2 = 46.1$. Figure 9.2(b) depicts a special case where the flow disturbance is introduced only at an upstream point, $x_d = -11.0$. Here, $\mathbf{W}^{1/2}$ is the $m \times 1$ vector discretizing $\exp(-(x - x_d)^2/(2\sigma^2))/(\sqrt{2\pi}\sigma)$. The optimal placement in this case is $x_a = -8.48$ and $x_s = -5.55$, yielding $\|\mathbf{G}\|_2 = 3.86$.

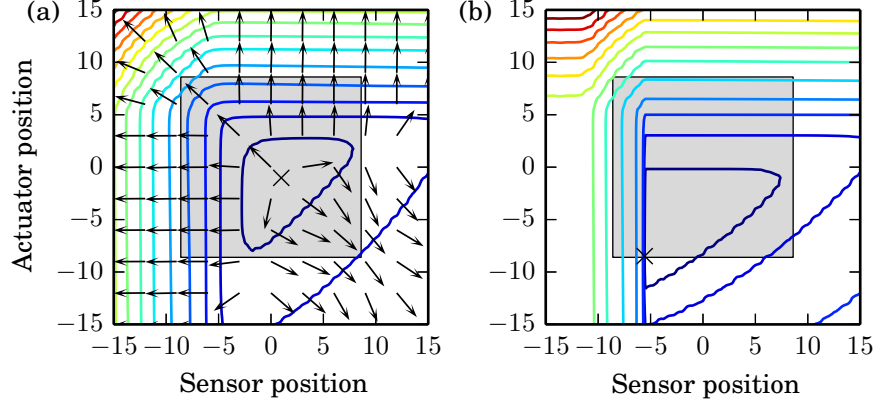


Figure 9.2: (a) Contours of $\log_{10} \|\mathbf{G}\|_2$ and the optimal placement (\times) for the linearized Ginzburg–Landau system. The contour levels are 2, 2.5, The amplification region is shaded gray. Direction vectors of $\nabla \|\mathbf{G}\|_2$ are overlaid as a demonstration. (b) As (a), but with flow disturbances introduced only at $x_d = -11.0$. The contour levels are 1, 1.5,




Figure 9.3: (a) Optimal actuator and sensor placements in the linearized Ginzburg–Landau system. Each row shows the solution for a different number of actuators and sensors. The amplification region is shaded gray. Values of $\|\mathbf{G}\|_2$ are shown on the right. (b) As (a), but with flow disturbances introduced only at $x_d = -11.0$. The top row has two collocated actuators.

In both cases, we point out three patterns. First, placing the actuator too far downstream leads to poor control. Since the Ginzburg–Landau system is highly convective, the actuator would be unable to influence much of the domain. Second, placing the sensor too far upstream also leads to poor control, since the disturbances it would detect would not primarily be the ones undergoing amplification. Third, placing the actuator too far upstream of the sensor is also detrimental, because it would effectively introduce a time lag in \mathbf{P} from \mathbf{u} to \mathbf{y} . With an upstream disturbance (Figure 9.2(b)), the steep increase in $\|\mathbf{G}\|_2$ upstream of the optimal sensor position indicates that the optimal position is very non-robust. This phenomenon occurs because a plant zero crosses into the right-half plane.

Although dense sampling yields insight into the efficacy of placements, it quickly becomes intractable as the number of actuators and sensors increases. Employing the gradient of Section 9.2.2 in a Broyden–Fletcher–Goldfarb–Shanno quasi-Newton algorithm, we compute the optimal placements for varying numbers of actuators and sensors; see Figure 9.3(a).

The resulting pattern is immediately obvious: it is optimal to place actuators and sensors in pairs—with each sensor slightly downstream of its corresponding actuator—and scatter these pairs throughout the amplification region. With the flow disturbance injected only at $x_d = -11.0$ (Figure 9.3(b)), the pattern is less obvious. Although it may seem intuitive to place a single sensor at the disturbance source, this actually leads to very poor control, because small estimation errors would enlarge greatly in the amplification region. Instead, it is more effective to put the single sensor in the amplification region. Only if a second sensor is available should it be located close to the disturbance source. As in Figure 9.3(a), actuators should not be too far upstream of sensors.

9.3.3 The Orr–Sommerfeld/Squire system

The Orr–Sommerfeld and Squire equations are also PDEs in one spatial dimension and in time. In this case, however, the spatial direction is transverse to, not aligned with, the dominant direction of the fluid flow. Let $y \in \mathbb{R}$ be the transverse direction, $\alpha \in \mathbb{R}$ be a streamwise wavenumber, $\beta \in \mathbb{R}$ be a spanwise wavenumber, and $k = \sqrt{\alpha^2 + \beta^2}$. Also, let $u_0(y)$ be the streamwise velocity of a base flow, and Re be the *Reynolds number* relating the velocity, length, density, and viscosity scales of the flow. The Orr–Sommerfeld/Squire operator dictates the vertical velocity perturbation mode amplitude $v(y, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$ and the vertical vorticity (i.e., velocity rotation) perturbation mode amplitude $\omega(y, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$ by

$$\mathcal{L} : \begin{bmatrix} \partial^2/\partial y^2 - k^2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix} \mapsto \begin{bmatrix} \mathcal{L}_{OS} & 0 \\ -i\beta u'_0 & \mathcal{L}_S \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix}, \quad (9.14a)$$

where

$$\mathcal{L}_{OS} = i\alpha \left(u''_0 - u_0 \left(\frac{\partial^2}{\partial y^2} - k^2 \right) \right) + Re^{-1} \left(\frac{\partial^4}{\partial y^4} - 2k^2 \frac{\partial^2}{\partial y^2} + k^4 \right) \quad (9.14b)$$

$$\mathcal{L}_S = -i\alpha u_0 + Re^{-1} \left(\frac{\partial^2}{\partial y^2} - k^2 \right). \quad (9.14c)$$

The Orr–Sommerfeld and Squire systems are respectively the top and bottom rows of (9.14a). At the endpoints of the domain, $v = v' = \omega = 0$. For more information about the equations and their stability analyses, consult Schmid and Henningson (2000).

In this example, our base flow of interest is the “plane Poiseuille” flow through a channel bounded by $y = \pm 1$, but is infinite in the streamwise and spanwise directions. The velocity profile is $u_0(y) = y^2 - 1$, and we choose $\sigma = 0.047$.

We model the Orr–Sommerfeld/Squire systems in discrete space using Chebyshev polynomials on the domain $y \in [-1, 1]$. With \mathbf{D}_n representing the n th Chebyshev pseudospectral derivative matrix, and \mathbf{U}_0 , \mathbf{U}'_0 , and \mathbf{U}''_0 representing $u_0(y)$, $u'_0(y)$, and $u''_0(y)$ on the Chebyshev nodes as diagonal matrices, the discrete operator in (9.1) is

$$\mathbf{A} = \begin{bmatrix} (\mathbf{D}_2 - k^2 \mathbf{I})^{-1} \mathbf{A}_{OS} & \mathbf{0} \\ -i\beta \mathbf{U}'_0 & \mathbf{A}_S \end{bmatrix}, \quad (9.15a)$$

where

$$\mathbf{A}_{OS} = i\alpha (\mathbf{U}''_0 - \mathbf{U}_0 (\mathbf{D}_2 - k^2 \mathbf{I})) + Re^{-1} (\mathbf{D}_4 - 2k^2 \mathbf{D}_2 + k^4 \mathbf{I}) \quad (9.15b)$$

$$\mathbf{A}_S = -i\alpha \mathbf{U}_0 + Re^{-1} (\mathbf{D}_2 - k^2 \mathbf{I}). \quad (9.15c)$$




Figure 9.4: (a) Contours of $\log_{10} \|\mathbf{G}\|_2$ and the optimal placements (+) for the decoupled Orr–Sommerfeld system. The contour levels are 1.55, 1.6, 1.7, 1.8, 2.0, 2.2, 2.6, 3.0. (b) As (a), but for the decoupled Squire system. The contour levels are 0.47, 0.48, \dots , 0.51.

Since the combined system includes both $v(y, t)$ and $\omega(y, t)$, the discrete-space state vector is $\mathbf{q} = [\mathbf{v}^H \ \boldsymbol{\omega}^H]^H$, with \mathbf{v} and $\boldsymbol{\omega}$ respectively representing $v(y, t)$ and $\omega(y, t)$. The actuation matrix \mathbf{B}_2 is similarly vertically stacked, and the sensing matrix \mathbf{C}_2 is horizontally stacked; each actuator and sensor acts on only one of \mathbf{v} and $\boldsymbol{\omega}$. The state penalty matrix is

$$\mathbf{Q}^{1/2} = \begin{bmatrix} \mathbf{H}^{1/2} & \mathbf{0} \\ \mathbf{0} & \mathbf{H}^{1/2} \end{bmatrix}, \quad (9.16)$$

and $\mathbf{R}^{1/2} = \mathbf{I}$, $\mathbf{W}^{1/2} = \mathbf{I}$, and $\mathbf{V}^{1/2} = 0.1\mathbf{I}$. With 125 nodes, the leading eigenvalues of \mathbf{A} match Fig. 3.1 of Schmid and Henningson (2000), and the optimal placements are well-converged.

We consider three sets of parameters in this section. When $Re = 10^4$, $\alpha = 1$, and $\beta = 0$, (9.14) models velocity and vorticity modes that oscillate in the streamwise direction, with no dependence along the span of the flow. In this special case, (9.14a) decouples into the Orr–Sommerfeld (9.14b) and Squire (9.14c) systems. Therefore, we treat the actuator and sensor placements separately for each system.

In the decoupled Orr–Sommerfeld system (Figure 9.4(a)), the optimal placement is $y_a = \pm 0.709$ and $y_s = \mp 0.017$, with $\|\mathbf{G}\|_2 = 31.8$. (Note that the dynamics (9.14) are symmetric about $y = 0$.) The actuator placement $y_a = 0$ is ineffective because it has little or no authority over many of the leading adjoint eigenmodes of $(\mathbf{D}_2 - k^2\mathbf{I})^{-1}\mathbf{A}_{OS}$. In the decoupled Squire system (Figure 9.4(b)), the optimal placement is $y_a = y_s = 0.000$, with $\|\mathbf{G}\|_2 = 2.92$. Simulations of the uncontrolled Squire system show that the point $y = 0$ acts like a “wavemaker,” sending waves toward $y = \pm 1$. Since waves propagate away from the origin, actuator and sensor placements away from this point are ineffective. We also comment that in this decoupled case, the Orr–Sommerfeld system is somewhat non-normal, but the Squire system is very close to normal.

In the next case, we consider $Re = 10^4$, $\alpha = 0$, and $\beta = 1$. This corresponds to “streamwise streaks” that oscillate in the spanwise direction, but are constant in the streamwise direction. Furthermore, this system is strongly non-normal. Because the Orr–Sommerfeld/Squire systems (9.14a) are now coupled, we place one actuator and one sensor in the Orr–Sommerfeld system, and another of each in the Squire system. The optimal placement is $y_a = \pm 0.443$ and $y_s = \mp 0.486$ in the Orr–Sommerfeld domain, and $y_a = \mp 0.485$ and $y_s = \pm 0.481$ in the Squire domain; here, $\|\mathbf{G}\|_2 = 539$.




Figure 9.5: (a) Leading direct (solid blue) and adjoint (dashed red) mode magnitudes of the decoupled Orr–Sommerfeld system, corresponding to the eigenvalue $\lambda = 0.004 - 0.238i$. (b) As (a), but for the decoupled Squire system; $\lambda = -0.007 - 0.993i$. The direct and adjoint parts overlap.

Finally, we consider $Re = 5 \cdot 10^3$ and $\alpha = \beta = 1$. The optimal placement is $y_a = \pm 0.628$ and $y_s = \pm 0.293$ in the Orr–Sommerfeld domain, and $y_a = \mp 0.728$ and $y_s = \mp 0.399$ in the Squire domain; $\|\mathbf{G}\|_2 = 103$ in this case.

9.4 Comparison with previous approaches

9.4.1 Eigenmode analysis

Naively, it may seem natural to place sensors where direct eigenmodes of \mathcal{L} are large, and actuators where the adjoint eigenmodes (as induced by (9.11)) are large. This seems reasonable in the sense that all control requires detectability and stabilizability. This is the approach that Åkervik et al. (2007) and Bagheri et al. (2009) use.

In the decoupled Orr–Sommerfeld/Squire system, where $Re = 10^4$, $\alpha = 1$, and $\beta = 0$, this method appears to be reasonably effective. As shown in Figure 9.5(a), it predicts an Orr–Sommerfeld optimal placement of $y_a = \pm 0.85$ and $y_s = 0$, which is close to the true optimum $y_a = \pm 0.709$ and $y_s = \mp 0.017$. In the Squire system (Figure 9.5(b)), it exactly predicts the true optimum $y_a = y_s = 0$.

Nonetheless, our previous work (Chen and Rowley, 2011) shows that this approach generally leads to suboptimal placements. The most unstable direct and adjoint eigenmodes of the linearized Ginzburg–Landau system are shown in Figure 9.6(a). This naive method predicts an optimal placement of $x_a = -7.28$ and $x_s = 7.28$, which is far from the true optimum $x_a = -1.03$ and $x_s = 0.98$. In the coupled Orr–Sommerfeld/Squire systems, the leading direct and adjoint eigenmodes also poorly predict the optimal placements, but we do not show these modes here.

There are at least two significant reasons for the eigenmode analysis’ typical failure. First, eigenmodes of non-normal operators tend to be nearly parallel. As a result, they generally fail to reveal the transient behavior of the operator (Schmid and Henningson, 2000). Thus, even if control at the eigenmode-based actuator and sensor placements is able to stabilize the dynamics for large times, they may not be able to suppress transient growths to any significant degree. This explains why the eigenmode analysis works well for the almost-normal decoupled Squire system, and somewhat well for the slightly-non-normal decoupled Orr–Sommerfeld system, but fails for the strongly non-normal Ginzburg–Landau and coupled Orr–Sommerfeld/Squire systems.

The second reason for the eigenmode analysis’ failure is that actuator and sensor placement are inherently coupled problems, and we should not expect to be able to solve them independently.




Figure 9.6: (a) As Figure 9.5(a), for the leading mode of the Ginzburg–Landau system, corresponding to the eigenvalue $\lambda = 0.012 - 0.648i$. The amplification region is shaded gray. (b) The corresponding sensitivity magnitude $|\zeta(x)|$.

This is particularly the case in the Ginzburg–Landau system, where x is in the direction of a highly convective flow. The eigenmode analysis places the actuator 14.6 units upstream of the sensor, because the non-normality creates a large separation between direct and adjoint eigenmodes. With an advection velocity of $\text{Re}(\nu) = 2.0$, this essentially creates a large time lag of 7.3 time units between the actuator and the sensor in the plant \mathbf{P} .

9.4.2 Gramian analysis

A Gramian argument similar to the eigenmode analysis is also possible (Ma et al., 2011). Consider a modified plant $\dot{\mathbf{q}} = \mathbf{A}\mathbf{q} + \mathbf{B}\mathbf{d}$, $\mathbf{y} = \mathbf{C}\mathbf{q}$, where \mathbf{A} is stable, $\mathbf{B} = \mathbf{I}$, and $\mathbf{C} = \mathbf{H}$. Essentially, we let external disturbances affect the uncontrolled plant uniformly, and the plant yields an integrated measure of the state size. To construct the Gramians, we use the adjoint matrices \mathbf{A}^* (representing \mathcal{L}^* , the adjoint of \mathcal{L} by (9.11)), along with $\mathbf{B}^* = \mathbf{B}^H \mathbf{H}$ and $\mathbf{C}^* = \mathbf{H}^{-1} \mathbf{C}^H$.

Intuitively, the leading eigenmodes of the controllability Gramian would indicate the regions in space where the state responds most to external disturbances, and this might yield a good sensor placement. Similarly, the leading eigenmodes of the observability Gramian would indicate the regions in space where perturbations in the state affect the integrated output the most, and this might yield a good actuator placement.

We experiment with a stable Ginzburg–Landau system, with $\mu(x) = 0.34 - 5 \cdot 10^{-3}x^2$ and all other parameters unchanged. We also test a stable decoupled Orr–Sommerfeld/Squire system, with $Re = 5 \cdot 10^3$, $\alpha = 1$, and $\beta = 0$. In both cases, the leading Gramian eigenmodes are close to the leading eigenmodes of \mathcal{L} (cf. Figures 9.5, 9.6(a) with unstable parameters). Therefore, Gramian eigenmodes are similarly unable to predict optimal placements. As with Section 9.4.1, the Gramian analysis decouples the actuator and sensor placement problems; it ignores the way actuators feed back into sensors via the plant dynamics.

9.4.3 Sensitivity analysis

We posit that an eigenvalue sensitivity analysis (Giannetti and Luchini, 2007; Schmid and Henningson, 2000, Ch. 3.3.2) may yield a viable heuristic for predicting optimal actuator and sensor positions. If the direct eigendecomposition $\mathcal{L}\phi(x) = \lambda\phi(x)$ and adjoint eigendecomposition $\mathcal{L}^*\psi(x) = \bar{\lambda}\psi(x)$ are given, then an infinitesimal perturbation of \mathcal{L} in the former yields

$$\mathcal{L} d\phi + d\mathcal{L} \phi = \lambda d\phi + d\lambda \phi. \quad (9.17)$$

Taking an inner product (9.11) of the above and ψ , and rearranging the terms, we obtain the eigenvalue sensitivity

$$d\lambda = \frac{\langle d\mathcal{L} \phi, \psi \rangle}{\langle \phi, \psi \rangle}. \quad (9.18a)$$

Let us assume that $d\mathcal{L} = \delta(x - \xi) ds$, i.e., the operator perturbation $d\mathcal{L}$ is a localized change in the dynamics at the point $x = \xi$. For instance, the change could be the application of localized feedback control. Then, we can define the sensitivity distribution

$$\zeta(\xi) = \frac{d\lambda}{ds}(\xi) = \frac{\bar{\psi}(\xi)\phi(\xi)}{\langle \phi, \psi \rangle}. \quad (9.18b)$$

Essentially, we expect localized feedback control to be effective in the “wavemaker”-like region where $\zeta(x)$ is large, and less so where $\zeta(x)$ is small.

Following this analysis, we find that the function $\zeta(x)$ for the Ginzburg–Landau operator’s unstable eigenvalue is largest at $x = 0$; see Figure 9.6(b). Given our foreknowledge that the optimal placement in the Ginzburg–Landau system involves placing actuators and sensors in pairs around $x = 0$, the sensitivity analysis appears to be successful.

The analysis is not as useful, however, on the Orr–Sommerfeld/Squire system. With the exception of the decoupled Squire equation, the optimal actuator and sensor placements are neither collocated nor nearly so. Therefore, we cannot predict the optimal placement using the above formulation. We are currently investigating heuristics that are applicable in this case.

9.5 Conclusion

In this paper, we consider the \mathcal{H}_2 optimal control of a distributed parameter system, where the actuator and sensor locations can be tuned in space. We present an efficient technique for computing the gradient of the closed-loop system’s squared \mathcal{H}_2 norm with respect to actuator and sensor locations. The gradient requires the solution of two Lyapunov equations; therefore, it is easier to solve than the \mathcal{H}_2 synthesis itself.

We then demonstrate the technique by iteratively solving the optimal actuator and sensor placements in the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire equations. The eigenmode and Gramian analyses are able to predict the optimal location only when the dynamics are almost normal. On the other hand, the eigenvalue sensitivity analysis may be able to predict the optimal location in some scenarios when actuators and sensors are collocated, or nearly so.

This work was supported by the National Science Foundation, and the first author was supported by a Department of Defense National Defense Science & Engineering Graduate fellowship.

Chapter 10

Heuristics for effective actuator and sensor placement in feedback flow control

Kevin K. Chen and Clarence W. Rowley

Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

Submitted to the proceedings of Active Flow and Combustion Control, Berlin, Germany, September 2014.

Actuator and sensor placement can be just as consequential for the performance of localized feedback flow control as controller design. Yet, effective placement is not well understood, and the use of suboptimal placements is common. This manuscript reports descriptions and characteristics of effective actuator and sensor placements for optimal flow control. We review \mathcal{H}_2 optimal placements in the linearized Ginzburg–Landau and Orr–Sommerfeld/Squire models of fluid flow. We then analyze the feedback control of these models by relating physical observations with mathematical tools. Although these tools do not fully predict optimal placements, they do reveal patterns that most or all effective placements share. Most notably, effective actuator–sensor placements provide good authority over unstable modes and transient growth, and avoid large time lags between inputs and outputs.

10.1 Introduction

Every application of feedback flow control must address a basic question: where should the actuators and sensors be located in the flow? This placement can impact the control performance just as much as the controller dictating actuation signals from sensor signals. By and large, research in actuator and sensor placement has trailed behind research in controller design. To fill this gap, this manuscript attempts to characterize effective placements for flow control.

A full review of actuator and sensor placement would be too expansive to provide here. Instead, we discuss varying points of view from which different communities have tackled this problem. Control theorists often approach placement theory in ways that are broadly applicable to partial differential equation systems (van de Wal and de Jager, 2001). Much of this research focuses on algorithms for optimizing actuator and sensor locations, but many key aspects are still unsolved. For instance, Morris (2011) and Darivandi et al. (2013) provide algorithms for converging toward

optimal actuator placements, but do not simultaneously solve for optimal sensor placements. The flexible structures community has put forth a sizable body of research (e.g., Hiramoto et al., 2000; Güney and Eşkinat, 2008), but the structures’ governing dynamics are typically described by normal linear operators, and are therefore better behaved than those of fluid flows, for which the linearized dynamics are often highly non-normal. In fluid mechanics, a number of researchers (Roussopoulos, 1993; Gillies, 2001; Lauga and Bewley, 2003; Giannetti and Luchini, 2007; Chen and Rowley, 2011) have corroborated the hypothesis that control is especially effective when actuators and sensors reside in absolutely unstable, or “wavemaker,” regions. Nevertheless, few have described effective actuator and sensor placements in general frameworks, and guess-and-check (e.g., Roussopoulos, 1993; Gillies, 2001; Strykowski and Sreenivasan, 1990) and other suboptimal placement methodologies (e.g., Åkervik et al., 2007; Bagheri et al., 2009) remain common.

The purpose of this manuscript is to provide qualitative characterizations of effective actuator and sensor placements in linear time-invariant (LTI) feedback flow control. We utilize a placement optimization technique described in our previous work (Chen and Rowley, 2011, 2014) to solve for \mathcal{H}_2 optimal placements in the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire models of fluid flow. We then describe features common to these optimal control systems. This analysis is based on mathematical concepts—such as right-half-plane poles and zeros, time lags, eigenmodes, sensitivity to spatially localized perturbations, optimal transient growth and disturbances, and impulse responses—as well as physical intuition.

We organize the manuscript as follows. Section 10.2 reviews \mathcal{H}_2 optimal control and placement. Section 10.3 then reviews the linearized Ginzburg–Landau and Orr–Sommerfeld/Squire models of fluid flow, along with the optimal placements in those systems. Section 10.4 describes common features of optimal actuator and sensor placements, and Section 10.5 provides concluding remarks.

10.2 \mathcal{H}_2 optimal control and placement

The chief goal of the theory in this section is the simultaneous optimization of closed-loop control performance over all actuator positions, sensor positions, and LTI controllers. This section is a concise summary of Chen and Rowley (2011, 2014).

We assume a linear model $\dot{q}(x, t) = \mathcal{L}q(x, t)$ of fluid flow for the state q (e.g., a velocity perturbation), spatial variable or variables x , time t , and dynamical operator \mathcal{L} . If the dynamics are subject to the effects of r actuators with spatial effect $b_j(x)$ and input signals $u_j(t)$, as well as some disturbances $d(x, t)$, then

$$\dot{q}(x, t) = \mathcal{L}q(x, t) + \sum_{j=1}^r b_j(x)u_j(t) + d(x, t). \quad (10.1)$$

Furthermore, if we assume p sensors with spatial form $c_k(x)$, sensor noise $n_k(t)$, and the inner product $\langle q_1(x), q_2(x) \rangle = \int_{\Omega} \bar{q}_2(x)q_1(x) dx$ over the domain Ω (where $\bar{(\cdot)}$ is the complex conjugate), then the sensor outputs $y_k(t)$ are

$$y_k(t) = \langle q(x, t), c_k(x) \rangle + n_k(t) \quad (10.2)$$

for $k = 1, \dots, p$. In this setup, the objective of \mathcal{H}_2 optimal control is to minimize a cost function on q and u_j .

In discrete space, we let $\mathbf{q}(t)$ be a vector representation of $q(x, t)$ over discrete points in the domain. We then bundle the disturbances and sensor noise in the vector $\mathbf{w}(t)$, and we let $\mathbf{z}(t)$ be the vector representation of the \mathcal{H}_2 optimal control cost function. We further let $\mathbf{u}(t)$ and $\mathbf{y}(t)$ be




Figure 10.1: The closed-loop system \mathbf{G} .

vectors respectively stacking the signals $u_j(t)$ and $y_k(t)$. The matrices \mathbf{A} , \mathbf{B}_1 , \mathbf{B}_2 , \mathbf{C}_1 , \mathbf{C}_2 , \mathbf{D}_{12} , and \mathbf{D}_{21} then relate these variables via the “plant” \mathbf{P} , given by

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \mathbf{z} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B}_1 & \mathbf{B}_2 \\ \mathbf{C}_1 & \mathbf{0} & \mathbf{D}_{12} \\ \mathbf{C}_2 & \mathbf{D}_{21} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \mathbf{w} \\ \mathbf{u} \end{bmatrix}. \quad (10.3)$$

We describe the values of these matrices in Section 10.3.

The \mathcal{H}_2 optimal controller is the LTI system \mathbf{K} , arranged as in Figure 10.1, that minimizes the \mathcal{H}_2 norm of the closed-loop system \mathbf{G} . This \mathcal{H}_2 norm $\|\mathbf{G}\|_2$ is essentially an average-case measure of the gain provided by \mathbf{G} from \mathbf{w} to \mathbf{z} . Given a unit white noise input $\mathbf{w}(t)$ and the expected value operator $E(\cdot)$, the norm satisfies $\|\mathbf{G}\|_2^2 = E\left(\lim_{T \rightarrow \infty} \int_{-T}^T \|\mathbf{z}(t)\|_2^2 dt / (2T)\right)$. The computation of \mathbf{K} and $\|\mathbf{G}\|_2$ is a standard procedure (see Skogestad and Postlethwaite, 2005).

Following Chen and Rowley (2011, 2014), we define the \mathcal{H}_2 optimal actuator and sensor locations to be the placements (i.e., the choice of \mathbf{B}_2 and \mathbf{C}_2 , given appropriate constraints) such that the resulting $\|\mathbf{G}\|_2$ is globally minimized. In Chen and Rowley (2011), we proposed an efficient method for computing the gradient of $\|\mathbf{G}\|_2$ with respect to actuator and sensor locations, allowing for gradient-based optimization. This technique is a more efficient variant of Hiramoto et al. (2000), but Chen and Rowley (2014) use adjoint operators to speed up the computation further. In this manuscript, we first sample the entire domain of actuator and sensor placements. Local minima in the sampling then serve as initial conditions for Broyden–Fletcher–Goldfarb–Shanno quasi-Newton minimization; we take the best of these results to be the globally optimal placement.

10.3 Fluid flow models

10.3.1 Overview

For computational tractability, we analyze the optimal actuator and sensor placements in the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire equations. These systems are complementary in the study of shear flows, as they respectively model streamwise and transverse variations in flow perturbations.

In our fluid models, we actuate and sense the native state variables directly. Although this is not particularly realistic, it simplifies the control problem and makes it generalizable. We model point actuators in (10.1) with the Gaussian functions $b_j(x) = \exp((x - x_{a,j})^2 / (2\sigma^2)) / (\sqrt{2\pi}\sigma)$, where $x_{a,j}$ is the j th actuator location, and σ is a small Gaussian width. Denoting the spatial discretization of $b_j(x)$ by \mathbf{b}_j , we set $\mathbf{B}_2 = [\mathbf{b}_1 \ \cdots \ \mathbf{b}_r]$. Similarly, we model point sensors in (10.2) with $c_k(x) = \exp((x - x_{s,k})^2 / (2\sigma^2)) / (\sqrt{2\pi}\sigma)$, where $x_{s,k}$ is the k th sensor location. If the inner product in (10.2) is equivalently $\langle \mathbf{q}_1, \mathbf{q}_2 \rangle = \mathbf{q}_2^H \mathbf{H} \mathbf{q}_1$ —with $(\cdot)^H$ the conjugate transpose,

and with the weighting matrix \mathbf{H} for trapezoidal integration—and \mathbf{c}_k is the discretization of $c_k(x)$, then $\mathbf{C}_2 = [\mathbf{c}_1 \ \cdots \ \mathbf{c}_p]^H \mathbf{H}$.

In our \mathcal{H}_2 optimal control setup, we assume a standard linear quadratic Gaussian (LQG) problem. Thus, the objective of the control is to minimize $E(\mathbf{q}^H \mathbf{Q} \mathbf{q} + \mathbf{u}^H \mathbf{R} \mathbf{u})$, given the disturbances and sensor noise in (10.1, 10.2) and appropriate choices of the state cost matrix $\mathbf{Q} \geq 0$ (i.e., positive semidefinite) and the input cost matrix $\mathbf{R} > 0$ (i.e., positive definite). We pick \mathbf{Q} such that $\mathbf{q}^H \mathbf{Q} \mathbf{q}$ is the kinetic energy of the state. Furthermore, if $\mathbf{d}(t)$ is the discretization of the disturbance $d(x, t)$ in (10.1), and the sensor noise $n_k(t)$ (see (10.2)) is stacked in the vector $\mathbf{n}(t)$, then we define the disturbance covariance $\mathbf{W} = E(\mathbf{d} \mathbf{d}^H) \geq 0$ and sensor noise covariance $\mathbf{V} = E(\mathbf{n} \mathbf{n}^H) > 0$. We pick \mathbf{W} so that the disturbance is effectively a unit white noise uncorrelated in space and time. Finally, in the plant (10.3), we assign $\mathbf{w} = [\mathbf{d}^H \ \mathbf{n}^H]^H$ and

$$\mathbf{B}_1 = [\mathbf{W}^{1/2} \ \mathbf{0}], \quad \mathbf{C}_1 = \begin{bmatrix} \mathbf{Q}^{1/2} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{D}_{12} = \begin{bmatrix} \mathbf{0} \\ \mathbf{R}^{1/2} \end{bmatrix}, \quad \mathbf{D}_{21} = [\mathbf{0} \ \mathbf{V}^{1/2}]. \quad (10.4)$$

The following sections describe the choice of \mathcal{L} in (10.1)—and consequently, \mathbf{A} in (10.3)—as well as \mathbf{Q} , \mathbf{R} , \mathbf{W} , and \mathbf{V} for the fluid flow models. They also report the optimal actuator and sensor locations in these systems. We employ pseudospectral differentiation with software based on Weideman and Reddy (2000).

10.3.2 Linearized Ginzburg–Landau equation

This section is a brief summary of Chen and Rowley (2011, 2014). The linearized Ginzburg–Landau equation models small fluid perturbations near a Hopf bifurcation (Cossu and Chomaz, 1997; Chomaz, 2005). Letting $x \in \mathbb{R}$ be the primary direction of the fluid flow, and choosing some $\mu(x) : \mathbb{R} \rightarrow \mathbb{R}$, $\nu \in \mathbb{C}$, and $\gamma \in \mathbb{C}$, the linearized Ginzburg–Landau operator is

$$\mathcal{L} = \mu(x) - \nu \frac{\partial}{\partial x} + \gamma \frac{\partial^2}{\partial x^2}. \quad (10.5)$$

The state variable is $q(x, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$, and $\text{Re}(q)$ represents a velocity perturbation in the flow. The boundary conditions are that $\lim_{x \rightarrow \pm\infty} q = 0$. For a comprehensive review of the equation and its control, see Bagheri et al. (2009). We choose $\mu(x) = \mu_0 - 0.04 - 5 \cdot 10^{-3} x^2$, $\mu_0 = 0.41$, $\nu = 2 + 0.4i$ (with i the imaginary unit), and $\gamma = 1 - i$, as in Bagheri et al. (2009), Chen and Rowley (2011), and Chen and Rowley (2014). This system is unstable, strongly non-normal, and highly convective in the positive x direction. The amplification region—where $\mu(x) > 0$ —is $[-8.60, 8.60]$, and we choose $\sigma = 0.4$.

To compute \mathbf{A} from \mathcal{L} , we employ pseudospectral differentiation with Hermite polynomials. Also, we choose $\mathbf{Q} = 49\mathbf{H}$, $\mathbf{R} = \mathbf{I}$, $\mathbf{W} = \mathbf{I}$, and $\mathbf{V} = 4 \cdot 10^{-8}\mathbf{I}$, thereby emphasizing the state disturbance and state cost. With 100 grid points, the leading eigenvalues of \mathbf{A} match the analytic eigenvalues of \mathcal{L} presented in Bagheri et al. (2009), and the optimal placement results are converged.

We employ the methodology of Chen and Rowley (2014) to solve for the \mathcal{H}_2 optimal actuator and sensor positions, with varying numbers of each; see Figure 10.2 and Figure 10.3(a). Figure 10.3(b) depicts the optimal positions for a system where the flow disturbance $d(x, t)$ is a single-channel white noise given over the shape $\exp((x - x_d)^2 / (2\sigma^2)) / (\sqrt{2\pi}\sigma)$, with the disturbance location $x_d = -11$.




Figure 10.2: Snapshots of $\text{Re}(q)$ (solid) and $\text{Im}(q)$ (dashed) in a neutrally stable linearized Ginzburg–Landau simulation ($\mu_0 = 0.398$) with white noise disturbances. In this and all following plots, the amplification region is shaded gray. (a) Without control. (b) With sensor noise and LQG control at the optimal actuator (\circ) and sensor placement (\times), shown below.




Figure 10.3: Optimal actuator and sensor placements in the linearized Ginzburg–Landau system. Each row shows the solution for a different number of actuators and sensors, and values of $\|\mathbf{G}\|_2$ are shown on the right. (a) Flow disturbances everywhere. (b) Flow disturbances introduced only at $x_d = -11.0$. The top row has two collocated actuators.

10.3.3 Orr–Sommerfeld/Squire equations

In shear flows, the Orr–Sommerfeld and Squire equations model transverse variations in exponential mode amplitudes (Schmid and Henningson, 2000). Let $y \in \mathbb{R}$ be the transverse spatial variable, $\alpha, \beta \in \mathbb{R}$ be streamwise and spanwise wavenumbers, and $k = \sqrt{\alpha^2 + \beta^2}$. Also, let $u_0(y)$ be the streamwise velocity of a base flow, and Re be the Reynolds number. With the vertical velocity and vorticity perturbation mode amplitudes $v(y, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$ and $\omega(y, t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$, we define

$$\mathcal{L}_{\text{OS}} = i\alpha \left(u_0'' - u_0 \left(\frac{\partial^2}{\partial y^2} - k^2 \right) \right) + Re^{-1} \left(\frac{\partial^4}{\partial y^4} - 2k^2 \frac{\partial^2}{\partial y^2} + k^4 \right) \quad (10.6a)$$

$$\mathcal{L}_{\text{S}} = -i\alpha u_0 + Re^{-1} \left(\frac{\partial^2}{\partial y^2} - k^2 \right), \quad (10.6b)$$

so that

$$\mathcal{L} : \begin{bmatrix} \frac{\partial^2}{\partial y^2} - k^2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix} \mapsto \begin{bmatrix} \mathcal{L}_{\text{OS}} & 0 \\ -i\beta u_0' & \mathcal{L}_{\text{S}} \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix}. \quad (10.6c)$$




Figure 10.4: Snapshots of $\text{Re}(\omega)$ (solid) and $\text{Im}(\omega)$ (dashed) in an Orr–Sommerfeld/Squire simulation with white noise disturbances, using $Re = 10^4$, $\alpha = 0$, and $\beta = 1$; v (not shown) is very small. (a) Without control. (b) With sensor noise and LQG control at the optimal placement (\circ : v actuator; \times : v sensor; \square : ω actuator; \star : ω sensor), shown below.

The Orr–Sommerfeld and Squire systems are respectively the top and bottom rows of (10.6c). At the endpoints of the domain, $v = v' = \omega = 0$. Our base flow is the plane Poiseuille flow through a channel bounded by $y = \pm 1$. Thus, $u_0(y) = y^2 - 1$, and we choose $\sigma = 0.047$.

We compute \mathbf{A} from \mathcal{L} using Chebyshev pseudospectral differentiation. The discrete-space state vector is $\mathbf{q} = [\mathbf{v}^H \quad \boldsymbol{\omega}^H]^H$, with \mathbf{v} and $\boldsymbol{\omega}$ representing $v(y, t)$ and $\omega(y, t)$. Similarly, we vertically stack \mathbf{B}_2 and horizontally stack \mathbf{C}_2 ; each actuator and sensor acts on only one of \mathbf{v} and $\boldsymbol{\omega}$. To penalize the kinetic energy of the flow perturbation in the LQG design, we set $\mathbf{Q} = \text{diag}(\mathbf{D}^H \mathbf{H} \mathbf{D} + k^2 \mathbf{H}, \mathbf{H}) / (2k^2)$, where \mathbf{D} is the Chebyshev differentiation matrix (Schmid and Henningson, 2000). In addition, we assume that the flow disturbances are an uncorrelated unit white noise in all three components of velocity; thus, $\mathbf{W} = \text{diag}(\mathbf{I}, k^2 \mathbf{I})$ (Schmid and Henningson, 2000). Finally, $\mathbf{R} = \mathbf{I}$ and $\mathbf{V} = 0.01 \mathbf{I}$. With 125 nodes, the leading eigenvalues of \mathbf{A} match Fig. 3.1 of Schmid and Henningson (2000), and the optimal placements are converged.

We consider three sets of parameters in this section. For the spanwise waves given by $Re = 10^4$, $\alpha = 1$, and $\beta = 0$, (10.6c) decouples into the Orr–Sommerfeld and Squire systems; we assume single-input, single-output (SISO) control in each. We also consider the streamwise streaks given by $Re = 10^4$, $\alpha = 0$, and $\beta = 1$ (see Figure 10.4), as well as the oblique waves given by $Re = 5 \cdot 10^3$ and $\alpha = \beta = 1$. In these cases, we consider two-input, two-output (2I2O) control, where the locations of the v -actuator and sensor are y_a^v and y_s^v , and the locations of the ω -actuator and sensor are y_a^ω and y_s^ω . The off-diagonal term in (10.6c) causes small changes in v to have large effects on ω . Therefore, even when $|v| \ll |\omega|$ (e.g., Figure 10.4), the v -actuator and sensor placement is still relevant.

Figure 10.5 shows the globally optimal actuator and sensor placements, as well as placements that are locally, and nearly globally, optimal. We also remark that this Orr–Sommerfeld/Squire system is symmetric about $y = 0$; for each optimal placement, we do not show the equivalent solution reflected about $y = 0$.

10.4 Observations and heuristics

In this section, we discuss various considerations, observations, and traits we observe in feedback flow control systems with effective actuator and sensor placements. These topics are not disjoint,




Figure 10.5: As Figure 10.3, but for the Orr–Sommerfeld/Squire system. (a) The decoupled SISO systems with $Re = 10^4$, $\alpha = 1$, and $\beta = 0$. (b) The coupled 2I2O system with $Re = 10^4$, $\alpha = 0$, and $\beta = 1$, showing both globally (top row) and locally optimal solutions. (c) As (b), but with $Re = 5 \cdot 10^3$ and $\alpha = \beta = 1$.

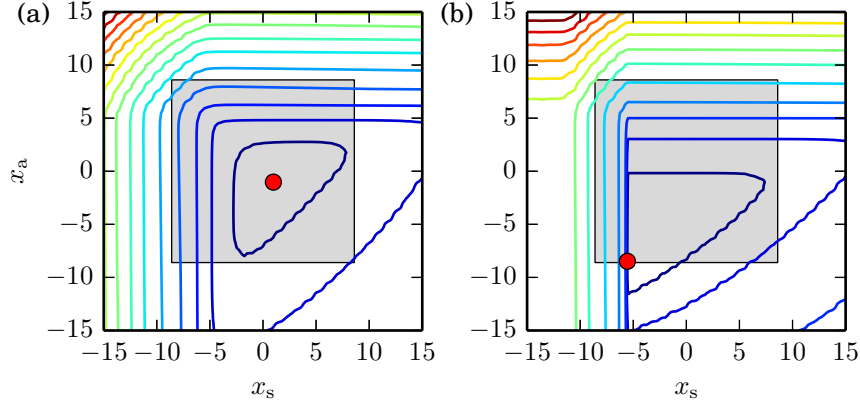


Figure 10.6: Optimal placement (●) and $\log_{10} \|\mathbf{G}\|_2$ contours in the linearized Ginzburg–Landau system, plotted against the actuator placement x_a and the sensor placement x_s . (a) With spatially uncorrelated flow disturbances; the contour levels are 2, 2.5, (b) Disturbances only at $x_d = -11.0$; the contour levels are 1, 1.5,

but are rather related in various ways. Each section within briefly describes a particular feature of interest.

10.4.1 Fundamental requirements

In rigorous control theory, it is possible to identify and characterize plants for which good control is never possible with any LTI controller. This section briefly summarizes ideas from Skogestad and Postlethwaite (2005).

Two basic characteristics that all plants must possess for closed-loop stability are *stabilizability* and *detectability*. That is, the actuators must be able to act on all unstable adjoint eigenmodes of \mathcal{L} , and the sensors must be able to observe all unstable direct eigenmodes. Comparing Figure 10.6(a) and Figure 10.7(a), for instance, we observe that the linearized Ginzburg–Landau plant is stabilizable and detectable at the optimal actuator and sensor placement. If the sensor is very far upstream or the actuator is very far downstream, however, they lose authority over the unstable modes, leading to high values of $\|\mathbf{G}\|_2$, and therefore, poor control. In the decoupled




Figure 10.7: (a) Magnitudes of leading direct (solid magenta) and adjoint (dashed green) modes of the linearized Ginzburg-Landau system; $\lambda = 0.012 - 0.648j$. (b) The corresponding sensitivity magnitude $|\zeta(x)|$.

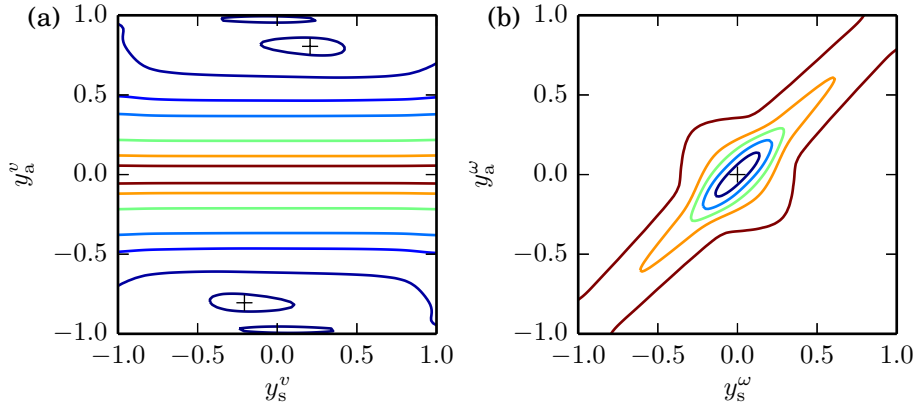


Figure 10.8: (a) The optimal placements (+) and $\log_{10} \|\mathbf{G}\|_2$ contours in the decoupled Orr-Sommerfeld system, with $Re = 10^4$, $\alpha = 1$, and $\beta = 0$. The contour levels are 2.04, 2.1, 2.3, 2.5, 3.0, 3.5, 4.0. (b) As (a), but with contours of $\|\mathbf{G}\|_2$ in the decoupled Squire system. The contour levels are 2.10, 2.15, \dots , 2.30.




Figure 10.9: Magnitudes of leading direct (solid magenta) and adjoint (dashed green) modes. (a) The decoupled Orr-Sommerfeld system; $\lambda = 0.004 - 0.238i$. (b) The decoupled Squire system; $\lambda = -0.007 - 0.993i$. The direct and adjoint eigenmodes overlap.

Orr-Sommerfeld and Squire systems, the optimal placement in Figure 10.8 has good actuator and sensor authority over the unstable (or least stable) modes in Figure 10.9. Other placements lead to poorer modal controllability and observability.

Plant poles and zeros can also limit the efficacy of LTI control. Given the transfer function $\mathbf{P}(s) = \mathbf{C}_2(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}_2$ of a plant (where we do not assume any direct feedthrough from \mathbf{u} to \mathbf{y}), the poles and zeros are the values of $s \in \mathbb{C}$ for which $\mathbf{P}(s)$ diverges or is rank-deficient, respectively. For good control to be possible, plant poles p and zeros z with positive real parts should satisfy $|p| < |z|/4$ (Skogestad and Postlethwaite, 2005). In Figure 10.6(b), a small upstream perturbation in the optimal sensor placement causes $\|\mathbf{G}\|_2$ to rise extremely quickly. This likely occurs because a plant zero crosses into the complex right half-plane, violating the pole-zero condition. Control designers should place actuators and sensors so as to avoid small right-half-plane zeros, but this discussion is beyond the scope of this manuscript.

Finally, if a system has right-half-plane poles p as well as time delays τ from actuators to sensors, then $\tau < 1/(2|p|)$ is generally required for effective control (Skogestad and Postlethwaite, 2005). Although our fluid models do not strictly contain time delays, the Ginzburg–Landau model does exhibit time lags, since convection predominantly drives the dynamics. Figure 10.6 confirms that when the actuator is far upstream of the sensor—leading to large time lags—the control performance decreases. In this system, it is generally preferable for the two to be nearly collocated, as Figure 10.3 shows.

10.4.2 Eigenmode-based analyses

Direct and adjoint eigenmodes

As a direct extension of the stabilizability and detectability discussion in Section 10.4.1, it seems natural to place actuators and sensors where the most unstable adjoint and direct eigenmodes are respectively large. This approach is common (Åkervik et al., 2007; Bagheri et al., 2009), and it works well for the decoupled Orr–Sommerfeld and Squire equations at $Re = 10^4$, $\alpha = 1$, and $\beta = 0$. Once again comparing Figure 10.8 and Figure 10.9, we confirm that the modal analysis exactly predicts the optimal vorticity placement of $y_a^\omega = y_s^\omega = 0$. Observing the peaks Figure 10.9(a), we also predict an optimal placement of $y_a^v = \pm 0.848$ and $y_s^v = 0$, yielding $\|\mathbf{G}\|_2 = 109.6$. This is not far from the true optimum¹ $y_a^v = \pm 0.789$ and $y_s^v = \pm 0.246$, with $\|\mathbf{G}\|_2 = 109.2$.

We must emphasize, however, that placement by eigenmode analysis is generally suboptimal in fluid systems. It is well established (Schmid and Henningson, 2000) that the eigenmodes of non-normal operators do not individually predict transient growth. Thus, even if a control system can prevent the long-time blow-up of unstable modes, it may be very poor at controlling transient growth. The linearized Ginzburg–Landau system’s eigenmodes predict an optimal SISO placement of $x_a = -7.28$ and $x_s = 7.28$, with $\|\mathbf{G}\| = 259$ (Figure 10.7(a)), but the true optimum is $x_a = -1.03$ and $x_s = 0.98$, with $\|\mathbf{G}\|_2 = 46.1$ (Figure 10.6(a)). The analysis also fails in the coupled Orr–Sommerfeld/Squire equations, where the comparatively large off-diagonal term in \mathcal{L} (10.6c) causes strong non-normality. The general failure of this approach necessitates alternate analyses, which we describe below.

Sensitivity analysis

The sensitivity analysis (Giannetti and Luchini, 2007; Chomaz, 2005) determines how much a certain eigenvalue of \mathcal{L} will move in the complex plane if \mathcal{L} experiences a point perturbation. This is related to the ideas of absolute instability and wavemakers in local flow stability theory. Given a perturbation strength ds and location $x = \xi$, and denoting the Dirac delta function by δ , we

¹The computation reports that $y_a^v = \pm 0.970$ and $y_s^v = \mp 0.046$ is the global optimum, with the slightly lower $\|\mathbf{G}\|_2 = 109.1$ (see Figure 10.8(a)). This placement appears to violate the boundary conditions considerably, however, so we reject this solution.

suppose that $d\mathcal{L} = \delta(x - \xi) ds$. Furthermore, we consider the eigendecomposition $\mathcal{L}\phi_j = \lambda_j\phi_j$, as well as the eigendecomposition $\mathcal{L}^*\psi_j = \bar{\lambda}_j\psi_j$ of the adjoint operator \mathcal{L}^* . The sensitivity function is then

$$\zeta_j(\xi) = \frac{d\lambda_j}{ds}(\xi) = \frac{\bar{\psi}_j(\xi)\phi_j(\xi)}{\langle\phi_j, \psi_j\rangle}. \quad (10.7)$$

This analysis is particularly useful with collocated actuators and sensors, where the feedback control exactly perturbs \mathcal{L} in our prescribed way. Therefore, the parts of the spatial domain where $\zeta_j(x)$ is large are the regions where feedback control can have a considerably favorable effect on the eigenvalues of the closed-loop system. In the linearized Ginzburg–Landau SISO system, we predict $x_a = x_s = 0$, with $\|\mathbf{G}\|_2 = 54.2$ (see Figure 10.7(b)). This is close to the true optimum $x_a = -1.03$ and $x_s = 0.98$, with $\|\mathbf{G}\|_2 = 46.1$ (see Figure 10.3(a) and Figure 10.6(a)).

The sensitivity analysis fails to predict the optimal placements in the Orr–Sommerfeld/Squire system, however. Figure 10.5 reveals that these optimal placements generally do not involve nearly collocating actuators and sensors, and this sensitivity analysis is therefore not directly applicable. It is necessary to consider alternate means of evaluating effective placements in non-normal systems, which the next section discusses.

10.4.3 Transient growth

Optimal growth and disturbances

The purpose of optimal growth theory (see Schmid and Henningson, 2000) is to analyze non-normal transient growth in ways that the eigenmode analysis (Section 10.4.2) cannot. The optimal energy growth $g(t)$ dictates the largest energy amplification that a dynamical operator \mathbf{A} (or equivalently, \mathcal{L}) can provide out of all possible initial conditions. If—given some \mathbf{M} — $\|\mathbf{q}\|_{\mathbf{E}} = \|\mathbf{M}^{1/2}\mathbf{q}\|_2$ is a norm whose square is the state’s energy, then

$$g(t) = \max_{\mathbf{q}(0) \neq \mathbf{0}} \frac{\|\mathbf{q}(t)\|_{\mathbf{E}}^2}{\|\mathbf{q}(0)\|_{\mathbf{E}}^2} = \|\mathbf{M}^{1/2}e^{\mathbf{A}t}\mathbf{M}^{-1/2}\|_2^2. \quad (10.8)$$

For the largest singular value of $\mathbf{M}^{1/2}e^{\mathbf{A}t}\mathbf{M}^{-1/2}$, the corresponding right singular vector is the optimal initial disturbance giving rise to the optimal growth, and the corresponding left singular vector is that disturbance at time t .

The optimal initial disturbances and the evolved disturbances hint at effective actuator and sensor placements, respectively. The reason for this is intuitive: the optimal initial disturbance indicates regions where actuation can have the largest short-term effect on the state, and the evolved disturbance represents the transient growth that the sensor can most easily detect. This is in contrast to the eigenmode analysis (Section 10.4.2), which only considers the infinite time horizon.

The choice of t , however, is not clear-cut. In the Orr–Sommerfeld/Squire system in Figure 10.10, it suffices to examine the time t that maximizes $g(t)$. Defining a characteristic time $t_c = 1/\min_j |\lambda_j|$, we use the scaled time $\hat{t} = t/t_c$ in Figure 10.10(a) and find that the optimal growth is largest at $\hat{t} = 5.31$. Observing the peaks in the optimal initial disturbance (Figure 10.10(b)), this predicts the optimal placement $y_a^v = 0.77$ or -0.77 , and $y_a^\omega = 0.70$ or -0.70 . Similarly, Figure 10.10(c) predicts $y_s^v = 0$ and $y_s^\omega = 0.75$ or -0.75 . The placement $y_a^v = \pm 0.77$, $y_a^\omega = \mp 0.70$, $y_s^v = 0$, and $y_s^\omega = \pm 0.75$ yields $\|\mathbf{G}\|_2 = 78.0$, which is close to the optimal $\|\mathbf{G}\|_2 = 72.2$ at $y_a^v = \pm 0.648$, $y_a^\omega = \mp 0.725$, $y_s^v = \pm 0.667$, and $y_s^\omega = \pm 0.361$ (Figure 10.5(c)).




Figure 10.10: The Orr–Sommerfeld/Squire system’s optimal growth with $Re = 5 \cdot 10^3$ and $\alpha = \beta = 1$. (a) The optimal energy growth $g(\hat{t})$. (b) The optimal initial disturbance for the optimal $\hat{t} = 5.31$, shown as $|v|$ (solid blue) and $|\omega|$ (dashed red). (c) The disturbance at $\hat{t} = 5.31$.




Figure 10.11: The linearized Ginzburg–Landau system’s optimal growth. (a) The optimal energy growth $g(\hat{t})$. (b) The optimal initial disturbance for $\hat{t} = 1$, shown as real (dashed) and imaginary (dotted) parts, and magnitude (solid). (c) The disturbance at $\hat{t} = 1$.

Our example linearized Ginzburg–Landau system is unstable, however, implying that $g(\hat{t})$ must diverge (Figure 10.11(a)). Nonetheless, we can still examine the optimal initial and evolved disturbance for the arbitrarily chosen $\hat{t} = 1$. The optimal initial disturbance (Figure 10.11(b)) predicts $x_a = -1.55$, and the disturbance at $\hat{t} = 1$ (Figure 10.11(c)) predicts $x_s = 1.55$, yielding $\|\mathbf{G}\|_2 = 48.3$. This is very close to the true optimum $x_a = -1.03$ and $x_s = 0.98$, with $\|\mathbf{G}\|_2 = 46.1$.

Impulse responses

Since we assume a linear framework in (10.1, 10.2), impulse responses from actuators and adjoint impulse responses from sensors are effectively Green’s functions. For non-normal systems, the impulse responses reveal each actuator’s ability to affect the state via the transient growth of its impulse; the adjoint impulses also reveal each sensor’s ability to detect transient dynamics. The impulse responses are given by (10.1, 10.2) with $u_j(t) = \delta(t)$ for one value of j , and $d(x, t)$, $n_k(t)$, and all other $u_j(t)$ equal to zero. To compute the adjoint impulse responses, we use the adjoint operator \mathcal{L}^* instead of \mathcal{L} , and we swap $c_k(x)$ and $b_j(x)$, as well as $u_j(t)$ and $y_k(t)$, with the impulse now on $y_k(t)$.

The impulse responses do not directly predict optimal actuator and sensor placements, but they do reveal why certain placements may be effective or ineffective. In the Orr–Sommerfeld/Squire system shown in Figure 10.12(a), we find that the impulse from the v -actuator at the optimal




Figure 10.12: Impulse responses of the Orr–Sommerfeld/Squire system with $Re = 10^4$, $\alpha = 0$, and $\beta = 1$ at the optimal placement, showing $\hat{t} = 0, 0.04, \dots, 0.2$. (a) Impulse in v from $y_a^v = 0.877$ (solid blue: $Re(v)$; dashed red: $Im(\omega)$). (b) Impulse in ω from $y_a^\omega = -0.488$ ($Re(\omega)$ shown). (c) Adjoint impulse in v from $y_s^v = -0.487$ ($Re(v)$ shown). (d) Adjoint impulse in ω from $y_s^\omega = 0.483$ (dashed blue: $Im(v)$; solid red: $Re(\omega)$).




Figure 10.13: Output traces in Figure 10.12. (a) $Im(\omega)$ sensor output at $y_s^\omega = 0.483$ from the actuator impulse at $y_a^v = 0.877$ (see Figure 10.12(a)). (b) $Im(v)$ adjoint actuator output at $y_a^v = 0.877$ from the adjoint sensor impulse at $y_s^\omega = 0.483$ (see Figure 10.12(d)).

$y_a^v = 0.877$ creates an extremely large transient growth in ω , as a result of the off-diagonal term in (10.6c). This implies that the actuator has significant authority over ω , using very little effort. Furthermore, the sensor at the optimal $y_s^\omega = 0.483$ registers a very large output from this impulse (Figure 10.13(a)), indicating strong feedback. Although the impulse from the ω -actuator at $y_a^\omega = -0.488$ (Figure 10.12(b)) does not appear to have a considerable effect, it does act in the $y < 0$ domain, where the v -actuator does not have as much authority over ω as in $y > 0$. We remark that neither actuator has an appreciable effect on ω at $y \approx 0$, but the underlying dynamics naturally

suppress ω there. Finally, we reach similar conclusions from the adjoint impulse responses at $y_s^v = -0.487$ and $y_s^\omega = 0.483$; see Figure 10.12(c, d) and Figure 10.13(b).

We also briefly mention that the input–output impulse responses in Figure 10.13 exhibit a nearly nonexistent time lag. This is a favorable characteristic in light of the fundamental requirements that Section 10.4.1 described.

10.4.4 Remarks

Physical intuition

Many of the conclusions we drew with mathematical tools are also possible using physical intuition. We highlight a few examples here.

First, Section 10.4.1 explained effective and ineffective placements in the linearized Ginzburg–Landau system (Figure 10.6) using stabilizability, detectability, and time lags. We intuit that in a convection-driven system, downstream actuators cannot influence much of the domain, and upstream sensors cannot detect the most relevant disturbances. Furthermore, sensors far downstream of actuators cannot detect actuator-induced growth until the growth has evolved and amplified for many convective time units. These are all recognizably unfavorable conditions.

Next, Section 10.4.2 and Section 10.4.3 justified the optimal placement in the decoupled Orr–Sommerfeld and Squire systems with $Re = 10^4$, $\alpha = 1$, and $\beta = 0$ using eigenmodes and optimal growth theory. Simulations with flow disturbances show that the dynamics propagate v waves from near the boundaries to the center of the domain, and ω waves from the center to the boundaries. This intuition supports the optimal placements shown in Figure 10.5(a). Section 10.4.3 also investigated the optimal placement in the Orr–Sommerfeld/Squire system with $Re = 10^4$, $\alpha = 0$, and $\beta = 1$ using impulse responses. Simulations (see Figure 10.4(a)) show that ω tends to be largest roughly near $y = \pm 0.5$; this intuitively supports the placement of actuators and sensor near these regions, as shown in Figure 10.5(b).

Other remarks

We briefly remark that in some optimal placements, certain actuators and sensors primarily exhibit feedforward behavior. This is the case, for instance, with the ω -actuator and the v -sensor in the globally optimal solution in Figure 10.5(b) and Figure 10.12. This is not necessarily detrimental, as Belson et al. (2013) describe. These feedforward actuators and sensors have authority over physical regions where other actuators and sensors do not, thereby improving the control performance.

Finally, we point out that it is common for the topography of $\|\mathbf{G}\|_2$ with respect to certain actuator and sensor placements to be relatively flat; see, for instance, the sensor placements in Figure 10.5(c). This affords the control designer additional flexibility in effective hardware placement.

10.5 Conclusion

This manuscript described several characteristics of effective actuator and sensor positions for the feedback control of fluid models. The \mathcal{H}_2 optimal control of the linearized Ginzburg–Landau and the Orr–Sommerfeld/Squire models served as test beds for this investigation.

Successful feedback control always requires actuator and sensor positions that allow the plant to stabilize and detect unstable eigenmodes. In addition, placements that cause small right-half-plane zeros, or large time lags or delays, typically lead to poor control. The placement of actuators and sensors using leading eigenmodes can be effective, but only if the dynamics are close to normal. For

non-normal dynamics with nearly collocated actuators and sensors, the sensitivity region better predicts advantageous placements. Optimal growth and disturbances, on the other hand, can predict optimal placements in non-normal systems without assuming collocation. Impulse responses also reveal the individual actuators' and sensors' authority over the state. Finally, physical intuition can often predict the same results that these mathematical tools reveal.

There currently does not exist a single, unifying method that is universally able to predict optimal actuator and sensor positions. The heuristics and observations in this manuscript provide some knowledge that may be useful in approaching this goal, however. The prediction of optimal placements remains a topic of future research.

This work was supported by the National Science Foundation's Graduate Research Fellowship Program, and grant CMMI-0932928.

Chapter 11

Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses

Kevin K. Chen, Jonathan H. Tu, and Clarence W. Rowley

Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

Appears in *Journal of Nonlinear Science* 22(6):887–915, December 2012.

Digital Object Identifier [doi:10.1007/s00332-012-9130-9](https://doi.org/10.1007/s00332-012-9130-9)

Dynamic mode decomposition (DMD) is an Arnoldi-like method based on the Koopman operator. It analyzes empirical data, typically generated by nonlinear dynamics, and computes eigenvalues and eigenmodes of an approximate linear model. Without explicit knowledge of the dynamical operator, it extracts frequencies, growth rates, and spatial structures for each mode. We show that expansion in DMD modes is unique under certain conditions. When constructing mode-based reduced-order models of partial differential equations, subtracting a mean from the data set is typically necessary to satisfy boundary conditions. Subtracting the mean of the data exactly reduces DMD to the temporal discrete Fourier transform (DFT); this is restrictive and generally undesirable. On the other hand, subtracting an equilibrium point generally preserves the DMD spectrum and modes. Next, we introduce an “optimized” DMD that computes an arbitrary number of dynamical modes from a data set. Compared to DMD, optimized DMD is superior at calculating physically relevant frequencies, and is less numerically sensitive. We test these decomposition methods on data from a two-dimensional cylinder fluid flow at a Reynolds number of 60. Time-varying modes computed from the DMD variants yield low projection errors.

11.1 Introduction

Scientists and engineers frequently employ data-based modal decomposition to investigate complex dynamical systems. When dynamical operators are too sophisticated to analyze directly—because of nonlinearity and high dimensionality, for instance—data-based techniques are often more practical. Such methods numerically analyze empirical data produced by experiments or simulations, yielding modes containing dynamically significant structures. In fluid flow control, for example, decompositions of an evolving flow field may identify coherent structures such as vortices and ed-

dies that are fundamental to the underlying flow physics (Holmes et al., 1996). A subset of the computed modes can then form the basis of a reduced-order model—that is, a low-dimensional approximation of the original dynamical system—using Galerkin projection. This is often critical in applications such as controller or observer design in a control system, where matrix operations on a high-dimensional dynamical system could be computationally intractable.

One common decomposition, proper orthogonal decomposition (POD, otherwise known as principal component analysis or Karhunen–Loève decomposition), generates orthogonal modes that optimally capture the vector energy of a given data set (Holmes et al., 1996). This method, while popular, suffers from a number of known issues. For instance, reduced-order models generated by POD may be inaccurate, because the principal directions in a set of data may not necessarily correspond with the dynamically important ones. As a result, the selection of POD modes to retain in a reduced-order model is nontrivial and potentially difficult (Ilak and Rowley, 2008). Improved methods are available, such as balanced truncation (Moore, 1981) or balanced POD (an approximation of balanced truncation for high-dimensional systems; Rowley, 2005), but both of these methods are applicable only to linear systems.

Dynamic mode decomposition (DMD) is a relatively recent development in the field of modal decomposition (Rowley et al., 2009; Schmid, 2010). Dynamic mode decomposition approximates the modes of the Koopman operator, which is a linear, infinite-dimensional operator that represents nonlinear, finite-dimensional dynamics without linearization (Mezić and Banaszuk, 2004; Mezić, 2005), and is the adjoint of the Perron-Frobenius operator. The method can be viewed as computing, from empirical data, eigenvalues and eigenvectors of a linear model that approximates the underlying dynamics, even if those dynamics are nonlinear. Unlike POD and balanced POD, this decomposition yields growth rates and frequencies associated with each mode, which can be found from the magnitude and phase of each corresponding eigenvalue. If a linear dynamical operator generates the data, then the decomposition recovers the leading eigenvalues and eigenvectors of that operator; if the data are periodic, then the decomposition is equivalent to a temporal discrete Fourier transform (DFT) (Rowley et al., 2009). Applications of DMD to experimental and numerical data can be found in Rowley et al. (2009), Schmid (2010, 2011), and Schmid et al. (2011).

In this paper, we present some new properties and some modifications of DMD, highlighting the relationship with traditional Fourier analysis. Section 11.2 briefly reviews the Koopman operator, DMD, and two algorithms for computing DMD. In Section 11.3, we prove the uniqueness of the decomposition under specific conditions. Section 11.4 discusses applications to reduced-order models of partial differential equations, and in particular compares methods to ensure that models based on DMD satisfy boundary conditions correctly. We explore the subtraction of either the data mean or an equilibrium point (commonly known as a “base flow” in fluid mechanics), and examine the consequences of each. In Section 11.5, we present advances toward an alternate, “optimized” DMD formulation, which tailors the decomposition to the user-specified number of modes, and is more accurate than truncated DMD. We provide examples in Section 11.6 by considering numerically generated data of a cylinder fluid flow, evolving via the nonlinear Navier–Stokes equations from an unstable equilibrium to a limit cycle.

11.2 The Koopman operator and dynamic mode decomposition

In this section, we briefly review the results of Rowley et al. (2009). The *Koopman operator* U is defined for a dynamical system

$$\boldsymbol{\xi}_{k+1} = \mathbf{f}(\boldsymbol{\xi}_k) \quad (11.1)$$

evolving on a finite-dimensional manifold M . (The dynamics need not be discrete in time, but we retain this form since we presume that the data to be analyzed are such.) It acts on scalar functions $g : M \rightarrow \mathbb{R}$ or \mathbb{C} according to

$$Ug(\boldsymbol{\xi}) \triangleq g(\mathbf{f}(\boldsymbol{\xi})). \quad (11.2)$$

Although the underlying dynamics are nonlinear and finite-dimensional, the Koopman operator is a linear infinite-dimensional operator; the linearity is a result of the definition in (11.2), not linearization. Given the eigendecomposition

$$U\phi_j(\boldsymbol{\xi}) = \lambda_j\phi_j(\boldsymbol{\xi}), \quad j = 1, 2, \dots \quad (11.3)$$

of U , we can generally express vector-valued “observables” $\mathbf{g} : M \rightarrow \mathbb{R}^n$ or \mathbb{C}^n of our choice in terms of the Koopman eigenfunctions ϕ_j by

$$\mathbf{g}(\boldsymbol{\xi}) = \sum_{j=1}^{\infty} \phi_j(\boldsymbol{\xi}) \mathbf{v}_j, \quad (11.4)$$

where $\{\mathbf{v}_j\}_{j=1}^{\infty}$ is a set of vector coefficients called *Koopman modes*. (Here, we assume that each component of \mathbf{g} lies within the span of the eigenfunctions.) Using (11.3) and the definition in (11.2), we can explicitly write

$$\mathbf{g}(\boldsymbol{\xi}_k) = \sum_{j=1}^{\infty} \lambda_j^k \phi_j(\boldsymbol{\xi}_0) \mathbf{v}_j. \quad (11.5)$$

The Koopman eigenvalues $\{\lambda_j\}_{j=1}^{\infty}$ therefore dictate the growth rate and frequency of each mode.¹ For simplicity, we will hereafter incorporate the constant $\phi_j(\boldsymbol{\xi}_0)$ into \mathbf{v}_j .

It is straightforward to show that if the dynamics (11.1) are linear, with $\mathbf{f}(\boldsymbol{\xi}) = \mathbf{A}\boldsymbol{\xi}$, then the eigenvalues λ_j of \mathbf{A} are also eigenvalues of the Koopman operator. Furthermore, if the observable is $\mathbf{g}(\boldsymbol{\xi}) = \boldsymbol{\xi}$, then the Koopman modes \mathbf{v}_j are the corresponding eigenvectors of \mathbf{A} (Rowley et al., 2009).

The practical idea behind the Koopman analysis is to collect a set of data $\{\boldsymbol{\xi}_k\}$, identify an observable \mathbf{g} of interest from the data, and express the observable in terms of Koopman modes and eigenvalues. For example, an experimenter may wish to collect fluid flow data from a wind tunnel experiment, measuring three-component velocity fields in a spatial region of interest, and decompose the fields into eigenvalues and modes containing spatial structures.

The DMD algorithm approximates Koopman modes and eigenvalues from a finite set of data. The algorithm itself, described in Rowley et al. (2009) and Schmid (2010), is a variant of a standard Arnoldi method, and we summarize it here. Suppose that the vector observable is $\mathbf{g}(\boldsymbol{\xi}_k) = \mathbf{x}_k$. The algorithm operates on the real or complex data set $\{\mathbf{x}_k\}_{k=0}^m$ and identifies complex Ritz values $\{\lambda_j\}_{j=1}^m$ and complex Ritz vectors $\{\mathbf{v}_j\}_{j=1}^m$ such that

$$\mathbf{x}_k = \sum_{j=1}^m \lambda_j^k \mathbf{v}_j, \quad k = 0, \dots, m-1 \quad (11.6a)$$

$$\mathbf{x}_m = \sum_{j=1}^m \lambda_j^m \mathbf{v}_j + \mathbf{r}, \quad \mathbf{r} \perp \text{span}(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}), \quad (11.6b)$$

¹In ergodic theory, we assume the map \mathbf{f} is measure preserving, making the Koopman operator unitary (Petersen, 1983). Here, we make no such assumption, allowing for the analysis of dynamics not lying on an attractor. As such, the growth rates given by $|\lambda_j|$ may differ from unity.

provided that $\{\lambda_j\}_{j=1}^m$ are distinct. This algorithm approximates the data—which a nonlinear process may have potentially generated—as the trajectory of a linear system. The first m data vectors are exactly represented by the Ritz values and vectors. Projecting the final data vector, however, results in a residual \mathbf{r} that is orthogonal to all previous data vectors, and hence to all Ritz vectors as well.

The computation of DMD modes proceeds as follows. Let $\mathbf{K} \triangleq [\mathbf{x}_0 \ \cdots \ \mathbf{x}_{m-1}]$ be a data matrix. Let $\mathbf{c} = [c_0 \ \cdots \ c_{m-1}]^T$ be a vector of coefficients that best constructs (in a least squares sense) the final data vector \mathbf{x}_m as a linear combination of all previous data vectors. That is,

$$\mathbf{x}_m = \mathbf{K}\mathbf{c} + \mathbf{r}, \quad \mathbf{r} \perp \text{span}(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}). \quad (11.7)$$

One solution for \mathbf{c} is

$$\mathbf{c} = \mathbf{K}^+ \mathbf{x}_m, \quad (11.8)$$

where $(\cdot)^+$ indicates the Moore-Penrose pseudoinverse. The solution for (11.7) is unique if and only if \mathbf{K} has linearly independent columns; in this case, we can use the pseudoinverse expansion $\mathbf{K}^+ = (\mathbf{K}^* \mathbf{K})^{-1} \mathbf{K}^*$. Now construct the companion matrix

$$\mathbf{C} \triangleq \begin{bmatrix} 0 & 0 & \cdots & 0 & c_0 \\ 1 & 0 & \cdots & 0 & c_1 \\ 0 & 1 & & 0 & c_2 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & & 1 & c_{m-1} \end{bmatrix}. \quad (11.9)$$

One possible diagonalization of this matrix is

$$\mathbf{C} = \mathbf{T}^{-1} \mathbf{\Lambda} \mathbf{T}, \quad (11.10)$$

where $\mathbf{\Lambda}$ is a diagonal matrix with $\lambda_1, \dots, \lambda_m$ on the diagonal, \mathbf{T} is the Vandermonde matrix defined by $T_{ij} \triangleq \lambda_i^{j-1}$, and $\{\lambda_j\}_{j=1}^m$ are distinct. The matrix $\mathbf{\Lambda}$ yields the Ritz values. Finally, the matrix $\mathbf{V} \triangleq [\mathbf{v}_1 \ \cdots \ \mathbf{v}_m]$ contains the Ritz vectors, where

$$\mathbf{V} = \mathbf{K} \mathbf{T}^{-1}. \quad (11.11)$$

To understand how this algorithm yields the decomposition in (11.6), consider that we could write an index-shifted data matrix $\mathbf{K}_* \triangleq [\mathbf{x}_1 \ \cdots \ \mathbf{x}_m]$ as

$$\mathbf{K}_* = \mathbf{K} \mathbf{C} + \mathbf{r} \mathbf{e}^T \quad (11.12a)$$

$$= \mathbf{K} \mathbf{T}^{-1} \mathbf{\Lambda} \mathbf{T} + \mathbf{r} \mathbf{e}^T, \quad (11.12b)$$

if $\mathbf{e} \triangleq [0 \ \cdots \ 0 \ 1]^T$. Finally, using (11.11), we obtain

$$\mathbf{K} = \mathbf{V} \mathbf{T} \quad (11.13a)$$

$$\mathbf{K}_* = \mathbf{V} \mathbf{\Lambda} \mathbf{T} + \mathbf{r} \mathbf{e}^T, \quad (11.13b)$$

which together are equivalent to (11.6).

The algorithm above is analytically correct, but may be ill-conditioned in practice. This may especially be the case when using noisy experimental data. Schmid (2010) recommends the following alternate algorithm, which is exact when \mathbf{K} is full-rank. Let $\mathbf{K} = \mathbf{U}\mathbf{\Sigma}\mathbf{W}^*$ be the economy-sized singular value decomposition (SVD) of \mathbf{K} . Solve the diagonalization problem

$$\mathbf{U}^*\mathbf{K}_*\mathbf{W}\mathbf{\Sigma}^{-1} = \mathbf{Y}\mathbf{\Lambda}\mathbf{Y}^{-1} \quad (11.14)$$

and set $\mathbf{V} = \mathbf{U}\mathbf{Y}$; $\mathbf{\Lambda}$ is as above. Finally, scale each column of \mathbf{V} by the appropriate complex scalar so that (11.6a) is satisfied.

For data sets with $\dim(\mathbf{x}_k) \gg m$, we recommend a method of snapshots approach. In fluid mechanics, for instance, \mathbf{K} may (very roughly) have a row dimension $\mathcal{O}(10^5)$ to $\mathcal{O}(10^{10})$ and a column dimension $\mathcal{O}(10^1)$ to $\mathcal{O}(10^3)$. In this case, $\mathbf{K}^*\mathbf{K}$ is much smaller than \mathbf{K} . The method of snapshots requires less memory, and may be faster to compute. Evaluate the matrix $\mathbf{K}^*\mathbf{K}$, and perform the diagonalization

$$\mathbf{K}^*\mathbf{K} = \mathbf{W}\mathbf{\Sigma}^2\mathbf{W}^*. \quad (11.15)$$

The diagonal matrix $\mathbf{\Sigma}$ with positive entries and the unitary matrix \mathbf{W} are as above. Complete the algorithm of Schmid above, starting at (11.14), with the substitution $\mathbf{U} = \mathbf{K}\mathbf{W}\mathbf{\Sigma}^{-1}$.

A key advantage of the SVD and method of snapshots approaches comes to light when \mathbf{K} is rank-deficient, or nearly so. In this case, we can approximate (11.6) with a smaller number of modes by truncating the parts of \mathbf{U} , $\mathbf{\Sigma}$, and \mathbf{W} corresponding to small singular values (Schmid, 2010).

11.3 Uniqueness of dynamic mode decomposition

Rowley et al. (2009) proved by construction that for a set of real or complex data $\{\mathbf{x}_k\}_{k=0}^m$, there exist complex $\{\lambda_j\}_{j=1}^m$ and $\{\mathbf{v}_j\}_{j=1}^m$ such that (11.6) is satisfied, provided that $\{\lambda_j\}_{j=1}^m$ are distinct. Here, we extend this result by showing that the decomposition is unique.

Theorem 11.1. *The choice of $\{\lambda_j\}_{j=1}^m$ and $\{\mathbf{v}_j\}_{j=1}^m$ in (11.6) is unique up to a reordering in j , if and only if $\{\mathbf{x}_k\}_{k=0}^{m-1}$ are linearly independent and $\{\lambda_j\}_{j=1}^m$ are distinct.*

Proof. We begin with the backward proof. Since λ_j are distinct, the Vandermonde matrix \mathbf{T} is full-rank. Recall that (11.6a) is equivalent to $\mathbf{K} = \mathbf{V}\mathbf{T}$, and that we assume \mathbf{K} is full-rank. It follows that \mathbf{V} is full-rank, and that $\{\mathbf{x}_k\}_{k=0}^{m-1}$ and $\{\mathbf{v}_j\}_{j=1}^m$ must each be linearly independent bases of the same space. Equations (11.6b, 11.7) therefore imply that \mathbf{x}_m can be decomposed into two additive parts,

$$\sum_{j=1}^m \lambda_j^m \mathbf{v}_j = \mathbf{K}\mathbf{c} \in \text{span}(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}) \quad (11.16)$$

and $\mathbf{r} \perp \text{span}(\mathbf{x}_0, \dots, \mathbf{x}_{m-1})$.

Let us rewrite (11.16) as

$$\mathbf{V} \begin{bmatrix} \lambda_1^m \\ \vdots \\ \lambda_m^m \end{bmatrix} = \mathbf{K}\mathbf{c} \quad (11.17a)$$

$$= \mathbf{V}\mathbf{T}\mathbf{c}. \quad (11.17b)$$

Given that the columns of \mathbf{V} are linearly independent, this then reduces to

$$\begin{bmatrix} \lambda_1^m \\ \vdots \\ \lambda_m^m \end{bmatrix} = \mathbf{T}\mathbf{c}, \quad (11.18)$$

from which we extract m scalar equations

$$\begin{aligned} \lambda_1^m &= c_0 + c_1\lambda_1 + c_2\lambda_1^2 + \cdots + c_{m-1}\lambda_1^{m-1} \\ &\vdots \\ \lambda_m^m &= c_0 + c_1\lambda_m + c_2\lambda_m^2 + \cdots + c_{m-1}\lambda_m^{m-1}. \end{aligned} \quad (11.19)$$

Therefore, the m distinct values of λ_j are precisely the m roots of the polynomial

$$p(\lambda; \mathbf{c}) \triangleq \lambda^m - \sum_{k=0}^{m-1} c_k \lambda^k. \quad (11.20)$$

(By extension, we may also write $p(\lambda; \mathbf{c}) = \prod_{j=1}^m (\lambda - \lambda_j)$; this will be used later.) Note from (11.8) that for linearly independent $\{\mathbf{x}_k\}_{k=0}^{m-1}$, $\{c_k\}_{k=0}^{m-1}$ exists and is unique. Thus, the m roots $\{\lambda_j\}_{j=1}^m$ of $p(\lambda; \mathbf{c})$ must be unique as well, up to a reordering in j . Finally, the vectors $\{\mathbf{v}_j\}_{j=1}^m$ must be unique, as a result of (11.11).

Next, we show that if $\{\mathbf{x}_k\}_{k=0}^{m-1}$ are linearly dependent or $\{\lambda_j\}_{j=1}^m$ are not distinct, then the choice of $\{\lambda_j\}_{j=1}^m$ and $\{\mathbf{v}_j\}_{j=1}^m$ is not unique, even beyond a reordering in j . The latter condition is straightforward to prove—suppose that $\lambda_q = \lambda_r$ for some $q \neq r$. Then for any arbitrary $\mathbf{v}_* \in \mathbb{C}^n$, \mathbf{v}_q may be replaced with $\mathbf{v}_q + \mathbf{v}_*$, and \mathbf{v}_r with $\mathbf{v}_r - \mathbf{v}_*$, and (11.6) will still hold.

Suppose now that $\{\mathbf{x}_k\}_{k=0}^{m-1}$ are linearly dependent, so $\text{rank}(\mathbf{K}) < m$. We maintain the condition that the λ_j be distinct; otherwise we refer back to the statement just proven. Thus, the Vandermonde matrix \mathbf{T} is full-rank, so $\text{rank}(\mathbf{V}) = \text{rank}(\mathbf{K}) < m$; furthermore, $\text{span}(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}) = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_m)$.

The remainder of this proof follows similarly to the backward proof. One difference, however, is that \mathbf{c} must be *non-unique* because \mathbf{K} does not have full column rank. To be precise, $\dim(\ker(\mathbf{K})) > 0$; if some \mathbf{c} satisfies (11.7), then $\mathbf{c} + \mathbf{c}_*$ does as well, for any $\mathbf{c}_* \in \ker(\mathbf{K})$. In addition, (11.18, 11.19) are now sufficient but not necessary for (11.17). In other words, the m solutions to the polynomial equation $p(\lambda; \mathbf{c}) = 0$, along with $\mathbf{V} = \mathbf{K}\mathbf{T}^{-1}$, are sufficient for satisfying (11.6).

Since \mathbf{c} is not unique, there must exist different coefficient sets $\{c_k\}_{k=0}^{m-1}$ and $\{d_k\}_{k=0}^{m-1}$ such that

$$\lambda^m - \sum_{k=0}^{m-1} c_k \lambda^k = \prod_{j=1}^m (\lambda - \lambda_j) = 0 \quad (11.21a)$$

$$\mu^m - \sum_{k=0}^{m-1} d_k \mu^k = \prod_{j=1}^m (\mu - \mu_j) = 0 \quad (11.21b)$$

produce roots $\{\lambda_j\}_{j=1}^m$ and $\{\mu_j\}_{j=1}^m$ each satisfying (11.6). If $\{\lambda_j\}_{j=1}^m$ and $\{\mu_j\}_{j=1}^m$ were identical, then according to the above relations, $\{c_k\}_{k=0}^{m-1}$ and $\{d_k\}_{k=0}^{m-1}$ must also be identical. Since the coefficient sets are different, however, there must exist distinct sets $\{\lambda_j\}_{j=1}^m$ and $\{\mu_j\}_{j=1}^m$ satisfying (11.6). \square

Remark. The polynomial $p(\lambda; \mathbf{c})$, defined in (11.20), is precisely the characteristic polynomial of the companion matrix in (11.9) used to construct the DMD. Thus, as in (11.10), the scalars $\{\lambda_j\}_{j=1}^m$ are the eigenvalues of \mathbf{C} .

11.4 Boundary condition handling

11.4.1 Overview

An important application of modal decompositions, such as POD or DMD, is to obtain low-dimensional approximations of partial differential equations. When used in this setting, the governing dynamics (11.1) become infinite dimensional. Nevertheless, we often use approximate numerical solutions (e.g., from finite-difference or spectral methods) in practice, and can treat these as systems with large, but finite, dimension. We may then use modal decompositions, such as POD and DMD, to obtain further approximations of surprisingly low dimension, by expanding the solution in terms of a linear combination of modes (Holmes et al., 1996).

Of course, partial differential equations differ fundamentally from initial value ordinary differential equations and maps in that they require boundary conditions. Subtleties of imposing boundary conditions on such low-dimensional models are frequently overlooked, for instance, as Noack et al. (2005) point out; in many cases, however, the situation is straightforward. For instance, in a fluid flow, the boundaries of the domain are often at a physical boundary such as a wall. There, the velocity is zero, and a simple Dirichlet boundary condition is appropriate. Other cases such as inflow or outflow conditions can be more difficult, but in practice, a Dirichlet or Neumann boundary condition often suffices for constructing low-dimensional models. When we use POD or DMD modes in these situations, we must pay careful attention to ensure that the low-dimensional models satisfy the boundary conditions, at least approximately.

In this section, we consider the case where the dynamics (11.1) represent an approximation of a partial differential equation, discretized in both time and space, and the observable $\mathbf{x} = \mathbf{g}(\xi)$ consists of the dependent variables in some spatial region Ω . This is a common situation in fluid flows, since we typically know flow information only in a limited region of space, whether we obtain the data from experiments or numerical simulations. We consider boundary conditions of the form

$$B(\mathbf{x}) = \gamma, \quad (11.22)$$

where B maps the observed variables \mathbf{x} to certain values on the boundary $\partial\Omega$. Suppose the boundary condition also satisfies the linearity condition

$$B(\alpha_1 \mathbf{x}_1 + \alpha_2 \mathbf{x}_2) = \alpha_1 B(\mathbf{x}_1) + \alpha_2 B(\mathbf{x}_2). \quad (11.23)$$

Given a set of modes $\{\mathbf{v}_j\}_{j=1}^m$, we may construct a low-dimensional approximation of the solution by expanding \mathbf{x} as a linear combination of modes. Using $\mathbf{x}(k)$ to denote \mathbf{x}_k from the previous section, we have the expansion

$$\mathbf{x}(k) = \mathbf{x}_b + \sum_{j=1}^m a_j(k) \mathbf{v}_j, \quad (11.24)$$

where the constant offset \mathbf{x}_b is commonly called a “base flow” in the context of fluid mechanics. It is then convenient to require that the base flow itself satisfy the boundary conditions, i.e.,

$$B(\mathbf{x}_b) = \gamma. \quad (11.25)$$

Equations (11.22–11.25) imply that for a set $\{\mathbf{x}(k)\}_{k=0}^m$ of observations each satisfying (11.22),

$$B(\mathbf{x}(k) - \mathbf{x}_b) = 0, \quad k = 0, \dots, m. \quad (11.26)$$

That is, if each solution $\mathbf{x}(k)$ satisfies the boundary conditions, then the subtraction of the base flow from each solution creates modified data with homogeneous boundary conditions. As a result, modes decomposed from the base-flow-subtracted solutions also acquire homogeneous boundary conditions.

Therefore, suppose we construct a set of modes $\{\mathbf{v}_j\}_{j=1}^m$ from base-flow-subtracted data; then as long as the base flow \mathbf{x}_b satisfies (11.25), the expansion (11.24) will also satisfy the boundary conditions (11.22), for any choice of $a_j(k)$. When employing DMD modes computed *without* first subtracting a base flow, it is generally not possible to satisfy boundary conditions. The modes will typically have nonzero boundary conditions, and any nonzero values along those mode boundaries will be scaled by $a_j(k)$.

In the application of POD, it is common to use the mean of a data set as the base flow (e.g., Noack et al., 2003). We show in the following sections, however, that the use of the data mean as the base flow exactly reduces DMD to the temporal DFT, and that this is typically undesirable. On the other hand, using an equilibrium point as the base flow leads to better-behaved decompositions.

11.4.2 Mean subtraction

The data mean we describe here is

$$\bar{\mathbf{x}} \triangleq \frac{1}{m+1} \sum_{k=0}^m \mathbf{x}_k, \quad (11.27a)$$

and the mean-subtracted data are

$$\mathbf{x}'_k \triangleq \mathbf{x}_k - \bar{\mathbf{x}}, \quad k = 0, \dots, m. \quad (11.27b)$$

We note briefly that the mean must include the vector \mathbf{x}_m . Otherwise, the mean-subtracted data would be equal to the rank-deficient matrix

$$\begin{bmatrix} \mathbf{x}_0 - \frac{1}{m} \sum_{k=0}^{m-1} \mathbf{x}_k & \cdots & \mathbf{x}_{m-1} - \frac{1}{m} \sum_{k=0}^{m-1} \mathbf{x}_k \end{bmatrix}; \quad (11.28)$$

according to Section 11.3, DMD would not be unique. Therefore, let $\mathbf{K}' \triangleq [\mathbf{x}'_0 \cdots \mathbf{x}'_{m-1}]$ be the mean-subtracted data matrix, with the mean given by (11.27a). Then

$$\mathbf{K}' = [\mathbf{x}_0 \cdots \mathbf{x}_m] \begin{bmatrix} 1 - \frac{1}{m+1} & & -\frac{1}{m+1} \\ & \ddots & \\ -\frac{1}{m+1} & & 1 - \frac{1}{m+1} \\ -\frac{1}{m+1} & \cdots & -\frac{1}{m+1} \end{bmatrix}. \quad (11.29)$$

The $(m+1) \times m$ rightmost matrix has a full column rank m . If $[\mathbf{x}_0 \cdots \mathbf{x}_m]$ also has a full column rank, then \mathbf{K}' does as well.

Subtracting $\bar{\mathbf{x}}$ from the data vectors has the unexpected result of determining the Ritz values in a way that is completely independent of the data vectors' content. If $\mathbf{c} = [-1 \cdots -1]^T$, then

$\mathbf{K}'\mathbf{c}$ reduces precisely to \mathbf{x}'_m . Therefore, the final snapshot \mathbf{x}'_m is exactly a linear combination of the previous mean-subtracted snapshots in \mathbf{K}' . With this value of \mathbf{c} , an eigendecomposition of the corresponding companion matrix \mathbf{C} yields Ritz values that satisfy

$$0 = \sum_{k=0}^m \lambda_j^k \quad (11.30a)$$

$$= \frac{1 - \lambda_j^{m+1}}{1 - \lambda_j}. \quad (11.30b)$$

(This is also evident from (11.20).) From here, it is easy to see that the Ritz values are the $(m+1)^{\text{th}}$ roots of unity, excluding unity itself. We write this explicitly as

$$\lambda_j = \exp\left(\frac{2\pi i j}{m+1}\right), \quad j = 1, \dots, m. \quad (11.31)$$

Thus, the Ritz values are completely independent of the data; we can compute them with only one parameter, m . Inserting this result into (11.6), we find

$$\mathbf{x}'_k = \sum_{j=1}^m \exp\left(\frac{2\pi i j k}{m+1}\right) \mathbf{v}_j, \quad k = 0, \dots, m. \quad (11.32)$$

Note that the residual typically found on $k = m$ is absent because of the perfect reconstruction of that snapshot.

11.4.3 Equivalence with temporal DFT and harmonic average

We define the temporal DFT of the mean-subtracted data $\{\mathbf{x}'_k\}_{k=0}^m$ by

$$\{\hat{\mathbf{x}}_j\}_{j=0}^m = \mathcal{F} [\{\mathbf{x}'_k\}_{k=0}^m] \triangleq \left\{ \frac{1}{m+1} \sum_{k=0}^m \exp\left(-\frac{2\pi i j k}{m+1}\right) \mathbf{x}'_k \right\}_{j=0}^m, \quad (11.33a)$$

and its inverse transform by

$$\{\mathbf{x}'_k\}_{k=0}^m = \mathcal{F}^{-1} [\{\hat{\mathbf{x}}_j\}_{j=0}^m] \triangleq \left\{ \sum_{j=0}^m \exp\left(\frac{2\pi i j k}{m+1}\right) \hat{\mathbf{x}}_j \right\}_{k=0}^m. \quad (11.33b)$$

(In this representation, we move the scaling term $1/(m+1)$ typically found on the inverse transform to the forward transform.) Note that (11.33a) implies $\hat{\mathbf{x}}_0 = \mathbf{0}$, since the set $\{\mathbf{x}'_k\}_{k=0}^m$ has a zero mean. Thus, (11.33b) reduces to

$$\mathbf{x}'_k = \sum_{j=1}^m \exp\left(\frac{2\pi i j k}{m+1}\right) \hat{\mathbf{x}}_j \quad (11.34)$$

for $k = 0, \dots, m$. The decompositions in (11.32, 11.34) are identical if we let $\mathbf{v}_j = \hat{\mathbf{x}}_j$ for $j = 1, \dots, m$. Appealing to the uniqueness of DMD shown in Section 11.3, we conclude that if a linearly independent data set has a zero mean, then DMD is exactly equivalent to a temporal DFT. This is perhaps surprising, since the DMD algorithm and its derivation are largely unrelated to those of the temporal DFT (see Section 11.2, Rowley et al. 2009, and Schmid 2010). Whereas the

temporal DFT model is based on a superposition of oscillating modes, the DMD model is based on the linear combination of eigenvectors that grow or decay according to their eigenvalues.

Mean-subtracted DMD and the temporal DFT are also equivalent to harmonic averaging, for the correct choices of averaging frequencies. Harmonic averaging is a concept that was applied to Koopman analysis (Mezić and Banaszuk, 2004; Mezić, 2005) before the DMD algorithms of Rowley et al. (2009) and Schmid (2010) were developed. Suppose we pick an oscillation rate $\omega_j = 2\pi j / ((m+1)\Delta t)$ for $j \in \{1, \dots, m\}$, along with time values $t_k = k\Delta t$, where Δt is the time interval between snapshots and $k = 0, \dots, m$. Multiplying \mathbf{x}'_k by this discrete-time oscillation $\exp(-i\omega_j t_k)$ and computing its average, we recover (11.33a). Hence, this choice of averaging frequencies exactly reproduces the temporal DFT. Thus, it is also equivalent to DMD with mean-subtracted data.

We also note briefly that if a data set has a nonzero mean, then a temporal DFT or a harmonic average yields $\hat{\mathbf{x}}_0 = \bar{\mathbf{x}}$. This is akin to a DMD analysis on mean-subtracted data, in which we separately account for the data mean.

11.4.4 Implications of mean subtraction

To make a clear distinction between DMD and mean-subtracted DMD, we will hereafter refer to the latter as a temporal DFT.

A key limitation of the temporal DFT is that its Ritz values depend only on the number of data points and not the data content; see (11.31). Although a Ritz vector energy would still dictate the degree to which that particular DFT mode is present in the data, the decomposition is only capable of outputting a predetermined set of frequencies. To ensure that a particular frequency is properly captured, the data must cover an integer number of corresponding periods. If multiple frequencies are of interest, then this constraint may be prohibitive, especially if the frequencies are unknown or are not related by a simple rational number.

In addition, this constraint implies that the longest period the temporal DFT can capture is the time span of the data set. We will see in the discussion of the cylinder flow example (Section 11.6.5) that DMD and a variant we will introduce in Section 11.5 are not subject to this limitation. Although DMD and its variant are still subject to the Nyquist frequency constraint, there is no theoretical lower bound on the frequencies they can compute.

Secondly, the temporal DFT is wholly incapable of determining modal growth rates. Equation (11.31) implies that all Ritz values have a unit magnitude. This—like the predetermined mode frequencies—is an artifact of the mean subtraction, and is independent of the dynamics underlying the data. Dynamic mode decomposition is predicated on the assumption that meaningful information about a trajectory can be inferred from related linear dynamics. When the mean is removed, important dynamical information is stripped away, and the modal decomposition necessarily yields a periodic trajectory.

Finally, we stress the well-known fact that a temporal DFT on non-periodic data produces a slow decay in the modal energy. We also demonstrate this on the cylinder flow example in Section 11.6.5. In this example, many modes need to be retained to reproduce non-periodic data correctly.

11.4.5 Subtracting an equilibrium

As an alternative to subtracting the mean from a given data set, one can consider subtracting the observable corresponding to an *equilibrium point* of the dynamics. For instance, suppose the dynamics are given by (11.1), with linear boundary conditions (11.22); suppose also that there is

an equilibrium solution $\boldsymbol{\xi}_e$ satisfying $\mathbf{f}(\boldsymbol{\xi}_e) = \boldsymbol{\xi}_e$, along with the boundary conditions $B(\mathbf{g}(\boldsymbol{\xi}_e)) = \gamma$. This is often the case for fluid flows, for instance, in which the equilibrium solution is a steady laminar flow, whereas the observed data consist of snapshots of an unsteady, possibly turbulent, flow that satisfies the same inflow and wall boundary conditions.

If we choose the “base flow” \mathbf{x}_b to be the observable corresponding to this equilibrium (i.e., $\mathbf{x}_b = \mathbf{x}_e = \mathbf{g}(\boldsymbol{\xi}_e)$), then the base flow satisfies the boundary conditions (11.25). Thus, the modes computed from base-flow-subtracted data will satisfy homogeneous boundary conditions, just as in the case where the base flow is the mean of the data. However, using the equilibrium ($\mathbf{x}_b = \mathbf{x}_e$) appears to have significant advantages over using the mean ($\mathbf{x}_b = \bar{\mathbf{x}}$, as in Section 11.4.2). The DMD algorithm no longer reduces to a traditional DFT, and we can once again obtain frequency and growth rate information from the data. We may view the equilibrium subtraction as merely translating the coordinate system so that the equilibrium is at the origin. As we will see in Section 11.6.6, this shift does not appear to change the DMD spectrum and modes in a significant way.

11.5 Optimized dynamic mode decomposition

We can think of the DMD formulation in (11.6) as a way to “curve-fit” $m+1$ nonlinearly-generated data vectors $\{\mathbf{x}_k\}_{k=0}^m$ to a linear model of m modes. The model achieves perfect accuracy for $k = 0, \dots, m-1$, as in (11.6a), and minimizes the error at $k = m$, as in (11.6b). This model, however, has potential pitfalls. First, if the user provides $m+1$ data points but requires fewer than m modes, then it is not entirely clear how to truncate the mode set—especially when the governing dynamics are not simple. This scenario is common in the construction of reduced-order models from data-based decompositions, where the desired number of modes is typically much less than the number of data vectors, and the dynamics are often quite complex.

Second, the DMD algorithm places a residual only on the final data vector \mathbf{x}_m , since the algorithm is built upon the reconstruction of \mathbf{x}_m from $\{\mathbf{x}_k\}_{k=0}^{m-1}$. Numerical experiments show that DMD results are more sensitive to variations in \mathbf{x}_m than to variations in other data vectors. This is because the Ritz values are the eigenvalues of the companion matrix \mathbf{C} (11.9), which dictates this reconstruction. The presence of noise in \mathbf{x}_m could drastically change the contents, and hence the eigenvalues, of the companion matrix.

To address these issues, we propose a new, optimized DMD, in which we fit p points on a trajectory to a linear model of m modes. We must satisfy the inequality $m < p$, but m and p are otherwise arbitrary. We allow the linear model to contain a residual at each of the p points, but the overall residual is minimized for the values of m and p that we choose.

More formally, suppose that $\{\mathbf{x}_k\}_{k=0}^{p-1}$ is a set of real or complex vectors. Given $m < p$, we seek complex scalars $\{\lambda_j\}_{j=1}^m$ and vectors $\{\mathbf{v}_j\}_{j=1}^m$ such that

$$\mathbf{x}_k = \sum_{j=1}^m \lambda_j^k \mathbf{v}_j + \mathbf{r}_k, \quad k = 0, \dots, p-1, \quad (11.35a)$$

and we minimize

$$\Gamma \triangleq \sum_{k=0}^{p-1} \|\mathbf{r}_k\|_2^2. \quad (11.35b)$$

To construct the optimized DMD, define

$$\mathbf{K} \triangleq \begin{bmatrix} \mathbf{x}_0 & \cdots & \mathbf{x}_{p-1} \end{bmatrix} \quad (11.36a)$$

$$\mathbf{V} \triangleq \begin{bmatrix} \mathbf{v}_1 & \cdots & \mathbf{v}_m \end{bmatrix} \quad (11.36b)$$

$$\mathbf{T} \triangleq \begin{bmatrix} 1 & \lambda_1 & \lambda_1^2 & \cdots & \lambda_1^{p-1} \\ 1 & \lambda_2 & \lambda_2^2 & \cdots & \lambda_2^{p-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & \lambda_m & \lambda_m^2 & \cdots & \lambda_m^{p-1} \end{bmatrix} \quad (11.36c)$$

$$\mathbf{r} \triangleq \begin{bmatrix} \mathbf{r}_1 & \cdots & \mathbf{r}_p \end{bmatrix}. \quad (11.36d)$$

We seek \mathbf{V} and the Vandermonde matrix \mathbf{T} such that $\mathbf{K} = \mathbf{V}\mathbf{T} + \mathbf{R}$, and we minimize the squared Frobenius norm $\Gamma = \|\mathbf{R}\|_{\text{F}}^2$. This is essentially a least-squares problem for \mathbf{V} . Of all the choices of \mathbf{V} that minimize $\|\mathbf{R}\|_{\text{F}}^2$, the one with the smallest Frobenius norm is

$$\mathbf{V} = \mathbf{K}\mathbf{T}^+. \quad (11.37)$$

(If $\{\lambda_j\}_{j=1}^m$ are distinct and therefore \mathbf{T} has full row rank, then $\mathbf{T}^+ = \mathbf{T}^* (\mathbf{T}\mathbf{T}^*)^{-1}$.) The residual is therefore $\mathbf{R} = \mathbf{K}(\mathbf{I} - \mathbf{T}^+\mathbf{T})$. Using the trace expansion of the Frobenius norm, we then find that

$$\Gamma = \text{tr}(\mathbf{K}^*\mathbf{K}) - \text{tr}(\mathbf{T}^+\mathbf{T}\mathbf{K}^*\mathbf{K}) \quad (11.38a)$$

$$= \|\mathbf{K}\|_{\text{F}}^2 - \|\mathbf{K}\mathbf{T}^+\mathbf{T}\|_{\text{F}}^2. \quad (11.38b)$$

Although the second form is more intuitive, the first form is faster to compute, since we can precompute and store the small matrix $\mathbf{K}^*\mathbf{K}$. The objective is to find a set of complex numbers $\{\lambda_j\}_{j=1}^m$ that minimizes (11.38); (11.37) then provides the modes.

At present, we do not have an analytic algorithm for computing optimized DMD. To find the choice of $\{\lambda_j\}_{j=1}^m$ that minimizes Γ , we employ a global optimization technique that combines simulated annealing and the Nelder-Mead simplex method (Press et al., 2007). At each iteration, the technique adds thermal fluctuations to the function evaluations at the points of the simplex. When it tests a new point, it subtracts a thermal fluctuation from the function evaluation; therefore, there is a calculable probability of accepting a worse point. As the annealing completes, the fluctuations vanish, and the technique reduces exactly to the Nelder-Mead method of local optimization. We also use the Broyden-Fletcher-Goldfarb-Shanno quasi-Newton iterator for purely local minimization. In this case, we apply the SVD-based DMD algorithm of Schmid (2010) and use select Ritz values as the initial condition for the iterator. The two iterators produce consistent results for $m = 3$, which we report in Section 11.6.4.

This minimization problem is similar to POD's, which minimizes

$$\|\mathbf{K}\|_{\text{F}}^2 - \|\mathbf{V}(\mathbf{V}^*\mathbf{V})^{-1}\mathbf{V}^*\mathbf{K}\|_{\text{F}}^2, \quad (11.39)$$

with the requirement that \mathbf{V} be unitary. In optimized DMD, \mathbf{T} is not unitary; rather, it obeys the Vandermonde structure given in (11.36c).

Finally, we present a property of optimized DMD as a theorem below.

Theorem 11.2. *Suppose the data $\{\mathbf{x}_k\}_{k=0}^{p-1}$ are linearly independent. Then the Ritz values $\{\lambda_j\}_{j=1}^m$ resulting from optimized DMD are distinct.*

Proof. The key idea behind this proof is that an orthogonal projection onto the rows of a Vandermonde matrix attains the maximum norm only when the rows are distinct. We proceed by contradiction.

Suppose that $\{\lambda_j\}_{j=1}^m$ are not distinct. In particular, suppose this set has $q < m$ distinct elements $\{\mu_j\}_{j=1}^q$. Let $\{\mu_j\}_{j=q+1}^m$ be any set of distinct complex scalars *not* in $\{\lambda_j\}_{j=1}^m$, so that $\{\mu_j\}_{j=1}^m$ is a distinct set of complex scalars. Construct the complex conjugate Vandermonde matrices

$$\mathbf{T}^* = \begin{bmatrix} 1 & \cdots & 1 \\ \bar{\lambda}_1 & \cdots & \bar{\lambda}_m \\ \vdots & & \vdots \\ \bar{\lambda}_1^{p-1} & \cdots & \bar{\lambda}_m^{p-1} \end{bmatrix} \quad \mathbf{S}^* \triangleq \begin{bmatrix} 1 & \cdots & 1 \\ \bar{\mu}_1 & \cdots & \bar{\mu}_m \\ \vdots & & \vdots \\ \bar{\mu}_1^{p-1} & \cdots & \bar{\mu}_m^{p-1} \end{bmatrix} \quad (11.40)$$

and the corresponding orthogonal projection matrices $\mathbf{P}_{\mathbf{T}^*} \triangleq \mathbf{T}^* (\mathbf{T}^*)^+$ and $\mathbf{P}_{\mathbf{S}^*} \triangleq \mathbf{S}^* (\mathbf{S}^*)^+$. Also, let $\{\mathbf{y}_k\}_{k=1}^n$ be the columns of \mathbf{K}^* , so that

$$\mathbf{K}^* = \begin{bmatrix} \mathbf{x}_0^* \\ \vdots \\ \mathbf{x}_{p-1}^* \end{bmatrix} = \begin{bmatrix} \mathbf{y}_1 & \cdots & \mathbf{y}_n \end{bmatrix}. \quad (11.41)$$

Thus, each vector \mathbf{y}_k contains the time history of the data at a particular index in \mathbf{x} .

First, for $k = 1, \dots, n$, consider the separation of \mathbf{y}_k into its projection onto the range of \mathbf{S}^* and the orthogonal complement,

$$\mathbf{y}_k = \mathbf{P}_{\mathbf{S}^*} \mathbf{y}_k + \boldsymbol{\zeta}_k, \quad \boldsymbol{\zeta}_k \notin \text{range}(\mathbf{S}^*). \quad (11.42)$$

Since $\text{range}(\mathbf{T}^*) \subset \text{range}(\mathbf{S}^*)$, we deduce that $\mathbf{P}_{\mathbf{T}^*} \boldsymbol{\zeta}_k = \mathbf{0}$. Thus, applying to the above equation a projection onto the range of \mathbf{T}^* , we obtain $\mathbf{P}_{\mathbf{T}^*} \mathbf{y}_k = \mathbf{P}_{\mathbf{T}^*} \mathbf{P}_{\mathbf{S}^*} \mathbf{y}_k$. Using this fact to separate $\mathbf{P}_{\mathbf{S}^*} \mathbf{y}_k$ into its projection onto \mathbf{T}^* and the orthogonal complement,

$$\mathbf{P}_{\mathbf{S}^*} \mathbf{y}_k = \mathbf{P}_{\mathbf{T}^*} \mathbf{y}_k + \boldsymbol{\eta}_k, \quad \mathbf{P}_{\mathbf{T}^*} \mathbf{y}_k \perp \boldsymbol{\eta}_k. \quad (11.43)$$

Furthermore, the Pythagorean theorem states that

$$\|\mathbf{P}_{\mathbf{S}^*} \mathbf{y}_k\|_2^2 = \|\mathbf{P}_{\mathbf{T}^*} \mathbf{y}_k\|_2^2 + \|\boldsymbol{\eta}_k\|_2^2. \quad (11.44)$$

Define $\mathbf{H} = [\boldsymbol{\eta}_1 \cdots \boldsymbol{\eta}_n]$; stack (11.43) and sum (11.44) to form

$$\mathbf{P}_{\mathbf{S}^*} \mathbf{K}^* = \mathbf{P}_{\mathbf{T}^*} \mathbf{K}^* + \mathbf{H} \quad (11.45a)$$

$$\|\mathbf{P}_{\mathbf{S}^*} \mathbf{K}^*\|_{\text{F}}^2 = \|\mathbf{P}_{\mathbf{T}^*} \mathbf{K}^*\|_{\text{F}}^2 + \|\mathbf{H}\|_{\text{F}}^2. \quad (11.45b)$$

The cost function in (11.38) is $\Gamma = \|\mathbf{K}^*\|_{\text{F}}^2 - \|\mathbf{P}_{\mathbf{T}^*} \mathbf{K}^*\|_{\text{F}}^2$. By the construction of optimized DMD, \mathbf{T}^* is the Vandermonde matrix that maximizes $\|\mathbf{P}_{\mathbf{T}^*} \mathbf{K}^*\|_{\text{F}}^2$. That is, $\|\mathbf{P}_{\mathbf{T}^*} \mathbf{K}^*\|_{\text{F}}^2 \geq \|\mathbf{P}_{\mathbf{S}^*} \mathbf{K}^*\|_{\text{F}}^2$ since \mathbf{S}^* is also a Vandermonde matrix, and therefore $\mathbf{H} = \mathbf{0}$. Equation (11.45a) reduces to $\mathbf{P}_{\mathbf{S}^*} \mathbf{K}^* = \mathbf{P}_{\mathbf{T}^*} \mathbf{K}^*$.

The rank property of Vandermonde matrices then yields

$$q = \text{rank}(\mathbf{P}_{\mathbf{T}^*}) = \text{rank}(\mathbf{P}_{\mathbf{T}^*} \mathbf{K}^*) \quad (11.46a)$$

$$m = \text{rank}(\mathbf{P}_{\mathbf{S}^*}) = \text{rank}(\mathbf{P}_{\mathbf{S}^*} \mathbf{K}^*); \quad (11.46b)$$




Figure 11.1: Color fields of vorticity (i.e., velocity rotation), with overlaid velocity vectors, for the cylinder flow at $Re = 60$. Clockwise vorticity is blue, and counterclockwise is red. An exterior uniform flow from left to right is present. (a): The unstable equilibrium of the cylinder flow. (b): A snapshot from the limit cycle.

the right-most side comes from the fact that \mathbf{K}^* has full row rank. We arrive at a contradiction: $\mathbf{P}_S^* \mathbf{K}^*$ and $\mathbf{P}_T^* \mathbf{K}^*$ are equal, yet they have different ranks. Therefore, $\{\lambda_j\}_{j=1}^m$ must be distinct. \square

11.6 Application to a cylinder flow

11.6.1 Overview

The two-dimensional incompressible flow past a cylinder is a canonical problem in fluid mechanics. Its dynamics are relatively well-understood (see Williamson 1996), and real-life examples with similar dynamics can be found in everyday objects such as flagpoles, smokestacks, and even entire skyscrapers. A number of recent studies have focused specifically on dynamical modeling, modal analysis, and flow control of the cylinder wake; see Roussopoulos (1993), Gillies (1998), Noack et al. (2003), Giannetti and Luchini (2007), Marquet et al. (2008), and Tadmor et al. (2010) for a short selection.

The continuity and Navier–Stokes equations,

$$\nabla \cdot \mathbf{u} = 0 \quad (11.47a)$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{\nabla p}{\rho} + \nu \nabla^2 \mathbf{u}, \quad (11.47b)$$

govern the dynamics of the incompressible cylinder flow. We denote the horizontal (free-stream) direction by x and the vertical (normal) direction by y . We assume the fluid density ρ and kinematic viscosity ν are constant, and the velocity $\mathbf{u} = [u_x(x, y, t) \ u_y(x, y, t)]^T$ and pressure $p(x, y, t)$ vary in space and time. Denoting the horizontal unit vector by \mathbf{e}_x , the boundary conditions are

$$\mathbf{u} = U \mathbf{e}_x, \quad x, y \rightarrow \pm\infty \quad (11.48a)$$

$$p = p_0, \quad x, y \rightarrow \pm\infty \quad (11.48b)$$

$$\mathbf{u} = \mathbf{0}, \quad \sqrt{x^2 + y^2} = \frac{D}{2}, \quad (11.48c)$$

where $U \mathbf{e}_x$ is the externally-imposed free-stream velocity, p_0 is the constant far-field pressure, and D is the cylinder diameter.

The qualities of the incompressible cylinder flow are characterizable by only one parameter, the Reynolds number $Re \triangleq UD/\nu$. At Reynolds numbers between 47 (Provansal et al., 1987) and roughly 194 (Williamson, 1996), there exists one unstable equilibrium (Figure 11.1a), consisting




Figure 11.2: (a): The lift coefficient magnitude as the cylinder flow evolves from the unstable equilibrium to the limit cycle. (b): The corresponding instantaneous Strouhal number.

of a long separation bubble with two counter-rotating vortices. There also exists one stable limit cycle (Figure 11.1b), known as a von Kármán vortex street, characterized by a periodic shedding of vortices from the top and bottom surfaces. To test our decomposition methods, we generate data from numerical simulations of a cylinder flow at $Re = 60$. We intentionally pick a Reynolds number close to 47—where the Hopf bifurcation occurs—so that dynamical parameters such as the shedding frequency do not change too drastically from the equilibrium to the limit cycle.

As the cylinder flow evolves from the unstable equilibrium to the limit cycle, we identify three distinct regimes (see Figure 11.2). Near the equilibrium state, perturbations in the flow from the equilibrium are small, and the dynamics are approximately linear. The flow oscillates at a frequency f given by a nondimensional Strouhal number of $St \triangleq fD/U = 0.126$. The lift coefficient—defined by $C_L \triangleq 2F_L/(\rho U^2 D)$ for a lift force per depth F_L —increases in oscillation amplitude by $1.054^{\hat{t}}$, where $\hat{t} \triangleq tU/D$ is the nondimensional advection time. In the transient regime, the Strouhal number increases and the C_L peaks grow more slowly than exponentially. Finally, near the limit cycle, $St = 0.141$ and the C_L peaks roughly reach their asymptotic values. This Strouhal number is consistent with findings by Roshko (1954).

11.6.2 Numerical method

We generate the data using a multi-domain immersed boundary method; see Taira and Colonius (2007) and Colonius and Taira (2008) for details. Each successively smaller domain covers half the streamwise and normal length, and samples the space twice as finely. The coarsest domain spans 48 diameters in the streamwise direction and 16 diameters in the normal direction. On the third and finest domain, shown in Figure 11.1, the computational domain spans 12 diameters in the streamwise direction and four diameters in the normal direction, with a resolution of 480 by 160 grid points. In this study, we only use velocity data from this innermost domain in DMD. We use a time step of $\Delta \hat{t} = 0.01875$, which satisfies the Courant–Friedrichs–Levy condition for the numerical scheme implemented. The two references above have verified the convergence of the scheme, as well as agreement with analytic models.

We use selective frequency damping (SFD; see Åkervik et al., 2006) to solve for the unstable equilibrium. This solution is used as the initial condition of the immersed boundary solver. Because time derivatives are extremely small at the start of the fluid simulation, the computed equilibrium is likely close to the true equilibrium. The numerical error in the SFD solution causes the trajectory to depart from the equilibrium, and ultimately to approach the limit cycle. Only one simulation, the one shown in Figures 11.1 and 11.2, is used in this study.

11.6.3 Dynamic mode decomposition of the cylinder flow

From each of the three regimes shown in Figure 11.2, we compute DMD on x - and y -velocity data spanning approximately two periods of oscillation. We retain every 25th time step from the immersed boundary solver, so the sampling time between successive data vectors is $\Delta\hat{t}_s = 25\Delta\hat{t} = 0.46875$. The data samples from the near-equilibrium, transient, and limit cycle regimes respectively cover $0 \leq \hat{t} < 15.94$, $95.63 \leq \hat{t} < 111.09$, and $210.94 \leq \hat{t} < 225.00$. To cover two periods, we respectively require 34, 33, and 30 snapshots. Although we feed an integer number of periods into DMD, our experience has been that DMD is generally insensitive to this condition. This is in contrast to other methods such as POD, which may yield unintended results when using non-integer periods of data.

Figure 11.3 shows the spectra of Ritz values. The Strouhal number is

$$St = \frac{\arg(\lambda)}{2\pi\Delta\hat{t}_s}, \quad (11.49)$$

and we normalize vector energies $\|\mathbf{v}\|_2$ by $\left\| \begin{bmatrix} \mathbf{K} & \mathbf{x}_m \end{bmatrix} \right\|_{\text{F}} / \sqrt{m+1}$, the root-mean-square 2-norm of the data. The most energetic mode typically has a Ritz value of $\lambda \approx 1$. This mode has the resemblance of (but is not exactly the same as) the mean of the data set; see Figures 11.4(a, b) and 11.6.

In the near-equilibrium regime (Figure 11.3a, b), where the dynamics are nearly linear, DMD easily recognizes the complex conjugate pair of eigenmodes primarily responsible for the oscillating exponential growth away from the unstable equilibrium. One mode from the pair is shown in Figure 11.4(c–e). The spatial oscillations are the result of instabilities that convect downstream. The velocity magnitude of this mode, as shown in Figure 11.4(e), attains its maximum value far downstream of the cylinder. The qualitative nature of this unstable mode is consistent with previous results, such as the eigenmode analyses of Giannetti and Luchini (2007) and Marquet et al. (2008) of the linearized Navier–Stokes equations. An additional unstable pair of modes is visible in Figure 11.3(a, b); we posit that this is a result of numerical noise.

We momentarily skip the transient regime and discuss the limit cycle, as shown in Figure 11.3(e, f). The spectrum appears well-ordered; all but the least energetic modes lie very close to the unit circle. The mode corresponding to $\lambda = 1$ (Figure 11.6) exhibits a shorter tail than the equivalent mode from near the equilibrium. This is expected, and is the basis of “shift mode” analyses of bluff body wakes, as in Noack et al. (2003) and Tadmor et al. (2010). Furthermore, the oscillatory mode frequencies are nearly integral multiples of the base frequency $St = 0.141$. An aliasing effect can be seen in Figure 11.3(f) as the multiples exceed the Nyquist frequency of 1.067. The integral multiples suggest that the cylinder flow’s periodic orbit exhibits harmonic-like behavior. The first and second harmonics are respectively shown in Figure 11.7(a–c) and (d–f). We observe that the first harmonic is symmetric in vorticity about the horizontal axis, and the second harmonic is antisymmetric. Thus, the former corresponds to the large-scale convection of fluid structures in the wake, and the latter corresponds to the top-bottom oscillation visible in Figure 11.1(b).

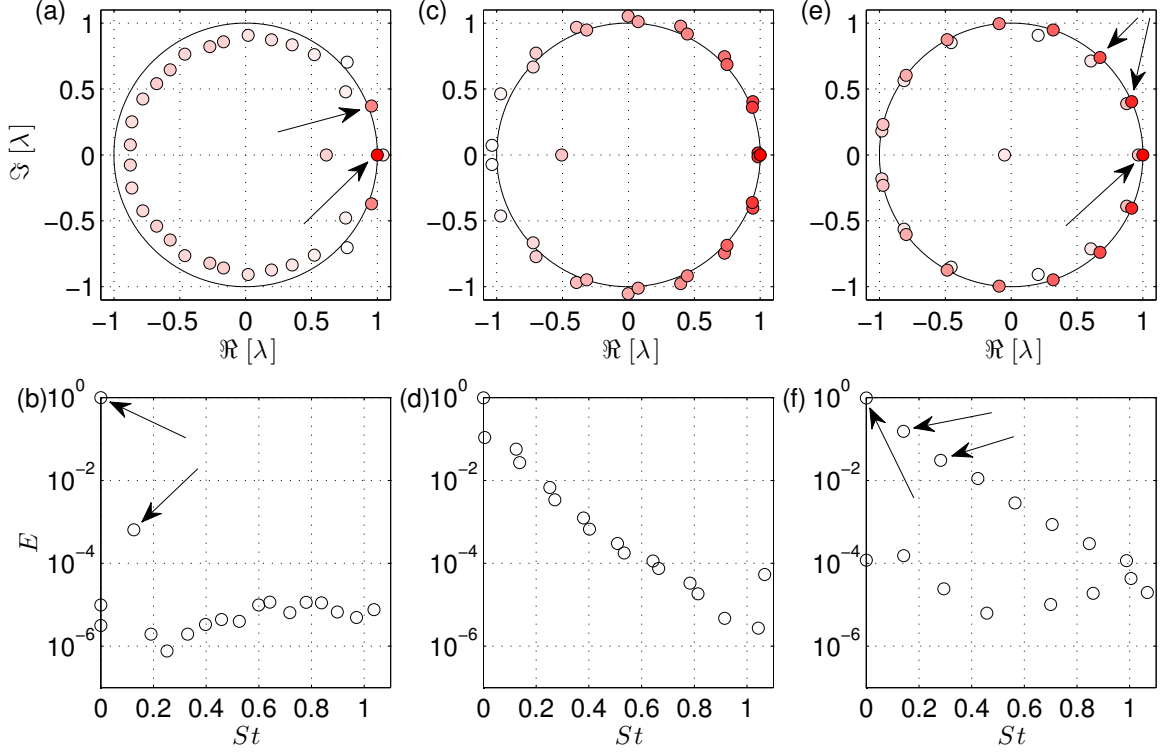


Figure 11.3: Ritz values λ computed from DMD on two periods of data. Top row: as the real and imaginary parts; the normalized vector energies $E = \sqrt{m+1} \|\mathbf{v}\|_2 / \|[\mathbf{K} \mathbf{x}_m]\|_F$ are indicated by the color of the circles, with darker colors indicating greater energies. Bottom row: as the normalized vector energy versus the Strouhal number. Left column: near-equilibrium regime; middle column: transient regime; right column: limit cycle regime. The arrows indicate Ritz values for which the corresponding modes are plotted in Figures 11.4, 11.6, and 11.7.

We also observe that the DMD modes from the limit cycle match very closely with POD modes (see Holmes et al., 1996). The parallel is not exact, because oscillatory DMD modes are complex, whereas POD modes from real data are always real. Nevertheless, let us index DMD modes in order of decreasing vector energy, so that \mathbf{v}_2 is the DMD mode corresponding to $\lambda_2 = 0.915 + 0.404i$ (Figure 11.7a–c), and \mathbf{v}_4 is the mode corresponding to $\lambda_4 = 0.674 + 0.739i$ (Figure 11.7d–f). Similarly, let ϕ_1, \dots, ϕ_4 be the first four POD modes from the same data set. We find that

$$\left| \frac{\langle \mathbf{v}_2, \phi_1 + i\phi_2 \rangle}{\|\mathbf{v}_2\|_2 \|\phi_1 + i\phi_2\|_2} \right|^2 = 0.9996 \quad (11.50a)$$

$$\left| \frac{\langle \mathbf{v}_4, \phi_3 + i\phi_4 \rangle}{\|\mathbf{v}_4\|_2 \|\phi_3 + i\phi_4\|_2} \right|^2 = 0.9999, \quad (11.50b)$$

where $\langle \mathbf{x}_j, \mathbf{x}_k \rangle \triangleq \mathbf{x}_k^* \mathbf{x}_j$ is an inner product. That is to say that in the limit cycle, POD modes are extremely similar to the real and imaginary parts of DMD modes, up to a complex multiplicative factor.

Of the three regimes, the transient one (Figure 11.3c, d) is the most difficult to analyze. Rowley et al. (2009) showed that the analysis of the Koopman operator, on which DMD is based, yields elegant solutions for linear dynamics and periodic dynamics. The transient regime, however, is

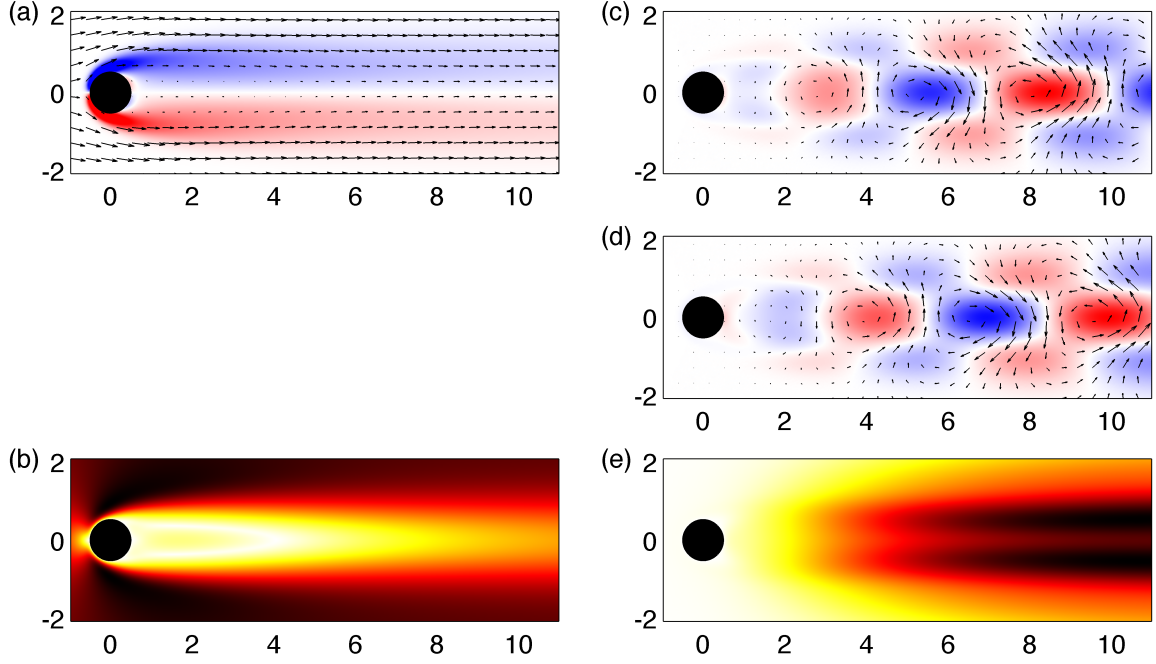


Figure 11.4: Modes \mathbf{v} computed from DMD on two periods of data in the near-equilibrium regime. Left column: $\lambda = 1.000$; right column: $\lambda = 0.956 + 0.371i$ (see Figure 11.3a, b). Top row: the real part $\text{Re}[\mathbf{v}]$; middle row: the imaginary part $\text{Im}[\mathbf{v}]$. Clockwise vorticity is blue, and counterclockwise is red. Note that the mode corresponding to $\lambda = 1.000$ is real. Bottom row: the magnitude of the complex velocity, with zero shown in white and the highest magnitude shown in black.




Figure 11.5: Modes \mathbf{v} computed from two periods of data in the transient regime. Left column: DMD mode at $\lambda = 0.940 + 0.361i$; right column: optimized DMD mode ($m = 3$) at $\lambda = 0.940 + 0.378i$ (see Figure 11.10). Top row: the real part $\text{Re}[\mathbf{v}]$; middle row: the imaginary part $\text{Im}[\mathbf{v}]$. Bottom row: the magnitude of the complex velocity.




Figure 11.6: The DMD mode corresponding to $\lambda = 1.000$, from two periods of data in the limit cycle regime (see Figure 11.3e, f). (a): Velocity and vorticity color fields. (b): The magnitude of the velocity.




Figure 11.7: Modes \mathbf{v} computed from DMD on two periods of data in the limit cycle regime. Left column: $\lambda = 0.915 + 0.404i$; right column: $\lambda = 0.674 + 0.739i$ (see Figure 11.3e, f). Top row: the real part $\text{Re}[\mathbf{v}]$; middle row: the imaginary part $\text{Im}[\mathbf{v}]$. Bottom row: the magnitude of the complex velocity.




Figure 11.8: The movement of the primary unstable Ritz value near the equilibrium (see Figures 11.3a and 11.4c–e) to the primary oscillatory Ritz value near the limit cycle (see Figures 11.3e and 11.7a–c), as a window of 16 snapshots (i.e., roughly one period) traverses the cylinder data. The unit circle is shown as a thick solid curve, and the movement of the Ritz value is shown as a thin solid curve with circles.




Figure 11.9: Ritz values λ computed from DMD on the entire transient regime. (a): As the real and imaginary parts. (b): As the normalized vector energy $E = \sqrt{m+1} \|\mathbf{v}\|_2 / \|\mathbf{K} \mathbf{x}_m\|_F$ versus the Strouhal number.

neither. Rather, it is the part of the highly nonlinear post-Hopf-bifurcation dynamics between the unstable equilibrium and the limit cycle. At $Re = 60$, a smooth “morphing” is visible in the DMD results when a window of snapshots traverses the cylinder data from the equilibrium to the limit cycle (Figure 11.8). Therefore, the Ritz values and vectors are loosely between those of the near-equilibrium regime and those of the limit cycle. For instance, compare the most energetic oscillatory mode, shown in Figure 11.5(a–c), to Figure 11.4(c–e) and Figure 11.7(a–c).

We obtain different results if we apply DMD to the entire transient regime, $78.21 \leq \hat{t} < 167.33$, instead of just two periods; see Figure 11.9. Interpreting this global analysis is especially challenging, because as Figure 11.2(b) implies, the frequency of large-scale oscillation is no longer approximately constant within this window. The most energetic mode exhibits a frequency of $St = 0.136$, which is between the near-equilibrium frequency of $St = 0.126$ and the limit cycle frequency of $St = 0.141$. In addition, each cluster of Ritz values in Figure 11.3(d) is now roughly an inverted parabola; there is now a characterizable spread of modes around each major frequency. Schmid (2010) observed a similar pattern of inverted parabolas when analyzing the flow of a jet between two cylinders, at a Reynolds number of 3,000 based on the volume flux velocity and the cylinder diameter. Although Schmid attributes the parabolas to the spread of spatial and temporal

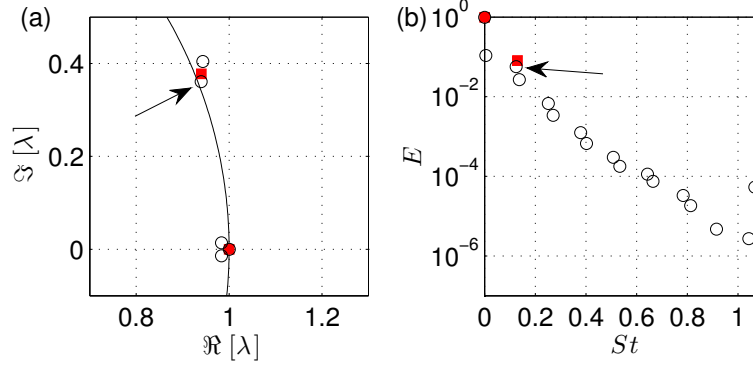


Figure 11.10: Comparison of the DMD (open black \circ) and optimized DMD ($m = 3$; filled red \blacksquare) spectra from two periods of data in the transient regime. (a): As the real and imaginary parts. (b): As the normalized vector energy ($E = \sqrt{m+1} \|\mathbf{v}\|_2 / \|\mathbf{K} \mathbf{x}_m\|_F$ for DMD, $E = \sqrt{p} \|\mathbf{v}\|_2 / \|\mathbf{K}\|_F$ for optimized DMD) versus the Strouhal number. The DMD data are reproduced from Figure 11.3(c, d). The arrows indicate the DMD Ritz value for which the corresponding mode is plotted in Figure 11.5(a–c). The nearby optimized DMD mode is plotted in Figure 11.5(d–f).

scales in the modes of advective-diffusive flows (e.g., in Schmid and Henningson, 2000), we believe that the increasing Strouhal number of the vortex shedding is the leading cause in this case.

11.6.4 Comparison with optimized DMD

Given a set of cylinder flow data, DMD and the m -mode optimized DMD typically yield very similar results, when we truncate the former to retain only the m most energetic modes. For this reason, we do not show the optimized DMD Ritz values and vectors of the entire cylinder flow. The transient regime, however, is an exception. Figure 11.10 compares the Ritz values from this regime, using DMD and a three-mode optimized DMD. Whereas DMD clusters Ritz values with similar vector energies (as in Figure 11.3c, d), optimized DMD combines these clusters into single modes. This is advantageous, since nearly as much dynamical information can be contained in three optimized DMD modes as in seven DMD modes. Furthermore, in Figure 11.14 and its surrounding discussion, we show that the DMD mode clusters may be problematic when truncating the mode set. Comparing Figure 11.5(a–c) and (d–f), we observe small but noticeable differences between the primary oscillatory DMD and optimized DMD modes.

To a large extent, the DMD algorithm is successful at extracting frequencies and growth rates from empirical data. Rowley et al. (2009) demonstrated this by comparing the numerical simulation of a three-dimensional jet in crossflow to its DMD results. The DMD algorithm is similarly successful for the cylinder flow (whose dynamics are simpler than those of the jet in crossflow), except in the nonlinear transient regime. Table 11.1 compares frequencies and growth rates calculated by three techniques. We extract one set of values from C_L oscillation peaks. In the case of DMD, we consider the most energetic oscillatory mode pair, as in Section 11.6.3. Three-mode optimized DMD recovers a similar complex conjugate pair of modes. Denoting the growth rate by β such that the growth follows $\beta^{\hat{t}}$, we have $\beta = |\lambda|^{1/\Delta \hat{t}_s}$.

The frequencies extracted from DMD and optimized DMD match lift force values to very good precision in the near-equilibrium and limit cycle regimes. In the transient regime, however, the three-mode optimized DMD captures the dominant Strouhal number better than DMD. This is because DMD imposes the constraint that the number of modes must be one fewer than the number of data vectors. By considering only three modes, we truncate the majority of the modal

Table 11.1: Comparison of Strouhal numbers St and growth rates β over two periods of data from each of the three regimes. The lift force data, DMD, and three-mode optimized DMD yield these results.

Regime	St			β		
	lift data	DMD	opt. DMD	lift data	DMD	opt. DMD
Near-equilibrium	0.126	0.126	0.126	1.054	1.054	1.054
Transient	0.130	0.124	0.130	1.044	1.017	1.028
Limit cycle	0.141	0.141	0.141	1.000	1.000	1.000

decomposition. In nonlinear dynamics, as in the transient regime, these truncated modes may contain important information that we should not carelessly throw out. On the other hand, the three-mode optimized DMD tailors the decomposition to find a single complex conjugate pair of modes. Thus, optimized DMD provides the “best” single-frequency representation of the transient flow.

Similarly, growth rates agree near the equilibrium, where dynamics are largely linear, and near the limit cycle, where $\beta = 1$ by definition. A larger variation is evident in the transient regime.

11.6.5 Comparison with the temporal DFT

Figure 11.11 shows the spectra produced by the temporal DFT in each of the three data sets. As (11.31) predicts, the Ritz values are exactly the $m + 1$ roots of unity, excluding unity itself. Therefore, to produce modes that correctly capture the features of the flow, it is imperative that the data set cover an integer number of oscillation periods. Comparing the temporal DFT vector energies to the DMD Ritz vector energies in Figure 11.3, we observe that the primary oscillation modes are well-captured. These DFT modes look the same as equivalent optimized DMD modes to the naked eye (Figures 11.4c–e, 11.5d–f, and 11.7).

The temporal DFT mode energies, however, have a much slower decay rate. This is expected in the near-equilibrium and transient regimes, since data sets there are non-periodic. The slow decay is generally undesirable, since the implication is that we may need to retain a large number of modes to construct an accurate solution. In addition, the low-energy modes in the transient regime are irregular and generally without meaningful physical interpretation. This is typically not the case with low-energy DMD and optimized DMD modes.

We highlight an additional advantage of DMD and optimized DMD over the temporal DFT in Figure 11.12. As Section 11.4.4 explains, the largest period that the temporal DFT can compute is the time span of the data set. Therefore, it would have limited utility in analyzing small data sets that cover less than a full period of oscillation. On the other hand, DMD and especially optimized DMD are not constrained as severely. In the cylinder’s limit cycle, where a full period is 15.1 snapshots, we compute DMD on two to 15 snapshots, and we compute a three-mode optimized DMD on four to 15 snapshots. The former computes the growth rate to within a 5% error of unity when eight or more snapshots are used, and the latter when any number of snapshots is used. Similarly, DMD computes the Strouhal number to within a 5% error of 0.141 when nine or more snapshots are used, and optimized DMD when eight or more are used. With only half a period of data, DMD can compute the modal decomposition modestly accurately; optimized DMD performs even better.

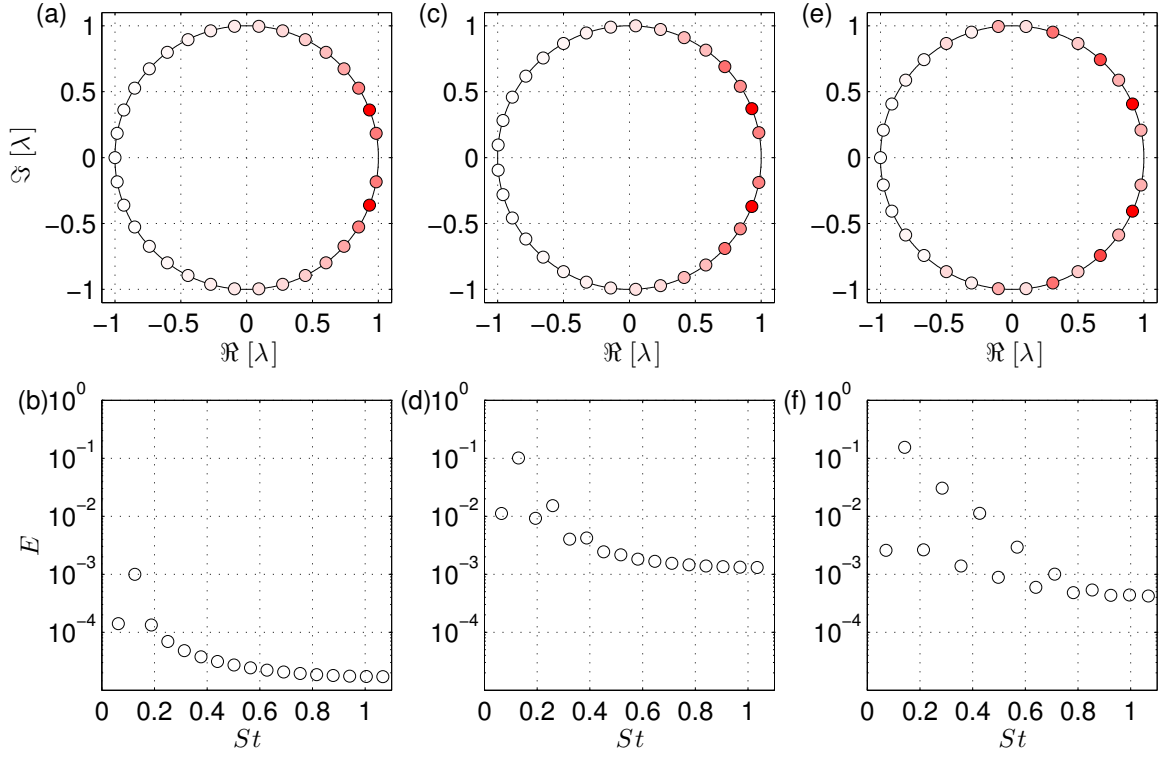


Figure 11.11: Ritz values λ computed from temporal DFT on two periods of data. Top row: as the real and imaginary parts. Bottom row: as the normalized vector energy $E = \sqrt{m+1} \|\mathbf{v}\|_2 / \|[\mathbf{K} \mathbf{x}_m]\|_F$ versus the Strouhal number. Left column: near-equilibrium regime; middle column: transient regime; right column: limit cycle regime.




Figure 11.12: The Ritz value of the limit cycle's primary oscillatory mode, using a full period of data or less in DMD (blue \circ) and three-mode optimized DMD (red $+$). A full period is 15.1 snapshots. The dashed line indicates the correct value. (a): As the growth rate β . (b): As the Strouhal number.

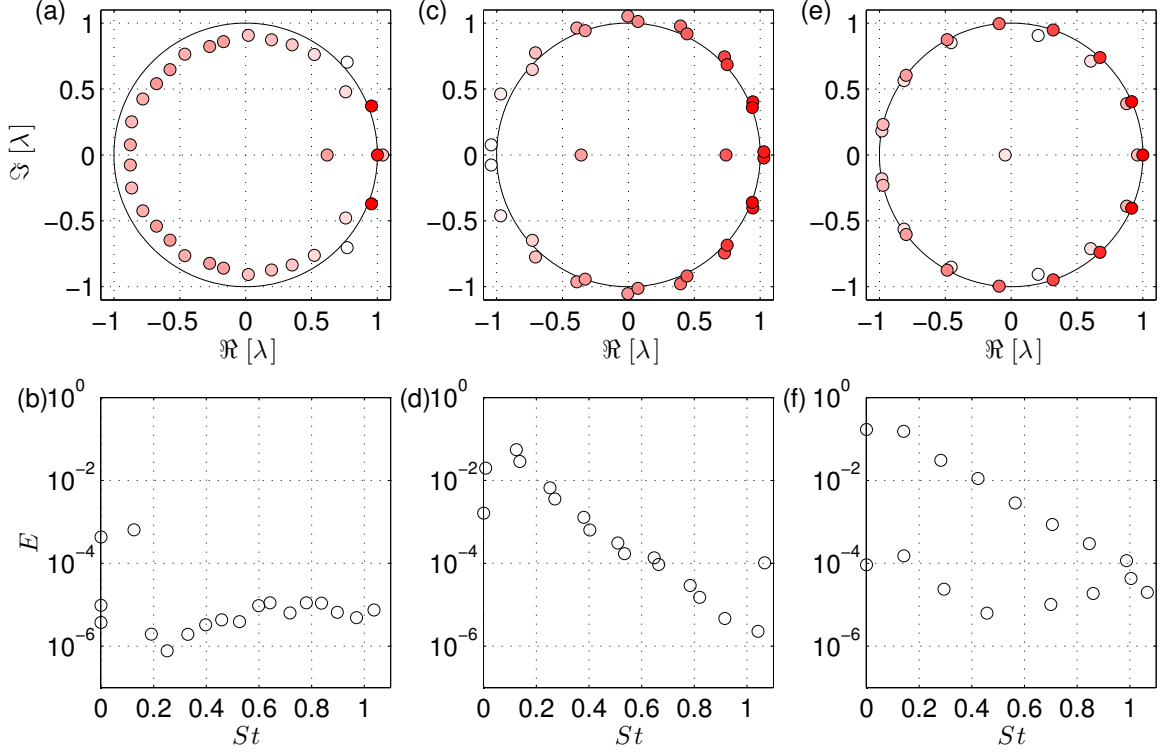


Figure 11.13: Ritz values λ computed from DMD on two periods of equilibrium-subtracted data. Top row: as the real and imaginary parts. Bottom row: as the normalized vector energy $E = \sqrt{m+1} \|\mathbf{v}\|_2 / \|[\mathbf{K} \mathbf{x}_m]\|_F$ versus the Strouhal number. Left column: near-equilibrium regime; middle column: transient regime; right column: limit cycle regime. Except for the Ritz value closest to $\lambda = 1$, the spectra are very close to ones in Figure 11.3, without the equilibrium subtraction.

11.6.6 Comparison with equilibrium subtraction

We show the DMD spectra of the three data sets—with the equilibrium subtracted—in Figure 11.13. Upon comparing these spectra with the ones from the original data sets (Figure 11.3), we notice that the results are nearly identical. One key difference is that the mode closest to $\lambda = 1$ is now significantly weaker. This is a natural consequence of the base flow subtraction. Since the dynamics of the cylinder flow at $Re = 60$ do not depart drastically from the unstable equilibrium, we could expect a stationary DMD mode to resemble the unstable equilibrium. Such a mode should be less present when we subtract the equilibrium from the data.

To the naked eye, these DMD modes from the three regimes look identical to the modes without base flow subtraction; see Figures 11.4, 11.5, and 11.7. We stress again, however, that the base flow subtraction is necessary to ensure that the modes have homogeneous boundary conditions. (The figures only show the innermost domain of the computational grid, and so the actual outer boundaries are not visible.)

11.6.7 Projection error

For the purpose of constructing reduced-order models (typically by Galerkin projection), we also assess the ability of the discussed modal decompositions to yield good projection subspaces. By construction, POD produces the modal subspace that retains the largest possible vector energy after projection (Holmes et al., 1996). Proper orthogonal decomposition, however, is fundamentally

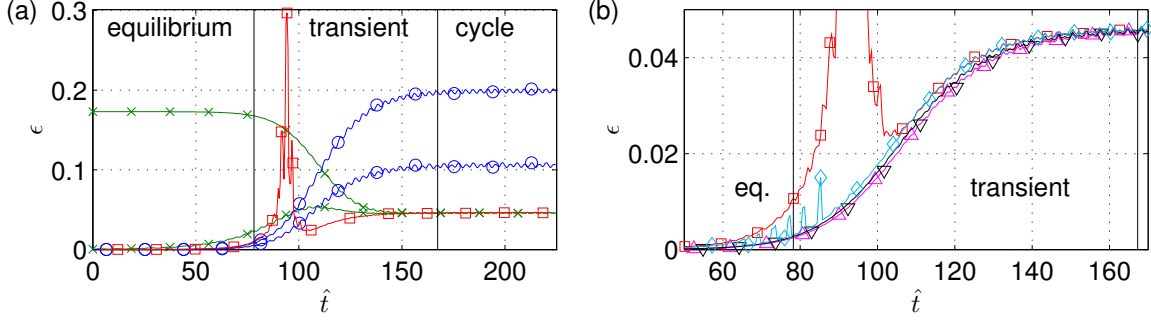


Figure 11.14: (a): The projection error using different DMD modes sets. (blue \circ , top): linear stability modes; (blue \circ , bottom): linear stability modes and the shift mode; (green \times , top): orbital modes; (green \times , bottom): orbital modes and the shift mode; (red \square): time-varying modes. (b): The projection error in the transient regime using different time-varying modes. (red \square): DMD, as in (a); (magenta Δ): temporal DFT; (black ∇): optimized DMD; (cyan \diamond): equilibrium-subtracted DMD. For reference, compare these two plots to Figure 11.2.

decoupled from any sense of time or dynamics, and therefore may not provide the best mode basis for constructing dynamical models.

Suppose that \mathbf{V} is a linearly independent set of modes, and $\mathbf{P}_\mathbf{V} = \mathbf{V}(\mathbf{V}^*\mathbf{V})^{-1}\mathbf{V}^*$ is the orthogonal projection onto the modes. We define the projection error of a data vector \mathbf{x} as

$$\epsilon(\mathbf{x}; \mathbf{V}) \triangleq \frac{\|\mathbf{x} - \mathbf{P}_\mathbf{V}\mathbf{x}\|_2}{\|\mathbf{x}\|_2}. \quad (11.51)$$

A low projection error indicates that the set of modes in \mathbf{V} captures a large amount of information in \mathbf{x} , from an energy standpoint. Again, we stress that modal decompositions with low projection errors do not necessarily create accurate reduced-order models. Nevertheless, this error is still a meaningful measure to consider. A large error suggests that reduced-order models created from the mode set cannot faithfully represent the true state of the dynamical system.

In Figure 11.14(a), we project the two-dimensional velocity data from the cylinder flow simulation onto different mode sets, and we show the resulting projection error $\epsilon(\mathbf{x}(\hat{t}); \mathbf{V})$. We compute linear stability modes from the near-equilibrium regime and orbital modes from the limit cycle, both using DMD. We retain three DMD modes—the $\lambda \approx 1$ mode and the most energetic complex conjugate pair—for projection, and truncate the rest (see Figures 11.4, 11.6, and 11.7). As expected, linear stability modes yield low projection errors near the equilibrium, and orbital modes yield low projection errors in the limit cycle. The errors become large, however, away from the regimes from which we compute each set of modes. Noack et al. (2003) shows that the inclusion of a shift mode, defined in this case as the limit cycle mean minus the unstable equilibrium state, improves the projection error. This is also shown in Figure 11.14(a). We comment that there is little change in projection error if we use the mean flow and two temporal DFT modes instead of three DMD modes; therefore, we do not show temporal DFT results.

We might expect that we could reduce the projection error by projecting data onto modes computed only from their local temporal neighborhood. Thus, we also project each data vector onto three time-varying modes, which we compute by DMD on a data set beginning half a cycle before the data vector and terminating half a cycle after. As shown in Figure 11.14, this is effective at minimizing projection errors, except around $\hat{t} = 94$. There, the error becomes large because there exist two DMD modes with Ritz values $\lambda \approx 1$, and we only retain whichever is closer to

$\lambda = 1$. Should we include the other, the spike in the error would not exist. Nonetheless, this shows that truncating DMD modes may cause unintended and harmful effects.

In Figure 11.14(b), we additionally compute the projection error onto three time-varying temporal DFT and optimized DMD modes. They are nearly equally good at minimizing the projection error, and they outperform any other set of modes tested. The figure also includes the projection error onto time-varying DMD modes with equilibrium subtraction. In this case, the projection is onto the span of the equilibrium, the $\lambda \approx 1$ mode, and the most energetic oscillatory pair of modes.

Finally, we remark that prior studies in the literature suggest the utility of DMD modes for Galerkin modeling. As Sections 11.2 and 11.6.3 mention, the DMD modes near the equilibrium are approximately the eigenmodes of the linearized dynamics, and DMD modes near the orbit are approximately the POD modes. A number of previous studies have already investigated the use of these modes in Galerkin modeling; see the first several references of Tadmor et al. (2010). For a comprehensive study specifically related to the Galerkin modeling of the cylinder wake, refer to Noack et al. (2003); this particular study shows that a combination of POD modes, linear stability modes, and the shift mode accurately models the wake.

11.7 Conclusion

The Koopman operator is a linear operator that represents nonlinear dynamics without linearization. Dynamic mode decomposition is a data-based method for estimating Koopman eigenvalues and modes. It approximates a trajectory, often from nonlinear dynamics, as the result of a linear process. The Ritz values of the decomposition yield growth rates and frequencies, and the Ritz vectors yield corresponding directions. Unlike many previous decomposition techniques such as POD and balanced POD, DMD is not tightly constrained; the data need be neither periodic nor from a linear process for DMD to construct a meaningful modal decomposition.

We prove that the DMD algorithm of Rowley et al. (2009) and Schmid (2010) produces a unique decomposition if and only if all vectors except the last in a given data set are linearly independent, and the Ritz values are distinct. If the data violate either condition, then non-unique Ritz values and vectors may still exist.

To use DMD modes in a reduced-order Galerkin model of a partial differential equation, we must typically include a “base flow” that satisfies the appropriate boundary conditions. The mean of the data is a common choice for a base flow, especially in POD-based analyses. Dynamic mode decomposition on a mean-subtracted set of data, however, is exactly equivalent to a temporal DFT and harmonic averaging. This implies that for $m + 1$ data vectors, the Ritz values are precisely the $m + 1$ roots of unity, excluding unity itself. The inability to compute growth rates, as well as the predetermination of frequencies regardless of data content, are fundamental properties of the temporal DFT. In addition, the temporal DFT cannot capture a particular frequency if the data set covers less than one corresponding period. Furthermore, the decay of temporal DFT vector energies is slow if the data are non-periodic. We can avoid these issues by choosing the equilibrium point instead as the base flow.

Next, we introduce an optimized DMD, which tailors the decomposition specifically to the desired number of modes. In this decomposition, the representation of any data vector may contain a residual, but this method minimizes the total residual over all data vectors. In the absence of an analytic algorithm, we use simulated annealing and quasi-Newton iterators to compute the optimized DMD. Aside from these iterators, this decomposition generally does not suffer from numerical sensitivity as DMD does. Optimized DMD is also superior for computing a small number

of modes, since the decomposition constructs a user-given number of modes and does not require arbitrary truncation.

We use the two-dimensional incompressible flow over a cylinder as a test bed for the aforementioned decomposition methods. Dynamic mode decomposition recovers the leading unstable eigenmodes from data near the unstable equilibrium. It also recovers POD-like modes from data near the limit cycle. When a small window of snapshots traverses the cylinder data from the equilibrium to the limit cycle, the Ritz values and vectors morph gradually.

The optimized DMD and equilibrium-subtracted DMD of the cylinder flow reproduce the DMD results to a large extent. A significant improvement is visible in the transient regime, however. Optimized DMD computes growth rates and frequencies more faithfully, and it avoids clusters of Ritz values, which are problematic for truncating mode sets. Equilibrium-subtracted DMD, on the other hand, largely preserves the DMD modes and spectra, while correctly accounting for boundary conditions. Temporal DFT modes of the cylinder flow look qualitatively like optimized DMD modes, but the roll-off in the Ritz vector energies is slow when the data are not from the periodic orbit. Unlike the temporal DFT, DMD is able to compute frequencies and growth rates from little over half a period of data; optimized DMD fares even better.

As the cylinder flow evolves from the unstable equilibrium to the stable limit cycle, time-varying temporal DFT, optimized DMD, and equilibrium-subtracted DMD modes are better able to represent (in the sense of energy) the data than time-varying DMD modes. These mode sets maintain a low error through the entire trajectory of the cylinder flow, from the equilibrium to the limit cycle. Based on these findings, we hypothesize that the use of these time-varying modes may yield accurate Galerkin models.

This work was supported by the Department of Defense National Defense Science & Engineering Graduate (DOD NDSEG) Fellowship, the National Science Foundation Graduate Research Fellowship Program (NSF GRFP), and AFOSR grant FA9550-09-1-0257. We thank Bernd Noack for discussions on mean subtraction.

Chapter 12

Normalized coprime robust stability and performance guarantees for reduced-order controllers

Kevin K. Chen and Clarence W. Rowley

Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

Appears in *IEEE Transactions on Automatic Control* 58(4):1068–1073, April 2013.

Digital Object Identifier [doi:10.1109/TAC.2012.2218152](https://doi.org/10.1109/TAC.2012.2218152)

Constructing a reduced-order controller from a high-dimensional plant is commonly necessary. The “reduce-then-design” approach constructs the controller from a reduced-order plant; “design-then-reduce” reduces a full-order controller. In both cases, we present sufficient conditions for the full-order plant and reduced-order controller to achieve closed-loop stability or performance. These conditions, motivated primarily by the ν -gap metric, reveal model reduction orders that guarantee stability or performance. The control of the linearized Ginzburg–Landau system provides validation.

12.1 Introduction

The last few decades have seen great development in linear time-invariant (LTI) state-space methods for computing powerful and effective controllers. Many real-life dynamical models, however, have a very large state dimension. Since controllers are often at least as large as the plant, it is frequently preferable to implement reduced-order controllers instead. When the matrix representing the plant dynamics is available, there exist several methods for designing such reduced-order controllers (Obinata and Anderson, 2001). In certain applications, however, the dynamical operator is so large that matrix computations, such as the solution of Riccati equations or linear matrix inequalities, are impractical or impossible. In the control of a linearized fluid flow, for instance, the number of states is typically on the order of the number of grid points used to represent the fluid dynamics. Such a number often ranges from hundreds of thousands to billions.

In the context of these very-large-dimensional plants, designers often first reduce the plant dynamics to an approximate system of much lower dimension. If a matrix representation of the dynamical operator is available, then standard methods include modal truncation, balanced

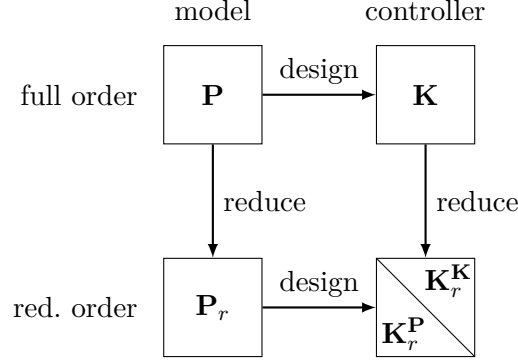


Figure 12.1: Two approaches for arriving at different reduced-order controllers \mathbf{K}_r from a full-order model \mathbf{P} , as used in this study (cf. Obinata and Anderson (2001), Figure 3.1.1).

truncation, optimal Hankel norm reduction, and coprime factor model reduction (Skogestad and Postlethwaite, 2005, Ch. 11). Should the matrix representation be unavailable, or too large to be tractable, designers may use data-based techniques such as proper orthogonal decomposition (Holmes et al., 1996) or balanced proper orthogonal decomposition (BPOD; Rowley, 2005), as well as system-identification methods such as the eigensystem realization algorithm (ERA; Juang and Pappa, 1985). In either case, the designer then constructs the controller from the reduced-order model (i.e., “reduce-then-design”), with the assumption that the reduced-order model is a sufficiently accurate representation of the original dynamics.

Alternatively, we may design the controller directly from the full-order plant. For large-dimensional plants, certain techniques approximate the full-order controller of interest directly from the original plant. For instance, Chandrasekhar’s method (Kailath, 1973) and recently-proposed methods such as the adjoint of the direct-adjoint (Bewley et al., 2011) approximate the \mathcal{H}_2 optimal controller. Whether we directly compute or approximate the full-order controller, we may then reduce the controller (i.e., “design-then-reduce”) to make their application more computationally tractable. There exists a wealth of literature that discusses effective methods for controller reduction. For instance, Goddard and Glover (1998), Obinata and Anderson (2001), and Gao et al. (2006) investigate frequency-weighted model reduction techniques, which are constructed from approximations of an internal-model-control-based transfer function.

Figure 12.1 depicts the two methods for designing low-order controllers. We denote the design-then-reduce controller $\mathbf{K}_r^{\mathbf{K}}$, and the reduce-then-design controller $\mathbf{K}_r^{\mathbf{P}}$. The notation \mathbf{K}_r may refer to either one.

A number of works have also specifically investigated the reduced-order control of distributed-parameter (e.g., partial differential equation) systems (e.g., Burns and King, 1998; King and Sachs, 2000; Atwell and King, 2004; Batten and Evans, 2010). Some of these references, as well as Obinata and Anderson (2001), loosely argue that reduce-then-design is inherently inferior to design-then-reduce. The latter delays the approximation step as much as possible, whereas the former allows approximation error to propagate.

In addition, certain works in the model reduction literature have addressed robust stability or performance in the context of controller reduction (e.g., Zhou and Chen, 1995; Goddard and Glover, 1998; King and Sachs, 2000; Wang and Huang, 2003). In particular, a recent work (Dehkordi and Boulet, 2011) analyzes multiplicative uncertainty using structured singular values and the $M\Delta$ -structure, focusing on balanced truncation for controller reduction. This model of controller reduction yields lower bounds on the controller order necessary for robust stability or performance.

In the present study, we connect concepts from robust stability and performance (such as the normalized coprime stability margin and the ν -gap metric) with concepts from model reduction (such as Hankel singular values and \mathcal{H}_∞ norms on model reduction error). By doing so, we derive simple theorems that guarantee the closed-loop stability or performance of full-order plants with reduced-order controllers. The primary merit of these theorems is that they are more generalizable and readily applicable to the design of reduced-order controllers—we believe—than other known robust stability or performance results. Furthermore, as in Dehkordi and Boulet (2011), and as we will see in the included example, the theorems provide a sense of the required model reduction accuracy for stability or performance. Control designers may therefore use them to guide the model reduction process. Lastly, our results are applicable to both the reduce-then-design and the design-then-reduce methods.

This paper is organized as follows. We briefly review elements of the normalized coprime stability margin and the ν -gap margin in Section 12.2. Section 12.3 presents the main results of the paper by proving sufficient conditions for the stability, or desired performance level, of full-order plants in closed-loop with reduced-order controllers. In Section 12.4, we show examples of these conditions using the control of the linearized Ginzburg–Landau system, and we provide concluding remarks in Section 12.5.

12.2 The normalized coprime stability margin and the ν -gap metric

In this paper, we use the following notation. We denote the conjugate transpose of a matrix or transfer function \mathbf{G} by \mathbf{G}^* . The maximum singular value of a matrix or transfer function \mathbf{G} is denoted $\bar{\sigma}(\mathbf{G})$. \mathcal{L}_∞ is the Lebesgue space of matrix functions essentially bounded on the imaginary axis, with norm $\|\mathbf{G}\|_\infty = \text{ess sup}_{\omega \in \mathbb{R}} \bar{\sigma}(\mathbf{G}(i\omega))$. The space of stable, continuous time, LTI transfer functions is \mathcal{H}_∞ , and $\mathcal{H}_\infty \subset \mathcal{L}_\infty$. \mathcal{R} is the space of rational functions, and $\mathcal{RL}_\infty \triangleq \mathcal{R} \cap \mathcal{L}_\infty$ and $\mathcal{RH}_\infty \triangleq \mathcal{R} \cap \mathcal{H}_\infty$. For $\mathbf{G} \in \mathcal{RL}_\infty$, $\|\mathbf{G}\|_\infty = \max_{\omega \in \mathbb{R}} \bar{\sigma}(\mathbf{G}(i\omega))$. We denote the k th largest Hankel singular value of $\mathbf{G} \in \mathcal{RH}_\infty$ by $\sigma_k(\mathbf{G})$. In the case of an unstable \mathbf{G} , we first separate \mathbf{G} into a stable and an anti-stable part; $\sigma_k(\mathbf{G})$ only includes the Hankel singular values of the stable part.

Next, we review the framework of the normalized coprime stability margin and the ν -gap metric. We base this discussion on Vinnicombe (1993) and Vinnicombe (2001); consult these references for a complete description.

The normalized coprime stability margin and the ν -gap metric provide a useful framework for providing bounds on closed-loop performance. Suppose the stability margin $b_{\mathbf{P}, \mathbf{K}}$ of a plant \mathbf{P} and a controller \mathbf{K} , as well as the ν -gap metric $\delta_\nu(\mathbf{P}, \mathbf{P}_p)$ between \mathbf{P} and a perturbed plant \mathbf{P}_p , are both known. Then, we may analytically solve for lower and upper bounds on $b_{\mathbf{P}_p, \mathbf{K}}$. The utility of this framework is multifaceted. First, we can potentially guarantee that the closed loop system $[\mathbf{P}_p, \mathbf{K}]$ will be stable and achieve a certain performance level—even though we constructed \mathbf{K} to control \mathbf{P} , not \mathbf{P}_p . Furthermore, we may determine this without actually testing $[\mathbf{P}_p, \mathbf{K}]$. In addition, the framework of normalized coprime factorization is more generalizable than, for instance, multiplicative or inverse multiplicative uncertainty alone. Finally, the framework for \mathbf{P} , \mathbf{K} , and a perturbed controller \mathbf{K}_p is exactly analogous.

We first introduce some definitions that we use in our theory.

Definition 12.1. The closed-loop system $[\mathbf{P}, \mathbf{K}]$ is *stable* if the eight transfer functions from \mathbf{v}_1 , \mathbf{v}_2 , \mathbf{v}_3 , and \mathbf{v}_4 to \mathbf{u} and \mathbf{y} in Figure 12.2 are stable (Vinnicombe, 2001, Definition 1.1). These eight




Figure 12.2: The standard feedback interconnection (Vinnicombe, 2001, Figure 1.1).

transfer function relations are

$$\begin{bmatrix} y \\ u \end{bmatrix} = \begin{bmatrix} \mathbf{P} \\ \mathbf{I} \end{bmatrix} (\mathbf{I} - \mathbf{KP})^{-1} \begin{bmatrix} -\mathbf{K} & \mathbf{I} \end{bmatrix} \begin{bmatrix} v_2 \\ v_1 \end{bmatrix} + \begin{bmatrix} \mathbf{I} \\ \mathbf{K} \end{bmatrix} (\mathbf{I} - \mathbf{PK})^{-1} \begin{bmatrix} \mathbf{I} & -\mathbf{P} \end{bmatrix} \begin{bmatrix} v_4 \\ v_3 \end{bmatrix}. \quad (12.1)$$

Definition 12.2. The *normalized right coprime factorization* of $\mathbf{G} \in \mathcal{R}^{p \times q}$ is the ordered pair (\mathbf{N}, \mathbf{M}) , where $\mathbf{N} \in \mathcal{RH}_{\infty}^{p \times q}$ and $\mathbf{M} \in \mathcal{RH}_{\infty}^{q \times q}$, $\mathbf{G} = \mathbf{NM}^{-1}$, and $\mathbf{N}^* \mathbf{N} + \mathbf{M}^* \mathbf{M} = \mathbf{I}$. The transfer functions \mathbf{N} and \mathbf{M} are unique up to a right multiplication by a unitary matrix. The *normalized left coprime factorization* is exactly analogous; see Vinnicombe (2001, Chapter 1.2.1).

Definition 12.3. For a $p \times q$ plant \mathbf{P} and a $q \times p$ controller \mathbf{K} , the *normalized coprime stability margin* is

$$b_{\mathbf{P}, \mathbf{K}} \triangleq \begin{cases} \left\| \begin{bmatrix} \mathbf{P} \\ \mathbf{I} \end{bmatrix} (\mathbf{I} - \mathbf{KP})^{-1} \begin{bmatrix} -\mathbf{K} & \mathbf{I} \end{bmatrix} \right\|_{\infty}^{-1}, & [\mathbf{P}, \mathbf{K}] \text{ stable} \\ 0, & \text{otherwise} \end{cases} \quad (12.2)$$

(see Vinnicombe, 2001, (2.1)). It is a measure of both the performance and the robustness of $[\mathbf{P}, \mathbf{K}]$. Furthermore, $b_{\mathbf{P}, \mathbf{K}} = b_{\mathbf{K}, \mathbf{P}}$.

Definition 12.4. For a system $\mathbf{G} \in \mathcal{R}^{p \times q}$ given by a normalized right coprime factorization (\mathbf{N}, \mathbf{M}) , and a perturbed system given by $\mathbf{G}_p = (\mathbf{N} + \Delta \mathbf{N})(\mathbf{M} + \Delta \mathbf{M})^{-1} \in \mathcal{R}^{p \times q}$,

$$\delta_{\nu}(\mathbf{G}, \mathbf{G}_p) \triangleq \inf_{\substack{\Delta \mathbf{N}, \Delta \mathbf{M} \in \mathcal{L}_{\infty} \\ \text{wno } |\mathbf{M} + \Delta \mathbf{M}| = \eta(\mathbf{G}_p)}} \left\| \begin{bmatrix} \Delta \mathbf{N} \\ \Delta \mathbf{M} \end{bmatrix} \right\|_{\infty} \quad (12.3)$$

(see Vinnicombe, 2001, Definition 1.8). Here, wno $g(s)$ is the counterclockwise winding number around 0 of the scalar function $g(s)$ as s follows the counterclockwise Nyquist D-contour (indented to the right of any pure imaginary poles or zeros). Also, $|\cdot|$ indicates the determinant, and $\eta(\mathbf{G})$ is the number of open right-half-plane poles of \mathbf{G} . The ν -gap metric is a metric on \mathcal{R} ; it obeys the relation $\delta_{\nu}(\mathbf{G}, \mathbf{G}_p) = \delta_{\nu}(\mathbf{G}_p, \mathbf{G})$.

We may also compute the ν -gap metric using a more readily calculable relation. First, we define the matrix square root.

Definition 12.5. Consider the matrix $\mathbf{X} \in \mathbb{C}^{n \times m}$. The Hermitian square root $(\mathbf{X}^* \mathbf{X})^{\frac{1}{2}}$ of the matrix $\mathbf{X}^* \mathbf{X}$ is the unique positive semidefinite matrix satisfying $((\mathbf{X}^* \mathbf{X})^{\frac{1}{2}})^2 = \mathbf{X}^* \mathbf{X}$ and $(\mathbf{X}^* \mathbf{X})^{\frac{1}{2}} = ((\mathbf{X}^* \mathbf{X})^{\frac{1}{2}})^*$. (See Vinnicombe, 2001, page 69.)

Given this definition, we may calculate the ν -gap metric by

$$\delta_\nu(\mathbf{G}, \mathbf{G}_p) = \left\| (\mathbf{I} + \mathbf{G}_p \mathbf{G}_p^*)^{-\frac{1}{2}} (\mathbf{G} - \mathbf{G}_p) (\mathbf{I} + \mathbf{G}^* \mathbf{G})^{-\frac{1}{2}} \right\|_\infty, \quad (12.4)$$

if $|\mathbf{I} + \mathbf{G}_p^*(i\omega)\mathbf{G}(i\omega)| \neq 0$ for all $\omega \in \mathbb{R}$, and $\text{wno}|\mathbf{I} + \mathbf{G}_p^*\mathbf{G}| + \eta(\mathbf{G}) - \eta(\mathbf{G}_p) - \eta_0(\mathbf{G}_p) = 0$, where $\eta_0(\mathbf{G}_p)$ is the number of pure imaginary poles of \mathbf{G}_p (Zhou and Doyle, 1998, Thm. 17.6). Otherwise, $\delta_\nu(\mathbf{G}, \mathbf{G}_p) = 1$.

The key relation that bounds the normalized coprime stability margin is

$$\sin^{-1} b_{\mathbf{P}_p, \mathbf{K}} \geq \sin^{-1} b_{\mathbf{P}, \mathbf{K}} - \sin^{-1} \delta_\nu(\mathbf{P}, \mathbf{P}_p). \quad (12.5a)$$

(See, for instance, Vinnicombe, 2001, (3.2)). For a more detailed discussion and proof, see Vinnicombe (1993, Theorem 4.2) or Vinnicombe (2001, Theorem 3.8). Since inputs to δ_ν commute, we have that

$$\sin^{-1} b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}} \geq \sin^{-1} b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}} - \sin^{-1} \delta_\nu(\mathbf{P}, \mathbf{P}_r). \quad (12.5b)$$

Furthermore,

$$\sin^{-1} b_{\mathbf{K}_r^{\mathbf{K}}, \mathbf{P}} \geq \sin^{-1} b_{\mathbf{K}, \mathbf{P}} - \sin^{-1} \delta_\nu(\mathbf{K}, \mathbf{K}_r^{\mathbf{K}}); \quad (12.5c)$$

since the inputs to b commute,

$$\sin^{-1} b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{K}}} \geq \sin^{-1} b_{\mathbf{P}, \mathbf{K}} - \sin^{-1} \delta_\nu(\mathbf{K}, \mathbf{K}_r^{\mathbf{K}}). \quad (12.5d)$$

12.3 Guarantees on the stability and performance of $[\mathbf{P}, \mathbf{K}_r]$

Equations (12.5b, 12.5d) motivate the stability and performance guarantees. If the right-hand side of these inequalities is positive, then $[\mathbf{P}, \mathbf{K}_r]$ must be stable. Additionally, if the right-hand side is greater than some positive value, then we can provide a lower bound on $b_{\mathbf{P}, \mathbf{K}_r}$ greater than zero.

The main results of this paper are readily applicable when the model reduction method of choice has an analytic upper bound on the \mathcal{H}_∞ norm of the error. For instance, balanced truncation has the upper bound $\|\mathbf{G} - \mathbf{G}_r\|_\infty \leq 2 \sum_{k=r+1}^n \sigma_k(\mathbf{G})$ (Skogestad and Postlethwaite, 2005, (11.17)), and a particular construction of optimal Hankel norm reduction has the bound $\|\mathbf{G} - \mathbf{G}_r\|_\infty \leq \sum_{k=r+1}^n \sigma_k(\mathbf{G})$ (Skogestad and Postlethwaite, 2005, (11.29–30)).

Our main results are based on the following basic theorem.

Theorem 12.1. *If $\mathbf{G}, \mathbf{G}_r \in \mathcal{RL}_\infty$ and $\delta_\nu(\mathbf{G}, \mathbf{G}_r) < 1$, then*

$$\delta_\nu(\mathbf{G}, \mathbf{G}_r) \leq \|\mathbf{G} - \mathbf{G}_r\|_\infty. \quad (12.6)$$

Proof. See, for instance, the proof of Proposition 3 in Jones and Kerrigan (2010). \square

We now show conditions that guarantee the closed-loop stability or performance of the full plant and reduced-order controller.

Theorem 12.2 (Condition for design-then-reduce performance). *If $\mathbf{K}_r^{\mathbf{K}}$ is an approximation of \mathbf{K} that satisfies $\|\mathbf{K} - \mathbf{K}_r^{\mathbf{K}}\|_\infty \leq \alpha$ for some $\alpha \in [0, 1)$, and $\delta_\nu(\mathbf{K}, \mathbf{K}_r^{\mathbf{K}}) < 1$, then the condition*

$$\sin^{-1} b_{\mathbf{P}, \mathbf{K}} > \sin^{-1} \alpha + \sin^{-1} b_d \quad (12.7)$$

is sufficient for satisfying the performance criterion $b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{K}}} > b_d$ for the desired $b_d \in [0, 1)$.

Proof. Given the stated conditions, (12.7) implies

$$\sin^{-1} b_{\mathbf{P}, \mathbf{K}} > \sin^{-1} \left\| \mathbf{K} - \mathbf{K}_r^{\mathbf{K}} \right\|_{\infty} + \sin^{-1} b_d, \quad (12.8a)$$

and Theorem 12.1 yields

$$\sin^{-1} b_{\mathbf{P}, \mathbf{K}} > \sin^{-1} \delta_{\nu}(\mathbf{K}, \mathbf{K}_r^{\mathbf{K}}) + \sin^{-1} b_d. \quad (12.8b)$$

Equations (12.5d, 12.8b) reduce to $b_{\mathbf{P}, \mathbf{K}^{\mathbf{K}}} > b_d$. \square

Corollary 12.1 (Weak condition for design-then-reduce stability). *If $\mathbf{K}_r^{\mathbf{K}}$ is an approximation of \mathbf{K} that satisfies $\left\| \mathbf{K} - \mathbf{K}_r^{\mathbf{K}} \right\|_{\infty} \leq \alpha$ for some $\alpha \in [0, 1)$, and $\delta_{\nu}(\mathbf{K}, \mathbf{K}_r^{\mathbf{K}}) < 1$, then the condition*

$$b_{\mathbf{P}, \mathbf{K}} > \alpha \quad (12.9)$$

is sufficient for the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{K}}]$.

Proof. This follows directly from Theorem 12.2, with the choice $b_d = 0$. \square

We may derive a stronger condition, however, from Obinata and Anderson (2001, Chapter 3.2).

Theorem 12.3 (Strong condition for design-then-reduce stability). *If $[\mathbf{P}, \mathbf{K}]$ is stable, $\mathbf{K}_r^{\mathbf{K}}$ is an approximation of \mathbf{K} that satisfies $\left\| \mathbf{K} - \mathbf{K}_r^{\mathbf{K}} \right\|_{\infty} \leq \alpha$ for some $\alpha \in [0, 1)$, and $\eta(\mathbf{K}) = \eta(\mathbf{K}_r^{\mathbf{K}})$, then*

$$\left\| \mathbf{P}(\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \right\|_{\infty}^{-1} > \alpha \quad (12.10)$$

is sufficient for the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{K}}]$.

Proof. An internal-model-control-like structure (Obinata and Anderson, 2001, Chapter 3.2) shows that if $\mathbf{K}, \mathbf{K}_r^{\mathbf{K}} \in \mathcal{RL}_{\infty}$, $[\mathbf{P}, \mathbf{K}]$ is stable, $\eta(\mathbf{K}) = \eta(\mathbf{K}_r^{\mathbf{K}})$, and

$$1 > \left\| (\mathbf{K} - \mathbf{K}_r^{\mathbf{K}})\mathbf{P}(\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \right\|_{\infty}, \quad (12.11a)$$

then $[\mathbf{P}, \mathbf{K}_r^{\mathbf{K}}]$ is stable. From the stated conditions and (12.10), we see that

$$1 > \left\| \mathbf{K} - \mathbf{K}_r^{\mathbf{K}} \right\|_{\infty} \left\| \mathbf{P}(\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \right\|_{\infty} \quad (12.11b)$$

$$\geq \left\| (\mathbf{K} - \mathbf{K}_r^{\mathbf{K}})\mathbf{P}(\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \right\|_{\infty}; \quad (12.11c)$$

therefore, $[\mathbf{P}, \mathbf{K}_r^{\mathbf{K}}]$ must be stable. \square

This condition is stronger than Corollary 12.1, because it can be shown using \mathcal{H}_{∞} norm properties that $\left\| \mathbf{P}(\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \right\|_{\infty}^{-1} \geq b_{\mathbf{P}, \mathbf{K}}$. The difference between the two sides of this inequality, however, may not be large in practice; see Figure 12.4 for an example comparison.

The design-then-reduce conditions presented above may be difficult or impossible to compute for very large systems. Reduce-then-design is a more common approach in the control of very large systems (e.g., fluid flow control); the conditions for reduce-then-design performance and stability—which we present below—are easy to compute. These results and their proofs are exactly analogous to the design-then-reduce counterparts.

Theorem 12.4 (Condition for reduce-then-design performance). *If \mathbf{P}_r is an approximation of \mathbf{P} that satisfies $\|\mathbf{P} - \mathbf{P}_r\|_\infty \leq \beta$ for some $\beta \in [0, 1)$, and $\delta_\nu(\mathbf{P}, \mathbf{P}_r) < 1$, then the condition*

$$\sin^{-1} b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}} > \sin^{-1} \beta + \sin^{-1} b_d \quad (12.12)$$

is sufficient for satisfying the performance criterion $b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}} > b_d$ for the desired $b_d \in [0, 1)$.

Proof. Given the stated conditions, (12.12) implies

$$\sin^{-1} b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}} > \sin^{-1} \|\mathbf{P} - \mathbf{P}_r\|_\infty + \sin^{-1} b_d, \quad (12.13a)$$

and Theorem 12.1 yields

$$\sin^{-1} b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}} > \sin^{-1} \delta_\nu(\mathbf{P}, \mathbf{P}_r) + \sin^{-1} b_d. \quad (12.13b)$$

Equations (12.5b, 12.13b) reduce to $b_{\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}} > b_d$. \square

Corollary 12.2 (Weak condition for reduce-then-design stability). *If \mathbf{P}_r is an approximation of \mathbf{P} that satisfies $\|\mathbf{P} - \mathbf{P}_r\|_\infty \leq \beta$ for some $\beta \in [0, 1)$, and $\delta_\nu(\mathbf{P}, \mathbf{P}_r) < 1$, then the condition*

$$b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}} > \beta \quad (12.14)$$

is sufficient for the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$.

Proof. This follows directly from Theorem 12.4, with the choice $b_d = 0$. \square

Next, we derive a stronger bound using Obinata and Anderson (2001, Chapter 3.2), as before.

Theorem 12.5 (Strong condition for reduce-then-design stability). *If $[\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}]$ is stable, \mathbf{P}_r is an approximation of \mathbf{P} that satisfies $\|\mathbf{P} - \mathbf{P}_r\|_\infty \leq \beta$ for some $\beta \in [0, 1)$, and $\eta(\mathbf{P}) = \eta(\mathbf{P}_r)$, then*

$$\left\| (\mathbf{I} - \mathbf{K}_r^{\mathbf{P}} \mathbf{P}_r)^{-1} \mathbf{K}_r^{\mathbf{P}} \right\|_\infty^{-1} > \beta \quad (12.15)$$

is sufficient for the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$.

Proof. We may modify the internal model control structure in Obinata and Anderson (2001, Chapter 3.2) to handle plant perturbations instead of controller perturbations. In this case, if $\mathbf{P}, \mathbf{P}_r \in \mathcal{RL}_\infty$, $[\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}]$ is stable, $\eta(\mathbf{P}) = \eta(\mathbf{P}_r)$, and

$$1 > \left\| (\mathbf{I} - \mathbf{K}_r^{\mathbf{P}} \mathbf{P}_r)^{-1} \mathbf{K}_r^{\mathbf{P}} (\mathbf{P} - \mathbf{P}_r) \right\|_\infty \quad (12.16a)$$

then $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$ is stable. From the stated conditions and (12.15),

$$1 > \left\| (\mathbf{I} - \mathbf{K}_r^{\mathbf{P}} \mathbf{P}_r)^{-1} \mathbf{K}_r^{\mathbf{P}} \right\|_\infty \|\mathbf{P} - \mathbf{P}_r\|_\infty \quad (12.16b)$$

$$\geq \left\| (\mathbf{I} - \mathbf{K}_r^{\mathbf{P}} \mathbf{P}_r)^{-1} \mathbf{K}_r^{\mathbf{P}} (\mathbf{P} - \mathbf{P}_r) \right\|_\infty; \quad (12.16c)$$

therefore, $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$ must be stable. \square

This condition is stronger than Corollary 12.2, because it can be shown using norm properties that $\left\| (\mathbf{I} - \mathbf{K}_r^{\mathbf{P}} \mathbf{P}_r)^{-1} \mathbf{K}_r^{\mathbf{P}} \right\|_\infty^{-1} \geq b_{\mathbf{P}_r, \mathbf{K}_r^{\mathbf{P}}}$. Again, the difference between the two sides of this inequality may not be large in practice; see Figure 12.6 for an example comparison.

Remark. Practically, the parameters α and β are the analytic upper bounds on the model reduction error's \mathcal{H}_∞ norm. For instance, if $\mathbf{K}_r^{\mathbf{K}}$ is the order r balanced truncation of \mathbf{K} , then Theorem 12.3 guarantees the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{K}}]$ when

$$\left\| \mathbf{P}(\mathbf{I} - \mathbf{K}\mathbf{P})^{-1} \right\|_\infty^{-1} > 2 \sum_{k=r+1}^n \sigma_k(\mathbf{K}). \quad (12.17a)$$

Alternatively, if \mathbf{P}_r is the order r balanced truncation of \mathbf{P} , then Theorem 12.5 guarantees the stability of $[\mathbf{P}, \mathbf{K}_r^{\mathbf{P}}]$ when

$$\left\| (\mathbf{I} - \mathbf{K}_r^{\mathbf{P}}\mathbf{P}_r)^{-1} \mathbf{K}_r^{\mathbf{P}} \right\|_\infty^{-1} > 2 \sum_{k=r+1}^n \sigma_k(\mathbf{P}). \quad (12.17b)$$

Remark. The conditions that $\eta(\mathbf{K}) = \eta(\mathbf{K}_r^{\mathbf{K}})$ (Theorem 12.3) and $\eta(\mathbf{P}) = \eta(\mathbf{P}_r)$ (Theorem 12.5) are often satisfied in practice. For instance, when performing balanced reduction or optimal Hankel norm reduction on an unstable system \mathbf{G} , we perform the decomposition $\mathbf{G} = \mathbf{G}_s + \mathbf{G}_a$, with \mathbf{G}_s stable and \mathbf{G}_a anti-stable, and reduce only \mathbf{G}_s . See, for example, the reduce-then-design results of Section 12.4.

12.4 Example: linearized Ginzburg–Landau control

The Ginzburg–Landau equation is a model for fluid flow perturbations, among many other applications. For a detailed discussion, see Chomaz et al. (1988); Cossu and Chomaz (1997); Chomaz (2005). Bagheri et al. (2009) provide a comprehensive review of Ginzburg–Landau control and model reduction, using conventional linear controllers designed from finite-dimensional numerical approximations of the governing PDE. This is the approach we take in the present paper. An alternative approach retains the infinite-dimensional PDE, but is restricted to boundary control; see Aamo et al. (2005), Aamo et al. (2007), Milovanovic and Aamo (2011) and references within, or Smyshlyaev and Krstić (2010), Chapter 6 for a summary of the technique. Our goal here is not to compare these control design approaches, but rather, to illustrate how the techniques of the previous section may be used to provide guaranteed bounds for reduced-order controllers.

The linearized system and control setup that we describe here are from Chen and Rowley (2011), which contains a more complete description. Given a time parameter t , a spatial coordinate x (typically the fluid flow direction), and a complex parameter $q(x, t)$ that represents a velocity or streamfunction perturbation, the linearized Ginzburg–Landau equation is

$$\frac{\partial q}{\partial t} = \mathcal{L}q, \quad (12.18a)$$

where

$$\mathcal{L} \triangleq \mu(x) - \nu \frac{\partial}{\partial x} + \gamma \frac{\partial^2}{\partial x^2}. \quad (12.18b)$$

With an infinite spatial domain $x \in (-\infty, \infty)$, the boundary conditions are $q(\pm\infty, t) = 0$ (Cossu and Chomaz, 1997; Chomaz, 2005).

The Ginzburg–Landau equation is a convective-diffusive model with amplification. The amplification function we choose is $\mu(x) = 0.37 - 0.005x^2$. The complex advection speed is $\nu = 2.0 + 0.4j$, and the complex diffusion parameter is $\gamma = 1.0 - 1.0j$. With this choice of parameters, \mathcal{L} has one unstable eigenvalue, at $0.0123 - 0.648j$.

To implement output feedback control of the linearized Ginzburg–Landau equation, we use single-input, single-output control with one localized actuator at $x = x_a$ and one localized sensor at $x = x_s$. For a scalar actuation signal $u(t)$ and a scalar sensor signal $y(t)$, the continuous-space state space is

$$\frac{\partial q}{\partial t}(x, t) = \mathcal{L}q(x, t) + \exp\left(-\frac{(x - x_a)^2}{2\sigma^2}\right) u(t) \quad (12.19a)$$

$$y(t) = \left\langle q(x, t), \exp\left(-\frac{(x - x_s)^2}{2\sigma^2}\right) \right\rangle, \quad (12.19b)$$

where $\langle f_1(x), f_2(x) \rangle \triangleq \int_{-\infty}^{\infty} \bar{f}_2(x) f_1(x) dx$ is a spatial inner product. For this study, we choose $\sigma = 0.4$.

To limit the Ginzburg–Landau model to a finite dimension, we employ a Hermite polynomial pseudospectral method. The software described in Weideman and Reddy (2000) computes the domain discretization $\{x_1, \dots, x_N\}$ and the derivatives. We choose $N = 100$ grid points; when $N = 200$, the results shown in Figs. 12.4–12.7 are nearly unchanged. The domain reaches from $x_1 = -56.1$ to $x_N = 56.1$; the region of amplification, where $\mu(x) > 0$, is $-8.6 < x < 8.6$.

With this discretization scheme, we define the state vector

$$\boldsymbol{\xi}(t) \triangleq \begin{bmatrix} q(x_1, t) & \cdots & q(x_N, t) \end{bmatrix}^T, \quad (12.20)$$

and we denote the pseudospectral discretization of \mathcal{L} by \mathbf{A} . The actuation matrix \mathbf{B} is the discretization of $\exp(-(x - x_a)^2/(2\sigma^2))$. The sensing matrix \mathbf{C} is trickier to compute, because of an implied spatial integration on an uneven grid. Define the trapezoidal integration matrix

$$\mathbf{H} \triangleq \frac{1}{2} \text{diag}(x_2 - x_1, x_3 - x_1, \dots, x_{j+1} - x_{j-1}, \dots, x_N - x_{N-2}, x_N - x_{N-1}). \quad (12.21)$$

With this weight, the discretization of the continuous-space inner product $\langle \mathbf{q}_1, \mathbf{q}_2 \rangle$ is given by $\langle \boldsymbol{\xi}_1, \boldsymbol{\xi}_2 \rangle = \boldsymbol{\xi}_2^* \mathbf{H} \boldsymbol{\xi}_1$. Thus, if \mathbf{f}_s is the discretization of $\exp(-(x - x_s)^2/(2\sigma^2))$, then $\mathbf{C} = \mathbf{f}_s^* \mathbf{H}$. Altogether, we define the Ginzburg–Landau plant

$$P : \begin{bmatrix} \dot{\boldsymbol{\xi}} \\ y \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi} \\ u \end{bmatrix}. \quad (12.22)$$

We implement a nearly optimal \mathcal{H}_∞ loop-shaping controller, using the negative-feedback \mathcal{H}_2 optimal controller as the plant weight w . As in our previous work (Chen and Rowley, 2011), we choose the state cost matrix $\mathbf{Q} = 49\mathbf{H}$ and the input cost matrix $R = 1$ for the \mathcal{H}_2 optimal controller. In addition, we choose the state disturbance covariance $\mathbf{W} = \mathbf{H}$, corresponding to a white noise uncorrelated and evenly distributed in time and space, and we choose the sensor noise covariance $V = 4 \cdot 10^{-8}$. These matrices, along with P , determine the \mathcal{H}_2 optimal controller for a given actuator and sensor location (Skogestad and Postlethwaite, 2005, Ch. 9.2).

We place the actuator and sensor at the \mathcal{H}_2 optimal placement—that is, the placement which minimizes the \mathcal{H}_2 norm from state disturbances and sensor noise to the cost on the state and input sizes, when P is in closed-loop with the \mathcal{H}_2 optimal controller. (See Chen and Rowley (2011) for complete details.) This occurs at $x_a = -1.0$ and $x_s = 1.0$. Setting the weight w equal to the negative-feedback \mathcal{H}_2 optimal controller at this placement, we denote the weighted plant by $P_w = Pw$.




Figure 12.3: Bode plots of the linearized complex Ginzburg–Landau input-output dynamics P and the weighted dynamics P_w .

Figure 12.3 shows the Bode plots of P and P_w . From a classical control standpoint, we see that the weight increases the loop gain at low frequencies, where we require disturbance rejection, and decreases the loop gain at high frequencies, where we require noise attenuation. Furthermore, the weighted plant has ample stability margins. Its Nyquist plot encircles the -1 point once, since it has one unstable pole, and it comes no closer than 0.61 units to the -1 point.

The true \mathcal{H}_∞ loop-shaping controller is $K_{\text{opt}} \in \mathcal{R}$ such that $b_{P_w, K_{\text{opt}}} = \sup_{K \in \mathcal{R}} b_{P_w, K}$. To simplify the computation of the controller, we first compute $\sup_{K \in \mathcal{R}} b_{P_w, K}$; then, we assemble a suboptimal controller K , where the target value of $b_{P_w, K}$ is $1/(1 + 10^{-5})$ times the optimal value (see Skogestad and Postlethwaite, 2005, Chapter 9.4.1).

In this example, we use balanced truncation to approximate the plant and controller. Therefore, we apply Theorems 12.2 and 12.3 with $\alpha = 2 \sum_{k=r+1}^n \sigma_k(K)$. In the reduce-then-design case (Theorems 12.4 and 12.5), the computation is more complicated because P_w has one right-half-plane pole. Therefore, we split P_w into two additive parts, with one stable and one anti-stable; we only reduce the stable part (see Remark 12.3). As a result, the lowest reduction order is $r = 1$. Furthermore, since $\sigma_k(P_w)$ only includes the Hankel singular values of the stable part, we use $\beta = 2 \sum_{k=r}^{n-1} \sigma_k(P_w)$.

In the design-then-reduce case, $b_{P_w, K} = 5.10 \cdot 10^{-1}$ and $\|P_w(I - KP_w)^{-1}\|_\infty^{-1} = 5.91 \cdot 10^{-1}$. Figure 12.4 depicts this, along with the ν -gap metric and the error upper bound α . From this figure, we observe that Theorem 12.3 guarantees the stability of $[P_w, K_r^K]$ when $r \geq 2$. If we directly apply the stability margin and ν -gap relation in (12.5d), then we also derive $r \geq 2$. In reality, $b_{P_w, K_r^K} > 0$, and hence $[P_w, K_r^K]$ is stable, when $r \geq 1$.

If the desired minimal performance is $b_d = 0.4$, then Theorem 12.2 guarantees this performance level when $r \geq 4$; see Figure 12.5. The direct application of the ν -gap metric in (12.8b) yields the better guarantee $r \geq 3$. In reality, $b_{P_w, K_r^K} > b_d$ when $r \geq 3$.

Figure 12.6 plots stability parameters for reduce-then-design control, along with the error upper bound β . Theorem 12.5 guarantees the stability of $[P_w, K_r^P]$ when $r \geq 4$. Directly applying the stability margin and ν -gap relation in (12.5b) gives the better guarantee $r \geq 3$. In reality, $b_{P_w, K_r^P} > 0$, and hence $[P_w, K_r^P]$ is stable, when $r \geq 3$.

Keeping the minimal performance specification $b_d = 0.4$, Theorem 12.4 and the direct application of the ν -gap metric in (12.13b) both guarantee this performance level when $r \geq 5$; see




Figure 12.4: Robust stability quantities as the balanced truncation order of the \mathcal{H}_∞ loop-shaping controller K varies. Dashed line: $\|P_w(I - KP_w)^{-1}\|_\infty^{-1}$ (constant); solid line: $b_{P_w, K}$ (constant); \circ : $\alpha = 2 \sum_{k=r+1}^n \sigma_k(K)$; \times : $\delta_\nu(K, K_r^K)$.




Figure 12.5: Robust performance quantities as the balanced truncation order of the \mathcal{H}_∞ loop-shaping controller K varies; here, $b_d = 0.4$. Solid line: $\sin^{-1} b_{P_w, K}$ (constant); \circ : $\sin^{-1} \alpha + \sin^{-1} b_d$; \times : $\sin^{-1} \delta_\nu(K, K_r^K) + \sin^{-1} b_d$.




Figure 12.6: Robust stability quantities as the balanced truncation order of the weighted plant P_w varies. Dashed line: $\|(I - K_r^P P_{wr})^{-1} K_r^P\|_\infty^{-1}$, where P_{wr} is the order r balanced truncation of P_w ; solid line: b_{P_{wr}, K_r^P} ; \circ : $\beta = 2 \sum_{k=r}^{n-1} \sigma_k(P_w)$; \times : $\delta_\nu(P_w, P_{wr})$.

Figure 12.7. This bound is actually tight; we verify that in reality, the closed-loop indeed meets this performance specification only for $r \geq 5$.

Finally, we remark briefly that for $r \leq 13$, design-then-reduce yields a stability margin b_{P, K_r^K} greater than the reduce-then-design stability margin b_{P, K_r^P} . At $r > 13$, the difference between these margins is negligible. This result is consistent with Obinata and Anderson (2001), which argues that effective control is best preserved by delaying the approximation step (i.e., the model reduction) as late in the design process as possible. Nonetheless, we must keep in mind that the




Figure 12.7: Robust performance quantities as the balanced truncation order of the weighted plant P_w varies; here, $b_d = 0.4$. Solid line: $\sin^{-1} b_{P_{wr}, K_r^P}$; \circ : $\sin^{-1} \beta + \sin^{-1} b_d$; \times : $\sin^{-1} \delta_\nu(P_w, P_{wr}) + \sin^{-1} b_d$.

design-then-reduce procedure is generally more computationally expensive than the reduce-then-design procedure, especially for large systems.

12.5 Conclusion

In the application of control theory to large systems, it is typically necessary to use a reduced-order controller \mathbf{K}_r in closed-loop with the full-order plant \mathbf{P} . This study presents conditions that are sufficient for guaranteeing the stability or required performance of $[\mathbf{P}, \mathbf{K}_r]$, whether \mathbf{K}_r is a reduced-order model of a full-order controller \mathbf{K} , or a controller designed from a reduced-order plant \mathbf{P}_r . The conditions do not place restrictions on the control design or model reduction, besides that the model reduction needs to have an upper bound on the \mathcal{H}_∞ norm of the error.

Theorems 12.3 and 12.5 state the stability conditions, and Theorems 12.2 and 12.4 state the performance conditions. The analytic bounds on the normalized coprime stability margin, the ν -gap metric, and model reduction error motivate the performance results. Although the performance guarantees yield analogous stability guarantees, we find stronger bounds by employing the internal-model-control-like structure in Obinata and Anderson (2001, Chapter 3.2).

A number of model reduction methods have known analytic upper bounds on the \mathcal{H}_∞ norm of the reduction error (e.g., modal truncation, balanced truncation, and optimal Hankel norm reduction). This \mathcal{H}_∞ norm, by itself, is not the best measure of robust stability—for robust stability, we only require that the model reduction be accurate near the crossover frequency. Therefore, many model reduction methods—particularly those that attempt to minimize the \mathcal{H}_∞ norm of the error—will tend to fit the reduced model to the original model over a larger frequency range than is necessary for robust stability.

Nevertheless, if the \mathcal{H}_∞ norm of the model reduction error is bounded from above, then the ν -gap metric shares the same upper bound (see Theorem 12.1). If the normalized coprime stability margin between the full-order plant and controller—or between the reduced-order plant and controller—is greater than the bound on the ν -gap metric, then the full-order plant will provably be stable in closed-loop with the reduced-order controller. We may conclude this without actually testing the closed-loop system. Furthermore, the bounds indicate how accurate a model approximation needs to be, in the sense of the error's \mathcal{H}_∞ norm, to be sufficient for guaranteeing normalized coprime robust stability or performance. This aids the control designer in choosing a model reduction order.

We demonstrate these sufficient conditions on the control of the linearized Ginzburg–Landau equation. For both the design-then-reduce and the reduce-then-design approaches, the conditions

correctly guarantee values of the controller order r for which $[\mathbf{P}, \mathbf{K}_r]$ is stable. Given a desired performance level, they also correctly guarantee the values of r for which $[\mathbf{P}, \mathbf{K}_r]$ meets that performance level. There generally exist lower values of r for which $[\mathbf{P}, \mathbf{K}_r]$ is stable or has satisfactory performance, but the theorems are *a priori* unable to guarantee such. Nonetheless, the bounds provided by the theorems and corollaries are fairly tight in this example.

These results are directly implementable in the control of large systems. For instance, the BPOD and ERA techniques are applicable to computer simulations of a linearized fluid flow, yielding approximations of the linearized dynamics' balanced truncation and Hankel singular values (Rowley, 2005; Ma et al., 2011). Theorem 12.5 can then predict the plant and controller reduction orders for which the reduced-order controller is guaranteed to stabilize the flow. Additionally, Theorem 12.4 can predict the orders for which the closed-loop system is guaranteed to achieve a desired performance level. This would be beneficial for the control designer, since fluid flows—and indeed, many other real-life systems—can be remarkably difficult to stabilize and control well.

The Department of Defense (DOD) supported this work through the National Defense Science & Engineering Graduate (NDSEG) fellowship.

Bibliography

- O. M. Aamo, A. Smyshlyaev, and M. Krstić. Boundary control of the linearized Ginzburg–Landau model of vortex shedding. *SIAM J. Control Optim.*, 43(6):1953–1971, 2005.
- O. M. Aamo, A. Smyshlyaev, M. Krstić, and B. A. Foss. Output feedback boundary control of a Ginzburg–Landau model of vortex shedding. *IEEE Trans. Autom. Control*, 52(4):742–748, Apr. 2007.
- E. Åkervik, L. Brandt, D. S. Henningson, J. Hoepffner, O. Marxen, and P. Schlatter. Steady solutions of the Navier–Stokes equations by selective frequency damping. *Phys. Fluids*, 18(6):068102, Jun. 2006.
- E. Åkervik, J. Hoepffner, U. Ehrenstein, and D. S. Henningson. Optimal growth, model reduction and control in a separated boundary-layer flow using global eigenmodes. *J. Fluid Mech.*, 579:305–314, May 2007.
- L. Armijo. Minimization of functions having Lipschitz continuous first partial derivatives. *Pac. J. Math.*, 16(1):1–3, Nov. 1966.
- J. A. Atwell and B. B. King. Reduced order controllers for spatially distributed systems via proper orthogonal decomposition. *SIAM J. Sci. Comput.*, 26(1):128–151, 2004.
- N. Aubry, W. Y. Lian, and E. S. Titi. Preserving symmetries in the proper orthogonal decomposition. *SIAM J. Sci. Comput.*, 14(2):483–505, Mar. 1993.
- S. Bagheri. Effects of small noise on the DMD/Koopman spectrum. Presented at the 66th Annual Meeting of the American Physical Society Division of Fluid Dynamics, Nov. 2013.
- S. Bagheri, D. S. Henningson, J. Hoepffner, and P. J. Schmid. Input–output analysis and control design applied to a linear model of spatially developing flows. *Appl. Mech. Rev.*, 62(2):020803, Mar. 2009.
- B. Bamieh and M. Dahleh. Energy amplification in channel flows with stochastic excitation. *Phys. Fluids*, 13(11):3258–3269, Nov. 2001.
- B. A. Batten and K. A. Evans. Reduced-order compensators via balancing and central control design for a structural control problem. *Int. J. Control*, 83(3):563–574, Mar. 2010.
- B. J. Bayly. Three-dimensional centrifugal-type instabilities in inviscid two-dimensional flows. *Phys. Fluids*, 31(1):56–64, Jan. 1988.
- B. A. Belson, O. Semeraro, C. W. Rowley, and D. S. Henningson. Feedback control of instabilities in the two-dimensional Blasius boundary layer: The role of sensors and actuators. *Phys. Fluids*, 25(5):054106, May 2013.

- T. R. Bewley, P. Luchini, and J. O. Pralits. Methods for solution of large optimal control problems that bypass open-loop model reduction. Submitted for publication, 2011. URL http://fccr.ucsd.edu/pubs/BLP_MCE_ADA_OSSI.pdf.
- J. A. Burns and B. B. King. A reduced basis approach to the design of low-order feedback controllers for nonlinear continuous systems. *J. Vib. Control*, 4(3):297–323, May 1998.
- K. M. Butler and B. F. Farrell. Three-dimensional optimal perturbations in viscous shear-flow. *Phys. Fluids A*, 4(8):1637–1650, Aug. 1992.
- L. N. Cattafesta and M. Sheplak. Actuators for active flow control. *Annu. Rev. Fluid Mech.*, 43:247–272, 2011.
- P. Chakraborty, S. Balachandar, and R. J. Adrian. On the relationships between local vortex identification schemes. *J. Fluid Mech.*, 535:189–214, Jul. 2005.
- K. K. Chen and C. W. Rowley. \mathcal{H}_2 optimal actuator and sensor placement in the linearised complex Ginzburg–Landau system. *J. Fluid Mech.*, 681:241–260, Aug. 2011.
- K. K. Chen and C. W. Rowley. Normalized coprime robust stability and performance guarantees for reduced-order controllers. *IEEE Trans. Autom. Control*, 58(4):1068–1073, Apr. 2013.
- K. K. Chen and C. W. Rowley. Fluid flow control applications of \mathcal{H}_2 optimal actuator and sensor placement. In *Proceedings of the American Control Conference (ACC)*, Portland, OR, USA, Jun. 2014. American Automatic Control Council.
- K. K. Chen, J. H. Tu, and C. W. Rowley. Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses. *J. Nonlinear Sci.*, 22(6):887–915, Dec. 2012.
- H. Choi, W.-P. Jeon, and J. Kim. Control of flow over a bluff body. *Annu. Rev. Fluid Mech.*, 40:113–139, 2008.
- H. Choi, J. Lee, and H. Park. Aerodynamics of heavy vehicles. *Annu. Rev. Fluid Mech.*, 46, 2014.
- J.-M. Chomaz. Global instabilities in spatially developing flows: non-normality and nonlinearity. *Annu. Rev. Fluid Mech.*, 37:357–392, 2005.
- J.-M. Chomaz, P. Huerre, and L. G. Redekopp. Bifurcations to local and global modes in spatially developing flows. *Phys. Rev. Lett.*, 60(1):25–28, Jan. 1988.
- J.-M. Chomaz, P. Huerre, and L. G. Redekopp. A frequency selection criterion in spatially developing flows. *Stud. Appl. Math.*, 84(2):119–144, Feb. 1991.
- K. Cohen, S. Siegel, D. Wetlesen, J. Cameron, and A. Sick. Effective sensor placements for the estimation of proper orthogonal decomposition mode coefficients in von Kármán vortex street. *J. Vib. Control*, 10(12):1857–1880, Dec. 2004.
- C. Colburn, D. Zhang, and T. Bewley. Gradient-based optimization methods for sensor & actuator placement in LTI systems. Submitted for publication, May 2011. URL http://fccr.ucsd.edu/pubs/CZB_placement.pdf.
- T. Colonius and K. Taira. A fast immersed boundary method using a nullspace approach and multi-domain far-field boundary conditions. *Comput. Meth. Appl. Mech. Eng.*, 197(25–28):2131–2146, Apr. 2008.

- C. Cossu and J.-M. Chomaz. Global measures of local convective instabilities. *Phys. Rev. Lett.*, 78(23):4387–4390, Jun. 1997.
- N. Darivandi, K. Morris, and A. Khajepour. An algorithm for LQ optimal actuator location. *Smart Mater. Struct.*, 22(3):035001, Mar. 2013.
- V. R. Dehkordi and B. Boulet. Robust controller order reduction. *Int. J. Control*, 84(5):985–997, May 2011.
- D. Doorly and S. Sherwin. *Cardiovascular Mathematics*, volume 1 of *Modeling, Simulation and Applications*, chapter 5. Springer, 2009.
- A. P. Dowling and A. S. Morgans. Feedback control of combustion oscillations. *Annu. Rev. Fluid Mech.*, 37:151–182, 2005.
- J. C. Doyle, K. Glover, P. P. Khargonekar, and B. A. Francis. State-space solutions to standard \mathcal{H}_2 and \mathcal{H}_∞ control problems. *IEEE Trans. Autom. Control*, 34(8):831–847, Aug. 1989.
- B. F. Farrell and P. J. Ioannou. Stochastic forcing of the linearized Navier–Stokes equations. *Phys. Fluids A*, 5(11):2600–2609, Nov. 1993.
- J. H. Ferziger and M. Perić. *Computational Methods for Fluid Dynamics*. Springer–Verlag, 3rd edition, 2002.
- H. Gao, J. Lam, and C. Wang. Controller reduction with \mathcal{H}_∞ error performance: Continuous- and discrete-time cases. *Int. J. Control*, 79(6):604–616, Jun. 2006.
- J. C. Geromel. Convex analysis and global optimization of joint actuator location and control problems. *IEEE Trans. Autom. Control*, 34(7):711–720, Jul. 1989.
- F. Giannetti and P. Luchini. Structural sensitivity of the first instability of the cylinder wake. *J. Fluid Mech.*, 581:167–197, Jun. 2007.
- F. Giannetti, S. Camarri, and P. Luchini. Structural sensitivity of the secondary instability in the wake of a circular cylinder. *J. Fluid Mech.*, 651:319–337, May 2010.
- E. A. Gillies. Low-dimensional control of the circular cylinder wake. *J. Fluid Mech.*, 371:157–178, Sep. 1998.
- E. A. Gillies. Multiple sensor control of vortex shedding. *AIAA J.*, 39(4):748–750, Apr. 2001.
- K. Glover. All optimal Hankel-norm approximations of linear-multivariable systems and their \mathcal{L}^∞ -error bounds. *Int. J. Control*, 39(6):1115–1193, 1984.
- K. Glover and D. McFarlane. Robust stabilization of normalized coprime factor plant descriptions with \mathcal{H}_∞ -bounded uncertainty. *IEEE Trans. Autom. Control*, 34(8):821–830, Aug. 1989.
- P. J. Goddard and K. Glover. Controller approximation: Approaches for preserving \mathcal{H}_∞ performance. *IEEE Trans. Autom. Control*, 43(7):858–871, Jul. 1998.
- R. Govindarajan and K. C. Sahu. Instabilities in viscosity-stratified flow. *Annu. Rev. Fluid Mech.*, 46:331–353, 2014.
- M. Güney and E. Eşkinat. Optimal actuator and sensor placement in flexible structures using closed-loop criteria. *J. Sound Vib.*, 312(1–2):210–233, Apr. 2008.

- D. Halim and S. O. R. Moheimani. An optimization approach to optimal placement of collocated piezoelectric actuators and sensors on a thin plate. *Mechatronics*, 13(1):27–47, Feb. 2003.
- M. G. Hall. Vortex breakdown. *Ann. Rev. Fluid Mech.*, 4:195–218, 1972.
- D. Haller, P. Woias, and N. Kockmann. Simulation and experimental investigation of pressure loss and heat transfer in microchannel networks containing bends and T-junctions. *Int. J. Heat Mass Transf.*, 52(11–12):2678–2689, May 2009.
- M. J. M. Hill. On a spherical vortex. *Philos. Trans. R. Soc. A*, 185:213–245, 1894.
- K. Hiramoto, H. Doki, and G. Obinata. Optimal sensor/actuator placement for active vibration control using explicit solution of algebraic Riccati equation. *J. Sound Vib.*, 229(5):1057–1075, Feb. 2000.
- P. Holmes, J. L. Lumley, and G. Berkooz. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, Cambridge, U.K., 1996.
- P. Huerre and P. A. Monkewitz. Local and global instabilities in spatially developing flows. *Annu. Rev. Fluid Mech.*, 22:473–537, 1990.
- M. Ilak and C. W. Rowley. Modeling of transitional channel flow using balanced proper orthogonal decomposition. *Phys. Fluids*, 20(3):034103, Mar. 2008.
- B. L. Jones and E. C. Kerrigan. When is the discretization of a spatially distributed system good enough for control? *Automatica*, 46(9):1462–1468, Sep. 2010.
- M. R. Jovanović and B. Bamieh. Componentwise energy amplification in channel flows. *J. Fluid Mech.*, 534:145–183, Jul. 2005.
- M. R. Jovanović, P. J. Schmid, and J. W. Nichols. Sparsity-promoting dynamic mode decomposition. *Phys. Fluids*, 26:024103, Feb. 2014.
- J.-N. Juang and R. S. Pappa. An eigensystem realization algorithm for modal parameter identification and model reduction. *J. Guid. Contr. Dyn.*, 8(5):620–627, Sep.–Oct. 1985.
- T. Kailath. Some new algorithms for recursive estimation in constant linear systems. *IEEE Trans. Inf. Theory*, 19(6):750–760, Nov. 1973.
- C. T. Kelley. *Iterative Methods for Linear and Nonlinear Equations*, volume 16 of *Frontiers in Applied Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1995.
- C. T. Kelley. *Solving Nonlinear Equations with Newton’s Method*, volume 1 of *Fundamentals of Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2003.
- J. Kim. Control of turbulent boundary layers. *Phys. Fluids*, 15(5):1093–1105, May 2003.
- J. Kim and T. R. Bewley. A linear systems approach to flow control. *Annu. Rev. Fluid Mech.*, 39:383–417, 2007.
- B. B. King and E. W. Sachs. Semidefinite programming techniques for reduced order systems with guaranteed stability margins. *Comput. Optim. Appl.*, 17(1):37–59, Oct. 2000.
- N. Kockmann, M. Engler, D. Haller, and P. Woias. Fluid dynamics and transfer processes in bended microchannels. *Heat Transf. Eng.*, 26(3):71–78, Apr. 2005.

- S. Kondoh, C. Yatomi, and K. Inoue. The positioning of sensors and actuators in the vibration control of flexible systems. *JSME Int. J., Series III*, 33(2):145–152, Jun. 1990.
- M. T. Landahl. A wave-guide model for turbulent shear flow. *J. Fluid Mech.*, 29(3):441–459, Sep. 1967.
- M. T. Landahl. Wave breakdown and turbulence. *SIAM J. Appl. Math.*, 28(4):735–756, 1975.
- E. Lauga and T. R. Bewley. The decay of stabilizability with Reynolds number in a linear model of spatially developing flows. *Proc. R. Soc. A*, 459(2036):2077–2095, Aug. 2003.
- E. Lauga and T. R. Bewley. Performance of a linear robust control strategy on a nonlinear model of spatially developing flows. *J. Fluid Mech.*, 512:343–374, Aug. 2004.
- P. Le Gal and V. Croquette. Visualization of the space–time impulse response of the subcritical wake of a cylinder. *Phys. Rev. E*, 62(3):4424–4426, Sep. 2000.
- A. Lifschitz and E. Hameiri. Local stability conditions in fluid dynamics. *Phys. Fluids A*, 3(11):2644–2651, Nov. 1991.
- Y. Lou and P. D. Christofides. Optimal actuator/sensor placement for nonlinear control of the Kuramoto–Sivashinsky equation. *IEEE Trans. Control Syst. Tech.*, 11(5):737–745, Sep. 2003.
- P. Luchini and A. Bottaro. Adjoint equations in stability analysis. *Annu. Rev. Fluid Mech.*, 46:493–517, 2014.
- P. Luchini, F. Giannetti, and V. Citro. Short-wave analysis of 3D and 2D instabilities in a driven cavity. Presented at the 66th Annual Meeting of the American Physical Society Division of Fluid Dynamics, Pittsburgh, PA, USA, Nov. 2013.
- J. Lumley and P. Blossey. Control of turbulence. *Annu. Rev. Fluid Mech.*, 30(1):311–327, 1998.
- Z. Ma, S. Ahuja, and C. W. Rowley. Reduced-order models for control of fluids using the eigensystem realization algorithm. *Theor. Comput. Fluid Dyn.*, 25(1–4):233–247, Jun. 2011.
- C. K. Mamun and L. S. Tuckerman. Asymmetry and Hopf bifurcation in spherical Couette flow. *Phys. Fluids*, 7(1):80–91, Jan. 1995.
- O. Marquet, D. Sipp, and L. Jacquin. Sensitivity analysis and passive control of cylinder flow. *J. Fluid Mech.*, 615:221–252, Nov. 2008.
- D. McFarlane and K. Glover. A loop shaping design procedure using \mathcal{H}_∞ synthesis. *IEEE Trans. Autom. Control*, 37(6):759–769, June 1992.
- D. McFarlane, K. Glover, and M. Vidyasagar. Reduced-order controller design using coprime factor model reduction. *IEEE Trans. Autom. Control*, 35(3):369–373, Mar 1990.
- D. G. Meyer and G. F. Franklin. A connection between normalized coprime factorizations and linear quadratic regulator theory. *IEEE Trans. Autom. Control*, 32(3):227–228, Mar. 1987.
- I. Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dyn.*, 41(1–3):309–325, Aug. 2005.
- I. Mezić. Analysis of fluid flows via spectral properties of the koopman operator. *Annu. Rev. Fluid Mech.*, 45:357–378, 2013.

- I. Mezić and A. Banaszuk. Comparison of systems with complex behavior. *Physica D*, 197(1–2): 101–133, Oct. 2004.
- M. Milovanovic and O. M. Aamo. Stabilization of 3D Ginzburg–Landau equation by model-based output feedback control. In *18th IFAC World Congress*, pages 14435–14439, Milano, Italy, Aug.–Sep. 2011. International Federation of Automatic Control.
- S. O. R. Moheimani and T. Ryall. Considerations on placement of piezoceramic actuators that are used in structural vibration control. In *Proceedings of the 38th Conference on Decision & Control*, pages 1118–1123. Institute of Electrical and Electronics Engineers, Dec. 1999.
- B. C. Moore. Principal component analysis in linear systems: controllability, observability, and model reduction. *IEEE Trans. Autom. Control*, 26(1):17–32, Feb 1981.
- K. Morris. Linear-quadratic optimal actuator location. *IEEE Trans. Autom. Control*, 56(1):113–124, Jan. 2011.
- B. R. Noack, K. Afanasiev, M. Morzyński, G. Tadmor, and F. Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.*, 497:335–363, Dec. 2003.
- B. R. Noack, P. Papas, and P. A. Monkewitz. The need for a pressure-term representation in empirical Galerkin models of incompressible shear flow. *J. Fluid Mech.*, 523:339–365, Jan. 2005.
- W. I. Norwood, J. K. Kirklin, and S. P. Sanders. Hypoplastic left heart syndrome: experience with palliative surgery. *Am. J. Cardiol.*, 45(1):87–91, Jan. 1980.
- G. Obinata and B. D. O. Anderson. *Model Reduction for Control System Design*. Springer, London, U.K., 2001.
- K. Petersen. *Ergodic Theory*. Cambridge University Press, 1983.
- J. O. Pralits and P. Luchini. *Seventh IUTAM Symposium on Laminar–Turbulent Transition*, pages 325–330. Springer, 2010.
- W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes*. Cambridge University Press, 3rd edition, 2007.
- M. Provansal, C. Mathis, and L. Boyer. Bénard–von Kármán instability: transient and forced regimes. *J. Fluid Mech.*, 182:1–22, Jan. 1987.
- 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. R. Soc. London*, 35:84–99, 1883.
- A. H. Ropper and M. A. Samuels. *Adams and Victor’s Principles of Neurology*, chapter 34. McGraw–Hill, 9th edition, 2009.
- A. Roshko. On the development of turbulent wakes from vortex streets. Technical Report NACA 1191, National Advisory Committee for Aeronautics, 1954.
- K. Roussopoulos. Feedback control of vortex shedding at low Reynolds numbers. *J. Fluid Mech.*, 248:267–296, Mar. 1993.

- K. Roussopoulos and P. A. Monkewitz. Nonlinear modelling of vortex shedding control in cylinder wakes. *Physica D*, 97(1–3):264–273, Oct. 1996.
- C. W. Rowley. Model reduction for fluids, using balanced proper orthogonal decomposition. *Int. J. Bifurc. Chaos*, 15(3):997–1013, Mar. 2005.
- C. W. Rowley and D. R. Williams. Dynamics and control of high-reynolds-number flow over open cavities. *Annu. Rev. Fluid Mech.*, 38:251–276, 2006.
- C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson. Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 641:115–127, Dec. 2009.
- Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving non-symmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856–869, Jul. 1986.
- M. G. Safonov and R. Y. Chiang. A Schur method for balanced-truncation model reduction. *IEEE Trans. Autom. Control*, 34(7):729–733, Jul. 1989.
- M. G. Safonov, R. Y. Chiang, and D. J. N. Limebeer. Optimal Hankel model reduction for non-minimal systems. *IEEE Trans. Autom. Control*, 35(4):496–502, Apr. 1990.
- P. Schmid and J. Sesterhenn. Dynamic Mode Decomposition of numerical and experimental data. Presented at the 61st Annual Meeting of the American Physical Society Division of Fluid Dynamics, Nov. 2008.
- P. J. Schmid. Nonmodal stability theory. *Annu. Rev. Fluid Mech.*, 39:129–162, 2007.
- P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.*, 656:5–28, Aug. 2010.
- P. J. Schmid. Application of the dynamic mode decomposition to experimental data. *Exp. Fluids*, 50(4):1123–1130, Apr. 2011.
- P. J. Schmid and D. S. Henningson. *Stability and Transition in Shear Flows*. Springer, 2000.
- P. J. Schmid, L. Li, M. P. Juniper, and O. Pust. Applications of the dynamic mode decomposition. *Theor. Comput. Fluid Dyn.*, 25(1–4):249–259, Jun. 2011.
- O. Semeraro, J. O. Pralits, C. W. Rowley, and D. S. Henningson. Riccati-less approach for optimal control and estimation: an application to two-dimensional boundary layers. *J. Fluid Mech.*, 731, Sep. 2013.
- S. M. Senn and D. Poulikakos. Laminar mixing, heat transfer and pressure drop in tree-like microchannel nets and their application for thermal management in polymer electrolyte fuel cells. *J. Power Sources*, 130(1–2):178–191, May 2004.
- A. S. Sharma, J. F. Morrison, B. J. Mckeen, D. J. N. Limebeer, W. H. Koberg, and S. J. Sherwin. Relaminarisation of $Re_\tau = 100$ channel flow with globally stabilising linear feedback control. *Phys. Fluids*, 23(12), Dec. 2011.
- L. Sirovich. Turbulence and the dynamics of coherent structures. *Q. Appl. Math.*, 45(3):561–590, Oct. 1987.

- S. Skogestad and I. Postlethwaite. *Multivariable Feedback Control: Analysis and Design*. John Wiley & Sons Ltd, West Sussex, U.K., 2nd edition, 2005.
- A. Smyshlyaev and M. Krstić. *Adaptive Control of Parabolic PDEs*. Princeton University Press, Princeton, NJ, USA, 2010.
- F. Sotiropoulos, Y. Ventikos, and T. C. Lackey. Chaotic advection in three-dimensional stationary vortex-breakdown bubbles: Šil’nikov’s chaos and the devil’s staircase. *J. Fluid Mech.*, 444:257–297, Oct. 2001.
- K. R. Sreenivasan, P. J. Strykowski, and D. J. Olinger. Hopf bifurcation, Landau equation, and vortex shedding behind circular cylinders. In K. N. Ghia, editor, *Proceedings of the Forum on Unsteady Flow Separation*, pages 1–13. American Society of Mechanical Engineers, Jun. 1987.
- P. J. Strykowski and K. R. Sreenivasan. On the formation and suppression of vortex ‘shedding’ at low reynolds numbers. *J. Fluid Mech.*, 218:71–107, Sep. 1990.
- G. Tadmor, O. Lehmann, B. R. Noack, and M. Morzyński. Mean field representation of the natural and actuated cylinder wake. *Phys. Fluids*, 22(3):034102, Mar. 2010.
- K. Taira and T. Colonius. The immersed boundary method: a projection approach. *J. Comput. Phys.*, 225(2):2118–2137, Aug. 2007.
- C. S. Tan, I. Day, S. Morris, and A. Wadia. Spike-type compressor stall inception, detection, and control. *Annu. Rev. Fluid Mech.*, 42:275–300, 2010.
- V. Theofilis. Global linear instability. *Annu. Rev. Fluid Mech.*, 43:319–352, 2011.
- P. T. Tokumaru and P. E. Dimotakis. Rotary oscillation control of a cylinder wake. *J. Fluid Mech.*, 224:77–90, Mar. 1991.
- L. N. Trefethen and D. Bau, III. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.
- L. N. Trefethen and M. Embree. *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press, Princeton, NJ, USA, 2005.
- L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. Hydrodynamic stability without eigenvalues. *Science*, 261(5121):578–584, Jul. 1993.
- J. H. Tu and C. W. Rowley. An improved algorithm for balanced pod through an analytic treatment of impulse response tails. *J. Comput. Phys.*, 231(16):5317–5333, Jun. 2012.
- M. van de Wal and B. de Jager. A review of methods for input/output selection. *Automatica*, 37(4):487–510, Apr. 2001.
- M. Vidyasagar. The graph metric for unstable plants and robustness estimates for feedback stability. *IEEE Trans. Autom. Control*, 29(5):403–417, May 1984.
- M. Vidyasagar. Normalized coprime factorizations for nonstrictly proper systems. *IEEE Trans. Autom. Control*, 33(3):300–301, Mar. 1988.
- D. Vigolo, I. M. Griffiths, S. Radl, and H. A. Stone. An experimental and theoretical investigation of particle-wall impacts in a T-junction. *J. Fluid Mech.*, 727:236–255, Jul. 2013.

- D. Vigolo, S. Radl, and Howard A. Stone. Unexpected trapping of particles at a bifurcation. *Proc. Natl. Acad. Sci.*, 2014. Accepted for publication.
- G. Vinnicombe. Frequency domain uncertainty and the graph topology. *IEEE Trans. Autom. Control*, 38(9):1371–1383, Sep. 1993.
- G. Vinnicombe. *Uncertainty and Feedback: \mathcal{H}_∞ Loop-shaping and the ν -gap Metric*. Imperial College Press, London, U.K., 2001.
- J.-Z. Wang and L. Huang. Controller order reduction with guaranteed performance via coprime factorization. *Int. J. Robust Nonlin. Control*, 13(6):501–517, May 2003.
- J. A. C. Weideman and S. C. Reddy. A MATLAB differentiation matrix suite. *ACM Trans. Math. Softw.*, 26(4):465–519, Dec. 2000.
- F. M. White. *Viscous Fluid Flow*. McGraw–Hill, Inc., 3rd edition, 2005.
- C. H. K. Williamson. Vortex dynamics in the cylinder wake. *Annu. Rev. Fluid Mech.*, 28:477–539, 1996.
- K. H. Winters. A bifurcation study of laminar flow in a curved tube of rectangular cross-section. *J. Fluid Mech.*, 180:343–369, Jul. 1987.
- P. K. Yeung and K. R. Sreenivasan. Progress and opportunities in direct numerical simulations at the next higher resolution. Presented at the 66th Annual Meeting of the American Physical Society Division of Fluid Dynamics, Pittsburgh, PA, USA, Nov. 2013.
- K. Zhou and J. Chen. Performance bounds for coprime factor controller reductions. *Syst. Control Lett.*, 26(2):119–127, Sep. 1995.
- K. Zhou and J. C. Doyle. *Essentials of Robust Control*. Prentice Hall, Upper Saddle River, NJ, U.S.A., 1998.