

Invariant Measures, Geometry, and Control of Hybrid and Nonholonomic Dynamical Systems

by
William Clark

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied and Interdisciplinary Mathematics)
in The University of Michigan
2020

Doctoral Committee:

Professor Anthony Bloch, Co-Chair
Professor Jessy Grizzle, Co-Chair
Professor Peter Miller
Associate Professor Shai Revzen
Professor Ralf Spatzier

William A. Clark

wiclark@umich.edu

ORCID iD: [0000-0002-0817-0309](https://orcid.org/0000-0002-0817-0309)

© William A. Clark 2020
All Rights Reserved

To Jade

ACKNOWLEDGEMENTS

I would like to first thank my two advisors, Anthony Bloch and Jessy Grizzle. Without your support and encouragement throughout the years, I would have never been able to get this far. Your help with growth and exploration will never be forgotten, thank you.

Next, I want to thank my two remarkable mentors during graduate school. To Leonardo Colombo: thank you for being there since my first day in Michigan, for introducing me to Anthony Bloch and hybrid systems, and in assisting me with my first publications and conferences. To Maani Ghaffari: thank you for introducing me to the field of computer vision and how wonderful it is to solve engineering problems with an overabundance of math.

I am grateful for MSRI and their summer schools which I was able to attend. These programs helped shape understanding of symplectic geometry, grow as a mathematician, and introduce me to many great peers. In particular, I would like to thank Thomas Melistas, Pedro Valentín De Jesús, Adam Yassine, Thomas Harris, and Jon Bush.

I would like to thank a number of my peers. I thank Matt Kvalheim, George Council, Michael Fisher, Ryan Sandberg, Zach Zwahlen, and Bruce Huang for many insightful conversations. Thank you to Marshall Scott, Jonathan Guzman, Anthony Della Pella, Jenia Rousseva, Jay Barraza, Jeremy Waters, and Jade Cain for making the best burgers group imaginable.

I would finally like to thank many of the fantastic professors from my experience at Michigan. Thank you Peter Miller, Alejandro Uribe, Ralf Spatzier, David Remy, and Gabor Orosz for your wonderful, thought-provoking classes and discussions. I also thank Anthony Bloch, Jessy Grizzle, Peter Miller, Shai Revzen, and Ralf Spatzier for serving on my thesis committee.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	viii
LIST OF TABLES	x
LIST OF SYMBOLS	xi
LIST OF APPENDICES	xii
ABSTRACT	xiii
CHAPTER	
I. Introduction	1
1.1 One-sided Constraints and Hybrid Systems	1
1.2 Nonholonomic Constraints	2
1.3 Structure of the Thesis	3
II. Hybrid Dynamical Systems	7
2.1 What is a Hybrid Dynamical System?	7
2.2 Solution Concept of a Hybrid Dynamical System	9
2.2.1 Hybrid Time Domains and Hybrid Arcs	9
2.2.2 Solutions to Hybrid Systems	11
2.2.3 Regularity of Solutions	12
2.3 General Hybrid Dynamical Systems	15
2.3.1 Directed Graphs	16
2.3.2 General Hybrid Systems	17
III. Geometric Mechanics and Nonholonomic Mechanics	20
3.1 Symplectic and Poisson Geometry	21
3.1.1 Symplectic Manifolds	21
3.1.2 Poisson Manifolds	22
3.2 Unconstrained Geometric Mechanics	24
3.2.1 Lagrangian Mechanics	24
3.2.2 Hamiltonian Mechanics	26
3.2.3 Symmetries in Hamiltonian Systems	27
3.3 Nonholonomic Mechanics	31
3.3.1 Constraint Distributions	31
3.3.2 Hamiltonian Nonholonomic Systems	32

3.3.3	Symmetries in Nonholonomic Systems	33
3.3.4	The Momentum Equation	35
3.4	Hamilton-Jacobi Theory	36
3.4.1	The Hamilton-Jacobi Equation	37
3.4.2	Lagrangian Submanifolds and Integrability	38
IV.	Stability and Asymptotic Behavior of Hybrid Systems	41
4.1	Limit Sets	41
4.2	Hybrid Poincaré Map and Floquet Theory	45
4.2.1	The Hybrid Poincaré Map	46
4.2.2	Floquet Theory	47
4.2.3	Example: Walking with Foot-slip	54
4.3	A Poincaré-Bendixson Theorem	61
4.3.1	Example: The rimless wheel	66
V.	Invariant Measures in Nonholonomic Systems	72
5.1	Global Nonholonomic Vector Fields	73
5.2	Symmetries in Nonholonomic Systems	76
5.3	Nonholonomic Volume Form	78
5.4	Divergence of a Nonholonomic System	79
5.4.1	Divergence Preliminaries	79
5.4.2	One Constraint	81
5.4.3	Multiple constraints	82
5.4.4	“Curvature”	82
5.5	Examples	83
5.5.1	Chaplygin Sleigh	83
5.5.2	Rolling Penny	85
5.6	Invariant Measures and the Cohomology Equation	87
5.6.1	The Cohomology Equation	87
5.6.2	Existence and Uniqueness	88
5.7	Special case: Measures Depending on Configuration	90
5.8	Invariant Measures Coming from a Momentum Map	95
5.8.1	The Chaplygin Sleigh Revisited	97
VI.	Mechanical Hybrid Systems	100
6.1	Mechanical Impact Maps	100
6.1.1	Holonomic Impacts	103
6.1.2	Nonholonomic Impacts	105
6.2	Regularity of Mechanical Hybrid Systems	111
6.3	Systems with Symmetry	112
6.3.1	Momentum Equation	112
6.3.2	The Mechanical Connection	113
6.3.3	Hybrid Reduction	114
6.3.4	Nonsymmetric Hybrid Impacts	115
6.4	Example: Nonholonomic Billiards - Rolling Disk	118
6.4.1	Dynamic nonholonomic equations	118
6.4.2	Impact map	119
6.4.3	The Momentum Equation	122
6.4.4	Numerical Results	124
VII.	Invariant Measures in Hybrid Systems	127

7.1	Invariant differential forms	127
7.1.1	A Formula for the Augmented Differential	129
7.1.2	Properties of Invariant Forms	130
7.1.3	Invariant Forms for Hybrid Mechanical Systems	132
7.2	Hybrid-Invariant Measures	133
7.3	Invariant Measures for Nonholonomic Hybrid Systems	135
7.3.1	The Hybrid Jacobian	135
7.4	The Zeno Issue in Measure Preserving Systems	141
7.5	Examples	143
7.5.1	Chaplygin Sleigh	144
7.5.2	Rolling Disk	145
VIII.	Hamilton-Jacobi Theory for Hybrid Systems	148
8.1	The Hybrid Hamilton-Jacobi Equation	148
8.1.1	Example: The Bouncing Ball	149
8.2	Integrability	152
8.2.1	Example: Harmonic Oscillator and Bouncing Ball	154
8.2.2	Example: Billiards on a Circular Table	156
IX.	Optimal Control of Mechanical Hybrid Systems	159
9.1	Hybrid Optimal Control Problem	160
9.2	Principle of Optimality	162
9.3	Dynamic Programming	164
9.3.1	Discretization of the Dynamics	164
9.3.2	Example: Bouncing Ball	166
9.4	The Strong Hybrid Hamilton-Jacobi-Bellman Equation	168
9.4.1	Inspiration of a Strong HJB	169
9.4.2	Impacts with Discontinuous Base	170
9.4.3	The Strong Hybrid HJB Equation	172
9.5	Maximum Principle	173
9.5.1	Optimal Control of a Mechanical System	174
9.6	Example: Bouncing Ball	175
9.6.1	Maximum Principle	176
9.6.2	The Hybrid Hamilton-Jacobi-Bellman Equation	177
9.6.3	Comparison with the Continuous Problem and Collision Loss	178
X.	Conclusions	182
10.1	Summary	182
10.2	Future Work	183
10.2.1	Hybrid Tongues	183
10.2.2	Measure Shaping Controls	184
10.2.3	Hybrid Integrators	186
10.2.4	Hybrid Bracket	186
	APPENDICES	188
	BIBLIOGRAPHY	197

LIST OF FIGURES

Figure

2.1	Graphical representation of a (simple) hybrid system. This can be viewed as a directed graph with one vertex and one edge.	16
4.1	The orbit of the periodic orbit for the system given by Theorem IV.20.	51
4.2	The vector $\delta x = F(y)\delta y + f(x)\delta t$, where the horizontal line is the tangent to S at the point x	52
4.3	The dowel-pin represents the center of mass for the pendulum.	54
4.4	The colored regions in the above figures show where a periodic orbit is found while the white is where a crash occurs. The left image shows ε small which is high friction while the right has large ε which corresponds to low friction.	59
4.5	The blue region corresponds to initial conditions that lie within the basin of attraction for the periodic orbit. The red region contains initial conditions that lead to a crash. The black asterisks indicate the forward orbit under the Poincaré map. The green line is a linear guess for the boundary of the red region.	60
4.6	The graph of a HDS where chaos can still occur.	63
4.7	An example of a node with outdegree 2. The set \mathcal{X}_i gets partitioned into two sets via the dashed line. Each of these now has outdegree 1.	66
4.8	Drawing of when both legs ℓ_i and ℓ_{i+1} are on the ground.	68
4.9	Both figures are for a 5 spoke wheel. Left: The plot of θ against time. Right: The plot of θ against time.	71
4.10	Both figures are for a 10 spoke wheel. Left: The plot of θ against time. Right: The plot of θ against time.	71
5.1	The Chaplygin sleigh is a rigid body attached to a sliding knife edge, §1.7 in [16].	84
5.2	Configuration of the rolling penny, §1.4 in [16].	86
6.1	The black curve is the continuous, piece-wise differentiable trajectory with an impact at t_1 while the red curve shows a virtual displacement.	103
6.2	The first 20 impacts for elastic/plastic impacts on a circular/elliptical table.	125
6.3	Plots of 100 different initial conditions at various times.	126
6.4	Left: The momentum for a circular table is 2-periodic which is representative of the order seen in Figure 6.2 (a). Center: The 2-periodicity is slightly destroyed with a slightly elliptical table. Right: All order in the momentum is lost which demonstrates the lack of order in Figure 6.2 (c).	126
7.1	Diagram for the proof of Theorem VII.1. The set $\tilde{S} := \Delta(S)$ and $T_{x_0}^S M := (\varphi_{-t_1})_* T_{y_0} S$	129
8.1	Left: The reachability bounds for the bouncing ball depending on the number of bounces. Right: Every trajectory that starts at $x = t = 0$ and ends at $x = 0.2$, $t = 5$ with no more than three bounces. The two solutions for each number of bounces corresponds to the two solutions to (8.3).	152
8.2	The plot of \mathcal{A}_k^\pm for k up to 5. The color-coding records the number of bounces: black means zero bounces, red (both light and dark) records a single bounce and so on.	153
8.3	Cross-sections of Figure 8.2 at various times. The corners on the left correspond to impacts while the cusps on the right are reachability bounds, corresponding to (8.4).	153

8.4	Left: The Lagrangian foliation for the harmonic oscillator. Right: The (multi-valued) action for the harmonic oscillator.	155
8.5	Left: The hybrid Lagrangian foliation for the bouncing ball. Right: The (multi-valued) action for the bouncing ball.	156
9.1	Optimal trajectories for the bouncing ball optimal control problem with various initial conditions x_0 with parameters $\alpha = 10$ and $\kappa = 0$. The effect of the noisy control is not visible in the above trajectories since the second-order nature of the dynamics mollifies the control input.	167
9.2	Optimal trajectories for the bouncing ball optimal control problem with various initial conditions x_0 with parameters $\alpha = 10$ and $\kappa = 1$	168
9.3	The black curve is the trajectory with a jump at t_1 while the red curve shows a virtual displacement.	171
9.4	The figures above compare the optimal trajectories computed via dynamic programming along with those found via the maximum principle. For this problem $m = 2$, $g = 1$, $t_0 = 0$, $t_f = 5$, and $\alpha = 10$	177
9.5	Left column: Plots of the function $J(x, p; t)$ at various time steps. Right column: Cross-sections of the solutions at $x = 0.5$	179
9.6	Left: A plot of $J(x, p; 0)$ as seen in Figure 9.5. Center: A plot of $J(x, p; 0)$ where the color corresponds to the number of impacts; zero impacts is dark blue, one impact is light blue, and two impacts is yellow. Right: A plot showing how many bounces is needed to obtain the minimum cost.	180
9.7	The optimal costs for the bouncing ball with various coefficients of restitution.	181
10.1	Left: Stability plot for the walker with foot-slip, cf. Figure 4.5. Right: Stability plot for Mathieu's equation $\ddot{x} + [\delta + \varepsilon \cdot \cos(t)]x = 0$, taken from §9.2 in [78].	184
10.2	Left: A modified version of the results from Figure 4.4 where stability is based off of the initial conditions $(\xi_0, \eta_0) = (-1, 0)$ rather than the adaptive initial conditions discussed in Appendix A. Right: A zoomed in version of the left figure.	185

LIST OF TABLES

Table

1.1	Chapter dependency chart. The dashed line between Chapters V and VI is due to the fact that Chapter VI only requires §5.1 and not the whole of Chapter V.	6
4.1	A random set of admissible parameters for the 5-spoked nonuniform rimless wheel with $\alpha = 0.5062$	70
4.2	A random set of admissible parameters for the 10-spoked nonuniform rimless wheel with $\alpha = 0.1752$	71
6.1	Four types of nonholonomic impact maps.	110
6.2	The chosen set of parameters for the time step example; a takes the first value for the elliptical cases and the latter for the circular ones.	124
9.1	The discretization used for the bouncing ball optimal control problem.	167

List of Symbols

$\mathcal{H} = (\mathcal{X}, S, f, \Delta)$	A hybrid dynamical system	7
$\varphi^{\mathcal{H}} : \mathbb{R}^+ \times \mathcal{X} \rightarrow \mathcal{X}$	The hybrid flow	12
$\mathcal{D} = (\mathcal{N}, \mathcal{E})$	A directed graph	16
$\mathcal{GH} = (\mathcal{N}, \mathcal{E}, \mathcal{X}, S, f, \Delta)$	A general hybrid dynamical system	17
TQ	Tangent bundle of Q	20
T^*Q	Cotangent bundle of Q	20
$\mathfrak{X}(Q)$	Vector fields over Q	20
$\Omega^k(Q)$	Differential k -forms on Q	20
$C^\infty(Q)$	Smooth functions $Q \rightarrow \mathbb{R}$	20
i_X	Interior Multiplication, $\Omega^k \rightarrow \Omega^{k-1}$	20
$[X, Y]$	Vector field bracket	20
\mathcal{L}_X	Lie derivative	20
$\Gamma(M)$	Sections of the bundle $M \rightarrow Q$	20
(M, ω)	A symplectic manifold	21
$(P, \{, \cdot\})$	A Poisson Manifold	23
(Q, g)	A Riemannian manifold	25
X_H	Hamiltonian vector field with Hamiltonian H	26
$J : M \rightarrow \mathfrak{g}^*$	A momentum map	28
$\mathcal{D} \subset TQ$	Constraint velocity distribution	31
$\eta^\alpha \in \Omega^1(Q)$	Constraint 1-forms	32
$\mathcal{D}^0 = \text{Ann}(\mathcal{D}) \subset T^*Q$	Annihilator of the constraint distribution $\mathcal{D} \subset TQ$	32
$W^\alpha \in \mathfrak{X}(Q)$	Dual constraint vector-fields	33
\mathcal{D}^*	Constraint momentum distribution	33
$X_H^{\mathcal{D}}$	Nonholonomic vector field with Hamiltonian H and constraint $\mathcal{D}^* \subset T^*Q$	33
$\Delta_*^f : T_x \mathcal{X} \rightarrow T_{\Delta(x)} \mathcal{X}$	The augmented differential	47
\mathcal{C}	Realization of the constraints	73
$\Xi_H^{\mathcal{C}}$	The global nonholonomic vector field from realization \mathcal{C}	73
$(m^{\alpha\beta})$	Constraint mass matrix with inverse $(m_{\alpha\beta})$	74
$\nu_H^{\mathcal{C}}$	The nonholonomic 1-form	76
$\mu^{\mathcal{C}}$	Nonholonomic volume form	78
$\text{div}_\Omega(X) \in C^\infty(M)$	Divergence of X with respect to the volume form Ω	79
$\vartheta_{\mathcal{C}}$	The nonholonomic density form	90
S^*	Hamiltonian impact surface	102
$S_{\mathcal{D}}^* := S^* \cap \mathcal{D}^*$	Nonholonomic Hamiltonian impact surface	102
$\mathbb{I} : \mathfrak{g} \rightarrow \mathfrak{g}^*$	The locked inertia tensor	116
$\mathcal{A}_s : M \rightarrow \mathfrak{g}$	The mechanical connection	116
$\langle\langle \eta, \mu \rangle\rangle_q = \eta(\mathbb{I}(q)^{-1}\mu)$	Mechanical metric on \mathfrak{g}^*	117
$\mathcal{A}_{\mathcal{H}} \subset \Omega(M)$	Hybrid-invariant differential forms	129
$\mathcal{J}_\mu(\Delta) \in C^\infty(S)$	The hybrid Jacobian of Δ	134
$C_{PW}^\infty(X, Y)$	Piecewise smooth functions $X \rightarrow Y$	160

LIST OF APPENDICES

Appendix

- A. Numerical Methods for the Foot-slip Model 189
 - A.1 Stable Parameters 190
 - A.2 Stable Initial Conditions 191
- B. Technical Lemmas from Continuous and Discrete Dynamics 192

ABSTRACT

Constraints are ubiquitous when studying mechanical systems and fall into two main categories: hybrid (1-sided, unilateral) and nonholonomic/holonomic (2-sided, bilateral) constraints. A hybrid constraint takes the form $h(x) \geq 0$. An example of a constraint of this nature is requiring a billiard ball to remain within the confines of a table-top. The notable feature of these constraints is that when the ball reaches the boundary of the table-top (i.e. when $h(x) = 0$), an *impact* occurs; this is a discontinuous jump in the dynamics. Dynamical systems that have this phenomenon generally fall under the domain of *hybrid dynamical systems*. On the other hand, nonholonomic constraints take the form $h(x) = 0$. Generally, h will depend on *both* the positions and velocities and cannot be integrated to only depend on the positions (when it can be integrated, the constraint is called holonomic). An example of a nonholonomic constraint is an ice skate: motion is not allowed perpendicular to the direction of the skate. It is common that these systems are studied using tools from differential geometry.

This thesis studies both hybrid and nonholonomic constraints together using the language of differential (specifically symplectic) geometry. However, due to the exotic nature of hybrid dynamics, some auxiliary results are found that pertain to the asymptotic nature of these systems. These include the idea of a hybrid limit-set, Floquet theory, and a Poincaré-Bendixson theorem for planar systems.

The bulk of this work focuses on finding (smooth) invariant measures for both

nonholonomic and hybrid systems (as well as systems involving *both* types of constraints). Necessary and sufficient conditions are found which guarantee the existence of an invariant measure for nonholonomic systems in which the density depends only on the configuration variables. Extending this idea to hybrid nonholonomic systems requires that the impact preserves the measure *as well*. To build towards this, relatively simple conditions to test whether or not a differential form is hybrid-invariant are derived. In the cases where the density depends on only the configuration variables, the measure is still invariant under the hybrid dynamics *independent of the choice of impacts*. The billiard problem with a vertical rolling disk as the billiard ball is one such system and is therefore recurrent for any choice of compact table-top.

This thesis concludes with optimal control of hybrid systems. First, Hamilton-Jacobi is extended to the hybrid setting (nonholonomic constraints are not considered here) and the idea of completely integrable hybrid systems is introduced. It is shown that the usual billiard problem on a circular table is completely integrable. Finally, the hybrid Hamilton-Jacobi theory is extended to a hybrid Hamilton-Jacobi-Bellman theory which allows for the study of optimal control problems.

CHAPTER I

Introduction

Constraints in mechanical systems fall under two broad categories: hybrid and nonholonomic. Let $h : M \rightarrow \mathbb{R}$ be a smooth function in which M is the manifold of all states; we can use this to impose a constraint in two different ways:

1. $h(x) \geq 0$, or
2. $h(x) = 0$.

In the words of Niemark and Fufaev [83], constraints of the form $h \geq 0$ are referred to as one-sided, nonrestrictive, or nonlimiting; and constraints of the form $h = 0$ are referred to as two-sided, restrictive, or limiting. In the language of this thesis, one-sided constraints will be referred to as *hybrid* while two-sided constraints will be called *nonholonomic*. We will not make a distinction between holonomic and nonholonomic constraints, as the former is a special case of the latter (both can be treated as nonholonomic without any issues).

1.1 One-sided Constraints and Hybrid Systems

When a one-sided constraint is imposed, $h(x) \geq 0$, the impact (when $h(x) = 0$) induces a discontinuity in the dynamics. The impact manifests as a discrete transition and we obtain the following:

- Continuous evolution when $h(x) > 0$, and
- Discrete evolution when $h(x) = 0$.

Dynamical systems which exhibit *both* continuous and discrete evolution are called *hybrid dynamical systems* (or hybrid systems for short).

It is important to note that the field of hybrid dynamical systems encompasses much more than systems with unilateral constraints. Hybrid dynamical systems are used to describe a wide variety of disciplines so a comprehensive list is nearly impossible to present. Some examples include: billiards [7, 25, 44, 97], bipedal robots [30, 54, 82, 96], biological systems [2, 21, 67, 84, 103], hybrid automata [3, 59, 106, 109, 111], and switched systems [35, 36, 51, 60].

Due to the extensive applicability of hybrid systems, there exist many different and nonequivalent formulations for these systems. Chapter II will outline the formulation used in this thesis; the choice for this version of hybrid systems is due to the motivating problem of understanding mechanical impacts. The definition of a hybrid dynamical system in Chapter II will be more general than needed, as it will be useful in Chapter IV, in which asymptotic properties of hybrid dynamical systems are studied.

1.2 Nonholonomic Constraints

Contrary to the previous section where $h(x) \geq 0$, suppose that the constraint is of the form $h(x) = 0$. For a mechanical (Lagrangian) system, the system evolves on a tangent bundle $M = TQ$, in which Q is called the configuration manifold. Choosing coordinates $q \in Q$ results in induced coordinates $(q, \dot{q}) \in TQ$. If we have a constraint depending *only* on the base variables, i.e. $h(q) = 0$, then the constraint is called holonomic. Differentiating this constraint induces a constraint involving velocities

as well: $dh_q(\dot{q}) = 0$. Therefore, holonomic constraints induce a velocity constraint given by the kernel of an exact 1-form. If, on the other hand, we have a constraint given by the kernel of a non-exact 1-form, the constraint is called nonintegrable or nonholonomic.

Nonholonomic constraints as explained above are inherently linear in the velocities. This restriction can be generalized to nonlinear constraints, cf. e.g. [37], but most physical examples will fall under the linear category. As such, the majority of the results in this work will assume that the constraints are linear in the velocity (although many results can be extended to the nonlinear case via Chetaev's rule).

1.3 Structure of the Thesis

There are three parts to this thesis: foundations, asymptotics and measure, and controls. A layout of the chapters into the corresponding parts is shown in Table 1.1.

As this thesis combines the theory of hybrid systems with geometric mechanics, there are two preliminary chapters. Chapter II outlines the version of hybrid systems studied here as well as the solution concept and basic regularity assumptions. Chapter III introduces the concepts from geometric mechanics that will be used throughout; in particular, symplectic geometry, Hamiltonian mechanics, symmetries, nonholonomic constraints, and Hamilton-Jacobi theory.

The second part of this thesis focuses on understanding some asymptotic properties of hybrid dynamical systems as well as trying to find invariant measures for both nonholonomic and hybrid systems. Chapter IV contains fundamental concepts on the stability and asymptotic behavior of hybrid systems. Specifically, this chapter defines and proves some properties of the hybrid limit set, introduces the augmented impact differential and uses this to determine the stability of hybrid periodic orbits in

a fashion similar to Floquet theory, and proves a version of the Poincaré-Bendixson theorem for hybrid systems. Chapter V examines the problem of the existence of smooth invariant measures for nonholonomic systems and presents necessary and sufficient conditions for an invariant density to exist that depends only on the base variables. It is shown that the rolling disk does preserve measure while the Chaplygin sleigh preserves no measures (with density depending only on the base variables). Chapter VI combines the principles of Chapter II and Chapter III to study mechanical hybrid systems. In particular, the billiard problem in which the billiard ball is a vertical rolling disk is presented along with a hybrid analogue of Noether's theorem [85] (or the momentum equation for nonholonomic systems [14]). The final chapter of this part, Chapter VII, extends results from Chapter V to the systems studied in Chapter VI. Necessary and sufficient conditions are presented for a differential form to be invariant under a hybrid flow; elastic hybrid holonomic systems do preserve the canonical symplectic form, in agreement with [91]. Finally, this chapter shows that if a nonholonomic system preserves a measure with density only depending on the base variables (cf. Chapter V), then the corresponding hybrid flow preserves the same measure, *independent of the choice of impact*. This result specializes to Corollary VII.22 which states that the nonholonomic billiard problem with the vertical rolling disk is recurrent for *any* choice of compact table-top.

The final part of this thesis explores Hamilton-Jacobi theory and its extension to Hamilton-Jacobi-Bellman theory for the purposes of optimal control of hybrid dynamical systems. Chapter VIII develops a preliminary Hamilton-Jacobi theory for hybrid systems; this is particularly for systems *without* nonholonomic constraints. We develop this theory with so-called "hybrid Lagrangian foliations" to avoid issues with a hybrid bracket. Using this context, we show that the elastic bouncing ball,

as well as, the circular billiard problem are both completely integrable. This chapter ends with a brief digression into issues with defining a “hybrid Poisson bracket.” Chapter IX develops some optimal control theory for hybrid systems (which *does* include systems subject to both types of constraints, in contrast with the previous chapter). The principle of optimality, dynamic programming, the Hamilton-Jacobi-Bellman equation, and Pontryagin’s maximum principle are discussed in the hybrid framework. This chapter concludes with a numerical example of the bouncing ball; in particular, when the controls are constrained to be weaker than gravity, optimal controls are able to utilize the impacts to gain height (a phenomenon which is impossible for the corresponding continuous system).

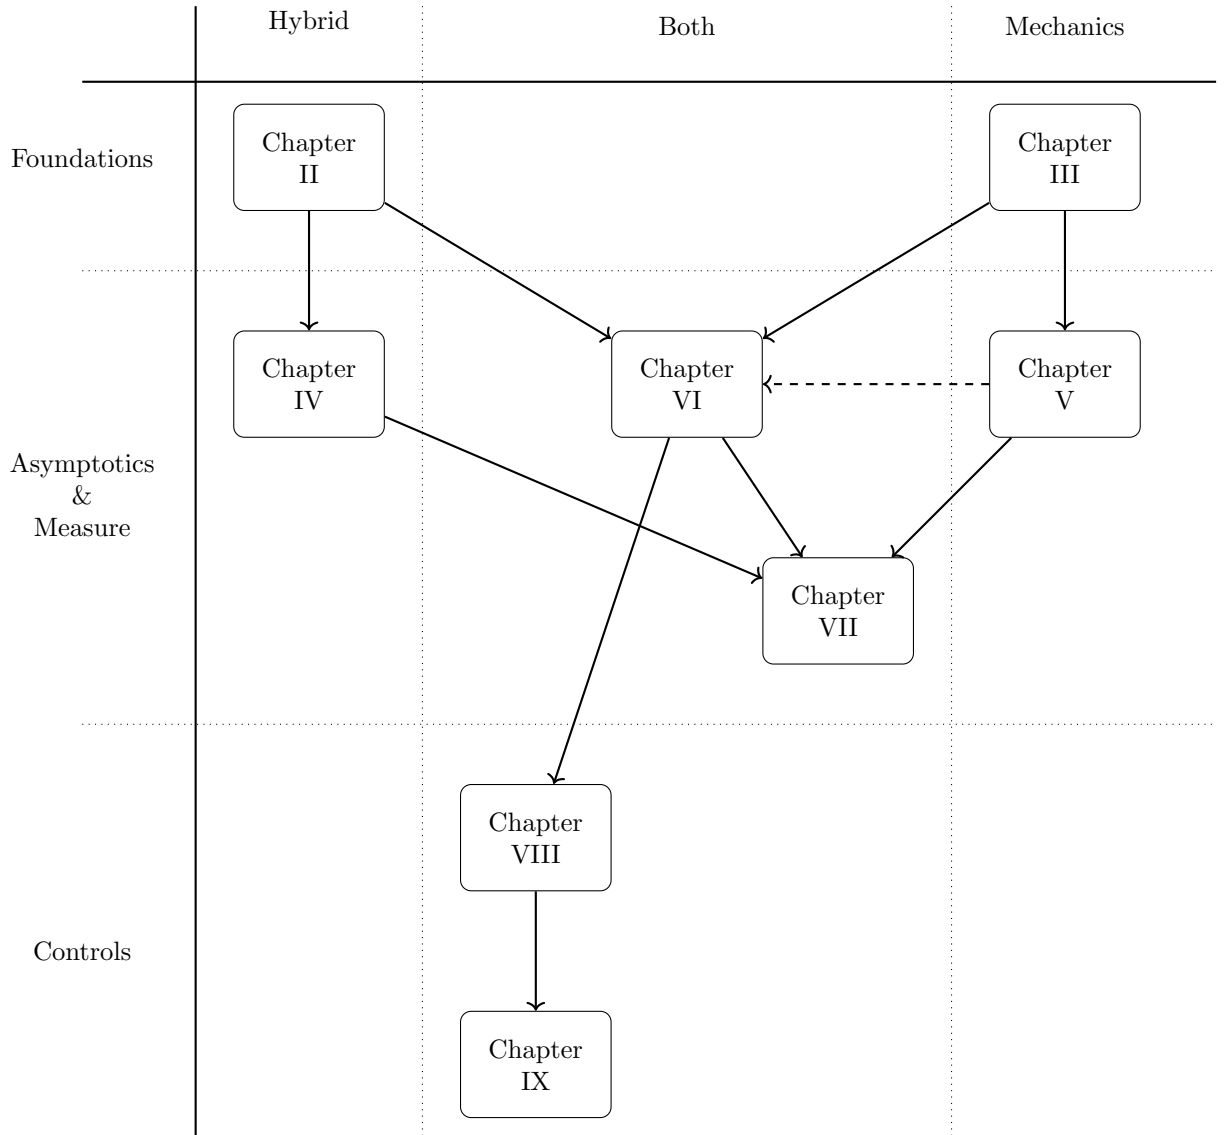


Table 1.1: Chapter dependency chart. The dashed line between Chapters V and VI is due to the fact that Chapter VI only requires §5.1 and not the whole of Chapter V.

CHAPTER II

Hybrid Dynamical Systems

2.1 What is a Hybrid Dynamical System?

A hybrid dynamical system is a dynamical system that experiences both continuous and discrete transitions. There exist many different, nonequivalent, ways to formalize this idea. However, as this thesis is concerned with modeling impact mechanics as hybrid systems, we will use the following definition for a hybrid system which depends on four pieces of data. Throughout, smooth will mean C^∞ .

Definition II.1. A hybrid dynamical system (HDS) is a 4-tuple $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ such that

(H.1) \mathcal{X} is a smooth (finite-dimensional) manifold,

(H.2) $S \subset \mathcal{X}$ is a smooth embedded submanifold with co-dimension 1,

(H.3) $f : \mathcal{X} \rightarrow T\mathcal{X}$ is a smooth vector field, and

(H.4) $\Delta : S \rightarrow \mathcal{X}$ is a smooth map.

Remark II.2. Using the symbols $(\mathcal{X}, S, f, \Delta)$ for an HDS is common notation in the field, [54, 81, 96]. As such, in Chapters II and IV a hybrid system will be referred to as $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$. However, to use notation more similar to differential geometry/geometric mechanics an HDS will usually be referred to as $\mathcal{H} = (M, S, X, \Delta)$ in the remaining chapters.

The manifold \mathcal{X} is called the state-space, S the impact surface, f the continuous dynamics, and Δ the impact map, discrete dynamics, or the reset map. The hybrid dynamics can be informally described as

$$(2.1) \quad \begin{cases} \dot{x} = f(x), & x \notin S, \\ x^+ = \Delta(x^-), & x^- \in S. \end{cases}$$

A hybrid dynamical system obeying only the conditions (H.1)-(H.4) is still capable of exhibiting troublesome behavior. If $S \cap \Delta(S) \neq \emptyset$, then multiple impacts can happen instantaneously, which is called *beating*. In order to discourage this phenomenon, we propose the following additional axiom.

Definition II.3. A hybrid dynamical system, $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$, is said to be *essentially non-beating* if, in addition to (H.1)-(H.4), it also satisfies

$$(H.5) \quad S \cap \Delta(S) = \emptyset \text{ and } \overline{\Delta(S)} \cap \overline{S} \subset \mathcal{X} \text{ has co-dimension of at least 2.}$$

Example II.4 (The bouncing ball). Consider the case of the bouncing ball where $(x, \dot{x}) \in \mathbb{R}^2 = \mathcal{X}$ with impact surface $S = \{(0, \dot{x}) : \dot{x} < 0\}$. Assuming that the impacts are elastic, the impact map is $\Delta(x, \dot{x}) = (x, -\dot{x})$ and $\Delta(S) = \{(0, \dot{x}) : \dot{x} > 0\}$. We see that $\overline{\Delta(S)} \cap \overline{S} = \{(0, 0)\}$ which has co-dimension 2 in \mathbb{R}^2 . For a general mechanical system with impacts, the set $\overline{\Delta(S)} \cap \overline{S}$ corresponds to “grazing” impacts, cf. e.g. §7.2 in [17]. The intention of this thesis is not to study this phenomenon and (H.5) justifies this avoidance by ensuring that it does not happen too often.

Remark II.5. Throughout this thesis, everything will be assumed smooth unless otherwise stated. Generally speaking, the continuous vector-field need only be Lipschitz and the impact surface be C^1 in order to guarantee basic regularity properties (like existence and uniqueness). Therefore, many of the smoothness assumptions can be relaxed.

It is important to note that the version of hybrid systems used here is similar to those studied in [17, 56, 82]. This is contrary to the differential inclusion formulation described by [52]; a hybrid dynamical system of this type is given by

$$\begin{cases} \dot{x} \in F(x), & x \in C, \\ x^+ \in G(x^-), & x^- \in D, \end{cases}$$

where F and G are set-valued maps. A reason for the added level of generality is Painlevé's paradox [104]. However, this thesis does not deal in situations where these paradoxes appear and as such will not use the differential inclusion approach.

2.2 Solution Concept of a Hybrid Dynamical System

Defining a solution to a hybrid system requires more care than for usual continuous or discrete dynamics. Part of this is due to three troublesome phenomena that can occur in hybrid dynamical systems: Zeno, beating, and confluence.

Most of the literature on hybrid dynamical systems comes from the controls community where arcs are defined to be absolutely continuous rather than smooth. Therefore a solution is only required to be locally absolutely continuous, but in practice it will be smooth.

2.2.1 Hybrid Time Domains and Hybrid Arcs

This subsection begins the treatment of a rigorous definition for a solution to a hybrid dynamical system. The fundamental objects required for this definition are hybrid time domains and hybrid arcs. Throughout, let $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ be a hybrid dynamical system.

Definition II.6 (Hybrid time domains, [52]). A subset $E \subset \mathbb{R}^+ \times \mathbb{N}$ is a compact

hybrid time domain if

$$E = \bigcup_{j=0}^{J-1} ([t_j, t_{j+1}], j).$$

for some finite sequence of times $0 = t_0 \leq t_1 \leq \dots \leq t_J$. A set $E \subset \mathbb{R}^+ \times \mathbb{N}$ is a hybrid time domain if for all $(T, J) \in E$, $E \cap ([0, T] \times \{0, \dots, J\})$ is a compact hybrid domain.

Hybrid time domains will be the domain for a hybrid trajectory in the same way that \mathbb{R} is the domain for a continuous system and \mathbb{Z} is for a discrete system. Solutions to hybrid dynamical systems are given by hybrid arcs.

Definition II.7 (Hybrid Arc, [52]). A function $\varphi : E \rightarrow \mathcal{X}$ is a hybrid arc if E is a hybrid time domain and if for each $j \in \mathbb{N}$, the function $t \mapsto \varphi(t, j)$ is locally absolutely continuous on the interval $I^j = \{t : (t, j) \in E\}$.

Before a solution of a hybrid dynamical system is defined, four important types of arcs will be introduced.

Definition II.8 (Types of hybrid arcs, [52, 56]). A hybrid arc, $\varphi : E \rightarrow \mathcal{X}$, is called

- nontrivial if E contains at least two points,
- complete if E is unbounded, i.e. $\sup_t E + \sup_j E = \infty$ where

$$\sup_t E = \sup \{t \in \mathbb{R}^+ : \exists j \in \mathbb{N} \text{ s.t. } (t, j) \in E\},$$

$$\sup_j E = \sup \{j \in \mathbb{N} : \exists t \in \mathbb{R}^+ \text{ s.t. } (t, j) \in E\},$$

- Zeno if φ is complete and $\sup_t E < \infty$,
- beating if there exists an interval, I^j , with empty interior.

Therefore, roughly speaking, a hybrid arc is Zeno if it experiences an infinite number of impacts in a finite amount of time and beating if multiple impacts occur in zero time. We are now able to define a solution to a hybrid dynamical system.

2.2.2 Solutions to Hybrid Systems

Definition II.9 (Solution to a hybrid dynamical system). A hybrid arc, $\varphi : E \rightarrow \mathcal{X}$, is a solution to the hybrid system $(\mathcal{X}, S, f, \Delta)$ if

1. for all $j \in \mathbb{N}$ such that $I^j := \{t : (t, j) \in E\}$ has nonempty interior,

$$\frac{d}{dt}\varphi(t, j) = f(\varphi(t, j)), \quad t \in \text{int}I^j,$$

2. for all $(t, j) \in E$ such that $(t, j + 1) \in E$,

$$\varphi(t, j) \in S, \quad \varphi(t, j + 1) = \Delta(\varphi(t, j)).$$

A hybrid time domain records both the time and the number of impacts a trajectory experiences. This work will not be interested in recording the number of impacts and therefore the following definition of a forgetful hybrid arc will be used throughout the remainder instead.

Definition II.10. For a hybrid arc, $\varphi : E \rightarrow \mathcal{X}$, its corresponding forgetful hybrid arc is a left-continuous map $\tilde{\varphi} : \mathcal{I} \rightarrow \mathcal{X}$ where

$$\mathcal{I} := [0, \sup_t E] \setminus \{t_1, \dots, t_{J-1}\}, \quad \tilde{\varphi}(t) = \varphi(t, j),$$

where the hybrid arc, φ , is assumed to not be beating.

The forgetful hybrid arc corresponding to a solution to $(\mathcal{X}, S, f, \Delta)$ has the following properties:

$$\begin{aligned} \frac{d}{dt}\tilde{\varphi}(t) &= f(\tilde{\varphi}(t)), \\ \lim_{\varepsilon \rightarrow 0^+} \tilde{\varphi}(t_k + \varepsilon) &= \Delta \left(\lim_{\varepsilon \rightarrow 0^-} \tilde{\varphi}(t_k + \varepsilon) \right), \end{aligned}$$

which “match” (2.1).

2.2.3 Regularity of Solutions

This section is concluded with a discussion on the regularity of solutions to hybrid dynamical systems. To ensure the impact times are well-posed, we make the following two assumptions [56]:

(A.1) If $\tilde{\varphi}(t) \in \overline{S} \setminus S$, then there exists $\varepsilon > 0$ such that for all $0 < \delta < \varepsilon$, $\tilde{\varphi}(t + \delta) \notin S$.

(A.2) If $\tilde{\varphi}(t) \in S$, then there exists $\varepsilon > 0$ such that for all $0 < \delta < \varepsilon$, $\tilde{\varphi}(t + \delta) \notin S$.

The first assumption, (A.1), prohibits the trajectory from entering S through \overline{S} while the second assumption, (A.2), requires that when a trajectory intersects S , it must instantaneously exit S . In terms of mechanical systems, (A.1) prohibits grazing while (A.2) states that impacts must move away from the wall. It can be seen that the condition (H.5) is weaker than (A.2) which is slightly weaker than requiring $\tilde{\varphi}$ to intersect S transversely. As will be seen in Theorem II.15, the conditions (A.1) and (H.5) imply (A.2).

The final regularity assumption we will examine in this section deals with continuous dependence of initial conditions. Due to the fact that the impacts are not continuous, true continuous dependence is impossible but quasi-continuous dependence is still achievable. Under the assumptions (A.1) and (A.2), there exists a unique forgetful hybrid arc corresponding to each initial condition $x_0 \in \mathcal{X}$: this association creates the *hybrid flow*.

Definition II.11. Assuming existence and uniqueness of hybrid solutions, the hybrid flow of a hybrid dynamical system, \mathcal{H} , is a map $\varphi^{\mathcal{H}} : \mathbb{R}^+ \times \mathcal{X} \rightarrow \mathcal{X}$, where $\varphi^{\mathcal{H}}(\cdot; x_0) : \mathbb{R}^+ \rightarrow \mathcal{X}$ is a forgetful hybrid arc with $\varphi^{\mathcal{H}}(0; x_0) = x_0$. Moreover, let $\mathcal{I}_{x_0} := \mathbb{R}^+ \setminus \{t_1, t_2, \dots\}$ where t_k are the impact times for $\varphi^{\mathcal{H}}(\cdot; x_0)$.

The hybrid flow allows for the statement of the third troublesome phenomenon of

hybrid dynamical systems: confluence.

Definition II.12. A hybrid dynamical system $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ has confluence if for some $t > 0$, the hybrid flow $\varphi_t^{\mathcal{H}} : \mathcal{X} \rightarrow \mathcal{X}$ fails to be injective.

Confluence results in time-irreversibility of the dynamics. Which, although undesirable, does not pose as many issues as either beating or Zeno solutions.

Due to the nature of hybrid systems, solutions will not be continuous in time (nor on choice of initial conditions). Rather than strict continuity, many hybrid systems will be *quasi-continuous*. To present this definition, we choose a distance metric, d , for \mathcal{X} .

Assumption II.13 (Quasi-continuous dependence property [41, 56]). *Consider the hybrid dynamical system $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ and let $\varphi^{\mathcal{H}}(t; x_0)$ be the hybrid flow and let $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a distance function. Then, for every $x_0 \in \mathcal{X} \setminus S$ and $\varepsilon > 0$ and $t \in \mathcal{I}_{x_0}$, there exists $\delta > 0$ such that if $d(x_0, y) < \delta$, then $d(\varphi^{\mathcal{H}}(t; x_0), \varphi^{\mathcal{H}}(t; y)) < \varepsilon$.*

Smoothness properties will be more important than continuous properties in this work. Therefore, we can strengthen Assumption II.13 to a quasi-smooth dependence property.

Assumption II.14 (Quasi-smooth dependence property). *Consider the hybrid system $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ and let $\varphi^{\mathcal{H}}(t; x_0)$ be the hybrid flow. Then for every $x_0 \in \mathcal{X} \setminus S$ and $t \in \mathcal{I}_{x_0}$, there exists an open neighborhood $x_0 \in U$ such that $U \cap S = \emptyset$ and the map $\varphi^{\mathcal{H}}(t; \cdot) : U \rightarrow \mathcal{X}$ is smooth.*

Theorem II.15. *Suppose a hybrid system $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ satisfies (H.1)-(H.5) and (A.1). If for all $x \in S$, we have $T_x \mathcal{X} = T_x S \oplus f(x) \mathbb{R}$, then \mathcal{H} also satisfies (A.2) and Assumption II.14.*

Proof. We need to show that \mathcal{H} satisfies (A.2) and Assumption II.14. Condition (A.2) follows from (H.5) from the following reason: Suppose $\tilde{\varphi}(t_k) \in S$, then

$$\tilde{\varphi}^+ := \lim_{\varepsilon \rightarrow 0^+} \tilde{\varphi}(t_k + \varepsilon) \notin S.$$

If $\varphi^+ \notin \overline{S}$, then we are done because $\mathcal{X} \setminus \overline{S}$ is an open set. If $\varphi^+ \in \overline{S} \setminus S$, then (A.1) provides such an ε .

It remains to show the quasi-smooth dependence property. Choose a $t \in \mathcal{I}_{x_0}$. If there are no impacts in the interval $[0, t]$ then the quasi-smooth dependence property is trivially satisfied. Without loss of generality, assume that there exists a single impact occurring at $t_1 \in (0, t)$. To show smoothness, we will compute the differential of $\varphi^{\mathcal{H}}(t; \cdot)$.

Choose an open set $x_0 \in U$ such that $U \cap S = \emptyset$ and let $\tau : U \rightarrow \mathbb{R}$ be the time until impact, e.g. $\tau(x_0) = t_1$. Then for $x \in U$, we can write the hybrid flow as

$$\varphi^{\mathcal{H}}(t; x) = \varphi(t - \tau(x), \Delta(\varphi(\tau(x); x))),$$

i.e. it can be written in terms of the continuous flow φ , the impact map Δ , and the return-time τ . The continuous flow (see e.g. §2.3 in [92]) and the impact map (by (H.4)) are smooth. Therefore, the result follows if τ is smooth. This is true by the following theorem, which is an application of the implicit function theorem. \square

Theorem II.16 ([29]). *Let $x_0 \in \mathcal{X} \setminus S$ be such that there exists a time, $T_0 > 0$, where $\varphi_{T_0}(x_0) \in S$. Additionally, assume that the flow intersects the impact surface transversely at $\varphi_{T_0}(x_0)$. Then there exists an $\varepsilon > 0$ and a smooth function $\tau : \mathcal{B}_\varepsilon \cap S \rightarrow \mathbb{R}^+$ such that for all $y \in \mathcal{B}_\varepsilon(x_0)$, $\varphi_{\tau(y)}(y) \in S$.*

Proof. Let $h : \mathcal{X} \rightarrow \mathbb{R}$ be such that zero is a regular value and $S = h^{-1}(0)$ (at least locally) and define the function $F : (0, +\infty) \times S \rightarrow \mathcal{X}$ by $F(t, x) = h(\varphi_t(x))$. It follows

from Theorem 1 in Section 2.5 in [92] that $(t, x) \mapsto \varphi_t(x)$ is $C^\infty(\mathbb{R} \times \mathcal{X})$. Combining this with the fact that h is a smooth function, we get that their composition is. Since $F \in C^\infty(\mathbb{R}^+ \times S)$, we can use the implicit function theorem. At our point $x_0 \in \mathcal{X}$, we know that the orbit enters the set S at some minimal future time, T_0 . This gives $F(T_0, x_0) = 0$. Differentiating F with respect to time yields:

$$\frac{\partial F}{\partial t}(T_0, x_0) = \frac{\partial h}{\partial y} \Big|_{y=\varphi_{T_0}(x_0) \in S} \cdot f(\varphi_{T_0}(x_0)) \neq 0.$$

The first factor is nonzero because zero is a regular value and the second is nonzero because we are away from a fixed point (if $f = 0$, it would not be transverse). Their inner product is nonzero because of the transversality condition. This lets us use the implicit function theorem (cf., e.g., Theorem 9.28 in [98] or Theorem B.2 in [107]), to show that there exists a neighborhood of $\mathcal{B}_\varepsilon(x_0)$ of x_0 and a smooth function τ with all the desired properties. \square

Notice that in the proof of Theorem II.15, (H.5) and (A.1) implies (A.2) while the quasi-smooth dependence property follows from impacts being transverse.

Definition II.17. A hybrid system, $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$, is called *smooth* if it satisfies the hypothesis of Theorem II.15.

2.3 General Hybrid Dynamical Systems

In what has been presented above, hybrid dynamical systems evolve according to a single manifold and a single impact. However, there exist cases where multiple different impacts may occur, e.g. multi-legged walking. This leads to the notion of a general hybrid dynamical system where there exist multiple different manifolds each containing continuous dynamics with impacts that move between these manifolds. This more general framework leads naturally to the theory of directed graphs where

each node is an individual state-space and each edge is an impact. This notion of a hybrid system is used in, e.g. [4, 74, 102]. General hybrid dynamical systems will appear in this thesis only in §4.3 where sufficient conditions are laid out to guarantee periodic asymptotic behavior as in [28]. The following exposition for general hybrid dynamical systems is similar to that as presented in [28].

2.3.1 Directed Graphs

Before the concept of a general hybrid system is introduced, hybrid systems will first be viewed through a different lens. An HDS consists of a single state-space and a single impact that maps back to the original state-space. This allows the HDS to be thought of as a directed graph (digraph) of order one with a single loop [23], see Fig 2.1.




Figure 2.1: Graphical representation of a (simple) hybrid system. This can be viewed as a directed graph with one vertex and one edge.

With this view of a hybrid system, it seems natural to extend the definition of a HDS to an arbitrary digraph. Before making this definition, a few important aspects of graph theory are reviewed.

Definition II.18. A directed graph, or digraph, $\mathcal{D} = (\mathcal{N}, \mathcal{E})$ consists of two components: a finite, nonempty set called the vertex set as well as a multi-set (a modified version of a set where multiple instances of an element are allowed) of ordered pairs of (not necessarily distinct) vertices called the (directed) edge set. The vertex set

if denoted by $V(\mathcal{D}) = \mathcal{N}$ and the directed edge set is denoted by $E(\mathcal{D}) = \mathcal{E}$. The order of a digraph is the cardinality of the vertex set.

This definition of digraphs allows for multiple edges to connect two vertices, u and v . This differs from the definition in [23] where this is not permitted. Because a directed edge may not be uniquely determined by an ordered pair of vertices, an element of \mathcal{E} will be denoted by (u, v) if only one edge connects u to v . If $e \in \mathcal{E}$ connects the vertex u to v , we will call u to be the *source* of e and v the *target* of e . This will be denoted by $u = \text{dom}(e)$ and $v = \text{cod}(e)$.

There are two remaining objects to define for digraphs: degree and cycles.

Definition II.19. Let $u \in \mathcal{N}$. If $e \in \mathcal{E}$ has $\text{dom}(e) = u$, e is said to be sourced from u . Likewise, if $\text{cod}(e) = u$, e is said to be targeted to u . The outdegree, $\text{od}(u)$, of u is the number of edges sourced from u . Additionally, the indegree, $\text{id}(u)$, of u is the number of edges targeted to u . The degree, $\text{deg}(u)$, of u is defined by

$$\text{deg}(u) := \text{od}(u) + \text{id}(u).$$

Definition II.20. A digraph of order n is a cycle if the vertices can be labeled as $i = 1, \dots, n$ and the n directed edges can be written as $(i, i + 1)$ (taken modulo n).

2.3.2 General Hybrid Systems

Definition II.1 will now be generalized to incorporate the directed graph structure as explained above, which now depends on six pieces of data.

Definition II.21. The six-tuple, $\mathcal{GH} = (\mathcal{N}, \mathcal{E}, \mathcal{X}, S, f, \Delta)$, is a general hybrid dynamical system (GHDS) if

(G.1) The pair, $(\mathcal{N}, \mathcal{E})$ is a directed graph per Definition II.18.

(G.2) \mathcal{X} is a collection of manifolds for each vertex. That is, for each $i \in \mathcal{N}$, \mathcal{X}_i is a smooth manifold and \mathcal{X} is the following disjoint union

$$\mathcal{X} = \bigsqcup_{i \in \mathcal{N}} \mathcal{X}_i.$$

(G.3) S is a collection of impact surfaces corresponding to the edges \mathcal{E} . Each impact surface, $S_e \subset \mathcal{X}_{\text{dom}(e)}$ is a nonempty embedded codimension 1 submanifold.

(G.4) Δ is a collection of impact maps. For each edge $e \in \mathcal{E}$, $\Delta_e : S_e \rightarrow \mathcal{X}_{\text{cod}(e)}$ and is smooth.

(G.5) f is a collection of smooth vector fields on \mathcal{X} . That is, for each $i \in \mathcal{N}$, $f_i : \mathcal{X}_i \rightarrow T\mathcal{X}_i$ is a smooth vector field.

Additionally, a general hybrid dynamical system, \mathcal{GH} , is said to be essentially nonbeating if, in addition to (G.1)-(G.5), it also satisfies

(G.6) $\overline{\Delta_e(S_e)} \cap \overline{S_{(\text{cod}(e), \cdot)}} \subset \mathcal{X}_{\text{cod}(e)}$ has codimension at least 2. Here,

$$S_{(\text{cod}(e), \cdot)} := \bigcup_{\substack{e' \in \mathcal{E} \\ \text{dom}(e') = \text{cod}(e)}} S_{e'}.$$

This section concludes with the definition of a solution to a GHDS (this will use the notion of a forgetful hybrid arc rather than a hybrid arc, but it can be defined for either).

Definition II.22. An arc, $x : [0, T) \rightarrow \mathcal{X}$ is a solution to $(\mathcal{N}, \mathcal{E}, \mathcal{X}, S, f, \Delta)$ if x is piecewise differentiable and

1. if for all $c \in (a, b)$, $x(c) \notin S_e$ for any $e \in \mathcal{E}$, then x is differentiable on (a, b) , $x(a, b) := \{x(t) : t \in (a, b)\}$ is contained in a single \mathcal{X}_i , and $\dot{x}(t) = f_i(x(t))$ for all $t \in (a, b)$, and
2. if $\lim_{t \rightarrow c^-} x(t) \in S_e$ for some $e \in \mathcal{E}$, then

$$\lim_{t \rightarrow c^+} x(t) = \Delta_e \left(\lim_{t \rightarrow c^-} x(t) \right).$$

Likewise, the notion of a hybrid flow per Definition II.11 can be extended to this case along with the regularity assumptions (A.1), (A.2) and Assumption II.13/II.14. A smooth GHDS is analogously defined as well.

CHAPTER III

Geometric Mechanics and Nonholonomic Mechanics

This chapter summarizes some basic concepts from geometric mechanics. We will assume some basic knowledge of differential topology / geometry [93, 107, 108], in particular the language of differential forms. This chapter starts with some basic concepts from symplectic and Poisson geometry [1, 11, 34, 77]. Next, (unconstrained) Lagrangian and Hamiltonian mechanics are introduced as well as the influence of symmetries [1, 6]. Finally, basics of nonholonomic systems are reviewed [16, 14, 83] with an emphasis on the Hamiltonian formalism [73, 79].

For notation, let Q be a (finite-dimensional) smooth manifold, TQ and T^*Q are its tangent and cotangent bundles respectively. $\mathfrak{X}(Q)$ is the set of all smooth vector fields over Q , $\Omega^k(Q)$ is the set of differential k -forms, and $C^\infty(Q)$ is the set of \mathbb{R} -valued smooth functions on Q . Moreover, let i_X be interior multiplication, $[X, Y]$ be the vector field bracket, and \mathcal{L}_X be the Lie derivative. If $M \rightarrow Q$ is a vector bundle over Q , $\Gamma(M)$ is the set of sections $Q \rightarrow M$, e.g. $\Gamma(TQ) = \mathfrak{X}(Q)$ and $\Gamma(T^*Q) = \Omega^1(Q)$. Also, Einstein summation will be used throughout: repeated indices implies summation so $a_i b^i = \sum_i a_i b^i$.

3.1 Symplectic and Poisson Geometry

Most of the content here can be found in [1, 11, 34, 77]. Symplectic and Poisson structures are central to the study of Hamiltonian mechanical systems so this section begins with their definitions and some basic properties.

3.1.1 Symplectic Manifolds

Let M be a smooth (finite-dimensional) real manifold. Endowing M with a symplectic structure means distinguishing a certain 2-form.

Definition III.1. A symplectic form, ω , is a closed and non-degenerate 2-form. If a manifold, M , admits a symplectic form then the pair (M, ω) is a symplectic manifold.

The symplectic form allows for an important identification $\mathfrak{X}(M) \rightarrow \Omega^1(M)$:

$$\omega^\flat : \mathfrak{X}(M) \rightarrow \Omega^1(M),$$

where $\omega^\flat(X) = \omega(X, \cdot)$. Because ω is non-degenerate, ω^\flat is invertible and its inverse is called ω^\sharp . In addition to the invertible association $\mathfrak{X} \rightarrow \Omega^1$, the nondegeneracy of the symplectic form induces a canonical volume form.

Proposition III.2 (3.1.5 in [1]). *Let $\omega \in \Omega^2(M)$ be a symplectic form. Then the form $\omega^n \in \Omega^{2n}(M)$ is a volume form.*

Although not every manifold possesses a symplectic form, the theorem below shows that an important class of manifolds do have a natural symplectic form.

Theorem III.3 (3.2.10 in [1]). *Let Q be a smooth manifold and $M := T^*Q$ be its cotangent bundle. Furthermore let $\pi_Q : M \rightarrow Q$ be its bundle projection map and $d\pi_Q : TM \rightarrow TQ$ be its tangent map. Let $\alpha_q \in M$ (where $q \in Q$) and $w_{\alpha_q} \in T_{\alpha_q}M$. Define the linear map $\theta_{\alpha_q} : T_{\alpha_q}M \rightarrow \mathbb{R}$ by $w_{\alpha_q} \mapsto \alpha_q(d\pi_Q \cdot w_{\alpha_q})$. Then the function $\theta_0 : \alpha_q \mapsto \theta_{\alpha_q}$ is a 1-form and $\omega_0 := -d\theta_0$ is a symplectic form.*

This theorem shows that every cotangent bundle has a natural symplectic structure. Additionally, the 1-form θ_0 is called the tautological or Liouville 1-form.

The symplectic form provides three types of special submanifolds, one of which will be important in Chapter VIII. We consider first symplectic vector spaces and then generalize to manifolds.

Definition III.4. Let (V, ω) be a symplectic vector space (so $\omega : V \times V \rightarrow \mathbb{R}$ is non-degenerate and skew) and let $U \subset V$ be a subspace. The orthogonal complement of U is given by

$$U^\perp := \{v \in V : \omega(u, v) = 0, \forall u \in U\}.$$

The subspace, U , is called

- Isotropic if $U \subset U^\perp$,
- Coisotropic if $U^\perp \subset U$, and
- Lagrangian if $U = U^\perp$, i.e. isotropic and coisotropic.

Definition III.5. A submanifold $S \subset M$ of a symplectic manifold (M, ω) is

- Isotropic if for all $x \in S$, $T_x S \subset T_x M$ is an isotropic subspace,
- Coisotropic if for all $x \in S$, $T_x S \subset T_x M$ is a coisotropic subspace, and
- Lagrangian if for all $x \in S$, $T_x S \subset T_x M$ is a Lagrangian subspace.

3.1.2 Poisson Manifolds

Poisson manifolds are a slightly more general version of symplectic manifolds as they do not possess the non-degeneracy assumption.

Definition III.6. Let P be a manifold and $C^\infty(P)$ be the set of real-valued smooth functions on P . Consider a bracket operation

$$\{\cdot, \cdot\} : C^\infty(P) \times C^\infty(P) \rightarrow C^\infty(P).$$

The pair, $(P, \{\cdot, \cdot\})$, is called a Poisson manifold if for any collection of numbers $\alpha_i, \beta_j \in \mathbb{R}$ and functions $f^i, g^j, f, g, h \in C^\infty(P)$ we have

$$(PB.1) \quad \{\alpha_i f^i, \beta_j g^j\} = \alpha_i \beta_j \{f^i, g^j\}, \text{ (\mathbb{R}-bilinearity)}$$

$$(PB.2) \quad \{f, g\} = -\{g, f\}, \text{ (skew-symmetry)}$$

$$(PB.3) \quad \{\{f, g\}, h\} + \{\{h, f\}, g\} + \{\{g, h\}, f\} = 0, \text{ (Jacobi's identity)}$$

$$(PB.4) \quad \{fg, h\} = f\{g, h\} + g\{f, h\}. \text{ (Leibniz's rule)}$$

Proposition III.7. *Symplectic manifolds are naturally Poisson.*

Proof. It is straight-forward to check that $\{f, g\} = \omega(\omega^\sharp(df), \omega^\sharp(dg))$ is a Poisson bracket. \square

Additionally, if the condition (PB.3) (Jacobi's identity) is relaxed, then the bracket is called *almost Poisson*. In this situation, the failure of Jacobi's identity is called the *Jacobiator*:

$$(3.1) \quad \text{Jac}(f, g, h) := \{\{f, g\}, h\} + \{\{h, f\}, g\} + \{\{g, h\}, f\}.$$

Due to the fact that all symplectic manifolds are automatically Poisson, is the converse true? This is true modulo degeneracy: ω must be nondegenerate but $\{\cdot, \cdot\}$ is allowed to be degenerate, i.e. the set of Casimir functions may be nontrivial:

$$Z(P) := \{f \in C^\infty(P) : \forall g \in C^\infty(P), \{f, g\} = 0\}.$$

We end this section with an explanation (when the bracket is nondegenerate) that (PB.3) is equivalent to ω being closed.

Theorem III.8 (Theorem 4.3 in [34]). *Let $\alpha \in \Omega^2(M)$ be a nondegenerate 2-form.*

Then the induced almost-Poisson bracket given by $\{f, g\} := \alpha(\alpha^\sharp(df), \alpha^\sharp(dg))$ has the following Jacobiator:

$$\text{Jac}(f, g, h) = d\alpha(\alpha^\sharp(df), \alpha^\sharp(dg), \alpha^\sharp(dh)).$$

In particular, Jacobi's identity is satisfied if and only if $d\alpha = 0$.

Proof. For this proof, we will use the notation $X_f := \alpha^\sharp(df)$. The first thing we will compute is $\alpha([X_f, X_g], X_h)$. Due to the Lie derivative being a derivation, we have

$$\begin{aligned} i_{[X_f, X_g]}\alpha &= \mathcal{L}_{X_f}i_{X_g}\alpha - i_{X_g}\mathcal{L}_{X_f}\alpha \\ &= \mathcal{L}_{X_f}dg - i_{X_g}(i_{X_f}d\alpha + di_{X_f}\alpha) \\ &= d(\mathcal{L}_{X_f}g) - i_{X_g}i_{X_f}d\alpha. \end{aligned}$$

This leads to

$$(3.2) \quad \alpha([X_f, X_g], X_h) = -\{\{f, g\}, h\} - d\alpha(X_f, X_g, X_h).$$

To finish the proof, we use the following identity for the differential of a 2-form:

$$(3.3) \quad \begin{aligned} d\alpha(X_f, X_g, X_h) &= \mathcal{L}_{X_f}\{X_g, X_h\} - \mathcal{L}_{X_g}\{X_f, X_h\} + \mathcal{L}_{X_h}\{X_f, X_g\} \\ &\quad - \alpha([X_f, X_g], X_h) + \alpha([X_f, X_h], X_g) - \alpha([X_g, X_h], X_f). \end{aligned}$$

Applying (3.2) to (3.3) and collecting terms yields the result. \square

3.2 Unconstrained Geometric Mechanics

We now introduce two main formulations of mechanics: Lagrangian and Hamiltonian. Lagrangian mechanics are described by variational principles and evolve on TQ (where Q is a smooth manifold) while Hamiltonian mechanics are described by symplectic principles and evolve on T^*Q (which always carries a symplectic/Poisson structure, Theorem III.3). We begin with an overview of the Lagrangian side and then the Hamiltonian. We note that the majority of this thesis will be concerned with the Hamiltonian side.

3.2.1 Lagrangian Mechanics

Lagrangian systems evolve on the tangent bundle of a smooth manifold. For our purposes here, we only consider *natural* Lagrangian systems (see §19-B of [6]).

Recall that a coordinate choice $q \in Q$ induces bundle coordinates $(q, \dot{q}) \in TQ$; so \dot{q} will be a coordinate in T_qQ . Moreover, for a curve $c : (-\varepsilon, \varepsilon) \rightarrow Q$ where $c(0) = q$, we can differentiate the curve to induce a point in the tangent space $\dot{c}(0) \in T_qQ$. Throughout this work, \dot{q} can refer to either notion.

Definition III.9. A Lagrangian function $L : TQ \rightarrow \mathbb{R}$ on a Riemannian manifold (Q, g) is natural if it is the difference between kinetic and potential energies, i.e.

$$L(q, \dot{q}) = \frac{1}{2}g_q(\dot{q}, \dot{q}) - V(q).$$

The Lagrangian generates dynamics on TQ by Hamilton's principle.

Definition III.10. Hamilton's principle singles out curves $q(t)$ by the condition that

$$(3.4) \quad \delta \int_a^b L(q(t), \dot{q}(t)) dt = 0,$$

where the variation is over smooth curves in Q with fixed end points.

A classical result in the calculus of variations states that (3.4) can be restated as a system of ordinary differential equations.

Proposition III.11. *Hamilton's principle is equivalent to the condition that the curve $q(t)$ satisfies the Euler-Lagrange equations*

$$(3.5) \quad \frac{d}{dt} \frac{\partial L}{\partial \dot{q}^i} - \frac{\partial L}{\partial q^i} = 0.$$

In the context where $V = 0$, L is only the kinetic energy and (3.5) is the geodesic equation.

3.2.2 Hamiltonian Mechanics

Given a natural Lagrangian function, $L : TQ \rightarrow \mathbb{R}$, there is an identification between TQ and T^*Q . This is given by the fiber derivative

$$\begin{aligned} \mathbb{F}L : TQ &\rightarrow T^*Q \\ \mathbb{F}L(v)(w) &:= \left. \frac{d}{dt} \right|_{t=0} L(q, v + tw) = g_q(v, w) \end{aligned}$$

The fiber derivative (which is a diffeomorphism when L is hyperregular; natural Lagrangians are hyperregular, cf. §3.6 in [1] for more details) offers a way to transform the function $L : TQ \rightarrow \mathbb{R}$ into a function $H : T^*Q \rightarrow \mathbb{R}$. This process is called the *Legendre transform*:

$$H(q, p) = \langle p, \dot{q} \rangle - L(q, \dot{q}), \quad p = \mathbb{F}L(\dot{q}).$$

While L generates dynamics on TQ variationally, H generates dynamics on T^*Q symplectically.

Definition III.12. Let (M, ω) be a symplectic manifold and $H : M \rightarrow \mathbb{R}$ a smooth function. A vector field X_H on M is called a Hamiltonian vector field with energy H if

$$(3.6) \quad i_{X_H} \omega = dH.$$

The triple (M, ω, H) is called a Hamiltonian system.

Remark III.13. An equivalent version of (3.6) using Poisson brackets is $X_H = \{\cdot, H\}$ where X_H is viewed as a derivation on $C^\infty(M)$.

An interesting property of Hamiltonian systems is that they are always volume preserving.

Proposition III.14. *The Hamiltonian dynamics preserve the symplectic form ω , and additionally preserve the induced volume form ω^n .*

Proof. This follows from Cartan's magic formula:

$$\mathcal{L}_{X_H}\omega = di_{X_H}\omega + i_{X_H}d\omega.$$

The first term disappears as $di_{X_H}\omega = ddH$ while the second vanishes as ω is closed.

□

As will be shown later in Chapter V, determining whether or not constrained Hamiltonian systems preserve measure is a much more involved question.

Finally, these two formulations, (3.5) and (3.6), are equivalent under FL.

Theorem III.15 (See Theorem 3.6.2 in [1]). *Let L be a natural Lagrangian on Q and H be its Legendre transform. Then the integral curves of (3.5) are mapped to the integral curves of (3.6) under FL. Furthermore, both systems have the same base integral curves.*

3.2.3 Symmetries in Hamiltonian Systems

Hamiltonian systems where the energy, H , is invariant under some symmetry offer a helpful way to find integrals, i.e. constants of motion. In what follows, we recap some fundamentals of symmetries in Hamiltonian mechanics, all of which can be found in Chapter 4 of [1], Appendix 5 of [6], and Chapter 4 of [11]. A similar approach can be used for symmetries in Poisson systems, but that will not be pursued here.

Symmetries

Throughout, we will assume that G is a Lie group.

Definition III.16. A Lie group, G , acts symplectically on (M, ω) if $\Phi : G \times M \rightarrow M$ is a group action, each $\Phi_g = \Phi(g, \cdot) : M \rightarrow M$ is a diffeomorphism, and $\Phi_g^* \omega = \omega$.

Definition III.17. Let G act on M by diffeomorphisms (not necessarily symplectically). Then for $\xi \in \mathfrak{g} = \text{Lie}(G)$, its *infinitesimal generator* $\xi_M \in \mathfrak{X}(M)$ is given by

$$(3.7) \quad \xi_M(x) = \left. \frac{d}{dt} \right|_{t=0} \Phi(\exp(t\xi), x).$$

When the group acts symplectically, the vector field has the property that $i_{\xi_M} \omega$ is closed. If $i_{\xi_M} \omega$ is exact, the infinitesimal generator has a corresponding Hamiltonian which is the action's *momentum map*.

Definition III.18. Let (M, ω) be a connected symplectic manifold and let $\Phi : G \times M \rightarrow M$ be a symplectic group action. We say that the map

$$(3.8) \quad J : M \rightarrow \mathfrak{g}^*,$$

is a momentum map for the action if for all $\xi \in \mathfrak{g}$ we have

$$d\hat{J}(\xi) = i_{\xi_M} \omega,$$

where $\hat{J}(\xi)(x) = \langle J(x), \xi \rangle$. The 4-tuple, (M, ω, Φ, J) , is called a Hamiltonian G -space.

In general, momentum maps need not exist for a given group action (which is because a symplectic vector field only requires that $i_{\xi_M} \omega$ is closed and not exact). However, in the special case where $M = T^*Q$ and the action is lifted from an action on Q , momentum maps exist and can be explicitly computed.

Definition III.19. Let $\Phi : G \times M \rightarrow M$ be a symplectic group action with momentum map $J : M \rightarrow \mathfrak{g}^*$. Then J is said to be Ad^* -equivariant if $J(\Phi_g(x)) = \text{Ad}_{g^{-1}}^* J(x)$ for

all $g \in G$ and $x \in M$. Equivalently, the following diagram commutes.

$$\begin{array}{ccc} M & \xrightarrow{\Phi_g} & M \\ \downarrow J & & \downarrow J \\ \mathfrak{g}^* & \xrightarrow{\text{Ad}_{g^{-1}}^*} & \mathfrak{g}^* \end{array}$$

Theorem III.20 (4.2.10 in [1]). *Let $\Phi : G \times M \rightarrow M$ be a symplectic group action. Assume that $\omega = -d\theta$ and $\Phi_g^*\theta = \theta$ for all $g \in G$. Then $J : M \rightarrow \mathfrak{g}^*$ defined by*

$$\langle J(x), \xi \rangle = (i_{\xi_M}\theta)(x),$$

is an Ad^ -equivariant momentum mapping for the action.*

Corollary III.21. *Let $\Phi : G \times Q \rightarrow Q$ be a smooth action and $\Psi : G \times T^*Q \rightarrow T^*Q$ be the lifted action. Then Ψ is symplectic with Ad^* -equivariant momentum map*

$$(3.9) \quad \hat{J}(\xi) = P(\xi_Q),$$

*where ξ_Q is the infinitesimal generator of Φ on Q and for a vector field $X \in \mathfrak{X}(Q)$, $P(X) : T^*Q \rightarrow \mathbb{R}$ is its momentum:*

$$P(X)(x, p) = \langle p, X(x) \rangle.$$

When the Hamiltonian is invariant under the symmetry, $H \circ \Phi_g = H$, the momentum map is a constant of the motion; this result is known as Noether's theorem [85] which states that each symmetry has a corresponding conservation law. We will state the version dealing with Hamiltonian systems, cf. §3.9 in [16]; for an alternate version dealing with Lagrangian systems, see §20 in [50].

Theorem III.22 (3.9.2 in [16]). *If H is a G -invariant Hamiltonian, then J is conserved along the trajectories of the Hamiltonian vector field X_H .*

Proof. Differentiating the invariance condition $H(\Phi_g(x)) = H(x)$ with respect to $g \in G$ gives

$$0 = dH(\xi_M) = \{H, \hat{J}(\xi)\} = -d\hat{J}(\xi)(X_H),$$

and so $\hat{J}(\xi)$ is conserved along the trajectories of X_H for every $\xi \in \mathfrak{g}$. \square

Reduction

If a Hamiltonian is invariant under the action of a symmetry, $H = H \circ \Phi_g$, then the momentum map is preserved under the flow. This allows us to replace the manifold, M , with $J^{-1}(\mu) \subset M$ and effectively reduce the dynamics to a lower-dimensional manifold. However when the momentum is Ad^* -equivariant, the dynamics can be reduced further by quotienting out the isotropy subgroup, $G_\mu := \{g \in G : \text{Ad}_{g^{-1}}^* \mu = \mu\}$. This works, in part, because $J(x) = J(\Phi_g(x)) = \mu$ for all $g \in G_\mu$. As it turns out, this new manifold $J^{-1}(\mu)/G_\mu$ is still a symplectic manifold as the theorem below shows.

Theorem III.23 (4.3.1 in [1]). *Let (M, ω) be a symplectic manifold and $\Phi : G \times M \rightarrow M$ be a symplectic group action with Ad^* -equivariant momentum map $J : M \rightarrow \mathfrak{g}^*$. Assume that $\mu \in \mathfrak{g}^*$ is a regular value of J and that the isotropy group G_μ acts freely and properly on $J^{-1}(\mu)$. Then $M_\mu := J^{-1}(\mu)/G_\mu$ has a unique symplectic form ω_μ with the property*

$$\pi_\mu^* \omega_\mu = \iota_\mu^* \omega,$$

where $\pi_\mu : J^{-1}(\mu) \rightarrow M_\mu$ is the canonical projection and $\iota_\mu : J^{-1}(\mu) \rightarrow M$ is the inclusion.

The above theorem shows how to reduce the manifold while the theorem below shows how to reduce the Hamiltonian.

Theorem III.24 (4.3.5 in [1]). *Under the assumptions of Theorem III.23, let $H : M \rightarrow \mathbb{R}$ be a Hamiltonian such that $H \circ \Phi_g = H$ for all $g \in G$. Then the flow of*

X_H leaves $J^{-1}(\mu)$ invariant and commutes with the action of G_μ on $J^{-1}(\mu)$ which induces a flow on M_μ . This flow is Hamiltonian on M_μ and its Hamiltonian, H_μ , is called the reduced Hamiltonian.

3.3 Nonholonomic Mechanics

Suppose that a Lagrangian system $L : TQ \rightarrow \mathbb{R}$ is subject to certain constraints, i.e. a figure skater who cannot slide perpendicular to the direction of her skate. Constraints involving the velocities of the system are known as nonholonomic constraints (holonomic constraints involve only the positions). Everything below that holds true for nonholonomic constraints also works for holonomic constraints. For the most part, we will assume that the constraints are linear in the velocities.

3.3.1 Constraint Distributions

Nonholonomic constraints are normally described as specifying a submanifold $\mathcal{D} \subset TQ$ that describes the restricted motion. When the constraints are linear in the velocities (which will be assumed throughout), the submanifold \mathcal{D} is a *distribution*.

Definition III.25. A *smooth distribution* on a manifold M is the assignment to each $x \in M$ a subspace $\mathcal{D}_x \subset T_x M$, i.e. $\mathcal{D} \subset TM$ is a vector sub-bundle. A distribution \mathcal{D} is *involutive* if for any two vector fields X, Y on M with values in \mathcal{D} , $[X, Y]$ also has values in \mathcal{D} . A distribution \mathcal{D} is *regular* if $\dim(\mathcal{D}_x)$ is the same for every $x \in M$.

Theorem III.26 (Frobenius' Theorem). \mathcal{D} is involutive if and only if there is a foliation on M whose tangent bundle equals \mathcal{D} .

If \mathcal{D} is involutive, it is said to be integrable and the constraints are called holonomic. When \mathcal{D} is *not* involutive, it is nonintegrable and the constraints are nonholonomic.

Finally, constraint distributions are usually described by a family of 1-forms η^α .

$$\mathcal{D} = \bigcap_{\alpha=1}^m \ker \eta^\alpha, \quad \eta^\alpha \in \Omega^1(Q).$$

In this situation, the distribution is integrable if the 1-forms η^α can be chosen such that they are all closed: $d\eta^\alpha = 0$. Integrability can also be viewed through the lens of Definition III.5 by examining what type of submanifold $\mathcal{D}^0 \subset T^*Q$ is where

$$\mathcal{D}^0 = \{(q, p) \in T^*Q : p(\mathcal{D}_q) = 0\}$$

is the annihilator of \mathcal{D} .

Proposition III.27. *A constraint distribution $\mathcal{D} \subset TQ$ is integrable if and only if its annihilator $\mathcal{D}^0 = \text{Ann}(\mathcal{D}) \subset T^*Q$ is coisotropic.*

Proof. Let $\{X_i\}$ be a local frame for \mathcal{D} . Then the annihilator is locally described by the vanishing of momenta:

$$\mathcal{D}^0 = \{(x, p) \in T^*Q : P(X_i) = 0, \forall i\}.$$

Therefore for \mathcal{D}^0 to be coisotropic, we need $\{P(X^i), P(X^j)\}$ to vanish on \mathcal{D}^0 as well. Indeed (cf. 4.2.12 in [1]), $\{P(X^i), P(X^j)\} = -P([X^i, X^j])$. Therefore, the bracket vanishes on \mathcal{D}^0 if and only if $[X^i, X^j]$ is a section of \mathcal{D} , i.e. it is integrable. \square

3.3.2 Hamiltonian Nonholonomic Systems

It is important to note that nonholonomic systems are not described by variational principles (on the Lagrangian side) nor are they symplectic (on the Hamiltonian side). Rather than obeying Hamilton's principle, nonholonomic systems follow the Lagrange-d'Alembert principle. In the Hamiltonian setting, this manifests as (see [73, 79] and §5.8 in [16]):

$$(3.10) \quad i_{X_H^{\mathcal{D}}} \omega = dH + \lambda_\alpha \pi_Q^* \eta^\alpha,$$

where $\pi_Q : T^*Q \rightarrow Q$ is the cotangent projection and the λ_α are multipliers to enforce the constraints.

Let g be the Riemannian metric underlying the natural Hamiltonian, H (a Hamiltonian is natural if it comes from a natural Lagrangian). For each constraining 1-form η^α , let $W^\alpha \in \mathfrak{X}(Q)$ such that $g(W^\alpha, \cdot) = \eta^\alpha$ (equivalently, $W^\alpha = \mathbb{F}L^{-1}\eta^\alpha$). This association is given by the musical isomorphisms, $\flat : \mathfrak{X}(Q) \rightarrow \Omega^1(Q)$ and $\sharp : \Omega^1(Q) \rightarrow \mathfrak{X}(Q)$. The constraint distribution $\mathcal{D} \subset TQ$ on the cotangent side becomes

$$(3.11) \quad \mathcal{D}^* = \{(x, p) \in T^*Q : P(W^\alpha)(x, p) = 0\}$$

Therefore the multipliers λ_α are chosen such that $X_H^{\mathcal{D}}$ is tangent to $\mathcal{D}^* \subset T^*Q$.

3.3.3 Symmetries in Nonholonomic Systems

We end this chapter with a section on how §3.2.3 relates to nonholonomic systems, cf. §5.4-5.7 of [16] as well as [14, 22] for a similar treatment. The following is for Lagrangian systems but can be carried over to Hamiltonian systems without too many issues.

Definition III.28. Let $L : TQ \rightarrow \mathbb{R}$ be a hyperregular Lagrangian (so $\mathbb{F}L$ is a diffeomorphism) and $\mathcal{D} \subset TQ$ be the constraint distribution defined by

$$\mathcal{D} := \{(q, \dot{q}) \in TQ : \eta^\alpha(\dot{q}) = 0, \forall \alpha\}.$$

(L.1) We say that the Lagrangian is **invariant** under the group action if $L \circ \Phi_g = L$ for all $g \in G$.

(L.2) We say that the Lagrangian is **infinitesimally invariant** if for any Lie algebra element $\xi \in \mathfrak{g}$, we have $dL(\xi_P) = 0$. Here, ξ_P is the infinitesimal generator of the lifted action to $P = TQ$.

(S.1) We say that the distribution, \mathcal{D} , is **invariant** if the subspace $\mathcal{D}_q \subset T_q Q$ is mapped by the tangent of the group action to the subspace $\mathcal{D}_{gq} \subset T_{gq} Q$, i.e. if $T_q \Phi_g \cdot \mathcal{D}_q \subset \mathcal{D}_{gq}$.

(S.2) A Lie algebra element ξ is said to act **horizontally** if $\xi_Q(q) \in \mathcal{D}_q$ for all $q \in Q$.

Proposition III.29. *Assume that \mathcal{D} is an invariant distribution in the sense of (S.1). Then the following are true:*

$$(3.12) \quad \mathcal{L}_{\Gamma(\mathfrak{g}_Q)} \Gamma(\mathcal{D}^0) \subset \Gamma(\mathcal{D}^0),$$

$$(3.13) \quad [\Gamma(\mathfrak{g}_Q), \Gamma(\mathcal{D})] \subset \Gamma(\mathcal{D}),$$

where $\mathcal{D}^0 \subset T^*Q$ is the annihilator of $\mathcal{D} \subset TQ$ and Γ denotes the sections of a vector bundle.

Proof. First, if $\eta \in \Gamma(\mathcal{D}^0)$, then (S.1) states that $\Phi_g^* \eta \in \Gamma(\mathcal{D}^0)$ as well. Differentiating this shows that $\mathcal{L}_{\xi_Q} \eta \in \mathcal{D}^0$. To show (3.13), we use Cartan's magic formula for the Lie derivative. Let $X \in \Gamma(\mathcal{D})$ and $\eta \in \Gamma(\mathcal{D}^0)$:

$$\begin{aligned} (\mathcal{L}_{\xi_Q} \eta)(X) &= i_{\xi_Q} d\eta(X) + di_{\xi_Q} \eta(X) \\ &= \mathcal{L}_{\xi_Q}(\eta(X)) - \mathcal{L}_X(\eta(\xi_Q)) - \eta([\xi_Q, X]) + \mathcal{L}_X(\eta(\xi_Q)) \\ &= -\eta([\xi_Q, X]) \end{aligned}$$

Due to the fact that we know $\mathcal{L}_{\xi_Q} \eta \in \Gamma(\mathcal{D}^0)$, it must be the case that $[\xi_Q, X] \in \mathcal{D}$. \square

Remark III.30. In the case of holonomic constraints, (S.2) implies (S.1). This can be seen in two ways. The first is noticing that the symmetries being horizontal means that $\Gamma(\mathfrak{g}_Q) \subset \Gamma(\mathcal{D})$ which gives (3.13) by Frobenius's theorem:

$$[\Gamma(\mathfrak{g}_Q), \Gamma(\mathcal{D})] \subset [\Gamma(\mathcal{D}), \Gamma(\mathcal{D})] \subset \Gamma(\mathcal{D}).$$

The second way of seeing this is from (3.12). Suppose that a holonomic constraint is given by $f = 0$. Then a symmetry is horizontal if $df(\xi_Q) = 0$. Computing the Lie

derivative gives

$$\mathcal{L}_{\xi_Q} df = di_{\xi_Q} df + i_{\xi_Q} ddf = 0 \in \mathcal{D}^0.$$

Moreover, if the constraint has the form gdf so it has the same kernel but is no longer exact, then

$$\mathcal{L}_{\xi_Q} gdf = g\cancel{\mathcal{L}_{\xi_Q} df} + (\mathcal{L}_{\xi_Q} g) df \in \mathcal{D}^0.$$

3.3.4 The Momentum Equation

It would be desirable to have a version of Noether's theorem for nonholonomic systems. This will generally not be exactly possible but a slightly weaker result, called the momentum equation, can still be achieved. First, we outline a few useful definitions.

Definition III.31. The intersection of the tangent space of the group orbit and the constraint distribution (in $T_q Q$, not $T_q^* Q$) is denoted by \mathcal{S}_q . i.e.

$$\mathcal{S}_q = \mathcal{D}_q \cap T_q(G \cdot q), \quad \mathcal{S} = \bigsqcup_{q \in Q} \mathcal{S}_q, \quad G \cdot q = \{\Phi_g(q) : g \in G\}.$$

Definition III.32. For each $q \in Q$, define the vector subspace $\mathfrak{g}^q \subset \mathfrak{g}$ to be all Lie algebra elements whose infinitesimal generators evaluated at q lie in \mathcal{S}_q :

$$\mathfrak{g}^q := \{\xi \in \mathfrak{g} : \xi_Q(q) \in \mathcal{S}_q\}.$$

The corresponding bundle over Q whose fibers are \mathfrak{g}^q is denoted by $\mathfrak{g}^{\mathcal{D}}$.

With the bundle $\mathfrak{g}^{\mathcal{D}}$, we can now modify (3.9) to a *nonholonomic momentum map*.

Definition III.33. The nonholonomic momentum map, $J^{\mathcal{D}}$, is the bundle map taking \mathcal{D}^* to the bundle $(\mathfrak{g}^{\mathcal{D}})^*$ whose fiber over the point $q \in Q$ is the dual of the vector space \mathfrak{g}^q that is defined by

$$(3.14) \quad J^{\mathcal{D}}(p_q)(\xi) = \langle p_q, \xi_Q(q) \rangle.$$

This can be thought of as a restricted version of the unconstrained momentum map, J . Due to the fact that symmetries may not be horizontal, $J^{\mathcal{D}}$ may not be preserved. Even though momentum is not conserved (so Noether's theorem does not hold), its dynamics can still be understood (see 5.5.4 and 5.5.5 in [16]).

Theorem III.34 (The Momentum Equation). *Assume that the Hamiltonian is invariant under the group action so $H \circ \Phi = H$ (which is equivalent to (L.2) via the Legendre transform) and that ξ^q is a section of $\mathfrak{g}^{\mathcal{D}}$. Then a solution to (3.10) must satisfy the momentum equation,*

$$(3.15) \quad \frac{d}{dt} (\hat{J}^{\mathcal{D}}(\xi^{q(t)})) = P \left(\left[\frac{d}{dt} \xi^{q(t)} \right]_Q \right).$$

Therefore, the failure for the momentum to be conserved is a direct consequence of ξ^q varying to remain in $\mathfrak{g}^{\mathcal{D}}$, i.e. failure of the symmetry to be horizontal. This leads to a direct corollary.

Corollary III.35. *If ξ is a horizontal symmetry, condition (S.2), then the following conservation law holds:*

$$\frac{d}{dt} (\hat{J}^{\mathcal{D}}(\xi^{q(t)})) = 0.$$

Remark III.36. In proving the above theorem and corollary, the assumption (S.1) is not required. However, when (S.1) is assumed, the momentum map is G -equivariant. See Remark 5 on page 263 of [16].

3.4 Hamilton-Jacobi Theory

Hamilton-Jacobi theory provides a powerful way to integrate Hamilton's equations. This is a well-studied field of classical mechanics and below we present some preliminaries in the spirit of [1, 77] and focus on the geometry of solutions. The ideas

presented here will be used to define a hybrid Hamilton-Jacobi theory which is the topic of Chapter VIII.

Everything in this section and Chapter VIII will deal with only unconstrained mechanics. While there is a Hamilton-Jacobi theory for nonholonomic systems [38, 63, 86], we will not pursue it here and relegate that to future research.

3.4.1 The Hamilton-Jacobi Equation

We introduce the Hamilton-Jacobi partial differential equation from a geometric viewpoint [1].

Theorem III.37 (Hamilton-Jacobi, cf. 5.2.18 in [1]). *Let $P = T^*Q$ and $X_H \in \mathfrak{X}(P)$ be a Hamiltonian vector-field. Let $S : Q \times \mathbb{R} \rightarrow \mathbb{R}$ be a smooth function. Then the following are equivalent:*

(HJ.1) *For every curve $c : I \rightarrow Q$ satisfying*

$$(3.16) \quad \dot{c}(t) = d\pi_Q \cdot X_H(dS_t(c(t))),$$

the curve $t \mapsto dS_t(c(t))$ is an integral curve of X_H .

(HJ.2) *The function S satisfies the Hamilton-Jacobi equation:*

$$(3.17) \quad H \circ dS_t + \frac{\partial}{\partial t} S_t = 0.$$

*Here, $\pi_Q : T^*Q \rightarrow Q$ is the usual cotangent projection and $S_t(x) = S(x, t)$.*

If $W : Q \rightarrow \mathbb{R}$ is a smooth function that solves $H \circ dW = E$ for some constant E , then the function

$$S(x, t) = W(x) - E \cdot t,$$

solves (3.17). Solutions of this type will be considered stationary. As we will see in the next subsection, stationary solutions correspond to invariant Lagrangian submanifolds.

3.4.2 Lagrangian Submanifolds and Integrability

Solutions to the partial differential equation (3.17) can be interpreted as Lagrangian submanifolds of (T^*Q, ω) where ω is the canonical symplectic form, cf. §5.3 in [1] or §9.4 in [77].

A Lagrangian submanifold, recall Definition III.5, $\mathfrak{L} \subset T^*Q$ has the property that if $\iota : \mathfrak{L} \hookrightarrow T^*Q$ is the inclusion, then $\iota^*\omega = 0$ and $\dim \mathfrak{L} = n$ (i.e. isotropic and maximal dimension). A broad class of these submanifolds have an elegant description based on graphs of differential forms.

Proposition III.38 (5.3.15 in [1]). *Let $\alpha \in \Omega^1(Q)$ and $\Gamma_\alpha \subset T^*Q$ be its graph. Then Γ_α is Lagrangian if and only if α is closed.*

Proof. Recall that the standard symplectic form on T^*Q is exact, $\omega = -d\theta$, where θ has the property $\alpha^*\theta = \alpha$. Then

$$d\alpha = \alpha^*d\theta = -\alpha^*\omega,$$

and α is closed if and only if $\alpha^*\omega = 0$, i.e. isotropic. The fact that $\dim \Gamma_\alpha = n$ finishes the proof. \square

If a Lagrangian submanifold projects diffeomorphically to Q (as the zero-section), we obtain a one-to-one correspondence between closed 1-forms and Lagrangian submanifolds.

Definition III.39. Let $\mathfrak{L} \subset T^*Q$ be a Lagrangian submanifold. \mathfrak{L} is called closed if $\mathfrak{L} = \Gamma_\alpha$ and exact if α is exact. Moreover, given that $\mathfrak{L} = \Gamma_\alpha$ on some neighborhood, α can be integrated and the function $S : Q \rightarrow \mathbb{R}$ such that $dS = \alpha$ is called the generating function for \mathfrak{L} .

These generating functions have a close relationship with the Hamilton-Jacobi partial differential equation which utilizes the following proposition.

Proposition III.40 (5.3.32 in [1]). *Let (M, ω) be a symplectic manifold and $\mathfrak{L} \subset M$ be a Lagrangian submanifold. Let $H \in C^\infty(M)$ and φ_t be the flow of the Hamiltonian vector-field X_H .*

1. *If H is constant on \mathfrak{L} , then \mathfrak{L} is invariant under φ_t .*
2. *$\varphi_t(\mathfrak{L})$ remains a Lagrangian submanifold.*

Suppose that \mathfrak{L}_0 has generating function S_0 . Then (locally and for a short time), the manifold $\mathfrak{L}_t = \varphi_t(\mathfrak{L}_0)$ has a generating function S_t . This results in the following observation (which follows from (HJ.1) and (HJ.2)):

The function $S(x, t) = S_t(x)$ described above solves the Hamilton-Jacobi equation.

Therefore, solutions to the Hamilton-Jacobi equation (3.17) are characterized by the orbits of Lagrangian submanifolds while stationary solutions correspond to invariant Lagrangian submanifolds. This leads us into the final definition of this chapter: completely integrable systems.

Definition III.41. Let (M, ω, H) be a Hamiltonian system and $f_1 = H, f_2, \dots, f_n$ be constants of motion. The constants are in involution if $\{f_i, f_j\} = 0$ for all i, j . The constants are independent if the set of critical points of $F = f_1 \times \dots \times f_n$ has measure zero in M . Finally, the system is completely integrable if there exists a set of n constants both in involution and independent such that $\dim M = 2n$.

In the language of Lagrangian submanifolds, a Hamiltonian system given by (M, ω, H) is completely integrable if there exists a Lagrangian foliation $\{\mathfrak{L}_\alpha\}$ such that each leaf has constant energy.

Definition III.42. A foliation $\mathcal{F} = \{\mathcal{L}_\alpha\}$ of dimension p on an m -dimensional manifold M is a decomposition of M into disjoint connected subsets $\{\mathcal{L}_\alpha\}$ called the leaves of the foliation such that each point of M has a neighborhood U and a system of coordinates $(x, y) : U \rightarrow \mathbb{R}^p \times \mathbb{R}^{m-p}$ such that for each leaf \mathcal{L}_α , the components of $U \cap \mathcal{L}_\alpha$ are described by

$$y_1 = \text{constant}, \dots, y_{m-p} = \text{constant}.$$

Recall that finding a Lagrangian submanifold with constant energy provides a stationary solution to the Hamilton-Jacobi equation. As we will see in Chapter VIII, it will be natural to define hybrid systems as being completely integrable in terms of a Lagrangian foliation.

CHAPTER IV

Stability and Asymptotic Behavior of Hybrid Systems

This chapter begins the second part and deals with some basics structure of solutions to HDSs. First, the notion of a hybrid ω -limit set is introduced and is shown to “usually” be invariant under the flow. Next, stability theory for hybrid systems is developed and investigated via a Poincaré map and Floquet theory. Finally, a version of the Poincaré-Bendixson theorem is proved for a general class of planar hybrid systems.

Throughout this chapter, the following notation will be used: $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ is a smooth hybrid dynamical system (recall Definition II.17) with hybrid flow $\varphi^{\mathcal{H}} : \mathbb{R}^+ \times \mathcal{X} \rightarrow \mathcal{X}$ (or $\varphi_t^{\mathcal{H}} : \mathcal{X} \rightarrow \mathcal{X}$ to emphasize the dependence on initial conditions). Additionally, $\varphi : \mathbb{R}^+ \times \mathcal{X} \rightarrow \mathcal{X}$ is the flow for just the continuous dynamics, $\dot{x} = f(x)$.

As this chapter deals mostly with periodic orbits, fixed points will generally be ignored. For a vector field $f : \mathcal{X} \rightarrow T\mathcal{X}$, $\text{fix}(f)$ will be the fixed points of the continuous dynamics $\dot{x} = f(x)$, that is

$$\text{fix}(f) := \{x \in \mathcal{X} : f(x) = 0\}.$$

4.1 Limit Sets

For a HDS, the ω -limit set is defined in an analogous fashion to the continuous case.

Definition IV.1 (Continuous ω -limit set). The (forward) *orbit* and the ω -limit set for the (complete) flow $\varphi_t(x)$ are given by

$$\begin{aligned} o^+(x) &:= \{\varphi_t(x) : t \in \mathbb{R}^+\}, \\ \omega(x) &:= \left\{ y \in \mathcal{X} : \exists t_n \rightarrow \infty \text{ s.t. } \lim_{n \rightarrow \infty} \varphi_{t_n}(x) = y \right\}. \end{aligned}$$

Definition IV.2 (Hybrid ω -limit set). The (forward) *orbit* and the ω -limit set for the (complete and non-Zeno) hybrid flow $\varphi_t^{\mathcal{H}}(x)$ are given by

$$\begin{aligned} o_{\mathcal{H}}^+(x) &:= \{\varphi_t^{\mathcal{H}}(x) : t \in \mathbb{R}^+\}, \\ \omega_{\mathcal{H}}(x) &:= \left\{ y \in \mathcal{X} : \exists t_n \rightarrow \infty \text{ s.t. } \lim_{n \rightarrow \infty} \varphi_{t_n}^{\mathcal{H}}(x) = y \right\}. \end{aligned}$$

Throughout this text, their continuous counterparts will be denoted by o_c^+ and ω_c .

Remark IV.3. Dropping the condition (H.4) on the HDS (which says that $\Delta : S \rightarrow \mathcal{X}$ is smooth) can result in systems where $\omega_{\mathcal{H}}(x)$ is not invariant under the flow. That is, if $p \in \omega_{\mathcal{H}}(x)$, then $o_{\mathcal{H}}^+(p) \not\subset \omega_{\mathcal{H}}(x)$. The following example shows this situation.

Example IV.4. Consider the following HDS: $\mathcal{X} = [0, 1] \times \mathbb{R} \subset \mathbb{R}^2$ and the continuous flow is determined by $\dot{x} = 1$ and $\dot{y} = -y^2$. Let the impact surface be $S = \{(1, y) : y \in \mathbb{R}\}$ and the impact map be given by

$$\Delta(1, y) = \begin{cases} (0, y), & y > 0 \\ (0, y - 1), & y \leq 0. \end{cases}$$

The $\omega_{\mathcal{H}}$ -limit set of the initial condition $(0, 1)$ is the interval $[0, 1] \times \{0\}$ which is clearly not invariant under the flow of the system because the impact moves the flow away from $\omega_{\mathcal{H}}(0, 1)$.

Two properties of the $\omega_{\mathcal{H}}$ -limit set are shown: sufficient conditions for $\omega_{\mathcal{H}}(x)$ to be nonempty and compact, and that $\omega_{\mathcal{H}}(x)$ is “usually” invariant under the flow.

Proposition IV.5 ([29]). *The $\omega_{\mathcal{H}}$ -limit set of a trajectory $o_{\mathcal{H}}^+(x)$ is a closed set. Additionally, if R is a compact and forward invariant set, then $\omega_{\mathcal{H}}(x)$ is nonempty and compact for $x \in R$.*

Proof. The proof follows the same arguments as for the continuous case (cf. [92] pp. 193). First, let us prove that $\omega_{\mathcal{H}}(x)$ is closed. Let $\{p_n\}_{n \in \mathbb{N}}$ be a sequence in $\omega_{\mathcal{H}}(x)$, such that $p_n \rightarrow p$ when $n \rightarrow \infty$. We want to show that $p \in \omega_{\mathcal{H}}(x)$. Since $p_n \in \omega_{\mathcal{H}}(x)$, there exists a sequence of times, $\{t_k^{(n)}\}$, such that $\varphi_{t_k^{(n)}}^{\mathcal{H}}(x) \rightarrow p_n$ when $t_k^{(n)} \rightarrow \infty$. Without loss of generality, consider $t_k^{(n+1)} > t_k^{(n)}$. Then, for all $n \geq 2$ there exists $K_n > K_{n-1}$ such that for all $k \geq K_n$

$$\left| \varphi_{t_k^{(n)}}^{\mathcal{H}}(x) - p_n \right| < \frac{1}{n}.$$

Choose a sequence of times $t_n = t_{K_n}^{(n)}$. Then, by the triangle inequality, as $t_n \rightarrow \infty$ we obtain that $\varphi_{t_n}^{\mathcal{H}}(x)$ converges to p , that is,

$$\left| \varphi_{t_n}^{\mathcal{H}}(x) - p \right| \leq \left| \varphi_{t_n}^{\mathcal{H}}(x) - p_n \right| + |p_n - p| \leq \frac{1}{n} + |p_n - p| \rightarrow 0 \text{ when } n \rightarrow \infty.$$

For the second part, we have that $\omega_{\mathcal{H}}(x) \subset R$, so it is compact since it is a closed subset of a compact set. To show that it is nonempty, we point out that the sequence $\{\varphi_n^{\mathcal{H}}(x)\}_{n \in \mathbb{N}}$ is in a compact set so by the Bolzano-Weierstrass Theorem, there exists a convergent subsequence. \square

Remark IV.6. Note that $\omega_{\mathcal{H}}(x)$ is closed while $o_{\mathcal{H}}^+(x)$ is not.

Before we prove that $\omega_{\mathcal{H}}(x)$ is an invariant set, we first define two objects.

Definition IV.7. Let $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ be a hybrid system and let d be a distance on \mathcal{X} . The hybrid distance, $d_{\mathcal{H}}$ is defined as the pseudometric on the quotient space

with equivalence classes

$$[x] = \begin{cases} \{x\}, & x \notin S \cup \Delta(S) \\ \Delta^{-1}(x), & x \in \Delta(S) \\ \Delta^{-1}(\Delta(x)), & x \in S. \end{cases}$$

That is,

$$d_{\mathcal{H}}(x, y) = \inf \sum_{i=1}^n d(z_i, w_i),$$

where $[x] = [p_1]$, $[y] = [q_n]$, and $[q_i] = [p_{i+1}]$. Additionally, for a set $Y \subset \mathcal{X}$, let $\mathcal{O}_t(Y)$ be the partial orbit:

$$\mathcal{O}_t(Y) := \{\varphi_{\tau}^{\mathcal{H}}(x) : \tau \in [0, t], x \in Y\}.$$

Remark IV.8. The hybrid distance will satisfy the identity of indiscernibles away from S , i.e. if $x, y \notin S$ then $d_{\mathcal{H}}(x, y) = 0$ implies $x = y$.

Under fairly weak conditions we can show invariance of the limit set, similar results are in [29, 56].

Proposition IV.9 (*). *If there exists an open set $U \subset S$ such that*

1. $\omega_{\mathcal{H}}(x) \cap S \subset U$, and

2. f is transverse to U ,

then $\omega_{\mathcal{H}}(x)$ is invariant under the flow, i.e. if $x \in \mathcal{X}$, for all $y \in \omega_{\mathcal{H}}(x)$, $o_{\mathcal{H}}^+(y) \subset \omega_{\mathcal{H}}(x)$.

Proof. Let $y \in \omega_{\mathcal{H}}(x)$ and $\{t_n\}_{n=0}^{\infty}$ be a sequence of times such that

$$\lim_{n \rightarrow \infty} \varphi_{t_n}^{\mathcal{H}}(x) = y.$$

*We thank Dr. Matthew Kvalheim for insightful conversations with this proposition and proof.

Suppose that $y \notin S$. Then for all $\varepsilon > 0$, and $t > 0$ (small enough that the trajectories remain within U on impacts) there exists $\delta > 0$ such that if $d_{\mathcal{H}}(y, z) < \delta$ then $d_{\mathcal{H}}(\varphi_t^{\mathcal{H}}(y), \varphi_t^{\mathcal{H}}(z)) < \varepsilon$. This is true because we can choose δ small enough that $\mathcal{B}_\delta(y) \cap S = \emptyset$ and $\mathcal{O}_t(\mathcal{B}_\delta(y)) \cap S \subset U$. Then continuity follows from the implicit function theorem, cf. Theorem II.16. We have

$$\varphi_t^{\mathcal{H}}(y) = \varphi_t^{\mathcal{H}}\left(\lim_{n \rightarrow \infty} \varphi_{t_n}^{\mathcal{H}}(x)\right) = \lim_{n \rightarrow \infty} \varphi_{t_n+t}^{\mathcal{H}}(x) \in \omega(x).$$

To finish the proof, we need to show that this works for not only small times and for $y \in S$. Arbitrary times follows from repeating the above argument. If $y \in S$, choose $\tilde{y} = \varphi_{-\varepsilon}(y)$ and choose $t > \varepsilon$. \square

Corollary IV.10. *The $\omega_{\mathcal{H}}$ -limit set is invariant for smooth hybrid systems.*

Remark IV.11. Strange things can happen when the flow is no longer transverse to the impact surface. However for every case considered here, the flow will be transverse.

Example IV.12 (A non-invariant $\omega_{\mathcal{H}}$ -limit set). Consider the hybrid system where $\mathcal{X} = \mathbb{R}^2$, $S = \{x_2 = 1\}$, $\Delta(x_1, x_2) = (x_1, x_2 - 3)$ and the continuous dynamics (in polar form)

$$\dot{\theta} = 1, \quad \dot{r} = 1 - r.$$

Then $\omega_{\mathcal{H}}(0.5, 0) = \{(x_1, x_2) : x_1^2 + x_2^2 = 1\}$. However, this is *not* invariant under the flow as it ignores the impact. The point of contention is $(0, 1)$ which is precisely where the vector field fails to be transverse to S .

4.2 Hybrid Poincaré Map and Floquet Theory

The bulk of the literature on stability for hybrid systems revolves around constructing and studying hybrid Lyapunov functions, e.g. [20, 24, 39, 48, 52, 57, 58,

62, 56, 105] to name a few. Rather than studying stability through this avenue, we will instead focus our attention on the stability of periodic orbits via Floquet theory [46]. By linearizing about the periodic orbit, we obtain a linear system with periodic coefficients whose stability guarantees stability for the original system. We will first discuss the Poincaré map (or first-return map) for a hybrid system and then show how Floquet theory makes it possible to numerically approximate the stability of this map. This theory is then applied to an example of a simple robotic walker experiencing foot-slip.

4.2.1 The Hybrid Poincaré Map

Given a periodic orbit for a HDS, we would like to determine its stability. We will study this by examining its Poincaré map [18, 26, 53, 82, 100].

Let $\gamma : [0, T] \rightarrow \mathcal{X}$ be a periodic orbit for a HDS which (for simplicity) crosses S once per period and without loss of generality assume that $\gamma(T) \in S$ (and therefore $\gamma(0) = \Delta(\gamma(T))$). To construct a Poincaré map, we need to take a section that is transverse to the flow so it seems natural to take S as the section [82]. The problem is that for a given $x \in S$, it is not immediate that this map should be differentiable. However, so long as the flow intersects S transversely, it will be differentiable by Theorem II.16. Recall that $\tau : \mathcal{B}_{x_0} \rightarrow \mathbb{R}$ is the time required to reach S , i.e. $\varphi_{\tau(y)}(y) \in S$. Throughout, we will be dealing with a single impact surface; for the case where the flow can intersect multiple impact surfaces simultaneously, see [19].

Definition IV.13. The map $P : \mathcal{B}_\varepsilon(x_0) \cap S \rightarrow S$ given by $y \mapsto \varphi_{\tau(\Delta(y))}(\Delta(y))$ is the Poincaré map or the first-return map.

The point x_0 lies on a periodic trajectory if and only if $P(x_0) = x_0$. This offers

a way to translate the hybrid problem to a discrete one; γ is stable under $\varphi^{\mathcal{H}}$ only if x_0 is stable under P (Theorem 13.1 [56]). It is therefore desirable to be able to differentiate P . We will proceed informally for intuitive purposes.

Let \tilde{x} be a (small) perturbation of γ by x_p . Then, we will use the following approximations to study the dynamics of x_p .

$$\begin{aligned}\dot{\tilde{x}} &= f(\gamma(t) + x_p(t)) \\ &\approx f(\gamma(t)) + \left. \frac{\partial f}{\partial x}(x) \right|_{x=\gamma(t)} \cdot x_p(t) \\ &= f(\gamma(t)) + A(t)x_p,\end{aligned}$$

$$\Delta(\tilde{x}_0) = \Delta(x_0 + x_p) \approx \Delta(x_0) + \Delta_*(x_0) \cdot x_p,$$

where $\Delta_*(x_0) : T_{x_0}S \rightarrow T_{\Delta(x_0)}\Delta(S)$ is its differential. The linearization of the impact requires that $x_p \in T_{x_0}S$. To adjust for $x_p \in T_{x_0}\mathcal{X}$, we augment Δ_* by including the map $f(x_0) \mapsto f(\Delta(x_0))$. Call this augmented linear map Δ_*^f .

Definition IV.14. Let $\mathcal{H} = (\mathcal{X}, S, f, \Delta)$ be a hybrid system and $x \in S$. Then the map $\Delta_*^f : T_x\mathcal{X} \rightarrow T_{\Delta(x)}\mathcal{X}$ is called the augmented differential where

$$\begin{aligned}\Delta_*^f \cdot u &= \Delta_* \cdot u, \quad u \in T_x S \subset T_x \mathcal{X}, \\ \Delta_*^f \cdot f(x) &= f(\Delta(x)).\end{aligned}$$

The linear map Δ_*^f is well-defined as long as f intersects S transversely. Equivalently suppose $S = h^{-1}(0)$, then Δ_*^f is well-defined as long as $\mathcal{L}_f h = dh(f) \neq 0$.

4.2.2 Floquet Theory

Consider that we have the following linear HDS:

$$(4.1) \quad \begin{cases} \dot{x} = A(t) \cdot x, & t \bmod T \neq 0 \\ x^+ = \Delta_*^f \cdot x^-, & t \bmod T = 0. \end{cases}$$

What can we say about the stability of the origin in terms of the matrices $A(t)$ and Δ_*^f ? For if the origin of this system is stable, then the periodic orbit of the original hybrid system will also be stable. The following theorem shows that the derivative of the Poincaré map is, indeed, given by the data in (4.1).

Remark IV.15. As is the case in both continuous and discrete dynamics, in order for eigenvalue analysis of the derivative of the Poincaré map to be helpful, it must be hyperbolic (no eigenvalues on the unit circle).

Theorem IV.16 ([26]). *Let $\mathcal{X} \subset \mathbb{R}^n$ and $\gamma : [0, T] \rightarrow \mathcal{X}$ be a periodic solution that impacts S exactly once. That is, $\gamma(T) \in S$ and $\gamma(0) = \Delta(\gamma(T))$. Call $x_0 = \gamma(0)$ and $x_f = \gamma(T)$. Assume, further, that the orbit is transverse to both S and $\Delta(S)$. Define the (nonorthogonal) projections*

$$\begin{aligned}\tilde{\pi} : T_{x_f} \mathcal{X} &= T_{x_f} S \oplus f(x_f) \mathbb{R} \rightarrow T_{x_f} S \\ \pi : T_{x_0} \mathcal{X} &= T_{x_0} \Delta(S) \oplus f(x_0) \mathbb{R} \rightarrow T_{x_0} \Delta(S),\end{aligned}$$

where \oplus is the direct sum. Then, the derivative of the Poincaré map is given by

$$(4.2) \quad P'(x_f) = \Delta_*(x_f) \cdot \tilde{\pi} \cdot \Phi(T) \Big|_{T_{x_0} S} = \pi \cdot \Delta_*^f \cdot \Phi(T) \Big|_{T_{x_0} S},$$

where $\Phi(t)$ is the principal fundamental matrix solution (see Chapter 6 of [31]) to $\dot{\Phi} = A(t)\Phi$.

Remark IV.17. Notice that perturbations in the flow direction result in an eigenvalue of 1 for the stroboscopic map of equation (4.1), $\Delta_*^f \cdot \Phi(T)$ (a stroboscopic map is a special case of the Poincaré map in which the return times are constant, cf., e.g. §13.3 of [65]). As a result, when calculating the (derivative of the) Poincaré map, it is normal to disregard the flow direction. That is the purpose of $\tilde{\pi}$, or π , in equation (4.2). Also, note that the requirement for the flow to intersect S and $\Delta(S)$

transversely is a requirement for $\tilde{\pi}$ and π to be defined. Therefore, \mathcal{H} being smooth is required to differentiate the Poincaré map.

Remark IV.18. The two right-sides of (4.1) are equivalent because the diagram below commutes.

$$\begin{array}{ccc} T_{x_f}M & \xrightarrow{\Delta_*^f} & T_{x_0}M \\ \downarrow \tilde{\pi} & & \downarrow \pi \\ T_{x_f}S & \xrightarrow{\Delta_*} & T_{x_0}S \end{array}$$

We now proceed with the proof for Theorem IV.16.

Proof. Let $\varphi_t(x_0)$ be the (continuous) flow. This means that

$$\frac{d}{dt}\varphi_t(x_0) = f(\varphi_t(x_0)).$$

Due to everything being smooth, we can use the Leibniz integration rule. Let \mathcal{B}_{x_0} be an open neighborhood of x_0 and $\tau : \mathcal{B}_{x_0} \rightarrow S$ the the time required to get to S . By Theorem II.16, τ is smooth. We have that $\tau(x_0) = T$ and recall that $\varphi_t(x_0) = \gamma(t)$.

$$\begin{aligned} (4.3) \quad \frac{\partial}{\partial x} \varphi_{\tau(x)}(x) \Big|_{x=x_0} &= \int_0^T \frac{\partial}{\partial x} f(\varphi_t(x_0)) dt + f(x_f) \cdot \tau'(x_0) \\ &= \int_0^T \frac{\partial f}{\partial y}(y) \Big|_{y=\gamma(t)} \cdot \frac{\partial}{\partial x} \varphi_t(x) \Big|_{x=x_0} dt + f(x_f) \cdot \tau'(x_0) \end{aligned}$$

Looking at the first part, we see that

$$\frac{\partial}{\partial x} \varphi_t(x) \Big|_{x=x_0} = \int_0^t A(s) \cdot \frac{\partial}{\partial x} \varphi_s(x) \Big|_{x=x_0} dt.$$

Therefore,

$$\frac{\partial}{\partial x} \varphi_t(x_0) = \Phi(t), \quad \dot{\Phi} = A(t)\Phi, \quad \Phi(0) = Id.$$

So, equation (4.3) becomes

$$\frac{\partial}{\partial x} \varphi_{\tau(x_0)}(x_0) = \Phi(T) + f(x_f) \cdot \tau'(x_0).$$

Let $v \in T_{x_0}\Delta(S)$. If $(\Phi(T) + f(x_f) \cdot \tau'(x_0))v \in T_{x_f}S$, we are done. Call $\tau'(x_0)v = \alpha$, then

$$(\Phi(T) + f(x_f) \cdot \tau'(x_0))v = \Phi(T)v + \alpha f(x_f).$$

Let $h : \mathcal{X} \rightarrow \mathbb{R}$ such that $h^{-1}(0) = S$ (which is locally possible by (H.2)). Then, we have that $h(\varphi_T(x_0)) = 0$. Differentiating with the chain rule yields

$$\left. \frac{\partial h}{\partial z} \right|_{z=x_f} \cdot (\Phi(T) + f(x_f) \cdot \tau'(x_0))v = 0.$$

Which tells us that

$$(\Phi(T) + f(x_f) \cdot \tau'(x_0))v \in \ker \left(\left. \frac{\partial h}{\partial z} \right|_{z=x_f} \right) = T_{x_f}S.$$

□

In order to understand the stability of a hybrid system, Δ_*^f is the correct object to study rather than Δ_* . This appears again in Chapter VII where hybrid-invariant differential forms are studied.

Generally the Poincaré map cannot be solved analytically and must be numerically calculated. However, in the special case where $\mathcal{X} = \mathbb{R}^2$, the function P' can be solved analytically. We first recall a helpful result about continuous flows. For what follows, $\mathcal{X} \subset \mathbb{R}^n$ and $\nabla \cdot f$ is the usual divergence.

Lemma IV.19 ([92], p. 86). *Let $\varphi_t(x_0)$ be the flow of $\dot{x} = f(x)$ with initial condition x_0 . Then,*

$$\det \left[\left. \frac{\partial}{\partial x} \varphi_t(x) \right|_{x=x_0} \right] = \exp \left(\int_0^t \nabla \cdot f(\varphi_s(x_0)) ds \right).$$

To understand the stability of our orbit, we want to look at the hybrid Poincaré return map, $P : \mathcal{B}_\varepsilon(x_0) \cap S \rightarrow S$. As in Theorem II.16, let $\tau : \Delta(\mathcal{B}_\varepsilon(x_0) \cap S) \rightarrow \mathbb{R}$ be the time required to return to the impact surface. Then, if we denote $y := \Delta(x)$, we

can write P as

$$P(x) = \varphi_{\tau(y)}(y) = \int_0^{\tau(y)} f(\varphi_s(y)) ds.$$

Theorem IV.20 ([29]). *Assume that $\mathcal{X} \subset \mathbb{R}^2$ and we have a hybrid periodic orbit that intersects S once. Let $x \in S$ and $y = \Delta(x)$ be the points of intersection on the orbit. Additionally, let θ be the angle $f(x)$ makes with the tangent of S at x and α be the angle of $f(y)$ with $\Delta(S)$. Assume that θ and α are not integer multiples of π . If we denote the continuous flow that connects y to x by $\gamma(t)$ and suppose that it takes time T to complete the loop, the derivative of the Poincaré map is*

$$(4.4) \quad P'(x) = \Delta'(x) \cdot \frac{\|f(y)\| \sin \alpha}{\|f(x)\| \sin \theta} \cdot \exp\left(\int_0^T \nabla \cdot f(\gamma(t)) dt\right).$$




Figure 4.1: The orbit of the periodic orbit for the system given by Theorem IV.20.

Proof. To differentiate P , let us first look at the continuous part (that is, starting at $y_0 = \Delta(x_0)$). Let $n(y)$ be the unit normal vector to $\Delta(S)$ at y and let p be the unit tangent vector. The assumption $\alpha \neq n \cdot \pi$ leads to $\langle f(y_0), n(y_0) \rangle \neq 0$.

$$\begin{aligned} \frac{\partial}{\partial p} \varphi_{\tau(y)}(y) \Big|_{y=y_0} &= \int_0^{\tau(y)} \frac{\partial}{\partial y} f(\varphi_s(y)) \Big|_{y=y_0} ds \cdot \frac{\partial y}{\partial p} + \frac{\partial}{\partial t} (\varphi_{\tau(y_0)}(y_0)) \cdot \frac{\partial t}{\partial p} \\ &= F(y_0) \cdot \delta y + G(y_0) \cdot \delta t \end{aligned}$$

Now call the flow $\varphi_t(y) =: \gamma(t)$, the time $T = \tau(y)$, and recall that the final point is $\varphi_{\tau(y)}(y) = x$. Then, $G(y) = f(x)$ and δy is the unit vector p rooted at y_0 . We need to figure out what δt and $F(y)$ are. By Lemma IV.19, we know the determinant of $F(y)$.

$$(4.5) \quad \det(F(y)) = \exp\left(\int_0^T \nabla \cdot f(\gamma(t)) dt\right)$$

To find F (in the direction of δy), we note that we know the derivative in the direction of the flow: $F(y) \cdot f(y) = f(x)$. Knowing the determinant and this direction, we can attempt to find $F(y)$ in the direction of δy . We first differentiate h along S , which is zero because S is a level set of h .

$$0 = \frac{\partial}{\partial p} h(\varphi_{\tau(y)}(y)) \Big|_{y=y_0} = \frac{\partial}{\partial x} h(x) \Big|_{x=x_0} \cdot (F(y) \cdot \delta y + f(x) \cdot \delta t) \Big|_{y=y_0}$$

This tells us that $F(y) \cdot \delta y + f(x) \cdot \delta t$ lies on the tangent to S at x . Let $V(u, v)$ be




Figure 4.2: The vector $\delta x = F(y)\delta y + f(x)\delta t$, where the horizontal line is the tangent to S at the point x .

the area of the parallelogram spanned by the two vectors u and v . Additionally, let $\Lambda = \det(F(y))$. Then, we have

$$\begin{aligned} V(f(y), \delta y) &= \|f(y)\| \cdot \|\delta y\| \sin \alpha \\ V(\underbrace{F(y) \cdot f(y)}_{=f(x)}, F(y) \cdot \delta y) &= \Lambda \|f(y)\| \cdot \|\delta y\| \sin \alpha \\ &= V(f(x), F(y) \cdot \delta y + f(x) \cdot \delta t) \\ &= \|f(x)\| \cdot \|\delta x\| \sin \theta. \end{aligned}$$

Collecting terms, we see that

$$\frac{\|\delta x\|}{\|\delta y\|} = \frac{\|f(y)\| \sin \alpha}{\|f(x)\| \sin \theta} \Lambda.$$

Combining this with equation (4.5), we arrive at equation (4.4). \square

Remark IV.21. If the impact surface, S , is a level-set of a function $h : \mathcal{X} \rightarrow \mathbb{R}$ (which can always be done locally), then (4.4) can be written as

$$P'(x) = \Delta'(x) \cdot \frac{\mathcal{L}_f h(y)}{\mathcal{L}_f h(x)} \cdot \exp\left(\int_0^T \nabla \cdot f(\gamma(t)) dt\right).$$

Corollary IV.22 ([29]). *Suppose now that we have a hybrid periodic orbit that intersects S n times. Let $x_1, \dots, x_n \in S$ and $y_i = \Delta(x_i)$. Additionally, let γ_i be the flow that connects y_i to x_{i+1} , i.e. $\gamma_i(0) = y_i$ and $\gamma_i(T_i) = x_{i+1}$. Also, let α_i be the angle $f(y_i)$ makes with $\Delta(S)$ and θ_i be the angle $f(x_i)$ makes with S . Then, the derivative of the Poincaré map is given by*

$$(P^n)'(x_1) = \prod_{i=1}^n \Delta'(x_i) \frac{\|f(y_i)\| \sin \alpha_i}{\|f(x_i)\| \sin \theta_i} \exp\left(\int_0^{T_i} \nabla \cdot f(\gamma_i(t)) dt\right).$$

This gives a precise test for determining the stability of planar hybrid orbits. We would like to extend this to higher dimensions, but we can only calculate $\det P'(x_0)$ and not its individual eigenvalues.

Theorem IV.23 ([29]). *Assume that $\mathcal{X} = \mathbb{R}^n$ and that $\gamma(\cdot)$ is a periodic orbit intersecting S once with $x \in S$ and $y = \Delta(x)$ and period length T . Let α and θ be described as in Theorem IV.20. If γ is stable, then*

$$(4.6) \quad \left| \det(\Delta'(x)) \frac{\|f(y)\| \sin \alpha}{\|f(x)\| \sin \theta} \cdot \exp\left(\int_0^T \nabla \cdot f(\gamma(t)) dt\right) \right| \leq 1.$$

Proof. Equation (4.6) is equal to $\det P'(x)$. Thus, if the determinant is greater than 1, it must have an eigenvalue greater than 1 and the system is unstable. \square

Corollary IV.24. *If the expression in (4.6) has value less than 1 and the orbit, $\gamma(t)$, is unstable, the point x_0 under P must be a saddle-type instability.*

4.2.3 Example: Walking with Foot-slip

We will demonstrate the above theory by examining a model of a simple passive walker experiencing foot-slip [26]. In this model, the mass of the foot is denoted by




Figure 4.3: The dowel-pin represents the center of mass for the pendulum.

m_f and the hip by m_h . The length is given by ℓ . The angles of the leg are restricted to $\theta \in [-\delta, \delta]$ (when θ hits the boundary, $-\delta$, a new step is taken and θ is reset to δ). The (Cartesian) coordinates of the center of mass will be given by (x_c, y_c) and the coordinates of the foot are (x_f, y_f) . Note that we impose the constraint $y_f = 0$. We will also assume that the frictional force exerted on the foot is proportional to its velocity, that is, $F_f = -b\dot{x}_f$. This is “viscous” friction rather than Coulomb friction. Interestingly, there is some evidence that viscous friction may be more appropriate for legged locomotion than previously thought [112].

We will now provide a few parameter relationships:

$$m = m_h + m_f, \quad I = \frac{\ell^2 m_f m_h}{m}, \quad \ell_f = \ell \frac{m_h}{m},$$

$$x_c = x_f - \ell_f \sin \theta, \quad y_c = y_f + \ell_f \cos \theta,$$

where I is the moment of inertia about the center of mass and ℓ_f is the distance from the foot to the center of mass. Notice that the constraint implies that $y_c = \ell_f \cos \theta$. Due to the hybrid nature of this model, we will split the derivation of the equations of motion into two parts: the continuous part and the discrete part.

Remark IV.25. The way we have set up our model, θ is constrained to be in the interval $[-\delta, \delta]$. If θ crosses the negative boundary, we say a new step occurs and θ is reset to δ . If θ crosses the positive boundary, specifically if $\theta = \pi/2$, we say that a crash has occurred. In this case, the model stops walking and we report a failure. This also implies that falling forwards is not permitted; the only way to crash is by falling backwards.

The Continuous Part

The Lagrangian for a (planar) rigid body is

$$L = \frac{1}{2}I\dot{\theta}^2 + \frac{1}{2}m(\dot{x}_c^2 + \dot{y}_c^2) - V(\theta, x_c, y_c).$$

The potential energy depends on the elevation of the center of mass.

$$V = mg[\ell_f \cos(\alpha - \theta) - (x_c + \ell_f \sin \theta) \sin \alpha].$$

This system is constrained by $y_c = \ell_f \cos \theta$. The forces exerted from the friction on the foot are given by

$$F_{x_c} = -b\dot{x}_f = -b(\dot{x}_c + \ell_f \dot{\theta} \cos \theta)$$

$$F_{y_c} = 0$$

$$F_\theta = -b\dot{x}_f(\ell_f \cos \theta) = -b\ell_f \cos \theta (\dot{x}_c + \ell_f \dot{\theta} \cos \theta).$$

Solving the (constrained) Euler-Lagrange equations yields the system of

$$m\ddot{x}_c - mg \sin \alpha = -b(\dot{x}_c + \ell_f \dot{\theta} \cos \theta)$$

$$m\ddot{y}_c - \lambda = 0$$

$$I\ddot{\theta} - \ell_f \sin \theta (\lambda + mg \cos \alpha) = -b\ell_f \cos \theta (\dot{x}_c + \ell_f \dot{\theta} \cos \theta).$$

Using the constraint above, we get the value for the Lagrange multiplier.

$$\lambda = -m(\ell_f \ddot{\theta} \sin \theta - \ell_f \dot{\theta}^2 \cos \theta).$$

The final, reduced, equations of motion are

$$m\ddot{x}_c - mg \sin \alpha = -b (\dot{x}_c + \ell_f \dot{\theta} \cos \theta)$$

$$I\ddot{\theta} - m\ell_f \sin \theta [g \cos \alpha - \ell_f \ddot{\theta} \sin \theta + \ell_f \dot{\theta}^2 \cos \theta] = -b\ell_f \cos \theta (\dot{x}_c + \ell_f \dot{\theta} \cos \theta).$$

Rearranging so the accelerations are all on the left side yields

$$m\ddot{x}_c = mg \sin \alpha - b (\dot{x}_c + \ell_f \dot{\theta} \cos \theta)$$

$$(I + m\ell_f^2 \sin^2 \theta) \ddot{\theta} = \ell_f m \sin \theta [g \cos \alpha + \ell_f \dot{\theta}^2 \cos \theta]$$

$$- b\ell_f \cos \theta (\dot{x}_c + \ell_f \dot{\theta} \cos \theta).$$

This system has the regularity condition of $m(I + m\ell_f^2 \sin^2 \theta) \neq 0$ which is always satisfied. Since the equations of motion are invariant with respect to x_c , we will not deal with tracking it. We will relabel $(\dot{x}_c, \dot{\theta}) \mapsto (v, \omega)$ to create three first-order equations

$$\dot{\theta} = \omega$$

$$(4.7) \quad \dot{\omega} = \frac{\ell_f}{I + m\ell_f^2 \sin^2 \theta} [m \sin \theta (g \cos \alpha + \omega^2 \cos \theta) - b\ell_f \cos \theta (v + \ell_f \omega \cos \theta)]$$

$$\dot{v} = g \sin \alpha - \frac{1}{m} b (v + \ell_f \omega \cos \theta).$$

The Impact Part

Before we can derive an impact map, we first need to determine what S is. In this example, we assume that the leg takes symmetric steps of angle δ , i.e. $\theta \in [\delta, -\delta]$. When $\theta = -\delta$, the angle is reset to δ (corresponding to a new step taking place and the swing legs switching). Therefore, we will take the impact set, S , to be $\{\theta = -\delta\}$.

To determine Δ , we will assume a rigid hip, i.e. the horizontal velocity of the foot is continuous at impacts. i.e.,

$$(4.8) \quad \dot{x}_f^+ = \dot{x}_f^-.$$

To determine the reset map for $\dot{\theta}$, we use conservation of momentum about the foot (see [100] for the case where there is no foot slip)

$$(4.9) \quad \dot{\theta}^+ = \cos(2\delta)\dot{\theta}^-.$$

The equation (4.9) is in terms of the coordinates used while equation (4.8), however, is not. To get a relationship for \dot{x}_c , we note that $\dot{x}_f = \dot{x}_c + \ell_f \dot{\theta} \cos \theta$.

$$\dot{x}_c^- + \ell_f \dot{\theta}^- \cos \theta^- = \dot{x}_c^+ + \ell_f \dot{\theta}^+ \cos \theta^+$$

Rearranging this (and relabeling as in (4.7)) gives a final impact map of

$$(4.10) \quad \Delta(\theta, \omega, v) = (\theta + 2\delta, \cos(2\delta)\omega, v + \ell_f \cos \delta [1 - \cos(2\delta)]\omega)$$

Nondimensionalization

We offer the following rescaling to nondimensionalize the system (4.7), (4.10): $\omega_0 \xi(\tau) = \omega(t)$, $v_0 \eta(\tau) = v(t)$, and $t_0 \tau = t$ where

$$t_0 = \frac{m}{b}, \quad v_0 = \frac{mg}{b}, \quad \omega_0 = \frac{mg}{b\ell_f}.$$

For the remainder of this example, \dot{f} will correspond to differentiating against t and f' will mean differentiation against τ . The equations of motion become

$$(4.11) \quad \begin{aligned} \theta' &= \varepsilon \xi \\ \xi' &= C [\sin \theta (\cos \alpha + \varepsilon \xi^2 \cos \theta) - \cos \theta (\eta + \xi \cos \theta)] \\ \eta' &= \sin \alpha - \eta - \xi \cos \theta \end{aligned}$$

where

$$C := \frac{\mu}{1 + \mu \sin^2 \theta}.$$

The parameters, (μ, ε) , are given by

$$\mu = \frac{m\ell_f^2}{I} = \frac{m_h}{m_f}, \quad \varepsilon = \frac{m^2 g}{b^2 \ell_f}.$$

Making the appropriate changes to the impact map, we have

$$\Delta(\theta, \xi, \eta) = (\theta + 2\delta, \cos(2\delta)\xi, \eta + \cos \delta [1 - \cos(2\delta)] \xi)$$

and the impact surface, S , is unchanged because we do not rescale θ . Our model for foot slip now depends on 4 parameters: ε , μ , α , and δ : ε contains the information relating gravity and frictional forces, μ is the mass ratio, α is how steep the floor is, and δ is the width of each step.

Dynamics on the Poincaré Section

We present two numerical calculations: First, we calculate the region in *parameter space*, $(\mu, \varepsilon, \alpha, \delta)$, where a periodic orbit exists and we then calculate its largest eigenvalue via Theorem IV.16. The second calculation involves determining which *initial conditions* for fixed parameters result in a periodic orbit, i.e. the region of attraction for a periodic orbit. For a discussion on the numerical methods used, see Appendix A. The results for the eigenvalue calculations are shown in Figure 4.4 and the stable initial conditions are shown in Figure 4.5.

Figure 4.4 contains two interesting features. The first is that larger steps (that is, larger δ) result in faster exponential stability of periodic orbits. The strange aspect of this is the boundary of the stability region cannot be computed by calculating when the modulus of the maximum eigenvalue crosses one. The second interesting feature is that for small ε (that is, a large coefficient of friction, so a small amount of slip), small steps result in crashes. It is an ongoing goal to understand why this happens.

Figure 4.5 contains two interesting features as well. The first is the emergence of tongues where the first tongue lies on the boundary between the blue and red regions. The emergence of tongues on stability diagrams is a common occurrence in




Figure 4.4: The colored regions in the above figures show where a periodic orbit is found while the white is where a crash occurs. The left image shows ε small which is high friction while the right has large ε which corresponds to low friction.

periodically forced systems, similar to [95]. The second is that it seems that the dynamics map tongues to tongues. This seems to imply that if we can (approximately) calculate the boundary between the red and blue regions, we can (approximately) determine the location of every tongue. The next section attempts this task.

Approximating the First Tongue

We would like to find the boundary of the red region in Figure 4.5. We will approximate this in the following way. Let x^* be the fixed point of (4.11) that corresponds to the pendulum being vertical. That is,

$$x^* = [\alpha, 0, \sin \alpha]^T.$$

Assuming that the stable manifold, $W^s(x^*)$, for this fixed point divides the state-space such that on one side the pendulum crashes and the other corresponds to where a step can be completed, calculating $W^s(x^*) \cap \Delta(S)$ would offer all we need to know.




Figure 4.5: The blue region corresponds to initial conditions that lie within the basin of attraction for the periodic orbit. The red region contains initial conditions that lead to a crash. The black asterisks indicate the forward orbit under the Poincaré map. The green line is a linear guess for the boundary of the red region.

Using the notation $\dot{x} = f(x)$, we calculate $J_f(x^*)$, the Jacobian of the flow, to be

$$J_f(x^*) = \begin{bmatrix} 0 & \varepsilon & 0 \\ C & -a^2C & -aC \\ 0 & -a & -1 \end{bmatrix},$$

$$C = \frac{\mu}{1 + \mu \sin^2 \alpha}, \quad a = \cos \alpha.$$

Its monic characteristic polynomial is

$$p(\lambda) = \lambda^3 + (a^2C + 1)\lambda^2 - C\varepsilon\lambda - C\varepsilon = 0$$

This polynomial has three roots, $\lambda_{1,2,3}$. Descartes' rule of signs tell us that there can be at most one root with positive real part. There is, in fact, always a positive real root. This is because $p(0) = -C\varepsilon < 0$ and $\lim_{\lambda \rightarrow \infty} p(\lambda) = +\infty$.

Let $W^s(x^*)$ be the stable manifold for the fixed-point x^* and $W^u(x^*)$ be its unstable manifold. By the previous discussion, $\dim(W^s(x^*)) = 2$ and $\dim(W^u(x^*)) =$

1. By the stable manifold theorem (§2.7 of [92]), $W^s(x^*) \approx E^s + x^*$ where E^s is the stable subspace of $J_f(x^*)$.

We will assume that $W^s(x^*)$ is the plane spanned by the stable subspaces of $J_f(x^*)$ (which we know is always 2-dimensional). The equation for $E^s + x^*$ is (where v_1 and v_2 are two independent stable eigenvectors)

$$\langle v_1 \times v_2, x - x^* \rangle = 0$$

Making this into $(E^s + x^*) \cap \Delta(S)$ results in

$$\langle v_1 \times v_2, [\delta - \alpha, \xi_0, \eta_0 - \sin \alpha]^T \rangle = 0$$

Which can be rearranged to

$$\xi_0 = \frac{n_3}{n_2}(\sin \alpha - \eta_0) + \frac{n_1}{n_2}(\alpha - \delta)$$

where $n = v_1 \times v_2$ (and as long as $n_2 \neq 0$). This relationship gives the green line in Figure 4.5.

4.3 A Poincaré-Bendixson Theorem

While the previous section dealt with the question of stability of periodic orbits, this section outlines sufficient conditions such that periodic motion is the *only* asymptotic behavior. We accomplish this by proving a version of the famous Poincaré-Bendixson Theorem [10] for hybrid systems. The classic version of the Poincaré-Bendixson Theorem states that if a planar trajectory is bounded and its limit set does not contain any fixed points, then the limit set is a periodic orbit (cf. §3.7 in [92]).

There have been attempts to prove an analogous version for hybrid systems, particularly [76, 102]. However, the results in [76] are restricted to the situation of

constant vector fields while in [102], the authors consider a particular class of systems with very strict assumptions that are difficult to check.

In this section, we present a different version of a hybrid Poincaré-Bendixson Theorem as developed in [28, 29]. The version proved in [29] is for (simple) hybrid dynamical systems while [28] extends the result to general hybrid dynamical systems. Because the simple version is a special case of the general case, we present only the result from [28].

Theorem IV.26 ([28]). *Let $(\mathcal{N}, \mathcal{E}, \mathcal{X}, S, \Delta, f)$ be an essentially non-beating general hybrid dynamical system. In addition, suppose the following:*

(Q.1) *For every $i \in \mathcal{N}$, $n_i = 2$, that is $\mathcal{X}_i \subset \mathbb{R}^2$.*

(Q.2) *Each map Δ_e is injective.*

(Q.3) *There exists a compact, forward invariant set $F \subset \mathcal{X}$. That is, $F = \sqcup_i F_i$ where $F_i \subset \mathcal{X}_i$ and each F_i is compact.*

(Q.4) *Let the outdegree of vertex n be $\text{od}(n)$. Then, each F_n can be written as a disjoint union of $\text{od}(n)$ many compact sets,*

$$F_n = \bigsqcup_{\substack{e \in \mathcal{E} \\ \text{dom}(e)=n}} F_n^e.$$

(Q.5) *The set $F_n^e \cap S_{e'}$ is only nonempty if $e = e'$ and $F_n^e \cap S_e$ is diffeomorphic to an interval.*

(Q.6) *The vector fields, f_n , are transverse to $F_n^e \cap S_e$ and $\Delta_{e'}(F_i^{e'} \cap S_{e'})$ (when $\text{cod}(e') = n$).*

Assume F_i and $\text{fix}(f_i)$ are disjoint for all $i \in \mathcal{N}$. Then $\omega_{\mathcal{H}}(x)$ is (the closure of) a periodic orbit, for all $x \in F$.

Remark IV.27. The images of the impact surfaces, $\Delta_e(F_n^e \cap S_e)$, are not guaranteed to be smooth manifolds (cusps can appear). In this case, transversality does not make sense and the condition (Q.6) automatically fails.

Before we prove this theorem, we first show how conditions (Q.4) and (Q.5) are required to combat the added complexity of GHDSs over HDSs.

Example IV.28. Consider the directed graph with three nodes as illustrated in Figure 4.6. We choose all three state-spaces to be identical, $\mathcal{X}_i = [0, 1]^2$ and we




Figure 4.6: The graph of a HDS where chaos can still occur.

define the impact surfaces as follows:

$$S_{(1,3)} = \{(1, y) : y \in (1/2, 1]\},$$

$$S_{(1,2)} = \{(1, y) : y \in [0, 1/2]\},$$

$$S_{(3,1)} = \{(1, y) : y \in [0, 1]\},$$

$$S_{(2,1)} = \{(1, y) : y \in [0, 1]\}.$$

Their corresponding impact maps are:

$$\Delta_{(1,3)}(1, y) = (0, 2(1 - y)),$$

$$\Delta_{(1,2)}(1, y) = (0, 2y),$$

$$\Delta_{(3,1)}(1, y) = (0, y),$$

$$\Delta_{(2,1)}(1, y) = (0, y).$$

Finally, we define the three vector fields to simply be uniform motion to the right, i.e. for all three state-spaces, we let

$$\dot{x} = 1, \quad \dot{y} = 0.$$

Notice that if we denote the first return map of $(0, y) \in \mathcal{X}_1$ by P , then $P(0, y) = (0, T(y))$ where $T(y)$ is the tent map and is chaotic [32]. Chaos occurs here precisely because there is mixing between the nodes.

Notice that this example violates (Q.4) and (Q.5) for the following reason: Take $F_i = \mathcal{X}_i$ and by (Q.4) we must decompose $F_1 = F_1^{(1,2)} \cup F_1^{(1,3)}$ such that each $F_1^{(1,2)}$ and $F_1^{(1,3)}$ are compact. This requires the decomposition to be

$$F_1^{(1,2)} = [0, 1] \times [0, 1/2 + \alpha],$$

$$F_2^{(1,3)} = [0, 1] \times [1/2 - \beta, 1],$$

for some positive numbers α, β . However, this violates (Q.5) because

$$F_1^{(1,3)} \cap S_{(1,2)} = \{1\} \times [1/2 - \beta, 1/2 + \alpha] \neq \emptyset.$$

We can now proceed to the proof of Theorem IV.26, see Appendix B for some technical lemmas.

Proof of Theorem IV.26. The proof will proceed in three steps.

Step 1: *Suppose $(\mathcal{N}, \mathcal{E})$ is a cycle.* Then, the flow of our GHDS will look like the following:

$$F_1 \xrightarrow{\varphi_t^1} S_{(1,2)} \xrightarrow{\Delta_{(1,2)}} F_2 \dots F_n \xrightarrow{\varphi_t^n} S_{(n,1)} \xrightarrow{\Delta_{(n,1)}} F_1$$

This allows us to collapse the dynamics to a first-return map $P : U \rightarrow S_{(1,2)}$ for some open set $U \subset S_{(1,2)}$. Then, define the set

$$S_{(1,2)}^\infty := \{x \in S_{(1,2)} : \forall t > 0, \exists T > t \text{ s.t. } \varphi_T^H(x) \in S_{(1,2)}\}.$$

This allows us to iterate the map P . So we have a discrete dynamical system

$$P : S_{(1,2)}^\infty \rightarrow S_{(1,2)}^\infty$$

This leads to two possibilities for $x \in S_{(1,2)}$: either $x \notin S_{(1,2)}^\infty$ or $x \in S_{(1,2)}^\infty$. If $x \notin S_{(1,2)}^\infty$, then there exists a time, t^* , large enough, that for all $t > t^*$, $\varphi_t^{\mathcal{H}}(x) = \varphi_{t-t^*}^k \circ \varphi_{t^*}^{\mathcal{H}}(x)$ for some k and thus, $\omega_{\mathcal{H}}(x) = \omega_c(\varphi_{t^*}^{\mathcal{H}}(x))$. This set being a periodic orbit follows from the classical version of the Poincaré-Bendixson Theorem (see Theorem 1 in Chapter 3.7 of [92]). Now, assume that $x \in S_{(1,2)}^\infty \neq \emptyset$. The set $S_{(1,2)}^\infty \cong [a, b]$ by Lemma B.3 from the appendix and the map, P , is C^1 by Lemma B.4. Using Lemma B.1, this shows that $\omega_{\mathcal{H}}(x)$ is a periodic orbit.

Step 2: *Suppose all vertices have outdegree less than or equal to 1.* This allows us to define dynamics on our vertices. Let $i \in \mathcal{N}$. Define $\eta : \mathcal{N} \rightarrow \mathcal{N}$ as

$$\eta(i) = \begin{cases} i, & \text{if } \text{od}(i) = 0 \\ j, & \text{if } (i, j) \in \mathcal{E} \end{cases}$$

Then this collapses to Step 1 because finite dynamics always leads to a periodic orbit, see Lemma B.2.

Step 3: *The general case.* This allows us to split up the vertices with outdegree greater than one to multiple vertices with outdegree all being one. For an example of splitting up a vertex with outdegree 2, see Figure 4.7.

□

We end this section with a brief discussion on how to check the conditions for Theorem IV.26 in practice. The conditions (Q.1), (Q.2) and (Q.6) are straightforward to verify. The difficult part is finding a compact, forward invariant set that also satisfies (Q.4) and (Q.5). The purpose for the second half of (Q.5) is just to




Figure 4.7: An example of a node with outdegree 2. The set \mathcal{X}_i gets partitioned into two sets via the dashed line. Each of these now has outdegree 1.

disqualify Kronecker flows whereas (Q.4) with the first half of (Q.5) is necessary to extract out an invariant cycle. This is generally quite difficult, but the example in §4.3.1 provides a case where this is possible.

4.3.1 Example: The rimless wheel

The (uniform) rimless wheel is a commonly studied object in hybrid systems, [100]. We next present a variation of this, the case of the nonuniform rimless wheel, as a demonstration of Theorem IV.26. We derive a cascade of inequalities that guarantee the existence of a periodic orbit. Later we demonstrate two numerical examples of randomized wheels that satisfy these inequalities.

Model setup

Here, we explicitly state what $(\mathcal{N}, \mathcal{E}, \mathcal{X}, S, \Delta, f)$ is for the nonuniform rimless wheel. The application of Theorem IV.26 is demonstrated in the next subsection. The underlying directed graph, $(\mathcal{N}, \mathcal{E})$, is simply a cycle with length equal to the number of spokes of the wheel. That is, $\mathcal{N} = \{1, \dots, n\}$ and $\mathcal{E} = \{(i, i+1) : i = 1, \dots, n\}$ where n is the number of spokes and it is understood that all indices are taken modulo n . The remaining data, $(\mathcal{X}, S, \Delta, f)$, depends on the geometry of the wheel.

Consider the rimless wheel with different spoke lengths and different inter-spoke angles. For $i = 1, \dots, n$, let δ_i and ℓ_i be half the angle between spokes i and $i + 1$ and the spoke length of spoke i respectively. We will assume that the grade of the slope being walked down has a constant angle of α , see Figure 4.8. In what follows, x_1 will be the angle of the wheel from vertical (with respect to the spoke in contact with the ground) and $x_2 = \dot{x}_1$ is the angular velocity. Each vector field, f_i , is that of an inverted pendulum of length ℓ_i ,

$$\dot{x} = f_i(x) = \begin{bmatrix} x_2 \\ \zeta_i \sin(x_1) \end{bmatrix},$$

where $\zeta_i = g/\ell_i$ and g is the acceleration due to gravity. The impact surface, $S_{(i,i+1)}$, is given by configurations where both legs i and $i + 1$ are in contact with the ground. That is,

$$S_{(i,i+1)} = \{(x_1, x_2) : x_1 = \pi/2 - \xi_i - 2\delta_i - \alpha\}.$$

Their corresponding impact maps, $\Delta_{(i,i+1)} : S_{(i,i+1)} \rightarrow \mathcal{X}_{i+1}$, are given by resetting the angle and projecting the angular velocity (see [100] for a discussion for the uniform rimless wheel case).

$$\Delta_{(i,i+1)}(x) = \begin{bmatrix} \pi/2 - \xi_i - \alpha \\ r_{i,i+1} \cos(2\delta_i) x_2 \end{bmatrix},$$

where $r_{i,i+1}$, ξ_i , and s_i are defined as (recall that all indices are taken modulo n)

$$\begin{aligned} r_{i,i+1} &= \ell_i/\ell_{i+1}, \\ s_i &= \sqrt{\ell_i^2 + \ell_{i+1}^2 - 2\ell_i\ell_{i+1}\cos(2\delta_i)}, \\ \xi_i &= \sin^{-1}\left(\sin(2\delta_i)\left(\frac{\ell_i}{s_i}\right)\right). \end{aligned}$$

The final pieces of information needed are the state-spaces, \mathcal{X}_i . These sets consist of angles where spoke i is the only spoke making contact with the ground,

$$(4.12) \quad \mathcal{X}_i = \{(x_1, x_2) : \pi/2 - \xi_{i-1} - \alpha \leq x_1 \leq \pi/2 - \xi_i - 2\delta_i - \alpha\}.$$

Formally, we need each \mathcal{X}_i to be open so we only need \mathcal{X}_i to contain each closed set specified in (4.12).

Remark IV.29. Even though the wheel continues to rotate as it descends the ramp, the angle x_1 will never go beyond $\pm\pi$. This is because each impact resets the angle to measure the current spoke.




Figure 4.8: Drawing of when both legs l_i and l_{i+1} are on the ground.

Existence of a periodic limit cycle

With the data, $(\mathcal{N}, \mathcal{E}, \mathcal{X}, S, \Delta, f)$, specified, we now examine the existence of periodic orbits. To use Theorem IV.26, we need to make sure (Q.1)-(Q.6) are satisfied. Conditions (Q.1) and (Q.2) are clearly met. For the meantime, let us assume that (Q.3) holds and that each F_i does not contain the origin. In fact, as we will show later, each F_i will be a rectangle which guarantees (Q.5) since each $S_{(i,i+1)}$ is a line. (Q.4) holds because the outdegrees of all the nodes are 1 (the directed graph for this system is a cycle). Finally, (Q.6) holds because the only place where transversality fails is at the origin - i.e. no motion is occurring. The only condition that needs to be checked is (Q.3): we need a compact, forward invariant set such that each component is rectangular and does not contain the identity. This will be done in the

same fashion as for the simple case, [29]. We will look at the potential energy gained over a cycle and compare it to the kinetic lost at all the impacts. For consistency, we will compare all energies at the beginning of a step, that is right after impacts occur. We will also set the reference point for the gravitational potential energy to be zero at the moment before impact occurs.

In order to show the existence of F , we will show that the energy of the system is bounded from above and below (when a system of inequalities are satisfied). First of all, the energy of the system is bounded from above because the amount of energy gained over each step is constant while the energy lost is linear, see equations (4.14) and (4.15). To obtain a lower bound, we do the following: Let K_i measure the kinetic energy lost between impacts. That is, let E^- be the energy in the system before impacting leg $i+1$. Then, $K_i(E^-) = E^+$ is the energy of the system immediately after impact occurs. Let us define two more quantities, P_i and R_i . P_i is the (constant) amount of potential energy accrued over the step between legs i and $i+1$. R_i is the minimum required energy at the beginning of a step to successfully make it to the next step. The requirement for a periodic orbit to exist is the satisfaction of the cascade of inequalities (4.13) (first, relabel so R_1 is the largest of all of the R_i 's). Define $\eta_i(E) := K_i(E + P_i)$.

$$\begin{aligned}
 & \eta_1(R_1) > R_2, \\
 & \eta_2 \circ \eta_1(R_1) > R_3, \\
 & \vdots \\
 & \eta_n \circ \dots \circ \eta_1(R_1) > R_{n+1} = R_1.
 \end{aligned}
 \tag{4.13}$$

For explicit formulas to compute η_i , we have the following:

$$R_i = mgl_i \left(1 - \cos \left(\frac{\pi}{2} - \alpha - \xi_{i-1} \right) \right),$$

$$(4.14) \quad P_i = mgl_i \left(\cos \left(\frac{\pi}{2} - \alpha - \xi_{i-1} \right) - \cos \left(\frac{\pi}{2} - \xi_i - 2\delta_i - \alpha \right) \right),$$

$$(4.15) \quad K_i(E) = \cos^2(2\delta_i) \cdot E.$$

Numerical results

For the remainder, we will take $m = g = 1$. Additionally, because we are modeling a wheel, we will add two more assumptions to the parameters. The first is

$$(4.16) \quad \sum_{i=1}^n \delta_i = \pi,$$

so the spokes line up to make a whole wheel. We also add the following spoke length requirements to ensure that no spokes are skipped while the wheel is rolling.

$$(4.17) \quad \ell_i > \frac{\ell_{i-1}\ell_{i+1} \sin(2\delta_{i-1} + 2\delta_i)}{\ell_{i-1} \sin(2\delta_{i-1}) + \ell_{i+1} \sin(2\delta_i)}.$$

Below we present two examples, one with 5 spokes and the other with 10. All parameters were randomly chosen, but do satisfy the inequalities (4.13) and the additional constraints: (4.16) and (4.17). Additionally, for the plots, Figures 4.9 and 4.10, instead of resetting θ after each impact, we track it around the whole cycle.

Example IV.30. Let $n = 5$. A set of admissible parameters are shown in Table 4.1.

The steady-state trajectory (over two revolutions) is given in Figure 4.9.

i	δ_i	ℓ_i	s_i	ξ_i
1	0.6516	1.7184	2.1981	0.8540
2	0.7124	1.8985	2.3220	0.9423
3	0.6410	1.6416	1.9313	0.9524
4	0.5639	1.5872	1.7153	0.9899
5	0.5728	1.6217	1.8121	0.9531

Table 4.1: A random set of admissible parameters for the 5-spoked nonuniform rimless wheel with $\alpha = 0.5062$.

Example IV.31. Let $n = 10$. A set of admissible parameters are shown in Table

4.2. The steady-state trajectory (over one revolution) is given in Figure 4.10.

i	δ_i	ℓ_i	s_i	ξ_i
1	0.3204	1.7662	1.0572	1.5234
2	0.2411	1.4658	0.7531	1.1261
3	0.1766	1.6226	0.5617	1.5354
4	0.3762	1.5025	1.1505	1.1029
5	0.3389	1.6158	1.1810	1.0311
6	0.3622	1.8655	1.2563	1.3920
7	0.2944	1.6204	1.0635	1.0088
8	0.3648	1.9142	1.2942	1.4033
9	0.3352	1.6427	1.1785	1.0472
10	0.3316	1.8764	1.1906	1.3268

Table 4.2: A random set of admissible parameters for the 10-spoked nonuniform rimless wheel with $\alpha = 0.1752$.




Figure 4.9: Both figures are for a 5 spoke wheel. Left: The plot of θ against time. Right: The plot of $\dot{\theta}$ against time.




Figure 4.10: Both figures are for a 10 spoke wheel. Left: The plot of θ against time. Right: The plot of $\dot{\theta}$ against time.

CHAPTER V

Invariant Measures in Nonholonomic Systems

An invariant measure is a powerful tool to understand the asymptotic nature of a dynamical system. In the case of nonholonomic systems, a smooth invariant measure offers two key insights. The first is the usual case in dynamical systems where an invariant measure allows for the use of the Birkhoff Ergodic Theorem (cf. e.g. 4.1.2 in [70]) as well as for recurrence. The other is unique to nonholonomic systems; even though nonholonomic systems are *not* Hamiltonian, “nonholonomic systems which do preserve volume are in a quantifiable sense closer to Hamiltonian systems than their volume changing counterparts,” [45]. Therefore, being able to find an invariant measure for a nonholonomic system allows for ergodic-like understanding of its asymptotic behavior as well as provide a way to “Hamiltonize” a nonholonomic system.

There has already been work done in finding invariant measures in systems where symmetries are present: Chaplygin systems are studied in, e.g. [64, 72, 80, 83], Euler-Poincaré-Suslov systems are studied in, e.g. [16, 66], and systems with internal degrees of freedom are studied in, e.g. [15, 16, 113]. This chapter, rather, uses an all-together different approach where *no* symmetries will be used. Additionally, in §5.7, we provide necessary and sufficient conditions for when an invariant measure exists

whose density depends only on the base variables, i.e. $f = \pi_Q^* g$ for some $g \in C^\infty(Q)$.

5.1 Global Nonholonomic Vector Fields

Given a constraint distribution, $\mathcal{D}^* \subset T^*Q$, we can determine the nonholonomic vector field $X_H^{\mathcal{D}} \in \mathfrak{X}(\mathcal{D}^*)$. Commonly local, noncanonical, coordinates are chosen for \mathcal{D}^* , §5.8 in [16] and [101]. However, we will instead work with the entire manifold T^*Q and define a *global* vector field $X_H^{global} \in \mathfrak{X}(T^*Q)$ such that $X_H^{global}|_{\mathcal{D}^*} = X_H^{\mathcal{D}}$. This section outlines an intrinsic way to determine such a vector field.

Definition V.1. For a given constraint submanifold $\mathcal{D}^* \subset T^*Q$ (\mathcal{D}^* need not be a distribution), a *realization* of \mathcal{D}^* is an ordered collection of functions $\mathcal{C} := \{g_i : T^*Q \rightarrow \mathbb{R}\}$ such that zero is a regular value of $G = g_1 \times \dots \times g_m$ and

$$\mathcal{D}^* = \bigcap_i g_i^{-1}(0).$$

If the functions g_i are given by momenta, i.e. $g_i = P(X^i)$ (cf. Corollary III.21), then the realization is called *natural*.

Remark V.2. Under the case where the Lagrangian is natural (which provides a Riemannian metric on Q) and the constraint submanifold is a distribution, we can choose the realization to be natural:

$$\mathcal{C} = \{P(W^1), \dots, P(W^m)\},$$

where $W^i = \mathbb{F}L^{-1}\eta^i = (\eta^i)^\sharp$.

By replacing \mathcal{D}^* with a realization \mathcal{C} , we can extend the nonholonomic vector field to a vector field on T^*Q that preserves the constraining functions g_i . Recall that the form of the nonholonomic vector field is $i_{X_H^{\mathcal{D}}}\omega = dH + \lambda_\alpha \pi_Q^* \eta^\alpha$. We construct the global nonholonomic vector field, $\Xi_H^{\mathcal{C}}$, by requiring that:

(NH.1) $i_{\Xi_H^{\mathcal{C}}}\omega = dH + \lambda_\alpha \pi_Q^* \eta^\alpha$ for smooth functions $\lambda_\alpha : T^*Q \rightarrow \mathbb{R}$, and

(NH.2) $\mathcal{L}_{\Xi_H^{\mathcal{C}}}g_i = 0$ for all $g_i \in \mathcal{C}$.

Under reasonable compatibility assumptions on \mathcal{C} (cf. §3.4.1 in [80]), such a vector field exists and is unique. However, given two different realizations, \mathcal{C} and \mathcal{C}' , of the same constraints \mathcal{D}^* , it is not generally true that $\Xi_H^{\mathcal{C}} = \Xi_H^{\mathcal{C}'}$, however $\Xi_H^{\mathcal{C}}|_{\mathcal{D}^*} = \Xi_H^{\mathcal{C}'}|_{\mathcal{D}^*}$. When both the Hamiltonian and realization are natural, the global field can be explicitly computed via the constraint mass matrix.

Remark V.3. The constraint manifold is given by the joint zero level-sets of the g_i while the realization provides additional irrelevant information off of the constraint manifold. This is why $\Xi_H^{\mathcal{C}} \neq \Xi_H^{\mathcal{C}'}$ but they agree once restricted.

Definition V.4. For a natural realization $\mathcal{C} = \{P(W^1), \dots, P(W^m)\}$ and natural Hamiltonian (so (Q, g) is Riemannian), the *constraint mass matrix*, $(m^{\alpha\beta})$, is given by orthogonally pairing the constraints, i.e.

$$m^{\alpha\beta} = g(W^\alpha, W^\beta) = \eta^\alpha(W^\beta).$$

Additionally, its inverse will be denoted by $(m_{\alpha\beta}) = (m^{\alpha\beta})^{-1}$.

Lemma V.5. *The constraint mass matrix is symmetric and positive-definite so long as all the constraints are linearly independent.*

Proof. This follows from the fact that $(m^{\alpha\beta})$ is a Gram matrix for a nondegenerate inner product. □

We can now write down a formula for $\Xi_H^{\mathcal{C}}$. Using (NH.1) and (NH.2), we get that

$$\begin{aligned}
\mathcal{L}_{\Xi_H^{\mathcal{C}}}P(W^\beta) &= i_{X_{P(W^\beta)}}\omega(\Xi_H^{\mathcal{C}}) \\
&= -i_{\Xi_H^{\mathcal{C}}}\omega(X_{P(W^\beta)}) \\
&= -dH(X_{P(W^\beta)}) - \lambda_\alpha \pi_Q^* \eta^\alpha(X_{P(W^\beta)}) \\
&= \{P(W^\beta), H\} - \lambda_\alpha \eta^\alpha(W^\beta) = 0 \\
&\implies \{P(W^\beta), H\} = m^{\alpha\beta} \lambda_\alpha.
\end{aligned}$$

Due to the constraint mass matrix being nondegenerate, the multipliers have a unique solution and the global nonholonomic vector field is given by

$$(5.1) \quad \boxed{i_{\Xi_H^{\mathcal{C}}}\omega = dH - m_{\alpha\beta} \{H, P(W^\alpha)\} \pi_Q^* \eta^\beta}$$

Remark V.6. The global nonholonomic vector field given by (5.1) can be extended to the case of nonlinear constraints via Chetaev's rule, which will give equivalent results to those in [37] where the ‘‘almost-tangent’’ structure of the tangent bundle is utilized. For Lagrangian systems, Chetaev's rule states that if we have a nonlinear constraint $f(q, \dot{q}) = 0$, then the constraint force takes the following form:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}} \right) - \frac{\partial L}{\partial q} = \lambda \cdot \frac{\partial f}{\partial \dot{q}} = \lambda \cdot \mathcal{S}^* df,$$

and $\mathcal{S} : T(TQ) \rightarrow T(TQ)$ is its almost-tangent structure. However as we are instead on the cotangent bundle, the object we will use will be the related to the almost-tangent structure through the fiber derivative,

$$\mathcal{C} : T(T^*Q) \rightarrow T(T^*Q)$$

$$\mathcal{C}^* (\alpha_i dx^i + \beta^j dp_j) = g_{ij} \beta^j dx^i.$$

For a constraint realization $\mathcal{C} = \{g^1, \dots, g^m\}$, the nonholonomic vector field is given by

$$i_{\Xi_H^{\mathcal{C}}}\omega = dH - m_{\alpha\beta} \{H, g^\alpha\} \mathcal{C}^* dg^\beta,$$

where the multipliers are given by $m^{\alpha\beta} = \mathcal{C}^* dg^\beta (X_{g^\alpha})$.

Definition V.7. The 1-form given by

$$\nu_H^{\mathcal{C}} := dH - m_{\alpha\beta} \{H, P(W^\alpha)\} \pi_Q^* \eta^\beta,$$

is called the *nonholonomic 1-form* with respect to the realization \mathcal{C} .

Proposition V.8. *Given two different natural realizations, \mathcal{C} and \mathcal{C}' , the global nonholonomic vector fields given by (5.1) agree on \mathcal{D}^* .*

Proof. Suppose that there is only a single constraint and $\mathcal{C} = \{P(W)\}$ and $\mathcal{C}' = \{fP(W)\}$ for some smooth f . By Leibniz's rule,

$$\begin{aligned} i_{\Xi_H^{\mathcal{C}'}} \omega &= dH + \frac{1}{f^2 g(W, W)} \{fP(W), H\} f \pi_Q^* \eta \\ &= dH + \frac{1}{f^2 g(W, W)} [f \{P(W), H\} + P(W) \{f, H\}] f \pi_Q^* \eta \\ &= dH + \frac{1}{g(W, W)} \{P(W), H\} \pi_Q^* \eta + \frac{P(W)}{fg(W, W)} \{f, H\} \pi_Q^* \eta. \end{aligned}$$

Therefore, we have

$$i_{\Xi_H^{\mathcal{C}'}} \omega - i_{\Xi_H^{\mathcal{C}}} \omega = \frac{P(W)}{fg(W, W)} \{f, H\} \pi_Q^* \eta,$$

which vanishes on \mathcal{D}^* . A similar argument works for multiple constraints. \square

Throughout the rest of this chapter, we will assume that \mathcal{C} is a natural realization.

This, in turn, requires that the constraints are linear in the velocities / momenta.

5.2 Symmetries in Nonholonomic Systems

Before we introduce the nonholonomic volume form to check for measure preservation, we take an aside to demonstrate how the global form (5.1) provides a clean proof of the momentum equation, Theorem III.34. We first show a version of the momentum equation where the Lie algebra element is constant and not assumed to be in $\mathfrak{g}^{\mathcal{D}}$.

Proposition V.9. *Let G act on Q by diffeomorphisms and lift its action to a symplectic action on $M = T^*Q$. Then, if $H : T^*Q \rightarrow \mathbb{R}$ is a Hamiltonian invariant under the G -action, for all $\xi \in \mathfrak{g}$ we have*

$$\dot{\hat{J}}(\xi) = \mathcal{L}_{\Xi_H^{\mathcal{G}}} \hat{J}(\xi) = m_{\alpha\beta} \{H, P(W^\beta)\} \langle J(\eta^\alpha), \xi \rangle.$$

Proof. We first notice that

$$\begin{aligned} \mathcal{L}_{\Xi_H^{\mathcal{G}}} \hat{J}(\xi) &= d\hat{J}(\xi)(\Xi_H^{\mathcal{G}}) \\ &= i_{\Xi_H^{\mathcal{G}}} i_{\xi_M} \omega = -\nu_H^{\mathcal{G}}(\xi_M). \end{aligned}$$

Because the Hamiltonian is G -invariant, infinitesimally this states that $dH(\xi_M) = 0$ and therefore the above simplifies to

$$\dot{\hat{J}}(\xi) = m_{\alpha\beta} \{H, P(W^\beta)\} (\pi_Q^* \eta^\alpha)(\xi_M).$$

To finish the proof, we see that $\pi_Q^* \eta^\alpha(\xi_M) = \eta^\alpha(\xi_Q)$. The result follows from applying equation (3.9) or (3.14). \square

We can now proceed with the proof of Theorem III.34.

Proof of Theorem III.34. The product rule gives us

$$\frac{d}{dt} \hat{J}^{nh}(\xi^{q(t)}) = m_{\alpha\beta} \{H, P(W^\beta)\} \langle J(\eta^\alpha), \xi^{q(t)} \rangle + P \left(\left[\frac{d}{dt} \xi^{q(t)} \right]_Q \right).$$

In order to prove the theorem, we will show that $\langle J(\eta^\alpha), \xi^{q(t)} \rangle$ vanishes. This follows directly from the following calculation:

$$\pi_Q^* \eta^\alpha(\xi_M) = \eta^\alpha(\xi_Q) = g(W^\alpha, \xi_Q).$$

Now, $g(W^\alpha(q(t)), \xi_Q^{q(t)}) = 0$ by definition of $\mathfrak{g}^{\mathcal{D}}$. \square

5.3 Nonholonomic Volume Form

The symplectic manifold T^*Q has a canonical volume form ω^n . However, the non-holonomic flow takes place on a submanifold $\mathcal{D}^* \subset T^*Q$ which is $2n - m$ dimensional. Therefore, ω^n is not a volume form on \mathcal{D}^* . Here, we construct a volume form on \mathcal{D}^* which is unique up to the choice of realization. The derivation of this will be similar to the construction of the volume form on an energy surface in §3.4 of [1]. For the realization $\mathcal{C} = \{P(W^1), \dots, P(W^m)\}$, define the m -form

$$\sigma_{\mathcal{C}} := dP(W^1) \wedge \dots \wedge dP(W^m).$$

Definition V.10. If we denote the inclusion map by $\iota : \mathcal{D}^* \hookrightarrow T^*Q$, then a nonholonomic volume, $\mu_{\mathcal{C}}$, is given by

$$\mu_{\mathcal{C}} = \iota^* \varepsilon, \quad \sigma_{\mathcal{C}} \wedge \varepsilon = \omega^n.$$

Proposition V.11. *Given an ordered collection of constraints, \mathcal{C} , the induced volume form $\mu_{\mathcal{C}}$ is unique.*

Proof. Suppose that ε and ε' are two forms satisfying $\sigma_{\mathcal{C}} \wedge \varepsilon = \omega^n$. Then

$$\varepsilon - \varepsilon' = \alpha, \quad \sigma_{\mathcal{C}} \wedge \alpha = 0.$$

Now let $\iota : \mathcal{D}^* \hookrightarrow T^*Q$ be the inclusion. Then from the above, we see that

$$\iota^* \eta = \iota^* \eta' + \iota^* \alpha.$$

The result will follow so long as $\iota^* \alpha = 0$. Suppose that $\iota^* \alpha \neq 0$ and choose vectors $v^1, \dots, v^{2n-m} \in T_q \mathcal{D}^* \subset T_q T^*Q$ such that $\alpha(v^1, \dots, v^{2n-m}) \neq 0$. Complete this collection of vectors to a basis of $T_q T^*Q$: $v^1, \dots, v^{2n-m}, v^{2n-m+1}, \dots, v^{2n}$ such that $\sigma_{\mathcal{C}}(v^{2n-m+1}, \dots, v^{2n}) \neq 0$. Then we have

$$\sigma_{\mathcal{C}} \wedge \alpha(v^1, \dots, v^{2n}) = (-1)^{(2n-m)m} \alpha(v^1, \dots, v^{2n-m}) \cdot \sigma_{\mathcal{C}}(v^{2n-m+1}, \dots, v^{2n}) \neq 0,$$

which is a contradiction. □

Remark V.12. Notice that for an ordered collection of constraints the volume form is unique. However, changing the order of the constraints changes the sign of the induced volume form and rescaling constraints rescales the volume form. In this sense, \mathcal{C} uniquely determines $\mu_{\mathcal{C}}$, but \mathcal{D}^* only determines $\mu_{\mathcal{C}}$ up to a multiple.

While examining the failure of Liouville's theorem (Proposition III.14) for non-holonomic systems, we will see when $\mu_{\mathcal{C}}$ is preserved under the flow of $X_H^{\mathcal{D}}$. More generally, we will consider the existence of a smooth density $f \in C^\infty(\mathcal{D}^*)$ when $f\mu_{\mathcal{C}}$ is preserved.

5.4 Divergence of a Nonholonomic System

Let $\omega = dq^i \wedge dp_i$ be the standard symplectic form on T^*Q . This in turn induces a volume form ω^n . It is a known result that Hamiltonian flows preserve this measure, however, nonholonomic flows generally do not.

5.4.1 Divergence Preliminaries

To understand volume preservation, we will use the notion of the divergence of a vector field, see §2.5 of [1] or §5.1 in [70].

Definition V.13. Let M be an orientable manifold with volume form Ω and X a vector field on M . Then the unique function $\operatorname{div}_\Omega(X) \in C^\infty(M)$ such that $\mathcal{L}_X\Omega = \operatorname{div}_\Omega(X)\Omega$ is called the divergence of X . The vector field X is called incompressible iff $\operatorname{div}_\Omega(X) = 0$.

This definition of divergence generalizes the familiar one from multivariate calcu-

lus in which $M = \mathbb{R}^n$ and $\Omega = dx^1 \wedge \dots \wedge dx^n$. Indeed,

$$\begin{aligned} \mathcal{L}_X \Omega &= di_X \Omega = d \left[\sum_{i=1}^n (-1)^{i-1} X^i \cdot dx^1 \wedge \dots \wedge dx^{i-1} \wedge dx^{i+1} \wedge \dots \wedge dx^n \right] \\ &= \left[\sum_{i=1}^n \frac{\partial X^i}{\partial x^i} \right] \cdot \Omega. \end{aligned}$$

Studying the divergence is a useful test to check volume preservation via the following proposition.

Proposition V.14 (2.5.25 in [1]). *Let M be a manifold with volume Ω and vector field X . Then X is incompressible iff every flow box of X is volume preserving.*

Liouville's theorem in this language states that for an unconstrained Hamiltonian system, $\operatorname{div}_{\omega^n}(X_H) = 0$. That is, Hamiltonian systems preserve the volume induced by the symplectic form. This is, in general, not the case for nonholonomic systems. The remainder of this section deals with computing $\operatorname{div}_{\mu_{\mathcal{C}}}(X_H^{\mathcal{D}})$ while the case of finding an a density f such that $\operatorname{div}_{f\mu_{\mathcal{C}}}(X_H^{\mathcal{D}}) = 0$ is postponed until §5.6.

Before we begin computing divergences, we first present a helpful lemma.

Lemma V.15. *If \mathcal{C} is a (natural) realization of a constraint $\mathcal{D}^* \subset T^*Q$, then*

$$(5.2) \quad \operatorname{div}_{\omega^n} \left(\Xi_H^{\mathcal{C}} \right) \Big|_{\mathcal{D}^*} = \operatorname{div}_{\mu_{\mathcal{C}}} \left(X_H^{\mathcal{D}} \right).$$

Proof. Leibniz's rule for the Lie derivative provides

$$\begin{aligned} \mathcal{L}_{\Xi_H^{\mathcal{C}}} \omega^n &= \mathcal{L}_{\Xi_H^{\mathcal{C}}} (\sigma_{\mathcal{C}} \wedge \varepsilon) \\ &= \left(\mathcal{L}_{\Xi_H^{\mathcal{C}}} \sigma_{\mathcal{C}} \right) \wedge \varepsilon + \sigma_{\mathcal{C}} \wedge \left(\mathcal{L}_{\Xi_H^{\mathcal{C}}} \varepsilon \right). \end{aligned}$$

However, $\mathcal{L}_{\Xi_H^{\mathcal{C}}} \sigma_{\mathcal{C}} = 0$ because the constraints are preserved under the flow. Applying this, we see that

$$\mathcal{L}_{\Xi_H^{\mathcal{C}}} \omega^n = \sigma_{\mathcal{C}} \wedge \left(\mathcal{L}_{\Xi_H^{\mathcal{C}}} \varepsilon \right),$$

which gives

$$\left(\operatorname{div}_{\omega^n} \left(\Xi_H^{\mathcal{C}} \right) \right) \sigma_{\mathcal{C}} \wedge \varepsilon = \sigma_{\mathcal{C}} \wedge \left(\mathcal{L}_{\Xi_H^{\mathcal{C}}} \varepsilon \right).$$

Due to the fact that the Lie derivative commutes with restriction, the result follows. \square

5.4.2 One Constraint

Before we compute the divergence of arbitrary nonholonomic systems, we first consider the simplified case where there is only a single constraint present, i.e. $\mathcal{C} = \{P(W)\}$. Here, we make the normalization $\eta(W) = 1$ to simplify equation (5.1). The divergence of $X_H^{\mathcal{D}}$ is given by

$$\mathcal{L}_{X_H^{\mathcal{D}}} \mu_{\mathcal{C}} = \operatorname{div}_{\mu_{\mathcal{C}}}(X_H^{\mathcal{D}}) \mu_{\mathcal{C}}.$$

In order to compute this, we will invoke Cartan's magic formula as well as Lemma V.15 (restricting to \mathcal{D}^* will occur at the end):

$$\begin{aligned} \mathcal{L}_{\Xi_H^{\mathcal{C}}}(\omega^n) &= i_{\Xi_H^{\mathcal{C}}} d\omega^n + di_{\Xi_H^{\mathcal{C}}}\omega^n \\ &= n \cdot d\left(i_{\Xi_H^{\mathcal{C}}}\omega \wedge \omega^{n-1}\right) \\ &= n \cdot d\left(i_{\Xi_H^{\mathcal{C}}}\omega\right) \wedge \omega^{n-1} - n \cdot \left(i_{\Xi_H^{\mathcal{C}}}\omega\right) \wedge d\omega^{n-1} \\ &= n \cdot \left(di_{\Xi_H^{\mathcal{C}}}\omega\right) \wedge \omega^{n-1}. \end{aligned}$$

The problem of computing the divergence collapses to calculating $di_{\Xi_H^{\mathcal{C}}}\omega$ (which captures how “non-symplectic” the flow is). Let N be difference between the non-holonomic and Hamiltonian vector fields:

$$N = \{H, P(W)\} \eta_k \frac{\partial}{\partial p_k}.$$

Then, from Hamilton's equations, we obtain

$$i_{\Xi_H^{\mathcal{C}}}\omega = dH + i_N\omega.$$

Returning to the divergence calculation, $di_{\overline{H}}\omega = di_N\omega$ where $i_N\omega = -\eta_i \{H, P(W)\} dq^i$.

Applying the exterior derivative yields:

$$di_N\omega = \left[-\frac{\partial\eta_i}{\partial q^k} \{H, P(W)\} - \eta_i \frac{\partial}{\partial q^k} \{H, P(W)\} \right] dq^k \wedge dq^i \\ - \eta_i \frac{\partial}{\partial p_\ell} \{H, P(W)\} dp_\ell \wedge dq^i$$

Notice that when we wedge $di_N\omega$ with ω^{n-1} , the entire first line vanishes and only the diagonal on the second survives. Combining everything, we see that

$$(5.3) \quad \operatorname{div}_{\mu_\ell} (X_H^{\mathcal{D}}) = n \cdot \eta_i \frac{\partial}{\partial p_i} \{H, P(W)\}.$$

5.4.3 Multiple constraints

The exact same procedure can be carried out when there are an arbitrary number of constraints. The divergence is then simply

$$(5.4) \quad \operatorname{div}_{\mu_\ell} (X_H^{\mathcal{D}}) = n \cdot m_{\alpha\beta} \cdot \eta_k^\alpha \frac{\partial}{\partial p_k} \{H, P(W^\beta)\}.$$

5.4.4 “Curvature”

This section concludes with an intrinsic way to interpret (5.3) and (5.4). Recall the cotangent projection $\pi_Q : T^*Q \rightarrow Q$ and the fact that

$$d\pi_Q \cdot X_f = \frac{\partial f}{\partial p_i} \frac{\partial}{\partial q^i}.$$

The divergence formula (5.3) becomes:

$$(5.5) \quad \operatorname{div}_{\mu_\ell} (X_H^{\mathcal{D}}) = n \cdot \pi_Q^* \eta \left(X_{\{H, P(W)\}} \right) \\ = -n \cdot \pi_Q^* \eta \left([X_H, X_{P(W)}] \right).$$

This can also be carried over to the multiple constraint case.

$$(5.6) \quad \boxed{\operatorname{div}_{\mu_\ell} (X_H^{\mathcal{D}}) = -n \cdot m_{\alpha\beta} \cdot \pi_Q^* \eta^\alpha \left([X_H, X_{P(W^\beta)}] \right)}.$$

The formulas (5.5) and (5.6) have a structure similar to the curvature of an Ehresmann connection. This is because these formulas have the structure of a projection composed with a vector field bracket. The main difference is that while the curvature of an Ehresmann connection is vertical-valued, these formulas are real-valued. We proceed with computing the divergence of two examples: the Chaplygin sleigh and the vertical rolling disk.

Remark V.16. In the same way that the nonholonomic 1-form can be extended to the case of nonlinear constraints via Chetaev's rule, see Remark V.6, the divergence described above by (5.6) can also be extended to the case of nonlinear constraints. The divergence is given by

$$\operatorname{div}_{\mu^{\mathcal{C}}} (X_H^{\mathcal{D}}) = -n \cdot m_{\alpha\beta} \cdot \mathcal{C}^* dg^\alpha \left([X_H, X_{g^\beta}] \right).$$

5.5 Examples

In this section, we pause from theory and compute the divergence for two commonly studied nonholonomic systems: the Chaplygin sleigh and the vertical rolling disk.

5.5.1 Chaplygin Sleigh

The Chaplygin sleigh is a nonholonomic system on the configuration space $Q = \operatorname{SE}_2 \cong \operatorname{SO}_2 \times \mathbb{R}^2$, the special Euclidean group in two dimensions, and has the following Lagrangian,

$$L = \frac{1}{2} \left(m\dot{x}^2 + m\dot{y}^2 + (I + ma^2)\dot{\theta}^2 - 2ma\dot{x}\dot{\theta} \sin \theta + 2ma\dot{y}\dot{\theta} \cos \theta \right),$$

where $(x, y) \in \mathbb{R}^2$ is the contact point of the knife edge, $\theta \in \operatorname{SO}_2$ is its orientation, m is the mass of the sleigh, I is the moment of inertia about the center of mass and a is the distance from the center of mass to the contact point; see Figure 5.1.




Figure 5.1: The Chaplygin sleigh is a rigid body attached to a sliding knife edge, §1.7 in [16].

The requirement that the knife edge can only slide in the direction it is pointing (no orthogonal sliding is allowed) is given by the constraint

$$\dot{y} \cos \theta - \dot{x} \sin \theta = 0.$$

Which corresponds to the 1-form $\eta = (\cos \theta)dy - (\sin \theta)dx$. As the problem is stated as Lagrangian, in order to use the results from above we must convert to the Hamiltonian side and transform the velocities to momenta, i.e. determine the musical isomorphism \sharp . (Where $g(\eta^\sharp, \cdot) = \eta$ which is the fiber derivative.)

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = M^{-1} \begin{bmatrix} p_x \\ p_y \\ p_\theta \end{bmatrix} = \begin{bmatrix} \frac{1}{m}p_x + \frac{1}{I} \left(ap_\theta \sin(\theta) - \frac{1}{2}a^2p_y \sin(2\theta) + a^2p_x \sin^2 \theta \right) \\ \frac{1}{m}p_y - \frac{1}{I} \left(ap_\theta \cos(\theta) + \frac{1}{2}a^2p_x \sin(2\theta) - a^2p_y \cos^2 \theta \right) \\ \frac{1}{I} (p_\theta - ap_y \cos(\theta) + ap_x \sin \theta) \end{bmatrix}$$

In terms of the momenta, the constraint (now on the cotangent bundle) is

$$(5.7) \quad P(W) = \frac{ma^2 + I}{Im} (p_y \cos \theta - p_x \sin \theta) - \frac{a}{I}p_\theta,$$

and the Hamiltonian is

$$H = \frac{1}{2m} (p_x^2 + p_y^2) + \frac{1}{2I} [p_\theta^2 + a^2 (p_y \cos \theta - p_x \sin \theta)^2 - 2ap_\theta p_y \cos \theta + 2ap_\theta p_x \sin \theta].$$

Applying equation (5.3), we see that

$$(5.8) \quad \operatorname{div}_{\mu_{\mathcal{E}}} (X_H^{\mathcal{D}}) = \frac{3Im}{I + ma^2} \left(-\sin \theta \frac{\partial}{\partial p_x} + \cos \theta \frac{\partial}{\partial p_y} \right) \{H, P(W)\}.$$

The Poisson bracket of the Hamiltonian with the constraint is

$$(5.9) \quad \{H, P(W)\} = \frac{1}{Im} (p_x \cos \theta + p_y \sin \theta) (p_\theta - ap_y \cos \theta + ap_x \sin \theta).$$

Combining equations (5.8) and (5.9), we get the following expression for the divergence.

$$(5.10) \quad \boxed{\operatorname{div}_{\mu_\phi}(X_H^{\mathcal{D}}) = -\frac{3a}{I + ma^2} (p_x \cos \theta + p_y \sin \theta)}$$

The well-known phenomenon of the dissipative nature of the Chaplygin sleigh, [83, 99], is encapsulated by the fact that $\operatorname{div}_{\mu_\phi}(X_H^{\mathcal{D}}) \neq 0$ unless $a = 0$.

Another way to interpret the divergence is via the sleigh's velocity,

$$v = \dot{x} \cos \theta + \dot{y} \sin \theta = \frac{1}{m} (p_x \cos \theta + p_y \sin \theta).$$

Substituting this into the formula for the divergence, we get

$$\operatorname{div}_{\mu_\phi}(X_H^{\mathcal{D}}) = -\frac{3mav}{I + ma^2}.$$

This shows that the vector field is contracting when the sleigh is moving forwards and expanding when it is moving backwards.

5.5.2 Rolling Penny

The next example we examine is that of the vertical rolling disk (penny). This is a system evolving on the configuration space $Q = \operatorname{SE}_2 \times S^1$ with coordinates $((x, y, \theta), \varphi)$; see Figure 5.2. The Lagrangian for this system is taken to be the kinetic energy,

$$L = \frac{1}{2}m(\dot{x}^2 + \dot{y}^2) + \frac{1}{2}I\dot{\theta}^2 + \frac{1}{2}J\dot{\varphi}^2.$$

Here, m is the mass of the penny, I is the moment of inertia of the disk about the axis perpendicular to the plane of the disk, and J is the moment of inertia about an




Figure 5.2: Configuration of the rolling penny, §1.4 in [16].

axis in the plane of the disk. If R is the radius of the penny, the conditions that the penny rolls without slipping is given by the two constraints

$$\eta^1(\dot{q}) = \dot{x} - R\dot{\theta} \cos \varphi = 0,$$

$$\eta^2(\dot{q}) = \dot{y} - R\dot{\theta} \sin \varphi = 0.$$

Transferring to the Hamiltonian side, we get the following expressions for the Hamiltonian and for the constraints:

$$\begin{aligned} H &= \frac{1}{2m} (p_x^2 + p_y^2) + \frac{1}{2I} p_\theta^2 + \frac{1}{2J} p_\varphi^2, \\ P(W^1) &= \frac{1}{m} p_x - \frac{1}{I} R p_\theta \cos \varphi, \\ P(W^2) &= \frac{1}{m} p_y - \frac{1}{I} R p_\theta \sin \varphi. \end{aligned}$$

The constraint mass matrix is

$$(m^{\alpha\beta}) = \begin{bmatrix} \frac{1}{m} + \frac{1}{I} R^2 \cos^2 \varphi & \frac{1}{I} R^2 \cos \varphi \sin \varphi \\ \frac{1}{I} R^2 \cos \varphi \sin \varphi & \frac{1}{m} + \frac{1}{I} R^2 \sin^2 \varphi \end{bmatrix}.$$

The Poisson brackets of the constraints are

$$\{H, P(W^1)\} = -\frac{1}{IJ} R p_\varphi p_\theta \sin \varphi,$$

$$\{H, P(W^2)\} = \frac{1}{IJ} R p_\varphi p_\theta \cos \varphi.$$

After a lengthy but straightforward computation with (5.4), we get that

$$(5.11) \quad \boxed{\operatorname{div}_{\mu^\phi} (X_H^{\mathcal{D}}) = 0}$$

Therefore, unlike the Chaplygin sleigh, the rolling penny *does* preserve measure.

5.6 Invariant Measures and the Cohomology Equation

In general, the divergence of a nonholonomic system does not vanish as (5.10) demonstrated. However, does there exist a *different* volume form that is invariant under the flow? i.e. does there exist a smooth $f > 0$ (a density) such that $\operatorname{div}_{f\mu_{\mathcal{E}}}(X_H^{\mathcal{D}}) = 0$? Finding such an f requires the solution of a certain type of partial differential equation (PDE) which is known as the smooth dynamical cohomology equation.

5.6.1 The Cohomology Equation

What conditions need to be met for f such that $f\mu_{\mathcal{E}}$ is an invariant volume form? Using the formula for the divergence as well as the fact that the Lie derivative is a derivation yields:

$$\operatorname{div}_{f\mu_{\mathcal{E}}}(X_H^{\mathcal{D}}) = \operatorname{div}_{\mu_{\mathcal{E}}}(X_H^{\mathcal{D}}) + \frac{1}{f}\mathcal{L}_{X_H^{\mathcal{D}}}(f).$$

Therefore the density, f , yields an invariant measure if and only if

$$(5.12) \quad \frac{1}{f}\mathcal{L}_{X_H^{\mathcal{D}}}(f) = -\operatorname{div}_{\mu_{\mathcal{E}}}(X_H^{\mathcal{D}}).$$

Notice that the left hand side of (5.12) can be integrated to

$$\frac{1}{f}\mathcal{L}_{X_H^{\mathcal{D}}}(f) = d(\ln f)(X_H^{\mathcal{D}}).$$

Calling $g = \ln f$, we have the following proposition.

Proposition V.17. *For a nonholonomic vector field, $X_H^{\mathcal{D}}$, there exists a smooth invariant volume, $f\mu_{\mathcal{E}}$, if there exists an exact 1-form $\alpha = dg$ such that*

$$(5.13) \quad \alpha(X_H^{\mathcal{D}}) = -\operatorname{div}_{\mu_{\mathcal{E}}}(X_H^{\mathcal{D}}).$$

Then the density is (up to a multiplicative constant) $f = e^g$.

Therefore the existence of invariant volumes boils down to finding global solutions to the PDE (5.13). The remainder of this section deals with uniqueness of solutions and a necessary condition for solutions to exist.

Remark V.18. PDEs of the form $dg(X) = f$ for a given smooth function f and vector field X are called cohomology equations [47, 75]. Thus the equation (5.13) is a cohomology equation.

5.6.2 Existence and Uniqueness

Uniqueness

Assuming that there exists a function $g \in C^\infty(\mathcal{D}^*)$ that solves (5.13), do there exist other solutions? Suppose that g_1 and g_2 both solve (5.13). Then their difference must be a first integral of the system: $\mathcal{L}_{X_H}(g_1 - g_2) = 0$. Solutions of (5.13) are then unique up to constants of motion. i.e. if g solves (5.13), then every invariant density has the form (again, up to a multiplicative constant)

$$f = \exp(g + \text{constant of motion}).$$

Therefore invariant measures can be thought of as an affine space with dimension being equal to the number of first integrals of the nonholonomic system.

Existence

For the purposes of this section, we examine an arbitrary cohomology equation $dg(X) = f$. Using the method of characteristics, if there exists a solution, it must take the form in the proposition below.

Proposition V.19. *Consider the first-order linear PDE $dg(X) = f$ where $X \in \mathfrak{X}(M)$, $f \in C^\infty(M)$ and denote the flow of $\dot{x} = X(x)$ by φ_t . Then g is a solution if and only if for all $x \in M$ and $t \in \mathbb{R}$*

$$(5.14) \quad g(\varphi_t(x)) - g(x) = \int_0^t f(\varphi_s(x)) ds.$$

Proof. First suppose that (5.14) holds. Then

$$\begin{aligned} dg_x(X) &= \mathcal{L}_X g(x) = \lim_{t \rightarrow 0} \frac{1}{t} [g(\varphi_t(x)) - g(x)] \\ &= \lim_{t \rightarrow 0} \frac{1}{t} \int_0^t f(\varphi_s(x)) ds \\ &= f(x) \end{aligned}$$

Now assume that $dg(X) = f$. Fix $x \in M$ and define the function $\lambda : \mathbb{R} \rightarrow \mathbb{R}$ by $\lambda(t) = g(\varphi_t(x))$. Then $\dot{\lambda}(t) = f(\varphi_t(x))$. Integrating this yields:

$$\lambda(t) - \lambda(0) = \int_0^t f(\varphi_s(x)) ds.$$

This is precisely (5.14). □

In fact, if there is a periodic orbit, $\varphi_t(x) = x$ for some t and x , we must have $g(\varphi_t(x)) - g(x) = 0$. Using (5.14) this can be written as a necessary condition for the existence of a solution.

Corollary V.20. *Consider the PDE $dg(X) = f$ on M . Then a necessary condition for a global solution is f integrated along periodic orbits must be zero. i.e. if there exists $x \in M$ and $t \in \mathbb{R}$ such that $\varphi_t(x) = x$, then*

$$\int_0^t f(\varphi_s(x)) ds = 0.$$

In the context of our original problem, in order for an invariant volume to exist, the divergence need not be zero but it must *average* to zero along periodic orbits. That is, the expanding and contracting parts must “cancel out.” Corollary V.20 offers a convenient test to show invariant measures are impossible, but does a converse exist? It turns out that with a certain other assumption the answer is yes, [55].

Theorem V.21 ([55]). *Let M be a compact manifold and X a smooth vector field on M whose associated flow is Anosov and transitive. Let f be a smooth function*

on M whose integral over every periodic orbit of X is zero. Then there exists a C^∞ function u such that $du(X) = f$.

Therefore, if we can show that the flow is Anosov and check that the integrals along periodic orbits vanish then an invariant measure must exist. However, this is generally quite difficult to check and will not be pursued here. Instead, the next section examines the case where invariant measures depend only on the configuration variables or variables.

5.7 Special case: Measures Depending on Configuration

In general, solving the cohomology equation (5.13) is quite difficult. It turns out, however, that it is *relatively* easy to determine necessary and sufficient conditions on the solvability when the density is assumed to depend only on the configuration variables.

Definition V.22. A density $f : T^*Q \rightarrow \mathbb{R}$ is said to depend on configuration if $f = \pi_Q^*g$ for some $g : Q \rightarrow \mathbb{R}$.

Under this assumption, (5.6) can be presented in a surprisingly nice way. In this case, the divergence can be described by an equivalence class of 1-forms. The density form, defined below, is a representative element from this class.

Definition V.23. Let \mathcal{C} be a (natural) realization of $\mathcal{D}^* \subset T^*Q$. Then, define the *density form* to be the following 1-form

$$\vartheta_{\mathcal{C}} = m_{\alpha\beta} \cdot \mathcal{L}_{W^\alpha} \eta^\beta.$$

Studying the 1-form, $\vartheta_{\mathcal{C}}$, provides necessary and sufficient conditions for the existence of densities depending only on configuration. Recall that $\mathcal{D}^0 = \text{Ann}(\mathcal{D}) \subset T^*Q$ is the annihilator of $\mathcal{D} \subset TQ$.

Theorem V.24. *There exists an invariant density depending on configuration if and only if there exists $\rho \in \Gamma(\mathcal{D}^0)$ such that $\vartheta_{\mathcal{C}} + \rho$ is exact.*

Proof. To show this, we will prove that $-n \cdot \pi_Q^* \vartheta_{\mathcal{C}}(X_H^{\mathcal{D}}) = \text{div}_{\mu_{\mathcal{C}}}(X_H^{\mathcal{D}})$. Recall that the differential of a 1-form is given by $d\alpha(X, Y) = X\alpha(Y) - Y\alpha(X) - \alpha([X, Y])$ and that $\pi_Q^* \eta^\beta(X_H)|_{\mathcal{D}^*} = 0$. Returning to (5.6), we have

$$\begin{aligned}
\text{div}_{\mu_{\mathcal{C}}}(X_H^{\mathcal{D}}) &= -n \cdot m_{\alpha\beta} \cdot \pi_Q^* \eta^\alpha([X_H, X_{P(W^\beta)}]) \\
&= -n \cdot m_{\alpha\beta} \cdot (X_H \pi_Q^* \eta^\alpha(X_{P(W^\beta)}) - X_{P(W^\beta)} \pi_Q^* \eta^\alpha(X_H) - d\pi_Q^* \eta^\alpha(X_H, X_{P(W^\beta)})) \\
&= -n \cdot m_{\alpha\beta} \cdot (X_H m^{\alpha\beta} - \pi_Q^* d\eta^\alpha(X_H, X_{P(W^\beta)})) \\
&= -n \cdot m_{\alpha\beta} \cdot (dm^{\alpha\beta}(\dot{q}) - d\eta^\alpha(\dot{q}, W^\beta)) \\
&= -n \cdot m_{\alpha\beta} \cdot (di_{W^\beta} \eta^\alpha + i_{W^\beta} d\eta^\alpha)(\dot{q}) \\
&= -n \cdot m_{\alpha\beta} \cdot \mathcal{L}_{W^\beta} \eta^\alpha(\dot{q}).
\end{aligned}$$

This computation shows that $\text{div}_{\mu_{\mathcal{C}}}(X_H^{\mathcal{D}}) = -n \cdot \vartheta_{\mathcal{C}}(\dot{q})$, but \dot{q} cannot be arbitrary as it must lie within \mathcal{D} . Therefore, we can add on an element of \mathcal{D}^0 to $\vartheta_{\mathcal{C}}$ without changing its value on \dot{q} : $\text{div}_{\mu_{\mathcal{C}}}(X_H^{\mathcal{D}}) = -n \cdot (\vartheta_{\mathcal{C}} + \rho)(\dot{q})$ for any $\rho \in \mathcal{D}^0$. Hence, a solution exists depending on configuration if $\vartheta_{\mathcal{C}} + \rho$ can be integrated, i.e. it is exact. \square

The above shows that exactness of $\vartheta_{\mathcal{C}}$ determines the existence of a density depending on configuration. How does this depend on the choice of \mathcal{C} to realize the constraints? It turns out the answer is independent of the choice of realization.

Theorem V.25. *Let \mathcal{C} and \mathcal{C}' both be realizations of the constraint \mathcal{D}^* . If $\vartheta_{\mathcal{C}} + \rho$ is exact, then there exists ρ' such that $\vartheta_{\mathcal{C}'} + \rho'$ is too. Moreover, if $\vartheta_{\mathcal{C}} + \rho = df$ and $\vartheta_{\mathcal{C}'} + \rho' = df'$, then $e^f \cdot \mu_{\mathcal{C}} = e^{f'} \cdot \mu_{\mathcal{C}'}$, modulo a constant of motion.*

Proof. Suppose, as was in the proof of Proposition V.8, that there is a single con-

straint such that $\mathcal{C} = \{P(W)\}$ and $\mathcal{C}' = \{h \cdot P(W)\}$. Computing $\vartheta_{\mathcal{C}'}$ gives:

$$\begin{aligned}\vartheta_{\mathcal{C}'} &= \frac{1}{h^2 m} \mathcal{L}_{hW} (h\eta) \\ &= \frac{1}{h^2 m} (h^2 \cdot \mathcal{L}_W \eta + hm \cdot dh + h \cdot dh(W)\eta) \\ &= \frac{1}{m} \mathcal{L}_W \eta + \frac{1}{h} dh + \frac{dh(W)}{hm} \eta \\ &= \vartheta_{\mathcal{C}} + d[\ln h] + \alpha \cdot \eta.\end{aligned}$$

This shows that $\vartheta_{\mathcal{C}'}$ and $\vartheta_{\mathcal{C}}$ differ by something exact and something living in \mathcal{D}^0 .

The component in \mathcal{D}^0 can be disregarded as it is absorbed into ρ' . Integrating gives

$f' = f + \ln h$ and it remains to prove that $h\mu_{\mathcal{C}'} = \mu_{\mathcal{C}}$. Recalling Definition V.10, we

have $\sigma_{\mathcal{C}} = dP(W)$ and $\sigma_{\mathcal{C}'} = P(W)dh + hdP(W)$, so

$$dP(W) \wedge \varepsilon = (P(W)dh + hdP(W)) \wedge \varepsilon' = \omega^n, \quad \mu_{\mathcal{C}} = \iota^* \varepsilon, \quad \mu_{\mathcal{C}'} = \iota^* \varepsilon'.$$

Using the fact that $P(W) = 0$ under the pullback of ι , this component can be ignored

and we have $\mu_{\mathcal{C}} = h\mu_{\mathcal{C}'}$. □

Remark V.26. It is only possible for $e^f \cdot \mu_{\mathcal{C}}$ and $e^{f'} \cdot \mu_{\mathcal{C}'}$ to be off by a constant of motion if there exists an exact form in $\Gamma(\mathcal{D}^0)$. This only happens if the constraints are not completely nonintegrable.

A reason why studying $\vartheta_{\mathcal{C}}$ is insightful is that it immediately demonstrates why holonomic systems are measure-preserving. This can be shown with the help of a useful lemma.

Lemma V.27 (*). $m_{\alpha\beta} \cdot dm^{\alpha\beta} = d[\ln \det(m^{\alpha\beta})]$.

Proof. It suffices to check along a curve in the manifold. Let $\gamma : I \rightarrow Q$ be a curve and let $A(t) = (m^{\alpha\beta}) \circ \gamma(t)$ be the mass matrix along the curve. Note that $A(t)$ is

*We thank Dr. Alexander Barvinok for help with this proof.

positive-definite and changes smoothly with t . We have

$$\frac{d}{dt} \ln \det A(t) = \frac{\frac{d}{dt} \det A(t)}{\det A(t)} = \sum_{i=1}^m \frac{\det A_i(t)}{\det A(t)},$$

where $A_i(t)$ is obtained from $A(t)$ by differentiating the i -th row and leaving all other rows intact, i.e.

$$A_i(t) = \begin{pmatrix} a_{11}(t) & \cdots & a_{1m}(t) \\ \vdots & \ddots & \vdots \\ a_{(i-1)1}(t) & \cdots & a_{(i-1)m}(t) \\ a'_{i1}(t) & \cdots & a'_{im}(t) \\ a_{(i+1)1}(t) & \cdots & a_{(i+1)m}(t) \\ \vdots & \ddots & \vdots \\ a_{m1}(t) & \cdots & a_{mm}(t) \end{pmatrix}.$$

Expanding $\det A_i(t)$ along the i -th row:

$$\det A_i(t) = \sum_{j=1}^m (-1)^{i+j-1} a'_{ij}(t) \det A_{ij}(t),$$

where $A_{ij}(t)$ is the $(m-1) \times (m-1)$ matrix obtained from $A_i(t)$ and hence from $A(t)$ by crossing out the i -th row and j -th column.

Next, observe that $(-1)^{i+j-1} \det A_{ij} / \det A(t)$ is the (j, i) -th entry of the inverse matrix $A^{-1}(t) = (b_{ij})(t)$, and since $A(t)$ is symmetric, is also the (i, j) -th entry of $(b_{ij})(t)$. Summarizing,

$$\frac{d}{dt} \ln \det A(t) = \sum_{i,j=1}^m a'_{ij}(t) b_{ij}(t).$$

□

Proposition V.28. *If the constraints are holonomic, then there exists a $\rho \in \Gamma(\mathcal{D}^0)$ such that $\vartheta_{\mathcal{E}} + \rho$ is exact. In particular, if \mathcal{E} is chosen such that all η^α are closed, $\vartheta_{\mathcal{E}}$ is exact.*

Proof. When the constraints are holonomic, the 1-forms η^α can be chosen such that they are closed. Then the density form is

$$\begin{aligned}\vartheta_{\mathcal{L}} &= m_{\alpha\beta} (di_{W^\beta}\eta^\alpha + i_{W^\beta}d\eta^\alpha) \\ &= m_{\alpha\beta} \cdot dm^{\alpha\beta} \\ &= d[\ln \det (m^{\alpha\beta})],\end{aligned}$$

which is exact by Lemma V.27. \square

Example V.29 (The Chaplygin Sleigh). We wish to compute $\vartheta_{\mathcal{L}}$ for the Chaplygin sleigh and show that no measures depending on configuration exist. For this example,

$$W = \frac{ma^2 + I}{Im} \left[\cos\theta \frac{\partial}{\partial y} - \sin\theta \frac{\partial}{\partial x} \right] - \frac{a}{I} \frac{\partial}{\partial \theta}, \quad \eta = (\cos\theta)dy - (\sin\theta)dx.$$

This gives us

$$\begin{aligned}\vartheta_{\mathcal{L}} &= \frac{1}{\eta(W)} \mathcal{L}_W \eta \\ &= \frac{ma}{ma^2 + I} [(\sin\theta)dy + (\cos\theta)dx].\end{aligned}$$

We want to show that for any $\tilde{\eta} \in \Gamma(\mathcal{D}^0)$, $\vartheta_{\mathcal{L}} + \tilde{\eta}$ is not exact. Because there is only one constraint, it suffices to show that there does not exist a smooth k such that $\vartheta_{\mathcal{L}} + k \cdot \eta$ is exact, i.e. it requires the following to be zero:

$$\begin{aligned}d(\vartheta_{\mathcal{L}} + k \cdot \eta) &= \frac{ma}{ma^2 + I} [(\cos\theta)d\theta \wedge dy - (\sin\theta)d\theta \wedge dx] \\ &\quad + \left(\frac{\partial k}{\partial x} \cos\theta + \frac{\partial k}{\partial y} \sin\theta \right) dx \wedge dy \\ &\quad + \left(\frac{\partial k}{\partial \theta} \cos\theta - k \sin\theta \right) d\theta \wedge dy \\ &\quad - \left(\frac{\partial k}{\partial \theta} \sin\theta + k \cos\theta \right) d\theta \wedge dx.\end{aligned}$$

Separating the above, we need the following three to vanish:

$$(5.15) \quad \begin{aligned}0 &= \frac{\partial k}{\partial x} \cos\theta + \frac{\partial k}{\partial y} \sin\theta, \\ 0 &= \frac{\partial k}{\partial \theta} \cos\theta - k \sin\theta + \frac{ma}{ma^2 + I} \cos\theta, \\ 0 &= \frac{\partial k}{\partial \theta} \sin\theta + k \cos\theta + \frac{ma}{ma^2 + I} \sin\theta.\end{aligned}$$

The second two lines of (5.15) are overdetermined for k in the θ -direction and are inconsistent (unless $a = 0$ and we obtain the trivial solution $k \equiv 0$). Therefore, there does not exist a smooth k such that $\vartheta_{\mathcal{E}} + k \cdot \eta$ is closed.

Remark V.30. The necessary condition $d(\vartheta_{\mathcal{E}} + \rho) = 0$ creates a system of $\binom{n}{2}$ partial differential equations with m unknowns. It is a straight-forward check to determine whether or not this system is inconsistent; if it is inconsistent, then automatically no invariant measures can exist. If this system turns out to be consistent, then actually solving it reduces to solving another cohomology equation.

5.8 Invariant Measures Coming from a Momentum Map

As can be seen in Example V.29, the Chaplygin sleigh does not have an invariant measure depending only on its base variables. Therefore, if any invariant measure were to exist, it would have to depend on the fiber variables as well. We present one way to search for solutions using the momentum map (3.8). In general for a Lie group G , this method produces a space of possible solutions of $\dim G < \infty$ while the space of all possible solutions is infinite. However, as we will see in §5.8.1, this technique *will* find an invariant measure for the Chaplygin sleigh.

Definition V.31. Let G be a Lie group which has a symplectic action on $P := T^*Q$. The nonholonomic system is said to have an invariant G -measure if there exists a $\xi \in \mathfrak{g}$ such that $\hat{J}(\xi)$ solves the cohomology equation (5.13), cf. Definition III.18.

As nice as it would be if a G -measure existed, it seems unlikely due to the fact that the space of smooth functions that *might* solve $dg(X) = f$ is infinite-dimensional while the dimension of momentum maps is $\dim G < \infty$. As such, we will widen our search by introducing integrating factors in the following sense: Suppose $\alpha \in \Omega^1(\mathcal{D}^*)$ solves $\alpha(X_H^{\mathcal{D}}) = \text{div}_{\mu_{\mathcal{E}}}(X_H^{\mathcal{D}})$ and define the vector field X_{α} via $i_{X_{\alpha}}\omega = \alpha$. Then α

being exact is equivalent to X_α being Hamiltonian and α being closed is equivalent to X_α being symplectic. We will therefore relax our search and only require X_α to be symplectic (i.e. closed solutions to (5.13)). Rewriting (5.13), we have to find a symplectic vector field X such that

$$(5.16) \quad \nu_H^{\mathcal{C}}(X) = \operatorname{div}_{\mu_{\mathcal{C}}} (X_H^{\mathcal{D}}),$$

where $\nu_H^{\mathcal{C}} = dH - m_{\alpha\beta}\{H, P(W^\beta)\}\pi_Q^*\eta^\alpha$ is the nonholonomic 1-form. For a given infinitesimal generator on $M = T^*Q$, ξ_M , we wish to determine an integrating factor such that $h\xi_M$ remains symplectic.

Lemma V.32. *Let $h \in C^\infty(M)$, then the vector field $h\xi_M$ is symplectic if and only if there exists a smooth function $k : \mathbb{R} \rightarrow \mathbb{R}$ such that $h = k \circ \hat{J}(\xi)$.*

Proof. In order for $h\xi_M$ to be symplectic, we need $hi_{\xi_M}\omega$ to be closed. A straightforward computation yields:

$$\begin{aligned} d(hi_{\xi_M}\omega) &= d(h \cdot d\hat{J}(\xi)) \\ &= dh \wedge d\hat{J}(\xi). \end{aligned}$$

Therefore, $h\xi_M$ is symplectic if and only if dh is linearly dependent on $d\hat{J}(\xi)$, i.e. $dh = k'(\hat{J}(\xi))d\hat{J}(\xi)$. \square

By replacing ξ_M with $h\xi_M$ with the help of Lemma V.32, we have drastically expanded the scope of our search.

Definition V.33. A nonholonomic system has a twisted G -measure if there is a smooth function $k : \mathbb{R} \rightarrow \mathbb{R}$ such that $k(\hat{J}(\xi))\xi_M$ solves (5.16).

As it turns out, there is a reasonable test to see if a twisted G -measure exists and the resulting symplectic vector field, $k \circ \hat{J}(\xi)\xi_M$, is Hamiltonian.

Theorem V.34. *There exists a twisted G -measure for (5.16) if and only if*

$$\{\ln(\nu_H^{\mathcal{C}}(\xi_M)), \hat{J}(\xi)\} = \{\ln(\operatorname{div}_{\mu_{\varphi}}(X_H^{\mathcal{D}})), \hat{J}(\xi)\},$$

for some $\xi \in \mathfrak{g}$. Moreover, the resulting symplectic vector field, $k(\hat{J}(\xi))\xi_M$ is Hamiltonian with Hamiltonian (the a is a constant of integration)

$$(5.17) \quad G(p, q) = \int_a^{\hat{J}(\xi)(p, q)} k(s) ds.$$

Proof. To simplify calculations, we will call $f := \operatorname{div}_{\mu_{\varphi}}(X_H^{\mathcal{D}})$. Let us suppose that $k(\hat{J}(\xi))\nu_H^{\mathcal{C}}(\xi_M) = f$. Differentiating gives

$$k'(\hat{J}(\xi))\nu_H^{\mathcal{C}}(\xi_M)d\hat{J}(\xi) + k(\hat{J}(\xi))d[\nu_H^{\mathcal{C}}(\xi_M)] = df.$$

This provides us with the two following equations:

$$k(\hat{J}(\xi))d[\nu_H^{\mathcal{C}}(\xi_M)](\xi_M) = df(\xi_M)$$

$$k(\hat{J}(\xi))\nu_H^{\mathcal{C}}(\xi_M) = f.$$

Dividing these equations gives

$$d[\ln(\nu_H^{\mathcal{C}}(\xi_M))](\xi_M) = d[\ln(f)](\xi_M).$$

The first result follows. To show the second result, suppose the symplectic vector field has a Hamiltonian of the form $h(\hat{J}(\xi))$. Differentiating this gives

$$h'(\hat{J}(\xi))d\hat{J}(\xi) = k(\hat{J}(\xi))d\hat{J}(\xi).$$

Therefore k is an antiderivative of h . □

5.8.1 The Chaplygin Sleigh Revisited

We conclude our work with proving the following theorem.

Theorem V.35. *The Chaplygin sleigh has a twisted S^1 -measure.*

Proof. Recall the Chaplygin sleigh where the Hamiltonian is

$$H = \frac{1}{2m} (p_x^2 + p_y^2) + \frac{1}{2I} [p_\theta^2 + a^2 (p_y \cos \theta - p_x \sin \theta)^2 - 2ap_\theta p_y \cos \theta + 2ap_\theta p_x \sin \theta],$$

with constraint

$$P(W) = \frac{I + ma^2}{Im} (p_y \cos \theta - p_x \sin \theta) - \frac{a}{I} p_\theta.$$

Also recall that the divergence of this system is

$$\operatorname{div}_{\mu^{\mathcal{C}}} (X_H^{\mathcal{D}}) = -\frac{3a}{I + ma^2} (p_x \cos \theta + p_y \sin \theta).$$

We will next introduce the S^1 action that will generate the invariant measure.

Let S^1 act on SE_2 in the following way: for $t \in \mathbb{S}^1 = \mathbb{R}/2\pi\mathbb{Z}$, and $(x, y, \theta) \in \operatorname{SE}_2$

$$t \cdot \begin{bmatrix} x \\ y \\ \theta \end{bmatrix} = \begin{bmatrix} -a \cdot \cos(\theta + t) + a \cdot \cos \theta + x \\ a \cdot \sin(\theta + t) - a \cdot \sin \theta + y \\ \theta + t \end{bmatrix}.$$

The infinitesimal generator for this action is

$$\xi_Q = a \cdot \sin \theta \frac{\partial}{\partial x} - a \cdot \cos \theta \frac{\partial}{\partial y} + \frac{\partial}{\partial \theta}.$$

The momentum map for the lifted action to $T^*\operatorname{SE}_2$ is

$$(5.18) \quad \hat{J}(u) = u [p_\theta - ap_y \cos \theta + ap_x \sin \theta],$$

where $u \in \operatorname{Lie}(S^1) = \mathbb{R}$. The action chosen here seems unnatural until it is realized that $\hat{J}(u) = u \cdot \mathbb{F}L(\dot{\theta})$, i.e. this is the action giving rotation about the contact point.

A computation provides:

$$\begin{aligned} \nu_H^{\mathcal{C}}(u_M) &= dH(u_M) + \frac{1}{\eta(W)} \{P(W), H\} \pi_Q^* \eta(u_M) \\ &= \{H, \hat{J}(u)\} - \frac{1}{ma^2 + I} (p_x \cos \theta + p_y \sin \theta) (p_\theta - ap_y \cos \theta + ap_x \sin \theta) \eta(u \cdot \xi_Q) \\ &= \frac{au}{ma^2 + I} (p_x \cos \theta + p_y \sin \theta) (p_\theta - ap_y \cos \theta + ap_x \sin \theta) \end{aligned}$$

Checking the condition of Theorem V.34, we see that

$$\begin{aligned} \{\ln(\nu_H^{\mathcal{C}}(u_M)), \hat{J}(u)\} &= u \cdot \frac{p_y \cos \theta - p_x \sin \theta}{p_x \cos \theta + p_y \sin \theta} \\ &= \{\ln(f), \hat{J}(u)\}. \end{aligned}$$

□

It remains a computation to find the invariant measure for the Chaplygin sleigh.

Theorem V.36. $\hat{J}(u)^{-3}$ is an invariant (singular) density for the Chaplygin sleigh.

Proof. We know that there exists a smooth function $k : \mathbb{R} \rightarrow \mathbb{R}$ such that $k \circ \hat{J}(u) \nu_H^{\mathcal{C}}(u_M) = f$. By inspection and using (5.17), we get that $-3 \ln(\hat{J}(u))$ is a solution to (5.13). Therefore, its exponential is the invariant density. □

Remark V.37. Notice that the measure is singular on the set $p_\theta - ap_y \cos \theta + ap_x \sin \theta = 0$. This corresponds to motion in a straight line with zero rotation - which is the set where the trajectories asymptotically approach. This is why asymptotic stability is still compatible with an invariant measure.

CHAPTER VI

Mechanical Hybrid Systems

This chapter begins to attempt to fuse the ideas of Chapters II and III, i.e. to analyze mechanical impacts using the theory of geometric mechanics. A mechanical hybrid system will have the form $\mathcal{H} = (M, S, X, \Delta)$ where $M = TQ$ or T^*Q (depending on Lagrangian/Hamiltonian), $S \subset M$ is the location of impact, X is a Lagrangian/Hamiltonian/nonholonomic vector field, and Δ is a mechanical impact map. Understanding the vector field was the object of study in Chapters III and V while this chapter deals with understanding Δ . In order to construct this map, we make the following assumption (cf. §3.5 in [17]):

Assumption VI.1. *A mechanical impact is the identity on the base and satisfies variational/Lagrange-d'Alembert principles on the fibers. In particular, the impact map will have the form $\Delta = (\text{Id}, \delta)$, e.g. for Hamiltonian systems we have $\pi_Q \circ \Delta = \pi_Q$ where $\pi_Q : T^*Q \rightarrow Q$ is the cotangent projection.*

6.1 Mechanical Impact Maps

With a the continuous dynamics understood from either (3.5), (3.6), or (5.1), we can develop the impact map. Our discussion considers first the case of unconstrained systems before examining nonholonomic systems. We first outline the definition of a hybrid mechanical system [5].

Definition VI.2. A hybrid Lagrangian system is a triple (Q, L, S) where $L : TQ \rightarrow \mathbb{R}$ is a Lagrangian and $S \subset Q$ is a smooth, embedded submanifold of codimension 1. Similarly, a hybrid Hamiltonian system is a triple (Q, H, S) where $H : T^*Q \rightarrow \mathbb{R}$ is the Legendre transform of L . A hybrid mechanical system is either a hybrid Lagrangian or hybrid Hamiltonian system.

Remark VI.3. Equivalently, a hybrid Lagrangian system can be given via a *realization*, h , of S where $h : Q \rightarrow \mathbb{R}$ such that $S = h^{-1}(0)$. It will be evident later on that the impact dynamics are derived using a realization but invariant under the choice of such a realization.

Hybrid mechanical systems, either Lagrangian or Hamiltonian, induce a hybrid dynamical system $\mathcal{H} = (M, S, X, \Delta)$ as described below. We abuse notation for $S \subset Q$ in the definition for a hybrid mechanical system and $S \subset TQ$ for a hybrid dynamical system; this notation will be straightened out in Remark VI.4.

- Lagrangian: $M = TQ$, $S = \{(q, \dot{q}) \in TQ : h(q) = 0, dh(\dot{q}) < 0\}$, and X is given by the Euler-Lagrange equations (3.5),
- Hamiltonian: $\mathcal{X} = T^*Q$, $S = \{(q, p) \in T^*Q : h(q) = 0, P(\nabla h) < 0\}$, and $X = X_H$ is the Hamiltonian vector field (3.6).

The map Δ is currently undefined and will be the focus of §6.1.1 and §6.1.2; the impact map will be fully determined by the data (Q, L, S) . The reason for the inequality $dh(\dot{q}) < 0$ (and analogously for $P(\nabla h) < 0$) for the Lagrangian impact surface is to say that an impact only occurs when the trajectory is moving outwards, i.e. S is the set of outward pointing vectors along the set $h^{-1}(0)$.

Remark VI.4. To assist with consistent notation, the following will be used for the remainder of this thesis:

- $S = \{q \in Q : h(q) = 0\} \subset Q$,
- $\hat{S} = \{(q, \dot{q}) \in TQ : h(q) = 0, dh(\dot{q}) < 0\} \subset TQ$, and
- $S^* = \{(q, p) \in T^*Q : h(q) = 0, P(\nabla h) < 0\} \subset T^*Q$.

This work is interested in understanding mechanical impacts while under the influence of nonholonomic constraints. For this reason we also define hybrid nonholonomic systems.

Definition VI.5. A nonholonomic hybrid Lagrangian system is a tuple (Q, L, S, \mathcal{D}) such that (Q, L, S) forms a hybrid Lagrangian system and $\mathcal{D} \subset TQ$ is a regular distribution. Additionally, a nonholonomic hybrid Hamiltonian system is a tuple (Q, H, S, \mathcal{D}^*) where H is the Legendre transform on L and $\mathcal{D}^* = \mathbb{F}L(\mathcal{D})$.

Remark VI.6. For notation purposes, we will use the following symbols for nonholonomic impact surfaces:

- $\hat{S}_{\mathcal{D}} := \hat{S} \cap \mathcal{D}$, and
- $S_{\mathcal{D}}^* := S^* \cap \mathcal{D}^*$.

Problems can emerge if $dh \in \mathcal{D}^0$ where $S = h^{-1}(0)$. If this happens h is preserved under the flow and impacts never (or constantly) occur, i.e. $S_{\mathcal{D}}^*$ no longer has codimension 1 because

$$S_{\mathcal{D}}^* = \{(q, p) \in T^*Q : h(q) = 0, P(\nabla h) > 0, p \in \mathcal{D}^*\},$$

and $p \in \mathcal{D}^*$ implies that $P(\nabla h) = 0$ which is contradictory. Therefore, for nonholonomic hybrid systems we will make the following assumption.

Assumption VI.7 (Nontrivial impact condition). *Suppose that $S \subset Q$ is given by $S = h^{-1}(0)$. Then $dh|_S \notin \mathcal{D}^0|_S$.*




Figure 6.1: The black curve is the continuous, piece-wise differentiable trajectory with an impact at t_1 while the red curve shows a virtual displacement.

Nonholonomic hybrid mechanical systems induce hybrid dynamical systems in the same way as their unconstrained counterparts.

6.1.1 Holonomic Impacts

We begin with the observation that the Euler-Lagrange equations are variational. With this, we assume that the impact map is variational as well. This is realized by the Weierstrass-Erdmann corner conditions, cf. e.g. §3.5 of [17] or §4.4 of [71]:

$$(6.1) \quad \begin{aligned} \mathbb{F}L^+ - \mathbb{F}L^- &= \varepsilon \cdot dh, \\ L^+ - \langle \mathbb{F}L^+, \dot{q}^+ \rangle &= L^- - \langle \mathbb{F}L^-, \dot{q}^- \rangle, \end{aligned}$$

where $S = \{q \in Q : h(q) = 0\}$ and the multiplier ε is chosen such that both equations are satisfied. These conditions have a cleaner interpretation on the Hamiltonian side:

$$(6.2) \quad \begin{aligned} p^+ &= p^- + \varepsilon \cdot dh, \\ H^+ &= H^-, \end{aligned}$$

i.e. energy is conserved and the change in momentum is perpendicular to the impact surface. This is precisely specular reflection. In the case where L is natural (the difference of kinetic and potential energy), the corner conditions can be explicitly solved. Recall that $\nabla h = dh^\sharp$, or equivalently $dh = g(\nabla h, \cdot)$.

Theorem VI.8. *Given a natural Lagrangian $L(q, \dot{q}) = \frac{1}{2}g_q(\dot{q}, \dot{q}) - V(q)$, the impact map is given by $\Delta : (q^-, \dot{q}^-) \mapsto (q^-, \delta(q^-, \dot{q}^-))$ where*

$$(6.3) \quad \delta(q, \dot{q}) = \dot{q} - 2 \frac{dh(\dot{q})}{g(\nabla h, \nabla h)} \nabla h.$$

Proof. For a Lagrangian having the form $L(q, \dot{q}) = \frac{1}{2}g_q(\dot{q}, \dot{q}) - V(q)$, the corresponding Hamiltonian is

$$H(q, p) = \frac{1}{2}\tilde{g}_q(p, p) + V(q), \quad \tilde{g}(p, p') = g(p^\sharp, p'^\sharp).$$

Using the fact that $q^+ = q^-$ along with energy conservation, we get

$$\tilde{g}(p^+, p^+) = \tilde{g}(p^-, p^-).$$

Coupling this with the first half of (6.2), we get

$$\begin{aligned} \tilde{g}(p^+, p^+) &= \tilde{g}(p^- + \varepsilon \cdot dh, p^- + \varepsilon \cdot dh) \\ &= \tilde{g}(p^-, p^-) + \varepsilon^2 \cdot \tilde{g}(dh, M^{-1}dh) + 2\varepsilon \cdot \tilde{g}(p^-, dh). \end{aligned}$$

This results in a quadratic equation in ε with a trivial solution $\varepsilon = 0$ as well as

$$\varepsilon = -2 \cdot \frac{\tilde{g}(p^-, dh)}{\tilde{g}(dh, dh)}.$$

In terms of the momenta, the impact map is

$$(6.4) \quad p^+ = p^- - 2 \cdot \frac{\tilde{g}(p^-, dh)}{\tilde{g}(dh, dh)} dh.$$

Flipping (6.4) back to the Lagrangian side with the musical isomorphism, $p = \dot{q}^\flat$ and $dh = \nabla h^\flat$, we see

$$\dot{q}^+ = \dot{q}^- - 2 \cdot \frac{dh(\dot{q}^-)}{g(\nabla h, \nabla h)} \nabla h.$$

This matches (6.3). □

Remark VI.9. The mechanical impact map, (6.3), is precisely a Householder matrix.

Plastic Holonomic Impacts

Notice that (6.3) matches the impact equation in [5] with the exception of the coefficient of restitution $0 \leq e \leq 1$, that is

$$(6.5) \quad \dot{q}^+ = \dot{q}^- - (1 + e) \frac{dh(\dot{q}^-)}{g(\nabla h, \nabla h)} \nabla h.$$

Remark VI.10. In equation (6.5), $e = 1$ is precisely the variational case (elastic). When $e = 0$, the impact is the orthogonal projection (with respect to g) $TQ \rightarrow \ker dh$ (plastic).

6.1.2 Nonholonomic Impacts

The study of nonholonomic impacts falls into two distinct categories: elastic and plastic. Elastic impacts arise from the Lagrange-d'Alembert principle while plastic impacts come from an orthogonal projection.

Elastic Impacts

Consider the variational impact equations (the corner conditions):

$$(\mathbb{F}L^+ - \mathbb{F}L^-) \delta q = 0,$$

$$(H^+ - H^-) \delta t = 0.$$

If the impact time is free, $\delta t \neq 0$, then we have conservation of energy which is the second corner condition. However, unlike before, the spatial variations δq must satisfy two sets of constraints: the nonholonomic constraints $\eta^k(\delta q) = 0$, and the impact constraint $dh(\delta q) = 0$. This leads to a modified version of (6.1),

$$(6.6) \quad \begin{aligned} \mathbb{F}L^+ - \mathbb{F}L^- &= \lambda_k \cdot \eta^k + \varepsilon \cdot dh, \\ L^+ - \langle \mathbb{F}L^+, \dot{q}^+ \rangle &= L^- - \langle \mathbb{F}L^-, \dot{q}^- \rangle, \\ \eta^k(\dot{q}^+) &= 0. \end{aligned}$$

Proposition VI.11. *Assuming that L is a natural Lagrangian and all η^k are linearly independent, there exists at most one nontrivial solution to (6.6).*

Proof. Using the fact that $p^+ = p^- + \lambda_k \cdot \eta^k + \varepsilon \cdot dh$ and substituting in the other constraints, we get m linear equations and one quadratic equation.

$$(6.7) \quad \lambda_k \cdot \tilde{g}(\eta^k, \eta^\ell) + \varepsilon \cdot \tilde{g}(dh, \eta^\ell) = 0, \quad \forall \ell,$$

$$(6.8) \quad 2\lambda_k \cdot \tilde{g}(p^-, \eta^k) + 2\varepsilon \cdot \tilde{g}(p^-, dh) + \lambda_i \lambda_j \cdot \tilde{g}(\eta^i, \eta^j) \\ + 2\lambda_k \varepsilon \cdot \tilde{g}(\eta^k, dh) + \varepsilon^2 \cdot \tilde{g}(dh, dh) = 0.$$

Assume that the linear system (6.7) has a solution, which then depends linearly on ε . When substituted into (6.8), the resulting equation is quadratic in a single variable, ε , with $\varepsilon = 0$ as a solution. Therefore there can be at most one nonzero solution. \square

Definition VI.12. Let (Q, g) be Riemannian and (Q, L, S, \mathcal{D}) a corresponding natural nonholonomic Lagrangian system. Then the impact map given by Proposition VI.11 is called the elastic nonholonomic impact map.

Due to the Riemannian structure, (6.6) is solvable and hence the elastic nonholonomic impact map is well-defined. As it turns out, the multipliers λ_k and ε can be explicitly solved for.

Theorem VI.13. *Suppose that there are constraints $P(W^\alpha) = 0$. Then the elastic nonholonomic impact map is*

$$(6.9) \quad \dot{q}^+ = \dot{q}^- + \lambda_k \cdot W^k + \varepsilon \cdot \nabla h,$$

where

$$(6.10) \quad \varepsilon = \frac{2 \cdot m_{\alpha\beta} \cdot g(\nabla h, W^\beta) P(W^\alpha) - 2 \cdot P(\nabla h)}{g(\nabla h, \nabla h) - m_{\alpha\beta} \cdot g(\nabla h, W^\alpha) g(\nabla h, W^\beta)},$$

$$(6.11) \quad \lambda_k = -m_{k\ell} \cdot g(\nabla h, W^\ell) \cdot \frac{2 \cdot m_{\alpha\beta} \cdot g(\nabla h, W^\beta) P(W^\alpha) - 2 \cdot P(\nabla h)}{g(\nabla h, \nabla h) - m_{\alpha\beta} \cdot g(\nabla h, W^\alpha) g(\nabla h, W^\beta)}.$$

Remark VI.14. When we apply the constraints $P(W^\alpha) = 0$, the above reduces to

$$(6.12) \quad \begin{aligned} \varepsilon &= \frac{-2 \cdot P(\nabla h)}{g(\nabla h, \nabla h) - m_{\alpha\beta} \cdot g(\nabla h, W^\alpha)g(\nabla h, W^\beta)}, \\ \lambda_k &= \frac{2 \cdot m_{k\ell} \cdot g(\nabla h, W^\ell)P(\nabla h)}{g(\nabla h, \nabla h) - m_{\alpha\beta} \cdot g(\nabla h, W^\alpha)g(\nabla h, W^\beta)}. \end{aligned}$$

Proof of Theorem VI.13. Solving (6.7) gives

$$(6.13) \quad \lambda_k = -m_{k\ell} \cdot \tilde{g}(dh, \eta^\ell) \varepsilon.$$

Substituting this into (6.8) to obtain a quadratic equation in ε provides

$$\begin{aligned} &-2 \cdot m_{k\ell} \cdot g(\nabla h, W^\ell) \tilde{g}(p, \eta^k) \cdot \varepsilon \\ &+ 2 \cdot \tilde{g}(p, dh) \cdot \varepsilon + m_{i\ell} m_{jr} m^{ij} \cdot g(\nabla h, W^\ell) g(\nabla h, W^r) \cdot \varepsilon^2 \\ &- 2 \cdot m_{k\ell} \cdot g(\nabla h, W^\ell) g(\nabla h, W^k) \cdot \varepsilon^2 + g(\nabla h, \nabla h) \cdot \varepsilon^2 = 0. \end{aligned}$$

Notice that $\tilde{g}(p, \eta^k) = P(W^k)$, $\tilde{g}(p, dh) = P(\nabla h)$, and $m_{i\ell} m_{jr} m^{ij} = m_{\ell r}$. Using this and collecting in terms of powers of ε , we obtain

$$\begin{aligned} 0 &= [2 \cdot P(\nabla h) - 2 \cdot m_{k\ell} \cdot g(\nabla h, W^\ell) P(W^k)] \varepsilon \\ &+ [m_{\ell r} \cdot g(\nabla h, W^\ell) g(\nabla h, W^r) - 2 \cdot m_{k\ell} \cdot g(\nabla h, W^\ell) g(\nabla h, W^k) + g(\nabla h, \nabla h)] \varepsilon^2. \end{aligned}$$

This has the trivial root $\varepsilon = 0$ as well as the nontrivial root (6.10). Finally, (6.11) comes from combining (6.10) and (6.13). \square

It is possible to rewrite (6.9) in a form that is strikingly similar to the unconstrained case, (6.3). To do this, we need the following linear algebra lemma.

Lemma VI.15. *Let $(V, \langle \cdot, \cdot \rangle)$ be an inner product space and $\{v^i\} \subset V$ be a basis. Define the matrix $m^{ij} = \langle v^i, v^j \rangle$. Then we have the following property:*

$$\langle x, y \rangle = m_{ij} \langle x, v^i \rangle \langle y, v^j \rangle.$$

Proof. This proof proceeds by expanding x and y in terms of the basis: $x = x_i v^i$ and $y = y_j v^j$. Taking their inner product yields

$$\langle x, y \rangle = x_i y_j \langle v^i, v^j \rangle = x_i y_j m^{ij}.$$

Expanding the other side gives

$$\begin{aligned} m_{ij} \langle x, v^i \rangle \langle y, v^j \rangle &= m_{ij} x_k \langle v^k, v^i \rangle y_\ell \langle v^\ell, v^j \rangle \\ &= x_k y_\ell m_{ij} m^{ki} m^{\ell j}. \end{aligned}$$

The two sides agree because $m_{ij} m^{ki} m^{\ell j} = m^{k\ell}$. \square

Proposition VI.16. *The elastic nonholonomic impact given in Theorem VI.13 can be written as*

$$(6.14) \quad p^+ = p^- - 2 \cdot \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} \pi_{\mathcal{D}}^* dh,$$

where $\pi_{\mathcal{D}} : TQ \rightarrow \mathcal{D}$ is the orthogonal projection.

Proof. By Lemma VI.15, $m_{\alpha\beta} \cdot g(\nabla h, W^\alpha) g(\nabla h, W^\beta) = g(\pi_{\mathcal{D}^\perp} \nabla h, \pi_{\mathcal{D}^\perp} \nabla h)$. Due to the fact that orthogonal projections are self-adjoint, we get that

$$g(\nabla h, \nabla h) - m_{\alpha\beta} \cdot g(\nabla h, W^\alpha) g(\nabla h, W^\beta) = g(\pi_{\mathcal{D}} \nabla h, \pi_{\mathcal{D}} \nabla h) = dh(\pi_{\mathcal{D}} \nabla h).$$

By similar reasoning, we have $\nabla h - m_{\alpha\beta} g(\nabla h, W^\beta) W^\alpha = \pi_{\mathcal{D}} \nabla h$. This allows for (6.9) to be written as

$$\begin{aligned} p^+ &= p^- + 2 \cdot m_{k\ell} \cdot g(\nabla h, W^\ell) \cdot \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} \eta^k - 2 \cdot \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} dh \\ &= p^- - 2 \cdot \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} (dh - m_{k\ell} \cdot dh(W^\ell) \eta^k) \\ &= p^- - 2 \cdot \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} \pi_{\mathcal{D}}^* dh, \end{aligned}$$

which matches (6.14). \square

Definition VI.17. The impact map given by (6.10) and (6.11) provides the *global* impact map $\Delta^{\mathcal{E}} : S^* \rightarrow T^*Q$ while the map given by (6.12) provides the *restricted* impact map $\Delta^{\mathcal{D}} : S_{\mathcal{D}}^* \rightarrow \mathcal{D}^*$.

The global impact map serves the same purpose as the global nonholonomic vector field. In the same way that the global nonholonomic vector field made computing the divergence straight-forward, the global impact map will make the corresponding “hybrid Jacobian” easier to compute in Chapter VII.

Plastic Impacts

Unconstrained elastic impacts are defined to be variational while nonholonomic elastic impacts are defined to satisfy the Lagrange-d’Alembert principle. Likewise, since unconstrained plastic impacts are given by an orthogonal projection, nonholonomic plastic impacts will be too.

Given the unconstrained impact, $\dot{q}^+ = \delta(q, \dot{q})$, it is not generally true that $\delta(q, \dot{q}) \in \mathcal{D}_q$. In order to enforce the constraint, we orthogonally project δ onto the subspace \mathcal{D} . This inspires the following definition.

Definition VI.18. Let (Q, g) be a Riemannian manifold and (Q, L, S, \mathcal{D}) be a corresponding natural nonholonomic Lagrangian. Denote $\pi_{\mathcal{D}} : TQ \rightarrow \mathcal{D}$ as the orthogonal (with respect to g) projection onto the constraint distribution. Then the plastic nonholonomic impact map is given by

$$(6.15) \quad (q, \dot{q}) \mapsto (q, \pi_{\mathcal{D}} \circ \delta(q, \dot{q})),$$

where δ is the variational impact given by (6.3).

So long as the Lagrangian is natural, we can determine an explicit formula for the plastic impact, $\pi_{\mathcal{D}} \circ \delta$.

	Elastic wall	Plastic wall
Elastic Constraint	Elastic	Impact plastic
Plastic Constraint	Nonholonomic plastic	Totally plastic

Table 6.1: Four types of nonholonomic impact maps.

Proposition VI.19. *The plastic nonholonomic impact map $\pi_{\mathcal{D}} \circ \delta$ is given by*

$$(6.16) \quad \dot{q} \mapsto \dot{q} - 2 \cdot \frac{dh(\dot{q})}{g(\nabla h, \nabla h)} (\nabla h - m_{ij} \eta^i (\nabla h) W^j).$$

Proof. This follows from the fact that the orthogonal projection is given by

$$\pi_{\mathcal{D}}(\dot{q}) = \dot{q} - m_{ij} \eta^i (\dot{q}) W^j,$$

and the fact that $\pi_{\mathcal{D}}(\dot{q}^-) = \dot{q}^-$ since the constraints are assumed to be satisfied before impact. □

To contrast the above with (6.14), the plastic impact map is

$$p^+ = p^- - 2 \cdot \frac{P(\nabla h)}{dh(\nabla h)} \pi_{\mathcal{D}}^* dh.$$

This plastic impact map is a projection onto \mathcal{D}^* while the plastic map given by (6.5) is a projection onto TS . During a nonholonomic impact, the reaction impulses can be viewed to fall into two categories: impulses from the constraints, η^α , and impulses from the wall, dh . As such, we can think of each being elastic or plastic (Definition VI.18 is for an elastic impact and plastic constraints). This prompts the following definition based on the four options in Table 6.1.

Definition VI.20. Let $\Delta : S^* \rightarrow T^*Q$ be an impact map of the form $\Delta(x, p) = (x, \delta(x, p))$. Δ is called

- Elastic if

$$\delta(x, p) = p - 2 \cdot \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} \pi_{\mathcal{D}}^* dh,$$

- Nonholonomic plastic if

$$\delta(x, p) = p - 2 \cdot \frac{P(\nabla h)}{dh(\nabla h)} \pi_{\mathcal{D}}^* dh,$$

- Impact plastic if

$$\delta(x, p) = p - \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} \pi_{\mathcal{D}}^* dh,$$

- Totally plastic if

$$\delta(x, p) = p - \frac{P(\nabla h)}{dh(\nabla h)} \pi_{\mathcal{D}}^* dh.$$

6.2 Regularity of Mechanical Hybrid Systems

This section proves that elastic mechanical hybrid systems are *smooth* as per Definition II.17. We first prove for unconstrained systems and nonholonomic systems immediately follow.

Proposition VI.21. *Let $\mathcal{H} = (T^*Q, S^*, X_H, \Delta)$ be a hybrid dynamical system induced by the hybrid Hamiltonian (Q, H, S) . Then \mathcal{H} is smooth.*

Proof. We must show that \mathcal{H} satisfies (H.1)-(H.5), (A.1), and the transverse condition. Conditions (H.1)-(H.4) are easily verified. For (H.5), we notice that

$$S^* = \{(q, p) \in T^*Q : h(q) = 0, P(\nabla h) < 0\},$$

$$\Delta(S^*) = \{(q, p) \in T^*Q : h(q) = 0, P(\nabla h) > 0\},$$

$$\overline{S^*} \cap \overline{\Delta(S^*)} = \{(q, p) \in T^*Q : h(q) = P(\nabla h) = 0\}.$$

Therefore $S^* \cap \Delta(S^*) = \emptyset$ and $\overline{S^*} \cap \overline{\Delta(S^*)}$ has codimension 2 so (H.5) is satisfied.

For (A.1) assume that $(q(0), p(0)) \in \overline{S^*} \cap \overline{\Delta(S^*)} = \overline{S^*} \setminus S^*$ and that there exists $\varepsilon > 0$ such that for all $\delta \in (0, \varepsilon)$, we have $(q(\delta), p(\delta)) \in S^*$. Since $q(\delta) \in S$ and $P(\nabla h)(q(\delta), p(\delta)) < 0$, $q(t)$ must intersect S transversely at δ . This leads to a contradiction.

To finish the proof, we need to show that for $(q, p) \in S^*$, we have the direct sum:

$$T_{(q,p)}T^*Q = T_{(q,p)}S^* \oplus X_H \cdot \mathbb{R},$$

i.e. X_H is not tangent to S^* . This follows from similar reasoning as for (A.1); let $\gamma(t)$ be a base curve of X_H , then γ intersects S transversely. \square

The above proof works exactly the same for nonholonomic hybrid systems.

Corollary VI.22. *Let $\mathcal{H} = (\mathcal{D}^*, S_{\mathcal{D}}^*, X_H^{\mathcal{D}}, \Delta^{\mathcal{D}})$ be a hybrid dynamical systems induced by the hybrid nonholonomic system (Q, H, S, \mathcal{D}^*) . Then \mathcal{H} is smooth.*

6.3 Systems with Symmetry

There are many different notions of symmetry in a hybrid system. Each of the flow, the impact surface, the impact map, or any combination of the above can experience a symmetry. When all three have the same symmetry, everything is straightforward [4]. We will be interested in the case where not all objects obey the same symmetry.

6.3.1 Momentum Equation

Before we discuss symmetries, we first examine how momentum maps evolve with time. Let $X \in \mathfrak{X}(Q)$ be a smooth vector field and call its momentum $P : T^*Q \rightarrow \mathbb{R}$.

Proposition VI.23. *Let $X \in \mathfrak{X}(Q)$ be a smooth vector field. Then the (hybrid) dynamics of its momentum is given by*

$$\begin{cases} \dot{P}(X) = \{H, P(X)\} - m_{\alpha\beta} \{H, P(W^\alpha)\} \eta^\beta(X), & h(q) \neq 0, \\ P^+(X) = P^-(X) - 2 \cdot \frac{P(\pi_{\mathcal{D}} \nabla h)}{dh(\pi_{\mathcal{D}} \nabla h)} \cdot dh(\pi_{\mathcal{D}} X), & h(q) = 0. \end{cases}$$

Another way to view the above is to use the linearity of momentum to transfer

the dynamics to X :

$$(6.17) \quad \begin{cases} \dot{P}(X) = \{H, P(\pi_{\mathcal{D}}X)\}, & h(q) \neq 0, \\ P^+(X) = P\left(X - 2 \cdot \frac{dh(\pi_{\mathcal{D}}X)}{dh(\pi_{\mathcal{D}}\nabla h)} \pi_{\mathcal{D}}\nabla h\right), & h(q) = 0. \end{cases}$$

The momentum dynamics (6.17) immediately show that each of the $P(W^\alpha)$ are conserved.

This also provides insight into Zeno states. Notice that a Zeno state occurs when $P(\nabla h) = 0$ during impact [74], i.e. Zeno states in mechanical systems occur during grazing impacts (the converse is not necessarily true). This says that understanding Zeno solutions requires understanding the following dynamics:

$$\begin{cases} \dot{P}(\nabla h) = \{H, P(\pi_{\mathcal{D}}\nabla h)\}, & h(q) \neq 0, \\ P^+(\nabla h) = P(\nabla h - 2 \cdot \pi_{\mathcal{D}}\nabla h), & h(q) = 0. \end{cases}$$

6.3.2 The Mechanical Connection

Studying the effect of symmetries in hybrid systems will, in effect, by studying (6.17) where $X = \xi_Q$ is the infinitesimal generator for some group action. In this sense $P(\xi_Q) = \hat{J}(\xi)$ will become the momentum map, recall Definition III.18. Therefore, we can view (6.17) as dynamics on \mathfrak{g}^* . In order to understand the impacts on \mathfrak{g}^* , we need to provide this space with a metric. This will be accomplished with the following two objects: the locked inertia tensor and the mechanical connection (see §3.12 in [16]).

Definition VI.24. Let (Q, g) be a Riemannian manifold and suppose that g is invariant under the group action. For each $q \in Q$, we define the locked inertia tensor to be the map $\mathbb{I}(q) : \mathfrak{g} \rightarrow \mathfrak{g}^*$ defined by

$$\mathbb{I}\eta(\zeta) = g(\eta_Q(q), \zeta_Q(q)).$$

Definition VI.25. Let $J : T^*Q \rightarrow \mathfrak{g}^*$ be the momentum map in Corollary III.21. The mechanical connection is the map $\mathcal{A}_s : T^*Q \rightarrow \mathfrak{g}$ given by

$$\mathcal{A}_s(q, p) = \mathbb{I}(q)^{-1}J(q, p).$$

As will be seen in §6.3.4, the mechanical connection is used to describe how the momentum changes across an impact.

6.3.3 Hybrid Reduction

We review Chapter 3 in [4] which is the case where the flow, impact surface, and impact map all satisfy the symmetries.

Definition VI.26. Let (M, S, X, Δ) be a hybrid dynamical system. A group action $\Phi : G \rightarrow \text{Diff}(M)$ is called a hybrid group action if $\Phi|_S$ is an action on S and $\Delta \circ \Phi_g|_S = \Phi_g \circ \Delta$.

Notice that this requires the impact surface, S , to be symmetric and for the impact map to be G -equivariant.

Definition VI.27. An Ad^* -equivariant momentum map $J : M \rightarrow \mathfrak{g}^*$ is called hybrid Ad^* -equivariant if $J \circ \Delta(s) = J(s)$ for all $s \in S$.

The above states that in addition to being Ad^* -equivariant, the impacts must also preserve the momentum.

Definition VI.28. The tuple $(M, \omega, S, \Delta, \Phi, J)$ is called a hybrid Hamiltonian G -space if

1. (M, ω, Φ, J) forms a Hamiltonian G -space where (M, ω) is a symplectic manifold, $\Phi : G \rightarrow \text{Diff}(M)$ is a symplectic group action with Ad^* -equivariant momentum map J ,

2. Φ is a hybrid group action, and
3. J is a hybrid Ad^* -equivariant momentum map.

This allows us to carry out hybrid reduction as described by Theorem 3.4 and 3.5 in [4]. First, we define what a hybrid regular value is.

Definition VI.29. Let $J : M \rightarrow \mathfrak{g}^*$ be a momentum map. Suppose that $\mu \in \mathfrak{g}^*$ is a regular value of J . We say that μ is a hybrid regular value if it is also a regular value of $J|_S$.

Theorem VI.30. Let $(M, \omega, S, \Delta, \Phi, J)$ be a hybrid Hamiltonian G -space with a G -invariant Hamiltonian $H : M \rightarrow \mathbb{R}$. Assume $\mu \in \mathfrak{g}^*$ is a hybrid regular value of the hybrid Ad^* -equivariant momentum map J and the action of G_μ on $J^{-1}(\mu)$ is free, proper, and hybrid. Then there is a reduced Hamiltonian hybrid system:

$$(M_\mu, S_\mu, X_\mu, \Delta_\mu),$$

where $M_\mu = \pi_\mu(J^{-1}(\mu))$ as in Theorem III.23, $S_\mu = \pi_\mu(S \cap J^{-1}(\mu))$, Δ_μ satisfies $\pi_\mu^* \Delta_\mu = \Delta$, and X_μ is defined by

$$dH_\mu = i_{X_\mu} \omega_\mu, \quad H_\mu \circ \pi_\mu = H \circ \iota_\mu.$$

6.3.4 Nonsymmetric Hybrid Impacts

In order to reduce the hybrid dynamics as in Theorem VI.30, it is needed that S inherits the same symmetry as Q and that the impact preserves the momentum. However, in practice this is not always the case. Assume that (M, ω, Φ, J) is a Hamiltonian G -space but is not hybrid. This means that momentum will not be conserved across impacts; although during the continuous portion, we still have $\dot{J} = 0$.

Now we are not assuming that the momentum is hybrid Ad^* -equivariant which means that the momentum can change across impacts.

Theorem VI.31. *Let $\eta : S \rightarrow \mathfrak{g}^*$ be such that $\eta(\xi) = dh(\xi_Q)$. Then the impact map for the momentum is given by*

$$(6.18) \quad J \circ \Delta = J - 2 \frac{\eta(\mathcal{A}_s)}{\eta(\mathbb{I}^{-1}\eta)} \eta,$$

where $\mathbb{I} : \mathfrak{g} \rightarrow \mathfrak{g}^*$ is the locked inertia tensor and $\mathcal{A}_s : M \rightarrow \mathfrak{g}$ is the mechanical connection.

Remark VI.32. When the group action is hybrid, we see that $\eta = 0$ and thus momentum is conserved across impacts.

In the nonsymmetric case, the flow between impacts is described on the reduced space M_μ but jumps to a different leaf, $M_{\delta(\mu)}$, upon impact. This induces dynamics between the different constant momentum leaves. When the impact dynamics (6.18) do not depend on the configuration, the dynamics end up being periodic.

Proposition VI.33. *Suppose that $\eta \neq 0$ and (6.18) does not depend on $q \in Q$. Then $J \circ \Delta = J$ if and only if $\eta(\mathbb{I}^{-1}J) = 0$. Moreover, whenever (6.18) does not depend on $q \in Q$, we have that $J \circ \Delta^2 = J$.*

Proof. The first part is just stating that $\eta(\mathcal{A}_s) = 0$. For the second, we notice that $J \circ \Delta = J + J(\lambda)\eta$ where $\lambda \in \mathfrak{g}$ and $\eta(\lambda) = -2$. Computing:

$$\begin{aligned} J \circ \Delta^2 &= J \circ \Delta + J(\lambda)\eta \circ \Delta \\ &= J + J(\lambda)\eta + J(\lambda) [\eta + \eta(\lambda)\eta] \\ &= J + J(\lambda)\eta - J(\lambda)\eta \\ &= J. \end{aligned}$$

Therefore, the momentum impact equation is 2-periodic. □

Nonholonomic Impacts

As a last note, we modify (6.18) to nonholonomic mechanical systems. In order to do this, we first construct a metric on \mathfrak{g}^* . Let $q \in Q$ and define $\langle\langle \cdot, \cdot \rangle\rangle_q : \mathfrak{g}^* \times \mathfrak{g}^* \rightarrow \mathbb{R}$ such that $\langle\langle \eta, \mu \rangle\rangle_q = \eta(\mathbb{I}(q)^{-1}\mu)$. Using this language with Theorem VI.13, we obtain the following for the momentum impact equation.

Theorem VI.34. *Let $\eta : S \rightarrow \mathfrak{g}^*$ such that $\eta(\xi) = dh(\xi_Q)$ and $\beta^k : S \rightarrow \mathfrak{g}^*$ such that $\beta^k(\xi) = \eta^k(\xi_Q)$. Then the nonholonomic impact map is given by*

$$(6.19) \quad J \circ \Delta^{\mathcal{C}} = J + \lambda_k \cdot \beta^k + \varepsilon \cdot \eta,$$

where

$$\varepsilon = \frac{-2 \cdot \eta(\mathcal{A}_s)}{\langle\langle \eta, \eta \rangle\rangle - m_{k\ell} \cdot \langle\langle \eta, \beta^k \rangle\rangle \langle\langle \eta, \beta^\ell \rangle\rangle},$$

$$\lambda_k = \frac{2 \cdot m_{k\ell} \cdot \langle\langle \eta, \beta^\ell \rangle\rangle \cdot \eta(\mathcal{A}_s)}{\langle\langle \eta, \eta \rangle\rangle - m_{ij} \cdot \langle\langle \eta, \beta^i \rangle\rangle \langle\langle \eta, \beta^j \rangle\rangle}.$$

Here the matrix $(m_{ij}) = (m^{ij})^{-1}$ where $m^{ij} = \langle\langle \beta^i, \beta^j \rangle\rangle$.

The expression (6.19) can be simplified when we use the nonholonomic momentum map.

Proposition VI.35. *Using the notation in Theorem VI.34, equation (6.19) becomes*

$$(6.20) \quad J^{\mathcal{D}} \circ \Delta^{\mathcal{D}} = J^{\mathcal{D}} + \varepsilon \cdot \eta,$$

where it is understood that η is restricted to $(\mathfrak{g}^q)^*$.

Proof. Let $\xi \in \mathfrak{g}^q$. Then we have $\beta^k(\xi) = \eta^k(\xi_Q) = 0$. □

The multiplier, ε , has a clean interpretation in the nonholonomic case.

Definition VI.36. Define the constrained locked inertia tensor as $\mathbb{I}^{\mathcal{D}} : \mathfrak{g}^{\mathcal{D}} \rightarrow (\mathfrak{g}^{\mathcal{D}})^*$

by

$$\mathbb{I}^{\mathcal{D}}(\xi)(\zeta) = g(\xi_Q, \zeta_Q).$$

Likewise, define a metric $\langle\langle \cdot, \cdot \rangle\rangle_{\mathcal{D}} : (\mathfrak{g}^{\mathcal{D}})^* \times (\mathfrak{g}^{\mathcal{D}})^* \rightarrow \mathbb{R}$ by $\mathbb{I}^{\mathcal{D}}$.

Theorem VI.37 (Nonholonomic Hybrid Momentum Equation). *Let $\eta : S \rightarrow (\mathfrak{g}^{\mathcal{D}})^*$ such that $\eta(\xi) = dh(\xi_Q)$ for all $\xi \in \mathfrak{g}^{\mathcal{D}}$. The nonholonomic impact map is given by*

$$(6.21) \quad J^{\mathcal{D}} \circ \Delta^{\mathcal{D}} = J^{\mathcal{D}} - 2 \cdot \frac{\langle\langle \eta, J^{\mathcal{D}} \rangle\rangle_{\mathcal{D}}}{\langle\langle \eta, \eta \rangle\rangle_{\mathcal{D}}} \eta.$$

Proof. Proving this requires showing that

$$\langle\langle \eta, \eta \rangle\rangle - m_{k\ell} \cdot \langle\langle \eta, \beta^k \rangle\rangle \langle\langle \eta, \beta^\ell \rangle\rangle = \langle\langle \eta|_{\mathcal{D}}, \eta|_{\mathcal{D}} \rangle\rangle_{\mathcal{D}}.$$

This follows from Lemma VI.15. □

6.4 Example: Nonholonomic Billiards - Rolling Disk

The vertical rolling disk is a simple example for illustrating nonholonomic dynamics and we will examine the billiard problem where the particle moves as the rolling disk in between impacts [27]. We will work with the somewhat nonphysical case where the disk is not permitted to tilt. The configuration space and local coordinates for the rolling disk are given by $q = (x, y, \theta, \varphi) \in Q = \mathbb{R}^2 \times S^1 \times S^1$, denoting the position of the contact point, the rotation angle of the disk, and the orientation of the disk, respectively, cf. Figure 5.2.

6.4.1 Dynamic nonholonomic equations

The Lagrangian for the vertical disk is taken to be the kinetic energy (no potential force will be included), i.e.

$$L = \frac{1}{2}m(\dot{x}^2 + \dot{y}^2) + \frac{1}{2}I\dot{\theta}^2 + \frac{1}{2}J\dot{\varphi}^2.$$

Here, m is the mass of the disk, I is the moment of inertia of the disk about the axis perpendicular to the plane of the disk, and J is the moment of inertia about an axis in the plane of the disk.

If $R > 0$ is the radius of the disk, the nonholonomic constraints for rolling *without* slipping are

$$(6.22) \quad \begin{aligned} \dot{x} &= R\dot{\theta} \cos \varphi, \\ \dot{y} &= R\dot{\theta} \sin \varphi. \end{aligned}$$

This can be expressed as $\mathcal{D} = \ker \eta^1 \cap \ker \eta^2$ where

$$(6.23) \quad \begin{aligned} \eta^1 &= dx - R(\cos \varphi) d\theta, \\ \eta^2 &= dy - R(\sin \varphi) d\theta. \end{aligned}$$

The equations of motions for this (uncontrolled) system are given by two dynamic equations and the two constraining equations (see §1.4 of [16] for the derivation).

$$\begin{aligned} J\ddot{\varphi} &= 0, \\ (I + mR^2)\ddot{\theta} &= 0, \\ \dot{x} &= R\dot{\theta} \cos \varphi, \\ \dot{y} &= R\dot{\theta} \sin \varphi. \end{aligned}$$

These equations can be easily integrated. Let the initial conditions be $(x_0, y_0, \theta_0, \varphi_0)$ and call $\omega = \dot{\varphi}$, $\Omega = \dot{\theta}$, which are constants. Then, the (continuous) equations of motion are

$$\begin{aligned} \varphi &= \omega t + \varphi_0, \\ \theta &= \Omega t + \theta_0, \\ x &= \frac{\Omega}{\omega} R \sin(\omega t + \varphi_0) - \frac{\Omega}{\omega} R \sin \varphi_0 + x_0, \\ y &= -\frac{\Omega}{\omega} R \cos(\omega t + \varphi_0) + \frac{\Omega}{\omega} R \cos \varphi_0 + y_0. \end{aligned}$$

6.4.2 Impact map

Below, we will construct both the elastic (6.9) and nonholonomic plastic (6.16) impact maps. Before we proceed with their derivations, we first lay out the impact

surface, S . Suppose that the edge of the table is described by the level-set $S' = \{(x, y) \in \mathbb{R}^2 : \tilde{h}(x, y) = 0\}$ for some smooth $\tilde{h} : \mathbb{R}^2 \rightarrow \mathbb{R}$. Then an impact occurs not when $\tilde{h}(x, y) = 0$, rather when

$$h^+(x, y, \theta\varphi) := \tilde{h}(x + R \cos \varphi, y + R \sin \varphi) = 0,$$

for when the leading edge makes contact and

$$h^-(x, y, \theta\varphi) := \tilde{h}(x - R \cos \varphi, y - R \sin \varphi) = 0,$$

for when the tailing edge makes contact. Each of these impacts will be the same (modulo a sign difference). As such, we will only construct the impact maps for when the leading edge makes contact, i.e. $h = h^+$.

Both impact maps require the constraining vector fields, W^i , as well as the constraint mass matrix, m_{ij} .

$$W^1 = \frac{1}{m} \frac{\partial}{\partial x} - \frac{R}{I} \cos \varphi \frac{\partial}{\partial \theta},$$

$$W^2 = \frac{1}{m} \frac{\partial}{\partial y} - \frac{R}{I} \sin \varphi \frac{\partial}{\partial \theta},$$

and

$$(m_{ij}) = \begin{bmatrix} m - K \cos^2 \varphi & -K \sin \varphi \cos \varphi \\ -K \sin \varphi \cos \varphi & m - K \sin^2 \varphi \end{bmatrix}, \quad K = \frac{m^2 R^2}{I + m R^2}.$$

Elastic Impact

Using (6.9), the elastic nonholonomic impact map is given by

$$\begin{aligned} \dot{x}^+ &= \dot{x}^- + \frac{1}{m} \lambda_1 + \frac{1}{m} \frac{\partial \tilde{h}}{\partial x} \varepsilon \\ \dot{y}^+ &= \dot{y}^- + \frac{1}{m} \lambda_2 + \frac{1}{m} \frac{\partial \tilde{h}}{\partial y} \varepsilon \\ \dot{\theta}^+ &= \dot{\theta}^- - \frac{R}{I} \cos \varphi \cdot \lambda_1 - \frac{R}{I} \sin \varphi \cdot \lambda_2 \\ \dot{\varphi}^+ &= \dot{\varphi}^- + \frac{R}{J} \left(\frac{\partial \tilde{h}}{\partial y} \cos \varphi - \frac{\partial \tilde{h}}{\partial x} \sin \varphi \right) \varepsilon, \end{aligned}$$

where the multipliers are chosen such that the constraint are preserved and energy is conserved, which are given by (6.12).

Nonholonomic Plastic Impact

Using (6.3) to determine the (pre-) impact map, we see that

$$(6.24) \quad \begin{aligned} \dot{x}^+ &= \dot{x}^- + \frac{C}{m} \frac{\partial \tilde{h}}{\partial x}, \\ \dot{y}^+ &= \dot{y}^- + \frac{C}{m} \frac{\partial \tilde{h}}{\partial y}, \\ \dot{\theta}^+ &= \dot{\theta}^-, \\ \dot{\varphi}^+ &= \dot{\varphi}^- + \frac{C}{J} R \left(\cos \varphi \frac{\partial \tilde{h}}{\partial y} - \sin \varphi \frac{\partial \tilde{h}}{\partial x} \right). \end{aligned}$$

Here, the number $C = C(\varphi, \dot{x}, \dot{y}, \dot{\varphi})$ has the value

$$C = \frac{-2 [\tilde{h}_x \dot{x} + \tilde{h}_y \dot{y} + R (\tilde{h}_y \cos \varphi - \tilde{h}_x \sin \varphi) \dot{\varphi}]}{\frac{1}{m} [\tilde{h}_x^2 + \tilde{h}_y^2] + \frac{1}{J} R^2 (\tilde{h}_y \cos \varphi - \tilde{h}_x \sin \varphi)^2},$$

where \tilde{h}_x and \tilde{h}_y are the x and y partial derivatives of \tilde{h} . Notice that at impact the rotation angle of the disk, $\dot{\theta}$, is unchanged. This is because the constraint of no sliding has not yet been imposed. In order to apply the constraints, we compute $\pi_{\mathcal{D}}$.

We can write $\pi_{\mathcal{D}}$ as a matrix with coordinates $(\dot{x}, \dot{y}, \dot{\theta}, \dot{\varphi})$:

$$(6.25) \quad (I + mR^2) \cdot \pi_{\mathcal{D}} = \begin{bmatrix} mR^2 \cos^2 \varphi & mR^2 \sin \varphi \cos \varphi & IR \cos \varphi & 0 \\ mR^2 \sin \varphi \cos \varphi & mR^2 \sin^2 \varphi & IR \sin \varphi & 0 \\ mR \cos \varphi & mR \sin \varphi & I & 0 \\ 0 & 0 & 0 & I + mR^2 \end{bmatrix}.$$

The impact map is then given by composing (6.25) with (6.24).

Remark VI.38. The projection map $\pi_{\mathcal{D}}$ given by (6.25) can be used to computing the two remaining types of impacts given by Table 6.1. The data we need, in particular,

are ∇h , $\pi_{\mathcal{D}}\nabla h$, and $\pi_{\mathcal{D}}^*dh$. These are:

$$\begin{aligned}\nabla h &= \frac{1}{m} \frac{\partial \tilde{h}}{\partial x} \frac{\partial}{\partial x} + \frac{1}{m} \frac{\partial \tilde{h}}{\partial y} \frac{\partial}{\partial y} + \frac{R}{I} \left(\frac{\partial \tilde{h}}{\partial y} \cos \varphi - \frac{\partial \tilde{h}}{\partial x} \sin \varphi \right) \frac{\partial}{\partial \varphi}, \\ \pi_{\mathcal{D}}\nabla h &= \frac{R^2}{2(I+mR^2)} \left[\frac{\partial \tilde{h}}{\partial x} (1 + \cos 2\varphi) + \frac{\partial \tilde{h}}{\partial y} \sin 2\varphi \right] \frac{\partial}{\partial x} \\ &\quad + \frac{R^2}{2(I+mR^2)} \left[\frac{\partial \tilde{h}}{\partial y} (1 - \sin 2\varphi) + \frac{\partial \tilde{h}}{\partial x} \cos 2\varphi \right] \frac{\partial}{\partial y} \\ &\quad + \frac{R}{I+mR^2} \left[\frac{\partial \tilde{h}}{\partial x} \cos \varphi + \frac{\partial \tilde{h}}{\partial y} \sin \varphi \right] \frac{\partial}{\partial \theta} + \frac{R}{I} \left[\frac{\partial \tilde{h}}{\partial y} \cos \varphi - \frac{\partial \tilde{h}}{\partial x} \sin \varphi \right] \frac{\partial}{\partial \varphi}, \\ \pi_{\mathcal{D}}^*dh &= \frac{mR^2}{2(I+mR^2)} \left[\frac{\partial \tilde{h}}{\partial x} (1 + \cos 2\varphi) + \frac{\partial \tilde{h}}{\partial y} \sin 2\varphi \right] dx \\ &\quad + \frac{mR^2}{2(I+mR^2)} \left[\frac{\partial \tilde{h}}{\partial y} (1 - \sin 2\varphi) + \frac{\partial \tilde{h}}{\partial x} \cos 2\varphi \right] dy \\ &\quad + \frac{IR}{I+mR^2} \left[\frac{\partial \tilde{h}}{\partial x} \cos \varphi + \frac{\partial \tilde{h}}{\partial y} \sin \varphi \right] d\theta + R \left[\frac{\partial \tilde{h}}{\partial y} \cos \varphi - \frac{\partial \tilde{h}}{\partial x} \sin \varphi \right] d\varphi.\end{aligned}$$

6.4.3 The Momentum Equation

Next, we derive the nonholonomic hybrid momentum equation (6.21). When analyzing the impacts, two different types can happen: the leading or tailing edge can make contact. In the same spirit as the preceding section, we will only consider the leading edge impact as the other case is nearly identical. In this example, $Q = \text{SE}_2 \times S^1$ which is itself a Lie group. Therefore, we can let $G = Q$ and let it act via left-translations. We identify $\mathfrak{g} \cong \mathbb{R}^4$ via the local coordinates $(x, y, \varphi; \theta)$. The corresponding infinitesimal generators are

$$(6.26) \quad \begin{aligned}(1, 0, 0, 0)_Q &= \frac{\partial}{\partial x}, \\ (0, 1, 0, 0)_Q &= \frac{\partial}{\partial y}, \\ (0, 0, 1, 0)_Q &= \frac{\partial}{\partial \varphi} - y \frac{\partial}{\partial x} + x \frac{\partial}{\partial y}, \\ (0, 0, 0, 1)_Q &= \frac{\partial}{\partial \theta}.\end{aligned}$$

To determine the bundle, $\mathfrak{g}^{\mathcal{D}}$, we use (6.26) to determine which elements of \mathfrak{g} map to the kernels of η^1 and η^2 in (6.23). Solving this system of equations produces

$$\mathfrak{g}^q = \text{span} \{ (y, -x, 1, 0), (R \cos \varphi, R \sin \varphi, 0, 1) \}.$$

Let us call $\xi^1 := (y, -x, 1, 0)$ and $\xi^2 := (R \cos \varphi, R \sin \varphi, 0, 1)$. To determine the constrained locked inertia tensor, $\mathbb{I}^{\mathcal{D}}$, we use the metric on Q , i.e. $\mathbb{I}^{\mathcal{D}}(\xi^i)(\xi^j) = g(\xi_Q^i, \xi_Q^j)$.

The matrix with respect to the basis $\mathcal{B} = \{\xi^1, \xi^2\}$ is

$$\mathbb{I}_{\mathcal{B}}^{\mathcal{D}} = \begin{bmatrix} J & 0 \\ 0 & mR^2 + I \end{bmatrix}.$$

The last piece of information required for the nonholonomic hybrid momentum equation is the map $\eta : \tilde{S} \rightarrow (\mathfrak{g}^{\mathcal{D}})^*$,

$$\begin{aligned} \eta(\xi^1) &= R \left(\cos \varphi \frac{\partial h}{\partial y} - \sin \varphi \frac{\partial h}{\partial x} \right), \\ \eta(\xi^2) &= R \left(\cos \varphi \frac{\partial h}{\partial x} + \sin \varphi \frac{\partial h}{\partial y} \right). \end{aligned}$$

To describe the evolution of $J^{\mathcal{D}} : T^*Q \rightarrow (\mathfrak{g}^{\mathcal{D}})^*$, let us choose a dual basis $\{\mu_1, \mu_2\}$ to $\{\xi^1, \xi^2\}$, so $\mu_i(\xi^j) = \delta_i^j$. With respect to these coordinates,

$$J_{\mathcal{B}}^{\mathcal{D}} = \begin{bmatrix} p_{\varphi} & Rp_x \cos \varphi + Rp_y \sin \varphi + p_{\theta} \end{bmatrix}.$$

By the momentum equation (3.15), the continuous evolution of the momentum is

$$\frac{d}{dt} J_{\mathcal{B}}^{\mathcal{D}} = \begin{bmatrix} 0 & Rp_y \cos \varphi - Rp_x \sin \varphi \end{bmatrix}.$$

Before we construct the impact map, we can simplify the expression for the momentum by substituting in the constraints (6.22). This simplifies the momentum to

$$(6.27) \quad J_{\mathcal{B}}^{\mathcal{D}} = \begin{bmatrix} p_{\varphi} & p_{\theta} \end{bmatrix}, \quad \frac{d}{dt} J_{\mathcal{B}}^{\mathcal{D}} = \begin{bmatrix} 0 & 0 \end{bmatrix}.$$

A slightly different derivation of (6.27) can be found in [14] and §5.6 of [16]. The impact map is linear so it can be written as $J_B^D \circ \Delta^D = J_B^D \cdot (I_2 - 2 \cdot R_B)$ where

$$(6.28) \quad R_B = \frac{1}{(mR^2 + I)c_1^2 + Jc_2^2} \begin{bmatrix} (mR^2 + I)c_1^2 & (mR^2 + I)c_1c_2 \\ Jc_1c_2 & Jc_2^2 \end{bmatrix}$$

$$c_1 = h_y \cos \varphi - h_x \sin \varphi,$$

$$c_2 = h_x \cos \varphi + h_y \sin \varphi.$$

This shows that the momentum is conserved during the continuous case while it is shuffled during impacts, which only depend on the angle the disk makes with the wall.

6.4.4 Numerical Results


Assume that the table-top is elliptical. i.e.

$$\tilde{h}(x, y) = \frac{x^2}{a^2} + \frac{y^2}{b^2} - 1.$$


We will also assume that the disk is homogeneous and thin so $I = 1/2mR^2$ and $J = 1/4mR^2$. For the remaining parameters, we would like $R \ll a, b$ so the coin has ample room to explore. The values are in the table below and are chosen to be similar to a US penny.

Parameter	Value
R	0.01 m
m	0.0025 kg
I	$1.25 \cdot 10^{-7}$ kg m ²
J	$6.25 \cdot 10^{-8}$ kg m ²
a	0.15 m or 0.20 m
b	0.20 m


Table 6.2: The chosen set of parameters for the time step example; a takes the first value for the elliptical cases and the latter for the circular ones.




(a) Elastic impacts on a circular table.



(b) Plastic impacts on a circular table.



(c) Elastic impacts on an elliptical table.



(d) Plastic impacts on an elliptical table.

Figure 6.2: The first 20 impacts for elastic/plastic impacts on a circular/elliptical table.

Plastic vs Elastic

For the simulations, Figure 6.2, the same initial conditions are chosen: $x_0 = y_0 = \theta_0 = 0$, $\varphi_0 = \pi/2$, $\dot{\theta}_0 = 10$, and $\dot{\varphi}_0 = 0.2$ rad/sec.

Hints of Chaos

In addition to comparing trajectories of the elastic and plastic impacts, we compare how changes in initial conditions propagate with time. All 100 initial conditions are taken to be those chosen in §6.4.4 except that $\dot{\theta}_0$ and $\dot{\varphi}_0$ are randomly perturbed by < 0.005 rad/sec. These results are shown in Figure 6.3.

The Momentum Equation

We finish this section with examining the momentum of the disk (6.27) and (6.28). Due to the fact that energy is conserved during both continuous and impact phases,

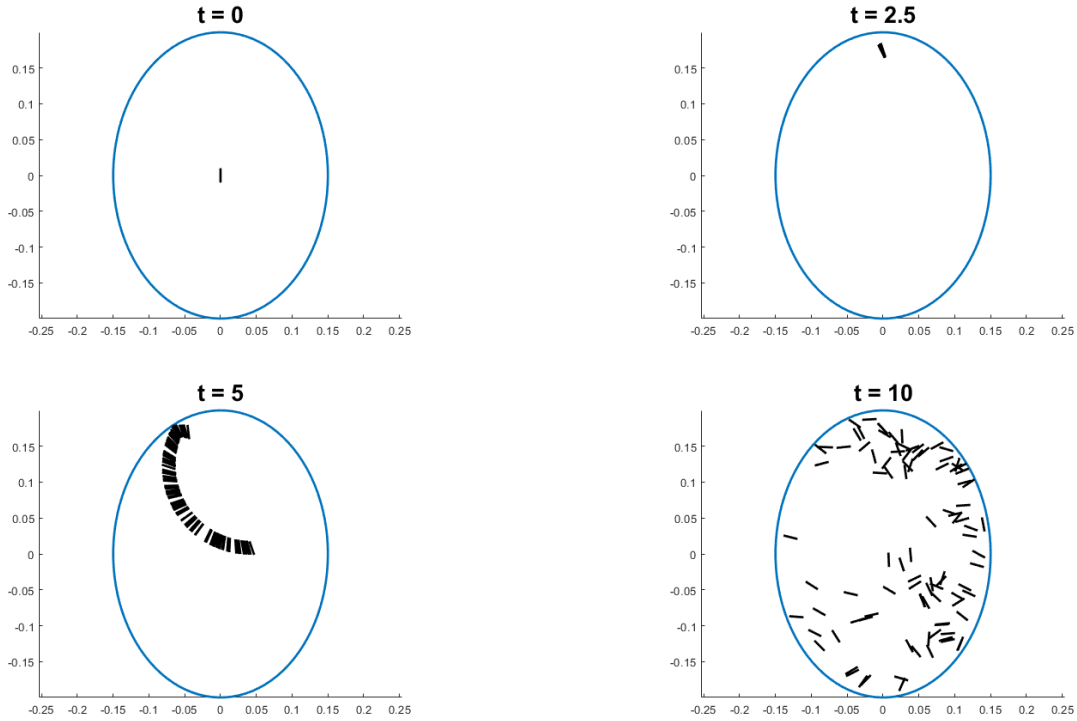


Figure 6.3: Plots of 100 different initial conditions at various times.

the momentum must remain on the ellipse of constant energy,

$$\left(\frac{mR^2 + I}{2I^2}\right)p_\theta^2 + \frac{1}{2J}p_\varphi^2 = \text{constant}.$$

Plots of the trajectories of the nonholonomic momentum are shown in Figure 6.4.




Figure 6.4: Left: The momentum for a circular table is 2-periodic which is representative of the order seen in Figure 6.2 (a). Center: The 2-periodicity is slightly destroyed with a slightly elliptical table. Right: All order in the momentum is lost which demonstrates the lack of order in Figure 6.2 (c).

CHAPTER VII

Invariant Measures in Hybrid Systems

In the same way that Chapter V examined the problem of invariant measures in nonholonomic systems, this chapter attempts to find invariant measures for hybrid dynamical systems. We first examine invariant differential forms before specializing to volume forms. This chapter concludes with a result that the hybrid (vertical) rolling disk is *always* volume-preserving independent of the choice of table-top.

We will tacitly assume that $\Delta(S) \subset M$ is also a smooth embedded submanifold.

7.1 Invariant differential forms

Before we deal with invariant volume forms, we first determine criteria for when an arbitrary differential form is preserved.

Theorem VII.1. *Let $\mathcal{H} = (M, S, X, \Delta)$ be a smooth hybrid system. Let $\varphi_t^{\mathcal{H}}$ be the hybrid flow. Then for a given $\alpha \in \Omega^m(M)$, we have $(\varphi_t^{\mathcal{H}})^* \alpha = \alpha$ if and only if $\mathcal{L}_X \alpha = 0$ and*

$$\alpha_{\Delta(x)} (\Delta_*^X \cdot v_1, \dots, \Delta_*^X \cdot v_m) = \alpha_x (v_1, \dots, v_m),$$

where Δ_*^X is the augmented differential, recall Definition IV.14.

Proof. For simplicity of calculations, we will assume that α is a 1-form. Let $x_0 \in M$,

then the condition that $(\varphi_T^{\mathcal{H}})^* \alpha = \alpha$ means

$$\alpha_{\varphi_T^{\mathcal{H}}(x_0)} \left((\varphi_T^{\mathcal{H}})_* v \right) = \alpha_{x_0}(v).$$

Choose x_0 and T such that a single impact occurs along the path $\{\varphi_t^{\mathcal{H}}(x_0) : t \in (0, T)\}$ and call this time t_1 and location y_0 , i.e. $y_0 = \varphi_{t_1}^{\mathcal{H}}(x_0) \in S$. Additionally, call $z_0 := \Delta(y_0)$ and $w_0 := \varphi_{T-t_1}^{\mathcal{H}}(z_0) = \varphi_T^{\mathcal{H}}(x_0)$. Because the vector field is transverse to S at y_0 , we can split up the tangent space at x_0 in the following way:

$$T_{x_0} M = T_{x_0}^S M \oplus X(x_0) \cdot \mathbb{R}, \quad (\varphi_{t_1})_* (T_{x_0}^S M) = T_{y_0} S.$$

To compute $(\varphi_T^{\mathcal{H}})_* v$, we split into the cases where $v \in T_{x_0}^S M$ and $v \in X(x_0) \cdot \mathbb{R}$ (which can be taken as $v = X(x_0)$ by linearity). See Figure 7.1 for an illustration of the setup.

Let $v \in T_{x_0}^S M$. Therefore, we can choose a curve $\gamma : (-\varepsilon, \varepsilon) \rightarrow M$ such that $\varphi_{t_1}(\gamma(s)) \in S$ for all $s \in (-\varepsilon, \varepsilon)$. Then $\varphi_T^{\mathcal{H}}(\gamma(s)) = \varphi_{T-t_1} \circ \Delta \circ \varphi_t(\gamma(s))$. Differentiating this provides

$$(\varphi_T^{\mathcal{H}})_* v = (\varphi_{T-t_1})_* \cdot \Delta_* \cdot (\varphi_{t_1})_* v.$$

Therefore, for $v \in T_{x_0}^S M$,

$$\alpha_{\varphi_T^{\mathcal{H}}(x_0)} \left((\varphi_T^{\mathcal{H}})_* v \right) = \alpha_{\varphi_T^{\mathcal{H}}(x_0)} \left((\varphi_{T-t_1})_* \cdot \Delta_* \cdot (\varphi_{t_1})_* v \right).$$

Which, if $\mathcal{L}_X \alpha = 0$, invariance is equivalent to $\alpha_{\Delta(y_0)}(\Delta_* \cdot v) = \alpha_{y_0}(v)$ for $v \in T_{y_0} S \subset T_{y_0} M$.

Let $v = X(x_0)$. To complete the proof, we need to show that $(\varphi_T^{\mathcal{H}})_* X(x_0) = X(\varphi_T^{\mathcal{H}}(x_0))$. Let $\gamma : (-\varepsilon, \varepsilon) \rightarrow M$ be given by $\gamma(t) = \varphi_t(x_0)$ such that $\varepsilon < t_1$ (so γ is




Figure 7.1: Diagram for the proof of Theorem VII.1. The set $\tilde{S} := \Delta(S)$ and $T_{x_0}^S M := (\varphi_{-t_1})_* T_{y_0} S$.

also the hybrid flow). Then we have

$$\begin{aligned} (\varphi_T^{\mathcal{H}})_* v &= \left. \frac{d}{dt} \right|_{t=0} \varphi_T^{\mathcal{H}} \circ \varphi_t^{\mathcal{H}}(x_0) \\ &= \left. \frac{d}{dt} \right|_{t=0} \varphi_{T+t}^{\mathcal{H}}(x_0) \\ &= X(w_0), \end{aligned}$$

which completes the proof. \square

Definition VII.2. A differential form α is called hybrid-invariant if $(\varphi_t^{\mathcal{H}})^* \alpha = \alpha$.

Let the set $\mathcal{A}_{\mathcal{H}} \subset \Omega(M)$ be all the hybrid-invariant forms.

7.1.1 A Formula for the Augmented Differential

We take a brief digression to write down an explicit formula for the augmented differential Δ_*^X . Let $v \in T_x M$ such that $h(x) = 0$ (so $x \in S$). If the impact map is $\Delta : S \rightarrow M$, then the augmented impact is given by

$$(7.1) \quad \Delta_*^X v = \frac{dh(v)}{dh(X(x))} X(\Delta(x)) + \Delta_* \left(v - \frac{dh(v)}{dh(X(x))} X(x) \right).$$

In practice computing Δ_*^X will be quite tedious. The next subsection provides necessary and sufficient conditions for $\alpha \in \mathcal{A}_{\mathcal{H}}$ without the need to explicitly compute Δ_*^X from (7.1).

7.1.2 Properties of Invariant Forms

Rather than having to compute Δ_*^X , the following theorem provides a way to test whether or not $\alpha \in \mathcal{A}_{\mathcal{H}}$ without computing Δ_*^X .

Theorem VII.3. *A differential form α is hybrid-invariant if and only if $\mathcal{L}_X\alpha = 0$ and*

$$(7.2) \quad \Delta^* \iota_{\tilde{S}}^* i_X \alpha = \iota_S^* i_X \alpha,$$

$$(7.3) \quad \Delta^* \iota_{\tilde{S}}^* \alpha = \iota_S^* \alpha,$$

where $\tilde{S} = \Delta(S)$ and $\iota_S : S \hookrightarrow M$, $\iota_{\tilde{S}} : \tilde{S} \hookrightarrow M$ are the inclusion maps.

Proof. Suppose that $\alpha \in \Omega^2(M)$ (the proof is almost identical for forms of different degrees). For $u, v \in T_{y_0}M$, decompose them in the following way:

$$u = a \cdot X(y_0) + \tilde{u}, \quad \tilde{u} \in T_{y_0}S$$

$$v = b \cdot X(y_0) + \tilde{v}, \quad \tilde{v} \in T_{y_0}S.$$

Under this decomposition, the augmented differential is

$$\Delta_*^X \cdot u = a \cdot X(z_0) + \Delta_* \tilde{u}$$

$$\Delta_*^X \cdot v = b \cdot X(z_0) + \Delta_* \tilde{v}.$$

Therefore, according to Theorem VII.1 invariance is equivalent to

$$\alpha_{y_0}(a \cdot X(y_0) + \tilde{u}, b \cdot X(y_0) + \tilde{v}) = \alpha_{z_0}(a \cdot X(z_0) + \Delta_* \tilde{u}, b \cdot X(z_0) + \Delta_* \tilde{v}).$$

Using the bi-linearity of α results in

$$\begin{aligned} 0 &= a \cdot \alpha(X(y_0), \tilde{v}) - a \cdot \alpha(X(z_0), \Delta_* \tilde{v}) \\ &\quad + b \cdot \alpha(\tilde{u}, X(y_0)) - b \cdot \alpha(\Delta_* \tilde{u}, X(z_0)) \\ &\quad + \alpha(\tilde{u}, \tilde{v}) - \alpha(\Delta_* \tilde{u}, \Delta_* \tilde{v}). \end{aligned}$$

This condition is equivalent to

$$\begin{aligned} i_{X(y_0)}\alpha(\tilde{v}) &= i_{X(z_0)}\alpha(\Delta_*\tilde{v}) \\ \alpha(\tilde{u}, \tilde{v}) &= \alpha(\Delta_*\tilde{u}, \Delta_*\tilde{v}) \end{aligned}$$

for all $\tilde{u}, \tilde{v} \in T_{y_0}S$. These are equivalent to (7.2) and (7.3). \square

It is interesting to point out that hybrid-invariance requires two additional conditions, not one. For reasons that will be apparent in §7.1.3, condition (7.2) will be called the **energy condition** while (7.3) will be called the **specular condition**.

The previous theorem provides a geometric way to check whether or not a given form is hybrid-invariant. This allows for a way to study some algebraic properties of $\mathcal{A}_{\mathcal{H}}$.

Corollary VII.4. *The set $\mathcal{A}_{\mathcal{H}} \subset \Omega(M)$ is a \wedge -subalgebra closed under d and i_X .*

Proof. If we denote $\mathcal{A} := \{\alpha \in \Omega(M) : \mathcal{L}_X\alpha = 0\}$, then it is already known that $\mathcal{A} \subset \Omega(M)$ is a \wedge -subalgebra closed under d and i_X (see Corollary 3.4.5 in [1]). Therefore, in order to prove the theorem, it suffices only to check (7.2) and (7.3). Let $\alpha, \beta \in \mathcal{A}_{\mathcal{H}}$. We only need to check that $d\alpha$, $i_X\alpha$, and $\alpha \wedge \beta$ obey (7.2) and (7.3).

Consider $i_X\alpha$. This satisfies (7.3) because α satisfies (7.2) and (7.2) is satisfied because $i_X i_X\alpha = 0$.

Consider $d\alpha$. Condition (7.3) follows from the fact that d commutes with pullbacks:

$$\Delta^* \iota_S^* d\alpha = d(\Delta^* \iota_S^* \alpha) = d(\iota_S^* \alpha) = \iota_S^* d\alpha.$$

Condition (7.2) requires Cartan's magic formula:

$$\begin{aligned} \Delta^* \iota_S^* i_X d\alpha &= \Delta^* \iota_S^* (\cancel{\mathcal{L}_X\alpha} - di_X\alpha) = -\Delta^* \iota_S^* di_X\alpha \\ &= -d(\Delta^* \iota_S^* i_X\alpha) = -d(\iota_S^* i_X\alpha) \\ &= -\iota_S^* di_X\alpha = \iota_S^* i_X d\alpha. \end{aligned}$$

Finally, consider $\alpha \wedge \beta$. The condition (7.3) holds because pullbacks distribute over the wedge product: $f^*(\alpha \wedge \beta) = f^*\alpha \wedge f^*\beta$. For condition (7.2), we see that

$$\begin{aligned}
\Delta^* \iota_{\tilde{S}}^* i_X (\alpha \wedge \beta) &= \Delta^* \iota_{\tilde{S}}^* (i_X \alpha \wedge \beta - \alpha \wedge i_X \beta) \\
&= (\Delta^* \iota_{\tilde{S}}^* i_X \alpha) \wedge (\Delta^* \iota_{\tilde{S}}^* \beta) - (\Delta^* \iota_{\tilde{S}}^* \alpha) \wedge (\Delta^* \iota_{\tilde{S}}^* i_X \beta) \\
&= (\iota_S^* i_X \alpha) \wedge (\iota_S^* \beta) - (\iota_S^* \alpha) \wedge (\iota_S^* i_X \beta) \\
&= \iota_S^* i_X (\alpha \wedge \beta),
\end{aligned}$$

which completes the proof. \square

7.1.3 Invariant Forms for Hybrid Mechanical Systems

It is known that (unconstrained) elastic mechanical impacts are symplectic [91]. In that work, the authors show symplecticity via variational principles. Here, on the contrary, we can show symplecticity directly via (7.2) and (7.3).

Theorem VII.5. *Elastic impacts are symplectic.*

Proof. The first condition is a consequence of conservation of energy:

$$\Delta^* \iota_{\tilde{S}}^* dH = \iota_S^* dH.$$

The second condition is a little more subtle. Choose coordinates such that $x^n = h$. In these coordinates,

$$\iota_{\tilde{S}}^* \omega = dx^1 \wedge dp_1 + \dots + dx^{n-1} \wedge dp_{n-1}.$$

The elastic impact map is the identity on every coordinate with the exception of p_n , but this is invisible to the restricted form $\iota_{\tilde{S}}^* \omega$. Therefore, it is preserved across impacts. \square

In order for the impacts to be symplectic, H need not be conserved, only dH . Another way of viewing this is symplecticity as a result of H being a *relative invariant*. This means that $H \circ \Delta = H + \text{constant}$ rather than $H \circ \Delta = H$, so energy can

be injected, or sapped, across impacts without breaking the symplectic structure. In particular, being symplectic is weaker than being variational.

Corollary VII.6. *Elastic impacts are volume preserving.*

Proof. Elastic impacts preserve the symplectic form and consequently preserves the volume form by Corollary VII.4. \square

Due to the fact that nonholonomic systems are generally not symplectic (this is actually a deep question, see e.g. [13] and [45] for an in-depth discussion), volume will generally not be preserved.

7.2 Hybrid-Invariant Measures

Before we address the issue of nonholonomic hybrid volume preservation, we first examine the problem of finding a smooth invariant measure for an *arbitrary* hybrid system. Instead of having to solve a (continuous) cohomology equation (5.13), we will now have to solve a “hybrid cohomology equation.”

An invariant measure for a smooth hybrid system $\mathcal{H} = (M, S, X, \Delta)$ is given by a hybrid-invariant volume form, i.e. μ is an invariant volume if $\mu \in \mathcal{A}_{\mathcal{H}} \cap \Omega^n(M)$. The following theorem gives conditions on such a μ .

Theorem VII.7. *Let $\mathcal{H} = (M, S, X, \Delta)$ be a smooth hybrid system. A volume form $\mu \in \Omega^n(M)$ is invariant under the hybrid dynamics if and only if*

$$\begin{cases} \operatorname{div}_{\mu}(X) = 0 \\ \Delta^* \iota_S^* i_X \mu = \iota_S^* i_X \mu \end{cases}$$

Proof. These conditions match those of Theorem VII.3 with the exception of the specular term (7.3). This is because μ is an n -form and $\dim S = n - 1$ and therefore $\iota_S^* \mu \equiv 0$ so the specular term is trivially satisfied. \square

In a manner similar to §5.6, suppose $\mu \in \Omega^n(M)$ but $\mu \notin \mathcal{A}_{\mathcal{H}}$. What conditions can be placed on a function $f : M \rightarrow \mathbb{R}$ to guarantee $f \cdot \mu \in \mathcal{A}_{\mathcal{H}}$?

Definition VII.8. Let $\mathcal{H} = (M, S, X, \Delta)$ be a smooth hybrid system and let μ be a volume form on M . The unique function $\mathcal{J}_{\mu}(\Delta) \in C^{\infty}(S)$ such that

$$(7.4) \quad \Delta^* \iota_{\tilde{S}}^* i_X \mu = \mathcal{J}_{\mu}(\Delta) \cdot \iota_S^* i_X \mu,$$

is called the hybrid Jacobian of Δ with respect to μ .

Proposition VII.9. *For a smooth hybrid system $\mathcal{H} = (M, S, X, \Delta)$, there exists a smooth hybrid-invariant volume, $f \cdot \mu$, if there exists a smooth function $g : M \rightarrow \mathbb{R}$ such that*

$$(7.5) \quad \begin{aligned} dg(X) &= -\operatorname{div}_{\mu}(X) \\ g \circ \Delta - g|_S &= -\ln(\mathcal{J}_{\mu}(\Delta)). \end{aligned}$$

Then the density is (up to a multiplicative constant) $f = e^g$.

Proof. The first condition is precisely (5.13). To see the second, we apply the energy condition to $f \cdot \mu$.

$$\begin{aligned} \Delta^* \iota_{\tilde{S}}^* i_X (f \cdot \mu) &= (\Delta^* \iota_{\tilde{S}}^* f) \cdot (\Delta^* \iota_{\tilde{S}}^* i_X \mu) \\ &= (f \circ \Delta) \cdot (\mathcal{J}_{\mu}(\Delta) \cdot \iota_S^* i_X \mu) \\ &= f|_S \cdot \iota_S^* i_X \mu. \end{aligned}$$

So the energy condition is satisfied if and only if $\mathcal{J}_{\mu}(\Delta) \cdot (f \circ \Delta) = f|_S$. Taking the logarithm of both sides yields the result. \square

The system of equations in Proposition VII.9 is a differential-algebraic equation which reflects the dual nature of hybrid systems.

Definition VII.10. Let $\mathcal{H} = (M, S, X, \Delta)$ be a smooth hybrid system. Suppose $f \in C^\infty(M)$ and $k \in C^\infty(S)$. The differential-algebraic equation

$$(7.6) \quad \begin{cases} dg(X) = f \\ g \circ \Delta - g|_S = k \end{cases}$$

is called a hybrid cohomology equation.

Just as in §5.6, finding an invariant measure is equivalent to solving a cohomology equation. In general, we expect solving (7.6) to be very difficult if not impossible. This is because merely the continuous cohomology equation is difficult to solve, let alone the additional algebraic boundary conditions that must be satisfied.

7.3 Invariant Measures for Nonholonomic Hybrid Systems

Theorem VII.5 shows that unconstrained hybrid mechanical systems automatically preserve measure. In the language of Proposition VII.9, $\text{div}_{\omega^n}(X_H) = 0$ and $\mathcal{J}_{\omega^n}(\Delta) = 1$ and the corresponding cohomology equation has a trivial solution. On the other hand, for a nonholonomic system it is no longer true that $\text{div}_{\mu_{\mathcal{E}}} (X_H^{\mathcal{D}}) = 0$ and determining invariant measures becomes a nontrivial task. Likewise, it is no longer obvious that $\mathcal{J}_{\mu_{\mathcal{E}}}(\Delta^{\mathcal{D}}) = 1$. In what follows, we compute the hybrid Jacobian and attempt to solve the corresponding hybrid cohomology equation.

7.3.1 The Hybrid Jacobian

In order to find invariant measures for nonholonomic hybrid systems, we need to be able to compute $\mathcal{J}_{\mu_{\mathcal{E}}}(\Delta^{\mathcal{D}})$. In order to calculate this, we will use a similar trick as we did to compute the divergence: calculate $\mathcal{J}_{\omega^n}(\Delta^{\mathcal{E}})$ and restrict to \mathcal{D}^* . This will require a tool similar to Lemma V.15.

Lemma VII.11. *Let $\mathcal{H}^\mathcal{C} = (T^*Q, S^*, \Xi_H^\mathcal{C}, \Delta^\mathcal{C})$ be the global version of the nonholonomic hybrid system $\mathcal{H} = (\mathcal{D}^*, S_{\mathcal{D}}^*, X_H^\mathcal{D}, \Delta^\mathcal{D})$ and let $\mu_\mathcal{C}$ be the nonholonomic volume form. Then*

$$\mathcal{J}_{\omega^n}(\Delta^\mathcal{C})|_{\mathcal{D}^*} = \mathcal{J}_{\mu_\mathcal{C}}(\Delta^\mathcal{D}).$$

Proof. The proof will be similar in spirit to Lemma V.15. A computation yields:

$$\begin{aligned} \Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* i_{\Xi_H^\mathcal{C}} \omega^n &= \Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* i_{\Xi_H^\mathcal{C}} (\sigma_\mathcal{C} \wedge \varepsilon) \\ &= \Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* \left(i_{\Xi_H^\mathcal{C}} \sigma_\mathcal{C} \wedge \varepsilon + (-1)^m \sigma_\mathcal{C} \wedge i_{\Xi_H^\mathcal{C}} \varepsilon \right) \\ &= (-1)^m \left(\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* \sigma_\mathcal{C} \right) \wedge \left(\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* i_{\Xi_H^\mathcal{C}} \varepsilon \right) \\ &= (-1)^m \left(\iota_S^* \sigma_\mathcal{C} \right) \wedge \left(\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* i_{\Xi_H^\mathcal{C}} \varepsilon \right), \end{aligned}$$

which uses the fact that the constraints are preserved under the flow. That is, $i_{\Xi_H^\mathcal{C}} \sigma_\mathcal{C} = 0$ and $\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* \sigma_\mathcal{C} = \iota_S^* \sigma_\mathcal{C}$. The right side of (7.4) produces

$$\begin{aligned} \mathcal{J}_{\omega^n}(\Delta^\mathcal{C}) \cdot \iota_S^* i_{\Xi_H^\mathcal{C}} (\sigma_\mathcal{C} \wedge \varepsilon) &= \mathcal{J}_{\omega^n}(\Delta^\mathcal{C}) \cdot \iota_S^* \left(i_{\Xi_H^\mathcal{C}} \sigma_\mathcal{C} \wedge \varepsilon + (-1)^m \sigma_\mathcal{C} \wedge i_{\Xi_H^\mathcal{C}} \varepsilon \right) \\ &= (-1)^m \mathcal{J}_{\omega^n}(\Delta^\mathcal{C}) \cdot \left(\iota_S^* \sigma_\mathcal{C} \right) \wedge \left(\iota_S^* i_{\Xi_H^\mathcal{C}} \varepsilon \right). \end{aligned}$$

Combining both of the above gives

$$\left(\iota_S^* \sigma_\mathcal{C} \right) \wedge \left(\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* i_{\Xi_H^\mathcal{C}} \varepsilon \right) = \mathcal{J}_{\omega^n}(\Delta^\mathcal{C}) \cdot \left(\iota_S^* \sigma_\mathcal{C} \right) \wedge \left(\iota_S^* i_{\Xi_H^\mathcal{C}} \varepsilon \right).$$

The result follows from restricting to \mathcal{D}^* . \square

Therefore, to calculate $\mathcal{J}_{\omega^n}(\Delta^\mathcal{C})$, we need to understand $\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* i_{\Xi_H^\mathcal{C}} \omega^n$. Expanding gives

$$\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* i_{\Xi_H^\mathcal{C}} \omega^n = \Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* \left(n \cdot \nu_H^\mathcal{C} \wedge \omega^{n-1} \right) = n \cdot \left(\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* \nu_H^\mathcal{C} \right) \wedge \left(\Delta^{\mathcal{C}*} \iota_{\tilde{S}}^* \omega \right)^{n-1}.$$

Therefore, the hybrid Jacobian is determined by how much the nonholonomic 1-form and the symplectic form fail the specular condition (7.3). We next present a helpful computational lemma which will be useful for computing the above.

Lemma VII.12. Let $x^1, \dots, x^{n-1}, p_1, \dots, p_n$ be local coordinates and let $A = (\alpha_j^i)$ be an $(n-1) \times n$ matrix. Then

$$(7.7) \quad \left(\sum_{j=1}^{n-1} \sum_{i=1}^n \alpha_j^i \cdot dx^j \wedge dp_i \right)^{n-1} = (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \sum_{k=1}^n \det(A_k) \cdot \Omega_k,$$

$$\Omega_k := dx^1 \wedge \dots \wedge dx^{n-1} \wedge dp_1 \wedge \dots \wedge \widehat{dp_k} \wedge \dots \wedge dp_n,$$

where A_k is the $(n-1) \times (n-1)$ matrix from deleting the k^{th} -column from A and the caret $\widehat{dp_k}$ means that dp_k is omitted from the wedge product.

Proof. Recall the multinomial theorem which states that

$$(7.8) \quad \left(\sum_{j=1}^{n-1} \sum_{i=1}^n \alpha_j^i \cdot dx^j \wedge dp_i \right)^{n-1} = \sum_{S(d_j^i)=n-1} \binom{n-1}{d_1^1, \dots, d_{n-1}^n} \prod_{ij} (\alpha_j^i \cdot dx^j \wedge dp_i)^{d_j^i},$$

where

$$S(d_j^i) = \sum_{j=1}^{n-1} \sum_{i=1}^n d_j^i, \quad \binom{n-1}{d_1^1, \dots, d_{n-1}^n} = \frac{n!}{d_1^1! \cdot \dots \cdot d_{n-1}^n!}.$$

Notice that for any i, j we have $(\alpha_j^i \cdot dx^j \wedge dp_i)^2 = 0$. This implies that the only nonzero terms in (7.8) have $d_j^i \in \{0, 1\}$. This simplifies (7.8) to

$$(7.9) \quad \left(\sum_{j=1}^{n-1} \sum_{i=1}^n \alpha_j^i \cdot dx^j \wedge dp_i \right)^{n-1} = (n-1)! \cdot \sum_{\substack{S(d_j^i)=n-1 \\ d_j^i \in \{0,1\}}} \prod_{ij} (\alpha_j^i \cdot dx^j \wedge dp_i)^{d_j^i}.$$

In order to evaluate (7.9), we wish to understand the structure of the matrices $\mathbf{d} = (d_j^i)$ that contribute a nonzero term. In addition to having coefficients in $\{0, 1\}$, they also have the following property: if $d_j^i = 1$, then $d_j^k = d_k^i = 0$ for all k . This is due to the fact that $(dx^j \wedge dp_i) \wedge (dx^\ell \wedge dp_k) = 0$ whenever $j = \ell$ or $i = k$. In other words, the matrix \mathbf{d} must have a single nonzero entry in each row and at most one in each column. The matrix \mathbf{d} is then given as a column permutation of the matrix

$$\mathbf{d}_0 = \begin{bmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{bmatrix}.$$

Let \mathcal{D} be the set of all such matrices and partition it as $\mathcal{D} = \sqcup_{k=1}^n \mathcal{D}_k$ where $\mathbf{d} \in \mathcal{D}_k$ if its k^{th} -column is identically zero. The expression (7.9) becomes

$$(7.10) \quad \left(\sum_{j=1}^{n-1} \sum_{i=1}^n \alpha_j^i \cdot dx^j \wedge dp_i \right)^{n-1} = (n-1)! \cdot \sum_{k=1}^n \sum_{(d_j^i) \in \mathcal{D}_k} \prod_{i,j:d_j^i \neq 0} \alpha_j^i \cdot dx^j \wedge dp_i.$$

By deleting the k^{th} -row from \mathcal{D}_k there is a natural isomorphism $S_{n-1} \rightarrow \mathcal{D}_k$, where S_{n-1} is the symmetric group of $n-1$ elements. A matrix $(d_j^i) \in \mathcal{D}_k$ if and only if there exists $\sigma \in S_{n-1}$ such that $d_j^i = 1$ if and only if $\sigma_k(j) = i$ where

$$\sigma_k(j) = \begin{cases} \sigma(j), & \sigma(j) < k \\ \sigma(j) + 1, & \sigma(j) > k. \end{cases}$$

(This modified permutation keeps track of the k^{th} -column deletion.) Before we finish the calculation of (7.7), we notice that (see 3.1.3 in [1])

$$\begin{aligned} \prod_{j=1}^{n-1} dx^j \wedge dp_{\sigma_k(j)} &= (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot dx^1 \wedge \dots \wedge dx^{n-1} \wedge dp_{\sigma_k(1)} \wedge \dots \wedge dp_{\sigma_k(n-1)} \\ &= (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot \text{sgn}(\sigma) \cdot \underbrace{dx^1 \wedge \dots \wedge dx^{n-1} \wedge dp_1 \wedge \dots \wedge \widehat{dp_k} \wedge \dots \wedge dp_n}_{=: \Omega_k}. \end{aligned}$$

Using this, we see that (7.10) becomes

$$\begin{aligned} \left(\sum_{j=1}^{n-1} \sum_{i=1}^n \alpha_j^i \cdot dx^j \wedge dp_i \right)^{n-1} &= (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \sum_{k=1}^n \sum_{\sigma \in S_{n-1}} \text{sgn}(\sigma) \prod_{j=1}^{n-1} \alpha_j^{\sigma_k(j)} \cdot \Omega_k \\ &= (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \sum_{k=1}^n \det(A_k) \cdot \Omega_k, \end{aligned}$$

which is precisely (7.7). □

Proposition VII.13. *In local coordinates where $h = x^n$, we have*

$$\iota_S^* (\nu_H^{\mathcal{C}} \wedge \omega^{n-1}) = (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \left(\frac{\partial H}{\partial p_n} \right) \cdot dx^1 \wedge \dots \wedge dx^{n-1} \wedge dp_1 \wedge \dots \wedge dp_n.$$

Proof. By Lemma VII.12 we can compute $(\iota_S^* \omega)^{n-1}$ where $\alpha_j^i = \delta_j^i$. This provides

$$(\iota_S^* \omega)^{n-1} = (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \Omega_n.$$

Due to the fact that Ω_n depends on every dx^j and dp_i with the exception of dp_n , the only component of $\nu_H^{\mathcal{C}}$ that wedges with $(\iota_S^* \omega)^{n-1}$ to produce a nonzero term is the p_n term, i.e.

$$\begin{aligned} \nu_H^{\mathcal{C}} \wedge (\iota_S^* \omega)^{n-1} &= \left(\frac{\partial H}{\partial p_n} \cdot dp_n \right) \wedge (\iota_S^* \omega)^{n-1} \\ &= (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \frac{\partial H}{\partial p_n} \cdot \Omega_n \wedge dp_n, \end{aligned}$$

where $\Omega_n \wedge dp_n = dp_n \wedge \Omega_n$ because Ω_n has even degree. \square

Corollary VII.14. *In coordinate-free language, we have*

$$\iota_S^* (\nu_H^{\mathcal{C}} \wedge \omega^{n-1}) = (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \pi_Q^* dh (\Xi_H^{\mathcal{C}}) \cdot \Omega_h,$$

where Ω_h is a volume on S^* given by

$$\Omega_h = \iota_S^* \varepsilon, \quad dh \wedge \varepsilon = (-1)^{n-1} \cdot dx^1 \wedge \dots \wedge dx^n \wedge dp_1 \wedge \dots \wedge dp_n.$$

We are now ready to proceed with calculating the hybrid Jacobian.

Theorem VII.15. *The hybrid Jacobian is given by*

$$(7.11) \quad \mathcal{J}_{\omega^n}(\Delta^{\mathcal{C}}) = \frac{(2 \cdot \pi_Q^* \pi_{\mathcal{D}}^* dh - \pi_Q^* dh)(\Xi_H^{\mathcal{C}})}{\pi_Q^* dh(\Xi_H^{\mathcal{C}})}.$$

In particular, $\mathcal{J}_{\mu_{\mathcal{C}}}(\Delta^{\mathcal{D}}) = 1$.

Proof. We will choose local coordinates such that $h = x^n$ in a manner similar to Proposition VII.13 and later translate to a coordinate-free language as in Corollary VII.14. We will first compute $(\Delta^{\mathcal{C}*} \iota_S^* \omega)^{n-1}$. In coordinates where $h = x^n$, this becomes

$$(\Delta^{\mathcal{C}*} \iota_S^* \omega)^{n-1} = (dx^1 \wedge d(p_1 \circ \Delta^{\mathcal{C}}) + \dots + dx^{n-1} \wedge d(p_{n-1} \circ \Delta^{\mathcal{C}}))^{n-1}.$$

The map $\Delta^{\mathcal{C}}$ given by Theorem VI.13 depends on both x and p :

$$d(p_j \circ \Delta^{\mathcal{C}}) = \alpha_j^i dp_i + \beta_{ij} dx^i.$$

With this notation, we have

$$\begin{aligned} (\Delta^{\mathcal{C}^*} \iota_S^* \omega)^{n-1} &= (\alpha_j^i \cdot dx^j \wedge dp_i + \beta_{ij} \cdot dx^j \wedge dx^i)^{n-1} \\ &= (\alpha_j^i \cdot dx^j \wedge dp_i)^{n-1}, \end{aligned}$$

where the $\beta_{ij} \cdot dx^j \wedge dx^i$ terms do not contribute because any piece containing them will necessarily have a repeated term. Therefore if we can determine the coefficients α_j^i , Lemma VII.12 shows how to compute the product. The expression (6.14) shows that the impacts are linear in the momentum and so the coefficients are

$$\alpha_j^i = \delta_j^i - 2 \frac{(\pi_{\mathcal{D}} \nabla h)^i}{dh(\pi_{\mathcal{D}} \nabla h)} (\pi_{\mathcal{D}}^* dh)_j,$$

where $(\pi_{\mathcal{D}} \nabla h)^i$ is the i^{th} -component of the vector $\pi_{\mathcal{D}} \nabla h$ and similarly for $\pi_{\mathcal{D}}^* dh$.

We must now calculate the determinants of the matrices A_k . For the remainder of the proof, we will deal with the $n = 4$ case but the general case works in the same way. For ease of notation, let $u := \pi_{\mathcal{D}} \nabla h$ and $v := \pi_{\mathcal{D}}^* dh$. Notice that in our choice of local coordinates,

$$dh(\pi_{\mathcal{D}} \nabla h) = dx^n(\pi_{\mathcal{D}} \nabla h) = (\pi_{\mathcal{D}} \nabla h)^n = u^n =: \frac{1}{\kappa}.$$

The matrix $A = (\alpha_j^i)$ is given by

$$A = \begin{bmatrix} 1 - 2\kappa u^1 v_1 & -2\kappa u^2 v_1 & -2\kappa u^3 v_1 & -2\kappa u^4 v_1 \\ -2\kappa u^1 v_2 & 1 - 2\kappa u^2 v_2 & -2\kappa u^3 v_2 & -2\kappa u^4 v_2 \\ -2\kappa u^1 v_3 & -2\kappa u^2 v_3 & 1 - 2\kappa u^3 v_3 & -2\kappa u^4 v_3 \end{bmatrix}.$$

The determinants $\det A_k$ are

$$\det A_1 = -2v_1,$$

$$\det A_2 = 2v_2,$$

$$\det A_3 = -2v_3,$$

$$\det A_4 = 1 - 2\kappa (u^1 v_1 + u^2 v_2 + u^3 v_3) = 2v^4 - 1.$$

Lemma VII.12 asserts that

$$(7.12) \quad (\Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* \omega)^3 = (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot 2 [(-v_1 \cdot \Omega_1 + v_2 \cdot \Omega_2 - v_3 \cdot \Omega_3 + v_4 \cdot \Omega_4) - \Omega_4].$$

To finish the theorem, we need to compute the wedge product of $\Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* \nu_H^{\mathcal{C}}$ with (7.12). It turns out that $\Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* \nu_H^{\mathcal{C}} \wedge (\Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* \omega)^3 = \iota_{\tilde{S}}^* \nu_H^{\mathcal{C}} \wedge (\Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* \omega)^3$. This is because

$$\begin{aligned} \Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* \nu_H^{\mathcal{C}} &= \Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* dH - \Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* (m_{\alpha\beta} \{H, P(W^\alpha)\} \pi_Q^* \eta^\beta) \\ &= \iota_{\tilde{S}}^* dH - \Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* (m_{\alpha\beta} \{H, P(W^\alpha)\} \pi_Q^* \eta^\beta), \end{aligned}$$

by conservation of energy. Notice that the second term above has the form $\gamma_i \cdot dx^i$ which pairs to zero when wedged with any Ω_k . Therefore,

$$\begin{aligned} \Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* (\nu_H^{\mathcal{C}} \wedge \omega^{n-1}) &= \iota_{\tilde{S}}^* dH \wedge (\Delta^{\mathcal{C}^*} \iota_{\tilde{S}}^* \omega)^{n-1} \\ &= (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \left[2 \left(v_j \cdot \frac{\partial H}{\partial p_j} \right) - \frac{\partial H}{\partial p_4} \right] \cdot \Omega \\ &= (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot \left[2 \cdot \pi_{\mathcal{D}}^* dh(\Xi_H^{\mathcal{C}}) - dh(\Xi_H^{\mathcal{C}}) \right] \cdot \Omega. \end{aligned}$$

The result follows from applying Proposition VII.13 which says

$$\iota_{\tilde{S}}^* (\nu_H^{\mathcal{C}} \wedge \omega^{n-1}) = (-1)^{\lfloor \frac{n-1}{2} \rfloor} \cdot (n-1)! \cdot dh(\Xi_H^{\mathcal{C}}) \cdot \Omega.$$

The quotient of coefficients is (7.11). □

Since $\mathcal{J}_{\mu_{\mathcal{C}}}(\Delta^{\mathcal{D}}) = 1$, we need an invariant density to be conserved across impacts: if $f \mu_{\mathcal{C}}$ is invariant then $f \circ \Delta^{\mathcal{D}} = f|_{S_{\mathcal{D}}^*}$. As it turns out, there is a clear qualitative difference between nonholonomic systems with measures depending on configurations versus those who do not. Before we demonstrate this with examples, we will quickly address the Zeno issue in measure-preserving systems.

7.4 The Zeno Issue in Measure Preserving Systems

Invariant measures provide information about asymptotic properties of flows. However, hybrid dynamical systems can exhibit a Zeno solution which prevents $\varphi_t^{\mathcal{H}}(x)$

making sense as $t \rightarrow \infty$. This section shows that, although they may exist, a given trajectory will almost never be Zeno.

Any Zeno issue will occur within the set $\mathcal{Z} := \overline{S} \cap \overline{\Delta(S)}$, which by (H.5) has codimension at least 2 (exactly 2 for mechanical systems). We will therefore focus our attention on trajectories that intersect this set; let \mathcal{N} be all points in M that eventually move to \mathcal{Z} ,

$$\mathcal{N} := \left\{ x \in M : \exists t > 0 \text{ s.t. } \lim_{s \rightarrow t^-} \varphi_s^{\mathcal{H}}(x) \in \mathcal{Z} \right\}.$$

Our goal is to show that \mathcal{N} has zero measure. A key to proving this is to introduce the following assumption.

Assumption VII.16 (Boundary identity property). *Consider a smooth hybrid dynamical system $\mathcal{H} = (M, S, X, \Delta)$. Then for any sequence $\{s_n\} \in S$ such that $s_n \rightarrow s \in \mathcal{Z}$ we have $\Delta(s_n) \rightarrow s$.*

This assumption is useful because it allows us to “complete” the hybrid flow in a manner similar to [5]. Essentially, suppose x_0 is Zeno so $\lim_{s \rightarrow t^-} \varphi_s^{\mathcal{H}}(x_0) = z_0 \in \mathcal{Z}$. Then we define $\varphi_t^{\mathcal{H}}(x_0) := z_0$ and we can extend it via assumption (A.1). Let $\varepsilon > 0$ such that $\varphi_t(z)$ does not intersect S for all $t \in (0, \varepsilon)$. We define the *completed* flow to be $\varphi_{t+\delta}^{\mathcal{H}}(x_0) = \varphi_\delta(z_0)$. If a hybrid flow is measure-preserving then its associated completed flow is too precisely due to the boundary identity property; we are ignoring any impacts at z_0 which comes from continuously extending the impact map from S to \overline{S} . We can now state the following theorem.

Theorem VII.17. *Suppose that $\mathcal{H} = (M, S, X, \Delta)$ is a smooth hybrid dynamical system with the boundary identity property, Assumption VII.16. If \mathcal{H} preserves a **smooth** measure μ , then $\mu(\mathcal{N}) = 0$.*

Proof. Partition \mathcal{Z} into a countable collection of compact sets, $\{V_\alpha\}$, and partition \mathcal{N} in the following way:

$$\mathcal{N}_{\alpha,\delta} = \{x \in \mathcal{N} : \exists t \in (0, \delta) \text{ s.t. } \varphi_t^{\mathcal{H}}(x) \in V_\alpha\}.$$

It follows that if each $\mathcal{N}_{\alpha,\delta}$ has zero measure then all of \mathcal{N} has zero measure, since a countable union of null sets is still a null set. In particular, we only need to prove that for all α , there exists δ such that $\mathcal{N}_{\alpha,\delta}$ has zero measure. This is because for $\delta > s$

$$\varphi_s^{\mathcal{H}}(\mathcal{N}_{\alpha,\delta} \setminus \mathcal{N}_{\alpha,s}) = \mathcal{N}_{\alpha,\delta-s}.$$

Fix an α . By (A.1), for each $z \in V_\alpha$, there exists $\varepsilon > 0$ such that $\varphi_t(z) \notin S$ for all $t \in (0, \varepsilon)$. Let δ be the infimum of all such ε which is positive due to the compactness of V_α . By the measure-preserving property of the flow, we get

$$\mu(\mathcal{N}_{\alpha,\delta/4}) = \mu\left(\varphi_{\delta/2}^{\mathcal{H}}(\mathcal{N}_{\alpha,\delta/4})\right) \leq \mu(\mathcal{O}(V_\alpha, \delta)),$$

where

$$\mathcal{O}(V_\alpha, \delta) = \bigcup_{t \in (0, \delta)} \varphi_t(V_\alpha).$$

Because zero impacts occur, the set $\mathcal{O}(V_\alpha, \delta)$ is a manifold with codimension at least 1 which necessarily has zero measure. \square

Hybrid mechanical systems are smooth and satisfy Assumption VII.16 so if a smooth invariant measure exists, Zeno states are negligible. We will use this result to justify that our ignorance of Zeno states is essentially benign.

7.5 Examples

For a nonholonomic hybrid system, a density f solves the hybrid cohomology equation (7.5) if and only if f solves the continuous cohomology equation (5.13) and

is invariant under impacts, $f \circ \Delta^{\mathcal{D}} = f|_{S_{\mathcal{D}}^*}$. It turns out that there is an important distinction between densities depending only on configuration versus those who do not. To demonstrate this we consider both the Chaplygin sleigh (with invariant density given by Theorem V.36) and the rolling disk (which trivially preserves measure, recall (5.11)).

7.5.1 Chaplygin Sleigh

An invariant density for the Chaplygin sleigh is

$$\hat{J}(u)^{-3} = \frac{1}{u^3} \cdot \frac{1}{(p_\theta - ap_y \cos \theta + ap_x \sin \theta)^3}.$$

Without loss of generality, we will take $Iu = I + ma^2$. The above can be simplified using the constraint (5.7).

$$\rho_{sleigh} := \hat{J} \left(\frac{I + ma^2}{I} \right)^{-3} = p_\theta^{-3}.$$

In order for measure to be conserved across impacts, we need p_θ to be conserved across impacts. Proposition VI.16 states that this is true if and only if $\pi_{\mathcal{D}}^* dh$ has zero p_θ component, i.e.

$$\begin{aligned} 0 &= \pi_{\mathcal{D}}^* dh \left(\frac{\partial}{\partial \theta} \right) = \left(dh - \frac{1}{\eta(W)} dh(W)\eta \right) \left(\frac{\partial}{\partial \theta} \right) \\ &= dh \left(\frac{\partial}{\partial \theta} \right) = \frac{\partial h}{\partial \theta}. \end{aligned}$$

In practical situations, this is impossible. Recall the construction of the impact map for the rolling disk in §6.4.2. If the boundary of the table (or rink if the sleigh is thought of as an ice-skate) is given by $\tilde{h}(x, y) = 0$, then an impact occurs when

$$h(x, y, \theta) := \tilde{h}(x + L \cos \theta, y + L \sin \theta) = 0,$$

where L is the length of the sleigh. Differentiating with respect to θ yields

$$\frac{\partial h}{\partial \theta} = L \left(\frac{\partial \tilde{h}}{\partial y} \cos \theta - \frac{\partial \tilde{h}}{\partial x} \sin \theta \right).$$

This is only zero when the sleigh impacts the wall perpendicularly, which cannot be guaranteed to always happen. Therefore, *the Chaplygin sleigh does not preserve hybrid measure*.

7.5.2 Rolling Disk

The issue with the Chaplygin sleigh is that even though there exists an invariant measure, it fails to satisfy the impact conditions, $f \circ \Delta^{\mathcal{D}} = f|_{S^*}$. When the non-holonomic system has an invariant measure depending only on configuration, this condition is automatically satisfied.

Proposition VII.18. *Let $\mathcal{H}^{\mathcal{C}} = (T^*Q, S^*, \Xi_H^{\mathcal{C}}, \Delta^{\mathcal{C}})$ be a global realization of $\mathcal{H} = (\mathcal{D}^*, S_{\mathcal{D}}^*, X_H^{\mathcal{D}}, \Delta^{\mathcal{D}})$. If the density form $\vartheta_{\mathcal{C}}$ can be integrated, i.e. there exists $\rho \in \Gamma(\mathcal{D}^0)$ such that $\vartheta_{\mathcal{C}} + \rho$ is exact, then the resulting density is hybrid-invariant.*

Proof. Let $dg = \vartheta_{\mathcal{C}} + \rho$. Then by Theorem V.24, g has the form $g = \tilde{g} \circ \pi_Q$. This is clearly a hybrid-invariant density because the impact is the identity on the base variables: $\pi_Q \circ \Delta^{\mathcal{C}} = \pi_Q$. \square

As we saw in §5.5.2, $\mu_{\mathcal{C}}$ is an invariant measure for the rolling disk. The corresponding density is $f \equiv 1$ which clearly depends on configuration. Therefore, the rolling disk always preserves measure *independent of the choice of impact surface* $S \subset M$. This allows for the following result. We first need to define the energy set.

Definition VII.19. Let (Q, S, H) be a hybrid Hamiltonian system. The energy set is

$$\Sigma_e := \{(x, p) \in T^*Q : H(x, p) = e\},$$

Moreover, let $\mu_{\mathcal{C}}^e$ be the measure on Σ_e given by

$$\mu_{\mathcal{C}}^e = \iota_e^* \varepsilon, \quad dH \wedge \varepsilon = \mu_{\mathcal{C}}.$$

where $\iota_e : \Sigma_e \hookrightarrow \mathcal{D}^*$ is the inclusion.

Lemma VII.20. *If $\mu_{\mathcal{E}}$ is an invariant measure for $(\mathcal{D}^*, S_{\mathcal{D}}^*, X_H^{\mathcal{D}}, \Delta^{\mathcal{D}})$, then $\mu_{\mathcal{E}}^e$ is an invariant measure for the system restricted to Σ_e .*

Proof. This proof will be similar essentially identical to 3.4.12 in [1]. Due to the fact that $\mu_{\mathcal{E}}$ is invariant under $X_H^{\mathcal{D}}$, we have

$$0 = \mathcal{L}_{X_H^{\mathcal{D}}} \mu_{\mathcal{E}} = dH \wedge \mathcal{L}_{X_H^{\mathcal{D}}} \varepsilon.$$

Thus $\mathcal{L}_{X_H^{\mathcal{D}}} \varepsilon$ must have the form $\mathcal{L}_{X_H^{\mathcal{D}}} \varepsilon = dH \wedge \tau$, for some τ . The volume $\mu_{\mathcal{E}}^e$ is invariant since

$$\mathcal{L}_{X_H^{\mathcal{D}}} \mu_{\mathcal{E}}^e = \iota_e^* \mathcal{L}_{X_H^{\mathcal{D}}} \varepsilon = \iota_e^* (dH \wedge \tau) = 0.$$

□

Theorem VII.21. *Let $f\mu_{\mathcal{E}}$ be an invariant measure depending only on configuration for $(\mathcal{D}^*, S_{\mathcal{D}}^*, X_H^{\mathcal{D}}, \Delta)$ and Σ_e be compact. Then for any open $E \subset \Sigma_e$ and time $T > 0$, there exists a time $t > T$ such that*

$$\varphi_t^{\mathcal{H}}(E) \cap E \neq \emptyset.$$

Proof. This proof is essentially the Poincaré recurrence theorem (see, e.g. 4.1.19 in [70], [40] for a version for impulsive systems, or [42] for a version for nonsmooth vector fields). The measure $f\mu_{\mathcal{E}}^e$ is invariant on the dynamics restricted to Σ_e and the result follows so long as Σ_e is compact. □

This allows for the following statement about nonholonomic billiards.

Corollary VII.22. *Let $S' \subset \mathbb{R}^2$ be a smooth closed curve. Then for any open set $E \in \Sigma_e$ and time $T > 0$ there exists a time $t > T$ such that*

$$\varphi_t^{\mathcal{H}}(E) \cap E \neq \emptyset,$$

where φ_t^H is the hybrid dynamics for the elastic rolling disk laid out in §6.4.

Proof. The form $\mu_{\mathcal{C}}$ is an invariant volume (so $f \equiv 1$ which clearly depends only on configuration). The energy set Σ_e is compact because H is convex and Q is compact. □

This result states that the rolling disk billiards is recurrent for *any* table-top so long as it is bounded.

CHAPTER VIII

Hamilton-Jacobi Theory for Hybrid Systems

This chapter begins the third and final part of this thesis and extends the ideas from §3.4 to hybrid mechanical systems in preparation for Chapter IX. In particular we will show a hybrid version of Theorem III.37, as well as provide a way to have completely integrable hybrid systems. For notation purposes, the action will be called \mathcal{A} rather than S as S is already used to describe the impact surface.

There has been work done with nonholonomic Hamilton-Jacobi theory [38, 63, 86], which could be extended to the hybrid situation as well (in the same spirit that Chapters VI and VII extend nonholonomic dynamics to include impacts). However, we postpone that for future work.

8.1 The Hybrid Hamilton-Jacobi Equation

Our first task is to extend Theorem III.37 to hybrid systems. This will essentially be (3.17) along with the impact conditions: $d\mathcal{A}_t^+ = \Delta(d\mathcal{A}_t^-)$. Solutions to this partial differential equation are necessarily multi-valued because for $x \in S$, $d\mathcal{A}$ has (at least) two distinct values.

Definition VIII.1. Let $\mathfrak{A} = \{\mathcal{A}_k\}$ be a family of smooth functions $\mathcal{A}_k : U_k \rightarrow \mathbb{R}$ where $U_k \subset (Q \setminus S) \times \mathbb{R}$ are open sets. The family \mathfrak{A} solves the hybrid Hamilton-Jacobi equation if:

1. For each k , \mathcal{A}_k solves the Hamilton-Jacobi equation on U_k .
2. Suppose $x \in \partial U_k \cap S$. Then there exists an m such that $x \in \partial U_m \cap S$ and

$$(8.1) \quad \lim_{y \rightarrow x} d\mathcal{A}_m(y) = \Delta \left(\lim_{y \rightarrow x} d\mathcal{A}_k(y) \right), \quad \lim_{y \rightarrow x} \mathcal{A}_m(y) = \lim_{y \rightarrow x} \mathcal{A}_k(y).$$

Remark VIII.2. We require continuity in the solutions at impact points because that will be a property of the optimal cost function discussed in Chapter IX.

As we will see in §8.1.1, the sets U_k can be nontrivial and nonequal. We can now state the hybrid Hamilton-Jacobi theorem. For a concise formulation, we will assume there is a single impact but it holds for any number of impacts.

Theorem VIII.3. *Let (Q, H, S) be a hybrid Hamiltonian system and $\mathfrak{A} = \{\mathcal{A}_k\}$ be a collection of smooth functions. Then the following are equivalent:*

(HHJ.1) *For every curve $c : (t_0, t_f) \rightarrow Q$ such that c only intersects S at $t^* \in (t_0, t_f)$ that satisfies*

$$\dot{c}(t) = d\pi_Q \cdot X_H(d\mathcal{A}_{k,t}(c(t))), \quad t < t^*$$

$$\dot{c}(t) = d\pi_Q \cdot X_H(d\mathcal{A}_{m,t}(c(t))), \quad t > t^*$$

where $\mathcal{A}_k(c(t^), t^*) = \mathcal{A}_m(c(t^*), t^*)$, the curve c is an integral curve of the hybrid dynamics.*

(HHJ.2) *The family \mathfrak{A} solves the hybrid Hamilton-Jacobi equation.*

Proof. Away from impacts, this follows directly from Theorem III.37. At impacts, we only need to apply (8.1). □

8.1.1 Example: The Bouncing Ball

To demonstrate the above theory, we will examine the case of the bouncing ball. Although this example is very simple, it offers the ability to be solved exactly so results can be compared.

No Impacts

Before we tackle the hybrid case, we first compute the action functional for the falling ball. The Lagrangian for this system is

$$L = \frac{1}{2}m\dot{x}^2 - mgx,$$

where m is the mass and g is gravity. Suppose that the ball starts at $(x_0, t_0) = (0, 0)$ and ends at a point (x, t) . The trajectory that connects these two points is

$$x(s) = -\frac{1}{2}gs^2 + v_0s, \quad v_0 = \frac{x}{t} + \frac{1}{2}gt, \quad s = t - t_0.$$

Integrating the Lagrangian over this path yields:

$$\mathcal{A}_0(x, t) = \frac{mx^2}{2t} - \frac{mgt}{24}(gt^2 + 12x).$$

The subscript indicates that zero impacts are present. Notice that \mathcal{A}_0 satisfies the Hamilton-Jacobi equation:

$$(8.2) \quad \frac{\partial \mathcal{A}_0}{\partial t} + \frac{1}{2m} \left(\frac{\partial \mathcal{A}_0}{\partial x} \right)^2 + mgx = 0.$$

Impacts

Assume that the ball makes k impacts between $(0, 0)$ and (x, t) . In order to achieve this, let the initial velocity be $v(0) = v_0^k > 0$. Let t^* be the time of the initial impact; the final impact will take place at $k \cdot t^*$. Solving the shooting problem for $(0, k \cdot t^*)$ to (x, t) gives the following system:

$$\begin{aligned} v_0^k &= \frac{x}{t - k \cdot t^*} + \frac{1}{2}g(t - k \cdot t^*), \\ t^* &= \frac{2v_0^k}{g}. \end{aligned}$$

Solving this system produces *two* solutions:

$$(8.3) \quad \begin{aligned} v_0^k &= \frac{(1 + 2k)gt \pm \sqrt{g^2t^2 - 8gk(1 + k)x}}{4k(k + 1)}, \\ t^* &= \frac{(1 + 2k)t \pm \sqrt{t^2 - 8k(k + 1)g^{-1}x}}{2k(k + 1)}. \end{aligned}$$

Remark VIII.4. The reason for two solutions in (8.3) is that (x, t) can be arrived at via either ascent or descent. This can be seen in Figure 8.1 where each color has two curves: one is rising and one is falling.

The problem of computing the action requires integrating the Lagrangian in the following sense:

$$\mathcal{A}_k^\pm(x, t) = k \cdot \int_0^{t^*} L dt + \int_0^{t-k \cdot t^*} L dt,$$

where the superscript of \pm indicates which solution in (8.3) we take. Calculating this integral yields:

$$\begin{aligned} \mathcal{A}_k^\pm(x, t) = & \frac{1}{k^2(k+1)^3} \left[\frac{mg^2\delta_\pm^3}{24} + \frac{m\alpha_\pm^2\delta_\pm}{64} - \frac{mg\alpha_\pm\delta_\pm^2}{16} \right] \\ & + \frac{mg^2\gamma_\pm^3}{3} + \frac{m\alpha_\pm^2\gamma_\pm}{32k^2(k+1)^2} - \frac{mg\alpha_\pm\gamma_\pm^2}{4k(k+1)}, \end{aligned}$$

where

$$\begin{aligned} \alpha_\pm &= g\delta_\pm, \\ \gamma_\pm &= t - \frac{k\delta_\pm}{2k(k+1)}, \\ \delta_\pm &= (1+2k)t \pm \sqrt{t^2 - 8k(k+1)xg^{-1}}. \end{aligned}$$

It is interesting to note that \mathcal{A}_k^\pm **does** solve the Hamilton-Jacobi equation (8.2) when $k \neq 0$. We can see that the family $\{\mathcal{A}_k^\pm\}$, along with \mathcal{A}_0 , solves the hybrid Hamilton-Jacobi equation as per Definition VIII.1.

Reachability

Notice that the solutions to (8.3) are not always real; there is an envelope where the solutions provide real answers:

$$(8.4) \quad x \leq \frac{gt^2}{8k(k+1)}.$$

This envelope defines the sets U_k :

$$U_k = \left\{ (x, t) \in \mathbb{R} \times \mathbb{R} : x \leq \frac{gt^2}{8k(k+1)} \right\}.$$

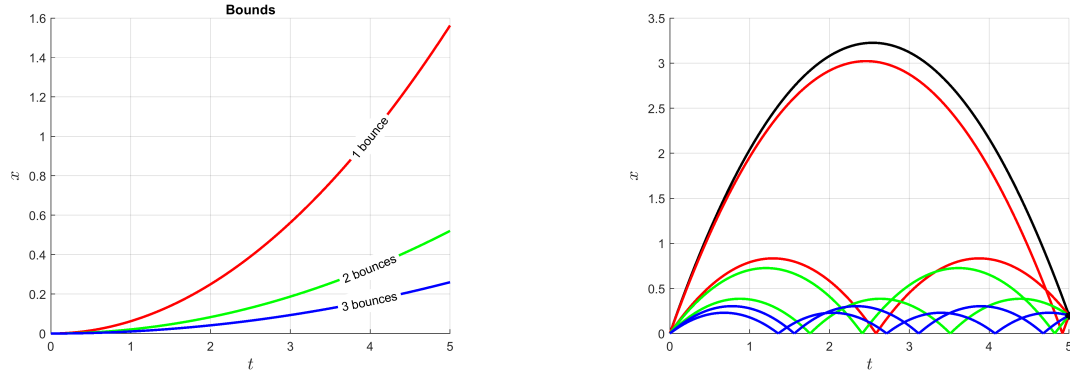


Figure 8.1: Left: The reachability bounds for the bouncing ball depending on the number of bounces. Right: Every trajectory that starts at $x = t = 0$ and ends at $x = 0.2$, $t = 5$ with no more than three bounces. The two solutions for each number of bounces corresponds to the two solutions to (8.3).

The Full Solution

We finish this example by plotting the whole solution. A plot of the functions \mathcal{A}_k^\pm for $k \leq 5$ is shown in Figure 8.2. To clarify the presentation of the plot, cross-sections depending on time are shown in Figure 8.3.

8.2 Integrability

Here, we extend the notion of Lagrangian submanifolds and complete integrability as described in §3.4.2. To show the applicability of this extension, we show that the bouncing ball is the hybrid equivalent of the harmonic oscillator as well as show the well-known result that the standard billiard problem on a circular table is integrable, see [69] and the references therein. We begin with defining a hybrid Lagrangian submanifold and then completely integrable hybrid systems.

Definition VIII.5. A submanifold (with boundary) $\mathcal{L} \subset M$ is a hybrid Lagrangian submanifold if

1. $\mathcal{L} \setminus \partial\mathcal{L} \subset M$ is a Lagrangian submanifold, cf. Definition III.5, and
2. $\pi_Q(\partial\mathcal{L}) \subset S$.




Figure 8.2: The plot of \mathcal{A}_k^\pm for k up to 5. The color-coding records the number of bounces: black means zero bounces, red (both light and dark) records a single bounce and so on.

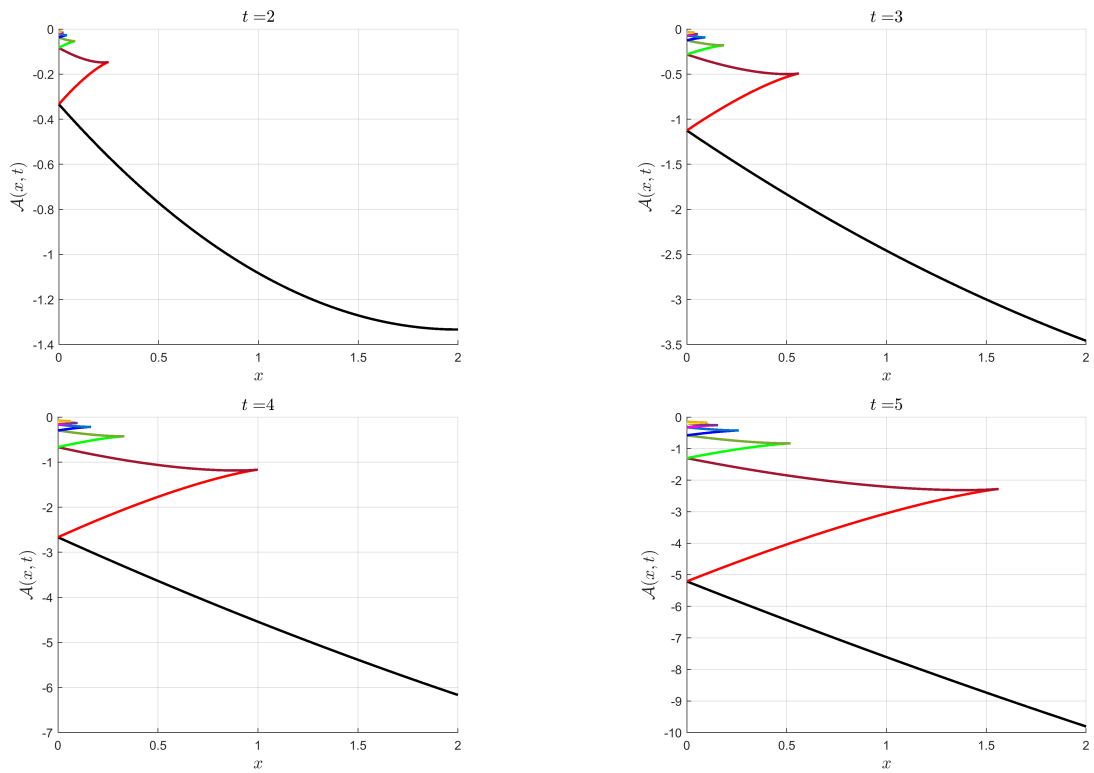


Figure 8.3: Cross-sections of Figure 8.2 at various times. The corners on the left correspond to impacts while the cusps on the right are reachability bounds, corresponding to (8.4).

This allows for the definition of an integrable hybrid system.

Definition VIII.6. A hybrid Hamiltonian system (Q, H, S) is completely integrable if there exists a hybrid Lagrangian foliation $\{\mathcal{L}_\alpha\}$ of T^*Q such that each leaf satisfies:

(HL.1) $\mathcal{L}_\alpha \subset H^{-1}(E)$ for some $E \in \mathbb{R}$, and

(HL.2) if $(x, p) \in \partial\mathcal{L} \cap S^*$, then $\Delta(x, p) \in \mathcal{L}$.

Remark VIII.7. We define integrability for hybrid systems via Lagrangian foliations rather than by constants of motion. This is due to the fact that the constants of motion must be in involution (Poisson commute) which requires some type of “hybrid bracket.”

Before we present examples, we first prove a hybrid version of Proposition III.40.

Proposition VIII.8. *Let (Q, H, S) be a hybrid Hamiltonian system and $\mathcal{L} \subset T^*Q$ be a hybrid Lagrangian submanifold. Call φ_t^H the flow of the hybrid Hamiltonian dynamics.*

1. *If \mathcal{L} satisfies (HL.1) and (HL.2), then \mathcal{L} is invariant under φ_t^H .*

2. *$\varphi_t^H(\mathcal{L})$ remains a Lagrangian submanifold.*

Proof. To show the first claim, we recall that (HL.1) implies that X_H is tangent to \mathcal{L} , cf. the proof of 5.3.32 in [1], X_H is tangent to \mathcal{L} . Invariance follows from this as well as closure under impacts, (HL.2).

The second follows immediately from the fact that φ_t^H is symplectic; recall Theorem VII.5. □

8.2.1 Example: Harmonic Oscillator and Bouncing Ball

Here, we show that the bouncing ball is essentially the hybrid version of the harmonic oscillator. We find a Lagrangian foliation for the 1-dimensional harmonic oscillator first for reference.




Figure 8.4: Left: The Lagrangian foliation for the harmonic oscillator. Right: The (multi-valued) action for the harmonic oscillator.

Harmonic Oscillator

Here, we examine the (non-hybrid) harmonic oscillator with $Q = \mathbb{R}$ and

$$H = \frac{1}{2m}p^2 + \frac{1}{2}kx^2.$$

The stationary Hamilton-Jacobi equation is

$$\frac{1}{2m} \left(\frac{\partial \mathcal{A}}{\partial x} \right)^2 + \frac{1}{2}kx^2 = E.$$

This can be solved by integrating

$$\frac{\partial \mathcal{A}}{\partial x} = \pm \sqrt{2mE - mkx^2}.$$

This produces a Lagrangian foliation of $T^*Q = \mathbb{R}^2$ into concentric circles.

Bouncing Ball

We repeat the same procedure as above but with the (hybrid) bouncing ball. The stationary Hamilton-Jacobi equation is

$$\frac{1}{2m} \left(\frac{\partial \mathcal{A}}{\partial x} \right)^2 + mgx = E.$$

Integrating this yields

$$\mathcal{A}(x) = -\frac{1}{3m^2g} (2Em - 2m^2gx)^{3/2} + C.$$




Figure 8.5: Left: The hybrid Lagrangian foliation for the bouncing ball. Right: The (multi-valued) action for the bouncing ball.

Instead of a foliation of concentric circles, the bouncing ball has a foliation of nested parabolas. Topologically, this is equivalent to circles once we glue the impacts together.

Remark VIII.9. We note that every hybrid Hamiltonian system with one degree of freedom is completely integrable in the same sense that their continuous counterparts are. Indeed, let \mathcal{A} be a solution to the Hamilton-Jacobi equation $H \circ d\mathcal{A} = E$. This means that the graph,

$$\Gamma_{d\mathcal{A}} = \{(q, d\mathcal{A}_q) : q \in Q\} \subset T^*Q$$

has constant energy. The corresponding Lagrangian foliation will be the foliation given by the level-sets of H . This foliation is also hybrid as energy is conserved across impacts.

8.2.2 Example: Billiards on a Circular Table

The classical billiard problem is known to be integrable when the table-top is circular or elliptical. In this example, we show that the circular billiard problem is integrable according to Definition VIII.6.

The circular billiard problem is given by the following hybrid Hamiltonian system

(Q, S, H) where

1. $Q = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq R^2\}$,
2. $S = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 = R^2\}$,
3. $H = \frac{1}{2m} (p_x^2 + p_y^2)$.

Theorem VI.8 states that the (elastic) impact map is given by

$$(8.5) \quad \begin{aligned} p_x &\mapsto p_x - \frac{2x}{R^2} (xp_x + yp_y), \\ p_y &\mapsto p_y - \frac{2y}{R^2} (xp_x + yp_y). \end{aligned}$$

In order to get complete integrability, we need to find a Lagrangian foliation of T^*Q that is invariant under the impact map (8.5). By construction the impact map preserves energy, which provides one first integral. For the other integral, we will choose angular momentum:

$$\begin{aligned} f_1 &= \frac{1}{2m} (p_x^2 + p_y^2), \\ f_2 &= xp_y - yp_x. \end{aligned}$$

These functions Poisson commute and therefore their joint level-set is Lagrangian,

$$\mathfrak{L}_{E,\nu} = \{(x, y, p_x, p_y) \in T^*Q : f_1 = E, f_2 = \nu\}.$$

These Lagrangian submanifolds satisfy both (HL.1) and (HL.2) since both energy and angular momentum are conserved across impacts (8.5). Therefore, the circular billiard problem is integrable.

Remark VIII.10. We need f_1 and f_2 to be in involution for their joint level-set to be Lagrangian. For this Lagrangian submanifold to be hybrid, we need f_1 and f_2 to be *preserved* under impacts. However, it is important to point out that the impact map is *constructed* from f_1 . Let this association be

$$\mathfrak{D} : C^\infty(T^*Q) \rightarrow C^\infty(S^*, T^*Q)$$

$$\mathfrak{D}(H) = \Delta_H,$$

where Δ_H satisfies (6.2). If the Lagrangian foliation defined by $f_1 := H = c_1$ and $f_2 = c_2$, to be hybrid we need

- $\{f_1, f_2\} = 0$, and
- $\mathfrak{D}(f_1)^* f_2 = f_2$.

This seems to imply a definition for a *hybrid bracket* via the following:

$$(8.6) \quad \{f, g\}_{\mathcal{H}} = \begin{cases} \{f, g\}, & x \notin S, \\ \mathfrak{D}(f)^* g - g, & x \in S. \end{cases}$$

This is problematic for (at least) four reasons.

(P.1) It is not clear that the hybrid bracket is smooth nor continuous.

(P.2) The hybrid bracket may not be skew.

(P.3) Certain conditions are needed for H to guarantee a single, nontrivial, solution to (6.2) (e.g. H being quadratic).

(P.4) Impacts might not occur. The Hamiltonian flow with Hamiltonian $xp_y - yp_x$ *never* impacts the circular table as the integral curves are circles.

It is an object of future research to address these problems.

CHAPTER IX

Optimal Control of Mechanical Hybrid Systems

This thesis concludes with a chapter on optimal control for hybrid dynamical systems. In particular we will show that the principle of optimality still holds for hybrid dynamical systems [56] which allows for dynamic programming. Moreover, we state the (strong) Hamilton-Jacobi-Bellman equation along with an accompanying maximum principle.

This problem has been covered extensively in the literature, e.g. [8, 33, 88, 89, 91, 110]. However, our approach will be somewhat different as we will focus on the geometric aspects of hybrid optimal control with an emphasis on systems whose continuous components are mechanical systems (Hamiltonian or nonholonomic). The main result from this chapter is the construction of the *strong hybrid Hamilton-Jacobi-Bellman equation*. This has been studied in [87, 90]; however the presentation here is much more geometric and is explicitly adapted for optimal control of mechanical systems.

The main restrictive assumption made throughout this chapter is that the controls do not influence the impact nor does the impact have any influence on the cost.

9.1 Hybrid Optimal Control Problem

We consider control problems of the form

$$\begin{cases} \dot{x} = f(x, u), & x \notin S, \\ x^+ = \Delta(x^-), & x \in S. \end{cases}$$

That is, we assume that there is no influence of the controls on the impact dynamics.

Definition IX.1 (Hybrid Control System). A hybrid control system is a 5-tuple $\mathcal{HC} = (M, \mathcal{U}, S, f, \Delta)$ such that

(HC.1) M is a smooth manifold,

(HC.2) $S \subset M$ is a smooth embedded submanifold with co-dimension 1,

(HC.3) $\mathcal{U} \subset \mathbb{R}^m$ a closed subset,

(HC.4) $f : M \times V \rightarrow TM$ is smooth where $U \subset V$ is an open neighborhood, and

(HC.5) $\Delta : S \rightarrow M$ is a smooth map.

The set $\mathcal{U} \subset \mathbb{R}^m$ is called the control set and a piece-wise smooth curve $u : \mathbb{R}^+ \rightarrow \mathcal{U}$ is an *admissible* control (here piece-wise smooth means u can be broken into a finite number of smooth curves). Let $C_{PW}^\infty(X, Y)$ be the set of all piece-wise smooth functions $X \rightarrow Y$. Therefore, an admissible control is $u \in C_{PW}^\infty(\mathbb{R}^+, \mathcal{U})$. The concept of a solution to a hybrid control system completely mimics that of the usual hybrid dynamical systems as laid out in §2.2.2.

Definition IX.2 (Performance Measure). A performance measure is a function $J : M \times C_{PW}^\infty([t_0, t_f], \mathcal{U}) \times [t_0, t_f] \rightarrow \mathbb{R}$ with the form

$$(9.1) \quad J(x_0, u(\cdot), s) = \int_s^{t_f} \ell(x(t), u(t), t) dt + g(x(t_f)),$$

$$\begin{cases} \dot{x}(t) = f(x(t), u(t)), & x \notin S, \\ x^+ = \Delta(x^-), & x \in S, \end{cases}$$

with the initial condition $x(s) = x_0$ and smooth functions $\ell : M \times \mathbb{R}^m \times \mathbb{R}^+ \rightarrow \mathbb{R}$ and $g : M \rightarrow \mathbb{R}$ called the continuous Lagrangian and terminal conditions, respectively.

Without changing the set up too much, we can also include a discrete Lagrangian, $\ell_d : S \times \mathbb{R}^m \times \mathbb{R}^+ \rightarrow \mathbb{R}$ and add an impact cost to (9.1). However, this will add subtle technicalities because the impact times are not free, rather predetermined by the initial condition and previous controls. This makes the impact times undefinable until *after* the controls are chosen. Another issue that a discrete Lagrangian introduces is that the performance measure will no longer be continuous along trajectories. We will postpone the study of these performance measures for future work.

Problem IX.3 (Optimal Control Problem). Given a hybrid control system, $\mathcal{HC} = (M, \mathcal{U}, S, f, \Delta)$, and a performance measure $J : M \times C_{PW}^\infty([t_0, t_f], \mathcal{U}) \times [t_0, t_f] \rightarrow \mathbb{R}$, the optimal control problem is to find, for a fixed $x_0 \in M$ and $s \in [t_0, t_f]$, an admissible control $u_{x_0, s}^* \in C_{PW}^\infty([t_0, t_f], \mathcal{U})$ which solves

$$(9.2) \quad u_{x_0, s}^* = \arg \min_{u(\cdot)} J(x_0, u(\cdot), s).$$

Remark IX.4. It is important to stress that existence and uniqueness of solutions to (9.2) does not always hold and can be difficult to determine. There is a natural way in which uniqueness fails; let $u_{x_0, s}^*$ be a solution to the optimal control problem, then any other $\tilde{u} \in C_{PW}^\infty([t_0, t_f], \mathcal{U})$ such that $u_{x_0, s}^*(\tau) = \tilde{u}(\tau)$ for $\tau > s$ is also clearly a solution. Because these are the same function when restricted to $[s, t_f]$ we will regard them as the same as to not further interfere with uniqueness, although other causes for non-uniqueness may still arise. As for existence, since $C_{PW}^\infty([t_0, t_f], \mathcal{U})$ is infinite-dimensional, minimizers may fail to exist e.g. weak solutions can appear.

If we assume that for a given $x_0 \in M$ and $s \in [t_0, t_f]$ there exists a unique solution $u_{x_0, s}^*$, then we can define the *optimal cost function* as

Definition IX.5. The optimal cost function, $J^* : M \times [t_0, t_f] \rightarrow \mathbb{R}$ is given by

$$J^*(x_0, s) = J(x_0, u_{x_0, s}^*, s).$$

Assumption IX.6. For all $x_0 \in M$ and $s \in [t_0, t_f]$, there exists a unique (in the sense of Remark IX.4) solution $u_{x_0, s}^*$ to the optimal control problem. Moreover, the optimal cost function is piece-wise smooth.

9.2 Principle of Optimality

The Bellman's principle of optimality allows us to solve for the optimal cost J^* backwards in time and is the fundamental principle behind dynamic programming. It states:

“An optimal policy has the property that whatever the initial state and initial decisions are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decisions,” [9].

In the language of this chapter, the principle of optimality takes the form of the following proposition.

Proposition IX.7 (Principle of Optimality). Let $u_{x_0, s}^*$ be an optimal solution. For an arbitrary $\xi \in [s, t_f]$ let $x_\xi = x(\xi)$ be the trajectory of the controlled system at time ξ . Then

$$u_{x_0, s}^* = \arg \min_{u(\cdot)} J(x_\xi, u(\cdot), \xi).$$

The proof of this depends on the structure of the performance measure and not on the dynamics; cf. §22.7.1 in [94] or §9.2 in [56] for a hybrid specific version. It is important to note that the converse is also true; if all tails are optimal then the

entire path must be too. A useful result of Proposition IX.7 is the “weak hybrid Hamilton-Jacobi-Bellman equation.”

Theorem IX.8 (Weak Hybrid Hamilton-Jacobi-Bellman). *Let $J^* : M \times [t_0, t_f] \rightarrow \mathbb{R}$ be the optimal cost function. Then*

$$(9.3) \quad 0 = \frac{\partial}{\partial t} J^*(x(t), t) + \min_{u \in \mathcal{U}} H(x(t), dJ(x(t), t), u, t), \quad x \notin S,$$

where dJ is the differential with respect to M only and $H : T^*Q \times \mathcal{U} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ where

$$H(x, p, u, t) = \ell(x, u, t) + \langle p, f(x, u) \rangle.$$

Moreover when $x \in S$, we have $J(\Delta(x), t) = J(x, t)$.

Proof. This proof is similar to that shown in [56]. Choose $(x(t), t) \in M \times [t_0, t_f]$ such that $x(t) \notin S$. Let $t' > t$ be small enough for no impact to occur on the interval $[t, t']$.

Then the principle of optimality states that

$$\begin{aligned} J^*(x(t), t) &= \min_{u(\cdot)} \int_t^{t_f} \ell(x(s), u(s), s) ds + g(x(t_f)) \\ &= \min_{u(\cdot)} \int_t^{t'} \ell(x(s), u(s), s) ds + \min_{u(\cdot)} \int_{t'}^{t_f} \ell(x(s), u(s), t) ds + g(x(t_f)) \\ &= \min_{u(\cdot)} \int_t^{t'} \ell(x(s), u(s), s) ds + J^*(x(t'), t'). \end{aligned}$$

Moving terms and dividing by $t' - t$ gives

$$0 = \frac{1}{t' - t} [J^*(x(t'), t') - J^*(x(t), t)] + \min_{u(\cdot)} \frac{1}{t' - t} \int_t^{t'} \ell(x(s), u(s), s) ds$$

Passing to the limit $t' \rightarrow t$ yields

$$0 = \frac{d}{dt} J^*(x(t), t) + \min_u \ell(x(t), u, t).$$

The result follows from the fact that

$$\frac{d}{dt} J^*(x(t), t) = \frac{\partial}{\partial t} J^*(x(t), t) + dJ(x(t), t) \cdot f(x, u).$$

The continuity condition $J(\Delta(x), t) = J(x, t)$ follows from the fact that integrating over a jump discontinuity still yields a continuous function. \square

The above theorem only describes what happens *away* from impacts and we are therefore lacking a boundary condition, thus the “weak” part. It is stated in [56] that J^* is continuous at impacts but stronger conditions can be found. Notice that (9.3) is precisely the continuous part of the hybrid Hamilton-Jacobi equation. This seems to suggest that $dJ^+ = \delta \circ dJ^-$ for some impact map δ . This will be postponed to §9.4 and used in §9.5 where we will work with a *hybrid maximum principle*.

9.3 Dynamic Programming

Dynamic programming is the method of solving the optimal control problem by discretizing the partial differential equation (9.3) with the terminal conditions $J(x, t_f) = h(x)$. Discretize $[t_0, t_f]$ with a sequence of times $t_0 < t_1 < \dots < t_N = t_f$ and call $J_K^*(\cdot) = J^*(\cdot, t_K)$. Then the principle of dynamic programming states that the optimal cost function satisfies the (backward) recurrence relation:

$$\begin{aligned} J_{N-k}^*(x_{N-k}) &= \min_{u_{N-k} \in \mathcal{U}} \ell_d(x_{N-k}, u_{N-k}, t_{N-k}) + J_{N-k+1}^*(a_d(x_{N-k}, u_{N-k})), \\ J_N^*(x_N) &= g(x_N), \end{aligned}$$

where $x_{k+1} = a_d(x_k, u_k)$ is a discretization of the dynamics and ℓ_d is a discretization of the Lagrangian.

9.3.1 Discretization of the Dynamics

For the purposes of this section, let us work within a coordinate chart so $x \in \mathbb{R}^n$. This allows for the continuous dynamics to be approximated by

$$x_{k+1} \approx x_k + \Delta t \cdot f(x_k, u_k).$$

Therefore, the discrete dynamics is given by $a_d(x_k, u_k) \approx x_k + \Delta t \cdot f(x_k, u_k)$ if no impact occurs. What happens if an impact occurs within $[t_k, t_{k+1}]$? If the impact is

given by $S = h^{-1}(0)$ then an impact occurs if

$$h(x_k) \cdot h(a_d(x_k, u_k)) < 0.$$

The impact time can be approximated via a first-order Taylor expansion:

$$\begin{aligned} 0 &= h(x_k + t^* \cdot f(x_k, u_k)) \approx h(x_k) + t^* \cdot dh(f(x_k, u_k)) \\ \implies &\boxed{t^* \approx -\frac{h(x_k)}{dh(f(x_k, u_k))}.} \end{aligned}$$

In order for the first-order approximation of the impact time to be finite, we need that the flow is transverse to the surface for *any* choice of controls. For dynamic control problems (as opposed to kinematic controls) we will see that this will not be a problem.

Using the above approximations, we get the discrete dynamics to be

$$(9.4) \quad a_d^{\mathcal{H}}(x_k, u_k) = \begin{cases} x_k + \Delta t_k \cdot f(x_k, u_k), \\ \Delta(\tilde{x}_k) + (\Delta t_k - t^*) \cdot f(\Delta(\tilde{x}_k), u_k), \end{cases}$$

$$\tilde{x}_k = x_k + t^* \cdot f(x_k, u_k)$$

where the first line is used if there is no impact and the second is used otherwise. In a similar manner, the discrete Lagrangian is

$$\ell_d^{\mathcal{H}}(x_k, u_k) = \begin{cases} \Delta t_k \cdot \ell(x_k, u_k), \\ t^* \cdot \ell(x_k, u_k) + (\Delta t_k - t^*) \ell(\Delta(x_k + t^* \cdot f(x_k, u_k)), u_k). \end{cases}$$

In practice due to approximations, $\tilde{x}_k \notin S$ and thus Δ needs to be defined on some tubular neighborhood of S .

We present an application of hybrid dynamic programming: the controlled bouncing ball.

9.3.2 Example: Bouncing Ball

Let the controlled bouncing ball be given by

$$\begin{aligned}\dot{x} &= \frac{1}{m}p, \\ \dot{p} &= -mg + u,\end{aligned}$$

along with the impact map $\Delta(x, p) = (x, -p)$ and impact surface $S = \{(0, p)\}$. Suppose the optimization problem is to minimize

$$(9.5) \quad J = \int_0^T \frac{1}{2} \kappa u^2 dt + \alpha (x(T) - 1)^2.$$

We want the final state to be $x(T) = 1$ while $p(T)$ is free and is chosen to minimize the L^2 -norm of the controller effort. We will solve this problem via dynamic programming. Using (9.4), the discrete dynamics are given by

$$\begin{aligned}x(t + \delta t) &= x(t) + \frac{\delta t}{m}p(t), \\ p(t + \delta t) &= p(t) + \delta t(u - mg),\end{aligned}$$

if no impact happens and

$$\begin{aligned}x(t + \delta t) &= -\left(\frac{p}{m} + \frac{x}{p}(mg - u)\right)\left(\delta t + \frac{mx}{p}\right), \\ p(t + \delta t) &= (u - mg)\left(\delta t + \frac{2mx}{p}\right) - p,\end{aligned}$$

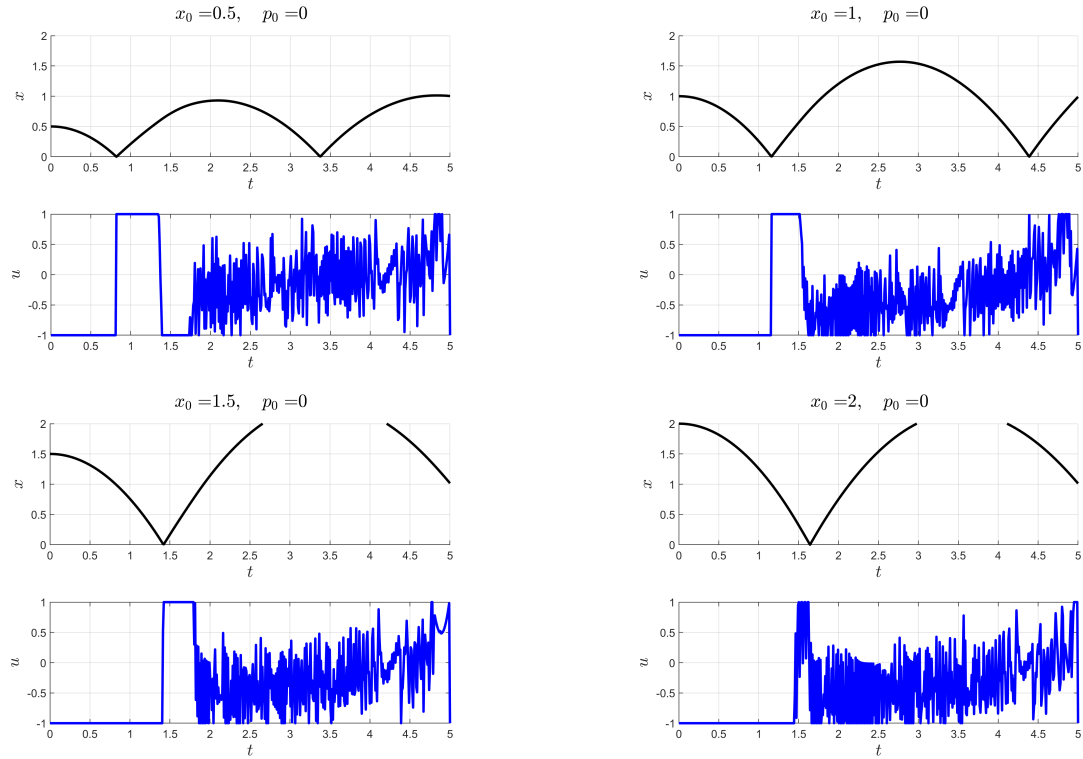
if an impact occurs.

Results

For this numerical example, we will take $\kappa = 0$ and $\alpha = 10$; the only performance we care about is that $x(T) = 1$. To make this problem interesting, and truly utilize the hybrid setting, we set $m = 2$, $g = 1$ and bound the controls so $u \in [-1, 1]$. This makes it impossible for the controls to actually support the ball, i.e. if this problem was purely continuous $x(t)$ must be monotone decreasing. Below, we demonstrate that if $x(0) < 1$ we can still have $x(T) = 1$.

Variable	Min	Max	# of samples
x	0.01	3	200
p	-5	5	200
u	-1	1	200
t	0	5	500

Table 9.1: The discretization used for the bouncing ball optimal control problem.

Figure 9.1: Optimal trajectories for the bouncing ball optimal control problem with various initial conditions x_0 with parameters $\alpha = 10$ and $\kappa = 0$. The effect of the noisy control is not visible in the above trajectories since the second-order nature of the dynamics mollifies the control input.

For the results in Figure 9.1, we choose $T = 5$ and use the uniform discretization as shown in Table 9.1. Notice that neither the number nor location of impacts are known a priori. The optimal controls seem to be bang-bang at the beginning and then become noisy. This is, in part, because the only objective is to have $x(T) = 1$ with no restriction on the controls. Choosing a positive κ has the property of “regularizing” the problem as can be seen in Figure 9.2 with $\kappa = 1$ (the discretization of Table 9.1 is still used).

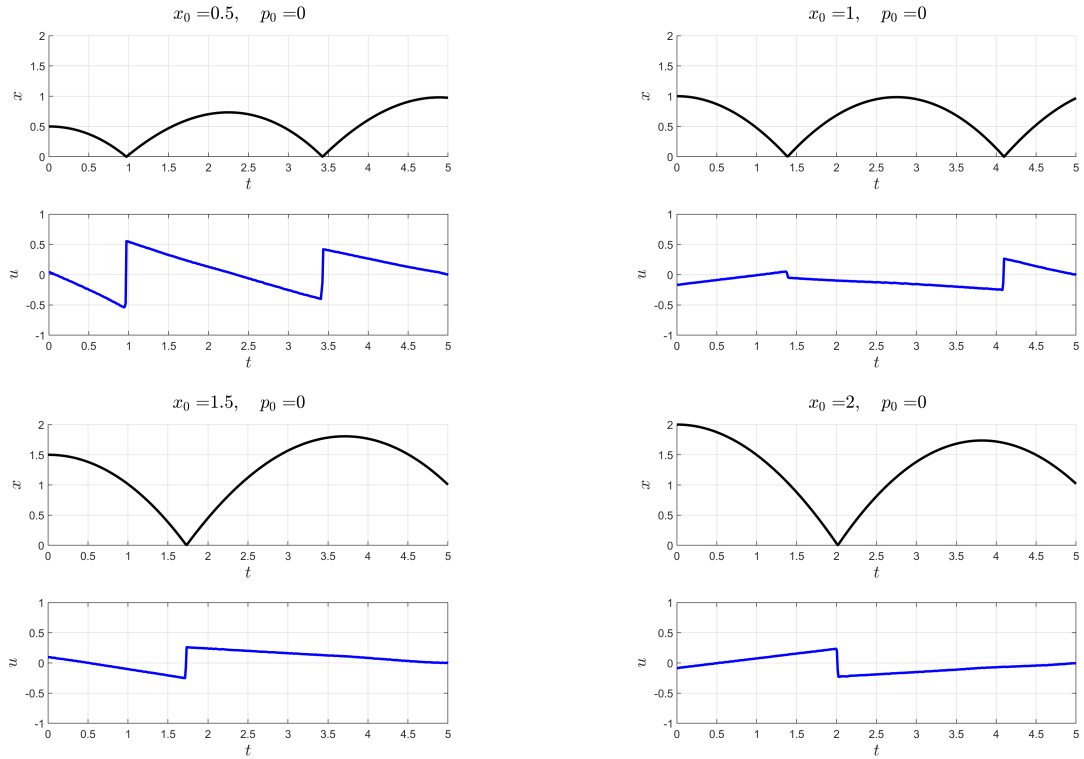


Figure 9.2: Optimal trajectories for the bouncing ball optimal control problem with various initial conditions x_0 with parameters $\alpha = 10$ and $\kappa = 1$.

9.4 The Strong Hybrid Hamilton-Jacobi-Bellman Equation

The principle of optimality, Proposition IX.7, allows for the optimal cost function to be described by the weak hybrid HJB equation (9.3). Discretizing this produces dynamic programming. Alternatively, we can solve (9.3) via the method of characteristics (cf. e.g. §3.2 in [43]) to produce a system of ordinary differential equations. The characteristics will provide the trajectories to the optimal control problem (Pontryagin's Maximum Principle).

Just as the continuous Hamilton-Jacobi-Bellman equation leads to continuous characteristics, the hybrid Hamilton-Jacobi-Bellman equation leads to hybrid characteristics. The continuous part of these characteristics can be determined from the weak hybrid HJB equation (9.3), but this offers no information about what happens

at impacts. To accomplish this, we will need a “strong” HJB equation, i.e. (9.3) along with a map $dJ^+ = \delta \circ dJ^-$ which encodes what happens at impacts.

9.4.1 Inspiration of a Strong HJB

To determine the map $dJ^+ = \delta \circ dJ^-$, we recall the principles that constructed this map for “usual” Hamiltonian systems studied in Chapter VI. Specifically, the map δ was defined variationally. Consider the case of a simple hybrid mechanical system: (Q, H, S) where Q is a smooth manifold, $H : T^*Q \rightarrow \mathbb{R}$ is a hyperregular Hamiltonian, and $S \subset Q$ is an embedded submanifold of codimension 1. Suppose $S = h^{-1}(0)$ for some function $h : Q \rightarrow \mathbb{R}$ such that 0 is a regular value. Then the impact is taken to be variational and given by the following rules:

$$(9.6) \quad \begin{aligned} p^+ &= p^- + \varepsilon \cdot dh, \\ H^+ &= H^-, \end{aligned}$$

which states that the change in momentum is proportional to dh and energy is conserved. In the language of the hybrid HJB equation, we get the following impact rules:

$$\begin{aligned} dJ^+ &= dJ^- + \varepsilon \cdot dh, \\ H^+ &= H^-. \end{aligned}$$

An implicit assumption made here is that the trajectory is still continuous at impact, i.e. $x^+ = x^-$. Unlike mechanical systems, this is no longer true for control systems where, generally, $\Delta(x) \neq x$.

9.4.2 Impacts with Discontinuous Base

We want to construct a generalization of (9.6) when impacts are *not* continuous on the base. Consider the case where we have a hybrid control system described by

$$(9.7) \quad \begin{cases} \dot{x} = f(x, u), & x \in M \setminus S, \\ x^+ = \Delta(x^-), & x \in S. \end{cases}$$

The characteristics to the hybrid HBJ (9.3) are given by

$$\begin{cases} \dot{x} = \frac{\partial H}{\partial p}, & \dot{p} = -\frac{\partial H}{\partial x}, & x \in M \setminus S, \\ x^+ = \Delta(x^-), & p^+ = \delta(x^-, p^-), & x \in S. \end{cases}$$

Determining the function δ will require a more general version of (9.6), in which jumps in the base are allowed.

Suppose that we have a trajectory with a forced discontinuity: $x : [t_0, t_f] \rightarrow M$ such that $t_1 \in (t_0, t_f)$ is the only point of discontinuity and

$$\lim_{t \rightarrow t_1^+} x(t) = \Delta \left(\lim_{t \rightarrow t_1^-} x(t) \right),$$

for some smooth map $\Delta : S \rightarrow M$. Determining the variations pre-jump can be done by considering the end point to be free in both space and time. However, the beginning of the second path is not free; it must obey the jump condition. Due to the infinitesimal nature of this problem, we can use

$$\Delta(x^*(t_1) + \delta x_1) \approx \Delta(x^*(t_1)) + \Delta_* (x^*(t_1)) \cdot \delta x_1.$$

Taking this into consideration, the variational ‘‘corner conditions’’ become (on the Lagrangian side)

$$\left[\frac{\partial L^-}{\partial \dot{x}} - \frac{\partial L^+}{\partial \dot{x}} \circ \Delta_* \right] \cdot \delta x = 0,$$

$$[E_L^- - E_L^+] \cdot \delta t = 0.$$




Figure 9.3: The black curve is the trajectory with a jump at t_1 while the red curve shows a virtual displacement.

The impact times are free so δt is arbitrary while the impact location is constrained to be on S so $\delta x \in TS$. This produces the general corner conditions as

$$\mathbb{F}L^- - \mathbb{F}L^+ \circ \Delta_* = \varepsilon \cdot dh,$$

$$E_L^- - E_L^+ = 0.$$

This is able to produce the jump conditions for the hybrid HJB equation:

$$(9.8) \quad dJ^+ \circ \Delta_* = dJ^- + \varepsilon \cdot dh,$$

$$H^+ = H^-.$$

In this derivation, we used Δ_* and not the augmented version Δ_*^f . As it turns out, as long as the impact is hit transversely, it does not matter which version is used. This is because the variations are assumed to be $\delta x \in TS$ and $\Delta_*^f(u) = \Delta_*(u)$ for all $u \in TS$. Any discrepancies in the transverse direction are handled via the multiplier ε in (9.8) as the following example demonstrates.

Example IX.9. Consider the simple case of the bouncing ball with $\dot{x} = v$ and $\dot{v} = -g$ with impact $\Delta(x, v) = (x, -v)$ when $x = 0$. For the purpose of this example, suppose

$H = p_x v - p_v g$. Using (7.1), the augmented differential is

$$\begin{aligned}\Delta_*^f \left(\alpha \frac{\partial}{\partial x} + \beta \frac{\partial}{\partial v} \right) &= \frac{\alpha}{v} \left(-v \frac{\partial}{\partial x} - g \frac{\partial}{\partial v} \right) + \Delta_* \left(\alpha \frac{\partial}{\partial x} + \beta \frac{\partial}{\partial v} - \frac{\alpha}{v} \left(v \frac{\partial}{\partial x} - g \frac{\partial}{\partial v} \right) \right) \\ &= -\alpha \frac{\partial}{\partial x} - \frac{\alpha g}{v} \frac{\partial}{\partial v} + \Delta_* \left(\left(\beta - \frac{\alpha g}{v} \right) \frac{\partial}{\partial v} \right) \\ &= -\alpha \frac{\partial}{\partial x} - \beta \frac{\partial}{\partial v}.\end{aligned}$$

The pre- and post-impact momenta are given by

$$\begin{aligned}p^-(v_x, v_v) &= p_x v_x + p_v v_v, \\ p^+ \circ \Delta_*^f(v_x, v_v) &= -p_x^+ v_x - p_v^+ v_v.\end{aligned}$$

The ‘‘specular’’ condition provides: $-p_x^+ v_x - p_v^+ v_v - p_x^- v_x - p_v^- v_v = \varepsilon \cdot v_x$. In particular, $p_v^+ = -p_v^-$ and $p_x^+ = -p_x^- - \varepsilon$. To determine the multiplier ε , we use conservation of energy which gives the impact map to be

$$\begin{aligned}p_x^+ &= -p_x^- + \frac{2g}{v} \cdot p_v^-, \\ p_v^+ &= -p_v^-.\end{aligned}$$

Likewise, if we use the ordinary differential, Δ_* , we see that

$$p^+ \circ \Delta_*(v_x, v_v) = p_x^+ v_x - p_v^+ v_v.$$

The new impact equations are then $p_x^+ = p_x^- + \varepsilon$ and $p_v^+ = -p_v^-$. Although the impact dynamics seem to be off by a sign, this is not a problem because the value of ε changes accordingly. Using conservation of energy, we see that

$$p_x^+ = p_x^- - \frac{2}{v} (p_x^- v - p_v^- g) = -p_x^- + \frac{2g}{v} \cdot p_v^-.$$

Therefore, we obtain the same impact map for either Δ_* or Δ_*^f .

9.4.3 The Strong Hybrid HJB Equation

The corner condition (9.8) allows for a strong version of the hybrid Hamilton-Jacobi-Bellman equation. We note that the following makes use of time-dependent

Hamiltonians which have not been studied thus far in this thesis. Hamilton's equations, (3.6) still hold although they are no longer autonomous.

Theorem IX.10 (Strong Hybrid Hamilton-Jacobi-Bellman). *Let $J^* : M \times [t_0, t_f] \rightarrow \mathbb{R}$ be the optimal cost function. Define the Hamiltonian $H : T^*Q \times \mathcal{U} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ where*

$$H(x, p, u, t) = \ell(x, u, t) + \langle p, f(x, u) \rangle,$$

and the optimal Hamiltonian, $H^ : T^*Q \times \mathbb{R}^+ \rightarrow \mathbb{R}$ where*

$$H^*(x, p, t) = \min_u H(x, p, u, t).$$

Then the optimal cost function satisfies

$$(9.9) \quad 0 = \frac{\partial}{\partial t} J^*(x, t) + H^*(x, dJ(x, t), t), \quad x \notin S,$$

$$(9.10) \quad \begin{cases} dJ(\Delta(x), t) \circ \Delta_* = dJ(x, t) + \varepsilon \cdot dh, \\ H^*(\Delta(x), dJ(\Delta(x), t)) = H^*(x, dJ(x, t)), \end{cases} \quad x \in S.$$

9.5 Maximum Principle

The maximum principle [49] is a way to solve the optimal control problem via the method of characteristics. For the continuous case, the characteristics of (9.9) are determined via the system of ordinary differential equations:

$$\dot{x} = \frac{\partial H^*}{\partial p}, \quad \dot{p} = -\frac{\partial H^*}{\partial x}.$$

Pontryagin's maximum principle states that necessary conditions for a trajectory to be optimal is for it to be a solution to the above characteristic equation. This is extended to the hybrid situation by incorporating (9.10).

Theorem IX.11 (Hybrid Maximum Principle). *Let u_{x_0, t_0}^* be an optimal control. Then the solution $x(t)$ satisfies*

$$(9.11) \quad \begin{cases} \dot{x} = \frac{\partial H^*}{\partial p}, & \dot{p} = -\frac{\partial H^*}{\partial x}, & x \notin S, \\ x^+ = \Delta(x^-), \\ p^+ \circ \Delta_* = p^- + \varepsilon \cdot dh, & x \in S. \\ (H^*)^+ = (H^*)^- \end{cases}$$

9.5.1 Optimal Control of a Mechanical System

Suppose now that the underlying hybrid system is of mechanical type; so Δ in (9.11) is given by (9.6). Also, let us assume that the Lagrangian is natural: kinetic minus potential. The impact map, Δ , is thus linear in the velocities / momentum. Here, we use coordinates $(q, y) \in M = T^*Q$ and $(q, y, p_q, p_y) \in T^*M$. The impact map is

$$\Delta(q, y) = (q, R(q)y),$$

with differential

$$\Delta_* = \begin{bmatrix} \text{Id} & 0 \\ R'y & R \end{bmatrix}, \quad R'y = \frac{\partial}{\partial q} (R(q)y).$$

Using the impact rule (9.11), we get an impact map of

$$\begin{aligned} q^+ &= q^- \\ y^+ &= R(q^-)y^- \\ [p_q^+, p_y^+] \begin{bmatrix} \text{Id} & 0 \\ R'y & R \end{bmatrix} &= [p_q^-, p_y^-] + \varepsilon \cdot [dh, 0]. \end{aligned}$$

Solving this last equation provides us with the following:

$$\begin{aligned}
 q &\mapsto q \\
 y &\mapsto R(q)y \\
 p_q &\mapsto p_q - p_y \cdot R^{-1}R'y + \varepsilon \cdot dh \\
 p_y &\mapsto p_y \cdot R^{-1}.
 \end{aligned}
 \tag{9.12}$$

Recall that the linear map $R(q)y$ (for unconstrained systems) is given by

$$R(q)y = y - 2 \frac{\tilde{g}(dh, y)}{\tilde{g}(dh, dh)} dh.$$

Recall that the matrix R is Householder which implies that $R = R^{-1}$. Writing (9.12) in terms of adjoints, we get the following for the impact map.

Proposition IX.12. *The impact map is given by*

$$\begin{aligned}
 q &\mapsto q \\
 y &\mapsto R(q)y \\
 p_q &\mapsto p_q - (R'y)^* R^* p_y + \varepsilon \cdot dh \\
 p_y &\mapsto R^* p_y.
 \end{aligned}$$

Definition IX.13. An impact is called flat if $R'y \equiv 0$.

Corollary IX.14. *If the impact is flat, the impact map simplifies to*

$$\begin{aligned}
 p_q &\mapsto p_q + \varepsilon \cdot dh, \\
 p_y &\mapsto R^* p_y.
 \end{aligned}$$

9.6 Example: Bouncing Ball

For this example, we revisit the controlled bouncing ball as discussed in §9.3.2. In particular, we will compute the trajectories of the maximum principle to show that the optimal controls are piece-wise linear, as suggested by Figure 9.2. We will then use the results from the maximum principle to compute the optimal cost function.

9.6.1 Maximum Principle

Recall the controlled bouncing ball dynamics (here, we relabel p as y),

$$\dot{x} = \frac{1}{m}y,$$

$$\dot{y} = -mg + u,$$

where impact occurs at $x = 0$ and the map is $Ry = -y$. Suppose we have the cost we wish to minimize (which is precisely (9.5) with $\kappa = 1$)

$$J = \int_{t_0}^{t_f} \frac{1}{2}u^2 dt + \alpha (x(t_f) - 1)^2.$$

The Hamiltonian is given by

$$\begin{aligned} H^*(x, y, p_x, p_y) &= \min_u \left[\frac{1}{2}u^2 + \frac{1}{m}p_x y + (u - mg)p_y \right] \\ &= -\frac{1}{2}p_y^2 + \frac{1}{m}p_x y - mgp_y. \end{aligned}$$

Therefore, the continuous dynamics are given by

$$(9.13) \quad \begin{aligned} \dot{x} &= \frac{1}{m}y, & \dot{p}_x &= 0, \\ \dot{y} &= -mg - p_y, & \dot{p}_y &= -\frac{1}{m}p_x. \end{aligned}$$

It remains to determine what the impact map is; the impact is flat so Corollary IX.14 provides the following.

$$(9.14) \quad \begin{aligned} x &\mapsto x, \\ y &\mapsto -y, \\ p_x &\mapsto -p_x + 2m^2 g \frac{p_y}{y} \\ p_y &\mapsto -p_y. \end{aligned}$$

Where the map for p_x comes from conservation of energy.

Remark IX.15. The optimal control is given by $u^* = -p_y$. During the continuous phase p_x is constant and p_y is linear, i.e. u^* is piece-wise linear with jumps during impacts. This explains the piece-wise linear nature of the controls in Figure 9.2.

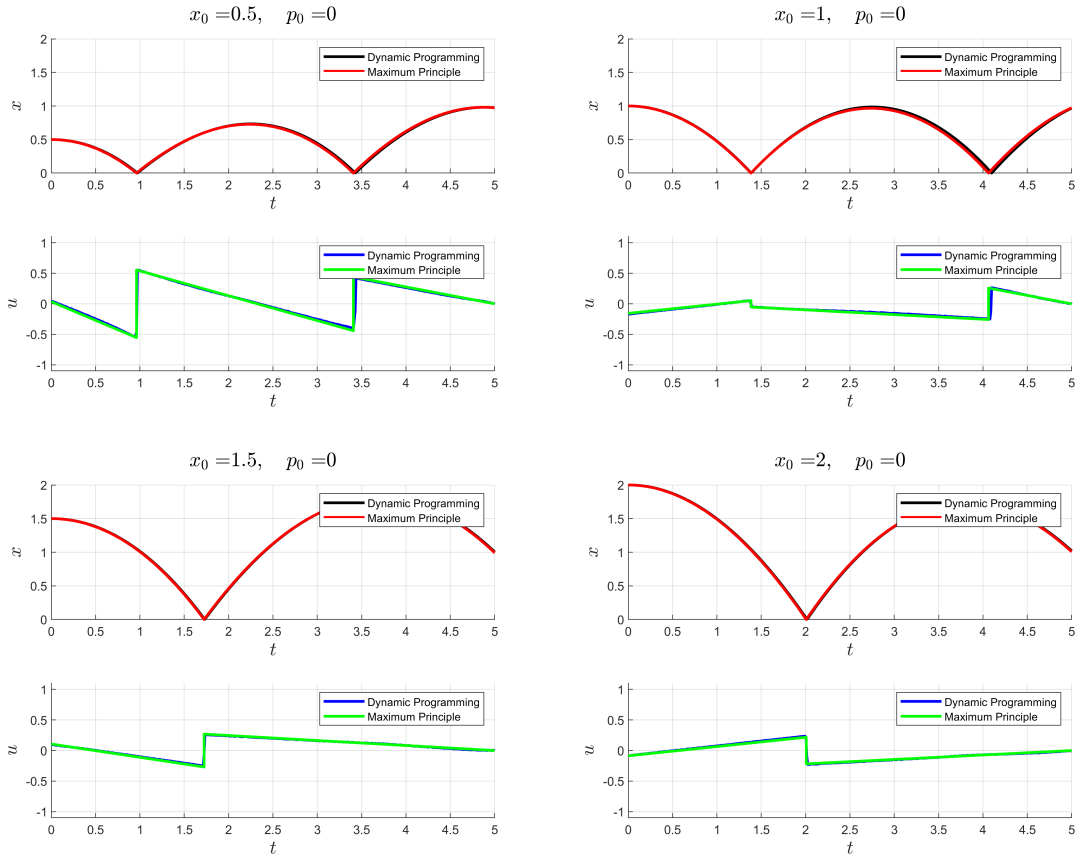


Figure 9.4: The figures above compare the optimal trajectories computed via dynamic programming along with those found via the maximum principle. For this problem $m = 2$, $g = 1$, $t_0 = 0$, $t_f = 5$, and $\alpha = 10$.

To solve the optimal control problem, we need to solve the following mixed boundary-value problem:

$$x(t_0) = x_0, \quad p_x(t_f) = 2\alpha (x(t_f) - 1),$$

$$y(t_0) = p_0, \quad p_y(t_f) = 0.$$

In Figure 9.4, we compare the results from dynamic programming in §9.3.2 with the results from the maximum principle.

9.6.2 The Hybrid Hamilton-Jacobi-Bellman Equation

The (strong) hybrid HJB equation for the bouncing ball is

$$0 = \frac{\partial J}{\partial t} - \frac{1}{2} \left(\frac{\partial J}{\partial y} \right)^2 + \frac{1}{m} y \left(\frac{\partial J}{\partial x} \right) - mg \left(\frac{\partial J}{\partial y} \right), \quad x \neq 0,$$

with the “impact condition,”

$$\begin{aligned} x \mapsto x, \quad \frac{\partial J}{\partial x} &\mapsto -\frac{\partial J}{\partial x} + \frac{2m^2g}{y} \left(\frac{\partial J}{\partial y} \right), \\ y \mapsto -y, \quad \frac{\partial J}{\partial y} &\mapsto -\frac{\partial J}{\partial y}, \end{aligned} \quad x = 0.$$

We will (numerically) solve this via the method of characteristics where the characteristics arise from the maximum principle, (9.13) and (9.14). We will compute trajectories *backwards* in time with final conditions

$$J(x, y; t_f) = \alpha(x - 1)^2, \quad \frac{\partial}{\partial x} J(x, y; t_f) = 2\alpha(x - 1), \quad \frac{\partial}{\partial y} J(x, y; t_f) = 0.$$

As before, we choose $m = 2$, $g = 1$, $t_0 = 0$, $t_f = 5$, and $\alpha = 10$. For the purposes of numerical analysis, we choose final conditions $x_f \in [0.9, 1.1]$ and $y_f \in [-4, 4]$ (that is, we backwards propagate this rectangle through time). The solutions are shown in Figure 9.5. Notice that as time recedes backwards, the optimal cost function becomes increasingly more multi-valued and complicated.

Optimal Number of Impacts

As we are able to compute the (multi-valued) optimal cost function, it is natural to ask how many impacts are in the optimal trajectory. For the parameters chosen for this example, a trajectory will require either none, one, or two impacts; see Figure 9.6.

9.6.3 Comparison with the Continuous Problem and Collision Loss

We have examined the problem of optimal control of a bouncing ball with elastic collisions. How does this compare to the purely continuous case or with inelastic collisions? To simplify the following computations, we will change the performance measure to remove the terminal condition but to now *require* that $x(T) = 1$.

$$u^* = \arg \min_u \int_0^T \frac{1}{2} u^2 dt, \quad x(T) = 1, \quad y(T) = 0.$$

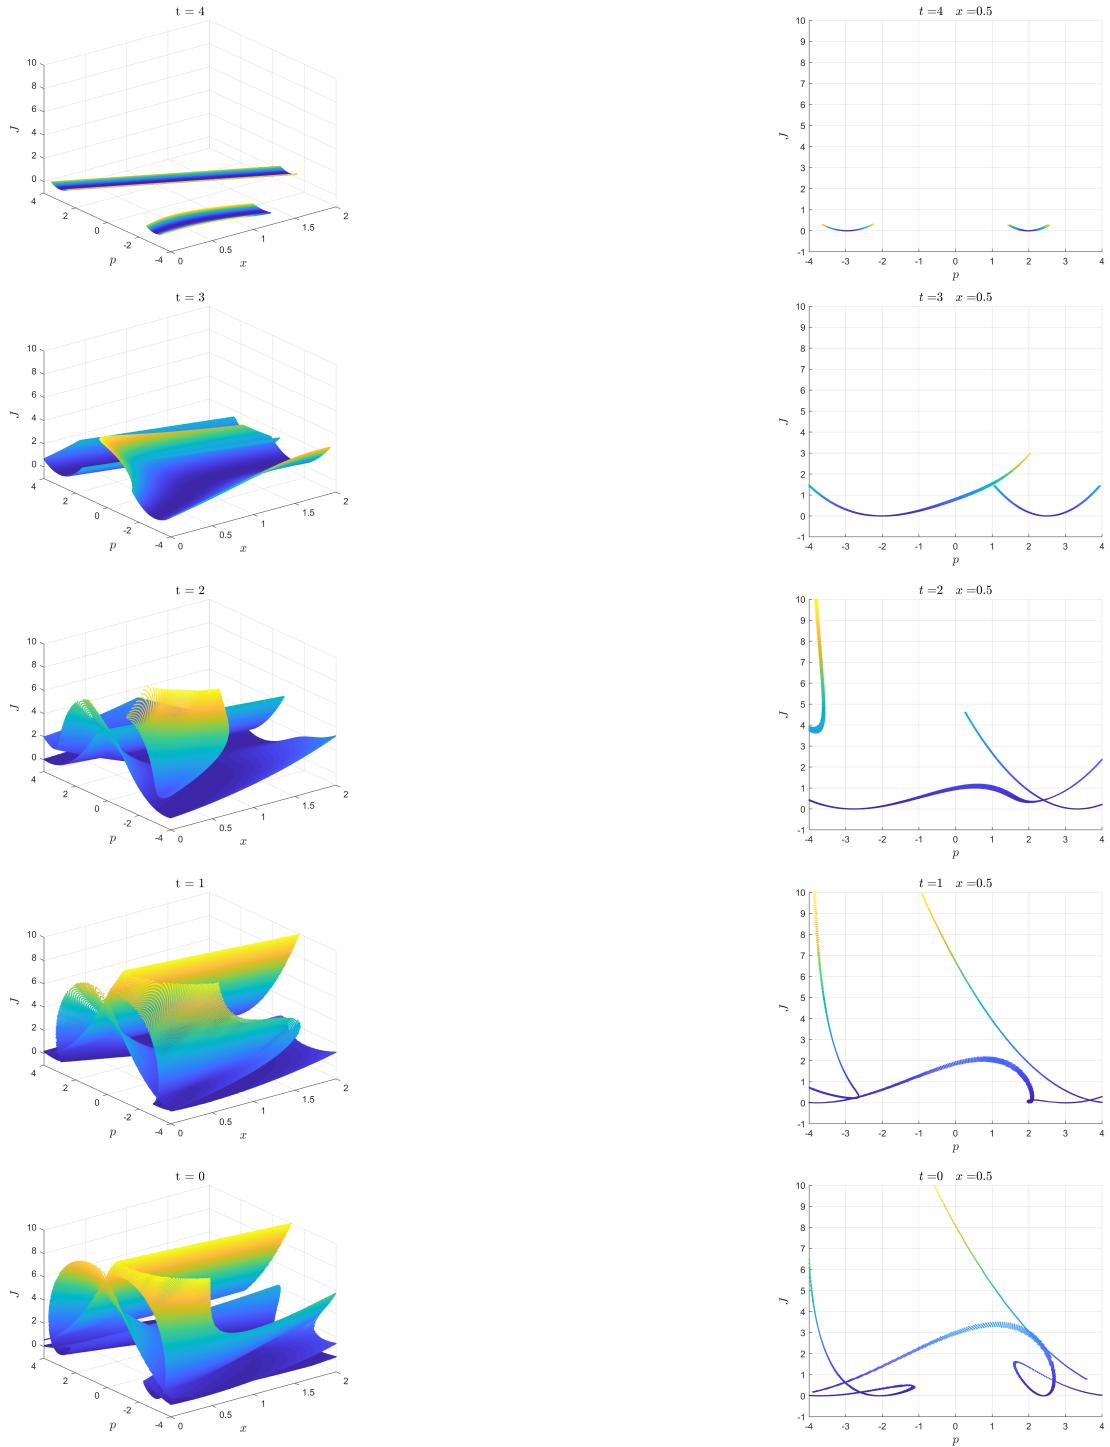


Figure 9.5: Left column: Plots of the function $J(x, p; t)$ at various time steps. Right column: Cross-sections of the solutions at $x = 0.5$.




Figure 9.6: Left: A plot of $J(x, p; 0)$ as seen in Figure 9.5. Center: A plot of $J(x, p; 0)$ where the color corresponds to the number of impacts; zero impacts is dark blue, one impact is light blue, and two impacts is yellow. Right: A plot showing how many bounces is needed to obtain the minimum cost.

Continuous Case

Let us first solve the above problem for the purely continuous case. The dynamics (9.13) can be exactly solved:

$$\begin{aligned}
 x(t) &= \frac{1}{m} \left[\frac{1}{2}(-mg - p_y^0)t^2 + \frac{1}{6m}p_x^0t^3 + y^0t \right] + x^0, \\
 y(t) &= (-mg - p_y^0)t + \frac{1}{2m}p_x^0t^2 + y^0, \\
 p_x(t) &= p_x^0, \\
 p_y(t) &= p_y^0 - \frac{1}{m}p_x^0t.
 \end{aligned}
 \tag{9.15}$$

Solving for p_x^0 and p_y^0 such that the boundary conditions are satisfied (i.e. $x(t_f) = 1$) gives $p_x^0 = 0$ and $p_y^0 = -mg$. This can be used to determine the cost since $u = -p_y = mg$:

$$J_{cont} = \frac{1}{2}m^2g^2T.$$

It would be interesting to see if $J_{hybrid} < J_{cont}$.

Inelastic Collisions

To determine the optimal cost, we must now solve the boundary value problem with the full hybrid dynamics rather than merely the continuous dynamics. Although this should be possible to do in closed-form with the help of (9.15), we will instead solve this numerically via shooting. Suppose now that energy is lost during each

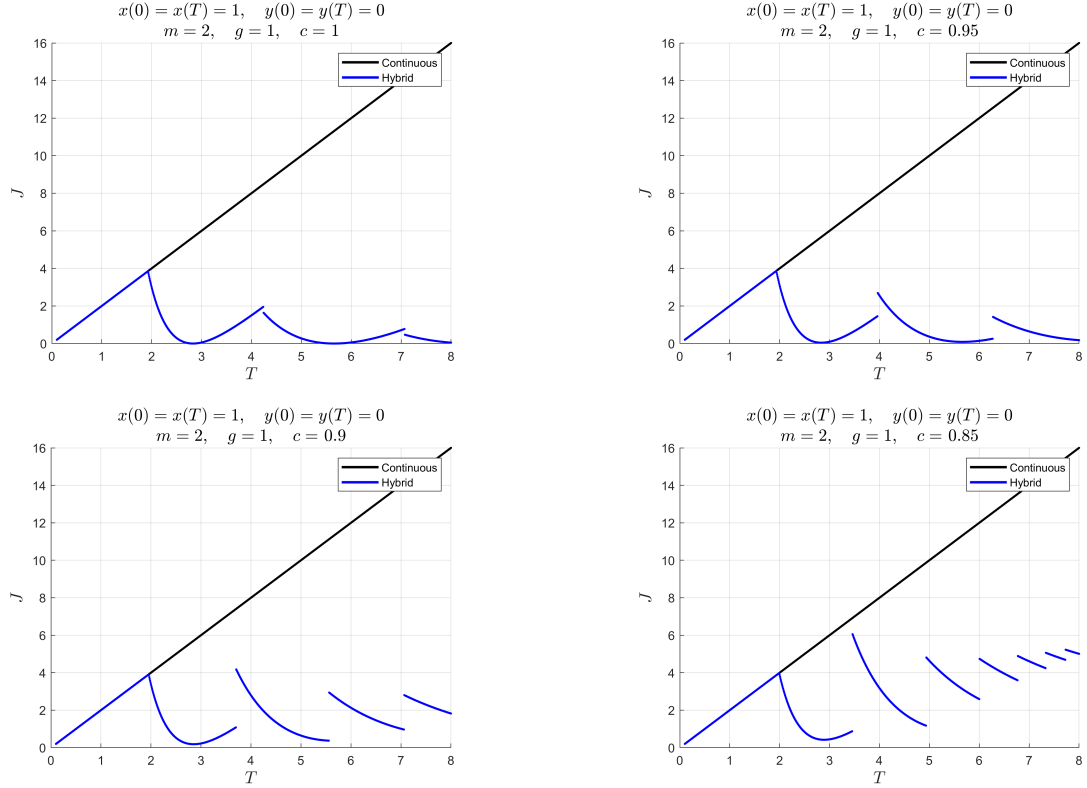


Figure 9.7: The optimal costs for the bouncing ball with various coefficients of restitution.

impact: $Ry = -c^2y$ for some $c \in [0, 1]$. The impact map now becomes:

$$\begin{aligned}
 q &\mapsto q, \\
 y &\mapsto -c^2y, \\
 p_q &\mapsto -\frac{1}{c^2}p_q + \frac{m}{2c^2} \frac{p_y^2}{y} (1 - c^{-4}) + \frac{m^2g}{c^2} \frac{p_y}{y} (1 + c^{-2}), \\
 p_y &\mapsto -\frac{1}{c^2}p_y.
 \end{aligned}$$

Although Figure 9.7 shows that hybrid controls outperform their continuous counterpart, the results shown are not *necessarily* the optimal cost for the hybrid problem; it only shows a local minimum for the multi-valued nature of J makes finding a true minimum difficult to find.

CHAPTER X

Conclusions

10.1 Summary

This thesis dealt with understanding the asymptotic nature, existence of invariant measures, and control for constrained systems - both unilateral (hybrid) and nonholonomic. However, as there are fewer results on the differential geometry of hybrid systems, more emphasis has been placed on translating hybrid systems to this context. To this end, Chapter IV is an auxiliary chapter which studied properties of hybrid dynamical systems (which are already standard for their continuous counterparts). These properties include: limit sets (Propositions IV.5 and IV.9), Floquet theory (Theorem IV.16), and an extension of the famous Poincaré-Bendixson theorem (Theorem IV.26).

The main results from this work dealt with finding (smooth) invariant measures for both nonholonomic and hybrid systems. To achieve this, Chapter VI outlined hybrid mechanical systems as well as systems experiencing *both* nonholonomic and hybrid constraints; in particular, hybrid versions of Noether's theorem are proved (Theorems VI.31 and VI.37). To find invariant measures, nonholonomic systems (without impacts) are discussed first in Chapter V in which reasonably testable, necessary, and sufficient conditions are laid out for the existence of invariant measures depen-

ding only on the configuration variables (Theorem V.24). Incorporating impacts, Chapter VII extends the idea of invariant measures to hybrid systems. Verifiable conditions are laid out (Theorem VII.3) to determine when a given differential form is preserved under the hybrid flow. This provides an extension to the nonholonomic case in which an invariant measure depending only on the configuration variables is preserved in the presence of *any* impact (Theorem VII.21).

The final part of this work is on integrability and control of hybrid systems. Chapter VIII extends Hamilton-Jacobi theory to hybrid systems, as well as, the notion of complete integrability. Both of these concepts utilize hybrid Lagrangian submanifolds. These submanifolds are the hybrid analogue of the usual Lagrangian submanifolds and obey many of the usual properties (Proposition VIII.8). To serve as a test-bed, the billiard problem on a circular table is proven to be integrable in this sense (§8.2.2) by constructing a hybrid Lagrangian foliation. Finally, Chapter IX extends the idea of hybrid Hamilton-Jacobi theory to optimal control problems with the hybrid Hamilton-Jacobi-Bellman equation (Theorem IX.10) and the hybrid maximum principle (Theorem IX.11).

10.2 Future Work

10.2.1 Hybrid Tongues

Continuous dynamical systems subject to periodic forcing, e.g. differential equations of the form $\ddot{x} + f(t)x = 0$ where f is periodic - known as Hill's equations [61], are known to have “tongues” in their stability diagrams. A similar structure appears in the stability domain for the hybrid model of walking with foot slip (Figure 4.5). This seems to make sense from the point of view that periodic walking induces periodic forcing (where we can think of $f(t)$ above as an impulse).

We note that the stability diagrams in Figure 10.1 do have some important diffe-




Figure 10.1: Left: Stability plot for the walker with foot-slip, cf. Figure 4.5. Right: Stability plot for Mathieu's equation $\ddot{x} + [\delta + \varepsilon \cdot \cos(t)]x = 0$, taken from §9.2 in [78].

rences; the stability plot for foot-slip is based on initial conditions and the tongues appear on the stable part while the stability plot for Mathieu's equation is based in parameter space and the tongues are regions of instability. To make the comparison more apt, Figure 10.2 shows the stability plot for the foot-slip model in parameter space based off of the fixed initial condition $(\xi_0, \eta_0) = (-1, 0)$ rather than the adaptive scheme discussed in Appendix A. However, the tongues are still backwards in the sense that they denote regions of stability rather than instability.

A powerful way to study Hill's equations (and in particular Mathieu's equation, cf. §9.2 in [78]) is by computing an asymptotic approximation for the solution. An object of future study is to explore this connection in better detail; particularly to apply ideas from asymptotic analysis to study stability domains in hybrid systems.

10.2.2 Measure Shaping Controls

Suppose we have a single constraint $\eta \in \Omega^1(Q)$. The density form is given by

$$\vartheta = \frac{1}{\eta(\eta^\sharp)} \mathcal{L}_{\eta^\sharp} \eta,$$

where $g(\eta^\sharp, \cdot) = \eta$ and g is the metric coming from the Lagrangian. Lagrangians have been chosen to obtain certain stability properties by shaping their energies [12], so




Figure 10.2: Left: A modified version of the results from Figure 4.4 where stability is based off of the initial conditions $(\xi_0, \eta_0) = (-1, 0)$ rather than the adaptive initial conditions discussed in Appendix A. Right: A zoomed in version of the left figure.

we will choose a Lagrangian that shapes measure instead, i.e. such that ϑ is exact.

This reduces to solving

$$\frac{1}{g^{ij}\eta_i\eta_j} \frac{\partial \eta_k}{\partial x^\ell} (g^{i\ell}\eta_i dx^k - g^{ik}\eta_i dx^\ell) = \frac{g^{i\ell}\eta_i}{g^{ij}\eta_i\eta_j} \left(\frac{\partial \eta_k}{\partial x^\ell} - \frac{\partial \eta_\ell}{\partial \eta_k} \right) dx^k =: d\alpha.$$

Then we have an invariant density

$$D_L := \eta(\eta^\sharp) \cdot \exp(\alpha),$$

which depends on the chosen Lagrangian. Denote the corresponding measure by μ_L and suppose that the flow is now *ergodic* on an invariant set E with $\mu_L(E) < \infty$.

Then the ergodic theorem states that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \chi_A(\varphi_t(x)) dt = \frac{\mu_L(A)}{\mu_L(E)},$$

where χ_A is the characteristic function for the set A . In particular if A is a null set, then the flow will (on average) almost never enter A . This can be used for safety purposes by the following recipe:

1. Let A be a set of “dangerous” states.
2. Find a Lagrangian with a corresponding density, D_L , which is “small” on A .

Another interesting application of measure-preserving flows is a result from [68] which states that for an affine control system, in which the drift vector field is measure-preserving and the state-space is Riemannian and compact, then accessibility implies controllability. This is helpful because if we can choose a Lagrangian such that measure is preserved, then if we subject the system to controls which implies accessibility we automatically get controllability.

10.2.3 Hybrid Integrators

The work done in [91] shows that elastic mechanical impacts preserve the symplectic form and constructs a hybrid integrator which preserves this form. An object of future work will be to attempt to construct numerical hybrid integrators which will preserve an arbitrary form $\alpha \in \mathcal{A}_{\mathcal{H}}$.

For nonholonomic systems, it is no longer generally true that the symplectic form is preserved but numerical integrators can still be constructed that preserve important structures, cf. Chapter 7 of [80] and the references therein. Specific examples of forms to be preserved under hybrid integrators include: ω and H for holonomic systems, and $f\mu_{\mathcal{E}}$ and H for nonholonomic systems (which have an invariant measure depending only on base variables such as the vertical rolling disk).

10.2.4 Hybrid Bracket

The definition of a completely integrable hybrid system, Definition VIII.6, involves hybrid Lagrangian submanifolds rather than constants of motion in involution as is normal with their continuous counterparts. The reason for this is that defining a hybrid bracket has some inherent difficulties which include; at least, the four issues at the end of §8.2.2. Uniqueness issues with (P.3) and (P.4) *might* be overcome by requiring skewness (P.2) to extend beyond quadratic Hamiltonians.

Skewness of the bracket in mechanics is paramount to providing energy-conservation which is a fundamental requirement for elastic impacts as in (6.2) for holonomic or (6.6) for nonholonomic. Therefore, we hope that we can resolve (P.2) in the case of only quadratic Hamiltonians so (P.3) and (P.4) do not interfere. The following example, however, demonstrates that this is *not* the case.

Example X.1. Let $Q = \mathbb{R}^2$ and $S = \{x_2 = 0\}$. Choose two Hamiltonians:

$$f(x^1, x^2, p_1, p_2) = \frac{1}{2}(p_1^2 + p_2^2),$$

$$g(x^1, x^2, p_1, p_2) = \frac{3}{2}(p_1^2 + p_2^2) - p_1 p_2.$$

As these are both regular Hamiltonians, we can compute their impact maps according to (6.2):

$$\mathfrak{D}(f)(x^1, x^2, p_1, p_2) = (x^1, x^2, p_1, -p_2),$$

$$\mathfrak{D}(g)(x^1, x^2, p_1, p_2) = (x^1, x^2, p_1, \frac{2}{3}p_1 - p_2).$$

Checking the “hybrid bracket,” we get

$$\mathfrak{D}(f)^* g - g = 2p_1 p_2,$$

$$\mathfrak{D}(g)^* f - f = \frac{2}{9}p_1^2 - \frac{2}{3}p_1 p_2,$$

which is clearly not skew.

This example shows that something more subtle is happening with the hybrid bracket. Therefore more work needs to be done to understand and hopefully to resolve this phenomenon.

APPENDICES

APPENDIX A

Numerical Methods for the Foot-slip Model

In this appendix, we explain the numerical methods and algorithms used to determine the dynamics on the Poincaré section for the foot-slip walking model studied in §4.2.3. The contents here can also be found in [26].

Recall that the hybrid dynamical system under consideration has continuous dynamics

$$\begin{aligned}\theta' &= \varepsilon\xi \\ \xi' &= C \left[\sin\theta (\cos\alpha + \varepsilon\xi^2 \cos\theta) - \cos\theta (\eta + \xi \cos\theta) \right] \\ \eta' &= \sin\alpha - \eta - \xi \cos\theta\end{aligned}$$

where

$$C = \frac{\mu}{1 + \mu \sin^2\theta}.$$

The impact occurs when $\theta = -\delta$ and the impact map is

$$\Delta(\theta, \xi, \eta) = (\theta + 2\delta, \cos(2\delta)\xi, \eta + \cos\delta [1 - \cos(2\delta)]\xi).$$

This model includes four parameters: ε , μ , α , and δ .

Below, we discuss the numerical implementation for finding regions in parameter space where a periodic orbit exists as well as finding the basin of attraction for a given periodic orbit exists, see Figures 4.4 and 4.5.

A.1 Stable Parameters

Determining which parameters result in the basin of attraction of a periodic orbit being non empty generally requires an expensive computation. This requires searching for fixed points of the Poincaré map (each evaluation requires integrating equation (4.11) which is expensive). Due to this issue, we will only test a single initial condition for each set of parameters. To accomplish this, we will need a systematic way to compute meaningful guesses for initial conditions that lie within the basin of attraction for the stable periodic orbits. In the case where the foot does not slide, the equations of motion can be expressed as

$$\ddot{\theta} = \frac{g}{\ell} \sin(\theta - \alpha) = \frac{g}{\ell_f} \left(\frac{\mu}{1 + \mu} \right) \sin(\theta - \alpha).$$

We would, of course, like to rescale time so this matches with equation (4.11). Doing so yields

$$\theta'' = \varepsilon \left(\frac{\mu}{1 + \mu} \right) \sin(\theta - \alpha).$$

This allows us to come up with our approximating hybrid system:

$$\begin{aligned} \theta' &= z \\ z' &= \frac{\varepsilon \mu}{1 + \mu} \sin(\theta - \alpha). \end{aligned}$$

With the impact surface and map being, $\tilde{S} = \{\theta = -\delta\}$ and

$$\tilde{\Delta}(\theta, z) = (\theta + 2\delta, \cos(2\delta)z).$$

Call the steady-state values for θ and z (immediately after impact) to be $\bar{\theta}$ and \bar{z} . Clearly $\bar{\theta} = \delta$. Using (4.11), we can turn z to ξ . Additionally, assuming that we start

with zero foot slip, we get the initial conditions to be:

$$\begin{aligned}\theta(0) &= \delta, \\ \xi(0) &= \frac{1}{\varepsilon} \bar{z}, \\ \eta(0) &= -\frac{1}{\varepsilon} \bar{z} \cos \delta.\end{aligned}$$

We will then explore parameter space to see when these initial conditions achieve a steady-state periodic orbit.

A.2 Stable Initial Conditions

We present two calculations: determining the region of initial conditions on $\Delta(S)$ (i.e. (ξ_0, η_0)) that result in stable walking, and finding the corresponding discrete dynamics $P : \Delta(S) \rightarrow \Delta(S)$.

To determine which initial conditions yield a (stable) periodic orbit, we implement the following numerical algorithm: Let $G : \mathbb{R}^2 \rightarrow \mathbb{N} \cup \{\infty\}$, where $G(\xi_0, \eta_0)$ is the number of steps the walker takes with initial conditions (δ, ξ_0, η_0) before it crashes. If no crashes occur, $G = \infty$. For practicality, we say that a periodic orbit exists if $G > 50$. These initial conditions are recorded as blue in Figure 4.5. On the other hand, if $G = 0$, a crash happens immediately, which is recorded as red. The white region corresponds to initial conditions where $0 < G < 50$.

APPENDIX B

Technical Lemmas from Continuous and Discrete Dynamics

In this appendix, we present a collection of technical lemmas that are used in proving Theorem IV.26 in §4.3.

Before we state the following lemmas, we need to define the discrete ω -limit set. Given a discrete dynamical system, $P : S \rightarrow S$, the discrete ω_d -limit set is defined as

$$\omega_d(x; P) = \left\{ y \in S : \exists N_n \rightarrow \infty \text{ s.t. } \lim_{n \rightarrow \infty} P^{N_n}(x) = y \right\}.$$

When the map P is understood, the limit-set will be called $\omega_d(x)$.

Lemma B.1 (See Lemma 4.2 in [29]). *Let $P : [a, b] \rightarrow [a, b]$ be C^1 and injective. Then for all $x \in [a, b]$, $\omega_d(x)$ is either a single point or two points. In either case, all trajectories approach a periodic orbit.*

Proof. By being injective, the map P is automatically monotone. Let $I := [a, b]$, and define the fixed-point set $F := \{x \in I : P(x) = x\}$, and consider $x \in I$. If $x \in F$ we are done, so assume $x \notin F$.

First, assume that P is nondecreasing. Then, since $P(x) \geq x$ implies $P^2(x) \geq P(x)$ and $P(x) \leq x$ implies $P^2(x) \leq P(x)$, it follows (since $x \notin F$) that $\{P^n(x)\}$ is a monotone and bounded sequence and therefore converges. Hence $\omega_d(x)$ is a singleton.

To complete the proof, consider the case where P is nonincreasing. Since $\omega_d(x; P) =$

$\omega_d(x; P^2) \cup \omega_d(P(x); P^2)$, the previous paragraph shows that $\omega_d(x)$ is either two points or a singleton. \square

Lemma B.2 (See Lemma 6.1 in [29]). *Let $S \subset \mathbb{R}$ be a finite set and $P : S \rightarrow S$. Then, for all $x \in S$ there exists $N \neq M$ large enough such that $P^N(x) = P^M(x)$ where*

$$P^N(x) = \underbrace{P \circ P \circ \dots \circ P}_{N \text{ times}}(x).$$

In particular, $\omega_d(x)$ is a periodic orbit.

Proof. This follows immediately from the Pigeonhole principle. We have a function $f_x : \mathbb{Z}^+ \rightarrow S$, $m \mapsto P^m(x)$, which cannot be injective. Therefore, there exists N and M such that $f_x(N) = f_x(M)$. \square

The following lemma is central to Theorem IV.26 so we will first review how this lemma fits into the theorem. Let $(\mathcal{N}, \mathcal{E}, \mathcal{X}, S, \Delta, f)$ be an essentially non-beating general hybrid dynamical system such that $(\mathcal{N}, \mathcal{E})$ is a cycle. Define the set

$$S_{(1,2)}^\infty := \{x \in S_{(1,2)} : \forall t > 0, \exists T > t \text{ s.t. } \varphi_T^{\mathcal{H}}(x) \in S_{(1,2)}\}.$$

For a compact, invariant set $F = \sqcup_i F_i \subset \mathcal{X}$, two assumptions in Theorem IV.26 that will appear in this proof are:

(Q.4) Let the outdegree of vertex n be $\text{od}(n)$. Then, each F_n can be written as a disjoint union of $\text{od}(n)$ many compact sets,

$$F_n = \bigsqcup_{\substack{e \in \mathcal{E} \\ \text{dom}(e)=n}} F_n^e.$$

(Q.5) The set $F_n^e \cap S_{e'}$ is only nonempty if $e = e'$ and $F_n^e \cap S_e$ is diffeomorphic to an interval.

Lemma B.3. *The set $S_{(1,2)}^\infty$ is either a point or an interval.*

Proof. We start by defining the sets $S_{(1,2)}^m$. Let P be the return map, $P : U \subset S_{(1,2)} \rightarrow S_{(1,2)}$. Then, $S_{(1,2)}^1 := \{x \in S_{(1,2)} : P(x) \in S_{(1,2)}\}$, that is, points in $S_{(1,2)}$ that return to $S_{(1,2)}$ *at least* once. Similarly, $S_{(1,2)}^m$ are points of $S_{(1,2)}$ that return to $S_{(1,2)}$ *at least* m times. This allows us to express $S_{(1,2)}^\infty$ as

$$S_{(1,2)}^\infty = \bigcap_{m=1}^{\infty} S_{(1,2)}^m.$$

If we can show that each $S_{(1,2)}^m$ is an interval, then the desired result follows due to nesting.

The base case is satisfied by (Q.4) via $S_{(1,2)}^0 := S_{(1,2)}$. We will continue by induction.

For each m , we can iterate by

$$S_{(1,2)}^{m+1} = S_{(1,2)}^m \cap P^{-1}\left(S_{(1,2)}^m\right).$$

Therefore, if $S_{(1,2)}^m$ and $P^{-1}\left(S_{(1,2)}^m\right)$ are both intervals, so is $S_{(1,2)}^{m+1}$. All that is left to prove is $P^{-1}\left(S_{(1,2)}^m\right)$ is an interval. Before we can study the structure of $P^{-1}\left(S_{(1,2)}^m\right)$, we first define the family of functions

$$\mu_i : \mathcal{D}_i \subset \mathcal{X}_i \rightarrow S_{(i,i+1)}.$$

Where $\mathcal{D}_i = \{x \in \mathcal{X}_i : \exists t > 0 \text{ with } \varphi_t^i(x) \in S_{(i,i+1)}\}$ and $\mu_i(x) = \varphi_t^i(x) \in S_{(i,i+1)}$. These functions can be thought of as “projecting” onto the sets $S_{(i,i+1)}$. What remains to show is that $\Delta_n(S_{(n,1)}) \cap \mathcal{D}_1$ is an interval. If we can show this, we are done because $\Delta_{(n,1)}^{-1}\left(\Delta_{(n,1)}(S_{(n,1)}) \cap \mathcal{D}_1\right)$ is also an interval, and we can march backwards around the cycle and end up on with a subinterval of $S_{(1,2)}$.

Since $S_{(n,1)}$ is diffeomorphic to an interval, so is the set $\Delta_{(n,1)}(S_{(n,1)})$. Let the map $h : \Delta_{(n,1)}(S_{(n,1)}) \rightarrow [a, b]$ be a diffeomorphism. Define the points \tilde{a} and \tilde{b} as

$$\tilde{b} = \max\{x \in [a, b] : o_c^+(h^{-1}(x)) \cap S_{(1,2)} \neq \emptyset\},$$

$$\tilde{a} = \min\{x \in [a, b] : o_c^+(h^{-1}(x)) \cap S_{(1,2)} \neq \emptyset\}.$$

We claim that $\Delta_{(n,1)}(S_{(n,1)}) \cap \mathcal{D}_1 = h^{-1}[\tilde{a}, \tilde{b}]$. By the choice of \tilde{a} and \tilde{b} , we know that $\Delta_{(n,1)}(S_{(n,1)}) \cap \mathcal{D}_1 \subset h^{-1}[\tilde{a}, \tilde{b}]$. To show the other direction, we consider the region bounded by the four curves: $h^{-1}[\tilde{a}, \tilde{b}]$, $\tilde{o}_c^+(h^{-1}(\tilde{a}))$, $\tilde{o}_c^+(h^{-1}(\tilde{b}))$, and $S_{(1,2)}$ (where $\tilde{o}^+(x)$ is the forward orbit of x until it impacts $S_{(1,2)}$). By assumption (Q.5) and uniqueness of solutions, any curve starting on $h^{-1}[\tilde{a}, \tilde{b}]$ must intersect $S_{(1,2)}$ before leaving the region described above (see Lemma 10 in [29]). Therefore, $h^{-1}[\tilde{a}, \tilde{b}] \subset \Delta_n(S_n) \cap \mathcal{D}_1$. \square

Lemma B.4. *Let $(\mathcal{N}, \mathcal{E}, \mathcal{X}, S, \Delta, f)$ be a GHDS where $(\mathcal{N}, \mathcal{E})$ is a cycle. Let $x_0 \in S_{(i-1,i)} \setminus \text{fix}(f_{i-1})$ be such that there exists a time, $T_0 > 0$, where $\varphi_{T_0}^i(\Delta_{(i-1,i)}(x_0)) \in S_{(i,i+1)}$ (φ_t^i is the continuous flow corresponding to $\dot{x} = f_i(x)$ in vertex i). Additionally, assume that the flow intersects the surface, $S_{(i,i+1)}$, transversely. Then there exists an $\varepsilon > 0$ and a C^1 function $\tau : \mathcal{B}_\varepsilon(x_0) \cap S_{(i-1,i)} \rightarrow \mathbb{R}^+$ such that for all $y \in \mathcal{B}_\varepsilon(x_0) \cap S_{(i-1,i)}$, $\varphi_{\tau(y)}^i(\Delta_{(i-1,i)}(y)) \in S_{(i,i+1)}$.*

Proof. Define the function

$$F_i : (0, +\infty) \times S_{(i-1,i)} \rightarrow \mathcal{X}_i,$$

by $F_i(t, x) = h_{(i,i+1)}(\varphi_t^i(\Delta_{(i-1,i)}(x)))$ where $h_{(i,i+1)}^{-1}(0) = S_{(i,i+1)}$ with zero as a regular value. It follows from Theorem 1 in Section 2.5 in [92] that $F_i \in C^1(\mathbb{R}^+ \times S_{(i-1,i)})$. This allows the use of the implicit function theorem. By the assumptions of the lemma, we know that $F_i(T_0, x_0) = 0$. Differentiating F_i with respect to time yields:

$$\frac{\partial F_i}{\partial t}(T_0, x_0) = \frac{\partial h_{(i,i+1)}}{\partial y} \Big|_{y=\varphi_{T_0}^i(\Delta_{(i-1,i)}(x_0)) \in S_{(i,i+1)}} \cdot f_i(\varphi_{T_0}^i(\Delta_{(i,i+1)}(x_0))) \neq 0.$$

The first factor is nonzero because zero is a regular value and the second is nonzero because we are away from fixed points of the continuous flow. Their inner product is nonzero because of the transversality condition. This allows the use of the implicit

function theorem (see Theorem 9.28 in [98]) to show that there exists a neighborhood of x_0 and a C^1 function τ with the desired properties. \square

Note that this implies that the maps $P_i : U_i \subset S_{(i-1,i)} \rightarrow S_{(i,i+1)}$ are all C^1 . Because the composition of C^1 maps is still C^1 , the map $P := P_1 \circ \dots \circ P_n : U \subset S_{(1,2)} \rightarrow S_{(1,2)}$ is C^1 .

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] R. Abraham and J.E. Marsden. Foundations of Mechanics. AMS Chelsea publishing. AMS Chelsea Pub./American Mathematical Society, 2008.
- [2] R. Alur, C. Belta, F. Ivančić, V. Kumar, M. Mintz, G. J. Pappas, H. Rubin, and J. Schug. Hybrid modeling and simulation of biomolecular networks. In International workshop on hybrid systems: Computation and control, pages 19–32. Springer, 2001.
- [3] 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, 138(1):3 – 34, 1995. Hybrid Systems.
- [4] A. Ames. A Categorical Theory of Hybrid Systems. PhD thesis, University of California, Berkeley, 2006.
- [5] A. D. Ames, , R. D. Gregg, and S. Sastry. Is there life after Zeno? taking executions past the breaking (Zeno) point. In 2006 American Control Conference, pages 6 pp.–, June 2006.
- [6] V.I. Arnold, K. Vogtmann, and A. Weinstein. Mathematical Methods of Classical Mechanics. Graduate Texts in Mathematics. Springer New York, 2013.
- [7] E. Artin. Ein mechanisches System mit quasiergodischen Bahnen. In Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, volume 3, pages 170–175. Springer, 1924.
- [8] V. Azhmyakov, S. A. Attia, and J. Raisch. On the maximum principle for impulsive hybrid systems. In International Workshop on Hybrid Systems: Computation and Control, pages 30–42. Springer, 2008.
- [9] R. Bellman. The theory of dynamic programming. Technical report, Rand corp santa monica ca, 1954.
- [10] I. Bendixson. Sur les courbes définies par des équations différentielles. Acta Math., 24:1–88, 1901.
- [11] R. Berndt, M. Klucznik, and American Mathematical Society. An Introduction to Symplectic Geometry. Graduate Studied in Mathematics. American Mathematical Society, 2001.
- [12] A. M. Bloch, D. E. Chang, N. E. Leonard, and J. E. Marsden. Controlled Lagrangians and the stabilization of mechanical systems. ii. potential shaping. IEEE Transactions on Automatic Control, 46(10):1556–1571, Oct 2001.
- [13] A. M. Bloch, O. E. Fernandez, and T. Mestdag. Hamiltonization of nonholonomic systems and the inverse problem of the calculus of variations. Reports on Mathematical Physics, 63(2):225–249, 2009.
- [14] A. M. Bloch, P. S. Krishnaprasad, J. E. Marsden, and R. M. Murray. Nonholonomic mechanical systems with symmetry. Archive for Rational Mechanics and Analysis, 136(1):21–99, Dec 1996.

- [15] A. M. Bloch, J. E. Marsden, and D. V. Zenkov. Quasivelocities and symmetries in non-holonomic systems. Dynamical systems, 24(2):187–222, 2009.
- [16] A.M. Bloch, J. Baillieul, P. Crouch, J.E. Marsden, D. Zenkov, P.S. Krishnaprasad, and R.M. Murray. Nonholonomic Mechanics and Control. Interdisciplinary Applied Mathematics. Springer New York, 2015.
- [17] B. Brogliato. Nonsmooth Mechanics. Communications and Control Engineering. Springer International Publishing, 2016.
- [18] S. A. Burden, S. Revzen, and S. S. Sastry. Model reduction near periodic orbits of hybrid dynamical systems. IEEE Transactions on Automatic Control, 60(10):2626–2639, 2015.
- [19] S. A. Burden, S. S. Sastry, D. E. Koditschek, and S. Revzen. Event–selected vector field discontinuities yield piecewise–differentiable flows. SIAM Journal on Applied Dynamical Systems, 15(2):1227–1267, 2016.
- [20] D. Carnevale, A. R. Teel, and D. Nesic. A Lyapunov proof of an improved maximum allowable transfer interval for networked control systems. IEEE Transactions on Automatic Control, 52(5):892–897, May 2007.
- [21] A. Casagrande, V. Mysore, C. Piazza, and B. Mishra. Independent dynamics hybrid automata in systems biology. In Proceedings of the First International Conference on Algebraic Biology (AB05), pages 61–73, 2005.
- [22] H. Cendra, J. E. Marsden, and T. S. Ratiu. Geometric Mechanics, Lagrangian Reduction, and Nonholonomic Systems, pages 221–273. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001.
- [23] G. Chartrand, L. Lesniak, and P. Zhang. Graphs & Digraphs, Fifth Edition. Chapman & Hall book. Taylor & Francis, 2010.
- [24] H. Chen, Z. Bell, R. Licitra, and W. Dixon. A switched systems approach to vision-based tracking control of wheeled mobile robots. In 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pages 4902–4907, Dec 2017.
- [25] N. Chernov and R. Markarian. Introduction to the Ergodic Theory of Chaotic Billiards. Publicações matemáticas. Instituto nacional de matematica pura e aplicada, IMPA, 2003.
- [26] W. Clark and A. Bloch. Stable orbits for a simple passive walker experiencing foot slip. In 2018 IEEE Conference on Decision and Control (CDC), pages 2366–2371. IEEE, 2018.
- [27] W. Clark and A. Bloch. The bouncing penny and nonholonomic impacts. In 2019 IEEE 58th Conference on Decision and Control (CDC), pages 2114–2119, Dec 2019.
- [28] W. Clark and A. Bloch. A Poincaré–Bendixson theorem for hybrid dynamical systems on directed graphs. Mathematics of Control, Signals, and Systems, 32(1):1, 2020.
- [29] W. Clark, A. Bloch, and L. Colombo. A Poincaré–Bendixson theorem for hybrid systems. Mathematical Control & Related Fields, 0(2156-8472.2019.0.2):0, 2019.
- [30] S. Collins, A. Ruina, R. Tedrake, and M. Wisse. Efficient bipedal robots based on passive-dynamic walkers. Science, 307(5712):1082–1085, 2005.
- [31] F. Colonius and W. Kliemann. Dynamical Systems and Linear Algebra:. Graduate Studies in Mathematics. American Mathematical Society, 2014.
- [32] M. Crampin and B. Heal. On the chaotic behaviour of the tent map. Teaching Mathematics and its Applications: An International Journal of the IMA, 13(2):83–89, 1994.

- [33] A. Cristofaro, C. Possieri, and M. Sassano. Time-optimal control for the hybrid double integrator with state-driven jumps. In 2019 IEEE Conference on Decision and Control (CDC), pages 6301–6306, Dec 2019.
- [34] A.C. da Silva, A. Weinstein, American Mathematical Society, Berkeley Center for Pure, and Applied Mathematics. Geometric Models for Noncommutative Algebras. Berkeley mathematics lecture notes. American Mathematical Society, 1999.
- [35] J. Daafouz, R. Postoyan, and P. Riedinger. Parameter and state estimation of switched affine systems. In 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pages 5316–5321, Dec 2017.
- [36] J. Daafouz, P. Riedinger, and C. Iung. Stability analysis and control synthesis for switched systems: a switched Lyapunov function approach. IEEE transactions on automatic control, 47(11):1883–1887, 2002.
- [37] M. de León, J. C. Marrero, and D. M. de Diego. Mechanical systems with nonlinear constraints. International Journal of Theoretical Physics, 36(4):979–995, Apr 1997.
- [38] M. De León, J. C. Marrero, and D. M. de Diego. Linear almost Poisson structures and Hamilton-Jacobi equation. applications to nonholonomic mechanics. arXiv preprint arXiv:0801.4358, 2008.
- [39] M. Della Rossa, A. Tanwani, and L. Zaccarian. Max-min Lyapunov functions for switching differential inclusions. In 2018 IEEE Conference on Decision and Control (CDC), pages 5664–5669, Dec 2018.
- [40] B. Ding, C. Ding, et al. Poincaré recurrence theorem in impulsive systems. Topological Methods in Nonlinear Analysis, 49(2):577–585, 2017.
- [41] A. B. Dishliev and D. D. Bainov. Continuous dependence on the initial condition of the solution of a system of differential equations with variable structure and with impulses. Publications of the Research Institute for Mathematical Sciences, 23(6):923–936, 1987.
- [42] R. D. Euzébio and M. R. A. Gouveia. Poincaré recurrence theorem for non-smooth vector fields. Zeitschrift für angewandte Mathematik und Physik, 68(2):40, Feb 2017.
- [43] L. C. Evans. Partial Differential Equations. Graduate Studies in Mathematics. American Mathematical Society, 2010.
- [44] R. Feres and G. Yablonsky. Knudsen’s cosine law and random billiards. Chemical Engineering Science, 59(7):1541 – 1556, 2004.
- [45] O. E. Fernandez. The Hamiltonization of nonholonomic systems and its applications. PhD thesis, University of Michigan, 2009.
- [46] G. Floquet. Sur les équations différentielles linéaires à coefficients périodiques. Annales scientifiques de l’École Normale Supérieure, 2e série, 12:47–88, 1883.
- [47] G. Forni. The cohomological equation for area-preserving flows on compact surfaces. Electronic Research Announcements of the American Mathematical Society, 1(3):114–123, 1995.
- [48] P. Forni and D. Angeli. Smooth Lyapunov functions for multistable hybrid systems on manifolds. In 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pages 5481–5486, Dec 2017.
- [49] A. T. Fuller. Bibliography of Pontryagm’s maximum principle. Journal of Electronics and Control, 15(5):513–517, 1963.

- [50] I.M. Gelfand and S.V. Fomin. Calculus of Variations. Dover Publications, 2000.
- [51] R. Goebel, R. G. Sanfelice, and A. R. Teel. Invariance principles for switching systems via hybrid systems techniques. Systems & Control Letters, 57(12):980–986, 2008.
- [52] R. Goebel, R. G. Sanfelice, and A. R. Teel. Hybrid dynamical systems. Princeton University Press, 2012.
- [53] J. Goodman and L. Colombo. On the Existence and Uniqueness of Poincaré Maps for Systems with Impulse Effects. arXiv e-prints, page arXiv:1810.05426, Oct 2018.
- [54] J. W. Grizzle, G. Abba, and F. Plestan. Asymptotically stable walking for biped robots: analysis via systems with impulse effects. IEEE Transactions on Automatic Control, 46(1):51–64, Jan 2001.
- [55] V. Guillemin and D. Kazhdan. On the cohomology of certain dynamical systems. Topology, 19(3):291 – 299, 1980.
- [56] W. M. Haddad, V.S. Chellaboina, and S. G. Nersesov. Impulsive and Hybrid Dynamical Systems: Stability, Dissipativity, and Control. Princeton Series in Applied Mathematics. Princeton University Press, 2006.
- [57] S. H. J. Heijmans, W. P. M. H. Heemels, D. Nešić, and R. Postoyan. A generalized hybrid Lyapunov proof for networked control systems: Improving the maximum allowable transmission interval. In 2018 IEEE Conference on Decision and Control (CDC), pages 428–433, Dec 2018.
- [58] S. H. J. Heijmans, R. Postoyan, D. Nei, and W. P. M. H. Heemels. Computing minimal and maximal allowable transmission intervals for networked control systems using the hybrid systems approach. IEEE Control Systems Letters, 1(1):56–61, July 2017.
- [59] T. A. Henzinger. The theory of hybrid automata. In Verification of digital and hybrid systems, pages 265–292. Springer, 2000.
- [60] J. P. Hespanha. Uniform stability of switched linear systems: Extensions of LaSalle’s invariance principle. IEEE Transactions on Automatic Control, 49(4):470–482, 2004.
- [61] G. W. Hill. On the part of the motion of the lunar perigee which is a function of the mean motions of the sun and moon. Acta Math., 8:1–36, 1886.
- [62] R. Iervolino, S. Trenn, and F. Vasca. Stability of piecewise affine systems through discontinuous piecewise quadratic Lyapunov functions. In 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pages 5894–5899, Dec 2017.
- [63] D. Iglesias-Ponte, M. De León, and D. M. de Diego. Towards a Hamilton–Jacobi theory for nonholonomic mechanical systems. Journal of Physics A: Mathematical and Theoretical, 41(1):015205, 2007.
- [64] I. Iliyev. On the conditions for the existence of the reducing chaplygin factor. Journal of Applied Mathematics and Mechanics, 49(3):295 – 301, 1985.
- [65] D. Jordan and P. Smith. Nonlinear Ordinary Differential Equations: An Introduction for Scientists and Engineers: An Introduction for Scientists and Engineers. Oxford Texts in Applied and Engineering Mathematics. OUP Oxford, 2007.
- [66] B. Jovanovic. Non-holonomic geodesic flows on lie groups and the integrable Suslov problem on SO(4). Journal of Physics A: Mathematical and General, 31(5):1415–1422, feb 1998.
- [67] A. Julius, J. Yin, and J. T. Wen. Time-optimal control for circadian entrainment for a model with circadian and sleep dynamics. In 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pages 4709–4714, Dec 2017.

- [68] V. Jurdjevic. Geometric Control Theory. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1996.
- [69] V. Kaloshin and A. Sorrentino. On the integrability of Birkhoff billiards. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 376(2131):20170419, 2018.
- [70] A. Katok and B. Hasselblatt. Introduction to the Modern Theory of Dynamical Systems. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1995.
- [71] D. E. Kirk. Optimal control theory: an introduction. Springer, 1970.
- [72] J. Koiller. Reduction of some classical non-holonomic systems with symmetry. Archive for Rational Mechanics and Analysis, 118(2):113–148, Jun 1992.
- [73] W. S. Koon and J. E. Marsden. The Hamiltonian and Lagrangian approaches to the dynamics of nonholonomic systems. Reports on Mathematical Physics, 40(1):21 – 62, 1997.
- [74] A. Lamperski and A. D. Ames. Lyapunov-like conditions for the existence of Zeno behavior in hybrid and Lagrangian hybrid systems. In 2007 46th IEEE Conference on Decision and Control, pages 115–120, Dec 2007.
- [75] A.N. Livšic. Cohomology of dynamical systems. Mathematics of the USSR-Izvestiya, 6(6):1278, 1972.
- [76] A. Matveev and A. Savkin. Qualitative Theory of Hybrid Dynamical Systems. Birkhauser Boston, Inc., 01 2000.
- [77] D. McDuff and D. Salamon. Introduction to Symplectic Topology. Oxford Mathematical Monographs. Oxford University Press, 2017.
- [78] P. D. Miller. Applied Asymptotic Analysis. Graduate Studies in Mathematics. American Mathematical Society, 2006.
- [79] M. Molina-Becerra, E. Freire, and J. Vioque. Equations of motion of nonholonomic Hamiltonian systems. Preprint obtained from http://www.matematicaaplicada2.es/data/pdf/1276179170_1811485430.pdf, 2010.
- [80] J. C. Monforte. Geometric, control and numerical aspects of nonholonomic systems. Springer, 2004.
- [81] B. Morris and J. W. Grizzle. Hybrid invariant manifolds in systems with impulse effects with application to periodic locomotion in bipedal robots. IEEE Transactions on Automatic Control, 54(8):1751–1764, Aug 2009.
- [82] B. Morris and J.W. Grizzle. A restricted Poincaré map for determining exponentially stable periodic orbits in systems with impulse effects: Application to bipedal robots. In Proceedings of the 44th IEEE Conference on Decision and Control, pages 4199 – 4206, 01 2006.
- [83] J.I. Neimark and N.A. Fufaev. Dynamics of Nonholonomic Systems. Translations of mathematical monographs. American Mathematical Society, 1972.
- [84] N. A. Neogi. Dynamic partitioning of large discrete event biological systems for hybrid simulation and analysis. In International Workshop on Hybrid Systems: Computation and Control, pages 463–476. Springer, 2004.
- [85] E. Noether. Invariante Variationsprobleme. Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse, 1918:235–257, 1918.
- [86] T. Ohsawa. Nonholonomic and Discrete Hamilton-Jacobi Theory. PhD thesis, University of Michigan, 2010.

- [87] A. Pakniyat. Optimal control of deterministic and stochastic hybrid systems: Theory and applications. PhD thesis, Ph. D. thesis, Department of Electrical & Computer Engineering, McGill , 2016.
- [88] A. Pakniyat and P. E. Caines. On the minimum principle and dynamic programming for hybrid systems. IFAC Proceedings Volumes, 47(3):9629–9634, 2014.
- [89] A. Pakniyat and P. E. Caines. On the minimum principle and dynamic programming for hybrid systems with low dimensional switching manifolds. In 2015 54th IEEE Conference on Decision and Control (CDC), pages 2567–2573. IEEE, 2015.
- [90] A. Pakniyat and P. E. Caines. On the relation between the minimum principle and dynamic programming for classical and hybrid control systems. IEEE Transactions on Automatic Control, 62(9):4347–4362, Sep. 2017.
- [91] D. N. Pekarek and J. E. Marsden. Variational collision integrators and optimal control. In Proceedings of the 18th International Symposium on Mathematical Theory of Networks and Systems (MTNS), 2008.
- [92] L. Perko. Differential Equations and Dynamical Systems. Texts in Applied Mathematics. Springer-Verlag New York, 1991.
- [93] A. Pollack and V. Guillemin. Differential Topology. AMS Chelsea publishing. AMS Chelsea Pub./American Mathematical Society, 1974.
- [94] A. S. Poznyak. Chapter 22 - variational calculus and optimal control. In A. S. Poznyak, editor, Advanced Mathematical Tools for Automatic Control Engineers: Deterministic Techniques, pages 647 – 711. Elsevier, Oxford, 2008.
- [95] H. Razavi, R. Gupta, F. Adams, and A. Bloch. Stability of a class of coupled Hill’s equations and the Lorentz oscillator model. SIAM Journal on Applied Dynamical Systems, 15(2):1104–1123, 2016.
- [96] S. H. Razavi. Symmetry Method for Limit Cycle Walking of Legged Robots. PhD thesis, University of Michigan, 2016.
- [97] S. Redner. A billiard-theoretic approach to elementary one-dimensional elastic collisions. American Journal of Physics, 72(12):1492–1498, 2004.
- [98] W. Rudin. Principles of Mathematical Analysis. International series in pure and applied mathematics. McGraw-Hill, 1976.
- [99] A. Ruina. Nonholonomic stability aspects of piecewise holonomic systems. Reports on Mathematical Physics, 42(1):91 – 100, 1998. Proceedings of the Pacific Institute of Mathematical Sciences Workshop on Nonholonomic Constraints in Dynamics.
- [100] C. O. Saglam, A. R. Teel, and K. Byl. Lyapunov-based versus Poincaré map analysis of the rimless wheel. In 53rd IEEE Conference on Decision and Control, pages 1514–1520. IEEE, 2014.
- [101] A.J. Van Der Schaft and B.M. Maschke. On the Hamiltonian formulation of nonholonomic mechanical systems. Reports on Mathematical Physics, 34(2):225 – 233, 1994.
- [102] S. N. Simic, S. Sastry, K. H. Johansson, and J. Lygeros. Hybrid limit cycles and hybrid Poincaré-Bendixson. IFAC Proceedings Volumes, 35(1):197 – 202, 2002. 15th IFAC World Congress.
- [103] S. E. Spagnolie, C. Wahl, J. Lukasik, and J. Thiffeault. Microorganism billiards. Physica D: Nonlinear Phenomena, 341:33 – 44, 2017.

- [104] D. E. Stewart. Existence of solutions to rigid body dynamics and the Painlevé paradoxes. Comptes Rendus de l'Académie des Sciences - Series I - Mathematics, 325(6):689 – 693, 1997.
- [105] S. Tarbouriech and A. Girard. LMI-based design of dynamic event-triggering mechanism for linear systems. In 2018 IEEE Conference on Decision and Control (CDC), pages 121–126, Dec 2018.
- [106] L. Tavernini. Differential automata and their discrete simulators. Nonlinear Analysis: Theory, Methods & Applications, 11(6):665 – 683, 1987.
- [107] L. W. Tu. An Introduction to Manifolds. Universitext. Springer-Verlag New York, 2011.
- [108] L. W. Tu. Differential Geometry. Graduate Texts in Mathematics. Springer International Publishing, 2017.
- [109] A. J. Van Der Schaft and J. M. Schumacher. An introduction to hybrid dynamical systems, volume 251. Springer London, 2000.
- [110] T. Westenbroek, X. Xiong, A.D. Ames, and S.S. Sastry. Optimal control of piecewise-smooth control systems via singular perturbations. In 2019 IEEE Conference on Decision and Control (CDC), pages 3046–3053, 2019.
- [111] H. Witsenhausen. A class of hybrid-state continuous-time dynamic systems. IEEE Transactions on Automatic Control, 11(2):161–167, April 1966.
- [112] Z. Wu, D. Zhao, and S. Revzen. Coulomb Friction Crawling Model Yields Linear Force-Velocity Profile. Journal of Applied Mechanics, 86(5), 03 2019. 054501.
- [113] D. V. Zenkov and A. M. Bloch. Invariant measures of nonholonomic flows with internal degrees of freedom. Nonlinearity, 16(5):1793–1807, 2003.