

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Symplectic invariants and moduli spaces of integrable systems

Permalink

<https://escholarship.org/uc/item/8fm2b234>

Author

Palmer, Joseph

Publication Date

2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Symplectic invariants and moduli spaces of integrable systems

A Dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Mathematics

by

Joseph Palmer

Committee in charge:

Professor Álvaro Pelayo, Chair
Professor Jorge Cortes
Professor Bruce Driver
Professor Daniel Kane
Professor Lei Ni

2016

Copyright
Joseph Palmer, 2016
All rights reserved.

The Dissertation of Joseph Palmer is approved, and it is acceptable in quality and form for publication on micro-film and electronically:

Chair

University of California, San Diego

2016

DEDICATION

To Cassie. Obviously.

EPIGRAPH

There are two ways to do great mathematics.

The first way is to be smarter than everybody else.

The second way is to be stupider than everybody else – but persistent.

—Raoul Bott

TABLE OF CONTENTS

| | | |
|--|---|------|
| Signature Page | | iii |
| Dedication | | iv |
| Epigraph | | v |
| Table of Contents | | vi |
| List of Figures | | viii |
| Acknowledgements | | ix |
| Vita | | x |
| Abstract of the Dissertation | | xi |
| Chapter 1 | Introduction | 1 |
| | 1.1 An informal example | 2 |
| | 1.2 The aim of this work | 4 |
| | 1.3 Outline | 5 |
| Chapter 2 | Preliminaries | 8 |
| | 2.1 Symplectic geometry | 8 |
| | 2.2 Integrable systems | 15 |
| Chapter 3 | The space of semitoric systems | 39 |
| | 3.1 Introduction | 39 |
| | 3.2 Construction of metric and statement of main theorem | 40 |
| | 3.3 The metric | 46 |
| | 3.4 The completion | 62 |
| | 3.5 Further questions | 76 |
| Chapter 4 | Classifying toric and semitoric fans | 78 |
| | 4.1 Introduction | 78 |
| | 4.2 Fans, symplectic geometry, and winding numbers | 80 |
| | 4.3 Algebraic set-up: matrices and $SL_2(\mathbb{Z})$ relations | 86 |
| | 4.4 Toric fans | 95 |
| | 4.5 Semitoric fans | 103 |
| Chapter 5 | Connected components of space of semitoric systems | 109 |
| | 5.1 Introduction | 109 |
| | 5.2 Preparations | 110 |
| | 5.3 The connected components of \mathcal{M}_{ST} | 110 |
| Chapter 6 | Semitoric helixes and minimal models | 118 |
| | 6.1 Introduction | 118 |
| | 6.2 An application to semitoric systems | 127 |
| | 6.3 From semitoric systems to semitoric helixes | 128 |
| | 6.4 Semitoric helixes and $SL_2(\mathbb{Z})$ | 131 |
| | 6.5 Standard form in $PSL_2(\mathbb{Z})$ and the winding number | 138 |
| | 6.6 Minimal models for semitoric helixes | 145 |

| | | |
|--------------|--|-----|
| Chapter 7 | Symplectic G -capacities | 151 |
| | 7.1 Introduction | 151 |
| | 7.2 Symplectic G -capacities | 153 |
| | 7.3 Hamiltonian $(\mathbb{T}^k \times \mathbb{R}^{n-k})$ -actions | 159 |
| | 7.4 Symplectic \mathbb{T}^n -capacities | 161 |
| | 7.5 Symplectic $(S^1 \times \mathbb{R})$ -capacities | 164 |
| Chapter 8 | Continuity of G -capacities | 170 |
| | 8.1 Introduction | 170 |
| | 8.2 Continuity of symplectic \mathbb{T}^n -capacities | 171 |
| | 8.3 Continuity of symplectic $(S^1 \times \mathbb{R})$ -capacities | 178 |
| Chapter 9 | Moduli spaces of maps | 189 |
| | 9.1 Introduction | 189 |
| | 9.2 The distance function | 190 |
| | 9.3 Main results | 194 |
| | 9.4 Definitions and preliminaries | 198 |
| | 9.5 Completeness of \mathcal{M}^\sim | 208 |
| | 9.6 Almost everywhere convergence and \mathcal{D} | 221 |
| | 9.7 Families with singular limits | 225 |
| | 9.8 Further questions | 229 |
| Index | | 233 |
| Bibliography | | 236 |

LIST OF FIGURES

Figure 2.1: The complete invariant for semitoric systems is a collection of these objects. 32
 Figure 2.2: Complete invariant of a semitoric system. 37

Figure 3.1: Image of Δ in proof of Lemma 3.3.8. 53
 Figure 3.2: The lines ℓ_1, ℓ_2, ℓ_3 , and triangle G 53
 Figure 3.3: The measure of the convex hull of a line segment and a point. 54
 Figure 3.4: Bounding the Radon-Nikodym derivative. 55
 Figure 3.5: A new polygon with the same ν -measure and top boundary $\{y = 0\}$ 61
 Figure 3.6: Finding a positive measure region which is always in the symmetric difference. 61
 Figure 3.7: Approximating an element of $\mathcal{P}'''_{m_f, [\vec{k}]}$ by one in $\mathcal{P}''_{m_f, [\vec{k}]}$ 68
 Figure 3.8: An arbitrary convex set can be approximated from the inside by a polygon. 70
 Figure 3.9: The action of $t_{\lambda_j^{k_m}}^1$ and $t_{\lambda_j^{k_n}}^1$ on a polygon. 73
 Figure 3.10: A continuous family in $\tilde{\mathbb{I}}$ in which one critical point passes over the other. 76

Figure 4.1: The three toric minimal models and a corner chopping of the Hirzebruch trapezoid. 81
 Figure 4.2: These vectors do not form a fan because they are not in counter-clockwise order. 89
 Figure 4.3: Any semitoric fan can be transformed into the standard fan. 105

Figure 5.1: Each fan transformation corresponds to a continuous path of semitoric polygons. 113
 Figure 5.2: Path from a compact primitive semitoric polygon to a noncompact one. 114
 Figure 5.3: Adding vertices to a noncompact primitive semitoric polygon. 116

Figure 6.1: The three possible minimal toric fans listed in Theorem 6.1.1. 120
 Figure 6.2: A minimal semitoric helix of length 6 and complexity 2. 121
 Figure 6.3: Minimal semitoric helixes of type (1)-(6) of Theorem 6.1.13. 127
 Figure 6.4: The semitoric helix is produced by unwinding the polygon on the fake corner p . 129
 Figure 6.5: The homotopy from the proof of Lemma 6.4.3. 133

Figure 7.1: A symplectic \mathbb{R} -embedding. 158
 Figure 7.2: Toric ball packing of S^2 by symplectic \mathbb{T}^2 -disks. 162

Figure 8.1: An ε -corner chop at a vertex x_0 of Δ for some $\varepsilon > 0$ 172
 Figure 8.2: (a) $\Delta(1) \subset \mathbb{R}^2$. (b) An admissible packing. 173
 Figure 8.3: Families of Delzant polygons on which \mathcal{T} is (a) continuous and (b) not continuous. 178
 Figure 8.4: An admissible semitoric packing. Here t denotes $t_{\vec{\lambda}}^{\vec{u}}$ 179
 Figure 8.5: Performing a corner chop on a hidden corner. 184
 Figure 8.6: Corner chop of a corner not used in the optimal packing. 187
 Figure 8.7: Families of primitive semitoric polygons on which $\mathcal{S}\mathcal{T}$ is (a) continuous and (b) not. 188

Figure 9.1: I will be considering maps from subsets of M to N 192
 Figure 9.2: A graphic representing the values of $p_{\phi_\psi}^d$ on $\mathcal{S} \subset M$ 193
 Figure 9.3: An image of $\Phi_{m,k}$ 206
 Figure 9.4: A few terms of $\{\phi_n\}$. It can be seen this does not converge pointwise. 207
 Figure 9.5: Applying the proof of Lemma 9.5.5 to Example 9.4.13. 213
 Figure 9.6: Examples of restricting the maps to find the limit ϕ_0^2 in the proof of Lemma 9.5.5. 213
 Figure 9.7: $\{\check{\phi}_t\}_{t \in (0,1)}$ is a family of maps into \mathbb{R}^k 220
 Figure 9.8: A figure of the relevant maps when defining a smooth family of embeddings. 225
 Figure 9.9: The strategy is to connect the embedding ϕ_{t_0} with the map ϕ 231

ACKNOWLEDGEMENTS

I am thankful for my advisor, Álvaro Pelayo, to whom I owe much of my success. He has again and again provided me with invaluable guidance and advice, and he has inspired me by showing me the true meaning of dedication.

I have been fortunate to work with many talented mathematicians throughout my career, both as my peers and as my teachers, from which I have learned many things. There far too many people in this category to name them all, but the highlights include, David Garth, Kevin Easley, the graduate students at Wash. U (such as Christopher Cox, Ben Passer, Brady Rocks, and Ryan Keast), Xiang Tang, the graduate students at UCSD, and Melvin Leok. The list goes on.

I have never been so indebted to someone as I am to Cassie Kling at this moment. She has been my inspiration all these years and has, time and time again, given me support and encouragement long after any rational person would have given up. I am grateful for my parents, Jane and Paul Palmer, who have done everything in their power to support my interest in math and science without a second thought as long as I can remember.

Chapter 3, in part, is comprised of material submitted for publication by the author of this dissertation as *Moduli spaces of semitoric systems*, currently available at arXiv:1502.07296 [63].

Chapters 4 and 5, in part, are comprised of material submitted for publication by the author of this dissertation, Daniel M. Kane, and Álvaro Pelayo as *Classifying toric and semitoric fans by lifting equations from $SL_2(\mathbb{Z})$* , currently available as arXiv:1502.07698 [47].

Chapter 6, in part, is comprised of material in preparation for submission by the author of this dissertation, Daniel M. Kane, and Álvaro Pelayo as *Minimal models in semitoric geometry*.

Chapters 7 and 8, in part, are comprised of material submitted for publication by the author of this dissertation, Alessio Figalli, and Álvaro Pelayo as *Symplectic G -capacities and integrable systems*, currently available as arXiv:1511.04499 [30]. Alessio Figalli is supported by NSF grants DMS-1262411 and DMS-1361122.

Chapter 9, in full, is comprised of material submitted for publication by the author of this dissertation as *Metrics and convergence in moduli spaces of maps*, currently available at arXiv:1406.4181 [62].

The work in this dissertation is supported by NSF grants DMS-1055897 and DMS-1518420.

VITA

| | |
|-----------|---|
| 2011 | Bachelor of Science in Mathematics, Bachelor of Science in Physics <i>magna cum laude</i> , Truman State University |
| 2011-2014 | Graduate Teaching/Research Assistant, Washington University in St. Louis |
| 2013 | Master of Science in Mathematics, Washington University in St. Louis |
| 2014-2016 | Graduate Teaching/Research Assistant, University of California, San Diego |
| 2016 | Doctor of Philosophy in Mathematics, University of California, San Diego |

PAPERS

Minimal models in semitoric geometry (with D. M. Kane and Á. Pelayo), *in preparation*.

Symplectic G -capacities and integrable systems (with A. Figalli and Á. Pelayo), *preprint*, arXiv:1511.04499.

Classifying toric and semitoric fans by lifting equations from $SL_2(\mathbb{Z})$ (with D. M. Kane and Á. Pelayo), *preprint*, arXiv:1502.07698.

Moduli spaces of semitoric systems, *preprint*, arXiv:1502.07296.

Metrics and convergence in the moduli space of maps, *preprint*, arXiv:1406.4181.

ABSTRACT OF THE DISSERTATION

Symplectic invariants and moduli spaces of integrable systems

by

Joseph Palmer

Doctor of Philosophy in Mathematics

University of California, San Diego, 2016

Professor Álvaro Pelayo, Chair

In this dissertation I prove a number of results about the symplectic geometry of finite dimensional integrable Hamiltonian systems, especially those of semitoric type. Integrable systems are, roughly, dynamical systems with the maximal amount of conserved quantities. Though the study of integrable systems goes back hundreds of years, the earliest general result in this field is the action-angle theorem of Arnold in 1963, which was later extended to a global version by Duistermaat. The results of Atiyah, Guillemin-Sternberg, and Delzant in the 1980s classified toric integrable systems, which are those produced by effective Hamiltonian \mathbb{T}^n -actions. Recently, Pelayo-Vũ Ngọc classified semitoric integrable systems, which generalize toric systems in dimension four, in terms of five symplectic invariants. Using this classification, I construct a metric on the space of semitoric integrable systems. To study continuous paths in this space produced via symplectic semitoric

blowups, I introduce an algebraic technique to study such systems by lifting matrix equations from the special linear group $SL_2(\mathbb{Z})$ to its preimage in the universal cover of $SL_2(\mathbb{R})$. With this method I determine the connected components of the space of semitoric integrable systems. Motivated by this algebraic technique, I introduce the notion of a semitoric helix; the natural combinatorial invariant of semitoric systems. By applying a refined version of the algebraic method to semitoric helices I classify all possible minimal semitoric integrable systems, which are those that do not admit a symplectic semitoric blowdown. I also produce invariants of integrable systems designed to respect the natural symmetries of such systems, especially toric and semitoric ones. For any Lie group G , I construct a G -equivariant analogue of the Ekeland-Hofer symplectic capacities. I give examples when $G = \mathbb{T}^k \times \mathbb{R}^{d-k}$, in which case the capacity is an invariant of integrable systems, and I study the continuity of these capacities using the metric I defined on semitoric systems. Finally, as a first step towards constructing a meaningful metric on general integrable systems, I provide a framework to study convergence properties of families of maps between manifolds which have distinct domains by defining a metric similar to the L^1 distance on such a collection.

Chapter 1

Introduction

A symplectic manifold is a pair (M, ω) where M is a smooth manifold and ω is a closed, non-degenerate 2-form on M . The study of symplectic manifolds was originally motivated by classical mechanics, where symplectic manifolds are used to model the phase space of certain physical dynamical systems. Such systems are called *integrable* if they possess the maximal number of independent quantities which are preserved by the dynamics.

Despite being a very classical subject going back to the 17th century or earlier, very few general results are known about integrable systems, although there is extensive literature about concrete examples. In the 20th century there have been a number seminal works which apply to general integrable systems, some of the most famous ones being the action-angle theorem by Arnold [2] and its global version by Duistermaat [22].

In the 1980s integrable systems of toric type were classified by Delzant [21], using as a stepping stone the convexity theorem of Atiyah [3] and Guillemin-Sternberg [40]. Toric integrable systems are those which have specific periodic symmetries in addition to having the conserved quantities of an integrable system, and thus can be viewed as an effective Hamiltonian \mathbb{T}^n -action. Delzant classified them in terms of the image of the associated momentum map, the joint map of the conserved quantities, and showed that in this case the image is what is now known as a Delzant polytope.

This subject has been recently revitalized by the Pelayo-Vũ Ngọc classification of semitoric

integrable systems [70], which generalize toric integrable systems in dimension 4. This classification is the culmination of several years of work constructing symplectic invariants of semitoric integrable systems [69, 76, 77]. Semitoric integrable systems are classified in terms of five invariants, one of which is analogous to the Delzant polytope while the other four are related to the focus-focus (or nodal) singularities which can arise in semitoric systems but not in toric ones. The classification of semitoric systems is much more delicate than that of toric systems because the presence of focus-focus singularities create Duistermaat monodromy [22] in the natural structure of the system.

We begin with a general outline of the relationship between symplectic geometry, dynamical systems, and integrable systems.

1.1 An informal example

If Q is a manifold which represents all possible configurations of a physical system, then each point of T^*Q , its cotangent bundle, is a state of the corresponding system, containing both location and momenta information. In this case, T^*Q is known as the phase space of the system. The phase space is also equipped with a function $H: T^*Q \rightarrow \mathbb{R}$, known as the Hamiltonian or energy function, which assigns to each state of the system the total energy of that state. For this illustration we assume that H is time-independent. Solving such a system amounts to extracting from H the dynamics of the system. That is, the goal is to use a function on the phase space to determine an \mathbb{R} -action, or flow, on the phase space which represents the way that the system transforms over time. This is where symplectic geometry becomes useful. We will see that every cotangent bundle is equipped with a natural symplectic form, and thus is a symplectic manifold.

If (M, ω) is a symplectic manifold and $f: M \rightarrow \mathbb{R}$ is any smooth function then there exists a unique vector field \mathcal{X}_f on M which satisfies

$$\omega(\mathcal{X}_f, Y) = -df(Y)$$

for all vector fields Y on M . The vector field \mathcal{X}_f is known as the *Hamiltonian vector field* of f . Thus, symplectic geometry provides a method of producing a vector field from a function. If M is a phase space and $f = H$ is the associated time-independent Hamiltonian then the flow of \mathcal{X}_H , when it is

complete, describes the dynamics of the system. In this case (M, ω, H) is known as a Hamiltonian system.

The problem of understanding the flow of \mathcal{X}_H can be a difficult one. In the situation we have described, H will be a conserved quantity of the system, that is, H is invariant under the flow of \mathcal{X}_H . This already gives us a hint about computing the dynamics, the flow lines must always lie on level sets of the function H . If more independent conserved quantities can be found, then the flow will be further restricted.

Roughly, a completely integrable Hamiltonian system, or simply an integrable system, is a Hamiltonian system on a symplectic manifold equipped with the maximum number of independent conserved quantities. Symplectic manifolds are even dimensional, and if M is $2n$ -dimensional then an integrable system on M will have a total of n conserved quantities. If an integrable system has conserved quantities

$$f_1 = H, f_2, \dots, f_n: M \rightarrow \mathbb{R}$$

then the joint map

$$F = (f_1, \dots, f_n): M \rightarrow \mathbb{R}^n$$

is known as the *momentum map*.

Suppose that F is proper, so each Hamiltonian vector field \mathcal{X}_{f_i} of the component f_i of F is complete. An important property of integrable systems is that the flows of $\mathcal{X}_{f_1}, \dots, \mathcal{X}_{f_n}$ commute, and thus produce an \mathbb{R}^n -action on M . The orbits of this \mathbb{R}^n -action are contained in level sets of F , since each component of F is preserved by this \mathbb{R}^n -action. Thus, in the study of integrable systems there is interplay between dynamics - the orbits - and a singular fibration - the level sets of F .

Now, we can give a better indication of the definitions of toric and semitoric integrable systems. An integrable system is toric if the flow of the Hamiltonian vector field associated to each of the components of its momentum map is periodic of period 2π . These systems are considered to be very interesting, but nonetheless this definition is very rigid. Delzant [21] showed that two toric integrable systems are diffeomorphic by a diffeomorphism which preserves the symplectic form and the momentum map if and only if the images of their momentum maps agree. That is, Delzant showed that toric integrable systems are classified by the image of their momentum map, which is

necessarily a specific type of polytope known as a Delzant polytope.

A semitoric integrable system is an integrable system in dimension 4 for which the singular points of the momentum map are of a specific type and one of the two components of the momentum map is proper with 2π -periodic flow. The requirements of this definition are much less rigid than those of a toric integrable system. Pelayo-Vũ Ngọc [69, 70] classified these systems in terms of a list of five invariants.

1.2 The aim of this work

The general goal of this dissertation is to construct invariants of integrable systems and to study continuous deformations of integrable systems. Special emphasis is placed on toric and semitoric systems. Questions addressed include the following:

1. Does there exist a meaningful topology or metric on the space of semitoric integrable systems?
2. Must the limit of semitoric systems be semitoric?
3. What steps can be taken towards a topology or metric on more general integrable systems?
4. What are the path-connected components of the space of semitoric systems?
5. Is there a reasonable way to define blowup/blowdown operations on semitoric systems? Can the systems that do not admit a blowdown be classified as they are in the toric case?
6. Is there a purely combinatorial invariant of semitoric systems, analogous to the fans of toric systems (and of toric varieties)?
7. What invariants of integrable systems can be produced which respect the natural symmetries of such systems?

With these goals in mind, in this dissertation I define a metric, and thus topology, on the moduli space of semitoric integrable systems which extends a known metric on toric integrable systems and I characterize the connected components of the resulting topological space. I introduce semitoric helixes, which are the natural combinatorial invariant of semitoric systems, and develop an equivariant version of symplectic capacities in order to construct new invariants of integrable systems.

Also, as a step towards a metric on more general spaces of integrable systems, I define and study a metric on spaces of maps which may not have the same domain.

1.3 Outline

- In **Chapter 2** I give a quick exposition of the background necessary for the rest of the dissertation. This chapter is broken into two parts. First is an introduction to symplectic geometry, which includes the basics of symplectic geometry along with discussions of symplectic invariants (such as symplectic capacities) and symplectic/Hamiltonian group actions. Secondly, there is an introduction to the integrable systems from a symplectic point of view. This includes the classifications of toric and semitoric integrable systems and a discussion of non-degenerate singular points of integrable systems.
- In **Chapter 3** I extend the metric on toric systems from [67] to semitoric systems. The metric on toric integrable systems is defined by declaring the distance between two systems to be the usual Lebesgue measure of the symmetric difference of their Delzant polytopes. I extend this distance to semitoric systems by considering each of the five semitoric invariants from the Pelayo-Vũ Ngọc classification [70]. The metric on semitoric systems is constructed in Section 3.2, and it depends on two parameters, a sequence and a measure on \mathbb{R}^2 . The main result of Chapter 3 is Theorem 3.2.12, which states that:
 1. the proposed function is indeed a metric;
 2. the topology of the metric space does not depend on parameters of the metric;
 3. the resulting metric space is incomplete,
 and describes the completion.
- In **Chapter 4** I develop an algebraic method to study an object I introduce in this chapter known as a semitoric fan. I study these objects in order to use them in Chapter 5 to determine the connected components of the space of semitoric integrable systems. Roughly, a *semitoric fan* is a collection of vectors in \mathbb{R}^2 which can be obtained as the inwards pointing normal vectors of an element of a semitoric polygon. Each subsequent pair of vectors in the fan determines

an element of $\mathrm{SL}_2(\mathbb{Z})$, and composing these matrices must equal a specific matrix which is determined by the number of focus-focus singular points in the corresponding semitoric system. This equation in $\mathrm{SL}_2(\mathbb{Z})$ does not contain enough information to classify such fans because it does not count the number of times the vectors in the fan wind around the origin. Thus, I study semitoric fans by lifting these equations from $\mathrm{SL}_2(\mathbb{Z})$ to its preimage in the universal cover of $\mathrm{SL}_2(\mathbb{R})$. I am then able to prove Theorem 4.1.1, which describes how semitoric fans can be related by a sequence specific transformations, and which is useful in Chapter 5.

- The goal of **Chapter 5** is to determine the connected components of the moduli space of semitoric systems from Chapter 3. From the definition of the metric, it is clear that any two semitoric systems with different numbers of focus-focus singularities and systems with different twisting index invariants must be in different components of the moduli space. In Chapter 5, using the algebraic tools developed in Chapter 4, I prove that this is the only obstruction to two systems being in the same component. That is, any two semitoric systems with the same number of focus-focus singular points and the same twisting index invariant may be continuously deformed into one another via a continuous path of semitoric systems with the same number of focus-focus singular points and the same twisting index invariant. This is the content of Theorem 5.1.1, which is proven using Theorem 4.1.1.
- In **Chapter 6** I introduce a combinatorial invariant of semitoric systems which I call the semitoric helix. Semitoric helices are a refinement of semitoric fans from Chapter 4 which generalize the toric fans of toric integrable systems. In this chapter I discuss the method of producing these invariants using the semitoric polygon invariant introduced by Vũ Ngọc [77]. I expand on the techniques from Chapter 4 to study minimal models of semitoric helices, which in turn gives a classification of minimal semitoric systems. A semitoric system is minimal if it is impossible to perform a symplectic blowdown on the manifold, and this corresponds to the case that each vector in the associated semitoric helix is not the sum of the adjacent two vectors. In this chapter, I give a complete classification of minimal helices in Theorem 6.1.13 and explain the relation between semitoric systems, semitoric fans, and semitoric helices.
- In **Chapter 7** I construct invariants of symplectic G -manifolds and define an equivariant

version of symplectic capacities. In particular, I study G -capacities when $G = \mathbb{T}^k \times \mathbb{R}^{d-k}$. Since integrable systems with complete Hamiltonian flows come equipped with a natural \mathbb{R}^n -action these can be used as invariants of integrable systems in which k of the components of F have periodic flows. I construct an equivariant version of the Gromov radius, which is proven to be a \mathbb{R}^n -capacity in Theorem 7.1.1 and I construct equivariant ball packing capacities for toric and semitoric systems, which are proven to be \mathbb{T}^2 and $S^1 \times \mathbb{R}$ -capacities, respectively, in Theorem 7.1.2. Roughly, these capacities measure the symplectic volume of the given toric or semitoric manifold which can be filled with equivariantly symplectically embedded balls.

- In **Chapter 8** I study the continuity of certain G -capacities. The topology on the space of toric integrable systems from [67] and the topology on the space of semitoric integrable systems in Chapter 3 create the possibility to study the continuity of symplectic invariants on these spaces. Theorem 8.1.1 specifies the regions of continuity of the packing capacities and the equivariant Gromov radius on toric and semitoric systems.
- In **Chapter 9** I study families of maps which have different domains. For manifolds M and N where M is equipped with a volume form I consider families of maps in the collection $\{(\phi, B_\phi) \mid B_\phi \subset M, \phi: B_\phi \rightarrow N \text{ with } B_\phi, \phi \text{ both measurable}\}$ and I define a distance function similar to a generalized L^1 distance on such a collection. Theorem 9.3.1 states that the proposed function is indeed a metric and that the resulting metric space is complete. The relationship between convergence in this metric and the natural notion of pointwise almost everywhere convergence in this context is given in Theorem 9.3.5. In Theorem 9.3.8 I use the new metric to study families of maps which converge after an arbitrarily small perturbation. This metric could in principle have many applications, but it was developed with integrable systems in mind. If (M, ω, F) and (M', ω', F') are integrable systems such that M and M' both symplectically embed into a symplectic manifold (N, η) then the distance from Chapter 9 can be used to compare M and M' .

Chapter 2

Preliminaries

In this section I provide a quick overview of the symplectic theory of integrable systems. Section 2.1 covers general symplectic geometry and Section 2.2 is concerned with integrable systems from a symplectic point of view.

There are many standard resources already available to supplement this chapter. For instance, the books [10, 55] together form a great introduction to symplectic geometry which includes much of the same information in Section 2.1. Additionally, the book [8] and survey paper [71] are good resources from which to study the symplectic theory of integrable systems to complement Section 2.2.

2.1 Symplectic geometry

A *symplectic manifold* is a pair (M, ω) where M is a smooth manifold and ω is a closed, non-degenerate 2-form on M . That is, $d\omega = 0$ and for each $p \in M$ if $v \in T_p M$ is such that $\omega_p(v, w) = 0$ for all $w \in T_p M$ then $v = 0$. In this case ω is known as the *symplectic form*. Given two symplectic manifolds (M, ω) and (M', ω') a diffeomorphism $\phi: M \rightarrow M'$ is known as a *symplectomorphism* if $\phi^*\omega' = \omega$, and in this case (M, ω) and (M', ω') are said to be *symplectomorphic*.

2.1.1 First consequences

The existence of ω puts some requirements on M . Let (M, ω) be a symplectic manifold and let $p \in M$. By using a skew-symmetric form of the Gram-Schmidt process, it can be seen that on the

tangent plane $T_p M$ there exists an integer $n > 0$ and linear coordinates $x_1, \dots, x_n, y_1, \dots, y_n$ such that

$$\omega_p(x_i, x_j) = \omega_p(y_i, y_j) = 0 \text{ and } \omega_p(x_i, y_j) = \delta_{ij}$$

for all $i, j = 1, \dots, n$. This implies that M must be even-dimensional. The degeneracy condition on ω implies that $\omega^n = \omega \wedge \dots \wedge \omega$ is a volume form, so M must be orientable.

If M is compact and $2n$ -dimensional, then its even dimensional de Rham cohomology groups $H^{2k}(M)$, for $k = 1, \dots, n$, must be non-trivial. This is because if $\omega^k = d\eta$ then

$$\int_M \omega^k = \int_M d\eta = \int_{\partial M} \eta = 0$$

by Stoke's theorem and the fact that $\partial M = \emptyset$. This contradicts the non-degeneracy of ω . In particular, the only sphere with a symplectic structure is the 2-sphere.

In dimension $2n = 2$ the definition of a symplectic form coincides with that of a volume form. Thus, any orientable surface is a symplectic manifold.

Recall Cartan's magic formula, which states that

$$\mathcal{L}_X \eta = d\iota_X \eta + \iota_X d\eta \tag{2.1}$$

where \mathcal{L} is the usual Lie derivative, X is any vector field on M , $\eta \in \Omega^k(M)$, and ι is the interior product defined by

$$\iota_X \eta(X_1, \dots, X_{k-1}) = \eta(X, X_1, \dots, X_{k-1})$$

for vector fields X_1, \dots, X_{k-1} on M . In the case that ω is a symplectic form then Equation (2.1) implies that

$$\mathcal{L}_X \omega = d\iota_X \omega$$

for any vector field X on M because ω is closed.

2.1.2 Local theory

An example of a symplectic manifold is \mathbb{R}^{2n} with linear coordinates $x_1, \dots, x_n, y_1, \dots, y_n$ and symplectic form

$$\omega_0 = \sum_{i=1}^n dx_i \wedge dy_i.$$

Given an open $U \subset M$ we pull back the symplectic form on M to U via the inclusion map in order to give U a symplectic structure.

Theorem 2.1.1 (Darboux). *Let (M, ω) be a symplectic manifold and let $p \in M$. Then there exists an open neighborhood of p which is symplectomorphic to an open subset of $(\mathbb{R}^{2n}, \omega_0)$.*

This theorem gives a local normal form for any point of all symplectic manifolds and thus implies that the only local symplectic invariant is the dimension. Therefore, the study of symplectic invariants by necessity is a study of global phenomena. This should be compared to Riemannian geometry in which there are local curvature invariants.

2.1.3 An example: Cotangent bundles

Let Q be any n -dimensional manifold and T^*Q its cotangent bundle. Let $\pi: T^*Q \rightarrow Q$ be the natural projection so $d\pi: T(T^*Q) \rightarrow TQ$. Any $x \in T^*(Q)$ is a pair $x = (q, \xi)$ where $q = \pi(x)$ and $\xi \in T_q^*Q$. Recall that if (U, q_1, \dots, q_n) is a chart for Q then $(dq_1)_q, \dots, (dq_n)_q$ form a basis of T_q^*Q . Thus, any $\xi \in T_q^*Q$ can be written as $\xi = \sum_{i=1}^n \xi_i (dq_i)_q$. This induces a chart, known as a cotangent chart, $(T^*U, q_1, \dots, q_n, \xi_1, \dots, \xi_n)$ on T^*Q about the point (q, ξ) .

Define a 1-form $\alpha \in \Omega^1(T^*Q)$ pointwise by

$$\alpha_{(q, \xi)} = \pi_{(q, \xi)}^* \xi.$$

In a cotangent chart $(T^*U, q_1, \dots, q_n, \xi_1, \dots, \xi_n)$ it can be seen that

$$\alpha = \sum_{i=1}^n \xi_i dx_i.$$

This form α is known as the tautological one-form on T^*Q .

Now define

$$\omega = -d\alpha$$

so locally

$$\omega = \sum_{i=1}^n dq_i \wedge d\xi_i.$$

We know that ω is intrinsically defined because α is, and it is immediate from its definition that ω is a closed, non-degenerate 2-form. That is, ω is a symplectic form on $T^*(Q)$.

2.1.4 Lagrangian submanifolds

Let (M, ω) be a $2n$ -dimensional symplectic manifold. A submanifold $X \subset M$ is said to be *isotropic* if $i^*\omega \equiv 0$ where $i: X \hookrightarrow M$ is the inclusion map. The dimension of X can be bounded above by symplectic linear algebra.

Proposition 2.1.2. *If X is a isotropic submanifold of (M, ω) then*

$$\dim(X) \leq \frac{1}{2} \dim(M)$$

Proof. Suppose that $X \subset M$ is isotropic and let $p \in M$ be any point. Then $T_p X \subset T_p M$ is such that $\omega_p(v, w) = 0$ for all $v, w \in T_p X$. Let

$$(T_p X)^\omega = \{v \in T_p M \mid \omega_p(v, w) = 0 \text{ for all } w \in T_p X\}$$

and $\dim(T_p X) + \dim((T_p X)^\omega) = \dim(T_p M)$ because the map

$$\begin{aligned} T_p V &\rightarrow T_p^* X \\ v &\mapsto \omega_p(v, \cdot) \end{aligned}$$

has rank $\dim(T_p X)$ and nullity $\dim((T_p X)^\omega)$. Since X is isotropic $T_p X \subset (T_p X)^\omega$, and we conclude that if X is isotropic then

$$2\dim(X) \leq \dim(M). \quad \square$$

Definition 2.1.3. A submanifold X of a symplectic manifold (M, ω) is a *Lagrangian submanifold* if X is isotropic, i.e. $i^*\omega = 0$ where $i: X \hookrightarrow M$ is the inclusion map, and $\dim(X) = \frac{1}{2}\dim(M)$.

Suppose $M = T^*Q$, the cotangent bundle of an n -dimensional manifold Q , equipped with the standard symplectic form on cotangent bundles as in Section 2.1.3. Then Q can be identified with the zero section of T^*Q which is given by

$$Q_0 = \{(q, \xi) \mid \xi = 0 \text{ in } T_q^*Q\}.$$

In a local cotangent chart $(T^*U, q_1, \dots, q_n, \xi_1, \dots, \xi_n)$ this submanifold is described as $\xi_1 = \dots = \xi_n = 0$, so the tautological 1-form $\alpha = \sum_{i=1}^n \xi_i dx_i$ vanishes on Q_0 . That is, if $i: Q_0 \hookrightarrow T^*Q$ is the inclusion map then $i^*\alpha = 0$. Thus $i^*\omega = -i^*d\alpha = 0$ so Q_0 is a Lagrangian submanifold.

Here we have seen that any manifold Q can be realized as a Lagrangian submanifold by identifying it with the zero section of its cotangent bundle. This is the content of "Weinstein's Lagrangian creed" that "everything is a Lagrangian submanifold".

2.1.5 Symplectic group actions

Let (M, ω) be a symplectic manifold and let G be a Lie group with Lie algebra $\text{Lie}(G)$ and dual Lie algebra $\text{Lie}(G)^*$. Then $\phi: G \times M \rightarrow M$ is a *symplectic G -action* if:

1. $\phi(g, \cdot): M \rightarrow M$ is a symplectomorphism for all $g \in G$;
2. $\phi(g, \phi(h, \cdot)) = \phi(gh, \cdot)$ for all $g, h \in G$.

If ϕ is a symplectic G -action on (M, ω) then the triple (M, ω, ϕ) is known as a symplectic G -manifold. Such an action is a *Hamiltonian G -action* if there exists a map $\mu: M \rightarrow \text{Lie}(G)^*$ such that

$$-d\langle \mu, X \rangle = \omega(X_M, \cdot)$$

for all $X \in \text{Lie}(G)$ where X_M denotes the vector field on M generated by X via the action of G . That is, for $X \in \text{Lie}(G)$ the vector field X_M is given by

$$(X_M)_p = \left. \frac{d}{dt} \phi(\exp(tX), p) \right|_{t=0}.$$

Such a quadruple (M, ω, ϕ, μ) is known as a Hamiltonian G -manifold and μ is known as the *momentum map*.

An as example, consider the sphere S^2 viewed as a submanifold of \mathbb{R}^3 with coordinates (θ, h) where $\theta \in [0, 2\pi)$ represents the angle and $h \in [-1, 1]$ represents the height (cylindrical coordinates on \mathbb{R}^3). Then S^2 is a symplectic manifold since it is an orientable surface and the symplectic form is given by $\omega = d\theta \wedge dh$. A symplectic action of \mathbb{R} on S^2 is given by

$$\begin{aligned} \phi: \mathbb{R} \times S^2 &\rightarrow S^2 \\ (t, (\theta, h)) &\mapsto (\theta + t, h). \end{aligned}$$

The vector field induced on S^2 by this action is $\frac{\partial}{\partial \theta}$ and

$$\omega \left(\frac{\partial}{\partial \theta}, \cdot \right) = dh$$

which means that the action ϕ is Hamiltonian with momentum map $-h: S^2 \rightarrow \mathbb{R}$.

Next consider a similar action on the torus. Let $\mathbb{T}^2 \cong S^1 \times S^1$ have coordinates (α_1, α_2) where $\alpha_i \in S^2 \subset \mathbb{C}$ and endow \mathbb{T}^2 with the symplectic form $\omega = d\alpha_1 \wedge d\alpha_2$. Define an action

$$\begin{aligned} \psi: \mathbb{R} \times \mathbb{T}^2 &\rightarrow \mathbb{T}^2 \\ (t, (\alpha_1, \alpha_2)) &\mapsto (\alpha_1 + t, \alpha_2). \end{aligned}$$

The induced vector field is $\frac{\partial}{\partial \alpha_1}$ and

$$\omega \left(\frac{\partial}{\partial \alpha_1}, \cdot \right) = d\alpha_2$$

which closed but not exact. Thus, this action is symplectic but not Hamiltonian.

Let ϕ be a symplectic G -action on (M, ω) and let X_M be the vector field on M generated by $X \in \text{Lie}(G)$. By Cartan's magic formula

$$\mathcal{L}_{X_M} \omega = d\iota_{X_M} \omega + \iota_{X_M} d\omega$$

which implies that

$$d\iota_{X_M}\omega = 0$$

because ω is closed and the G -action is symplectic. The action is Hamiltonian if and only if $\iota_{X_M}\omega$ is exact for all $X \in \text{Lie}(G)$. Thus, we see that if $H^1(M) = 0$ then all symplectic actions are Hamiltonian.

2.1.6 Symplectic capacities

By Theorem 2.1.1 there are no local invariants in symplectic geometry, but we will see that there are many interesting global invariants. An important class of such invariants are known as symplectic capacities. Let

$$B^{2n}(r) = \{(x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}^{2n} \mid x_1^2 + y_1^2 + \dots + x_n^2 + y_n^2 < r^2\}$$

denote the $2n$ -dimensional symplectic ball of radius r and let $Z^{2n}(r) = B^{2n}(r) \times \mathbb{R}^{2n-2}$ denote the $2n$ -dimensional symplectic cylinder. Both inherit a symplectic structure from the natural embedding into \mathbb{R}^{2n} . A *symplectic embedding* is an embedding which preserves the symplectic form. In the 1980s Gromov proved the following influential theorem.

Theorem 2.1.4 (Non-squeezing theorem [35]). *There exists a symplectic embedding $\rho: B^{2n}(r) \hookrightarrow Z^{2n}(R)$ if and only if $r \leq R$.*

This should be contrasted to the case of volume-preserving embeddings. For all $r, R > 0$ there exists a volume preserving embedding $B^{2n}(r) \hookrightarrow Z^{2n}(R)$.

Theorem 2.1.4 motivates the following definition.

Definition 2.1.5. The *Gromov radius* of a $2n$ -dimensional symplectic manifold (M, ω) is given by

$$c_B(M) = \sup\{r > 0 \mid \text{there exists a symplectic embedding } B^{2n}(r) \hookrightarrow M\}$$

If $\phi: (M, \omega) \rightarrow (M', \omega')$ is a symplectomorphism then $\rho: B^{2n}(r) \hookrightarrow M$ is a symplectic embedding if and only if $\phi \circ \rho$ is a symplectic embedding, so the Gromov radius is a symplectic

invariant. The non-squeezing theorem asserts that

$$c_B(\mathbb{Z}^{2n}(r)) = r.$$

Let Symp^{2n} be the category of symplectic $2n$ -dimensional manifolds with symplectic embeddings as morphisms. A *symplectic category* is a subcategory \mathcal{C} of Symp^{2n} such that $(M, \omega) \in \mathcal{C}$ implies $(M, \lambda\omega) \in \mathcal{C}$ for all $\lambda \in \mathbb{R} \setminus \{0\}$. Let $\mathcal{C} \subset \text{Symp}^{2n}$ be a symplectic category.

The following fundamental notion of symplectic invariant is due to Ekeland and Hofer.

Definition 2.1.6 ([26, 44]). A *generalized symplectic capacity* on \mathcal{C} is a map $c: \mathcal{C} \rightarrow [0, \infty]$ satisfying:

1. *Monotonicity*: if $(M, \omega), (M', \omega') \in \mathcal{C}$ and there exists a symplectic embedding $M \hookrightarrow M'$ then $c(M, \omega) \leq c(M', \omega')$;
2. *Conformality*: if $\lambda \in \mathbb{R} \setminus \{0\}$ and $(M, \omega) \in \mathcal{C}$ then $c(M, \lambda\omega) = |\lambda|c(M, \omega)$.

If additionally $B^{2n}, \mathbb{Z}^{2n} \in \mathcal{C}$ and c satisfies:

3. *Non-triviality*: $0 < c(\mathbb{Z}^{2n}, \omega_0) < \infty$ and $0 < c(B^{2n}, \omega_0) < \infty$;

then c is a *symplectic capacity*.

It is clear that the Gromov radius satisfies items (1) and (2), and by the non-squeezing theorem we can see that it also satisfies item (3). Thus, the Gromov radius is a symplectic capacity on the category of all symplectic manifolds. Since this definition was formulated, there have been many more examples of symplectic capacities, see [13].

2.2 Integrable systems

Let (M, ω) be a symplectic manifold and let $\mathfrak{X}(M)$ denote the collection of vector fields on M . Since ω is non-degenerate the mapping

$$\begin{aligned} \mathfrak{X}(M) &\rightarrow \Omega^1(M) \\ X &\mapsto \omega(X, \cdot) \end{aligned}$$

is an isomorphism. This means that given any 1-form η there exists a unique $X \in \mathfrak{X}(M)$ such that $\eta = \omega(X, \cdot)$. In particular, for each $f: M \rightarrow \mathbb{R}$ there exists a unique vector field, called the *Hamiltonian vector field of f* and denoted \mathcal{X}_f , on M which satisfies $-df = \omega(\mathcal{X}_f, \cdot)$. Hamiltonian vector fields preserve the symplectic form because

$$\mathcal{L}_{\mathcal{X}_f}\omega = d\iota_{\mathcal{X}_f}\omega + \iota_{\mathcal{X}_f}d\omega = -d(df) = 0$$

making use of Cartan's magic formula and the flow of \mathcal{X}_f preserves f because

$$\mathcal{L}_{\mathcal{X}_f}f = \mathcal{X}_f(f) = df(\mathcal{X}_f) = -(\iota_{\mathcal{X}_f}\omega)(\mathcal{X}_f) = -\omega(\mathcal{X}_f, \mathcal{X}_f) = 0.$$

2.2.1 Classical mechanics

A Hamiltonian dynamical system is a triple (M, ω, H) where $H: M \rightarrow \mathbb{R}$ is the Hamiltonian function and represents the total energy of the system. Such a system evolves by flowing along \mathcal{X}_H .

As a first example, consider \mathbb{R}^{2n} with coordinates $(q_1, \dots, q_n, p_1, \dots, p_n)$ and the usual symplectic form $\omega = \sum_{i=1}^n dq_i \wedge dp_i$. Let us compute \mathcal{X}_H in this case. Let $\mathcal{X}_H = \sum_{i=1}^n a_i \frac{\partial}{\partial q_i} + b_i \frac{\partial}{\partial p_i}$. Then $\iota_{\mathcal{X}_H}\omega = H$ implies that

$$\sum_{i=1}^n a_i dp_i - b_i dq_i = \sum_{i=1}^n \frac{\partial H}{\partial q_i} dq_i + \frac{\partial H}{\partial p_i} dp_i$$

and thus $\mathcal{X}_H = \sum_{i=1}^n \frac{\partial H}{\partial p_i} \frac{\partial}{\partial q_i} - \frac{\partial H}{\partial q_i} \frac{\partial}{\partial p_i}$. This means that a curve $\gamma(t) = (q(t), p(t))$ is an integral curve for \mathcal{X}_H if and only if

$$\begin{aligned} \frac{dq_i}{dt} &= \frac{\partial H}{\partial p_i}; \\ \frac{dp_i}{dt} &= -\frac{\partial H}{\partial q_i}, \end{aligned}$$

recovering Hamilton's equations from classical mechanics.

2.2.2 Integrable systems

Recall that a Poisson bracket on $C^\infty(M)$ is a bilinear map

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

satisfying the following for $f, g, h \in C^\infty(M)$:

1. *skew symmetry*: $\{f, g\} = -\{g, f\}$;
2. *Jacobi identity*: $\{f, \{g, h\}\} = \{\{f, g\}, h\} + \{g, \{f, h\}\}$;
3. *Leibniz rule*: $\{f, gh\} = \{f, g\}h + g\{f, h\}$.

The symplectic form on M determines a natural Poisson bracket on $C^\infty(M)$, given by

$$\{f, g\} = -\omega(\mathcal{X}_f, \mathcal{X}_g).$$

Proposition 2.2.1. *Let $f, H \in C^\infty(M)$ where (M, ω) is a symplectic manifold. Then $\{f, H\} = 0$ if and only if f is preserved by the flow of \mathcal{X}_H .*

Proof. This is because

$$\mathcal{L}_{\mathcal{X}_H}(f) = \mathcal{X}_H(f) = df(\mathcal{X}_H) = (-\iota_{\mathcal{X}_f}\omega)(\mathcal{X}_H) = -\omega(\mathcal{X}_f, \mathcal{X}_H) = \{H, f\}. \quad \square$$

Definition 2.2.2. An *integrable system* is a triple (M, ω, F) where (M, ω) is a $2n$ -dimensional symplectic manifold and $F = (f_1, \dots, f_n): M \rightarrow \mathbb{R}^n$ satisfies

1. $(df_1)_p, \dots, (df_n)_p$ are linearly independent in T_p^*M for almost all $p \in M$;
2. f_1, \dots, f_n Poisson commute. That is, $\{f_i, f_j\} = 0$ for all $i, j < n$.

It is important to notice that the differentials df_1, \dots, df_n may be dependent on a set of measure zero in M . These points are the singularities of the integrable system and they are precisely the points at which the most interesting dynamics occur.

The map F is known as the *momentum map* and its components, f_1, \dots, f_n , represent conserved quantities of the corresponding physical system. By Proposition 2.2.1 we see that $\{f_i, f_j\} =$

0 implies that f_i is preserved by the flow of \mathcal{X}_{f_j} for each i, j . If $H = f_1$ is the Hamiltonian function of the system then each f_i is preserved under the flow of \mathcal{X}_H .

Remark 2.2.3. *A note on momentum maps:* The momentum map of an integrable system $F: M \rightarrow \mathbb{R}^n$ and the momentum map of a Hamiltonian G -manifold $\mu: M \rightarrow \text{Lie}(G)^*$ are related but not the same. It is unlucky that they have the same name. Let (M, ω, F) be any integrable system such that $\mathcal{X}_{f_1}, \dots, \mathcal{X}_{f_n}$ are complete. This is automatically true, for instance, if F is proper. Let $B: \text{Lie}(G)^* \rightarrow \mathbb{R}^n$ be any isomorphism of vector spaces and $\phi: \mathbb{R}^n \times M \rightarrow M$ take $((t_1, \dots, t_n), m)$ to the point on M obtained by flowing m along \mathcal{X}_{f_i} for time t_i , $i = 1, \dots, n$. The order of the flows doesn't matter because the functions Poisson commute. Then $(M, \omega, \phi, B \circ F)$ is a Hamiltonian \mathbb{R}^n -manifold. \circlearrowright

2.2.3 Example

Consider \mathbb{R}^{2n} with the standard symplectic form

$$\omega = \sum_{i=1}^n dx_i \wedge dy_i$$

and $F = (f_1, \dots, f_n): \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$ given by

$$f_i = \frac{1}{2}(x_i^2 + y_i^2).$$

Then

$$\mathcal{X}_{f_i} = x_i \frac{\partial}{\partial x_i} + y_i \frac{\partial}{\partial y_i}$$

for $i = 1, \dots, n$. The flow of \mathcal{X}_{f_i} is circular motion in a plane parallel to $\{x_i = y_i = 0\}$, and the origin is a fixed point of the flow of \mathcal{X}_{f_i} for all i . By taking $H: \mathbb{R}^{2n} \rightarrow \mathbb{R}$ to be given by

$$H = f_1 + \dots + f_n$$

one obtains a model of the the motion of a particle in a harmonic oscillator centered at the origin. The coordinates x_1, \dots, x_n describe the location of the particle and the coordinates y_1, \dots, y_n describe

its momenta.

2.2.4 Regular points of integrable systems

If $p \in M$ is a regular point of F , where $(df_1)_p, \dots, (df_n)_p$ are independent, then

$$(\mathcal{X}_{f_1})_p, \dots, (\mathcal{X}_{f_n})_p \in T_p M$$

are also linearly independent. Also,

$$\omega_p((\mathcal{X}_{f_i})_p, (\mathcal{X}_{f_j})_p) = \{f_i, f_j\}(p) = 0$$

which means that $\text{span}((\mathcal{X}_{f_1})_p, \dots, (\mathcal{X}_{f_n})_p) \subset T_p M$ is a isotropic subspace of $T_p M$. Thus, since it is dimension n , this means that for any regular value $c \in \mathbb{R}^n$ the submanifold $F^{-1}(c)$ of M is a Lagrangian submanifold. From this computation we can also see that n is the maximal number of such functions which can exist by Proposition 2.1.2 which states that a isotropic subspace can have dimension at most n . An integrable system may be viewed as a dynamical system with the maximal amount of independent Poisson-commuting conserved quantities.

Lemma 2.2.4. *Let (M, ω, F) be an integrable system and let $c \in \mathbb{R}^n$ be a regular value of F . If $\mathcal{X}_{f_1}, \dots, \mathcal{X}_{f_n}$ are complete on $F^{-1}(c)$ then each connected component of $F^{-1}(c)$ is diffeomorphic to $\mathbb{R}^{n-k} \times \mathbb{T}^k$ for some k .*

Proof. Use the flows of \mathcal{X}_{f_i} to produce charts on $F^{-1}(c)$. □

The following theorem completely describes the local dynamics at regular points of the momentum map. Identify $T^*\mathbb{T}^n$, the cotangent bundle of the n -torus, with $\mathbb{T}^n \times \mathbb{R}^n$ by choosing coordinates $(x_1, \dots, x_n, \xi_1, \dots, \xi_n)$ chosen such that the canonical 1-form is written $\alpha = \sum_{i=1}^n \xi_i \wedge dx_i$.

Theorem 2.2.5 (Liouville-Arnold Theorem [2]). *Let (M, ω, F) be an integrable system with F proper and let $c \in \mathbb{R}^n$ be a regular point of F . Then there exists a symplectomorphism χ from a neighborhood of the zero section of $T^*\mathbb{T}^n$ to a neighborhood of each connected component of $F^{-1}(c)$ such that $F \circ \chi = \phi \circ (\xi_1, \dots, \xi_n)$ where ϕ is some local diffeomorphism of \mathbb{R}^n .*

This theorem, known as the Action-Angle Theorem, states that if F is proper then the preimage of a regular point is always a torus, known as a Liouville torus when they exist, on which the flow of the Hamiltonian vector fields of each f_i is linear and moreover that in a neighborhood of that fiber there exist coordinates complimentary to those on the torus which create a Darboux chart.

Notice that this result is of a *semiglobal* nature. That is, it gives a local normal form for a neighborhood of a fiber $F^{-1}(c)$. It is not global but it gives more information than a local theory, which would only give a normal form for a neighborhood of a point $p \in M$.

Let $p \in M$ be such that $F(p) = c$ is a regular point of F and as before view (M, ω, F) as a Hamiltonian system by taking $H = f_1$. Then, by Theorem 2.2.5, the corresponding Hamiltonian system evolves from an initial condition of p along a linear path in the torus $F^{-1}(c)$. The dynamics at regular points of an integrable system are very well understood. We will see that the dynamics around the points that are not regular, the singular points, can be very complicated and the existence of certain types of singular points can have surprising global effects on the system.

2.2.5 Singular points of integrable systems

The point $p \in M$ is a critical (or singular) point of $F: M \rightarrow \mathbb{R}^n$ and $c = F(p)$ is a critical (or singular) value if $(df_1)_p, \dots, (df_n)_p$ are not linearly independent. This is equivalent to the map $d_p F: T_p M \rightarrow \mathbb{R}^n$ having rank less than n .

After covering the required background, we will define a notion of non-degenerate singular points of integrable systems which is analogous to non-degenerate singularities in Morse theory. In short, a singular point $p \in M$ is non-degenerate if the Hessians of f_1, \dots, f_n span a Cartan subalgebra of the space of quadratic forms on $T_p M$. In this section, we explain this definition in detail and describe two important results about non-degenerate singular points: Williamson's pointwise classification of non-degenerate points and Eliasson's local normal forms for non-degenerate points.

This description is standard, and can for instance be found in [8, 79, 87].

Background: $\mathcal{Q}(\mathbb{T}_p M)$ and $\mathfrak{sp}(2n, \mathbb{R})$

Here we quickly review how the space of quadratic forms on a linear vector space can be identified with the Lie algebra $\mathfrak{sp}(2n, \mathbb{R})$ and mention several other results which will be necessary for studying singular points of integrable systems.

Let (M, ω, F) be an n -dimensional integrable system and let $p \in M$. Fix symplectic coordinates $x_1, y_1, \dots, x_n, y_n$ on $\mathbb{T}_p M$ and view the elements $x \in \mathbb{T}_p M$ as column vectors. Denote by $\mathcal{Q}(\mathbb{T}_p M)$ the quadratic forms on M . That is,

$$\mathcal{Q}(\mathbb{T}_p M) = \left\{ q \in C^\infty(\mathbb{T}_p M) \mid q(x) = \frac{1}{2} x^t A_q x \text{ for } A_q \text{ a symmetric } 2n \times 2n \text{ real matrix} \right\}$$

where x^t denotes the transpose of x . The matrix A_q is the Hessian of the function $q: \mathbb{T}_p M \rightarrow \mathbb{R}$. That is, in coordinates $w_1 = x_1, \dots, w_n = x_n, w_{n+1} = y_1, \dots, w_{2n} = y_n$,

$$(A_q)_{ij} = \frac{\partial^2 q}{\partial w_i \partial w_j}.$$

Since $(\mathbb{T}_p M, \omega_p)$ is a symplectic vector space, and thus a symplectic manifold, it comes equipped with a Poisson bracket on $C^\infty(\mathbb{T}_p M)$, which we denote $\{\cdot, \cdot\}_p$. Thus, $(\mathcal{Q}(\mathbb{T}_p M), \{\cdot, \cdot\}_p)$ is a Lie algebra. Let $\mathfrak{sp}(2n, \mathbb{R})$ denote the Lie algebra of real Hamiltonian matrices. That is,

$$\mathfrak{sp}(2n, \mathbb{R}) = \{A \in \mathfrak{gl}(2n, \mathbb{R}) \mid A^t J^t = JA\}$$

where

$$J = \begin{pmatrix} 0 & -I_n \\ I_n & 0 \end{pmatrix}$$

and I_n is the $n \times n$ identity matrix. The Lie bracket on $\mathfrak{sp}(2n, \mathbb{R})$ is given by the usual matrix commutator $[A, B] = AB - BA$.

Proposition 2.2.6. *The Lie algebras $\mathcal{Q}(\mathbb{T}_p M, \{\cdot, \cdot\}_p)$ and $(\mathfrak{sp}(2n, \mathbb{R}), [\cdot, \cdot])$ are isomorphic.*

Proof. It is clear that $A_q = A_{q'}$ implies that $q = q'$. A matrix B is in $\mathfrak{sp}(2n, \mathbb{R})$ if and only if $B = J^{-1}A$ for a symmetric matrix A , since $B^t J^t = JB$ if and only if $(JB)^t = JB$. The isomorphism from $\mathcal{Q}(\mathbb{T}_p M)$ to $\mathfrak{sp}(2n, \mathbb{R})$ is given by $q \mapsto J^{-1}A_q$.

Next we will show that this preserves brackets. Use symplectic coordinates $w_1 = x_1, \dots, w_n = x_n, w_{n+1} = y_1, \dots, w_{2n} = y_n$ and write $A_q = (a_{ij})$ and $A_{q'} = (a'_{ij})$. Then use

$$\begin{aligned} \{q, q'\}_p &= \sum_{k=1}^n \frac{\partial q}{\partial x_k} \frac{\partial q'}{\partial y_k} - \frac{\partial q}{\partial y_k} \frac{\partial q'}{\partial x_k} \\ &= \sum_{k=1}^n \left(\left(\sum_{i=1}^{2n} a_{ki} w_i \right) \left(\sum_{i=1}^{2n} a'_{(k+n)i} w_i \right) - \left(\sum_{i=1}^{2n} a_{(k+n)i} w_i \right) \left(\sum_{i=1}^{2n} a'_{ki} w_i \right) \right) \\ &= \sum_{k=1}^n \left(\sum_{i,j=1}^{2n} (a_{ki} a'_{(k+n)j} - a_{(k+n)i} a'_{kj}) w_i w_j \right) \\ &= \sum_{i,j=1}^{2n} \left(\sum_{k=1}^n a_{ki} a'_{(k+n)j} - a_{(k+n)i} a'_{kj} \right) w_i w_j \end{aligned}$$

and write out

$$[J^{-1}A_q, J^{-1}A_{q'}] = [J^{-1}(a_{ij}), J^{-1}(a'_{ij})]$$

in coordinates to conclude that

$$J^{-1}A_{\{q,q'\}_p} = [J^{-1}A_q, J^{-1}A_{q'}]$$

□

Given any function $f: M \rightarrow \mathbb{R}$ such that $d_p f = 0$ the *Hessian of f at p* is the quadratic form $\mathcal{H}_p(f) \in \mathcal{Q}(\mathbb{T}_p M)$ with associated matrix

$$A_{\mathcal{H}_p(f)} = \left(\frac{\partial^2 f}{\partial w_i \partial w_j} (p) \right)$$

where w_1, \dots, w_{2n} are now coordinates on M .

Proposition 2.2.7. *If $f, g \in C^\infty(M)$ and $p \in M$ is a singular point of both f and g then*

$$\mathcal{H}_p\{f, g\} = \{\mathcal{H}_p(f), \mathcal{H}_p(g)\}_p.$$

Proof. Write out both sides in coordinates. □

The following will be useful when studying integrable systems.

Corollary 2.2.8. *Let $f, g \in C^\infty(M)$ be such that $d_p f = d_p g = 0$. If $\{f, g\} = 0$ then*

$$\{\mathcal{H}_p(f), \mathcal{H}_p(g)\} = 0.$$

Background: Cartan subalgebras of $\mathcal{Q}(T_p M)$

Here we very briefly indicate the results we will need regarding Cartan subalgebras of $\mathcal{Q}(T_p M)$ which, in light of Proposition 2.2.6, we identify with $\mathfrak{sp}(2n, \mathbb{R})$.

If \mathfrak{g} is any Lie algebra then the lower central series is the series

$$\mathfrak{g} = \mathfrak{g}^1 \supseteq \mathfrak{g}^2 \supseteq \dots$$

defined by $\mathfrak{g}^n = [\mathfrak{g}, \mathfrak{g}^{n-1}]$. A subalgebra $\mathfrak{h} \leq \mathfrak{g}$ is *nilpotent* if $\mathfrak{h}^n = 0$ for some n . The *normalizer* of a subalgebra $\mathfrak{h} \leq \mathfrak{g}$ is given by

$$N_{\mathfrak{g}}(\mathfrak{h}) = \{a \in \mathfrak{g} \mid [a, \mathfrak{h}] \subset \mathfrak{h}\}.$$

Definition 2.2.9. A subalgebra $\mathfrak{h} \leq \mathfrak{g}$ is a *Cartan subalgebra* if it is nilpotent and self-normalizing (i.e. $\mathfrak{h} = N_{\mathfrak{g}}(\mathfrak{h})$).

For any element $a \in \mathfrak{g}$, denote by $C(a)$ the commutator of a given by

$$C(a) = \{b \in \mathfrak{g} \mid ab = ba\}.$$

An element of a Lie algebra is *regular* if its commutator is of minimal dimension among commutators of elements of \mathfrak{g} . In the case of a matrix Lie algebra, such as $\mathfrak{sp}(2n, \mathbb{R})$, this amounts to a having all distinct eigenvalues, in which case $\dim(C(a)) = n$.

Non-degenerate singular points

We now have enough machinery to define the singular points which can arise in the integrable systems we study. First consider the case of singular points with rank 0.

Definition 2.2.10. Let $(M, \omega, F = (f_1, \dots, f_n))$ be an integrable system. If $p \in M$ is such that

$d_p F \equiv 0$ then p is a *non-degenerate singular point* if

$$\text{span}(\mathcal{H}_p(f_1), \dots, \mathcal{H}_p(f_n)) \subset \mathcal{Q}(\mathbb{T}_p M)$$

is a Cartan subalgebra.

Suppose that $p \in M$ is a singular point of F with rank $n - k$ for some $0 < k < n$. Let K denote the kernel of $d_p F: \mathbb{T}_p M \rightarrow \mathbb{T}_{F(p)} \mathbb{R}^n$ and let

$$\mathcal{J} = \text{span}((\mathcal{X}_{f_1})_p, \dots, (\mathcal{X}_{f_n})_p) \subset \mathbb{T}_p M.$$

Then \mathcal{J} is a maximally isotropic subspace of (K, ω_p) which means that

$$\bar{\omega}_p(v + \mathcal{J}, w + \mathcal{J}) := \omega_p(v, w)$$

is a well-defined symplectic form $\bar{\omega}_p$ on K/\mathcal{J} . The quotient K/\mathcal{J} can be identified with a subspace $R \subset \mathbb{T}_p M$. The dimension of R is $2k$.

Definition 2.2.11. A singular point $p \in M$ of rank $n - k$ is non-degenerate if $\mathcal{H}_p(f_1), \dots, \mathcal{H}_p(f_n)$ span a Cartan subalgebra of $\mathcal{Q}(R)$.

Suppose that (M, ω, F) is an integrable system where $F = (f_1, \dots, f_n)$ and suppose that $d_p F = 0$ for some $p \in M$. Since $\{f_i, f_j\} = 0$ we know that

$$\{\mathcal{H}_p(f_i), \mathcal{H}_p(f_j)\}_p = 0$$

by Corollary 2.2.8. This means that $\mathcal{H}_p(f_1), \dots, \mathcal{H}_p(f_n)$ span an abelian subalgebra of $\mathcal{Q}(\mathbb{T}_p M)$. If they are independent, so the span is of dimension n , and their span includes a regular element then p is non-degenerate. This idea is used to prove the following Proposition, which is a useful characterization of singular points of any rank.

Proposition 2.2.12 (Bolsinov-Fomenko [8]). *Let $p \in M$ be a singular point such that the rank of $d_p F$ is k . Then there exists functions g_1, \dots, g_n related to f_1, \dots, f_n by a linear transformation such that:*

1. $d_p g_1 = \dots = d_p g_{n-k} = 0$;
2. $(\mathcal{X}_{g_{n-k+1}})_p, \dots, (\mathcal{X}_{g_n})_p$ are linearly independent in $T_p M$.

Let $L = \text{span}\{(\mathcal{X}_{g_1})_p, \dots, (\mathcal{X}_{g_{n-k}})_p\} \subset T_p M$ and let

$$L^{\omega_p} = \{v \in T_p M \mid \omega_p(v, w) = 0 \text{ for all } w \in L\}$$

The point p is non-degenerate if and only if $\mathcal{H}_p(g_1), \dots, \mathcal{H}_p(g_{n-k})$ are independent and there exists a matrix

$$A \in \text{span}\{\mathcal{H}_p(g_{n-k+1}), \dots, \mathcal{H}_p(g_n)\} \subset \mathcal{Q}(T_p M)$$

such that $J^{-1}A$ has $2(n-k)$ distinct eigenvalues in L^J .

Since every Cartan subalgebra of $\mathcal{Q}(T_p M)$ is generated as the centralizer of a regular element (denoted A in Proposition 2.2.12), we can classify such subalgebras by the types of eigenvalues which appear in the associated regular element.

Pointwise classification of singular points

To classify non-degenerate singular points we appeal to a classification of Cartan subalgebras of $\mathfrak{sp}(2n, \mathbb{R})$ due to Williamson [82]. Let $q \in \mathcal{Q}(T_p M)$ be a regular element. Then Williamson [82] showed the following. The eigenvalues of a are divided into three types of groups:

1. Pairs of imaginary roots $\pm i\alpha$ (elliptic block);
2. Pairs of real roots $\pm\beta$ (hyperbolic block);
3. Quadruples of roots $\pm\alpha \pm i\beta$ (focus-focus block).

The matrix $J^{-1}A_q$ corresponding to q can be written in blocks. For each elliptic pair of eigenvalues $\pm i\alpha$ there is a 2×2 block

$$\begin{pmatrix} 0 & \alpha \\ -\alpha & 0 \end{pmatrix},$$

for each hyperbolic pair $\pm\beta$ there is a 2×2 block

$$\begin{pmatrix} \beta & 0 \\ 0 & -\beta \end{pmatrix},$$

and for each focus-focus quadruple $\pm\alpha \pm i\beta$ there is a 4×4 block

$$\begin{pmatrix} \alpha & \beta & 0 & 0 \\ -\beta & \alpha & 0 & 0 \\ 0 & 0 & -\alpha & \beta \\ 0 & 0 & -\beta & -\alpha \end{pmatrix}.$$

The matrices which commute with A_q preserve these blocks while scaling them. Translating these matrices back into quadratic operators implies the following classification of Cartan subalgebras of $\mathcal{Q}(T_pM)$.

Theorem 2.2.13 (Williamson [82]). *For any Cartan subalgebra $\mathfrak{h} \subset \mathcal{Q}(T_pM)$ there exist symplectic coordinates $x_1, y_1, \dots, x_n, y_n$ on T_pM and a basis q_1, \dots, q_n of \mathfrak{h} such that each q_i is one of:*

1. *elliptic type:* $q_i = x_i^2 + y_i^2$;
2. *hyperbolic type:* $q_i = x_i y_i$;
3. *focus-focus type:* $\begin{cases} q_i = x_i y_{i+1} - x_{i+1} y_i; \\ q_{i+1} = x_i y_i + x_{i+1} y_{i+1}; \end{cases}$

Theorem 2.2.13 implies a classification of singular points for integrable system: the *Williamson type* [87] of a singular point is given by the integers (m_e, m_h, m_f) where the Hessians span a Cartan subalgebra whose basis from Theorem 2.2.13 has m_e elliptic blocks, m_h hyperbolic blocks, and m_f focus-focus blocks. The rank of the singular point is by $m_e + m_h + 2m_f$.

Local normal forms for singular points

The definition of non-degenerate and the Williamson type of a singular point are both *pointwise* definitions. Here we review a result of Eliasson [28, 29] which states that the Williamson classification can be used to construct *local* normal forms of non-degenerate singular points.

Theorem 2.2.14 (Eliasson [28, 29]). *If $p \in M$ is a non-degenerate singular point of an integrable system (M, ω, F) then there exists local symplectic coordinates $(x_1, \dots, x_n, y_1, \dots, y_n)$ about p such that there exist $q_1, \dots, q_n: M \rightarrow \mathbb{R}$ where each q_i is given by one of:*

1. *elliptic:* $q_i = \frac{x_i^2 + y_i^2}{2}$;
2. *hyperbolic:* $q_i = x_i y_i$;
3. *focus-focus:* $\begin{cases} q_i = x_i y_{i+1} - x_{i+1} y_i; \\ q_{i+1} = x_i y_i x_{i+1} y_{i+1}; \end{cases}$
4. *non-singular:* $q_i = y_i$,

such that $\{f_i, q_j\} = 0$ for all $i, j = 1, \dots, n$. If p has no hyperbolic blocks then the assertion that $\{f_i, q_j\} = 0$ for all $i, j = 1, \dots, n$ may be replaced by the equality

$$(F - F(p)) \circ \phi = g \circ q$$

where $q = (q_1, \dots, q_n)$, $\phi = (x_1, \dots, x_n, y_1, \dots, y_n)^{-1}$, and g is a local diffeomorphism of \mathbb{R}^n which fixes the origin.

2.2.6 Toric integrable systems

An integrable system (M, ω, F) is *toric* if M is compact and the Hamiltonian flows associated to the components of F are all periodic of the same period. That is, if $B: \text{Lie}(\mathbb{T}^n)^* \rightarrow \mathbb{R}^n$ is an isomorphism of vector spaces then $B^{-1} \circ F$ is the momentum map (in the group action sense) for a Hamiltonian \mathbb{T}^n -action. Such systems are classified, up to isomorphism, by the image of their momentum map, which is necessarily a specific type of convex polytope known as a Delzant polytope. Two toric integrable systems (M, ω, F) and (M', ω', F') are isomorphic if there exists a symplectomorphism $\phi: M \rightarrow M'$ such that $\phi^* F' = F$.

A *Delzant polytope* $\Delta \subset \mathbb{R}^n$ is a convex polytope which is also:

1. *rational:* each face of the polytope has a vector in \mathbb{Z}^n which is perpendicular to it (i.e. it has rational slope);
2. *simple:* n edges meet at each vertex;

3. *smooth*: for any vertex p of Δ there exists $v_1, \dots, v_n \in \mathbb{Z}^n$ each an inwards pointing normal vector for a face which is adjacent to p such that v_1, \dots, v_n spans \mathbb{Z}^n .

Denote the set of Delzant polytopes in \mathbb{R}^n by \mathcal{P}_T .

Using as a stepping stone the convexity theorem of Atiyah [3], Guillemin-Sternberg [40], Delzant proved the following.

Theorem 2.2.15 (Delzant classification [21]). *If (M, ω, F) is a toric integrable system then $F(M) \subset \mathbb{R}^n$ is a Delzant polytope. Furthermore, given any Delzant polytope $\Delta \subset \mathbb{R}^n$ there exists a toric integrable system (M, ω, F) such that $F(M) = \Delta$ and such a system is unique up to isomorphism.*

All singular points of toric integrable systems are automatically non-degenerate singular points with only elliptic blocks. The singular points of rank n are precisely the preimages under F of the vertices of the Delzant polytope, $F(M)$.

Toric fans and minimal models

In addition to Theorem 2.2.15, Delzant also showed that every toric integrable system is naturally a toric variety. Toric varieties [16, 17, 19, 33, 56, 61] have been extensively studied in algebraic and differential geometry and so have their symplectic analogues.

The relationship between toric integrable systems and toric varieties has been understood since the 1980s, see for instance Delzant [21], Guillemin [38, 39]. The article [24] contains a coordinate description of this relation.

Associated to every toric variety is a fan, which can be recovered from the associated toric integrable system as the collection of inwards pointing integral normal vectors to the faces of the Delzant polytope. We will specialize to dimension $2n = 4$.

A *toric fan* is a collection of vectors $(v_1, \dots, v_{d-1}) \in \mathbb{Z}^d$ which are arranged in counter-clockwise order and satisfy $\det(v_i, v_{i+1}) = 1$ for $i = 0, \dots, d-1$ where v_d denotes v_0 . The inwards pointing normal vectors of any Delzant polytope in \mathbb{R}^2 form a toric fan and any toric fan can be achieved in this way. What we call a toric fan is the type of fan which is associated to a non-singular, complete toric surface.

If (v_0, \dots, v_{d-1}) is a toric fan then $(v_0, \dots, v_i, v_i + v_{i+1}, v_{i+1}, \dots, v_{d-1})$ is also a toric fan and is known as the *blowup* of the original fan. The inverse of this operation, removing a vector

which is the sum of two adjacent vectors, is known as a *blowdown*. A toric fan is *minimal* if it does not admit a blowdown. That is, (v_0, \dots, v_{d-1}) is minimal if

$$v_i \neq v_{i-1} + v_{i+1}$$

for $i = 0, \dots, d-1$.

Minimal toric fans have a simple classification. I provide an alternative proof for the following well-known result in Section 4.4.

Theorem 2.2.16 (Fulton [33], page 44). *Up to the action of $\mathrm{SL}_2(\mathbb{Z})$ all toric fans are of one of the following three types:*

1. $v_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, v_2 = \begin{pmatrix} -1 \\ -1 \end{pmatrix};$
2. $v_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, v_2 = \begin{pmatrix} -1 \\ 0 \end{pmatrix}, v_3 = \begin{pmatrix} 0 \\ -1 \end{pmatrix};$
3. $v_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, v_2 = \begin{pmatrix} -1 \\ k \end{pmatrix}, v_3 = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$ for $k \in \mathbb{Z}, k \neq 0$.

This implies that any toric fan may be obtained from one of the three minimal models by a finite sequence of blowups. Blowup of the toric fan correspond to the removal of an equivariantly embedded ball from the associated toric integrable system. By equivariant embedded ball we mean the image in M of an embedding $\rho: \mathbb{B}^{2n}(r) \hookrightarrow M$ for some r such that rotations of the coordinates of the ball (as a subset of \mathbb{C}^n) agree with the \mathbb{T}^n -action on M up to an element of $\mathrm{Aut}(\mathbb{T}^n)$.

Metric on toric integrable systems

In [67] the authors construct a natural metric on the moduli space of toric integrable systems making use of the Delzant classification, Theorem 2.2.15. For $A, B \subset \mathbb{R}^n$ let $A \ominus B = (A \setminus B) \cup (B \setminus A)$ denote the symmetric difference and let λ denote the Lebesgue measure on \mathbb{R}^n . Given two toric integrable systems (M, ω, F) and (M', ω', F') with associated Delzant polytopes Δ and Δ' , respectively, the distance between (M, ω, F) and (M', ω', F') is given by $\Delta \ominus \Delta'$.

2.2.7 Semitoric integrable systems

A *semitoric integrable system* [69], or *symplectic semitoric manifold*, is a 4-dimensional, connected, symplectic manifold (M, ω) with a momentum map $F = (J, H) : M \rightarrow \mathbb{R}^2$ such that:

1. the flow of the Hamiltonian vector field \mathcal{X}_J is periodic;
2. J is proper;
3. F has only non-degenerate singularities without real-hyperbolic blocks (as in Section 2.2.5).

Notice that though semitoric systems are required to be 4-dimensional there is much more freedom in the choice of momentum map compared to toric systems and M is not required to be compact. Since the singularities of semitoric systems are non-degenerate and not of hyperbolic type by condition (3), Theorem 2.2.14 implies any critical point of F has one of three possible forms. If $p \in M$ is a critical point of F then there exists g a local diffeomorphism of \mathbb{R}^n which fixes the origin, Darboux coordinates (x_1, x_2, ξ_1, ξ_2) about p , in which p is represented by $(0, 0, 0, 0)$, and $q = (q_1, q_2) : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that

$$(F - F(p)) \circ (x_1, x_2, \xi_1, \xi_2)^{-1} = g \circ (q_1, q_2)$$

and q_1, q_2 are in one of the following forms:

1. *elliptic-elliptic*: $q_1 = x_1^2 + \xi_1^2/2$ and $q_2 = x_2^2 + \xi_2^2/2$;
2. *transversely-elliptic*: $q_1 = x_1^2 + \xi_1^2/2$ and $q_2 = \xi_2$;
3. *focus-focus*: $q_1 = x_1\xi_2 - x_2\xi_1$ and $q_2 = x_1\xi_1 + x_2\xi_2$.

A semitoric integrable system $(M, \omega, F = (J, H))$ is said to be *simple* if there is at most one focus-focus critical point in $J^{-1}(x)$ for all $x \in \mathbb{R}$. A similar (but weaker) assumption is generic according to Zung [87], that each fiber $F^{-1}(c)$ for $c \in \mathbb{R}^2$ contains at most one critical point $p \in M$. Any semitoric system has only finitely many focus-focus critical points [77] so we will denote them by $m_1, \dots, m_{m_f} \in M$ and the associated singular values are denoted $c_j = F(m_j)$, $j = 1, \dots, m_f$. All semitoric systems studied in this dissertation are assumed to be simple and we assume $J(m_1) < \dots < J(m_{m_f})$. Suppose that $(M_i, \omega_i, F_i = (J_i, H_i))$ is a semitoric system for $i = 1, 2$. An *isomorphism of semitoric systems* is a symplectomorphism $\phi : M_1 \rightarrow M_2$ such that $\phi^*(J_2, H_2) = (J_1, f(J_1, H_1))$ where $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a smooth function such that $\frac{\partial f}{\partial H_1}$ nowhere vanishes.

We denote by \mathcal{M}_{ST} the space of simple semitoric systems modulo isomorphism. Semitoric integrable systems have, in addition to the elliptic singularities of toric systems, also singularities of focus-focus type.

Semitoric systems have been studied by mathematicians and physicists in the past decade, and there have been contributions to their study from many points of view, including mathematical physics (eg. see Babelon-Douçot [4, 5], Dullin [25]), symplectic topology (eg. see Eliashberg-Polterovich [27] or Leung-Symington [53]), and mirror symmetry (see Gross and Siebert [36, 37]). While this dissertation is focused on classical integrable systems, much of the work on these systems is motivated by inverse spectral problems about quantum integrable systems as pioneered in the work of Colin de Verdière [14, 15] and others, and which also has been the subject of recent works [11, 86].

Semitoric integrable systems are of interest in mathematical physics, symplectic geometry and spectral theory, because they exhibit rich features from dynamical, geometric, and topological viewpoints. Examples of semitoric systems permeate the physics literature. For instance the Jaynes-Cummings system [18, 46], which is one of the most thoroughly studied examples of semitoric system [73], models simple physical phenomena. It is obtained by coupling a spin and an oscillator, and its phase space is $S^2 \times \mathbb{R}^2$.

In [69, 70], Pelayo-Vũ Ngọc provide a complete classification for semitoric systems in terms of a collection of several invariants. We have included this as Theorem 2.2.27. While compact toric integrable systems are classified in terms of Delzant polytopes, in the semitoric case a polygon plays a role but the complete invariant must contain more information. Loosely speaking, the complete invariant of semitoric systems is a collection of convex polygons in \mathbb{R}^2 (which may not be compact) each with a finite number of distinguished points corresponding to the focus-focus singularities labeled by a Taylor series and an integer (See Figure 2.1).

The number of singular points invariant

In [77, Theorem 1] Vũ Ngọc proves that a semitoric system has finitely many focus-focus singular points. Thus, to a system we may associate an integer $0 \leq m_f < \infty$ which is the total number of focus-focus points in the system. The singular points are preserved by isomorphism so this is an invariant of the system. For any nonnegative integer $m_f \in \mathbb{Z}_{\geq 0}$ let $\mathcal{M}_{\text{ST}, m_f}$ denote the




Figure 2.1: The complete semitoric invariant is a collection of polygons with distinguished points $\{c_1, \dots, c_{m_f}\}$ each labeled with extra information: a Taylor series $(S_j)^\infty$, an integer k_j known as the twisting index, and an element $\epsilon_j \in \{-1, +1\}$ known as the cut direction. There is one polygon in the family for each possible choice of cut directions and each allowed choice of twisting indices.

collection of simple semitoric systems with m_f focus-focus points modulo semitoric isomorphism.

The Taylor series invariant

The next invariant we will study completely classifies the structure of a focus-focus critical point in the neighborhood of a fiber up to isomorphism [76]. It is defined in terms of the length of certain flow lines of the Hamiltonian vector fields for the components of the momentum map. The details can be found in [76, 71].

Definition 2.2.17. Let $\mathbb{R}[[X, Y]]$ refer to the algebra of real formal power series in two variables and let $\mathbb{R}[[X, Y]]_0 \subset \mathbb{R}[[X, Y]]$ be the subspace of series $\sum_{i,j \geq 0} \sigma_{i,j} X^i Y^j$ which have $\sigma_{0,0} = 0$ and $\sigma_{0,1} \in [0, 2\pi)$.

The Taylor series invariant is one element of $\mathbb{R}[[X, Y]]_0$ for each of the m_f focus-focus points.

The affine invariant and the twisting index invariant

The affine invariant is similar to the polygon from Delzant's result, except in this case we instead have a family of polygons related by specific linear transformations. The twisting index describes how each critical point sits with respect to a privileged momentum map. These two invariants will be described together because the twisting indices which label each critical point will

be defined only up to the addition of a common integer related to the choice of polygon.

A *convex polygon* is the intersection in \mathbb{R}^2 of (finitely or infinitely many) closed half planes such that on each compact subset of \mathbb{R}^2 there are at most finitely many corner points. A convex polygon is rational if each edge is directed along a vector with rational coefficients. We denote the set of all rational convex polygons by $\text{Polyg}(\mathbb{R}^2)$. For $\lambda \in \mathbb{R}$ let $\ell_\lambda = \{(x, y) \in \mathbb{R}^2 \mid x = \lambda\}$ and let $\text{Vert}(\mathbb{R}^2) = \{\ell_\lambda \mid \lambda \in \mathbb{R}\}$.

Definition 2.2.18. A *labeled weighted polygon of complexity $m_f \in \mathbb{Z}_{\geq 0}$* is an element

$$(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f}) \in \text{Polyg}(\mathbb{R}^2) \times (\text{Vert}(\mathbb{R}^2) \times \{-1, +1\} \times \mathbb{Z})^{m_f}$$

with

$$\min_{s \in \Delta} \pi_1(s) < \lambda_1 < \lambda_2 < \dots < \lambda_{m_f} < \max_{s \in \Delta} \pi_1(s)$$

where $\pi_1 : \mathbb{R}^2 \rightarrow \mathbb{R}$ is projection onto the x -coordinate. We denote the space of labeled weighted polygons of complexity m_f by $\mathcal{LWPolyg}_{m_f}(\mathbb{R}^2)$.

Notice there is a triple $(\ell_{\lambda_j}, \epsilon_j, k_j)$ associated with the singular point c_j for each $j = 1, \dots, m_f$. These are related to the critical points of the semitoric system as follows: $\lambda_j = \pi_1(c_j)$; ϵ_j is the cut direction at c_j ; and k_j is the twisting index of c_j .

Here we will briefly review how the affine invariant is produced in [77]. Consider the set $F(M) \subset \mathbb{R}^2$. In the toric case this is the Delzant polygon. Let $c_1, \dots, c_{m_f} \in F(M)$ denote the images of the focus-focus points and let $B_r = \text{Int}(F(M)) \setminus \{c_1, \dots, c_{m_f}\}$ which is precisely the regular values of F [69, Remark 3.2]. For each $j = 1, \dots, m_f$ remove from B_r the line segment $\ell_{\lambda_j}^{\epsilon_j}$ which starts at c_j and goes upwards if $\epsilon_j = 1$ and downwards if $\epsilon_j = -1$ to form the set $B_r^{\vec{\epsilon}}$, where $\vec{\epsilon} = (\epsilon_1, \dots, \epsilon_{m_f})$. Now, $B_r^{\vec{\epsilon}}$ is a simply connected set of regular values of F so define a global toric momentum map

$$F_{\text{toric}} : F^{-1}(B_r^{\vec{\epsilon}}) \rightarrow \mathbb{R}^2$$

and define $\Delta = \overline{F_{\text{toric}}(B_r^{\vec{\epsilon}})}$, the closure. The polygon produced depends on the choice of $(\epsilon_j)_{j=1}^{m_f}$ and of the toric momentum map on $B_r^{\vec{\epsilon}}$. The distinguished points in each polygon are the image of the focus-focus singular points under F_{toric} . Of course, we are omitting many details in this explanation.

Again, the interested reader should see [69, 70].

Let $T \in \mathrm{SL}_2(\mathbb{Z})$ be given by

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \in \mathrm{SL}(2, \mathbb{Z}). \quad (2.2)$$

Definition 2.2.19. Let $\Delta \in \mathrm{Polyg}(\mathbb{R}^2)$ be a rational convex polygon. We say that a *vertex* of Δ is a point in the boundary $\partial\Delta$ where the meeting edges are not co-linear. A point is said to be in the *top-boundary* of Δ if it is the top end of a vertical segment formed by intersecting Δ with a vertical line. Suppose that z is a vertex of Δ and (u, v) are a pair of primitive integral vectors starting at z and extending along the direction of the edges which meet at z in the order such that $\det(u, v) \geq 0$. Then the point z is called

1. a *Delzant corner* when $\det(u, v) = 1$;
2. a *hidden Delzant corner* when it belongs to the top boundary and $\det(u, T^1v) = 1$;
3. a *fake corner* when it belongs to the top boundary and $\det(u, T^1v) = 0$.

The action of $G_{m_f} \times \mathcal{G}$

In order for isomorphic systems to produce the same invariants, we must consider the collection of invariants we have so far modulo a group action.

Notation 2.2.20. Throughout this dissertation when referring to an m_f -tuple such as $(k_j)_{j=1}^{m_f}$ or $(\epsilon_j)_{j=1}^{m_f}$ for simplicity we will sometimes use vector notation. That is, we may refer to these m_f -tuples as \vec{k} and $\vec{\epsilon}$, respectively. These vectors will always have length m_f .

Let $G_{m_f} = \{-1, +1\}^{m_f}$ and $\mathcal{G} = \{(T^t)^k \mid k \in \mathbb{Z}\}$ where T^t is the transpose of the matrix T from Equation (2.2). Given $\ell \in \mathrm{Vert}(\mathbb{R}^2)$ define $t_\ell^k : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by

$$t_\ell^k(x, y) = \begin{cases} (x, y), & x \leq \lambda \\ (x, k(x - \lambda) + y), & x > \lambda. \end{cases}$$

That is, t_ℓ^k acts as the identity on the left of ℓ and, after a translation of coordinates which moves

the origin onto ℓ , acts as $(T^t)^k$ to the right of ℓ . For $\vec{u} \in \mathbb{Z}^{m_f}$ let $t_{\vec{u}} = t_{\ell_1}^{u_1} \circ \dots \circ t_{\ell_{m_f}}^{u_{m_f}}$ where $\ell_j = \ell_{\lambda_j}$.

We define the action of $G_{m_f} \times \mathcal{G}$ on $\mathcal{LWPolyg}_{m_f}(\mathbb{R}^2)$ by

$$((\epsilon'_j)_{j=1}^{m_f}, (T^t)^k) \cdot (\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f}) = (t_{\vec{u}}((T^t)^k \Delta), (\ell_{\lambda_j}, \epsilon'_j, k + k_j)_{j=0}^{m_f}) \quad (2.3)$$

where $\vec{u} = ((\epsilon_j - \epsilon'_j)/2)_{j=1}^{m_f}$.

Remark 2.2.21. Notice that if $(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f})$ is changed via the action of G_{m_f} to have $\epsilon'_j \in \{-1, 1\}$ instead of ϵ_j for each $j = 1, \dots, m_f$ then the new polygon is $t_{\vec{u}}(\Delta)$ where $\vec{u} = (\epsilon_j - \epsilon'_j/2)_{j=1}^{m_f} \in \{-1, 0, 1\}^{m_f}$. Thus, the orbit of Δ under the action of G_{m_f} may be written as $(t_{\vec{u}}(\Delta))_{\vec{u} \in \{0, 1\}^{m_f}}$ if Δ is the polygon with $\epsilon_j = +1$ for all $j = 1, \dots, m_f$. \circlearrowright

The orbit under this action is the appropriate invariant. The choice of cut direction and constant by which to shift the twisting indices parameterize the collection of all polygons in a given orbit. Notice that the action of $t_{\vec{u}}$ does not necessarily preserve convexity, but it will in the case of the polygons we are interested in (Proposition 2.2.24).

Definition 2.2.22. A *labeled Delzant semitoric polygon* is the equivalence class

$$[\Delta_w] \in \mathcal{LWPolyg}_{m_f}(\mathbb{R}^2)/(G_{m_f} \times \mathcal{G})$$

of an element $\Delta_w = (\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})$ satisfying the following.

1. The intersection of Δ and any vertical line is either compact or empty;
2. each ℓ_{λ_j} intersects the top boundary of Δ ;
3. each point in the top boundary which is also in some ℓ_{λ_j} is either a hidden or a fake corner;
4. all other corners are Delzant corners.

The space of labeled Delzant semitoric polygons is denoted by

$$\mathcal{DPolyg}_{m_f}(\mathbb{R}^2) \subset \mathcal{LWPolyg}_{m_f}(\mathbb{R}^2)/(G_{m_f} \times \mathcal{G}).$$

Remark 2.2.23. The twisting index is not a unique integer assigned to each focus-focus singular point because such integers are not preserved under the action of $G_{m_f} \times \mathcal{G}$, but the relative twisting index between two points is preserved. \circlearrowright

Any set satisfying Condition (1) of Definition 2.2.22 is said to have *everywhere finite height*. The following Proposition is a restatement of [70, Lemma 4.2]. Since a preferred representative Δ can be chosen with $\vec{\epsilon} = (+1, \dots, +1)$ we see that it says that the orbit of Δ under G_{m_f} is a subset of $\text{Polyg}(\mathbb{R}^2)$.

Proposition 2.2.24. *Suppose $\Delta_w \in \mathcal{LWPolyg}_{m_f}(\mathbb{R}^2)$ satisfies items (1)-(4) in Definition 2.2.22 and $\Delta_w = [(\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})]$. Then for each $\vec{u} \in \{0, 1\}^{m_f}$ the set $t_{\vec{u}}(\Delta_w)$ is convex.*

The volume invariant

The action of $G_{m_f} \times \mathcal{G}$ can change the vertical position of the images of the focus-focus points, but their height with respect to the bottom of the polygon is preserved.

Definition 2.2.25. Suppose $[\Delta_w] \in \mathcal{DPolyg}_{m_f}(\mathbb{R}^2)$ with associated toric momentum map F . For $j = 1, \dots, m_f$ we define $0 < h_j < \text{length}(\Delta_w \cap \ell_{\lambda_j})$ by

$$h_j = F(m_j) - \min_{s \in \Delta \cap \ell_{\lambda_j}} \{\pi_2(s)\}$$

where $\pi_2 : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is projection onto the second coordinate and $(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f}) \in [\Delta_w]$ is any representative.

This is well defined for any choice of polygon in the same equivalence class by [69, Lemma 5.1]. The word ‘‘volume’’ is used here because h_j can also be viewed as the volume of a specific submanifold of M [69].

The classification theorem

Now that we have defined all of the invariants we can state the result of Pelayo-Vũ Ngọc found in [69, 70].

Definition 2.2.26 (Pelayo-Vũ Ngọc [70]). *A semitoric list of ingredients is*




Figure 2.2: Complete invariant of a semitoric system.

1. a nonnegative integer m_f ;
2. a labeled Delzant semitoric polygon $[\Delta_w] = [(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f})]$ of complexity m_f ;
3. a collection of m_f real numbers $h_1, \dots, h_{m_f} \in \mathbb{R}$ such that $0 < h_j < \text{length}(\pi_2(\Delta \cap \ell_{\lambda_j}))$ for each $j = 1, \dots, m_f$; and
4. a collection of m_f Taylor series $(S_1)^\infty, \dots, (S_{m_f})^\infty \in \mathbb{R}[[X, Y]]_0$.

In other words, a semitoric list of ingredients is a nonnegative integer m_f and an element of $\mathcal{DPolyg}_{m_f}(\mathbb{R}^2) \times \mathbb{R}^{m_f} \times \mathbb{R}[[X, Y]]_0^{m_f}$ where j^{th} element of \mathbb{R} must be in the interval $(0, \text{length}(\pi_2(\Delta \cap \ell_{\lambda_j})))$. Let \mathbb{I} denote the collection of all semitoric lists of ingredients and let \mathbb{I}_{m_f} be lists of ingredients with Ingredient (1) equal to the nonnegative integer m_f .

Notice how the ingredients interact in Definition 2.2.26. Ingredient (1) determines the number of copies of each other ingredient and Ingredient (3) is in an interval determined by Ingredient (2).

Theorem 2.2.27 (Pelayo-Vũ Ngọc [70, Theorem 4.6]). *There exists a bijection between the set of simple semitoric integrable systems modulo semitoric isomorphism and \mathbb{I} , the set of semitoric lists of*

ingredients. In particular,

$$\begin{aligned} \Phi: \mathcal{M}_{\text{ST}} &\rightarrow \mathbb{I} \\ [(M, \omega, (H, J))] &\mapsto [(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f}), (h_j)_{j=1}^{m_f}, ((S_j)^\infty)_{j=1}^{m_f}] \end{aligned}$$

is bijective.

Chapter 3

The space of semitoric systems

3.1 Introduction

In [67] the authors define a metric on the space of Delzant polytopes via the volume of the symmetric difference and pull this back to produce a metric on the moduli space of toric integrable systems, as we described in Section 2.2.6. The construction of this metric is related to the Duistermaat-Heckman measure [23].

The goal of this chapter is to define a metric on the space of semitoric invariants from Theorem 2.2.27 and thus induce a metric on the moduli space of semitoric systems, \mathcal{M}_{ST} , thereby addressing Problem 2.43 from [72], in which the authors ask for a description of the topology of the moduli space of semitoric systems. Problems 2.44 and 2.45 in the same article are related to the closure of \mathcal{M}_{ST} in the moduli space of all integrable systems, so in this chapter I also compute the completion of the space of invariants, which corresponds to the completion of \mathcal{M}_{ST} , in order to lay the foundation to begin work on these problems. The main result of this chapter, Theorem 3.2.12, states that the function I propose is a metric on \mathcal{M}_{ST} and describes the completion of the space of invariants. Theorem 3.2.12 is stated in Section 3.2 after I have defined the metric.

3.2 Construction of metric and statement of main theorem

The invariants of semitoric integrable systems are discussed in Section 2.2.7. To define a metric on \mathcal{M}_{ST} I will first define a metric on each invariant and then I will combine all of these metrics to form a metric on \mathbb{I} . Finally, I will pull this metric back by the map in Theorem 2.2.27 to produce a metric on the space of semitoric systems. This is the same strategy used in [67].

3.2.1 Comparing the Taylor series invariant

First we will define a metric on the Taylor series invariant. For $\sum_{i,j \geq 0} \sigma_{i,j} X^i Y^j \in \mathbb{R}[[X, Y]]_0$ note that the term $\sigma_{0,1}$ should actually be regarded as an element of $S^1 = \mathbb{R}/2\pi\mathbb{Z}$. This can be seen from the construction in [76].

Definition 3.2.1. Suppose that $\{b_n\}_{n=0}^\infty$ is a sequence such that $b_n \in (0, \infty)$ for each $n \in \mathbb{Z}_{\geq 0}$ and $\sum_{n=0}^\infty n b_n < \infty$. We will say that such a sequence is *linear summable*. Now we define

$$d_0^{\{b_n\}_{n=0}^\infty} : \mathbb{R}[[X, Y]]_0 \times \mathbb{R}[[X, Y]]_0 \rightarrow \mathbb{R}$$

to be given by

$$d_0^{\{b_n\}_{n=0}^\infty}(s, s') = \sum_{\substack{i,j \geq 0 \\ (i,j) \neq (0,1)}}^\infty \min \left\{ |\sigma_{i,j} - \sigma'_{i,j}|, b_{i+j} \right\} + \min \left\{ |\sigma_{0,1} - \sigma'_{0,1}|, 2\pi - |\sigma_{0,1} - \sigma'_{0,1}|, b_1 \right\}$$

where $s = \sum_{i,j \geq 0} \sigma_{i,j} X^i Y^j$ and $s' = \sum_{i,j \geq 0} \sigma'_{i,j} X^i Y^j$.

Notice that two series which agree up to a high order will be very close in the metric space and two series which agree only on the high order terms will be distant, as one would expect. In Section 3.3.1 we develop a similar metric on $\mathbb{R}[[X, Y]]$, which could be of independent interest.

Proposition 3.2.2. *For any choice of linear summable sequence $\{b_n\}_{n=0}^\infty$ the space*

$(\mathbb{R}[[X, Y]]_0, d_0^{\{b_n\}_{n=0}^\infty})$ is a complete path-connected metric space and a sequence of Taylor series converges if and only if the coefficient of Y converges in $\mathbb{R}/2\pi\mathbb{Z}$ and all other terms converge in \mathbb{R} . Thus, the topology of $(\mathbb{R}[[X, Y]]_0, d_0^{\{b_n\}_{n=0}^\infty})$ does not depend on the choice of $\{b_n\}_{n=0}^\infty$.

Proposition 3.2.2 follows from the proof of Proposition 3.3.2 in Section 3.3.1.

3.2.2 Comparing the volume invariant

Since the volume invariant h_j is a real number we simply use the standard metric on \mathbb{R} .

3.2.3 Comparing the affine invariant

The topology of spaces of polygons have been studied by many authors. For example, in [41, 42] the authors study polygons with a fixed number of edges up to translations and positive homotheties in Euclidean space and in [48] the authors study polygons in \mathbb{R}^2 with fixed side length up to orientation preserving isometries. For this chapter we will use a topology on polygons related to the Duistermaat-Heckman measure [23] similar to what is done in [67]. A natural way to define a metric on closed subsets of \mathbb{R}^2 is to use the volume of the symmetric difference. Let \ominus denote the symmetric difference of sets. That is, for $A, B \subset \mathbb{R}^2$ let

$$A \ominus B = (A \setminus B) \cup (B \setminus A).$$

In order to define a metric on labeled Delzant semitoric polygons we would like to use the volume of the symmetric difference of the polygons (as is done in [67]) but there are two problems. First, the polygons here are not required to be compact, so the symmetric difference may have infinite volume, and second there are many polygons to choose from. To solve the first problem we will define a non-standard measure on \mathbb{R}^2 . A natural choice would be a probability measure on \mathbb{R}^2 but the structure of $\mathcal{DPolyg}_{m_f}(\mathbb{R}^2)$ is such that vertical translation should not affect the measure. This is because the elements of $\mathcal{DPolyg}_{m_f}(\mathbb{R}^2)$ are only unique up to specific vertical transformations.

Definition 3.2.3. We say that a measure ν on \mathbb{R}^2 is *admissible* if:

1. it is in the same measure class as μ , the Lebesgue measure on \mathbb{R}^2 (i.e. $\mu \ll \nu$ and $\nu \ll \mu$);
2. its Radon-Nikodym derivative with respect to Lebesgue measure only depends on the x -coordinate, i.e. there exists a $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $d\nu/d\mu(x, y) = g(x)$ for all $(x, y) \in \mathbb{R}^2$;
3. this function g satisfies $xg \in L^1(\mu, \mathbb{R})$ and g is bounded and bounded away from zero on any compact interval.

Equation (3.1), in Section 3.3.2, is an example of such a measure. When only considering compact semitoric systems one can use the Lebesgue measure on \mathbb{R}^2 instead to produce a metric which induces the same topology, see Remark 3.3.16.

We say that a measurable map $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a *vertical transformation* if it is of the form $T(x, y) = (x, y + f(x))$ for some $f : \mathbb{R} \rightarrow \mathbb{R}$. Part (2) of Definition 3.2.3 implies that the measure is invariant under vertical transformations and part (3) will force convex sets which have a finite height at every x -value to have finite measure.

Proposition 3.2.4. *Suppose that ν is an admissible measure on \mathbb{R}^2 and $\Delta \in \text{Polyg}(\mathbb{R}^2)$. Then Δ has everywhere finite height if and only if $\nu(\Delta) < \infty$.*

Proposition 3.2.4 is proven in Section 3.3.2. Let S^{m_f} denote the symmetric group on m_f elements. For $p \in S^{m_f}$ let the action of p on a vector $\vec{v} = (v_j)_{j=1}^{m_f}$ by permuting the elements be denoted by $p(\vec{v}) = (v_{p(j)})_{j=1}^{m_f}$.

Definition 3.2.5. Suppose $\vec{k}, \vec{k}' \in \mathbb{Z}^{m_f}$ for some nonnegative integer m_f . Then we say $\vec{k} \sim \vec{k}'$ if there exists a constant $c \in \mathbb{Z}$ and a permutation $p \in S^{m_f}$ such that $k_j = k'_{p(j)} + c$ for all $j = 1, \dots, m_f$. We denote by $[\vec{k}]$ the equivalence class of \vec{k} in \mathbb{Z}^{m_f} / \sim .

Definition 3.2.6. Fix any $\vec{k}, \vec{k}' \in \mathbb{Z}^{m_f}$ such that $\vec{k} \sim \vec{k}'$. Let

$$S_{\vec{k}, \vec{k}'}^{m_f} = \left\{ p \in S^{m_f} \left| \begin{array}{l} \text{there exists } c \in \mathbb{Z} \text{ such that} \\ k_j = k'_{p(j)} + c \text{ for all } j = 1, \dots, m_f \end{array} \right. \right\}.$$

Notice that $\vec{k} \sim \vec{k}'$ is equivalent to $S_{\vec{k}, \vec{k}'}^{m_f} \neq \emptyset$. The elements of $S_{\vec{k}, \vec{k}'}^{m_f}$ will be called *appropriate permutations* for \vec{k} and \vec{k}' .

Now, assume that two labeled weighted polygons have the same number of focus-focus points and twisting indices related by \sim . We can shift the twisting index of one of the labeled weighted polygons by the action of an element of \mathcal{G} such that after the shift the two labeled weighted polygons in question will have the same twisting index modulo the ordering. Once the twisting indices are fixed we still have a family of polygons which depends on the choice of $\vec{\epsilon} \in \{-1, +1\}^{m_f}$. The number of possible choices of $\vec{\epsilon}$ is finite so we will simply sum up the symmetric difference of each pair of polygons for each choice of $\vec{\epsilon}$. Using Remark 2.2.21 we can concisely write this in the following way.

Definition 3.2.7. Suppose that for $i = 1, 2$ we have $[\Delta_w]_i = [(\Delta^i, (\ell_{\lambda_j}^i, +1, k_j^i)_{j=1}^{m_f})] \in \mathcal{DPolyg}_{m_f}(\mathbb{R}^2)$ for some $m_f > 0$ and with $\vec{k} \sim \vec{k}'$, so $\mathcal{S}_{\vec{k}, \vec{k}'}^{m_f} \neq \emptyset$. For $p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$ define

$$d_{\mathbb{P}}^{p, \nu}([\Delta_w]_1, [\Delta_w]_2) = \sum_{\vec{u} \in \{0, 1\}^{m_f}} \nu(t_{\vec{u}}(\Delta) \ominus t_{p(\vec{u})}((T^t)^{-c}(\Delta')))$$

where $c \in \mathbb{Z}$ is the unique integer such that $k_j - k'_{p(j)} = c$ for all $j = 1, \dots, m_f$. In the case that $[(\Delta)], [(\Delta')] \in \mathcal{DPolyg}_0(\mathbb{R}^2)$ define

$$d_{\mathbb{P}}^{\nu}([(\Delta)], [(\Delta')]) = \nu(\Delta \ominus \Delta').$$

If $m_f = 0$ the labeled weighted polygon becomes a single polygon. The definition of $d_{\mathbb{P}}^{p, \nu}$ in this case should be thought of as the same formula as the $m_f > 0$ case and it is only treated separately because the sum in the more general formula would be empty if $m_f = 0$.

Notice that $d_{\mathbb{P}}^{p, \nu}$ is not a metric if $p \neq p^{-1}$ in \mathcal{S}^{m_f} because it is not symmetric. We will remove the dependence on a choice of permutation in the next section when we define the final version of the metric. There are many ways to choose a representative from each equivalence class which have matching twisting indices, but the volume of the symmetric difference will not actually depend on that choice (see Proposition 3.3.5) so this function is well-defined on orbits of $G_{m_f} \times \mathcal{G}$.

3.2.4 Definition of metric and main result

We assume that systems with a different number of singular points are in different components of \mathcal{M}_{ST} . Additionally, since the invariant \vec{k} is discrete one might expect that different values of \vec{k} would not be comparable; this is not correct. If $\vec{k} \sim \vec{k}'$ then systems with these twisting indices can be compared via the metric we are about to define but they are in different connected components (Remark 3.3.19).

Definition 3.2.8. Suppose that $m_f \in \mathbb{Z}_{>0}$ and $\vec{k} \in \mathbb{Z}^{m_f}$. Then we define $\mathcal{M}_{\text{ST}, m_f, \vec{k}} \subset \mathcal{M}_{\text{ST}, m_f}$ to be those elements with twisting index exactly \vec{k} and define

$$\mathcal{M}_{\text{ST}, m_f, [\vec{k}]} = \bigcup_{\vec{k}' \in [\vec{k}]} \mathcal{M}_{\text{ST}, m_f, \vec{k}'}$$

Furthermore, define

$$\mathcal{DPolyg}_{m_f, [\vec{k}]}(\mathbb{R}^2) = \{(\Delta, (\ell_{\lambda_j}, \epsilon_j, k'_j)_{j=1}^{m_f}) \in \mathcal{DPolyg}_{m_f}(\mathbb{R}^2) \mid \vec{k} \sim \vec{k}'\}$$

and

$$\mathbb{I}_{m_f, [\vec{k}]} = \mathbb{I}_{m_f} \cap \left(\mathcal{DPolyg}_{m_f, [\vec{k}]}(\mathbb{R}^2) \times \mathbb{R}^{m_f} \times \mathbb{R}[[X, Y]]_0^{m_f} \right).$$

Notice that

$$\mathcal{M}_{\text{ST}} = \bigcup_{\substack{m_f \in \mathbb{Z}_{\geq 0} \\ \vec{k} \in \mathbb{Z}^{m_f}}} \mathcal{M}_{\text{ST}, m_f, \vec{k}}.$$

This union, and the union in Definition 3.2.8, are not disjoint unions only because they have repeated terms. For instance, since the action of \mathcal{G} can shift all of the twisting indices, we have that

$$\mathcal{M}_{\text{ST}, m_f, \vec{k}} = \mathcal{M}_{\text{ST}, m_f, (k_j + c)_{j=1}^{m_f}}$$

for any $c \in \mathbb{Z}$.

From Sections 3.2.1, 3.2.2, and 3.2.3, given some fixed appropriate permutation we already know how to define a “distance” function on two systems with specified twisting index. To produce a metric which does not depend on fixing a permutation we will take the minimum of each possibility.

Definition 3.2.9. Let $m_f \in \mathbb{Z}_{\geq 0}$ and $\vec{k} \in \mathbb{Z}^{m_f}$ and suppose that $m, m' \in \mathbb{I}_{m_f, [\vec{k}]}$ with $m = ([\Delta_w], (h_j)_{j=1}^{m_f}, ((S_j)^\infty)_{j=1}^{m_f})$ and $m' = ([\Delta'_w], (h'_j)_{j=1}^{m_f}, ((S'_j)^\infty)_{j=1}^{m_f})$. Let ν be an admissible measure, $\{b_n\}_{n=0}^\infty$ be a linear summable sequence, and $p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$. We define:

1. the *comparison with alignment* p to be

$$d_{m_f, [\vec{k}]}^{p, \nu, \{b_n\}_{n=0}^\infty}(m, m') = d_{\text{P}}^{p, \nu}([\Delta_w], [\Delta'_w]) + \sum_{j=1}^{m_f} \left(d_0^{\{b_n\}_{n=0}^\infty}((S_j)^\infty, (S'_{p(j)})^\infty) + |h_j - h'_{p(j)}| \right);$$

2. the *distance between m and m'* to be

$$d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^\infty}(m, m') = \min_{p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}} \left\{ d_{m_f, [\vec{k}]}^{p, \nu, \{b_n\}_{n=0}^\infty}(m, m') \right\}.$$

A minimum of even a finite number of metrics is not a metric in general, but we will see in Theorem 3.2.12 that $d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^{\infty}}$ is a metric in this case. Now we use this distance defined on each component to induce a distance on the whole space which can be pulled back to produce a metric on \mathcal{M}_{ST} .

Definition 3.2.10. Let ν be an admissible measure and $\{b_n\}_{n=0}^{\infty}$ be a linear summable sequence. Then we define

1. the *distance on \mathbb{I}* by

$$d^{\nu, \{b_n\}_{n=0}^{\infty}}(m, m') = \begin{cases} d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^{\infty}}(m, m') & , \text{ if } m, m' \in \mathbb{I}_{m_f, [\vec{k}]} \text{ for some } m_f \in \mathbb{Z}, \vec{k} \in \mathbb{Z}^{m_f} \\ 1 & , \text{ otherwise} \end{cases}$$

for $m, m' \in \mathbb{I}$;

2. the *distance on \mathcal{M}_{ST}* by $\mathcal{D}^{\nu, \{b_n\}_{n=0}^{\infty}} = \Phi^* d^{\nu, \{b_n\}_{n=0}^{\infty}}$ where $\Phi : \mathcal{M}_{\text{ST}} \rightarrow \mathbb{I}$ is the bijective correspondence from Theorem 2.2.27.

To state the main theorem we will have to first define the completion.

Definition 3.2.11. For any choice of $m_f \in \mathbb{Z}_{\geq 0}$ and $\vec{k} \in \mathbb{Z}^{m_f}$ we define

$$\tilde{\mathbb{I}}_{m_f, [\vec{k}]} = \widetilde{\mathcal{DPoly}}_{m_f, [\vec{k}]}(\mathbb{R}^2) \times [0, 1]^{m_f} \times \mathbb{R}[[X, Y]]_0^{m_f}$$

and

$$\tilde{\mathbb{I}} = \bigcup_{\substack{m_f \in \mathbb{Z}_{\geq 0} \\ \vec{k} \in \mathbb{Z}^{m_f}}} \tilde{\mathbb{I}}_{m_f, [\vec{k}]}$$

where the critical points satisfy the ordering convention from Remark 3.4.7 and $\widetilde{\mathcal{DPoly}}_{m_f, [\vec{k}]}(\mathbb{R}^2)$ is defined as in Definition 3.4.9.

Theorem 3.2.12. For any choice of

1. a linear summable sequence $\{b_n\}_{n=0}^{\infty}$;
2. an admissible measure ν ;

the space $(\mathcal{M}_{\text{ST}}, \mathcal{D}^{\nu, \{b_n\}_{n=0}^{\infty}})$ is a non-complete metric space whose completion corresponds to $\widetilde{\mathbb{I}}$. Moreover, the topology of $(\mathcal{M}_{\text{ST}}, \mathcal{D}^{\nu, \{b_n\}_{n=0}^{\infty}})$ is independent of the choice of ν and $\{b_n\}_{n=0}^{\infty}$.

Remark 3.2.13. There are several important facts to notice about Theorem 3.2.12:

1. This distance induces a unique topology on \mathcal{M}_{ST} and thus Theorem 3.2.12 completely resolves Problem 2.43 from [72].
2. In special cases a less complicated form of the metric can be used. The metric

$$\mathcal{D}^{\text{Id}} = \Phi^* d_{m_f, \tilde{k}}^{\text{Id}, \nu, \{b_n\}_{n=0}^{\infty}}$$

is easier to work with and induces the same topology as $\mathcal{D}^{\nu, \{b_n\}_{n=0}^{\infty}}$ (Proposition 3.3.18) so this should be used to study topological properties of \mathcal{M}_{ST} . Additionally, when studying compact semitoric systems the admissible measure on \mathbb{R}^2 can be instead replaced by the standard Lebesgue measure without changing the topology (Remark 3.3.16). See Example 3.4.16 for an explanation of why $\mathcal{D}^{\nu, \{b_n\}_{n=0}^{\infty}}$ produces the appropriate metric space structure on \mathcal{M}_{ST} .

3. Since toric integrable systems fall into the broader category of semitoric systems it is natural to wonder if the metric defined in this chapter is compatible with the metric on toric systems from [67]. Because we must choose an admissible measure to apply to the more general cases the metric induced by d does not exactly match the metric defined on toric systems but they do induce the same topology, see Section 3.3.5.
4. Since all metric spaces are Tychonoff (completely regular and Hausdorff) we know that \mathcal{M}_{ST} is Tychonoff. Thus the Stone-Cěch compactification [78, 75] applies to \mathcal{M}_{ST} so it admits a Hausdorff compactification (just as in [67]).

◻

3.3 The metric

In this section we fill in the details of constructing the metric and prove that it is a metric.

3.3.1 Metrics on Taylor series

Let $\mathbb{R}[[X, Y]]$ refer to the algebra of real formal power series in two variables, X and Y .

Definition 3.3.1. Suppose that $\{b_n\}_{n=0}^\infty$ is any linear summable sequence. Then we define the *distance on Taylor series* to be the function

$$d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty} : \mathbb{R}[[X, Y]] \times \mathbb{R}[[X, Y]] \rightarrow \mathbb{R}$$

given by

$$d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty} \left(\sum_{i,j \geq 0} \sigma_{i,j} X^i Y^j, \sum_{i,j \geq 0} \sigma'_{i,j} X^i Y^j \right) = \sum_{i,j=0}^\infty \min \{ |\sigma_{i,j} - \sigma'_{i,j}|, b_{i+j} \}.$$

Proposition 3.3.2. *The space $(\mathbb{R}[[X, Y]], d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty})$ is a complete path-connected metric space and a sequence of Taylor series $(s_k = \sum_{i,j \geq 0} \sigma_{i,j}^k X^i Y^j)_k$ converges if and only if each sequence of terms $(\sigma_{i,j}^k)_k$ converges.*

Proof. First notice that the sum in the definition of the distance always converges. This is because

$$d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty} \left(\sum_{i,j \geq 0} \sigma_{i,j} X^i Y^j, \sum_{i,j \geq 0} \sigma'_{i,j} X^i Y^j \right) \leq \sum_{i,j=0}^\infty b_{i+j} = \sum_{n=0}^\infty (n+1)b_n < \infty$$

for any pair of Taylor series by the choice of $\{b_n\}_{n=0}^\infty$. It is also clear that $d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty}$ is symmetric and positive definite. It satisfies the triangle inequality because that inequality is satisfied for each term and thus we can see that $(\mathbb{R}[[X, Y]], d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty})$ is a metric space.

Next we will prove the condition on convergence. Suppose that

$$\lim_{k \rightarrow \infty} d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty} (s_k, s_0) = 0$$

with $s_k, s_0 \in \mathbb{R}[[X, Y]]$ as in the statement of the Proposition. Fix any $I, J \in \mathbb{Z}_{\geq 0}$ and we will show that $\sigma_{I,J}^k \xrightarrow{k \rightarrow \infty} \sigma_{I,J}^0$. Fix $\varepsilon > 0$ and find K such that $k > K$ implies that

$$\sum_{i,j=0}^\infty \min \{ |\sigma_{i,j}^k - \sigma_{i,j}^0|, b_{i+j} \} < \varepsilon$$

because we may assume that $\varepsilon < b_{I+J}$. Then we can see that $|\sigma_{I,J}^k - \sigma_{I,J}^0| < \varepsilon$ so the result follows.

Now we will show the converse. Suppose that

$$\lim_{k \rightarrow \infty} |\sigma_{i,j}^k - \sigma_{i,j}^0| = 0$$

for all $i, j \in \mathbb{Z}_{\geq 0}$. Fix $\varepsilon > 0$, let $N \in \mathbb{N}$ be such that

$$\sum_{n \geq N} (n+1)b_n < \varepsilon/2,$$

and let $K \in \mathbb{Z}$ be such that $k > K$ implies that

$$|\sigma_{i,j}^k - \sigma_{i,j}^0| < \frac{\varepsilon}{N(N+1)}$$

for each $i, j \in \mathbb{Z}_{\geq 0}$ such that $i+j < N$. Notice it is possible to do this simultaneously because there are only finitely many such pairs (i, j) . For any $k > K$ we have that

$$\begin{aligned} d_{\mathbb{R}[[X,Y]]}^{\{b_n\}_{n=0}^{\infty}}(s_k, s_0) &\leq \sum_{i+j < N} |\sigma_{i,j}^k - \sigma_{i,j}^0| + \sum_{i+j \geq N} b_{i+j} \\ &< \frac{\varepsilon}{N(N+1)} \sum_{i+k < N} 1 + \sum_{n \geq N} (n+1)b_n \\ &< \frac{\varepsilon}{N(N+1)} \frac{N(N+1)}{2} + \frac{\varepsilon}{2} = \varepsilon. \end{aligned}$$

This proves the convergence condition.

Any element of this space may be continuously transformed into any other linearly in each term, so it is path-connected. To finish the proof we will show that this space is complete. Suppose that $(s_k)_{k=0}^{\infty}$ is a Cauchy sequence in $\mathbb{R}[[X, Y]]$. Using an argument similar to the one for convergence, we can see that the sequence $\{\sigma_{i,j}^k\}_{k=0}^{\infty}$ is Cauchy for each $i, j \in \mathbb{Z}_{\geq 0}$ and therefore $\sigma_{i,j}^k \xrightarrow{k \rightarrow \infty} \sigma_{i,j}^0$ for some $\sigma_{i,j}^0 \in \mathbb{R}$. Since it converges in each term, we can use the convergence condition to conclude that

$$\lim_{k \rightarrow \infty} d_{\mathbb{R}[[X,Y]]}^{\{b_n\}_{n=0}^{\infty}} \left(s_k, \sum_{i,j \geq 0} \sigma_{i,j}^0 X^i Y^j \right) = 0$$

and so all Cauchy sequences have limits. \square

We have characterized convergence in this space in a way which is independent of the sequence $\{b_n\}_{n=0}^\infty$. Since the topology of a metrizable space is completely determined by its convergent sequences we have the following result.

Corollary 3.3.3. *The topology on $\mathbb{R}[[X, Y]]$ determined by $d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty}$ does not depend on the choice of linear summable sequence $\{b_n\}_{n=0}^\infty$.*

Notice that $\mathbb{R}[[X, Y]]_0$ is not a closed subset of $(\mathbb{R}[[X, Y]], d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty})$ and $(\mathbb{R}[[X, Y]]_0, d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty})$ with the restricted metric is not a complete metric space. To see this consider any collection of Taylor series in which $\sigma_2 \rightarrow 2\pi$. This does not accurately describe the structure of the semitoric systems and thus we use the altered metric from Definition 3.2.1. Proposition 3.2.2 follows from a slightly altered version of the proof of Proposition 3.3.2.

Remark 3.3.4. A similar construction to $d_{\mathbb{R}[[X, Y]]}^{\{b_n\}_{n=0}^\infty}$ can be used to produce such a metric on Taylor series in any number of variables. The only difference is that to produce a metric on Taylor series in m variables the sequence $\{b_n\}_{n=0}^\infty$ would be required to satisfy

$$\sum_{n=0}^{\infty} \binom{n+m-1}{n} b_n < \infty$$

because there are $\binom{n+m-1}{n}$ terms of degree n in a Taylor series on m variables. ◊

3.3.2 Metrics on labeled weighted polygons

We start this section with a proof.

Proof of Proposition 3.2.4. By definition Δ is the intersection of half-spaces and since it is assumed to have everywhere finite height we can see that this collection of half spaces must include at least two which are not completely vertical, i.e. not of the form $\{x \geq c\}$ or $\{x \leq c\}$ for $c \in \mathbb{R}$. Let B denote the intersection of these two half planes. Then by definition $\Delta \subset B$ and thus $\nu(\Delta) < \nu(B)$. If the two half planes are parallel of a distance c apart then

$$\nu(B) = \int_B \frac{d\nu}{d\mu} d\mu = \int_{\mathbb{R}} cg d\mu < \infty$$

because $xg \in L^1(\mu, \mathbb{R})$ implies that $g \in L^1(\mu, \mathbb{R})$. If the spaces are not parallel then their boundaries intersect at some point (x_0, y_0) . Let m be the absolute value of the difference in the slopes of the two boundaries. Then for each value $(x, y) \in \mathbb{R}^2$ the height of B at that x -coordinate is $m|x - x_0|$ and the sign of $x - x_0$ is the same for each $(x, y) \in B$. Assume that $x - x_0 \geq 0$ for all $(x, y) \in B$ so we have

$$\nu(B) = \int_B \frac{d\nu}{d\mu} d\mu = \int_{x_0}^{\infty} m(x - x_0)g(x) d\mu = m \int_{x_0}^{\infty} xg d\mu - mx_0 \int_{x_0}^{\infty} g d\mu < \infty$$

because $g \in L^1(\mu, \mathbb{R})$ and $xg \in L^1(\mu, \mathbb{R})$. The computation is similar if $x - x_0 \leq 0$ for all $(x, y) \in B$.

Any compact set without everywhere finite height will have infinite ν -measure. This is because a compact set which does not have everywhere finite height either is a vertical line, which is not a polygon, or includes a subset of the form $\{(x, y) \mid a_1 < x < a_2\}$ for some $a_1 < a_2$. Such a subset has infinite ν -measure because ν is invariant under vertical translations. \square

Even once we have fixed the cut directions there are many polygons to choose from based on the choice of the twisting index (i.e. the orbit of the action of \mathcal{G}) but if the same choice is made for each pair of polygons this choice does not change the volume of the symmetric difference.

Proposition 3.3.5. *Let $m_f \in \mathbb{Z}_{\geq 0}$, $p \in \mathcal{S}^{m_f}$, and let $\mathcal{P}^p \subset (\mathcal{D}Poly_{m_f, [\vec{k}]}(\mathbb{R}^2))^2$ be given by*

$$\mathcal{P}^p = \left\{ \left([(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f})], [(\Delta', (\ell_{\lambda'_j}, \epsilon'_j, k'_j)_{j=1}^{m_f})] \right) \in (\mathcal{D}Poly_{m_f, [\vec{k}]}(\mathbb{R}^2))^2 \mid p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f} \right\}.$$

Then the function $d_P^{p, \nu} : \mathcal{P}^p \rightarrow \mathbb{R}$ is well defined.

Proof. Suppose that

$$\Delta_w^1 = (\Delta^1, (\ell_{\lambda_j}^1, +1, k_j^1)_{j=1}^{m_f}), \Delta_w^2 = (\Delta^2, (\ell_{\lambda_j}^2, +1, k_j^2)_{j=1}^{m_f}) \in [(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f})]$$

and $\Delta'_w = (\Delta', (\ell_{\lambda'_j}, +1, k'_j)_{j=1}^{m_f})$. Then there exists some $d \in \mathbb{Z}$ such that $k_j^1 = k_j^2 - d$ for $j = 1, \dots, m_f$ and $\Delta^1 = T^d(\Delta^2)$. Since $p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$, there exists $c \in \mathbb{Z}$ such that $k_j^1 - k'_{p(j)} = c$ for all j and notice

that this means that $k_j^2 - k'_{p(j)} = c + d$ for all j . Therefore,

$$\begin{aligned}
d_{\mathbb{P}}^{p,\nu}([\Delta_w^1], [\Delta_w']) &= \sum_{\bar{u} \in \{0,1\}^{mf}} \nu(t_{\bar{u}}(\Delta^1) \ominus t_{p(\bar{u})}(T^{-c}(\Delta'))) \\
&= \sum_{\bar{u} \in \{0,1\}^{mf}} \nu(t_{\bar{u}}(T^d(\Delta^2)) \ominus t_{p(\bar{u})}(T^{-c}(\Delta'))) \\
&= \sum_{\bar{u} \in \{0,1\}^{mf}} \nu(T^d(t_{\bar{u}}(\Delta^2) \ominus t_{p(\bar{u})}(T^{-c-d}(\Delta')))) \\
&= \sum_{\bar{u} \in \{0,1\}^{mf}} \nu(t_{\bar{u}}(\Delta^2) \ominus t_{p(\bar{u})}(T^{-(c+d)}(\Delta'))) \\
&= d_{\mathbb{P}}^{p,\nu}([\Delta_w^2], [\Delta_w'])
\end{aligned}$$

because admissible measures are invariant under vertical transformations such as T^d . The argument that this function is well defined in the second input is similar. \square

An example of an admissible measure on \mathbb{R}^2 is the following. Define ν_0 so that

$$\frac{d\nu_0}{d\mu}(x, y) = \begin{cases} 1 & , \text{ if } |x| < 1 \\ \frac{1}{x^3} & , \text{ else.} \end{cases} \quad (3.1)$$

Notice $x^{-2} \in L^1(\mu, \mathbb{R})$ and $g_0 = \frac{d\nu_0}{d\mu}(x, 0)$ is bounded and bounded away from zero on compact intervals. This proves the following.

Proposition 3.3.6. *The measure ν_0 is an admissible measure on \mathbb{R}^2 .*

3.3.3 Choice of ν does not change the topology

While the choice of admissible measure will change the metric it does not change the topology induced by that metric.

Lemma 3.3.7. *Suppose that ν is an admissible measure and $\Delta_k, \Delta \in \text{Polyg}(\mathbb{R}^2)$ for $k \in \mathbb{N}$ are such that $\nu(\Delta_k \ominus \Delta) \xrightarrow{k \rightarrow \infty} 0$. Then there exists a proper vertical segment $A = \{x_0\} \times [y_0, y_1]$, with $y_0 < y_1$, and $K > 0$ such that $A \subset \Delta_k \cap \Delta$ for all $k > K$.*

Proof. Fix any $N > 0$ such that $\Delta \cap ([-N, N] \times \mathbb{R})$ has non-zero measure with respect to ν , and thus also with respect to μ . Since ν is admissible we can find some $c > 0$ such that $d\nu/d\mu > c$ on

$[-N, N] \times \mathbb{R}$.

For each $\varepsilon > 0$ let

$$U_\varepsilon = \{p \in \mathbb{R}^2 \mid B_\varepsilon(p) \subset (-N, N) \times \mathbb{R}\} \cap \Delta$$

where $B_\varepsilon(p)$ is the standard open ball of radius ε centered at p and $\text{int}(A)$ denotes the interior of the set A .

Fix any $k \in \mathbb{N}$ and suppose $U_\varepsilon \setminus \Delta_k \neq \emptyset$. Because Δ_k is the intersection of closed half-planes its complement, Δ_k^c , is the union of open half-planes. If $q \in U_\varepsilon \setminus \Delta_k$ then there exists some open half-plane with boundary including q which is a subset of Δ_k^c . Let H_q be the intersection of one such half-plane with $B_\varepsilon(q)$ so $H_q \subset \Delta \setminus \Delta_k$. Then, since $H_q \subset (-N, N) \times \mathbb{R}$,

$$\nu(H_p) = \int_{H_p} \frac{d\nu}{d\mu} d\mu > c\mu(H_p) = \frac{c}{2}\mu(B_\varepsilon(p)) = \frac{c\pi}{2}\varepsilon^2.$$

Thus, if $U_\varepsilon \setminus \Delta_k$ is non-empty then $\nu(\Delta_k \ominus \Delta) > \frac{c\pi}{2}\varepsilon^2$.

Now choose ε small enough that U_ε is non-empty and choose $K > 0$ such that $k > K$ implies that $\nu(\Delta \ominus \Delta_k) < \frac{c\pi}{2}\varepsilon^2$. If $U_\varepsilon \setminus \Delta_k \neq \emptyset$ then $\nu(\Delta_k \ominus \Delta) > \frac{c\pi}{2}\varepsilon^2$, so we conclude $U_\varepsilon \subset \Delta_k$ for $k > K$. The set U_ε has nonempty interior so we can find the set A as in the statement of the Lemma. \square

Now we will use Lemma 3.3.7 to prove Lemma 3.3.8, which says that the same sequences of polygons converge with respect to any admissible measure.

Lemma 3.3.8. *Suppose that ν_1, ν_2 are admissible measures and that $\Delta_k, \Delta \in \text{Polyg}(\mathbb{R}^2)$ for $k \in \mathbb{N}$ have $\nu_1(\Delta), \nu_1(\Delta_k) < \infty$. If $\nu_1(\Delta_k \ominus \Delta) \xrightarrow{k \rightarrow \infty} 0$ then $\nu_2(\Delta_k \ominus \Delta) \xrightarrow{k \rightarrow \infty} 0$.*

Proof. Suppose that $\nu_1(\Delta_k \ominus \Delta) \xrightarrow{k \rightarrow \infty} 0$ and let A, x_0, y_0 , and y_1 be as in Lemma 3.3.7. We know that the line $\{x = x_0\}$ intersects Δ so it must intersect the top boundary of Δ , since Δ has everywhere finite height by Proposition 3.2.4. Since a convex set is the intersection of half-planes there must exist a line ℓ_1 which goes through the point where $\{x = x_0\}$ intersects the top boundary such that all of Δ is in a closed half-plane bounded by ℓ_1 (as in Figure 3.1). Such a line may not be unique if there is a vertex with x -coordinate equal to x_0 , but any choice of such a line will do.




Figure 3.1: Since Δ is convex it must all lie on the same side of ℓ_1 .

The situation we describe next is shown in Figure 3.2. Let m denote the slope of ℓ_1 and let ℓ_2 be the line through (x_0, y_1) with slope $m + 1$. Let m' denote the slope of the line through the point (x_0, y_0) and the point which is the intersection of ℓ_1 with ℓ_2 . Finally let ℓ_3 be the line through (x_0, y_0) with slope $(m+m')/2$. Since the slope of ℓ_3 is greater than the slope of ℓ_2 these two lines must intersect at some x -coordinate greater than x_0 , but since the slope of ℓ_3 is less than m' we know that the intersection of ℓ_2 and ℓ_3 must be to the right of the intersection of ℓ_1 and ℓ_2 . Thus the lines ℓ_1, ℓ_2 , and ℓ_3 bound a triangle which we will denote by G , as is shown in Figure 3.2. Let $N_1 = \max_{s \in G} \pi_1(s)$. Since Δ is on one side of ℓ_1 and G is on the other we conclude that $G \cap \Delta = \emptyset$.

For any $N > x_0$ let E_N^1 denote the region of \mathbb{R}^2 which has $x > N$ and is above or on ℓ_2 . Now suppose that k is large enough so that $A \subset \Delta_k$ and let $p \in E_{N_1}^1 \cap \ell_2$. Then $p \in \Delta_k$ implies that $G \subset \Delta_k \ominus \Delta$ because Δ_k is convex and $\Delta \cap G = \emptyset$. Similarly, if p is any other point in $E_{N_1}^1$ we can conclude that some ν_1 -preserving transformation of G must be contained in $\Delta_k \ominus \Delta$. This is because moving p vertically will result in acting on G by some matrix T^r (as in Equation (2.2)) with $r \in \mathbb{R}$ with origin on the line $\{x = x_0\}$ (see Figure 3.3). In any case, if $\Delta_k \cap E_{N_1}^1$ is nonempty




Figure 3.2: The lines ℓ_1, ℓ_2, ℓ_3 , and triangle G .




Figure 3.3: Notice that for a fixed vertical line segment $A \subset \mathbb{R}^2$ the measure of the convex hull of A and $p \in \mathbb{R}^2$ only depends on the x -component of p . This is because if $p_1, p_2 \in \mathbb{R}^2$ with $\pi_1(p_1) = \pi_1(p_2)$ then the convex hulls are related by a vertical transformation.

and k is large enough so that $A \subset \Delta_k$ then we can conclude that $\nu_1(\Delta \ominus \Delta_k) \geq \nu_1(G) > 0$. Since $\nu_1(\Delta \ominus \Delta_k) \xrightarrow{k \rightarrow \infty} 0$ we can conclude that for large enough k the set $\Delta_k \cap E_N^1$ is empty.

Using a similar argument, one can define sets E_N^i for $i = 2, 3, 4$ that must also be disjoint from Δ_k for large enough k and N ; these are shown in Figure 3.4. The sets E_N^1 and E_N^2 are bounded to the left by the line $\{x = N\}$ and the sets E_N^3 and E_N^4 are bounded to the right by $\{x = -N\}$. The sets E_N^1 and E_N^4 are bounded below by lines and the sets E_N^2 and E_N^3 are bounded above by lines. Let $E_N = \cup_{i=1}^4 E_N^i$ and let $N_2 > N_1$ be large enough so that for large enough k we have that $\Delta_k \cap E_{N_2} = \emptyset$. Let $D_N = [-N, N] \times \mathbb{R}$ for $N \in \mathbb{R}$ and let $S_N = \mathbb{R}^2 \setminus (E_N \cup D_N)$.

Fix $\varepsilon > 0$. Notice that for each $N > 0$ the set S_N is of finite ν_2 -measure. Since $\{S_N\}_{N>0}$ are nested we conclude that $\lim_{N \rightarrow \infty} \nu_2(S_N) = 0$. Now choose some fixed $N_3 > N_2$ and $K_1 > 0$ such that $\nu_2(S_{N_3}) < \varepsilon$ and $k > K_1$ implies that $\Delta_k \cap E_{N_3} = \emptyset$. Since both ν_1 and ν_2 are admissible measures we know that their Radon-Nikodym derivative is bounded on D_{N_3} . This is because

$$\frac{d\nu_2}{d\nu_1} = \frac{d\nu_2}{d\mu} \left(\frac{d\nu_1}{d\mu} \right)^{-1},$$

which are both bounded on D_{N_3} . Let $c > 0$ be such that $d\nu_2/d\nu_1 < c$ on D_{N_3} . Now choose $K_2 > K_1$




Figure 3.4: For large choices of N and k the set S_N is small and the set E_N has empty intersection with Δ_k . Then we can concentrate on the set D_N , on which the Radon-Nikodym derivative $d\nu_2/d\nu_1$ is bounded.

such that $k > K_2$ implies $\nu_1(\Delta \ominus \Delta_k) < \varepsilon$. Finally, for $k > K_2$ we have

$$\begin{aligned}
\nu_2(\Delta_k \ominus \Delta) &= \int_{\mathbb{R}^2} |\chi_{\Delta_k} - \chi_{\Delta}| \, d\nu_2 \\
&= \int_{S_{N_3}} |\chi_{\Delta_k} - \chi_{\Delta}| \, d\nu_2 + \int_{E_{N_3}} |\chi_{\Delta_k} - \chi_{\Delta}| \, d\nu_2 + \int_{D_{N_3}} |\chi_{\Delta_k} - \chi_{\Delta}| \, d\nu_2 \\
&\leq \nu_2(S_{N_3}) + 0 + \int_{D_{N_3}} |\chi_{\Delta_k} - \chi_{\Delta}| \frac{d\nu_2}{d\nu_1} \, d\nu_1 \\
&< \varepsilon + c\nu_1(\Delta_k \ominus \Delta) \\
&< (1 + c)\varepsilon,
\end{aligned}$$

which can be made arbitrarily small. \square

By combining Lemma 3.3.8 and Proposition 3.2.2 we have the following corollary.

Corollary 3.3.9. *Fix a nonnegative integer $m_f \in \mathbb{Z}_{\geq 0}$, a vector $\vec{k} \in \mathbb{Z}^{m_f}$, any two linearly summable sequences $\{b_n\}_{n=0}^{\infty}$ and $\{b'_n\}_{n=0}^{\infty}$, and two admissible measures ν and ν' . Then the metric spaces $(\mathbb{I}_{m_f, [\vec{k}]}, d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^{\infty}})$ and $(\mathbb{I}_{m_f, [\vec{k}]}, d_{m_f, [\vec{k}]}^{\nu', \{b'_n\}_{n=0}^{\infty}})$ have the same topology generated by their respective metrics.*

3.3.4 d is a metric

While it does not hold in general that the minimum of even a finite collection of metrics will be itself a metric, it does hold in this particular case. For this section fix an admissible measure

ν , a linear summable sequence $\{b_n\}_{n=0}^\infty$, a nonnegative integer m_f , and $\vec{k}, \vec{k}' \in \mathbb{Z}^{m_f}$. Let d denote $d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^\infty}$ and let d^p denote $d_{m_f, [\vec{k}]}^{p, \nu, \{b_n\}_{n=0}^\infty}$, as given in Definition 3.2.9. It is clear that d is positive definite and it is symmetric because $\mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$ is closed under inverses so we must only show that the triangle inequality holds. We show this in Lemma 3.3.12 but first we must prove two lemmas.

Lemma 3.3.10. *Fix $\vec{k}, \vec{k}', \vec{k}'' \in \mathbb{Z}^{m_f}$ and let $\mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$ be as in Definition 3.2.6. Then for any fixed $q \in \mathcal{S}_{\vec{k}, \vec{k}''}^{m_f}$ we have that $\mathcal{S}_{\vec{k}'', \vec{k}'}^{m_f} = \{p \circ q^{-1} | p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}\}$.*

Proof. Let $r \in \mathcal{S}_{\vec{k}'', \vec{k}'}^{m_f}$. Then there exist constants $c_1, c_2 \in \mathbb{Z}$ such that

$$k_j - k_{q(j)}'' = c_1 \text{ and } k_j'' - k_{r(j)}' = c_2$$

for all $j = 1, \dots, m_f$. In particular, for $i = q(j)$ we have

$$k_j - k_{r(q(j))}' = (c_1 + k_i'') - (k_i'' - c_2) = c_1 + c_2$$

and so we conclude that $p = r \circ q \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$ and clearly $r = p \circ q^{-1}$ so $\mathcal{S}_{\vec{k}'', \vec{k}'}^{m_f} \subset \{p \circ q^{-1} | p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}\}$.

Now let $p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$ and $q \in \mathcal{S}_{\vec{k}, \vec{k}''}^{m_f}$ so there are constants $c, c_1 \in \mathbb{Z}$ such that

$$k_j - k_{p(j)}' = c \text{ and } k_j - k_{q(j)}'' = c_1.$$

Subtracting these two equations gives $k_{q(j)}'' - k_{p(j)}' = c - c_1$ and thus $p \circ q^{-1} \in \mathcal{S}_{\vec{k}'', \vec{k}'}^{m_f}$. □

Lemma 3.3.11. *Let $m, m', m'' \in \mathbb{I}_{m_f, [\vec{k}]}$ and suppose $p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$ and $q \in \mathcal{S}_{\vec{k}, \vec{k}''}^{m_f}$. Then*

$$d^p(m, m') \leq d^q(m, m'') + d^{p \circ q^{-1}}(m'', m').$$

Proof. The $m_f = 0$ case is trivial so assume $m_f > 0$. Since $p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}$ and $q \in \mathcal{S}_{\vec{k}, \vec{k}''}^{m_f}$ there must be constants $c, c_1 \in \mathbb{Z}$ such that

$$k_j - k_{p(j)}' = c \text{ and } k_j - k_{q(j)}'' = c_1.$$

Because d^p is a sum of distances we can use the triangle inequality for each term with an appropriate

permutation on the elements:

$$\begin{aligned}
d^p(m, m') &= \sum_{\bar{u} \in \{0,1\}^{m_f}} \nu(t_{\bar{u}}(\Delta) \ominus t_{p(\bar{u})}(T^{-c}(\Delta'))) \\
&\quad + \sum_{j=1}^{m_f} \left(d_0^{\{b_n\}_{n=0}^\infty}((S_j)^\infty, (S'_{p(j)})^\infty) + |h_j - h'_{p(j)}| \right) \\
&\leq \sum_{\bar{u} \in \{0,1\}^{m_f}} \left[\nu(t_{\bar{u}}(\Delta) \ominus t_{q(\bar{u})}(T^{-c_1}(\Delta''))) + \nu(t_{q(\bar{u})}(T^{-c_1}(\Delta'')) \ominus t_{p(\bar{u})}(T^{-c}(\Delta'))) \right] \\
&\quad + \sum_{j=1}^{m_f} \left(d_0^{\{b_n\}_{n=0}^\infty}((S_j)^\infty, (S''_{q(j)})^\infty) + d_0^{\{b_n\}_{n=0}^\infty}((S''_{q(j)})^\infty, (S'_{p(j)})^\infty) \right. \\
&\quad \left. + |h_j - h''_{q(j)}| + |h''_{q(j)} - h'_{p(j)}| \right) \\
&= d^q(m, m'') + d^{p \circ q^{-1}}(m'', m').
\end{aligned}$$

□

Notice that in the case that $p = q = \text{Id}$ this gives a proof of the triangle inequality for d^{Id} .

Lemma 3.3.12. *The triangle inequality holds for d .*

Proof. Let $m, m', m'' \in \mathbb{I}_{m_f, [\bar{k}]}$. There exists some $q \in \mathcal{S}_{\bar{k}, \bar{k}'}^{m_f}$ such that $d(m, m'') = d^q(m, m'')$ and by Lemma 3.3.10 we know that

$$\min_{p \in \mathcal{S}_{\bar{k}, \bar{k}'}^{m_f}} \{d^{p \circ q^{-1}}(m'', m')\} = d(m'', m').$$

Now, using the inequality from Lemma 3.3.11 we have that

$$\begin{aligned}
d(m, m') &= \min_{p \in \mathcal{S}_{\bar{k}, \bar{k}'}^{m_f}} \{d^p(m, m')\} \\
&\leq \min_{p \in \mathcal{S}_{\bar{k}, \bar{k}'}^{m_f}} \{d^q(m, m'') + d^{p \circ q^{-1}}(m'', m')\} \\
&= d^q(m, m'') + \min_{p \in \mathcal{S}_{\bar{k}, \bar{k}'}^{m_f}} \{d^{p \circ q^{-1}}(m'', m')\} \\
&= d(m, m'') + d(m'', m')
\end{aligned}$$

as desired. □

Combining the arguments in Sections 3.3.1 and 3.3.2 with the present section, in particular Proposition 3.2.2 and Lemma 3.3.12, we get the following.

Proposition 3.3.13. *Let $m_f \in \mathbb{Z}_{\geq 0}$, $\vec{k} \in \mathbb{Z}^{m_f}$, $\{b_n\}_{n=0}^{\infty}$ be a linear summable sequence, and ν an admissible measure. Then the space $(\mathbb{I}_{m_f, [\vec{k}]}, d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^{\infty}})$ is a metric space.*

3.3.5 Relation to the metric on the moduli space of toric systems

In [67] the authors construct a metric on the moduli space of (compact) toric integrable systems which we denote by \mathcal{M}_T . Recall there is a one-to-one correspondence between elements of \mathcal{M}_T and Delzant polytopes. The authors of [67] define a metric on \mathcal{M}_T by pulling back the natural metric on the space of Delzant polytopes given by the Lebesgue measure of the symmetric difference.

Toric integrable systems can also be viewed as compact semitoric systems with no focus-focus singularities. If $m_f = 0$ then $G_{m_f} \times \mathcal{G} = \emptyset$ and thus the affine invariant is a unique polygon, the Delzant polytope. To compare two such systems the semitoric metric defined in the present chapter takes the ν -measure of the symmetric difference of the polygons for some admissible measure ν , as opposed to using the standard Lebesgue measure on \mathbb{R}^2 as is done in [67]. Notice also that \mathcal{M}_T is not equal to $\mathcal{M}_{ST,0}$ because, for instance, there are elements of $\mathcal{M}_{ST,0}$ which are not compact.

Moreover it is possible for two toric systems to be isomorphic as semitoric systems but not isomorphic as toric systems. This is because if $(M, \omega, (J, H))$ and $(M', \omega', (J', H'))$ are two choices of 4 dimensional toric systems then a diffeomorphism $\phi : M \rightarrow M'$ is an isomorphism of toric systems if $\phi^*(J', H') = (J, H)$. This corresponds to taking f to be the identity in the definition of semitoric isomorphisms. Thus we see that if \sim represents the equivalence induced by semitoric isomorphisms we have that $\mathcal{M}_T / \sim \subset \mathcal{M}_{ST,0}$ so the metric on \mathcal{M}_T produces a topology on a subset of $\mathcal{M}_{ST,0}$ via the quotient topology.

In $\mathcal{M}_{ST,0}$ the semitoric invariant is a unique polygon so to conclude that the metrics produce the same topology it is sufficient to show that the same sequences of convex compact polygons converge with respect to both the Lebesgue measure and any admissible measure.

Lemma 3.3.14. *Let $\Delta_k, \Delta \subset \mathbb{R}^2$ be convex compact sets for each $k \in \mathbb{N}$, let μ denote the Lebesgue measure on \mathbb{R}^2 , and let ν be any admissible measure. Then $\lim_{k \rightarrow \infty} \mu(\Delta \ominus \Delta_k) = 0$ if and only if $\lim_{k \rightarrow \infty} \nu(\Delta \ominus \Delta_k) = 0$.*

Proof. If $\lim_{k \rightarrow \infty} \mu(\Delta \ominus \Delta_k) = 0$ we can see that $\lim_{k \rightarrow \infty} \nu_0(\Delta \ominus \Delta_k) = 0$ where ν_0 is the example of an admissible measure from Section 3.3.2. This is because $\nu_0(A) < \mu(A)$ for any set $A \subset \mathbb{R}^2$. Thus we conclude that $\lim_{k \rightarrow \infty} \nu(\Delta \ominus \Delta_k) = 0$ by Lemma 3.3.8.

Now we will show the other direction. Suppose $\lim_{k \rightarrow \infty} \nu(\Delta \ominus \Delta_k) = 0$ and fix $\varepsilon > 0$. Choose some $L > 0$ such that $\pi_1(\Delta) \subset [-L, L]$. By Lemma 3.3.7 we know there exists $x_0, y_0, y_1 \in \mathbb{R}$ with $y_0 < y_1$ and $x_0 \in [-L, L]$ such that the set $\{x_0\} \times [y_0, y_1] \subset \Delta$ is a subset of Δ_k for $k > K_1$ for some fixed $K_1 \in \mathbb{N}$. Now, suppose that $k > K_1$ and $p \in \Delta_k$ has $\pi_1(p) > L + 1$. Then, since Δ_k is convex, the triangle with vertices $(x_0, y_0), (x_0, y_1), p$, which we will denote by G_p , must be a subset of Δ_k . Since $\pi_1(\Delta) \subset [-L, L]$ we know that $G_p \setminus \pi_1^{-1}([-L, L]) \subset \Delta \ominus \Delta_k$ and the ν -measure of any such triangle G_p defined by a point $p \in \mathbb{R}^2$ with $\pi_1(p) > L$ is bounded below by a constant $c_1 = \nu(G_{p_0}) > 0$ where $p_0 = (L + 1, 0)$. This is because any triangle G_p where $\pi_1(p) > L$ contains a triangle $G_{(L+1, y)}$ for some $y \in \mathbb{R}$ and any such triangle is the image under a vertical, and thus ν -preserving, transformation of G_{p_0} . Similarly, $p \in \Delta_k$ for $k > K_1$ with $\pi_1(p) < -L$ would imply that $\nu(\Delta \ominus \Delta_k) > c_2$ for some constant $c_2 > 0$. Thus, since $\lim_{k \rightarrow \infty} \nu(\Delta \ominus \Delta_k) = 0$ we conclude that there exists some $K_2 > K_1$ such that $k > K_2$ implies that $\Delta_k \subset \pi_1^{-1}([-L, L])$. Since ν is admissible we know that there exists some $c_3 > 0$ such that $d\mu/d\nu < c_3$ on $\pi_1^{-1}([-L, L])$. Choose $K_3 > K_2$ such that $k > K_3$ implies that $\nu(\Delta \ominus \Delta_k) < \varepsilon/c_3$ and notice that

$$\mu(\Delta \ominus \Delta_k) = \frac{d\mu}{d\nu} \nu(\Delta \ominus \Delta_k) < c_3 \nu(\Delta \ominus \Delta_k) < \varepsilon,$$

because while the Radon-Nikodym derivative is not bounded on all of \mathbb{R}^2 it is bounded on the set $\Delta \ominus \Delta_k$ for large enough k . □

Corollary 3.3.15. *The metric d induces the same topology on \mathcal{M}_T as the metric defined in [67] does.*

Corollary 3.3.15 follows from Lemma 3.3.14. This result is concerning compact polygons. Of course, if we consider non-compact sets these metrics will not induce the same topology.

Remark 3.3.16. Let $\mathcal{M}_{ST}^{\text{cpt}} \subset \mathcal{M}_{ST}$ be the collection of *compact* semitoric integrable systems. Then the polygons produced will always be compact and thus Lemma 3.3.14 applies. So we can conclude that when restricting to $\mathcal{M}_{ST}^{\text{cpt}}$ the standard Lebesgue measure can be used in place of the choice of

admissible measure and the same topology will be produced. \circlearrowright

3.3.6 d and d^{Id} induce the same topology

Let

$$\mathbb{I}_{m_f, \vec{k}} = \{m \in \mathbb{I}_{m_f, [\vec{k}]} \mid m \text{ has twisting index } \vec{k}\}$$

and define $d^{\text{Id}} = d^{\text{Id}, \nu, \{b_n\}_{n=0}^\infty}$ on \mathbb{I} by

$$d^{\text{Id}}(m, m') = \begin{cases} d_{m_f, \vec{k}}^{\text{Id}, \nu, \{b_n\}_{n=0}^\infty}(m, m') & \text{if } m, m' \in \mathbb{I}_{m_f, \vec{k}} \text{ for some } m_f \in \mathbb{Z}_{\geq 0}, \vec{k} \in \mathbb{Z}^{m_f} \\ 1 & \text{otherwise.} \end{cases}$$

Both d and d^{Id} are defined on \mathbb{I} and the main result of this section will be that both of these metrics induce the same topology on \mathbb{I} .

Lemma 3.3.17. *Let $m, m_n \in \mathbb{I}$ for $n \in \mathbb{N}$. Then $d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^\infty}(m, m_n) \xrightarrow{n \rightarrow \infty} 0$ implies that $\lambda_j^n \xrightarrow{n \rightarrow \infty} \lambda_j$ for all $j = 1, \dots, m_f$.*

Proof. Again we use d to denote $d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^\infty}$ and d^p to denote $d_{m_f, [\vec{k}]}^{p, \nu, \{b_n\}_{n=0}^\infty}$.

Step 1: Let $p_n \in \mathbb{S}^{m_f}$ satisfy $d(m, m_n) = d^{p_n}(m, m_n)$ for each $n \in \mathbb{N}$. For the first step of this proof we will argue that $\lambda_{p_n(j)}^n \xrightarrow{n \rightarrow \infty} \lambda_j$ by contrapositive. Suppose there exists some $j \in 1, \dots, m_f$ such that $\lambda_{p_n(j)}^n \not\xrightarrow{n \rightarrow \infty} \lambda_j$ as $n \rightarrow \infty$. This means there exists $a > 0$ and a subsequence $(n_i)_{i=0}^\infty$ such that

$$\left| \lambda_{p_{n_i}(j)}^{n_i} - \lambda_j \right| > a \text{ for all } i \in \mathbb{N}.$$

Now let $t_j = t_{\ell_{\lambda_j}}^1$ and $t_j^{n_i} = t_{\ell_{\lambda_{p_{n_i}(j)}^{n_i}}}^1$. Let Δ be a polygon which represents a choice of $\vec{\varepsilon} = \{+1, \dots, +1\}$ for m . We must show that $\nu(t_j(\Delta) \ominus t_j^{n_i}(\Delta))$ is bounded away from zero. We may assume that a is less than the horizontal distance from λ_j to the edge of the polygon Δ because $\lambda_j \in \text{int}(\pi_1(\Delta))$. Let $b = \min_{x \in [\lambda_j - a, \lambda_j + a]} \{\text{length}(\Delta \cap \ell_x)\}$ and notice that since Δ is a convex polygon we must have that $b > 0$.

The set Δ may be shifted by a vertical transformation so that $\max\{\pi_2(\Delta \cap \ell_x)\} = 0$ for each $x \in \mathbb{R}$ to form a new set $\Delta' \subset \mathbb{R}^2$, as is shown in Figure 3.5. Let $A: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the composition of these transformations so $A(\Delta) = \Delta'$. This new set may not be convex but since ν is invariant under




Figure 3.5: Without changing the ν -measure we can produce a new polygon which has $\{y = 0\}$ as its top boundary.




Figure 3.6: Either $\lambda_j^{n_i} < \lambda_j - a$ or $\lambda_j^{n_i} > \lambda_j + a$. Each case is shown above and in either case there is some positive measure region which is always in the symmetric difference. This causes convergence to be impossible.

vertical translations we have that $\nu(\Delta) = \nu(\Delta')$. Notice that $\mathcal{B} = [\lambda_j - a, \lambda_j + a] \times [-b, 0]$ satisfies $\mathcal{B} \subset \Delta'$.

Now there are two cases, both shown in Figure 3.6. If $\lambda_j < \lambda_j^{n_i}$ then $t_j(\mathcal{B}) \cap \{y > 0\} \subset t_j(\Delta') \ominus t_j^{n_i}(\Delta')$. This is because $t_j^{n_i}$ is the identity on points where $x \in [\lambda_j - a, \lambda_j + a]$ and so for x in this interval Δ' does not intersect the open upper half plane. The set $t_j(\mathcal{B}) \cap \{y > 0\}$ always contains the rectangle $[\lambda_j + a/2, \lambda_j + a] \times [0, a/2]$, as in Figure 3.6. Let $c_1 = \nu([\lambda_j + a/2, \lambda_j + a] \times [0, a/2])$.

Now suppose that $\lambda_j > \lambda_j^{n_i}$. In this case the symmetric difference always contains the region $t_j([\lambda_j, \lambda_j + a] \times [a - b, a])$ which has the same measure as $[\lambda_j, \lambda_j + a] \times [a - b, a]$; see Figure 3.6. Let $c_2 = \nu([\lambda_j, \lambda_j + a] \times [a - b, a])$ and let $c = \min\{c_1, c_2\}$. So in any case we have that $\nu(t_j(\Delta) \ominus t_j^{n_i}(\Delta)) \geq c > 0$.

Assume that $\lim_{n \rightarrow \infty} d(m, m_n) = 0$. This implies $\lim_{i \rightarrow \infty} \nu(\Delta \ominus \Delta^{n_i}) = 0$. In this case fix

$\varepsilon > 0$ such that $\varepsilon < c$, and find $I > 0$ such that $i > I$ implies that $\nu(\Delta \ominus \Delta^{n_i}) < \varepsilon$. Then for $i > I$ we have that

$$\begin{aligned} \nu[t_j(\Delta) \ominus t_j^{n_i}(\Delta)] &\leq \nu[t_j(\Delta) \ominus t_j^{n_i}(\Delta^{n_i})] + \nu[t_j^{n_i}(\Delta^{n_i}) \ominus t_j^{n_i}(\Delta)] \\ &= \nu[t_j(\Delta) \ominus t_j^{n_i}(\Delta^{n_i})] + \nu[\Delta^{n_i} \ominus \Delta], \end{aligned}$$

which implies

$$\nu[t_j(\Delta) \ominus t_j^{n_i}(\Delta^{n_i})] \geq \nu[t_j(\Delta) \ominus t_j^{n_i}(\Delta)] - \nu[\Delta^{n_i} \ominus \Delta] > c - \varepsilon.$$

Thus $\lim_{n \rightarrow \infty} \nu[t_j(\Delta) \ominus t_j^n(\Delta^n)] = 0$ is impossible, but this is a term in $d(m, m_n)$ so $d(m, m_n) \rightarrow 0$ is impossible as well. We conclude that $\lambda_{p_n(j)}^n \rightarrow \lambda_j$ for all $j = 1, \dots, m_f$.

Step 2: From Step 1 we know that $\lambda_{p_n(j)}^n \xrightarrow{n \rightarrow \infty} \lambda_j$ for each $j = 1, \dots, m_f$. Let $D = \min\{|\lambda_j - \lambda_{j'}| \mid j, j' \in \{1, \dots, m_f\}, j \neq j'\}$. Then there exists some $N > 0$ such that $n > N$ implies that $|\lambda_{p_n(j)}^n - \lambda_j| < d/2$. Thus, for $n > N$ we have that $p_n = \text{Id}$ and the result follows. \square

Proposition 3.3.18. *Let $m_f \in \mathbb{Z}_{\geq 0}$, $\vec{k} \in \mathbb{Z}^{m_f}$, $\{b_n\}_{n=0}^\infty$ be a linear summable sequence, and ν be an admissible measure. Then $d_{m_f, \vec{k}}^{\nu, \{b_n\}_{n=0}^\infty}$ and $d_{m_f, \vec{k}}^{\text{Id}, \nu, \{b_n\}_{n=0}^\infty}$ induce the same topology on \mathbb{I} .*

Proof. Any sequence which converges for d^{Id} will converge for d because $d < d^{\text{Id}}$. Suppose that $(m_n)_{n=1}^\infty$ is a sequence in \mathbb{I} which converges to $m \in \mathbb{I}$ with respect to d . Then by Step 2 of the proof of Lemma 3.3.17 we know there exists some $N > 0$ such that for $n > N$ we have that $d(m, m_n) = d^{\text{Id}}(\tau, \tau_n)$. Thus, we see that the sequence $d^{\text{Id}}(m, m_n)$ is eventually equal to a sequence which converges to zero, so we conclude that $d^{\text{Id}}(m, m_n) \xrightarrow{n \rightarrow \infty} 0$. \square

Remark 3.3.19. Each $\mathbb{I}_{m_f, \vec{k}} \subset \mathbb{I}$ is in a separate component of (\mathbb{I}, d) . This is because these are defined to be in different components for d^{Id} and we have just shown that d^{Id} and d induce the same topology. \circlearrowright

3.4 The completion

In this section we compute the completion of the space of semitoric ingredients \mathbb{I} which corresponds to the completion of \mathcal{M}_{ST} by Theorem 2.2.27. We will show that the completion of \mathbb{I} is

$\widetilde{\mathbb{I}}$, where $\widetilde{\mathbb{I}}$ is as is described in Definition 3.2.11 and Definition 3.4.9. The completion of an open interval in \mathbb{R} with the usual metric is the corresponding closed interval and we have already stated that $\mathbb{R}[[X, Y]]_0$ is complete (Proposition 3.2.2), so to produce the completion of \mathbb{I} it seems the only difficulty will be with the weighted polygons. This is not the case since in fact defining the distance as a minimum of permutations has intertwined the metrics on these different spaces so we can not consider them separately. This section has similar arguments to those in [67] except that in our case we must consider a whole family of polygons all at once instead of only one polygon. For the remainder of this section fix some admissible measure ν , some linear summable sequence $\{b_n\}_{n=0}^{\infty}$, a nonnegative integer m_f , and a vector $\vec{k} \in \mathbb{Z}^{m_f}$. For simplicity we will use d and d^p to refer to $d_{m_f, [\vec{k}]}^{\nu, \{b_n\}_{n=0}^{\infty}}$ and $d_{m_f, [\vec{k}]}^{p, \nu, \{b_n\}_{n=0}^{\infty}}$ (from Definition 3.2.9) respectively, where $p \in \mathcal{S}^{m_f}$.

In Section 3.4.1 we show that the completion must contain $\widetilde{\mathbb{I}}$ and in the remaining subsections we show that $\widetilde{\mathbb{I}}$ is complete. In Section 3.4.2 we prove several Lemmas about Cauchy sequences which are used in Section 3.4.3 to conclude that $\widetilde{\mathbb{I}}$ is in fact the completion of \mathbb{I} .

There is no way for elements of \mathbb{I} with different numbers of focus-focus points or twisting indices that are not equivalent (under the equivalence from Definition 3.2.5) to be close to one another because the distance between any two such systems is always 1 (see Definition 3.2.10). Thus, we will work with the components $\mathbb{I}_{m_f, [\vec{k}]}$ of \mathbb{I} .

First, notice that the definition of d from Definition 3.2.9 holds on $\widetilde{\mathbb{I}}$ as well. That is, extend the definition of d in the following way:

Definition 3.4.1. Suppose that $m, m' \in \widetilde{\mathbb{I}}$ denote

$$([A_w], (h_j)_{j=1}^{m_f}, ((S_j)^\infty)_{j=1}^{m_f}) \text{ and } ([A'_w], (h'_j)_{j=1}^{m'_f}, ((S'_j)^\infty)_{j=1}^{m'_f})$$

, respectively. Then:

1. the *comparison with alignment* p is

$$d^p(m, m') = d_{\mathbb{P}}^{p, \nu}([A_w], [A'_w]) + \sum_{j=1}^{m_f} \left(d_0^{\{b_n\}_{n=0}^{\infty}}((S_j)^\infty, (S_{p(j)})^\infty) + |h_j, h'_{p(j)}| \right);$$

2. the distance between m and m' is

$$d(m, m') = \min_{p \in \mathcal{S}_{\vec{k}, \vec{k}'}^{m_f}} \{d^p(m, m')\}.$$

Proposition 3.4.2. d is a metric on $\widetilde{\mathbb{I}}$.

This proposition follows from the proof of Proposition 3.3.13.

Remark 3.4.3. Notice that d^{ld} is not a metric on $\widetilde{\mathbb{I}}$ because it does not satisfy the triangle inequality.

This can be seen in Example 3.4.16. \circlearrowright

Throughout Section 3.4.1 each space we examine can be viewed as a subspace of $\widetilde{\mathbb{I}}$ and we will endow them with the structure of a metric subspace.

Remark 3.4.4. The space \mathbb{I} can be viewed as a subspace of $\widetilde{\mathbb{I}}$ because there is a natural correspondence between the elements of \mathbb{I} and the elements of a subset of $\widetilde{\mathbb{I}}$. This is because there is at most one element of \mathbb{I} in each equivalence class in $\widetilde{\mathbb{I}}$ so the space \mathbb{I} corresponds to the subset $\{[m] \mid m \in \mathbb{I}\}$. \circlearrowright

3.4.1 The completion must contain $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$

In the next few lemmas we start with $\mathbb{I}_{m_f, [\vec{k}]}$ and build up to $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$ in several steps, showing that each inclusion is dense. First we will show that the completion of $\mathbb{I}_{m_f, [\vec{k}]}$ must include at least all rational labeled polygons which satisfy the convexity requirements.

Lemma 3.4.5. Let $\mathcal{P}'_{m_f, [\vec{k}]} \subset \widetilde{\mathcal{DPoly}}_{m_f, [\vec{k}]}(\mathbb{R}^2)$ be given by

$$\mathcal{P}'_{m_f, [\vec{k}]} = \left\{ [(\Delta, (\ell_{\lambda_j}, +1, k'_j)_{j=1}^{m_f})] \left| \begin{array}{l} t_{\vec{u}}(\Delta) \in \text{Poly}(\mathbb{R}^2) \text{ for any } \vec{u} \in \{0, 1\}^{m_f}, \\ \vec{k} \sim \vec{k}', \nu(\Delta) < \infty, \text{ and} \\ \min_{s \in \Delta} \pi_1(s) < \lambda_1 < \dots < \lambda_{m_f} < \min_{s \in \Delta} \pi_1(s) \end{array} \right. \right\}$$

and let

$$\mathbb{I}'_{m_f, [\vec{k}]} = \mathcal{P}'_{m_f, [\vec{k}]} \times [0, 1]^{m_f} \times \mathbb{R}[[X, Y]]_0^{m_f}.$$

Then the inclusion $\mathbb{I}_{m_f, [\vec{k}]} \subset \mathbb{I}'_{m_f, [\vec{k}]}$ is dense.

Proof. Fix any element $m = ([\Delta_w], (h_j)_{j=1}^{m_f}, ((S_j)^\infty)_{j=1}^{m_f}) \in \mathbb{I}'_{m_f, [\vec{k}]}$. Since $d \leq d^{\text{ld}}$ we will show there exists an element $m' \in \mathbb{I}_{m_f, [\vec{k}]}$ arbitrarily close to m with respect to the function d^{ld} . Clearly we will have no problems with making the volume invariant or the Taylor series arbitrarily close so just consider the polygons.

Let $[\Delta_w] = [(\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})] \in \mathcal{P}'_{m_f, [\vec{k}]}$ and fix $\varepsilon > 0$. We will show there exists some element $[\Delta'_w] \in \mathcal{DPolyg}_{m_f, [\vec{k}]}(\mathbb{R}^2)$ such that $d_{\mathbb{P}}^{\text{ld}, \nu}([\Delta_w], [\Delta'_w]) < \varepsilon$. We will choose this element of $\mathcal{DPolyg}_{m_f, [\vec{k}]}(\mathbb{R}^2)$ to have the same λ_j values as Δ_w . Since the action of $t_{\vec{u}}$, $\vec{u} \in \{0, 1\}^{m_f}$, does not change the volume of sets we have

$$d_{\mathbb{P}}^{\text{ld}}([\Delta_w], [\Delta'_w]) \leq 2^{m_f} \nu(\Delta \ominus \Delta')$$

where $(\Delta', (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f}) \in [\Delta'_w]$. To complete the proof it suffices to show that there exists an element $(\Delta', (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f}) \in \mathcal{DPolyg}_{m_f, [\vec{k}]}(\mathbb{R}^2)$ such that Δ and Δ' are equal except on a set of ν -measure less than $2^{-m_f} \varepsilon$.

For $j = 1, \dots, m_f$ let $p_j \in \mathbb{R}^2$ be the intersection of ℓ_{λ_j} with the top boundary of Δ . Let $U \subset \mathbb{R}^2$ be a union of disjoint neighborhoods around each corner of Δ which is not an element of $\{p_j\}_{j=1}^{m_f}$ such that $\nu(U) < \varepsilon/2^{m_f+1}$. Also, let $V \subset \mathbb{R}^2 \setminus U$ be a union of disjoint neighborhoods around each point p_j for each $j = 1, \dots, m_f$ and $\nu(V) < \varepsilon/2^{m_f+1}$. We will define Δ' in several stages, editing it several times. Start by assuming that $\Delta' = \Delta$. By [67, Remark 23] we can edit Δ' on the set U so that every vertex is Delzant except possibly the ones in V .

Now, recall that for a *semitoric* polygon to be Delzant the points p_j must all either be fake or hidden Delzant corners. This is equivalent to saying that the corners on the top boundary of $t_{\vec{u}}(\Delta')$ must all be Delzant for $\vec{u} = \langle 1, \dots, 1 \rangle$. Since $t_{\vec{u}}(\Delta')$ is a convex polygon and $t_{\vec{u}}(V)$ is a neighborhood of the edges $t_{\vec{u}}(p_j)$ we can again use [67, Remark 23] to conclude that we may edit $t_{\vec{u}}(\Delta')$ inside of the set V such that all of the vertices on the top boundary are Delzant. Now we have finished defining $t_{\vec{u}}(\Delta')$ and since this map is invertible we have also defined Δ' . Notice that for $j = 1, \dots, m_f$ each point $t_{\vec{u}}(p_j)$ is either a Delzant corner, which would make p_j a hidden Delzant corner, or it is not a vertex at all, in which case p_j would be a fake corner. Also, it is easy to check that any new Delzant corner we had to define in $t_{\vec{u}}(V)$ which is not on the point $t_{\vec{u}}(p_j)$ for some $j = 1, \dots, m_f$ gets transformed by $t_{\vec{u}}^{-1}$ to form a Delzant corner on Δ' . In conclusion, $[\Delta'_w]$ is a

Delzant semitoric polygon and each of the 2^{m_f} polygons in the equivalence class is equal to each polygon in $[\Delta_w]$ except on a set of ν -measure less than $\varepsilon/2^{m_f}$. \square

So from the above Lemma we conclude that the completion of $\mathbb{I}_{m_f, [\vec{k}]}$ must contain $\mathbb{I}'_{m_f, [\vec{k}]}$. In the next Lemma we show it must contain a larger set. The only difference between $\mathcal{P}'_{m_f, [\vec{k}]}$ and $\mathcal{P}''_{m_f, [\vec{k}]}$ is that $\mathcal{P}''_{m_f, [\vec{k}]}$ allows irrational polygons.

Lemma 3.4.6. *Let*

$$\mathcal{P}''_{m_f, [\vec{k}]} = \left\{ \left[(\Delta, (\ell_{\lambda_j}, +1, k'_j)_{j=1}^{m_f}) \mid \begin{array}{l} t_{\vec{u}}(\Delta) \text{ is a convex polygon for any } \vec{u} \in \{0, 1\}^{m_f}, \\ 0 < \nu(\Delta) < \infty, \vec{k} \sim \vec{k}' \text{ and} \\ \min_{s \in \Delta} \pi_1(s) < \lambda_1 < \dots < \lambda_{m_f} < \max_{s \in \Delta} \pi_1(s) \end{array} \right. \right\}$$

and let

$$\mathbb{I}''_{m_f, [\vec{k}]} = \mathcal{P}''_{m_f, [\vec{k}]} \times [0, 1]^{m_f} \times \mathbb{R}[[X, Y]]_0^{m_f}.$$

Then the inclusion $\mathbb{I}'_{m_f, [\vec{k}]} \subset \mathbb{I}''_{m_f, [\vec{k}]}$ is dense.

Proof. Just as in the proof of Lemma 3.4.5 we can see that we only need to consider the polygons. Suppose that $[\Delta_w] \in \mathcal{P}''_{m_f, [\vec{k}]}$ and $(\Delta, (\ell_{\lambda_j}, +1, k'_j)_{j=1}^{m_f}) \in [\Delta_w]$. Given any $\varepsilon > 0$ we can find an open neighborhood of the boundary of Δ which has ν -measure less than ε (since the boundary has measure zero and ν is regular) and we may approximate Δ by a rational polygon with boundary inside of this neighborhood. In the case that Δ is compact this can be done by approximating the irrational slopes with rational ones (exactly as done in [67]).

This strategy will work even if Δ is not compact. For the faces of Δ which are non-compact with irrational slope (if there are any) we can still approximate these with a line of rational slope because of the properties of the admissible measure ν . Suppose there is a non-compact face of Δ which has irrational slope $r \in \mathbb{R} \setminus \mathbb{Q}$. Then choose $q \in \mathbb{Q}$ such that $q < r$ and $\nu(\{qx < y < rx\}) < \varepsilon$ and let the edge on the rational polygon have slope q . Such a slope can be chosen because if the measure of that set is always finite and replacing q by $q_2 = q+r/2$ will produce a wedge with half the measure of the original. \square

Remark 3.4.7. Since it is possible for $\lambda_j = \lambda_{j+1}$ for some $j \in 1, \dots, m_f - 1$ the order in which the

critical points are labeled in a system cannot be made unique by only considering those λ values. This means that there could be two elements in $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$ which have the same invariants except labeled in a different order. Of course, we do not want this because these two elements should be the same, so we use the other invariants to create a unique ordering on the critical points of any element of $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$. We fix the order so that if $\lambda_j = \lambda_{j+1}$ for some $j = 1, \dots, m_f - 1$ then we require that $h_j \leq h_{j+1}$. In the case that $\lambda_j = \lambda_{j+1}$ and $h_j = h_{j+1}$ we look to the Taylor series. In this situation we require that the coefficient of X of the Taylor series $(S_j)^\infty$ is less than or equal to the coefficient of X in $(S_{j+1})^\infty$ and if those are equal we look to the coefficient of Y and continue in this fashion. Now given any system with critical points there is a unique order in which to label them which is essentially the lexicographic order on the invariants. \circlearrowright

For the next Lemma we only slightly change the restrictions on the $(\lambda_j)_{j=1}^{m_f}$. Notice that we allow (positive only) infinite values for the λ_j . This can only happen in the case that the polygon is non-compact. If $\lambda_j = +\infty$ then we define t_j^1 to be the identity because all of \mathbb{R}^2 is to the left of this value.

Lemma 3.4.8. *Let*

$$\mathcal{P}'''_{m_f, [\vec{k}]} = \left\{ \left[(\Delta, (\ell_{\lambda_j}, +1, k'_j)_{j=1}^{m_f}) \right] \left| \begin{array}{l} t_{\vec{u}}(\Delta) \text{ is a convex polygon for any } \vec{u} \in \{0, 1\}^{m_f}, \\ 0 < \nu(\Delta) < \infty, \vec{k} \sim \vec{k}', \\ \lambda_j \in \mathbb{R} \cup \{\infty\} \text{ for } j = 1, \dots, m_f, \text{ and} \\ \min_{s \in \Delta} \pi_1(s) \leq \lambda_1 \leq \dots \leq \lambda_{m_f} \leq \max_{s \in \Delta} \pi_1(s) \end{array} \right. \right\}$$

and let

$$\mathbb{I}'''_{m_f, [\vec{k}]} = \mathcal{P}'''_{m_f, [\vec{k}]} \times [0, 1]^{m_f} \times \mathbb{R}[[X, Y]]_0^{m_f}.$$

The inclusion $\mathbb{I}''_{m_f, [\vec{k}]} \subset \mathbb{I}'''_{m_f, [\vec{k}]}$ is dense.

Proof. Again, we only need to consider the polygons. We will prove this Lemma in two steps. First, suppose that $[\Delta_w] \in \mathcal{P}'''_{m_f, [\vec{k}]}$ has $\lambda_j < \infty$ for each $j = 1, \dots, m_f$ so the only thing that is keeping $[\Delta_w]$ from being in $\mathcal{P}''_{m_f, [\vec{k}]}$ is the possibility that $\lambda_j = \lambda_{j+1}$ for some fixed $j \in \{1, \dots, m_f - 1\}$. Let \vec{u} be all zeros except for a 1 in the j^{th} and $(j+1)^{\text{st}}$ positions. Then $[\Delta_w] \in \mathcal{P}'''_{m_f, [\vec{k}]}$ implies that $t_{\vec{u}}(\Delta)$ is convex so we know that there is a vertex of Δ on the top boundary with x -coordinate λ_j .




Figure 3.7: By cutting the corner and adjusting the values of λ_j and λ_{j+1} of an element in $\mathcal{P}'''_{m_f, [\vec{k}]}$ we can produce an element of $\mathcal{P}''_{m_f, [\vec{k}]}$ which is very close.

Let m_1 denote the slope of the edge to the left of this vertex and let m_2 denote the slope to the right. Then we can see that the convexity of $t_{\vec{u}}(\Delta)$ implies that $m_1 \geq m_2 + 2$. Now we want to show that there exists some $[\Delta'_w] \in \mathcal{P}''_{m_f, [\vec{k}]}$ arbitrarily close in d^{ld} to $[\Delta_w]$. Let $[\Delta'_w]$ be equal to $[\Delta_w]$ except that $\lambda'_j < \lambda_j < \lambda'_{j+1}$ and that the top boundary of Δ' has slope $m_1 - 1$ on the interval $x \in (\lambda'_j, \lambda'_{j+1})$. So, as is shown in Figure 3.7, we have cut the corner off of Δ to produce Δ' and clearly this cut can be made as small as desired. This process can be repeated for each instance of $\lambda_j = \lambda_{j+1}$ for $j \in \{1, \dots, m_f\}$.

Now we proceed to step two. Assume that $[\Delta_w] = [(\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})] \in \widetilde{\mathcal{DPolyg}}_{m_f, [\vec{k}]}(\mathbb{R}^2)$ has $\lambda_{m_f} = +\infty$ (and $\lambda_j < \infty$ for $j = 1, \dots, m_f - 1$) and we will construct a sequence with $[\Delta_w]$ as its limit. Let $N = \max_{j=1, \dots, m_f-1} |\lambda_j|$ and for any $n \in \mathbb{N}$ which satisfies $n > N$ define a set $\Delta^n = \Delta \cap [-n, n]$ with $\lambda_{m_f} = n$. That is

$$[\Delta_w^n] = (\Delta^n, (\ell_{\lambda_j}, +1)_{j=1}^{m_f-1}, (\ell_n, +1)).$$

Notice that each polygon in each family $[\Delta_w^n]$ is convex because it is the intersection of two convex sets. Then $d_{\mathbb{P}}^{\text{c}}([\Delta_w], [\Delta_w^n]) \rightarrow 0$. Clearly a similar process can be used to produce sets which have multiple λ values which are infinite.

□

Next we would like to consider arbitrary convex sets, but there is a subtlety. We must instead consider all sets which are convex up to measure zero corrections (as is done in [67]). So far we have only been working with polygons and if the symmetric difference of two polygons has zero measure in ν , and therefore also in μ , those polygons are the same set. This is not true for arbitrary subsets of \mathbb{R}^2 . Recall that ν and the Lebesgue measure μ have exactly the same measure zero sets, so the equivalence relation in the following definition does not depend on the choice of admissible measure.

Definition 3.4.9. Let

$$\mathcal{C}_{m_f, [\vec{k}]} = \left\{ \left[(A, (\ell_{\lambda_j}, +1, k'_j)_{j=1}^{m_f}) \right] \left| \begin{array}{l} A \subset \mathbb{R}^2, \lambda_j \in \mathbb{R} \cup \{\infty\} \text{ for } j = 1, \dots, m_f, \\ t_{\vec{u}}(A) \text{ is a convex set for any } \vec{u} \in \{0, 1\}^{m_f}, \\ \vec{k} \sim \vec{k}', 0 < \nu(\Delta) < \infty, \text{ and} \\ \min_{s \in A} \pi_1(s) \leq \lambda_1 \leq \dots \leq \lambda_{m_f} \leq \max_{s \in A} \pi_1(s) \end{array} \right. \right\}.$$

Further, for any measurable sets $A, B \subset \mathbb{R}^2$ we say $A \simeq B$ if and only if $\mu(A \ominus B) = 0$ and let $[A]$ denote the equivalence class of A with respect to this relation. Finally, let

$$\widetilde{\mathcal{DPolyg}}_{m_f, [\vec{k}]}(\mathbb{R}^2) = \left\{ \left[([A], (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f}) \right] \left| \begin{array}{l} [(A, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})] \in \mathcal{C}_{m_f, [\vec{k}]} \text{ or} \\ \nu(A) = 0 \text{ and } \lambda_j = 0 \text{ for } j = 1, \dots, m_f \end{array} \right. \right\}.$$

Here it is important to notice that we have included one extra element in each $\widetilde{\mathcal{DPolyg}}_{m_f, [\vec{k}]}(\mathbb{R}^2)$, the equivalence class of the empty set. For this element the values of λ_j are unimportant so we set them all equal to zero (in fact, any fixed number will work). For the last Lemma in this section we will show that the inclusion in $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$, which is defined in Definition 3.2.11, is also dense.

Lemma 3.4.10. *The inclusion $\mathbb{I}'''_{m_f, [\vec{k}]} \subset \widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$ is dense¹.*

Proof. Once more, we only have to consider the labeled weighted convex sets since it is easy to align the volume invariant and Taylor series invariant. Let $[(A, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})] = [\Delta_w] \in \widetilde{\mathcal{DPolyg}}_{m_f, [\vec{k}]}(\mathbb{R}^2)$. Now pick $[(B, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})] \in \mathcal{P}'''_{m_f, [\vec{k}]}$ and notice that they have the same λ values, so if A and B are close then so are all of the other polygons. Simply approximate A by a family of disjoint

¹The explanation of how $\mathbb{I}'''_{m_f, [\vec{k}]}$ can be viewed as a subspace of $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$ is in Remark 3.4.4.




Figure 3.8: An arbitrary convex set can be approximated from the inside by a polygon. The convexity requirements will be met as long as the vertices on the top boundary at $\{x = \lambda_j\}$ for each $j = 1, \dots, m_f$ are included in the polygon.

rectangles contained in A . We need to be sure that $t_{\vec{u}}(B)$ is convex for any choice of $\vec{u} \in \{0, 1\}^{m_f}$ so take B to be the convex hull of the rectangles which approximate A from the inside and the points in the top boundary of A which have x -value equal to λ_j for some $j \in \{1, \dots, m_f\}$. Since $B \subset A$ and $t_{\vec{u}}(A)$ is convex around $x = \lambda_j$ for each $j = 1, \dots, m_f$ we know that $t_{\vec{u}}(B)$ is convex (Figure 3.8).

□

From the results of Lemma 3.4.5, Lemma 3.4.6, Lemma 3.4.8, and Lemma 3.4.10, the following lemma is immediate.

Lemma 3.4.11. *The completion of $\mathbb{I}_{m_f, [\vec{k}]}$ must contain $\tilde{\mathbb{I}}_{m_f, [\vec{k}]}$.*

3.4.2 Cauchy sequences for d and d^{Id}

In this section we investigate the relationship between Cauchy sequences in d^{Id} and d^p . This will be used to prove Lemma 3.4.14; that $\tilde{\mathbb{I}}$ is complete.

Lemma 3.4.12. *Let $m_n \in \tilde{\mathbb{I}}_{m_f, [\vec{k}]}$ for $n = 1, \dots, \infty$. If $(m_n)_{n=1}^{\infty}$ is Cauchy with respect to d then there exists a subsequence $(m_{n_i})_{i=1}^{\infty}$ which is Cauchy with respect to d^{Id} .*

Proof. Let $(m_n)_{n=1}^{\infty}$ be as in the statement of the Lemma. Let $A_0 = \mathbb{N}$ and let $M_0 = 0$. We will define A_n and M_n recursively for each $n \in \mathbb{N}$. Suppose that $|A_{n-1}| = \infty$ and $M_{n-1} \in A_{n-1}$. Let $\varepsilon_n = 2^{-n}$. Find some $M > 0$ such that $k, l > M$ implies that $d(m_k, m_l) > \varepsilon_n/2$. Now let M_n be any element of A_{n-1} which is greater than M and M_{n-1} . This means $d(m_{M_n}, m_l) < \varepsilon_n/2$ for any $l > M_n$. For $p \in S^{m_f}$

let $\mathcal{B}_p^n = \{l \in A_{n-1} \mid l > M_n, d^p(m_{M_n}, m_l) < \varepsilon_n/2\}$. Notice that $\bigcap_{p \in \mathcal{S}^{m_f}} \mathcal{B}_p^n = \{l \in A_{n-1} \mid l > M_n\}$ by the definition of M_n . The union of this finite number of sets has infinite cardinality so at least one of those sets must also have infinite cardinality. Choose any $p_n \in \mathcal{S}^{m_f}$ such that $|\mathcal{B}_{p_n}^n| = \infty$ (there may be several possible choices). Now define $A_n = (A_{n-1} \cap [0, M_n]) \cup \mathcal{B}_{p_n}^n$. Notice that $M_n \in A_n$ and $|A_n| = \infty$.

Now let $A = \bigcap_{n \in \mathbb{N}} A_n$ and notice that $|A| = \infty$ because $\{M_n \mid n \in \mathbb{N}\} \subset A$.

So $(m_a)_{a \in A}$ is a subsequence of $(m_n)_{n=1}^\infty$. We will show that this subsequence is Cauchy with respect to d^{Id} . Fix any $\varepsilon > 0$ and find $n \in \mathbb{N}$ such that $\varepsilon_n < \varepsilon$. Now pick any $k, l > M_n$ with $k, l \in A$. Then $k, l \in A_n$ implies that $k, l \in \mathcal{B}_{p_n}^n$ so $d^{p_n}(m_{M_n}, m_k), d^{p_n}(m_{M_n}, m_l) < \varepsilon_n/2 < \varepsilon/2$. Also notice that p_n being an appropriate permutation to compare m_k with m_{M_n} and also appropriate to compare m_l with m_{M_n} implies that $\text{Id} \in \mathcal{S}^{m_f}$ is an appropriate permutation to compare m_k and m_l . Thus

$$d^{\text{Id}}(m_k, m_l) \leq d^{p_n}(m_k, m_{M_n}) + d^{p_n}(m_l, m_{M_n}) < \varepsilon$$

by Lemma 3.3.11. □

Lemma 3.4.13. *Suppose that $(m_n)_{n=1}^\infty$ is a sequence of elements of $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$ which is Cauchy with respect to the function d^{Id} . Then there exists some $p \in \mathcal{S}^{m_f}$ and $m \in \widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$ such that*

$$\lim_{n \rightarrow \infty} d^p(m_n, m) = 0.$$

Proof. For $A, B \subset \mathbb{R}^2$ say $A \simeq B$ if and only if $\nu(A \ominus B) = 0$ and let \mathcal{F} denote the subsets of \mathbb{R}^2 with finite ν -measure modulo \simeq . Now let $\mathcal{E} = \{[A] \in \mathcal{F} \mid \text{there exists } B \in [A] \text{ which is convex}\}$ and let $d_{\mathcal{E}}$ be the metric on this space given by the ν -measure of the symmetric difference. For simplicity we will write $A \in \mathcal{E}$ instead of $[A] \in \mathcal{E}$. We will show that this metric space is complete. Let χ_A denote the characteristic function of the set $A \in \mathcal{E}$. Then for $A, B \in \mathcal{E}$ we can see that

$$d_{\mathcal{E}}(A, B) = \int_{\mathbb{R}^2} |\chi_A - \chi_B| d\nu = \|\chi_A - \chi_B\|_{L^1}$$

the L^1 norm on (\mathbb{R}^2, ν) . Now suppose that $(A^k)_{k=1}^\infty$ is a Cauchy sequence in $(\mathcal{E}, d_{\mathcal{E}})$ and by measure zero adjustments we can assume that each A^k is convex. Then $(\chi_{A^k})_{k=1}^\infty$ is Cauchy in $L^1(\mathbb{R}^2, \nu)$ and

thus there must exist some function $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined up to measure zero such that

$$\lim_{k \rightarrow \infty} \|g - \chi_{A^k}\|_{L^1} = 0$$

because L^1 is complete.

The functions $(\chi_{A^k})_{k=1}^\infty$ converge to g in L^1 so we know that there is a subsequence $(\chi_{A^{k_n}})_{n=1}^\infty$ which converges to g pointwise off of some measure zero set S . Let

$$A = \{x \in \mathbb{R}^2 \setminus S \mid g(x) = 1\}$$

and now we will show that A is almost everywhere equal to a convex set so \mathcal{E} is complete. Let A' be the convex hull of A and we will show that $\nu(A \ominus A') = 0$. Let $p \in A'$ which means there exists $q, r \in A$ and $t \in [0, 1]$ such that $p = (1-t)q + tr$. Since the subsequence $(\chi_{A^{k_n}})_{n=1}^\infty$ converges pointwise to χ_A at the points q and r (since $q, r \in A$ and A is disjoint from S) this means that there exists some $N > 0$ such that $n > N$ implies $q, r \in A^{k_n}$. Thus, since each A^k is convex we see that for $n > N$ we have $p \in A^{k_n}$. We conclude that $p \in A \cup S$ and thus $A \ominus A' \subset S$ so $\nu(A \ominus A') = 0$. Also notice $\nu(A^k, A) \rightarrow 0$ as $k \rightarrow \infty$ implies that $\nu(A) < \infty$. This means $A \in \mathcal{E}$ so $(\mathcal{E}, d_{\mathcal{E}})$ is a complete metric space.

Let $([A_w^k])_{k=1}^\infty$ be a Cauchy sequence in $(\widetilde{\mathcal{DPoly}}_{m_f, [\vec{k}]}(\mathbb{R}^2), d_{\mathbb{P}}^{\text{Id}, \nu})$. Let

$$[A_w^k] = [(A^k, (\ell_{\lambda_j^k}, +1, k_j)_{j=1}^{m_f})] \text{ and let } A_{\vec{u}}^k = t_{\vec{u}}^k(A^k)$$

for each $\vec{u} \in \{-1, 1\}^{m_f}$ with $u_j = \frac{1-\epsilon_j}{2}$. Since this sequence is Cauchy we also know that the sequence $(A_{\vec{u}}^k)_{k=1}^\infty$ is a Cauchy sequence in $(\mathcal{E}, d_{\mathcal{E}})$. Thus for each $\vec{u} \in \{-1, 1\}^{m_f}$ there exists some convex $A_{\vec{u}} \in \mathcal{E}$ which is the limit of $(A_{\vec{u}}^k)_{k=1}^\infty$ in \mathcal{E} . Let $A = A_{(1, \dots, 1)}$. We have produced a family of convex, ν -finite sets which could be the limit, but we still need to check that there is some choice of $(\Lambda_j)_{j=1}^{m_f}$ such that $A_{\vec{u}} = t_{\vec{u}}(A_0)$ in \mathcal{E} for each $j = 1, \dots, m_f$.

Fix some $j \in \{1, \dots, m_f\}$ and let $A_j = A_{\epsilon}$ where $\epsilon_j = -1$ and $\epsilon_i = +1$ for $i \neq j$ and let t_k




Figure 3.9: The action of $t_{\lambda_j^{k_m}}^1$ and $t_{\lambda_j^{k_n}}^1$ on a polygon. It can be seen that $\mathcal{R}_{n,m}$ is a subset of the symmetric difference and has measure which is nonzero if $|\lambda_j^{k_n} - \lambda_j^{k_m}| \neq 0$.

denote $t_{\lambda_j^k}^1$. Since ν is invariant under vertical translations we have that

$$d_{\mathcal{E}}(t_k(A), t_{\bar{a}}(A^k)) = d_{\mathcal{E}}(A, A^k)$$

so both go to zero as $k \rightarrow \infty$. By the triangle inequality we can see that

$$d_{\mathcal{E}}(t_k(A), A_j) \leq d_{\mathcal{E}}(t_k(A), t_k(A^k)) + d_{\mathcal{E}}(t_k(A^k), A_j)$$

so we conclude that

$$d_{\mathcal{E}}(t_k(A), A_j) \rightarrow 0 \text{ as } k \rightarrow \infty. \quad (3.2)$$

If $(\lambda_j^k)_{k=1}^{\infty}$ diverges to $+\infty$ or converges to $\sup(\pi_1(A))$ then we are done. This is because in this case $d_{\mathcal{E}}(t^k(A), A_j) \rightarrow 0$ as $k \rightarrow \infty$ implies that A and A_j represent the same element in \mathcal{E} (i.e. they are equal almost everywhere) and t^{Λ} acts as the identity on A_0 if Λ is the rightmost value of A_0 .

Otherwise we can find some $x_0, a \in \mathbb{R}$ with $a > 0$ such that $[x_0, x_0 + 2a] \subset \pi_1(A)$ and there exists a subsequence $(\lambda_j^{k_n})_{n=1}^{\infty}$ such that $\lambda_j^{k_n} < x_0$ for all n . Notice that $A \cap \ell_x$ is an interval for any $x \in \pi_1(A)$ because A is convex. Let $\delta_1 = \text{length}(A \cap \ell_{x_0})$ and $\delta_2 = \text{length}(A \cap \ell_{x_0+a})$ and notice

that $\delta_1, \delta_2 < \infty$ because otherwise we would have $\nu(A) = \infty$ or $\nu(A) = 0$ because A is convex and ν is invariant under vertical translations. Also notice that $\text{length}(A \cap \ell_x) \geq \min\{\delta_1, \delta_2\}$ for any $x \in [x_0, x_0 + a]$ because A is convex. Pick any $n, m \in \mathbb{N}$ and we can see that $t_{\lambda_j^{k_n}}^1$ and $t_{\lambda_j^{k_m}}^1$ only differ by a vertical translation when acting on $A \cap \pi_1^{-1}([x_0, x_0 + a])$ (see Figure 3.9). This guarantees that there is a region $\mathcal{R}_{n,m}$ in the symmetric difference $t_{\lambda_j^{k_n}}^1(A) \ominus t_{\lambda_j^{k_m}}^1(A)$ which has the same measure as a rectangle of length a and height $\min\{\delta_1, \delta_2, |\lambda_j^{k_n} - \lambda_j^{k_m}|\}$ positioned between the x -values of x_0 and $x_0 + a$ (since ν is translation invariant). If R stands for the measure of a rectangle from $x = x_0$ to $x = x_0 + a$ of unit height we can see that

$$\nu(\mathcal{R}_{n,m}) = \min\left\{\delta_1, \delta_2, \left|\lambda_j^{k_n} - \lambda_j^{k_m}\right|\right\} R$$

and since $\mathcal{R}_{n,m} \subset t_{\lambda_j^{k_n}}^1(A) \ominus t_{\lambda_j^{k_m}}^1(A)$ we know that

$$\min\{\delta_1, \delta_2, |\lambda_j^{k_n} - \lambda_j^{k_m}|\} R \leq \nu(t_{\lambda_j^{k_n}}^1(A) \ominus t_{\lambda_j^{k_m}}^1(A)). \quad (3.3)$$

The right side of Equation (3.3) is Cauchy with respect to m and n because $(t_{\lambda_j^{k_n}}^1)_{n=1}^\infty$ converges by Equation (3.2) and thus the left side is Cauchy as well. This means that $(\lambda_j^{k_n})_{n=1}^\infty$ is a Cauchy sequence of real numbers and thus must converge. Call its limit $\Lambda_j \in \mathbb{R}$. To complete the proof we must only show that $\nu(t_{\Lambda_j}^1(A) \ominus A_j) = 0$. This is clear because

$$\nu(t_{\Lambda_j}^1(A) \ominus A_j) \leq \nu(t_{\Lambda_j}^1(A) \ominus t_{\lambda_j^{k_n}}^1(A)) + \nu(t_{\lambda_j^{k_n}}^1(A) \ominus A_j)$$

and the right side goes to zero as $n \rightarrow \infty$. So we conclude that the original Cauchy sequence converges to $[(A, (\ell_{\Lambda_j}, +1, k_j)_{j=1}^{m_f})]$. Clearly the elements of each copy of $\mathbb{R}[[X, Y]]_0$ and $[0, 1]$ can be made to converge. The only problem is that possibly this limit does not have the critical points labeled in the correct order according to Remark 3.4.7 to be an element of $\widetilde{\mathcal{DPoly}}_{m_f, [\vec{k}]}(\mathbb{R}^2)$ so we reorder it by some permutation $p \in \mathcal{S}^{m_f}$ and the result follows. \square

3.4.3 $\widetilde{\mathbb{I}}$ is complete

Lemma 3.4.14. $\widetilde{\mathbb{I}}_{m_f, [\vec{k}]}$ is complete.

Proof. Any Cauchy sequence in $\widetilde{\mathbb{I}}_{m_f, [\bar{k}]}$ must have a subsequence which is Cauchy with respect to d^{Id} by Lemma 3.4.12. By Lemma 3.4.13 that sequence must converge with respect to d^p for some fixed $p \in S^{m_f}$, which in particular means that it must converge with respect to d . A Cauchy sequence with a subsequence which converges must converge. \square

Now Lemma 3.4.11 and Lemma 3.4.14 imply the main result of this section.

Proposition 3.4.15. *Given an admissible measure ν and a linear summable sequence the completion of $(\mathbb{I}, d^\nu, \{b_n\}_{n=0}^\infty)$ is $(\widetilde{\mathbb{I}}, d^\nu, \{b_n\}_{n=0}^\infty)$.*

Example 3.4.16. The reason to use d instead of d^{Id} can be seen by the examining structure of the completion. Let

$$[\Delta_w^l] = \begin{cases} [(\Delta_l, (\lambda_1 = 0, \epsilon_1 = 1, k_1 = 0), (\lambda_2 = l, \epsilon_2 = 1, k_2 = 0))] & \text{if } l > 0 \\ [(\Delta_l, (\lambda_1 = l, \epsilon_1 = 1, k_1 = 0), (\lambda_2 = 0, \epsilon_2 = 1, k_2 = 0))] & \text{if } l < 0 \end{cases}$$

and suppose that $m_l \in \mathbb{I}$ is a system given by

$$m_l = \begin{cases} ([\Delta_w^l], ((S_1)^\infty, h_1), ((S_2)^\infty, h_2)) & \text{if } l > 0 \\ ([\Delta_w^l], ((S_2)^\infty, h_2), ((S_1)^\infty, h_1)) & \text{if } l < 0 \end{cases}$$

for $l \in [-1, 1] \setminus \{0\}$ such that $\lim_{l \rightarrow 0} m_l$ exists in $(\widetilde{\mathbb{I}}, d)$. This can be thought of as one of the critical points being fixed and the other passing over it at $l = 0$ as is shown in Figure 3.10. The complications in defining this come from the fact that the order of the critical points switches at $l = 0$ so the labeling has to switch. Now we can see the problem with using d^{Id} , which is that $\lim_{l \rightarrow 0^+} m_l \neq \lim_{l \rightarrow 0^-} m_l$ with respect to d^{Id} . \circlearrowright

Finally, Theorem 3.2.12 is produced by combining Proposition 3.3.13, Corollary 3.3.9, and Proposition 3.4.15.




Figure 3.10: A continuous family in $\tilde{\mathbb{I}}$ in which one critical point passes over the other as h increases from negative to positive.

3.5 Further questions

Now that we have defined a metric, and in particular a topology, on \mathcal{M}_{ST} there are several questions that would be natural to address. First of all, one may be interested extending the metric defined in this chapter in the way that this chapter has extended the metric from [67]. To produce such an extension to a larger class of integrable systems one would first have to classify those systems with invariants in a way which extends the classification from [69, 70]. Also, one can now ask what are the connected components of \mathcal{M}_{ST} , which is the topic of Chapter 4. Furthermore, with a topology on \mathcal{M}_{ST} we can consider Problem 2.45 from [72], which asks what the closure of the set of semitoric integrable systems would be when considered as a subset of $C^\infty(M, \mathbb{R}^2)$. To address this problem an appropriate topology on $C^\infty(M, \mathbb{R}^2)$ would have to be defined. This situation is much more general than the systems which are the focus of this chapter so it may be best to study metrics constructed in a more general case such as in Chapter 9.

This chapter is partially motivated by the desire to understand limits of semitoric systems which are themselves not semitoric. One method to do this is to study the elements of $\tilde{\mathbb{I}} \setminus \mathbb{I}$ in relation to integrable systems. Perhaps some subset of this can be interpreted as corresponding to non-simple semitoric systems or to some other type of integrable system not included in the classification by Pelayo-Vũ Ngọc [69, 70]. Problem 2.44 from [72] asks if some integrable systems may be expressed as the limit of semitoric systems in an appropriate topology and the study of $\tilde{\mathbb{I}} \setminus \mathbb{I}$ may make some progress on this question.

Acknowledgements. Chapter 3, in part, is comprised of material submitted for publication

by the author of this dissertation as *Moduli spaces of semitoric systems*, currently available at arXiv:1502.07296 [63].

Chapter 4

Classifying toric and semitoric fans

4.1 Introduction

In this chapter I present an algebraic viewpoint to study four-dimensional toric integrable systems, based on the study of matrix relations in the special linear group $\mathrm{SL}_2(\mathbb{Z})$. Recall, one can associate to a rational convex Delzant polygon Δ the collection of primitive integer inwards pointing normal vectors to its faces, called a *toric fan*. This is a d -tuple

$$(v_0 = v_d, v_1, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d,$$

where $d \in \mathbb{Z}$ is the number of faces and $\det(v_i, v_{i+1}) = 1$ for each $0 \leq i \leq d-1$. This determinant condition forces the vectors to satisfy the linear equations $a_i v_i = v_{i-1} + v_{i+1}$, for $i = 0, \dots, d-1$ where $v_{-1} = v_{d-1}$, which are parameterized by integers $a_0, \dots, a_{d-1} \in \mathbb{Z}$ (Lemma 4.3.3). These integers satisfy

$$\begin{pmatrix} 0 & -1 \\ 1 & a_0 \end{pmatrix} \begin{pmatrix} 0 & -1 \\ 1 & a_1 \end{pmatrix} \cdots \begin{pmatrix} 0 & -1 \\ 1 & a_{d-1} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad (4.1)$$

(this equation appears in [33], page 44) but in fact not all integers satisfying Equation (4.1) correspond to a Delzant polygon. This is an equation in $\mathrm{SL}_2(\mathbb{Z})$ and in this chapter I lift it to the group G presented as

$$G = \langle S, T \mid STS = T^{-1}ST^{-1} \rangle$$

where $\mathrm{SL}_2(\mathbb{Z}) \cong G/(S^4)$. The group G is the pre-image of $\mathrm{SL}_2(\mathbb{Z})$ in the universal cover of $\mathrm{SL}_2(\mathbb{R})$, and thus I can define what I call the *winding number* of an element of $g \in G$ that evaluates to the identity in $\mathrm{SL}_2(\mathbb{Z})$. Roughly speaking, we view g as a word in S and T and by applying this word to a vector one term at a time I produce a path around the origin. I define the winding number of g to be the winding number of this path in the classical sense.

The equation in G analogous to Equation (4.1) has the property that a collection of integers a_0, \dots, a_{d-1} satisfy the equation if and only if they correspond to a toric fan, and from the integers it is straightforward to recover the fan. As such, this is a method to translate problems about toric fans into equivalent problems about the algebraic structure of the group G . This allows us to simplify the proofs of some classical results about toric fans (Section 4.4), and generalize these results to the semitoric case (Section 4.5). Associated to a semitoric system there is also a collection of vectors (v_0, \dots, v_{d-1}) which satisfy more complicated equations (given explicitly in Definition 4.2.5) known as a semitoric fan. A semitoric fan can be thought of as a toric fan for which the relations between some pairs of adjacent vectors have been changed as a result of the presence of the focus-focus singularities of semitoric integrable systems. Roughly speaking, semitoric fans encode aspects of the singular affine structure induced by the singular fibration associated to a semitoric integrable system.

I present the following theorem in this chapter which is an application of the algebraic method I introduce. In Chapter 5 I will use this theorem to describe the path-connected components of the moduli space of semitoric integrable systems with a fixed number of focus-focus singular points.

Theorem 4.1.1. *Any semitoric fan may be obtained from a standard semitoric fan in a finite number of steps using four standard transformations.*

A more detailed description of Theorem 4.1.1 is given in Theorem 4.2.8 and the definitions of the standard semitoric fans and the four standard transformations used in Theorem 4.1.1 are given in Definition 4.2.7.

The machinery in this chapter is developed with the goal of proving Theorem 5.1.1 in the next chapter, which describes the connected components of the moduli space of semitoric systems from Chapter 3.

4.2 Fans, symplectic geometry, and winding numbers

4.2.1 Toric fans

A toric variety is a variety which contains an algebraic torus as a dense open subset such that the standard action of the torus on itself can be extended to the whole variety. That is, a toric variety is the closure of an algebraic torus orbit [56]. By an algebraic torus we mean the product $\mathbb{C}^* \times \dots \times \mathbb{C}^*$, where $\mathbb{C}^* = \mathbb{C} \setminus \{0\}$. It is well known that the geometry of a toric variety is completely classified by the associated fan. In general, a fan is set of rational strongly convex cones in a real vector space such that the face of each cone is also a cone and the intersection of any two cones is a face of each. In this chapter we will be concerned with two-dimensional nonsingular complete toric varieties and their associated fans, which we will simply call toric fans. As described in Section 2.2.6, these fans are given by a sequence of lattice points

$$(v_0 = v_d, v_1, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$$

labeled in counter-clockwise order such that each pair of adjacent vectors generates all of \mathbb{Z}^2 and the angle between any two adjacent vectors is less than π radians. That is, $\det(v_i, v_{i+1}) = 1$ for $i = 0, \dots, d - 1$.

Recall a Delzant polygon is a convex polygon Δ in \mathbb{R}^2 which is simple, rational, and smooth. A toric fan may be produced from a Delzant polygon considering the collection of inwards pointing normal vectors of the polygon.

Here I adapt Theorem 2.2.16 to relate it to Delzant polygons.

Theorem 4.2.1 (Fulton [33], page 44). *Up to the action of $\mathrm{SL}_2(\mathbb{Z})$, every Delzant polygon can be obtained from a Delzant triangle, rectangle, or Hirzebruch trapezoid by a finite number of corner chops.*

The minimal models (the Delzant triangle, rectangle, and Hirzebruch trapezoid) are defined in Definition 4.4.9 and depicted in Figure 4.1. The corner chop operation is defined in Definition 4.2.7 and is a standard operation in algebraic and symplectic geometry (see, for instance, [49]). The proof of Theorem 4.2.1 sketched by Fulton in [33], which uses only two-dimensional geometry and




Figure 4.1: (a) The three minimal models from Theorem 4.2.1. (b) An illustration of a Delzant polygon produced by corner chopping the Hirzebruch trapezoid.

basic combinatorial arguments, is relatively long and does not immediately generalize. In Section 4.4 we provide an alternative proof using $\mathrm{SL}_2(\mathbb{Z})$ -relations. This proof may be easily extended to the semitoric case.

As a consequence of Theorem 4.2.1 in [67] it was recently proved:

Theorem 4.2.2 ([67]). *The moduli space of toric polygons is path-connected.*

That is, any two toric polygons may be deformed onto each other continuously via a path of toric polygons. One shows this by first knowing how to generate all toric polygons as in Theorem 4.2.1. Then one shows, using elementary analysis, that the four minimal models can be continuously transformed into one another and that the corner chop operation is continuous. Again, let us emphasize that the results about toric fans are not new, we have just used a new viewpoint to arrive at classical results. We will see that this new viewpoint allows us to generalize the known results.

4.2.2 Semitoric fans

In analogy with semitoric polygons (originally defined in [70, Definition 2.5]) we define semitoric fans. Recall

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}).$$

Definition 4.2.3. Let $v, w \in \mathbb{Z}^2$. The ordered pair (v, w) of vectors:

1. is *on the top boundary* if both vectors are in the open lower half-plane;
2. satisfies the *Delzant condition* if $\det(v, w) = 1$;
3. satisfies the *hidden condition* if $\det(v, Tw) = 1$; and
4. satisfies the *fake condition* if $\det(v, Tw) = 0$.

Remark 4.2.4. Notice that a pair (v, w) satisfies both the fake and Delzant conditions if and only if

$$v = \begin{pmatrix} k + \varepsilon \\ \varepsilon \end{pmatrix} \text{ and } w = \begin{pmatrix} k \\ \varepsilon \end{pmatrix}$$

for some $k \in \mathbb{Z}$ and $\varepsilon \in \{-1, +1\}$ and in order for such a pair to be in the top boundary we can only have the case in which $\varepsilon = -1$. ◊

Definition 4.2.5. Let $d \in \mathbb{Z}$ with $d > 2$. A *semitoric fan* is a collection of primitive vectors $(v_0 = v_d, v_1, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ labeled in counter-clockwise order such that each pair of adjacent vectors (v_i, v_{i+1}) for $i \in \{0, \dots, d-1\}$ is labeled as a *Delzant*, *fake*, or *hidden* corner. We require that each labeled pair of vectors satisfies the corresponding condition from Definition 4.2.3 and we further require that all fake and hidden corners be on the top boundary. The *complexity* of a semitoric fan is the number of corners which are either fake or hidden.

Notice that the labeling of the pairs is required only because of the case described in Remark 4.2.4 in which a pair can satisfy both the fake and Delzant conditions. In all other cases the corner type of a pair of vectors can be uniquely determined by inspecting the vectors involved.

Definition 4.2.5 is inspired by the toric case. Theorem 4.2.1 states that any toric fan can be produced from a minimal model using only corner chops. Similarly, our goal is to use a series of transformations to relate any semitoric fan to a standard form up to the action of the appropriate symmetry group.

Definition 4.2.6. The *symmetry group of semitoric fans* is given by

$$\mathcal{G}' = \{T^k \mid k \in \mathbb{Z}\}$$

where \mathcal{G}' acts on a semitoric fan by acting on each vector in the fan.

Definition 4.2.7.

1. Let $c \in \mathbb{Z}_{\geq 0}$. The *standard semitoric fan of complexity c* is the fan $(u_0, \dots, u_{c+3}) \in (\mathbb{Z}^2)^{c+4}$

given by

$$u_0 = \begin{pmatrix} 0 \\ -1 \end{pmatrix}, u_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, u_2 = \begin{pmatrix} c \\ 1 \end{pmatrix}, u_3 = \begin{pmatrix} -1 \\ 0 \end{pmatrix},$$

and

$$u_{4+n} = \begin{pmatrix} -c+n \\ -1 \end{pmatrix}$$

for $n = 0, \dots, c-1$ in which the first four pairs of vectors are Delzant corners and the rest are fake corners.

2. Let $(v_0 = v_d, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ be a semitoric fan. The following are called the *four fan transformations*:

- (a) Suppose that (v_i, v_{i+1}) is a Delzant corner for some $i \in \{0, \dots, d-1\}$. Then

$$(v_0, \dots, v_i, v_i + v_{i+1}, v_{i+1}, \dots, v_{d-1}) \in (\mathbb{Z}^2)^{d+1}$$

obtained by inserting the sum of two adjacent vectors between them. The new pairs $(v_i, v_i + v_{i+1})$ and $(v_i + v_{i+1}, v_{i+1})$ are both Delzant corners. The process of producing this new fan from the original is known as *corner chopping* [49].

- (b) A *reverse corner chop* is the procedure by which a single vector v_i , $i \in \{0, \dots, d-1\}$, which is the sum of its adjacent vectors is removed from the fan. This is the inverse of a corner chop and can only be performed if both pairs in the original fan involving v_i are Delzant corners.

- (c) Suppose that the pair (v_i, v_{i+1}) is a hidden corner. Then

$$(v_0, \dots, v_i, Tv_{i+1}, v_{i+1}, \dots, v_{d-1}) \in (\mathbb{Z}^2)^{d+1}$$

is a semitoric fan with (v_i, Tv_{i+1}) a Delzant corner and (Tv_{i+1}, v_{i+1}) a fake corner (Lemma

4.5.1). The process of producing this fan is known as *removing the hidden corner* (v_i, v_{i+1}) .

- (d) Suppose that the pair (v_i, v_{i+1}) is a fake corner and the pair (v_{i+1}, v_{i+2}) is a Delzant corner. Then

$$(v_0, \dots, v_i, Tv_{i+2}, v_{i+2}, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$$

is a semitoric fan with (v_i, Tv_{i+2}) a Delzant corner and (Tv_{i+2}, v_{i+2}) a fake corner (Lemma 4.5.2). The process of producing this fan is known as *commuting a fake and a Delzant corner*.

Using the algebraic results from Section 4.3 we show the following.

Theorem 4.2.8. *Let $d \geq 1$ be an integer. Any semitoric fan $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ of complexity $c \in \mathbb{Z}_{\geq 0}$ may be transformed into a semitoric fan \mathcal{G}' -equivalent to the standard semitoric fan of complexity c by using the four fan transformations.*

Remark 4.2.9. The method we are using to study semitoric integrable systems is analogous to the method we use to study toric integrable systems. Theorem 4.2.1 explains how to generate the Delzant polygons and is used to prove that the space of Delzant polygons is path-connected (Theorem 4.2.2) which implies that the space of toric integrable systems is connected. Similarly, Theorem 4.2.8 shows how to generate the semitoric polygons, and as an application we prove Lemma 5.3.5 which describes the connected components in the space of semitoric ingredients (Definition 2.2.26) and this implies Theorem 5.1.1, which describes the connected components of the moduli space of semitoric systems. ◊

4.2.3 Algebraic tools: the winding number

It is shown in Lemma 4.3.1 that the special linear group $\mathrm{SL}_2(\mathbb{Z})$ may be presented as

$$\mathrm{SL}_2(\mathbb{Z}) = \langle S, T \mid T^{-1}ST^{-1} = STS, S^4 = I \rangle$$

where

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \text{ and } T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

Thus Equation (4.1) becomes

$$ST^{a_0} \dots ST^{a_{d-1}} = I$$

where $a_0, \dots, a_{d-1} \in \mathbb{Z}$ and I denotes the 2×2 identity matrix. Given $v_0, v_1 \in \mathbb{Z}^2$ with $\det(v_1, v_2) = 1$ a set of vectors

$$(v_0, v_1, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$$

may be produced by

$$v_{i+2} = -v_i + a_i v_{i+1}$$

for $i = 0, \dots, d-1$ where we define $v_d = v_0$ and $v_{d+1} = v_1$. In this way associated to each list of integers satisfying Equation (4.1) there is an ordered collection of vectors unique up to $\mathrm{SL}_2(\mathbb{Z})$. It can be seen that the determinant between any adjacent pair of these vectors is one and thus if these vectors are labeled in counter-clockwise order, then they are a toric fan. The reason that not all sequences of integers which satisfy Equation (4.1) correspond to a toric fan is that the vectors $v_0, \dots, v_{d-1} \in \mathbb{Z}^2$ may circle more than once around the origin, and thus not be labeled in counter-clockwise order (see Figure 4.2). Thus, we see that viewing $ST^{a_0} \dots ST^{a_{d-1}}$ as an element of $\mathrm{SL}_2(\mathbb{Z})$ is losing too much information.

Let $K = \ker(\langle S, T \rangle \rightarrow \mathrm{SL}_2(\mathbb{Z}))$ where $\langle S, T \rangle$ denotes the free group with generators S and T and the map $\langle S, T \rangle \rightarrow \mathrm{SL}_2(\mathbb{Z})$ is the natural projection. For any word in K a sequence of vectors may be produced by letting the word act on a vector $v \in \mathbb{Z}^2$ one term at a time. We know we will end back at v , but the sequence of vectors produced contains more information about the word. This sequence can be used to define a path in $\mathbb{R}^2 \setminus \{(0, 0)\}$. Of particular interest, especially when studying toric and semitoric fans, is the *winding number* of such a path. That is, the number of times the path, and hence the collection of vectors, circles the origin. This construction is explained in detail in Section 4.4, and in particular Definition 4.4.4 given a precise definition of the number of times an ordered collection of vectors circles the origin. Let $w : \langle S, T \rangle \rightarrow \mathbb{Z}$ be given by

$$w(\sigma) = 3s - t$$

for $\sigma \in \langle S, T \rangle$ where s is the number of appearances of S in σ and t is the number of appearances

of T in σ . We find that if $\sigma \in K$ then $w(\sigma)$ is a multiple of 12 and $w(\sigma)/12$ is the winding number associated to the word σ . We present the group

$$G = \langle S, T \mid STS = T^{-1}ST^{-1} \rangle$$

on which the w descends to a well-defined function $w_G : G \rightarrow \mathbb{Z}$. In fact, G is isomorphic to the pre-image of $\mathrm{SL}_2(\mathbb{Z})$ in the universal cover of $\mathrm{SL}_2(\mathbb{R})$ (Proposition 4.3.7). Thus, if K' is the image of K projected to G , then given some $g \in K'$ there is an associated closed loop in $\mathrm{SL}_2(\mathbb{R})$. The fundamental group of $\mathrm{SL}_2(\mathbb{R})$ is \mathbb{Z} and the classical winding number of this loop in $\mathrm{SL}_2(\mathbb{R})$ coincides with $W(\sigma)$ for any $\sigma \in \langle S, T \rangle$ which projects to g . Finally, in Corollary 4.4.6 we show that integers $a_0, \dots, a_{d-1} \in \mathbb{Z}$ correspond to a toric fan if and only if the equality

$$ST^{a_0} \dots ST^{a_{d-1}} = S^4$$

is satisfied in G . This correspondence is the basis of our method to study toric and semitoric fans.

4.3 Algebraic set-up: matrices and $\mathrm{SL}_2(\mathbb{Z})$ relations

The 2×2 special linear group over the integers, $\mathrm{SL}_2(\mathbb{Z})$, is generated by the matrices

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \text{ and } T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

We will see that to each toric (resp. semitoric) integrable system there is an associated toric (resp. semitoric) fan and we will use the algebraic structure of $\mathrm{SL}_2(\mathbb{Z})$ to study these fans. For our purposes, the following presentation of $\mathrm{SL}_2(\mathbb{Z})$ will be the most natural way to view the group.

Lemma 4.3.1. *The 2×2 special linear group over the integers, $\mathrm{SL}_2(\mathbb{Z})$, may be presented as*

$$\mathrm{SL}_2(\mathbb{Z}) = \langle S, T \mid T^{-1}ST^{-1} = STS, S^4 = I \rangle.$$

Proof. It is well-known that

$$\mathrm{SL}_2(\mathbb{Z}) = \langle S, T \mid (ST)^3 = S^2 = -I \rangle \quad (4.2)$$

(see for instance [12]) where $(-I)^2 = I$ and $-I$ is in the center of $\mathrm{SL}_2(\mathbb{Z})$. To obtain the relations in the statement of the lemma from those in Equation (4.2) notice

$$(ST)^3 = S^2 \Leftrightarrow STS = T^{-1}ST^{-1} \text{ and } S^2 = -I \Rightarrow S^4 = I.$$

To obtain the relations of Equation (4.2) from those in the statement of the lemma we only have to show that S^2 squares to the identity and is in the center of the group to conclude that $S^2 = I$. We have that $(S^2)^2 = I$ and that S^2 commutes with S . To show that S^2 commutes with T notice

$$\begin{aligned} TS^2 &= TS(TSTST) \\ &= (TSTST)ST \\ &= S^2T. \end{aligned}$$

This concludes the proof. □

For $v, w \in \mathbb{Z}^2$ let $[v, w]$ denote the 2×2 matrix with v as the first column and w as the second and let $\det(v, w)$ denote the determinant of the matrix $[v, w]$.

Lemma 4.3.2. *Let $u, v, w \in \mathbb{Z}^2$ and $\det(u, v) = 1$. Then $\det(v, w) = 1$ if and only if there exists some $a \in \mathbb{Z}$ such that $w = -u + av$.*

Proof. In the basis (u, v) we know that $v = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. Write $w = \begin{pmatrix} b \\ a \end{pmatrix}$ for some $a, b \in \mathbb{Z}$. Then we can see that $\det(v, w) = -b$ so $\det(v, w) = 1$ if and only if $b = -1$. That is, $w = -u + av$. □

The result of Lemma 4.3.2 can be easily summarized in a matrix equation, as we will now show. Let

$$(v_0 = v_d, v_1 = v_{d+1}, v_2, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$$

be a toric fan and define $A_i = [v_i, v_{i+1}]$ for $i = 0, \dots, d$. Note that $A_d = A_0$.

Lemma 4.3.3. *For each $i \in 0, \dots, d-1$ there exists an integer $a_i \in \mathbb{Z}$ such that $A_{i+1} = A_i ST^{a_i}$.*

Proof. By the definition of a toric fan we know that for each $0 \leq i < d-2$ we have that

$$\det(v_i, v_{i+1}) = \det(v_{i+1}, v_{i+2}) = 1$$

so by Lemma 4.3.2 there exists $a_i \in \mathbb{Z}$ such that $v_{i+2} = -v_i + a_i v_{i+1}$. Then

$$\begin{aligned} A_i ST^{a_i} &= [v_{i+1}, -v_i + a_i v_i] \\ &= [v_{i+1}, v_{i+2}] \\ &= A_{i+1}, \end{aligned}$$

and this concludes the proof. □

It follows that

$$A_d = A_{d-1} ST^{a_{d-1}} = A_{d-2} ST^{a_{d-2}} ST^{a_{d-1}} = \dots = A_0 ST^{a_0} \dots ST^{a_{d-1}}$$

which means

$$A_0 = A_d = A_0 ST^{a_0} \dots ST^{a_{d-1}},$$

and so

$$ST^{a_0} \dots ST^{a_{d-1}} = I. \tag{4.3}$$

This is a restatement of Equation (4.1) which is from [33]. So to each toric fan of d vectors there is an associated d -tuple of integers which satisfy Equation (4.3), but having a tuple of integers which satisfy Equation (4.3) is not enough to assure that they correspond to an toric fan. The determinant of the vectors will be correct but, roughly speaking, if the vectors wind around the origin more than once then they will not be labeled in the correct order to be a toric fan, as it occurs in the following example.

Example 4.3.4. Consider the sequence of integers $a_0 = -1, a_1 = -1, a_2 = -2, a_3 = -1, a_4 =$




Figure 4.2: These vectors do not form a fan because they are not labeled in counter-clockwise order.

$-1, a_5 = 0$ and notice that

$$ST^{-1}ST^{-1}ST^{-2}ST^{-1}ST^{-1}ST^0 = I$$

so Equation (4.3) is satisfied for $c = 0$ but these integers do not correspond to a toric fan. This is because the vectors they produce:

$$v_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, v_2 = \begin{pmatrix} -1 \\ -1 \end{pmatrix},$$

$$v_3 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_4 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}, v_5 = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$$

travel twice around the origin¹, see Figure 4.2. ⊙

So we need extra information that is not captured by viewing this word in S and T as an element of $\mathrm{SL}_2(\mathbb{Z})$. For a more obvious example notice that even though they are equal in $\mathrm{SL}_2(\mathbb{Z})$ we can see that S^4 corresponds to a toric fan while S^8 does not. From [33] we know that integers $(a_0, \dots, a_d) \in \mathbb{Z}^d$ which satisfy Equation (4.1) correspond to a toric fan if and only if

$$a_0 + \dots + a_{d-1} = 3d - 12$$

¹For a formal definition of this, see Definition 4.4.4.

so we would like to prove that

$$\frac{3d - \sum_{i=0}^{d-1} a_i}{12}$$

is the number of times that the vectors corresponding to (a_0, \dots, a_{d-1}) circle the origin. In order to prove this we will need some more terminology, and in order to keep track of the extra information about circling the origin we will need to consider a group which is larger than $\mathrm{SL}_2(\mathbb{Z})$.

Consider instead the free group with generators S and T , denoted $\langle S, T \rangle$. We know $\mathrm{SL}_2(\mathbb{Z})$ is a quotient of $\langle S, T \rangle$ by Lemma 4.3.1 so there exists a natural projection map $\pi_1 : \langle S, T \rangle \rightarrow \mathrm{SL}_2(\mathbb{Z})$. Also define a map $w : \langle S, T \rangle \rightarrow \mathbb{Z}$ by

$$w(S^{b_0}T^{a_0} \dots S^{b_\ell}T^{a_\ell}) = 3 \sum_{i=0}^{\ell} b_i - \sum_{i=0}^{\ell} a_i$$

for any $S^{b_0}T^{a_0} \dots S^{b_\ell}T^{a_\ell} \in \langle S, T \rangle$ where $a_0, \dots, a_\ell, b_0, \dots, b_\ell \in \mathbb{Z}$. Given a toric fan with associated integers $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ we will show that $w(ST^{a_0} \dots ST^{a_{d-1}}) = 12$. Both π_1 and w factor over the same group G which is the fiber product of $\mathrm{SL}_2(\mathbb{Z})$ and \mathbb{Z} over $\mathbb{Z}/(12)$. Now we can see that we wanted the particular presentation of $\mathrm{SL}_2(\mathbb{Z})$ from Lemma 4.3.1 so that the relationship between G and $\mathrm{SL}_2(\mathbb{Z})$ would be clear. This discussion is made precise in the following proposition.

Proposition 4.3.5. *The following diagram commutes.*

$$\begin{array}{ccc}
 & \langle S, T \rangle & \\
 \pi_1 \swarrow & & \searrow \pi_2 \\
 \mathrm{SL}_2(\mathbb{Z}) & \xleftarrow{\pi_3} & G \\
 w_{\mathrm{SL}_2(\mathbb{Z})} \downarrow & & \downarrow w_G \\
 \mathbb{Z}/(12) & \xleftarrow{\pi_4} & \mathbb{Z}
 \end{array}
 \quad (4.4)$$

The group G is the fiber product of $\mathrm{SL}_2(\mathbb{Z})$ and \mathbb{Z} over $\mathbb{Z}/(12)$ and is given by

$$G = \langle S, T \mid STS = T^{-1}ST^{-1} \rangle,$$

each of π_1, π_2, π_3 , and π_4 is a projection, and $w : \langle S, T \rangle \rightarrow \mathbb{Z}$, $w_G : G \rightarrow \mathbb{Z}$, $w_{\mathrm{SL}_2(\mathbb{Z})} : \mathrm{SL}_2(\mathbb{Z}) \rightarrow \mathbb{Z}/(12)$

are given by the same formal expression

$$S^{b_0}T^{a_0} \dots S^{b_\ell}T^{a_\ell} \mapsto 3 \sum_{i=0}^{\ell} b_i - \sum_{i=0}^{\ell} a_i. \quad (4.5)$$

Proof. It can be seen that the map $w_{\mathrm{SL}_2(\mathbb{Z})} : \mathrm{SL}_2(\mathbb{Z}) \rightarrow \mathbb{Z}/(12)$ is well-defined by noting that both relations in $\mathrm{SL}_2(\mathbb{Z})$ as presented in Lemma 4.3.1 preserve the value of the formula (Equation (4.5)) up to a multiple of 12. Similarly, since the relation $ST^{-1}S = STS$ preserves the value of the equation we know that w_G is well-defined. Since each of these functions to \mathbb{Z} or $\mathbb{Z}/(12)$ is given by the same formal expression and since each π is a quotient map, the diagram commutes.

To show that G with the associated maps is the fiber product of $\mathrm{SL}_2(\mathbb{Z})$ and \mathbb{Z} over $\mathbb{Z}/(12)$ we must only show that w_G restricted to the fibers is bijective. That is, we must show that

$$w_G \upharpoonright_{\pi_3^{-1}(A)} : \pi_3^{-1}(A) \rightarrow \pi_4^{-1}(w_{\mathrm{SL}_2(\mathbb{Z})}(A))$$

is a bijection for each $A \in \mathrm{SL}_2(\mathbb{Z})$. To show it is surjective, notice that for any $g \in G$

$$\pi_3(S^{4k}g) = \pi_3(g) \text{ and } w_G(S^{4k}g) = w_G(g) + 12k$$

for any $k \in \mathbb{Z}$. To show it is injective it is sufficient to consider only $A = I$. Since S^4 is in the center of G we know that $\mathrm{SL}_2(\mathbb{Z}) = G/(S^4)$ so $\pi_3^{-1}(I) = \{S^{4k} \mid k \in \mathbb{Z}\}$. Since $w_G(S^{4k}) = 12k$ we know for each choice of k this maps to a distinct element of \mathbb{Z} .

□

Notation 4.3.6. We have several groups with generators S and T . To denote the different equalities in these groups we will use an equal sign with the group in question as a subscript. That is, if an equality holds in the group H we will write $=_H$. For example, $S^4 =_{\mathrm{SL}_2(\mathbb{Z})} I$ but $S^4 \neq_G I$.

There is another useful sense in which G is an unwinding of $\mathrm{SL}_2(\mathbb{Z})$. While $\mathrm{SL}_2(\mathbb{Z})$ is discrete, and thus does not have a natural cover, it sits inside the group $\mathrm{SL}_2(\mathbb{R})$, which has a universal cover. We claim that G is the preimage of $\mathrm{SL}_2(\mathbb{Z})$ inside of the universal cover of $\mathrm{SL}_2(\mathbb{R})$.

Proposition 4.3.7. *The group G is isomorphic to the preimage of $\mathrm{SL}_2(\mathbb{Z})$ within the universal cover of $\mathrm{SL}_2(\mathbb{R})$.*

Proof. Let G' be the preimage of $\mathrm{SL}_2(\mathbb{Z})$ in the universal cover of $\mathrm{SL}_2(\mathbb{R})$. We note that there exists a homomorphism, ϕ from G to G' defined by

$$\phi(S) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}_{0 \leq \theta \leq \pi/2}, \quad \phi(T) = \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}_{0 \leq x \leq 1}$$

where the paths given are to represent elements of the universal cover of $\mathrm{SL}_2(\mathbb{R})$. It is easy to show that $\phi(S)\phi(T)\phi(S)$ equals $\phi(T)^{-1}\phi(S)\phi(T)^{-1}$, and thus ϕ does actually define a homomorphism. We have left to show that ϕ defines an isomorphism. To show this, we note that each of G and G' have obvious surjections π and π' to $\mathrm{SL}_2(\mathbb{Z})$. Furthermore it is clear that $\pi = \pi' \circ \phi$. Thus, to show that ϕ is an isomorphism, it suffices to show that $\phi : \ker(\pi) \rightarrow \ker(\pi')$ is an isomorphism.

However, it is clear that $\ker(\pi)$ is $\langle S^4 \rangle$. On the other hand, $\ker(\pi') \cong \pi_1(\mathrm{SL}_2(\mathbb{R})) = \mathbb{Z}$, and is generated by $\phi(S)^4$. This completes the proof. \square

We will see that there is a one to one correspondence between toric fans up to the action of $\mathrm{SL}_2(\mathbb{Z})$ and lists of integers $a_0, \dots, a_{d-1} \in \mathbb{Z}$ satisfying

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^4 \tag{4.6}$$

(Corollary 4.4.6). Equation (4.6) is a refinement of Equation (4.3) which implies both that the successive pairs of vectors form a basis of \mathbb{Z}^2 and that the vectors are labeled in counter-clockwise order. In Proposition 4.5.4 we produce an analogous equation for semitoric fans.

Now we would like to simplify these toric fans. We will understand which integers a_0, \dots, a_{d-1} are possible in an element $ST^{a_0} \dots ST^{a_{d-1}} \in G$ corresponding to toric fan by studying $\mathrm{PSL}_2(\mathbb{Z}) = \mathrm{SL}_2(\mathbb{Z})/(-I)$. The following lemma is important for this and will also be useful later on when classifying semitoric fans. If $ST^{a_0} \dots ST^{a_{d-1}} \in G$ projects to the identity in $\mathrm{SL}_2(\mathbb{Z})$ then by Lemma 4.3.8 we can see when one of the exponents must be in the set $\{-1, 0, 1\}$. In any of these cases, we will be able to use relations in G to help simplify the expression.

Lemma 4.3.8. *Suppose that*

$$ST^{a_0} \dots ST^{a_{d-1}} =_{\mathrm{PSL}_2(\mathbb{Z})} I \tag{4.7}$$

for some $d \in \mathbb{Z}$, $d > 0$. Then if $d \geq 3$ there exist $i, j \in \mathbb{Z}$ satisfying $0 \leq i < j \leq d - 1$ such that $a_i, a_j \in \{-1, 0, 1\}$. Furthermore:

1. If $d > 3$ then i, j can be chosen such that $i \neq j - 1$ and $(i, j) \neq (0, d - 1)$.
2. If $d = 3$ then $a_0 = a_1 = a_2 = 1$ or $a_0 = a_1 = a_2 = -1$.
3. If $d = 2$ then $a_0 = a_1 = 0$.
4. If $d = 1$ then Equation (4.7) cannot hold.

Note that part (1) is the statement that i and j are not consecutive in the cyclic group $\mathbb{Z}/(d)$. Of course, it is clear that such i and j may be chosen if three or more elements of the list a_0, \dots, a_{d-1} are in the set $\{-1, 0, 1\}$.

Proof. It is well known that $\text{PSL}_2(\mathbb{Z})$ acts on the real projective line $\mathbb{R} \cup \{\infty\}$ by linear fractional transformations:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} (x) = \frac{ax + b}{cx + d} \text{ for } x \in \mathbb{R} \text{ and } \begin{pmatrix} a & b \\ c & d \end{pmatrix} (\infty) = \frac{a}{c}.$$

Let $d > 3$. Suppose that at most two of a_0, \dots, a_{d-1} are in $\{-1, 0, 1\}$ and if there are two in $\{-1, 0, 1\}$ that they are consecutive or indexed by 0 and $d - 1$. Notice that

$$ST^{a_0} \dots ST^{a_{d-1}} =_{\text{PSL}_2(\mathbb{Z})} I \text{ implies that } ST^{a_1} \dots ST^{a_{d-1}} ST^{a_0} =_{\text{PSL}_2(\mathbb{Z})} I$$

by conjugating each side with ST^{a_0} . This conjugation method and renumbering the integers can be used to assure that $a_i \notin \{-1, 0, 1\}$ for $i = 1, \dots, d - 2$. Since this expression is equal to the identity in $\text{PSL}_2(\mathbb{Z})$ it acts trivially on $\mathbb{R} \cup \{\infty\}$. In particular, we have

$$ST^{a_0} \dots ST^{a_{d-1}}(\infty) = \infty.$$

Notice that $S(x) = -1/x$ and $T^a(x) = x + a$ for $a \in \mathbb{Z}$. Further notice that for any $a \in \mathbb{Z} \setminus \{-1, 0, 1\}$ and $x \in (-1, 1) \setminus \{0\}$ we have $ST^a(x) \in (-1, 1) \setminus \{0\}$. We see that $ST^{a_{d-1}}(\infty) = 0$ and since

$a_{d-2} \notin \{-1, 0, 1\}$ we know $ST^{a_{d-2}}(0) \in (-1, 1) \setminus \{0\}$. Putting these facts together we have

$$\begin{aligned}
ST^{a_0} \dots ST^{a_{d-1}}(\infty) &= ST^{a_0} \dots ST^{a_{d-2}}(0) \\
&= ST^{a_0} \dots ST^{a_{d-3}}(x) && \text{for some } x \in (-1, 1) \setminus \{0\} \\
&= ST^{a_0}(y) && \text{for some } y \in (-1, 1) \setminus \{0\} \\
&= \frac{-1}{y + a_0} \neq \infty
\end{aligned}$$

This contradiction finishes the $d > 3$ case.

If $d = 3$ essentially the same result holds except that it is not possible to choose two elements that are non-consecutive. Notice

$$ST^{a_0} ST^{a_1} ST^{a_2}(\infty) = ST^{a_0} ST^{a_1}(0) = -\left(\frac{-1}{a_1} + a_0\right)^{-1}.$$

For this function to be the identity we would need

$$\frac{-1}{a_1} + a_0 = 0.$$

This implies that $a_0 a_1 = 1$ so since they are both integers we have $a_0 = a_1 = \epsilon$ where $\epsilon \in \{-1, 1\}$.

Conjugating by ST^{a_0} , we find symmetrically, that $a_1 a_2 = 1$, and thus that $a_0 = a_1 = a_2 = \pm 1$.

If $d = 2$ then we have

$$ST^{a_0} ST^{a_1}(\infty) = ST^{a_0}(0) = \frac{-1}{a_0}$$

so we must have $a_0 = 0$ and then

$$ST^0 ST^{a_1}(x) = S^2 T^{a_1}(x) = x + a_1,$$

so we are also forced to have that $a_1 = 0$, as stated in the lemma. If $d = 1$ then $ST^{a_0}(\infty) = 0$ for any choice of $a_0 \in \mathbb{Z}$, so there are no solutions. \square

4.4 Toric fans

Let $a_0, a_1, \dots, a_{d-1} \in \mathbb{Z}$ be a collection of integers such that

$$ST^{a_0} \dots ST^{a_{d-1}} =_{\mathrm{SL}_2(\mathbb{Z})} I.$$

This means that

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^{4k} \text{ for some } k \in \mathbb{Z}$$

by Proposition 4.3.5. We claim that these integers correspond to a toric fan if and only if $k = 1$.

The idea is that $k = 1$ precisely when the vectors in the corresponding fan are labeled in counter-clockwise order, and the only relation in the group G , which is $STS = T^{-1}ST^{-1}$, preserves the number of times the vectors circle the origin. Now we make this idea precise.

Lemma 4.4.1. *Let $g \in \ker(\pi_3)$. Then $\frac{w_G(g)}{12} \in \mathbb{Z}$.*

Proof. Since $\pi_3(g) = I$ we know $w_{\mathrm{SL}_2(\mathbb{Z})} \circ \pi_3(g) = 0$, so by Proposition 4.3.5 $\pi_4 \circ w_G(g) = 0$. Thus $w_G(g) \in \ker(\pi_4) = \{12k \mid k \in \mathbb{Z}\}$. \square

Recall that G is isomorphic to the preimage of $\mathrm{SL}_2(\mathbb{Z})$ in the universal cover of $\mathrm{SL}_2(\mathbb{R})$ by Proposition 4.3.7. Let ϕ be the isomorphism from G to its image in the universal cover of $\mathrm{SL}_2(\mathbb{R})$ with

$$\phi(S) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}_{0 \leq \theta \leq \pi/2}, \quad \phi(T) = \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}_{0 \leq x \leq 1}$$

This means to each element of the kernel of π_3 we can associate a closed loop based at the identity in $\mathrm{SL}_2(\mathbb{R})$ denoted $\phi(g)$. The fundamental group $\pi_1(\mathrm{SL}_2(\mathbb{R}))$ is isomorphic to \mathbb{Z} and is generated as $\langle \phi(S^4) \rangle$, so let $\psi : \pi_1(\mathrm{SL}_2(\mathbb{R})) \rightarrow \mathbb{Z}$ be the isomorphism with $\psi(\phi(S^4)) = 1$.

Lemma 4.4.2. *Let $g \in \ker(\pi_3 : G \rightarrow \mathrm{SL}_2(\mathbb{Z}))$. Then*

$$\psi \circ \phi(g) = \frac{w_G(g)}{12}.$$

Proof. Since $\ker(\pi_3)$ is generated by S^4 , it suffices to check that $\psi(\phi(S^4)) = \frac{w_G(S^4)}{12} = 1$, but this holds by definition.

□

Definition 4.4.3. Define $W : \ker(\pi_3) \rightarrow \mathbb{Z}$ by

$$W(g) = \frac{w_G(g)}{12}.$$

We call $W(g)$ the *winding number* of $g \in \ker(\pi_3)$.

Definition 4.4.4. Let

$$(v_0 = v_d, v_1, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$$

with $\det(v_i, v_{i+1}) > 0$ for $i = 0, \dots, d-1$. We define *the number of times* (v_0, \dots, v_{d-1}) *circles the origin* to be the winding number of the piecewise linear path in $(\mathbb{R}^2)^* = \mathbb{R}^2 \setminus \{(0, 0)\}$ produced by concatenating the linear paths between v_i and v_{i+1} for $i = 0, \dots, d-1$.

Lemma 4.4.5. Let $a_0, \dots, a_{d-1} \in \mathbb{Z}$ such that $ST^{a_0} \dots ST^{a_{d-1}} =_{\text{SL}_2(\mathbb{Z})} I$ and let $v_0, v_1 \in \mathbb{Z}^2$ such that $\det(v_0, v_1) = 1$. Define v_2, \dots, v_{d-1} by

$$v_{i+2} = -v_i + a_i v_{i+1}$$

where $v_d = v_0$ and $v_{d+1} = v_1$. Then the winding number $W(ST^{a_0} \dots ST^{a_{d-1}}) \in \mathbb{Z}$ is the number of times that $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ circles the origin.

Proof. Let A_i be the matrix $[v_i, v_{i+1}]$ and recall that $A_{i+1} = A_i ST^{a_i}$. Thinking of the elements ST^{a_i} as elements of G , and thus as elements of the universal cover of $\text{SL}_2(\mathbb{R})$, this gives a path from A_i to A_{i+1} , and concatenating these paths gives a path from A_0 to itself in $\text{SL}_2(\mathbb{R})$. Projecting this path into the first column vector of the appropriate matrix gives a path in $(\mathbb{R}^2)^*$. We claim that this path is homotopic to the path formed by taking line segments between v_i and v_{i+1} . This is easily verified because both paths between v_i and v_{i+1} travel counterclockwise less than a full rotation.

We know that $W(ST^{a_0} \dots ST^{a_{d-1}})$ equals ψ of the path in $\text{SL}_2(\mathbb{R})$, and now we need to show that this equals the winding number in $(\mathbb{R}^2)^*$ of the first column vectors. To show this we note that the element of $\pi_1((\mathbb{R}^2)^*)$ (the group is abelian, so we may ignore basepoint) is given by the image of the element of $\pi_1(\text{SL}_2(\mathbb{R}))$ under the natural projection map. Thus, we merely need to

show that this projection acts correctly on a generator of $\pi_1(\mathrm{SL}_2(\mathbb{R}))$, but it is easy to see that the image of $\phi(S^4)$ yields a path of winding number 1.

□

Corollary 4.4.6. *There exists a bijection from the set of all sequences $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$, $d > 0$, which satisfy*

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^4$$

to the collection of all toric fans modulo the action of $\mathrm{SL}_2(\mathbb{Z})$. This bijection sends $(a_0, \dots, a_d) \in \mathbb{Z}^d$ to the equivalence class of fans

$$\{(v_0 = v_d, v_1 = v_{d+1}, v_2, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d \mid v_0, v_1 \in \mathbb{Z}^2, \det(v_0, v_1) = 1\}$$

in which

$$v_{i+2} = -v_i + a_i v_{i+1}$$

for $i = 0, \dots, d-1$.

Proof. Let $(v_0 = v_d, v_1, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ be a toric fan. That is, $\det(v_i, v_{i+1}) = 1$ for each $i = 0, \dots, d-1$ and the vectors are labeled in counter-clockwise order. It is shown in Section 4.3 Equation (4.3) that associated integers $(a_0, \dots, a_{d-1}) \in (\mathbb{Z})^d$ exist such that

$$ST^{a_0} \dots ST^{a_{d-1}} =_{\mathrm{SL}_2(\mathbb{Z})} I$$

which means

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^{4k}$$

for some $k \in \mathbb{Z}$ with $k \geq 0$. By Lemma 4.4.5 we know

$$W(ST^{a_0} \dots ST^{a_{d-1}}) = 1$$

so that the vectors will be labeled in the correct order for it to be a fan. Thus, $W(S^{4k}) = 1$ but $W(S^{4k}) = k$ so $k = 1$. Notice such a construction is well-defined on equivalence classes of toric fans

because the integers are prescribed via linear equations and fans in a common equivalence class are related by a linear map.

Now suppose that $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ satisfy $ST^{a_0} \dots ST^{a_{d-1}} =_G S^4$ and define $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ by

$$v_{i+2} = -v_i + a_i v_{i+1}$$

where $v_0, v_1 \in \mathbb{Z}^2$ are any two vectors for which $\det(v_0, v_1) = 1$. Then for each $i = 0, \dots, d-1$ we have

$$\det(v_{i+1}, v_{i+2}) = \det \begin{pmatrix} 0 & -1 \\ 1 & a_i \end{pmatrix} \det(v_i, v_{i+1}) = \det(v_i, v_{i+1}),$$

so by induction all of these determinants are 1. By Lemma 4.4.5, the path connecting adjacent vectors wraps around the origin only once, and since each v_{i+1} is located counterclockwise of v_i , we have that the v_i 's must be sorted in counterclockwise order.

It is straightforward to see that these constructions are inverses of one another.

□

Now that we have set up the algebraic framework the following results are straightforward to prove. First we prove that any fan with more than four vectors can be reduced to a fan with fewer vectors.

Lemma 4.4.7. *If $(v_0 = v_d, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ is a toric fan with $d > 4$ then there exists some $i \in \{0, \dots, d-1\}$ such that $v_i = v_{i-1} + v_{i+1}$.*

Proof. By Corollary 4.4.6 we know that to the fan $(v_0 = v_d, \dots, v_{d-1})$ there is an associated list of integers $a_0, \dots, a_{d-1} \in \mathbb{Z}$ such that $v_{i+2} = -v_i + a_i v_{i+1}$ and

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^4. \tag{4.8}$$

We must only show that for some $i \in \mathbb{Z}$ we have $a_i = 1$. Since $S^4 =_{\text{PSL}_2(\mathbb{Z})} I$ we can use Lemma 4.3.8 to conclude that there exist $i, j \in \mathbb{Z}$ satisfying $0 \leq i < j-1 \leq d-2$ such that $a_i, a_j \in \{-1, 0, 1\}$ and $(i, j) \neq (0, d-1)$. By way of contradiction assume that $a_i, a_j \in \{-1, 0\}$. Conjugate Equation (4.8) by ST^{a_n} for varying $n \in \mathbb{Z}$ to assure that $i \neq 0$ and $j \neq d-1$. Then at each of these values we may

use either $ST^0S =_G S^2$ or $ST^{-1}S =_G S^2TST$ to reduce the number of ST -pairs by one or two and produce a factor of S^2 . These reductions do not interfere with one another because the values in question are not adjacent. So we end up with

$$S^4 ST^{b_0} \dots ST^{b_{\ell-1}} =_G S^4$$

where $\ell \geq 1$ because we started with at least five ST -pairs and have reduced by at most four. This means

$$ST^{b_0} \dots ST^{b_{\ell-1}} =_G I$$

with $\ell \geq 1$. In fact, $\ell > 1$ because $ST^{b_0} =_G I$ is impossible for any choice of $b_0 \in \mathbb{Z}$.

This implies that $W(ST^{b_0} \dots ST^{b_{\ell-1}}) = 0$ and thus, by Lemma 4.4.5, that the corresponding collection of vectors winds no times about the origin. However, this is impossible since for such a sequence of vectors v_{i+1} is always counterclockwise of v_i and $\ell > 1$. \square

The case in which a vector in the fan is the sum of the adjacent vectors is important because this means the fan is the result of corner chopping a fan with fewer vectors in it. Now that we have the proper algebraic tools, we will be clear about the specifics of the corner chopping and reverse corner chopping operations.

Suppose $(v_0 = v_d, v_2, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ is a toric fan with associated integers $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$. Then

$$v_{i+2} = -v_i + a_i v_{i+1}$$

so if $a_i = 1$ then we have that $v_{i+1} = v_i + v_{i+2}$. Now we see that in this case

$$\det(v_i, v_{i+2}) = \det(v_i, -v_i) + \det(v_i, v_{i+1}) = 1$$

so

$$(w_0 = v_0, \dots, w_i = v_i, w_{i+1} = v_{i+2}, \dots, w_{i-2} = v_{d-1}) \in (\mathbb{Z}^2)^{d-1}$$

is also a fan. Next notice

$$-w_i + (a_{i+1} - 1)w_{i+1} = -(v_i + v_{i+2}) + a_{i+1}v_{i+2} = -v_{i+1} + a_{i+1}v_{i+2} = w_{i+2}$$

and

$$-w_{i-1} + (a_{i-1} - 1)w_i = (-v_{i-1} + a_{i-1}v_i) - v_i = v_{i+1} - v_i = w_{i+1}$$

so this new fan has associated to it the tuple of integers $(a_0, \dots, a_{i-1} - 1, a_{i+1} - 1, \dots, a_{d-1}) \in \mathbb{Z}^{d-1}$. An occurrence of 1 from the original tuple of integers has been removed and the adjacent integers have been reduced by 1. Algebraically, this move corresponds to the relation $STS =_G T^{-1}ST^{-1}$. Geometrically this move corresponds to the inverse of chopping a corner from the associated polygon (as is shown in Figure 4.1). The corner chopping of a toric polygon is done such that the new face of the polygon produced has inwards pointing normal vector given by the sum of the adjacent inwards pointing primitive integer normal vectors.

Now we can see that Lemma 4.4.7 tells us that fans with five or more vectors are the result of corner chopping a fan with fewer vectors. We will next classify all possible fans with fewer than five vectors.

Lemma 4.4.8. *Suppose that integers $a_0, \dots, a_{d-1} \in \mathbb{Z}$ satisfy*

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^4 \tag{4.9}$$

for some $d \in \mathbb{Z}$, $d \geq 0$.

1. If $d = 4$ then up to a cyclic reordering the set of integer quadruples which satisfy this equation is exactly $a_0 = 0, a_1 = k, a_2 = 0, a_3 = -k$ for each $k \in \mathbb{Z}$.
2. If $d = 3$ then $a_0 = a_1 = a_2 = -1$.
3. If $d < 3$ then there do not exist integers satisfying Equation (4.9).

Proof. Notice $ST^{a_0} \dots ST^{a_{d-1}} =_G S^4$ implies that

$$ST^{a_0} \dots ST^{a_{d-1}} =_{\text{PSL}_2(\mathbb{Z})} I.$$

From Lemma 4.3.8 we know that if $d = 1$ then this equality is impossible and if $d = 2$ the only possibility is $ST^0ST^0 =_{\text{PSL}_2(\mathbb{Z})} I$ but $ST^0ST^0 \neq_G S^4$. Now assume that $d = 3$. Again by Lemma 4.3.8 we know the only possibilities are $a_0 = a_1 = a_2 = \pm 1$. If $a_0 = a_1 = a_2 = 1$ then notice that

$$STSTST =_G S^2 \neq_G S^4.$$

Next notice that

$$ST^{-1}ST^{-1}ST^{-1} =_G S^4$$

so that is the only possibility for $d = 3$.

Now suppose that $d = 4$. Lemma 4.3.8 tells us that at least one of the a_i is in the set $\{-1, 0, 1\}$. By conjugation (which cyclically permutes the order of the integers) we may assume that $a_0 \in \{-1, 0, 1\}$. If $a_0 = 1$ then

$$STST^{a_1}ST^{a_2}ST^{a_3} =_G ST^{a_1-1}ST^{a_2}ST^{a_3-1}$$

so for this to equal S^4 in G we must have $a_1 - 1 = a_2 = a_3 - 1 = -1$ by the $d = 3$ argument above. It is straightforward to check that $STSST^{-1}S =_G S^4$ so we have found the required solution.

If $a_0 = -1$ then notice

$$ST^{-1}ST^{a_1}ST^{a_2}ST^{a_3} =_G S^4 \text{ implies } ST^{a_1+1}ST^{a_2}ST^{a_3+1} =_G S^2 =_{\text{PSL}_2(\mathbb{Z})} I$$

so by Lemma 4.3.8 we must have $a_1 + 1 = a_2 = a_3 + 1 = \pm 1$. This time, if $a_1 + 1 = a_2 = a_3 + 1 = -1$ then Equation (4.9) does not hold, since the left side will equal S^6 , but if $a_1 + 1 = a_2 = a_3 + 1 = 1$ then the equation holds. So we have found another solution, $ST^{-1}SSTS =_G S^4$, which has the form described in the statement of the Lemma .

Finally, suppose that $a_0 = 0$. Notice

$$ST^0ST_{a_1}ST^{a_2}ST^{a_3} =_G S^4 \text{ implies } ST^{a_2}ST^{a_1+a_3} =_G S^2 =_{\text{PSL}_2(\mathbb{Z})} I$$

so we can use Lemma 4.3.8 to conclude that we need $a_2 = a_1 + a_3 = 0$. Let $a_1 = k \in \mathbb{Z}$. Now we

have that

$$ST^0 ST^k ST^0 ST^{-k} =_G S^4$$

for any $k \in \mathbb{Z}$. Finally, observe that the other two possibilities we derived in the $d = 4$ case are just reorderings of this one with $k = 1$. \square

Definition 4.4.9. A *Delzant triangle* is the convex hull of the points $(0, 0)$, $(0, \lambda)$, $(\lambda, 0)$ in \mathbb{R}^2 for any $\lambda > 0$. A *Hirzebruch trapezoid* with parameter $k \in \mathbb{Z}_{\geq 0}$ is the convex hull of $(0, 0)$, $(0, a)$, (b, a) , and $(b + ak, 0)$ in \mathbb{R}^2 where $a, b > 0$. A Hirzebruch trapezoid with parameter zero is a *rectangle*.

These are shown in Figure 4.1. So we see that the fan corresponding to any Delzant triangle is

$$\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ -1 \end{pmatrix} \right)$$

with associated integers $(-1, -1, -1)$ and the fan corresponding to a Hirzebruch trapezoid with parameter k is

$$\left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ -k \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right)$$

with associated integers $(0, k, 0, -k)$. The following Theorem is immediate from Lemma 4.4.7 and Lemma 4.4.8.

Theorem 4.4.10 ([33]). *Every Delzant polygon can be obtained from a polygon $\mathrm{SL}_2(\mathbb{Z})$ -equivalent to a Delzant triangle, a rectangle, or a Hirzebruch trapezoid by a finite number of corner chops.*

Proof. Let Δ be any Delzant polygon with d edges, let $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ be the associated fan of inwards pointing primitive normal vectors, and let $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ be the integers associated to this fan. By Lemma 4.4.7 if $d > 4$ then $a_i = 1$ for some $i \in \{0, \dots, d-1\}$ so the fan is the result of a corner chop for some fan with $d-1$ vectors. That is, Δ is the result of a corner chop of some Delzant polygon with $d-1$ edges. If $d < 5$ then Lemma 4.4.8 lists each possibility. If $d = 4$ and $a_0 = a_1 = a_2 = a_3 = 0$ then Δ is $\mathrm{SL}_2(\mathbb{Z})$ -equivalent to a rectangle, if $a_0 = 0, a_1 = k, a_2 = 0$, and $a_3 = -k$ for $k \in \mathbb{Z} \setminus \{0\}$ then Δ is $\mathrm{SL}_2(\mathbb{Z})$ -equivalent to a Hirzebruch trapezoid, and if $d = 3$ with $a_0 = a_1 = a_2 = -1$ then Δ is $\mathrm{SL}_2(\mathbb{Z})$ -equivalent to a Delzant triangle. \square

4.5 Semitoric fans

Now we will apply the method from Section 4.4 to classify semitoric fans (Definition 4.2.5). The first step in the classification is given by a series of lemmas which we will use to manipulate the semitoric fans in a standard form.

Lemma 4.5.1. *If $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ is a semitoric fan and (v_i, v_{i+1}) is a hidden corner, then*

$$(w_0 = v_0, \dots, w_i = v_i, w_{i+1} = Tv_{i+1}, w_{i+2} = v_{i+1}, \dots, w_d = v_{d-1}) \in (\mathbb{Z}^2)^{d+1}$$

is a semitoric fan in which (w_i, w_{i+1}) is a Delzant corner and (w_{i+1}, w_{i+2}) is a fake corner.

Proof. We know that $\det(v_i, Tv_{i+1}) = 1$ because that pair of vectors forms a hidden corner. Notice that

$$\det(w_i, w_{i+1}) = \det(v_i, Tv_{i+1}) = 1$$

and

$$\det(w_{i+1}, Tw_{i+2}) = \det(Tv_{i+1}, Tv_{i+1}) = 0,$$

which concludes the proof. □

Lemma 4.5.2. *If $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ is a semitoric fan and (v_i, v_{i+1}) is a fake corner and (v_{i+1}, v_{i+2}) is a Delzant corner, then*

$$(w_0 = v_0, \dots, w_i = v_i, w_{i+1} = Tv_{i+2}, w_{i+2} = v_{i+2}, \dots, w_{d-1} = v_{d-1}) \in (\mathbb{Z}^2)^d$$

is a semitoric fan in which (w_i, w_{i+1}) is a Delzant corner and (w_{i+1}, w_{i+2}) is a fake corner.

Proof. We know $\det(v_i, Tv_{i+1}) = 0$ so $v_i = Tv_{i+1}$ since they are both on the top boundary and we also know $\det(v_{i+1}, v_{i+2}) = 1$. Now we can check that

$$\det(w_i, w_{i+1}) = \det(v_i, Tv_{i+2}) = \det(Tv_{i+1}, Tv_{i+2}) = \det(v_{i+1}, v_{i+2}) = 1$$

and

$$\det(w_{i+1}, Tw_{i+2}) = \det(Tv_{i+2}, Tv_{i+2}) = 0,$$

which concludes the proof. \square

In Lemma 4.5.1 we have described the process of removing a hidden corner and in Lemma 4.5.2 we have described the process of commuting a fake and Delzant corner. Both of these processes are defined in Definition 4.2.7.

Lemma 4.5.3. *Suppose that $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ is a semitoric fan. Then after a finite number of corner choppings the fan will be \mathcal{G}' -equivalent to one in which two adjacent vectors are $\begin{pmatrix} 0 \\ -1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$.*

Proof. Let $v_d = v_0$. If

$$v_i = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

for some $i \in \{0, \dots, d-1\}$ then notice

$$v_{i-1} = \begin{pmatrix} a \\ -1 \end{pmatrix}$$

for some $a \in \mathbb{Z}$. This is because (v_{i-1}, v_i) is not on the upper boundary so it must be a Delzant corner. Then by the action of $T^{-a} \in \mathcal{G}'$, which does not change v_i , we can attain the required pair of vectors.

Otherwise, renumber so that v_0 is in the lower half plane and v_1 is in the upper half plane. Then insert the vector $v_0 + v_1$ between them. This new vector will have a second component with a smaller magnitude than that of v_0 or of v_1 . Repeat this process until the new vector lies on the x -axis. Since it is a primitive vector it must be $\begin{pmatrix} \pm 1 \\ 0 \end{pmatrix}$. However since it is the sum of two vectors of opposite sides of the x -axis with the one above being counterclockwise about the origin of the one on bottom, it must be $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$. \square

Now we can put Lemmas 4.5.1, 4.5.2, and 4.5.3 together to produce a standard form for semitoric fans (see Figure 4.3). This standard form will be important to us in Section 5.2 because




Figure 4.3: Any semitoric fan with complexity $c \in \mathbb{Z}_{\geq 0}$ can be transformed into the standard fan of complexity c . This image has $c = 3$.

it can be obtained from any semitoric fan of complexity c by only using transformations which are continuous in the space of semitoric polygons.

Proposition 4.5.4. *Let $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ be a semitoric fan of complexity $c \in \mathbb{Z}$.*

1. *By only corner chopping, removing hidden corners, and commuting fake and Delzant corners we can obtain a new fan $(w_0 = w_{\ell+c}, \dots, w_{\ell+c-1}) \in (\mathbb{Z}^2)^{\ell+c}$ with $\ell + c \geq d$ such that*

- $w_0 = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$ and $w_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$;
- each corner (w_i, w_{i+1}) for $i = 0, \dots, \ell - 1$ is Delzant;
- each corner (w_i, w_{i+1}) for $i = \ell, \dots, \ell + c - 1$ is fake; and
- $w_\ell = T^c w_0$ so $\det(w_{\ell-1}, T^c w_0) = 1$.

2. *The fan $(w_0, \dots, w_{c+\ell-1})$ has associated integers $b_0, \dots, b_{\ell-1}$ such that*

$$w_2 = -w_\ell + b_0 w_1,$$

$$w_{i+2} = -w_i + b_i w_{i+1} \text{ for } i = 1, \dots, \ell - 2, \text{ and}$$

$$w_1 = -w_{\ell-1} + b_{\ell-1} w_\ell.$$

These integers satisfy

$$ST^{b_0} ST^{b_1} \dots ST^{b_{\ell-1}} =_G S^4.$$

3. The fan $(w_0, \dots, w_{c+\ell-1})$ can be obtained via a finite number of corner chops and reverse corner chops from a fan $(u_0, \dots, u_{c+3}) \in (\mathbb{Z}^2)^{c+4}$ where

$$u_0 = \begin{pmatrix} 0 \\ -1 \end{pmatrix}, u_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, u_2 = \begin{pmatrix} c \\ 1 \end{pmatrix}, u_3 = \begin{pmatrix} -1 \\ 0 \end{pmatrix},$$

and

$$u_{4+n} = \begin{pmatrix} -c+n \\ -1 \end{pmatrix}$$

for $n = 0, \dots, c-1$. In this fan the first four pairs are Delzant corners and the rest are fake corners.

Proof. The first part is immediate from Lemmas 4.5.1, 4.5.2, and 4.5.3. By Lemma 4.5.3 we know after a finite number of cuts and renumbering it can be arranged that the first two vectors in the fan are

$$\begin{pmatrix} 0 \\ -1 \end{pmatrix} \text{ and } \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Then we invoke Lemma 4.5.1 to remove all of the hidden corners, and finally use Lemma 4.5.2 commute all of the fake corners to be adjacent and arrive at the fan $(w_0, \dots, w_{\ell+c-1}) \in (\mathbb{Z}^2)^{\ell+c}$ in the statement of the proposition. Notice that (w_i, w_{i+1}) being fake for $i = \ell, \dots, \ell+c-1$ implies that

$$\det(w_\ell, Tw_{\ell+1}) = \dots = \det(w_{\ell+c-1}, Tw_0) = 0.$$

Now, we know both vectors in a fake corner must have negative second component by definition, so this implies that

$$w_\ell = Tw_{\ell+1} = T^2w_{\ell+2} = \dots = T^cw_0.$$

Now $(T^cw_0 = w_\ell, w_1, \dots, w_{\ell-1}) \in (\mathbb{Z}^2)^\ell$ is a toric fan (we know the vectors are in counter-clockwise order because we started with a semitoric fan) so there must exist $b_0, \dots, b_{\ell-1} \in \mathbb{Z}$ as in Part 2 of the statement of this theorem. Since

$$ST^{b_0}ST^{b_1}\dots ST^{b_{\ell-1}} =_G S^4$$

if $\ell > 4$ we can use Lemma 4.3.8 to conclude two nonconsecutive exponents are in $\{-1, 0, 1\}$ and one of these exponents is not b_0 or $b_{\ell-1}$. Let $b_i \in \{-1, 0, 1\}$ for some $0 < i < \ell - 1$. If $b_i = 1$ then we can remove the vector w_{i+1} from the semitoric fan to produce a new semitoric fan via the reverse corner chop operation. In this way we remove a vector from the fan and reduce the number of ST -pairs in the corresponding element of G via $STS =_G T^{-1}ST^{-1}$. If $b_i = -1$ then we use the relation $ST^{-1}S =_G SSTST$ and since

$$SSTST =_G S(STS)T =_G S(T^{-1}ST^{-1})T =_G ST^{-1}S$$

this relation can actually be realized by a corner chop, which we know corresponds to a legal move at the level of fans. If $b_i = 0$ then we actually have a factor of S^2 in the word. Notice that $TSTSTST$ reduces to either TSS or SST via a corner chop depending on where it is cut. In particular, if our word contains the subword $T^{a+1}SST^b$, we can perform a corner chop to obtain $T^{a+1}STSTST^{b+1}$, and then a reverse corner chop to reduce to T^aSST^{b+1} . Note that this can be done even if the T^a was in the ST^{b_0} term. By repeating this operation as necessary, we can move any factors of SS to the front of our word. So we see that in any case we can do an algebraic reduction which will reduce the number of ST -pairs by one and also corresponds to a fan transformation.

Repeat this process until there are only four ST -pairs. Then by Lemma 4.4.8 we know we must have reduced this equation to

$$ST^k S^2 T^{-k} S^2 =_G S^4$$

and thus we use the relation $TSS = SST$ to end up with S^4 . This means the corresponding integers are $b'_0 = 0, b'_1 = 0, b'_2 = 0$, and $b'_3 = 0$ which produces the desired fan (u_0, \dots, u_{c+3}) .

If $\ell < 4$ by Lemma 4.4.8 we must have $\ell = 3$ and $b_0 = -1, b_1 = -1, b_2 = -1$. In this case we can do a reverse corner chop to again achieve S^4 . □

Theorem 4.2.8 is immediate from Proposition 4.5.4.

Remark 4.5.5. Notice that Theorem 4.4.10 is different from Theorem 4.2.8 because in Theorem 4.4.10 the minimal models of the Delzant polygons may be achieved through only corner chops. In

Theorem 4.2.8 we use instead a variety of transformations (all of which are continuous, as we show in Section 5.2). \circlearrowright

Acknowledgements. Chapter 4, in part, is comprised of material submitted for publication by the author of this dissertation, Daniel M. Kane, and Álvaro Pelayo as *Classifying toric and semitoric fans by lifting equations from $SL_2(\mathbb{Z})$* , currently available as arXiv:1502.07698 [47].

Chapter 5

Connected components of space of semitoric systems

5.1 Introduction

Recall \mathcal{M}_{ST} , the moduli space of semitoric integrable systems which is defined in Section 2.2.7 and endowed with a metric, and thus topology, in Chapter 3. In this short chapter we use the results of Chapter 4 to prove the following:

Theorem 5.1.1. *If (M, ω, F) and (M', ω', F') are simple semitoric integrable systems such that:*

- (i) *they have the same number of focus-focus singularities;*
- (ii) *they have the same sequence of twisting indices,*

then there exists a continuous (with respect to the topology defined in [63]) path of semitoric systems with the same number of focus-focus points and same twisting indices between them. That is, the space of semitoric systems with fixed number of focus-focus points and twisting index invariant is path-connected.

This is analogous to Theorem 4.2.2, which states that the moduli space of symplectic toric manifolds is path-connected.

The Jaynes-Cummings system is an important example of a semitoric system with precisely one focus-focus point and is studied for example in [73] (systems with exactly one focus-focus point are referred to as systems of *Jaynes-Cummings type* for this reason). Since the twisting index is trivial when there is precisely one focus-focus point Theorem 5.1.1 implies the following.

Corollary 5.1.2. *Any semitoric system with precisely one focus-focus singular point may be continuously deformed into the Jaynes-Cummings system via a path of semitoric systems.*

5.2 Preparations

The results of Section 4.5 have an interpretation in symplectic geometry of toric manifolds and semitoric integrable systems.

To define see semitoric fans from this point of view we will need to examine single elements of the Delzant semitoric polygons defined in Definition 2.2.22. An element $(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f}) \in \mathcal{LWPolyg}_{m_f}(\mathbb{R}^2)$ is a *primitive semitoric polygon* if

1. Δ has everywhere finite height;
2. $\epsilon_j = +1$ for $j = 1, \dots, m_f$;
3. each ℓ_{λ_j} intersects the top boundary of Δ ;
4. any point in $\partial^{\text{top}}\Delta \cap \ell_{\lambda_j}$ for some $j \in \{1, \dots, m_f\}$ satisfies either the hidden or fake condition;
and
5. all other corners satisfy the Delzant condition.

Remark 5.2.1. Though a semitoric polygon is a family of polygons, it is determined by choosing a single primitive semitoric polygon. ◊

5.3 The connected components of \mathcal{M}_{ST}

Recall semitoric fans from Definition 4.2.5. Here we describe how they are related to Delzant semitoric polygons.

Definition 5.3.1. Let $\Delta \in \mathcal{LWPoly}_{m_f}(\mathbb{R}^2)$ be a primitive semitoric polygon. Then the *associated semitoric fan* is the semitoric fan \mathcal{F} formed by the inwards pointing primitive integer normal vectors to the edges of Δ in which the pair of vectors in \mathcal{F} are labeled as fake, hidden, or Delzant to correspond with the labeling of the corners of Δ .

Lemma 5.3.2. *Each relation in Theorem 4.2.8 corresponds to some continuous transformation of the polygons. More specifically, suppose that two fans $\mathcal{F}_0, \mathcal{F}_1 \in (\mathbb{Z}^2)^d$ are related by*

1. *performing corner chops;*
2. *performing reverse corner chops;*
3. *removing hidden corners; or*
4. *commuting fake and Delzant corners;*

(see Definition 4.2.7). Then there exists a continuous family of (compact) primitive semitoric polygons Δ_t , $t \in [0, 1]$, such that the fan associated to Δ_0 is \mathcal{F}_0 and the fan associated to Δ_1 is \mathcal{F}_1 .

Proof. Suppose that Δ is a primitive semitoric polygon with associated fan $\mathcal{F} = (v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ and fix some $i \in \{0, \dots, d-1\}$. Let $v_{-1} := v_{d-1}$ and $v_d := v_0$ so that the formulas used in this proof will be valid if $i = 0$ or $i = d-1$. Throughout, let $p \in \Delta$ be the vertex situated between the edges with inwards pointing normal vectors v_i and v_{i+1} . Let $u_i \in \mathbb{Z}^2$, $i = 1, 2$, denote the primitive vectors along which the edges adjacent to p are aligned, ordered so that $\det(u_1, u_2) > 0$.

For $w_1, w_2 \in \mathbb{Z}^2$ let $\mathcal{H}_p^\varepsilon(w_1, w_2)$ denote the half-space given by

$$\mathcal{H}_p^\varepsilon(w_1, w_2) = \{p + t_1 w_1 + t_2 w_2 : t_1 + t_2 \geq \varepsilon\}.$$

First we consider the corner chop operation. Suppose p is a Delzant corner. Fix some $\varepsilon_0 > 0$ smaller than the length of the edges incident at p .

For $t \in [0, 1]$ let

$$\Delta_t = \Delta \cap \mathcal{H}_p^{t\varepsilon_0}(u_1, u_2).$$

We see that Δ_t is a continuous family and since the edges of Δ_t are parallel to the edges of Δ except for the new edge with inwards pointing normal vector given by $v_i + v_{i+1}$ we see that the fan of Δ_t

is the corner chop of the fan for Δ for $t \in (0, 1]$. Since a reverse corner chop is the inverse of this operation, we can use the same path backwards.

Now suppose that p is a hidden corner. Let $\alpha \in \mathbb{Z}$ be the second component of v_{i+1} . For $\varepsilon_0 > 0$ smaller than the length of the adjacent edges and $t \in [0, 1]$ let

$$\Delta_t = \Delta \cap \mathcal{H}_p^{t\varepsilon_0}(u_1, \alpha^2 u_2).$$

It is straightforward to see that the normal vector to the new edge of the polygon is Tv_{i+1} . We see that Δ_t is a continuous family and by construction it has the desired fan. Thus Δ_t is the required family for the operation of removing hidden corners.

Finally, suppose that p is a fake corner and the next corner, which has adjacent edges which have inwards pointing normal vectors v_{i+1} and v_{i+2} , is a Delzant corner. Here we can see that this fan is the result of removing the hidden corner (v_i, v_{i+2}) from the fan

$$\mathcal{F}' = (v_0, \dots, v_i, v_{i+2}, \dots, v_{d-1}) \in (\mathbb{Z})^{d-1}.$$

We know (v_i, v_{i+2}) is a hidden corner because it is given that (v_i, v_{i+1}) is fake, which means $v_i = Tv_{i+1}$. Then we compute

$$\det(v_i, Tv_{i+2}) = \det(Tv_{i+1}, Tv_{i+2}) = 1$$

because (v_{i+1}, v_{i+2}) is Delzant. Thus there is a continuous path from any polygon with fan \mathcal{F} to any polygon with fan \mathcal{F}' . Let Δ' have fan \mathcal{F}' and let $p' \in \Delta'$ be the corner with adjacent edges which have v_i and v_{i+2} as inwards pointing normal vectors. Now, for some $\varepsilon_0 > 0$ small enough, we can consider

$$\Delta_t = \Delta' \cap \mathcal{H}_{p'}^{t\varepsilon_0}(\beta^2 u'_1, u'_2)$$

where u'_1 and u'_2 are the primitive integral vectors directing the edges adjacent to p' , ordered so that $\det(u'_1, u'_2) > 0$, and β^2 is the second component of v_i . This is a continuous path to a polygon with the required fan. This completes the proof.

□




Figure 5.1: For each fan transformation there is a continuous path of semitoric polygons which transitions between the fans. In (a) we show the corner chop and in (b) we show the removal of a hidden corner (which is replaced by a fake and a Delzant corner).

Recall that toric polygons are precisely the compact primitive semitoric polygons with complexity zero. Thus, by Proposition 4.5.4 and Lemma 5.3.2 we have recovered Theorem 4.2.2.

In light of Proposition 4.5.4 and Proposition 5.3.2 the only difficulty remaining is to prove the following lemma is incorporating the case of semitoric systems which have noncompact polygons as invariants. Recall that $\text{Polyg}_{\text{ST}}^{m_f, \vec{k}}(\mathbb{R}^2)$ denotes the set of labeled semitoric polygons which have m_f focus-focus points and twisting index \vec{k} .

Lemma 5.3.3. *Let $m_f \in \mathbb{Z}_{\geq 0}$ and $\vec{k} \in \mathbb{Z}^{m_f}$. Then $\text{Polyg}_{\text{ST}}^{m_f, \vec{k}}(\mathbb{R}^2)$ is path-connected.*

Proof. We must only consider the primitive semitoric polygons, because if the primitive polygons converge so do all of the polygons in the family. Any two compact primitive semitoric polygons in $\text{Polyg}_{\text{ST}}^{m_f, \vec{k}}(\mathbb{R}^2)$ with the same fan can be connected by a continuous path. This path is made by continuously changing the lengths of the edges because the angles of the two polygons must all be the same since they have the same fan. So, by Proposition 4.5.4 given two elements of $\text{Polyg}_{\text{ST}}^{m_f, \vec{k}}(\mathbb{R}^2)$ which are compact we know that the corresponding fans are related by the moves listed in that proposition and then by Lemma 5.3.2 we know these moves correspond to continuous paths of polygons. So we have established that any two compact elements of $\text{Polyg}_{\text{ST}}^{m_f, \vec{k}}(\mathbb{R}^2)$ are connected by a continuous path.




Figure 5.2: The continuous path from a compact primitive semitoric polygon to a noncompact primitive semitoric polygon with finitely many vertices.

Next assume that

$$[\Delta_w] = [(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f})] \in \text{Poly}_{\text{GST}}^{m_f, \vec{k}}(\mathbb{R}^2)$$

is such that Δ is noncompact but has only finitely many vertices. Choose $N \in \mathbb{R}$ such that all of the vertices of Δ are in the region $\{(x, y) \in \mathbb{R}^2 \mid -N < x < N\}$. The set $\Delta \cap [-N - 1, N + 1]$ is a polygonal subset of \mathbb{R}^2 but the corners which intersect $\ell_{N+1} \cup \ell_{-N-1}$ may not be Delzant. By [67, Remark 23] we may change the set on arbitrarily small neighborhoods of these corners to produce a new set, Δ' , which is equal to $\Delta \cap [-N - 1, N + 1]$ outside of those small neighborhoods and has only Delzant corners inside of those neighborhoods. Thus Δ' is a primitive semitoric polygon and by choosing the neighborhoods small enough we can assure that $\Delta \cap [-N, N] = \Delta' \cap [-N, N]$. For $t \in [0, 1)$ let $\Delta(t)$ be the polygon with the same fan as Δ' , the property that

$$\Delta(t) \cap [-N, N] = \Delta' \cap [-N, N],$$

and which has all of the same edge lengths as Δ' with the exception of the two or four edges which intersect $\ell_N \cup \ell_{-N}$. These edges are extended horizontally by a length of $\frac{1}{t-1}$. By this we mean that if an edge of Δ' which intersected $\{(x, y) \in \mathbb{R}^2 \mid x = N\}$ had as one of its endpoints (x_0, y_0) with $x_0 > N$ then the corresponding edge of $\Delta(t)$ would have as its endpoint $(x_0 + \frac{1}{t-1}, y_0 + m\frac{1}{t-1})$, where m is the slope of the edge in question. Then define $\Delta(1) = \Delta$ and we can see that $\Delta(t)$ for $t \in [0, 1]$ is a path from Δ' , which is compact, to Δ so $[(\Delta(t), (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})]$ is a continuous path which connects a compact semitoric polygon to $[\Delta_w]$. This process is shown in Figure 5.2.

Now we have connected all of the elements except for those with an infinite amount of

vertices. Suppose that

$$[\Delta_w] = [(\Delta, (\ell_{\lambda_j}, \epsilon_j, k_j)_{j=1}^{m_f})] \in \text{Polyg}_{\text{ST}}^{m_f, \vec{k}}(\mathbb{R}^2)$$

is such that Δ is noncompact and has infinitely many vertices. We will connect $[\Delta_w]$ to a polygon which has only finitely many vertices to finish the proof. Since Δ has everywhere finite height and is the intersection of infinitely many half-planes we can choose two of these planes which are not horizontal and are not parallel to one another. Denote the intersection of these two half-planes by A and notice $\Delta \subset A$. Since the boundaries of these two half-planes must intersect we can see that A can only be unbounded in either the positive or negative x -direction, but not both. Without loss of generality assume that A is unbounded in the positive x -direction.

Let ν be any admissible measure. For any $n \in \mathbb{Z}_{\geq 0}$ since $\nu(A) < \infty$ we know there exists some $x_n \in \mathbb{R}$ such that

$$\nu(A \cap [x_n, \infty)) < 1/n,$$

Δ does not have a vertex on the line ℓ_{x_n} , and $x_n > |\lambda_j|$ for all $j = 1, \dots, m_f$. Let Δ_n denote the polygon which satisfies

$$\Delta_n \cap [-\infty, x_n] = \Delta \cap [-\infty, x_n]$$

and has no vertices with x -coordinate greater than x_n .

For each $n \in \mathbb{Z}_{\geq 0}$ and $t \in (0, 1]$ define $\Delta_n(t)$ to have the same fan as Δ_{n+1} and to have all the same edge lengths as Δ_{n+1} except for the two edges which intersect ℓ_{x_n} . Extend those two edges horizontally by $1/t - 1$. Define $\Delta_n(0) = \Delta_n$. Now $\Delta_n(t)$ for $t \in [0, 1]$ is a C^0 path which takes Δ_n to Δ_{n+1} . Moreover,

$$\Delta \ominus \Delta_n(t) \subset A \cap [x_n, \infty) \text{ so } \nu(\Delta \ominus \Delta_n(t)) < 1/n$$

for each $t \in [0, 1]$. Each of these paths for $n \in \mathbb{Z}_{\geq 0}$ can be concatenated to form a continuous path Δ_t , $t \in [0, 1]$, from Δ_0 to Δ and we know that Δ_0 has only finitely many vertices. It is important not only that each Δ_n be getting closer to Δ but also that the path from Δ_n to Δ_{n+1} stays close to Δ . Then we define $[(\Delta(t), (\lambda_j, +1, k_j)_{j=1}^{m_f})]$ which is a continuous path from a semitoric polygon with finitely many vertices to $[\Delta_w]$. This is shown in Figure 5.3.

□




Figure 5.3: The continuous path which adds new vertices to a noncompact primitive semitoric polygon. This process is repeated to form a path to systems with infinitely many vertices.

Remark 5.3.4. According to [77, Theorem 3] all of the polygons in $\text{Poly}_{\text{ST}}^{m_f, \vec{k}}(\mathbb{R}^2)$ are compact if $m_f > 1$. ⊗

Now we can classify the connected components of \mathcal{M}_{ST} and \mathcal{M}_{T} . Recall that $\mathbb{I}_{m_f, \vec{k}} = \mathbb{I}_{m_f, \vec{k}'}$ if $k_j = k'_j + c$ for some $c \in \mathbb{Z}$ so when stating the following lemma we require that the first component of the twisting index be 0. This is done only to make sure that there are no repeats in the list of components. Recall \mathbb{I}_0 is the collection of semitoric lists of ingredients with $m_f = 0$ and $\mathcal{M}_{\text{ST},0} = \Phi^{-1}(\mathbb{I}_0)$ is the collection of semitoric systems with no focus-focus singularities.

Lemma 5.3.5. *The connected components of \mathbb{I} are*

$$\{\mathbb{I}_{m_f, \vec{k}} \mid m_f \in \mathbb{Z}_{>0}, \vec{k} \in \mathbb{Z}^{m_f} \text{ with } k_1 = 0\} \cup \mathbb{I}_0$$

and they are each path-connected.

Proof. It is sufficient to prove that $\mathbb{I}_{m_f, \vec{k}}$ is path-connected for each choice of $m_f \in \mathbb{Z}_{\geq 0}$ and $\vec{k} \in \mathbb{Z}^{m_f}$.

Let $m, m' \in \mathbb{I}_{m_f, \vec{k}}$ with

$$m = ([\Delta_w], (h_j)_{j=1}^{m_f}, ((S_j)^\infty)_{j=1}^{m_f}) \text{ and } m' = ([\Delta'_w], (h'_j)_{j=1}^{m_f}, ((S'_j)^\infty)_{j=1}^{m_f}).$$

By Lemma 5.3.3 we know there exists a continuous path

$$[\Delta_w(t)] = [(\Delta(t), (\ell_{\lambda_j(t)}, +1, k_j)_{j=1}^{m_f})],$$

$t \in [0, 1]$, from $[\Delta_w]$ to $[\Delta'_w]$ and by Proposition 3.2.2 we know $\mathbb{R}[[X, Y]]_0$ is path-connected so there

exists a continuous path $(S_j(t))^\infty$ from $(S_j)^\infty$ to $(S'_j)^\infty$ for each $j = 1, \dots, m_f$. For $j = 1, \dots, m_f$ let

$$\text{len}_j = \text{length}(\pi_2(\Delta \cap \ell_{\lambda_j}))$$

$$\text{len}'_j = \text{length}(\pi_2(\Delta' \cap \ell_{\lambda'_j}))$$

$$\text{len}_j(t) = \text{length}(\pi_2(\Delta(t) \cap \ell_{\lambda_j(t)}))$$

for $t \in [0, 1]$ and define

$$h_j(t) = \left(\frac{(1-t)h_j}{\text{len}_j} + \frac{th'_j}{\text{len}'_j} \right) \text{len}_j(t).$$

Now we have that $0 < h_j(t) < \text{len}_j(t)$ and $t \mapsto h_j(t)$ is a continuous function from $[0, 1]$ to \mathbb{R} because it is impossible for a semitoric polygon to have a vertical boundary at ℓ_{λ_j} for any $j \in \{1, \dots, m_f\}$.

Now define

$$m(t) = ([\Delta_w(t)], (h_j(t))_{j=1}^{m_f}, ((S_j(t))^\infty)_{j=1}^{m_f})$$

for $t \in [0, 1]$ which is a continuous path from m to m' . □

Thus we have established the following result.

Theorem 5.3.6. *The set of connected components of \mathcal{M}_{ST} is*

$$\{\mathcal{M}_{\text{ST}, m_f, \vec{k}} \mid m_f \in \mathbb{Z}_{>0}, \vec{k} \in \mathbb{Z}^{m_f} \text{ with } k_1 = 0\} \cup \mathcal{M}_{\text{ST}, 0}$$

and each $\mathcal{M}_{\text{ST}, m_f, \vec{k}}$ is path-connected.

Theorem 5.1.1 is equivalent to Theorem 5.3.6.

Acknowledgements. Chapter 5, in part, is comprised of material submitted for publication by the author of this dissertation, Daniel M. Kane, and Álvaro Pelayo as *Classifying toric and semitoric fans by lifting equations from $\text{SL}_2(\mathbb{Z})$* , currently available as arXiv:1502.07698 [47].

Chapter 6

Semitoric helixes and minimal models

6.1 Introduction

In this chapter I introduce a new combinatorial symplectic invariant of compact symplectic semitoric manifolds which I call a semitoric helix. This invariant encodes information about the singular affine structure around the preimage of the boundary of the momentum map by correcting for the effects of the Duistermaat monodromy from the focus-focus singular points which effects semitoric fans. Using this new invariant I give a complete classification of the minimal models of semitoric integrable systems; that is, I give an exact list of seven models depending on parameters which do not admit a symplectic semitoric blowdown, and any semitoric integrable system can be obtained from one of these by a sequence of semitoric blowups. This clarifies the relationship between blowdowns and the previously known invariants of compact semitoric integrable systems described in Section 2.2.7. I classify the minimal models of compact 4-dimensional symplectic semitoric manifolds by characterizing the minimal models that can be obtained from a semitoric helix by a finite sequence of blowdowns. A surprising fact is that the proof is purely algebraic.

6.1.1 Minimal models of symplectic toric manifolds

Recall that a toric fan $(v_0, \dots, v_{d-1}) \in (\mathbb{Z}^2)^d$ is *minimal* if

$$v_i \neq v_{i-1} + v_{i+1}$$

for $i = 0, \dots, d-1$. Since a blowup on a toric fan is inserting the vector $v_k + v_{k+1}$ into the fan and a blowdown is the opposite of this operation, minimal toric fans are exactly those on which a blowdown cannot be performed. Thus, any toric fan can be reduced to a minimal toric fan by performing blowdowns until no more are possible. On the other hand, this implies that any toric fan may be obtained from a minimal toric fan by a finite sequence of blowups. This occurs exactly when the associated toric integrable system does not admit a symplectic toric blowdown.

Let (M, ω, F) be a 4-dimensional toric integrable system and recall that M comes equipped with a natural \mathbb{T}^2 -action. A symplectic toric blowup is performed by removing an equivariantly embedded open ball $B^4(r)$ from M and collapsing the boundary via the Hopf map. A symplectic toric blowdown is the inverse of this operation.

We can, once again, restate Theorem 2.2.16 to suit the motivation of this chapter.

Theorem 6.1.1 (Fulton [33]). *A toric manifold is minimal if its fan is one of the following up to the action of $\mathrm{SL}_2(\mathbb{Z})$:*

1. $v_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, v_2 = \begin{pmatrix} -1 \\ -1 \end{pmatrix};$
2. $v_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, v_2 = \begin{pmatrix} -1 \\ 0 \end{pmatrix}, v_3 = \begin{pmatrix} 0 \\ -1 \end{pmatrix};$
3. $v_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, v_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, v_2 = \begin{pmatrix} -1 \\ k \end{pmatrix}, v_3 = \begin{pmatrix} 0 \\ -1 \end{pmatrix}$ for $k \in \mathbb{Z}, k \neq 0$.

Respectively, these are known as the *Delzant triangle*, the *square*, and the *Hirzebruch trapezoid*.

These fans are shown in Figure 6.1. Respectively, they correspond to $\mathbb{C}\mathbb{P}^2$, $\mathbb{C}\mathbb{P}^1 \times \mathbb{C}\mathbb{P}^1$, and a Hirzebruch surface.




Figure 6.1: The three possible minimal toric fans listed in Theorem 6.1.1, where $k \in \mathbb{Z}$ is the parameter for the Hirzebruch trapezoid.

6.1.2 Symplectic semitoric manifolds and helixes

Around any elliptic-elliptic point of a semitoric integrable system there exists a toric momentum map as is used when producing the associated semitoric polygon (see Section 2.2.7). A symplectic semitoric blowup or blowdown is performed by performing a symplectic toric blowup or blowdown, respectively, on this toric momentum map around the elliptic-elliptic point as described in Section 6.1.1.

Definition 6.1.2. A symplectic semitoric manifold (M, ω, F) is *minimal* if there does not exist any (M', ω', F') such that (M, ω, F) can be obtained from (M', ω', F') by a symplectic semitoric blowup.

In this chapter I construct a combinatorial invariant of semitoric manifolds, the semitoric helix, by correcting for the effects of the Duistermaat-monodromy of the focus-focus points at the cost of the periodicity.

Definition 6.1.3. A *semitoric helix* of length $d \in \mathbb{Z}_{>0}$ and complexity $c \in \mathbb{Z}_{\geq 0}$ is an equivalence class $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ where $\{v_i\}_{i \in \mathbb{Z}} \subset \mathbb{Z}^2$ is a collection of vectors such that:

1. $\det(v_i, v_{i+1}) = 1$ for all $i \in \mathbb{Z}$;
2. v_0, \dots, v_{d-1} are arranged in counter-clockwise order;
3. $T^c v_i = v_{i+d}$ for all $i \in \mathbb{Z}$.

The equivalence relation on such collections of vectors is given by $\{v_i\}_{i \in \mathbb{Z}} \sim \{w_i\}_{i \in \mathbb{Z}}$ if and only if




Figure 6.2: A minimal semitoric helix of length 6 and complexity 2. In the classification from Theorem 6.1.13 this is a type (7) minimal semitoric helix with $A_0 = ST^2ST^2$.

$v_i = T^k w_{i+\ell}$ for all $i \in \mathbb{Z}$ and for some fixed $k, \ell \in \mathbb{Z}$. A semitoric helix is *minimal* if

$$v_i \neq v_{i-1} + v_{i+1}$$

for all $i \in \mathbb{Z}$.

A minimal semitoric helix is shown in Figure 6.2. The semitoric helix is better suited to studying blowups and blowdowns than the semitoric fan is, and it is a more natural object because the vectors in a semitoric fan are effected by the artificial corners which are created by making cuts when producing the semitoric polygon.

Definition 6.1.4. Let $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ be a semitoric helix. If $v_j = v_{j-1} + v_{j+1}$ for some $j \in \mathbb{Z}$ then a new semitoric helix of the same complexity and one less length can be produced by removing $\{v_{j+nd}\}_{n \in \mathbb{Z}}$ from $\{v_i\}_{i \in \mathbb{Z}}$. The new helix is known as the *blowdown of \mathcal{H} at v_j* . The opposite of this operation, adding in the sum of two adjacent vectors, is known as a *blowup*.

Remark 6.1.5. Performing a blowup/blowdown on a semitoric helix \mathcal{H} corresponds to performing a blowup/blowdown on the associated semitoric manifold, so minimal semitoric integrable systems are precisely those associated to minimal semitoric helices. \circlearrowright

6.1.3 From semitoric systems to semitoric helixes

There is a unique semitoric helix related to each semitoric integrable system. We outline the construction here, and it is given in full detail in Section 6.3. Let (M, ω, F) be a semitoric integrable system.

1. **Construct polygon:** Associated to (M, ω, F) is semitoric polygon which consists of an infinite family of polygons [77] which I have described in Section 2.2.7. Fix a polygon Δ from this family;
2. **Construct semitoric fan:** Δ is rational, so take the collection of inwards pointing integer normal vectors w_0, \dots, w_{m-1} of minimal length to its edges. If all of the cuts in Δ are up, this corresponds to considering the associated semitoric fan (see Chapter 4);
3. **Correct for monodromy effect:** For each j such that (w_j, w_{j+1}) is either hidden or fake replace w_{j+1}, \dots, w_{m-1} by

$$Tw_{j+1}, \dots, Tw_{m-1}.$$

Label the new list of vectors w'_0, \dots, w'_{m-1} ;

4. **Remove repeated vectors:** Now each pair (w'_i, w'_{i+1}) either satisfies $\det(w'_i, w'_{i+1}) = 1$ or $w'_i = Tw'_{i+1}$. For each j such that $w'_j = w'_{j+1}$ remove w'_{j+1} from the list (these come from the fake corners) and denote the remaining vectors by v_0, \dots, v_{d-1} . Notice $\det(v_i, v_{i+1}) = 1$ for all $i = 0, \dots, d-2$;
5. **Extend to helix:** Extend this finite list of vectors to a semitoric helix \mathcal{H} of length d and complexity c (the number of focus-focus points of the original semitoric manifold) by forcing condition (3) from Definition 6.1.3.

Given a symplectic system (M, ω, F) the semitoric helix \mathcal{H} produced by the above construction is known as the semitoric helix associated to (M, ω, F) . The construction procedure includes several choices by the resulting semitoric helix is unique.

Lemma 6.1.6. *Given a semitoric integrable system (M, ω, F) there exists exactly one associated semitoric helix.*

Lemma 6.1.6 follows from Lemma 6.3.1.

Definition 6.1.7. The map ϕ assigns to each semitoric integrable system (M, ω, F) a semitoric helix $\phi(M, \omega, F) = \mathcal{H}$, where \mathcal{H} is the semitoric helix associated to (M, ω, F) .

Lemma 6.1.6 shows that ϕ is well-defined.

The main result of this chapter is Theorem 6.1.13, which is the classification of all minimal semitoric systems in terms of the associated semitoric helix. To prove Theorem 6.1.13 I use a refined version of the algebraic tools used in Chapter 4.

6.1.4 The algebraic technique

Recall the groups

$$\mathrm{SL}_2(\mathbb{Z}) = \langle S, T \mid STS = T^{-1}ST^{-1}, S^4 = I \rangle,$$

$$\mathrm{PSL}_2(\mathbb{Z}) = \langle S, T \mid STS = T^{-1}ST^{-1}, S^2 = I \rangle,$$

and

$$G = \langle S, T \mid STS = T^{-1}ST^{-1} \rangle$$

Let $(\mathbb{R}^2)^* := \mathbb{R}^2 \setminus \{(0, 0)\}$. Here we set some notation.

Definition 6.1.8. Given any closed loop $\tilde{\gamma}: [0, 1] \rightarrow (\mathbb{R}^2)^*$, $\tilde{\gamma}(0) = \tilde{\gamma}(1)$, we denote by $\mathrm{wind}(\tilde{\gamma}) \in \mathbb{Z}$ the usual winding number of $\tilde{\gamma}$.

Define $\mathrm{pr}: \mathrm{SL}_2(\mathbb{Z}) \rightarrow (\mathbb{R}^2)^*$ by

$$\mathrm{pr}: \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mapsto \begin{pmatrix} a \\ c \end{pmatrix}.$$

Definition 6.1.9. Given any loop $\gamma: [0, 1] \rightarrow \mathrm{SL}_2(\mathbb{Z})$, $\gamma(0) = \gamma(1)$, we define the *winding number* of γ , denoted $\mathrm{wind}(\gamma)$, by

$$\mathrm{wind}(\gamma) := \mathrm{wind}(\mathrm{pr}(\gamma)).$$

Thus, there exists a map

$$\ker(G \rightarrow \mathrm{SL}_2(\mathbb{Z})) \rightarrow \mathbb{Z}$$

which takes each element $\sigma \in \ker(G \rightarrow \mathrm{SL}_2(\mathbb{Z}))$ to $\mathrm{wind}(\rho(\sigma))$. The winding number

$$W : \langle S, T \rangle \rightarrow \frac{1}{12}\mathbb{Z},$$

defined in Chapter 4, extends this map to all of $\langle S, T \rangle$ and descends to a map on G which we also denote W . Recall that given $\sigma \in G$,

$$W(\sigma) = \mathrm{wind}(\mathrm{pr} \circ \rho(\sigma)).$$

For a semitoric helix $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ and $B \in \mathrm{SL}_2(\mathbb{Z})$ we use the notation

$$B\mathcal{H} := [\{Bv_i\}_{i \in \mathbb{Z}}].$$

A cyclic permutation of a list $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ of integers is given by

$$(a_{k \bmod d}, a_{k+1 \bmod d}, \dots, a_{k+d-1 \bmod d}).$$

The following is the helix version of Corollary 4.4.6.

Proposition 6.1.10. *Associated to any semitoric helix of length d and complexity c there is a list of integers $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ which satisfy*

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^4 X^{-1} T^c X$$

for some $X \in G$. This list of integers is unique up to cyclic permutation, and any such list of integers is associated to some semitoric helix. Semitoric helices \mathcal{H} and \mathcal{H}' have the same length, complexity, and associated integers if and only if there exists $B \in \mathrm{SL}_2(\mathbb{Z})$ such that

$$\mathcal{H} = B\mathcal{H}'.$$

Proposition 6.1.10 is proven in Section 6.4.

An obstruction to producing semitoric helixes is producing words in S and T which not only evaluate to the correct matrix in $\mathrm{SL}_2(\mathbb{Z})$ but also which have low winding number. In order to study these we develop a theory of reduced forms of words; those words with only positive powers of S which evaluate to a given matrix, up to sign, in $\mathrm{SL}_2(\mathbb{Z})$ and which have smallest possible winding number. It turns out there is a relatively simple characterization of these reduced forms and we are able to take advantage of the fact that the words corresponding to minimal semitoric helixes are nearly reduced. By understanding the difference between these and their reduced forms we classify the possibilities.

Thus, the first step towards classifying minimal semitoric helixes is to produce a standard form of minimal winding number for elements of $\mathrm{PSL}_2(\mathbb{Z})$.

Definition 6.1.11. A word in S and T is *S-positive* if it can be written using only non-negative powers of S , T , and T^{-1} .

Theorem 6.1.12 (Standard form in $\mathrm{PSL}_2(\mathbb{Z})$). *If $X \in \mathrm{SL}_2(\mathbb{Z})$ there exists a unique string $\overline{X} \in \langle S, T \rangle$ such that $X =_{\mathrm{PSL}_2(\mathbb{Z})} \overline{X}$ and*

$$\overline{X} =_{\langle S, T \rangle} T^b S T^{a_0} \dots S T^{a_{d-1}}$$

where $a_i > 1$ for $i = 0, \dots, d-2$. Moreover,

$$W(\overline{X}) \leq W(\eta)$$

for all *S-positive* $\eta \in \langle S, T \rangle$ satisfying $\eta =_{\mathrm{PSL}_2(\mathbb{Z})} X$.

We call \overline{X} the *standard form* of X . Theorem 6.1.12 is proven in Section 6.5.

6.1.5 Main result: minimal models of semitoric integrable systems

Lemma 6.6.3 has the following consequence, which is proven in Section 6.6 and is the main result of this chapter. Let

$$\mathcal{S} = \{A \in \mathrm{SL}_2(\mathbb{Z}) \mid \overline{A} = S T^{a_0} \dots S T^{a_{d-1}}, \text{ such that } d > 1, a_{d-1} \notin \{0, 1\}\}. \quad (6.1)$$

A semitoric helix of length d is determined by specifying the complexity and any d consecutive vectors in any representative of the helix.

Theorem 6.1.13. *Suppose that (M, ω, F) is a compact minimal semitoric integrable system with associated semitoric helix $\mathcal{H} = \phi(M, \omega, F)$ of length d and complexity $c > 0$. Write $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$. If $d < 5$ then the representative $\{v_i\}_{i \in \mathbb{Z}}$ can be chosen to be exactly one of the following:*

| type | length | v_0, \dots, v_{d-1} | complexity |
|------|---------|--|------------|
| (1) | $d = 2$ | $\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ -2 \end{pmatrix}$ | $c = 1$ |
| (2) | $d = 2$ | $\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ -1 \end{pmatrix}$ | $c = 2$ |
| (3) | $d = 3$ | $\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ k \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \end{pmatrix}$ $k \neq \pm 2$ | $c = 1$ |
| (4) | $d = 3$ | $\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ -1 \end{pmatrix}$ | $c \neq 2$ |
| (5) | $d = 4$ | $\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ k \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \end{pmatrix}$ $k \neq \pm 1$ | $c \neq 1$ |
| (6) | $d = 4$ | $\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \end{pmatrix}, \begin{pmatrix} k \\ -1 \end{pmatrix}$ $k \neq -1$ $k \neq 1 - c$ | $c > 0$ |

Otherwise, $d > 5$ and we say that \mathcal{H} is of type (7). In this case there exists some unique $A_0 \in \mathcal{S}$ such that $[v_0, v_1] = A_0$ and

$$v_j = a_{j-2}v_{j-1} - v_{j-2}$$

determines $\{v_i\}_{i \in \mathbb{Z}} \in \mathcal{H}$ where a_0, \dots, a_{d-1} are defined by

$$S^2 \overline{A_0^{-1}} T^c \overline{A_0} = ST^{a_0} \dots ST^{a_{d-1}}.$$

Additionally, for each $A_0 \in \mathcal{S}$ there exists such a semitoric helix.

Types (1)-(6) are shown in Figure 6.3 and a representative example of type (7) is shown in Figure 6.2.

Idea of proof of Theorem 6.1.13

In the proof of Theorem 6.1.12 we use a reduction algorithm with four steps. Three of these steps reduce the winding number by $1/2$ and the remaining step, which corresponds to a blowdown,




Figure 6.3: Minimal semitoric helices of type (1)-(6) of Theorem 6.1.13.

does not change the winding number. We will see, by Lemmas 6.5.4 and 6.5.6, that if a_0, \dots, a_{d-1} is associated to a semitoric helix then

$$W(ST^{a_0} \dots ST^{a_{d-1}}) - W(\overline{ST^{a_0} \dots ST^{a_{d-1}}}) = \begin{cases} 1, & X = \mathrm{PSL}_2(\mathbb{Z}) T^k \\ \frac{1}{2}, & \text{otherwise,} \end{cases}$$

and thus we know that $ST^{a_0} \dots ST^{a_{d-1}}$ can be reduced to the standard form from Theorem 6.1.12 by using only one or two of the moves which reduce W along with any number of blowdowns.

This observation allows us to prove Lemma 6.6.3, which classifies all minimal words satisfying Equation (6.3). This implies Theorem 6.1.13, which is proven in Section 6.6.

6.2 An application to semitoric systems

Theorem 6.1.13 has the following surprising consequence in the study of symplectic semitoric manifolds, which follows from Lemma 6.6.4.

Theorem 6.2.1. *If a compact semitoric integrable system $(M, \omega, (J, H))$ has more than 2 focus-focus singular points, then J has either a non-unique maximum or a non-unique minimum.*

In [77, Theorem 3] Vũ Ngọc uses an analytic argument related to the Duistermaat-Heckman measure on semitoric manifolds to prove a result analogous to Theorem 6.2.1 in the case that the

manifold is non-compact with 2 or more focus-focus singular points. Combining these two results we conclude that if (M, ω, F) is a semitoric manifold, compact or not, with greater than 2 focus-focus points then J either has a non-unique maximum or a non-unique minimum.

Remark 6.2.2. Type (7) semitoric helices always have a specific form. If $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ is a semitoric helix then

1. $d > 5$;
2. $v_2 = -v_0$;
3. $v_k = \begin{pmatrix} \pm 1 \\ 0 \end{pmatrix}$ for some $k \in \{3, \dots, d-1\}$ (Lemma 6.6.4).

◊

6.3 From semitoric systems to semitoric helices

In this section we explain how to produce a semitoric helix from a semitoric integrable system using the associated semitoric polygon.

6.3.1 Producing a semitoric helix

Let $[\Delta_w]$ be a compact Delzant semitoric polygon of complexity $c \in \mathbb{Z}_{\geq 0}$ (i.e. it has c focus-focus points) and write

$$\Delta_w = (\Delta, (\ell_{\lambda_j})_{j=1}^c, (\epsilon_j)_{j=1}^c).$$

The hidden and fake corners come from the existence of focus-focus singular points in the system. To create a semitoric helix we remove the effect of these points.

Suppose that Δ has $\ell > 0$ edges and let $w_0, \dots, w_{\ell-1}, w_\ell = w_0 \in \mathbb{Z}^2$ denote the inwards pointing normal vectors of minimal length to those edges ordered cyclically. That is, choose w_0 to be an inwards pointing normal vector to any edge of Δ and order $w_1, \dots, w_{\ell-1}, w_\ell$ so that w_i and w_{i+1} are normal vectors to edges adjacent to the same vertex and $\det(w_i, w_{i+1}) > 0$. Then the vectors are ordered in counter-clockwise order and $w_0 = w_\ell$ because they are normal to the same edge. We




Figure 6.4: The semitoric helix is produced by unwinding the polygon on the fake corner p .

call a pair (w_i, w_{i+1}) *Delzant*, *fake*, or *hidden* depending on the corner type of the vertex in between them. Notice there are exactly c vertices of Δ which are not Delzant.

Define

$$f(m) = |\{i < m \mid (w_i, w_{i+1}) \text{ is not Delzant}\}|$$

and consider the sequence

$$(T^{f(0)}w_0, T^{f(1)}w_1, \dots, T^{f(\ell-1)}w_{\ell-1}, T^{f(\ell)}w_\ell) \in (\mathbb{Z}^2)^{\ell+1}.$$

Notice that any two adjacent vectors in this sequence are either a Delzant pair or equal to one another. Remove any vector that is equal to an adjacent one and renumber in the same order to produce

$$(v_0, \dots, v_{d-1}, v_d) \in (\mathbb{Z}^2)^{d+1}$$

which are again in counter-clockwise order. Now $v_d = T^c v_0$ where c is the total number of corners that are not Delzant. Extend $(v_i)_{i=0}^{d-1}$ into a semitoric helix \mathcal{H} of length d and complexity c by using the relation $T^c v_i = v_{i+d}$. The semitoric helix \mathcal{H} is said to be associated to Δ_w .

The idea of this construction is shown in Figure 6.4.

Lemma 6.3.1. *Suppose $[\Delta_w]$ and $[\Delta'_w]$ are compact Delzant semitoric polygons such that $[\Delta_w] = [\Delta'_w]$.*

If \mathcal{H} and \mathcal{H}' are associated to Δ_w and Δ'_w , respectively, then $\mathcal{H} = \mathcal{H}'$.

Proof. Check that the semitoric helix produced does not depend on which element of $[\Delta_w]$ is chosen or on any choices made during the construction. The element of $[\Delta_w]$ is defined up to the action of $G_c \times \mathcal{G}$ and the only choice made during the construction is which edge to place w_0 on. \square

Lemma 6.1.6 follows from Lemma 6.3.1.

Lemma 6.3.2. *Given any semitoric helix \mathcal{H} there exists a Delzant semitoric polygon $[\Delta_w]$ such that*

$$\phi([\Delta_w]) = \mathcal{H}.$$

Proof. Let $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ be a semitoric helix of length d and complexity c . Define a collection of vectors w_0, \dots, w_{d+c} by

$$w_i = v_i \text{ for } i = 0, \dots, d-1$$

and

$$w_i = T^{i-d+1}v_{d-1} \text{ for } i = d, \dots, d+c.$$

Then $\det(w_i, w_{i+1}) = 1$ for $i = 0, \dots, d-1$,

$$\det(Tw_i, w_{i+1}) = \det(Tw_i, Tw_i) = 0$$

for $i = d, \dots, d+c-1$, and $w_0 = w_{d+c}$ by the periodicity requirement on the helix \mathcal{H} . The vectors w_0, \dots, w_{d+c-1} are arranged counter-clockwise so there exists a polygon $\Delta \subset \mathbb{R}^2$ with $d+c$ edges which has these as inwards pointing normal vectors. The polygon Δ has d Delzant corners c fake corners, and since T does not change the y -value of a vector we see that either all of the fake corners are on the top boundary of Δ or all of the fake corners are on the bottom boundary of Δ . Let λ_i be the horizontal position of the i^{th} fake corner and we may number these so that

$$\lambda_1 < \lambda_2 < \dots < \lambda_c$$

Since each vertical line intersects the top and bottom boundaries at most once each. If the fake corners are on the top boundary let $\epsilon_j = +1$ for $j = 1, \dots, c$ and otherwise let $\epsilon_j = -1$ for $j = 1, \dots, c$.

Then,

$$[(\Delta, (\ell_{\lambda_j})_{j=1}^c, (\epsilon_j)_{j=1}^c)]$$

is a Delzant semitoric polygon with associated semitoric helix \mathcal{H} . □

Lemma 6.3.2 shows that ϕ is surjective by producing a right inverse, but ϕ is not injective. This is because the semitoric helix does not encode information about the length of the edges of the associated polygon or about the position of the focus-focus points, the values $\lambda_1, \dots, \lambda_c$.

If Δ_w is such that $\epsilon_j = 1$ for $j = 1, \dots, c$ then the inwards pointing normal vectors to Δ_w form a semitoric fan as in Chapter 4.

6.4 Semitoric helices and $\mathrm{SL}_2(\mathbb{Z})$

In this Section we prove Proposition 6.1.10, which is the tool we use to translate questions about semitoric helices into questions about words on letters S and T .

Lemma 6.4.1. *Given any semitoric helix $\mathcal{H} = \{v_i\}_{i \in \mathbb{Z}}$ of length d there is a list of integers $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ such that*

$$a_i v_{i+1} = v_i + v_{i+2} \tag{6.2}$$

for $i = 0, \dots, d-1$. Furthermore, given v_0, v_1 , and (a_0, \dots, a_{d-1}) the helix can be recovered.

Proof. Let $\mathcal{H} = \{v_i\}_{i \in \mathbb{Z}}$ be a semitoric helix of length d and complexity c . Let $A_i = [v_i, v_{i+1}]$ and write v_{i+2} in the $\{v_i, v_{i+1}\}$ basis as

$$v_{i+2} = b_i v_i + a_i v_{i+1},$$

for $a_i, b_i \in \mathbb{Z}$. Thus,

$$A_i \begin{pmatrix} 0 & b_i \\ 1 & a_i \end{pmatrix} = A_{i+1}$$

and since $A_i, A_{i+1} \in \mathrm{SL}_2(\mathbb{Z})$ we see the determinant of each side is 1 so $b_i = -1$. The result follows. □

Definition 6.4.2. The a_0, \dots, a_{d-1} in Lemma 6.4.1 are the *associated integers* to the helix \mathcal{H} .

Recall $\widetilde{\mathrm{SL}}_2(\mathbb{R})$ denotes the universal cover of $\mathrm{SL}_2(\mathbb{R})$, so $\alpha \in \widetilde{\mathrm{SL}}_2(\mathbb{R})$ is a continuous map $\alpha: [0, 1] \rightarrow \mathrm{SL}_2(\mathbb{R})$ satisfying $\alpha(0) = I$. Recall G is isomorphic to the preimage of $\mathrm{SL}_2(\mathbb{Z})$ in the universal cover of $\mathrm{SL}_2(\mathbb{R})$, denoted $\widetilde{\mathrm{SL}}_2(\mathbb{R})$, by Proposition 4.3.7 via the homomorphism $\rho: G \rightarrow \widetilde{\mathrm{SL}}_2(\mathbb{Z})$ generated by

$$(\rho(T))(t) = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad (\rho(S))(t) = \begin{pmatrix} \cos(\frac{\pi t}{2}) & -\sin(\frac{\pi t}{2}) \\ \sin(\frac{\pi t}{2}) & \cos(\frac{\pi t}{2}) \end{pmatrix}.$$

The operation in G is concatenation and the operation in $\widetilde{\mathrm{SL}}_2(\mathbb{R})$ is given by the following. if $\alpha, \beta \in \mathrm{SL}_2(\mathbb{R})$ then

$$\alpha, \beta: [0, 1] \rightarrow \mathrm{SL}_2(\mathbb{R})$$

and we define $\alpha\beta \in \widetilde{\mathrm{SL}}_2(\mathbb{R})$ by

$$\alpha\beta(t) = \begin{cases} \alpha(2t), & 0 \leq t \leq 1/2 \\ \alpha(1)\beta(2t-1), & 1/2 \leq t \leq 1. \end{cases}$$

That is, the path $\alpha\beta$ is obtained by traveling first along the path α and then along the path produced by multiplying each element of the path β on the left by $\alpha(1)$. It turns out that the path produced by traveling first along β and then along α multiplied on the right by $\beta(1)$ is homotopic.

Lemma 6.4.3. *If $\alpha, \beta \in \widetilde{\mathrm{SL}}_2(\mathbb{R})$ then the paths in $\mathrm{SL}_2(\mathbb{R})$ from I to $\alpha(1)\beta(1)$ given by*

$$\gamma_0(t) = \begin{cases} \beta(2t), & 0 \leq t \leq 1/2 \\ \alpha(2t-1)\beta(1), & 1/2 < t \leq 1 \end{cases}$$

and

$$\gamma_1(t) = \begin{cases} \alpha(2t), & 0 \leq t \leq 1/2 \\ \alpha(1)\beta(2t-1), & 1/2 < t \leq 1 \end{cases}$$

are homotopic.




Figure 6.5: The homotopy from the proof of Lemma 6.4.3. A point (x, y) in the above plane represents $\alpha(x)\beta(y) \in \mathrm{SL}_2(\mathbb{R})$.

Proof. A continuous homotopy between them is given by

$$\gamma_s(t) = \begin{cases} \alpha(2t), & 0 \leq t \leq s/2 \\ \alpha(s)\beta(2t-s), & s/2 \leq t \leq \frac{1+s}{2} \\ \alpha(2t-1)\beta(1), & \frac{1+s}{2} \leq t \leq 1 \end{cases}$$

for $0 \leq s \leq 1$, which is shown in Figure 6.5.

□

Recall the map $\mathrm{pr}: \mathrm{SL}_2(\mathbb{R}) \rightarrow (\mathbb{R}^2)^*$, where $(\mathbb{R}^2)^* = \mathbb{R}^2 \setminus \{(0, 0)\}$, given by $\mathrm{pr}([v_1, v_2]) = v_2$.

Since

$$\pi_1(\mathrm{SL}_2(\mathbb{R})) \cong \pi_1((\mathbb{R}^2)^*) \cong \mathbb{Z}$$

and

$$\mathrm{pr} \begin{pmatrix} \cos(2\pi t) & -\sin(2\pi t) \\ \sin(2\pi t) & \cos(2\pi t) \end{pmatrix} = \begin{pmatrix} \cos(2\pi t) \\ \sin(2\pi t) \end{pmatrix}$$

for $t \in \mathbb{R}$ we see pr sends a generator of $\pi_1(\mathrm{SL}_2(\mathbb{R}))$ to a generator of $\pi_1((\mathbb{R}^2)^*)$, so $\mathrm{pr}^*: \pi_1(\mathrm{SL}_2(\mathbb{R})) \rightarrow \pi_1((\mathbb{R}^2)^*)$ is an isomorphism.

We say a path $\gamma: [0, 1] \rightarrow (\mathbb{R}^2)^*$ travels only counter-clockwise at most one full rotation if $t \mapsto \theta(\gamma(t))$ is an increasing function where $\theta: (\mathbb{R}^2)^* \rightarrow [0, 2\pi)$ is the angle coordinate with

$$\theta(\gamma(0)) = 0.$$

Lemma 6.4.4. *Let $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ be a semitoric helix of length d and complexity c and let $A_0 = [v_0, v_1]$. If $\sigma \in G$ is given by*

$$\sigma = ST^{a_0} \dots ST^{a_{d-1}}$$

then $\text{pr}(A_0 \rho(\sigma))$ is homotopic to a path from v_0 to v_{d-1} which travels counter-clockwise at most one full rotation.

Proof. Let $A_i = [v_i, v_{i+1}]$ for $1 \leq i \leq d-1$ and recall $A_i = A_{i-1}ST^{a_{i-1}}$. Thus,

$$\text{pr}(A_{i-1} \rho(ST^{a_{i-1}}))$$

is a path from v_i to v_{i+1} which is homotopic to

$$\begin{aligned} \gamma_i(t) &= \text{pr} \left(A_{i-1} \begin{pmatrix} \cos\left(\frac{\pi t}{2}\right) & -\sin\left(\frac{\pi t}{2}\right) + ta_{i-1}\cos\left(\frac{\pi t}{2}\right) \\ \sin\left(\frac{\pi t}{2}\right) & \cos\left(\frac{\pi t}{2}\right) + ta_{i-1}\sin\left(\frac{\pi t}{2}\right) \end{pmatrix} \right) \\ &= \cos\left(\frac{\pi t}{2}\right) v_{i-1} + \sin\left(\frac{\pi t}{2}\right) v_i. \end{aligned}$$

The path γ_i travels only counter-clockwise at most one full rotation from v_{i-1} to v_i so the composition of paths $\gamma_1, \dots, \gamma_{d-1}$ travels counter-clockwise from v_0 to v_{d-1} . The result follows because v_0, \dots, v_{d-1} are arranged in counter-clockwise order. \square

Lemma 6.4.5. *The integers $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ are associated to a semitoric helix of complexity $c \geq 0$ if and only if*

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^4 X^{-1} T^c X \tag{6.3}$$

for some $X \in G$. If $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ is a semitoric helix with associated integers (a_0, \dots, a_{d-1}) then $A_0 = [v_0, v_1]$ satisfies $X =_G A_0$.

Proof. Let $A_i = [v_i, v_{i+1}]$. By Lemma 6.4.1 and the fact that

$$\begin{pmatrix} 0 & -1 \\ 1 & a_i \end{pmatrix} = ST^{a_i}$$

we find that $A_{i+1} =_{\text{SL}_2(\mathbb{Z})} A_i ST^{a_i}$ for all $i \in \mathbb{Z}$. We conclude that

$$A_d =_{\text{SL}_2(\mathbb{Z})} A_0 ST^{a_0} \dots ST^{a_{d-1}}$$

and since $T^c A_0 =_{\text{SL}_2(\mathbb{Z})} A_d$ this implies that

$$ST^{a_0} \dots ST^{a_{d-1}} =_{\text{SL}_2(\mathbb{Z})} A_0^{-1} T^c A_0.$$

Since S^4 generates the kernel of the projection $G \rightarrow \text{SL}_2(\mathbb{Z})$ this means that

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^{4k} A_0^{-1} T^c A_0$$

for some $k \in \mathbb{Z}$. This is because when reducing an element of $\text{SL}_2(\mathbb{Z})$ we can assume that the relation $S^4 =_{\text{SL}_2(\mathbb{Z})} I$ is not used until the last step. Rearranging we have

$$A_0 ST^{a_0} \dots ST^{a_{d-1}} A_0^{-1} T^{-c} =_G S^{4k}. \quad (6.4)$$

To complete the proof we must only show that $k = 1$ in Equation (6.4).

Let $\sigma, \eta \in G$ be given by

$$\sigma =_G ST^{a_0} \dots ST^{a_{d-1}} \text{ and } \eta =_G A_0 \sigma A_0^{-1} T^{-c}.$$

Since $W(S^{4k})$ it is sufficient to show that $W(\eta) = 1$. By Lemma 4.4.2 $\sigma =_{\text{SL}_2(\mathbb{Z})} I$ implies that $W(\eta) = \text{wind}(\rho(\eta))$. By Lemma 6.4.3 write

$$(\rho(\eta))(t) = \begin{cases} \rho(A_0^{-1})(4t), & 0 \leq t \leq 1/4 \\ [(\rho(\sigma))(4t-1)] A_0^{-1}, & 1/4 \leq t \leq 1/2 \\ \sigma A_0^{-1} [(\rho(T^{-c}))(4t-2)], & 1/2 \leq t \leq 3/4 \\ [(\rho(A_0))(4t-3)] \sigma A_0^{-1} T^{-c}, & 3/4 \leq t \leq 1. \end{cases}$$

Let $\gamma_0: [0, 1] \rightarrow \mathrm{SL}_2(\mathbb{R})$ be the path from A_0^{-1} to itself given by

$$\gamma_0(t) = \begin{cases} [(\rho(\sigma))(2t)]A_0^{-1}, & 0 \leq t \leq 1/2 \\ \sigma A_0^{-1} [(\rho(T^{-c}))(2t-1)], & 1/2 \leq t \leq 1. \end{cases} \quad (6.5)$$

The paths γ_0 and $\rho(\eta)$ are homotopic via the homotopy

$$(\rho(\eta))(t) = \begin{cases} \rho(A_0^{-1})\left(\frac{4t}{s}\right), & 0 \leq t \leq s/4 \\ \left[(\rho(\sigma))\left(\frac{4t-s}{2-s}\right)\right]A_0^{-1}, & s/4 \leq t \leq 1/2 \\ \sigma A_0^{-1} \left[\left(\rho(T^{-c})\right)\left(\frac{4t-2}{2-s}\right)\right], & 1/2 \leq t \leq \frac{4-s}{4} \\ \left[(\rho(A_0))\left(\frac{4t+s-4}{s}\right)\right]\sigma A_0^{-1}T^{-c}, & \frac{4-s}{4} \leq t \leq 1 \end{cases}$$

for $0 < s \leq 1$ where γ_0 is defined as above. Thus, to complete the proof we only must show $\mathrm{wind}(\gamma_0) = 1$ where γ_0 is as in Equation 6.5. By Lemma 6.4.4, the path

$$\left(\mathrm{pr}(\rho(\sigma)(2t))\right)_{0 \leq t \leq 1/2}$$

is homotopic to a path which travels counter-clockwise at most one full rotation. The path

$$\left[\mathrm{pr}(\sigma A_0^{-1} \rho(T^{-c})(2t-1))\right]_{1/2 \leq t \leq 1},$$

travels only counter-clockwise and cannot cross the line $\{y = 0\}$, so it completes at most one half-rotation. Since $\sigma A_0^{-1}T^{-c} =_{\mathrm{SL}_2(\mathbb{Z})} A_0^{-1}$, the path γ_0 thus circles the origin an integer number of times, so we conclude that $\mathrm{wind}(\mathrm{pr}(\gamma_0)) = \mathrm{wind}(\gamma_0) = 1$. This completes the proof. \square

Proof of Proposition 6.1.10. Let $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ be a semitoric helix of length d and complexity c . Then there exists associated integers $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ as in Definition 6.4.2 by Lemma 6.4.1. If $\{w_i\}_{i \in \mathbb{Z}} \in \mathcal{H}$ then by Definition 6.1.3 there exists some $k, \ell \in \mathbb{Z}$ such that $v_i = T^k w_{i+\ell}$ for all $i \in \mathbb{Z}$. In this case $a_i v_{i+1} = v_i + v_{i+2}$ implies that

$$a_i w_{i+1+\ell} = w_{i+\ell} + w_{i+2+\ell}$$

and denoting $a_j := a_{j \bmod d}$ this implies that

$$a_{i-\ell} w_{i+1} = w_i + w_{i+2}.$$

Thus, the associated integers for $\{w_i\}_{i \in \mathbb{Z}}$ are given by

$$(a_{-\ell}, a_{1-\ell}, \dots, a_{d-1-\ell})$$

which agrees with those integers for $\{v_i\}_{i \in \mathbb{Z}}$ up to cyclic permutation. Thus, the map is well-defined.

Suppose $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$ is a list of integers satisfying

$$ST^{a_0} \dots ST^{a_{d-1}} =_G S^d X^{-1} T^{-c} X$$

for some $c \in \mathbb{Z}_{>0}$. Let $A_0 \in \mathrm{SL}_2(\mathbb{Z})$ be any matrix satisfying $X =_{\mathrm{SL}_2(\mathbb{Z})} A_0$ and define $v_0, v_1 \in \mathbb{Z}^2$ so that $A_0 = [v_0, v_1]$. Then define v_2, \dots, v_{d-1} by

$$v_i = a_{i-2} v_{i-1} - v_{i-2}$$

for $i = 2, \dots, d-1$. Use the relationship $v_{i+d} = T^c v_i$ to extend v_0, \dots, v_{d-1} to $\{v_i\}_{i \in \mathbb{Z}}$. Since $W(ST^{a_0} \dots ST^{a_{d-1}}) = 1$, the vectors v_0, \dots, v_{d-1} are in counter-clockwise order and by construction $\det(v_i, v_{i+1}) = 1$ for all $i \in \mathbb{Z}$, so $[\{v_i\}_{i \in \mathbb{Z}}]$ is a semitoric helix with the prescribed associated integers.

If \mathcal{H} and \mathcal{H}' are semitoric helices of the same length and complexity which are such that $\mathcal{H} = B \mathcal{H}'$ for some $B \in \mathrm{SL}_2(\mathbb{Z})$ then they have the same associated integers since those integers are defined by a linear equation, Equation (6.2), which is invariant under the action of B .

Conversely, suppose that $\mathcal{H} = [\{v_i\}_{i \in \mathbb{Z}}]$ and $\mathcal{H}' = [\{v'_i\}_{i \in \mathbb{Z}}]$ are semitoric helices of the same length and complexity. Let $B \in \mathrm{SL}_2(\mathbb{Z})$ be the matrix which satisfies

$$[v_0, v_1] = B[v'_0, v'_1].$$

Then act on both sides of Equation (6.2) with B^{-1} to discover that \mathcal{H}' has the same associated integers. □

6.5 Standard form in $\mathrm{PSL}_2(\mathbb{Z})$ and the winding number

First we prove several lemmas which will be needed in the proof of Theorem 6.1.12.

Lemma 6.5.1. *If $\sigma \in \langle S, T \rangle$ is S -positive and $\sigma =_{\mathrm{PSL}_2(\mathbb{Z})} I$ then $W(\sigma) \geq 0$ where $W(\sigma) = 0$ if and only if σ is the empty word.*

Proof. Since σ is S -positive up to conjugation by T we may write it as

$$\sigma =_{\langle S, T \rangle} ST^{a_0} \dots ST^{a_{d-1}}$$

for some $a_0, \dots, a_{d-1} \in \mathbb{Z}$. We define a sequence of vectors

$$v_0, \dots, v_{d-1} \in \mathbb{Z}^2$$

by choosing any $v_0, v_1 \in \mathbb{Z}^2$ with $\det(v_0, v_1) = 1$ and defining v_2, \dots, v_{d-1} by

$$v_{i+2} = -v_i + a_i v_{i+1}$$

for $i = 0, \dots, d-3$. Let $\gamma: [0, 1] \rightarrow (\mathbb{R}^2)^*$ be a path which connects v_0, \dots, v_{d-1} in order and travels only counter-clockwise. Then, $W(\sigma) = \mathrm{wind}(\gamma)$ and $\mathrm{wind}(\gamma) > 0$ because γ must travel at least once around the origin to move only counter-clockwise and return to $\gamma(0)$. \square

This implies that each element of $\mathrm{PSL}_2(\mathbb{Z})$ has a representation in S and T with minimal winding number.

Lemma 6.5.2. *If $X \in \mathrm{PSL}_2(\mathbb{Z})$ then there exists some $q \in \frac{1}{12}\mathbb{Z}$ such that $w(\sigma) \geq q$ for all $\sigma \in \langle S, T \rangle$ which are S -positive and satisfy $\sigma =_{\mathrm{PSL}_2(\mathbb{Z})} X$.*

Proof. Since $S =_{\mathrm{PSL}_2(\mathbb{Z})} S^{-1}$ every element of $\mathrm{PSL}_2(\mathbb{Z})$ has an S -positive representation. Fix some S -positive $\eta \in \langle S, T \rangle$ such that $\eta =_{\mathrm{PSL}_2(\mathbb{Z})} X^{-1}$ and let $q = -W(\eta)$. Let σ be any S -positive element of $\langle S, T \rangle$ such that $\sigma =_{\mathrm{PSL}_2(\mathbb{Z})} X$. Now $\sigma\eta =_{\mathrm{PSL}_2(\mathbb{Z})} I$, so $W(\sigma\eta) \geq 0$ by Lemma 6.5.1. This means $W(\sigma) \geq q$ and the result follows because $W(\eta)$ does not depend on the choice of σ . \square

6.5.1 Standard form for elements of $\mathrm{PSL}_2(\mathbb{Z})$

In this section we prove Theorem 6.1.12.

Lemma 6.5.3. $ST^{-n}S =_{\mathrm{PSL}_2(\mathbb{Z})} (TST)^n$ for $n \geq 0$.

Proof. First notice $STS =_{\mathrm{PSL}_2(\mathbb{Z})} T^{-1}ST^{-1}$ implies $S =_{\mathrm{PSL}_2(\mathbb{Z})} TSTST$ so

$$ST^{-1}S =_{\mathrm{PSL}_2(\mathbb{Z})} TST$$

since $S =_{\mathrm{PSL}_2(\mathbb{Z})} S^{-1}$. Now,

$$ST^{-n}S =_{\mathrm{PSL}_2(\mathbb{Z})} (ST^{-1}S)^n =_{\mathrm{PSL}_2(\mathbb{Z})} (TST)^n$$

for $n > 0$, and if $n = 0$ the claim reduces to $S^2 =_{\mathrm{PSL}_2(\mathbb{Z})} I$. □

Proof of Theorem 6.1.12. Let $\sigma \in \langle S, T \rangle$ any S -positive word with $\sigma =_{\mathrm{PSL}_2(\mathbb{Z})} X$. There are four steps to the reduction algorithm we will use on σ .

1. replace each S^2 with I ;
2. replace each $ST^{-n}S$ with $(TST)^n$, for some $n > 0$;
3. replace each STS with $T^{-1}ST^{-1}$;
4. go back to Step (1) if Step (1) or Step (2) is possible.

Each of these reductions preserves the value of σ in $\mathrm{PSL}_2(\mathbb{Z})$ and recall that the winding number cannot decrease indefinitely by Lemma 6.5.2. Steps (1) and (2) reduce the winding number while Step (4) can only be performed if the winding number is not already minimal and Step (3) preserves the winding number but reduces the number of times S appears in the word, which is also bounded below. Thus, this process must terminate and after the reduction the word will be of the required form.

Now we will show uniqueness. Suppose that

$$\sigma, \eta \in \langle S, T \rangle$$

with $\sigma =_{\text{PSL}_2(\mathbb{Z})} \eta$ and

$$\begin{aligned}\sigma &=_{\langle S, T \rangle} T^b ST^{a_0} \dots ST^{a_{d-1}}, \\ \eta &=_{\langle S, T \rangle} T^{b'} ST^{a'_0} \dots ST^{a'_{d'-1}},\end{aligned}$$

where $a_i, a'_j > 1$ for $i = 0, \dots, d-2$ and $j = 0, \dots, d'-2$. First assume $\min(d, d') \leq 1$, and in this case assume $d \geq d'$.

If $d' = 0$ then

$$T^{b-b'} ST^{a_0} \dots ST^{a_{d-1}} =_{\text{PSL}_2(\mathbb{Z})} I$$

which contradicts Lemma 4.3.8 unless $d = 0$, in which case $T^{b-b'} =_{\text{PSL}_2(\mathbb{Z})} I$ so $b = b'$.

If $d' = 1$ then

$$T^{b-b'} ST^{a_0} \dots ST^{a_{d-1}-a'_0} ST^0 =_{\text{PSL}_2(\mathbb{Z})} I$$

so $a_{d-1} - a'_0 \in \{0, \pm 1\}$ by Lemma 4.3.8. Consider the cases if $d > 1$. If $a_{d-1} - a'_0 = 0$ then

$$T^{b-b'} ST^{a_0} \dots ST^{a_{d-2}} S^2 =_{\text{PSL}_2(\mathbb{Z})} I$$

which contradicts Lemma 4.3.8 after replacing S^2 by I . If $a_{d-1} - a'_0 = -1$ then

$$T^{b-b'} ST^{a_0} \dots ST^{a_{d-2}} ST^{-1} S =_{\text{PSL}_2(\mathbb{Z})} I$$

which contradicts Lemma 4.3.8 after replacing $ST^{-1}S$ by TST . Finally, if $a_{d-1} - a'_0 = 1$, then

$$T^{b-b'} ST^{a_0} \dots ST^{a_{d-2}-1} ST^{-1} =_{\text{PSL}_2(\mathbb{Z})} I$$

which contradicts Lemma 4.3.8 unless $a_{d-2} = 2$. This process is repeated to conclude that $a_0 = \dots = a_{d-2} = 2$ so

$$T^{b-b'} (ST^2)^{d-1} STS =_{\text{PSL}_2(\mathbb{Z})} I$$

which implies

$$T^{b-b'-1} ST^{-d} =_{\text{PSL}_2(\mathbb{Z})} I.$$

By Lemma 4.3.8 this can not hold. Thus, $d = 1$, in which case

$$T^{b-b'} ST^{a_0-a'_0} S =_{\text{PSL}_2(\mathbb{Z})} I,$$

so $b - b' = 0$ and $a_0 - a'_0 = 0$ by Lemma 4.3.8.

Finally, assume $d, d' > 1$ and assume that $a_{d-1} \neq a'_{d-1}$, otherwise cancel them from both sides. In this case we see that $XY^{-1} =_{\text{PSL}_2(\mathbb{Z})} I$ implies

$$T^b ST^{a_0} \dots ST^{a_{d-1}} T^{a'_{d-1}} ST^{a'_{d-2}} \dots ST^{a'_0} S =_{\text{PSL}_2(\mathbb{Z})} I$$

so $a_{d-1} - a'_{d-1} = \pm 1$, by Lemma 4.3.8, since we have assumed $a_{d-1} - a'_{d-1} \neq 0$. Assume

$$a_{d-1} - a'_{d-1} = 1,$$

otherwise exchange X and Y . Then choose maximal $k \in \mathbb{Z}_{\geq 0}$ such that

$$a_{d-2} = a_{d-3} = \dots = a_{d-2-(k-1)} = 2.$$

where $k = 0$ if $a_{d-2} \neq 2$. If $k < d - 1$ then

$$\begin{aligned} XY^{-1} &=_{\text{PSL}_2(\mathbb{Z})} T^b ST^{a_0} \dots ST^{a_{d-2-k}} (ST^2)^k (STS) T^{-a'_{d-2}} \dots ST^{-a'_0} ST^{-b'} \\ &=_{\text{PSL}_2(\mathbb{Z})} T^b ST^{a_0} \dots ST^{a_{d-2-k-1}} (TST)^k ST^{-a'_{d-2}-1} \dots ST^{-a'_0} ST^{-b'} \\ &=_{\text{PSL}_2(\mathbb{Z})} T^b ST^{a_0} \dots ST^{a_{d-2-k-1}} ST^{-a'_{d-2}-k-1} \dots ST^{-a'_0} ST^{-b'}. \end{aligned}$$

Since $a_{d-2-k} - 1 > 1$ and $-a'_{d-2} - k - 1 < -1$ this expression cannot evaluate to the identity in $\text{PSL}_2(\mathbb{Z})$ by Lemma 4.3.8.

Otherwise, $k = d - 1$, in which case

$$\begin{aligned} XY^{-1} &=_{\text{PSL}_2(\mathbb{Z})} T^b (ST^2)^{d-1} (STS) T^{-a'_{d-2}} \dots ST^{-a'_0} ST^{-b'} \\ &=_{\text{PSL}_2(\mathbb{Z})} T^{b-1} (TST)^{d-1} ST^{-a'_{d-2}-1} \dots ST^{-a'_0} ST^{-b'} \\ &=_{\text{PSL}_2(\mathbb{Z})} T^{b-1} ST^{-a'_{d-2}-d} \dots ST^{-a'_0} ST^{-b'}. \end{aligned}$$

which again cannot evaluate to the identity in $\text{PSL}_2(\mathbb{Z})$ by Lemma 4.3.8. This completes the proof of uniqueness.

Lastly, we will show the standard form has minimal winding number. Let $X \in \text{PSL}_2(\mathbb{Z})$ and suppose $\eta \in \langle S, T \rangle$ is S -positive with $\eta =_{\text{PSL}_2(\mathbb{Z})} X$. Then η can be reduced to the standard form of X , denoted \bar{X} , by following the reduction algorithm at the beginning of the proof. Since each step of the algorithm either preserves or reduces the winding number, $W(\bar{X}) \leq W(\eta)$. \square

6.5.2 Standard forms and the winding number

Recall that given any $X \in \text{PSL}_2(\mathbb{Z})$ we denote by

$$\bar{X} \in \langle S, T \rangle$$

the standard form of X , as given in Theorem 6.1.12.

Lemma 6.5.4. *If $X \in \text{PSL}_2(\mathbb{Z}) \setminus \{T^k\}_{k \in \mathbb{Z}}$ then*

$$W(\bar{X}) + W(\bar{X}^{-1}) = \frac{1}{2}.$$

Proof. Write

$$\bar{X} = T^b ST^{a_0} \dots ST^{a_{d-1}}$$

and since $X \neq T^k$ for any $k \in \mathbb{Z}$, $d > 0$. Now, $W((\bar{X})^{-1}) = -W(\bar{X})$ where

$$(\bar{X})^{-1} = S^{-1} T^{-a_{d-1}} \dots S^{-1} T^{-a_0} S^{-1} T^{-b}.$$

We will reduce $(\bar{X})^{-1}$ to standard form and keep track of the winding number. Replacing each

S^{-1} by S increases the winding number by $d/2$. Now replace each $ST^{-a_i}S$ with $(TST)^{a_i}$ for each even index i which at most increases the odd indexed powers of T by 2. Since each $a_i \geq 2$ for $i = 0, \dots, d-2$ we do the replacement $ST^{-a_i+2}S = (TST)^{a_i-2}$ for odd $0 < i < d-3$ and the replacement $ST^{-a_i+1}S = (TST)^{a_i-1}$ for $i = 1$ and the highest odd $i \leq d-2$. Thus we have now used $ST^{-n}S = (TST)^n$, for varying values of $n > 0$, a total of $d-1$ times decreasing W by $1/2$ each time. The word produced in this way is now in standard form so it is equal to $\overline{X^{-1}}$ and

$$W(\overline{X^{-1}}) = -W(\overline{X}) + \frac{d}{2} - \frac{d-1}{2} = -W(\overline{X}) + \frac{1}{2}$$

as desired. □

We can now prove that in many cases the first power of T in \overline{X} and the last power of T in $\overline{X^{-1}}$ must sum to 1.

Lemma 6.5.5. *For $X \in \text{PSL}_2(\mathbb{Z})$ write*

$$\overline{X} =_{\langle S, T \rangle} T^b ST^{a_0} \dots ST^{a_{d-1}}$$

and

$$\overline{X^{-1}} =_{\langle S, T \rangle} T^{b'} ST^{a'_0} \dots ST^{a'_{d-1}}.$$

Then

$$a_{d-1} + b' = a'_{d-1} + b = 0$$

if $X =_{\text{PSL}_2(\mathbb{Z})} T^k ST^a$ or $X =_{\text{PSL}_2(\mathbb{Z})} T^k$ for some $k, a \in \mathbb{Z}$, and

$$a_{d-1} + b' = a'_{d-1} + b = 1$$

otherwise.

Proof. The cases of $X =_{\text{PSL}_2(\mathbb{Z})} T^k ST^a$ and $X =_{\text{PSL}_2(\mathbb{Z})} T^k$ are easily checked. Suppose X is not of that form. Since $\overline{X^{-1}} \overline{X} =_{\text{SL}_2(\mathbb{Z})} I$ by Lemma 4.3.8 some power of T that is not at the front or end of the word must be $-1, 1$, or 0 . Since \overline{X} and $\overline{X^{-1}}$ are in standard form, $X \neq_{\text{PSL}_2(\mathbb{Z})} T^k ST^q$, and $X \neq_{\text{PSL}_2(\mathbb{Z})} T^k$ this means that $a'_{d-1} + b \in \{\pm 1, 0\}$.

If $a'_{d'-1} + b = 0$ then S^2 is a subword of $\overline{X^{-1}X}$ which can be replaced by I and if $a'_{d'-1} + b = -1$ then $ST^{-1}S$ is a subword of $\overline{X^{-1}X}$ which can be replaced by TST . In either case this means that

$$w(\overline{X^{-1}X}) \leq W(\overline{X^{-1}}) + W(\overline{X}) - \frac{1}{2} = 0$$

where the last equality is by Lemma 6.5.4. By Lemma 6.5.1 $w(\overline{X^{-1}X}) \geq 0$ with equality only when $X = I$. Since $X \neq I$ we must have $a'_{d'-1} + b = 1$. The same analysis on $\overline{X} \overline{X^{-1}}$ implies that $a_{d-1} + b' = 1$. \square

Lemma 6.5.6. *Let $X \in \text{PSL}_2(\mathbb{Z})$ and $c \in \mathbb{Z}_{>0}$. Then $W(\overline{X^{-1}T^cX}) = W(\overline{X^{-1}T^cX})$.*

Proof. If $X =_{\text{PSL}_2(\mathbb{Z})} T^k$ for some $k \in \mathbb{Z}$ then

$$\overline{X^{-1}T^cX} =_{\langle S, T \rangle} \overline{X^{-1}T^cX} =_{\langle S, T \rangle} T^c$$

so the result holds. If $X =_{\text{PSL}_2(\mathbb{Z})} T^k ST^a$ for some $k, a \in \mathbb{Z}$ then there are two cases. If $c > 1$ then

$$\overline{X^{-1}T^cX} =_{\langle S, T \rangle} \overline{X^{-1}T^cX} =_{\langle S, T \rangle} T^k ST^c ST^{-k}$$

so the result holds. If $c = 1$ then

$$\overline{X^{-1}T^cX} =_{\langle S, T \rangle} T^{k-1} ST^{-k-1}$$

while

$$\overline{X^{-1}T^cX} =_{\langle S, T \rangle} T^k STST^{-k}$$

and the result still holds.

If $X \neq T^k$ and $X \neq T^k ST^a$ for all $k, a \in \mathbb{Z}$, then $\overline{X^{-1}T^cX}$ is already in standard form for any $c > 0$ by Lemma 6.5.5, so

$$\overline{X^{-1}T^cX} =_{\langle S, T \rangle} \overline{X^{-1}T^cX}.$$

\square

6.6 Minimal models for semitoric helixes

Definition 6.6.1. An S -positive word with no leading T

$$ST^{a_0} \dots ST^{a_{d-1}} \in \langle S, T \rangle$$

is *minimal* if $a_0, \dots, a_{d-1} \neq 1$.

Minimal words are those associated to minimal helixes.

Lemma 6.6.2. Suppose $\sigma = ST^{a_0} \dots ST^{a_{d-1}} \in \langle S, T \rangle$ is minimal and there exists $X \in G \setminus \{S^{2\ell}T^k\}_{\ell, k \in \mathbb{Z}}$ such that

$$\sigma =_G S^4 X^{-1} T^c X.$$

Then, after cyclically reordering a_0, \dots, a_{d-1} if necessary, $a_0 \leq 0$ and one of the following hold:

- (i) $a_0 = 0$ and $\bar{\sigma} =_{\langle S, T \rangle} T^{a_1} ST^{a_2} \dots ST^{a_{d-1}}$;
- (ii) $a_0 < 0$ and $\bar{\sigma} =_{\langle S, T \rangle} (TST)^n T^{a_1} ST^{a_2} \dots ST^{a_{d-1}}$.

Proof. Notice

$$W(\sigma) = W(S^4 X^{-1} T^c X) = 1 - \frac{c}{12}$$

while

$$\begin{aligned} W(\bar{\sigma}) &= W(\overline{X^{-1} T^c X}) = W(\overline{X^{-1} T^c X}) \\ &= W(\overline{X^{-1}}) + W(\overline{X}) - \frac{c}{12} = \frac{1}{12} - \frac{c}{12} \end{aligned}$$

by Lemmas 6.5.6 and 6.5.4 since $X \notin_{\text{PSL}_2(\mathbb{Z})} T^k$ for any $k \in \mathbb{Z}$. Thus, $W(\sigma) \neq W(\bar{\sigma})$ so σ is not in standard form. This means that $a_j \leq 1$ for some fixed $j \in \{0, \dots, d-2\}$ and since σ is minimal this implies $a_j \leq 0$.

If $a_j = 0$ for any $j \in \{0, \dots, d-2\}$ then reorder so that $a_0 = 0$ so

$$\sigma = S^2 T^{a_1} ST^{a_2} \dots ST^{a_{d-1}}$$

and notice that

$$\eta = T^{a_1}ST^{a_2} \dots ST^{a_{d-1}}$$

satisfies $W(\eta) = W(\bar{\sigma})$. All steps in the reduction algorithm in the proof of Theorem 6.1.12 reduce the winding number, except for the blowdown $STS \rightarrow T^{-1}ST^{-1}$, so the only possible step to reduce η into standard form is a blowdown. For a blowdown to be possible we must have $a_j = 1$ for some $j \in \{1, \dots, d-1\}$, contradicting the minimality of σ . Thus, $\bar{\sigma} =_{\langle S, T \rangle} \eta$.

Otherwise, $a_j \neq 0$ for all $j \in \{0, \dots, d-1\}$ so, after cyclically reordering, $a_0 = -n < 0$. In this case let

$$\eta' = (TST)^n T^{a_1}ST^{a_2} \dots ST^{a_1}$$

and notice $\eta' =_{\text{PSL}_2(\mathbb{Z})} \sigma$. Again, $W(\eta') = W(\bar{\sigma})$ so the only possible reduction move would be a blowdown, but if a blowdown could be performed on η' that would contradict the minimality of σ , except in the case that $a_1 = 0$, which we have assumed does not occur. Thus, $\bar{\sigma} = \eta'$. \square

Here we classify all words associated to minimal semitoric helices. Recall \mathcal{S} from Equation (6.1).

Lemma 6.6.3 (Classification of minimal words). *The word $\sigma \in \langle S, T \rangle$ is associated to a minimal semitoric helix \mathcal{H} of complexity $c > 0$ if and only if σ is exactly one of the following, where $A_0 = [v_0, v_1]$ and $\{v_i\}_{i \in \mathbb{Z}} \in \mathcal{H}$:*

| type | $\sigma \in \langle S, T \rangle$ | c | A_0 |
|------|--|------------|--|
| (1) | $\sigma = ST^{-1}ST^{-4}$ | $c = 1$ | $ST^{-2} = \begin{pmatrix} 0 & -1 \\ 1 & -2 \end{pmatrix}$ |
| (2) | $\sigma = ST^{-2}ST^{-2}$ | $c = 2$ | $ST^{-1} = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}$ |
| (3) | $\sigma = S^2T^aST^{-a-2}, a \neq 1, -3$ | $c = 1$ | $ST^{-a-1} = \begin{pmatrix} 0 & -1 \\ 1 & -a-1 \end{pmatrix}$ |
| (4) | $\sigma = ST^{-1}ST^{-1}ST^{c-1}$ | $c \neq 2$ | $I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ |
| (5) | $\sigma = S^2T^aST^cST^{-a}, a \neq \pm 1$ | $c \neq 1$ | $ST^{-a} = \begin{pmatrix} 0 & -1 \\ 1 & -a \end{pmatrix}$ |
| (6) | $\sigma = S^2T^aS^2T^{c-a}, a \neq 1, c-1$ | $c > 0$ | $I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ |
| (7) | $\sigma = S^2\overline{A_0^{-1}}T^c\overline{A_0}$ | $c > 0$ | $A_0 \in \mathcal{S}$ |

where $a \in \mathbb{Z}$ is a parameter.

Proof. Suppose that σ is minimal and associated to a semitoric helix \mathcal{H} of length d and complexity $c > 0$. By Lemma 6.4.5 there exists some $X \in G$ such that

$$\sigma =_G S^4 X^{-1} T^c X. \quad (6.6)$$

We will proceed by cases on X .

Case I: $X =_{\text{PSL}_2(\mathbb{Z})} T^k$ for some $k \in \mathbb{Z}$. This implies that

$$ST^{a_0} \dots ST^{a_{d-1}-c} =_G S^4,$$

and so $a_0, \dots, a_{d-2}, a_{d-1} - c$ are associated to a minimal toric fan. Such words are completely classified in Lemma 4.4.8 and we conclude either $d = 3$ and $a_0 = a_1 = a_2 - c = -1$, which is minimal only when $c \neq 2$, or $d = 4$ and $a_0 = a_2 = 0$, $a_3 = c - a_1$, which is minimal only when $a \neq 1, c - 1$. Thus σ is either of type (4) or (6).

Case II: $X \neq_{\text{PSL}_2(\mathbb{Z})} T^k$ for all $k \in \mathbb{Z}$. In light of Equation (6.6) apply Lemma 6.6.2 to σ and conclude that, after passing to an equivalent helix by cyclically permuting,

$$\sigma =_{\langle S, T \rangle} ST^{a_0} \dots ST^{a_{d-1}}$$

satisfies either

1. $a_0 = 0$; or
2. $a_j \neq 0$ for all $j = 0, \dots, d-1$ and $a_0 = -n < 0$.

If $a_0 = 0$ then

$$\bar{\sigma} = T^{a_1} ST^{a_2} \dots ST^{a_{d-1}}$$

and otherwise

$$\bar{\sigma} = (TST)^n T^{a_1} ST^{a_2} \dots ST^{a_{d-1}}.$$

We now have three further cases on X .

Case IIa: $X =_{\text{PSL}_2(\mathbb{Z})} T^k ST^a$ for $k, a \in \mathbb{Z}$. First assume $a_0 = 0$. If $c = 1$ then

$$\overline{X^{-1}T^cX} =_{\langle S, T \rangle} T^{-a-1}ST^{a-1}$$

so $\sigma = S^2T^{-a-1}ST^{a-1}$ which is minimal for $a \neq \pm 2$ and is of type (3). If $c \neq 1$ then

$$\overline{X^{-1}T^cX} = T^{-a}ST^cST^a$$

so $\sigma = S^2T^{-a}ST^a$ which is minimal if $a \neq \pm 1$ and is of type (5).

Now suppose $a_0 \neq 0$. Then

$$\bar{\sigma} = (TST)^n T^{a_1} ST^{a_2} \dots ST^{a_{d-1}}$$

for some $n > 0$. If $c = 1$ then $\bar{\sigma} = T^{-a-1}ST^{a-1}$ so $a = -2$ and thus $\bar{\sigma} = (TST)T^{-4}$ so $\sigma = ST^{-1}ST^{-4}$, which is of type (1). If $c = 2$, then $\bar{\sigma} = T^{-a}ST^2ST^a$ which means $a = -1$ and $a_0 = -2$ so $\sigma = ST^{-2}ST^{-2}$, which is of type (2). If $c > 2$, then $\bar{\sigma} = T^{-a}ST^cST^a$ which means $a = -1$ and $a_0 = -1$ so $\sigma = ST^{-1}ST^{c-1}ST^{-1}$, which is of type (4).

Case IIb: $X \neq_{\text{PSL}_2(\mathbb{Z})} S^{2\ell}T^k, S^{2\ell}T^kST^a$ for all $\ell, k \in \mathbb{Z}$ and $a_i \neq 0$ for all $i = 0, \dots, d-1$.

In this case $a_0 = -n < 0$. If $d = 2$, then $\sigma = ST^{-n}ST^{a_1}$ and $\bar{\sigma} = (TST)^n T^{a_1}$ which means $(TST)^n T^{a_1} =_{\langle S, T \rangle} \overline{X^{-1}T^cX}$. Since $\overline{X^{-1}T^cX}$ starts with TS it must end with S by Lemma 6.5.5 so $a_1 = -1$. Now $W(\sigma) = 1 - c/12$ implies that $n = 5 - c$ so we have

$$(TST)^{5-c}T^{-1} =_{\langle S, T \rangle} \overline{X^{-1}T^cX}. \quad (6.7)$$

The right side of Equation (6.7) contains T^{c+1} while the highest power of T on the left side is T^2 , so $c = 1$. Thus we obtain $\sigma = ST^{-4}ST^{-1}$ and have type (3).

If $d > 2$ then

$$(TST)^n T^{a_1} ST^{a_2} \dots ST^{a_{d-1}} =_{\langle S, T \rangle} \overline{X^{-1}T^cX}$$

implies that $\overline{X^{-1}T^cX}$ must end with S by Lemma 6.5.5, so $a_{d-1} = 0$ which contradicts our assumption in this case.

Case IIc: $X \notin_{\text{PSL}_2(\mathbb{Z})} S^{2\ell}T^k, S^{2\ell}T^kST^a$ for all $\ell, k \in \mathbb{Z}$ and $a_i = 0$ for some $i \in \{0, \dots, d-1\}$.

If $a_0 \neq 0$ then $a_j = 0$ for some $j \neq 0$ which contradicts the forms of the minimal word prescribed in Lemma 6.6.2. Thus, $a_0 = 0$ and so

$$\sigma = S^2\bar{\sigma} = S^2\overline{X^{-1}T^c\bar{X}}$$

which is minimal if \bar{X} does not end with ST and $\overline{X^{-1}}$ does not begin with TS , and is of type (7). □

Proof of Theorem 6.1.13. Let \mathcal{H} be a minimal semitoric helix of length d with associated integers $(a_0, \dots, a_{d-1}) \in \mathbb{Z}^d$. Then $\sigma = ST^{a_0} \dots ST^{a_{d-1}}$ is a minimal word and, passing to an equivalent helix if necessary, we conclude σ must be of some type (1)-(7) in Lemma 6.6.3. Types (1)-(6) for σ in Lemma 6.6.3 correspond to types (1)-(6) for \mathcal{H} in Theorem 6.1.13. Notice these each have length $d < 5$.

Otherwise, σ must be of type (7), which means there exists some $X =_G A_0$, where $\mathcal{H} = \{v_i\}_{i \in \mathbb{Z}}$ and $A_0 = [v_0, v_1] \in \mathcal{S}$, such that

$$\sigma =_{\langle S, T \rangle} S^2\overline{X^{-1}T^c\bar{X}}.$$

Since $A_0 \in \mathcal{S}$ notice that $A_0 = ST^{a_0} \dots ST^{a_{\ell-1}}$ with $\ell \geq 2$, which implies that σ has at least six occurrences of S , so the length d of \mathcal{H} satisfies $d \geq 6$. □

Lemma 6.6.4. *Any semitoric helix of complexity $c > 2$ includes the vector $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ or its negative.*

Proof. We will show that this vector is in every minimal semitoric helix of complexity $c > 2$, and since every semitoric helix can be produced by a sequence of blowups on a minimal semitoric helix, and since blowups do not remove vectors from the helix or change the complexity, the result will follow.

Since $c > 2$ the only possible types for minimal models are types (4)-(7). By Theorem 6.1.13 we see that (4), (5), and (6) include the required vector. Let $\mathcal{H} = \{v_i\}_{i \in \mathbb{Z}}$ be a semitoric helix of

type (7) with associated integers a_0, \dots, a_{d-1} and let $A_0 = [v_0, v_1]$. Then

$$S^2 \overline{A_0^{-1}} T^c \overline{A_0} =_{\langle S, T \rangle} ST^{a_0} \dots ST^{a_{d-1}}, \quad (6.8)$$

implies

$$S^2 \overline{A_0^{-1}} T^c =_{\langle S, T \rangle} ST^{a_0} \dots ST^{a_{k+1}} \quad (6.9)$$

for some $k \in \mathbb{Z}$, since $\overline{A_0}$ starts with S . By the recurrence relation

$$a_i v_{i+1} = v_i + v_{i+2}$$

we see

$$[v_{k+2}, v_{k+3}] =_{\text{SL}_2(\mathbb{Z})} A_0 ST^{a_0} \dots ST^{a_{k+1}}. \quad (6.10)$$

Combining Equations (6.9) and (6.10) yields

$$[v_{k+2}, v_{k+3}] =_{\text{PSL}_2(\mathbb{Z})} A_0 S^2 \overline{A_0^{-1}} T^c$$

which implies

$$[v_{k+2}, v_{k+3}] =_{\text{PSL}_2(\mathbb{Z})} T^c =_{\text{PSL}_2(\mathbb{Z})} \begin{pmatrix} 1 & c \\ 0 & 1 \end{pmatrix}$$

so v_{k+2} is the required vector. □

Theorem 6.2.1 follows from Lemma 6.6.4 because the existence of a horizontal vector in a semitoric helix implies that there must be a vertical edge on the moment polygon for the associated semitoric manifold.

Acknowledgements. Chapter 6, in part, is comprised of material in preparation for submission by the author of this dissertation, Daniel M. Kane, and Álvaro Pelayo as *Minimal models in semitoric geometry*.

Chapter 7

Symplectic G -capacities

7.1 Introduction

In this chapter I give a notion of symplectic capacity for symplectic G -manifolds, where G is any Lie group, which I call a *symplectic G -capacity*, and give nontrivial examples. Such a capacity retains the properties of a symplectic capacity (monotonicity, conformality, and an analogue of non-triviality) with respect to symplectic G -embeddings. Symplectic capacities are examples of symplectic G -capacities in the case that G is trivial. In analogy with symplectic capacities, symplectic G -capacities distinguish the symplectic G -type of symplectic G -manifolds. As a first example I construct an equivariant analogue of the Gromov radius where $G = \mathbb{R}^k$ as follows. Let $\text{Symp}^{2n,G}$ denote the category of $2n$ -dimensional symplectic G -manifolds. That is, an element of $\text{Symp}^{2n,G}$ is a triple (M, ω, ϕ) where (M, ω) is a symplectic manifold and $\phi: G \times M \rightarrow M$ is a symplectic G -action. Given integers $0 \leq k \leq m \leq n$ I define the (m, k) -equivariant Gromov radius

$$\begin{aligned} c_{\mathbb{B}}^{m,k} : \text{Symp}^{2n, \mathbb{R}^k} &\rightarrow [0, \infty] \\ (M, \omega, \phi) &\mapsto \sup\{ r > 0 \mid \mathbb{B}^{2m}(r) \xrightarrow{\mathbb{R}^k} M \}, \end{aligned} \tag{7.1}$$

where $\xrightarrow{\mathbb{R}^k}$ denotes a symplectic \mathbb{R}^k -embedding and $\mathbb{B}^{2m}(r) \subset \mathbb{C}^m$ is the standard $2m$ -dimensional ball of radius $r > 0$ with \mathbb{R}^k -action given by rotation of the first k coordinates.

Theorem 7.1.1. *If $k \geq 1$, the (m, k) -equivariant Gromov radius $c_B^{m,k} : \text{Symp}^{2n, \mathbb{R}^k} \rightarrow [0, \infty]$ is a symplectic \mathbb{R}^k -capacity.*

I prove Theorem 7.1.1 in Section 7.2.3. Thanks to the added structure of the \mathbb{R}^k -action the proof is elementary.

Let $\text{Symp}_T^{2n, \mathbb{T}^n}$ be the category of \mathbb{T}^n -manifolds which can be obtained from toric integrable systems. The morphisms of $\text{Symp}_T^{2n, \mathbb{T}^n}$ are \mathbb{T}^n -equivariant embeddings. Similarly, let $\text{Symp}_{\text{ST}}^{4, S^1 \times \mathbb{R}}$ be the category of $(S^1 \times \mathbb{R})$ -manifolds which can be obtained from semitoric integrable systems with $(S^1 \times \mathbb{R})$ -equivariant embeddings as the morphisms. Notice that \mathcal{M}_T is isomorphic to the quotient of $\text{Symp}_T^{2n, \mathbb{T}^n}$ by toric isomorphisms and translations of the momentum map while \mathcal{M}_{ST} is the quotient of $\text{Symp}_{\text{ST}}^{4, S^1 \times \mathbb{R}}$ by semitoric isomorphisms and translations of the momentum map.

As an application of symplectic G -capacities to integrable systems I define the *toric packing capacity*

$$\begin{aligned} \mathcal{J} : \text{Symp}_T^{2n, \mathbb{T}^n} &\rightarrow [0, \infty] \\ (M, \omega, \phi) &\mapsto \left(\frac{\sup\{\text{vol}(P) \mid P \text{ is a toric ball packing of } M\}}{\text{vol}(B^{2n})} \right)^{\frac{1}{2n}}, \end{aligned} \tag{7.2}$$

where $\text{vol}(E)$ denotes the symplectic volume of a subset E of a symplectic manifold, B^{2n} is the standard symplectic unit $2n$ -ball, $\text{Symp}_T^{2n, \mathbb{T}^n}$ is the category of $2n$ -dimensional symplectic toric manifolds, and a toric ball packing P of M is given by a disjoint collection of symplectically and \mathbb{T}^n -equivariantly embedded balls. In analogy I define the *semitoric packing capacity*

$$\mathcal{S}\mathcal{J} : \text{Symp}_{\text{ST}}^{4, S^1 \times \mathbb{R}} \rightarrow [0, \infty]$$

on $\text{Symp}_{\text{ST}}^{4, S^1 \times \mathbb{R}}$, the category of semitoric manifolds (see Section 2.2.7), where P in (7.2) is replaced by a semitoric ball packing of M (Definition 7.5.2).

Theorem 7.1.2. *The toric packing capacity $\mathcal{J} : \text{Symp}_T^{2n, \mathbb{T}^n} \rightarrow [0, \infty]$ is a symplectic \mathbb{T}^n -capacity and the semitoric packing capacity $\mathcal{S}\mathcal{J} : \text{Symp}_{\text{ST}}^{4, S^1 \times \mathbb{R}} \rightarrow [0, \infty]$ is a symplectic $(S^1 \times \mathbb{R})$ -capacity.*

In Chapter 8, I will study the continuity of capacities which are defined in the present chapter using the metric on toric integrable systems from [67] and the metric on semitoric integrable systems

from Chapter 3.

7.2 Symplectic G -capacities

Recall for $n \geq 1$ and $r > 0$, $B^{2n}(r) \subset \mathbb{C}^n$ denotes the $2n$ -dimensional open symplectic ball of radius r and

$$Z^{2n}(r) = \{ (z_i)_{i=1}^n \in \mathbb{C}^n \mid |z_1| < r \}$$

is the $2n$ -dimensional open symplectic cylinder of radius r . Both inherit a symplectic structure from their embedding as a subset of \mathbb{C}^n with symplectic form $\omega_0 = \frac{i}{2} \sum_{j=1}^n dz_j \wedge d\bar{z}_j$. We write $B^{2n} = B^{2n}(1)$, $Z^{2n} = Z^{2n}(1)$, and in this chapter \hookrightarrow will always denote a symplectic embedding.

7.2.1 Symplectic G -capacities

Let G be a Lie group and let $\text{Sympl}(M)$ denote the group of symplectomorphisms of the symplectic manifold (M, ω) . Recall, a smooth G -action $\phi: G \times M \rightarrow M$ is *symplectic* if $\phi(g, \cdot) \in \text{Sympl}(M)$ for each $g \in G$. The triple (M, ω, ϕ) is a *symplectic G -manifold*. A *symplectic G -embedding* $\rho: (M_1, \omega_1, \phi_1) \hookrightarrow (M_2, \omega_2, \phi_2)$ is a symplectic embedding for which there exists an automorphism $\Lambda: G \rightarrow G$ of G such that $\rho(\phi_1(g, p)) = \phi_2(\Lambda(g), \rho(p))$ for all $p \in M_1$, $g \in G$, in which case we say that ρ is a symplectic G -embedding with respect to Λ . We write \xrightarrow{G} to denote a symplectic G -embedding. We denote the collection of all $2n$ -dimensional symplectic G -manifolds by $\text{Sympl}^{2n, G}$. The set $\text{Sympl}^{2n, G}$ is a category with morphisms given by symplectic G -embeddings. We call a subcategory \mathcal{C}_G of $\text{Sympl}^{2n, G}$ a *symplectic G -category* if $(M, \omega, \phi) \in \mathcal{C}_G$ implies $(M, \lambda\omega, \phi) \in \mathcal{C}_G$ for any $\lambda \in \mathbb{R} \setminus \{0\}$. Let $\mathcal{C}_G \subset \text{Sympl}^{2n, G}$ be a symplectic G -category.

Definition 7.2.1. A *generalized symplectic G -capacity* on \mathcal{C}_G is a map $c: \mathcal{C}_G \rightarrow [0, \infty]$ satisfying:

1. *Monotonicity:* if $(M, \omega, \phi), (M', \omega', \phi') \in \mathcal{C}_G$ and there exists a symplectic G -embedding $M \xrightarrow{G} M'$ then $c(M, \omega, \phi) \leq c(M', \omega', \phi')$;
2. *Conformality:* if $\lambda \in \mathbb{R} \setminus \{0\}$ and $(M, \omega, \phi) \in \mathcal{C}_G$ then $c(M, \lambda\omega, \phi) = |\lambda| c(M, \omega, \phi)$.

When the symplectic form and G -action are understood we often write $c(M)$ for $c(M, \omega, \phi)$.

Let c be a generalized symplectic G -capacity on a symplectic G -category \mathcal{C}_G .

Definition 7.2.2. For $(N, \omega_N, \phi_N) \in \mathcal{C}_G$ we say that c satisfies N -*non-triviality* or is *non-trivial on* N if $0 < c(N) < \infty$.

Definition 7.2.3. We say that c is *tamed* by $(N, \omega_N, \phi_N) \in \text{Symp}^{2n,G}$ if there exists some $a \in (0, \infty)$ such that the following two properties hold:

- (1) if $M \in \mathcal{C}_G$ and there exists a symplectic G -embedding $M \xrightarrow{G} N$ then $c(M) \leq a$;
- (2) if $P \in \mathcal{C}_G$ and there exists a symplectic G -embedding $N \xrightarrow{G} P$ then $a \leq c(P)$.

The non-triviality condition in Definition 2.1.6 requires that $B^{2n}, Z^{2n} \in \mathcal{C}_G$ and $0 < c(B^{2n}) \leq c(Z^{2n}) < \infty$, and tameness encodes this second condition without necessarily including the first one. If c is a generalized symplectic G -capacity on $\mathcal{C}_G \subset \text{Symp}^{2n,G}$ we define

$$\text{Symp}_0^{2n,G}(c) = \{ N \in \text{Symp}^{2n,G} \mid \inf\{ c(P) \mid P \in \mathcal{C}_G, N \xrightarrow{G} P \} = 0 \},$$

$$\text{Symp}_\infty^{2n,G}(c) = \{ N \in \text{Symp}^{2n,G} \mid \sup\{ c(M) \mid M \in \mathcal{C}_G, M \xrightarrow{G} N \} = \infty \},$$

$$\text{Symp}_{\text{tame}}^{2n,G}(c) = \{ N \in \text{Symp}^{2n,G} \mid c \text{ is tamed by } N \}.$$

A generalized symplectic G -capacity gives rise to a decomposition of $\text{Symp}^{2n,G}$.

Proposition 7.2.4. *Let c be a generalized symplectic G -capacity on a symplectic G -category \mathcal{C}_G .*

Then:

- (a) $\text{Symp}^{2n,G} = \text{Symp}_0^{2n,G}(c) \cup \text{Symp}_\infty^{2n,G}(c) \cup \text{Symp}_{\text{tame}}^{2n,G}(c)$;
- (b) *the union in part (a) is pairwise disjoint;*
- (c) *c is non-trivial on $N \in \text{Symp}^{2n,G}$ if and only if $N \in \mathcal{C}_G \cap \text{Symp}_{\text{tame}}^{2n,G}(c)$.*

Proof. In order to prove item (a) we show that if $N \in \text{Symp}^{2n,G}$ is not in $\text{Symp}_0^{2n,G}(c) \cup \text{Symp}_\infty^{2n,G}(c)$ then it is in $\text{Symp}_{\text{tame}}^{2n,G}(c)$. If $M \xrightarrow{G} N \xrightarrow{G} P$ for some $M, P \in \mathcal{C}_G$ then $M \xrightarrow{G} P$ so $c(M) \leq c(P)$. Let $a_1 = \sup\{ c(M) \mid M \xrightarrow{G} N \}$ and $a_2 = \inf\{ c(P) \mid N \xrightarrow{G} P \}$. Since $N \notin \text{Symp}_0^{2n,G}(c) \cup \text{Symp}_\infty^{2n,G}(c)$ we have that $0 < a_1 \leq a_2 < \infty$. Pick $a \in [a_1, a_2]$. If $M \in \mathcal{C}_G$ and $M \xrightarrow{G} N$ then $c(M) \leq a_1 \leq a$ and if $P \in \mathcal{C}_G$ and $N \xrightarrow{G} P$ then $c(P) \geq a_2 \geq a$ so $N \in \text{Symp}_{\text{tame}}^{2n,G}(c)$. Item (b) follows from a similar argument and (c) is immediate. \square

In light of item (c) we view $\text{Symp}_{\text{tame}}^{2n,G}(c)$ as an extension of the set of elements of $\text{Symp}^{2n,G}$ on which c is non-trivial to include those elements outside of the domain of c .

7.2.2 Symplectic $(\mathbb{T}^k \times \mathbb{R}^{d-k})$ -capacities

For $1 \leq d \leq n$ the standard action of \mathbb{T}^d on \mathbb{C}^n is given by

$$\phi_{\mathbb{C}^n}((\alpha_i)_{i=1}^d, (z_i)_{i=1}^n) = (\alpha_1 z_1, \dots, \alpha_d z_d, z_{d+1}, \dots, z_n).$$

This action induces actions of $\mathbb{T}^d = \mathbb{T}^k \times \mathbb{T}^{d-k}$ on B^{2n} and Z^{2n} , which in turn induce the standard actions of $\mathbb{T}^k \times \mathbb{R}^{d-k}$ on B^{2n} and Z^{2n} for $k \leq d$. The action of an element of $\mathbb{T}^k \times \mathbb{R}^{d-k}$ is the action of its image under the quotient map $\mathbb{T}^k \times \mathbb{R}^{d-k} \rightarrow \mathbb{T}^d$. In the following we endow B^{2n} and Z^{2n} with the standard actions.

Definition 7.2.5. A generalized symplectic $(\mathbb{T}^k \times \mathbb{R}^{d-k})$ -capacity is a *symplectic $(\mathbb{T}^k \times \mathbb{R}^{d-k})$ -capacity* if it is tamed by B^{2n} and Z^{2n} .

7.2.3 A first example

The *Gromov radius* $c_B: \text{Symp}^{2n} \rightarrow (0, \infty]$ is given by

$$c_B(M) := \sup\{r > 0 \mid B^{2n}(r) \hookrightarrow M\}.$$

Fix $0 \leq k \leq m \leq n$ and let $c_B^{m,k}$ be as in Equation (7.1). If $k = 0$ and $m = n$ then $c_B = c_B^{m,k}$.

Proof of Theorem 7.1.1. Parts (1) and (2) of Definition 7.2.1 are immediate. By the standard inclusion map $c_B^{m,k}(B^{2n}) \geq 1$ so we only must show that $c_B^{m,k}(Z^{2n}) \leq 1$. Suppose that for $r > 1$ $\rho: B^{2m}(r) \xrightarrow{\mathbb{R}^k} Z^{2n}$ is a symplectic \mathbb{R}^k -embedding with respect to some $\Lambda \in \text{Aut}(\mathbb{R}^k)$. Let

$$(\eta_1, \dots, \eta_k) = \Lambda^{-1}(1, 0, \dots, 0).$$

Since Λ is an automorphism $\eta_{j_0} \neq 0$ for some $j_0 \in \{1, \dots, k\}$. Pick

$$w = (0, \dots, 0, w_{j_0}, 0, \dots, 0) \in B^{2m}(r)$$

with entries all zero except in the j_0^{th} position and such that $|w_{j_0}| > 1$. Let $u = (u_1, \dots, u_n) = \rho(w)$ and note $|u_1| < 1$. Let $\iota: \mathbb{R} \hookrightarrow \mathbb{R}^k$ be given by $\iota(x) = (x, 0, \dots, 0)$. Let $\phi_B: \mathbb{R}^k \times B^{2m}(r) \rightarrow B^{2m}(r)$

and $\phi_Z: \mathbb{R}^k \times Z^{2n} \rightarrow Z^{2n}$ be the standard actions of \mathbb{R}^k . Then for $x \in \mathbb{R}$

$$\rho(\phi_B(\Lambda^{-1} \circ \iota(x), w)) = \phi_Z(\iota(x), \rho(w)) = \phi_Z(\iota(x), u).$$

Thus

$$\rho(\{(0, \dots, e^{2ix\eta_{j_0}} w_{j_0}, 0, \dots, 0) \mid x \in \mathbb{R}\}) = \{(e^{2ix} u_1, u_2, \dots, u_n) \mid x \in \mathbb{R}\} \quad (7.3)$$

and since ρ is injective and $\eta_{j_0} \neq 0$ this means that $u_1 \neq 0$. Let

$$S_B = \{(0, \dots, 0, \alpha, 0, \dots, 0) \in B^{2m}(r) \mid |\alpha| < |w_{j_0}|\}$$

where α is in the j_0^{th} position and

$$S_Z = \{(\beta, u_2, \dots, u_n) \in Z^{2n} \mid |\beta| < |u_1|\}.$$

Equation (7.3) implies that $\rho(\partial S_B) = \partial S_Z$ and since ρ is an embedding this means $\partial(\rho(S_B)) = \partial S_Z$.

Since $\rho(S_B)$ and S_Z have the same boundary, ω_Z is closed, and Z^{2n} has trivial second homotopy group,

$$\int_{\rho(S_B)} \omega_Z = \int_{S_Z} \omega_Z.$$

Finally, integrating over z we have

$$\frac{i}{2} \int_{|z| < |w_j|} dz \wedge d\bar{z} = \int_{S_B} w_B = \int_{S_B} \rho^* \omega_Z = \int_{\rho(S_B)} \omega_Z = \int_{S_Z} \omega_Z = \frac{i}{2} \int_{|z| < |u_1|} dz \wedge d\bar{z}.$$

This implies that $1 < |w_j| = |u_1| < 1$, which is a contradiction. \square

It follows from the proof that $c_B^{m,k}(B^{2n}) = c_B^{m,k}(Z^{2n}) = 1$.

Proposition 7.2.6. *Let $M = (S^2)^n$ with symplectic form $\omega_M = \frac{1}{2} \sum_{i=1}^n dh_i \wedge d\theta_i$ where $h_i \in [-1, 1]$, $\theta_i \in [0, 2\pi)$, $i = 1, \dots, n$, are the standard height and angle coordinates. Let \mathbb{R}^k , $1 \leq k \leq n$, act on M by rotating the first k components. Then*

$$c_B^{m,k}(M) = \sqrt{2}$$

for all $m, k \in \mathbb{Z}$ with $1 \leq k \leq m \leq n$.

Proof. The map $\rho: B^{2n}(\sqrt{2}) \xrightarrow{\mathbb{R}^n} M$ given by

$$\rho(r_1 e^{i\theta_1}, \dots, r_n e^{i\theta_n}) = (\theta_1, r_1^2 - 1, \dots, \theta_n, r_n^2 - 1)$$

is a symplectic \mathbb{R}^n -embedding, so $c_B^{n,n}(M) \geq \sqrt{2}$.

Fix $k, m, n \in \mathbb{Z}$ satisfying $0 < k \leq m \leq n$ and let $\rho: B^{2m}(r) \xrightarrow{\mathbb{R}^k} M$ be a symplectic \mathbb{R}^k -embedding for some $r > 0$. Let

$$B_j = \{ (h_i, \theta_i)_{i=1}^n \in M \mid h_i \in \{\pm 1\} \text{ if } i \leq k \text{ and } i \neq j \}$$

for $j = 1, \dots, k$. For $R \in (0, r)$ let

$$A_R = \{ (z, 0, \dots, 0) \in B^{2m}(r) \mid |z| < R \}.$$

Every point in A_R , except at the identity, has the same $(k-1)$ -dimensional stabilizer in \mathbb{R}^k so there exists $j_0 \leq k$ such that $\rho(A_R) \subset B_{j_0}$ for all $R \in (0, r)$. Write $\rho = (H_i, \Theta_i)_{i=1}^n$ and consider coordinates (r, θ) on A_R given by $(re^{i\theta}, 0, \dots, 0) \rightarrow (r, \theta)$. For $i \neq j_0$ this means that H_i is constant if $i \leq k$ and the \mathbb{R}^n -equivariance of ρ implies that H_i and Θ_i are independent of θ if $i > k$. Thus if $i \in \{1, \dots, n\}$ and $i \neq j_0$ then

$$\int_{\rho(A_R)} dh_i \wedge d\theta_i = \int_{A_R} dH_i \wedge d\Theta_i = 0$$

for $R \in (0, r)$. Therefore,

$$\pi R^2 = \int_{A_R} \omega_B = \int_{\rho(A_R)} \omega_M = \frac{1}{2} \int_{\rho(A_R)} dh_{j_0} \wedge d\theta_{j_0} + \frac{1}{2} \sum_{i \neq j_0} \left(\int_{\rho(A_R)} dh_i \wedge d\theta_i \right) \leq \frac{1}{2} \int_{S^2} dh \wedge d\theta = 2\pi$$

for any $R \in (0, r)$. This implies that $r \leq \sqrt{2}$ so

$$\sqrt{2} \leq c_B^{n,n}(M) \leq c_B^{m,k}(M) \leq \sqrt{2}$$

for any $k, m, n \in \mathbb{Z}$ satisfying $0 < k \leq m \leq n$. □




Figure 7.1: A symplectic \mathbb{R} -embedding.

Example 7.2.7. For $k, n \in \mathbb{Z}_{>0}$ with $k < n$ let $M = \mathbb{Z}^{2n}$ with the standard symplectic form. There are two natural ways in which \mathbb{R}^k can act symplectically on M given by

$$\phi_1((t_i)_{i=1}^k, (z_i)_{i=1}^n) = (e^{2it_1} z_1, e^{2it_2} z_2, \dots, e^{2it_k} z_k, z_{k+1}, \dots, z_n)$$

and

$$\phi_2((t_i)_{i=1}^k, (z_i)_{i=1}^n) = (z_1, e^{2it_1} z_2, \dots, e^{2it_k} z_{k+1}, z_{k+2}, \dots, z_n)$$

where $\phi_i: \mathbb{R}^k \times M \rightarrow M$ for $i = 1, 2$. Let $\rho: M \rightarrow M$ be given by

$$\rho((z_i)_{i=1}^n) = \left(\frac{z_{k+1}}{1 + |z_{k+1}|}, \frac{z_1}{1 - |z_1|}, z_2, \dots, z_k, z_{k+2}, \dots, z_n \right)$$

similar to the map shown in Figure 7.1. The map ρ is well-defined because $|z_1| < 1$ and it is an \mathbb{R}^k -equivariant diffeomorphism because

$$\begin{aligned} \rho(\phi_1((t_i)_{i=1}^k, (z_i)_{i=1}^n)) &= \left(\frac{z_{k+1}}{1 + |z_{k+1}|}, e^{2it_1} \frac{z_1}{1 - |z_1|}, e^{2it_2} z_2, \dots, e^{2it_k} z_k, z_{k+2}, \dots, z_n \right) \\ &= \phi_2((t_i)_{i=1}^k, \rho((z_i)_{i=1}^n)) \end{aligned}$$

for all $t_1, \dots, t_k \in \mathbb{R}$. Thus the symplectic \mathbb{R}^k -manifolds (M, ω, ϕ_1) and (M, ω, ϕ_2) are symplectomorphic via the identity map and \mathbb{R}^k -equivariantly diffeomorphic via ρ but they are not \mathbb{R}^k -equivariantly symplectomorphic because $c_B^{1,1}(M, \omega, \phi_1) = 1$ and $c_B^{1,1}(M, \omega, \phi_2) = \infty$. ⊙

7.3 Hamiltonian $(\mathbb{T}^k \times \mathbb{R}^{n-k})$ -actions

In this section we review the facts we need for the remainder of the chapter about Hamiltonian $(\mathbb{T}^k \times \mathbb{R}^{n-k})$ -actions and their relation to toric and semitoric systems. Let (M, ω) be a symplectic manifold and G a Lie group with Lie algebra $\text{Lie}(G)$ and dual Lie algebra $\text{Lie}(G)^*$. Recall, a symplectic G -action is *Hamiltonian* if there exists a map $\mu: M \rightarrow \text{Lie}(G)^*$, known as the *momentum map*, such that

$$-d\langle \mu, X \rangle = \omega(X_M, \cdot)$$

for all $X \in \text{Lie}(G)$ where X_M denotes the vector field on M generated by X via the action of G . A *Hamiltonian G -manifold* is a quadruple (M, ω, ϕ, μ) where (M, ω, ϕ) is a symplectic G -manifold for which the action of G is Hamiltonian with momentum map μ . Let $\text{Ham}^{2n, G}$ denote the category of $2n$ -dimensional Hamiltonian G -manifolds with morphisms given by symplectic G -embeddings which intertwine the momentum maps. Let \mathcal{J}^{2n} denote the set of all $2n$ -dimensional integrable systems and define an equivalence relation $\sim_{\mathcal{J}}$ on this space by declaring (M, ω, F) and (M', ω', F') to be equivalent if there exists a symplectomorphism $\phi: M \rightarrow M'$ such that $F - \phi^*F': M \rightarrow \mathbb{R}^n$ is constant.

7.3.1 Hamiltonian \mathbb{R}^n -actions and integrable systems

Let $(M, \omega, F = (f_1, \dots, f_n))$ be an integrable system and for $i = 1, \dots, n$ let $\psi_i^t: M \rightarrow M$ denote the flow along \mathcal{X}_{f_i} , the Hamiltonian vector field of f_i . The *Hamiltonian flow action* $\phi_F: \mathbb{R}^n \times M \rightarrow M$, given by $\phi_F((t_1, \dots, t_n), p) = \psi_1^{t_1} \circ \dots \circ \psi_n^{t_n}(p)$, defines a Hamiltonian \mathbb{R}^n -action on M . The action of G on M is *almost everywhere locally free* if the stabilizer of p is discrete for almost all $p \in M$. Let $\mathcal{FSymp}^{2n, \mathbb{R}^n}$ be the space of \mathbb{R}^n -manifolds on which the action of \mathbb{R}^n is Hamiltonian and almost everywhere locally free and let $\sim_{\mathbb{R}^n}$ denote equivalence by \mathbb{R}^n -equivariant symplectomorphisms.

Lemma 7.3.1. *Let $\mathcal{X}_1, \dots, \mathcal{X}_n$ be vector fields with commuting flows on an m -manifold M , with $n \leq m$. Let \mathbb{R}^n act on M by $\phi((t_1, \dots, t_n), p) = \psi_1^{t_1} \circ \dots \circ \psi_n^{t_n}(p)$ where ψ_i^t is the flow of \mathcal{X}_i . Then, for $p \in M$, the vectors $(\mathcal{X}_1)_p, \dots, (\mathcal{X}_n)_p \in T_p M$ are linearly independent if and only if the stabilizer of p under the action ϕ is discrete.*

Proof. If $(\mathcal{X}_1)_p, \dots, (\mathcal{X}_n)_p$ are linearly independent then, since they have commuting flows, there is a chart (U, g) , with $U \subset M$ and $g: U \rightarrow \mathbb{R}^m$, such that $g^{-1}: g(U) \rightarrow U$ satisfies

$$g^{-1}(t_1, \dots, t_n, 0, \dots, 0) = \phi((t_1, \dots, t_n), p)$$

for any $(t_1, \dots, t_n, 0, \dots, 0) \in g(U)$. Thus $g(U)$ is an open neighborhood of the identity in \mathbb{R}^n and there exists no non-zero point in $g(U)$ which fixes p , so the stabilizer of p under the action of \mathbb{R}^n is discrete. On the other hand, if $(\mathcal{X}_1)_p, \dots, (\mathcal{X}_n)_p$ are linearly dependent, there exist $t_1, \dots, t_n \in \mathbb{R}$ not all zero such that $\sum_{i=1}^n t_i (\mathcal{X}_i)_p = 0$. Thus $(\alpha t_1, \dots, \alpha t_n) \in \mathbb{R}^n$ fixes p for all $\alpha \in \mathbb{R}$ and so the stabilizer of p is not discrete. \square

Proposition 7.3.2. *Let ψ be the map which takes an integrable system on M to M equipped with its Hamiltonian flow action. Then*

$$\psi: \mathcal{J}^{2n}/\sim_{\mathcal{J}} \rightarrow \mathcal{FSymp}^{2n, \mathbb{R}^n}/\sim_{\mathbb{R}^n}$$

is a bijection.

Proof. By Lemma 7.3.1 we know that the Hamiltonian flow action must be almost everywhere locally free because the Hamiltonian vector fields of an integrable system are by definition independent almost everywhere. Next suppose that \mathbb{R}^n acts Hamiltonianly on M in such a way that the action is almost everywhere locally free. Since the action is Hamiltonian there exists a momentum map $\mu: M \rightarrow \text{Lie}(\mathbb{R}^n)^*$. Define $F = (f_1, \dots, f_n): M \rightarrow \mathbb{R}^n$ by $F = A \circ \mu$ where $A: \text{Lie}(\mathbb{R}^n)^* \rightarrow \mathbb{R}^n$ is the standard identification which is induced by the standard basis $\{e_1, \dots, e_n\}$ of \mathbb{R}^n . These functions Poisson commute because action by the components of \mathbb{R}^n commute and are linearly independent at almost all points because the group action is almost everywhere locally free (Lemma 7.3.1). Thus, (M, ω, F) is an integrable Hamiltonian system. Let $\{v_1, \dots, v_n\}$ be the standard basis of $\text{Lie}(\mathbb{R}^n) \cong \mathbb{R}^n$ induced by the standard basis of \mathbb{R}^n . Let v_M denote the vector field on M generated by $v \in \text{Lie}(\mathbb{R}^n)$ via the action of G . Then $\langle \mu, v_i \rangle = f_i: M \rightarrow \mathbb{R}$ so $df_i = \omega((v_i)_M, \cdot)$ which means that the Hamiltonian vector field associated to f_i is $(v_i)_M$. Thus the Hamiltonian flow action related to F is the original action of \mathbb{R}^n . \square

Here we fix the identification between $\text{Lie}(\mathbb{T}^n)^*$ and \mathbb{R}^n that we will use for the remainder of the chapter. We specify our convention by choosing an epimorphism from \mathbb{R} to \mathbb{T}^1 , which we take to be $x \mapsto e^{2\sqrt{-1}x}$.

7.3.2 Hamiltonian \mathbb{T}^k -actions

We denote by $\text{Ham}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ the category of $2n$ -dimensional symplectic toric manifolds with morphisms as symplectic \mathbb{T}^n -embeddings and we denote equivalence by toric isomorphism by $\approx_{\mathbb{T}}$. In general being an invariant is weaker than being monotonic, but in the case of toric manifolds these are equivalent because symplectic \mathbb{T}^n -embeddings between toric manifolds are automatically \mathbb{T}^n -equivariant symplectomorphisms. Delzant proved [21] that in this case $\mu(M)$ is a *Delzant polytope*, i.e. simple, rational, and smooth, and that

$$\begin{aligned} \Psi: \text{Ham}_{\mathbb{T}}^{2n, \mathbb{T}^n} / \approx_{\mathbb{T}} &\rightarrow \mathcal{P}_{\mathbb{T}} \\ [(M, \omega, \phi, \mu)] &\mapsto \mu(M) \end{aligned}$$

is a bijection, where $\mathcal{P}_{\mathbb{T}}$ denotes the set of n -dimensional Delzant polytopes. Let $\text{Ham}^{2n, \mathbb{T}^n} \rightarrow \text{Symp}^{2n, \mathbb{T}^n}$ be given by $(M, \omega, \phi, \mu) \mapsto (M, \omega, \phi)$ and let $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ denote the image of $\text{Ham}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ under this map. Also let $\sim_{\mathbb{T}}$ denote equivalence on $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ by \mathbb{T}^n -equivariant symplectomorphisms.

7.3.3 Hamiltonian $(S^1 \times \mathbb{R})$ -actions

Let $(M_i, \omega_i, F_i = (J_i, H_i))$ be a semitoric integrable system for $i = 1, 2$. Let $\text{Ham}_{S^1 \times \mathbb{R}}^{4, S^1 \times \mathbb{R}}$ denote the category of simple semitoric systems and let $\approx_{S^1 \times \mathbb{R}}$ denote equivalence by semitoric isomorphism. Let $\text{Symp}_{S^1 \times \mathbb{R}}^{4, S^1 \times \mathbb{R}}$ denote the image of $\text{Ham}_{S^1 \times \mathbb{R}}^{4, S^1 \times \mathbb{R}}$ under the map $\text{Ham}^{4, S^1 \times \mathbb{R}} \rightarrow \text{Symp}^{4, S^1 \times \mathbb{R}}$ given by $(M, \omega, \phi, \mu) \mapsto (M, \omega, \phi)$ and let $\sim_{S^1 \times \mathbb{R}}$ denote the equivalence on $\text{Symp}_{S^1 \times \mathbb{R}}^{4, S^1 \times \mathbb{R}}$ inherited from $\sim_{S^1 \times \mathbb{R}}$ on $\text{Ham}_{S^1 \times \mathbb{R}}^{4, S^1 \times \mathbb{R}}$.

7.4 Symplectic \mathbb{T}^n -capacities

In this section we construct a symplectic \mathbb{T}^n -capacity on the space of symplectic toric manifolds. Recall $\mathcal{M}_{\mathbb{T}} \cong \text{Ham}_{\mathbb{T}}^{2n, \mathbb{T}^n} / \approx_{\mathbb{T}}$ is the moduli space of $2n$ -dimensional toric integrable systems

up to \mathbb{T}^n -equivariant symplectomorphisms which preserve the moment map. In [31, 64, 65, 68] the authors study the *toric optimal density function* $\Omega: \text{Ham}_{\mathbb{T}^n}^{2n, \mathbb{T}^n} / \approx_{\mathbb{T}^n} \rightarrow (0, 1]$, which assigns to each toric integrable system the fraction of that manifold which can be filled by equivariantly embedded disjoint open balls. This function is not a capacity because it is not monotonic or conformal. Next we study a modified version of this function which is a capacity.

For $M \in \text{Symp}^{2n, \mathbb{T}^n}$ by a \mathbb{T}^n -equivariantly embedded ball we mean the image $\phi(B^{2n}(r))$ of a symplectic \mathbb{T}^n -embedding $\phi: B^{2n}(r) \xrightarrow{\mathbb{T}^n} M$ for some $r > 0$. A *toric ball packing* of M [64] is a disjoint union $P = \bigsqcup_{\alpha \in \mathcal{A}} B_{\alpha}$ where $B_{\alpha} \subset M$ is a symplectically and \mathbb{T}^n -equivariantly embedded ball in M for each $\alpha \in \mathcal{A}$, where \mathcal{A} is some index set. That is, for each $\alpha \in \mathcal{A}$ there exists some $r_{\alpha} > 0$ and some symplectic \mathbb{T}^n -embedding $\phi_{\alpha}: B^{2n}(r_{\alpha}) \xrightarrow{\mathbb{T}^n} M$ such that

$$\phi_{\alpha}(B^{2n}(r_{\alpha})) = B_{\alpha}.$$

An example is shown in Figure 7.2. Recall the toric packing capacity $\mathcal{J}: \text{Symp}_{\mathbb{T}^n}^{2n, \mathbb{T}^n} \rightarrow [0, \infty]$ defined in Equation (7.2). In the following for $M \in \text{Symp}^{2n, \mathbb{T}^n}$ let $c_{\mathbb{B}}^{n, n}(M)$ be defined by first lifting the action of \mathbb{T}^n on M to an action of \mathbb{R}^n and applying the usual $c_{\mathbb{B}}^{n, n}$ to the resulting symplectic \mathbb{R}^n -manifold.




Figure 7.2: Toric ball packing of S^2 by symplectic \mathbb{T}^2 -disks.

Lemma 7.4.1. *Let $M \in \text{Symp}_{\mathbb{T}^n}^{2n, \mathbb{T}^n}$, $N \in \text{Symp}^{2n, \mathbb{T}^n}$ be such that the \mathbb{T}^n -action on N has $\ell \in \mathbb{Z}_{\geq 0}$ fixed points. If there is a symplectic \mathbb{T}^n -embedding $M \xrightarrow{\mathbb{T}^n} N$ then $\mathcal{J}(M) \leq \ell^{1/2n} c_{\mathbb{B}}^{n, n}(N)$.*

Proof. Since the center of $B^{2n}(r)$, $r > 0$, is a fixed point of the \mathbb{T}^n -action we see that the maximal number of such balls that can be simultaneously equivariantly embedded with disjoint images into M is the Euler characteristic $\chi(M)$ of M , which is the number of fixed points of the \mathbb{T}^n -action on M . Each of these balls has radius at most $c_{\mathbb{B}}^{n, n}(M)$. For $r > 0$ we have that $\text{vol}(B^{2n}(r)) = r^{2n} \text{vol}(B^{2n})$.

Therefore

$$(\mathcal{J}(M))^{2n} \text{vol}(\mathbb{B}^{2n}) \leq \chi(M) \text{vol}(\mathbb{B}^{2n}(c_{\mathbb{B}}^{n,n}(M))) = \chi(M) (c_{\mathbb{B}}^{n,n}(M))^{2n} \text{vol}(\mathbb{B}^{2n}).$$

Since \mathbb{T}^n -embeddings send fixed points to fixed points and $M \xrightarrow{\mathbb{T}^n} N$ we know that $\chi(M) \leq \ell$. Furthermore, since $M \xrightarrow{\mathbb{T}^n} N$ and $c_{\mathbb{B}}^{n,n}$ is a symplectic \mathbb{T}^n -capacity by Theorem 7.1.1 we have that $c_{\mathbb{B}}^{n,n}(M) \leq c_{\mathbb{B}}^{n,n}(N)$. Hence $\mathcal{J}(M) \leq \ell^{1/2n} c_{\mathbb{B}}^{n,n}(N)$. \square

Proposition 7.4.2. *The toric packing capacity is a symplectic \mathbb{T}^n -capacity on $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$.*

Proof. Let $M \in \text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ with $\chi(M) \in \mathbb{Z}_{\geq 0}$ fixed points and fix any ordering of these points. Notice that $\mathcal{J}(M)$ is the supremum of

$$\{ \|\vec{r}\|_{2n} \mid \vec{r} \in \mathbb{R}^{\chi(M)}, P_M(\vec{r}) \subset M \text{ is a toric packing} \}$$

where $\vec{r} = (r_1, \dots, r_{\chi(M)}) \in \mathbb{R}^{\chi(M)}$,

$$\|\vec{r}\|_{2n} = \left(\sum_{j=1}^{\chi(M)} r_j^{2n} \right)^{1/2n}$$

is the standard ℓ^{2n} -norm, and $P_M(\vec{r}) \subset M$ is the toric ball packing of M in which $\mathbb{B}^{2n}(r_j)$ is embedded at the j^{th} fixed point of M for $j = 1, \dots, \chi(M)$. Suppose that $\rho: \mathbb{B}^{2n}(r) \xrightarrow{\mathbb{T}^n} M$ is a symplectic \mathbb{T}^n -embedding into (M, ω, ϕ) for some $r > 0$. Then for any $\lambda \in \mathbb{R} \setminus \{0\}$ the map $\rho_\lambda: \mathbb{B}^{2n}(|\lambda|r) \xrightarrow{\mathbb{T}^n} M$ given by

$$\rho_\lambda(z) = \rho(z/|\lambda|)$$

is a symplectic \mathbb{T}^n -embedding into $(M, \lambda\omega, \phi)$. Thus if $P_M(\vec{r})$ is a toric packing of (M, ω, ϕ) then $P_M(|\lambda|r_1, \dots, |\lambda|r_{\chi(M)})$ is a toric ball packing of $(M, \lambda\omega, \phi)$ for any $\lambda \in \mathbb{R} \setminus \{0\}$. This and the fact that $\|\lambda r\|_{2n} = |\lambda| \|r\|_{2n}$ for all $r \in \mathbb{R}^{\chi(M)}$ and $\lambda \in \mathbb{R}$ imply that \mathcal{J} is conformal. Now suppose that $M, M' \in \text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ and $\rho: M \xrightarrow{\mathbb{T}^n} M'$. If $P \subset M$ is a toric ball packing of M then $\rho(P) \subset M'$ is a toric ball packing of M' of the same volume so $\mathcal{J}(M) \leq \mathcal{J}(M')$ and we see that \mathcal{J} is monotonic. Finally, suppose that there is a symplectic \mathbb{T}^n -embedding $M \xrightarrow{\mathbb{T}^n} Z^{2n}$. Then, since Z^{2n} has only one

point fixed by the \mathbb{T}^n -action and recalling that $c_{\mathbb{B}}^{n,n}(Z^{2n}) = 1$, it follows from Lemma 7.4.1 that

$$\mathcal{J}(M) \leq (1)^{1/2n} c_{\mathbb{B}}^{n,n}(Z^{2n}) = 1.$$

Finally, suppose that $\rho: B^{2n} \xrightarrow{\mathbb{T}^n} M$ is a symplectic \mathbb{T}^n -embedding. Then $P = \rho(B^{2n}) \subset M$ is a toric ball packing of M and thus

$$\mathcal{J}(M) \geq \left(\frac{\text{vol}(P)}{\text{vol}(B^{2n})} \right)^{1/2n} = 1.$$

Hence \mathcal{J} is tame. □

Example 7.4.3. Let $M \in \text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$. In [65] it is shown that there exists a \mathbb{Z} -valued function $\text{Emb}_M: \mathbb{R}_{\geq 0} \rightarrow [0, n! \chi(M)]$ such that the homotopy type of the space of symplectic \mathbb{T}^n -embeddings from $B^{2n}(r)$ into M is given by the disjoint union of $\text{Emb}_M(r)$ copies of \mathbb{T}^n . Thus, for each $r \in \mathbb{R}_{\geq 0}$ we may define a symplectic \mathbb{T}^n -capacity \mathcal{E}_r on $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ given by

$$\begin{aligned} \mathcal{E}_r: \text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n} &\rightarrow [0, \infty] \\ (M, \omega, \phi) &\mapsto (\text{vol}(M))^{\frac{1}{n}} \text{Emb}_M((\text{vol}(M))^{\frac{1}{n}} r). \end{aligned}$$

Since Emb_M is invariant up to \mathbb{T}^n -equivariant symplectomorphisms [65] and symplectic embeddings in $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ are automatically symplectomorphisms we see that \mathcal{E}_r is monotonic and it is an exercise to check that it is conformal. It is tame because the space of symplectic \mathbb{T}^n -embeddings of B^{2n} into Z^{2n} is homotopic to $n!$ disjoint copies of \mathbb{T}^n . ⊗

7.5 Symplectic $(S^1 \times \mathbb{R})$ -capacities

In this section we construct a symplectic $(S^1 \times \mathbb{R})$ -capacity on the space of semitoric integrable systems. Let $(M, \omega, F = (J, H))$ be a simple semitoric integrable system with m_f focus-focus singular points and let $\{\lambda_j\}_{j=1}^{m_f} \subset \mathbb{R}$ be the image under J of these points ordered so that $\lambda_1 < \lambda_2 < \dots < \lambda_{m_f}$. Let (λ_j, y_j) be the image under F of the j^{th} focus-focus singular point and for $\epsilon \in \{\pm 1\}$ let $\ell_{\lambda_j}^{\epsilon}$ be

those $(\lambda_j, y) \in \ell_{\lambda_j}$ such that $\epsilon y > \epsilon y_j$. Let $\ell^{\vec{\epsilon}} = \ell_{\lambda_1}^{\epsilon} \cup \dots \cup \ell_{\lambda_{m_f}}^{\epsilon}$. A homeomorphism

$$f: F(M) \rightarrow f(F(M)) \subset \mathbb{R}^2$$

is a *straightening map for M* [77] if for some choice of $\vec{\epsilon} \in \{\pm 1\}^{m_f}$ we have the following: $f|_{F(M) \setminus \ell^{\vec{\epsilon}}}$ is a diffeomorphism onto its image; $f|_{F(M) \setminus \ell^{\vec{\epsilon}}}$ is affine with respect to the affine structure $F(M)$ inherits from action-angle coordinates on M and the affine structure $f(F(M))$ inherits as a subset of \mathbb{R}^2 ; f preserves J , i.e. $f(x, y) = (x, f^{(2)}(x, y))$; $f|_{F(M) \setminus \ell^{\vec{\epsilon}}}$ extends to a smooth multi-valued map from $F(M)$ to \mathbb{R}^2 such that for any $c = (x_0, y_0) \in \ell^{\vec{\epsilon}}$ we have

$$\lim_{\substack{(x,y) \rightarrow c \\ x < x_0}} df(x, y) = T \lim_{\substack{(x,y) \rightarrow c \\ x > x_0}} df(x, y);$$

and the image of f is a rational convex polygon. Recall that T is the matrix given in Equation (2.2).

We say f is associated to $\vec{\epsilon}$.

Let $\mathfrak{T} \subset \text{AGL}_2(\mathbb{Z})$ be the subgroup including powers of T composed with vertical translations. It was proved in [77] that a semitoric system (M, ω, F) has a straightening map $f: M \rightarrow \mathbb{R}^2$ associated to each $\vec{\epsilon} \in \{\pm 1\}^{m_f}$, unique up to left composition with an element of \mathfrak{T} . Define

$$\mathcal{F}_M = \{ f \circ F \mid f \text{ is a straightening map for } M \}. \quad (7.4)$$

If $V_a: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ denotes vertical translation by $a \in \mathbb{R}$, then

$$\{ \tilde{F}(M) \mid \tilde{F} \in \mathcal{F}_M \} = \{ V_a(\Delta) \subset \mathbb{R}^2 \mid \Delta \text{ is associated to } M \text{ and } a \in \mathbb{R} \}$$

where a polygon is associated to M if it is an element of the affine invariant of M . Up to vertical translations the set \mathcal{F}_M is the orbit of a single non-unique function under the action of $G_{m_f} \times \mathcal{G}$. If $\tilde{F} \in \mathcal{F}_M$ then there exists some $\vec{\epsilon} \in \{-1, +1\}^{m_f}$ such that $\tilde{F}|_{M^{\vec{\epsilon}}}: M^{\vec{\epsilon}} \rightarrow \mathbb{R}^2$ is a momentum map for a \mathbb{T}^2 -action $\phi_{\tilde{F}}: \mathbb{T}^2 \times M^{\vec{\epsilon}} \rightarrow M^{\vec{\epsilon}}$ where $M^{\vec{\epsilon}} = M \setminus F^{-1}(\ell^{\vec{\epsilon}})$.

Corollary 7.5.1. *The manifold $M^{\vec{\epsilon}}$ has on it a momentum map for a Hamiltonian \mathbb{T}^2 -action unique up to \mathcal{G} . Thus $M^{\vec{\epsilon}} \in \text{Symp}^{4, \mathbb{T}^2}$ and the given \mathbb{T}^2 -action is unique up to composing the associated*

momentum map with an element of \mathcal{G} .

We call such actions of \mathbb{T}^2 on $M^{\bar{\epsilon}}$ *induced actions of \mathbb{T}^2* . Given any $\rho: N \rightarrow M$ with $\rho(N) \subset M^{\bar{\epsilon}}$ define $\rho_{\bar{\epsilon}}: N \rightarrow M^{\bar{\epsilon}}$ by $\rho_{\bar{\epsilon}}(p) = \rho(p)$ for $p \in N$.

Definition 7.5.2. Let (M, ω, F) be a semitoric integrable system and let $(N, \omega_N, \phi) \in \text{Symp}^{4, \mathbb{T}^2}$. A symplectic embedding $\rho: N \hookrightarrow M$ is a *semitoric embedding* if there exists $\bar{\epsilon} \in \{\pm 1\}^{m_f}$ and an induced action $\phi_{\bar{\epsilon}}: \mathbb{T}^2 \times M_{\bar{\epsilon}} \rightarrow M_{\bar{\epsilon}}$ such that $\rho(N) \subset M^{\bar{\epsilon}}$ and $\rho_{\bar{\epsilon}}: (N, \omega_N, \phi) \xrightarrow{\mathbb{T}^2} (M^{\bar{\epsilon}}, \omega, \phi_{\bar{\epsilon}})$ is a symplectic \mathbb{T}^2 -embedding.

Let (M, ω, F) be a semitoric manifold. A *semitoric ball packing* of M is a disjoint union $P = \bigsqcup_{\alpha \in \mathcal{A}} B_\alpha$ where $B_\alpha \subset M$ is a semitorically embedded ball in M . The *semitoric packing capacity* $\mathcal{S}\mathcal{T}: \text{Symp}_{\mathbb{S}\mathbb{T}}^{4, S^1 \times \mathbb{R}} \rightarrow [0, \infty]$ is given by

$$\mathcal{S}\mathcal{T}(M) = \left(\frac{\sup\{\text{vol}(P) \mid P \subset M \text{ is a semitoric ball packing of } M\}}{\text{vol}(\mathbb{B}^4)} \right)^{\frac{1}{4}}.$$

In order to show that $\mathcal{S}\mathcal{T}$ is a $(S^1 \times \mathbb{R})$ -capacity we need the following lemmas.

Lemma 7.5.3. For $i = 1, 2$ let (M_i, ω_i) be a symplectic integrable system, let $f_i: M_i \rightarrow \mathbb{R}$ be a function, and let \mathcal{X}_{f_i} denote the Hamiltonian vector field of f_i on M_i . If $\rho: M_1 \rightarrow M_2$ is a symplectomorphism such that $\rho_* \mathcal{X}_{f_1} = \mathcal{X}_{f_2}$ then $f_1 - \rho^* f_2: M_1 \rightarrow \mathbb{R}$ is constant.

Proof. Notice that

$$\begin{aligned} d(\rho^* f_2) &= \rho^*(df_2) = \rho^*(\iota_{\mathcal{X}_{f_2}} \omega_2) = \rho^*(\iota_{\rho_* \mathcal{X}_{f_1}} \omega_2) \\ &= \omega_2(\rho_* \mathcal{X}_{f_1}, \rho_*(\cdot)) = (\rho^* \omega_2)(\mathcal{X}_{f_1}, \cdot) = \iota_{\mathcal{X}_{f_1}} \omega_1 = df_1, \end{aligned}$$

thus f_1 and $\rho^* f_2$ differ by a constant. □

Lemma 7.5.4. Let $(M_i, \omega_i, F_i = (J_i, H_i))$ be semitoric integrable systems for $i = 1, 2$. If

$$\rho: M_1 \xrightarrow{S^1 \times \mathbb{R}} M_2$$

is a symplectic $(S^1 \times \mathbb{R})$ -embedding with respect to the Hamiltonian flow action on each system, then

$$\rho^* J_2 = eJ_1 + c_J \quad \text{and} \quad \rho^* H_2 = aJ_1 + bH_1 + c_H$$

for some $e \in \{\pm 1\}$ and $a, b, c_J, c_H \in \mathbb{R}$ such that $b \neq 0$.

Proof. Since ρ is $S^1 \times \mathbb{R}$ -equivariant there exists $\Lambda \in \text{Aut}(S^1 \times \mathbb{R})$ such that $\rho(\phi(g, m_1)) = \phi(\Lambda(g), \rho(m_1))$ for all $g \in S^1 \times \mathbb{R}$ and $m_1 \in M_1$. Associate $S^1 \times \mathbb{R}$ with $\mathbb{R}/\mathbb{Z} \times \mathbb{R}$ and give it coordinates $(x, y) \in \mathbb{R}^2$. Then $\Lambda \in \text{Aut}(S^1 \times \mathbb{R})$ and Λ continuous means that Λ descends from a linear invertible map from \mathbb{R}^2 to itself, which we will also denote $\Lambda \in \text{GL}_2(\mathbb{R})$. Write $\Lambda = (\Lambda_{ij})$ for $\Lambda_{ij} \in \mathbb{R}$ and $i, j \in \{1, 2\}$. The automorphism Λ sends the identity to itself so $\Lambda \begin{pmatrix} n \\ 0 \end{pmatrix} \in \mathbb{Z} \times \{0\}$ for all choices of $n \in \mathbb{Z}$. This implies that $\Lambda_{11} \in \mathbb{Z}$ and $\Lambda_{21} = 0$. Since Λ is invertible and $\Lambda^{-1} \in \text{Aut}(S^1 \times \mathbb{R})$ we see that $(\Lambda_{11})^{-1} \in \mathbb{Z}$ and so $\Lambda_{11} = \pm 1$. Since Λ is invertible and upper triangular we know that $\Lambda_{22} \neq 0$.

For a function $f: M_i \rightarrow \mathbb{R}$ let \mathcal{X}_f denote the associated Hamiltonian vector field on M_i , $i = 1, 2$. Also, for $v \in \mathfrak{g} = \text{Lie}(S^1 \times \mathbb{R})$, thought of as the tangent space to the identity, let v_{M_i} denote the vector field on M_i generated by v by the group action. Endow \mathfrak{g} with the coordinates (α, β) so that the exponential map will send $(\alpha, \beta) \in \mathfrak{g}$ to $(\alpha, \beta) \in \mathbb{R}/\mathbb{Z} \times \mathbb{R}$. Now notice that $\mathcal{X}_{J_1} = (1, 0)_{M_1}$ and $\mathcal{X}_{H_1} = (0, 1)_{M_1}$.

For $m_i \in M_i$, $i = 1, 2$, such that $\rho(m_1) = m_2$ we have

$$\rho_* \mathcal{X}_{J_1}(m_2) = \frac{d}{dt} \Big|_{t=0} \left(\rho(\phi((t, 0), m_1)) \right) = \frac{d}{dt} \Big|_{t=0} \left(\phi(\Lambda[(t, 0)], m_2) \right) = (T\Lambda(1, 0))_{M_2}(m_2)$$

Notice that $T_{(1,0)} = (\Lambda_{11}, 0) \in \mathfrak{g}$. Then $\rho_* \mathcal{X}_{J_1} = (T\Lambda(1, 0))_{M_2} = \Lambda_{11}(1, 0)_{M_2} = \Lambda_{11} \mathcal{X}_{J_2}$. Similarly we see that $\rho_* \mathcal{X}_{H_1} = \Lambda_{12} \mathcal{X}_{J_2} + \Lambda_{22} \mathcal{X}_{H_2}$. By Lemma 7.5.3 this implies that

$$\rho^* J_2 = \frac{1}{\Lambda_{11}} J_1 + c_J \quad \text{and} \quad \rho^* H_2 = \frac{-\Lambda_{12}}{\Lambda_{11} \Lambda_{22}} J_1 + \frac{1}{\Lambda_{22}} H_1 + c_H$$

for some $c_J, c_H \in \mathbb{R}$. Recalling that $\Lambda_{11} \in \{\pm 1\}$ and $\Lambda_{11}, \Lambda_{22} \neq 0$ take $e = (\Lambda_{11})^{-1}$, $a = \frac{-\Lambda_{12}}{\Lambda_{11} \Lambda_{22}}$, and $b = (\Lambda_{22})^{-1}$ to complete the proof. \square

Proposition 7.5.5. *The semitoric packing capacity, \mathcal{ST} , is a symplectic $(S^1 \times \mathbb{R})$ -capacity on $\text{Symp}_{\text{ST}}^{4, S^1 \times \mathbb{R}}$.*

Proof. The proof that \mathcal{ST} is conformal and non-trivial is analogous to the proof of Proposition 7.4.2, so we must only show that \mathcal{ST} is monotonic. Let (M_i, ω_i, F_i) be semitoric for $i = 1, 2$ and suppose $\phi: M_1 \xrightarrow{S^1 \times \mathbb{R}} M_2$ is a symplectic $(S^1 \times \mathbb{R})$ -embedding. Recall that action-angle coordinates are local Darboux charts in which the flow of the Hamiltonian vector fields are linear. Since ϕ is symplectic, $(S^1 \times \mathbb{R})$ -equivariant, and $\phi^*(F_2) = A \circ F_1$ where $A: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is affine (Lemma 7.5.4), this means that ϕ sends action-angle coordinates to action-angle coordinates. Since semitoric embeddings are those which respect the action-angle coordinates, given any semitoric embedding $\rho: B^{2n}(r) \hookrightarrow M_1$ the map $\phi \circ \rho: B^{2n}(r) \hookrightarrow M_2$ is a semitoric embedding. It follows that $\mathcal{ST}(M_1) \leq \mathcal{ST}(M_2)$. \square

Theorem 7.1.2 follows from Propositions 7.4.2 and 7.5.5.

Remark 7.5.6. There are many examples of classical symplectic capacities (see for instance [13]), and it would be of interest to adapt these capacities to the equivariant category. It would also be useful to construct symplectic G -capacities for more general integrable systems. In particular, integrable systems where a complete list of invariants is not known (that is, the vast majority).

In [34] the authors give a lower bound on the number of fixed points of a circle action on a compact almost complex manifold M with nonempty fixed point set, under the condition that the Chern number $c_1 c_{n-1}[M]$ vanishes. These results apply to a class of manifolds which do not support any Hamiltonian circle action with isolated fixed points, and which includes all symplectic Calabi-Yau manifolds [84] (see [34, Proposition 2.15]). The class of symplectic Calabi-Yau manifolds is thus of particular interest because they do not admit integrable systems of toric or semitoric type. Also, there is work extending the classification in [70] and related results to higher dimensions [80], so one could extend the semitoric packing capacity to higher dimensional semitoric systems, for which there is currently no classification.

Another interesting direction would be to generalize the work in [51] to our setting. There, the author constructs infinite dimensional symplectic capacities for a general class of Hamiltonian PDEs. In case the PDEs preserves some G -action, one may expect to construct also G -capacities in such infinite dimensional setting, and this may give new interesting result on the long time behavior of solutions.

Symplectic capacities are also of interest from a physical view point, for instance in [20] the authors describe interrelations between symplectic capacities and the uncertainty principle. It would be interesting to explore similar connections to symplectic G -capacities. \circlearrowright

Remark 7.5.7. In this chapter G can be a compact Lie group (like in the case of symplectic toric manifolds) or a non-compact Lie group (like in the case of semitoric systems). In general there are obstructions to the existence of effective G -actions on compact and non-compact manifolds, even in the case that the G -action is only required to be smooth. For instance, in [83, Corollary in page 242] it is proved that if N is an n -dimensional manifold on which a compact connected Lie group G acts effectively and there are $\sigma_1, \dots, \sigma_n \in \mathbf{H}^1(M, \mathbb{Q})$ such that $\sigma_1 \cup \dots \cup \sigma_n \neq 0$ then G is a torus and the G -action is locally free. In [83] Yau also proves several other results giving restrictions on G , M , and the fixed point set M^G . If the G -action is moreover assumed to be symplectic or Kähler, there are even more non-trivial constraints. Therefore the class of symplectic manifolds for which one can define a notion of symplectic G -capacity with G non-trivial is in general much more restrictive than the class of all symplectic manifolds. \circlearrowright

Acknowledgements. Chapter 7, in part, is comprised of material submitted for publication by the author of this dissertation, Alessio Figalli, and Álvaro Pelayo as *Symplectic G -capacities and integrable systems*, currently available as arXiv:1511.04499 [30]. Alessio Figalli is supported by NSF grants DMS-1262411 and DMS-1361122.

Chapter 8

Continuity of G -capacities

8.1 Introduction

In this Chapter I use the metric on toric integrable systems [67] and the metric on semitoric integrable systems from Chapter 3 to study the continuity of the symplectic invariants defined in Chapter 7.

The continuity of symplectic capacities is discussed in [6, 13, 26, 85]. The semitoric and toric packing capacities are each defined on categories of integrable systems which have a natural topology, but we can only discuss the continuity of the (m, k) -equivariant Gromov radius on a subcategory of its domain which has a topology, so I restrict to the case of $(m, k) = (n, n)$. The \mathbb{T}^n -action on a symplectic toric manifold may be lifted to an action of \mathbb{R}^n . Let $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{R}^n}$ be the symplectic category of symplectic toric manifolds each of which is endowed with the \mathbb{R}^n -action obtained by lifting the given \mathbb{T}^n -action which is a subcategory of $\text{Symp}^{2n, \mathbb{R}^n}$.

Theorem 8.1.1 (Continuity of capacities). *The following hold:*

- (i) *The toric packing capacity $\mathcal{J}: \text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n} \rightarrow [0, \infty]$ is everywhere discontinuous and the restriction of \mathcal{J} to the space of symplectic toric $2n$ -dimensional manifolds with exactly N fixed points of the \mathbb{T}^n -action is continuous for any choice of $N \geq 0$;*
- (ii) *The semitoric packing capacity $\mathcal{S}\mathcal{J}: \text{Symp}_{\mathbb{S}\mathbb{T}}^{4, \mathbb{S}^1 \times \mathbb{R}} \rightarrow [0, \infty]$ is everywhere discontinuous and the restriction of $\mathcal{S}\mathcal{J}$ to the space of semitoric manifolds with exactly N elliptic-elliptic fixed points*

of the associated $(S^1 \times \mathbb{R})$ -action is continuous for any choice of $N \geq 0$;

(iii) The (n, n) -equivariant Gromov radius restricted to the space of symplectic toric manifolds

$$c_B^{n,n}|_{\text{Symp}_T^{2n, \mathbb{R}^n}} : \text{Symp}_T^{2n, \mathbb{R}^n} \rightarrow [0, \infty]$$

is everywhere discontinuous and the restriction of $c_B^{n,n}|_{\text{Symp}_T^{2n, \mathbb{R}^n}}$ to the space of symplectic toric $2n$ -dimensional manifolds with exactly N fixed points of the \mathbb{R}^n -action is continuous for any choice of $N \geq 0$.

Theorem 8.1.1 generalizes [31, Theorem A], which deals with 4-manifolds, and solves [67, Problem 30].

8.2 Continuity of symplectic \mathbb{T}^n -capacities

In this section we study the continuity of the symplectic \mathbb{T}^n -capacity constructed in Section 7.4. The metric on $\mathcal{M}_T \cong \text{Ham}_T^{2n, \mathbb{T}^n} / \approx_T$ from [67] is described in Section 2.2.6.


The map

$$\text{Ham}_T^{2n, \mathbb{T}^n} / \approx_T \rightarrow \text{Symp}_T^{2n, \mathbb{T}^n} / \sim_T$$

given by $[(M, \omega, \phi, \mu)] \mapsto [(M, \omega, \phi)]$ is a quotient map and thus we can endow $\text{Symp}_T^{2n, \mathbb{T}^n} / \sim_T$ with the quotient topology. Since $\text{Symp}_T^{2n, \mathbb{T}^n} / \sim_T$ is a quotient of $\text{Symp}_T^{2n, \mathbb{T}^n}$ we can pull the topology up from $\text{Symp}_T^{2n, \mathbb{T}^n} / \sim_T$ to $\text{Symp}_T^{2n, \mathbb{T}^n}$ by declaring that a set in $\text{Symp}_T^{2n, \mathbb{T}^n}$ is open if and only if it is the preimage of an open set from $\text{Symp}_T^{2n, \mathbb{T}^n} / \sim_T$ under the natural projection. Two points in $\text{Symp}_T^{2n, \mathbb{T}^n}$ are not separable if and only if they are \mathbb{T}^n -equivariantly symplectomorphic. Thus, a map $c: \text{Symp}_T^{2n, \mathbb{T}^n} \rightarrow [0, \infty]$ which descends to a well-defined map ϕ on $\text{Symp}_T^{2n, \mathbb{T}^n} / \sim_T$ is continuous if and only if the map

$$\hat{c}: \text{Ham}_T^{2n, \mathbb{T}^n} / \approx_T \rightarrow [0, \infty]$$

is continuous, where \hat{c} is defined by the following commutative diagram:



Next we define an operation on Delzant polytopes. Let $n \in \mathbb{Z}_{>0}$. For $x_0 \in \mathbb{R}^n$, $w_1, \dots, w_n \in \mathbb{Z}^n$, and $\varepsilon > 0$ define

$$\mathcal{H}_{x_0}^\varepsilon(w_1, \dots, w_n) = \{ x_0 + \sum_j t_j w_j \mid t_1, \dots, t_n \in \mathbb{R}_{\geq 0}, \sum_j t_j \geq \varepsilon \}. \tag{8.1}$$

Recall that $\mathcal{P}_{\mathbb{T}}$ denotes the set of Delzant polytopes (Section 2.2.6). Suppose that $\Delta \in \mathcal{P}_{\mathbb{T}}$ and $x_0 \in \mathbb{R}^n$ is a vertex of Δ . Let $u_i \in \mathbb{Z}^n$, $i = 1, \dots, n$, denote the primitive vectors along which the edges adjacent to x_0 are aligned. The ε -corner chop of Δ at x_0 is the polygon $\Delta_{x_0}^\varepsilon \in \mathcal{P}_{\mathbb{T}}$ given by $\Delta_{x_0}^\varepsilon = \Delta \cap \mathcal{H}_{x_0}^\varepsilon(u_1, \dots, u_n)$ where ε is sufficiently small so that $\Delta_{x_0}^\varepsilon$ has exactly one more face than Δ does as is shown in Figure 8.1. One can check that if $\Delta \in \mathcal{P}_{\mathbb{T}}$ then $\Delta_{x_0}^\varepsilon \in \mathcal{P}_{\mathbb{T}}$. Notice that




Figure 8.1: An ε -corner chop at a vertex x_0 of Δ for some $\varepsilon > 0$.

$\lim_{\varepsilon \rightarrow 0} d_{\mathcal{P}}(\Delta, \Delta_{x_0}^\varepsilon) = 0$. This means that given any element of $\mathcal{P}_{\mathbb{T}}$ with N vertices, corner chopping can be used to produce other polygons which are close in $d_{\mathcal{P}}$ and all polygons produced in this way will have more than N vertices. Let $\mathcal{P}_{\mathbb{T}}^N$ denote the set of Delzant polygons in \mathbb{R}^n with exactly N vertices. We will later need the following.

Proposition 8.2.1 ([31]). *Let $N \in \mathbb{Z}_{>0}$ and $\Delta \in \mathcal{P}_{\mathbb{T}}^N$. Any sufficiently small neighborhood of Δ is a subset of $\cup_{(N' \geq N)} \mathcal{P}_{\mathbb{T}}^{N'}$.*

We study ball packing problems about symplectic toric manifolds by instead studying packings of the associated Delzant polygon. Let $\Delta \in \mathcal{P}_{\mathbb{T}}$ be a Delzant polytope. Let $\text{AGL}_n(\mathbb{Z}) = \text{GL}_n(\mathbb{Z}) \ltimes \mathbb{R}^n$




Figure 8.2: (a) An image of $\Delta(1) \subset \mathbb{R}^2$. (b) An image of an admissible, but not maximal, packing.

denote the group of affine transformations in \mathbb{R}^n with linear part in $GL_n(\mathbb{Z})$. For $r > 0$ let $\Delta(r) = \text{Conv}\{re_1, \dots, re_n, 0\} \setminus \text{Conv}\{re_1, \dots, re_n\}$ where $\text{Conv}(E)$ denotes the convex hull of the set $E \subset \mathbb{R}^n$ and $\{e_1, \dots, e_n\}$ denote the standard basis vectors in \mathbb{R}^n . Following [64], a subset Σ of Δ is an *admissible simplex of radius $r > 0$ with center at a vertex x_0 of Δ* if there exists some $A \in \text{AGL}_n(\mathbb{Z})$ such that:

- (1) $A(\Delta(r^{1/2})) = \Sigma$;
- (2) $A(0) = x_0$;
- (3) A takes the edges of $\Delta(r^{1/2})$ meeting at the origin to the edges of Δ meeting at x_0 .

An *admissible packing* of Δ is a disjoint union $R = \bigsqcup_{\alpha \in \mathcal{A}} \Sigma_\alpha \subset \Delta$ where each Σ_α is an admissible simplex for Δ . This is illustrated in Figure 8.2. The half-plane $\mathcal{H}_{x_0}^\varepsilon$ given in Equation (8.1) is designed so that that an ε -corner chop on a Delzant polytope corresponds to the removal of an admissible simplex of radius ε .

The function $\Omega: \text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n} / \sim_{\mathbb{T}} \rightarrow (0, 1]$ given by

$$\Omega(M) = \frac{\sup\{\text{vol}(P) \mid P \text{ is a toric ball packing of } M\}}{\text{vol}(M)},$$

known as the *optimal toric density function*, has been studied in [31, 64, 68]. In particular, in [31] Pelayo-Figalli studied the regions of continuity of Ω and proved the $n = 2$ case of Theorem 8.1.1 part (i). They stated the theorem in terms of Ω , while we state it in terms of \mathcal{J} .

Let $\text{vol}: \text{Symp}_{\mathbb{T}}^{2n, \mathbb{T}^n} \rightarrow \mathbb{R}$ denote the total symplectic volume of a symplectic toric manifold and let $\text{vol}_{\mathcal{P}}: \mathcal{P}_{\mathbb{T}} \rightarrow \mathbb{R}$ denote Euclidean volume function of a polytope in \mathbb{R}^n . Let

$$(B^{2n}(r), \omega_B, \phi_B, \mu_B) \in \text{Ham}_{\mathbb{T}}^{2n, \mathbb{T}^n}$$

denote the standard ball of radius $r > 0$ in \mathbb{C}^n with the standard action of \mathbb{T}^n and suppose

that $(M, \omega, \phi, \mu) \in \text{Ham}_{\mathbb{T}}^{2n, \mathbb{T}^n}$. Let $\Delta_B = \mu_B(B^{2n}(r))$ and $\Delta = \mu(M)$. Then, as shown in [64], $\text{vol}(M) = n! \pi^n \text{vol}_{\mathcal{P}}(\Delta)$ and if $f: B^{2n}(r) \xrightarrow{\mathbb{T}^n} M$ is a symplectic \mathbb{T}^n -embedding then

$$\text{vol}(B^{2n}(r)) = \text{vol}(f(B^{2n}(r))) = n! \pi^n \text{vol}_{\mathcal{P}}(\mu \circ f(B^{2n}(r))) = n! \pi^n \text{vol}_{\mathcal{P}}(\Delta_B).$$

Theorem 8.2.2 ([64]). *Let $(M, \omega, \phi, \mu) \in \text{Ham}_{\mathbb{T}}^{2n, \mathbb{T}^n}$ and let $\Delta = \mu(M)$. Suppose $\phi: B^{2n}(r) \hookrightarrow M$ is a symplectic \mathbb{T}^n -embedding for some $r > 0$. Then $\mu(\phi(B^{2n}(r))) \subset \Delta$ is an admissible simplex of radius r^2 . Conversely, if $\Sigma \subset \Delta$ is an admissible simplex of radius r^2 then there exists a symplectic \mathbb{T}^n -embedding $\phi: B^{2n}(r) \hookrightarrow M$ such that $\mu(\phi(B^{2n}(r))) = \Sigma$.*

Moreover, if P is a toric ball packing of M , then $\mu(P) \subset \Delta$ is an admissible packing of Δ . Conversely, if R is an admissible packing of Δ then there exists a toric ball packing P of M such that $\mu(P) = R$.

Since there is a toric ball packing P of M related to an admissible packing R of Δ by $\mu(P) = R$, it follows that $\text{vol}(P) = n! \pi^n \text{vol}_{\mathcal{P}}(R)$. To study packing of the manifold we will study packing of the polygon. Thus, we define $\pi_{\mathbb{T}}: \mathcal{P}_{\mathbb{T}} \rightarrow (0, \infty)$ by

$$\pi_{\mathbb{T}}(\Delta) = \sup\{\text{vol}_{\mathcal{P}}(R) \mid R \text{ is an admissible packing of } \Delta\}.$$

Suppose that $\Delta \in \mathcal{P}_{\mathbb{T}}^N$ with vertices $v_1, \dots, v_N \in \mathbb{R}^n$ and let $\pi_{\mathbb{T}}^i(\Delta)$ be the supremum of $\text{vol}_{\mathcal{P}}(R)$ over all admissible packings \mathcal{R} of Δ in which $v_i \notin \mathcal{R}$.

The following result generalizes [31, Theorem 7.1] to the case $n \geq 3$.

Theorem 8.2.3. *Fix $n \in \mathbb{Z}_{>0}$. For $N \in \mathbb{Z}_{\geq 1}$ and let $\mathcal{P}_{\mathbb{T}}^N$ denote the set of Delzant polygons in \mathbb{R}^n with exactly N vertices. Then:*

1. $\pi_{\mathbb{T}}$ is discontinuous at each point in $\mathcal{P}_{\mathbb{T}}$;
2. the restriction $\pi_{\mathbb{T}}|_{\mathcal{P}_{\mathbb{T}}^N}$ is continuous for each $N \geq 1$;
3. if $\Delta \in \mathcal{P}_{\mathbb{T}}^N$ then $\mathcal{P}_{\mathbb{T}}^N$ is the largest neighborhood of Δ in $\mathcal{P}_{\mathbb{T}}$ in which $\pi_{\mathbb{T}}$ is continuous if and only if $\pi_{\mathbb{T}}^i(\Delta) < \pi_{\mathbb{T}}(\Delta)$ for all $1 \leq i \leq N$.

Proof. First we show (1). Let $\Delta \in \mathcal{P}_{\mathbb{T}}^N$ and for any small enough $\varepsilon > 0$ perform an ε -corner chop (as in Section 8.2) at each corner to produce $\Delta_{\varepsilon} \in \mathcal{P}_{\mathbb{T}}^{2N}$. Any admissible packing of Δ_{ε} can have at most

$2N$ simplices and each simplex must have one side with length at most ε while the other sides are universally bounded by the maximal side length of Δ . The size of such simplices decreases to zero as ε does, so $\lim_{\varepsilon \rightarrow 0} \pi_T(\Delta_\varepsilon) = 0$. Hence

$$\lim_{\varepsilon \rightarrow 0} d_{\mathcal{P}}(\Delta, \Delta_\varepsilon) = 0$$

but

$$\lim_{\varepsilon \rightarrow 0} |\pi_T(\Delta) - \pi_T(\Delta_\varepsilon)| = \pi_T(\Delta) > 0,$$

so π_T is discontinuous at Δ .

Now we prepare to show part (2). For any $v_1, \dots, v_n \in \mathbb{Z}^n$ let $[v_1, \dots, v_n]$ denote the $n \times n$ integer matrix with i^{th} column given by v_i for $i = 1, \dots, n$. Let $\eta: \text{SL}_n(\mathbb{Z}) \rightarrow \text{GL}_n(\mathbb{R})$ given by

$$\eta([v_1, \dots, v_n]) = \left[\frac{v_1}{|v_1|}, \dots, \frac{v_n}{|v_n|} \right]$$

take a nonsingular integer matrix to its *column normalization*. Notice for any $A = [v_1, \dots, v_n] \in \text{SL}_n(\mathbb{Z})$ that

$$\det(A) = |v_1| \cdots |v_n| \cdot \det(\eta(A)).$$

Suppose $\Delta \in \mathcal{P}_T$ is n -dimensional. In a neighborhood around each vertex the polytope is described by a collection of vectors $v_1, \dots, v_n \in \mathbb{Z}^n$ with $\det(v_1, \dots, v_n) = 1$ along which the edges adjacent to this vertex are directed. So, associated to any vertex of a Delzant polytope, there is a matrix $A \in \text{SL}_n(\mathbb{Z})$ given by $A = [v_1, \dots, v_n]$ which is unique up to even permutations of its columns and thus, though A is not unique, the values determined by $\det(A)$ and $\det(\eta(A))$ associated to a vertex are well-defined. Fix $\Delta \in \mathcal{P}_T^N$ and $\{\Delta_j\}_{j=1}^\infty \subset \mathcal{P}_T^N$ such that

$$\lim_{j \rightarrow \infty} d_{\mathcal{P}}(\Delta, \Delta_j) = 0. \tag{8.2}$$

For j large enough for each vertex V of Δ there must be a corresponding vertex V_j of Δ_j so that $V_j \rightarrow V$ as $j \rightarrow \infty$. Let $A \in \text{SL}_n(\mathbb{Z})$ be a matrix corresponding to V and let $A_j \in \text{SL}_n(\mathbb{Z})$ be a matrix corresponding to V_j for $j \in \mathbb{Z}$ large enough. In particular, convergence in $d_{\mathcal{P}}$, which is convergence in

$L^1(\mathbb{R}^n)$, implies that locally these vertices must converge, so Equation (8.2) implies that

$$\lim_{j \rightarrow \infty} |\det(\eta(A)) - \det(\eta(A_j))| = 0.$$

Now we are ready to prove (2) by showing that the collection of possible vertices of Delzant polytopes is discrete. Fix $\Delta \in \mathcal{P}_T^N$ with a vertex V at the origin and let $\varepsilon > 0$. Choose $\delta > 0$ small enough so that if $\Delta' \in \mathcal{P}_T^N$ with a vertex V' at the origin then $d_{\mathcal{P}}(\Delta, \Delta') < \delta$ implies that

$$|\det(\eta(A)) - \det(\eta(A'))| < \varepsilon, \tag{8.3}$$

where $A \in \mathrm{SL}_n(\mathbb{Z})$ is a matrix associated to V and $A' \in \mathrm{SL}_n(\mathbb{Z})$ is a matrix associated to V' . Suppose that $\varepsilon < \det(\eta(A))$. Now let $A' = [w_1, \dots, w_n]$ for $w_i \in \mathbb{Z}^n$, $i = 1, \dots, n$. These are all nonzero integer vectors so $|w_i| \geq 1$ for $i = 1, \dots, n$. For each i we have

$$1 = \det(A') = |w_1| |w_2| \dots |w_n| \det(\eta(A')) \geq |w_i| \det(\eta(A'))$$

and so by Equation (8.3)

$$|w_i| \leq \frac{1}{\det(\eta(A'))} \leq \frac{1}{\det(\eta(A)) - \varepsilon}.$$

Thus each $w_i \in \mathbb{Z}^n$ has length at most $(\det(\eta(A)) - \varepsilon)^{-1}$, a value which does not depend on Δ' , and so to be within δ of Δ the vectors directing the edges coming out from the vertex V' of Δ' must be chosen from only finitely many options. This means the set of possible local neighborhoods of vertices is discrete. Thus, for small enough $\delta > 0$ we conclude that $d_{\mathcal{P}}(\Delta, \Delta') < \delta$ implies that there exist open sets $U, U' \subset \mathbb{R}^n$ around the vertices V and V' such that

$$\Delta \cap U = F_c(\Delta' \cap U')$$

where $F_c: \mathbb{R} \rightarrow \mathbb{R}$ is a translation by some fixed $c \in \mathbb{R}^n$. Now, let $\Delta \in \mathcal{P}_T^N$ be any Delzant polytope in \mathbb{R}^n with N vertices. In a sufficiently small $d_{\mathcal{P}}$ -neighborhood of Δ all polytopes must have the same angles at the finitely many vertices by the argument above. Thus they are all related to Δ by translating its faces in a parallel way, which continuously changes π_T . This proves (2) because π_T is

continuous on such families.

Finally we show (3). Let $\Delta \in \mathcal{P}_T^N$ and assume that $\pi_T(\Delta) = \pi_T^i(\Delta)$ for some $i \in \{1, \dots, N\}$. Then there is an optimal packing of Δ which avoids the i^{th} vertex. For $\varepsilon > 0$ let $\Delta_\varepsilon \in \mathcal{P}_T^{N+1}$ be the ε -corner chop of Δ at the i^{th} vertex. Since the optimal packing of Δ avoids the i^{th} vertex, we see that $\lim_{\varepsilon \rightarrow 0} d_{\mathcal{P}}(\Delta, \Delta_\varepsilon) = 0$ and $\lim_{\varepsilon \rightarrow 0} \pi_T(\Delta_\varepsilon) = \pi_T(\Delta)$ so there is a set larger than \mathcal{P}_T^N on which π_T is continuous around Δ .

Conversely assume that $\Delta \in \mathcal{P}_T^N$ satisfies $\pi_T^i(\Delta) < \pi_T(\Delta)$ for all $i = 1, \dots, n$. By Proposition 8.2.1 we know that any small enough neighborhood of Δ only includes polytopes with N vertices and polytopes with more than N vertices, which are produced from corner chops of Δ . We must now only show that π_T cannot be continuous on any neighborhood of Δ which includes any such polygons. For $\varepsilon > 0$ let $\Delta_\varepsilon \in \mathcal{P}_T^{N+1}$ be the ε -corner chop of Δ at the i^{th} vertex. Then $\lim_{\varepsilon \rightarrow 0} \pi_T(\Delta_\varepsilon) = \pi_T^i(\Delta) < \pi_T(\Delta)$ so for small enough corner chops $\pi_T(\Delta_\varepsilon)$ is bounded away from $\pi_T(\Delta)$. Thus any set on which π_T is continuous around Δ cannot include any corner chops of Δ . From this we conclude that any such set cannot include polytopes with greater than N vertices. The result follows since π_T is continuous on all of \mathcal{P}_T^N . □

Theorem 8.1.1 part (i) follows from Theorem 8.2.2 and Theorem 8.2.3. In addition, these Theorems also imply the following result. Let $N \geq 1$ and let $\text{Symp}_{T,N}^{2n, \mathbb{T}^n}$ denote the set of symplectic toric manifolds with exactly N points fixed by the \mathbb{T}^n -action. For $(M, \omega, \phi) \in \text{Symp}_{T,N}^{2n, \mathbb{T}^n}$ with fixed points $p_1, \dots, p_N \in M$ let

$$\mathcal{J}^i(M) = \left(\frac{\sup\{\text{vol}(P) \mid P \text{ is a toric ball packing of } M \text{ such that } p_i \notin P\}}{\text{vol}(\mathbb{B}^{2n})} \right)^{\frac{1}{2n}}.$$

Proposition 8.2.4. *The space $\text{Symp}_{T,N}^{2n, \mathbb{T}^n}$ is the largest neighborhood of M in $\text{Symp}_T^{2n, \mathbb{T}^n}$ in which \mathcal{J} is continuous if and only if $\mathcal{J}^i(M) < \mathcal{J}(M)$ for every $1 \leq i \leq N$.*

Theorem 8.1.1 part (i) and Proposition 8.2.4 are illustrated in Figure 8.3. If $n = 2$ Proposition 8.2.4 was proved in [31].




Figure 8.3: Continuous families of Delzant polygons on which (a) \mathcal{T} is continuous and (b) \mathcal{T} is not continuous.


8.3 Continuity of symplectic $(S^1 \times \mathbb{R})$ -capacities

In this section we study the continuity of the symplectic $(S^1 \times \mathbb{R})$ -capacity constructed in Section 7.5 relative to the metric defined in Chapter 3. We are only interested in the topology of $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \sim_{\text{ST}}$ so, as is suggested in Remark 3.2.13, we will use the simplified version of the metric

$$\mathcal{D}^{\text{Id}} = \Phi^* d_{m_f, \vec{k}}^{\text{Id}, \nu, \{b_n\}_{n=0}^\infty}.$$

Recall that while \mathcal{D}^{Id} produces a different metric space structure on $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \sim_{\text{ST}}$ it induces the same topology on $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \sim_{\text{ST}}$ as the full metric (Proposition 3.3.18). Fix an admissible measure ν and let $d_{\text{P}}^{\text{ST}} = d_{\text{P}}^{\text{Id}, \nu}$ denote the metric on semitoric polygons from Definition 3.2.7 relative to the identity permutation.

Since $\text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \approx_{\text{ST}}$ is a quotient of $\text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}}$ we can pull the topology up from $\text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \approx_{\text{ST}}$ to $\text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}}$ by declaring that a set in $\text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}}$ is open if and only if it is the preimage of an open set from $\text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \approx_{\text{ST}}$ under the natural projection. We endow $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}}$ with the quotient topology relative to the map $\text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}} \rightarrow \text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}}$ which forgets the momentum map. Thus a map $c: \text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}} \rightarrow [0, \infty]$ which descends to a well-defined map ϕ on $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \sim_{\text{ST}}$ is continuous if and only if the map $\hat{c}: \text{Ham}_{\text{ST}}^{4,S^1 \times \mathbb{R}} / \approx_{\text{ST}} \rightarrow [0, \infty]$ is continuous where \hat{c} is defined by the commutative diagram:



Let $\Delta_w = (\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})$ be a primitive semitoric polygon, and let $v \in \Delta$ be a vertex.

Definition 8.3.1. An admissible semitoric simplex of radius $r > 0$ with center at v is a subset Σ of Δ such that there exist some $A \in \text{AGL}_2(\mathbb{Z})$ and $\vec{u} \in \{0, 1\}^{m_f}$ satisfying:

- $A(\Delta(r^{1/2})) = t_{\vec{\lambda}}^{\vec{u}}(\Sigma)$;
- $A(0) = t_{\vec{\lambda}}^{\vec{u}}(v)$;
- A takes the edges of $\Delta(r^{1/2})$ meeting at the origin to the edges of $t_{\vec{\lambda}}^{\vec{u}}(\Delta)$ meeting at $t_{\vec{\lambda}}^{\vec{u}}(v)$;
- $\Sigma \subset \Delta^{\vec{u}}$ where

$$\Delta^{\vec{u}} = \Delta \setminus \left\{ (x, y) \in \Delta \mid \begin{array}{l} x = \lambda_j \text{ and } (-2\vec{u} + 1)y \geq \min_{(\lambda_j, y_0)} y_0 + h_j \\ \text{for some } j \in \{1, \dots, m_f\} \end{array} \right\}.$$

An admissible semitoric packing of Δ_w is a disjoint union $R = \bigsqcup_{\alpha \in \mathcal{A}} \Sigma_\alpha$ where each Σ_α is an admissible simplex of some radius, where the radii of the simplices are allowed to be different.

Such a simplex cannot exist at a fake corner.




Figure 8.4: An admissible semitoric packing. Here t denotes $t_{\vec{\lambda}}^{\vec{u}}$.

Lemma 8.3.2 ([65]). Let F^B be a momentum map for the usual \mathbb{T}^n -action on $B^{2n}(r)$, $r > 0$, and let (M, ω, ϕ, F) be a Hamiltonian \mathbb{T}^n -manifold of dimension $2n$. If $\rho: B^{2n}(r) \hookrightarrow M$ is a symplectic \mathbb{T}^n -embedding with respect to some $\Lambda \in \text{Aut}(\mathbb{T}^n)$ then there exists some $x \in \mathbb{R}^n$ such that the following

diagram commutes:

$$\begin{array}{ccc} \mathbb{B}^{2n}(r) & \xrightarrow{\rho} & M \\ F^B \downarrow & & \downarrow F \\ \mathbb{R}^2 & \xrightarrow{(\Lambda^t)^{-1}+x} & \mathbb{R}^2 \end{array}$$

where $(\Lambda^t)^{-1} + x$ is the affine map with linear part $(\Lambda^t)^{-1}$ which takes 0 to x .

In [50] a *proper Hamiltonian \mathbb{T}^n -manifold* is a quadruple $(Q, \omega^Q, F^Q, \Gamma)$ where (Q, ω^Q) is a connected $2n$ -dimensional symplectic manifold with momentum map F^Q for an action of \mathbb{T}^n and $\Gamma \subset \text{Lie}(\mathbb{T}^n)^*$ is an open convex subset with $F^Q(Q) \subset \Gamma$ and such that F^Q is proper as a map to Γ . A proper Hamiltonian \mathbb{T}^n -manifold is *centered about* $p \in \Gamma$ if p is an element of each component of $F^Q(Q^K)$ for each subgroup $K \subset \mathbb{T}^n$, where Q^K is the set of all points in Q which are fixed by the action of all elements of K .

Lemma 8.3.3 ([50]). *Let $(Q, \omega^Q, F^Q, \Gamma)$ be a proper Hamiltonian \mathbb{T}^n -manifold of dimension $2n$. If $(Q, \omega^Q, F^Q, \Gamma)$ is centered about $p \in \Gamma$ and $(F^Q)^{-1}(\{p\}) = \{q\}$, then Q is equivariantly symplectomorphic to $\{z \in \mathbb{C}^n \mid p + \sum_{j=1}^n |z_j|^2 \eta_j^q \in \Gamma\}$, where $\eta_1^q, \dots, \eta_n^q \in \text{Lie}(\mathbb{T}^n)^*$ are the weights of the isotropy representation of \mathbb{T}^n on $T_q Q$.*

We use Lemma 8.3.2 and Lemma 8.3.3 to prove the following.

Proposition 8.3.4. *Let $(M, \omega, F = (J, H))$ be a semitoric manifold such that*

$$\Phi((M, \omega, F)) = (m_f, ((S_j)^\infty)_{j=1}^{m_f}, [\Delta_w], (h_j)_{j=1}^\infty)$$

where $\Delta_w = (\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})$ is primitive with associated momentum map $\tilde{F} \in \mathcal{F}_M$ such that $\tilde{F}(M) = \Delta$. Then:

1. Suppose $\rho: \mathbb{B}^4(r) \hookrightarrow M$ is a semitoric embedding for some $r > 0$. Then $\tilde{F}(\rho(\mathbb{B}^4(r))) \subset \Delta$ is an admissible semitoric simplex with radius r^2 . Conversely, if $\Sigma \subset \Delta$ is an admissible semitoric simplex with radius r^2 then there exists a semitoric embedding $\rho: \mathbb{B}^4(r) \hookrightarrow M$ such that $\tilde{F}(\rho(\mathbb{B}^4(r))) = \Sigma$.
2. Let P be a semitoric ball packing of M . Then $\tilde{F}(P) \subset \Delta$ is an admissible packing of Δ_w . Conversely, if R is an admissible packing of Δ_w then there exists a semitoric ball packing P of M such that $\tilde{F}(P) = R$.

Proof. Part (2) follows from Part (1) since the semitoric simplices associated to disjoint semitorically embedded balls are disjoint. This follows from the fact that $\tilde{F}^{-1}(p)$ is a 2-dimensional submanifold of M for any regular point $p \in \Delta$ and the embedded balls are 2-dimensional.

Suppose that $B \subset M$ is a semitorically embedded ball of radius $r > 0$. Then for some $\vec{\epsilon} \in \{-1, +1\}^{m_f}$ the map $\rho_{\vec{\epsilon}}: B^4(r) \hookrightarrow M^{\vec{\epsilon}}$ is a \mathbb{T}^2 -embedding with respect to some $\Lambda \in \text{Aut}(\mathbb{T}^2)$. Recall $M^{\vec{\epsilon}}$ is a Hamiltonian \mathbb{T}^2 -manifold and denote a momentum map for this action by $F^{\vec{\epsilon}}$. Let $p = F^{\vec{\epsilon}}(\rho(0))$ and let $\Delta^{\vec{\epsilon}} = F^{\vec{\epsilon}}(\Delta)$. Hence by Lemma 8.3.2 the diagram

$$\begin{array}{ccc} B^4(r) & \xrightarrow{\rho} & M^{\vec{\epsilon}} \\ F_B \downarrow & & \downarrow F^{\vec{\epsilon}} \\ \Delta_B & \xrightarrow{(\Lambda^t)^{-1} + x} & \Delta^{\vec{\epsilon}} \end{array}$$

commutes for some $x \in \text{Lie}(\mathbb{T}^2)^*$. Since Λ is an automorphism so is $(\Lambda^t)^{-1}$, hence it sends the weights of the isotropy representation of \mathbb{T}^2 on $T_0(B^4(r))$ to the weights of the isotropy representation on $T_p M$. Since $(\Lambda^t)^{-1}$ is linear and Δ_B is the convex hull of the isotropy weights of the representation on $T_0(B^4(r))$ and the origin, we find that

$$\Sigma^{\vec{\epsilon}} := [(\Lambda^t)^{-1} + x](\Delta_B)$$

is the convex hull of p , $p + r^2\alpha_1$, and $p + r^2\alpha_2$, minus the convex hull of $p + r^2\alpha_1$ and $p + r^2\alpha_2$, where α_1 and α_2 are the weights of the isotropy representation of \mathbb{T}^2 on $T_p M$. For $\vec{u} = \frac{1}{2}(1 - \vec{\epsilon})$ recall that $t_{\vec{\lambda}}^{\vec{u}}(\Delta) = \Delta^{\vec{\epsilon}}$ and let $\Sigma = (t_{\vec{\lambda}}^{\vec{u}})^{-1}(\Sigma^{\vec{\epsilon}})$. Notice that $\Sigma = \tilde{F}(\rho(B^4(r))) \subset \Delta$ and is an admissible semitoric simplex.

To prove the converse let $\Sigma \subset \Delta$ be an admissible semitoric simplex. This means that there exists some $\vec{\epsilon} \in \{-1, +1\}^{m_f}$ such that

$$\Sigma' := t_{\vec{\lambda}}^{\vec{u}}(\Sigma)$$

satisfies the requirements of Definition 8.3.1, where $\vec{u} = \frac{1}{2}(1 - \vec{\epsilon})$. Let $\Delta' = t_{\vec{\lambda}}^{\vec{u}}(\Delta)$. Let p be the unique vertex of Σ' . Thus, Σ' is the convex hull of p , $p + r^2\alpha_1$, and $p + r^2\alpha_2$, minus the convex hull of $p + r^2\alpha_1$ and $p + r^2\alpha_2$, for some $\alpha_i \in \mathbb{R}^2$, $i = 1, 2$. Let $\Gamma \subset \mathbb{R}^2$ be the unique open half plane satisfying $\Gamma \cup \Delta' = \Sigma'$. Let $N = \tilde{F}^{-1}(\Sigma)$ and let $\omega^N = \omega|_N$. We can see that $N \subset M$ is open and by

the proof of the Atiyah-Guillemin-Sternberg Convexity Theorem [3, 40] we know that N is connected. The map \tilde{F} is proper because its first component, J , is proper and thus $\tilde{F}^N := t_{\tilde{\lambda}}^{\tilde{q}}(\tilde{F}|_N) : N \rightarrow \Sigma'$ is proper. Therefore $\tilde{F}^N : N \rightarrow \Gamma$ is proper because $(\tilde{F}^N)^{-1}(\Gamma \setminus \Sigma') = \emptyset$, and hence $(N, \omega^N, \tilde{F}^N, \Gamma)$ is a proper Hamiltonian \mathbb{T}^2 -manifold. Since $(N, \omega^N, \tilde{F}^N, \Gamma)$ is centered about $p \in \mathbb{R}^2$ by Lemma 8.3.3 we conclude that N is equivariantly symplectomorphic to

$$\{z \in \mathbb{C}^2 \mid p + |z_1|^2 \alpha_1 + |z_2|^2 \alpha_2 \in \Gamma\} = B^4(r).$$

It follows that there exists a symplectic \mathbb{T}^2 -embedding $\rho: B^4(r) \hookrightarrow M^\varepsilon$ with image N so $\tilde{F}(\rho(B^4(r))) = \tilde{F}(N) = \Sigma$. \square

Define the *optimal semitoric polygon packing function* $\pi_{\text{ST}}: \text{Polyg}_{\text{ST}}(\mathbb{R}^2) \rightarrow [0, \infty]$ by

$$\pi_{\text{ST}}([\Delta_w]) = \sup\{\text{vol}_{\mathcal{P}}(P) \mid P \text{ is an admissible semitoric packing of } \Delta_w\}.$$

It is well-defined because any two primitive semitoric polygons in the same orbit are related to one another by a transformation in $G_{m_f} \times \mathcal{G}$ which sends semitoric packings to semitoric packings and preserves volume.

Definition 8.3.5. We call $\alpha \in (0, \pi)$ a *smooth angle* if it can be obtained as an angle in a Delzant polygon.

Equivalently, $\alpha \in (0, \pi)$ is smooth if and only if it is the angle at the origin of $A_\alpha(\Delta(1))$ for some $A_\alpha \in \text{SL}_2(\mathbb{Z})$.

Lemma 8.3.6. *The set of smooth angles is discrete in $(0, \pi) \subset \mathbb{R}$.*

Proof. Fix a smooth angle $\alpha \in (0, \pi)$ and fix some $\varepsilon > 0$ small enough so that $(\alpha - \varepsilon, \alpha + \varepsilon) \subset (0, \pi)$.

Let

$$B_\varepsilon(\alpha) = \{\beta \in (0, \pi) \mid \beta \text{ is a smooth angle and } |\alpha - \beta| < \varepsilon\}$$

and let $\delta_\varepsilon > 0$ be such that if $\beta \in B_\varepsilon(\alpha)$ then $|\sin(\alpha) - \sin(\beta)| < \delta_\varepsilon$. Now fix any $\beta \in B_\varepsilon(\alpha)$. This means there exists some $A_\beta \in \text{SL}_2(\mathbb{Z})$ such that β is the angle at the origin of $\Delta = A_\beta(\Delta(1))$. Let $\ell_1, \ell_2 \in \mathbb{R}$ denote the lengths of two edges of the simplex Δ which are adjacent to the vertex at the

origin. These each represent the magnitude of a vector in \mathbb{Z}^n so $\ell_i \geq 1$ for $i = 1, 2$. By the choice of δ_ε we have that $\sin(\beta) > \sin(\alpha) - \delta_\varepsilon$. Since Δ has area $1/2$ we know that $\frac{\ell_1 \ell_2 \sin(\beta)}{2} = \frac{1}{2}$ and so for $i = 1, 2$ we conclude that $1 = \ell_1 \ell_2 \sin(\beta) \geq \ell_i \sin(\beta)$ which implies that

$$\ell_i \leq \frac{1}{\sin(\beta)} < \frac{1}{\sin(\alpha) - \delta_\varepsilon}.$$

Therefore associated to each $\beta \in B_\varepsilon(\alpha)$ there is a pair of vectors in \mathbb{Z}^2 each with length less than $(\sin(\alpha) - \delta_\varepsilon)^{-1}$, a value which does not depend on β . There are only finitely many such vectors. \square

The proof of Lemma 8.3.6 is taken from the proof of [31, Theorem 7.1] and is a two-dimensional version of the strategy used in Theorem 8.2.3. Let $\alpha \in (0, \pi)$ be called a *hidden smooth angle* if it can be obtained as a hidden corner in a primitive semitoric polygon.

Corollary 8.3.7. *The set of hidden smooth angles is discrete in $(0, \pi) \subset \mathbb{R}$.*

It is important to notice that a sequence of smooth angles can approach π . This must be the case, for example, if a semitoric polygon has infinitely many vertices.

Definition 8.3.8. We say that a vertex v of $(\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})$ is *non-fake* if it is either Delzant or hidden in one, and hence all, elements of the affine invariant. For $N \geq 1$ let $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)_0$ denote the set of primitive polygons with exactly N non-fake vertices and let $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ denote the set of $(G_{m_f} \times \mathcal{G})$ -orbits of elements of $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)_0$. Let \mathbb{I}^N be the set of all semitoric ingredients for which the affine invariant is an element of $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ and let

$$\text{Symp}_{\text{ST}, N}^{4, S^1 \times \mathbb{R}} = \Phi^{-1}(\mathbb{I}^N)$$

where Φ is as in Theorem 2.2.27.

Recall $\mathcal{H}_p^\varepsilon(v)$ defined in Equation (8.1). The following are two operations which can be performed on $[\Delta_w]$ to produce a new element of $\text{Polyg}_{\text{ST}}(\mathbb{R}^2)_0$.

Definition 8.3.9. Let $\Delta_w = (\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})$. Let $p \in \Delta$ be a vertex and let $v_1, v_2 \in \mathbb{Z}^2$ be the primitive inwards pointing normal vectors to the two edges which meet at p ordered so that $\det(v_1, v_2) > 0$.

If p is a Delzant vertex of Δ_w then the ε -corner chop of Δ_w at p is the primitive semitoric polygon

$$\Delta_w^{p,\varepsilon} = (\Delta \cap \mathcal{H}_p^\varepsilon(v_1 + v_2), (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f}).$$

Similarly, given $[\Delta_w]$ we say that $[\Delta_w^{p,\varepsilon}]$ is the ε -corner chop of $[\Delta_w]$ at p .

Suppose p is a hidden corner of Δ_w and thus there exists $j \in \{1, \dots, m_f\}$ such that $p \in \ell_{\lambda_j}$.

The ε -hidden corner chop of Δ_w at p is the primitive semitoric polygon

$$\Delta_w^{p,\varepsilon} = (\Delta \cap t_{\ell_{\lambda_j}}^{-1}(\mathcal{H}_p^\varepsilon(v_1 + v_2)), (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f}).$$

We say that $[\Delta_w^{p,\varepsilon}]$ is the ε -hidden corner chop of $[\Delta_w]$ at p .

The hidden corner chop of a hidden corner amounts to acting on the polygon with $t_{\ell_{\lambda_j}}^1$ to transform the hidden corner into a Delzant corner, performing the usual corner chop on this Delzant corner, and then transforming the polygon back with $t_{\ell_{\lambda_j}}^{-1}$. This is shown in Figure 8.5.




Figure 8.5: In (a) a hidden corner is shown. In (b) we unfold it by reversing the sign of the associated ε_i resulting in a Delzant corner. In (c) we perform corner chop on this corner and in (d) the ε_i returns to its original sign.

Lemma 8.3.10. Fix $N \in \mathbb{Z}_{\geq 0}$. Each $[\Delta_w] \in \text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ has an open neighborhood in $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ which consists exclusively of transformations of $[\Delta_w]$ in which its sides are moved in a parallel way. Moreover, any sufficiently small neighborhood of $[\Delta_w]$ in $\text{Polyg}_{\text{ST}}(\mathbb{R}^2)$ is contained in $\cup_{(N' \geq N)} \text{Polyg}_{\text{ST}}^{N'}(\mathbb{R}^2)$.

Proof. The angles of non-fake corners are discrete by Lemma 8.3.6 and Corollary 8.3.7. This means that there exists a neighborhood of $[\Delta_w]$ in which all elements which have N non-fake vertices must have all of the same angles as $[\Delta_w]$. This is the open neighborhood described in the Lemma. Any semitoric polygon with fewer non-fake vertices than $[\Delta_w]$ is bounded away from $[\Delta_w]$ because the only ways to change the number of non-fake vertices are a corner chop or introducing a smooth angle

into an edge of infinite length but by Lemma 8.3.6 smooth angles are discrete. \square

Lemma 8.3.11. *The map $\pi_{\text{ST}}: \text{Polyg}_{\text{ST}}(\mathbb{R}^2) \rightarrow [0, \infty]$ is discontinuous at every point.*

Proof. Primitive semitoric polygons must have at least one non-fake vertex. Let

$$[\Delta_w] = [(\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})]$$

be a semitoric polygon. First assume that $[\Delta_w] \in \text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ for some $N \geq 1$ and that $\pi_{\text{ST}}([\Delta_w]) < \infty$. Then for $\varepsilon > 0$ small enough define $[\Delta_w^\varepsilon]$ to be the semitoric polygon produced by performing an ε -corner chop at each non-fake vertex of $[\Delta_w]$. We have that

$$\lim_{\varepsilon \rightarrow 0} d_{\text{P}}^{\text{ST}}([\Delta], [\Delta_w^\varepsilon]) = 0. \quad (8.4)$$

A packing of $[\Delta_w^\varepsilon]$ has at most $2N$ disjoint admissible simplices. Since their side lengths are determined by the lengths of the adjacent edges, one of which is length ε , we have that $\lim_{\varepsilon \rightarrow 0} \pi_{\text{ST}}([\Delta_w^\varepsilon]) = 0$. Since every semitoric polygon has positive optimal packing we have

$$\lim_{\varepsilon \rightarrow 0} |\pi_{\text{ST}}([\Delta_w]) - \pi_{\text{ST}}([\Delta_w^\varepsilon])| = \pi_{\text{ST}}([\Delta_w]) > 0$$

and thus, in light of Equation (8.4), π_{ST} is discontinuous at $[\Delta_w]$.

Suppose $[\Delta_w] \in \text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ for some $N \geq 1$ and $\pi_{\text{ST}}([\Delta_w]) = \infty$. Since $[\Delta_w]$ has only finitely many non-fake vertices, any admissible packing has only finitely many admissible simplices. Hence there is a vertex at which an arbitrarily large simplex fits. The only possible case is that $N = 1$ and the polygon is of complexity zero. Taking a corner chop of any size at the single non-fake vertex produces a polygon on which π_{ST} evaluates to a finite number, so π_{ST} is discontinuous at $[\Delta_w]$.

Now suppose that $\pi_{\text{ST}}([\Delta_w]) < \infty$ and $[\Delta_w] \in \text{Polyg}_{\text{ST}}(\mathbb{R}^2) \setminus \bigcup_{N \geq 1} \text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$. For $i \in \mathbb{Z}_{\geq 1}$ let $I_i \subset \mathbb{R}$ be given by $I_i = [-n, n] \setminus (-(n-1), n-1)$ and let $N_i \in \mathbb{Z}_{\geq 0}$ denote the number of non-fake vertices of $[\Delta_w]$ with x -coordinate in I_i . This number is finite by the definition of a convex polygon and it is invariant under the action of $G_{m_f} \times \mathcal{G}$. For $\varepsilon > 0$ small enough let $[\Delta_w^\varepsilon]$ be a semitoric polygon which has a small corner chop at each non-fake vertex such that, at each vertex in I_i for $i \in \mathbb{Z}_{\geq 1}$, the largest possible admissible simplex that can fit into that vertex has volume at

most $\varepsilon/(N_i S^{i+1})$. Then an admissible packing R of $[\Delta_w^\varepsilon]$ satisfies

$$\text{vol}_{\mathcal{P}}(R) \leq \sum_{i=1}^{\infty} \frac{\varepsilon}{N_i 2^{i+1}} 2N_i = \varepsilon.$$

Therefore

$$\lim_{\varepsilon \rightarrow 0} d_{\mathcal{P}}^{\text{ST}}([\Delta_w], [\Delta_w^\varepsilon]) = 0$$

while

$$\lim_{\varepsilon \rightarrow 0} |\pi_{\text{ST}}([\Delta_w]) - \pi_{\text{ST}}([\Delta_w^\varepsilon])| = \pi_{\text{ST}}([\Delta_w]) > 0$$

and thus π_{ST} is not continuous at $[\Delta_w]$. \square

For $[\Delta_w] = [(\Delta, (\ell_{\lambda_j}, +1, k_j)_{j=1}^{m_f})] \in \text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ with non-fake vertices v_1, \dots, v_N , let $\pi_{\text{ST}}^{\mathcal{P}, i}(\Delta)$ be the total volume of the optimal packing excluding all packings which have a simplex centered at v_i .

Theorem 8.3.12. *Let $\pi_{\text{ST}}: \text{Polyg}_{\text{ST}}(\mathbb{R}^2) \rightarrow [0, \infty]$ be the optimal semitoric polygon packing function.*

Then:

1. π_{ST} is discontinuous at each point in $\text{Polyg}_{\text{ST}}(\mathbb{R}^2)$;
2. the restriction $\pi_{\text{ST}}|_{\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)}$ is continuous for each $N \in \mathbb{Z}_{\geq 1}$;
3. if $[\Delta_w] \in \text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ then $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ is the largest neighborhood of Δ_w in $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ in which π_{ST} is continuous if and only if $\pi_{\text{ST}}^i([\Delta_w]) < \pi_{\text{ST}}([\Delta_w])$ for all $1 \leq i \leq N$.

Proof. Part (1) is the content of Lemma 8.3.11.

By Lemma 8.3.10, given any $[\Delta_w] \in \text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$, there exists a neighborhood of $[\Delta_w]$ in $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ containing exclusively orbits of polygons formed by translating the sides of Δ_w in a parallel way. Hence part (2) follows from this because π_{ST} is continuous on such transformations.

For Part (3) suppose first that $\pi_{\text{ST}}([\Delta_w]) = \pi_{\text{ST}}^i([\Delta_w])$ for some $i \in \{1, \dots, N\}$. This means that there exists some optimal packing avoiding the i^{th} non-fake vertex. For $\varepsilon > 0$ let $[\Delta_w^\varepsilon]$ be the result of an ε -corner chop at the i^{th} vertex and notice that $\lim_{\varepsilon \rightarrow 0} d_{\mathcal{P}}^{\text{ST}}([\Delta_w], [\Delta_w^\varepsilon]) = 0$ and $\lim_{\varepsilon \rightarrow 0} \pi_{\text{ST}}([\Delta_w^\varepsilon]) = \pi_{\text{ST}}([\Delta_w])$. Thus there exists some set larger than $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)$ on which π_{ST} is continuous, as shown in Figure 8.6.




Figure 8.6: Corner chop of a corner not used in the optimal packing.

Finally, to show the converse assume that $[\Delta_w]$ satisfies $\pi_{\text{ST}}^{\mathcal{P},i}([\Delta_w]) < \pi_{\text{ST}}^{\mathcal{P}}([\Delta_w])$ for all $1 \leq i \leq N$. By Lemma 8.3.10 there is an open set around $[\Delta_w]$ in which the only elements not in $\text{Polyg}_{\text{ST}}^N(\mathbb{R}^2)_0$ are obtained from $[\Delta_w]$ by iterations of corner chops, parallel translations of the edges, and introducing a smooth angle into an edge of infinite length. For $\varepsilon > 0$ let $[\Delta_w^\varepsilon]$ be any ε -corner chop at the i^{th} non-fake vertex of $[\Delta_w]$. Then

$$\lim_{\varepsilon \rightarrow 0} \pi_{\text{ST}}([\Delta_w^\varepsilon]) = \pi_{\text{ST}}^i([\Delta_w]) < \pi_{\text{ST}}([\Delta_w])$$

and the result follows. \square

Notice that the quotient map $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}} \rightarrow \text{Polyg}_{\text{ST}}(\mathbb{R}^2)$ is continuous and the metric on $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}}$ is the sum of the metric on $\text{Polyg}_{\text{ST}}(\mathbb{R}^2)$ and the metric on the remaining components. Thus, Theorem 8.1.1 part (ii) follows from Theorem 8.3.12. For $(M, \omega, F) \in \text{Symp}_{\text{ST},N}^{4,S^1 \times \mathbb{R}}$ with fixed points $p_1, \dots, p_N \in M$ let

$$\mathcal{ST}^i(M) = \left(\frac{\sup\{\text{vol}(P) \mid P \subset M \text{ is a semitoric ball packing of } M \text{ and } p_i \notin P\}}{\text{vol}(\mathbb{B}^4)} \right)^{\frac{1}{4}}.$$

Proposition 8.3.13. *Let $N \geq 1$. If $(M, \omega, F) \in \text{Symp}_{\text{ST},N}^{4,S^1 \times \mathbb{R}}$ then $\text{Symp}_{\text{ST},N}^{4,S^1 \times \mathbb{R}}$ is the largest neighborhood of M in $\text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}}$ in which \mathcal{ST} is continuous if and only if $\mathcal{ST}^i(M) < \mathcal{ST}(M)$ for all $1 \leq i \leq N$.*

Theorem 8.1.1 part (ii) and Proposition 8.3.13 are illustrated in Figure 8.7.

Definition 8.3.14. The *semitoric radius capacity* is the symplectic $(S^1 \times \mathbb{R})$ -capacity

$$\mathcal{ST}_{\text{rad}} : \text{Symp}_{\text{ST}}^{4,S^1 \times \mathbb{R}} \rightarrow [0, \infty]$$




Figure 8.7: Continuous families of primitive semitoric polygons on which (a) $\mathcal{S}\mathcal{T}$ is continuous and (b) $\mathcal{S}\mathcal{T}$ is not continuous.

given by

$$\mathcal{S}\mathcal{T}_{\text{rad}}(M) = \sup\{r > 0 \mid \text{there exists a semitoric embedding } B^4(r) \hookrightarrow M\}.$$

It can be shown that $\mathcal{S}\mathcal{T}_{\text{rad}}$ is a $(S^1 \times \mathbb{R})$ -capacity in the same way that it was shown that $\mathcal{S}\mathcal{T}$ is a $(S^1 \times \mathbb{R})$ -capacity. Recall that $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{R}^n}$ is the symplectic \mathbb{R}^n -category which is the collection of toric manifolds with their \mathbb{T}^n -action lifted to an \mathbb{R}^n -action. Let $\text{Symp}_{\mathbb{T}, N}^{2n, \mathbb{R}^n}$ denote those systems with exactly N points fixed by the \mathbb{R}^n -action. By repeating the proofs of the continuity results Theorem 8.1.1 part (i), Proposition 8.2.4, Theorem 8.1.1 part (ii), and Proposition 8.3.13 we immediately have the following result, that yields Theorem 8.1.1 part (iii).

Theorem 8.3.15. *The maps $c_{\mathbb{B}}^{n, n}|_{\text{Symp}_{\mathbb{T}}^{2n, \mathbb{R}^n}}$ and $\mathcal{S}\mathcal{T}_{\text{rad}}$ are discontinuous everywhere on their domains and the restrictions $c_{\mathbb{B}}^{n, n}|_{\text{Symp}_{\mathbb{T}, N}^{2n, \mathbb{R}^n}}$ and $\mathcal{S}\mathcal{T}_{\text{rad}}|_{\text{Symp}_{\mathbb{S}\mathbb{T}, N}^{4, S^1 \times \mathbb{R}}}$ are both continuous. For*

$$(M, \omega, F) \in \text{Symp}_{\mathbb{T}, N}^{2n, \mathbb{R}^n}$$

the set $\text{Symp}_{\mathbb{T}, N}^{2n, \mathbb{R}^n}$ is not the largest neighborhood of M in $\text{Symp}_{\mathbb{T}}^{2n, \mathbb{R}^n}$ in which $c_{\mathbb{B}}^{n, n}|_{\text{Symp}_{\mathbb{T}}^{2n, \mathbb{R}^n}}$ is continuous and for $(M, \omega, F) \in \text{Symp}_{\mathbb{S}\mathbb{T}, N}^{4, S^1 \times \mathbb{R}}$ the set $\text{Symp}_{\mathbb{S}\mathbb{T}, N}^{4, S^1 \times \mathbb{R}}$ is the largest neighborhood of M in $\text{Symp}_{\mathbb{S}\mathbb{T}}^{4, S^1 \times \mathbb{R}}$ in which $\mathcal{S}\mathcal{T}_{\text{rad}}$ is continuous if and only if $N = 1$.

Acknowledgements. Chapter 8, in part, is comprised of material submitted for publication by the author of this dissertation, Alessio Figalli, and Álvaro Pelayo as *Symplectic G-capacities and integrable systems*, currently available as arXiv:1511.04499 [30]. Alessio Figalli is supported by NSF grants DMS-1262411 and DMS-1361122.

Chapter 9

Moduli spaces of maps

9.1 Introduction

In [74] the authors shown that if M and N are symplectic manifolds with $B_t \subset M$ for each $t \in (a, b)$ and

$$\{(\phi_t, B_t) \mid t \in (a, b) \text{ and } \phi_t: B_t \hookrightarrow N\}$$

is a smooth (see Definition 9.7.1) family of symplectic embeddings such that

1. each B_t is open and simply connected;
2. if $s < t$ then $\overline{B_t} \subset B_s$;
3. for all $t, s \in (a, b)$ the set $\bigcup_{v \in [t, s]} \phi_v(B_v)$ is relatively compact in N ,

then there exists a symplectic embedding

$$\phi_0: \bigcup_{t \in (a, b)} B_t \hookrightarrow N.$$

Given a collection of embeddings which satisfy certain conditions not related to convergence, this result assures the existence of an embedding from the union of their domains, which takes the place of the limit. In the present chapter, given a collection of embeddings which does not converge, I ask how much each embedding needs to be perturbed in order to produce a convergent collection. In particular, I am interested in situations in which each element of the collection may be perturbed

by an arbitrarily small amount in order to produce a converging family. In this case, unlike in the result above, I am more interested in the nature of the family of embeddings than the existence of such a limiting embedding. To formalize what I mean by a small perturbation, we define a distance function on maps which do not necessarily have the same domain. An L^p space is a space of maps with a natural metric, and we explain the relationship between the distance defined in the present chapter and the L^1 norm in Remark 9.3.3. Another example of examining a metric on a space of maps is the study of symplectic energy [43] which is defined in terms of a specific metric on the space of compactly supported Hamiltonian symplectomorphisms.

The study of collections of maps between smooth manifolds, particularly of embeddings or diffeomorphisms, has recently attracted a lot of interest [1, 7, 65, 74, 66]. Having a distance function defined on a collection of such mappings gives the collections the structure of a metric space about which new questions may be posed, as it is for instance done in [67]. It is the goal of this chapter to define a distance function on collections of maps with distinct domains, which are subsets of the same manifold, and study the properties of the resultant metric space.

9.1.1 Outline of chapter

In Section 9.2 I define the space of maps over which we will be working and the distance function. I state the main results of this chapter in Section 9.3. In Section 9.4 I prove several properties of the distance function including some parts of Theorem 9.3.1, and in Section 9.5 we prove the rest of Theorem 9.3.1. Next, in Section 9.6 I examine the convergence properties of the distance and prove Theorem 9.3.5. Finally, in Section 9.7 I use what we have established in the preceding sections to study families of embeddings which do not converge to an embedding and prove Theorem 9.3.8. In the last section, Section 9.8, I comment on how the ideas from this chapter can be used to further study such families and mention some other possibilities for applications of this distance.

9.2 The distance function

Considering families of maps with different domains is essential for applications, see for instance the work of Pelayo-Vũ Ngọc [74, 66]. Suppose that the maps are defined on subsets of a

smooth manifold M with a volume form \mathcal{V} and map to a complete Riemannian manifold N with natural distance d . By this we mean that if g is the Riemannian metric on N and $y_1, y_2 \in N$ then

$$d(y_1, y_2) = \inf \left\{ \int_0^1 \sqrt{g(\gamma'(t), \gamma'(t))} dt \left| \begin{array}{l} \gamma : [0, 1] \rightarrow N \text{ is piecewise } C^1 \text{ with} \\ \gamma(0) = y_1 \text{ and } \gamma(1) = y_2 \end{array} \right. \right\}.$$

We will soon see that the properties of the distance will not depend on the choice of metric g and it is known that any smooth manifold admits a complete Riemannian metric, so we are not making any assumptions on N . Throughout the chapter by *metric* we will always mean a metric function on the space and if referring to a metric tensor we will always specify the Riemannian metric. Also, it is well known (see the Hopf-Rinow Theorem [45, Satz I]) that (N, g) is a geodesically complete Riemannian manifold if and only if (N, d) is a complete metric space, so throughout this chapter we will call such a manifold complete without specifying. Let $\mu_{\mathcal{V}}$ be the measure on M induced by \mathcal{V} . That is, for any $A \subset M$ we have $\mu_{\mathcal{V}}(A) = \int_A \mathcal{V}$. Now we will define the set of maps we will be working with (shown in Figure 9.1).

Definition 9.2.1. Let

$$\mathcal{M}(M, N) := \left\{ (\phi, B_\phi) \left| \begin{array}{l} B_\phi \subset M \text{ a nonempty measurable set and} \\ \phi : B_\phi \rightarrow N \text{ a measurable function} \end{array} \right. \right\}$$

which we will frequently denote by \mathcal{M} when M and N are understood and we will also frequently write only ϕ where the associated domain is understood to be denoted by B_ϕ . Also let

$$\mathcal{F}(\mathcal{M}) = \{ \{(\phi_t, B_t)\}_{t \in (a, b)} \subset \mathcal{M} \mid a, b \in \mathbb{R} \text{ with } a < b \}.$$

For the remaining chapter we will denote by $\mathcal{F}(\mathcal{S})$ the collection of one parameter families in a set \mathcal{S} indexed by an open interval in \mathbb{R} .

Recall the symmetric difference of sets A and B is given by $A \Delta B = (A \setminus B) \cup (B \setminus A)$.




Figure 9.1: I will be considering maps from subsets of M to N .

Definition 9.2.2. For $(\phi, B_\phi), (\psi, B_\psi) \in \mathcal{M}$ we define the *penalty function* $p_{\phi\psi}^d : M \rightarrow [0, 1]$ by

$$p_{\phi\psi}^d(x) = \begin{cases} 1 & \text{if } x \in B_\phi \Delta B_\psi; \\ \min\{1, d(\phi(x), \psi(x))\} & \text{if } x \in B_\phi \cap B_\psi; \\ 0 & \text{otherwise,} \end{cases}$$

and we define

$$\mathcal{D}_M^d((\phi, B_\phi), (\psi, B_\psi)) = \int_M p_{\phi\psi}^d d\mu_\nu.$$

A reasonable first guess for the “distance” between two elements in \mathcal{M} would be to integrate a penalty function over M . That is, we start with a function which assigns a penalty at each point in M depending on how different the mappings are at that point, and then compute the “distance” between the two mappings by adding up all of these penalties via integration. For each point in the symmetric difference, we know that one mapping acts on it while the other does not, so we assign it a maximum penalty of 1. For each point which is in the intersection of the domains, we simply find the distance between where each map sends the point, cut off to not exceed a maximum value of 1, and use this as the penalty.

Notice that we need the minimum in Definition 9.2.2 to make sure that any point on which both mappings act is not penalized more than the points which are only acted on by one mapping. It is worth noting that even though the choice of the constant 1 may seem arbitrary it is shown in Proposition 9.4.3 that any positive constant may be used instead and the induced distance will be strongly equivalent (see Definition 9.4.1). Also, if d is chosen so that the metric space (N, d) is complete (which can always be done [60, Theorem 1]) the choice of d will not change the properties




Figure 9.2: A graphic representing the values of $p_{\phi\psi}^d$ on $S \subset M$.

of the induced metric.

However while \mathcal{D}_M^d is the natural “distance” it turns out to not be a distance function on \mathcal{M} . There are two main problems. First, it is possible that \mathcal{D}_M^d will evaluate to zero on two distinct elements of \mathcal{M} and second it might be that \mathcal{D}_M^d evaluates to infinity. The first problem is a common one and can be addressed in the standard way, by having \mathcal{D}_M^d act on equivalence classes of maps, but the second problem will require a more delicate solution.

The problem of \mathcal{D}_M^d evaluating to infinity is even worse than it seems. Suppose that $\phi_t(x) = (x, t)$ takes \mathbb{R} into \mathbb{R}^2 for all $t \in (0, 1)$. Using the notation from above in this case we have that $M = B_{\phi_t} = \mathbb{R}$ for all $t \in (0, 1)$ and $N = \mathbb{R}^2$ with $d_{\mathbb{R}^2}$ the usual distance. Then ϕ_t has a pointwise limit of $\phi_0(x) := (x, 0)$ as $t \rightarrow 0$, but despite this we have that $\mathcal{D}_M^d(\phi_t, \phi_0)$ is infinite for all $t \in (0, 1)$. This example shows that \mathcal{D}_M^d is not always able to capture when a family of maps is converging. I am able to solve this problem by observing \mathcal{D}_M^d restricted to various subsets of M .

Definition 9.2.3. I define \mathcal{D} restricted to a measurable set $S \subset M$ by

$$\mathcal{D}_S^d((\phi, B_\phi), (\psi, B_\psi)) = \int_S p_{\phi\psi}^d \, d\mu_V.$$

Figure 9.2 shows a good way to visualize computing \mathcal{D}_S^d . Now each \mathcal{D}_S^d contains all of the information about \mathcal{D}_M^d on the set S and, as long as S is chosen to be of finite volume, \mathcal{D}_S^d cannot evaluate to infinity. The problem now, of course, is that we no longer have just a single metric with

information about all of M but instead have an infinite family of metrics which each have information about only one finite volume subset of M . I solve this last problem by recalling that any manifold admits a nested exhaustion by compact sets, which must each have finite volume. For the remaining portion of this chapter by exhaustion we will always mean a countable nested exhaustion by finite volume sets. In the following definition we set up the framework for this chapter. I will write $\nu_{\{S_n\}}$ in place of $\nu_{\{S_n\}_{n=1}^\infty}$ and $\mathcal{D}_{\{S_n\}}^d$ in place of $\mathcal{D}_{\{S_n\}_{n=1}^\infty}^d$ for simplicity.

Definition 9.2.4. Let M and N be manifolds with d a metric on N induced by a Riemannian metric.

1. Let $\{S_n\}_{n=1}^\infty$ be a exhaustion of M by nested finite volume sets and let $\nu_{\{S_n\}}$ be the measure on M given by

$$\nu_{\{S_n\}}(A) = \sum_{n=1}^{\infty} 2^{-n} \frac{\mu_{\mathcal{V}}(A \cap S_n)}{\mu_{\mathcal{V}}(S_n)}$$

for $A \subset M$. Notice that $\nu_{\{S_n\}}(M) = 1$ so $\nu_{\{S_n\}}$ is a probability measure. Then define

$$\mathcal{D}_{\{S_n\}}^d(\phi, \psi) = \int_M p_{\phi\psi}^d d\nu_{\{S_n\}}.$$

2. If $\mathcal{D}_{\{S_n\}}^d(\phi, \psi) = 0$ for one choice of exhaustion then, by Corollary 9.4.7, it equals zero for all choices of exhaustion and complete metrics d , so in that case we write $\mathcal{D}(\phi, \psi) = 0$.
3. Let

$$\mathcal{M}^\sim(M, N) := \mathcal{M}(M, N) / \sim$$

where $(\phi, B_\phi) \sim (\psi, B_\psi)$ if and only if $\mathcal{D}(\phi, \psi) = 0$. As before we will frequently shorten this to \mathcal{M}^\sim and we denote by $[\phi, B_\phi]$ the equivalence class of $(\phi, B_\phi) \in \mathcal{M}$.

There is an equivalent definition of $\mathcal{D}_{\{S_n\}}^d$ given in Proposition 9.4.4 which is used in some of the proofs in this chapter and explicitly shows the relation between $\mathcal{D}_{\{S_n\}}^d$ and \mathcal{D}_S^d .

9.3 Main results

Now we have enough notation to state our first result. Let M and N be manifolds and \mathcal{V} a volume form on M .

Theorem 9.3.1. *For any choice of a metric d on N induced by a complete Riemannian metric and a countable exhaustion $\{\mathcal{S}_n\}_{n=1}^\infty$ of M by nested finite volume sets, the space $(\mathcal{M}^\sim, \mathcal{D}_{\{\mathcal{S}_n\}}^d)$ is a complete metric space. Moreover, such a metric and exhaustion always exist and if d' and $\{\mathcal{S}'_n\}_{n=1}^\infty$ are other such choices then $\mathcal{D}_{\{\mathcal{S}'_n\}}^{d'}$ induces the same topology as $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ on \mathcal{M}^\sim .*

In light of Theorem 9.3.1 we can now make the following definitions. Recall that $\mathcal{F}(\mathcal{M})$ denotes the collection of one-parameter families of \mathcal{M} indexed by an interval $(a, b) \subset \mathbb{R}$.

Definition 9.3.2. Let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Also let $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ and $\phi_0 \in \mathcal{M}$.

1. Let $S \subset M$ be any subset. If $\lim_{t \rightarrow c} \mathcal{D}_S^d(\phi_t, \phi_0) = 0$ we write

$$\phi_t \xrightarrow{\mathcal{D}_S^d} \phi_0 \text{ as } t \rightarrow c.$$

2. If $\lim_{t \rightarrow c} \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t, \phi_0) = 0$ for one, and hence all, choices of $\{\mathcal{S}_n\}_{n=1}^\infty$ and d , we write

$$\phi_t \xrightarrow{\mathcal{D}} \phi_0 \text{ as } t \rightarrow c.$$

3. Since all metrics $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ generate the same topology on the set \mathcal{M}^\sim we denote this set with such topology as $(\mathcal{M}^\sim, \mathcal{D})$.

Thus \mathcal{M}^\sim is a metric space with metric $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ for any choice of exhaustion and complete metric and the metric spaces for different choices of exhaustion are all equivalent topologically. Notice that all of the information about $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ is contained in \mathcal{D}_M^d if M is finite volume, and in this case we will only have to consider \mathcal{D}_M^d , see Remark 9.4.12.

Remark 9.3.3. Recall that L^p spaces are collections of maps from a fixed measure set to \mathbb{R} . Since \mathcal{M} is a collection of all maps between fixed manifolds we can see that in some sense \mathcal{M} is a generalization of L^p spaces. The function $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ is similar to the L^1 norm, but there are several differences. It is noteworthy that *any* measurable mapping from M to N is “integrable” with respect to $\mathcal{D}_{\{\mathcal{S}_n\}}^d$, in the sense that the distance between any two measurable mappings is finite. This is why \mathcal{M} includes all measurable maps, while L^p includes only functions which satisfy a growth restriction.

In Example 9.4.13 we work out a specific case which does not converge in L^p for any p but does converge with respect to the distance defined in this chapter. \circlearrowright

Now that there is a metric defined on \mathcal{M}^\sim we can explore families in $\mathcal{F}(\mathcal{M}^\sim)$ which converge with respect to that metric. In Section 9.6 we study another type of convergence and we explore the connection between these two natural forms of convergence on \mathcal{M}^\sim . The limit inferior and limit superior of a family of sets are reviewed in Equations (9.6) and (9.7) in Section 9.6.

Definition 9.3.4. Let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Let $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ and suppose there exists some measurable $B \subset M$ satisfying

$$B \subset \left\{ x \in \varliminf_{t \rightarrow c} B_t \mid \lim_{t \rightarrow c} \phi_t(x) \text{ exists} \right\}$$

and $\mu_\nu(\overline{\lim}_{t \rightarrow c} B_t \setminus B) = 0$. This in particular requires that the domains converge as sets as is described in Definition 9.6.1. Then, with

$$\begin{aligned} \phi : B &\rightarrow N \\ x &\mapsto \lim_{t \rightarrow c} \phi_t(x). \end{aligned}$$

we say that $\{(\phi_t, B_t)\}_{t \in (a, b)}$ converges to (ϕ, B) almost everywhere pointwise as $t \rightarrow c$ in \mathcal{M} and we write $\phi_t \xrightarrow{\text{a.e.}} \phi$ as $t \rightarrow c$.

Theorem 9.3.5. Let $a, b, c \in \mathbb{R}$ such that $a < b$ and $c \in [a, b]$. Suppose $\{(\phi_t, B_t)\}_{t \in (a, b)}$ is a family such that $(\phi_t, B_t) \in \mathcal{M}$ for $t \in (a, b)$ and let $(\phi, B) \in \mathcal{M}$. If $\phi_t \xrightarrow{\text{a.e.}} \phi$ as $t \rightarrow c$ then $\phi_t \xrightarrow{\mathcal{D}} \phi$ as $t \rightarrow c$.

There are many different directions one could head from this point, but since there is research already being done regarding the convergence properties of families of embeddings [66, 74] we will pursue an application in that field. I will use \mathcal{D} to study families of embeddings which do not converge to an embedding and quantify how far they are from converging. With this in mind we make the following definitions.

Definition 9.3.6. Define $\text{Emb}_\subset(M, N) \subset \mathcal{M}$ to be those elements $(\phi, B) \in \mathcal{M}$ such that $B \subset M$

is a submanifold and $\phi: B \hookrightarrow N$ is an embedding. Also define $\text{Emb}_{\subset}^{\sim}(M, N) \subset \mathcal{M}^{\sim}$ to be those equivalence classes $[\phi, B] \in \mathcal{M}^{\sim}$ such that $[\phi, B]$ includes an element of $\text{Emb}_{\subset}(M, N)$.

Definition 9.3.7. Let $a, b \in \mathbb{R}$ with $a < b$, $\varepsilon \geq 0$, and $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$. A smooth family $\{(\tilde{\phi}_t, \tilde{B}_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ is a *convergent ε -perturbation* (with respect to $\mathcal{D}_{\{\mathcal{S}_n\}}^d$) of $\{(\phi_t, B_t)\}_{t \in (a, b)}$ if

1. there exists $(\tilde{\phi}, \tilde{B}) \in \text{Emb}_{\subset}(M, N)$ such that $\tilde{\phi}_t \xrightarrow{\text{a.e.}} \tilde{\phi}$ as $t \rightarrow a$;
2. $B_t = \tilde{B}_t$ for all $t \in (a, b)$ and

$$\varliminf_{t \rightarrow c} \tilde{B}_t \subset \tilde{B};$$

3. for all $t \in (a, b)$ we have that $\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t, \tilde{\phi}_t) \leq \varepsilon$.

The function

$$r_{\{\mathcal{S}_n\}}^d : \mathcal{F}(\mathcal{M}) \rightarrow [0, \infty]$$

takes a family in $\mathcal{F}(\mathcal{M})$ to its *radius of convergence* given by

$$r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}_{t \in (a, b)}) := \inf \left\{ \varepsilon \geq 0 \left| \begin{array}{l} \text{there exists a smooth convergent} \\ \varepsilon\text{-perturbation of } \{(\phi_t, B_t)\}_{t \in (a, b)} \end{array} \right. \right\}.$$

In part 2 of Definition 9.3.7 we make a requirement on the domains. This is so that the singular points cannot simply be removed from the domain to form a convergent ε -perturbation. It is important to notice that, unlike many of the properties we have introduced so far, $r_{\{\mathcal{S}_n\}}^d$ does depend on the choice of d and $\{\mathcal{S}_n\}_{n=1}^{\infty}$. I am most interested in the $r_{\{\mathcal{S}_n\}}^d = 0$ case, where an arbitrarily small perturbation can cause the family to converge to an embedding. It is natural to wonder whether a family can have radius of convergence zero but still not converge to any element of \mathcal{M} . The following Theorem addresses this.

Theorem 9.3.8. Let $a, b \in \mathbb{R}$ with $a < b$, $\{(\phi_t, B_t)\}_{t \in (a, b)}$ be such that $(\phi_t, B_t) \in \mathcal{M}$ for each $t \in (a, b)$, and let $r_{\{\mathcal{S}_n\}}^d$ be the radius of convergence function associated to a complete Riemannian distance d on N and an exhaustion of finite volume nested sets $\{\mathcal{S}_n\}_{n=1}^{\infty}$ of M . If $r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}_{t \in (a, b)}) = 0$ then there exists $(\phi, B) \in \mathcal{M}$ unique up to \sim such that $\phi_t \xrightarrow{\mathcal{D}} \phi$ as $t \rightarrow a$. Furthermore, the converse holds if there exists some $T \in (a, b)$ such that $s < t < T$ implies $B_s \subset B_t$.

This theorem is important in the study of families with $r_{\{S_n\}}^d = 0$ because to characterize such families we may assume right away that there exists some limit ϕ_0 and study its properties in order to understand the family we started with. In the final section we explore some ideas about the open questions regarding $r_{\{S_n\}}^d$ including restricting to embeddings with specific properties and considering a converse of Theorem 9.3.8 in the case in which the domains do not eventually shrink or stabilize.

9.4 Definitions and preliminaries

9.4.1 Basic properties of the distance

Let M be an orientable smooth manifold with volume form \mathcal{V} and let N be a smooth Riemannian manifold with natural distance function d . Again let $\mu_{\mathcal{V}}$ be the measure on M induced by the volume form \mathcal{V} . In this section we will prove all but the completeness statement in Theorem 9.3.1, which is postponed to Section 9.5. Recall the different notions of equivalent metrics. The use of these terms varies, but for this chapter we will use the following conventions.

Definition 9.4.1. Let d_1 and d_2 be metrics on a set X . Then we say that d_1 and d_2 are:

1. *topologically equivalent* if they induce the same topology on X ;
2. *weakly equivalent* if they induce the same topology on X and exactly the same collection of Cauchy sequences;
3. *strongly equivalent* if there exist $c_1, c_2 > 0$ such that

$$c_1 d_1 \leq d_2 \leq c_2 d_1.$$

Now we define the following function.

Definition 9.4.2. Let $(\phi, B_\phi), (\psi, B_\psi) \in \mathcal{M}$. For $\alpha > 0$ and a finite volume subset $S \subset M$ define

$$\mathcal{D}_S^{d,\alpha}((\phi, B_\phi), (\psi, B_\psi)) = \int_S p_{\phi\psi}^{d,\alpha} d\mu_{\mathcal{V}}$$

where

$$p_{\phi\psi}^{d,\alpha}(x) = \begin{cases} \alpha & \text{if } x \in B_\phi \Delta B_\psi; \\ \min\{\alpha, d(\phi(x), \psi(x))\} & \text{if } x \in B_\phi \cap B_\psi; \\ 0 & \text{otherwise.} \end{cases}$$

In Definition 9.4.2 we have a family of functions depending on the choice of $\alpha > 0$, but in fact these will induce strongly equivalent metrics.

Proposition 9.4.3. *Let \mathcal{S} be a finite volume subset of M . If $\beta > \alpha > 0$ then*

$$\mathcal{D}_S^{d,\alpha} \leq \mathcal{D}_S^{d,\beta} \leq \frac{\beta}{\alpha} \mathcal{D}_S^{d,\alpha}.$$

Proof. Notice

$$\begin{aligned} \mathcal{D}_S^{d,\alpha}(\phi, \psi) &= \int_{B_\phi \cap B_\psi \cap \mathcal{S}} \min\{\alpha, d(\phi, \psi)\} d\mu_V + \alpha \mu_V((B_\phi \Delta B_\psi) \cap \mathcal{S}) \\ &\leq \int_{B_\phi \cap B_\psi \cap \mathcal{S}} \min\{\beta, d(\phi, \psi)\} d\mu_V + \beta \mu_V((B_\phi \Delta B_\psi) \cap \mathcal{S}) \\ &= \mathcal{D}_S^{d,\beta}(\phi, \psi) \end{aligned}$$

and also notice that

$$\begin{aligned} \mathcal{D}_S^{d,\beta}(\phi, \psi) &= \int_{B_\phi \cap B_\psi \cap \mathcal{S}} \min\{\beta, d(\phi, \psi)\} d\mu_V + \beta \mu_V((B_\phi \Delta B_\psi) \cap \mathcal{S}) \\ &\leq \int_{B_\phi \cap B_\psi \cap \mathcal{S}} \min\{\beta, \frac{\beta}{\alpha} d(\phi, \psi)\} d\mu_V + \beta \mu_V((B_\phi \Delta B_\psi) \cap \mathcal{S}) \\ &= \frac{\beta}{\alpha} \mathcal{D}_S^{d,\alpha}(\phi, \psi). \end{aligned}$$

□

So Proposition 9.4.3 means that the choice of $\alpha > 0$ will not matter when we use $\mathcal{D}_S^{d,\alpha}$ to define a metric, so henceforth we will assume that $\alpha = 1$. That is, for any finite volume subset $\mathcal{S} \subset M$ we have \mathcal{D}_S^d as defined in Definition 9.2.3. In the above proof we wrote out the definition of \mathcal{D}_S^d in

a way which did not explicitly use the penalty function $p_{\phi\psi}^d$. We can now notice that there is an equivalent definition of $\mathcal{D}_{\mathcal{S}}^d$ which will be useful for several of the proofs.

Proposition 9.4.4. *Let M and N be manifolds with a volume form \mathcal{V} on M , d a distance on N induced by a Riemannian metric, $\mathcal{S} \subset M$ a compact subset, and $\{\mathcal{S}_n\}_{n=1}^{\infty}$ a nested exhaustion of M by finite volume sets. The function $\mathcal{D}_{\mathcal{S}}^d$ given in Definition 9.2.3 can be written*

$$\mathcal{D}_{\mathcal{S}}^d(\phi, \psi) = \int_{B_{\phi} \cap B_{\psi} \cap \mathcal{S}} \min\{1, d(\phi, \psi)\} d\mu_{\mathcal{V}} + \mu_{\mathcal{V}}((B_{\phi} \Delta B_{\psi}) \cap \mathcal{S}).$$

and the function $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ from Definition 9.2.4 satisfies

$$\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi, \psi) = \sum_{n=1}^{\infty} 2^{-n} \frac{\mathcal{D}_{\mathcal{S}_n}^d(\phi, \psi)}{\mu_{\mathcal{V}}(\mathcal{S}_n)}.$$

This proposition has a trivial proof. Before the next Proposition we have a definition.

Definition 9.4.5. Suppose $a, b \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. For a set X and a function

$$F : X \times X \rightarrow [0, \infty]$$

we say that a family $\{a_t\}_{t \in (a,b)} \subset X$ is *Cauchy with respect to F as $t \rightarrow c$* if for all $\varepsilon > 0$ there exists some $\delta > 0$ such that $s, t \in (c - \delta, c + \delta) \cap (a, b)$ implies $F(a_t, a_s) < \varepsilon$.

Below are several important properties of $\mathcal{D}_{\{\mathcal{S}_n\}}^d$, which is defined in Definition 9.2.4.

Proposition 9.4.6. *Let $a, b \in \mathbb{R}$ with $a < b$, $\{(\phi_t, B_t)\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M})$, and $\phi, \psi \in \mathcal{M}$. Further suppose that d is a metric on N induced by a Riemannian metric and $\{\mathcal{S}_n\}_{n=1}^{\infty}$ is an exhaustion of M by nested finite volume sets. The function $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ has the following properties.*

1. $\{(\phi_t, B_t)\}_{t \in (a,b)}$ is Cauchy with respect to $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ as $t \rightarrow c$ if and only if it is Cauchy with respect to $\mathcal{D}_{\mathcal{S}}^d$ as $t \rightarrow c$ for all compact $\mathcal{S} \subset M$.
2. $\lim_{t \rightarrow c} \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t, \phi) = 0$ if and only if $\phi_t \xrightarrow{\mathcal{D}_{\mathcal{S}}^d} \phi$ as $t \rightarrow c$ for all compact $\mathcal{S} \subset M$.
3. $\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi, \psi) = 0$ if and only if $\mathcal{D}_{\mathcal{S}}^d(\phi, \psi) = 0$ for all compact $\mathcal{S} \subset M$ if and only if $\mu_{\mathcal{V}}((B_{\phi} \Delta B_{\psi}) \cap \mathcal{S}) = 0$ for every compact $\mathcal{S} \subset M$ and $\phi = \psi$ almost everywhere on $B_{\phi} \cap B_{\psi}$.

Proof. Let $\varepsilon > 0$ and fix some compact subset $\mathcal{S} \subset M$. Then $\mathcal{S} \subset \bigcup_{n=1}^{\infty} \mathcal{S}_n = M$ and since \mathcal{S} has finite volume and the \mathcal{S}_n are nested we can find some $I \in \mathbb{N}$ such that $\mu_{\mathcal{V}}(\mathcal{S} \setminus \mathcal{S}_I) < \varepsilon$. This means that

$$\mathcal{D}_{\mathcal{S}}^d \leq \mathcal{D}_{\mathcal{S}_I}^d + \varepsilon.$$

Now that we have this fact we will prove the three properties.

(1) It is sufficient to assume that $a = c = 0$ and $b = 1$. Suppose that $\{(\phi_t, B_t)\}_{t \in (0,1)}$ is Cauchy with respect to $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ as $t \rightarrow 0$ and fix some compact $\mathcal{S} \subset M$. Let $\varepsilon > 0$.

From the above fact we can find some $I \in \mathbb{N}$ such that $\mathcal{D}_{\mathcal{S}}^d \leq \mathcal{D}_{\mathcal{S}_I}^d + \varepsilon/2$. Now, since this family is Cauchy with respect to $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ we can find some $\delta \in (0,1)$ such that $s, t < \delta$ implies

$$\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t, \phi_s) < \frac{\varepsilon}{2^{I+1}\mu_{\mathcal{V}}(\mathcal{S}_I)}.$$

Using the expression for $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ from Proposition 9.4.4 we have that

$$\sum_{n=1}^{\infty} 2^{-n} \frac{\mathcal{D}_{\mathcal{S}_n}^d(\phi_t, \phi_s)}{\mu_{\mathcal{V}}(\mathcal{S}_n)} < \frac{\varepsilon}{2^{I+1}\mu_{\mathcal{V}}(\mathcal{S}_I)}$$

which in particular means

$$2^{-I} \frac{\mathcal{D}_{\mathcal{S}_I}^d(\phi_t, \phi_s)}{\mu_{\mathcal{V}}(\mathcal{S}_I)} < \frac{\varepsilon}{2^{I+1}\mu_{\mathcal{V}}(\mathcal{S}_I)}$$

so $\mathcal{D}_{\mathcal{S}_I}^d(\phi_t, \psi_t) < \varepsilon/2$.

Finally, we have that for $s, t < \delta$

$$\mathcal{D}_{\mathcal{S}}^d(\phi_t, \phi_s) \leq \mathcal{D}_{\mathcal{S}_I}^d(\phi_t, \phi_s) + \frac{\varepsilon}{2} < \varepsilon.$$

The converse is easy and the proof of (2) is similar to the proof of (1).

(3) Suppose $\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi, \psi) = 0$ and fix some compact $\mathcal{S} \subset M$. Notice that this means that $\mathcal{D}_{\mathcal{S}_n}^d(\phi, \psi) = 0$ for all n . For any $\varepsilon > 0$ from the fact above we know we can choose some I such that

$$\mathcal{D}_{\mathcal{S}}^d(\phi, \psi) \leq \mathcal{D}_{\mathcal{S}_I}^d(\phi, \psi) + \varepsilon = \varepsilon$$

so we may conclude that $\mathcal{D}_S^d(\phi, \psi) = 0$.

Next, assume that $\mathcal{D}_S^d(\phi, \psi) = 0$ for all compact $S \subset M$, which means that

$$\mu_V((B_\phi \Delta B_\psi) \cap S) = 0$$

because this is a term in \mathcal{D}_S^d . Suppose that there is some set of positive measure in $B_\phi \cap B_\psi$ for which $\phi \neq \psi$. Then since manifolds are inner regular there exists some compact subset of positive measure K on which they are not equal. But this implies that $\mathcal{D}_K^d(\phi, \psi) \neq 0$.

Now since $\mu_V((B_\phi \Delta B_\psi) \cap S) = 0$ for every compact $S \subset M$ and $\phi = \psi$ almost everywhere on $B_\phi \cap B_\psi$ it is clear that $\mathcal{D}_{\{S_n\}}^d(\phi, \psi) = 0$. \square

Corollary 9.4.7. *Let $\{S_n\}_{n=1}^\infty$ be an exhaustion of M and let d be a metric on N induced by a Riemannian metric. Suppose that $(\phi, B_\phi), (\psi, B_\psi) \in \mathcal{M}$ such that $\mathcal{D}_{\{S_n\}}^d(\phi, \psi) = 0$. Then for any such parameters $\{S'_n\}_{n=1}^\infty$ and d' we have that $\mathcal{D}_{\{S'_n\}}^{d'}(\phi, \psi) = 0$ as well.*

Given the new information in Proposition 9.4.6 we can prove the following important Proposition.

Proposition 9.4.8. *For any choice of an exhaustion of M by finite volume sets $\{S_n\}_{n=1}^\infty$ we have that $\mathcal{D}_{\{S_n\}}^d$ is well defined and is a distance function on \mathcal{M}^\sim . Also, if $\{S'_n\}_{n=1}^\infty$ is another such choice of exhaustion then $\mathcal{D}_{\{S_n\}}^d$ and $\mathcal{D}_{\{S'_n\}}^d$ are weakly equivalent metrics on \mathcal{M}^\sim .*

Proof. Fix some $\{S_n\}_{n=1}^\infty$ a compact exhaustion of M and let $\phi, \rho, \psi \in \mathcal{M}$. It is a straightforward exercise to show that

$$p_{\phi\psi}^d(x) \leq p_{\phi\rho}^d(x) + p_{\rho\psi}^d(x)$$

for each $x \in M$ and thus

$$\mathcal{D}_{\{S_n\}}^d(\phi, \psi) \leq \mathcal{D}_{\{S_n\}}^d(\phi, \rho) + \mathcal{D}_{\{S_n\}}^d(\rho, \psi).$$

It should be noted that this inequality would not hold without the minimum in $p_{\phi\psi}^d$. From here we can see that if $\phi \sim \rho$ then

$$\mathcal{D}_{\{S_n\}}^d(\phi, \psi) \leq \mathcal{D}_{\{S_n\}}^d(\rho, \psi)$$

and similarly the opposite inequality is true as well. So

$$\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi, \psi) = \mathcal{D}_{\{\mathcal{S}_n\}}^d(\rho, \psi)$$

and thus $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ is well defined on \mathcal{M}^\sim .

Now $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ is positive definite on \mathcal{M}^\sim because it is positive on \mathcal{M} and by definition $\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi, \rho) = 0$ implies $\phi \sim \rho$. Since $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ is well defined on \mathcal{M}^\sim and satisfies the triangle inequality on \mathcal{M} we know that it satisfies the triangle inequality on \mathcal{M}^\sim and similarly we know that $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ is symmetric on \mathcal{M}^\sim .

Proposition 9.4.6 parts (1) and (2) characterize both convergent and Cauchy sequences of $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ in a way which is independent of the choice of $\{\mathcal{S}_n\}_{n=1}^\infty$. This means that different choices of $\{\mathcal{S}_n\}_{n=1}^\infty$ will produce weakly equivalent metrics $\mathcal{D}_{\{\mathcal{S}_n\}}^d$. \square

9.4.2 Independence of Riemannian structure

We have seen that \mathcal{M}^\sim is a metric space with metric $\mathcal{D}_{\{\mathcal{S}_n\}}^d$ for any choice of compact exhaustion and the metric spaces for different choices of exhaustion are all weakly equivalent. Now we will show that this construction is actually independent of the choice of Riemannian metric on N as well. For the remaining portion of the chapter we will use $\|\cdot\|$ to denote the usual norm in \mathbb{R}^k and $d_{\mathbb{R}^k}$ to denote the usual distance on \mathbb{R}^k .

Lemma 9.4.9. *Fix any measurable finite volume subset $\mathcal{S} \subset M$ and let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Now let $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ and $(\phi, B) \in \mathcal{M}$. Suppose that $\phi_t \xrightarrow{\mathcal{D}_{\mathcal{S}}^d} \phi \in \mathcal{M}$ as $t \rightarrow c$ and $\mathcal{R} : B \cap \mathcal{S} \rightarrow (0, \infty)$ is any function. Then*

$$\lim_{t \rightarrow c} \mu_V(\{x \in B_t \cap B \cap \mathcal{S} \mid d(\phi(x), \phi_t(x)) > \mathcal{R}(x)\}) = 0.$$

Proof. It is sufficient to prove for $a = c = 0$ and $b = 1$. First, for $t \in (0, 1)$ let $C_t = \{x \in B_t \cap B \cap \mathcal{S} \mid$

$d(\phi(x), \phi_t(x)) > \mathcal{R}(x)$. Since $C_t \subset \mathcal{S}$ we notice that

$$\begin{aligned} \mathcal{D}_{\mathcal{S}}^d(\phi_t, \phi) &\geq \int_{C_t} \min\{1, d(\phi, \phi_t)\} d\mu_{\mathcal{V}} \\ &\geq \int_{C_t} \min\{1, \mathcal{R}\} d\mu_{\mathcal{V}}. \end{aligned}$$

Now for each $n \in \mathbb{N}$ let $D^n = \{x \in B \cap \mathcal{S} \mid \mathcal{R}(x) > 2^{-n}\}$ and notice that

$$\begin{aligned} \int_{C_t} \min\{1, \mathcal{R}\} d\mu_{\mathcal{V}} &\geq \int_{D^n \cap C_t} \min\{1, \mathcal{R}\} d\mu_{\mathcal{V}} \\ &\geq 2^{-n} \cdot \mu_{\mathcal{V}}(D^n \cap C_t). \end{aligned}$$

Now combining the above facts we have that $\mathcal{D}_{\mathcal{S}}^d(\phi_t, \phi) \geq 2^{-n} \cdot \mu_{\mathcal{V}}(D^n \cap C_t)$ for any choice of $n \in \mathbb{N}$ so

$$\lim_{t \rightarrow 0} \mu_{\mathcal{V}}(D^n \cap C_t) = 0 \tag{9.1}$$

for all $n \in \mathbb{N}$.

Finally fix $\varepsilon > 0$. Since $\mathcal{R}(x) > 0$ for all $x \in B \cap \mathcal{S}$ we know that the collection $\{D^n\}_{n=1}^{\infty}$ covers $B \cap \mathcal{S}$. Since $B \cap \mathcal{S}$ has finite volume we know there exists some $N \in \mathbb{N}$ such that $\mu_{\mathcal{V}}((B \cap \mathcal{S}) \setminus D^N) < \varepsilon/2$. This implies that for all $t \in (0, 1)$ we have that $\mu_{\mathcal{V}}(C_t \setminus D^N) < \varepsilon/2$. By Equation (9.1) we conclude that we can choose some T such that $t < T$ implies that $\mu_{\mathcal{V}}(C_t \cap D^N) < \varepsilon/2$. Now for $t < T$ we have that $\mu_{\mathcal{V}}(C_t) = \mu_{\mathcal{V}}(C_t \setminus D^N) + \mu_{\mathcal{V}}(C_t \cap D^N) < \varepsilon$. \square

Now we show that any choice of continuous metric on N will produce a weakly equivalent metric on \mathcal{M}^{\sim} .

Lemma 9.4.10. *Suppose that d_1 and d_2 are topologically equivalent metrics on N each induced by a Riemannian metric and let $\{\mathcal{S}_n\}_{n=1}^{\infty}$ be any exhaustion of M by finite volume sets. Then $\mathcal{D}_{\{\mathcal{S}_n\}}^{d_1}$ and $\mathcal{D}_{\{\mathcal{S}_n\}}^{d_2}$ are topologically equivalent metrics on \mathcal{M}^{\sim} .*

Proof. Fix finite volume $\mathcal{S} \subset M$. If we show $\mathcal{D}_{\mathcal{S}}^{d_1}$ and $\mathcal{D}_{\mathcal{S}}^{d_2}$ are topologically equivalent then we have proved the lemma by Proposition 9.4.6. It is sufficient to show that the same families indexed by $(0, 1)$ converge so suppose $\{(\phi_t, B_t)\}_{t \in (0, 1)} \in \mathcal{F}(\mathcal{M})$ and $(\phi_0, B_0) \in \mathcal{M}$ such that $\phi_t \xrightarrow{\mathcal{D}_{\mathcal{S}}^{d_1}} \phi_0$ as $t \rightarrow 0$ and we will show that $\phi_t \xrightarrow{\mathcal{D}_{\mathcal{S}}^{d_2}} \phi_0$ as $t \rightarrow 0$. Fix $\varepsilon > 0$ and without loss of generality assume that

$\varepsilon < \mu_{\mathcal{V}}(\mathcal{S})$. Let

$$C_t^2 = \left\{ x \in B_t \cap B_0 \cap \mathcal{S} \mid d_2(\phi_0(x), \phi_t(x)) > \frac{\varepsilon}{3\mu_{\mathcal{V}}(\mathcal{S})} \right\}.$$

Let $b_{y_0}^i(r) = \{y \in N \mid d_i(y, y_0) < r\}$ for $i = 1, 2$. Since d_1 and d_2 are weakly equivalent metrics for each $y \in N$ there exists some radius $r_y > 0$ such that the ball with respect to d_1 of radius r_y centered at y is a subset of the ball with respect to d_2 of radius $\varepsilon/3\mu_{\mathcal{V}}(\mathcal{S})$ centered at y . Thus there exists some $\mathcal{R} : B_0 \cap \mathcal{S} \rightarrow (0, \infty)$ such that

$$b_{\phi_0(x)}^1(\mathcal{R}(x)) \subset b_{\phi_0(x)}^2\left(\frac{\varepsilon}{3\mu_{\mathcal{V}}(\mathcal{S})}\right) \text{ for all } x \in B_0 \cap \mathcal{S}. \quad (9.2)$$

Define $C_t^1 = \{x \in B_t \cap B_0 \cap \mathcal{S} \mid d_1(\phi_0(x), \phi_t(x)) > \mathcal{R}(x)\}$ and notice that Equation (9.2) implies that $C_t^2 \subset C_t^1$. By Lemma 9.4.9 since $\phi_t \xrightarrow{\mathcal{D}_s^{d_1}} \phi_0$ as $t \rightarrow 0$ we know that $\lim_{t \rightarrow 0} \mu_{\mathcal{V}}(C_t^1) = 0$ and so we can conclude that

$$\lim_{t \rightarrow 0} \mu_{\mathcal{V}}(C_t^2) = 0.$$

Now we can find some $T \in (0, 1)$ such that if $t < T$ then $\mu_{\mathcal{V}}(C_t^2) < \varepsilon/3$ and also $\mu_{\mathcal{V}}((B_t \Delta B_0) \cap \mathcal{S}) < \varepsilon/3$. Then

$$\begin{aligned} \mathcal{D}_s^{d_2}(\phi_t, \phi_0) &= \int_{B_t \cap B_0 \cap \mathcal{S}} \min\{1, d_2(\phi_t, \phi_0)\} d\mu_{\mathcal{V}} + \mu_{\mathcal{V}}((B_t \Delta B_0) \cap \mathcal{S}) \\ &\leq \int_{(B_t \cap B_0 \cap \mathcal{S}) \setminus C_t^2} \min\{1, d_2(\phi_t, \phi_0)\} d\mu_{\mathcal{V}} + \int_{C_t^2} \min\{1, d_2(\phi_t, \phi_0)\} d\mu_{\mathcal{V}} + \mu_{\mathcal{V}}((B_t \Delta B_0) \cap \mathcal{S}) \\ &\leq \int_{\mathcal{S}} \frac{\varepsilon}{3\mu_{\mathcal{V}}(\mathcal{S})} d\mu_{\mathcal{V}} + \mu_{\mathcal{V}}(C_t^2) + \mu_{\mathcal{V}}((B_t \Delta B_0) \cap \mathcal{S}) \\ &< \varepsilon/3 + \varepsilon/3 + \varepsilon/3 = \varepsilon. \end{aligned}$$

□

We conclude this section with the following lemma.

Lemma 9.4.11. *Let $\{\mathcal{S}_n\}_{n=1}^{\infty}$ be a nested exhaustion of M by finite volume sets and suppose that d_1 and d_2 are metrics on N induced by smooth Riemannian metrics. Then $\mathcal{D}_{\{\mathcal{S}_n\}}^{d_1}$ and $\mathcal{D}_{\{\mathcal{S}_n\}}^{d_2}$ are topologically equivalent metrics on \mathcal{M}^{\sim} .*

Proof. Both d_1 and d_2 are continuous with respect to the given topology on N . This means that




Figure 9.3: An image of $\Phi_{m,k}$.

they are topologically equivalent metrics and so by Lemma 9.4.10 the result follows. \square

Remark 9.4.12. If M is finite volume, such as in the case that M is compact, then there is an obvious preferred choice to make when choosing the exhaustion, namely simply $\{M\}$ itself. In such a case we will always use

$$\mathcal{D}_M^d(\phi, \psi) = \int_M p_{\phi\psi}^d d\mu_\nu = \int_{B_\phi \cap B_\psi} \min\{1, d(\phi, \psi)\} d\mu_\nu + \mu_\nu(B_\phi \Delta B_\psi).$$

There are also no choices now when defining convergent ε -perturbations or the radius of convergence except for the choice of metric on N . \circlearrowright

9.4.3 A representative example

To conclude Section 9.4 we work out an important example which will be referenced throughout the chapter.

Example 9.4.13. Let $\Phi_{m,k} : (0, 1) \rightarrow \mathbb{R}$ by

$$\Phi_{m,k}(x) = m \cdot \chi_{(k/m, k+1/m)}(x)$$

(shown in Figure 9.3) for $k, m \in \mathbb{N}$ with $k < m$ where χ_S is the indicator function for the set $S \subset (0, 1)$.

We can see that

$$\int_{(0,1)} \Phi_{m,k} = 1$$

for all possible values of k and m . We will use these functions to construct an example which is similar to the “traveling wave” example that is common in introductory analysis [32] except that our example changes height so it always integrates to 1.

Consider the sequence

$$\phi_1 = \Phi_{0,1}, \phi_2 = \Phi_{0,2}, \phi_3 = \Phi_{1,2}, \phi_4 = \Phi_{0,3}, \phi_5 = \Phi_{1,3}, \phi_6 = \Phi_{2,3}, \phi_7 = \Phi_{0,4}, \dots$$

(as shown in Figure 9.4) and let $\phi_0: (0,1) \rightarrow \mathbb{R}$ by

$$\phi_0(x) = 0 \text{ for all } x \in (0,1).$$

Notice that this sequence does not converge pointwise to ϕ_0 for any point $x \in (0,1)$. Also notice




Figure 9.4: A few terms of $\{\phi_n\}$. It can be seen that each integrates to 1 and the “traveling waves” pass over every point infinitely many times, so pointwise convergence is impossible.

that since the integral of any element in this sequence is 1 we can conclude that this sequence does not converge in L^1 (or L^p for any $p \in [1, \infty)$) either (as is mentioned in Remark 9.3.3), but it will converge with respect to \mathcal{D} . This is because the measure of values in the domain which get sent to a number other than zero is becoming arbitrarily small, so we can conclude that

$$\lim_{n \rightarrow \infty} \mathcal{D}_{(0,1)}^{d_{\mathbb{R}}}(\phi_n, \phi_0) = 0.$$

This example shows a case in which we have a family which does not behave well pointwise almost everywhere or with respect to the L^p norm, but it does behave well with respect to \mathcal{D} .

Of course, if we replace the indicator function with a bump function we can produce a sequence of smooth functions which has the same essential properties as these functions. In fact, for this example we have considered a sequence of functions instead of a continuous family of functions because it made it easier to describe the sequence, but we could easily extend this sequence to a smooth (see Definition 9.7.1) family of smooth embeddings of $(0, 1)$ into $(0, 1) \times \mathbb{R}$ indexed by $t \in (0, 1)$ which has the same properties. \circlearrowright

9.5 Completeness of \mathcal{M}^\sim

9.5.1 Preparation

Below is a collection of various technical Lemmas which are needed for Section 9.5.2. In this section we will frequently use the alternative expression for \mathcal{D} given in Proposition 9.4.4.

Lemma 9.5.1. *Let $A \subset M$ be a measurable finite volume set and let $a, b, c \in \mathbb{R}$ such that $a < b$ and $c \in [a, b]$. Suppose that some family of measurable functions $\{f_t : A \rightarrow \mathbb{R}^k\}_{t \in (a, b)}$, is Cauchy with respect to $\int_A \|f_t - f_s\| \, d\mu_\nu$ as $t \rightarrow c$. Then there exists some $f : A \rightarrow \mathbb{R}^k$ such that $f_t \xrightarrow{\mathcal{D}_A^{d_{\mathbb{R}^k}}} f$ as $t \rightarrow c$ where $d_{\mathbb{R}^k}$ is the usual metric on \mathbb{R}^k .*

Proof. It is sufficient to show the result in the case that $a = c = 0$ and $b = 1$. For $t \in (0, 1)$ we know that f_t maps into \mathbb{R}^k so we may write it into components. Write

$$f_t(x) = (f_t^1(x), f_t^2(x), \dots, f_t^k(x)).$$

Notice for any fixed $j \in \{1, 2, \dots, k\}$ that

$$\begin{aligned} \int_A \|f_t - f_s\| \, d\mu_\nu &= \int_A \left(\sum_{i=1}^k (f_t^i - f_s^i)^2 \right)^{1/2} \, d\mu_\nu \\ &\geq \int_A |f_t^j - f_s^j| \, d\mu_\nu \end{aligned}$$

so we conclude that $\{f_t^j\}_{t \in (0,1)}$ is Cauchy in $L^1(A)$ for each $j \in \{1, 2, \dots, k\}$. Since $L^1(A)$ is complete we know for each $j \in \{1, 2, \dots, k\}$ there exists some function $f^j : A \rightarrow \mathbb{R}$ such that

$$\lim_{t \rightarrow 0} \int_A |f_t^j - f^j| d\mu_\nu = 0.$$

So we define $f(x) = (f^1(x), f^2(x), \dots, f^k(x))$ for $x \in A$. Now, notice that for any $x \in A$ we have

$$\|f_t(x) - f(x)\| \leq \sum_{i=1}^k |f_t^i(x) - f^i(x)|.$$

Finally, notice

$$\begin{aligned} \mathcal{D}_A^{d_{\mathbb{R}^k}}(f_t, f) &= \int_A \min\{1, \|f_t - f\|\} d\mu_\nu \\ &\leq \int_A \min\{1, \sum_{i=1}^k |f_t^i - f^i|\} d\mu_\nu \\ &\leq \sum_{i=1}^k \int_A \min\{1, |f_t^i - f^i|\} d\mu_\nu \\ &\leq \sum_{i=1}^k \int_A |f_t^i - f^i| d\mu_\nu. \end{aligned}$$

Since $\int_A |f_t^j - f^j| d\mu_\nu$ goes to 0 as t goes to 0 for any choice of $j \in \{1, 2, \dots, k\}$ the result follows. \square

Lemma 9.5.2. *Let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Let $\{(\phi_t, B_t)\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M}^\sim(M, \mathbb{R}^k))$ and $(\phi, B) \in \mathcal{M}^\sim(M, \mathbb{R}^k)$ be such that $\phi_t \xrightarrow{\mathcal{D}} \phi$ as $t \rightarrow c$ and suppose there exists a fixed closed subset $P \subset \mathbb{R}^k$ such that $\phi_t(B_t) \subset P$ for all $t \in (a, b)$. Then*

$$\mu_\nu(\{x \in B \mid \phi(x) \notin P\}) = 0$$

and thus there exists some $(\phi', B') \sim (\phi, B)$ such that $\phi'(B) \subset P$.

Proof. Without loss of generality assume that $a = c = 0$ and $b = 1$. Since P is closed notice that for $y \in \mathbb{R}^k$ we have that

$$\inf_{p \in P} \{d_{\mathbb{R}^k}(y, p)\} = 0 \text{ implies } y \in P$$

where $d_{\mathbb{R}^k}$ is the standard metric on \mathbb{R}^k . Thus, if we let $C = \{x \in B \mid \phi(x) \notin P\}$ and

$$C_n = \{x \in B \mid \inf_{p \in P} \{d_{\mathbb{R}^k}(\phi(x), p)\} > 2^{-n}\}$$

for each $n \in \mathbb{N}$ then we have that

$$C = \bigcup_{n=1}^{\infty} C_n.$$

So it will be sufficient to prove that $\mu_{\mathcal{V}}(C_n) = 0$ for each $n \in \mathbb{N}$.

Let $\mathcal{S} \subset M$ be compact and notice that $\phi_t \xrightarrow{\mathcal{D}} \phi$ as $t \rightarrow 0$ implies that

$$\lim_{t \rightarrow 0} \mathcal{D}_{C_n \cap \mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_t, \phi) = 0.$$

We know

$$\begin{aligned} \mathcal{D}_{C_n \cap \mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_t, \phi) &= \int_{B_t \cap C_n \cap \mathcal{S}} \min\{1, d_{\mathbb{R}^k}(\phi_t, \phi)\} d\mu_{\mathcal{V}} + \mu_{\mathcal{V}}((B_t \Delta B) \cap C_n \cap \mathcal{S}) \\ &> 2^{-n} \cdot \mu_{\mathcal{V}}(B_t \cap C_n \cap \mathcal{S}) + \mu_{\mathcal{V}}((C_n \setminus B_t) \cap \mathcal{S}) \\ &\geq 2^{-n} \cdot \mu_{\mathcal{V}}(C_n \cap \mathcal{S}) \geq 0. \end{aligned}$$

This implies that

$$\lim_{t \rightarrow 0} (2^{-n} \cdot \mu_{\mathcal{V}}(C_n \cap \mathcal{S})) = 0$$

for any choice of compact $\mathcal{S} \subset M$ which of course means $\mu_{\mathcal{V}}(C_n) = 0$ for each $n \in \mathbb{N}$. \square

Lemma 9.5.3. *Suppose that $\rho : N \rightarrow \mathbb{R}^k$ is an isometric embedding of Riemannian manifolds (ie, it preserves the metric tensor) where \mathbb{R}^k is equipped with the standard Riemannian metric and $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Then given some family $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ and $\phi \in \mathcal{M}$ we have that $\phi_t \xrightarrow{\mathcal{D}} \phi$ as $t \rightarrow c$ if and only if $(\rho \circ \phi_t) \xrightarrow{\mathcal{D}} (\rho \circ \phi)$ as $t \rightarrow c$.*

Proof. Let d_N be the natural distance function on N and let $d_{\mathbb{R}^k}$ be the standard distance on \mathbb{R}^k .

Then we may define a second distance function d_2 on N by

$$d_2 : N \times N \rightarrow \mathbb{R}$$

$$(y_1, y_2) \mapsto d_{\mathbb{R}^k}(\rho(y_1), \rho(y_2))$$

If we can show that these are topologically equivalent metrics on N then the result will follow by Lemma 9.4.10. Fix some $y_0 \in N$ and let

$$b(r) = \{y \in N \mid d_N(y, y_0) < r\} \text{ and } b^2(r) = \{y \in N \mid d_2(y, y_0) < r\}.$$

Now notice that in general $d_2 \leq d_N$ (see Remark 9.5.4), so we must only show that given some arbitrary $R > 0$ we can find some $r > 0$ such that $b^2(r) \subset b(R)$.

Since $b(R) \subset N$ is an open set and ρ is an embedding we can find some open set $U \subset \mathbb{R}^k$ such that $U \cap \rho(N) = \rho(b(R))$. Now since U is open and $\rho(y_0) \in U$ we can find some $r > 0$ such that

$$\{z \in \mathbb{R}^k \mid d_{\mathbb{R}^k}(z, \rho(y_0)) < r\} \subset U. \tag{9.3}$$

Now let $y \in b^2(r)$. Then we can see that Equation (9.3) tells us that $\rho(y) \in U$. Clearly $\rho(y) \in \rho(N)$ so $\rho(y) \in U \cap \rho(N) = \rho(b(R))$. Since ρ is injective we now know that $y \in b(R)$. Thus $b^2(r) \subset b(R)$.

□

Remark 9.5.4. An isometric embedding of Riemannian manifolds preserves the metric at each point, so it will preserve the length of curves, but often the shortest path between two points in $\rho(N) \subset \mathbb{R}^k$ (a straight line) is not contained in $\rho(N)$. This means that even though ρ preserves the metric the images of two points in \mathbb{R}^k may be closer than those two points are in N and this is why $d_2 \leq d_N$ in the proof above. ◊

9.5.2 Proof that $(\mathcal{M}^\sim, \mathcal{D})$ is complete

The goal of this section is to prove that $(\mathcal{M}^\sim, \mathcal{D}_{\{\mathcal{S}_n\}}^d)$ is a complete metric space for any choice of an exhaustion of M by finite volume sets $\{\mathcal{S}_n\}_{n=1}^\infty$ and metric on N induced by a complete Riemannian metric d . To do this we first have to prove several lemmas. We will start by considering mappings restricted to a compact set and indexed by $(0, 1)$, but later it will be easy to generalize this to all of M by using a compact exhaustion and to arbitrary intervals. The first lemma proves the theorem in the special case that $N = \mathbb{R}^k$ and all maps have the same domain.

Lemma 9.5.5. *Fix some compact set $\mathcal{S} \subset M$ and let $\{(\phi_t, \mathcal{S})\}_{t \in (0,1)} \in \mathcal{F}(\mathcal{M}(M, \mathbb{R}^k))$ be a family which is Cauchy with respect to $\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}$ as $t \rightarrow 0$. Then there exists some $\phi_0 : \mathcal{S} \rightarrow \mathbb{R}^k$, unique up to \sim , such that $\phi_t \xrightarrow{\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}} \phi_0$ as $t \rightarrow 0$.*

Proof. The proof has five steps. Figures 9.5 and 9.6 show how the proof works in a specific case.

Step 1: First we will define a new family $\{(\overline{\phi}_t^n, \mathcal{S})\}_{t \in (0,1)} \in \mathcal{F}(\mathcal{M})$ for each $n \in \mathbb{N}$. Since $\{(\phi_t, \mathcal{S})\}_{t \in (0,1)}$ is Cauchy with respect to $\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}$ for each $n \in \mathbb{N}$ pick some $T_n \in (0, 1)$ such that

$$t \leq T_n \implies \mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_t, \phi_{T_n}) < 2^{-n}. \quad (9.4)$$

Now for each $n \in \mathbb{N}$ we can define a new family $\{(\overline{\phi}_t^n, \mathcal{S})\}_{t \in (0, T_n)}$ by

$$\overline{\phi}_t^n(x) = \begin{cases} \phi_t(x) & \text{if } \|\phi_t(x) - \phi_{T_n}(x)\| \leq 1/2; \\ \phi_{T_n}(x) + \frac{\phi_t(x) - \phi_{T_n}(x)}{2\|\phi_t(x) - \phi_{T_n}(x)\|} & \text{otherwise.} \end{cases}$$

Step 2: Next we will show that each family $\{(\overline{\phi}_t^n, \mathcal{S})\}_{t \in (0, T_n)}$ converges in $\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}$. Notice for any $t, s < T_n$ we have that $\|\overline{\phi}_t^n(x) - \overline{\phi}_s^n(x)\| \leq 1$ so in fact we have that

$$\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\overline{\phi}_t^n, \overline{\phi}_s^n) = \int_{\mathcal{S}} \|\overline{\phi}_t^n - \overline{\phi}_s^n\| \, d\mu_{\mathcal{V}}.$$

¹Recall that the functions in Example 9.4.13 are labeled in the opposite order for convenience.




Figure 9.5: Applying the proof of Lemma 9.5.5 to Example 9.4.13. In Step 1 we choose $T_2 = 12$ because¹ in this case ϕ_{12} satisfies Equation (9.4) for $n = 2$ and we restrict each mapping to have values within the shaded area (within a distance of $1/2$ from ϕ_{T_2}) to produce the family $\{(\overline{\phi_t^n}, \mathcal{S})\}$. In Step 2 we find the limit of those functions to define $\overline{\phi_0^2}$. At the points in which this function takes values on the boundary of the shaded area we can see that the family is approaching a value outside of the shaded area, so in Step 3 we remove these points from the domain to form ϕ_0^2 .




Figure 9.6: Two examples in which the maps are restricted to find the limit $\overline{\phi_0^2}$ in Steps 1 and 2 of the proof of Lemma 9.5.5. In each case we start with ϕ_t and create $\overline{\phi_t}$ by changing the function to have only values with a distance less than $1/2$ to ϕ_{T_2} .

Since for any $x \in \mathcal{S}$ we have that $\|\phi_t(x) - \phi_s(x)\| \geq \|\overline{\phi_t^n}(x) - \overline{\phi_s^n}(x)\|$ we know

$$\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_t, \phi_s) \geq \mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\overline{\phi_t^n}, \overline{\phi_s^n}) = \int_{\mathcal{S}} \|\overline{\phi_t^n} - \overline{\phi_s^n}\| \, d\mu_{\mathcal{V}}$$

and since $\{(\phi_t, \mathcal{S})\}_{t \in (0,1)}$ is Cauchy with respect to $\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}$ we now know that $\{(\overline{\phi_t^n}, \mathcal{S})\}_{t \in (0, T_n)}$ is Cauchy with respect to $\int_{\mathcal{S}} \|\overline{\phi_t^n} - \overline{\phi_s^n}\| \, d\mu_{\mathcal{V}}$. Thus by Lemma 9.5.1 we know that for each $n \in \mathbb{N}$ there exists a map $\overline{\phi_0^n} : \mathcal{S} \rightarrow \mathbb{R}^k$ such that $\overline{\phi_t^n} \xrightarrow{\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}} \overline{\phi_0^n}$ as $t \rightarrow 0$.

Step 3: In this step we will define ϕ_0^n for each n on all but a subset of measure less than 2^{-n+2} of \mathcal{S} . Let

$$B_0^n = \{x \in \mathcal{S} \mid \|\phi_{T_n}(x) - \overline{\phi_0^n}(x)\| < 1/4\}.$$

Now define

$$\phi_0^n = \overline{\phi_0^n}|_{B_0^n} : B_0^n \rightarrow \mathbb{R}^k.$$

Now we will show that ϕ_0^n is defined on all but a small subset of \mathcal{S} .

Let $\varepsilon > 0$ and pick some $t < T_n$ such that $\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\overline{\phi_t^n}, \overline{\phi_0^n}) < \varepsilon$. Then

$$\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_{T_n}, \overline{\phi_0^n}) \leq \mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_{T_n}, \overline{\phi_t^n}) + \mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\overline{\phi_t^n}, \overline{\phi_0^n}) < 2^{-n} + \varepsilon$$

for all $\varepsilon > 0$ so we may conclude that $\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_{T_n}, \overline{\phi_0^n}) \leq 2^{-n}$. Next also notice that since $\mathcal{S} \setminus B_0^n \subset \mathcal{S}$ we know that

$$\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\phi_{T_n}, \overline{\phi_0^n}) \geq \int_{\mathcal{S} \setminus B_0^n} \min\{1, \|\phi_{T_n} - \overline{\phi_0^n}\|\} \, d\mu_{\mathcal{V}} \geq \frac{1}{4} \mu_{\mathcal{V}}(\mathcal{S} \setminus B_0^n).$$

This means that $\mu_{\mathcal{V}}(\mathcal{S} \setminus B_0^n) \leq 2^{-n+2}$. Since $B_0^n \subset \mathcal{S}$ we conclude that $\mu_{\mathcal{V}}(B_0^n) \geq \mu_{\mathcal{V}}(\mathcal{S}) - 2^{-n+2}$.

So if

$$\mu_{\mathcal{V}}\left(\mathcal{S} \setminus \bigcup_{n=1}^{\infty} B_0^n\right) = \alpha > 0$$

then we would have a contradiction because we can choose some $n \in \mathbb{N}$ such that $2^{-n+2} < \alpha$. Thus we have that

$$\mu_{\mathcal{V}}\left(\mathcal{S} \setminus \bigcup_{n=1}^{\infty} B_0^n\right) = 0.$$

Step 4: Next we must show that the limiting functions are equal on the overlap of their

domains. That is, we must show for any $m, n \in \mathbb{N}$ that $\phi_0^m(x) = \phi_0^n(x)$ for almost every $x \in B_0^m \cap B_0^n$.

Our first step towards this goal is to define

$$E_t^n = \{x \in B_0^n \mid \|\phi_{T_n}(x) - \overline{\phi_t^n}(x)\| \geq 1/2\}$$

and show that

$$\lim_{t \rightarrow 0} \mu_{\mathcal{V}}(E_t^n) = 0$$

for all $n \in \mathbb{N}$.

Since $E_t^n \subset B_0^n$ we know that for any $x \in E_t^n$ we have that

$$\|\phi_{T_n}(x) - \overline{\phi_t^n}(x)\| \geq 1/2$$

and also that

$$\|\phi_{T_n}(x) - \overline{\phi_0^n}(x)\| < 1/4.$$

Thus we may apply the triangle inequality to notice that

$$\|\overline{\phi_t^n}(x) - \overline{\phi_0^n}(x)\| \geq 1/4 \text{ for } x \in E_t^n.$$

Now we just notice that since $E_t^n \subset \mathcal{S}$ we have

$$\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}(\overline{\phi_t^n}, \overline{\phi_0^n}) \geq \int_{E_t^n} \|\overline{\phi_t^n} - \overline{\phi_0^n}\| \, d\mu_{\mathcal{V}} \geq \frac{1}{4} \mu_{\mathcal{V}}(E_t^n).$$

Thus we conclude that $\lim_{t \rightarrow 0} \mu_{\mathcal{V}}(E_t^n) = 0$, as desired.

Now let $C = \{x \in B^n \cap B^m \mid \phi_0^n(x) \neq \phi_0^m(x)\}$ and we will show that $\mu_{\mathcal{V}}(C) = 0$ to complete this step. Notice that for any $x \in C \setminus (E_t^n \cup E_t^m)$ we have that $\overline{\phi_t^n}(x) = \overline{\phi_t^m}(x) = \phi_t(x)$. Notice

$$\begin{aligned} \int_C \min\{1, \|\overline{\phi_t^n} - \overline{\phi_t^m}\|\} \, d\mu_{\mathcal{V}} &\leq \mu_{\mathcal{V}}(E_t^n) + \mu_{\mathcal{V}}(E_t^m) + \int_{C \setminus (E_t^n \cup E_t^m)} \min\{1, \|\overline{\phi_t^n} - \overline{\phi_t^m}\|\} \, d\mu_{\mathcal{V}} \\ &= \mu_{\mathcal{V}}(E_t^n) + \mu_{\mathcal{V}}(E_t^m). \end{aligned}$$

Then, since $0 < \int_C \min\{1, \|\overline{\phi}_t^n - \overline{\phi}_t^m\|\} d\mu_\nu \leq \mu_\nu(E_t^n) + \mu_\nu(E_t^m)$ and the right side decreases to zero as $t \rightarrow 0$ we conclude that

$$\lim_{t \rightarrow 0} \int_C \min\{1, \|\overline{\phi}_t^n - \overline{\phi}_t^m\|\} d\mu_\nu = 0$$

so $\lim_{t \rightarrow 0} \mathcal{D}_C^{d_{\mathbb{R}^k}}(\overline{\phi}_t^n, \overline{\phi}_t^m) = 0$.

Finally, by the triangle inequality

$$\mathcal{D}_C^{d_{\mathbb{R}^k}}(\phi_0^n, \phi_0^m) \leq \mathcal{D}_C^{d_{\mathbb{R}^k}}(\phi_0^n, \overline{\phi}_t^n) + \mathcal{D}_C^{d_{\mathbb{R}^k}}(\overline{\phi}_t^n, \overline{\phi}_t^m) + \mathcal{D}_C^{d_{\mathbb{R}^k}}(\overline{\phi}_t^m, \phi_0^m)$$

and we know each term on the right goes to zero as $t \rightarrow 0$. Since t does not appear on the left side we may conclude that

$$\mathcal{D}_C^{d_{\mathbb{R}^k}}(\phi_0^n, \phi_0^m) = \int_C \min\{1, \|\phi_0^n - \phi_0^m\|\} d\mu_\nu = 0.$$

Notice that the function $\min\{1, \|\phi_0^n(x) - \phi_0^m(x)\|\}$ is strictly positive on C , so since integrating it over C yields zero we conclude that $\mu_\nu(C) = 0$.

Step 5: In this step we will define the map $\phi_0 : \mathcal{S} \rightarrow \mathbb{R}^k$ and show that it is the unique limit.

Now define

$$\phi_0(x) = \phi_0^n(x) \text{ for any } n \text{ such that } x \in B_0^n.$$

This map is well defined almost everywhere because the ϕ_0^n are equal almost everywhere on the overlap of their domains and $\cup_{n=1}^\infty B_0^n$ covers almost all of \mathcal{S} .

Now we must show this is the limit. Since we already know that $\{(\phi_t, \mathcal{S})\}_{t \in (0,1)}$ is Cauchy it is sufficient to choose a subsequence and show it converges to ϕ_0 . We will consider the sequence $\{(\phi_{T_n}, \mathcal{S})\}_{n=1}^\infty$. Fix some $\varepsilon > 0$ and pick $N \in \mathbb{N}$ such that $2^{-n+2} < \varepsilon/3$ for all $n > N$. Now for each

$n > N$ pick $t_n \in (0, T_n)$ such that $\mathcal{D}_S^{d_{\mathbb{R}^k}}(\overline{\phi_{t_n}^n}, \overline{\phi_0^n}) < \varepsilon/3$. Then for any $n > N$ we have

$$\begin{aligned} \mathcal{D}_S^{d_{\mathbb{R}^k}}(\phi_{T_n}, \phi_0) &\leq \mathcal{D}_S^{d_{\mathbb{R}^k}}(\phi_{T_n}, \overline{\phi_{t_n}^n}) + \mathcal{D}_S^{d_{\mathbb{R}^k}}(\overline{\phi_{t_n}^n}, \overline{\phi_0^n}) + \mathcal{D}_S^{d_{\mathbb{R}^k}}(\overline{\phi_0^n}, \phi_0) \\ &< 2^{-n} + \varepsilon/3 + 2^{-n+2} \\ &< \varepsilon. \end{aligned}$$

Thus we conclude that $\phi_{T_n} \xrightarrow{\mathcal{D}_S^{d_{\mathbb{R}^k}}} \phi_0$ as $t \rightarrow 0$ and thus $\phi_t \xrightarrow{\mathcal{D}_S^{d_{\mathbb{R}^k}}} \phi_0$ as $t \rightarrow 0$. To show that this is unique suppose that there exists some other $\phi'_0 : \mathcal{S} \rightarrow \mathbb{R}^k$ such that $\phi_t \xrightarrow{\mathcal{D}_S^{d_{\mathbb{R}^k}}} \phi'_0$ as $t \rightarrow 0$. Then for any compact set $S' \subset M$ and $t \in (0, 1)$ we have that

$$\mathcal{D}_S^{d_{\mathbb{R}^k}}(\phi_0, \phi'_0) \leq \mathcal{D}_S^{d_{\mathbb{R}^k}}(\phi_0, \phi_t) + \mathcal{D}_S^{d_{\mathbb{R}^k}}(\phi_t, \phi'_0) \rightarrow 0 \text{ as } t \rightarrow 0$$

since both have domain \mathcal{S} , so $\phi_0 \sim \phi'_0$.

□

For the next step we will continue to focus on a single compact set and the case in which $N = \mathbb{R}^k$, but this time we will allow the domains of the functions to vary.

Lemma 9.5.6. *Fix some compact subset $\mathcal{S} \subset M$. Any family $\{(\phi_t, B_t)\}_{t \in (0,1)} \in \mathcal{F}(\mathcal{M}(M, \mathbb{R}^k))$ which is Cauchy with respect to $\mathcal{D}_S^{d_{\mathbb{R}^k}}$ as $t \rightarrow 0$ also converges with respect to $\mathcal{D}_S^{d_{\mathbb{R}^k}}$ as $t \rightarrow 0$ to some $\phi_0 : B_0 \rightarrow \mathbb{R}^k$ where $B_0 \subset \mathcal{S}$. Moreover, among maps in \mathcal{M} with domains a subset of \mathcal{S} that share this property, (ϕ_0, B_0) is unique up to \sim .*

Proof. Let $\{(\phi_t, B_t)\}_{t \in (0,1)} \in \mathcal{F}(\mathcal{M}^\sim(M, \mathbb{R}^k))$ be a family which is Cauchy as $t \rightarrow 0$ and define $\pi : \mathbb{R}^k \rightarrow \mathbb{R}^{k+1}$ via $\pi(x_1, \dots, x_k) = (0, x_1, \dots, x_k)$. Now, for each $t \in (0, 1)$ define $\widehat{\phi}_t : \mathcal{S} \rightarrow \mathbb{R}^{k+1}$ by

$$\widehat{\phi}_t(x) = \begin{cases} (1, 0, \dots, 0) & \text{if } x \notin B_t \\ \pi(\phi_t(x)) & \text{if } x \in B_t. \end{cases}$$

Notice that for all $s, t \in (0, 1)$

$$\min\{1, \|\widehat{\phi}_t(x) - \widehat{\phi}_s(x)\|\} = \begin{cases} \min\{1, \|\phi_t(x) - \phi_s(x)\|\} & \text{if } x \in B_t \cap B_s \\ 1 & \text{if } x \in B_t \Delta B_s \\ 0 & \text{if } x \notin B_t \cup B_s. \end{cases}$$

Thus for $s, t \in (0, 1)$ we have that

$$\begin{aligned} \mathcal{D}_S^{d_{\mathbb{R}^{k+1}}}(\widehat{\phi}_t, \widehat{\phi}_s) &= \int_S \min\{1, \|\widehat{\phi}_t - \widehat{\phi}_s\|\} d\mu_V \\ &= \int_{B_t \cap B_s \cap S} \min\{1, \|\widehat{\phi}_t - \widehat{\phi}_s\|\} d\mu_V + \int_{(B_t \Delta B_s) \cap S} 1 d\mu_V + \int_{(B_t \cup B_s) \setminus S} 0 d\mu_V \\ &= \int_{B_t \cap B_s \cap S} \min\{1, \|\widehat{\phi}_t - \widehat{\phi}_s\|\} d\mu_V + \mu_V((B_t \Delta B_s) \cap S) \\ &= \mathcal{D}_S^{d_{\mathbb{R}^k}}(\phi_t, \phi_s). \end{aligned}$$

Now we can see that $\{\widehat{\phi}_t, S\}_{t \in (0, 1)}$ must be Cauchy as well. Since these are all functions into \mathbb{R}^{k+1} with the same domain we can invoke Lemma 9.5.5 to conclude that there exists some limit $\widehat{\phi}_0 : S \rightarrow \mathbb{R}^{k+1}$ which is unique up to \sim such that $\widehat{\phi}_t \xrightarrow{\mathcal{D}_S^{d_{\mathbb{R}^{k+1}}}} \widehat{\phi}_0$ as $t \rightarrow 0$. Since $K = \pi(\mathbb{R}^k) \cup \{(1, 0, \dots, 0)\}$ is a closed subset of \mathbb{R}^{k+1} we can invoke Lemma 9.5.2 to conclude that we may assume that $\widehat{\phi}_0(S) \subset K$.

This allows us to define (ϕ_0, B_0) in the following way. Let

$$B_0 = \{x \in S \mid \widehat{\phi}_0(x) \neq (1, 0, \dots, 0)\}.$$

So for any $x \in B_0$ we know that $\widehat{\phi}_0(x) \in \pi(\mathbb{R}^k)$, which means that we can define

$$\begin{aligned} \phi_0 : B_0 &\rightarrow \mathbb{R}^k \\ x &\mapsto \pi^{-1}(\widehat{\phi}_0(x)). \end{aligned}$$

Since $\mathcal{D}_S^{d_{\mathbb{R}^k}}(\phi_t, \phi_0) = \mathcal{D}_S^{d_{\mathbb{R}^{k+1}}}(\widehat{\phi}_t, \widehat{\phi}_0)$ we can see that $\widehat{\phi}_t \xrightarrow{\mathcal{D}_S^{d_{\mathbb{R}^{k+1}}}} \widehat{\phi}_0$ implies that $\phi_t \xrightarrow{\mathcal{D}_S^{d_{\mathbb{R}^k}}} \phi_0$ and we know that ϕ_0 is unique up to \sim because $\widehat{\phi}_0$ is.

□

Finally we will expand to consider all of M instead of just a single compact set, but we will still only consider $N = \mathbb{R}^k$.

Lemma 9.5.7. $(\mathcal{M}^\sim(M, \mathbb{R}^k), \mathcal{D})$ is complete. That is, if $\{\mathcal{S}_n\}_{n=1}^\infty$ is any nested finite exhaustion of M then $(\mathcal{M}^\sim(M, \mathbb{R}^k), \mathcal{D}_{\{\mathcal{S}_n\}}^{d_{\mathbb{R}^k}})$ is a complete metric space where $d_{\mathbb{R}^k}$ is the standard metric on \mathbb{R}^k .

Proof. It is sufficient to show that families indexed by $(0, 1)$ which are Cauchy as $t \rightarrow 0$ also converge as $t \rightarrow 0$. Let $\{(\phi_t, B_t)\}_{t \in (0, 1)} \in \mathcal{F}(\mathcal{M}^\sim(M, \mathbb{R}^k))$ be Cauchy with respect to $\mathcal{D}_{\{\mathcal{S}_n\}}^{d_{\mathbb{R}^k}}$. From Proposition 9.4.6 we know that this means for all compact $\mathcal{S} \subset M$ this sequence is Cauchy with respect to $\mathcal{D}_{\mathcal{S}}^{d_{\mathbb{R}^k}}$ and from Lemma 9.5.6 we know that this means for each compact $\mathcal{S} \subset M$ we have some $\phi_0^{\mathcal{S}} : B_0^{\mathcal{S}} \rightarrow \mathbb{R}^k$ where $B_0^{\mathcal{S}} \subset \mathcal{S}$ such that $(\phi_0^{\mathcal{S}}, B_0^{\mathcal{S}})$ is unique up to \sim . Let $\{\mathcal{S}_n\}_{n=1}^\infty$ be a nested compact exhaustion of M and now we would like to conclude that for $n < m$ we have that

$$(\phi_0^{\mathcal{S}_m}|_{\mathcal{S}_n}, B_0^{\mathcal{S}_m} \cap \mathcal{S}_n) \sim (\phi_0^{\mathcal{S}_n}, B_0^{\mathcal{S}_n}).$$

Notice that

$$\mathcal{D}_{\mathcal{S}_n}^{d_{\mathbb{R}^k}}(\phi_0^{\mathcal{S}_m}|_{\mathcal{S}_n}, \phi_t) = \mathcal{D}_{\mathcal{S}_n}^{d_{\mathbb{R}^k}}(\phi_0^{\mathcal{S}_m}, \phi_t) \leq \mathcal{D}_{\mathcal{S}_m}^{d_{\mathbb{R}^k}}(\phi_0^{\mathcal{S}_m}, \phi_t).$$

because $\mathcal{S}_n \subset \mathcal{S}_m$. Since $\mathcal{D}_{\mathcal{S}_m}^{d_{\mathbb{R}^k}}(\phi_0^{\mathcal{S}_m}, \phi_t) \rightarrow 0$ as $t \rightarrow 0$ we know that $\mathcal{D}_{\mathcal{S}_n}^{d_{\mathbb{R}^k}}(\phi_0^{\mathcal{S}_m}|_{\mathcal{S}_n}, \phi_t) \rightarrow 0$ as $t \rightarrow 0$. From Lemma 9.5.6 we know that such a limit with domain a subset of \mathcal{S}_n is unique up to \sim . Thus we conclude that $\phi_0^{\mathcal{S}_m}|_{\mathcal{S}_n} \sim \phi_0^{\mathcal{S}_n}$. This means that the symmetric difference of their domains has zero volume, so $\mu_V((B_0^{\mathcal{S}_n} \Delta B_0^{\mathcal{S}_m}) \cap \mathcal{S}_n) = 0$, and also they are equal almost everywhere on the overlap of their domains. So now we can define $B_0 = \bigcup_{n=1}^\infty B_0^{\mathcal{S}_n}$ and $\phi_0 : B_0 \rightarrow \mathbb{R}^k$ almost everywhere by

$$\phi_0(x) = \phi_0^{\mathcal{S}_n}(x) \text{ where } x \in \mathcal{S}_n$$

and this is well defined. Since $\phi_t \xrightarrow{\mathcal{D}_{\mathcal{S}_n}^{d_{\mathbb{R}^k}}} \phi_0$ as $t \rightarrow 0$ for all \mathcal{S}_n in a compact exhaustion of M we know by definition that $\phi_t \xrightarrow{\mathcal{D}} \phi_0$ as $t \rightarrow 0$. \square

Now we are ready to prove that $(\mathcal{M}^\sim, \mathcal{D})$ is complete.

Lemma 9.5.8. Suppose that $\{\mathcal{S}_n\}_{n=1}^\infty$ is a nested exhaustion of M by finite measure sets and that d is a metric on N induced by a Riemannian metric. Then $(\mathcal{M}^\sim, \mathcal{D}_{\{\mathcal{S}_n\}}^d)$ is complete if and only if (N, d) is complete.




Figure 9.7: $\{\check{\phi}_t\}_{t \in (0,1)}$ is a family of maps into \mathbb{R}^k .

Proof. It is sufficient to show that Cauchy families indexed by $(0, 1)$ converge as $t \rightarrow 0$. First assume that (N, d) is complete and let $\{(\phi_t, B_t)\}_{t \in (0,1)} \in \mathcal{F}(M)$ be Cauchy as $t \rightarrow 0$. Now, by the Nash embedding theorem [59, Theorem 3] we know there exists an isometric embedding $\rho : N \rightarrow \mathbb{R}^k$ for some $k \in \mathbb{N}$. In fact, since N is a complete Riemannian manifold we can choose ρ to have a closed image [58, Theorem 0.2].

Let d_N denote the distance on N induced by the metric and let $d_{\mathbb{R}^k}$ denote the standard distance on \mathbb{R}^k . Notice for $y_1, y_2 \in N$ we have that

$$d_{\mathbb{R}^k}(\rho(y_1), \rho(y_2)) \leq d_N(y_1, y_2). \quad (9.5)$$

(See Remark 9.5.4). Define

$$\check{\phi}_t := \rho \circ \phi_t : B_t \rightarrow \mathbb{R}^k$$

as is shown in Figure 9.7. From Equation (9.5) above we know that

$$\mathcal{D}_S^{d_{\mathbb{R}^k}}(\check{\phi}_t, \check{\phi}_s) \leq \mathcal{D}_S^d(\phi_t, \phi_s)$$

for all compact $S \subset M$ so we can conclude that $\{(\check{\phi}_t, B_t)\}_{t \in (0,1)}$ is also Cauchy with respect to \mathcal{D} . By Lemma 9.5.7 we know that there exists some $\check{\phi}_0 : B_0 \rightarrow \mathbb{R}^k$ such that $\check{\phi}_t \xrightarrow{\mathcal{D}} \check{\phi}_0$ as $t \rightarrow 0$ and by

Lemma 9.5.2 we can conclude, up to measure zero corrections, that

$$\check{\phi}_0(B_0) \subset \rho(N).$$

Thus we may define

$$\phi_0 := \rho^{-1} \circ \check{\phi}_0 : B_0 \rightarrow N.$$

By Lemma 9.5.3 we know that $\check{\phi}_t \xrightarrow{\mathcal{D}} \check{\phi}_0$ as $t \rightarrow 0$ implies that $\phi_t \xrightarrow{\mathcal{D}} \phi_0$ as $t \rightarrow 0$ and so we can conclude that the Cauchy sequence converges.

It is easy to see that if N is not complete then \mathcal{M}^\sim is not complete. Consider a sequence of constant functions $\{\phi_t : M \rightarrow N\}_{t \in (0,1)}$ such that

$$\phi_t(x) = y_t$$

where y_t is a Cauchy family in N which does not converge.

□

Theorem 9.3.1 follows from Proposition 9.4.8, Lemma 9.4.11, Lemma 9.5.8, and the fact that every manifold admits a complete Riemannian metric [60, Theorem 1].

9.6 Almost everywhere convergence and \mathcal{D}

We already have a definition of convergence in distance, so in this section we will define and explore the properties of a way in which these maps can converge pointwise almost everywhere. To talk about convergence of a family in $\mathcal{F}(\mathcal{M})$ we must have both the domains and the mappings converge. First, we will describe the convergence of the domains.

Let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Now let $\{B_t \subset M\}_{t \in (a,b)}$ be a collection of measurable subsets of M . Recall the limit inferior and limit superior of a family of sets, given by

$$\underline{\lim}_{t \rightarrow c} (B_t) := \bigcup_{\delta \in (0,1)} \left(\bigcap_{\substack{t \in (a,b), \\ |t-c| < \delta}} B_t \right) \quad (9.6)$$

and

$$\overline{\lim}_{t \rightarrow c}(B_t) := \bigcap_{\delta \in (0,1)} \left(\bigcup_{\substack{t \in (a,b), \\ |t-c| < \delta}} B_t \right) \quad (9.7)$$

respectively. So the limit inferior of the family is the collection of all points which are eventually in every B_t as $t \rightarrow c$ and the limit superior is the collection of all points which are not eventually outside of every B_t . Clearly it can be seen that $\underline{\lim}(B_t) \subset \overline{\lim}(B_t)$. We say that the family converges if these two sets only differ by a set of measure zero. That is,

Definition 9.6.1. Let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$ and let $\{(\phi_t, B_t)\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M})$. If

$$\mu_{\mathcal{V}} \left\{ \overline{\lim}_{t \rightarrow c}(B_t) \setminus \underline{\lim}_{t \rightarrow c}(B_t) \right\} = 0$$

we say that the collection of sets $\{B_t\}_{t \in (a,b)}$ converges to $\underline{\lim}_{t \rightarrow c}(B_t)$ as $t \rightarrow c$ or $\{(\phi_t, B_t)\}_{t \in (a,b)}$ has *converging domains* as $t \rightarrow c$. Furthermore, if $\{[\phi_t, B_t]\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M}^{\sim})$ is such that $\{B_t\}_{t \in (a,b)}$ converges for one choice of representative we say it has *converging domains*.

Remark 9.6.2. Notice that any nested family of subsets will converge by this definition. For $a, b \in \mathbb{R}$ with $a < b$ let $\{B_t\}_{t \in (a,b)}$ be a family of subsets such that for $s, t \in (a, b)$ we have that $s < t$ implies $B_t \subset B_s$. Then

$$\underline{\lim}_{t \rightarrow a} B_t = \overline{\lim}_{t \rightarrow a} B_t = \bigcup_{t \in (a,b)} B_t.$$

⊙

Remark 9.6.3. Notice that if $\{[\phi_t, B_t]\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M}^{\sim})$ has converging domains as $t \rightarrow c$, for $a, b, c \in \mathbb{R}$, $a < b$, $c \in [a, b]$, then we can always choose some collection of representatives $\{(\phi'_t, B'_t) \in [\phi_t, B_t]\}_{t \in (a,b)}$ such that $\underline{\lim} B'_t = \overline{\lim} B'_t$ where both limits are taken as $t \rightarrow c$. ⊙

Now that we understand the convergence of domains we are prepared to describe almost everywhere convergence in \mathcal{M} . Let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Notice that if $x \in \underline{\lim}_{t \rightarrow c}(B_t)$ then there exists some $\delta > 0$ such that if $t \in (a, b)$ and $|t - c| < \delta$ then $x \in B_t$. This means that $\phi_t(x)$ exists for such t so we may ask if $\{\phi_t(x)\}_{t \in (a,b) \cap (c-\delta, c+\delta)}$ converges as a family of points in N

as $t \rightarrow c$. If it does converge than we have a limit

$$\lim_{t \rightarrow c} \phi_t(x)$$

and thus we arrive at Definition 9.3.4.

Remark 9.6.4. Here it is important to notice that the limit (ϕ, B) from Definition 9.3.4 is not unique in \mathcal{M} but by Corollary 9.6.7 we know it does represent a unique element in \mathcal{M}^\sim . Furthermore, given $\{[\phi_t, B_t]\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M}^\sim)$ we can create a family in $\mathcal{F}(\mathcal{M})$ by making a choice of representative for each $t \in (a, b)$. If a choice exists such that the resulting family in $\mathcal{F}(\mathcal{M})$ converges than we say that $\{[\phi_t, B_t]\}_{t \in (a,b)}$ converges almost everywhere pointwise. Corollary 9.6.6 shows that any limit computed in this way gives the same element of \mathcal{M}^\sim . In such a case we would write $[\phi_t] \xrightarrow{\text{a.e.}} [\phi_0]$ as $t \rightarrow c$. Note that the existence of one choice of representatives which converges does not guarantee that all choices will converge. \circlearrowright

We are now ready to prove Theorem 9.3.5.

Proof of Theorem 9.3.5. It is sufficient to prove for families indexed by $(0, 1)$ and limits as $t \rightarrow 0$. Let $\{\mathcal{S}_n\}_{n=1}^\infty$ be a nested exhaustion of M by finite volume sets, $(\phi, B) \in \mathcal{M}$, and $\{(\phi_t, B_t)\}_{t \in (0,1)} \in \mathcal{F}(\mathcal{M})$ such that $\phi_t \xrightarrow{\text{a.e.}} \phi$ as $t \rightarrow 0$. For the duration of this proof let $\underline{\lim}(B_t)$ denote $\underline{\lim}_{t \rightarrow 0}(B_t)$ and $\overline{\lim} B_t$ denote $\overline{\lim}_{t \rightarrow 0} B_t$.

Recall that for $x \in B$ we have that $x \in \underline{\lim} B_t$ and $\phi_t(x) \rightarrow \phi(x)$ as $t \rightarrow 0$ by Definition 9.3.4.

Thus

$$\lim_{t \rightarrow 0} p_{\phi_t, \phi}^d(x) = \lim_{t \rightarrow 0} \min\{1, d(\phi_t(x), \phi(x))\} = \min\left\{1, d\left(\lim_{t \rightarrow 0} \phi_t(x), \phi(x)\right)\right\} = 0.$$

Also notice that for any $x \in M \setminus \overline{\lim} B_t$ we know that $x \notin B$ and also for small enough t we know $x \notin B_t$. That is, there exists some $T \in (0, 1)$ such that $t < T$ implies that $x \notin B_t$ so for such t we have that $x \notin B \cup B_t$. This means that for $t < T$ we have that $p_{\phi_t, \phi}^d(x) = 0$. Thus

$$\lim_{t \rightarrow 0} p_{\phi_t, \phi}^d(x) = 0$$

for any $x \in M \setminus \overline{\lim} B$ as well. Every $x \in \mathcal{S}$ must either

1. be in B or $M \setminus \overline{\lim} B_t$ and thus satisfy $\lim p_{\phi_t, \phi}^d(x) = 0$ as $t \rightarrow 0$;
2. be in $\overline{\lim} B_t \setminus B_0$, which is a set of measure zero.

This means that $p_{\phi_t, \phi}^d \rightarrow 0$ as $t \rightarrow 0$ pointwise almost everywhere. Also notice that each $p_{\phi_t, \phi}^d$ is bounded by the constant function 1, which is integrable on M because $\nu_{\{\mathcal{S}_n\}}(M) = 1$. These two facts allow us to invoke the Lebesgue Dominated Convergence Theorem to conclude that

$$\lim_{t \rightarrow 0} \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t, \phi) = \lim_{t \rightarrow 0} \int_M p_{\phi_t, \phi}^d d\nu_{\{\mathcal{S}_n\}} = \int_M \lim_{t \rightarrow 0} p_{\phi_t, \phi}^d d\nu_{\{\mathcal{S}_n\}} = 0. \quad \square$$

Remark 9.6.5. Notice that the converse of Theorem 9.3.5 does not hold. We know because of Example 9.4.13 in which the family converges in \mathcal{D} but not pointwise almost everywhere. \circlearrowright

The following two results are a consequence of Theorem 9.3.5 and the fact that $(\mathcal{M}^\sim, \mathcal{D}_{\{\mathcal{S}_n\}}^d)$ is a metric space.

Corollary 9.6.6. *Almost everywhere pointwise limits of families in $\mathcal{F}(\mathcal{M}^\sim)$ are unique in \mathcal{M}^\sim . That is, let $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$. Now suppose $\{[\phi_t, B_t]\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M}^\sim)$, $(\phi_t^1, B_t^1), (\phi_t^2, B_t^2) \in [\phi_t, B_t]$ for $t \in (a, b)$, and $(\phi^1, B^1), (\phi^2, B^2) \in \mathcal{M}$ such that $(\phi_t^i, B_t^i) \xrightarrow{a.e.} (\phi^i, B^i)$ as $t \rightarrow c$ for $i = 1, 2$. Then $[\phi^1, B^1] = [\phi^2, B^2]$ in \mathcal{M}^\sim .*

Proof. Let $\{[\phi_t, B_t]\}_{t \in (a, b)}, (\phi_t^1, B_t^1), (\phi_t^2, B_t^2), (\phi^1, B^1)$, and (ϕ^2, B^2) be as in the statement of the Corollary. Thus for any choice of a nested exhaustion of M by finite volume sets $\{\mathcal{S}_n\}_{n=1}^\infty$, a complete metric d on N which is induced by a Riemannian metric, and $t \in (a, b)$ we have that

$$0 \leq \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi^1, \phi^2) \leq \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi^1, \phi_t^1) + \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t^1, \phi_t^2) + \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t^2, \phi^2).$$

The middle term on the right side is zero because $(\phi_t^1, B_t^1) \sim (\phi_t^2, B_t^2)$ and the remaining terms both approach zero as $t \rightarrow c$ because $(\phi_t^i, B_t^i) \xrightarrow{\mathcal{D}} (\phi^i, B^i) \in \mathcal{M}^\sim$ as $t \rightarrow c$ by Theorem 9.3.5. \square

Corollary 9.6.7. *Almost everywhere pointwise limits of families in $\mathcal{F}(\mathcal{M})$ are unique up to \sim . That is, suppose that $a, b, c \in \mathbb{R}$ with $a < b$ and $c \in [a, b]$ and further suppose that $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ and $\phi, \phi' \in \mathcal{M}$. If $\phi_t \xrightarrow{a.e.} \phi$ and $\phi_t \xrightarrow{a.e.} \phi'$ then $\phi \sim \phi'$.*




Figure 9.8: A figure of the relevant maps when defining a smooth family of embeddings.

9.7 Families with singular limits

We will be considering one parameter families of mappings in $\mathcal{F}(\mathcal{M}^\sim)$. For this type of family we can adapt the definition of smoothness from [74] which is visualized in Figure 9.8.

Definition 9.7.1. Let $a, b \in \mathbb{R}$ with $a < b$. We say that a family of smooth maps $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ is *smooth* if:

1. each element of $\{B_t\}$ is a submanifold of M ;
2. there exists a smooth manifold B and a smooth map $g : (a, b) \times B \rightarrow M$ such that
 - (a) the mapping $g_t : b \mapsto g(t, b)$ is a smooth immersion;
 - (b) for each $t \in (a, b)$ we have $g_t(B) = B_t$.
3. the map $(t, b) \mapsto \phi_t \circ g_t(b)$ is smooth.

Despite the choice of terminology, it is unknown if this sense of smoothness implies that the family is continuous with respect to the topology on \mathcal{M} .

Definition 9.7.2. We say that a smooth family $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\text{Emb}_\subset(M, N))$ has a *singular limit* if either

1. the family does not converge in \mathcal{D} as $t \rightarrow a$;

2. there exists some $\phi_0 \in \mathcal{M}$ such that $\phi_t \xrightarrow{\mathcal{D}} \phi_0$ as $t \rightarrow a$ but $[\phi_0] \notin \text{Emb}_{\mathcal{C}}^{\sim}(M, N)$.

In Case 1 we say that the singularity is *essential* because Theorem 9.3.8 assures it cannot be removed by arbitrarily small changes.

Recall the function $r_{\{\mathcal{S}_n\}}^d : \mathcal{F}(\mathcal{M}) \rightarrow [0, \infty]$ from Definition 9.3.7. This function quantifies how far a family is from converging by measuring how much each embedding must be changed in order to create a new family which does converge. It is straightforward to show that r is surjective.

Proposition 9.7.3. *For any $q \in [0, \infty]$ there exists some choice of manifolds M and N , an exhaustion $\{\mathcal{S}_n\}_{n=1}^{\infty}$ of M , a distance d induced by a complete Riemannian metric on N , $a, b \in \mathbb{R}$ such that $a < b$, and a smooth family $\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$ for which $r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}_{t \in (a, b)}) = q$.*

Proof. From the existence of families of embeddings which do converge we know that 0 is in the image of r . Also, notice that if $\phi_t : (0, 1) \rightarrow \mathbb{R}$, $\phi_t(x) = (1/t) \sin(1/t)$ then $r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}_{t \in (a, b)}) = \infty$.

Pick some $q \in (0, 1)$ and let $\phi_t : (0, 3q) \rightarrow \mathbb{R}$ for $t \in (0, 1)$ via

$$\phi_t(x) = \frac{x}{9q} + \frac{1}{3} \sin(1/t).$$

So in this case $B_t = (0, 3q)$ for all t , $a = 0$, $b = 1$, $M = (0, 3q)$ with the usual measure inherited from \mathbb{R} , and $N = \mathbb{R}$ with the usual distance. Since M is finite throughout this example let $\mathcal{D} := \mathcal{D}_{\{M\}}^d$ and $r := r_{\{M\}}^d$ where d is the standard distance on \mathbb{R} . Notice that if we perturbed this family to converge to some limit which did not have $(0, 3q)$ as its domain we could change the domain of the limit to $(0, 3q)$ and have a smaller perturbation. So we can assume that the domain of the limit is $(0, 3q)$. Suppose that we wanted to change this family so it converged to some map $\phi_0 : (0, 3q) \rightarrow \mathbb{R}$. We can see that the ϕ_t oscillate to the left and right, so let

$$\phi_L(x) = \frac{x}{9q} - 1/3$$

and

$$\phi_R(x) = \frac{x}{9q} + 1/3.$$

Now let

$$l_n = \frac{2}{(4n+1)\pi} \text{ and } r_n = \frac{2}{(4n+3)\pi}$$

so that $\phi_{l_n} = \phi_L$ and $\phi_{r_n} = \phi_R$ for all $n \in \mathbb{N}$. Notice

$$d(\phi_L(x), \phi_R(x)) = 2/3$$

for all $x \in (0, 3q)$ so

$$d(\phi_L(x), \phi_0(x)) + d(\phi_0(x), \phi_R(x)) \geq 2/3.$$

Clearly this implies that

$$\min\{1, d(\phi_L(x), \phi_0(x))\} + \min\{1, d(\phi_0(x), \phi_R(x))\} \geq 2/3$$

and so integrating each side over $(0, 3q)$ gives

$$\mathcal{D}(\phi_L, \phi_0) + \mathcal{D}(\phi_0, \phi_R) \geq 2q$$

so one of the two terms must be greater than or equal to q . Without loss of generality suppose that $\mathcal{D}(\phi_L, \phi_0) \geq q$. In such a case choose any $\varepsilon > 0$ and find some $T \in (0, 1)$ such that $t < T$ implies $\mathcal{D}(\tilde{\phi}_t, \phi_0) < \varepsilon$ where $\{\tilde{\phi}_t\}$ is any family which converges to ϕ_0 . Then pick some $n \in \mathbb{N}$ such that $l_n < T$ and let $t = l_n$. Now

$$\mathcal{D}(\phi_t, \tilde{\phi}_t) + \mathcal{D}(\tilde{\phi}_t, \phi_0) \geq \mathcal{D}(\phi_t, \phi_0)$$

so $\mathcal{D}(\phi_t, \tilde{\phi}_t) \geq q - \varepsilon$ for all $\varepsilon > 0$. This allows us to conclude that $r(\{(\phi_t, B_t)\}) \geq q$.

Now let $\tilde{\phi}_t : (0, 3q) \rightarrow \mathbb{R}$ with $\tilde{\phi}_t(x) = \frac{x}{9q}$ be a family of maps which is clearly smooth and has limit $\phi_0(x) = \frac{x}{9q}$. Now notice

$$\mathcal{D}(\phi_t, \tilde{\phi}_t) = \int_{(0, 3q)} \min\{1, d(\phi_t, \tilde{\phi}_t)\} d\mu_{\mathcal{V}} = q |\sin(1/t)| \leq q$$

and it is important to notice that $\mathcal{D}(\phi_t, \tilde{\phi}_t) = q$ is achieved infinitely often. Thus we know that $r(\{(\phi_t, B_t)\}) \leq q$ so in fact we know that $r(\{(\phi_t, B_t)\}) = q$. \square

It is unknown if having zero radius of convergence is independent of the choice of parameters d and $\{\mathcal{S}_n\}_{n=1}^\infty$ (see Section 9.8.2). In the case that $r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}) = 0$ we say that the family has a *removable singularity with respect to $\mathcal{D}_{\{\mathcal{S}_n\}}^d$* . Now we are prepared to prove Theorem 9.3.8.

Proof of Theorem 9.3.8. Suppose that $\{(\phi_t, B_t)\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M})$ and $r(\{(\phi_t, B_t)\}_{t \in (a,b)}) = 0$. Fix some compact $\mathcal{S} \subset M$ and we will show that $\{(\phi_t, B_t)\}_{t \in (a,b)}$ is Cauchy with respect to $\mathcal{D}_{\mathcal{S}}$.

Fix some $\varepsilon > 0$. Let $\delta = \varepsilon/4$ and since $r(\{(\phi_t, B_t)\}_{t \in (a,b)}) = 0$ define some other family $\{(\widetilde{\phi}_t^\delta, \widetilde{B}_t^\delta)\}$ such that

1. $\mathcal{D}(\phi_t, \widetilde{\phi}_t^\delta) < \delta$ for all $t \in (a, b)$;
2. $B_t = \widetilde{B}_t^\delta$ for all $t \in (a, b)$;
3. there exists some $\widetilde{\phi}^\delta \in \text{Emb}_{\mathcal{C}}(M, N)$ such that $\widetilde{\phi}_t^\delta \xrightarrow{\text{a.e.}} \widetilde{\phi}^\delta$ as $t \rightarrow a$.

From Theorem 9.3.5 and item (3) above we know that

$$\widetilde{\phi}_t^\delta \xrightarrow{\mathcal{D}_{\mathcal{S}}^d} \widetilde{\phi}_0^\delta$$

as $t \rightarrow a$ so we can choose some $T \in (a, b)$ such that $t < T$ implies $\mathcal{D}_{\mathcal{S}}(\widetilde{\phi}_t^\delta, \widetilde{\phi}^\delta) < \delta$. Finally, we can conclude that for any $t, s < T$ we have that

$$\begin{aligned} \mathcal{D}_{\mathcal{S}}(\phi_t, \phi_s) &\leq \mathcal{D}_{\mathcal{S}}(\phi_t, \widetilde{\phi}_t^\delta) + \mathcal{D}_{\mathcal{S}}(\widetilde{\phi}_t^\delta, \widetilde{\phi}^\delta) + \mathcal{D}_{\mathcal{S}}(\widetilde{\phi}^\delta, \widetilde{\phi}_s^\delta) + \mathcal{D}_{\mathcal{S}}(\widetilde{\phi}_s^\delta, \phi_s) \\ &< 4\delta = \varepsilon. \end{aligned}$$

This means that $\{(\phi_t, B_t)\}_{t \in (a,b)}$ is Cauchy as $t \rightarrow a$ for each $\mathcal{D}_{\mathcal{S}}$ so by Proposition 9.4.6 we know that it is Cauchy with respect to \mathcal{D} as $t \rightarrow a$. Finally, since $(\mathcal{M}^\sim, \mathcal{D})$ is complete by Theorem 9.3.1 we can come to the first conclusion of this Theorem.

Now we will show the second claim. Suppose that the domains satisfy the required property for $T \in (a, b)$ and that $\phi_t \xrightarrow{\mathcal{D}} \phi_0$ as $t \rightarrow a$. Fix $\varepsilon > 0$ and find some $T_1 \in (a, T)$ such that $s, t < T_1$ implies that $\mathcal{D}(\phi_t, \phi_s) < \varepsilon$. Now let $\mathcal{B} : (a, b) \rightarrow [0, 1]$ be a smooth bump function such that $\mathcal{B}(t) = 0$ for $t \geq T_1$ and $\mathcal{B}(t) = 1$ for $t < T_1 + a/2$. Now define $f : (a, b) \rightarrow [T_1 + a/2, b)$ via

$$f(t) = (1 - \mathcal{B}(t))t + \mathcal{B}(t)\frac{T_1 + a}{2}.$$

Finally let

$$\tilde{\phi}_t = \phi_{f(t)}|_{B_t}$$

and notice that $\{(\tilde{\phi}_t, B_t)\}$ is a smooth family satisfying $\tilde{\phi}_t \xrightarrow{\text{a.e.}} \phi_{T_1/2}$ as $t \rightarrow a$. By the choice of T_1 we can see that for all $t \in (a, b)$ we have $\mathcal{D}(\phi_t, \tilde{\phi}_t) < \varepsilon$. Also, because of the requirement on the domains we know that $B_t \subset B_{f(t)}$ and thus $\tilde{\phi}_t : B_t \rightarrow N$ is defined on all of B_t .

□

Remark 9.7.4. It is natural to wonder if $r(\{(\phi_t, B_t)\}_{t \in (a,b)}) = 0$ implies the family must in fact converge pointwise almost everywhere in \mathcal{M}^\sim . The answer to this question is no; again consider Example 9.4.13. The functions in Example 9.4.13 converge in \mathcal{D} and all have the same domain so we know that $r_{\{\mathcal{S}_n\}}^d = 0$ for these functions, but we also know that they do not converge pointwise almost everywhere. ◊

9.8 Further questions

9.8.1 Approaches to prove a converse to Theorem 9.3.8

Now we have set up all of the machinery to begin to explore the converse of Theorem 9.3.8 in the case that the domains are not restricted to shrink or stabilize eventually. That is, we will outline some potential avenues to answer the following question.

Question 9.8.1 Is it true that $\{(\phi_t, B_t)\}_{t \in (a,b)} \xrightarrow{\mathcal{D}} \phi_0$ implies that $r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}_{t \in (a,b)}) = 0$?

There are two approaches in the general case: we can attempt to extend embeddings or we can smooth singular limits by understanding the singularities locally.

Extending embeddings to remove singularities

This follows the method used to prove the partial converse direction of Theorem 9.3.8 given in the statement of the theorem. The idea is that if $\{(\phi_t, B_t)\}_{t \in (a,b)} \xrightarrow{\mathcal{D}} \phi_0$ as $t \rightarrow c$ (for $a, b, c \in \mathbb{R}, a < b, c \in [a, b]$) then in order to get an ε -perturbation of $\{(\phi_t, B_t)\}_{t \in (a,b)}$ we choose

some $T \in (a, b)$ such that $s, t < T$ implies that $\mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t, \phi_s) < \varepsilon$. Then, just as in the proof of Theorem 9.3.8, we must smoothly change the family so that $t < \frac{T+a}{2}$ implies that $\tilde{\phi}_t = \phi_{\frac{T+a}{2}}$. The difficulty here is dealing with the domains. If $\varinjlim B_t \not\subset B_{\frac{T}{2}}$ then this idea we have outlined will not define an embedding with domain all of B_0 , so this embedding would have to be extended. It is important to notice that $\mu_V(B_0 \Delta B_{\frac{T+a}{2}}) < \varepsilon$ and so the embedding can be defined in any way on the extension, as long as it does not change on $B_{\frac{T+a}{2}}$. Thus, this question comes down to asking when an embedding of some subset of M can be extended to a larger domain in M . Extending embeddings or smooth maps has been of independent interest for many years. See for example the Tietze Extension Theorem [32, Theorem 4.16], the Whitney Extension Theorem [81, Theorem I], and the Extension Lemma [52, Lemma 2.27]. For a collection of more recent work in extension problems see [9].

Removing singularities locally

The basic strategy is the following. Suppose for $a, b, c \in \mathbb{R}, a < b, c \in [a, b]$ that

$$\{(\phi_t, B_t)\}_{t \in (a, b)} \in \mathcal{F}(\mathcal{M})$$

satisfies $\phi_t \xrightarrow{\mathcal{D}} \phi$ as $t \rightarrow a$ for some $\phi \in \mathcal{M}$ and suppose further that $\mathcal{S} \subset B$ is a closed subset of M containing all of the singular points of the limiting map ϕ and that eventually $\mathcal{S} \subset B_t$ for all t . That is, we assume that

$$\phi|_{B \setminus \mathcal{S}} : B \setminus \mathcal{S} \hookrightarrow N$$

is an embedding and there exists some $T \in (a, b)$ such that $t < T$ implies $\mathcal{S} \subset B_t$. Then for some neighborhood of \mathcal{S} we can define $\tilde{\phi}$ by ϕ_{t_0} restricted to that neighborhood for some small enough $t_0 \in (a, b)$. Then to define $\tilde{\phi}$ outside of a slightly larger neighborhood of \mathcal{S} we simply use ϕ unchanged. We must then connect these two pieces in a way which makes the result an embedding. Finally each ϕ_t can then be changed on a neighborhood of \mathcal{S} to converge to ϕ_{t_0} and outside of that neighborhood they converge to $\phi = \tilde{\phi}$ already. A schematic of this idea is shown in Figure 9.9. The difficulty comes when we must connect the two embeddings; it is well known that partition of unity type arguments can be used to smoothly transition between two smooth maps [52] but in this case we must also




Figure 9.9: The strategy is to connect the embedding ϕ_{t_0} with the map ϕ which is an embedding away from \mathcal{S} . In this way we are able to avoid the singular part of ϕ while only changing it slightly on a small set.

preserve the embedding structure.

9.8.2 Implications of a positive answer to Question 9.8.1

If the answer to Question 9.8.1 were yes, then there are several implications. First, we will have a new characterization of families with removable singularities, namely these are exactly the families which converge in \mathcal{D} . Second, there is then an easy proof that $r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}) = 0$ does not depend on the choices of $\{\mathcal{S}_n\}_{n=1}^\infty$ and d . The proof is the following:

Let $\{\mathcal{S}_n\}_{n=1}^\infty$, $\{\mathcal{S}'_n\}_{n=1}^\infty$, d , and d' be choices of finite exhaustion and metric. Suppose that $a, b \in \mathbb{R}$ with $a < b$ and $\{(\phi_t, B_t)\}_{t \in (a,b)} \in \mathcal{F}(\mathcal{M})$ is a smooth family such that

$$r_{\{\mathcal{S}_n\}}^d(\{(\phi_t, B_t)\}_{t \in (a,b)}) = 0.$$

Then by Theorem 9.3.8 we know that $\lim \mathcal{D}_{\{\mathcal{S}_n\}}^d(\phi_t, \phi) = 0$ as $t \rightarrow a$ for some $\phi \in \mathcal{M}$. By Theorem 9.3.1 this means that $\lim \mathcal{D}_{\{\mathcal{S}'_n\}}^{d'}(\phi_t, \phi) = 0$ as $t \rightarrow a$ and thus by the assumed positive answer to Question 9.8.1 we know that $r_{\{\mathcal{S}'_n\}}^{d'}(\{(\phi_t, B_t)\}_{t \in (a,b)}) = 0$.

9.8.3 Applications in symplectic and Riemannian geometry

In [67] the authors produce a metric on the space of toric integrable systems. An integrable system is a $2n$ -dimensional symplectic manifold (M, ω) along with a map $F: M \rightarrow \mathbb{R}^n$ such that the components of F Poisson commute and are independent almost everywhere. The metric defined in

the present chapter could potentially be used to define a metric on more general spaces of integrable systems, as long as those systems could be viewed as subsets of the same manifold.

It would be interesting to study Question 9.8.1 restricted to a specific type of embedding. For example, thinking back to the original motivation from Section 9.1, one could consider whether this is true for the collection of symplectic embeddings where the original smooth family $\{(\phi_t, B_t)\}_{t \in (a,b)}$ consists exclusively of symplectic embeddings and the perturbed family $\{(\tilde{\phi}_t, \tilde{B}_t)\}_{t \in (a,b)}$ from the definition of the radius of convergence is also required to be symplectic. Resolving singular points of symplectic manifolds is related to this in spirit and has been studied extensively such as in [54]. Symplectic manifolds have been shown to admit a high degree of flexibility (see for example Moser's Theorem [57] or Darboux's Theorem [10]) although Gromov's nonsqueezing theorem [35] represents a level of rigidity that symplectic embeddings do need to respect. One could also consider the case of isometric embeddings of Riemannian manifolds, even in the case of $\mathcal{M}(\mathbb{R}, \mathbb{R}^2)$. Clearly studying further types of embeddings would be enlightening as it would allow us to gain a greater understanding of the rigidity of these structures. Indeed, it is the purpose of this chapter to create a foundation off of which many types of families of embeddings may be studied.

Acknowledgements. Chapter 9, in full, is comprised of material submitted for publication by the author of this dissertation as *Metrics and convergence in moduli spaces of maps*. currently available at arXiv:1406.4181 [62].

Index

- Action-Angle theorem, 19
- admissible measure, 41
- admissible packing, 173
- admissible semitoric packing, 179
- affine invariant, *see* polygon invariant
- almost everywhere locally free, 159

- blowdown, 28
- blowup, 28
- blowup/down, 120

- commute fake and Delzant corners, 84
- convergence
 - almost everywhere pointwise, 196, 221
 - convergent perturbation, 196
 - in \mathcal{D} , 195
 - sets, 222
- corner chop, 83, 172
- Cotangent bundle, 10

- Darboux theorem, 10
- Delzant
 - corner, 34
 - polytope, 27
 - triangle, 102

- equivariant Gromov radius, 151, 155

- everywhere finite height, 36

- fake corner, 34
- fan, *see also* toric fan, 80
- four fan transformations, 83

- $G_{m_f} \times \mathcal{G}$, 35
- Gromov radius, 14, 15
 - equivariant, 155

- Hamiltonian
 - dynamical system, 16
 - flow action, 159
 - group action, 12, 159
 - vector field, 16
- harmonic oscillator, 18
- Hessian, \mathcal{H}_p , 22
- hidden corner, 34
- Hirzebruch trapezoid, 102

- integrable system, 17
 - and group actions, 159
 - semitoric, 30
 - simple semitoric, 30
 - toric, 27, 58, 161

- labeled Delzant semitoric polygon, 35

- labeled weighted polygon, 33
- Lagrangian submanifold, 11
- linear summable sequence, 40
- m_f , 31
- metric on maps, 194
- minimal models
 - toric, 80, 102
- moduli space of maps, 191
- momentum map
 - actions vs. integrable systems, 18
 - of a group action, 13, 159
 - of an integrable system, 17
- non-degenerate singular point, 23
- non-squeezing theorem, 14, 15
- penalty function, 191
- Poisson bracket, 17
- polygon invariant, 32
- $\mathcal{Q}(T_p M)$, 21
- quadratic form, 21
- removal of hidden corner, 84
- reverse corner chop, 83
- semitoric
 - ball packing, 166
 - blowup/down, 120
 - completion of moduli space, 45, 62
 - embedding, 166
 - fan, 81, 111
 - integrable system, 30
 - invariants of system, 36
 - isomorphism, 30
 - list of ingredients, 36
 - manifold, 30
 - metric, 45
 - optimal packing function, 182
 - packing capacity, 152, 166
 - radius capacity, 187
 - special cases of metric, 46
 - singularity
 - elliptic, 30
 - focus-focus, 30, 32
 - smooth angle, 182
 - standard semitoric fan, 83, 105
 - straightening map, 165
 - symplectic
 - ball, 14, 153
 - capacity, 15
 - category, 15
 - cylinder, 14, 153
 - embedding, 14, 231
 - form, 8
 - G -capacity, 169
 - G -embedding, 153
 - generalized G -capacity, 153
 - group action or G -action, 12, 153
 - manifold, 8
 - $(\mathbb{T}^k \times \mathbb{R}^{d-k})$ -capacity, 155

symplectomorphism, 8

tame, 154

tautological form, 10

Taylor series invariant, 32

toric

 ball packing, 162

 blowup/down, 119

 fan, 28, 78, 80, 97

 integrable system, 27

 optimal density function, 162

 variety, 80

twisting index invariant, 32

 equivalence of, 42

volume invariant, 36

Williamson type, 26

winding number, 96, 123

winding number (of vectors), 85

zero section, 12

Bibliography

- [1] M. Abrue, *Topology of symplectomorphism groups of $S^2 \times S^2$* , *Inventiones mathematicae* **131** (1996), 1–24.
- [2] V. I. Arnold, *A theorem of Liouville concerning integrable problems of dynamics.*, *Sibirsk. Mat. Ž.* **4** (1963), 471–474.
- [3] M. Atiyah, *Convexity and commuting Hamiltonians*, *Bull. London Math. Soc.* **14** (1982), 1–15.
- [4] O. Babelon and B. Douçot, *Classical Bethe Ansatz and normal forms in an integrable version of the Dicke model*, *Phys. D* **241** (2012), no. 23-24, 2095–2108.
- [5] ———, *Higher index focus-focus singularities in the Jaynes-Cummings-Gaudin model: symplectic invariants and monodromy*, *J. Geom. Phys.* **87** (2015), 3–29.
- [6] S. M. Bates, *Some simple continuity properties of symplectic capacities*, *The Floer memorial volume*, *Progr. Math.*, vol. 133, Birkhäuser, Basel, 1995, pp. 185–193.
- [7] P. Biran, *Connectedness of spaces of symplectic embeddings*, *Int. Math. Res. Lett.* **10** (1996), 487–491.
- [8] A. V. Bolsinov and A. T. Fomenko, *Integrable Hamiltonian systems*, Chapman & Hall/CRC, Boca Raton, FL, 2004, *Geometry, topology, classification*, Translated from the 1999 Russian original.
- [9] A. Brudnyi and Y. Brudnyi, *Methods of geometric analysis in extension and trace problems*, *Monographs in Mathematics*, vol. 102–103, Springer-Birkhäuser, 2012.
- [10] A. Cannas da Silva, *Lectures on symplectic geometry*, Springer-Verlag, Berlin, 2008.
- [11] L. Charles, Á. Pelayo, and S. Vũ Ngọc, *Isospectrality for quantum toric integrable systems*, *Ann. Sci. Éc. Norm. Supér. (4)* **46** (2013), no. 5, 815–849.
- [12] K. Christian and V. Turaev, *Presentations of $SL_2(z)$ and $PSL_2(z)$* , *Braid Groups*, *Graduate Texts in Mathematics*, vol. 247, Springer New York, 2008, pp. 311–314 (English).
- [13] K. Cieliebak, H. Hofer, J. Latschev, and F. Schlenk, *Quantitative symplectic geometry*, *Dynamics, ergodic theory, and geometry*, *Math. Sci. Res. Inst. Publ.*, vol. 54, Cambridge Univ. Press, Cambridge, 2007, pp. 1–44.
- [14] Y. Colin de Verdière, *Spectre conjoint d’opérateurs pseudo-différentiels qui commutent. I. Le cas non intégrable*, *Duke Math. J.* **46** (1979), no. 1, 169–182.
- [15] ———, *Spectre conjoint d’opérateurs pseudo-différentiels qui commutent. II. Le cas intégrable*, *Math. Z.* **171** (1980), no. 1, 51–73.

- [16] D. Cox, *What is a toric variety?*, Topics in algebraic geometry and geometric modeling, Contemp. Math., vol. 334, Amer. Math. Soc., Providence, RI, 2003, pp. 203–223.
- [17] D.A. Cox, J.B. Little, and H.K. Schenck, *Toric varieties*, Graduate Studies in Mathematics, vol. 124, American Mathematical Society, Providence, RI, 2011.
- [18] F. W. Cummings, *Stimulated emission of radiation in a single mode*, Phys. Rev. **140** (1965), A1051–A1056.
- [19] V.I. Danilov, *The geometry of toric varieties*, Russian Mathematical Surveys **33** (1978), no. 2, 97–154.
- [20] M. de Gosson and F. Luef, *Symplectic capacities and the geometry of uncertainty: the irruption of symplectic topology in classical and quantum mechanics*, Phys. Rep. **484** (2009), no. 5, 131–179.
- [21] T. Delzant, *Hamiltoniens périodiques et image convexe de l'application moment*, Bull. Soc. Math. France **116** (1988), 315–339.
- [22] J. J. Duistermaat, *On global action-angle coordinates*, Comm. Pure Appl. Math. **33** (1980), no. 6, 687–706.
- [23] J.J. Duistermaat and G.J. Heckman, *On the variation in the cohomology of the symplectic form of the reduced phase space*, Invent. Math. **69** (1982), 259–268.
- [24] J.J. Duistermaat and Á. Pelayo, *Reduced phase space and toric variety coordinatizations of Delzant spaces*, Math. Proc. Cambridge Philos. Soc. **146** (2009), no. 3, 695–718.
- [25] H. Dullin, *Semi-global symplectic invariants of the spherical pendulum*, J. Differential Equations **254** (2013), no. 7, 2942–2963.
- [26] I. Ekeland and H. Hofer, *Symplectic topology and Hamiltonian dynamics*, Math. Z. **200** (1989), no. 3, 355–378.
- [27] Y. Eliashberg and L. Polterovich, *Symplectic quasi-states on the quadric surface and lagrangian submanifolds*, arXiv:1006.2501v1.
- [28] L. H. Eliasson, *Hamiltonian systems with poisson commuting integrals*, Ph.D. thesis, University of Stockholm, 1984.
- [29] ———, *Normal forms for Hamiltonian systems with poisson commuting integrals—elliptic case*, Comment. Math. Helv. **65** (1990), 4–35.
- [30] A. Figalli, J. Palmer, and Á. Pelayo, *Symplectic g-capacities and integrable systems*, arXiv:1511.04499.
- [31] A. Figalli and Á. Pelayo, *Continuity of ball packing density on moduli spaces of toric manifolds*, Advances in Geometry, (to appear) arXiv:1408.1462.
- [32] G. Folland, *Real analysis: Modern techniques and their applications*, Wiley, 1999.
- [33] W. Fulton, *Introduction to toric varieties*, Annals of Mathematics Studies, no. 131, Princeton University Press, 1993.
- [34] L. Godinho, Á. Pelayo, and S. Sabatini, *Fermat and the number of fixed points of periodic flows*, arXiv:1404.4541.
- [35] M. Gromov, *Pseudoholomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347.

- [36] M. Gross and B. Siebert, *Mirror symmetry via logarithmic degeneration data, II*, J. Algebraic Geom. **19** (2010), no. 4, 679–780.
- [37] ———, *From real affine geometry to complex geometry*, Ann. of Math. (2) **174** (2011), no. 3, 1301–1428.
- [38] V. Guillemin, *Kaehler structures on toric varieties*, J. Differential Geom. **40** (1994), no. 2, 285–309.
- [39] ———, *Moment maps and combinatorial invariants of Hamiltonian T^n -spaces*, Progress in Mathematics, vol. 122, Birkhäuser Boston, Inc., Boston, MA, 1994.
- [40] V. Guillemin and S. Sternberg, *Convexity properties of the moment mapping*, Invent. Math. **67** (1982), 491–513.
- [41] J.C. Hausmann and A. Knutson, *Polygon spaces and Grassmannians*, Enseign. Math. (2) **43** (1997), no. 1-2, 173–198.
- [42] ———, *The cohomology ring of polygon spaces*, Ann. Inst. Fourier (Grenoble) **48** (1998), no. 1, 281–321.
- [43] H. Hofer, *On the topological properties of symplectic maps*, Proceedings of the Royal Society of Edinburgh **115** (1990), 25–38.
- [44] ———, *Symplectic capacities*, Geometry of low-dimensional manifolds, 2 (Durham, 1989), London Math. Soc. Lecture Note Ser., vol. 151, Cambridge Univ. Press, Cambridge, 1990, pp. 15–34.
- [45] H. Hopf and W. Rinow, *Über den begriff der vollständigen differentialgeometrischen fläche*, Comment. Math. Helv. **3** (1931), no. 1, 209–225.
- [46] E.T. Jaynes and F.W. Cummings, *Comparison of quantum and semiclassical radiation theories with application to the beam maser*, Proceedings of the IEEE **51** (1963), no. 1, 89–109.
- [47] D. M. Kane, J. Palmer, and Á. Pelayo, *Classifying toric and semitoric fans by lifting equations from $SL_2(\mathbb{Z})$* , arXiv:1502.07698.
- [48] M. Kapovich and J. Millson, *On the moduli space of polygons in the Euclidean plane*, J. Differential Geom. **42** (1995), no. 2, 430–464.
- [49] Y. Karshon and L. Kessler, *Circle and torus actions on equal symplectic blow-ups of $\mathbb{C}P^2$* , Math. Res. Lett. **14** (2007), no. 5, 807–823.
- [50] Y. Karshon and S. Tolman, *The Gromov width of complex Grassmannians*, Alg. and Geo. Topology **5** (2005), 911–922.
- [51] S. Kuksin, *Infinite-dimensional symplectic capacities and a squeezing theorem for Hamiltonian PDEs*, Comm. Math. Phys. **167** (1995), no. 3, 531–552.
- [52] J. M. Lee, *Introduction to smooth manifolds*, Springer, 2006.
- [53] N. C. Leung and M. Symington, *Almost toric symplectic four-manifolds*, J. Symplectic Geom. **8** (2010), no. 2, 143–187.
- [54] J. D. McCarthy and J. G. Wolfson, *Symplectic gluing along hypersurfaces and resolution of isolated orbifold singularities*, Inventiones mathematicae **119** (1995), no. 1, 129–154.
- [55] D. McDuff and D. Salamon, *Introduction to symplectic topology*, second ed., Oxford Mathematical Monographs, The Clarendon Press, Oxford University Press, New York, 1998.

- [56] E. Miller, *What is...a toric variety?*, Notices Amer. Math. Soc. **55** (2008), no. 5, 586–587.
- [57] J. Moser, *On the volume elements on a manifold*, Transactions of the AMS **120** (1965), 286–294.
- [58] O. Müller, *A note on closed isometric embeddings*, Journal of Mathematical Analysis and Applications **349** (2008), no. 1, 297–298.
- [59] J. Nash, *The imbedding problem for Riemannian manifolds*, Annals of Math. (2) **63** (1956), no. 1, 20–63.
- [60] K. Nomizu and H. Ozenki, *The existence of complete Riemannian metrics*, Proceedings of the AMS **12** (1961), no. 6, 889–891.
- [61] T. Oda, *Convex bodies and algebraic geometry*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 15, Springer-Verlag, Berlin, 1988, An introduction to the theory of toric varieties, Translated from the Japanese.
- [62] J. Palmer, *Metrics and convergence in moduli spaces of maps*, arXiv:1406.4181.
- [63] J. Palmer, *Moduli spaces of semitoric systems*, arXiv:1502.07296.
- [64] Á Pelayo, *Toric symplectic ball packing*, Topology and its Appl. **157** (2006), 3633–3644.
- [65] Á. Pelayo, *Topology of spaces of equivariant symplectic embeddings*, Proc. Amer. Math. Soc. **135** (2007), no. 1, 277–288.
- [66] Á. Pelayo and S. Vũ Ngọc, *Sharp symplectic embeddings of cylinders*, Indag. Math. (N.S.) **27** (2016), no. 1, 307–317.
- [67] Á. Pelayo, A.R. Pires, T. Ratiu, and S. Sabatini, *Moduli spaces of toric manifolds*, Geometriae Dedicata **169** (2014), 323–341.
- [68] Á Pelayo and B. Schmidt, *Maximal ball packings of symplectic-toric manifolds*, Intern. Math. Res. Not. (2008), 24p, ID rnm139.
- [69] Á. Pelayo and S. Vũ Ngọc, *Semitoric integrable systems on symplectic 4-manifolds*, Invent. Math. **177** (2009), 571–597.
- [70] ———, *Constructing integrable systems of semitoric type*, Acta Math. **206** (2011), 93–125.
- [71] ———, *Symplectic theory of completely integrable Hamiltonian systems*, Bull. Amer. Math. Soc. **48** (2011), 409–455.
- [72] ———, *First steps in symplectic and spectral theory of integrable systems*, Discrete and Cont. Dyn. Syst., Series A **32** (2012), 3325–3377.
- [73] ———, *Hamiltonian dynamics and spectral theory for spin-oscillators*, Comm. Math. Phys. **309** (2012), 123–154.
- [74] ———, *Hofer’s question on intermediate symplectic capacities*, Proc. Lond. Math. Soc. (3) **110** (2015), no. 4, 787–804.
- [75] M. H. Stone, *Applications of the theory of Boolean rings to general topology*, Trans. Amer. Math. Soc. **41** (1937), 375–481.
- [76] S. Vũ Ngọc, *On semi-global invariants of focus-focus singularities*, Topology **42** (2003), no. 2, 365–380.
- [77] ———, *Moment polytopes for symplectic manifolds with monodromy*, Adv. Math. **208** (2007), no. 2, 909–934.

- [78] E. Čech, *On bicomact spaces*, Ann. of Math. **38** (1937), no. 2, 823–844.
- [79] S. Vũ Ngọc, *Systèmes intégrables semi-classiques: du local au global*, Panoramas et Synthèses [Panoramas and Syntheses], vol. 22, Société Mathématique de France, Paris, 2006.
- [80] C. Wacheux, *Asymptotics of action variables near semi-toric singularities*, J. Geom. Phys. **98** (2015), 28–39.
- [81] H. Whitney, *Analytic extensions of functions defined on closed sets*, Transactions of the AMS **36** (1934), no. 1, 63–89.
- [82] J. Williamson, *On the algebraic problem concerning the normal form of linear dynamical systems*, Amer. J. Math. **58** (1936), 141–163.
- [83] S. T. Yau, *Remarks on the group of isometries of a Riemannian manifold*, Topology **16** (1977), no. 3, 239–247.
- [84] ———, *A survey of Calabi-Yau manifolds*, Geometry and analysis. No. 2, Adv. Lect. Math. (ALM), vol. 18, Int. Press, Somerville, MA, 2011, pp. 521–563.
- [85] K. Zehmisch and F. Ziltener, *Discontinuous symplectic capacities*, J. Fixed Point Theory Appl. **14** (2013), no. 1, 299–307.
- [86] S. Zelditch, *The inverse spectral problem for surfaces of revolution*, J. Differential Geom. **49** (1998), no. 2, 207–264.
- [87] N. T. Zung, *Symplectic topology of integrable Hamiltonian systems, I: Arnold-Liouville with singularities*, Compositio Math. **101** (1996), no. 2, 179–215.