Hierarchies, Spikes, and Hybrid Systems: Physiologically Inspired Control Problems

Thesis by

Andrew Lamperski

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2011

(Defended May 27, 2011)

© 2011
Andrew Lamperski
All Rights Reserved

To my parents, Curtis and Gloria Lamperski.

Acknowledgements

First of all, I would like to thank my advisor, John Doyle. His unrelenting zeal for knowledge and quirky approach to science have been infectious. Through John's mentorship and broad perspective, I learned to approach fundamental problems with confidence. I would also like to extend my thanks to the other members of my committee: Sanjay Lall who generously set aside time to meet with me during my visits to Stanford and gave me important insights into decentralized control, Pietro Perona who provided a much needed neuroscience perspective to thesis discussions and quickly found the boundaries of my theories, and Richard Murray who provided a wealth of insight, both technical and practical, throughout my years at Caltech and kept me and the rest of the committee in line.

I want to express my gratitude to my undergraduate advisor, Noah Cowan, who got me so pumped about robots and control that I had no choice but to do a PhD. From my first real research experiences, to now, he has regularly provided interesting discussion and excellent advice.

Credit goes out to Aaron Ames, my first productive collaborator at Caltech. It was with Aaron that I cut my teeth on theoretical research, and I will always be indebted to him for helping me get the ball rolling.

I am grateful to Genti Buzi and Dennice Gayme for carefully reading the half-baked drafts I sent to them, listening to my latest ideas and darkest laments, and above all, being good friends. Vanessa Jönsson and Jerry Cruz, whether it was discussing work, running, or playing music, you have made my time at Caltech more enjoyable. Ufuk Topcu and Nader Motee, thank you for the stimulating conversations and for giving me a glimpse of life after the PhD.

I want to acknowledge all the members of the Doyle group, both past and present, for their regular feedback and support. I am grateful to the rest of my CDS comrades for creating a positive and productive atmosphere. Anissa Scott, Gloria Bain, and the rest of the CDS administrative staff, thank you for making my time go smoothly.

I am grateful to my parents, Curtis and Gloria Lamperski, for their support and encouragement. Much appreciation goes to my sister, Minda Apel, brother-in-law, Chris Apel, and their children, Christian, Ryan, Julianna, and Kelsea, for making every visit to Pittsburgh fun.

Finally, I want to thank my fiancé, Rica Enriquez. Throughout the course of my PhD, she provided a reasonably steady stream of the best cookies on the planet and a strong incentive to

vacuum my floor. More importantly, she has been the ultimate supporter and friend, in good times and in bad. The best I can do is be grateful and strive to be at least a fraction as caring and supportive for her as she has always been for me.

Abstract

In animal motor control and locomotion, neurons process information, muscles are the actuators, and the body is the plant. Control theory is the standard mathematical language for describing motor control and locomotion, but many phenomena in physiological control remain outside of the scope of control theoretic reasoning. Unlike traditional engineering control systems, nearly all the components of physiological control systems have complex dynamics. Instead of a fast centralized computer, an animal implements controllers using a distributed network of slow, nonlinear, and noisy neurons. Rather than having linear plants and actuators, the animal must control limbs with nonlinear and hybrid dynamics.

This dissertation develops basic control theory motivated by physiological systems. Dynamical phenomena that arise in physiology but remain outside the scope of mathematical methods are isolated and studied in general control theoretic frameworks. In particular, three problems are discussed: distributed linear quadratic Gaussian (LQG) control with communication delays, control over communication channels modeled after spiking neurons, and Zeno stability of hybrid systems.

Motivated by the presence of delays in the human motor system, Chapter 2 explores the architecture of distributed LQG controllers when communication between subsystems is limited by delays. Sensory and motor command information is processed in several different regions throughout the nervous system. Since processing speed in neurons is limited, information from different sensory and motor regions can only be integrated after a time delay. In spite of this difficulty, humans make efficient and reliable motions that are well-described by optimal feedback control. Optimal delay compensation is studied in a distributed LQG framework. The structure that emerges as the result of optimization resembles a management hierarchy, bearing similarities with the organization of the motor system.

Networked control systems, in which communication between the controller and the plant occurs over a special neuron-inspired channel, are analyzed in Chapter 3. In addition to being the basic computing elements, neurons are the long-range communication channels of the body. Neurons transmit information in the form of short-lived voltage spikes, called action potentials. Sufficient conditions for stable control over the spiking channel are presented, along with bounds on tracking error and data rates.

The final technical chapter studies the connections between Zeno behavior and Lyapunov stability. Zeno behavior occurs in a hybrid system when an infinite number of discrete transitions occurs in a finite amount of time. While Zeno behavior results from modeling abstractions, it is commonly observed in models of mechanical systems undergoing impacts, including models important for locomotion. Often, Zeno behavior is associated with dynamical mode transitions, such as knee locking and the transition between bouncing and sliding. To reason about such transitions without modifying the models, the chapter on hybrid systems gives Lyapunov-like sufficient conditions for Zeno behavior. A technique for constructing the Lyapunov-like certificates is presented for a general class of mechanical systems undergoing impacts.

Contents

List of Figures

Chapter 1

Introduction

To make progress in any field, easy problems must be mastered first. Traditional control systems are typically actuated with components such as motors and jet engines that have simple input-output behavior. Likewise, vehicles such as cars and airplanes are designed to follow simple trajectories. Including more complex components or requiring elaborate behavior increases the difficulty of design, construction, and testing. If the components and behavior of a system become too complex to accurately characterize, undesirable behavior could result. If an airplane is designed to safely transport people from city to city, emphasis should be placed on reliable, well-characterized behavior. In this case, there is little reason to consider more complex components, such as flapping wings.

Just as the physical components of engineering control systems are typically designed for simplicity and reliability, so are the computational and communication structures. While control theory has roots that predate computers, digital implementations of control laws quickly became standard. In the traditional application domains for control, such as the military and aerospace industries, digital control confers many benefits. Computers can be easily programmed and reprogrammed, eliminating the need to modify circuitry each time the control laws are changed. Digital computation is also more reliable than analog circuitry. Compared to the cost of keeping a jet in the air, the computational energy required for implementing control laws is tiny. Thus, for many of the driving applications of control, there is little incentive to keep computational costs low, or to consider computing architectures other than high-speed digital computers. Similarly, communication within common control systems, like vehicles, occurs via high-speed, noise-resistant wiring. Again, because of the comparatively low cost of good wires and the high cost of failure, there is little incentive to explore low cost alternatives to wires that may be slower and less reliable.

In contrast, animal motor control and locomotion are characterized by numerous hardware constraints that are absent from traditional control applications. Despite hundreds of millions of years of evolution, producing animals with wide variations in features and form, animals have not grown wheels, propellers, or jet engines. To move about, animals walk on legs, swim with fins, and fly by flapping wings. In contrast to wheels, legs have complex nonlinear and hybrid dynamics. Similarly,

animals process information with brains composed of neurons, as opposed to digital computers. In contrast to digital computers used for control, neurons are slow, noisy, and highly parallelized.

Industrial applications have been instrumental in the development of control theory. Without simple application problems and cheap digital computation, it is unlikely that control theory would be as sophisticated and successful as it is today. Unfortunately, many physiological phenomena are too still complex for tractable control theoretic analysis.

1.1 Objective and Approach

This thesis develops results in basic control theory, motivated by physiological systems. Much more research has been devoted to applying control theory to physiological systems than has been devoted to developing new control theoretic results from physiological considerations (Figure 1.1). This situation fits with the view that scientific information flows linearly from more general fields to less general (chemistry is applied physics, biology is applied chemistry, etc.). Of course, scientific development is not so simple, and a theory is only as good as the examples it describes. As described above, control theory developed in parallel with its engineering applications, which have been designed to avoid many of the inherent difficulties of physiology. To increase the scope of control theoretic reasoning, this dissertation advocates the development of basic control theory to explain physiological phenomena that are typically overlooked in engineering. While physiologically inspired control theory is not a new idea (see, for example [4, 5]), a large number of interesting problems remain to be explored.

Each of the main chapters focuses on a single physiological phenomenon and studies its theoretical implications in isolation. The phenomena are isolated for theoretical tractability, at the expense of biological realism. It is hoped that understanding dynamic phenomena in isolation can lend insight into the behavior of physiological systems.

1.2 Problems and Contributions

The results in this thesis are divided into three, essentially independent, technical chapters, which can be read in any order. Chapters 2 and 3 study feedback control problems motivated by neuroscience. Chapter 4 studies a hybrid system phenomenon, known as Zeno behavior, motivated by locomotion problems.

1.2.1 Distributed Optimal Control and the Motor System

Imagine a baseball player running to catch a fly ball. Motor commands to the arms and legs must be coordinated with visual data about the ball's trajectory. Throughout the process, information

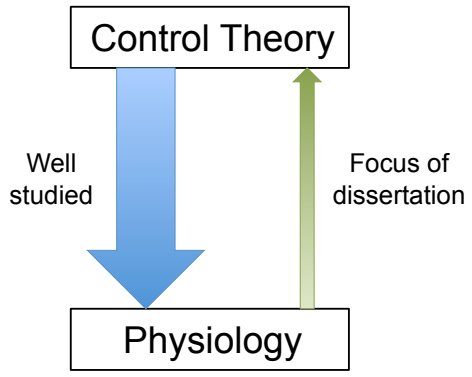

Figure 1.1: A great deal of work applies control theoretic methods to physiological problems. Considerably less work has been done to derive new control theory to explain physiology. As a result, physiological phenomena that do not resemble traditional control theoretic models are difficult to analyze. This dissertation presents basic results in physiologically motivated control theory.

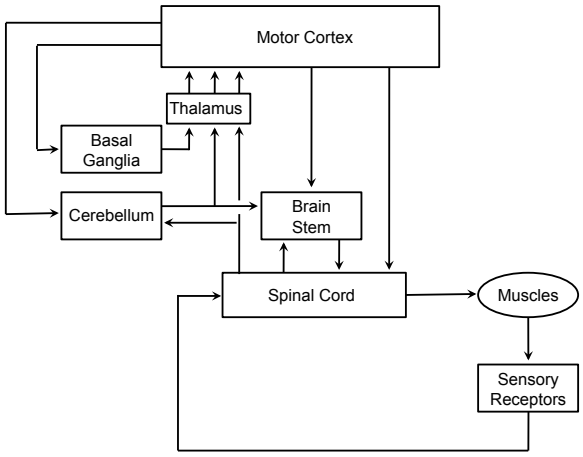

Figure 1.2: The motor system is a hierarchical control system with three main levels, the spinal cord, the brain stem, and the motor cortex. The cerebellum, basal ganglia, and thalmamus are brain regions providing feedback between the levels. Sensory information is integrated with motor commands at every level of the hierarchy.

from stretch receptors in the muscles must be combined with data from the inner ear to keep the player from falling. The information is processed by a hierarchical neural controller called the motor system (Figure 1.2[1]). All of the subsystems of the motor hierarchy, and all the connections between them are composed of neurons. Neurons process and transmit information rather slowly, and thus there are noticeable delays between the various subsystems of the motor system.

To gain a better understanding of how the motor system produces reliable and precise movements in spite of delays, Chapter 2 examines the structure of optimal distributed controllers for systems with communication delays. In the particular problem studied, a group of players works together to solve a linear quadratic Gaussian (LQG) control problem. Each player can directly measure a part of the state, but they can only communicate their state information with delays.

Contributions. While the problem studied in Chapter 2 has been solved by semidefinite programming [6], the structure of the optimal controllers is not apparent from the solution. The main contribution of Chapter 2 is the derivation of the explicit structures of the optimal distributed controllers, through a novel dynamic programming argument. A graphical structure, reminiscent of a management hierarchy, emerges from the optimization (Figure 1.3). At the lowest level, players have immediate access to local state information. After a delay the state information of neighboring workers is integrated at the next level up by a manager. At the top level, an executive integrates all of the information from the managers, after another delay. For control, the executive applies the standard linear quadratic regulator gain to the delayed state. The managers and workers then apply corrections to the input based on their newer state information. A computationally efficient algorithm to construct the hierarchical optimal controller is given for a broad class of systems with delays.

Though more work is needed to strengthen the connection between distributed LQG and the motor system, the structural results are promising. Indeed, there is already a basic resemblance in the organization: a hierarchical control scheme with sensory data integrated at all levels, with varying amounts of delay. In order to further the connection, observations from motor control experiments and anatomical knowledge must be incorporated into the theory developed in Chapter 2.

1.2.2 Spiking Neurons in Feedback Loops

In addition to being the computing elements, neurons are also the long range communication channels of the body. The main components of the neuron are sketched in Figure 1.4[2]. Input currents at

[1] Reproduced with permission of The McGraw-Hill Companies from Kandel et al, Principles of Neural Science, 3rd Edition ©1992, Elsevier [7].

[2] This picture was modified from http://commons.wikimedia.org/wiki/File:Neuron.svg, which is an image from the Public Domain work "Anatomy and Physiology" by the US National Cancer Institute's Surveillance, Epidemiology and End Results (SEER) Program., redrawn by User Dhp1080. It is available under the Creative Commons Attribution-

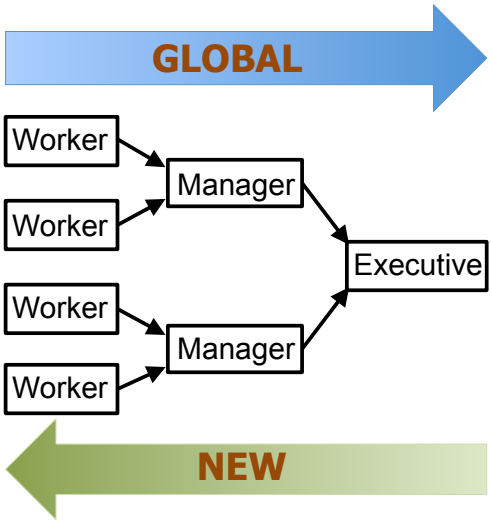

Figure 1.3: Schematic interpretation of the optimal controllers from Chapter 2. At the leftmost nodes of the graph, the workers have exact, up-to-date, knowledge of a portion of the state of the system. They do not, however, have access to information known to the other workers. The workers then communicate their states with a delay to the managers at the next level up. In this example, each manager knows the states of two of workers, but its information is one time step old. Finally, the managers give their knowledge of the state to the top-level executive. So, the executive knows the state of the whole system is, at a delay of two time steps. The optimal control policy is such that the executive applies a centralized control law to its old, global data, while the managers and workers use the newer information to reduce errors caused by the delays.

the dendrites increase or decrease the voltage across the cell membrane. When the voltage across the membrane reaches a threshold value, it rapidly increases, producing a spike, called an action potential (Figure 1.5). The action potential travels across the axon, producing outputs at the axon terminals.[3] All of the information sent through a neuron is encoded in the rate and timing of the spikes [7, 8].

Chapter 3 studies feedback control over communication channels modeled after spiking neurons. Communication in control systems, such as the connections between the plant and the controller, is typically assumed to occur instantaneously over wires. Even if a communication channel is explicitly modeled in a control system, it is typically assumed that the channel transmits packets of numerical data at periodically sampled time intervals. In contrast, neural communication employs no periodic sampling, and the individual spikes carry no numerical information. The numerical information must be encoded in spike rates and spike timing.

Contributions. The main contribution of Chapter 3 is the stability analysis for networked control systems in which the plant and the controller communicate via a spiking mechanism. Bounds on tracking error and spike rates are also derived.

Share Alike 3.0 Unported license.

[3] This description is highly simplified. See [7] for a more precise discussion.

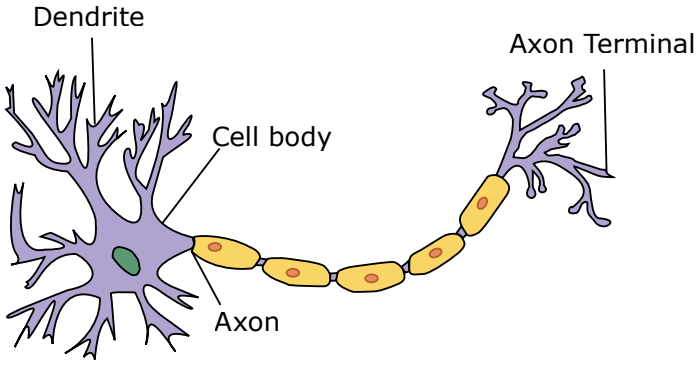

Figure 1.4: A basic neuron. The neuron takes inputs at the dendrites, sends signals down the axon, and produces outputs at the axon terminals.

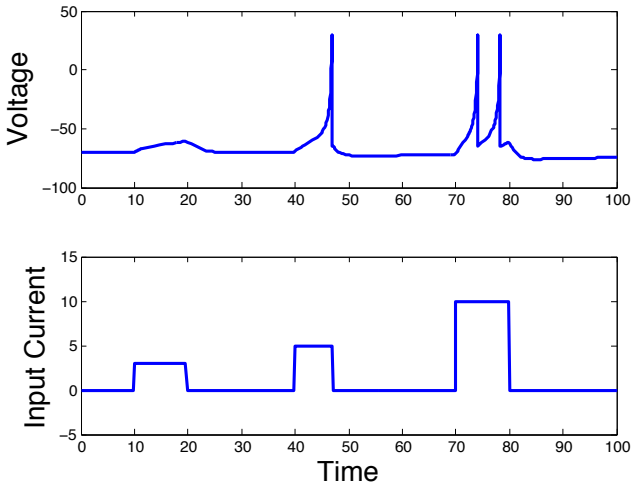

Figure 1.5: A simulation of the membrane potential (the voltage across the cell membrane) in response to current steps. Note that for small currents, no spikes occur. For sufficiently large currents, spikes can occur. As the current increases the rate of spiking increases.

The relationship between networked control and neuroscience is natural, but largely unexplored. Indeed, a great deal of neurons function as communication channels in feedback loops. The connection can be studied further by relating the coding strategies in neural control systems, such as the motor system, to engineering approaches to networked control.

1.2.3 Zeno Behavior and Lyapunov Stability

Locomotion behaviors are best modeled by nonlinear hybrid systems. Human limbs move by rotating bones around joints, resulting in inherently nonlinear dynamics. Furthermore, in locomotion, the limbs undergo impacts and the body moves through several different dynamic modes (such as having both feet down, or one foot down and the other swinging).

Motivated by problems in locomotion, bipedal walking, in particular, Chapter 4 studies a phe-

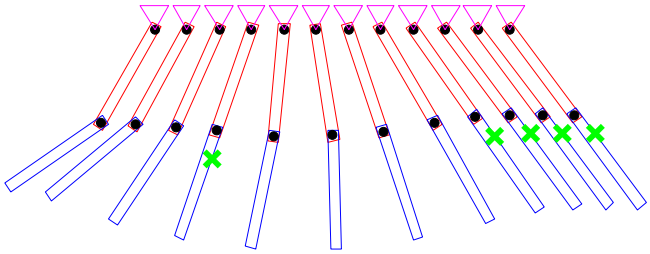

Figure 1.6: A simple knee joint model is given by a double pendulum with a mechanical stop. When the lower leg reaches the stop it rebounds based on Newtonian impact equations. Impacts are denoted by an "x" on the figure. Zeno behavior is observed as the knee converges to a locked position (on the right).

nomenon of hybrid systems known as Zeno behavior. Zeno behavior arises when an infinite number of discrete transitions occur in a finite amount of time. Many models of mechanical systems undergoing impacts, such as the bouncing ball and simple knee joint models (Figure 1.6[4]), display Zeno behavior. While Zeno behavior is typically the result of modeling abstractions, and does not occur in reality, it is often associated with transitions between dynamical modes in physical systems. For instance, in bouncing ball models, Zeno behavior occurs at the transition between bouncing and sliding, while in the knee joint model, Zeno behavior occurs as the knee converges to a locked position. Perhaps less obviously, Zeno behavior is closely related to Lyapunov stability. In the one-dimensional bouncing ball, all trajectories exhibit Zeno behavior as they converge to a resting position. More interestingly, all trajectories starting sufficiently close to certain locked configurations of the knee joint model of Figure 1.6 exhibit Zeno behavior and converge to nearby locked configurations.

Contributions. The main contribution of Chapter 4 is the development of the connection between Zeno behavior and Lyapunov stability, and in particular, Lyapunov-like sufficient conditions for Zeno behavior. A general method for constructing the Lyapunov-like certificates is given for a class of models for mechanical systems undergoing impacts. That construction is then used to derive algebraic sufficient conditions for Zeno behavior in the mechanical models.

The work in Chapter 4 restricts the class of hybrid systems studied in order to isolate the mechanisms causing Zeno behavior. To extend the range of applicability of the work in Chapter 4, these restrictions must be weakened or removed. Work to this end is already under way [9, 10], and the results in Chapter 4 have since been applied in bipedal walking applications [11].

[4]Generated using code from Yizhar Or.

Chapter 2

Distributed LQG with Applications to Management and Human Motor Control

2.1 Introduction

Imagine a large event, such as a conference, is being planned. A group of people must move chairs and tables, prepare and serve food, and set up audio-visual equipment. In addition, typically, there are also people who do little or none of the physical work, but whose main function is coordination. Someone oversees the food, while someone else might coordinate the audio-visual equipment. At the top level, there is often an individual, or group, that manages the coordinators. Management hierarchies, such as the one just described, are common, even though they are not actually necessary to perform the actions needed for the event.

On the other hand, consider one of the most familiar motor coordination tasks: walking. Sensory data from the eyes, inner ear, and muscles gets integrated with motor commands to produce robust, efficient locomotion. Of course, all communication occurs over neurons, which are rather slow. Given constraints on communication speed, what optimal strategy arises?

This chapter studies how hierarchical control structures can arise as optimal methods to deal with communication delays. While the problem studied is a simple variant of distributed linear quadratic Gaussian (LQG) control, the results have some intuitive similarities with the event planning and walking examples discussed above. In the problems studied, a group of players works together to minimize a quadratic cost. The players have access to local state information, but can only communicate their state with a delay. This chapter shows that for such problems, a hierarchical control structure emerges as the optimal strategy. In particular, the optimal controller can be decomposed based on a hierarchical graph structure in which the lowest level nodes represent the players (which do the physical work), while higher level nodes are used for coordination (Figure 1.3).

2.1.1 Related Work

The focus of this chapter, LQG control with communication delays, is a basic problem in distributed or decentralized control. Decentralized control has a long history, [12, 13, 14, 15, 16], but computationally tractable solutions to nontrivial problems have been rare, until recently. Notably, in the past ten years, certain decentralized optimal control problems were shown to be convex [17, 18]. More recently, computationally efficient solutions to some of the convex problems have been found [6, 19, 20, 21, 22]. Of the work cited, Rantzer's paper on linear quadratic teams is the most closely related [6]. In this paper, Rantzer solves the problem studied in this chapter using semidefinite programming (SDP), but does not explore the structure of the solution. The solution techniques used in this chapter are closely related to the dynamic programming methods used in [21]. Also related are the works of [14, 15, 16], which give solutions to the problem of Section 2.3. The work in Section 2.3 differs from these works, in that it naturally leads to generalizations.

2.1.2 Motivation

The main motivation for studying the present problem is neural motor control. Humans effortlessly execute motor control tasks such as walking, tool manipulation, and speech. Other animals can fly, run, and swim with greater efficiency, precision, and reliability than the best engineered systems. Furthermore, the data processing is carried out by a massively distributed network of noisy, and rather slow neurons. Motor control research aims at discovering the strategies employed by humans and other animals to execute control tasks. While animals are not expected to be using the exact control strategies from engineering, interpreting experimentally observed behavior in terms of established control theory has been fruitful.

In the past few decades, ideas from optimal control [23, 24, 25, 26, 27, 28] and estimation [29, 30, 31, 32] have influenced theories about neural motor control systems. Optimal control has been proposed to explain why humans and animals execute the particular trajectories that are observed. Ideas from estimation are used to describe how the brain integrates sensory data with motor commands.

While the nervous system is a distributed system, aside from a few works such as [33, 34, 35], most studies of feedback in motor control focus on centralized controllers. Luckily, the confluence of theory and experiment have created a good environment to explore decentralized controllers in neural systems. In the past decade, control theoretic ideas have been applied to motor control with increasing sophistication. (See, for example [25, 27, 28, 30].) Centralized controllers provide good explanations for much of the motor behavior observed in humans and animals, but in order to reason about information processing across the motor control hierarchy (see Figure 1.2), decentralized control must be used. As discussed in the previous subsection, recent theoretical advances show

that optimal decentralized controllers can be efficiently computed for some nontrivial systems.

This chapter attempts to drive the theory of distributed and decentralized control towards motor control applications. In particular, this chapter studies optimal control with communication delays, with the hope that the resulting control structures might lend intuition to the structures used to compensate for delays in neural control systems.

2.1.3 Contributions

Given the focus on intuition, the primary contribution of this chapter is the explicit structure of the optimal controllers found. Even though a computationally efficient solution for the problem is already known from [6], no hierarchical structure is evident from the corresponding SDPs. Using a novel derivation, this chapter shows that by simply assuming that communication between players is delayed, a control hierarchy arises as the optimal solution. The control structure has a simple management interpretation, but also gives intuition about hierarchies in neuroscience. Furthermore, a computationally efficient method for constructing the the hierarchical controller is provided for a general class of delay structures.

While not emphasized in the chapter, the method for computing the optimal controllers from this chapter appears to be more efficient than the method from [6]. The computational savings arise from the fact that the only numerical optimization required is the solution of a single Riccati equation that depends on the total state dimension but not the number of players. The method from [6], on the other hand, relies on an SDP that grows both in state dimension and in the number of players.

2.1.4 Overview

The chapter is structured as follows. Section 2.2 defines the general problem studied. Section 2.3 gives a full solution for a special case termed the two-player problem. This section develops basic ideas on decoupling information into independent terms, and the associated dynamic programming problem. Comparisons to centralized control and a numerical example of pole balancing are also given in Section 2.3. Next, Section 2.4 presents a solution to the general problem of the chapter, based on ideas from Section 2.3. To solve this problem, a systematic method for decoupling information based on how it is shared by the players is introduced. Conclusions and future work are outlined in Section 2.5. The basic results and definitions for partially nested systems are given in the appendix to this chapter, Section 2.6. This section is included to justify the assumption that the controllers of this chapter are linear.

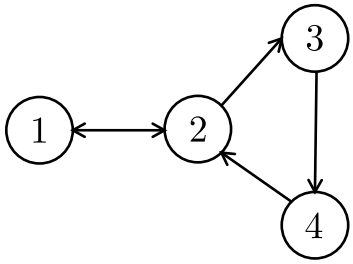

Figure 2.1: A delay structure graph with four nodes. Each edge corresponds to a single step delay. So, one time-step is required for information to travel between nodes 1 and 2. Two time-steps are needed for information to travel from node 1 to node 3, and so on. Associated to each node i is a player which chooses an input u_i.

2.2 Problem Statement

Notation. The expected value of a random variable, x, is denoted by $\mathbb{E}[x]$. The conditional expectation of x given y is denoted by $\mathbb{E}[x|y]$. Let $x(0:t)$ denote the stacked sequence of vectors:

$$x(0:t) = \begin{bmatrix} x(0) \\ x(1) \\ \vdots \\ x(t) \end{bmatrix}.$$

Consider a directed graph $G = (V, E)$ with $|V| = n$, called a *delay structure graph*. Throughout this section, the graph in Figure 2.1 will be used as an example. It is assumed that one time-step is required for any piece of information to travel across an edge in the delay graph. Thus, if the shortest path from node i to node j has length d, then d time-steps are required for information to flow from node i to node j.

Associate a state vector $x_i \in \mathbb{R}^{k_i}$, an input vector $u_i \in \mathbb{R}^{p_i}$, and a process noise vector $w_i \in \mathbb{R}^{k_i}$ to each node in $i \in V$. The state vector is updated according to the following discrete-time dynamic equations:

$$x_i(t+1) = A_{ii}x_i(t) + \sum_{\{j:(j,i)\in E\}} A_{ij}x_j(t) + B_{ii}u_i(t) + w_i(t). \tag{2.1}$$

In Equation (2.1), A_{ij} and B_{ii} are matrices of appropriate dimension.

For all $(i,j) \notin E$, let A_{ij} be the zero matrix of dimension $k_i \times k_j$. Then define the matrices A and B by

$$A = \begin{bmatrix} A_{11} & \cdots & A_{1n} \\ \vdots & \ddots & \vdots \\ A_{n1} & \cdots & A_{nn} \end{bmatrix}, \quad B = \begin{bmatrix} B_{11} & & \\ & \ddots & \\ & & B_{nn} \end{bmatrix}.$$

By stacking x_i, u_i, and w_i into larger vectors,

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}, \quad u = \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix}, \quad w = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix},$$

Equation (2.1) can be written in the more compact form,

$$x(t+1) = Ax(t) + Bu(t) + w(t). \tag{2.2}$$

For the graph in Figure 2.1, A has the structure

$$A = \begin{bmatrix} A_{11} & A_{12} & 0 & 0 \\ A_{21} & A_{22} & 0 & A_{24} \\ 0 & A_{32} & A_{33} & 0 \\ 0 & 0 & A_{43} & A_{44} \end{bmatrix}.$$

To see how information flows around the graph based on the structure of A, consider a vector w with sparsity pattern given by

$$w = \begin{bmatrix} 0 \\ 0 \\ * \\ 0 \end{bmatrix}.$$

The $*$ is used to indicate that the particular value of w is not important.

$$w = \begin{bmatrix} 0 \\ 0 \\ * \\ 0 \end{bmatrix}, \quad Aw = \begin{bmatrix} 0 \\ 0 \\ * \\ * \end{bmatrix}, \quad A^2w = \begin{bmatrix} 0 \\ * \\ * \\ * \end{bmatrix}, \quad A^3w = \begin{bmatrix} * \\ * \\ * \\ * \end{bmatrix}.$$

Returning to the general case, it will be assumed that $x(0) = 0$ and (A, B) is stabilizable. The process noise is Gaussian white noise, with terms corresponding to different nodes assumed to be uncorrelated: $\mathbb{E}[w_i w_j^T] = 0$, when $i \neq j$. So the covariance of the noise, w, is given by

$$\mathbb{E}[ww^T] = \begin{bmatrix} W_1 & & \\ & \ddots & \\ & & W_n \end{bmatrix}.$$

Assume that the graph G is strongly connected. Let $d_{ii} = 0$, and let d_{ij} be the length of the shortest path from node i to node j. The control problem is to minimize

$$\lim_{t\to\infty} \mathbb{E}\left[x(t)^T Q x(t) + u(t)^T R u(t)\right] \tag{2.3}$$

with inputs of the form

$$u_i(t) = \gamma_{i,t}(x_1(0:t-d_{1i}), \ldots, x_n(0:t-d_{ni})), \tag{2.4}$$

where $\gamma_{i,t}$ are Borel-measurable functions to be chosen in the optimization procedure.

For the graph in Figure 2.1, the constraints on the input are given by

$$\begin{aligned}
u_1(t) &= \gamma_{1,t}(x_1(0:t), x_2(0:t-1), x_3(0:t-3), x_4(0:t-2)) \\
u_2(t) &= \gamma_{2,t}(x_1(0:t-1), x_2(0:t), x_3(0:t-2), x_4(0:t-1)) \\
u_3(t) &= \gamma_{3,t}(x_1(0:t-2), x_2(0:t-1), x_3(0:t), x_4(0:t-2)) \\
u_4(t) &= \gamma_{4,t}(x_1(0:t-3), x_2(0:t-2), x_3(0:t-1), x_4(0:t)).
\end{aligned}$$

The weight matrices Q and R are assumed to be partitioned into blocks, $Q = (Q_{ij})_{i,j\in V}$ and $R = (R_{ij})_{i,j\in V}$, conforming to the partitions of x and u, respectively. The matrix Q is positive semidefinite, and R is positive definite. To guarantee a stabilizing solution to the corresponding algebraic Riccati equation, $(\sqrt{Q}, A)$ will be assumed to be detectable. Aside from that, no other assumptions are made about Q and R.

In the deriving the optimal controller, the following finite-horizon variant of the control problem is studied. Minimize

$$\mathbb{E}\left[\sum_{t=0}^{N-1} \left(x(t)^T Q x(t) + u(t)^T R u(t)\right) + x(N)^T \Lambda x(N)\right] \tag{2.5}$$

with inputs of the form of Equation (2.4). Here Λ is a positive semidefinite matrix of appropriate dimensions, corresponding to a terminal cost. If Λ is positive definite, then as $N \to \infty$, the optimal controller for this finite-horizon problem approaches the steady-state controller.

Note that the assumptions about the structure of input and the sparsity structures of A and B guarantee that communication between the players choosing u_i occurs as least as fast as information travels through the plant. As explained in the appendix to this chapter (Section 2.6), this assumption implies that the information structure (the set of input constraints) is *partially nested*, which in turn implies that optimal inputs are linear in the associated information. In [6], the condition on the communication delays is described by saying that there is no incentive to signal through the plant.

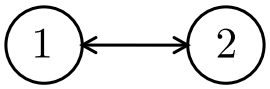

Figure 2.2: Two-player graph

The assumption is also closely related to the quadratic invariance condition of [18].

2.3 Two-Player Problem

In order to demonstrate the main ideas of the problem solution without the notational details, this section studies the simplest nontrivial instance, referred to as the *two-player problem*.

The two-player problem is the special case of the problem in this chapter defined by the graph in Figure 2.2. For the two-player problem, the state matrices have sparsity structure

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \quad B = \begin{bmatrix} B_{11} & 0 \\ 0 & B_{22} \end{bmatrix},$$

and inputs are restricted to the form

$$\begin{aligned} u_1(t) &= \gamma_{1,t}(x_1(0:t), x_2(0:t-1)) \\ u_2(t) &= \gamma_{2,t}(x_1(0:t-1), x_2(0:t)). \end{aligned} \tag{2.6}$$

These constraints on the inputs are often referred to in the literature as the "one-step delay information pattern."

The process noise $w(t)$ is Gaussian white noise with covariance given by

$$\mathbb{E}[ww^T] = \mathbb{E}\left[\begin{bmatrix} w_1 \\ w_2 \end{bmatrix}\begin{bmatrix} w_1^T & w_2^T \end{bmatrix}\right] = \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix}.$$

Dynamic programming solutions for the two-player problem have been known since the 1970s [14, 15, 16]. The approaches in the cited work are all rather direct and it is not immediately clear how to generalize them to other delay structures. The method of this section is to decompose the information into independent components, a priori, and use this decomposition to decouple the dynamic programming problem into independent subproblems. Later, in Section 2.4, it will be shown how this information decoupling method extends naturally to the general problem of this chapter.

The section is organized as follows. First the optimal controller is presented in Subsection 2.3.1. Subsection 2.3.2 derives the optimal solution to a finite-horizon version of the optimal control problem. Next,the optimal controller is derived in Subsection 2.3.3 by applying a limiting argument. To place the results of this section in the context of classical results, the properties of the two-player

controller are compared with centralized controllers in Subsection 2.3.4. Finally, to recap the ideas in a concrete setting, Subsection 2.3.5 applies results of this section to a pole balancing problem.

2.3.1 Two-Player Problem: Optimal Solution

This subsection presents the optimal solution to the steady state minimization for the two-player problem.

In order to find a structure for the optimal controller, decompose $x(t)$ into three independent terms

$$x(t) = \begin{bmatrix} \zeta_1(t) \\ \zeta_2(t) \end{bmatrix} + \hat{x}(t),$$

where $\hat{x}(t) = \mathbb{E}[x(t)|x(0:t-1)]$. Since the input $u(t-1)$ depends on $x(0:t-1)$, it follows that

$$\begin{aligned} \hat{x}(t) &= Ax(t-1) + Bu(t-1) \\ \zeta_1(t) &= w_1(t-1) \\ \zeta_2(t) &= w_2(t-1). \end{aligned}$$

Thus $\hat{x}(t)$, $\zeta_1(t)$, and $\zeta_2(t)$ are, indeed, pairwise independent. The term $\hat{x}(t)$ denotes the expected value of $x(t)$ given the information shared by both player 1 and player 2. The term $\zeta_1(t)$ depends on the information available only to player 1, and similarly $\zeta_2(t)$ depends on the information available only to player 2.

Furthermore, note that there are some matrices $M_1(t)$ and $M_2(t)$ such that $\hat{x}_1(t) = M_1(t)x(0:t-1)$ and $\hat{x}_2(t) = M_2(t)x(0:t-1)$. It follows that there is an invertible linear mapping between $[\zeta_1(t)^T, x(0:t-1)^T]^T$ and $[x_1(t)^T, x(0:t-1)^T]^T$:

$$\begin{bmatrix} \zeta_1(t) \\ x(0:t-1) \end{bmatrix} = \begin{bmatrix} x_1(t) - \hat{x}_1(t) \\ x(0:t-1) \end{bmatrix} = \begin{bmatrix} x_1(t) - M_1(t)x(0:t-1) \\ x(0:t-1) \end{bmatrix} = \begin{bmatrix} I & -M_1(t) \\ 0 & I \end{bmatrix} \begin{bmatrix} x_1(t) \\ x(0:t-1) \end{bmatrix}$$

Similarly

$$\begin{bmatrix} \zeta_2(t) \\ x(0:t-1) \end{bmatrix} = \begin{bmatrix} I & -M_2(t) \\ 0 & I \end{bmatrix} \begin{bmatrix} x_2(t) \\ x(0:t-1) \end{bmatrix}.$$

Therefore, if $u_1(t)$ can be shown to be a function of $\zeta_1(t)$ and $x(0:t-1)$ (or a function of $\zeta_1(t)$ and $\hat{x}(t)$ in particular), then it must satisfy the information constraint of Equation (2.6). Similarly if $u_2(t)$ is a function of $\zeta_2(t)$ and $x(0:t-1)$ (or a function of $\zeta_2(t)$ and $\hat{x}(t)$), then it satisfies the information constraint.

Theorem 1. *There exist matrices K, H_1, H_2, X_1, and X_2, such that the optimal controller for the*

two-player problem is given by

$$u(t) = -\begin{bmatrix} H_1\zeta_1(t) \\ H_2\zeta_2(t) \end{bmatrix} - K\hat{x}(t),$$

and the optimal cost is given by

$$\mathrm{Tr}(W_1X_1) + \mathrm{Tr}(W_2X_2).$$

***Remark* 1.** The input $-K\hat{x}(t)$ could be interpreted as a command sent by a "manager" based information $x(0:t-1)$. Player 1 then applies a correction term $-H_1\zeta_1(t)$, based on information unavailable to the "manager." Similarly, $-H_2\zeta_2(t)$ represents player 2's correction term.

The gains, as well as the costs are specified by the stabilizing solution to the algebraic Riccati equation, S:

$$S = Q + A^TSA - A^TSB(R + B^TSB)^{-1}B^TSA.$$

For more compact notation, define the block columns of A and B as

$$\left[\begin{array}{c|c} A_1 & A_2 \end{array}\right] = \left[\begin{array}{c|c} A_{11} & A_{12} \\ A_{21} & A_{22} \end{array}\right], \qquad \left[\begin{array}{c|c} B_1 & B_2 \end{array}\right] = \left[\begin{array}{c|c} B_{11} & 0 \\ 0 & B_{22} \end{array}\right]. \tag{2.7}$$

The gains are then given by

$$\begin{aligned} K &= (R + B^TSB)^{-1}B^TSA \\ H_1 &= (R_{11} + B_1^TSB_1)^{-1}B_1^TSA_1 \\ H_2 &= (R_{22} + B_2^TSB_2)^{-1}B_2^TSA_2, \end{aligned}$$

and the cost matrices, X_1 and X_2, are given by

$$\begin{aligned} X_1 &= Q_{11} + A_1^TSA_1 - A_1^TSB_1(R_{11} + B_1^TSB_1)^{-1}B_1^TSA_1 \\ X_2 &= Q_{22} + A_2^TSA_2 - A_2^TSB_2(R_{22} + B_2^TSB_2)^{-1}B_2^TSA_2. \end{aligned}$$

The mapping from x to u can be rearranged to give a dynamic state-space controller in the standard form:

$$\begin{aligned} \hat{x}(t+1) &= B\left(\begin{bmatrix} H_1 & 0 \\ 0 & H_2 \end{bmatrix} - K\right)\hat{x}(t) + \left(A - B\begin{bmatrix} H_1 & 0 \\ 0 & H_2 \end{bmatrix}\right)x(t) \\ u(t) &= \left(\begin{bmatrix} H_1 & 0 \\ 0 & H_2 \end{bmatrix} - K\right)\hat{x}(t) - \begin{bmatrix} H_1 & 0 \\ 0 & H_2 \end{bmatrix}x(t). \end{aligned}$$

Note that the direct feedthrough term is block diagonal, which is required since $u_1(t)$ cannot depend on $x_2(t)$ and $u_2(t)$ cannot depend on $x_1(t)$.

2.3.2 Two-Player Problem: Finite-Horizon Derivation

In order to derive the optimal controller, the finite-horizon version, with cost given by Equation 2.5, will be solved, and the infinite-horizon version follows by taking limits.

The following lemma shows how an input structure based on the distribution of information between the players can be assumed.

Lemma 1. *The optimal input can be decomposed as*

$$u(t) = \begin{bmatrix} \varphi_1(t) \\ \varphi_2(t) \end{bmatrix} + \hat{u}(t),$$

where $\varphi_1(t)$, $\varphi_2(t)$, and $\hat{u}(t)$ are independent random variables which are linear functions of $\zeta_1(t)$, $\zeta_2(t)$, and $x(0:t-1)$, respectively.

Proof. By Lemma 5 in the appendix to this chapter, the information structure given by $u_1(t) = \gamma_{1,t}(x_1(0:t), x_2(0:t-1))$ and $u_2(t) = \gamma_{2,t}(x_1(0:t-1), x_2(0:t))$ is partially nested. Thus, by Theorem 4 in the appendix to this chapter, the optimal inputs must be linear functions of relevant information. Note that $(x_1(0:t), x_2(0:t-1))$ is just a rearrangement of terms in $(x(0:t-1), x_1(t))$. Furthermore, it was shown above that there is an invertible linear transformation between $(x(0:t-1), x_1(t))$ and $(x(0:t-1), \zeta_1(t))$. Therefore, there is an invertible linear transformation between $(x_1(0:t), x_2(0:t-1))$ and $(x(0:t-1), \zeta_1(t))$. Likewise, there is an invertible transformation between $(x_1(0:t-1), x_2(0:t))$ and $(x(0:t-1), \zeta_2(t))$. It follows that there are matrices $\Gamma_1(t)$, $\Gamma_2(t)$, $\chi_1(t)$, and $\chi_2(t)$, such that

$$\begin{aligned} u_1(t) &= \Gamma_1(t)x(0:t-1) + \chi_1(t)\zeta_1(t) \\ u_2(t) &= \Gamma_2(t)x(0:t-1) + \chi_2(t)\zeta_2(t). \end{aligned}$$

The proof is completed by defining

$$\hat{u}(t) = \begin{bmatrix} \Gamma_1(t) \\ \Gamma_2(t) \end{bmatrix} x(0:t-1),$$

$$\varphi_1(t) = \chi_1(t)\zeta_1(t), \text{ and } \varphi_2(t) = \chi_2(t)\zeta_2(t).$$

Independence of the terms follows from the pairwise independence of $x(0:t-1)$, $\zeta_1(t)$, and $\zeta_2(t)$. □

The lemma combined with the decomposition of the state in terms of $\hat{x}$, ζ_1, and ζ_2 implies that the summand of the cost function can be decomposed as

$$\begin{aligned}\mathbb{E}\left[x(t)^T Q x(t) + u(t)^T R u(t)\right] &= \mathbb{E}\left[\hat{x}(t)^T Q \hat{x}(t) + \hat{u}(t)^T R \hat{u}(t)\right] \\ &\quad + \mathbb{E}\left[\zeta_1(t)^T Q_{11} \zeta_1(t) + \varphi_1(t)^T R_{11} \varphi_1(t)\right] \\ &\quad + \mathbb{E}\left[\zeta_2(t)^T Q_{22} \zeta_2(t) + \varphi_2(t)^T R_{22} \varphi_2(t)\right].\end{aligned} \tag{2.8}$$

The solution will proceed via a dynamic programming argument. Let $\mathbb{E}[J(\hat{x}, \zeta_1, \zeta_2, t)]$ denote the optimal expected cost-to-go function, when the state is decomposed as $\hat{x}$, ζ_1, and ζ_2 at time t. By independence, $\mathbb{E}[J(\hat{x}, \zeta_1, \zeta_2, N)]$ can be decoupled as

$$\mathbb{E}[J(\hat{x}, \zeta_1, \zeta_2, N)] = \mathbb{E}\left[\hat{x}^T \Lambda \hat{x}\right] + \mathbb{E}\left[\zeta_1^T \Lambda_{11} \zeta_1\right] + \mathbb{E}\left[\zeta_2^T \Lambda_{22} \zeta_2\right].$$

Let $S(N) = \Lambda$, $X_1(N) = \Lambda_{11}$, and $X_2(N) = \Lambda_{22}$. For $t \leq N$, it will be shown that $J(\hat{x}, \zeta_1, \zeta_2, t)$ has the form

$$J(\hat{x}, \zeta_1, \zeta_2, t) = \hat{x}^T S(t) \hat{x} + \zeta_1^T X_1(t) \zeta_1 + \zeta_2^T X_2(t) \zeta_2 + \sum_{j=t+1}^{N} \left(\mathrm{Tr}(W_1 X(j)) + \mathrm{Tr}(W_2 Y X_2(j))\right), \tag{2.9}$$

for some matrices $S(t)$, $X_1(t)$, and $X_2(t)$ to be specified.

Inductively assume that $J(\hat{x}, \zeta_1, \zeta_2, t+1)$ has the form given in Equation (2.9). Then $\mathbb{E}[J(\hat{x}, \zeta_1, \zeta_2, t)]$ is given by the Bellman equation:

$$\mathbb{E}[J(\hat{x}, \zeta_1, \zeta_2, t)] = \min_{\hat{u}, \zeta_1, \zeta_2} \mathbb{E}\left[x^T Q x + u^T R u + J(Ax + Bu, w_1, w_2, t+1)\right].$$

Note that $J(Ax + Bu, w_1, w_2, t+1)$ can be expanded as

$$\begin{aligned}J(Ax + Bu, w_1, w_2, t+1) &= (Ax + Bu)^T S(t+1)(Ax + Bu) \\ &\quad + w_1^T X_1(t+1) w_1 + w_2^T X_2(t+1) w_2 + \sum_{j=t+2}^{N} \left(\mathrm{Tr}(W_1 X_1(j)) + \mathrm{Tr}(W_2 X_2(j))\right).\end{aligned} \tag{2.10}$$

The expected value of the second line in Equation (2.10) can be grouped as

$$\sum_{j=t+1}^{N} \left(\mathrm{Tr}(W_1 X_1(j)) + \mathrm{Tr}(W_2 X_2(j))\right).$$

Recalling the expansion of x in terms of $\hat{x}$, ζ_1, and ζ_2 and the expansion of u in terms of $\hat{u}$, φ_1, and φ_2, the expected value of the first term on the right-hand side of Equation (2.10) can be

expanded further as follows:

$$\begin{aligned}\mathbb{E}[(Ax+Bu)^T S(t+1)(Ax+Bu)] &= \mathbb{E}[(A\hat{x}+B\hat{u})^T S(t+1)(A\hat{x}+B\hat{u})] \\ &\quad +\mathbb{E}\left[(A_1\zeta_1+B_1\varphi_1)^T S(t+1)(A_1\zeta_1+B_1\varphi_1)\right] \\ &\quad +\mathbb{E}\left[(A_2\zeta_2+B_2\varphi_2)^T S(t+1)(A_2\zeta_2+B_2\varphi_2)\right].\end{aligned} \tag{2.11}$$

Here A_1, A_2, B_1, and B_2 are the block columns defined in Equation (2.7). Note that in the expansion, independent cross terms are set to zero.

Combining Equations (2.8) and (2.11) shows that the right-hand side of the Bellman equation can be decomposed into three independent minimizations, plus a constant term:

$$\begin{aligned}&\min_{\hat{u},\zeta_1,\zeta_2} \mathbb{E}\left[x^TQx+u^TRu+J(Ax+Bu,w_1,w_2,t+1)\right] = \\ &\quad \min_{\hat{u}} \mathbb{E}\left[\hat{x}^TQ\hat{x}+\hat{u}^TR\hat{u}+(A\hat{x}+B\hat{u})^TS(t+1)(A\hat{x}+B\hat{u})\right] + \\ &\quad \min_{\varphi_1} \mathbb{E}\left[\zeta_1^TQ_{11}\zeta_1+\varphi_1^TR_{11}\varphi_1+(A_1\zeta_1+B_1\varphi_1)^TS(t+1)(A_1\zeta_1+B_1\varphi_1)\right] + \\ &\quad \min_{\varphi_2} \mathbb{E}\left[\zeta_2^TQ_{22}\zeta_2+\varphi_2^TR_{22}\varphi_2+(A_2\zeta_2+B_2\varphi_2)^TS(t+1)(A_2\zeta_2+B_2\varphi_2)\right] + \\ &\quad \sum_{j=t+1}^{N}(\mathrm{Tr}(W_1X_1(j))+\mathrm{Tr}(W_2X_2(j))).\end{aligned}$$

Quadratic minimization shows that the optimal inputs are given by

$$\begin{aligned}\hat{u}(t) &= -K(t)\hat{x}(t) \\ \varphi_1(t) &= -H_1(t)\zeta_1(t) \\ \varphi_2(t) &= -H_2(t)\zeta_2(t),\end{aligned}$$

where the gains are given by

$$\begin{aligned}K(t) &= (R+B^TS(t+1)B)^{-1}B^TS(t+1)A \\ H_1(t) &= (R_{11}+B_1^TS(t+1)B_1)^{-1}B_1^TS(t+1)A_1 \\ H_2(t) &= (R_{22}+B_2^TS(t+1)B_2)^{-1}B_2^TS(t+1)A_2.\end{aligned}$$

Finally, the matrices $S(t)$, $X_1(t)$, and $X_2(t)$ are computed recursively as follows:

$$\begin{aligned}S(t) &= Q+A^TS(t+1)A-A^TS(t+1)B(R+B^TS(t+1)B)^{-1}B^TS(t+1)A \\ X_1(t) &= Q_{11}+A_1^TS(t+1)A_1-A_1^TS(t+1)B_1(R_{11}+B_1^TS(t+1)B_1)^{-1}B_1^TS(t+1)A_1 \\ X_2(t) &= Q_{22}+A_2^TS(t+1)A_2-A_2^TS(t+1)B_2(R_{22}+B_2^TS(t+1)B_2)^{-1}B_2^TS(t+1)A_2.\end{aligned}$$

By construction, $J(\hat{x}, \zeta_1, \zeta_2, t)$ satisfies the Bellman equation for all $t \leq N$. Thus, since $\mathbb{E}[J(\hat{x}, \zeta_1, \zeta_2, N)]$ is the optimal expected cost-to-go at time N, it follows inductively that $\mathbb{E}[J(\hat{x}, \zeta_1, \zeta_t, t)]$ is the optimal expected cost-to-go for all $t \leq N$, and the optimal control has been found. Noting that $x(0) = 0$, the optimal expected cost is given by

$$\sum_{t=1}^{N} \left(\mathrm{Tr}(W_1 X_1(t)) + \mathrm{Tr}(W_2 X_2(t)) \right). \tag{2.12}$$

2.3.3 Two-Player Problem: Steady State

To derive the steady state regulator from the finite-horizon regulator, assume that as N approaches ∞, $S(t)$ approaches the stabilizing solution of the corresponding algebraic Riccati equation. Then $K(t)$, $H_1(t)$, $H_2(t)$, $X_1(t)$, and $X_2(t)$ will approach the values of K, H_1, H_2, X_1, and X_2 specified by the theorem and the derivation of the controller is complete.

To compute the steady state cost, note that the average cost approaches the steady state cost as $N \to \infty$:

$$\begin{aligned} &\lim_{t\to\infty} \mathbb{E}\left[x(t)^T Q x(t) + u(t)^T R u(t)\right] = \\ &\quad \lim_{N\to\infty} \tfrac{1}{N} \mathbb{E}\left[\sum_{t=0}^{N-1} \left(x(t)^T Q x(t) + u(t)^T R u(t)\right) + x(N)^T \Lambda x(N)\right]. \end{aligned}$$

Recall that for fixed N, the cost is given by $\sum_{t=1}^{N} \left(\mathrm{Tr}(W_1 X_1(t)) + \mathrm{Tr}(W_2 X_2(t)) \right)$. Thus, dividing by N and taking a limit gives the steady state cost:

$$\lim_{N\to\infty} \frac{1}{N} \sum_{t=1}^{N} \left(\mathrm{Tr}(W_1 X_1(t)) + \mathrm{Tr}(W_2 X_2(t)) \right) = \mathrm{Tr}(W_1 X_1) + \mathrm{Tr}(W_2 X_2).$$

2.3.4 Cost Comparisons

This subsection places the results of Theorem 1 in the context of more well known results. In particular, the controller for the two-player problem will be compared to controllers for two centralized information structures: state feedback with and without delay. Expressions for the steady state costs in the various scenarios will be compared.

Consider the following three information structures:

$$\begin{array}{ll} \text{1) Centralized} & \begin{cases} u_1(t) &= \gamma_{1,t}(x_1(0:t), x_2(0:t)) \\ u_2(t) &= \gamma_{2,t}(x_1(0:t), x_2(0:t)) \end{cases} \\ \text{2) Two-Player} & \begin{cases} u_1(t) &= \gamma_{1,t}(x_1(0:t), x_2(0:t-1)) \\ u_2(t) &= \gamma_{2,t}(x_1(0:t-1), x_2(0:t)) \end{cases} \\ \text{3) Delayed} & \begin{cases} u_1(t) &= \gamma_{1,t}(x_1(0:t-1), x_2(0:t-1)) \\ u_2(t) &= \gamma_{2,t}(x_2(0:t-1), x_2(0:t-1)), \end{cases} \end{array}$$

with optimal steady state costs, c_{cen}, c_{dec}, and c_{del}, respectively. Note that any delayed control law can be implemented by a two-player controller, and any two-player controller can be implemented by a centralized controller, and thus

$$c_{cen} \leq c_{dec} \leq c_{del}.$$

To get more explicit comparisons, c_{cen} and c_{del} will be derived. The following argument is based on a classical solutions for centralized LQG (see [36]).

Assume that x and u are in steady state, and define c by

$$c = \mathbb{E}\left[x^T Q x + u^T R u\right]. \tag{2.13}$$

By the steady state assumption, x and $Ax + Bu + w$ must have the same covariance, and thus

$$\begin{aligned} \mathbb{E}\left[x^T S x\right] &= \mathbb{E}\left[(Ax + Bu + w)^T S (Ax + Bu + w)\right] \\ &= \mathbb{E}\left[(Ax + Bu)^T S (Ax + Bu)\right] + \mathrm{Tr}(WS). \end{aligned} \tag{2.14}$$

Here S is solution to the LQR Riccati equation used in Subsection 2.3.1. Adding Equations (2.13) and (2.14) gives

$$\mathbb{E}\left[x^T S x\right] + c = \mathbb{E}\left[x^T Q x + u^T R u + (Ax + Bu)^T S (Ax + Bu)\right] + \mathrm{Tr}(WS). \tag{2.15}$$

Let K be the LQR gain, from Subsection 2.3.1, and let $\Omega = R + B^T S B$. Completing the square on the first term on the right-hand side gives

$$x^T Q x + u^T R u + (Ax + Bu)^T S (Ax + Bu) = x^T S x + (Kx + u)^T \Omega (Kx + u).$$

Therefore, Equation (2.15) can be rewritten as

$$\mathbb{E}\left[x^T S x\right] + c = \mathbb{E}\left[x^T S x\right] + \mathbb{E}\left[(Kx+u)^T \Omega (Kx+u)\right] + \mathrm{Tr}(WS)$$

Subtracting the $\mathbb{E}\left[x^T S x\right]$ term from both sides gives the general expression for the cost:

$$c = \mathrm{Tr}(WS) + \mathbb{E}\left[(Kx+u)^T \Omega (Kx+u)\right]. \tag{2.16}$$

Aside from the steady state assumption, no other assumptions about the input u has been made. In the case of centralized state-feedback, x is available to both players and $u = -Kx$ is the optimal control and the optimal cost is given by

$$c_{cen} = \mathrm{Tr}(WS).$$

In the delayed centralized case, let $\hat{x}(t) = \mathbb{E}[x(t)|x(0:t-1)]$, as in Subsection 2.3.1. Recall that $x(t) - \hat{x}(t) = w(t)$. Since the input $u(t)$ can only depend on $x(0:t-1)$, it follows that $Kx(t) + u(t)$ can be decomposed into two independent terms as

$$Kx(t) + u(t) = K(x(t) - \hat{x}(t)) + (K\hat{x}(t) + u(t)) = Kw(t-1) + (K\hat{x}(t) + u(t)).$$

Thus, for the case of delayed centralized feedback, the cost can be further decomposed as

$$\begin{aligned} c &= \mathrm{Tr}(WS) + +\mathbb{E}\left[w(t-1)^T K^T \Omega K w(t-1)\right] + \mathbb{E}\left[(K\hat{x}(t) + \hat{u}(t))^T \Omega (K\hat{x}(t) + \hat{u}(t))\right] \\ &= \mathrm{Tr}(WS) + \mathrm{Tr}(WK^T \Omega K) + \mathrm{E}\left[(K\hat{x}(t) + u(t))^T \Omega (K\hat{x}(t) + u(t))\right]. \end{aligned}$$

The optimal control is seen to be $u(t) = -K\hat{x}(t)$ with optimal cost

$$c_{del} = \mathrm{Tr}(WS) + \mathrm{Tr}(WK^T \Omega K).$$

The decomposition employed to derive c_{del} is a special case of the classical separation principle in output feedback control. In particular if $\hat{x}(t)$ is viewed as the filter state, and P is the steady state covariance of $x(t) - \hat{x}(t)$, then the cost becomes $c_{out} = \mathrm{Tr}(WS) + \mathrm{Tr}(PK^T \Omega K)$. In the current case, $P = W$.

***Remark* 2.** Note that the optimal input in the delay case is given by $u(t) = -K\hat{x}(t)$, while the

optimal input for the two-player case is given by

$$u(t) = -\begin{bmatrix} H_1\zeta_1(t) \\ H_2\zeta_2(t) \end{bmatrix} - K\hat{x}(t).$$

Thus, the terms $-H_1\zeta_1(t)$ and $-H_2\zeta_2(t)$ can be viewed as corrections, based on local information, to the optimal delayed controller.

The relationships between the costs can be seen most readily after a few algebraic rearrangements. First, c_{cen} can be rewritten as

$$\begin{aligned} c_{cen} &= \mathrm{Tr}(WS) \\ &= \mathrm{Tr}(W(Q + A^TSA - A^TSB\Omega^{-1}B^TSA)) \\ &= \mathrm{Tr}(W(Q + A^TSA)) - \mathrm{Tr}(W_1A_1^TSB\Omega^{-1}B^TSA_1) - \mathrm{Tr}(W_2A_2^TSB\Omega^{-1}B^TSA_2) \\ &= \mathrm{Tr}(W(Q + A^TSA)) - \mathrm{Tr}(B^TSA_1W_1A_1^TSB\Omega^{-1}) - \mathrm{Tr}(B^TSA_2W_2A_2^TSB\Omega^{-1}). \end{aligned}$$

Next, the two-player cost can be rewritten:

$$\begin{aligned} c_{dec} &= \mathrm{Tr}(W_1X_1) + \mathrm{Tr}(W_2X_2) \\ &= \mathrm{Tr}(W_1(Q_{11} + A_1^TSA_1 - A_1^TSB_1\Omega_{11}^{-1}B_1^TSA_1)) + \\ &\quad \mathrm{Tr}(W_2(Q_{22} + A_2^TSA_2 - A_2^TSB_2\Omega_{22}^{-1}B_2^TSA_2)) \\ &= \mathrm{Tr}(W(Q + A^TSA)) - \mathrm{Tr}\left(W_1A_1^TSB\begin{bmatrix} \Omega_{11}^{-1} & 0 \\ 0 & 0 \end{bmatrix}B^TSA_1\right) \\ &\quad -\mathrm{Tr}\left(W_2A_2^TSB\begin{bmatrix} 0 & 0 \\ 0 & \Omega_{22}^{-1} \end{bmatrix}B^TSA_2\right) \\ &= \mathrm{Tr}(W(Q + A^TSA)) - \mathrm{Tr}\left(B^TSA_1W_1A_1^TSB\begin{bmatrix} \Omega_{11}^{-1} & 0 \\ 0 & 0 \end{bmatrix}\right) \\ &\quad -\mathrm{Tr}\left(B^TSA_2W_2A_2^TSB\begin{bmatrix} 0 & 0 \\ 0 & \Omega_{22}^{-1} \end{bmatrix}\right) \end{aligned}$$

The third equality follows because

$$B_1\Omega_{11}^{-1}B_1^T = \begin{bmatrix} B_{11} \\ 0 \end{bmatrix}\Omega_{11}^{-1}\begin{bmatrix} B_{11}^T & 0 \end{bmatrix} = \begin{bmatrix} B_{11} & 0 \\ 0 & B_{22} \end{bmatrix}\begin{bmatrix} \Omega_{11}^{-1} & 0 \\ 0 & 0 \end{bmatrix}\begin{bmatrix} B_{11}^T & 0 \\ 0 & B_{22}^T \end{bmatrix} = B\begin{bmatrix} \Omega_{11}^{-1} & 0 \\ 0 & 0 \end{bmatrix}B^T$$

and

$$B_2\Omega_{22}^{-1}B_2^T = \begin{bmatrix} 0 \\ B_{22} \end{bmatrix} \Omega_{22}^{-1} \begin{bmatrix} 0 & B_{22}^T \end{bmatrix} = \begin{bmatrix} B_{11} & 0 \\ 0 & B_{22} \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 0 & \Omega_{22}^{-1} \end{bmatrix} \begin{bmatrix} B_{11}^T & 0 \\ 0 & B_{22}^T \end{bmatrix} = B \begin{bmatrix} 0 & 0 \\ 0 & \Omega_{22}^{-1} \end{bmatrix} B^T.$$

Finally, c_{del} can be rewritten as

$$\begin{aligned} c_{del} &= \mathrm{Tr}(WS) + \mathrm{Tr}(WK^T\Omega K) \\ &= \mathrm{Tr}(W(S + A^TSB\Omega^{-1}B^TSA)) \\ &= \mathrm{Tr}(W(Q + A^TSA)). \end{aligned}$$

Now the inequalities $c_{cen} \leq c_{dec} \leq c_{del}$ can be written as

$$\begin{matrix} \mathrm{Tr}(W(Q+A^TSA)) \\ -\mathrm{Tr}(B^TSA_1W_1A_1^TSB\Omega^{-1}) \\ -\mathrm{Tr}(B^TSA_2W_2A_2^TSB\Omega^{-1}) \end{matrix} \leq \begin{matrix} \mathrm{Tr}(W(Q+A^TSA)) \\ -\mathrm{Tr}\left(B^TSA_1W_1A_1^TSB \begin{bmatrix} \Omega_{11}^{-1} & 0 \\ 0 & 0 \end{bmatrix}\right) \\ -\mathrm{Tr}\left(B^TSA_2W_2A_2^TSB \begin{bmatrix} 0 & 0 \\ 0 & \Omega_{22}^{-1} \end{bmatrix}\right) \end{matrix} \leq \mathrm{Tr}(W(Q+A^TSA)).$$

So the cost of the delayed controller appears in each term, with the subtracted terms of c_{dec} and c_{cen} corresponding to benefits of extra information.

To see why $c_{cen} \leq c_{dec}$, note that

$$\begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{bmatrix}^{-1} = \begin{bmatrix} \Omega_{11}^{-1}\Omega_{12} \\ -I \end{bmatrix} (\Omega_{22} - \Omega_{21}\Omega_{11}^{-1}\Omega_{12})^{-1} \begin{bmatrix} \Omega_{21}\Omega_{11}^{-1} & -I \end{bmatrix} + \begin{bmatrix} \Omega_{11}^{-1} & 0 \\ 0 & 0 \end{bmatrix} \succeq \begin{bmatrix} \Omega_{11}^{-1} & 0 \\ 0 & 0 \end{bmatrix}.$$

A similar argument shows that

$$\begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{bmatrix}^{-1} \succeq \begin{bmatrix} 0 & 0 \\ 0 & \Omega_{22}^{-1} \end{bmatrix}.$$

***Remark* 3.** The comparisons of the various costs demonstrates the benefits of using all available information. It is interesting to compare the result of the comparisons with intuition about motor learning. When faced with a new task, motor commands must be processed consciously in "high-level" brain regions. If it is assumed that there is a large computational delay for conscious processing, then this control policy may be analogous to the centralized delay case. Eventually, gains and correction terms are learned in "lower" brain regions and the spinal cord. At this point, the control strategy becomes decentralized, but performance increases. While the connection is

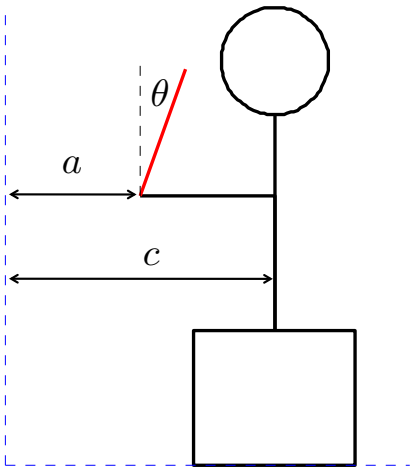

Figure 2.3: An arm balancing a pole sits on top of a moving cart. Two inputs are chosen, a force on the arm and force on the cart. The player choosing the arm forces is assumed to have instantaneous state information about the arm and the pendulum (through "vision" and "proprioception"). The player choosing the cart forces is assumed to have instantaneous state information only about the cart. Communication between the players occurs via a one-step delay.

speculative, it may lend insight into the function of low-level feedback loops in the motor system.

2.3.5 Pole Balancing Example

To see how the concepts of this section arise in physical systems, consider the setup in Figure 2.3. Here a pole is balanced using an arm, which is mechanically coupled to a moving cart. The inputs are forces applied to the arm and forces applied to the cart. The dynamics of the system are defined by

$$\begin{bmatrix} m_c & 0 & 0 \\ 0 & m_a + m_p & m_p l \cos\theta \\ 0 & m_p l \cos\theta & m_p l^2 \end{bmatrix} \begin{bmatrix} \ddot{c} \\ \ddot{a} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} -kc + ka + u_c + w_c \\ kc - ka + m_p l \dot{\theta}^2 \sin\theta + u_a + w_a \\ m_p g l \sin\theta + w_\theta \end{bmatrix}.$$

Here c and a are the positions of the arm and the cart, respectively. The variable θ is the angle of the pendulum. The cart and arm input forces are given by u_c and u_a, respectively, while the noises are given by w_c, w_a, and w_θ. The masses of the cart, arm, and pendulum are given by m_c, m_a, and m_p, respectively. The length of the pendulum is given by l, the gravitational constant is given by g. The spring constant, k, describes the coupling between the cart and the arm.

Linearizing the dynamics and writing the equations in first-order form gives:

$$\dot{x} = A_c x + B_c u + G_c w,$$

where the variables are given by

$$x = \begin{bmatrix} c \\ \dot{c} \\ a \\ \dot{a} \\ \theta \\ \dot{\theta} \end{bmatrix}, \qquad u = \begin{bmatrix} u_c \\ u_a \end{bmatrix}, \qquad w = \begin{bmatrix} w_c \\ w_a \\ w_\theta \end{bmatrix}$$

and the matrices are given by

$$A_c = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ -\frac{k}{m_c} & 0 & \frac{k}{m_c} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ \frac{k}{m_a} & 0 & -\frac{k}{m_a} & 0 & -\frac{m_p g}{m_a} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ -\frac{k}{m_a l} & 0 & \frac{k}{m_a l} & 0 & \frac{(m_a+m_p)g}{m_a l} & 0 \end{bmatrix}, \quad B_c = \begin{bmatrix} 0 & 0 \\ \frac{1}{m_c} & 0 \\ 0 & 0 \\ 0 & \frac{1}{m_a} \\ 0 & 0 \\ 0 & -\frac{1}{m_a l} \end{bmatrix},$$

$$G_c = \begin{bmatrix} 0 & 0 & 0 \\ \frac{1}{m_c} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & \frac{1}{m_a} & -\frac{1}{m_a l} \\ 0 & 0 & 0 \\ 0 & -\frac{1}{m_a l} & \frac{(m_a+m_p)g}{m_a l} \end{bmatrix}.$$

Discretizing the dynamics by a first-order Euler approximation with time-step τ gives state matrices

$$A = I + \tau A_c, \qquad B = \tau B_c, \qquad G = \tau G_c.$$

Assume that the noise w (applied through G) has identity covariance. Then the discrete-time dynamics can be equivalently written as

$$x(t+1) = Ax(t) + Bu(t) + w(t),$$

where w has covariance $W = GG^T$.

Grouping the state variables as

$$x_1 = \begin{bmatrix} c \\ \dot{c} \end{bmatrix} \qquad \text{and} \qquad x_2 = \begin{bmatrix} a \\ \dot{a} \\ \theta \\ \dot{\theta} \end{bmatrix},$$

and set $u_1 = u_c$ and $u_2 = u_a$. It follows that B and W are block-diagonal with respect to this partitioning of the state and input variables. Thus the discrete-time dynamics fit the sparsity constraints for the two-player problem:

$$B = \left[\begin{array}{c|c} 0 & 0 \\ \frac{\tau}{m_c} & 0 \\ \hline 0 & 0 \\ 0 & \frac{\tau}{m_a} \\ 0 & 0 \\ 0 & -\frac{\tau}{m_a l} \end{array}\right], \qquad W = \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix},$$

with

$$W_1 = \begin{bmatrix} 0 \\ \frac{\tau}{m_c} \end{bmatrix} \begin{bmatrix} 0 & \frac{\tau}{m_c} \end{bmatrix}, \qquad W_2 = \begin{bmatrix} 0 & 0 \\ \frac{\tau}{m_a} & -\frac{\tau}{m_a l} \\ 0 & 0 \\ -\frac{\tau}{m_a l} & \frac{(m_a+m_p)g}{m_a l} \end{bmatrix} \begin{bmatrix} 0 & \frac{\tau}{m_a} & 0 & -\frac{\tau}{m_a l} \\ 0 & -\frac{\tau}{m_a l} & 0 & \frac{(m_a+m_p)g}{m_a l} \end{bmatrix}.$$

To see how the different control constraints lead to different steady state costs, see Figure 2.4. By increasing the spring constant k, the coupling between the cart and arm subsystems becomes stronger and the system becomes harder to control. With the stronger coupling, the penalty for delay increases. Thus the cost grows fastest for the delayed case. The cost of the two-player policy grows with intermediate speed, and the cost of the centralized case grows the slowest.

2.4 General Case

This section extends the method from the two-player problem to derive optimal controllers for delay structures specified by any strongly connected graph. The general method for solving the optimal control problems follows the pattern from the two-player case:

- Decompose the information available to the players into independent components, based on "who knows what."

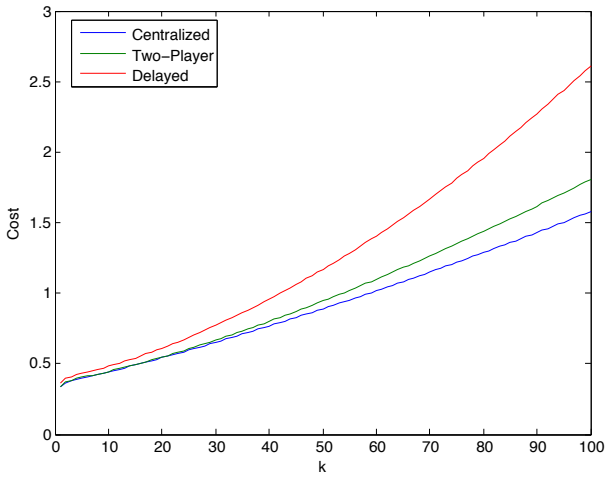

Figure 2.4: Steady state costs for different values of the spring constant, k. Note that as k increases, the coupling between x_1 and x_2 gets stronger. The steady state cost grows in all cases, but it grows fastest in the delay case and slowest in the centralized case. The other parameters are set to $\tau = 0.1$, $m_p = m_c = m_a = 1$, $g = 1$, and $l = 10$. Both Q and R are set to identity matrices of appropriate size.

- Use the information decomposition to decouple the input and the state into independent components.
- Find update equations for the decoupled state components.
- Set up the dynamic programming problem.
- Use independence to decompose the dynamic programming problem into independent subproblems.
- Solve the independent subproblems to find the optimal control and optimal cost.

This section is organized as follows. The information, input, and state are decoupled in Subsection 2.4.1. In that subsection, the dynamics of the decoupled state variables are computed. Next, in Subsection 2.4.2, the optimal solution to the general problem posed in this chapter is stated. In Subsection 2.4.3, a finite-horizon version of the problem is solved via dynamic programming. Finally, in Subsection 2.4.4, the steady state controller and optimal cost are derived by limiting arguments.

Notation. For a vector partitioned into blocks

$$\begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix}$$

and $v \subset \{1, \ldots, n\}$, let $z^v = (z_i)_{i \in v}$. For instance, if $n = 5$ and $v = \{1, 3, 5\}$, then z^v is given by

$$z^{\{1,3,5\}} = \begin{bmatrix} z_1 \\ z_3 \\ z_5 \end{bmatrix}.$$

For a matrix partitioned into blocks

$$M = \begin{bmatrix} M_{11} & \cdots & M_{1n} \\ \vdots & \ddots & \vdots \\ M_{n1} & \cdots & M_{nn} \end{bmatrix}$$

and $s, v \subset \{1, \ldots, n\}$, let $M^{s,v} = (M_{i,j})_{i \in s, j \in v}$. For instance, if $n = 3$, $s = \{1, 2, 3\}$, and $v = \{1, 2\}$, then $M^{s,v}$ is given by

$$M^{\{1,2,3\},\{1,2\}} = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \\ M_{31} & M_{32} \end{bmatrix}.$$

2.4.1 Decoupled State Dynamics

This subsection describes a method for decoupling the information available to the players, based on an auxiliary graph, termed the *information hierarchy graph*. Once the information has been decoupled, the state and inputs are decomposed into independent terms. Finally, the dynamic equations for updating the decoupled state terms are given.

Let $G = (V, E)$ be the graph describing the delay structure, with $V = \{1, \ldots, n\}$. The information hierarchy graph $\mathscr{I} = (\mathscr{V}, \mathscr{E})$ is a graph describing the flow of information through G as constructed in Algorithm 1. See Figure 2.5 for a few examples of information hierarchy graphs constructed from their delay structure graphs.

Some of the more useful properties information hierarchy graphs are now listed. All of the properties are direct consequences of Algorithm 1.

- Each node has exactly one outgoing edge.
- Nodes $\{1\}, \ldots, \{n\}$ are the only nodes with no incoming edges.
- If there is a path of length k from node $\{i\}$ to node v in $\mathscr{I}$, then v is exactly the set of nodes reachable from node i within k steps in graph G.
- Since G is strongly connected, V is always a node in $\mathscr{V}$. Furthermore, the outgoing edge of V is a self-loop: $(V, V) \in \mathscr{E}$.

Algorithm 1 Information Hierarchy Graph Construction Algorithm

Start with $G = (V, E)$ and assume that $V = \{1, \ldots, n\}$.
Set $\mathscr{V} = \{\{1\}, \ldots, \{n\}\}$
Set $\mathscr{E} = \emptyset$
while There is a vertex $v \in \mathscr{V}$ with no outgoing edge **do**
  Pick $v \in \mathscr{V}$ with no outgoing edge
  Set $s = v$
  {Add to s all nodes reachable from nodes in v in one step}
  for all $i \in v$ **do**
    for all j such that $(i, j) \in E$ **do**
      if $j \notin s$ **then**
        Add j to s
      end if
    end for
  end for
  if $s \notin \mathscr{V}$ **then**
    Add s to $\mathscr{V}$
  end if
  Add edge (v, s) to $\mathscr{E}$
end while
return $\mathscr{I} = (\mathscr{V}, \mathscr{E})$

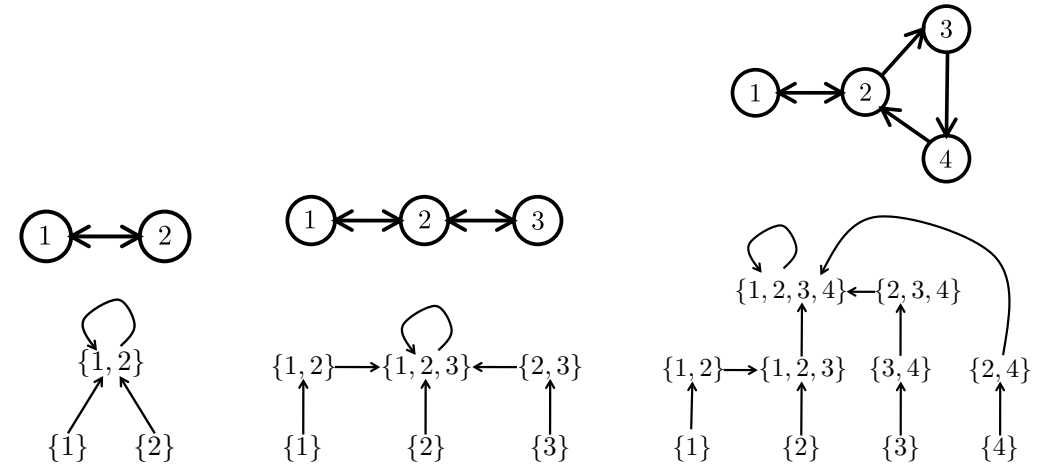

(a) Two-Player Graph (b) Three-Player Chain (c) Four-Player Example

Figure 2.5: Each subfigure depicts a delay structure graphs on the top with the associated information hierarchy graph on the bottom.

- If $(v, s) \in \mathscr{E}$ and $v \neq V$, then v is a strict subset of s.
- If every node in G is reachable from every other node in at most d steps, then for $i = 1, \ldots, n$ there is a path from $\{i\}$ to V in $\mathscr{I}$ of length at most d.
- $|\mathscr{V}| = |\mathscr{E}| \leq n(d+1)$

Algorithm 2 Information Hierarchy Graph Labeling

Label nodes $\{1\}, \ldots, \{n\}$ with $\mathscr{L}_{\{1\}}(t) = w_1(t-1), \ldots, \mathscr{L}_{\{n\}}(t) = w_n(t-1)$, respectively.
while There is a node $s \in \mathscr{V} \setminus \{V\}$ that has not been labeled **do**
  Pick $s \in \mathscr{V} \setminus \{V\}$ such that v is labeled for all v with $(v, s) \in \mathscr{E}$
  for all v such that $(v, s) \in \mathscr{E}$ **do**
    if The label for s has not been created **then**
      Set $\mathscr{L}_s(t) = \mathscr{L}_v(t-1)$
    else
      Set $\mathscr{L}_s(t) = \begin{bmatrix} \mathscr{L}_s(t) \\ \mathscr{L}_v(t-1) \end{bmatrix}$
    end if
  end for
end while
for $i = 1, \ldots, n$ **do**
  Find s and k such that $(s, V) \in \mathscr{E}$ and $w_i(t-k)$ appears in $\mathscr{L}_s(t)$ {s and k will be unique}
  Set $d_i = k$
end for
Set $\mathscr{L}_V(t) = \begin{bmatrix} x_1(0 : t-d_1) \\ \vdots \\ x_n(0 : t-d_n) \end{bmatrix}$

The main reason for defining the information hierarchy graph is that it gives a graphical method for decomposing the information available to the various players. In particular, Algorithm 2 shows how to label each node $v \in \mathscr{V}$ with a vector of information $\mathscr{L}_v(t)$ that will be useful for decomposing the state and input vectors. See Figure 2.6 for examples of labeled information hierarchy graphs.

Once the labels are defined, they can associated with the players. Let $\chi_i(t) = (\mathscr{L}_v(t))_{v:i \in v}$ be the collection of labels corresponding to player i. For example, in the three-player chain (Figure

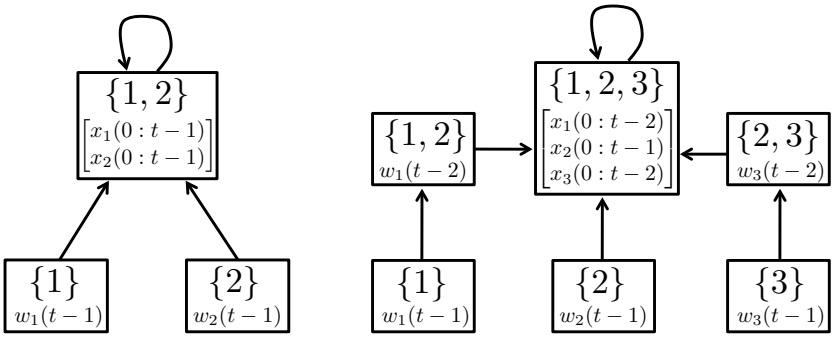

(a) Two-Player Problem

(b) Three-Player Chain

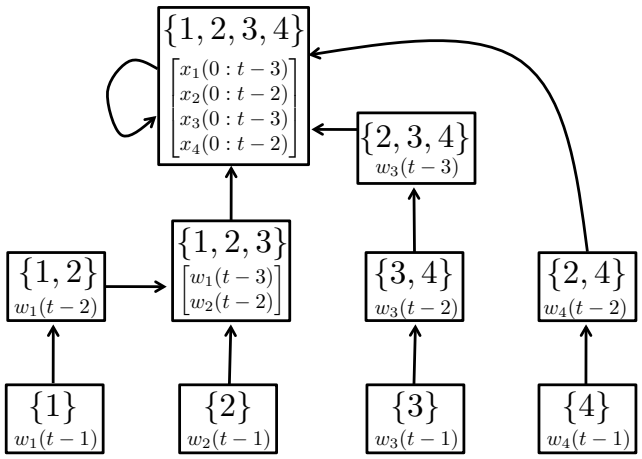

(c) Four-Player Example

Figure 2.6: Labeled information hierarchy graphs from Figure 2.5. The labels are pairwise independent and correspond to information available to all players in the corresponding node, but none of the other players.

2.6(b)), $\chi_1(t)$, $\chi_2(t)$, and $\chi_3(t)$ are given by

$$\chi_1(t) = \begin{bmatrix} \mathscr{L}_{\{1\}}(t) \\ \mathscr{L}_{\{1,2\}}(t) \\ \mathscr{L}_{\{1,2,3\}}(t) \end{bmatrix} = \begin{bmatrix} w_1(t-1) \\ w_1(t-2) \\ x_1(0:t-2) \\ x_2(0:t-1) \\ x_3(0:t-2) \end{bmatrix} \qquad \chi_2(t) = \begin{bmatrix} \mathscr{L}_{\{2\}}(t) \\ \mathscr{L}_{\{1,2\}}(t) \\ \mathscr{L}_{\{2,3\}}(t) \\ \mathscr{L}_{\{1,2,3\}}(t) \end{bmatrix} = \begin{bmatrix} w_2(t-1) \\ w_1(t-2) \\ w_3(t-2) \\ x_1(0:t-2) \\ x_2(0:t-1) \\ x_3(0:t-2) \end{bmatrix}$$

$$\chi_3(t) = \begin{bmatrix} \mathscr{L}_{\{3\}}(t) \\ \mathscr{L}_{\{2,3\}}(t) \\ \mathscr{L}_{\{1,2,3\}}(t) \end{bmatrix} = \begin{bmatrix} w_3(t-1) \\ w_3(t-2) \\ x_1(0:t-2) \\ x_2(0:t-1) \\ x_3(0:t-2) \end{bmatrix}.$$

Define $\theta_i(t)$ to be the information available to player i at time t. In other words $\theta_i(t) = [x_1(0 : t-d_{1i})^T, \ldots, x_n(0 : t-d_{ni})^T]^T$. The following lemma shows that the labels are pairwise independent and that the groupings $\chi_i(t)$ can be used in place of the information vectors $\theta_i(t)$. The lemma assumes that the input is linear. By partial nestedness, there is no loss of generality, since the optimal input must be a linear function of the available information.

Lemma 2.

1. *The labels $\mathscr{L}_v(t)$ are pairwise independent.*
2. *Assume that the each input $u_i(t)$ is a linear function of $\theta_i(t)$. For $i = 1, \ldots, n$ there is an invertible linear mapping $\Pi_i(t)$ such that $\chi_i(t) = \Pi_i(t)\theta_i(t)$.*

Proof. **Item 1.** The proof of Item 1 follows a few steps. First it is shown that for all i and k such that $1 \le k \le d_i$, there is a unique node $v \neq V$ such that $w_i(t-k)$ is a component of $\mathscr{L}_v(t)$. Next it is shown that for all $v \neq V$, the label $\mathscr{L}_v(t)$ is composed of noise terms $w_i(t-k)$ with $1 \le k \le d_i$. Once the steps above are shown, the proof of Item 1 will be complete since all such $w_i(t-k)$ are pairwise independent and also independent of

$$\mathscr{L}_V(t) = \begin{bmatrix} x_1(0:t-d_1) \\ \vdots \\ x_n(0:t-d_n) \end{bmatrix}$$

by construction.

Take any i and k such that $1 \leq k \leq d_i$. Let $v \in \mathscr{V}$ be the unique node such that there is a path from $\{i\}$ to v in $\mathscr{I}$ of length $k-1$. By construction $w_i(t-k)$ is a component of $\mathscr{L}_v(t)$. Furthermore, the path from $\{i\}$ to V has length d_i, so $v \neq V$. Uniqueness of the path implies that $\mathscr{L}_v(t)$ is the unique label containing $w_i(t-k)$.

Now say that $w_i(t-k)$ is a component of $\mathscr{L}_v(t)$ with $v \neq V$. By construction $k \geq 1$. Algorithm 2 implies that there must be a path from $\{i\}$ to v of length $k-1$. Now because the path from $\{i\}$ to V has length d_i, and $v \neq V$, it follows that $k-1 \leq d_i - 1$. In other words, $1 \leq k \leq d_i$.

The proof of Item 1 is now completed by applying assumptions of independence.

Item 2. First note that $\mathscr{L}_V(t)$ is a component of both $\chi_i(t)$ and $\theta_i(t)$ (by permuting entries). Furthermore, note that the terms in $\theta_i(t)$ that are not in $\mathscr{L}_V(t)$ are given by $x_j(t-k+1)$ for all k such that $d_{ji} < k \leq d_j$. Similarly, $w_j(t-k)$ is a component of $\chi_i(t)$ if and only if $d_{ji} < k \leq d_j$. Indeed, note that $w_j(t-k)$ appears in $\chi_i(t)$ if and only if $k \leq d_j$ and there is a path from j to i in G of length at most $k-1$. The shortest such path has length d_{ji}. Therefore $d_{ji} \leq k-1$ and the equivalence holds.

To prove the existence of an invertible linear mapping $\Pi_i(t)$ such that $\chi_i(t) = \Pi_i(t)\theta_i(t)$, consider the following sequence of equalities:

$$\Gamma_1 \theta_i(t) = \Gamma_1 \begin{bmatrix} x_1(0:t-d_{1i}) \\ \vdots \\ x_n(0:t-d_{ni}) \end{bmatrix} = \begin{bmatrix} w_1(0:t-d_{1i}-1) \\ \vdots \\ w_n(0:t-d_{ni}-1) \end{bmatrix}$$

$$\Gamma_2 \begin{bmatrix} w_1(0:t-d_{1i}-1) \\ \vdots \\ w_n(0:t-d_{ni}-1) \end{bmatrix} = \begin{bmatrix} w_1(t-d_1:t-d_{1i}-1) \\ \vdots \\ w_n(t-d_n:t-d_{ni}-1) \\ w_1(0:t-d_1-1) \\ \vdots \\ w_n(0:t-d_n-1) \end{bmatrix}$$

$$\Gamma_3 \begin{bmatrix} w_1(t-d_1:t-d_{1i}-1) \\ \vdots \\ w_n(t-d_n:t-d_{ni}-1) \\ w_1(0:t-d_1-1) \\ \vdots \\ w_n(0:t-d_n-1) \end{bmatrix} = \begin{bmatrix} w_1(t-d_1:t-d_{1i}-1) \\ \vdots \\ w_n(t-d_n:t-d_{ni}-1) \\ x_1(0:t-d_1) \\ \vdots \\ x_n(0:t-d_n) \end{bmatrix} = \chi_i(t).$$

Here Γ_1, Γ_2, and Γ_3 are all invertible linear mappings. The mappings Γ_1 and Γ_3 are guaranteed by partial nestedness and linearity of the inputs. Indeed partial nestedness implies that there are

invertible linear transformations between the information available to player i and the corresponding noise at all time steps. The matrix Γ_2 is simply a permutation. The proof is completed by defining $\Pi_i(t) = \Gamma_3\Gamma_2\Gamma_1$. □

The following lemma demonstrates that the information decomposition from Lemma 2 can be used to decouple the input into independent terms.

Lemma 3. *The optimal input $u(t)$ can be decomposed as a sum*

$$u(t) = \sum_{s\in\mathscr{V}} I_u^{V,s}\varphi_s(t), \tag{2.17}$$

where I_u is the identity matrix partitioned into blocks conforming to the partition of u, and $\varphi_s(t)$ is a linear function of $\mathscr{L}_s(t)$ of appropriate size.

Before getting into the proof, an example of the notation will be given. Say that $V = \{1,\dots,4\}$. In this case

$$I_u^{V,\{1,4\}}\varphi_{\{1,4\}}(t) = \begin{bmatrix} I & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \varphi_{\{1,4\}}(t)_1 \\ \varphi_{\{1,4\}}(t)_2 \end{bmatrix} = \begin{bmatrix} \varphi_{\{1,4\}}(t)_1 \\ 0 \\ 0 \\ \varphi_{\{1,4\}}(t)_2 \end{bmatrix}.$$

Proof. By Lemma 2, the information available to player i at time t can be decomposed into independent vectors as $\chi_i(t) = (\mathscr{L}_s(t))_{s:i\in s}$. By linearity of the optimal solution, there exist matrices $H_{i,s(t)}$ such that the optimal input is given by

$$u_i(t) = \sum_{s\in\mathscr{V}:i\in s} H_{i,s}(t)\mathscr{L}_s(t).$$

For each s let $s = \{i_1,\dots,i_s\}$, with $i_1 < \cdots < i_s$. Define $\varphi_s(t)$ by

$$\varphi_s(t) = \begin{bmatrix} H_{i_1,s}(t) \\ \vdots \\ H_{i_s,s}(t) \end{bmatrix} \mathscr{L}_s(t).$$

The following chain of equalities completes the proof:

$$\begin{aligned}\left(\sum_{s\in\mathscr{V}} I_u^{V,s}\varphi_s(t)\right)^{\{i\}} &= \sum_{s\in\mathscr{V}} I_u^{\{i\},s}\varphi_s(t)\\ &= \sum_{s\in\mathscr{V}:i\in s} I_u^{\{i\},s}\varphi_s(t)\\ &= \sum_{s\in\mathscr{V}:i\in s} H_{i,s}(t)\mathscr{L}_s(t)\\ &= u_i(t).\end{aligned}$$

□

Now that the input has been decomposed into independent terms, the state $x(t)$ can be similarly decomposed. Let $\zeta_r(t)$ be vectors, of the same dimension as $x^r(t)$, defined by the following dynamics:

$$\begin{aligned}\zeta_r(t+1) &= \textstyle\sum_{s:(s,r)\in\mathscr{E}}\left(A^{r,s}\zeta_s(t)+B^{r,s}\varphi_s(t)\right) && \text{for } r\in\mathscr{V} \text{ with } |r|>1\\ \zeta_{\{i\}}(t+1) &= w_i(t) && \text{for } i=1,\ldots,n\end{aligned} \tag{2.18}$$

with initial conditions $\zeta_s(0)=0$ for all $s\in\mathscr{V}$.

Lemma 4. *The state vector can be decomposed as a sum*

$$x(t)=\sum_{s\in\mathscr{V}} I_x^{V,s}\zeta_s(t), \tag{2.19}$$

where I_x is the identity partitioned into blocks conforming to the partition of x, and $\zeta_s(t)$ is defined by Equations (2.18) with initial condition $\zeta_s(0)=0$. Furthermore, $\zeta_s(t)$ is a linear function of $\mathscr{L}_s(t)$.

Proof. The lemma will be proved by induction. By the initial conditions, $\zeta_s(0)=0$ and $x(0)=0$, Equation (2.19) holds at $t=0$. Furthermore, the property that $\zeta_s(0)$ is a linear function of $\mathscr{L}_s(0)$ trivially holds.

Now, inductively assume that Equation (2.19) holds at time t and that $\zeta_s(t)$ is a linear function of $\mathscr{L}_s(t)$. Plugging Equations (2.19), (2.17) and the update equations for $\zeta_{\{i\}}$, into the dynamic equations shows that $x(t+1)$ is updated as follows:

$$\begin{aligned}x(t+1) &= Ax(t)+Bu(t)+w(t)\\ &= \sum_{s\in\mathscr{V}}\left(AI_x^{V,s}\zeta_s(t)+BI_u^{V,s}\varphi_s(t)\right)+\sum_{i=1}^{n} I_x^{V,\{i\}}\zeta_{\{i\}}(t+1).\end{aligned} \tag{2.20}$$

Now it is claimed that $AI_x^{V,s}=I_x^{V,r}A^{r,s}$, where $r\in\mathscr{V}$ is the unique node such that $(s,r)\in\mathscr{E}$.

Assume that $s = \{j_1, \ldots, j_s\}$ with $j_1 < \cdots < j_s$ and $r = \{i_1, \ldots, i_r\}$ with $i_1 < \cdots < i_r$. The product $AI_x^{V,s}$ is computed as

$$AI_x^{V,s} = \begin{bmatrix} A_{11} & \cdots & A_{1n} \\ \vdots & \ddots & \vdots \\ A_{n1} & \cdots & A_{nn} \end{bmatrix} \begin{bmatrix} I_x^{V,\{j_1\}} & \cdots & I_x^{V,\{j_s\}} \end{bmatrix} = \begin{bmatrix} A_{1j_1} & \cdots & A_{1j_s} \\ \vdots & \ddots & \vdots \\ A_{nj_1} & \cdots & A_{nj_s}. \end{bmatrix} \tag{2.21}$$

On the other hand, the product $I_x^{V,r}A^{r,s}$ is computed as

$$I_x^{V,r}A^{r,s} = \begin{bmatrix} I^{V,\{i_1\}} & \cdots & I^{V,\{i_r\}} \end{bmatrix} \begin{bmatrix} A_{i_1j_1} & \cdots & A_{i_1j_s} \\ \vdots & \ddots & \vdots \\ A_{i_rj_1} & \cdots & A_{i_rj_s} \end{bmatrix}. \tag{2.22}$$

To see that Equations (2.21) and (2.22) give the same result, note that if $i \notin r$, then $A_{ij} = 0$ for all $j \in s$. Indeed, if $j \in s$ then $j \in r$ and for all i such that $(j, i) \in E$, it must be that $i \in r$. Therefore, if $A_{ij} \neq 0$, then $i \in r$. By contrapositive, $i \notin r$ implies that $A_{ij} = 0$.

A similar argument shows that $BI_u^{V,s} = I_x^{V,r}B^{r,s}$.

For all $s \in \mathscr{V}$, define $\eta(s)$ to be the unique node r such that $(s, r) \in \mathscr{E}$. Plugging the identities $AI_x^{V,s} = I_x^{V,\eta(s)}A^{\eta(s),s}$ and $BI_u^{V,s} = I^{V,\eta(s)}B^{\eta(s),s}$ into Equation (2.20) and applying Equation (2.18) to update ζ_r shows that

$$\begin{aligned} x(t+1) &= \sum_{s\in\mathscr{V}} \left(I_x^{V,\eta(s)}A^{\eta(s),s}\zeta_s(t) + I_x^{V,\eta(s)}B^{\eta(s),s}\varphi_s(t) \right) + \sum_{i=1}^{n} I_x^{V,\{i\}}\zeta_{\{i\}}(t+1) \\ &= \sum_{|r|>1} \sum_{s:(s,r)\in\mathscr{E}} I_x^{V,r} \left(A^{r,s}\zeta_s(t) + B^{r,s}\varphi_s(t) \right) + \sum_{i=1}^{n} I_x^{V,\{i\}}\zeta_{\{i\}}(t+1) \\ &= \sum_{|r|>1} I_x^{V,r}\zeta_r(t+1) + \sum_{i=1}^{n} I_x^{V,\{i\}}\zeta_{\{i\}}(t+1) \\ &= \sum_{s\in\mathscr{V}} I_x^{V,s}\zeta_s(t+1). \end{aligned}$$

The only part that remains to be shown is that $\zeta_s(t+1)$ is a linear function of $\mathscr{L}_s(t+1)$. Note that $\zeta_{\{i\}}(t+1) = w_i(t) = \mathscr{L}_{\{i\}}(t+1)$ for $i = 1, \ldots, n$. Consider a node r with $|r| > 1$, and assume that $\zeta_s(t)$ and $\varphi_s(t)$ are linear functions of $\mathscr{L}_s(t)$ for all s such that $(s, r) \in \mathscr{E}$. Then by Equation (2.18), $\zeta_r(t+1)$ must be a linear function of the $\mathscr{L}_s(t)$ terms. By construction, $\mathscr{L}_r(t+1)$ is composed entirely of terms of the form $\mathscr{L}_s(t)$. Thus the result follows. □

2.4.2 General Case: Optimal Solution

The controller gains and steady state cost are defined by propagating the solution to a standard Riccati equation through the information hierarchy graph. Let X_V be the stabilizing solution to the discrete-time algebraic Riccati equation:

$$X_V = S = Q + A^T SA - A^T SB(R + B^T SB)^{-1} B^T SA. \tag{2.23}$$

Define the gain K_V by the standard LQR gain:

$$K_V = (R + B^T X_V B)^{-1} B^T X_V A. \tag{2.24}$$

For $s \neq V$, let $\eta(s)$ be the unique node r such that $(s, r) \in \mathscr{E}$. Assume that $X_{\eta(s)}$ has already been defined and define X_s by

$$\begin{aligned} X_s &= Q^{s,s} + A^{\eta(s),s^T} X_{\eta(s)} A^{\eta(s),s} \\ &\quad - A^{\eta(s),s^T} X_{\eta(s)} B^{\eta(s),s} \left(R^{s,s} + B^{\eta(s),s^T} X_{\eta(s)} B^{\eta(s),s} \right)^{-1} B^{\eta(s),s^T} X_{\eta(s)} A^{\eta(s),s}. \end{aligned} \tag{2.25}$$

Define the gain K_s by

$$K_s = \left(R^{s,s} + B^{\eta(s),s^T} X_{\eta(s)} B^{\eta(s),s} \right)^{-1} B^{\eta(s),s^T} X_{\eta(s)} A^{\eta(s),s}. \tag{2.26}$$

Theorem 2. *The optimal controller for the general problem defined in Section 2.2 is given by*

$$u(t) = -\sum_{s \in \mathscr{V}} I_u^{V,s} K_s \zeta_s(t),$$

and the steady state cost is given by

$$\sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}).$$

Here K_s and X_s are defined by Equations (2.23)–(2.26).

As with the two-player problem, the proof will consist of a finite-horizon derivation and a limiting argument to prove the final result.

2.4.3 General Case: Finite-Horizon Derivation

Assume that the optimal expected cost-to-go function is of the form $\mathbb{E}[J(\zeta, t)]$. Recalling the finite-horizon cost function and plugging in the state decomposition of Equation (2.19), $\mathbb{E}[J(\zeta, N)]$ is given

by

$$\begin{aligned}\mathbb{E}[J(\zeta,N)] &= \mathbb{E}\left[x^T\Lambda x\right] \\ &= \mathbb{E}\left[\left(\sum_{s\in\mathscr{V}} I_x^{V,s}\zeta_s\right)^T \Lambda \left(\sum_{s\in\mathscr{V}} I_x^{V,s}\zeta_s\right)\right] \\ &= \sum_{s\in\mathscr{V}} \mathbb{E}\left[\zeta_s^T\Lambda^{s,s}\zeta_s\right].\end{aligned}$$

The last equality follows from the pairwise independence of ζ_s.

Set $X_s(N) = \Lambda^{s,s}$ for all $s \in \mathscr{V}$ and define $J(\zeta, N)$ to be $J(\zeta,N) = \sum_{s\in\mathscr{V}} \zeta_s^T X_s(N)\zeta_s$. Inductively assume that for some $t+1 \le N$, $J(\zeta, t+1)$ is defined by

$$J(\zeta,t+1) = \sum_{s\in\mathscr{V}} \zeta_s^T X_s(t+1)\zeta_s + \sum_{k=t+2}^{N}\sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}(k)). \tag{2.27}$$

The optimal expected cost-to-go function at time t is computed by solving the Bellman equation:

$$\mathbb{E}[J(\zeta,t)] = \min_{\varphi} \mathbb{E}\left[x^TQx + u^TRu + J(\zeta',t+1)\right], \tag{2.28}$$

where ζ_s' are the variables ζ_s, updated according to Equation (2.18).

Substituting the decompositions for x and u shows that the first two terms on the right-hand side can be decoupled as follows:

$$\begin{aligned}\mathbb{E}\left[x^TQx + u^TRu\right] &= \mathbb{E}\left[\left(\sum_{s\in\mathscr{V}} I_x^{V,s}\zeta_s\right)^T Q \left(\sum_{s\in\mathscr{V}} I_x^{V,s}\zeta_s\right) + \left(\sum_{s\in\mathscr{V}} I_u^{V,s}\varphi_s\right)^T R \left(\sum_{s\in\mathscr{V}} I_u^{V,s}\varphi_s\right)\right] \\ &= \sum_{s\in\mathscr{V}} \left[\zeta_s^T Q^{s,s}\zeta_s + \varphi_s^T R^{s,s}\varphi_s\right] \end{aligned} \tag{2.29}$$

Combining Equations (2.18) and (2.27) shows that $\mathbb{E}[J(\zeta', t+1)]$ can be expanded as

$$\begin{aligned}\mathbb{E}[J(\zeta',t+1)] &= \sum_{|r|>1} \mathbb{E}\left[\left(\sum_{s:(s,r)\in\mathscr{E}} (A^{r,s}\zeta_s + B^{r,s}\varphi_s)\right)^T X_r(t+1)\left(\sum_{s:(s,r)\in\mathscr{E}} (A^{r,s}\zeta_s + B^{r,s}\varphi_s)\right)\right] \\ &\quad + \sum_{i=1}^{n} \mathbb{E}\left[w_i^T X_{\{i\}}(t+1) w_i\right] + \sum_{k=t+2}^{N}\sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}(k)) \\ &= \sum_{s\in\mathscr{V}} \mathbb{E}\left[(A^{\eta(s),s}\zeta_s + B^{\eta(s),s}\varphi_s)^T X_{\eta(s)}(t+1)(A^{\eta(s),s}\zeta_s + B^{\eta(s),s}\varphi_s)\right] \\ &\quad + \sum_{k=t+1}^{N}\sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}(k)) \end{aligned} \tag{2.30}$$

Combining Equations (2.29) and (2.30) shows that the right-hand side of the Bellman equation

can be decomposed into a sum of independent terms, plus a constant term:

$$\min_{\varphi} \mathbb{E}\left[x^T Q x + u^T R u + J(\zeta', t+1)\right] =$$
$$\sum_{s\in\mathscr{V}} \min_{\varphi_s} \mathbb{E}\left[\zeta_s^T Q^{s,s}\zeta_s + \varphi_s^T R^{s,s}\varphi_s + (A^{\eta(s),s}\zeta_s + B^{\eta(s),s}\varphi_s)^T X_{\eta(s)}(t+1)(A^{\eta(s),s}\zeta_s + B^{\eta(s),s}\varphi_s)\right]$$
$$+ \sum_{k=t+1}^{N}\sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}(k)).$$

Standard quadratic minimization arguments show that the optimal inputs are given by

$$\varphi_s = -K_s(t)\zeta_s$$

with gains $K_s(t)$ computed as

$$K_s(t) = \left(R^{s,s} + {B^{\eta(s),s}}^T X_{\eta(s)}(t+1)B^{\eta(s),s}\right)^{-1} {B^{\eta(s),s}}^T X_{\eta(s)}(t+1)A^{\eta(s),s}.$$

Plugging in the inputs $\varphi_s = -K_s(t)\zeta_s(t)$ shows that $J(\zeta, t)$ has the form

$$J(\zeta, t) = \sum_{s\in\mathscr{V}} \zeta_s^T X_s(t)\zeta_s + \sum_{k=t+1}^{N}\sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}(k))$$

where the matrices $X_s(t)$ are computed as follows (denoting $X_{\eta(s)}(t+1)$ by $X'_{\eta(s)}$ to save space):

$$\begin{aligned} X_s(t) &= Q^{s,s} + {A^{\eta(s),s}}^T X'_{\eta(s)} A^{\eta(s),s} \\ &\quad - {A^{\eta(s),s}}^T X'_{\eta(s)} B^{\eta(s),s}\left(R^{s,s} + {B^{\eta(s),s}}^T X'_{\eta(s)} B^{\eta(s),s}\right)^{-1} {B^{\eta(s),s}}^T X'_{\eta(s)} A^{\eta(s),s}. \end{aligned}$$

Since $\mathbb{E}[J(\zeta, t+1)]$ was the optimal expected cost-to-go at time $t+1$, it follows inductively that $\mathbb{E}[J(\zeta, t)]$ is the optimal expected cost-to-go at time t, and the form of $J(\zeta, t)$ is valid for all $t \leq N$. Finally, since $x(0) = 0$, the total cost is calculated to be

$$\sum_{t=1}^{N}\sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}(t)).$$

2.4.4 General Case: Steady State

Since $X_V(t)$ is just the solution to the centralized LQR Riccati equation, as $N \to \infty$, $X_V(t) \to X_V = S$, the stabilizing solution of the algebraic Riccati equation. Since all the other matrices, $K_V(t)$, $X_s(t)$, and $K_s(t)$, are specified by $X_V(t)$, they respectively converge to the matrices K_V, X_s, and K_s, as defined in Equations (2.24), (2.25), and (2.26), as $X_V(t) \to X_V$. Thus, the optimal gains and Riccati solutions have been found.

The steady state cost is calculated by noting that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{t=1}^{N} \sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}(t)) = \sum_{i=1}^{n} \mathrm{Tr}(W_i X_{\{i\}}).$$

2.5 Conclusion

This chapter presents Riccati-based solutions for a class of decentralized linear control problems with communication delays. The structure of the controllers is dictated by a decomposition of the information based on a specially constructed graph, referred to as the information hierarchy graph. The controllers can be interpreted as simple management schemes. In these schemes, a top level "executive" generates an input, based on delayed global information. The input is modified using newer, more localized information as it gets passed down the chain of command.

In the case of the simple two-player architecture, the optimal control scheme is compared to centralized controllers, both with and without delays. It is found that the performance is always at least as good as centralized control with a single-step delay, but can never be better than centralized control. Explicit comparisons of the costs are given.

This chapter is intended to serve as groundwork for studying the connections between motor control and distributed control. Many theoretical and biological questions can be posed to follow up.

On the biological end, it would be interesting to explore the use of feedforward signaling between motor regions [37, 38, 39] in terms of partially nested systems. Lesion studies could lend insight into the distributed architecture of the motor system. Finally, feedback and feedforward processing in the spinal cord must be studied in greater depth.

A simple next step would be to utilize the algorithmic solution to the general problem of this chapter to study biologically motivated control problems with complex delay structures. Using the relatively straightforward construction of information hierarchy graph, the optimal control hierarchies can be extracted and compared to the existing control schemes.

The results in this chapter relied on state feedback, but the associated output feedback problem must be solved. Biological sensors may be noisy, and measurements of all states may not be available. Luckily, Rantzer's method for this chapter's problem extends naturally to output feedback [19]. As in the state feedback case, the controller structure is not apparent from the current solution. The challenge will be to see if a similar hierarchical structure emerges via a different derivation.

Recently, it has been shown that humans display risk-sensitive control policies [28]. Future work will attempt to derive related distributed optimal control laws for risk-sensitive cost functions [40] and other biologically motivated cost functions.

Finally, connections to social sciences should be explored. The mathematical work in this chapter

originated in management science in the study of teams [41]. Explicit structures for nontrivial decentralized optimal controllers, as described in this chapter and similar works [20, 22], have been discovered only recently. The structures arising mathematically typically have simple management interpretations. Thus, it would be interesting to see how well the theory predicts management structures of real organization.

2.6 Appendix to Chapter 2

The control problem studied in this chapter, however, has constraints on the input from Equation (2.4) that cannot be handled with centralized control methods. The only assumption about the controllers made in the derivations is that they are linear. The objective of this appendix is to justify the linearity assumption.

In the 1960s, Witsenhausen showed that some decentralized linear quadratic Gaussian control problems have nonlinear optimal controllers [13]. Not long after, Ho and Chu defined a class of decentralized control systems, termed partially nested systems, whose optimal controllers are linear [12]. This chapter will give a brief introduction to the theory of partially nested systems. It will be shown that the systems from this chapter are partially nested (Lemma 5) and thus admit linear optimal controllers.

This appendix just presents the basic definitions and results required in the chapter. Ho and Chu's original paper gives a short and readable introduction to the theory of partially nested systems [12].

2.6.1 General Form of Finite Horizon LQG Control

Let $\xi \in \mathbb{R}^m$ be a zero-mean Gaussian random variable ξ with covariance X. Let $Q \in \mathbb{R}^{p\times p}$ be a positive definite matrix and let $S \in \mathbb{R}^{p\times m}$. The basic optimization problem is to choose an input $u = [u_1^T, \ldots, u_n^T]^T \in \mathbb{R}^p$, with $u_i \in \mathbb{R}^{p_i}$, to minimize the quadratic cost

$$\mathbb{E}\left[\frac{1}{2}u^T Q u + u^T S\xi\right], \tag{2.31}$$

subject to the constraint that each input u_i is a (Borel-measurable) function of its available measurement:

$$z_i = H_i\xi + \sum_j D_{ij}u_j \in \mathbb{R}^{q_i}. \tag{2.32}$$

More precisely, a control policy $\gamma(z)$ is said to be *admissible* if

$$\gamma(z) = \begin{bmatrix} \gamma_1(z_1) \\ \vdots \\ \gamma_n(z_n) \end{bmatrix}, \tag{2.33}$$

where each $\gamma_i : \mathbb{R}^{q_i} \to \mathbb{R}^{p_i}$ is a Borel-measurable function. Then the optimization problem is to find measurable functions $\gamma_i^* : \mathbb{R}^{q_i} \to \mathbb{R}^{p_i}$ such that

$$\mathbb{E}\left[\frac{1}{2}\gamma(z)^T Q\gamma(z) + \gamma(z)^T S\xi\right] \geq \mathbb{E}\left[\frac{1}{2}\gamma^*(z)^T Q\gamma^*(z) + \gamma^*(z)^T S\xi\right] \text{ for all admissible } \gamma.$$

Definition 1. The form of the outputs, Equation (2.32), is called an *information structure*.

It will be assumed that each player knows all the problem data H_i, D_{ij}, Q, and S. Furthermore, when a control policy γ is chosen, it is assumed that all players know the functions γ_i.

***Remark* 4.** Note that the finite horizon problem from Equations (2.2), (2.4), and (2.5) can be reduced to the problem described above. The noise $w(0 : N-1)$ plays the role of ξ. By plugging in the dynamics, the state can be computed as

$$x(t) = \sum_{k=0}^{t-1} A^{t-1-k} w(k) + \sum_{k=0}^{t-1} A^{t-1-k} Bu(k).$$

Substituting the value of $x(t)$ into Equation (2.5) and ignoring the terms that do not depend on u, this problem can be put in the form of (2.31), (2.32), and (2.33).

2.6.2 Partially Nested Information Structures

This subsection describes a special class of information structures, known as partially nested information structures. First, a few remarks on a simpler class of information structures, known as static information structures are given. Problems with static information structures have linear optimal controllers. The key idea of [12] is that problems with partially nested information structures can be reduced to problems with static information structures.

Definition 2. If $D_{ij} = 0$ for all i and j in Equation (2.32), then the information structure is called *static*.

Theorem 3 ([42])**.** *If a problem defined by Equations (2.31), (2.32), and (2.33) has a static information structure, then it has a unique optimal solution which is linear. In other words, there are matrices K_i such that $u_i = K_i z_i$ is optimal.*

The main idea behind the definition of partially nested systems is that if player i could deduce

u_j whenever $D_{ij} \neq 0$, then it could subtract off the effects of u_j, leaving the measurement associated with the static information structure:

$$\tilde{z}_i = z_i - \sum_j D_{ij} u_j = H_i \xi.$$

Then the optimization problem can be solved as though the information structure were static.

Now, the notion of what it means for player i to be able to deduce u_j must be formalized. Note how in Equation (2.32) input u_j can affect the measurement z_i if $D_{ij} \neq 0$. Likewise, if $D_{jk} \neq 0$, then u_k can affect measurement z_j. Since u_j depends on z_j, which depends on u_k, it follows that u_k can affect z_i, as well. This flow of influence is captured by a graph called a precedence diagram.

Definition 3. The *precedence diagram* for the information structure defined by Equation (2.32) is a directed graph with nodes $\{1, \ldots, n\}$ such that there is an edge (j, i) for each i and j such that $D_{ij} \neq 0$.

Note that u_k can influence measurement z_i if and only if there is a path from node k to node i in the precedence diagram. It will be assumed that the precedence diagram is acyclic.

Let the random variable $\xi \in \mathbb{R}^n$ be defined on the probability space $(\mathbb{R}^n, \mathcal{F}, P)$. Here $\mathcal{F}$ is a σ-algebra and P is a probability measure. If γ is an admissible controller then the assumption that the precedence diagram is acyclic implies that the measurement

$$z_i = H_i \xi + \sum_j D_{ij} \gamma_j(z_j) \tag{2.34}$$

is a measurable function from $\mathbb{R}^m$ to $\mathbb{R}^{q_i}$. In other words, once γ is fixed, the measurement z_i becomes a function of ξ. Indeed, if i has no incoming edges in the precedence diagram (which must hold for some node, since it is acyclic), then $z_i = H_i \xi$. Now, inductively assume that z_j is a measurable function of ξ for j such that $D_{ij} \neq 0$. Then measurability of γ_j and Equation (2.34) imply that z_i is a measurable function of ξ.

The map, $\xi \mapsto z_i$, induces a subalgebra $\mathcal{Z}_i \subset \mathcal{F}$. With the notions of the influence diagram and the induced subalgebras, partial nestedness can finally be defined.

Definition 4. An information structure is partially nested if $\mathcal{Z}_j \subset \mathcal{Z}_i$ for all admissible controllers and all j and i such that there is a path from j to i in the precedence diagram.

The intuitive meaning of Definition 4 is that whenever input u_j can influence measurement z_i, then z_j can be deduced from z_i. Then, since $u_j = \gamma_j(z_j)$, and player i knows γ_j, player i can deduce u_j. In particular, if the information structure is partially nested and $D_{ij} \neq 0$, then player i can deduce z_j and thus u_j.

The main theorem on partially nested systems and its application to the problem in this chapter are now stated.

Theorem 4 ([12])**.** *If the problem defined by Equations (2.31) and (2.32) has a partially nested information structure, then the optimal controller is linear. In other words, there are matrices K_i such that $u_i = K_i z_i$ is optimal.*

Lemma 5. *The information structure defined by Equation (2.4) is partially nested.*

Proof. Recall that $u_i(t) = \gamma_{i,t}(x_1(0 : t - d_{1i}), \ldots, x_n(0 : t - d_{ni}))$. Note that $u_j(t - d_{ji} - 1)$ is the newest input from player j that can affect $u_i(t)$. It is claimed that at time t, player i has access to all the information that player j had at time $t - d_{ji} - 1$. Indeed, note that $u_j(t - d_{ji} - 1)$ is given by

$$u_j(t - d_{ji} - 1) = \gamma_{j,t-d_{ji}-1}(x_1(0 : t - d_{ji} - 1 - d_{1j}), \ldots, x_n(0 : t - d_{ji} - 1 - d_{nj})).$$

Now for any k, the information about x_k available to player i at time t is $x_k(0 : t - d_{ki})$ while the information about x_k available to player j at time $t - d_{ji} - 1$ is $x_k(0 : t - d_{ji} - 1 - d_{kj})$. From the definition of d_{ji} as the length of the shortest path from node j to node i, it follows that

$$d_{ki} \leq d_{kj} + d_{ji} < d_{kj} + d_{ji} + 1.$$

Thus the claim follows, and the proof is complete. □

Chapter 3

Control Over Spiking Neuron Channels

3.1 Introduction

This chapter continues the study of feedback in the nervous system, focusing on the most notable aspect of the dynamics of individual neurons: action potentials. As discussed in the Chapter 1, an action potential is a short-lived voltage spike that occurs as a response to an input current (Figure 1.5).

Similar to Chapter 2, this chapter aims to explore the connections between one research trend in neuroscience and another in control theory. Whereas Chapter 2 focused on the use of optimal feedback control in the motor system and distributed control, the goal of this chapter is to find connections between the study of spiking neurons as communications channels and the study of communication channels within feedback loops.

Neurons are the high-speed communication channels of the body. Not long after Shannon's 1948 paper [43], physiologists began to study information theoretic properties of neurons [44]. Much later, in the early 1990s, Bialek et al.'s seminal paper [1] set off a wave of research on the connections between information theory and neuroscience [45, 46, 47, 48].

Meanwhile, networked control systems, control systems in which the plant and controller communicate over a network, began to gain attention in the late 1990s. Control theorists increasingly studied the effects of control over noisy or data-limited communication channels [49, 50, 51, 52, 53, 54, 55, 56]. One of the primary goals in networked control research has been determining the amount of information that must be sent across a communication channel in order to guarantee stability. Most of the works mentioned deal with discrete-time systems, and thus they abstract away some of the difficulties of sending data in real-time. Physical systems operate in continuous time, but when digital control is used, control signals can only be applied at discrete time increments. To address communication-limited control in continuous time, researchers also began to study the maximum amount of time

between control inputs that systems could tolerate while maintaining stability [57, 58, 59, 60]. See [61] and [62] for surveys.

While neurons are studied as communication channels, and communication channels are studied within control loops, neurons are quite different from the channels typically studied in engineering. A common assumption in engineering is that communication occurs at periodically sampled time instants. In the case that communication and control signals are generated by a digital computer, the periodic sampling assumption is reasonable. In neural control, there is no periodic sampling. A neuron sends a spike signal when the voltage across the cell membrane reaches a certain threshold. Thus the control scheme is event-triggered, as in [59, 60]. Traditional networked control methods require that the information transmitted across the communication network be packets of (possibly quantized) numerical data. In neural control, the basic unit of communication is the spike, which conveys minimal numerical information. Instead, information is conveyed in the timing and the rate of spiking of neurons [8].

This chapter studies continuous-time networked control in which signals to and from the plant pass through a novel communication channel, termed the *spike channel*, that is modeled after spiking neurons. In particular, the channel was designed based on methods for reconstructing the input current to neurons by applying a linear filter to their spike sequences [1, 45, 63]. The spike channel (Figure 3.1) operates through the following sequence of events. A continuous-time input is sent through a low-pass filter. When the state of the filter reaches a threshold value (either positive or negative) the state of the filter is set to zero, and a delta function of a fixed magnitude and appropriate sign is sent to an identical low-pass filter. Then the output of the channel is the output of the second filter. Surprisingly, even with this strong spiking nonlinearity, the spike channel behaves like a low-pass filter, up to a bounded additive disturbance (Figures 3.5 and 3.6).

The main results of this chapter describe stability, tracking performance, and data rate for feedback control of a continuous-time linear time invariant (LTI) single-input single-output (SISO) plant by a continuous-time LTI SISO controller when signals to and from the plant pass through the spike channel (Figures 3.2 and 3.3). In particular, if the nominal feedback loop is internally stable, then spike channel parameters can be chosen so that internal stability is preserved. By appropriate choice of parameters, tracking performance in the spike channel system can be made arbitrarily close to that of the nominal system. The price of good tracking and stability properties manifests in higher spike rates.

3.2 Preliminaries

This section defines the notation used throughout the chapter. Next, it describes the problem of interest and gives a neurobiological motivation for the communication channel introduced. Finally

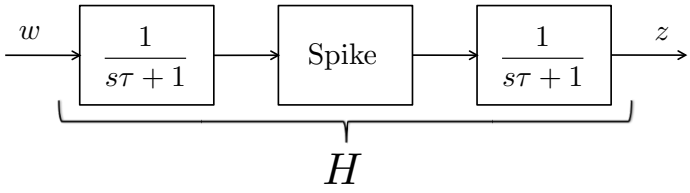

Figure 3.1: The spike channel consists of a first-order low-pass filter, followed by a spiking nonlinearity, which is, in turn, followed by a low-pass filter identical to the first filter. The channel is denoted by H.

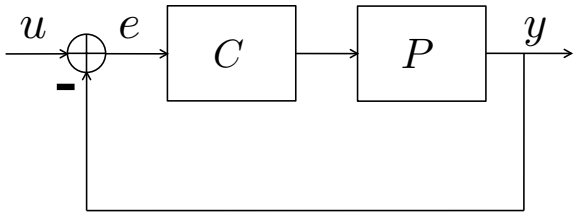

Figure 3.2: The nominal feedback loop. Here C is a continuous-time proper SISO transfer function and P is a continuous-time strictly proper SISO transfer function.

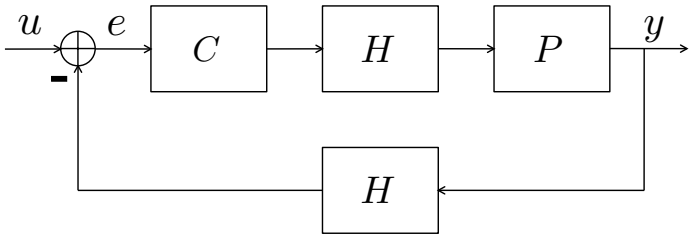

Figure 3.3: A standard linear feedback loop modified so that signals to and from the plant must pass through a nonlinear communication channel H (Figure 3.1)

the communication channel of interest is defined.

3.2.1 Notation

The real numbers are denoted by $\mathbb{R}$. For a function, $x : \mathbb{R} \to \mathbb{R}$, the L_1 and L_∞ norms are denoted by $\|x\|_1$ and $\|x\|_\infty$, respectively. Let $x(t^-) = \lim_{s \uparrow t} x(s)$ and $x(t^+) = \lim_{s \downarrow t} x(s)$ if the limits exist. The unit imaginary number is denoted by $j = \sqrt{-1}$.

3.2.2 Problem Formulation

This chapter studies a networked control system in which communications to and from the plant occur via the spike channel. Consider the feedback loop depicted in Figure 3.2. As is standard in classical control theory, P is a strictly proper continuous-time SISO transfer function, and C is a proper continuous time SISO transfer function. As is common in networked control, the plant is assumed to be separated from the controller, and thus all signals to and from the plant must pass through communication channels (Figure 3.3). The particular communication channel used is the spike channel, which is formally defined in Subsection 3.2.4. Results on stability, tracking performance, and data rate of this networked control system are studied in Section 3.3.

3.2.3 Neurobiological Motivation

The channel studied in this chapter is motivated by research on neural decoding [45, 1, 63]. While it is widely accepted that information is encoded in the spike rate of neurons (the number of spikes over a given interval), it is less clear how much information is conveyed by a small number of spikes. Bialek et al. provided evidence that neural decisions may result from only a few spikes [1]. They studied a motion sensitive neuron in bowflies, called H1, which has a maximum spike rate between 100 and 200 spikes per second. However, given that bowflies can make course corrections to visual stimuli in about 30 ms, in the time span of a course correction, H1 can only send between three and six spikes. To examine how accurately input signals could be reconstructed from sequences of action potentials, called *spike trains*, the authors proposed a linear filtering scheme which could reconstruct input stimuli on the basis of only a few spikes.

A schematic of the reconstruction method from [1] is shown in Figure 3.4. (See also [45, 63].) On the left is the input signal. The fly has two H1 neurons on either side of its head, sensitive to positive and negative stimuli, respectively. This is depicted by two neurons with opposite sensitivities in the center. The graph on the right shows the reconstruction (red) of the input signal (blue) based on linear filtering of the two spike trains (below).

The work of Bialek et al. demonstrated that simple linear filtering techniques can give accurate input reconstructions, even if spikes are sparsely distributed. The rest of this chapter theoretically

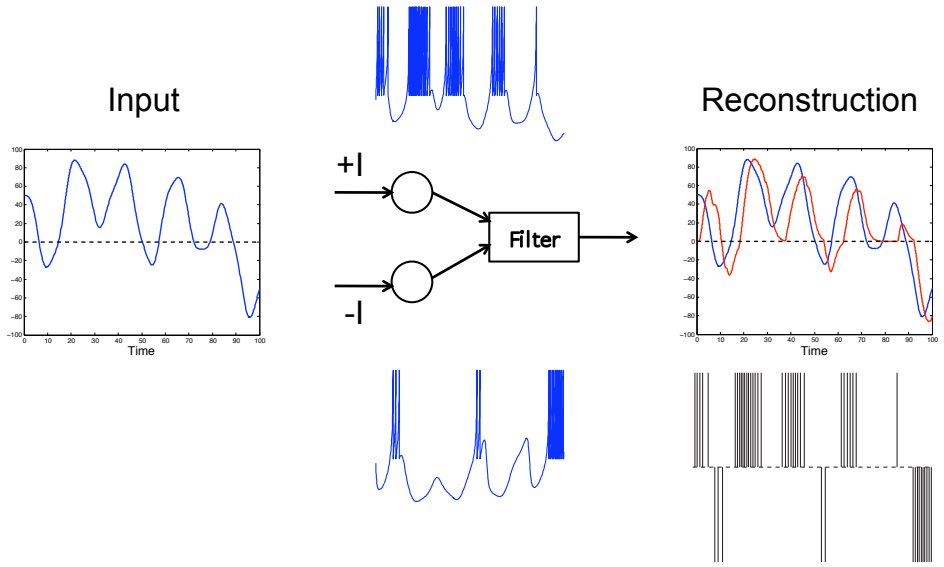

Figure 3.4: A simulation of the experiment from [1]. An input is passed to opposing neurons, with the top neuron sensitive to $+I$, and the bottom neuron sensitive to $-I$. The input signal is reconstructed by a linear filter. Note that the reconstruction process captures transient behavior on the basis of only one or two spikes. For a more realistic simulation in the figure, the neural model from [2] is used in place of the leaky integrate-and-fire model studied in this chapter.

investigates the application of linear spike train reconstruction when a neuron-like channel is used for communication.

To study the reconstruction process formally, a neuron model must be specified. A common model for a single neuron, known as the leaky integrate-and-fire model, is given by the following modified RC-circuit equations:

$$\begin{aligned} \dot{V}(t) &= -\tfrac{1}{RC}V(t) + \tfrac{1}{C}I(t) && \text{if } V(t) < V_{th} \\ V(t^+) &= 0 && \text{if } V(t^-) = V_{th} \\ V_{\text{out}}(t) &= \textstyle\sum_{\{\tilde{t} \le t : V(\tilde{t}^-) = V_{th}\}} \delta(t - \tilde{t}). \end{aligned} \tag{3.1}$$

The state variable V corresponds to the electrical potential across the cell membrane of the neuron, and I is an input current. Thus, the neuron integrates the potential with a leak term proportional to $-\frac{1}{RC}V$. When the potential reaches some threshold, the electrical potential across the membrane is set to zero and a spike (approximated as a delta function) is sent as an output.

The leaky integrate-and-fire model can be viewed as an input-output mapping M_{LIF} that takes an input current I and outputs a sequence of spikes V_{out}. Note that the operator M_{LIF} only generates outputs for positive currents. Just as the experiment of Bialek et al. [1] used opposing neurons to sense positive and negative stimuli, information about positive and negative input current can be obtained by examining opposing leaky integrate-and-fire neurons, $M_{\text{LIF}}(I) - M_{\text{LIF}}(-I)$. The

reconstruction technique of [1, 45] applies equally well to simulated neurons [63]. Thus, a linear filter G can be constructed such that, for well behaved input signals, $G(M_{\text{LIF}}(I) - M_{\text{LIF}}(-I)) \approx I$.

***Remark* 5.** Many treatments of leaky integrate-and-fire neurons include an extra dynamical mode, called a refractory period. If a spike occurs at time t, then $V(t^+) = 0$, and in models with a refractory period, V is held constant at 0 for the interval $(t, t + \tau_{ref}]$. In the model above, which has no refractory period, the membrane begins to integrate current immediately after time t. While including a refractory period makes the neuron model more biologically realistic, this chapter neglects refractory periods in the interest of analytical simplicity. See [8, 63, 64] for more on neuron models with refractory periods.

3.2.4 The Spike Channel

The spike channel is defined is an input-output mapping H taking input $w(t)$ to output $z(t)$ based on the following rules:

$$\begin{aligned} \dot{x}_1(t) &= -\tfrac{1}{\tau}x_1(t) + \tfrac{1}{\tau}w(t) \\ \dot{x}_2(t) &= -\tfrac{1}{\tau}x_2(t) \end{aligned} \qquad \text{if } |x_1(t)| < r \tag{3.2}$$

$$\begin{aligned} x_1(t^+) &= 0 \\ x_2(t^+) &= x_2(t^-) + x_1(t^-) \end{aligned} \qquad \text{if } |x_1(t^-)| = r \tag{3.3}$$

$$z(t) = x_2(t). \tag{3.4}$$

So, the spike channel consists of two first-order low-pass filters with equal time constant τ, such that when the magnitude of the state of the filter reaches the threshold, r, it is immediately set to zero and a delta function (or "spike") of magnitude $r\tau$ is applied to the second filter (Figure 3.1).

To understand the spike behavior more explicitly, assume that a spike occurs at time t. Then the second line of equation (3.3) can be viewed as the application of a delta function as follows:

$$x_2(t^+) = x_2(t^-) + \frac{1}{\tau}\int_{-\infty}^{t} e^{-(t-\sigma)/\tau}\tau x_1(t^-)\delta(\sigma - t)d\sigma.$$

Since $|x_1(t^-)| = r$, the spike has magnitude $r\tau$, and it has the same sign as $x_1(t^-)$.

The threshold r is analogous to V_{th} in the leaky integrate-and-fire model, Equation (3.1). The time constant τ plays the role of RC. In actual neuron models, these parameters would be set based on biological considerations. In the current chapter, they are merely viewed as channel parameters that can affect the stability and performance of the corresponding feedback loop. The input w plays the role of the input current I, and z can be thought of as the approximately reconstructed input. The special form of the spike channel leads to a straightforward quantitative analysis showing that H behaves like a low-pass filter, up to a bounded additive disturbance (Figures 3.5 and 3.6).

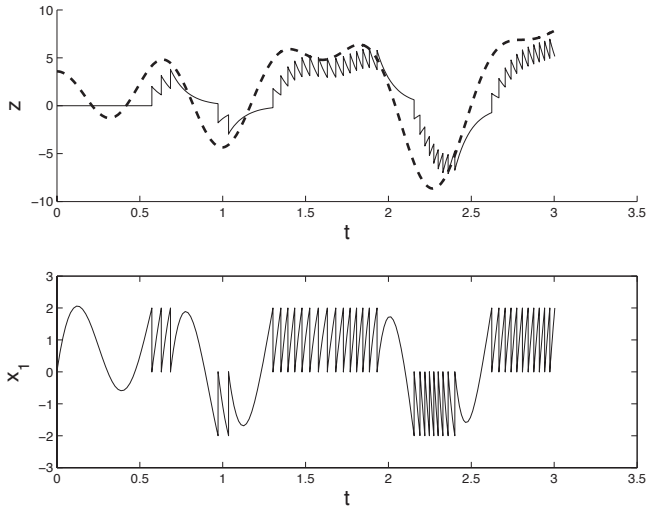

Figure 3.5: Response of the spike channel to a randomly generated input. The parameters were set to $\tau = 0.1$ and $r = 2$. Top. The input is the dashed line and the output is the solid line. Bottom. The signal x_1

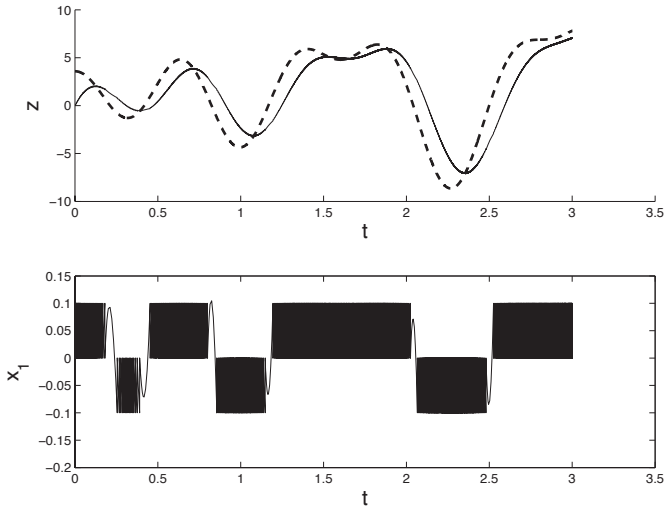

Figure 3.6: Response of the spike channel to input from Figure 3.5 with parameters set to $\tau = 0.1$ and $r = 0.1$. Note that with r smaller than in Figure 3.5, the output is smoother and tracks the input more accurately at the expense of a higher spike rate. Top. The dashed line is the input while the solid line is the output. Bottom. The signal x_1

3.3 Results

This section presents some key lemmas about the spike channel and uses them to derive the main results about feedback control with communication between the controller and plant occurring over spike channels.

3.3.1 Spike Channel Lemmas

This subsection presents two preliminary results about the spike channel that facilitate the analysis of spike-channel-based feedback schemes. The first lemma gives a bound on the *spike rate* (the number of spikes per unit time) based on the threshold r, the time constant τ, and the size of the input to the channel. The second lemma (Lemma 7), the most important preliminary result for this chapter, shows that the output of the spike channel differs from the output of a low-pass filter by at most r.

Lemma 6. *If $\|w\|_\infty = m$, then the spike rate is bounded above by*

$$f(m) = \begin{cases} \frac{1}{\tau \ln \frac{m}{m-r}} & \text{if } m > r \\ 0 & \text{if } m \leq r. \end{cases}$$

Furthermore, $f(m) \leq \frac{m}{r\tau}$ for all $m \geq 0$, and $\lim_{m\to\infty} \frac{f(m)}{m} = \frac{1}{r\tau}$.

Lemma 6 is proved in the appendix of this chapter (Section 3.5).

Lemma 7. *If w is bounded and $y(t) = \frac{1}{\tau}\int_{-\infty}^{t} e^{-(t-\sigma)/\tau} w(\sigma)d\sigma$ is the output of a first-order low-pass filter with time constant τ, then*

$$y(t) = x_1(t) + x_2(t).$$

In particular, $|y(t) - z(t)| \leq r$, for all $t \in \mathbb{R}$.

Proof. To find an expression for $x_1(t) + x_2(t)$, it is useful to have expressions for each term, individually. Two cases arise: either an infinite or finite number of spikes have occurred up to time t. Only the infinite spike case will be proven since the finite case is similar. Let $\ldots < t_{-2} < t_{-1} < t_0 \leq t$ be the times at which spikes occurred, up prior to time t.

Since x_1 is reset to 0 at time t_0, $x_1(t)$ is calculated to be

$$x_1(t) = \frac{1}{\tau}\int_{t_0}^{t} e^{-(t-\sigma)/\tau} w(\sigma)d\sigma. \tag{3.5}$$

On the other hand, from the spiking behavior defined in Equation (3.3), the input to the second filter is

$$\tau \sum_{k \leq 0} x_1(t_k^-)\delta(t - t_k).$$

Therefore, the output of the second filter is

$$\begin{aligned} x_2(t) &= \frac{1}{\tau}\int_{-\infty}^{t} e^{-(t-\sigma)/\tau}\tau \sum_{k\leq 0} x_1(t_k^-)\delta(\sigma - t_k)d\sigma \\ &= \sum_{k\leq 0} e^{-(t-t_k)/\tau}x_1(t_k^-). \end{aligned} \tag{3.6}$$

To make the expression for $x_2(t)$ independent of $x_1(t_k^-)$, the summand of equation (3.6) can be expanded as

$$\begin{aligned} & e^{-(t-t_k)/\tau}x_1(t_k^-) \\ &= e^{-(t-t_k)/\tau}\frac{1}{\tau}\int_{t_{k-1}}^{t_k} e^{-(t_k-\sigma)/\tau}w(\sigma)d\sigma \\ &= \frac{1}{\tau}\int_{t_{k-1}}^{t_k} e^{-(t-\sigma)/\tau}w(\sigma)d\sigma. \end{aligned} \tag{3.7}$$

Finally, combining equations (3.5)–(3.7) gives

$$\begin{aligned} x_1(t) + x_2(t) &= \frac{1}{\tau}\int_{t_0}^{t} e^{-(t-\sigma)/\tau}w(\sigma)d\sigma + \\ &\qquad \frac{1}{\tau}\sum_{k\leq 0}\int_{t_{k-1}}^{t_k} e^{-(t-\sigma)/\tau}w(\sigma)d\sigma \\ &= \frac{1}{\tau}\int_{-\infty}^{t} e^{-(t-\sigma)/\tau}w(\sigma)d\sigma. \end{aligned}$$

The second equality follows from the fact that $t_k - t_{k-1} \geq \frac{r\tau}{\|w\|_\infty} > 0$ (by Lemma 6), and thus $\lim_{k\to-\infty} t_k = -\infty$. Note that since spikes occur, it must be that $\|w\|_\infty > 0$. □

3.3.2 Main Results

Lemma 7 implies that the spike channel can be conservatively approximated by a low-pass filter followed by an additive disturbance which is bounded by r (Figure 3.7). By studying the feedback loop depicted in Figure 3.7, results about the stability, tracking, and data rate (because it is bounded by signal size) of feedback loop with spike channels can be inferred.

A feedback loop is said to be *internally stable* if whenever all the inputs (including disturbances) are bounded, then all the signals in the loop are bounded. The first result states that for small enough τ, the spike channel preserves internal stability of the feedback loop.

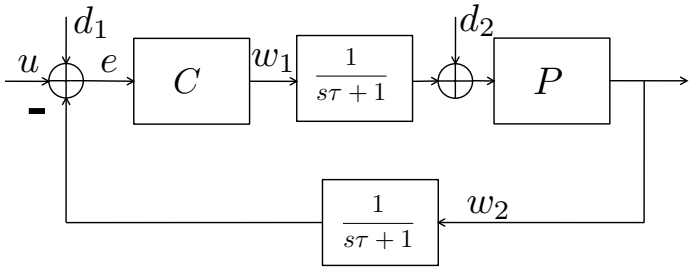

Figure 3.7: The overapproximation of the feedback loop in Figure 3.3 obtained by replacing the spike channel with a low-pass filter and additive disturbance. Note that the disturbances are bounded as $|d_i(t)| \leq r$ for all $t \in \mathbb{R}$.

Theorem 5. *If the nominal feedback loop from Figure 3.2 is internally stable, then there exists $T > 0$ such that for all $\tau \in (0, T]$, the feedback loop with spike channels from Figure 3.3 is internally stable.*

Theorem 5 follows immediately from Lemma 7 and the following lemma (which is proved in the appendix of this chapter, Section 3.5).

Lemma 8. *If the nominal feedback loop from Figure 3.2 is internally stable, then there exists $T > 0$ such that the disturbance feedback loop from Figure 3.7 is internally stable for all τ such that $0 \leq \tau \leq T$.*

The next result states that by choosing r and τ small enough, the tracking error does not significantly degrade. Let G^{τ}_{eu} be the impulse response of the mapping from input u to tracking error e for the system in Figure 3.7. Define $G^{\tau}_{ed_1}$ and $G^{\tau}_{ed_2}$ similarly. Let G^{nom}_{eu} be the nominal mapping from input u to tracking error e.

Theorem 6. *If e_{nom} is the nominal tracking error, e is the tracking error from the spike channel feedback loop and τ is such that the disturbance loop from Figure 3.7 is internally stable, then*

$$\|e_{\text{nom}} - e\|_{\infty} \leq \|G^{\tau}_{eu} - G^{\text{nom}}_{eu}\|_1 \|u\|_{\infty} + (\|G^{\tau}_{ed_1}\|_1 + \|G^{\tau}_{ed_2}\|_1) r. \tag{3.8}$$

Theorem 6 is an immediate consequence of Lemmas 7 and 8 combined with standard gain bounds for the L_{∞} norm [65].

The final result gives an upper bound on the spike rate required for control in terms of the channel parameters r and τ and the system gains. Define $G^{\tau}_{w_1 u}$, $G^{\tau}_{w_1 d_1}$, $G^{\tau}_{w_1 d_2}$, $G^{\tau}_{w_2 u}$, $G^{\tau}_{w_2 d_1}$, and $G^{\tau}_{w_2 d_2}$ to be the input-output mappings for the corresponding signals in Figure 3.7.

Theorem 7. *If the nominal feedback is internally stable and τ is small enough that the disturbance loop from Figure 3.7 is internally stable, then the total number of spikes per unit time from both channels in Figure 3.3 is bounded above by*

$$\frac{\alpha(\tau)}{r\tau}\|u\|_\infty + \frac{\beta(\tau)}{\tau},$$

where

$$\begin{aligned}\alpha(\tau) &= \|G^\tau_{w_1u}\|_1 + \|G^\tau_{w_2u}\|_1 \\ \beta(\tau) &= \|G^\tau_{w_1d_1}\|_1 + \|G^\tau_{w_1d_2}\|_1 \\ &\quad + \|G^\tau_{w_2d_1}\|_1 + \|G^\tau_{w_2d_2}\|_1.\end{aligned}$$

Proof. By Lemma 6, the total spike rate is bounded by

$$f(\|w_1\|_\infty) + f(\|w_2\|_\infty) \leq \frac{\|w_1\|_\infty + \|w_2\|_\infty}{r\tau}.$$

Now applying Lemmas 7 and 8, and L_∞ gain bounds gives the result. □

Theorem 7 implies that if small τ and r are chosen in order to maintain internal stability and good tracking, then the data rate could become large.

Note that large values of $\alpha(0)$ or $\beta(0)$ correspond to large gains in the nominal system, whereas rapid growth of α or β corresponds to sensitivity to the perturbation caused by inserting low-pass filters into the loop. If α or β are large, then the data rates could be high, even when r and τ are large. Precise bounds on how α and β vary with τ are beyond the scope of this chapter.

See Figure 3.8 for examples of the tracking response of the spike channel feedback loop. The top plot depicts the response of a system with the unstable plant $P = \frac{1}{s-1}$ and controller $C = \frac{s+2}{s+1}$. The bottom plot depicts the response of the integrator $P = \frac{1}{s}$ with unity controller $C = 1$. The inputs are both the same, but the figures look different because of the larger tracking errors for the unstable plant. Furthermore, the system with the unstable plant requires a much higher spike rate (136.4 spikes per unit time) than the integrator system (21.3 spikes per unit time) to track the same input. Since the channel parameters, r and τ, as well as the input were identical, the difference in spike rate must be due to differences in loop gains.

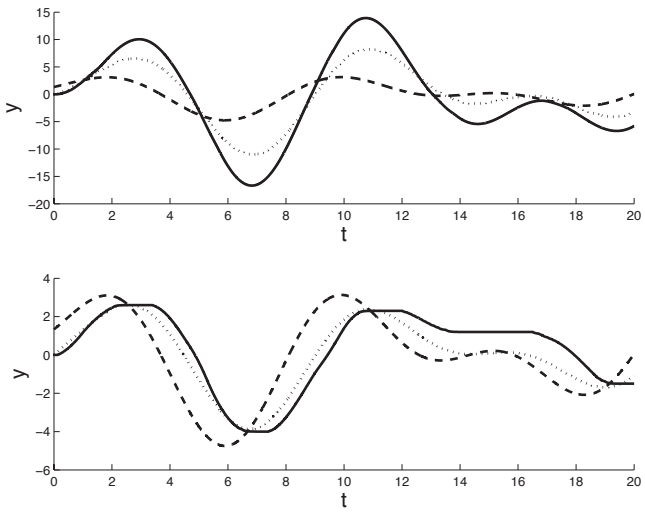

Figure 3.8: The tracking response for an unstable plant, and an integrator. In both cases $r = 1$ and $\tau = 0.1$ and the input was identical. The dashed line is the input, the dotted line is the output of the nominal feedback loop, and the solid line is the output of the feedback loop with spike channels. Top. $P = \frac{1}{s-1}$ and $C = \frac{s+2}{s+1}$. In this case the total spike rate (from both channels) was 136.4 spikes per unit time. Bottom. $P = \frac{1}{s}$ and $C = 1$. Here the total spike rate was 21.3 spikes per unit time.

3.4 Conclusion

This chapter describes feedback control over a novel communication channel, termed the spike channel. The spike channel is patterned after a configuration of neurons that allows simple reconstruction of the input current signal via linear filtering. It was shown that the spike channel can be conservatively approximated by a low-pass filter plus a bounded additive disturbance. Using this approximation, sufficient conditions for stability and good tracking performance were obtained. Data rate bounds based on signal gains and channel parameters were also obtained. Because all of the analysis in the chapter relied on approximations of the spike channel, it is unclear how conservative the results might be.

The work in this chapter represents just one of many possible research directions in the study of control over neuron-inspired communication channels. Future research will include varying the neuron models used, studying more sophisticated neural communication networks, and strengthening the connections to biology.

While the spike channel model admits a simple approximation analysis, precise bounds on the spike rate required for stability are difficult to obtain. By considering different neuron models, it may be possible to calculate tighter data rate bounds.

Another interesting variation on the channel would be to consider more realistic neuron models, such as the model proposed by [2]. It seems likely that stability could be proved for systems with

communication occurring over more realistic models through the use of Lyapunov arguments similar to those used in event-triggered control [59, 60].

While the opposing neuron construction of [1, 45], which motivated the spike channel, has been widely studied and admits theoretical analysis, it is unclear how widely such communication schemes are used in biology. In particular, in vertebrates, motor commands are encoded by large populations of neurons, instead of the single neurons studied in this chapter. It would be interesting to interpret motor control coding strategies in terms of networked control.

Finally, simple experiments can be designed to test the data rates used in human motor control. Studies such as [54] and [56] have exposed a general relationship between the the rate of information processing and the magnitude of the unstable poles of the plant (in discrete time). In order to assess information processing rates in humans, subjects could attempt to stabilize a plant (such as a flight simulator) with unstable poles varied by the experimenter.

3.5 Appendix to Chapter 3

Proof of Lemma 6. The maximum spike rate is given by $1/T_{\min}$, where $T_{\min}$ is the minimum time between spikes. If a spike ever occurred, without loss of generality, assume that a spike occurred at $t = 0$. Then until the next spike, the magnitude of the first filter is bounded as

$$\begin{aligned} |x_1(t)| &= \left| \frac{1}{\tau} \int_0^t e^{-(t-\sigma)/\tau} w(\sigma) d\sigma \right| \\ &\leq \frac{m}{\tau} \int_0^t e^{-(t-\sigma)/\tau} d\sigma \\ &= m \left(1 - e^{-t/\tau}\right). \end{aligned}$$

If $m \leq r$, then $|x_1(t)| < r$ for all $t > 0$, and thus no more spikes occur, so $f(m) = 0$. On the other hand, if $m > r$, then by solving $m\left(1 - e^{-t/\tau}\right) = r$ for t, $T_{\min}$ is calculated to be

$$T_{\min} = \tau \ln \frac{m}{m-r}.$$

Thus the maximum spike rate is calculated to be

$$f(m) = \begin{cases} \frac{1}{\tau \ln \frac{m}{m-r}} & \text{if } m > r \\ 0 & \text{if } m \leq r. \end{cases}$$

To calculate $\lim_{m\to\infty} \frac{f(m)}{m}$, let $\lambda = \frac{1}{m-r}$. Then since $\lim_{\lambda\to 0} \frac{\ln(1+r\lambda)}{\lambda} = r$, the limit of the

denominator of $\frac{f(m)}{m}$ can be simplified as follows:

$$\begin{aligned}\lim_{m\to\infty} m\ln\frac{m}{m-r} &= \lim_{\lambda\to 0}\frac{1+r\lambda}{\lambda}\ln(1+r\lambda)\\ &= \lim_{\lambda\to 0}\frac{\ln(1+r\lambda)}{\lambda}+\lim_{\lambda\to 0} r\ln(1+r\lambda)\\ &= r.\end{aligned}$$

Therefore $\lim_{m\to\infty}\frac{f(m)}{m}$ is calculated to be

$$\lim_{m\to\infty}\frac{1}{\tau m\ln\frac{m}{m-r}}=\frac{1}{\tau r}.$$

To see that $f(m)\le\frac{m}{\tau r}$ for all $m>r$, note that

$$f(m)\le\frac{m}{\tau r}\quad\text{iff}\quad m\ln\frac{m}{m-r}\ge r.$$

The calculations above show that $\lim_{m\to\infty} m\ln\frac{m}{m-r}=r$, and furthermore $\lim_{m\to r} m\ln\frac{m}{m-r}=\infty$. Thus, the proof can be completed by showing that $m\ln\frac{m}{m-r}$ is monotonically decreasing.

$$\begin{aligned}\frac{d}{dm}m\ln\frac{m}{m-r} &= \ln\frac{m}{m-r}+m\left(\frac{1}{m}-\frac{1}{m-r}\right)\\ &= \ln\frac{m}{m-r}-\frac{r}{m-r}.\end{aligned}$$

From the expression of the derivative, $m\ln\frac{m}{m-r}$ is monotonically decreasing if and only if $(m-r)\ln\frac{m}{m-r}<r$ for all $m>r$. Applying the change of variables $\lambda=\frac{1}{m-r}$ gives

$$(m-r)\ln\frac{m}{m-r}=\frac{\ln(1+r\lambda)}{\lambda}<r,$$

where the inequality follows from the first order necessary conditions for concavity. Thus $f(m)\le\frac{m}{r\tau}$ for all $m\ge 0$ and the proof is complete. □

Proof of Lemma 8. First note that the the closed loop poles of the disturbance loop are the zeros of

$$1+\frac{1}{(s\tau+1)^2}PC.$$

Note that the term $\frac{1}{(s\tau+1)^2}$ does not introduce any unstable open loop poles or zeros into the feedback loop. The idea of the proof is to show that for small enough τ, the Nyquist plots of $\frac{1}{(s\tau+1)^2}PC$ and PC encircle -1 the same number of times, since that would prove that the number of unstable closed loop poles remains unchanged.

Consider a Nyquist plot of PC. Since PC is strictly proper, there exists $M>0$ such that

$|P(j\omega)C(j\omega)| \leq \frac{1}{2}$ whenever $|\omega| \geq M$. Thus all the encirclements of -1 of the Nyquist plot occur in the image of $[-jM, jM]$. Furthermore, for all ω with $|\omega| \geq M$,

$$\begin{aligned} & \left| \frac{1}{(j\omega\tau+1)^2} P(j\omega)C(j\omega) \right| \\ = \quad & \frac{1}{\omega^2\tau^2+1} |P(j\omega)C(j\omega)| \\ \leq \quad & |P(j\omega)C(j\omega)| \\ \leq \quad & \frac{1}{2}. \end{aligned}$$

So, similarly, all encirclements of -1 in the Nyquist plot of $\frac{1}{(s\tau+1)^2}PC$ occur in the image of $[-jM, jM]$.

By internal stability and continuity of a continuous function over a compact domain, there exists $\gamma > 0$ such that $|1 + P(j\omega)C(j\omega)| \geq \gamma$ for all $\omega \in [-M, M]$. Furthermore, by continuity, T can be chosen small enough such that for all $\tau \in [0, T]$, and all $\omega \in [-M, M]$,

$$\left| 1 + \frac{1}{(j\omega\tau+1)^2} P(j\omega)C(j\omega) \right| \geq \frac{\gamma}{2}.$$

Thus, for all $\tau \in [0, T]$ the Nyquist plot of $\frac{1}{(s\tau+1)^2}PC$ does not pass through -1. Therefore the Nyquist plots of PC and $\frac{1}{(s\tau+1)^2}PC$ must encircle -1 the same number of times. $\square$

Chapter 4

Hybrid Systems and Local Zeno Stability

4.1 Introduction

This final technical chapter discusses some of the intricacies arising from hybrid dynamics. A hybrid system is a dynamical system that incorporates both continuous and discrete dynamics. Discrete dynamics could include instantaneous jumps in a continuous state variable, as well switches to completely different dynamical modes. The spiking neuron models from Chapter 3 are hybrid systems because once the current reaches threshold, a spike is sent and the current is reset to its resting value. In that chapter, the aim was to reduce the analysis of the hybrid system to the analysis of classical control systems. In contrast, for some systems, hybrid phenomena may be too pervasive to ignore, and in others, hybrid dynamics may actually be useful. For instance, it has been suggested that the computing power and energy efficiency of the brain may be at least partially due to a sophisticated interaction between low-power analog processing and discrete spikes [66, 67, 68].

On a more basic and obvious level, hybrid dynamics are important for locomotion in animals and robots [69, 70, 71]. Walking, for instance, incorporates both mode switches and rapid variable changes (Figure 4.1). During a stride, the left foot might start out on the ground while the right leg swings. During the leg swing, the knee rotates until the joint limit is reached and the knee locks. During each mode switch, some of the continuous variables change rapidly. For example, when the knee locks, the knee rotation speed rapidly decreases to zero.

This chapter studies a hybrid phenomenon known as Zeno behavior, and relates it to Lyapunov stability. Zeno behavior occurs in hybrid systems when an execution (or solution) undergoes infinitely many discrete transitions in a finite amount of time. Zeno behavior often occurs in models of mechanical systems undergoing impacts, including models important for locomotion. While Zeno behavior can be attributed to insufficiently modeling the complex dynamics of a system, i.e., it can be attributed to a "modeling pathology," it is present even in elementary examples such as the

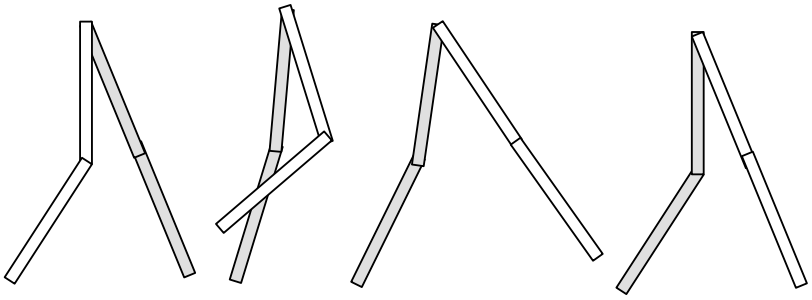

Figure 4.1: Dynamical modes in walking. The figure starts with both feet on the ground. In the next mode, the left foot is on the ground while the right leg swings with bent knee. After that, the left foot remains on the ground but the right leg swings with a locked knee. Finally both feet touch the ground again. The transition from a bent knee to a locked knee involves a rapid change in the rotational speed of the knee joint, while the transition from a swinging leg to a fixed leg corresponds to a rapid change in foot swing.

ubiquitous bouncing ball. Formally, it is important to understand Zeno behavior as it is indicative of phenomena unique to hybrid systems due to the complex interaction between the discrete and continuous (even if this is a result of the abstractions that yield hybrid models). From a practical perspective, Zeno behavior can stall simulations and lead to unexpected behavior if a hybrid control law admits Zeno behavior. In physical models, Zeno behavior occurs in mode switches, such as in the transition from bouncing to sliding. In such systems, being able to reason about the simplified models that exhibit Zeno behavior might be preferable to introducing extra terms to eliminate the behavior. Finally, if Zeno behavior can be understood, this understanding can be used as a preventative measure to eliminate the negative effects of this behavior *a priori*.

4.1.1 Simple Zeno Hybrid Systems

Since Zeno behavior may be unfamiliar to many readers, it will be illustrated in a few concrete examples before proceeding further.

***Example* 1.** Consider a ball bouncing with Newtonian impacts, defined as follows:

$$\left.\begin{array}{rl} \ddot{x} = -g & \text{for } x \geq 0 \\ x^+ = 0 & \\ \dot{x}^+ = -e\dot{x} & \text{for } x = 0 \text{ and } \dot{x} \leq 0. \end{array}\right\} \tag{4.1}$$

Here $g > 0$ denotes gravitational acceleration and $e > 0$ is the coefficient of restitution. Assume that $x(0) = 0$ and $\dot{x}(0) = \gamma > 0$. Integration shows that

$$\begin{aligned} \dot{x}(t) &= \gamma - gt \\ x(t) &= \gamma t - \frac{1}{2}gt^2. \end{aligned}$$

To find the time of the first impact, the second equation is solved for $x(t) = 0$ and $t > 0$. Explicitly

$$x(t) = t\left(\gamma - \frac{1}{2}gt\right) = 0,$$

and thus, the first impact occurs at $\tau_1 = 2\gamma/g$. The velocity at the first impact is given by

$$\dot{x}(\tau_1^-) = \gamma - g\frac{2\gamma}{g} = -\gamma.$$

Here τ_1^- denotes the left limit. Similarly let τ_1^+ denote the right limit. Applying the impact equation from Equation (4.1) gives the new state

$$x(\tau_1^+) = 0, \quad \dot{x}(\tau_1^+) = \gamma e.$$

Let τ_2 denote the time of the second impact. The argument above shows that $\tau_2 - \tau_1 = 2\gamma e/g$ and after the second impact the state is

$$x(\tau_2^+) = 0, \quad \dot{x}(\tau_2^+) = \gamma e^2.$$

Similarly, if τ_k denotes the time of the kth impact, then

$$\tau_k - \tau_{k-1} = \frac{2\gamma}{g}e^k, \quad x(\tau_k^+) = 0, \quad \dot{x}(\tau_k) = \gamma e^k.$$

Note that both the time between impacts and also the velocity after impacts decrease geometrically. The key thing to note is that an infinite number of impacts occur in a finite amount of time. Let $\tau_\infty = \lim_{k\to\infty} \tau_k$. Let $\tau_0 = 0$. It follows that

$$\begin{aligned}
\tau_\infty &= \lim_{k\to\infty} \tau_k \\
&= \sum_{k=1}^{\infty} (\tau_k - \tau_{k-1}) \\
&= \sum_{k=0}^{\infty} \frac{2\gamma}{g} e^k \\
&= \frac{2\gamma}{g}\frac{1}{1-e}.
\end{aligned}$$

It is also important to note that as $t \to \tau_\infty$, the state variables $x(t)$ and $\dot{x}(t)$ both converge to the zero (Figure 4.2).

In the bouncing ball example, τ_∞ is called the *Zeno time*. The property that all trajectories approach the origin as t approaches the Zeno time is called *Zeno stability*. The origin is referred

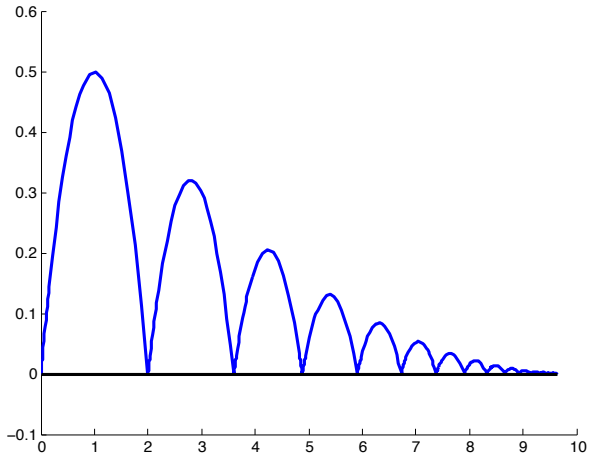

Figure 4.2: A plot of the trajectory of the bouncing ball. The time between impacts and the velocity decrease geometrically in the number of impacts. As such, the Zeno behavior displayed appears to be linked to stability.

to as a *Zeno equilibrium*. A Zeno equilibrium is a fixed point of the discrete dynamics, but not a fixed point of the continuous dynamics. The importance of Zeno equilibria will be discussed later. Contrast these ideas with classical stability, in which the state approaches an equilibrium point as $t \to \infty$.

***Example* 2.** Another canonical example of Zeno behavior occurs in the water tank system, shown in Figure 4.3. Two water tanks leak at rates $v_1 > 0$ and $v_2 > 0$, respectively. A pipe is used to fill the tanks by pouring water at a rate w. The pipe, however, can only pour water into one tank at a time. While $x_2 \geq 0$, the pipe pours water into tank 1. When the water level in tank 2 reaches zero, the pipe switches from tank 1 to tank 2. Then, when tank 1 reaches zero, the pipe switches back to tank 1. In other words, the dynamics are given by

$$\begin{aligned} \text{Mode 1}: &\quad \left.\begin{aligned} \dot{x}_1 &= -v_1 \\ \dot{x}_2 &= w - v_2 \end{aligned}\right\} \quad \text{while } x_1 \geq 0 \\ \text{Mode 2}: &\quad \left.\begin{aligned} \dot{x}_1 &= w - v_1 \\ \dot{x}_2 &= -v2 \end{aligned}\right\} \quad \text{while } x_2 \geq 0. \end{aligned}$$

Note that if $w > v_1$ and $w > v_2$, then while the pipe is filling each tank, the level increases. The fact that the draining rates v_1 and v_2 are positive guarantees that while one tank is filled, the other eventually drains and the pipe must switch. Thus, an infinite number of switches must occur. Now it will be shown how Zeno behavior can occur in the water tank system. Assume that $x_1(0) = \gamma > 0$

and $x_2(0) = 0$ and the system starts in Mode 1. Tank 1 is empty when

$$x_1(\tau_1) = \gamma - v_1\tau_1 = 0,$$

in which case $\tau_1 = \gamma/v_1$. At this point $x_2(\tau_1)$ is given by

$$x_2(\tau_1) = \frac{w - v_2}{v_1}\gamma.$$

Now the pipe switches and tank 2 drains until it reaches zero at time τ_2. A similar calculation to the one just performed shows that

$$\tau_2 - \tau_1 = \frac{x_2(\tau_1)}{v_2}.$$

The state at τ_2 is given by $x_2(\tau_2) = 0$ and

$$x_1(\tau_2) = \frac{w - v_1}{v_2}x_2(\tau_1) = \frac{w - v_1}{v_2}\frac{w - v_2}{v_1}\gamma = \alpha\gamma.$$

Consider the case when $\alpha < 1$. As in the bouncing ball example, let τ_k be the time of the kth mode switch. When $k = 2j$, a straightforward calculation shows that $x_2(\tau_{2j}) = 0$ and

$$x_1(\tau_{2j}) = \alpha^j\gamma.$$

Thus, if $0 < \alpha < 1$, then the state size decreases geometrically.

Note that τ_2 is the time required to return to Mode 1, and it is calculated as

$$\tau_2 = \tau_1 + \frac{x_2(\tau_1)}{v_2} = \left(\frac{1}{v_1} + \frac{w - v_2}{v_1}\frac{1}{v_2}\right)\gamma.$$

It follows, similarly, that $\tau_{2j+2} - \tau_{2j}$ is given by

$$\left(\frac{1}{v_1} + \frac{w - v_2}{v_1}\frac{1}{v_2}\right)\gamma\alpha^j.$$

Thus the time between mode switches also decreases geometrically. Similar to the bouncing ball, an infinite number of transitions must occur in a finite amount of time.

The bouncing ball and water tank examples serve to illustrate what Zeno behavior is and why it is associated with stability. In Zeno behavior, the time between discrete transitions must converge to zero. Excluding pathological cases such as having infinitely many discrete modes or vector fields that blow up, the temporal convergence implies that the state converges spatially.

The weakness of these examples lies in the fact that proofs of Zeno behavior and spatial convergence both rely heavily on closed form solutions to the hybrid systems. Prior to the introduction of

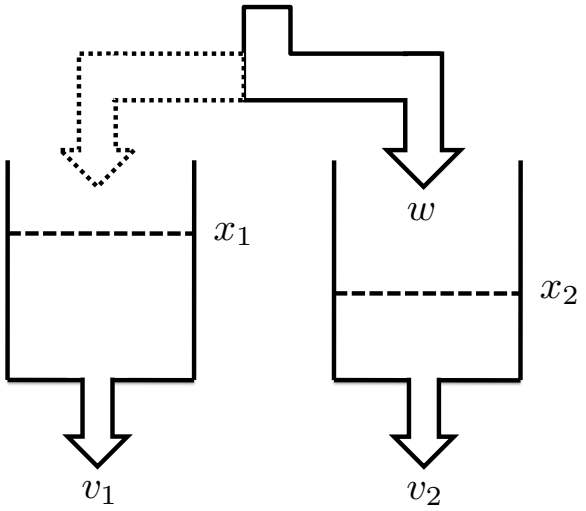

Figure 4.3: The water tank system, as adapted from [3]. Water drains from tank 1 at a rate v_1 and drains from tank 2 at a rate v_2. The current level in each tank is given by x_1 and x_2. The goal is to keep water in both tanks. Water flows into tank 1 at a rate w until the water level of tank 2 drops to zero. At that point, the pipe switches to fill tank 2 at rate $w > 0$. Similarly when tank 1 reaches zero, the pipe switches back to tank 1. This represents a switch from one dynamical mode to another. If $x_1(0) > 0$, $x_2(0) > 0$ and $\max v_1, v_2 < w < v_1 + v_2$ then an infinite number of switches occurs in finite amount of time.

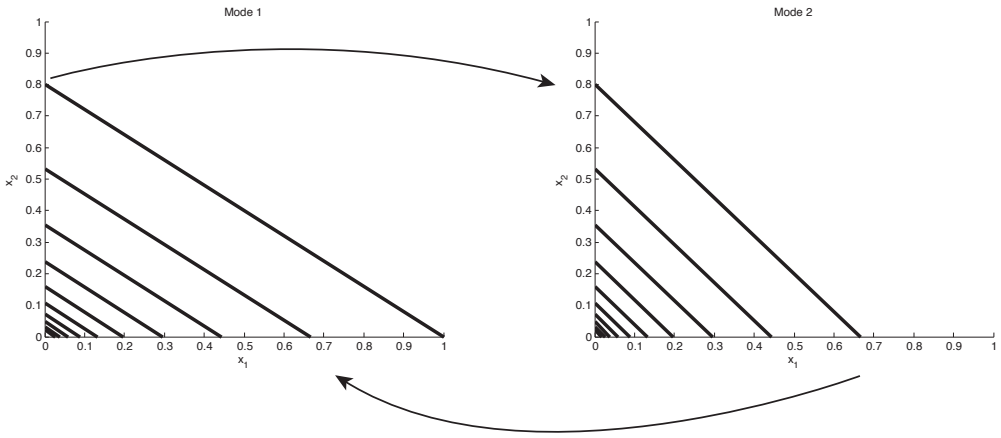

Figure 4.4: This shows a solution for the water tank system. As in the bouncing ball system as t approaches the Zeno time, the solution converges to a Zeno equilibrium. In this case the Zeno equilibrium is the set containing the origin from mode 1 and the origin from mode 2.

Lyapunov and local approximation methods (both of which will be touched in this chapter), Zeno behavior was rarely proved in hybrid systems that did not admit closed form solutions. In order to develop a theory of hybrid systems that handle complex dynamics for locomotion, understanding Zeno behavior for systems without closed form solutions is useful.

4.1.2 Summary of Contributions

With some concrete examples in mind, the concepts of this chapter can be discussed in greater depth. The main contributions of this chapter are: 1) a close connection between asymptotic Zeno stability and the geometry of Zeno equilibria; 2) Lyapunov-like sufficient conditions for local Zeno stability for hybrid systems over cycles; 3) easily verifiable sufficient conditions for Zeno stability of Lagrangian hybrid systems, which model mechanical systems undergoing impacts.

The first main contribution relates to the Zeno stability of Zeno equilibria. A Zeno equilibrium is a special type of invariant set unique to hybrid systems consisting of a set of points (with one point in each discrete domain) that is invariant under the discrete dynamics of the hybrid system but not the continuous dynamics. In the context of Zeno stability, it is shown that a Zeno equilibrium is asymptotically Zeno stable if and only if it is isolated (each point in each domain is isolated). This result highlights a major difference between classical Lyapunov stability and Zeno stability. It also clarifies the limitations of the most recent results, [72, 73, 74, 75], which focus either on isolated Zeno equilibria or asymptotic convergence.

The next main contribution is Lyapunov-like sufficient conditions for Zeno stability that apply to both isolated and non-isolated Zeno equilibria. The classical Lyapunov theorem uses a Lyapunov function to map solutions of a complex differential equation down to the solution of a simple one-dimensional differential inclusion, and then uses the structure of the Lyapunov function to prove that the original system inherits the stability properties of the one-dimensional system. The approach to Zeno stability in this chapter is similarly inspired. Lyapunov-like functions map executions of a complex hybrid system down to executions of simple two-dimensional differential inclusion hybrid systems, and then the structure of the Lyapunov-like functions is used to prove that the original system inherits some Zeno stability properties of the two-dimensional system. In contrast to existing results, this theorem applies equally well to isolated and non-isolated Zeno equilibria and, as will be seen, it thus applies to both asymptotic and non-asymptotic Zeno stability.

The final contribution applies the Lyapunov-like theorem to Lagrangian hybrid systems (which model mechanical systems undergoing impacts). While the technical machinery of hybrid systems is not needed to develop the theory of mechanical systems with impacts, the aim is to show how the theory of this chapter covers this important special case. Zeno stability in Lagrangian hybrid systems is proved by constructing a general form for a Lyapunov-like function that applies to any Lagrangian hybrid system whose vector field satisfies simple algebraic conditions at a single point

(based upon the *unilateral constraint function* defining the discrete component of the Lagrangian hybrid system). The strength of the theorem lies in its applications, thus several examples will be given. The result has been extended and refined in [76, 9, 10].

4.1.3 Relationship with Previous Results

Given the pathological nature of Zeno behavior, most early research on the existence of Zeno behavior focused on ruling it out, i.e., focused on necessary conditions, but left numerous open questions. Some results relied on strong structural assumptions about hybrid automata [3, 77], while others were simply difficult to verify [78]. Later research gained traction by studying Zeno behavior in restricted classes of hybrid systems. Linear complementarity systems researchers developed some of the first verifiable conditions to rule out Zeno behavior in a nontrivial class of hybrid systems [79, 80]. Early papers on sufficient conditions for Zeno behavior relied heavily on closed-form solutions of the associated vector fields [81, 82]. In particular, [81] provides a full characterization of Zeno behavior in a class of simple hybrid systems.

Sufficient conditions for Zeno behavior reached a new level of maturity based upon local approximations [72], [74], and connections with Zeno stability, which is conceptually the same as Lyapunov stability, except that all executions are required to be Zeno, [73], [75], [83]. The paper [75] gives a full characterization of asymptotic Zeno stability in a general class of systems. The papers mentioned above, with exception of [75], study Zeno behavior around Zeno equilibria.

This chapter's results on Zeno stability in mechanical system have been preceded by similar results in the mechanics literature, some dating back to the early 1990s [84, 85]. It should also be noted that simulation techniques for Lagrangian hybrid systems exist that seamlessly handle Zeno behavior because they do not need to explicitly calculate impact events [86, 87]. Thus, in the special case of mechanical systems, Zeno behavior causes fewer problems in simulation.

Despite recent advances, the state-of-the-art sufficient conditions cannot handle many simple and important examples of hybrid systems that appear to have Zeno executions. While the bouncing ball can be proved to be Zeno by integrating trajectories, balls bouncing on moving surfaces or mechanical systems with even slightly more complex geometries cannot be handled by the results mentioned above. Similarly, none of the previous sufficient conditions can explain Zeno behavior in a double pendulum with a mechanical stop, which is used by roboticists as a simplified model of the leg of a bipedal walker with a knee joint [88, 89]. Bouncing on irregular surfaces and the knee joint problem will be treated in this chapter.

4.2 Definitions and Geometric Results

This section introduces the basic terminology used throughout this chapter, such as hybrid systems, executions, and Zeno behavior. The terminology used in this chapter is different from other chapters. This is largely an artifact of the history of hybrid systems. Early incarnations of hybrid systems were finite automata augmented with simple continuous dynamics [90, 91, 92, 93]. To match automata-theoretic terminology, solutions to the continuous and discrete dynamics are called *executions*. Similarly, automata are traditionally defined in terms of tuples consisting of a state space, a transition function, etc. Since hybrid automata are extensions of finite automata, they are defined by adding more terms to the tuple. The size and complication of tuple definitions and the associated solution concepts for hybrid systems is, admittedly, unfortunate. More recent formulations of hybrid systems have lead to more compact notation [94, 95], but this chapter will adhere to the notation used in the original publication of this work [73, 96].

This chapter focuses on a restricted class of hybrid automata that strips away the nondeterminism and complicated graph structures in order to focus on consequences of the continuous dynamics. For more on the dynamic aspects of hybrid automata see [97].

Definition 5. A *hybrid system on a cycle* is a tuple:

$$\mathscr{H} = (\Gamma, D, G, R, F),$$

where

- $\Gamma = (Q, E)$ is a directed cycle, with

$$\begin{aligned} Q &= \{q_0, \ldots, q_{k-1}\}, \\ E &= \{e_0 = (q_0, q_1), e_1 = (q_1, q_2), \ldots, e_{k-1} = (q_{k-1}, q_0)\}. \end{aligned}$$

  We denote the source of an edge $e \in E$ by $\text{source}(e)$ and the target of an edge by $\text{target}(e)$.

- $D = \{D_q\}_{q \in Q}$ is a set of *continuous domains*, where D_q is a smooth manifold.

- $G = \{G_e\}_{e \in E}$ is a set of *guards*, where $G_e \subseteq D_{\text{source}(e)}$ is an embedded submanifold of $D_{\text{source}(e)}$.

- $R = \{R_e\}_{e \in E}$ is a set of *reset maps*, where $R_e : G_e \subseteq D_{\text{source}(e)} \to D_{\text{target}(e)}$ is a smooth map.

- $F = \{f_q\}_{q \in Q}$, where $f_q : D_q \to TD_q$ is a Lipschitz vector field on D_q.

Definition 6. An *execution* (or *solution*) of a hybrid system $\mathscr{H} = (\Gamma, D, G, R, F)$ is a tuple:

$$\chi = (\Lambda, I, \rho, C)$$

where

- $\Lambda = \{0, 1, 2, \ldots\} \subseteq \mathbb{N}$ is a finite or infinite indexing set.
- $I = \{I_i\}_{i \in \Lambda}$ where for each $i \in \Lambda$, I_i is defined as follows: $I_i = [\tau_i, \tau_{i+1}]$ if $i, i+1 \in \Lambda$ and $I_{N-1} = [\tau_{N-1}, \tau_N]$ or $[\tau_{N-1}, \tau_N)$ or $[\tau_{N-1}, \infty)$ if $|\Lambda| = N$, N finite. Here, for all $i, i+1 \in \Lambda$, $\tau_i \leq \tau_{i+1}$ with $\tau_i, \tau_{i+1} \in \mathbb{R}$, and $\tau_{N-1} \leq \tau_N$ with $\tau_{N-1}, \tau_N \in \mathbb{R}$. We set $\tau_0 = 0$ for notational simplicity.
- $\rho : \Lambda \to Q$ is a map such that for all $i, i+1 \in \Lambda$, $(\rho(i), \rho(i+1)) \in E$. This is the *discrete component* of the execution.
- $C = \{c_i\}_{i \in \Lambda}$ is a set of *continuous trajectories*, and they must satisfy $\dot{c}_i(t) = f_{\rho(i)}(c_i(t))$ for $t \in I_i$.

We require that when i, $i+1 \in \Lambda$,

$$\begin{aligned} &\text{(i)} && c_i(t) \in D_{\rho(i)} \;\forall\; t \in I_i \\ &\text{(ii)} && c_i(\tau_{i+1}) \in G_{(\rho(i),\rho(i+1))} \\ &\text{(iii)} && R_{(\rho(i),\rho(i+1))}(c_i(\tau_{i+1})) = c_{i+1}(\tau_{i+1}). \end{aligned} \tag{4.2}$$

When $i = |\Lambda| - 1$, we still require that (i) holds.

We call $c_0(0) \in D_{\rho(0)}$ the *continuous initial condition* of χ. Likewise $\rho(0)$ is the discrete initial condition of χ.

This chapter studies Zeno executions, defined as follows:

Definition 7. An execution χ is *Zeno* if $\Lambda = \mathbb{N}$ and

$$\lim_{i \to \infty} \tau_i = \sum_{i=0}^{\infty} \tau_{i+1} - \tau_i = \tau_\infty < \infty.$$

Here τ_∞ is called the *Zeno time*.

A hybrid system $\mathscr{H}$ is *Zeno*[1] if there exists a Zeno execution χ such that $\tau_{i+1} - \tau_i \neq 0$ for some $i \in \mathbb{N}$.

[1] This definition is motivated to exclude the possibility that a hybrid system is "trivially" Zeno, i.e., the only Zeno executions are executions that begin at a Zeno equilibrium.

***Remark* 6.** Note that if a hybrid system over a finite graph displays Zeno behavior, the graph must contain a cycle. Indeed, since an infinite number of transitions must occur in a finite graph, the graph must have a cycle. (See [77] and [3] for similar structural conditions on Zeno behavior.) Therefore, beginning with hybrid systems defined on cycles greatly simplifies the analysis, while still capturing characteristic types of Zeno behavior. Future research will examine Zeno behavior in hybrid systems with more complex graph structures.

Zeno behavior displays strong connections with Lyapunov stability [73, 75]. Just as classical stability focuses on equilibria, much of the interesting Zeno behavior occurs near a special type of invariant set, termed Zeno equilibria.

Definition 8. A *Zeno equilibrium* of a hybrid system $\mathscr{H} = (\Gamma, D, G, R, F)$ is a set $z = \{z_q\}_{q\in Q}$ satisfying the following conditions for all $q \in Q$:

- For the unique edge $e = (q, q') \in E$
  - $z_q \in G_e$,
  - $R_e(z_q) = z_{q'}$,
- $f_q(z_q) \neq 0$.

A Zeno equilibrium $z = \{z_q\}_{q\in Q}$ is *isolated* if there is a collection of open sets $\{W_q\}_{q\in Q}$ such that $z_q \in W_q \subset D_q$, and $\{W_q\}_{q\in Q}$ contains no Zeno equilibria other than z. Otherwise, z is *non-isolated*.

Note that, in particular, the conditions given in Definition 8 imply that for all $i \in \{0, \ldots, k-1\}$,

$$R_{e_{i-1}} \circ \cdots \circ R_{e_0} \circ R_{e_{k-1}} \circ \cdots \circ R_{e_i}(z_i) = z_i.$$

That is, the element z_i is a fixed point under the reset maps composed in a cyclic manner. Furthermore, any infinite execution with initial condition $c_0(0) \in z$ must be *instantaneously* Zeno (that is, $\tau_i = 0$ for all $i \in \mathbb{N}$).

The condition that $f_q(z_q) \neq 0$ is made for technical reasons. Convergent, non-chattering Zeno executions (those with $\tau_i < \tau_{i+1}$ for infinitely many i) must converge to a Zeno equilibrium unless the domains have geometric pathologies, such as cusps, or the vector fields are not locally lipschitz. See [78] Proposition 4.4 for a proof. See [75] and [78] for examples of Zeno hybrid systems defined on cusps which do not have Zeno equilibria.

Finally, the following definitions connect Zeno behavior to Lyapunov stability.

Definition 9. An execution $\chi = (\Lambda, I, \rho, C)$ is *maximal* if for all executions $\hat{\chi} = (\hat{\Lambda}, \hat{I}, \hat{\rho}, \hat{C})$ such that

$$\Lambda \subset \hat{\Lambda}, \qquad \bigcup_{j \in \Lambda} I_j \subset \bigcup_{j \in \hat{\Lambda}} \hat{I}_j,$$

and $c_j(t) = \hat{c}_j(t)$ for all $j \in \Lambda$ and $t \in I_j$, it follows that $\hat{\chi} = \chi$.

Definition 10. A Zeno equilibrium $z = \{z_q\}_{q \in Q}$ of a hybrid system $\mathscr{H} = (\Gamma, D, G, R, F)$ is:

- *bounded-time locally Zeno stable* if for every collection of open sets $\{U_q\}_{q \in Q}$ with $z_q \in U_q \subset D_q$ and every $\epsilon > 0$, there is another collection of open sets $\{W_q\}_{q \in Q}$ with $z_q \in W_q \subset U_q$ such that if χ is a maximal execution with $c_0(0) \in W_{\rho(0)}$, then χ is Zeno with $\tau_\infty < \epsilon$ and $c_i(t) \in U_{\rho(i)}$ for all $i \in \mathbb{N}$ and all $t \in I$.

- *bounded-time asymptotically Zeno stable* if it is bounded-time locally Zeno stable and there is a collection of open sets $\{W_q\}_{q \in Q}$ such that $z_q \in W_q \subset D_q$ and every Zeno execution $\chi = (\Lambda, I, \rho, C)$ with $c_0(0) \in W_{\rho(0)}$ converges to z as $i \to \infty$. More precisely, for any collection of open sets $\{U_q\}_{q \in Q}$ with $z_q \in U_q \subset D_q$, there is $N \in \mathbb{N}$ such that if $i \geq N$, then $c_i(t) \in U_{\rho(i)}$ for all $t \in I_i$.

- *bounded-time non-asymptotically Zeno stable* if it is bounded-time locally Zeno stable but not bounded-time asymptotically Zeno stable.

- *bounded-time globally asymptotically Zeno stable* if it is bounded-time asymptotically Zeno stable and every maximal execution is Zeno and converges to z.

The following structural fact shows that isolatedness of a Zeno equilibrium dictates the type of Zeno stability properties it can display. While the theorem is independent of the other main results of the chapter, it clarifies the existing sufficient conditions for Zeno stability and adds context to the current work.

Theorem 8. *Let $z = \{z_q\}_{q \in Q}$ be a bounded-time locally stable Zeno equilibrium. Then z is bounded-time asymptotically Zeno stable if and only if z is isolated.*

Note the sharp contrast between Theorem 8 and classical stability theory. The standard theory of continuous dynamical systems focuses nearly exclusively on isolated equilibria without much apparent conceptual loss. In Zeno hybrid systems, however, non-isolated Zeno equilibria must be studied just to describe the non-asymptotic analog of Lyapunov stability.

The theorem shows that many of the recent sufficient conditions for Zeno stability have similar limitations, but for different reasons. The work in [74] and [75] requires bounded-time asymptotic Zeno stability (or the stronger global version), while [72] and [73] assume that the hybrid systems

studied have isolated Zeno equilibria. None of the conditions in the papers listed above could apply to the mechanical systems in this chapter, precisely because the more complex systems have no bounded-time asymptotically Zeno stable Zeno equilibria. (This will be discussed in greater detail in Section 4.5.2.)

Proof. Let z be an isolated Zeno equilibrium. By continuity, there is a collection of bounded neighborhoods $\{U_q\}_{q\in Q}$ containing no Zeno equilibria other than z, such that for all $q \in Q$, $z_q \in U_q \subset D_q$ and $f_q(x) \neq 0$ for all $x \in U_q$. From bounded-time local Zeno stability, there is another collection of neighborhoods $\{W_q\}_{q\in Q}$ such that all maximal executions with initial conditions in $\{W_q\}_{q\in Q}$ are all Zeno and never leave $\{U_q\}_{q\in Q}$. Let χ be any maximal execution such that $c_0(0) \in W_q$ for some $q \in Q$. Since χ is Zeno and bounded, Proposition 4.3 of [78] implies that there is a collection of points $\hat{z} = \{\hat{z}_q\}_{q\in Q}$ such that:

- $\hat{z}_q \in G_{(q,q')} \cap U_q$ for all $(q, q') \in E$,
- $R_{(q,q')}(\hat{z}_q) = \hat{z}_{q'}$ for all $(q, q') \in E$,
- $c_i(t) \to \hat{z}_{\rho(i)}$ as $i \to \infty$.

Since $\hat{z}_q \in U_q$, it follows that $f_q(\hat{z}_q) \neq 0$ for all $q \in Q$. Therefore $\hat{z}$ is a Zeno equilibrium contained in $\{U_q\}_{q\in Q}$. The construction of U_q implies that $\hat{z} = z$, and thus χ converges to z. It follows that z is bounded-time asymptotically Zeno stable.

Conversely, let z be a non-isolated Zeno equilibrium. Then for any collection of neighborhood $\{U_q\}_{q\in Q}$, there is a Zeno equilibrium $\hat{z}$ with $\hat{z} \neq z$ and $\hat{z} \subset \{U_q\}_{q\in Q}$. Furthermore, any maximal execution with $c_0(0) \in \hat{z} \subset \{U_q\}_{q\in Q}$ is Zeno but does not converge to z. Therefore, z is not bounded-time asymptotically Zeno stable. □

***Example* 3** (Bouncing Ball on a Circle)**.** The definitions and concepts above, as well as the theorems to follow are illustrated by studying a ball bouncing on a circular surface (Figure 4.5).

Formally, bouncing ball on a circular surface is modeled by the hybrid system

$$\mathscr{H}_{\mathbb{B}} = (\Gamma = (\{q\}, \{(q, q)\}, \{D_{\mathbf{B}}\}, \{G_{\mathbf{B}}\}, \{R_{\mathbf{B}}\}, \{f_{\mathbf{B}}\}),$$

where

$$\begin{aligned} D_{\mathbf{B}} &= \{(x, \dot{x}) \in \mathbb{R}^2 \times \mathbb{R}^2 : \|x\| \geq 1\}, \\ G_{\mathbf{B}} &= \{(x, \dot{x}) \in \mathbb{R}^2 \times \mathbb{R}^2 : \|x\| = 1, \text{ and } x^T\dot{x} \leq 0\}, \end{aligned}$$

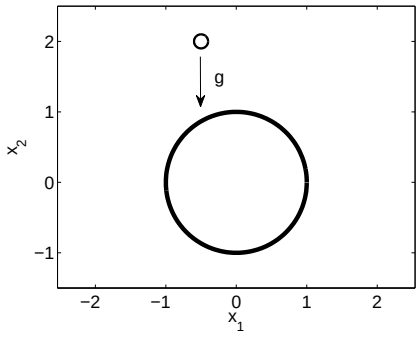

Figure 4.5: A ball moving through the plane under gravitational acceleration that bounces on a fixed circular surface. This simple system demonstrates most of the important phenomena discussed in this chapter.

$$R_{\mathbf{B}}(x,\dot{x}) = \begin{pmatrix} x \\ \dot{x} - (1+e)(x^T\dot{x})x \end{pmatrix}, \qquad f_{\mathbf{B}}(x,\dot{x}) = \begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ 0 \\ -g \end{pmatrix}.$$

Here the reset map, $R_{\mathbf{B}}$, is a Newtonian impact model, with a coefficient of restitution $0 \leq e \leq 1$, that describes an instantaneous jump in velocity when the ball impacts the circle. The vector field $f_{\mathbf{B}}$ models flight under gravitational acceleration.

Since $f_{\mathbf{B}}(x,\dot{x}) \neq 0$ on the entire continuous domain, the Zeno equilibria are exactly the fixed points of the reset map:

$$Z_{\mathbf{B}} = \{(x,\dot{x}) \in \mathbb{R}^2 \times \mathbb{R}^2 : \|x\| = 1, \text{ and } x^T\dot{x} = 0\}.$$

Note that $Z_{\mathbf{B}}$ is an infinite, connected set. Therefore, $\mathscr{H}_{\mathbb{B}}$ has no isolated Zeno equilibria. From Theorem 8, this bouncing ball system has no bounded-time asymptotically stable Zeno equilibria.

Turning to Zeno stability, the theory developed in this chapter predicts that whenever $0 < e < 1$ and $(x^*,\dot{x}^*) \in \mathbb{R}^2 \times \mathbb{R}^2$ satisfies the following algebraic conditions:

$$\|x^*\| = 1, \qquad {x^*}^T\dot{x}^* = 0, \qquad \|\dot{x}^*\|^2 < g x_2^*, \tag{4.3}$$

the singleton set $\{(x^*,\dot{x}^*)\}$ is a bounded-time non-asymptotically Zeno stable Zeno equilibrium. Note how the conditions guarantee a noncompact continuum of bounded-time locally Zeno stable sets along the entire open upper half circle, even at points with nearly vertical tangent spaces.

The ball on a circular surface also captures the fundamentally local nature of the conditions in this chapter. Many executions will never hit the circle at all, and simply free fall for all time.

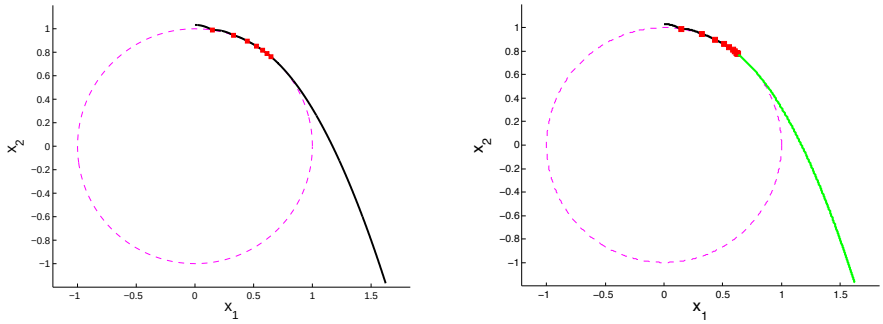

(a) (b)

Figure 4.6: Two simulations of the bouncing ball system with $g = 1$, $e = 1/2$ and initial conditions slightly varied. 4.6(a) With initial condition $c_0(0) = (x_1, x_2, \dot{x}_1, \dot{x}_2)^T = (0, 1.033, .5, 0)^T$, the execution bounces several times before free falling to infinity. 4.6(b) Shifting x_2 down a small amount, so that $c_0(0) = (x_1, x_2, \dot{x}_1, \dot{x}_2)^T = (0, 1.032, .5, 0)^T$, the execution becomes Zeno. After the Zeno behavior occurs, the green line depicts how the ball rolls along the surface before eventually falling off.

Others will make a finite number of collisions before escaping to a free fall. Finally, some executions can be Zeno. The theory developed in this chapter can be used to numerically distinguish between executions that take several bounces before free fall, and Zeno executions (Figure 4.6).

4.3 First Quadrant Interval Hybrid Systems

This section gives conditions for Zeno stability in a simple class of hybrid systems termed first-quadrant interval hybrid systems. These systems are easy to analyze, yet flexible enough to capture important characteristics of nontrivial systems. Indeed, first-quadrant interval systems serve as targets for Lyapunov-like reductions. First-quadrant interval hybrid systems are a variant on first-quadrant hybrid systems studied in [82] and [72]. The term "interval" is used since both the vector fields and reset maps are interval valued. See [98] and [99] for more on set valued functions and differential inclusions.

Definition 11. A *first-quadrant interval* FQI *hybrid system* is a tuple

$$\mathscr{H}_{FQI} = (\Gamma, D, G, R, F)$$

where

- $\Gamma = (Q, E)$ is a directed cycle as in Definition 5.

- $D = \{D_q\}_{q\in Q}$ where for all $q \in Q$,

$$D_q = \mathbb{R}^2_{\geq 0} = \{(x_1, x_2)^T \in \mathbb{R}^2 : x_1 \geq 0,\ \ x_2 \geq 0\}.$$

- $G = \{G_e\}_{q\in Q}$ where for all $e \in E$,

$$G_e = \{(x_1, x_2)^T \in \mathbb{R}^2_{\geq 0} : x_1 = 0,\ \ x_2 \geq 0\}.$$

- $R = \{R_e\}_{e\in E}$ where for all $e \in E$, R_e is a set valued function defined by

$$R_e(0, x_2) = \{(y_1, y_2)^T \in D_{q'} : y_2 = 0,\ \ y_1 \in [\gamma_e^l x_2, \gamma_e^u x_2]\},$$

  for $\gamma_e^u \geq \gamma_e^l > 0$ and for all $(0, x_2)^T \in G_e$.

- $F = \{f_q\}_{q\in Q}$ where for all $q \in Q$, f_q is the (constant) set-valued function defined by

$$f_q(x) = \{(y_1, y_2)^T \in \mathbb{R}^2 : y_1 \in [\alpha_q^l, \alpha_q^u],\ \ y_2 \in [\beta_q^l, \beta_q^u]\}.$$

Definition 12. An *execution* of a first-quadrant interval system, $\mathscr{H}_{FQI}$ is a tuple $\chi_{FQI} = (\Lambda, I, \rho, C)$ where

- Λ, I and ρ are defined as in Definition 6.
- $C = \{c_i\}_{i\in\Lambda}$ is a set of *continuous trajectories* that satisfy the differential inclusion $\dot{c}_i(t) \in f_{\rho(i)}(c_i(t))$ for $t \in I_i$.

When i, $i+1 \in \Lambda$, the conditions at the resets are given by

$$\begin{aligned} &\text{(i)} && c_i(t) \in D_{\rho(i)} \ \forall\ t \in I_i \\ &\text{(ii)} && c_i(\tau_{i+1}) \in G_{(\rho(i),\rho(i+1))} \\ &\text{(iii)} && c_{i+1}(\tau_{i+1}) \in R_{(\rho(i),\rho(i+1))}(c_i(\tau_{i+1})). \end{aligned} \tag{4.4}$$

When $i = |\Lambda| - 1$, condition (i) must still hold.

Theorem 9. *Let $\mathscr{H}_{FQI} = (\Gamma, D, G, R, F)$ be a first-quadrant interval hybrid system. If $\alpha_q^u < 0 < \beta_q^l$ for all $q \in Q$, $\gamma_e^l > 0$ for all $e \in E$ and*

$$\prod_{i=0}^{|Q|-1} \left|\gamma_{e_i}^u \frac{\beta_{q_i}^u}{\alpha_{q_i}^u}\right| < 1$$

then the origin $\{0_q\}_{q\in Q}$ *is bounded-time globally asymptotically Zeno stable.*

From the definition of bounded-time global asymptotic stability, there is a function $T_{\text{Zeno}} : \mathbb{R} \to \mathbb{R}$ such that the Zeno time of any maximal execution satisfies $\tau_\infty \leq T_{\text{Zeno}}(\|c_0(0)\|)$ and $T_{\text{Zeno}}(\|c_0(0)\|) \to 0$ as $\|c_0(0)\| \to 0$.

Proof. Define $0 < \zeta < 1$ by

$$\zeta := \prod_{i=0}^{|Q|-1} \left| \gamma^u_{e_i} \frac{\beta^u_{q_i}}{\alpha^u_{q_i}} \right|. \tag{4.5}$$

Let χ_{FQI} be an execution of $\mathscr{H}_{FQI}$. Without loss of generality, assume that $c_0(0) \in D_{q_0}$. Since $f_q(x)_2 \geq \beta^l_q > 0$, the continuous trajectories travel upwards, away from the x_1-axis. Likewise, $f_q(x)_1 \leq \alpha^u_q < 0$ implies that the continuous trajectories travel left, towards the x_2-axis. Therefore, by construction, events are always guaranteed, so Λ can be assumed to be $\mathbb{N}$ without loss of generality. For simplicity, assume that $c_0(0)_2 = 0$. Dropping this assumption changes little, though the proofs become messier.

The hypothesis $\alpha^u_{\rho(i)} < 0$ implies that $c_i(t)_1 \leq c_i(\tau_i)_1 + \alpha^u_{\rho(i)}(t - \tau_i)$, and therefore

$$\tau_{i+1} - \tau_i \leq \left| \frac{c_i(\tau_i)_1}{\alpha^u_{\rho(i)}} \right|, \tag{4.6}$$

for all $i \geq 0$. The continuous state at the first event must satisfy

$$c_0(\tau_1)_2 \leq \beta^u_0(\tau_1 - \tau_0) + c_0(\tau_0)_2 \leq c_0(\tau_0)_1 \left| \frac{\beta^u_0}{\alpha^u_0} \right| + c_0(\tau_0)_2.$$

Thus, after the first event the continuous state satisfies

$$c_1(\tau_1)_1 \leq \gamma^u_{\rho(0,1)} c_0(0)_1 \left| \frac{\beta^u_0}{\alpha^u_0} \right| + \gamma^u_{\rho(0,1)} c_0(0)_2. \tag{4.7}$$

Furthermore, $c_i(\tau_i)_2 = 0$ for all $i \geq 1$. Stability, asymptotic convergence, and the bound on the Zeno time all follow from bounds on $c_i(\tau_i)_1$.

It is claimed that

$$c_i(\tau_i)_1 \leq c_1(\tau_1)_1 \prod_{j=1}^{i-1} \left| \gamma^u_{(\rho(j),\rho(j+1))} \frac{\beta^u_{\rho(j)}}{\alpha^u_{\rho(j)}} \right| \tag{4.8}$$

for all $i \in \mathbb{N}$. Interpreting the product to be 1 when $i = 1$, Equation (4.8) holds for $i = 1$. Now inductively assume that equation (4.8) holds for some $i \geq 1$. Combining Equation (4.6) with the form of f_q gives an upper bound on $c_i(\tau_{i+1})_2$,

$$c_i(\tau_{i+1})_2 \leq \beta^u_{\rho(i)}(\tau_{i+1} - \tau_i) \leq c_i(\tau_i)_1 \left| \frac{\beta^u_{\rho(i)}}{\alpha^u_{\rho(i)}} \right|. \tag{4.9}$$

Using the form of the reset maps, the claim follows:

$$c_{i+1}(\tau_{i+1})_1 \leq c_i(\tau_i)_1 \left| \gamma^u_{(\rho(i),\rho(i+1))} \frac{\beta^u_{\rho(i)}}{\alpha^u_{\rho(i)}} \right| \leq c_1(\tau_1)_1 \prod_{j=1}^{i} \left| \gamma^u_{(\rho(j),\rho(j+1))} \frac{\beta^u_{\rho(j)}}{\alpha^u_{\rho(j)}} \right|.$$

To prove stability and asymptotic convergence note that $\alpha^u_q < 0 < \beta^l_q$ implies that $c_i(t)_1 \leq c_i(\tau_i)_1$ and $c_i(t)_2 \leq c_i(\tau_{i+1})_2$ for all $t \in I_i$. Combining Equations (4.7), (4.8), and (4.9) gives the bound

$$\begin{aligned}
\|c_i(t)\| &\leq c_i(\tau_i)_1 + c_i(\tau_{i+1})_2 \\
&\leq \left(1 + \left| \frac{\beta^u_{\rho(i)}}{\alpha^u_{\rho(i)}} \right| \right) c_i(\tau_i)_1 \\
&\leq \left(1 + \left| \frac{\beta^u_{\rho(i)}}{\alpha^u_{\rho(i)}} \right| \right) c_1(\tau_1)_1 \prod_{j=1}^{i-1} \left| \gamma^u_{(\rho(j),\rho(j+1))} \frac{\beta^u_{\rho(j)}}{\alpha^u_{\rho(j)}} \right| \\
&\leq \left(1 + \left| \frac{\beta^u_{\rho(i)}}{\alpha^u_{\rho(i)}} \right| \right) \gamma^u_{\rho_{(0,1)}} \left(c_0(0)_1 \left| \frac{\beta^u_0}{\alpha^u_0} \right| + c_0(0)_2 \right) \prod_{j=1}^{i-1} \left| \gamma^u_{(\rho(j),\rho(j+1))} \frac{\beta^u_{\rho(j)}}{\alpha^u_{\rho(j)}} \right|.
\end{aligned}$$

Since the product in the last inequality converges to 0 as $i \to \infty$, executions with $\|c_0(0)\|$ small must remain near the origin, and $c_i(t) \to 0_{\rho(i)}$ as $i \to \infty$.

Combining Equations (4.5), (4.6), and (4.8) and proves that χ is Zeno:

$$\begin{aligned}
\sum_{i=0}^{\infty} \tau_{i+1} - \tau_i &\leq c_0(0)_1 \sum_{i=0}^{\infty} \frac{1}{|\alpha^u_{\rho(i)}|} \prod_{j=0}^{i-1} \left| \gamma^u_{(\rho(j),\rho(j+1))} \frac{\beta^u_{\rho(j)}}{\alpha^u_{\rho(j)}} \right| \\
&= c_0(0)_1 \left(\sum_{j=0}^{|Q|-1} \frac{1}{|\alpha^u_{q_j}|} \prod_{k=0}^{j-1} \left| \gamma^u_{e_k} \frac{\beta^u_{q_k}}{\alpha^u_{q_k}} \right| \right) \cdot \left(\sum_{i=0}^{\infty} \zeta^i \right) \\
&< \infty.
\end{aligned}$$

Furthermore, note that the bound on the Zeno time goes to zero as $c_1(0)_1 \to 0$. □

Theorem 9 can also be proved using Lyapunov methods from [75], but the close relationship between spatial convergence and temporal convergence exploited in the proof above is used to prove Theorem 10, particularly for Lemma 10.

4.4 Sufficient Conditions for Zeno Stability through Reduction to FQI Hybrid Systems

This section gives the second main result of this paper, sufficient conditions for bounded-time local Zeno stability of hybrid systems via reduction to FQI hybrid systems. The theorem uses special Lyapunov-like functions to map executions of complex hybrid systems down to executions of FQI

hybrid systems, thus transferring some Zeno stability properties from Theorem 9.

In the theorem below, Lyapunov-like functions map neighborhoods around Zeno equilibria to the first quadrant. The theorem applies to both isolated and non-isolated Zeno equilibria. Therefore, by Theorem 8, the sufficient conditions below can imply bounded-time asymptotic or non-asymptotic Zeno stability, depending on the type of Zeno equilibrium in question.

Assumption. In this section, assume that each D_q is a subset of $\mathbb{R}^{n_q}$ with $n_q = \dim(D_q)$ and $z_q = 0$. No generality is lost because the results are local, and coordinate charts charts can be used.

Reduction conditions. Let $z = \{z_q\}_{q\in Q}$ be a Zeno equilibrium (not necessarily isolated) of a hybrid system $\mathscr{H} = (\Gamma, D, G, R, F)$, $\{W_q\}_{q\in Q}$ be a collection of sets with $z_q \in W_q \subseteq D_q$ and $\{\psi_q\}_{q\in Q}$ be a collection of C^1 maps; these are "Lyapunov-like" functions, with

$$\psi_q : W_q \subseteq D_q \to \mathbb{R}^2_{\geq 0}.$$

Consider the following conditions:

R1: $\psi_q(z_q) = 0$ for all $q \in Q$.

R2: If $(q, q') \in E$, then $\psi_q(x)_1 = 0$ if and only if $x \in G_{(q,q')} \cap W_q$.

R3: $d\psi_q(z_q)_1 f_q(z_q) < 0 < d\psi_q(z_q)_2 f_q(z_q)$ for all $q \in Q$.

R4: $\psi_{q'}(R_{(q,q')}(x))_2 = 0$ and there exist constants $0 < \gamma_e^l \leq \gamma_e^u$ such that

$$\psi_{q'}(R_{(q,q')}(x))_1 \in \left[\gamma^l_{(q,q')}\psi_q(x)_2, \gamma^u_{(q,q')}\psi_q(x)_2\right]$$

for all $x \in G_{(q,q')} \cap W_q$ and all $(q, q') \in E$.

R5:

$$\prod_{i=0}^{|Q|-1} \left|\gamma^u_{e_i} \frac{d\psi_{q_i}(z_{q_i})_2 f_{q_i}(z_{q_i})}{d\psi_{q_i}(z_{q_i})_1 f_{q_i}(z_{q_i})}\right| < 1.$$

R6: There exists $K \geq 0$ such that

$$\|R_{(q,q')}(x) - z_{q'}\| \leq \|x - z_q\| + K\psi_q(x)_2$$

for all $x \in G_{(q,q')} \cap W_q$ and all $(q, q') \in E$.

Theorem 10. *Let $\mathscr{H}$ be a hybrid system with a Zeno equilibria $z = \{z_q\}_{q\in Q}$. If there exists a collection of sets $\{W_q\}_{q\in Q}$ with $z_q \in W_q \subseteq D_q$ and maps $\{\psi_q\}_{q\in Q}$ satisfying conditions* **R***1-***R***6, then z is bounded-time locally Zeno stable.*

Before getting to the proof of the theorem, note that Theorems 8 and 10 immediately imply the following corollary.

Corollary 1. *Let $\mathscr{H}$ be a hybrid system with a Zeno equilibria $z = \{z_q\}_{q\in Q}$ satisfying the conditions of Theorem 10. If z is an isolated Zeno equilibrium, then z is bounded-time asymptotically Zeno stable. Otherwise, if z is a non-isolated Zeno equilibrium, then z is bounded-time non-asymptotically Zeno stable.*

Theorem 10 is proved as follows:

1. Construct a Zeno first-quadrant interval system $\mathscr{H}_{\text{FQI}}$ from the hybrid system $\mathscr{H}$ and map executions of the hybrid system to executions of the FQI hybrid system (Lemma 9).
2. Prove that executions of $\mathscr{H}$ stay "close" to the Zeno equilibria for a bounded period of time (Lemma 10).
3. Use 2) and 1) to show that $\mathscr{H}$ is Zeno because $\mathscr{H}_{\text{FQI}}$ is Zeno due to conditions **R**1–**R**6.

Constructing a FQI hybrid system. A first-quadrant interval system $\mathscr{H}_{\text{FQI}}$ can be defined from a hybrid system $\mathscr{H}$ based on the reduction conditions. Assume that $\mathscr{H}$ is a hybrid system satisfying **R**1–**R**5. Pick α_q^l, α_q^u, β_q^l, and β_q^u such that

$$\alpha_q^l < d\psi_q(0)_1 f_q(0) < \alpha_q^u < 0 < \beta_q^l < d\psi_q(0)_2 f_q(0) < \beta_q^u$$

for all $q \in Q$ and

$$\prod_{i=0}^{|Q|-1} \left| \gamma_{e_i}^u \frac{\beta_{q_i}^u}{\alpha_{q_i}^u} \right| < 1,$$

where $\gamma_{e_i}^u$ is given by **R**4. The constants α_q^l, α_q^u, β_q^l, β_q^u, $\gamma_{(q,q')}^l$, and $\gamma_{(q,q')}^u$ (with $\gamma_{(q,q')}^l$ also given by **R**4) thus define a first-quadrant interval system $\mathscr{H}_{FQI}$, on the same graph as $\mathscr{H}$, satisfying the conditions of Theorem 9 due to conditions **R**3–**R**5. Thus all executions of $\mathscr{H}_{FQI}$ extend to Zeno executions.

The following lemma shows how an execution of $\mathscr{H}$ remaining near the Zeno equilibria gives rise to an execution of $\mathscr{H}_{FQI}$.

Lemma 9. *Suppose $\mathscr{H}$ is a hybrid system satisfying the conditions of Theorem 10. Then there exists $\mu > 0$ such that if $\chi = (\Lambda, \rho, I, C)$ is an execution of $\mathscr{H}$ with $\|c_i(t)\| < \mu$ for all $t \in I_i$ and all $i \in \Lambda$, then $\chi_{FQI} = (\Lambda, \rho, I, \Psi \circ C)$, where $\Psi \circ C = \{\psi_{\rho(i)} \circ c_i\}_{i \in \Lambda}$, is an execution of of $\mathscr{H}_{FQI}$.*

Proof. By continuity, there exists $\mu > 0$ such that for all $q \in Q$ and for all $x \in W_q$ with $\|x\| < \mu$,

$$\alpha_q^l < d\psi_q(x)_1 f_q(x) < \alpha_q^u < 0 < \beta_q^l < d\psi_q(x)_2 f_q(x) < \beta_q^u,$$

wherein it follows that χ_{FQI} satisfies the conditions of $\mathscr{H}_{FQI}$ by construction. Indeed, $\psi_{\rho(i)}(c_i(t))$ satisfies the differential inclusion:

$$\frac{d}{dt}\psi_{\rho(i)}(c_i(t)) \in \{(x_1, x_2)^T \in \mathbb{R}^2 : x_1 \in [\alpha_{\rho(i)}^l, \alpha_{\rho(i)}^u],\ \ x_2 \in [\beta_{\rho(i)}^l, \beta_{\rho(i)}^u]\}.$$

Condition **R**2 guarantees that an event of χ_{FQI} occurs if and only if an event occurs in χ. Condition **R**4 guarantees that χ_{FQI} satisfies the first-quadrant interval system reset condition defined by γ_e^l and γ_e^u. Therefore, χ_{FQI} is an execution of $\mathscr{H}_{FQI}$. □

Lemma 10. *Let $\mathscr{H}$ satisfy the conditions of Theorem 10. Then there exists a function $T_{\text{escape}} : \mathbb{R}^2 \to \mathbb{R}$ such that if $\mu > 0$ is sufficiently small and $\eta > 0$ is sufficiently smaller than μ, then $T_{\text{escape}}(\eta, \mu) > 0$ and any execution χ of $\mathscr{H}$ with $\|c_0(0)\| < \eta$ satisfies $\|c_i(t)\| < \mu$ for all $t \in I_i$ with $t < T_{\text{escape}}(\eta, \mu)$.*

Furthermore, if $\hat{\eta} \leq \eta$, then $T_{\text{escape}}(\hat{\eta}, \mu) \geq T_{\text{escape}}(\eta, \mu)$.

Proof. Pick μ such that $B_q(\mu) \subset W_q$, where $B_q(\mu)$ is a ball of radius μ around the origin of dimension $\dim(D_q)$.

Say χ is an execution with $\|c_0(0)\| < \eta$ such that $\|c_i(t)\| \geq \mu$ for some $t \in I_i$ and $i \in \Lambda$. Define τ and i^* by

$$\begin{aligned} \tau &= \inf\{t : t \in I_i, i \in \Lambda, \|c_i(t)\| \geq \mu\} \\ i^* &= \min\{i \in \Lambda : \tau \in I_i\}. \end{aligned}$$

Minimality of i^* implies that either $\tau_{i^*} < \tau$ or $\tau = 0$. Indeed, say $\tau = \tau_{i^*}$, with $i^* > 0$ (recall that $\tau_0 = 0$). Then $\tau \in [\tau_{i^*-1}, \tau_{i^*}] = I_{i^*-1}$, contradicting the minimality of i^*. First, consider the case $\tau_{i^*} < \tau$ and examine the truncated execution $\hat{\chi} = (\hat{\Lambda}, \hat{I}, \hat{\rho}, \hat{C})$ defined by

- $\hat{\Lambda} = \{0, 1, \ldots, i^*\}$,

- $\hat{I} = \{\hat{I}_i\}_{i\in\hat{\Lambda}}$ with $\hat{I}_i = I_i$ for $i < i^*$ and $\hat{I}_{i^*} = [\tau_{i^*}, \tau)$,
- $\hat{\rho} = \rho|_{\hat{\Lambda}}$,
- $\hat{C} = \{\hat{c}_i\}_{i\in\hat{\Lambda}}$ with $\hat{c}_i = c_i|_{\hat{I}_i}$ for all $i \in \hat{\Lambda}$.

To simplify notation, identify χ with $\hat{\chi}$. By continuity of the vector fields f_q, there exists $M > 0$ such that $\|f_q(x)\| \leq M$ for all $x \in B_q(\mu)$ and all $q \in Q$. Moreover, integrating the vector fields f_q and applying the bound gives

$$\begin{aligned} \|c_{i^*}(\tau)\| &\leq \|c_{i^*}(\tau_{i^*})\| + M(\tau - \tau_{i^*}) \\ \|c_i(\tau_{i+1})\| &\leq \|c_i(\tau_i)\| + M(\tau_{i+1} - \tau_i), \end{aligned}$$

for $i = 0, \ldots, i^* - 1$.

A telescoping series argument combined with **R**6 gives

$$\begin{aligned} M\tau &\geq \|c_{i^*}(\tau)\| - \|c_0(0)\| - \sum_{i=0}^{i^*-1} \|c_{i+1}(\tau_{i+1})\| - \|c_i(\tau_{i+1})\| \\ &\geq \|c_{i^*}(\tau)\| - \|c_0(0)\| - K\sum_{i=0}^{i^*-1} \psi_{\rho(i)}(c_i(\tau_{i+1}))_2. \end{aligned} \tag{4.10}$$

Now last sum will be bounded. Lemma 9 implies that $\chi_{FQI} = (\Lambda, \rho, I, \Psi \circ C)$ is an execution of $\mathscr{H}_{FQI}$. Arguing as in the proof of Theorem 9, calculate an upper bound on $\psi_{\rho(i)}(c_i(\tau_{i+1}))_2$ as

$$\psi_{\rho(i)}(c_i(\tau_{i+1}))_2 \leq \psi_{\rho(i)}(c_i(\tau_i))_1 \left|\frac{\beta^u_{\rho(i)}}{\alpha^u_{\rho(i)}}\right| \leq \psi_{\rho(0)}(c_0(0))_1 \left|\frac{\beta^u_{\rho(i)}}{\alpha^u_{\rho(i)}}\right| \left|\prod_{j=0}^{i-1} \gamma^u_{(\rho(j),\rho(j+1))} \frac{\beta^u_{\rho(j)}}{\alpha^u_{\rho(j)}}\right|. \tag{4.11}$$

Summing the terms from Equation (4.11) and recalling the definition of T_{Zeno} from the proof Theorem 9 gives the bound

$$\sum_{i=0}^{i^*-1} \psi_{\rho(i)}(c_i(\tau_{i+1}))_2 \leq T_{\text{Zeno}}\left(\psi_{\rho(0)}(c_0(0))_1 \beta_{\max}\right) \leq T_{\text{Zeno}}(g(\eta)\beta_{\max}), \tag{4.12}$$

where $\beta_{\max} = \max_{q\in Q} \beta^u_q$ and $g(\eta) = \max_{\|x\|\leq\eta, q\in Q} \|\psi_q(x)\|$.

Now, a lower bound the escape time is derived in terms of η and μ. Pick η small enough so that $\max\{\eta, KT_{\text{Zeno}}(g(\eta)\beta_{\max})\} < \mu/2$. Combining Equations (4.10) and (4.12) gives the lower bound

$$\begin{aligned} \tau &\geq \frac{1}{M}(\|c_{i^*}(\tau)\| - \|c_0(0)\| - KT_{\text{Zeno}}(g(\eta)\beta_{\max})) \\ &\geq \frac{\mu - \eta - \frac{\mu}{2}}{M} \\ &= \frac{\frac{\mu}{2} - \eta}{M}. \end{aligned}$$

Define T_{escape} by $T_{\text{escape}}(\eta,\mu) := \frac{\frac{\mu}{2}-\eta}{M}$. Clearly T_{escape} increases as η decreases.

Finally, the possibility that $\tau = 0$ must be ruled out. To this end, let $i' = \min\{i \in \Lambda : \|c_i(\tau)\| \geq \mu\}$. Note that $i' > i^* = 0$. A contradiciton follows by computing

$$\begin{aligned}\|c_{i'}(\tau)\| &\leq \|c_0(\tau)\| + K\sum_{i=0}^{i'-1}\psi_{\rho(i)}(c_i(\tau))_2 \\ &\leq \|c_0(\tau)\| + KT_{\text{Zeno}}(g(\eta)\beta_{\max}) \\ &< \frac{\mu}{2}+\frac{\mu}{2}.\end{aligned}$$

The respective inequalities follow from repeated application of **R**6, Equation (4.12), and the choice of η. □

Proof of Theorem 10. Because the vector fields f_q, $q \in Q$, are Lipschitz, the continuous dynamics are always well-defined on each domain. Furthermore, since Γ is a directed cycle the dynamics of $\mathscr{H}$ are completely deterministic. Thus given $x_0 \in D_q$, there exists a unique execution χ of $\mathscr{H}$ with $c_0(0) = x_0$ such that either χ is defined for all $t \geq 0$ or χ is Zeno.

Assume that $z_q = 0$ for all $q \in Q$. Given $\epsilon > 0$ and a collection of neighborhoods $\{U\}_{q\in Q}$ with $0 \in U_q \subset D_q$, pick small constants η and μ such that

- $\{x \in D_q : \|x\| < \mu\} \subset U_q$ for all $q \in Q$.
- $T_{\text{Zeno}}(g(\eta)) < \epsilon$, where $g(\eta) = \max_{\|x\|\leq\eta, q\in Q}\|\psi_q(x)\|$.
- $T_{\text{Zeno}}(g(\eta)) < T_{\text{escape}}(\eta,\mu)$.

First it is shown that all maximal executions starting near z are Zeno. Assume for the sake of contradiction that there is an execution $\chi = (\mathbb{N}, I, \rho, C)$ with $\|c_0(0)\| < \eta$ that is not Zeno. Let $\hat{\chi}$ be the execution χ restricted to $t < T_{\text{escape}}(\eta,\mu)$. To be more precise, define i^* by

$$i^* := \min\{i \in \Lambda : T_{\text{escape}}(\eta,\mu) \in I_i\}.$$

Define $\hat{\chi} = (\hat{\lambda}, \hat{I}, \hat{\rho}, \hat{C})$ by

- $\hat{\Lambda} = \{0, 1, \ldots, i^*\}$,
- $\hat{I} = \{\hat{I}_i\}_{i\in\hat{\Lambda}}$ with $\hat{I}_i = I_i$ for $i < i^*$ and $\hat{I}_{i^*} = [\tau_{i^*}, T_{\text{escape}}(\eta,\mu))$,
- $\hat{\rho} = \rho|_{\hat{\Lambda}}$,
- $\hat{C} = \{\hat{c}_i\}_{i\in\hat{\Lambda}}$ with $\hat{c}_i = c_i|_{\hat{I}_i}$ for all $i \in \hat{\Lambda}$.

If μ is sufficiently small, then Lemma 9 combined with Lemma 10 implies that $\hat{\chi}$ gives rise to an execution χ_{FQI} of $\mathscr{H}_{FQI}$ that is defined for all $t \in [0, T_{\text{escape}}(\eta,\mu))$. Furthermore, the initial

condition of χ_{FQI} satisfies $\|\psi_{\rho(0)}(c_0(0))\| \leq g(\eta)$. Now Theorem 9 implies that the Zeno time for χ_{FQI} is at most $T_{\text{Zeno}}(g(\eta))$. But recalling that $T_{\text{Zeno}}(g(\eta)) < T_{\text{escape}}(\eta, \mu)$, we find that χ_{FQI} is defined past its Zeno time, a contradiction.

Finally, by the choice of η and μ, if the initial condition of a maximal execution, χ, satisfies $\|c_0(0)\| < \eta$, then $c_i(t) \in U_{\rho(i)}$ for all $i \in \Lambda$ and $t \in I_i$, and the Zeno time satisfies

$$\tau_\infty \leq T_{\text{Zeno}}(g(\eta)) < \epsilon.$$

Therefore z is bounded-time locally Zeno stable. □

First-quadrant hybrid systems. Theorem 10 immediately generalizes the sufficient conditions from [72] for Zeno behavior in hybrid systems defined on the first quadrant of $\mathbb{R}^2$.

Definition 13. A *first-quadrant hybrid system* is a tuple

$$\mathscr{H}_{FQ} = (\Gamma, D, G, R, F)$$

where

- $\Gamma = (Q, E)$ is a directed cycle as in Definition 5.
- $D = \{D_q\}_{q \in Q}$ where for all $q \in Q$, $D_q = \mathbb{R}^2_{\geq 0} = \{(x_1, x_2)^T \in \mathbb{R}^2 : x_1 \geq 0,\ x_2 \geq 0\}$.
- $G = \{G_e\}_{q \in Q}$ where for all $e \in E$, $G_e = \{(x_1, x_2)^T \in \mathbb{R}^2_{\geq 0} : x_1 = 0,\ x_2 \geq 0\}$.
- $R = \{R_e\}_{e \in E}$ where for all $e \in E$, $R_e(0, x_2) = (r_e(x_2), 0)$ and $r_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$.
- $F = \{f_q\}_{q \in Q}$ where for all $q \in Q$, f_q is a vector field on $D_q = \mathbb{R}^2_{\geq 0}$.

Corollary 2. *Let $\mathscr{H} = (\Gamma, D, R, F)$ be a first-quadrant hybrid system. If $r'_e(0) > 0$ for all $e \in E$, $f_q(0)_1 < 0 < f_q(0)_2$ for all $q \in Q$ and*

$$\prod_{i=0}^{|Q|-1} \left| r'_{e_i}(0) \frac{f_{q_i}(0)_2}{f_{q_i}(0)_1} \right| < 1,$$

then $\{0_q\}_{q \in Q}$ bounded-time asymptotically Zeno stable, where 0_q is the origin of D_q.

Proof. Let ψ_q be the identity for all $q \in Q$. Let γ^l_e and γ^u_e be such that $0 < \gamma^l_e < r'_e(0) < \gamma^u_e$ and

$$\prod_{i=0}^{|Q|-1} \left| \gamma^u_{e_i} \frac{f_{q_i}(0)_2}{f_{q_i}(0)_1} \right| < 1.$$

Let $K = \max_{e \in E} \gamma_e^u$. Routine calculations verify that the conditions of Theorem 10 hold on a sufficiently small neighborhood of the origin. Since the origin is an isolated Zeno equilibrium, Corollary 1 implies asymptotic convergence. □

4.5 Application to Simple Hybrid Mechanical Systems

This section develops Zeno stability theory for a simple model of mechanical systems undergoing impacts, known as Lagrangian hybrid systems. First, Lagrangian hybrid systems are defined. Then, Theorem 10 is applied to give sufficient conditions for Zeno behavior in Lagrangian hybrid systems based on the the value of the vector field at a single point. Finally, examples illustrate the theory. For more on Lagrangian hybrid systems, see [9, 10, 84, 85, 86, 100, 101].

4.5.1 Lagrangian Hybrid Systems

Lagrangians. Consider a configuration space[2] Θ and a Lagrangian $L : T\Theta \to \mathbb{R}$ given in coordinates by:

$$L(\theta, \dot{\theta}) = \frac{1}{2}\dot{\theta}^T M(\theta)\dot{\theta} - U(\theta) \tag{4.13}$$

where $M(\theta)$ is positive definite and symmetric and $U(\theta)$ is the potential energy. For the sake of simplicity, assume $\Theta \subset \mathbb{R}^n$ since all the results can be proven in a coordinate chart. The equations of motion are then given in coordinates by the Euler-Lagrange equations,

$$\frac{d}{dt}\frac{\partial L}{\partial \dot{\theta}} - \frac{\partial L}{\partial \theta} = 0.$$

In the case of Lagrangians of the form given in (4.13), the Lagrangian vector field, f_L, associated to L takes the familiar form

$$\dot{x} = f_L(x) = \begin{pmatrix} \dot{\theta} \\ M(\theta)^{-1}(-C(\theta, \dot{\theta})\dot{\theta} - N(\theta)) \end{pmatrix}, \tag{4.14}$$

where $x = (\theta^T, \dot{\theta}^T)^T$, $C(\theta, \dot{\theta})$ is the *Coriolis matrix* and $N(\theta) = \frac{\partial U}{\partial \theta}(\theta)$.

This process of associating a dynamical system to a Lagrangian will be mirrored in the setting of hybrid systems. First, hybrid Lagrangians will be introduced.

Definition 14. A *hybrid Lagrangian* is a tuple, $\text{Ł} = (\Theta, L, h)$, where

- $\Theta \subset \mathbb{R}^n$ is the configuration space,

[2]Note that the configuration space is written as Θ rather than Q, due to the fact that Q denotes the vertices of the graph of a hybrid system.

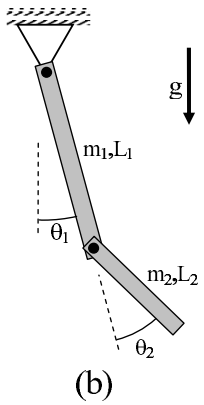

Figure 4.7: The double pendulum with mechanical stop

- $L : T\Theta \to \mathbb{R}$ is a Lagrangian of the form given in (4.13),
- $h : \Theta \to \mathbb{R}$ is a *unilateral constraint function*, where 0 is assumed to be a regular value of h (to ensure that $h^{-1}(\{0\})$ is a smooth manifold).

To concretely illustrate the hybrid Lagrangian concepts of the rest of the paper, consider a double pendulum with a mechanical stop (Figure 4.7[3]).

***Example* 4** (Double Pendulum)**.** The double pendulum consists of two rigid links of masses m_1, m_2, lengths L_1, L_2, and uniform mass distribution, which are attached by passive joints, while a mechanical stop dictates the range of motion of the second link. The example serves as a simplified model of a leg with a passive knee and a mechanical stop, which is widely investigated in the robotics literature in the context of passive dynamics of bipedal walkers (see [88] and [89]). In this case

$$\mathbf{P} = (\Theta_{\mathbf{P}}, L_{\mathbf{P}}, h_{\mathbf{P}}),$$

where $\Theta_{\mathbf{P}} = \mathbb{S}^1 \times \mathbb{S}^1 = \mathbb{T}^2$, $q = (\theta_1, \theta_2)$, and

$$L_{\mathbf{P}}(\theta, \dot{\theta}) = \frac{1}{2}\dot{\theta}^T M(q)\dot{\theta} + \left(\frac{1}{2}m_1 L_1 + m_2 L_1\right) g \cos\theta_1 + \frac{1}{2}m_2 L_2 g \cos(\theta_1 + \theta_2),$$

with the 2×2 inertia matrix $M(\theta)$ given by

$$M(\theta) = \begin{pmatrix} m_1 L_1^2/3 + m_2(L_1^2 + L_2^2/3 + L_1 L_2 \cos\theta_2) & m_2(3L_1 L_2 \cos\theta_2 + 2L_2^2)/6 \\ m_2(3L_1 L_2 \cos\theta_2 + 2L_2^2)/6 & m_2 L_2^2/3 \end{pmatrix}. \tag{4.15}$$

Finally, the constraint that represents the mechanical stop is given by $h_{\mathbf{P}}(q) = \theta_2 \geq 0$. So, for this example, there are nontrivial dynamics and a trivial unilateral constraint function.

[3]Figure by Yizhar Or

Domains from constraints. Given a smooth (unilateral constraint) function $h : \Theta \to \mathbb{R}$ on a configuration space Θ such that 0 is a regular value of h, a domain and a guard can be explicitly constructed. Define the domain, D_h, as the manifold (with boundary):

$$D_h = \{(\theta, \dot{\theta}) \in T\Theta : h(\theta) \geq 0\}.$$

Similarly, there is an associated guard, G_h, defined as the following submanifold of D_h:

$$G_h = \{(\theta, \dot{\theta}) \in T\Theta : h(\theta) = 0 \text{ and } dh(\theta)\dot{\theta} \leq 0\},$$

where $dh(\theta) = \begin{pmatrix} \frac{\partial h}{\partial \theta_1}(\theta) & \cdots & \frac{\partial h}{\partial \theta_n}(\theta) \end{pmatrix}$. Note that 0 is a regular value of h if and only if $dh(\theta) \neq 0$ whenever $h(\theta) = 0$.

Lagrangian Hybrid Systems. Given a hybrid Lagrangian $Ł = (\Theta, L, h)$, the *Lagrangian hybrid system associated to* Ł is the hybrid system

$$\mathscr{H}_Ł = (\Gamma = (\{q\}, \{(q, q)\}), D_Ł, G_Ł, R_Ł, F_Ł),$$

where $D_Ł = \{D_h\}$, $F_Ł = \{f_L\}$, $G_Ł = \{G_h\}$, and $R_Ł = \{R_h\}$, with the reset map given by the Newtonian impact equation $R_h(\theta, \dot{\theta}) = (\theta, P(\theta, \dot{\theta}))$, with

$$P(\theta, \dot{\theta}) = \dot{\theta} - (1 + e)\frac{dh(\theta)\dot{\theta}}{dh(\theta)M(\theta)^{-1}dh(\theta)^T} M(\theta)^{-1}dh(\theta)^T. \tag{4.16}$$

Here $0 \leq e \leq 1$ is the coefficient of restitution.

***Example* 5** (Double Pendulum)**.** From the hybrid Lagrangian, $\mathbf{P} = (\Theta_\mathbf{P}, L_\mathbf{P}, h_\mathbf{P})$, construct the hybrid system

$$\mathscr{H}_\mathbf{P} = (\Gamma = (\{q\}, \{(q, q)\}), D_\mathbf{P}, G_\mathbf{P}, R_\mathbf{P}, F_\mathbf{P}),$$

where

$$\begin{aligned} D_{h_\mathbf{P}} &= \{(\theta, \dot{\theta}) \in \mathbb{T}^2 \times \mathbb{R}^2 : \theta_2 \geq 0\}, \\ G_{h_\mathbf{P}} &= \{(\theta, \dot{\theta}) \in \mathbb{T}^2 \times \mathbb{R}^2 : \theta_2 = 0 \text{ and } \dot{\theta}_2 \leq 0\}, \end{aligned}$$

and $R_{h_\mathbf{P}}(\theta, \dot{\theta}) = (\theta, P_{h_\mathbf{P}}(\theta, \dot{\theta}))$ is computed on $G_{h_\mathbf{P}}$ to be

$$P_{h_\mathbf{P}}(\theta, \dot{\theta}) = \begin{pmatrix} \dot{\theta}_1 + \rho\dot{\theta}_2 \\ -e\dot{\theta}_2 \end{pmatrix} \quad \text{with} \quad \rho = (1 + e)\frac{3m_2L_1L_2 + 2m_2L_2^2}{2m_1L_1^2 + 6m_2L_1^2 + 2m_2L_2^2 + 6m_2L_1L_2}.$$

The vector field is computed as

$$f_{L_{\mathbf{P}}}(\theta,\dot{\theta}) = \begin{pmatrix} \dot{\theta} \\ M(\theta)^{-1} \begin{pmatrix} \alpha\dot{\theta}_2^2 \sin\theta_2 - \beta\sin(\theta_1+\theta_2) - \gamma\sin\theta_1 \\ -\alpha\dot{\theta}_1^2 \sin\theta_2 - \beta\sin(\theta_1+\theta_2) \end{pmatrix} \end{pmatrix}$$

where $M(\theta)$ is the mass matrix from Equation (4.15) and the constants α, β, and γ are defined by

$$\alpha = \frac{1}{2}m_2 L_1 L_2, \qquad \beta = \frac{1}{2}m_2 L_2 g, \qquad \gamma = \left(\frac{1}{2}m_1 L_1 + m_2 L_1\right) g.$$

4.5.2 Sufficient Conditions for Zeno Behavior in Lagrangian Hybrid Systems

This subsection presents sufficient conditions for bounded-time local Zeno stability of Lagrangian hybrid systems, based on an explicitly constructed Lyapunov-like function. The paper [84] proves a special case of the main result in this section, Theorem 11, for a class of Lagrangian hybrid systems with configuration manifolds of dimension two. If the potential energy is a convex function and the domain specified by the unilateral constraint is a convex set, global Zeno stability results have been proved in [85]. Of course, the convexity assumptions preclude the fundamentally local phenomena occurring in the examples of this chapter.

First, however, the Zeno equilibria of Lagrangian hybrid systems are studied, observing that isolated Zeno equilibria only occur in systems with one-dimensional configuration manifolds. Thus, no Lagrangian hybrid system with configuration manifold of dimension greater than one can have bounded-time asymptotically stable Zeno equilibria.

Zeno equilibria in Lagrangian hybrid systems. If $\mathscr{H}_{\mathrm{L}}$ is a Lagrangian hybrid system, then applying the definition of Zeno equilibria and examining the special form of the reset maps shows that $z = \{(\theta^*, \dot{\theta}^*)\}$ is a Zeno equilibrium if and only if

$$f_{\mathrm{L}}(\theta^*,\dot{\theta}^*) \neq 0, \quad h(\theta^*) = 0, \quad dh(\theta^*)\dot{\theta}^* \leq 0, \quad \dot{\theta}^* = P(\theta^*,\dot{\theta}^*).$$

Furthermore, the form of P implies that $\dot{\theta}^* = P(\theta^*,\dot{\theta}^*)$ holds if and only if $dh(\theta^*)\dot{\theta}^* = 0$. Therefore the set of all Zeno equilibria for a Lagrangian hybrid system is given by the surfaces in $T\Theta$:

$$Z_h = \{(\theta,\dot{\theta}) \in T\Theta : f_{\mathrm{L}}(\theta,\dot{\theta}) \neq 0,\ h(\theta) = 0,\ dh(\theta)\dot{\theta} = 0\}.$$

Note that if $\dim(\Theta) > 1$, the Lagrangian hybrid system has no isolated Zeno equilibria.

Theorem 11. *Let $\mathscr{H}_L$ be a Lagrangian hybrid system and $(\theta^*, \dot{\theta}^*) \in D_h$. If the coefficient of restitution satisfies $0 < e < 1$ and $(\theta^*, \dot{\theta}^*)$ satisfies*

$$h(\theta^*) = 0, \quad \dot{h}(\theta^*, \dot{\theta}^*) = 0, \quad \ddot{h}(\theta^*, \dot{\theta}^*) < 0,$$

then $\{(\theta^*, \dot{\theta}^*)\}$ is a bounded-time locally stable Zeno equilibrium.

Here $\dot{h}(\theta^*, \dot{\theta}^*) = dh(\theta^*)\dot{\theta}^*$ and

$$\ddot{h}(\theta^*, \dot{\theta}^*) = (\dot{\theta}^*)^T H(h(\theta^*))\dot{\theta}^* + dh(\theta^*)M(\theta^*)^{-1}(-C(\theta^*, \dot{\theta}^*)\dot{\theta}^* - N(\theta^*)),$$

where $H(h(\theta^*))$ is the Hessian of h at θ^*.

Proof. First note that $\{(\theta^*, \dot{\theta}^*)\}$ is a Zeno equilibrium. Indeed, $d\dot{h}(\theta^*, \dot{\theta}^*)f_{\mathbf{L}}(\theta^*, \dot{\theta}^*) = \ddot{h}(\theta^*, \dot{\theta}^*) \neq 0$ implies that $f_{\mathbf{L}}(\theta^*, \dot{\theta}^*) \neq 0$. Then the conditions $h(\theta^*) = 0$ and $\dot{h}(\theta^*, \dot{\theta}^*) = 0$ imply that $\{(\theta^*, \dot{\theta}^*)\}$ is a Zeno equilibrium.

Let V be a small neighborhood of $(\theta^*, \dot{\theta}^*)$ and assume (by passing to a coordinate chart) that $V \subset \mathbb{R}^{2n}$ with Euclidean norm. Let K satisfy

$$K > \frac{1+e}{2} \frac{\|M(\theta^*)^{-1}dh(\theta^*)^T\|}{dh(\theta^*)M(\theta^*)^{-1}dh(\theta^*)^T}.$$

The proof proceeds by verifying that the constants $\gamma_h^u = \gamma_h^l = e$, K and the function

$$\psi_h(\theta, \dot{\theta}) = \begin{pmatrix} \dot{h}(\theta, \dot{\theta}) + \sqrt{\dot{h}(\theta, \dot{\theta})^2 + 2h(\theta)} \\ -\dot{h}(\theta, \dot{\theta}) + \sqrt{\dot{h}(\theta, \dot{\theta})^2 + 2h(\theta)} \end{pmatrix} \tag{4.17}$$

satisfy conditions **R**1–**R**6 on V.

R1: Since $(\theta^*, \dot{\theta}^*)$ is a Zeno equilibrium, $h(\theta^*) = 0$ and $\dot{h}(\theta^*, \dot{\theta}^*) = dh(\theta^*)\dot{\theta}^* = 0$. Thus $\psi_h(\theta^*, \dot{\theta}^*) = 0$.

R2: Since $h(\theta) \geq 0$, $\psi_h(\theta, \dot{\theta})_1 = 0$ if and only if $h(\theta) = 0$ and $\dot{h}(\theta, \dot{\theta}) \leq 0$. So $\psi_h(\theta, \dot{\theta})_1 = 0$ if and only if $(\theta, \dot{\theta}) \in G_h$.

R3: The square root in the definition of ψ_h creates some differentiability problems at the Zeno equilibrium.

Assume V is small enough that V contains no equilibria of $f_{\mathbf{L}}$. Then $(D_h \setminus Z_h) \cap V$ has the form

$$(D_h \setminus Z_h) \cap V = \{(\theta, \dot{\theta}) \in V : h(\theta) > 0 \text{ or } \dot{h}(\theta, \dot{\theta}) \neq 0\},$$

and that ψ_h is continuously differentiable on $(D_h \setminus Z_h) \cap V$ with Lie derivative given by

$$\dot{\psi}_h(\theta,\dot{\theta}) = \begin{pmatrix} \ddot{h}(\theta,\dot{\theta}) + \frac{\dot{h}(\theta,\dot{\theta})}{\sqrt{\dot{h}(\theta,\dot{\theta})^2+2h(\theta)}}\left(\ddot{h}(\theta,\dot{\theta})+1\right) \\ -\ddot{h}(\theta,\dot{\theta}) + \frac{\dot{h}(\theta,\dot{\theta})}{\sqrt{\dot{h}(\theta,\dot{\theta})^2+2h(\theta)}}\left(\ddot{h}(\theta,\dot{\theta})+1\right) \end{pmatrix}. \tag{4.18}$$

Recall that $\ddot{h}(\theta^*,\dot{\theta}^*) < 0$. It follows from the definitions that scaling h by a positive constant does not change D_h, G_h, or R_h. Therefore it can be assumed that $\ddot{h}(\theta^*,\dot{\theta}^*) = -1$.

While the function $\frac{\dot{h}(\theta,\dot{\theta})}{\sqrt{\dot{h}(\theta,\dot{\theta})^2+2h(\theta)}}$ may not have a unique limit as $(\theta,\dot{\theta}) \to (\theta^*,\dot{\theta}^*)$ it remains bounded on $(D_h \setminus Z_h) \cap V$:

$$\frac{|\dot{h}(\theta,\dot{\theta})|}{\sqrt{\dot{h}(\theta,\dot{\theta})^2 + 2h(\theta)}} \leq 1.$$

Therefore, $\dot{\psi}_h$ has the well-defined limit

$$\lim_{(\theta,\dot{\theta})\in(D_h\setminus Z_h)\cap V,\ (\theta,\dot{\theta})\to(\theta^*,\dot{\theta}^*)} \dot{\psi}_h(\theta,\dot{\theta}) = \begin{pmatrix} -1 \\ 1 \end{pmatrix}. \tag{4.19}$$

Since the differentiability problems only arise on the guard, and in particular only on the Zeno equilibria, the limit in Equation (4.19) suffices for the evaluation in **R**3.

R4: Let $(\theta,\dot{\theta}) \in G_h$. Then $h(\theta) = 0$ and $\dot{h}(\theta,\dot{\theta}) \leq 0$. So, substitution into Equation (4.17) gives $\psi_h(\theta,\dot{\theta}) = (0, 2|\dot{h}(\theta,\dot{\theta})|)^T$.

Multiplying both sides of Equation (4.16) on the left by $dh(\theta)$ and the definition of $\dot{h}(\theta,\dot{\theta})$ gives $\dot{h}(R_h(\theta,\dot{\theta})) = -e\dot{h}(\theta,\dot{\theta})$.

Therefore $\psi_h(R_h(\theta,\dot{\theta})) = (2e|\dot{h}(\theta,\dot{\theta})|, 0)^T$. So if $\gamma_h^l = \gamma_h^u = e$, **R**4 holds with $\psi_h(R_h(\theta,\dot{\theta}))_1 \in [e\psi_h(\theta,\dot{\theta})_2, e\psi_h(\theta,\dot{\theta})_2]$.

R5:

$$\left|\gamma_h^u \frac{d\psi_h(\theta^*,\dot{\theta}^*)_2 f_L(\theta^*,\dot{\theta}^*)}{d\psi_h(\theta^*,\dot{\theta}^*)_1 f_L(\theta^*,\dot{\theta}^*)}\right| = \left|e\frac{1}{-1}\right| = e < 1.$$

R6: Take a point $(\theta,\dot{\theta}) \in G_h \cap V$ (this is the only step that requires a norm, and hence the coordinate chart on V). The growth due to the reset map can be bounded as follows,

$$\begin{aligned} &\|R_h(\theta,\dot{\theta}) - (\theta^*,\dot{\theta}^*)\| \\ &= \left\|(\theta,\dot{\theta}) - (\theta^*,\dot{\theta}^*) - \left(0, (1+e)\frac{dh(\theta)\dot{\theta}}{dh(\theta)M(\theta)^{-1}dh(\theta)^T}M(\theta)^{-1}dh(\theta)^T\right)\right\| \\ &\leq \|(\theta,\dot{\theta}) - (\theta^*,\dot{\theta}^*)\| + (1+e)\frac{|dh(\theta)\dot{\theta}|}{dh(\theta)M(\theta)^{-1}dh(\theta)^T}\|M(\theta)^{-1}dh(\theta)^T\|. \end{aligned}$$

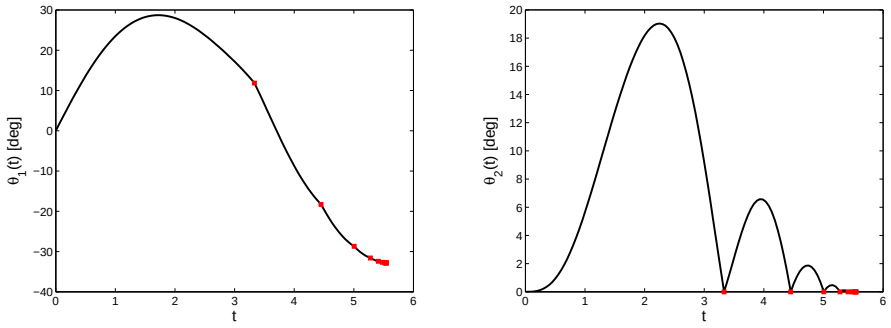

Figure 4.8: A simulation of an execution of the double pendulum with a mechanical stop that appears to be Zeno. See Remark 7 for a discussion on numerical proofs Zeno behavior in simulations.

Recall that $\psi_h(\theta,\dot\theta)_2 = 2|dh(\theta)\dot\theta|$. Plugging in the definition of K proves **R**6:

$$\|R_h(\theta,\dot\theta) - (\theta^*,\dot\theta^*)\| \leq \|(\theta,\dot\theta) - (\theta^*,\dot\theta^*)\| + K\psi_h(\theta,\dot\theta)_2.$$

Since **R**1–**R**6 hold, Theorem 10 implies that there is a neighborhood W of $(\theta^*,\dot\theta^*)$ with $W \subset V$ such that there is a unique Zeno execution with $c_0(0) = x$ for all $x \in W$. □

***Example* 6** (Double Pendulum)**.** Recall that the double pendulum system, $\mathscr{H}_{\mathbf{P}}$ has a trivial unilateral constraint: $h_{\mathbf{P}}(\theta) = \theta_2$ and $\dot h_{\mathbf{P}} = \dot\theta_2$. Whenever $\theta_2 = 0$ and $\dot\theta_2 = 0$ the second derivative of the unilateral constraint reduces to

$$\ddot h_{\mathbf{P}}(\theta_1,\theta_2,\dot\theta_1,\dot\theta_2) = \frac{g\sin\theta_1}{\tilde L} < 0, \text{ where } \tilde L = \tfrac{(4m_1+3m_2)L_1L_2}{3(m_1(L_1+2L_2)m_2L_2)}.$$

Thus, if $(\theta_1^*,\theta_2^*,\dot\theta_1^*,\dot\theta_2^*)$ satisfies $\theta_2^* = 0$, $\dot\theta_2^* = 0$, and $\sin(\theta_1^*) < 0$, then all executions with initial conditions near $(\theta_1^*,\theta_2^*,\dot\theta_1^*,\dot\theta_2^*)$ are Zeno (Figure 4.8).

***Example* 7** (Ball on a Circle)**.** With Zeno stability tools in hand, return to the ball bouncing on a circle from Example 3. Basic calculations show that the bouncing ball hybrid system, $\mathscr{H}_{\mathbb{B}}$, is the Lagrangian hybrid system associated to the hybrid Lagrangian $\mathbb{B} = (\mathbb{R}^2, L_{\mathbb{B}}, h_{\mathbb{B}})$, where

$$L_{\mathbb{B}}(x,\dot x) = \frac{1}{2}m\|\dot x\|^2 - mgx_2, \qquad h_{\mathbb{B}}(x) = \|x\|^2 - 1.$$

The conditions in Equation (4.3) for Zeno stability follow from Theorem 11.

***Remark* 7.** Given the simple dynamics of the bouncing ball system, $\mathscr{H}_{\mathbb{B}}$, conservative estimates of T_{escape} and T_{Zeno} can be numerically computed to prove Zeno behavior in experiments. Indeed, the execution from Figure 4.6(b) was proved to be Zeno, numerically, using tighter bounds on T_{Zeno}

from [102].

The double pendulum, $\mathscr{H}_{\mathbf{P}}$ has complicated dynamics, and Zeno behavior in the execution depicted in Figure 4.8 was assumed to be Zeno, heuristically, based the value of $\dot{h}_{\mathbf{P}}$ at collisions.

4.6 Conclusion

In this chapter, it was observed that a bounded-time locally stable Zeno equilibrium displayed asymptotic stability if and only if it was isolated. The tight link between non-asymptotic Zeno stability and non-isolated Zeno equilibria highlighted the differences between classical Lyapunov stability theory and Zeno stability theory. In particular, the theory of dynamical systems near isolated equilibria usually provides enough traction that researchers and practitioners rarely need to consider the complications of non-isolated equilibria. Contrast this to the situation that is found in hybrid systems; to study non-asymptotic Zeno stability, there is no choice but to examine non-isolated Zeno equilibria. Because most of the existing conditions for Zeno behavior required either isolated Zeno equilibria or asymptotically stable Zeno equilibria, they all had similar limitations.

To reason about both asymptotic and non-asymptotic Zeno stability, Lyapunov-like sufficient conditions for bounded-time local Zeno stability were presented that flexibly apply to isolated and non-isolated Zeno equilibria. The proof methodology had two main components. First, a class of hybrid systems with simple conditions for Zeno stability was defined. Then, special structured (Lyapunov-like) functions were proposed to map executions of interesting hybrid systems to executions of the simple Zeno hybrid systems.

The main subtlety of the Lyapunov-like theorem arises from its local nature. In particular, executions must remain "close enough" to the Zeno equilibrium in question, so that the reduction conditions remain valid. The locality turns out to be crucial for examples, such as the bouncing ball on a circular surface, in which Zeno and non-Zeno executions are separated by slight perturbations in initial conditions.

Applications to Lagrangian hybrid systems showed that the sufficient conditions for local Zeno stability can handle some nontrivial, high-dimensional hybrid systems. Furthermore, the Lyapunov-like sufficient conditions specialize to algebraic constraints on the Zeno equilibria. In particular, in Lagrangian hybrid systems, Zeno stability properties are inferred from the zero-order approximation to the vector fields at the Zeno equilibria, similar to the local approximation results of [72] and [74].

The extensions of this work on Lagrangian hybrid systems, [9, 10], show that Zeno stability can be used, in a practical sense, to study the transition from bouncing to sliding in mechanical systems. To study more complex systems, such as bipedal walkers, rigid bodies, and other locomotion systems, the results must be extended to cover more complex graph structures, and perhaps nonsmooth unilateral constraints.

Chapter 5

Conclusion

The areas of motor control, spiking neuron dynamics, and locomotion are ideal for control theorists. Each of those areas employs sophisticated control theory but leaves many dynamical phenomena unexplored. In many cases, as in this dissertation, the unexplored phenomena can be handled with elementary techniques. The work in this dissertation represents just a few directions in which progress can be made. Furthermore, all the results in this dissertation pave the way for related research. Chapter 2 studies distributed LQG in order to understand how humans make efficient, reliable movements, in spite of the fact that the underlying control architecture is a distributed network of relatively slow subsystems. This chapter only solved the state feedback case. Output feedback, as well as other cost functions should be studied to increase biological relevance. Chapter 3 explores the use of spike-based communication schemes for networked control. To improve on the work of Chapter 3, the coding strategies of known feedback loops in the body should be studied. Finally, Chapter 4 studies the relationship between Zeno behavior and Lyapunov stability in order to understand mode transitions in mechanical systems. Explorations of the connections between the work of Chapter 4 and bipedal walking are already underway, and should continue.

To make progress, the work in this thesis uses the approach of isolating phenomena and studying them in an abstract setting, removed from the physiological motivations. As a consequence, the physiological implications of the theory developed are not immediately obvious. Nonetheless, the phenomena studied are inherent in the motivating problems. The problems studied in this thesis are chosen as reasonable, and perhaps necessary, steps toward a deeper mathematical understanding of motor control and locomotion.

One goal of this dissertation, which is far from being achieved, is to use physiologically inspired control problems to increase the dialog between control theorists and biological experts who study related problems. For instance, in neuroscience, it is taken for granted that control in the motor system is distributed across several regions throughout the nervous system. It is also widely recognized that feedback control theory is the natural framework for reasoning about human motor control. In spite of these facts, little is known about how results from distributed control apply to the motor

system. It is hoped that the work in Chapter 2 can serve as a starting point for conversations between distributed control theorists and motor control experts. Much in the same way, the work in Chapter 3 is intended to serve as a bridge between networked control theorists and spiking neuron experts.

To improve the dialog between the locomotion and the hybrid systems communities, dynamical phenomena from locomotion should be isolated and studied in a general hybrid systems framework. Locomotion experts, both in biology and robotics, understand that their models are hybrid systems, but the technical challenges of locomotion mainstream hybrid systems topics, such as formal verification and switched linear systems, are quite different. Chapter 4 isolates the mechanisms producing Zeno behavior in mechanical models and studies them in a hybrid systems framework. It is likely that taking a similar approach to other locomotion phenomena, such as the exploitation of passive dynamics in walking, would be fruitful for both communities.

Above all, this thesis demonstrates, through examples, the abundance of beautiful mathematical structures resulting from physiological phenomena. In the coming years, insight into the nervous system's control strategies is bound to give rise to countless unimagined (and currently unimaginable) mathematical structures within control theory. Perhaps more importantly, these insights may lead to more precise and efficient control strategies in technological systems.

Bibliography

[1] W. Bialek, F. Rieke, R. R. de Ruyter van Stevenink, and D. Warland, "Reading a neural code," *Science*, vol. 252, no. 5014, pp. 1854–1857, 1991.

[2] E. M. Izhikevich, "Simple model of spiking neurons," *IEEE Transactions on Neural Networks*, vol. 14, no. 6, pp. 1569–1572, 2003.

[3] J. Zhang, K. H. Johansson, J. Lygeros, and S. Sastry, "Zeno hybrid systems," *International Journal Robust and Nonlinear Control*, vol. 11, no. 2, pp. 435–451, 2001.

[4] E. D. Sontag, "Some new directions in control theory inspired by systems biology," *Systems Biology*, vol. 1, no. 1, 2004.

[5] E. Todorov and W. Li, "A generalized iterative lgq method for locally-optimal feedback control of constrained nonlinear systems," in *American Control Conference*, 2005.

[6] A. Rantzer, "Linear quadratic team theory revisited," in *American Control Conference*, 2006.

[7] E. R. Kandel, J. H. Schwartz, and T. M. Jessell, *Principles of Neural Science*. Elsevier Science Publishing, third ed., 1991.

[8] W. Gerstner and W. M. Kistler, *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, 2002.

[9] Y. Or and A. D. Ames, "Formal and practical completion of Lagrangian hybrid systems," in *ASME/IEEE American Control Conference*, 2009.

[10] Y. Or and A. D. Ames, "Existence of periodic orbits with Zeno behavior in completed Lagrangian hybrid systems," in *Hybrid Systems: Computation and Control*, 2009.

[11] A. D. Ames, "Characterizing knee-bounce in bipedal robotic walking: A zeno behavior approach," in *Hybrid Systems: Computation and Control*, 2011.

[12] Y.-C. Ho and K.-C. Chu, "Team decision theory and information structures in optimal control problems—part i," *IEEE Transactions on Automatic Control*, vol. 17, no. 1, 1972.

[13] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal of Control*, vol. 6, no. 1, 1968.

[14] N. R. Sandell and M. Athans, "Solution of some nonclassical lqg stochastic decision problems," *IEEE Transactions on Automatic Control*, vol. 19, no. 2, pp. 108–116, 1974.

[15] B.-Z. Kurtaran and R. Sivan, "Linear-quadratic-gaussian control with one-step-delay sharing pattern," *IEEE Transactions on Automatic Control*, vol. 19, no. 5, pp. 571–574, 1974.

[16] T. Yoshikawa, "Dynamic programming approach to decentralized stochastic control problems," *IEEE Transactions on Automatic Control*, vol. 20, no. 6, pp. 796–797, 1975.

[17] X. Qi, M. V. Salapaka, P. G. Voulgaris, and M. Khammash, "Structured optimal and robust control with multiple criteria: A convex solution," *IEEE Transactions on Automatic Control*, vol. 49, pp. 1623–1640, 2004.

[18] M. Rotkowitz and S. Lall, "A characterization of covex problems in decentralized control," *IEEE Transactions on Automatic Control*, vol. 51, no. 2, pp. 1984–1996, 2006.

[19] A. Rantzer, "A separation principle for distributed control," in *IEEE Conference on Decision and Control*, 2006.

[20] J. Swigart and S. Lall, "An explicit state-space solution for a decentralized two-player optimal linear-quadratic regulator," in *American Control Conference*, pp. 6385–6390, 2010.

[21] J. Swigart and S. Lall, "An explicit dynamic programming solution for a decentralized two-player optimal linear-quadratic regulator," in *Symposium on the Mathematical Theory of Networks*, pp. 1443 – 1447, 2010.

[22] P. Shah and P. Parrilo, "$\mathcal{H}_2$-optimal decentralized control over posets: A state space solution for state-feedback," in *IEEE Conference on Decision and Control*, 2010.

[23] J. He, W. S. Leving, and G. E. Loeb, "Feedback gains for correcting small perturbations to standing posture," *IEEE Transactions on Automatic Control*, vol. 36, no. 3, pp. 322–332, 1991.

[24] A. D. Kuo, "An optimal control model for analyzing human postural balance," *IEEE Transactions on Biomedical Engineering*, vol. 42, no. 1, pp. 87–101, 1995.

[25] E. Todorov and M. I. Jordan, "Optimal feedback control as a theory of motor coordination," *Nature Neuroscience*, vol. 5, no. 11, 2002.

[26] S. H. Scott, "Optimal feedback control and the neural basis of volitional motor control," *Nature Reviews Neuroscience*, vol. 5, pp. 534–546, 2004.

[27] R. Shadmehr and J. W. Krakauer, "A computational neuroanatomy for motor control," *Experimental Brain Research*, vol. 185, pp. 359–381, 2008.

[28] A. J. Nagengast, D. A. Braun, and D. M. Wolpert, "Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty," *PLoS Computational Biology*, vol. 6, 2010.

[29] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan, "An internal model for sensorimotor integration," *Science*, vol. 269, no. 5232, pp. 1880–1882, 1995.

[30] K. Körding and D. M. Wolpert, "Bayesian integration in sensorimotor learning," *Nature*, vol. 427, pp. 244–247, 2004.

[31] D. M. Wolpert, K. Doya, and M. Kawato, "A unifying computational framework for motor control and social interaction," *Philosphical Transactions of the Royal Society of London B*, vol. 358, pp. 593–602, 2003.

[32] R. Grush, "The emulation theory of representation: Motor control, imagery, and perception," *Behavioral and Brain Sciences*, vol. 27, pp. 377–442, 2004.

[33] G. E. Loeb, I. E. Brown, and E. J. Cheng, "A hierarchical foundation for models of sensorimotor control," *Experimental Brain Reseasrch*, vol. 126, pp. 1–18, 1999.

[34] W. Li, E. Todorov, and X. Pan, "Hierarchical optimal control of redundant biomechanical systems," in *International Conference of the IEEE Engineering in Medicine and Biology Society*, 2004.

[35] J. C. Spagna, D. I. goldman, P. C. Lin, D. E. Koditschek, and R. J. Full, "Distributed mechanical feedback in arthropods and robots simlifies control of rapid running on challenging terrain," *Bioinspiration and Biomimetics*, vol. 2, pp. 9–18, 2007.

[36] F. L. Lewis, L. Xie, and D. Popa, *Optimal and Robust Estimation: With an Introduction to Stochastic Control Theory*. CRC Press, second ed., 2008.

[37] E. von Holst and H. Mittlestädt, "Das reafferenzprinzip," *Die Naturawissenschaften*, vol. 37, no. 20, pp. 464–476, 1954.

[38] E. von Holst, "Relations between the central nervous system and the peripheral organs," *British Journal of Animal Behaviour*, vol. 2, pp. 89–94, 1954.

[39] S.-J. Blakemore, D. M. Wolpert, and C. D. Frith, "Central cancellation of self-produced tickle sensation," *Nature Neuroscience*, vol. 1, pp. 635–640, 1998.

[40] P. Whittle, "Risk-sensitive linear/quadratic/gaussian control," *Advances in Applied Probability*, vol. 13, no. 4, pp. 764–777, 1981.

[41] J. Marschak, "Elements for a theory of teams," *Management Science*, vol. 1, no. 2, 1955.

[42] R. Radner, "Team decision problems," *The Annals of Mathematical Statistics*, vol. 33, no. 3, pp. 857–881, 1962.

[43] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656, 1948.

[44] D. M. MacKay and W. S. McCulloch, "The limiting information capacity of a neuronal link," *Bulletin of Mathematical Biophysics*, vol. 14, pp. 127–135, 1952.

[45] F. Rieke, D. Warland, R. de Ruyer van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code*. The MIT Press, 1997.

[46] F. Gabbiani, "Coding of time-varying signals in spike trains of linear and half-wave rectifying neurons," *Computation in Neural Systems*, vol. 7, pp. 61–85, 1996.

[47] F. Gabbiani and C. Koch, "Coding of time-varying signals in spike trains of integrate-and-fire neurons with random threshold," *Neural Computation*, vol. 8, no. 1, pp. 44–66, 1996.

[48] A. Borst and F. E. Theunissen, "Information theory and neural coding," *Nature Neuroscience*, vol. 2, no. 11, pp. 947–957, 1999.

[49] W. S. Wong and R. W. Brockett, "Systems with finite communication bandwidth constraints—part i: State estimation problems," *IEEE Transactions on Automatic Control*, vol. 42, no. 9, pp. 1294–1299, 1997.

[50] W. S. Wong and R. W. Brockett, "Systems with finite communication bandwith constraints—ii: Stabilization with limited information feedback," *IEEE Transactions on Automatic Control*, vol. 44, no. 5, pp. 1049–1053, 1999.

[51] G. N. Nair and R. J. Evans, "Stabilization with data-rate-limited feedback: Tightest attainable bounds," *Systems & Control Letters*, vol. 41, no. 1, pp. 49–56, 2000.

[52] N. Elia and S. K. Mitter, "Stabilization of linear systems with limited information," *IEEE Transactions on Automatic Control*, vol. 46, pp. 1384–1400, September 2001.

[53] S. Tatikonda and S. Mitter, "Control under communication constraints," *IEEE Transactions on Automatic Control*, vol. 49, no. 7, pp. 1056–1068, 2004.

[54] A. Sahai and S. Mitter, "The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link—part i: Scalar systems," *Transactions on Information Theory*, vol. 52, no. 8, 2006.

[55] V. Gupta, B. Hassibi, and R. M. Murray, "Optimal LQG control across packet-dropping links," *Systems & Control Letters*, vol. 56, no. 6, pp. 439–446, 2007.

[56] N. C. Martins and M. A. Dahleh, "Freedback control in the presence of noisy channels: "Bode-like" fundamental limitations of performance," *IEEE Transactions on Automatic Control*, vol. 53, August 2008.

[57] D. Nešić, A. R. Teel, and P. V. Kokotović, "Sufficient conditions for stabilization of sampled-data nonlinear systems via discrete-time approximations," *Systems & Control Letters*, vol. 38, no. 4–5, pp. 259–270, 1999.

[58] D. Nešić and A. R. Teel, "Input-output stability properties of networked control systems," *IEEE Transactions on Automatic Control*, vol. 49, no. 10, pp. 1650–1667, 2004.

[59] P. Tabuada, "Event-triggered real-time scheduling of stabilizing control tasks," *IEEE Transactions on Automatic Control*, vol. 52, no. 9, 2007.

[60] X. Wang and M. D. Lemmon, "Self-triggered feedback control systems with finite-gain $\mathcal{L}_2$ stability," *IEEE Transactions on Automatic Control*, vol. 54, no. 3, pp. 452–467, 2009.

[61] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 138–162, 2007.

[62] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, "Feedback control under data rate constraints: An overview," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 108–137, 2007.

[63] C. Eliasmith and C. H. Anderson, *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems.* The MIT Press, 2003.

[64] P. Dayan and L. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems.* The MIT Press, 2001.

[65] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory.* Macmillan, 1990.

[66] C. Koch, *Biophysics of Computation: Information Processing in Single Neurons.* Oxford University Press, 1999.

[67] R. Sarpeshkar, "Analog versus digital: Extrapolating from electronics to neurobiology," *Neural Computation*, vol. 10, no. 7, pp. 1601–1638, 1998.

[68] R. Sarpeshkar and M. O'Halloran, "Scalable hybrid computation with spikes," *Neural Computation*, vol. 14, pp. 2003–2038, 2002.

[69] S. Collins, A. Ruina, R. Tedrake, and M. Wisse, "Efficient bipedal robots based on passive-dynamic walkers," *Science*, vol. 307, pp. 1082–1085, 2005.

[70] J. Schmitt and P. Holmes, "Mechanical models for insect locomotion: Dynamics and stability in the horizontal plane I. theory," *Biological Cybernetics*, vol. 83, pp. 501–515, 2000.

[71] E. R. Westervelt, J. W. Grizzle, and D. E. Koditschek, "Hybrid zero dynamics of planar biped walkers," *IEEE Transactions on Automatic Control*, vol. 48, no. 1, pp. 42–56, 2003.

[72] A. D. Ames, A. Abate, and S. Sastry, "Sufficient conditions for the existence of Zeno behavior in nonlinear hybrid systems via constant approximations," in *IEEE Conference on Decision and Control*, 2007.

[73] A. Lamperski and A. D. Ames, "Lyapunov-like conditions for the existence of Zeno behavior in hybrid and Lagrangian hybrid systems," in *IEEE Conference on Decision and Control*, 2007.

[74] R. Goebel and A. R. Teel, "Zeno behavior in homogeneous hybrid systems," in *IEEE Conference on Decision and Control*, 2008.

[75] R. Goebel and A. R. Teel, "Lyapnuov characterization of Zeno behavior in hybrid systems," in *IEEE Conference on Decision and Control*, 2008.

[76] Y. Or and A. D. Ames, "Stability of Zeno equilibria in Lagrangian hybrid systems," in *IEEE Conference on Decision and Control*, 2008.

[77] K. H. Johansson, J. Lygeros, S. Sastry, and M. Egerstedt, "Simulation of Zeno hybrid automata," in *Proceedings of the 38th IEEE Conference on Decision and Control*, (Phoenix, AZ), 1999.

[78] S. Simic, K. H. Johansson, S. Sastry, and J. Lygeros, "Towards a geometric theory of hybrid systems," *Dyanmics of Discrete, Continuous, and Impulsive Systems, Series B*, vol. 12, no. 5–6, pp. 649–687, 2005.

[79] M. K. Camlibel and J. M. Schumacher, "On the Zeno behavior of linear complementarity systems," in *40th IEEE Conference on Decision and Control*, 2001.

[80] J. Shen and J.-S. Pang, "Linear complementarity systems: Zeno states," *SIAM Journal on Control and Optimization*, vol. 44, no. 3, pp. 1040–1066, 2005.

[81] M. Heymann, F. Lin, G. Meyer, and S. Resmerita, "Analysis of Zeno behaviors in a class of hybrid systems," *IEEE Transactions on Automatic Control*, vol. 50, no. 3, pp. 376–384, 2005.

[82] A. D. Ames, A. Abate, and S. Sastry, "Sufficient conditions for the existence of Zeno behavior," in *44th IEEE Conference on Decision and Control and European Control Conference*, (Seville, Spain), 2005.

[83] A. D. Ames, P. Tabuada, and S. Sastry, "On the stability of Zeno equilibria," in *Hybrid Systems: Computation and Control* (J. Hespanha and A. Tiwari, eds.), vol. 3927 of *Lecture Notes in Computer Science*, pp. 34–48, Springer-Verlag, 2006.

[84] Y. Wang, "Dynamic modeling and stability analysis of mechanical systems with time-varying topologies," *Journal of Mechanical Design*, vol. 115, pp. 808–816, 1993.

[85] A. Cabot and L. Paoli, "Asymptotics for some vibro-impact problems with a linear dissipation term," *Journal de Mathmatiques Pures et Appliqus*, vol. 87, no. 3, pp. 291–323, 2007.

[86] P. Ballard, "The dynamics of discrete mechanical systems with perfect unilateral constraints," *Archive for Rational Mechanics and Analysis*, vol. 154, pp. 199–274, 2000.

[87] D. E. Stewart, "Convergence of a time-stepping scheme for rigid-body dynamics and resolution of Painlevé's problem," *Archive for Rational Mechanics and Analysis*, vol. 145, pp. 215–260, 1998.

[88] T. McGeer, "Passive walking with knees," in *IEEE International Conference on Robotics and Automation*, 1990.

[89] J. Pratt and G. A. Pratt, "Exploiting natural dynamics in the control of a planar bipedal walking robot," in *Allerton Conference on Communications, Control and Computing*, 1998.

[90] R. Alur and D. L. Dill, "A theory of timed automata," *Theoretical Computer Science*, vol. 126, pp. 183–235, 1994.

[91] R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine, "The algorithmic analysis of hybrid systems," *Theoretical Computer Science*, vol. 138, no. 1, pp. 3–34, 1995.

[92] T. A. Henzinger, "The theory of hybrid automata," in *IEEE Symposium on Logic in Computer Science*, pp. 278–292, 1996.

[93] T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya, "What's decidable about hybrid automata," *Journal of Computer and System Sciences*, vol. 57, pp. 94–124, 1998.

[94] P. Collins, "A trajectory-space approach to hybrid systems," in *Proceedings of the 16th International Symposium on the Mathematical Theory of Networks & Systems*, 2004.

[95] R. Goebel and A. R. Teel, "Solutions to hybrid inclusions via set and graphical convergence with stability theory applications," *Automatica*, vol. 42, pp. 573–587, 2006.

[96] A. Lamperski and A. D. Ames, "On the existence of zeno behavior in hybrid systems with non-isolated Zeno equilibria," in *IEEE Conference on Decision and Control*, 2008.

[97] J. Lygeros, K. H. Johansson, S. Simic, J. Zhang, and S. Sastry, "Dynamical properties of hybrid automata," *IEEE Transactions on Automatic Control*, vol. 48, pp. 2–17, 2003.

[98] J. P. Aubin and H. Frankowska, *Set Valued Analysis*. Birkhäuser, 1990.

[99] A. Filippov, *Differential equations with discontinuous right-hand sides*. Kluwer Academic Publishers, 1988.

[100] B. Brogliato, *Nonsmooth Mechanics*. Springer-Verlag, 1999.

[101] A. D. Ames and S. Sastry, "Routhian reduction of hybrid Lagrangians and Lagrangian hybrid systems," in *American Control Conference*, 2006.

[102] Y. Or and A. R. Teel, "Zeno stability of the set-valued bouncing ball," *IEEE Transactions on Automatic Control*, vol. 56, pp. 447–452, 2011.