


Optimal and Cooperative Control of Vehicle Formations

Thesis by
J. Alexander Fax

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2002

(Defended November 20, 2001)

Acknowledgements

Five years ago, in an act of shocking naivete, I decided to pursue a doctorate at Caltech while living in Los Angeles and raising a family. That I emerged with a diploma in hand and my sanity (relatively) intact is due to the help, support and friendship of many people.

Firstly, I wish to thank my advisor, Prof. Richard Murray, for being an outstanding mentor, colleague and friend. In addition to exposing me to his wide range of research interests and helping to guide my own research, he inspired me, by word and example, to define new challenges and pursue them successfully. On more than one occasion, I began a meeting with him frustrated at my own lack of progress, only to emerge newly energized and inspired to continue. His inclusive, unassuming model of leadership did much to help me realize my potential at Caltech. He is truly the model advisor.

I would also like to thank the other members of my committee: Prof. Jerry Marsden, Prof. Joel Burdick, and Prof. Jason Hickey. In particular, I enjoyed working with Prof. Marsden on various research projects, and I benefited greatly from the clarity of vision he brings to his work.

I and my family owe a great debt of gratitude to the women of the ARCS Foundation, whose generosity enabled me to complete my degree while supporting a growing family. Their dedication to communal and national service is a model I will carry with me for many years. They also know how to throw a mean luncheon.

My time in graduate school was made much more enjoyable by the many friends I made along the way: Alfredo Martinez, Pablo Parrilo, and Sergey Pekarsky, with whom I survived the first year and qualifying exams; fellow researchers Dong-Eui Chang, David Chichka, Bill Dunbar, Wang-Sang Koon, and Mark Milam; and my officemate *par excellence*, Kristi Morgansen.

My parents, Gene and Ruth Fax, have always encouraged me to pursue excellence and supported me in my every endeavor. To them I owe my love of knowledge and desire to excel. They have also been a great source of love, support, and understanding for me and my family, despite the great distance separating them from their grandchildren.

Julie's parents, Natan and Bronetta Gruenbaum, have helped me and my family in ways too numerous to mention. I do not know how we would have coped with the demands of school and children without their love and unending support.

I owe the greatest debt of gratitude to my wife, Julie. Julie has been at my

side throughout graduate school, enduring the erratic hours, the sporadic income, and my own graduate student angst, all the while raising a wonderful family which I am blessed to be a part of. Our life together is my greatest joy, and it, above all, is my greatest source of strength.

Optimal and Cooperative Control of Vehicle Formations

by

J. Alexander Fax

In Partial Fulfillment of the
Requirements for the Degree of
Doctor of Philosophy

Abstract

Control of vehicle formations has emerged as a topic of significant interest to the controls community. In applications such as microsatellites and underwater vehicles, formations have the potential for greater functionality and versatility than individual vehicles. In this thesis, we investigate two topics relevant to control of vehicle formations: optimal vehicle control and cooperative control.

The framework of optimal control is often employed to generate vehicle trajectories. We use tools from geometric mechanics to specialize the two classical approaches to optimal control, namely the calculus of variations and the Hamilton-Jacobi-Bellman (HJB) equation, to the case of vehicle dynamics. We employ the formalism of the covariant derivative, useful in geometric representations of vehicle dynamics, to relate variations of position to variations of velocity. When variations are computed in this setting, the evolution of the adjoint variables is shown to be governed by the covariant derivative, thus inheriting the geometric structure of the vehicle dynamics. To simplify the HJB equation, we develop the concept of time scalability enjoyed by many vehicle systems. We employ this property to eliminate time from the HJB equation, yielding a purely spatial PDE whose solution supplies both finite-time optimal trajectories and a time-invariant stabilizing control law.

Cooperation among vehicles in formation depends on intervehicle communication. However, vehicle communication is often subject to disruption, especially in an adversarial setting. We apply tools from graph theory to relate the topology of the communication network to formation stability. We prove a Nyquist criterion

that uses the eigenvalues of the graph Laplacian matrix to determine the effect of the graph on formation stability. We also propose a method for decentralized information exchange between vehicles. This approach realizes a dynamical system that supplies each vehicle with a common reference to be used for cooperative motion. We prove a separation principle that states that formation stability is achieved if the information flow is stable for the given graph and if the local controller stabilizes the vehicle. The information flow can be rendered highly robust to changes in the graph, thus enabling tight formation control despite limitations in intervehicle communication capability.

Contents

1	Introduction	1
1.1	Vehicle Formations	1
1.2	Thesis Outline	3
1.2.1	Optimal Control of Vehicles	4
1.2.2	Cooperative Control of Vehicle Formations	6
1.3	Statement of Contributions	11
2	Optimal Control of Affine Connection Control Systems	13
2.1	Mathematical Preliminaries	14
2.1.1	Theory of Affine Connections	14
2.1.2	Calculus of Variations	19
2.2	Optimal Control of Affine Connection Control Systems	21
2.2.1	Affine Connection Control Systems	21
2.2.2	Cost Functions	22
2.2.3	Problem Statement	22
2.2.4	Problem Solution: Lagrange Multipliers	23
2.2.5	Variations of u : stationarity condition	27
2.2.6	Endpoint Conditions	28
2.3	Examples	29
2.3.1	Splines on Manifolds	29
2.3.2	Planar Rigid Body	30
2.4	Conclusions	32
3	Optimal Control of Time-Scalable Systems	33
3.1	Time Scaling	33
3.1.1	Definitions	33
3.1.2	Properties of Time-Scalable Systems	35
3.2	Optimal Control of Time-Scalable Systems	38

3.2.1	Application of Time-Scalability to the HJB Equation	38
3.2.2	Driftless Systems	39
3.2.3	Affine Connection Control Systems	41
3.2.4	Example: Wheeled Locomotion	42
3.3	Conclusions	47
4	Graph Theory	48
4.1	Introduction	48
4.2	Introductory Graph Theory	49
4.2.1	Basic Definitions	49
4.2.2	Connectivity in Directed Graphs	49
4.3	Algebraic Graph Theory	50
4.3.1	Graph Laplacians	50
4.4	Theory of Nonnegative Matrices	52
4.4.1	Nonnegative Matrices and Graph Theory	52
4.4.2	Perron-Frobenius Theorem for Irreducible Matrices	53
4.4.3	Reducible Matrices	54
4.5	Eigenvalues of Laplacians	55
4.5.1	Elementary Results	55
4.5.2	Additional Results	57
5	Stabilization of Vehicle Formations	59
5.1	Relative Position Control in Vehicle Formations	59
5.1.1	Formation Equations of Motion	60
5.1.2	Commanded Vehicle Offsets	61
5.1.3	Graph Laplacians and Formation Dynamics	62
5.2	Stabilization of Vehicle Formations	62
5.2.1	The Role of The Laplacian in Formation Stability	62
5.2.2	Location of Equilibrium Points	64
5.2.3	Stability Theorems	65
5.2.4	Mixed Absolute/Relative Sensing	66
5.2.5	Example: Double Integrator with Time Delay	66
5.3	Evaluating Formations via Laplacian Eigenvalues	67
5.3.1	Graph Periodicity and Formation Stability	69
5.3.2	Stability of Sparsely Connected Formations	71
5.3.3	Weighted Graphs	73

6	Information Flow in Vehicle Formations	74
6.1	Introduction: A Motivating Example	74
6.2	An Information Flow Paradigm	75
6.2.1	Problem Setup	75
6.2.2	Convergence of the Information Flow Loop	77
6.2.3	Shaping the Information Flow	79
6.2.4	Example	84
6.2.5	Information Flow in Weakly Connected Graphs	84
6.2.6	Mixed Absolute/Relative Sensing	86
6.3	Information Flow in the Loop	86
6.4	Examples	92
6.4.1	Hovercraft Formations	92
6.4.2	Satellite Reconfiguration	93
6.5	Information Flow and String Stability	98
6.6	Conclusions	102
7	Conclusions and Future Work	104
7.1	Thesis Summary	104
7.2	Future Work in Optimal Control of Vehicles	105
7.2.1	Geometry of Optimal Control Equations	105
7.2.2	Approximation of Optimal Trajectories	106
7.2.3	HJB Equation and Sub-Riemannian Geometry	107
7.3	Future Work in Cooperative Control of Vehicle Formations	107
7.3.1	Graph Periodicity and Laplacian Eigenvalues	108
7.3.2	Information-Rich Control	108
7.3.3	Vehicles with Nonlinear Dynamics	109
7.3.4	Information Flow and Software-Enabled Control	110
7.3.5	Vehicle Control Over Networks	111
7.3.6	Other Vehicle Formation Tasks	111
7.4	Optimal Control of Vehicle Formations	113

List of Figures

2.1	Diagram of Planar Rigid Body.	31
3.1	Diagram of Kinematic Cart.	43
3.2	Finite-Time Optimal (left) and Stabilizing Trajectories (right), $x_0 = (2, -3, 4)$	46
3.3	Finite-Time Optimal (left) and Stabilizing Trajectories (right), $x_0 = (0, 0, 4)$	46
4.1	Sample Graph \mathcal{G}	51
4.2	Induced Subgraphs of Components of \mathcal{G}	51
4.3	Graph of Components of \mathcal{G}	51
5.1	Formation Graph.	68
5.2	Formation Nyquist Plot.	68
5.3	Inclusion Regions for $-\lambda(L)$	72
5.4	Inclusion Regions for $-\lambda(L)^{-1}$	72
6.1	Hexagon Acquisition, No Information Flow.	75
6.2	Block Diagram of Information Flow in the Loop.	81
6.3	Information Filter Nyquist Plots.	85
6.4	Information Filter Convergence.	85
6.5	Block Diagram of Information Flow with Feedforward Correction.	89
6.6	Hexagon Acquisition with Info Flow, no Info Pre-Convergence.	94
6.7	Hexagon Acquisition with Info Flow, Info Pre-Convergence.	94
6.8	y -axis Transients of Formation Reconfiguration, no Info Flow.	95
6.9	y -axis Transients of Formation Reconfiguration, with Info Flow.	95
6.10	Target Acquisition, no Info Flow.	96
6.11	Target Acquisition, with Info Flow.	96
6.12	Nyquist Plot for Satellite Info Flow.	99
6.13	Convergence of Satellite y -axis Reference Data.	99

6.14 Satellite Reconfiguration Trajectories.	100
6.15 Information Flow Law Bode Plot, $I - H_1(z), I - H_2(z)$	103

List of Tables

5.1	Sample Graphs, Spectra, and Nyquist Locations.	70
-----	--	----

Chapter 1

Introduction

1.1 Vehicle Formations

Recent technological advances have spurred a broad interest in autonomous, adaptable vehicle formations. The development of powerful control techniques for single vehicles, the explosion in computation and communication capabilities, and the advent of miniaturization technologies have elevated interest in vehicles which can interact autonomously with the environment and other vehicles to perform, in the presence of uncertainty and adversity, tasks beyond the ability of individual vehicles. Research in vehicle formation control is currently progressing in multiple fields. Some examples, along with the controls challenges they pose, are discussed in the paragraphs below.

Microsatellite Clusters In recent years, researchers have begun to consider the advantages of microsatellite clusters over large, complex, single-purpose satellites. The use of satellite clusters has the potential to expand functionality, distribute risk, and reduce cost. One example of a microsatellite cluster which has attracted attention recently is the Air Force's Techsat21 mission, which is investigating the ability of a satellite cluster to perform high-resolution imaging through the distribution of microsatellites along a lattice of points and jointly processing the interferometric data [18]. Through this technique, known as sparse aperture radar (SAR), the cluster can realize an effective antenna larger than can be deployed on a single satellite. Many of NASA's missions planned for the upcoming decades also involve satellite clusters [61].

Successful deployment of satellite clusters faces multiple technological challenges. For example, Earth's gravitational field tends to cause satellite formations

to diverge, and the constraints of fuel and power, present in any satellite mission, are particularly restrictive for microsatellites. For this reason, many researchers in the controls community have focused in recent years on understanding satellite formation dynamics in the vicinity of the Earth and implementation of fuel-efficient control laws [21, 88, 107, 108].

Unmanned Aerial Vehicles Advances in avionics, GPS-based navigation, and flight control techniques have brought unmanned aerial vehicle (UAV) technology to a point where it is routinely used in commercial and military applications, leading to renewed interest in UAV formation flight. Applications of this technology include coordinated military maneuvers [71, 72] and drag reduction via close formation flight [19, 22, 46, 103]. In a battlefield environment, unforeseen threats, electromagnetic countermeasures, and vehicle damage can impact both formation goals and UAV communication capability. The ability of UAVs to reliably exchange information and achieve consensus as to formation goals in an uncertain and adversarial environment is critical for UAV formation flight.

Autonomous Underwater Vehicles Over the last decade, autonomous underwater vehicles (AUVs) have moved from the laboratory to commercial, scientific, and military applications. (See [109, 110] and the references therein for AUV applications and research groups.) Recently, researchers have turned to AUV formations to accomplish more challenging tasks. Potential applications of underwater vehicle formations include oceanographic sampling and minesweeping [30, 49].

Unlike satellites, underwater vehicles face significant uncertainty in their dynamics due to ocean currents. Design constraints such as power limitations and the harsh sea environment have led researchers to consider underwater gliders, which effect motion in the fluid via manipulation of vehicle orientation or center of gravity [47, 92]. Also, underwater communication is very limited compared to above ground, forcing designers to consider carefully AUV communication architectures.

Automated Highway Systems The prospect of computer-controlled automobiles autonomously navigating the nation's highways has intrigued engineers and researchers for decades, and advances in computation and machine vision are bringing this technology closer to fruition. (See [1, 90] and the references therein for an overview of this technology). A predominant concern to autonomous traffic flow is collision avoidance. In that respect, vehicles must be aware not only of

the environment but of the other vehicles as well. The interaction of individual vehicles on the highway induces formation-level dynamics that can exhibit undesirable properties despite the simplicity of local vehicle control laws. Properties such as formation disturbance rejection often depend more heavily on the nature of the information available to each vehicle controller than on the controller itself [95, 96, 106].

While each of these areas poses its own unique challenges, several common threads can be found. In most cases, the vehicles are *dynamically decoupled*, meaning the motion of one does not directly affect the others. Instead, the vehicles are coupled through the task they are trying to accomplish jointly. The tasks must be accomplished in the face of *nontrivial vehicle dynamics*. Decisions must be made by each vehicle using only *limited information* about the other vehicles — information which may be subject to uncertainty and transmission delay. The reaction of a vehicle to other vehicles' motions renders the formation an *interconnected dynamical system* whose behavior depends not only on the individual vehicle dynamics, but on the nature of their interconnection. *Environmental factors* can impact the overall formation goal, the actions of individual vehicles within the formation, and the ability of vehicles to communicate. The overarching goal of the formation is *autonomy*, meaning the ability to accomplish its goal in the face of significant uncertainty without human intervention.

1.2 Thesis Outline

Clearly, vehicle formation control poses many interesting research challenges. This thesis explores two issues within vehicle formation control: *optimal control of vehicles* and *cooperative control of vehicle formations*. Rather than focus on a single application area, we consider vehicle dynamics generally. Both areas of study use tools outside the domain of linear control theory. In the former, we apply tools from *differential geometry* to specialize general results in optimal control theory to the case of vehicle dynamics. In the latter, we use tools from *graph theory* to understand the effect of communication topologies on formation stability and performance.

In the following sections, we motivate these two topics, discuss our research methodology, and review the relevant literature.

1.2.1 Optimal Control of Vehicles

One task a formation may face is reconfiguration, meaning the repositioning of the vehicles relative to one another or to a target. Change in formation goals, loss or damage of a vehicle, and emergence of an external threat are all situations which may necessitate formation reconfiguration. Reconfiguration involves both determination of the desired final configuration and the derivation of trajectories and control policies to bring each vehicle to its desired position. Generation of trajectories is often accomplished by posing the problem as an optimization constrained by vehicle dynamics and possibly by limits on position and control effort. The optimized quantity will be situation dependent. For example, satellite formation reconfiguration demands a fuel-optimal trajectory to preserve mission life and is constrained by the limited thrust available. Indeed, it was the satellite formation reconfiguration problem that originally motivated this research.

Solving optimal control problems is often a challenging mathematical problem. One can often gain insight into the structure of solutions through examination of the underlying dynamics. While numerical approaches have made significant progress in recent years [77, 89], they benefit substantially from good initial guesses as to the optimal solution. Those approximate solutions are often based on an analysis of the dynamics of the underlying system. With that in mind, the first two chapters of this thesis are devoted to exploring optimal control of vehicle dynamics from the perspective of the two classical formulations of the optimal control problem, namely the Euler-Lagrange equations and the Hamilton-Jacobi-Bellman equation.

Central to our methodology is the use of tools from geometric mechanics to formulate vehicle dynamics, which are often mechanical in nature. Many exciting advances have been made in recent decades at the juncture of nonlinear control theory, differential geometry, and geometric mechanics: see [7, 52, 57, 80, 82] for a few significant examples. In the modeling of mechanical control systems, one tool that has emerged is the affine connection and the associated notion of covariant differentiation. See [35, 58] for extensive development of these concepts and their relation to Riemannian geometry. The affine connection appears within Lagrangian systems as the Levi-Civita connection, and can also be used to model mechanical control systems with nonholonomic constraints. This use of the affine connection, which can be thought of as a covariant formulation of $F = ma$, is proposed in Bloch and Crouch [4] as a means of modeling mechanical systems

with nonholonomic constraints, an idea which is furthered by Lewis [65]. The role of the connection in symmetry and reduction of nonholonomic mechanical systems is exploited to great effect by Bloch et al. [7]. From the perspective of control theory, the affine connection has been used to derive controllability tests for mechanical control systems [67], motion control algorithms [17], and series expansions of trajectories of mechanical systems [16]. The work presented in this thesis continues in that spirit by applying this formulation to the optimal control problem.

In Chapter 2, we consider the optimal control problem as a constrained optimization problem and use the technique of Lagrange multipliers to derive the equations of motion. Mechanical systems, which include many examples of vehicle dynamics, are naturally second-order. Rather than rewrite the equations of motion as a first-order system, thereby doubling the number of constraints, we include them in the more natural form, namely as a second-order constraint of the optimization. In optimizing, we consider how variations of position induce variations in the velocity, rather than taking independent variations of the two. Central to these computations is the use of the affine connection and covariant differentiation. The resulting equations of motion of the adjoint variables are second-order differential equations defined using the covariant derivative, thus inheriting the geometric structure of the vehicle dynamics. This approach has a potential computational advantage in that the number of adjoint variables has been halved. Additionally, recently developed integration schemes based on variational principles [101] or which exploit the presence of an affine connection [16] can potentially be used in the approximation of optimal trajectories. Finally, this formulation affords greater insight into the role geometric notions such as symmetry and reduction play in the construction of optimal trajectories.

In Chapter 3, we consider the Hamilton-Jacobi-Bellman (HJB) equation for optimal control of vehicles which are not subject to external forces such as gravity. Examples include satellite orientation, AUV motion, and robotic locomotion. These systems have the property that they are *time scalable*, meaning the equations of motion remain the same when time is reparameterized. For time-scalable systems, we show that if the cost function is compatible with the time scaling, one can deduce *a priori* the time dependence of the value function associated with the HJB partial differential equation (PDE). This results in conversion of that PDE, which depends on space and time, into a purely spatial PDE. The solution of this PDE, combined with the known time dependency, yields the control law

which generates the optimal finite-time trajectories. If the time variation of the control law is omitted, the same value function can produce a stabilizing (though not optimal) control law, with the value function playing the role of Lyapunov function. These results are applied to driftless systems as well, which are time scalable though not mechanical.

Optimal control is a venerable mathematical field, with antecedents far older than control theory itself. A survey of the field is beyond the scope of this introduction; we mention only [13, 93, 111] as recent texts covering aspects of modern optimal control theory. Optimal control and geometric mechanics share intimate links due to their origins in the calculus of variations [70]. This link is explored in [60] in the context of optimal control of mechanical systems with symmetry and by Bloch and Crouch [6], who examine the equivalence of optimal control problems with higher-order variational problems. The variational approach to optimal trajectory generation for systems defined by an affine connection is considered in [29, 83] in the context of finding force-minimizing arcs on Riemannian manifolds. In that case, however, the underlying system was fully actuated, meaning that the constraints could be directly substituted into the cost function, thus eliminating the need for adjoint variables. The results presented here are more general, in that they apply to underactuated systems, systems with drift, and a broader class of cost functions. The equations we derive match those derived recently in [66], which considers the same problem from a different perspective. In that paper, the first-order optimal control equations given by the Pontryagin Maximum Principle are recast as second-order equations using splittings derived from the affine connection. Our work complements those results, in that it shows how to interpret the resulting equations from the perspective of variational calculus.

1.2.2 Cooperative Control of Vehicle Formations

As we noted earlier, vehicles in formation depend on information from one another to accomplish their objective. When the formation is dynamically coupled, that coupling constrains, or at least naturally suggests, what information must be available to each component of the decentralized controller. In the case of cooperative vehicle control, no such architecture is necessarily suggested. In some cases, such as automated highway systems [99], the task may suggest which vehicles ought to have information about one another, but for many tasks this is not the case. As such, central to any discussion of cooperative control of vehicles is a determination

of the nature of the *information flow* throughout the formation. We will distinguish between two types of information flow: *sensed information*, meaning the ability of a single vehicle to sense some information (e.g., relative position) about another vehicle in a way which involves no action on the part of that vehicle, and *transmitted information*, meaning transfer of information between two vehicles which requires some action on the part of both the sender and recipient. Sensing and transmission, or “seeing” and “hearing,” as we refer to them colloquially, together are the means by which each vehicle acquires the information necessary to perform its task within the formation.

Several observations about the information flow within a formation make clear the need to consider its impact on the formation performance. The first is that as a rule, no vehicle will be able to see or hear the entire formation. Having each vehicle simultaneously solve a centralized control problem using complete information is therefore infeasible; some form of decentralized control is required. The second observation is the sensing and transmitted information topologies are themselves *dynamic*, meaning they are subject to disruption by external influences or changes in the formation itself. As such, a control law which is optimized for one topology may exhibit poor performance, or even instability, for another topology.

One possible approach to vehicle formation control is to implement a centralized controller or decision maker, and to overcome the limitations in the information flow topologies by having each vehicle transmit all information it possesses, thereby allowing all necessary information to eventually arrive at one vehicle, who then transmits the results of his centralized computation. While this architecture may be appropriate in certain cases, it possesses certain deficiencies which render it infeasible in a dynamic and adversarial environment. This architecture, which requires *maximal* information flow, is necessarily slow and expensive in terms of bandwidth. It is fragile, in that it depends on the reliable transmission of large amounts of data over potentially unreliable and dynamic communication channels. Instead, we intend to research the implementation of decentralized control laws augmented by *minimal* information flow. A minimal information flow paradigm has the potential to balance the performance improvements achieved through information sharing with the requirements of reliability and stealth.

With that in mind, we devote the second section of the thesis to exploring the role which the communication topologies play in cooperative control of vehicle formations. Our approach is to model the communication topologies as graph; each vehicle is a node of a (directed) graph, and an arc is drawn from node i to

node j if vehicle i receives information from vehicle j . By merging ideas from graph theory, control theory, and dynamical systems theory, we are able to study the interplay between the communication network and vehicle dynamics, and to propose strategies for information exchange which mitigate those effects. We limit our focus to linear dynamics and control systems in order to elucidate the role of the graph in the system behavior. In Chapter 7 we discuss extensions to nonlinear systems.

In our investigation of the interplay between the communication topology and formation performance, we will not focus on how the communication network is physically realized. Indeed, the connectivity of the graph may not coincide with the physical realization of the communication network. For example, the Internet can be thought of as a complete graph due to the ability of any two nodes to communicate, despite the obvious lack of physical connection between all the nodes. In our case, intervehicle communication may take place over various media, including wireless networks. Nonetheless, we do not wish to assume that any two vehicles communicating over a network are necessarily connected. One reason this assumption cannot be made is that vehicles need to make decisions and implement control laws in real time. Any communication network possesses time delays, and these delays may render two vehicles effectively unconnected for the purpose of real time control. A second reason is that the problem of computational complexity may require vehicles to ignore some information they receive so that the necessary computations can be carried out in the required time. A third reason is that the demands of network capacity and/or stealth may force the network to limit the amount of transmission taking place. For these reasons, we will model the communication network as an arbitrary graph and set aside the issue of how that graph is physically realized.

The use of graphs in analysis of interconnected systems is not new. Mason's gain formula, which can be found in standard introductory control texts (e.g., [42]), uses the notion of a signal flow graph to compute the transfer function of an interconnected system. When decentralized control became an area of study in the 1970s, researchers used graph-theoretic ideas in modeling interconnections [27, 91, 100]. Many researchers focused on arbitrary interconnected systems, for which decentralized controller synthesis procedures are difficult to derive. The current broad interest in vehicle formations has revived an interest in graph-theoretic ideas [46, 75, 94, 97].

A significant source of research on vehicle formations is the mobile robotics community (see [20] for a recent survey). Within this community, several approaches to formation control have emerged. Many of these researchers focus on “leader-follower” formations. Two recent papers exploring graph-theoretic ideas in the context of a leader-follower architecture are Desai, et al. [33] and Tabuada et al. [97]. This approach has the advantage of simplicity in that a reference trajectory is clearly defined by the leader, and no cycles exist in the graph which complicate the dynamics. However, they possess disadvantages which lead us to study the behavior of vehicle formations whose interconnection topology is not acyclic. First, leader-follower architectures are known to have poor disturbance rejection properties (see, e.g., [106]). Secondly, enforcing an acyclic architecture requires some global knowledge of the graph, which may not be available to a given vehicle deciding whom to follow. Finally, a leader-follower architecture depends heavily on the leader for achieving its goal, and over-reliance on a single vehicle in the formation may be undesirable, especially in adversarial environments.

Another approach is the “virtual leader” approach [38, 63, 92], in which vehicles in the formation jointly synthesize a single fictitious leader vehicle whose trajectory acts as a leader for the group. This approach avoids the problems with disturbance rejection inherent in the leader-follower approach, but at the expense of high communication and computation requirements needed to synthesize the virtual leader and communicate its position. From a graph-theoretic perspective, this approach requires a complete graph, in that the virtual leader is computed using information from all vehicles. Some researchers [63] use techniques such as artificial potentials and nearest neighbor techniques to synthesize the control laws to circumvent the requirement of global communication and computation. This approach has yielded some success and continues to bear fruit; however, it has not yet reached the point where it can be applied to arbitrary communication topologies and to arbitrary vehicle dynamics.

What these approaches have in common is an assumption about the underlying topology of the graph which enables the use of a particular formation control methodology. As discussed above, this assumption cannot be maintained in many application areas, especially as the number of vehicles in the formation increases. For these reasons, we wish to consider a broader range of vehicle interconnection possibilities. The first challenge in this case is formation stabilization. When cycles are present in the graph, they induce global formation dynamics which preclude local stability analysis. A second issue that must be confronted is the ability of

the formation to agree upon an effective leader in the situation where no leader is defined by the graph. We explore these two topics in the second half of this thesis, with a goal of deriving information exchange strategies which improve formation stability and performance and are robust to changes in the graph.

In Chapter 5, we examine formation stability in the case where no transmitted information exists between vehicles. We demonstrate how the Laplacian matrix of the graph appears in the formation equations of motion. In the case where the vehicles have identical dynamics and controllers, we derive stability criteria involving the eigenvalues of the Laplacian. This approach yields an elegant generalization of the Nyquist criterion in which the negative inverse of the Laplacian eigenvalues replace the -1 point about which one computes encirclements. The criterion allows one to design a decentralized vehicle controller at the local level and determine if it will stabilize the formation. The Laplacian is an object of study within algebraic graph theory [24, 39, 73]; one area of research is the correlation of Laplacian eigenvalue locations to structural properties of the underlying graph. Using those insights, we evaluate desirable graph properties in terms of formation stability.

One area of research which focused on similar issues is the study of synchronization of chaotic oscillators. The work of [48, 84] identified the Laplacian eigenvalues as an important object of study in synchronization. Other researchers took a more control-theoretic approach, using tools such as Lyapunov stability [105] and the circle criterion [104] to derive sufficient conditions for nonlinear chaotic oscillators to synchronize. In particular, the observation in [105] that the eigenvalues must be “negative enough” to achieve stability is similar in spirit to the work presented here. Our work differs in that we restrict our focus to stabilization of formations with linear dynamics, and as such we are able to make more precise statements about the role of Laplacian eigenvalues in determining formation stability. Our results are useful for controller design and also lay the foundation for our investigation of the role of intervehicle communication.

In Chapter 6, we turn to the flow of transmitted information between vehicles. Our approach is to model the information flow as a dynamical system, and to construct a dynamical system with desirable stability and convergence properties. This development will rely heavily on ideas from Perron-Frobenius theory, which are closely tied to graph-theoretic concepts. We propose an information flow law which, upon convergence, supplies each vehicle with an agreed-upon formation center. The stability of the information flow law is analyzed using the tools from

the previous chapter. The common reference supplied by the information flow law supplies the formation with an effective leader, thereby facilitating coordinated formation motion. A feedforward information correction law is also proposed, which prevents the information flow law from lagging vehicle motion. When this term is included, a stability separation principle is derived, wherein stability of the formation is achieved if the information flow law stabilizes the graph independent of the vehicle dynamics, and the local controller stabilizes the plant independent of the graph. This striking result renders formation stability and performance largely independent of the underlying graph. The dynamical systems approach to achieving consensus to formation center renders the information flow law highly robust to changes in the communication topology. Finally, this approach is shown to have good string stability properties.

1.3 Statement of Contributions

A brief restatement of the thesis contributions by chapter is found in the following paragraphs.

- *Chapter 2:* In this chapter we review the affine connection and its role in modeling mechanical control systems. We derive an alternate formulation of the Euler-Lagrange equations using a variational approach appropriately tailored to our setting. The resulting equations for the adjoint variable are also governed by the affine connection, revealing the parallel structure of the optimal control equations.
- *Chapter 3:* In this chapter we develop the notion of time scalability as a property of mechanical control systems which do not include external forces such as potential gradients or damping. Time scalability is used to eliminate time from the HJB equation, leaving a purely spatial PDE. The resulting PDE, when solved, supplies both finite-horizon optimal trajectories and a stabilizing control law.
- *Chapter 4:* In this chapter we review relevant ideas from graph theory and Perron-Frobenius theory in preparation for our discussion of cooperative control of vehicle formations.
- *Chapter 5:* In this chapter we demonstrate the role of the Laplacian of the graph in determining formation stability. A Nyquist criterion is derived

for formation stability using Laplacian eigenvalues which facilitates design of local vehicle control laws which stabilized the overall formation. The spectral properties of the Laplacian are used to evaluate desirable properties of formation graphs.

- *Chapter 6:* In this chapter we propose an information flow law which supplies each vehicle with an agreed-upon formation center. This information flow law can be made robust to changes in the graph. When the output of the information flow law is used as an input to the vehicle controller, together with a feedforward compensation term, the resulting system exhibits a stability separation principle which renders the overall system highly robust to changes in the graph. This approach is also shown to have good string stability properties.
- *Chapter 7:* In this chapter we review the results presented in this thesis and discuss avenues for future research.

Chapter 2

Optimal Control of Affine Connection Control Systems

In this chapter we derive optimal control equation for vehicle motion using the machinery of affine connections. This approach exposes the parallel geometric structure between the vehicle dynamics and the dynamics of the adjoint variables. The approach in this chapter is motivated by recent work using affine connection in the context of finding force-minimizing arcs on Riemannian manifolds [29, 83]. The results here extend that approach to the general optimal control problem, including underactuated mechanical systems, systems with damping, and a broad class of cost functions. The results presented in this chapter confirm the work of Lewis [66], which considers the same problem from the perspective of the Pontryagin Maximum Principle. In that paper, the standard first-order equations are recast as second-order equations using splittings derived from the affine connection. Our approach is complementary, in that it shows how to interpret the resulting equations from the perspective of variational calculus.

This chapter is organized as follows. In Section 2.1, we introduce the affine connection and recall necessary ideas from the calculus of variations. In Section 2.2, we define an affine connection control system, pose the optimal control problem, and derive the optimal control equations. Section 2.3 contains several examples. For consistency, we employ the notation of [66].

2.1 Mathematical Preliminaries

2.1.1 Theory of Affine Connections

In this section, we introduce terminology which we will use throughout the chapter. For a more thorough introduction to the topic, see [35, 58].

Connections and associated Tensors

Let Q be a manifold, and TQ (T^*Q) be the (co-)tangent bundle of Q . Letting $\{q^1, \dots, q^i, \dots\}$ be local coordinates on Q , the associated coordinate vector field is denoted $\frac{\partial}{\partial q^i}$ and the associated coordinate one-form field is denoted dq^i . The natural pairing between vectors and one forms is denoted $\langle \cdot, \cdot \rangle$. Recall that $\langle \frac{\partial}{\partial q^i}, dq^j \rangle = \delta_j^i$. Let $\mathcal{X}(Q)$ be the set of all smooth vector fields on Q , and let $[\cdot, \cdot]$ denote the Lie bracket of vector fields on Q . Using these definitions, we can define the affine connection:

Definition 2.1. An affine connection on Q is a map $\nabla : \mathcal{X}(Q) \times \mathcal{X}(Q) \rightarrow \mathcal{X}(Q)$, denoted $\nabla : X, Y \mapsto \nabla_X Y$, which satisfies the following properties:

1. $\nabla_{fX+gY} Z = f\nabla_X Z + g\nabla_Y Z$
2. $\nabla_X(Y + Z) = \nabla_X Y + \nabla_X Z$
3. $\nabla_X(fY) = f\nabla_X Y + X(f)Y$

where f, g are smooth functions on Q .

The affine connection allows us to introduce the notion of covariant differentiation, which is the differentiation of vectors (or arbitrary tensors, as we shall see) along a path $c(t)$ in Q :

Proposition 2.1. Let $c(t)$ be a differentiable curve in Q . There exists a unique correspondence which associates to a vector field V defined along $c(t)$ another vector field $\frac{DV}{dt}$, called the covariant derivative of V along c , such that

1. $\frac{D}{dt}(V + W) = \frac{DV}{dt} + \frac{DW}{dt}$
2. $\frac{D}{dt}(fV) = \frac{df}{dt}V + f\frac{DV}{dt}$
3. If V is induced by a vector field $Y \in \mathcal{X}(Q)$, i.e., $V(t) = Y(c(t))$, then $\frac{DV}{dt} = \nabla_{\frac{dc}{dt}} Y$.

Note that for the covariant derivative of X to be defined at a point $c(t)$, we need know only $X(c(t))$, $\frac{c(t)}{dt}$, and the rate of change of X along $c(t)$. The affine connection and the covariant derivative are often used interchangeably in the literature; we will endeavor to use the affine connection when the operands are elements of $\mathcal{X}(Q)$ and the covariant derivative when the operands are vector fields defined only along curves. Note that the covariant derivative satisfies linearity and a product rule, thus making it a derivative in the algebraic sense. A salient feature of the covariant derivative is that it returns an element of TQ rather than an element of TTQ .

Let X, Y be vector fields whose representation in coordinates are $X = X^i \frac{\partial}{\partial q^i}$, $Y = Y^i \frac{\partial}{\partial q^i}$, where X^i, Y^i are smooth functions on Q . The affine connection, written in coordinates, takes the form

$$\nabla_X Y = \left[X^i Y^j \Gamma_{ij}^k + X(Y^k) \right] \frac{\partial}{\partial q^k} \quad (2.1)$$

where Γ_{ij}^k are functions on Q known as the *Christoffel symbols* of the connection. Note that we employ the Einstein summation convention when performing computations in coordinates.

While the connection ∇ is not a tensor, it has two tensors associated with it which will be used later in the paper. The first is the curvature form, which is a (1,3) tensor defined to be

$$R(X, Y)Z = \nabla_X \nabla_Y Z - \nabla_Y \nabla_X Z - \nabla_{[X, Y]} Z \quad (2.2)$$

and the second is the torsion form, a (1,2) tensor defined to be

$$T(X, Y) = \nabla_X Y - \nabla_Y X - [X, Y]. \quad (2.3)$$

We can write these tensors in coordinates using the Christoffel symbols Γ_{ij}^k associated with ∇ :

$$T_{ij}^s = \Gamma_{ij}^s - \Gamma_{ji}^s \quad (2.4)$$

$$R_{ijk}^s = \frac{\partial \Gamma_{jk}^s}{\partial q^i} - \frac{\partial \Gamma_{ik}^s}{\partial q^j} + \Gamma_{jk}^l \Gamma_{il}^s - \Gamma_{ik}^l \Gamma_{jl}^s, \quad (2.5)$$

where the index notation is given by

$$T(X, Y) = T_{ij}^s X^i Y^j \frac{\partial}{\partial q^s} \quad (2.6)$$

$$R(X, Y)Z = R_{ijk}^s X^i Y^j Z^k \frac{\partial}{\partial q^s}. \quad (2.7)$$

Note that although the definitions in Equations (2.2),(2.3) use the affine connection, the tensors themselves are functions of the vectors at a given point, since the elements of right-hand side of Equations (2.2),(2.3) which involve the local variation of the operands are internally canceled. When the torsion tensor is zero, the affine connection is said to be *symmetric*. In this case, the affine connection satisfies the property

$$[X, Y] = \nabla_X Y - \nabla_Y X. \quad (2.8)$$

Differentiation of Tensors

In this section, we present the formulas for covariant differentiation of an arbitrary covariant tensor. Let A be a tensor of order $(0, r)$. Then its covariant differential ∇A is a tensor of order $(0, r + 1)$, defined by ([35], p. 102)

$$\begin{aligned} (\nabla A)(X_1, \dots, X_r, Z) = & \quad (2.9) \\ Z[A(X_1, \dots, X_r)] - A(\nabla_Z X_1, \dots, X_r) - \dots - A(X_1, \dots, \nabla_Z X_r) \end{aligned}$$

and the covariant derivative of A in the direction of Z is an $(0, r)$ tensor $\nabla_Z A$ defined by

$$(\nabla_Z A)(X_1, \dots, X_r) = (\nabla A)(X_1, \dots, X_r, Z). \quad (2.10)$$

In particular, if we wish to differentiate a one-form α , we see that

$$\langle X, \nabla_Z \alpha \rangle = Z(\langle X, \alpha \rangle) - \langle \nabla_Z X, \alpha \rangle. \quad (2.11)$$

If Z is the tangent vector to a path $c(t)$, we can rewrite this identity as

$$\frac{d}{dt} \langle X, \alpha \rangle = \left\langle \frac{DX}{dt}, \alpha \right\rangle + \left\langle X, \frac{D\alpha}{dt} \right\rangle. \quad (2.12)$$

We see that our definition for covariant differentiation of tensors leads to a product rule as one would expect. This identity is true for any connection, unlike the product rule to be introduced in the next section. This equation will be significant

for two reasons. Firstly, we will use it to integrate by parts expressions involving covariant derivatives. Secondly, if we replace t with ϵ , we see that this expression will be useful in understanding how variational principles enter equations with covariant derivatives.

Riemannian Geometry

The development until now has not relied on the existence of a metric on Q . Let us now consider the case where such a metric exists. Suppose Q is a Riemannian manifold, meaning it is endowed with a symmetric, positive definite two-form, denoted $\langle\langle X, Y \rangle\rangle$. As with all two-forms, the Riemannian metric can be thought of as a map $g : TQ \rightarrow T^*Q$ which satisfies the property $\langle g(X), Y \rangle = \langle\langle X, Y \rangle\rangle$ for all $X, Y \in TQ$. Clearly, g is invertible.

A landmark result in Riemannian geometry is the following:

Theorem 2.1 (Levi-Civita). *A Riemannian manifold Q possesses a unique affine connection ∇ which satisfies the following properties:*

1. ∇ is symmetric.
2. Given vector fields X, Y defined along a differentiable curve $c(t)$ in Q ,

$$\frac{d}{dt} \langle\langle X, Y \rangle\rangle = \left\langle\left\langle \frac{DX}{dt}, Y \right\rangle\right\rangle + \left\langle\left\langle X, \frac{DY}{dt} \right\rangle\right\rangle. \quad (2.13)$$

This connection, known as the Levi-Civita connection, is the connection which satisfies a product rule with regard to the Riemannian metric. It can be defined in coordinates in the following way:

$$\Gamma_{ij}^k = \frac{1}{2} g^{mk} \left(\frac{\partial g_{mj}}{\partial q^i} + \frac{\partial g_{mi}}{\partial q^j} - \frac{\partial g_{ij}}{\partial q^m} \right), \quad (2.14)$$

where g_{jk}, g^{jk} are g and g^{-1} in coordinates. Note that defining the Christoffel symbols uniquely defines the affine connection.

When employing the Levi-Civita connection in computations, the following result is often useful:

Proposition 2.2. *Let ∇ be the Levi-Civita connection associated with Riemannian metric g , as defined above. Let X and α be a smooth vector field and one form defined along a differentiable curve $c(t)$ in Q . Then g and g^{-1} commute with ∇ ,*

that is:

$$g\left(\frac{DX}{dt}\right) = \frac{D}{dt}(gX) \quad (2.15)$$

$$g^{-1}\left(\frac{D\alpha}{dt}\right) = \frac{D}{dt}(g^{-1}\alpha) \quad (2.16)$$

Proof. Consider how the two terms act on another vector field, Z . We have

$$g\left(\frac{DX}{dt}\right)(Z) = \left\langle\left\langle\frac{DX}{dt}, Z\right\rangle\right\rangle \quad (2.17)$$

and, from Equation (2.11), we have

$$\frac{D}{dt}(gX)(Z) = \frac{d}{dt}\langle gX, Z \rangle - gX\left(\frac{DZ}{dt}\right) \quad (2.18)$$

or, using the metric notation,

$$\frac{D}{dt}(gX)(Z) = \frac{d}{dt}\langle\langle X, Z \rangle\rangle - \left\langle\left\langle X, \frac{DZ}{dt} \right\rangle\right\rangle. \quad (2.19)$$

Subtracting Equation (2.17) from this equation yields

$$\frac{D}{dt}(gX)(Z) - g\left(\frac{DX}{dt}\right)(Z) = \frac{d}{dt}\langle\langle X, Z \rangle\rangle - \left\langle\left\langle X, \frac{DZ}{dt} \right\rangle\right\rangle - \left\langle\left\langle \frac{DX}{dt}, Z \right\rangle\right\rangle. \quad (2.20)$$

The right-hand side of this equation is zero by Theorem 2.1 if ∇ is the Levi-Civita connection, which proves our assertion. The second assertion of the proposition is shown to be true if X is replaced with $g^{-1}\alpha$ in the first assertion. ■

An equivalent statement, which can be seen by examining Equation 2.9 and seeing that the right-hand side is zero, is that $\nabla g = 0$.

Additional Definitions

We now introduce some additional terminology, borrowed from [66], which will be useful. Given an $(1, r)$ tensor A , we associate with it an $(2, r - 1)$ tensor A^* which satisfies the identity

$$\langle A(X_1, \dots, X_r), \alpha \rangle = \langle X_1, A^*(\alpha, X_2, \dots, X_r) \rangle. \quad (2.21)$$

We will also associate with any $(0, r)$ tensor A an $(1, r - 1)$ tensor \hat{A} such that

$$A(X_1, \dots, X_{r-1}, Y) = \langle Y, \hat{A}(X_1, \dots, X_{r-1}) \rangle. \quad (2.22)$$

Finally, we will associate with any $(r, 2)$ tensor A another $(r, 2)$ tensor \tilde{A} which is the same with the indices for the vector inputs reversed, that is:

$$A(X, Y) = \tilde{A}(Y, X). \quad (2.23)$$

2.1.2 Calculus of Variations

In this section, we recall how variations are defined, and we derive properties of the variations which will be used later on. We consider a family of trajectories $q(t, \epsilon)$ defined on some interval $[t_0, t_f] \times [-\epsilon_0, \epsilon_0]$, and we denote $q(t, 0)$ as $q(t)$. We can also freeze t and consider a path dependent on ϵ . We now define the variations of $q(t)$ in the standard way:

$$\delta q(t) = \left. \frac{\partial q}{\partial \epsilon} \right|_{\epsilon=0} \in T_{q(t)}Q. \quad (2.24)$$

We also note that the velocity vector field is given by

$$V(t, \epsilon) = \frac{\partial q}{\partial t} \in T_{q(t, \epsilon)}Q. \quad (2.25)$$

When $\epsilon = 0$, we denote the velocity as $V(t)$. We see that $\delta q(t)$ can be thought of as a vector field defined along a path $q(t)$, and if we hold t fixed, $V(t, \epsilon)$ can be thought of as a vector field defined along a path $q(t, \epsilon)$ parametrized by ϵ . With that in mind, we see that $\frac{DV}{d\epsilon}$ is well-defined at $\epsilon = 0$ and $\frac{D\delta q}{dt}$ is also well defined for all t . We now prove three propositions using the coordinate definition of covariant differentiation.

Proposition 2.3. $\frac{\partial V^k}{\partial \epsilon} = \frac{\partial \delta q^k}{\partial t}$.

Proof. This follows immediately from the definitions and the equality of mixed partials. ■

Proposition 2.4. $T(\delta q(t), V(t)) = \frac{DV}{d\epsilon} - \frac{D\delta q}{dt}$.

Proof. The definition of covariant differentiation in coordinates is

$$\frac{DV}{d\epsilon} = \left(\frac{\partial V^k}{\partial \epsilon} + V^j \delta q^i \Gamma_{ij}^k \right) \frac{\partial}{\partial q^k} \quad (2.26)$$

and

$$\frac{D\delta q}{dt} = \left(\frac{\partial \delta q^k}{\partial t} + \delta q^j V^i \Gamma_{ij}^k \right) \frac{\partial}{\partial q^k}. \quad (2.27)$$

Subtracting and applying Proposition 2.3 yields

$$\frac{DV}{d\epsilon} - \frac{D\delta q}{dt} = V^j \delta q^i \left(\Gamma_{ij}^k - \Gamma_{ji}^k \right) \frac{\partial}{\partial q^k} \quad (2.28)$$

which is the coordinate definition of $T(\delta q, V)$ given in Equation (2.4). ■

Proposition 2.5. $R(\delta q(t), V(t))Z = \left(\frac{D}{d\epsilon} \frac{D}{dt} - \frac{D}{dt} \frac{D}{d\epsilon} \right) Z$.

Proof. The proof follows in exactly the same fashion as Proposition 2.4. If one expands the right-hand side in coordinates and cancels like terms via Proposition 2.3, one is left with the coordinate definition of $R(\delta q(t), V(t))Z$ as given in Equation (2.5). ■

In the case where $\delta q, V$ can be written as vector fields on some open set containing $q(t, 0)$, then it is possible to prove the above by showing that $[\delta q(t), V(t)] = 0$ and applying this to the intrinsic definition of the curvature and torsion forms. The difficulty is that $\delta q(t)$ and $V(t)$ are not defined as vector fields, may not be locally extendable to form vector fields, in which case $[\delta q(t), V(t)]$ is not defined.

We now discuss how covariant derivatives enter the picture. If we fix t , we can consider $V(t, \epsilon)$ as a vector field varying along a trajectory $q(t, \epsilon)$. Suppose we had a one-form $\alpha(t, \epsilon)$ similarly defined, and we consider the variation of

$$\int \langle V, \alpha \rangle dt. \quad (2.29)$$

Using Equation (2.12), we see that

$$\begin{aligned} \delta \int \langle V, \alpha \rangle dt &= \int \frac{d}{d\epsilon} \langle V, \alpha \rangle dt \\ &= \int \left\langle \frac{DV}{d\epsilon}, \alpha \right\rangle + \left\langle V, \frac{D\alpha}{d\epsilon} \right\rangle dt. \end{aligned} \quad (2.30)$$

Indeed, the expression $\frac{DV}{d\epsilon}$ is precisely the variation of V as ϵ is varied. It is

important to note that variations of V arise solely from the family of trajectories $q(t, \epsilon)$.

2.2 Optimal Control of Affine Connection Control Systems

2.2.1 Affine Connection Control Systems

Having developed the machinery of affine connections, we now discuss its application to modeling of vehicle dynamics. Essentially, the affine connection facilitates a coordinate-independent formulation of Newtonian dynamics. When the equations of motion are written in terms of $\frac{DV}{dt}$ rather than \ddot{q} , the resulting equations are coordinate-independent, with the connection encompassing the Coriolis terms present in moving frames. As such the affine connection is a useful tool for analysis of vehicle dynamics.

Many mechanical systems admit a Lagrangian formulation, where the kinetic energy term is in fact a Riemannian metric. As we have seen, the affine connection is closely related to the Riemannian metric, which should make clear its utility in representing Lagrangian systems. In fact, Lagrangian systems with nonholonomic constraints can also be represented via the affine connection, though not the Levi-Civita connection (see [4, 65]). From the perspective of control theory, the affine connection has been used to derive controllability tests for mechanical control systems [67] and to design motion control algorithms [17]. For an extensive discussion of modeling of vehicle dynamics via affine connections, see the work of Lewis [64] and Bullo [14].

Definition 2.2. An *affine connection control system* on a manifold Q endowed with an affine connection ∇ is a (1,1) tensor $F : TQ \rightarrow TQ$, a vector field P , a set of vector fields $Y_i \in TQ$, $i = 1 \dots m$, and a set of controls $u^i(t) : \mathbb{R}^+ \rightarrow \mathbb{R}$. Together, these define the following differential equation:

$$\frac{DV}{dt} = Y_i(q)u^i + F(V) + P(q). \quad (2.31)$$

The vector fields Y_i map the controls to TQ . The vector field P represents the drift vector field (in a second-order sense). In a mechanical setting, we would write this as the gradient of a potential function, but we need not make this restriction here. The term $F(V)$ represents dissipation. The latter two terms are

not generally included in discussions of affine connection control systems, but here they have been included, since they do not complicate the theory in this case. As before, the velocity \dot{q} is denoted by V . Note that despite the fact that all vector fields are defined on Q , they govern the evolution of the system on TQ , and the affine connection facilitates that formulation.

2.2.2 Cost Functions

We now need to define a cost function to complete the optimal control problem. To do so, we borrow the framework of [66]. Let $i \in \{1, \dots, s\}$, r_i be a nonnegative integer, and A_i be a symmetric \mathbb{R}^m -dependent $(0, r_i)$ tensor field on Q . Our integral cost function J is therefore

$$J(q, u, V) = \sum_{i=1}^s A_i(q, u)(V, \dots, V), \quad (2.32)$$

where V is repeated r_i times as an operand for each A_i . Thus, when $r_i = 0$, the tensor represents a cost term associated with position and control effort alone. If $r_i = 2$, then the cost is quadratic in velocity and possibly dependent on q and u .

Also, let $\phi_i(q)$ be a symmetric $(0, b_i)$ tensor field on Q . We define a terminal cost function J_f as

$$J_f(q, V) = \sum_{i=1}^w \phi_i(q)(V, \dots, V). \quad (2.33)$$

2.2.3 Problem Statement

We can now state two optimal control problems:

P1 (Fixed Final State) Given an initial condition $(q_0, V_0) \in TQ$ and a final condition $(q_f, V_f) \in TQ$, determine trajectories that minimize

$$\int_{t_0}^{t_f} J(q, u, V) dt \quad (2.34)$$

subject to the constraint of Equation (2.31).

P2 (Free Final State) Given an initial condition $(q_0, V_0) \in TQ$, determine trajectories that minimize

$$\int_{t_0}^{t_f} J(q, u, V) dt + J_f(q(t_f), V(t_f)) \quad (2.35)$$

subject to the constraint of Equation (2.31).

2.2.4 Problem Solution: Lagrange Multipliers

To solve this problem we use the technique of Lagrange multipliers and calculus of variations. Recall that using this approach, the problem of optimizing of $J(x, u)$ when constrained by $f(x, u) = 0$ is done by solving the unconstrained minimization of $J + \lambda f$, where λ is an additional variable, known as the adjoint to x , whose dimension is equal to that of the number of equality constraints. When this optimization is carried out, the result equations, known as the Euler-Lagrange equations, supply necessary conditions on the optimal trajectory in x and λ . When inequality constraints on the states or controls are present, these conditions may not be satisfied at all points along the trajectory. For a comprehensive study of this approach to optimal control, see Bryson and Ho [13] or Stengel [93]. In this study, we will limit our scope to the case where the optimization is only constrained by the dynamics, in which case the resulting equations can be obtained via a variational approach.

The standard Euler-Lagrange equations are derived for a system whose constraints are of the form $\dot{x} = f(x, u)$. It is possible to represent an affine connection control system in that form by introducing velocities as independent states, and considering the variations of positions and velocities independently. This approach, while valid, has several disadvantages. One is that it necessitates doubling the number of constraints. Additionally, by rendering the system first-order, the underlying geometric structure is lost. Our goal is to derive the ‘‘Euler-Lagrange’’ equations which govern the optimal control problem in a way which preserves that structure. Thus, rather than artificially doubling the constraints to conform to the standard structure, we consider only the given constraint. The unconstrained cost function takes the following form:

$$\int_{t_0}^{t_f} J(q, u, V) + \left\langle \lambda, -\frac{DV}{dt} + F(V) + P + Y_i u^i \right\rangle dt + J_f(q(t_f), V(t_f)). \quad (2.36)$$

We see that our constraint lives in TQ , meaning the Lagrange multiplier is a one-

form field on Q . When we take the variation, we consider the variation of u and q . Unlike the first-order case, we do not consider the variation of V independently of that of q . Instead, we look at how variations of q affect the velocity vector field, and we do so using the covariant derivative, as discussed above.

Our governing equation is now

$$\delta \left(\int_{t_0}^{t_f} J(q, u, V) + \left\langle \lambda, -\frac{DV}{dt} + F(V) + P + Y_i u^i \right\rangle dt + J_f(q(t_f), V(t_f)) \right) = 0 \quad (2.37)$$

or

$$\int_{t_0}^{t_f} \frac{d}{d\epsilon} \left(J(q, u, V) + \left\langle \lambda, -\frac{DV}{dt} + F(V) + P + Y_i u^i \right\rangle \right) dt + \frac{d}{d\epsilon} J_f(q(t_f), V(t_f)) = 0. \quad (2.38)$$

We now consider the variation of each term in this expression with respect to q . In the following section, the term V^{r_i} indicates V repeated r_i times as an operand.

1. $\frac{d}{d\epsilon} J(q, u, V)$

Suppose J consists of a single tensor A_i . We will calculate the variation of this term, and sum over the resulting expression for the case where J consists of multiple tensors. We assume $r_i \neq 0$, apply Equation (2.9) and exploit the symmetry of A_i to arrive at

$$\frac{dJ}{d\epsilon} = \nabla A_i(V^{r_i}, \delta q) + r_i A_i \left(\frac{DV}{d\epsilon}, V^{r_i-1} \right) \quad (2.39)$$

which, using Proposition 2.4, becomes

$$\frac{dJ}{d\epsilon} = \nabla A_i(V^{r_i}, \delta q) + r_i A_i \left(\frac{D\delta q}{dt} + T(\delta q, V), V^{r_i-1} \right). \quad (2.40)$$

Using the notation in Equation (2.22), we rewrite this as

$$\frac{dJ}{d\epsilon} = \left\langle \widehat{\nabla} A_i(V^{r_i}), \delta q \right\rangle + r_i \left\langle \frac{D\delta q}{dt} + T(\delta q, V), \widehat{A}_i(V^{r_i-1}) \right\rangle. \quad (2.41)$$

We split the latter term into two and, using the terminology of Equation (2.21), we write this as

$$\frac{dJ}{d\epsilon} = \left\langle \widehat{\nabla} A_i(V^{r_i}), \delta q \right\rangle + r_i \left\langle \frac{D\delta q}{dt}, \widehat{A}_i(V^{r_i-1}) \right\rangle + r_i \left\langle T^*(\widehat{A}_i(V^{r_i-1}), V), \delta q \right\rangle. \quad (2.42)$$

The second term of this equation is integrated by parts via Equation (2.12):

$$\begin{aligned} \frac{dJ}{d\epsilon} = & \left\langle \widehat{\nabla} \widehat{A}_i(V^{r_i}), \delta q \right\rangle - r_i \left\langle \delta q, \nabla_V \widehat{A}_i(V^{r_i-1}) \right\rangle + r_i \left\langle T^*(\widehat{A}_i(V^{r_i-1}), V), \delta q \right\rangle \\ & + \frac{d}{dt} r_i \left\langle \delta q, \widehat{A}_i(V^{r_i-1}) \right\rangle. \end{aligned} \quad (2.43)$$

The last term can be removed from the integrand, and will be considered later. Finally, we differentiate the second term on the right-hand side and collect terms using the symmetry of A_i , as we did before, to arrive at

$$\begin{aligned} \frac{dJ}{d\epsilon} = & \left\langle \widehat{\nabla} \widehat{A}_i(V^{r_i}), \delta q \right\rangle - r_i \left\langle \delta q, \left(\nabla_V \widehat{A}_i \right) (V^{r_i-1}) \right\rangle \\ & + r_i \left\langle T^*(\widehat{A}_i(V^{r_i-1}), V), \delta q \right\rangle - r_i(r_i - 1) \left\langle \delta q, \widehat{A}_i\left(\frac{DV}{dt}, V^{r_i-2}\right) \right\rangle. \end{aligned} \quad (2.44)$$

If we collect terms, we arrive at

$$\begin{aligned} \frac{dJ}{d\epsilon} = & \left\langle \left(\widehat{\nabla} \widehat{A}_i - r_i \nabla_V \widehat{A}_i \right) (V^{r_i}) \right. \\ & \left. - r_i(r_i - 1) \widehat{A}_i\left(\frac{DV}{dt}, V^{r_i-2}\right) + r_i T^*(\widehat{A}_i(V^{r_i-1}), V), \delta q \right\rangle. \end{aligned} \quad (2.45)$$

Of course, we can substitute in for $\frac{DV}{dt}$ using Equation (2.31). Note that if $r_i = 0$, this reduces to ∇A_i , which is simply the gradient of the function A_i , so the notation is consistent even for this case.

2. $\frac{d}{d\epsilon} \left\langle \lambda, \frac{DV}{dt} \right\rangle$

As before, we write

$$\frac{d}{d\epsilon} \left\langle \lambda, \frac{DV}{dt} \right\rangle = \left\langle \lambda, \frac{D}{d\epsilon} \frac{DV}{dt} \right\rangle. \quad (2.46)$$

Applying Proposition 2.5, we rewrite this as

$$\frac{d}{d\epsilon} \left\langle \lambda, \frac{DV}{dt} \right\rangle = \left\langle \lambda, \frac{D}{dt} \frac{DV}{d\epsilon} + R(\delta q, V)V \right\rangle. \quad (2.47)$$

We integrate the first term by parts, and rewrite the resulting expression using Equation (2.21) and Proposition 2.4:

$$\frac{d}{d\epsilon} \left\langle \lambda, \frac{DV}{dt} \right\rangle = - \left\langle \frac{D\lambda}{dt}, \frac{D\delta q}{dt} + T(\delta q, V) \right\rangle + \left\langle \delta q, R^*(\lambda, V)V \right\rangle + \frac{d}{dt} \left\langle \lambda, \frac{DV}{d\epsilon} \right\rangle. \quad (2.48)$$

We integrate by parts again and use Equation (2.21) to write this term as

$$\begin{aligned} \frac{d}{d\epsilon} \left\langle \lambda, \frac{DV}{dt} \right\rangle &= \left\langle \frac{D^2\lambda}{dt^2} - T^*\left(\frac{D\lambda}{dt}, V\right) + R^*(\lambda, V)V, \delta q \right\rangle \\ &\quad + \frac{d}{dt} \left(\left\langle \lambda, \frac{DV}{d\epsilon} \right\rangle - \left\langle \frac{D\lambda}{dt}, \delta q \right\rangle \right). \end{aligned} \quad (2.49)$$

Again, the final term can be integrated directly, and will be addressed later. The remaining term in the integrand is

$$\frac{d}{d\epsilon} \left\langle \lambda, \frac{DV}{dt} \right\rangle = \left\langle \frac{D^2\lambda}{dt^2} - T^*\left(\frac{D\lambda}{dt}, V\right) + R^*(\lambda, V)V, \delta q \right\rangle. \quad (2.50)$$

3. $\frac{d}{d\epsilon} \langle \lambda, F(V) \rangle$

Following the steps used earlier, we have

$$\begin{aligned} \frac{d}{d\epsilon} \langle \lambda, F(V) \rangle &= \left\langle \lambda, \nabla F(V, \delta q) + F\left(\frac{DV}{d\epsilon}\right) \right\rangle \\ &= \left\langle \lambda, \widetilde{\nabla F}(\delta q, V) + F\left(\frac{D\delta q}{dt} + T(\delta q, V)\right) \right\rangle \\ &= \left\langle \widetilde{\nabla F}^*(\lambda, V), \delta q \right\rangle + \left\langle F^*(\lambda), \frac{D\delta q}{dt} + T(\delta q, V) \right\rangle \\ &= \frac{d}{dt} \langle F^*(\lambda), \delta q \rangle + \left\langle \widetilde{\nabla F}^*(\lambda, V) + T^*(F^*(\lambda), V) \right. \\ &\quad \left. - \nabla F^*(\lambda, V) - F^*\left(\frac{D\lambda}{dt}\right), \delta q \right\rangle. \end{aligned} \quad (2.51)$$

As before, the final term is integrated directly, and the remaining term in the integrand is

$$\frac{d}{d\epsilon} \langle \lambda, F(V) \rangle = \left\langle \widetilde{\nabla F}^*(\lambda, V) + T^*(F^*(\lambda), V) - \nabla F^*(\lambda, V) - F^*\left(\frac{D\lambda}{dt}\right), \delta q \right\rangle. \quad (2.52)$$

4. $\frac{d}{d\epsilon} \langle \lambda, P(q) \rangle$

This term evaluates to

$$\begin{aligned} \frac{d}{d\epsilon} \langle \lambda, P(q) \rangle &= \langle \lambda, \nabla P(\delta q) \rangle \\ &= \langle \delta q, \nabla P^*(\lambda) \rangle. \end{aligned} \quad (2.53)$$

$$5. \frac{d}{d\epsilon} \langle \lambda, Y_i(q)u^i \rangle$$

This term evaluates similarly to the previous one to be

$$\frac{d}{d\epsilon} \langle \lambda, Y_i(q)u^i \rangle = \langle \delta q, \nabla Y_i^* u^i(\lambda) \rangle. \quad (2.54)$$

We now have expressed the variation of each term in the form $\langle \cdot, \delta q \rangle$. If we collect the expressions in Equations (2.45),(2.50),(2.52),(2.53), and (2.54) and set the integrand to zero by setting $\delta q = 0$, we arrive at the following equations of motion for λ . The equation is stated explicitly below in Theorem 2.2.

2.2.5 Variations of u : stationarity condition

We now turn to the requirement that cost function be a critical point with respect to variations of u . Since $u \in \mathbb{R}^m$, this variation is far simpler to calculate. The variation with respect to u^i yields the equation

$$\frac{\partial J}{\partial u^i} \delta u^i + \langle \lambda, Y_i \delta u^i \rangle = 0. \quad (2.55)$$

Since δu^i is a scalar, we can pull it outside the pairing, collect terms, and arrive at

$$\frac{\partial J}{\partial u^i} + \langle \lambda, Y_i \rangle = 0. \quad (2.56)$$

We thus have m conditions which define u as a function of q, V, λ . If we assume that the cost function is smooth and convex with respect to u , and that there are no constraints on the system states or controls, then Equations (2.31),(2.57) and (2.56) together represent necessary conditions which the optimal trajectory must satisfy.

These results are summarized in the following theorem:

Theorem 2.2. *A necessary condition for a trajectory $x(t)$ to optimize $J(q, u, V)$ is the existence of $\lambda(t)$ and $u(t)$ satisfying*

$$\begin{aligned} \frac{D^2 \lambda}{dt^2} - T^* \left(\frac{D\lambda}{dt}, V \right) + R^*(\lambda, V)V &= \sum_i \left[\left(\widehat{\nabla} \widehat{A}_i - r_i \nabla_V \widehat{A}_i \right) (V^{r_i}) \right. \\ &\quad \left. - r_i (r_i - 1) \widehat{A}_i (F(q, V) + P(q) + Y_i(q)u^i, V^{r_i-2}) + r_i T^* (\widehat{A}_i (V^{r_i-1}), V) \right] \\ &+ \widehat{\nabla} \widehat{F}^*(\lambda, V) + T^*(F^*(\lambda), V) - \nabla F^*(\lambda, V) - F^* \left(\frac{D\lambda}{dt} \right) + \nabla P^*(\lambda) \\ &+ \sum_{i=1}^m (\nabla Y_i^* u^i(\lambda)) \end{aligned} \quad (2.57)$$

and Equation (2.56).

As desired, we arrive at a second order equation on T^*Q rather than a first-order equation on TT^*Q . While the notation is cumbersome, it is not particularly complicated. In many cases, the dissipation and/or torsion terms will be zero, which will simplify the expansion significantly.

2.2.6 Endpoint Conditions

In this section, we state the endpoint conditions needed to solve the differential equations which govern the optimal control system. The optimal control ODE is now two second-order differential equations, one on Q , the other on T^*Q . We therefore need four endpoint conditions to solve for the controller which generates the optimal trajectory. We consider the two optimal control problems separately:

P1: In this case, the initial conditions $q(t_0), V(t_0)$ and the final time conditions $q(t_f), V(t_f)$ are fixed. These conditions completely specify optimal control and trajectory. Because the endpoints are fixed, the variations of the terminal cost, as well as the endpoint terms generated by the integrations by parts, are zero.

P2: In this case, the final time conditions $q(t_f), V(t_f)$ are not supplied. This means that the variation of the trajectory at t_f is nonzero, and therefore we derive the final time condition by considering the variations of the terminal cost and the endpoint terms derived from the integrations by parts. Specifically, using Equation (2.9), we have

$$\delta J_f = \sum_{i=1}^w \left\langle \widehat{\nabla} \widehat{\phi}_i(V^{b_i}), \delta q \right\rangle + b_r \left\langle \widehat{\phi}_i(V^{b_i-1}), \frac{DV}{D\epsilon} \right\rangle \quad (2.58)$$

with all terms evaluated at t_f . The integrations by parts of Equations (2.12),(2.49), and (2.51) extract several terms from the integrand, and we consider their variations now. Their variations are also zero at t_0 , but at the final time, we have

$$\delta J|_{t=t_f} = \sum_{i=1}^s r_i \left\langle \widehat{A}_i(V^{r_i-1}), \delta q \right\rangle + \left\langle \lambda, \frac{DV}{d\epsilon} \right\rangle - \left\langle \frac{D\lambda}{dt}, \delta q \right\rangle + \langle F^*(\lambda), \delta q \rangle. \quad (2.59)$$

Summing these and setting to zero, we have

$$\begin{aligned} & \left\langle \sum_{i=1}^w \widehat{\nabla} \widehat{\phi}_i(V^{b_i}) + \sum_{i=1}^s r_i \widehat{A}_i(V^{r_i-1}) - \frac{D\lambda}{dt} + F^*(\lambda), \delta q \right\rangle + \\ & \left\langle \lambda + \sum_{i=1}^w b_r \widehat{\phi}_i(V^{b_i-1}), \frac{DV}{d\epsilon} \right\rangle = 0. \end{aligned} \quad (2.60)$$

At a given point in time, the term $\frac{DV}{d\epsilon}$ is independent of δq . We thus have two final time conditions:

$$\lambda(t_f) = - \sum_{i=1}^w b_r \widehat{\phi}_i(V^{b_i-1}) \quad (2.61)$$

$$\frac{D\lambda}{dt}(t_f) = \sum_{i=1}^w \widehat{\nabla} \widehat{\phi}_i(V^{b_i}) + \sum_{i=1}^s r_i \widehat{A}_i(V^{r_i-1}) + F^*(\lambda). \quad (2.62)$$

We thus recover two final-time conditions for $\lambda, \frac{D\lambda}{dt}$ which, when paired with $q(t_0), V(t_0)$, provide endpoint conditions to solve for the optimizing controller.

2.3 Examples

2.3.1 Splines on Manifolds

One example to which this theory can be applied is the problem of calculating force-minimizing curves which link two points in the tangent space of some Riemannian manifold Q . While this is not a vehicle control problem, it is a problem which has historically been studied in this context [29, 83]. In this example, the system is fully actuated, so the problem can be stated as an unconstrained higher-order variational problem, and there is no need to use adjoint variables. Nonetheless, the theory developed here covers this case.

The problem statement is to minimize

$$J = \frac{1}{2} \int_{t_0}^{t_f} \langle\langle u, u \rangle\rangle dt \quad (2.63)$$

subject to

$$\frac{DV}{dt} = u. \quad (2.64)$$

Substituting these into Equations (2.57),(2.56) yields the conditions:

$$\frac{D^2\lambda}{dt^2} + R^*(\lambda, V)V = 0 \quad (2.65)$$

(recall that the Riemannian metric is torsion-free and has the property $\nabla g = 0$), and

$$\lambda = g(u). \quad (2.66)$$

Substituting in the equation for λ yields

$$\frac{D^2}{dt^2} \left(g \frac{DV}{dt} \right) + R^* \left(g \frac{DV}{dt}, V \right) V = 0. \quad (2.67)$$

Using the cyclic properties of the curvature tensor and Proposition 2.2, we can rewrite this equation as

$$g \left(\frac{D^3 V}{dt^3} \right) + g \left(R \left(\frac{DV}{dt}, V \right) V \right) = 0 \quad (2.68)$$

which is equivalent to

$$\frac{D^3 V}{dt^3} + R \left(\frac{DV}{dt}, V \right) V = 0 \quad (2.69)$$

which matches the results of [29, 83].

2.3.2 Planar Rigid Body

In this example, we consider a rigid body free to move in the plane, actuated by a force applied in the body frame applied at a distance h from the center of mass. A diagram is shown in Figure 2.1. The Riemannian metric for this system is

$$\langle\langle V, V \rangle\rangle = \frac{1}{2} V^T \begin{pmatrix} m_1 \cos^2 \theta + m_2 \sin^2 \theta & (m_1 - m_2) \cos \theta \sin \theta & 0 \\ (m_1 - m_2) \cos \theta \sin \theta & m_1 \sin^2 \theta + m_2 \cos^2 \theta & 0 \\ 0 & 0 & J \end{pmatrix} V \quad (2.70)$$

where $V = [\dot{x}, \dot{y}, \dot{\theta}]$, the velocity in the inertial frame. When $m_1 = m_2 = m$, metric corresponds to a hovercraft of mass m . When the masses are unequal, it corresponds to the motion of an underwater body, with different effective masses in the directions of motion depending on the orientation of the body relative to its motion in the fluid. We note in passing that the Riemannian metric, and hence the associated Lagrangian, is $SE(2)$ -invariant, and the equations of motion admit a simpler characterization using reduction (see [62] for a thorough discussion of full $SE(3)$ underwater motion). Since our goal is to expose the role of the connection, we consider the undreduced vehicle motion.

The governing equation of motion is

$$\frac{DV}{dt} = Y_i u^i, \quad (2.71)$$




Figure 2.1: Diagram of Planar Rigid Body.

where

$$Y_1 = \begin{pmatrix} \frac{\cos \theta}{m_1} \\ \frac{\sin \theta}{m_1} \\ 0 \end{pmatrix}, \quad Y_2 = \begin{pmatrix} \frac{-\sin \theta}{m_2} \\ \frac{\cos \theta}{m_2} \\ -\frac{h}{J} \end{pmatrix} \quad (2.72)$$

are the control vector fields. Again, we choose our cost function to be force minimization:

$$J = \langle\langle Y_1 u^1, Y_1 u^1 \rangle\rangle + \langle\langle Y_2 u^2, Y_2 u^2 \rangle\rangle = \frac{(u^1)^2}{2m_1} + \frac{(J + m_2 h^2)(u^2)^2}{2Jm_2}. \quad (2.73)$$

When $m_1 = m_2$, the coefficients of g are all constant, and hence the Christoffel symbols are $\Gamma_{ij}^k = 0$, and thus $\frac{DV}{dt} = [\ddot{x}, \ddot{y}, \ddot{\theta}]$. When $m_1 \neq m_2$, the Christoffel symbols are nontrivial. The torsion tensor is zero because we use the Levi-Civita connection, but the curvature tensor is not. Hence, applying Equation (2.57), we recover a simple equation for the adjoint variables:

$$\frac{D^2 \lambda}{dt^2} + R^*(\lambda, V)V = \nabla Y_i^*(\lambda) u^i. \quad (2.74)$$

When $m_1 = m_2$, the curvature tensor is zero, and the dynamics evaluate in coordinates to

$$\ddot{\lambda}_x = 0 \quad (2.75)$$

$$\ddot{\lambda}_y = 0 \quad (2.76)$$

$$\ddot{\lambda}_\theta = \frac{1}{m} [(-\lambda_x \sin \theta + \lambda_y \cos \theta) u_1 + (-\lambda_x \cos \theta - \lambda_y \sin \theta) u_2] \quad (2.77)$$

and the optimal control is given by

$$u_1 = -(\lambda_x \cos \theta + \lambda_y \sin \theta) \quad (2.78)$$

$$u_2 = \frac{Jm}{J + mh^2} \left(\frac{\lambda_x}{m} \sin \theta - \frac{\lambda_y}{m} \cos \theta + \frac{h}{J} \lambda_\theta \right). \quad (2.79)$$

2.4 Conclusions

The equations derived in this chapter reveal the relevance of the affine connection in formulating the optimal control equations of motion. Because they are derived using a variational approach, they are less general than the equations derived by Lewis in [66], which proceed directly from the Pontryagin Maximum Principle. However, understanding the relationship between the affine connection and the calculus of variation provides insight into the structure of these equations and lays the framework for future analysis of topics such as the role of symmetry and generating approximate solutions. Topics for future work will be discussed in detail in Chapter 7.

Chapter 3

Optimal Control of Time-Scalable Systems

In this chapter, we turn our attention to the other classical approach to optimal control, namely the Hamilton-Jacobi-Bellman (HJB) equation. As in the previous chapter, our goal is to understand the role which the structure of vehicle dynamics plays in this formulation of the optimal control problem. In this case, we will restrict ourselves to the class of systems which are time scalable, meaning the equations of motion remain the same when time is reparametrized. Examples of time scalable dynamic systems include driftless systems and mechanical systems without potential energy. Vehicle applications which fall into this category include locomotion of systems with nonholonomic constraints, control of underwater vehicles, and satellite reorientation [7, 17, 56].

The chapter is organized as follows. In Section 3.1, we define time-scalable control systems and derive relevant properties. Specific attention is paid to properties of functions which are compatible with the time scaling. In Section 3.2, we apply these ideas to the HJB equation. We consider the derivation of both finite-time optimal trajectories and stabilizing control laws from the HJB value function. The example of the kinematic wheeled locomotion is examined in detail.

3.1 Time Scaling

3.1.1 Definitions

Let $x \in \mathbb{R}^p$ $u \in \mathbb{R}^s$. Let an n th order control system be a dynamical system of the form $x^{(n)} = f(x, \dot{x}, \dots, x^{(n-1)}, u)$, where the superscript in parentheses indicates the order of differentiation. Let $\bar{x} \in \mathbb{R}^{pn}$ be a point of the form $\bar{x} = (x, \dot{x}, \dots, x^{(n-1)})$.

Definition 3.1. An n th order control system is said to be *affinely time scalable*

if, for any affine scaling of time $\tau = \alpha t + \beta$, there exists a scaling of u such that the scaled equations of motion are identical to the original equations. Let $K : \mathbb{R}^m \times \mathbb{R}^s \times \mathbb{R} \rightarrow \mathbb{R}^m \times \mathbb{R}^s \times \mathbb{R}$ denote the time scaling operator associated with a given time scaling. That is, $K(\bar{x}, u, t)$ maps each of the operands to its time-scaled counterpart.

We can easily write how K acts on the various operands by applying the chain rule. We see that

$$\frac{dx}{dt} = \frac{dx}{d\tau} \frac{d\tau}{dt} \quad (3.1)$$

$$= \alpha \frac{dx}{d\tau} \quad (3.2)$$

and thus we see that K maps \dot{x} to \dot{x}/α . More generally, we can see that K maps $x^{(j)}$ to $x^{(j)}/\alpha^j$. Because x is independent of time, it does not vary by time scaling. To determine if a control system is time scalable, one simply makes this substitution and then sees if u can be compatibly scaled such that α can be factored out of the equation.

Remark Because K acts on each operand independently (as well as acting on each $x^{(i)}$ independently), we will use, for example, $K(u)$ to denote the image of u under the time scaling.

Example 3.1 (Driftless Systems). First-order driftless systems, meaning systems of the form $\dot{x} = f(x)u$, are time scalable. If we scale time and substitute, we arrive at the equation

$$\frac{dx}{d\tau} = f(x) \frac{u}{\alpha}. \quad (3.3)$$

If we define $K(u) = u/\alpha$, and substitute, we recover the original equation. More generally, systems of the form $x^{(n)} = f(x)u$, are also time scalable, with the substitution $K(u) = u/\alpha^n$.

Example 3.2 (Affine Connection Control Systems). The affine connection control system of Equation 2.31, when $F, P = 0$, is

$$\frac{DV}{dt} = \sum_i Y^i(q)u^i. \quad (3.4)$$

Because the covariant derivative is linear in \ddot{q} and quadratic in \dot{q} (see Equation (2.1)), this system is time scalable. When $u = 0$, the time-scalability corresponds

with the well-known time scalability of geodesics [35]. When u is nonzero, it can still be scaled by the procedure noted above.

Definition 3.2. A point \bar{x} which exhibits the property $K(\bar{x}) = \bar{x}$ is said to be *scale invariant*. A scaling K for which $K(T) = T$, where $T \in \mathbb{R}$, is said to be *T-invariant*.

Remark The scale invariant points are clearly those for which all derivatives are zero. For a first-order control system, these consist of the entire space. A time scaling is T -invariant if $\beta = (1 - \alpha)T$. We see that the function

$$r(t) = \frac{1}{(T - t)^n} \quad (3.5)$$

is compatible with any T -invariant time scaling, and the scale factor is $f(\alpha) = \alpha^{-n}$.

We now introduce a class of functions whose behavior under time scaling we wish to investigate:

Definition 3.3. A function $L(x, \dot{x}, \dots, x^{\{n-1\}}, u, t)$ is said to be *compatible* with a time-scaling operation K if the following holds there exists a function $f : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$L(K(x, \dot{x}, \dots, x^{\{n-1\}}, u, t)) = f(\alpha)L(x, \dot{x}, \dots, x^{\{n-1\}}, u, t). \quad (3.6)$$

Remark An obvious property of $f(\alpha)$ is that $f(1) = 1$. Also, it is easy to see that $f(\alpha^{-1}) = f(\alpha)^{-1}$.

We see that the class of time-scalable systems is fairly rich, and includes systems for which significant work has been done in the areas of controllability and trajectory generation. To our knowledge, however, the time scaling property of these systems has never explicitly been exploited. In the development which follows, we discuss the implications time scalability has for the optimal control problem.

3.1.2 Properties of Time-Scalable Systems

The following proposition follows immediately:

Proposition 3.1. *Given a time-scalable control system, initial conditions \bar{x}_i at initial time $t = t_i$, a control time history $u(t)$ defined on the interval $t \in [t_i, t_f]$, and a resulting trajectory $\bar{x}(t)$ defined on the interval, the time scaled control $K(u)$*

defined on the interval $K([t_i, t_f])$, when applied to the time-scaled initial condition $K(\bar{x}_i)$ at time $K(t_i)$, will produce the time-scaled trajectory $K(\bar{x}(K(t)))$.

Remark In other words, the time scaling of an integral curve yields another integral curve generated by the time-scaled control history. In particular, if the initial and final points are scale invariant, then a control time history which connects the two points in a given time interval can be scaled to connect those two points in any interval of time.

We now apply Proposition 3.1 to derive a property of integrals of compatible functions.

Proposition 3.2. *Given a time-scalable control system, a time scaling K for which $\alpha \geq 0$, a function $L(\bar{x}, u)$ which is compatible with K and for which $f(\alpha) \geq 0$, and a control history $u^*(t)$, $t \in [t_i, t_f]$ which drives the system between $m_i \in M$ at time $t = t_i$ and $m_f \in M$ at time $t = t_f$ and minimizes the quantity*

$$J = \int_{t_i}^{t_f} L(\bar{x}, u) dt. \quad (3.7)$$

Then the time-scaled control history $K(u^(K(t)))$ minimizes J between the time-scaled endpoints.*

Proof. To see the correspondence of J with its scaled counterpart, let us evaluate

$$\tilde{J} = \int_{K(t_i)}^{K(t_f)} L(\bar{x}, v) d\tau \quad (3.8)$$

with endpoints $K(m_i)$ and $K(m_f)$. We can evaluate the integral in the unscaled interval by changing coordinates through K^{-1} . Letting $t = K^{-1}(\tau)$, we apply the formula for change of coordinates of an integral to arrive at

$$\tilde{J} = \int_{t_i}^{t_f} L(K(\bar{x}, v)) \alpha dt, \quad (3.9)$$

where the extra α term is the Jacobian of the coordinate transformation. Replacing v with $K(u)$ and applying the definition of compatibility yields

$$\tilde{J} = \int_{t_i}^{t_f} \alpha f(\alpha) L(\bar{x}, u) dt \quad (3.10)$$

or

$$\tilde{J} = \alpha f(\alpha) J \quad (3.11)$$

which is clearly minimized by $u = u^*$, or $v = K(u^*)$, since the quantity $\alpha f(\alpha)$ is assumed positive. ■

Proposition 3.2 guarantees that $K(u^*)$ does in fact drive the system to the desired endpoint, and hence is an admissible control history. Furthermore, it guarantees that every control history which is admissible for the scaled function is also admissible for the original function. Hence no other control history could exist which yields a lower J than $v = K(u^*)$.

We now consider the properties of the integral function J for a special case. Suppose the final point m_f of the previous proposition is scale invariant, and let us fix the final time t_f . The function J can be thought of as a function of the initial point \bar{x} and the initial time $t_i \in (-\infty, t_f)$. It is not an explicit function of u , since u is chosen according to the initial position and time. From the proof above, the following proposition holds:

Proposition 3.3. *Given the above conditions, $J(\bar{x}, t_i)$ is compatible with the set of all t_f -invariant scalings. Specifically,*

$$J(K(\bar{x}, t_i)) = \alpha f(\alpha) J(\bar{x}, t_i). \quad (3.12)$$

Finally, we are able to use compatibility to prove the following about the partial derivative of compatible functions:

Proposition 3.4. *Let $L(\bar{x}, u, t)$ be a function which is compatible with the set of all T -invariant time scalings. For simplicity, assume $K(u) = u/\alpha^n$. Then, for all $t_0 \neq T$,*

$$\left. \frac{\partial L}{\partial t} \right|_{t=t_0} = \frac{1}{t_0 - T} \left(\left. \frac{\partial f}{\partial \alpha} \right|_{\alpha=1} L + \sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right). \quad (3.13)$$

Proof. Recall the definition of the partial derivative:

$$\left. \frac{\partial L}{\partial t} \right|_{t=t_0} = \lim_{\Delta t \rightarrow 0} \frac{L(\bar{x}, u, t_0 + \Delta t) - L(\bar{x}, u, t_0)}{\Delta t}. \quad (3.14)$$

We can map the first term on the right-hand side using the time scaling operator. If we set

$$\alpha = \frac{t_0 - T}{t_0 + \Delta t - T} \quad (3.15)$$

$$\beta = \frac{T \Delta t}{t_0 + \Delta t - T}, \quad (3.16)$$

we recover the unique T -invariant time-scaling for which $K(t_0 + \Delta t) = t_0$. Thus,

$$L(\bar{x}, u, t_0 + \Delta t) = \frac{1}{f(\alpha)} L \left(x, \frac{\dot{x}}{\alpha}, \dots, \frac{x^{(n-1)}}{\alpha^{n-1}}, \frac{u}{\alpha^n}, t_0 \right). \quad (3.17)$$

Now $\alpha^{-1} = 1 + \Delta t / (t_0 - T)$, and $\alpha^{-n} = 1 + n\Delta t / (t_0 - T) + \mathcal{O}(\Delta t^2)$. Using this, we can Taylor expand the right-hand side to first order, to arrive at

$$L(\bar{x}, u, t_0 + \Delta t) = \frac{1}{f(\alpha)} \left[L + \left(\sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right) \frac{\Delta t}{t_0 - T} \right] + \mathcal{O}(\Delta t^2), \quad (3.18)$$

where the right-hand side is evaluated at $(x, \dot{x}, \dots, x^{(n-1)}, u, t_0)$. If we recall that at $\Delta t = 0$, we have $\alpha = 1$ and thus $f(\alpha)$ goes to 1 in the limit. Substituting into Equation (3.14) and taking the limit, we have

$$\frac{\partial L}{\partial t} \Big|_{t=t_0} = \frac{1}{t_0 - T} \left(\sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right) + \lim_{\Delta t \rightarrow 0} \frac{1 - f(\alpha)}{f(\alpha)\Delta t} L. \quad (3.19)$$

The final term can be simplified by noting that $\alpha = 1 - \frac{\Delta t}{t_0 + \Delta t - T}$, Taylor expanding $f(\alpha)$ to first order, and taking the limit as before. This results in the following:

$$\frac{\partial L}{\partial t} \Big|_{t=t_0} = \frac{1}{t_0 - T} \left(\frac{\partial f}{\partial \alpha} \Big|_{\alpha=1} L + \sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right), \quad (3.20)$$

which proves our proposition. ■

The significance of this proposition is that it shows how the *temporal* partial derivative can be converted to a *spatial* quantity.

3.2 Optimal Control of Time-Scalable Systems

3.2.1 Application of Time-Scalability to the HJB Equation

We are now ready to consider the relationship between time scaling and the optimal control problem. Recall that the value function $V(\bar{x}, t)$ associated with the HJB equation is defined as the minimum cost necessary to drive a system from a point \bar{x} at time t to a desired final condition at time (not necessarily finite) T . To apply the above theory to this problem, we consider a finite time optimal control problem with fixed final state which is scale invariant, and a cost function $L(\bar{x}, u, t)$ which

is compatible with all time scalings. Given this, we see that the value function is given by

$$V(\bar{x}, t) = \int_t^T L(\bar{x}, u, \tau) d\tau. \quad (3.21)$$

Applying Proposition 3.3, we see that $V(\bar{x}, t)$ is compatible with all T -invariant time scalings, and hence Proposition 3.4 applies as well.

The HJB equation for this optimal control problem is given by [13, 93]

$$\frac{\partial V}{\partial t} = - \min_u \left[L(\bar{x}, u, t) + \frac{\partial V}{\partial \bar{x}} g(\bar{x}, u) \right]. \quad (3.22)$$

Proposition 3.4 allows us to replace the left-hand side with a term involving V and its spatial partial derivatives. Solving this spatial PDE yields the value function at one point in time, and the value function at any other time can be found using the time scaling. The presence of the $\frac{1}{t-T}$ term in Equation (3.13) indicates that as t approaches T , the value function will approach infinity at all points except the desired endpoint. This is logical — as time runs out, the cost associated with driving the system to the endpoint increases.

Ordinarily, the HJB equation is solved by defining $V(\bar{x}, T)$ using a terminal cost function and propagating that function backwards in time. In our case, the $\bar{x}(T)$ is fixed, so a terminal cost constraint is not meaningful. Nonetheless, the Principle of Optimality from which the HJB equation is derived is still applicable, so the use of that equation is still valid.

In the next sections, we consider the cases of first-order systems and affine connection control systems in greater detail. For the former, we shall show that solving the spatial PDE yields a stabilizing controller in addition to solving the finite-horizon optimal control problem.

3.2.2 Driftless Systems

Let us consider a system of the form

$$\dot{x} = g(x)u, \quad (3.23)$$

where $g(x)$ is C^1 , and a cost function

$$J = \frac{1}{2} \int_{t_0}^T u^T u dt, \quad (3.24)$$

and suppose the constraint was that $x(T) = 0$. Evaluating the HJB equation for this system yields

$$V_t = \frac{1}{2}V_x^T g(x)g^T(x)V_x \quad (3.25)$$

and optimal control

$$u^*(x, t) = -g^T(x)V_x, \quad (3.26)$$

where the subscripts on V indicate partial derivatives. We now turn to replacing V_t with its spatial counterpart. The cost function $u^T u$ is compatible with any time scaling, and the scaling function is $f(\alpha) = \alpha^{-2}$. The scaling for V , following Proposition 4, is $f(\alpha) = \alpha^{-1}$. Applying Proposition 3.4, we have

$$V_t = \frac{1}{T-t}V. \quad (3.27)$$

Furthermore, applying the definition of time scalability, we see that

$$V(x, t) = \frac{T-t_0}{T-t}V(x, t_0). \quad (3.28)$$

Letting $t_0 = 0$ and defining $\tilde{V}(x) = V(x, 0)$, we can substitute into Equation (3.25) and derive the following spatial PDE:

$$\frac{1}{T}\tilde{V} = \frac{1}{2}\tilde{V}_x^T g(x)g^T(x)\tilde{V}_x. \quad (3.29)$$

Solving this PDE yields the optimal finite-horizon control law by finding

$$\tilde{u}(x) = -g^T(x)\tilde{V}_x \quad (3.30)$$

and scaling $\tilde{u}(x)$ appropriately in time to arrive at $u^*(x, t)$. Note that t was eliminated from the equation via the substitution, as expected.

We now consider the possibility of using $\tilde{u}(x)$ as a control law without time scaling, and asking whether that control law stabilizes the system. When V is smooth, it is easy to show that it acts as a Lyapunov function for the control law $\tilde{u}(x)$:

$$\begin{aligned} \dot{\tilde{V}} &= \tilde{V}_x^T \dot{x} \\ &= \tilde{V}_x^T g(x)u^*(x) \\ &= -\tilde{V}_x^T g(x)g^T(x)\tilde{V}_x \\ &= -\frac{2}{T}\tilde{V}. \end{aligned} \quad (3.31)$$

Since \tilde{V} is a value function, $\dot{\tilde{V}}$ must be negative at all points $x \neq 0$. Thus, solving a single spatial PDE yields both optimal finite-horizon trajectories and an asymptotically stabilizing controller whose gain can be set. Conceptually, a controller of this form can be thought of as applying the optimal control one would use if the final time were always $T-t_0$ away, reminiscent of receding-horizon control concepts used in model predictive control [45, 86]. Note further that by varying the value of t_0 used to define \tilde{V} , we can influence the rate of convergence of \tilde{V} , i.e., the gain of the controller. Whether or not the convergence of the system is exponential will depend on whether the value function is locally quadratic, a fact which will depend on the underlying system.

Finally, we note in passing that Equation (3.24) is not the only compatible cost function. A function of the form

$$J = \int_{t_0}^T \frac{a}{(T-\tau)^2} x^T x + bu^T u d\tau \quad (3.32)$$

is also compatible. This cost function can be viewed as forcing $x(t)$ to approach zero sufficiently quickly to avoid the cost function becoming infinite.

3.2.3 Affine Connection Control Systems

In this section, we consider systems of the form

$$\ddot{q}^i + \Gamma_{jk}^i \dot{q}^j \dot{q}^k = g^{ij}(q) u_j \quad (3.33)$$

which is Equation (3.4) in coordinates, with the cost function in Equation (3.24). In this case, the scaling function is $f(\alpha) = \alpha^{-4}$, and the scaling function for V is $\alpha f(\alpha) = \alpha^{-3}$. Thus, following Proposition 3.4, we have

$$V_t = \frac{1}{T-t} (3V - V_{\dot{q}} \dot{q}). \quad (3.34)$$

As in the first-order case, we can derive the following equation for V :

$$V(q, \dot{q}, t) = \frac{(T-t_0)^3}{(T-t)^3} V(q, \frac{T-t}{T-t_0} \dot{q}, t_0). \quad (3.35)$$

It is easily verified that this form for V satisfies the equation for V_t . Again, we set $t_0 = 0$ and denote the spatial PDE as $\tilde{V}(x, \dot{x}) = V(q, \dot{q}, 0)$. We use x for the operands of \tilde{V} since the operands of V and \tilde{V} are not the same. The relationship

between the two is

$$x = q, \quad \dot{x} = \frac{T-t}{T}\dot{q}, \quad (3.36)$$

a fact we shall use shortly, but we want to make clear the difference between partial differentiation with respect to the second operand of V , i.e., x_2 , and with respect to q_2 .

The HJB PDE for this system is

$$V_t = -V_q\dot{q} + V_{\dot{q}}\Gamma_{jk}^i\dot{q}^j\dot{q}^k + \frac{1}{2}V_{\dot{q}}^T g(q)g^T(q)V_{\dot{q}}. \quad (3.37)$$

We take the partials of Equation (3.35) with respect to the relevant terms and substitute in Equations (3.35),(3.36), we arrive at

$$\frac{3}{T}\tilde{V} - \frac{1}{T}\tilde{V}_x\dot{x} = -\tilde{V}_x\dot{x} + \tilde{V}_x\Gamma_{jk}^i\dot{x}^j\dot{x}^k + \frac{1}{2}\tilde{V}_x^T g(x)g^T(x)\tilde{V}_x. \quad (3.38)$$

If we attempt to use this value function as a Lyapunov function as done earlier, we arrive at the following equation:

$$\dot{\tilde{V}} = -\frac{3}{T}\tilde{V} - \frac{1}{2}\tilde{V}_x^T g(x)g^T(x)\tilde{V}_x + \frac{1}{T}\tilde{V}_x\dot{x}. \quad (3.39)$$

The above equation is not necessarily negative due to the last term. Given a solution \tilde{V} , this equation can be evaluated to determine if \tilde{V} is in fact a Lyapunov function.

3.2.4 Example: Wheeled Locomotion

As an example, we consider optimal control of a planar cart actuated by two wheels whose angular rate can be specified. A diagram of such a cart is shown in Figure 3.1. The (normalized) equations of motion for this cart are

$$\begin{aligned} \dot{x} &= v_1 \cos \theta \\ \dot{y} &= v_1 \sin \theta \\ \dot{\theta} &= v_2 \end{aligned} \quad (3.40)$$

where v_r, v_l are the angular velocities of the left and right wheels, $v_1 = \frac{1}{2}(v_r + v_l)$, and $v_2 = \frac{1}{2}(v_r - v_l)$. This system is driftless and underactuated. It is known (see, e.g., [79]) that this system can be locally transformed to the following form:




Figure 3.1: Diagram of Kinematic Cart.

$$\begin{aligned}
 \dot{x}_1 &= u_1 \\
 \dot{x}_2 &= u_2 \\
 \dot{x}_3 &= x_1 u_2 - x_2 u_1.
 \end{aligned} \tag{3.41}$$

This system is the well-known first-order nonholonomic integrator [11]. When we employ the cost function

$$J = \frac{1}{2} \int_{t_0}^T u_1^2 + u_2^2 dt \tag{3.42}$$

and the constraint that the states are all zero at $t = T$, then the optimal trajectories of this system, derived using the Euler-Lagrange method, are known to take the form

$$\begin{aligned}
 u_1 &= k_1 \sin \mu t + k_2 \cos \mu t \\
 u_2 &= k_1 \cos \mu t - k_2 \sin \mu t,
 \end{aligned} \tag{3.43}$$

where k_1, k_2, μ are parameters chosen to meet the endpoint conditions. The HJB equation for this system is

$$\frac{\partial V}{\partial t} = \frac{1}{2} \frac{\partial V}{\partial x} \begin{pmatrix} 1 & 0 & -x_2 \\ 0 & 1 & x_1 \\ -x_2 & x_1 & x_1^2 + x_2^2 \end{pmatrix} \frac{\partial V}{\partial x}. \tag{3.44}$$

Following the development in Section 3.2.2, we can rewrite this as

$$\frac{1}{T}\tilde{V} = \frac{1}{2}\frac{\partial\tilde{V}}{\partial x}^T \begin{pmatrix} 1 & 0 & -x_2 \\ 0 & 1 & x_1 \\ -x_2 & x_1 & x_1^2 + x_2^2 \end{pmatrix} \frac{\partial\tilde{V}}{\partial x}. \quad (3.45)$$

We now seek a functional form for \tilde{V} which is consistent with this equation. If we assume that $\tilde{V} = \tilde{V}(d, x_3)$ where $d = x_1^2 + x_2^2$ (a fact which can be deduced from the symmetry of the problem), we can substitute and simplify the equation to

$$\frac{1}{T}\tilde{V} = 2d \left(\frac{\partial\tilde{V}}{\partial d} \right)^2 + \frac{d}{2} \left(\frac{\partial\tilde{V}}{\partial x_3} \right)^2. \quad (3.46)$$

If we further assume that V has the form

$$\tilde{V} = \frac{d}{2T}f(r) \quad (3.47)$$

where $r = x_3/d$, we can reduce the equation to an ODE:

$$f = f^2 - 2rf f' + \left(r^2 + \frac{1}{4} \right) f'^2. \quad (3.48)$$

This equation can be factored to form two ODEs, namely

$$f' = \frac{2 \left(2rf \pm \sqrt{4r^2 f^2 - (4r^2 + 1)(f^2 - f)} \right)}{4r^2 + 1}. \quad (3.49)$$

To solve either ODE, we need an initial condition. Note that when $r = 0$, this is equivalent initial conditions in the x_1x_2 plane. Reviewing the equations of motion, we see that this case can be solved directly using constant $u_1 = x_1/T(T-t), u_2 = x_2/(T-t)$. This corresponds to $V(d, 0) = \frac{d}{2T}$, meaning $f(0) = 1$. Note that for $f = 1, r \geq 0$, the ODE corresponding to the positive square root is a fixed point of the system, and for $r \leq 0$ the other one is. Thus, by solving the initial value problem for both positive and negative r using the nontrivial ODE, we derive $f(r)$, and thus the value function as well.

Solving Equation (3.49) reveals that $\lim_{r \rightarrow \infty} f(r) = \infty$. However, $\lim_{r \rightarrow \pm\infty} V(d, r)$ is finite, as is the limit of $f'(r)$. This allows us to define V at points where $x_1, x_2 = 0$.

In fact, it can be computed that

$$\tilde{V}(0, 0, x_3) = \frac{\pi|x_3|}{T}. \quad (3.50)$$

However, \tilde{V} is only C^0 (and Lipschitz) at these points. To proceed further requires tools from nonsmooth analysis, which are beyond the scope of this discussion. It will suffice to say that it can be verified that our solution for \tilde{V} is a viscosity solution of the HJB PDE in the sense of [28].

At points away from the x_3 axis, we have the following control law:

$$\begin{aligned} \tilde{u}_1 &= -x_1(f - f'r) + \frac{1}{2}x_2f' \\ \tilde{u}_2 &= -x_2(f - f'r) - \frac{1}{2}x_1f'. \end{aligned} \quad (3.51)$$

Along the x_3 axis, the control laws cannot be derived using Equation (3.30). The limit $\lim_{r \rightarrow \infty} u_i$ exists but is path-dependent. It can be shown that along the x_3 axis a control law of the form

$$\begin{aligned} \tilde{u}_1 &= \sin \phi \sqrt{\frac{2\pi|x_3|}{T}} \\ \tilde{u}_2 &= \cos \phi \sqrt{\frac{2\pi|x_3|}{T}} \end{aligned} \quad (3.52)$$

for arbitrary ϕ yields a control law which is continuous along trajectories.¹ Note that the discontinuity of the control law is necessary for systems of this form, as proven by Brockett [12].

Figures 3.2 and 3.3 show optimal finite-time trajectories and stabilizing trajectories for two initial conditions. The stabilizing trajectories were generated using the value function at $t_0 = 9$ to give faster convergence. The initial conditions in Figure 3.2 are at a point where the value function is smooth. Those for Figure 3.3 are at a nonsmooth point. Both show the trajectory reaching zero at the appropriate time for the finite-horizon case, and decaying to zero in the infinite horizon case. For first-order systems such as this, the finite and infinite time trajectories are actually the same. Because the stabilizing controller does not have the $1/(T-t)$ term to increase the control effort, it takes infinite time to reach the origin.

The resulting trajectories are consistent with derived by applying the control law of Equation (refnhu1). We also note that for certain initial conditions, multiple values of k_1, k_2, μ satisfy those equations and endpoint conditions. The value

¹This control law is the boundary of the control law derived using the subdifferential of \tilde{V} in Equation (3.30).




Figure 3.2: Finite-Time Optimal (left) and Stabilizing Trajectories (right), $x_0 = (2, -3, 4)$.

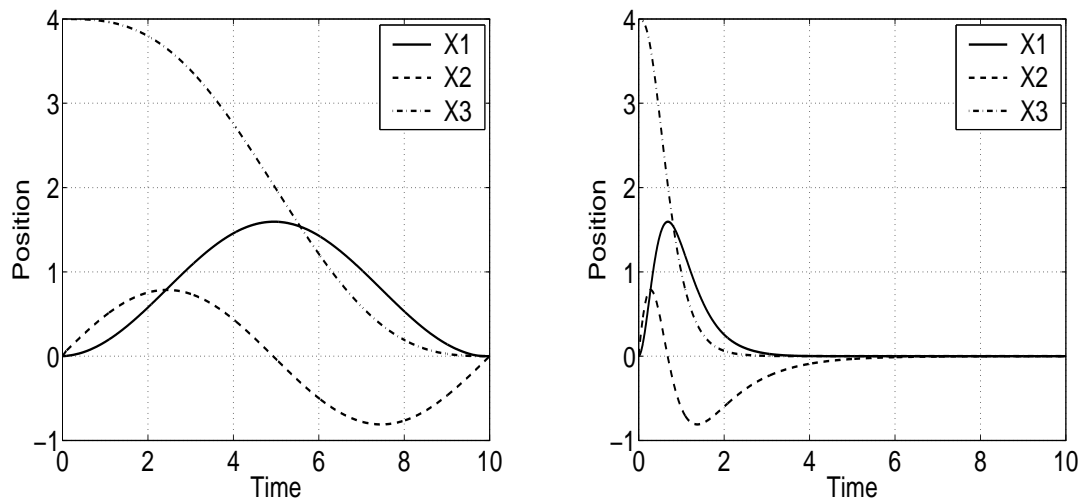


Figure 3.3: Finite-Time Optimal (left) and Stabilizing Trajectories (right), $x_0 = (0, 0, 4)$.

function derived above returns the globally minimizing trajectory.

3.3 Conclusions

In this chapter, we showed how the HJB equation can be simplified through *a priori* knowledge of the solution structure corresponding to the structure of the underlying dynamical system and cost function. For this class of systems, solving a single simplified PDE yields both finite time trajectories and a stabilizing control law. In practice, this PDE may not be easier to solve than the original one. In Chapter 7, we will discuss the utility of this formulation and how it may be exploited in practice.

Chapter 4

Graph Theory

4.1 Introduction

In this chapter, we introduce notation and basic concepts from graph theory that will be useful in our analysis of vehicle formation structure. In Section 4.2, we introduce basic terminology from graph theory and focus on connectivity structure in graphs. In Section 4.3, we introduce ideas from algebraic graph theory that will be useful in linking graph-theoretic and control-theoretic concepts. Specifically, we define the Laplacian matrix of a graph, which we will later identify in Section 5.1 as the matrix of interest in vehicle formation control. In Section 4.4, we review ideas from Perron-Frobenius theory. Finally, in Section 4.5, we discuss the spectral structure of the Laplacian using ideas from the previous section.

Many excellent texts on graph theory exist; a recent example is [34]. Recent results regarding the Laplacian and its spectral structure can be found in the work of Merris [73, 74] and Chung [24]. Perron-Frobenius theory can be found in many texts; the presentation here is based on material in [2, 51, 98]. The terminology regarding connectivity in directed graphs is borrowed from [87], and that terminology is used to derive Theorem 4.5, which is a specialized case of the topic addressed in that paper.

We now introduce some notation which we will use in the remainder of the thesis. Variables in lower case refer to scalars, vectors or elements of sets; the distinction will be clear from context. Variables in upper case to matrices, and calligraphed letters refer to sets or graphs. When v is a vector, v_i refers to the i th element of that vector, and when v is a set, v_i refers to the i th indexed element of that set. $|\mathcal{G}|$ denotes the cardinality of the set \mathcal{G} . A_{ij} refers to the element occupying the i th row and j th column of A . I_n refers to the $n \times n$ identity matrix.

4.2 Introductory Graph Theory

4.2.1 Basic Definitions

A *directed graph* \mathcal{G} consists of a set of vertices, or nodes, denoted \mathcal{V} , and a set of arcs $\mathcal{A} \subset \mathcal{V}^2$, where $a = (v, w) \in \mathcal{A}$ and $v, w \in \mathcal{V}$. The first element of a is denoted $\text{tail}(a)$, and the second is denoted the $\text{head}(a)$. It is said that a points from v to w . We will assume that $\text{tail}(a) \neq \text{head}(a)$, meaning that the graph has no loops. We also assume that each element of \mathcal{A} is unique. A graph with the property that for any $(v, w) \in \mathcal{A}$, the arc $(w, v) \in \mathcal{A}$ as well is said to be *undirected*; in undirected graphs the pair of arcs is often modeled as a single edge with no direction associated to it. The *in(out)-degree* of a vertex v is the number of arcs with v as its head (tail). For an undirected graph, the in-degree and out-degree of a given vertex are equal. If all vertices have the same in(out)-degree, the graph is said to be *in(out)-regular*. If every possible arc exists, the graph is said to be *complete*.

A *path* on \mathcal{G} of length N from v_0 to v_N is an ordered set of distinct vertices $\{v_0, v_i, \dots, v_N\}$ such that $(v_{i-1}, v_i) \in \mathcal{A} \forall i \in [1, N]$. An *N -cycle* on \mathcal{G} is defined the same as a path except that $v_0 = v_N$, meaning the path rejoins itself. A graph without cycles is said to be *acyclic*. A graph with the property that the set of all cycle lengths has a common divisor k greater than one is said to be *k -periodic*.

Let $\mathcal{V}_1 \subset \mathcal{V}$ and $\mathcal{A}_1 \subset \mathcal{A}$, where each $a \in \mathcal{A}_1$ has its head and tail in \mathcal{V}_1 . The directed graph $\mathcal{G}_1 = (\mathcal{V}_1, \mathcal{A}_1)$ is termed a *subgraph* of \mathcal{G} . If \mathcal{A}_1 contains every arc in \mathcal{A} whose head and tail are in \mathcal{V}_1 , then \mathcal{G}_1 is termed an *induced subgraph* of \mathcal{G} .

4.2.2 Connectivity in Directed Graphs

If a path exists from v_i to v_j , it is said that v_i has *access* to v_j . A graph with the property that every vertex has access to every other vertex is said to be *strongly connected*. (A graph consisting of a single vertex with no arcs is also considered strongly connected.) A graph in which disjoint subsets of vertices exists whose elements do not have access to one another is termed *disconnected*. Note an undirected graph is either strongly connected or disconnected.

Two vertices which have access to one another are said to *communicate*. Communication is an equivalence relation, and the equivalence classes of \mathcal{V} induced by the communication relation are termed *components* of \mathcal{G} . Note that the induced subgraph of a component of \mathcal{G} is strongly connected. The component structure

also generates a second graph, where each equivalence class is represented by a vertex, and an arc joins two vertices if elements of the first equivalence class have access to the second. This graph is necessarily acyclic. In such a graph, a vertex which is not the head (tail) of any arc is said to be *initial* (*final*).

Example 4.1. Figure 4.1 shows a graph with twelve enumerated vertices. In this graph, vertices 1 through 6 communicate with one another, as do vertices 7, 8, 9, 11, and 12. The latter set has access to the former. Vertex 10 has access to all vertices, but no vertex has access to it. As such, this graph has three components. Figure 4.2 shows the induced subgraphs of these components. Figure 4.3 shows the acyclic graph induced by the component structure. In that graph, the leftmost vertex is final and the rightmost is initial.

4.3 Algebraic Graph Theory

4.3.1 Graph Laplacians

One area of graph theory which will be of significant interest to us is algebraic graph theory, which studies relationships between the structure of graphs and different matrix representations of graphs. For the purpose of defining graphs, we assume that the vertices of \mathcal{G} are enumerated, and each is denoted v_i . The *adjacency matrix* of a graph, denoted $A(\mathcal{G})$, is a square matrix of size $|\mathcal{V}|$, defined by $A_{ij} = 1$ if $(v_i, v_j) \in \mathcal{A}$, and is zero otherwise. When the graph in question is clear, the adjacency matrix will be denoted as A . Note that A uniquely specifies a graph, although A itself is not unique for a given graph, as it depends on the enumeration of the vertices. However, two adjacency matrices of the same graph are necessarily similar to one another via a permutation matrix. As such, it is clear that the eigenvalues of A are uniquely specified by the graph (though the converse is not true), and early research in algebraic graph theory focused on the relationship between eigenvalues of $A(\mathcal{G})$ and graph-theoretic properties of \mathcal{G} [31, 32].

Our work will make use of a different graph which has been the object of study more recently. Let D be the matrix with the out-degree of each vertex along the diagonal. The *Laplacian* of the graph is defined as¹

$$L = D^{-1}(D - A). \quad (4.1)$$

¹Some references define L as $D - A$. Others use the transpose of A to define the Laplacian of the directed graph. This distinction is of little consequence in terms of the theory, but the definition stated above better suits our purposes.




Figure 4.1: Sample Graph \mathcal{G} .




Figure 4.2: Induced Subgraphs of Components of \mathcal{G} .




Figure 4.3: Graph of Components of \mathcal{G} .

In the event that D is singular due to a vertex v_i with zero out-degree, set $D_{ii}^{-1} = 0$ to complete the definition. We will further denote the weighted adjacency matrix $I - L$ as G . For example, the Laplacian of the leftmost graph in Figure 4.2 is

$$L = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & -1 \\ -\frac{1}{4} & 1 & -\frac{1}{4} & -\frac{1}{4} & -\frac{1}{4} & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 \\ -\frac{1}{2} & 0 & 0 & 0 & -\frac{1}{2} & 1 \end{pmatrix}. \quad (4.2)$$

L can be viewed as a normalized version of the adjacency matrix. In this construction, each arc leading into a given vertex is weighted equally such that the weights sum to one. More generally, it is possible to work with weighted graphs, in which the off-diagonal elements of L are unequal yet still sum to -1 . Most of the results of the following sections do not depend on the arcs being weighted equally, though we will assume that for convenience.

Interest in the Laplacian matrix has increased in recent years; see [24, 74, 73, 78]. In particular, researchers identified many links between the eigenvalues of L and various graph-theoretic properties. We will return to this in Section 4.5, but before that we digress slightly into a discussion of nonnegative matrices.

4.4 Theory of Nonnegative Matrices

A property of the Laplacian is that G is nonnegative by construction. The theory of nonnegative matrices, much of which derives from the celebrated Perron-Frobenius theorem, will be quite useful in understanding the links between graph theory and vehicle formation control. The results of this section can be found in [2, 51, 98].

4.4.1 Nonnegative Matrices and Graph Theory

A matrix A is *positive* (*nonnegative*) if each element is positive (nonnegative). Given two nonnegative $n \times n$ matrices A, B , we say $A > B$ ($A \geq B$) if $A - B$ is positive (nonnegative). A square nonnegative matrix A is *reducible* if there exists a permutation matrix P such that PAP^T can be represented

$$PAP^T = \begin{pmatrix} A_{11} & 0 \\ A_{21} & A_{22} \end{pmatrix}, \quad (4.3)$$

where A_{11}, A_{22} are square, or if $n=1$ and $A = 0$. A matrix which is not reducible is said to be *irreducible*.

Just as in Section 4.3 a directed graph was used to define a matrix, a matrix can also be used to define a directed graph. This graph, denoted $\mathcal{G}(A)$, has n vertices denoted v_1, \dots, v_n , and $(v_i, v_j) \in \mathcal{A}(A)$ if $A_{ij} \neq 0$. Unlike the directed graphs discussed above, $\mathcal{G}(A)$ may contain loops if $A_{ii} \neq 0$. It should be clear that the directed graph associated with $G(\mathcal{G})$ is in fact \mathcal{G} itself.

The following theorem relates a nonnegative matrix with its directed graph, as well as supplying an algebraic characterization:

Theorem 4.1. *Given a nonnegative $n \times n$ matrix A , the following are equivalent:*

1. A is irreducible.
2. A^T is irreducible.
3. $\mathcal{G}(A)$ is strongly connected.
4. $(I_n + A)^{n-1} > 0$.

If $\mathcal{G}(A)$ is aperiodic, then A is termed *primitive*. If $\mathcal{G}(A)$ is k -periodic, it is termed imprimitive, or *cyclic of index k* . Primitivity admits the following algebraic characterization:

Theorem 4.2. *A is primitive if and only if there exists m such that $A^m > 0$.*

Often, this is taken as the definition of primitivity, and the relationship to graph theory can be deduced from it. Note that all the results discussed thus far concern only the nonzero entries of A , but not their actual value. Thus, these results can be used to relate a graph \mathcal{G} to its weighted adjacency matrix G .

4.4.2 Perron-Frobenius Theorem for Irreducible Matrices

Thus far, we have addressed the structure of nonnegative matrices and its relationship to related directed graphs. We now turn to the spectral structure of nonnegative matrices. The following celebrated theorem was proven for positive matrices by Perron and extended to irreducible matrices by Frobenius. We will denote the spectral radius of a matrix A as $\rho(A)$.

Theorem 4.3 (Perron-Frobenius). *Let A be a nonnegative, irreducible matrix. The following are true:*

1. $\rho(A) > 0$.
2. $\rho(A)$ is a simple eigenvalue of A , and any eigenvalue of A of the same modulus is also simple.
3. A has a positive eigenvector x corresponding to $\rho(A)$.
4. $B > A \Rightarrow \rho(B) > \rho(A)$.

Furthermore, if A is primitive, then all eigenvalues of A other than $\rho(A)$ have modulus strictly less than $\rho(A)$.

The Perron-Frobenius theorem states that the spectral radius of a nonnegative matrix is in fact an eigenvalue, often known as the Perron root. Associated with the Perron root is a positive eigenvector, known as the Perron vector. When A is irreducible, it will have positive left and right Perron vectors. When A is reducible, a more complicated picture emerges, as we shall see shortly.

If A is not primitive, the eigenvalues of A have an interesting structure:

Theorem 4.4. *Let A be a nonnegative, irreducible matrix which is cyclic of index k . Then A has k eigenvalues of modulus $\rho(A)$, equal to*

$$\lambda_i = \rho(A)e^{\frac{2\pi j}{k}i}, \quad i = 0, \dots, k-1. \quad (4.4)$$

4.4.3 Reducible Matrices

When a matrix A is reducible, the graph structure of the components of $\mathcal{G}(A)$, as shown in the example in Figure 4.3, will prove useful in extending the ideas of Perron-Frobenius theory. It is possible to enumerate the components such that no component has access to any component behind it in the ordering. When this is done (and the vertices within each component are ordered arbitrarily), then A has the following block structure:

$$A = \begin{pmatrix} A_{1,1} & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ A_{m-1,1} & \dots & A_{m-1,m-1} & 0 \\ A_{m,1} & \dots & A_{m,m-1} & A_{m,m} \end{pmatrix}, \quad (4.5)$$

where $A_{i,i}$ are the principal submatrices corresponding to each component. The component corresponding to $i = 1$ is final, and the component corresponding to $i = m$ is initial, though they need not be the only initial and final components. Each $A_{i,i}$ is irreducible or else a 1×1 zero matrix. As such, each irreducible $A_{i,i}$ has a Perron root. Clearly, the eigenvalues of A are the union of eigenvalues of $A_{i,i}$. A component is termed *basic* if its Perron root equals $\rho(A)$.

Rothblum [87] analyzed the eigenstructure of an arbitrary reducible matrix, using the location of the basic components within the graph generated by the component structure. In the next section, we specialize those results to the Laplacian.

4.5 Eigenvalues of Laplacians

4.5.1 Elementary Results

We now return to the structure of the spectrum of the Laplacian. We begin with the observation that the rows of L sum to zero by definition, which implies that

Proposition 4.1. *Zero is an eigenvalue of L .*

Furthermore, this condition implies that 1^T is the eigenvector associated with this eigenvalue. Of course, any eigenvalue λ of L corresponds to an eigenvalue $1 - \lambda$ of G . Thus, G has an eigenvalue of 1. The fact that this eigenvalue has a positive eigenvector implies that it is, in fact, the Perron root of G . We therefore conclude from the Perron-Frobenius theorem that

Proposition 4.2. *All eigenvalues of L lie in a disk of radius 1 centered at the point $1 + 0j$ in the complex plane.*

We denote this region the *Perron disk*. We can apply further ideas from Perron-Frobenius theory:

Proposition 4.3. *If \mathcal{G} is strongly connected, the zero eigenvalue of L is simple. If, in addition, \mathcal{G} is aperiodic, all nonzero eigenvalues lie in the interior of the Perron disk. If \mathcal{G} is k -periodic, L has k evenly spaced eigenvalues on the boundary of the Perron disk.*

If \mathcal{G} is undirected, then L is similar to $I - D^{1/2}AD^{1/2}$, which is clearly symmetric, from which we conclude that

Proposition 4.4. *If \mathcal{G} is undirected, then all eigenvalues of L are real.*

When \mathcal{G} is undirected, the graph is either strongly connected or disconnected. Since we will be concerned with Laplacians of directed graphs, we wish to consider the eigenvalue/eigenvector structure of L for directed graphs. To do so, we specialize the results of [87], and derive the following:²

Proposition 4.5. *The multiplicity m of the zero eigenvalue of L is equal to the number of final components of \mathcal{G} . The dimension of the kernel of L also equal to m , and is spanned by a basis of m nonnegative vectors.*

Proof. We begin by proving the first assertion. Suppose \mathcal{G} has m components, and let G be partitioned as in Equation (4.5). The final components of \mathcal{G} have the property that if component j is final, $G_{j,i} = 0$ for all $i < j$. Thus, the rows of $G_{j,j}$ sum to one, implying that each final component is basic. We now consider the non-final components. Let $\tilde{\mathcal{G}}_j$ be the induced subgraph of the j th component, presumed non-final, and let \tilde{G}_j its weighted adjacency matrix. Because $\tilde{\mathcal{G}}_j$ is non-final, it must contain at least one vertex whose out-degree is less than that of the corresponding vertex in \mathcal{G} . Obviously, no vertex can have a higher out-degree than its counterpart. We see therefore that $G_{j,j} < \tilde{G}_j$, which implies by Theorem 4.3 that the Perron root of $G_{j,j}$ is less than 1. Thus, only the final components of \mathcal{G} are basic, meaning the multiplicity of 1 as an eigenvalue is equal to the number of basic components.

We now construct m linearly independent vectors which lie in the kernel of L . Suppose there are r non-final components. If $r = 0$, the graph consists of m disconnected components, and the result is trivial. If $r \neq 0$, let $g_{i,j}$ denote the vector populated by the row sum of $G_{i,j}$, and let $g_k = [g_{m+1,k}^T, \dots, g_{m+r,k}^T]^T$. Let G' be the lower right-hand quadrant of G when partitioned into final and non-final components. Let $v_k = [0_1, \dots, 1_k, \dots, 0_m, w_k^T]^T$, where the first m terms are sized compatibly with their components. and w is sized compatibly with the r non-final components. We see that

$$Gv_k = [0_1, \dots, 1_k, \dots, 0_m, g_k^T]^T + [0_1, \dots, 0_k, \dots, 0_m, (G'w_k)^T]^T. \quad (4.6)$$

Recall that $\rho(G') < 1$, as discussed above, so $I - G'$ is invertible. In order for v_k

²Rothblum uses the term “class” to refer to the subset of rows whose corresponding vertices induce a component of the associate graph. For simplicity, we will use component to refer both to the subgraph and the rows of the matrix which correspond to it.

to be an eigenvector of G , we must choose w_k such that

$$g_k + G'w_k = w_k \tag{4.7}$$

To do so, choose $w_k = (I - G')^{-1}g_k$. Clearly the v_k are linearly independent of one another. We thus construct m linearly independent eigenvectors corresponding to the 1 eigenvector of G , and hence the kernel of L . Of course, there cannot be more than m linearly independent eigenvectors in the kernel of L , since the multiplicity of the zero eigenvalue is m .

Finally, we prove that w_k is nonnegative. Each g_k is nonnegative by construction. $I - G'$ has the property that all off-diagonal elements are non-positive and all eigenvalues have positive real part. Therefore ([41], Theorem 4.3), $I - G'$ is known as a K-matrix, and its inverse (known as an M-matrix) is nonnegative. Thus, w_k , and hence v_k , are nonnegative. ■

Of course, $\sum_{k=1}^m v_k = 1^T$, since this is known to be in the kernel of L . We see that when multiple final components exist, the positive Perron eigenvector is “split up” among m vectors, each of which can be associated with a final component.

4.5.2 Additional Results

If \mathcal{G} consists of two disconnected components, it is clear that the multiplicity of the zero eigenvalue of L is two. If, instead, the two components are connected by a small number of arcs, it follows from a perturbation argument that L will have an eigenvalue near zero. For this reason, Fiedler [39] termed this eigenvalue the *algebraic connectivity* of the graph, and began a program of research which continues to this day relating this parameter to graph-theoretic concepts, including measures of connectivity in graphs. Fiedler [40] also showed the associated eigenvector, often termed the Fiedler vector, to be useful in partitioning graphs into densely connected subgraphs. While his paper applies to undirected graphs, the heuristic technique he developed can be applied to directed graphs as well; see [44] for one such example.

When \mathcal{G} is undirected, the eigenvalues of the Laplacian can be viewed as solutions to a minimization or maximization problem following the Rayleigh-Ritz theorem [51]. This characterization of the eigenvalues has led to many bounds on the eigenvalues of L and links to other topics in graph theory; see [24] for an excellent exposition of this topic. For example, finding optimal cuts in graphs has

led to the definition of *Cheeger constants* for graphs in an analogous fashion to their definition for Riemannian manifolds, which in turn bound the eigenvalues of L ; again, see [24], and [3, 43, 81] for more recent results. Again, though these results are for undirected graphs, they can sometimes be extended to bound the real component of eigenvalues of directed graphs.

Chapter 5

Stabilization of Vehicle Formations

In this chapter, we use the ideas from graph theory developed in the previous chapter to address stability of vehicle formations. We identify the Laplacian as the graph-theoretic object of interest in the context of stabilization of relative position in a vehicle formation. This leads to formation stability criteria involving only the local plant and control law and the eigenvalues of the Laplacian. The negative inverse of the Laplacian eigenvalues play the role of the -1 point in the Nyquist stability criterion. We use graph-theoretic characterizations of Laplacian eigenvalue locations to evaluate formation interconnection strategies from the perspective of formation stability. Our approach in this section is similar to the approach in [48, 84], but from a control-theoretic perspective.

The chapter is structured as follows. In Section 5.1, we present the formation equations of motion and identify the role of the Laplacian. In Section 5.2, we present several stability theorems. In Section 5.3, we discuss the implication of different structures in the graph for formation stability.

5.1 Relative Position Control in Vehicle Formations

Vehicle formations can be assigned a wide range of tasks; in this chapter we consider one such task to gain insight into the effect of network topologies on formation dynamics. The problem we consider in this chapter is the stabilization of the relative position of a set of vehicles with identical linear dynamics. This problem arises in the context of automated highway systems [106] and satellite formations [108]. At the end of the chapter we will comment on mixed absolute/relative position measurements, as might occur in scenarios such as vehicle pursuit and obstacle avoidance.

5.1.1 Formation Equations of Motion

We consider a set of N vehicles, whose (identical) linear dynamics are denoted

$$\dot{x}_i = P_A x_i + P_B u_i, \quad (5.1)$$

where $i \in [1, N]$ is the index for the vehicles in the flock. Note that each vehicle's dynamics are decoupled from the vehicles around it. Each vehicle's sensed information is defined as follows:

$$y_i = P_{C_1} x_i \quad (5.2)$$

$$z_{ij} = P_{C_2}(x_i - x_j), \quad j \in \mathcal{J}_i, \quad (5.3)$$

where the set $\mathcal{J}_i \subset [1, N] \setminus \{i\}$ represents the set of vehicles which vehicle i can sense. Thus, y_i represents internal state measurements, and z_{ij} represents external state measurements relative to other vehicles. To obtain relative state measurements, a vehicle must have access, in this case via some form of sensing, to other vehicles' states. We assume that $\mathcal{J}_i \neq \emptyset$, meaning each vehicle can see at least one other vehicle. Note that a single vehicle cannot drive all the z_{ij} terms to zero simultaneously; the errors must be synthesized into a single signal. For simplicity, we will assume that all relative state measurements are weighted equally to form one error measurement:

$$z_i = \frac{1}{|\mathcal{J}_i|} \sum_{j \in \mathcal{J}_i} z_{ij}. \quad (5.4)$$

Of course, we could also weight different measurements unequally if desired. We also define a decentralized control law $K(s)$ which maps y_i, z_i to u_i , represented in state-space form by

$$\begin{aligned} \dot{v}_i &= K_A v_i + K_{B_1} y_i + K_{B_2} z_i \\ u_i &= K_C v_i + K_{D_1} y_i + K_{D_2} z_i. \end{aligned} \quad (5.5)$$

We now consider the system of all N vehicles together. We use the hat notation, for example \widehat{A} , to represent the matrix A repeated N times along the diagonal, or $\widehat{A} = I_N \otimes A$, where \otimes represents the Kronecker product. Using this notation, the total system dynamics is represented as follows:

$$\begin{aligned} \dot{x} &= \widehat{P}_A x + \widehat{P}_B \widehat{K}_{D_1} \widehat{P}_{C_1} x + \widehat{P}_B \widehat{K}_{D_2} \widehat{P}_{C_2} L_{(n)} x + \widehat{P}_B \widehat{K}_C v \\ \dot{v} &= \widehat{K}_A v + \widehat{K}_{B_1} \widehat{P}_{C_1} x + \widehat{K}_{B_2} \widehat{P}_{C_2} L_{(n)} x \end{aligned} \quad (5.6)$$

or

$$\begin{pmatrix} \dot{x} \\ \dot{v} \end{pmatrix} = \begin{pmatrix} \hat{P}_A + \hat{P}_B \hat{K}_{D_1} \hat{P}_{C_1} + \hat{P}_B \hat{K}_{D_2} \hat{P}_{C_2} L_{(n)} & \hat{P}_B \hat{K}_C \\ \hat{K}_{B_1} \hat{P}_{C_1} + \hat{K}_{B_2} \hat{P}_{C_2} L_{(n)} & \hat{K}_A \end{pmatrix} \begin{pmatrix} x \\ v \end{pmatrix}. \quad (5.7)$$

The resulting system is block diagonal with the exception of the matrix $L_{(n)}$, which contains the relative sensing information. We now turn our attention to this matrix. Define L in the following way:

$$L_{ii} = 1 \quad (5.8)$$

$$L_{ij} = \begin{cases} -\frac{1}{|J_i|}, & j \in J_i \\ 0, & j \notin J_i. \end{cases} \quad (5.9)$$

Of course, L is the Laplacian of a graph to be defined in the next section. Letting n be the dimension of x_i , $L_{(n)}$ is of dimension $Nn \times Nn$ and is defined by replacing each element of L with that element multiplied by I_n (i.e., $L_{(n)} = L \otimes I_n$), thus generating a version of L dimensionally compatible with x_i .

5.1.2 Commanded Vehicle Offsets

The goal of the controller as defined above is to drive the states (or at least a subset of them) to a common value. In this problem definition, we are not concerned about the final value so long as the vehicles share it. For some applications, such as orienting underwater vehicles, this is an understandable goal. For other applications, such as relative satellite positioning, it is necessary to add an offset term to z_{ij} to achieve the desired intervehicle spacing. We define a time-varying offset function $h : [1, N] \times [1, N] \times \mathbb{R} \rightarrow \mathbb{R}^m, i, j, t \mapsto h_{ij}(t)$, where m is the dimension of z_{ij} , which defines the intervehicle spacing. We assume that h is defined so that for all i, j, k , $h_{ij} + h_{jk} = h_{ik}$. This definition means that it is possible to position each vehicle such that $z_{ij} = h_{ij}$ for all i, j . One way to generate such a function is to define an offset $h_{i0}(t)$ for each vehicle relative to an arbitrary reference. Letting $h_0(t)$ be the vector of $h_{i0}(t)$ offsets, the error signal is then defined by

$$z(t) = L_{(m)}(y(t) - h_0(t)). \quad (5.10)$$

The offset function acts as an input to the dynamical system. We assume that $h_{ij}(t)$ is bounded. Because BIBO stability is implied by internal stability for LTI systems, the actual value of the h_{ij} terms does not play a role in the stability

analysis, and will henceforth be omitted. In general, h_{ij} will be chosen to be consistent with the open-loop dynamics of the vehicles in formation. We will not consider the case where h is also a function of the measurements, which is the case for variable spacing policies discussed in [96, 106].

5.1.3 Graph Laplacians and Formation Dynamics

We are now able to identify the role of the sensing graph in the formation dynamics. The vehicles and their sensing indices \mathcal{J}_i together form a graph, where each node represents a vehicle and an arc leads from node j to node i if $j \in \mathcal{J}_i$. Our assumption that each vehicle can sense at least one other vehicle implies that the out-degree of each vertex is at least 1. The matrix L defined in Equation (5.8) is none other than the Laplacian of the graph, defined in Section 4.3. The normalization of the Laplacian is equivalent in our setting to our averaging of the z_{ij} terms so as not to add a gain term. This stands in distinction to other examples in the literature which use the Laplacian-like matrices in analyzing stability of interconnected systems [84].

The classification of graphs can now be translated into more familiar terms. A final component corresponds to one which is sensed by others but does not sense others, often known as a “leader.” Similarly, initial components correspond to “followers” in the formation. Note, however, that the leader-follower distinction is not a function of one’s position relative to other vehicles, but a function of one’s position within the topology of the sensed information flow. (The two, of course, may coincide.) When a graph is irreducible, no leaders or followers are distinguishable within the formation. As we discussed earlier, leader-follower formations are simpler from a stability perspective in that the motion of the vehicle “ahead” can be treated as a disturbance, whereas strongly connected formations must somehow take into account the global picture. In the next section, we will make all these ideas precise using the Laplacian.

5.2 Stabilization of Vehicle Formations

5.2.1 The Role of The Laplacian in Formation Stability

We now consider the relationship between graph Laplacians and formation stability. We show the following to be true:

Theorem 5.1. *A local controller $K(s)$ stabilizes the formation dynamics in Equation (5.7) iff it simultaneously stabilizes the set of N systems*

$$\begin{aligned}\dot{x} &= P_A x + P_B u \\ y &= P_{C_1} x \\ z &= \lambda_i P_{C_2} x\end{aligned}\tag{5.11}$$

where λ_i are the eigenvalues of L .

Proof. We will show the above to be true by transforming the closed-loop dynamics in the following way: Let T be a Schur transformation of L , meaning the unitary matrix such that $U = T^{-1}LT$ is upper triangular with the eigenvalues of L along the diagonal [51]. Clearly, $T_{(n)}$ is a Schur transformation of $L_{(n)}$. This transformation has the following useful property, a clear consequence of the block structure of the relevant matrices:

Lemma 5.2. *Let X be an $r \times s$ matrix, and Y be an $N \times N$ matrix. Then*

$$\hat{X}Y_{(s)} = Y_{(r)}\hat{X}.\tag{5.12}$$

Proof. Using Kronecker product algebra, both sides can be shown to be equal to $Y \otimes X$. ■

Applying this property to the system dynamics, we see if we let $\tilde{x} = T_{(n)}x$, and $\tilde{v} = T_{(m)}v$, we can rewrite Equation (5.7) as

$$\begin{pmatrix} \dot{\tilde{x}} \\ \dot{\tilde{v}} \end{pmatrix} = \begin{pmatrix} \hat{P}_A + \hat{P}_B \hat{K}_{D_1} \hat{P}_{C_1} + \hat{P}_B \hat{K}_{D_2} \hat{P}_{C_2} U_{(n)} & \hat{P}_B \hat{K}_C \\ \hat{K}_{B_1} \hat{P}_{C_1} + \hat{K}_{B_2} \hat{P}_{C_2} U_{(n)} & \hat{K}_A \end{pmatrix} \begin{pmatrix} \tilde{x} \\ \tilde{v} \end{pmatrix}.\tag{5.13}$$

The elements of the transformed system matrix are either diagonal or upper triangular. This means that stability of this system is equivalent to the stability of the systems along the diagonal, i.e.:

$$\begin{aligned}\dot{\tilde{x}}_i &= (P_A + P_B K_{D_1} P_{C_1} + \lambda_i P_B K_{D_2} P_{C_2}) \tilde{x}_i + P_B K_C \tilde{v}_i \\ \dot{\tilde{v}}_i &= (K_{B_1} P_{C_1} + \lambda_i K_{B_2} P_{C_2}) \tilde{x}_i + K_A \tilde{v}_i\end{aligned}\tag{5.14}$$

which is equivalent to the controller $K(s)$ stabilizing the system

$$\begin{aligned}\dot{x} &= P_A x + P_B u \\ y &= P_{C_1} x \\ z &= \lambda_i P_{C_2} x.\end{aligned}\tag{5.15}$$

■

We thus identify the Laplacian eigenvalues as the object of interest in understanding formation stability. In this context, the zero eigenvalue of L can be interpreted as the unobservability of absolute motion of the formation in the measurements z_i . It seems that a prudent controller design strategy is to close an inner loop around y_i such that the result system is stable, and then to close an outer loop around z_i which achieves desired formation performance. For the remainder of this paper, we concern ourselves solely with the outer loop. Hence, we assume from now on that P_{C_1} is empty and that P_A has no eigenvalue in the open RHP. We do not wish to exclude eigenvalues along the $j\omega$ axis because many vehicle formations (e.g., vehicle platoons, satellite clusters) possess those, and the presence of unobservable secular or periodic motion of the formation may be tolerable in those cases. If $K(s)$ stabilizes the system in Equation (5.15) for all λ_i other than the zero eigenvalue, we will say that it stabilizes the *relative formation dynamics*.

5.2.2 Location of Equilibrium Points

In addition to verifying stability, we wish to verify that the (relative) equilibrium point about which it is stable is consistent with the goals of the formation. The goal of the formation is to have z_{ij} be equal to zero for all i, j , not just for $j \in \mathcal{J}_i$, meaning each vehicle is positioned properly relative to every other vehicle. However, only the measurement z_i is sent to the controller. We must verify conditions under which $z_i = 0$ for all i implies $z_{ij} = 0$ for all i, j .

As seen in Equation (5.10), z will be zero if $y - h_0$ lies in the kernel of L . The trivial solution is, of course, $y = h_0$, which satisfies the formation goals by construction of h_0 . The remainder depends on the kernel of L . If the graph is strongly connected or has a single leader component, the kernel of L is spanned by 1^T . In this case, possible equilibria take the form $y = h_0 + k1^T$. This entire subspace is consistent with the formation goals, since it still satisfies $y_i - y_j = h_{i0} - h_{j0} = h_{ij}$.

If multiple leader components exist, then the kernel of L will include a non-constant eigenvector, and the formation will likely stabilize about a point which is not consistent with the formation goals. This is because of the unobservability of each leader component by the other, and the inability of the follower components to distinguish between the two. We will therefore assume that our formation has no more than one leader component, implying that G is irreducible and that the Perron eigenvalue is simple. We will see again in the next chapter that formations with multiple leader components exhibit, unsurprisingly, very poor performance.

5.2.3 Stability Theorems

We now return to the topic of formation stabilization. In general, proving simultaneous stabilization results can be difficult. This set of systems is special, in that they differ only by a complex scalar. For SISO systems, we can state a second version of Theorem 5.1 which is useful for stability and robustness analysis:

Theorem 5.3. *Suppose $P(s) = P_{C_2}(sI - P_A)^{-1}P_B$ is SISO. Then $K(s)$ stabilizes the relative formation dynamics iff the net encirclement of $-\lambda_i^{-1}$ by the Nyquist plot of $K(s)P(s)$ is zero for all nonzero λ_i .*

Proof. The Nyquist criterion states that stability of the closed loop system in Theorem 5.1 is equivalent to the number of CCW encirclements of $-1 + j0$ by the forward loop $\lambda_i P(j\omega)K(j\omega)$ being equal to the number of RHP poles of $P(s)$, which is assumed to be zero. This criterion is equivalent to the number of encirclements of $-\lambda_i^{-1}$ by $P(j\omega)K(j\omega)$ being zero. ■

In the case where $P(s)$ is MIMO, the formation can be thought of as a structured uncertainty of the type scalar time identity (see [111]) where the scalars are the Laplacian eigenvalues. More specifically, we shall write the eigenvalues as $\lambda_i = 1 + \mu_i$ and consider bounds on μ_i . Suppose it is known that $|\mu_i| \leq M$ for all nonzero λ_i . If we close the loop around the unity block and leave $\mu_i I$ as an uncertainty, the resulting lower block is $C(s) = P(s)K(s)(I + P(s)K(s))^{-1}$, which is assumed to be stable. The following result from robust control theory then applies:

Theorem 5.4. *$K(s)$ stabilizes the relative formation dynamics of the MIMO vehicle $P(s)$ if*

$$\rho(C(j\omega)) < M^{-1} \quad \forall \omega \in (-\infty, \infty) \quad (5.16)$$

This formulation enables the control designer to use tools from robust control theory to design a controller, assuming the perturbation of the Laplacian eigenvalues can be bounded. Not considered here is the case where the graph itself is a function of vehicle position, and therefore must be included in the representation of the formation dynamics. Such an analysis would require tools from hybrid systems analysis, and is outside the scope of this thesis.

5.2.4 Mixed Absolute/Relative Sensing

The results we developed thus far have focused on a formation where only relative position measurements exist. We now briefly address the case where mixed absolute and relative measurements exist. This would apply to a formation in which some vehicles could sense the target but needed to approach it in formation, or a case where a formation needed to avoid an obstacle which some of the vehicles could sense.

Insofar as the target is a source of sensing information, it can be included in the graph. Of course, it is a final component. The dynamics of the target are not relevant to formation stability (assuming they are bounded), so our stability analysis concerns only the submatrix of L containing the vehicles. In this case, the error signal of vehicle i takes the form (assuming equal weighting of all measurements),

$$z_i = \frac{1}{|\mathcal{J}_i| + 1} \left(z_{i0} + \sum_{j \in \mathcal{J}_i} z_{ij} \right). \quad (5.17)$$

In this case, the relevant submatrix of L will still have 1 along the diagonal, but the off-diagonal terms will sum to $-\frac{|\mathcal{J}_i|}{|\mathcal{J}_i| + 1}$. Assuming all vehicles have access to the target, then $\rho(G) < 1$ (Theorem 4.3), so L does not have a zero eigenvalue. This corresponds to the fact that unobservability of bulk formation motion is no longer relevant when the formation can view the target. From this perspective, the presence of absolute measurements tightens the bound on $\lambda(L)$, which improves the stability picture. However, it does not guarantee that individual eigenvalues will not be nearer to the bound than they were in the case where only relative measurements were present.

5.2.5 Example: Double Integrator with Time Delay

Let $P(s) = \frac{e^{-sT}}{s^2}$ and $K(s) = K_p + K_d s$. This corresponds to a double integrator with a time delay being controlled by a PD controller. This is equivalent to the

hovercraft of Section 2.3.2 where the force offset is zero and rotational dynamics are ignored. Figure 5.1 shows a formation graph, and Figure 5.2 the Nyquist plot of $K(s)P(s)$ with the Laplacian eigenvalues. The black ‘o’ locations in Figure 5.2 correspond to the eigenvalues of the graph defined by the black arcs in Figure 5.1, and the ‘x’ locations are for eigenvalues of the graph when the dashed arc is included as well. This example clearly shows the effect the formation has on stability margins. The standard Nyquist plot reveals a system with reasonable stability margins — about 8 dB and 45 degrees. When one accounts for the effects of the formation, however, one sees that for the ‘o’ formation, the stability margins are substantially degraded, and for the ‘x’ formation, the system is in fact unstable. Interestingly, the formation is rendered unstable when additional information (its position relative to vehicle 6) is used by vehicle 1. We shall return to this point shortly.

5.3 Evaluating Formations via Laplacian Eigenvalues

The location of Laplacian eigenvalues has emerged as the parameter which enables formation stability to be analyzed on the local level. We now turn to the question of bounding or predicting eigenvalue location based on properties of the graph. We begin by considering simple formation structures and their eigenvalue placement. Examples of these graphs are shown in Table 5.1, where sample graphs, their nonzero spectra, and the locations on the Nyquist plot are shown.

1. Complete graph. The complete graph is one where every possible arc exists. In this case, the eigenvalues of a graph with N vertices can be analytically determined to be zero and $1 + \frac{1}{N-1}$, the latter repeated $N - 1$ times. For large N , stabilization of the complete graph is equivalent to stabilizing an individual vehicle. Of course, a complete interconnection structure can place an enormous burden on each vehicle’s sensing and computational capacities.
2. Acyclic (directed) graph. This graph has the 1 eigenvalue repeated N times. This can be seen from the fact that the vertices can be ordered such that L is upper triangular with ones along the diagonal. This is the “leader-follower” architecture discussed earlier. In this case, vehicle stabilization is truly a local result, since other vehicles’ dynamics enter only as a disturbance. However, this architecture has drawbacks regarding disturbance rejection.




Figure 5.1: Formation Graph.




Figure 5.2: Formation Nyquist Plot.

3. Single directed cycle. This graph is periodic, and therefore has eigenvalues at $1 - e^{j(i-1)/2\pi}$, $i \in [1, N]$ according to Proposition 4.3. These eigenvalues lie on the boundary of the Perron disk in which all the eigenvalues must lie. Note that the negative inverse of these points lie on the -0.5 vertical in the complex plane.
4. Two-cyclic undirected graph. A graph of this type would include a vehicle platoon with bidirectional position measurement. This graph will have an eigenvalue at 2, due to its periodicity, and all other eigenvalues will be real, due to the symmetry of the graph.

Figure 5.3 shows various eigenvalue regions for $-L$, and Figure 5.4 shows the corresponding regions for $-L^{-1}$. The region bounded by the solid line is the Perron disk in which all eigenvalues must lie. Its inverse is the LHP shifted by -0.5. The dashed region is a bound in the magnitude of the nonzero eigenvalues of L . It corresponds to a shifted circle on the right-hand side of Figure 5.4. Finally, the dash-dot line corresponds to a bound on the real component of the eigenvalues. The inverse of this bound corresponds to a circle which touches the origin. The shaded region represents the “desirable” region, in which the eigenvalues’ locations do not differ substantially from -1 .

As we discuss in Chapter 4, significant efforts have been made by graph theorists to relate eigenvalue location to various graph-theoretic properties. These results were largely derived for undirected graphs, whose eigenvalues are real and can be bounded using variational techniques. For directed graphs, the problem is more challenging. Nonetheless, we can make qualitative statements about the effect of graph structure on eigenvalue placement, and hence stability.

5.3.1 Graph Periodicity and Formation Stability

If we consider the complete graph and the single directed cycle graph of Table 5.1 as representing two extremes — one with all eigenvalues at a single location, the other with eigenvalues maximally dispersed, we see that eigenvalue placement can be related to the rate of mixing of information through the network. When the graph is highly connected, the global component of an individual vehicle’s dynamics are rapidly averaged out through the rest of the graph, and thus has only a minor effect on stability. When the graph is periodic, the global component of the dynamics introduces periodic forcing of the vehicle, and the rest of the network never averages it out. This is represented on the Nyquist plot by putting

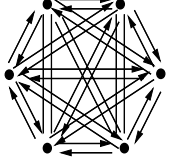
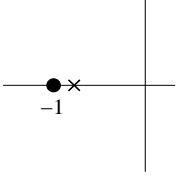
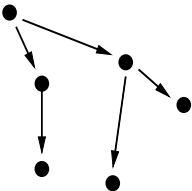
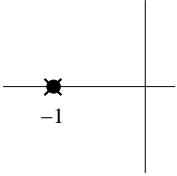
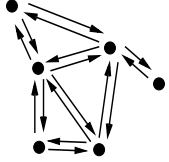
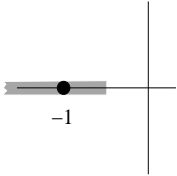
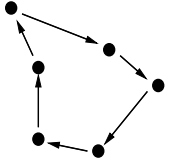
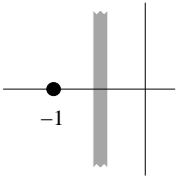
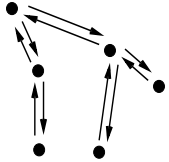
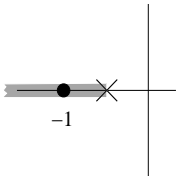
	$\lambda(L) = \{1 + \frac{1}{N-1}\}$	
	$\lambda(L) = \{1\}$	
	$\lambda(L) \subset [0, 2]$	
	$\lambda_i(L) = 1 - e^{\frac{2\pi j}{N}i}$	
	$\lambda(L) \ni 2$	

Table 5.1: Sample Graphs, Spectra, and Nyquist Locations.

the inverse eigenvalues nearer to the imaginary axis, thus diminishing stability margins.

We see that aperiodicity is a desirable property of formation interconnection topologies. With this insight, we can see why the system in Figure 5.1 loses stability margin when a link is added. The “solid” graph possesses two 3-cycles and two 2-cycles. When the dashed link is added, an additional 3-cycle is created, rendering the graph more nearly 3-periodic. This drives two of the eigenvalues nearer to the positions they would occupy if the graph were truly periodic, i.e., the -0.5 vertical. An interesting, and to this author’s knowledge, open challenge is to quantify this insight: to define a measure of periodicity within a graph which can be correlated to bounds on the eigenvalue locations.

5.3.2 Stability of Sparsely Connected Formations

Another aspect of eigenvalue placement which deserves mention is that of weakly connected graphs. As we discussed in Section 4.5, if we consider two disconnected subgraphs, then clearly the Laplacian has two eigenvalues at zero. If we weakly connect the graphs, one of the eigenvalues moves slightly away from zero. From a controls perspective, these near-zero eigenvalues represent weakly observable modes. On the formation Nyquist plot, they will be far from the origin, but they have the potential to destabilize a graph if its Nyquist plot possesses an upper gain margin. Generally speaking, they will not impact formation stability. However, one cannot deduce stability of the formation by examining separately the eigenvalues of the subgraphs: connecting the subgraphs perturbs all the eigenvalues, not only the zero eigenvalue. In [23], bounds on the perturbation of these eigenvalues are derived as a function of graph density parameters derived therein.

Most of the work cited earlier regarding eigenvalue bounding focuses on the algebraic connectivity and its relationship to substructures in the graph. As we have seen, the algebraic connectivity does not directly impact stability margins. Of greater interest to this application are bounds such as λ in Figure 5.4, which identify the region in which the majority of eigenvalues are located, with the possible exception of isolated eigenvalues near the origin. Examining the structure of the cycles in the graph may lead to results in this area.




Figure 5.3: Inclusion Regions for $-\lambda(L)$.




Figure 5.4: Inclusion Regions for $-\lambda(L)^{-1}$.

5.3.3 Weighted Graphs

It is important to note that these results do not depend on the equal weighting of the sensed signals used to define the Laplacian in Equation (5.8). The concepts of connectivity and periodicity explored in this (and the following) chapter depend only on the nonzero elements of the Laplacian, not on their actual values. As such, these results are equally valid for the case of weighted Laplacians, where the off-diagonal elements of the Laplacian are not equal. However, we will maintain the restriction that the off-diagonal elements sum to -1 . This restriction will maintain the Perron root of G at 1, and it implies that the averaging of sensed signals performed by each vehicle does not add gain in the control loop. Maintaining this constraint will be even more important in the next chapter.

This observation leads one to consider the possibility of a weighting algorithm, by which each vehicle varies the weighting of each input it sees in order to augment stability or improve some performance criterion. This topic is an interesting avenue of research, but lies outside the scope of this thesis: in the absence of a weighting algorithm, we have restricted our attention to the simplest case of equal weighting. As we have shown by example, naive decisions by an individual vehicle can have negative consequences for the formation. To be effective, any weighting algorithm will require some information from other vehicles. The identification of such parameters to be communicated between vehicles is a natural first step in such a research effort.

Chapter 6

Information Flow in Vehicle Formations

6.1 Introduction: A Motivating Example

In the previous chapter, we discussed how sparseness in the sensing graph can lead to poor stability margins. In this chapter, we begin our discussion of information flow design by examining the effect of sparseness on performance. Returning to the example in Section 5.2.5, suppose six vehicles whose dynamics are double integrators in the plane are asked to take up positions on the points of a regular hexagon relative to one another. (In this case, the offset $h_{i0} = [\cos \frac{\pi i}{3}, \sin \frac{\pi i}{3}]$). Figure 6.1 shows the trajectories followed by the six vehicles as they approach their target positions. The initial positions are marked with an ‘o,’ and the final positions are marked with an ‘x.’ While the formation is stable and achieves its desired position (as verified in the previous chapter), the trajectories followed by the vehicles are very circuitous and far from optimal. This is a clear consequence of the fact that individual vehicles do not have global knowledge of the behavior of the formation. Specifically, the vehicles do not have knowledge of the formation center, which would enable them to position themselves relative to the formation center rather than relative to the small number of vehicles visible to each.

In this chapter, we will explore strategies for flow of information which enable the vehicles to acquire the information they lack. In doing so, we must keep in mind several areas of concern. The first is the interaction of the information flow loops with the vehicle dynamics and the possible consequences for stability. The tools for formation stability analysis developed in the previous chapter will be useful in this regard. The second is the need of the information flow policy to be robust to changes in the graph which defines the flow of transmitted information, as well as lack of knowledge at the local level of global structures in the graph.




Figure 6.1: Hexagon Acquisition, No Information Flow.

Finally, we wish to pursue information flow policies that minimize the amount of information exchange yet enable robust formation control.

There is no limit to the complexity that information exchange policies could exhibit. Because our interest is the interaction of the information exchange with vehicle dynamics, we will eschew the approach to consensus among distributed systems found in references such as Lynch [69], and instead pursue a “bottom-up” approach, in which we begin with the simplest forms of information exchange which remain amenable to stability analysis. As we shall demonstrate, even simple information exchange, when properly designed, can yield significant results. Hopefully, the methodology used in this and the previous chapter contains elements which can be applied to analyzing the interaction of more complex information exchange protocols with dynamical systems.

6.2 An Information Flow Paradigm

6.2.1 Problem Setup

In the previous chapter, we assumed that sensed information was available instantaneously, and we used a continuous-time model of the vehicle dynamics. In this case, we will assume that information takes a fixed time T to travel between vehicles. To facilitate analysis, we also model our vehicle as a discrete time dynamical

system:

$$\begin{aligned} x_{k+1}^i &= P_A x_k^i + P_B u_k^i \\ y_k^i &= P_C x_k^i + P_D u_k^i \end{aligned} \quad (6.1)$$

where k is the time step of duration T and i is the vehicle index. As in the previous chapter, the error signal used by each vehicle is

$$z_k^i = \frac{1}{|\mathcal{J}_i^S|} \sum_{j \in \mathcal{J}_i^S} y_k^i - y_k^j, \quad (6.2)$$

where the index set \mathcal{J}_i^S represents the set of vehicles sensed by vehicle i , and we can form the sensing graph based on those sets. Note that the stability results of Chapter 5 can be reproduced for discrete time systems by plotting the response of the discrete-time transfer function for $z = e^{j\omega}$ and applying the Nyquist criterion.

Broadly speaking, any information flow consists of vehicles receiving a transmission from other vehicles and performing some computation using that information, information from previous transmissions, and sensed information. Each vehicle then transmits the results of their computation to other vehicles. We can view this process as a discrete-time dynamical system where the states are the information at each vehicle. Thus, a generic information flow law can be represented

$$p_{k+1}^i = f(p_k^i, p_{k-1}^i, \dots, \{p_k^j, p_{k-1}^j, \dots | j \in \mathcal{J}_i^T\}, z_{k+1}^i, z_k^i \dots), \quad (6.3)$$

where \mathcal{J}_i^T is the set which determines the transmitted information topology. We will assume that $\mathcal{J}_i^T = \mathcal{J}_i^S$ for all i , meaning the sensed information graph and transmitted information graph are identical. We will henceforth omit the superscript.

As discussed in the introduction, the graph which defined the information flow topology need not coincide with the physical realization of the communication network. Such considerations are outside the scope of this thesis. It should be emphasized, however, that conventional notions of network connectivity may not be applicable in this case. In our case, an arc is present in the graph if a transmission can reach vehicle i from vehicle j with in time Δt . Many network systems achieve high levels of connectivity with little guarantees on latency, and this is often a limiting factor in control over networks. For this reason, and for others touched upon in the introduction, we consider the communication topology as uncertain, regardless of the means by which communication is realized.

The information flow law we are going to investigate will mimic the structure

of the sensed information, taking the following form:

$$p_{k+1}^i = \sum_{j \in \mathcal{J}_i} p_k^j + (y_k^i - y_k^j) \quad (6.4)$$

or, in vector form:

$$p_{k+1} = G_{(m)} p_k + L_{(m)} y_k, \quad (6.5)$$

where $G_{(m)}$ and $L_{(m)}$ are the directed adjacency matrix and Laplacian of the graph, as defined in 4.3, dimensioned compatibly with the measurement vector y_k^i whose dimension is denoted m . Henceforth, we shall assume that $m = 1$, and dispense with the extra notation. As discussed in the previous chapter, the commutation result of Lemma 5.2 implies that the dimension of y is not relevant. For the information flow laws to be derived, one can replicate all the results by replacing the given transfer functions with the same transfer function repeated m times along the diagonal.

6.2.2 Convergence of the Information Flow Loop

Let us explore this information flow paradigm in some detail. The information flow component is a discrete time dynamical system which, as discussed above, is neutrally stable due to the Perron root of G . We begin by determining the steady-state behavior of the information flow loop. In preparation, we introduce some definitions and related lemmas.

Let e_r denote the right Perron eigenvector of G , and e_l its left Perron eigenvector, normalized such that $e_r^T e_l = 1$. If G is irreducible, both e_r and e_l are positive (Theorem 4.3), so such a scaling must exist. Let $E = e_r e_l^T$. The following relationships between G and E are known to be true: (See [51], p. 498, and recall that the Perron eigenvalue is 1.)

Lemma 6.1. $G^j = E + (G - E)^j$.

Lemma 6.2. *The eigenvalues of $G - E$ are the eigenvalues of G with the Perron eigenvalue replaced with a zero eigenvalue.*

We now state and prove the following theorem:

Theorem 6.3. *Suppose the directed graph $\mathcal{G}(G)$ is strongly connected and aperiodic, and let the input y_k be fixed in time. The steady state value of the dynamical*

system in Equation (6.5), when $p_0 = 0$, is

$$p_{ss}^i = y^i - \sum_{j=1}^N e_l^i y^j \quad (6.6)$$

where e_l^i is the i th element of the left Perron eigenvector of G , scaled so that $\sum e_l^i = 1$.

Proof. Consider the evolution of Equation (6.5):

$$p_k = G^k p_0 + \left(\sum_{j=0}^k G^j \right) Ly. \quad (6.7)$$

We assume that $p_0 = 0$, and we wish to find

$$p_{ss} = \lim_{k \rightarrow \infty} p_k, \quad (6.8)$$

if such a limit exists.

Substituting into Equation (6.7) via Lemma 6.1, we have

$$p_k = \left(\sum_{j=0}^k E^j + (G - E)^j \right) Ly. \quad (6.9)$$

Recalling that $E = e_r e_l^T$, and that L shares eigenvectors with G , we see that e_r and e_l are the eigenvectors of L corresponding to the zero eigenvalue. Therefore, $EL = e_r e_l^T L = e_r \cdot 0 = 0$, and we can rewrite p_k as

$$p_k = \left(\sum_{j=0}^k (G - E)^j \right) Ly. \quad (6.10)$$

Because G is assumed irreducible and aperiodic, all non-Perron eigenvalues of G have modulus strictly less than one (Theorem 4.3). Therefore, by Lemma 6.2, we see that $\rho(G - E) < 1$. The infinite expansion of p_{ss} therefore converges ([51], p. 301) and can be written as follows:

$$\begin{aligned} p_{ss} &= \left(\sum_{j=0}^{\infty} (G - E)^j \right) Ly \\ &= (I - G + E)^{-1} Ly \end{aligned}$$

$$\begin{aligned}
&= (L + E)^{-1}Ly \\
&= (L + E)^{-1}(L + E - E)y \\
&= (I - (L + E)^{-1}E)y.
\end{aligned} \tag{6.11}$$

Now $Le_r = 0$, and $Ee_r = (e_re_l^T)e_r = e_r(e_l^Te_r) = e_r$, so $(L + E)e_r = e_r \Rightarrow (L + E)^{-1}e_r = e_r$, and the above equation can be rewritten

$$\begin{aligned}
p_{ss} &= (I - (L + E)^{-1}E)y \\
&= (I - (L + E)^{-1}e_re_l^T)y \\
&= (I - e_re_l^T)y \\
&= (I - E)y.
\end{aligned} \tag{6.12}$$

We now interpret the above equation. The eigenvector e_r is known to be 1^T . The eigenvector e_l is positive, and is scaled such that $\sum e_l^i = 1$. The columns of E are therefore constant, and the rows are each e_l^T . Therefore, Equation (6.12) can be written

$$p_{ss}^i = y^i - \sum_{j=0}^N e_l^i y^j. \tag{6.13}$$

■

The information flow loop therefore has the effect of having the formation track the formation center, where the center is defined according to a weighting given by the graph. In this architecture, the weighting cannot be chosen, though in principle it could be set by unevenly weighting the information when performing the averaging. However, this would require global knowledge of the graph, which is assumed not to be available.

6.2.3 Shaping the Information Flow

In the above section, we looked at the response of the information flow law to a constant input. Of course, the input to the information flow law need not be constant; it will reflect the dynamics of the formation. The designer may wish to know the response of the filter to different inputs, and to design it to be within certain tolerances over the range of expected inputs. With that in mind,

we consider a more general form for the information flow filter:

$$\begin{aligned} q_{k+1} &= \sum_{j=0}^R a_j q_{k-j} + G \sum_{j=0}^R b_j q_{k-j} + Ly_k \\ p_k &= \sum_{j=0}^R c_j q_{k-j}. \end{aligned} \quad (6.14)$$

In this version, we are computing our current information using information from previous time steps as well as information received from other vehicles through a filter. This formulation can also be used to account for the presence of additional delays in data transmission. As in the previous case, we wish to determine the steady-state value of the filter for a constant input to understand the effects of the filter.

We begin by checking stability of the information flow law using the tools from Chapter 5:

Theorem 6.4. *The system in Equation (6.14) is (neutrally) stable if the transfer function*

$$F(z) = \frac{\sum_{j=0}^R b_j z^{R-j}}{z^{R+1} - \sum_{j=0}^R (a_j + b_j) z^{R-j}} \quad (6.15)$$

is (neutrally) stable and its Nyquist plot avoids encirclement of the negative inverse of any of the nonzero eigenvalues of L .

Proof. We can take the z -transform of Equation (6.14), setting aside the input, and rewrite it as follows:

$$\begin{aligned} zq(z) &= \sum_{j=0}^R a_j z^{-j} q(z) + \sum_{j=0}^R b_j G z^{-j} q(z) \\ &= \sum_{j=0}^R (a_j + b_j) z^{-j} q(z) - \sum_{j=0}^R b_j L z^{-j} q(z) \end{aligned}$$

or, if we collect terms not including L and multiply both sides by z^R ,

$$q(z) = -\frac{\sum_{j=0}^R b_j z^{R-j}}{z^{R+1} - \sum_{j=0}^R (a_j + b_j) z^{R-j}} Lq(z). \quad (6.16)$$

The transfer function in the above equation is $\hat{F}(z)$, and this equation is equivalent to the lower loop shown in Figure 6.2. This block diagram has the same structure as the system of vehicle formations examined in Chapter 5, where it was shown in Theorems 5.1 and 5.3 that the stability of this system is given by the Nyquist criterion stated above. Because one set of eigenvalues of this system corresponds




Figure 6.2: Block Diagram of Information Flow in the Loop.

to the open-loop dynamics, this system can be at best neutrally stable if $F(z)$ is itself neutrally stable. ■

We now turn to the steady-state performance of the information flow law. We assume that $c_j = b_j$, which will ensure that the information flow law does not add gain to the loop, and which will be useful in the stability proofs of Section 6.3. Additionally, if $c_j = b_j$, then it is only necessary for each vehicle to transmit $p_k = \sum_{j=0}^R b_j q_{k-j}^i$ to its neighbors. We also assume that $F(z)$ has all poles on the interior of the unit circle with the possible exception of a simple pole at 1. Finally, we assume that the polynomial $\sum_{i=0}^R a_i z^{R-i}$ has roots in the interior of the unit circle.

Theorem 6.5. *If $F(z)$ stabilizes L in the sense of Theorem 6.4, and under the above assumptions,*

$$p_{ss} = c \left(I - cE - (1 - c)(I - c(G - E))^{-1} G \right) y \quad (6.17)$$

where $a = \sum_{j=0}^R a_j$, $b = \sum_{j=0}^R b_j$, and $c = \frac{b}{1-a}$.

Proof. If we transform Equation (6.14) to the z -domain, it can be written as follows:

$$zq(z) = \sum_{j=0}^R a_j z^{-j} q(z) + \sum_{j=0}^R b_j G z^{-j} q(z) + L \frac{z}{z-1} y \quad (6.18)$$

The term $\frac{z}{z-1}y$ is the z -transform of a constant y . We recall the Final Value Theorem for discrete time systems [42]:

$$\lim_{k \rightarrow \infty} x(k) = \lim_{z \rightarrow 1} (z-1)X(z). \quad (6.19)$$

We can rewrite Equation (6.18) as

$$(z-1)q(z) = \left(zI_N - \sum_{j=0}^R a_j z^{-j} - \sum_{j=0}^R b_j G z^{-j} \right)^{-1} zLy. \quad (6.20)$$

The assumed neutral stability of the system implies that the matrix to be inverted is in fact nonsingular for all $|z| > 1$, with $z = 1$ the only possible pole of modulus 1. We can therefore set $z = 1 + \epsilon$, and be insured of the existence of the inverse for arbitrarily small ϵ . We will analyze the expansion for nonzero ϵ to verify the existence of the inverse when we take the limit, setting $\epsilon = 0$.

We begin by extracting a z from the inverted matrix to cancel the z which postmultiplies it. The resulting expression is

$$(z-1)q(z) = \left(I_N - \sum_{j=0}^R a_j z^{-j-1} - \sum_{j=0}^R b_j G z^{-j-1} \right)^{-1} Ly. \quad (6.21)$$

The invertibility of the matrix implies we can again use the expansion of the inverse, Lemma 6.2 and the fact that $EL = 0$ to rewrite our expression as

$$(z-1)q(z) = \left[I_N - \sum_{r=1}^{\infty} \left(\sum_{j=0}^R a_j z^{-j-1} - \sum_{j=0}^R b_j (G-E) z^{-j-1} \right)^r \right] Ly. \quad (6.22)$$

We now examine the eigenvalues of our system with G replaced by $G - E$ as ϵ goes to zero. The eigenvalues of this system are the poles of $F(z)$ with its loop closed about each eigenvalue of L . By assumption, they are all stable with the possible exception of a pole at 1, corresponding to closing the loop about the zero eigenvalue of L . When we replace G with $G - E$, we leave the eigenvalues unchanged with the exception of the Perron eigenvalue, which is now zero. The corresponding eigenvalues of our system are the eigenvalues of Equation (6.22) with the $G - E$

term removed. These eigenvalues are the roots of $\sum_{j=0}^R a_j z^{R-j}$, which are in the interior of the unit disk by assumption. We therefore conclude that the series in Equation (6.22) converges even when $\epsilon = 0$, and represent the equation in inverted form:

$$\lim_{k \rightarrow \infty} q(k) = (I_N - a - b(G - E))^{-1} Ly \quad (6.23)$$

where $a = \sum_{j=0}^R a_j$ and $b = \sum_{j=0}^R b_j$. Letting $c = \frac{b}{1-a}$, we rewrite this as

$$\begin{aligned} \lim_{k \rightarrow \infty} q(k) &= \frac{1}{1-a} (I_N - c(G - E))^{-1} Ly \\ &= \frac{1}{1-a} \left((I_N - c(G - E))^{-1} (I_N - c(G - E) - (1-c)G - cE) \right) y \\ &= \frac{1}{1-a} \left(I_N - (I_N - c(G - E))^{-1} (1-c)G + cE \right) y \\ &= \frac{1}{1-a} \left(I_N - cE - (1-c) (I_N - c(G - E))^{-1} G \right) y. \end{aligned} \quad (6.24)$$

The output $p(z)$ is (again, via Final Value Theorem)

$$\lim_{k \rightarrow \infty} p(k) = c \left(I_N - cE - (1-c) (I_N - c(G - E))^{-1} G \right) y. \quad (6.25)$$

■

Note that $c = 1$ corresponds to $a + b = 1$, which implies that the system has a pole at 1. When $c = 1$, we recover the steady-state result of Theorem 6.3, only we now see it to be true for any information flow filter with a pole at 1 (and which stabilizes the graph). When $c < 1$, the steady-state is offset by an additional term. Note that when $c = 1$, the vehicles all agree on the location of the formation center (expressed in each vehicle's coordinates), while when $c < 1$, they do not. We can say that when $c = 1$, the vehicles achieve *consensus* on formation center. From this perspective, having $c = 1$ appears to be a desirable property of the information flow filter. However, when $c = 1$, the system is only neutrally stable. The reason the filter converges to a steady state is because the input passed through L , whose kernel is equal to the Perron eigenvector of G . However, the presence of noise or sensor errors has the potential to introduce drift. Additionally, the eigenvalue at 1 means that old information never decays out, rendering the system sensitive to initial conditions. Of course, the initial conditions of the information flow law can be set (or reset) by the vehicle, assuming the existence of a protocol which guaranteed that this could be done without disrupting the formation.

6.2.4 Example

To understand the effects of shaping the information flow, we will consider two examples. The first filter is in Equation (6.5). In this case, following Equation (6.15), $F_1(z) = \frac{1}{z-1}$. The second filter is given by

$$p_{k+1} = 1.0625p_k - 0.2313p_{k-1} + 0.1875Gp_k - 0.0188Gp_{k-1} \quad (6.26)$$

$$q_k = 0.1875p_k - 0.0188p_{k-1}. \quad (6.27)$$

This corresponds to $F_2(z) = \frac{0.1875(z-0.1)}{(z-0.25)(z-1)}$. The pole at 1 means that $c = 1$ in both cases. Figure 6.3 shows the Nyquist plot for these two filters. The first lies along the -0.5 vertical. Points on that line correspond to periodic graphs (see Section 4.5), which confirms Theorem 6.3. The second lies entirely to the right of the -0.5 vertical. Recalling from Section 5.3 that the Nyquist plot points generated by the Laplacian, $-\lambda_i(L)^{-1}$, all must lie on or to the left of this vertical, we conclude that this information flow law stabilizes any graph.

Figure 6.4 shows the response of the two filters to a step response for a sparsely connected graph. While both settle in approximately 0.5 sec (using a time step of 0.02 sec), the first filter exhibits ringing due to the proximity of the closed loop poles to the unit circle. The second filter has a much smoother response. We see how the information flow filter can be designed to achieve desirable responses and robustness to uncertainty in the graph. The information flow filter should also be designed to have good tracking properties over the frequency range of the vehicle dynamics.

6.2.5 Information Flow in Weakly Connected Graphs

The results derived above assumed strong connectivity in the formation graph. In fact, we can generalize these results to the case of weakly connected graphs with a single leader component, as discussed in Section 5.2. The results derived above depended solely the Perron root being an isolated eigenvalue with the corresponding right eigenvector being 1^T . This is still the case when the graph is weakly connected with a single leader component. The difference will be in the weighting of the vehicles in calculation of the formation center. By examining the representation of G in Equation 4.5, it is easily verified that the left Perron eigenvector of G is equal to $[e_l^1, 0, 0, 0, \dots, 0]$, where e_l^1 is the left Perron eigenvector of the leader component. We can thus interpret the weighting of Theorem 6.3 to mean

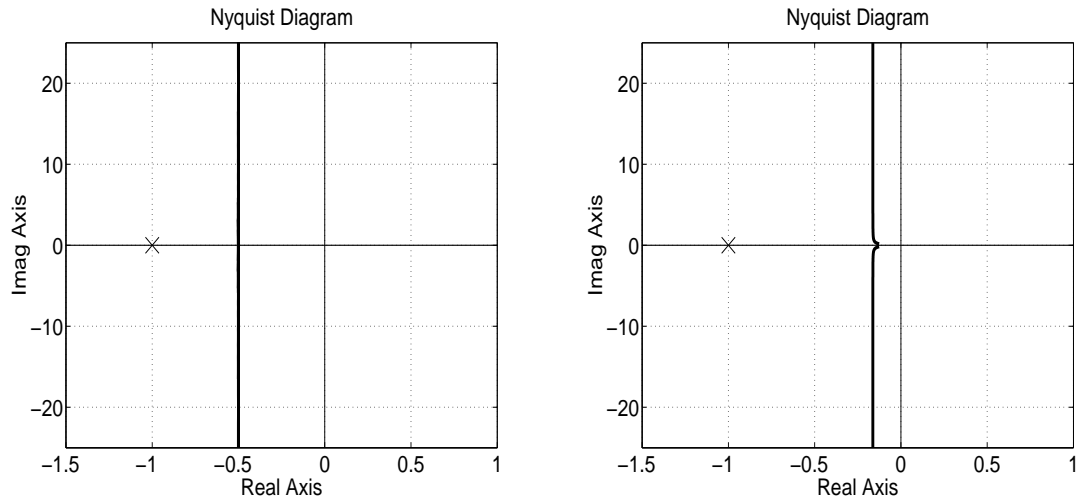


Figure 6.3: Information Filter Nyquist Plots.




Figure 6.4: Information Filter Convergence.

that only the position of the vehicles in the leader component are included in the determination of the formation center. This is clearly necessary, since the leader vehicles receive no information regarding the follower vehicles. What is significant is that the formation still achieves consensus on the formation center despite the lack of strong connectivity.

When multiple leader components exist, the situation is, unsurprisingly, far worse. In the extreme case, where G consists of two disconnected subgraphs, clearly each component will determine its own reference position, and the two components will act independently. When a follower component has access to two leader components, there will be two Perron eigenvectors, and 1^T will lie in the subspace spanned by them (Theorem 4.5). The reasoning used in Theorem (6.3) for removing the Perron eigenvalue can still be applied, and the information flow will achieve steady state. However, the formation will not achieve consensus: the reference supplied to a vehicle in the follower component will coincide neither with the reference in either leader component, nor with the reference supplied to another vehicle in the same component. This is because achieving consensus relied on the uniqueness of 1^T as the Perron eigenvector.

6.2.6 Mixed Absolute/Relative Sensing

The results of the above section give insight into extending these results to the case where some vehicles can sense their position relative to the target or to an absolute reference frame. In that case, the desired graph structure is that the target be the sole leader component of the graph to which all other vehicles have access. (Obviously, the target cannot be anything other than a leader, since it does not have access to the formation.) In this case, the Perron eigenvector, as discussed above, will be $[1, 0, 0, \dots]^T$, meaning each vehicle's information converges to its position relative to the target — a desirable result. Of course, the target is not transmitting any information, so the vehicles that sense the target must generate the information the target would be transmitting. Since the target is by definition at its desired location, the target's transmitted information is just zero.

6.3 Information Flow in the Loop

The information flow filter supplies each vehicle with the information it cannot sense: a formation center about which to do control. In this setting p is the input to the controller $K(z)$. A block diagram for this architecture is shown in Figure

6.2. As before, we can analyze stability with respect to uncertainties in the graph by isolating L and applying the Nyquist criterion as in Theorem 5.3. In this case, one determines stability by analyzing the Nyquist plot of

$$F(z)(1 + K(z)P(z)). \quad (6.28)$$

For a given plant and controller, the information flow loop can be designed to provide desirable margins. However, care must be taken in interpreting the stability margins derived from this plot. The gain and phase margins of this plot do not correspond to uncertainties in the plant in the typical fashion due to the location of $P(z)$ in the transfer function. Instead, they correspond more directly to uncertainties in L . Small variations in $P(z)$ can produce unexpected perturbations of the Nyquist plot. A reasonable design methodology is to design $K(z)$ to stabilize $P(z)$, without regard to the graph (remember that stabilizing the formation is never easier than stabilizing an individual vehicle) and then design $F(z)$ to stabilize L . However, the coupling between the dynamics of the two can produce unexpected results. In this section, we explore a means to improve this situation.

The information flow algorithm presented earlier is necessarily reactive; it does not anticipate the motion of the cluster. A logical means of improving performance is to supply the information flow loop with feedforward information regarding the expected motion of the formation.

Recalling that the information represents an averaged position of the vehicles' positions, a logical choice for a feedforward signal is the anticipated change in vehicle position. This can be calculated by using each vehicles' control signal $u(z)$ as the input to a model of the plant, denoted $\tilde{P}(z)$, and differencing that. The resulting signal

$$w^i(z) = (1 - z^{-1})\tilde{P}(z)u^i(z) \quad (6.29)$$

is then transmitted in addition to the signal $q(z)$ and used by each vehicle as a correction term to p . For example, we would replace Equation (6.5) with the following information flow law:

$$\begin{aligned} q_{k+1} &= G(q_k + w_k) + Ly_k \\ p_k &= q_k + w_k \end{aligned} \quad (6.30)$$

In this case, the transmitted quantity is $p_k^i = q_k^i + w_k^i$, as is clear from is premul-

tiplication by G . Of course, this feedforward correction term is only current if the control signal is delayed by a time step before application to the plant to allow a time step for the information to reach the other vehicles. Alternatively, each vehicle could delay the use of its sensed information until it receives the transmitted information from that vehicle.

To allow for information flow laws more general than Equation (6.29), we will let $w(z)$ take on the more general form

$$w(z) = H(z)\tilde{P}(z)u(z). \quad (6.31)$$

The information flow block diagram for this architecture is shown in Figure 6.5. When $H(z)$ is chosen properly, the following result can be derived:

Theorem 6.6. *Choose $H(z)$ to be*

$$H(z) = \frac{1}{F(z) + 1}, \quad (6.32)$$

and suppose the feedback interconnection of $P(z)$ and $K(z)$ is well-posed. Then the relative formation dynamics are stabilized if and only if $F(z)$ stabilizes L in the sense of Theorem 6.4 and $K(z)$ stabilizes $P(z)$.

Proof. By construction, $F(z)$ is biproper. Using our definition of $F(z)$, we can write $H(z)$ as

$$H(z) = \frac{z^{R+1} - \sum_{j=0}^R (a_j + b_j)z^{R-j}}{z^{R+1} - \sum_{j=0}^R a_j z^{R-j}}. \quad (6.33)$$

Note that $H(z)$ is stable due to the assumptions of Theorem 6.4.

We prove the presence of a separation principle for the system of equations, through the use of a transformation of coordinates that isolates the subsystems whose stability implies stability of the overall system. To do this, we first present the system of equations in state-space form. The state-space equations of motion for the plant are given in Equation (6.1). The predictor $\tilde{P}(z)$ is presumed to be identical to the plant $P(z)$, and has the same equations of motion with x, y replaced by \tilde{x}, \tilde{y} . The dynamics of the controller will be represented as

$$\begin{aligned} v_{k+1}^i &= K_A v_k^i + K_{BP}^i \\ u_k^i &= K_C v_k^i + K_{DP}^i. \end{aligned} \quad (6.34)$$

The information flow filter $F(z)$ is defined as found in Equation (6.14), but with

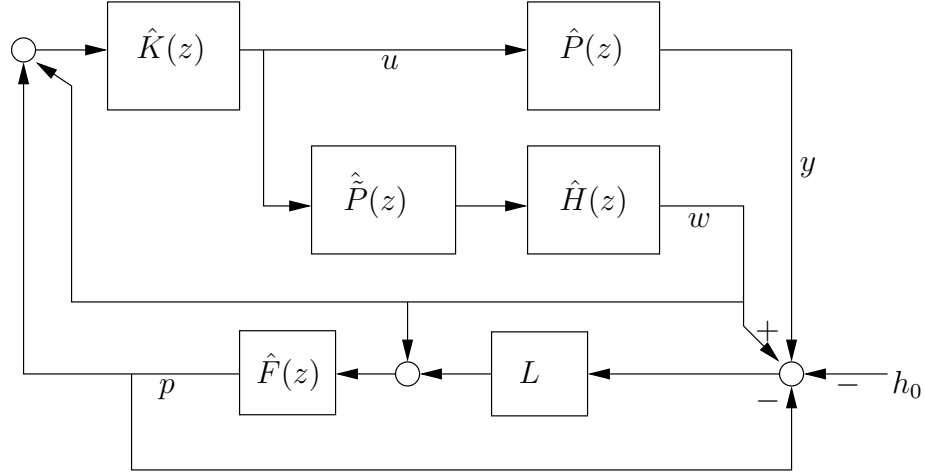


Figure 6.5: Block Diagram of Information Flow with Feedforward Correction.

the feedforward correction term added:

$$\begin{aligned} q_{k+1} &= \sum_{j=0}^R a_j q_{k-j} + G \left(\sum_{j=0}^R b_j q_{k-j} + w_k \right) + Ly_k \\ p_k &= \sum_{j=0}^R b_j q_{k-j} + w_k. \end{aligned} \quad (6.35)$$

Once again, it should be clear from the position of the quantity $p_k = \sum_{j=0}^R b_j q_{k-j} + w_k$ that it is the information which is transmitted by each vehicle.

Finally, the state space representation of the feedforward correction term, found in Equation (6.33), is

$$\begin{aligned} r_{k+1}^i &= \sum_j^R a_j r_{k-j}^i + \tilde{y}_k^i \\ w_k^i &= -\sum_j^R b_{k-j} r_k^i + \tilde{y}_k^i. \end{aligned} \quad (6.36)$$

To simplify the representation of Equations (6.35) and (6.36) in state-space notation, we introduce the following notation. Let

$$H_A = \begin{pmatrix} 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \\ a_R & a_{R-1} & \cdots & a_0 \end{pmatrix}, \quad (6.37)$$

let

$$H_B = \begin{bmatrix} 0 & 0 & \cdots & 1 \end{bmatrix}^T, \quad (6.38)$$

and let

$$H_C = \begin{bmatrix} b_R & \cdots & b_0 \end{bmatrix}, \quad (6.39)$$

where H_B is dimensioned compatibly with H_A . For the information flow law and feedforward term, we use \bar{q}^i to denote $[q_{k-R}^i, q_{k-R+1}^i, \dots, q_k^i]^T$, and similarly for \bar{r} . The state-space representation of Equation (6.35) can thus be represented as

$$\begin{aligned} \bar{q}_{k+1} &= \hat{H}_A \bar{q}_k + \hat{H}_B G (\hat{H}_C \bar{q}_k + w_k) + \hat{H}_B L y_k \\ p_k &= \hat{H}_C \bar{q}_k + w_k \end{aligned} \quad (6.40)$$

and of Equation (6.36) as

$$\begin{aligned} \bar{r}_{k+1} &= \hat{H}_A \bar{r}_k + \hat{H}_B \tilde{y}_k \\ w_k &= -\hat{H}_C \bar{r}_k + \tilde{y}_k. \end{aligned} \quad (6.41)$$

If one solves Equations (6.1), (6.34), (6.40), and (6.41) for the states, the resulting system of equations is

$$X_{k+1} = \Psi X_k, \quad (6.42)$$

where $X_k = [x_k, v_k, \tilde{x}_k, \bar{r}_k, \bar{q}_k]$ and

$$\Psi = \begin{pmatrix} \hat{P}_A & \hat{P}_B \Delta \hat{K}_C & \hat{P}_B \hat{K}_D \Delta \hat{P}_C & -\hat{P}_B \hat{K}_D \Delta \hat{H}_C & \hat{P}_B \hat{K}_D \Delta \hat{H}_C \\ 0 & \hat{K}_A + \hat{K}_B \hat{P}_D \Delta \hat{K}_C & \hat{K}_B \Delta \hat{P}_C & -\hat{K}_B \Delta \hat{H}_C & \hat{K}_B \Delta \hat{H}_C \\ 0 & \hat{P}_B \Delta \hat{K}_C & \hat{P}_A + \hat{P}_B \hat{K}_D \Delta \hat{P}_C & -\hat{P}_B \hat{K}_D \Delta \hat{H}_C & \hat{P}_B \hat{K}_D \Delta \hat{H}_C \\ 0 & \hat{H}_B \hat{P}_D \Delta \hat{K}_C & \hat{H}_B \Delta \hat{P}_C & \hat{H}_A - \hat{H}_B \hat{P}_D \hat{K}_D \Delta \hat{H}_C & \hat{H}_B \hat{P}_D \hat{K}_D \Delta \hat{H}_C \\ \hat{H}_B L \hat{P}_C & \hat{H}_B \hat{P}_D \Delta \hat{K}_C & \phi \hat{P}_C & -\phi \hat{H}_C & \hat{H}_A + \phi \hat{H}_C \end{pmatrix} \quad (6.43)$$

where $\phi = \hat{H}_B (\hat{P}_D \hat{K}_D \Delta + G)$, and $\Delta = (I - \hat{P}_D \hat{K}_D)^{-1}$, which is invertible by assumption of well-posedness of the interconnection. If we apply the transformation

$$T = \begin{pmatrix} I & 0 & 0 & I & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & 0 & 0 & I & 0 \\ 0 & I & 0 & 0 & I \\ 0 & 0 & 0 & 0 & I \end{pmatrix} \quad (6.44)$$

to the system matrix, we recover the matrix

$$T^{-1}\Psi T = \begin{pmatrix} \boxed{\hat{P}_A} & 0 & 0 & 0 & 0 \\ -\hat{H}_B L \hat{P}_C & \hat{H}_A + \hat{H}_B G \hat{H}_C & 0 & 0 & 0 \\ 0 & -\hat{K}_B \Delta \hat{H}_C & \hat{K}_A + \hat{K}_B \hat{P}_D \Delta \hat{K}_C & \hat{K}_B \Delta \hat{P}_C & 0 \\ 0 & -\hat{P}_B \hat{K}_D \Delta \hat{H}_C & \hat{P}_B \Delta \hat{K}_c & \hat{P}_A + \hat{P}_B \hat{K}_D \Delta \hat{K}_C & 0 \\ \hat{H}_B L \hat{P}_C & -\phi \hat{H}_C & \hat{H}_B \hat{P}_D \Delta \hat{K}_C & \hat{H}_B \Delta \hat{P}_C & \boxed{\hat{H}_A} \end{pmatrix}. \quad (6.45)$$

Stability of the system is equivalent to stability of the blocks along the diagonal. The first, P_A , is neutrally stable by assumption. The assumption that the information flow law stabilizes the graph is equivalent to the second block, $H_A + H_B G H_C$, being stable. The third block along the diagonal, which comprises the third and fourth columns/rows, is stable if $K(z)$ stabilizes $P(z)$. (The reader will verify that this is the matrix derived when $K(z)$ and $P(z)$ are interconnected directly via feedback.) The final block represents the states of $H(z)$, which is stable by the assumption in Theorem 6.4. We thus derive a separation principle for our formation which demonstrates that design of the individual vehicle controller $K(z)$ and the information flow filter $F(z)$ can be decoupled. ■

Remarks Equation (6.45) can be interpreted in the following way. The first set of states are open-loop copies of the vehicles' dynamics, and represent mismatches in initial conditions between the predictor and the actual vehicle. The second set is identical to the dynamics of Equation (6.14), whose stability and convergence properties were studied above. The output of this set of states acts as a reference to N more copies of the vehicle dynamics in feedback interconnection with the local controllers, found in the third and fourth rows. We see, therefore, that the effect of this architecture is to supply the local controllers with a reference signal which, if implemented properly, represents the error of that vehicle relative to a common reference trajectory whose dynamics obey the open loop dynamics of an individual vehicle. The final set of states represent the feedforward component. These states are unobservable in the motion of the vehicles, but are stable by design.

Several observations can be made regarding implementation. The first is that the motion of the formation is sensitive to mismatches between initial conditions of the vehicle and predictor. This can lead to drift of the cluster if the mismatch is in velocities. It should be possible to improve upon this through the use of an observer which will prevent the vehicle and predictor from diverging. Another

solution is to initialize the predictor using earlier measurements.

The second is that if $c \neq 1$, meaning the information flow loop does not converge to a common reference, then the vehicles' final positions will incorporate those errors as well (although the system is stable in this case as well). The position of the vehicles will also depend on the ability of the information flow filter to track the natural motion of the vehicles. When the natural motion of the vehicles is at rest, we have seen that it achieves a proper steady state when $c = 1$. When the natural motion is secular drift or oscillation (corresponding to poles at the origin or along the $j\omega$ axis), the quality of the reference signal will depend on the ability of the information flow filter to track signals at the relevant frequencies.

We also note that the model of the plant $\tilde{P}(z)$ is not an observer, but a predictor of vehicle motion. The zero at 1 in $H(z)$ corresponds to differencing the input, which generally amplifies signal noise. However, the input to $H(z)$ is derived by integrating $u(z)$, so no net differencing takes place in the filter. In fact, it is possible to compress $P(z)$ and $H(z)$ into a single filter, but it is easier not to do so when proving stability.

Finally, we note that unlike the results of the previous chapter, this separation principle does *not* rely on the vehicles having identical plants or controllers. It merely relies on each vehicle's predictor matching the vehicle dynamics and on each vehicle implementing the same information flow and feedforward correction computation. This eliminates a significant obstacle to implementation. A minor consequence is that when the vehicles have the same dynamics, the bulk motion of the formation itself obeys the dynamics of a single vehicle, while when the vehicles have different dynamics, that motion will be more complex.

6.4 Examples

6.4.1 Hovercraft Formations

Formation Acquisition

We return to the case with which we opened the chapter. If the information flow law together with feedforward compensation is enabled, the vehicles follow the trajectories shown in Figure 6.6. The trajectories are smoother, but still show some curving due to action of the control law prior to convergence of the information flow law. Figure 6.7 shows the trajectories followed by the vehicles when the information flow law is enabled one second prior to enabling the control loop. In

this case, the vehicles follow straight lines to their targets. Note that the formation center is identical in the two cases despite the differing trajectories. This is due to the decoupling of the information flow law from expected formation motion.

Formation Reconfiguration

In this example, the formation is already in a hexagonal formation and traveling in the positive y direction, when a command is issued for the formation to rotate counterclockwise. The transients for the motion of each vehicle in the y direction are shown in Figure 6.8 for the case where information flow is disabled and in Figure 6.9 for the case where information flow is enabled. The use of information flow reduces the transients associated with the reconfiguration as it did with the formation acquisition.



Target Acquisition

In this example a target becomes visible to a single vehicle as the formation is acquiring the hexagon. The vehicle which views the target includes that information as described in Section 5.2.4, and attempts to bring the vehicles into formation with the target at the center of the hexagon. Figure 6.10 shows the formation motion with information flow disabled. In this case, the vehicle which can view the target has to reconcile conflicting information: its position relative to the target and relative to the other vehicles, which are unaware of the target's existence. This causes the formation to overshoot the target, marked with a diamond, and to slowly settle into the desired position.

Figure 6.11 shows the same situation with information flow enabled. In this case, the information flow loop disseminates the target information to the other vehicles, causing the information loop to treat the target as the formation leader and use its position as the common reference, as discussed in Section 6.2.6. In this case, the formation gracefully changes course and quickly acquires the target.

6.4.2 Satellite Reconfiguration

As discussed in the introduction, a current area of research is control of relative satellite motion. The relative motion of a second satellite about a reference satellite can be approximated by linearizing the Keplerian orbital mechanics about the reference trajectory. These equations are known as Hill's equations due to Hill's study of lunar motion [50] and as the Clohessy-Wiltshire equations due to their

Figure 6.8: y -axis Transients of Formation Reconfiguration, no Info Flow.Figure 6.9: y -axis Transients of Formation Reconfiguration, with Info Flow.




Figure 6.10: Target Acquisition, no Info Flow.




Figure 6.11: Target Acquisition, with Info Flow.

work on orbital rendezvous [26]. The equations of motion about a circular reference orbit are

$$\begin{aligned}\ddot{x} &= 3n^2x + 2n\dot{y} + a_x \\ \ddot{y} &= -2n\dot{x} + a_y \\ \ddot{z} &= -n^2z + a_z\end{aligned}\tag{6.46}$$

where x point in the radial direction, y along track, and z out of plane. The orbit rate is given by n . The vector $[a_x, a_y, a_z]^T$ represents external accelerations, either due to environmental disturbances or applied thrust. The equations, when solved, reveal families of periodic orbits about the origin as well as secular drift along y . The periodic solutions include a 2×1 inclined ellipse whose projection onto the Earth (the yz plane) is a circle. The ground track of these satellites remains fixed relative to one another and rotates at orbit rate. These orbits are attractive for space-based interferometry, such as the TechSat21 mission, and were part of the impetus for exploring missions of this type [59]. We also note that interferometry requires accurate knowledge of relative satellite position, but drift in absolute position is more tolerable. Hill's equations are often used as a coarse model for relative satellite motion despite the absence in the model of external forces and perturbations of Earth's gravitational field.

Note that the xy -dynamics are decoupled from the z -dynamics. Setting $n = 1$, a family of unforced solutions to the xy dynamics is given by

$$\begin{aligned}x(t) &= A \cos(t + \phi) \\ y(t) &= y_0 - 2A \sin(t + \phi).\end{aligned}\tag{6.47}$$

Consider a set of six satellites, evenly spaced initially along the y -axis, that are asked to take up stations along a Hill's ellipse given by $A = 1$ at evenly spaced ϕ . Each satellite can measure the full relative states of a subset of other satellites, and an LQR controller has been designed. The offset h_{i0} is given by Equation (6.47) and its derivative, with $\phi = \frac{\pi i}{3}$, and $y_0 = 0$. We begin by designing the information flow law. In this case, the reference signal which the vehicles must determine follows a periodic trajectory. To ensure good tracking, we place poles of the information flow law at the (discretized) frequency locations, along with a pole at 1 so that the $c = 1$ condition is satisfied. A candidate information flow law is given by

$$F(z) = \frac{z^2 - 1.6575z + 0.7225}{z^3 - 2.9975z^2 + 2.9975z - 1}.\tag{6.48}$$

The Nyquist plot for this information flow law is found in Figure 6.12. In this

case, the desire for good tracking of the reference signal places limits on the range of graphs which the information flow law stabilizes. Nonetheless, the encircled region of Figure 6.12, which offsets the encirclement at infinity, leading to zero net encirclements, gives reasonable latitude around the -1 point. Once $F(z)$ is designed, then the feedforward term $H(z)$ is derived automatically. Figure 6.13 shows the reference signal supplied to each satellite for measurements of y converging to a common trajectory, and Figure 6.14 shows the motion of the satellites in the xy plane. The initial positions of the satellite are at the center, and the final positions are marked with an ‘x.’

As in the previous case, our information flow approach greatly enhances stability. It is important to note that while the information flow law was restricted to those with good tracking performance at the reference trajectory frequencies, this is far less restrictive than the design of a controller which stabilized the vehicle dynamics. The difference is most noticeable in the presence of plant zeros, as would occur when only a subset of states are measurable.

6.5 Information Flow and String Stability

Thus far, we have seen that the proper design of information flow leads to improved stability due to the separation principle of Theorem 6.4, and improved vehicle trajectories due to the achievement of consensus among the vehicles as to the formation center. In this section, we turn our attention to disturbance rejection in the formation. A well-known area of concern within leader-follower formations is the possibility that a following vehicle can amplify disturbances of a leading vehicle. Depending on the length of the chain of vehicles, this can lead to unacceptably large disturbances of the vehicles at the end of the chain. If we posit an infinite chain of vehicles, and define the difference between the positions of vehicles i and $i + 1$ at time step k as $e^i(k) = y_k^{i+1}(t) - y_k^i(t)$, the effect of disturbance of the lead vehicle on the behavior of follower vehicles can be determined by looking at the sequence

$$\{\|e^1(k)\|, \|e^2(k)\|, \dots, \|e^i(k)\|, \dots\}. \quad (6.49)$$

If the sequence converges to zero for any bounded input disturbance to the lead vehicle, the formation is said to be *string stable*. If the sequence remains bounded, it is said to be *weakly string stable*. When the graph consists of a single chain,

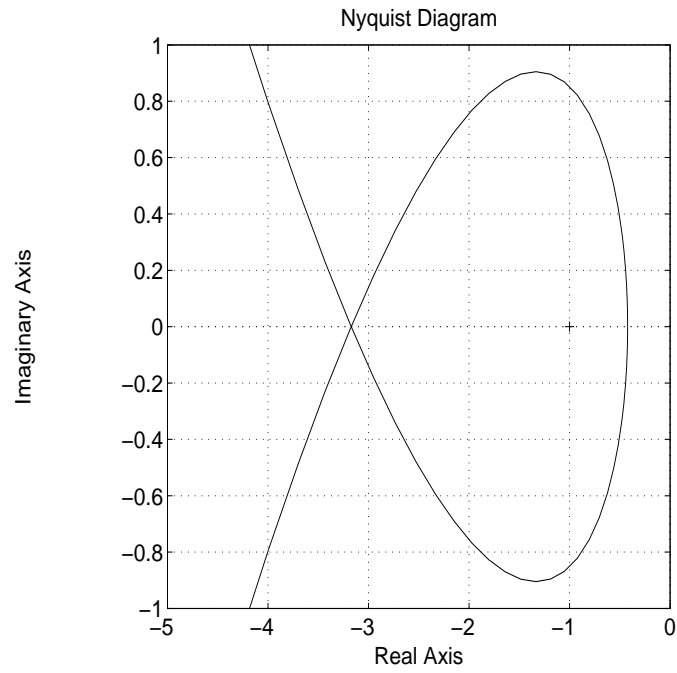



Figure 6.12: Nyquist Plot for Satellite Info Flow.

Figure 6.13: Convergence of Satellite y -axis Reference Data.




Figure 6.14: Satellite Reconfiguration Trajectories.

string stability is often determined by deriving the transfer function such that

$$e^{i+1}(z) = T(z)e^i(z) \quad (6.50)$$

and checking the infinity norm of $T(z)$. If it is less than one, then the sequence converges under the 2-norm, and thus the formation is string stable. If it is one, then the sequence remains bounded under the 2-norm, and the formation is weakly string stable. See [96] and the references therein for more formal definitions of string stability, and see [95] for a nonlinear approach.

A well-known property of leader-follower formation is that with only relative state knowledge of the immediately preceding vehicle, even weak string stability cannot be achieved [96, 106]. Candidate solutions explored in these references include having some knowledge of the lead vehicle's position or employing a variable speed spacing policy (i.e., having $h_{i,i+1} = h_{i,i+1}^0 + h_{i,i+1}^1 \dot{e}^i$). The former appears to be the better strategy, but it relies on the ability to transmit the lead vehicle's position down the chain infinitely quickly.

With this in mind, we will analyze the string stability of a single chain formation with information flow and feedforward correction enabled. A complete analysis of

this approach from the perspective of string stability is beyond the scope of this thesis, but insofar as the single chain is a known worst case, we rely on it to demonstrate the utility of the method. Because our system involves two pieces of information, namely position and transmitted information, we will need to consider the behavior of both as they are transmitted down the chain. We will denote relative position as

$$\Delta y_k^i = y_k^{i+1} - y_k^i \quad (6.51)$$

and similarly for the other variables, e.g., Δp_k^i representing the difference between p_k^{i+1} and p_k^i . From Equations (6.1), (6.34), (6.40), and (6.41), we can derive the following:

$$\Delta y^i(z) = P(z)\Delta u^i(z) \quad (6.52)$$

$$\Delta \tilde{y}^i(z) = P(z)\Delta u^i(z) \quad (6.53)$$

$$\Delta u^i(z) = K(z)\Delta p^i(z) \quad (6.54)$$

$$\Delta p^i(z) = \Delta w^i(z) + (I - H(z))(\Delta p^{i-1}(z) + \Delta y^i(z) - \Delta y^{i-1}(z)) \quad (6.55)$$

$$\Delta w^i(z) = H(z)\Delta \tilde{y}^i(z), \quad (6.56)$$

where

$$P(z) = P_C(zI - P_A)^{-1}P_B + P_D \quad (6.57)$$

$$K(z) = K_C(zI - K_A)^{-1}K_B + K_D \quad (6.58)$$

$$H(z) = -H_C(zI - H_A)^{-1}H_B + I. \quad (6.59)$$

These equations can be simplified to

$$\Delta y^i(z) = P(z)K(z) [(I - H(z)) (\Delta p^{i-1}(z) - \Delta y^{i-1}(z)) + \Delta y^i] \quad (6.60)$$

$$\Delta p^i(z) = (I - H(z)) (\Delta p^{i-1}(z) - \Delta y^{i-1}(z)) + \Delta y^i, \quad (6.61)$$

which can further be simplified to

$$\Delta y^i(z) = C(z)(I - H(z))(\Delta p^{i-1}(z) - \Delta y^{i-1}(z)) \quad (6.62)$$

$$\Delta p^i(z) - \Delta y^i(z) = (I - H(z))(\Delta p^{i-1}(z) - \Delta y^{i-1}(z)), \quad (6.63)$$

where $C(z) = (I - P(z)K(z))^{-1}P(z)K(z)$. It is now clear that string stability of the quantity $\Delta p^i(z) - \Delta y^i(z)$ depends solely on the ∞ -norm of $I - H(z)$. Further-

more, the presumed stability of the closed loop transfer function $C(z)$ implies that if $\Delta p_k^i - \Delta y_k^i$ is bounded, then Δy_k^i will be bounded, and if $\Delta p_k^i - \Delta y_k^i$ decays to zero, then so will Δy_k^i . We thus arrive at the intriguing result that string stability is independent of the plant and controller, and depends solely on the design of the information flow filter. Recalling that $H(z)$ has a zero at $z = 1$, it follows that $I - H(z) = 1$ at $z = 1$, and thus the ∞ -norm cannot be less than 1, so weak string stability is the best that can be achieved. However, design of information flow laws which achieve weak string stability does not appear difficult: Figure 6.15 shows the Bode plots of $I - H_1(z)$ and $I - H_2(z)$ as defined relative to $F_1(z)$, $F_2(z)$ in Section 6.2.4. In both cases, the transfer functions achieve weak string stability. The first information flow law has $|I - H_1(z)| = 1$ at all frequencies, which does not imply particularly good disturbance rejection. The second has better string stability properties, realizing a gain of one only at DC.

Swaroop and Hedrick [95] argue that when lead vehicle information is not available, weak string stability is the best result than can be achieved, and these results do not contradict that. Additionally, because only weak string stability is achieved, it cannot be assumed that this condition is robust to errors in the predictor model. However, it is significant that the information flow methodology presented here provides a systematic method for achieving the best possible (theoretical) result in a fashion which does not make assumptions on the plant or controller. While this result is only derived here for a single chain formation, because the result is essentially a consequence of separation principle, it should hold for more complex leader-follower architectures.

6.6 Conclusions

The information flow architecture presented in this chapter relies on two key ideas. The first is the use of dynamical systems as a paradigm for understanding information exchange between vehicles, and the design of a dynamical system which enables the vehicles to achieve consensus on the formation center. The second is the use of feedforward compensation to render the sensed and transmitted information timely. The resulting architecture achieves improvements in stability, vehicle trajectories, disturbance rejection, and robustness to changes in the interconnection structure. The architecture is flexible in that it does not rely on uniform vehicle dynamics, nor does it rely on a vehicle having any global knowledge of the information flow graph.




Figure 6.15: Information Flow Law Bode Plot, $I - H_1(z)$, $I - H_2(z)$.

We have also seen limitations to the method in the course of the derivation. The first limitation is the need for an exact model of the vehicle dynamics. The sensitivity of the method to modeling errors has not been analyzed, nor has it been validated in an experimental setting. Simulation results do not expose high sensitivity to modeling errors. Another limitation is the sensitivity to mismatches in initial conditions, particularly in velocities, between the vehicle and the predictor. It may be possible to improve on this through the use of an observer rather than a predictor. A third limitation is the constraint that $c = 1$ in the information flow law. The need for consensus among vehicles forces the information flow law to be neutrally stable, which means that information never decays out. This renders the system sensitive to sensor errors that cause the vector of measurements Ly_k to have a component which lies along the Perron vector, leading to secular drift of the information flow. One possibility for improving on this is having a protocol for resetting the information to zero periodically or in response to an event as a means of limiting any drift. Such a protocol could lie in a higher layer in the control architecture, and may itself require stability analysis. All these issues represent avenues of future research for improving an already potent method for vehicle formation control.

Chapter 7

Conclusions and Future Work

7.1 Thesis Summary

In this thesis, we have explored two topics relevant to control of vehicle formations. The first is optimal control of vehicle motion, and the second is cooperative control of vehicle formations. In both cases, we employed tools outside the domain of linear control to find an analytical framework in which to address these issues. For the former, that framework was differential geometry and the use of the affine connection. For the latter, it was the use of graph theory and Perron-Frobenius theory. The results of each section, which we will summarize shortly, open up new avenues for research into vehicle control.

In Chapters 2 and 3, we examined the two classical approaches to optimal control, namely the Euler-Lagrange equations and the Hamilton-Jacobi-Bellman equation from the perspective of vehicle control. Our modeling framework was that of the affine connection, which facilitates a geometric expression of the equations of motion. We derived a novel form of the Euler-Lagrange equations using the affine connection and revealed the role of the affine connection in the evolution of the adjoint variables. When we restricted our view to vehicles not affected by external forces such as potential gradients or damping, the vehicle dynamics were observed to be time-scalable, a property which had significant implications for the HJB approach to optimal control. We showed how time can be eliminated from the PDE, leading to a purely spatial computation, which defines both the optimal finite-time trajectories as well as, potentially, a stabilizing control law.

In Chapters 4–6, we explored the role of communication links in cooperative control of vehicle formations. Our observation that the ability of vehicles to sense or communicate with one another is itself a source of uncertainty led to an inves-

tigation of the effects of the formation graph on stability. We proved a Nyquist criterion for formation stability that is useful in design of stabilizing controllers and that provides insight into how the formation graph affects stability. Information exchange between vehicles gives the formation the possibility to overcome these limitations, but raises concerns that the information exchange itself may not be robust to changes in the communication topologies. To overcome this, we explored a decentralized information flow paradigm that converges to an acceptable solution for almost any graph. The information flow law represents a design parameter that can be chosen to have suitable transient properties and accuracy over the frequency range of interest. When the information flow law is coupled with a vehicle motion predictor, we achieved the striking result that formation stability decouples into two separate problems: the stabilization of an individual vehicle by the controller and the stabilization of the graph by the information flow law. This architecture has superior stability properties, in that the loss of stability margin to the graph is recovered. It also exhibits smoother vehicle trajectories due to rapid convergence of the information flow loop and good disturbance rejection properties in terms of string stability. Central to this was a dynamical systems approach to information exchange, whereby the information held by each vehicle was treated as the state of a system whose response could be shaped.

7.2 Future Work in Optimal Control of Vehicles

The results presented in this thesis raise many questions and open interesting avenues of research. In the case of optimal vehicle control, the incorporation of our framework into an effective tool for computing optimal trajectories is, as yet, an unrealized goal. In the case of cooperative formation control, there are many possible extensions to the method, as well as challenges to overcome in order to realize its application within a real-world setting. In the sections which follow, we outline research directions which spring from this thesis.

7.2.1 Geometry of Optimal Control Equations

In Chapter 2, we presented a geometric formulation of the optimal control equations of motion. Exploration of the geometry of those equations within the affine connection framework has not been performed. A traditional area of interest within this setting is the role of symmetry in the equations. Preliminary investigation has revealed that the equations of motion for the adjoint variable can be reduced in a

fashion similar to that of the state variables, but the structure of that reduction and the role of the affine connection within it has not been fully explored. Once the reduction framework is revealed, it will be possible to analyze how optimal trajectories within the reduced space generate motion in the unreduced space. Resulting methods for optimal locomotion of affine connection control systems would be a worthwhile extension of the methods developed in this thesis.

7.2.2 Approximation of Optimal Trajectories

In general, the geometric formulation of the optimal control equations does not alter the fact that solving them involves solving a two-point boundary-value problem, which is computationally difficult. One approach that can simplify this is to approximate the solution of the ODE, thereby generating a map from initial to final conditions which could then be solved more easily. The geometric formulation is well-suited to such an approach. In the next two paragraphs, we discuss two methods for approximate solution which exploit the framework developed in this thesis.

Variational Integrators One active area of research has been the development of integrators for equations of motion which exploit the variational structure of the equations of motion. This approach to integration, developed by Marsden and coworkers [54, 55, 101], is attractive in that it preserves both conserved quantities and the underlying geometric (in this case symplectic) structure of the equations. In this thesis, we have developed an understanding of the variational approach to the optimal control equations for mechanical systems. As such, this work is well positioned to facilitate an application of variational integration techniques to approximate solution of the equations of motion.

Series Expansions Another recent area of research, pursued by Bullo [15, 16], is the derivation of series expansions of solutions to ODEs defined via affine connections. Homogeneity properties of the affine connection are exploited to show that series expansions such as the Campbell-Backer-Hausdorff expansion require only finite levels of bracketing of the associated vector fields. Since we have shown that the optimal control equations are also defined by an affine connection, it should be possible to apply these series expansions to these equations and arrive at an approximate solution.

7.2.3 HJB Equation and Sub-Riemannian Geometry

The approach to simplifying the HJB equation developed in Chapter 3 was successful in eliminating time from the PDE, but the problem still relies on solution of a PDE. When the vehicle being controlled is underactuated, the value function will likely exhibit nonsmoothness, as we observed in the example of the nonholonomic integrator. Because of that, the PDE is even more difficult to solve; it cannot even be locally approximated by a quadratic. What facilitated solution of the PDE in that case was *a priori* knowledge of the points of nonsmoothness. In general, those will not be known. We also observed that the optimal trajectories originating from those points are not unique. These phenomena have been observed before and are often associated with the relationship between the optimal control problem and problems in sub-Riemannian geometry [12]. Geodesics of a sub-Riemannian manifold exhibit local nonuniqueness, entirely analagous to the nonunique optimal trajectories we observed. Recent research [25] has addressed the computation of the so-called conjugate loci of the sub-Riemannian manifold. An interesting research direction would be to employ those tools within our framework to identify to points of nonsmoothness, and then to identify other methods for solution or approximation within the smooth regions.

7.3 Future Work in Cooperative Control of Vehicle Formations

Research into cooperative control of vehicle formation is very much in its infancy. Success in this arena will involve merging tools from computer science such as reliable protocols, algorithm design, and networking with insights from controls regarding dynamics, uncertainty and design. Our goal in this thesis was to suggest a framework in which ideas from both fields — in this case, graph theory and stability analysis — can coexist. To do so, we constructed an idealized setting in which the links between the two could be identified and exploited. Advancing these ideas to a point where they can be implemented in a realistic setting will require relaxing many of these simplifications and verifying if the core ideas of this methodology can still be exploited. One source of insight will undoubtedly be laboratory experimentation. In the sections which follow, we outline directions in which the theory can be advanced to pave the way for implementation.

7.3.1 Graph Periodicity and Laplacian Eigenvalues

In Section 5.3.1, we discussed how periodic graphs represent a “worst-case” scenario in terms of formation stability margin. Furthermore, the near-periodicity of a graph was meaningful in characterizing the effects of adding or removing sources of information in the sensing graph. At present, we lack tools to make that notion precise. An interesting avenue of research would be to define a *measure* of periodicity of a graph (as opposed to a graph being either periodic or aperiodic), which could then be correlated to Laplacian eigenvalue locations. One reason this has not been addressed to date is the general focus on undirected graphs. Another is the focus on algebraic connectivity and its relation to substructures in the graph. While algebraic connectivity may not be related to graph periodicity, the location of the other eigenvalues, which impact stability more directly, may submit to such a characterization.

The ability to characterize eigenvalue location would be a significant advance of the methods developed in this thesis. For example, we currently lack an edge-weighting methodology to improve stability margins, short of direct computation on the graph. The ability of individual vehicles to determine which edges augment stability margins would be quite useful. Insofar as periodicity is a quasi-local property (each vehicle need only know about its neighbors k edges away to know if it possesses any k -cycles), it may be possible to implement an edge-weighting algorithm which does not require global knowledge of the graph.

7.3.2 Information-Rich Control

In this thesis, we examined formation stability for two transmission graphs: $\mathcal{J}_i^T = \emptyset$, meaning no transmission, and $\mathcal{J}_i^T = \mathcal{J}_i^S$, meaning the transmission and sensing graphs are identical. We have not examined the case where $\mathcal{J}_i^T \supset \mathcal{J}_i^S$, meaning each vehicle receives information from vehicles it cannot directly sense. This, too, is a realistic setting, since the ability to sense other vehicles is often more limited than the ability to receive transmissions. Another area into which the methodology can be extended is this information-rich setting. Because the results of Chapter 6 depend on equal sensing and transmission graphs for convergence of the information flow filter, the extension to the information-rich case is not immediately apparent. Research on methods to incorporate additional information while retaining the stability separation principle would therefore be valuable in extending these techniques.

7.3.3 Vehicles with Nonlinear Dynamics

In this thesis, we focused on vehicles with linear dynamics for the purpose of elucidating the role which graph theory plays in the analysis of formation dynamics. Of course, many vehicles exhibit nonlinear dynamics. A significant research direction is the extension of the ideas developed in this thesis to vehicles which exhibit nonlinear dynamics. The results of Chapter 5, with their reliance on eigenvalues, may not readily exhibit a nonlinear characterization; the results of Chapter 6, however, appear more promising. It is reasonable to expect that the separation principle of Theorem 6.6 will survive when the vehicle and control law are nonlinear. Exploration of this extension is a logical next step in advancing the ideas of this thesis.

Within the universe of nonlinear controller design tools, the method of controlled Lagrangians developed by Bloch, Leonard, Marsden and their students [5, 8, 9, 10] stands out as a natural choice for implementation in this setting. While we did not focus on the geometry of vehicle formations in Chapters 5 and 6, the geometry of many vehicle systems corresponds nicely to the setting employed in the method of controlled Lagrangians. As we discussed in Chapter 2, many vehicles admit a Lagrangian formulation of their dynamics. Often, the Lagrangian is cyclic with respect to variables corresponding to position or orientation of the vehicle, which in turn are the directions along which sensing and actuation occur. It is precisely this setting for which the method of controlled Lagrangians was developed, in which one exploits this geometry and Lagrangian structure in designing controllers which tend to exhibit large regions of attraction. The marriage of those tools with the information flow paradigm put forth in Chapter 6 would both expand the range of systems to which information flow can be applied and facilitate the systematic implementation of the method of controlled Lagrangians to formation control.

One research area needed to accomplish this is the extension of the separation principle to nonlinear systems, as discussed above. Another is the extension of the information flow method from the case where the sensed information is in \mathbb{R}^m to the case where it lives in an arbitrary manifold. Many vehicles' dynamics evolve on Lie groups, such as $SO(3)$ for satellite orientation or $SE(3)$ for underwater vehicles. The results of Chapter 6 rely on the Perron-Frobenius properties of specific matrices whose analogs when acting on Lie groups are not immediately clear. Even designing an information flow loop whose input is in S^1 is not fully

understood. This is a topic of particular importance, since it would enable the vehicles to achieve consensus as to relative orientation as well as relative position. To accomplish this, one must consider the geometric analog of the cone-invariance property exhibited by Perron-Frobenius operators (see [2]).

7.3.4 Information Flow and Software-Enabled Control

The linear controllers employed in this thesis are the simplest form of vehicle control possible, and certainly does not represent the state of the art in vehicle control. Just as the methods of this thesis should be extended to vehicles with nonlinear dynamics, they should also be investigated with regard to more advanced control paradigms. In particular, recent research has exploited advances in computational power to control vehicles via real-time optimization. Recent results include Lyapunov-based stability proofs for receding-horizon controllers [53, 85] as well as real-time implementation of aggressive trajectory tracking via optimization based control [37, 76]. Work on applying this methodology to control of vehicle formations is in its infancy. Distributed optimization faces the same obstacle addressed in this thesis: the need for each vehicle to perform a computation (in this case trajectory optimization) with only limited information about the state of the formation.

The methods developed in this thesis may be of use in addressing this problem. For example, formation trajectory optimization is often complicated by the fact that one must optimize not only over trajectories but over the endpoints as well. It is optimization over endpoints which couples each vehicle's optimization problem. One approach which would incorporate ideas from this thesis is for the formation to achieve a suboptimal consensus as to the endpoint locations, and then for each vehicle to optimize its trajectory independently. Of course, stability using the class of controller in this setting would have to be verified using suitable extensions of the tools developed in this thesis. Another approach currently under investigation is for each vehicle to optimize its trajectory using a simplified model of the other vehicles' motion [36]. Again, the stability analysis of such an approach will depend on what limited information is available to each vehicle, which in turn incorporates ideas explored in this thesis.

7.3.5 Vehicle Control Over Networks

One of the simplifications made in the development of the information flow algorithm was the assumption that information is transmitted and received at evenly spaced time steps. Means of transmission such as networks often exhibit unknown latencies, which represent a significant obstacle in controlling across them. A recent area of research is in determining the nature of time delays and generating control laws which stabilize over long random delays. [68, 99, 102]. To make our approach useful for vehicles which share information over networks, it must be extended to the case of variable time delays. While the analytical setting for doing so may be different, there is reason to believe that such an extension will be successful. The feedforward correction term through which one achieves the stability separation principle can overcome the latency in transmission. By predicting the expected change in the transmitted information, it renders the different sources of information (sensed and transmitted, in this case) coincident in time. When this is achieved, the graph ceases to impact vehicle stability. It should be possible to realize this in a setting where time delays are variable. For example, if each vehicle transmits its expected motion over some time horizon (as is often computed in optimization-based control methods), the recipient vehicle could use that information to synchronize its various sources of information so long as the latency does not exceed the transmitted horizon. The ability to synchronize information sources would then act as a test which each vehicle could implement to decide whether to incorporate that new information into its decision-making framework.

7.3.6 Other Vehicle Formation Tasks

Relative position maintenance is only one of the tasks a formation may undertake. In a realistic setting, the formation will have to perform more complex tasks. The position controller would represent only one layer of the control architecture. Because that layer operates at the highest frequency, it is the most sensitive to uncertainty and changes in the communication topology, and therefore represented a logical starting point for analysis. However, future research must investigate the interaction between this layer of the formation control architecture and events, less frequent, which force changes in the architecture. The ability of the formation to handle these tasks on a decentralized level of the control hierarchy will likely run into many of the issues examined in this thesis. Several examples of more complex formation tasks are examined below.

Role Selection In this thesis, we assumed that each vehicle's assignment in the formation was predetermined and known to each other vehicle. In many settings, vehicle assignments may be interchangeable; so long as each station of the formation is occupied, it does not matter which vehicle occupies it. Furthermore, external influences or a change in formation goal may effect a change in vehicle assignments. The ability of the formation to assign roles to each vehicle in a distributed setting is an interesting area of research. For example, two vehicles in communication with one another could agree to switch roles, so long as they can communicate to the rest of the formation. Again, the ability of a distributed role-switching algorithm to converge to an agreeable role distribution would have to be examined. Unlike the stability analyses of this thesis, role assignments do not live on a manifold, so analyses may involve more traditional algorithmic analysis using tools from computer science. Nonetheless, the insights from this thesis will still be relevant.

Collision Avoidance The information flow algorithm developed in this thesis allows each vehicle to react to the overall formation motion rather than the motion of the vehicles it can sense. While this is generally a desirable goal, in some instances, such as collision avoidance, a vehicle must ignore the rest of the formation and react only to its local situation. As we have seen, such an action may have significant consequences for the information flow algorithm and for stability. One possibility is for the vehicle to communicate its change in control policy to nearby vehicles, which can then ignore its motion. Because collision avoidance often requires large control effort and fast response times, the ability of the formation to rapidly disseminate this new information will be essential.

Formation Merging/Splitting Another common objective is the merging or splitting of formations. For example, a formation which encounters an obstacle may choose to have different vehicles in the formation circumvent the obstacle in different directions, and reunite on the far side of the obstacle. Again, the vehicles must come to a consensus as to which vehicles belong to the same formation so the problem of multiple leader components does not occur. The decision for formations to merge may depend on sufficient level of inter-formation communication to avoid the stability issues associated with sparseness in formation graphs.

7.4 Optimal Control of Vehicle Formations

Throughout this thesis, we have treated optimal vehicle control and cooperative control of vehicle formations as distinct topics. The ultimate goal of this research effort remains a synthesis of the two topics. We close the thesis with a discussion of how the two main topics of this thesis could be jointly employed.

Any sensible approach to vehicle formation control involves the partitioning of the formation control problem into tractable subproblems. In the case of optimal formation control, the dynamic decoupling discussed in the introduction facilitates the partitioning of the optimal formation control problem into a set of N single vehicle optimal control problems which can be solved independently and whose solution would involve the methods explored in the first section of this thesis. Of course, the formation control problem is not totally decoupled — achieving the formation goal and preventing collision avoidance couple the trajectory generation problem. In our development, these joint issues are handled by the information flow. In chapter 6, the “joint” component, meaning determining a formation center, is handled by the information flow, while each vehicle’s trajectory is determined by a local control law. These ideas can be extended by replacing the linear control laws of Chapters 5 and 6 with nonlinear optimal control laws and relying on the information flow to supply each vehicle with the information it needs (e.g. trajectory endpoints, constraints, cost functions) to perform its local computation. Of course, extending these ideas to a nonlinear optimal control setting is far from trivial, but we expect that tools developed here for achieving robust, distributed information sharing with good stability properties will be essential in that setting as well.

Bibliography

- [1] J. G. Bender. An overview of systems studies of automated highway systems. *IEEE Transactions on Vehicular Technology*, 40(1):82–99, 1991.
- [2] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, 1979.
- [3] A. Berman and X.D. Zhang. Lower bounds for the eigenvalues of Laplacian matrices. *Linear Algebra and Applications*, 316(1-3):13–20, 2000.
- [4] A. Bloch and P. Crouch. Nonholonomic control systems on Riemannian manifolds. *SIAM Journal of Control and Optimization*, 33(1):126–148, 1995.
- [5] A. M. Bloch, D.-E. Chang, N. E. Leonard, and J. E. Marsden. Controlled Lagrangians and the stabilization of mechanical systems II: Potential shaping. *IEEE Transactions on Automatic Control*, 46(10):1556–1571, 2001.
- [6] A. M. Bloch and P. E. Crouch. On the equivalence of higher order variational problems and optimal control problems. In *Proceedings of the 35th IEEE Conference on Decision and Control*, pages 1648–1653, 1996.
- [7] A. M. Bloch, P. Krishnaprasad, J. E. Marsden, and R. M. Murray. Nonholonomic mechanical systems with symmetry. *Arch. Rat. Mech. Anal.*, 136:21–99, 1996.
- [8] A. M. Bloch, N. E. Leonard, and J. E. Marsden. Stabilization of mechanical systems using controlled Lagrangians. In *Proceedings of the 36th IEEE Conference on Decision and Control*, pages 2356–2361, 1997.
- [9] A. M. Bloch, N. E. Leonard, and J. E. Marsden. Matching and stabilization by the method of controlled Lagrangians. In *Proceedings of the 37th IEEE Conference on Decision and Control*, pages 1446–1451, 1998.

- [10] A. M. Bloch, N. E. Leonard, and J. E. Marsden. Controlled Lagrangians and the stabilization of mechanical systems I: the first matching theorem. *IEEE Transactions on Automatic Control*, 45(12):2253–2270, 2000.
- [11] R. W. Brockett. Control theory and singular Riemannian geometry. In P.J. Hilton and G.S. Young, editors, *New Directions in Applied Mathematics*. Springer-Verlag, 1981.
- [12] R. W. Brockett. Asymptotic stability and feedback stabilization. In R. W. Brockett, R.S. Millman, and H.J. Sussmann, editors, *Differential Geometric Control Theory*. Birkhauser, 1983.
- [13] A. E. Bryson and Y. Ho. *Applied Optimal Control*. Ginn and Company, 1969.
- [14] F. Bullo. *Nonlinear Control of Mechanical Systems: A Riemannian geometry approach*. Ph.D. thesis, California Institute of Technology, 1999.
- [15] F. Bullo. Series expansions for nonlinear analytic systems with additive controls. Submitted to *Automatica*, 2000.
- [16] F. Bullo. Series expansions for the evolution of mechanical systems. *SIAM Journal of Control and Optimization*, 40(1):166–190, 2001.
- [17] F. Bullo, N. Leonard, and A. D. Lewis. Controllability and motion algorithms for underactuated Lagrangian systems on Lie groups. *IEEE Transactions on Automatic Control*, 45(8):1437–1454, 2000.
- [18] R. Burns et al. TechSat21: Formation design, control, and simulation. In *Proceedings of the IEEE Aerospace Conference*, pages 19–25, 2000.
- [19] L. E. Buzogany, M. Pachter, and J. J. d’Azzo. Automated control of aircraft in formation flight. In *Proceedings of the AIAA Conference on Guidance, Navigation, and Control*, pages 1349–1370, 1993.
- [20] Y. U. Cao, A. S. Fukunaga, and A. B. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4:7–27, 1997.
- [21] D. F. Chichka. Satellite clusters with constant apparent distribution. *AIAA Journal of Guidance, Control, and Dynamics*, 24(1):117–122, 2001.
- [22] D. F. Chichka, J. L. Speyer, and C. G. Park. Peak-seeking control with application to formation flight. In *Proceedings of the 38th IEEE Conference on Decision and Control*, pages 2463–2470, 1999.

- [23] J. H. Chow and P. V. Kokotovic. Time scale modeling of sparse dynamic networks. *IEEE Transactions on Automatic Control*, 30(8):714–722, 1985.
- [24] F. R. K. Chung. *Spectral Graph Theory*, volume 92 of *Regional Conference Series in Mathematics*. American Mathematical Soc., 1997.
- [25] M. Chyba. Numerical method for sub-Riemannian geometry. In *Proceedings of the 38th IEEE Conference on Decision and Control*, pages 7–12, 1999.
- [26] W. H. Clohessy and R. S. Wiltshire. Terminal guidance system for satellite rendezvous. *Journal of the Aerospace Sciences*, 27(9):653–658, 1960.
- [27] J. Corfmat and A. Morse. Decentralized control of linear multivariate systems. *Automatica*, 12:476–495, 1976.
- [28] M. G. Crandall, L. C. Evans, and P. L. Lions. Some properties of viscosity solutions of Hamilton-Jacobi equations. *Transactions of the American Mathematical Society*, 282(2):487–502, 1984.
- [29] P. Crouch and F. Silva Leite. The dynamic interpolation problem: on Riemannian manifolds, Lie groups, and symmetric spaces. *Journal of Dynamical and Control Systems*, 1(2):177–202, 1995.
- [30] T. B. Curtin, J. G. Bellingham, J. Catipovic, and D. Webb. Autonomous oceanographic sampling networks. *Oceanography*, 6:86–94, 1993.
- [31] D. Cvetkovic, M. Doob, and H. Sachs. *Spectra of Graphs*, volume 87 of *Pure and Applied Mathematics*. Academic Press, Inc., 1980.
- [32] D. Cvetkovic, P. Rowlinson, and S. Simic. *Eigenspaces of Graphs*, volume 66 of *Encyclopedia of Mathematics and Applications*. Cambridge University Press, 1997.
- [33] J. Desai, J. P. Ostrowski, and V. Kumar. Modeling and control of formations of nonholonomic mobile robots. *IEEE Transactions on Robotics and Automations*, 17(6):905–908, 2001. submitted.
- [34] R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, 1997.
- [35] M. P. do Carmo. *Riemannian Geometry*. Birkhauser, Boston, 1992.

- [36] W. B. Dunbar. Model predictive control: extension to coordinated multi-vehicle formations and real-time implementation. Technical Report CDS 01-016, California Institute of Technology, 2001.
- [37] W. B. Dunbar, M. B. Milam, R. Franz, and R. M. Murray. Model predictive control of a thrust-vectoring flight control experiment. To appear, *Proceedings of the 15th IFAC World Congress*, 2002.
- [38] M. Egerstedt, X. Hu, and A. Stotsky. Control of mobile platforms using a virtual vehicle approach. *IEEE Transactions on Automatic Control*, 46(11):1777–1782, 2001.
- [39] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23:298–305, 1973.
- [40] M. Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its applications to graph theory. *Czechoslovak Mathematical Journal*, 25(100):619–633, 1975.
- [41] M. Fiedler and V. Ptak. On matrices with non-positive off-diagonal elements and positive principal minors. *Czechoslovak Mathematical Journal*, 12(87):382–400, 1962.
- [42] G. F. Franklin and J. D. Powell. *Digital Control of Dynamic Systems*. Addison-Wesley, 1980.
- [43] S. Friedland and R. Nabben. On Cheeger-type inequalities for weighted graphs. Preprint.
- [44] G. Froyland and M. Dellnitz. Detecting and locating near-optimal almost-invariant sets and cycles. Preprint, 2001.
- [45] C. E. Garcia, D. M. Prett, and M. Morari. Model predictive control: Theory and practice. a survey. *Automatica*, 25(3):335–348, 1989.
- [46] F. Giulletti, L. Pollini, and M. Innocenti. Autonomous formation flight. *IEEE Control Systems Magazine*, pages 34–44, December 2000.
- [47] J. Graver, J. Liu, C. Woolsey, and N. E. Leonard. Design and analysis of an underwater vehicle for controlled gliding. In *Proceedings of the 32nd Conference of Information Sciences and Systems*, 1998.

- [48] J. F. Heagy, T. L. Carroll, and L. M. Pecora. Synchronous chaos in coupled oscillator systems. *Physical Review E*, 50(3):1874–1885, 1994.
- [49] A. J. Healey. Application of formation control for multi-vehicle robotic minesweeping. Submitted to *Proceedings of the 40th IEEE Conference on Decision and Control*, 2001.
- [50] G. W. Hill. Researches in the lunar theory. *American Journal of Mathematics*, 1(1):5–26, 1878.
- [51] R. Horn and C. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- [52] A. Isidori. *Nonlinear Control Systems*. Springer-Verlag, New York, NY, third edition, 1995.
- [53] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46(5):776–783, 2001.
- [54] C. Kane, J. E. Marsden, and M. Ortiz. Symplectic energy–momentum integrators. *Journal of Mathematical Physics*, 40:3353–3371, 1999.
- [55] C. Kane, J. E. Marsden, M. Ortiz, and M. West. Variational integrators and the Newmark algorithm for conservative and dissipative mechanical systems. *International Journal for Numerical Methods in Engineering*, 49(10):1295–1325, 2000.
- [56] S. Kelly and R. M. Murray. Geometric phases and locomotion. *Journal of Robotic Systems*, 12(6):417–431, 1995.
- [57] H. Khalil. *Nonlinear Systems*. Macmillan Publishing Co., New York, NY, 1992.
- [58] S. Kobayashi and K. Nomizu. *Foundations of Differential Geometry, Vol I*. Interscience Publishers, New York, 1963.
- [59] E. M. C. Kong. Optimal trajectories and optimal design for separated spacecraft interferometry. Master’s thesis, Massachusetts Institute of Technology, 1999.
- [60] W.-S. Koon and J.E. Marsden. Optimal control for holonomic and nonholonomic mechanical systems with symmetry and Lagrangian reduction. *SIAM Journal on Control and Optimization*, 35(3):901–929, 1997.

- [61] K. Lau et al. An innovative deep space application of GPS technology for formation flying spacecraft. In *Proceedings of the AIAA Conference on Guidance, Navigation, and Control*, 1996. AIAA 96-3819.
- [62] N. E. Leonard. Stabilization of underwater vehicle dynamics with symmetry-breaking potentials. *Systems & Control Letters*, 32:35–42, 1997.
- [63] N. E. Leonard and E. Fiorelli. Virtual leaders, artificial potentials and coordinated control of groups. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 2968–2973, 2001.
- [64] A. D. Lewis. *Aspects of Geometric Mechanics and Control of Mechanical Systems*. Ph.D. thesis, California Institute of Technology, 1995.
- [65] A. D. Lewis. Affine connections and distributions with applications to non-holonomic mechanics. *Reports on Mathematical Physics*, 42(1/2):135–164, 1998.
- [66] A. D. Lewis. The geometry of optimal control for affine connection control systems. Preprint, 1999.
- [67] A. D. Lewis and R. Murray. Controllability of simple mechanical control systems. *SIAM Review*, 41(3):555–574, 1999.
- [68] B. Lincoln and B. Bernhardsson. Optimal control over networks with long random delays. In *Proceedings of 14th International Symposium on Mathematical Theory of Networks and Systems*, 2000.
- [69] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufman Publishers, 1996.
- [70] J. E. Marsden and T. S. Ratiu. *Introduction to Mechanics and Symmetry*, volume 17 of *Texts in Applied Mathematics*. Springer-Verlag, New York, NY, 1994.
- [71] T. W. McLain and R. W. Beard. Trajectory planning for coordinated rendezvous of unmanned air vehicles. In *Proceedings of the AIAA Conference on Guidance, Navigation, and Control*, 2000.
- [72] T. W. McLain, P. R. Chandler, and M. Pachter. A decomposition strategy for optimal coordination of unmanned air vehicles. In *Proceedings of the the American Control Conference*, pages 369–373, 2000.

- [73] R. Merris. Laplacian matrices of graphs: A survey. *Linear Algebra and its Applications*, 197,198:143–176, 1994.
- [74] R. Merris. A survey of graph Laplacians. *Linear and Multilinear Algebra*, 39:19–31, 1995.
- [75] M. Mesbahi and F. Hadaegh. Formation flying of multiple spacecraft via graphs, matrix inequalities, and switching. *AIAA Journal of Guidance, Control and Dynamics*, 24(2):369–377, 2001.
- [76] M. B. Milam, R. Franz, and R. M. Murray. Real-time constrained trajectory generation applied to a flight control experiment. To appear, *Proceedings of the 15th IFAC World Congress*, 2002.
- [77] M. B. Milam, K. Mushambi, and R. M. Murray. A computational approach to real-time trajectory generation for constrained mechanical systems. In *Proceedings of the 39th IEEE Conference on Decision and Control*, pages 845–851, 2000.
- [78] B. Mohar. Laplace eigenvalues of graphs — a survey. *Discrete Mathematics*, 109:171–183, 1992.
- [79] K. Morgansen. *Temporal Patterns in Learning and Control*. Ph.D. thesis, Harvard University, 1999.
- [80] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, Boca Raton, FL, 1994.
- [81] R. Nabben. Improved upper bounds for the real part of nonmaximal eigenvalues of nonnegative matrices. *SIAM Journal of Matrix Analysis and Applications*, 22(2):574–579, 2000.
- [82] H. Nijmeijer and A. van der Schaft. *Nonlinear Dynamical Control Systems*. Springer-Verlag, New York, NY, 1990.
- [83] L. Noakes, G. Heinzinger, and B. Paden. Cubic splines on curved spaces. *IMA J. Math. Control and Inf.*, 6(4):465–473, 1989.
- [84] L. M. Pecora and T. L. Carroll. Master stability functions for synchronized coupled systems. *Physics Review Letters*, 80(10):2109–2112, 1998.

- [85] J. Primbs, V. Nevistic, and J. C. Doyle. Nonlinear optimal control: a control lyapunov function and receding horizon perspective. *Asia Journal of Control*, 1(1):1–11, 1999.
- [86] J. A. Primbs. *Nonlinear Optimal Control: A Receding Horizon Approach*. Ph.D. thesis, California Institute of Technology, Pasadena CA, 1999.
- [87] U. G. Rothblum. Algebraic eigenspaces of nonnegative matrices. *Linear Algebra and its Applications*, 12:281–292, 1975.
- [88] H. Schaub et al. Spacecraft formation flying control using mean orbital elements. *Journal of the Astronautical Sciences*, 48(1):69–87, 2000.
- [89] R. Serban and L. R. Petzold. COOPT — a software package for optimal control of large-scale differential-algebraic equation systems. *Journal of Mathematics and Computers in Simulation*, 56(2):187–203, 2001.
- [90] S. E. Shladover et al. Automatic vehicle control developments in the PATH program. *IEEE Transactions on Vehicular Technology*, 40(1):114–130, 1991.
- [91] D. D. Siljak. *Decentralized Control of Complex Systems*, volume 184 of *Mathematics in Science and Engineering*. Academic Press, Boston, MA, 1990.
- [92] T. R. Smith, H. Hanßmann, and N. E. Leonard. Orientation control of multiple underwater vehicles. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 4598–4603, 2001.
- [93] R. F. Stengel. *Stochastic Optimal Control*. John Wiley & Sons, 1986.
- [94] D. J. Stilwell and B. E. Bishop. Platoons of underwater vehicles. *IEEE Control Systems Magazine*, pages 45–52, December 2000.
- [95] D. Swaroop and J. K. Hedrick. String stability of interconnected systems. *IEEE Transactions on Automatic Control*, 41(3):349–356, 1996.
- [96] D. Swaroop and J. K. Hedrick. Constant spacing strategies for platooning in automated highway systems. *ASME Journal of Dynamic Systems, Measurement and Control*, 121:462–470, 1999.
- [97] P. Tabuada, G. J. Pappas, and P. Lima. Feasible formations of multi-agent systems. In *Proceedings of the American Control Conference*, pages 56–61, 2001.

- [98] R. S. Varga. *Matrix Iterative Analysis*, volume 27 of *Springer Series in Computational Mathematics*. Springer-Verlag, 2 edition, 1991.
- [99] G. C. Walsh, H. Ye, and L. Bushnell. Stability analysis of networked control systems. In *Proceedings of the American Control Conference*, pages 2876–2880, 1999.
- [100] S. Wang and E. Davison. On the stabilization of decentralized control systems. *IEEE Transactions on Automatic Control*, AC-18:473–478, 1973.
- [101] J. M. Wendlandt and J. E. Marsden. Mechanical integrators derived from a discrete variational principle. *Physica D*, 106:223–246, 1997.
- [102] B. Wittenmark, B. Bastian, and J. Nilsson. Analysis of time delays in synchronous and asynchronous control loops. In *Proceedings of 37th IEEE Conference on Decision and Control*, pages 283–288, 1998.
- [103] J. D. Wolfe, D. F. Chichka, and J. L. Speyer. Decentralized controllers for unmanned aerial vehicle formation flight. In *Proceedings of AIAA Conference on Guidance, Navigation, and Control*, 1996. AIAA Paper 96-3833.
- [104] C. W. Wu. Synchronization in arrays of coupled nonlinear systems: passivity, circle criterion, and observer design. *IEEE Transaction on Circuits and Systems I—Fundamental Theory and Applications*, 48(10):1257–1261, 2001.
- [105] C. W. Wu and L. O. Chua. Synchronization in an array of linearly coupled dynamical systems. *IEEE Transactions on Circuits and Systems — I. Fundamental Theory and Applications*, 42(8):430–447, 1995.
- [106] D. Yanakiev and I. Kanellakopoulos. A simplified framework for string stability analysis in AHS. In *Proceedings of 13th IFAC World Congress*, volume Q, pages 177–182, San Francisco, CA, 1996.
- [107] H.-H. Yeh, E. Nelson, and A. Sparks. Nonlinear tracking control for satellite formations. In *Proceedings of the 39th IEEE Conference on Decision and Control*, pages 328–333, 2000.
- [108] H.-H. Yeh and A. Sparks. Geometry and control of satellite formations. In *Proceedings of the American Control Conference*, pages 384–388, June 2000.
- [109] J. Yuh. Developments in underwater robotics. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1862–1867, 1995.

- [110] J. Yuh. Underwater robotics. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 932–937, 2000.
- [111] K. Zhou and J. C. Doyle. *Essentials of Robust Control*. Prentice Hall, New Jersey, 1998.