

Symplectic–prequantum structures and dynamics on the codimension-2 shape space

by

Sadashige Ishida

October, 2025

*A thesis submitted to the
Graduate School
of the
Institute of Science and Technology Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy*

Committee in charge:

Julian Fischer, Chair

Chris Wojtan

Albert Chern

Jan Maas

Cornelia Vizman

The thesis of Sadashige Ishida, titled *Symplectic–prequantum structures and dynamics on the codimension-2 shape space*, is approved by:

Supervisor: Chris Wojtan, ISTA, Klosterneuburg, Austria

Signature: _____

Co-supervisor: Albert Chern, UCSD, California, USA

Signature: _____

Committee Member: Jan Maas, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Cornelia Vizman, West University of Timișoara, Romania

Signature: _____

Defense Chair: Julian Fischer, ISTA, Klosterneuburg, Austria

Signature: _____

Signed page is on file

© by Sadashige Ishida, October, 2025

CC BY 4.0 The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution 4.0 International License. Under this license, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author.

ISTA Thesis, ISSN: 2663-337X

ISBN: 978-3-99078-070-1

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I accept full responsibility for the content and factual accuracy of this work, including the data and their analysis and presentation, and the text and citation of other work.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: _____

Sadashige Ishida
October, 2025

Abstract

The space of codimension-2 shapes, such as curves in 3D and surfaces in 4D, is an infinite-dimensional manifold. This thesis explores geometric structures and dynamics on this space, with emphasis on their implications for physics, particularly hydrodynamics.

Our investigation ranges from theoretical studies of infinite-dimensional symplectic and prequantum geometry to numerical computation of the time evolution of shapes. The thesis presents four main contributions.

In the first part, we introduce implicit representations of codimension-2 shapes using a class of complex-valued functions, and prove that the space of these implicit representations forms a prequantum bundle over the codimension-2 shape space. This reveals a new geometric interpretation of the canonical symplectic structure on the codimension-2 shape space.

In the second part, we use implicit representations to develop a simulation method for the dynamics of space curves. To handle chaotic systems such as vortex filaments in hydrodynamics, we exploit the infinite degrees of freedom, hidden in both the configuration and dynamics of implicit representations.

In the third part, we introduce new symplectic structures on the space of space curves, which generalize the only previously known symplectic structure on this space, allowing for new Hamiltonian dynamics of space curves.

In the fourth part, we apply a symplectic viewpoint to a differential geometric problem with practical applications. We derive a new area formula for spherical polygons via prequantization.

Acknowledgements

My PhD journey started (actually didn't start) when I finished my master's study in mathematics with poor grades and no chance of doing a PhD. But after some time, I'm now completing a PhD with a thesis that lies somewhat in the union of math, physics, and computer graphics.¹ During these times, I had a great (not only PhD) life.

These are all due to the people who made it possible. Here I want to acknowledge them.

My first thanks go to those who maintain the great working environment and the beautiful campus of ISTA. Specifically, I thank the Cleaning staff, Construction, IT, Admin, Campus services, Gradschool office, ISTA Family, and everyone who contributes to keeping ISTA a great place to work. I love chatting over free coffee in a clean common space. But moments like this are not automatic. Their hard work and dedication let it happen. I am truly grateful for their commitment, which has let me focus on research and have fun in such a pleasant environment.

Next I thank my advisors and collaborators. I am doing research across multiple fields. This was made possible thanks to the huge support I received from my advisors and collaborators. They opened the doors of research for me, who was (and still is) a novice in these fields, welcomed me as a colleague, guided me through, and shared their expertise. They spent an immense amount of time for me. Thanks to their support, I am now completing my PhD.

I thank Albert Chern. His influence on me is huge and will remain throughout my scientific life. He does science seamlessly across multiple fields without being trapped by labels like mathematics, graphics, physics, and engineering. I'm always impressed by his command of mathematics, which is not just fancy mathematics for itself, but is strongly driven by his deep physical and geometric intuition, and the way he can talk about it using the language of different fields. I want to do science this way.

I thank Hugo Lavenant, who guided me through my first research experience in the field of optimal transport. His responses to my questions and ideas were always quick and spot-on. I also learned from him how to write a mathematical paper in a reader-friendly way.

I thank Martin Bauer. He welcomed me into the field of infinite-dimensional geometry, showing that one can naturally enter research in this seemingly background-heavy and intimidating field. Among the many mathematicians I have met, he stood out for his low barrier to working with. I learned a lot from his exceptional ability to find interesting research questions and make reasonable decisions.

I thank Peter W. Michor, a math giant and a founder of the modern theory of infinite-dimensional geometry. I have always been amazed by the way he can turn my vague thoughts

¹On the way, I got some thoughts about math, computer graphics, and science in general. If you are interested, see Chapter 8.

into theorems, and never hesitates to get his hands dirty with manual computations. I also like coffee you make for us after lunch.

I thank Fumika Suzuki. At a time when I had a difficult time keeping motivation for my PhD and didn't know what to do, she invited me to a project on quantum mechanics, a field I was always interested in but didn't know what to start from. Her open-mindedness and hands-on support saved me during that period. I look forward to completing our suspended project with you.

I thank Toshiya Hachisuka. My very first research started with him. He showed me how to perform scientific research while respecting my original ideas. Toshiya is known to be a photo-realistic rendering guy, but to me he is like a scientist, who can think outside the Cornell box. This inspires me to become a scientist rather than either of a mathematician, physicist, or computer scientist.

I thank Ryoichi Ando. He was my first teacher in fluid dynamics and physics simulation, from whom I learned how fluids behave intuitively and visually. I hope to become an advisor like him who is always accessible and offers help in an hands-on manner.

I thank my official advisor, Chris Wojtan. People in the field know Chris Wojtan is a great scientist. But I was very lucky to be in a position where I could learn about him as a person too. In my opinion, he is a rare figure in academia who maintains a great sense of humanity and fairness.

I learned from him many things like how to turn ideas into tractable problems, how to present them appropriately depending on the audience—whether sharing with potential collaborators, explaining to experts, or entertaining a broader audience—and how to set reasonable goals in research. But that's not all. I was not a typical graphics student. In particular, I was trying to research many things even beyond graphics, which I think made it hard to advise me. Yet he celebrated my ambition and supported me in trying to make it come true, both scientifically and non-scientifically. This made me want to become an advisor like Chris who can commit to realizing someone's dreams come true.

I thank my former supervisors, Antti Herva at Remedy Entertainment and Toru Fujii at Nikon Corporation. They always tried to find a way for me to engage in work that I could be passionate about and enjoy, and backed me in doing a PhD.

I thank my coauthors, Peter Heiss-Synak, Fumiya Narita, Masafumi Yamamoto. I learned a lot through collaborating with them, especially the joy of sharing excitement, challenges, fun, and sometimes even pain.

I owe a lot to the members of the visual computing group at ISTA during my PhD—Peter, Manas, Alexei, Christian, Mickaël, Evgeny, Georg, Carmen, Cristina, Dian, Michaela, Mathilde, Samara, Mau, Camille, Arian, Ruslan, Malina, Stefan, Tomáš, Thomas, Martin, Eder, Ihor, Desmond, Benjamin, Dafne, and Christine. My kind of research was quite different from the group's main direction and I didn't actively join group activities like daily lunch. But I never felt alone or isolated. I always felt cared, liked, supported, respected, and accepted. I loved those spontaneous moments of just dropping by someone's office and procrastinating over chats. With you guys, I could stay honest and be myself, and had a wonderful PhD life.

I thank the students in my cohort. I especially enjoyed the first PhD year with Alice and Anton, who shared the same office. I remember doing the analysis course homework together—not only did it help me learn analysis, but I also had so much fun with you.

I thank my friends at and outside ISTA, in Klosterneuburg, and across many countries. However, due to space limitations, it is unfortunately impossible to mention everyone by name. Thanks to them, my life has been so much richer.

I thank my (extended) family—Sadao, Hiroko, Akira, Dominique, Tomoaki, Miku, Mari, Naoki, Satomi, Mio, Yuri, Ichizo, Teiko—for their constant support. This includes their financial help. This made it possible for me to receive a university education and study overseas, which eventually enabled me to get a job abroad. And whenever we go back to Tokyo, they always welcome us and let us stay. Having a place I can always return to has given me a great feeling of comfort. This extends to Linda, Emil, and Luna, friends of both my family and me, who made us feel that Klosterneuburg, Austria is also our home.

I would also like to acknowledge my grandfather, Tsunehiro, who had a huge influence on the way I am now. He was almost the only person in my family involved in science or engineering and traveling around the world. His sense of art and science, as well as his honesty and love for everyone, shaped me.

I think my PhD was quite different from a typical one for several reasons. A major factor was my children, Ai Amelia and Yuu Noa, who made my PhD actually more challenging and full of hassle, while making my life happier, adding new colors and adventures. Science was just one of the things I learned during my PhD. I thank Ai Amelia and Yuu Noa, from whom I am learning a lot—about being a parent, being a human, a form of happiness, and about life.

And Shizuka Claire. For me, she is like a wizard. From her, I learned that magic is not a supernatural power. Magic is love for everyone and everything. I admire how your love always makes people happy and smile.

And thank you for always trusting me. I tend to be unplanned and spontaneously change the direction of life. You let me take on the adventure of leaving a stable job and moved to Austria with me while having a nine month old baby. Thanks to that, we are now in Klosterneuburg, Austria, having our time with people here and sharing moments as our children grow.

今この瞬間、今の日々をあなたと共有できていることが嬉しいです。いつも一緒に冒険してくれてありがとう。これからもあなたと楽しい時も苦しい時も人生の冒険を続けてゆきたいです。

2025年10月 定繁

Grant support Projects contained in this thesis were financially supported in part by the European Research Council with grants 1. ERC Consolidator Grant 101045083 CoDiNA, and 2. the European Union's Horizon 2020 research and innovation programme under grant agreement No. 638176.

About the Author

Before starting the PhD program at ISTA, Sadashige Ishida studied mathematics in Japan and Switzerland, then worked in the optics and game industries. At ISTA, he researched the union of mathematics, physics, and computer graphics arising from natural phenomena. He especially likes to describe and study dynamics through the lens of infinite-dimensional geometry. This spans purely theoretical to applied aspects of dynamics like fluids, optimal transport, deformation of shapes, and geometric structures behind these dynamics. Sadashige is also passionate about scientific and non-scientific education. An example is the foundation of the course *Geometric Mechanics* at ISTA. Aside from research, he enjoys interacting with nature, like plants and animals (including humans) around Kierling, Klosterneuburg where he lives with his family.

List of Collaborators and Publications

This thesis comprises the following articles:

1. Albert Chern and Sadashige Ishida. Area formula for spherical polygons via prequantization. *SIAM Journal on Applied Algebra and Geometry*, 8(3), 2024
2. Sadashige Ishida, Chris Wojtan, and Albert Chern. Hidden degrees of freedom in implicit vortex filaments. *ACM Transactions on Graphics*, 41(6), 2022
3. Albert Chern and Sadashige Ishida. Implicit representations of codimension-2 submanifolds and their prequantum structures. *Preprint. arXiv:2507.11727*, 2025
4. Martin Bauer, Sadashige Ishida, and Peter W. Michor. Symplectic structures on the space of space curves. *Preprint. arXiv:2407.19908*, 2024

The following works were also produced during the author's PhD studies:

5. Sadashige Ishida and Hugo Lavenant. Quantitative convergence of a discretization of dynamic optimal transport using the dual formulation. *Foundations of Computational Mathematics*, 2024
6. Sadashige Ishida*, Peter Synak*, Fumiya Narita, Toshiya Hachisuka, and Chris Wojtan. A model for soap film dynamics with evolving thickness. *ACM Transactions on Graphics*, 39(4), 2020

I am the first author of the article 2 and a joint first author with Peter Synak (now Peter Heiss-Synak) of the article 6, both published at ACM Transactions on Graphics. For the other articles, authors are listed alphabetically, following the convention of mathematics journals.

Table of Contents

Abstract	vii
Acknowledgements	viii
About the Author	xiii
List of Collaborators and Publications	xiv
Table of Contents	xv
1 Introduction	1
1.1 Background	1
1.2 Research projects	3
2 Preliminary: Symplectic geometry on the space of codimension-2 shapes	9
2.1 Diffeomorphism action on currents	10
2.2 The space of codimension-2 shapes	13
2.3 Symplectic structure on the codimension-2 shape space	15
2.4 Appendix: Liouville form via the tilde calculus	24
3 Implicit representations of codimension-2 submanifolds	25
3.1 Introduction	25
3.2 Implicit representations of codimension-2 shapes	26
3.3 Geometry of fibers	28
4 Prequantum structure on the space of implicit representations	33
4.1 Overview	33
4.2 Liouville form in the prequantum sense	35
4.3 Prequantum structure	42
4.4 Additional structures	47
4.5 Appendix: Curvature as monodromy	51
5 Hidden degrees of freedom in implicit vortex filaments	53
5.1 Introduction	54
5.2 Related work	55
5.3 Representations for evolving curves	56
5.4 Untwisted Clebsch variables and non-swirling dynamics	60
5.5 Algorithms	67
5.6 Applications	69
5.7 Conclusion and future work	79

6	Symplectic structures on the space of space curves	81
6.1	Introduction	81
6.2	Liouville structures and (pre)symplectic structures	84
6.3	Symplectic structures induced by conformal factors	90
6.4	Symplectic structures induced by length weighted metrics	98
6.5	Presymplectic structures induced by curvature weighted Riemannian metrics	104
6.6	Numerical illustrations	107
6.7	Appendix: Approach via an almost complex structure	109
7	Area formula for spherical polygons via prequantization	115
7.1	Introduction	115
7.2	Area of a spherical polygon	117
7.3	The area formula via the Hopf fibration	120
7.4	Derivation of the classical formula by $SO(3)$ as a prequantum bundle . . .	122
7.5	Numerical examples	123
7.6	Appendix: Prequantum bundles \mathbb{S}^3 and $SO(3)$ over \mathbb{S}^2	127
8	Postface: Mathematics, physics, and computer graphics	129
	Bibliography	133

Introduction

This thesis explores geometric structures and dynamics on the codimension-2 shape space, with emphasis on their implications for physics.

Our investigation spans theoretical studies of infinite-dimensional symplectic and prequantum geometry to numerical computation of the time evolution of shapes.

1.1 Background

Codimensional shapes and physics The history of much of physical science is the history of the study of shapes. Humans have long sought to describe physical phenomena in terms of the deformation of shapes. For example, some models—mathematical descriptions such as equations or rules that aim to replicate nature—express the motion of soap bubbles or biological cell membranes as the time evolution of surfaces in 3D. These surfaces are called codimension-1 because they are two-dimensional objects living in three-dimensional space; that is, the difference between the shape's dimension and the ambient space's dimension is one.



Figure 1.1: Evolution of soap bubbles modeled as deformation of surfaces.

Physical phenomena also involve shapes of other codimensions. Consider codimension-0, for instance: the motion of water droplets or elastic bodies that occupy a full 3D volume, so the codimension is zero. On the other hand, codimension-2 appears in situations such as electrically charged particles moving on a surface or tiny balls rolling on a billiard table—here the objects are zero-dimensional points moving within a 2D domain.

Hydrodynamics, the science of the motions of fluids, namely liquids and gases, provides particularly interesting examples involving codimension-1 and codimension-2 shapes. A fluid state is often described as a velocity field over a 3D region like a box. A quantity called

vorticity measures how much the fluid velocity is swirling locally. In common situations,¹ the velocity field can be recovered from the vorticity. Mathematically, this is expressed by the relation $d\eta = \omega$, where η is the velocity field represented as a differential 1-form, and ω is a differential 2-form representing vorticity. Hence, knowing how vorticity evolves over time allows us to track how the entire fluid state changes.

This draws our attention back to codimensional shapes. Vorticity often concentrates around “thin” codimension-1 or codimension-2 objects like surfaces or curves in 3D. This observation led to idealized models known as vortex sheets and vortex filaments—localizations of vorticity on surfaces and space curves, respectively. Since velocity can be recovered from vorticity, tracking the evolution of these shapes effectively tracks the evolution of the whole fluid state, illustrated as in Figure 1.2. This shows how understanding the time evolution of shapes, especially those of various codimensions, offers descriptions of physical phenomena.



Figure 1.2: Vorticity of a jet of fluid represented as a collection of space curves (left), inducing a velocity field which can transport smoke density (right). Simulation and visualization were performed using a method presented in Chapter 5.

Among these, codimension-2 shapes are significant in symplectic geometry, a framework for describing dynamics geometrically. I explain this now.

Symplectic geometry and physics Symplectic geometry is a playground for Hamiltonian systems, which are descriptions of how physical objects evolve over time while conserving “energy”. Here, energy—also called the Hamiltonian—is a quantity assigned to the system that we expect to remain constant in time, like kinetic energy, potential energy, or something more abstract, depending on the type of phenomenon we want to describe. Rather than tracking the motion of objects directly, symplectic geometry views such evolution as a path in a static space of possible physical states.

A key ingredient in this framework is a geometric object called a *symplectic structure*. It takes the Hamiltonian as input and outputs a dynamical system on the state space.

Mathematically, this idea is formulated as follows. A symplectic manifold X is a manifold equipped with a closed and non-degenerate 2-form ω called a symplectic structure. Once we give the Hamiltonian H as a function on X , the symplectic structure ω generates a vector field V_H on X through the relation $dH = \iota_{V_H}\omega$. Then the Hamiltonian system is the flow along the vector field V_H , that is, solutions to the differential equation $\partial_t x(t) = V_H(x(t))$ where x represents a state of the object as a point of X . Then, the time evolution is a path on which the Hamiltonian is constant.

The symplectic manifold X here typically represents the collection of all the possible states of a physical phenomenon. For example, if we consider particles moving in a domain M , then the

¹Technically, when the first homology group of the ambient space is trivial.

space X is the cotangent bundle T^*M^n , which contains all possible positions and momenta of n particles. In quantum mechanics, it is often the space of wave functions, representing all possible quantum states. These spaces are known to be symplectic manifolds, on which phenomena like celestial mechanics and the Schrödinger equation are modeled as Hamiltonian systems.

In this way, if the space of possible states is a symplectic manifold, we can model dynamics as Hamiltonian systems. This setup offers further advantages. For example, the Hamiltonian is often not the only constant in motion. If the system has symmetries, the symplectic structure reveals other conserved quantities. For instance, if the Hamiltonian is invariant under rotations, a quantity often called angular momentum remains constant in time. This symmetry-to-conservation link is known as Noether's theorem.

The codimension-2 shape space as an infinite-dimensional symplectic manifold

Earlier, we saw that many physical phenomena can be phrased in terms of the deformation of codimensional shapes. Among these, the case of codimension-2 stands out, as the space of codimension-2 shapes is an infinite-dimensional symplectic manifold, equipped with the so-called Marsden–Weinstein (MW) symplectic structure. We postpone the precise definitions of these notions to the next chapter.

Many phenomena in fields such as fluid dynamics and elasticity can be modeled as Hamiltonian systems with respect to the MW symplectic structure. Important examples include vortex filaments and their localized approximation called the binormal equation. Using the framework of symplectic geometry, we can study properties of these dynamics, such as conserved quantities, integrability, periodic orbits, and connections between seemingly unrelated dynamics.

This rich interplay motivates the present thesis to explore the geometric structures and dynamics on the codimension-2 shape space. By doing so from several angles, the author hopes to offer a glimpse into the beautiful bridge between shapes and dynamics.

1.2 Research projects

1.2.1 Implicit representations

The first major part of this thesis is concerned with *implicit representations* of shapes.

A submanifold (mostly synonymous with shape in this thesis ²) in an ambient space M is often described *explicitly* as an embedding of a manifold S into M , modulo reparametrizations.

An alternative is the *implicit* representation. For example, a codimension-1 shape, like a surface in \mathbb{R}^3 , can be represented as a preimage of a function, often called the level set function. This representation naturally extends to higher codimensions using multiple functions. When the codimension is 2, submanifolds can be expressed as the zero sets of complex-valued functions (Figure 1.3).

The implicit representation for each codimension-2 submanifold is not unique, as multiple (in fact, infinitely many) complex-valued functions can share the same zero set. The following two projects investigate this redundancy of implicit representations.

²Whenever unnecessary, we stay imprecise about whether a submanifold or shape may be an immersion or must be an embedding, and whether it is parametrized or unparametrized.



Figure 1.3: Linked curves, known as the Hopf link (left) and the level sets of a complex function $\psi = a + ib$ consisting of real-valued functions a and b (right). The blue and red surfaces are zero level sets of a and b respectively. Their intersection, the zero level set of ψ , agrees with the curves.

Project 1. Implicit representations of codimension-2 shapes and their prequantum structures

This part is based on the article [CI25].

Implicit representations via complex functions have been widely applied to study the dynamics of submanifolds. However, the geometry of the space of these representations is largely unexplored. In this project, we investigate its geometric structures, focusing on symplectic geometric aspects.

The non-uniqueness of implicit representations for each shape makes the implicit shape space a fiber bundle over the space of codimension-2 submanifolds. Our main contribution is showing that a particular quotient space of this fiber bundle forms a prequantum bundle.

Roughly speaking, a prequantum bundle is a specific kind of fiber bundle over a base symplectic manifold. Each fiber over a point carries the additional information of quantum phase, in such a way that Hamiltonian flows on the base manifold are naturally and uniquely lifted onto this bundle. In this way, a prequantum bundle is a sandbox for quantum mechanics, aligned with the base symplectic manifold as a sandbox for classical mechanics.

We do not, however, perform any quantum mechanical analysis using this framework. Rather, we focus on the geometric insights offered by our prequantum bundle. As illustrated in Figure 1.4, we present a new interpretation of the Marsden–Weinstein form:

Each complex function ψ representing a codimension-2 shape γ carries phase information $\phi = \psi/|\psi| \in \mathbb{S}^1$. The level sets $\{\phi^{-1}(s)\}_{s \in \mathbb{S}^1}$ of the phase function ϕ define a family of hypersurfaces in the ambient space, each bounded by γ . Any motion of the base shape γ induces motions of these hypersurfaces, which sweep out volumetric domains. The MW form arises as the curvature of a connection of our prequantum bundle, which measures the average volume swept by these hypersurfaces.

In the limiting case where the phase of the complex function is constant everywhere except for a 2π jump across a single hypersurface, this result reduces to the swept volume of that hypersurface, without explicit reference to the complex function.

This prequantum viewpoint thus unifies explicit and implicit descriptions of codimension-2 shapes.

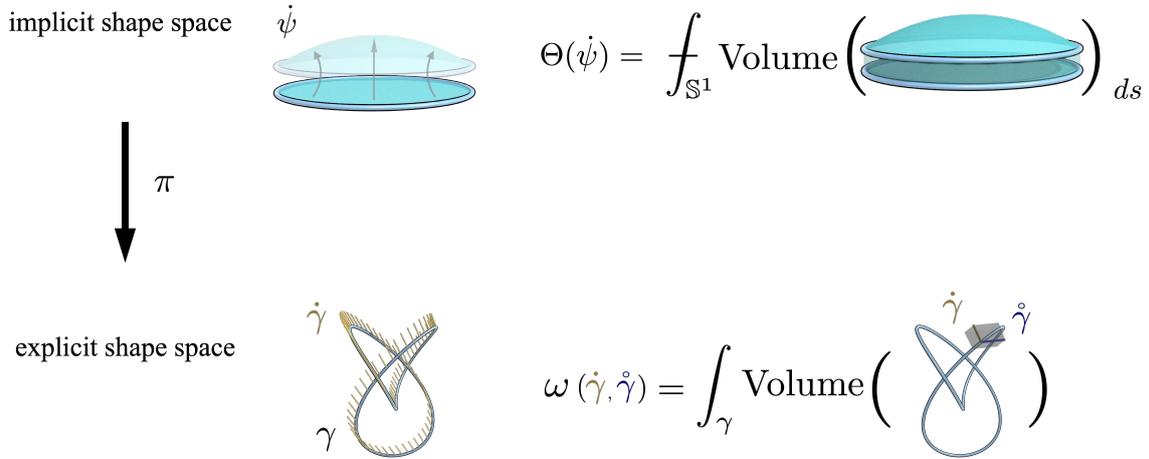


Figure 1.4: A schematic image of our prequantum bundle. The fiber bundle is the space of implicit representations (top row) and the base space, referred to as the explicit shape space, is the space of codimension-2 shapes (bottom row). The tangent vector $\dot{\psi}$ at each implicit representation ψ can be visualized as the deformation of the phase hypersurface $\phi^{-1}(s)$, and the connection form Θ measures the infinitesimal swept volumes of these phase hypersurfaces, averaged over all the phases. The tangent vector $\dot{\gamma}$ at each codimension-2 shape γ is a vector field on γ , representing a velocity of γ . The MW symplectic form ω measures the volume spanned by the velocity fields $\dot{\gamma}, \dot{\gamma}$ and the tangent of γ , integrated over γ . The 2-form ω is the curvature of the connection form Θ , that is, $\pi^*\omega = d\Theta$.

Project 2. Hidden degrees of freedom in implicit vortex filaments

This part is based on the article [IWC22].

The deformation of space curves is rooted in several fields like physics, biology, and mathematics, and has been studied both theoretically and numerically.

In this project, we develop a numerical method for space curve dynamics using implicit representations. Instead of simulating space curves explicitly, we compute the time evolution of implicit representations.

Our approach, in particular, exploits the non-uniqueness of these implicit representations in both their configurations and dynamics. As noted, multiple complex functions share the same zero level set. Such non-uniqueness also exists in the dynamics. The time evolution of curves represented by a complex function ψ can be described by the transport of ψ along a vector field, and the choice of the vector field yielding the same motion of the zero set is not unique. In fact, there are infinite degrees of freedom within the collection of both possible complex functions and vector fields.³

Within these redundancies, we can make a specific choice for a particular purpose, such as improving the robustness of numerical simulation. Specifically, to handle chaotic dynamics like vortex filaments in fluid dynamics, we introduce *untwisted Clebsch variables* and *non-swirling dynamics*. They successfully reduce the numerical instability, taming the twisting modes around the filaments.

³Strictly speaking, the infinite degrees of freedom become finite when we perform a finite-dimensional approximation, a.k.a. discretization, for numerical purposes.

Additionally, the level set description by implicit representations inherently supports topological changes of curves. Built on these features, our resulting method stably simulates the dynamics of vortex filaments that split and merge (Figure 1.2 and Figure 1.5).

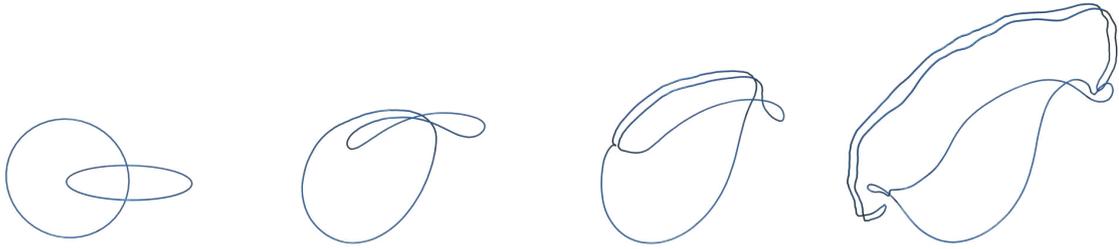


Figure 1.5: Time evolution of vortex filaments undergoing topological changes (left to right), simulated using our method.

In this project, we do not restrict ourselves to Hamiltonian systems. Our method can accommodate any first-order dynamics in time. As an example, we simulate the curve shortening flow, which is not a Hamiltonian flow (energy preserving flow) of the Marsden–Weinstein structure but rather a gradient flow (energy descent flow) with respect to the L^2 -Riemannian metric.

1.2.2 No more symplectic structures?

We now go back to explicit representations of shapes and focus on the space of space curves, which is the simplest infinite-dimensional instance of the codimension-2 shape space.

The Marsden–Weinstein (MW) structure is known to be a canonical symplectic structure on this space. But are there no other symplectic structures? This following project answers this question.

Project 3. Symplectic structures on the space of space curves

This part is based on the article [BIM24].

We derive new symplectic structures on the space of space curves. These new structures generalize the MW structure, which was the only previously studied symplectic structure on this space.

Our approach builds on two key ingredients. First, the MW form ω admits a Liouville 1-form η , *i.e.*, $d\eta = \omega$, and this η can be expressed in terms of the standard L^2 Riemannian metric on the space of space curves. Second, in mathematical shape analysis, many alternative Riemannian metrics on this space have been developed by integrating a suitable operator on the tangent bundle of the space.

Motivated by this, we construct new 1-forms by modifying the standard Liouville form with such operators. Taking their exterior derivatives yields closed 2-forms, and we verify that these forms are non-degenerate, thus defining symplectic structures.

We then derive Hamiltonian systems induced by these new symplectic structures, as different symplectic structures ω' induce distinct Hamiltonian flows V_H' from the same input Hamiltonian function H , according to the relation $dH = \iota_{V_H'}\omega'$ (Figure 1.6). We also numerically illustrate simple examples.

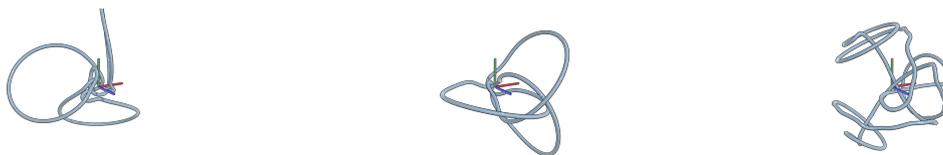


Figure 1.6: Snapshots of a space curve evolving under Hamiltonian flows of the same Hamiltonian function, but induced by the standard MW structure (left) and different symplectic structures (middle and right), exhibiting distinct behaviors.

1.2.3 Prequantum geometry for polygons on the sphere

Symplectic and prequantum geometry are frameworks for studying dynamics, but they can also be useful for purely geometric problems. This project is such an instance. Unlike the above projects which concerned dynamics and infinite-dimensional manifolds, the following project deals with a purely geometric, finite-dimensional problem.

Project 4: Area formula for spherical polygons via prequantization

This part is based on the article [BIM24].

In this project, we derive a new formula for the area of spherical polygons via prequantization. A spherical polygon is a finite sequence of ordered points on \mathbb{S}^2 , connected by geodesics. Calculating the solid angle of the region enclosed by such a polygon has practical applications in fields like fluid dynamics, light transport theory, and geography (Figure 1.7).

The areas of spherical polygons are often computed using a formula based on the Gauss-Bonnet theorem. However, this formula requires the non-degeneracy assumption, meaning that no two consecutive points can lie on the same location. This constraint renders the formula unusable or numerically unstable in certain situations.

Our new formula mimics Green's formula, which is not directly applicable on \mathbb{S}^2 because the standard area form is not exact. We circumvent this issue using a prequantum bundle. We lift the area form onto the bundle, where the resulting 2-form becomes exact. This converts an area integral on \mathbb{S}^2 into a line integral on the bundle.

Through this procedure, we derive a formula which is applicable to a wider range of degenerate spherical curves and polygons.

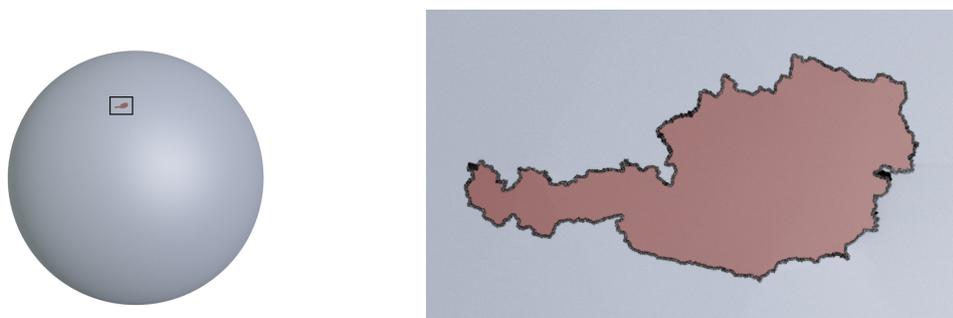


Figure 1.7: A spherical polygon representing the shape of Austria (left), and a zoomed-in view (right). Our formula can robustly compute the area of such an intricate shape.

1.2.4 Organization of the thesis

Chapters 2–4 are based on the article [CI25].

Chapter 2 reviews the symplectic geometry of codimension-2 shape spaces using the language of currents from geometric measure theory. This allows a unified description of results for both the classical setting with explicit representations and the implicit representations developed in the subsequent chapters.

Chapter 3 introduces implicit representations of codimension-2 shapes using a class of complex-valued functions. We study the geometry of the space of such representations as a fiber bundle over the space of explicit representations. In particular, we show that over each codimension-2 shape, the fiber may have multiple connected components indexed by $H_{dR}^1(M, \mathbb{Z})$, the first integral de Rham cohomology group of the ambient manifold.

Chapter 4 constructs a prequantum structure over the codimension-2 shape space. We show that a quotient bundle of the space of implicit representations forms a prequantum G -bundle with structure group $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$. As a corollary, we obtain a new geometric interpretation of the Marsden–Weinstein structure.

Chapter 5 is the content of the article [IWC22]. Using implicit representations, we develop a numerical method for simulating space curve dynamics. We leverage the infinite degrees of freedom, hidden in implicit representations of both the configuration and the dynamics. This results in a simulation method that can handle otherwise numerically instable scenarios like vortex filaments from hydrodynamics, while automatically processing topological changes.

Chapter 6 is the content of the article [BIM24]. We introduce new symplectic structures on the space of space curves by modifying the Liouville 1-form of the Marsden–Weinstein structure, inspired by a recent trend in shape analysis. We also derive the corresponding Hamiltonian vector fields for several Hamiltonian functions with respect to these new symplectic structures, and numerically illustrate a few examples.

Chapter 7 is the content of the article [CI24]. We derive a formula for the signed area of a spherical polygon, which is a piecewise geodesic on \mathbb{S}^2 . For this purpose, we use a prequantum bundle. This gives rise to a version of Green’s theorem that translates the area integral on the base manifold into a line integral along a lifted perimeter in the bundle.

Chapter 8 is independent of the rest of the thesis. I present my thoughts on mathematics, physics, and computer graphics, which I have developed over the past few years.

Dependencies of chapters Chapter 2 is a prerequisite for Chapter 3 and Chapter 4. The remaining chapters Chapter 5, Chapter 6, and Chapter 7 are written in a self-contained manner, so readers can jump directly to any of them. For a deeper understanding, note that Chapter 5 is an application of implicit representations introduced in Chapter 3, and Chapter 6 is a generalization of the Marsden–Weinstein structure explained in Chapter 2.

Preliminary: Symplectic geometry on the space of codimension-2 shapes

Chapter 2, Chapter 3, and Chapter 4 are based on the article:

Albert Chern and Sadashige Ishida. Implicit representations of codimension-2 submanifolds and their prequantum structures. *Preprint. arXiv:2507.11727*, 2025

In this chapter, we review basic results on the symplectic geometry of codimension-2 shape space, as the topics in this thesis revolve around this subject. We first lay out the preliminary settings for the shape space of codimension-2 submanifolds, and then explain that this space is a symplectic manifold equipped with the so-called Marsden–Weinstein symplectic structure.

We formulate these known results using currents from geometric measure theory, which are the continuous dual of differential forms. Currents allow a unified description of both the classical results reviewed here and the new results for implicit representations that we introduce and study in Chapter 3 and Chapter 4.

Alongside, we extend a previous result. The existence of a Liouville form η for the MW form ω (i.e., 1-form η such that $d\eta = \omega$) was proven only for space curves. We extend this to a general closed and oriented codimension-2 manifold, embedded in an ambient manifold of arbitrary dimension, equipped with an exact volume form.

Dependencies of the other chapters on this chapter The main purpose of this chapter is to provide an overview of the codimension-2 shape space as an infinite-dimensional symplectic manifold, formulated using the language of currents. This builds a foundation for Chapter 3 and Chapter 4, where implicit representations of shapes are introduced and studied.

On the other hand, Chapter 5, Chapter 6, and Chapter 7 are written in a self-contained manner and do not strictly rely on the results and the descriptions in this chapter. Hence the readers interested in these chapters can directly jump to any of them.

Organization of the chapter In Section 2.1, we review currents from geometric measure theory and introduce the action of diffeomorphism groups on currents, as well as the Lie derivative as its infinitesimal generator.

In Section 2.2, we define the codimension-2 shape space as embeddings of a manifold into an ambient manifold modulo reparametrization. We do so in terms of diffeomorphism actions on currents, highlighting that the deformation of a shape can be seen as the transport of a current along a vector field.

In Section 2.3, we review the Marsden–Weinstein symplectic structure on the codimension-2 shape space. Additionally, we extend the previous result on the existence of a Liouville form for space curves to a more general setting in arbitrary dimension.

2.1 Diffeomorphism action on currents

We briefly review the notion of currents from geometric measure theory as a preliminary and define a diffeomorphism action on them. For readers interested in the general theory of currents, we refer to the literature [Fed14, dR84, Mor08]. Currents are a natural language for shapes embedded in an ambient manifold as they allow us to treat differential forms, submanifolds, and their generalizations in a unified manner. Using currents, we will describe relations between quantities defined on an ambient manifold and a submanifold (possibly with a boundary). This includes, for example, the flux of a vector field through a surface, and a superposition of the fluxes through infinitely many surfaces.

Let M be an oriented m -dimensional manifold and $\Omega_c^k(M)$ be compactly supported differential k -forms. We say k -currents $\mathcal{D}_k(M)$, also denoted by $\Omega_c^k(M)^*$, are linear functionals on $\Omega_c^k(M)$ that are continuous in the sense of distributions. For details, see [CSdR12, Alb06] for example.

We note that $\Omega^{m-k}(M) \subset \mathcal{D}_k(M)$ via

$$\langle h, \alpha \rangle := \int_M h \wedge \alpha, \quad h \in \Omega^{m-k}(M), \alpha \in \Omega_c^k(M). \quad (2.1.1)$$

The space of differential forms $\Omega^{m-k}(M)$ is strictly smaller than the space of current $\mathcal{D}_k(M)$, but $\mathcal{D}_k(M)$ is attained as the closure of $\Omega^{m-k}(M)$ with respect to the locally convex topology. Keeping this in mind, we will formally write $\int_M h \wedge \alpha$ also for h in $\mathcal{D}_k(M)$ but not in $\Omega^{m-k}(M)$ as if h was a $m-k$ form.

For $h \in \mathcal{D}_k(M)$, we define the boundary operator $\partial_k: \mathcal{D}_k(M) \rightarrow \mathcal{D}_{k-1}(M)$ by

$$\langle \partial_k h, \alpha \rangle := \langle h, d_{k-1} \alpha \rangle, \quad \forall \alpha \in \Omega_c^k(M)$$

using the exterior derivative $d_{k-1}: \Omega^{k-1}(M) \rightarrow \Omega^k(M)$. Unless necessary, we will simply write ∂ for ∂_k and d for d_k . Notice that ∂ works for $h \in \Omega^{m-k}(M) \subset \mathcal{D}_k(M)$ as the exterior derivative d up-to sign change with $\text{sgn} = (-1)^{m-k+1}$,

$$\begin{array}{ccc} \text{sgn } d : \Omega^{m-k}(M) & \longrightarrow & \Omega^{m-k+1}(M) \\ & \cap & \cap \\ \partial : \mathcal{D}_k(M) & \longrightarrow & \mathcal{D}_{k-1}(M). \end{array}$$

We say that the current homology, which is the dual of de-Rham cohomology, is defined by

$$H_k^{CR}(M) := Z_k^{CR}(M) / B_k^{CR}(M)$$

where $Z_k^{CR}(M) = \ker \partial_k$ and $B_k^{CR}(M) = \text{im } \partial_{k+1}$.

2.1.1 de Rham-Dirac currents

We call a special class of currents de Rham-Dirac currents (or simply de Rham currents). These currents extend the Dirac measures for points (codimension- m geometry) to other codimensions. For a k -dimensional submanifold Σ of M possibly with a boundary, we define a k -de Rham current $\delta_\Sigma \in \mathcal{D}_k(M)$ by

$$\langle \delta_\Sigma, \alpha \rangle = \int_\Sigma \alpha, \quad \alpha \in \Omega^k(M).$$

We denote the collection of de Rham- k current by $\mathcal{D}_k^{dR}(M)$.

The boundary operators ∂ for the de Rham currents and for submanifolds commute with the mapping of submanifolds to currents *i.e.*, $\partial\delta_\Sigma = \delta_{\partial\Sigma}$ as

$$\langle \partial\delta_\Sigma, \alpha \rangle = \int_\Sigma d\alpha = \int_{\partial\Sigma} \alpha = \langle \delta_{\partial\Sigma}, \alpha \rangle, \quad \forall \alpha \in \Omega^k(M).$$

We will also formally write this as

$$\int_M \delta_{\partial\Sigma} \wedge \alpha = \int_M \delta_\Sigma \wedge d\alpha$$

with the mind of (2.1.1).

Remark 2.1.1. de Rham currents can be defined from rectifiable sets on M , which form a much broader class containing submanifolds and singular chains of M [Mor08]. We focus on the ones induced from submanifolds as it suffices for our purpose.

2.1.2 Orbits under diffeomorphism groups

We represent unparametrized shapes by currents and describe the deformation of shapes in terms of diffeomorphism actions. To do so, we define the action of diffeomorphism groups on currents and the Lie derivative as its infinitesimal generator.

Let us denote by $\text{Diff}(M)$ the space of diffeomorphisms over M . For each $f \in \text{Diff}(M)$, we define the action $f_*: \mathcal{D}_k(M) \rightarrow \mathcal{D}_k(M)$ by pushforward as the adjoint of pullback $f^*: \Omega^k(M) \rightarrow \Omega^k(M)$,

$$\langle f_*h, \alpha \rangle := \langle h, f^*\alpha \rangle, \quad h \in \mathcal{D}_k(M), \alpha \in \Omega^k(M).$$

For $h \in \Omega^{m-k}(M) \subset \mathcal{D}_k(M)$, the pushforward action f_* on h as a current agrees with the pullback action f^{-1*} on h as a differential form. For a de-Rham current $\delta_\Sigma \in \mathcal{D}_k^{dR}(M)$, we have $f_*\delta_\Sigma = \delta_{f \circ \Sigma}$. Note also that f_* and ∂ commute similarly to the commutativity between d and f^* for differential forms.

On an orbit \mathcal{O} under the $\text{Diff}(M)$ -action by pushforward, the tangent space is ¹

$$T_h\mathcal{O}(M) = \{-\mathcal{L}_v h \mid v \in \text{diff}(M)\}. \quad (2.1.2)$$

¹When the orbit \mathcal{O} does not carry a manifold structure, the tangent space at a current h is defined only formally as the image of fundamental vector field mapping of the $\text{Diff}(M)$ -action. If h is a de Rham current of an embedded submanifold, \mathcal{O} is a manifold.

Here, $\text{diff}(M)$ is the Lie algebra of $\text{Diff}(M)$, which is the space of smooth vector fields on M , and we say that Lie derivative $\mathcal{L}_v: \mathcal{D}(M) \rightarrow \mathcal{D}(M)$ for currents is the skew-adjoint of the Lie derivative for differential forms,

$$\langle -\mathcal{L}_v h, \alpha \rangle := \langle h, \mathcal{L}_v \alpha \rangle, \quad \alpha \in \Omega^k(M).$$

The minus sign of $-\mathcal{L}_v h$ in (2.1.2) expresses the feeling of *advection* of h along the velocity field v by the formal transport equation $\partial_t h + \mathcal{L}_v h = 0$. Similarly to differential forms, we have $\mathcal{L}_{[u,v]} h = \mathcal{L}_u \mathcal{L}_v h - \mathcal{L}_v \mathcal{L}_u h$ for currents, which can be verified by direct computation.

Lastly, we have a current analogy of the relation $\mathcal{L}_u \alpha = \frac{d}{dt} \Big|_{t=0} \text{Fl}_u^t * \alpha$ for differential forms where $\text{Fl}_u^t \in \text{Diff}(M)$ is the time- t flow map of $u \in \text{diff}(M)$ defined as the solution to the ODE,

$$\begin{aligned} \partial_t(\text{Fl}_u^t(x)) &= u(\text{Fl}_u^t(x)), \quad x \in M \\ \text{Fl}_u^0 &= \text{id}_M. \end{aligned}$$

For a current h , we have $-\mathcal{L}_u h = \frac{d}{dt} \Big|_{t=0} \text{Fl}_{u*}^t h$ as

$$\frac{d}{dt} \Big|_{t=0} \langle \text{Fl}_{u*}^t h, \alpha \rangle = \frac{d}{dt} \Big|_{t=0} \langle h, \text{Fl}_u^t * \alpha \rangle = \langle h, \mathcal{L}_u \alpha \rangle = \langle -\mathcal{L}_u h, \alpha \rangle, \quad \forall \alpha \in \Omega^k(M).$$

2.1.3 Currents as functionals on the space of vector fields

In later sections, we will consider fluxes of vector fields through possibly infinitely many hypersurfaces. Currents provide a natural description for this.

Let $\text{diff}(M)$ and $\text{sdiff}(M)$ denote the spaces of smooth vector fields and divergence free-vector fields on M equipped with a volume form μ . They are Lie algebras of $\text{Diff}(M)$ and $\text{SDiff}(M)$, the groups of diffeomorphisms and volume-preserving diffeomorphisms on M . Let us also define $\text{exdiff}(M)$, a subspace of $\text{sdiff}(M)$, consisting of *exact divergence-free* vector fields, that is, $\iota_u \mu$ is exact.

Definition 2.1.2 (Flux for $m-1$ currents and exact $m-2$ currents). For $h \in \mathcal{D}_{m-1}(M)$, we say that the flux of a vector field $u \in \text{diff}(M)$ through h is

$$\text{Flux}_{m-1}(h, u) := \langle h, \iota_u \mu \rangle.$$

For $\beta \in B_{m-2}^{CR}(M)$, we say that the flux of $v \in \text{exdiff}(M)$ through β is

$$\text{Flux}_{m-2}^{\text{ex}}(\beta, v) := \text{Flux}_{m-1}(h, v)$$

with any $h \in \partial^{-1} \beta \subset \mathcal{D}_{m-1}(M)$.

By design $\text{Flux}_{m-2}^{\text{ex}}$ is independent of the choice of h in $\partial^{-1} \beta$. Via the notion of flux, we can now regard $m-1$ and exact $m-2$ currents as linear functions on $\text{diff}(M)$ and $\text{exdiff}(M)$.

For a de Rham current $\delta_\Sigma \in \mathcal{D}_{m-1}^{dR}(M)$, $\text{Flux}_{m-1}(\delta_\Sigma, v)$ is indeed the flux of v through Σ . But this notion of flux does not require h to be a de Rham current. For example, h may be a formal infinite sum of de Rham currents representing infinitely many hypersurfaces. This observation plays an important role when we consider a Liouville form for a symplectic structure on the space of implicit representations in Section 4.2.

2.2 The space of codimension-2 shapes

This section provides a preliminary review of the space of *explicit representations* for codimension-2 shapes, namely, embeddings of an n -dimensional manifold S into an $(n + 2)$ -dimensional ambient manifold M . We revisit earlier work on the canonical symplectic structure on this space, the Marsden–Weinstein (MW) structure, as studied in [MW83, HV03, Tab17, PCK⁺19]. Our presentation is framed in terms of currents, which will play a central role in the next section when we introduce implicit representations (Chapter 3) and relate them to the explicit framework.

In addition, we obtain new results (Theorem 2.3.3, Theorem 2.3.6): the MW structure on the codimension-2 shape space is exact if the volume form μ of the ambient manifold is exact. This extends earlier results for closed curves in \mathbb{R}^3 [Tab17, PCK⁺19] to a broader class of submanifolds in arbitrary dimensions.

We conclude this section by considering the shape space of all codimension-2 submanifolds. That is, we allow the submanifold S to range over all oriented and closed n -manifolds, rather than fixing one S in advance. This broader viewpoint naturally leads to the introduction of implicit representations in the next chapter.

2.2.1 Spaces of parametrized and unparametrized codimension-2 submanifolds

We begin with the case where the dimension of the ambient manifold is greater than 2, postponing remarks on the 2-dimensional case to the end of this section. Let $m > 2$, and let M be an m -dimensional manifold equipped with a volume form μ . Let S be a closed (i.e., compact and without boundary), oriented manifold of dimension $n = m - 2$.

Consider the space of smooth embeddings

$$\text{Emb}(S, M) := \{\tilde{\gamma} \in C^\infty(S, M) \mid \text{rank}(d\tilde{\gamma}) = n, \tilde{\gamma}(s) = \tilde{\gamma}(s') \Rightarrow s = s'\},$$

which is an infinite dimensional manifold with the Fréchet topology [BBM14, Mic19]. Its tangent space at each $\tilde{\gamma} \in \text{Emb}(S, M)$ is given by sections of the pullback bundle:

$$T_{\tilde{\gamma}} \text{Emb}(S, M) = \Gamma(\tilde{\gamma}^*TM).$$

That is, a tangent vector $\dot{\tilde{\gamma}} \in T_{\tilde{\gamma}} \text{Emb}(S, M)$ assigns to each point $s \in S$ a vector $\dot{\tilde{\gamma}}(s) \in T_{\tilde{\gamma}(s)}M$. When $M = \mathbb{R}^m$, the tangent space $T_{\tilde{\gamma}} \text{Emb}(S, \mathbb{R}^m)$ is identified with $C^\infty(S, \mathbb{R}^m)$.

On $\text{Emb}(S, M)$ we define a right action of the orientation-preserving diffeomorphism group $\text{Diff}^+(S)$ representing reparametrizations. Taking the quotient by this action, we obtain the *shape space of unparametrized shapes*:

$$\text{UEmb}(S, M) := \text{Emb}(S, M) / \text{Diff}^+(S).$$

The space $\text{UEmb}(S, M)$ is also referred to as the *nonlinear Grassmannian of type S* [HV03].

Note that $\text{Diff}^+(S)$ may have multiple connected components. For example, if $S = \bigsqcup_{i=1}^{k_1} \mathbb{S}^2 \sqcup \bigsqcup_{i=1}^{k_2} \mathbb{T}^2$, then $\text{Diff}^+(S)$ is isomorphic to the semi-direct product of $\text{Diff}^+(\mathbb{S}^2)^{k_1} \times \text{Diff}^+(\mathbb{T}^2)^{k_2}$ and the symmetric groups $\text{Sym}(k_1) \times \text{Sym}(k_2)$. The operation of permuting identical copies of \mathbb{S}^2 or \mathbb{T}^2 is included in $\text{Diff}^+(S)$.

The shape space $\text{UEmb}(S, M)$ is an infinite-dimensional manifold [BBM14]. In what follows, we denote an element of $\text{UEmb}(S, M)$ by γ and any representative of γ by $\tilde{\gamma} \in \text{Emb}(S, M)$.

With the fibration $\pi: \text{Emb}(S, M) \rightarrow \text{UEmb}(S, M)$ the tangent space is written as

$$T_{\pi(\tilde{\gamma})}\text{UEmb}(S, M) = d\pi|_{\tilde{\gamma}}(T_{\tilde{\gamma}}\text{Emb}(S, M)).$$

Note that $\ker d\pi|_{\tilde{\gamma}} = d\tilde{\gamma}(\text{diff}(S))$ where $\text{diff}(S)$ is the space of smooth vector fields on S , which is the Lie algebra of $\text{Diff}^+(S)$. Hence $\ker d\pi$ consists of the components of tangent vectors which do not change *the shape* of γ .

Example 2.2.1 (Tangent space at a unparametrized space curve). When $S = \mathbb{S}^1$ and $M = \mathbb{R}^3$ equipped with the standard Euclidean metric, we have $\ker d\pi|_{\tilde{\gamma}} = \{a\partial_s\tilde{\gamma} \mid a \in C^\infty(S)\}$. Hence $T_{\tilde{\gamma}}\text{UEmb}(S, M) = \{\dot{\tilde{\gamma}} + d\tilde{\gamma}\text{diff}(S) \mid \dot{\tilde{\gamma}} \in T_{\tilde{\gamma}}\text{Emb}(S, M)\}$ is identified with $\{\dot{\tilde{\gamma}}: S \rightarrow \mathbb{R}^3 \mid \dot{\tilde{\gamma}}(s) \perp \partial_s\tilde{\gamma}(s), \forall s \in \mathbb{S}^1\}$, vector fields on $\tilde{\gamma}(S)$ that are everywhere perpendicular to the tangent vector $\partial_s\tilde{\gamma}$.

We now define a left action of $\text{Diff}(M)$ on $\text{UEmb}(S, M)$ by

$$f \triangleright \gamma = \pi(f \circ \tilde{\gamma}), \quad f \in \text{Diff}_0(M), \gamma \in \text{UEmb}(S, M)$$

where π is the projection $\pi: \text{Emb}(S, M) \rightarrow \text{UEmb}(S, M)$ and $\tilde{\gamma} \in \pi^{-1}\gamma$ is any parametrization of γ .

Each connected component of $\text{UEmb}(S, M)$ is an orbit of the action of $\text{Diff}_0(M)$, the connected component of $\text{Diff}(M)$ containing id_M . This is a classical result due to Thom (see, e.g., [Hir12]). Consequently, any tangent vector $\dot{\gamma} \in T_{\gamma}\text{UEmb}(S, M)$ can be written as $\dot{\gamma} = v \circ \gamma := d\pi|_{\tilde{\gamma}}(v \circ \tilde{\gamma})$ for some $v \in \text{diff}(M)$. With this in mind, we will restrict our attention to a single $\text{Diff}_0(M)$ -orbit in $\text{UEmb}(S, M)$, denoted by \mathcal{O} .

In fact, the action by the subgroup $\text{SDiff}_0(M) \subset \text{Diff}_0(M)$ of volume-preserving diffeomorphisms is transitive on each $\text{Diff}_0(M)$ -orbit in $\text{UEmb}(S, M)$ [HV03, Proposition 2]. Much of the theory in this section remains the same after restricting the diffeomorphism group $\text{Diff}_0(M)$ to $\text{SDiff}_0(M)$ and the space $\text{diff}(M)$ of vector fields to the space $\text{sdiff}(M)$ of divergence-free vector fields.

Remark 2.2.2 (Labeled and unlabeled shapes). There are two standard definitions of unparametrized shapes $\text{UEmb}(S, M)$. One defines $\text{UEmb}(S, M)$ as the quotient of $\text{Emb}(S, M)$ by the orientation-preserving diffeomorphism group $\text{Diff}^+(S)$, and the other uses $\text{Diff}_0(S)$, the connected component in $\text{Diff}(S)$ containing the identity map [BBM14].

The quotient by $\text{Diff}_0(S)$ retains labels on identical components of the same shape. For example, if $S = \mathbb{S}^1 \sqcup \mathbb{S}^1$, then two embeddings $\gamma_a = (\gamma_1, \gamma_2)$ and $\gamma_b = (\gamma_2, \gamma_1)$, with non-intersecting embeddings γ_1 and γ_2 of \mathbb{S}^1 , represent distinct elements in $\text{Emb}(S, M)/\text{Diff}_0(S)$ but the same element in $\text{Emb}(S, M)/\text{Diff}^+(S)$.

In this thesis, we consider *unlabeled shapes*, defined as $\text{Emb}(S, M)/\text{Diff}^+(S)$. This convention better aligns with the implicit representations of shapes introduced in Section 3.2 (See also Remark 3.2.1). However, we note that the theorems and propositions in this section hold under both definitions.

2.2.2 Shapes as de Rham currents

Since the shape space $\text{UEmb}(S, M)$ is a subset of the space of $(m - 2)$ -dimensional submanifolds in M , there is a natural injection into the space of de Rham currents:

$$\begin{aligned} \mathcal{I}: \text{UEmb}(S, M) &\rightarrow \mathcal{D}_{m-2}^{dR}(M) \\ \gamma &\mapsto \delta_\gamma. \end{aligned}$$

We denote the image of \mathcal{I} by $\mathcal{D}_{\text{UEmb}(S, M)}$, and for a $\text{Diff}_0(M)$ -orbit $\mathcal{O} \subset \text{UEmb}(S, M)$, we write $\mathcal{D}_\mathcal{O} := \mathcal{I}(\mathcal{O})$.

The tangent space at $\delta_\gamma \in \mathcal{D}_\mathcal{O}$ is given by the pushforward of the tangent space at γ :

$$T_{\delta_\gamma} \mathcal{O} = d\mathcal{I}|_\gamma (T_\gamma \mathcal{O}).$$

More explicitly,

$$d\mathcal{I}|_\gamma: \dot{\gamma} = v \circ \gamma \mapsto -\mathcal{L}_v \delta_\gamma, \quad \text{for some } v \in \text{diff}(M),$$

where \mathcal{L}_v denotes the Lie derivative in the sense of currents, as defined in Section 2.1.2. Therefore, the tangent space at δ_γ is

$$T_{\delta_\gamma} \mathcal{D}_\mathcal{O} = \{-\mathcal{L}_v \delta_\gamma \mid v \in \text{diff}(M)\}.$$

Note also that the injection \mathcal{I} commutes with the $\text{Diff}_0(M)$ -action in the sense that $f_* \mathcal{I} = \mathcal{I} \circ f$ for all $f \in \text{Diff}_0(M)$, as discussed in Section 2.1.2. In the remainder of the chapter, we will occasionally identify \mathcal{O} with $\mathcal{D}_\mathcal{O}$, and $\text{UEmb}(S, M)$ with $\mathcal{D}_{\text{UEmb}(S, M)}$, and their tangent bundles, without explicitly stating so.

2.3 Symplectic structure on the codimension-2 shape space

The space $\mathcal{O} \cong \mathcal{D}_\mathcal{O}$ is a *weak symplectic manifold*. That is, it is equipped with a closed and *weakly non-degenerate* 2-form. The Marsden–Weinstein (MW) form on \mathcal{O} is given by

$$\omega_{\delta_\gamma}(\dot{\delta}_\gamma, \mathring{\delta}_\gamma) = \langle \delta_\gamma, \iota_v \iota_u \mu \rangle = \int_\gamma \iota_v \iota_u \mu, \quad (2.3.1)$$

where $\dot{\delta}_\gamma = -\mathcal{L}_u \delta_\gamma$ and $\mathring{\delta}_\gamma = -\mathcal{L}_v \delta_\gamma$ for vector fields $u, v \in \Gamma(TM)$. When convenient, we also write $\omega_\gamma(\dot{\gamma}, \mathring{\gamma}) := \omega_{\delta_\gamma}(\dot{\delta}_\gamma, \mathring{\delta}_\gamma)$ for $\dot{\gamma} = u \circ \gamma$ and $\mathring{\gamma} = v \circ \gamma$.

The MW form ω is closed and weakly-nondegenerate in the sense that the associated flat operator

$$\begin{aligned} \flat^\omega: T\mathcal{O} &\rightarrow T^*\mathcal{O} \\ \dot{\gamma} &\mapsto \iota_{\dot{\gamma}} \omega \end{aligned}$$

is injective. In this chapter and the next few chapters, we refer to such weak symplectic forms simply as *symplectic*. For background on weak symplectic geometry, we refer the reader to [Mic84, Chapter VI] and [BIM24, Appendix A].

On the parametrized shape space $\tilde{\mathcal{O}} := \pi^{-1}(\mathcal{O})$, where $\pi: \text{Emb}(S, M) \rightarrow \text{UEmb}(S, M)$ is the natural projection, the pullback 2-form $\tilde{\omega} := \pi^*\omega \in \Omega^2(\tilde{\mathcal{O}})$ is written as

$$\tilde{\omega}_{\tilde{\gamma}}(\dot{\tilde{\gamma}}, \ddot{\tilde{\gamma}}) = \int_S \mu(\dot{\tilde{\gamma}}, \ddot{\tilde{\gamma}}, \partial_{s_1}\tilde{\gamma}, \dots, \partial_{s_n}\tilde{\gamma}) ds_1 \dots ds_n$$

in local coordinates $s = (s_1, \dots, s_n)$ on S . From this expression, we see that $\tilde{\omega}$ has a nontrivial kernel consisting of infinitesimal reparametrizations, that is, $\ker \tilde{\omega}|_{\tilde{\gamma}} = \ker d\pi|_{\tilde{\gamma}} = d\tilde{\gamma}(\text{diff}(S))$ at each $\tilde{\gamma} \in \tilde{\mathcal{O}}$. Thus, $\tilde{\omega}$ is not symplectic, but merely presymplectic *i.e.*, a closed 2-form.

Example 2.3.1 (MW structure on the space of space curves). On the space of space curves $\text{UEmb}(\mathbb{S}^1, \mathbb{R}^3)$ where \mathbb{R}^3 is equipped with the standard Euclidean volume form $\mu = dx_1 \wedge \dots \wedge dx_3$, the MW form takes the expression

$$\omega_{\gamma}(\dot{\gamma}, \ddot{\gamma}) = \int_{\mathbb{S}^1} \det(\partial_s \tilde{\gamma}, \dot{\tilde{\gamma}}, \ddot{\tilde{\gamma}}) ds,$$

where $\tilde{\gamma} \in \text{Emb}(S, M)$ is any parametrization of $\gamma \in \text{UEmb}(S, M)$ and the integrand denotes the determinant of the 3×3 matrix whose columns are the vectors $\partial_s \tilde{\gamma}(s)$, $\dot{\tilde{\gamma}}(s)$, and $\ddot{\tilde{\gamma}}(s)$ in \mathbb{R}^3 . Note that $\ker \tilde{\omega}|_{\tilde{\gamma}}$ consists of infinitesimal reparametrizations of the form $a \partial_s \tilde{\gamma}$ with $a \in C^\infty(\mathbb{S}^1)$, corresponding to tangent vector fields along the curve γ viewed as a submanifold of \mathbb{R}^3 .

2.3.1 Hamiltonian vector fields

For a function $\mathcal{H}: \mathcal{O} \rightarrow \mathbb{R}$ with $d\mathcal{H} \in \text{im } \flat\omega$, the Hamiltonian vector field $X_{\mathcal{H}} \in \Gamma(T\mathcal{O})$ with respect to the MW form ω is the unique vector field satisfying

$$d\mathcal{H} = \iota_{X_{\mathcal{H}}}\omega$$

and the Hamiltonian system is the flow along $X_{\mathcal{H}}$. That is,

$$\partial_t \gamma = X_{\mathcal{H}}(\gamma).$$

Example 2.3.2. On $\text{UEmb}(\mathbb{S}^1, \mathbb{R}^3)$ where \mathbb{R}^3 is equipped with the standard Euclidean metric, the length function as the Hamiltonian yields the binormal equation,

$$\partial_t \gamma = D_s \gamma \times D_s^2 \gamma. \tag{2.3.2}$$

The right-hand side is understood as $d\pi|_{\tilde{\gamma}}(D_s \tilde{\gamma} \times D_s^2 \tilde{\gamma})$ with fibration $\pi: \text{Emb}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \text{UEmb}(\mathbb{S}^1, \mathbb{R}^3)$ and any parametrization $\tilde{\gamma} \in \pi^{-1}(\gamma)$, where $D_s = \partial_s / |\partial_s \tilde{\gamma}|$ is the derivative with respect to the arc-length parameter.

More generally, on $\text{UEmb}(S, \mathbb{R}^m)$ with an oriented closed n -manifold S , choosing the n -dimensional volume of γ as the Hamiltonian \mathcal{H} yields the Hamiltonian vector field $X_{\mathcal{H}}$ as the mean curvature normal on the submanifold $\gamma(S)$, rotated by 90 degrees in the normal bundle $N\gamma(S)$ [HV03, Khe12]. While the binormal equation (6.4.2) is known to be an infinite-dimensional integrable system forming a KdV-type hierarchy [CKPP20], the integrability of its higher-dimensional analogues remains less understood.

2.3.2 Liouville form

On a symplectic manifold, a *Liouville form* is a 1-form whose exterior derivative equals to the symplectic form. Previous work [PCK⁺19, Tab17] showed that the MW structure ω on the space of space curves $\text{UEmb}(\mathbb{S}^1, \mathbb{R}^3)$ is exact by providing an explicit Liouville form. The 1-form η defined by

$$\eta_\gamma(\dot{\gamma}) = \frac{1}{3} \int_{\mathbb{S}^1} \det(\tilde{\gamma}, \dot{\tilde{\gamma}}, \partial_s \tilde{\gamma}) ds \quad (2.3.3)$$

where $\tilde{\gamma}$ is any parametrization of γ , satisfies $\omega = d\eta$.

A natural question is to consider its higher-dimensional generalization. However, the proof in [Tab17] relies on integration by parts using an explicit parametrization of \mathbb{S}^1 , which does not directly extend to higher dimensions.

Here, we present a new result for a generic closed and oriented n -dimensional manifold S and $M = \mathbb{R}^m$ equipped with the Euclidean volume form $\mu = dx_1 \wedge \dots \wedge dx_m$ (Theorem 2.3.3), and then extend it to a slightly more general setting (Theorem 2.3.6).

For $\dot{\delta}_\gamma = -\mathcal{L}_u \delta_\gamma \in T_{\delta_\gamma} \mathcal{D}_\mathcal{O}$, define

$$\eta_{\delta_\gamma}(\dot{\delta}_\gamma) = \frac{1}{m} \langle \dot{\delta}_\gamma, \iota_u \iota_x \mu \rangle = \frac{1}{m} \int_\gamma \iota_u \iota_x \mu. \quad (2.3.4)$$

This 1-form (2.3.4) admits an explicit representation analogous to (2.3.3):

$$\eta_\gamma(\dot{\gamma}) = \frac{1}{m} \int_S \det(\tilde{\gamma}, \dot{\tilde{\gamma}}, \partial_{s_1} \tilde{\gamma}, \dots, \partial_{s_{m-2}} \tilde{\gamma}) ds_1 \dots ds_{m-2}$$

where $\dot{\gamma} = u \circ \gamma$ and $s = (s_1, \dots, s_n)$ is local coordinates.

Theorem 2.3.3 (Liouville form on \mathbb{R}^m). *Suppose that \mathbb{R}^m is equipped with the standard volume form $\mu = dx_1 \wedge \dots \wedge dx_m$. Then we have $d\eta = \omega$.*

A special case of Theorem 2.3.3 for $\text{Emb}(\mathbb{S}^1, \mathbb{R}^3)$ is given in [Tab17, Proposition 2.1], where the proof is based on integration by parts using an explicit parametrization of a space curve.

However, this approach does not extend to the general setting $\text{Emb}(S, M)$ unless S admits an explicit global parametrization, such as $S = \mathbb{T}^n$. In contrast, we prove the general case of Theorem 2.3.3 in a parametrization-free fashion.

We formulate the proof using currents, as this computational routine will be reused throughout the present and the next chapters. To that end, we introduce an auxiliary result. First, note that each vector field $u \in \Gamma(TM)$ induces a vector field X^u on the $\text{Diff}_0(M)$ -orbit $\mathcal{D}_\mathcal{O}$ in $\mathcal{D}_{m-2}(M)$, defined by

$$X^u(\delta_\gamma) = -\mathcal{L}_u \delta_\gamma.$$

In other words, X^u is the fundamental vector field associated with the $\text{Diff}_0(M)$ -action on $\mathcal{D}_\mathcal{O}$, since $-\mathcal{L}_u \delta_\gamma = \frac{d}{dt} \Big|_{t=0} \text{Fl}_{u*}^t \delta_\gamma$, where $\text{Fl}_u^t \in \text{Diff}_0(M)$ is the time- t flow of u on M (see Section 2.1.2). This defines the following homomorphism;

Lemma 2.3.4. *Let $\hat{\mathcal{D}}$ be a $\text{Diff}_0(M)$ -orbit in the space of k -currents $\mathcal{D}_k(M)$. Then the map*

$$\begin{aligned} \Gamma(TM) &\rightarrow \Gamma(T\hat{\mathcal{D}}) \\ u &\mapsto X^u \end{aligned}$$

is a Lie algebra homomorphism. That is, for $u, v \in \Gamma(TM)$, we have

$$[X^u, X^v] = X^{[u,v]}.$$

Proof. To compute $[X^u, X^v] = \mathcal{L}_{X^u}X^v$, we use the formula for the Lie derivative:

$$\mathcal{L}_{X^u}X^v(h) = \left. \frac{d}{dt} \right|_{t=0} D\Phi_{X^u}^{-t} [X^v(\Phi_{X^u}^t(h))], \quad h \in \hat{\mathcal{D}}.$$

Here $\Phi_X^t: \hat{\mathcal{D}} \rightarrow \hat{\mathcal{D}}$ denotes the flow map on $\hat{\mathcal{D}}$ along a given vector field $X \in \Gamma(T\hat{\mathcal{D}})$, defined by the ODE

$$\begin{aligned} \frac{d}{dt}(\Phi_X^t(h)) &= X(\Phi_X^t(h)), \\ \Phi_X^0 &= \text{id}_{\hat{\mathcal{D}}} \end{aligned}$$

and $D\Phi_X^t|_h: T_h\hat{\mathcal{D}} \rightarrow T_{\Phi_X^t(h)}\hat{\mathcal{D}}$ is the differential of Φ_X^t at h .

For the fundamental vector field X^u , the flow is given by $\Phi_{X^u}^t(h) = \text{Fl}_{u*}^t h$ by definition. Using this, we have

$$X^v(\Phi_{X^u}^t(h)) = \left. \frac{d}{ds} \right|_{s=0} \text{Fl}_{v*}^s(\Phi_{X^u}^t(h)) = \left. \frac{d}{ds} \right|_{s=0} \text{Fl}_{v*}^s \text{Fl}_{u*}^t h,$$

and obtain

$$\begin{aligned} \left. \frac{d}{dt} \right|_{t=0} D\Phi_{X^u}^{-t} [X^v(\Phi_{X^u}^t(h))] &= \left. \frac{d}{dt} \right|_{t=0} \left. \frac{d}{ds} \right|_{s=0} \Phi_{X^u}^{-t} (\text{Fl}_{v*}^s \text{Fl}_{u*}^t h) \\ &= \left. \frac{d}{dt} \right|_{t=0} \left. \frac{d}{ds} \right|_{s=0} \text{Fl}_{u*}^{-t} \text{Fl}_{v*}^s \text{Fl}_{u*}^t h \\ &= -\mathcal{L}_u \mathcal{L}_v h + \mathcal{L}_v \mathcal{L}_u h \\ &= -\mathcal{L}_{[u,v]} h = X^{[u,v]}(h), \end{aligned}$$

which completes the proof. \square

Remark 2.3.5. At first glance, Lemma 2.3.4 may seem to contradict the standard result that the fundamental vector field mapping $\mathfrak{g} \rightarrow \Gamma(TN)$ on a manifold N is an *anti* Lie algebra homomorphism when the G -action on N is a left action [Lee12, Theorem 20.18]. The sign difference in our setting occurs because of our identification of vector fields $\text{diff}(M)$ with derivations $\text{Der}(C^\infty(M)) = \Gamma(TM)$, where the mapping $\text{der}: \text{diff}(M) \rightarrow \text{Der}(C^\infty(M))$ by $u \mapsto \mathcal{L}_u$ is an anti Lie algebra homomorphism.

Proof of Theorem 2.3.3. For $\dot{\delta}_\gamma = -\mathcal{L}_u \delta_\gamma$, $\dot{\delta}_\gamma = -\mathcal{L}_v \delta_\gamma$ with some $u, v \in \text{diff}(M)$, we compute $d\eta_{\delta_\gamma}(\dot{\delta}_\gamma, \dot{\delta}_\gamma) = d\eta_{\delta_\gamma}(X^u, X^v)$ explicitly, using the fundamental vector fields $X^u, X^v \in \Gamma(T\mathcal{D}_\mathcal{O})$ given by $X^u(\delta_\gamma) = -\mathcal{L}_u \delta_\gamma$, $X^v(\delta_\gamma) = -\mathcal{L}_v \delta_\gamma$. We use the formula

$$d\eta_{\delta_\gamma}(X^u, X^v) = \mathcal{L}_{X^u} \iota_{X^v} \eta_{\delta_\gamma} - \mathcal{L}_{X^v} \iota_{X^u} \eta_{\delta_\gamma} - \iota_{[X^u, X^v]} \eta_{\delta_\gamma}.$$

First, by Lemma 2.3.4, we have

$$\eta_{\delta_\gamma}([X^u, X^v]) = \eta_{\delta_\gamma}(X^{[u,v]}) = \eta_{\delta_\gamma}(-\mathcal{L}_{[u,v]}\delta_\gamma) = \frac{1}{m}\langle \delta_\gamma, \iota_{[u,v]}\iota_x\mu \rangle.$$

Next, we compute

$$\begin{aligned} \mathcal{L}_{X^u}\iota_{X^v}\eta_{\delta_\gamma} &= \left. \frac{d}{dt} \right|_{t=0} \eta(X^v(\Phi_{X^u}^t(\delta_\gamma))) = \left. \frac{1}{m} \frac{d}{dt} \right|_{t=0} \langle \text{Fl}_{u*}^t \delta_\gamma, \iota_v \iota_x \mu \rangle \\ &= \left. \frac{1}{m} \frac{d}{dt} \right|_{t=0} \langle \delta_\gamma, \text{Fl}_u^{t*} \iota_v \iota_x \mu \rangle = \frac{1}{m} \langle \delta_\gamma, \mathcal{L}_u \iota_v \iota_x \mu \rangle, \end{aligned}$$

where $\Phi_{X^u}^t$ is the time- t flow map of X^u , given by $\Phi_{X^u}^t(\delta_\gamma) = (\text{Fl}_u^t)_* \delta_\gamma$, as explained in the proof of Lemma 2.3.4. Similarly, we have $\mathcal{L}_{X^v}\iota_{X^u}\eta_{\delta_\gamma} = \frac{1}{m}\langle \delta_\gamma, \mathcal{L}_v \iota_u \iota_x \mu \rangle$.

Combining these results, we obtain

$$\begin{aligned} d\eta_{\delta_\gamma}(X^u, X^v) &= \frac{1}{m} (\langle \delta_\gamma, \mathcal{L}_u \iota_v \iota_x \mu \rangle - \langle \delta_\gamma, \mathcal{L}_v \iota_u \iota_x \mu \rangle - \langle \delta_\gamma, \iota_{[u,v]}\iota_x \mu \rangle) \\ &= \frac{1}{m} (\langle \delta_\gamma, \mathcal{L}_u \iota_v \iota_x \mu - \mathcal{L}_v \iota_u \iota_x \mu - \iota_{[u,v]}\iota_x \mu \rangle). \end{aligned}$$

By direct computation, we simplify the expression inside the dual pairing:

$$\begin{aligned} \mathcal{L}_u \iota_v \iota_x \mu - \mathcal{L}_v \iota_u \iota_x \mu - \iota_{[u,v]}\iota_x \mu &= \mathcal{L}_u \iota_v \iota_x \mu - \mathcal{L}_v \iota_u \iota_x \mu - \mathcal{L}_u \iota_v \iota_x \mu + \iota_v \mathcal{L}_u \iota_x \mu \\ &= \iota_v \mathcal{L}_u \iota_x \mu - \mathcal{L}_v \iota_u \iota_x \mu \\ &= \iota_v d\iota_u \iota_x \mu + \iota_v \iota_u d\iota_x \mu - d\iota_v \iota_u \iota_x \mu - \iota_v d\iota_u \iota_x \mu \\ &= \iota_v \iota_u \mathcal{L}_x \mu - d\iota_v \iota_u \iota_x \mu. \end{aligned}$$

Note that $\mathcal{L}_x \mu = m\mu$ for the volume form $\mu = dx_1 \wedge \cdots \wedge dx_m$, and that $\partial\delta_\gamma = \delta_{\partial\gamma} = 0$. Therefore, we obtain

$$\begin{aligned} d\eta_{\delta_\gamma}(X^u, X^v) &= \frac{1}{m} \langle \delta_\gamma, \iota_v \iota_u \mathcal{L}_x \mu - d\iota_v \iota_u \iota_x \mu \rangle \\ &= \langle \delta_\gamma, \iota_v \iota_u \mu \rangle - \frac{1}{m} \langle \partial\delta_\gamma, \iota_v \iota_u \iota_x \mu \rangle \\ &= \omega_{\delta_\gamma}(X^u, X^v) - 0. \end{aligned}$$

Thus we have shown $d\eta = \omega$. □

Theorem 2.3.3 can be extended from \mathbb{R}^m to a general manifold M equipped with an exact volume form;

Theorem 2.3.6 (Liouville form on M with an exact volume form). *Let μ be a volume form on M such that $\mu = d\nu$ for some $\nu \in \Omega^{m-1}(M)$. Define a 1-form η on \mathcal{D}_O by*

$$\eta_{\delta_\gamma}(-\mathcal{L}_u \delta_\gamma) = \langle \delta_\gamma, \iota_u \nu \rangle.$$

Then $d\eta = \omega$.

Example 2.3.7. The Liouville form η in Theorem 2.3.3 is a special case of Theorem 2.3.6 with $(M, \mu) = (\mathbb{R}^m, dx_1 \wedge \cdots \wedge dx_m)$ and

$$\nu = \frac{1}{m} \sum_i x_i dx_{i+1} \wedge \cdots \wedge dx_{i+m-1},$$

where the indices are taken modulo m .

Proof of Theorem 2.3.6. Let $\dot{\delta}_\gamma = -\mathcal{L}_u \delta_\gamma$, $\mathring{\delta}_\gamma = -\mathcal{L}_v \delta_\gamma$ with some $u, v \in \text{diff}(M)$. Using the same computational routine as in the proof of Theorem 2.3.3, with the fundamental vector fields $X^u, X^v \in \Gamma(T\mathcal{D}_\mathcal{O})$, we compute:

$$\begin{aligned} d\eta_{\delta_\gamma}(\dot{\delta}_\gamma, \mathring{\delta}_\gamma) &= \mathcal{L}_{X^u} \iota_{X^v} \eta - \mathcal{L}_{X^v} \iota_{X^u} \eta - \iota_{[X^u, X^v]} \eta \\ &= \langle \dot{\delta}_\gamma, \mathcal{L}_u \iota_v \nu \rangle - \langle \delta_\gamma, \mathcal{L}_v \iota_u \nu \rangle - \langle \delta_\gamma, \mathcal{L}_{[u, v]} \nu \rangle \\ &= \langle \dot{\delta}_\gamma, \iota_v \iota_u d\nu - d\iota_v \iota_u \nu \rangle \\ &= \langle \dot{\delta}_\gamma, \iota_v \iota_u \mu \rangle - \langle \partial \delta_\gamma, \iota_v \iota_u \nu \rangle \\ &= \omega_{\delta_\gamma}(\dot{\delta}_\gamma, \mathring{\delta}_\gamma) - 0. \end{aligned}$$

□

Remark 2.3.8 (Proof via tilde calculus). Theorem 2.3.6 can alternatively be proved using the tilde calculus introduced by Haller and Vizman [HV03, Viz11]. We explain this approach in the appendix of this chapter (Section 2.4).

Remark 2.3.9 (Liouville form on a closed manifold). We are not aware of the existence or non-existence of a Liouville form on a general closed manifold (M, μ) . There are a few special cases where exactness has been proved. For example, the space of two distinct points on the sphere, equipped with the MW form (see Section 2.3.5), is symplectomorphic to the cotangent bundle of a certain space and therefore admits a Liouville form [OU13].

However, we speculate that exactness of the MW form is unlikely in the general case when the volume form μ is not exact, or at least that a Liouville form may not admit an explicit expression.

In contrast, we will explicitly define a Liouville form in the prequantum sense (Definition 4.2.2) on the space of implicit representations of submanifolds in a generic closed ambient manifold. Our result does not, however, directly extend to unbounded manifolds, including simple cases such as \mathbb{R}^m .

Remark 2.3.10 (The Liouville form η is not a connection form). The fibration $\pi: \text{Emb}(S, M) \rightarrow \text{UEmb}(S, M)$ is a principal G -bundle, where the structure group G is the group of reparametrizations $\text{Diff}^+(S)$. Using the Liouville form η defined in Theorem 2.3.6, we define a 1-form $\tilde{\eta} = \pi^* \eta$ on $\text{Emb}(S, M)$. By construction, $d\tilde{\eta} = \pi^* \omega$, and $\tilde{\eta}$ is equivariant under the $\text{Diff}^+(S)$ -action. However, it is not a connection form, as it lacks vertical reproducibility: since $\ker d\pi \subset \ker \tilde{\eta}$, we have $\tilde{\eta}_{\tilde{\gamma}}(\hat{\xi}) = 0$ for any $\xi \in \text{diff}(S)$, where $\hat{\xi}(\tilde{\gamma}) = d\tilde{\gamma}(\xi)$ is the fundamental vector field corresponding to ξ at $\tilde{\gamma}$.

On the other hand, the Liouville form on the space of implicit representations that we will define does yield a connection form (Proposition 4.3.6).

2.3.3 Momentum maps on $\text{UEmb}(S, \mathbb{R}^m)$

The Liouville form η of the MW form ω can be used to describe conserved quantities of Hamiltonian systems on \mathcal{O} . As in finite-dimensional symplectic geometry, when a Lie group G acts on \mathcal{O} , a map $J: \mathcal{O} \rightarrow \mathfrak{g}^*$ is called a momentum map if

$$d\langle J(\cdot), \xi \rangle|_\gamma = \iota_\xi \omega|_\gamma$$

for any $\gamma \in \mathcal{O}$, $\xi \in \mathfrak{g}$, and its fundamental vector field $\hat{\xi} \in \Gamma(T\mathcal{O})$. If a given Hamiltonian function $\mathcal{H}: \mathcal{O} \rightarrow \mathbb{R}$ is invariant under the G -action, then $J(\gamma_t) \in \mathfrak{g}^*$ is conserved along the Hamiltonian flow γ_t .

When the Liouville form η is also invariant under the G -action, the momentum map can be expressed as

$$d\langle J(\gamma), \xi \rangle = \iota_{\hat{\xi}}\omega = \mathcal{L}_{\hat{\xi}}\eta - d\iota_{\hat{\xi}}\eta = -d(\eta(\hat{\xi})),$$

so that $\langle J(\gamma), \xi \rangle = -\eta(\hat{\xi})$, up to addition of a constant.

Example 2.3.11. Consider the rotation action of $\mathrm{SO}(3)$ on $\mathrm{UEmb}(\mathbb{S}^1, \mathbb{R}^3) \cong \mathcal{D}_{\mathrm{UEmb}(\mathbb{S}^1, \mathbb{R}^3)}$. The fundamental vector field $\hat{\xi} \in \Gamma(T\mathcal{D}_{\mathrm{UEmb}(\mathbb{S}^1, \mathbb{R}^3)})$ associated with $\xi \in \mathfrak{so}(3)$ is given at each δ_γ by $\hat{\xi}(\delta_\gamma) = -\mathcal{L}_{R_\xi}\delta_\gamma$ where $R_\xi \in \mathrm{sdiff}(M)$ is the rotational vector field defined by $R_\xi(x) = \xi \times x$ on \mathbb{R}^3 .

Since the Liouville form η (2.3.4) is invariant under rotation, we have $\langle J(\cdot), \xi \rangle = -\eta(\hat{\xi})$, up to an additive constant. Therefore, the *angular momentum* $J(\delta_\gamma) \in \mathfrak{so}^*(3)$ at δ_γ can be computed from R_ξ and the expression (2.3.3). Since $R_\xi(x) = \xi \times x = \mathrm{curl}\left(-\frac{1}{3}x \times (\xi \times x)\right)$, we have

$$\begin{aligned} \langle J(\gamma), \xi \rangle &= -\frac{1}{3} \int_{\mathbb{S}^1} \det(\partial_s \tilde{\gamma}, \tilde{\gamma}, \hat{\xi} \circ \tilde{\gamma}) ds \\ &= \int_\gamma \left(-\frac{1}{3} x \times (\xi \times x) \right)^{\flat} \\ &= \mathrm{Flux}_{m-2}^{\mathrm{ex}}(\delta_\gamma, R_\xi), \quad \xi \in \mathfrak{so}(3), \end{aligned}$$

where $\tilde{\gamma}$ is any parametrization of γ and the flat operator $\flat: \mathrm{diff}(\mathbb{R}^3) \rightarrow \Omega^1(\mathbb{R}^3)$ is defined with respect to the Euclidean metric.

The quantity $\langle J(\gamma), \xi \rangle$ can be interpreted as the signed volume of the surface of revolution obtained by rotating γ about the axis $\xi/|\xi| \in \mathfrak{so}(3) \cong \mathbb{R}^3$, multiplied by $|\xi|$.

On \mathbb{R}^m , the Liouville form η is not invariant under $\mathrm{SDiff}_0(\mathbb{R}^m)$; for instance, its value depends on the choice of origin and is therefore not translation-invariant. It is preserved only by volume-preserving linear transforms $x \mapsto fx$ with $f \in \mathrm{SL}(m)$, which form a finite-dimensional subgroup of $\mathrm{SDiff}^+(\mathbb{R}^m)$.

In contrast, the Liouville form on the space of implicit representations we will introduce in Section 4.2 is invariant under $\mathrm{SDiff}^+(M)$ for any closed manifold M (see Section 4.4.1).

2.3.4 Riemannian and formal Kähler structure

So far, we introduced the MW structure using only a volume form on the ambient manifold M . When a Riemannian metric is given on M , we can additionally equip \mathcal{O} with a Riemannian structure.

Assume now that M is equipped with a Riemannian metric g inducing the volume form μ . We first define an almost complex structure, that is, a vector bundle homomorphism $\mathcal{J}: T\mathcal{O} \rightarrow T\mathcal{O}$ satisfying $\mathcal{J}^2 = -\mathrm{id}$.

Let $\tilde{\mathcal{O}} := \pi^{-1}\mathcal{O}$ be the parametrization space of \mathcal{O} via embeddings. We define an operator $\tilde{\mathcal{J}}: T\tilde{\mathcal{O}} \rightarrow T\tilde{\mathcal{O}}$ by projecting each $\dot{\tilde{\gamma}} \in T_{\tilde{\gamma}}\tilde{\mathcal{O}}$ onto the normal bundle $N_{\tilde{\gamma}}\tilde{\mathcal{O}}$, and then applying a 90-degree rotation in $N_{\tilde{\gamma}}\tilde{\mathcal{O}}$ with respect to the metric g , oriented so that

$$\mu(\partial_{s_1}\tilde{\gamma}(s), \dots, \partial_{s_n}\tilde{\gamma}(s), \tilde{\mathcal{J}}\dot{\tilde{\gamma}}(s), \dot{\tilde{\gamma}}(s)) \geq 0,$$

where $s = (s_1, \dots, s_n)$ is any local chart consistent with the orientation of $\tilde{\gamma}$.

Since $\ker \tilde{\mathcal{J}} = \ker d\pi$ and $\tilde{\mathcal{J}}$ is invariant under reparametrization, it descends to an almost complex structure \mathcal{J} on \mathcal{O} , defined by

$$\mathcal{J}\dot{\gamma} = d\pi|_{\tilde{\gamma}}(\tilde{\mathcal{J}}\dot{\tilde{\gamma}})$$

with any representatives $\tilde{\gamma}$ of γ and $\dot{\tilde{\gamma}}$ of $\dot{\gamma}$, respectively.

With \mathcal{J} and the MW form ω , we define a Riemannian metric on \mathcal{O} by

$$\mathcal{G}_\gamma(\dot{\gamma}, \dot{\gamma}) = \omega_\gamma(\dot{\gamma}, \mathcal{J}\dot{\gamma}).$$

In terms of de Rham currents, this becomes

$$\mathcal{G}_{\delta_\gamma}(-\mathcal{L}_u\delta_\gamma, -\mathcal{L}_v\delta_\gamma) = \langle \delta_\gamma, \iota_{\hat{v}}\iota_u\mu \rangle,$$

where \hat{v} is any vector field on M satisfying $\hat{v} \circ \gamma = \mathcal{J}(v \circ \gamma)$.

We remark that the triple $(\mathcal{G}, \omega, \mathcal{J})$ defines a formal Kähler structure in the sense that

$$\mathcal{G}(\mathcal{J}\cdot, \cdot) = \omega(\cdot, \cdot).$$

Strictly speaking, this is not a Kähler structure in the classical sense, which additionally requires a complex structure *i.e.*, the existence of holomorphic coordinates [Lem93, MZ96].

In some infinite-dimensional settings, including our case, the Nijenhuis tensor

$$\mathcal{N}(\dot{\gamma}, \dot{\gamma}) := [\mathcal{J}\dot{\gamma}, \mathcal{J}\dot{\gamma}] - \mathcal{J}[\mathcal{J}\dot{\gamma}, \dot{\gamma}] - \mathcal{J}[\dot{\gamma}, \mathcal{J}\dot{\gamma}] - [\dot{\gamma}, \dot{\gamma}], \quad \dot{\gamma}, \dot{\gamma} \in T_\gamma\mathcal{O}$$

vanishes [Bry09, Theorem 3.4.3][Hen09, Theorem 2.5], but this does not imply the existence of a complex structure [Lem93, Theorem 10.5]. This contrasts with the finite-dimensional case, where vanishing of the Nijenhuis tensor is equivalent to integrability. For details, see [Hen09, Chapter 2] for example.

2.3.5 Base Dimension 2

We briefly discuss the case where the dimension m of the ambient manifold M is 2, which is a finite-dimensional setting in contrast to $m > 2$. Although many aspects are simplified in this case, the symplectic structure requires additional care: orientations of points must be supplied explicitly, as they are not encoded in the embeddings themselves.

We regard the domain S as a collection of N distinct points. The space of embeddings is

$$\text{Emb}(S, M) := M^N \setminus \Delta = \{(x_1, \dots, x_N) \in M^N \mid x_j \neq x_k \text{ for } j \neq k\},$$

which is a smooth $2N$ -dimensional manifold.

In contrast to higher-dimensional settings, this space does not carry orientation data, that is, whether each point is embedded positively or negatively. Hence we need to manually attach the orientations $I = (I_1, \dots, I_N) \in \{\pm 1\}^N$ to the symplectic structure. Namely we define

$$\tilde{\omega}_\gamma(\dot{\gamma}, \mathring{\gamma}) := \sum_{j=1}^N I_j \pi_j^* \mu(\dot{\gamma}, \mathring{\gamma}) = \sum_{j=1}^N I_j \mu(\dot{\gamma}_j, \mathring{\gamma}_j), \quad (2.3.6)$$

where $\dot{\gamma}, \mathring{\gamma} \in T_\gamma \text{Emb}(S, M) \cong T_{x_1} M \times \dots \times T_{x_N} M$, and π_j denotes projection onto the j -th factor. Since the volume form μ is symplectic on a 2-dimensional manifold M , $\tilde{\omega}_\gamma$ is non-degenerate and thus already symplectic.

We may still consider the unparametrized shape space

$$\text{UEmb}(S, M) := \text{Emb}(S, M) / \text{Diff}^+(S),$$

by quotienting out the reparametrization group $\text{Diff}^+(S) := \text{Sym}(k_1) \times \text{Sym}(k_2)$, which permutes the k_1 positively oriented and k_2 negatively oriented points, with $k_1 + k_2 = N$.

Using currents (which are simply distributions in base dimension 2), we can encode orientations directly to the shape $\delta_\gamma \in \mathcal{D}_{\text{UEmb}(S, M)} \subset \mathcal{D}_0(M)$ by

$$\delta_\gamma := \sum_{j=1}^N I_j \delta_{x_j}$$

with the Dirac delta distribution δ_{x_j} at each x_j .

Then, for $\dot{\delta}_\gamma = -\mathcal{L}_u \delta_\gamma$ and $\mathring{\delta}_\gamma = -\mathcal{L}_v \delta_\gamma$ with vector fields $u, v \in \Gamma(TM)$, the MW structure is defined

$$\omega_{\delta_\gamma}(-\mathcal{L}_u \delta_\gamma, -\mathcal{L}_v \delta_\gamma) = \langle \delta_\gamma, \iota_v \iota_u \mu \rangle$$

in the same way as for higher dimensional cases. This evaluates to

$$\langle \delta_\gamma, \iota_v \iota_u \mu \rangle = \sum_{j=1}^N I_j \langle \delta_{x_j}, \mu(u, v) \rangle = \sum_{j=1}^N I_j \mu(\dot{\gamma}_j, \mathring{\gamma}_j),$$

which agrees with the pointwise expression of the MW form (2.3.6).

2.3.6 Toward implicit representations: the space of all the codimension-2 shapes

We considered the explicit shape space consisting of embeddings modulo reparametrization for each fixed manifold S . We conclude this section by defining their disjoint union:

$$\mathcal{X} := \bigsqcup_S \text{UEmb}(S, M)$$

where S runs through all oriented, closed n -dimensional manifolds. Thus, the explicit shape space \mathcal{X} is the set of all closed n -dimensional submanifolds of M , also known as the space of n -dimensional nonlinear Grassmannians [HV03]. Each $\gamma \in \mathcal{X}$ lies in a unique $\text{UEmb}(S, M)$, and we have $T_\gamma \mathcal{X} = T_\gamma \text{UEmb}(S, M)$. Hence, the symplectic structure ω defined on each $\text{UEmb}(S, M)$ naturally extends to \mathcal{X} . For simplicity, we will identify \mathcal{X} with its current counterpart $\mathcal{D}_\mathcal{X}$ via the injection $\mathcal{I}: \mathcal{X} \rightarrow \mathcal{D}_n(M)$ without further mention.

Example 2.3.12. If $m = 2$, the space $\mathcal{X} = \bigsqcup_{k \in \mathbb{N}} \left\{ \sum_{j=1}^k I_j \delta_{x_j} \mid I_j \in \{-1, 1\}, x \in M^k \setminus \Delta \right\}$ is all the possible configurations of finitely many unordered oriented points. If $m = 3$, $\mathcal{X} = \bigsqcup_{k \in \mathbb{N}} \text{UEmb} \left(\bigsqcup^k \mathbb{S}^1, M \right)$ consists of all smooth links with finitely many knot components. If $m > 3$, then S may be a disjoint union of different types of manifolds.

2.4 Appendix: Liouville form via the tilde calculus

An alternative approach to prove Theorem 2.3.6 is using the so-called tilde calculus, also known as the transgression of differential forms, presented in [HV03, Viz11]. The tilde calculus is driven by the tilde operator which takes as input a differential form in a finite-dimensional manifold and outputs another differential form in an infinite-dimensional manifold.

Let M be an m -dimensional oriented manifold and S be $n (< m)$ dimensional closed (i.e., compact & boundary-less) oriented manifold. Let us define the tilde operator

$$\begin{aligned} \sim : \Omega^k(M) &\rightarrow \Omega^{k-n}(\text{UEmb}(S, M)) \\ \alpha &\mapsto \tilde{\alpha} \end{aligned}$$

for $k > n$ by

$$\tilde{\alpha}_\gamma(u_1 \circ \gamma, \dots, u_{k-n} \circ \gamma) = \int_S \gamma^*(\iota_{u_{k-n}} \dots \iota_{u_1} \alpha)$$

where $\gamma \in \text{UEmb}(S, M)$ and $u_1 \circ \gamma, \dots, u_{k-n} \circ \gamma \in T_\gamma \text{UEmb}(S, M)$ with some $u_1, \dots, u_{k-n} \in \Gamma(TM)$.

The framework of the tilde calculus allows a simple expression of the Marsden–Weinstein structure:

Example 2.4.1 (The Marsden–Weinstein form). Let μ be a volume form on the ambient manifold M and let the dimension of S be $n = m - 2$. Then the Marsden–Weinstein form (2.3.1) is written as $\tilde{\mu} \in \Omega^2(\text{UEmb}(S, M))$. Explicitly, it is

$$\tilde{\mu}_\gamma(u \circ \gamma, v \circ \gamma) = \int_M \gamma^* \iota_v \iota_u \mu.$$

The tilde calculus has useful properties. In particular, we can use the commutativity between d and \sim [HV03, Lemma 1] to show the exactness of the Marsden–Weinstein form.

Proposition 2.4.2. *We have $d\tilde{\alpha} = \tilde{d\alpha}$.*

With this, the exactness of the Marsden–Weinstein form $\tilde{\mu}$ is a direct consequence of the exactness of μ .

An alternative proof of Theorem 2.3.6. The 1-form η given by $\eta_{\delta_\gamma}(-\mathcal{L}_u \delta_\gamma) = \langle \delta_\gamma, \iota_u \mu \rangle$ in Theorem 2.3.6 agrees with $\tilde{\mu}$. Applying to Proposition 2.4.2 to $\mu = d\nu$, we get

$$\tilde{\mu} = \tilde{d\nu} = d\tilde{\nu}.$$

□

Implicit representations of codimension-2 submanifolds

In the previous chapter, we reviewed the *explicit* representations of codimension-2 submanifolds as embeddings of a manifold into an ambient manifold. In this chapter, we introduce *implicit* representations via a class of complex-valued functions and study the geometric structures of the space of these implicit representations.

This leads to further study of its relation with the Marsden–Weinstein symplectic structure of the codimension-2 shape space (Chapter 4), and its exploitation for numerical simulation of curve dynamics (Chapter 5).

3.1 Introduction

Codimension-1 submanifolds—such as curves on a surface and surfaces in a three dimensional space—can be implicitly represented using a level set function. This representation naturally extends to multi-codimensional submanifolds using multiple level set functions. Specifically, when the codimension is 2, submanifolds—such as curves in a three dimensional space and surfaces in a four dimensional space—can be expressed as the zero sets of complex-valued functions.

Such implicit representations via complex functions are useful for studying the geometry and dynamics of submanifolds. For instance, prior work has exploited these implicit representations to study the dynamics of space curves, such as curve-shortening flow [RMXO01] and quantum vortex filaments governed by the Gross-Pitaevskii equation [OTH02, VKPS16, JS18].

In the present and the next chapter, we investigate the space of these implicit representations for codimension-2 submanifolds. In particular, we study its geometric structures related to the Marsden–Weinstein (MW) structure.

As in the previous chapter, let us denote \mathcal{O} a $\text{Diff}_0(M)$ -orbit in the space of codimension-2 submanifolds, represented by embeddings modulo reparametrization. We call \mathcal{O} the *explicit shape space*. The implicit representation we define for each element in \mathcal{O} is not unique, as multiple complex-valued functions can share the same zero set. This redundancy makes the

implicit shape space $\mathcal{F}_{\mathcal{O}}$ a fiber bundle over the explicit shape space \mathcal{O} .¹ Each element $\psi \in \mathcal{F}_{\mathcal{O}}$ is a complex-valued function over M , while each point $\gamma \in \mathcal{O}$ is a codimension-2 submanifold in M . The projection from $\mathcal{F}_{\mathcal{O}}$ to \mathcal{O} is to extract the zero set and the orientation.

This chapter builds a foundation of implicit representations for our further investigation in the subsequent chapters, which leverages their non-uniqueness. In Chapter 4, we study the geometry of the implicit shape space related to the Marsden-Weinstein structure. Specifically, we reveal a prequantum structure on a certain quotient space of the implicit shape space $\mathcal{F}_{\mathcal{O}}$ over the explicit shape space \mathcal{O} . In Chapter 5, we exploit this redundancy to develop a numerical method for space curve dynamics.

Conventions In the last chapter, we introduced the explicit shape space in terms of embeddings modulo reparametrization, and in terms of de Rham currents. Throughout this and the following chapter, we identify the explicit shape space \mathcal{X} of all the codimension-2 submanifolds with its counterpart $\mathcal{D}_{\mathcal{X}}$ in terms of de Rham currents, as well as each $\text{Diff}_0(M)$ -orbit \mathcal{O} with the corresponding currents $\mathcal{D}_{\mathcal{O}}$, without explicitly restating this identification.

3.2 Implicit representations of codimension-2 shapes

We define the space \mathcal{F} of implicit representations of submanifolds as a subset of $C^\infty(M, \mathbb{C})$ as follows. The space \mathcal{F} consists of functions ψ such that the zero set $\psi^{-1}(0)$ is nonempty, and the differential $d\psi|_x : T_x M \rightarrow \mathbb{C}$ is surjective at each $x \in \psi^{-1}(0)$.

Define a fibration $\Pi: \mathcal{F} \rightarrow \mathcal{X}$ by assigning to each $\psi \in \mathcal{F}$ a unique $\gamma \in \mathcal{X}$ such that $\text{im } \gamma = \psi^{-1}(0)$, with the orientation of ψ determined as follows. For $\gamma \in \text{UEmb}(S, M)$ with some S having k connected components, there are 2^k possible orientations. Among these, we select the unique one satisfying

$$\int_{\partial \Sigma} \text{Im} \frac{d\psi}{\psi} = 2\pi \Sigma \cap \gamma(S)$$

for any oriented $(m-1)$ -dimensional topological disk Σ intersecting transversely with $\gamma(S)$, where $\Sigma \cap \gamma(S)$ denotes the signed intersection number between Σ and $\gamma(S)$.

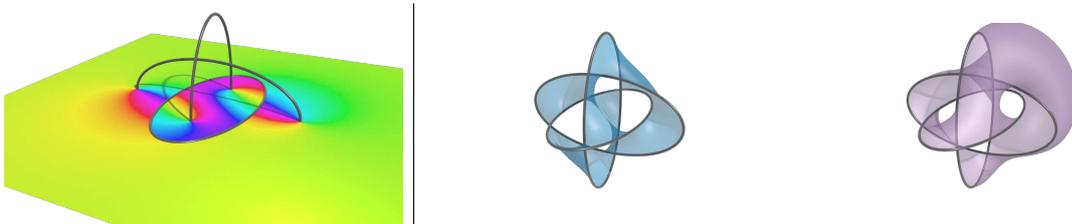


Figure 3.1: Left: Visualization of an implicit representation $\psi \in \mathcal{F}$ on the plane $z = 0$ in \mathbb{R}^3 for a Borromean ring $\gamma \in \mathcal{X}$. Each color represents a phase *i.e.*, the preimage of an element $s \in \mathbb{S}^1$ under the phase map $\phi = \psi/|\psi|$. Middle and right: Visualizations of the preimages of two distinct elements in \mathbb{S}^1 as surfaces bordered by $\text{im } \gamma = \psi^{-1}(0)$. These preimages have codimension 1, and their intersection has codimension 2, representing $\text{im } \gamma$.

¹Similar to \mathcal{O} , the space $\mathcal{F}_{\mathcal{O}}$ is more precisely described as a subset of the space of complex-valued functions, characterized by the orbits under a certain group action.

Existence of implicit representations The fibration $\Pi: \mathcal{F} \rightarrow \mathcal{X}$ is not surjective. More precisely, our implicit representation of shapes is limited to codimension-2 submanifolds that are exact in homology; that is, each γ must be the boundary of an oriented codimension-1 submanifold Σ .

In this case, we can construct a map $\theta: M \setminus \text{im } \gamma \rightarrow \mathbb{T}^1$ as the solution to the Dirichlet boundary problem:

$$\begin{aligned} \Delta\theta(x) &= 0 & \text{for } x \notin \Sigma, \\ \theta(x) &\rightarrow 2\pi & \text{as } x \rightarrow \Sigma_f, \\ \theta(x) &\rightarrow 0 & \text{as } x \rightarrow \Sigma_b, \end{aligned}$$

where Σ_f and Σ_b denote the front and back sides of Σ , respectively. Given any smooth non-negative function ρ such that $\rho^{-1}(0) = \text{im } \gamma$, we can define an implicit representation ψ by $\psi(x) = \rho(x)e^{i\theta(x)}$. For $M = \mathbb{S}^m$, viewed as the one-point compactification of \mathbb{R}^m , the so-called solid angle field (visualized in Figure 3.1) is a special case of such a phase field $e^{i\theta}$ [BDR20, CI24].

Remark 3.2.1. In this article, we consider *unlabeled shapes* by taking the quotient of $\text{Emb}(S, M)$ by $\text{Diff}^+(S)$, rather than by $\text{Diff}_0(S)$, as noted in Remark 2.2.2. Our implicit representations are naturally suited to this setting, since the level set of each ψ is unlabeled and unparametrized, yet having the orientation as explained above.

Remark 3.2.2 (Relation with open book decomposition). Each $\psi \in \mathcal{F}$ can be interpreted as a (relaxed) instance of an *open book decomposition* in contact geometry: a decomposition of the ambient manifold M into an $m - 2$ dimensional submanifold B called the *binder*, and a family of $m - 1$ dimensional submanifolds $\{\Sigma_s\}_{s \in \mathbb{S}^1}$, called *pages*, such that $\partial\Sigma_s = B$ for each $s \in \mathbb{S}^1$ [Gei08, Etn05].

For a given $\psi \in \mathcal{F}$, if the phase map $\phi: M \setminus \psi^{-1}(0) \rightarrow \mathbb{S}^1$, defined by $\phi(x) = \frac{\psi(x)}{|\psi(x)|}$ is a submersion, then the zero set $\{\psi = 0\}$ serves as the binder, and the level sets $\{\phi^{-1}(s)\}_{s \in \mathbb{S}^1}$ form the pages. Even if ϕ fails to be a submersion, its level sets still form a foliation of M , and almost every $s \in \mathbb{S}^1$ is a binder by Sard's theorem.

Orbit space The full space \mathcal{F} is large and, in particular, consists of multiple connected components. Just as in the case of the explicit representation space \mathcal{X} , where each connected component is a $\text{Diff}_0(M)$ -orbit \mathcal{O} , we focus here on a single orbit $\mathcal{U} \subset \mathcal{F}$ under the action of the following group.

We define the semi-direct product group

$$\text{DC} := \text{Diff}_0(M) \ltimes \text{Exp}(C^\infty(M, \mathbb{C})),$$

where $\text{Diff}_0(M)$ acts on $\text{Exp}(C^\infty(M, \mathbb{C})) := \{e^\varphi \mid \varphi \in C^\infty(M, \mathbb{C})\}$ via

$$f \triangleright e^\varphi = e^\varphi \circ f^{-1}, \quad f \in \text{Diff}_0(M), e^\varphi \in \text{Exp}(C^\infty(M, \mathbb{C})).$$

The Lie group DC has a Lie algebra $\mathfrak{DC} = \text{diff}(M) \ltimes C^\infty(M, \mathbb{C})$ with Lie bracket

$$[(u, a), (v, b)] = ([u, v], -\mathcal{L}_u b + \mathcal{L}_v a), \quad u, v \in \text{diff}(M), a, b \in C^\infty(M, \mathbb{C}), \quad (3.2.1)$$

where $[u, v]$ denotes the Lie bracket of vector fields on M .

Proposition 3.2.3. *The operation $\mathfrak{DC} \times \mathfrak{DC} \rightarrow \mathfrak{DC}$ given in (3.2.1) defines a Lie bracket i.e., it is a bilinear skew-symmetric form satisfying the Jacobi identity.*

Proof. These properties can be verified by a straightforward computation. \square

We now define an DC-action on \mathcal{F} by

$$(f, e^\varphi) \triangleright \psi = (\psi \circ f^{-1}) \cdot e^\varphi, \quad (f, e^\varphi) \in \text{DC}, \psi \in \mathcal{F}.$$

Let $\mathcal{U}_\psi \subset \mathcal{F}$ denote the DC-orbit containing $\psi \in \mathcal{F}$. The tangent space at ψ is

$$T_\psi \mathcal{U}_\psi = \{-\mathcal{L}_u \psi + \varphi \psi \mid (u, \varphi) \in \mathfrak{DC}\}.$$

Given $\gamma = \Pi\psi \in \mathcal{X}$ and $\dot{\psi} = -\mathcal{L}_u \psi + \varphi \psi \in T_\psi \mathcal{U}_\psi$, we have $d\Pi(\dot{\psi}) = u \circ \gamma \in T_\gamma \mathcal{X}$. Using the identification of \mathcal{X} with de Rham currents $\mathcal{D}_\mathcal{X}$, this becomes $d\Pi(\dot{\psi}) = -\mathcal{L}_u \delta_\gamma \in T_{\delta_\gamma} \mathcal{D}_\mathcal{X}$. Hence, we obtain

$$\begin{aligned} \Pi(\mathcal{U}_\psi) &= \mathcal{O}_{\Pi\psi}, \\ d\Pi(T_\psi \mathcal{U}_\psi) &= T_{\Pi\psi} \mathcal{O}_{\Pi\psi}, \end{aligned}$$

where $\mathcal{O}_{\Pi\psi}$ is the $\text{Diff}_0(M)$ -orbit of $\Pi\psi \in \mathcal{X}$.

For the remainder of the present and the next chapter, we denote these orbits simply by $\mathcal{U} \subset \mathcal{F}$ and $\mathcal{O} \subset \mathcal{X}$ (or $\mathcal{D}_\mathcal{O} \subset \mathcal{D}_\mathcal{X}$ under the identification), without specifying a subscript ψ or γ for indicating a particular orbit.

Remark 3.2.4. The semi-direct product group $\text{DC} = \text{Diff}_0(M) \ltimes \text{Exp}(C^\infty(M, \mathbb{C}))$ is the simplest extension of $\text{Exp}(C^\infty(M, \mathbb{C}))$ by $\text{Diff}_0(M)$. That is, there is a split short exact sequence of groups

$$1 \rightarrow \text{Exp}(C^\infty(M, \mathbb{C})) \rightarrow \text{DC} \rightarrow \text{Diff}_0(M) \rightarrow \text{id}_M.$$

On \mathcal{F} , the group DC acts analogously to Euclidean transformations, combining rotation and translation. The composition rule is

$$(f_2, e^{\varphi_2}) \circ (f_1, e^{\varphi_1}) = (f_2 \circ f_1, (e^{\varphi_1} \circ f_2^{-1}) \cdot e^{\varphi_2}),$$

and $\text{id}_M \ltimes \text{Exp}(C^\infty(M, \mathbb{C}))$ is a normal subgroup in DC.

3.3 Geometry of fibers

In this subsection, we study the geometry of the fiber bundle $\Pi^{-1}\mathcal{O}$ consisting of implicit representations over a $\text{Diff}_0(M)$ -orbit \mathcal{O} of explicit representations, as well as the geometry of a DC-orbit \mathcal{U} within $\Pi^{-1}\mathcal{O}$.² In particular, we show that while \mathcal{U} is path-connected, each fiber may have multiple connected components, and the number of these components equals the rank of the first integral de Rham cohomology group of M .

²Clearly, $\Pi^{-1}\mathcal{O}$ contains a DC-orbit \mathcal{U} , but at this point we do not know whether $\mathcal{U} = \Pi^{-1}\mathcal{O}$. For further discussion, see Conjecture 3.3.5.

We define a fiber-preserving group action on $\Pi^{-1}\mathcal{O}$. For each $\gamma \in \mathcal{O}$, let $\text{Diff}_\gamma(M)$ be the connected component of the stabilizer of γ containing id_M . This is the subgroup of $\text{Diff}_0(M)$ consisting of diffeomorphisms f for which there exists a path $\{f_t\}_{t \in [0,1]} \subset \text{Diff}_0(M)$ such that $f_t \circ \gamma = \gamma$ for all $t \in [0,1]$, with $f_0 = \text{id}_M$ and $f_1 = f$. Then $\text{DC}_\gamma := \text{Diff}_\gamma(M) \times \text{Exp}(C^\infty(M, \mathbb{C}))$ acts on \mathcal{U} as a subgroup of DC , preserving each fiber $\Pi^{-1}\gamma$. This action is clearly free on each fiber, and it turns out to be transitive on each connected component of the fiber.

We emphasize that both components of the DC_γ -action, $\text{Diff}_\gamma(M)$ and $\text{Exp}(C^\infty(M, \mathbb{C}))$, must work together to achieve connected component-wise transitivity. Indeed, $\mathcal{U}_\gamma := \Pi^{-1}\gamma \cap \mathcal{U}$ (and hence $\Pi^{-1}\gamma$) may not be attained by the action of either $\text{Exp}(C^\infty(M, \mathbb{C}))$ or $\text{Diff}_\gamma(M)$ alone, as illustrated in the following examples.

Example 3.3.1 (Non-transitivity of $\text{Diff}_\gamma(M)$ -action). Let γ be four points $\{z_k\}_{k=1}^4$ on $M = \mathbb{R}^2$ where z_1, z_2 have positive orientations and z_3, z_4 have negative orientations. We then let ψ_0 be an implicit representation of γ defined by $\psi_0(z) = (z - z_1)(z - z_2)(\overline{z - z_3})(\overline{z - z_4})$, and $\psi_1 := \psi_0 e^{i\pi/2}$ as illustrated in Figure 3.2. These two functions clearly lie in the same fiber $\Pi^{-1}\gamma$, but there is no diffeomorphism f such that $\psi_1 = \psi_0 \circ f$ resolving the topological differences of the phase level sets. In contrast, the group action of $e^{i\pi/2} \in \text{Exp}(C^\infty(M, \mathbb{C}))$ joins ψ_0 and ψ_1 .



Figure 3.2: Implicit representations ψ_0 (left) and $\psi_1 = \psi_0 e^{i\pi/2}$ (right) of four points in \mathbb{R}^2 . Each color indicates a phase value of $\phi = \psi/|\psi|$. The highlighted light-blue and red curves correspond to the level sets $\phi^{-1}(1)$ and $\phi^{-1}(e^{i\pi})$, respectively. Clearly, there is no diffeomorphism f such that $\psi_0 \circ f = \psi_1$ that can handle the topological changes in these level sets.

Example 3.3.2 (Non-transitivity of $\text{Exp}(C^\infty(M, \mathbb{C}))$ -action). Let $M = \mathbb{R}^2$, and consider implicit representations ψ_0, ψ_1 defined by $\psi_0(x, y) = x + iy$ and $\psi_1(x, y) = x + y + iy$. These functions share the same zero at $(x, y) = (0, 0)$ and the same orientation. However, there is no nowhere-vanishing smooth function φ such that $\psi_1 = \varphi\psi_0$. Indeed, the quotient $\frac{\psi_1}{\psi_0} = 1 + \frac{xy}{x^2+y^2} - i\frac{y^2}{x^2+y^2}$ is discontinuous at the origin. On the other hand, with the diffeomorphism $f(x, y) = (x - y, y)$, we attain $\psi_1 = \psi_0 \circ f$.

The non-transitivity of the $\text{Diff}_\gamma(M)$ -action and the $\text{Exp}(C^\infty(M, \mathbb{C}))$ -action separately holds on any ambient manifold M , as it is caused solely by local information of implicit representations near the zero level set. For instance, on $M = \mathbb{S}^2$ regarded as the one-point compactification of \mathbb{R}^2 , the same choice of pairs ψ_0, ψ_1 as in Example 3.3.1 and Example 3.3.2 works, with the modification that $|\psi_0|, |\psi_1| \rightarrow 1$ at infinity so that they are smooth functions on \mathbb{S}^2 .

We now show that the action of DC_γ is indeed transitive on each connected component of the fiber $\Pi^{-1}\gamma$, and that the number of connected components is determined by the topology

of the ambient manifold. Let $H_{dR}^1(M, \mathbb{Z})$ be the first integral de Rham cohomology defined by

$$H_{dR}^1(M, \mathbb{Z}) = \left\{ [\eta] \in H_{dR}^1(M, \mathbb{R}) \mid \int_{\sigma} \eta \in \mathbb{Z} \text{ for any } [\sigma] \in H_1(M, \mathbb{Z}) \right\}.$$

Then we have the following relation;

Proposition 3.3.3. *For each $\gamma \in \mathcal{O}$, we have a bijection*

$$\Pi^{-1}\gamma / \text{DC}_{\gamma} \cong H_{dR}^1(M, \mathbb{Z}).$$

To prove the proposition and highlight the roles of $\text{Diff}_{\gamma}(M)$ -action and $\text{Exp}(C^{\infty}(M, \mathbb{C}))$ -action separately, we introduce two equivalence classes.

Definition 3.3.4 (Twist class and conformal class). We say $\psi_0, \psi_1 \in \Pi^{-1}\gamma$ are in the same *twist class* if there is a diffeomorphism $f \in \text{Diff}_{\gamma}(M)$ such that $\psi_1 = \psi_0 \circ f$.

We say $\psi_0, \psi_1 \in \Pi^{-1}\gamma$ are in the same *conformal class* if there is a nowhere vanishing function $\zeta \in C^{\infty}(M, \mathbb{C}^{\times})$ such that $\psi_1 = \zeta\psi_0$.

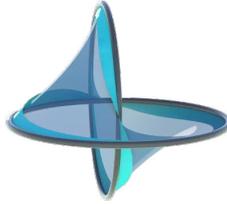


Figure 3.3: A ribbon of an implicit representation ψ for the Hopf link γ . For each regular value $s \in \mathbb{S}^1$ of the phase field $\phi = \psi/|\psi|$, the ribbon R_s (opaque cyan) is defined as the intersection of $\phi^{-1}(s)$ (translucent blue) and a small tubular neighborhood of $\text{im } \gamma$.

Intuitively, ψ_0 and ψ_1 are in the same twist class if their *ribbons* (Figure 5.3) have the same total twists and hence are diffeomorphic to each other. On the other hand, ψ_0 and ψ_1 are in the same conformal class if they share the same shear states (*i.e.*, the rate of phase change) around the zeros so that ψ_1/ψ_0 exists and smooth on the zeros.

Proof of Proposition 3.3.3. We first claim that for each pair $\psi_0, \psi_1 \in \Pi^{-1}\gamma$, there exist functions $e^{\varphi_0}, e^{\varphi_1} \in \text{Exp}(C^{\infty}(M, \mathbb{C}))$ and a diffeomorphism $f \in \text{Diff}_{\gamma}(M)$ such that $(e^{\varphi_0}\psi_0) \circ f$ and $e^{\varphi_1}\psi_1$ lie in the same conformal class.

To see this, take some $\varphi_0, \varphi_1 \in C^{\infty}(M, \mathbb{C})$ so that, letting $\hat{\psi}_0 := e^{\varphi_0}\psi_0$ and $\hat{\psi}_1 := e^{\varphi_1}\psi_1$, we have $|\hat{\psi}_0| = |\hat{\psi}_1|$ in a small tubular neighborhood B of $\text{im } \gamma$, and the phase field $\hat{\phi}_j = \frac{\hat{\psi}_j}{|\hat{\psi}_j|} : B \setminus \text{im } \gamma \rightarrow \mathbb{S}^1$ is a submersion for $j = 0, 1$.

We can take coordinates (p, z) in B with $p \in \text{im } \gamma$ and $z := \hat{\psi}_0(p, \cdot)$. Then consider a tubular sub-neighborhood $B_R = \{(p, z) \mid p \in \text{im } \gamma, |z| \leq R\}$ small enough so that $\hat{\psi}_1$ is injective on each circle $S_p^{\epsilon} = \{(p, z) \mid |z| = \epsilon\}$ for $p \in \text{im } \gamma$ and $\epsilon \leq R$. Then for a sufficiently small $r < R$, there exists $f \in \text{DC}_{\gamma}$ such that $\hat{\psi}_0 \circ f = \tau\hat{\psi}_1$ inside the smaller tubular neighborhood $B_r = \{(p, z) \mid p \in \text{im } \gamma, |z| \leq r\}$ and $f = \text{id}$ outside B_R , as illustrated in Figure 3.4. Here τ is a smooth function on B_r of the form $\tau(p, z) = e^{i\theta(p)}$ with some \mathbb{T}^1 -valued function θ .

Outside B_r , the function $\frac{\hat{\psi}_1}{\hat{\psi}_0 \circ f}$ is smooth and nowhere vanishing. Hence there exists a smooth function $\zeta \in C^\infty(M, \mathbb{C}^\times)$ such that $\hat{\psi}_0 \circ f = \zeta \hat{\psi}_1$ and $\zeta = \tau$ inside B_r , showing that $\psi_0 \circ f$ and ψ_1 are in the same conformal class.

We now note that the following is an exact sequence:

$$0 \rightarrow \mathbb{Z} \xrightarrow{k \mapsto 2\pi i k} C^\infty(M, \mathbb{C}) \xrightarrow{\text{exp}} C^\infty(M, \mathbb{C}^\times) \xrightarrow{\mathfrak{h}} H_{dR}^1(M, \mathbb{Z}) \rightarrow 0$$

where $\mathfrak{h}: C^\infty(M, \mathbb{C}^\times) \rightarrow H_{dR}^1(M, \mathbb{Z})$ is defined by $\mathfrak{h}(\kappa) = [\text{Im} \frac{1}{2\pi} \frac{d\kappa}{\kappa}]$. Therefore ζ defines an element $\mathfrak{h}(\zeta) \in H_{dR}^1(M, \mathbb{Z})$, and the functions $\hat{\psi}_0 \circ f$ and $\hat{\psi}_1$ are in the same twist class if and only if $\mathfrak{h}(\zeta) = [0]$.

On the other hand, each pair of elements $[\eta_1], [\eta_2] \in H_{dR}^1(M, \mathbb{Z})$ define elements $[\zeta_1], [\zeta_2] \in C^\infty(M, \mathbb{C}^\times) / \text{Exp}(C^\infty(M, \mathbb{C}))$ such that $\zeta_1 \psi$ and $\zeta_2 \psi$ for any choice of $\psi \in \Pi^{-1}\gamma$, and the representatives ζ_1, ζ_2 are in the same DC_γ -orbit if and only if $[\eta_1] = [\eta_2]$. This completes the proof. \square

Hence we have shown that, while a DC -orbit \mathcal{U} is path connected, each fiber $\Pi^{-1}\gamma$ consists of a disjoint copies of a DC_γ -orbit associated with the discrete group $H_{dR}^1(M, \mathbb{Z})$.

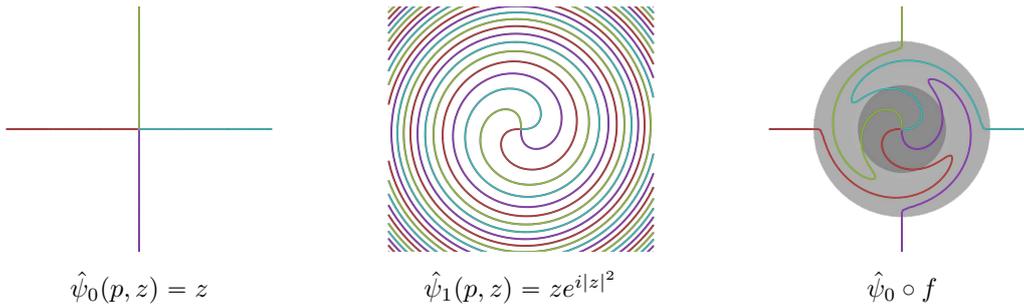


Figure 3.4: Level curves $\hat{\phi}^{-1}(1), \hat{\phi}^{-1}(e^{i\frac{\pi}{2}}), \hat{\phi}^{-1}(-1), \hat{\phi}^{-1}(e^{i\frac{3\pi}{2}})$ of the phase maps $\hat{\phi} = \hat{\psi}/|\hat{\psi}|$ of implicit representations $\hat{\psi}_0(p, z) = z$ (left), $\hat{\psi}_1(p, z) = ze^{i|z|^2}$ (middle), and $\hat{\psi}_0 \circ f$ (right), on the tubular domain B sliced into a disc at fixed $p \in \text{im } \gamma$. The diffeomorphism f rotates these level lines on each circle $S_p^\epsilon = \{(p, z) \mid |z| = \epsilon\}$ for $\epsilon \leq R$ so that $\hat{\psi}_0 \circ f$ agrees with $\hat{\psi}_1$ inside the smaller domain B_r (dark gray) and with $\hat{\psi}_0$ outside the larger domain B_R (light gray).

Is the fiber bundle $\mathcal{F}_\mathcal{O} := \Pi^{-1}\mathcal{O}$ larger than a single orbit \mathcal{U} ? Since each fiber $\Pi^{-1}\gamma$ is an orbit of $\text{Diff}_\gamma \times C^\infty(M, \mathbb{C}^\times)$, the entire fiber bundle $\mathcal{F}_\mathcal{O} := \Pi^{-1}\mathcal{O}$ over a fixed \mathcal{O} is a $\text{Diff}_0(M) \times C^\infty(M, \mathbb{C}^\times)$ -orbit. As each fiber has multiple connected components when $H_{dR}^1(M, \mathbb{Z}) \neq 0$, the fiber bundle $\mathcal{F}_\mathcal{O}$ appears to be strictly larger than \mathcal{U} , a single orbit under the action of $\text{DC} = \text{Diff}_0(M) \times \text{Exp}(C^\infty(M, \mathbb{C}))$. It, however, turns out that $\mathcal{F}_\mathcal{O} = \mathcal{U}$ at last for the case of the base dimension 2 as illustrated in Figure 3.5. By moving around zeros while multiplying functions in $\text{Exp}(C^\infty(M, \mathbb{C}))$, we can overcome topological changes of phase level lines. We expect that a similar argument holds in higher dimensional cases and the result remains valid:

Conjecture 3.3.5. *Let \mathcal{O} be a $\text{Diff}_0(M)$ -orbit in the explicit shape space \mathcal{X} , $\mathcal{F}_\mathcal{O} := \Pi^{-1}\mathcal{O}$ be the entire fiber bundle over \mathcal{O} , and \mathcal{U} be a DC -orbit in $\mathcal{F}_\mathcal{O}$. Suppose each connected component of M intersects with $\text{im } \gamma$ for $\gamma \in \mathcal{O}$. Then $\mathcal{F}_\mathcal{O} = \mathcal{U}$.*

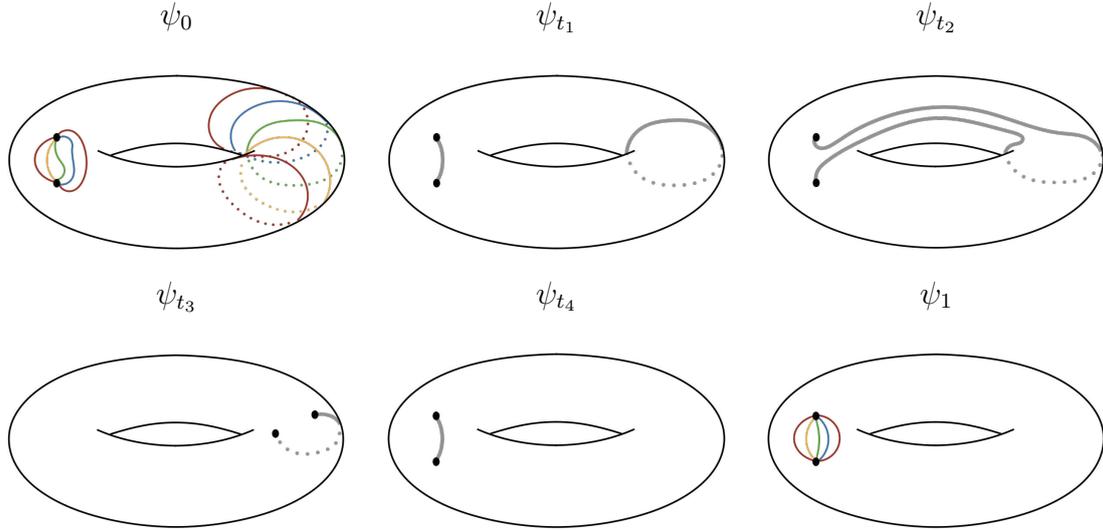


Figure 3.5: A transition of the implicit representation from ψ_0 (top left) to ψ_1 (bottom right) for two points γ on the torus \mathbb{T}^2 . In both ψ_0 and ψ_1 , the colored curves represent level sets $\phi^{-1}(s)$ of some phases. Clearly, ψ_0 and ψ_1 belong to different connected components of the fiber $\pi^{-1}\gamma$. Nevertheless, they lie on the same DC-orbit. Starting from ψ_0 , we first multiply by some $e^{\varphi_1} \in \text{Exp}(\mathbb{C}^\infty(M, \mathbb{C}))$ to *compress* the support of $d\phi$ into narrow bands, visualized as thick gray curves in ψ_{t_1} . Next, multiplying ψ_{t_1} by $e^{\varphi_2} \in \text{Exp}(\mathbb{C}^\infty(M, \mathbb{C}))$ merges two level curves, producing ψ_{t_2} . Applying a diffeomorphism f_1 moves the level lines, yielding ψ_{t_3} , and a second diffeomorphism f_2 transforms ψ_{t_3} into ψ_{t_4} . Finally, multiplying by another $e^{\varphi_3} \in \text{Exp}(\mathbb{C}^\infty(M, \mathbb{C}))$ decompresses the support of $d\phi$, resulting in ψ_1 .

We conclude this subsection by observing that an $\text{Diff}_0(M) \times C^\infty(M, \mathbb{C}^\times)$ -orbit $\mathcal{F}_\mathcal{O}$ is locally a DC-orbit \mathcal{U} and they share the same tangent space at each ψ even if $\mathcal{U} \subsetneq \mathcal{F}_\mathcal{O}$. In this section and the next section, we still focus on local arguments on \mathcal{U} , which are all available on $\mathcal{F}_\mathcal{O}$ when we consider prequantum G -bundles in Section 4.3.

Prequantum structure on the space of implicit representations

In this chapter, we further study the geometry of the implicit shape space $\mathcal{F}_{\mathcal{O}}$, related to the Marsden–Weinstein symplectic structure on the explicit shape space \mathcal{O} . In particular, we construct a prequantum bundle over \mathcal{O} as a certain quotient space of $\mathcal{F}_{\mathcal{O}}$. This reveals a new geometric interpretation of the Marsden–Weinstein form.

4.1 Overview

As we saw in Chapter 3, the fiber bundle $\mathcal{F}_{\mathcal{O}}$ of implicit representations has a natural geometric interpretation. Each complex function $\psi \in \mathcal{F}_{\mathcal{O}}$ representing a codimension-2 submanifold $\gamma \in \mathcal{O}$ carries additional information in the form of its complex phase. The level sets of this phase function constitute a \mathbb{S}^1 -family of hypersurfaces in the ambient manifold M , that are bordered by γ . For the motion of each of these hypersurfaces, we can consider the swept volume. Hence any motion of ψ induces the average volume swept out by these phase hypersurfaces.

From this observation, we introduce a quotient space of $\mathcal{F}_{\mathcal{O}}$. In each fiber, we define two hypersurface configurations as equivalent (\sim) if they can be continuously deformed into each other while keeping their boundaries fixed and maintaining zero net volume change throughout the motion. The space of equivalence classes is denoted by $\mathcal{P} := \mathcal{F}_{\mathcal{O}} / \sim$, which is still a fiber bundle over \mathcal{O} . We call \mathcal{P} the *volume class* for the implicit shape space.

It turns out that the volume class \mathcal{P} over \mathcal{O} carries a *prequantum bundle* structure, where each fiber consists of multiple copies of a circle, reflecting the first integral de Rham cohomology $H_{dR}^1(M, \mathbb{Z})$. This is a relaxation of a standard definition of a prequantum bundle which usually requires each fiber to be a complex line or a circle, while each fiber of our bundle is still one dimensional.

The connection form $\Theta_{\mathcal{P}}$ on this bundle is defined such that the horizontal lift of any path in \mathcal{O} corresponds to a motion of phase hypersurfaces that sweep out zero net volume along the path. With this setup, we have the following theorem.

Theorem 4.1.1 (Theorem 4.3.7). *The fibration $\Pi_{\mathcal{P}}: (\mathcal{P}, \Theta_{\mathcal{P}}) \rightarrow (\mathcal{O}, \omega)$ is a prequantum G -bundle where ω is the MW symplectic form and the structure group is $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$.¹*

In other words, the curvature of the connection $\Theta_{\mathcal{P}}$ agrees with the MW symplectic form. As a consequence, this prequantum bundle allows the following geometric interpretation of the MW form:

Corollary 4.1.2 (Corollary 4.3.9). *Consider a closed path $\partial\Sigma$ in \mathcal{O} that bounds a 2-dimensional disk Σ , representing a cyclic motion of a codimension-2 submanifold $\gamma_t \subset M$ for $0 \leq t \leq 1$, with $\gamma_0 = \gamma_1$. Let $[\psi_t]_{\mathcal{P}}$ be a horizontal lift over γ_t and $\phi_t := \psi_t/|\psi_t|: M \setminus \psi_t^{-1}(0) \rightarrow \mathbb{S}^1$ be the phase map of a representative $\psi_t \in [\psi_t]_{\mathcal{P}}$. That is, γ_t bounds a family of hypersurfaces $\{\sigma_t^s\}_{s \in \mathbb{S}^1}$, defined by $\sigma_t^s = \phi_t^{-1}(s)$ and the average volume swept out by σ_t^s remains zero at each t , i.e.,*

$$\frac{1}{2\pi} \int_{\mathbb{S}^1} \int_{\sigma_t^s} \iota_{\partial_t \sigma_t^s} \mu \, ds = 0,$$

where μ is the volume form on M .

Then, the volume enclosed between σ_0^s and σ_1^s , averaged over $s \in \mathbb{S}^1$, equals to $\iint_{\Sigma} \omega$, where ω is the MW form on \mathcal{O} .

In the limiting case of Corollary 4.1.2, where the phase of ψ_t becomes constant except a 2π jump at a single hypersurface σ_t bounding γ_t , the result simplifies to the following corollary. This version has no explicit reference to the complex function ψ_t .

Corollary 4.1.3 (Corollary 4.3.10). *Let Σ be a disc and $\{\gamma_t\}_{t \in [0,1]}$ be a path in \mathcal{O} as in Corollary 4.1.2. Suppose that each γ_t bounds a hypersurface, i.e., $\gamma_t = \partial\sigma_t$, and that the volume swept out by σ_t remains zero at each t . Then, the volume enclosed between σ_0 and σ_1 agrees with $\iint_{\Sigma} \omega$.*

Related work on prequantum bundles over the codimension-2 shape spaces There are other prequantum bundles over (\mathcal{O}, ω) . In [HV03], Haller and Vizman constructed a prequantum circle bundle by covering \mathcal{O} with a collection of open sets $\{U_i\}_i$ and gluing the sets $\{\mathbb{S}^1 \times U_i\}_i$ using carefully chosen transition functions that ensure a prequantum bundle structure. Another approach, utilizing differential characters, was introduced in [DJNV20]. Presently we are not aware of any explicit relationships between their proposed prequantum bundles and ours.

However, we emphasize that, unlike these previously established prequantum bundle structures, our framework is inherently local. Specifically, we focus on submanifolds that are exact in homology, meaning they must bound codimension-1 domains. Consequently, the existence of our prequantum bundle is restricted to such exact submanifolds as explained in Section 3.2.

Future work The codimension-2 shape space admits not only the MW structure but also other symplectic structures, as recently shown in [BIM24]. In this article, we focus specifically on the MW structure. Extending our framework to a wider class of symplectic structures is an interesting direction for future work.

¹Precisely, we normalize Θ and ω in Theorem 4.3.7 so that the normalized Θ is a connection form in a standard definition of principal bundle. For details, see Section 4.3.

Secondly, we restrict our attention to the space of embeddings modulo reparametrization, meaning that self-intersecting shapes are outside the scope of this study, even though the MW form is defined on the space of immersions. On the other hand, the tools we employ from geometric measure theory, such as currents, are capable of handling such singularities. In the future, we aim to extend our framework to accommodate topological changes and reconnection events in shapes.

Organization of the chapter In Section 4.2, we define a 1-form Θ on $\mathcal{F}_{\mathcal{O}}$ and show that it defines a Liouville form for the MW form $\omega \in \Omega^2(\mathcal{O})$ in the prequantum sense, i.e., $\Pi^*\omega = d\Theta$ via the fibration $\Pi: \mathcal{F}_{\mathcal{O}} \rightarrow \mathcal{O}$. Our proof, using geometric measure theory, offers a geometric interpretation of Θ in terms of the swept volumes of phase hypersurfaces.

In Section 4.3, we define the volume class \mathcal{P} as explained above. We show that Θ factored onto \mathcal{P} becomes a connection form $\Theta_{\mathcal{P}}$ forming a prequantum G -bundle $\Pi_{\mathcal{P}}: (\mathcal{P}, \Theta_{\mathcal{P}}) \rightarrow (\mathcal{O}, \omega)$ with the structure group $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$.

In Section 4.4, we study additional structures that implicit representations and prequantum structures reveal, such as horizontal Hamiltonian vector fields on \mathcal{P} with respect to the connection $\Theta_{\mathcal{P}}$, and a degenerate formal Kähler structure on $\mathcal{F}_{\mathcal{O}}$.

4.2 Liouville form in the prequantum sense

In this section, we present one of our main results: the MW symplectic form ω on \mathcal{O} over a general closed manifold M admits a Liouville form in the *prequantum sense*. In the rest of the chapter, we assume that the ambient manifold M is closed, in contrast to the setting in Section 2.2 for the explicit shape space.

Definition 4.2.1. Define the following 1-form on each DC-orbit \mathcal{U} of \mathcal{F} :

$$\Theta_{\psi}(\dot{\psi}) = \frac{1}{2\pi} \int_M \operatorname{Im} \frac{\dot{\psi} \bar{\psi}}{|\psi|^2} \mu, \quad (4.2.1)$$

where $\dot{\psi} = -\mathcal{L}_u \psi + \varphi \psi \in T_{\psi} \mathcal{U}$ for some $(u, \varphi) \in \mathfrak{DC}$.

At first glance, the integral in (4.2.1) may appear divergent as the integrand is unbounded near the zeros of ψ . However, we will show that it is in fact finite. We also demonstrate that the value $\Theta_{\psi}(\dot{\psi})$ has a clear physical interpretation (Remark 4.2.15).

Definition 4.2.2 (Formal prequantization and Liouville form). Let (B, β) be a symplectic manifold, and let E be a manifold equipped with a 1-form α . We say that a fibration $\pi: E \rightarrow B$ is a *formal prequantization*, and that α is a *Liouville form* of β in the *prequantum sense*, if

$$d\alpha = \pi^* \beta.$$

Theorem 4.2.3 (Liouville form for the MW structure in the prequantum sense). *Suppose M is a closed manifold equipped with a volume form μ . Then the fibration $\Pi: (\mathcal{U}, \Theta) \rightarrow (\mathcal{O}, \omega)$ is a formal prequantization.*

Proof strategy To prove the theorem, that is, to show $d\Theta = \Pi^*\omega$, we combine a standard approach in differential geometry with the coarea formula from geometric measure theory. We compute the exterior derivative of Θ_ψ by the formula

$$d\Theta_\psi(\dot{\psi}, \dot{\psi}) = \mathcal{L}_{Y_1} \iota_{Y_2} \Theta_\psi - \mathcal{L}_{Y_2} \iota_{Y_1} \Theta_\psi - \iota_{[Y_1, Y_2]} \Theta_\psi \quad (4.2.2)$$

for each $\psi \in \mathcal{U}$ and $\dot{\psi}, \dot{\psi} \in T_\psi \mathcal{U}$. Here, the vector fields $Y_1, Y_2 \in \Gamma(T\mathcal{U})$ are arbitrary extensions of the vectors $\dot{\psi}, \dot{\psi}$ respectively; that is, $Y_1|_\psi = \dot{\psi}$ and $Y_2|_\psi = \dot{\psi}$.

To perform this calculation, we first choose specific vector fields Y_1 and Y_2 and explicitly compute each term in (4.2.2). This calculation leads to a simple expression of $d\Theta$ (Lemma 4.2.12). As the final step of the proof, we show that the fibration $\Pi : \mathcal{U} \rightarrow \mathcal{D}_\mathcal{O} \subset \mathcal{D}_{m-2}(M)$ is decomposed into $\Pi = \partial \circ \Lambda$ (Lemma 4.2.13) using the coarea formula, where the map $\Lambda : \mathcal{U} \rightarrow \mathcal{D}_{m-1}(M)$ assigns to each ψ its *circle differential* $(m-1)$ -current, as defined below (Definition 4.2.4), and $\partial : \mathcal{D}_{m-1}(M) \rightarrow \mathcal{D}_{m-2}(M)$ is the boundary operator for currents.

Definition 4.2.4 (Circle differential 1-form and $(m-1)$ -current). Using the normalization map $\tau : \mathbb{C}^\times \rightarrow \mathbb{S}^1$ given by $z \mapsto \frac{z}{|z|}$ and the standard Haar measure $\sigma \in \Omega^1(\mathbb{S}^1)$ on \mathbb{S}^1 with $\int_{\mathbb{S}^1} \sigma = 2\pi$, define the *circle differential 1-form* $\lambda_\psi \in \Omega^1(M \setminus \psi^{-1}(0))$ by

$$\lambda_\psi := \psi^* \tau^* \sigma.$$

This induces the map $\Lambda : \mathcal{F} \rightarrow \mathcal{D}_{m-1}(M)$ by

$$\langle \Lambda(\psi), \beta \rangle := \int_M \lambda_\psi \wedge \beta, \quad \beta \in \Omega^{m-1}(M). \quad (4.2.3)$$

We call $\Lambda(\psi)$ the *circle differential current*.

The circle differential λ_ψ represents the gradient of the phase of ψ . In fact, there is a small neighborhood N_x around each point $x \in M \setminus \psi^{-1}(0)$ such that ψ is locally represented by $\psi(y) = r(y) \cdot e^{i\phi(y)}$ with some function $r : N_x \rightarrow \mathbb{R}_{>0}$ and $\phi : N_x \rightarrow \mathbb{R}$ with $d\phi = \lambda_\psi$.

The integral (4.2.3) is well-defined because the integrand $\lambda_\psi \wedge \alpha$ is defined almost everywhere on M with respect to μ , in particular, except on the codimension-2 set $\psi^{-1}(0)$. In fact, $\Lambda(\psi)$ is indeed a current *i.e.*, a continuous linear functional on $\Omega^{m-1}(M)$, and namely the value (4.2.3) is always finite;

Lemma 4.2.5. *Let M be a closed manifold and $\psi \in \mathcal{F}$. Then $\Lambda(\psi) \in \mathcal{D}_{m-1}(M)$.*

This result is a direct consequence of the coarea formula from geometric measure theory, which is a nonlinear version of the Fubini theorem.

Proposition 4.2.6 (Smooth coarea formula). *Let X be an oriented smooth manifold and Y be an oriented compact smooth manifold with dimension n_X and $n_Y < n_X$ respectively. For a submersion $f : X \rightarrow Y$ and forms $\beta \in \Omega^{n_X - n_Y}(Y), \sigma \in \Omega^{n_Y}(Y)$, we have*

$$\int_X f^* \sigma \wedge \beta = \int_Y \langle \delta_{f^{-1}(\cdot)}, \beta \rangle \sigma.$$

Remark 4.2.7. There are different versions of the coarea formula. Here, we use the version from [Dem97, Chapter I.3], as it is formulated in terms of differential forms and does not rely on Hausdorff measure or a Riemannian metric, making it well-suited for our context.

For standard references on the coarea formula, readers may find works like [Fed14] and [Nic11] helpful.

Proof of Lemma 4.2.5. Clearly $\Lambda(\psi)$ is linear in $\mathcal{D}_{m-1}(M)$. We now show that $\Lambda(\psi)$ is a bounded functional. By the coarea formula (Proposition 4.2.6), we have for any $\beta \in \Omega^{m-1}(M)$ that,

$$\langle \Lambda(\psi), \beta \rangle = \int_M \lambda_\psi \wedge \beta = \int_{M^{\text{reg } \phi}} \lambda_\psi \wedge \beta = \int_{\mathbb{S}^1} \langle \delta_{(\phi^{-1}(\cdot) \cap M^{\text{reg } \phi})}, \beta \rangle \sigma$$

where $\phi = \tau \circ \psi: M \setminus \psi^{-1}(0) \rightarrow \mathbb{S}^1$ and $M^{\text{reg } \phi}$ denotes the set of regular points of ϕ . Note that the set $\phi^{-1}(s) \cap M^{\text{reg } \phi}$ for each $s \in \mathbb{S}^1$ is a $m-1$ dimensional submanifold and namely $\delta_{\partial(\phi^{-1}(s) \cap M^{\text{reg } \phi})}$ is a $m-1$ -current. In particular, the operator norm $\|\delta_{\partial(\phi^{-1}(s) \cap M^{\text{reg } \phi})}\|_{\mathcal{D}_{m-1}(M)}$ is bounded. Therefore,

$$\|\Lambda(\psi)\|_{\mathcal{D}_{m-1}(M)} \leq \int_{\mathbb{S}^1} \sigma \sup_{s \in \mathbb{S}^1} \|\delta_{\partial(\phi^{-1}(s) \cap M^{\text{reg } \phi})}\|_{\mathcal{D}_{m-1}(M)} = 2\pi \sup_{s \in \mathbb{S}^1} \|\delta_{\partial(\phi^{-1}(s) \cap M^{\text{reg } \phi})}\|_{\mathcal{D}_{m-1}(M)}$$

is bounded, from which we conclude that $\Lambda(\psi)$ is a current. \square

As we see from the proof, the current $\Lambda(\psi)$ is in fact the superposition of a \mathbb{S}^1 -family of hypersurfaces given as the phase level sets of ψ .

Vector fields $Y_1, Y_2 \in \Gamma(T\mathcal{U})$ We now choose the vector fields Y_1 and Y_2 that extend $\dot{\psi}$ and $\dot{\psi}$ to evaluate (4.2.2). As in our approach for Theorem 2.3.3, we use the fundamental vector field associated with the action of $\text{DC} = \text{Diff}_0(M) \times \text{Exp}(C^\infty(M, \mathbb{C}))$ on \mathcal{U} . Recalling that the Lie algebra of DC is $\mathfrak{DC} = \text{diff}(M) \times C^\infty(M, \mathbb{C})$, we define $Y^{(u,a)} \in \Gamma(T\mathcal{U})$ for each $(u, a) \in \mathfrak{DC}$ by

$$Y^{(u,a)}(\psi) := \left. \frac{d}{dt} \right|_{t=0} \text{Fl}_u^{-t*} \psi \cdot e^{ta} = -\mathcal{L}_u \psi + a\psi. \quad (4.2.4)$$

We first explicitly compute their Lie bracket;

Lemma 4.2.8. For $(u, a), (v, b) \in \mathfrak{DC}$ and $\psi \in \mathcal{U}$, we have

$$[Y^{(u,a)}, Y^{(v,b)}](\psi) = -\mathcal{L}_{[u,v]}\psi + (\mathcal{L}_u b - \mathcal{L}_v a)\psi.$$

Proof. To compute the Lie derivative of vector fields $[Y^{(u,a)}, Y^{(v,b)}] = \mathcal{L}_{Y^{(u,a)}} Y^{(v,b)}$, we use the formula:

$$\mathcal{L}_{Y^{(u,a)}} Y^{(v,b)}(\psi) = \left. \frac{d}{dt} \right|_{t=0} D\Phi_{Y^{(u,a)}}^{-t} (Y^{(v,b)}(\Phi_{Y^{(u,a)}}^t(\psi))), \quad \psi \in \mathcal{U}$$

where $\Phi_Y^t: \mathcal{U} \rightarrow \mathcal{U}$ is the flow map along $Y \in \Gamma(T\mathcal{U})$ defined by the ODE

$$\begin{aligned} \frac{d}{dt}(\Phi_Y^t(\psi)) &= Y(\Phi_Y^t(\psi)), \\ \Phi_Y^0 &= \text{id}_{\mathcal{U}}, \end{aligned}$$

and $D\Phi_Y^t|_\psi: T_\psi \mathcal{U} \rightarrow T_{\Phi_Y^t(\psi)} \mathcal{U}$ is its differential at ψ . For the fundamental vector field $Y^{(u,a)}$ associated to $(u, a) \in \mathfrak{DC}$, we have by definition $\Phi_{Y^{(u,a)}}^t(\psi) = \text{Fl}_u^{-t*} \psi \cdot e^{ta}$. Since

$$Y^{(v,b)}(\Phi_{Y^{(u,a)}}^t(\psi)) = \left. \frac{d}{ds} \right|_{s=0} \Phi_{Y^{(v,b)}}^s (\Phi_{Y^{(u,a)}}^t(\psi)),$$

we have

$$\begin{aligned} \frac{d}{dt} \Big|_{t=0} D\Phi_{Y^{(u,a)}}^{-t} \left(Y^{(v,b)}(\Phi_{Y^{(u,a)}}^t(\psi)) \right) &= \frac{d}{dt} \Big|_{t=0} \frac{d}{ds} \Big|_{s=0} \Phi_{Y^{(u,a)}}^{-t} \left(\Phi_{Y^{(v,b)}}^s \left(\Phi_{Y^{(u,a)}}^t(\psi) \right) \right) \\ &= \frac{d}{dt} \Big|_{t=0} \frac{d}{ds} \Big|_{s=0} \Phi_{Y^{(u,a)}}^{-t} \left(\Phi_{Y^{(v,b)}}^s(\psi) \right) + \Phi_{Y^{(v,b)}}^s \left(\Phi_{Y^{(u,a)}}^t(\psi) \right) \end{aligned}$$

using the product rule. We then compute the first term:

$$\begin{aligned} \frac{d}{dt} \Big|_{t=0} \frac{d}{ds} \Big|_{s=0} \Phi_{Y^{(u,a)}}^{-t} \left(\Phi_{Y^{(v,b)}}^s(\psi) \right) &= \frac{d}{dt} \Big|_{t=0} \frac{d}{ds} \Big|_{s=0} \text{Fl}_u^t * \left(\text{Fl}_v^{-s*} \psi \cdot e^{sb} \right) \cdot e^{-ta} \\ &= \frac{d}{dt} \Big|_{t=0} \frac{d}{ds} \Big|_{s=0} \text{Fl}_u^t * \left(\text{Fl}_v^{-s*} \psi \right) \cdot \text{Fl}_u^t * e^{sb} \cdot e^{-ta} \\ &= \frac{d}{ds} \Big|_{s=0} \mathcal{L}_u \left(\text{Fl}_v^{-s*} \psi \right) \cdot e^{sb} + \left(\text{Fl}_v^{-s*} \psi \right) \mathcal{L}_u \left(e^{sb} \right) - \left(\text{Fl}_v^{-s*} \psi \right) \cdot e^{sb} \cdot a \\ &= -\mathcal{L}_u \mathcal{L}_v \psi + \mathcal{L}_u \psi \cdot b - \mathcal{L}_v \psi \cdot 0 + \psi \mathcal{L}_u b + \mathcal{L}_v \psi \cdot a - \psi \cdot b \cdot a \\ &= -\mathcal{L}_u \mathcal{L}_v \psi + b \mathcal{L}_u \psi + a \mathcal{L}_v \psi + \psi \mathcal{L}_u b - ab\psi \end{aligned}$$

where we used $\mathcal{L}_u(e^{sb}) = s\mathcal{L}_u b \cdot e^{sb}$. Similarly, we have

$$\frac{d}{dt} \Big|_{t=0} \frac{d}{ds} \Big|_{s=0} \Phi_{Y^{(v,b)}}^s \left(\Phi_{Y^{(u,a)}}^t(\psi) \right) = \mathcal{L}_v \mathcal{L}_u \psi - a \mathcal{L}_v \psi - b \mathcal{L}_u \psi - \psi \mathcal{L}_v a + ab\psi.$$

Summing these results, we have

$$\begin{aligned} \frac{d}{dt} \Big|_{t=0} \frac{d}{ds} \Big|_{s=0} \Phi_{Y^{(u,a)}}^{-t} \left(\Phi_{Y^{(v,b)}}^s(\psi) \right) + \Phi_{Y^{(v,b)}}^s \left(\Phi_{Y^{(u,a)}}^t(\psi) \right) &= -\mathcal{L}_u \mathcal{L}_v \psi + \mathcal{L}_v \mathcal{L}_u \psi + \psi \mathcal{L}_v b - \psi \mathcal{L}_v a \\ &= -\mathcal{L}_{[u,v]} \psi + (\mathcal{L}_u b - \mathcal{L}_v a) \psi. \end{aligned}$$

We thus obtained the stated expression. \square

Remark 4.2.9. The fundamental vector field mapping $\mathfrak{DC} \rightarrow \Gamma(T\mathcal{U})$ given by (4.2.4) is in fact a Lie algebra anti-homomorphism. To see this, notice that the mapping $\text{diff}(M) \rightarrow \text{Der}(C^\infty(M, \mathbb{C}))$ by $u \rightarrow \mathcal{L}_u$ is an anti-homomorphism as noted in Remark 2.3.5. Hence we have

$$Y^{[(u,a),(v,b)]} = Y^{([u,v], -\mathcal{L}_u b + \mathcal{L}_v a)} = \mathcal{L}_{[u,v]} \psi + (\mathcal{L}_v a - \mathcal{L}_u b) \psi = -[Y^{(u,a)}, Y^{(v,b)}]$$

where the last equality is the application of Lemma 4.2.8.

Evaluation of Θ and $d\Theta$ We now evaluate each term in (4.2.2) with $Y_1 = Y^{(u,a)}$ and $Y_2 = Y^{(v,b)}$. First, we express the term $\Theta_\psi(Y^{(u,a)}) = \Theta_\psi(-\mathcal{L}_u \psi + a\psi)$ more explicitly.

Lemma 4.2.10 (Evaluation of Θ_ψ). *For any $(u, a) \in \mathfrak{DC}$, we have*

$$\begin{aligned} \Theta_\psi(-\mathcal{L}_u \psi + a\psi) &= \Theta_\psi(-\mathcal{L}_u \psi) + \Theta_\psi(a\psi) \\ &= \frac{1}{2\pi} \left(-\langle \Lambda(\psi), \iota_u \mu \rangle + \int \text{Im } a \mu \right), \end{aligned}$$

where $\Lambda(\psi)$, defined in (4.2.3), is the current associated to the circle differential of ψ , and μ is the volume form on M .

Proof. By direct computation we get

$$\Theta_\psi(a\psi) = \frac{1}{2\pi} \int \operatorname{Im} \frac{a\psi\bar{\psi}}{|\psi|^2} \mu = \frac{1}{2\pi} \int \operatorname{Im} a \mu.$$

We then compute $\Theta_\psi(-\mathcal{L}_u\psi)$. Using the local expression $\psi = re^{i\phi}$ with $r = |\psi|$ and a function ϕ around each $x \in M \setminus \psi^{-1}(0)$, we have locally

$$d\psi = e^{i\phi} dr + re^{i\phi} i d\phi$$

and

$$\operatorname{Im} \frac{\iota_u d\psi \cdot \bar{\psi}}{|\psi|^2} = \operatorname{Im} \frac{\iota_u dr}{r} + \operatorname{Im}(i\iota_u d\phi) = 0 + \iota_u \lambda_\psi.$$

Hence, we obtain using the product rule for the interior product that,

$$\begin{aligned} \Theta_\psi(-\mathcal{L}_u\psi) &= -\frac{1}{2\pi} \int \operatorname{Im} \frac{\iota_u d\psi \cdot \bar{\psi}}{|\psi|^2} \mu = -\frac{1}{2\pi} \int \iota_u \lambda_\psi \mu \\ &= -\frac{1}{2\pi} \int \lambda_\psi \wedge \iota_u \mu = -\frac{1}{2\pi} \langle \Lambda(\psi), \iota_u \mu \rangle. \end{aligned}$$

□

We next need an auxiliary result describing how Θ varies under the DC-action.

Lemma 4.2.11. *Let $\psi \in \mathcal{U}$, $(f, e^\varphi) \in \text{DC}$, and $\hat{\psi} := (f, e^\varphi) \triangleright \psi = \psi \circ f^{-1} \cdot e^\varphi$. For the fundamental vector field of $Y^{(u,a)} \in \Gamma(T\mathcal{U})$ of $(u, a) \in \mathfrak{DC}$, we have*

$$\Theta_{\hat{\psi}}(Y^{(u,a)}) = \frac{1}{2\pi} \left(- \int \iota_u (f^{-1*} \lambda_\psi + \operatorname{Im} d\varphi) \mu + \int a \mu \right).$$

Additionally, if u is divergence-free, it simplifies to

$$\Theta_{\hat{\psi}}(Y^{(u,a)}) = \frac{1}{2\pi} \left(- \int \iota_u f^{-1*} \lambda_\psi + \int a \mu \right).$$

Proof. Clearly we have $\Theta_{\hat{\psi}}(a\hat{\psi}) = \int a \mu$. We now compute $\Theta_{\hat{\psi}}(-\mathcal{L}_u\hat{\psi})$. From

$$d\hat{\psi} = d(\psi \circ f^{-1} \cdot e^\varphi) = e^\varphi \cdot d(\psi \circ f^{-1}) + \psi \circ f^{-1} \cdot e^\varphi d\varphi = e^\varphi d(\psi \circ f^{-1}) + \hat{\psi} d\varphi$$

it follows that,

$$\begin{aligned} 2\pi \Theta_{\hat{\psi}}(-\mathcal{L}_u\hat{\psi}) &= - \int \operatorname{Im} \frac{\iota_u d\hat{\psi} \cdot \bar{\hat{\psi}}}{|\hat{\psi}|^2} \mu = - \int \operatorname{Im} \frac{e^\varphi \iota_u d(\psi \circ f^{-1}) + \hat{\psi} \hat{\psi} \iota_u d\varphi}{|\hat{\psi}|^2} \mu \\ &= - \int \operatorname{Im} \frac{\hat{\psi} \hat{\psi} \iota_u \lambda_{\psi \circ f} + |\hat{\psi}|^2 \iota_u d\varphi}{|\hat{\psi}|^2} \mu = - \int (\iota_u \lambda_{\psi \circ f} + \iota_u \operatorname{Im} d\varphi) \mu. \end{aligned}$$

We have by definition that $\lambda_{\psi \circ f} = (\psi \circ f)^* \tau^* \sigma = f^* \psi^* \tau^* \sigma = f^* \lambda_\psi$, which gives the stated expression.

Finally, if u is divergence-free, we get by integral by parts that

$$\int \iota_u d \operatorname{Im} \varphi \mu = - \int \operatorname{Im} \varphi \mathcal{L}_u \mu = 0.$$

□

Using these results, we can now explicitly evaluate each term in (4.2.2) and compute $d\Theta$.

Lemma 4.2.12. *We have*

$$d\Theta_\psi(Y^{(u,a)}, Y^{(v,b)}) = \frac{1}{2\pi} \langle \partial\Lambda(\psi), \iota_v \iota_u \mu \rangle. \quad (4.2.5)$$

where $\partial: \mathcal{D}_{m-1}(M) \rightarrow \mathcal{D}_{m-2}(M)$ is the boundary operator for currents.

Proof. First, we have

$$\Theta_\psi([Y^{(u,a)}, Y^{(v,b)}]) = \frac{1}{2\pi} \int (-\iota_{[u,v]}\lambda_\psi + \mathcal{L}_u \text{Im } b - \mathcal{L}_v \text{Im } a) \mu$$

using the expressions of $[Y^{(u,a)}, Y^{(v,b)}]$ (Lemma 4.2.8) and of Θ_ψ (Lemma 4.2.10).

We then compute $\mathcal{L}_{Y^{(u,a)}} \iota_{Y^{(v,b)}} \Theta_\psi$. Applying Lemma 4.2.11 to the time- t flow map of $Y^{(u,a)}$ given by $\Phi_{Y^{(u,a)}}^t(\psi) = \text{Fl}_u^{-t*} \psi \cdot e^{ta}$, we have

$$\begin{aligned} \mathcal{L}_{Y^{(u,a)}} \Theta(Y^{(v,b)}(\psi)) &= \left. \frac{d}{dt} \right|_{t=0} \Theta(Y^{(v,b)}(\Phi_{Y^{(u,a)}}^t(\psi))) \\ &= \frac{1}{2\pi} \left(\left. \frac{d}{dt} \right|_{t=0} \int -(\iota_v \text{Fl}_u^{-t*} \lambda_\psi + \iota_v d \text{Im}(ta)) \mu + \int \text{Im } b \mu \right) \\ &= \frac{1}{2\pi} \int (\iota_v \mathcal{L}_u \lambda_\psi - \mathcal{L}_v \text{Im } a) \mu \end{aligned}$$

and $\mathcal{L}_{Y^{(v,b)}} \Theta(Y^{(u,a)}(\psi)) = \frac{1}{2\pi} \int (\iota_u \mathcal{L}_v \lambda_\psi - \mathcal{L}_u \text{Im } b) \mu$ in the same way.

We can now compute

$$\begin{aligned} \iota_{Y^{(v,b)}} \iota_{Y^{(u,a)}} d\Theta_\psi &= \mathcal{L}_{Y^{(u,a)}} \iota_{Y^{(v,b)}} \Theta_\psi - \mathcal{L}_{Y^{(v,b)}} \iota_{Y^{(u,a)}} \Theta_\psi - \iota_{[Y^{(u,a)}, Y^{(v,b)}]} \Theta_\psi \\ &= \frac{1}{2\pi} \int \left(\iota_v \mathcal{L}_u \lambda_\psi - \mathcal{L}_v \text{Im } a - \iota_u \mathcal{L}_v \lambda_\psi + \mathcal{L}_u \text{Im } b \right. \\ &\quad \left. + \iota_{[u,v]} \lambda_\psi + \mathcal{L}_v \text{Im } a - \mathcal{L}_u \text{Im } b \right) \mu \\ &= \frac{1}{2\pi} \int (\iota_v \mathcal{L}_u \lambda_\psi - \iota_u \mathcal{L}_v \lambda_\psi + \iota_{[u,v]} \lambda_\psi) \mu. \end{aligned}$$

We then simplify the integrand:

$$\begin{aligned} \iota_v \mathcal{L}_u \lambda_\psi - \iota_u \mathcal{L}_v \lambda_\psi - \iota_{[v,u]} \lambda_\psi &= (\iota_v \mathcal{L}_u - \iota_u \mathcal{L}_v - \mathcal{L}_v \iota_u + \iota_u \mathcal{L}_v) \lambda_\psi \\ &= (\iota_v \mathcal{L}_u - \mathcal{L}_v \iota_u) \lambda_\psi \\ &= (\iota_v d \iota_u + \iota_v \iota_u d - d \iota_v \iota_u - \iota_v d \iota_u) \lambda_\psi \\ &= (\iota_v \iota_u d - d \iota_v \iota_u) \lambda_\psi \\ &= \iota_v \iota_u d \lambda_\psi - 0. \end{aligned}$$

Hence we obtain by the product rule that,

$$\begin{aligned} \iota_{Y^{(v,b)}} \iota_{Y^{(u,a)}} d\Theta_\psi &= \frac{1}{2\pi} \int \iota_v \iota_u d \lambda_\psi \mu = -\frac{1}{2\pi} \int d \lambda_\psi \wedge \iota_u \iota_v \mu = -\frac{1}{2\pi} \int \lambda_\psi \wedge d \iota_u \iota_v \mu \\ &= \frac{1}{2\pi} \langle \Lambda(\psi), d \iota_v \iota_u \mu \rangle = \frac{1}{2\pi} \langle \partial\Lambda(\psi), \iota_v \iota_u \mu \rangle \end{aligned}$$

and get the stated expression. \square

We are now one step away from proving Theorem 4.2.3. To complete the proof, we show the next lemma asserting that $\frac{1}{2\pi}\partial\Lambda(\psi)$ in (4.2.5) is $\delta_{\Pi\psi}$.

Lemma 4.2.13. *Let $\Lambda: \mathcal{U} \rightarrow \mathcal{D}_{m-1}(M)$ be the circle differential map as defined in Definition 4.2.4 and $\partial: \mathcal{D}_{m-1}(M) \rightarrow \mathcal{D}_{m-2}(M)$ be the boundary operator for currents. Then we have $\frac{1}{2\pi}\partial \circ \Lambda = \Pi: \mathcal{U} \rightarrow \mathcal{D}_{\mathcal{O}}$. That is, for any $\psi \in \Pi^{-1}\gamma$ for some $\gamma \in \mathcal{O}$, we have $\frac{1}{2\pi}\partial(\Lambda(\psi)) = \delta_{\gamma}$.*

Proof. We have for $\alpha \in \Omega^{m-2}(M)$ that

$$\langle \partial\Lambda(\psi), \alpha \rangle = \langle \Lambda(\psi), d\alpha \rangle = \int_M \lambda_{\psi} \wedge d\alpha.$$

As in the proof of Lemma 4.2.5, we have by the coarea formula (Proposition 4.2.6) that,

$$\int_M \lambda_{\psi} \wedge d\alpha = \int_{\mathbb{S}^1} \langle \delta_{(\phi^{-1}(\cdot) \cap M^{\text{reg}} \phi)}, d\alpha \rangle \sigma = \int_{\mathbb{S}^1} \langle \delta_{\partial(\phi^{-1}(\cdot) \cap M^{\text{reg}} \phi)}, \alpha \rangle \sigma$$

where $\phi: M \setminus \psi^{-1}(0) \rightarrow \mathbb{S}^1$ is the phase map and $M^{\text{reg}} \phi$ is the set of regular points of ϕ as in Lemma 4.2.5.

Due to Sard's theorem, almost σ -every point of \mathbb{S}^1 is a regular value of ϕ . Hence we have,

$$\begin{aligned} \int_{\mathbb{S}^1} \langle \delta_{\partial(\phi^{-1}(\cdot) \cap M^{\text{reg}} \phi)}, \alpha \rangle \sigma &= \int_{\mathbb{S}_{\text{reg } \phi}^1} \langle \delta_{\partial(\phi^{-1}(\cdot))}, \alpha \rangle \sigma \\ &= \int_{\mathbb{S}_{\text{reg } \phi}^1} \langle \delta_{\gamma}, \alpha \rangle \sigma = \langle \delta_{\gamma}, \alpha \rangle \cdot \int_{\mathbb{S}_{\text{reg } \phi}^1} \sigma = 2\pi \langle \delta_{\gamma}, \alpha \rangle \end{aligned}$$

where $\mathbb{S}_{\text{reg } \phi}^1$ is the set of regular values of ϕ . □

We now complete the proof of the theorem.

Proof of Theorem 4.2.3. By Lemma 4.2.12 and Lemma 4.2.13 we have for $\psi \in \Pi^{-1}\gamma$ and $\dot{\psi} = -\mathcal{L}_u\psi + a\psi$, $\dot{\psi} = -\mathcal{L}_v\psi + b\psi$ that,

$$d\Theta_{\psi}(\dot{\psi}, \dot{\psi}) = \frac{1}{2\pi} \langle \partial\Lambda(\psi), \iota_v \iota_u \mu \rangle = \langle \delta_{\gamma}, \iota_v \iota_u \mu \rangle = \omega_{\delta_{\gamma}}(-\mathcal{L}_u \delta_{\gamma}, -\mathcal{L}_v \delta_{\gamma}) = \Pi^* \omega_{\psi}(\dot{\psi}, \dot{\psi}).$$

□

Lemma 4.2.13 also gives a physical interpretation of the Liouville form Θ in the following sense.

Corollary 4.2.14 (Liouville form as flux). *We have for $\psi \in \mathcal{U}$, $u \in \text{diff}(M)$,*

$$\Theta_{\psi}(-\mathcal{L}_u\psi) = -\frac{1}{2\pi} \text{Flux}_{m-1}(\Lambda(\psi), u) = -\frac{1}{2\pi} \int_{\mathbb{S}_{\text{reg } \phi}^1} \text{Flux}_{m-1}(\delta_{\phi^{-1}(\cdot)}, u) \sigma.$$

Moreover, if u is exact divergence-free, we have

$$\Theta_{\psi}(-\mathcal{L}_u\psi) = -\text{Flux}_{m-1}(\delta_{\phi^{-1}(s)}, u) = -\text{Flux}_{m-2}^{\text{ex}}(\delta_{\Pi\psi}, u), \quad \forall s \in \mathbb{S}_{\text{reg } \phi}^1.$$

Here, Flux_{m-1} and $\text{Flux}_{m-2}^{\text{ex}}$ are flux for $m-1$ currents and $m-2$ exact currents respectively (Definition 2.1.2).

Proof. The first equality is immediate from Lemma 4.2.13. The second equality follows from the exactness of the $m - 1$ form $\iota_u \mu$. \square

Remark 4.2.15 (Physical interpretation of Θ_ψ). By Sard's theorem, σ -almost every point of \mathbb{S}^1 is a regular value of the phase map ϕ , and each such point s defines a hypersurface $\phi^{-1}(s)$. The first expression in Corollary 4.2.14 indicates that $\Theta(-\mathcal{L}_u \psi)$ measures the average flux of the vector field u through these phase level hypersurfaces. The second expression in the corollary shows that if the vector field u is exact divergence-free, the flux through all these level surfaces are actually the same.

Together with Lemma 4.2.10, we now observe that $\Theta(\dot{\psi})$ captures the infinitesimal phase shift of ψ over the space and the flux of vector fields through phase hypersurfaces. In other words, it is the average of the swept volume by the \mathbb{S}^1 -family of hypersurfaces $\{\phi^{-1}(s)\}_s$.

In this section, we have shown that Θ is a Liouville form for ω in the prequantum sense, that is, $d\Theta = \Pi^* \omega$. We conclude this section by noting that Θ does not directly induce a Liouville form for the symplectic structure ω on \mathcal{O} in the classical sense. This is because we cannot factor Θ onto \mathcal{O} , as the kernel of $d\Pi$ for $\Pi: \mathcal{U} \rightarrow \mathcal{O}$ is not contained within the kernel of Θ . For example, for $\dot{\psi} = e^{ic_1} \psi$, $\dot{\psi} = e^{ic_2} \psi \in T_\psi \mathcal{U}$ with some real constants $c_1 \neq c_2$, we have $d\Pi(\dot{\psi}) = d\Pi(\dot{\psi}) = 0 \in T_{\Pi(\psi)} \mathcal{O}$, but the values $\Theta(\dot{\psi}) = \frac{\text{Vol}(M)c_1}{2\pi}$ and $\Theta(\dot{\psi}) = \frac{\text{Vol}(M)c_2}{2\pi}$ are different.

However, it is still possible to take a quotient space of \mathcal{U} where the Liouville form can descend onto. This construction actually leads to a prequantum bundle structure, as we will explain in the next section.

4.3 Prequantum structure

We have shown that each DC-orbit \mathcal{U} in the fiber bundle $\mathcal{F}_\mathcal{O} = \Pi^{-1}\mathcal{O}$ admits a Liouville form in the prequantum sense for the MW structure on \mathcal{O} . The full fiber $\mathcal{F}_\mathcal{O}$ inherits these symplectic and Liouville structures as $\mathcal{F}_\mathcal{O}$ is a $\text{Diff}_0(M) \times C^\infty(M, \mathbb{C}^\times)$ -orbit, foliated by DC-orbits, as discussed in Section 3.3.

Building on this, we now construct a prequantum bundle as a quotient bundle of $\mathcal{F}_\mathcal{O}$, equipped with a connection form whose curvature form recovers the MW symplectic form on \mathcal{O} . This framework enables us to define a unique horizontal lift in $\mathcal{F}_\mathcal{O}$ over a path in \mathcal{O} .

Definition 4.3.1 (Prequantum G -bundle). A *prequantum G -bundle* over a symplectic manifold (B, β) is a principal G -bundle $\pi: E \rightarrow B$ equipped with a connection form $\alpha \in \Omega^1(E; \mathfrak{g})$ satisfying

$$\pi^* \beta = d\alpha + \frac{1}{2}[\alpha \wedge \alpha].$$

Remark 4.3.2. Definition 4.3.1 relaxes the standard definition of a prequantum bundle, which often requires the structure group to be a circle or a (complex) line. In fact, the prequantum bundle we construct in this section, the structure group G is not a circle but copies of a circle unless the ambient space M has trivial first cohomology. In any either, G is one-dimensional and hence we have $\pi^* \beta = d\alpha$ without the term $\frac{1}{2}[\alpha \wedge \alpha]$, as in the case of a circle bundle.

For a prequantum G -bundle, the vertical and horizontal distributions are defined in the standard way for principal bundles:

Definition 4.3.3 (Vertical and horizontal distributions, and horizontal lift). Let $\pi: (E, \alpha) \rightarrow (B, \beta)$ be a prequantum G -bundle. At each point $x \in E$, the tangent space splits as $T_x E = V_x E \oplus H_x E$, where the vertical distribution is defined by $V_x E = \ker d\pi|_x$, and the horizontal distribution is given by $H_x E = \ker \alpha|_x$. A *horizontal lift* of a path $\{\gamma_t\}_t \subset B$ is a path $\{\ell_t\}_t \subset E$ satisfying $\pi \circ \ell_t = \gamma_t$ and $\partial_t \ell_t \in H_{\ell_t} E$ for all t . Such a lift is uniquely determined by the choice of the initial point ℓ_0 .

Theorem 4.2.3 states that the fibration $\Pi: (\mathcal{F}_\mathcal{O}, \Theta) \rightarrow (\mathcal{O}, \omega)$ is a formal prequantization *i.e.*, $\Pi^* \omega = d\Theta$. However, it is not a prequantum G -bundle since Θ is not a connection form. Consequently, it does not define a unique horizontal lift: for a path $\{\gamma_t\}_t \subset \mathcal{O}$, there exist infinitely many paths $\{\ell_t\}_t \subset \mathcal{F}_\mathcal{O}$ with fixed initial point ℓ_0 satisfying $d\Pi(\partial_t \ell_t) = \partial_t \gamma_t$ and $\partial_t \ell_t \in \ker \Theta_{\ell_t}$ for all t . To define a bundle on which Θ becomes a genuine connection 1-form, we take the quotient of the tangent space at each $\psi \in \mathcal{F}_\mathcal{O}$ by the intersection $\ker \Theta \cap \ker d\Pi$. This quotient process is characterized by the following equivalence relation.

Definition 4.3.4 (Volume class and volume bundle). For $\psi_0, \psi_1 \in \mathcal{F}_\mathcal{O}$, we say $\psi_0 \sim_{\mathcal{P}} \psi_1$ if there is a path $\{\psi_t\}_{t \in [0,1]}$ in a DC-orbit joining ψ_0 and ψ_1 such that $\partial_t \psi_t \in \ker d\Pi_{\psi_t} \cap \ker \Theta_{\psi_t}$ for any t . We call each equivalence class $[\psi]_{\mathcal{P}}$ a *volume class* and the resulting quotient space $\mathcal{P} := \mathcal{F}_\mathcal{O} / \sim_{\mathcal{P}}$ the *volume bundle*. The tangent space at each $[\psi]_{\mathcal{P}}$ is $T_{[\psi]_{\mathcal{P}}} \mathcal{P} = T_{\psi} \mathcal{F}_\mathcal{O} / \sim_{\mathcal{P}}$ where $\dot{\psi} \sim_{\mathcal{P}} \dot{\psi}'$ if $\dot{\psi} - \dot{\psi}' \in \ker d\Pi|_{\psi} \cap \ker \Theta|_{\psi}$.

We note that the projection $\Pi: \mathcal{F}_\mathcal{O} \rightarrow \mathcal{O}$ decomposes into two projections: $\pi_{\mathcal{P}}: \mathcal{F}_\mathcal{O} \rightarrow \mathcal{P}$ and $\Pi_{\mathcal{P}}: \mathcal{P} \rightarrow \mathcal{O}$. By construction, the Liouville form Θ descends to a 1-form $\Theta_{\mathcal{P}} \in \Omega^1(\mathcal{P})$ by $\pi_{\mathcal{P}}^* \Theta_{\mathcal{P}} = \Theta$. Since $d\Theta = \Pi^* \omega$, we have $d\Theta_{\mathcal{P}} = \Pi_{\mathcal{P}}^* \omega$, and therefore the fibration $\Pi_{\mathcal{P}}: (\mathcal{P}, \Theta_{\mathcal{P}}) \rightarrow (\mathcal{O}, \omega)$ is a formal prequantization.

In addition, the volume bundle \mathcal{P} forms a principal G -bundle over \mathcal{O} , where the structure group is $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$, acting on \mathcal{P} as follows. We first define a circle action by constant phase shift:

$$e^{ic} \triangleright [\psi]_{\mathcal{P}} := [\psi e^{ic}]_{\mathcal{P}}, \quad \text{for } \psi \in \mathcal{F}, \quad e^{ic} \in \mathbb{S}^1. \quad (4.3.1)$$

As shown in Proposition 3.3.3, each fiber $\Pi^{-1}\gamma$ may contain multiple connected components indexed by the discrete group $H_{dR}^1(M, \mathbb{Z})$. The \mathbb{S}^1 -action (4.3.1) is free and transverse within each connected component.

We may further specify a group action of $H_{dR}^1(M, \mathbb{Z})$ on \mathcal{P} as a deck transformation between the connected components of each fiber. That is, it is a fiber-preserving action that maps points from one connected component of a fiber to another component. We fix a lattice basis $[\eta_1], \dots, [\eta_n]$ for $H_{dR}^1(M, \mathbb{Z})$, where n is the first Betti number of M . Recall that there is an isomorphism

$$\begin{aligned} \mathfrak{i}: C^\infty(M, \mathbb{C}^\times) / \text{Exp}(C^\infty(M, \mathbb{C}^\times)) &\rightarrow H_{dR}^1(M, \mathbb{Z}) \\ [\kappa] &\mapsto \left[\text{Im} \frac{1}{2\pi} \frac{d\kappa}{\kappa} \right]. \end{aligned}$$

Choose arbitrary representatives $\tau_i \in C^\infty(M, \mathbb{C}^\times)$ of $\mathfrak{i}^{-1}[\eta_i]$ for $i = 1, \dots, n$. Using these functions τ_i , we define the action of $H_{dR}^1(M, \mathbb{Z})$ by

$$[\eta] \triangleright [\psi]_{\mathcal{P}} := [\tau_1^{c_1} \cdots \tau_n^{c_n} \cdot \psi]_{\mathcal{P}}, \quad \psi \in \mathcal{F}, \quad [\eta] = [c_1 \eta_1 + \dots + c_n \eta_n] \in H_{dR}^1(M, \mathbb{Z}).$$

These \mathbb{S}^1 - and $H_{dR}^1(M, \mathbb{Z})$ -actions together define a $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$ -action

$$(e^{ic}, [\eta]) \triangleright [\psi]_{\mathcal{P}} = [e^{ic} \tau_1^{c_1} \cdots \tau_n^{c_n} \psi]_{\mathcal{P}}, \quad e^{ic} \in \mathbb{S}^1, [\eta] \in H_{dR}^1(M, \mathbb{Z}), \quad (4.3.4)$$

which is free, transverse, and fiber preserving. We thus obtain the following result.

Proposition 4.3.5. *The fibration $\pi_{\mathcal{P}}: \mathcal{F} \rightarrow \mathcal{O}$ equipped with the above action of $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$ is a principal G -bundle.*

We now define scaled versions of the MW form $\omega \in \Omega^2(\mathcal{O})$ and the Liouville form $\Theta_{\mathcal{P}} \in \Omega^1(\mathcal{P})$ and show that the latter defines a connection form on \mathcal{P} . Let us define the normalized MW form and Liouville form $\omega^{\mathcal{N}} := \mathcal{N}\omega$ and $\Theta_{\mathcal{P}}^{\mathcal{N}} := \mathcal{N}\Theta_{\mathcal{P}}$ with the normalization factor $\mathcal{N} = 2\pi/\text{Vol}(M)$.

Proposition 4.3.6. *On the principal G -bundle $\Pi_{\mathcal{P}}: \mathcal{P} \rightarrow \mathcal{O}$ defined above, the 1-form $\Theta_{\mathcal{P}}^{\mathcal{N}}$ is a connection form. That is, $\Theta_{\mathcal{P}}^{\mathcal{N}}$ satisfies the following two properties:*

1. *Equivariance under the G -action $\Phi_g: \mathcal{P} \rightarrow \mathcal{P}$; i.e., $\Phi_g^* \Theta_{\mathcal{P}} = \Theta_{\mathcal{P}}$ for every $g \in G$.*
2. *Vertical reproducibility; i.e., there exists a nonzero constant $C \neq 0$ such that $\Theta_{\mathcal{P}}(\hat{\xi}) = C\xi$ for any $\xi \in \mathfrak{g}$ and its fundamental vector field $\hat{\xi} \in \Gamma(T\mathcal{P})$.*

Proof. We first verify equivariance. Let $g = (e^{ic}, [\sum_i c_i \eta_i]) \in G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$, and denote by Φ_g the g -action (4.3.4). By direct computation, we have

$$\begin{aligned} \Theta_{\mathcal{P}}^{\mathcal{N}}|_{\Phi_g([\psi]_{\mathcal{P}})}([\Phi_{g*}\dot{\psi}]_{\mathcal{P}}) &= \frac{1}{\text{Vol}(M)} \int \text{Im} \frac{\dot{\psi} e^{ic} \tau_1^{c_1} \cdots \tau_n^{c_n} \cdot \overline{\psi e^{ic} \tau_1^{c_1} \cdots \tau_n^{c_n}}}{|\psi e^{ic} \tau_1^{c_1} \cdots \tau_n^{c_n}|^2} \mu \\ &= \frac{1}{\text{Vol}(M)} \int \text{Im} \frac{\dot{\psi} \cdot \bar{\psi}}{|\psi|^2} \mu = \Theta_{\mathcal{P}}|_{[\psi]_{\mathcal{P}}}([\dot{\psi}]_{\mathcal{P}}), \end{aligned}$$

where $\psi \in \mathcal{F}_{\mathcal{O}}$ and $\dot{\psi} \in T_{\psi}\mathcal{F}_{\mathcal{O}}$ are any representatives of $[\psi]_{\mathcal{P}}$ and $[\dot{\psi}]_{\mathcal{P}}$ respectively.

Next, we verify vertical reproducibility. Let $\xi \in \mathfrak{g} = T_{(1,0]}G = T_1\mathbb{S}^1 = \mathbb{R}$, and let $\hat{\xi} \in \Gamma(T\mathcal{P})$ be the corresponding fundamental vector field, defined by

$$\hat{\xi}([\psi]_{\mathcal{P}}) = \left[\frac{d}{dt} \Big|_{t=0} e^{i\xi t} \psi \right]_{\mathcal{P}} = [i\xi \psi]_{\mathcal{P}}.$$

Then, we have

$$\Theta_{\mathcal{P}}^{\mathcal{N}}|_{[\psi]_{\mathcal{P}}}(\hat{\xi}) = \frac{1}{\text{Vol}(M)} \int \text{Im} \frac{i\xi \psi \bar{\psi}}{|\psi|^2} \mu = \xi$$

□

Since $\Theta_{\mathcal{P}}^{\mathcal{N}}$ and $\omega^{\mathcal{N}}$ are constant multiples of $\Theta_{\mathcal{P}}$ and ω , all the results established so far applies except Proposition 4.3.6. Therefore, combining the results that $\Pi_{\mathcal{P}}: (\mathcal{P}, \Theta_{\mathcal{P}}) \rightarrow (\mathcal{O}, \omega)$ is a formal prequantization (Theorem 4.2.3) and a principal G -bundle with $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$ (Proposition 4.3.5), and that $\Theta_{\mathcal{P}}^{\mathcal{N}}$ is a connection form (Proposition 4.3.6), we obtain the main theorem.

Theorem 4.3.7. *The fibration $\Pi_{\mathcal{P}}: (\mathcal{P}, \Theta_{\mathcal{P}}^N) \rightarrow (\mathcal{O}, \omega^N)$ is a prequantum G -bundle with $G = \mathbb{S}^1 \times H_{dR}^1(M, \mathbb{Z})$.*

Remark 4.3.8 (Role of normalization). We normalized the MW form ω and the Liouville form $\Theta_{\mathcal{P}}$ for two reasons, while noting that our results are essentially unaffected by the presence or absence of normalization.

The first reason is to align our setting with the standard definition of vertical reproducibility of a connection form, $\Theta(\hat{\xi}) = \xi$ for $\xi \in \mathfrak{g}$. Without normalization we have $\Theta(\hat{\xi}) = \frac{1}{N}\xi$ instead, which is harmless: one may relax the definition to allow a constant factor, and all of our results remain valid.

Second, normalization makes ω physically natural in terms of physical units. Note that the value $\omega(-\mathcal{L}_u\delta_\gamma, -\mathcal{L}_v\delta_\gamma) = \int_\gamma \iota_v \iota_u \mu$ carries the dimension of $\text{length}^m / \text{time}^2$ as u, v have length/time, so $\int_\Sigma \omega$ for a topological disc $\Sigma \subset \mathcal{O}$ has unit length^m . Thus, multiplying ω by the factor $2\pi / \text{Vol}(M)$ makes its unit radius, canceling length^m . By relaxing the definition of vertical reproducibility as above, one could even use $1/\text{Vol}(M)$ as the normalization factor to make ω dimension-less.

Additionally the *prequantizability* of a symplectic form on \mathcal{O} is often expressed in the literature [HV03]: “ ω is prequantizable if $\text{Vol}(M) = 2\pi$ or if it is of integral class times 2π ”. However, the specific value of $\text{Vol}(M)$ depends on the unit of measurement e.g., cubic meter m^3 and cubic yard yd^3 in 3D. We believe that the choice of unit for $\text{Vol}(M)$ should not affect the prequantizability of the mechanics on the space. This motivates our normalization to remove this artificial dependence.

Using either the normalized or the unnormalized connection form $\Theta_{\mathcal{P}}$, each tangent space $T_{[\psi]_{\mathcal{P}}}\mathcal{P}$ splits into the vertical distribution $V_{[\psi]_{\mathcal{P}}}\mathcal{P} = \ker d\Pi_{\mathcal{P}}|_{[\psi]_{\mathcal{P}}}$ and the horizontal distribution $H_{[\psi]_{\mathcal{P}}}\mathcal{P} = \ker \Theta_{\mathcal{P}}|_{[\psi]_{\mathcal{P}}}$. Since $\ker d\Pi_{\mathcal{P}}|_{[\psi]_{\mathcal{P}}} = \ker d\Pi / \ker \Theta|_{\psi}$ for any representative $\psi \in \pi_{\mathcal{P}}^{-1}([\psi]_{\mathcal{P}})$, and $\Theta \in \Omega^1(\mathcal{F}_{\mathcal{O}})$ decomposes the tangent space $T_\psi\mathcal{F}_{\mathcal{O}}$ into codimension-1 hyperplanes as the level sets of Θ_ψ , the vertical distribution $V_{[\psi]_{\mathcal{P}}}\mathcal{P}$ has precisely one dimension.

Hence we can now define a unique horizontal lift on \mathcal{P} (Definition 4.3.3) over a path in \mathcal{O} . In light of Remark 4.2.15, this horizontal lift can be interpreted as the evolution of the implicit representation $[\psi]_{\mathcal{P}}$ such that the average swept volume of phase level hypersurfaces remains zero at all times. That is, for any representative ψ_t of a horizontal lift $[\psi_t]_{\mathcal{P}}$, we have for the family of phase hypersurfaces $\{\sigma_t^s\}_{s \in \mathbb{S}^1}$ defined by $\sigma_t^s = \phi_t^{-1}(s)$,

$$\frac{1}{2\pi} \int_{\mathbb{S}^1} \int_{\sigma_t^s} \iota_{\partial_t \sigma_t^s} \mu \, ds = 0$$

at each t where $\partial_t \sigma_t^s$ is the velocity defined on each hypersurface σ_t^s .

This also reveals a geometric interpretation of the MW form as the curvature form of the Liouville form $\Theta_{\mathcal{P}}$, measuring the holonomy induced by parallel transport on \mathcal{P} over a closed path in \mathcal{O} :

Corollary 4.3.9 (Average swept volume). *Consider a closed path $\partial\Sigma$ in \mathcal{O} that bounds a 2-dimensional disk Σ , representing a cyclic motion of a codimension-2 submanifold $\gamma_t \subset M$ for $0 \leq t \leq 1$, with $\gamma_0 = \gamma_1$ and let $[\psi_t]_{\mathcal{P}}$ be a horizontal lift over γ_t with respect to $\Theta_{\mathcal{P}}$. Suppose $[\psi_0]_{\mathcal{P}}$ and $[\psi_1]_{\mathcal{P}}$ are on the same connected component of the fiber $\Pi_{\mathcal{P}}^{-1}\gamma$. Then, for representatives ψ_0 and ψ_1 of $[\psi_0]_{\mathcal{P}}$ and $[\psi_1]_{\mathcal{P}}$, the volume enclosed between phase hypersurface σ_0^s and σ_1^s , averaged over $s \in \mathbb{S}^1$, equals to $\iint_{\Sigma} \omega$, where ω is the MW form on \mathcal{O} .*

Proof. The integral of the curvature is given by the monodromy of the parallel transport in \mathcal{P} over a closed path bounding Σ in \mathcal{O} , which is measured by the connection form $\Theta_{\mathcal{P}}^{\mathcal{N}}$ (and unnormalized form $\Theta_{\mathcal{P}}$ corresponding to ω). That is,

$$\int_{\Sigma} \omega^{\mathcal{N}} = \int_0^1 \Theta_{\mathcal{P}}^{\mathcal{N}} \left(\frac{d}{dt} (e^{i \operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}}) t} \triangleright [\psi_1]_{\mathcal{P}}) \right) dt = \operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}}) \pmod{2\pi},$$

where $\operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}}) \in \mathbb{R}/2\pi\mathbb{Z}$ is defined by $e^{i \operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}})} \triangleright [\psi_1]_{\mathcal{P}} = [\psi_0]_{\mathcal{P}}$ when $[\psi_0]_{\mathcal{P}}$ and $[\psi_1]_{\mathcal{P}}$ are in the same connected component of a fiber. Explicitly,

$$\operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}}) = \arg \int \bar{\psi}_1 \psi_0 \mu$$

with any representatives ψ_0, ψ_1 .

For unnormalized forms ω and $\Theta_{\mathcal{P}}$, the result is multiplied by $1/\mathcal{N} = \operatorname{Vol}(M)/2\pi$,

$$\int_{\Sigma} \omega = \int_0^1 \Theta_{\mathcal{P}} \left(\frac{d}{dt} (e^{i \operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}}) t} \triangleright [\psi_1]_{\mathcal{P}}) \right) dt = \frac{\operatorname{Vol}(M)}{2\pi} \operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}}) \pmod{\operatorname{Vol}(M)},$$

Let us now consider any path $\{\psi_{t'}\}_{t' \in [0,1]} \subset \Pi^{-1}\gamma_0$ joining representatives ψ_1 and ψ_0 within the fiber. As observed in Remark 4.2.15, the Liouville form Θ measures the infinitesimal swept volume of the phase level hypersurfaces σ_t^s averaged in \mathbb{S}^1 . Therefore, the integral of Θ along $\psi_{t'}$ is the average signed volume enclosed by pairs $\{(\sigma_1^s, \sigma_0^s)\}_s$, which is formally

$$\frac{\operatorname{Vol}(M)}{2\pi} \operatorname{Arg}([\psi_1]_{\mathcal{P}}, [\psi_0]_{\mathcal{P}}) = \int_0^1 \Theta(\dot{\psi}_{t'}) dt' = \int_{\mathbb{S}^1} \operatorname{Vol}(\partial^{-1}(\sigma_s^0 - \sigma_s^1)) ds \pmod{\operatorname{Vol}(M)}.$$

□

In the proof, we used a known fact that the integral of curvature form is the monodromy of parallel transport measured by the connection form. In Section 4.5, we gave a precise statement and a proof.

As illustrated in Figure 3.4, we may compress the phase field so that it is non-constant only within narrow bands. By considering a limiting case of Corollary 4.3.9, where $\arg \psi_t := M \setminus \psi^{-1}(0) \rightarrow \mathbb{T}^1$ becomes constant except a 2π jump at on a single hypersurface σ_t bounding γ_t , we obtain the following:

Corollary 4.3.10 (Swept volume by a hypersurface). *Let $\Sigma \subset \mathcal{O}$ and $\{\gamma_t\}_{t \in [0,1]} \subset M$ be as in Corollary 4.3.9. Suppose that each γ_t bounds a hypersurface, i.e., $\gamma_t = \partial \sigma_t$, and that the volume swept out by σ_t remains zero at each t , meaning $\int_{\sigma_t} \iota_{\partial_t \sigma_t} \mu = 0$. Then, the volume enclosed between σ_0 and σ_1 is given by $\iint_{\Sigma} \omega$.*

Note that the interpretation of the MW form in Corollary 4.3.10 reduces to the swept volume of a single surface, no longer explicitly involving the complex function.

Remark 4.3.11. Corollary 4.3.10 can also be shown directly in the framework for the explicit shape space (Section 2.3.2). Let us take an $m-1$ dimensional submanifold σ_0 of M bounding some $\gamma_0 \in \mathcal{O}$, and consider the orbit $\mathcal{S} := \operatorname{Diff}_0(M) \triangleright \sigma_0$, where the action is defined by $f \triangleright \sigma = f \circ \sigma$ for $f \in \operatorname{Diff}_0(M)$ and $\sigma \in \mathcal{S}$. Then $\pi: \mathcal{S} \rightarrow \mathcal{O}$ is a fiber bundle where the

fibration is given by the boundary operator, i.e., $\pi(\sigma) = \partial\sigma$, and the tangent space at each σ is $T_\sigma\mathcal{S} = \{u \circ \sigma \mid u \in \text{diff}(M)\}$.

Define a 1-form η on \mathcal{S} by $\eta_\sigma(u \circ \sigma) = \text{Flux}_{m-1}(\delta_\sigma, u)$ where δ_σ is a de Rham $m-1$ current of σ . Then η serves as a formal prequantization, i.e., $d\eta = \pi^*\omega$, which can be shown in a manner similar to the proof of Theorem 2.3.3. For a path $\{\gamma_t\} \subset \mathcal{O}$, there exist infinitely many lifts $\{\ell_t\} \subset \mathcal{S}$ such that $\partial_t \ell_t \in \ker \eta$ for all t , but the notion of no swept volume still makes sense, and we recover Corollary 4.3.10.

4.4 Additional structures

Our implicit representations of shapes and the resulting prequantum bundle structure reveal some additional structures.

4.4.1 Hamiltonian systems on the space of implicit representations

In Section 2.3.1, we reviewed Hamiltonian vector fields on each $\text{Diff}_0(M)$ -orbit \mathcal{O} in the explicit shape space with respect to the MW form ω . We now consider Hamiltonian vector fields in the implicit shape space $\mathcal{F}_\mathcal{O} = \Pi^{-1}\mathcal{O}$.

Define a closed 2-form \tilde{Z} on $\mathcal{F}_\mathcal{O}$ by $\tilde{Z} := \Pi^*\omega$. The closedness of \tilde{Z} follows from that of ω . Note that \tilde{Z} is degenerate with $\ker \tilde{Z} = \ker d\Pi$, which is infinite-dimensional. Hence, \tilde{Z} is not symplectic but merely a *presymplectic* form, in contrast to ω being symplectic.

Consequently, Hamiltonian vector fields are defined only up to $\ker \tilde{Z}$. Let $\mathcal{H}: \mathcal{O} \rightarrow \mathbb{R}$ be a function with $d\mathcal{H} \in \text{im } \flat^\omega$, and define $H: \mathcal{F}_\mathcal{O} \rightarrow \mathbb{R}$ by $H = \mathcal{H} \circ \Pi$. We say that any vector field $X_H \in \Gamma(T\mathcal{F}_\mathcal{O})$ satisfying $dH = \iota_{X_H} \tilde{Z}$ is a Hamiltonian vector field of H , despite the non-uniqueness.

Horizontal Hamiltonian vector fields

The prequantum structure we have built allows a canonical choice among these Hamiltonian vector fields. We can define a unique *horizontal Hamiltonian* vector field on the prequantum bundle \mathcal{P} over a Hamiltonian vector field on \mathcal{O} , using $\Theta_\mathcal{P}$ as the connection form.

Since $H = \mathcal{H} \circ \Pi$ is by construction constant along each fiber $\Pi^{-1}\gamma$ over $\gamma \in \mathcal{O}$, it descends onto \mathcal{P} as a function $H_\mathcal{P} = \mathcal{H} \circ \Pi_\mathcal{P}$ defined using $\Pi_\mathcal{P}: \mathcal{P} \rightarrow \mathcal{O}$. Then there is one degree of freedom within the vector fields $X_{H_\mathcal{P}} \in \Gamma(T\mathcal{P})$ satisfying $dH_\mathcal{P} = \iota_{X_{H_\mathcal{P}}} Z_\mathcal{P}$ where $Z_\mathcal{P} := \Pi_\mathcal{P}^*\omega$ is a presymplectic structure on \mathcal{P} . Among these, the horizontal Hamiltonian vector field is the unique one additionally satisfying the horizontality condition $X_{H_\mathcal{P}} \in H_{[\psi]_\mathcal{P}} \mathcal{P} = \ker \Theta_\mathcal{P}|_{[\psi]_\mathcal{P}}$ at each $[\psi]_\mathcal{P} \in \mathcal{P}$.

Momentum maps

Momentum maps can also be defined for presymplectic structures $\tilde{Z} = \Pi^*\omega \in \Omega^2(\mathcal{F}_\mathcal{O})$ in much the same way as for the MW symplectic structure ω explained in Section 2.3.3. Given an action of a Lie group G on $\mathcal{F}_\mathcal{O}$, we say that $J: \mathcal{F}_\mathcal{O} \rightarrow \mathfrak{g}^*$ is a momentum map if $d\langle J(\cdot), \xi \rangle = \iota_{\hat{\xi}} \tilde{Z}$ for any $\xi \in \mathfrak{g}$ and its corresponding fundamental vector field $\hat{\xi} \in \Gamma(T\mathcal{F}_\mathcal{O})$. When the Hamiltonian H is invariant under the G -action, $J(\psi_t) \in \mathfrak{g}^*$ is conserved along the Hamiltonian flow ψ_t .

Similar to the case of the Liouville form η for the MW symplectic form ω on the explicit shape space \mathcal{O} , the Liouville form Θ of the presymplectic form \tilde{Z} in the implicit shape space $\mathcal{F}_\mathcal{O}$ can

describe conserved quantities of Hamiltonian flows. When Θ is invariant under the G -action, the corresponding conserved quantity $J(\psi)$ satisfies $\langle J(\psi), \xi \rangle = -\Theta(\hat{\xi}(\psi))$, up to an additive constant.

As observed in Section 2.3.3, the Liouville form η for the MW form ω on $\text{UEmb}(S, \mathbb{R}^m)$ is not invariant under the action of the entire $\text{SDiff}^+(\mathbb{R}^m)$, but only under the volume-preserving linear transformations in $\text{SL}(\mathbb{R}^m)$. In contrast, the Liouville form Θ in the implicit shape space $\mathcal{F}_{\mathcal{O}}$ is invariant under the entire group $\text{SDiff}^+(M)$. We, however, acknowledge that their settings differ: η on \mathcal{O} is the Liouville form in the classical sense *i.e.*, $d\eta = \omega$ and the ambient manifold M has an exact volume form, whereas Θ on $\mathcal{F}_{\mathcal{O}}$ is defined in the prequantum sense *i.e.*, $d\Theta = \pi^*\omega$ and the ambient manifold is a closed manifold M .

To see the invariance of Θ under the $\text{SDiff}^+(M)$ -action, denote by $R_f: \psi \mapsto \psi \circ f^{-1}$ the action of $f \in \text{SDiff}^+(M)$ on $\psi \in \mathcal{F}$. A direct computation shows that $R_{f*}\dot{\psi} = \dot{\psi} \circ f^{-1}$ for $\dot{\psi} \in T_{\psi}\mathcal{F}_{\mathcal{O}}$. Then, using $f^*\mu = \mu$, we have

$$\begin{aligned} \Theta_{R_f(\psi)}(R_{f*}\dot{\psi}) &= \Theta_{\psi \circ f^{-1}}(\dot{\psi} \circ f^{-1}) \\ &= \frac{1}{2\pi} \int \text{Im} \frac{\dot{\psi} \circ f^{-1} \cdot \overline{\dot{\psi} \circ f^{-1}}}{|\dot{\psi} \circ f^{-1}|^2} \mu \\ &= \frac{1}{2\pi} \int \text{Im} \frac{\dot{\psi} \cdot \overline{\dot{\psi}}}{|\dot{\psi}|^2} f^*\mu = \Theta_{\psi}(\dot{\psi}). \end{aligned}$$

Note that the fundamental vector field $\hat{u} \in \Gamma(T\mathcal{F}_{\mathcal{O}})$ of $u \in \text{sdiff}(M)$ with respect to the $\text{SDiff}^+(M)$ -action on $\mathcal{F}_{\mathcal{O}}$ is given by $\hat{u}(\psi) = -\mathcal{L}_u\psi$. Therefore, the corresponding momentum map $J: \mathcal{F}_{\mathcal{O}} \rightarrow \text{sdiff}(M)^*$ is computed as

$$\langle J(\psi), u \rangle = -\Theta_{\psi}(-\mathcal{L}_u\psi) = \frac{1}{2\pi} \text{Flux}_{m-1}(\Lambda(\psi), u),$$

where $\Lambda(\psi) \in \mathcal{D}_{m-1}(M)$ is the circle differential current (Lemma 4.2.13) and Flux_{m-1} is the flux functional (Definition 2.1.2). In the case where u is exact divergence-free *i.e.*, $\iota_u\mu$ is exact, this further reduces to $\langle J(\psi), u \rangle = \text{Flux}_{m-2}^{\text{ex}}(\delta_{\Pi\psi}, u)$. This refers only to the explicit shape $\Pi\psi \in \mathcal{O}$ and coincides with the momentum map on \mathcal{O} given in Example 2.3.11.

4.4.2 Marsden–Weinstein structure in terms of implicit representations

Just as $\text{Emb}(S, M)$ provides parametrizations of unparametrized shapes $\text{UEmb}(S, M)$, we may regard $\mathcal{F}_{\mathcal{O}} = \Pi^{-1}\mathcal{O}$ as another parametrization space for an orbit \mathcal{O} in $\text{UEmb}(S, M)$. This perspective allows us to express the MW form ω in terms of implicit representations.

Let $\tilde{Z} = \Pi^*\omega$ be the presymplectic form on $\mathcal{F}_{\mathcal{O}}$ defined in the previous subsection, and let us now suppose that (M, μ) is equipped with a Riemannian metric inducing the volume form μ . Then \tilde{Z} admits an explicit expression, as we describe below.

At each zero $p \in \psi^{-1}(0)$, the tangent space decomposes as $T_pM = T_p\gamma(S) \oplus N_p\gamma(S)$ with respect to the given Riemannian metric g . We can choose a unit frame field (n_1, n_2) over the normal bundle $N\gamma(S)$ such that $\mu(\partial_{s_1}\gamma, \dots, \partial_{s_n}\gamma, n_1, n_2) > 0$ for any local coordinates (s_1, \dots, s_n) consistent with the orientation of S .

Let $d\psi_N: N_p\gamma(S) \rightarrow \mathbb{R}^2$ be the restriction of the differential $d\psi: T_pM \rightarrow \mathbb{C} \simeq \mathbb{R}^2$ to the normal space $N_p\gamma(S)$. This is a linear map between real 2-dimensional vector spaces, and

is invertible since $d\psi$ is surjective on $\psi^{-1}(0)$ by the definition of \mathcal{F} . We define the value $\det(d\psi_N) \in \mathbb{R}^\times$ by

$$\det(d\psi_N) \nu_N(\cdot, \cdot) = \nu_{\mathbb{C}}(d\psi_N \cdot, d\psi_N \cdot)$$

using the area form ν_N on $N\gamma(S)$ induced by g , and the standard area form $\nu_{\mathbb{C}}$ on \mathbb{C} , defined by $\nu_{\mathbb{C}}(h, k) = \operatorname{Re}(h\bar{i}k)$ for $h, k \in \mathbb{C}$. Note that $\det d\psi_N$ depends on the choice of Riemannian metric g , but is independent of the choice of the frame field (n_1, n_2) . The following proposition gives an explicit expression for \tilde{Z}_ψ .

Proposition 4.4.1. *Given a Riemannian metric g on M , the presymplectic form \tilde{Z}_ψ is explicitly written as*

$$\tilde{Z}_\psi(\dot{\psi}, \dot{\psi}) = \int_{\{\psi=0\}} \det(d\psi_N^{-1}) \operatorname{Re}(\dot{\psi} i \bar{\dot{\psi}}) d\mathcal{H}^n, \quad (4.4.1)$$

where $d\mathcal{H}^n$ is the n -dimensional Hausdorff measure induced by g .

We note that in (4.4.1), both $\det(d\psi_N^{-1})$ and $d\mathcal{H}^n$ depend on the metric g , but their product is independent of g .

Proof. Let $\gamma = \Pi\psi$ and $\dot{\psi} = -\mathcal{L}_u\psi + a\psi$, $\dot{\psi} = -\mathcal{L}_v\psi + b\psi$ with $u, v \in \operatorname{Diff}(M)$ and $a, b \in C^\infty(M; \mathbb{C})$. From $d\Pi(\dot{\psi}) = -\mathcal{L}_u\delta_\gamma \in T_{\delta_\gamma}\mathcal{D}_{\mathcal{O}}$ and $d\Pi(\dot{\psi}) = -\mathcal{L}_v\delta_\gamma$, we have

$$\Pi^*\omega(\dot{\psi}, \dot{\psi}) = \omega(-\mathcal{L}_u\delta_\gamma, -\mathcal{L}_v\delta_\gamma) = \int_\gamma \iota_v \iota_u \mu = \int_\gamma \iota_{v^\perp} \iota_{u^\perp} \mu$$

where u^\perp, v^\perp denote the orthogonal projections of u, v onto the normal space with respect to g on the zeros, which are defined on $\psi^{-1}(0)$.

From $u^\perp = d\psi_N^{-1}\dot{\psi}$ and $v^\perp = d\psi_N^{-1}\dot{\psi}$, we have

$$\begin{aligned} \int_\gamma \iota_{v^\perp} \iota_{u^\perp} \mu &= \int_\gamma \mu(d\psi_N^{-1}\dot{\psi}, d\psi_N^{-1}\dot{\psi}, \dots) \\ &= \int_\gamma \nu_N(d\psi_N^{-1}\dot{\psi}, d\psi_N^{-1}\dot{\psi}) \iota_{n_2} \iota_{n_1} \mu \\ &= \int_\gamma \det(d\psi_N^{-1}) \nu_{\mathbb{C}}(\dot{\psi}, \dot{\psi}) \iota_{n_2} \iota_{n_1} \mu, \end{aligned}$$

which reads the stated expression. □

Riemannian and Kähler structures

We can also describe Riemannian and formal Kähler structures we defined in Section 2.3.4 in terms of implicit representations.

We first observe a relation between the presymplectic form \tilde{Z} and the standard symplectic structure on the space of complex valued functions. The expression (4.4.1) of \tilde{Z} in Proposition 4.4.1 appears similar to the symplectic structure on $C^\infty(M, \mathbb{C})$ given by

$$\mathfrak{Z}(\dot{\psi}, \dot{\psi}) = \int_M \operatorname{Re}(\dot{\psi} i \bar{\dot{\psi}}) \mu$$

for $\dot{\psi}, \dot{\bar{\psi}} \in T_\psi C^\infty(M, \mathbb{C}) = C^\infty(M, \mathbb{C})$, which forms a Kähler structure, together with a Riemannian metric

$$\mathfrak{G}(\dot{\psi}, \dot{\bar{\psi}}) = \int_M \operatorname{Re}(\dot{\psi} \dot{\bar{\psi}}) \mu$$

and the almost complex structure $\tilde{\mathcal{J}}_i: \dot{\psi} \mapsto i\dot{\psi}$ on $C^\infty(M, \mathbb{C})$. In light of this, it is tempting to define a degenerate Riemannian metric \tilde{B} on $\mathcal{F}_\mathcal{O}$ by $\tilde{B}(\cdot, \cdot) = \tilde{Z}(\cdot, \tilde{\mathcal{J}}_i \cdot)$. Explicitly,

$$\tilde{B}_\psi(\dot{\psi}, \dot{\bar{\psi}}) = \int_{\{\psi=0\}} \det d\psi_N^{-1} \operatorname{Re}(\dot{\psi} \dot{\bar{\psi}}) d\mathcal{H}^n.$$

Then $(\tilde{B}, \tilde{Z}, \tilde{\mathcal{J}}_i)$ is formally a degenerate Kähler structure. At a glance, this setting seems natural and intrinsic as it does not require any Riemannian metric of the ambient manifold unlike the formal Kähler structure $(\mathcal{G}, \omega, \mathcal{J})$ on the explicit shape space, explained in Section 2.3.4.

However, we emphasize that the triple $(\tilde{B}, \tilde{Z}, \tilde{\mathcal{J}}_i)$ on $\mathcal{F}_\mathcal{O}$ does not descend onto the base space \mathcal{O} . To see this, first note that the presymplectic form \tilde{Z} is by design invariant under the action of $\operatorname{DC}_\gamma = \operatorname{Diff}_\gamma(M) \times C^\infty(M, \mathbb{C})$ on each fiber $\Pi^{-1}\gamma$ over $\gamma \in \mathcal{O}$ and $\ker \tilde{Z}_\psi$ agrees with

$$\ker d\Pi_\psi = \{-\mathcal{L}_u \psi + \varphi \psi \mid (u, \varphi) \in \operatorname{diff}_0(M) \times C^\infty(M, \mathbb{C}) \text{ s.t. } -\mathcal{L}_u \delta_\gamma = 0\}.$$

Hence \tilde{Z} descends onto \mathcal{O} and defines the MW symplectic structure ω by $\tilde{Z} = \Pi^* \omega$.

However, both \tilde{B} and $\tilde{\mathcal{J}}_i$ do not descend onto \mathcal{O} , and hence fail to define a Riemannian metric and an almost complex structure on \mathcal{O} , as shown in the following example. In particular, the diagram

$$\begin{array}{ccc} T\mathcal{F}_\mathcal{O} & \xrightarrow{\tilde{\mathcal{J}}_i} & T\mathcal{F}_\mathcal{O} \\ \downarrow d\Pi & & \downarrow d\Pi \\ T\mathcal{O} & \xrightarrow{\mathcal{J}} & T\mathcal{O} \end{array}$$

does not commute.

Example 4.4.2. The almost complex structure $\tilde{\mathcal{J}}_i$ behaves differently at ψ_0 and ψ_1 in the same fiber $\Pi^{-1}\gamma$ if they are in different conformal classes (Definition 3.3.4).

Consider the setting in Example 3.3.2 where $\psi_0(x, y) = x + iy$ and $\psi_1(x, y) = x + y + iy$ on \mathbb{R}^2 so $\Pi\psi_j = \delta_0$ for $j = 0, 1$. With a fixed vector field $u(x, y) = (1, 0)$ on \mathbb{R}^2 , we have $\mathcal{L}_u \psi_j = 1$ and $\tilde{\mathcal{J}}_i(\mathcal{L}_u \psi_j) = i \cdot 1 = i$ for both $j = 1, 2$. However, observe that $\tilde{\mathcal{J}}_i(\mathcal{L}_u \psi_j) = i = \mathcal{L}_{u_j} \psi_j$ with different fixed vector fields $u_0(x, y) = (0, 1)$ and $u_1(x, y) = (1, 1)$. Namely $d\Pi(\tilde{\mathcal{J}}_i(\mathcal{L}_u \psi_j)) = d\Pi|_{\psi_j}(i) = -\mathcal{L}_{u_j} \delta_0$ for $j = 1, 2$ yield different motions of the zero i.e., $-\mathcal{L}_{u_0} \delta_0 \neq -\mathcal{L}_{u_1} \delta_0$.

The ill-definedness of $\tilde{\mathcal{J}}_i$ for different conformal classes also affects the degenerate Riemannian structure \tilde{B} . For the above ψ_0 and ψ_1 we have

$$\det(d\psi_j) = \det \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = 1$$

for both $j = 0, 1$ with respect to the standard Euclidean metric, and hence $\tilde{B}_{\psi_j}(\dot{\psi}_j, \dot{\bar{\psi}}_j) = \operatorname{Re}(\dot{\psi}_j \dot{\bar{\psi}}_j)$. With a vector field $v = (0, 1)$, we have $\mathcal{L}_v \psi_0 = i$ and $\mathcal{L}_v \psi_1 = 1 + i$, namely

$$\begin{aligned} \tilde{B}_{\psi_0}(\mathcal{L}_v \psi_0, \mathcal{L}_v \psi_0) &= \operatorname{Re}(i \cdot (-i)) = 1, \\ \tilde{B}_{\psi_1}(\mathcal{L}_v \psi_1, \mathcal{L}_v \psi_1) &= \operatorname{Re}((1 + i) \cdot (1 - i)) = 2. \end{aligned}$$

Thus neither of $\tilde{\mathcal{J}}_i$ nor \tilde{B} descends to \mathcal{O} .

We may still construct a degenerate Riemannian metric on $\mathcal{F}_\mathcal{O}$ that descends onto \mathcal{O} . Let us define a degenerate metric \tilde{G} on $\mathcal{F}_\mathcal{O}$ by pulling back the metric \mathcal{G} on \mathcal{O} defined in Section 2.3.4:

$$\tilde{G} = \Pi^* \mathcal{G}.$$

Explicitly, we have

$$\tilde{G}_\psi(-\mathcal{L}_u\psi + a\psi, -\mathcal{L}_v\psi + b\psi) = \mathcal{G}_{\delta_\gamma}(-\mathcal{L}_u\delta_\gamma, -\mathcal{L}_v\delta_\gamma) = \int_\gamma \iota_{\hat{v}} \iota_u \mu,$$

for $-\mathcal{L}_u\psi + a\psi, -\mathcal{L}_v\psi + b\psi \in T_\psi \mathcal{F}_\mathcal{O}$ where $\gamma = \Pi\psi$, and \hat{v} is any vector field satisfying $\hat{v} \circ \gamma = \mathcal{J}(v \circ \gamma)$ for the almost complex structure $\mathcal{J}: T\mathcal{O} \rightarrow T\mathcal{O}$ defined on \mathcal{O} .

The degenerate metric \tilde{G} is compatible with the presymplectic form \tilde{Z} through the homomorphism $\tilde{J}: T\mathcal{F}_\mathcal{O} \rightarrow T\mathcal{F}_\mathcal{O}$ defined by

$$\tilde{J}_\psi(-\mathcal{L}_v\psi + \varphi\psi) = -\mathcal{L}_{\hat{v}}\psi + i\varphi\psi,$$

where, again, \hat{v} is chosen so that $\hat{v} \circ \gamma = \mathcal{J}(v \circ \gamma)$. In general, \tilde{J} may not satisfy $\tilde{J}^2 = -\text{id}$ pointwise, so it may fail to be an almost complex structure on $\mathcal{F}_\mathcal{O}$. However, the failure is controlled: the difference $\tilde{J} \circ \tilde{J}(\dot{\psi}) - (-\text{id})(\dot{\psi}) = (\tilde{J} \circ \tilde{J} + \text{id})(\dot{\psi})$ lies in $\ker d\Pi|_\psi$ for any $\dot{\psi} \in T_\psi \mathcal{F}_\mathcal{O}$. Thus, the following diagram commutes:

$$\begin{array}{ccc} T\mathcal{F}_\mathcal{O} & \xrightarrow{\tilde{J}} & T\mathcal{F}_\mathcal{O} \\ \downarrow d\Pi & & \downarrow d\Pi \\ T\mathcal{O} & \xrightarrow{\mathcal{J}} & T\mathcal{O} \end{array}$$

By construction, both \tilde{G} and \tilde{J} descend to the shape space \mathcal{O} , defining the Riemannian and the almost complex structure \mathcal{G} and \mathcal{J} , hence forming a formal Kähler structure $(\mathcal{G}, \omega, \mathcal{J})$ on \mathcal{O} .

4.5 Appendix: Curvature as monodromy

In the proof of Corollary 4.3.9 in this chapter and Corollary 7.3.5 in Chapter 7, we use the fact that the integral of curvature form is given as the monodromy of a horizontal lift, measured by a connection form. This is formalized as the following proposition.

Proposition 4.5.1. *Let $\pi: (E, \alpha) \rightarrow (B, \beta)$ be a prequantum circle bundle and Σ be a topological disc in B . Consider any smooth path $x: [0, 1] \rightarrow E$ such that it is horizontal with respect to α and $\pi \circ x$ parametrizes the closed curve $\partial\Sigma \subset B$. Then we have*

$$\int_\Sigma \beta = \text{Arg}(x(1), x(0)) \pmod{2\pi}$$

where $\text{Arg}(x(1), x(0)) \in [0, 2\pi)$ is defined by $e^{i \text{Arg}(x(1), x(0))} \triangleright x(1) = x(0)$.

Proof. First consider any smooth lift ℓ of Σ to E which is diffeomorphic in its image and contains $x(0)$. Let $y: [0, 1] \rightarrow E$ be the parametrization of $\partial(\ell(\Sigma))$ such that $\pi(x(t)) =$

$\pi(y(t))$ for any t . Since y is a closed curve bounding a topological disc $\ell(\Sigma)$, we have by the Stokes theorem that

$$\int_{\Sigma} \beta = \int_{\pi(\ell(\Sigma))} \beta = \int_{\ell(\Sigma)} \pi^* \beta = \int_{\ell(\Sigma)} d\alpha = \int_{\partial(\ell(\Sigma))} \alpha = \int_0^1 y^* \alpha.$$

We now define a non-closed path $\tilde{y}: [0, 1] \rightarrow E$ by

$$\tilde{y}(t) = e^{-i \int_0^t y^* \alpha} \triangleright y(t).$$

The path \tilde{y} is horizontal with respect to α , which can be verified by direct computation.

Since $y(1) = y(0) = \tilde{y}(0)$ by definition, we have

$$\tilde{y}(1) = e^{-i \int_0^1 y^* \alpha} \triangleright y(0) = e^{-i \int_0^1 y^* \alpha} \triangleright \tilde{y}(0)$$

and hence

$$\int_{\Sigma} \beta = \int_0^1 y^* \alpha = \text{Arg}(\tilde{y}(1), \tilde{y}(0)).$$

Finally, since $\tilde{x}(0) = \tilde{y}(0)$ by design, we have $x = \tilde{y}$ due to the uniqueness of the horizontal lift for a given initial point, which concludes the proof. \square

As in Corollary 4.3.9, when the connection form α has scaling factor in vertical reproducibility, the result of Proposition 4.5.1 is modified. For a connection α with a scaling factor $1/\mathcal{N}$ in vertical reproducibility *i.e.*, $\alpha(\hat{\xi}) = \frac{1}{\mathcal{N}}\xi$ for $\xi \in \mathfrak{g}$ and its fundamental vector field $\hat{\xi}$, we have

$$\int_{\Sigma} \beta = \frac{1}{\mathcal{N}} \text{Arg}(x(1), x(0)) \pmod{\frac{2\pi}{\mathcal{N}}}$$

since the horizontal vector field \tilde{y} in terms of a path y in the proof becomes $\tilde{y}(t) = e^{-i\mathcal{N} \int_0^t y^* \alpha} \triangleright y(t)$. Consequently, we have $\tilde{y}(1) = e^{-i\mathcal{N} \int_0^1 y^* \alpha} \triangleright \tilde{y}(0)$ and hence $\frac{1}{\mathcal{N}} \text{Arg}(\tilde{y}(1), \tilde{y}(0)) = \int_0^1 y^* \alpha = \int_{\Sigma} \beta$.

Hidden degrees of freedom in implicit vortex filaments

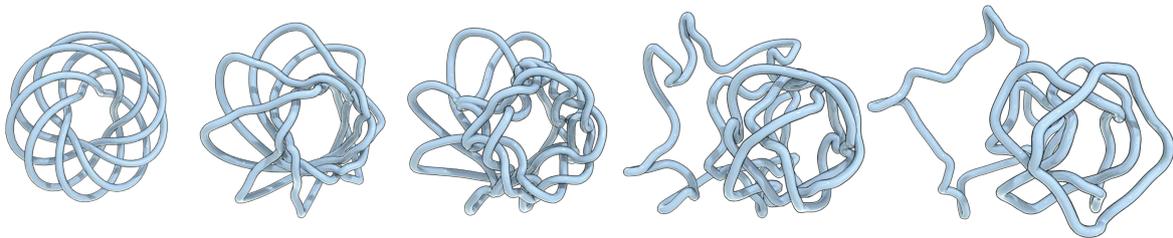


Figure 5.1: The time evolution of highly knotted vortex filaments. Our implicit description represents the dynamics of filaments as the zero levelsets of space-time complex-valued functions, which automatically handles topological changes of curves without yielding singularities.

This chapter is largely a reprint of the article:

Sadashige Ishida, Chris Wojtan, and Albert Chern. Hidden degrees of freedom in implicit vortex filaments. *ACM Transactions on Graphics*, 41(6), 2022.

This chapter develops a computational method for the dynamics of space curves, with emphasis on its applications to vortex filaments in fluid dynamics. Instead of representing these filaments with explicit curve geometry and Lagrangian equations of motion, we describe them using implicit representations via a class of complex functions. We leverage their redundant degrees of freedom in both the configuration and the dynamics, which can be tailored specifically to mitigate numerical instability that native approaches suffer from.

For handling chaotic systems specifically, we introduce *untwisted* Clebsch variables and *non-swirling* dynamics, which regularize sources of numerical instability, particularly in the twisting modes around curve filaments. Built on these, our simulation method robustly produces the dynamics of large numbers of interacting vortex filaments and effortlessly handles topological changes and re-connection events.

Changes from the original publication This chapter contains a few updates from the article [IWC22]. We provide a more precise definition of the equivalence class of complex functions that represent the same space curve. The discussion and results concerning the twist of complex functions have also been revised to reflect this refined definition.

Dependency of the chapter This chapter is self-contained, as it reintroduces the necessary concepts from previous chapters. Readers interested in the theoretical foundations of implicit representations may find Chapter 3 useful.

5.1 Introduction

The deformation of space curves is an interesting topic in many subjects such as differential geometry, low-dimensional topology, classical and quantum fluid mechanics, and electromagnetism. One example from fluid mechanics is the dynamics of vortex filaments. In a nearly inviscid fluid, vorticity originates from *codimension-1* interfaces or obstacle surfaces. The vortex sheets subsequently roll up into *codimension-2* vortex filaments, due to the Kelvin–Helmholtz instability. Hence, most physical inviscid fluids have their vorticity concentrated into a sparse set of space curves, rather than distributed evenly throughout space. Based on this observation, certain physical equations model fluids only with dynamically deforming space curves. Many fluid simulation methods take advantage of this sparsity structure.

One major challenge for an explicit (Lagrangian) filament-based fluid solver is to handle reconnection events when filaments collide. Without any reconnection, the total length of filaments can grow exponentially, exploding the computational cost and halting the solver. Hence, existing explicit filament simulators include a tedious process of collision detection followed by non-differentiable heuristic geometry surgeries.

To that end, implicit (Eulerian) curve representations are more appealing. The recently emerging *Clebsch representation* expresses vortex lines as level sets of a 2-dimensional-valued function called *Clebsch variables* [Cle59]. Like any level set method, topological changes of level set geometries occur gracefully. The difficulty, however, in a Clebsch-based fluid solver is in the dynamics of the Clebsch variables. The Clebsch variables satisfy the transport equation advected with the fluid velocity, which unfortunately behaves in a swirling motion with a high-spatial frequency and singularities near the vortex filaments. Such a rough transporting vector field is hard to resolve accurately in a computational grid. Even if the transport equation is computed accurately, the level set function will quickly evolve into a twisted and distorted function that is difficult to deal with.

This chapter develops a new approach for describing the geometry and dynamics of filaments with implicit curve functions. Our main insight is that the problem has a huge number of redundant degrees of freedom: both the velocity field and the level set function (i.e., the Clebsch representation) can be varied in ways that do not change the solution. We exploit these additional degrees of freedom to ensure stable numerical simulation and automatic handling of topological changes, without sacrificing accuracy. In particular, we choose an *untwisted* Clebsch representation for the level set geometry, and *non-swirling* dynamics for advecting vortex filaments. We regularize these functions by identifying and constraining hidden degrees of freedom in their representations, allowing us to greatly improve numerical robustness compared to naive implementations.

Our algorithm is the first method for animating implicit vortex filament geometry with automatic topological changes. Our mathematical formulation offers new tools for future research on fluid simulation and curve geometry processing, and our results show greatly improved stability compared to a standard level set implementation, with fewer user parameters than for explicit Lagrangian filament techniques.

5.2 Related work

We review previous work, highlighting numerical methods for simulating the time evolution of vorticity in incompressible fluids. These vortex-capturing schemes seek a useful representation of the vorticity and solve its governing equations of motion.

Explicit vortex methods In previous vortex methods, vortices are represented either as particles [GLG95, PK05, SRF05, ZB14, Ang17], filaments [CK⁺00, AN05, WP09, WP10, PCK⁺19], segments [Cho90, XTZ⁺21], sheets [BKB12, PTG12, DBWG15] or volumes [ETK⁺07, ZBG15]. Vortex particle methods represent vortices as a disconnected point cloud. However, the strength of vortex per particle or per unit volume [ZBG15] undergoes a numerically unstable *vortex stretching*, requiring an artificial clamping or diffusion that sacrifices accuracy. The stretching problem is avoided by representing vorticity per filament, segment, or sheet, or per unit area using differential 2-forms [ETK⁺07]. However, describing vortex explicitly (Lagrangian method) with filaments, segments and sheets comes with a cost of sophisticated and heuristic treatment for changes of vortex topology. Volumetric (Eulerian) methods [ETK⁺07] do not require managing topological changes, but they do not have a handle on codimensional structures. Our codimension-2 level set method is an Eulerian method that can represent sharp filament structures without any additional difficulty from vortex reconnection.

Clebsch representations Another Eulerian representation of vorticity is to describe vortex lines as the level sets of a 2D-valued function known as the Clebsch variable [Cle59, Lam95]. The representation was not widely adopted since all \mathbb{R}^2 -valued Clebsch variables can only describe fluids with zero *helicity* [CKPS17]. The helicity problem is solved by using a sphere-valued Clebsch variable [KM80] which can represent nonzero (though quantized) values of helicity. Since the recognition of these “spherical” Clebsch maps, they have become an established method for vortex representation in computer graphics [CKP⁺16, CKPS17, YXZ⁺21]. However, Clebsch variables represent a *smooth vorticity field* by the continuum of level sets of all values. It remains a challenge to represent a *sharp codimension-2 filament*, especially with a limited grid resolution. By contrast, our method represents a sharp filament just as the *zero set* of a complex-valued function. Moreover, we show in Section 5.4.1 that our representation can actually resolve continuous values of helicity (as opposed to the quantized values mentioned earlier).

Dynamics of Clebsch variables In addition to the implicit representation of vorticity, Clebsch variables also play significant roles in a variational and Hamiltonian formulation for the incompressible Euler equation [Cle59, Lam95, Mor98, Che17]. In short, one (of many possible) governing equation(s) for the Clebsch variable is exceedingly simple: the Clebsch variable is advected by the fluid velocity. This equation of motion is recently adopted by [YXZ⁺21]. However, referred to as the *Lagrangian chaos*, a direct transportation by the fluid velocity quickly stirs and twists any variable to a distorted one unresolved by the finite computational grid [QZG⁺19]. The Clebsch variable is no exception under such dynamics. The method of *Schrödinger’s Smoke* [CKP⁺16] bypassed the Lagrangian chaos: its total energy (Hamiltonian) includes the Dirichlet energy of the Clebsch variable, which is therefore bounded for all time. However, while the dynamics of *Schrödinger’s Smoke* appear to be similar to that of Euler’s equation, it is only an approximation to the Euler fluid. There is still a large degree of freedom in the Clebsch representation and its dynamical system. Finding an equation of motion for the Clebsch variable that both describes the correct Euler fluid without Lagrangian chaos is

an unexplored research topic. This chapter describes an instance of a Lagrangian-chaos-free dynamical system for an implicit representation of vortex filaments.

Implicit filament representations We represent filaments as the zero set of a complex-valued function. These zero sets of a complex phase field are widely studied in condensed matter physics as *topological defects* appearing in superfluids and superconductors [BBH⁺94, Pis99]. These topological defect models also facilitate singularity placements in flow analysis and geometry processing [WPS14, SVB17, PBS20]. Complex phase field models are taken more generally as high-codimensional level set representations by [AS96, RMXO01, BCMO01, Min04]. However, in the physics and geometry processing literature, the level set functions have specific physical meanings such as the phase of a wave, leaving little room for a smoother representative. In most cases, the phase field has norm 1 except for a sudden dip to 0 near the filaments, creating a configuration that is difficult to resolve efficiently on a computational grid. In the level set method literature, the norm-1 condition is often adopted for (re)initializing the level set functions despite the discontinuity [RMXO01]. The complex phase is constructed locally with little discussion about global topological obstruction. [BCMO01] addressed the challenges in reinitializing the multi-component level set functions; by mimicking the codimension-1 signed distance functions, they propose a sophisticated reinitialization by solving a “manifold eikonal equation” along the isosurface of each function component. Unfortunately the process will not resolve the twists of the framed curve. To our knowledge, there has not been a thorough discussion about the degrees of freedom in the implicit filament representations or in their dynamics until now.

In this chapter, we explore the degrees of freedom of both the codimension-2 level set functions and their equations of motion. We further provide a simple reinitialization method comparable to the codimension-1 signed distance function. Comparisons show that exploiting these degrees of freedom are essential to a robust simulation.

5.3 Representations for evolving curves

We now begin our description of implicit filament dynamics. The main mathematical object is a union of closed space curves. This section describes an implicit representation for these curves and their dynamics, as well as the degrees of freedom in the representation.

5.3.1 Representations for curve configuration

Let the physical domain be an open region $M \subset \mathbb{R}^3$. We use γ to represent a collection of m closed curves

$$\gamma: \bigsqcup_{i=1}^m \mathbb{S}^1 \rightarrow M,$$

where \bigsqcup denotes a disjoint union, \mathbb{S}^1 is the topological circle, m is the number of filaments, and γ is the mapping from the 1D curves into 3D. The configuration space \mathcal{X} of these filaments is the space of all possible placements of these curves:¹

$$\mathcal{X} = \bigsqcup_{m=1}^{\infty} \{ \gamma: \bigsqcup_{i=1}^m \mathbb{S}^1 \rightarrow M \} / \text{reparametrizations.} \quad (5.3.1)$$

¹Precisely, reparametrizations mean the quotient by the orientation preserving diffeomorphisms on $\bigsqcup_{i=1}^m \mathbb{S}^1$ for each fixed m (See Section 2.2.1 for detail).

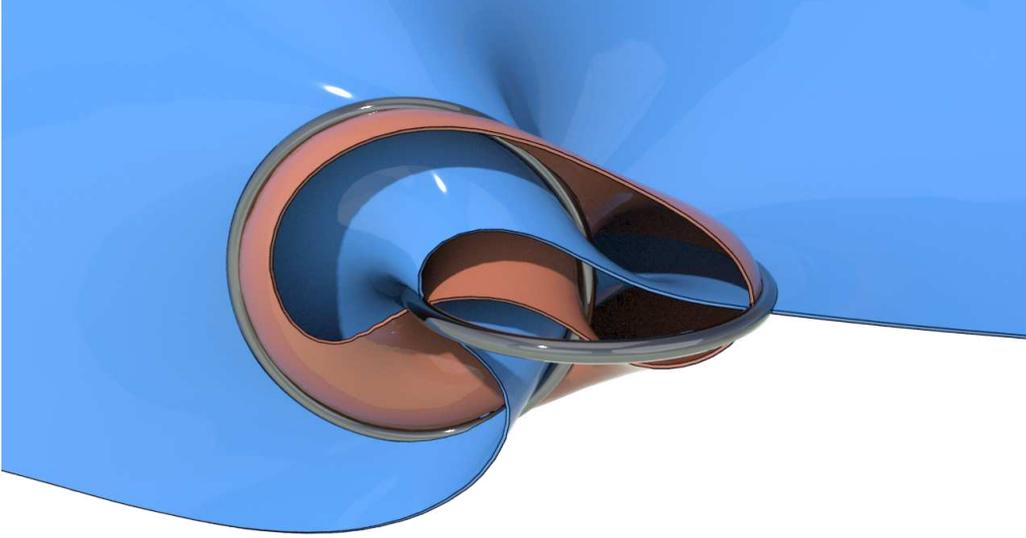


Figure 5.2: An example of level surfaces of two linked curves. The blue and red surfaces are cross sections of $\{\text{Im}\psi = 0\}$ and $\{\text{Re}\psi = 0\}$ respectively.

Although it is straightforward to represent curves *explicitly* via objects γ in (5.3.1) as parameterized curves or their discrete counterparts, topological changes of curves such as splitting or merging are difficult to describe mathematically and algorithmically. For instance, the number m of components can change when curves reconnect or split apart.

Instead of relying on an explicit curve representation, our work adopts an *implicit* representation for the elements in \mathcal{X} . We model every collection of closed curves in M as the zero set of a complex-valued level set function $\psi: M \rightarrow \mathbb{C}$:

$$\gamma = \{p \in M \mid \psi(p) = 0\} = \{\text{Re}\psi = 0\} \cap \{\text{Im}\psi = 0\},$$

In other words, an alternative definition for γ is the set of all points p where *both* the real and imaginary components of a level set function ψ evaluate to zero as in Figure 5.2. We find it useful to draw an analogy to the scalar-valued level sets commonly used in computer graphics applications: the zero level-set of a scalar-valued function represents shapes of codimension 1 (*a.k.a.* surfaces), while our complex-valued level set has twice as many variables and thus represents shapes of codimension 2 (curves).

We next note that different functions ψ can represent the same collection of curves if they share the same zeros and orientation. To formulate this redundancy precisely, we define the following equivalence relation \sim on the function space $(M \rightarrow \mathbb{C})$. We say $\psi_1 \sim \psi_2$ if their zeros agree *i.e.*, $\psi_1^{-1}(0) = \psi_2^{-1}(0)$, and share the same orientation, that is,

$$\int_{\partial\Sigma} \text{Im} \frac{d\psi_1}{\psi_1} = \int_{\partial\Sigma} \text{Im} \frac{d\psi_2}{\psi_2}$$

for any oriented topological disc Σ intersecting transversely with γ . For details, see Section 3.2.

For the implicit representation to work properly, we make a regularity assumption that the level set functions ψ are sufficiently “nice” around the zero set. Precisely, we assume ψ is smooth, and the differential $d\psi|_p: T_p M \rightarrow \mathbb{C} \cong \mathbb{R}^2$ is surjective, *i.e.* its matrix form has rank 2 at any point p on the zero-set.

To summarize, our configuration space of filaments \mathcal{X} from (5.3.1) is replaced by the space of complex-valued functions ψ modulo the above equivalence class:

$$\mathcal{F} := \{\psi: M \rightarrow \mathbb{C}\} / \sim .$$

While \mathcal{X} and \mathcal{F} describe the same configuration space of filaments, the objects in \mathcal{F} are much more continuous compared to the disjoint spaces of \mathcal{X} .

Relation to Clebsch representations The representation of codimension-2 curve geometries in 3D is known in fluid dynamics as *Clebsch representations* [Cle59, Lam95, CKPS17, YXZ⁺21]. For a fluid flow with a smooth vorticity field, the vortices are geometrically depicted as fibrous vortex lines diffusely distributed over the fluid domain. A Clebsch representation aims at an implicit representation for such fibrous structure. The representation uses a map $s: M \rightarrow \Sigma$ from the 3D fluid domain M to a 2-dimensional manifold Σ with a measure σ to describe the *vortex lines* as preimages $s^{-1}\{p\}$ of points $p \in \Sigma$, and the density of the vortex lines as the pullback $s^*(\sigma)$ of the measure σ . A smooth Clebsch map s and a smooth measure σ yields a smooth distribution of vortex lines. To achieve a more singular and concentrated vorticity field such as vortex filaments, one would consider s with larger derivatives [MW83] (s sweeps out more measure σ over a small area in M).

In our setup, we want to represent singular curves with a Dirac- δ density, instead of diffused distribution of vortex lines. Previous considerations in Clebsch representations would set the Clebsch map s with enormous derivative. By contrast, we obtain such concentrated filaments by setting σ singular while keeping the Clebsch map smooth. Our complex level function $\psi: M \rightarrow \mathbb{C}$ is a Clebsch map with the target space $\Sigma = \mathbb{C}$ equipped with a δ -measure $\sigma = \delta_0$ at the origin.

An important discussion about Clebsch representations [CKPS17] is whether or not a fluid configuration can be represented with the choice of Σ and σ . Previous Clebsch representations [CKPS17, YXZ⁺21] adopt $\Sigma = \mathbb{S}^2$, since the more straightforward choice of $\Sigma = \mathbb{R}^2 \cong \mathbb{C}$ with the standard area measure σ can only represent fluid flows with zero *helicity*. The helicity obstruction is reduced for $\Sigma = \mathbb{S}^2$ as it can admit a discrete set of nonzero helicity [CKPS17]. Our Clebsch representation can represent any space curve without any obstruction. In particular, the helicity of a vortex filament is proportional to its *writhe* [AK21] which can take any real value.

5.3.2 Representations for curve dynamics

In the explicit representation, a first-order time evolution of curves $\gamma: (\bigsqcup^m \mathbb{S}^1) \times \mathbb{R} \rightarrow M$ can be described by an equation of the form:

$$\frac{\partial \gamma}{\partial t}(s, t) = V_{\gamma_t}(s), \quad s \in \bigsqcup \mathbb{S}^1, t \in \mathbb{R}. \quad (5.3.2)$$

Here, s is the parameterization of the curve, t is time, and the *velocity* $V_{(\cdot)}: \mathcal{X} \rightarrow (\bigsqcup \mathbb{S}^1 \rightarrow \mathbb{R}^3)$ is a dynamical model that tells the filament how to move based on the current filament shape and position.

Example 5.3.1. In the context of fluid dynamics, important examples for the velocity model V are the ones that govern the motion of vortex filaments. When $M = \mathbb{R}^3$, *i.e.* there are no

obstacles or boundaries, the velocity models are the Biot–Savart model

$$V_{\gamma}^{\text{BS}}(s) := \frac{\Gamma}{4\pi} \oint \gamma'(\tilde{s}) \times \frac{\gamma(s) - \gamma(\tilde{s})}{|\gamma(s) - \gamma(\tilde{s})|^3} d\tilde{s} \quad (5.3.3)$$

and the more regular Rosenhead–Moore model [Saf93, pp. 213]

$$V_{\gamma}^{\text{RM}}(s) := \frac{\Gamma}{4\pi} \oint \gamma'(\tilde{s}) \times \frac{\gamma(s) - \gamma(\tilde{s})}{\sqrt{e^{-3/2}a^2 + |\gamma(s) - \gamma(\tilde{s})|^2}^3} d\tilde{s} \quad (5.3.4)$$

where the constants Γ and a are the vortex strength and vortex thickness respectively, and the integrating measure $d\tilde{s}$ is the arclength element (set $a = 0$ for the Biot–Savart model). Note that (5.3.3) and (5.3.4) are the restrictions at the curve of the entire fluid velocity field over the 3D domain

$$U_{\gamma}^{\text{RM}}(x) := \frac{\Gamma}{4\pi} \oint \frac{\gamma'(\tilde{s}) \times (x - \gamma(\tilde{s}))}{\sqrt{e^{-3/2}a^2 + |x - \gamma(\tilde{s})|^2}^3} d\tilde{s}, \quad x \in \mathbb{R}^3 \quad (5.3.5)$$

That is, $V_{\gamma}^{\text{RM}}(s) = U_{\gamma}^{\text{RM}}(\gamma(s))$.

Now, we translate the dynamical system (5.3.2) into an evolution equation for a time-dependent complex level function ψ . First, we note that the evolution of ψ around the zeros is given as the *transport equation* along some vector field, which is formally the following lemma.

Lemma 5.3.2. *For any time-dependent complex level set function ψ with the regularity assumptions in Section 5.3.2, there exists a neighborhood $U \subset M$ of the zero set of ψ and a vector field $v: U \rightarrow \mathbb{R}^3$ such that*

$$\frac{\partial \psi}{\partial t} + v \cdot \nabla \psi = 0 \quad \text{in } U. \quad (5.3.6)$$

Proof. By the regularity assumptions in Section 5.3.2, there exists a neighborhood U of the zero set of ψ where the 3D-to-2D linear map $d\psi$ is full-rank and thus surjective. Hence at every point $x \in U$ and for any value of $\partial\psi/\partial t(x)$, there exists a vector $v_x \in \mathbb{R}^3$ such that $d\psi|_x(v_x) = -\partial\psi/\partial t(x)$. With this construction, we obtain a vector field $v: U \rightarrow \mathbb{R}^3$ satisfying $\partial\psi/\partial t + v \cdot \nabla \psi = 0$ in U . \square

Hence the representation of the dynamics for ψ can be encapsulated into a *vector field*. Observe that the zero level curve γ of ψ evolves by $\partial\gamma/\partial t = v|_{\gamma}$ under the transport equation (5.3.6), and that $\ker(d\psi)$ at each point on γ is spanned by the tangent γ' as γ' lies on the tangent spaces of both $\{\text{Im}\psi = 0\}$ and $\{\text{Re}\psi = 0\}$. By matching these induced curve dynamics in Lemma 5.3.2 with (5.3.2) we conclude:

Theorem 5.3.3. *The zero level curve γ of ψ evolves according to (5.3.2) if and only if ψ satisfies (5.3.6) for some vector field v that agrees with the curve velocity at the curve:*

$$v(\gamma(s)) = V_{\gamma}(s) + f(s)\gamma'(s) \quad (5.3.7)$$

where $f\gamma'$ is the tangent vector γ' multiplied by an arbitrary scalar function f . The degrees of freedom of the dynamics for ψ are the degrees of freedom for choosing v with the condition (5.3.7).

Essentially, the only velocities that really matter for the evolution of the curves are the velocities located *on the zero level set*. Furthermore, the locations of the curves will not change if we slide them around their tangent direction (like spinning a circle around its axis of symmetry), so we only need to pin down their *normal and bi-normal components*. So we have a huge number of velocity variables (3 for each point in the 3D domain) with very few constraints (2 for each point on the 1D curves). This under-determined system gives us a *redundancy* in possible velocity fields for curve dynamics, which is largely unexplored by previous work.

Let us apply Theorem 5.3.3 to vortex filament dynamics (*cf.* Example 5.3.1): Plugging in (5.3.5) for v gives us the following dynamics for ψ :

$$\frac{\partial}{\partial t}\psi + v \cdot \nabla\psi = 0, \quad v(x) = U_\gamma^{\text{RM}}(x). \quad (5.3.8)$$

This is the most straightforward way to do it: simply advect the level set ψ in the exact same way as the rest of the fluid. However, we know that $v = U_\gamma^{\text{RM}}$ is an extremely sensitive function to deal with numerically — it tends to infinity near γ , has unbounded derivatives, and rapidly changes direction in very tight swirls. Small errors in γ inevitably create huge errors in velocity, making simulations unstable, as demonstrated in our accompanying video. Fortunately, according to Theorem 5.3.3 we now know that there are infinitely many velocity fields that will all theoretically give us the same filament motions; our mission in the next section is to swap out this unstable Biot-Savart velocity field for one that is much more numerically robust.

Remark 5.3.4. When a curve has a self-intersection, that is, there is a point $p = \gamma(s) = \gamma(s')$ for some $s \neq s'$, the differential $d\psi_p$ fails to be surjective. Consequently, Lemma 5.3.2 and hence Theorem 5.3.3 do not apply. Intuitively, the time evolution of a self-intersecting curve is not described by the advection of ψ along a single vector field as the overlapping points $\gamma(s)$ and $\gamma(s')$ may have different velocities.

In the numerical scheme we present in Section 5.5, we handle this issue by introducing a mechanism that automatically resolves self-intersections as soon as they appear.

5.4 Untwisted Clebsch variables and non-swirling dynamics

5.4.1 Untwisted Clebsch variables

Like the common codimension-1 real-valued level set methods, the implicit representation benefits from the regularity of the level set function. There, a level set function is well-conditioned if the magnitude of the gradient is close to one. For that reason, the level set function is typically initialized as the *signed distance function*, and this property is typically maintained as the level set is evolved (called *re-distancing*).

For our codimension-2 complex level set representation, we shall also characterize a set of desirable qualities of the complex level set function $\psi: M \rightarrow \mathbb{C}$. Due to the higher codimension, the discussion involves the notion of *twist* from the mathematical *ribbon theory*. Finally, we describe a concrete construction of ψ that will be used for initialization and re-distancing.

Conditioning of a complex level set function

We want ψ to be continuous everywhere and non-zero outside the curves $\gamma = \{\psi = 0\}$.

Near the curves γ , we want the differential $d\psi$ to be well-conditioned: if we only consider the function restricted to the plane spanned by the curve normal and bi-normal, $d\psi|_{\gamma^\perp}$ is close to an isometry; that is $d\psi|_{\gamma^\perp}: \gamma^\perp \rightarrow \mathbb{C} \cong \mathbb{R}^2$ has singular values $\approx (1, 1)$. Notice, however, that even when $d\psi|_{\gamma^\perp}$ is well-behaved on each normal plane, the level set function ψ can still exhibit significant *shearing* if the complex phase varies significantly along the tangent direction.

This variation of complex phase along the curve's tangent can be seen more intuitively via the geometry of the surface S_ψ formed by the level sets of the positive real part of ψ ($S_\psi = \{\operatorname{Re}(\psi) < 0, \operatorname{Im}(\psi) = 0\} = \{\arg(\psi) = \pi\}$). Note that the boundary of the surface S_ψ is γ , as illustrated in Figure 5.3. The tangential variation of the phase $\arg(\psi)$ is embodied by the twist of S_ψ at its boundary γ .

For the specific method presented in this chapter, we take a simple construction of ψ that is known to have a smooth S_ψ with little twisting at γ . We first lay out a few geometric and topological properties about the twist. Readers who are only concerned about the simulation method can skip ahead to the last part of this subsection, *Solid-angle distance function*.

The twist of a complex level set function

To discuss the twist precisely, we first review the standard definition of twist for a ribbon of a curve. Then, we apply this definition to the ribbon induced by ψ , which is tangent to the surface S_ψ .

Definition 5.4.1 (Ribbon and twist). Let U be a unit normal vector field over a collection of non self-intersecting space curves γ . A *ribbon* R of γ is defined as another collection of curves $R = \gamma + \epsilon U$ where $\epsilon > 0$ is taken sufficiently small so that R does not intersect with itself and γ .

The *twist* of the pair (γ, R) at each point on γ is given by

$$\omega := U' \cdot (\gamma' \times U)$$

where $(\cdot)'$ is the derivative along curve's tangent relative to arclength.

Then the *total twist* of (γ, R) is defined as

$$\operatorname{Tw}(\gamma, R) := \frac{1}{2\pi} \oint_{\gamma} \omega.$$

We now consider the ribbon induced by ψ .

Definition 5.4.2 (ψ -induced framing). Each complex level set function $\psi: M \rightarrow \mathbb{C}$ for γ gives rise to a normal vector field $U_\psi: \gamma \rightarrow \mathbb{R}^3$, $|U_\psi| = 1$, $\langle \gamma', U_\psi \rangle = 0$ that points to the direction where ψ is real and positive. Explicitly, at each point on γ ,

$$U_\psi := \frac{(d\psi|_{\gamma^\perp})^{-1}(-1)}{|(d\psi|_{\gamma^\perp})^{-1}(-1)|}.$$

This defines the ribbon $R_\psi := \gamma + \epsilon U_\psi$ with a small enough $\epsilon > 0$ as in Definition 5.4.1. Using this, we define the twist ω_ψ and the total twist $\operatorname{Tw}(\psi)$ as those of the ribbon (γ, R_ψ) .

When $d\psi|_{\gamma^\perp} : \gamma^\perp \rightarrow \mathbb{C}$ is conformal (which is true, for the solid-angle distance function we will introduce in this subsection, since it is isometric), the twist directly relates to the derivative of ψ along the curve:

$$\omega_\psi = -(d\theta)(\gamma') = -\operatorname{Re} \left(\frac{(d\psi)(\gamma')}{i\psi} \right), \quad \theta = \arg \psi,$$

in a small neighborhood of γ .

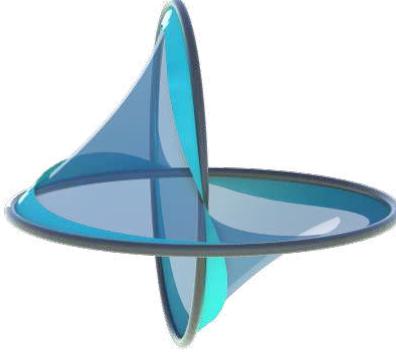


Figure 5.3: The surface S_ψ (translucent blue) and the ribbon U_ψ (opaque cyan).

Now, a natural consideration for designing ψ is to minimize the twist. For example, one may try to construct *Bishop's parallel frame* [BWR⁺08], which has no twist. However, one would find the construction *impossible* for a closed curve and a union of closed curves in general. This is due to the following theorem that the total twist $\operatorname{Tw}(\psi)$ is fixed by the curve geometry $\gamma \in \mathcal{X}$ or the equivalent class of level functions $[\psi] \in \mathcal{F}$, and hence non-zero unless $\operatorname{Tw}(\gamma) = 0$.

Theorem 5.4.3 (Invariance of total twist). *Suppose $M = \mathbb{R}^3$. Let $[\psi] \in \mathcal{F}$. Then any representative $\psi \in [\psi]$ has the same total twist $\operatorname{Tw}(\psi)$.*

To prove the theorem, we use an auxiliary setting.

Definition 5.4.4 (Seifert surface and Seifert framing). A Seifert surface Σ of γ is an oriented and compact surface such that $\partial\Sigma = \gamma$.² A ribbon R of γ is called a Seifert framing if there is a Seifert surface Σ such that R is tangent to Σ on $\gamma = \partial\Sigma$.

Definition 5.4.5 (Ribbon helicity). Let γ_i be each connected component of γ and R_i be a ribbon of γ_i . Then the *ribbon helicity* of γ_i is

$$H(\gamma_i, R_i) = \operatorname{Link}(\gamma_i, R_i) + \sum_{j \neq i} \operatorname{Link}(\gamma_i, \gamma_j)$$

where Link is the linking number between non-intersecting curves.

The linking number in Definition 5.4.5 measures how much two curves are linked. For details we refer to [AK21].

When the ribbon is a Seifert framing, the helicity of the ribbon vanishes ([DWCWR21, Theorem 3]):

²Note that a standard definition of Seifert surface additionally requires connectedness [Sei35, Mur96], but we drop this condition for our purposes.

Lemma 5.4.6 (Zero ribbon helicity of Seifert framing). *If the ribbon R of a collection of curves γ is a Seifert framing, we have*

$$H(\gamma_i, R_i) = 0, \quad \forall i$$

where (γ_i, R_i) are each connected component of γ and its associated component of R .

Another key ingredient is the Călugăreanu theorem:

Lemma 5.4.7 (Călugăreanu). *For a single curve γ and a ribbon R , we have*

$$\text{Link}(\gamma, R) = \text{Tw}(\gamma, R) + \text{Wr}(\gamma)$$

where $\text{Wr}(\gamma)$ is the writh of γ (for details, see [AK21]).

We are now ready to prove the theorem.

Proof of Theorem 5.4.3. Without loss of generality, we can assume a constant phase-shift to ψ so that S_ψ is a Seifert surface of γ .³ Hence for the ribbon R_ψ , we have $H(\gamma_i, R_{\psi,i}) = 0$ and namely $\text{Link}(\gamma_i, R_{\psi,i}) = -\sum_{j \neq i} \text{Link}(\gamma_i, \gamma_j)$ for each i , due to Lemma 5.4.6. Applying the Călugăreanu theorem (Lemma 5.4.7), we have

$$\begin{aligned} \text{Tw}(\psi) &= \text{Tw}(\gamma, R_\psi) = \sum_i \text{Tw}(\gamma_i, R_{\psi,i}) \\ &= \sum_i -\text{Wr}(\gamma_i) + \text{Link}(\gamma_i, R_{\psi,i}) \\ &= -\sum_i \left(\text{Wr}(\gamma_i) + \sum_{j \neq i} \text{Link}(\gamma_i, \gamma_j) \right). \end{aligned}$$

This show that $\text{Tw}(\psi)$ depends only on γ . □

Note that even though the total twist is fixed among possible choices of ψ , strong *local* twist can still be present.

Solid-angle distance function

As explained above, we desire a complex level set function ψ that is numerically well-conditioned: it should be nearly isometric near the curve γ and non-zero outside γ . In order to isolate these properties and enforce them explicitly, we model ψ with a complex wave function:

$$\psi(x) = r(x)e^{i\theta(x)}.$$

Similar to *signed distance functions* for the codimension-1 level set functions, we set

$$|\psi(x)| = r(x) := \text{dist}(\gamma, x), \quad x \in M \subset \mathbb{R}^3.$$

This setup ensures that the value $\psi(x)$ is non-zero for $x \notin \gamma$. What remains is a choice for the complex phase $\theta: M \setminus \gamma \rightarrow \mathbb{S}^1 = \mathbb{R}_{\text{mod } 2\pi}$, which we set to the *half solid-angle subtended by γ* :

$$\theta(x) := \frac{1}{2} \text{SolidAngle}(\gamma; x) \pmod{2\pi}$$

where $\text{SolidAngle}(\gamma; x) \in \mathbb{R}_{\text{mod } 4\pi}$ is a dimensionless quantity given by the signed spherical area⁴ enclosed by the projection of γ on the unit sphere centered at x i.e. $\text{Proj } \gamma_x(s) :=$

³The connectedness of S_ψ is not necessarily guaranteed, but Definition 5.4.4 and Lemma 5.4.6 do not require connectedness.

⁴Our definition of the solid angle is modulated by 4π as the spherical polygon of projected curves may cover the sphere multiple times.

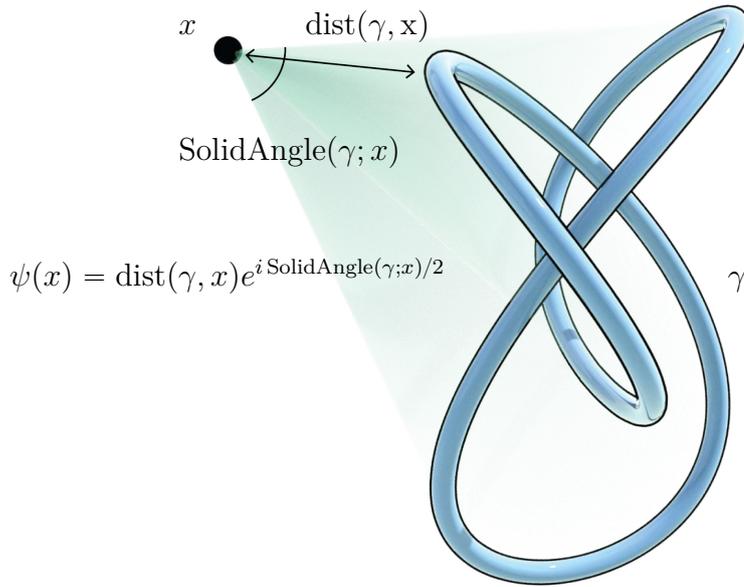


Figure 5.4: The solid-angle distance function ψ for a space curve γ is constructed by the distance and the angle subtended by the curve.

$(\gamma(s) - x)/|\gamma(s) - x|$. Combining these two, we now have the *solid-angle distance function*,

$$\psi(x) = \text{dist}(\gamma, x)e^{i \text{SolidAngle}(\gamma, x)/2} \quad (5.4.1)$$

illustrated in Figure 5.4.

The solid-angle distance function meets our desired conditions for the codimension-2 level set representation. On each normal plane $\theta|_{\gamma^\perp}$ is asymptotically the 2D angle function about the zero γ , so that $d\psi|_{\gamma^\perp}: \gamma^\perp \rightarrow \mathbb{C}$ is close to an isometry.

Moreover, as studied by [BA18], the surface $S_\psi = \{\theta = 0\}$ features little twist at γ . Figure 5.5 illustrates this concept.

Remark 5.4.8 (Properties of untwisted Clebsch variable). We note that our choice of the Clebsch variables ψ is an instance of many reasonable ones that exploit degrees of freedom in curve representations rather than the optimal one for a specific dynamical system such as vortex filaments.

Nevertheless, our ψ has a number of desirable properties. For example ψ is uniquely computed from given curves with an explicit formula, and is locally little twisted with $d\psi|_{\gamma^\perp}$ being locally nearly-isometry, as we saw in this section. In addition, $\arg \psi$ is harmonic and the zeros of the real and the imaginary parts intersect orthogonally, which ensures the zeros to be always codimension-2 and makes numerically extracting zeros robust.

Remark 5.4.9 (Open curves). Our construction of ψ is not limited to closed curves but is also valid for open curves with end points located on the boundary or obstacles subject to the following integrability condition. The orientation of each curve assigns a positive or negative signature for its two end points. The collection of curves is said to be integrable if all ends of the curves lie on the boundary and each connected component of the boundary contains an equal number of positive and negative ends. Note that vortex filaments must

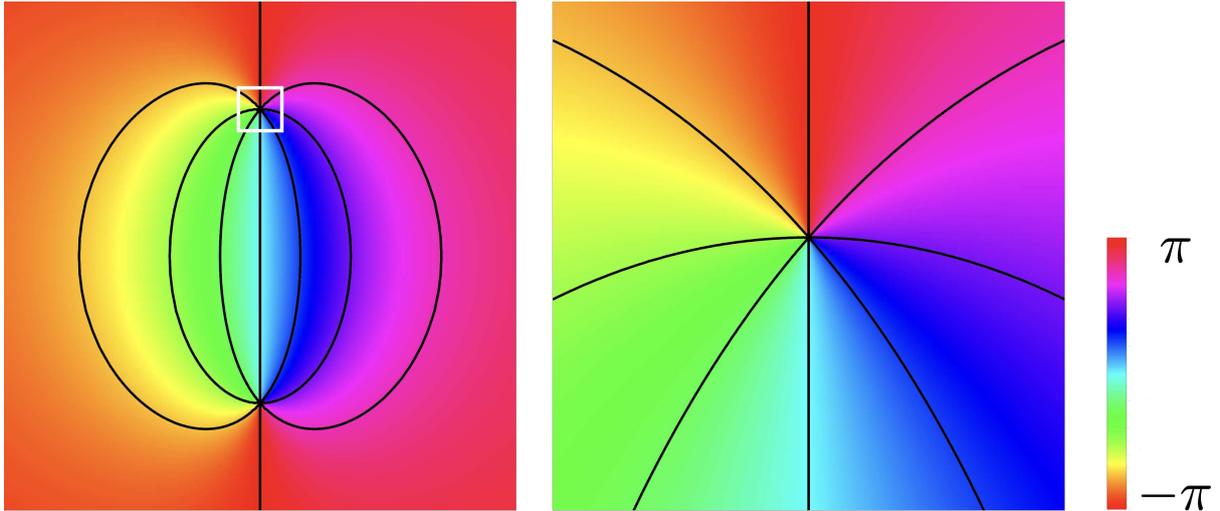


Figure 5.5: Plotting θ on a 2D plane which intersects a circular vortex ring at two points (left). The color indicates the value of the θ , and the black lines are its level curves. The curves meet where the filament intersects the plane. Zooming into the white box (right) shows evenly-spaced curves closer to the filament, where $e^{i\theta}$ resembles the complex plane \mathbb{C} .

be integrable since this integrability condition is precisely the integrability condition for curl: The vorticity 2-form is the exterior derivative of a velocity 1-form if and only if it is closed (divergence-free) and its restriction to every boundary component has zero total flux. In this case, the integrability of curl ensures that the endpoints of open filaments landing on boundaries must give equal positive and negative ends per boundary component. Therefore, it is possible to pair the endpoints along the boundary and complete the filaments as closed curves, from which we know how to construct ψ .

5.4.2 Non-swirling dynamics

In Section 5.4.1, we leveraged the degrees of freedom in the complex level set function ψ to design a sufficiently regular implicit representation. Here, we exploit similar degrees of freedom to construct evolution equations that produce theoretically equivalent dynamics but are more numerically robust.

Many dynamical systems for curves already come with a known physical evolution equation. For example, the vortex filament dynamics can be simulated with (5.3.8), *i.e.* by advecting ψ using the Biot–Savart or the Rosenhead–Moore flow (Example 5.3.1). Hence, redesigning the equation may seem unnecessary. However, when it comes to numerically advancing the variables, the highly oscillatory or discontinuous nature of the Biot–Savart and the Rosenhead–Moore flows near the vortex core (as illustrated in Figure 5.6) can cause significant interpolation error. We point out that these errors are avoidable by redesigning the flow of the advection.

We consider dynamical systems as discussed in Section 5.3.2. Suppose the evolution of the curve is given by $\frac{\partial \gamma}{\partial t} = V_\gamma$ where V_γ is the velocity field defined on the curve. The evolution for ψ must be an advection by an extension v of the velocity field V_γ in a neighborhood of the curve (*cf.* (5.3.6)). A straightforward construction of v is a constant extrapolation. That is, $v(x)$ is set to V_γ at the closest point on γ from x . This extrapolation is, however, singular where closest points are not unique as in the middle row of Figure 5.6.

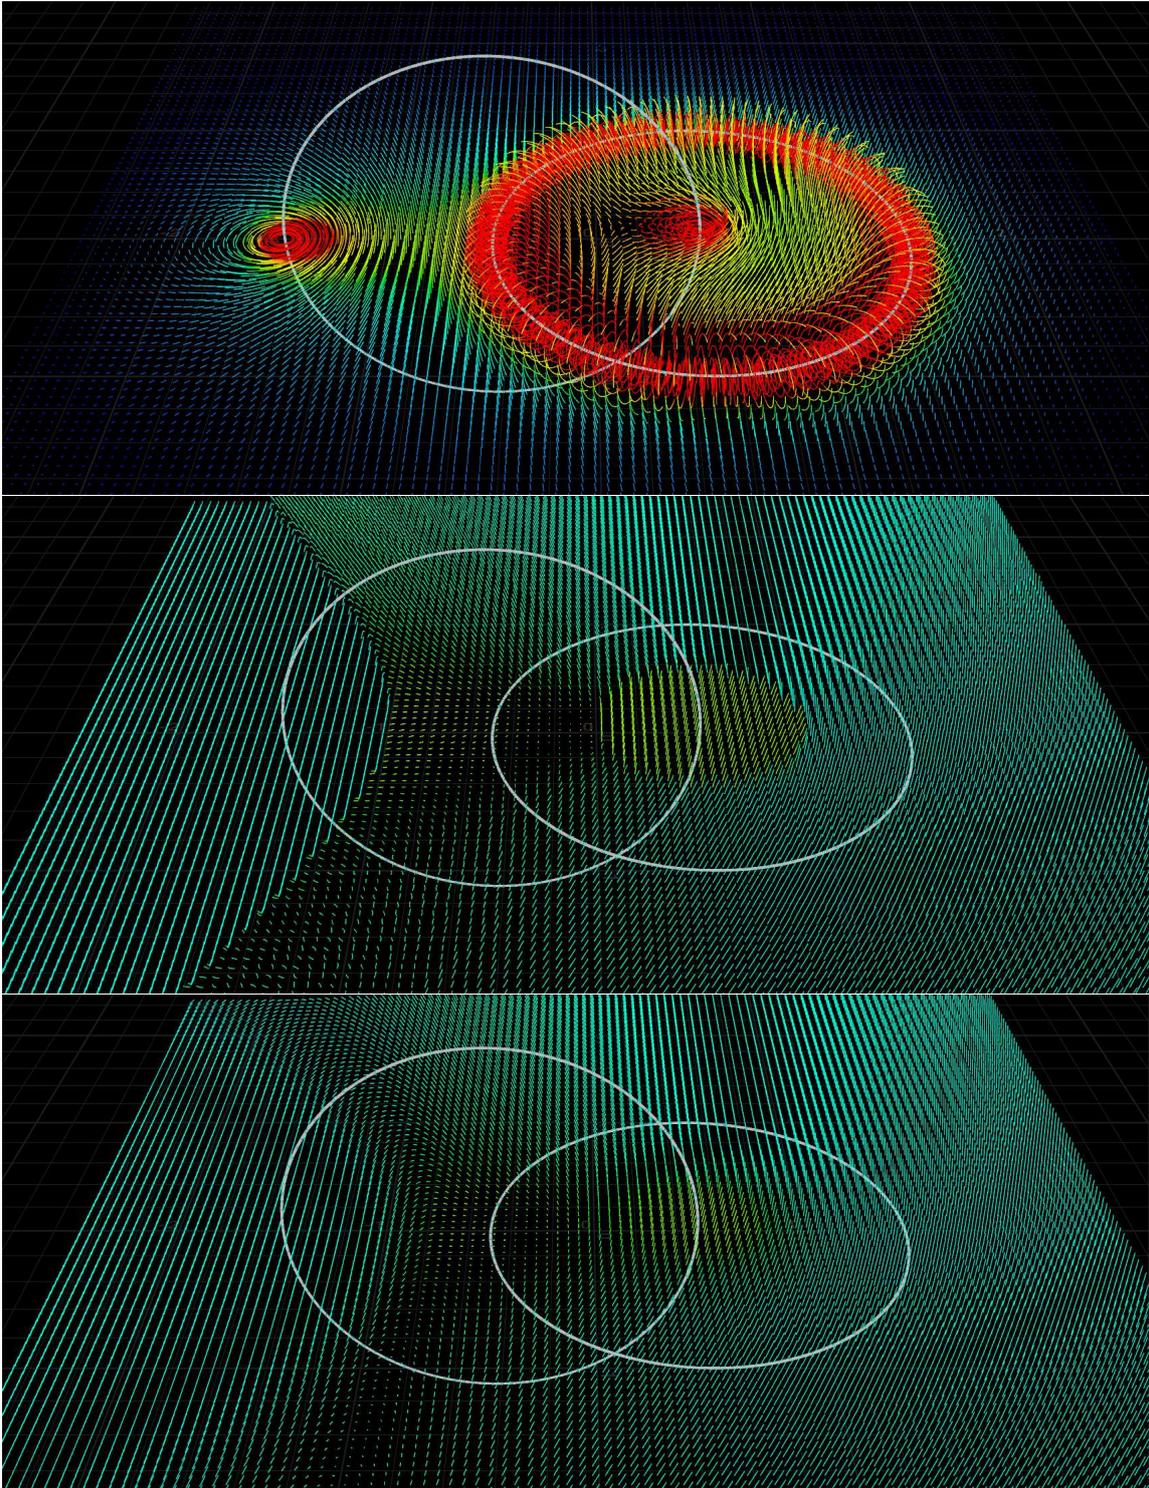


Figure 5.6: Different vector fields for two linked rings. The original Rosenhead–Moore model (top), the nearest point velocity field (middle), and a smooth weighted average field (bottom). While these three velocity fields coincide on the filaments, they differ significantly outside the filaments.

To gain continuity without changing the velocity on the filaments, we smooth away these singularities by taking the weighted average of the filament velocity as

$$v(x) := \frac{1}{\mathcal{N}(x)} \oint_{\gamma} V_{\gamma}(\gamma(s)) w(x, \gamma(s)) ds.$$

Here, $\mathcal{N}(x)$ is the normalization factor

$$\mathcal{N}(x) = \oint_{\gamma} w(x, \gamma(s)) ds,$$

and w is some weight function that applies less smoothing as it gets closer to the filament, i.e. $w(x, \gamma(s))/\mathcal{N}(x) \rightarrow \delta(\gamma^{-1}(x) - s)$ as $\text{dist}(x, \gamma) \rightarrow 0$. For example, we observed that a Gaussian function with distance-dependent variance works stably:

$$w(x, s) = \exp\left(-\frac{|x - \gamma(s)|^2}{\sigma^2 \text{dist}(x, \gamma)}\right), \quad \sigma \text{ some constant.}$$

To accelerate computation for v , we can further multiply the integrand of Figure 5.4.2 by a smooth cutoff function which equals to 1 near γ and 0 far away from γ . Then v is non-vanishing only near γ . By applying this smooth cutoff, we only need to evaluate v in a narrow band close to the filament.

Note that v is in general not divergence-free. The velocity near the curve is determined according to the curve velocities so that the motion of the zeros of ψ emulates the motion of the curves. Imposing an additional constraint like incompressibility to the velocity field may trade off the fidelity to the original curve motion or the numerical smoothness of the surrounding vector field.

5.5 Algorithms

In this section, we describe an algorithm for simulating filament dynamics.

Throughout the simulation, we maintain a complex level set function ψ . The main algorithm computes the transport equation of ψ along a velocity field v in a neighborhood of the zeros of ψ . This main algorithm is accompanied by a few subroutines for evaluating the velocity and redistancing: one subroutine extracts the zero set γ of ψ ; another subroutine constructs the solid-angle distance function ψ from γ (5.4.1); a third subroutine evaluates the filament motion V_{γ} using γ ; and the last subroutine extends V_{γ} to a velocity field v in a neighborhood of γ (Section 5.4.2).

We store the level set function ψ and the velocity v on a 3D lattice and discretize the curves γ as oriented collections of line segments.

5.5.1 Details of subroutines

1. Advection

To advect ψ with a given flow v , one can adopt any Eulerian advection scheme. In our implementation, we use the modified MacCormack method [SFK⁺08] with 4th order Runge–Kutta back-tracing.

Algorithm 5.1: The main time integration

Input: Initial filament $\gamma \in \mathcal{X}$;
 1 $\psi \leftarrow$ construct ψ from γ ; ▷ Subroutine 3
 2 **while** simulating **do**
 3 $v \leftarrow$ evaluate V_γ and extend it on grids near γ ; ▷ Subroutine 4
 4 $\psi \leftarrow$ advect ψ along v ; ▷ Subroutine 1
 5 $\gamma \leftarrow$ extract the zero set of ψ ; ▷ Algorithm 5.2
 6 $\psi \leftarrow$ construct ψ from γ ; ▷ Subroutine 3
 7 **end**

Algorithm 5.2: Extract the zero curve γ from ψ

1 **foreach** cell c **do**
 2 **foreach** face f in c **do**
 3 Compute incidence $n_f \in \{-1, 0, 1\}$; ▷ Eq. (5.5.1)
 4 **if** $|\vec{x}^{(k)} - \vec{x}^{(k-1)}| < \epsilon$ **then**
 5 Find $p_f^\pm = \psi^{-1}(0) \in \mathbb{R}^3$ in f ; ▷ bilinear interp.
 6 **end**
 7 **end**
 8 Connect p_f^- and p_g^+ of some faces f, g in c ;
 9 **end**

2. Construction of γ from ψ

After updating ψ , we need to update the filaments γ by extracting $\{\psi = 0\}$. We summarize this subroutine in Algorithm 5.2, which is adopted from [WPS14].

In our setting, each vertex of γ lives on a face f of the volumetric grid. We first evaluate for each face f the $\{-1, 0, +1\}$ -valued *signed intersection* n_f with the zero curve of ψ using the *argument principle*: If the vertices of a face f are i, j, k, ℓ in an oriented order, then

$$n_f = \frac{1}{2\pi} \left(\arg\left(\frac{\psi_j}{\psi_i}\right) + \arg\left(\frac{\psi_k}{\psi_j}\right) + \arg\left(\frac{\psi_\ell}{\psi_k}\right) + \arg\left(\frac{\psi_i}{\psi_\ell}\right) \right) \quad (5.5.1)$$

using the principal branch $-\pi < \arg(\cdot) \leq \pi$. Geometrically, (5.5.1) describes how many times the quadrilateral $\psi_i, \psi_j, \psi_k, \psi_\ell \in \mathbb{C}$ winds around the origin. For each face where $n_f \neq 0$ we evaluate the more precise location of the zero using a bilinear interpolation. That is, we regard $f = [0, 1]^2$ by scaling and ψ is bilinearly interpolated as

$$\begin{aligned} \psi_f(x, y) := & (1-x)(1-y)\psi(0, 0) + x(1-y)\psi(1, 0) \\ & + (1-x)y\psi(0, 1) + xy\psi(1, 1), \end{aligned}$$

and $\psi_f : f \rightarrow \mathbb{C}$ has the inverse when $n_f = \pm 1$. The location $\psi_f^{-1}(0)$ is a vertex p_f of the curve γ . Finally, we build the edges of γ by running over the grid cells where we connect the pairs of zeros on the face with a consistent orientation. Each cube c may have up to two pairs of zeros with positive and negative n_f . A cube with two pairs of vertices has an ambiguity similarly to the marching cube algorithm [LC87]. We resolve this ambiguity by connecting vertices arbitrarily in a way that preserves curve orientation. We have not investigated higher order algorithms to connect curves more accurately at the sub-grid scale.

3. Construction of ψ from γ

Given γ , we construct ψ as the solid-angle distance function (5.4.1). For each grid point x near γ , we evaluate $|\psi(x)| = \text{dist}(\gamma, x)$ as the distance to the closest polygon edge. We evaluate $\arg(\psi(x)) = \theta(x) = \frac{1}{2} \text{SolidAngle}(\gamma; x)$ by computing the signed area of a spherical polygon with vertices $\{p_i\}_{i=1}^N$ that are the discrete points of γ projected onto the unit sphere centered at x . To ensure differentiability and the local isometry properties of $d\psi|_{\gamma^\perp}$, the sign of θ must depend on curve orientation: we compute the *signed* area by introducing a pole $Z := (1, 0, 0)$:

$$\text{SignedArea}(\{p_i\}_{i=1}^N) = \sum_{i=1}^N \text{sign}((p_i \times p_{i+1}) \cdot Z) \text{Area}(p_i, p_{i+1}, Z)$$

where the *unsigned* area of each spherical triangle $\{q_1, q_2, q_3\}$ is computed using a standard area formula [BC87, Arv95a],

$$\text{Area}(q_1, q_2, q_3) = -\pi + \sum_{i=1}^3 \arccos\left(\frac{(q_{i-1} \times q_i) \cdot (q_i \times q_{i+1})}{\|q_{i-1} \times q_i\| \|q_i \times q_{i+1}\|}\right).$$

These equations assume cyclic vertex indexing, so $p_{N+1} = p_1$, $q_{3+1} = q_1$, and $q_{1-1} = q_3$. Alternatively, we can use the area formula for spherical polygons introduced in Chapter 7, which is robust even when the polygon $\{p_i\}_{i=0}^N$ has degenerate points.

4. Extending velocity to grid

In order to advect ψ , we need to extend the velocity field defined on γ to the grid points near the curves. To produce non-swirling dynamics, we used a smooth average field (Equation 5.4.2, 5.4.2, and 5.4.2) for the examples in this chapter, unless otherwise explained. In Figure 5.4.2, we observed that σ ranging from 0.1 to 10 times the grid size works stably without smoothing out detailed dynamics.

5.6 Applications

This section discusses applications of our approach, specifically applied to vortex filament dynamics. We implemented our algorithms on Houdini 18.5.759 and ran all simulations on a MacBook Pro (13-inch, 2020) with a 2.3 GHz Quad-Core Intel Core i7 processor. For an example implementation, see <https://github.com/sdsgisd/ImplicitVortexFilaments>

We use our algorithm to animate two ‘‘leapfrogging’’ vortex rings in Figure 5.7 and our accompanying video. We note that the system remains stable and highly symmetric even at the end of a long simulation with several high-speed ring interactions. Figure 5.8 visualizes two vortex rings colliding with one another at right angles, reconnecting, and detaching into two new rings. We note that the final rings retain plenty of energy after the collision event, in contrast to Eulerian simulations of this phenomenon which tend to damp out over time. The visual detail in our simulations is also practically independent of grid resolution, as the motion of marker particles are described analytically by velocities (Equation 5.3.5), instead of a vector field stored on a coarse grid.

Figure 5.9 illustrates a jet of smoke created by generating a new smoke ring at the left side of the domain every three time steps. For transporting smoke as a scalar field stored on grid points, we again used Equation 5.3.5. This simulation shows the robustness of our

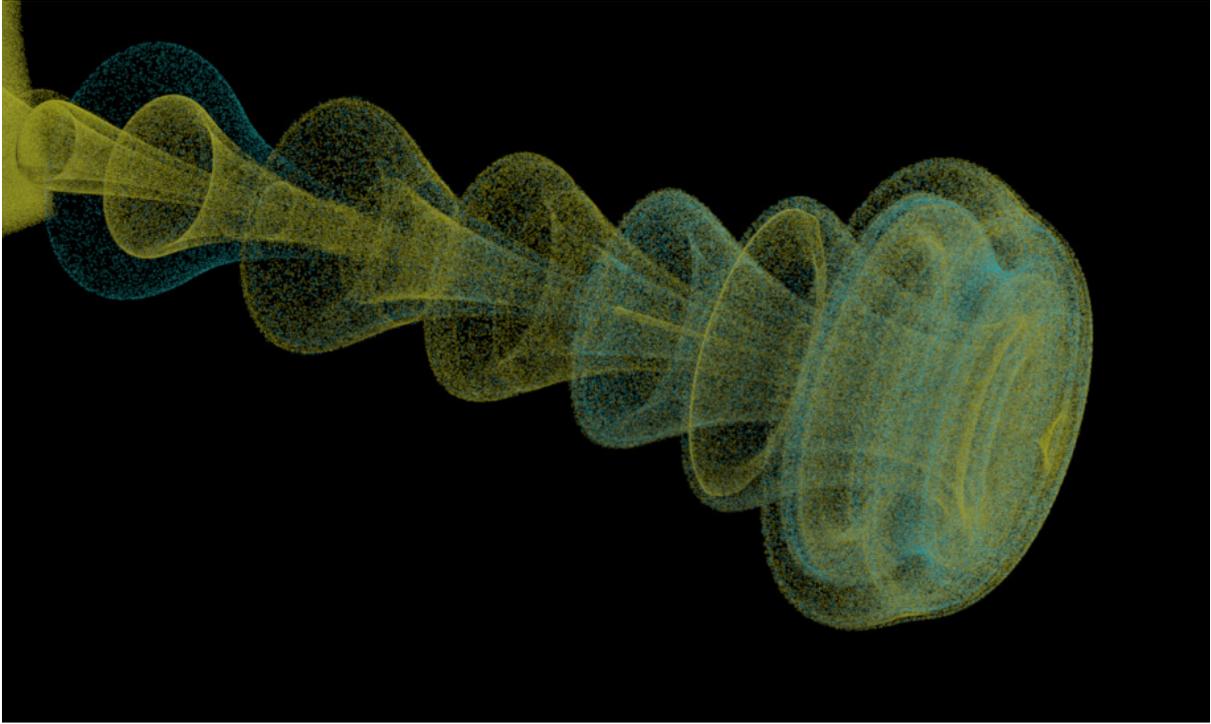


Figure 5.7: Two filaments leapfrog through one another, dragging marker particles into the shape of a mushroom cloud. The initial filament geometry consists of two co-planar vortex rings, one with half the radius of the other.

topology changes: each re-connection event is the result of a curve extraction from a level set function, so there is no possibility of any unexpected edge cases, and no need for any geometric intersection code. The simple set-up creates a large variety of chaotic motions resulting from fast leap-frogging rings squeezing in between others and reconnecting filaments causing sudden changes in direction. When rings shrink smaller than the grid resolution, our algorithm deletes them (similar to codimension-1 level set methods).

Lastly, Figure 5.1 illustrates how our method can evolve intricate filament geometry, specifically the (5,8)-torus knot defined by

$$\gamma(s) = ((\cos(qs) + 2) \cos(ps), (\cos(qs) + 2) \sin(ps), -\sin(qs))$$

with $(p, q) = (5, 8)$ and $s \in [0, 2\pi)$.

5.6.1 Influence of numerical parameters

Figure 5.10 demonstrates the importance of each step in our approach by selectively removing or modifying different algorithmic components and illustrating the consequences. First we illustrate what happens when we vary the free degrees of freedom in the velocity field v used to advect the filaments. In agreement with the discussion in Section 5.3.2, we see that setting v to the fluid velocity (based on Biot-Savart kernels) field causes the level set function ψ to rapidly twist up and become unstable. Setting v to equal the velocity at the nearest point on the filament creates similar noise, presumably due to spatial discontinuities in the field. Compare these results to the smooth geometry generated by our velocity field at the bottom of Figure 5.10.

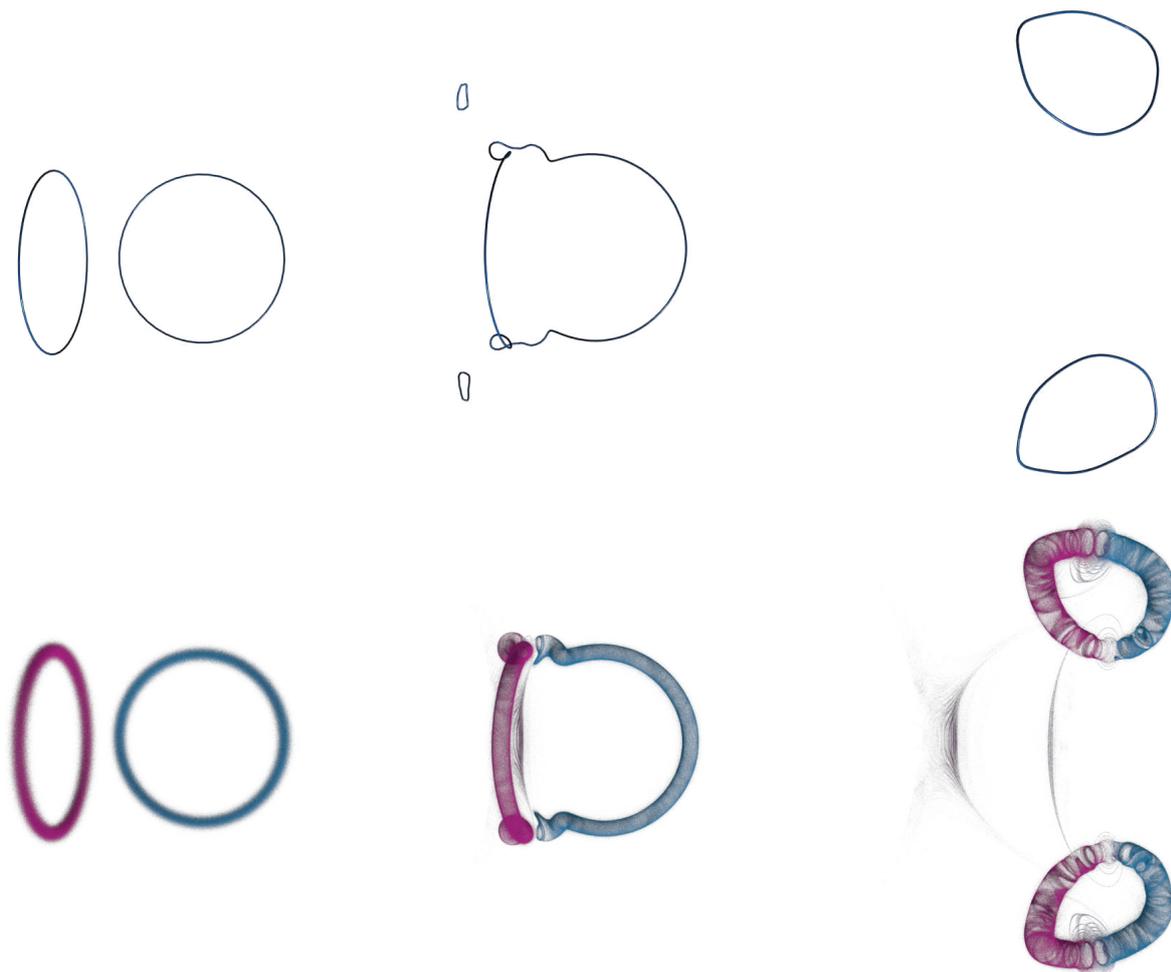


Figure 5.8: Two smoke rings (left) colliding at orthogonal angles (middle) and re-connecting (right). Rendered as filaments (top) and marker particles (bottom). The colliding rings leave swirly trails of smoke particles after their collision and reconnection.

Lower down in the same figure, we illustrate the effect of varying the free degrees of freedom in our level set function ψ . Starting with an initially smooth ψ and advecting it without any re-distancing or regularization works well at the beginning, but it eventually accumulates topological noise. To illustrate the impact of ψ 's free parameters on numerical stability and accuracy, the fourth row in Figure 5.10 replaces our smooth choice of ψ with one that is intentionally twisted by a phase shift of $\Delta\theta(x) := 0.05 \text{ dist}(x, \gamma)$; the twisted ψ causes high-frequency geometric noise and artificially shrinks the filaments. Again, we can compare these results to the smooth geometry generated by our un-twisted ψ at the bottom of Figure 5.10.

Another important numerical parameter is the spatial resolution of the filament. Lagrangian methods constrain the curve resolution by subdividing and collapsing edges when they become too long or short. In contrast, our method controls the curve details via the resolution of the grid used for the level set ψ . Figure 5.11 illustrates a simulation of an evolving trefoil knot on a $50 \times 50 \times 60$ grid and one twice as detailed at $100 \times 100 \times 120$. As expected, higher resolution grids create filaments with sharper details. Note, however, that the common way to visualize fluids is with marker particles or smoke densities, not by visualizing the filaments themselves. Thus, even very low resolution vortex filaments can still produce high resolution visual details. More detailed filaments make themselves evident via their more detailed velocity

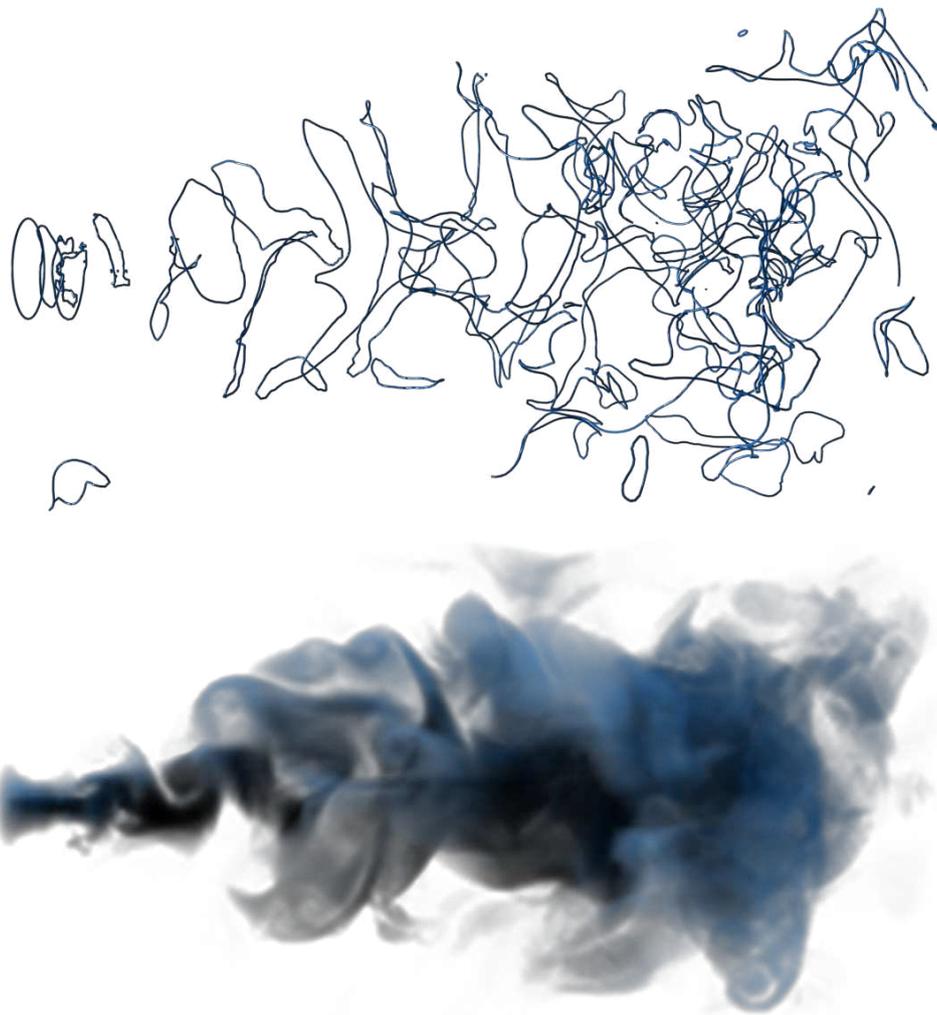


Figure 5.9: A jet of smoke, rendered as raw filament geometry (top) and an advected smoke density function (bottom).

fields and complex smoke dynamics.

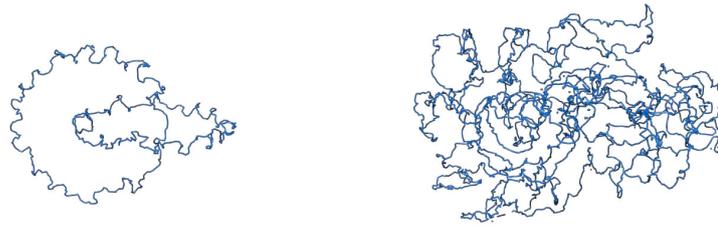
5.6.2 Comparisons with Lagrangian filaments

Next, we qualitatively compare implicit and explicit representations of curve dynamics using our algorithm and the Lagrangian vortex filament technique of Weißmann and Pinkall [WP10], as implemented in Houdini software by SideFX. Figure 5.12 shows the evolution of a knotted vortex filament with both methods. The filament is initialized as

$$\gamma(s) = (\sin(s) + 2 \sin(2s), \cos(s) - 2 \cos(2s), -\sin(3s))$$

with $s \in [0, 2\pi)$. Aside from some small differences arising from the particulars of how filaments break apart and reconnect, the two methods produce roughly the same dynamics.

On the other hand, the two methods have significantly different mechanisms for handling topological changes, which can produce divergent results. The Lagrangian method depends on user parameters like the thresholds for distances and angles between curves; the only relevant user parameter for changing topology in our method is the grid spacing, which prescribes



Advection along the Biot-Savart velocity field



Advection along the nearest point velocity field

Without redistancing ψ Artificially twisted ψ 

Our proposed setting

Figure 5.10: Comparisons with different settings. Initial state is the 1st row of Figure 5.13. The left and right columns show states at frames 20 and 45.

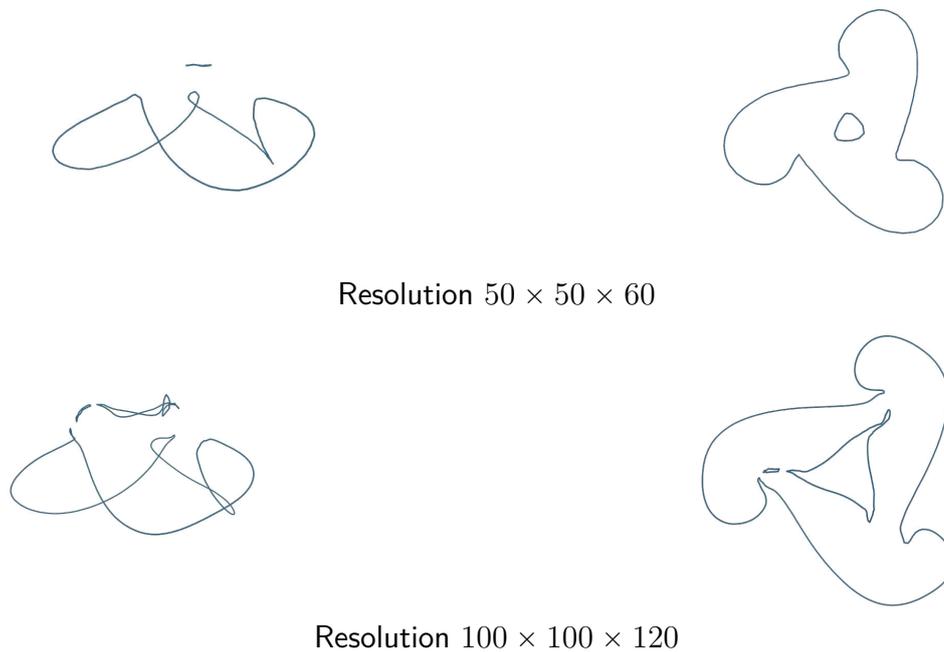


Figure 5.11: Simulation of a trefoil knot with different grid resolutions, as viewed from the side (left column) and front (right column).

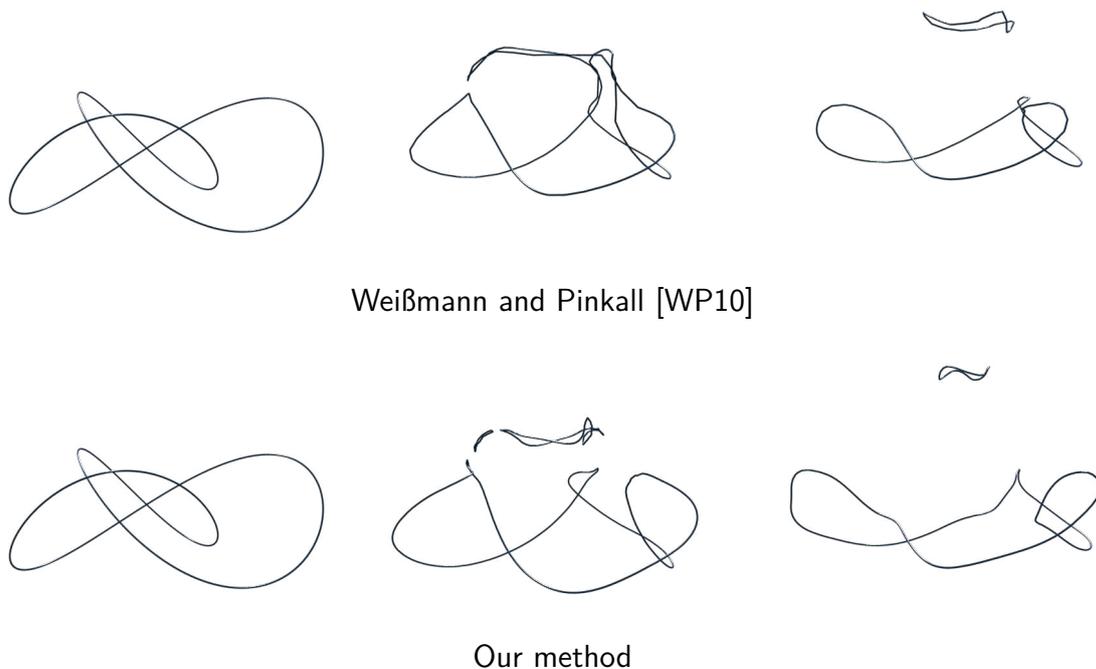


Figure 5.12: Comparing a buoyant trefoil knot simulation by Weißmann and Pinkall [WP10] to ours. The simulations evolve from left to right.

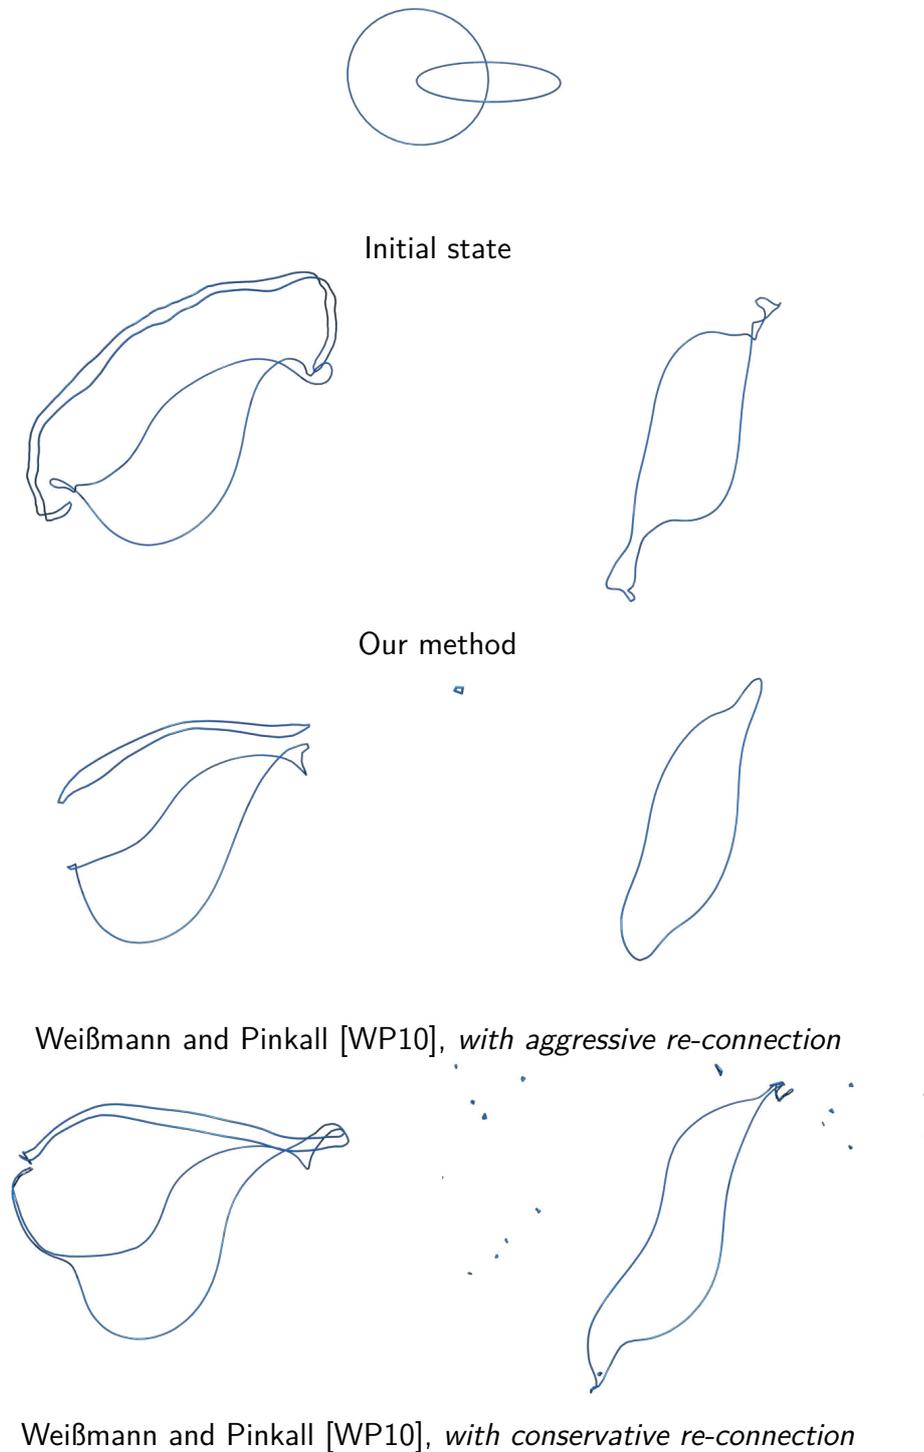


Figure 5.13: Simulating two linked vortex rings (top) with our method (2nd panel), and with the explicit filament approach of Weißmann and Pinkall [WP10] (bottom two panels). Lagrangian methods can be sensitive to numerical parameters for topological changes.

the resolution of the level set ψ . Figure 5.13 shows a simulation of two linked rings: our approach both preserves long thin tendrils and filters out topological noise. The Lagrangian simulation is sensitive to re-connection parameters: setting these parameters too aggressively leads to smooth geometry but loses thin features, while setting parameters too conservatively preserves thin features but creates noisy, persistent, high-speed “ringlets” that dominate the fluid velocity field.

We stress here that the purpose of this comparison is to show how explicit and implicit descriptions handle topological changes of curves result in qualitatively different ways; the accuracy of these methods are not easily comparable as they have different mechanisms, and the accuracy depends on the type of curve dynamics.

5.6.3 Obstacles

Like many methods for vortex dynamics, our method can also make filaments circumvent obstacles. A typical approach is to find a smooth harmonic potential ϕ such that for a given obstacle $B \subset M$, it solves

$$\begin{aligned} \langle v_\gamma - v_{\partial B} - \nabla\phi, n \rangle_{\mathbb{R}^3} &= 0 \text{ on } \partial B \\ \phi &\rightarrow 0 \text{ at infinity} \end{aligned}$$

where $v_{\partial B}$ is the boundary velocity at each point of ∂B . Weißmann and Pinkall [WP12] construct such a potential by regarding points inside B as sources of localized potentials so their weighted sum solves Section 5.6.3, and Brochu et al. [BKB12] and Zhang et al. [ZB14] solve a similar system using boundary element techniques. Nabizadeh et al. [NCR21] address an equivalent problem by solving linear PDEs on infinite domains using the Kelvin transform.

Another approach by Weißmann and Pinkall [WP10] is more specialized to vortex filaments and phenomena like vortex shedding; it regards the obstacle as a collection of artificial filaments γ_M such that the normal component of the velocity is zero. Other approaches, like that of Park and Kim [PK05] and Da et al. [DBWG15] add point constraints to the boundary which zero out both the normal and tangential velocity components.

Our vortex filament algorithm is compatible with any of these obstacle-handling methods; our particular implementation uses [NCR21], as seen in Figure 5.14. We can observe that filaments near the obstacle are accelerated due to the induced mirrored image of themselves in the obstacle (or vortex sheet on the surface).

Note that for the evolution of curves, only the evaluation of the velocity field needs the treatment of the boundary. Other components of the algorithm including the construction of levelset function ignore the boundary and do not require any additional treatments.

5.6.4 Other types of filament dynamics

Our idea of implicit representation of filament dynamics is not limited to vortex filaments. In theory, this should be applicable to any first-order time evolution of curves. Figure 5.15 shows an example with the curve-shortening flow,

$$\frac{\partial \gamma}{\partial t}(s, t) = -\gamma''(s, t).$$

As expected, this velocity causes the filament to shrink over time, similar to mean-curvature flow for surfaces [OS88].

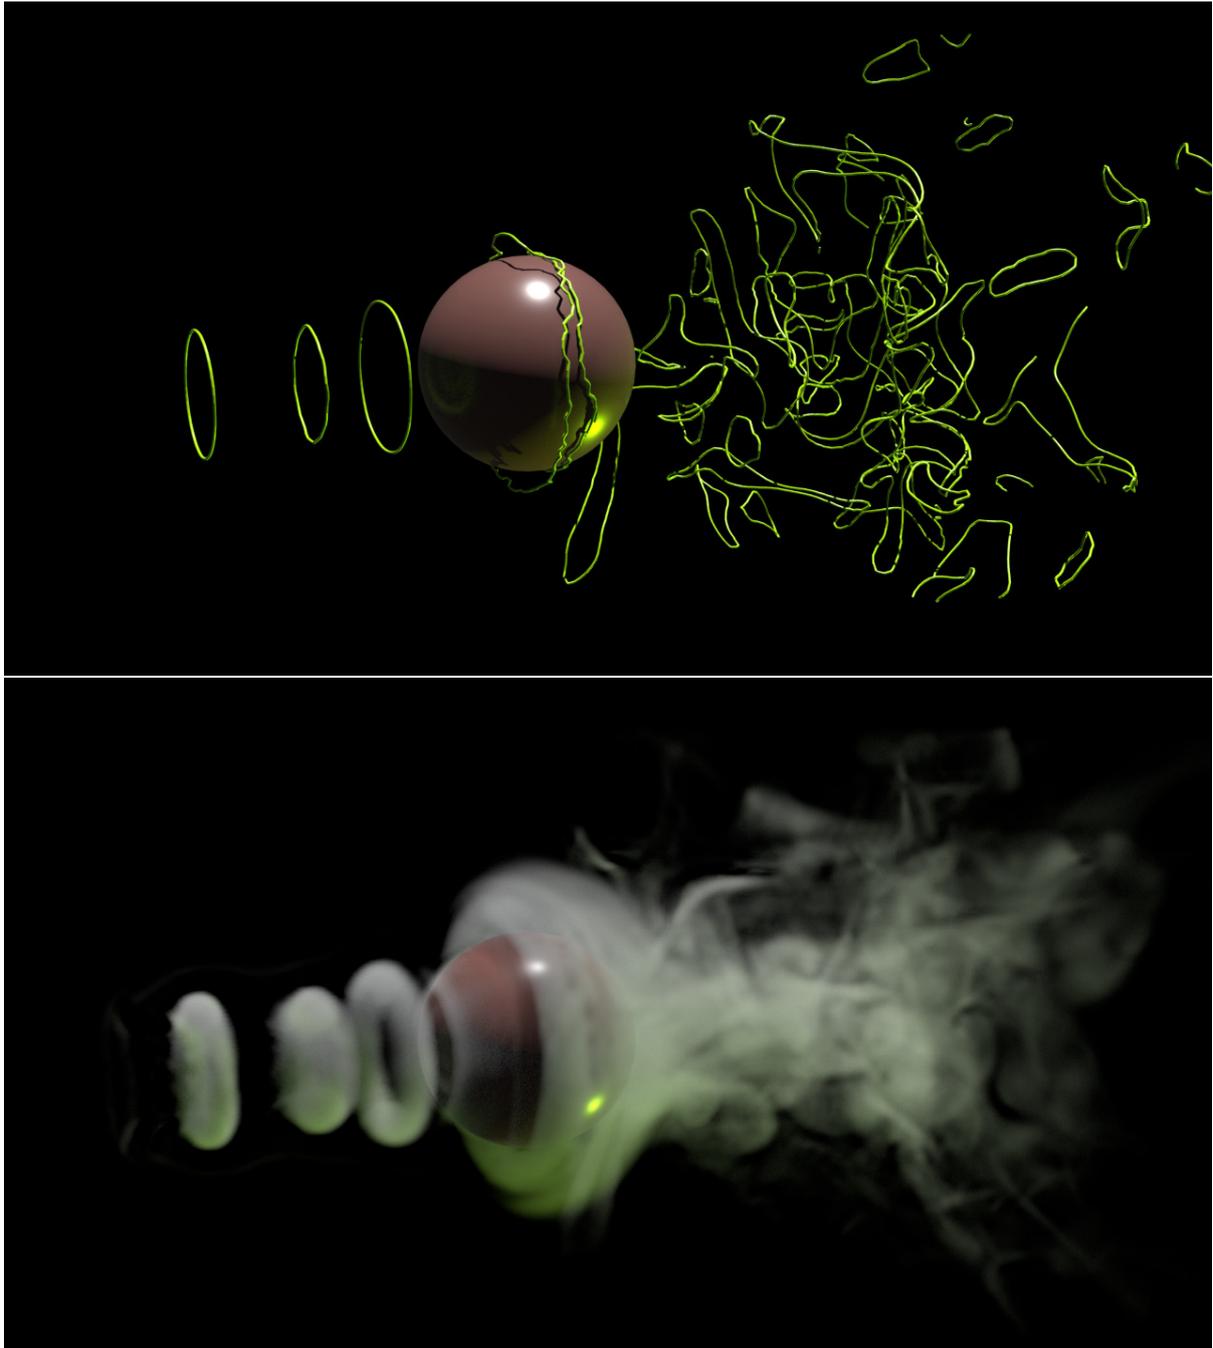


Figure 5.14: Turbulence caused by a spherical obstacle. We generate a new vortex filament ring every 10 frames.

5.6.5 Discussion

Our method represents a fundamental new way to animate fluids and vortex filament dynamics. Its unique features give it some strengths and weaknesses relative to existing approaches.

First of all, we represent our filaments with implicit functions ψ . While there are many ways to encode an implicit function, our implementation uses a regular grid, which implies a finite bounding box. While common for Eulerian fluid simulations, bounded domains are a constraint not shared by Lagrangian methods. This constraint could be mitigated if we use sparse grids or trees [MAB19].



Figure 5.15: Curves evolving under the curve shortening flow, from left to right.

Table 5.1: Parameters and timing breakdown per frame for all simulations in this chapter and the accompanying video. Symbols Γ and a are the intensity and the thickness of filaments. All simulations are 24fps and the time step size $1.0/24s$ except for Figure 5.15, which we used 120fps. Average timings are taken over the entire simulation. The “Other” column includes the remaining operations, including obstacle handling and generation of new curves at the sources in Figure 5.9 and Figure 5.14.

Scene	Resolution	Γ	a	Total time	Evaluate v	Advect ψ	Construct γ	Construct ψ	Other
Figure 5.13	$100 \times 100 \times 100$	1.0	0.05	0.159s	0.041s	0.017s	0.041s	0.056s	0.002s
Figure 5.12	$100 \times 100 \times 120$	1.0	0.05	0.274s	0.056s	0.026s	0.049s	0.102s	0.001s
Figure 5.7	$60 \times 60 \times 200$	2.0	0.08	0.090s	0.013s	0.021s	0.023s	0.031s	0.001s
Figure 5.8	$80 \times 120 \times 140$	1.0	0.05	0.152s	0.031s	0.030s	0.046s	0.044s	0.001s
Figure 5.1	$100 \times 100 \times 140$	1.0	0.05	0.655s	0.144s	0.085s	0.091s	0.332s	0.002s
Figure 5.9	$100 \times 100 \times 180$	2.0	0.08	1.423s	0.177s	0.099s	0.812s	0.324s	0.011s
Figure 5.14	$100 \times 100 \times 160$	2.0	0.08	0.969s	0.047s	0.087s	0.473s	0.226s	0.135s
Figure 5.15	$100 \times 100 \times 100$	N/A	N/A	0.152s	0.028s	0.012s	0.068s	0.043s	0.001s

Table 5.2: Computational timings compared with Houdini’s built-in implementation of Weissmann & Pinkall [WP10]. “Same DOF” refers to simulations with approximately the same number of computational degrees of freedom as our method: we set the relevant parameters (re-connection distance, minimum and maximum edge lengths) so that the number of explicit curve vertices are similar to ours.

Scene description	W&P	W&P (Same DOF)	Ours
Linked rings (Figure 5.13)	0.052s	0.065s	0.159s
Trefoil knot (Figure 5.12)	0.045s	0.101s	0.274s

The main parameter in our method is the grid resolution. As discussed earlier, this parameter affects the geometric detail and topology of our filaments, influencing the velocity field directly, but influencing the final visual results only indirectly. The grid resolution also directly influences topological changes; the only mechanism for topological changes in our algorithm is to merge curves when they intersect the same grid cell. We believe this automatic and robust method for handling topological changes is a strength of our method; it minimizes the need to fine-tune parameters (especially the relationship between minimum/maximum edge length and topological change interaction lengths) and seems immune to the types of numerical blow-ups that we have seen in Lagrangian methods when filaments get close or exhibit near-degenerate geometry.

On the other hand, this grid-based method for re-connecting curves will also delete small features when they shrink below the grid cell size (similar to level set methods for surfaces). This behavior is most evident in our jet example (Figure 5.9). We believe that the deletion of small features can be reduced in the future in a number of ways. More accurate advection schemes will probably preserve higher frequency features of ψ better without deleting them. Also, although we took care to introduce a v and ψ which alleviates egregious numerical stability problems, we have not quite optimized for accuracy or geometric durability for the

specific dynamics of vortex filaments. Finding a pair of v and ψ that is more suitable for each curve dynamics in the degrees of freedom discussed in Section 5.3 would lead to even better numerical performance.

Although our method relies on Eulerian advection to evolve the geometry, it *does not* suffer from the artificial viscosity typically associated with Eulerian fluid simulations. Our fluid velocity is reconstructed from filament dynamics, so the velocity field is not recursively re-sampled and does not accumulate damping errors. Consequently, our method produces swirly and energetic fluid flows even at low grid resolutions.

Table 5.1 lists the simulation parameters and timing breakdown for each of the simulations in this chapter and accompanying video. We stress that our prototype implementation is meant as a proof of concept, and it has plenty of room for optimization. Our current implementation employs regular grids for ease of implementation; future implementations can make great use of sparse grids, since our filaments only use a 1-dimensional path through the 3D grid. Our implementation of “Evaluate v ” and “Construct ψ ” iterates over the entire curve geometry for each point in the narrow band where γ and ψ are required; future implementations could use a sparse data structure and approximated fast summation via tree-codes like the fast multipole method to reduce total evaluation time. Additionally, our current implementation of “Construct γ ” redundantly doubles the work per curve vertex, so that step can be sped up by at least a factor of 2.

With these inefficiencies in mind, our implementation appears to run modestly slower than Houdini’s optimized implementation of [WP10] for the scenarios we tested. The two methods have completely different numerical degrees of freedom, so we find it difficult to compare them directly. Table 5.2 compares the computational cost of the two methods for figures in this chapter. We aimed to keep the number of curve vertices roughly the same as our method’s in the “Same DOF” simulations, so we believe these are the most relevant for comparing timings.

In our examples of vortex filaments, we assumed inviscid filaments, but we can handle viscous motions too once the time-derivative of viscous curves is specified. We also assumed a uniform vortex strength. Just as in the explicit description, our implicit description requires a well-defined equation of motion for curves, and it outputs only new curve configurations without processing additional quantities like vortex strengths. Simulating vortex filaments with different strengths is an independent challenge. Merging filaments with different strengths would create a graph of filaments rather than just disjoint closed curves, which would require a new implicit representation for filament graphs.

5.7 Conclusion and future work

We have shown that implicit representations of geometric curves exhibit large degrees of freedom in both their mathematical representation, as well as their dynamics. We then took advantage of these redundant degrees of freedom to improve the stability of vortex filament simulations.

We see a number of avenues that can be explored in future work. Our current strategy exploits redundancy in the 3D velocity field and level set outside the filaments, but we can also incorporate free degrees of freedom in the tangential components of the velocity on the filament itself. More generally, one can attempt to formalize the regularization of v and ψ for various purposes. A possible instance is an optimization problem for certain energies aiming provable guarantees on numerical accuracy and stability. Besides, the use of redundancy in

dynamics should also be possible for other codimensional cases such as level surfaces in 3D or level sets of an arbitrary codimension in a higher dimensional space.

Our numerical scheme can be made more sophisticated as well: higher order advection schemes and geometric curve representations, as well as sparse and adaptive grids can make our method both more efficient and more numerically accurate.

Finally, this chapter explores vortex filaments and curve shortening flows, but our ideas are not limited to these specific dynamical systems. We expect that the idea of exploiting hidden degrees of freedom in implicitly represented curve dynamics will generalize to many more types of dynamics appearing in both scientific fields and engineering applications.

Symplectic structures on the space of space curves

The content of this chapter is the article: Martin Bauer, Sadashige Ishida, and Peter W. Michor. Symplectic structures on the space of space curves. *Preprint. arXiv:2407.19908*, 2024.

In the previous chapters, we studied symplectic geometry of codimension-2 shape space with respect to the Marsden–Weinstein (MW) structure. To the best of the author’s knowledge, MW structure was the only symplectic structure previously found and studied on this space.

In this chapter, we focus on the case of base dimension 3, namely the space of space curves. We introduce new symplectic structures on this space, which generalize the MW structure.

Our method integrates the Liouville 1-form of the MW structure with Riemannian structures that have been introduced in mathematical shape analysis. We also derive Hamiltonian vector fields for several classical Hamiltonian functions with respect to these new symplectic structures, leading to new Hamiltonian dynamics.

Dependency of the chapter This chapter is self-contained, as it reintroduces key concepts appearing in previous chapters. For a deeper understanding of the general background on the symplectic geometry of the codimension-2 shape space, readers may find Chapter 2 useful.

6.1 Introduction

Motivation and background The space of unparametrized space curves

$$\text{UImm}(\mathbb{S}^1, \mathbb{R}^3) := \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) / \text{Diff}^+(\mathbb{S}^1)$$

as an infinite dimensional orbifold is known to have a symplectic structure called the Marsden–Weinstein structure (MW-structure) [MW83]. Here $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ denotes the space of immersions of \mathbb{S}^1 into \mathbb{R}^3 and $\text{Diff}^+(\mathbb{S}^1)$ is the orientation preserving diffeomorphisms acting on immersions as reparametrizations.

The MW structure is thought of as a *canonical* symplectic structure as it is formally a Kirillov–Kostant–Souriau form by regarding space curves as linear functionals on the space of divergence-free vector fields in \mathbb{R}^3 ; see eg. [MW83, Theorem 4.2] and [AK21, Chapter VI,

Proposition 3.6]. Another incentive for studying the MW symplectic structure can be found in its appearance in mathematical fluid dynamics: for example, one can interpret vortex filaments as the MW flow of the kinetic energy of the velocity field induced by vorticity concentrated on the curve. Via so-called localized induction approximation vortex filaments reduce to the binormal flow, which is a completely integrable system and is again an MW flow for the length functional as the Hamiltonian, see eg. [Saf93, Chapter 11] or [MB01, Chapter 7] and the references therein.

To the best of the authors' knowledge, to date no symplectic structures other than the MW form have been studied on the space of unparametrized space curves. Riemannian structures on this space, on the other hand, have attracted a significant amount of interest; primarily due to their relevance to mathematical shape analysis [You10, SK16, BBM14]. The arguably most natural such metric, the reparametrization invariant L^2 -metric admits a surprising degeneracy: the geodesic distance between any pair of curves vanishes on both the space of parametrized and unparametrized curves [MM05, BBHM12]. This result renders the L^2 -metric unsuited as a basis for mathematical shape analysis and thus started a quest for stronger Riemannian metrics, which induce a non-degenerate distance function and consequently can be used for applications in these areas, see eg. [MM06, YM05, SKJJ10, BBHMA17] and the references therein. In these previous works, a typical approach to define a new Riemannian metric is to incorporate a suitable operator L into the aforementioned standard L^2 -metric G .

On the other hand, the MW-form $\bar{\Omega}^{\text{MW}}$ has a Liouville form $\bar{\Theta}^{\text{MW}}$ (*i.e.*, 1-form Θ s.t. $-d\Theta = \bar{\Omega}$) as we saw in Section 2.3.2.¹ Our key observation is that both the MW form $\bar{\Omega}^{\text{MW}}$ and the Liouville form $\bar{\Theta}^{\text{MW}}$ are related to the L^2 -metric G via an almost complex structure, induced on the shape space by the cross product with the unit tangent vector of the curve c , *i.e.*, $J_c(h) := \frac{c\theta}{|c\theta|} \times h$; here $c : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ is a space curve and $h : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ is a tangent vector to c . We have,

$$\bar{\Omega}^{\text{MW}}_{\bar{c}}(\bar{h}, \bar{k}) = G_c(J_c(h), k), \quad \bar{\Theta}^{\text{MW}}_{\bar{c}}(\bar{h}) := -\frac{1}{3}G_c(J_c(c), h)$$

where \bar{c} is an element and \bar{h}, \bar{k} are tangent vectors on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ related to c and $h, k \in T_c\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ by the projection from $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ to $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$.

Main contributions These relations between Riemannian geometry and symplectic geometry on the space of space curves are the starting point of the present chapter: our principal goal is to construct new symplectic structures on the space of unparametrized curves by combining the above classical construction with more recent advances in Riemannian geometry of these spaces, *i.e.*, we construct new presymplectic structures by modifying the Liouville form of the MW form using different Riemannian metrics from mathematical shape analysis. This construction automatically leads to a closed 2-form (and thus a presymplectic form) on the space of parametrized curves. Under certain assumptions on the Riemannian metric this form then descends to a presymplectic structure on the space of unparametrized space curves and thus it only remains to check the non-degeneracy of this 2-form to conclude that it is (weakly) symplectic. Proving this property turns out to be surprisingly difficult and provides the main technical contribution of the present chapter. Interestingly, in some cases the presymplectic

¹In Section 2.3.2, we used the sign convention $\Omega = d\Theta$. In this chapter we flip the sign for computational convenience. Additionally we change the notation for the Liouville form from η used in Chapter 2 to Θ so that both the MW form and the Liouville form are written with capital letters. But we emphasize that the Liouville form Θ in this section is different from the Liouville form in the prequantum sense introduced in Chapter 4 although they are both denoted by Θ .

form still has a nontrivial kernel on the shape space, but becomes symplectic when the quotient by a further 2-dimensional foliation is taken.

We also derive formulae for Hamiltonian vector fields of several classical Hamiltonian functions generated by our new symplectic structures and provide numerical illustrations to qualitatively show a few simple examples among these new Hamiltonian flows. The Riemannian counterparts can be found in the area of geometric gradient flows on the space of curves, where the investigations of gradient flows for certain well-known Energy functionals (e.g. the entropy or length energy functional) for Riemannian metrics other than the L^2 -Riemannian metric has been recently initiated, see e.g. the work of Okabe, Schrader, Wheeler and Wheeler [OSWW23, SWW23]. In our investigations we observe that, for certain choices of Hamiltonian and symplectic structure, we obtain a new representation of well-known Hamiltonian flows, *i.e.*, we may reproduce Hamiltonian flows of the MW symplectic structure from a different pair of a symplectic structure and a Hamiltonian function. For other examples, we obtain genuinely new Hamiltonian flows, which do not seem to be represented as a Hamiltonian flow for the MW symplectic structure.

A seemingly more straightforward approach to obtain new symplectic structures is directly defining a new skew-symmetric 2-form via alternating the Riemannian metric and combining it with the almost complex structure \mathcal{J} from the MW symplectic structure. This approach turns out to be unsuccessful as the resulting skew-symmetric 2-form is usually not closed and thus not even presymplectic. We discuss this approach and the resulting 2-forms in the appendix of this chapter (Section 6.7). This highlights the non-trivial challenge of finding a 2-form which is both closed and non-degenerate, rather than one being merely closed or non-degenerate. This observation was our original incentive to follow the slightly more complicated procedure described above.

Future directions In this chapter, we introduced new (pre)symplectic structures on the shape space of space curves. Our procedure of modifying the Liouville form of a (pre)symplectic form and taking the exterior derivative is not limited to such shape spaces. It would be interesting to apply the same machinery for other infinite-dimensional (weak-)symplectic manifolds that admit Liouville forms such as the space of complex functions on a domain or the cotangent bundle of an infinite-dimensional Riemannian manifold.

At this point, the connection between new Hamiltonian flows and existing physical theories seems to remain unclear. Hence we are also keen to use this new framework to find new interpretations of physically relevant quantities as Hamiltonian flows, in a similar way that certain compressible fluids are modeled as geodesic flows of higher order metrics [MM13].

Organization of the chapter In Section 6.2, we introduce Liouville forms via the modification of the L^2 -Riemannian metric, and then compute presymplectic forms by taking the exterior derivative. In Section 6.3, we show that, a class of presymplectic structures attained by conformal factors on the shape space are indeed weakly symplectic. We also derive Hamiltonian vector fields with respect to these weak symplectic structures. In Section 6.4, we describe more concretely symplectic structures induced by the length function as a special case of conformal factors and provide several examples of Hamiltonian vector fields. In Section 6.5, we discuss the presymplectic structure induced by the curvature-weighted metric, where we leave the non-degeneracy open for future research. In Section 6.6, we numerically illustrate simple Hamiltonian flows with respect to symplectic structures induced by length functions.

Finally, in Section 6.7, we explain an alternative approach via almost symplectic structures and the resulting conformal symplectic structures.

Notational conventions In this chapter, we adopt some notations tailored to its specific content, which deviates from the conventions used in previous chapters.

For instance, from Chapter 2 to Chapter 4, we wrote parametrized shapes with a tilde (e.g., $\tilde{\gamma} \in \text{Emb}(S, M)$), and unparametrized shapes with a raw alphabet (e.g., $\gamma \in \text{Emb}(S, M)/\text{Diff}^+(S)$). We made this choice because we performed our analysis in those chapters in a parameterization-free manner, and simpler notation helped with readability.

In this chapter, this convention is reversed. Parametrized shapes and their related quantities will use plain letters (e.g., $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$), while unparametrized shapes and their associated quantities will have an overbar (e.g., $\bar{c} \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/\text{Diff}^+(\mathbb{S}^1)$). We make this change because much of the analysis in this chapter uses parametrization of curves, so using the raw alphabets for them makes the exposition clearer.

In addition, we denote differential forms (including symplectic forms) and Riemannian metrics with capital characters. With this, we aim to ensure clarity, especially with the frequent use of multiple subscripts and superscripts. For example, we denote the MW structure at $\bar{c} \in \text{UImm}$ by $\Omega_{\bar{c}}^{\text{MW}}$. The subscript ‘MW’ is added to identify this specific structure, as several different symplectic structures appear in this chapter. We do so to avoid confusion caused by using lowercase characters like ω . When combined with various subscripts like ‘MW’, they can appear similar in font size to the main symbol, making it difficult to distinguish between the primary object and its subscript.

6.2 Liouville structures and (pre)symplectic structures

Marsden–Weinstein form on the space of unparametrized curves

We briefly review the MW structure on the shape space of unparametrized space curves as an infinite-dimensional orbifold. This is essentially a special case of the generic setting for codimension-2 submanifolds, explained in Chapter 2. Nevertheless, we outline the setting here to make this chapter self-contained. The only difference from the setting given in Chapter 2 is that we allow self-intersections of curves in this chapter, namely deal with immersions rather than embeddings.

We consider the space of regular space curves:

$$\text{Imm}(\mathbb{S}^1, \mathbb{R}^3) := \{c \in C^\infty(\mathbb{S}^1, \mathbb{R}^3) : |\partial_\theta c| \neq 0\},$$

which consists of immersions of \mathbb{S}^1 into \mathbb{R}^3 . The space $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ is an open subset of the vector space $C^\infty(\mathbb{S}^1, \mathbb{R}^3)$ and thus, similar as in finite dimensions, it is a manifold with tangent space given by the sections of the pullback bundle by c ,

$$T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) = \Gamma(c^* T\mathbb{R}^3).$$

Since \mathbb{R}^3 is a vector space, $\Gamma(c^* T\mathbb{R}^3)$ is identified with $C^\infty(\mathbb{S}^1, \mathbb{R}^3)$. In light of this, we treat $T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ as $C^\infty(\mathbb{S}^1, \mathbb{R}^3)$ throughout this chapter.

We will denote differentiation with respect to θ by a subscript, *i.e.*, we write $\partial_\theta c = c_\theta$ and $\partial_\theta h = h_\theta$ for $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ and $h \in T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$. Furthermore, we will occasionally

consider constant vector fields on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ obtained by extending tangent vectors $h \in T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ at some c to the entire $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$. We will denote this vector field also by h for simplicity, without explicitly stating so whenever it is clear from the context.

On the manifold of immersions we consider the action of the group of orientation-preserving diffeomorphisms $\text{Diff}^+(\mathbb{S}^1)$ by composition from the right. This leads us to consider the quotient (shape) space via the projection

$$\pi : \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) := \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) / \text{Diff}^+(\mathbb{S}^1),$$

which is an infinite dimensional orbifold with finite cyclic groups at the orbifold singularities, see [CMM91] and [Mic08, 7.3].

The vertical direction of the tangent space $T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ with respect to the natural fibration $\pi : \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ consists exactly of all fields h that are tangent to its foot point c , i.e., $h = ac_\theta$ with some $a \in C^\infty(\mathbb{S}^1)$. Hence the tangent space at $\bar{c} = \pi(c)$ is

$$\begin{aligned} T_{\bar{c}} \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) &= T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) / \ker d\pi|_c \\ &= C^\infty(\mathbb{S}^1, \mathbb{R}^3) / \{ac_\theta \in C^\infty(\mathbb{S}^1, \mathbb{R}^3) \mid a \in C^\infty(\mathbb{S}^1)\}. \end{aligned}$$

Intuitively, we mod out all the tangent vectors that do not change the shape \bar{c} .

In Section 2.3, we saw that $\text{UEmb}(\mathbb{S}^1, \mathbb{R}^3) = \text{Emb}(\mathbb{S}^1, \mathbb{R}^3)$ is a symplectic manifold equipped with the MW symplectic form. The MW form extends to the larger space $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ and can be explicitly written using the pullback 2-form $\Omega^{\text{MW}} := \pi^* \bar{\Omega}^{\text{MW}}$ on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$:

$$\Omega_c^{\text{MW}}(h, k) = \int_{\mathbb{S}^1} \det(\partial_\theta c, h, k) d\theta, \quad h, k \in T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3).$$

Note that $\bar{\Omega}^{\text{MW}}$ has a Liouville 1-form $\bar{\Theta}^{\text{MW}}$ i.e., $\bar{\Omega}^{\text{MW}} = -d\bar{\Theta}^{\text{MW}}$, which also has an explicit expression via the pullback 1-form $\Theta^{\text{MW}} := \pi^* \bar{\Theta}^{\text{MW}}$:

$$\Theta_c^{\text{MW}}(h) = \frac{1}{3} \int_{\mathbb{S}^1} \det(\partial_\theta c, c, h) d\theta.$$

We will find new symplectic structures by modifying the Liouville form $\bar{\Theta}^{\text{MW}}$ in a way inspired by the recent study of Riemannian metrics in shape analysis, which we explain now.

Reparametrization invariant Riemannian metrics on spaces of curves

On the space of parametrized curves we will consider reparameterization invariant (weak)-Riemannian metrics of the form:

$$G_c^L(h, k) = \int_{\mathbb{S}^1} \langle L_c h, k \rangle |c_\theta| d\theta = \int_{\mathbb{S}^1} \langle h, L_c k \rangle |c_\theta| d\theta$$

where $L \in \Gamma(\text{End}(T \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)))$ is an operator field, depending smoothly on $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ such that for each fixed curve c the operator

$$L_c : T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) = C^\infty(\mathbb{S}^1, \mathbb{R}^3) \rightarrow T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) = C^\infty(\mathbb{S}^1, \mathbb{R}^3)$$

is an elliptic pseudo differential operator that is equivariant under the right action of the diffeomorphism group $\text{Diff}^+(\mathbb{S}^1)$ and also under the left action of $SO(3)$, and which is also self-adjoint with respect to the L^2 -metric, i.e.,

$$L_{c \circ \varphi}(h \circ \varphi) = (L_c(h)) \circ \varphi \quad \text{and} \quad \int \langle L_c h, k \rangle ds = \int \langle h, L_c k \rangle ds.$$

Remark 6.2.1 (Sobolev metrics). An important class of such metrics is the class of Sobolev H^m -metrics, where $L = (1 - (-1)^m D_s^{2m})$ with $D_s = \frac{1}{|c_\theta|} \partial_\theta$ being the arclength derivative. Using the notation $ds = |c_\theta| d\theta$ for the arclength measure we obtain for $m = 0$ the metric

$$G_c^{\text{id}}(h, k) = \int_{\mathbb{S}^1} \langle h, k \rangle |c_\theta| d\theta = \int_{\mathbb{S}^1} \langle h, k \rangle ds$$

and for $m = 1$ the metric

$$G_c^{\text{id} - D_s^2}(h, k) = \int_{\mathbb{S}^1} \langle h, k \rangle + \langle -D_s^2 h, k \rangle ds = \int_{\mathbb{S}^1} \langle h, k \rangle + \langle D_s h, D_s k \rangle ds.$$

All these metrics can be written in terms of arc-length derivative $D_s = \frac{1}{|c_\theta|} \partial_\theta$ and arc-length integration $ds = |c_\theta| d\theta$ only. It has been shown that each such metric induces a corresponding metric on the shape space $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ such that the projection $\pi : \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ is a Riemannian submersion [MM07]. In finite dimension this would follow directly from the invariance of the metric, but in this infinite dimensional situation one has to show in addition the existence of the horizontal complement (with respect to the Riemannian metric). We will see, however, that this particular class of metrics will not be suited for the purpose of the present chapter, as the induced symplectic structure will not descend to a symplectic structure on the quotient space.

The induced Liouville 1-form

Next we will use the metric G^L to define a (Liouville) 1-form on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$. Therefore we consider for $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ and $h \in T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ the 1-form:

$$\Theta_c^L(h) := G_c^L(c \times D_s c, h) = \int \langle c \times D_s c, L_c h \rangle ds = \int \det(c, D_s c, L_c h) ds,$$

where \times denotes the vector cross product on \mathbb{R}^3 . We have the following result concerning its invariance properties:

Lemma 6.2.2 (Liouville 1-form). *For any inertia operator L , which is equivariant under the right action of the group of all orientation preserving diffeomorphisms and the left action of the rotation group $SO(3)$, the induced Liouville 1-form Θ^L is invariant under the right action of $\text{Diff}^+(\mathbb{S}^1)$ and the left action of $SO(3)$, i.e., for any $c \in \text{Imm}$, $h \in T_c \text{Imm}$, $\varphi \in \text{Diff}^+(\mathbb{S}^1)$ and $O \in SO(3)$ we have*

$$\Theta_{O(c \circ \varphi)}^L(O(h \circ \varphi)) = \Theta_c^L(h).$$

Proof. We will only show the reparametrization invariance, the invariance under $SO(3)$ is similar but easier. Using the equivariance of both L and D_s we calculate

$$\Theta_{c \circ \varphi}^L(h \circ \varphi) = \int \langle c \circ \varphi \times (D_s c) \circ \varphi, (L_c h) \circ \varphi \rangle |c_\theta| \circ \varphi |\varphi'| d\theta = \int \langle c \times D_s c, L_c h \rangle ds = \Theta_c^L(h). \quad \square$$

Remark 6.2.3. If L is equivariant under the left action of not only $SO(3)$ but of the larger group $SL(3) = \{M \in GL(3, \mathbb{R}) \mid \det(M) = 1\}$, then also Θ^L is invariant under $SL(3)$. This is the case for the Marsden-Weinstein structure $L = \text{id}$ (see Remark 6.2.5), but in general not for the inertia operators we deal with in this chapter.

The induced presymplectic form on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^2)$

Once we have defined the 1-form Θ we can formally consider the induced symplectic form

$$\Omega_c^L(h, k) := -d\Theta_c^L(h, k) = -D_{c,h}\Theta_c^L(k) + D_{c,k}\Theta_c^L(h) + \Theta_c^L([h, k]),$$

where d denotes the exterior derivative, $D_{c,h}$ denotes the directional derivative at $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ in the direction h , and when applied to a function $f: \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}$, we have $D_{c,h}f = \mathcal{L}_h f(c)$. The bracket $[h, k]$ is the Lie-bracket in $\mathfrak{X}(\text{Imm}(\mathbb{S}^1, \mathbb{R}^3))$ given by $[h, k] = D_{c,h}k - D_{c,k}h$.

In the following theorem we calculate this 2-form explicitly:

Theorem 6.2.4 (The (pre)symplectic form Ω^L on parametrized curves). *Let $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ and $h, k \in T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$. We have*

$$\begin{aligned} \Omega_c^L(h, k) = \int \left(\langle D_s c, L_c h \times k + h \times L_c k \rangle - \langle c, D_s h \times L_c k + L_c h \times D_s k \rangle \right. \\ \left. + \langle c \times D_s c, (D_{c,k} L_c) h - (D_{c,h} L_c) k \rangle \right) ds \end{aligned} \quad (6.2.1)$$

Furthermore, Ω^L is invariant under the right action of $\text{Diff}^+(\mathbb{S}^1)$ and under the left action of $SO(3)$.

Remark 6.2.5 (Marsden-Weinstein symplectic structure). It is known that for the invariant L^2 -metric, i.e., $L = \text{id}$, one obtains three times the Marsden-Weinstein (weak)-symplectic structure with this procedure (See [Tab17, PCK⁺19] for example), i.e.,

$$3\Omega_c^{\text{MW}}(h, k) := \Omega_c^{\text{id}}(h, k) = 3 \int_{\mathbb{S}^1} \langle D_s c \times h, k \rangle ds = 3 \int \det(D_s c, h, k) ds.$$

Its kernel consists exactly of all vector fields along c which are tangent to c , so by reduction it induces a presymplectic structure on shape space $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/\text{Diff}^+(\mathbb{S}^1)$ which is easily seen to be weakly non-degenerate and thus is a symplectic structure there.

Proof of Theorem 6.2.4. To prove the formula for Ω^L we first collect several variational formulas, see eg. [MM06] for a proof:

$$\begin{aligned} ds &= |c_\theta| d\theta, & D_{c,h} ds &= \frac{\langle h_\theta, c_\theta \rangle}{|c_\theta|} d\theta = \langle D_s h, D_s c \rangle ds \\ D_s &= \frac{1}{|c_\theta|} \partial\theta, & D_{c,h} D_s &= \frac{-\langle h_\theta, c_\theta \rangle}{|c_\theta|^3} \partial\theta = -\langle D_s h, D_s c \rangle D_s. \end{aligned}$$

Since $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ is open in $C^\infty(\mathbb{S}^1, \mathbb{R}^3)$, we can choose globally constant h, k i.e., independent of the location c on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$, namely $D_{c,h}(k) = D_{c,k}(h) = 0$ and $[h, k] = 0$. Using $D_{c,h}(L_c k) = (D_{c,h} L_c)k + L_c(h(k)) = (D_{c,h} L_c)k$, we compute

$$\begin{aligned} D_{c,h}\Theta_c^L(k) &= \int \left(\det(h, D_s c, L_c k) - \langle D_s h, D_s c \rangle \det(c, D_s c, L_c k) + \det(c, D_s h, L_c k) \right. \\ &\quad \left. + \det(c, D_s c, (D_{c,h} L_c)k) + \langle D_s h, D_s c \rangle \det(c, D_s c, L_c k) \right) ds \\ &= \int \left(\det(h, D_s c, L_c k) + \det(c, D_s h, L_c k) + \det(c, D_s c, (D_{c,h} L_c)k) \right) ds. \end{aligned}$$

Thus we get for Ω^L :

$$\begin{aligned}\Omega_c^L(h, k) &= -D_{c,h}\Theta_c^L(k) + D_{c,k}\Theta_c^L(h) + 0 \\ &= \int \left(-\det(h, D_s c, L_c k) + \det(k, D_s c, L_c h) - \det(c, D_s h, L_c k) + \det(c, D_s k, L_c h) \right. \\ &\quad \left. - \det(c, D_s c, (D_{c,h}L_c)k - (D_{c,k}L_c)h) \right) ds \\ &= \int \left(\langle D_s c, L_c h \times k + h \times L_c k \rangle - \langle c, D_s h \times L_c k - D_s k \times L_c h \rangle \right. \\ &\quad \left. - \langle c \times D_s c, (D_{c,h}L_c)k - (D_{c,k}L_c)h \rangle \right) ds,\end{aligned}$$

which yields the desired formula for Ω^L . The invariance properties of Ω^L follow directly from the corresponding invariance properties of Θ^L . \square

The induced symplectic structure on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$

In the previous part we have calculated a (pre)symplectic form on the space of parametrized curves $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$; we are, however, rather interested to construct symplectic structures on the shape space of geometric curves $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$. The following result contains necessary and sufficient conditions for the forms Θ^L and Ω^L to descend to this quotient space:

Theorem 6.2.6 (The (pre)symplectic structure on unparametrized curves). *The form Ω^L factors to a (pre)symplectic form $\bar{\Omega}^L$ on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ if the inertia operator L is equivariant under the $\text{Diff}^+(\mathbb{S}^1)$ -action and maps vertical tangent vectors to $\text{span}\{c, c_\theta\}$, i.e., if for all $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ and $a \in C^\infty(\mathbb{S}^1)$ we have $L_c(a \cdot c_\theta) = a_1 c_\theta + a_2 c$ for some functions $a_i \in C^\infty(\mathbb{S}^1)$.*

Proof. The Liouville form Θ^L on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ factors to a smooth 1-form $\bar{\Theta}^L$ on shape space $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ with $\Theta^L = \pi^* \bar{\Theta}^L$ if and only if Θ^L is invariant under the reparameterization group $\text{Diff}^+(\mathbb{S}^1)$ and is *horizontal* in the sense that it vanishes on each vertical tangent vector $h = a \cdot c_\theta$ for a in $C^\infty(\mathbb{S}^1, \mathbb{R})$.

Since Θ^L is invariant under the reparameterization group $\text{Diff}^+(\mathbb{S}^1)$ by construction it only remains to determine a condition on L such that Θ^L vanishes on all vertical h , i.e., we want

$$\Theta_c^L(a c_\theta) = \int \langle c \times D_s c, L_c(a c_\theta) \rangle ds = 0.$$

From here it is clear that this holds if $L_c(a \cdot c_\theta) = a_1 c_\theta + a_2 c$ for some functions $a_i \in C^\infty(\mathbb{S}^1)$.

In that case also its exterior derivative satisfies

$$\Omega^L = -d\Theta^L = -d\pi^* \bar{\Theta}^L = -\pi^* d\bar{\Theta}^L =: \pi^* \bar{\Omega}^L$$

for the presymplectic form $\bar{\Omega}^L = -d\bar{\Theta}^L$ on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$. \square

Example 6.2.7 (Inertia operators with a prescribed horizontal bundle). There are several different examples of operators that satisfy these conditions, including in particular the class of almost local metrics:

$$\begin{aligned}L_c(h) &= F(c) \cdot h \text{ for } F \in C^\infty(\text{Imm}(\mathbb{S}^1, \mathbb{R}^3), \mathbb{R}_{>0}), \text{ for example} \\ L_c(h) &= \Phi(\ell(c))h, \quad L_c(h) = \Phi\left(\int_{\mathbb{S}^1} \frac{\kappa^2}{2} ds\right)h, \quad L_c(h) = (1 + A\kappa^2)h,\end{aligned}$$

where $\kappa_c = |D_s^2 c|$ denotes the curvature and $\Phi: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{> 0}$ is a suitable smooth function. Note, that the class of Sobolev metrics, as introduced in Remark 6.2.1 does not satisfy the conditions of the above theorem. Thus these metrics do not induce a (pre)symplectic form on the quotient space. By including a projection operator in their definition one can, however, modify these higher-order metrics to still respect the vertical bundle:

$$L_c h = \left(\text{pr}_c (1 - (-1)^k D_s^{2k}) \text{pr}_c + (1 - \bar{\text{pr}}_c) (1 - (-1)^k D_s^{2k}) (1 - \text{pr}_c) \right) h,$$

where $\text{pr}_c h = \langle D_s c, h \rangle D_s c$ is the L^2 -orthogonal projection to the vertical bundle. For more details see [BH15], where metrics of this form were studied in detail.

Remark 6.2.8 (Horizontal Ω^L -Hamiltonian vector fields and $\bar{\Omega}^L$ -Hamiltonian vector fields). In the following we assume that the inertia operator $L \in \Gamma(\text{End}(T\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)))$ induces a (weak) symplectic structure on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$, i.e., it satisfies the conditions of Theorem 6.2.6 and is moreover weakly non-degenerate in the sense that $\bar{\Omega}^L: T\text{UImm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow T^*\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ is injective. Since $T_c^* \pi \circ \bar{\Omega}_{\pi(c)}^L \circ T_c \pi = \Omega_c^L$, this is equivalent to the kernel of $\Omega_c^L: T_c \text{Imm} \rightarrow T_c^* \text{Imm}$ being equal to the tangent space to the $\text{Diff}^+(\mathbb{S}^1)$ -orbit $c \circ \text{Diff}^+(\mathbb{S}^1)$ for all c . Thus Ω_c^L restricted to the G^L -orthogonal complement of $T_c(c \circ \text{Diff}^+(\mathbb{S}^1))$ is injective. See [KM97, Section 48] for more details.

Assume that H is a $\text{Diff}^+(\mathbb{S}^1)$ -invariant smooth function on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$. Then H induces a Hamiltonian function \bar{H} on the quotient space $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ with $\bar{H} \circ \pi = H$. Since the 2-form Ω^L on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ is only presymplectic it does not directly define a Hamiltonian vector field. However, if each dH_c lies in the image of $\Omega^L: T\text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow T^*\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$, then a unique smooth *horizontal Hamiltonian vector field* $X \in \mathfrak{X}(\text{Imm})$ is determined by

$$dH = \iota_X \Omega^L = \Omega^L(X, \cdot) \text{ and } G_c^L(X_c, T_c Y) = 0, \quad \forall Y \in \mathfrak{X}(\mathbb{S}^1)$$

which we will denote by $\text{hgrad}^{\Omega^L}(H)$. Obviously we then have

$$\text{grad}^{\bar{\Omega}^L}(\bar{H}) \circ \pi = T\pi \circ \text{hgrad}^{\Omega^L}(H).$$

Here and in the rest of the chapter, we write $\text{grad}^A E$ for the vector field satisfying $A(\text{grad}^A E, \cdot) = dE$ for a given non-degenerate bilinear form A such as a Riemannian metric and a symplectic form, and a function E .

Sometimes the kernel of Ω^L will be larger than the tangent spaces to the $\text{Diff}^+(\mathbb{S}^1)$ -orbits; then $\text{hgrad}^{\Omega^L}(H)$ will be chosen G^L -perpendicular to the kernel of Ω^L . This will happen in Theorem 6.3.2, for example, where L is a function of c such that Θ_c^L is also invariant under scaling. The Hamiltonian H factors to the corresponding space $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/\ker \Omega^L$ (which denotes the quotient by the foliation generated by $\ker \Omega^L$) if H is additionally invariant under each vector in $\ker \Omega^L$.

Remark 6.2.9. For the Marsden-Weinstein structure $\Omega^{\text{MW}} = -d\Theta^{\frac{1}{3}\text{id}}$, we have

$$\text{hgrad}^{\Omega^{\text{MW}}} H = -D_s c \times \text{grad}^{G^{\text{id}}} H$$

since

$$G^{\text{id}}(D_s c \times \cdot, \cdot) = \Omega^{\text{MW}}(\cdot, \cdot).$$

Remark 6.2.10 (Momentum mappings). If a Lie group \mathcal{G} acts on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ and preserves Θ^L , the corresponding momentum mapping J can be expressed in terms of Θ^L and the fundamental vector field mapping $\zeta : \mathfrak{g} \rightarrow \mathfrak{X}(\text{Imm}(\mathbb{S}^1, \mathbb{R}^3))$. For $Y \in \mathfrak{g}$, we have

$$\langle J(c), Y \rangle = \Theta^L(\zeta_Y)_c = \int \langle c \times D_s c, L_c \zeta_Y \rangle ds,$$

where $\langle \cdot, \cdot \rangle : \mathfrak{g}^* \times \mathfrak{g} \rightarrow \mathbb{R}$ is the duality product and ζ_Y is the fundamental vector field generated by Y . Namely,

$$d\Theta^L(\zeta_Y) = dt_{\zeta_Y} \Theta^L = \mathcal{L}_{\zeta_Y} \Theta^L - \iota_{\zeta_Y} d\Theta^L = 0 - \iota_{\zeta_Y} \Omega^L.$$

Lemma 6.2.2 asserts that Θ^L is invariant under the right action of $\text{Diff}^+(\mathbb{S}^1)$ and the left action of $SO(3)$.

Thus for $X = a \cdot \partial\theta \in \mathfrak{X}(\mathbb{S}^1) = C^\infty(\mathbb{S}^1) \partial\theta$ the *reparameterization momentum* is given as follows:

$$\begin{aligned} \zeta_{a \cdot \partial\theta}(c) &= D_{c, a \cdot c_\theta} \quad \text{as derivation at } c \text{ on } C^\infty(\text{Imm}, \mathbb{R}) \\ &= a \cdot c_\theta = a \cdot |c_\theta| D_s c \in T_c \text{Imm} = C^\infty(\mathbb{S}^1, \mathbb{R}^3) \\ L_{c \circ \varphi}(h \circ \varphi) &= (L_c h) \circ \varphi \implies (D_{c, a \cdot c_\theta} L_c)(h) + L_c(a \cdot h_\theta) = a \cdot (L_c h)_\theta \\ \langle J^{\text{Diff}^+(\mathbb{S}^1)}(c), a \cdot \partial\theta \rangle &= \Theta_c^L(\zeta_{a \cdot \partial\theta}(c)) = \Theta_c^L(a \cdot c_\theta) = \int \langle c \times D_s c, L_c(a \cdot c_\theta) \rangle ds \\ &= \int \langle c \times D_s c, a \cdot (L_c c)_\theta - (D_{c, a \cdot c_\theta} L_c)(c) \rangle ds. \end{aligned}$$

For $Y \in \mathfrak{so}(3)$ the *angular momentum* is

$$\begin{aligned} \langle J^{SO(3)}(c), Y \rangle &= \Theta^L(Y \circ c) = \int \langle c \times D_s c, L_c(Y \circ c) \rangle ds \\ &= \int \langle c \times D_s c, Y \circ L_c(c) - D_{c, Y \circ c} L_c(c) \rangle ds \end{aligned}$$

where $Y \circ c = \zeta_Y(c)$ denotes the multiplication of Y as a matrix with c as a vector. For a correct interpretation of the angular momentum recall (from [Mic08, 4.31], e.g.) that the action of $Y \in \mathbb{R}^3 \cong \mathfrak{so}(3) \cong L_{\text{skew}}(\mathbb{R}^3, \mathbb{R}^3)$ on \mathbb{R}^3 is given by $X \mapsto 2Y \times X$.

If L is also invariant under translations, then the *linear momentum*, for $y \in \mathbb{R}^3$, is

$$\langle J^{\mathbb{R}^3}(c), y \rangle = \Theta_c^L(y) = \int \langle c \times D_s c, L_c(y) \rangle ds.$$

Note that the above also furnishes conserved quantities on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$, if $\bar{\Omega}^L$ is non-degenerate.

6.3 Symplectic structures induced by conformal factors

In this section we consider symplectic structures induced by Riemannian metrics, that are conformally equivalent to the L^2 -metric, *i.e.*, we consider the G^L metric for $L_c = \lambda(c)$ where $\lambda : \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}_{>0}$ is invariant under reparametrization. Thus λ factors to a function $\bar{\lambda} : \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}_{>0}$ by $\pi^* \bar{\lambda} = \bar{\lambda} \circ \pi = \lambda$. Moreover, if $\text{grad}_c^{G^{\text{id}}} \lambda$ exists (which we assume) it is pointwise perpendicular to $D_s c$.

We first study the scale invariance of the corresponding Liouville 1-form, which will be of importance for the calculation of the induced (pre)symplectic structure. We say Θ^L is scale-invariant at $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ if $\mathcal{L}_I \Theta_c^L = 0$ where $I \in \Gamma(T\text{Imm}(\mathbb{S}^1, \mathbb{R}^3))$ is the scaling vector field $I_c := c$ with flow $\text{Fl}_t^I(c) = e^t \cdot c$. Depending on the context, we use both I and c for scaling as a tangent vector in this chapter.

Lemma 6.3.1 (Scale invariance of Θ^λ). *Let $L_c = \lambda(c) \text{id}$. Then the following are equivalent:*

- (a) Θ^λ is invariant under scalings.
- (b) $3\lambda(c) + \mathcal{L}_I \lambda(c) = 3\lambda(c) + D_{c,c} \lambda(c) = 0$ for all $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$.
- (c) $\lambda(c) = \Lambda(c/\ell(c)) \cdot \ell(c)^{-3}$ for a smooth function $\Lambda : \{c \in \text{Imm} : \ell(c) = 1\} \rightarrow \mathbb{R}_{>0}$,

where $\ell(c)$ is the length of c .

Proof. We have the following equivalences.

(a) \iff (b):

$$\begin{aligned} \mathcal{L}_I \Theta^\lambda &= \mathcal{L}_I(\lambda \Theta^{\text{id}}) = d\iota_I(\lambda \Theta^{\text{id}}) + \iota_I d(\lambda \Theta^{\text{id}}) = 0 + \iota_I(d\lambda \wedge \Theta^{\text{id}} + \lambda d\Theta^{\text{id}}) \\ &= \iota_I d\lambda \wedge \Theta^{\text{id}} + 0 + \lambda \iota_I d\Theta^{\text{id}} = (\iota_I d\lambda) \Theta^{\text{id}} - \lambda \iota_I \Omega^{\text{id}} = (\iota_I d\lambda + 3\lambda) \Theta^{\text{id}}. \end{aligned}$$

(b) \iff (c): Let $\ell(c) = 1$.

$$\begin{aligned} \partial_t \lambda(tc) &= d\lambda_{tc} = D_{c,tc} \lambda = \frac{1}{t} D_{tc,tc} \lambda = \frac{-3}{t} \lambda(tc) \\ \iff \partial_t \log(\lambda(tc)) &= \frac{-3}{t} \iff \log(\lambda(tc)) = \log(\Lambda(c)t^{-3}) \iff \lambda(tc) = \Lambda(c) \cdot t^{-3}. \quad \square \end{aligned}$$

Equipped with the above Lemma we are now ready to calculate the induced symplectic structure Ω^λ , where we will distinguish between the scale-invariant and non-invariant case.

Theorem 6.3.2 (The (pre)symplectic structure Ω^λ). *Let $L_c = \lambda(c) \text{id}$ be $\text{Diff}^+(\mathbb{S}^1)$ -invariant. Then the induced (pre)symplectic structure on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ is given by*

$$\Omega^\lambda = \lambda \Omega^{\text{id}} + \Theta^{\text{id}} \wedge d\lambda. \quad (6.3.1)$$

Furthermore we have

- (a) *If $3\lambda(c) + (\mathcal{L}_I \lambda)(c) = 3\lambda(c) + D_{c,c} \lambda(c) \neq 0$ on any open subset of $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$, then Ω^λ induces a non-degenerate 2-form on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$, which is thus symplectic.*
- (b) *Assume in addition, that $X := \text{hgrad}^{\Omega^{\text{id}}} \lambda$ exists, is smooth, admits a flow, and that $3\lambda(c) + (\mathcal{L}_I \lambda)(c) = 0$ for all c . Denote by \mathcal{F} the involutive 2-dimensional vector subbundle spanned by the vector fields I and $\text{hgrad}^{\Omega^{\text{id}}} \lambda$. Then Ω^λ induces a non-degenerate 2-form on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/(\text{Diff}^+(\mathbb{S}^1) \times \mathcal{F})$. If $\mathcal{L}_X \ell = 0$, it is also non-degenerate on $\{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\ell}_{\bar{c}} = 1\}/\text{span}(\text{grad}^{\bar{\Omega}^{\text{id}}} \bar{\lambda})$ where $\ell = \bar{\ell} \circ \pi$ denotes the length function $\bar{\ell}$ on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$.*

Remark 6.3.3 (Smooth structure of the orbit space). In case (b), the vector field $X := \text{hgrad}^{\Omega^{\text{id}}} \lambda$ exists in $\mathfrak{X}(\text{Imm}(\mathbb{S}^1, \mathbb{R}^3))$ if and only if $\text{grad}^{G^{\text{id}}} \lambda$ exists and is smooth as we have $\text{hgrad}^{\Omega^{\text{id}}} \lambda = \text{hgrad}^{3\Omega^{\text{MW}}} \lambda = -\frac{1}{3} D_s c \times \text{grad}^{G^{\text{id}}} \lambda$. This is equivalent to $\bar{\lambda} \in C^\infty(\text{UImm}(\mathbb{S}^1, \mathbb{R}^3), \mathbb{R})$. For details, see [BIM24, Appendix]. Moreover, the vector fields I and $\text{hgrad}^{\Omega^{\text{id}}} \lambda (= \frac{1}{3} \text{hgrad}^{\Omega^{\text{MW}}} \lambda)$ are linearly independent at any c because $\Omega_c^{\text{id}}(\text{hgrad}_c^{\Omega^{\text{id}}} \lambda, I_c) = \iota_I d\lambda(c) = -3\lambda(c) \neq 0$ by assumption. So the dimension of \mathcal{F} is always 2. We project to the leaf space of the 2-dimensional distribution if it is integrable. This is the case, if the flow of $X = \text{hgrad}_c^{\Omega^{\text{id}}} \lambda$ and thus also of $\text{grad}^{\bar{\Omega}^{\text{id}}} \bar{\lambda}$ exists; then the flows of I and $\text{grad}^{\bar{\Omega}^{\text{id}}} \bar{\lambda}$ combine to a 2-dimensional $(ax + b)$ -group acting on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$. We assume that this is the case; to prove existence of the flow one has first to specify λ and then solve a non-linear PDE.

Furthermore, the smooth structure of the quotient space $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/(\text{Diff}^+(\mathbb{S}^1) \times \mathcal{F})$ is slightly subtle: it is always a Frölicher space with tangent bundle, see [KM97, Section 23]. If local smooth sections of the projection to the leaf space exist and if the $(ax + b)$ -orbits admit slices, then we get a principal bundle with structure group the $(ax + b)$ -group in the category of orbifolds, so the leaf space is also an orbifold.

Proof. The formula directly follows from the product rule applied to $d(\Theta^\lambda) = d(\lambda\Theta^{\text{id}})$.

Case (a): We now show the non-degeneracy; if a tangent vector h satisfies $h \perp D_s c$ pointwise and $\Omega_c^\lambda(h, k) = 0$ for any k , then $h = 0$. First, choosing $k = a.c$ with some non-zero constant $a \in \mathbb{R}^\times$ we get from $a.c \in \ker \Theta^{\text{id}}$ that,

$$0 = \Omega_c^\lambda(h, a.c) = \lambda\Omega^{\text{id}}(h, a.c) + \Theta^{\text{id}}(h)\iota_{a.c}d\lambda - 0 = a[3\lambda + D_{c,c}\lambda]\Theta^{\text{id}}(h).$$

With our assumption $3\lambda + D_{c,c}\lambda \neq 0$ we see $h \in \ker \Theta^{\text{id}}$.

Next, we test for $h \in \ker \Theta^{\text{id}}$ and $k = a.c$ with some function $a \in C^\infty(\mathbb{S}^1)$ to see

$$\Omega_c^\lambda(h, a.c) = \lambda\Omega^{\text{id}}(h, a.c) = 3\lambda \int a \langle c \times D_s c, h \rangle ds.$$

If this vanishes for any function a , we have $\langle c \times D_s c, h \rangle = 0$ everywhere. We now consider the regions:

- (i) The open subset $U = \{\theta \in \mathbb{S}^1 : c(\theta) \times D_s c(\theta) \neq 0\}$,
- (ii) The closed set $\mathbb{S}^1 \setminus U = \{\theta \in \mathbb{S}^1 : c(\theta) \times D_s c(\theta) = 0\}$.

Any h satisfying both $h \perp D_s c$ and $h \perp (c \times D_s c)$ pointwise is of the form $h = b.c + v$ with a function $b \in C^\infty(\mathbb{S}^1)$ supported on U and a vector field $v \in C^\infty(\mathbb{S}^1, \mathbb{R}^3)$ supported on $\mathbb{S}^1 \setminus U$ and $v \perp D_s c$ (and hence $v \perp c$ as well). Then we have

$$\begin{aligned} \Omega_c^\lambda(h, k) &= \lambda\Omega^{\text{id}}(h, k) + \Theta^{\text{id}}(h)\iota_k d\lambda - \Theta^{\text{id}}(k)\iota_h d\lambda \\ &= \lambda\Omega^{\text{id}}(b.c, k) + 0 - \Theta^{\text{id}}(k)\iota_{b.c} d\lambda \\ &\quad + \lambda\Omega^{\text{id}}(v, k) + 0 - \Theta^{\text{id}}(k)\iota_v d\lambda \\ &= \int_{\mathbb{S}^1} \langle (3\lambda.b + D_{c,b.c}\lambda + D_{c,v}\lambda)D_s c \times c + 3\lambda.D_s c \times v, k \rangle ds. \end{aligned}$$

We assumed that $\Omega_c^\lambda(h, k) = 0$ for all k , in particular, for ones supported on $\mathbb{S}^1 \setminus U$. Hence we have $v \equiv 0$. With this we have

$$\Omega_c^\lambda(h, k) = \int_U (3\lambda.b + D_{c,b,c}\lambda) \langle D_s c \times c, k \rangle ds.$$

In order that $\Omega_c^\lambda(h, k) = 0$ for any k , we must have $3\lambda.b + D_{c,b,c}\lambda \equiv 0$ on U . Since $D_{c,b,c}\lambda \in \mathbb{R}$ is constant, b is constant. Hence we have $b(3\lambda + D_{c,c}\lambda) \equiv 0$ and get $b \equiv 0$ from our assumption $3\lambda + D_{c,c}\lambda \neq 0$. Thus we obtained $h = 0$.

Case (b): By assumption $X := \text{hgrad}^{\Omega^{\text{id}}} \lambda$ exist; i.e., $d\lambda$ is in the image of $\Omega^{\text{id}} : T\text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow T^*\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ and satisfies $d\lambda = \iota_X \Omega^{\text{id}}$ and $\langle X, D_s c \rangle = 0$. Then we see $X \in \ker \Omega_c^{\text{id}}$ by direct computation using the assumed condition $3\lambda_c + D_{c,c}\lambda = 0$;

$$(\iota_X \Theta^{\text{id}})_c = \int \langle c \times D_s c, X \rangle ds = \frac{1}{3} \Omega_c^{\text{id}}(X_c, c) = \frac{1}{3} \iota_I d\lambda_c = -\lambda(c) \text{ Lemma 6.3.1.}$$

$$\begin{aligned} (\iota_X \Omega^\lambda)_c &= \iota_{X_c} (\lambda \cdot \Omega^{\text{id}} + \Theta^{\text{id}} \wedge d\lambda)_c = \lambda(c) \cdot \iota_{X_c} \Omega_c^{\text{id}} + \Theta_c^{\text{id}}(X_c) \cdot d\lambda_c - \iota_{X_c} d\lambda(c) \cdot \Theta_c^{\text{id}} \\ &= \lambda(c) \cdot d\lambda_c - \lambda(c) \cdot d\lambda_c - 0 = 0. \end{aligned}$$

Note also that the scaling field $I_c := c$ with flow $\text{Fl}_t^I(c) = e^t \cdot c$ is in the kernel of Ω_c^λ as we have

$$\iota_I \Omega_c^\lambda = \lambda \iota_I \Omega_c^{\text{id}} + \Theta^{\text{id}}(c) d\lambda - D_{c,c}\lambda \Theta^{\text{id}} = -3\lambda \Theta^{\text{id}} + 0 - D_{c,c}\lambda \Theta^{\text{id}} = 0.$$

Thus \bar{I} and \bar{X} , the π -related versions of I and X , are in the kernel of $\bar{\Omega}^\lambda$.

$$\begin{aligned} (\mathcal{L}_I \lambda)(c) &= d\lambda(c) = -3\lambda(c) \\ \mathcal{L}_I \Theta^{\text{id}} &= \iota_I d\Theta^{\text{id}} = -\iota_I \Omega^{\text{id}} = 3\Theta^{\text{id}} \\ \mathcal{L}_I \Omega^{\text{id}} &= -\mathcal{L}_I d\Theta^{\text{id}} = -d\mathcal{L}_I \Theta^{\text{id}} = -d(3\Theta^{\text{id}}) = 3\Omega^{\text{id}} \\ -3d\lambda &= \mathcal{L}_I d\lambda = \mathcal{L}_I (\iota_X \Omega^{\text{id}}) = (\iota_X \mathcal{L}_I + \iota_{[I,X]}) \Omega^{\text{id}} = 3\iota_X \Omega^{\text{id}} + \iota_{[I,X]} \Omega^{\text{id}} \\ \iota_{[I,X]} \Omega^{\text{id}} &= -6d\lambda = -6\iota_X \Omega^{\text{id}} \end{aligned}$$

Thus $\iota_{[I,X] + 6X} \Omega^{\text{id}} = 0$, so $[I, X] + 6X$ is in the kernel of Ω^{id} . Their π -related version $[\bar{I}, \bar{X}] + 6\bar{X}$ is in the kernel of $\bar{\Omega}^{\text{id}}$ which is weakly non-degenerate on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$. So $[\bar{I}, \bar{X}] = -6\bar{X}$ and also $[I, X] = -6X$. Thus if the Frobenius integrability theorem applies in this situation (equivalently, if the local flow of X exists), then the fields I and X span an integrable distribution, and the leaf space exists.

Now we shall make use of $\bar{\lambda}(\bar{c}) = \Lambda(c/\ell(c)) \cdot \ell(c)^{-3}$, where $\bar{c} = \pi(c) \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$. The function Λ is defined on the ℓ -unit sphere $\{c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) : \ell(c) = 1\}$. To simplify notation, extend it constantly to $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ so that $\Lambda(c) = \Lambda|_{\{\ell=1\}}(c/\ell(c))$, and we let $\lambda = \bar{\lambda} \circ \pi$

and $\Lambda = \bar{\Lambda} \circ \pi$. Then we have

$$\begin{aligned} d\lambda_c(h) &= \ell(c)^{-3} \left(d\Lambda_c(h) - 3\Lambda(c) \frac{1}{\ell(c)} \int \langle D_s h, D_s c \rangle ds \right) \\ &= d\Lambda\left(\frac{c}{\ell(c)}\right) \left(-\ell(c)^{-2} \int \langle D_s h, D_s c \rangle ds \cdot c + \ell(c)^{-1} h \right) - 3\Lambda\left(\frac{c}{\ell(c)}\right) \ell(c)^{-4} \int \langle D_s h, D_s c \rangle ds \end{aligned}$$

$$\begin{aligned} \Omega_c^\lambda(h, k) &= \lambda(c) \Omega_c^{\text{id}}(h, k) + (\Theta_c^{\text{id}} \wedge d\lambda_c)(h, k) \\ &= \Lambda(c/\ell(c)) \cdot \ell(c)^{-3} \cdot \Omega_c^{\text{id}}(h, k) + \ell(c)^{-3} \Theta_c^{\text{id}}(h) \cdot \left(d\Lambda_c(k) - 3\Lambda(c) \frac{1}{\ell(c)} \int \langle D_s k, D_s c \rangle ds \right) \\ &\quad - \ell(c)^{-3} \left(d\Lambda_c(h) - 3\Lambda(c) \frac{1}{\ell(c)} \int \langle D_s h, D_s c \rangle ds \right) \cdot \Theta_c^{\text{id}}(k). \end{aligned}$$

We have diffeomorphisms which are equivariant under scalings

$$\begin{aligned} \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) / \text{Diff}^+(\mathbb{S}^1) &\cong \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) / (\text{Diff}^+(\mathbb{S}^1) \times \mathbb{R}_{>0}) \times \mathbb{R}_{>0} \\ &\cong \{\bar{c} \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) / \text{Diff}^+(\mathbb{S}^1) : \bar{\ell}(\bar{c}) = 1\} \times \mathbb{R}_{>0} \\ &\cong \{\bar{c} \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) / \text{Diff}^+(\mathbb{S}^1) : \bar{\lambda}(\bar{c}) = 1\} \times \mathbb{R}_{>0} \\ \bar{c} &\longleftrightarrow \left(\frac{1}{\bar{\ell}(\bar{c})} \bar{c}, \bar{\ell}(\bar{c}) \right) \longleftrightarrow \left(\bar{\Lambda}(\bar{c}/\bar{\ell}(\bar{c}))^3 \bar{c}, \bar{\ell}(\bar{c}) \right) \end{aligned}$$

and presymplectomorphisms

$$\{(\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\ell}(\bar{c}) = 1\}, \bar{\Omega}^\lambda) \ni \bar{c} \mapsto F(\bar{c}) = \bar{\Lambda}(\bar{c})^{1/3} \bar{c} \in (\{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\lambda}(\bar{c}) = 1\}, \bar{\Omega}^{\text{id}})$$

$$\{(\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\ell}(\bar{c}) = 1\}, \bar{\Omega}^\lambda) \xrightarrow{\iota_\lambda} (\text{UImm}(\mathbb{S}^1, \mathbb{R}^3), \bar{\Omega}^\lambda)$$

$$\{(\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\lambda}(\bar{c}) = 1\}, \bar{\Omega}^{\text{id}}) \xrightarrow{\iota_\lambda} (\text{UImm}(\mathbb{S}^1, \mathbb{R}^3), \bar{\Omega}^\lambda) \quad \text{since}$$

$$dF_c(k) = \frac{1}{3} \Lambda(c)^{-2/3} d\Lambda(c)(k) \cdot c + \Lambda(c)^{1/3} k$$

$$D_s F(c) = \frac{1}{3} \Lambda(c)^{-2/3} d\Lambda(c)(D_s c) \cdot c + \Lambda(c)^{1/3} D_s c$$

$$(F^* \bar{\Omega}^{\text{id}})_c(h, k) = \Omega_{F(c)}^{\text{id}}(dF_c(h), dF_c(k))$$

$$\begin{aligned} &= 3 \int \left\langle \left(\frac{1}{3} \Lambda(c)^{-2/3} d\Lambda_c(D_s c) \cdot c + \Lambda(c)^{1/3} D_s c \right) \times \left(\frac{1}{3} \Lambda(c)^{-2/3} d\Lambda_c(h) \cdot c + \Lambda(c)^{1/3} h \right), \right. \\ &\quad \left. \left(\frac{1}{3} \Lambda(c)^{-2/3} d\Lambda_c(k) \cdot c + \Lambda(c)^{1/3} k \right) \right\rangle ds \end{aligned}$$

$$= \Lambda(c) \Omega_c^{\text{id}}(h, k) + \Theta_c^{\text{id}}(h) \cdot d\Lambda(c)(k) - \Theta_c^{\text{id}}(k) \cdot d\Lambda(c)(h)$$

Since $(\text{UImm}(\mathbb{S}^1, \mathbb{R}^3), \bar{\Omega}^{\text{id}})$ is weakly symplectic and $\{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\lambda}(\bar{c}) = 1\}$ is a codimension 1 sub-orbifold diffeomorphic to $\{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\ell}(\bar{c}) = 1\}$, the kernel of $(\iota_\lambda^* \bar{\Omega}_c^{\text{id}})$ is 1-dimensional, and we have already found it as $\bar{X} = \text{grad}^{\bar{\Omega}^{\text{id}}} \bar{\lambda}$ which is tangent to $\{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\lambda}(\bar{c}) = 1\}$. \square

Remark 6.3.4 (Symplectic reduction). Our reduction of the space $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ in the second case of Theorem 6.3.2 can be seen as an infinite-dimensional instance of the Marsden-Weinstein-Meyer symplectic reduction. To see this, let us set $\bar{X} := \text{grad}_c^{\bar{\Omega}^{\text{id}}} \bar{\lambda}$ and take the momentum map $\bar{J}: \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}$ by $\bar{J}(\bar{c}) := \bar{\lambda}(\bar{c})$, with the corresponding group action

being the time- t flow of \bar{X} with \bar{c} as initial data. We have shown that $\bar{\Omega}^\lambda$ is degenerate on the codimension-1 sub-orbifold $\bar{J}^{-1}(1) = \{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) \mid \bar{\lambda}(\bar{c}) = 1\}$, and that it becomes symplectic when factored onto the codimension-2 sub-orbifold $\bar{J}^{-1}(1)/\text{grad}^{\bar{\Omega}^{\text{id}}} \bar{\lambda}$.

We also remark that the dual product for the momentum map \bar{J} is just the multiplication of scalar values as we have

$$\langle \bar{J}(\bar{c}), t \rangle = -\bar{\Theta}_c^{\text{id}}(t.\bar{X}) = t.\bar{\lambda}(\bar{c})$$

for $t \in \mathbb{R}$ such that the time- t flow map of \bar{X} exists. Here we used the invariance of $\bar{\Theta}_c^{\text{id}}$ under the flow of \bar{X} , which is shown by $\mathcal{L}_{\bar{X}} \bar{\Theta}_c^{\text{id}} = d\iota_{\bar{X}} \bar{\Theta}_c^{\text{id}} + \iota_{\bar{X}} \bar{\Omega}_c^{\text{id}} = -d\bar{\lambda} + d\bar{\lambda} = 0$ mimicking computations in the proof of Theorem 6.3.2. We may get the same result also using $\bar{\Theta}^\lambda$ and $\text{grad}_c^{\bar{\Omega}^\lambda} \bar{\lambda} = T_c \pi(\text{hgrad}_c^{\Omega^\lambda} \lambda)$ (Proposition 6.3.6) instead of $\bar{\Theta}_c^{\text{id}}$ and $\bar{X} = \text{grad}_c^{\bar{\Omega}^{\text{id}}} \bar{\lambda}$.

Remark 6.3.5 (A pseudo-Riemannian metric via Ω^L and $\mathcal{J} = D_s c \times \cdot$). Using the presymplectic form Ω^L and the standard almost complex structure

$$\begin{aligned} \mathcal{J}: T_c \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) &\rightarrow T_c \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) \\ \bar{h} &\mapsto \overline{D_s c \times h}, \end{aligned}$$

we may define a pseudo-Riemannian metric \bar{R} , which is compatible with $\bar{\Omega}^L$ via \mathcal{J} . Note that such \bar{R} is different from the Riemannian metric G^L factored onto $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ we used to define the Liouville form Θ^L .

We here compute \bar{R} for the conformal factor $L_c = \lambda(c)$. In the computation, we identify the tangent space at \bar{c} of $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ with the space of tangent vectors h in $T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$, such that $\langle D_s c, h \rangle = 0$, and denote $J = D_s c \times \cdot$.

We then have

$$\bar{R}_{\bar{c}}(\bar{h}, \bar{k}) := \Omega_c^\lambda(h, Jk) = \lambda(c) \Omega_c^{\text{id}}(h, Jk) + \Theta_c(h) \mathcal{L}_{Jk} \lambda(c) - \Theta_c(Jk) \mathcal{L}_h \lambda(c).$$

By design \bar{R} is non-degenerate. The symmetry follows from $\Omega_c^\lambda(Jh, k) = -\Omega_c^\lambda(h, Jk)$. It is, however, not clear if \bar{R} is positive-definite, *i.e.*, if it is a Riemannian metric. We leave this question open for future research.

6.3.1 Hamiltonian vector fields

Now we compute the horizontal Hamiltonian vector field $\text{hgrad}^{\Omega^\lambda} H$ for a given reparametrization-invariant Hamiltonian H . We express $\text{hgrad}^{\Omega^\lambda} H$ in terms of $\text{grad}^{G^{\text{id}}} H$ since the latter is in general relatively easy to obtain.

Proposition 6.3.6 (Horizontal Hamiltonian vector fields for Ω^λ). *Assume that $\text{grad}^{G^{\text{id}}} \lambda$ exists.*

(a) Consider a $\text{Diff}^+(\mathbb{S}^1)$ -invariant Hamiltonian $H: \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}^3$. If $3\lambda + \mathcal{L}_I \lambda \neq 0$ on any open subset of $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ then

$$\begin{aligned} \text{hgrad}^{\Omega^\lambda} H = & -\frac{1}{3\lambda(c)} \left\{ D_s c \times \text{grad}^{G^{\text{id}}} H \right. \\ & + \frac{1}{3\lambda(c) + D_{c,c}\lambda(c)} \left[\langle \text{grad}_c^{G^{\text{id}}} \lambda, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} D_s c \times (D_s c \times c) \right. \\ & \left. \left. - \langle c, \text{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} D_s c \times \text{grad}_c^{G^{\text{id}}} \lambda \right] \right\}. \end{aligned}$$

(b) Consider a Hamiltonian $H: \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}^3$ invariant under $\text{Diff}^+(\mathbb{S}^1)$ and the flows of the scaling vector field I and $\text{hgrad}^{\Omega^{\text{MW}}} \lambda = -D_s c \times \text{grad}^{G^{\text{id}}} \lambda$. If $3\lambda(c) + \mathcal{L}_{c,I}\lambda(c) = 0$ for all c then $\text{hgrad}_c^{\Omega^\lambda} H$ is the orthonormal projection of

$$X_c^H = -\frac{1}{3\lambda_c} D_s c \times \text{grad}_c^{G^{\text{id}}} H = \frac{1}{\lambda_c} \text{hgrad}_c^{\Omega^{\text{id}}} H$$

to the G_c^{id} -orthogonal complement of the kernel of Ω^λ , which is spanned by I , $\text{hgrad}^{\Omega^{\text{MW}}} \lambda$, and $\{a \cdot D_s c \mid a \in C^\infty(\mathbb{S}^1)\}$, namely

$$\text{hgrad}_c^{\Omega^\lambda} H = \frac{1}{3\lambda(c)} \left(-D_s c \times \text{grad}_c^{G^{\text{id}}} H + a_c \cdot (1 - \text{pr}_c) I_c - b_c \cdot D_s c \times \text{grad}^{G^{\text{id}}} \lambda \right)$$

where the pair $(a_c, b_c) \in \mathbb{R}^2$ is given by

$$\begin{pmatrix} a_c \\ b_c \end{pmatrix} = \begin{pmatrix} \langle v, v \rangle_{L_2} & \langle v, w \rangle_{L_2} \\ \langle v, w \rangle_{L_2} & \langle w, w \rangle_{L_2} \end{pmatrix}^{-1} \begin{pmatrix} \langle u, v \rangle_{L_2} \\ \langle u, w \rangle_{L_2} \end{pmatrix}.$$

with

$$\begin{aligned} u &= -D_s c \times \text{grad}^{G^{\text{id}}} H = \text{hgrad}^{\Omega^{\text{MW}}} H \\ v &= (1 - \text{pr}_c) I_c \\ w &= -D_s c \times \text{grad}^{G^{\text{id}}} \lambda = \text{hgrad}^{\Omega^{\text{MW}}} \lambda \end{aligned}$$

where the matrix appearing here is invertible because v_c and w_c are linearly independent at every $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$.

Note that in the scale-invariant case (Case (b)), the flow of the field Y^H projects to the Hamiltonian flow of \bar{H} on $\{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\lambda}(\bar{c}) = 1\} / \text{grad}^{\bar{\Omega}^{\text{id}}} \bar{\lambda}$ with respect to a multiple of the Marsden-Weinstein symplectic structure.

Proof. Let us denote for simplicity $A := \text{grad}^{G^{\text{id}}} \lambda$ and $X_H := \text{hgrad}^{\Omega^\lambda} H$. We can isolate out k from $\Omega_c^\lambda(X_H, k)$ by

$$\begin{aligned} \Omega^\lambda(X_H, k) &= \lambda \Omega^{\text{id}}(X_H, k) + \Theta^{\text{id}}(X_H) D_{c,k} \lambda - \Theta^{\text{id}}(k) D_{c,X_H} \lambda \\ &= \int_{\mathbb{S}^1} \langle 3\lambda \cdot D_s c \times X_H - D_{c,X_H} \lambda \cdot c \times D_s c + \Theta^{\text{id}}(X_H) A, k \rangle ds. \end{aligned}$$

Using $\Omega_c^\lambda(X_H, k) = dH(k) = G^{\text{id}}(\text{grad}^{G^{\text{id}}} H, k)$, we get

$$\begin{aligned} 0 &= \Omega^\lambda(X_H, k) - dH(k) \\ &= \int_{S^1} \langle 3\lambda.D_s c \times X_H - D_{c, X_H} \lambda.c \times D_s c + \Theta^{\text{id}}(X_H)A - \text{grad}^{G^{\text{id}}} H, k \rangle ds. \end{aligned}$$

This must be satisfied for any k , namely we have

$$3\lambda.D_s c \times X_H - D_{c, X_H} \lambda.c \times D_s c + \Theta^{\text{id}}(X_H)A - \text{grad}^{G^{\text{id}}} H = 0. \quad (6.3.2)$$

Our goal is to solve this for X_H . Applying $-D_s c \times$ reads

$$3\lambda.X_H - D_{c, X_H} \lambda.D_s c \times (D_s c \times c) - \Theta^{\text{id}}(X_H)D_s c \times A + D_s c \times \text{grad}^{G^{\text{id}}} H = 0.$$

Let us set

$$X_H = \frac{-1}{3\lambda} D_s c \times \text{grad}^{G^{\text{id}}} H + K_1 D_s c \times (D_s c \times c) + K_2 D_s c \times A_c \quad (6.3.3)$$

with some coefficients K_1, K_2 to be determined.

From

$$D_{c, X_H} \lambda = \int \langle A_c, X_H \rangle ds, \quad \Theta^{\text{id}}(X_H) = \int \langle c \times D_s c, X_H \rangle ds,$$

we get

$$\begin{aligned} 0 &= 3\lambda K_1.D_s c \times (D_s c \times c) + 3\lambda K_2.D_s c \times A_c \\ &\quad - \int \langle A_c, \frac{-1}{3\lambda} D_s c \times \text{grad}^{G^{\text{id}}} H + K_1 D_s c \times (D_s c \times c) \rangle ds.D_s c \times (D_s c \times c) \\ &\quad - \int \langle A_c, \frac{-1}{3\lambda} D_s c \times \text{grad}^{G^{\text{id}}} H + K_2 D_s c \times A_c \rangle ds.D_s c \times A_c \\ &= \left[K_1 \left(3\lambda - \int \langle A_c, D_s c \times (D_s c \times c) \rangle ds \right) \right. \\ &\quad \left. + \frac{1}{3\lambda} \int \langle A_c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle ds \right] D_s c \times (D_s c \times c) \\ &\quad + \left[K_2 \left(3\lambda + \int \langle D_s c \times c, D_s c \times A_c \rangle ds \right) \right. \\ &\quad \left. - \frac{1}{3\lambda} \int \langle D_s c \times c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle ds \right] D_s c \times A_c \\ &= \left[K_1 (3\lambda + D_{c,c}\lambda) + \frac{1}{3\lambda} \int \langle A_c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle ds \right] D_s c \times (D_s c \times c) \quad (6.3.4) \\ &\quad + \left[K_2 (3\lambda + D_{c,c}\lambda) - \frac{1}{3\lambda} \int \langle D_s c \times c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle ds \right] D_s c \times A_c. \end{aligned}$$

In the last step we used

$$- \int \langle D_s c \times (D_s c \times c), A_c \rangle ds = \int \langle D_s c \times c, D_s c \times A_c \rangle ds = D_{c, (1-\text{pr}_c)c} \lambda = D_{c,c} \lambda$$

where the last equality is due to the reparametrization-invariance of λ .

Case (a): Observe that

$$\begin{aligned} K_1 &= -\frac{1}{(3\lambda + D_{c,c}\lambda)3\lambda} \int \langle A_c, D_s c \times (D_s c \times c) \rangle ds, \\ K_2 &= \frac{1}{(3\lambda + D_{c,c}\lambda)3\lambda} \int \langle D_s c \times c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle ds \\ &= \frac{1}{(3\lambda + D_{c,c}\lambda)3\lambda} \int \langle c, \text{grad}^{G^{\text{id}}} H \rangle ds \quad \text{since } \text{grad}^{G^{\text{id}}} H \perp D_s c. \end{aligned}$$

satisfy the equality. Substituting K_1 and K_2 to (6.3.3), we obtain the stated formula. Note that the choice of the pair (K_1, K_2) is unique since $-(1 - \text{pr}_c)I_c = D_s c \times (D_s c \times c)$ and $-\text{hgrad}^{\Omega^{\text{MW}}} \lambda = D_s c \times A_c$ are linearly independent at least for some θ , namely in a small neighborhood. This follows from the linear independence of these two tangent vectors on $T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$, which is seen by the argument in the comment after Theorem 6.3.2 (b) with the reparametrization invariance of Ω^{id} .

Case (b): By assumption $3\lambda + D_{c,c}\lambda = 0$ we see from (6.3.4) that,

$$\begin{aligned} 0 &= \left[\int \langle A_c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle ds \right] D_s c \times (D_s c \times c) - \\ &\quad - \left[\int \langle D_s c \times c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle ds \right] D_s c \times A_c. \end{aligned}$$

Using this equality, it is easy to check that

$$X_H := \frac{-1}{3\lambda} D_s c \times \text{grad}^{G^{\text{id}}} H$$

satisfies (6.3.2). At this point there are up to two degrees of freedom in vector fields that satisfy (6.3.2). We can make X_H the unique horizontal lift of $\text{grad}^{\bar{\Omega}^\lambda} \bar{H}$ by performing the G_c^{id} -orthogonal projection with respect to $(1 - \text{pr}_c)I_c$ and $\text{hgrad}_c^{\Omega^{\text{MW}}} \lambda$, and hence obtain the stated expression. The resulting vector field X_H is G^{id} -orthogonal to $\{a.D_s c \mid a \in C^\infty(\mathbb{S}^1)\}$, $\text{hgrad}^{\Omega^{\text{MW}}} \lambda$ and I_c . \square

6.4 Symplectic structures induced by length weighted metrics

Next we study a special class of symplectic structures induced by conformal factors introduced in the previous section; namely we consider length-weighted metrics as studied in [YM05, MM06, Sha08]. More precisely, we consider operators of the form $L_c = \Phi(\ell(c))$ where $\ell(c) = \int_{\mathbb{S}^1} |c_\theta| d\theta$ denotes the length of the curve c and $\Phi : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ is a suitable function. Using Theorem 6.3.2 we obtain the following result concerning the induced symplectic structure $\Omega^{\Phi(\ell)}$:

Corollary 6.4.1 (The (pre)symplectic structure $\Omega^{\Phi(\ell)}$). *Let $\Phi \in C^\infty(\mathbb{R}_{>0}, \mathbb{R}_{>0})$. The induced*

(pre)symplectic structure of the $G^{\Phi(\ell)}$ -metric is given by:

$$\begin{aligned}\Omega_c^{\Phi(\ell)}(h, k) &= \Phi(\ell(c))\Omega^{\text{id}}(h, k) - \Phi'(\ell(c)) \left(\int_{\mathbb{S}^1} \langle D_s h, D_s c \rangle ds \Theta^{\text{id}}(k) - \int_{\mathbb{S}^1} \langle D_s k, D_s c \rangle ds \Theta^{\text{id}}(h) \right) \\ &= \Phi(\ell(c))\Omega^{\text{id}}(h, k) + \Phi'(\ell(c)) \left(\int_{\mathbb{S}^1} \langle h, D_s^2 c \rangle ds \Theta^{\text{id}}(k) - \int_{\mathbb{S}^1} \langle k, D_s^2 c \rangle ds \Theta^{\text{id}}(h) \right).\end{aligned}$$

Furthermore, we have:

- (a) If $\Phi(\ell) \neq C\ell^{-3}$ then the presymplectic structure $\bar{\Omega}^{\Phi(\ell)}$ on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ is non-degenerate and thus symplectic.
- (b) If $\Phi(\ell) = C\ell^{-3}$, then Ω^λ induces a non-degenerate 2-form on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/(\text{Diff}^+(\mathbb{S}^1) \times \mathcal{F}) \simeq \{\bar{c} \in \text{UImm}(\mathbb{S}^1, \mathbb{R}^3) : \bar{\ell} = 1\} / \text{span}(\text{grad}^{\bar{\Omega}^{\text{id}}} \bar{\ell})$, where it agrees with a multiple of the Marsden-Weinstein symplectic structure. Here \mathcal{F} is the 2-dimensional vector subbundle spanned by the scaling vector field I and $\text{hgrad}^{\Omega^{\text{MW}}} \ell = D_s c \times D_s^2 c$.

The Liouville form $\Theta^{C\ell^{-3}}$ is invariant under the scaling action $c \mapsto a.c$ for $a \in \mathbb{R}_{>0}$, which is equivalent to $\mathcal{L}_I \Theta^{C\ell^{-3}} = 0$. Note also that we have a diffeomorphism which is equivariant under scalings:

$$\begin{aligned}\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/\text{Diff}^+(\mathbb{S}^1) &\cong \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/(\text{Diff}^+(\mathbb{S}^1) \times \mathbb{R}_{>0}) \times \mathbb{R}_{>0} \\ &\cong \{\bar{c} \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)/\text{Diff}^+(\mathbb{S}^1) : \bar{\ell}(\bar{c}) = 1\} \times \mathbb{R}_{>0} \\ \bar{c} &\longleftrightarrow \left(\frac{1}{\bar{\ell}(\bar{c})} \bar{c}, \bar{\ell}(\bar{c}) \right)\end{aligned}$$

Proof. To calculate the formula for $\Omega^{\Phi(\ell)}$ we first need to calculate the variation of the length $\ell(c)$. We have:

$$D_{c,h} \ell(c) = \int_{\mathbb{S}^1} \langle D_s h, D_s c \rangle ds, \quad D_{c,h} \Phi(\ell(c)) = \Phi'(\ell(c)) \int_{\mathbb{S}^1} \langle D_s h, D_s c \rangle ds.$$

Applying this to (6.2.1) using integration by parts, we get

$$\begin{aligned}\Omega_c^{\Phi(\ell)}(h, k) &= \int_{\mathbb{S}^1} 2\Phi(\ell(c)) \langle D_s c, h \times k \rangle - \Phi(\ell(c)) \langle c, D_s h \times k - D_s k \times h \rangle ds \\ &\quad - \int_{\mathbb{S}^1} \langle c \times D_s c, (D_{c,h} \Phi(\ell(c))) k \rangle ds + \int_{\mathbb{S}^1} \langle c \times D_s c, (D_{c,k} \Phi(\ell(c))) h \rangle ds \\ &= 3\Phi(\ell(c)) \int_{\mathbb{S}^1} \langle D_s c, h \times k \rangle ds - \Phi'(\ell(c)) \int_{\mathbb{S}^1} \langle D_s h, D_s c \rangle ds \int_{\mathbb{S}^1} \langle c \times D_s c, k \rangle ds \\ &\quad + \Phi'(\ell(c)) \int_{\mathbb{S}^1} \langle D_s k, D_s c \rangle ds \int_{\mathbb{S}^1} \langle c \times D_s c, h \rangle ds \\ &= \Phi(\ell(c))\Omega^{\text{id}}(h, k) - \Phi'(\ell(c)) \int_{\mathbb{S}^1} \langle D_s h, D_s c \rangle ds \Theta^{\text{id}}(k) + \Phi'(\ell(c)) \int_{\mathbb{S}^1} \langle D_s k, D_s c \rangle ds \Theta^{\text{id}}(h),\end{aligned}$$

which proves the first formula for Ω . We may directly draw the last expression applying (6.3.1) to $\lambda = \Phi(\ell)$.

Case (a): It follows from Theorem 6.3.2 (a) that $\ker \Omega^{\Phi(\ell)} = \{a.D_s c \mid a \in C^\infty(\mathbb{S}^1)\}$, namely $\Omega^{\Phi(\ell)}$ induces a symplectic form $\bar{\Omega}^{\Phi(\ell)}$ on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$.

Case (b): By direct computation, we have $\text{hgrad}^{\Omega^{\text{id}}} C \ell^p = 3Cp\ell^{p-1}D_s c \times D_s^2 c$, which is a constant multiple of the Marsden-Weinstein flow $\text{hgrad}^{\Omega^{\text{MW}}} \ell = D_s c \times D_s^2 c$, so these two vector fields span the same distribution. Now the statements follow directly from Theorem 6.3.2 (b). \square

Now we will compute Hamiltonian vector fields. Therefore we note that the conditions of Remark 6.2.8 are satisfied, which allows us to obtain the following result:

Corollary 6.4.2 (Horizontal Hamiltonian Vector Fields for $\Omega^{\Phi(\ell)}$). *Consider a $\text{Diff}^+(\mathbb{S}^1)$ -invariant Hamiltonian $H: \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}$.*

(a) *If $\Phi(\ell) \neq C\ell^{-3}$, then:*

$$\begin{aligned} \text{hgrad}_c^{\Omega^{\Phi(\ell)}} H = \frac{1}{3\Phi(\ell(c))} & \left\{ -D_s c \times \text{grad}^{G^{\text{id}}} H \right. \\ & + \frac{\Phi'(\ell(c))}{3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c)} \left[\langle D_s^2 c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} D_s c \times (D_s c \times c) \right. \\ & \left. \left. + \langle c, \text{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} D_s c \times D_s^2 c \right] \right\}. \end{aligned} \quad (6.4.1)$$

(b) *If $\Phi(\ell) = C\ell^{-3}$, and if the Hamiltonian $H: \text{Imm}(\mathbb{S}^1, \mathbb{R}^3) \rightarrow \mathbb{R}^3$ invariant under $\text{Diff}^+(\mathbb{S}^1)$ and the flows of I and $\text{hgrad}^{\Omega^{\text{MW}}} \ell = D_s c \times D_s^2 c$, then $\text{hgrad}_c^{\Omega^\lambda} H$ is the orthonormal projection of*

$$X_c^H = -\frac{\ell^3}{3C} D_s c \times \text{grad}^{G^{\text{id}}} H$$

to the G_c^{id} -orthogonal complement of the kernel of Ω^λ , which is spanned by I and $\text{hgrad}^{\Omega^{\text{MW}}} \ell$, and $\{a.D_s c \mid a \in C^\infty(\mathbb{S}^1)\}$.

Proof. The stated formula follows from Proposition 6.3.6 with $\text{grad}_c^{G^{\text{id}}} \Phi(\ell(c)) = -\Phi'(\ell(c))D_s^2 c$ and that $\text{hgrad}^{\Omega^{\text{MW}}} C \ell^p$ is a constant multiple of $\text{hgrad}^{\Omega^{\text{MW}}} \ell$. \square

Remark 6.4.3. From the above Proposition it follows $\text{hgrad}^{\Omega^{\Phi(\ell)}} H$ agrees with $\text{hgrad}^{\Omega^{\text{id}}} H$ up to a constant scaling if $\Phi'(\ell(c)) = 0$. If $\Phi'(\ell(c)) \neq 0$ and $\langle D_s^2 c, \text{grad}^{\Omega^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} \neq 0$ then it is, however, genuinely different, *i.e.*, it does not seem realizable as a Hamiltonian vector field for the Marsden-Weinstein form Ω^{MW} . To formally prove that a given vector field X_H is never attained by the Marsden-Weinstein structure one needs to show that $\mathcal{L}_{X_H} \Omega^{\text{MW}} \neq 0$. Using the closeness of Ω^{MW} and Cartan's formula, this can be reduced to show that $d\iota_{X_H} \Omega^{\text{MW}} \neq 0$. However the necessary computations for this turn out to become extremely cumbersome and not very insightful. We refrain from providing them here.

Next we will consider several explicit examples, that will further highlight the statement of the above remark. We acknowledge that many of the Hamiltonian functions we consider were studied for the Marsden-Weinstein structure in [CKPP20].

Example 6.4.4 (Length function). We start with the arguably simplest Hamiltonian, namely we assume that H is a function of the total length ℓ , i.e., $H(c) = f \circ \ell(c)$ for some function f . In this case we calculate:

$$dH_c(k) = d[f \circ \ell]_c(k) = D_{c,k}f(\ell(c)) = f'(\ell(c)) \int \langle D_s k, D_s c \rangle ds = -f'(\ell(c)) \int \langle D_s^2 c, k \rangle ds,$$

hence

$$\text{grad}_c^{G^{\text{id}}} H = -f'(\ell(c)) D_s^2 c.$$

Using Corollary 6.4.2, we thus have

$$\text{hgrad}_c^{\Omega^{\Phi(\ell)}} H = \frac{f'(\ell(c))}{3\Phi(\ell(c))} \left(1 + \frac{\Phi'(\ell(c))\ell(c)}{3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c)} \right) D_s c \times D_s^2 c.$$

If $f'(\ell(c)) = 0$ for the initial length of the curve $\ell(c)$, it is a zero vector field. If $f'(\ell(c)) \neq 0$, then the length $\ell(c)$ is conserved along the flow as $H = f \circ \ell$ is conserved. Note that $\text{hgrad}_c^{\Omega^{\Phi(\ell)}} H$ is a constant multiple of the binormal equation (also known as the vortex filament equation),

$$\text{hgrad}_c^{\Omega^{\text{MW}}} \ell = D_s c \times D_s^2 c \quad (6.4.2)$$

using the Marsden-Weinstein symplectic structure.

Thus we have seen that the Hamiltonian vector field of the the symplectic structure $\Omega^{\Phi(\ell)}$ is a constant multiple of the Hamiltonian vector field of the Marsden-Weinstein symplectic structure. Note, that this constant factor, i.e., *the relative speed* with respect to the standard binormal equation, depends on the initial length $\ell(c)$.

Example 6.4.5 (Flux of a divergence-free vector field on \mathbb{R}^3 though a Seifert surface). Our next examples of Hamiltonians are the fluxes of vector fields through Seifert surfaces. We consider for any divergence-free vector field $V \in \Gamma(T\mathbb{R}^3)$ the closed 2-form $\times_{\iota_V} := \iota_V(dx \wedge dy \wedge dz)$. We can then define the corresponding *flux* by

$$E_V := \int_{D^2} \langle V \circ \Sigma, n \rangle = \int_{\Sigma(D^2)} \xi_V$$

where $\Sigma: D^2 \rightarrow \mathbb{R}^3$ is a smooth Seifert surface, i.e., an oriented and connected surface with $\Sigma|_{\partial D^2} = c$, and n is the unit surface normal.

We remark that E_V is independent of the choice of Σ . To see this, first notice that there is a unique 1-form α_V (up to addition of an exact 1-form) such that $d\alpha_V = \xi_V$ as $H_{dR}^1(\mathbb{R}^3) = 0$ and $H_{dR}^2(\mathbb{R}^3) = 0$. By Stokes theorem we have,

$$\int_{\Sigma(D^2)} \xi_V = \int_{D^2} \Sigma^* d\alpha_V = \int_{\Sigma(\partial D^2)} \alpha_V = \int_{c(S^1)} \alpha_V$$

where Σ^* denotes the pullback by Σ . For E_V , we have the following formulas from [CKPP20, Theorem 4]:

$$\begin{aligned} \text{grad}_c^{G^{\text{id}}} E_V &= D_s c \times (V \circ c), \\ \text{hgrad}_c^{\Omega^{\text{MW}}} E_V &= V \circ c. \end{aligned}$$

We consider E_V for two specific choices of V , where we use an analogous notation as in [CKPP20]: the translation $V_{-1} = v$ by some $v \in \mathbb{R}^3$ and the rotation $V_{-2}(x) = v \times x$ with some unit $v \in \mathbb{R}^3$ and we denote the corresponding fluxes by $H_{-1} = E_{V_{-1}}$ and $H_{-2} = E_{V_{-2}}$. Next we compute the horizontal Hamiltonian vector fields. From the computation

$$\begin{aligned} \langle D_s^2 c, D_s c \times (D_s c \times v) \rangle_{L^2(ds)} &= 0, \\ \langle c, D_s c \times v \rangle_{L^2(ds)} &= \int_{\mathbb{S}^1} \langle D_s c, v \times c \rangle ds = \int_{\mathbb{S}^1} \langle D_s c, 2 \operatorname{curl}(v) \circ c \rangle ds = 2H_{-1}(c), \end{aligned}$$

and

$$\begin{aligned} \langle D_s^2 c, D_s c \times (D_s c \times (v \times c)) \rangle_{L^2(ds)} &= 0, \\ \langle c, D_s c \times (v \times c) \rangle_{L^2(ds)} &= \int_{\mathbb{S}^1} \langle D_s c, (v \times c) \times c \rangle ds = \int_{\mathbb{S}^1} \langle D_s c, 3 \operatorname{curl}(v \times x) \circ c \rangle ds = 3H_{-2}(c), \end{aligned}$$

we obtain for $i \in \{-1, -2\}$,

$$\begin{aligned} \operatorname{hgrad}^{\Omega^{\Phi(\ell)}} H_i &= \frac{w_i}{3\Phi(\ell(c))} + \frac{C_i H_i(c) \Phi'(\ell(c))}{3\Phi(\ell(c))(3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c))} D_s c \times D_s^2 c \quad (6.4.3) \\ &= \frac{1}{3\Phi(\ell(c))} \operatorname{hgrad}^{\Omega^{\text{MW}}} H_i + \frac{C_i H_i(c) \Phi'(\ell(c))}{3\Phi(\ell(c))(3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c))} \operatorname{hgrad}^{\Omega^{\text{MW}}} \ell \end{aligned}$$

where $w_{-1} = v, w_{-2} = v \times c$ and $C_{-1} = 2, C_{-2} = 3$ respectively.

Since all of the three quantities ℓ, H_{-1} , and H_{-2} are constants in motion along the fields $\operatorname{hgrad}^{\Omega^{\text{id}}} H_i$ and $\operatorname{hgrad}^{\Omega^{\text{id}}} \ell$ [CKPP20, Corollary 1], the coefficients of both terms in (6.4.3) do not change along $\operatorname{hgrad}^{\Omega^{\Phi(\ell)}} H_i$. Hence the Hamiltonian fields $\operatorname{hgrad}^{\Omega^{\Phi(\ell)}} H_i$ are weighted sums of the Marsden-Weinstein Hamiltonian fields of ℓ and H_{-1} (or H_{-2} respectively).

Example 6.4.6 (Squared curvature). We next compute the Hamiltonian vector field for the squared curvature

$$H(c) := \frac{1}{2} \int \kappa^2 ds.$$

We have according to [CKPP20],

$$\operatorname{grad}^{G^{\text{id}}} H = D_s \left(D_s^3 c + \frac{3}{2} \kappa^2 D_s c \right), \quad D_s c \times \operatorname{grad}^{G^{\text{id}}} H = D_s c \times D_s^4 c + \frac{3}{2} \kappa^2 D_s c \times D_s^2 c.$$

Then, from

$$\begin{aligned} \langle D_s^2 c, D_s c \times \operatorname{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} &= \langle D_s^2 c, D_s c \times D_s^4 c \rangle_{L_{ds}^2(\mathbb{S}^1)} + 0 = 0, \\ \langle c, \operatorname{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} &= \int \kappa^2 - \frac{3}{2} \kappa^2 ds = -H(c), \end{aligned}$$

we have

$$\begin{aligned} \operatorname{hgrad}^{\Omega^{\Phi(\ell)}} H &= \frac{1}{3\Phi(\ell(c))} \left\{ -D_s c \times D_s^4 c - \frac{3}{2} \kappa^2 D_s c \times D_s^2 c - \frac{H\Phi'(\ell(c))}{3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c)} D_s c \times D_s^2 c \right\} \\ &= \frac{1}{3\Phi(\ell(c))} \left\{ \operatorname{hgrad}^{\Omega^{\text{MW}}} H - \frac{H\Phi'(\ell(c))}{3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c)} \operatorname{hgrad}^{\Omega^{\text{MW}}} \ell \right\}. \end{aligned}$$

Since both H and ℓ are again constants in motion along both $\operatorname{hgrad}^{\Omega^{\text{MW}}} \ell$ and $\operatorname{hgrad}^{\Omega^{\text{MW}}} H$ [CKPP20], $\operatorname{hgrad}^{\Omega^{\Phi(\ell)}} H$ is also realized as a Hamiltonian vector field of Ω^{id} .

Example 6.4.7 (Total torsion). We next consider the total torsion

$$H(c) := \int \tau ds.$$

Using the results [CKPP20, Theorem 2]

$$\begin{aligned} \text{grad}^{G^{\text{id}}} H &= -D_s c \times D_s^3 c, \\ D_s c \times \text{grad}^{G^{\text{id}}} H &= -D_s c \times (D_s c \times D_s^3 c), \end{aligned}$$

we compute

$$\begin{aligned} \langle D_s^2 c, D_s c \times \text{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} &= -\frac{1}{2} \int D_s \kappa^2 ds = 0 \\ \langle c, \text{grad}^{G^{\text{id}}} H \rangle_{L_{ds}^2(\mathbb{S}^1)} &= \langle D_s c, D_s c \times D_s^2 c \rangle_{L_{ds}^2(\mathbb{S}^1)} + \langle c, D_s^2 c \times D_s^2 c \rangle_{L_{ds}^2(\mathbb{S}^1)} = 0 + 0. \end{aligned}$$

Then we get

$$\text{hgrad}^{\Omega^{\Phi(\ell)}} H = \frac{1}{3\Phi(\ell(c))} D_s c \times (D_s c \times D_s^3 c) = \frac{1}{3\Phi(\ell(c))} \text{hgrad}^{\Omega^{\text{MW}}} H,$$

which is a scaled version of the Marsden-Weinstein gradient flow.

Example 6.4.8 (Squared scale). Next we consider the squared scale

$$E(c) := \frac{1}{2} \int |c|^2 ds,$$

as a Hamiltonian function. This is seen as the total kinetic energy of a moving particle in a periodic orbit in \mathbb{R}^3 .

We first get by a direct computation that,

$$\begin{aligned} \text{grad}^{G^{\text{id}}} E &= c - \langle c, D_s c \rangle D_s c - \frac{1}{2} |c|^2 D_s^2 c = (1 - \text{pr}_c) c - \frac{1}{2} |c|^2 D_s^2 c, \\ D_s c \times \text{grad}^{G^{\text{id}}} E &= D_s c \times c - \frac{1}{2} |c|^2 D_s c \times D_s^2 c, \end{aligned}$$

and

$$\begin{aligned} \langle D_s^2 c, D_s c \times \text{grad}^{G^{\text{id}}} E \rangle_{L_{ds}^2(\mathbb{S}^1)} &= -\Theta_c^{\text{id}}(D_s^2 c), \\ \langle c, \text{grad}^{G^{\text{id}}} E \rangle_{L_{ds}^2(\mathbb{S}^1)} &= \|D_s c \times c\|_{L_{ds}^2(\mathbb{S}^1)}^2 - \frac{1}{2} \langle c, |c|^2 D_s^2 c \rangle_{L_{ds}^2(\mathbb{S}^1)} \\ &= \|D_s c \times c\|_{L_{ds}^2(\mathbb{S}^1)}^2 + E(c). \end{aligned}$$

Using them with Corollary 6.4.2 gives us;

$$\begin{aligned}
 \text{hgrad}^{\Omega^{\Phi(\ell)}} E &= \frac{1}{3\Phi(\ell(c))} \left\{ -D_s c \times c + \frac{1}{2}|c|^2 D_s c \times D_s^2 c \right. \\
 &\quad + \frac{\Phi'(\ell(c))}{3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c)} \left[-\Theta_c^{\text{id}}(D_s^2 c) D_s c \times (D_s c \times c) \right. \\
 &\quad \quad \left. \left. - \left(\|D_s c \times c\|_{L_{ds}^2(\mathbb{S}^1)}^2 - \frac{1}{2} \langle c, |c|^2 D_s^2 c \rangle_{L_{ds}^2(\mathbb{S}^1)} \right) D_s c \times D_s^2 c \right] \right\} \\
 &= \frac{1}{3\Phi(\ell(c))} \left\{ -D_s c \times c + \frac{1}{2}|c|^2 D_s c \times D_s^2 c \right. \\
 &\quad + \frac{\Phi'(\ell(c))}{3\Phi(\ell(c)) + \Phi'(\ell(c))\ell(c)} \left[\Theta_c^{\text{id}}(D_s^2 c)(1 - \text{pr}_c)c \right. \\
 &\quad \quad \left. \left. - \left(\|D_s c \times c\|_{L_{ds}^2(\mathbb{S}^1)}^2 + E(c) \right) D_s c \times D_s^2 c \right] \right\}.
 \end{aligned} \tag{6.4.4}$$

Example 6.4.9 (Product of length and total squared curvature). Our last example is the Hamiltonian given by

$$H(c) = \ell(c)K(c)$$

where $K(c) = \int_{\mathbb{S}^1} \kappa^2 ds$ is the total squared curvature. This somewhat unusual Hamiltonian is the only one among our examples that satisfies the condition required in Corollary 6.4.2 (b) the scale-invariant case. That is, H is invariant under the both flows of $I = c$ and $Y := \text{hgrad}^{\Omega^{\text{MW}}} \ell = D_s c \times D_s^2 c$. To see this, let us compute

$$\mathcal{L}_Y H = K \mathcal{L}_Y \ell + \ell \mathcal{L}_Y K = K \cdot 0 + \ell \cdot 0 = 0$$

as ℓ is the Hamiltonian of Y and the last equality follows from a direct computation using (6.5.2). This shows the existence of a Hamiltonian vector field horizontal in the sense of Corollary 6.4.2 (b).

Question 6.4.10. We know from the above examples that some vector fields are realized as Hamiltonian vector fields of *both* $\bar{\Omega}^{\text{MW}}$ and $\bar{\Omega}^{\Phi(\ell)}$. We still do not know whether the spaces of all Hamiltonian vector fields generated by these two symplectic structures coincide, or if one is contained in the other. More generally, the coverage of Hamiltonian vector fields of $\bar{\Omega}^L$ for a given operator L is an independent question, which we have not investigated in this chapter.

6.5 Presymplectic structures induced by curvature weighted Riemannian metrics

In this section we will consider the special case of symplectic structures, that are induced by curvature weighted metrics, *i.e.*, we consider the Riemannian metric

$$G_c^{1+\kappa^2}(h, k) = \int_{\mathbb{S}^1} (1 + \kappa^2) \langle h, k \rangle ds,$$

where $\kappa = \kappa_c$ denotes the curvature of the curve c . Note, that in the notation of the previous sections, this metric corresponds to the G^L metric with $L = 1 + \kappa^2$. This metric, which is sometimes also called the Michor-Mumford metric, has been originally introduced in [MM06] to overcome the vanishing distance phenomenon of L^2 -metric, see also [MM05].

Remark 6.5.1 (Relations to the Frenet-Serret formulas). Given $c \in \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ we consider the open subset $U = \{\kappa > 0\} = \{D_s^2 c \neq 0\} \subset \mathbb{S}^1$. Note that $\kappa = 0$ on the boundary $\overline{U} \setminus U$, and is also 0 on the open complement $\mathbb{S}^1 \setminus \overline{U}$ which is a union of at most countably many open intervals in \mathbb{S}^1 ; on each of these intervals c is straight line segment since $D_s c$ is constant there. So we may assume that the torsion τ is defined and 0 on $\mathbb{S}^1 \setminus U$. On U the moving frame and the Frenet-Serret formulas are given by

$$\begin{aligned} T &= D_s c, & N &= \kappa^{-1} D_s^2 c, & B &= T \times N = \kappa^{-1} D_s c \times D_s^2 c \\ D_s T &= \kappa \cdot N = D_s^2 c, \\ D_s N &= -D_s \kappa \cdot \kappa^{-2} \cdot D_s^2 c + \kappa^{-1} D_s^3 c = -\kappa \cdot T + \tau \cdot B = -\kappa \cdot D_s c + \tau \cdot \kappa^{-1} D_s c \times D_s^2 c \\ D_s B &= -D_s \kappa \cdot \kappa^{-2} \cdot D_s c \times D_s^2 c = -\tau \cdot N = -\tau \cdot \kappa^{-1} D_s^2 c \end{aligned}$$

This implies the following which are valid on the whole of \mathbb{S}^1 since both sides vanish on $\mathbb{S}^1 \setminus U$:

$$\begin{aligned} D_s^3 c &= \langle D_s^3 c, T \rangle T + \langle D_s^3 c, N \rangle N + \langle D_s^3 c, B \rangle B \quad \text{valid on } U \\ &= \langle D_s^3 c, D_s c \rangle D_s c + \kappa^{-2} \langle D_s^3 c, D_s^2 c \rangle D_s^2 c + \kappa^{-2} \langle D_s^3 c, D_s c \times D_s^2 c \rangle D_s c \times D_s^2 c \quad \text{on } \mathbb{S}^1 \\ &= -\kappa^2 D_s c + D_s \kappa \cdot \kappa^{-1} \cdot D_s^2 c + \tau \cdot D_s c \times D_s^2 c \quad \text{valid on } U \text{ but extends smoothly to } \mathbb{S}^1 \\ \implies \langle D_s^3 c, D_s c \rangle &= -\kappa^2, & \langle D_s^3 c, D_s^2 c \rangle &= D_s \kappa \cdot \kappa, & \langle D_s^3 c, D_s c \times D_s^2 c \rangle &= \tau \cdot \kappa^2 \quad \text{valid on } \mathbb{S}^1 \\ \tau &= \kappa^{-2} \langle D_s^3 c, D_s c \times D_s^2 c \rangle \quad \text{valid on } \mathbb{S}^1. \end{aligned}$$

Remark 6.5.2. Similarly to Remark 6.2.10 we obtain again conserved quantities and corresponding momentum mappings. Here we want to specifically highlight the momentum map $J^{SO(3)}$: as an element of $\mathbb{R}^3 \approx \mathfrak{so}^*(3)$, the angular momentum $J^{SO(3)}$ is given by

$$\langle J^{SO(3)}(c), Y \rangle = \int (1 + \kappa^2) \langle c \times D_s c, Y \circ c \rangle ds,$$

which can be understood as the angular momentum of a thickened curve where the thickness (or mass) at each point is a function of $1 + \kappa^2$. Note, that this is in stark contrast to the previous section, *i.e.*, the length weighted case, where the angular momentum for $\Omega^{\Phi(\ell)}$ is just the $\Phi(\ell)$ -scaled version of the angular momentum for $\Omega^{\text{id}} = 3\Omega^{\text{MW}}$.

We have the following result concerning the induced presymplectic structure:

Theorem 6.5.3 (The presymplectic structure $\Omega^{1+\kappa^2}$). *The induced (pre)symplectic structure of the $G^{1+\kappa^2}$ -metric is given by:*

$$\begin{aligned} \Omega_c^{1+\kappa^2}(h, k) &= \int 3(1 + \kappa^2) \langle D_s c, h \times k \rangle + (D_s \kappa^2) \langle c, h \times k \rangle + 4\kappa^2 \langle D_s h, D_s c \rangle \langle c \times D_s c, k \rangle \\ &\quad - 2 \langle D_s^2 h, D_s^2 c \rangle \langle c \times D_s c, k \rangle - 4\kappa^2 \langle D_s k, D_s c \rangle \langle c \times D_s c, h \rangle + 2 \langle D_s^2 k, D_s^2 c \rangle \langle c \times D_s c, h \rangle ds, \end{aligned} \tag{6.5.1}$$

and the vertical vectors $\{a \cdot D_s c \mid a \in C^\infty(\mathbb{S}^1)\} \subset T_c \text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$ is in the kernel.

Proof of Theorem 6.5.3. To calculate the formula for $\Omega^{1+\kappa^2}$ we first need the variation of $\kappa^2 = \langle D_s^2 c, D_s^2 c \rangle$. Using, that $D_{c,h} D_s = -\langle D_s h, D_s c \rangle D_s$, cf. the proof of Theorem 6.2.4, we

calculate:

$$\begin{aligned}
 D_{c,h}(D_s^2 c) &= (D_{c,h} D_s) \cdot D_s c + D_s ((D_{c,h} D_s) c) + D_s^2 h \\
 &= -\langle D_s h, D_s c \rangle \cdot D_s^2 c - D_s (\langle D_s h, D_s c \rangle D_s c) + D_s^2 h \\
 &= -\langle D_s h, D_s c \rangle \cdot D_s^2 c - (D_s \langle D_s h, D_s c \rangle) D_s c - \langle D_s h, D_s c \rangle D_s^2 c + D_s^2 h \\
 &= -2\langle D_s h, D_s c \rangle \cdot D_s^2 c - (D_s \langle D_s h, D_s c \rangle) D_s c + D_s^2 h
 \end{aligned}$$

Thus we obtain

$$D_{c,h} \kappa^2 = -4\langle D_s h, D_s c \rangle \kappa^2 - 0 + 2\langle D_s^2 h, D_s^2 c \rangle. \quad (6.5.2)$$

Next we note that

$$\Omega_c^{1+\kappa^2}(h, k) = \Omega_c^{\text{id}}(h, k) + \Omega_c^{\kappa^2}(h, k)$$

as the operation $L_c \mapsto \Theta_c^L$ is linear in L_c . Using (6.2.1), we then calculate

$$\begin{aligned}
 \Omega_c^{\kappa^2}(h, k) &= \int \langle D_s c, \kappa^2 h \times k + h \times \kappa^2 k \rangle - \langle c, D_s h \times \kappa^2 k - D_s k \times \kappa^2 h \rangle \\
 &\quad - \langle c \times D_s c, (D_{c,h} \kappa^2) k - (D_{c,k} \kappa^2) h \rangle ds, \\
 &= \int 2\kappa^2 \langle D_s c, h \times k \rangle - \kappa^2 \langle c, D_s h \times k \rangle - \langle \kappa^2 c, h \times D_s k \rangle \\
 &\quad - D_{c,h} \kappa^2 \langle c \times D_s c, k \rangle + D_{c,k} \kappa^2 \langle c \times D_s c, h \rangle ds \\
 &= \int 2\kappa^2 \langle D_s c, h \times k \rangle - \kappa^2 \langle c, D_s h \times k \rangle \\
 &\quad + \langle D_s(\kappa^2 c), h \times k \rangle + \kappa^2 \langle c, D_s h \times k \rangle \\
 &\quad - D_{c,h} \kappa^2 \langle c \times D_s c, k \rangle + D_{c,k} \kappa^2 \langle c \times D_s c, h \rangle ds \\
 &= \int 3\kappa^2 \langle D_s c, h \times k \rangle + (D_s \kappa^2) \langle c, h \times k \rangle \\
 &\quad - D_{c,h} \kappa^2 \langle c \times D_s c, k \rangle + D_{c,k} \kappa^2 \langle c \times D_s c, h \rangle ds.
 \end{aligned}$$

Hence

$$\begin{aligned}
 \Omega^{1+\kappa^2}(h, k) &= \int 3(1 + \kappa^2) \langle D_s c, h \times k \rangle + (D_s \kappa^2) \langle c, h \times k \rangle \\
 &\quad - D_{c,h} \kappa^2 \langle c \times D_s c, k \rangle + D_{c,k} \kappa^2 \langle c \times D_s c, h \rangle ds.
 \end{aligned}$$

and (6.5.1) follows by using the variation formula (6.5.2) for κ^2 .

That Ω descends to a form on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ follows again from Theorem 6.2.6; alternatively we can also see this directly from the above formula: a straightforward calculation shows that $h = a \cdot D_s c$ is indeed in the kernel of $\Omega_c^{1+\kappa^2}$. \square

Question 6.5.4. It remains open if the presymplectic structure $\bar{\Omega}^{1+\kappa^2}$ on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ is non-degenerate and thus symplectic. Therefore it remains to show that tangent vectors of the form $a D_s c$ are the whole kernel of $\Omega_c^{1+\kappa^2}$. It seems natural to employ a similar strategy as in the previous section for length weighted metrics, *i.e.*, for given h we test with all k of the form $k = ac$ for $a \in C^\infty(\mathbb{S}^1)$. This leads to reducing the degeneracy of $\Omega^{1+\kappa^2}$ to solving the equation $P_c(a) = f$ for any given $f \in C^\infty(\mathbb{S}^1)$, where

$$P_c(a) := 2\langle D_s^2 c, c \rangle D_s^2 a - 4\langle D_s c, c \rangle \kappa^2 \cdot D_s a + (3 + \kappa^2) a;$$

The existence of periodic solutions for the above equation is, however, non-trivial. Note, that the coefficient functions are in general degenerate, *e.g.*, $\langle D_s^2 c, c \rangle$ can vanish somewhere.

Question 6.5.5. We may consider a more general version. Suppose $L_c: h \mapsto f_c \cdot h$ where f_c is a positive function for any c and is of form $f_c(\theta) = \rho(c(\theta), D_s c(\theta), D_s^2 c(\theta), \dots, D_s^N c(\theta))$ with some finite N and a function $\rho: \mathbb{R}^{3N} \rightarrow \mathbb{R}_{\geq 0}$. We expect that $\bar{\Omega}^L$ is symplectic on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ if Θ^L is not scale-invariant, or on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)/\mathcal{F}$ with a 2-dimensional distribution \mathcal{F} if Θ^L is scale-invariant (Theorem 6.3.2).

6.6 Numerical illustrations

In this section we numerically illustrate two Hamiltonian flows with respect to the new symplectic structures introduced in this chapter. For interested readers, we share video footage of the simulations shown in Figure 6.1 and Figure 6.2; see <https://youtu.be/nu09lwRK-tY>.

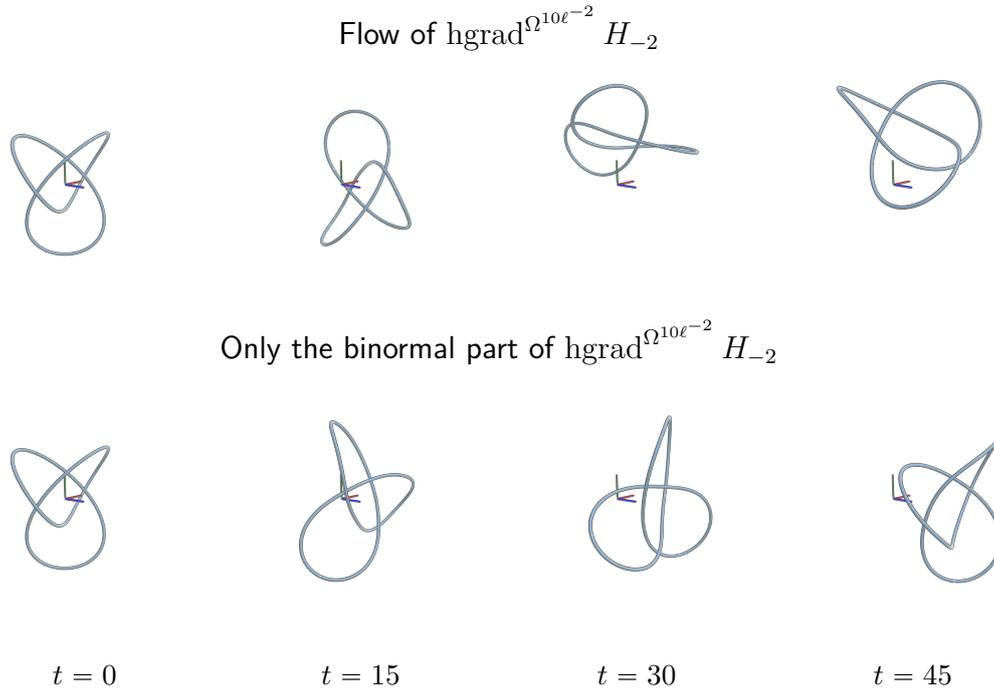


Figure 6.1: Hamiltonian flow of H_{-2} , the flux of a rotational vector field from Example 6.4.5 using $\Phi(\ell(c)) = 10\ell(c)^{-2}$ (top), and the flow only with its binormal component (bottom). The red, green, and blue axes are the x, y, z axes respectively.

For the numerical simulations, we discretized each curve as an ordered sequence of points in \mathbb{R}^3 . To approximate terms involving spatial derivatives, such as the binormal vector and the curvature, we follow the methods of discrete differential geometry, see [Bob15]. We then compute the time integration of each Hamiltonian vector field using the explicit Runge-Kutta method of fourth-order in time. We want to emphasize that our numerical examples are only for illustrative purposes and we do not guarantee any correctness of (even short-time) behaviors of the curve dynamics.

In our experiments, we use length-weighted presymplectic structures $\Omega^{\Phi(\ell)}$ (and symplectic structures $\bar{\Omega}^{\Phi(\ell)}$ for unparametrized curves) as derived in Section 6.4. That is, we use functions of the form $\Phi(\ell) = C\ell^p$ with some $C > 0$ and $p \in \mathbb{R}$. Note that C only works as time-scaling and does not change the orbit under the Hamiltonian flow. This is because in the expression of the field $\text{hgrad}^{\Omega^{\Phi(\ell)}} H$, cf. equation (6.4.1), the coefficient C appears only in the factor $\frac{1}{3\Phi(\ell)} = \frac{1}{C\ell^p}$ shared by all the terms and the factor $\frac{\Phi'(\ell)}{3\Phi(\ell) + \Phi'(\ell)\ell} = \frac{p}{(3+p)\ell}$ does not depend on C .

We choose C to run each simulation with a reasonable discrete timestep, but it essentially does not affect the dynamics.

We simulate two Hamiltonian flows (Example 6.4.5 and Example 6.4.8) from Section 6.4. These two examples involve only up to second-order spatial derivatives. Simulating other Hamiltonian flows, such as those discussed in Example 6.4.6 and Example 6.4.7 having third or higher-order derivatives is more challenging as one would have to discretize these higher-order derivatives more carefully.

As for the initial curve, we consider the trefoil

$$c(\theta) = ((2 + \cos(2\theta)) \cos(3\theta), (2 + \cos(2\theta)) \sin(3\theta), \sin(4\theta)), \quad \theta \in \mathbb{S}^1 = \mathbb{R}/2\pi\mathbb{Z} \quad (6.6.1)$$

in both of our examples.

Example 6.6.1 (Flux of a vector field). We first simulate the Hamiltonian flow for the Hamiltonian that is defined as the flux of a vector field through a Seifert surface whose boundary is the curve c , cf. Example 6.4.5. We chose the vector field of a rigid body rotation $V(x) = v \times x$ with the rotation axis $v = \frac{1}{\sqrt{3}}(1, 1, 1) \in \mathbb{R}^3$. This amounts to the Hamiltonian H_{-2} in Example 6.4.5.

The horizontal Hamiltonian field (6.4.3) is a weighted sum of the rotation $\text{hgrad}^{\Omega^{\text{MW}}} H_{-2} = v \times c$ and the binormal field $\text{hgrad}^{\Omega^{\text{MW}}} \ell = D_s^2 c \times D_s c$ with time-constant coefficients. Since these two flows are Poisson commutative, we can simulate the flow by evolving the curves under the binormal equation and rotating it at each time, *i.e.*, $c_t = \exp(t_1 \hat{v}) c_{t_2}^{\text{Binormal}}$ where $\hat{v} \in \mathfrak{so}(3)$ corresponds v and t_1, t_2 are time t weighted by the coefficients in (6.4.3). Figure 6.1 illustrates our simulation using $\Phi(\ell(c)) = 10\ell(c)^{-2}$. The top row is the flow of $\text{hgrad}^{\Omega^{\Phi(\ell)}} H_{-2}$ and the bottom row is the flow by only the binormal equation part where the curve moves toward the z -direction while showing a rotational motion around the z -axis.

Example 6.6.2 (Total squared scale). Our next example is the squared scale functional E (Example 6.4.8). Here we test three different choices of $\Phi(\ell) = C\ell^p$. Note again that we vary C only for computational purposes and this does not change the trajectory. The simulation results are shown in Figure 6.2.

We first compute for $\Phi(\ell) = \frac{1}{20}$, which corresponds to (a constant multiple of) the Marsden-Weinstein flow $\text{hgrad}^{\Omega^{\text{MW}}} E$. The curve moves back and forth in the z -direction, but curve points tend to get stuck once they come closer to the origin as both the term $-D_s c \times c$ and the term $\frac{1}{2}|c|^2 D_s c \times D_s c$ decrease as c goes to zero. As a result these parts form a complex shape around the origin. The next case is $\Phi(\ell) = \frac{1}{20}\ell^{-1/10}$. This shows a behavior similar to the first case, but points do not get stuck near the origin due to the additional term in (6.4.4). While moving back and forth, the curve does not become as entangled as in the previous case and seems to alternately transform between a trefoil and a trivial knot. The last case is $\Phi(\ell) = 10^{-5}\ell^2$. This shows a very different evolution. Unlike the other test cases, the curve does not globally translate in the z -direction but forms a complex spiral shape while shrinking slowly. In all three cases, the symmetry of the trefoil, *i.e.*, that rotation of 120 degrees around the z -axis does not change the shape, seems to be preserved in time.

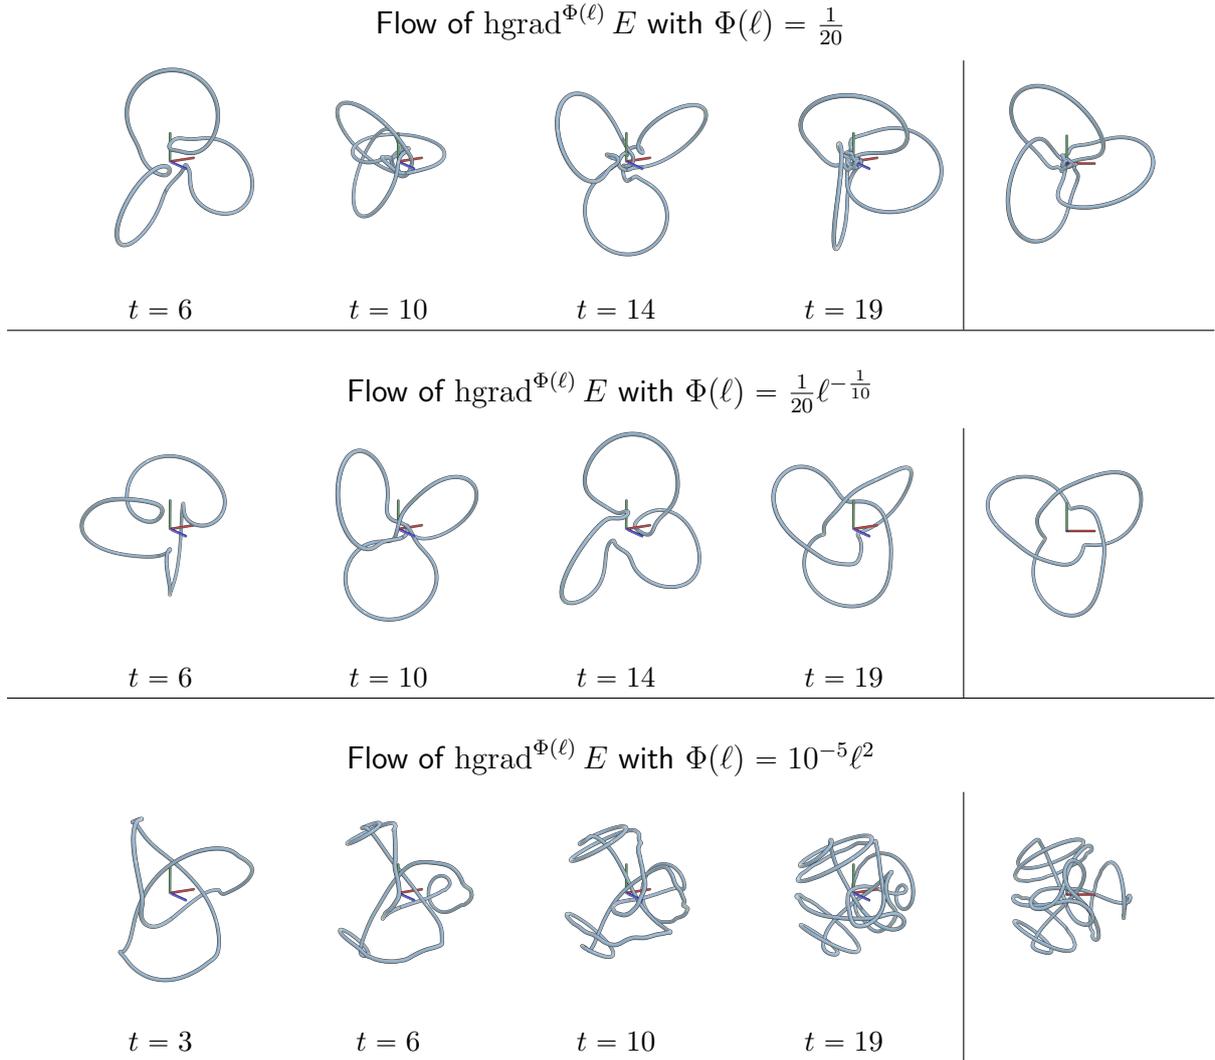


Figure 6.2: Hamiltonian flow of $\text{hgrad}^{\Omega^{\Phi(\ell)}} E$ with different choices of $\Phi(\ell)$. In each row the initial curve, which is not shown, corresponds to the trefoil (6.6.1). The right-most images are the front-view of the last configurations of curves showing high symmetry for the 120-degree rotation around the z-axis.

6.7 Appendix: Approach via an almost complex structure

In the main part of this chapter we constructed new symplectic structures on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$, by alternating the Liouville 1-form of the MW symplectic form. Doing so one obtained a skew symmetric 2-form on $\text{Imm}(\mathbb{S}^1, \mathbb{R}^3)$, which is automatically closed. Consequently to obtain a new symplectic structure it only remained to verify the non-degeneracy of this 2-form.

One may find this approach somewhat artificial or ad-hoc and could imagine, that it would be easier to construct new symplectic structures via directly alternating the MW symplectic structure instead of its Liouville 1-form. Following this alternate strategy one would arrive directly at a non-degenerate 2-form, but would instead need to prove its closedness. In this appendix we will discuss that this approach fails in the sense that we were not able to construct any closed forms following this procedure; nevertheless it results in an interesting class of 2-forms and will discuss some of their properties in more details.

On the shape space $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$, the mapping of 90 degrees rotation given by

$$\begin{aligned} \mathcal{J}: \text{TUImm}(\mathbb{S}^1, \mathbb{R}^3) &\rightarrow \text{TUImm}(\mathbb{S}^1, \mathbb{R}^3) \\ \bar{h} &\mapsto \overline{D_s c \times \bar{h}} \end{aligned}$$

is an almost complex structure, *i.e.*, an isomorphism with $\mathcal{J}^2 = -1$. The L^2 -Riemannian metric \bar{G}^{id} , the Marsden-Weinstein symplectic structure $\bar{\Omega}^{\text{MW}}$, and the almost complex structure \mathcal{J} formally define an almost Kähler structure (also called a compatible triple),

$$\bar{\Omega}^{\text{MW}}(\bar{h}, \bar{k}) = \bar{G}^{\text{id}}(\mathcal{J}(\bar{h}), \bar{k})$$

on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$ ².

This observation suggests to define a family of almost symplectic structures \bar{Z}^L via

$$\bar{Z}^L(\bar{h}, \bar{k}) = \bar{G}^L(\mathcal{J}(\bar{h}), \bar{k}). \quad (6.7.1)$$

If L is non-degenerate, self-adjoint with respect to \bar{G}^{id} and commutes with \mathcal{J} , then \bar{Z}^L is by construction *an almost symplectic structure*, *i.e.*, \bar{Z}^L is skew-symmetric and non-degenerate. To show that the induced forms are indeed symplectic it suffices thus to check the closedness of \bar{Z}^L .

At a first glance this approach seems promising and simpler than the approach presented in the main part of the chapter. However, as we will see in the following proposition, the 2-form \bar{Z}^L fails to be closed at least for all Riemannian metrics that are conformally equivalent (but not equal) to the L^2 -metric: The non-closedness of \bar{Z}^λ is due to the following proposition.

Proposition 6.7.1. *Let (M, ω) be a symplectic manifold (possibly orbifold) whose dimension is greater than 2 (possibly infinite-dimensional). Then the only symplectic structures in the conformal class of ω are the connected component-wise constant multiples of ω .*

Proof. Let $\omega^\lambda := \lambda\omega$ be a 2-form with conformal factor $\lambda: M \rightarrow \mathbb{R}_{>0}$. Then we have

$$d\omega^\lambda = \lambda d\omega + d\lambda \wedge \omega = d\lambda \wedge \omega.$$

The 3-form $d\lambda \wedge \omega$ is not identically zero unless λ is constant on each connected component. To see this let us denote $X_H = \text{grad}^\omega H$ for a given function $H: M \rightarrow \mathbb{R}$. Then we have,

$$d\lambda \wedge \omega = 0 \iff 0 = \iota_{X_H}(d\lambda \wedge \omega) = \iota_{X_H}d\lambda \wedge \omega - d\lambda \wedge \iota_{X_H}\omega = (\mathcal{L}_{X_H}\lambda) \cdot \omega - d\lambda \wedge dH \quad \forall H.$$

Since the two terms $(\mathcal{L}_{X_H}\lambda) \cdot \omega^{\text{id}}$ and $d\lambda \wedge dH$ have different ranks, they are both zero. Namely, $\mathcal{L}_{X_H}\lambda$ must be zero. Since at each point x any tangent vector $h \in T_x M$ is locally realized as $X_H(x)$ by choosing a suitable Hamiltonian H , λ must be constant. \square

While the above discussion is limited to Riemannian metrics that are conformally equivalent to the L^2 -metric, it seems that a similar phenomenon is also true for more complicated (higher

²This is not a Kähler structure in the classical sense, which additionally requires a complex structure *i.e.*, the existence of holomorphic coordinates. Indeed the Marsden-Weinstein symplectic structure does not admit a complex structure [Lem93]; it has been shown that on the space of isometric mappings of a circle into \mathbb{R}^3 modulo Euclidean transformations there is indeed a Kähler structure closely related to the Marsden-Weinstein structure, but with a more complicated almost complex structure than \mathcal{J} [MZ96], see also the comments in [Nee18].

order) metrics. In particular, we were not able to construct any pair of an almost complex structure \mathcal{J} and a non-conformal operator L which satisfy the required invariance conditions and leads to a closed form \tilde{Z}^L on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$.

This observation is the main reason why we proceeded to define our symplectic structures by altering the Liouville form, thereby ensuring closeness of the corresponding 2-form. We may also reach types of symplectic forms by solving the non-closeness issue of the approach in this section via correcting the obtained 2-form, e.g. in the conformal case when λ is not a constant and thus Z^λ is not closed, we could add some $W \in d^{-1}(d\lambda \wedge \Omega^{\text{id}})$ so that $Z^\lambda + W$ is closed. By doing this, the non-degeneracy property may be lost and thus one needs to check this again. Our approach for constructing a symplectic form from the Liouville form Θ^λ amounts to choosing $W = -d\lambda \wedge \Theta^{\text{id}}$. We emphasize that there is a large degree of freedom in $d^{-1}(d\lambda \wedge \Omega^{\text{id}})$ and our choice is not the unique one that makes the resulting form symplectic.

(Locally) conformal symplectic structure

Conformal symplectic geometry is a subbranch of symplectic geometry, which arose around 70's and has been studied mostly in finite-dimensional settings. For references on the general theory of (locally) conformal symplectic geometries and their induced dynamics we refer the interested reader to [Vai85, YM05, AA24].

Following these references an almost symplectic structure $\tilde{\omega}$ is said to be locally conformal symplectic if any point has an open neighborhood U such that there is a function f_U with $f_U \tilde{\omega}$ being symplectic in U . This condition is equivalent to $d\tilde{\omega} = \alpha \wedge \tilde{\omega}$ for some closed 1-form α , called the Lee 1-form. If α satisfies the additional condition that $\alpha = df$ for some globally defined function f then $f\tilde{\omega}$ is symplectic. In this case $\tilde{\omega}$ is simply called conformally symplectic.

While the approach using (6.7.1) with a Riemannian metric that is conformally equivalent to the L^2 -metric is unsuccessful for constructing new symplectic structures, this procedure exactly leads to conformal symplectic structures. This raises the following open question:

Question 6.7.2 (Locally but not globally conformal symplectic structures on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$). A natural question one may ask in this context concerns the existence of a locally but not globally conformal symplectic structure. We note that $H^1(\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)) \cong \mathbb{R}$ since for the fundamental group we have $\pi_1(\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)) = \mathbb{Z}$; this follows from [Sma59] and some standard arguments on fibrations and loop spaces; see [Ada93] for details. This observation suggests the existence of a non-degenerate ω s.t. $d\omega = \alpha \wedge \omega$ for a closed but non-exact Lie form α . We are, however, unaware of the existence/non-existence of such an only locally conformal structure on $\text{UImm}(\mathbb{S}^1, \mathbb{R}^3)$.

6.7.1 Hamiltonian vector fields of conformally symplectic structures

Here we consider the analogue of Hamiltonian systems for conformally symplectic structures. We begin by recalling the definitions and basic properties of Hamiltonian systems induced by a locally conformally symplectic structure. In this setting two types of Hamiltonian vector fields have been introduced: the first only preserves the Hamiltonian function but not the conformal symplectic structure, whereas the second one only preserves the conformal symplectic structure (up to conformal changes), but not the Hamiltonian function. Note, that this in contrast to Hamiltonian systems induced by a genuine symplectic structure, which preserve both of them.

To define these two types of Hamiltonian vector fields let ω be a locally conformal symplectic structure on a manifold M with a Lee form α , *i.e.*, α is a closed 1-form such that $d\omega = \alpha \wedge \omega$. Similar as for a classical Hamiltonian system we can define the field

$$\iota_{X_H}\omega = dH$$

Indeed, the flow along X_H preserves the Hamiltonian H but does not preserve ω , see e.g. [WL98]. The second type of Hamiltonian vector fields can be defined via

$$\iota_{Y_H}\omega = d_\alpha H$$

where d_α is the so-called twisted de Rham differential given by $d_\alpha = d - \alpha \wedge \cdot$. The vector field Y_H preserves the locally conformal symplectic structure up to conformal change given by $\mathcal{L}_{Y_H}\omega = \alpha(Y_H)\omega$. In general, it does not preserve the Hamiltonian except under restrictive circumstances such as H being constant, see [Vai85, MS17]. This type of Hamiltonian system can be used to model energy dissipating systems as studied in [AA24, BFM16].

Next we will derive these two types of Hamiltonian systems in our case, *i.e.*, for symplectic structures conformal to the MW structure on the space of space curves.

Horizontal Hamiltonian vector field via the standard differential

We first compute the first type of Hamiltonian vector field X_H via $\iota_{X_H}Z^\lambda = dH$. Using the same computational routine as in Section 6.3.1, we obtain the horizontal Hamiltonian vector field with respect to the standard L^2 metric G^{id} ,

$$X_H = -\frac{1}{\lambda}D_s c \times \text{grad}^{G^{\text{id}}} H = \frac{1}{\lambda} \text{hgrad}^{\Omega^{\text{MW}}} H,$$

which is a multiple of the Marsden–Weinstein flow with scaling factor λ . Note that H is preserved in time by design and that

$$\begin{aligned} \mathcal{L}_{X_H}Z^\lambda &= \iota_{X_H}d(\lambda\Omega^{\text{MW}}) + d(\iota_{X_H}(\lambda\Omega^{\text{MW}})) \\ &= \iota_{X_H}(d\lambda \wedge \Omega^{\text{MW}}) + ddH = \mathcal{L}_{X_H}\lambda \cdot \Omega^{\text{MW}} - \frac{d\lambda}{\lambda} \wedge dH, \end{aligned}$$

which is not preserved (even conformally) unless $d\lambda$ and dH are linearly dependent at each c .

Horizontal Hamiltonian vector field via the twisted differential

Next we discuss horizontal Hamiltonian dynamics of the second type, *i.e.*, we calculate the vector field Y_H via the relation $\iota_{Y_H}Z^\lambda = d_{d\lambda}H$ with the twisted differential $d_{d\lambda} = d - d\lambda \wedge \cdot$ mentioned above. Using the relation

$$dH - Hd\lambda = d_{d\lambda}H = \iota_{Y_H}Z^\lambda = \lambda \iota_{Y_H}\Omega^{\text{MW}}$$

and the computational routine in Section 6.3.1, we get

$$\text{grad}^{G^{\text{id}}} H - H \text{grad}^{G^{\text{id}}} \lambda = \lambda D_s c \times Y_H,$$

and hence

$$\begin{aligned} Y_H &= \frac{1}{\lambda} \left(-D_s c \times \text{grad}^{G^{\text{id}}} H + H \cdot D_s c \times \text{grad}^{G^{\text{id}}} \lambda \right) \\ &= \frac{1}{\lambda} \left(\text{hgrad}^{\Omega^{\text{MW}}} H - H \text{hgrad}^{\Omega^{\text{MW}}} \lambda \right). \end{aligned}$$

By direct computation, we see that

$$\mathcal{L}_{Y_H} H = H \mathcal{L}_{Y_H} \lambda, \quad \mathcal{L}_{Y_H} Z^\lambda = \mathcal{L}_{Y_H} \lambda \cdot Z^\lambda$$

and

$$\mathcal{L}_{Y_H} \lambda = \frac{1}{\lambda} (\mathcal{L}_{\text{hgrad}^{\Omega^{\text{MW}}}_H} \lambda - H \mathcal{L}_{\text{hgrad}^{\Omega^{\text{MW}}}_\lambda} \lambda) = \frac{1}{\lambda} \mathcal{L}_{\text{hgrad}^{\Omega^{\text{MW}}}_H} \lambda.$$

Hence, if λ is a conserved quantity under $\text{hgrad}^{\Omega^{\text{MW}}}_H$ (equivalently H is conserved along $\text{hgrad}^{\Omega^{\text{MW}}}_\lambda$), then H, λ and Z^λ are exactly preserved and the flow of Y_H behave like the Marsden–Weinstein Hamiltonian flow. Note, that there are infinitely many commutative Hamiltonian systems with respect to the MW structure forming a KdV-type hierarchy, cf. [CKPP20]. Therefore if we pick up any combination of the quantities within this hierarchy, we can find Hamiltonian systems induced by a conformal symplectic structure that resemble Hamiltonian systems induced by the MW structure.

Area formula for spherical polygons via prequantization

This chapter is largely a reprint of the article:

Albert Chern and Sadashige Ishida. Area formula for spherical polygons via prequantization. *SIAM Journal on Applied Algebra and Geometry*, 8(3), 2024

In the previous chapters, we studied symplectic geometry of codimension-2 shape space related to dynamics on that space. This chapter, on the other hand, applies a symplectic viewpoint to a purely geometric problem with practical applications.

We present a formula for the signed area of a spherical polygon via prequantization. In contrast to the traditional formula based on the Gauss–Bonnet theorem that requires measuring angles, the new formula mimics Green’s theorem and is applicable to a wider range of degenerate spherical curves and polygons.

Dependency of the chapter This chapter is self-contained, so readers can read it on its own.

7.1 Introduction

A spherical polygon is a finite number of ordered points on \mathbb{S}^2 connected by geodesics. Computing the solid angle of the region enclosed by a spherical polygon is important in many subjects, which we briefly summarize at the end of this section.

For a given polygon $\Gamma = (p_0, \dots, p_{n-1})$, its surface area is often computed using the formula

$$\text{Area}(\Gamma) = 2\pi - \sum_i \vartheta_i, \quad (7.1.1)$$

derived from the Gauss–Bonnet theorem. This area formula involves the evaluation of the exterior angle ϑ_i at each vertex

$$\vartheta_i = \text{sign}(\det(p_{i-1}, p_i, p_{i+1})) \arccos\left(\frac{p_{i-1} \times p_i}{|p_{i-1} \times p_i|} \cdot \frac{p_i \times p_{i+1}}{|p_i \times p_{i+1}|}\right), \quad (7.1.2)$$

which is a function of three points. However, this formula requires non-degeneracy assumptions about these points that render it unavailable or numerically unstable in certain situations. For example, two consecutive points of the polygon cannot lie on the same location as the exterior angle is undefined. We may consider removing such points from the polygon, but judging precisely whether two points are at the same location or just nearby locations is not always possible in numerical computation, especially when the points are obtained after some computational operations.

On the other hand, an area formula for polygons in \mathbb{R}^2 does not have such limitations. For a polygon $\Gamma_{\mathbb{R}^2} = ((x_0, y_0), \dots, (x_{n-1}, y_{n-1}))$, the area is

$$\text{Area}(\Gamma_{\mathbb{R}^2}) = \sum_i \frac{1}{2}(x_{i+1}y_i - x_iy_{i+1}). \quad (7.1.3)$$

Unlike the classical formula for spherical polygons (7.1.1, 7.1.2), this planar area formula involves adding numerically stable edge quantities. This formula can be derived by converting the area integral over the region into a line integral along the polygon curve using Green's formula. This is the same technique that enables *planimeters* (Figure 7.1). The key that allows Green's formula is that the area form of \mathbb{R}^2 is an exact differential form. Unfortunately, this strategy does not work directly for spherical polygons as the area form of \mathbb{S}^2 is not exact.

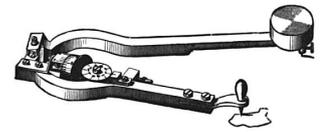


Figure 7.1: A planimeter measures the area of a planar region by tracing and accumulating quantities along its perimeter.

We circumvent this issue in this work via so-called *prequantization*, which is a preliminary setup to transform a classical mechanical system into a quantum mechanical system [BW97]. We lift the area form of the sphere onto a space where the resulting 2-form is exact. More precisely, we utilize a prequantum bundle, which is a principal circle bundle where the lifted 2-form is exact. This gives rise to a version of Green's theorem that translates the area integral on the base manifold into a line integral along a lifted perimeter in the bundle.

By choosing a specific prequantum bundle over \mathbb{S}^2 , we can obtain an explicit expression of the line integral. We use the Hopf fibration $\pi: \mathbb{S}^3 \rightarrow \mathbb{S}^2$ with a specific connection 1-form and derive a formula for the area of a spherical polygon. Unlike the classical formula (7.1.1), the new formula does not involve numerically unstable evaluation of angles of three consecutive vertices. Instead, it only involves a sum of edge-wise measurement resembling (7.1.3). We also recover the classical formula (7.1.1) by choosing $\text{SO}(3)$ as the prequantum bundle with a specific lift of the polygon. Finally, we present numerical examples comparing our formula and the classical formula.

Relation to quantum mechanics Our formula for the spherical area via prequantization is similar to evaluating the *Berry phase* for a quantum mechanical system. In this analogy, the points on \mathbb{S}^3 (or any prequantum circle bundle) represent the possible quantum states of a spinor (or 2 qubits), while their projections on \mathbb{S}^2 are their classical representations on the *Bloch sphere* [BZ06].

Applications and computations of solid angles Computing the area, or solid angle, of a spherical region occurs in multiple disciplines. They are either analytically computed by the angle-based formula (7.1.1, 7.1.2) [BC87], or approximated by point counting [FOM06, Gon10] and Monte-Carlo methods [Arv95b]. A standard routine in geological survey is to find the

area of an irregularly shaped region on the nearly-spherical earth [BC87, Gon10]. In quantum mechanics, the solid angle of a spherical curve generated by a Brownian motion of a particle describes the phase of two-state quantum systems undergoing random evolution [SW89, KSS00]. The so-called Majorana representation of polarized light also involves the solid angle of certain spherical quadrilaterals [Han98].

Solid angles give rise to another geometric concept, *the solid angle field of a space curve*. Given a space curve in \mathbb{R}^3 , the solid angle field is an \mathbb{S}^1 -valued function over \mathbb{R}^3 whose value at each point $x \in \mathbb{R}^3$ is half of the solid angle subtended by the closed curve at x . In contact geometry, a so-called *open-book decomposition* can be constructed by the solid angle field, with the level sets of the \mathbb{S}^1 -valued solid angle field being the *pages* and the space curve being the *binder* [Gei08]. Each page can also serve as a *Seifert surface*, an oriented surface bordered by a given knot [Dan16, BA18]. Solid angle fields also play important roles in fluid dynamics and electrodynamics, as the gradients of the solid angle fields are the *Biot-Savart fields* of space curves [Som52, pp.118], which represent the magnetic fields induced by electric currents, and the velocity fields corresponding to given vortex filaments. In addition, solid angle fields are applied to constructing implicit representations for space curves for simulating vortex dynamics [IWC22], visualization of nematic disclinations in electromagnetism [BA18], and radiosity illumination in rendering [WCZR20].

7.2 Area of a spherical polygon

We begin with the problem of seeking a line integral formula for the areas of spherical polygons. Next, we introduce the notion of prequantum bundles. While we only use the essential properties of this notion in a self-contained manner, the readers may find backgrounds in principal bundles [Nak03, KN63] and geometric quantization [BW97, Kos70] useful. Finally, we derive a line integral formula for spherical areas via prequantization and apply it to spherical polygons.

An oriented spherical polygon Γ is a finite cyclic ordered list of spherical points $\Gamma = (p_0, \dots, p_{n-1})$, $p_i \in \mathbb{S}^2$, $i \in \mathbb{Z}_n = \mathbb{Z}/(n\mathbb{Z})$. Each edge, *i.e.* each pair of adjacent points (p_i, p_{i+1}) , $i \in \mathbb{Z}_n$ (including (p_{n-1}, p_0) using the modulo arithmetic of $i \in \mathbb{Z}_n$), is joined by the shortest connecting path on \mathbb{S}^2 . This edge path is a constant point when $p_i = p_{i+1}$, and is otherwise a part of the great circle containing p_i, p_{i+1} . To ensure uniqueness of the shortest edge path, we assume that $p_{i+1} \neq -p_i$ for each $i \in \mathbb{Z}_n$, which does not take away generality as one may add intermediate points if the polygon contains any antipodal edge.

By concatenating all the edge paths, the spherical polygon Γ is naturally associated with a continuous closed path $C_\Gamma: \mathbb{S}^1 \rightarrow \mathbb{S}^2$. Let $S_\Gamma: \mathbb{D}^2 \rightarrow \mathbb{S}^2$ be any smooth extension¹ of $C_\Gamma: \mathbb{S}^1 = \partial\mathbb{D}^2 \rightarrow \mathbb{S}^2$ to the unit disk \mathbb{D}^2 . We call S_Γ the *enclosed region* of the spherical polygon Γ , which is unique up to an integer number of full wrappings around the entire sphere. Define the *signed area* of Γ as the total signed area of S_Γ , which is well-defined modulo 4π :

$$\text{Area}(\Gamma) := \iint_{S_\Gamma} \sigma = \iint_{\mathbb{D}^2} S_\Gamma^* \sigma \in \mathbb{R}/(4\pi\mathbb{Z}). \quad (7.2.1)$$

¹A smooth extension $S_\Gamma: \mathbb{D}^2 \rightarrow \mathbb{S}^2$ exists for any smoothly parameterized path $C_\Gamma: \mathbb{S}^1 \rightarrow \mathbb{S}^2$. A smooth path C_Γ can be constructed by concatenating smooth parameterizations of the great circular edges with vanishing derivatives $0 = \partial_t C_\Gamma = \partial_t^2 C_\Gamma = \dots$ at the vertices.

Here S_Γ^* denotes the pullback via S_Γ , and $\sigma \in \Omega^2(\mathbb{S}^2)$ is the standard area form of the unit sphere $\mathbb{S}^2 \subset \mathbb{R}^3$ induced by the Euclidean metric. The area form σ can be explicitly written as $\sigma = \sin \theta d\theta \wedge d\phi$ using a spherical coordinate chart, or as $\sigma = (x dy \wedge dz + y dz \wedge dx + z dx \wedge dy)/(x^2 + y^2 + z^2)^{3/2} |_{\mathbb{S}^2}$ using the Cartesian coordinates in \mathbb{R}^3 . This definition is valid for self-intersecting polygons (See Figure 7.2) and degenerate polygons which may contain edges of zero lengths or consecutive edges that fold back onto each other.

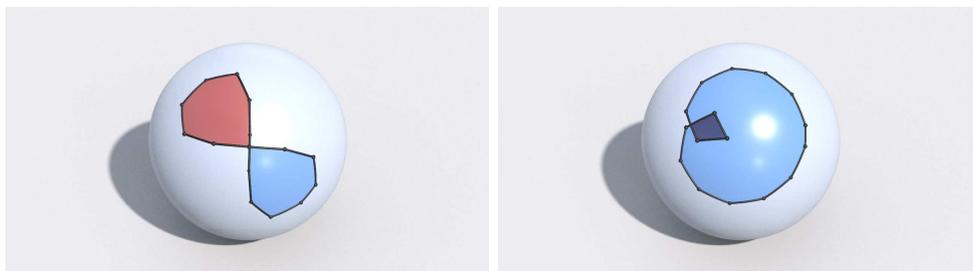


Figure 7.2: Examples of spherical polygons. Left: the two regions with different colors contribute positively or negatively to the signed area. Right: the darker region contributes to the signed area twice.

7.2.1 Areas as line integrals

Our goal is to find a numerically robust formula for (7.2.1) in terms of the vertex positions $\Gamma = (p_0, \dots, p_{n-1})$. More precisely, an ideal area formula we look for is a line integral “ $\text{Area}(\Gamma) = \oint_{C_\Gamma} \alpha$ ” for some smooth differential 1-form α , analogous to Green’s Theorem in the plane. The reason for this desire goes as follows. Once (7.2.1) becomes a line integral along the polygonal curve, we can derive the formula for $\text{Area}(\Gamma)$ by summing the explicit integrals of α along the great circular arc of each edge. Such a line-integral-based formula would be applicable to degenerate cases: The angles between consecutive edges would never appear in the formula; the line integral of a smooth 1-form along an edge shrinks to zero gracefully if the edge length shrinks to zero.

Such a line integral formula “ $\text{Area}(\Gamma) = \oint_{C_\Gamma} \alpha$ ” appears to rely on the exactness of σ , which is the existence of a smooth 1-form α so that $d\alpha = \sigma$. If α exists, then by Stokes’ Theorem $\text{Area}(\Gamma) = \iint_{S_\Gamma} \sigma = \iint_{S_\Gamma} d\alpha = \oint_{C_\Gamma} \alpha$. However, the *spherical area form σ is not exact*. Fortunately, and perhaps surprisingly, a line integral formula does not require the exactness of σ as described below.

Definition 7.2.1 (Prequantum bundle). Let $\beta \in \Omega^2(\Sigma)$ be a closed 2-form on a manifold Σ . A *prequantum bundle* over (Σ, β) is a principal circle bundle $\pi: Q \rightarrow \Sigma$ equipped with an equivariant 1-form $\alpha \in \Omega^1(Q)$ with $d\alpha = \pi^*\beta$.

Intuitively, a principal circle bundle Q over the base manifold Σ is a $(\dim(\Sigma) + 1)$ -dimensional space where a circle is attached to each point of Σ . The equivariance of $\alpha \in \Omega^1(Q)$ means that α is invariant under a uniform rotation in the circle dimension.

Remark 7.2.2. In a classical definition for a prequantum bundle, α is a *connection 1-form*, which has an additional requirement that $\alpha(V) = 1$ where V is the generator of the rotation action. In general, every compact symplectic manifold (Σ, β) admits a principal circle bundle (Q, α) with a connection 1-form α satisfying $d\alpha = \pi^*\beta$ upon a rescaling of β so that $\beta/2\pi \in H_2(\Sigma, \mathbb{Z})$ [BW58, Theorem 3] [Kob63, Proposition 9].

Proposition 7.2.3 (Lifted Green's theorem). *Let $\pi: (Q, \alpha) \rightarrow (\Sigma, \beta)$ be a prequantum bundle. For each surface $S: \mathbb{D}^2 \rightarrow \Sigma$ consider an arbitrary lift $\tilde{S}: \mathbb{D}^2 \rightarrow Q$, $\pi \circ \tilde{S} = S$. Then*

$$\iint_S \beta = \oint_{\partial \tilde{S}} \alpha.$$

Proof. $\oint_{\partial \tilde{S}} \alpha = \iint_{\tilde{S}} d\alpha = \iint_{\tilde{S}} \pi^* \beta = \iint_{\pi \circ \tilde{S}} \beta = \iint_S \beta.$ \square

The lifted Green's theorem enables line integral formula for the area of a spherical curve. Suppose we have a prequantum bundle $\pi: (Q, \alpha) \rightarrow (\mathbb{S}^2, \sigma)$ over the sphere \mathbb{S}^2 . For each closed curve C_Γ on \mathbb{S}^2 , construct an arbitrary lift $\tilde{C}_\Gamma: \mathbb{S}^1 \rightarrow Q$, $\pi \circ \tilde{C}_\Gamma = C_\Gamma$. Then

$$\text{Area}(\Gamma) = \oint_{\tilde{C}_\Gamma} \alpha \pmod{4\pi}.$$

Remark 7.2.4. While all the examples we provide in this chapter are prequantum bundles, for Proposition 7.2.3 we only need any smooth map $\pi: Q \rightarrow \Sigma$ and $\alpha \in \Omega^1(Q)$ that satisfies $d\alpha = \pi^* \beta$ and has a general liftability of a topological disk S to \tilde{S} . In particular, we do not need π to be a circle bundle or α to be equivariant.

7.2.2 The Hopf fibration

The Hopf fibration $\pi: \mathbb{S}^3 \rightarrow \mathbb{S}^2$ is a prequantum bundle over \mathbb{S}^2 . We provide its explicit expressions in the quaternion coordinate $\mathbb{H} = \{x\mathbf{i} + y\mathbf{j} + z\mathbf{k} + w \mid (x, y, z, w) \in \mathbb{R}^4\}$. Using the quaternion coordinates $q: \mathbb{S}^3 \hookrightarrow \mathbb{H} \cong \mathbb{R}^4$ and $p: \mathbb{S}^2 \hookrightarrow \text{Im } \mathbb{H} \cong \mathbb{R}^3$, Hopf's bundle projection is given by

$$\pi: \mathbb{S}^3 \rightarrow \mathbb{S}^2, \quad \pi(q) := q\mathbf{i}\bar{q}. \quad (7.2.2)$$

Note that π is a principal circle bundle with action $\triangleleft: \mathbb{S}^1 \times \mathbb{S}^3 \rightarrow \mathbb{S}^3$, $(\triangleleft e^{-i\theta})q := qe^{-i\theta}$. For q and $q' := qe^{-i\theta}$ on a same fiber, we write $\arg(\bar{q}'q) = \theta$.

The 1-form $\alpha \in \Omega^1(\mathbb{S}^3)$ for the prequantization is expressed as

$$\alpha = 2 \text{Re}(\mathbf{i}\bar{q}dq) = -2 \text{Re}(d\bar{q}q\mathbf{i}) = 2 \text{Re}(dq\mathbf{i}\bar{q}) = -2 \text{Re}(q\mathbf{i}d\bar{q}). \quad (7.2.3)$$

The 1-form $\alpha/2$ is a connection form as it is equivariant under \mathbb{S}^1 actions: $(\triangleleft e^{i\theta})^* \alpha = \alpha$ and $\frac{1}{2}\alpha(\frac{d}{d\theta}|_{\theta=0}(\triangleleft e^{-i\theta})q) = 1$.

Proposition 7.2.5. *The 1-form $\alpha \in \Omega^1(\mathbb{S}^3)$ defined in (7.2.3) and the map $\pi: \mathbb{S}^3 \rightarrow \mathbb{S}^2$ defined in (7.2.2) satisfy $d\alpha = \pi^* \sigma$, where $\sigma \in \Omega^2(\mathbb{S}^2)$ is the standard area form on the unit sphere. That is, (\mathbb{S}^3, α) is a prequantum bundle over (\mathbb{S}^2, σ) .*

This result is known ([BZ06], [CKP⁺16, Theorem 1]), but we give a proof for completeness in the appendix of this chapter (Section 7.6).

With this setting, Proposition 7.2.3 is now written more concretely for a spherical polygon:

Corollary 7.2.6. *The area $\text{Area}(\Gamma)$ of a spherical polygon $\Gamma = \{p_i\}_i$ can be evaluated by a line integral*

$$\text{Area}(\Gamma) = \sum_i \int_{\tilde{C}_\Gamma([t_i, t_{i+1}])} \alpha \pmod{4\pi} \quad (7.2.4)$$

where α is given by (7.2.3) and \tilde{C}_Γ is an arbitrary lift $\tilde{C}_\Gamma: \mathbb{S}^1 \rightarrow \mathbb{S}^3$ of the polygon curve $C_\Gamma = \pi \circ \tilde{C}_\Gamma: \mathbb{S}^1 \rightarrow \mathbb{S}^2$ with $C_\Gamma(t_i) = p_i$ for each i .

7.3 The area formula via the Hopf fibration

In this section, we obtain an explicit formula by evaluating each piece of the line integral in (7.2.4) along a great circular arc lifted onto \mathbb{S}^3 . We first define *dihedral* for a pair of spherical points.

Definition 7.3.1. For each $p, p' \in \mathbb{S}^2$, $p \neq -p'$, define

$$\text{Dihedral}(p, p') := \sqrt{\frac{1 + \langle p, p' \rangle}{2}} + \frac{p \times p'}{\sqrt{2 + 2\langle p, p' \rangle}} \in \mathbb{S}^3 \subset \mathbb{H}.$$

which is the unit quaternion that represents the minimal rotation that rotates p to p' i.e. $rp\bar{r} = p'$, $r = \text{Dihedral}(p, p')$.

There is a unique (unnormalized) *rotation axis* $v \in T_1\mathbb{S}^3 (= \text{Im } \mathbb{H})$ for the minimal rotation for non-antipodal $p, p' \in \mathbb{S}^2$. This axis v is parallel to $p \times p'$ and $\text{Dihedral}(p, p') = e^{\frac{v}{2}}$ by the exponential map on \mathbb{S}^3 . We also note that any $q \in \pi^{-1}p$, the point $\text{Dihedral}(p, p')q$ is on $\pi^{-1}p'$.

We will see that each summand in (7.2.4) can be explicitly expressed in terms of the dihedral after recognizing how the dihedral represents the horizontal lift over a great circular arc.

Proposition 7.3.2 (Horizontal lift on (\mathbb{S}^3, α)). *Consider the Hopf fibration $\pi: (\mathbb{S}^3, \alpha) \rightarrow (\mathbb{S}^2, \sigma)$ described in Section 7.2.2. Let p_0, p_1 be two arbitrary non-antipodal points on \mathbb{S}^2 , let $\gamma: [0, 1] \rightarrow \mathbb{S}^2$ be the great circular arc joining p_0 and p_1 , and let $v \in T_1\mathbb{S}^3$ be the imaginary quaternion such that $e^{\frac{v}{2}} = \text{Dihedral}(p_0, p_1)$. Then for each given point $q \in \pi^{-1}p_0$, the horizontal lift $\tilde{\gamma}_H$ of γ with $\tilde{\gamma}_H(0) = q$ is given by*

$$\tilde{\gamma}_H(t) = e^{\frac{vt}{2}} q.$$

Proof. It follows from $\gamma(t) = e^{\frac{vt}{2}} p_0 e^{-\frac{vt}{2}}$ that $\tilde{\gamma}_H$ is a lift over γ with respect to π . We now show that $\tilde{\gamma}_H$ is horizontal. At each $q \in \mathbb{S}^3$, the tangent space and the horizontal subspace with respect to the connection 1-form $\alpha/2$ are $T_q\mathbb{S}^3 = \text{Span}(q\mathfrak{i}, q\mathfrak{j}, q\mathfrak{k})$ and $H_q\mathbb{S}^3 = \text{Span}(q\mathfrak{j}, q\mathfrak{k})$ respectively. We have for each t that

$$\frac{d}{dt} \tilde{\gamma}_H(t) = \frac{1}{2} e^{\frac{vt}{2}} vq.$$

Here we note that the conjugation $h \mapsto qh\bar{q}$ is an isometry in $\text{Im}(\mathbb{H})$, from which it follows that $v \in q(H_q\mathbb{S}^3)\bar{q}$ as

$$0 = \langle v, p \rangle = \langle q(\bar{q}vq)\bar{q}, q\mathfrak{i}\bar{q} \rangle.$$

Therefore, $e^{\frac{vt}{2}} vq$ is horizontal. □

Lemma 7.3.3. *Assume $\pi: (\mathbb{S}^3, \alpha) \rightarrow (\mathbb{S}^2, \sigma)$, $p_0, p_1 \in \mathbb{S}^2$, and $\gamma: [0, 1] \rightarrow \mathbb{S}^2$ as in Proposition 7.3.2. Then for any lift $\tilde{\gamma}: [0, 1] \rightarrow \mathbb{S}^3$ of γ , the line integral $\int_{\tilde{\gamma}} \alpha$ is explicitly given by*

$$\int_{\tilde{\gamma}} \alpha = 2 \arg(\bar{q}_1 \text{Dihedral}(p_0, p_1) q_0),$$

where $q_0 := \tilde{\gamma}(0)$, $q_1 := \tilde{\gamma}(1)$.

Proof. Proposition 7.3.2 asserts that $q_0 \in \pi^{-1}p_0$ and $q_H := \text{Dihedral}(p_0, p_1)q_0 \in \pi^{-1}p_1$ are connected by the unique horizontal lift $\tilde{\gamma}_H$. Let us consider a parametric surface $\Delta \subset \mathbb{S}^3$ given by

$$\Delta := \{ \tilde{\gamma}_H(t)e^{i\theta} \in \mathbb{S}^3 \mid t \in [0, 1], \theta \in [0, \theta_t] \},$$

where θ_t for each t is the angular difference $\theta_t := \arg(\overline{\tilde{\gamma}(t)}\tilde{\gamma}_H(t))$. We first obtain that

$$\int_{\partial\Delta} \alpha = \int_{\Delta} d\alpha = \int_{\Delta} \pi^* \sigma = 0,$$

as $\pi(\Delta) = \gamma([0, 1])$.

Note that the boundary $\partial\Delta$ is a closed path consisting of three segments: (I). the lifted path $\tilde{\gamma}$ from q_0 to q_1 ; (II). the vertical path $\{q_1e^{i\theta} \mid \theta \in [0, \theta_1]\}$ from q_1 to q_H ; (III). the horizontal lift $\tilde{\gamma}_H$ from q_H to q_0 . Since the integral of α along the third path makes no contribution, we have

$$\begin{aligned} \int_{\tilde{\gamma}} \alpha &= - \int_{\{q_1e^{i\theta} \mid \theta \in [0, \theta_1]\}} \alpha = \int_{\{q_1e^{-i\theta} \mid \theta \in [-\theta_1, 0]\}} \alpha = \int_{-\theta_1}^0 \alpha \left(\frac{d}{d\theta} (\langle e^{-i\theta} \rangle q_1) \right) d\theta \\ &= \int_{-\theta_1}^0 2d\theta = 2 \arg(\overline{q_1}q_H), \end{aligned}$$

which concludes the proof. \square

As a direct result of Corollary 7.2.6 and Lemma 7.3.3, we obtain our main theorem:

Theorem 7.3.4 (Area formula via the Hopf fibration). *Let $\Gamma = (p_0, \dots, p_{n-1})$, $p_i \in \mathbb{S}^2$, $i \in \mathbb{Z}_n$, be a spherical polygon. For each $i \in \mathbb{Z}_n$, pick an arbitrary lift $q_i \in \mathbb{S}^3$ i.e. $\pi(q_i) = p_i$. Then,*

$$\text{Area}(\Gamma) = 2 \sum_{i=0}^{n-1} \arg(\overline{q_{i+1}} \text{Dihedral}(p_i, p_{i+1})q_i) \pmod{4\pi}. \quad (7.3.1)$$

The formula requires an arbitrary lift $q_i \in \mathbb{S}^3$ of the vertex positions $p_i \in \mathbb{S}^2$ for $i \in \mathbb{Z}_n$. An example is

$$q_i = \begin{cases} \text{Dihedral}(\mathfrak{i}, p_i), & \langle p_i, \mathfrak{i} \rangle \geq 0, \\ \text{Dihedral}(-\mathfrak{i}, p_i)\mathfrak{j}, & \langle p_i, \mathfrak{i} \rangle < 0, \end{cases} \quad (7.3.2)$$

which is uniquely defined globally.

As a special case of Theorem 7.3.4, choosing the horizontal lift results in no contribution of α except for the endpoint of the polygon.

Corollary 7.3.5 (Area formula by the horizontal lift). *Let q_0 be a point in the fiber $\pi^{-1}p_0$ and let us inductively define $q_{i+1} := \text{Dihedral}(p_i, p_{i+1})q_i$ for $i = 0, \dots, n-1$. Then we have*

$$\text{Area}(\Gamma) = 2 \arg(\overline{q_0}q_n).$$

At the end of the section, we make a remark regarding the numerical stability of the formula (7.3.1).

Remark 7.3.6. The branching discontinuities in the “arg” function in (7.3.1) and the “if” statement in (7.3.2) are smooth in the mod- 4π arithmetic of (7.3.1). The only calculation that can be numerically unstable is the evaluation of the Dihedral function when the two arguments are close to antipodal. This antipodal dihedral evaluation is avoided by the choice (7.3.2). The entire evaluation of our area formula (7.3.1) with (7.3.2) is numerically stable as long as we do not have antipodal edges where $\langle p_i, p_{i+1} \rangle \approx -1$, which is easily preventable by inserting a midpoint to any close-to-antipodal edge.

7.4 Derivation of the classical formula by $SO(3)$ as a prequantum bundle

The Hopf fibration structure can also be seen in the group $SO(3)$ of 3D rotations. In fact, the classical formula (7.1.1) can be interpreted as a special case of the lifted Green’s theorem on $SO(3)$ using a specific lift not as numerically stable as our formula (Theorem 7.3.4) using either the lift (7.3.2) or the horizontal lift (Corollary 7.3.5). We see that $SO(3)$ as the unit tangent bundle over \mathbb{S}^2 is also a prequantum bundle with a specific connection form. As \mathbb{S}^3 is a double cover of $SO(3)$, the Hopf fibration $\pi : \mathbb{S}^3 \rightarrow \mathbb{S}^2$ has a decomposition $\pi = \pi_2 \circ \pi_1$ given by,

$$\begin{aligned} \pi_1 : \mathbb{S}^3 &\rightarrow SO(3) \\ q &\mapsto (q \mathring{i} \bar{q}, q \mathring{j} \bar{q}, q \mathring{k} \bar{q}), \end{aligned}$$

and

$$\begin{aligned} \pi_2 : SO(3) &\rightarrow \mathbb{S}^2 \\ (p_1, p_2, p_3) &\mapsto p_1, \end{aligned} \tag{7.4.1}$$

where each element of $SO(3)$ is represented by three column vectors.

The tangent space $T_P SO(3)$ at each $P \in SO(3)$ is $dL_P \mathfrak{so}(3) = \{PW | W^T = -W\}$. We identify each $W = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix} \in \mathfrak{so}(3)$ with $\omega = (\omega_1, \omega_2, \omega_3) \in \mathbb{R}^3$. We define a 1-form η by

$$\eta|_P(PW) := -\omega_1. \tag{7.4.2}$$

Then (π_2, η) is a principal circle bundle with \mathbb{S}^1 action $\triangleleft : \mathbb{S}^1 \times SO(3) \rightarrow SO(3)$, by $(\triangleleft e^{i\theta})P = P \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}$.

Proposition 7.4.1. *The 1-form $\eta \in \Omega^1(SO(3))$ defined in (7.4.2) and the map $\pi_2 : SO(3) \rightarrow \mathbb{S}^2$ defined in (7.4.1) satisfy $d\eta = \pi_2^* \sigma$, where $\sigma \in \Omega^2(\mathbb{S}^2)$ is the standard area form on the unit sphere. That is, $(SO(3), \eta)$ is a prequantum bundle over (\mathbb{S}^2, σ) .*

We give a proof in the appendix of this chapter (Section 7.6). To recover the classical formula via $\pi_2 : (SO(3), \eta) \rightarrow (\mathbb{S}^2, \sigma)$, we define a lift as follows. We take for each p_i the forward velocity from p_i to p_{i+1} on \mathbb{S}^2 . It is given at $p_i =: \gamma(t_i)$ by

$$v(p_i) := \lim_{h \rightarrow +0} \frac{\gamma(t_i + h) - \gamma(t_i)}{h} \in T_{p_i} \mathbb{S}^2,$$

which is a positive multiple of $-p_i \times (p_i \times p_{i+1})$. Then

$$\tilde{\gamma}(t_i) := \left(p_i, \frac{v(p_i)}{|v(p_i)|}, p_i \times \frac{v(p_i)}{|v(p_i)|} \right)$$

defines a lift $\tilde{\gamma}: \mathbb{S}^1 \rightarrow \text{SO}(3)$ of γ . For $\tilde{\gamma}$, we have

$$\int_{\tilde{\gamma}} \eta = - \sum_i \vartheta_i \quad (7.4.3)$$

with the exterior angles ϑ_i given in (7.1.2). One has to be cautious when drawing conclusion from (7.4.3) about the area formula using Proposition 7.2.3. In fact, $\text{Area}(\Gamma) = 2\pi + \int_{\tilde{\gamma}} \eta = 2\pi - \sum_i \vartheta_i$ noting the extra term of 2π (cf. (7.1.1)). This is because $\tilde{\gamma}$ is a non-contractible loop in $\text{SO}(3)$ which is not the boundary of a disk. To obtain the classical formula (7.1.1), lift $\tilde{\gamma}: \mathbb{S}^1 \rightarrow \text{SO}(3)$ to $\hat{\gamma}: \mathbb{S}^1 \rightarrow \mathbb{S}^3$ by the universal cover $\pi_1: \mathbb{S}^3 \rightarrow \text{SO}(3)$. Note that $\alpha = \pi_1^* \eta$, and that $\hat{\gamma}(0) = \hat{\gamma}(2\pi)$ and $\lim_{t \nearrow 2\pi} \hat{\gamma}(t)$ has an angle difference of π in the fiber $\pi^{-1}(\gamma(0))$.

The classical formula has a variant that locates a pole $Z \in \mathbb{S}^2$ and sums up the signed area of triangles (p_i, p_{i+1}, Z) . This formula is given as,

$$\text{Area}(\Gamma) = \sum_i \text{sign}(\det(p_i, p_{i+1}, Z)) \text{UnsignedArea}(p_i, p_{i+1}, Z), \quad (7.4.4)$$

where the unsigned area of each spherical triangle (x_0, x_1, x_2) is computed as,

$$\text{UnsignedArea}(x_0, x_1, x_2) = -\pi + \sum_{i \in \mathbb{Z}_3} \arccos \left(\frac{x_{i-1} \times x_i}{|x_{i-1} \times x_i|} \cdot \frac{x_i \times x_{i+1}}{|x_i \times x_{i+1}|} \right).$$

This formula can also be recovered by the Hopf fibration. Setting the lift $q_i := \text{Dihedral}(Z, p_i)$ for each p_i , we obtain (7.4.4). Numerically, this formula is unstable if any of the vertices is close to Z or $-Z$, which is explained by the numerical sensitivity of $\text{Dihedral}(Z, p_i)$ and $\frac{p_i \times Z}{\|p_i \times Z\|}$.

7.5 Numerical examples

In this section, we present numerical examples of area computation for spherical polygons using our formula (7.3.1). Moreover, we demonstrate how this formula can also be utilized to determine the total torsion of a space curve, which differs from 2π exactly by the enclosed area of the spherical curve traced out by the tangents of the curve. By comparing the results obtained using both our formula and the classical formula (7.1.1), we show that our formula produces consistent and converging solutions, even for singular curves. This improved numerical robustness allows for more accurate measurements of spherical areas and total torsion.

We employ the horizontal lift approach (Corollary 7.3.5) for computation in all of our examples. For a given closed spherical curve $\gamma: [0, 2\pi) \rightarrow \mathbb{S}^2$, we use a uniform division $\{t_i := \frac{2\pi i}{n}\}_{i=0}^{n-1}$ of the interval $[0, 2\pi)$ to specify the vertices $\{\gamma(t_i)\}_{i=0}^{n-1}$ with some positive integer n . This process turns the spherical curve into a spherical polygon.

7.5.1 Spherical cardioid

We compute the area of a spherical curve γ given by the stereographic projection image $\gamma = P \circ \gamma_{\mathbb{R}^2}$ of a planar cardioid $\gamma_{\mathbb{R}^2}$ (Figure 7.3). Explicitly, the planar cardioid is parametrically

given by

$$\gamma_{\mathbb{R}^2}(t) = (2(1 - \cos(t)) \cos(t), 2(1 - \cos(t)) \sin(t)),$$

and the stereographic projection from the plane to the sphere is

$$P: (x, y) \mapsto \frac{1}{x^2 + y^2 + 1} (2x, 2y, x^2 + y^2 - 1).$$

We compute the area with various numbers n of vertices using our formula and the classical formula (7.1.1,7.1.2). Figure 7.4 (left) shows their numerical results. Note that this spherical cardioid has a cusp, *i.e.* $\partial_t \gamma$ changes sign at $t = 0$. As the polygon refines ($n \rightarrow \infty$), the edge lengths adjacent to $\gamma(t_0)$ decrease to zero superlinearly, and $\frac{\gamma(t_{n-1}) \times \gamma(t_0)}{|\gamma(t_{n-1}) \times \gamma(t_0)|} \cdot \frac{\gamma(t_0) \times \gamma(t_1)}{|\gamma(t_0) \times \gamma(t_1)|} \rightarrow -1$. These conditions make the classical formula numerically unstable as observed in Figure 7.4, left. In contrast, our formula is numerically stable despite the presence of the cusp.



Figure 7.3: Cardioid (left) stereographically projected on the sphere (right).

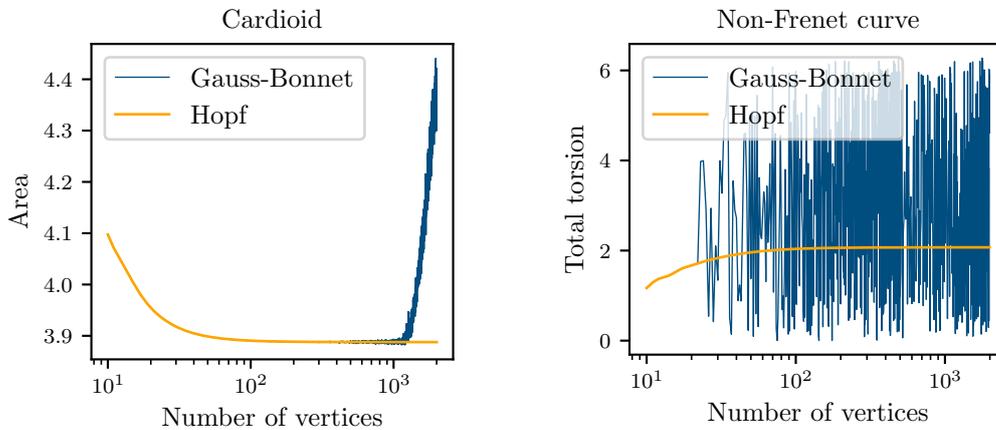


Figure 7.4: Numerical values of the signed areas of spherical curves discretized into spherical polygons with increasing number of vertices, computed using the classical formula (Gauss–Bonnet) and our formula (Hopf). Both formulae give consistent values when the number of vertices is small, but the classical formula becomes unstable as the number of vertices increases. Left: The area enclosed by a spherical cardioid (Section 7.5.1). Right: The total torsion ($2\pi - \text{Area}(\gamma')$) of a non-Frenet space curve (Section 7.5.3).

7.5.2 Total torsions of space curves

Our next examples are about total torsions of closed space curves. For a space curve $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$, the total torsion [PG24, Ch. 1 Sec. 5.1] (equivalently, the total writhe up to a



Figure 7.5: Figure-eight knot (left) and its unit velocity map (right).



Figure 7.6: Trefoil (left) and its unit velocity map (right).

minus sign and a multiple of 2π) can be evaluated as,

$$\text{Torsion}(\gamma) = 2\pi - \text{Area}(\gamma') \pmod{2\pi}, \quad (7.5.1)$$

where $\text{Area}(\gamma')$ is the signed area of the unit velocity map given by $\gamma' = \partial_t \gamma / |\partial_t \gamma|$. This notion of total torsion also works for a space polygon $\{\gamma(t_i)\}_i$. For a space polygon, the unit velocity γ' is given by the normalized edge vector

$$\gamma'(t_i) = \frac{\gamma(t_{i+1}) - \gamma(t_i)}{|\gamma(t_{i+1}) - \gamma(t_i)|},$$

which forms a spherical polygon, whose signed area can be evaluated by our formula. We compute the total torsions of the figure-eight knot (Figure 7.5)

$$\gamma(t) = ((2 + \cos(2t)) \cos(3t), (2 + \cos(2t)) \sin(3t), \sin(4t)),$$

and the trefoil knot (Figure 7.6),

$$\gamma(t) = (\sin(t) + 2 \sin(2t), \cos(t) - 2 \cos(2t), -\sin(3t)).$$

With sufficiently many vertices, our results converge to numbers that agree with the results in a previous study [NVR21]: -0.5423 of the figure-eight knot and 2.2250 of the trefoil.

7.5.3 Total torsion of a non-Frenet curve

In this example we compute the total torsion of a regular closed space curve with a singular (infinitely oscillatory) Frenet–Serret frame. Note that the mod- 2π total torsion (7.5.1) and the writhe only require the curve to be regular *i.e.*, $\partial_t \gamma(t) \neq 0$. In particular, the curve does not need to possess a regular Frenet torsion. A smooth regular curve without a regular Frenet–Serret frame is called a non-Frenet curve (See [PG24, Ch. 1 Sec. 5.6] for a detailed discussion). Consider the following example of a smooth non-Frenet closed space curve

$$\gamma(t') = \frac{1}{e^{-2t'^2} + t'^2} \left(e^{-t'^2} \cos(e^{t'}), e^{-t'^2} \sin(e^{t'}), t' \right),$$

where $t' \in [-\infty, \infty)$ is a reparametrization of $t \in [0, 2\pi)$ by $t' = \tan\left(\frac{t-\pi}{2}\right)$. The spherical curve γ' traced out by the unit velocity displays an exponential spiral about $t = \pi$ with an infinite turning number and an unbounded geodesic curvature (Figure 7.7). Note that the total Frenet torsion of γ is the total turning angle of the spherical curve γ' , which is divergent. Despite the divergence of the total Frenet torsion, the mod- 2π torsion is well-defined since the area enclosed by γ' is bounded.



Figure 7.7: A curve with a spiral Frenet-Serre frame (left), its unit velocity map (middle), and the close-up view of the spiral (right).

In this example, it is crucial to avoid the classical angle-based formula (7.1.1–7.1.2) for evaluating $\text{Area}(\gamma')$ due to the divergent turning angle in γ' . In fact, evaluating the total torsion using the classical formula is equivalent to sampling and summing the Frenet torsion (exterior angle of the spherical polygon γ'). The process produces a result that diverges as the number of sample points $n \rightarrow \infty$ (Figure 7.4, right). In contrast, our formula (7.3.1) is able to robustly evaluate the total torsion of this non-Frenet curve.

7.5.4 Area of a region on the earth

The next example applies to geography. We compute an approximate area of Austria using the data from the Database of Global Administrative Areas (GADM) [are12]. The data contains a sequence of latitudes and longitudes (Figure 7.8) forming a spherical polygon. We treat the earth as a round sphere while acknowledging that we neglect its slight ellipsoidal figure and terrains.

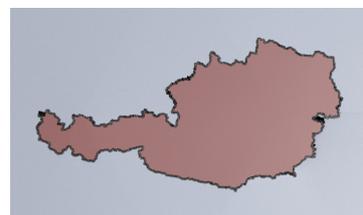


Figure 7.8: Austria plotted on the sphere.

The solid angle value we computed via our formula was 2.06206×10^{-3} on the unit sphere. By multiplying the square of the arithmetic mean radius $R := (2R_E + R_P)/3 \approx 6,371$ km [Mor00], where R_E, R_P are the equatorial and polar radii, we obtain $83,882 \text{ km}^2$, which is a descent approximation of the official area $83,871 \text{ km}^2$ with 0.013% relative error.

7.5.5 Solid angle fields and Seifert surfaces

Our last example demonstrates the construction of the solid angle fields of given space curves (Section 7.1). We consider three rectangular loops linked into the topological configuration of Borromean rings (Figure 7.9). Let $\tilde{\gamma}: \sqcup^3 \mathbb{S}^1 \rightarrow \mathbb{R}^3$ denote this triplet of space polygons. For each point $x \in \mathbb{R}^3 \setminus \tilde{\gamma}$, we let $\Omega(x)$ be half the solid angle subtended by $\tilde{\gamma}$ at x . Explicitly, consider the spherical curves $\gamma^x: \sqcup^3 \mathbb{S}^1 \rightarrow \mathbb{S}^2$ given by projecting $\tilde{\gamma}$ on the unit sphere centered at x :

$$\gamma^x(s) := \frac{\tilde{\gamma}(s) - x}{|\tilde{\gamma}(s) - x|}, \quad s \in \sqcup^3 \mathbb{S}^1.$$

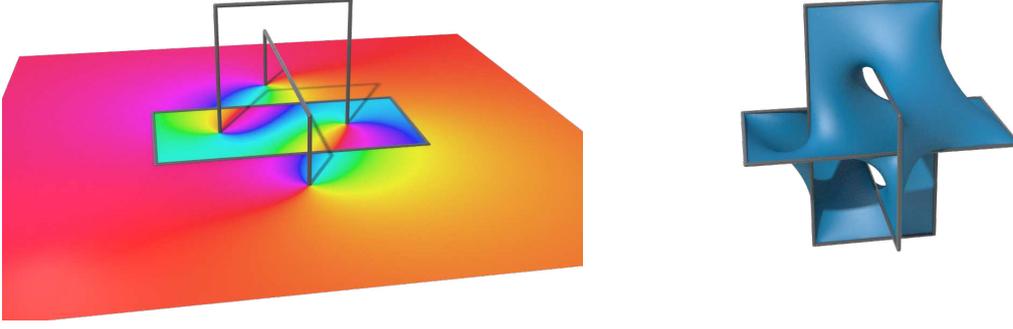


Figure 7.9: The solid angle field of the Borromean rings visualized on the $z = 0$ plane (left) where each color corresponds to a value in \mathbb{S}^1 . A levelset of the solid angle field in \mathbb{R}^3 is a Seifert surface (right).

The solid angle field $\Omega: \mathbb{R}^3 \setminus \tilde{\gamma} \rightarrow \mathbb{R}/(2\pi\mathbb{Z})$ is defined by

$$\Omega(x) := \frac{1}{2} \text{Area}(\gamma^x).$$

Note that the projected spherical curve γ^x is degenerate for x on any extended tangent line of $\tilde{\gamma}$. Despite this unavoidable degeneracy, our formula robustly handle the solid angle computation for all x (Figure 7.9, left). By extracting a levelset of the solid angle field, we construct a smooth Seifert surface (Figure 7.9, right).

Concluding remark and outlook

In this chapter, we derived area formulae using Green's theorem on prequantum bundles \mathbb{S}^3 and $\mathrm{SO}(3)$ over \mathbb{S}^2 . These formulae avoid relying on angle calculation, unlike the classical formula that fails on degenerate cases. As the prequantum version of Green's theorem is available for any compact symplectic manifold (Remrak 7.2.2), one may investigate area formulae or integral of symplectic form of polygons in other manifolds.

For example, in quantum information; the complex projective space $\mathbb{C}\mathbb{P}^{2^n-1}$ is regarded as the space of possible states of n -qubits [BZ06] and a closed path is a periodic orbit. We hope that finding an explicit expression of its enclosed area (geometric phase) may lead to practical applications in quantum computation.

7.6 Appendix: Prequantum bundles \mathbb{S}^3 and $\mathrm{SO}(3)$ over \mathbb{S}^2

Proof of Proposition 7.2.5. In terms of the quaternion coordinate, the area form of \mathbb{S}^2 is written as

$$\sigma = -\frac{1}{2} \text{Re}(p \text{Im}(dp \wedge dp)).$$

On the other hand, the differential of $\pi = q\mathfrak{i}\bar{q}$ is $d\pi = dq\mathfrak{i}\bar{q} + q\mathfrak{i}d\bar{q}$. Hence, the pullback area form is computed as

$$\begin{aligned}\pi^*\sigma &= -\frac{1}{2}\operatorname{Re}(\pi\operatorname{Im}(d\pi \wedge d\pi)) = -\frac{1}{2}\operatorname{Re}(q\mathfrak{i}\bar{q}\operatorname{Im}(d\pi \wedge d\pi)) \\ &= -\frac{1}{2}\operatorname{Re}(\mathfrak{i}\bar{q}\operatorname{Im}(d\pi \wedge d\pi)q) = -\frac{1}{2}\operatorname{Re}(\mathfrak{i}\operatorname{Im}(\bar{q}d\pi \wedge d\pi q)) \\ &= -\frac{1}{2}\operatorname{Re}(\mathfrak{i}\operatorname{Im}((\bar{q}dq\mathfrak{i}\bar{q} + \mathfrak{i}d\bar{q}) \wedge (dq\mathfrak{i} + q\mathfrak{i}d\bar{q}q))) \\ &= -\frac{1}{2}\operatorname{Re}(\mathfrak{i}\bar{q}dq\mathfrak{i} \wedge \bar{q}dq\mathfrak{i} + \mathfrak{i}dq \wedge d\bar{q} - \mathfrak{i}d\bar{q} \wedge dq - d\bar{q}q\mathfrak{i} \wedge d\bar{q}q),\end{aligned}$$

which agrees with

$$d\alpha = -2\operatorname{Re}(dq \wedge \mathfrak{i}d\bar{q}).$$

Here, we have applied $d\bar{q}q = d|q|^2 - \bar{q}dq = -\bar{q}dq$, which holds on \mathbb{S}^3 where $|q|^2 = 1$. \square

Proof of Proposition 7.4.1. Let $V, W \in \mathfrak{so}(3)$, and $\omega = (\omega_1, \omega_2, \omega_3)$, $\nu = (\nu_1, \nu_2, \nu_3)$ be their coefficients with respect to the standard basis of $\mathfrak{so}(3)$ as in Section 7.4. For $PV, PW \in T_P\operatorname{SO}(3)$ on each $P \in \operatorname{SO}(3)$, we have,

$$d\eta|_P(PV, PW) = -\eta|_P([PV, PW]) = -\eta|_I([V, W]) = \nu_2\omega_3 - \omega_2\nu_3,$$

due to the left-invariance of the vector fields PV, PW under $\operatorname{SO}(3)$. Here $[\cdot, \cdot]$ denotes the Lie bracket. Now we compute $\pi_2^*\sigma$. Let us write W, V and P column-wise as $W = (w_1 w_2 w_3)$, $V = (v_1 v_2 v_3)$, and $P = (p_1 p_2 p_3)$. We have,

$$\begin{aligned}\pi_2^*\sigma|_P(PV, PW) &= \sigma|_{p_1}(d\pi_2(PV), d\pi_2(PW)) = \sigma|_{p_1}(Pv_1, Pw_1) \\ &= \sigma|_{\mathfrak{i}}(v_1, w_1) = dy \wedge dz(v_1, w_1) = \nu_2\omega_3 - \omega_2\nu_3,\end{aligned}$$

where we used the invariance of σ under $\operatorname{SO}(3)$ and the expression of σ using the Cartesian coordinates. \square

Postface: Mathematics, physics, and computer graphics

All models are wrong, but some are useful.

— George E. P. Box

This is a personal essay, almost entirely unrelated to the main contents of this thesis. It might be more like a piece of prose, collecting unorganized thoughts.

The beginning of my PhD was when I finished my master's study in mathematics with poor grades and no chance of continuing to a PhD. So, it actually didn't begin. Then I spent some time working in the optics and game industries. Around that time, I came across the field of computer graphics. After some time, I started a PhD and now complete it, with research in the union of mathematics, physics, and computer graphics.

Through this journey, I've got some thoughts about mathematics, physics, and computer graphics. I'd like to leave some of them here. My path had some twists, but looking back, I think my aim has always been to express how I see nature and the world in a mathematical way. In this regard, I've been looking at math and CG through their connection to physics. Additionally, math and CG are often regarded as opposites, like the "pure" and "applied" ends of science, but I think they actually share a lot in common. I'd like to write a bit about some aspects that I think math and CG share.

No rigid definitions First, when it comes to what mathematics and computer graphics actually are, I think there are no rigid definitions, perhaps even more so than many other fields. I'll briefly describe my view of modern computer graphics, because outside of this field, people often have little idea of what this can be. I probably wouldn't have known it too without doing research in graphics. Modern graphics is a much broader field than many people imagine. It's not just about creating gorgeous imageries that appear in movies of Pixar or Disney, although it is an important part.

For instance, research of fluids, like water and smoke, in graphics can involve many directions and tastes. It is probably easy to imagine that improving numerical methods for existing fluid equations, like the Navier–Stokes equation, can be a research topic. But that's not all there is to it. [YNW⁺23, ZYC25] pointed out that the existing formulation of the Navier–Stokes equations in terms of fluid vorticity was incorrect, and derived the correct formulation

that aligns with the velocity-based formulation. As a consequence, they found closed-form expressions for non-stationary solutions of fluids. Another valid direction is proposing new equations of motion for physical phenomena. For instance, I geometrically formulated the dynamics of soap bubbles and films [IYAH17]. Many people would probably view these types of work as mathematics or physics. But in fact, they can stand as computer graphics work. Basically, if a paper involves either computer or graphics in some way, it can be a computer graphics work. In my opinion, even that condition isn't strictly necessary. ¹

Similarly, "mathematics" is just a vague and undefined jargon. For me mathematics isn't really a field, but rather an activity or a practice. It sets up axioms then deduces what can happen within them through logic. It also involves describing vaguely perceived concepts in a language that everyone understands in a unique way, or developing that universal language, like Newton and Leibniz introduced *differentiation*, a language which describes infinitesimal changes of quantities. I think all these activities together are what is called mathematics. Therefore, mathematics appears in many works that aren't "mathematics" papers. A mathematics paper is just one in which the claimed main contribution is this activity itself. In the first place, whether a piece of work counts as a "math" paper or not may hardly matter mathematically. ²

Lenses for physics Since ancient times, physics research has been conducted by observing natural phenomena, making hypotheses to explain them, and verifying them through experiments. Much of modern physics research is, however, less intuitive to me as it often requires interpreting or reasoning the numbers obtained in experiments. In the old physics, on the other hand, the results of experiments were more direct and clearer. For example, Michael Faraday actually *saw* his discovery when a coil was rotated by the invisible magnetic field provoked by an electric field. For me, physics simulation in CG can bring back this intuitive aspect of physics research, allowing us to visually see the results of experiments based on the theoretical models we build.

But unfortunately, it seems to me that the recent CG industry is getting trapped by the curse of "realistic" or "physically-based". Here "physically-based" just means, deriving from or combining existing physics equations. Creating new equations out of nowhere, based on observing natural phenomena, is "not physically accurate". Yet that is precisely what physics is about—observing nature and extracting its simple essence. ³ I feel like the papers by Newton or Faraday, who modeled natural phenomena as principles via observation rather than deriving from existing equations, would likely be rejected in today's CG venues as "physically unplausible". Being

¹In practice, however, as of 2025, papers in computer graphics journals seem to me largely limited to work that presents a *method* achieving or solving something. It seems very difficult to publish a paper about pure discoveries that are common in other fields—like finding a new species, observing or predicting interesting physical phenomena (like black holes) through theory, numerical simulations, or physical experiments, or proving a particular underlying law. In the current graphics venues, remarks or proofs are often considered something that should be pushed to appendices rather than the primary contribution. Still, it seems (and I hope) that gradually more people are recognizing such discoveries or interpretations themselves as valid contributions.

²In the past, such categories as "mathematics" or "physics" were quite vague, and there were no jargons like "pure math" or "applied math". Everything was more mixed, and today's standard mathematical rigor was not required, depending on the type of work. So one could be more free in what one can publish as a math work—not only rigorous proofs of statements, but also new concepts without rigorous analysis. Many papers by Poincaré, who introduced what is now called topology, did not meet today's standard rigor, and so they likely wouldn't have places in today's "mathematics" papers. I find this narrowing of math sad.

³For example, this sand model paper [NO93] practices this spirit and was accepted in Physical Review Letters, a top journal in physics.

caught by “physically-based” or “realistic” attitudes appears to me, ironically, the opposite of physical science.

Okay, let’s get back on track. Like CG is a modern experimental ground for physics, mathematics is a sandbox for physics too. Creating and validating physical models, and deducing phenomena from them, works the same way. It’s just a matter of whether the experiments are theoretical or numerical involving visualization. For example, Einstein tried to validate his theory of gravity by calculating the precession of Mercury’s perihelion. I think such a work should fit as a graphics paper today.

No need to stay in the real world In both computer graphics and mathematics, we do not have to be bounded by the laws or phenomena of real-world physics. In game development, for instance, physics can be designed for the sake of entertainment, without following real-world mechanics. Likewise, mathematics allows one to freely explore imaginary objects. So both fields offer the freedom to create and explore fascinating phenomena within fictional worlds governed by their own rules. ⁴

Do not have to be science or technology. Can be just art CG and math don’t even have to be science or technology. They can be just art. In a sense, one could even say that they are arts of fake.

Obviously, computer graphics is an effort for making fake imagery that mimics reality. Mathematics can also be seen as a form of fake physics. Mathematical results are considered very rigorous, but they are so only mathematically. They may not be rigorous at all in different senses.

Many mathematical models start from drastic simplifications and approximations of physical phenomena, yet humans often do not even know to what extent they are justified. They also begin with presumed principles like Newton’s laws, or with hypothetical setups to make them tractable within mathematical frameworks, like elementary particles, Hilbert spaces, or the concept of continuum. But actual natural phenomena do not respect these artificial setups.

The quote by George E. P. Box at the opening of this essay — “*All models are wrong, but some are useful*” — was originally about statistical models, but I feel it applies to many sciences including physics.

And I think this is okay. Models are wrong by definition. They aren’t nature itself. They are, like elementary particles for instance, just theoretical accounts that try to explain a class of (but not all) observations seemingly without contradiction. They are never truth, but they are the only and best we can do to interpret how the world looks and behaves.

Different people see the world in different ways. One can start from — how I see this world — as their own principle and express that view through (fake) arts of science like mathematics and computer graphics.

I hope modern mathematics and computer graphics are and remain grounds that accommodate and celebrate such diversity of perspectives.

A side note: A personal note on my journey Here is a brief note on how I came to have these views.

⁴Just as mentioned above, such freedom seems to have been slipping away as the community gets caught up in “realistic” or “physically-accurate”. It’s a shame that CG is throwing away its virtue.

During my master's studies, I wanted to do research in mathematics and I had some vague ideas. But I had no clue how to start concretely or what I should learn for this. I think I was preoccupied with mathematical expressions and formulas, and following the correctness of logic, like many math students feel they are forced to do so. And like many other students again, I was not able to develop an intuitive understanding of what these concepts are. For instance, I didn't even realize that the Fourier transform are waves. Even though I said I liked mathematics, I guess I had lost sight of how I liked it.

After finishing my master's degree and before starting my PhD, I spent some time in the optics and game industry. There, I learned what mathematical concepts intuitively are, how they look, and how they affect the real world—for example, how the Fourier transform influences pictures seen through a camera lens.

Then, I came across research in computer graphics, in which I could visually perceive how the mathematics I design affects the motion and shapes of objects, and how it builds the laws of a (real and imaginary) world. Through this process, mathematics came to life inside me. Observing motions and shapes of physical phenomena in CG research also brought me many questions that I wanted to investigate. CG turned out not to be just a place to apply fancy mathematics, but a source of mathematical questions and it made me love mathematics again.

Hopefully, this journey goes on, but I wanted to record these thoughts and feelings of mine, at this moment.

Bibliography

- [AA24] Simon Allais and Marie-Claude Arnaud. The dynamics of conformal Hamiltonian flows: dissipativity and conservativity. *Revista matemática iberoamericana*, 40(3):987–1021, 2024.
- [Ada93] Masahisa Adachi. *Embeddings and immersions. Translated from the Japanese*, volume 124 of *Transl. Math. Monogr.* Providence, RI: American Mathematical Society, 1993.
- [AK21] Vladimir I. Arnold and Boris A. Khesin. *Topological methods in hydrodynamics.* Applied Mathematical Sciences. Springer, Cham, second edition, 2021.
- [Alb06] G. Alberti. Geometric measure theory. In *Encyclopedia of Mathematical Physics*, pages 520–528. Academic Press, Oxford, 2006.
- [AN05] Alexis Angelidis and Fabrice Neyret. Simulation of smoke based on vortex filament primitives. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 87–96, 2005.
- [Ang17] Alexis Angelidis. Multi-scale vorticle fluids. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017.
- [are12] GADM, database of global administrative areas. *Global Administrative Areas*, 2012.
- [Arv95a] James Arvo. Applications of irradiance tensors to the simulation of non-lambertian phenomena. *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 335–342, 1995.
- [Arv95b] James Arvo. Stratified sampling of spherical triangles. *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 1995.
- [AS96] Luigi Ambrosio and Halil Mete Soner. Level set approach to mean curvature flow in arbitrary codimension. *Journal of differential geometry*, 43(4):693–737, 1996.
- [BA18] Jack Binysh and Gareth P Alexander. Maxwell’s theory of solid angle and the construction of knotted fields. *Journal of Physics A: Mathematical and Theoretical*, 51(38):385202, aug 2018.
- [BBH⁺94] Fabrice Bethuel, Haïm Brezis, Frédéric Hélein, et al. *Ginzburg-landau vortices*, volume 13. Springer, 1994.

- [BBHM12] Martin Bauer, Martins Bruveris, Philipp Harms, and Peter W Michor. Vanishing geodesic distance for the riemannian metric with geodesic equation the KdV-equation. *Annals of Global Analysis and Geometry*, 41:461–472, 2012.
- [BBHMA17] Martin Bauer, Martins Bruveris, Philipp Harms, and Jakob Moller-Andersen. A numerical framework for Sobolev metrics on the space of curves. *SIAM Journal on Imaging Sciences*, 10(1):47–73, 2017.
- [BBM14] Martin Bauer, Martins Bruveris, and Peter W Michor. Overview of the geometries of shape spaces and diffeomorphism groups. *Journal of Mathematical Imaging and Vision*, 50:60–97, 2014.
- [BC87] Michael G. Bevis and Greg Cambareri. Computing the area of a spherical polygon of arbitrary shape. *Mathematical Geology*, 19:335–346, 1987.
- [BCMO01] Paul Burchard, Li-Tien Cheng, Barry Merriman, and Stanley Osher. Motion of curves in three spatial dimensions using a level set approach. *Journal of Computational Physics*, 170(2):720–741, 2001.
- [BDR20] Maciej Borodzik, Supredee Dangskul, and Andrew Ranicki. Solid angles and seifert hypersurfaces. *Annals of Global Analysis and Geometry*, 57(3):415–454, 2020.
- [BFM16] Ashish Bhatt, Dwayne Floyd, and Brian E Moore. Second order conformal symplectic schemes for damped Hamiltonian systems. *Journal of Scientific Computing*, 66(3):1234–1259, 2016.
- [BH15] Martin Bauer and Philipp Harms. Metrics on spaces of immersions where horizontality equals normality. *Differential Geometry and its Applications*, 39:166–183, 2015.
- [BIM24] Martin Bauer, Sadashige Ishida, and Peter W. Michor. Symplectic structures on the space of space curves. *Preprint. arXiv:2407.19908*, 2024.
- [BKB12] Tyson Brochu, Todd Keeler, and Robert Bridson. Linear-time smoke animation with vortex sheet meshes. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 87–95. Citeseer, 2012.
- [Bob15] Alexander I Bobenko. *Geometry II: Discrete differential geometry*. 2015.
- [Bry09] J.L. Brylinski. *Loop Spaces, Characteristic Classes and Geometric Quantization*. Modern Birkhäuser Classics. Birkhäuser Boston, 2009.
- [BW58] W. M. Boothby and H. C. Wang. On contact manifolds. *Annals of Mathematics*, 68(3):721–734, 1958.
- [BW97] S. Bates and A. Weinstein. *Lectures on the Geometry of Quantization*. Berkeley mathematics lecture notes. American Mathematical Society, 1997.
- [BWR⁺08] Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. Discrete elastic rods. 2008.
- [BZ06] Ingemar Bengtsson and Karol Zyczkowski. *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Cambridge University Press, 2006.

- [Che17] Albert Chern. *Fluid dynamics with incompressible Schrödinger flow*. PhD thesis, California Institute of Technology, 2017.
- [Cho90] Alexandre Joel Chorin. Hairpin removal in vortex interactions. *Journal of Computational Physics*, 91(1):1–21, 1990.
- [CI24] Albert Chern and Sadashige Ishida. Area formula for spherical polygons via prequantization. *SIAM Journal on Applied Algebra and Geometry*, 8(3), 2024.
- [CI25] Albert Chern and Sadashige Ishida. Implicit representations of codimension-2 submanifolds and their prequantum structures. *Preprint. arXiv:2507.11727*, 2025.
- [CK⁺00] Georges-Henri Cottet, Petros D Koumoutsakos, et al. *Vortex methods: theory and practice*, volume 8. Cambridge university press Cambridge, 2000.
- [CKP⁺16] Albert Chern, Felix Knöppel, Ulrich Pinkall, Peter Schröder, and Steffen Weißmann. Schrödinger’s smoke. *ACM Trans. Graph.*, 35(4), jul 2016.
- [CKPP20] Albert Chern, Felix Knöppel, Franz Pedit, and Ulrich Pinkall. *Commuting Hamiltonian Flows of Curves in Real Space Forms*, volume 1 of *London Mathematical Society Lecture Note Series*, page 291–328. Cambridge University Press, 2020.
- [CKPS17] Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. Inside fluids: Clebsch maps for visualization and processing. *ACM Transactions on Graphics (TOG)*, 36(4):1–11, 2017.
- [Cle59] A. Clebsch. Ueber die integration der hydrodynamischen gleichungen. *Journal für die reine und angewandte Mathematik*, 56:1–10, 1859. English translation by D. H. Delphenich, http://www.neo-classical-physics.info/uploads/3/4/3/6/34363841/clebsch_-_clebsch_variables.pdf.
- [CMM91] Vicente Cervera, Francisca Mascaro, and Peter W Michor. The action of the diffeomorphism group on the space of immersions. *Differential Geometry and its Applications*, 1(4):391–401, 1991.
- [CSdR12] S.S. Chern, F.R. Smith, and G. de Rham. *Differentiable Manifolds: Forms, Currents, Harmonic Forms*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2012.
- [Dan16] Supreede Dangsukul. *Construction of Seifert surfaces by differential geometry*. PhD thesis, The University of Edinburgh, 2016.
- [DBWG15] Fang Da, Christopher Batty, Chris Wojtan, and Eitan Grinspun. Double bubbles sans toil and trouble: Discrete circulation-preserving vortex sheets for soap films and foams. *ACM Transactions on Graphics (TOG)*, 34(4):1–9, 2015.
- [Dem97] J.P. Demailly. *Complex Analytic and Differential Geometry*. Université de Grenoble I, 1997.
- [DJNV20] Tobias Diez, Bas Janssens, Karl-Hermann Neeb, and Cornelia Vizman. Induced differential characters on nonlinear Graßmannians. *Annales de l’Institut Fourier*, 2020.

- [dR84] G. de Rham. *Differentiable Manifolds: Forms, Currents, Harmonic Forms*. Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen mit besonderer Berücksichtigung der Anwendungsgebiete. Springer-Verlag, 1984.
- [DWCWR21] L Sumners De Witt, Irma I Cruz-White, and Renzo L Ricca. Zero helicity of seifert framed defects. *Journal of Physics A: Mathematical and Theoretical*, 54(29):295203, 2021.
- [ETK⁺07] Sharif Elcott, Yiyong Tong, Eva Kanso, Peter Schröder, and Mathieu Desbrun. Stable, circulation-preserving, simplicial fluids. *ACM Transactions on Graphics (TOG)*, 26(1):4–es, 2007.
- [Etn05] John B. Etnyre. Lectures on open book decompositions and contact structures. 2005. <https://arxiv.org/abs/math/0409402>.
- [Fed14] H. Federer. *Geometric Measure Theory*. Classics in Mathematics. Springer Berlin Heidelberg, 2014.
- [FOM06] Shaojun Feng, Washington Y Ochieng, and Rainer Mautz. An area computation based method for raim holes assessment. *Journal of Global Positioning Systems*, 5(1-2):11–16, 2006.
- [Gei08] Hansjörg Geiges. *An Introduction to Contact Topology*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2008.
- [GLG95] Manuel Noronha Gamito, Pedro Faria Lopes, and Mário Rui Gomes. Two-dimensional simulation of gaseous phenomena using vortex particles. In *Computer Animation and Simulation'95*, pages 3–15. Springer, 1995.
- [Gon10] Álvaro González. Measurement of areas on a sphere using fibonacci and latitude–longitude lattices. *Mathematical geosciences*, 42:49–64, 01 2010.
- [Han98] J. H. Hannay. The majorana representation of polarization, and the berry phase of light. *Journal of Modern Optics*, 45(5):1001–1008, 1998.
- [Hen09] Falk-Florian Henrich. *Loop spaces of Riemannian manifolds*. Doctor thesis, Technische Universität Berlin, 2009.
- [Hir12] M.W. Hirsch. *Differential Topology*. Graduate Texts in Mathematics. Springer New York, 2012.
- [HV03] Stefan Haller and Cornelia Vizman. Non-linear grassmannians as coadjoint orbits. *Mathematische Annalen*, 329:771–785, 2003.
- [IL24] Sadashige Ishida and Hugo Lavenant. Quantitative convergence of a discretization of dynamic optimal transport using the dual formulation. *Foundations of Computational Mathematics*, 2024.
- [ISN⁺20] Sadashige Ishida*, Peter Synak*, Fumiya Narita, Toshiya Hachisuka, and Chris Wojtan. A model for soap film dynamics with evolving thickness. *ACM Transactions on Graphics*, 39(4), 2020.
- [IWC22] Sadashige Ishida, Chris Wojtan, and Albert Chern. Hidden degrees of freedom in implicit vortex filaments. *ACM Transactions on Graphics*, 41(6), 2022.

- [IYAH17] Sadashige Ishida, Masafumi Yamamoto, Ryoichi Ando, and Toshiya Hachisuka. A hyperbolic geometric flow for evolving films and foams. *ACM Transactions on Graphics (TOG)*, 36(6):1–11, 2017.
- [JS18] Robert L Jerrard and Didier Smets. Leapfrogging vortex rings for the three dimensional gross-pitaevskii equation. *Annals of PDE*, 4:1–48, 2018.
- [Khe12] Boris Khesin. Symplectic structures and dynamics on vortex membranes. *Moscow Mathematical Journal*, 12, 01 2012.
- [KM80] Evgenii A Kuznetsov and Aleksandr V Mikhailov. On the topological meaning of canonical clebsch variables. *Physics Letters A*, 77(1):37–38, 1980.
- [KM97] A. Kriegel and P.W. Michor. *The Convenient Setting of Global Analysis*. Mathematical Surveys. American Mathematical Society, 1997.
- [KN63] S. Kobayashi and K. Nomizu. *Foundations of Differential Geometry*. Number Bd. 1 in Foundations of Differential Geometry. Interscience Publishers, 1963.
- [Kob63] Shoshichi Kobayashi. Topology of positively pinched Kaehler manifolds. *Tohoku Mathematical Journal*, 15(2):121 – 139, 1963.
- [Kos70] Bertram Kostant. Quantization and unitary representations. In C. T. Taam, editor, *Lectures in Modern Analysis and Applications III*, pages 87–208, Berlin, Heidelberg, 1970. Springer Berlin Heidelberg.
- [KSS00] M M G Krishna, Joseph Samuel, and Supurna Sinha. Brownian motion on a sphere: distribution of solid angles. *Journal of Physics A: Mathematical and General*, 33(34):5965, sep 2000.
- [Lam95] Horace Lamb. *Hydrodynamics*. Cambridge University Press, 1895.
- [LC87] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph.*, 21(4):163–169, aug 1987.
- [Lee12] J. Lee. *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer New York, 2012.
- [Lem93] László Lempert. Loop spaces as complex manifolds. *Journal of Differential Geometry*, 38:519–543, 1993.
- [MAB19] Ken Museth, Nick Avramoussis, and Dan Bailey. Openvdb. In *ACM SIGGRAPH 2019 Courses*, pages 1–56. 2019.
- [MB01] Andrew J. Majda and Andrea L. Bertozzi. *Vorticity and Incompressible Flow*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2001.
- [Mic84] Peter Michor. A convenient setting for differential geometry and global analysis ii. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 25(2):113–178, 1984.
- [Mic08] Peter W. Michor. *Topics in differential geometry*, volume 93 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2008.

- [Mic19] Peter W. Michor. Manifolds of mappings for continuum mechanics. *Advances in Mechanics and Mathematics*, 2019.
- [Min04] Chohong Min. Local level set method in high dimension and codimension. *Journal of computational physics*, 200(1):368–382, 2004.
- [MM05] Peter W Michor and David Mumford. Vanishing geodesic distance on spaces of submanifolds and diffeomorphisms. *Documenta Mathematica*, 10:217–245, 2005.
- [MM06] Peter W. Michor and David Mumford. Riemannian geometries on spaces of plane curves. *Journal of the European Mathematical Society*, 008(1):1–48, 2006.
- [MM07] Peter W. Michor and David Mumford. An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. *Applied and Computational Harmonic Analysis*, 23(1):74–113, 2007. Special Issue on Mathematical Imaging.
- [MM13] David Mumford and Peter W. Michor. On Euler’s equation and ‘EPDiff’. *Journal of Geometric Mechanics*, 5(3):319–344, 2013.
- [Mor98] Philip J Morrison. Hamiltonian description of the ideal fluid. *Reviews of modern physics*, 70(2):467, 1998.
- [Mor00] Geodetic Reference System 1980 by H. Moritz. *Journal of Geodesy*, 74(1):128–162, March 2000.
- [Mor08] F. Morgan. *Geometric Measure Theory: A Beginner’s Guide*. Academic Press, 2008.
- [MS17] Stefano Marò and Alfonso Sorrentino. Aubry–Mather theory for conformally symplectic systems. *Communications in Mathematical Physics*, 354(2):775–808, 2017.
- [Mur96] Kunio Murasugi. *Knot theory & its applications*. Modern Birkhäuser classics. Birkhäuser, Boston, 2008 - 1996.
- [MW83] Jerrold Marsden and Alan Weinstein. Coadjoint orbits, vortices, and Clebsch variables for incompressible fluids. *Physica D: Nonlinear Phenomena*, 7(1):305–323, 1983.
- [MZ96] John J. Millson and Brett Zombro. A Kähler structure on moduli space of isometric maps of a circle into Euclidean space. *Inventiones mathematicae*, 123(1):35–60, 1996.
- [Nak03] M. Nakahara. *Geometry, Topology and Physics, Second Edition*. Graduate student series in physics. Taylor & Francis, 2003.
- [NCR21] Mohammad Sina Nabizadeh, Albert Chern, and Ravi Ramamoorthi. Kelvin transformations for simulations on infinite domains. *ACM Transactions on Graphics (TOG)*, 40(4):97:1–97:15, 2021.
- [Nee18] Tom Needham. Kähler structures on spaces of framed curves. *Annals of Global Analysis and Geometry*, 54(1):123–153, 2018.

- [Nic11] Liviu I Nicolaescu. The coarea formula. In *seminar notes*. Citeseer, 2011.
- [NO93] Hiraku Nishimori and Noriyuki Ouchi. Formation of ripple patterns and dunes by wind-blown sand. *Physical Review Letters*, 71(1):197, 1993.
- [NVR21] Marija S Najdanović, Ljubica S Velimirović, and Svetozar R Rančić. The total torsion of knots under second order infinitesimal bending. *Applicable Analysis and Discrete Mathematics*, 15(2):283–294, 2021.
- [OS88] Stanley Osher and James A Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations. *Journal of computational physics*, 79(1):12–49, 1988.
- [OSWW23] Shinya Okabe, Philip Schrader, Valentina Wheeler, and Glen Wheeler. A Sobolev gradient flow for the area-normalised Dirichlet energy of H^1 maps. *arXiv preprint arXiv:2310.05459*, 2023.
- [OTH02] Shin-ichiro Ogawa, Makoto Tsubota, and Yuji Hattori. Study of reconnection and acoustic emission of quantized vortices in superfluid by the numerical analysis of the gross-pitaevskii equation. *Journal of the Physical Society of Japan*, 71(3):813–821, 2002.
- [OU13] Joel Oakley and Michael Usher. On certain lagrangian submanifolds of $S^2 \times S^2$ and $\mathbb{C}P^n$. *Algebraic & Geometric Topology*, 16:149–209, 2013.
- [PBS20] David Palmer, David Bommes, and Justin Solomon. Algebraic representations for volumetric frame fields. *ACM Transactions on Graphics (TOG)*, 39(2):1–17, 2020.
- [PCK⁺19] Marcel Padilla, Albert Chern, Felix Knöppel, Ulrich Pinkall, and Peter Schröder. On bubble rings and ink chandeliers. *ACM Trans. Graph.*, 38(4), 2019.
- [PG24] Ulrich Pinkall and Oliver Gross. *Differential Geometry: From Elastic Curves to Willmore Surfaces*. Springer Nature, 2024.
- [Pis99] Len M Pismen. *Vortices in nonlinear fields: from liquid crystals to superfluids, from non-equilibrium patterns to cosmic strings*, volume 100. Oxford University Press, 1999.
- [PK05] Sang Il Park and Myoung Jun Kim. Vortex fluid for gaseous phenomena. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 261–270, 2005.
- [PTG12] Tobias Pfaff, Nils Thuerey, and Markus Gross. Lagrangian vortex sheets for animating fluids. *ACM Transactions on Graphics (TOG)*, 31(4):1–8, 2012.
- [QZG⁺19] Ziyin Qu, Xinxin Zhang, Ming Gao, Chenfanfu Jiang, and Baoquan Chen. Efficient and conservative fluids using bidirectional mapping. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019.
- [RMXO01] Steven J Ruuth, Barry Merriman, Jack Xin, and Stanley Osher. Diffusion-generated motion by mean curvature for filaments. *Journal of Nonlinear Science*, 11:473–493, 2001.

- [Saf93] P. G. Saffman. *Vortex Dynamics*. Cambridge Monographs on Mechanics. Cambridge University Press, 1993.
- [Sei35] H. Seifert. über das geschlecht von knoten. *Math. Ann.*, 110(1):571–592, 1935.
- [SFK⁺08] Andrew Selle, Ronald Fedkiw, Byungmoon Kim, Yingjie Liu, and Jarek Rossignac. An unconditionally stable maccormack method. *Journal of Scientific Computing*, 35(2):350–371, 2008.
- [Sha08] Jayant Shah. H^0 -type Riemannian metrics on the space of planar curves. *Quarterly of Applied Mathematics*, 66(1):123–137, 2008.
- [SK16] Anuj Srivastava and Eric P Klassen. *Functional and shape data analysis*, volume 1. Springer, 2016.
- [SKJJ10] Anuj Srivastava, Eric Klassen, Shantanu H Joshi, and Ian H Jermyn. Shape analysis of elastic curves in euclidean spaces. *IEEE transactions on pattern analysis and machine intelligence*, 33(7):1415–1428, 2010.
- [Sma59] Stephen Smale. The classification of immersions of spheres in Euclidean spaces. *Ann. Math. (2)*, 69:327–344, 1959.
- [Som52] Arnold Sommerfeld. *Electrodynamics: lectures on theoretical physics, vol. 3*, volume 3. Academic Press, 1952.
- [SRF05] Andrew Selle, Nick Rasmussen, and Ronald Fedkiw. A vortex particle method for smoke, water and explosions. In *ACM SIGGRAPH 2005 Papers*, pages 910–914. 2005.
- [SVB17] Justin Solomon, Amir Vaxman, and David Bommes. Boundary element octahedral fields in volumes. *ACM Transactions on Graphics (TOG)*, 36(4):1, 2017.
- [SW89] A. Shapere and F. Wilczek. *Geometric Phases In Physics*. Advanced Series In Mathematical Physics. World Scientific Publishing Company, 1989.
- [SWW23] Philip Schrader, Glen Wheeler, and Valentina-Mira Wheeler. On the $H^1(ds)$ -gradient flow for the length functional. *The Journal of geometric analysis*, 33(9):297, 2023.
- [Tab17] Serge Tabachnikov. On the bicycle transformation and the filament equation: results and conjectures. *Journal of Geometry and Physics*, 115:116–123, 2017.
- [Vai85] Izu Vaisman. Locally conformal symplectic manifolds. *International Journal of Mathematics and Mathematical Sciences*, 8(3):521–536, 1985.
- [Viz11] Cornelia Vizman. Induced differential forms on manifolds of functions. *Archivum Mathematicum*, 047(3):201–215, 2011.
- [VKPS16] Alberto Villois, Giorgio Krstulovic, Davide Proment, and Hayder Salman. A vortex filament tracking method for the gross–pitaevskii model of a superfluid. *Journal of Physics A: Mathematical and Theoretical*, 49(41):415502, 2016.

- [WCZR20] Lifan Wu, Guangyan Cai, Shuang Zhao, and Ravi Ramamoorthi. Analytic spherical harmonic gradients for real-time rendering with many polygonal area lights. *ACM Trans. Graph.*, 39(4), aug 2020.
- [WL98] Maciej P Wojtkowski and Carlangelo Liverani. Conformally symplectic dynamics and symmetry of the Lyapunov spectrum. *Communications in Mathematical Physics*, 194(1):47–60, 1998.
- [WP09] Steffen Weißmann and Ulrich Pinkall. Real-time Interactive Simulation of Smoke Using Discrete Integrable Vortex Filaments. In *Virtual Reality Interactions and Physical Simulation "VRIPHYS" (2009)*, 2009.
- [WP10] Steffen Weißmann and Ulrich Pinkall. Filament-based smoke with vortex shedding and variational reconnection. *ACM Trans. Graph.*, 29(4), jul 2010.
- [WP12] Steffen Weißmann and Ulrich Pinkall. Underwater rigid body dynamics. *ACM Trans. Graph.*, 31(4), jul 2012.
- [WPS14] Steffen Weißmann, Ulrich Pinkall, and Peter Schröder. Smoke rings from smoke. *ACM Trans. Graph.*, 33(4), jul 2014.
- [XTZ⁺21] Shiyong Xiong, Rui Tao, Yaorui Zhang, Fan Feng, and Bo Zhu. Incompressible flow simulation on vortex segment clouds. *ACM Transactions on Graphics (TOG)*, 40(4):98:1–98:11, 2021.
- [YM05] Anthony Yezzi and Andrea Mennucci. Conformal metrics and true "gradient flows" for curves. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 1, pages 913–919. IEEE, 2005.
- [YNW⁺23] Hang Yin, Mohammad Sina Nabizadeh, Baichuan Wu, Stephanie Wang, and Albert Chern. Fluid cohomology. *ACM Trans. Graph.*, 42(4), July 2023.
- [You10] Laurent Younes. *Shapes and diffeomorphisms*, volume 171. Springer, 2010.
- [YXZ⁺21] Shuqi Yang, Shiyong Xiong, Yaorui Zhang, Fan Feng, Jinyuan Liu, and Bo Zhu. Clebsch gauge fluid. *ACM Transactions on Graphics (TOG)*, 40(4):1–11, 2021.
- [ZB14] Xinxin Zhang and Robert Bridson. A PPPM fast summation method for fluids and beyond. *ACM Transactions on Graphics (TOG)*, 33(6):1–11, 2014.
- [ZBG15] Xinxin Zhang, Robert Bridson, and Chen Greif. Restoring the missing vorticity in advection-projection fluid solvers. *ACM Transactions on Graphics (TOG)*, 34(4):1–8, 2015.
- [ZYC25] Cuncheng Zhu, Hang Yin, and Albert Chern. Viscous vortex dynamics on surfaces. *ACM Trans. Graph.*, 44(6), November 2025.

