


Geometric Techniques in PDE Theory, Fluid Dynamics, and Relativity

Lewis William Napper 

Thesis submitted to the University of Surrey
for the degree of Doctor of Philosophy

Supervised by: Dr. Martin Wolf and Prof. Ian Roulstone



*School of Maths, Physics, and Space,
University of Surrey, Guildford, GU2 7XH, United Kingdom.*

23rd December 2024

Copyright © 2024 by Lewis Napper All rights reserved.


E-mail address: lewis.napper@surrey.ac.uk

*For my grandad, Leslie Bennett.
I'm sorry you never got to see me graduate.*

Abstract

This thesis presents a number of recently published results [1–3] in the field of geometry, motivated by the study of partial differential equations arising in fluid dynamics and relativity. In this format, shared background material is consolidated and the results are supplemented by additional exposition.

This thesis is divided into two parts. In the first, we introduce a general framework for studying incompressible Navier–Stokes flows on a Riemannian manifold via Monge–Ampère geometry, inspired by multi-symplectic techniques. In two dimensions, the Poisson equation for the pressure, which acts as a diagnostic equation for the balance of vorticity and rate-of-strain, can be encoded in terms of a Monge–Ampère structure on the cotangent bundle and we show that, in higher dimensions, this structure generalises to a pair of multi-symplectic forms. Submanifolds of the cotangent bundle on which these forms vanish are subsequently studied in lieu of solutions and are equipped with a metric whose signature and curvature are dictated by the accumulation of vorticity. This description admits a (multi-)symplectic reduction principle for three dimensional flows with symmetry and allows topological information about the flow to be deduced. We conclude by discussing ongoing work to classify the higher dimensional generalisation of a Monge–Ampère structure. In the second part of this thesis, we derive analogues of Alexandrov’s patchwork theorem and Toponogov’s theorem for globalising curvature bounds defined by comparison methods, in the rapidly developing field of Lorentzian pre-length spaces. Along the way, we derive a number of supplementary results, including Lorentzian analogues of the Bonnet–Myers theorem and the Lebesgue number lemma. We also highlight several small, but not insignificant, modifications to the definition of a comparison neighbourhood to account for points with infinite time separation (e.g. in Anti-de Sitter space). We conclude by discussing the open problem of the relationship between Lorentzian pre-length spaces, causal sets, and Gromov–Hausdorff convergence.



Lay Summary

The art of geometry has developed immensely since Euclid's *Elements*. Key among these developments was the introduction of non-Euclidean geometry, where, in a plane, a pair of straight lines which are both perpendicular to some third line need not remain a fixed distance apart and may instead converge or diverge. This thesis deals with the application of non-Euclidean geometry to the study of partial differential equations, which relate several variables to their rates of change.

In the first part, we investigate a problem associated with the Navier–Stokes equations for a fluid flow, namely the identification of vortices and turbulence. The Navier–Stokes equations are notoriously resilient to analysis, so much so that proving the existence of solutions to the equations holds a spot on the list of Millennium prize problems. However, by utilising the language of differential geometry, it is possible to encode the Navier–Stokes equations as a Monge–Ampère structure, which constrains two/three-dimensional submanifolds in the four/six-dimensional phase space of the fluid. Properties of these submanifolds, such as their curvature, can then be related to properties of fluid flows, like the balance of vorticity and rate-of-strain.

In the second part, we direct our attention to the problem of describing the geometry associated with low-regularity solutions to Einstein's equations, which can describe gravitational waves, black holes, or cosmic strings. These solutions may not be sufficiently regular to permit differentiation, so while regular solutions correspond to Lorentzian manifolds, that is not the most suitable language here. Fortunately, the recently introduced Lorentzian pre-length spaces allow us to describe Lorentzian geometry via distances and angles, similarly to how metric spaces generalise Riemannian manifolds. As such, we focus on transferring several fundamental results from metric geometry to this new setting, to provide a language with which to study low-regularity solutions.

Acknowledgements

This Ph.D. thesis was supported by the STFC grant ST/V507118/1. The author would like to thank the French Embassy in the United Kingdom for their support in the form of CampusFrance fellowship 156255R during the final stages of this work, as well as the Fields Institute, Toronto, for their hospitality at the *Workshop on Mathematical Relativity, Scalar Curvature and Synthetic Lorentzian Geometry* in October 2022.

The author would like to thank their collaborators, Dr. Martin Wolf, Prof. Ian Roulstone, Prof. Volodya Rubtsov, Dr. Tobias Beran, Dr. Felix Rott, Dr. John Harvey, and Dr. Vladimir Salnikov for being such amenable, hard-working, and supportive colleagues throughout the author's doctoral studies. In particular, thanks go to Dr. Tobias Beran and Dr. Felix Rott for being open to collaborating following the author's spontaneous emails.

Additional thanks go to the author's supervisors, Dr. Martin Wolf and Prof. Ian Roulstone, for their kind advice, guidance, and support on all things academic, bureaucratic, and personal, as well as to Prof. Volodya Rubtsov, Dr. Vladimir Salnikov, and Dr. John Harvey for providing visibility, assistance, and endorsement to the author when engaging with the wider scientific community. The author would also like to thank Prof. Tom Bridges, Dr. Roberto D'Onofrio, Dr. Jan Gutowski, Dr. Giovanni Ortenzi, Dr. Paul Skerritt, and the PGR community at the University of Surrey for many helpful comments and insightful discussions through the years.

Finally, a personal thank you goes to my parents for putting up with the quirks of living with a doctoral student and for supporting me in every way possible in spite of it all. This wouldn't have been possible without you.

Declaration

This thesis was produced during the author's registration on the Mathematics Ph.D. programme at the University of Surrey and has not been submitted in whole or in part for any other academic degree or professional qualification. In submitting this thesis, the author declares that the work contained within is a result of their own efforts and that the demarcation both below and within the text is an accurate reflection of their involvement in any collaborative works contributing to this text.

Part I of this thesis is based on work published in [1] as well as unpublished work with the same authors. The author of this thesis produced all figures and examples for this work and contributed to calculations throughout, focusing in particular on the properties of Lychagin–Rubtsov metrics, generalisations of Monge–Ampère geometry to higher dimensions, and topological implications in the setting of fluid dynamics.

Part II of this thesis is based on work published in [2,3], for which the author proposed the collaboration and its key research questions, in addition to providing appropriate background material, lemmata, and contributions to proofs throughout.

Data and Licence Management

As a theoretical study, no practical data was collected or recorded to produce this report, hence no concerns about the ethical, safe storage/sharing of data arose. There are no immediate ethical concerns about the impact of the computations contained within this report. No additional research data beyond that presented and cited within this work is required to validate the findings herein. Citations to any known relevant sources are provided within the thesis text, bibliography, or footnotes, with any material knowingly drawn from the work of others (whether published or unpublished, including content generated by deep learning/artificial intelligence tools) being fully identified in this manner. For the purposes of open access, the authors apply a Creative Commons Attribution (CC BY) license to this document and any Author Accepted Manuscript version arising from the included material.

Keywords: Monge–Ampère geometry, Navier–Stokes equations, higher symplectic geometry, vortex topology, Lorentzian length spaces, synthetic curvature bounds, globalisation, null distance

AMS Classification Codes: 76D05, 53B35, 53D05, 53D20, 53D12, 53C50, 53C23, 53B30, 51K10, 53C80

Contents

I (Higher) Monge–Ampère Geometry and Fluid Dynamics	1
1. Overview of Part I	3
2. An Introduction to Monge–Ampère Geometry	9
2.1. Symplectic Monge–Ampère Equations and their Solutions	10
2.1.1. Multisymplectic Geometry	10
2.1.2. Monge–Ampère Structures	13
2.1.3. Classical and Generalised Solutions	15
2.2. Almost (Para-)Complex Structures and Lychagin–Rubtsov Metrics	19
2.2.1. Monge–Ampère Structures and Metrics in Two Dimensions	21
2.2.2. Monge–Ampère Structures and Metrics in Higher Dimensions	26
2.3. Chapter Summary	28
3. Monge–Ampère Geometry of Two-Dimensional Flows	29
3.1. The Incompressible Navier–Stokes Equations and Pressure	30
3.1.1. Incompressible Fluid Flows on \mathbb{R}^m and the Weiss Criterion	30
3.1.2. Incompressible Fluid Flows on Manifolds	34
3.2. Geometric Properties of Two-Dimensional Incompressible Fluid Flows	38
3.2.1. Curvature and Topology of Two-Dimensional Incompressible Fluid Flows	42
3.3. Incompressible Fluids on the Euclidean Plane — Examples	46
3.3.1. Simplified Formulæ in the Euclidean Case	46
3.3.2. The Taylor–Green Vortex	48
3.4. Chapter Summary	50
4. (Higher) Monge–Ampère Geometry of Fluid Flows	53
4.1. A Bridge to Higher Dimensions	53
4.2. Geometric Properties of Higher-Dimensional Incompressible Fluid Flows	56
4.2.1. Almost (Para-)Complex Structures For Three-Dimensional Fluid Flows	59
4.2.2. Topology of Three-Dimensional Incompressible Fluid Flows	60
4.3. Reduction of Three-Dimensional Flows with Symmetry	62
4.3.1. General Setup	62
4.3.2. Symplectic Reduction	65
4.3.3. k-Plectic Reduction	68
4.4. Incompressible Fluids in Euclidean Space — Examples	70
4.4.1. Arnold–Beltrami–Childress Flows (\mathbb{R} Symmetry)	70
4.4.2. Hicks–Moffatt Vortex (\mathbb{S}^1 Symmetry)	74
4.5. Chapter Summary	78
5. Part I Outlook	81

II	Globalisation of Curvature Bounds in Lorentzian Length Spaces	93
6.	Overview of Part II	95
7.	A Brief Review of Metric Geometry	101
7.1.	Fundamentals of Metric Geometry	101
7.2.	Globalisation of Curvature Bounds in Metric Spaces	109
7.2.1.	Curvature Bounded Above	110
7.2.2.	Curvature Bounded Below	113
7.3.	Chapter Summary	115
8.	An Introduction to Lorentzian Pre-length Spaces	117
8.1.	Fundamentals of Lorentzian Pre-length Spaces	117
8.1.1.	Topology of Lorentzian Pre-length Spaces	127
8.1.2.	Timelike Curvature Bounds via Triangle Comparison	131
8.2.	Angles in Lorentzian Pre-length Spaces and Curvature Bounds	139
8.2.1.	Comparison and Hyperbolic Angles	139
8.2.2.	Alternative Definitions of Timelike Curvature Bounds	143
8.3.	Chapter Summary	146
9.	Globalisation of Timelike Curvature Bounded Above	149
9.1.	Preliminary Results	149
9.1.1.	The Gluing Lemma for Curvature Bounded Above	150
9.1.2.	Continuously Varying Distance Realisers	151
9.2.	Alexandrov’s Patchwork	153
9.3.	Chapter Summary	162
10.	Globalisation of Timelike Curvature Bounded Below	163
10.1.	Preliminary Results	164
10.1.1.	Time Functions and Null Distance	164
10.1.2.	The Gluing Lemma for Curvature Bounded Below	171
10.2.	Toponogov Globalisation	176
10.3.	The Bonnet–Myers Theorem	184
10.4.	Chapter Summary	192
11.	Part II Outlook	195

III	Appendices	201
A.	Locally-a-section Lagrangian Submanifolds	203
B.	Quaternionic Structures and Integrability in Two Dimensions	207
	B.1. Quaternionic Structures and Lychagin–Rubtsov Metrics	207
	B.2. The Poisson Equation for Pressure and Integrability	209
C.	Connections and Curvature	211
	C.1. Pull-back Metric in Two Dimensions	211
	C.1.1. Levi-Civita Connection of the Pull-back Metric	211
	C.1.2. Ricci Curvature Scalar of the Pull-back Metric	212
	C.2. Lychagin–Rubtsov Metric with Arbitrary Background Dimension	214
	C.2.1. Vielbein formalism	214
	C.2.2. Levi-Civita Connection of the Lychagin–Rubtsov Metric	215
	C.2.3. Ricci Curvature Scalar of the Lychagin–Rubtsov Metric	216
D.	The Cat’s Cradle for Lorentzian Length Spaces	217
	References	224

Part I

(Higher) Monge–Ampère Geometry
and Fluid Dynamics



1

Overview of Part I

Let M be a m -dimensional manifold with Riemannian metric \dot{g} and local coordinates $\{x^i\}_{i=1}^m$. A (Contact) Monge–Ampère Equation for a function $\psi \in \mathcal{C}^\infty(M)$ is a second-order non-linear differential equation on M , given by a quasi-linear combination of the minor determinants of the Hessian matrix for ψ . The minor determinants of order $k = 0, 1, \dots, m - 1$ for an $m \times m$ matrix are the determinants of the sub-matrices formed by removing $(m - k)$ rows and $(m - k)$ columns from the matrix, while the Hessian matrix is given by $\overset{\circ}{\nabla}^i \partial_j \psi$, where $\overset{\circ}{\nabla}^i$ is the Levi-Civita connection for \dot{g} . In particular, the coefficients of these quasi-linear combinations of minor determinants may themselves be non-linear in the coordinates x^i , the function ψ , and the first order partial derivatives of ψ . If the coefficients do not depend on ψ , then we call the equation a Symplectic Monge–Ampère Equation.

Such equations arise naturally in models of physical systems, including, but not limited to:

- $\psi^\lambda \psi_{xx} + (\lambda \psi^{\lambda-1} \psi_x - \psi_t + F(\psi)) = 0$ — The reaction-diffusion equation in one spatial dimension, which can be used to model the spread of both chemicals and biological populations [5].
- $\psi_{tt} + \psi_{yy} + \psi_{zz} - \psi_{tx} + (\psi_t)^2 = 0$ — The Khokhlov–Zabolotskaya equation in three spatial dimensions, which models the propagation of a bounded sound beam through a non-linear medium [6].
- $(\psi_{xx} \psi_{yy} - \psi_{xy}^2) + \psi_{zz} = 0$ — The Chynoweth–Sewell equation in three spatial dimensions, which can be solved for the geopotential of some incompressible semi-geostrophic flow [7,8].
- $\psi_{tt} + \psi_{xz} \psi_{ty} - \psi_{xy} \psi_{tz} = 0$ — The Grant equation in three spatial dimensions, whose solu-

tions yield metrics with self-dual curvature tensor, which have applications in the study of Einstein's equations and quantum gravity [9,10] (see also Plebanski's Heavenly equations).

Following on from the work of Sophus Lie, Lychagin [11] observes that all contact Monge–Ampère equations can be interpreted as constraints on submanifolds of the manifold of 1-jets J^1M , with the constraints obtained by imposing that certain differential forms $\alpha \in \Omega^m(J^1M)$ vanish when pulled back to said submanifolds. In the case where the Monge–Ampère equation is symplectic and does not depend on ψ , the differential form α in fact lives on the cotangent bundle T^*M . It is these symplectic Monge–Ampère equations which we consider in the first part of this thesis.

As an example, consider the Euclidean plane \mathbb{R}^2 with coordinates x, y and cotangent bundle $T^*\mathbb{R}^2$. Equip the cotangent bundle with local coordinates x, y, q_1, q_2 . Pulling back the differential form

$$\alpha = dq_1 \wedge dy - dq_2 \wedge dx, \quad (1.0.1)$$

to the submanifold $\iota : L \hookrightarrow T^*\mathbb{R}^2$, with

$$\iota(L) := \{(x, y, \psi_x, \psi_y) \mid (x, y) \in \mathbb{R}^2\} \subset T^*\mathbb{R}^2 \quad (1.0.2)$$

for some function $\psi \in \mathcal{C}^\infty(\mathbb{R}^2)$ yields

$$\iota^*\alpha = (\psi_{xx} + \psi_{yy}) dx \wedge dy. \quad (1.0.3)$$

In particular, since $dx \wedge dy$ is a volume form on \mathbb{R}^2 and therefore does not vanish, the constraint $\iota^*\alpha = 0$ corresponds to the Laplace equation $\Delta\psi = \psi_{xx} + \psi_{yy} = 0$, which is a Monge–Ampère equation.

The submanifolds to which we pull-back take the form of the generalised solutions introduced in [12–14] and correspond to solutions of the original Monge–Ampère equation, which are additionally allowed to be multivalued or defined only on a subset of the domain of the equation. As such, we may infer properties of solutions to the Monge–Ampère equation by studying the properties of these generalised solution submanifolds. For example, when classifying Monge–Ampère equations in two independent variables, [15, Theorem 1.5] introduce an almost (para-)complex structure on T^*M , whose type is related to the elliptic/hyperbolic nature of the underlying equation.

It was with this perspective in mind that Roulstone, Rubtsov, et al. [16–19] took to studying fluid dynamical flows, first considering the semi-geostrophic equations, shortly followed by the Poisson equation for the pressure of incompressible Navier–Stokes flows. Of particular interest

to op. cit. was the formation and topology of vortices; the Poisson equation for the pressure can be viewed as a non-linear relation between the pressure, vorticity, and rate-of-strain of a Navier–Stokes flow, hence is a key diagnostic quantity in this investigation [20]. The equation also plays a role in calculations of turbulence in three dimensions, where vorticity tends to accumulate in tubes [21]. These vortex tubes are ubiquitous features of turbulent flows and have been dubbed ‘the sinews of turbulence’ in [22], however, their topology can become highly complicated. Consequently, extracting such topological information from the Navier–Stokes equations has proven to be a resilient problem.

While the Poisson equation for the pressure of a fluid flow in \mathbb{R}^n can be written in terms of the components of the velocity as

$$\Delta p = -(\partial_i v_j)(\partial_j v_i), \quad (1.0.4)$$

in two dimensions, where the velocity can be expressed in terms of the derivatives of a stream function $\psi \in \mathcal{C}^\infty(\mathbb{R}^2)$, we may view (1.0.4) as the Monge–Ampère equation

$$\frac{1}{2}\Delta p = \partial_{xx}\psi\partial_{yy}\psi - (\psi_{xy})^2 \quad (1.0.5)$$

for the stream function, provided Δp is known. Consequently, it is possible to encode the Poisson equation for the pressure in two dimensions as a differential two-form $\alpha \in \Omega^2(T^*\mathbb{R}^2)$, which we can study using the so-called Monge–Ampère geometry of Lychagin. Indeed, [17, 19] were able to show that when $\Delta p > 0$, the equation is associated to an almost complex structure and is therefore elliptic, which implies that vorticity dominates over rate-of-strain. This provides a geometric description of the so-called Weiss criterion derived by [23–25], which can be used to determine the dominance of vorticity and rate-of-strain via the sign of the Laplacian of pressure. Despite not having a Monge–Ampère equation in three dimensions, [18] were able to find a differential three-form $\alpha \in \Omega^3(T^*\mathbb{R}^3)$ which encodes (1.0.4) and whose associated Hitchin almost (para-)complex structure [26, Section 2.1] plays the same role (we shall return to these points in greater detail later in the thesis).

The reformulation of (1.0.4) on an arbitrary Riemannian manifold is the main focal point for this part of the thesis. By approaching our investigation in a covariant manner, we are able to unify the results of [16–19] in a framework for studying incompressible fluid flows in both arbitrary dimension and on curved backgrounds, via the multisymplectic geometry of [27–29]. In particular, we are able to show that there is a choice of multisymplectic form on the cotangent bundle T^*M that encodes the incompressibility, which was assumed a priori in earlier works.

Organisation of Part I An outline of this part of the thesis, which is based on the work [1], is as follows. We start in [Chapter 2](#) with a brief review of the theory of Monge–Ampère geometry for symplectic Monge–Ampère equations, beginning by motivating the definition of a Monge–Ampère structure, which consists of a symplectic form ω and differential m -form α (where $m = \dim(M)$) on T^*M . We then define what it means for submanifolds $\iota : L \hookrightarrow T^*M$ to be generalised solutions of a Monge–Ampère structure and provide a description of classical solutions to Monge–Ampère equations in terms of these submanifolds. For concreteness, we then specialise to Monge–Ampère equations in two dimensions, providing explicit formulae for these equations and their corresponding structures, as well as the famous classification result of Lychagin and Rubtsov [15, Theorem 1.5]. [Chapter 2](#) concludes with the introduction of an entire family of almost (para-)Hermitian metrics associated to Monge–Ampère equations, along with a proof that such a metric always exists in the two-dimensional case. This family of metrics generalises the invariant quantities defined in two dimensions by [16, Equation 4.11] and three dimensions by [15, Equation 2.1].

In [Chapter 3](#) we move on to considering how this theory can be applied to the study of incompressible, two-dimensional fluid flows. We begin in [Section 3.1](#) with a review of the incompressible Navier–Stokes equations in Euclidean domains and introduce the physical notions of vorticity, strain, and divergence. After discussing the Weiss criterion for the dominance of vorticity and strain in two and three dimensions, we present a formulation of the Navier–Stokes equations on a Riemannian manifold M and demonstrate how to obtain the Poisson equation for the pressure in this setting. In particular, the Poisson equation for the pressure contains additional terms which depend on the Ricci curvature of the underlying manifold in this case. Passing to two dimensions, in [Section 3.2](#) the (now covariant) Poisson equation for the pressure becomes a Monge–Ampère equation for the stream function and we revisit the work of [17, 19], in order to provide a geometric justification for the Weiss criterion in the curved setting. We highlight the existence of an associated almost (para-)Hermitian metric on the cotangent bundle, which is distinct from that previously studied in [19, Section 3.2] and whose signature, along with the signature of its pull-back to classical solutions, is dictated by the relationship between vorticity and rate-of-strain. As a consequence, the Ricci scalar curvature of the pull-back metric can be described physically in terms of the accumulation of vorticity, rate-of-strain, and their gradients, which can then be related to topological information about the flow by means of the Gauß–Bonnet theorem and Euler characteristic. Finally, in [Section 3.3](#), an example of the foregoing theory in the (far simpler) Euclidean setting is presented, in order to demonstrate how our new tools are consistent with existing techniques and indicate how they might be applied.

In [Chapter 4](#), we consider how this approach can be generalised to fluid flows in higher

dimensions. Concretely, in [Section 4.1](#), we demonstrate there exists a modified Monge–Ampère structure in two dimensions, which follows from choosing a different symplectic structure on the cotangent bundle. This new choice of symplectic form encodes the incompressibility of the flow, coupling it to the Poisson equation for the pressure, which is otherwise unchanged. The choice is characterised by taking the Hodge-dual of differential forms on the background manifold, so while it is rather inconsequential in two dimensions, in higher dimensions this duality leads to the symplectic form in our structure being promoted to an $(m - 1)$ -plectic form, as we shall see in [Section 4.2](#). Furthermore, in arbitrary dimensions, we can again define an almost (para-)Hermitian metric on the cotangent bundle, whose pull-back to solutions has curvature described by gradients of vorticity and strain. Specialising to flows in three dimensions, the Gauß–Bonnet theorem no longer provides meaningful topological information via the Euler characteristic. However, in [Section 4.2.2](#) we note that the helicity, a much studied invariant of incompressible fluid flows [[30–32](#)], can be formulated in terms of our geometry and is associated with topological invariants of linked/knotted vortex tubes. In the penultimate section of this chapter, [Section 4.3](#), we show how three-dimensional, incompressible fluid flows admitting certain symmetries can be reduced, by means of both the Marsden–Weinstein reduction process [[33](#)] and the multisymplectic reduction of [[34](#)], to a comparatively easy to visualise two-dimensional flow. The resulting fluid flow need not be incompressible, but can again be studied using the tools established in [Chapter 3](#). This reduction process replicates the famous Lundgren transformation for three-dimensional fluid flows with symmetry and the symplectic variant was first illustrated in the context of Burgers’ vortex in [[19](#), [Section 6](#)]. Finally, in [Section 4.4](#), we provide examples of these reductions in action, in the form of an Arnol’d–Beltrami–Childress flow (which has \mathbb{R} symmetry) and Hicks–Moffatt vortices (which have \mathbb{S}^1 symmetry).

We conclude in [Chapter 5](#) by presenting an overview of ongoing work and future endeavours. We take particular interest in the classification of systems of equations given by the pull-back of differential forms to submanifolds of the cotangent space, such as the Poisson equation for the pressure and the divergence-free condition in dimension greater than two. These systems of equations appear to have some relationship to multisymplectic geometry and in that sense can be seen as a “higher category” of Monge–Ampère equation.



An Introduction to Monge–Ampère Geometry

In this chapter, we provide a pedagogical introduction to the study of Monge–Ampère equations via geometry. We begin by giving a general overview of Monge–Ampère structures [35, Section 1] and their generalised solutions [11, Section 3.4], before moving on to describe several tools which have been developed specifically to study Monge–Ampère equations in two dimensions. While the majority of the concepts we introduce in the first two sections of this chapter are well-known and appear in existing literature (in particular, we follow the textbook [4] and direct the reader there for additional details), their inclusion allows us to provide context for our latter results and address some statements which are typically treated as folklore.

The final section of this chapter introduces a new family of almost (para-)Hermitian metrics, called Lychagin–Rubtsov metrics, which generalise the pseudo-metric defined in three dimensions in [15, Equation 2.1] and provide several perspectives from which to study the geometry of generalised solutions. The observation that there exists such a family of metrics was first implied in the work-in-progress monograph [36, Definition 3.19, Definition 3.30], where they study the so-called Monge–Ampère metric: the Lychagin–Rubtsov metric which uniquely satisfies some additional self-duality properties. This metric has also appeared in the same form in [37, Definition 3.36, Definition 3.41] and is distinct from the metric described in, say [19, Section 3.2]. The fact that this observation is novel might be surprising, especially given the consideration of Kähler structures in [16, Section 4.c] and [38, Section 3]. However, the metric has historically been fixed and used to define the corresponding Hermitian form as opposed to the converse. Our approach of fixing the Kähler form first has been a boon in the study of the Poisson equation for the pressure of an incompressible fluid flow, since it allows us to provide a unified geometric

description of fluid flows in two and three dimensions, as we shall see in Chapter 3 and Chapter 4.

2.1 Symplectic Monge–Ampère Equations and their Solutions

As in the introduction, let M be an m -dimensional Riemannian manifold with metric \mathring{g} and local coordinates $\{x^i\}_{i=1}^m$. In particular, the coordinates represent the independent variables on which solutions to our Monge–Ampère equations depend. We refer to this manifold as the *background* or *configuration space* for our problem. The *cotangent bundle* of M is denoted T^*M and equipped with local coordinates $\{x^i, q_i\}_{i=1}^m$, where the x^i are the local coordinates we defined on M and q_i are the local fibre coordinates defined on the bundle. We may refer to T^*M as the *phase space* of our problem. Throughout the remainder of this part of the thesis, we use Einstein’s summation convention, with indices lowered and raised using \mathring{g}_{ij} and its inverse \mathring{g}^{ij} respectively, unless otherwise stated.

In this section, we will describe the process of defining a geometric structure on the cotangent bundle T^*M which encodes a given Monge–Ampère equation on M . We shall also observe how effectiveness of differential forms can be used to remove redundancy in this definition and demonstrate how a certain class of submanifolds of T^*M provides a natural generalisation of a classical solution to a Monge–Ampère equation.

2.1.1 Multisymplectic Geometry

In order to define such a *Monge–Ampère structure* on the cotangent bundle, we first require some tools from differential geometry, namely contraction maps and symplectic forms. In this subsection we shall define these objects and describe some of their more useful properties, in particular the existence of Darboux coordinates and a canonical symplectic form.

Although we will only require symplectic geometry when studying the class of Monge–Ampère equations defined in the introduction, here we couch our definitions in terms of multisymplectic (or k -plectic) geometry. The reason for this will become clear in Chapter 4 when we study the Poisson equation for the pressure of a three-dimensional incompressible fluid flow, which fails to be a symplectic Monge–Ampère equation and is hence better described by a pair of 2-plectic forms. For further details on multisymplectic geometry, we direct the interested reader to [27–29].

To begin, let us define the contraction map for a $(k + 1)$ -form:

Definition 2.1.1 (Contraction Map)

Let V be a real vector space and $\varpi \in \bigwedge^{k+1} V^*$ be a $(k + 1)$ -form on V . The *contraction map* for ϖ is the map

$$\lrcorner : V \rightarrow \bigwedge^k V^* ; \quad X \lrcorner \varpi = \varpi(X, \dots), \quad (2.1.1)$$

which takes vector fields X and returns the k -form given by ϖ with X as its first argument.

Definition 2.1.2 (Non-Degenerate Forms and k -Plectic Vector Spaces)

Let V be a real vector space. A $(k + 1)$ -form $\varpi \in \bigwedge^{k+1} V^*$ is called non-degenerate if its contraction map is injective on V . The pair (V, ϖ) is then called a k -plectic vector space. In the case where $k = 1$, we call (V, ϖ) a symplectic vector space.

In general, the contraction map for a given $(k + 1)$ -form ϖ is not surjective. However, when $k = 1$, the rank-nullity theorem implies that the contraction map $\lrcorner : V \rightarrow V^* \cong V$ for ϖ is surjective precisely when it is injective. That is, if (V, ϖ) is a symplectic vector space, the contraction map for ϖ is bijective.

Furthermore, it is a standard exercise in linear algebra to show that symplectic vector spaces are necessarily even-dimensional and admit a distinguished basis $\{e^i, f_i\}_{i=1}^m$, known as the Darboux basis, in which

$$\varpi(e^i, e^j) = 0, \quad \varpi(f_i, f_j) = 0, \quad \text{and} \quad \varpi(e^i, f_j) = \delta_j^i, \quad (2.1.2)$$

where δ_j^i is the Kronecker delta, and takes the value $\delta_j^i = 1$ if $i = j$ and zero otherwise. With respect to the corresponding dual basis $\{e_i^*, f^{*i}\}_{i=1}^m$ on V^* , we can therefore write

$$\varpi = e_i^* \wedge f^{*i}. \quad (2.1.3)$$

It is then straightforward to show that a 2-form ϖ on a $2m$ -dimensional vector space V is non-degenerate if and only if the corresponding m -th exterior power is a non-vanishing top-form $\varpi \wedge \dots \wedge \varpi \neq 0$. For further discussion of these properties in the context of differential geometry, see for example [39, Section 1].

Definition 2.1.3 (ℓ -Lagrangian Subspace)

Let (V, ϖ) be a k -plectic vector space and let $U \subseteq V$ be a subspace of V . The ℓ -th orthogonal complement $U^{\perp, \ell}$ of U with respect to ϖ is defined by

$$U^{\perp, \ell} := \{v \in V \mid v \lrcorner u_1 \lrcorner \dots \lrcorner u_\ell \lrcorner \varpi = 0 \text{ for all } u_1, \dots, u_\ell \in U\}, \quad (2.1.4)$$

for $\ell = 1, \dots, k$. If $U = U^{\perp, \ell}$, then we call U an ℓ -Lagrangian submanifold.

In the case where $k = 1$ and our vector space is symplectic, there are only 1-Lagrangian, or more simply Lagrangian, subspaces and these subspaces all have dimension $\frac{1}{2} \dim(V)$. For $k > 1$, the ℓ -Lagrangian subspaces may have different dimension. Furthermore, the condition

that $U \subseteq U^{\perp,k}$ is equivalent to the restriction $\varpi|_U = 0$ for $k \geq 1$.

We denote by $\Omega^k(N)$ the space of (smooth) differential k -forms on a manifold N , where $\Omega^0(N) \cong \mathcal{C}^\infty(N)$. By considering the tangent spaces to N point-wise as vector spaces, we can extend Definition 2.1.2 to differential forms on N as follows:

Definition 2.1.4 (k -Plectic Manifolds)

Let N be a (smooth) manifold and let $\varpi \in \Omega^{k+1}(N)$ be a differential $(k+1)$ -form on N . If ϖ is point-wise non-degenerate, then it is called an almost k -plectic form and the pair (M, ϖ) is an almost k -plectic manifold.

If, in addition, ϖ is closed (that is, $d\varpi = 0$), then it is called a k -plectic form and the pair (M, ϖ) is a k -plectic manifold. When $k = 1$, we replace the term 1-plectic with the more familiar symplectic, as above.

Definition 2.1.5 (Local Diffeomorphisms and Symplectomorphisms)

Let L and M be Riemannian manifolds. A function $h : L \rightarrow M$ is called a local diffeomorphism if, for all $y \in L$, there exists some open neighbourhood $V_y \subseteq L$ of y such that $h(V_y)$ is open in M and the restriction $h|_{V_y} : V_y \rightarrow h(V_y)$ of h to V_y is a diffeomorphism.

Now let (L, ϖ_1) and (M, ϖ_2) be k -plectic manifolds, for the same fixed k . A local diffeomorphism $h : L \rightarrow M$ is called a local k -plectomorphism if it additionally preserves the k -plectic structure, that is, if $h^*\varpi_2 = \varpi_1$, on each of the V_y . The corresponding definition of a (global) k -plectomorphism also holds.

Since a symplectic manifold must have even dimensional tangent spaces, it is necessary that the manifold itself be even dimensional. By the Darboux theorem [39, Theorem 1.9], not only do we have a Darboux basis on each tangent space, but there always exists local coordinates $\{x^i, q_i\}_{i=1}^m$ on N , called Darboux coordinates, such the symplectic form can be written as

$$\omega = dq_i \wedge dx^i. \quad (2.1.5)$$

We call this the canonical symplectic form and denote it by ω to signify its standing. In particular, for the cotangent bundle $N = T^*M$ of our m -dimensional manifold M , the x^i are local coordinates on M and the q_i are local coordinates on the fibres.

In contrast, for $k > 1$, k -plectic manifolds need not admit coordinates which give their k -plectic forms some canonical presentation [40, Section 5]. While this shall not prove problematic for our study of higher dimensional fluids in Chapter 4, since we shall instead write the k -plectic form in symplectic Darboux coordinates, it does prove to be an obstacle in the classification

of equations arising from such k -plectic forms. We shall discuss ongoing work concerning this classification in more detail in Chapter 5.

Finally, we require the following definition:

Definition 2.1.6 (ℓ -Lagrangian Submanifolds)

Let (N, ϖ) be a k -plectic manifold. A submanifold $\iota : L \hookrightarrow N$, where L is a manifold and ι is an immersion, is called an ℓ -Lagrangian submanifold of N if and only if $TL = TL^{\perp, \ell}$, where¹

$$TL^{\perp, \ell} := \bigcup_{p \in L} \{(p, X_p) \mid X_p \in (T_p L)^{\perp, \ell}\}, \quad (2.1.6)$$

for some $\ell = 1, \dots, k$.

As for ℓ -Lagrangian subspaces, when $k = 1$ and (N, ϖ) is a symplectic manifold, there are only 1-Lagrangian, or more simply Lagrangian, submanifolds and these have dimension $\frac{1}{2} \dim(N)$. When $N = T^*M$ and $k = 1$, we find $\dim(L) = \dim(M)$. Furthermore, the condition that $TL \subseteq TL^{\perp, k}$ is equivalent to the requirement that the pull-back of the k -plectic form to L vanishes, i.e. $\iota^* \varpi = 0$, for $k \geq 1$.

For the remainder of this thesis, submanifolds shall be assumed to be smoothly embedded and consequently equipped with the induced topology from T^*M , unless otherwise specified.

2.1.2 Monge–Ampère Structures

In (1.0.3), we saw how to obtain a Monge–Ampère equation from the pull-back of a differential 2-form α , on $T^*\mathbb{R}^2$, to submanifolds $\iota : L \hookrightarrow T^*\mathbb{R}^2$ of the form

$$\iota(L) := \{(x, y, \psi_x, \psi_y) \mid (x, y) \in \mathbb{R}^2\}, \quad (2.1.7)$$

for some $\psi \in \mathcal{C}^\infty(\mathbb{R}^2)$. However, it is possible for two differential forms to give the same equation if their difference is a differential form whose pull-back to L vanishes. In particular, the canonical symplectic form (2.1.5) on $T^*\mathbb{R}^2$ trivially vanishes on L given by (2.1.7), so

$$\alpha + F\omega = dq_1 \wedge dy - dq_2 \wedge dx + F(dq_1 \wedge dx + dq_2 \wedge dy) \quad (2.1.8)$$

also yields the Laplace equation on \mathbb{R}^2 when we impose $\iota^*(\alpha + F\omega) = 0$ and this works for any function $F \in \mathcal{C}^\infty(T^*\mathbb{R}^2)$. Therefore, if we wish to study the geometry associated to some Monge–Ampère equation on a manifold M , which differential form on T^*M should we choose?

¹Since ι is an immersion, the pushforward $\iota_*(T_p L)$ is a subset of the tangent space $T_{\iota(p)}N$ and it is in this sense that $T_p L$ should be understood.

In order to answer this question, we introduce the concept of an effective differential form, following [11, Definition 1.3] and [35, Section 1].

Definition 2.1.7 (Effective Differential m -Forms)

Let (N, ϖ) be a $2m$ -dimensional symplectic manifold. A differential m -form α on N is called ϖ -effective if $\alpha \wedge \varpi = 0$.

The following theorem of Hodge, Lepage, and Lychagin [11, Theorem 1.4] then tells us that any differential m -form on N can be uniquely decomposed into a ϖ -effective piece and a multiple of the symplectic form ϖ . For a more comprehensive treatment of the full result, we direct the reader to [4, Section 11.2].

Theorem 2.1.8 (Hodge–Lepage–Lychagin Theorem)

Let (N, ϖ) be a $2m$ -dimensional symplectic manifold. Any differential m -form $\tilde{\alpha}$ on N has a unique decomposition

$$\tilde{\alpha} = \alpha + \alpha_0 \wedge \varpi, \quad (2.1.9)$$

where α is ϖ -effective and α_0 is some differential $(m-2)$ -form. Furthermore, if two ϖ -effective m -forms vanish on the same Lagrangian submanifolds of N , then the m -forms must be proportional.

A differential form $\tilde{\alpha}$ is therefore ϖ -effective if and only if $\alpha_0 \wedge \varpi = 0$, by uniqueness of the decomposition. Hence, Theorem 2.1.8 defines an equivalence relation on differential m -forms, where each class contains a single ϖ -effective form. Furthermore, given a submanifold $\iota : L \hookrightarrow N$ which is Lagrangian with respect to ϖ , we find that the pull-back of $\tilde{\alpha}$ to L is

$$\iota^* \tilde{\alpha} = \iota^* \alpha + \iota^* \alpha_0 \wedge \iota^* \varpi = \iota^* \alpha, \quad (2.1.10)$$

so for every $\tilde{\alpha}$ in the equivalence class $[\alpha]$, the constraint $\iota^* \tilde{\alpha} = 0$ is the same as $\iota^* \alpha = 0$. In particular, for our example of $(T^*\mathbb{R}^2, \omega)$ with $\alpha = dq_1 \wedge dy - dq_2 \wedge dx$, it is easy to verify $\alpha \wedge \omega = 0$, so α is the sole ω -effective form in the equivalence class defined by (2.1.8).

Consequently, if we wish to study Monge–Ampère equations which arise in this manner, it will only be necessary to study the ϖ -effective forms α . Putting this all together, we now define Monge–Ampère structures on T^*M [35, Section 1]:

Definition 2.1.9 (Monge–Ampère Structure)

Let M be an m -dimensional manifold with cotangent bundle T^*M . Let ϖ be a symplectic form and α be a ϖ -effective m -form on T^*M . The pair (ϖ, α) is called a Monge–Ampère structure

on T^*M and α is referred to as the Monge–Ampère form.

So far in this section, we have been pulling back our example Monge–Ampère structure $(T^*\mathbb{R}^2, \omega)$ to the class of Lagrangian submanifolds (2.1.7), given by the graphs of gradients of functions $\psi \in \mathcal{C}^\infty(\mathbb{R}^2)$ in $T^*\mathbb{R}^2$, in order to generate a Monge–Ampère equation with solution ψ . However, we are also free to take the pull-back of a Monge–Ampère structure to an arbitrary Lagrangian submanifold. In the next subsection, we show that doing so allows us to consider a more general class of solutions which includes multivalued functions.

2.1.3 Classical and Generalised Solutions

Let us start by providing a formal definition of a generalised solution [11, Section 3.4]:

Definition 2.1.10 (Generalised Solutions of Monge–Ampère Structures)

A generalised solution of a Monge–Ampère structure (ϖ, α) is a submanifold $\iota : L \hookrightarrow T^*M$ which is Lagrangian with respect to ϖ , i.e. $\dim(L) = \dim(M)$ and $\iota^*\varpi = 0$, and for which the additional constraint $\iota^*\alpha = 0$ holds.

Now fix the Monge–Ampère structure (ω, α) on T^*M , where ω is the canonical symplectic form. Consider a (global) section $d\psi : M \rightarrow T^*M$ given by the differential of some function $\psi \in \mathcal{C}^\infty(M)$ and described by $x^i \mapsto (x^i, \partial_i\psi)$. This defines a Lagrangian submanifold with respect to the canonical symplectic form ω and $\iota(L) := d\psi(M) \subset T^*M$, as in (2.1.7). As we have seen in our example, the requirement that $\iota^*\alpha = 0$ then yields a Monge–Ampère equation

$$\Delta_\alpha(\psi) := (d\psi)^*\alpha = 0, \quad (2.1.11)$$

for which ψ is a classical solution. We will often abuse terminology and refer to Lagrangian submanifolds given by $d\psi(M)$ as classical solutions of the Monge–Ampère structure (ω, α) . The key here is that *classical solutions are trivially Lagrangian with respect to the symplectic form from our Monge–Ampère structure*.

Remark 2.1.11 (Classical Solutions in Non-Canonical Coordinates)

Since generalised solutions are Lagrangian submanifolds, the precise requirements for a generalised solution of a Monge–Ampère structure to correspond to a classical solution of a Monge–Ampère equation depend on the choice of symplectic form. For example, if we take the symplectic form $\varpi = dq_1 \wedge dx^2 - dq_2 \wedge dx^1$ on $T^*\mathbb{R}^2$, the submanifolds $d\psi(M)$ are not trivially Lagrangian and $(d\psi)^*\varpi = 0$ yields the Laplace equation as a constraint on ψ .² Instead, consider submani-

²In particular, if ψ satisfies both $(d\psi)^*\varpi = 0$ and $(d\psi)^*\alpha = 0$, then $d\psi(M)$ is a *generalised* solution for (ϖ, α) .

folds $\iota : L \hookrightarrow T^*\mathbb{R}^2$ given by

$$\iota(L) = \star d\psi(\mathbb{R}^2) = \{(x, y, -\psi_y, \psi_x) \mid (x, y) \in \mathbb{R}^2\}, \quad (2.1.12)$$

where \star is the Hodge star operator on \mathbb{R}^2 . These are trivially Lagrangian with respect to ϖ and $(\star d\psi)^*\varpi = 0$ is automatic.

More generally, let $(\varpi, \tilde{\alpha})$ be a Monge–Ampère structure on T^*M . Then by the Darboux theorem, there exists a local symplectomorphism

$$F : (T^*M, \omega) \rightarrow (T^*M, \varpi); \quad F^*\varpi = \omega, \quad (2.1.13)$$

which takes our symplectic form ϖ to the canonical one in a neighbourhood of each point. This function also yields an ω -effective Monge–Ampère form $\alpha = F^*\tilde{\alpha}$ via $(F^*\tilde{\alpha}) \wedge \omega = F^*(\tilde{\alpha} \wedge \varpi) = 0$. It is straightforward to check that a submanifold $\iota : L \hookrightarrow T^*M$ is a generalised solution for (ω, α) if and only if $\tilde{\iota} := F \circ \iota : L \hookrightarrow T^*M$ is a generalised solution for $(\varpi, \tilde{\alpha})$. In particular,

$$(d\psi)^*\alpha = (F \circ d\psi)^*\tilde{\alpha}, \quad (2.1.14)$$

so while classical solutions $\psi \in \mathcal{C}^\infty(M)$ of our Monge–Ampère equation correspond to generalised solutions of (ω, α) with the form $d\psi(M)$, they correspond to generalised solutions of $(\varpi, \tilde{\alpha})$ with the form $F \circ d\psi(M)$. It is easy to verify that these are precisely the submanifolds of T^*M on which ϖ trivially vanishes:

$$(F \circ d\psi)^*\varpi = (d\psi)^*(F^*\varpi) = d\psi^*\omega = 0. \quad (2.1.15)$$

In summary, a classical solution of a Monge–Ampère structure is a solution which would be given by a section $d\psi$ if we were in Darboux coordinates.

△

So far we have seen how pulling back differential forms on (T^*M, ϖ) to Lagrangian submanifolds can yield Monge–Ampère equations and how differential forms with the same ϖ -effective piece produce the same equation. In [15], it is observed that, as a corollary of Theorem 2.1.8, there is a bijection between Monge–Ampère equations on a manifold M and ω -effective forms on T^*M ; using the previous remark, we may now formulate this in the language of Monge–Ampère structures:

Corollary 2.1.12 (Monge–Ampère Structures and Equations)

*Let M be an m -dimensional manifold, T^*M be its cotangent bundle, and ϖ be a fixed symplectic*

form on T^*M . There is a bijection between Monge–Ampère structures $(\varpi, \tilde{\alpha})$ and Monge–Ampère equations, with Monge–Ampère equations being given by

$$(F \circ d\psi)^* \tilde{\alpha} = (d\psi)^*(F^* \tilde{\alpha}) = 0, \quad (2.1.16)$$

where $F : (T^*M, \omega) \rightarrow (T^*M, \varpi)$ is a local symplectomorphism satisfying $F^* \varpi = \omega$.

In essence, Remark 2.1.11 and Corollary 2.1.12 state that there is a one-to-one correspondence between Monge–Ampère equations and Monge–Ampère structures with fixed symplectic form, regardless of the choice of symplectic form. Consequently, we can consider all Monge–Ampère equations on a manifold M with respect to whichever symplectic form we prefer. Furthermore, the second part of Theorem 2.1.8 tells us that if two Monge–Ampère structures (ϖ, α_1) and (ϖ, α_2) on T^*M have the same set of generalised solutions, then the Monge–Ampère forms α_1 and α_2 are proportional. That is, their corresponding Monge–Ampère equations differ only by a scale factor $F(x^i, \partial_i \psi)$. In particular, this allows us to freely scale our Monge–Ampère form as necessary. Should we wish to remove this redundancy, it is possible to fix a scale for our Monge–Ampère form, as we shall see in Section 2.2.1.

In what follows, we fix our symplectic form ω to be the canonical one for simplicity. Given that the class of generalised solutions of a Monge–Ampère structure (ω, α) includes submanifolds $d\psi(M) \subset T^*M$ which correspond to classical solutions $\psi \in \mathcal{C}^\infty(M)$ of a Monge–Ampère equation, this begs the question: what do the generalised but non-classical solutions of a Monge–Ampère structure correspond to, in terms of functions which solve a Monge–Ampère equation? In [35, Section 1], Banos provides conditions under which a generalised solution may be described locally by classical solutions, which goes part-way to answering this question. These conditions are described precisely by the following definition and proposition:

Definition 2.1.13 (Locally-a-Section Submanifold)

A submanifold $\iota : L \hookrightarrow T^*M$ is called locally-a-section if, for all $y \in L$, there exists a neighbourhood $V_y \subseteq L$ of y , an open set $U_y \subseteq M$, and a function $\psi_y \in \mathcal{C}^\infty(U_y)$ such that $\iota(V_y) = d\psi_y(U_y)$.

Proposition 2.1.14 (Locally-a-Section iff Local Diffeomorphism)

Let (M, g) be a Riemannian manifold and let its cotangent bundle be equipped with the canonical symplectic form ω . Let $\pi : T^*M \rightarrow M$ be the canonical projection on the cotangent bundle T^*M . A Lagrangian submanifold $\iota : L \hookrightarrow T^*M$ with respect to ω is locally-a-section if and only if the projection $\pi|_L := \pi \circ \iota : L \rightarrow M$ is a local diffeomorphism.

Since [35] does not provide a proof of Proposition 2.1.14, one can be found in Appendix A.

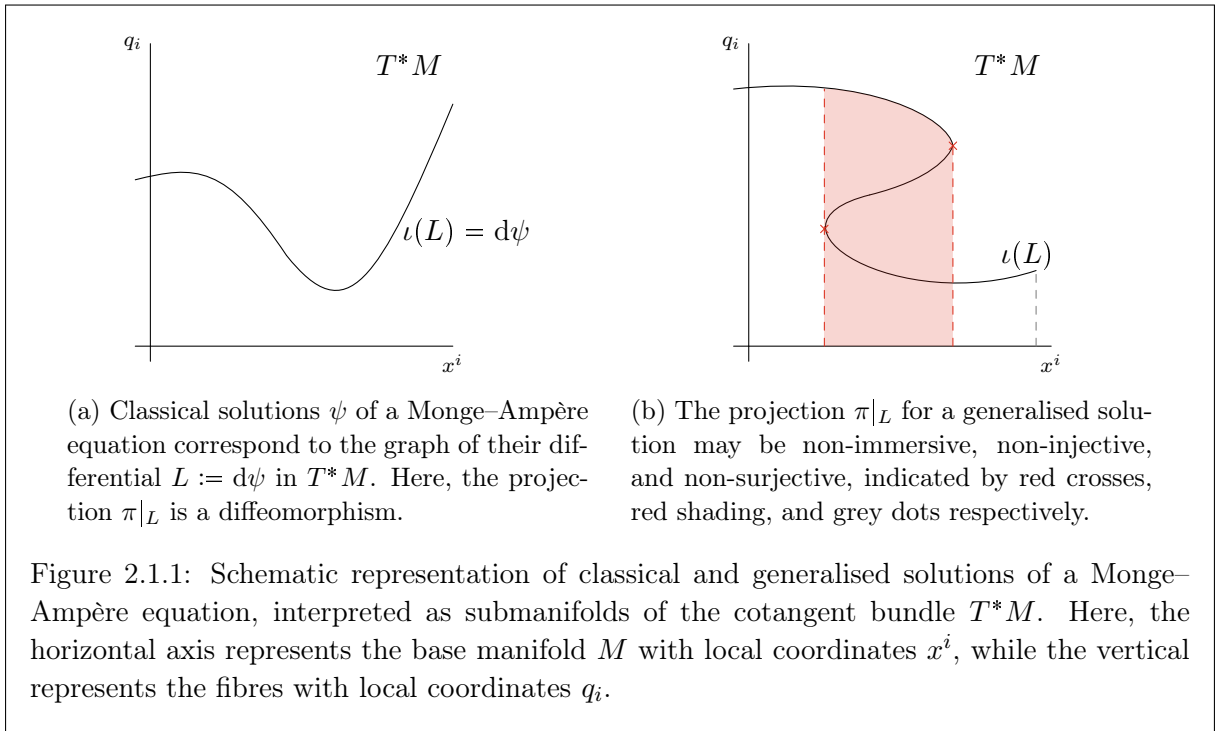
For now, we note that [Proposition 2.1.14](#) implies, for locally-a-section submanifolds $\iota : L \hookrightarrow T^*M$, there is a neighbourhood $V_y \subseteq L$ of each $y \in L$ which can be given the local coordinates from some open set $U_y \subseteq M$, with respect to which $\pi|_y := \pi|_{V_y} : V_y \rightarrow U_y$ is locally the identity. Furthermore, if L corresponds to a classical solution, i.e. is a global section, then this property holds with $V_y = L$ and $U_y = M$, see [Figure 2.1.1](#).

Although $\pi|_L$ is a local diffeomorphism when L is locally-a-section, the projection need not be injective or surjective onto M , though $\pi|_L$ must be an immersion by [[41](#), [Proposition 4.8\(a\)](#)]. In particular, patching together the functions ψ_y may result in a multivalued function on M , or a function which is only defined on a subset of the domain; [Definition 2.1.5](#) and [Definition 2.1.13](#) only require properties to hold in a neighbourhood of each $y \in L$ and are not interested in how many times each point in M is used to satisfy these properties. Furthermore, arbitrary generalised solutions (those which are not locally-a-section) may exhibit points where $\pi|_L$ is non-immersive, which we call projection singularities. Around such singular points, we cannot describe L solely using local coordinates x^i from M and instead require the language of generating functions [[42](#), [43](#)]. [Figure 2.1.1b](#) provides a schematic of a generalised solution exhibiting each of these pathologies.

Since locally, around generic points (where $\pi|_L$ is immersive), generalised solutions can be described by a classical solution, we shall predominantly consider classical solutions moving forward. Recent work of [[44](#)] showed that, in the context of the semi-geostrophic equations, projection singularities can produce degeneracies in metrics on L and that these can be related to shocks or wave-fronts. Considering only (locally) classical solutions to [\(1.0.5\)](#) will allow us to investigate the relationship between metrics on L and the properties of fluid flows, without having to account for such additional degeneracies. The study of fully generalised solutions to [\(1.0.5\)](#) and the degeneracies arising from projection singularities is a subject for future research.

Remark 2.1.15 (Sufficiently Smooth Solutions)

While classical solutions to a ℓ -th order partial differential equation are generally only required to be as differentiable as the the maximum number ℓ of derivatives taken, we restrict ourselves to considering smooth classical solutions above. This ensures that, for an m -dimensional manifold M , the *Monge–Ampère operator* $\Delta_\alpha : \mathcal{C}^\infty(M) \rightarrow \Omega^m(M) \cong \mathcal{C}^\infty(M)$ from [\(2.1.11\)](#) takes smooth functions to smooth functions. One approach to properly considering all classical solutions would be to treat Monge–Ampère operators of the form $\Delta_\alpha : \mathcal{C}^k(M) \rightarrow \Omega_{k-\ell}^m(M) \cong \mathcal{C}^{k-\ell}(M)$ with $k \geq \ell$, where $\Omega_{k-\ell}^m(M)$ denotes the space of m -forms on M with only $(k - \ell)$ -times differentiable coefficients. The Lagrangian submanifolds corresponding to our generalised solutions would then not be required to be smooth. Smooth solutions are sufficient for our purposes, so we leave further investigation of this idea for future research.



△

2.2 Almost (Para-)Complex Structures and Lychagin–Rubtsov Metrics

In this section, we introduce the notion of an almost (para-)complex structure on a manifold N and illustrate how said structure can be used to define a metric on N . In particular, we wish to consider almost (para-)complex structures on T^*M which can be constructed from the differential forms in our Monge–Ampère structure, such that the resulting metric depends on our choice of Monge–Ampère equation. For more information on the almost complex case, we direct the reader to the textbook [45, Chapter 8]. There is far less literature concerning almost para-complex structures, however [46] provide an excellent review of the fundamentals.

In what follows, let N be a manifold with tangent bundle TN and let $\Gamma(TN)$ denote the space of global sections of TN , that is, the space of vector fields on N . Furthermore, let \mathbb{I} denote the identity on TN .

Definition 2.2.1 (Almost (Para-)Complex Structure)

Let N be a manifold. An endomorphism $J : TN \rightarrow TN$ is called an

- almost complex structure if $J^2 = -\mathbb{I}$,

- almost para-complex structure if $J^2 = \mathbb{I}$ and $\text{tr}(J) = 0$.

Accordingly, we call the pair (N, J) with $J^2 = \pm\mathbb{I}$ and $\text{tr}(J) = 0$ an almost (para-)complex manifold.

If J is an almost complex structure on N , then its eigenvalues are $\pm i$, where $i := \sqrt{-1}$. The corresponding eigenspaces, which we call the holomorphic and anti-holomorphic tangent bundles, are respectively given by

$$T^{1,0}N := \{X - iJX \mid X \in TN\} \quad \text{and} \quad T^{0,1}N := \{X + iJX \mid X \in TN\}. \quad (2.2.1)$$

These bundles define a decomposition of the complexified tangent bundle $T_{\mathbb{C}}N := TN \otimes \mathbb{C} \cong T^{1,0}N \oplus T^{0,1}N$. Furthermore, Binet's theorem implies

$$(-1)^{\dim(N)} = \det(J^2) = (\det J)^2 > 0, \quad (2.2.2)$$

so almost complex manifolds are necessarily even dimensional.

Analogously, in the almost para-complex case, the eigenvalues of J are ± 1 and the corresponding eigenspaces

$$T^{1,0}N := \{X - JX \mid X \in TN\} \quad \text{and} \quad T^{0,1}N := \{X + JX \mid X \in TN\} \quad (2.2.3)$$

define a decomposition of the *real* tangent bundle $TN \cong T^{1,0}N \oplus T^{0,1}N$. As $\text{tr}(J) = 0$, the eigenvalues ± 1 must appear with equal multiplicities (i.e. $T^{1,0}N$ and $T^{0,1}N$ are of the same rank), so almost para-complex manifolds are also necessarily even dimensional.

Definition 2.2.2 (Integrable (Para-)Complex Structures)

Let J be an almost (para-)complex structure on a manifold N . We call J integrable if and only if the bundles $T^{1,0}N$ and $T^{0,1}N$ are both integrable in the sense of Frobenius, that is, if

$$[\Gamma(T^{1,0}N), \Gamma(T^{1,0}N)] \subseteq \Gamma(T^{1,0}N) \quad \text{and} \quad [\Gamma(T^{0,1}N), \Gamma(T^{0,1}N)] \subseteq \Gamma(T^{0,1}N), \quad (2.2.4)$$

where $\Gamma(\mathcal{D})$ denotes the space of smooth sections of a bundle \mathcal{D} over M and $[\cdot, \cdot]$ is the usual Lie bracket on vector fields. In this case, we may refer to J as a (para-)complex structure.

It is straightforward to verify that an algebraic condition equivalent to the Frobenius integrability of the bundles $T^{1,0}N$ and $T^{0,1}N$ is the vanishing of the Nijenhuis tensor

$$N_J(X, Y) := -J^2[X, Y] + J[X, JY] + J[JX, Y] - [JX, JY], \quad (2.2.5)$$

for all $X, Y \in TN$. If J is almost complex, integrability of $T^{1,0}N$ is equivalent to integrability of $T^{0,1}N$ by complex conjugation, however this is not true in the almost para-complex case. Indeed, for J almost para-complex, the condition that $T^{1,0}N$ (resp. $T^{0,1}N$) is integrable in the sense of Frobenius is equivalent to N_J vanishing on $T^{0,1}N$ (resp. $T^{1,0}N$) and these two requirements are independent.³

By duality of the tangent and cotangent bundles, when J is an almost complex structure, we also obtain a decomposition of the space of complex 1-forms $\Omega_{\mathbb{C}}^1(N) := \Omega^1(N) \otimes \mathbb{C} \cong \Omega^{1,0}(N) \oplus \Omega^{0,1}(N)$. Taking exterior products and linear combinations, it is then possible to define the space of *smooth* differential (p, q) -forms $\Omega^{p,q}(N)$ on N which are complex r -forms (where $r = p + q$) that act non-trivially on p elements of the $T^{1,0}N$ and q elements of $T^{0,1}N$. Similarly, when J is an almost para-complex structure, we obtain a decomposition of the space of real 1-forms $\Omega^1(N) \cong \Omega^{1,0}(N) \oplus \Omega^{0,1}(N)$ and $\Omega^{p,q}(N)$ is the analogous space of (real) (p, q) -forms.

Definition 2.2.3 (Almost (Para-)Hermitian Form and Metric)

Let (N, J) be an almost (para-)complex manifold. A real differential $(1, 1)$ -form on N is called an almost (para-)Hermitian form for J . Since $(1, 1)$ forms satisfy

$$K(JX, Y) = -K(X, JY), \quad (2.2.6)$$

for all $X, Y \in \Gamma(TT^*M)$, we are naturally led to a metric $g(X, Y) = K(X, JY)$ on N , which we call an almost (para-) Hermitian metric.

If the differential $(1, 1)$ -form K is also closed, we shall call K an almost (para-)Kähler form and g an almost (para-) Kähler metric. In the case where J is integrable, then we drop the word “almost” from the qualifiers. A priori, almost (para-)Hermitian metrics can be any signature. However, [51, Section 1] show, using $g(X, X) = -g(JX, JX)$ and isotropy of the bundles $T^{1,0}N$ and $T^{0,1}N$, that almost para-Hermitian metrics are necessarily Kleinian, i.e. they have signature (m, m) on a $2m$ dimensional manifold. In contrast, almost Hermitian metrics satisfy $g(X, Y) = g(JX, JY)$ and we can make no such statement.

2.2.1 Monge–Ampère Structures and Metrics in Two Dimensions

Let us now apply the theory we described above to the case of two-dimensional Monge–Ampère structures. For more details concerning Monge–Ampère equations and their structures in two dimensions, we direct the interested reader to [4].

³For more information on how the integrability of the structure J is related to the existence of holomorphic coordinates on N , see the Newlander–Nirenberg theorem e.g. [47, 48]. In the para-complex case, see [49] and [50, Section 2.1].

Let $N = T^*M$ be the four-dimensional cotangent bundle of some two-dimensional Riemannian manifold M and let (ϖ, α) be some Monge–Ampère structure on T^*M . The symplectic form ϖ is a non-degenerate two-form and therefore, by our discussion in Section 2.1.1, defines a volume form $\varpi \wedge \varpi \neq 0$ on T^*M . We shall call the scaled volume form $\frac{1}{2}\varpi \wedge \varpi$ the *Liouville volume form* in four dimensions. The following quantity encodes the ratio between the wedge-square of the Monge–Ampère form α and the symplectic form ϖ .

Definition 2.2.4 (Pfaffian)

Let M be a two-dimensional Riemannian manifold with cotangent bundle T^*M and let (ϖ, α) be a Monge–Ampère structure on T^*M . The Pfaffian of α with respect to ϖ is given by

$$\text{Pf}_\varpi : \Omega^2(T^*M) \rightarrow \mathcal{C}^\infty(T^*M); \quad \alpha \wedge \alpha := \text{Pf}_\varpi(\alpha)\varpi \wedge \varpi. \quad (2.2.7)$$

It should be clear that the Pfaffian of a Monge–Ampère form α vanishes at a point $p \in T^*M$ if and only if α is degenerate at p , in which case we call the Monge–Ampère form and its corresponding equation *parabolic*. When $\text{Pf}_\varpi(\alpha) > 0$, we call the Monge–Ampère form α and equation *elliptic*, while if $\text{Pf}_\varpi(\alpha) < 0$, we call them *hyperbolic*.

Remark 2.2.5 (Classification of Monge–Ampère PDEs)

The fact that our choice of terminology here coincides with the classification of non-linear second order PDEs in two variables is no coincidence, as we shall now show. The family of Monge–Ampère equations for functions $\psi \in \mathcal{C}^\infty(\mathbb{R}^2)$, where \mathbb{R}^2 is equipped with coordinates x, y , can be written as

$$A\psi_{xx} + 2B\psi_{xy} + C\psi_{yy} + D(\psi_{xx}\psi_{yy} - (\psi_{xy})^2) + E = 0, \quad (2.2.8)$$

where subscripts indicate partial derivatives with respect to x and y , and A, B, C, D , and E are smooth functions of x, y, ψ_x and ψ_y (in the contact case they may also depend on ψ). The type of this equation is given by the type of its linearisation $\frac{d}{d\varepsilon}|_{\varepsilon=0}\Delta_\alpha(\psi + \varepsilon\phi) = 0$ about a solution ψ . This, in turn, is given by the determinant of the matrix of coefficients

$$\ell := \begin{pmatrix} A + D\psi_{xx} & B - D\psi_{xy} \\ B - D\psi_{xy} & C + D\psi_{yy} \end{pmatrix}, \quad (2.2.9)$$

of the second order terms in the linearised equation. More precisely,

$$\det(\ell) = AC - B^2 - DE, \quad (2.2.10)$$

and we call the equation *elliptic* if $\det(\ell) > 0$, *hyperbolic* if $\det(\ell) < 0$, and *parabolic* if $\det(\ell) = 0$. For more details on the classification of second order PDEs, see for example [52, Chapter 2].

Now let $T^*\mathbb{R}^2$ be equipped with the coordinates x, y, q_1 and q_2 , as in Chapter 1, and fix ω to be the canonical symplectic form in these coordinates. The ω -effective forms on $T^*\mathbb{R}^2$ are

$$\alpha = \tilde{A}dq_1 \wedge dy - \tilde{C}dq_2 \wedge dx + \tilde{B}(dq_2 \wedge dy - dq_1 \wedge dx) + \tilde{D}dq_1 \wedge dq_2 + \tilde{E}dx \wedge dy, \quad (2.2.11)$$

where $\tilde{A}, \tilde{B}, \tilde{C}, \tilde{D}$, and \tilde{E} are smooth functions of x, y, q_1 , and q_2 . The equations (2.2.8) are then obtained by pulling back the differential forms α to classical solutions, i.e. by considering $(d\psi)^*\alpha = 0$ and fixing the notation $A = (d\psi)^*\tilde{A}$, $B = (d\psi)^*\tilde{B}$ and so forth. The Pfaffian of α with respect to ω is then given by

$$\text{Pf}_\omega(\alpha) = \tilde{A}\tilde{C} - \tilde{B}^2 - \tilde{D}\tilde{E}. \quad (2.2.12)$$

The pull-back of this quantity to a classical solution ψ is then precisely the discriminating quantity (2.2.10)

△

Given a Monge–Ampère structure (ϖ, α) on T^*M , whenever the $\text{Pf}_\varpi(\alpha)$ is non-zero, following [15, Section 1], we may define an endomorphism $J_\alpha : T(T^*M) \rightarrow T(T^*M)$ by

$$\frac{\alpha}{\sqrt{|\text{Pf}_\varpi(\alpha)|}} =: J_\alpha \lrcorner \varpi. \quad (2.2.13)$$

The ϖ -effectiveness of α then yields the identity

$$\sqrt{|\text{Pf}_\varpi(\alpha)|} (J_\alpha^2 \lrcorner \varpi + \text{sgn}(\text{Pf}_\varpi(\alpha)) \varpi) \wedge \varpi = J_\alpha \lrcorner (\alpha \wedge \varpi) = 0. \quad (2.2.14)$$

Upon combining this identity with the non-degeneracy of ω , we immediately see that $J_\alpha^2 = -\text{sgn}(\text{Pf}_\varpi(\alpha))\mathbb{I}$. Letting ϖ^{-1} denote the poly-vector dual to ϖ , i.e. $\varpi^{-1} \lrcorner \varpi = 1$, [4, Theorem 5.2.3] implies that a differential 2-form $\alpha \in \Omega^2(T^*M)$ is ϖ -effective in the sense of Definition 2.1.7 if and only if $\varpi^{-1} \lrcorner \alpha = 0$. Therefore, it follows from

$$\text{tr}(J_\alpha) = \varpi^{-1} \lrcorner J_\alpha \lrcorner \varpi = \varpi^{-1} \lrcorner \alpha = 0, \quad (2.2.15)$$

that J_α is traceless. Consequently, J_α is an almost complex (resp. para-complex) structure when $\text{Pf}_\varpi(\alpha) > 0$ (resp. $\text{Pf}_\varpi(\alpha) < 0$).

The Pfaffian transforms as $\text{Pf}_\varpi(\kappa\alpha) = \kappa^2\text{Pf}_\varpi(\alpha)$ under scaling of the form α by a non-zero function $\kappa \in \mathcal{C}^\infty(T^*M)$. In particular, $\text{Pf}_\varpi(J_\alpha \lrcorner \varpi) = \text{sgn}(\text{Pf}_\varpi(\alpha))$. Therefore, all multiples of a Monge–Ampère form α are assigned the same almost (para-)complex structure J_α , up to an overall sign given by $\text{sgn}(\kappa)$, and (2.2.13) defines a canonical scale for our Monge–Ampère forms,

namely that in which the Pfaffian has modulus 1.

The following fundamental theorem from [15, Theorem 1.5] (see also [53, Theorem 1]) provides us with necessary and sufficient conditions on the Monge–Ampère structure (ϖ, α) for J to be integrable:

Theorem 2.2.6 (The Lychagin–Rubtsov Theorem)

Let M be a two-dimensional Riemannian manifold with cotangent bundle T^*M and let (ϖ, α) be a Monge–Ampère structure on T^*M . Let J_α denote the almost (para-)complex structure (2.2.13). The following are equivalent:

- J_α is an integrable (para-)complex structure.
- $d(J_\alpha \lrcorner \varpi) = 0$.
- The Monge–Ampère equation corresponding to α is locally symplectomorphic to the Laplace equation if $\text{Pf}_\varpi(\alpha) > 0$ and the Wave equation if $\text{Pf}_\varpi(\alpha) < 0$.

Furthermore, the following result from [1, Proposition B.2.] (which also appears in [37, Lemma 3.8] and the work in progress monograph [36, Lemma 3.14]) shows that there always exists an almost (para-)Hermitian form with respect to J_α on T^*M , which we denote by K . As ϖ and $J_\alpha \lrcorner \varpi$ define non-degenerate differential $(2, 0)$ - and $(0, 2)$ -forms on T^*M with respect to J_α , the triple ϖ , α , and K are linearly independent.⁴

Proposition 2.2.7 (Existence of Differential $(1, 1)$ -Forms)

For J_α as defined in (2.2.13), there exists a differential $(1, 1)$ -form $K \in \Omega^2(T^*M)$ such that $K \wedge K \neq 0$, $K \wedge \varpi = 0$, and $K \wedge (J_\alpha \lrcorner \varpi) = 0$.

Given an almost (para-)Hermitian form on T^*M with respect to the almost (para-)complex structure described by a Monge–Ampère structure (ϖ, α) as in (2.2.13), we may define an almost (para-)Hermitian metric $g(X, Y) = K(X, JY)$ as in Definition 2.2.3. We call an almost (para-)Hermitian metric constructed from a Monge–Ampère structure in this way a Lychagin–Rubtsov metric.

Remark 2.2.8 (Choices of Lychagin–Rubtsov Metric)

At this stage, it is worth noting how our consideration of the almost (para-)Hermitian metrics associated to a Monge–Ampère structure differs from earlier works. More precisely, we are free to choose any $(1, 1)$ -form with respect to J_α in order to construct our metric, while [16, Equation

⁴The $(2, 0)$ - and $(0, 2)$ -forms are given by $\varpi \pm iJ_\alpha \lrcorner \varpi$ in the almost complex case and by $\varpi \pm J_\alpha \lrcorner \varpi$ in the almost para-complex case.

4.11] (see also [19, Section 3.2] and [38, Section 2]) fix the non-degenerate bilinear form

$$g_\alpha(X, Y) := \frac{[(X \lrcorner \alpha) \wedge (Y \lrcorner \varpi) + (Y \lrcorner \alpha) \wedge (X \lrcorner \varpi)] \wedge \text{vol}_M}{\frac{1}{2}\varpi^2} \quad (2.2.16)$$

on T^*M to be their metric and define their $(1, 1)$ -form to be the one which yields this metric, namely $K(X, Y) = g_\alpha(\hat{\mathcal{J}}X, Y)$ (up to a choice of conformal factor which varies depending on the text). In [54], the metric (2.2.16) has been linked to a metric occurring in the theory of optimal mass transport with respect to which optimal maps are characterised via volume-maximising Lagrangian submanifolds.

Up to a conformal factor obtained by scaling the $(1, 1)$ -form K , it can be seen that the choice of Lychagin–Rubtsov metric g corresponds to a choice of almost (pseudo-)quaternionic structure on T^*M . Indeed, consider the triple of non-degenerate 2-forms (ϖ, α, K) , where α and K are scaled such that their Pfaffians with respect to ϖ have modulus 1, i.e.

$$\varpi \wedge \varpi = \text{Pf}_\varpi(\alpha) \alpha \wedge \alpha = \text{Pf}_\varpi(K) K \wedge K. \quad (2.2.17)$$

The identity (2.2.13) can be used to define an almost (para-)complex structure for any pair of non-degenerate 2-forms, so we define

$$\alpha =: J \lrcorner \varpi, \quad K =: R \lrcorner \alpha, \quad \text{and} \quad \varpi =: S \lrcorner K, \quad (2.2.18)$$

with $J^2 = -\text{Pf}_\varpi(\alpha)\mathbb{I}$, $S^2 = -\text{Pf}_\varpi(K)\mathbb{I}$, and $R^2 = -\text{Pf}_\varpi(\alpha)\text{Pf}_\varpi(K)\mathbb{I}$. These three almost (para-) complex structures define an almost (pseudo-)quaternionic structure on T^*M , see Appendix B.1 for details. In particular, changing K while keeping the Monge–Ampère structure (ϖ, α) fixed changes R and S but not J .

Finally, in [36, Proposition 3.15] (see also [37, Proposition 3.13]), extra constraints are placed on the Lychagin–Rubtsov metric associated with the Monge–Ampère structure (ϖ, α) in order to make the choice unique (up to a sign). In particular, K is chosen such that $\text{Pf}_\varpi(K) = \text{Pf}_\varpi(\alpha)$, so R is always an almost complex structure, and such that (ϖ, α, K) are all anti-self-dual with respect to the resulting metric g . They call this metric the *Monge–Ampère metric*. We shall not make use of the Monge–Ampère metric in this thesis and instead make a choice of Lychagin–Rubtsov metric that brings the geometric structures for fluid flows in two and three dimensions into alignment.

△

2.2.2 Monge–Ampère Structures and Metrics in Higher Dimensions

Let us now make a brief comment on the analogous construction in higher dimensions, focusing on the three dimensional case. Although the Poisson equation for the pressure we study in later chapters fails to be a Monge–Ampère equation in dimension higher than two, we shall still take inspiration from the following techniques in that setting.

We begin by noting that, while the complex structure J_α in (2.2.13) is defined for any Monge–Ampère structure (ϖ, α) in two dimensions (this is easily seen by taking the matrix representation of the definition and using the non-degeneracy of ϖ), in higher dimensions the forms ϖ and α are of different orders and the definition is nonsensical. Consequently, in three dimensions, we require an alternative way to construct an almost (para-)complex structure from our Monge–Ampère structure.

Fortunately for us, the work of Hitchin [26] provides a resolution (see also [35, Section 3]). In particular, they note that there is an isomorphism $\Omega^5(T^*M) \cong \Gamma(TT^*M) \otimes \Omega^6(T^*M)$ induced by the natural exterior product pairing $\Omega^1(T^*M) \otimes \Omega^5(T^*M) \rightarrow \Omega^6(T^*M)$. Explicitly, we have

$$\phi : \Omega^5(T^*M) \rightarrow \Gamma(TT^*M) \otimes \Omega^6(T^*M), \quad \phi(\rho)(\lambda, X_1, \dots, X_6) := X_1 \lrcorner \dots \lrcorner X_6 \lrcorner (\rho \wedge \lambda) \quad (2.2.19)$$

for all $\rho \in \Omega^5(T^*M)$, $\lambda \in \Omega^1(T^*M)$, and $X_1, \dots, X_6 \in \Gamma(TT^*M)$. Consequently, we may define the following:

Definition 2.2.9 (Hitchin Endomorphism and Pfaffian)

Let M be a three-dimensional manifold with cotangent bundle T^*M . Fix a volume form vol on T^*M and let $\alpha \in \Omega^3(T^*M)$. The Hitchin endomorphism $\hat{\mathcal{J}}_\alpha^{\text{vol}} : T(T^*M) \rightarrow T(T^*M)$ is given by

$$\hat{\mathcal{J}}_\alpha^{\text{vol}}(X) \text{vol} := \phi(\alpha \wedge X \lrcorner \alpha). \quad (2.2.20)$$

In particular, $(\hat{\mathcal{J}}_\alpha^{\text{vol}})^2 = -\text{HPf}_{\text{vol}}(\alpha)\mathbb{I}$ for some function $\text{HPf}_{\text{vol}}(\alpha) \in \mathcal{C}^\infty(T^*M)$, which we call the Hitchin Pfaffian [26, Section 2] and $\text{tr}(\hat{\mathcal{J}}_\alpha^{\text{vol}}) = 0$.

Therefore, we may define an almost (para-)complex structure by scaling the Hitchin endomorphism with respect to the Hitchin Pfaffian, whenever the Hitchin Pfaffian is non-zero. More precisely,

$$\hat{\mathcal{J}}_\alpha := \frac{1}{\sqrt{|\text{HPf}_{\text{vol}}(\alpha)|}} \hat{\mathcal{J}}_\alpha^{\text{vol}}. \quad (2.2.21)$$

is an almost complex structure when $\text{HPf}_{\text{vol}}(\alpha) > 0$ and an almost para-complex structure when $\text{HPf}_{\text{vol}}(\alpha) < 0$. This definition is independent of the choice of volume form, up to an overall sign, since scaling vol in (2.2.20) by a non-zero function $\kappa \in \mathcal{C}^\infty(T^*M)$ causes the Hitchin Pfaffian

to scale as $\kappa^2 \text{HPf}_{\kappa \text{vol}}(\alpha) = \text{HPf}_{\text{vol}}(\alpha)$. As we will only be interested in the sign of the Hitchin Pfaffian and this is invariant under scaling of vol , we omit the subscript on $\text{HPf}(\alpha)$ moving forward.

Now let (ϖ, α) be a Monge–Ampère structure on T^*M , with M a three-dimensional manifold. Using the above, we may define an almost (para-)complex structure $\hat{\mathcal{J}}_\alpha$ associated with the Monge–Ampère structure. Furthermore, given a differential 2-form $\hat{\mathcal{K}}$, which is $(1, 1)$ with respect to $\hat{\mathcal{J}}_\alpha$ and satisfies $\hat{\mathcal{K}}(\hat{\mathcal{J}}_\alpha X, Y) = -\hat{\mathcal{K}}(X, \hat{\mathcal{J}}_\alpha Y)$, we may again define a metric

$$\hat{g}(X, Y) := \hat{\mathcal{K}}(X, \hat{\mathcal{J}}_\alpha Y), \quad (2.2.22)$$

which we call a *Lychagin–Rubtsov metric* for (ϖ, α) . Furthermore, following [35, Definition 4.1], we call the Monge–Ampère equation associated with α *elliptic* if $\text{HPf}(\alpha) > 0$, *hyperbolic* if $\text{HPf}(\alpha) < 0$, and *parabolic* if $\text{HPf}(\alpha) = 0$, in complete analogy with the two-dimensional case.

The standard choice of metric when studying Monge–Ampère structures (ϖ, α) in three dimensions is that given by the Lychagin–Rubtsov invariant [15, Equation 2.1] (see also [19, Section 4.1]), namely:

$$g_\alpha(X, Y) = \frac{(X \lrcorner \alpha) \wedge (Y \lrcorner \alpha) \wedge \varpi}{\text{vol}}, \quad (2.2.23)$$

where we typically choose the Liouville volume form $\text{vol} = \frac{1}{3!} \varpi \wedge \varpi \wedge \varpi$ for concreteness. It has been noted by Banos, see [35, Proposition 3.9], that this metric is almost (para-)Kähler, with closed $(1, 1)$ -form given by ϖ . That is,

$$g_\alpha(X, Y) = \varpi(\hat{\mathcal{J}}X, Y), \quad (2.2.24)$$

so the metric given by the Lychagin–Rubtsov invariant is indeed of Lychagin–Rubtsov type, hence the nomenclature.

Remark 2.2.10 (Even Further Beyond)

Beyond three dimensions, the methods for generating almost (para-)complex structures associated to Monge–Ampère structures get ever more involved and bear little resemblance to (2.2.13). However, the Hitchin endomorphism construction (2.2.20) does generalise to $2m + 1$ dimensional manifolds M — for a more detailed discussion of this construction, see [37, Section 3.5.2] (material from which first appeared in [36, Section 3.1.7, Section 5])

△

2.3 Chapter Summary

In this chapter, we have provided an overview of the tools from (multi-)symplectic geometry which we shall require in subsequent chapters when studying the Poisson equation for the pressure of an incompressible fluid flow. Motivated by the simple example of the Laplace equation, we built up to a formal definition of a Monge–Ampère structure in [Definition 2.1.9](#). We also introduced the notion of generalised solutions to a Monge–Ampère structure in [Definition 2.1.10](#), with these taking the form of $\dim(M)$ -dimensional submanifolds of T^*M on which the Monge–Ampère structure vanishes. Generalised solutions correspond to solutions of Monge–Ampère equations which are permitted to be multivalued or undefined at points in M and may even exhibit singular behaviour. For simplicity, in the remainder of this thesis we consider solutions which are (locally) given by classical solutions; such solutions preclude any singular behaviour.

In [Section 2.2](#), we reviewed some concepts from (para-)complex geometry which naturally arise when studying Monge–Ampère structures in two and three dimensions. In particular, we demonstrated how the associated Monge–Ampère equation can be classified via a geometric quantity, known as the Pfaffian (see [Definition 2.2.4](#)) and Hitchin Pfaffian (see [Definition 2.2.9](#)) in two and three dimensions respectively and how an almost (para-)complex structure can be constructed from the Monge–Ampère structure in each of these cases (see [\(2.2.13\)](#) and [\(2.2.20\)](#)). Furthermore, given a differential two-form which is $(1, 1)$ with respect to said almost (para-)complex structure, we may define an almost (para-)Hermitian metric associated with the Monge–Ampère structure, which we call the Lychagin–Rubtsov metric. Here we observe explicitly for the first time that we are free to make a choice of $(1, 1)$ -form, hence have a family of such metrics to choose from. This is in contrast to earlier works which take the metric to have fixed form given by either [\(2.2.16\)](#) or [\(2.2.23\)](#), depending on the dimension.

In the next chapter, we shall apply the techniques we have outlined here to the Poisson equation for the pressure of an incompressible fluid flow in two dimensions. In doing so, we aim to illustrate that a specific choice of Lychagin–Rubtsov metric encodes the balance of vorticity and strain in the flow, with its signature being related to the ellipticity/hyperbolicity of the equation. We shall also discuss how this geometric picture can be used to obtain topological information about the fluid flow, from the curvature of our metric.



Monge–Ampère Geometry of Two-Dimensional Flows

Now that we have established the geometric tools at our disposal when studying Monge–Ampère equations, let us address the use case proposed in [Chapter 1](#): incompressible fluid flows.

One of the enduring challenges in fluid mechanics is understanding the topology of fluid flows (see the review [55]). In particular, we wish to know how vortices form and what types of topological artefact they introduce. However, no systematic method has been developed to extract such topological information from the governing system of partial differential equations. Indeed, there is not even a universally applicable definition of a vortex, only a number of loosely connected quantitative and qualitative criterion [56–58]. However, in [20], it is noted that the Poisson equation for the pressure (1.0.4) “is by no means fully understood and locally holds the key to the formation of vortex structures through the sign of the Laplacian of the pressure.”

In this chapter, we follow the work of [19, 18, 17] and study the Poisson equation for the pressure in two dimensions through the lens of Monge–Ampère geometry, in order to make first steps towards clarifying the relationship between the the accumulation of vorticity, the Laplacian of pressure, and the topology of fluid flows. By beginning from a covariant formulation of the incompressible Navier–Stokes equations, in contrast to earlier works, which only considered flows on Euclidean background, we are able to extend a number of known results to flows on arbitrary Riemannian manifolds. In particular, we study solutions to the Poisson equation for the pressure via their associated Lagrangian submanifolds of T^*M , as in [Section 2.1.3](#), which we equip with the pull-back of some suitable Lychagin–Rubtsov metric, as per [Remark 2.2.8](#). The curvature of this metric and the Gauß–Bonnet theorem then enable us to describe the topology of regions of the flow.

3.1 The Incompressible Navier–Stokes Equations and Pressure

Before we apply the tools introduced in the previous chapter, let us review some concepts from fluid dynamics, to provide further context for our analysis. We begin with a discussion of the role the sign of the Laplacian of the pressure has in the dominance of vorticity over rate-of-strain in the Euclidean setting, before deriving the Poisson equation for the pressure on an arbitrary Riemannian manifold. This enables us to address the impact of the curvature of the background manifold on the dominance of vorticity. While much of the material in this section is classical and well known, we appreciate that this thesis predominantly appeals to an audience familiar with geometry, who may not have experience in the field of fluid dynamics. As such, we include the discussion here for completeness.

3.1.1 Incompressible Fluid Flows on \mathbb{R}^m and the Weiss Criterion

Let us start by stating the constraints which should be satisfied by a fluid in \mathbb{R}^m , equipped with the standard Euclidean metric $\hat{g}_{ij} = \delta_{ij}$. As in the previous chapter, give \mathbb{R}^m the local coordinates $\{x^i\}_{i=1}^m$. The viscosity of the fluid is given by some scalar $\nu \in \mathbb{R}$, while the pressure is given by some function $p \in \mathcal{C}^\infty(\mathbb{R}^m)$. The (unknown) velocity of the fluid flow is described by a (co-)vector field $v \in \Omega^1(\mathbb{R}^m)$ with components v^i . For simplicity we make the following additional assumptions:

- The fluid is homogeneous, that is, its density is uniform in space and constant in time. Without loss of generality, we set the density to be identically equal to 1.
- The net external (body) force acting on the fluid is divergence-free for all time. This force is represented by another (co-)vector field $c \in \Omega^1(\mathbb{R}^m)$, whose components c^i satisfy $\partial_i c^i = 0$. Recall here that we use Einstein summation convention.

The incompressible Euclidean Navier–Stokes equations are then given by

$$\frac{\partial v^i}{\partial t} = -v^j \partial_j v^i - \partial^i p + \nu \Delta v^i + c^i \quad (3.1.1a)$$

and

$$\partial_i v^i = 0, \quad (3.1.1b)$$

where $\Delta := \partial^i \partial_i$ is the standard Euclidean Laplacian. When the flow is inviscid, $\nu = 0$, and we have the so-called incompressible Euler equations.

Regardless of the value of ν , taking the divergence of (3.1.1a), followed by the trace and an

application of (3.1.1b), yields a Poisson equation for the pressure $p(x)$

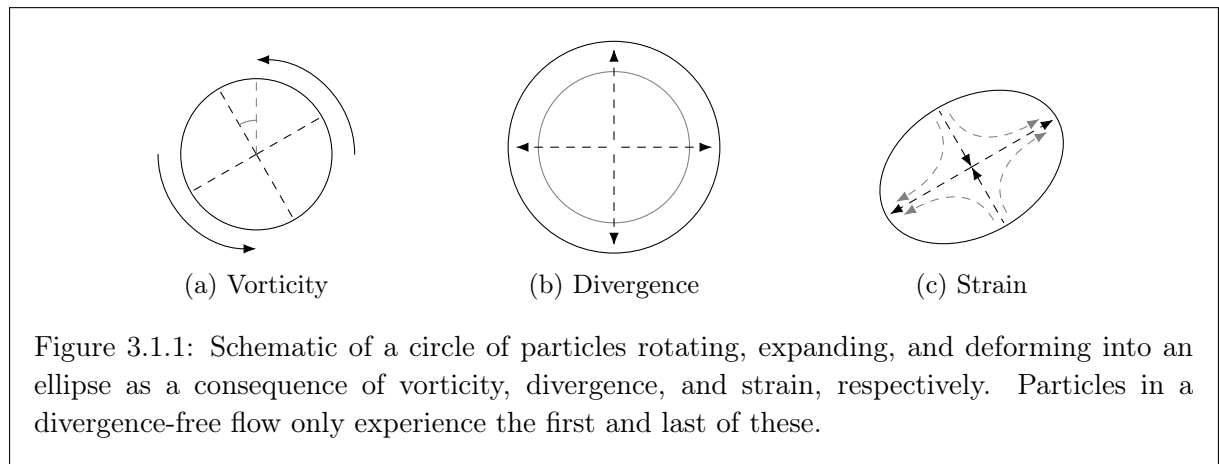
$$\Delta p = -(\partial_i v_j)(\partial_j v_i) = \zeta_{ij}\zeta^{ij} - S_{ij}S^{ij}, \quad (3.1.1c)$$

in terms of the vorticity two-form and rate-of-strain tensor, which are respectively given by

$$\zeta_{ij} := \partial_{[i}v_{j]} \quad \text{and} \quad S_{ij} := \partial_{(i}v_{j)}. \quad (3.1.2)$$

Here and in the following, parentheses (respectively, square brackets) denote normalised symmetrisation (respectively, anti-symmetrisation) of the enclosed indices.

In regions where $\Delta p > 0$, the vorticity term in (3.1.1c) dominates, while in regions where $\Delta p < 0$, the strain term dominates. In order to understand what this tells us about the physical behaviour of the fluid, let us consider what it means for a two-dimensional fluid flow to (locally) have divergence, vorticity, and strain. This description is inspired by [59, Figure 1], which is recreated in Figure 3.1.1 for convenience. The divergence of a fluid flow can be thought of as the amount a fluid “spreads out” or dilates over time — consider a small circle of particles dropped into a region of fluid with divergence d (and negligible vorticity/strain); the radius of the circle will increase by a factor of $(1 + \frac{d}{2})$ per unit time, with the radius of the circle remaining constant in a divergence free flow. The vorticity two-form has a single independent component $\zeta := 2\zeta_{12}$ in two dimensions, which is known as the vorticity. If we drop our small circle of particles into a region of fluid with vorticity ζ , then the circle will rotate through an angle of $\frac{\zeta}{2}$ per unit time. That is, vorticity can be associated with rotation of the fluid. Finally, for an incompressible fluid flow $\partial_1 v_1 + \partial_2 v_2 = 0$, so $S_{11} = \partial_1 v_1 = -\partial_2 v_2 = -S_{22}$ and the rate-of-strain tensor only has two independent components, S_{11} and S_{12} . These correspond to the deformation of our circle of particles into an ellipse with major axis at an angle $\frac{1}{2} \arctan\left(\frac{S_{12}}{S_{11}}\right)$ to the positive x^1 axis.



In two dimensions, Weiss [25] proposed, via dynamical and numerical arguments, local criterion for identifying regions of elliptic/hyperbolic flow in fluids modelled by the incompressible Navier–Stokes equations (in particular, by (3.1.1c)). The so-called *Weiss criterion*, or *Q-criterion*, states that, when $\Delta p > 0$ and vorticity dominates over strain in (3.1.1c), the fluid flow is *elliptic* and tends towards exhibiting rotation. On the other hand, when $\Delta p < 0$ and strain dominates over vorticity, the fluid flow is *hyperbolic* and tends towards exhibiting shearing along some preferred axis (the asymptote of the hyperbola). In essence, the criterion makes explicit the association between the physical descriptions of vorticity and strain above and the behaviour of trajectories around elliptic and hyperbolic fixed points of the stability matrix of (3.1.1c), which is given by the *velocity gradient tensor* $A_{ij} = \partial_j v_i$ and obtained by linearising (3.1.1c).

A geometric argument for the Weiss criterion was later provided by Larchevêque in [23, 24], where it was noted that, for two-dimensional, incompressible flows, the divergence free constraint (3.1.1b) can be used to write the components of the velocity $v(x^1, x^2)$ in terms of a function $\psi \in \mathcal{C}^\infty(\mathbb{R}^2)$, called the *stream function*:

$$v_1 = -\partial_y \psi \quad \text{and} \quad v_2 = \partial_x \psi, \quad (3.1.3)$$

and (3.1.1c) becomes the following Monge–Ampère equation for ψ :

$$\frac{1}{2}\Delta p = \partial_x^2 \psi \partial_y^2 \psi - (\partial_x \partial_y \psi)^2. \quad (3.1.4)$$

As discussed in Remark 2.2.5, the equation type of a Monge–Ampère equation in two dimensions is determined by the sign of its Pfaffian, which is given by $\frac{\Delta p}{2}$ in this case. Consequently, when the strain term in (3.1.1c) dominates, the Laplacian of pressure is negative and the Monge–Ampère equation (3.1.4) is of hyperbolic type. Similarly, when the vorticity term dominates, the Laplacian of pressure is positive, and the Monge–Ampère equation is of elliptic type. Phrases of this form will become a mantra throughout this thesis.

Larchevêque and Weiss both note that, at fixed time t , with $(x^1, x^2) = (x, y)$, the sign of the Laplacian of pressure coincides with the sign of the Gaußian curvature of the graph $S := \{(x, y, z) \in \mathbb{R}^3 \mid z = \psi(x, y)\}$ described by the stream function $\psi(x, y)$. Hence, given a simply connected, open subset $V \subseteq \mathbb{R}^2$ on which vorticity dominates, the Gaußian curvature of the graph S over V , is positive. Recall that the Gaußian curvature is given by the ratio between the first and second fundamental forms on S . Since the induced metric on S is Riemannian, the first fundamental form is positive, so positivity of the Gaußian curvature implies positivity of the second fundamental form. As a result, $S|_V$ must be convex by [60, Lemma 2.4] and the Tietze–Nakajima theorem [61, 62] (see also [63, Theorem 1]). Furthermore, if the boundary of V

is given by a closed streamline $\psi = \text{constant}$, then all streamlines contained within V are closed, convex contours, as we might expect of a vortex.

In three dimensions, vorticity no longer accumulates in elliptic regions, but rather along curves called *vortex lines* or on thin quasi two-dimensional sets which can manifest as *vortex sheets* or as *vortex tubes*. This increase in variety when moving from two to three dimensions results in an increase in complexity in the number of ways “vortex structures” might interact.¹ In order to account for this, a wider cast of criteria for vortices need to be considered. These include the λ_2 -*criterion* [56, Section 4], which considers vortices to be regions where the middle eigenvalue of $\zeta_{ij}\zeta^{jk} + S_{ij}S^{jk}$ is negative, the Δ -*criterion* [58, Section 3], which states that, for points within the core of a vortex, the local streamlines are closed surfaces, closed contours, or spirals in the reference frame of the point, and the Q -*criterion* [64], which considers vortices to be defined by the dominance of the vorticity term in (3.1.1c), as in two dimensions. The relationships between these criteria are explained in more detail in [57, 65].

Consequently, the Poisson equation for the pressure is also employed in studies of the accumulation of vorticity in three dimensions, though is only part of the overall picture. Furthermore, it is no longer possible to invoke the existence of a stream function and derive a Monge–Ampère equation whose ellipticity/hyperbolicity can be studied à la Larchevêque in this case. However, it is still possible to encode (3.1.1c) in terms of differential forms on $T^*\mathbb{R}^3$, as noted by [18, 17]. We shall invoke this observation in Chapter 4, in order to investigate the geometry associated with vortex structures in three dimensions.

Remark 3.1.1 (Evolution of Incompressible Fluid Flows)

Recall that the Poisson equation for the pressure (3.1.1c) holds for solutions of both the incompressible Euler ($\nu = 0$) and incompressible Navier–Stokes equations (3.1.1). We now wish to emphasize the following point: *the converse is NOT true in general*. That is, it is possible to find solutions of (3.1.1c) which are not solutions to the full system (3.1.1). In particular, the Poisson equation for the pressure depends only on time as a parameter and needs to be coupled with an equation describing the evolution of the flow, in order to recover solutions of the full Navier–Stokes or Euler equations. One choice of evolution equation is the *vorticity equation*, given by

$$\frac{D\zeta}{Dt} = \nu\Delta\zeta, \quad (3.1.5)$$

in two dimensions and by

$$\frac{D\zeta^i}{Dt} = (\zeta^j\partial_j)v^i + \nu\Delta\zeta^i, \quad (3.1.6)$$

¹As we shall discuss in Section 4.2.2, these interactions now include linking and knotting together.

in three dimensions. Here $\frac{D}{Dt} := \frac{\partial}{\partial t} + v^i \nabla_i$ is the *material derivative* and $\zeta^i = \frac{1}{2} \varepsilon^{ijk} \zeta_{jk}$ are the independent components of the vorticity in three dimensions, with ε^{ijk} the Levi-Civita symbol. This equation reintroduces the system's dependence on the viscosity ν and holds for the Euler and Navier–Stokes equations in both two and three dimensions. It is an open question as to how we might encode the evolution of solutions within the framework of Monge–Ampère geometry, with geometric flow of Lagrangian submanifolds being a strong candidate, though this falls outside the scope of the thesis.

However, since the set of solutions to the full Navier–Stokes equations is a subset of the solutions to the Poisson equation for the pressure, if we derive properties which hold for all solutions to the latter, then they also hold for solutions to the former, without considering the evolution of the flow. As such, for the remainder of this thesis, we focus on characterising solutions, rather than searching for new solutions themselves.

△

3.1.2 Incompressible Fluid Flows on Manifolds

Now let us consider the case where the domain of our fluid flow is given by a more general m -dimensional Riemannian manifold M with metric² \mathring{g} . In what follows, we reintroduce many of the objects from Section 3.1.1 in covariant form and show how to derive the analogue of (3.1.1c) in this setting.

Denote by d the *exterior derivative* on M and by $\star_{\mathring{g}}$ the *Hodge star* operator with respect to the metric \mathring{g} . Furthermore, let the *codifferential* acting on differential k -forms and the *Hodge Laplacian* be given by

$$\mathring{\delta} := (-1)^{m(k-1)+1} \star_{\mathring{g}} d \star_{\mathring{g}} \quad \text{and} \quad \mathring{\Delta}_H := \mathring{\delta} d + d \mathring{\delta}, \quad (3.1.7)$$

respectively. Finally, define the *norm square* of a differential k -form $\eta \in \Omega^k(M)$ by

$$|\eta|^2 := \frac{\eta \wedge \star_{\mathring{g}} \eta}{\text{vol}_M}, \quad (3.1.8)$$

where vol_M is the *volume form* on M induced by \mathring{g} .

The *velocity (co-)vector field* on M describing our fluid flow is now given by a one-parameter family of differential one-forms $v \in \Omega^1(M)$, parametrised by time $t \in \mathbb{R}$. As in the Euclidean case, let the viscosity be given by a scalar $\nu \in \mathbb{R}$ and the pressure be given by a function

²The metric \mathring{g} may depend on parameter time t , however, since we do not consider the evolution of the flow, we suppress this dependence for simplicity.

$p \in \mathcal{C}^\infty(M)$. We again make the assumptions concerning the net external force and density of the fluid introduced at the start of Section 3.1.1, except the divergence-free net external force is now represented by a family of differential one-forms $c(t) \in \Omega^1(M)$ which satisfy $\mathring{\delta}c \equiv 0$. The incompressible, covariant Navier–Stokes equations for v on M are then given by a system consisting of a flow equation for v :

$$\frac{\partial v}{\partial t} = -(-1)^m \star_{\mathring{g}}(v \wedge \star_{\mathring{g}} dv) - \frac{1}{2} d|v|^2 - dp - \nu \mathring{\Delta}_H v + c, \quad (3.1.9a)$$

supplemented by the divergence-free constraint:

$$\mathring{\delta}v = 0, \quad (3.1.9b)$$

which follows from the continuity equation under our assumptions. Applying the codifferential to (3.1.9a) and using (3.1.9b), we recover the so-called pressure equation, which describes the pressure solely in terms of the velocity field and the background metric:

$$\mathring{\Delta}_H p = -|dv|^2 + \star_{\mathring{g}}(v \wedge \star_{\mathring{g}} \mathring{\Delta}_H v) - \frac{1}{2} \mathring{\Delta}_H |v|^2. \quad (3.1.9c)$$

In order to present our equations in coordinate form, let M have coordinates denoted $\{x^i\}_{i=1}^m$. Recall from Chapter 2 that we lower and raise indices using the metric \mathring{g}_{ij} and its inverse \mathring{g}^{ij} respectively. Let the Levi-Civita connection associated with the metric \mathring{g}_{ij} be denoted by $\mathring{\nabla}_i$ and let $\mathring{\Gamma}_{ij}^k$ be the corresponding Christoffel symbols. Furthermore, the components of the associated Ricci curvature tensor and Riemann curvature tensor are respectively given by

$$\mathring{R}_{ij} := \mathring{R}_{kij}{}^k \quad \text{and} \quad \mathring{R}_{ijk}{}^l := \partial_i \mathring{\Gamma}_{jk}{}^l - \partial_j \mathring{\Gamma}_{ik}{}^l - \mathring{\Gamma}_{ik}{}^m \mathring{\Gamma}_{jm}{}^l + \mathring{\Gamma}_{jk}{}^m \mathring{\Gamma}_{im}{}^l. \quad (3.1.10)$$

The Beltrami Laplacian is

$$\mathring{\Delta}_B := \mathring{g}^{ij} \mathring{\nabla}_i \mathring{\nabla}_j = \mathring{\nabla}^i \mathring{\nabla}_i \quad (3.1.11)$$

and is related to the Hodge Laplacian by the Weitzenböck formula for k -forms $\eta = \frac{1}{k!} \eta_{i_1 \dots i_k} dx^{i_1} \wedge \dots \wedge dx^{i_k}$, i.e.

$$(\mathring{\Delta}_H \rho)_{i_1 \dots i_k} = -\mathring{\Delta}_B \rho_{i_1 \dots i_k} + k \mathring{R}_{j[i_1} \rho^j{}_{i_2 \dots i_k]} + \frac{1}{2} k(k-1) \mathring{R}_{jk[i_1 i_2} \rho^{jk}{}_{i_3 \dots i_k]}. \quad (3.1.12)$$

Finally, the velocity field is given by $v = v_i dx^i$ with $v_i = v_i(t, x^1, \dots, x^m)$ and similarly for the external force c .

The equations (3.1.9) then become

$$\frac{\partial v^i}{\partial t} = -v^j \overset{\circ}{\nabla}_j v^i - \partial^i p + \nu(\overset{\circ}{\Delta}_B v^i - \overset{\circ}{R}^{ij} v_j) + c^i \quad (3.1.13a)$$

and

$$\overset{\circ}{\nabla}_i v^i = 0, \quad (3.1.13b)$$

with the pressure equation given by

$$2f := \overset{\circ}{\Delta}_B p + v^i v^j \overset{\circ}{R}_{ij} = -(\overset{\circ}{\nabla}_i v_j)(\overset{\circ}{\nabla}^j v^i). \quad (3.1.13c)$$

It should be clear that, upon fixing $M = \mathbb{R}^m$ with the standard Euclidean metric $\overset{\circ}{g}_{ij} = \delta_{ij}$, the equations (3.1.9) simplify to the more familiar (3.1.1) introduced in Section 3.1.1.

We define the vorticity two-form and the rate-of-strain tensor on M respectively by

$$\zeta_{ij} := \overset{\circ}{\nabla}_{[i} v_{j]} = \partial_{[i} v_{j]} \quad \text{and} \quad S_{ij} := \overset{\circ}{\nabla}_{(i} v_{j)}, \quad (3.1.14)$$

replacing the partial derivatives in (3.1.2) with covariant ones. By definition, the rate-of-strain tensor vanishes if and only if the velocity vector field is a Killing vector field. Furthermore, define the velocity gradient tensor by

$$A_{ij} := \overset{\circ}{\nabla}_j v_i = S_{ij} - \zeta_{ij}. \quad (3.1.15)$$

and recall that the decomposition of any matrix $B = B_{\text{sym}} + B_{\text{anti}}$, into symmetric and anti-symmetric parts $B_{\text{sym}} = \frac{1}{2}(B + B^T)$ and $B_{\text{anti}} = \frac{1}{2}(B - B^T)$, satisfies the following identity:

$$\text{tr}[B^2] = \text{tr}[B_{\text{sym}} B_{\text{sym}}] - \text{tr}[B_{\text{anti}}^T B_{\text{anti}}]. \quad (3.1.16)$$

Hence, it is possible to rewrite equation (3.1.13c) in the more descriptive form

$$\overset{\circ}{\Delta}_B p + v^i v^j \overset{\circ}{R}_{ij} = -A_j^i A_i^j = \zeta_{ij} \zeta^{ij} - S_{ij} S^{ij}, \quad (3.1.17)$$

from which it is clear that the Laplacian of the pressure of an incompressible fluid flow on an arbitrary Riemannian manifold depends on the vorticity and strain of the flow, in addition to the Ricci curvature of the underlying manifold.

Finally, as a result of the Poincaré lemma, it is always possible to locally solve the divergence-free constraint (3.1.9b) on an open and contractible set $U \subseteq M$, with the velocity

field being given by

$$v = \star_{\mathring{g}} d\psi \quad \text{for } \psi \in \Omega^{m-2}(U) \quad \leftrightarrow \quad v^i = \frac{\sqrt{\det(\mathring{g})}}{(m-2)!} \varepsilon^{i_1 \dots i_{m-1} i} \partial_{i_1} \psi_{i_2 \dots i_{m-1}}, \quad (3.1.18)$$

where $\varepsilon_{i_1 \dots i_m}$ is the Levi-Civita symbol with $\varepsilon_{1 \dots m} = 1$; note that $\varepsilon^{1 \dots m} = \frac{1}{\det(\mathring{g})} \varepsilon_{1 \dots m}$. Upon substituting this expression into (3.1.13c), we obtain a Monge–Ampère type equation for the components of the differential form ψ . Generally, we may refer to $\psi \in \Omega^{m-2}(U)$ as the *stream $(m-2)$ -form*. In the case where $m = 2$, $\psi \in \mathcal{C}^\infty(U)$ is known as the *stream function* and we obtain a genuine Monge–Ampère equation on M , which we shall study in more detail in Section 3.2. From this point onward, we assume that M is equipped with a *good cover*, in which finite intersections of open sets are contractible, to facilitate the use of the Poincaré lemma. We also assume that M is orientable, so that the volume form vol_M exists.

Remark 3.1.2 (Non Divergence-Free External Forces)

It is also possible to apply the equations (3.1.9b) and (3.1.9c) to the case where the net external force c is not divergence-free. In this case, (3.1.13c) becomes

$$\mathring{\Delta}_B p + v^i v^j \mathring{R}_{ij} - \mathring{\nabla}_i c^i = -(\mathring{\nabla}_i v_j)(\mathring{\nabla}^j v^i) \quad (3.1.19)$$

and we redefine f accordingly. We suppress such forces for the remainder of the thesis, since the additional $-\mathring{\nabla}_i c^i$ term remains grouped with the Laplacian of pressure in all subsequent computations and is easily reinstated. Do note, however, that the dominance of vorticity and strain is swayed by external forces which are not divergence free — if the divergence of the external force is negative, the left side of (3.1.19) increases, so the vorticity is more likely to dominate over the strain. The converse holds if the divergence is positive.

△

Remark 3.1.3 (Alternative Covariantisations of Navier–Stokes)

The viscosity term in the Navier–Stokes equations (3.1.9a) may be modified to

$$\frac{\partial v}{\partial t} = -(-1)^m \star_{\mathring{g}}(v \wedge \star_{\mathring{g}} dv) - \frac{1}{2} d|v|^2 - dp - \nu[\mathring{\Delta}_H v + \tilde{c} \mathring{\text{Ric}}(v)], \quad (3.1.20)$$

where $\tilde{c} \in \mathbb{R}$, $\mathring{\text{Ric}}(\rho) := \mathring{R}_i^j \rho_j dx^i$ for any $\rho \in \Omega^1(M)$ and we omit any external forces. In coordinates, (3.1.13a) then becomes

$$\frac{\partial v^i}{\partial t} = -v^j \mathring{\nabla}_j v^i - \partial^i p + \nu[\mathring{\Delta}_B v^i - (\tilde{c} + 1) \mathring{R}^{ij} v_j]. \quad (3.1.21)$$

Evidently, in the flat case, the extra term vanishes and this modified equation again reduces to the standard equation (3.1.1a). When $\tilde{c} = 0$, we return to the version of the covariant Navier–Stokes equations given by (3.1.13a), which was perhaps first studied in the seminal work [66] and more recently in e.g. [67, 68]. The case $\tilde{c} = -1$ was discussed in e.g. [69] and the case $\tilde{c} = -2$ in e.g. [66, 70, 71].

Under the assumption of the divergence-free constraint (3.1.13b), when $\tilde{c} = -2$, for example, it can straightforwardly be seen that the viscosity term is given in terms of the rate-of-strain tensor S_{ij} as

$$\mathring{\Delta}_B v^i + \mathring{R}^{ij} v_j = 2 \mathring{\nabla}_j S^{ij} . \quad (3.1.22)$$

Hence, the viscosity term drops out from the Navier–Stokes equations for velocity fields that preserve the metric (i.e. that are Killing), in this case. Generally, with the \tilde{c} -term switched on, the pressure equation (3.1.13c) takes the form

$$\mathring{\Delta}_B p = -(\mathring{\nabla}_i v_j)(\mathring{\nabla}^j v^i) - v^i v^j \mathring{R}_{ij} - \nu \tilde{c} (\mathring{R}_{ij} S^{ij} + \frac{1}{2} v^i \partial_i \mathring{R}) , \quad (3.1.23)$$

where $\mathring{R} := \mathring{g}^{ij} \mathring{R}_{ij}$ is the curvature scalar and we have used the well-known identity

$$\mathring{\nabla}^i (\mathring{R}_{ij} - \frac{1}{2} \mathring{g}_{ij} \mathring{R}) = 0 . \quad (3.1.24)$$

In particular, when $\tilde{c} \neq 0$, we gain an additional term which depends on both the curvature of M and the viscosity of the fluid, (3.1.23) distinguishes between the Euler and Navier–Stokes equations.

Henceforth, we shall *always* assume that $\tilde{c} = 0$, so that we can address both the viscous and inviscid case simultaneously. Our results and conclusions remain unchanged for non-vanishing \tilde{c} and all of the formulæ can be readily adjusted to accommodate this, via a straightforward redefinition of f in (3.1.13c) to include the viscosity term.

△

3.2 Geometric Properties of Two-Dimensional Incompressible Fluid Flows

Now let us specialise once more to incompressible fluid flows on two-dimensional Riemannian manifolds and show how to obtain a Monge–Ampère equation in this case. For a two-dimensional manifold M , the components (3.1.10) of the Ricci and Riemann curvature tensors simplify to

$$\mathring{R}_{ij} = \frac{\mathring{R}}{2} \mathring{g}_{ij} \quad \text{and} \quad \mathring{R}_{ijk}{}^l = \mathring{R} \mathring{g}_{k[j} \delta_{i]}{}^l \quad (3.2.1)$$

respectively, where \mathring{R} is the Ricci curvature scalar.

As discussed around (3.1.18), applying the Poincaré lemma to the divergence-free constraint (3.1.9b) in two dimensions yields the velocity components v^i in terms of the stream function $\psi \in \mathcal{C}^\infty(U)$, on open and contractible sets $U \subseteq M$. Explicitly, we find

$$v^i = -\sqrt{\det(\mathring{g})}\varepsilon^{ij}\partial_j\psi, \quad (3.2.2)$$

and the pressure equation (3.1.13c) becomes³

$$\frac{1}{2}\mathring{\Delta}_{BP} + \frac{\mathring{R}}{4}|\mathrm{d}\psi|^2 = \det(\mathring{g}^{-1}\mathrm{Hess}(\psi)) \iff \frac{1}{2}\mathring{\nabla}^i\partial_i p + \frac{\mathring{R}}{4}(\mathring{\nabla}^i\psi)(\partial_i\psi) = \det(\mathring{\nabla}^i\partial_j\psi) \quad (3.2.3)$$

on open and contractible sets $U \subseteq M$, where $\mathrm{Hess}(\psi)$ denotes the *Hessian* of the function $\psi \in \mathcal{C}^\infty(U)$ and is given in local coordinates by $\mathrm{Hess}(\psi) = \mathring{\nabla}_i\partial_j\psi$. Hence, solving (3.2.3) for the stream function and substituting back into (3.2.2) allows us to locally describe the velocity vector field. Furthermore, (3.2.3) can be understood as a Monge–Ampère equation for the stream function and reduces to (3.1.4) in the Euclidean case, as expected.

Importantly, it is possible to find Monge–Ampère structures on T^*M from which the Poisson equation for the pressure (3.2.3) arises. Indeed, upon fixing the notation

$$\hat{f} := \frac{1}{2}\mathring{\Delta}_{BP} + \frac{\mathring{R}}{4}|q|^2 \quad \text{and} \quad \mathring{\nabla}q_i := \mathrm{d}q_i - \mathrm{d}x^j\mathring{\Gamma}_{ji}{}^k q_k, \quad (3.2.4a)$$

with $\mathring{\Gamma}_{ij}{}^k$ the Christoffel symbols for \mathring{g}_{ij} , it is readily checked that the differential forms

$$\begin{aligned} \omega &:= \mathring{\nabla}q_i \wedge \mathrm{d}x^i = \mathrm{d}q_i \wedge \mathrm{d}x^i, \\ \alpha &:= \frac{\sqrt{\det(\mathring{g})}}{2}[\varepsilon^{ij}\mathring{\nabla}q_i \wedge \mathring{\nabla}q_j - \hat{f}\varepsilon_{ij}\mathrm{d}x^i \wedge \mathrm{d}x^j] \end{aligned} \quad (3.2.4b)$$

form a Monge–Ampère structure on T^*M . Indeed, ω is the canonical symplectic form and we shall confirm in Section 4.2 that α is both closed and ω -effective. Furthermore, while $\iota^*\omega = 0$ is automatic on Lagrangian submanifolds $\iota : L \hookrightarrow T^*M$, on submanifolds which are given by a section $\mathrm{d}\psi : M \rightarrow T^*M$; $\iota(L) = \mathrm{d}\psi(M)$, the condition $\iota^*\alpha = 0$ is equivalent to requiring that the function $\psi \in \mathcal{C}^\infty(M)$ satisfies the Monge–Ampère equation (3.2.3). That is, ψ is a classical solution to the Poisson equation for the pressure (the corresponding statement holds for locally-a-section submanifolds). In conclusion, the Monge–Ampère equation (3.2.3) arises from the Monge–Ampère structure (3.2.4b).

Recall that $\frac{1}{2}\omega \wedge \omega$ is the Liouville volume form with respect to ω and is therefore non-

³The determinant of a matrix-like quantity $A^i{}_j$ given by $\det(A^i{}_j) = \frac{1}{2}\det(\mathring{g})\varepsilon_{i_1,i_2}A^{i_1}{}_{j_1}A^{i_2}{}_{j_2}\varepsilon^{j_1,j_2}$.

vanishing. It then follows from the equation

$$\alpha \wedge \alpha = \hat{f} \overset{\circ}{\nabla} q_i \wedge dx^i \wedge \overset{\circ}{\nabla} q_j \wedge dx^j = \hat{f} \omega \wedge \omega, \quad (3.2.5)$$

that $\text{Pf}_\omega(\alpha) = \hat{f}$ and α is non-degenerate if and only if $\hat{f} \neq 0$. Furthermore, pulling back α via the velocity components $q_i \rightarrow v_i = (\star_g d\psi)_i$, i.e. to submanifolds $\iota : L \hookrightarrow T^*M$ with $\iota(L) = \star_g d\psi(M)$, again yields (3.2.3) for incompressible fluid flows; This observation informs the alternative choice of Monge–Ampère structure chosen in Section 4.1, which more naturally generalises to higher dimensions.

Next, as discussed around (2.2.13), we associate with the Monge–Ampère structure (3.2.4b) an endomorphism $\hat{J} : T(T^*M) \rightarrow T(T^*M)$ of vector fields defined by

$$\frac{\alpha}{\sqrt{|\hat{f}|}} =: \hat{J} \lrcorner \omega, \quad (3.2.6)$$

where \lrcorner is the interior product with respect to ω defined in Definition 2.1.1, \hat{f} is defined as in (3.2.4a), and we assume that \hat{f} does not vanish. Furthermore, \hat{J} is an almost complex structure when $\hat{f} > 0$, in which case the Monge–Ampère equation (3.2.3) is elliptic, and an almost para-complex structure when $\hat{f} < 0$, in which case the Monge–Ampère equation (3.2.3) is hyperbolic (see Section 2.2.1 for details).

Applying Theorem 2.2.6 to α tells us that the endomorphism (3.2.6) is integrable if and only if \hat{f} is constant, in which case \hat{J} is a (para-)complex structure and the Poisson equation for the pressure (3.2.3) is locally symplectomorphic to the Laplace equation if $\hat{f} > 0$ or wave equation if $\hat{f} < 0$ (see Appendix B.2 for the computation). Therefore, when \hat{f} is constant, solutions of the Poisson equation for the pressure are locally given by solutions to the Laplace or wave equation. However, \hat{f} being constant implies that M is flat⁴ and $\hat{f} = \frac{1}{2} \overset{\circ}{\Delta}_B p$. By contraposition, if M is non-flat, then \hat{f} is non-constant and \hat{J} is not integrable. Consequently, solutions to the Poisson equation for the pressure (3.2.3) are not locally given by solutions to the Laplace or wave equation for fluid flows on curved surfaces.

Furthermore, by Proposition 2.2.7, we can always find a differential two-form \hat{K} which is of type (1, 1) with respect to \hat{J} , such that $\hat{K} \wedge \omega = 0$, $\hat{K} \wedge (\hat{J} \lrcorner \omega) = 0$, and $\hat{K} \wedge \hat{K} \neq 0$. Explicitly, we may take

$$\hat{K} := -\sqrt{|\hat{f}|} \overset{\circ}{\nabla} q_i \wedge \star_g dx^i. \quad (3.2.7)$$

where \hat{f} is again assumed to be non-vanishing, such that \hat{K} is well defined and satisfies the

⁴Since $\hat{f} = \frac{1}{2} \overset{\circ}{\Delta}_B p + \frac{\hat{R}}{4} |q|^2$ only depends on q via its norm square, which is not constant, \hat{R} must vanish.

aforementioned constraints. The corresponding Lychagin–Rubtsov metric $\hat{g}(X, Y) := \hat{K}(X, \hat{J}Y)$ on T^*M is then explicitly given by

$$\hat{g} = \frac{1}{2}\hat{f}\hat{g}_{ij}dx^i \odot dx^j + \frac{1}{2}\hat{g}^{ij}\overset{\circ}{\nabla}q_i \odot \overset{\circ}{\nabla}q_j . \quad (3.2.8)$$

Evidently, in the elliptic case, when $\hat{f} > 0$, the metric \hat{g} is Riemannian, whilst in the hyperbolic case, when $\hat{f} < 0$, the metric is Kleinian. As discussed in Remark 2.2.8 and illustrated in Appendix B.1, the triple $(\omega, \alpha, \hat{K})$ define an almost (pseudo-)quaternionic structure. Additionally, (3.2.7) is promoted to a (para-)Kähler form if and only if it is closed, which occurs precisely when \hat{f} is constant. It follows that \hat{f} being constant is a necessary and sufficient condition for (3.2.8) to be a (para-)Kähler metric and $(\omega, \alpha, \hat{K})$ to define a (pseudo-)quaternionic Kähler structure à la [16].

In two dimensions, the vorticity two-form (3.1.14) can be written as

$$\zeta_{ij} = \frac{1}{2}\sqrt{\det(\hat{g})}\varepsilon_{ij}\zeta \quad \text{with} \quad \zeta := \overset{\circ}{\Delta}_B\psi \quad \implies \quad \zeta_{ij}\zeta^{ij} = \frac{1}{2}\zeta^2 , \quad (3.2.9)$$

where ζ is referred to as the *vorticity* of the flow, in the spirit of the Euclidean discussion following (3.1.2). Using this, it can easily be verified that the pull-back $g := \iota^*\hat{g}$ of the Lychagin–Rubtsov metric (3.2.8), to a Lagrangian submanifold $\iota : L \hookrightarrow T^*M$ described by a section $\iota(L) = d\psi(M)$, has the form

$$g = \frac{1}{2}g_{ij}dx^i \odot dx^j \quad \text{with} \quad g_{ij} := \zeta\overset{\circ}{\nabla}_i\partial_j\psi , \quad (3.2.10)$$

where $dx^i \odot dx^j = dx^i \otimes dx^j + dx^j \otimes dx^i$ is the *symmetric tensor product* and we have used that

$$\det\left(\overset{\circ}{\nabla}^k\partial_l\psi\right)g_{ij} + \overset{\circ}{g}^{kl}(\overset{\circ}{\nabla}_i\partial_k\psi)(\overset{\circ}{\nabla}_j\partial_l\psi) = \overset{\circ}{\Delta}_B\psi\overset{\circ}{\nabla}_i\partial_j\psi , \quad (3.2.11)$$

and

$$f := \iota^*\hat{f} = \frac{1}{2}\overset{\circ}{\Delta}_B\psi + \frac{\overset{\circ}{R}}{4}|d\psi|^2 = \det\left(\overset{\circ}{\nabla}^i\partial_j\psi\right) . \quad (3.2.12)$$

When both $\text{tr}(\overset{\circ}{g}^{ik}g_{kj}) > 0$ and $\det(\overset{\circ}{g}^{ik}g_{kj}) > 0$, it follows that g is Riemannian.⁵ As $\text{tr}(\overset{\circ}{g}^{ik}g_{kj}) = \zeta^2$, the former condition is always satisfied, while

$$\det(\overset{\circ}{g}^{ik}g_{kj}) = \zeta^2 \det\left(\overset{\circ}{\nabla}^i\partial_j\psi\right) = \zeta^2 f \quad (3.2.13)$$

implies that the latter is satisfied if and only if $f > 0$. Similarly, g is Kleinian when $f < 0$. Hence, the signature of g is independent of the sign of the vorticity (3.2.9) and only depends on

⁵Recall that the trace and determinant are defined for matrix-like quantities A^i_j .

the sign of f . Furthermore, the pull-back metric degenerates on regions where $f = 0$ or $\zeta = 0$.

Upon comparing (3.1.17) and (3.2.12), we find that $f = \frac{1}{2}(\zeta_{ij}\zeta^{ij} - S_{ij}S^{ij})$. Hence, when $f > 0$ and the metric g is Riemannian, vorticity dominates and the Monge–Ampère equation (3.2.3) is elliptic. On the other hand, when $f < 0$ and g is Kleinian, strain dominates and (3.2.3) is hyperbolic. This statement is essentially a covariantisation of the Weiss criterion for a vortex, as given in [23, 24] and discussed in Section 3.1.1, which allows us to diagnose the dominance of vorticity and strain in fluid flows on arbitrary two-dimensional Riemannian manifolds (and, as we shall see in Section 4.2, also in higher dimensions), while accounting for the curvature of said background manifold. In particular, as $f = \frac{1}{2}\hat{\Delta}_B p + \frac{\hat{R}}{4}|\mathrm{d}\psi|^2$, if the curvature of M is positive, then vorticity tends towards dominating over rate-of-strain, when compared to a fluid flow on a flat background with the same Laplacian of pressure. The converse holds if M is negatively curved. This makes intuitive sense; anecdotally, removing the plug from a curved basin is more likely to result in a large vortex than removing the plug from a flat-bottomed sink.

3.2.1 Curvature and Topology of Two-Dimensional Incompressible Fluid Flows

Now that we have a Weiss-type criterion for fluid flows on two-dimensional manifolds, let us consider how we might use the geometry we have defined to obtain topological information about the flow, beginning with the Ricci curvature scalars associated with the Lychagin–Rubtsov metric (3.2.8) and its pull-back (3.2.10). For more detail regarding the calculation of the curvatures, see Appendix C.

The Ricci curvature scalar for the metric (3.2.8) is given by (C.2.12b), which, when specialising to two dimensions, simplifies to:

$$\begin{aligned} \hat{R} = & \frac{1}{\hat{f}}\hat{R} - \frac{1}{4\hat{f}}\hat{R}_{ijk}{}^l\hat{R}^{ijkm}q_kq_m - \hat{\Delta}_B \log(|\hat{f}|) \\ & - \hat{g}_{ij} \left[\frac{\partial^2}{\partial q_i \partial q_j} \log(|\hat{f}|) - \frac{1}{2} \frac{\partial}{\partial q_i} \log(|\hat{f}|) \frac{\partial}{\partial q_j} \log(|\hat{f}|) \right], \end{aligned} \quad (3.2.14)$$

where $\hat{\Delta}_B$ is the Beltrami Laplacian with respect to \hat{g} , defined analogously to (3.1.11).

Writing down the Ricci curvature scalar for the pull-back metric (3.2.10) requires a little more in the way of notation. In particular, we write

$$\psi_{i_1 \dots i_n} := \hat{\nabla}_{(i_1} \cdots \hat{\nabla}_{i_{n-1}} \partial_{i_n)} \psi. \quad (3.2.15)$$

for $n \in \mathbb{N}$. For $n > 1$, a brief calculation shows that $\psi_{i_1 \dots i_n}$ can be expressed in terms of the

rate-of-strain tensor (3.1.14) and the vorticity (3.2.9) as

$$\psi_{i_1 \dots i_n} = -\sqrt{\det(\mathring{g})} \mathring{g}^{jk} \varepsilon_{j(i_1} \mathring{\nabla}_{i_2} \cdots \mathring{\nabla}_{i_{n-1}} S_{i_n)k} + \frac{1}{2} \mathring{g}_{(i_1 i_2} \mathring{\nabla}_{i_3} \cdots \mathring{\nabla}_{i_{n-1}} \partial_{i_n)} \zeta. \quad (3.2.16)$$

Using (3.2.9), we can then write the metric (3.2.10) as $g_{ij} = \zeta \tilde{g}_{ij}$ with $\tilde{g}_{ij} := \psi_{ij}$, which gives the pull-back metric g a conformal structure with conformal factor $|\zeta|$, when $\zeta \neq 0$. Hence, the Christoffel symbols $\Gamma_{ij}{}^k$ for g_{ij} take the form

$$\Gamma_{ij}{}^k = \tilde{\Gamma}_{ij}{}^k + \partial_{(i} \delta_{j)}{}^k \log(|\zeta|) - \frac{1}{2} \tilde{g}_{ij} \tilde{g}^{kl} \partial_l \log(|\zeta|), \quad (3.2.17a)$$

where \tilde{g}^{ij} and $\tilde{\Gamma}_{ij}{}^k$ denote the inverse and Christoffel symbols of the Hessian metric \tilde{g}_{ij} respectively, with the latter given by

$$\tilde{\Gamma}_{ij}{}^k = \mathring{\Gamma}_{ij}{}^k + \frac{1}{2} \Upsilon_{ijl} \tilde{g}^{lk} \quad \text{with} \quad \Upsilon_{ijk} := \psi_{ijk} + \frac{4}{3} \psi_l \mathring{R}_{k(ij)}{}^l. \quad (3.2.17b)$$

Consequently, the curvature scalar of g_{ij} is given by

$$R = \frac{1}{\zeta} \left\{ \tilde{R} - \frac{1}{\sqrt{|\det(\tilde{g})|}} \partial_i \left[\sqrt{|\det(\tilde{g})|} \tilde{g}^{ij} \partial_j \log(|\zeta|) \right] \right\}, \quad (3.2.18a)$$

where \tilde{R} is the curvature scalar of the Hessian metric \tilde{g}_{ij} ,

$$\begin{aligned} \tilde{R} = & \frac{1}{2} \tilde{g}^{ij} \mathring{g}_{ij} \mathring{R} - \frac{1}{4} \tilde{g}^{ij} \tilde{g}^{kl} \tilde{g}^{mn} (\Upsilon_{ijm} \Upsilon_{kln} - \Upsilon_{ikm} \Upsilon_{jln}) \\ & + \frac{2}{3} \tilde{g}^{ij} \tilde{g}^{kl} [\psi_{mn} (\delta_i^m \mathring{R}_{j(kl)}{}^n - \delta_j^m \mathring{R}_{l(ik)}{}^n) + \psi_m (\mathring{\nabla}_i \mathring{R}_{j(kl)}{}^m - \mathring{\nabla}_j \mathring{R}_{l(ik)}{}^m)]. \end{aligned} \quad (3.2.18b)$$

Importantly, no fourth-order derivatives of the stream function appear, and in that sense, the curvature scalar of the pull-back metric (3.2.10) is generated by gradients of vorticity and strain (which are given by third-order derivatives of ψ , see (3.2.16)). Additionally, the first-order partial derivatives of the stream function occur explicitly in (3.2.18b), hence the curvature scalar also depends on the components of the velocity directly.

Recall from (3.2.1) that, in two dimensions, the Riemann curvature tensor is fully determined by the Ricci curvature scalar. Hence, all *curvature singularities* are given by singularities in the Ricci curvature scalar. We can therefore use (3.2.18) to assess⁶ whether degeneracies of the pull-back metric (3.2.10) correspond to curvature singularities, hence persist under coordinate changes, or if they are the result of a poor choice of coordinates i.e. correspond to *coordinate singularities*. The conformal structure of the metric (3.2.10) isolates a factor of $\frac{1}{\zeta}$

⁶Due to the complexity of the derived formulæ, these assessments will generally be made on an example-by-example basis, see Section 3.3.

in (3.2.18a), suggesting that contours along which vorticity vanishes may produce curvature singularities. However, this does not hold in general — see Section 3.3.2 for a counterexample, where vanishing vorticity induces degeneracies in the pull-back metric (3.2.10), but not in the curvature (3.2.18a). Evidently then, whether or not we have curvature singularities depends on which solution of (3.2.3) we consider. For further discussion on the singularity structures that can occur in metrics of Lychagin–Rubtsov type and their pull-backs, see [44] where such metrics are studied on generalised solutions, in the context of semi-geostrophic theory.

Now consider an incompressible Navier–Stokes flow on a Riemannian background manifold M and let $\iota : L \hookrightarrow T^*M$ be a locally-a-section submanifold, which is additionally a generalised solution of the Monge–Ampère structure (ω, α) given in (3.2.4b). Then for each $y \in L$, there exists open subsets $U \subseteq M$ and $V \subseteq L$, with $y \in V$, along with a function $\psi \in \mathcal{C}^\infty(U)$, such that $\iota(V) = d\psi(U)$ and ψ is a stream function solving (3.2.3) on $U \subseteq M$. As discussed in Appendix A and Section 2.1.3, the canonical projection $\pi : T^*M \rightarrow M$ then defines a diffeomorphism $\pi|_V := \pi \circ \iota : V \rightarrow U$, with inverse given by $\iota^{-1} \circ d\psi$. Suppose that, for some choice of y , there exists a compact region $\Sigma \subseteq U$ on which $f > 0$. We can then define a corresponding compact region $L_\Sigma \subseteq V \subseteq L$ by $\iota(L_\Sigma) := d\psi(\Sigma)$. In particular, if $\iota(L)$ is described globally by a section $d\psi : M \rightarrow T^*M$, then we can do this for any compact subset of $\Sigma \subseteq M$.

As L_Σ is given by a section over Σ , we can use (3.2.18a) to compute the Ricci curvature scalar associated with the pull-back metric (3.2.10) on L_Σ . It is therefore natural to consider the question of how we might use the local Gauß–Bonnet theorem to relate the geometry of L_Σ , as given by the Ricci scalar curvature, to its topology, as given by the Euler characteristic $\chi(L_\Sigma)$. For convenience, we state a version of this theorem below and direct the reader to [72, Theorem 4.2] for further details.

Theorem 3.2.1 (Local Gauß–Bonnet Theorem)

Let Σ be a two-dimensional, compact, oriented manifold with Riemannian metric g . Suppose that Σ has a boundary composed of disjoint, simple, closed, piecewise regular, piecewise arc-length parametrised curves γ_α , such that $\partial\Sigma = \bigcup_\alpha \gamma_\alpha$. Let R be the Ricci curvature scalar of the metric g , vol_Σ the volume form on Σ , and κ the geodesic curvature. Furthermore, let φ_β be the exterior angles at the non-smooth points of the boundary $\partial\Sigma$. Then, the Euler number $\chi(\Sigma)$ of Σ is given by

$$\frac{1}{2} \int_\Sigma \text{vol}_\Sigma R + \sum_\alpha \int_{\gamma_\alpha} ds \kappa(\gamma_\alpha(s)) + \sum_\beta \varphi_\beta = 2\pi\chi(\Sigma). \quad (3.2.19)$$

Given an arc-length parametrised curve $\gamma : s \rightarrow (y^1(s), y^2(s))$ in two dimensions, we may

express its *geodesic curvature* κ at $\gamma(s)$ using Beltrami's formula:

$$\kappa(\gamma(s)) = \sqrt{|\det(g(y(s)))|} \varepsilon_{ij} \dot{y}^i(s) [\ddot{y}^j(s) + \Gamma_{kl}^j(y(s)) \dot{y}^k(s) \dot{y}^l(s)] , \quad (3.2.20)$$

where superposed dots indicate derivatives with respect to the arc-length parameter s .

Let $\Sigma \subseteq U \subseteq M$ and $L_\Sigma \subseteq V \subseteq L$ be as defined before [Theorem 3.2.1](#). As U and V are diffeomorphic, it follows that $\chi(\Sigma) = \chi(L_\Sigma)$. Consider the case where Σ is bounded by a simple, regular, closed curve $c : \mathbb{R} \rightarrow U$, such as a closed, isovortical contour or a closed streamline, for example. Then Σ is homeomorphic to a disc and so, $\chi(L_\Sigma) = 1$. As $\iota^{-1} \circ d\psi : U \rightarrow V$ is a diffeomorphism, it follows that the boundary of $L_\Sigma = \iota^{-1} \circ d\psi(\Sigma)$ is given by $\gamma := (\iota^{-1} \circ d\psi) \circ c : \mathbb{R} \rightarrow V$, which is also a simple, regular, closed curve and may be assumed to be arc-length parametrised without loss of generality. Recall that we impose $f > 0$ on Σ , such that the metric (3.2.10) remains Riemannian on L_Σ . Putting this all together, (3.2.19) evaluates to

$$\int_\gamma ds \kappa(\gamma(s)) = 2\pi - \frac{1}{2} \int_{L_\Sigma} \text{vol}_{L_\Sigma} R , \quad (3.2.21)$$

on L_Σ , where R is given by (3.2.18a).

That is, the mean curvature of the boundary of L_Σ is determined by the average curvature of its interior. Noting (3.2.18a), (3.2.18b), and (3.2.20), we remark that at a formal qualitative level, the local Gauß–Bonnet relation (3.2.21) is a statement to the effect of⁷

$$\begin{aligned} \text{mean curvature of the boundary of } L_\Sigma &= \\ &= 2\pi - \text{mean gradients of vorticity and strain} . \end{aligned} \quad (3.2.22)$$

In this sense, we can use Monge–Ampère geometry to assign a topological quantity to a ‘vortex.’ Here, the vortex is described by the graph of the gradient of the stream function ψ over some compact region of M , on which $f > 0$ and which is bounded by a closed streamline. This result illuminates the relationship between physical properties of a flow, such as the gradients of vorticity and strain, and the topology of what we might call a vortex. The Gauß–Bonnet theorem can also be extended to cases where the pull-back metric (3.2.10) is Kleinian, under certain restrictions pertaining to the boundary ∂L_Σ — e.g. it should have no null segments — however, the link between the Gauß–Bonnet theorem and topology as quantified by the Euler characteristic becomes more tenuous [73, 74]. On regions where the projection $\pi|_L$ fails to be a (local) diffeomorphism and L is no longer locally-a-section, then singular behaviour should be

⁷Recall here that the boundary of L_Σ is given by the image of the boundary of Σ , that is, the image of the closed stream-line bounding a candidate vortex in M , under $d\psi$.

anticipated in the pull-back metric g and Ricci curvature scalar R , as observed by [44] in the context of the semi-geostrophic equations for meteorological flows.

3.3 Incompressible Fluids on the Euclidean Plane — Examples

In this section, we adopt the notation $x := x^1$ and $y := x^2$ and consider the case of flows in \mathbb{R}^2 with Euclidean background metric $\mathring{g}_{ij} = \delta_{ij}$, to verify consistency with known results. We begin by providing the simplified formulæ in this setting, before computing the Lychagin–Rubtsov metric and its pull-back for the classical example of a two-dimensional Taylor–Green vortex. While not all solutions to the Poisson equation for the pressure (3.1.1c) are solutions to the Navier–Stokes equations (3.1.1), our example of the Taylor–Green vortex is in fact a solution to the full system. For more examples, see [1, Section 2.3].

3.3.1 Simplified Formulæ in the Euclidean Case

For reference, we summarise here the relevant simplified formulæ in the Euclidean case. Evidently, setting $\mathring{g}_{ij} = \delta_{ij}$ implies $\mathring{R} = 0$, so for \hat{f} as given in (3.2.4a) and f as given in (3.1.13c), we find

$$\hat{f} = \frac{1}{2}\Delta p = \partial_x^2 \psi \partial_y^2 \psi - (\partial_x \partial_y \psi)^2 = f \quad \text{with} \quad \Delta := \partial_x^2 + \partial_y^2. \quad (3.3.1)$$

Hence, the Lychagin–Rubtsov metric (3.2.8) on $T^*\mathbb{R}^2$ takes the form

$$\hat{g} = \text{diag}(f\delta_{ij}, \delta^{ij}), \quad (3.3.2)$$

with its signature dictated by the sign of f . This metric is degenerate if and only if $f = 0$ and the corresponding curvature scalar (3.2.14) becomes

$$\hat{R} = \frac{1}{f^3}(\partial_x f \partial_x f + \partial_y f \partial_y f - f \Delta f). \quad (3.3.3)$$

Thus, at a stationary point of f , where $\partial_x f = \partial_y f = 0$, the sign of \hat{R} is determined solely by the sign of Δf . Consequently, when f accumulates at a local maximum, $\Delta f < 0$ and $\hat{R} > 0$.

The vorticity (3.2.9) is simply $\zeta = \Delta \psi$, where $\psi = \psi(x, y)$ is the stream function, so the pull-back metric (3.2.10) becomes

$$g = \zeta \begin{pmatrix} \partial_x^2 \psi & \partial_x \partial_y \psi \\ \partial_x \partial_y \psi & \partial_y^2 \psi \end{pmatrix} = \frac{\zeta}{2} \begin{pmatrix} \zeta + 2S_{xy} & -2S_{xx} \\ -2S_{xx} & \zeta - 2S_{xy} \end{pmatrix}, \quad (3.3.4)$$

where $S_{xx} = -S_{yy}$ and S_{xy} are the components of rate-of-strain tensor (3.1.14) which describe a shearing deformation at an angle of $\frac{1}{2} \arctan\left(\frac{S_{xy}}{S_{xx}}\right)$, as discussed in Section 3.1.1. In contrast

to the metric (2.2.16), which was studied in [16, 19] and whose pull-back to classical solutions is simply the Hessian of ψ , the vorticity appears as a conformal factor in the metric g . As a result, our metric g is degenerate when the vorticity vanishes, in addition to when the Hessian matrix is degenerate, i.e. where $f = 0$. We also note that when f depends on time t , then the metric (3.3.2) will depend on t as a parameter. The same is true for (3.3.4), via the time-dependence of vorticity and rate-of-strain. The one-parameter families of metrics (3.3.2) and (3.3.4) will thus evolve according to either the Euler or the Navier–Stokes equations, as discussed in Remark 3.1.1.

The eigenvalues of (3.3.4) are given by

$$E_{\pm} = \frac{1}{2}(\zeta^2 \pm |\zeta|D_R) \quad \text{with} \quad D_R^2 := 4(\partial_x \partial_y \psi)^2 + (\partial_x^2 \psi - \partial_y^2 \psi)^2, \quad (3.3.5)$$

where D_R is the *resultant deformation*, which [59] note is a (rotational) invariant of the velocity gradient tensor $A_{ij} = \partial_j v_i$ defined in (3.1.15). Furthermore,

$$D_R^2 = \zeta^2 - 4f, \quad (3.3.6)$$

so the eigenvalues take the same sign for $f > 0$ and opposite sign for $f < 0$, provided they are both non-zero. This tells us that the degeneracies of (3.3.4) depend solely on ζ and f .

Finally, the curvature scalars (3.2.18) reduce to

$$R = \frac{1}{\zeta} \left\{ \tilde{R} - \frac{1}{\sqrt{|\det(\tilde{g})|}} \partial_i [\sqrt{|\det(\tilde{g})|} \tilde{g}^{ij} \partial_j \log(|\zeta|)] \right\}, \quad (3.3.7a)$$

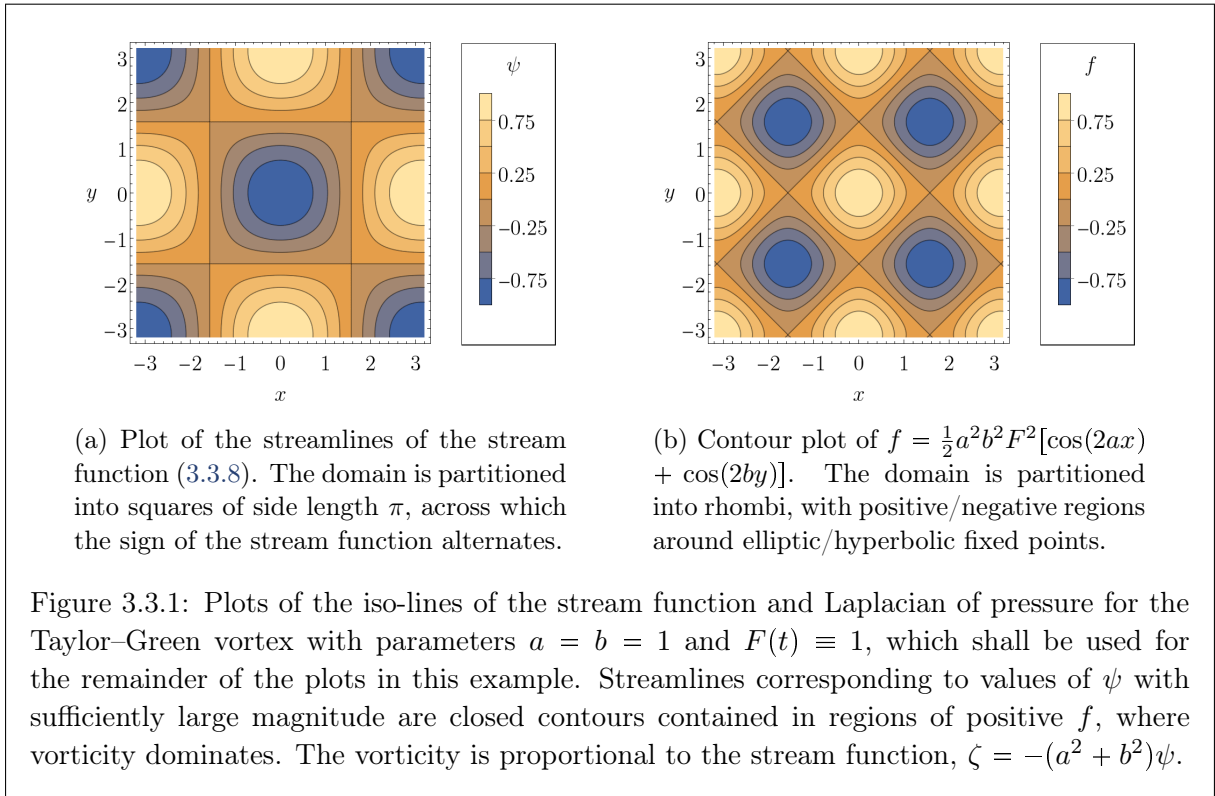
where

$$\tilde{g}^{-1} = \frac{1}{f} \begin{pmatrix} \partial_y^2 \psi & -\partial_x \partial_y \psi \\ -\partial_x \partial_y \psi & \partial_x^2 \psi \end{pmatrix} \quad (3.3.7b)$$

and

$$\tilde{R} = -\frac{1}{4} \tilde{g}^{ij} \tilde{g}^{kl} \tilde{g}^{mn} (\partial_i \partial_j \partial_m \psi \partial_k \partial_l \partial_n \psi - \partial_i \partial_k \partial_m \psi \partial_j \partial_l \partial_n \psi). \quad (3.3.7c)$$

As discussed following (3.2.18b), the curvature (3.3.7) involves derivatives of the vorticity and rate-of-strain. These gradients can become large in turbulent flows featuring fine-scale accumulation of vorticity along vortex filaments [20], which in turn could present a challenge when computing the Ricci curvature scalar (3.3.7) in numerical simulations. However, as the following example demonstrates, when the metric structure degenerates and/or the scalar curvature becomes singular (due to blow-up in the gradient of vorticity, for example) we can infer the presence of topological changes in the fluid flow.



3.3.2 The Taylor–Green Vortex

The Taylor–Green vortex [75] in two dimensions is described by the function

$$\psi(t, x, y) := -F(t) \cos(ax) \cos(by), \quad (3.3.8)$$

where F is a function of time t alone and $a, b \in \mathbb{R}$ are parameters. See Figure 3.3.1a. The diagnostic quantity (3.3.1) is given by

$$f = \frac{1}{2}a^2b^2F^2[\cos(2ax) + \cos(2by)], \quad (3.3.9)$$

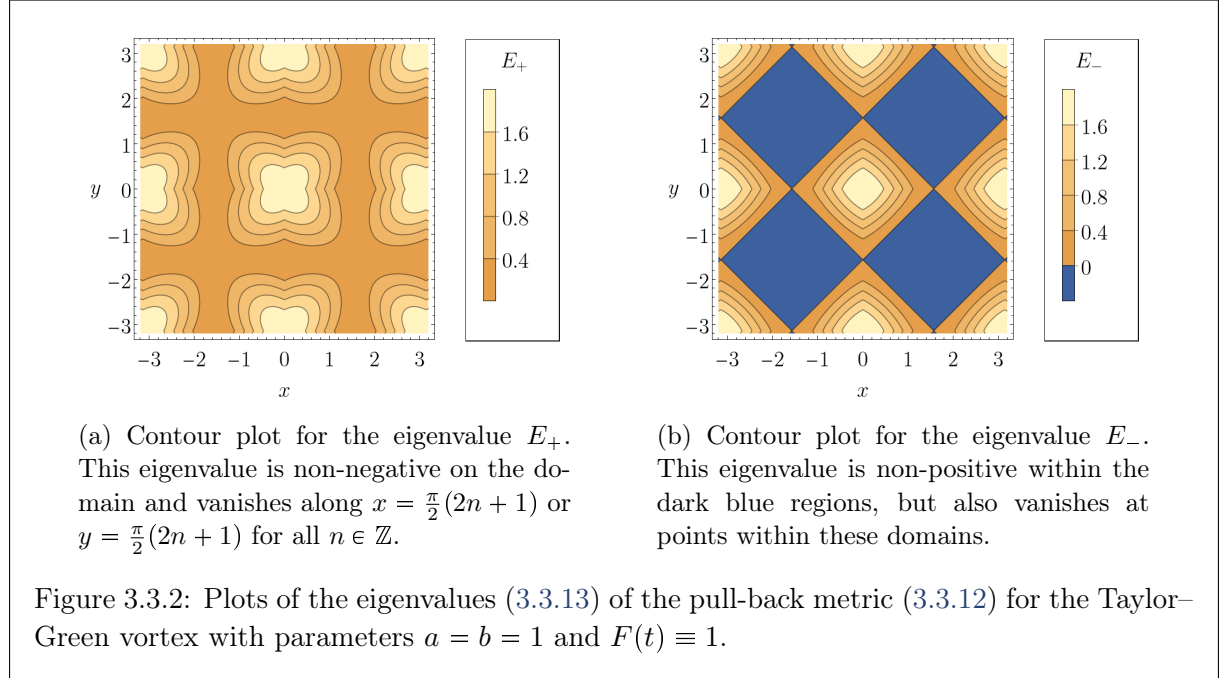
see Figure 3.3.1b. Therefore, the Lychagin–Rubtsov metric (3.3.2) has curvature scalar (3.3.3) given by

$$\hat{R} = \frac{8(a^2 + b^2)[1 + \cos(2ax) \cos(2by)]}{a^2b^2F^2[\cos(2ax) + \cos(2by)]^3}, \quad (3.3.10)$$

and the sign of \hat{R} coincides with that of f . See Figure 3.3.3a.

Consequently, when $\cos(2ax) + \cos(2by) > 0$, the metric is Riemannian with a positive curvature scalar and vorticity dominates. When $\cos(2ax) + \cos(2by) < 0$ the metric is Kleinian with negative curvature scalar and rate-of-strain dominates. Both the metric and curvature

scalar are singular when $abF = 0$ and along the lines $y = \frac{a}{b}x + \frac{\pi}{2b}(2n + 1)$ for all $n \in \mathbb{Z}$ (when $\cos(2ax) + \cos(2by) = 0$), corresponding to where $f = 0$.



Furthermore, the vorticity is given by

$$\zeta = (a^2 + b^2)F \cos(ax) \cos(by) \quad (3.3.11)$$

and the pull-back metric (3.3.4) becomes

$$g = \frac{(a^2 + b^2)F^2}{4} \begin{pmatrix} a^2[1 + \cos(2ax)][1 + \cos(2by)] & -ab \sin(2ax) \sin(2by) \\ -ab \sin(2ax) \sin(2by) & b^2[1 + \cos(2ax)][1 + \cos(2by)] \end{pmatrix}. \quad (3.3.12)$$

Its eigenvalues (3.3.5) are

$$E_{\pm} = \frac{F^2(a^2 + b^2)}{4} \left[2(a^2 + b^2) \cos^2(ax) \cos^2(by) \pm |\cos(ax) \cos(by)| \sqrt{\tilde{E}} \right] \quad (3.3.13a)$$

with

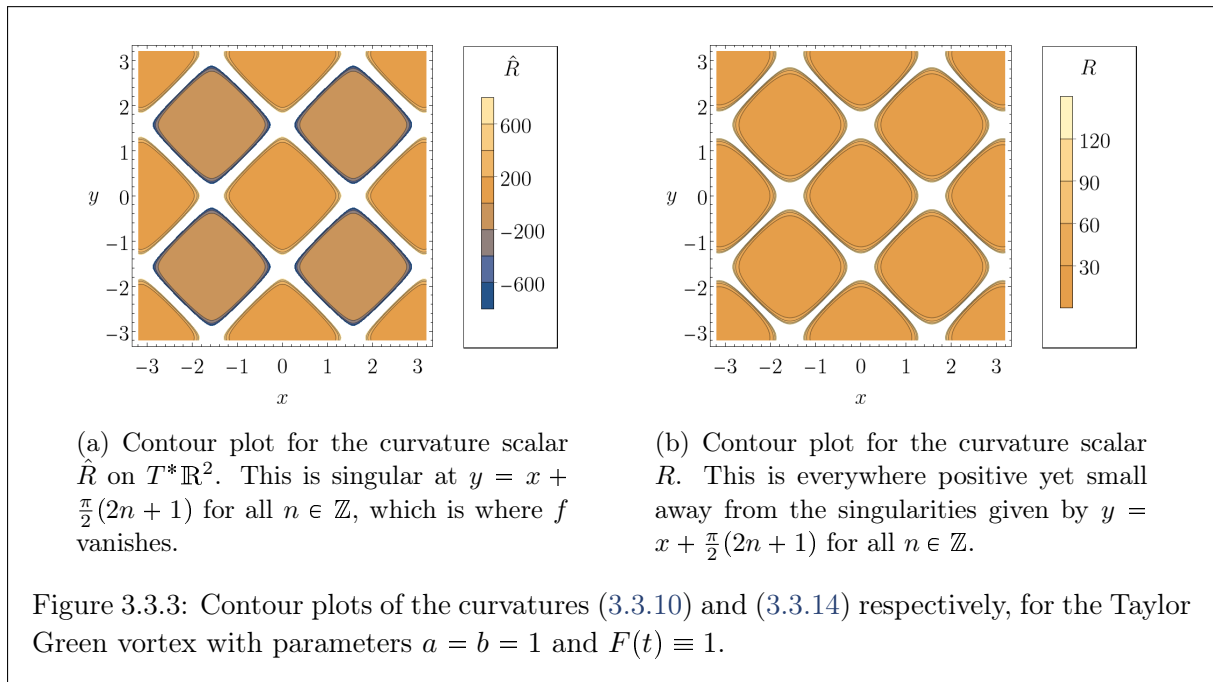
$$\tilde{E} := (a^4 - 6a^2b^2 + b^4)[\cos(2ax) + \cos(2by)] + (a^2 + b^2)^2[1 + \cos(2ax) \cos(2by)], \quad (3.3.13b)$$

as shown in Figure 3.3.2. The corresponding curvature scalars (3.3.7) are

$$R = \frac{8}{F^2(a^2 + b^2)[\cos(2ax) + \cos(2by)]^2} \quad \text{and} \quad \tilde{R} = 0. \quad (3.3.14)$$

Observe that E_+ is everywhere non-negative, so the signature of the metric (3.3.12) is determined by the sign of E_- . It is clear from Figure 3.3.2b that, when $\cos(2ax) + \cos(2by) \geq 0$, we have $E_- \geq 0$ and the metric g is Riemannian/Kleinian, with vorticity/strain dominating based on the sign of f , as given by (3.3.9). Also, $E_- = 0$ when $\cos(2ax) + \cos(2by) = 0$ and by (3.3.14), these contours are curvature singularities. Furthermore, the vorticity changes sign as the contours $x = \frac{\pi}{2a}(2n+1)$ or $y = \frac{\pi}{2b}(2n+1)$ are crossed, with the metric (3.3.12) degenerating along the contours but remaining Kleinian on both sides; this is consistent with an observation by [24], who states that the vorticity must have constant sign within simply connected domains on which $f := \det(\partial_i \partial_j \psi) > 0$, i.e. where the vorticity dominates and (by the discussion following (3.2.13)) where the pull-back metric is Riemannian. However, the curvature (3.3.14) is not singular along these contours, so they arise as a result of our coordinate choice — it is possible to find a coordinate change where the conditions $\zeta = 0$ and $f = 0$ coincide, for example.

Finally, when $F \equiv 0$ or $a = b = 0$, we have $\psi(x, y; t) = -F(t)$, the velocity components vanish and the flow is stationary. Similarly, when $a = 0$ or $b = 0$, the stream function depends on only one coordinate and the flow is one-dimensional. In each of these cases, $f \equiv 0$, so both the Lychagin–Rubtsov metric (3.2.8) and its pull-back (3.2.10) degenerate everywhere.



3.4 Chapter Summary

In Section 3.1 of this chapter, we reviewed some known properties of incompressible fluid flows in Euclidean domains and observed that the divergence-free condition induces a Poisson equa-

tion (3.1.1c) for the pressure, which holds regardless of the viscosity of the flow. In particular, for flows in the Euclidean plane, we saw that the dominance of vorticity and strain is determined by the sign of the Laplacian of the pressure, which the Weiss criterion relates to the equation type of (3.1.4), viewed as a Monge–Ampère equation for the stream function. Subsequently, we demonstrated how an analogous, covariant Poisson equation for the pressure (3.1.13c) could be derived for incompressible Navier–Stokes flows on an arbitrary Riemannian manifold and that, once again, the equation could be viewed as a Monge–Ampère equation for a stream function in two dimensions. We also noted in Remark 3.1.1 that solutions to the full Navier–Stokes equations are a subset of the solutions of the Poisson equation, so general properties of solutions to the latter are also properties of solutions to the former.

In Section 3.2, we demonstrated that the Poisson equation for the pressure (3.2.3) could be encoded by the Monge–Ampère structure (3.2.4b) in two dimensions. Furthermore, by utilising the results of [15] introduced in Section 2.2.1, an almost (para-)complex structure (3.2.6) and almost (para-)Hermitian form (3.2.7) can be associated with the Monge–Ampère structure and used to induce the Lychagin–Rubtsov metric (3.3.2) on the cotangent bundle. The Lychagin–Rubtsov theorem (Theorem 2.2.6) then informs us that the Poisson equation for the pressure is locally equivalent to the Laplace or wave equation precisely when its Pfaffian \hat{f} is constant, which can only occur if M is flat.

Considering classical solutions $\iota : L \hookrightarrow T^*M$ to the Monge–Ampère structure (3.2.4b), we observe the following relationship: when vorticity dominates over strain, $f = \iota^*\hat{f} = \frac{1}{2}\mathring{\Delta}_B p + \frac{\mathring{R}}{4}|\mathrm{d}\psi|^2$ is positive, the Monge–Ampère equation (3.2.3) is of elliptic type (as determined by the sign of the Pfaffian) and both the Lychagin–Rubtsov metric and its pull-back to L are Riemannian. Similar associations hold for the rate-of-strain and these relationships are summarised in Table 3.4.1. This correspondence is, in essence, an extension of the Weiss criterion to the general Riemannian setting, taking account of the additional term introduced in (3.2.3) by the curvature of the underlying manifold. We closed the section by showing that the curvature of the pull-back metric is given by the gradients of vorticity and rate-of-strain, and can be used to obtain topological information about the flow, via the Gauß–Bonnet theorem, for example.

Finally, in Section 3.3, a simplified geometric description in the Euclidean setting was presented. The Lychagin–Rubtsov metric, its pull-back to classical solutions, and their associated curvatures were shown to reflect the standard Weiss criterion in this case, as illustrated by the example of the Taylor–Green vortex.

In the next chapter, we shall show how, through a subtle modification to the Monge–Ampère structure (3.2.4b), this analysis can be extended to incompressible fluid flows in dimension higher than two, where (3.1.17) no longer takes the form of a Monge–Ampère equation. Furthermore, in

Dominance	Vorticity	Strain	None
f	0	< 0	$= 0$
$\alpha _L = 0$	Elliptic	Hyperbolic	Parabolic
\hat{J}^2	$-\mathbb{I}$	\mathbb{I}	Singular
\hat{g}	Riemannian	Kleinian	Degenerate
g	Riemannian	Kleinian	Degenerate

Table 3.4.1: Summary of the relationships between fluid dynamical, PDE theoretical, and geometric quantities.

three dimensions, we will again highlight how this might be used to obtain topological information regarding vortices, before considering the application of dimensional reductions to flows with symmetries.



4

(Higher) Monge–Ampère Geometry of Fluid Flows

Having discussed incompressible fluid flows on two-dimensional Riemannian manifolds, we now wish to consider fluid flows in three (or even higher) dimensions. In this case, the Poisson equation for the pressure (3.1.13c) no longer takes the form of a Monge–Ampère equation, since in general such flows are not described by stream functions, but rather by stream forms¹ (3.1.18).

However, subtly modifying the Monge–Ampère structure (3.2.4b) by means of choosing a different symplectic form, we can obtain a pair of differential forms on T^*M which not only encode the Poisson equation for the pressure (3.2.4a), but also the divergence-free constraint (3.1.13b). Indeed, in the previous chapter, incompressibility never arose as a consequence of the geometry, rather, we assumed that our flow was divergence-free a priori to induce the existence of a stream function. Furthermore, while this alternative view makes little difference to the geometric picture in two dimensions, it naturally generalises to higher dimensions and provides a mechanism for unifying the otherwise quite distinct geometric descriptions of two-dimensional and three-dimensional fluid flows described in [19].

4.1 A Bridge to Higher Dimensions

Before discussing fluid flows in higher dimensions, let us make precise the aforementioned subtle modification in two dimensions. In Section 3.2, we saw that the Monge–Ampère structure (3.2.4b) encodes incompressible fluid flows on a two-dimensional Riemannian manifold (M, \mathring{g}) . As we did then, let $\{x^i\}_{i=1}^2$ and $\{x^i, q_i\}_{i=1}^2$ be coordinates on M and T^*M respectively. However, this time,

¹In Section 4.3, we shall present some examples of three-dimensional flows with symmetry, which can be described in terms of the stream function for some associated two-dimensional fluid flow.

instead of using the canonical symplectic form (2.1.5) in the Monge–Ampère structure (3.2.4b), we propose taking the form

$$\varpi := \overset{\circ}{\nabla} q_i \wedge \star_{\dot{g}} dx^i, \quad (4.1.1)$$

that is, (3.2.7) without the pre-factor. It should be evident that ϖ is non-degenerate² and we shall demonstrate in Section 4.2 that it is closed, therefore it is indeed a symplectic form. It is easily seen that, when $\iota : L \hookrightarrow T^*M$ is a submanifold of T^*M given by

$$\iota : x^i \mapsto (x^i, q_i) := (x^i, v_i(x)), \quad (4.1.2)$$

with $v_i = v_i(x)$ the components of the velocity (co-)vector field, the condition $\iota^*\varpi = 0$ is equivalent to requiring the divergence-free constraint (3.1.13b). Thus, our fluid flow is again described by a Lagrangian submanifold L of T^*M , this time with respect to ϖ in place of ω .

Moreover, using that $\star_{\dot{g}}(dx^i \wedge dx^j) = \sqrt{\det(\dot{g})} \varepsilon^{ij}$ in two dimensions and denoting the volume form on M by

$$\text{vol}_M := \frac{\sqrt{\det(\dot{g})}}{2} \varepsilon_{ij} dx^i \wedge dx^j, \quad (4.1.3)$$

we may rewrite the Monge–Ampère form α defined in (3.2.4b) as

$$\alpha = \frac{1}{2} \overset{\circ}{\nabla} q_i \wedge \overset{\circ}{\nabla} q_j \wedge \star_{\dot{g}}(dx^i \wedge dx^j) - \hat{f} \text{vol}_M. \quad (4.1.4)$$

As discussed around (3.2.5), α is non-degenerate precisely when $\hat{f} \neq 0$ and as for ϖ , we shall demonstrate in Section 4.2 that α is closed in arbitrary dimension. Since we also have $\alpha \wedge \varpi = 0$, the pair (ϖ, α) does indeed form a Monge–Ampère structure. Indeed, the structure given by (4.1.1) and (4.1.4) is precisely the covariantisation of the structure studied in [18, 19]. Additionally, the requirement that $\iota^*\alpha = 0$ under (4.1.2) is equivalent to the Poisson equation for the pressure (3.1.13c), written in terms of the velocity components, provided that we simultaneously demand $\iota^*\varpi = 0$. As for the structure (ω, α) , the Pfaffian of α with respect to ϖ is given by the function \hat{f} defined in (3.2.4a).

Recall from (3.1.18) that, for incompressible fluid flows in two dimensions, it is possible to find stream functions $\psi \in \mathcal{C}^\infty(U)$ on open contractible subsets $U \subseteq M$, such that $v = \star_{\dot{g}} d\psi$. Upon describing the velocity in this manner, the constraint $\iota^*\varpi = 0$ becomes trivial and $\iota^*\alpha = 0$ yields the Poisson equation for the pressure in the form (3.2.3), that is, as a Monge–Ampère equation for the stream function on U . In particular, submanifolds of the form

$$\iota : x^i \mapsto (x^i, q_i) := (x^i, (\star_{\dot{g}} d\psi)_i), \quad (4.1.5)$$

²Indeed, the form \mathcal{K} from (3.2.7) is non-degenerate when $\hat{f} \neq 0$ and equal to $-\varpi$ when $\hat{f} = 1$.

additionally satisfying $\iota^*\alpha = 0$, correspond to smooth classical solutions of the Monge–Ampère equation (3.2.3), as discussed in Remark 2.1.11. Note that the form α is unchanged under the symplectomorphism which takes ϖ to ω .

We may now follow our discussion in Section 3.2 and define an endomorphism $\hat{\mathcal{J}} : \Gamma(TT^*M) \rightarrow \Gamma(TT^*M)$ of vector fields on T^*M by

$$\frac{\alpha}{\sqrt{|\hat{f}|}} =: \hat{\mathcal{J}} \lrcorner \varpi, \quad (4.1.6)$$

under the assumption that \hat{f} does not vanish. As before, $\hat{\mathcal{J}}$ is an almost complex structure when $\hat{f} > 0$, an almost para-complex structure when $\hat{f} < 0$, and integrable if and only if \hat{f} is constant. In light of our future applications, note that (4.1.6) can be rewritten as

$$\hat{\mathcal{J}}X = \frac{1}{\sqrt{|\hat{f}|}} \varepsilon \lrcorner (\varpi \wedge X \lrcorner \alpha), \quad (4.1.7)$$

where ε is the poly-vector dual to the Liouville volume form $\frac{1}{2}\omega \wedge \omega = \frac{1}{2}\varpi \wedge \varpi$. Again, Proposition 2.2.7 guarantees the existence of a differential two-form $\hat{\mathcal{K}}$ of type (1, 1) with respect to $\hat{\mathcal{J}}$ such that $\hat{\mathcal{K}} \wedge \varpi = 0$, $\hat{\mathcal{K}} \wedge (\hat{\mathcal{J}} \lrcorner \varpi) = 0$, and $\hat{\mathcal{K}} \wedge \hat{\mathcal{K}} \neq 0$. In particular, we choose

$$\hat{\mathcal{K}} := \sqrt{|\hat{f}|} \overset{\circ}{\nabla} q_i \wedge dx^i, \quad (4.1.8)$$

that is, the canonical symplectic structure (3.2.4b) scaled with respect to the Pfaffian \hat{f} of α . In essence, we have interchanged the roles of ω and ϖ in the construction. Importantly, the compatibility condition $\hat{\mathcal{K}}(\hat{\mathcal{J}}X, Y) = -\hat{\mathcal{K}}(X, \hat{\mathcal{J}}Y)$ again yields the metric (3.2.8). Furthermore, pulling the metric (3.2.8) along (4.1.5) also yields the pull-back metric (3.2.10) on Lagrangian submanifolds which correspond to classical solutions.

In conclusion, the Monge–Ampère structure (ϖ, α) , with ϖ as defined by (4.1.1) and α written as (4.1.4) represents an alternative means to describe two-dimensional, incompressible fluids. Whilst the Monge–Ampère structure (3.2.4b) yields explicitly a description of the fluid flow in terms of the stream function, the alternative structure yields two advantages: firstly, that the divergence-free constraint (3.1.13b) is now encoded by ϖ in the geometry and secondly that the differential forms (4.1.1) and (4.1.4) naturally generalise to higher dimensions. However, in $m > 2$ dimensions, we leave the realm of symplectic geometry, as we shall see shortly.

4.2 Geometric Properties of Higher-Dimensional Incompressible Fluid Flows

Having introduced our modified Monge–Ampère structure in two dimensions, we can now make precise the description of higher-dimensional incompressible fluid flows. We begin by providing a formulation of the Poisson equation for the pressure (3.1.13c) and divergence-free constraint (3.1.13b) in terms of differential forms which vanish on submanifolds of the cotangent bundle of an m -dimensional Riemannian manifold, as we did in two dimensions.

In light of this, let M be a smooth, m -dimensional, Riemannian manifold with metric \mathring{g} and consider the pair of differential m -forms on T^*M , given by

$$\begin{aligned}\varpi &:= \mathring{\nabla}q_i \wedge \star_{\mathring{g}}dx^i, \\ \alpha &:= \frac{1}{2}\mathring{\nabla}q_i \wedge \mathring{\nabla}q_j \wedge \star_{\mathring{g}}(dx^i \wedge dx^j) - \hat{f}\text{vol}_M,\end{aligned}\tag{4.2.1a}$$

where indices now run from 1 to m , the volume form on M is given by

$$\text{vol}_M := \frac{\sqrt{\det(\mathring{g})}}{3!}\varepsilon_{ijk}dx^i \wedge dx^j \wedge dx^k,\tag{4.2.1b}$$

and as in two dimensions, we set

$$\hat{f} := \frac{1}{2}(\mathring{\Delta}_{BP} + \mathring{R}^{ij}q_iq_j).\tag{4.2.1c}$$

In three dimensions, these choices can be understood as a covariantisation of the differential forms previously presented by [18, 19]. Importantly, the formulation (4.2.1) makes it clear that this construction works in any dimension $m > 1$. Indeed, we simply need to take the appropriate volume form in α and the function \hat{f} is the same in any dimension. In this sense, (ϖ, α) is a kind of “higher Monge–Ampère structure” associated to the equations (3.1.13).

It is straightforward to verify that ϖ is non-degenerate for $m \geq 2$, while α is non-degenerate for $m > 2$ (and $m = 2$ if $\hat{f} \neq 0$). Furthermore, by using

$$d(\mathring{\nabla}q_i) = \frac{1}{2}dx^l \wedge dx^k \mathring{R}_{kli}{}^j q_j + dx^j \mathring{\Gamma}_{ji}{}^k \wedge \mathring{\nabla}q_k\tag{4.2.2a}$$

and

$$d\star_{\mathring{g}}dx^i = -\mathring{g}^{jk}\mathring{\Gamma}_{jk}{}^i \text{vol}_M,\tag{4.2.2b}$$

we may verify that ϖ is closed. It follows that ϖ is an $(m - 1)$ -plectic structure on T^*M for $m \geq 2$. Similarly, it follows from (4.2.2a),

$$d\star_{\mathring{g}}(dx^i \wedge dx^j) = 2\mathring{g}^{kl}\mathring{\Gamma}_{kl}{}^{[i}\star_{\mathring{g}}dx^{j]} + 2\mathring{g}^{k[i}\mathring{\Gamma}_{kl}{}^{j]}\star_{\mathring{g}}dx^l\tag{4.2.3a}$$

and

$$dx^k \wedge \star_{\hat{g}}(dx^i \wedge dx^j) = -2\hat{g}^{k[i} \star_{\hat{g}} dx^j], \quad (4.2.3b)$$

that α is closed, hence is also an $(m-1)$ -plectic structure on T^*M for $m > 2$.

Since neither ϖ nor α is a symplectic form for $m > 2$, the pair (ϖ, α) , does not define a Monge–Ampère structure in the sense of [Definition 2.1.9](#). However, we may still consider submanifolds $\iota : L \hookrightarrow T^*M$ on which $\iota^*\varpi = 0$ and $\iota^*\alpha = 0$, in the spirit of generalised solutions. Indeed, taking

$$\iota : x^i \mapsto (x^i, q_i) := (x^i, v_i(x)), \quad (4.2.4)$$

where $v_i = v_i(x)$ are interpreted as the components of the velocity (co-)vector field, then $\iota^*\varpi = 0$ corresponds to the divergence-free constraint [\(3.1.13b\)](#), while $\iota^*\alpha = 0$ corresponds to the Poisson equation for the pressure in terms of the velocity components [\(3.1.13c\)](#). As such, we shall consider [\(4.2.4\)](#) in lieu of classical solutions.

For the canonical symplectic structure

$$\omega = \overset{\circ}{\nabla} q_i \wedge dx^i, \quad (4.2.5)$$

on T^*M , both $\varpi \wedge \omega$ and $\alpha \wedge \omega = 0$, so ϖ and α are both Monge–Ampère forms for ω . In contrast, $\alpha \wedge \varpi = 0$ if and only if $m \neq 3$. However, we do not consider the pair of coupled Monge–Ampère structures (ω, ϖ) and (ω, α) , since this would require imposing $\iota^*\omega = 0$ and

$$\iota^*\omega = dv = \frac{1}{2}\zeta_{ij} dx^j \wedge dx^i \quad (4.2.6)$$

along [\(4.2.4\)](#), which implies that the vorticity two-form vanishes everywhere. As we wish to use our geometry to study vorticity, this would not be very productive!

Following our construction in two dimensions, we can again equip T^*M with the metric

$$\hat{g} = \frac{1}{2}\hat{f}\hat{g}_{ij}dx^i \odot dx^j + \frac{1}{2}\hat{g}^{ij}\overset{\circ}{\nabla}q_i \odot \overset{\circ}{\nabla}q_j, \quad (4.2.7)$$

whose signature is determined by the sign of the quantity \hat{f} . Let $\hat{\Delta}_B$ be the Beltrami Laplacian

for \hat{g} . The Ricci curvature scalar for this metric is derived in [Appendix C.2.3](#) and takes the form

$$\begin{aligned} \hat{R} = & \frac{1}{\hat{f}} \mathring{R} - \frac{1}{4\hat{f}^2} \mathring{R}_{ijk}{}^l \mathring{R}^{ijkm} q_k q_m - (m-1) \hat{\Delta}_B \log(|\hat{f}|) - \mathring{g}_{ij} \frac{\partial^2}{\partial q_i \partial q_j} \log(|\hat{f}|) \\ & + \frac{1}{4\hat{f}} (m-1)(m-2) \mathring{g}^{ij} \left(\frac{\partial}{\partial x^i} + \mathring{\Gamma}_{ik}{}^l q_l \frac{\partial}{\partial q_k} \right) \log(|\hat{f}|) \left(\frac{\partial}{\partial x^j} + \mathring{\Gamma}_{jm}{}^n q_n \frac{\partial}{\partial q_m} \right) \log(|\hat{f}|) \\ & + \frac{1}{4} m(m-3) \mathring{g}_{ij} \frac{\partial}{\partial q_i} \log(|\hat{f}|) \frac{\partial}{\partial q_j} \log(|\hat{f}|), \end{aligned} \quad (4.2.8)$$

in any dimension. The occurrence of the term $\hat{\Delta}_B \log(|\hat{f}|)$ again suggests that the accumulation of \hat{f} determines the sign of the Ricci scalar curvature near fixed points, as it does in the two-dimensional case.

Finally, let us present a formula for the pull-back of the metric (4.2.7) via (4.2.4) in arbitrary dimension, in terms of fluid dynamical quantities. Pulling back \hat{f} via (4.2.4) and applying the pressure equation (3.1.13c) (in the form (3.1.17)) yields

$$f = \iota^* \hat{f} = -\frac{1}{2} A_j^i A_i^j, \quad (4.2.9)$$

where $A_i^j = A_{ik} \mathring{g}^{kj}$ and A_{ij} is the velocity gradient tensor defined in (3.1.15). Additionally, observe that the pull-back of $\mathring{\nabla} q_i$ via (4.2.4) is given by

$$\iota^*(\mathring{\nabla} q_i) = \mathring{\nabla}_j v_i dx^j = A_{ij} dx^j. \quad (4.2.10)$$

It follows that the pull-back of (4.2.7) is

$$g = \frac{1}{2} g_{ij} dx^i \odot dx^j \quad \text{with} \quad g_{ij} := \mathring{g}^{kl} A_{ki} A_{lj} - \frac{1}{2} A_k^l A_l^k \mathring{g}_{ij}. \quad (4.2.11)$$

As in the two-dimensional case, the diagnostic nature of the quantity f in (3.1.17) results in the pull-back metric being a quadratic function of the velocity gradient tensor, with curvature generated by gradients of vorticity and rate-of-strain.

In the case of two-dimensional incompressible flows, it is possible to recover the metric (3.2.8) from (4.2.11) by recalling that, on open contractible neighbourhoods $U \subseteq M$, the divergence free constraint can be solved as in (3.1.18) to obtain

$$v^i = \sqrt{\det(\mathring{g})} \varepsilon^{ij} \partial_j \psi, \quad (4.2.12)$$

while the velocity gradient tensor is given by $A_{ij} = \mathring{\nabla}_j v_i$. However, in contrast to the two-dimensional case, where we could directly relate the signature of (3.2.8) to the sign of (4.2.9),

the correspondence in higher dimensions is not so clear. Indeed, [1, Section 3.4] provide an example where the pull-back metric (4.2.11) is not singular along the curve given by $f = 0$. In particular, the final row of Table 3.4.1 does not hold for dimension higher than 2 in general (while the remainder of the table does). In Section 4.4 we shall demonstrate that the sign of (4.2.9) also fails to directly determine the signature of (3.2.8) for two-dimensional compressible fluid flows generated by symmetry reduction, although it is not yet fully understood why this is the case.

4.2.1 Almost (Para-)Complex Structures For Three-Dimensional Fluid Flows

We now wish to know the relationship between the metric (4.2.7), which is of Lychagin–Rubtsov type in dimension two, and the classical Lychagin–Rubtsov metric (2.2.16) in three dimension, which has been studied in the context of incompressible fluid flows on Euclidean background in [19], for example. In particular, we wish to know whether (4.2.7) is a Lychagin–Rubtsov metric in dimension greater than two, i.e. is the metric almost (para-)Hermitian for some almost (para-)complex structure arising from our “higher Monge–Ampère structure.”

Let us specialise to the case $m = 3$ for concreteness. Unlike for a standard Monge–Ampère structure in three dimensions, ϖ and α are of the same order, so an identity of the form

$$\alpha =: J_\alpha \lrcorner \varpi, \quad (4.2.13)$$

as in (2.2.13), makes sense at a glance. However, as we shall see in Chapter 5 when we consider the classification of such higher structures, not all pairs of differential three-forms (ϖ, α) define an endomorphism J in this way. In particular, the structure (4.2.1) does not do so.

In light of this, we will again use the Hitchin endomorphism (2.2.20) to construct our almost (para-)complex structure. A priori we have two choices, as we could construct the endomorphism for either ϖ or α , however it turns out that $\text{HPf}(\varpi) = 0$ so (2.2.21) is ill-defined for ϖ . Consequently, we shall compute the structure associated with α ; this is convenient since α has explicit dependence on our diagnostic quantity \hat{f} . In order to do so explicitly, fix $\text{vol} = \frac{1}{3!}\omega^3$ to be the Liouville volume form with respect to the standard symplectic form ω on T^*M and let ε denote the poly-vector field dual to the Liouville volume form, that is, $\varepsilon \lrcorner \frac{1}{3!}\omega^3 = 1$. Then, under the assumption that \hat{f} does not vanish, we may use (2.2.21) to associate with the differential three-form α defined in (4.2.1) the endomorphism

$$\hat{\mathcal{J}}X := -\frac{1}{2\sqrt{|\hat{f}|}} \varepsilon \lrcorner (\alpha \wedge X \lrcorner \alpha) \quad \text{for all } X \in T(T^*M). \quad (4.2.14)$$

In particular, $\text{HPf}(\alpha) = \text{sgn}(\hat{f})$, from which it follows that $\hat{\mathcal{J}}$ is an almost complex structure on

T^*M when $\hat{f} > 0$ and an almost para-complex structure when $\hat{f} < 0$. Since α is closed and can also be taken as the imaginary part of a holomorphic top form with respect to (4.2.14), this choice of endomorphism, in fact, defines what is known as a *nearly (para-) Calabi–Yau structure* [76,77].

Furthermore, the differential two-form $\hat{\mathcal{K}}$ defined in (4.1.8), now with i running from one to three, is a $(1,1)$ -form for (4.2.14) and satisfies $\hat{\mathcal{K}}(\hat{\mathcal{J}}X, Y) = -\hat{\mathcal{K}}(X, \hat{\mathcal{J}}Y)$ for all $X, Y \in \Gamma(TT^*M)$. In fact, the almost (para-)Hermitian metric associated with $\hat{\mathcal{K}}$ via $\hat{g}(X, Y) := \hat{\mathcal{K}}(X, \hat{\mathcal{J}}Y)$ is precisely the metric (4.2.7) in three dimensions. In particular, it is a conformal scaling of the classical Lychagin–Rubtsov metric (2.2.23) for α , with conformal factor given by $\sqrt{|\hat{f}|}$ from (4.1.8), and is essentially covariantisation of the metric studied in [19, Section 4] in the context of three-dimensional incompressible flows on a Euclidean background. In contrast to earlier works, where the metrics in two and three dimensions were defined independently, our metric (4.2.7) is a direct generalisation of the metric (3.2.8) to dimension greater than two and relates the two- and three-dimensional descriptions. Fundamentally, this presentation provides a unified, geometric perspective from which to study incompressible fluid flows in arbitrary dimension.

Remark 4.2.1 (Rescaled Sasaki Metrics)

We observe that the metric (4.2.7) on T^*M has the same overall form as the *rescaled Sasaki metrics*, see e.g. [78, Definition 3.1]. The main difference here is that our rescaling function \hat{f} is a function on T^*M rather than on M . This results in our metric (4.2.7) being able to change type across T^*M . In [78, Section 3], the Ricci scalar curvature of the rescaled Sasaki metric is computed and is consistent with our findings in Appendix C.2. Furthermore, earlier work [79] has focused on constructing almost para-Nordenian manifolds in the case $\hat{f} > 0$, preferentially selecting an endomorphism J which satisfies the Nordenian property³ $\hat{g}(\hat{\mathcal{J}}X, Y) = \hat{g}(X, \hat{\mathcal{J}}Y)$ and is almost para-complex instead of almost complex. Further investigation into the relationship between the geometric properties of incompressible fluid flows, rescaled Sasaki metrics, and para-Nordenian manifolds is a promising subject for future research.

△

4.2.2 Topology of Three-Dimensional Incompressible Fluid Flows

Now let us turn to the problem of describing the topological artefacts associated with the presence of vortex structures in a fluid flow. In two dimensions, we utilised the local Gauß–Bonnet theorem (3.2.19) in order to relate the geometry of fluid flows within compact regions $U \subseteq M$, as described by the curvature scalar (3.2.18), to a topological invariant, namely the Euler

³Note this property differs by a sign from the compatibility condition our almost para-complex structures satisfy

characteristic. In three dimensions, where the Euler characteristic of a compact manifold is always zero, it quickly becomes apparent that this approach is not suitable and we require an alternative topological quantity.

Let M be a three-dimensional Riemannian manifold. Recall from (4.2.5) that the pull-back of the canonical symplectic form along (4.2.4) is given by $\iota^*\omega = \frac{1}{2}\zeta_{ij}dx^i \wedge dx^j$. Similarly, the associated tautological one-form $\theta := q_i dx^i$, which satisfies $\omega = d\theta$, has pull-back $\iota^*\theta = v = v_i dx^i$ along (4.2.4). It follows from the distributive property of the pull-back over the wedge product that

$$\iota^*(\theta \wedge d\theta) = v_i \zeta^i \text{vol}_{M^3}, \quad \text{with} \quad \zeta^i := \sqrt{\det(\hat{g}_3)} \varepsilon^{ijk} \zeta_{jk} \quad (4.2.15)$$

the vorticity in three dimensions, given in terms of the vorticity two-form (3.1.14). Integrals of quantities of the form (4.2.15), over a compact region $U \subseteq M$, are referred to as *helicity* [30, 80]. Hence, in our context, $v_i \zeta^i$ may be referred to as the *helicity density* or *helicity per volume*.

Consider an inviscid, incompressible fluid described by the Euler equations (3.1.13) ($\nu = 0$) on a compact region $U \subseteq M$. Suppose also that U is a volume contained inside some closed orientable surface, which is moving with the fluid and has (continuous, outward) unit normal n with components denoted n_i . It is shown in [30, Equation 6] that, provided the distribution of vorticity is local and continuous, and $n_i \zeta^i = 0$, then the integral of (4.2.15) over $U \subseteq M$ is an invariant of the Euler equations and the vorticity two-form in U is conserved,⁴ while [81] demonstrate that helicity is an isotopy invariant of the volume. Furthermore, for discrete vortex filaments, [31, 82] associate the helicity with the topological invariants given by the *Gauß linking number* and *Călugăreanu invariant* [83, 84]. Perhaps more significantly, a recent work [85] has managed to demonstrate that, in ideal conditions, helicity-type quantities can be reinterpreted as Abelian Chern–Simons actions. Hence, the helicity density (4.2.15) of a configuration of discrete vortex filaments can be related to the Jones polynomial of the knots they form.

Consequently, in addition to the interpretation of the pull-backs of (4.2.1) under (4.2.4) as the divergence-free constraint and the pressure equation, in three dimensions, we also have that the corresponding pull-back of the canonical symplectic form encodes the helicity. Furthermore, earlier work relating helicity to various topological invariants demonstrates that our geometric construction may be used to investigate the topological properties of fluid flows in three dimensions (as we did in two dimensions).

⁴In the context of magneto-hydrodynamics, the analogous result was presented in [80, Equation 3].

4.3 Reduction of Three-Dimensional Flows with Symmetry

Earlier works [18, 19] considered a class of solutions to the three-dimensional incompressible Navier–Stokes equations on Euclidean space which take the form [86]

$$(\dot{x}^1, \dot{x}^2, \dot{x}^3) := (v_1(x^1, x^2, t), v_2(x^1, x^2, t), x^3 \gamma(x^1, x^2, t) + W(x^1, x^2, t)), \quad (4.3.1)$$

for some functions γ and W , where the superposed dot refers to the derivative with respect to the time parameter t . Such flows are referred to as ‘two-and-a-half-dimensional’ flows and are a generalisation of what is known as ‘columnar flow’ [87], where only the vertical component of the velocity v_3 depends on the vertical coordinate x^3 and does so linearly. One example of such a two-and-a-half-dimensional flow is Burgers’ vortex [88], where $W \equiv 0$ and $\gamma = \gamma(t)$. Burgers’ vortex was treated by [19, Section 5] in the following manner: assume that \mathbb{R} acts as a Lie group on the phase space $T^*\mathbb{R}^3$ via translation of x^3 and q_3 , with the infinitesimal generator of the action given by the vector field $\partial_{x^3} + \gamma \partial_{q_3}$. Given that such an action preserves the Monge–Ampère structure (4.2.1) and Burgers’ vortex is a classical solution of the Poisson equation for the pressure (3.1.13c) in the sense of (4.2.4), then there exists a two-dimensional Monge–Ampère equation with a classical solution which provides an equivalent description of the flow. By performing this reduction via our geometric formulation, the resulting Monge–Ampère equation can be studied in the sense of Section 3.2.

In this section, we shall consider solutions of the form (4.3.1) with $\gamma \equiv 0$ and $W = W(x^1, x^2, t)$ in more detail, as well as providing a generalisation of the technique laid out by [19] to fluid flows on an arbitrary Riemannian manifold. That is, we consider the case when the three-dimensional background manifold M exhibits some symmetry in one particular coordinate x^3 , along which we reduce. As a consequence, we find that the three-dimensional, covariant, incompressible Navier–Stokes equations on spaces with such symmetry reduce to equations for some (potentially compressible) two-dimensional flow. While we shall drop the parameter time t for brevity moving forward, all functions and constants should be interpreted as being ‘at fixed t ,’ so are in general allowed to vary in time.

4.3.1 General Setup

Let us begin by reviewing some fundamental definitions. We base our formulation here on the work of [28, 34] in the multisymplectic setting, while more information specific to the symplectic setting can be found in [41, Chapters 7 and 8]

Definition 4.3.1 (k -Plectic Vector Fields)

Let (N, ϖ) be a k -plectic manifold. A vector field $X \in \Gamma(TN)$ is called a k -plectic vector field for ϖ if the flow along X preserves ϖ , i.e. ⁵

$$\mathcal{L}_X \varpi = d(X \lrcorner \varpi) = 0. \quad (4.3.2)$$

If $X \lrcorner \varpi$ is not only closed, but exact, with

$$X \lrcorner \varpi = dH_X \quad (4.3.3)$$

for some $H_X \in \Omega^{k-1}(N)$, then we call X a Hamiltonian vector field with Hamiltonian H_X .

Let \mathbf{G} be a connected Lie group with Lie algebra \mathfrak{g} , whose dual is denoted \mathfrak{g}^* . For each $\xi \in \mathfrak{g}$, let X_ξ denote the corresponding vector field given by

$$X_\xi(x) = \left(\frac{d}{dt} \exp(t\xi) \cdot x \right) \Big|_{t=0}, \quad (4.3.4)$$

where $\exp : \mathfrak{g} \rightarrow \mathbf{G}$ is the exponential map, \cdot is the action of \mathbf{G} on N , and $x \in N$. We call the action of \mathbf{G} on N k -plectic (resp. Hamiltonian) if, for all $\xi \in \mathfrak{g}$, the corresponding vector field X_ξ is a k -plectic (resp. Hamiltonian) vector field with respect to ϖ .

Definition 4.3.2 (Moment Map)

Let (N, ϖ) be a k -plectic manifold and let \mathbf{G} be a connected Lie group whose action on N is Hamiltonian. The moment map corresponding to this action is a map $\mu : N \rightarrow \bigwedge^{k-1} T^*N \otimes \mathfrak{g}^*$ such that

$$d(\mu(\cdot)(\xi)) = X_\xi \lrcorner \varpi, \quad (4.3.5)$$

for all $\xi \in \mathfrak{g}$. Since the action of \mathbf{G} is Hamiltonian on (N, ϖ) , we have $d(\mu(\cdot)(\xi)) = dH_{X_\xi}$ by (4.3.3). In particular, $\mu(\cdot)(\xi) = H_{X_\xi}$, up to a shift by an exact $(k-1)$ -form.

We now wish to consider the class of incompressible fluid flows on a three-dimensional, background manifold M^3 , for which the velocity components have the form (4.3.1) with $W \equiv 0$ and show that these flows admit one-dimensional symplectic/ k -plectic reductions. In order to facilitate this, consider a three-dimensional manifold M^3 given by a warped-product of a two-dimensional manifold M^2 with Riemannian metric \dot{g}_2 and a one-dimensional manifold N with local coordinates x^3 . Then, there exists a metric on M^3 of the form

$$\dot{g}_3 = \dot{g}_2 + e^{2\varphi} dx^3 \otimes dx^3, \quad (4.3.6)$$

⁵The first equality here follows from the closure of ϖ and Cartan's homotopy formula [89, Theorem 20.10].

where we refer to $\varphi \in \mathcal{C}^\infty(M^2)$ as the *warping factor*. For the remainder of this chapter, we consider lower case indices $i, j, \dots = 1, 2$ so that, for example,

$$\mathring{g}_2 = \frac{1}{2} \mathring{g}_{ij} dx^i \odot dx^j. \quad (4.3.7)$$

Hence, the only non-vanishing Christoffel symbols for \mathring{g}_3 are $\mathring{\Gamma}_{ij}^k$, which are precisely the Christoffel symbols for \mathring{g}_2 , and

$$\mathring{\Gamma}_{33}^i = e^{2\varphi} \mathring{g}^{ij} \partial_j \varphi \quad \text{and} \quad \mathring{\Gamma}_{i3}^3 = \partial_i \varphi. \quad (4.3.8)$$

Next, for $m = 2, 3$, consider the differential forms (4.2.1) on T^*M^m and denote these by ϖ_m and α_m respectively. We shall use similar notation for other quantities occurring in both M^3 and M^2 . Then, under the assumption that the pressure is independent of the x^3 coordinate, that is, $p \in \mathcal{C}^\infty(M^2)$, some algebra reveals

$$\begin{aligned} \varpi_3 &= e^\varphi \varpi_2 \wedge dx^3 + e^{-\varphi} \text{vol}_{M^2} \wedge \mathring{\nabla} q_3, \\ \alpha_3 &= e^\varphi (\alpha_2 - \hat{h}_+ \text{vol}_{M^2}) \wedge dx^3 + e^{-\varphi} (\varpi_2 - q_3 dx^3 \wedge \star_{\mathring{g}_2} d\varphi) \wedge \mathring{\nabla} q_3, \end{aligned} \quad (4.3.9a)$$

with

$$\hat{h}_\pm := \frac{1}{2} [\mathring{\nabla}^i \varphi \partial_i p - (\mathring{\nabla}^i \mathring{\nabla}^j \varphi \pm \mathring{\nabla}^i \varphi \mathring{\nabla}^j \varphi) q_i q_j - e^{-2\varphi} (\mathring{\Delta}_B \varphi \pm \mathring{\nabla}^i \varphi \partial_i \varphi) q_3^2], \quad (4.3.9b)$$

where all differential operators in \hat{h}_\pm are taken with respect to the metric \mathring{g}_2 . Unless indicated otherwise, this shall hold true for all formulae in the remainder of this chapter. Furthermore,

$$\begin{aligned} \varpi'_2 &:= \partial_{x^3} \lrcorner \varpi_3 \\ &= e^\varphi (\varpi_2 + q_i \mathring{\nabla}^i \varphi \text{vol}_{M^2}), \\ \alpha'_2 &:= \partial_{x^3} \lrcorner \alpha_3 \\ &= e^\varphi [\alpha_2 - (\hat{h}_+ + e^{-2\varphi} \mathring{\nabla}^i \varphi \partial_i \varphi q_3^2) \text{vol}_{M^2} + q_i \mathring{\nabla}^i \varphi \varpi_2] + e^{-\varphi} q_3 dq_3 \wedge \star_{\mathring{g}_2} d\varphi. \end{aligned} \quad (4.3.10)$$

A short calculation shows that both ϖ'_2 and α'_2 are closed. In fact, using (4.2.2b), we also find

$$\varpi'_2 = d(\star_{\mathring{g}_2} e^\varphi q_i dx^i). \quad (4.3.11)$$

Consequently, ∂_{x^3} is a 2-plectic vector field with respect to both ϖ_3 and α_3 , as well as being a Hamiltonian vector field for ϖ_3 with Hamiltonian ϖ'_2 , hence the geometric flow along ∂_{x^3} preserves our Monge–Ampère structure (4.2.1). Furthermore, if ∂_{x^3} is the infinitesimal generator for some action of a one-dimensional, connected Lie group G on T^*M^3 , then said action is

2-plectic for α_3 and Hamiltonian for ϖ_3 . Crucially, such a Lie group must be isomorphic to either \mathbb{R} or \mathbb{S}^1 . Choosing different local coordinates x^3 on N corresponds to a different choice of metric (4.3.6) and hence results a different infinitesimal generator.

In summary, given an incompressible fluid flow on a three dimensional manifold M^3 with metric (4.3.6), pressure $p \in \mathcal{C}^\infty(M^2)$, and a one-dimensional, connected Lie group \mathbf{G} whose action on T^*M^3 is generated by ∂_{x^3} , then the Monge–Ampère structure (4.2.1) is preserved by the flow along ∂_{x^3} . This suggests that such flows admit a dimensional reduction along x^3 . In the subsequent two subsections, we present two different approaches to this reduction, using symplectic and 2-plectic geometry in turn. This is where the prematurely introduced moment map (see Definition 4.3.2) will come into play.

4.3.2 Symplectic Reduction

We begin this subsection by recalling the Marsden–Weinstein reduction process [33, 90], a well known tool in symplectic geometry, used to reduce spaces with symmetries. Concretely, this reduction process can be summarised as follows:

Theorem 4.3.3 (Marsden–Weinstein Reduction Process)

Let (N, ω) be a symplectic manifold and \mathbf{G} be a Lie group whose action on (N, ω) is Hamiltonian. Let $\mu : N \rightarrow \mathfrak{g}^$ be the moment map for this action, with \mathfrak{g} the Lie algebra of \mathbf{G} . Furthermore, let $c \in \mathfrak{g}^*$ be a regular value of μ and $\mathbf{G}_c \subseteq \mathbf{G}$ be the (coadjoint) stabiliser group of c . We assume that \mathbf{G}_c acts freely and properly on $\mu^{-1}(\{c\})$. Set $N_c := \mu^{-1}(\{c\})/\mathbf{G}_c$, with $\mathbf{p} : \mu^{-1}(\{c\}) \rightarrow N_c$ the natural projection, and consider,*

$$\begin{array}{ccc} \mu^{-1}(\{c\}) & \xleftarrow{\mathbf{i}} & N \\ \downarrow \mathbf{p} & & \\ N_c & & \end{array} \quad (4.3.12)$$

Then, there exists a unique symplectic form ω_c on N_c such that $\mathbf{p}^\omega_c = \mathbf{i}^*\omega$.*

As in Section 4.3.1, let M^3 be a three-dimensional manifold equipped with the metric (4.3.6). Let T^*M^3 be the corresponding cotangent bundle, which admits the following ‘twisted’ symplectic form:

$$\omega_3 := dq_i \wedge dx^i - d(\lambda q_3) \wedge dx^3, \quad (4.3.13)$$

where $\lambda \in \mathcal{C}^\infty(M^2)$ is non-vanishing. Evidently, $\partial_{x^3} \lrcorner \omega_3 = d(\lambda q_3)$ and ω_3 is closed, so ∂_{x^3} is Hamiltonian for ω_3 . Additionally, given a one-dimensional, connected Lie group \mathbf{G} whose action on T^*M^3 is generated by ∂_{x^3} , said action is Hamiltonian with respect to ω_3 . Consequently, we can consider a dimensional reduction of (T^*M^3, ω_3) under the action of \mathbf{G} , following Theorem 4.3.3.

As $\partial_{x^3} \lrcorner \omega_3$ is exact, we can define the moment map (4.3.5) on (T^*M, ω_3) by $\mu(x, q) = \lambda q_3$. Hence,

$$\mu^{-1}(\{c\}) = \left\{ (x, q) \mid q_3 = \frac{c}{\lambda} =: v_3(x^i) \right\} \quad (4.3.14)$$

for any regular value $c \in \mathbb{R}$. There is a natural embedding of $\mu^{-1}(\{c\})/\mathbb{G}_c$ into $\mu^{-1}(\{c\})$ and hence into T^*M^3 , which is given by $(x^i, x^3, q_i, q_3) = (x^i, C, q_i, q_3 = v_3(x^i))$ for any constant C . The geometric flow generated by our Hamiltonian corresponds to varying C and generates a foliation of $\mu^{-1}(\{c\})$ with leaves $\mu^{-1}(\{c\})/\mathbb{G}_c$. Consequently, $\mu^{-1}(\{c\})/\mathbb{G}_c \cong T^*M^2$ and can be given local coordinates (x^i, q_i) by Theorem 4.3.3. Furthermore, Theorem 4.3.3 also yields the unique symplectic form $\omega_c := dq_i \wedge dx^i$ on $\mu^{-1}(\{c\})/\mathbb{G}_c \cong T^*M^2$ satisfying $\mathbf{p}^*\omega_c = \mathbf{i}^*\omega_3$, as well as two closed differential two-forms given by

$$\begin{aligned} \tilde{\omega}_2 &:= e^\varphi (\varpi_2 + q_i \overset{\circ}{\nabla}^i \varphi \text{vol}_{M^2}) , \\ \tilde{\alpha}_2 &:= e^\varphi \left\{ \alpha_2 - [\hat{h}_+ + e^{-2\varphi} (\overset{\circ}{\nabla}^i \varphi \partial_i \varphi v_3^2 - v_3 \overset{\circ}{\nabla}^i \varphi \partial_i v_3)] \text{vol}_{M^2} + q_i \overset{\circ}{\nabla}^i \varphi \varpi_2 \right\} , \end{aligned} \quad (4.3.15)$$

which are simply those from (4.3.10) with $q_3 = v_3(x^1, x^2)$.

Upon requiring the pull-backs of $\tilde{\omega}_2$ and $\tilde{\alpha}_2$ along (4.2.4) to vanish, we obtain

$$\begin{aligned} \overset{\circ}{\nabla}_i v^i &= -v^i \partial_i \varphi , \\ \overset{\circ}{\Delta}_{\text{B}} p + \overset{\circ}{\nabla}_i v^j \overset{\circ}{\nabla}_j v^i + \frac{1}{2} |v|^2 \overset{\circ}{R} &= -\overset{\circ}{g}^{ij} \partial_i \varphi \partial_j p + v^i v^j \overset{\circ}{\nabla}_i \partial_j \varphi \\ &\quad + e^{-2\varphi} [(\overset{\circ}{\Delta}_{\text{B}} \varphi - \overset{\circ}{g}^{ij} \partial_i \varphi \partial_j \varphi) v_3^2 + 2v_3 \overset{\circ}{g}^{ij} \partial_i \varphi \partial_j v_3] . \end{aligned} \quad (4.3.16)$$

These are precisely the divergence-free constraint (3.1.13b) and the pressure equation (3.1.13c), adapted to the warped product metric (4.3.6) and under the assumption that p is independent of x^3 . Additionally, the first equation of (4.3.16) can be rewritten as $\overset{\circ}{\nabla}^i (e^\varphi v_i) = 0$, hence by the Poincaré lemma, any solution is locally of the form

$$v_i = -\sqrt{\det(\overset{\circ}{g}_2)} e^{-\varphi} \varepsilon_{ij} \overset{\circ}{g}^{jk} \partial_k \psi , \quad (4.3.17)$$

for some $\psi \in \mathcal{C}^\infty(M^2)$. That is, we have a modified divergence-free constraint and stream function for the two-dimensional fluid flow produced by our reduction and this flow may be compressible. Evidently, when $\varphi = 0$, we find that v_3 is unconstrained by (4.3.16), x^3 coordinatises \mathbb{R} , and we obtain, from our reduced flow, the standard situation of an incompressible fluid in two dimensions, as discussed in Section 3.2.

Next, let X be a vector field on $\mu^{-1}(\{c\})/\mathbb{G}_c \cong T^*M^2$ and consider its horizontal lift \tilde{X} to

T^*M^3 using the Levi-Civita connection for the metric (4.3.6),

$$\tilde{X} := X + X \lrcorner dx^i \mathring{\Gamma}_{i3}^3 q_3 \frac{\partial}{\partial q_3} = X + X \lrcorner d\varphi q_3 \frac{\partial}{\partial q_3}. \quad (4.3.18)$$

Noting that $\tilde{X} \lrcorner \mathring{\nabla} q_3 = 0$ and $\varpi_2 \wedge (\alpha_2 - \hat{h}_+ \text{vol}_{M^2}) = 0$, we obtain from (4.3.9) that $\alpha_3 \wedge (\tilde{X} \lrcorner \alpha_3) = -2\varpi_2 \wedge X \lrcorner (\alpha_2 - \hat{h}_+ \text{vol}_{M^2}) \wedge \mathring{\nabla} q_3 \wedge dx^3$. Consequently, the endomorphism (4.2.14) becomes

$$\hat{\mathcal{J}}_3 \tilde{X} = \frac{1}{\sqrt{|\hat{f}_2 + \hat{h}_+|}} \varepsilon_2 \lrcorner [\varpi_2 \wedge X \lrcorner (\alpha_2 - \hat{h}_+ \text{vol}_{M^2})], \quad (4.3.19)$$

where ε_2 is the dual to the Liouville volume form on T^*M^2 . Hence, we obtain an endomorphism on $\mu^{-1}(\{c\})/\mathbb{G}_c$, which we denote by $\hat{\mathcal{J}}_2$, that is of the form (4.1.7) for the Monge–Ampère structure $(\varpi_2, \alpha_2 - \hat{h}_+ \text{vol}_{M^2})$.⁶ This makes clear the relationship between the Hitchin almost (para-)complex structure (4.2.14) and the almost (para-)complex structure (4.1.6) for two-dimensional, incompressible fluid flows; indeed, when $\varphi = 0$, then $\hat{h}_+ = 0$, we recover the Monge–Ampère structure (ϖ_2, α_2) on T^*M^2 , and (4.3.19) is precisely (4.1.6). Note that $\alpha_2 - \hat{h}_+ \text{vol}_{M^2}$ is simply α_2 with \hat{f}_2 replaced by $\hat{f}_2 + \hat{h}_+$. Furthermore, whilst ϖ_2 is closed, $\alpha_2 - \hat{h}_+ \text{vol}_{M^2}$ is not closed in general.

Mirroring (4.1.8) we set

$$\hat{\mathcal{K}}_2 := \text{sgn}(\hat{f}_2 + \hat{h}_+) \sqrt{|\hat{f}_2 + \hat{h}_+|} \mathring{\nabla} q_i \wedge dx^i. \quad (4.3.20)$$

Then, as before, $\hat{\mathcal{K}}_2(\hat{\mathcal{J}}_2 X, Y) = -\hat{\mathcal{K}}_2(X, \hat{\mathcal{J}}_2 Y)$ for all vector fields X and Y on $\mu^{-1}(\{c\})/\mathbb{G}_c$, so $\hat{g}_2(X, Y) := \hat{\mathcal{K}}_2(X, \hat{\mathcal{J}}_2 Y)$ is a Lychagin–Rubtsov metric on $\mu^{-1}(\{c\})/\mathbb{G}_c$. Explicitly,

$$\hat{g}_2 = \frac{1}{2}(\hat{f}_2 + \hat{h}_+) \mathring{g}_{ij} dx^i \odot dx^j + \frac{1}{2} \mathring{g}^{ij} \mathring{\nabla} q_i \odot \mathring{\nabla} q_j. \quad (4.3.21)$$

The pull-back of the metric (4.3.21) along

$$\iota : x^i \mapsto (x^i, q_i) := (x^i, -\sqrt{\det(\mathring{g}_2)} e^{-\varphi} \varepsilon_{ij} \mathring{g}^{jk} \partial_k \psi), \quad (4.3.22)$$

obtained by combining (4.2.4) and (4.3.17) is therefore

$$g_2 = \frac{1}{2} (\mathring{\Delta}_B \psi \mathring{\nabla}_i \partial_j \psi + T_{ij}) e^{-2\varphi} dx^i \odot dx^j, \quad (4.3.23a)$$

⁶Here, \hat{h}_+ is understood as a function of x^1, x^2, q_1 , and q_2 only, since on $\mu^{-1}(\{c\})/\mathbb{G}_c$, $q_3 = v_3(x^1, x^2)$.

with

$$\begin{aligned}
T_{ij} := & \dot{g}_{ij} \{ \dot{\nabla}^l \varphi \partial_l \psi (\dot{\nabla}^k \varphi \partial_k \psi - \dot{\Delta}_B \psi) - (\dot{\nabla}^k \varphi \partial_k \varphi) (\dot{\nabla}^l \psi \partial_l \psi) \\
& + \dot{\nabla}^k \varphi [\dot{\nabla}^l \psi \dot{\nabla}_k \partial_l \psi + v_3 (\partial_k v_3 - v_3 \partial_k \varphi)] \} \\
& + \partial_i \varphi \partial_j \varphi (\dot{\nabla}^k \psi \partial_k \psi) - \dot{\nabla}^k \psi [\partial_i \varphi \dot{\nabla}_j \partial_k \psi + \partial_j \varphi \dot{\nabla}_i \partial_k \psi].
\end{aligned} \tag{4.3.23b}$$

Evidently, when $\varphi = 0$ then $T_{ij} = 0$ and we recover the metric (3.2.8). It should be noted, however, that when $\varphi \neq 0$, the correspondence between the signature of g_2 and the sign of the pull-back of $(\hat{f}_2 + \hat{h}_+)$ via (4.3.22) is no longer a direct assignment, as we shall see in Section 4.4. This observation echoes the discussion following (4.2.12) for incompressible flows in dimension higher than two.

4.3.3 k-Plectic Reduction

We saw in Section 4.2 that multisymplectic geometry is an appropriate language with which to formulate higher dimensional flows. It can therefore be reasonably assumed that a multisymplectic generalisation of Theorem 4.3.3 would be an appropriate tool when considering symmetry reductions. Fortunately for us, [34] have recently produced such a generalisation:

Theorem 4.3.4 (Blacker–Marsden–Weinstein Reduction Process)

Let (N, ϖ) be a k -plectic manifold and \mathbf{G} be a Lie group whose action on (N, ϖ) is Hamiltonian. Let $\mu : N \rightarrow \bigwedge^{k-1} T^*N \otimes \mathfrak{g}^*$ be the moment map for this action, with \mathfrak{g} the Lie algebra of \mathbf{G} . Furthermore, let $c \in \Omega^{k-1}(N, \mathfrak{g}^*)$ be closed and define

$$\begin{aligned}
\mu^{-1}(\{c\}) & := \{x \in N \mid \mu(x) = c_x\}, \\
\mathbf{G}_c & := \{g \in \mathbf{G} \mid g_*^{-1} X_1 \lrcorner \dots \lrcorner g_*^{-1} X_{k-1} \lrcorner \text{Ad}_g^* c_{g^{-1}x} = X_1 \lrcorner \dots \lrcorner X_{k-1} \lrcorner c_x \\
& \text{for all } x \in N \text{ and for all } X_1, \dots, X_{k-1} \in T_x N \}.
\end{aligned} \tag{4.3.24}$$

Suppose that $\mu^{-1}(\{c\})$ is a submanifold of N with (smooth) embedding $\mathfrak{i} : \mu^{-1}(\{c\}) \hookrightarrow N$ and that \mathbf{G}_c acts freely and properly on $\mu^{-1}(\{c\})$. Set $N_c := \mu^{-1}(\{c\})/\mathbf{G}_c$, again with natural projection $\mathfrak{p} : \mu^{-1}(\{c\}) \rightarrow N_c$ and consider,

$$\begin{array}{ccc}
\mu^{-1}(\{c\}) & \xleftarrow{\mathfrak{i}} & N \\
\downarrow \mathfrak{p} & & \\
N_c & &
\end{array} \tag{4.3.25}$$

Then, there exists a unique, closed differential form $\varpi_c \in \Omega^{k+1}(N_c)$ on N_c such that $\mathfrak{p}^* \varpi_c = \mathfrak{i}^* \varpi$.

Evidently, for $k = 1$ this result reduces to Theorem 4.3.3. It is important to stress that for $k > 1$, the differential form $\varpi_c \in \Omega^{k+1}(N_c)$ might be degenerate.

Again, let M^3 be a three-dimensional manifold equipped with the metric (4.3.6). Recall from (4.3.10) that ∂_{x^3} is a 2-plectic vector field for α_3 and a Hamiltonian vector field for ϖ_3 . Additionally, given a one-dimensional, connected Lie group \mathbf{G} whose action on T^*M^3 is generated by ∂_{x^3} , said action is Hamiltonian with respect to ϖ_3 . Consequently, we can consider a dimensional reduction of (T^*M^3, ϖ_3) under the action of \mathbf{G} , following Theorem 4.3.4.

In (4.3.11), we demonstrated that $\partial_{x^3} \lrcorner \varpi_3$ is exact, hence we can define the moment map on T^*M^3 with respect to ϖ_3 by

$$\mu(x, q) = \star_{\hat{g}_2} e^\varphi q_i dx^i. \quad (4.3.26)$$

For $\psi \in \mathcal{C}^\infty(M^2)$, the submanifold $\mu^{-1}(\{-d\psi\})$ is non-empty and given by

$$\mu^{-1}(\{-d\psi\}) = \left\{ (x, q) \mid q_i = -\sqrt{\det(\hat{g}_2)} e^{-\varphi} \varepsilon_{ij} \hat{g}^{jk} \partial_k \psi \right\}. \quad (4.3.27)$$

Furthermore, there is a natural embedding of $\mu^{-1}(\{-d\psi\})/\mathbf{G}_{-d\psi}$ into $\mu^{-1}(\{-d\psi\})$ and hence into T^*M^3 , which is given by $(x^i, x^3, q_i, q_3) = (x^i, C, -\sqrt{\det(\hat{g}_2)} e^{-\varphi} \varepsilon_{ij} \hat{g}^{jk} \partial_k \psi, q_3)$ for any constant C . The geometric flow generated by our Hamiltonian again corresponds to varying C and generates a foliation of $\mu^{-1}(\{-d\psi\})$ with leaves given by $\mu^{-1}(\{-d\psi\})/\mathbf{G}_{-d\psi}$. Consequently, $\mu^{-1}(\{-d\psi\})/\mathbf{G}_{-d\psi}$ can be given local coordinates (x^i, q_3) . Furthermore, there exists a closed differential form (in fact a volume form) $\varpi_{-d\psi} := e^{-\varphi} \text{vol}_{M^2} \wedge dq_3$ on $\mu^{-1}(\{-d\psi\})/\mathbf{G}_{-d\psi}$, which satisfies $\mathbf{p}^* \varpi_{-d\psi} = \mathbf{i}^* \varpi_3$ by Theorem 4.3.4. Returning to (4.3.10), while the pull-back of ϖ'_2 to $\mu^{-1}(\{-d\psi\})/\mathbf{G}_{-d\psi}$ vanishes identically, the equation $\mathbf{p}^* \alpha_{-d\psi} = \mathbf{i}^* \alpha'_2$ is satisfied by the two-form

$$\alpha_{-d\psi} := e^\varphi \left[\det(\hat{\nabla}^i q_j) - (\hat{f}_2 + \hat{h}_-) \right] \Big|_{q_i = -\sqrt{\det(\hat{g}_2)} e^{-\varphi} \varepsilon_{ij} \hat{g}^{jk} \partial_k \psi} \text{vol}_{M^2} + e^{-\varphi} q_3 dq_3 \wedge \star_{\hat{g}_2} d\varphi \quad (4.3.28)$$

on $\mu^{-1}(\{-d\psi\})/\mathbf{G}_{-d\psi}$, where \hat{h}_- is as defined in (4.3.9b). Finally, requiring that the pull-back of $\alpha_{-d\psi}$ along

$$\iota : x^i \mapsto (x^i, q_3) := (x^i, v_3(x^i)), \quad (4.3.29)$$

vanishes, the second equation of (4.3.16), with v_i given by $v_i = -\sqrt{\det(\hat{g}_2)} e^{-\varphi} \varepsilon_{ij} \hat{g}^{jk} \partial_k \psi$, is obtained, trivially satisfying the first equation from (4.3.16), as discussed around (4.3.17). Observe that, in contrast to the symplectic reduction where Poincaré's lemma is required, the 2-plectic reduction of the Monge–Ampère structure directly yields a two-dimensional flow (which may not be incompressible) in terms of a stream function ψ , with the trade off that the metric (4.3.21) is not obtained. As a result, the k -plectic reduction is a more elegant, compact tool, should a description of the reduced kinematics be all that is required.

Remark 4.3.5 (Non-zero γ)

The reductions presented in this section assume a one-dimensional symmetry of the underlying manifold M^3 , generated by ∂_{x^3} and characterised by complete x^3 independence of the velocity components and pressure. This corresponds to the class of fluid flows (4.3.1) with $\gamma = 0$ and $W = W(x^1, x^2, t)$. However, it should be clear that the above approach may also be applied to more general cases where the symmetry lies in $\Gamma(TT^*M^3)$, in order to obtain the full class of flows (4.3.1), which are characterised by linear dependence of v_3 on x^3 . Indeed, it has already been shown in [19] that the infinitesimal generator $\partial_{x^3} + \gamma(t)\partial_{q_3}$ yields a flow with $W \equiv 0$ and $\gamma = \gamma(t)$ via symplectic reduction, namely Burgers' vortex. In this sense, the reduction of the Monge–Ampère structure (4.2.1) in three dimensions can be considered a geometric description of the Lundgren transformation [91], which reduces certain three-dimensional, incompressible fluid flows on Euclidean background with symmetry to equivalent flows in two dimensions. Furthermore, by encoding this in a geometric framework, we have extended the Lundgren transformation to fluid flows on Riemannian background.

△

4.4 Incompressible Fluids in Euclidean Space — Examples

Let us now discuss some examples of incompressible fluid dynamical flows on three-dimensional Euclidean space, focusing on the reduction process outlined in Section 4.3. We provide an example for each of the connected, one-dimensional Lie groups whose actions can be generated by the vector field ∂_{x^3} , namely \mathbb{R} and \mathbb{S}^1 .

4.4.1 Arnold–Beltrami–Childress Flows (\mathbb{R} Symmetry)

Let us consider flows on $M_3 := \mathbb{R}^3$ equipped with the standard Euclidean metric

$$\hat{g}_3 := \hat{g}_2 + dz \otimes dz \quad \text{with} \quad \hat{g}_2 := dx \otimes dx + dy \otimes dy, \quad (4.4.1)$$

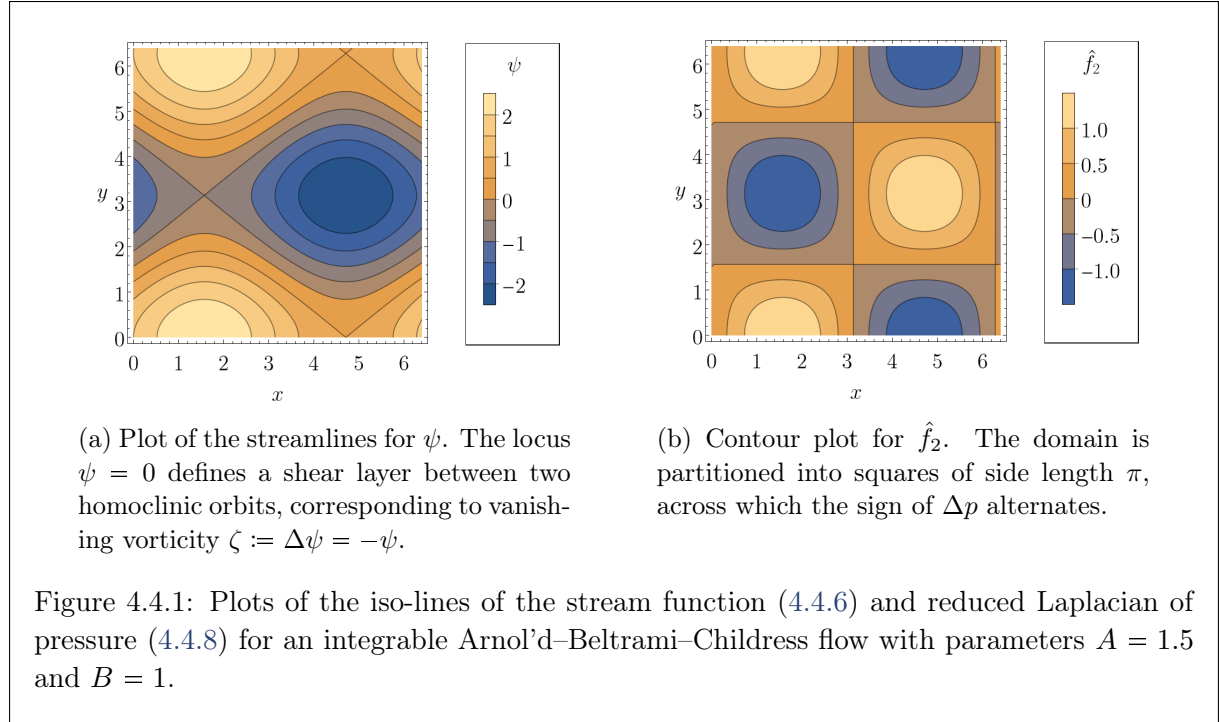
which corresponds to (4.3.6) when $\varphi = 0$. Note that $\hat{h}_\pm = 0$. In this case, the infinitesimal generator $\partial_{x^3} = \partial_z$ corresponds to the choice of Lie group $\mathbf{G} = \mathbb{R}$ acting by translation in z . Our symplectic reduction yields an incompressible fluid flow in two dimensions, on an Euclidean background. In summary, the equations (4.3.10) reduce to $\tilde{\varpi}_2 = \varpi_2$, and $\tilde{\alpha}_2 = \alpha_2$, with the divergence-free constraint and the pressure equation (4.3.16) respectively given by

$$\partial_x v_x + \partial_y v_y = 0 \quad (4.4.2a)$$

and

$$\Delta p = 2(\partial_x v_x \partial_y v_y - \partial_x v_y \partial_y v_x) \quad \text{with} \quad \Delta := \partial_x^2 + \partial_y^2, \quad (4.4.2b)$$

where v_x and v_y are functions of x and y only.



Additionally, performing the 2-plectic reduction from Section 4.3.3 to obtain velocity components v_x and v_y satisfying (4.4.2a), in terms of a stream function in two dimensions, yields the same result as applying the Poincaré lemma to (4.4.2a) itself, that is,

$$q_x := v_x = -\partial_y \psi \quad \text{and} \quad q_y := v_y = \partial_x \psi \quad (4.4.3)$$

for some stream function $\psi = \psi(x, y)$. The differential form corresponding to (4.3.28) is

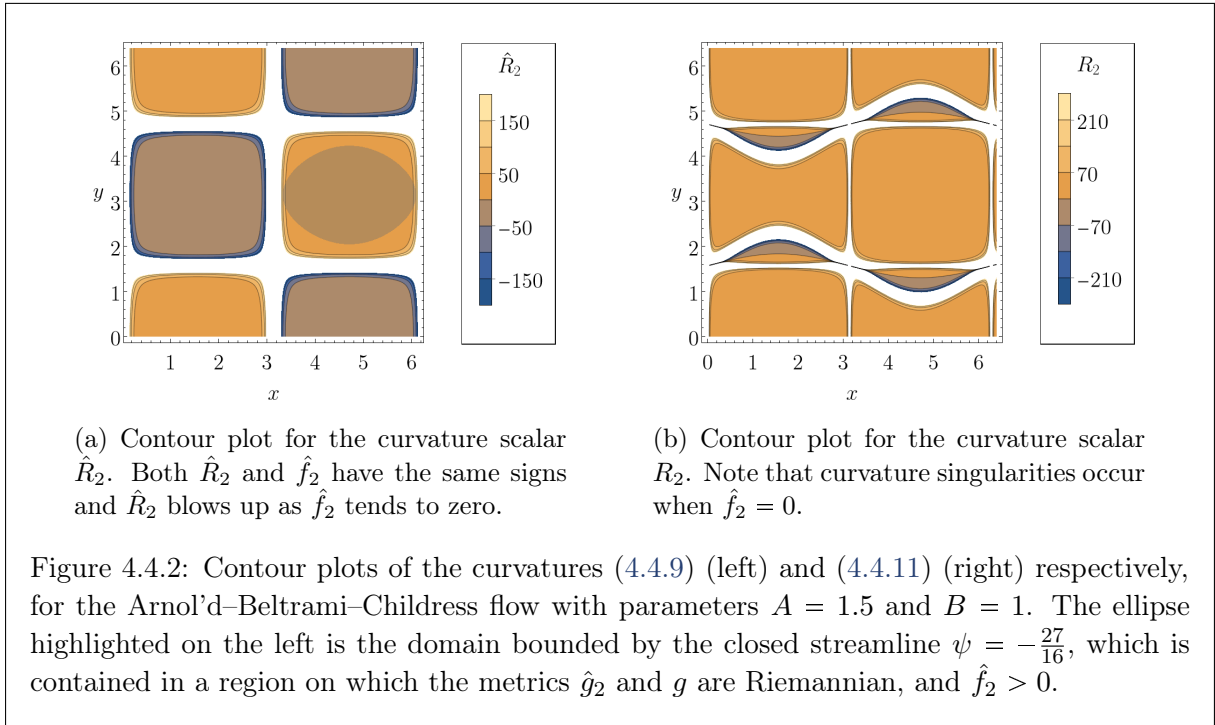
$$\alpha_{-d\psi} = \left[\partial_x^2 \psi \partial_y^2 \psi - (\partial_x \partial_y \psi)^2 - \frac{1}{2} \Delta p \right] dx \wedge dy, \quad (4.4.4)$$

which is unchanged when pulled back along $(x, y) \mapsto (x, y, q_z) := (x, y, v_z(x, y))$. Imposing that said pull-back vanishes is therefore equivalent to the Monge–Ampère equation

$$\frac{1}{2} \Delta p = \partial_x^2 \psi \partial_y^2 \psi - (\partial_x \partial_y \psi)^2, \quad (4.4.5)$$

which is, in turn, precisely (4.4.2b) with v_x and v_y evaluated as per (4.4.3). Hence, upon making the free choice of a pair of z -independent functions ψ and v_z , an incompressible fluid flow in \mathbb{R}^3

that reduces to an incompressible flow on the (x, y) -plane is recovered.



Making the choice

$$v_z(x, y) = \psi(x, y) := A \cos(y) + B \sin(x) \quad (4.4.6)$$

for $A, B \in \mathbb{R}$ some constants (see Figure 4.4.1a) and computing (4.4.3), we recover the velocity field for the integrable case of *Arnol'd–Beltrami–Childress flow* [92],

$$(v_x, v_y, v_z) = (\dot{x}, \dot{y}, \dot{z}) = (A \sin(y), B \cos(x), A \cos(y) + B \sin(x)). \quad (4.4.7)$$

Next, following [92], we take the quotient of v_x and v_y to deduce that this system integrates to $v_z = A \cos(y) + B \sin(x) = \text{const}$. Furthermore, (4.4.5) becomes

$$\hat{f}_2 = \frac{1}{2} \Delta p = AB \sin(x) \cos(y), \quad (4.4.8)$$

as displayed in Figure 4.4.1b.

Since $\hat{h}_+ = 0$ and $M^2 = \mathbb{R}^2$, it follows that the metric (4.3.21) on the reduced phase space $\mu^{-1}(\{c\})/G_c \cong T^*\mathbb{R}^2$ is precisely (3.3.2). Hence, we may follow exactly the treatment from Section 3.3 to obtain the curvature scalar \hat{R}_2 for the metric (4.3.21). In particular, (3.3.3)

yields:

$$\hat{R}_2 = \frac{\sin^2(x) + \cos^2(y)}{AB \sin^3(x) \cos^3(y)}, \quad (4.4.9)$$

and as in previous examples, when $\hat{f}_2 \geq 0$ the metric \hat{g}_2 is Riemannian/Kleinian and the associated curvature is positive/negative. Again, when $\hat{f}_2 = 0$, both the metric and the curvature scalar are singular.

In turn, the pull-back metric (4.3.23), with v_x and v_y as given in (4.4.7), is

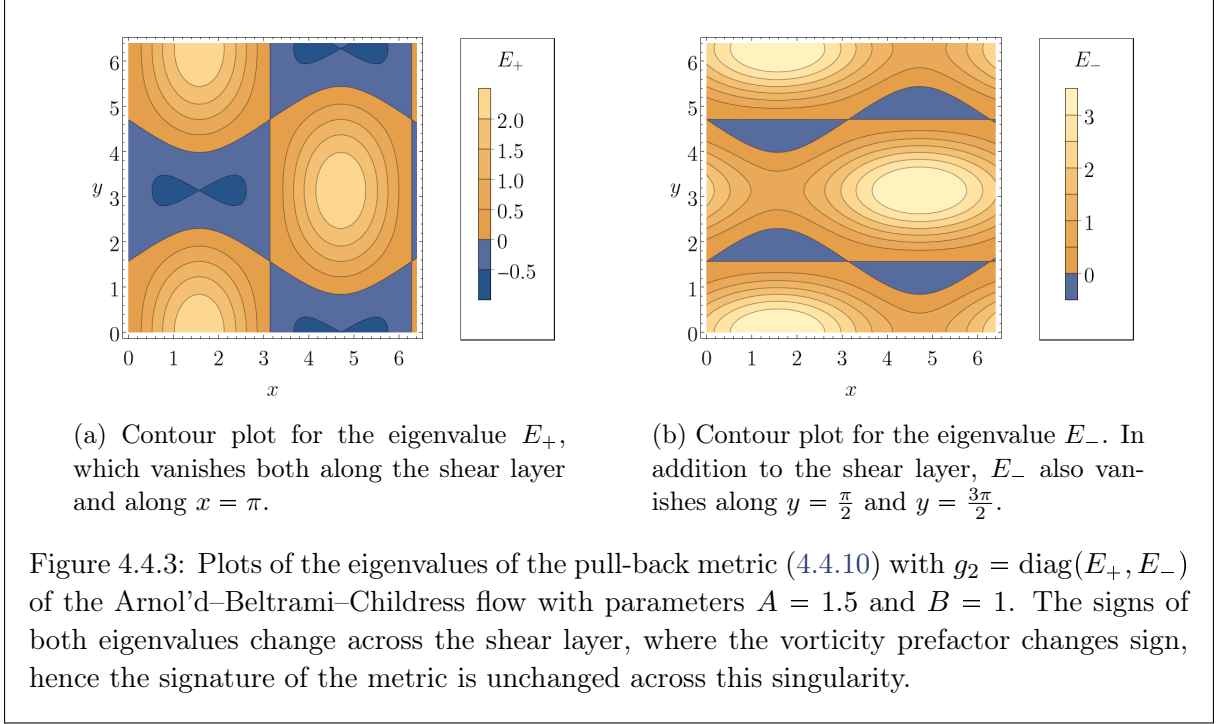
$$g_2 = [A \cos(y) + B \sin(x)] \begin{pmatrix} B \sin(x) & 0 \\ 0 & A \cos(y) \end{pmatrix}, \quad (4.4.10)$$

with vorticity is $\zeta = -A \cos(y) - B \sin(x)$. This metric is again degenerate when $\hat{f}_2 = 0$. There is a further degeneracy when $A \cos(y) + B \sin(x) = 0$, precisely along the shear layer featuring in the streamlines of Figure 4.4.1a, which corresponds to the vorticity (3.2.9) vanishing. The curvature scalar R_2 associated with (4.4.10) is

$$R_2 = \frac{B \sin(x) [\sin^2(x) + 3 \cos^2(y)] + A \cos(y) [\cos^2(y) + 3 \sin^2(x)]}{2 \sin^2(x) \cos^2(y) [B \sin(x) + A \cos(y)]^3}. \quad (4.4.11)$$

Observe that the lines $x = n\pi$ and $y = (n + \frac{1}{2})\pi$ for all $n \in \mathbb{Z}$, along which $f_2 = 0$, are singularities of both the metric g and its curvature R , as was the case for the metric (4.3.21). Additionally, the presence of $A \cos(y) + B \sin(x)$ in the denominator illustrates that the shear layer is a curvature singularity, see Figure 4.4.2b. This curvature singularity arises due to the vanishing vorticity and is otherwise unseen by the Q-criterion discussed in Section 3.1.1. The shear layer is a separatrix between topologically distinct flows.

Furthermore, the metric g_2 is globally degenerate when $A = 0$ or $B = 0$ independently, in addition to when both A and B vanish (in which case both the vorticity and the Hessian part of the metric vanish). Whilst the latter choice corresponds to $\psi = 0$, where there is no flow, the former two choices correspond to cases where the flow is trivial due to an additional symmetry: let $B = 0$ (resp. $A = 0$) such that the stream function (4.4.6) depends only on y (resp. x). The streamlines are then simply $y = \text{const.}$ (resp. $x = \text{const.}$) and after our (multi-)symplectic reduction along the z -axis, we have a one-dimensional flow in the (x, y) -plane. Compare this to the behaviour we saw in Section 3.3.2 when global degeneracy occurred.



4.4.2 Hicks–Moffatt Vortex (\mathbb{S}^1 Symmetry)

Let us now consider flows on $M := (\mathbb{R}^+ \times \mathbb{R}) \times_{r,2} \mathbb{S}^1$, equipped with the metric

$$\hat{g}_3 := \hat{g}_2 + r^2 d\theta \otimes d\theta \quad \text{with} \quad \hat{g}_2 := dr \otimes dr + dz \otimes dz, \quad (4.4.12)$$

where $r \in \mathbb{R}^+$, $z \in \mathbb{R}$, and $\theta \in [0, 2\pi)$, that is, standard cylindrical coordinates. Then,

$$\varphi = \log(r) \quad \text{and} \quad \hat{h}_+ = \frac{1}{2r} \partial_r p, \quad (4.4.13)$$

with $p = p(r, z)$ in (4.3.10). In this case, the infinitesimal generator $\partial_{x^3} = \partial_\theta$ corresponds to the choice of Lie group $\mathbf{G} = \mathbb{S}^1$, acting by translation in θ modulo 2π . That is, we consider the case of rotational symmetry.

Hence, the equations (4.3.16) become

$$\begin{aligned} \tilde{\omega}_2 &= r \left(\varpi_2 + \frac{1}{r} q_r dr \wedge dz \right), \\ \tilde{\alpha}_2 &= r \left\{ \alpha_2 - \left[\frac{1}{2r} \partial_r p + \frac{1}{r^2} \left(\frac{1}{r^2} q_\theta^2 - \frac{1}{r} q_\theta \partial_r q_\theta \right) \right] dr \wedge dz + \frac{1}{r} q_r \varpi_2 \right\}. \end{aligned} \quad (4.4.14)$$

Furthermore, the requirements that the pull-backs of $\tilde{\omega}_2$ and $\tilde{\alpha}_2$ under (4.1.2) vanish become

$$\frac{1}{r} \partial_r (r v_r) + \partial_z v_z = 0, \quad (4.4.15a)$$

and

$$\frac{1}{r}\partial_r(r\partial_r p) + \partial_z^2 p = 2\left[\partial_r v_r \partial_z v_z - \partial_r v_z \partial_z v_r - \frac{1}{r^2}v_r^2 - \frac{1}{r^4}(v_\theta^2 - \frac{r}{2}\partial_r v_\theta^2)\right], \quad (4.4.15b)$$

which are the equations (4.3.16) for the metric (4.4.12), with $v_\theta = v_\theta(r, z)$ arbitrary.

Turning to the 2-plectic reduction from Section 4.3.3, the moment map (4.3.26) is given by

$$\mu(x, q) = r q_r dz - r q_z dr. \quad (4.4.16)$$

It then follows that, locally on $\mu^{-1}(\{-d\psi\})/G_{-d\psi}$, we have

$$q_r := v_r = -\frac{1}{r}\partial_z \psi \quad \text{and} \quad q_z := v_z = \frac{1}{r}\partial_r \psi, \quad (4.4.17)$$

which can be interpreted as expressions for the velocity components in the r and z directions, in terms of a stream function $\psi = \psi(r, z)$ in two dimensions. Consequently, these solve the adapted divergence-free constraint (4.4.15a). In fact, imposing that the pull-back of the closed differential form (4.3.28) along $(r, z) \mapsto (r, z, q_\theta) := (r, z, v_\theta(r, z))$ vanishes, we find

$$\begin{aligned} \frac{1}{2}\left[\frac{1}{r}\partial_r(r\partial_r p) + \partial_z^2 p\right] &= \frac{1}{r^2}\left[\partial_r^2 \psi \partial_z^2 \psi - (\partial_r \partial_z \psi)^2\right] - \frac{1}{r^4}(\partial_z \psi)^2 \\ &+ \frac{1}{r^3}(\partial_z \psi \partial_r \partial_z \psi - \partial_r \psi \partial_z^2 \psi) - \frac{1}{r^4}(v_\theta^2 - \frac{r}{2}\partial_r v_\theta^2), \end{aligned} \quad (4.4.18)$$

which is precisely (4.4.15) with v_r and v_z given in terms of ψ , as in (4.4.17). Again, we are free to choose ψ and v_θ , provided they satisfy (4.4.18). Furthermore, (4.4.15a) is trivially satisfied for any such choices, given (4.4.17).

Consequently, appropriate choices of v_θ and ψ allow us to study another important class of examples — *Hicks–Moffatt vortices* — a family of spherical vortices with *swirl parameter* κ , studied by [93] and rediscovered by [30]. For an in-depth review of such vortices, we direct the interested reader to [94] (see also [95] in the context of magneto-hydrodynamics). We obtain these examples by first considering a unit sphere in \mathbb{R}^3 and setting $\sigma(r, z) := \sqrt{r^2 + z^2}$. Recall from (4.4.12) that we are working in cylindrical polar coordinates. We normalise the scale of the sphere for convenience and call the sets of points with $\sigma^2 \gtrless 1$ the exterior/interior of the vortex, respectively. Next, fix the angular velocity to be

$$v_{\theta, \kappa}(r, z) = \frac{\kappa \psi}{r}, \quad (4.4.19)$$

on the whole domain. The swirl parameter κ scales the angular velocity and flows without swirl ($\kappa = 0$) therefore have vanishing angular component to their velocity field.

We fix the stream function on the interior and exterior of the sphere separately, under the

condition that they coincide on the boundary given by $\sigma^2 = 1$. In particular, on the interior we set

$$\psi_{\text{int},\kappa}(r, z) := \frac{3}{2}r^2 \left(b(\kappa) - c(\kappa) \frac{J_{\frac{3}{2}}(\kappa\sigma)}{(\kappa\sigma)^{\frac{3}{2}}} \right), \quad (4.4.20a)$$

with

$$b(\kappa) := \frac{J_{\frac{3}{2}}(\kappa)}{\kappa J_{\frac{5}{2}}(\kappa)} \quad \text{and} \quad c(\kappa) := \frac{\sqrt{\kappa}}{J_{\frac{5}{2}}(\kappa)}, \quad (4.4.20b)$$

where $J_n(x)$ is the n -th order *Bessel function* with argument x . On the exterior of the sphere, we choose the stream function independently of the swirl parameter; for example, we choose

$$\psi_{\text{ext}}(r, z) := \frac{1}{2}r^2 \left(1 - \frac{1}{\sigma^3} \right), \quad (4.4.21)$$

so that far from the sphere, the (non-angular) velocity is uniform, with unit speed directed along the z -axis. The flow has non-zero velocity in the θ direction when $\kappa \neq 0$. It is important to note here, in the context of [Section 4.2.2](#), that it is possible for a Hicks–Moffatt vortex to have vanishing helicity but non-vanishing vorticity when $\kappa = 0$.

Let us focus on this limiting case, which is referred to as *Hill's spherical vortex* [96]. The θ -component of velocity (4.4.19) vanishes in this case, while the stream function for the exterior solution (4.4.21) remains unchanged. Henceforth, we focus our attention on the interior solution, whose stream function (4.4.20) is given by

$$\psi_{\text{int},0}(r, z) := \frac{3}{4}r^2(r^2 + z^2 - 1), \quad (4.4.22)$$

see [Figure 4.4.4a](#). Upon applying (4.4.15a), we obtain the velocity components

$$v_r = -\frac{3}{2}rz \quad \text{and} \quad v_z = \frac{3}{2}(2r^2 + z^2 - 1), \quad (4.4.23)$$

It then follows from (4.4.18) that the Laplacian of pressure⁷ is given by

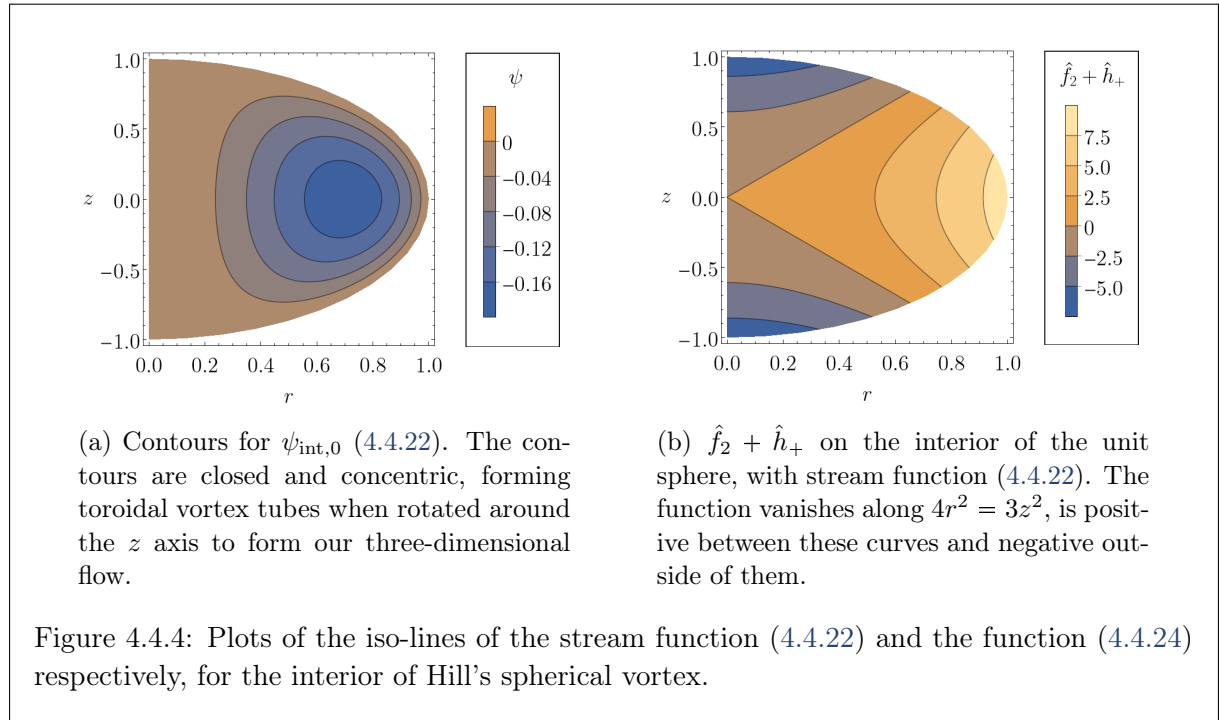
$$\hat{f}_2 + \hat{h}_+ = \frac{1}{2}(\partial_r^2 p + \partial_z^2 p + \frac{1}{r}\partial_r p) = \frac{9}{4}(4r^2 - 3z^2), \quad (4.4.24)$$

see [Figure 4.4.4b](#).

The metric (4.3.21) takes the form

$$\hat{g}_2 = \begin{pmatrix} (\hat{f}_2 + \hat{h}_+) \mathbb{1}_2 & 0 \\ 0 & \mathbb{1}_2 \end{pmatrix}. \quad (4.4.25)$$

⁷Note here that $\hat{f}_2 + \hat{h}_+$ has no dependence on q_i, q_θ , or θ , hence is unchanged when pulled back along (4.1.2).



and \hat{R}_2 is given by (3.3.3), with f replaced by $\hat{f}_2 + \hat{h}_+$. Namely

$$\hat{R}_2 = \frac{56(4r^2 + 3z^2)}{9(4r^2 - 3z^2)^3}, \quad (4.4.26)$$

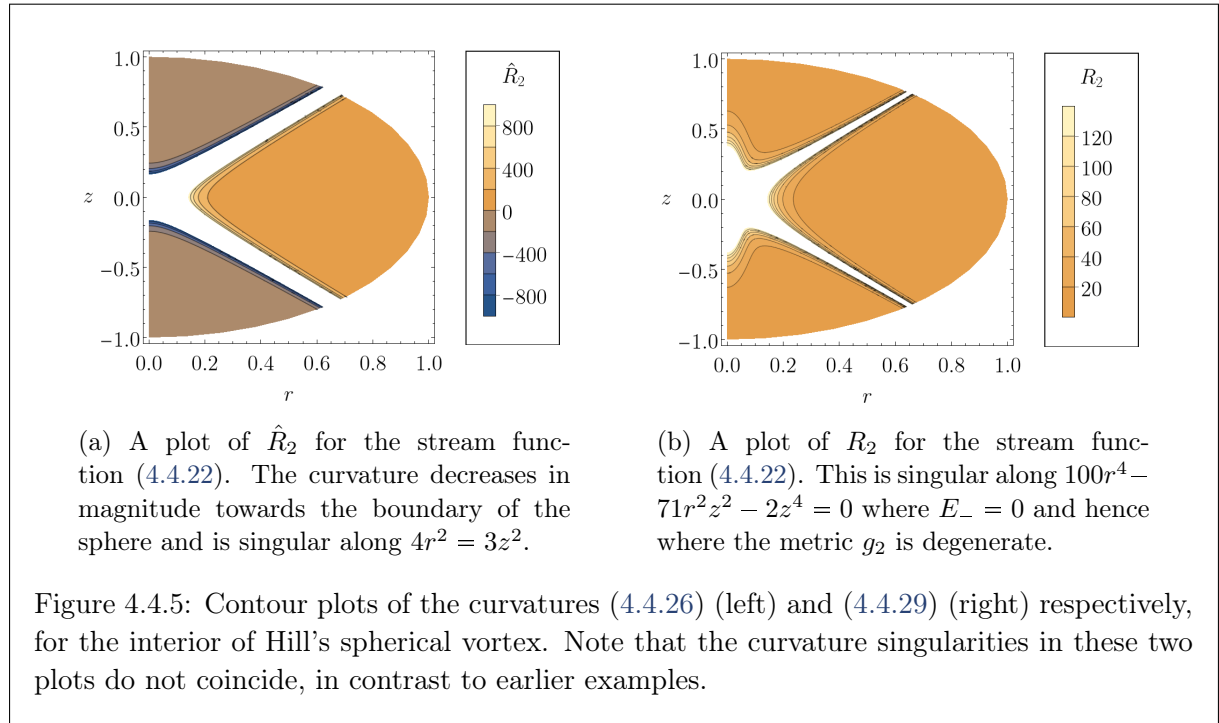
which is plotted in Figure 4.4.5a. When $4r^2 > 3z^2$, then $\hat{f}_2 + \hat{h}_+ > 0$, the metric is Riemannian, and the Ricci curvature scalar is positive. Similarly, the metric is Kleinian and the curvature scalar negative when $\hat{f}_2 + \hat{h}_+ < 0$ and $4r^2 < 3z^2$. Furthermore, the metric is singular when $4r^2 = 3z^2$, that is, when $\hat{f}_2 + \hat{h}_+ = 0$, and it is clear that this singularity is also one for the curvature.

The pull-back metric (4.3.23) becomes

$$g_2 = \frac{9}{4} \begin{pmatrix} 20r^2 - 2z^2 & 9rz \\ 9rz & 5r^2 + z^2 \end{pmatrix}. \quad (4.4.27)$$

Its eigenvalues, displayed in Figure 4.4.6, are given by

$$E_{\pm} = \frac{9}{8} \left(25r^2 - z^2 \pm 3\sigma \sqrt{(25r^2 + z^2)} \right). \quad (4.4.28)$$



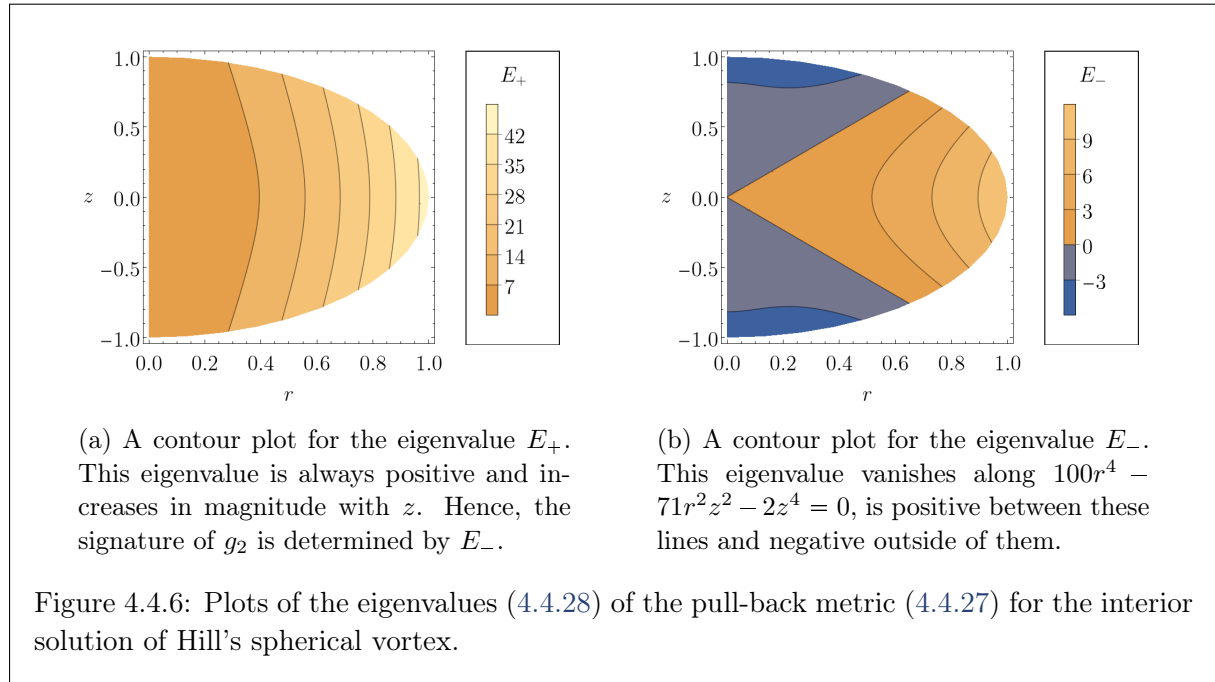
Furthermore, the curvature scalar R_2 associated with (4.4.27) is

$$R_2 = \frac{28(50r^4 + z^4)}{9(100r^4 - 71r^2z^2 - 2z^4)^2}, \quad (4.4.29)$$

see Figure 4.4.5b. Both the curvature R_2 and eigenvalue E_+ are non-negative, with E_+ vanishing only at the origin. Additionally, R_2 is singular precisely when E_- vanishes, that is, where the pull-back metric (4.4.27) is singular. However, this does not coincide with when $\hat{f}_2 + \hat{h}_+ = 0$. In particular, the region on which $\hat{f}_2 + \hat{h}_+ > 0$ falls within the region where g_2 is Riemannian. Hence, upon noting that (4.3.9) implies $\hat{f}_3 = \hat{f}_2 + \hat{h}_+$, it follows from (3.1.17) that dominance of vorticity over strain implies the pull-back metric is Riemannian, but not vice-versa. Similarly, the pull-back metric being Kleinian implies that strain dominates, however strain may still dominate on Riemannian regions of the solution.

4.5 Chapter Summary

In this chapter, we have demonstrated how the Monge–Ampère geometric formulation of the Poisson equation for the pressure of a two-dimensional incompressible fluid flow can be extended to fluid flows in (arbitrary) higher dimension. We began in Section 4.1 by illustrating that taking a non-canonical choice of symplectic form (4.1.1) in our Monge–Ampère structure, while leaving the Monge–Ampère form unchanged, allows us to encode both the divergence-free con-



straint (3.1.13b) and the Poisson equation for the pressure (3.1.13c).

Subsequently, in Section 4.2, we observed that the differential forms (ϖ, α) from our modified Monge–Ampère structure generalise to a pair of $(m - 1)$ -plectic forms on T^*M when M is m -dimensional and that this pair of forms also yields (3.1.13b) and (3.1.13c) in the general case. As in Section 3.2, we were then able to construct a Lychagin–Rubtsov metric (4.2.7) on T^*M whose signature is determined by the quantity $\hat{f} = \frac{1}{2}(\hat{\Delta}_B p + q_i q_j \hat{R}^{ij})$, though not by its sign alone in general. The block form of this metric is independent of m and the pull-back of the metric along (4.2.4) can be given solely in terms of the underlying metric \hat{g} on M and the velocity gradient tensor A_{ij} of the fluid flow. In this sense, the geometry we construct is determined by the gradients of vorticity and strain of the fluid and the signatures of the Lychagin–Rubtsov metric and its pull-back can be seen as a generalisation of the Q-criterion for a vortex. Specialising to the three-dimensional case, in Section 4.2.1, we show that this metric is a conformal scaling of the classical Lychagin–Rubtsov metric (2.2.16) for α , while in Section 4.2.2, we demonstrated how topological quantities can be inferred from our geometric framework, as in the two-dimensional case.

Finally, in Section 4.3, we considered a large subclass of the so-called ‘two-and-a-half-dimensional flows’ (4.3.1), with symmetry generated by ∂_{x^3} . Such flows were shown to admit dimensional reduction via connected, one-dimensional Lie groups, using both the symplectic and 2-plectic Marsden–Weinstein processes, resulting in a two-dimensional Navier–Stokes flow, which need not be incompressible. This reduction is in the same spirit as the Lundgren trans-

formation and extends the results of [19, Section 5] beyond Burgers' vortex. Furthermore, the equations (4.3.16) for the reduced flow can be encoded by a Monge–Ampère structure (4.3.10) and can hence be studied in the same manner as the incompressible two-dimensional flows of Section 3.2, in that they admit a corresponding Lychagin–Rubtsov metric (4.3.21) whose signature is associated with the gradients of vorticity and strain. We closed the chapter in Section 4.4 by presenting a pair of examples, namely the Arnold–Beltrami–Childress flow, which has symmetry corresponding to $\mathbf{G} = \mathbb{R}$, and the Hicks–Moffatt vortex, which has symmetry corresponding to $\mathbf{G} = \mathbb{S}^1$.

In the conclusion to this part of the thesis, we shall discuss ongoing work to develop a theory of higher Monge–Ampère structures, which include the structure (4.2.1) arising from incompressible fluid flows in dimension greater than two. We take particular interest in their classification, providing necessary conditions for the identity (4.2.13) to define an endomorphism, before making first steps towards a Lychagin–Rubtsov type theorem in the three dimensional case.



5

Part I Outlook

In this part of the thesis, we have established a Monge–Ampère geometric framework for studying incompressible fluid dynamical flows, inspired by techniques from multisymplectic geometry, with the differential forms (4.2.1) taking the place of the Navier–Stokes equations (3.1.9) and submanifolds (4.2.4) being considered in lieu of solutions. This approach provides a unified description of incompressible fluids on Riemannian manifolds of arbitrary dimension $m \in \mathbb{N} \setminus \{1\}$. Consequently, our framework is not only capable of recovering known results, such as the Weiss criterion on a two-dimensional Euclidean background, see Section 3.1.1, but also extending these results to curved backgrounds and higher dimensions by means of the signature of the Lychagin–Rubtsov metric and type of the endomorphism J , as detailed in Table 3.4.1. Furthermore, we highlighted how topological information about vortices might be obtained from our geometry, using the Gauß–Bonnet theorem in Section 3.2.1 and helicity in Section 4.2.2, for example. Finally, in Section 4.3.1, we illustrated the application of the symplectic and 2-plectic reduction principles to incompressible fluid flows in three dimensions admitting a one-dimensional Lie symmetry, in order to obtain a two-dimensional fluid flow, which need not be incompressible. Such a reduction might shed light on the relationship between vortices in two and three dimensions, since their distinct qualitative behaviour, as described in Section 3.1.1, is hitherto poorly understood.

In the remainder of this chapter, we provide a brief overview of a preprint-in-progress, written by the author in collaboration with I. Roulstone, V. Rubtsov, and M. Wolf, which aims to develop a wider range of geometric tools for studying partial differential equations, such as those arising from our “higher Monge–Ampère structures.” We pay particular attention to the classification of said structures in what follows.

Let M be an m -dimensional Riemannian manifold with local coordinates $\{x^i\}_{i=1}^m$ and let T^*M be its cotangent bundle, equipped with coordinates $\{x^i, q_i\}_{i=1}^m$ as in earlier sections. Pulling back an arbitrary differential m -form

$$\beta = \sum_{k=0}^m \beta(x, q)^{i_1, \dots, i_k}_{j_{k+1}, \dots, j_m} dq_{i_1} \wedge \dots \wedge dq_{i_k} \wedge dx^{j_{k+1}} \wedge \dots \wedge dx^{j_m} \quad (5.0.1)$$

on T^*M along $\iota : x^i \mapsto (x^i, q_i) = (x^i, v_i(x))$, as in (4.2.4), yields the partial differential equation

$$\sum_{k=0}^m \beta(x, v(x))^{i_1, \dots, i_k}_{j_{k+1}, \dots, j_m} (\partial_{j_1} v_{i_1}) \dots (\partial_{j_k} v_{i_k}) \varepsilon^{j_1 \dots j_m} = 0. \quad (5.0.2)$$

Observe how the sum in (5.0.2) is a $\mathcal{C}^\infty(M)$ -linear combination of the minor determinants of the Jacobian of $v(x)$, hence when $m = 2$, pairs of linearly independent differential 2-forms (ϖ, α) are referred to as *Jacobi systems* of partial differential equations, see [4, Chapter 15.2]. Given a Jacobi system (ϖ, α) where both 2-forms are non-degenerate,

$$\alpha = A \lrcorner \varpi \quad (5.0.3)$$

defines an endomorphism A on TT^*M which is proportional to an almost (para-)complex structure J_α if and only if α is ϖ -effective.¹ Furthermore, we recover the class of Monge–Ampère equations by fixing one of the forms to be symplectic and applying the Poincaré lemma to obtain $v(x)$ in terms of the derivatives of some function $\psi \in \mathcal{C}^\infty(M)$, as in Section 2.1.3. In fact, any Jacobi system admitting a non-degenerate conservation law is locally given by a Monge–Ampère structure [4, Theorem 15.5.1] and we can apply the results of Chapter 2 in this case. We direct the interested reader to [4, Chapter 15] for more details on Jacobi systems.

Most significantly for us, the Poisson equation for the pressure (3.1.9c) and divergence-free constraint (3.1.9b) have the form (5.0.2) in arbitrary dimension, so constitute a Jacobi system in dimension two. It is this fact we exploited in Section 4.1 when constructing our alternative Monge–Ampère structure. However, in dimensions greater than two, there is no established theory for studying pairs of the form (5.0.2), which we formally called higher Monge–Ampère structures in earlier chapters. In particular, the equation (5.0.3) does not always define an endomorphism when $m > 2$; in Section 4.2.1 we invoked the Hitchin endomorphism to avoid this issue, however this approach only takes one of the two differential forms into account and is therefore not reflective of the whole structure. This raises the following questions for $m > 2$:

¹If A_α is proportional to an almost (para-)complex structure and α is not ϖ -effective, then $\alpha = \alpha_0 + f\varpi$ for some ϖ -effective α_0 and $f \in \mathcal{C}^\infty(T^*M)$. Then $\alpha_0 = (A_\alpha - fI_4) \lrcorner \varpi$ and by (2.2.14), $(A_\alpha - fI_4)^2$ is proportional to the identity, from which it follows that $f = 0$ and we have a contradiction.

- (i) When is the endomorphism A_α in (5.0.3) defined for m -forms (ϖ, α) ?
- (ii) When is A_α proportional to an almost (para-)complex structure and what takes the place of the Pfaffian (Definition 2.2.4) and effectiveness (Definition 2.1.7) in this case?
- (iii) Is there an analogue of the Lychagin–Rubtsov theorem (Theorem 2.2.6) for classifying systems of the form (5.0.2)?

Let us begin addressing these questions by presenting some properties of the identity (5.0.3) in arbitrary dimensions, before specialising to three dimensions, where more progress can be made. Our first proposition demonstrates that any endomorphism obtained from the identity (5.0.3) *must* be *compatible* with each of the forms used to define it, where compatibility of a differential form and endomorphism pair is defined as follows:

Definition 5.0.1 (Compatibility)

Let N be a smooth manifold, $\varpi \in \Omega^k(N)$ a differential k -form on N , and A an endomorphism on TN . The pair (ϖ, A) are called compatible if

$$\varpi(AX_1, X_2, \dots, X_k) = \varpi(X_1, AX_2, \dots, X_k) = \dots = \varpi(X_1, X_2, \dots, X_{k-1}, AX_k), \quad (5.0.4)$$

for all $X_1, \dots, X_k \in TN$.

Proposition 5.0.2 (Differential Forms and Compatibility)

Let N be a smooth manifold, $\varpi \in \Omega^k(N)$ a differential k -form on N , and A an endomorphism on TN . Then (ϖ, A) are compatible if and only if $\alpha := A \lrcorner \varpi$ is a differential k -form on TN , in which case (α, A) are also compatible.

Proof. Let ϖ, A , and α be as in the statement. Then we find

$$\begin{aligned} \alpha(X_1, \dots, X_k) &= \varpi(AX_1, X_2, \dots, X_k) = -\varpi(X_\ell, X_2, \dots, X_{\ell-1}, AX_1, X_{\ell+1} \dots X_k) \\ -\alpha(X_\ell, X_2, \dots, X_{\ell-1}, X_1, X_{\ell+1}, \dots, X_k) &= -\varpi(AX_\ell, X_2 \dots X_{\ell-1}, X_1, X_{\ell+1} \dots X_k), \end{aligned} \quad (5.0.5)$$

for all $X_1, \dots, X_k \in TN$ and each $\ell = 1, 2, \dots, k$. The left hand sides of the equations (5.0.5) are equal if and only if α is a (smooth, k -linear) skew-symmetric operator, while the right hand sides are equal if and only if (ϖ, A) are compatible, as required. Compatibility of (α, A) then follows directly from the definition of α . \square

Consequently, given a compatible k -form and endomorphism pair, we can always define a second (compatible) differential k -form via (5.0.3). Conversely, while a pair of differential k -forms need not define an endomorphism via (5.0.3), if they do so, then the endomorphism is

necessarily compatible with each of the k -forms. This illustrates the two approaches we may take when presenting results concerning the identity (5.0.3): we either begin with a compatible differential form and endomorphism pair (ϖ, A) and relate the properties of A to the properties of the differential form $\alpha := A \lrcorner \varpi$ (which we know exists by Proposition 5.0.2), or we begin with a pair of differential forms (ϖ, α) and relate the properties of α to the properties of the endomorphism A defined via (5.0.3), in the knowledge that such an A might not exist. We choose to present our results in the former manner, due to its elegance, leaving the latter for the surrounding contextual discourse.

Our next result provides necessary and sufficient conditions for an endomorphism defined via (5.0.3) to be an isomorphism, under the assumption that one of the defining differential forms is non-degenerate.

Proposition 5.0.3 (Non-Degeneracy and Compatibility)

Let N be a smooth manifold, $\varpi \in \Omega^k(N)$ a differential k -form on N , and A an endomorphism on TN , such that (ϖ, A) are compatible. Assume that ϖ is non-degenerate. Then $\alpha := A \lrcorner \varpi$ is non-degenerate if and only if A is an isomorphism. The analogous result holds if we exchange the roles of ϖ and α .

Proof. Let (ϖ, A) be a compatible k -form and endomorphism pair and define $\alpha = A \lrcorner \varpi$. Assume that ϖ is non-degenerate. Then, by direct computation, we have

$$X \lrcorner \alpha = X \lrcorner A \lrcorner \varpi = A(X) \lrcorner \varpi. \quad (5.0.6)$$

Since ϖ is non-degenerate, Definition 2.1.2 implies that the contraction map with ϖ is injective. Combining this with (5.0.6), we have

$$X \lrcorner \alpha = 0 \Leftrightarrow A(X) \lrcorner \varpi = 0 \Leftrightarrow A(X) = 0, \quad (5.0.7)$$

from which our result follows.² □

Thus far, we have shown a number of properties that an endomorphism A must possess if it is defined via (5.0.3), with (ϖ, α) both non-degenerate differential k -forms. In particular, A must be an isomorphism which is compatible with both ϖ and α . These necessary conditions contribute towards answering question (i) above by placing constraints on the A that can exist for a given pair of differential forms (ϖ, α) .

²Recall that a linear map is an isomorphism if and only if it is a monomorphism.

If we assume that the endomorphism defined via (5.0.3) exists and is an almost (para-) complex structure, we might ask what the necessary and sufficient conditions are for said almost (para-)complex structure to be integrable. The first of the following two results appears in [36, Proposition 2.8] in essentially the same form, with ϖ being referred to as an (almost) $(k-1)$ -plectic form on N , cf. Definition 2.1.4.

Proposition 5.0.4 (Closure Necessary for Integrability)

Let N be a smooth manifold, $\varpi \in \Omega^k(N)$ a non-degenerate k -form on N , and J an almost (para-) complex structure on TN , such that (ϖ, J) are compatible. Then, J is integrable if ϖ and $J \lrcorner \varpi$ are both closed.

Proof. Recall from our discussion around (2.2.5) that J is integrable if and only if the Nijenhuis tensor

$$N_J(X, Y) := -\varepsilon[X, Y] + J[X, JY] + J[JX, Y] - [JX, JY], \quad (5.0.8)$$

vanishes for all vector fields $X, Y \in TN$, where $J^2 =: \varepsilon \mathbb{I}$ and \mathbb{I} is the identity on TN . It is precisely this condition we wish to verify when $d\varpi = 0$ and $d(J \lrcorner \varpi) = 0$. Let us begin by recalling the definition of the Lie bracket and Cartan's homotopy formula, which are respectively given by

$$[X, Y] \lrcorner \varpi = \mathcal{L}_X(Y \lrcorner \varpi) - Y \lrcorner \mathcal{L}_X(\varpi) \quad \text{and} \quad \mathcal{L}_X(\varpi) = X \lrcorner d\varpi + d(X \lrcorner \varpi), \quad (5.0.9)$$

where $\mathcal{L}_X(\varpi)$ is the Lie derivative of ϖ along X . Combining these identities, we find

$$[X, Y] \lrcorner \varpi = X \lrcorner d(Y \lrcorner \varpi) - Y \lrcorner d(X \lrcorner \varpi) - Y \lrcorner X \lrcorner d\varpi + d(X \lrcorner Y \lrcorner \varpi). \quad (5.0.10)$$

Taking $\varpi \rightarrow J \lrcorner \varpi$ in the above, we also obtain

$$\begin{aligned} J[X, Y] \lrcorner \varpi &= [X, Y] \lrcorner (J \lrcorner \varpi) \\ &= X \lrcorner d(JY \lrcorner \varpi) - Y \lrcorner d(JX \lrcorner \varpi) - Y \lrcorner X \lrcorner d(J \lrcorner \varpi) + d(X \lrcorner JY \lrcorner \varpi). \end{aligned} \quad (5.0.11)$$

Substituting (5.0.10) and (5.0.11) into (5.0.8) and recalling that $J^2 = \pm 1$ when J is almost (para-)complex, we find

$$N_J(X, Y) \lrcorner \varpi = \varepsilon Y \lrcorner X \lrcorner d\varpi + JY \lrcorner JX \lrcorner d\varpi - Y \lrcorner JX \lrcorner d(J \lrcorner \varpi) - JY \lrcorner X \lrcorner d(J \lrcorner \varpi). \quad (5.0.12)$$

Finally, when $d\varpi = 0$ and $d(J \lrcorner \varpi) = 0$, we have $N_J(X, Y) \lrcorner \varpi = 0$ for all $X, Y \in TN$. Since ϖ is non-degenerate, this can only be true if $N_J \equiv 0$, that is, if J is integrable. \square

Theorem 5.0.5 (Partial Higher Lychagin–Rubtsov Theorem)

Let N be a smooth manifold with $\dim(N) = 2m$, $\varpi \in \Omega^m(N)$ a closed, non-degenerate m -form on N , and J an almost (para-)complex structure on TN , such that (ϖ, J) are compatible. Then, J is integrable if and only if $J \lrcorner \varpi$ is closed.

Proof. By Proposition 5.0.4, J is integrable if $J \lrcorner \varpi$ is closed, so it remains to check the converse. Let J be a complex structure and define

$$\Omega := J \lrcorner \varpi + \mathbf{i}\varpi, \quad (5.0.13)$$

where $\mathbf{i} := \sqrt{-1}$. Since ϖ and J are both non-degenerate on TN , we know that Ω is non-degenerate on $T^{1,0}N$. Furthermore, as $J \lrcorner \Omega = \mathbf{i}\Omega$, it follows from compatibility of (ϖ, J) that Ω is an $(m, 0)$ -form. Denoting the holomorphic and anti-holomorphic derivatives with respect to J by ∂ and $\bar{\partial}$ respectively, we find $\partial\Omega = 0$ and

$$d(J \lrcorner \varpi) = d\Omega = \bar{\partial}\Omega, \quad (5.0.14)$$

where we use that $d = \partial + \bar{\partial}$ and $d\varpi = 0$. Therefore, $J \lrcorner \varpi$ is closed if and only if $\bar{\partial}\Omega = 0$.

Since Ω is a non-degenerate $(m, 0)$ -form on $T^{1,0}N$, it defines a smooth global section of the canonical bundle $\det(T^{1,0}N)$, which is consequently trivial as a smooth bundle. This implies that global sections of $\det(T^{1,0}N)$ only differ by overall factors which are smooth complex functions on N . As J is an integrable complex structure, by the Newlander–Nirenberg theorem [47, 48], there exists charts $\{U_a\}$ covering N , local holomorphic complex coordinates (z_a^1, \dots, z_a^m) on U_a , and a global section Ω_0 of $\det(T^{1,0}N)$ such that $\Omega_0|_{U_a} = dz_a^1 \wedge \dots \wedge dz_a^m$ for each a . In particular, $\Omega = f\Omega_0$ for some $f \in \mathcal{C}^\infty(M, \mathbb{C})$. Therefore,

$$d(J \lrcorner \varpi) = \bar{\partial}\Omega = (\bar{\partial}f) \wedge \Omega_0. \quad (5.0.15)$$

Furthermore, as ϖ is closed and J is integrable, (5.0.12) becomes

$$N_J(X, Y) \lrcorner \varpi = JX \lrcorner Y \lrcorner d(J \lrcorner \varpi) + X \lrcorner JY \lrcorner d(J \lrcorner \varpi) = 0. \quad (5.0.16)$$

Evaluating (5.0.16) on holomorphic vector fields X and Y , which satisfy $J(X) = \mathbf{i}X$ and $J(Y) = \mathbf{i}Y$, we find

$$X \lrcorner Y \lrcorner ((\bar{\partial}f) \wedge \Omega_0) = X \lrcorner Y \lrcorner d(J \lrcorner \varpi) = 0, \quad (5.0.17)$$

however, as $\bar{\partial}f$ is a $(0,1)$ -form, $X \lrcorner \bar{\partial}f = 0 = Y \lrcorner \bar{\partial}f$, for X and Y holomorphic. Hence,

$$(\bar{\partial}f) \wedge X \lrcorner Y \lrcorner \Omega_0 = 0. \quad (5.0.18)$$

As Ω_0 is a non-degenerate $(m,0)$ -form on $T^{1,0}N$, it follows that $\bar{\partial}f = 0$. Plugging this back into (5.0.15) implies $d(J \lrcorner \varpi) = 0$, as required.

The proof follows similarly in the case where J is para-complex, provided we use the adapted para-holomorphic coordinates of [50, Section 2.1]. \square

Setting $N = T^*M$ for some smooth manifold M of dimension m , such that linearly independent pairs of differential forms (ϖ, α) can be interpreted as equations of the form (5.0.2), we see that Theorem 5.0.5 is a generalisation of part of the Lychagin–Rubtsov theorem, as stated in Theorem 2.2.6, to dimension greater than two. That is, we have a partial answer to question (iii) from our list above. In particular, Theorem 5.0.5 relates the closure of a higher Monge–Ampère structure to the integrability of the almost (para-)complex structure J defined via (5.0.3), if such a structure exists. However, the partial theorem does not relate these conditions to the local (multisymplectomorphic) equivalence of the equations (5.0.2) and their (as yet unknown) canonical forms, akin to the Laplace and Wave equation in Theorem 2.2.6. In order to complete the generalisation of Theorem 2.2.6, we first need to fully resolve questions (i) and (ii), namely when the endomorphism A exists and is (proportional to) an almost (para-)complex structure — in the two dimensional case, this was given precisely by the effectiveness and yielded the Pfaffian as the proportionality factor.

Let us specify to the case when N is six-dimensional (so that M is three-dimensional when $N = T^*M$) and consider pairs of 3-forms (ϖ, α) on N . In light of Proposition 5.0.3 and as our eventual goal is to obtain almost (para-)complex structures, which are isomorphisms, we restrict ourselves to the case when ϖ and α are both non-degenerate. In [97, Appendix A], Bryant kindly provides us with the normal forms of the non-degenerate differential 3-forms in six dimensions.³ More precisely, given a non-degenerate 3-form $\varpi \in \Omega^3(N)$, there exists (locally) a basis $\{e^i\}_{i=1}^6$ of T^*N such that

$$\begin{aligned} \eta_- &:= e^{123} + e^{456} & \text{HPf}(\eta_-) &< 0, \\ \eta_+ &:= e^{136} + e^{426} + e^{235} + e^{145} & \text{HPf}(\eta_+) &> 0, \\ \eta_0 &:= e^{135} + e^{416} + e^{326} & \text{HPf}(\eta_0) &= 0, \end{aligned} \quad (5.0.19)$$

³Recall that in Definition 2.2.9, we take $(J_\alpha^{\text{vol}})^2 = -\text{HPf}(\alpha)\mathbb{I}$, so our Hitchin Pfaffian has opposite sign to [97]. It is also worth noting that Banos [98] simplifies these normal forms further under the assumptions that N admits a symplectic form ω , as it does when $N = T^*M$, and that the 3-forms are ω -effective. Since our 3-forms need not be ω -effective, we use the more general formulation here.

where $e^{ijk} := e^i \wedge e^j \wedge e^k$. Furthermore, $\eta_{\pm,0}$ form disjoint $\text{GL}(TN)$ -orbits, which are distinguished by the sign of the Hitchin Pfaffian of the 3-form, i.e. the sign of the Hitchin Pfaffian is invariant under changes of basis on TN (hence on T^*N). Using these normal forms, we have the following result:

Proposition 5.0.6 (Normal Forms and Isomorphisms)

Let N be a smooth manifold, (ϖ, A) a compatible non-degenerate differential 3-form and isomorphism pair on N , and $\alpha := A \lrcorner \varpi$. Assume that $A \neq \mathbb{I}$. Then α and ϖ cannot be simultaneously brought into normal form.

Proof. First, consider the case where ϖ and α have the same normal form. Then, in the basis $\{e^i\}_{i=1}^6$, we have $\alpha = A \lrcorner \varpi = \varpi$. Since ϖ is non-degenerate, this implies $A = \mathbb{I}$ and we have a contradiction.

Let $\{E_i\}_{i=1}^6$ denote the basis on TN dual to $\{e^i\}_{i=1}^6$. Consider the case where $\varpi = \eta_+$ and $\alpha = \eta_0$ in normal form and assume that they take this form simultaneously. Then we have

$$(\eta_0)_{ijk} = \alpha(E_i, E_j, E_k) = \varpi(A(E_i), E_j, E_k) = \varpi(A_i^d E_d, E_j, E_k) = A_i^d (\eta_+)_{dj k}, \quad (5.0.20)$$

which implies that

$$-1 = (\eta_0)_{236} = A_2^d (\eta_+)_{d36} = A_2^1 = A_2^d (\eta_+)_{d45} = (\eta_0)_{245} = 0, \quad (5.0.21)$$

which is a contradiction. Hence, α does not have the form η_0 or η_+ when ϖ has the form η_+ .

Next we check the case where $\varpi = \eta_+$ and $\alpha = \eta_-$ when simultaneously in normal form. As above, we find

$$(\eta_-)_{ijk} = \alpha(E_i, E_j, E_k) = \varpi(A_i^d E_d, E_j, E_k) = A_i^d (\eta_+)_{djk}, \quad (5.0.22)$$

which implies that

$$1 = (\eta_-)_{312} = A_3^d (\eta_+)_{d12} = 0, \quad (5.0.23)$$

and again we have a contradiction. Consequently, if ϖ has normal form η_+ , then ϖ and α cannot be brought into their normal forms simultaneously (unless $A = \mathbb{I}$).

As A is an isomorphism, we can write $\varpi = A^{-1} \lrcorner \alpha$ and α is a non-degenerate 3-form. Consequently, the above cases hold with the normal forms of ϖ and α swapped. Hence, it remains to show the case where $\varpi = \eta_-$ and $\alpha = \eta_0$ when simultaneously in normal form. As

above, we find

$$(\eta_0)_{ijk} = \alpha(E_i, E_j, E_k) = \varpi(A_i^d E_d, E_j, E_k) = A_i^d (\eta_-)_{dj k}, \quad (5.0.24)$$

which implies that

$$1 = (\eta_0)_{135} = A_1^d (\eta_-)_{d35} = 0, \quad (5.0.25)$$

and we have a contradiction for a third and final time. \square

Consequently, if a pair of non-degenerate differential 3-forms (ϖ, α) , such that $\varpi \neq \alpha$, can be brought into normal form simultaneously, then those forms do not define an isomorphism (or in fact an endomorphism, by Proposition 5.0.3). In particular, this includes the differential forms (4.2.1) describing our incompressible fluid flows in three dimensions, which is precisely why we used the Hitchin endomorphism of α directly in Section 4.2.1. Indeed, ϖ has normal form η_0 , while the choice of normal form for α depends on the sign of $\hat{f} = \text{HPf}(\alpha)$, and (ϖ, α) can be brought into their normal forms simultaneously. For example, when $\hat{f} > 0$ and $N = T^*\mathbb{R}^3$, we can take

$$\begin{aligned} e^1 &= -\hat{f}^{-\frac{1}{3}} dx^1, & e^2 &= \hat{f}^{-\frac{5}{6}} dq_1, & e^3 &= -\hat{f}^{\frac{2}{3}} dx^2, \\ e^4 &= \hat{f}^{\frac{1}{6}} dq_2, & e^5 &= -\hat{f}^{\frac{1}{6}} dq_3, & e^6 &= -\hat{f}^{\frac{2}{3}} dx^3, \end{aligned} \quad (5.0.26)$$

with analogous choices when $N = T^*M$ for some smooth Riemannian manifold M or when $\hat{f} < 0$. We previously discarded the case when $\hat{f} = 0$, since the almost (para-)complex structure given by the Hitchin endomorphism of α became degenerate, however this should be considered more closely in future work, since both ϖ and α have the same normal form when $\hat{f} = 0$.

Returning to the general setting of pairs of non-degenerate 3-forms (ϖ, α) on a six-dimensional manifold N , we can take the idea of Proposition 5.0.6 a bit further and consider what happens ϖ and α have different normal forms, but cannot necessarily be brought into these forms simultaneously.

Proposition 5.0.7 (Normal Forms and Almost (Para-)Complex Structures)

Let N be a smooth manifold, (ϖ, A) a compatible non-degenerate differential 3-form and isomorphism pair on N , and $\alpha := A \lrcorner \varpi$. Assume that A is proportional to an almost (para-)complex structure, i.e. $A^2 = \kappa \mathbb{I}$ for some $\kappa \in \mathcal{C}^\infty(N)$ with $\kappa \neq 0$. Then α has the same normal form as ϖ .

Proof. As the normal form of a 3-form in six dimensions is distinguished by the sign of the Hitchin Pfaffian of said form, we must check that $\text{sgn}(\text{HPf}(\alpha)) = \text{sgn}(\text{HPf}(\varpi))$. We begin by

noting that

$$\alpha = J \lrcorner \varpi = \varpi(A \cdot, \cdot, \cdot) = \frac{1}{\kappa} \varpi(A^3 \cdot, \cdot, \cdot) = \frac{1}{\kappa} \varpi(A \cdot, A \cdot, A \cdot), \quad (5.0.27)$$

where in the last step we use the compatibility of ϖ and A . Recall from [Definition 2.2.9](#) that the Hitchin Pfaffian scales as $\text{HPf}(\frac{1}{\kappa}\alpha) = \frac{1}{\kappa^4} \text{HPf}(\alpha)$, so from the above we obtain

$$\text{HPf}(\alpha) = \frac{1}{\kappa^4} \text{HPf}(\varpi(A \cdot, A \cdot, A \cdot)). \quad (5.0.28)$$

As A is an isomorphism on TN , $\kappa^4 > 0$, and the sign of the Hitchin Pfaffian is $\text{GL}(TN)$ -invariant, it then follows that

$$\text{sgn}(\text{HPf}(\alpha)) = \text{sgn}(\text{HPf}(\varpi)), \quad (5.0.29)$$

so ϖ and α have the same normal form, as required. \square

Consequently, if a pair of differential 3-forms on a six-dimensional manifold have different normal forms [\(5.0.19\)](#), they cannot define an endomorphism which is proportional to an almost (para-)complex structure via [\(5.0.3\)](#). Furthermore, it is shown in [\[36, Proposition 2.7\]](#) that every almost (para-)complex structure on N , and therefore every endomorphism proportional to an almost (para-)complex structure on N , admits a pair of compatible non-degenerate 3-forms (ϖ, α) , which satisfy $\alpha = A \lrcorner \varpi^4$ and we now know that the 3-forms in this pair must have the same normal form. However, the converse problem of finding sufficient conditions for pairs of differential forms (ϖ, α) to define an almost (para-)complex structure is still open. Similarly worth investigating is whether or not we may obtain all almost (para-)complex structures from [\(5.0.3\)](#) when restricting to just the class of differential forms corresponding to a given normal form and whether complex structure deformations respect these classes.

In summary, we have presented first steps towards answering the questions proposed at the start of this outlook, by providing a series of necessary conditions on pairs of differential 3-forms (ϖ, α) in six dimensions (and more generally differential m -forms in $2m$ -dimensions) for the endomorphism A in [\(5.0.3\)](#) to possess certain desirable properties, if said endomorphism exists. These properties include compatibility, non-degeneracy, being an almost (para-)complex structure, and being integrable. These results rule out a large number of pairs of differential forms from defining almost (para-)complex structures (and indeed endomorphisms) via [\(5.0.3\)](#) and allowed us to derive a partial generalisation of the Lychagin–Rubtsov theorem in [Theorem 5.0.5](#).

⁴In fact, they show that there exists a pair of compatible non-degenerate m -forms (ϖ, α) which satisfy $\alpha = A \lrcorner \varpi$ on any $2m$ -dimensional almost (para-)complex manifold N . It is also possible to infer from the proof of [\[36, Proposition 2.7\]](#) that $\alpha \wedge \varpi = 0$ if and only if m is even, otherwise $\alpha \wedge \varpi$ is a volume form.

Moving forward, we aim to find not only necessary conditions for a pair of differential forms to define an endomorphism/isomorphism/almost (para-)complex structure, but necessary *and* sufficient conditions for this to be the case. We also hope to investigate whether or not $A \lrcorner \varpi$ yields all differential forms with the same normal form as ϖ . While the direct approach of testing properties has been fruitful for ruling out pairs of differential forms with(out) certain properties, one line of inquiry we have not probed here is what happens to the Hitchin Pfaffian when taking linear combinations of differential 3-forms in six dimensions. In particular, given (ϖ, α) such that ϖ and α have different normal forms, when does $\varpi + f\alpha$ have the same normal form as α ? Since (ϖ, α) and $(\varpi + f\alpha, \alpha)$ yield equivalent systems of equations (5.0.2) for any $f \in \mathcal{C}^\infty(N)$, this naturally leads us to the question of whether or not there is a multisymplectic generalisation of the Hodge–Lepage–Lychagin theorem, cf. Theorem 2.1.8, for removing such redundant pairs of differential forms. Such a theorem would enable us to identify a “canonical” pair of differential forms for a given system of two “higher Monge–Ampère equations” (5.0.2) as the pair satisfying some higher notion of effectiveness. As the standard notion of effectiveness in two dimensions yields the equation (2.2.14), which describes when an endomorphism A obtained from (5.0.3) squares to a multiple of the identity operator, it is not too far-fetched to expect a similarly useful identity might be obtained from its multisymplectic generalisation.

In the context of fluid dynamics, the end goal is to obtain (if possible, or prove otherwise) a pair of differential forms (ϖ, α) , which not only encode the divergence-free constraint and Poisson equation for the pressure (3.1.9) in dimension greater than two, as the forms (4.2.1) do, but also define an almost (para-)complex structure via (5.0.3), as in two dimensions. The resulting almost (para-)complex structure would be a more natural generalisation of the two-dimensional case than the Hitchin endomorphism (4.2.14) of α alone, since it would reflect changes in both of the defining equations. Hence, it would be interesting to compare the geometry resulting from the two choices. Furthermore, any progress towards Hodge–Lepage–Lychagin or Lychagin–Rubtsov style theorems could be used to study the equations accordingly.

Part II

Globalisation of Curvature Bounds in
Lorentzian Length Spaces



6

Overview of Part II

Given a smooth manifold M , if we wish to consider the lengths of curves and the angles between them as we would in Euclidean space, we equip M with a Riemannian metric g , making the pair (M, g) a Riemannian manifold. The length of a curve $\gamma : [a, b] \rightarrow M$, whose derivative is defined almost everywhere and is non-zero, is then given with respect to g by

$$L_g(\gamma) := \int_a^b \sqrt{g(\gamma'(t), \gamma'(t))} dt. \quad (6.0.1)$$

In fact, sufficiently smooth¹ Riemannian metrics have the tendency to sing like canaries and not only provide information about the curves on M , but also about M itself. As we saw in [Part I](#) of this thesis, if the metric g is at least $\mathcal{C}^2(M)$, we may write down the Ricci and Riemann curvature tensors, denoted R_{ij} and $R_{ijk}{}^l$ respectively, in terms of derivatives of the metric (see (3.1.10)), with the differentiability of g being crucial to the definition. We may also define the so-called sectional curvature

$$S(u, v) := \frac{g(R(u, v)v, u)}{g(u, u)g(v, v) - g(u, v)^2} \quad (6.0.2)$$

of the tangent plane at $p \in M$, spanned by vectors $u, v \in T_pM$. Here, $R(u, v)w = u^i v^j w^k R_{ijk}{}^l \partial_l$ is given by the Riemann curvature tensor at $p \in M$, for $u, v, w \in T_pM$.

While Riemannian manifolds have predominantly been studied using this analytical, differential approach to geometry since their introduction by Riemann in 1853 [99], it turns out that many of their properties arise as a consequence of the underlying metric space structure given

¹Recall that a metric g is called smooth if it varies smoothly with $p \in M$, that is, if its components g_{ij} are smooth functions on M .

by the distance function

$$d_g(x, y) = \inf\{L_g(\gamma) \mid \gamma \text{ is an admissible curve from } x \text{ to } y \text{ in } M\}, \quad (6.0.3)$$

between each pair of points x, y in M . In fact, studying Riemannian manifolds via this complementary, synthetic description, using the incidence of points and lines rather than relying on coordinates and differentiability, allows us to abstract many of the fundamental properties of smooth manifolds to structures of lower regularity. Indeed, given an arbitrary metric space (X, d) , X need only be a set as opposed to a smooth Riemannian manifold, while the distance function d need only be continuous in the metric topology it defines.²

One such fundamental property that extends to metric spaces is the notion of curvature. While we can write down the curvature (be it Ricci, Riemann, or sectional) of a Riemannian manifold explicitly in terms of the associated Riemannian metric, for arbitrary metric spaces, which may not have a Riemannian metric, we instead describe the curvature by bounding the values it may take. A metric space X is said to have curvature bounded above/below by $k \in \mathbb{R}$ if the triangles in a neighbourhood of each point satisfy some “*comparison condition*” with respect to triangles with the same side-lengths in the Riemannian manifold of constant (sectional) curvature k . We shall make precise what we mean by a comparison condition in [Definition 7.1.12](#), however, note that there are several different choices we could make which do not require any smooth structure and which are shown to be equivalent in [\[100, Theorem 8.30\]](#). Observe that our comparison conditions are applied locally, to small triangles in a neighbourhood of each point, rather than on the entire space — if a comparison condition holds on the whole of X , that is, for arbitrarily large triangles, we say that X has *global* curvature bounded above/below. The pertinent question concerning the relationship between local and global curvature bounds was answered for metric spaces with curvature bounded above by Alexandrov [\[101\]](#), who showed that the existence of unique distance realisers which vary continuously with their end-points is necessary and sufficient for the two notions to coincide for fixed k . Building on the work of Toponogov [\[102\]](#), who answered the question for Riemannian manifolds with curvature bounded below, Perelman [\[103\]](#) demonstrated that all complete metric length spaces with local curvature bounded below by k have the same bound globally.

Spurred onward by the work of Gromov [\[104, 105\]](#), metric spaces with curvature bound have seen a wide range of applications which are not confined to geometry. On one hand, it has been shown that global lower curvature bounds on metric spaces are stable under Gromov–Hausdorff convergence [\[103, 106\]](#), hence the limit of a sequence of Riemannian manifolds with curvature

²Given that M is used to denote manifolds throughout this thesis, we have chosen to denote metric spaces by X . Although X previously referred to a tangent vector, we shall not need this meaning again moving forward.

bounded below by k is a metric space with the same bound [107], while on the other hand, it is known that any complete metric length space with curvature bounded below by zero is isometric to the quotient space formed by the action of some crystallographic group on \mathbb{R}^n , see [108]. Furthermore, at the intersection of group theory and algebraic topology, we have the statement that the fundamental group of any complete metric space with curvature bounded above by zero has no non-trivial finite subgroups [109, Corollary 9.3.2.]. Other applications include the study of gradient flows and optimal transport [110–112].

Meanwhile, in the realm of Lorentzian geometry, the necessity of low regularity geometry has again made itself known, this time in the study of Einstein’s equations of general relativity

$$R_{\mu\nu} + \left(\Lambda - \frac{1}{2}R \right) g_{\mu\nu} = CT_{\mu\nu}, \quad (6.0.4)$$

where Λ is the cosmological constant, C is the gravitational constant, and $T_{\mu\nu}$ is the energy momentum tensor. Indeed, while we usually assume a priori that a metric $g_{\mu\nu}$ is $\mathcal{C}^2(M)$ when writing down its Ricci curvature tensor $R_{\mu\nu}$ and scalar R , the vacuum Einstein equations (where $\Lambda, T_{\mu\nu} = 0$) are known to admit solutions of Sobolev regularity H_{loc}^s , if the induced Riemannian metric on spacelike slices is H_{loc}^s and $s > \frac{5}{2}$, see [113]. Low regularity solutions to (6.0.4) also arise in models of (conical) cosmic strings, which possess distributional energy momentum tensor $T_{\mu\nu}$ and are described by quasi-regular singularities [114]. Another physically relevant phenomenon described by low regularity geometry is the impulsive gravitational wave, which may be described by a continuous or distributional metric [115]. Since the Lorentzian metric in (6.0.4) may be continuous, distributional, or singular in these settings, the geometry they describe is not a (smooth) Lorentzian manifold.

It was with this problem in mind that Kunzinger and Sämann [116] introduced the concept of the Lorentzian pre-length space as a Lorentzian analogue of the metric space. By building on the causal space defined in [117], these Lorentzian pre-length spaces admit a causality theory much like a Lorentzian manifold, yet by mirroring the structure of a metric space, the time separation between points may be defined without the need for a smooth metric. Furthermore, the triangle comparison conditions developed by [118] for timelike triangles in Lorentzian manifolds were also extended to the synthetic setting,³ allowing us to bound the curvature of a Lorentzian pre-length space. Consequently, since its inception in 2018, the study of Lorentzian pre-length spaces has seen rapid growth, akin to that of the metric spaces which came before it. Much of this progress has been motivated by the application to general relativity discussed

³It should be noted that the comparison methods of [119] for triangles of arbitrary causal character in a semi-Riemannian manifold do not extend quite so completely, since the Lorentzian pre-length space does not have a natural way of measuring the length of a spacelike curve.

above, with [120] relating inextendibility of a Lorentzian pre-length space to the existence of a curvature singularity (i.e. the curvature of the space cannot be bounded), while [121] prove singularity theorems for Lorentzian pre-length spaces given by warped products, which underlie the Friedmann–Lemaître–Robertson–Walker model of cosmology [122, Chapter 12].

However, Lorentzian pre-length spaces are also interesting from a purely geometric perspective, as Lorentzian analogues of metric spaces and it is not outlandish to expect that they would admit similarly deep geometry. Consequently, a menagerie of additional comparison conditions have been introduced, based around the notions of hyperbolic angles, quadrilaterals, and convexity [123–125], and the curvature bounds associated with these conditions are understood to behave like their metric counterparts. This understanding has been crucial for deriving synthetic Lorentzian versions of many significant results from metric geometry, including the Reshetnyak gluing theorem [126, 127], the splitting theorem [128], and the notion of Hausdorff measure [129]. Nevertheless, none of these earlier works address the relationship between local and global curvature bounds on Lorentzian pre-length spaces. In this part of the thesis, we take up this problem and provide Lorentzian analogues of both Alexandrov’s patchwork and the Toponogov globalisation theorem, in addition to a number of smaller results, including the Lebesgue number lemma and Bonnet–Myers theorem. The goal of this programme of Lorentzification is to widen the appeal of the Lorentzian pre-length space framework and facilitate its application to problems outside of general relativity, including those concerning the convergence of Lorentzian manifolds, causal set theory, and group theory.

Organisation of Part II An outline of this part of the thesis, which is based on the works [2, 3], is as follows. We begin in [Chapter 7](#) with a brief review of the theory of metric spaces, building from the notion of lengths of curves and distance realisers in [Section 7.1](#), to globalisation theorems for spaces with curvature bounds in [Section 7.2](#). In particular, we provide statements of Alexandrov’s patchwork globalisation theorem for spaces with curvature bounded above and the Toponogov globalisation theorem for spaces with curvature bounded below. While the content in this chapter is by no means novel, by starting in the metric setting, we are able to introduce the essential tools of synthetic geometry without the additional complications of Lorentzian artefacts, that is, in a way which may be understood physically, with curves drawn on a piece of paper or the surface of a ball, for example. Furthermore, it allows us to highlight the subtle differences between the metric space and Lorentzian pre-length space constructions, which result in additional assumptions being required to derive the globalisation theorems in later chapters. These additional assumptions primarily revolve around the notions of topology, continuity, and finiteness.

In [Chapter 8](#), we provide an overview of the properties of Lorentzian pre-length spaces, based on the seminal work [\[116\]](#) and subsequent research in [\[123, 125\]](#). In addition to defining appropriate notions for length and angle, it will also be necessary to discuss the causality properties Lorentzian pre-length spaces may exhibit, such as strong causality, which guarantees that the Alexandrov topology defined by timelike diamonds coincides with the topology defined by the underlying metric space structure. We subsequently introduce a new, generalised notion of curvature bounds via triangle comparison, based on the original definition provided by [\[116, Definition 4.7\]](#), but with some modifications necessary for the model space with constant negative sectional curvature, anti-de Sitter space, to be considered a space with curvature bounded above and below. We also provide conditions under which the two notions of curvature bounds coincide. While the main purpose of this chapter is to allow for the thesis to be relatively self contained, along the way we prove refinements of several other results from both smooth and synthetic Lorentzian geometry. These are signposted as appropriate. In the second half of this chapter, [Section 8.2](#), we review several basic yet crucial properties of hyperbolic angles and how they can be used to describe curvature bounds, concluding with a statement of the conditions under which the curvature bounds defined by angle, hinge, and triangle comparison are equivalent.

The first substantial new results appear in [Chapter 9](#), in which we prove an analogue of Alexandrov's patchwork globalisation for Lorentzian pre-length spaces with curvature bounded above. After reiterating the statement of the gluing lemma for Lorentzian pre-length spaces provided by [\[126, Corollary 4.3.2\]](#), we define what it means for timelike distance realisers to vary continuously with their endpoints and demonstrate that this may be encoded by the so-called geodesic map. The bulk of the chapter is dedicated to the proof of the main result: using the geodesic map to decompose an arbitrary timelike triangle, in a (suitably nice) Lorentzian pre-length space with local curvature bounded above, into sub-triangles which are all timelike and satisfy the triangle comparison condition, such that the gluing lemma may be used to deduce that the initial triangle also satisfied said condition. The key difficulty here is ensuring that all of the curves in our decomposition remain timelike, rather than causal or spacelike, so that the gluing lemma can be applied, in contrast to the metric setting where all curves are fair game. The chapter ends with the observation that in Lorentzian pre-length spaces with global curvature bounded above by K , distance realisers which can be realised in the model space of constant curvature K are unique.

Following on from our proof of the globalisation theorem for Lorentzian pre-length spaces with curvature bounded above, in [Chapter 10](#) we prove an analogue of the Toponogov globalisation theorem for Lorentzian pre-length spaces with curvature bounded below. The chapter

opens with a series of supplementary results concerning time functions and the null distance, streamlining those proven in [130–132], and building towards a Lorentzian analogue of the Lebesgue number lemma for the null distance. We subsequently derive a decomposition result for triangles in Lorentzian pre-length spaces with curvature bounded below, in the spirit of the gluing lemma for curvature bounded above. Finally, by performing a construction derived from the “cat’s cradle” of Lang and Schroeder [133] in the metric setting, combined with our decomposition result, we can recursively show that triangles of increasingly large null distance diameter satisfy the angle comparison condition. The Lebesgue number lemma ensures that this process terminates and does so once all triangles satisfy said condition. While our pair of globalisation theorems respectively consider the triangle and angle comparison conditions, the globalisation property passes to a wide array of equivalent comparison conditions (detailed in [123, Theorem 5.1]), as in the metric setting. We close the chapter by considering some geometric consequences of the Toponogov globalisation theorem, beginning with a Lorentzian analogue of the Bonnet–Myers theorem bounding the finite diameter of a Lorentzian pre-length space. We then provide conditions under which the diameter of a Lorentzian pre-length space is not attained, which in turn enables us to show that, in Lorentzian pre-length spaces which satisfy the assumptions of the Bonnet–Myers theorem, all triangles satisfy size-bounds (and therefore admit comparison).

We conclude in [Chapter 11](#) by presenting an overview of some open problems, including how the scope of some of our results might be generalised. We take particular interest in the relationship between curvature bounds and Gromov–Hausdorff convergence; while several notions of Gromov–Hausdorff convergence for Lorentzian spaces have been proposed [132, 134, 135], we expect that curvature bounds should remain stable under those which are sensible, as in the metric setting. In a similar vein, it is known that compact length spaces are given by the Gromov–Hausdorff limit of polyhedral spaces, or more precisely, graphs. The analogous Lorentzian question would be whether or not we can discretise Lorentzian pre-length spaces and describe them as the limit of causal sets. We shall also discuss the impact of curvature bounds on graphs and causal sets in some detail.



7

A Brief Review of Metric Geometry

In this chapter, we provide a brief review metric geometry, beginning with some core concepts, including intrinsic metrics and distance realisers, before moving on to a more detailed description of triangle comparison and curvature bounds. We close this chapter by taking a closer look at the globalisation theorems for spaces with curvature bounded above/below by $k \in \mathbb{R}$. While we provide statements of all relevant results, we do not provide details of their proofs. Instead, we shall direct the reader to appropriate references and make comparisons with the Lorentzian case in due course. For a more extensive discussion of metric geometry as a whole, we direct the reader to the excellent textbooks [109, 136] and the more recent review [100].

Although this part of the thesis is primarily focused on proving the globalisation of curvature bounds in the setting of Lorentzian pre-length spaces, we feel that there is value in providing an overview of the metric problem, where the statements of these results are known. In particular, by studying our problem’s “Riemannian signature” counterpart, we are able to get a feel for the structure of the results we wish to emulate in the synthetic Lorentzian world and highlight the obstacles that arise as a consequence of this change in “signature” in later chapters. Indeed, even before the end of Chapter 8, we shall experience issues caused by time separation in Lorentzian pre-length spaces not needing to be finite, in contrast to distances in metric spaces.

7.1 Fundamentals of Metric Geometry

In this section, we present some fundamental definitions from the setting of metric geometry, so that we may subsequently highlight the key differences between these and their synthetic Lorentzian analogues. Most significantly, we introduce the notion of the length of a curve, the

concept of distance realisers, and the technique of triangle comparison. Much of this discussion can be found in abridged form in [2, Section 2.2] and in more detail in the aforementioned textbooks [109, 136, 100].

Firstly, let us remind the reader of the conditions under which a set X is a metric space:

Definition 7.1.1 (Metric Space)

The pair (X, d) , where X is a set and $d : X \times X \rightarrow \mathbb{R}$ is a function on X , is called a metric space if, for all points $x, y, z \in X$, we have

(i) *Positive definite:* $d(x, y) > 0$ for all $x \neq y$ and $d(x, x) = 0$.

(ii) *Symmetry:* $d(x, y) = d(y, x)$.

(iii) *Triangle inequality:* $d(x, z) \leq d(x, y) + d(y, z)$.

The function d is then called a distance function on X , with $d(x, y)$ the distance between the points x and y in X .

As mentioned in the introduction, it is straightforward to show that the distance function is continuous in the metric topology. Furthermore, we may assume without loss of generality that distance functions are finite valued, as any metric space with $[0, \infty]$ -valued distance function can be decomposed into components each with finite valued distance function, see [100, Section 2.A]. We shall also need the notion of a metric space's diameter:

Definition 7.1.2 (Diameter)

The diameter of a metric space (X, d) is given by

$$\text{diam}(X) := \sup\{d(x, y) \mid x, y \in X\}. \quad (7.1.1)$$

That is, points in X may be at most a distance $\text{diam}(X)$ apart, though $\text{diam}(X)$ may be infinite.

In what follows, we call any connected subset $I \subseteq \mathbb{R}$ an interval. In particular, an interval may be (half-)closed, (half-)open, finite, infinite, or a single point. Let (X, d) be a metric space. The technique of triangle comparison will require us to construct triangles from curves in X and compare them to triangles in some well understood model spaces. Taking $I \subseteq \mathbb{R}$ to be an interval, a curve γ in X is a map $\gamma : I \rightarrow X$ which is continuous with respect to the distance function d on X .¹ A curve $\gamma : [a, b] \rightarrow X$ from the point x to the point y , which we may denote by γ_{xy} , additionally satisfies $\gamma(a) = x$ and $\gamma(b) = y$. The length of such a curve is defined in terms of the metric as follows:

¹Here the metric on \mathbb{R} is taken to be the standard Euclidean metric.

Definition 7.1.3 (Length)

Let (X, d) be a metric space. The d -length (or simply length where the metric is clear) of a curve $\gamma : [a, b] \rightarrow X$ is given by

$$L_d(\gamma) := \sup \left\{ \sum_{i=0}^{n-1} d(\gamma(t_i), \gamma(t_{i+1})) \mid a = t_0 < t_1 < \dots < t_n = b, n \in \mathbb{N} \right\}. \quad (7.1.2)$$

Curves γ with finite d -length are called d -rectifiable (or simply rectifiable). We call a metric space (X, d) a length space and the distance function d intrinsic if the metric is given by

$$d(x, y) := \inf \{ L_d(\gamma) \mid \gamma : [a, b] \rightarrow X, \gamma(a) = x, \gamma(b) = y \}, \quad (7.1.3)$$

for all pairs $x, y \in X$.

We may occasionally choose to restrict ourselves to considering only those curves from some stricter admissible class, such as the class of rectifiable curves (see [109, Chapter 2] for example). This choice is often more convenient when working with parametrised constructions, since every rectifiable curve admits a 1-Lipschitz parametrisation, namely the arclength parametrisation. Hence, if $\gamma : [a, b] \rightarrow X$ is a curve of finite length $L_d(\gamma) := L$, there exists a reparametrisation $\sigma : [0, L] \rightarrow [a, b]$ such that $\gamma \circ \sigma : [0, L] \rightarrow X$ and $L_d(\gamma|_{[\sigma(t), \sigma(s)]}) = |s - t|$ for all $s, t \in [0, L]$. Furthermore, we may scale σ such that $\gamma \circ \sigma$ is instead parametrised on $[0, 1]$ and $L_d(\gamma|_{[\sigma(t), \sigma(s)]}) = L|s - t|$ [137, Section 1.1.].




Figure 7.1.1: If we allowed arbitrary curves between points, both the left and right drawings would be triangles in Euclidean space with vertices x, y, z , even though γ_{yz} is far longer in the right drawing than the left one. As distance realisers in Euclidean space are straight lines, we are choosing to only admit the left case.

Since there may be multiple curves of differing length between each pair of points in X , defining a triangle to be a triple of points pairwise connected by *arbitrary* curves would result in triangles with different side-lengths but the same vertices, see Figure 7.1.1. If we wish to compare triangles in different spaces, we will therefore need to control the side-lengths of the triangles associated with each triple of points. We do so by considering only distance realising

curves; then, even if there are multiple triangles associated with a given triple of points, the triangles must have the same side-lengths.

Definition 7.1.4 (Shortest Curves and Distance Realisers)

Let (X, d) be a metric space and γ be a curve from x to y in X . The curve γ is called a shortest curve from x to y if, for all other curves $\tilde{\gamma}$ from x to y in X , we have $L_d(\tilde{\gamma}) \geq L_d(\gamma)$. If $L_d(\gamma) = d(x, y)$, then γ is called a distance realiser.

Applying the triangle inequality to the sum in (7.1.2) tells us that $L_d(\gamma) \geq d(\gamma(a), \gamma(b))$ for all curves γ . Hence, any distance realiser is automatically a shortest curve. Furthermore, in a length space, if a shortest curve exists then the infimum in (7.1.3) is attained by said curve, which is therefore a distance realiser. That is, in length spaces, shortest curves are precisely distance realisers. Since the metric d is finite valued, distance realisers are rectifiable and therefore we may assume that any distance realiser γ_{xy} is constant speed parametrised on $[0, 1]$, that is, $d(\gamma(s), \gamma(t)) = d(x, y)|s - t|$ for all $s, t \in [0, 1]$

Remark 7.1.5 (A Note on Terminology)

The reader may find that in the wider literature, the concepts of shortest curves and distance realisers are each referred to as geodesics. We choose not to use this term here, in order to avoid confusion with the notion of geodesics on a smooth manifold, since these need only be locally shortest curves. Similarly, we refer to d as a distance function and not a metric in order to avoid confusion with the (pseudo-)Riemannian metrics on smooth manifolds we discussed in Part I and Chapter 6.

△

Let us briefly address the questions of existence and uniqueness of distance realisers in metric spaces, beginning with the latter. In Definition 7.1.4, we call γ a distance realiser and not *the* distance realiser. This is because distance realisers need not be unique, as shown by the following example:

Example 7.1.6 (Distance Realisers on the Circle)

Consider the unit circle $X := \mathbb{S}^1$. Equip X with the intrinsic distance function d , where the distance between each pair of points is given by the (Euclidean) length of the shorter of the two arcs between them. By definition, the shorter arcs are distance realisers² and X is a length space. Now consider any pair of antipodal points in X . Both of the arcs between them are of equal length and are therefore both distance realisers.

²The length of curves in X with respect to d is precisely their Euclidean arclength. Furthermore, the longer of the two arcs is only locally a shortest curve (a geodesic).

△

Concerning the existence of distance realisers, it is also fairly straightforward to construct examples of metric spaces which do not admit distance realisers between some pairs of points, as shown by the following example. A similar construction can be found in [138, Chapter 2].

Example 7.1.7 (Distance Realisers on the Punctured Euclidean Plane)

Recall that distance realisers in the Euclidean plane are uniquely given by straight lines between their endpoints. Take any two points whose unique distance realiser in the Euclidean plane is a straight line through the origin, e.g. $p = (-1, 0)$ and $q = (1, 0)$. Now consider the punctured Euclidean plane $X = \mathbb{R}^2 \setminus \{(0, 0)\}$, equipped with the natural restriction of the Euclidean distance function, and assume there exists a distance realiser between p and q which is contained in X . Since this distance realiser cannot pass through the origin, we now have two distance realisers in the Euclidean plane which connect these points, contradicting uniqueness. Hence, there is no distance realiser between p and q in X .

△

Consequently, we now define the class of *geodesic spaces* — spaces for which distance realisers always exist between pairs of points. Restricting our consideration to such spaces guarantees that, for any triple of points, we may construct a triangle with said points as vertices and with sides given by distance realisers. We have already seen one such a space in Example 7.1.6, which highlights that the distance realisers may still not be unique.

Definition 7.1.8 (Geodesic Spaces)

A metric space (X, d) is called an ℓ -geodesic space if there exists a distance realiser between every pair of points $x, y \in X$ for which $d(x, y) < \ell$, that is, if the infimum in (7.1.3) exists for such pairs of points. If a distance realiser exists between every pair of points in X , then we call X a geodesic space.

The metric space (X, d) will be called uniquely ℓ -geodesic (resp. uniquely geodesic) if there exists a unique distance realiser between pairs of points $x, y \in X$ for which $d(x, y) < \ell$ (resp. between every pair of points in X).

Now that we understand what it means for a curve in a metric space to be a distance realiser, we may define a triangle in a metric space (X, d) to be a collection of three points x, y, z in X (the vertices) and three pairwise connecting distance realisers $\gamma_{xy}, \gamma_{xz}, \gamma_{yz}$ (the sides) between them. We will denote such a triangle by $\Delta(x, y, z)$. As highlighted in Example 7.1.6, distance realisers between points in X need not be unique, so where the choice of distance realiser

is important, we shall write our selection explicitly. If not explicitly stated, constructions are valid for any choice of distance realiser and the reader is left to pick their favourite. We write $p \in \Delta(x, y, z)$ to indicate that a point $p \in X$ lies on the triangle and $p \in \gamma_{xy}$ to specify that p lies on the side between x and y , for example.

Next, we introduce the spaces we wish to compare against, namely the Riemannian model spaces of constant sectional curvature. Let $\mathbb{R}^{m,n}$ with $m, n \in \mathbb{N}_0$ denote the vector space \mathbb{R}^{m+n} equipped with the indefinite inner product

$$b^{m,n}(x, y) := - \sum_{i=1}^m x_i y_i + \sum_{i=1}^n x_{m+i} y_{m+i} \quad (7.1.4)$$

for $x = (x_1, x_2, \dots, x_{m+n})$ and $y = (y_1, y_2, \dots, y_{m+n})$. The Riemannian model spaces may be described in terms of quadric surfaces in suitable $\mathbb{R}^{m,n}$ — for more detail regarding the construction of these model spaces, including their uniqueness and choices of metric, see [136, Chapters I.2, I.6] and [122, Proposition 4.29].

Definition 7.1.9 (Riemannian Model Spaces)

The Riemannian model space of constant sectional curvature k is the unique (up to local isometry) complete, simply connected, two-dimensional Riemannian manifold with sectional curvature k , which we denote by $M^2(k)$. In particular,

$$M^2(k) := \begin{cases} \mathbb{S}^2(r) & k = \frac{1}{r^2} > 0 \\ \mathbb{R}^2 & k = 0 \\ \mathbb{H}^2(r) & k = -\frac{1}{r^2} < 0 \end{cases} . \quad (7.1.5)$$

Here, $\mathbb{S}^2(r) := \{x \in \mathbb{R}^{0,3} \mid b^{0,3}(x, x) = r^2\}$ is the 2-sphere with radius r and distance function induced by the restriction of $b^{0,3}$. Similarly, $\mathbb{H}^2(r) := \{x \in \mathbb{R}^{1,2} \mid b^{1,2}(x, x) = -r^2 \text{ and } x_1 > 0\}$ is the hyperbolic or Lobachevsky plane of radius r , as described by the upper sheet of a two-sheeted hyperboloid a distance³ r from the origin in $\mathbb{R}^{1,2}$. We equip $\mathbb{H}^2(r)$ with the distance function induced by the restriction of $b^{1,2}$ and \mathbb{R}^2 with its usual Euclidean distance function $b^{0,2}$.

The diameters (7.1.1) of the Riemannian model spaces are

$$\text{diam}(M^2(k)) = \begin{cases} \frac{\pi}{\sqrt{k}} & k > 0 \\ \infty & k \leq 0 \end{cases} . \quad (7.1.6)$$

³Strictly speaking, this “radius” is the fixed *proper time* between points on $\mathbb{H}^2(r)$ and the origin in $\mathbb{R}^{1,2}$.

Definition 7.1.10 (Comparison Triangles and Points)

Let (X, d) be a metric space and $M^2(k)$ be the model space of constant sectional curvature k , equipped with its natural distance function d_k . Let $\Delta(x, y, z)$ be a triangle in X . A triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in $M^2(k)$ is called a comparison triangle for $\Delta(x, y, z)$ if

$$d(x, y) = d_k(\bar{x}, \bar{y}), \quad d(y, z) = d_k(\bar{y}, \bar{z}), \quad \text{and} \quad d(x, z) = d_k(\bar{x}, \bar{z}), \quad (7.1.7)$$

that is, if the lengths of the sides of the two triangles in their respective spaces are equal.

Let $p \in \gamma_{xy}$ be a point on one of the sides of $\Delta(x, y, z)$ (analogous statements hold for γ_{xz} and γ_{yz}). The comparison point for p in a comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ is the unique point $\bar{p} \in \gamma_{\bar{x}\bar{y}}$ satisfying

$$d(p, x) = d_k(\bar{p}, \bar{x}) \quad \text{and} \quad d(p, y) = d_k(\bar{p}, \bar{y}). \quad (7.1.8)$$

Moving forward, we shall drop the subscript from d_k where it is clear which model space we refer to and use the same notation d for the distance function on X and on $M^2(k)$. Which distance function we refer to will be clear from context, since points in the model space will be decorated by bars e.g. \bar{x} . Furthermore, when we say distances between points in X and points in $M^2(k)$ are equal, we always apply the unspoken caveat “with respect to their associated distance functions.”

When defining a comparison triangle, we assumed a priori that there exists points $\bar{x}, \bar{y}, \bar{z}$ whose pairwise distances realise the lengths of the sides of $\Delta(x, y, z)$. However, it is possible that we could construct a triangle in X for which no such points exist in a given model space. For example, we could draw a triangle Δ in the Euclidean plane with longest side of length 4. Since the unit sphere has diameter π (see (7.1.6)), there are no points in $M^2(1) = \mathbb{S}^2(1)$ whose distance realises the length of the longest side of Δ . Hence we cannot construct a comparison triangle for Δ on the unit sphere.

As such, when comparing a metric space X to a model space $M^2(k)$, we should only consider triangles in X for which we can construct comparison triangles. In order to guarantee the existence of a comparison triangle, we introduce size-bounds on X .

Definition 7.1.11 (Metric Size-Bounds)

Let (X, d) be a metric space. A triangle $\Delta(x, y, z)$ in X is said to satisfy size-bounds for the model space $M^2(k)$ if

$$d(x, y) + d(y, z) + d(x, z) < 2 \operatorname{diam}(M^2(k)). \quad (7.1.9)$$

Applying triangle inequality to (7.1.9) says that $d(x, y) < \text{diam}(M^2(k))$ and analogously for the other two distances. That is, if a triangle in X satisfies size-bounds for a given model space, then each of its sides have lengths which can be realised in the model space. In fact, if a triangle in X satisfies size-bounds for $M^2(k)$, then a unique comparison triangle exists in $M^2(k)$ (up to isometry), see [136, Lemma I.2.14]. For the remainder of this chapter, we assume that all triangles the appropriate satisfy size-bounds; by (7.1.6), this is only a restriction for $k > 0$.

We now introduce the technique of triangle comparison. In line with the previous discussion, we only impose conditions on points which are closer than $\text{diam}(M^2(k))$.

Definition 7.1.12 (Curvature Bounds by Triangle Comparison)

Let (X, d) be a metric space. An open subset $U \subseteq X$ is called a $(\geq k)$ -comparison neighbourhood (resp. $(\leq k)$ -comparison neighbourhood) if

- (i) For all pairs of points $x, y \in U$ which are a distance $d(x, y) < \text{diam}(M^2(k))$ apart, there exists a distance realiser connecting them which is contained entirely within U , that is, $(U, d|_{U \times U})$ is a $\text{diam}(M^2(k))$ -geodesic metric space.
- (ii) For all triangles $\Delta(x, y, z)$ in U and all pairs of points $p, q \in \Delta(x, y, z)$ the following holds: given a comparison triangle $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ in $M^2(k)$ corresponding to $\Delta(x, y, z)$ and comparison points $\bar{p}, \bar{q} \in \bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ for p and q respectively, then

$$d(p, q) \geq d(\bar{p}, \bar{q}) \quad (\text{resp. } d(p, q) \leq d(\bar{p}, \bar{q})). \quad (7.1.10)$$

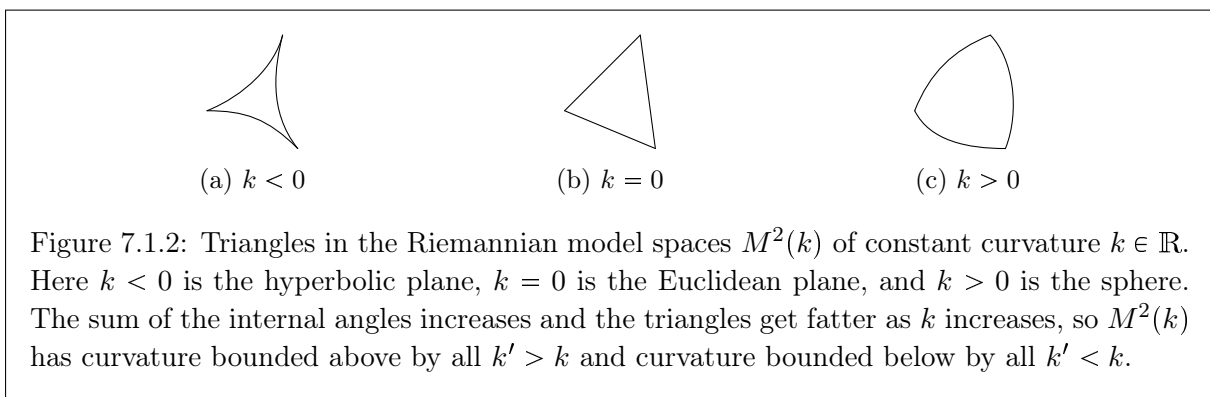
We say X has curvature bounded below by k if it is covered by $(\geq k)$ -comparison neighbourhoods. Likewise, X has curvature bounded above by k if it is covered by $(\leq k)$ -comparison neighbourhoods. We may say that such spaces have local curvature bounds to emphasize that (7.1.10) only holds on a neighbourhood of each point.

We say X has global curvature bounded below by k or is a CBB(k) space if X is a $(\geq k)$ -comparison neighbourhood. Similarly, we say X has global curvature bounded above by k or is a CAT(k) space if X is a $(\leq k)$ -comparison neighbourhood.

Recall that, for manifolds M , the sectional curvature (6.0.2) of the tangent plane at $p \in M$ spanned by $u, v \in T_p M$ also corresponds to the Gaussian curvature of the two-dimensional surface swept out by distance realisers emerging from p , with initial directions u, v . Furthermore, in regions where space is positively curved, these distance realisers are observed to spread apart faster than in regions of negative curvature [136, Appendix 1A]. From this, it should be clear that the notion of curvature which is bounded by comparison techniques, such as Definition 7.1.12, is precisely the sectional curvature, written in a synthetic manner using its description in terms

of distance realisers. In particular, the two-dimensional surfaces swept out by distance realisers are precisely triangles $\Delta(x, y, z)$ and we are bounding the curvature of such surfaces if they are contained in certain neighbourhoods.

In essence, a metric space (X, d) has curvature bounded below by k if “small” triangles in X are fatter than their comparison triangles in $M^2(k)$ and curvature bounded above by k if “small” triangles in X are thinner than their comparison triangles in $M^2(k)$. If we impose a global curvature bound, we instead consider arbitrary triangles in X .



It is worth noting here that there are several different comparison conditions we could have used to define curvature bounds, however we choose to introduce the concept via the triangle comparison condition (7.1.10) due to its intuitive description in terms of the thickness of a triangle. For completeness, [136, Proposition II.1.7] lists a selection of these conditions and demonstrates that they are equivalent for arbitrary metric spaces with curvature bounded above, while [100, Theorem 8.30] shows that they are equivalent for complete length spaces with curvature bounded below. When it comes to addressing curvature bounds in the Lorentzian setting in Chapter 8, we will explicitly introduce the various definitions which we require when proving our main results.

7.2 Globalisation of Curvature Bounds in Metric Spaces

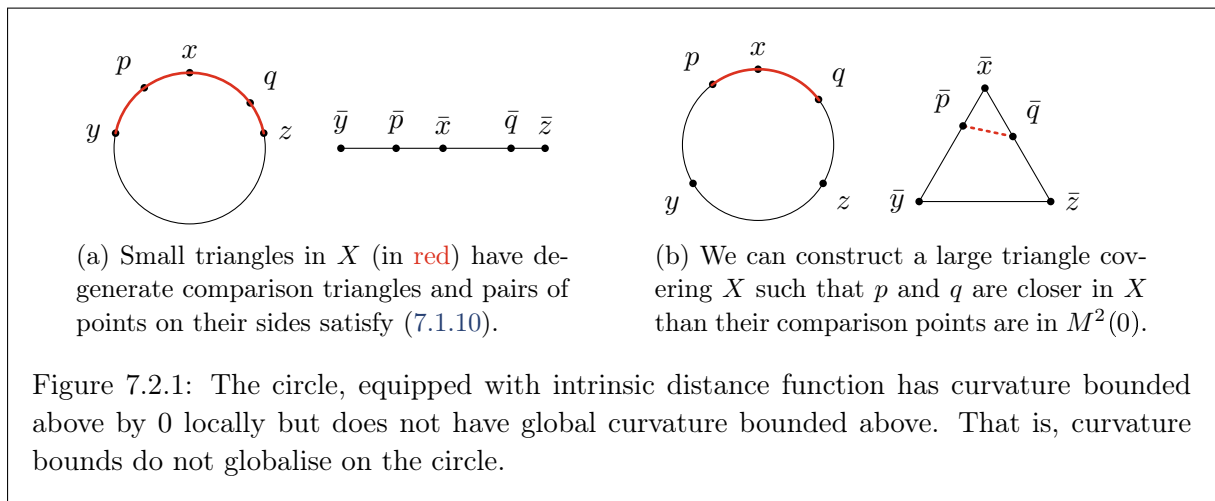
In this section, we address the globalisation of curvature bounds on metric spaces. By globalisation, we mean finding conditions under which imposing a local constraint, in this case curvature bounded above/below, automatically implies the corresponding global constraint. More precisely, if we have a space where “small triangles” are thinner/fatter than their comparison triangles in the model space $M^2(k)$, what additional assumptions are required for arbitrarily large triangles to also satisfy this comparison condition?

We will begin by looking at Alexandrov's patchwork globalisation theorem for spaces with curvature bounded above, before moving on to the Toponogov globalisation theorem for curvature bounded below. Much of this material can be found in [2, Chapter 3], however we stress that the results in this section are by no means original and are included as a point of inspiration for the analogous synthetic Lorentzian results described in Chapter 9. References for points of interest and further discussion will be included throughout.

Recall from the previous chapter that while curvature bounds may be defined via various comparison conditions, these conditions are equivalent under some additional mild assumptions. As it turns out, these same assumptions are also necessary for the globalisation theorems to hold. As such, the different definitions of curvature bounds globalise simultaneously and we need not specify which we use.

7.2.1 Curvature Bounded Above

Let us begin by looking at an example of a space which has (local) curvature bounded above, but does not have global curvature bounded above.



Example 7.2.1 (Curvature Bounds on the Circle (1))

Again let $X := \mathbb{S}^1$ be the unit circle with intrinsic distance function d , as in Example 7.1.6. Locally, X is isometric to a line segment, hence has curvature bounded above by 0. We can see this by considering triangles in X which do not cover the whole circle, for example $\Delta(x, y, z)$ in Figure 7.2.1a. The corresponding comparison triangle in Euclidean space degenerates to a straight line and we have equality in (7.1.10) for any points $p, q \in \Delta(x, y, z)$. In particular, all triangles contained within any open ball B of radius $\frac{\pi}{2}$ in X are of this type and it is

straightforward to see that B is a geodesic space, so B is a (≤ 0) -comparison neighbourhood, c.f. Definition 7.1.12. Furthermore, we can cover X with three of these balls, from which it follows that X has curvature (locally) bounded below by 0.

However, X itself is not a (≤ 0) -comparison neighbourhood: Consider a large triangle defined by three equidistant points in X , such that it covers the whole circle, see $\Delta(x, y, z)$ in Figure 7.2.1b, for example. Since $\text{diam}(M_0) = \infty$, the triangle $\Delta(x, y, z)$ satisfies size-bounds for $k = 0$ and the corresponding comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in the Euclidean plane is also an equilateral. Let p and q be two points on different sides of $\Delta(x, y, z)$, again as in Figure 7.2.1b, from which it follows that:

$$d(p, q) = d(p, x) + d(x, q) = d(\bar{p}, \bar{x}) + d(\bar{x}, \bar{q}) > d(\bar{p}, \bar{q}), \quad (7.2.1)$$

hence there exists a triangle in X for which (7.1.10) fails to hold, so X does not have global curvature bounded above by 0.

△

We see from this example that globalisation of curvature bounded above is not an automatic phenomenon, so our metric space will need to satisfy some additional assumptions, which \mathbb{S}^1 with intrinsic metric should not satisfy. The required assumptions were first written down in [101] by Alexandrov and are reiterated below for convenience. The proof of Theorem 7.2.2 involves decomposing arbitrarily large triangles into a patchwork of smaller triangles which are contained within $(\leq k)$ -comparison neighbourhoods. For brevity we direct the reader to [136, Proposition II.4.9] for the proof.

Theorem 7.2.2 (Alexandrov's Patchwork)

Let (X, d) be a uniquely $\text{diam}(M^2(k))$ -geodesic metric space with curvature bounded above by k . That is, for each pair of points in X a distance less than $\text{diam}(M^2(k))$ apart, there exists a unique distance realiser in X joining them. If these distance realisers vary continuously with their endpoints, then X has global curvature bounded above by k .

Let us make precise what we mean by the second requirement:

Definition 7.2.3 (Continuously Varying Distance Realisers)

Let (X, d) be a uniquely ℓ -geodesic metric space and let $x, y \in X$ be a pair of points such that $d(x, y) < \ell$. The distance realiser γ_{xy} from x to y in X is said to vary continuously with its endpoints if for every pair of sequences $x_n \rightarrow x$ and $y_n \rightarrow y$, such that there exists a unique distance realiser $\gamma_{x_n y_n}$ for each $n \in \mathbb{N}$, we have $\gamma_{x_n y_n} \rightarrow \gamma_{xy}$ uniformly.

The converse statement to [Theorem 7.2.2](#) also holds, hence the additional assumptions are not only sufficient, but necessary for upper curvature bounds to globalise (see [[136](#), Proposition II.1.4 (1)] for a proof):

Proposition 7.2.4 (Assumptions are Necessary)

Let (X, d) be a metric space with global curvature bounded above by k . Then (X, d) has (local) curvature bounded above by k by definition, distance realisers of length less than $\text{diam}(M^2(k))$ exist and are unique, and these distance realisers vary continuously with their endpoints.




Figure 7.2.2: It is possible to construct a sequence of distance realisers γ_n in the circle which do not vary continuously with their endpoint. In particular, if one endpoint is fixed at say, x , then if the other endpoint y_n varies through the antipodal point \tilde{x} of x , then the distance realisers γ_{n+1} (red) and γ_n (blue) lie on opposite sides of x and \tilde{x} .

Example 7.2.5 (Curvature Bounds on the Circle (2))

We saw in [Example 7.1.6](#) that the unit circle, denoted X , does not have unique distance realisers between points at distance π when equipped with its intrinsic metric. Recall that the Euclidean plane $M^2(0)$ has diameter $\text{diam}(M^2(0)) = \infty$, hence by [Proposition 7.2.4](#), the circle X cannot have global curvature bounded above by $k = 0$, as we expected.

The circle also fails to meet the other criteria of [Theorem 7.2.2](#) for $k = 0$, that is, the circle does not have continuously varying distance realisers between points a finite distance apart. Take any pair of points $x, y \in X$ which are not antipodal and set $x_n = x$ for all n . Any sequence $y_n \rightarrow y$ which passes through the point antipodal to x yields a sequence of $\gamma_{x_n y_n}$ which do not converge uniformly to γ_{xy} . In particular, the pair of distance realisers $\gamma_{x y \pm \varepsilon}$ would be on opposite sides of y , see [Figure 7.2.2](#).

△

In summary, a metric space X has global curvature bounded above by k if and only if X has (local) curvature bounded below, is uniquely $\text{diam}(M^2(k))$ -geodesic, and all distance realisers in X of length less than $\text{diam}(M^2(k))$ vary continuously with their endpoints. In particular, there

is a unique triangle in X for every triple of points satisfying (7.1.9).

7.2.2 Curvature Bounded Below

The progression of globalisation theorems in the case of curvature bounded below is rather a lot more involved than the case of curvature bounded above, with suitable conditions in two dimensions first being found by Pizzetti in [139] (see [140] for a review of Pizzetti's role in the history of comparison geometry). However, it was not until the late 1950s, when these conditions were independently re-derived by Alexandrov [141,101] and generalised to Riemannian manifolds of arbitrary dimension by Toponogov [102], that the idea of globalising lower curvature bounds garnered much interest. More recently however, as part of a push to develop the field of synthetic geometry, a large number of advancements have been made, with [103, 142, 143] extending the Toponogov globalisation theorem to arbitrary complete length spaces. Petrunin has also shown [144] that the completions of geodesic metric spaces with curvature bounded below similarly admit globalisation.

Since no additional conditions are imposed to prove the globalisation theorem for curvature bounded below, other than those required for the various comparison conditions to be equivalent, we simply provide the statement of the theorem in its most general form.

Theorem 7.2.6 (Toponogov's Globalisation Theorem)

Let X be a complete length space with curvature bounded below by k . Then X has global curvature bounded below by k .

Several nice proofs of the above result exist in the literature, including in [100, Theorem 8.31] and [109, Theorem 10.3.1], with the latter being made more accessible by a local compactness assumption. The proof we shall be most interested in for the purpose of globalising curvature bounds in the Lorentzian setting is that of Lang and Schroeder [133] since, like the Alexandrov's patchwork for curvature bounded above, their proof utilises a rather visual argument of constructing arbitrarily large triangles from triangles living inside comparison neighbourhoods.

The authors of [109] explicitly exclude certain one-dimensional spaces from the class of non-negatively curved spaces. More precisely, for X to have curvature bounded below by $k > 0$, X must not be isometric to \mathbb{R} , $(0, \infty)$, $[0, B]$ for any $B > \frac{\pi}{\sqrt{k}}$, or any circle with radius greater than $\frac{1}{\sqrt{k}}$. This is not due to the definitions themselves failing for these spaces, but rather a consequence of the following addendum to Toponogov's theorem:

Theorem 7.2.7 (Synthetic Bonnet–Myers Theorem)

Let (X, d) be a complete length space with (local) curvature bounded below by some $k > 0$, which is

not one of the aforementioned one-dimensional spaces. Then $\text{diam}(X) \leq \text{diam}(M^2(k)) := \frac{\pi}{\sqrt{k}}$.

This result was first proven by Bonnet for complete, two-dimensional, Riemannian manifolds, with Myers [145] formulating the result in higher dimensions. Myers later demonstrated that using a positive lower Ricci curvature bound in place of the sectional curvature bound is still sufficient to obtain the associated upper bound on the diameter [146]. In the metric setting, [100, Theorem 8.44] show that the aforementioned pathological, one-dimensional spaces are precisely the spaces with curvature bounded below for which the Bonnet–Myers theorem fails. Therefore, by following the approach of [133] when developing a synthetic Lorentzian Toponogov theorem in Section 10.2, we will not need to exclude one-dimensional spaces from consideration. However, when addressing the Bonnet–Myers theorem in the synthetic Lorentzian setting, we shall take the approach of [109, Theorem 10.4.1], first defining a degeneracy property which corresponds to the space being one-dimensional and then deriving a diameter bound for the remaining non-degenerate spaces, see Theorem 10.3.3.

The above theorem states that if a complete length space X (which is not one of the pathological ones) has local curvature bounded below by $k \in \mathbb{R}$, then for any distance realiser in X , there is a corresponding distance realiser of the same length in $M^2(k)$. As we saw in Definition 7.1.11, a triangle in X is said to satisfy size-bounds for $M^2(k)$ if the length of its perimeter is less than twice the diameter of $M^2(k)$. For $k \leq 0$, the perimeter simply has to be finite and this is no restriction at all, however for $k > 0$ it may be true a priori that triangles which don't satisfy size-bounds exist in X . The following corollary of Theorem 7.2.7, which is proven in [109, Theorem 10.4.2], states that the only triangles which don't satisfy size-bounds are those whose perimeters are exactly $2 \text{diam}(M^2(k))$, for which comparison triangles exist in $M^2(k)$ but may not be unique.

Corollary 7.2.8 (Lower Curvature Bounds and Size-Bounds)

Let (X, d) be a complete length space with (local) curvature bounded below by some $k > 0$, which is not one of the aforementioned one-dimensional spaces. Then for every triangle $\Delta(x, y, z)$ in X , we have

$$d(x, y) + d(y, z) + d(x, z) \leq 2 \text{diam}(M^2(k)). \quad (7.2.2)$$

In particular, triangles either satisfy size-bounds for k or have perimeter $2 \text{diam}(M^2(k))$.

In summary, a metric space X has global curvature bounded below by k if it is a complete length space with (local) curvature bounded below by k . Additionally, such spaces have diameter bounded above by $\text{diam}(M^2(k))$ and all triangles in them admit comparison triangles in $M^2(k)$, provided X is not locally isomorphic to certain pathological one-dimensional spaces.

7.3 Chapter Summary

In this chapter we have introduced the fundamental concepts of synthetic geometry, using the distances between points and curves to describe the curvature of an arbitrary metric space. We have discussed the properties of distance realisers and stated the conditions under which local bounds on the curvature of a metric space can be extended to global bounds. We also noted that there were various different comparison conditions that we could use to define curvature bounds and while these are equivalent for complete length spaces with curvature bounded below, they also coincide for generic metric spaces with curvature bounded above.


We observed how metric spaces with global curvature bounded above by k are precisely the uniquely $\text{diam}(M^2(k))$ -geodesic metric spaces with local curvature bounded above by k , whose distance realisers of length less than $\text{diam}(M^2(k))$ vary continuously with their endpoints. In particular, the Alexandrov's patchwork theorem says that a unique triangle must exist for every triple of sufficiently close points in X , in order for us to decompose large triangles in such a way that they satisfy a comparison condition.

Similarly, a metric space has global curvature bounded below by k if it has (local) curvature bounded below by k and is a complete length space. Unlike the curvature bounded above case, here the implication is only one-way and we do not recover the supplementary assumptions of completeness and an intrinsic distance function. However, these assumptions are still rather natural as they guarantee that the various comparison conditions used to define curvature bounds are equivalent. Hence, the globalisation theorem applies to all of these conditions simultaneously. Further contrasting the above/below cases, for curvature bounded below we do not even need to assume distance realisers exist, let alone (unique) triangles; Toponogov's theorem states that if I am given a big triangle Δ then there must exist sufficiently many small triangles to decompose Δ in such a way that it satisfies a comparison condition.

Finally, we investigated the synthetic Bonnet–Myers theorem, which bounds the diameter of spaces with curvature bounded below by $k > 0$, provided that the space is not in a particular pathological class. The Bonnet–Myers theorem implies that in spaces with curvature bounded below, all triangles satisfy size-bounds or have perimeter $2 \text{diam}(M^2(k))$. Hence, Toponogov's theorem for curvature bounded below considers all triangles whose perimeters are not equal to $2 \text{diam}(M^2(k))$, while Alexandrov's patchwork only considers triangles which are not “too large” when compared to the size of the model space.

In the next chapter, we introduce the Lorentzian pre-length space, the synthetic Lorentzian analogue of the metric space. While distance realisers, length, and triangle comparison make a return, they will have taken on a rather more tricky persona — that of an underlying Lorentzian

signature. Among our friends, a few new faces, including angles, Alexandrov's Lemma, and time functions, take on key technical roles. While many of these new concepts do not have an immediate metric analogue, some of them do and although we have been fairly expository so far, we feel that it would be rather cumbersome to introduce every dual concept in both its metric and its synthetic Lorentzian guise. However, we shall direct the reader to relevant resources when these comparisons can be made and highlight any that we deem especially important. We recommend that the knowledgeable reader keeps these comparisons in mind, since it will make the technical details in our main results more apparent.



8

An Introduction to Lorentzian Pre-length Spaces

This chapter is split into two main sections. In [Section 8.1](#), we introduce the notion of a Lorentzian pre-length space as our principle synthetic Lorentzian structure, alongside the Lorentzian model spaces of constant curvature, against which we make our comparisons. We then move on to defining curvature bounds via triangle comparison and illustrate, with reference to the definition's metric counterpart, why said comparison technique should only require the existence of distance realising curves up to a certain length, in contrast to the definition provided by [\[116, Definition 4.7\]](#), which requires distance a realising curve between every pair of points. In [Section 8.2](#), we review the concepts of angles and hinges in a Lorentzian pre-length space, before providing definitions of curvature bounds associated with each of these notions. While there are a wide array of comparison conditions we could choose from, see [\[123\]](#), we shall only require triangle, hinge, and angle comparison to prove the main results in this part of the thesis. Under relatively mild assumptions on our spaces, which we shall state at the end of this chapter, these three comparison conditions are known to be equivalent.

8.1 Fundamentals of Lorentzian Pre-length Spaces

We begin this chapter by describing some fundamental properties of the Lorentzian pre-length space framework and contrasting them with the metric setting. In particular, we discuss the causal classification of curves, the Alexandrov topology, and the definition of curvature bounds via triangle comparison. Along the way we introduce several rungs of the causal ladder [\[127, 147\]](#) which will be necessary in later proofs. For a more extensive exposition regarding the framework as a whole, we direct the interested reader to [\[116\]](#).

Mirroring the previous chapter, let us start by defining the synthetic Lorentzian analogue of metric spaces, namely Lorentzian pre-length spaces [116, Definition 2.8]:

Definition 8.1.1 (Lorentzian Pre-length Space)

Let (X, d) be a metric space equipped with two relations denoted \leq, \ll and a function $\tau : X \times X \rightarrow [0, \infty]$. The quintuple (X, d, \leq, \ll, τ) is called a Lorentzian pre-length space if it satisfies the following:

- (i) *Causal space:* \leq is a reflexive and transitive relation, and \ll is a transitive relation contained in \leq .
- (ii) *Semi-continuity:* τ is lower semi-continuous with respect to d .
- (iii) *Timelike positivity:* $\tau(x, y) > 0 \iff x \ll y$.
- (iv) *Reverse triangle inequality:* $\tau(x, z) \geq \tau(x, y) + \tau(y, z)$ for $x \leq y \leq z$.

The function τ is then called the time separation function, with \leq and \ll referred to as the causal relation and timelike relation, respectively. For brevity, we shall denote a Lorentzian pre-length space by its associated set X when the supplementary structures are clear from the context.

Let us make the comparison to metric geometry more explicit. The time separation function τ takes the place of the distance function and now satisfies a reverse triangle inequality in place of the usual triangle inequality, reflecting the change in signature in the corresponding smooth settings. Positivity in Definition 7.1.1 is replaced with the notion of timelike positivity here, with the time separation of points x and y being non-negative in general and positive when the points are timelike related. This is an inherently non-symmetric construction between timelike related points: a pair of points $x, y \in X$ satisfying $x \ll y$ and $\tau(x, y) > 0$ need not satisfy $y \ll x$ a priori, so $\tau(y, x)$ may vanish. Indeed, [116, Proposition 2.14] shows that for $x \neq y$, if $\tau(x, y)$ is positive and finite, then $\tau(y, x) = 0$, so $x \ll y$ but $y \not\ll x$.

Furthermore, unlike distance functions, the time separation τ need not be finite nor continuous, with only semi-continuity being enforced via the underlying metric space (X, d) in point (ii) of Definition 8.1.1. The main role of the distance function d here is to provide a topology, hence, when we refer to topological notions such as neighbourhoods and continuity, these are defined with respect to d . This mimics how any \mathcal{C}^k -Lorentzian spacetime, being a smooth manifold equipped with a \mathcal{C}^k -Lorentzian metric ($k \geq 0$), admits a Riemannian metric which provides the manifold topology [148, Proposition A.3.]. The synthetic Lorentzian construction has a synthetic Riemannian construction underlying it, in the form of a metric space, in the same way the smooth Lorentzian manifold has a smooth Riemannian manifold underlying it. Indeed, for smooth spacetimes, the Lorentzian and Riemannian metrics induce a Lorentzian pre-length

space structure:

Example 8.1.2 (Spacetimes as Lorentzian pre-length spaces)

Let (M, g) be a smooth spacetime and assume that g has Lorentzian signature $(-, +, +, \dots)$. Recall that a curve γ in a spacetime is called timelike if the tangent vector $\gamma'(t)$ at each point on the curve satisfies $g(\gamma'(t), \gamma'(t)) < 0$. Similarly, γ is called null if $g(\gamma'(t), \gamma'(t)) = 0$ and causal if it is timelike or null at each point on the curve.

The manifold topology on M is induced by any choice of Riemannian metric h on M , and as in (6.0.3), we may define a distance function d_h on M . Define the timelike (resp. causal) relation as follows: $x \ll y$ (resp. $x \leq y$) if there exists a future-directed timelike (resp. causal) curve from x to y . Finally, taking the length of a curve $\gamma : [a, b] \rightarrow M$ to be as in (6.0.1), the time separation of $x, y \in M$ is induced by the Lorentzian metric g similarly to how a Riemannian metric induces a distance function, namely

$$\tau_g(x, y) := \sup\{L_g(\gamma) \mid \gamma \text{ a future directed causal curve from } x \text{ to } y\}, \quad (8.1.1)$$

for $x \leq y$ and zero otherwise. It is then straightforward to verify that these choices satisfy the required properties in Definition 8.1.1, see [116, Examples 2.2, 2.11].

△

The structure of a Lorentzian pre-length space draws inspiration not only from the definition of a metric space, but also from that of a causal space cf. [117, Section 1.2], as indicated by the aforementioned lack of symmetry and by condition (i) in Definition 8.1.1. This enables us to capture the inherently Lorentzian notions of causality and chronology via relations describing the ordering of points in time. In particular, we have the push-up condition [116, Lemma 2.10]

$$x \leq y \ll z \Rightarrow x \ll z \text{ (similarly for } x \ll y \leq z). \quad (8.1.2)$$

However, the causal space as defined by Kunzinger and Sämann is actually a slight generalisation of that found in [117]. This is because the definition of a Lorentzian pre-length space does not force the timelike relation to be irreflexive, that is, we admit points $x \in X$ such that $x \ll x$. In particular, $x \ll x$ if and only if $\tau(x, x) = \infty$, otherwise $\tau(x, x) = 0$ [116, Proposition 2.14]. We call spaces for which \ll is irreflexive *chronological*. Furthermore, the causal relation need not be a partial order, that is, there may exist $x, y \in X$ such that $x \leq y$ and $y \leq x$, but $x \neq y$. Spaces for which $x \leq y \leq x$ implies $x = y$ are called *causal-ordered*.¹

¹The seminal paper [116, Definition 2.35] calls this property *causal*, but we feel that the terminology “causal causal space” is rather awkward and could lead to confusion.

As in the metric setting, an interval refers to any connected subset $I \subseteq \mathbb{R}$. Given an interval I , a curve $\gamma : I \rightarrow X$ in a Lorentzian pre-length space X is a locally 1-Lipschitz continuous map with respect to the distance function d , i.e. for each compact subset $K \subseteq I$ there exists a constant $C_K \geq 0$ such that for all $a, b \in K$, we have

$$d(\gamma(a), \gamma(b)) \leq C_K |b - a|. \quad (8.1.3)$$

As before, γ_{xy} is a curve from x to y with $\gamma(a) = x$ and $\gamma(b) = y$.

Here, we assume curves to be locally 1-Lipschitz continuous a priori, rather than just continuous as in the metric case. This again takes inspiration from the study of \mathcal{C}^k -spacetimes, where locally 1-Lipschitz curves with respect to the underlying Riemannian metric are generally considered due to their integrability [148, Section 2.1]. Whilst it could be argued that continuous curves are more natural for synthetic geometry due to continuity being the weaker condition, defining curves to be locally 1-Lipschitz also has practical benefits in Lorentzian pre-length spaces, with the limit curve theorem [116, Theorem 3.7] demonstrating that sequences of (locally 1-Lipschitz) causal curves in compact sets have subsequences which converge to curves of the same type.

As discussed in 8.1.2, curves on a smooth Lorentzian manifold admit an additional classification depending on whether or not the points along the curve are all timelike, causal, null, or spacelike related. We can similarly classify timelike, causal, and null curves using the causal space structure (X, \leq, \ll) and time separation function τ .

Definition 8.1.3 (Causal and Timelike Curves)

Let X be a Lorentzian pre-length space. A non-constant² curve $\gamma : [a, b] \rightarrow X$ is called future-directed causal (resp. timelike) if $\gamma(s) \leq \gamma(t)$ (resp. $\gamma(s) \ll \gamma(t)$) for all $s < t$ in $[a, b]$. Similarly, a curve is called past-directed causal (resp. timelike) if $\gamma(s) \leq \gamma(t)$ (resp. $\gamma(s) \ll \gamma(t)$) for all $s > t$ in $[a, b]$. For brevity and without loss of generality, we assume that causal curves (inclusive of timelike curves) are future-directed unless explicitly stated otherwise. A null curve is then a causal curve along which no two points are related by \ll .

While it is possible to define spacelike curves as those curves $\gamma : [a, b] \rightarrow X$ whose points are not causally related in either direction, the utility of such a class in Lorentzian pre-length spaces is minimal, since we cannot compute their length with respect to the time separation function. In the words of [116] “there is no built-in notion of a spacelike curve in our setting.” Consequently, when considering synthetic curvature bounds on Lorentzian pre-length spaces,

²Note that this still admits curves which are constant on some subinterval of $[a, b]$.

we can only describe triangles with sides given by causal curves (in fact, we require at least two sides to be timelike), hence only bound the curvature of timelike planes. The work [149] demonstrates that for sufficiently well behaved spacetimes, this is equivalent to a bound on the timelike sectional curvature.

Definition 8.1.4 (τ -length)

Let X be a Lorentzian pre-length space. The τ -length of a causal curve $\gamma : [a, b] \rightarrow X$ from x and y in X is given by

$$L_\tau(\gamma) := \inf \left\{ \sum_{i=0}^{n-1} \tau(\gamma(t_i), \gamma(t_{i+1})) \mid a = t_0 < t_1 < \dots < t_n = b, n \in \mathbb{N} \right\}. \quad (8.1.4)$$

A causal curve $\gamma : [a, b] \rightarrow X$ is called τ -rectifiable if $L_\tau(\gamma|_{[s,t]}) > 0$ for all $a \leq s < t \leq b$.

The τ -length of a (future-directed) causal curve $\gamma : [a, b] \rightarrow X$ is additive by [116, Lemma 2.25]; for $c \in (a, b)$ we have

$$L_\tau(\gamma) = L_\tau(\gamma|_{[a,c]}) + L_\tau(\gamma|_{[c,b]}). \quad (8.1.5)$$

Recall that, in the metric setting, applying the triangle inequality for the distance function d to the sum in (7.1.2) tells us that approximating the d -length of a curve using increasingly fine partitions of the curve's domain results in a non-decreasing sequence. In particular, $L_d(\gamma) \geq d(\gamma(a), \gamma(b))$. As such, the length of a curve is defined as the supremum over such partitions.

Performing a similar exercise here and applying the reverse triangle inequality for the time separation τ to the sum in (8.1.4), we find that approximating the τ -length of a causal curve using increasingly fine partitions of the curve's domain results in a non-increasing sequence. In particular, $L_\tau(\gamma) \leq \tau(\gamma(a), \gamma(b))$ for any causal curve $\gamma : [a, b] \rightarrow X$. As such, the length of a curve is defined as an infimum in this setting. Furthermore, causal distance realisers in spaces with Lorentzian signature (be they \mathcal{C}^k -spacetimes or Lorentzian pre-length spaces) are then an example of a ‘‘longest curve’’ and are defined as follows:

Definition 8.1.5 (Longest Curves and Distance Realisers)

Let X be a Lorentzian pre-length space and let $\gamma : [a, b] \rightarrow X$ be a causal curve from x to y in X . The curve γ is called a longest curve from x to y if, for all other causal curves $\tilde{\gamma}$ from x to y in X , we have $L_\tau(\gamma) \geq L_\tau(\tilde{\gamma})$. If $L_\tau(\gamma) = \tau(x, y)$, then γ is called a distance realiser.

Timelike *distance realisers* are τ -rectifiable causal curves by [116, Proposition 2.34] and τ -rectifiable causal curves are timelike *curves* by [116, Lemma 2.30]. In particular, a distance

realiser is timelike if and only if it is τ -rectifiable.

Definition 8.1.6 (Geodesic Space)

A Lorentzian pre-length space X is called an ℓ -geodesic space if there exists a distance realiser (in the sense of Definition 8.1.5) between every pair of causally related points $x, y \in X$ for which $\tau(x, y) < \ell$. If a distance realiser exists between every pair of causally related points in X , then we call X a geodesic space. Furthermore, X is called uniquely ℓ -geodesic (resp. uniquely geodesic) if there exists a unique distance realiser between pairs of causally related points $x, y \in X$ for which $\tau(x, y) < \ell$ (resp. between every pair of points in X).

As in the metric setting, distance realisers between points in X need not exist or be unique (see Remark 8.1.27), so where the choice of distance realiser is important, we shall write our selection explicitly. Otherwise the reader is free to pick their favourite and we shall simply write γ_{xy} , for example. Since we do not use metric and Lorentzian distance realisers simultaneously, we generally do not state which of the two notions we refer to, since context will be a sufficient indicator.

Now that we have decided upon an appropriate class of curves, namely causal distance realisers, we turn our attention to constructing triangles from said curves.

Definition 8.1.7 (Admissible Causal Triangles)

Let X be an Lorentzian pre-length space. An admissible causal triangle in X consists of a collection of three points $x, y, z \in X$ (the vertices), which must satisfy $\tau(x, z) < \infty$ and either $x \ll y \leq z$ or $x \leq y \ll z$. Furthermore, the time separations of the vertices, if non-zero, are realised by pairwise connecting distance realisers $\gamma_{xy}, \gamma_{xz}, \gamma_{yz}$ (the sides). Vanishing time separations may be realised by either null or constant curves. We will denote such a triangle by $\Delta(x, y, z)$, where points are written according to their causal order. If we additionally have $x \ll y \ll z$, then the points and distance realisers comprise a timelike triangle, cf. [116, Lemma 4.4].

By the reverse triangle inequality in Definition 8.1.1, the side γ_{xz} of an admissible causal triangle $\Delta(x, y, z)$ is said to be the longest side, with the γ_{xy} and γ_{yz} being referred to as the short sides. A side between timelike related points shall be referred to as a timelike side, even though the realising curve may merely be causal. We write $p \in \Delta(x, y, z)$ to indicate that a point $p \in X$ lies on the triangle and $p \in \gamma_{xy}$ to specify that p lies on the side between x and y , for example.

We call an admissible causal triangle $\Delta(x, y, z)$ non-degenerate if the reverse triangle inequality is strict for the vertices of the triangle, i.e. if $\tau(x, z) > \tau(x, y) + \tau(y, z)$. An ad-

missible causal triangle is therefore called *degenerate* if it satisfies triangle equality $\tau(x, z) = \tau(x, y) + \tau(y, z)$. In particular, this holds if the triangle is given by a single distance realiser, that is, if γ_{xy} and γ_{yz} are segments of γ_{xz} , by (8.1.5).

Let us now introduce the key components of comparison geometry in Lorentzian pre-length spaces: the Lorentzian model spaces of constant curvature and comparison triangles in them. As in Section 7.1, we write $\mathbb{R}^{m,n}$ with $m, n \in \mathbb{N}_0$ for the vector space \mathbb{R}^{m+n} equipped with indefinite inner product

$$b^{m,n}(x, y) := -\sum_{i=1}^m x_i y_i + \sum_{i=1}^n x_{m+i} y_{m+i}, \quad (8.1.6)$$

for $x = (x_1, x_2, \dots, x_{m+n})$ and $y = (y_1, y_2, \dots, y_{m+n})$ and describe the Lorentzian model spaces in terms of quadric surfaces in suitable \mathbb{R}^{m+n} . We refer to the directions with indices $1, 2, \dots, m$ as *time directions*, while the remaining indices correspond to *spacial directions*, so $\mathbb{R}^{m,n}$ has m time directions and n spacial directions. More detail on the explicit constructions in the Lorentzian case can be found in [125, Definition 1.11] and [125, Proposition 4.29]. There should be no confusion between the Riemannian and Lorentzian model spaces, since we shall not consider them simultaneously, however we shall use K to denote the curvature of a (synthetic) Lorentzian space, as opposed to k in the metric case, cf. Definition 7.1.9.

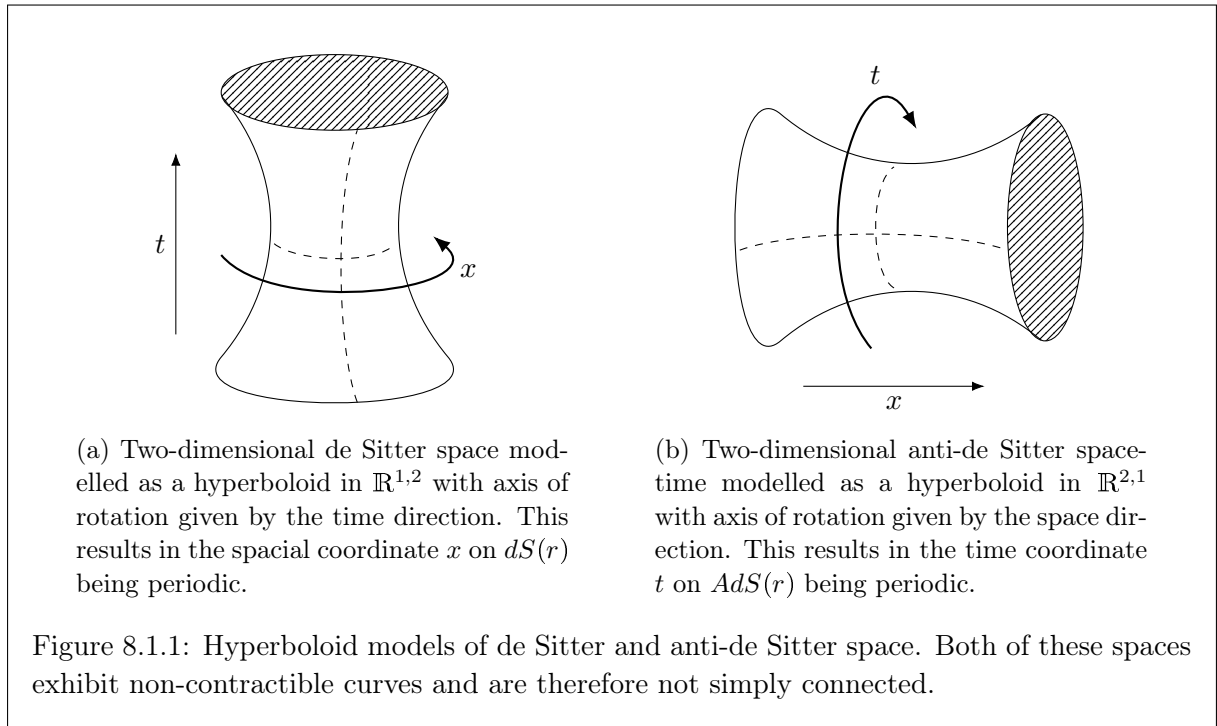
Definition 8.1.8 (Lorentzian Model Spaces)

The Lorentzian model space of constant curvature K is the unique, simply connected, two-dimensional, Lorentzian manifold with sectional curvature K , which we denote by $\mathbb{L}^2(K)$. In particular,

$$\mathbb{L}^2(K) = \begin{cases} dS(r) & K = \frac{1}{r^2} > 0 \\ \mathbb{R}^{1,1} & K = 0 \\ AdS(r) & K = -\frac{1}{r^2} < 0 \end{cases} \quad (8.1.7)$$

Here, $dS(r)$ denotes the universal cover of 2-dimensional de Sitter space $\{x \in \mathbb{R}^{1,2} \mid b^{1,2}(x, x) = r^2\}$ with scale factor r , which can be thought of as a one-sheeted hyperboloid with skirt radius r and axis of rotation given by the time direction in $\mathbb{R}^{1,2}$. Similarly, $AdS(r)$ denotes the universal cover of 2-dimensional anti-de Sitter space $\{x \in \mathbb{R}^{2,1} \mid b^{2,1}(x, x) = -r^2\}$ with scale factor r , which can be thought of as a one-sheeted hyperboloid with skirt radius r and axis of rotation given by the space direction in $\mathbb{R}^{2,1}$. Finally, $\mathbb{R}^{1,1}$ is the well-known Minkowski space.

Each of the Lorentzian model spaces can be viewed as a smooth Lorentzian manifold M , equipped with the Lorentzian metric given by the restriction of the pseudo-Riemannian met-



ric (8.1.6) (where x and y are viewed as vector fields on M) on their respective ambient spaces. Viewed in this way, the model spaces are spacetimes, cf. Example 8.1.2, hence Lorentzian pre-length spaces.

Unlike in the Riemannian case, when considering the Lorentzian model spaces of constant non-zero curvature, we are required to consider the universal covering of a quadric surface, rather than the surface itself. This is because both de Sitter and anti-de Sitter, as constructed above, are one-sheeted hyperboloids, which are not simply connected, since one can draw a non-contractible circle around their neck³, see Figure 8.1.1. The process of taking the universal cover effectively “rolls out” the hyperboloid so that it is no longer periodic and does not admit such curves. During the remainder of this thesis, when we refer to (anti-)de Sitter space it is understood that we mean the universal covering thereof.

Definition 8.1.9 (Comparison Triangles and Points)

Let X be a Lorentzian pre-length space and $\mathbb{L}^2(K)$ the model space of constant curvature K , equipped with its natural time separation function τ_K . Let $\Delta(x, y, z)$ be an admissible causal triangle in X . An admissible causal triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in $\mathbb{L}^2(K)$ is called a comparison triangle

³In the case of anti-de Sitter, these non-contractible curves are examples of closed timelike curves [150].

for $\Delta(x, y, z)$ if

$$\tau(x, y) = \tau_K(\bar{x}, \bar{y}), \quad \tau(y, z) = \tau_K(\bar{y}, \bar{z}), \quad \text{and} \quad \tau(x, z) = \tau_K(\bar{x}, \bar{z}), \quad (8.1.8)$$

that is, if the τ -lengths of the sides of the two triangles in their respective spaces are equal.

Assume that the side γ_{xy} of the triangle $\Delta(x, y, z)$ has timelike related endpoints $x \ll y$ and let $p \in \gamma_{xy}$ be a point on that side (analogously for γ_{xz} if $x \ll z$ and γ_{yz} if $y \ll z$). The comparison point for p in a comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ is the unique point $\bar{p} \in \gamma_{\bar{x}\bar{y}}$ satisfying

$$\tau(x, p) = \tau_K(\bar{x}, \bar{p}) \quad \text{and} \quad \tau(p, y) = \tau_K(\bar{p}, \bar{y}). \quad (8.1.9)$$

In contrast to [Definition 7.1.10](#), the order of the points in (8.1.8) and (8.1.9) are important — we write points following the causal ordering of the curve on which they lie.

As when we discussed triangle comparison in metric geometry, we shall drop the subscript from τ_K where it is clear which model space we refer to and denote the time separation on both X and $\mathbb{L}^2(K)$ by τ . Points in the model space will be decorated with bars e.g. \bar{x} , from which it should be clear when we refer to the time separation on the model space. Furthermore, when we say time separation between points in X and points in $\mathbb{L}^2(K)$ are equal, we omit the caveat “in their respective time separations.”

In [Definition 8.1.7](#), we required admissible causal triangles to have finite time separation between their vertices, with $\tau(x, z) < \infty$ being imposed and the reverse triangle inequality implying $\tau(x, y), \tau(y, z) < \infty$. For the purposes of triangle comparison, however, this restriction is mostly a formality since we shall only be considering admissible causal triangles for which a corresponding comparison triangle exists. In particular, the values of the time separation between the vertices of the admissible causal triangle in X must be attained by the time separation on $\mathbb{L}^2(K)$. As in [Definition 7.1.11](#), these conditions shall be referred to as size-bounds.

However, if we use the τ -diameter (or timelike diameter) defined for Lorentzian pre-length spaces in complete analogy to [Definition 7.1.2](#), that is,

$$\text{diam}_\tau(X) := \sup\{\tau(x, y) \mid x, y \in X\}, \quad (8.1.10)$$

we run into issues when $K < 0$ and we are make comparisons to $\mathbb{L}^2(K) = AdS\left(\frac{1}{\sqrt{-K}}\right)$. In particular, the time separation function (8.1.1) on $AdS\left(\frac{1}{\sqrt{-K}}\right)$ takes values in $\left[0, \frac{\pi}{\sqrt{-K}}\right] \cup \{\infty\}$ for $K = -\frac{1}{r^2}$. We shall discuss how this pathological behaviour occurs in [Remark 8.1.27](#), where it results in another, more nuanced problem. For now, note that admissible causal triangles in a

Lorentzian pre-length space X with (at least one) side γ such that $L_\tau(\gamma) \in \left(\frac{\pi}{\sqrt{-K}}, \infty\right)$ have no comparison triangle in $AdS\left(\frac{1}{\sqrt{-K}}\right)$, since no distance realisers of τ -length $L_\tau(\gamma)$ exist in that model space.

Consequently, we now introduce the finite timelike diameter [125, Definition 1.11], which is the supremal value of the time separation between points which are not infinitely far apart.

Definition 8.1.10 (Finite Timelike Diameter)

Let X be a Lorentzian pre-length space.

(i) The finite timelike diameter of X is given by

$$\text{diam}_{\text{fin}}(X) := \sup(\{\tau(x, y) \mid x, y \in X\} \setminus \{\infty\}), \quad (8.1.11)$$

i.e., the supremum of all values τ takes except ∞ .

(ii) We denote the finite timelike diameter of $\mathbb{L}^2(K)$ by D_K . In particular,

$$D_K := \text{diam}_{\text{fin}}(\mathbb{L}^2(K)) = \begin{cases} \infty & K \geq 0, \\ \frac{\pi}{\sqrt{-K}} & K < 0. \end{cases} \quad (8.1.12)$$

Definition 8.1.11 (Timelike Size-Bounds)

Let X be a Lorentzian pre-length space. The vertices x, y, z of an admissible causal triangle $\Delta(x, y, z)$ in X are said to satisfy size-bounds for the model space $\mathbb{L}^2(K)$ if

$$\tau(x, z) < D_K. \quad (8.1.13)$$

In particular, x, y, z satisfy the conditions of the realisability lemma [116, Lemma 4.6]⁴, hence admit a unique comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in $\mathbb{L}^2(K)$. For the remainder of this thesis, we tacitly state that the triangle itself satisfies size-bounds and assume that the vertices of admissible causal triangles satisfy size-bounds unless explicitly stated otherwise.

At this stage, we highlight why it is necessary to allow the non-timelike sides of admissible causal triangles to be realised by a constant curve. Consider the degenerate admissible causal triangle $\Delta(x, y, z)$ with $\tau(x, y) = 0$ (γ_{xy} may be null or constant) and $\tau(y, z) = \tau(x, z) > 0$. Let $\Delta(\bar{x}, \bar{y}, \bar{z})$ be a comparison triangle for $\Delta(x, y, z)$ in $\mathbb{L}^2(0)$. The side $\gamma_{\bar{x}\bar{z}}$ is given by the unique distance realiser from \bar{x} to \bar{z} in $\mathbb{L}^2(0)$, which is described by a timelike line. However, by (8.1.5),

⁴With the additional constraint that at most one of the τ -lengths may vanish.

the curve given by the concatenation of $\gamma_{\bar{x}\bar{y}} + \gamma_{\bar{y}\bar{z}}$ has length

$$\tau(\bar{x}, \bar{y}) + \tau(\bar{y}, \bar{z}) = \tau(x, z) = \tau(\bar{x}, \bar{z}), \quad (8.1.14)$$

and is therefore also a distance realiser from \bar{x} to \bar{z} . Hence, $\gamma_{\bar{x}\bar{z}} = \gamma_{\bar{x}\bar{y}} + \gamma_{\bar{y}\bar{z}}$ by uniqueness. Since $\gamma_{\bar{x}\bar{z}}$ is a timelike line, $\gamma_{\bar{x}\bar{y}}$ cannot be a null curve, from which it follows that it must be a constant curve, with $\bar{x} = \bar{y}$. This process illustrates the fact that the unique comparison triangle in $\mathbb{L}^2(K)$ for a degenerate admissible causal triangle is a timelike line.

In summary, if we did not allow non-timelike sides of admissible causal triangles to be realised by constant curves, then some admissible causal triangles with null sides would not have admissible causal comparison triangles in one or more of the model spaces. In particular, if an admissible causal triangle has a side realised by a constant curve, then the τ -lengths of the two timelike sides are equal, since they are both distance realisers. Furthermore, on a uniquely geodesic space, these distance realisers must coincide and all admissible causal triangles with a side given by a constant curve degenerate to a timelike line.

The attentive reader may have noticed the absence of a particular concept from our introductory foray into synthetic Lorentzian geometry thus far, namely that of a *Lorentzian length space* with an intrinsic time separation function. Such spaces are analogous to metric length spaces with intrinsic distance function, however their construction is a little more involved, since we must respect the causal character of curves. In order to provide their definition, we first need to take a brief detour up the causal ladder.

8.1.1 Topology of Lorentzian Pre-length Spaces

In this subsection, we provide a minimal overview of the relevant rungs of the causal ladder required for our globalisation theorems, focusing on the topological constraints of strong causality and global hyperbolicity. Along the way, we also introduce the notion of a Lorentzian length space. Let us begin our detour with another observation by [117, Section 1.2]; the causal (\preceq) and timelike (\ll) relations can be used to define neighbourhoods on a causal space, which in turn may be used to describe its topology.

Definition 8.1.12 (Causal and Timelike Pasts and Futures)

Let (X, \preceq, \ll) be a causal space as defined in Definition 8.1.1.

(i) The causal future (resp. past) of a point $x \in X$ is given by

$$J^+(x) := \{y \in X \mid x \preceq y\} \text{ (resp. } J^-(x) := \{y \in X \mid y \preceq x\}). \quad (8.1.15)$$

(ii) The timelike future (resp. past) of a point $x \in X$ is given by

$$I^+(x) := \{y \in X \mid x \ll y\} \text{ (resp. } I^-(x) := \{y \in X \mid y \ll x\}). \quad (8.1.16)$$

(iii) The causal (resp. timelike) diamond with past governing point x and future governing point y is given by

$$\begin{aligned} J(x, y) &:= J^+(x) \cap J^-(y) = \{z \in X \mid x \leq z \leq y\} \\ \text{(resp. } I(x, y) &:= I^+(x) \cap I^-(y) = \{z \in X \mid x \ll z \ll y\}). \end{aligned} \quad (8.1.17)$$

Since the causal and timelike relations are both transitive, (8.1.17) implies that a causal (resp. timelike) diamond is empty if the the past and future governing points are not causally (resp. timelike) related in that order. The causal relation is also reflexive, so if the governing points of a causal diamond $J(x, y)$ are causally related, we have

$$x \leq x \leq y \quad \text{and} \quad x \leq y \leq y. \quad (8.1.18)$$

Hence, the governing points of a causal diamond are contained within the diamond and the converse statement also holds — a causal diamond is empty if and only if its governing points are not causally related. For timelike diamonds, it is not generally true that the governing points are contained within the diamond, so timelike related governing points do not imply that a timelike diamond is non-empty. In the remainder of this thesis, diamonds are non-empty (and hence have causal/timelike related governing points) unless explicitly stated otherwise.

Timelike diamonds have some additional properties which make them favourable over arbitrary neighbourhoods of points when studying the time separation function. In particular, the governing points of a timelike diamond may be used to control the time separation between points within the diamond.

Lemma 8.1.13 (Timelike Diamond Properties)

Let X be a Lorentzian pre-length space and $I(x_-, x_+)$ be a timelike diamond with governing past governing point x_- and future governing point x_+ . The following properties hold for all $y, z \in I(x_-, x_+)$ such that $y \leq z$:

(i) Causal convexity: $J(y, z) \subseteq I(x_-, x_+)$.

(ii) Bounded time separation: $\tau(y, z) \leq \tau(x_-, x_+)$.

Proof. Let us prove each property in turn.

- (i) First note that $x_- \ll y \leq z \ll x_+$. Since $J(y, z)$ is non-empty by definition, consider an arbitrary point $p \in J(y, z)$, i.e. $y \leq p \leq z$. We therefore have $x_- \ll y \leq p$ and from (8.1.2) it follows that $x_- \ll p$. Analogously, we find $p \ll x_+$, so $p \in I(x_-, x_+)$ as required. In particular, if γ_{yz} is a causal curve, then $\gamma_{yz} \subseteq J(y, z) \subseteq I(x_-, x_+)$.
- (ii) Since $x_- \ll y \leq z \ll x_+$, our result follows from repeated applications of the reverse triangle inequality:

$$\tau(x_-, x_+) \geq \tau(x_-, y) + \tau(y, x_+) \geq \tau(x_-, y) + \tau(y, z) + \tau(z, x_+), \quad (8.1.19)$$

with the non-negativity of τ yielding $\tau(x_-, x_+) \geq \tau(y, z)$.

□

Now let (X, d, \leq, \ll, τ) be a Lorentzian pre-length space. Recall how, in Definition 8.1.1 and Definition 8.1.3, we imposed semi-continuity and Lipschitz continuity with respect to the topology induced by the (seemingly arbitrary) distance function d on X . As an alternative, we may use timelike diamonds as a means of defining topology on X :

Definition 8.1.14 (Alexandrov Topology)

The Alexandrov topology with respect to \ll on a causal space (X, \leq, \ll) is the topology generated by taking the set of all timelike diamonds $\mathcal{I} := \{I(x, y) \mid x, y \in X\}$ as a subbase. In particular, every open set in the Alexandrov topology can be given as the union of finite intersections of timelike diamonds.

Being able to describe the topology of a Lorentzian pre-length space in terms of timelike diamonds is useful because it enables us to apply Lemma 8.1.13 within a neighbourhood of each point. As such, let us distinguish Lorentzian pre-length spaces for which the *metric topology* and the Alexandrov topology coincide, following the conventions of [116, Definition 2.35]:

Definition 8.1.15 (Strong Causality)

A Lorentzian pre-length space X is called strongly causal if the metric topology induced by d coincides with the Alexandrov topology with respect to \ll . That is, $\mathcal{I} := \{I(x, y) \mid x, y \in X\}$ is a subbase for the metric topology induced by d .

Let us now turn to defining the notion of a Lorentzian length space, that is, a Lorentzian pre-length space with an intrinsic time separation function. As with the definition of a Lorentzian pre-length space, we must ensure that we respect both the causal structure defined by \leq, \ll and the topology defined by d . We begin by introducing the notion of a (locally) causally closed Lorentzian pre-length space, for which sequences of causally related points converge to

another pair of causally related points. This property also arises in the definition of an angle in a Lorentzian pre-length space as given by [125] (see also Section 8.2).

Definition 8.1.16 (Locally Causally Closed)

Let X be a Lorentzian pre-length space. A neighbourhood U of a point $x \in X$ is called causally closed if, for pairs of sequences $p_n \rightarrow p \in \bar{U}$ and $q_n \rightarrow q \in \bar{U}$ such that $p_n, q_n \in U$ and $p_n \preceq q_n$ for all $n \in \mathbb{N}$, we have $p \preceq q$. That is, \preceq is closed with respect to d in $\bar{U} \times \bar{U}$, where \bar{U} is the closure of U . A Lorentzian pre-length space X is called locally causally closed if every point $x \in X$ has a causally closed neighbourhood.

Definition 8.1.17 (Lorentzian Length Space)

Let X be a locally causally closed Lorentzian pre-length space. Then X is called a Lorentzian length space if additionally:

- (i) X is localisable in the sense of [116, Definition 3.16].
- (ii) X is causally path connected: For $x, y \in X$ such that $x \neq y$, if $x \preceq y$ (resp. $x \ll y$) then there exists a future directed causal (resp. timelike) curve from x to y .
- (iii) τ is intrinsic: For all $x, y \in X$, the time separation is given by

$$\tau(x, y) := \sup\{L_\tau(\gamma) \mid \gamma \text{ are future-directed causal curves from } x \text{ to } y\}, \quad (8.1.20)$$

where we define $\tau(x, y) = 0$ if there are no future-directed causal curves from x to y .

Recall from Definition 8.1.5 that distance realisers in a Lorentzian pre-length space are longest curves. As for the distance function in the metric setting, imposing that the time separation is intrinsic yields the converse statement. That is, if a longest curve exists, then it attains the supremum in (8.1.20) and the curve is a distance realiser.

Definition 8.1.18 (Globally Hyperbolic)

A Lorentzian pre-length space X is called globally hyperbolic if the causal diamonds $J(x, y)$ are compact for all $x, y \in X$ and for every compact set $K \subseteq X$ there exists a constant C_K such that $L_d(\gamma) \leq C_K$ for all causal curves γ contained in K .

When used tactfully, the globally hyperbolic Lorentzian length space is the Swiss army knife of synthetic Lorentzian spaces, incorporating all of the tools we have gathered this far into one neat package. In particular, the causal ladder [116, Theorem 3.26] tells us that strongly causal Lorentzian length spaces are chronological, while globally hyperbolic Lorentzian length spaces are strongly causal. In fact, by [116, Theorems 3.28, 3.30], globally hyperbolic Lorentzian length

spaces are geodesic and have finite, continuous time separation function. Furthermore, we have the following result:

Proposition 8.1.19 (Global Hyperbolicity Implies Causal Closure)

Let X be a globally hyperbolic Lorentzian length space. Then X is causally closed.

Proof. First recall that, for globally hyperbolic Lorentzian length spaces X , the causal pasts and futures $J^\pm(x)$ are closed for all $x \in X$, see [147, Proposition 3.12]. The remainder of the proof follows [151, Theorem 4.12] in the setting of globally hyperbolic spacetimes.

Let $p_n \rightarrow p$ and $q_n \rightarrow q$ be a pair of sequences in X with $p_n \preceq q_n$. We want to show $p \preceq q$. Let $r \in I^+(q)$ such that for sufficiently large n , we have $q_n \in I^-(r)$ (recall here that $I^-(r)$ is open). Consequently, we find $p_n \preceq q_n \ll r$, which in turn implies $p_n \ll r$. Hence, for large enough n , both $p_n, q_n \in I^-(r) \subseteq J^-(r)$. Since $J^-(r)$ is closed, this then implies that $p \in J^-(r)$ and $r \in J^+(p)$.

Since $q \in J^+(q)$, there exists a sequence of such r , say $\{r_k\}_{k \in \mathbb{N}} \subseteq I^+(q)$, such that $r_k \rightarrow q$ (see [147, Propositions 3.11, 3.12] for more details). In particular, $r_k \in J^+(p)$ for all k . As $J^+(p)$ is closed, it follows that $q \in J^+(p)$, i.e. $p \preceq q$, as required. \square

Lorentzian length spaces also possess a property known as non-timelike local isolation, which we shall introduce in the next section.

8.1.2 Timelike Curvature Bounds via Triangle Comparison

In this subsection, we will finally provide the definition of a Lorentzian pre-length space with timelike curvature bounds. However, the knowledgeable reader will notice some subtle differences between our definition and the original definition provided in [116, Definition 4.7]. Therefore, the bulk of this subsection is dedicated to clarifying these technical details (first proposed by the authors in [2]), as well as several others we have brushed under the carpet thus far.

Mirroring the metric setting, cf. Definition 7.1.12, we consider timelike triangles satisfying size-bounds (see Definition 8.1.11) for $K \in \mathbb{R}$, since these are the triangles for which comparison triangles in $\mathbb{L}^2(K)$ exist. However, in contrast to the metric setting, the original definition of timelike curvature bounds requires the existence of a distance realiser between points an arbitrarily large time separation apart; the Lorentzian pre-length space is assumed to be geodesic. The first modification we make is weakening this constraint to D_K -geodesic, in-line with the metric setting.

The second modification we make concerns the time separation function τ on the Lorentzian pre-length space X . In [116, Definition 4.7], τ is required to be finite and continuous on X . This

constraint has no analogue in the metric setting, but is instead intended to mimic the fact that the distance function on a metric space is continuous and can be assumed to be finite a priori.⁵ However, in the synthetic Lorentzian setting, if we are only interested in points a finite time separation less than D_K apart, then it is not necessary to impose that τ is finite, nor indeed that it is continuous on the whole domain (rather on some suitable sub-domain consisting of points which are not too far apart).

In summary, *all* properties of a comparison neighbourhood should respect the appropriate range of values for τ , with respect to the model space $\mathbb{L}^2(K)$. In particular, we are not interested in how the Lorentzian pre-length space behaves at time separations which cannot be realised in the model space. The resulting definition of timelike curvature bounds is as follows:

Definition 8.1.20 (Curvature Bounds by Triangle Comparison)

Let X be a Lorentzian pre-length space. An open subset U is called a $(\geq K)$ -comparison neighbourhood (resp. $(\leq K)$ -comparison neighbourhood) if:

- (i) τ is continuous on $(U \times U) \cap \tau^{-1}([0, D_K])$ and $(U \times U) \cap \tau^{-1}([0, D_K])$ is open (in the product topology on $X \times X$).
- (ii) For all $x, y \in U$ with $x \ll y$ and $\tau(x, y) < D_K$, there exists a distance realiser connecting them which is contained entirely in U . That is, U is D_K -geodesic with respect to $\tau|_{U \times U}$ in the sense of Definition 8.1.6.
- (iii) Let $\Delta(x, y, z)$ be a timelike triangle in U , with p, q two points on the sides of $\Delta(x, y, z)$. Let $\bar{\Delta}(\bar{x}, \bar{y}, \bar{z})$ be a comparison triangle in $\mathbb{L}^2(K)$ for $\Delta(x, y, z)$ and \bar{p}, \bar{q} comparison points for p and q , respectively. Then

$$\tau(p, q) \leq \tau(\bar{p}, \bar{q}) \quad (\text{resp. } \tau(p, q) \geq \tau(\bar{p}, \bar{q})). \quad (8.1.21)$$

We say X has timelike curvature bounded below by K if it is covered by timelike $(\geq K)$ -comparison neighbourhoods. Likewise, X has timelike curvature bounded above by K if it is covered by timelike $(\leq K)$ -comparison neighbourhoods.

We say X has global timelike curvature bounded below by K if X itself is a $(\geq K)$ -comparison neighbourhood. Similarly, X has global timelike curvature bounded above by K if X is a $(\leq K)$ -comparison neighbourhood.

The following proposition illustrates that this new definition of curvature bounds is a generalisation of the that provided in [116, Definition 4.7].

⁵See the discussion following Definition 7.1.1.

Proposition 8.1.21 (Generalised Curvature Bounds)

Let X be a Lorentzian pre-length space. All open $(\leq K)$ -comparison neighbourhoods in the sense of [116, Definition 4.7] are $(\leq K)$ -comparison neighbourhoods in the sense of Definition 8.1.20. The analogous statement holds for the $(\geq K)$ case. In particular, if X has a global curvature bound in the former sense, it also has that global curvature bound in the latter sense.

Proof. Let U be an open comparison neighbourhood in the sense of [116, Definition 4.7] and consider each part of Definition 8.1.20 in turn.

- (i) By definition, the time separation τ is finite and continuous on $U \times U$. Hence, τ is continuous on $(U \times U) \cap \tau^{-1}([0, D_K])$ as a subspace. Since $[0, D_K)$ is open in $[0, \infty)$ and $\tau : U \times U \rightarrow [0, \infty)$ is continuous, it follows that $(U \times U) \cap \tau^{-1}([0, D_K])$ is open in $U \times U$ and hence in $X \times X$.
- (ii) Distance realisers exist between any pair of timelike related points in U and are contained within U by definition. Hence such distance realisers exist between timelike related points $x, y \in U$ such that $\tau(x, y) < D_K$.
- (iii) Since we consider only timelike triangles satisfying size- bounds (see Definition 8.1.11), the sides of triangles have τ -length at most D_K . Consequently, this point is unchanged between definitions.

Setting $U = X$ yields the statement concerning global curvature bounds. □

Under the assumption that X is a chronological Lorentzian pre-length space, we may recover, in part, the restrictions on the finiteness and continuity of τ from [116, Definition 4.7]. We call the time separation function τ *locally continuous*⁶ if, for every point $x \in X$, there is a neighbourhood $V \subseteq X$ of x on which $\tau|_{V \times V}$ is continuous. Similarly, we call τ *locally finite* if there is such a collection of neighbourhoods V on which $\tau|_{V \times V}$ is finite. Any Lorentzian pre-length space with locally finite τ is chronological, since for any $x \in X$, $\tau(x, x)$ must be finite and therefore $\tau(x, x) = 0$, cf. the discussion following (8.1.2), from which it follows $x \ll x$. Conversely we have the following lemma, which should be considered as a non-intrinsic alternative to [120, Lemma 4.3], where it is shown that strongly causal Lorentzian length spaces have locally finite and locally continuous time separation τ .

Lemma 8.1.22 (Locally Finite and Locally Continuous Time Separation)

Let $K \in \mathbb{R}$ and X be a chronological Lorentzian pre-length space which admits a covering of

⁶Note that local continuity in this sense does not imply continuity, since τ may not be continuous on neighbourhoods of $(x, y) \in X \times X$ for $x \neq y$.

neighbourhoods U satisfying the property (i) from Definition 8.1.20, i.e.

$$\tau \text{ is continuous on } (U \times U) \cap \tau^{-1}([0, D_K]), \text{ which is an open set.} \quad (8.1.22)$$

Then X is covered by neighbourhoods V such that $\tau|_{V \times V}$ is finite and continuous. In particular, $\tau|_{V \times V} < D_K$.

Proof. Let X be a chronological Lorentzian pre-length space covered by neighbourhoods U as detailed above. Consider $U \subseteq X$ be a neighbourhood of some $x \in X$ which satisfies (8.1.22) for some $K \in \mathbb{R}$. Fix $\tilde{U} := (U \times U) \cap \tau^{-1}([0, D_K])$. By chronology, we have $\tau(x, x) = 0$, from which it follows that $(x, x) \in \tilde{U}$. Furthermore, \tilde{U} is open in the product topology, hence there exists an open neighbourhood $V \subseteq U \subseteq X$ of x such that $V \times V \subseteq (U \times U) \cap \tau^{-1}([0, D_K])$.⁷ In particular, $\tau|_{V \times V} < D_K$ and since τ is continuous on \tilde{U} , it is continuous on $V \times V$. If X is covered by neighbourhoods U , then there exists such a V for each $x \in X$, as required. \square

The above lemma demonstrates that the time separation function τ , on a chronological Lorentzian pre-length space with timelike curvature bound, is locally finite and locally continuous. What we may not yet do, however, is assume that τ is finite and continuous on comparison neighbourhoods a priori, as in [116, Definition 4.7]; that is, the neighbourhoods V may not be comparison neighbourhoods, as there need not exist distance realisers between points in V which remain in V (cf. (ii) of Definition 8.1.20). One way of guaranteeing that such distance realisers exist is to replace V with some causally convex set contained in V , such that τ remains finite and continuous. Since timelike diamonds are causally convex by (i) of Lemma 8.1.13 and arise naturally from the causal structure of X , they are a reasonable candidate.

Recall that the set $\mathcal{I} := \{I(x, y) \mid x, y \in X\}$ of timelike diamonds on a strongly causal Lorentzian pre-length space X form a subbase for the topology by Definition 8.1.15. In particular, any point $x \in X$ is contained in at least one timelike diamond and has non-empty past and future in X . However, there need not be a timelike diamond contained in V which contains x — to ensure that such a diamond exists, we need \mathcal{I} to be a basis for the topology. As shown in [127, Lemma 3.5], X being a *non-timelike locally isolating* Lorentzian pre-length space (cf. [126, Definition 3.2.3]) is a sufficient additional requirement to promote \mathcal{I} from subbase to basis:

Definition 8.1.23 (Non-timelike Locally Isolating Spaces)

Let X be a Lorentzian pre-length space and let $A \subseteq X$ be a subset of X . Then A is called non-future (resp. non-past) locally isolating if for all $a \in A$ with $I^+(a) \neq \emptyset$ (resp. $I^-(a) \neq \emptyset$) and

⁷Open sets of the form $V \times V$ are a basis for the product topology.

for all neighbourhoods $U_a \subseteq A$ of a there exists $b_+ \in U_a$, such that $a \ll b_+$ (resp. $b_- \in U_a$, such that $b_- \ll a$). We say A is non-timelike locally isolating if it satisfies both properties.

Note that all Lorentzian length spaces are non-timelike locally isolating as a consequence of [126, Remark 3.2.4].

Proposition 8.1.24 (Diamonds Form a Basis)

Let X be a strongly causal and non-timelike locally isolating Lorentzian pre-length space. Then the set $\mathcal{I} := \{I(x, y) \mid x, y \in X\}$ forms a basis for the topology. In particular, given any neighbourhood of any point, we can construct a timelike diamond containing the point, such that the diamond and its governing points are also contained in the neighbourhood.

Now that we have sufficient conditions under which there exist timelike diamonds within the neighbourhoods V from Lemma 8.1.22, we may provide criteria for the existence of a covering of X by comparison neighbourhoods with respect to which the time separation function is locally finite and locally continuous. The following result may be seen as a refinement of [152, Remark 2.2.12].

Lemma 8.1.25 (Local Automatic Size-Bounds)

Let X be a strongly causal, non-timelike locally isolating, and chronological Lorentzian pre-length space with timelike curvature bounded above/below by some $K \in \mathbb{R}$. Then the set

$$\tilde{\mathcal{I}} := \{W \in \mathcal{I} \mid \tau|_{W \times W} < D_K, \text{ } W \text{ is a comparison neighbourhood}\}, \quad (8.1.23)$$

of timelike diamonds W which satisfy $\tau|_{W \times W} < D_K$ and are curvature comparison neighbourhoods (in the sense of either Definition 8.1.20 or [116, Definition 4.7]), form a basis for the topology on X . In particular, X is covered by such diamonds and all timelike triangles contained in such a W satisfy size-bounds cf. Definition 8.1.11.

Proof. Let \tilde{U} be a neighbourhood of some $x \in X$. Since X has a timelike curvature bound, there exists some (\leq (resp. \geq) K)-comparison neighbourhood U (in the sense of Definition 8.1.20) containing x , that is, $x \in U \cap \tilde{U}$. Furthermore, by Lemma 8.1.22, there exists a neighbourhood $V \subseteq U$ of x on which $\tau|_{V \times V} < D_K$. Hence $x \in V \cap \tilde{U}$. Since X is strongly causal and non-timelike locally isolating, Proposition 8.1.24 implies that there exists a timelike diamond $W = I(x_-, x_+) \subset V \cap \tilde{U}$ such that $x, x_{\pm} \in V \cap \tilde{U}$ and $x_- \ll x \ll x_+$. Since $W \times W \subset V \times V$, we have $\tau|_{W \times W} < D_K$.

It remains to verify that W is a ($\geq K$)-comparison (resp. ($\leq K$)-comparison) neighbourhood in the sense of [116, Definition 4.7] (hence also in the sense of Definition 8.1.20 by Propo-

ition 8.1.21):

- (i) By construction, τ is finite on W . Furthermore, τ is continuous on $(U \times U) \cap \tau^{-1}([0, D_k])$, since U is a comparison neighbourhood, therefore τ is continuous on $W \times W$ as a subspace.
- (ii) By Lemma 8.1.13 (i), timelike diamonds are causally convex, hence distance realisers between points in W remain in W . Furthermore, since $W \subset U$ and U is D_K -geodesic, distance realisers between points in W exist, so W also is D_K -geodesic.⁸
- (iii) This coincides with property (iii) of Definition 8.1.20, which W inherits from the comparison neighbourhood U , since timelike triangles in W are also timelike triangles in U .

Therefore, for any neighbourhood \tilde{U} of any point $x \in X$, we can construct a timelike diamond $W \subset U$ whose governing points are contained in U and which is a comparison neighbourhood satisfying $\tau|_{W \times W} < D_K$. \square

Consequently, on strongly causal, non-timelike locally isolating, and chronological Lorentzian pre-length spaces X , we may choose comparison neighbourhoods on which τ is continuous and $\tau < D_K$, without loss of generality. We also have the following corollary:

Corollary 8.1.26 (Equivalence of Curvature Bounds)

A chronological, strongly causal, and non-timelike locally isolating Lorentzian pre-length space X has (local) timelike curvature bounds in the sense of Definition 8.1.20 if and only if it has (local) timelike curvature bounds in the sense of [116, Definition 4.7].

Proof. The fact that local curvature bounds in the sense of Definition 8.1.20 imply those in the sense of [116, Definition 4.7] follows directly from Lemma 8.1.25 — every neighbourhood of every point contains a timelike diamond which is a comparison neighbourhood in the latter sense.

For the converse implication, let $x \in X$ and U be a curvature comparison neighbourhood of x in the sense of [116, Definition 4.7]. Then by Proposition 8.1.24, there exists a timelike diamond $V \subseteq U$ which contains x . By a similar argument to Lemma 8.1.25, we can show that V is a comparison neighbourhood in the sense of [116, Definition 4.7]. Furthermore, since timelike diamonds are open by strong causality, Proposition 8.1.21 yields that V is a comparison neighbourhood in the sense of Definition 8.1.20. \square

However, a (\leq (resp. \geq) K)-comparison neighbourhood in the sense of Definition 8.1.20 need not be a comparison neighbourhood in the sense of [116, Definition 4.7], it must merely contain one. In particular, it is possible for X to have global curvature bounds in the former

⁸While V also satisfies (i) and (iii), it need not be causally convex, so distance realisers between points in V need not stay in V , hence the need to construct a diamond where we can guarantee this property.

sense but not the latter, since, even with our additional assumptions, Definition 8.1.20 only requires X to be D_K -geodesic and contain a neighbourhood on which τ is finite and continuous, as opposed to being geodesic with (globally) finite and continuous τ . Most importantly for the globalisation theorems, while chronological, strongly causal, non-timelike locally isolating Lorentzian pre-length spaces with curvature bounded above/below by K admit coverings of comparison neighbourhoods in which triangles automatically satisfy size-bounds, arbitrary triangles not contained in one of these neighbourhoods may still not satisfy size-bounds in general.

Remark 8.1.27 (Global Curvature Bounds and Anti-de Sitter Space)

Let us return to the problem with the exotic behaviour exhibited by anti-de Sitter spaces, namely that geodesics (in the smooth sense) in $AdS(\frac{1}{\sqrt{-K}})$ stop being longest curves when they exceed length $\frac{\pi}{\sqrt{-K}}$. Before we proceed, recall that geodesics in anti-de Sitter space arise from the intersection of anti-de Sitter space with a plane which passes through the origin of the ambient space $\mathbb{R}^{2,1}$ and the endpoints of the geodesic [153, Figure 11].

We may visualise this non-maximising behaviour by considering a pair of non-antipodal points x, y on the skirt of $AdS(\frac{1}{\sqrt{-K}})$ and connecting them via the longer timelike geodesic between them. We may then create timelike curves of increasing length (given by the Lorentzian metric as in (6.0.1)) by going out from x towards infinity on one side of the skirt and then returning to y , see [154, Figure 1]. In particular, the time separation travelled may be made arbitrarily large by travelling further out along curves which are increasingly null. Consequently, $\tau(x, y) = \infty$ and there is no longest curve/distance realiser between x and y .⁹ Furthermore, if we let \tilde{y} be the antipodal point to x , then there are two timelike distance realisers of length $\frac{\pi}{\sqrt{-K}}$ between them, along the skirt of the space. Hence, perturbing from \tilde{y} to some non-antipodal y on the skirt of the space, the time separation jumps from $\frac{\pi}{\sqrt{-K}}$ to ∞ .

Therefore, τ is not finite or continuous on anti-de Sitter space, and distance realisers between points $D_K = \frac{\pi}{\sqrt{-K}}$ or further apart need not exist or be unique. It follows that the anti-de Sitter spaces do not satisfy a global curvature bound with respect to [116, Definition 4.7], but since we only consider points closer than the finite diameter D_K , anti-de Sitter space $AdS(\frac{1}{\sqrt{-K}})$ does satisfy curvature bound both above and below by $-K$ in the sense of Definition 8.1.20.

This is similar to the behaviour described by the maximal diameter theorem in Riemannian geometry [155, Theorem 6.5] in the metric setting. More precisely, in the metric model spaces M_k with $k > 0$, there exist infinitely many geodesics between points which are exactly a distance $\frac{\pi}{\sqrt{k}}$ apart (compare the antipodal points on AdS with those on the sphere).

⁹The path consisting of two timelike rays of infinite length to and from infinity fails to be a distance realiser as it is not Lipschitz continuous and therefore fails to be a curve.

△

Before we turn our attention to curvature bounds defined in terms of angles, let us comment on how to parametrise distance realisers in Lorentzian pre-length spaces. In [116, Section 3.7], it is shown that for chronological Lorentzian pre-length spaces X with (locally) continuous time separation function, any timelike distance realiser $\gamma : [a, b] \rightarrow X$ with finite τ -length $L_\tau(\gamma) := L$ admits τ -arclength parametrisation. That is, there exists some $\sigma : [0, L] \rightarrow [a, b]$ such that $\gamma \circ \sigma : [0, L] \rightarrow X$ and

$$L_\tau(\gamma|_{[\sigma(t), \sigma(s)]}) = \tau(\gamma(\sigma(t)), \gamma(\sigma(s))) = |s - t|, \quad (8.1.24)$$

for all $s < t \in [0, L]$. As in the metric setting (cf. the discussion following Definition 7.1.3), we may scale this parametrisation σ such that $\gamma \circ \sigma$ is instead parametrised with constant τ -speed on $[0, 1]$ and

$$L_\tau(\gamma|_{[\sigma(t), \sigma(s)]}) = \tau(\gamma(\sigma(t)), \gamma(\sigma(s))) = \tau(\gamma(a), \gamma(b))|s - t|. \quad (8.1.25)$$

These parametrisations $\gamma \circ \sigma$ have the same τ -length and causal character as γ by [116, Lemma 2.28] and [116, Lemma 3.32] respectively. However, they may not be locally 1-Lipschitz with respect to the distance function d (despite being 1-Lipschitz with respect to τ). Therefore, $\gamma \circ \sigma$ may not be a causal curve in the sense of Definition 8.1.3.

Since Lemma 8.1.22 implies τ is locally continuous on chronological Lorentzian pre-length spaces with timelike curvature bound and Definition 8.1.20 only considers points with finite time separation, unless stated otherwise we assume that timelike distance realisers on chronological Lorentzian pre-length spaces with curvature bound are constant τ -speed parametrised on $[0, 1]$. In particular, let X be a strongly causal, non-timelike locally isolating, and chronological Lorentzian pre-length space with timelike curvature bound. Then timelike triangles in X which are contained within a comparison neighbourhood and whose sides are given by timelike (not just causal) distance realisers, can be assumed to have sides parametrised by constant τ -speed on $[0, 1]$, due to Lemma 8.1.25.

As observed in [116, Section 4.5], there is no such natural parametrisation on $[0, 1]$ for causal curves with null pieces. Hence, it is important to highlight here that even if two points $x, y \in X$ with $x \ll y$ are joined by a distance realiser γ_{xy} , this does not mean that we have $x \ll p \ll y$ for all $p \in \gamma$, since γ_{xy} will be a timelike curve if and only if it is τ -rectifiable, see our discussion following Definition 8.1.5. In particular, distance realisers between timelike related points may contain a null piece and therefore merely be causal curves, as shown by [116, Example 3.19].

In order to ensure that the sides of timelike triangles are in fact timelike, we will consider Lorentzian pre-length spaces which are well-behaved, in the sense that distance realisers between timelike related points are always timelike. In order to encode this as a property of the space, we now introduce the notion of regularity.

Definition 8.1.28 (Regular Lorentzian pre-length space)

Let X be a Lorentzian pre-length space. We call X regular if for all $x, y \in X$ such that $x \ll y$, all distance realisers joining x and y are timelike. That is, distance realisers between timelike related points do not contain a null piece.

8.2 Angles in Lorentzian Pre-length Spaces and Curvature Bounds

Recall from the discussion following [Definition 7.1.12](#) that, in the metric setting, there are several different comparison conditions we can use to define curvature bounds, with each alternative involving the replacement of the constraint on points $p, q \in \Delta(x, y, z)$ on the sides of triangles with a constraint on some other geometric quantity. Furthermore, these conditions are shown to be equivalent under the assumptions of the globalisation theorems, cf. [Section 7.2](#).

Analogously, in this section we introduce several definitions of timelike curvature bounds on Lorentzian pre-length spaces, drawing from the pool of definitions collated in [\[123\]](#) and showing that under relatively mild conditions which we have already introduced, these definitions are equivalent. We limit ourselves to discussing only the definitions which will be used in the globalisation theorems, namely curvature bounds via angle and hinge comparison. Before we do so, however, let us first discuss how to define an angle in a Lorentzian pre-length space.

8.2.1 Comparison and Hyperbolic Angles

The concept of a (hyperbolic) angle in a Lorentzian pre-length space was independently introduced in [\[124, 125\]](#) with the intention of providing further characterisation of such spaces under curvature bounds. In what follows, we predominantly follow the conventions of [\[123, 125\]](#) for parity with the author's papers [\[2, 3\]](#).

In this subsection, we will refer to certain triples of points as vertices of an admissible causal triangles, however it will not be necessary for the distance realisers between the vertices to exist when we do so. We use this terminology because the triples of points required to define an angle on a Lorentzian pre-length space need to satisfy the same causal relations as the vertices of an admissible causal triangle. In particular, since we assume the vertices of admissible causal triangles satisfy size-bounds for $K \in \mathbb{R}$, a comparison triangle exists in $\mathbb{L}^2(K)$ cf. [Definition 8.1.11](#). Furthermore, in [Chapter 9](#) and [Chapter 10](#), we will generally work in a

setting where distance realisers between causally related points do exist, in which case these points define an admissible causal triangle as expected.

We begin with the observation that hyperbolic angles are defined in the model spaces $\mathbb{L}^2(K)$ when they are interpreted as smooth spacetimes via [Example 8.1.2](#). As such, we may define the angle at some vertex of an admissible causal triangle in a Lorentzian pre-length space X , to be the angle at the corresponding vertex of the comparison triangle in a given model space $\mathbb{L}^2(K)$. We call this type of angle a K -comparison angle:

Definition 8.2.1 (K -Comparison Angle)

Let X be a Lorentzian pre-length space, $K \in \mathbb{R}$ and x, y, z (in the stated order) be the vertices of an admissible causal triangle in X . Let $\Delta(\bar{x}, \bar{y}, \bar{z})$ be a comparison triangle in $\mathbb{L}^2(K)$ for $\Delta(x, y, z)$ and assume that the sides adjacent to \bar{x} are timelike (analogously for \bar{y} and \bar{z}). Then:

- (i) The K -comparison angle at x is given by the hyperbolic angle at \bar{x} and between \bar{y}, \bar{z} in $\mathbb{L}^2(K)$. Namely:

$$\tilde{\chi}_x^K(y, z) := \chi_{\bar{x}}^{\mathbb{L}^2(K)}(\bar{y}, \bar{z}) = \operatorname{arcosh}(|g(\gamma'_{\bar{x}\bar{y}}(0), \gamma'_{\bar{x}\bar{z}}(0))|), \quad (8.2.1)$$

where g is the Lorentzian metric on $\mathbb{L}^2(K)$, the distance realisers $\gamma_{\bar{x}\bar{y}}, \gamma_{\bar{x}\bar{z}}$ are unit speed parametrised so that their derivatives are normalised, and $\gamma_{\bar{x}\bar{y}}(0) = \gamma_{\bar{y}\bar{z}}(0) = \bar{x}$.

- (ii) The sign σ of the K -comparison angle at x is given by the sign of $g(\gamma'_{\bar{x}\bar{y}}(0), \gamma'_{\bar{x}\bar{z}}(0))$.

- (iii) The signed K -comparison angle at x is then

$$\tilde{\chi}_x^{K,S}(y, z) := \sigma \tilde{\chi}_x^K(y, z). \quad (8.2.2)$$

K -comparison angles are not defined at vertices of a triangle which are adjacent to a non-timelike side. Furthermore, the sign of the K -comparison angle at x or z (where it is defined) is always $\sigma = -1$, while the sign of the K -comparison angle at y is always $\sigma = +1$.

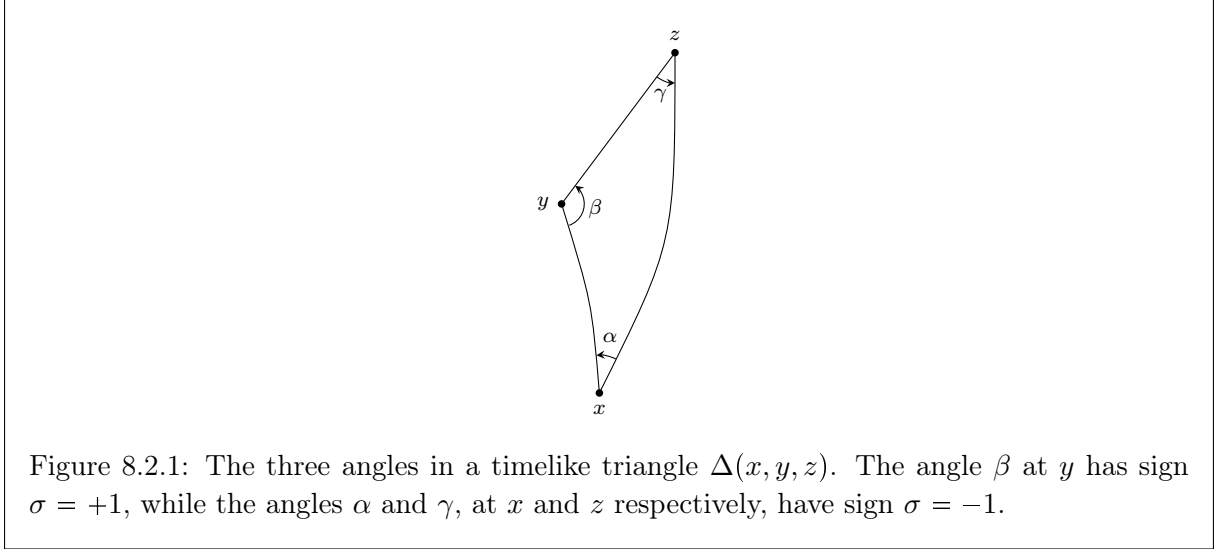
The following two results are direct corollaries of the law of (hyperbolic) cosines [[125](#), [Lemma 2.3](#)] and provide us with several properties of K -comparison angles which will be useful throughout the remainder of this thesis.

Corollary 8.2.2 (Law of (Hyperbolic) Cosines is Monotone)

Consider an admissible causal triangle $\bar{\Delta}$ in the Lorentzian model space $\mathbb{L}^2(K)$, for some $K \in \mathbb{R}$.

- (i) If the τ -lengths of the two short sides are fixed, then hyperbolic angles in $\bar{\Delta}$ are strictly monotonically increasing with the τ -length of the long side.

- (ii) If the τ -lengths of the longest side and one of the short sides are fixed, then hyperbolic angles in $\bar{\Delta}$ are strictly monotonically decreasing with the τ -length of the remaining side.



Corollary 8.2.3 (Non-degeneracy is Equivalent to Non-Zero Comparison Angles)

Let X be a Lorentzian pre-length space, $K \in \mathbb{R}$, and $\Delta(x, y, z)$ be an admissible causal triangle in X satisfying size-bounds for K . The triangle $\Delta(x, y, z)$ is non-degenerate if and only if the K -comparison angles at each of its vertices are positive, where they are defined. In particular, if any of the K -comparison angles are zero, then all of the K -comparison angles are zero and the triangle is degenerate.

Observe that the K -comparison angle at x depends on the points y and z chosen in the formation of the admissible causal triangle $\Delta(x, y, z)$. In order to define the angle between two timelike curves, we consider limit of the K -comparison angles in a sequence of such triangles, with the points y and z getting progressively closer to x .

Definition 8.2.4 (Angles)

Let X be a Lorentzian pre-length space and let $\alpha : [0, a] \rightarrow X$, $\beta : [0, b] \rightarrow X$ be two timelike curves of arbitrary time orientation emanating from $x := \alpha(0) = \beta(0)$.

- (i) The upper angle between α and β at x is given by

$$\angle_x(\alpha, \beta) := \limsup_{\substack{(s,t) \in D \\ s,t \rightarrow 0}} \tilde{\angle}_x^K(\alpha(s), \beta(t)), \quad (8.2.3)$$

where $D \subseteq (0, a] \times (0, b]$ consists of the pairs (s, t) such that $x, \alpha(s), \beta(t)$, in some order,

are the vertices of an admissible causal triangle.

(ii) If the upper angle is finite and given by the limit

$$\angle_x(\alpha, \beta) = \lim_{\substack{(s,t) \in D \\ s,t \rightarrow 0}} \tilde{\angle}_x^K(\alpha(s), \beta(t)), \quad (8.2.4)$$

then we call $\angle_x(\alpha, \beta)$ the angle between α and β at x .

It is shown in [125, Proposition 2.14] that the limit (8.2.4) is independent of the choice of model space, hence from here onward we drop the decoration K in K -comparison angle and simply write $\tilde{\angle}_x(y, z)$, where the chosen model space will be clear from context.

The sign of the comparison angle $\tilde{\angle}_x(\alpha(s), \beta(t))$ is independent of the choice of $(s, t) \in D$, since the curves α and β are timelike and therefore of fixed time-orientation. Therefore, we define the sign σ of the (upper) angle to be the sign of the corresponding set of comparison angles. The signed (upper) angle is then defined as $\angle_x^S(\alpha, \beta) := \sigma \angle_x(\alpha, \beta)$. As for comparison angles, $\sigma = -1$ if α and β have the same time orientation and 1 otherwise.

We also introduce the notion of a hinge, which is, in essence, an angle accompanied by its adjacent curves:

Definition 8.2.5 (Hinges)

Let X be a Lorentzian pre-length space and let $\alpha : [0, a] \rightarrow X$, $\beta : [0, b] \rightarrow X$ be two timelike distance realisers of arbitrary time orientation emanating at $\alpha(0) = \beta(0) =: x$.

(i) The distance realisers α, β along with their associated angle $\theta := \angle_x(\alpha, \beta)$ at x (if it exists) are called a hinge which we denote by $(\alpha, \beta; \theta)$.

(ii) Let (α, β) be a hinge at $x \in X$ with finite angle $\theta := \angle_x(\alpha, \beta)$ and $K \in \mathbb{R}$. A comparison hinge for $(\alpha, \beta; \theta)$ in $\mathbb{L}^2(K)$ is a hinge $(\tilde{\alpha}, \tilde{\beta}; \tilde{\theta})$ in $\mathbb{L}^2(K)$ consisting of two timelike distance realisers $\tilde{\alpha}$ and $\tilde{\beta}$ with the same time orientation and τ -length as α and β respectively, such that $\tilde{\alpha}(0) = \tilde{\beta}(0) = \tilde{x}$ and the angle $\tilde{\theta} := \angle_{\tilde{x}}^{\mathbb{L}^2(K)}(\tilde{\alpha}, \tilde{\beta}) = \angle_x(\alpha, \beta) = \theta$.

We conclude this subsection with the following useful fact about angles, which provides sufficient conditions for adjacent angles along a distance realiser to have equal magnitude. This property was proven in [125, Corollary 4.6], with existence of the two angles proven in [125, Lemma 4.10], and is analogous to angles along distance realisers in metric geometry summing to π , cf. [133, Lemma 1.3]. We call such distance realisers balanced.

Proposition 8.2.6 (Balanced Distance Realisers)

Let X be a strongly causal and locally causally closed Lorentzian pre-length space with timelike

curvature bounded below by $K \in \mathbb{R}$, and let $\alpha : [0, 1] \rightarrow X$ be a timelike distance realiser. Let $x = \alpha(t)$ for $t \in (0, 1)$ and consider the restrictions $\alpha_-(s) = \alpha|_{[0, t]}(t-s)$ and $\alpha_+(s) = \alpha|_{[t, 1]}(t+s)$ as past-directed and future-directed distance realisers emanating from x , respectively. Let β be a timelike distance realiser emanating from x . Then $\angle_x(\alpha_-, \beta) = \angle_x(\alpha_+, \beta)$.

8.2.2 Alternative Definitions of Timelike Curvature Bounds

As previously discussed, we now introduce alternative definitions of timelike curvature bounds which feature comparison conditions in terms of angles and hinges. Each of these definitions has been introduced in the context of Lorentzian pre-length spaces in earlier works (predominantly [125] and [124]) and takes inspiration from the associated metric statement. We direct the curious reader to [123] for details regarding the wide array of available comparison conditions and follow the formulation provided by op. cit. here. We begin by stating the definition of timelike curvature bounds via angle comparison.

Definition 8.2.7 (Curvature Bounds by Angle Comparison)

Let X be a regular Lorentzian pre-length space and $K \in \mathbb{R}$. An open subset U is called a $(\geq K)$ - (resp. $(\leq K)$ -) comparison neighbourhood in the sense of angle comparison if:

- (i) τ is continuous on $(U \times U) \cap \tau^{-1}([0, D_K])$ and $(U \times U) \cap \tau^{-1}([0, D_K])$ is open.
- (ii) U is D_K -geodesic.
- (iii) Let $\alpha : [0, a] \rightarrow U$ and $\beta : [0, b] \rightarrow U$ be timelike distance realisers with arbitrary time orientation, such that $x := \alpha(0) = \beta(0)$ and $x, \alpha(a), \beta(b)$, in some causal order, are the vertices of an admissible causal triangle in U . Then we require:

$$\angle_x^S(\alpha, \beta) \leq \tilde{\angle}_x^{K,S}(\alpha(a), \beta(b)) \quad (\text{resp. } \angle_x^S(\alpha, \beta) \geq \tilde{\angle}_x^{K,S}(\alpha(a), \beta(b))). \quad (8.2.5)$$

- (iv) For $(\geq K)$ -comparison neighbourhoods only: If $\alpha, \beta, \gamma : [0, \varepsilon] \rightarrow U$ are three timelike distance realisers of finite τ -length, with $x := \alpha(0) = \beta(0) = \gamma(0)$ and α, γ both having the opposite time orientation to β , then the following triangle inequality of angles must hold:

$$\angle_x(\alpha, \gamma) \leq \angle_x(\alpha, \beta) + \angle_x(\beta, \gamma). \quad (8.2.6)$$

We say that X has timelike curvature bounded below by K (resp. timelike curvature bounded above by K) in the sense of angle comparison if every point in X has a $(\geq K)$ -comparison (resp. $(\leq K)$ -comparison) neighbourhood in the sense of angle comparison.

If X itself is a $(\geq K)$ -comparison (resp. $(\leq K)$ -comparison) neighbourhood in the sense of

angle comparison, then we say that X has global timelike curvature bounded below (resp. above) by K in the sense of angle comparison.

Observe that points (i) and (ii) in the definition above are precisely the first two requirements for a space to have curvature bounds in the sense of Definition 8.1.20, while (iii) and regularity of X (alongside (iv) in the lower curvature bound case) are imposed in place of the third requirement in our earlier definition. As we shall see shortly, it is possible to write down conditions under which these constraints also coincide.

The astute reader may notice that, due to the similarity between the definition of angles and hinges, we may also define curvature bounds using the latter by replacing point (iii) in Definition 8.2.7:

Definition 8.2.8 (Curvature Bounds by Hinge Comparison)

Let X be a regular Lorentzian pre-length space and $K \in \mathbb{R}$. An open subset U is called a $(\geq K)$ - (resp. $(\leq K)$ -)comparison neighbourhood in the sense of hinge comparison if:

- (i) τ is continuous on $(U \times U) \cap \tau^{-1}([0, D_K])$ and $(U \times U) \cap \tau^{-1}([0, D_K])$ is open.
- (ii) U is D_K -geodesic.
- (iii) Let $\alpha : [0, a] \rightarrow U$, $\beta : [0, b] \rightarrow U$ be distance realisers such that

$$L_\tau(\alpha), L_\tau(\beta), \tau(\alpha(a), \beta(b)), \tau(\beta(b), \alpha(a)) < D_K \quad (8.2.7)$$

and $\theta := \angle_x(\alpha, \beta)$ is finite.¹⁰ Let $(\tilde{\alpha}, \tilde{\beta}; \tilde{\theta})$ be a comparison hinge for $(\alpha, \beta; \theta)$ in $\mathbb{L}^2(K)$. Then

$$\tau(\alpha(a), \beta(b)) \geq \tau(\tilde{\alpha}(a), \tilde{\beta}(b)) \quad (\text{resp. } \tau(\alpha(a), \beta(b)) \leq \tau(\tilde{\alpha}(a), \tilde{\beta}(b))). \quad (8.2.8)$$

- (iv) For $(\geq K)$ -comparison neighbourhoods only: If $\alpha, \beta, \gamma : [0, \varepsilon] \rightarrow U$ are three timelike distance realisers of finite τ -length, with $x := \alpha(0) = \beta(0) = \gamma(0)$ and α, γ both having the opposite time orientation to β , then (8.2.6) must hold.

(Global) timelike curvature bounded below/above by K in the sense of hinge comparison is then defined in complete analogy to Definition 8.2.7.

The conditions (8.2.7) amount to assuming that, should $x, \alpha(a), \beta(b)$ form an admissible

¹⁰[123, Definition 3.14] introduces a minor technical assumption in the limiting case where this angle is infinite: in the $(\geq K)$ - (resp. $(\leq K)$ -)comparison case, $\angle_x(\alpha, \beta)$ is required to be finite if α and β have different (resp. the same) time orientations. While this prevents certain pathological spaces from exhibiting curvature bounds, we omit this assumption from the statement for simplicity.

causal triangle, said triangle must satisfy size-bounds cf. Definition 8.1.11, though (8.2.8) holds even if $\alpha(a)$ and $\beta(b)$ are not causally related in either direction. In this case, α and β have the same time orientation and (8.2.5) is trivially satisfied. The relationship between (8.2.8) and (8.2.5) is the subject of [119, Lemma 2.2] and is associated to the monotonicity of the law of (hyperbolic) cosines.

In order to ensure that we are always working in a setting where we may interchange between Definition 8.1.20, Definition 8.2.7, and Definition 8.2.8 freely, we now state conditions under which these definitions are equivalent. The following proposition is an abridged version of [123, Theorem 5.1], which shows the equivalence of a far greater number of comparison conditions:

Proposition 8.2.9 (Equivalence of Curvature Bounds)

Let X be a regular and chronological Lorentzian pre-length space.

- (i) ($\leq K$)-case: Triangle, Hinge, and Angle comparison are equivalent.*
- (ii) ($\geq K$)-case: Angle comparison and Hinge comparison are equivalent. Furthermore, Angle comparison implies Triangle comparison and if X satisfies the triangle inequality for angles (8.2.6), then all three definitions of curvature bounds coincide.*

Remark 8.2.10 (Required Assumptions)

Following [123, Figure 4], the reader could be forgiven for expecting the above proposition to require X to be strongly causal and locally D_K -geodesic, in order for angle comparison to imply triangle comparison. Looking more closely, op. cit proves this implication by passing through two other comparison conditions, namely the four-point condition and the monotonicity condition, loosely stating that

$$\text{angle comparison neighbourhood} \Rightarrow \text{four-point comparison} \tag{8.2.9}$$

$$\Rightarrow \text{monotonicity comparison neighbourhood} \tag{8.2.10}$$

$$\Rightarrow \text{triangle comparison neighbourhood.} \tag{8.2.11}$$

However, in [123], these three steps are decoupled and the properties of the angle comparison neighbourhood in the first step [123, Proposition 4.15] are forgotten prior to the second step. The additional assumptions are then brought in for the second step [123, Proposition 4.17, 4.18] to ensure that the four-point comparison neighbourhood satisfies property (ii) of a monotonicity comparison neighbourhood [123, Definition 3.8]. If we instead choose to only consider four-point comparison neighbourhoods coming from angle comparison neighbourhoods, as in the first step, then our comparison neighbourhoods automatically satisfy (ii) of Definition 8.2.7, which is

precisely property (ii) of a monotonicity comparison neighbourhood [123, Definition 3.8].

Regarding the remaining assumptions, considering chronological Lorentzian pre-length spaces excludes spaces with time separation $\tau \equiv \infty$, which would otherwise trivially satisfy curvature bounds, since no points would satisfy $\tau(x, y) < D_K$ and hence there would be no triples of points admitting comparison triangles. Regularity is required by the definitions of both angle and hinge comparison comparison. Finally, in contrast to the metric setting, we do not require τ to be intrinsic in the curvature bounded below case, instead imposing the triangle inequality of angles (8.2.6), which is a consequence of the causal character of curves (namely that causal curves emanating from a point may either be in the past or future of the point). Most interestingly, there is not a substantial difference between the assumptions required for curvature bounded above compared to curvature bounded below and the assumptions we impose are relatively mild (for the comparison conditions chosen).

△

8.3 Chapter Summary

In this chapter, we introduced the concept of a Lorentzian pre-length space and its associated time separation function, comparing and contrasting these objects to their Riemannian analogues, the metric space and distance function. In [Example 8.1.2](#), we illustrated that smooth spacetimes are examples of Lorentzian pre-length spaces. However, we note here that not all spacetimes with continuous Lorentzian metrics are Lorentzian pre-length spaces; in particular, the time separation (8.1.1) is not lower semi-continuous on so-called bubbling spaces [116, Example 5.2]. Hence, Lorentzian pre-length spaces are a generalisation of smooth spacetimes which include lower regularity structures but exclude some pathological behaviour [156].

In addition to summarising some key properties of Lorentzian pre-length spaces from the literature, in [Section 8.1.2](#), we provided novel modifications to the definition of curvature bounds first detailed in [116, Definition 4.7], which have since been developed further by the author's collaborators in [123]. These modifications ensure that all of the criteria of a comparison neighbourhood respect the finite (timelike) diameter of the space, so that only points whose separation is realised by the length of a distance realiser in a given model space are considered. This brings the definition in-line with the metric setting and the behaviour of Anti-de Sitter space now mimics that of the sphere. Furthermore, we provided an alternative to [120, Lemma 4.3] in the form of [Lemma 8.1.22](#), demonstrating that chronological Lorentzian pre-length spaces with curvature bounds in our new sense still retain locally finite and locally continuous time separation, though not necessarily on comparison neighbourhoods. In particular, timelike distance realisers may be

τ -arclength parametrised on such spaces. In the second half of the chapter, we reviewed the concept of hyperbolic angles on Lorentzian pre-length spaces as well as how they can be used to define timelike curvature bounds in the sense of angle and hinge comparison, following the work of [125]. Finally, in Proposition 8.2.9 and Remark 8.2.10, we commented on how the assumptions of the equivalence of definitions result quoted in [3, Proposition 2.7] and derived in [123, Theorem 5.1] may be relaxed for the comparison conditions we consider.

In the next two chapters we will derive conditions under which (local) timelike curvature bounds on Lorentzian pre-length spaces imply global timelike curvature bounds on the space, beginning with the case of timelike curvature bounded above. Our method for deriving this result will follow the approach of Alexandrov's Patchwork and hence require the so called gluing lemma in order to construct arbitrarily large timelike triangles satisfying curvature bounds from smaller triangles with the same bound. We shall also introduce the geodesic map for Lorentzian pre-length spaces as a means of measuring the variation of distance realisers.



9

Globalisation of Timelike Curvature Bounded Above

As previously discussed, the goal of this part of the thesis is to derive conditions under which timelike curvature bounds on Lorentzian pre-length spaces globalise. In this chapter, we prove the first of these results by following the approach of Alexandrov’s patchwork [Theorem 7.2.2](#). Loosely speaking, the proof considers arbitrarily large triangles in a space with timelike curvature bounded above by some $K \in \mathbb{R}$ and decomposes them into triangles which are contained inside comparison neighbourhoods. These triangles therefore satisfy timelike curvature bounds and may be “glued” together to show that the large triangle also satisfies the bounds. Recall that, in the metric setting, it is necessary and sufficient to assume that the space is uniquely $\text{diam}(M^2(k))$ -geodesic and distance realisers vary continuously with their endpoints.

This approach translates fairly directly from the metric to the synthetic Lorentzian setting, with surprisingly few obstacles once a suitable decomposition has been found. There are, however, some technical details to be aware of in the choice of decomposition. In particular, given an arbitrary covering of a triangle by comparison neighbourhoods, it may not be possible to obtain a decomposition consisting of timelike triangles where each triangle is contained within one of the comparison neighbourhoods. Both the covering and the sub-triangles need to be carefully chosen.

9.1 Preliminary Results

Let us begin our derivation by providing a number of preliminary lemmata, starting with the so called gluing lemma for timelike triangles and then moving on to discuss what it means for a distance realiser in a Lorentzian pre-length space to vary continuously with its endpoints.

9.1.1 The Gluing Lemma for Curvature Bounded Above

Gluing techniques for Lorentzian pre-length spaces were first introduced in [126] and further studied in the context of the causal ladder in [127]. The main focus of the former paper was a Lorentzian analogue to the Reshetnyak gluing theorem [126, Theorem 5.2.1], which states that the amalgamation (along non-timelike locally isolating and closed subsets) of two strongly causal, smooth spacetimes with timelike curvature bounded above by $K \in \mathbb{R}$ is a Lorentzian pre-length space with timelike curvature bounded above by K .

The main tool used in the proof of [126, Theorem 5.2.1] is the gluing lemma [126, Lemma 4.3.1, Corollary 4.3.2], which states that if a timelike triangle can be decomposed into two sub-triangles satisfying a curvature bound, then the original triangle also satisfies said bounds. This is precisely the property we require. In the following statement of the gluing lemma, recall that timelike triangles are assumed to satisfy size-bounds for $\mathbb{L}^2(K)$, such that the required comparison triangles exist.

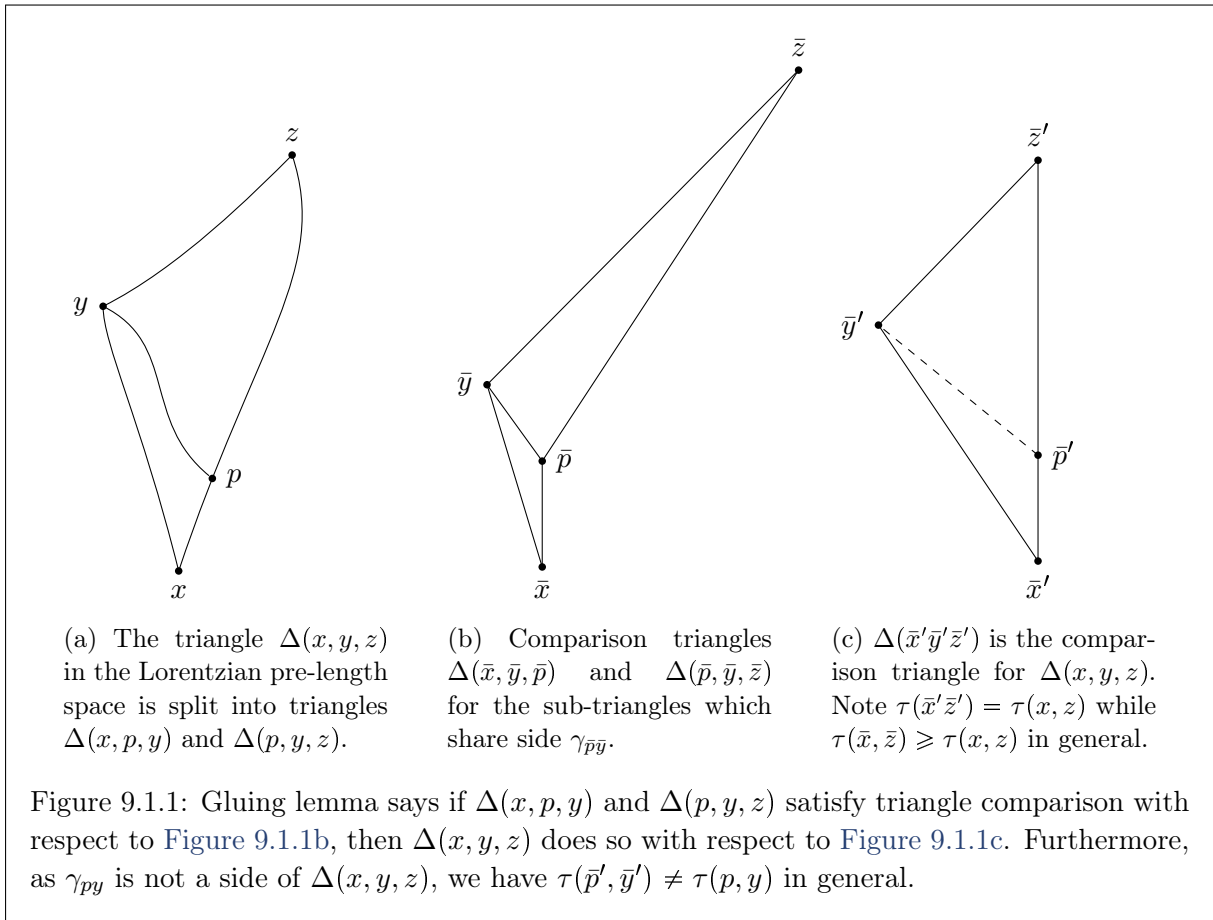
Lemma 9.1.1 (Gluing Lemma: Curvature Bounded Above)

Let $K \in \mathbb{R}$, X be a Lorentzian pre-length space, and $U \subseteq X$ be an open subset satisfying points (i) and (ii) in Definition 8.1.20. Let $\Delta(x, y, z)$ be a timelike triangle in U and fix a point $p \in \gamma_{xz}$ with $p \ll y$, such that $\Delta(x, p, y)$ and $\Delta(p, y, z)$ are also timelike triangles, see Figure 9.1.1. If $\Delta(x, p, y)$ and $\Delta(p, y, z)$ have timelike curvature bounded above by K , i.e. they satisfy the second inequality in (8.1.21), then the same is true for $\Delta(x, y, z)$.

Furthermore, the analogous statement holds if we fix $p \in \gamma_{xz}$ with $y \ll p$, or if p is on either of the shorter two sides (in which case we automatically have the timelike relation $p \ll z$ or $x \ll p$ between p and the opposite vertex, by (8.1.2)).

While the above lemma was originally stated with the definition [116, Definition 4.7] of curvature bounds in mind, the proof remains valid using Definition 8.1.20, where we assume U is chronological so that τ remains finite and continuous in a neighbourhood of each point. The proof is rather lengthy however, so we choose not to re-derive the result here.

We now know that, given a decomposition of an arbitrary timelike triangle into multiple sub-triangles, each of which are contained in an $(\leq K)$ -comparison neighbourhood for some fixed $K \in \mathbb{R}$, then the original triangle must satisfy (iii) of Definition 8.1.20. Therefore, it remains to find conditions which guarantee that the sub-triangles can each be chosen to be within a comparison neighbourhood. Lemma 8.1.25 already gives us a solution — we can break down the starting triangle in such a way that the sub-triangles are contained in comparison neighbourhoods given by timelike diamonds on which $\tau < D_K$, since such diamonds form a (neighbourhood) basis.



9.1.2 Continuously Varying Distance Realisers

Recall that in the metric Alexandrov Patchwork Theorem 7.2.2, we required the notion of continuous variation of distance realisers with their endpoints, in the sense of Definition 7.2.3. In essence, the proof creates a continuous variation of distance realisers in a triangle $\Delta(x, y, z)$ by fixing one endpoint, say x and varying the other endpoint along γ_{yz} , which then allows the triangle to be broken up into a fan of sub-triangles, see Figure 9.2.1.

In order to introduce an analogous notion for distance realisers in Lorentzian pre-length spaces, we must first bear in mind their causal character. Given a timelike triangle $\Delta(x, y, z)$ and a point $p \in \gamma_{yz}$, we see in Figure 9.1.1 that $x \ll p$ is automatic. Hence, we only need to consider continuous variation of distance realisers between timelike related points. While we assume that X is uniquely ℓ -geodesic to mirror the metric setting, we also require the space to be regular. This will ensure that $\Delta(x, y, p)$ and $\Delta(x, p, z)$ possess timelike sides in a later step of the proof.

Definition 9.1.2 (Continuously Varying Timelike Distance Realisers)

Let X be a uniquely ℓ -geodesic, regular Lorentzian pre-length space and let $x, y \in X$ be a pair of points such that $x \ll y$ and $\tau(x, y) < \ell$. The distance realiser γ_{xy} from x to y in X is said to vary continuously with its endpoints if, for every pair of sequences $x_n \rightarrow x$ and $y_n \rightarrow y$ such that there exists a unique distance realiser $\gamma_{x_n y_n}$ and $x_n \ll y_n$ for each $n \in \mathbb{N}$, we have $\gamma_{x_n y_n} \rightarrow \gamma_{xy}$ uniformly.

Since we will want to select specific distance realisers from our continuous variation, having a description of the variation in terms of a function of the endpoints and curve parameter will be useful. It is with this in mind that we introduce the so-called geodesic map, which is closely related to the *line-of-sight map* defined for metric spaces in [100, Definition 9.32], however we allow both endpoints of the curve to vary, so long as they remain timelike related (otherwise we would require separate line-of-sight maps for the past and future of the fixed point). Continuous variation of distance realisers can then be written in terms of the continuity of the geodesic map.

Definition 9.1.3 (Geodesic Map)

Let X be a chronological, regular Lorentzian pre-length space with locally continuous τ . The geodesic map into X is then formally defined as

$$G : \tau^{-1}((0, \infty)) \times [0, 1] \rightarrow X, \quad G(x, y, t) := \gamma_{xy}(t). \quad (9.1.1)$$

when there exists a unique timelike distance realiser γ_{xy} from x to y . Here $\gamma_{xy}(t)$ is taken in constant τ -speed parametrisation on $[0, 1]$.

In the above definition, the subset $\tau^{-1}((0, \infty)) \subseteq X \times X$ is the set of pairs of timelike related points (x, y) with finite time separation $\tau(x, y)$. Under the assumption that X is uniquely ℓ -geodesic, G is defined on the whole of $\tau^{-1}((0, \ell))$.

Proposition 9.1.4 (Equivalent Notions of Continuous Variation)

Let X be a uniquely ℓ -geodesic, chronological, regular Lorentzian pre-length space with locally continuous τ and assume that $\tau^{-1}([0, \ell])$ is open. Then $G|_{\tau^{-1}((0, \ell)) \times [0, 1]}$ is continuous if and only if timelike distance realisers with τ -length less than ℓ vary continuously with their endpoints in the sense of Definition 9.1.2.

Proof. We begin by proving the reverse direction and assume that timelike distance realisers with τ -length less than ℓ vary continuously with their endpoints. We want to show that $G(x_n, y_n, t_n) \rightarrow G(x, y, t)$ for all sequences $x_n \rightarrow x$, $y_n \rightarrow y$, $t_n \rightarrow t$ such that $(x_n, y_n) \in \tau^{-1}((0, \ell))$. This amounts to showing $\gamma_{x_n y_n}(t_n) \rightarrow \gamma_{xy}(t)$ by (9.1.1). By triangle inequality on

the underlying distance function d , we have

$$d(\gamma_{x_n y_n}(t_n), \gamma_{xy}(t)) \leq d(\gamma_{x_n y_n}(t_n), \gamma_{xy}(t_n)) + d(\gamma_{xy}(t_n), \gamma_{xy}(t)). \quad (9.1.2)$$

Since γ_{xy} varies continuously with its endpoints, it follows that $\gamma_{x_n y_n} \rightarrow \gamma_{xy}$ uniformly, hence the first term on the right of (9.1.2) goes to zero. Furthermore, as distance realisers are continuous by definition, $G(x, y, t) = \gamma_{xy}(t)$ is continuous in t , so the second term on the right also vanishes and the implication follows.

We now turn to the forward implication and begin by noting that, given sequences $x_n \rightarrow x$ and $y_n \rightarrow y$ with $\tau(x, y) < \ell$, since $\tau^{-1}([0, \ell])$ is open and our sequences converge with respect to d , we have $\tau(x_n, y_n) < \ell$ for large enough n . Consequently, we need only consider sequences for which $\tau(x_n, y_n) < \ell$ in what follows. Suppose that $G|_{\tau^{-1}((0, \ell)) \times [0, 1]}$ is continuous and consider sequences x_n, y_n as described above. Then we have

$$\gamma_{xy}(t) = G(x, y, t) = \lim_{n \rightarrow \infty} G(x_n, y_n, t) = \lim_{n \rightarrow \infty} \gamma_{x_n y_n}(t), \quad (9.1.3)$$

so the sequence of curves converges pointwise. It remains to show that we have uniform convergence. The set

$$S := (\{(x_n, y_n) \mid n \in \mathbb{N}\} \cup \{(x, y)\}) \times [0, 1] \quad (9.1.4)$$

is a compact subset of $\tau^{-1}((0, \ell)) \times [0, 1]$ on which G is continuous, hence $G|_S$ is uniformly continuous. In particular, the sequence $\gamma_{x_n y_n}(\cdot) = G(x_n, y_n, \cdot)$ forms a uniformly equicontinuous family of functions on the compact interval $[0, 1]$. Thus we can apply the Arzelà–Ascoli Theorem [136, Lemma 3.10] to get that the sequence converges to its limit curve uniformly.¹ \square

In the metric setting, the condition that $d^{-1}([0, \ell])$ is open is automatic by the continuity of the distance function d . Since τ is merely lower semi-continuous, we must explicitly constrain $\tau^{-1}([0, \ell])$ to be open in order to obtain the above equivalence in the Lorentzian setting.

9.2 Alexandrov's Patchwork

We can now prove the Lorentzian version of the Alexandrov's Patchwork globalisation theorem. Before we begin, recall from Lemma 8.1.22 that a chronological Lorentzian pre-length space X with timelike curvature bound has locally continuous time separation function τ . In particular, timelike distance realisers of finite length in X admit constant τ -speed parametrisation on $[0, 1]$.

As in the metric setting, we will be assuming that our Lorentzian pre-length space X is

¹Note we do not need compactness to do this, since we have already shown that there is a limit curve pointwise.

uniquely D_K -geodesic a priori and working towards globalising the triangle comparison property (iii) of Definition 8.1.20. However, we also need to address the globalisation of property (i) of Definition 8.1.20, which concerns the continuity of τ on $\tau^{-1}([0, D_K])$. Since distance functions are continuous and finite, this property does not appear in the definition of curvature bounds in the metric setting (cf. Definition 7.1.12), so we cannot look there a resolution. For the purpose of this thesis we shall take the simpler approach and assume that this condition also holds a priori. We direct the curious reader to [2, Remark 4.9], where continuity of τ is discussed in more detail and it is shown that, under the same assumptions as the Lorentzian version of Alexandrov's Patchwork,² we have continuity of τ on $\tau^{-1}((0, D_K))$. However, the argument included therein does not extend to the set $\tau^{-1}([0, D_K])$ we require, since it utilises the geodesic map G , which is only defined for points joined by a distance realiser with constant τ -speed parametrisation on $[0, 1]$, which excludes points $x, y \in X$ for which $\tau(x, y) = 0$.

With that matter out of the way, we may now focus on our main task of globalising property (iii). In order to break down the proof, we first state the following lemma, which enables us to cover finite τ -length timelike distance realisers using finitely many timelike diamonds. Eventually, this will provide us with specific finite open covers of our triangles, see Figure 9.2.2

Lemma 9.2.1 (Finite Covering by Timelike Diamonds)

Let X be a strongly causal, non-timelike locally isolating, chronological Lorentzian pre-length space which has (local) timelike curvature bounded above by $K \in \mathbb{R}$. Then any timelike distance realiser γ_{xy} with finite τ -length has a finite cover of timelike diamonds I_i which are comparison neighbourhoods, for $i = 0, \dots, n$. Furthermore,

$$I_i \cap I_j \neq \emptyset \iff |i - j| \leq 1, \quad (9.2.1)$$

and both of the governing points of each diamond lie on γ_{xy} , unless the diamond contains x or y in which case only the future/past governing point does, respectively.

Proof. Let γ_{xy} be a timelike distance realiser from x to y with finite τ -length, parametrised by constant τ -speed on $[0, 1]$. For each $s \in [0, 1]$, we find a timelike diamond $I(x_s, y_s)$ that is a comparison neighbourhood of $\gamma_{xy}(s)$ by Lemma 8.1.25. Since distance realisers, e.g. γ_{xy} , are continuous, there is a neighbourhood $N_s \subseteq [0, 1]$ of s such that $\gamma_{xy}(N_s) \subseteq I(x_s, y_s)$. Therefore, we find $s^- < s < s^+$ in N_s , such that $\gamma_{xy}(s^\pm) \in I(x_s, y_s)$, and $\gamma_{xy}(s^-) \ll \gamma_{xy}(s^+)$. By causal convexity, see Lemma 8.1.13, we then find that the timelike diamond $I_s := I(\gamma_{xy}(s^-), \gamma_{xy}(s^+)) \subseteq I(x_s, y_s)$ which is also a comparison neighbourhood and whose governing points lie on γ_{xy} . When $s = 0, 1$, this is not possible as these parameters correspond to the ends of the distance realiser,

²Bar continuity of τ on $\tau^{-1}([0, D_K])$ itself, of course.

however we may still force one of the governing points to be on γ_{xy} , i.e. set $I_0 := I(p_0, \gamma_{xy}(0^+))$ and $I_1 := I(\gamma_{xy}(1^-), q_1)$ for some $p_0, q_1 \in X$ with $p_0 \ll x$ and $y \ll q_1$.

It follows that the set $\{I_s \mid s \in [0, 1]\}$ forms an open cover of $\gamma_{xy}([0, 1])$, which is the continuous image of a compact set and therefore compact. Hence, we can extract a finite subcover $\{I_{s_k} \mid k = 0, \dots, n \text{ and } s_k \in [0, 1]\}$, which must contain I_0 and I_1 , since $\gamma_{xy}(0)$ and $\gamma_{xy}(1)$ are respectively not contained in any of the other I_s by definition. We order this subcover by (the parameters of) their future governing points, i.e. $s_k^+ < s_{k+1}^+$ for all k . By assuming that the cover is minimal in the sense that no diamond can be removed while preserving the covering property, no diamond can be contained within another, so the bottom governing points can be ordered similarly, i.e. $s_k^- < s_{k+1}^-$ for all k . Then $I_{s_0} = I_0$ and $I_{s_n} = I_1$. Furthermore, diamonds which are subsequent overlap and only such diamonds do so, i.e. (9.2.1) holds, as required. \square

Theorem 9.2.2 (Synthetic Lorentzian Alexandrov's Patchwork Globalization)

Let X be a strongly causal, non-timelike locally isolating, chronological, and regular Lorentzian pre-length space which has (local) timelike curvature bounded above by $K \in \mathbb{R}$. Assume that X is uniquely D_K -geodesic, timelike distance realisers with τ -length less than D_K vary continuously with their endpoints, and τ is continuous on $\tau^{-1}([0, D_K])$ which is assumed to be open. Then X satisfies (iii) of Definition 8.1.20, in particular X has global curvature bounded above by K .

Proof. Our Lorentzian pre-length space X satisfies properties (i) and (ii) of Definition 8.1.20, by assumption. Hence, showing that X has global curvature bounded above by K amounts to showing triangle comparison (point (iii) of Definition 8.1.20).

Let $\Delta(x, y, z)$ be a timelike triangle in X which satisfies size-bounds and let G denote the geodesic map restricted to $\tau^{-1}((0, D_K)) \times [0, 1]$. By Lemma 8.1.22 τ is locally continuous and by (i) of Definition 8.1.20 $\tau^{-1}([0, D_K])$ is open, so it follows from Proposition 9.1.4 that G is continuous. For each $t \in [0, 1]$, let

$$\beta_t := \gamma_{x\gamma_{yz}(t)} : [0, 1] \rightarrow X \tag{9.2.2}$$

be the unique timelike distance realiser from x to $\gamma_{yz}(t)$.³ As G and γ_{yz} are continuous, the map

$$F(s, t) := \beta_t(s) = G(x, \gamma_{yz}(t), s) \tag{9.2.3}$$

can then be viewed as a continuous geodesic variation with starting point x which spans the $\Delta(x, y, z)$, see Figure 9.2.1 In particular, this “filled in” triangle is compact, since it is the image of the continuous function F on the compact set $[0, 1] \times [0, 1]$.

³We have existence and uniqueness by assumption, while timelike follows from regularity and (8.1.2).




Figure 9.2.1: Taking a continuous variation of distance realisers from γ_{xy} to γ_{xz} , yields distance realisers from x to $\gamma_{yz}(t)$ for each $t \in [0, 1]$. A finite cover of the filled in $\Delta(x, y, z)$ can then be obtained by covering finitely many $\gamma_{yz}(t_l)$, shown in red and blue.

Fix $t \in [0, 1]$ and consider an individual distance realiser $\beta_t = \gamma_{x\gamma_{yz}(t)}$ from this variation. Applying Lemma 9.2.1, we cover β_t with finitely many timelike diamonds $\{I_{s_k}^t \mid k = 0, \dots, n_t \text{ and } s_k \in [0, 1]\}$ whose governing points lie on β_t and such that only subsequent diamonds overlap. As F is continuous and $\bigcup_{k=0}^{n_t} I_{s_k}^t$ is a neighbourhood of $\beta_t([0, 1])$, there exists an open neighbourhood $J_t \subseteq [0, 1]$ of t , such that $F([0, 1], J_t) \subseteq \bigcup_{k=0}^{n_t} I_{s_k}^t$. By shrinking J_t if necessary, we can assume that $\gamma_{yz}(J_t) \subseteq I_1^t$. Visually, this says that $\bigcup_{k=0}^{n_t} I_{s_k}^t$ covers not only β_t but the distance realisers $\beta_{t'}$ for all t' in a neighbourhood of t and each $\beta_{t'}$ ends in the top diamond I_1 . Executing the above procedure for each $t \in [0, 1]$ yields an open cover $\{J_t \mid t \in [0, 1]\}$ of $[0, 1]$. By compactness of $[0, 1]$, we may extract a finite subcover $\{J_{t_\ell} \mid \ell = 0, \dots, m \text{ and } t_\ell \in [0, 1]\}$. Order these intervals by their right endpoint and again assume the cover is minimal, from which it follows the left endpoints have the same order. In particular, subsequent sets overlap, these are the only sets that do so, and we have $0 \in J_0$, $1 \in J_1$, i.e. $J_{t_\ell} \cap J_{t_{\ell+1}} \neq \emptyset$ for all $\ell = 0, \dots, m-1$.

Set $n_\ell := n_{t_\ell}$. Recall that $F([0, 1], J_{t_\ell}) \subseteq \bigcup_{k=0}^{n_\ell} I_{s_k}^{t_\ell}$ for each t_ℓ , from which it follows that all distance realisers from x which end in $\gamma_{yz}(J_{t_\ell})$ are contained in $\bigcup_{k=0}^{n_\ell} I_{s_k}^{t_\ell}$. Furthermore,

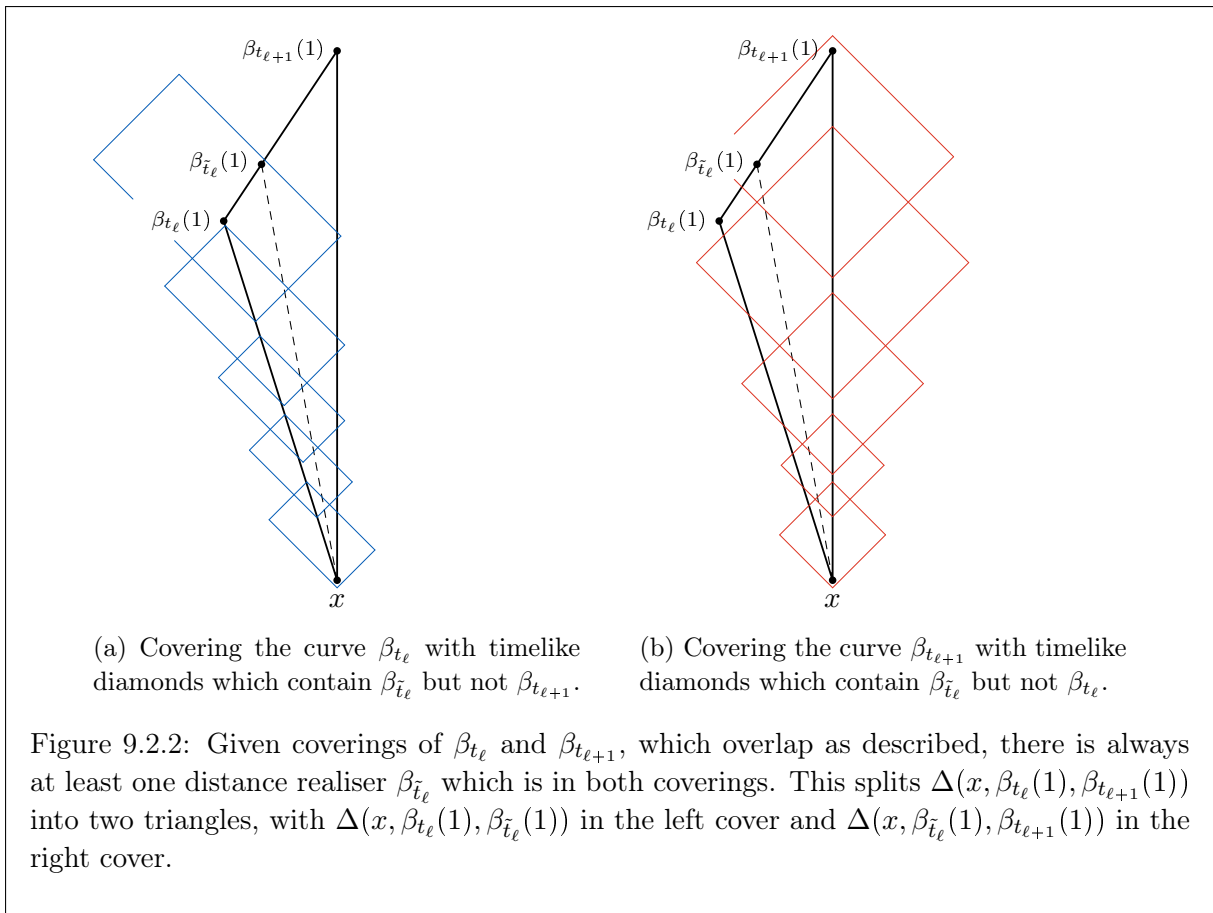
subsequent J_{t_ℓ} overlap, hence subsequent $\bigcup_{k=0}^n I_{s_k}^t$ also overlap. In total, we obtain

$$\bigcup_{\ell=0}^m \bigcup_{k=0}^{n_\ell} I_{s_k}^{t_\ell} \supseteq F([0, 1], [0, 1]). \tag{9.2.4}$$

The reward for this tedious construction is the following: The triangle $\Delta(x, y, z)$ may be viewed as a fan consisting of m pieces $\Delta(x, \beta_{t_\ell}(1), \beta_{t_{\ell+1}}(1))$ and there exists some $\tilde{t}_\ell \in J_{t_\ell} \cap J_{t_{\ell+1}}$ such that $t_\ell < \tilde{t}_\ell < t_{\ell+1}$ and

$$\beta_{\tilde{t}_\ell}([0, 1]) \subseteq \left(\bigcup_{k=0}^{n_\ell} I_{s_k}^{t_\ell} \right) \cap \left(\bigcup_{k=0}^{n_{\ell+1}} I_{s_k}^{t_{\ell+1}} \right), \tag{9.2.5}$$

for $\ell = 0, \dots, m - 1$, see Figure 9.2.2



We now wish to triangulate the “long and slim” triangles $\Delta(x, \beta_{\tilde{t}_\ell}(1), \beta_{t_{\ell+1}}(1))$ such that each of the sub-triangles is contained within one of the $I_{s_k}^{t_{\ell+1}}$. By construction, both $\beta_{\tilde{t}_\ell}$ and $\beta_{t_{\ell+1}}$ end in $\gamma_{yz}(J_{t_{\ell+1}}) \subseteq I_{s_{n_{\ell+1}}}^{t_{\ell+1}}$ and enter that set via $I_{s_{n_{\ell+1}}}^{t_{\ell+1}} \cap I_{s_{n_{\ell+1}-1}}^{t_{\ell+1}}$, that is, the distance realisers pass through the intersection of the ultimate and penultimate diamonds covering $\beta_{t_{\ell+1}}$. In particular,

there exists \tilde{r}_1 such that $\beta_{\tilde{t}_\ell}(\tilde{r}_1)$ is contained in said intersection and $\beta_{\tilde{t}_\ell}(\tilde{r}_1) \ll \beta_{t_{\ell+1}}(s_{n_{\ell+1}-1}^+)$ (recall that $s_n^+ \in [0, 1]$ is the parameter corresponding to the future governing point of the n -th diamond in the cover of β_t). Since $\tau^{-1}((0, \infty])$ is open (cf. [116, Proposition 2.13]), we may move slightly before $\beta_{t_{\ell+1}}(s_{n_{\ell+1}-1}^+)$ and stay in the timelike future of $\beta_{\tilde{t}_\ell}(\tilde{r}_1)$. In particular, there exists $r_1 \in [0, 1]$ such that $\beta_{t_{\ell+1}}(r_1) \in I_{s_{n_{\ell+1}}}^{t_{\ell+1}} \cap I_{s_{n_{\ell+1}-1}}^{t_{\ell+1}}$ and $\beta_{\tilde{t}_\ell}(\tilde{r}_1) \ll \beta_{t_{\ell+1}}(r_1)$. Both $\beta_{\tilde{t}_\ell}(r_1)$ and $\beta_{\tilde{t}_\ell}(\tilde{r}_1)$ are contained in the intersection of the last two diamonds, as is their adjoining distance realiser by causal convexity.

This process produces a quadrilateral $\beta_{\tilde{t}_\ell}(\tilde{r}_1), \beta_{\tilde{t}_\ell}(1), \beta_{t_{\ell+1}}(r_1)$, and $\beta_{t_{\ell+1}}(1)$ which is completely contained within $I_{s_{n_{\ell+1}}}^{t_{\ell+1}}$ and whose sides are timelike distance realisers. By transitivity of \ll , we also have $\beta_{\tilde{t}_\ell}(\tilde{r}_1) \ll \beta_{t_{\ell+1}}(1)$, hence by regularity there is a distance realiser between the “past most” and “future most” points of the quadrilateral. Therefore, we have two timelike triangles $\Delta(\beta_{\tilde{t}_\ell}(\tilde{r}_1), \beta_{t_{\ell+1}}(r_1), \beta_{t_{\ell+1}}(1))$ and $\Delta(\beta_{\tilde{t}_\ell}(\tilde{r}_1), \beta_{\tilde{t}_\ell}(1), \beta_{t_{\ell+1}}(1))$. Both of these triangles are entirely contained in the timelike diamond $I_{s_{n_{\ell+1}}}^{t_{\ell+1}}$, which is a comparison neighbourhood by assumption, hence have curvature bounded above by K , i.e. they satisfy the second inequality in (8.1.21). As $\beta_{\tilde{t}_\ell}(\tilde{r}_1) \ll \beta_{t_{\ell+1}}(r_1)$ are also contained in $I_{s_{n_{\ell+1}-1}}^{t_{\ell+1}}$, we can repeat the previous step and find $\beta_{\tilde{t}_\ell}(\tilde{r}_2) \ll \beta_{t_{\ell+1}}(r_2)$ in $I_{s_{n_{\ell+1}-1}}^{t_{\ell+1}} \cap I_{s_{n_{\ell+1}-2}}^{t_{\ell+1}}$, yielding another quadrilateral. We can continue this procedure iteratively and after $n_{\ell+1} - 1$ steps, we end up with $\beta_{\tilde{t}_\ell}(\tilde{r}_{n_{\ell+1}-1}) \ll \beta_{t_{\ell+1}}(r_{n_{\ell+1}-1})$ in $I_{s_1}^{t_{\ell+1}}$, where also x lies. That is, in the end we have $n_{\ell+1} - 1$ quadrilaterals, each of which we can split into two timelike triangles, and one additional timelike triangle $\Delta(x, \beta_{\tilde{t}_\ell}(\tilde{r}_{n_{\ell+1}-1}), \beta_{t_{\ell+1}}(r_{n_{\ell+1}-1}))$ at the bottom of $\Delta(x, y, z)$, see Figure 9.2.3.

Each of these timelike triangles is contained in one of the diamonds $\{I_{s_k}^{t_{\ell+1}} \mid k = 0, \dots, n_{\ell+1} \text{ and } \ell = 0, \dots, m-1\}$ which are comparison neighbourhoods, hence has curvature bounded above by K . Therefore, $2(n_{\ell+1} - 1)$ applications of the gluing lemma (Lemma 9.1.1), starting with $\Delta(x, \beta_{\tilde{t}_\ell}(\tilde{r}_{n_{\ell+1}-1}), \beta_{t_{\ell+1}}(r_{n_{\ell+1}-1}))$ and working up $\Delta(x, \beta_{\tilde{t}_\ell}(1), \beta_{t_{\ell+1}}(1))$ (so that the result is always another timelike triangle) yields that the “long and slim” triangle $\Delta(x, \beta_{\tilde{t}_\ell}(1), \beta_{t_{\ell+1}}(1))$ also satisfies the curvature bound.

The same process can be applied to “long and slim” triangles of the form $\Delta(x, \beta_{t_\ell}(1), \beta_{\tilde{t}_{\ell+1}}(1))$, provided we account for the different time orientation of top side relative to the geodesic β_{t_ℓ} around which the covering is centred. Aside from reversing the timelike relation, the process remains unchanged. This results in $2m - 1$ “long and slim” triangles which have curvature bounded above by K . Hence, $2m - 2$ more applications of Lemma 9.1.1 yields that the original triangle $\Delta(x, y, z)$ satisfies the desired curvature bound.

□

Recall that while the above theorem explicitly globalises curvature bounded above in terms




Figure 9.2.3: Subdividing the slim triangle $\Delta(x, \beta_{i_\ell}(1), \beta_{t_{\ell+1}}(1))$ into triangles by constructing a quadrilateral in all but the first timelike diamond. Each sub-triangle is then contained in a comparison neighbourhood and the Gluing Lemma can be applied.

of the triangle comparison condition [Definition 8.1.20](#), since X is assumed to be chronological and regular, we may apply [Proposition 8.2.9](#) to obtain that curvature bounds in the sense of angle and hinge comparison also globalise.

As in the metric setting, it is possible to show that some of the assumptions made in [Theorem 9.2.2](#) are not only sufficient, but necessary. We begin this process by providing an illustrative example where curvature bounded above by $K = 0$ does not globalise, namely the Lorentzian cylinder. In particular, this space has neither unique nor continuously varying distance realisers (cf. the metric circle from [Example 7.2.1](#)).

Example 9.2.3 (The Lorentzian Cylinder)

Let $X := \mathbb{R} \times \mathbb{S}^1$ be the Lorentzian pre-length space given by the strip $\mathbb{R} \times [0, 2\pi]$ in the Minkowski plane $\mathbb{L}^2(0)$ with the timelike boundary glued in-line with its time-orientation, see [Figure 9.2.4](#). We call this space the Lorentzian cylinder. Locally, this space is isometric to $\mathbb{L}^2(0)$, hence has timelike curvature bounded above by 0.

However, X itself is not a (≤ 0) -comparison neighbourhood. Firstly, recall that $D_0 = \infty$, so we need to consider triangles of arbitrary side lengths. Note that triangles which wrap around the whole cylinder do not look like triangles in $\mathbb{L}^2(0)$ and are impacted by the periodic structure. We construct such a triangle as follows. Take two vertical lines on the cylinder which are opposite




Figure 9.2.4: The triangle $\Delta(x, y, z)$ in the Lorentzian cylinder and its comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in the Minkowski plane. $\Delta(\bar{x}, \bar{y}, \bar{z})$ is degenerate, so $\Delta(x, y, z)$ fails to have curvature bounded above.

one another, for example the dotted lines in Figure 9.2.4. Then, given two timelike related points $x \ll z$ with one on each of the vertical lines, there exist precisely two distance realisers between them, which wrap around the right and left of the cylinder respectively. We may choose a third point y along one of the two distance realisers to form a triangle $\Delta(x, y, z)$ for which $\tau(x, z) = \tau(x, y) + \tau(y, z)$. It is then clear that the comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in $\mathbb{L}^2(0)$ is degenerate. Furthermore, we may choose a pair of points p, q with one on each of the distance realisers from x to z , such that $p \not\ll q$ and $\tau(x, p) \neq \tau(x, q)$.⁴ We then have $\tau(p, q) = 0 < \tau(\bar{p}, \bar{q})$, as in Figure 9.2.4, from which it follows that X does not have curvature globally bounded above by 0.

In particular, X does not have unique distance realisers between points with separation $\geq \pi$ and we can show it does not have continuously varying distance realisers in the same manner as for the circle in Example 7.2.5.

△

It should now be clear that, without unique distance realisers between timelike related points with τ -separation less than D_K , which also vary continuously with their endpoints, globalisation of curvature bounded above by K is by no means automatic. In fact, our next result shows that (timelike) distance realisers with length less than D_K are necessarily unique in spaces with global curvature bounded above by K , as was the case in the metric setting (cf. Proposition 7.2.4).

Proposition 9.2.4 (Unique Distance Realisers)

⁴Assume here that both distance realisers are parametrised on $[0, 1]$ with constant speed.

Let X be a chronological, strongly causal, and regular Lorentzian pre-length space with timelike curvature bounded above by $K \in \mathbb{R}$. Let $x, y \in U \subseteq X$ with $x \ll y$, $\tau(x, y) < D_K$, and U a comparison neighbourhood. Then there exists a unique distance realiser from x to y contained in U . In particular, if X has global timelike curvature bounded above by K , distance realisers between timelike related points in X with τ -distance less than D_K are unique.

Proof. We know from the definition of a comparison neighbourhood that there exists at least one distance realiser from x to y in U . Suppose for contradiction that there are (at least) two distance realisers from x to y . Denote these curves by α_1 and α_2 and parametrise them with constant speed on $[0, 1]$. Let $p = \alpha_1(s)$ for some $s \in (0, 1)$ and consider the timelike triangle $\Delta(x, p, y)$ with sides α_2 , $\alpha_1|_{[0, s]}$, and $\alpha_2|_{[s, 1]}$. This satisfies size-bounds for $\mathbb{L}^2(K)$ since $\tau(x, y) < D_K$.

As for the Lorentzian cylinder in [Example 9.2.3](#), it is clear that the comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ in $\mathbb{L}^2(K)$ is degenerate. Furthermore, since $\alpha_1 \neq \alpha_2$, there exists $t \in (0, 1)$ such that $\alpha_1(t) \neq \alpha_2(t)$. Since metric spaces are Hausdorff, there then exists neighbourhoods V_1 of $\alpha_1(t)$ and V_2 of $\alpha_2(t)$ with $V_1 \cap V_2 = \emptyset$. As X is strongly causal, there are points $x_1, y_1, \dots, x_n, y_n$ and $p_1, q_1, \dots, p_m, q_m$ such that $\alpha_1(t) \in U_1 := \bigcap_{i=1}^n I_1(x_i, y_i) \subseteq V_1$ and $\alpha_2(t) \in U_2 := \bigcap_{j=1}^m I_1(p_j, q_j) \subseteq V_2$, where the $I(p, q)$ are timelike diamonds. By continuity of the α_i , there is some neighbourhood $J_i \subseteq [0, 1]$ of t such that $\alpha_i(J_i) \subseteq U_i$. In particular, there exists $\varepsilon > 0$ such that $\alpha_i((t - \varepsilon, t + \varepsilon)) \subseteq U_i$.

By [Lemma 8.1.13](#), the intersection of timelike diamonds is causally convex, so the U_i are causally convex and contain the diamonds $D_i := I(\alpha_i(t - \varepsilon), \alpha_i(t + \varepsilon)) \subseteq U_i$ respectively. As $V_1 \cap V_2 = \emptyset$ and $U_i \subseteq V_i$, it follows that $\alpha_2(t) \notin D_1$ (and $\alpha_1(t) \notin D_2$). Hence, either $\alpha_1(t - \varepsilon) \not\ll \alpha_2(t)$ or $\alpha_2(t) \not\ll \alpha_1(t + \varepsilon)$. However, in $\mathbb{L}^2(K)$ we have $\bar{\alpha}_1(s) = \bar{\alpha}_2(s)$ for all $s \in [0, 1]$ since the comparison triangle is degenerate, from which it follows that, $\bar{\alpha}_1(t - \varepsilon) \ll \bar{\alpha}_2(t) \ll \bar{\alpha}_1(t + \varepsilon)$. In particular, we have either $\tau(\alpha_1(t - \varepsilon), \alpha_2(t)) = 0 < \tau(\bar{\alpha}_1(t - \varepsilon), \bar{\alpha}_2(t))$ or $\tau(\alpha_2(t), \alpha_1(t + \varepsilon)) = 0 < \tau(\bar{\alpha}_2(t), \bar{\alpha}_1(t + \varepsilon))$, each of which contradict point (iii) of [Definition 8.1.20](#), which holds on U due to the curvature bound. \square

Recall that for metric spaces with global curvature bounded above by k , we also obtained that the assumption distance realisers of length less than $\text{diam}(M^2(k))$ vary continuously with their endpoints was necessary, see [Proposition 7.2.4](#). While we have shown above that timelike distance realisers of length less than D_K exist and are unique in well behaved Lorentzian pre-length spaces with global timelike curvature bounded above by K , in contrast to the metric setting it is rather complicated to obtain that these distance realisers vary continuously with their endpoints. In [\[2, Proposition 4.8\]](#), the authors prove that distance realisers exhibit this behaviour in all globally hyperbolic Lorentzian length spaces with global curvature bounded

above by K ,⁵ however this is a marked restriction of the chronological, strongly causal, and regular Lorentzian pre-length spaces considered in this chapter thus far. Hence, it is expected that, as with the continuity of $\tau^{-1}([0, D_K])$ there is some room for refinement in this result and we shall not discuss it further here.

9.3 Chapter Summary

In this chapter we have proven the first globalisation theorem for curvature bounds on Lorentzian pre-length spaces. In particular, we derived [Theorem 9.2.2](#) for Lorentzian pre-length spaces with curvature bounded above, in the spirit of Alexandrov’s patchwork from [Theorem 7.2.2](#).

We began by restating the gluing lemma [126] in a more compact manner, which allows us to infer that arbitrary timelike triangles have curvature bounded above if they can be decomposed into smaller timelike triangles with such bounds. Subsequently, in [Definition 9.1.2](#), we introduced a notion of continuously varying distance realisers between timelike related points and showed, in [Proposition 9.1.4](#), that under some relatively mild assumptions this is equivalent to the geodesic map being continuous. This allowed us to build out our main result — by assuming X to be uniquely D_K -geodesic, τ to be continuous on the open set $\tau^{-1}([0, D_K])$,⁶ and distance realisers to vary continuously with their endpoints, we were able to globalise the triangle comparison condition (point (iii) of [Definition 8.1.20](#)) in [Theorem 9.2.2](#). This precisely mirrors the statement of the Alexandrov’s Patchwork theorem in the metric setting. Furthermore, in [Proposition 9.2.4](#) we demonstrated that assuming X to be uniquely D_K -geodesic is not only sufficient, but necessary in the sense that every well-behaved Lorentzian pre-length space with global curvature bounded above by $K \in \mathbb{R}$ is uniquely D_K -geodesic. It should be noted that while we prove these results for curvature bounds in the sense of triangle comparison, [Proposition 8.2.9](#) implies that curvature bounded above with respect to angle comparison and hinge comparison globalise under the same assumptions.

In the next chapter, we address globalisation of curvature bounds for Lorentzian pre-length spaces with curvature bounded below. Unlike the curvature bounded above case, where many of the necessary tools had already been derived (e.g. the gluing lemma from [Lemma 9.1.1](#)) or could be drawn from the metric setting (e.g. [Definition 9.1.2](#)), for curvature bounded below we will need to develop several new tools including a Lebesgue number lemma for Lorentzian pre-length spaces and a gluing lemma for curvature bounded below. This increase in complexity turns out to be a boon however, providing us with avenues for new research, as we will see in [Chapter 11](#).

⁵In fact, the result holds for chronological, strongly causal, non-timelike locally isolating, non-totally imprisoning, locally causally closed, regular Lorentzian pre-length spaces with global curvature bounded above.

⁶These assumptions correspond to points (i) and (ii) of [Definition 8.1.20](#).



Globalisation of Timelike Curvature Bounded Below

Where the previous chapter treated Lorentzian pre-length spaces with timelike curvature bounded above, in this chapter we focus on the notion of Lorentzian pre-length spaces with timelike curvature bounded below. In particular, we work towards proving the lower bound counterpart to the Alexandrov's Patchwork, namely the Toponogov globalisation result, [Theorem 7.2.6](#), in the synthetic Lorentzian setting. We also formulate a theorem in the spirit of Bonnet–Myers, see [Theorem 7.2.7](#), bounding the (finite, timelike) diameter of spaces with timelike curvature bounded below by $K < 0$. We prove this result by first considering spaces with global timelike curvature bounded below by $K < 0$ and then apply the Toponogov globalisation theorem to extend to spaces with local curvature bounds.

In contrast to the case of curvature bounded above, where we globalised the comparison condition and used equivalence of definitions to extend the result to angle and hinge comparison, in this section we will work directly with the angle condition. We do this in part to keep parity with the existing literature; we follow the approach of [\[133\]](#) who use the angle condition in the synthetic metric setting, while [\[118\]](#) globalise the angle condition in the smooth Lorentzian setting. Equally, using the angle condition enables us to obtain globalisation theorems for the triangle, hinge, and angle conditions under the additional assumption that the triangle inequality for angles [\(8.2.6\)](#) holds, by [Proposition 8.2.9](#), without us having to carry around this supplementary condition in the statements of our results.

We formulate many of the results in this chapter in terms of globally hyperbolic Lorentzian length spaces, which we previously described as the Swiss army knife of synthetic Lorentzian spaces (see the discussion after [Definition 8.1.18](#)). We do so because, while we do not use all of the

properties of these spaces, we use sufficiently many that writing down the most general statement for “*causally path connected, piecewise causal path connected, regular, strongly causal, globally hyperbolic, chronological, non-timelike locally isolating, uniquely D_K -geodesic... Lorentzian pre-length spaces with time function*” becomes less useful than a succinct statement. We highlight this here because the catch-all quality of globally hyperbolic Lorentzian length space can sometimes be viewed as a hammer for forcing results to work, if it is unclear precisely which properties are necessary.

10.1 Preliminary Results

Let us begin our derivation by providing a series of preliminary results, collected in the following two subsections. In the first subsection, we discuss the time function and null distance introduced by Sormani and Vega [131] in the smooth setting and extended to the synthetic Lorentzian setting by Kunzinger and Steinbauer [132]. In particular, we highlight how the somewhat artificial notion of a *piecewise causal path connected* Lorentzian pre-length space is simply a connected Lorentzian pre-length space, under some relatively mild topological assumptions which hold for all Lorentzian length spaces. This culminates in a rather interesting observation: the Lebesgue number lemma [157, Lemma 7.2] can be applied to “well-behaved” null distances and on connected, globally hyperbolic Lorentzian length spaces, this allows us to ensure that timelike triangles which are sufficiently small with respect to the null distance lie within a comparison neighbourhood.

In the second subsection, we recall the Lorentzian version of the Alexandrov lemma and use it to derive a decomposition result, in the spirit of the gluing lemma Lemma 9.1.1 from the curvature bounded above case. This result allows us to break triangles which do not adhere to the curvature bound down into triangles which are small with respect to the null distance. Furthermore, one of these small triangles must also not adhere to the curvature bound. Combining this with the Lebesgue number lemma for the null distance, we may obtain the contradiction necessary for the Toponogov globalisation theorem to hold in the synthetic Lorentzian setting.

10.1.1 Time Functions and Null Distance

The null distance d_T induced by a time function T was originally introduced in [131] as a convenient way of equipping a spacetime with a natural metric space structure which induces the manifold topology and is compatible with the causal structure.¹ This concept was later extended to the setting of Lorentzian pre-length spaces in [132] and it is this reference to which we refer the reader for more detail.

¹Recall that the function (6.0.3) fails to be a true distance for Lorentzian metrics g and the manifold topology is usually obtained from a Riemannian metric h , as in Example 8.1.2

With regard to our goal of deriving a globalisation theorem, the null distance provides us with an ideal way of describing the “size” of a timelike triangle. Unlike in the metric setting, where size is measured exclusively using the distance function, in a Lorentzian pre-length space there are two different functions which we need to control: the time separation τ which is used in the definition of timelike curvature bounds, cf. [Definition 8.2.7](#), and the distance function d which defines the underlying topology and dictates whether or not a triangle is in a comparison neighbourhood. As we shall see shortly, when applied to the covering constructed in [Lemma 8.1.25](#), the null distance allows us to control both perspectives simultaneously.

Recall that null distances are induced by so-called *time functions*. Before writing down a precise definition of the former, let us introduce the latter:

Definition 10.1.1 (Time Functions)

A generalised time function T on a Lorentzian pre-length space X is a map $T : X \rightarrow \mathbb{R}$ which is strictly increasing along every (non-constant) future-directed causal curve. If T is additionally continuous with respect to d , then it is called a time function.

Throughout the remainder of this chapter, we will directly assume that our space possesses a time function. However, we would like to draw the readers’ attentions to the following result which demonstrates that, in our simplified setting, second countability of the space is sufficient to ensure that a time function exists.

Proposition 10.1.2 (Existence of Time Functions)

Let X be a second countable and globally hyperbolic Lorentzian length space. Then X possesses a time function T .

Proof. In [[130](#), Theorem 3.2] it is shown that second countable, locally compact, stably causal Lorentzian pre-length space X with limit curves possess a time function T . It remains to show that a globally hyperbolic Lorentzian length space is locally compact, stably causal, and has limit curves. We do so by combining the following results.

- (i) Lorentzian length spaces are causally path connected, locally causally closed, and by localisability, $I^\pm(x) \neq \emptyset$ for all $x \in X$, see [Definition 8.1.17](#).
- (ii) Causally path connected, locally causally closed, and globally hyperbolic Lorentzian pre-length spaces such that $I^\pm(x) \neq \emptyset$ for all $x \in X$ are locally compact, see [[127](#), Lemma 3.8].
- (iii) Globally hyperbolic Lorentzian pre-length spaces are stably causal, see [[147](#), Theorem 3.20].

- (iv) Locally causally closed Lorentzian pre-length spaces possess limit curves, see [116, Theorem 3.7].

□

Before we define the null distance, we require the following class of curves between (not necessarily causally related) points in X .

Definition 10.1.3 (Piecewise Causal Curves)

Let X be a Lorentzian pre-length space. A piecewise causal curve $\gamma : [a, b] \rightarrow X$ in X is a continuous curve consisting only of future-directed causal, past-directed causal, and constant pieces. That is, there exists a partition $a = s_0 < s_1 < \dots < s_k = b$ of $[a, b]$, for some $k \in \mathbb{N}$, such that $\gamma|_{[s_i, s_{i+1}]}$ is constant, future-directed causal, or past-directed causal for each $i = 0, \dots, k-1$.

Definition 10.1.4 (Null Distance)

Let X be a Lorentzian pre-length space with generalised time function T . Let γ be a piecewise causal curve with partition as in Definition 10.1.3 and set $x_i := \gamma(s_i)$. The null length of γ is given by

$$L_T(\gamma) := \sum_{i=1}^k |T(x_i) - T(x_{i-1})|, \quad (10.1.1)$$

and the null distance between any pair of points $x, y \in X$ is

$$d_T(x, y) := \inf\{L_T(\gamma) \mid \gamma \text{ piecewise causal from } x \text{ to } y\}. \quad (10.1.2)$$

The null length of a piecewise causal curve is independent of the choice of partition of $[a, b]$, as long as the partition results in sub-curves $\gamma|_{[s_{i-1}, s_i]}$ which are either causal or constant (in particular, the sub-curves are not allowed to be piecewise causal).

While the definition of the null length ensures that the null distance is always symmetric, non-negative, and satisfies the triangle inequality, cf. Definition 7.1.1, the null distance need not be a true distance function a priori. In particular, [131, Theorem 4.6] and [132, Theorem 3.12] demonstrate in their respective settings that T being locally anti-Lipschitz is a necessary and sufficient condition for d_T to be a genuine distance function. For our purposes, the null distance is sufficiently well-behaved if it is merely a finite, continuous pseudo-distance, that is, d_T may not distinguish points i.e. $d_T(x, x) = 0$ but $d_T(x, y) = 0$ does not imply $x = y$ in general. We now provide conditions under which the null distance is of this type.

Given two points $x, y \in X$ which are not joined by a piecewise causal curve, $d(x, y) = \infty$. Consequently, for the null distance to be finite, it is necessary (but not in general sufficient) for

every pair of points to be joined by a piecewise causal curve; we begin by addressing when a Lorentzian pre-length space has this level of connectedness, aided by the following terminology from [132, Definition 3.4]:

Definition 10.1.5 (Causal Connectedness)

Let X be a Lorentzian pre-length space. We call X piecewise causal path connected if there exists a piecewise causal curve between each pair of points $x, y \in X$. Furthermore, if X is causally path connected (cf. Definition 8.1.17), path connected, and every point $x \in X$ lies on some timelike curve, then we call X sufficiently causally connected.

The following lemma highlights that, for causally path connected Lorentzian pre-length spaces, the notion of piecewise causal path connectedness is equivalent to the far more familiar notion of connectedness. The equivalence between path-connectedness and piecewise causal path connectedness was also noted by [132, Lemma 3.5] and [131, Lemma 3.5] in their respective settings. In particular, Lemma 10.1.6 holds for Lorentzian length spaces.

Lemma 10.1.6 (Path-connected Lorentzian pre-length spaces)

Let X be a causally path connected Lorentzian pre-length space such that for each $x \in X$ either $I^+(x)$ or $I^-(x)$ is non-empty. Then the following are equivalent:

- (i) X is connected.
- (ii) X is path connected.
- (iii) X is piecewise causal path connected.

Proof. It is well known that path connected spaces are connected. In a similar vein, it is clear that piecewise causal path connected spaces are path connected — simply take the path between the points to be the piecewise causal one. Hence, it remains to show that for X as described above, connectedness implies piecewise causal path connectedness.

Let $p \in X$ and R_p be the set of all points which are connected to p via a piecewise causal path. By assumption, for each $q \in R_p$, there exists an r such that either $r \ll q$ or $q \ll r$. Assume the former without loss of generality. Since X is causally path connected, there exists a causal curve between q and r , from which it follows that there is a piecewise causal curve from p to r and $r \in R_p$. Similarly, each point in $I^+(r)$ (resp. $I^-(r)$) is connected to r and therefore p by a piecewise causal curve, so $I^+(r) \subseteq R_p$ (resp. $I^-(r) \subseteq R_p$) is an open neighbourhood of q contained in R_p . As q was arbitrary, it follows that R_p is open.

Furthermore, given any pair of points $p, q \in X$, the sets R_p and R_q are either equal or disjoint. Consequently, $\{R_p \mid p \in X\}$ gives an open partition of X . Then $\{R_p \mid p \in X\}$ can be

reduced to open partition of X (by discarding duplicates). However, connected spaces cannot be represented as the union of two or more non-empty, disjoint open subsets, so R_p must consist of precisely one element, namely $R_p = X$ for all $p \in X$. It follows that X is piecewise causal path connected, as required. \square

Lemma 10.1.6 implies that sufficiently causally connected Lorentzian pre-length spaces are piecewise causal path connected, hence the null distance between pairs of points is not trivially infinite. In fact, the null distance is a finite continuous pseudo-distance on such spaces, which include connected Lorentzian length spaces, as shown by the following proposition:

Proposition 10.1.7 (Null Distance is a Finite Continuous Pseudo-distance)

Let X be a sufficiently causally connected Lorentzian pre-length space with generalised time function T and distance function d . Then the null distance d_T induced by T is a finite pseudo-distance which is continuous (with respect to d). Moreover,

$$d_T(x, y) = 0 \Rightarrow T(x) = T(y) \quad (10.1.3a)$$

and

$$x \preceq y \Rightarrow d_T(x, y) = T(y) - T(x). \quad (10.1.3b)$$

In particular, the above holds when X is a connected Lorentzian length space.

Proof. For sufficiently causally connected Lorentzian pre-length spaces, this result was proven in [132]. More precisely,

- (i) d_T is a finite pseudo-distance by [132, Lemma 3.7].
- (ii) d_T is continuous with respect to d by [132, Proposition 3.9].
- (iii) (10.1.3a) and (10.1.3b) hold by [132, Proposition 3.8.(i), (ii)] respectively.

It remains to show that connected Lorentzian length spaces are sufficiently causally connected. By definition, Lorentzian length spaces are causally path connected, while connectedness is equivalent to path connectedness in this setting by Lemma 10.1.6. Finally, by localisability of Lorentzian length spaces, see [116, Definitions 3.16, 3.22], $I^\pm(x)$ are non-empty for each $x \in X$. As Lorentzian length spaces are causally path connected, there exists a future directed timelike curve between x and some other point. Hence every point $x \in X$ lies on (is the endpoint of) some timelike curve, from which it follows that Lorentzian length spaces are sufficiently causally connected. \square

As we did with distance functions in Definition 7.1.1, we may define the diameter of a set

with respect to a pseudo-distance via Definition 7.1.2. In particular, the d_T -diameter is

$$\text{diam}_T(U) := \sup\{d_T(x, y) \mid x, y, \in U\}. \quad (10.1.4)$$

Interpreting an admissible causal triangle $\Delta(x, y, z)$ as the one-dimensional space formed by the union of the curves $\gamma_{xy}, \gamma_{yz}, \gamma_{xz}$, we find $\text{diam}_T(\Delta(x, y, z)) = T(z) - T(x)$. It is in this sense that the d_T -diameter of an admissible causal triangle represents its “size.” From the metric point of view, any admissible causal triangle is degenerate with respect to d_T by (10.1.3b), i.e.

$$d_T(x, z) = d_T(x, y) + d_T(y, z). \quad (10.1.5)$$

A similar result holds for causal and timelike diamonds, as described in the following lemma, which is a strengthening of [132, Proposition 3.8(iv)].

Lemma 10.1.8 (Null Distance Diameter of Diamonds)

Let X be a sufficiently causally connected Lorentzian pre-length space with generalised time function T and associated null distance d_T . Then the d_T -diameter of a non-empty causal diamond $J(x, y)$ is given by $\text{diam}_T(J(x, y)) = T(y) - T(x)$, while for a non-empty timelike diamond $I(x, y)$ we have $\text{diam}_T(I(x, y)) \leq T(y) - T(x)$. In particular, this holds when X is a connected Lorentzian length space.

Proof. Let $J(x, y) \subseteq X$ be a non-empty causal diamond in X . Then $x \leq y$ and (10.1.3b) implies that $d_T(x, y) = T(y) - T(x)$. If $x = y$, then $J(x, y) = \{x\}$ and we are done. Otherwise, there are at least two points in $J(x, y)$ and we must show that $d_T(p, q) \leq T(y) - T(x)$ for all $p, q \in J(x, y)$.

So let $p, q \in J(x, y)$ and consider the causal (or constant) curves $\gamma_{xp}, \gamma_{qx}, \gamma_{py}, \gamma_{yq}$, which exist by causal path connectedness of X . By (10.1.1), $L_T(\gamma_{xp}) = T(p) - T(x)$ and similarly for the other curves. Define the concatenated curves

$$\begin{aligned} \gamma_{pq}^+ &= \gamma_{py} + \gamma_{yq} \\ \gamma_{qp}^- &= \gamma_{qx} + \gamma_{xp} \\ \gamma &= \gamma_{xp} + \gamma_{py} + \gamma_{yq} + \gamma_{qx}. \end{aligned} \quad (10.1.6)$$

Since the null length (10.1.1) is additive by definition, we find

$$L_T(\gamma_{pq}^+) + L_T(\gamma_{qp}^-) = L_T(\gamma) = 2(T(y) - T(x)), \quad (10.1.7)$$

from which it follows that one of γ_{pq}^+ and γ_{qp}^- must have null length less than $T(y) - T(x)$ (it is not possible for both curves to contribute more than half the length). Since in Definition 10.1.4,

the null distance between p and q is given by the infimum of the null lengths of all piecewise causal curves between them and one of these curves has null length at most $T(y) - T(x)$, it follows that $d_T(p, q) \leq T(y) - T(x) = d_T(x, y)$ as required.

Now consider a non-empty timelike diamond $I(x, y) \subseteq J(x, y) \subseteq X$. Since all points $p, q \in I(x, y)$ are also in $J(x, y)$, we know $d_T(p, q) \leq d_T(x, y)$ from above. It then follows from $\text{diam}_T(I(x, y)) = \sup\{d_T(p, q) \mid p, q \in I(x, y)\}$ that $\text{diam}_T(I(x, y)) \leq d_T(x, y)$. \square

We now turn to the main and final result of this subsection, the Lebesgue number lemma for the null distance, which allows us to confine sufficiently small subsets of causal diamonds within a timelike diamond which is additionally a comparison neighbourhood.

Lemma 10.1.9 (Lebesgue Number Lemma for the Null Distance)

Let X be a connected, globally hyperbolic, Lorentzian length space with time function T and associated null distance d_T . Then for any causal diamond $J(x, y) \subseteq X$ and open cover $\{D_i\}_{i=1}^n$ of $J(x, y)$, there exists an $\varepsilon > 0$ such that any causal (hence any timelike) diamond with d_T -diameter less than ε which is contained in $J(x, y)$ is also contained in an element of the covering.

In particular, there exists an $\varepsilon > 0$ such that any causal/timelike diamond with d_T -diameter less than ε which is contained in $J(x, y)$ is contained inside a timelike diamond. If X additionally has curvature bounded above/below by $K \in \mathbb{R}$, there exists an $\varepsilon > 0$ such that the enclosing timelike diamond is a comparison neighbourhood.²

Proof. The proof of this result is similar to the classical proof of the Lebesgue number lemma for a metric space [157, Lemma 7.2], with the added observation that the diameter can be taken with respect to a pseudo-distance d_T and this pseudo-distance need not coincide with the distance function d with respect to which we define the topology on (X, d) . Let $x, y \in X$. Since $J(x, y)$ is compact by global hyperbolicity, we assume without loss of generality that $\{D_i\}_{i=1}^n$ is a *finite* open cover of $J(x, y)$.

Firstly, if $J(x, y) \subseteq D_i$ for some i then we can choose arbitrary ε and we are done. Otherwise, denote by $C_i := J(x, y) \setminus D_i$ the complement of D_i in $J(x, y)$. Define a function $f : J(x, y) \rightarrow \mathbb{R}$ via

$$f(z) := \max_{i \in \{1, 2, \dots, n\}} d_T(z, C_i \cap (J^+(z) \cup J^-(z))) = \max_{i \in \{1, 2, \dots, n\}} \inf\{d_T(z, p) \mid p \in C_i \cap (J^+(z) \cup J^-(z))\}. \quad (10.1.8)$$

Since X is globally hyperbolic, $J^\pm(z)$ are closed by [147, Theorem 3.16, Definition 3.7]. Hence $C_i \cap (J^+(z) \cup J^-(z))$ is a closed subset of $J(x, y)$ and therefore compact, from which it follows

²Here curvature bounds are taken with respect to triangle comparison, but can be cast to angle or hinge comparison via Proposition 8.2.9.

that the infimum in (10.1.8) is attained.

We now show that $f(z) \in (0, \infty)$ for all p . If $f(z) = 0$ for some $z \in J(x, y)$ fixed, this implies $d_T(z, C_i \cap (J^+(z) \cup J^-(z))) = 0$ for all $i = 1, \dots, n$. As the infimum in (10.1.8) is attained, this implies that for each $i = 1, \dots, n$, there exists a $p_i \in C_i \cap (J^+(z) \cup J^-(z))$, such that $d_T(z, p_i) = 0$. In particular, by (10.1.3a), $T(z) = T(p_i)$. Furthermore, as $p_i \in J^+(z) \cup J^-(z)$, we have either $z \preceq p_i$ or $p_i \preceq z$ and as X is causally path connected, cf. Definition 8.1.17, there exists a future-directed (resp. past-directed) causal curve from z to p_i . Since time functions are strictly increasing (resp. decreasing) on such curves by definition, this implies $z = p_i$ for all $i = 1, \dots, n$. Therefore, $z \in C_i$ or, equivalently, $z \notin D_i$ for all $i = 1, \dots, n$. As the D_i cover $J(x, y)$, we find $z \notin J(x, y)$, which is a contradiction. If $f(z) = \infty$ for some z , then there exists some i such that $C_i \cap (J^+(z) \cup J^-(z)) = \emptyset$. Indeed, as all of these sets are compact and the null distance is finite valued, the maximum of finitely many infima can only be infinite if (at least) one of the sets is empty. Thus, $J(x, y) \cap (J^+(z) \cup J^-(z)) \subseteq D_i$, and hence $x, y \in D_i$. As D_i is a timelike diamond and therefore causally convex, this implies $J(x, y) \subseteq D_i$, which we have already treated.

As the sets $C_i \cap (J^+(z) \cup J^-(z))$ and $\{1, \dots, n\}$ are all compact and the null distance is continuous, it follows that f is a continuous function³ on a compact set, hence attains its minimum value. Consequently, we set $\varepsilon := \min_{z \in J(x, y)} f(z) > 0$. Now let $p, q \in J(x, y)$ with $p \preceq q$ and $\text{diam}_T(J(p, q)) = d_T(p, q) < \varepsilon$. As $f(p) \geq \varepsilon$, there exists i such that $d_T(p, C_i \cap (J^+(p) \cup J^-(p))) \geq \varepsilon$. Then clearly, $p \notin C_i$. Furthermore, since $p \preceq q$ and $d_T(p, q) < \varepsilon$, we have $q \notin C_i$. Thus, $p, q \in D_i$ and by the causal convexity of diamonds, also $J(p, q) \subseteq D_i$.

Since X is a globally hyperbolic Lorentzian length space, it is strongly causal, non-timelike locally isolating, and chronological, so the addenda follow by taking the open cover to be the basis of diamonds from Proposition 8.1.24 and Lemma 8.1.25 respectively. \square

10.1.2 The Gluing Lemma for Curvature Bounded Below

Recall that globally hyperbolic Lorentzian length spaces X are geodesic, with finite and continuous time separation function, see the discussion following Definition 8.1.18. Consequently, for the remainder of this chapter, when we speak of a distance realiser between causally related points, we may rest assured that such a distance realiser exists. Furthermore, conditions (i) and (ii) of Definition 8.2.7 hold in the case $U = X$ and for our globalisation theorem, we again only need to consider points (iii) and (iv). While point (iv) of Definition 8.2.7 is required in the case of curvature bounded below, its globalisation is automatic: the condition considers the

³Recall that if $h : X \times Y \rightarrow \mathbb{R}$ is a continuous function, with Y compact, then $f(x) = \inf\{h(x, y) \mid y \in Y\}$ is continuous (and similarly for the maximum)

germs of distance realisers at each point $x \in U$ and we have a covering of X by U on which the inequality (8.2.6) holds, from which it follows that the condition holds at each $x \in X$. With the other points addressed, we now turn our focus to point (iii).

In order to simplify the discussion surrounding the angle comparison condition (8.2.5), we first introduce some terminology to describe when a given angle in an admissible causal triangle does or does not satisfy the constraint.

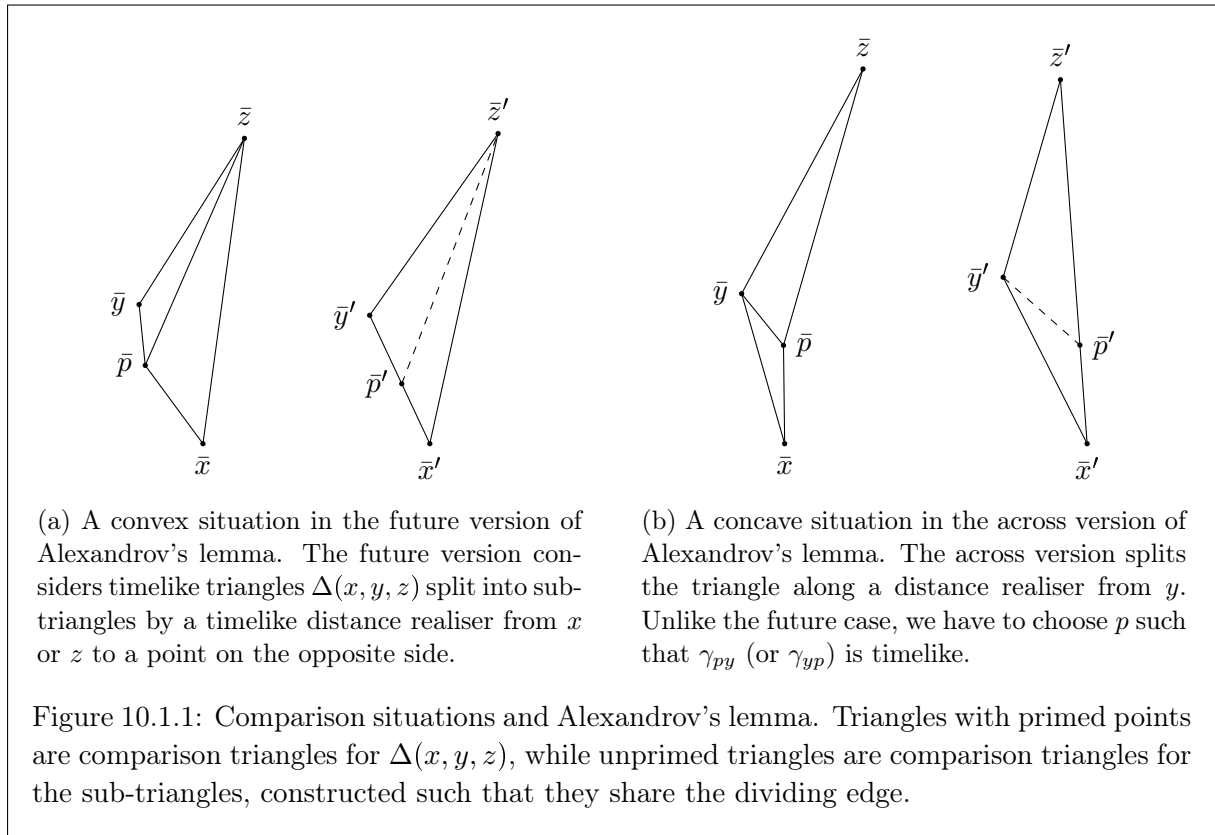
Definition 10.1.10 (Angle Condition Holds/Fails)

Let X be a regular Lorentzian pre-length space with timelike curvature bounded below by $K \in \mathbb{R}$ in the sense of Definition 8.2.7. Let $\alpha : [0, a] \rightarrow X$, $\beta : [0, b] \rightarrow X$ be timelike distance realisers of arbitrary time-orientation, such that $x := \alpha(0) = \beta(0)$ and $x, \alpha(a), \beta(b)$, in some causal order, are the vertices of an admissible causal triangle in X .⁴ We say that the angle condition holds at x if (8.2.5) is satisfied at x , with respect to the curvature bound K on X . Similarly, we say that the angle condition fails to hold at x if (8.2.5) is not satisfied at x . In particular, the angle condition can be said to hold/fail at vertices between timelike sides of an admissible causal triangle.

A key component of the proof of the globalisation theorem (Theorem 9.2.2) in the case of curvature bounded above was the so-called gluing lemma (Lemma 9.1.1), which roughly states that when two timelike sub-triangles satisfy curvature bounds, then the large timelike triangle formed by combining the two triangles must also satisfy that curvature bound. While this version of the gluing lemma, which utilises the triangle comparison condition from Definition 8.1.20, does not hold in the case of curvature bounded below, we propose an alternative result in the same spirit, instead using the aforementioned angle comparison condition. We propose that if the angle condition fails to hold at some vertex in a timelike triangle, then splitting the triangle into two timelike sub-triangles along one of the adjacent sides results in at least one sub-triangle with failing angle condition. In particular, there are three candidate angles at which the angle condition may fail. In our eventual proof of the synthetic Lorentzian Toponogov globalisation theorem, we will use this to induce a contradiction, starting with an arbitrary admissible causal triangle for which the angle condition fails at a given vertex and showing that there must exist a sub-triangle contained within a comparison neighbourhood, for which the angle condition also fails at some vertex.

As with the gluing lemma for spaces with curvature bounded above, the proof of result will utilise the twin Lorentzian versions of the Alexandrov's lemma to compare the properties of the two sub-triangles. Recall that we obtain two Alexandrov lemmas in Lorentzian signature, since

⁴Recall we assume that the vertices of admissible causal triangles satisfy size-bounds cf. Definition 8.1.11.



we may divide a timelike triangle along its long or short side (resulting in the across or future version of the lemma respectively). In what follows, we provide condensed statements of these lemmata and direct the reader to [128, Proposition 2.42, 2.43] and [126, Lemma 4.2.1, 4.2.2] for their proofs.⁵

Proposition 10.1.11 (Alexandrov's Lemma: Future Version)

Let X be a Lorentzian pre-length space, $\Delta(x, y, z)$ be a timelike triangle satisfying size-bounds for $K \in \mathbb{R}$, and $\Delta(\bar{x}', \bar{y}', \bar{z}')$ be a comparison triangle for $\Delta(x, y, z)$ in $\mathbb{L}^2(K)$.

Let $p \in \gamma_{xy}$ such that the timelike distance realiser γ_{pz} exists. Consider the sub-triangles $\Delta(x, p, z)$ and $\Delta(p, y, z)$ and construct their comparison triangles $\Delta(\bar{x}, \bar{p}, \bar{z})$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ in $\mathbb{L}^2(K)$, such that they share the side γ_{pz} , see Figure 10.1.1a. Finally, let \bar{p}' be a comparison point for p in $\Delta(\bar{x}', \bar{y}', \bar{z}')$. Then:

- (i) The configuration of $\Delta(\bar{x}, \bar{p}, \bar{z})$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ is convex (i.e. $\tau(p, z) = \tau(\bar{p}, \bar{z}) \leq \tau(\bar{p}', \bar{z}')$) if and only if $\angle_{\bar{p}}(\bar{y}, \bar{z}) \leq \angle_{\bar{p}}(\bar{x}, \bar{z})$.

⁵While these references prove the results for $K = 0$, generalising to non-zero K is straightforward under the appropriate size-bounds.

- (ii) The configuration of $\Delta(\bar{x}, \bar{p}, \bar{z})$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ is concave (i.e. $\tau(p, z) = \tau(\bar{p}, \bar{z}) \geq \tau(\bar{p}', \bar{z}')$) if and only if $\angle_{\bar{p}}(\bar{y}, \bar{z}) \geq \angle_{\bar{p}}(\bar{x}, \bar{z})$.

The analogous result with $p \in \gamma_{yz}$ and a timelike distance realiser γ_{xp} also holds.

Proposition 10.1.12 (Alexandrov's Lemma: Across Version)

Let X be a Lorentzian pre-length space, $\Delta(x, y, z)$ be a timelike triangle satisfying size-bounds for $K \in \mathbb{R}$, and $\Delta(\bar{x}', \bar{y}', \bar{z}')$ be a comparison triangle for $\Delta(x, y, z)$ in $\mathbb{L}^2(K)$.

Let $p \in \gamma_{xz}$ with $p \ll y$ such that the timelike distance realiser γ_{py} exists. Consider the sub-triangles $\Delta(x, p, y)$ and $\Delta(p, y, z)$ and construct their comparison triangles $\Delta(\bar{x}, \bar{p}, \bar{y})$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ in $\mathbb{L}^2(K)$, such that they share the side γ_{py} , see Figure 10.1.1b. Finally, let \bar{p}' be a comparison point for p in $\Delta(\bar{x}', \bar{y}', \bar{z}')$. Then:

- (i) The configuration of $\Delta(\bar{x}, \bar{p}, \bar{y})$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ is convex (i.e. $\tau(p, y) = \tau(\bar{p}, \bar{y}) \leq \tau(\bar{p}', \bar{y}')$) if and only if $\angle_{\bar{p}}(\bar{y}, \bar{z}) \leq \angle_{\bar{p}}(\bar{x}, \bar{y})$.
- (ii) The configuration of $\Delta(\bar{x}, \bar{p}, \bar{y})$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ is concave (i.e. $\tau(p, y) = \tau(\bar{p}, \bar{y}) \geq \tau(\bar{p}', \bar{y}')$) if and only if $\angle_{\bar{p}}(\bar{y}, \bar{z}) \geq \angle_{\bar{p}}(\bar{x}, \bar{y})$.

The analogous result with $y \ll p$ and a timelike distance realiser γ_{yp} also holds.

Note that the inequalities in the above lemmata are strict, except in the case where the configuration of triangles is simultaneously convex and concave, when all inequalities become equality and $\Delta(x, y, z)$ is degenerate. Furthermore, the convexity/concavity conditions are respectively automatically satisfied if X has timelike curvature bounded below/above by K and $\Delta(x, y, z)$ is contained within a comparison neighbourhood.




Figure 10.1.2: The gluing lemma for curvature bounded below: if the angle condition at x in $\Delta(x, y, z)$ fails to hold (shown in black), then at least one of the three angles conditions at p or x in the smaller triangles fail to hold (shown in red).

Lemma 10.1.13 (Gluing Lemma: Curvature Bounded Below)

Let X be a strongly causal, locally causally closed, geodesic, and regular Lorentzian pre-length space with curvature bounded below by $K \in \mathbb{R}$ in the sense of angle comparison. Let $\Delta(x, y, z)$ be an arbitrary timelike triangle in X and $\Delta(\bar{x}', \bar{y}', \bar{z}')$ be a comparison triangle for $\Delta(x, y, z)$ in $\mathbb{L}^2(K)$. Assume that the angle condition fails to hold at x in $\Delta(x, y, z)$. Then given a point

- (i) p on γ_{xy} , the angle condition fails for at least one of the following: p on $\Delta(x, p, z)$, x in $\Delta(x, p, z)$, or p in $\Delta(p, y, z)$.
- (ii) p on γ_{xz} with $y \ll p$ (or with p and y in reverse order), the angle condition fails for at least one of the following: p in $\Delta(x, y, p)$, x in $\Delta(x, y, p)$, or p in $\Delta(y, p, z)$.

Analogous statements hold if the angle condition fails at y in $\Delta(x, y, z)$ and p is on γ_{xy} or γ_{yz} , or the angle condition fails at z in $\Delta(x, y, z)$ and p is on γ_{yz} or γ_{xz} (and is timelike related to y).

Proof. We prove point (i) of the result in the case where the angle condition fails to hold at x in $\Delta(x, y, z)$. Each of the remaining statements can be proven similarly using the various configurations of Proposition 10.1.12 and Proposition 10.1.11 corresponding to the situation produced.

Denote by α the distance realiser (which exists since X is geodesic) from p to z and let β_- and β_+ the parts of γ_{xy} which go from p to x and p to y , respectively. Consider the comparison triangles $\Delta(\bar{x}, \bar{p}, \bar{z})$ for $\Delta(x, p, z)$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ for $\Delta(p, y, z)$ in $\mathbb{L}^2(K)$, such that they share the edge $\gamma_{\bar{p}\bar{z}}$. Also consider the comparison triangle $\Delta(\bar{x}', \bar{y}', \bar{z}')$ for $\Delta(x, y, z)$ and the comparison point \bar{p}' for p in $\Delta(\bar{x}', \bar{y}', \bar{z}')$. Assume the angle condition at x in $\Delta(x, p, z)$ holds, i.e. $\angle_x(\gamma_{xz}, \beta_-) \geq \angle_{\bar{x}}(\bar{p}, \bar{z})$ (here the sign of the angle is -1), otherwise we are done. We now show the angle condition at p must fail in either $\Delta(x, p, z)$ or $\Delta(p, y, z)$.

For the angle at \bar{x} in $\Delta(\bar{x}, \bar{p}, \bar{z})$ and \bar{x}' in the sub-triangle⁶ $\Delta(\bar{x}', \bar{p}', \bar{z}')$, we know

$$\angle_{\bar{x}}(\bar{p}, \bar{z}) \leq \angle_x(\gamma_{xz}, \gamma_{xy}) < \angle_{\bar{x}'}(\bar{y}', \bar{z}') = \angle_{\bar{x}'}(\bar{p}', \bar{z}'), \quad (10.1.9)$$

since the angle condition holds at x in $\Delta(x, p, z)$ and fails at x in $\Delta(x, y, z)$. Furthermore, as \bar{p}' is a comparison point for p , we have $\tau(\bar{x}, \bar{z}) = \tau(\bar{x}', \bar{z}')$ and $\tau(\bar{x}, \bar{p}) = \tau(\bar{x}', \bar{p}')$ and Corollary 8.2.2 yields

$$\tau(p, z) = \tau(\bar{p}, \bar{z}) > \tau(\bar{p}', \bar{z}'), \quad (10.1.10)$$

since the length of short sides is monotonically decreasing with respect to the opposite angle. By using the version of Alexandrov's Lemma as given in Proposition 10.1.11, the comparison

⁶Recall that in general, the sub-triangle $\Delta(\bar{x}', \bar{p}', \bar{z}') \neq \Delta(\bar{x}, \bar{p}, \bar{z})$.

triangles $\Delta(\bar{x}, \bar{p}, \bar{z})$ and $\Delta(\bar{p}, \bar{y}, \bar{z})$ therefore form a concave situation and

$$\angle_{\bar{p}}(\bar{x}, \bar{z}) < \angle_{\bar{p}}(\bar{y}, \bar{z}). \quad (10.1.11)$$

Moreover, by [Proposition 8.2.6](#), we have $\angle_x(\beta_-, \alpha) = \angle_x(\alpha, \beta_+)$. If the angle condition were to hold at p in both $\Delta(x, p, z)$ and $\Delta(p, y, z)$, we would have

$$\angle_{\bar{p}}(\bar{x}, \bar{z}) \geq \angle_p(\beta_-, \alpha) = \angle_p(\alpha, \beta_+) \geq \angle_{\bar{p}}(\bar{y}, \bar{z}), \quad (10.1.12)$$

contradicting (10.1.11). Hence, the angle condition must fail at p either in $\Delta(x, p, z)$ or $\Delta(p, y, z)$, if it does not fail at x in $\Delta(x, p, z)$. \square

In particular, since Lorentzian length spaces are locally causally closed and globally hyperbolic Lorentzian length spaces are both geodesic and strongly causal, the above result holds in our setting of globally hyperbolic, regular Lorentzian length spaces.

10.2 Toponogov Globalisation

The main goal of this section is to prove a synthetic Lorentzian analogue of Toponogov's globalisation theorem (see [Theorem 7.2.6](#)) for lower timelike curvature bounds. As previously discussed, we will be working in the setting of connected, globally hyperbolic, regular Lorentzian length spaces with time function, with second countability being a sufficient condition to ensure the existence of a time function by [Proposition 10.1.2](#). The set-up we have detailed also contains the assumptions of the Toponogov-style globalisation theorem for globally hyperbolic Lorentzian manifolds established in [[118](#), [Theorem 4.4](#)]. In particular, [[116](#), [Example 3.24](#)] shows that strongly causal spacetimes are regular Lorentzian length spaces and [[131](#)] note that globally hyperbolic spacetimes admit time functions.⁷

Our approach follows that of [[133](#)] in the metric setting and will hence consider curvature bounded below in the sense of the angle condition, as in [Definition 8.2.7](#). We proceed roughly as follows: consider an arbitrarily large admissible causal triangle in a space with curvature bounded below by some $K \in \mathbb{R}$ which satisfies size-bounds for K and for which angle comparison fails, as in [Definition 10.1.10](#). By contradiction, we will prove that such a triangle cannot exist; if this triangle did exist, it could be broken into smaller admissible causal sub-triangles using [Lemma 10.1.13](#), resulting in a sub-triangle contained in a comparison neighbourhood for which angle comparison fails at some vertex, contradicting the curvature bound. Here the size

⁷Strictly speaking, [[118](#)] consider smooth timelike sectional curvature bounded above, however for spacetimes this is equivalent to synthetic timelike curvature bounded below, see [[149](#)]. We will discuss this in more detail in [Section 10.3](#).

of timelike triangles is measured with respect to the null distance from Definition 10.1.4 and corresponds to its null distance diameter (10.1.4). Note that since the null distance is finite by Proposition 10.1.7, the d_T -diameter of a triangle is always finite and the notion of “smaller” triangles makes sense. Moreover, by [123, Remark 3.12], if there exists an admissible causal triangle with a vertex at which the angle condition fails, then there exists a timelike triangle with a vertex at which the angle condition fails. Hence, by contraposition, it is sufficient to show that no timelike triangles possess vertices at which the angle condition fails. In particular, we need not consider triangles with a null side at any stage.

The proof itself is split into four main parts, with the first three parts yielding properties of the angle condition at the vertices of an arbitrary timelike triangle $\Delta(x, y, z)$, provided the angle condition holds at all vertices of all timelike triangles which are almost as large with respect to the null distance. The final part demonstrates that (local) curvature bounded below by $K \in \mathbb{R}$ is sufficient to guarantee that all timelike triangles below a certain size with respect to the null distance do not have any vertices at which the angle condition fails, yielding our contradiction.

Our first result demonstrates that, if an angle condition fails at some vertex of a timelike triangle $\Delta(x, y, z)$, but does not fail at any vertex in any timelike triangle which is almost as large as $\Delta(x, y, z)$, then there exists a timelike triangle $\Delta(x'', y'', z'')$ of comparable size to $\Delta(x, y, z)$ for which the angle condition fails at a vertex corresponding to an angle of sign $\sigma = +1$. Consequently, to prove that the angle condition does not fail at any vertex of any triangle, it suffices to check that the angle condition does not fail at vertices corresponding to an angle of sign $\sigma = +1$.

Proposition 10.2.1 (Sub-Triangles with Failing ($\sigma = +1$)-Angles Exist)

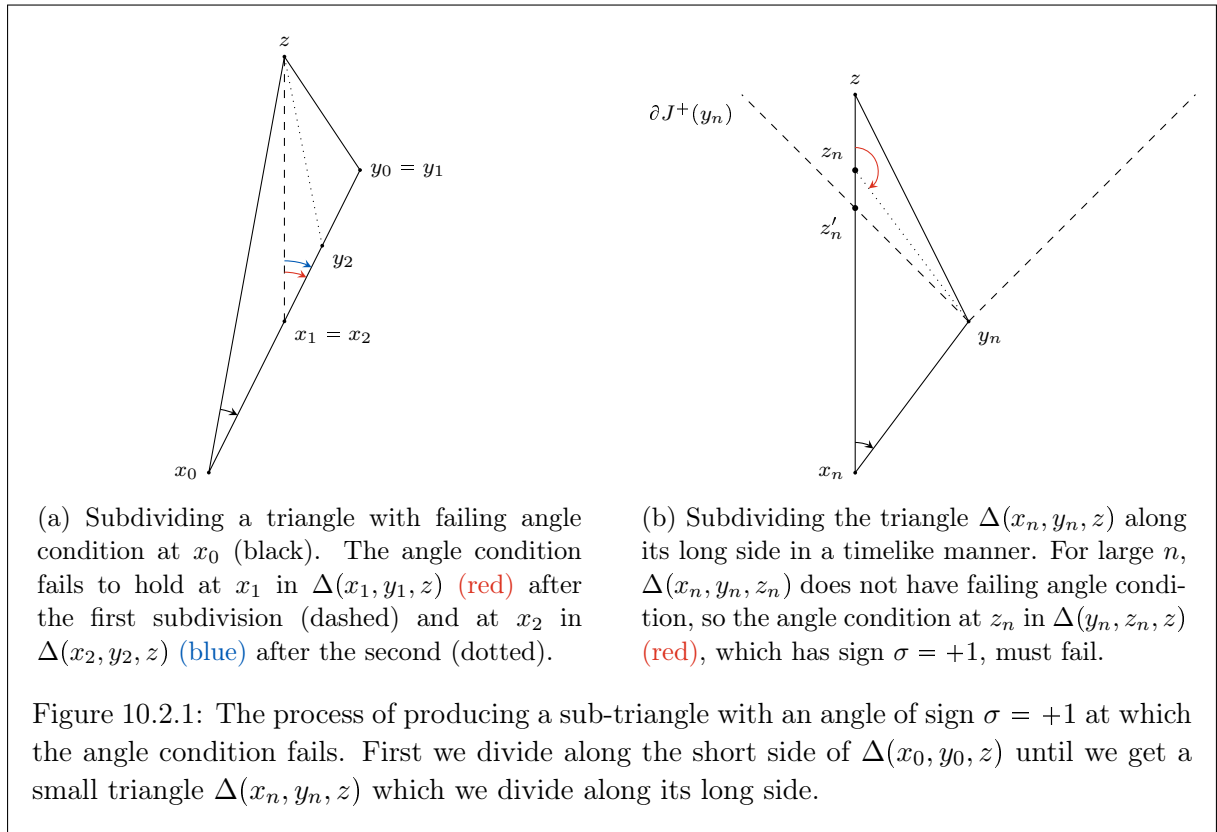
Let X be a connected, globally hyperbolic, regular Lorentzian length space with time function T and curvature bounded below by $K \in \mathbb{R}$. Fix $0 < \varepsilon < 1$ and let $\Delta(x, y, z)$ be a timelike triangle in X for which the angle condition fails at some vertex. If the angle condition holds at each angle in every timelike triangle $\Delta(x', y', z')$ satisfying both

$$(i) \quad x \leq x' \ll y' \ll z' \leq z$$

$$(ii) \quad d_T(x', z') \leq (1 - \varepsilon)d_T(x, z)$$

then there is at least one timelike triangle $\Delta(x'', y'', z'')$ with $x \leq x'' \ll y'' \ll z'' \leq z$ such that the angle condition fails at y'' , i.e. $\angle_{y''}(x'', z'') > \angle_{\bar{y}''}(\bar{x}'', \bar{z}'')$.

Proof. If the angle condition fails at y in $\Delta(x, y, z)$ then we are done and $\Delta(x'', y'', z'') = \Delta(x, y, z)$. So assume that the angle condition fails at either x or z in $\Delta(x, y, z)$, say at x without



loss of generality (the case where the angle condition fails at z is analogous under reversal of the time orientation).

We split the side γ_{xy} at its midpoint with respect to the null distance d_T , that is, at $p_1 \in \gamma_{xy}$ such that $d_T(x, p_1) = d_T(p_1, y) = \frac{1}{2}d_T(x, y)$. The distance realiser⁸ $\gamma_{p_1 z}$ then divides $\Delta(x, y, z)$ into two timelike sub-triangles $\Delta(x, p_1, z)$ and $\Delta(p_1, y, z)$. By Lemma 10.1.13 (the gluing lemma), we get that either the angle condition fails at p in $\Delta(x, p_1, z)$, in which case the result follows with $\Delta(x'', y'', z'') = \Delta(x, p_1, z)$, or the angle condition fails at either x in $\Delta(x, p_1, z)$ or p_1 in $\Delta(p_1, y, z)$. In either of the latter two cases, rename the triangle where the angle condition fails to $\Delta(x_1, y_1, z)$ with the angle condition now failing at x_1 (in the case where both triangles have failing angle condition, simply pick your favourite).

We may then take $\Delta(x_1, y_1, z)$ and repeat this process arbitrarily many times as in Figure 10.2.1a, terminating the process if we obtain a triangle with an angle of sign $\sigma = +1$ at which the angle condition fails. If we never find such a triangle, then we obtain a sequence of pairs $x_n \ll y_n$ on the side γ_{xy} , such that the angle condition fails at x_n in $\Delta(x_n, y_n, z)$ for each n . In

⁸This distance realiser need not be unique, since our space is not uniquely geodesic in general. If this is the case, the reader is free to make a choice.

particular, $d_T(x_n, y_n) = T(y_n) - T(x_n)$ by (10.1.3b). Furthermore, since the triangle $\Delta(x_n, y_n, z)$ is always subdivided according to the d_T -midpoint p_{n+1} of $\gamma_{x_n y_n}$ (which becomes either x_{n+1} or y_{n+1} in the subsequent triangle), then $d_T(x_n, y_n) = T(y_n) - T(x_n) \rightarrow 0$. As the pairs (x_n, y_n) lie on the distance realiser γ_{xy} and time functions are strictly monotone on causal curves, this implies that the sequences x_n and y_n have a common limit point $x^* \in \gamma_{xy}$, with $x_n \nearrow x^*$ and $y_n \searrow x^*$.

We now consider two cases. Firstly, if $d_T(x^*, z) < (1 - \varepsilon)d_T(x, z)$, then for sufficiently large n , we have $d_T(x_n, z) \leq (1 - \varepsilon)d_T(x, z)$ and the triangle $\Delta(x_n, y_n, z)$ is sufficiently small that, by assumption, the angle condition must hold at each of its angles. Since the angle condition fails at x_n in $\Delta(x_n, y_n, z)$ by construction, we have a contradiction.

Secondly, if $d_T(x^*, z) \geq (1 - \varepsilon)d_T(x, z)$, the triangle $\Delta(x_n, y_n, z)$ is always too large with respect to d_T for the previous argument to apply. It therefore becomes necessary to split some $\Delta(x_n, y_n, z)$ along its long side $\gamma_{x_n z}$ in the following manner, in order to obtain sub-triangle with an angle of sign $\sigma = +1$ at which the angle condition fails. Let z'_n be the point on the intersection of the distance realiser $\gamma_{x_n z}$ with the boundary $\partial J^+(y_n)$, i.e. $x_n \ll z'_n \ll z$, $y_n \leq z'_n$, and $y_n \not\ll z'_n$, see Figure 10.2.1b. This point is unique by regularity of X .⁹ By global hyperbolicity of X , the causal diamond $J(x, z)$ is compact and we may assume that $z'_n \rightarrow z^*$ in X , where we pass to a subsequence if necessary.

We claim that $z^* = x^*$. By construction, $y_n \not\ll z'_n$, so $\tau(y_n, z'_n) = 0$. Therefore, by continuity of τ (see [116, Theorem 3.30]), in the limit $n \rightarrow \infty$ (where $y_n \searrow x^*$), we obtain $\tau(x^*, z^*) = 0$. Furthermore, since $y_n \leq z'_n$ and X is causally closed by Proposition 8.1.19, we have $x^* \leq z^*$. Additionally, since x_n, z'_n , and z all lie along the distance realiser $\gamma_{x_n z}$, additivity of length (8.1.5) yields

$$0 < \tau(x_n, z) = \tau(x_n, z'_n) + \tau(z'_n, z), \quad (10.2.1)$$

with a further application of the continuity of τ implying that

$$0 < \tau(x^*, z) = \tau(x^*, z^*) + \tau(z^*, z) = \tau(z^*, z) \quad (10.2.2)$$

in the limit $n \rightarrow \infty$, where positivity of $\tau(x^*, z)$ follows from the fact that $x^* \in \gamma_{xy}$. In particular, since X is geodesic (see [116, Theorem 3.28]), the points x^*, z^* , and z lie along a distance realiser $\gamma_{x^* z}$, where the segment from x^* to z^* contributes zero τ -length. Since $x^* \ll z$, regularity implies that this distance realiser is timelike, so $\tau(x^*, z^*) = 0 \implies x^* = z^*$. In summary, as $n \rightarrow \infty$, all of the sequences x_n, y_n , and z'_n tend to x^* .

⁹Indeed, assume that there exists another such point $\tilde{z}_n \in \gamma_{x_n z}$. Since $\gamma_{x_n z}$ is timelike, by regularity either $z'_n \ll \tilde{z}_n$ or $\tilde{z}_n \ll z'_n$. Without loss of generality, choose $z'_n \ll \tilde{z}_n$, which implies $y_n \leq z'_n \ll \tilde{z}_n$. By push up (8.1.2), $y_n \ll \tilde{z}_n$, yielding a contradiction

Returning to the triangle $\Delta(x_n, y_n, z)$, in which the angle condition fails at x_n , we find a timelike distance realiser $\gamma_{y_n z_n}$ from y_n to a point z_n on $\gamma_{x_n z}$ which is slightly in the future of z'_n . Next we use $\gamma_{y_n z_n}$ to divide $\Delta(x_n, y_n, z)$ into two timelike sub-triangles, $\Delta(x_n, y_n, z_n)$ and $\Delta(y_n, z_n, z)$. Applying Lemma 10.1.13 then tells us that either the angle condition fails at z_n in $\Delta(y_n, z_n, z)$, or the angle condition fails at x_n or z_n in $\Delta(x_n, y_n, z_n)$. However, for large enough n the latter case cannot occur, since the points x_n, y_n , and z_n are all close enough to x^* that we have $d_T(x_n, z_n) \leq (1 - \varepsilon)d_T(x, z)$ and no angle conditions fail in $\Delta(x_n, y_n, z_n)$, by assumption. Consequently, the angle condition fails at z_n in $\Delta(y_n, z_n, z)$ for large enough n and the result follows with $\Delta(x'', y'', z'') = \Delta(y_n, z_n, z)$ \square

From this point onwards, we therefore only consider triangles which have failing angle condition at a vertex corresponding to an angle of sign $\sigma = +1$. Namely, the angle condition fails at the vertex opposite their longest side, and we divide these triangles along their shorter sides. This avoids the issue where splitting a timelike triangle along its longest side may not result in timelike sub-triangles.

Following the work of Plaut across two papers [142, 143], Lang and Schroeder [133] provide a “cat’s cradle” construction for use in proving Toponogov’s theorem for metric length spaces. This cat’s cradle construction is a recursive decomposition of an arbitrary triangle $\Delta(x, y, z)$ into smaller triangles in such a way that we can guarantee certain angle conditions hold in $\Delta(x, y, z)$. In the second result of this section, we prove that the cat’s cradle can also be performed in the Lorentzian setting. This is in spite of the challenge posed by sub-triangles not needing to be small with respect to the metric topology induced by d , even if their sides have small τ -length, in contrast to the metric setting. We circumvent this issue by performing the decomposition in such a way that subsequent triangles are instead small with respect to the null distance d_T .

Proposition 10.2.2 (Cat’s Cradle)

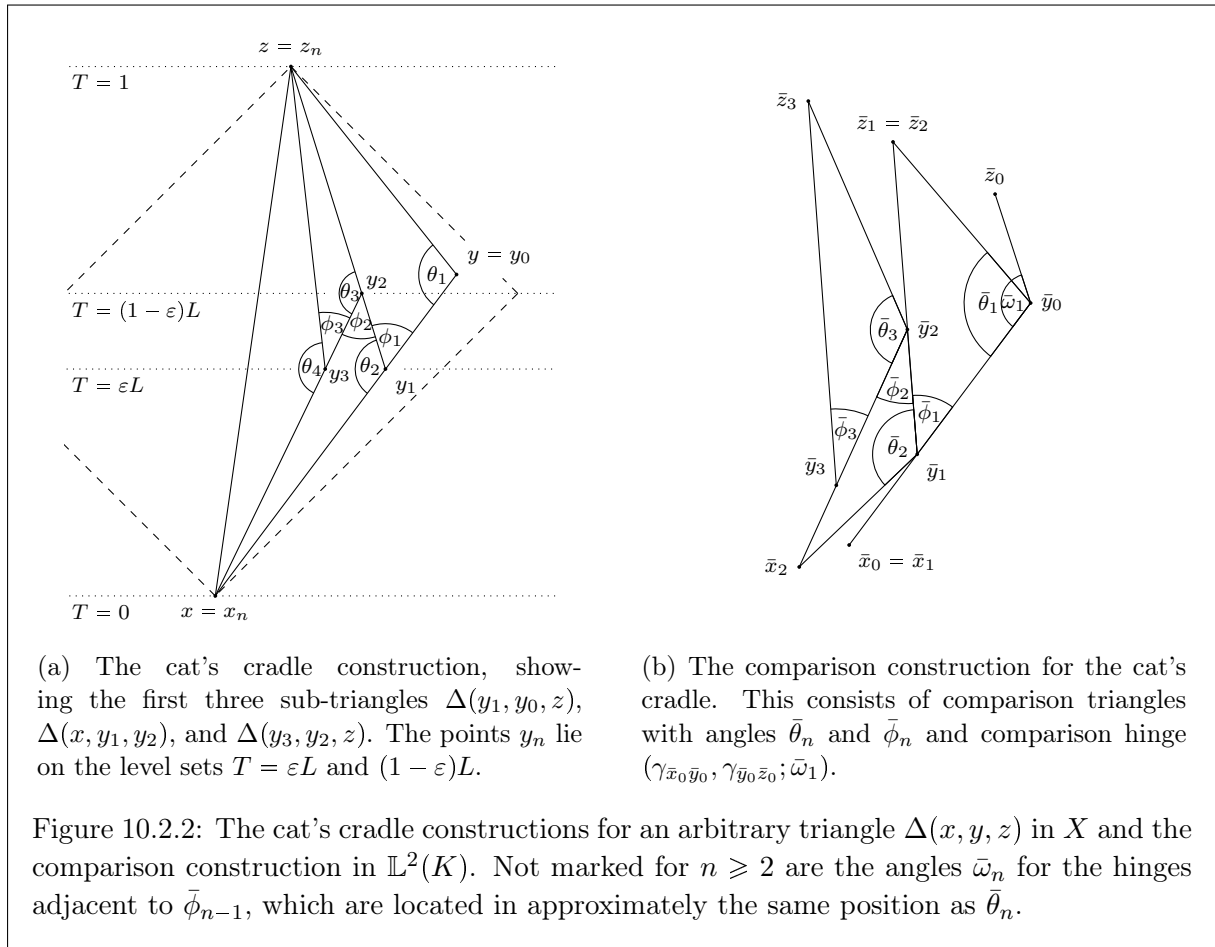
Let X be a connected, globally hyperbolic, regular Lorentzian length space with time function T and curvature bounded below by $K \in \mathbb{R}$. Fix $0 < \varepsilon < \frac{1}{2}$ and let $\Delta(x, y, z)$ be a timelike triangle in X . If, for every timelike triangle $\Delta(x', y', z')$ satisfying both

$$(i) \quad x \leq x' \ll y' \ll z' \leq z$$

$$(ii) \quad d_T(x', z') \leq (1 - \varepsilon)d_T(x, z)$$

the angle condition holds at all vertices of $\Delta(x', y', z')$, then the angle condition also holds at y in $\Delta(x, y, z)$.

Proof. (sketch) Since the proof of Proposition 10.2.2 is both extensive and technical, we provide only a brief overview here and instead direct the reader to Appendix D for full details.



We begin by considering an arbitrary triangle $\Delta(x, y, z)$ of null distance diameter $d_T(x, z) = L$ and performing a recursive decomposition of said triangle. We do so by carving off triangles whose null distance diameter (10.1.4) is $(1 - \varepsilon)L$ and hence for whom the angle condition holds at all vertices, see Figure 10.2.2a. From this construction, we obtain the sequence of inequalities

$$0 < l_0 \leq l_1 \leq \dots \leq \tau(x, z). \tag{10.2.3}$$

Here $l_n := \tau(x, y_n) + \tau(y_n, z)$ measures how much time separation it takes to get from x to z via the short sides of the triangle once we have removed n pieces. By then performing a corresponding recursive construction in the model space $\mathbb{L}^2(K)$ using both comparison hinges and comparison triangles, see Figure 10.2.2b, we obtain a second sequence of inequalities

$$\tau(\bar{x}_0, \bar{z}_0) \leq \tau(\bar{x}_1, \bar{z}_1) \leq \dots \tag{10.2.4}$$

By showing that these two sequences converge to the same limit, we find that the hinge compar-

ison condition $\tau(x, z) \geq \tau(\bar{x}_0, \bar{z}_0)$ holds at y in $\Delta(x, y, z)$. Then by monotonicity of the law of hyperbolic cosines, this implies that the angle at \bar{y} in the comparison triangle $\Delta(\bar{x}, \bar{y}, \bar{z})$ is bigger than the angle at y in $\Delta(x, y, z)$, so the angle condition holds at y . \square

Collecting the previous two propositions, we can deduce that the angle conditions hold in arbitrarily large timelike triangles so long as they hold for all triangles which are not quite as large. In particular, [Proposition 10.2.1](#) says any angle condition fails in a given triangle $\Delta(x, y, z)$, then we can find a triangle $\Delta(x', y', z')$ of comparable size and position for which the angle condition fails at a vertex corresponding to an angle with sign $\sigma = +1$. On the other hand, [Proposition 10.2.2](#) says if the angle conditions hold in all triangles which are almost as large as $\Delta(x', y', z')$ then the angle condition holds at vertices of $\Delta(x', y', z')$ corresponding to angles of sign $\sigma = +1$, yielding a contradiction. We formulate this more precisely as follows:

Corollary 10.2.3 (Angle Conditions Hold for Slightly Larger Triangles)

Let X be a connected, globally hyperbolic, regular Lorentzian length space with time function T and curvature bounded below by $K \in \mathbb{R}$ in the sense of angle comparison. Fix $0 < \varepsilon < \frac{1}{2}$ and let $\Delta(x, y, z)$ be a timelike triangle in X which satisfies the size-bounds for K . If, for every timelike triangle $\Delta(x', y', z')$ satisfying both

$$(i) \quad x \leq x' \ll y' \ll z' \leq z \text{ and}$$

$$(ii) \quad d_T(x', z') \leq (1 - \varepsilon)d_T(x, z),$$

the angle condition holds at all vertices of $\Delta(x', y', z')$, then the angle condition also holds at each angle in $\Delta(x, y, z)$.

Proof. First, observe that our assumptions include the criteria for [Proposition 10.2.2](#) to hold. In particular, the angle condition must not fail at y in $\Delta(x, y, z)$. Now assume for a contradiction that the angle condition fails at either x or z in $\Delta(x, y, z)$. Then by [Proposition 10.2.1](#), there exists a timelike triangle $\Delta(x'', y'', z'')$ with $x \leq x'' \ll y'' \ll z'' \leq z$, such that the angle condition fails at y'' .

Furthermore, since $x \leq x'' \ll z'' \leq z$, the points x'' and z'' are contained in the timelike diamond $I(x, z)$, so by [Lemma 10.1.8](#) we have $d_T(x'', z'') \leq d_T(x, z)$. Suppose now that $\Delta(x', y', z')$ is some timelike triangle with $x'' \leq x' \ll y' \ll z' \leq z''$ and $d_T(x', z') \leq (1 - \varepsilon)d_T(x'', z'')$. Then we also have $x \leq x' \ll y' \ll z' \leq z$ and $d_T(x', z') \leq (1 - \varepsilon)d_T(x, z)$, hence by our assumptions,¹⁰ the angle condition holds at all vertices of all such $\Delta(x', y', z')$. Therefore, the result of Pro-

¹⁰Here we show that [Proposition 10.2.2](#) holds for $\Delta(x'', y'', z'')$ with the same ε as $\Delta(x, y, z)$. In fact, as $d_T(x'', z'') = \delta d_T(x, z)$ for $\delta \in (1 - \varepsilon, 1]$, if $\Delta(x, y, z)$ satisfies the assumptions of the proposition for some ε , then $\Delta(x'', y'', z'')$ does so for any value in $[1 - \frac{1-\varepsilon}{\delta}, \frac{1}{2})$.

position 10.2.2 holds not only for $\Delta(x, y, z)$ but also for $\Delta(x'', y'', z'')$. In particular, the angle condition does not fail at y'' in $\Delta(x'', y'', z'')$, yielding a contradiction. Consequently, the angle condition must not fail at x or z in $\Delta(x, y, z)$, from which our result follows. \square

Given that all timelike triangles below a certain size (with respect to the null distance) possess no vertices with failing angle condition, the previous proposition says that we may expand this a little bit and say that slightly larger triangles also do not have vertices with failing angle condition. By repeating this process, we may then infer that no timelike triangles exhibit failing angle condition. It remains to show that (local) curvature bounded below by K provide us with a positive lower bound on the size of timelike triangles which have a vertex at which the angle condition fails, that is, a size below which triangles do not possess vertices with failing angle condition.

Theorem 10.2.4 (Synthetic Lorentzian Toponogov Globalisation)

Let X be a connected, globally hyperbolic, regular Lorentzian length space with a time function T and timelike curvature bounded below by $K \in \mathbb{R}$. Then X has timelike curvature globally bounded below by K .

Proof. First recall that X having timelike curvature globally bounded below by K (in the sense of angle comparison) corresponds to X being a $(\geq K)$ -comparison neighbourhood with respect to the conditions given in Definition 8.2.7. However, condition (iv) is a local property, only requiring the germs of curves, hence it globalises trivially. Furthermore, recall that globally hyperbolic Lorentzian length spaces are geodesic with continuous time separation τ by [116, Theorems 3.28, 3.30], so properties (i) and (ii) are also satisfied globally by assumption. It remains to show that property (iii) holds, that is, arbitrarily large admissible causal triangles in X do not have vertices at which the angle condition fails. Following our discussion at the start of this section, it is sufficient to check that arbitrarily large timelike triangles do not have vertices at which the angle condition fails, since the existence of an admissible causal triangle with a failing angle condition implies the existence of a timelike triangle with a failing angle condition.

So assume for contradiction that $\Delta(x, y, z)$ is a timelike triangle in X which has failing angle condition at some vertex (this also permits triangles where the angle condition fails at multiple vertices). The triangle $\Delta(x, y, z)$ is contained in the diamond $J(x, z)$, which is compact by global hyperbolicity of X . Suppose that $\delta > 0$ is a greatest lower bound on the d_T -diameter of timelike triangles in $J(x, z)$ which exhibit a failing angle condition. In particular, all the vertices of any timelike triangle contained in $J(x, z)$ with d_T -diameter (10.1.4) less than δ must satisfy the angle condition, but there exist triangles with d_T -diameter greater than yet arbitrarily close

to δ which possess at least one vertex with failing angle condition.¹¹ Applying Corollary 10.2.3 to the triangles with d_T -diameter greater than but arbitrarily close to δ , then yields that they cannot have failing angle condition, which is a contradiction.

We now need to establish the existence of a positive greatest lower bound δ on the d_T -diameter of timelike triangles in $J(x, z)$ with failing angle condition. Define

$$A := \{ \text{diam}_T(\Delta(x', y', z')) \mid \Delta(x', y', z') \subset J(x, z) \text{ and exhibits failing angle condition} \}. \quad (10.2.5)$$

By assumption, the angle condition fails in $\Delta(x, y, z)$, so $\text{diam}_T(\Delta(x, y, z)) = d_T(x, z) \in A \neq \emptyset$. It follows that A has a greatest lower bound, which we verify is positive by demonstrating the existence of some positive lower bound. By the Lebesgue number lemma for the null distance (see Lemma 10.1.9), there exists some $\delta' > 0$ such that any causal diamond with d_T -diameter less than δ' which is contained in $J(x, z)$ is contained within a timelike diamond which is a comparison neighbourhood. In particular, any timelike triangle $\Delta(x', y', z') \subset J(x', z') \subseteq J(x, z)$ with d_T -diameter $d_T(x', z') < \delta'$ is contained in a comparison neighbourhood and so has vertices with failing angle condition. It follows that δ' is a positive lower bound for A , as required. \square

While the above theorem proves that curvature bounded below by $K \in \mathbb{R}$ in the sense of angle comparison globalises, an application of Proposition 8.2.9 yields that curvature bounded below in the senses of hinge and triangle comparison also globalise, provided we additionally assume that (8.2.6) also holds.

10.3 The Bonnet–Myers Theorem

The significance of the synthetic Lorentzian Toponogov globalisation theorem lies in its use to extend known results for spaces with global timelike curvature bounds, to those with local timelike curvature bounds, mirroring the use cases of Theorem 7.2.6 in the metric setting. Among these direct corollaries is the result we present in this section, namely a synthetic Lorentzian analogue of the Bonnet–Myers Theorem, in complete analogy to Theorem 7.2.7. We begin the section by deriving a bound on the finite diameter of a Lorentzian pre-length space, under the assumption of global curvature bounded below by some $K < 0$, then apply Theorem 10.2.4 to obtain said diameter bound on spaces with the same curvature bound locally. Further results which follow from Theorem 10.2.4 will be discussed in Chapter 11.

While we discussed the development of the Bonnet–Myers theorem for Riemannian geometry in Section 7.2.2, there are also a number of existing results in the Lorentzian literature

¹¹It is not strictly necessary that δ be a *greatest* lower bound. This restriction allows us to apply our earlier propositions for arbitrarily small constant $\varepsilon > 0$, however we only require $\varepsilon \in (0, \frac{1}{2})$.

which we feel it is important to address. In the smooth setting, Beem and Ehrlich [158, Theorem 9.5] have shown that globally hyperbolic spacetimes with timelike (sectional) curvature bounded below by some $K < 0$ have $\text{diam}(M) \leq \frac{\pi}{\sqrt{-K}}$, where the diameter is defined in terms of the time separation (8.1.1). More recently, in the synthetic setting, Cavalletti and Mondino [159, Proposition 5.10] have shown using optimal transport that Lorentzian pre-length spaces with suitable timelike measure contraction property, such as that implied by lower Ricci curvature bounds, possess a corresponding upper bound on their diameter, with [160] deriving a corresponding result in the setting of low-regularity spacetimes. While in the metric setting, we could be content with a result utilising bounds on the Ricci curvature (since Ricci curvature bounds are known to be a weaker requirement than sectional curvature bounds, see [161]), in the synthetic Lorentzian setting the hierarchy of timelike Ricci curvature bounds and timelike (sectional) curvature bounds via triangle comparison is still an open problem. Consequently, our approach to bounding the timelike diameter of a Lorentzian pre-length space via triangle comparison stands independently of the existing literature on Ricci bounds and introduces a novel method whose details have merit in their own right. Regardless, the similarity between the assumptions used in Theorem 10.3.4 and [159, Proposition 5.10], in particular global hyperbolicity, local causal closure, and geodesy, is certainly striking and hints at a strong interplay between the two notions of curvature in the synthetic Lorentzian setting, akin to their metric analogues.

At this stage it is worth raising a detail which we tacitly overlooked when stating the definition of synthetic timelike curvature bounds via triangle comparison way back in Definition 8.1.20. When initially defining synthetic timelike curvature bounds via triangle comparison, [116] extrapolated from the work of [119] and said that a spacetime should have synthetic timelike curvature bounded *above* by K when it has smooth sectional curvature bounded *above* by K and extrapolated from there. This is nice from the perspective of this text, where we compare metric and synthetic Lorentzian pictures, since it means, for example, that the Toponogov approach to globalisation is applied to curvature bounded below in both settings. However, it has one flaw when comparing to the smooth Lorentzian picture: [119] defined semi-Riemannian manifolds like spacetimes as having smooth sectional curvature bounded above by K when they have smooth timelike sectional curvature bounded *below* by K . Consequently, where we speak about spaces with synthetic timelike curvature bounded below by some $K < 0$, [158, 159] discuss spaces with smooth timelike sectional curvature bounded above by some $K < 0$. Furthermore, while in the metric and smooth Lorentzian literature, curvature bounded below (resp. above) by k implies curvature bounded below by all $k' < k$ (resp. above by all $k' > k$), synthetic timelike curvature bounds have the strange property that curvature bounded below (resp. above) by k implies curvature bounded below by all $k' > k$ (resp. above by all $k' < k$).

Before diving into the theorem proper, we provide a technical lemma which yields a non-degeneracy condition for adjacent timelike sub-triangles:

Lemma 10.3.1 (Non-degeneracy Condition)

Let X be a strongly causal, locally causally closed, regular, and geodesic Lorentzian pre-length space and let U be a comparison neighbourhood in X . Let $a \ll z$ in U and let γ_{az} be a timelike distance realiser from a to z in U . Let $x = \gamma_{az}(t)$ and let $y \in I(x, z)$. Assume that both $\Delta(a, x, y)$ and $\Delta(x, y, z)$ satisfy size-bounds. Let γ_{xy} be a timelike distance realiser from x to y , and denote by $\alpha_- = \gamma_{az}|_{[0,t]}$ and $\alpha_+ = \gamma_{az}|_{[t,1]}$ the parts of γ_{az} in the past and future of x , respectively.

If X has timelike curvature bounded below by K and $\Delta(x, y, z)$ is non-degenerate, then $\Delta(a, x, y)$ is also non-degenerate and the angles $\angle_x(\alpha_+, \gamma_{xy})$ and $\angle_x(\alpha_-, \gamma_{xy})$ are equal and positive.

Proof. As $\Delta(x, y, z)$ is non-degenerate, the corresponding comparison triangle in $\mathbb{L}^2(K)$ is non-degenerate, with the comparison angle at \bar{x} given by $\tilde{\angle}_x(y, z) > 0$ (see Corollary 8.2.3). Then by angle comparison, see Definition 8.2.7, it follows that $0 < \tilde{\angle}_x(y, z) \leq \angle_x(\alpha_+, \gamma_{xy})$. Since X is strongly causal, locally causally closed, and has timelike curvature bounded below, we can apply Proposition 8.2.6, yielding $\angle_x(\alpha_-, \gamma_{xy}) = \angle_x(\alpha_+, \gamma_{xy}) > 0$. Finally, by applying angle comparison again we obtain that $\tilde{\angle}_x(a, y) \geq \angle_x(\alpha_-, \gamma_{xy}) > 0$ and it follows from Corollary 8.2.3 that $\Delta(a, x, y)$ is non-degenerate. \square

Now let us turn to the Bonnet–Myers theorem for Lorentzian pre-length spaces itself. In light of the absence of a synthetic Lorentzian Toponogov globalisation theorem when this result was first derived, our first variation of the synthetic Lorentzian Bonnet–Myers theorem requires the assumption of global timelike curvature bounded below by some $K < 0$. While the original formulation of Theorem 10.3.3 found in [2, Theorem 4.12] considers the finite timelike diameter, the following result shows that we may consider the full τ -diameter from the offset, in-line with [158, 159]:

Lemma 10.3.2 (Time Separation is Finite)

Let X be a chronological and geodesic Lorentzian pre-length space with timelike curvature bounded below/above by $K \in \mathbb{R}$. Then τ is finite on X and the finite timelike diameter coincides with the τ -diameter. In particular, the τ -diameter may still be infinite but is not attained on X if so.

Proof. Consider any two points $x, y \in X$ with $x \ll y$ and let $\gamma : [a, b] \rightarrow X$ be the distance realiser between them. Since X is chronological and has timelike curvature bound, Lemma 8.1.22 implies that τ is locally finite. It follows that we may cover γ with with neighbourhoods on which τ is finite. In fact, we may use finitely many such neighbourhoods since γ is the continuous image of

a compact set. The segments of γ within each neighbourhood have finite τ -length and by (8.1.5), the τ -length of γ is the finite sum of these finite lengths, hence finite.

Since X is geodesic, we may perform this exercise for any pair of timelike related points and obtain a distance realiser of finite τ -length. Consequently, the time separation between any pair of timelike related points must be finite. It then follows that the finite timelike diameter is simply the τ -diameter. \square

In the statement of the following theorem, we impose a further non-degeneracy condition on the space X : for each pair of points $x \ll z$ in X , there exists a $y \in X$ such that $\Delta(x, y, z)$ is a non-degenerate timelike triangle. This allows us to apply Lemma 10.3.1 along timelike distance realisers and ensures that the space is not locally one-dimensional, reflecting our discussion concerning the Bonnet–Myers theorem for metric spaces in Section 7.2.2, where spaces isomorphic to \mathbb{R} , $(0, \infty)$, $[0, B]$ for all $B > \frac{\pi}{\sqrt{k}}$, or circles of radius greater than $\frac{1}{\sqrt{k}}$ are explicitly excluded from the result. By utilising the approach of [100, Proposition 8.44], it may also be possible to classify these pathological one-dimensional Lorentzian pre-length spaces to which the diameter bound does not apply, though this is an open problem left for future work.

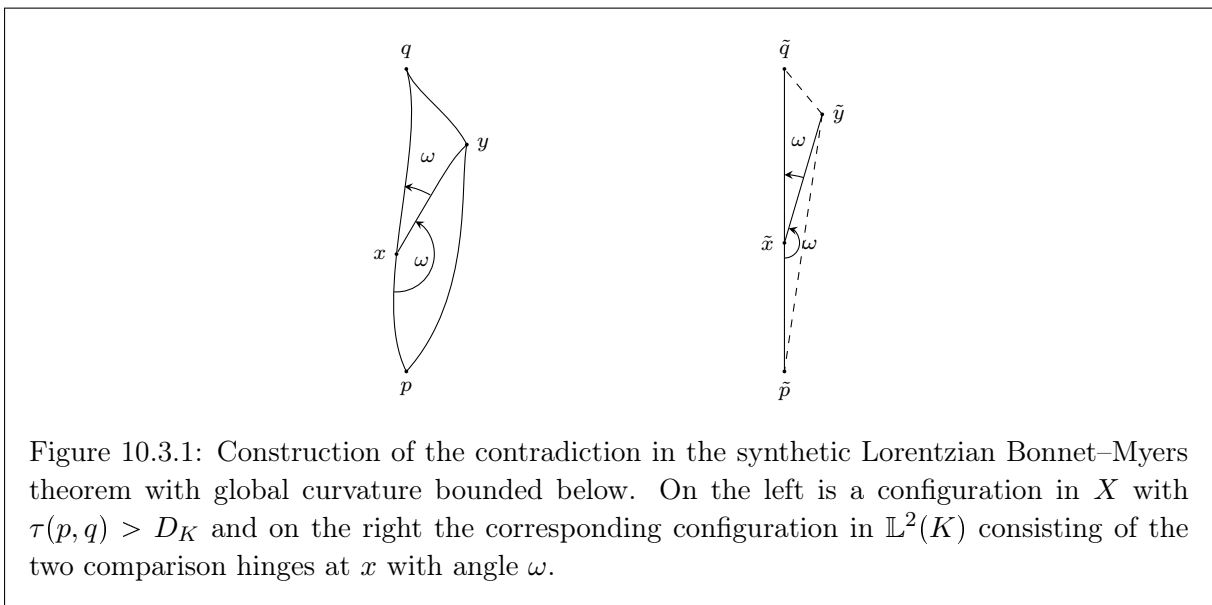


Figure 10.3.1: Construction of the contradiction in the synthetic Lorentzian Bonnet–Myers theorem with global curvature bounded below. On the left is a configuration in X with $\tau(p, q) > D_K$ and on the right the corresponding configuration in $\mathbb{L}^2(K)$ consisting of the two comparison hinges at x with angle ω .

Theorem 10.3.3 (Synthetic Lorentzian Bonnet–Myers: Global Curvature Bounds)

Let X be a chronological, strongly causal, locally causally closed, regular, and geodesic Lorentzian

pre-length space with global curvature bounded below¹² by $K < 0$. Assume that for each pair of points $x \ll z$ in X , there exists a point $y \in X$ such that $\Delta(x, y, z)$ is a non-degenerate timelike triangle. Then $\text{diam}(X) \leq D_K$.

Proof. Without loss of generality, we consider the case $K = -1$. All other cases can be recovered from this one by imposing the appropriate size-bounds and scaling.

Consider $p, q \in X$ with $\tau(p, q) = \pi + \varepsilon$ for some small $\varepsilon > 0$, such that $\tau(p, q) > \pi = D_{-1}$. Furthermore, let $\alpha : [0, \pi + \varepsilon] \rightarrow X$ be a timelike distance realiser from p to q parametrised by τ -arclength, which is possible by the discussion before [Definition 8.1.28](#). Now let $x := \alpha(t_-)$ and $z = \alpha(t_+)$ for $t_- = \frac{\pi}{2} + \frac{\varepsilon}{2}$ (the midpoint of α) and $t_+ = \frac{\pi}{2} + \frac{\pi}{8}$. The specific value of t_+ is not especially relevant, so long as it lies within (t_-, π) , since the point $z = \alpha(t_+)$ is only required in the construction of a triangle which satisfies size-bounds for $\mathbb{L}^2(-1)$ later in the proof. Finally, let $\alpha_- := \alpha|_{[0, t_-]}$ and $\alpha|_{[t_-, \pi + \varepsilon]}$ be the parts of α in the past and future of x , respectively.

By our non-degeneracy assumption, there exists a point $y \in I(x, z)$ such that $\Delta(x, y, z)$ is non-degenerate. Let γ_{xy} be a distance realiser from x to y . Then by [Lemma 10.3.1](#), we get that both $\Delta(p, x, y)$ and $\Delta(x, y, z)$ are non-degenerate. Furthermore, $\omega := \angle_x(\alpha_-, \gamma_{xy}) = \angle_x(\alpha_+, \gamma_{xy}) > 0$. We now claim that $\tau(p, q) < \tau(p, y) + \tau(y, q)$, contradicting the reverse triangle inequality. We label the lengths as follows

$$t := \tau(x, q) = \tau(p, x) = t_-, \quad a := \tau(p, y), \quad b := \tau(y, q), \quad \text{and} \quad m := \tau(x, y), \quad (10.3.1)$$

so our claim reads

$$2t < a + b. \quad (10.3.2)$$

We now construct the following comparison situation in $\mathbb{L}^2(-1)$ via hinge comparison, see [Definition 8.2.8](#): Let $(\tilde{\alpha}_-, \gamma_{\tilde{x}\tilde{y}}; \omega)$ and $(\tilde{\alpha}_+, \gamma_{\tilde{x}\tilde{y}}; \omega)$ be comparison hinges for $(\alpha_-, \gamma_{xy}; \omega)$ and $(\alpha_+, \gamma_{xy}; \omega)$ respectively, constructed such that they share the distance realiser $\gamma_{\tilde{x}\tilde{y}}$, see [Figure 10.3.1](#). In particular, these hinges have angle $\omega > 0$ and therefore yield non-degenerate timelike triangles $\Delta(\tilde{p}, \tilde{x}, \tilde{y})$ and $\Delta(\tilde{x}, \tilde{y}, \tilde{q})$ in $\mathbb{L}^2(-1)$. Note these triangles are *not* comparison triangles. Label the unknown side-lengths $\tilde{a} := \tau(\tilde{p}, \tilde{y})$ and $\tilde{b} := \tau(\tilde{y}, \tilde{q})$. By hinge comparison, we then obtain

$$a = \tau(p, x) \geq \tau(\tilde{p}, \tilde{x}) = \tilde{a} \quad (10.3.3)$$

$$b = \tau(x, q) \geq \tau(\tilde{x}, \tilde{q}) = \tilde{b}, \quad (10.3.4)$$

¹²Global curvature bounds guarantee the existence of a distance realiser in X between all pairs $a \ll b$ with $\tau(a, b) < D_K$. In this context, however, we require the existence of distance realisers for all timelike related pairs of points slightly further than D_K . X being geodesic is a sufficient condition for this.

so $a + b \geq \tilde{a} + \tilde{b}$. For (10.3.2) to hold, it remains to show that $\tilde{a} + \tilde{b} > 2t$.

We now prove this new claim. By applying reverse triangle inequality to the non-degenerate triangles $\Delta(\tilde{p}, \tilde{x}, \tilde{y})$ and $\Delta(\tilde{x}, \tilde{y}, \tilde{q})$ respectively, we obtain

$$m + \tilde{b} < t \quad \text{and} \quad 2m + \tilde{b} < m + t < \tilde{a}, \quad (10.3.5)$$

from which it follows $2m < \tilde{a} - \tilde{b}$. Recall that

$$\tilde{a} \leq a = \tau(p, y) \leq \tau(p, z) = \frac{\pi}{2} + \frac{\pi}{8}, \quad (10.3.6)$$

and since $\tilde{b} \geq 0$, it follows that

$$0 < 2m < \tilde{a} - \tilde{b} \leq \frac{\pi}{2} + \frac{\pi}{8} < \pi. \quad (10.3.7)$$

In particular, as cosine is non-negative and strictly decreasing on $[0, \frac{\pi}{2}]$, we have

$$0 < \cos\left(\frac{\tilde{a} - \tilde{b}}{2}\right) < \cos(m). \quad (10.3.8)$$

We now write down equations for $\omega = \angle_{\tilde{x}}(\tilde{p}, \tilde{y}) = \angle_{\tilde{x}}(\tilde{y}, \tilde{q})$ using the hyperbolic law of cosines, see Corollary 8.2.2 and [125, Lemma 2.4] with $K = -1$. Namely, we find

$$\cos(m) \cos(t) - \sin(m) \sin(t) \cosh(\omega) = \cos(\tilde{a}), \quad (10.3.9)$$

$$\cos(m) \cos(t) + \sin(m) \sin(t) \cosh(\omega) = \cos(\tilde{b}). \quad (10.3.10)$$

Adding these two equations together allows us to eliminate ω , yielding

$$2 \cos(m) \cos(t) = \cos(\tilde{a}) + \cos(\tilde{b}) = 2 \cos\left(\frac{\tilde{a} + \tilde{b}}{2}\right) \cos\left(\frac{\tilde{a} - \tilde{b}}{2}\right), \quad (10.3.11)$$

where we have used the cosine addition formula in the second step. Since ε is small, we have $\frac{\pi}{2} < t = \frac{\pi}{2} + \frac{\varepsilon}{2} < \pi$ and therefore $\cos(t) < 0$. Furthermore, by (10.3.8), we know $\cos(m) > \cos\left(\frac{\tilde{a} - \tilde{b}}{2}\right) > 0$, so

$$0 > \cos\left(\frac{\tilde{a} - \tilde{b}}{2}\right) \cos(t) > \cos(m) \cos(t) = \cos\left(\frac{\tilde{a} + \tilde{b}}{2}\right) \cos\left(\frac{\tilde{a} - \tilde{b}}{2}\right) \quad (10.3.12)$$

from which it follows that

$$\cos(t) > \cos\left(\frac{\tilde{a} + \tilde{b}}{2}\right) \quad (10.3.13)$$

Since the cosine function is strictly monotonically decreasing on $[\frac{\pi}{2}, \pi]$, we obtain

$$\tilde{a} + \tilde{b} > 2t, \quad (10.3.14)$$

as required. □

Applying [Theorem 10.2.4](#) allows us to directly re-frame the above result in terms of local timelike curvature bounds, at the cost of requiring X to adhere to the assumptions of the globalisation theorem.

Theorem 10.3.4 (Synthetic Lorentzian Bonnet–Myers: Local Curvature Bounds)

Let X be a connected, globally hyperbolic, and regular Lorentzian length space which has a time function T and local curvature bounded below by some $K < 0$. Assume that X possesses the following non-degeneracy condition: for each pair of points $x \ll z$ in X we find $y \in X$ such that $\Delta(x, y, z)$ is a non-degenerate timelike triangle. Then the diameter satisfies $\text{diam}(X) \leq D_K$.

When written in this way, it is not too hard to see that our Bonnet–Myers style results are direct synthetic extensions of the result for globally hyperbolic spacetimes with smooth timelike sectional curvature bounds in [[158](#), Theorem 9.5]. Indeed, globally hyperbolic spacetimes are regular Lorentzian length spaces with time function (see [[116](#), Example 3.24] and [[131](#)]) and it has recently been shown in [[149](#), Section 3] that, for strongly causal spacetimes, the notions of smooth timelike sectional curvature bounds and synthetic timelike curvature bounds agree.¹³

The following lemma, which was previously presented by [[158](#), Proposition 9.4] for spacetimes, gives us conditions under which the diameter of a Lorentzian pre-length space is not attained by a distance realiser. Following the publication of [Theorem 10.3.3](#) in [[2](#), Theorem 4.12], it was shown in [[162](#)] that additionally assuming D_K is attained by an *open* distance realiser $\gamma : (a, b) \rightarrow X$ forces the space to take the form of a Lorentzian warped product.

Lemma 10.3.5 (Unattainable Diameter)

Let X be a strongly causal Lorentzian pre-length space. If $\text{diam}(X)$ is finite, then it is not attained by a distance realiser on X . Furthermore, if X is a globally hyperbolic Lorentzian length space, then $\text{diam}(X)$ is never attained by a distance realiser on X , independently of whether it is finite.

¹³Up to the disparity in terminology discussed earlier, where smooth timelike bounds above are synthetic timelike bounds below and vice versa.

Proof. Let X be a strongly causal Lorentzian pre-length space. Assume for contradiction that $\text{diam}(X)$ is finite and attained by some $p, q \in X$, that is, $\tau(p, q) = \text{diam}(X)$. Then, by strong causality, q is contained within some timelike diamond $I(q_-, q_+)$ with $q_- \ll q \ll q_+$. In particular, since the diameter of the space is finite, $q \neq q_+$, as $\tau(q, q)$ is only non-zero if it is infinite. Therefore,

$$\tau(p, q_+) \geq \tau(p, q) + \tau(q, q_+) > \tau(p, q) = \text{diam}(X), \quad (10.3.15)$$

contradicting the definition of the diameter.

Now assume that X is a globally hyperbolic Lorentzian length space. Recall from Lemma 10.3.2 (see also the discussion following Definition 8.1.18) that the time separation function is finite on such a space, so if the diameter is infinite it cannot be attained. Furthermore, the assumptions of the previous part still hold, hence $\text{diam}(X)$ can never be attained. \square

Recall from Corollary 7.2.8 that there is an immediate consequence of the metric version of the Bonnet–Myers theorem from Theorem 7.2.7, which states that the perimeter of any triangle in a metric space with curvature bounded below is less than or equal to $\frac{2\pi}{\sqrt{k}}$. In the metric setting, this is argued using hinge comparison, however in the synthetic Lorentzian setting the result follows directly from the reverse triangle inequality and the previous lemma. In fact we obtain a stronger result, with a strict inequality. This is of note as a rarity in the process of transferring results from metric geometry to synthetic Lorentzian geometry; as we have seen, the derivations of Lorentzian analogues to metric results are usually as difficult as, if not harder than, their metric counterparts.

Corollary 10.3.6 (Lower Curvature Bounds and Size-Bounds)

Let X be a chronological, strongly causal, geodesic Lorentzian pre-length space with curvature bounded below by $K \in \mathbb{R}$. If $K < 0$ then additionally assume that X satisfies the assumptions of Theorem 10.3.3 (or Theorem 10.3.4). Then all triangles in X satisfy size-bounds for K and have perimeter less than $2D_K$.

Proof. The case where $K \geq 0$ is relatively straightforward. We have $\text{diam}(X) \leq D_K = \infty$ which cannot be attained on X since τ is finite on X by Lemma 10.3.2. In particular, the perimeter of every triangle in X is finite.

Now let $K < 0$. By Theorem 10.3.3, we know that $\text{diam}(X) \leq D_K$, which is finite. By Lemma 10.3.5, this is not attained on X . Now consider a triangle $\Delta(x, y, z)$ in X . Since D_K is not attained on X , this triangle must satisfy size-bounds for $\mathbb{L}^2(K)$. Furthermore, we have

$$D_K > \tau(x, z) \geq \tau(x, y) + \tau(y, z), \quad (10.3.16)$$

from which it follows that $\tau(x, y) + \tau(y, z) + \tau(x, z) < 2D_K$ as required. Furthermore, since the assumptions for [Theorem 10.3.4](#) imply those for [Theorem 10.3.3](#), the result holds under these assumptions as well. \square

Throughout this part of the thesis, we have assumed that triangles satisfy the appropriate size-bounds for K , such that their comparison triangle is realisable in $\mathbb{L}^2(K)$, cf. [[119](#), Lemma 2.1]. The importance of [Corollary 10.3.6](#) is that it tells us we never needed to do so when deriving the Toponogov globalisation theorem for Lorentzian length spaces in [Theorem 10.2.4](#) or the Bonnet–Myers style theorem in its final form in [Theorem 10.3.4](#); all triangles we considered automatically satisfied the size-bounds to begin with.

10.4 Chapter Summary

In this chapter we have proven the second globalization theorem for curvature bounds on Lorentzian pre-length spaces and the counterpart of our Alexandrov’s patchwork from [Chapter 9](#) in the case of curvature bounded below. In particular, we derived the result of [Theorem 10.2.4](#) for globally hyperbolic Lorentzian length spaces with curvature bounded below, in the spirit of Toponogov’s globalisation theorem, from [Theorem 7.2.6](#).

To begin, we recalled several properties of time functions and the null distance, showing that second countability is sufficient to ensure the existence of the former in our setting. We also showed that *piecewise causal path connectedness*, as introduced by [[132](#)], is equivalent to the more standard notion of connectedness on causally path connected spaces. In [Lemma 10.1.8](#), we strengthened the bound on the null distance diameter of causal diamonds provided by *op. cit.* These results culminated in a variation of the Lebesgue number lemma for the null distance in [Lemma 10.1.9](#). This lemma states that, for any causal diamond $J(x, y)$ in X , there exists an $\varepsilon > 0$ such that any causal diamond with null distance diameter less than ε , which is contained within $J(x, y)$, is also contained within some timelike diamond. Furthermore, if X has a curvature bound, then ε can be chosen such that the timelike diamonds are comparison neighbourhoods.

We then moved on to deriving a gluing lemma for spaces with curvature bounded below, as a mirror of [Lemma 9.1.1](#) from the curvature bounded above case. While the exact statement from the curvature bounded above case does not apply here, we did succeed in showing that if the angle condition fails at a vertex of a triangle $\Delta(x, y, z)$, then the angle condition must fail in one of the two sub-triangles generated by splitting the triangle in a timelike manner through either of the other two vertices. This allowed us to perform a cat’s cradle decomposition of an arbitrary timelike triangle in a globally hyperbolic Lorentzian length space with (local) timelike curvature bounded below by $K \in \mathbb{R}$, following the approach of [[133](#)] in the metric setting. In

particular, we show that if a timelike triangle has failing angle condition, then there is always a sub-triangle (after some number of decompositions) which possesses a vertex at which the angle has sign $\sigma = +1$ and the angle condition fails. With a careful choice of decomposition, such a sub-triangle may be found in a comparison neighbourhood, yielding a contradiction. In summary, no timelike (and hence no admissible causal) triangles may have a vertex at which the angle condition fails — an analogue of Toponogov’s globalisation theorem in the synthetic Lorentzian setting. As in the curvature bounded above case, [Proposition 8.2.9](#) implies that curvature bounded below with respect to triangle comparison and hinge comparison globalise under the same assumptions, if [\(8.2.6\)](#) also holds.

In the final section of this chapter, [Section 10.3](#), we considered one of the fundamental consequences of the Toponogov globalisation theorem, namely the Bonnet–Myers theorem, first deriving a bound on the finite timelike diameter of a Lorentzian pre-length space under the assumption of global curvature bounded below by some $K < 0$ in [Theorem 10.3.3](#), before applying [Theorem 10.2.4](#) to obtain a bound on globally hyperbolic Lorentzian length spaces with local curvature bounded below by $K < 0$ in [Theorem 10.3.4](#). This result should be viewed as the extension of [[158](#), [Theorem 9.5](#)] for globally hyperbolic spacetimes. As in the metric setting, it was necessary to exclude a number of spaces satisfying some degeneracy condition, corresponding to local one-dimensionality. As a consequence of [Corollary 10.3.6](#), we obtain that all triangles in chronological, strongly causal, and geodesic Lorentzian pre-length spaces with curvature bounded below by $K \geq 0$ and all triangles in Lorentzian pre-length spaces satisfying the assumptions of either [Theorem 10.3.3](#) or [Theorem 10.3.4](#) for $K < 0$ satisfy size bounds for K .

In the conclusion to this part of the thesis, we shall discuss a number of further consequences of the globalisation theorems for Lorentzian pre-length spaces derived herein, including their application to causal set theory and Gromov–Hausdorff convergence. For now, however, we leave the reader with one small open problem: given that the one-dimensional spaces excluded from the Bonnet–Myers theorem in the metric setting are isometric to \mathbb{R} , $(0, \infty)$, $[0, B]$ for $B > \frac{\pi}{\sqrt{k}}$, and circles of radius greater than $\frac{1}{\sqrt{k}}$, what structures do the locally one-dimensional spaces correspond to in the synthetic Lorentzian setting?



11

Part II Outlook

In this part of the thesis, we have proven a pair of globalisation theorems for Lorentzian pre-length spaces with timelike curvature bound, namely an Alexandrov’s patchwork theorem for spaces with curvature bounded above, [Theorem 9.2.2](#), and a Toponogov globalisation theorem for spaces with curvature bounded below, [Theorem 10.2.4](#). Along the way, we also derived a gluing lemma for Lorentzian pre-length spaces with timelike curvature bounded below, [Lemma 10.1.13](#), and synthetic Lorentzian analogues of both the Bonnet–Myers theorem [Theorem 10.3.4](#), which bounds the diameter of a Lorentzian pre-length space with curvature bounded below, and the Lebesgue number lemma [Lemma 10.1.9](#). While the proofs of these results generally follow their counterparts from metric geometry, a not insignificant amount of difficulty was introduced by the causal structure we had to adhere to. In particular, only being able to use admissible causal triangles in our decompositions and having to account for points with infinite time separation proved to be tricky obstacles. However, in addressing these issues, we were not only able to refine the definition of curvature bounds to admit anti-de Sitter space in [Remark 8.1.27](#), but also demonstrate that in Lorentzian length spaces, global hyperbolicity implies causal closure in [Proposition 8.1.19](#), as it does for spacetimes.¹

Let us now discuss future research directions stemming from our work, beginning with how the results might be strengthened. First and foremost, the metric Toponogov globalisation theorem is proven in both [\[133\]](#) and [\[103\]](#) for complete metric length spaces. In contrast, [Theorem 10.2.4](#) considers globally hyperbolic Lorentzian length spaces; while having a Lorentzian length space is the natural analogue to a metric length space, the global hyperbolicity assump-

¹This is a seemingly straightforward result that does not appear elsewhere in the literature.

tion is not especially mild in the context of the causal ladder (see [147, 127] or the discussion following Definition 8.1.18). Since the publication of Theorem 10.2.4 in [3, Theorem 3.6], the four-point condition used to prove the metric Toponogov globalisation theorem in [103] has been adapted to the Lorentzian setting [123, Definition 4.6]. Consequently, rather than following [133] and using global hyperbolicity to perform the cat's cradle construction, we are optimistic that a Lorentzian version of the proof found in [103] might allow us to obtain a more general result. By extension, this would result in the requirements for the Bonnet–Myers theorem (Theorem 10.3.4) and splitting theorem [3, Theorem 4.5]² for Lorentzian length spaces with local curvature bounded below becoming less restrictive. In particular, the assumptions under which the Bonnet–Myers theorem holds for Lorentzian length spaces with global curvature bounds Theorem 10.3.3 are currently weaker than those required for Lorentzian length spaces with local curvature bounds Theorem 10.3.4, aside from the bounds themselves.

In a similar vein, when investigating the Lorentzification of Alexandrov's patchwork globalisation in Section 9.2, we forwent the process of globalising the continuity of τ , instead assuming that τ was globally continuous a priori and referring the reader to the preliminary result [2, Remark 4.9], where it is shown that τ is continuous on $\tau^{-1}((0, D_K))$ under the assumption of curvature (locally) bounded above by K . Extending this result to prove that τ is globally continuous would allow us to remove the corresponding assumption from the statement of Theorem 9.2.2. Equally, while it is shown in [2, Proposition 4.8] that timelike distance realisers of length less than D_K vary continuously with their endpoints, when we additionally assume that the Lorentzian pre-length space is locally causally closed and non-totally imprisoning, the corresponding result from metric geometry, Proposition 7.2.4, does not require such additional assumptions. Since the constraints of local causal closure and non-total imprisonment are involved in the application of the limit curve theorem [116, Theorem 3.7], any further investigation in this direction may also provide a better understanding of said limit curves. Both of these improvements would bring Theorem 9.2.2 more in line with the metric version of Alexandrov's patchwork globalisation Theorem 7.2.2.

Moving further afield now, let us consider the problem of Gromov–Hausdorff convergence of Lorentzian manifolds and pre-length spaces. It is a well known result of Riemannian geometry that the limits (in the sense of Gromov–Hausdorff) of Riemannian manifolds with sectional curvature bounded below by k are metric length spaces with the same lower curvature bound in the sense of triangle comparison [107], with the Toponogov globalisation theorem [103] guaranteeing that the size of the triangle comparison neighbourhoods in the manifolds does not collapse to zero in the limit. More generally, it is known that lower curvature bounds on metric spaces

²See [128, Theorem 1.4] for the statement of the theorem under global curvature bounds.

are stable under Gromov–Hausdorff convergence, though not all metric spaces with curvature bounded below can be obtained from such a limit [106]. These statements have been used to prove a wide range of finiteness theorems [163–165] which have even been generalised to the setting of Riemannian orbifolds [166].

Typically, Gromov–Hausdorff convergence is first defined for compact spaces, then extended to the non-compact case in the form of *pointed Gromov–Hausdorff convergence*, by picking a distinguished point in each space and considering the Gromov–Hausdorff convergence of closed metric balls of radius R , around said points, for each $R > 0$ [167]. However, in Lorentzian geometry, most spaces of interest are non-compact; indeed, a compact Lorentzian manifold must contain a closed timelike curve [122, Lemma 14.10], violating causality [116, Lemma 3.3]. Consequently, it is difficult to establish a suitably general notion of Gromov–Hausdorff convergence in the Lorentzian setting and the convergence of sequences of Lorentzian pre-length spaces is still an open problem, in spite of the various definitions proposed [135, 159, 168]. Regardless, by reflecting on the metric case, we should expect timelike lower curvature bounds to be stable under any reasonable notion of Gromov–Hausdorff convergence. In [132, Proposition 4.17], Kunzinger and Steinbauer show that this is true for the standard notion of pointed Gromov–Hausdorff convergence on Lorentzian length spaces given by warped products, while in [134, Theorem 6.7], Minguzzi and Suhr demonstrate that curvature bounds on so-called “*bounded Lorentzian length spaces*” remain stable under a specially tailored notion of Gromov–Hausdorff convergence [134, Theorem 4.34]. This general principle has been described for globally hyperbolic Lorentzian length spaces by the author and collaborators in [3, Section 4.1], using the notion of convergence introduced by Minguzzi and Suhr as an example; it is this example that we summarise below.

A *bounded Lorentzian metric space* is a set X equipped with a function $\tilde{\tau} : X \times X \rightarrow [0, \infty)$ which satisfies the reverse triangle inequality, is continuous in a topology where $\{(p, q) \mid \tilde{\tau}(p, q) \geq \varepsilon\}$ is compact for all $\varepsilon > 0$, and distinguishes points,³ see [134, Definition 1.1]. If, in addition, timelike related points are connected by distance realisers, then we call X a *bounded Lorentzian length space*. Recalling that causal diamonds in globally hyperbolic Lorentzian pre-length spaces are compact, Harvey proposes the following two results [3, Lemma 4.1, Theorem 4.2], which bring together the Toponogov globalisation theorem for Lorentzian length spaces with the convergence of bounded Lorentzian length spaces in the sense of Minguzzi–Suhr:

Lemma 11.0.1 (Causal Diamonds Are (Almost) Bounded Lorentzian Length Spaces)

Let X be a globally hyperbolic, regular Lorentzian length space and let $J(p, q)$ be a causal diamond in X . Let S be the spacelike boundary of $J(p, q)$ – the set of points in $J(p, q)$ which are not timelike

³Recall that $\tilde{\tau}$ is said to distinguish $x, y \in X$ if $\tilde{\tau}(x, y) = 0$ implies $x = y$.

related to any other point. Then $J(p, q) \setminus S$ is a bounded Lorentzian length space.

Theorem 11.0.2 (Stability of Lower Curvature Bounds)

Let X_i be a sequence of connected, globally hyperbolic, and regular Lorentzian length spaces with time functions and timelike curvature bounded below by $K \in \mathbb{R}$ in the sense of Definition 8.1.20. Let $J_i = J(p_i, q_i)$ be a sequence of causal diamonds in X_i with spacelike boundary S_i . If the sequence $J_i \setminus S_i$ converges in the sense of Minguzzi–Suhr [134, Theorem 4.34] to some J , then J is a bounded Lorentzian length space with sectional curvature bounded below by K in the sense of Minguzzi–Suhr [134, Definition 6.9].

Another open question in the same spirit, initially posed by [116, Section 5.3], asks whether Lorentzian length spaces (and by extension spacetimes) can be discretised using “causal sets”. *Causal sets* are sets equipped with locally finite partial order and are used to model discrete spacetimes in the theory of quantum gravity [169, 170]. They may be illustrated by locally finite, transitively reduced, acyclic graphs called *Hasse diagrams*, where the additional qualifiers reflect the implicit Lorentzian structure. In that sense, causal sets may be considered the Lorentzian counterpart to graphs and one might expect that they can be given the structure of a Lorentzian pre-length space, similarly to how a graph can be given the structure of a metric space. However, unlike when we consider metric graphs (see [109, Section 3.2.2]), a causal set should not be thought of as including the edges of its corresponding Hasse diagram; the edges merely represent the partial ordering on the set. Indeed, any vertex of valency > 2 in a graph produces a branching curve given by its adjacent edges, so Hasse diagrams which are not given by a path graph violate lower semi-continuity of the time separation function, cf. [116, Lemma 2.12], and cannot be given the structure of a Lorentzian pre-length space. Instead, [116, Section 5.3] show that a causal set may be given the structure of a Lorentzian pre-length space (though not a Lorentzian length space), when interpreted as an “ordered point-cloud.”

In metric geometry, it is known that every compact length space is the Gromov–Hausdorff limit of a sequence of finite graphs [109, Proposition 7.5.5], amounting to the discretisation of continuous space. When we consider the limit of a sequence of causal sets, we fall victim to the same issue as when we considered the limit of Lorentzian manifolds — the spaces we are interested are, in general, not compact. Some progress may be made, however, as Minguzzi and Suhr show that causal sets are precisely finite bounded Lorentzian metric spaces [134, Proposition 2.3] and every (not necessarily finite) bounded metric space is the limit of causal sets [134, Corollary 4.32]. We may therefore obtain the following result:

Proposition 11.0.3 (Local Approximation by Causal Sets)

Globally hyperbolic, regular Lorentzian length spaces are locally given by the limit of causal sets,

in the sense of Minguzzi–Suhr [134, Theorem 4.34].

Proof. Let X be a globally hyperbolic, regular Lorentzian length space and consider an arbitrary point $x \in X$. Since globally hyperbolic Lorentzian length spaces are strongly causal and non-timelike locally isolating, Proposition 8.1.24 implies that there exists points $p, q \in X$ such that $x \in I(p, q) \subseteq J(p, q)$ where $I(p, q)$ and $J(p, q)$ are timelike and causal diamonds respectively. In particular, $x \in I(p, q) \subseteq J(p, q) \setminus S$ for S the spacelike boundary of $J(p, q)$. As $J(p, q) \setminus S$ is a bounded Lorentzian length space by Lemma 11.0.1, it follows from [134, Corollary 4.32] that $J(p, q) \setminus S$ is given by the limit of causal sets. Consequently, for each $x \in X$, there exists an open set $I(p, q)$ containing x , which is itself contained in a set $J(p, q) \setminus S$ given by the limit of causal sets. \square

It is unclear from the above result whether or not globally hyperbolic, regular Lorentzian length spaces are globally given by the limit of causal sets, since gluing together the local approximations at a given step may not yield a causal set; we need to ensure that the result of this gluing remains locally finite and acyclic.⁴

The converse problem of whether all causal sets have continuum approximation given by spacetimes is known to be false, see the excellent review paper [171]. However, it is plausible that, in general, causal sets are instead well-approximated by sequences of (non-finite, non-discrete) Lorentzian pre-length spaces. Resolving both of these issues would strengthen the correspondence between Lorentzian pre-length spaces and causal sets, enabling results to be shared between the two perspectives.

⁴Recall that the partial order is both antisymmetric and transitive.

Part III

Appendices



A

Locally-a-section Lagrangian Submanifolds

In this appendix, we shall fulfil our promise from [Section 2.1.3](#) and provide a proof of the statement of Banos [35] given in [Proposition 2.1.14](#). As in earlier chapters, T^*M will denote the cotangent bundle of some manifold M . Let us begin by restating the relevant definitions here for convenience:

Definition 2.1.5 (Local Diffeomorphism)

A function $h : L \rightarrow M$, between manifolds L and M , is called a local diffeomorphism if, for all $y \in L$, there exists some open neighbourhood $V_y \subseteq L$ of y such that $h(V)$ is open in M and the restriction $h|_{V_y}$ of h to V_y is a diffeomorphism onto its image.

Definition 2.1.13 (Locally-a-Section Submanifold)

A submanifold $\iota : L \hookrightarrow T^*M$ is called locally-a-section if, for all $y \in L$, there exists a neighbourhood $V_y \subseteq L$ of y , an open set $U_y \subseteq M$, and a function $\psi_y \in \mathcal{C}^\infty(U_y)$ such that $\iota(V_y) = d\psi_y(U_y)$.

Now let M be an m -dimensional, connected, Riemannian manifold with cotangent bundle denoted by T^*M . Denote the canonical projection by $\pi : T^*M \rightarrow M$. Let x^i and q_i , with $i = 1, \dots, m$, denote local coordinates on M and the fibres of T^*M over M , respectively. Consider Lagrangian submanifolds $\iota : L \hookrightarrow T^*M$ with respect to the canonical symplectic form $\omega := dq_i \wedge dx^i$. We then have the following result:

Proposition 2.1.14 (Locally-a-Section iff Local Diffeomorphism)

Let (M, g) be a Riemannian manifold and let its cotangent bundle be equipped with the canonical symplectic form ω . Let $\pi : T^*M \rightarrow M$ be the canonical projection on the cotangent bundle T^*M .

A Lagrangian submanifold $\iota : L \hookrightarrow T^*M$ with respect to ω is locally-a-section if and only if the projection $\pi|_L := \pi \circ \iota : L \rightarrow M$ is a local diffeomorphism.

Proof. Let $\pi|_L : L \rightarrow M$ be a local diffeomorphism and let y be an arbitrary point in L . We want to show that L is locally-a-section. As $\pi|_L$ is a local diffeomorphism, there exists some open neighbourhood $V \subseteq L$ of y such that $\pi|_L(V) =: U$ is open in M and $\pi|_V$ is a diffeomorphism onto its image. In particular, U can be taken to be contractible.¹ Let y^i be local coordinates on V . Then, $V \ni y^i \xrightarrow{\pi|_L} x^i(y) \in U$. Again by the local diffeomorphism properties of $\pi|_L$, we have locally the invertibility of the Jacobian $\frac{\partial x^i}{\partial y^j}$ and the inverse relation $y^i = y^i(x)$. Hence, the embedding $\iota : L \hookrightarrow T^*M$ becomes $i : y^i \mapsto (x^i(y), q_i(y)) = (x^i, q_i(y(x))) =: (x^i, p_i(x))$ in local coordinates. Furthermore, since L is Lagrangian with respect to $\omega = dq_i \wedge dx^i$, we find

$$\frac{\partial x^i}{\partial y^j} \frac{\partial q_i}{\partial y^k} = \frac{\partial x^i}{\partial y^k} \frac{\partial q_i}{\partial y^j} \quad (\text{A.0.1})$$

upon computing $\iota^*\omega = 0$.² Hence,

$$\begin{aligned} \frac{\partial x^i}{\partial y^l} \left(\frac{\partial p_j}{\partial x^i} - \frac{\partial p_i}{\partial x^j} \right) &= \frac{\partial x^i}{\partial y^l} \left(\frac{\partial y^k}{\partial x^i} \frac{\partial q_j}{\partial y^k} - \frac{\partial y^k}{\partial x^j} \frac{\partial q_i}{\partial y^k} \right) \\ &= \frac{\partial q_j}{\partial y^l} - \frac{\partial y^k}{\partial x^j} \frac{\partial x^i}{\partial y^l} \frac{\partial q_i}{\partial y^k} \\ &= \frac{\partial q_j}{\partial y^l} - \frac{\partial y^k}{\partial x^j} \frac{\partial x^i}{\partial y^k} \frac{\partial q_i}{\partial y^l} \\ &= \frac{\partial q_j}{\partial y^l} - \frac{\partial q_j}{\partial y^l} \\ &= 0. \end{aligned} \quad (\text{A.0.2})$$

Therefore,

$$\frac{\partial p_j}{\partial x^i} - \frac{\partial p_i}{\partial x^j} = 0 \quad (\text{A.0.3})$$

that is, the one-form $\eta := p_i dx^i$ on $U \subseteq M$ is closed. Consequently, by the Poincaré lemma, there is a function $\psi \in \mathcal{C}^\infty(U)$ so that $\eta = d\psi$ (and therefore $p_i = \partial_i \psi$). It follows that $d\psi(U) = \iota(V)$. In summary, L is locally-a-section.

Conversely, let L be locally-a-section, such that, for all $y \in L$, there exists $V \subseteq L$ open, $U \subseteq M$ open, and $\psi \in \mathcal{C}^\infty(U)$ such that $d\psi(U) = \iota(V)$. We wish to show that $\pi|_L(V)$ is open in M and $\pi|_V := \pi|_L|_V = \pi|_{\iota(V)} \circ \iota$ is a diffeomorphism onto its image. As $\iota : L \hookrightarrow T^*M$ is a smoothly embedded submanifold, ι is a topological diffeomorphism, hence V and $\iota(V)$ are

¹If U is not contractible, take an open ball $B \subset M$ in the metric topology induced by g , such that $\pi|_L(y) \in B \subset U$. Since metric balls are convex, B is contractible. Then take $\tilde{U} = B$, $\tilde{V} = \iota^{-1}(B) \subset V$.

²Note that $\iota_* \left(\frac{\partial}{\partial y^i} \right) = \frac{\partial x^j}{\partial y^i} \frac{\partial}{\partial x^j} + \frac{\partial q_j}{\partial y^i} \frac{\partial}{\partial q_j}$.

diffeomorphic, with $\iota(V)$ carrying the subspace topology in T^*M . Furthermore, for $x \in U$, we have

$$\pi \circ d\psi(x) = \pi((x^i, \partial_i \psi(x))) = x, \quad (\text{A.0.4})$$

so $\pi \circ d\psi$ is the identity map on U . Similarly, for $(x^i, \partial_i \psi(x)) \in d\psi(U) \subset T^*M$, we find

$$d\psi \circ \pi((x^i, \partial_i \psi(x))) = d\psi(x^i) = (x^i, \partial_i \psi(x)), \quad (\text{A.0.5})$$

so $d\psi \circ \pi$ is the identity map on $d\psi(U)$. Consequently, $\pi|_{d\psi(U)}$ is the inverse of $d\psi : U \rightarrow d\psi(U)$ and U is diffeomorphic to $d\psi(U)$, with $d\psi(U)$ carrying the subspace topology in T^*M . Combining this with the locally-a-section property $d\psi(U) = \iota(V)$, we find that the composition $\pi|_V = \pi|_{\iota(V)} \circ \iota = \pi|_{d\psi(U)} \circ \iota$ is a diffeomorphism onto its image, with inverse $\iota^{-1} \circ d\psi : \pi|_V(V) \rightarrow V$. Since $\pi|_L(V) = \pi|_V(V) = U$ is known to be open in M , it follows that $\pi|_L$ is a local diffeomorphism from L to M . \square

Consequently, for locally-a-section submanifolds and suitably nice sets $V \subseteq L$, $U = \pi|_L(V) \subseteq M$, we can explicitly write the inverse $\iota^{-1} \circ d\psi : U \rightarrow V$ of the map $\pi|_V$. By the local diffeomorphism property of $\pi|_L$, we may therefore use the same coordinates on these corresponding pairs of nice sets; in these coordinates, $\pi|_V$ becomes the identity. Conversely, if a Lagrangian submanifold L has an open neighbourhood around each point on which one can take coordinates from M , it follows that the submanifold is locally-a-section.

For submanifolds which are given by global sections, the above proof implies that $\pi|_L$ is a diffeomorphism and we may use the coordinates from M on L globally. However, the converse does not hold; given that $\pi|_L$ is a diffeomorphism, we can only infer that L is locally-a-section, since our proof uses local arguments, such as the Poincaré lemma.

Observe, as we did in Section 2.1.3, that while a local diffeomorphism must be an immersion, it need not be injective or surjective globally. Consequently, given a generalised solution $\iota : L \hookrightarrow T^*M$ to a Monge–Ampère structure which is also locally-a-section, the projection $\pi|_L$ may be non-injective and $\iota(L)$ may intersect several times with any given fibre of T^*M over M , resulting in multiple “branches” of $\iota(L)$ (see the interior of the red region in Figure 2.1.1b). Each of the branches can be interpreted as a separate classical solution, hence generalised solutions which are locally-a-section with non-injective projection $\pi|_L$ correspond to multivalued functions ψ which solve the Monge–Ampère equation. Similarly, if the projection $\pi|_L$ is non-surjective, then the submanifold L may not intersect with some of the fibres of T^*M over M at all, corresponding to solutions of the Monge–Ampère equation whose derivatives are not defined on the whole domain M (see the grey dashed line in Figure 2.1.1b). This includes functions ψ which are themselves only defined on a subset of M .



\mathcal{B}

Quaternionic Structures and Integrability in Two Dimensions

B.1 Quaternionic Structures and Lychagin–Rubtsov Metrics

As promised in [Remark 2.2.8](#), we now demonstrate how choosing a Lychagin–Rubtsov metric in two dimensions corresponds, up to conformal scaling of the almost (para-)Hermitian form, to a choice of almost (pseudo-)quaternionic structure on T^*M .

Let M be a two-dimensional Riemannian manifold with cotangent bundle T^*M and (ϖ, α) be a Monge–Ampère structure on T^*M . Assume that α is non-degenerate. Without loss of generality, we can then scale α so that its Pfaffian with respect to ϖ has modulus 1, i.e. $\text{Pf}_{\varpi}(\alpha) = \pm 1$ pointwise. We can do so because the definition of the almost (para-)complex structure (2.2.13) would otherwise perform this scaling for us. Let K denote one of the (1,1) forms implied by [Proposition 2.2.7](#). Again, scale K such that its Pfaffian with respect to ϖ has modulus 1, i.e. $\text{Pf}_{\varpi}(K) = \pm 1$. This second rescaling corresponds to our freedom to choose a conformal factor for our Lychagin–Rubtsov metric. In summary, we have

$$\text{Pf}_{\varpi}(K)K \wedge K = \varpi \wedge \varpi = \text{Pf}_{\varpi}(\alpha)\alpha \wedge \alpha, \quad (\text{B.1.1})$$

$$\alpha \wedge \varpi = 0, \quad K \wedge \varpi = 0, \quad \text{and} \quad K \wedge \alpha = 0, \quad (\text{B.1.2})$$

where we can move the Pfaffians off of $\varpi \wedge \varpi$ because they have modulus 1.

The formula (2.2.13) defining the almost (para-)complex structure associated with a Monge–Ampère structure holds for any two pairwise effective, non-degenerate forms, without the additional requirement that one of them is closed (read: symplectic). Consequently, we may define

three almost (para-)complex structures:

$$\alpha =: J \lrcorner \varpi, \quad K =: R \lrcorner \alpha, \quad \text{and} \quad \varpi =: S \lrcorner K, \quad (\text{B.1.3})$$

which satisfy

$$J^2 = -\text{Pf}_\varpi(\alpha)\mathbb{I}, \quad R^2 = -\text{Pf}_\varpi(\alpha)\text{Pf}_\varpi(K)\mathbb{I}, \quad \text{and} \quad S^2 = -\text{Pf}_\varpi(K)\mathbb{I}. \quad (\text{B.1.4})$$

Using the matrix representations of (B.1.3) and (B.1.4), it can be verified that the compositions of J , R , and S obey the following Cayley table:

	\mathbb{I}	J	R	S
\mathbb{I}	\mathbb{I}	J	R	S
J	J	$-\text{Pf}_\varpi(\alpha)\mathbb{I}$	$-\text{Pf}_\varpi(K)S$	$\text{Pf}_\varpi(\alpha)\text{Pf}_\varpi(K)R$
R	R	$\text{Pf}_\varpi(K)S$	$-\text{Pf}_\varpi(\alpha)\text{Pf}_\varpi(K)\mathbb{I}$	$-\text{Pf}_\varpi(\alpha)J$
S	S	$-\text{Pf}_\varpi(\alpha)\text{Pf}_\varpi(K)R$	$\text{Pf}_\varpi(\alpha)J$	$-\text{Pf}_\varpi(K)\mathbb{I}$

Hence, when $\text{Pf}_\varpi(\alpha)$ and $\text{Pf}_\varpi(K)$ both equal 1, it follows that (J, R, S) is a triple of almost complex structures satisfying the quaternionic relations and we call them an *almost quaternionic structure* on T^*M . However, when one or both of $\text{Pf}_\varpi(\alpha)$ and $\text{Pf}_\varpi(K)$ is equal to -1 , two of (J, R, S) become almost para-complex and the triple satisfies the pseudo-quaternionic relations, in which case we call (J, R, S) an *almost pseudo-quaternionic structure* on T^*M (see [16, Section 4] where a similar construction is performed in the context of the semi-geostrophic equations).

By construction, K is an almost (para-)Hermitian form for J , ϖ is an almost (para-)Hermitian form for R , and α is an almost (para-)Hermitian form for S . Furthermore, again using the matrix representations of (B.1.3) and (B.1.4), it is straightforward to check that

$$g(X, Y) = K(X, JY) = \text{Pf}_\varpi(\alpha) \varpi(X, RY) = \text{Pf}_\varpi(K) \alpha(X, SY), \quad (\text{B.1.5})$$

hence g is almost (para-)Hermitian with respect to all three of our almost (para-)complex structures simultaneously, that is, g is almost (pseudo-)quaternionic Hermitian. Most importantly, making a different choice of K while keeping (ϖ, α) fixed leads to different (para-)complex structures R and S which satisfy the construction above. Consequently, a choice of Lychagin–Rubtsov metric is, up to a conformal factor, a choice of almost (pseudo-)quaternionic structure on T^*M , with the type depending on $\text{Pf}_\varpi(\alpha)$ and $\text{Pf}_\varpi(K)$.

Furthermore, ϖ is symplectic by assumption and therefore closed, which implies ϖ is (up to a sign) an almost (para-)Kähler form for g . If α is closed, then α also becomes an almost (para-

)Kähler form for g , the structure J is integrable by [Theorem 2.2.6](#), and K is a (para-)Hermitian form for g . Similarly, if K is closed, then K becomes an almost (para-)Kähler form for g , the structure S is integrable, and α is a Hermitian form for g . Finally, if both K and α are closed, then (J, R, S) are all integrable and (ϖ, α, K) are all, up to a sign, (para-)Kähler for g , in which case, g is a (pseudo-)quaternionic Kähler metric.

Consider now our fluid dynamical structures $(\varpi, \alpha, K = \omega)$ from [\(4.1.1\)](#), [\(4.1.4\)](#), and [\(4.1.8\)](#), with scaling such that their Pfaffians have modulus 1. In particular, $\text{Pf}_{\varpi}(\alpha) = \text{sgn}(\hat{f})$ and $\text{Pf}_{\varpi}(\omega) = 1$. The corresponding triple of almost (para-)complex structures $(J, R, S) = (\hat{\mathcal{J}}, \hat{\mathcal{J}}^{-1} = \text{sgn}(\hat{f})\hat{\mathcal{J}}, S)$ are therefore quaternionic when $\hat{f} > 0$ and pseudo-quaternionic when $\hat{f} < 0$, with g the corresponding almost (pseudo-) quaternionic Hermitian metric. Furthermore, (J, R, S) are simultaneously integrable if and only if f is constant, in which case g is a (pseudo-)quaternionic Kähler metric. As a result of our discussion above, we obtain the same Lychagin–Rubtsov metric in [Section 3.2](#) as in [Section 4.1](#) since, aside from their scaling, the triple of forms used to define the Lychagin–Rubtsov metric is unchanged and merely reordered, resulting in the same triple of almost (para-)complex structures. In contrast to the hyper-(para-)Kähler structures constructed from the semi-geostrophic equations in [\[172\]](#) and the group structure of the analogous tensors presented in [\[19\]](#) for incompressible Navier–Stokes flows, the group structure of our almost (para-)complex structures exhibits a type change from quaternionic to pseudo-quaternionic when \hat{f} goes from positive to negative. This suggests that our choice of Lychagin–Rubtsov metric is more natural for studying incompressible Navier–Stokes flows, as all associated geometric quantities are sensitive to changes in the dominance of vorticity and strain.

B.2 The Poisson Equation for Pressure and Integrability

Let (M, \hat{g}) be a two-dimensional Riemannian manifold with coordinates $\{x^i\}_{i=1}^2$ and equip its cotangent bundle with coordinates $\{x^i, q_i\}_{i=1}^2$. Let (ϖ, α) be the Monge–Ampère structure given by [\(4.1.1\)](#) and [\(4.1.4\)](#). We wish to show that the almost (para-)complex structure $\hat{\mathcal{J}}$ in [\(4.1.6\)](#) is integrable if and only if \hat{f} is constant. Note that the following computation implies the same is true for the Monge–Ampère structure [\(3.2.4b\)](#) and the almost (para-) complex structure $\hat{\mathcal{J}}$ from [\(3.2.6\)](#), since $\text{Pf}_{\omega}(\alpha) = \text{Pf}_{\varpi}(\alpha)$.

Recall that the Lychagin–Rubtsov theorem (see [Theorem 2.2.6](#)) tells us that, given a Monge–Ampère structure (ϖ, α) for which $\text{Pf}_{\varpi}(\alpha) \neq 0$, the corresponding almost (para-)complex structure

$$\frac{\alpha}{\sqrt{|\text{Pf}_{\varpi}(\alpha)|}} =: J_{\alpha} \lrcorner \varpi, \tag{B.2.1}$$

is integrable if and only if

$$d\left(\frac{1}{\sqrt{|\text{Pf}_\varpi(\alpha)|}}\alpha\right) = 0. \quad (\text{B.2.2})$$

For our Monge–Ampère structure, we have $d\alpha = 0$ and $\text{Pf}_\varpi(\alpha) = \hat{f}$, which we assume to be non-zero, hence

$$d\left(\frac{\alpha}{\sqrt{|\text{Pf}_\varpi(\alpha)|}}\right) = -\frac{1}{2\hat{f}|\hat{f}|^{\frac{1}{2}}}d\hat{f} \wedge \alpha. \quad (\text{B.2.3})$$

so in order to find out when $\hat{\mathcal{J}}$ and \hat{J} are integrable, it remains to check when $d\hat{f} \wedge \alpha = 0$.

From (3.2.4a), we have a basis of 1-forms on T^*M given by

$$dx^i \quad \text{and} \quad \overset{\circ}{\nabla}q_i := dq_i - dx^j \overset{\circ}{\Gamma}_{ji}^k q_k. \quad (\text{B.2.4})$$

Denote the basis of $\Gamma(TT^*M)$ dual to (B.2.4) by E_i and \tilde{E}^i respectively. Then we may write

$$d\hat{f} = (\tilde{E}^i \hat{f}) \overset{\circ}{\nabla}q_i + (E_i \hat{f}) dx^i, \quad (\text{B.2.5})$$

from which it follows that

$$d\hat{f} \wedge \alpha = \frac{\sqrt{\det(\overset{\circ}{g})}}{2} \left[(E_i \hat{f}) \varepsilon^{jk} dx^i \wedge \overset{\circ}{\nabla}q_j \wedge \overset{\circ}{\nabla}q_k - (\tilde{E}^i \hat{f}) \hat{f} \varepsilon_{jk} \overset{\circ}{\nabla}q_i \wedge dx^j \wedge dx^k \right]. \quad (\text{B.2.6})$$

Since $\hat{f} \neq 0$, this implies $d\hat{f} \wedge \alpha = 0$ if and only if $E_i \hat{f} = 0$ and $\tilde{E}^i \hat{f} = 0$ for $i = 1, 2$. That is, $d\hat{f} = 0$ as required.

Finally, note that when $\hat{f} := \frac{1}{2}\overset{\circ}{\Delta}_{\text{BP}} - \frac{1}{4}\overset{\circ}{R}|q|^2$ is a non-zero constant, $\partial_{q_i} \hat{f} = 0$ implies $\partial_{q_i}(\frac{1}{4}\overset{\circ}{R}|q|^2) = 0$, since the first term in \hat{f} is independent of q . However, as $\partial_{q_i}|q|^2 \neq 0$, this implies $\overset{\circ}{R} = 0$. That is, if \hat{f} is constant for some two-dimensional, incompressible fluid flow, then the background of the flow M is Ricci-flat. By contraposition, if M is not Ricci-flat, then \hat{f} is not constant, which in turn implies (by the above calculation) that \hat{J} and $\hat{\mathcal{J}}$ are not integrable, that is, they are never (para-)complex on non-flat M .



Connections and Curvature

C.1 Pull-back Metric in Two Dimensions

In what follows, we shall provide some more details on the computation of the Levi-Civita connection and curvature scalar associated with the metric (3.2.10) from Section 3.2. Firstly, recall that using (3.2.9), the metric (3.2.10) can be written in the form

$$g_{ij} = \zeta \tilde{g}_{ij} \quad \text{with} \quad \tilde{g}_{ij} = \psi_{ij} , \quad (\text{C.1.1})$$

where the indices on $\psi \in \mathcal{C}^\infty(M)$ are interpreted via (3.2.15). When $\zeta \neq 0$, g is a conformal scaling of the metric $\text{sgn}(\zeta)\tilde{g}$ with conformal factor $|\zeta|$, where \tilde{g} is the Hessian metric with respect to ψ . We wish to exploit this conformal nature to write the Levi-Civita connection and Ricci curvature scalar of g in terms of those of \tilde{g} .

C.1.1 Levi-Civita Connection of the Pull-back Metric

We begin by observing that a consequence of the Ricci identity

$$[\overset{\circ}{\nabla}_i, \overset{\circ}{\nabla}_j]\eta_k = -\overset{\circ}{R}_{ijk}{}^l \eta_l , \quad (\text{C.1.2})$$

for one-forms $\eta = \eta_i dx^i$, is the following expression for the triple derivative of ψ in terms of the totally symmetrised triple derivative and curvature terms:

$$\overset{\circ}{\nabla}_i \psi_{jk} = \psi_{ijk} + \frac{1}{3}([\overset{\circ}{\nabla}_i, \overset{\circ}{\nabla}_j]\psi_k + [\overset{\circ}{\nabla}_i, \overset{\circ}{\nabla}_k]\psi_j) = \psi_{ijk} - \frac{2}{3}\overset{\circ}{R}_{i(jk)}{}^l \psi_l . \quad (\text{C.1.3})$$

Upon applying the first Bianchi identity, $\mathring{R}_{ijk}{}^l + \mathring{R}_{jki}{}^l + \mathring{R}_{kij}{}^l = 0$, we find

$$\mathring{\nabla}_i \psi_{jk} + \mathring{\nabla}_j \psi_{ik} - \mathring{\nabla}_k \psi_{ij} = \psi_{ijk} + \frac{4}{3} \mathring{R}_{k(ij)}{}^l \psi_l. \quad (\text{C.1.4})$$

Consequently, the Christoffel symbols for \tilde{g} are given by

$$\begin{aligned} \tilde{\Gamma}_{ij}{}^k &= \frac{1}{2} \tilde{g}^{kl} (\partial_i \tilde{g}_{jl} + \partial_j \tilde{g}_{il} - \partial_l \tilde{g}_{ij}) \\ &= \mathring{\Gamma}_{ij}{}^k + \frac{1}{2} \tilde{g}^{kl} (\mathring{\nabla}_i \psi_{jl} + \mathring{\nabla}_j \psi_{il} - \mathring{\nabla}_l \psi_{ij}) \\ &= \mathring{\Gamma}_{ij}{}^k + \frac{1}{2} \tilde{g}^{lk} \left(\psi_{ijl} + \frac{4}{3} \mathring{R}_{l(ij)}{}^r \psi_r \right), \end{aligned} \quad (\text{C.1.5a})$$

where we have used (C.1.4) in the second step. Introducing the notation

$$\Upsilon_{ijk} := \psi_{ijk} + \frac{4}{3} \mathring{R}_{k(ij)}{}^l \psi_l. \quad (\text{C.1.5b})$$

yields (3.2.17a). In general, when a metric is changed by an overall sign, its Christoffel symbols are unchanged, hence (C.1.5a) are also the Christoffel symbols for $\text{sgn}(\zeta)\tilde{g}$, when $\zeta \neq 0$. The following proposition, see e.g. [173, Theorem 1.159], therefore yields the Christoffel symbols (3.2.17a), upon setting $\phi = \frac{1}{2} \log(|\zeta|)$ and $g' = \text{sgn}(\zeta)\tilde{g}$.

Proposition C.1.1 (Conformal Scaling of the Levi-Civita Connection)

Let M be a smooth (pseudo-)Riemannian manifold with metric $g = e^{2\phi}g'$, where $\phi \in \mathcal{C}^\infty(M)$ and g' is another metric on M , to which g is conformal. Let $\Gamma'_{ij}{}^k$ denote the Christoffel symbols of the second kind associated with g' . Then the Christoffel symbols of the second kind associated with g are given in terms of g' as

$$\Gamma_{ij}{}^k = \Gamma'_{ij}{}^k + (\partial_i \phi) \delta_j{}^k + (\partial_j \phi) \delta_i{}^k - (\partial_l \phi) (g')^{lk} g'_{ij}, \quad (\text{C.1.6})$$

where $(g')^{ij}$ denotes the inverse of the metric g'_{ij} .

C.1.2 Ricci Curvature Scalar of the Pull-back Metric

Let us now compute the curvature scalar for (3.2.10). Firstly, we note that

$$\begin{aligned} \tilde{R}_{ijk}{}^l &= \partial_i \tilde{\Gamma}_{jk}{}^l - \partial_j \tilde{\Gamma}_{ik}{}^l - \tilde{\Gamma}_{ik}{}^m \tilde{\Gamma}_{jm}{}^l + \tilde{\Gamma}_{jk}{}^m \tilde{\Gamma}_{im}{}^l \\ &= \mathring{R}_{ijk}{}^l + \frac{1}{2} (\mathring{\nabla}_i \Upsilon_{jk}{}^l - \mathring{\nabla}_j \Upsilon_{ik}{}^l - \frac{1}{2} \Upsilon_{ik}{}^m \Upsilon_{jm}{}^l + \frac{1}{2} \Upsilon_{jk}{}^m \Upsilon_{im}{}^l), \end{aligned} \quad (\text{C.1.7})$$

where we have used (C.1.5a) and set $\Upsilon_{ij}{}^k := \Upsilon_{ijl}{}^k g'^{lk}$. Next, using (C.1.3), it is straightforward to show that

$$\mathring{\nabla}_i \tilde{g}^{jk} = -\tilde{g}^{jl} \tilde{g}^{km} (\psi_{ilm} - \frac{2}{3} \mathring{R}_{i(lm)}{}^n \psi_n). \quad (\text{C.1.8})$$

and

$$\overset{\circ}{\nabla}_i \psi_{jkl} = \psi_{ijkl} - 2\overset{\circ}{R}_{i(jk}{}^m \psi_{l)m} - \frac{1}{2}\psi_m \overset{\circ}{\nabla}_{(j} \overset{\circ}{R}_{|i|kl)}{}^m. \quad (\text{C.1.9})$$

Using these two relations, we find that

$$\begin{aligned} \overset{\circ}{\nabla}_i \Upsilon_{jk}{}^l &= -\tilde{g}^{lr} (\psi_{irs} - \frac{2}{3}\overset{\circ}{R}_{i(rs)}{}^n \psi_n) \Upsilon_{jk}{}^s \\ &\quad + \frac{4}{3}\tilde{g}^{lm} (\overset{\circ}{R}_{m(jk)}{}^n \psi_{in} + \psi_n \overset{\circ}{\nabla}_i \overset{\circ}{R}_{m(jk)}{}^n) \\ &\quad + \tilde{g}^{lm} (\psi_{ijkm} - \frac{1}{2}\psi_n \overset{\circ}{\nabla}_{(j} \overset{\circ}{R}_{|i|km)}{}^n - 2\overset{\circ}{R}_{i(jk}{}^n \psi_{m)n}). \end{aligned} \quad (\text{C.1.10})$$

Upon substituting this expression and (C.1.5b) into (C.1.7), the curvature scalar (3.2.18b) then follows directly upon taking the traces $\tilde{R} = \tilde{g}^{ij} \tilde{R}_{kij}{}^k$.

Setting $\phi = \frac{1}{2} \log(|\zeta|)$, the following proposition from e.g. [173, Theorem 1.159] yields the Ricci curvature scalar of the pull-back metric g in terms of the Ricci curvature scalar of $g' = \text{sgn}(\zeta)\tilde{g}$ in two dimensions.

Proposition C.1.2 (Conformal Scaling of the Ricci Curvature Scalar)

Let M be a smooth, m -dimensional, (pseudo-)Riemannian manifold with metric $g = e^{2\phi}g'$, where $\phi \in \mathcal{C}^\infty(M)$ and g' is another metric on M , to which g is conformal. Let R' denote the Ricci curvature scalar associated with g' . Then the Ricci curvature scalar associated with g is given in terms of g' as

$$R = e^{-2\phi} [R' - 2(m-1)\Delta'\phi - (m-2)(m-1)|d\phi|^2], \quad (\text{C.1.11})$$

where g'^{ij} is the inverse of g'_{ij} and Δ' denotes the Beltrami Laplacian with respect to g' .

Recall that the Beltrami Laplacian with respect to a metric \tilde{g} can be written as

$$\tilde{\Delta}\phi = \frac{1}{\sqrt{|\det(\tilde{g})|}} \partial_i \left[\sqrt{|\det(\tilde{g})|} \tilde{g}^{ij} \partial_j \phi \right]. \quad (\text{C.1.12})$$

Furthermore, when the metric \tilde{g} is changed by an overall sign, both its Beltrami Laplacian and Ricci curvature scalar change by precisely the same overall sign, due to the presence of \tilde{g}^{ij} in their respective formulæ. Combining these observations with (C.1.11) in two dimensions, again setting $\phi = \frac{1}{2} \log(|\zeta|)$, we find

$$R = \frac{1}{\zeta} \left[\tilde{R} - \tilde{\Delta}(\log(|\zeta|)) \right], \quad (\text{C.1.13})$$

from which (3.2.18a) then follows.

C.2 Lychagin–Rubtsov Metric with Arbitrary Background Dimension

We now wish to compute the curvature related to the metric (4.2.7). It will be more efficient to work in the vielbein formalism than a coordinate basis, so we summarise this below.

C.2.1 Vielbein formalism

Let (M, g) be an m -dimensional (semi-)Riemannian manifold with metric g and coordinatised by x^i with $i, j, \dots = 1, \dots, m$. Then, we can write $g = \frac{1}{2}g_{ij}dx^i \odot dx^j$. The vielbein tangent vector fields are denoted by $E_a \in \mathfrak{X}(M)$ for $a, b, \dots = 1, \dots, m$ and can be written in terms of the coordinate basis ∂_i as $E_a = E_a^i \partial_i$, where $(E_a^i) \in \mathcal{C}^\infty(M, \text{GL}(m))$ is an $m \times m$ matrix of smooth functions on M . Dually, we have $e^a = dx^i e_i^a \in \Omega^1(M)$ with $(e_i^a) \in \mathcal{C}^\infty(M, \text{GL}(m))$, satisfying the identities $E_a \lrcorner e^b = \delta_a^b$, $E_a^i e_i^b = \delta_a^b$, and $e_i^a E_a^j = \delta_i^j$. The metric can then be written in the vielbein frame as $g = \frac{1}{2}e^b \odot e^a \eta_{ab}$, with $\eta_{ab} = \text{diag}(-1, \dots, -1, 1, \dots, 1)$.

The structure functions $C_{ab}^c \in \mathcal{C}^\infty(M)$ are given by

$$[E_a, E_b] = C_{ab}^c E_c, \quad (\text{C.2.1a})$$

or, dually,

$$de^a = \frac{1}{2}e^c \wedge e^b C_{bc}^a. \quad (\text{C.2.1b})$$

The torsion and curvature two-forms,

$$T^a = \frac{1}{2}e^c \wedge e^b T_{bc}^a \quad \text{and} \quad R_a^b = \frac{1}{2}e^d \wedge e^c R_{cda}^b, \quad (\text{C.2.2a})$$

are then defined by the Cartan structure equations

$$de^a - e^b \wedge \omega_b^a =: -T^a \quad \text{and} \quad d\omega_a^b - \omega_a^c \wedge \omega_c^b =: -R_a^b, \quad (\text{C.2.2b})$$

where $\omega_a^b = e^c \omega_{ca}^b$ is the connection one-form. The associated Ricci tensor and scalar are

$$R_{ab} := R_{cab}^c \quad \text{and} \quad R := \eta^{ba} R_{ab}. \quad (\text{C.2.2c})$$

Furthermore, metric compatibility amounts to requiring

$$\omega_{ab} = -\omega_{ba} \quad \text{with} \quad \omega_{ab} := \omega_a^c \eta_{cb}. \quad (\text{C.2.3})$$

We obtain the Levi-Civita connection by imposing the metric compatibility condition (C.2.3) and the torsion-free constraint $T^a = 0$. A short calculation enables us to write

this connection explicitly as

$$\omega_{ab}{}^c = \frac{1}{2}(C^c{}_{ab} + C^c{}_{ba} + C_{ab}{}^c), \quad (\text{C.2.4})$$

with indices raised and lowered using η_{ab} . In this case, the curvature scalar (C.2.5) is

$$R = 2E_a C^a{}_b{}^b - C_{ab}{}^b C^a{}_c{}^c - \frac{1}{2}C_{abc} C^{acb} - \frac{1}{4}C_{abc} C^{abc}. \quad (\text{C.2.5})$$

C.2.2 Levi-Civita Connection of the Lychagin–Rubtsov Metric

Let (M, \hat{g}) be a Riemannian manifold and consider the metric (4.2.7) on T^*M , now assumed to be $2m$ -dimensional. Furthermore, let

$$\mathring{E}_a := \mathring{E}_a{}^i \frac{\partial}{\partial x^i} \quad \text{and} \quad \mathring{e}^a := dx^i \mathring{e}_i{}^a \quad (\text{C.2.6})$$

be the vielbeins and dual vielbeins on (M, \hat{g}) , with structure functions $\mathring{C}_{ab}{}^c$, and set

$$\begin{aligned} (\mathring{e}^A) &= (\mathring{e}^a, \mathring{e}_a) := \left(\sqrt{|\hat{f}|} dx^i \mathring{e}_i{}^a, \mathring{E}_a{}^i \nabla q_i \right), \\ (\mathring{\eta}_{AB}) &= \begin{pmatrix} \mathring{\eta}_{ab} & \mathring{\eta}_a{}^b \\ \mathring{\eta}^a{}_b & \mathring{\eta}^{ab} \end{pmatrix} := \begin{pmatrix} \text{sgn}(\hat{f}) \mathbb{1}_m & 0 \\ 0 & \mathbb{1}_m \end{pmatrix} \end{aligned} \quad (\text{C.2.7})$$

for multi-indices A, B . Then, the metric (4.2.7) becomes

$$\hat{g} = \frac{1}{2} \mathring{e}^B \odot \mathring{e}^A \mathring{\eta}_{AB}. \quad (\text{C.2.8})$$

Here, $\mathring{e}_i{}^a$ and $\mathring{E}_a{}^i$ only depend on the base coordinates x^i and not on the fibre coordinates q_i .

Next, dually, we have $\hat{E}_A \lrcorner \mathring{e}^B = \delta_A{}^B$ with $(\hat{E}_A) = (\hat{E}_a, \hat{E}^a)$ and

$$\hat{E}_a := \frac{1}{\sqrt{|\hat{f}|}} \mathring{E}_a{}^i \left(\frac{\partial}{\partial x^i} + \mathring{\Gamma}_{ij}{}^k q_k \frac{\partial}{\partial q_j} \right) \quad \text{and} \quad \hat{E}^a := \mathring{e}_i{}^a \frac{\partial}{\partial q_i}. \quad (\text{C.2.9})$$

A straightforward calculation then yields, for $[\hat{E}_A, \hat{E}_B] = \hat{C}_{AB}{}^C \hat{E}_C$, the relations

$$[\hat{E}_a, \hat{E}_b] = \frac{1}{\sqrt{|\hat{f}|}} \mathring{C}_{ab}{}^c \hat{E}_c - \hat{E}_{[a} \log(|\hat{f}|) \hat{E}_{b]} + \frac{1}{|\hat{f}|} \mathring{R}_{abc}{}^d q_d \hat{E}^c, \quad (\text{C.2.10a})$$

$$[\hat{E}_a, \hat{E}^b] = \frac{1}{2} \hat{E}^b \log(|\hat{f}|) \hat{E}_a - \frac{1}{\sqrt{|\hat{f}|}} \mathring{\omega}_{ac}{}^b \hat{E}^c, \quad (\text{C.2.10b})$$

$$[\hat{E}^a, \hat{E}^b] = 0, \quad (\text{C.2.10c})$$

where we have set $q_a := \hat{E}_a^i q_i$ and used the identities

$$\hat{\omega}_{ab}{}^c = \hat{E}_a^i \hat{E}_b^j \left(\hat{\Gamma}_{ij}{}^k \hat{e}_k{}^c - \frac{\partial}{\partial x^i} \hat{e}_j{}^c \right) \quad \text{and} \quad \hat{R}_{abc}{}^d = \hat{E}_a^i \hat{E}_b^j \hat{E}_c^k \hat{R}_{ijk}{}^l \hat{e}_l{}^d. \quad (\text{C.2.10d})$$

Reading off the structure functions $\hat{C}_{AB}{}^C$ from these relations and using the formula (C.2.4), the Levi-Civita connection $\hat{\omega}_{AB}{}^C$ for the metric (4.2.7) is given, in terms of the Levi-Civita connection $\hat{\omega}_{ab}{}^c$ for the background metric \hat{g} , by

$$\hat{\omega}_{AB}{}^C = \frac{1}{2}(\hat{C}^C{}_{AB} + \hat{C}^C{}_{BA} + \hat{C}_{AB}{}^C). \quad (\text{C.2.11})$$

C.2.3 Ricci Curvature Scalar of the Lychagin–Rubtsov Metric

Upon combining (C.2.10) and (C.2.11) with (C.2.5), the curvature scalar of the metric (4.2.7) is

$$\begin{aligned} \hat{R} &= \frac{1}{\hat{f}} \hat{R} - \frac{1}{4\hat{f}^2} \hat{R}_{abc}{}^d \hat{R}^{abce} q_d q_e - (m-1) \hat{\Delta}_B \log(|\hat{f}|) - \delta_{ab} \hat{E}^a \hat{E}^b \log(|\hat{f}|) \\ &\quad + \frac{\text{sgn}(\hat{f})}{4} (m-1)(m-2) \delta^{ab} \hat{E}_a \log(|\hat{f}|) \hat{E}_b \log(|\hat{f}|) \\ &\quad + \frac{1}{4} m(m-3) \delta_{ab} \hat{E}^a \log(|\hat{f}|) \hat{E}^b \log(|\hat{f}|), \end{aligned} \quad (\text{C.2.12a})$$

where $\hat{\Delta}_B$ is the Beltrami Laplacian for \hat{g} . Here, $\hat{R}_{abc}{}^d$ is the Riemann curvature tensor for the background metric \hat{g} and \hat{R} is the Ricci curvature scalar. In our coordinate basis, we find

$$\begin{aligned} \hat{R} &= \frac{1}{\hat{f}} \hat{R} - \frac{1}{4\hat{f}^2} \hat{R}_{ijk}{}^l \hat{R}^{ijklm} q_k q_m - (m-1) \hat{\Delta}_B \log(|\hat{f}|) - \hat{g}_{ij} \frac{\partial^2}{\partial q_i \partial q_j} \log(|\hat{f}|) \\ &\quad + \frac{1}{4\hat{f}} (m-1)(m-2) \hat{g}^{ij} \left(\frac{\partial}{\partial x^i} + \hat{\Gamma}_{ik}{}^l q_l \frac{\partial}{\partial q_k} \right) \log(|\hat{f}|) \left(\frac{\partial}{\partial x^j} + \hat{\Gamma}_{jm}{}^n q_n \frac{\partial}{\partial q_m} \right) \log(|\hat{f}|) \\ &\quad + \frac{1}{4} m(m-3) \hat{g}_{ij} \frac{\partial}{\partial q_i} \log(|\hat{f}|) \frac{\partial}{\partial q_j} \log(|\hat{f}|), \end{aligned} \quad (\text{C.2.12b})$$

where we have used (C.2.9). This verifies (3.2.14) and (4.2.8).

Finally, for the standard Euclidean background metric $\hat{g}_{ij} = \delta_{ij}$, we have $\hat{f} = f = \frac{1}{2} \Delta p$, with Δ the standard Laplacian on \mathbb{R}^m , so the formula (C.2.12b) simplifies to

$$\hat{R} = \frac{m-1}{4f^3} [(6-m) \partial_i f \partial^i f - 4f \Delta f]. \quad (\text{C.2.13})$$



The Cat's Cradle for Lorentzian Length Spaces

In Section 10.2, we provided the statement of the cat's cradle construction for globally hyperbolic Lorentzian length spaces, but stopped short of providing a full proof, instead discussing the idea of the process. Here we provide full details of the derivation of this construction, as originally found in [3, Proposition 3.4].

Since this proof is still rather extensive, we first offer a brief overview. The cat's cradle construction (see Figure D.0.2) is a recursive decomposition of a triangle $\Delta(x, y, z)$ into smaller triangles chosen such that the angle condition holds at angles of sign $\sigma = +1$, that is, the angles opposite their longest sides. From this construction, we infer a sequence of inequalities (D.0.3). We then continue with a similarly recursive construction in the model space, assembling a sequence of comparison triangles and hinges to infer a sequence of inequalities (D.0.7). Finally we show that the two sequences of inequalities converge to the same limit, which implies that hinge comparison at y cannot fail.

Proposition 10.2.2 (Cat's Cradle)

Let X be a connected, globally hyperbolic, regular Lorentzian length space with time function T and curvature bounded below by $K \in \mathbb{R}$. Fix $0 < \varepsilon < \frac{1}{2}$ and let $\Delta = \Delta(x, y, z)$ be a timelike triangle in X . If, for every timelike triangle $\Delta(x', y', z')$ satisfying both of the following

- (i) $x \leq x' \ll y' \ll z' \leq z$,
- (ii) $d_T(x', z') \leq (1 - \varepsilon)d_T(x, z)$,

the angle condition holds at all vertices of $\Delta(x', y', z')$, then the angle condition also holds at y in $\Delta(x, y, z)$.

Proof. To begin, set $L := d_T(x, z)$ and $y_0 := y$. Assume without loss of generality that $d_T(x, y_0) \geq d_T(y_0, z)$, otherwise the roles of x and z should be interchanged. Let y_1 be the point on the timelike distance realiser γ_{xy_0} such that $d_T(x, y_1) = \varepsilon L$, which exists since $\varepsilon < \frac{1}{2}$ and $d_T(x, y_0) \geq \frac{1}{2}L$. As timelike triangles are degenerate in the null distance, see (10.1.5), it follows that $d_T(y_1, z) = (1 - \varepsilon)L$. The new triangle $\Delta_1 = \Delta(y_1, y_0, z)$, see Figure D.0.1, is a timelike triangle satisfying assumptions (i) and (ii) of our statement, hence the angle condition holds at all vertices in Δ_1 by construction.

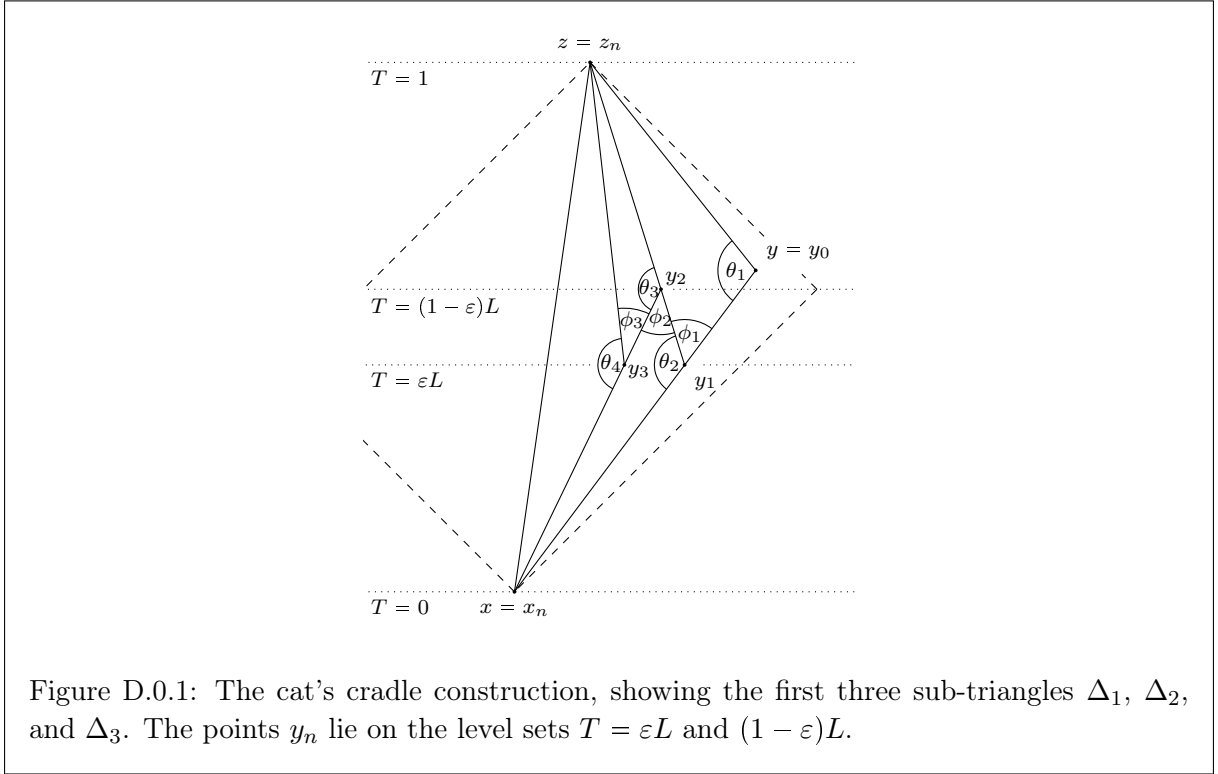


Figure D.0.1: The cat's cradle construction, showing the first three sub-triangles Δ_1 , Δ_2 , and Δ_3 . The points y_n lie on the level sets $T = \varepsilon L$ and $(1 - \varepsilon)L$.

We continue this construction recursively, picking points y_n to form new triangles. For even n , pick y_n on the distance realiser $\gamma_{y_{n-1}z}$ such that $d_T(y_n, z) = \varepsilon L$ and $d_T(x, y_n) = (1 - \varepsilon)L$. Again, such a y_n exists, as

$$d_T(y_{n-1}, z) = (1 - \varepsilon)L > \varepsilon L = d_T(y_n, z). \quad (\text{D.0.1})$$

This defines a new triangle $\Delta_n = \Delta(x, y_{n-1}, y_n)$ for $n \geq 1$ even. Similarly, for odd $n \geq 1$, pick y_n on the distance realiser $\gamma_{xy_{n-1}}$ and define $\Delta_n = \Delta(y_n, y_{n-1}, z)$. In both cases, Δ_n satisfies assumptions (i) and (ii) from our statement, so the angle condition holds at all vertices of Δ_n by construction.

Now consider the angles in Δ_n . Let $\theta_n := \sphericalangle_{y_{n-1}}(x, z)$ be the angle at y_{n-1} in Δ_n , which is equal to the angle $\sphericalangle_{y_{n-1}}(x, y_n)$ when n is even and $\sphericalangle_{y_{n-1}}(y_n, z)$ when n is odd. In particular, $\theta_1 = \sphericalangle_y(x, z)$ is equal to the angle at y in $\Delta(x, y, z)$ — the angle we are interested in. Denote by ϕ_n the angle at y_n in Δ_n , which is adjacent to θ_{n+1} in the subsequent triangle. For even n , $\phi_n = \sphericalangle_{y_n}(y_{n-1}, x)$, while when n is odd, $\phi_n = \sphericalangle_{y_n}(y_{n-1}, z)$. In either case, by Proposition 8.2.6 we find $\phi_n = \theta_{n+1}$, but with opposite signs σ .

Now set $l_n := \tau(x, y_n) + \tau(y_n, z)$, for $n \geq 0$. For odd n , we find

$$\begin{aligned} l_{n-1} &= \tau(x, y_{n-1}) + \tau(y_{n-1}, z) \\ &= \tau(x, y_n) + \tau(y_n, y_{n-1}) + \tau(y_{n-1}, z) \\ &\leq \tau(x, y_n) + \tau(y_n, z) = l_n, \end{aligned} \tag{D.0.2}$$

where in the first step we use that x, y_n , and y_{n-1} all lie along a distance realiser and the additivity of length (8.1.5), while in the second step we use reverse triangle inequality on Δ_n . For even $n > 0$, we find $l_{n-1} \leq l_n$ by a similar argument. Consequently, we obtain the sequence

$$0 < l_0 \leq l_1 \leq \dots \leq \tau(x, z). \tag{D.0.3}$$

The initial strict inequality is due to $\Delta(x, y_0, z)$ being a non-degenerate timelike triangle, while the final inequality in the chain follows from applying reverse triangle inequality to $\Delta(x, y_n, z)$ (we can terminate the sequence at any n in this way).

The sequence $\{l_n\}_{n \geq 0}$ in (D.0.3) is a Cauchy sequence, as it is monotone increasing and bounded above by $\tau(x, z)$ (which is finite as $\Delta(x, y, z)$ satisfies size-bounds). Therefore, $l_n - l_{n-1} \rightarrow 0$ as $n \rightarrow \infty$. For odd $n \geq 1$, we find

$$\begin{aligned} l_n - l_{n-1} &= \tau(y_n, z) - \tau(y_{n-1}, z) + \tau(x, y_n) - \tau(x, y_{n-1}) \\ &= \tau(y_n, z) - \tau(y_{n-1}, z) - \tau(y_n, y_{n-1}), \end{aligned} \tag{D.0.4}$$

again using that x, y_n , and y_{n-1} lie on a distance realiser in that order and the additivity of length (8.1.5). Similarly, for even $n \geq 1$ we find

$$l_n - l_{n-1} = \tau(x, y_n) - \tau(x, y_{n-1}) - \tau(y_{n-1}, y_n). \tag{D.0.5}$$

Hence, this value can be visualised as the excess in the triangle Δ_n , that is, the value by which the longest side exceeds the sum of the two shortest sides, see Figure D.0.1.

Let $x_n := x$ and $z_n := z$ for all $n \geq 0$. We now carry out a similar construction in the model space $\mathbb{L}^2(K)$ by arranging comparison triangles $\bar{\Delta}_n$ for Δ_n . Since, in general, the angles

in $\bar{\Delta}_n$ do not match those in Δ_n , the construction in $\mathbb{L}^2(K)$ does not fit together quite as neatly, see Figure D.0.2. In contrast to the construction in X , here we begin by considering the comparison hinge $(\gamma_{\bar{y}_0\bar{x}_0}, \gamma_{\bar{y}_0\bar{z}_0}; \bar{\omega}_1)$ for $(\gamma_{y_0x_0}, \gamma_{y_0z_0}; \theta_1)$ in $\mathbb{L}^2(K)$. Recall that, by the definition of a comparison hinge, Definition 8.2.1, $(\bar{x}_0, \bar{y}_0, \bar{z}_0)$ is a triple of points such that $\tau(\bar{x}_0, \bar{y}_0) = \tau(x_0, y_0)$, $\tau(\bar{y}_0, \bar{z}_0) = \tau(y_0, z_0)$, and $\bar{\omega}_1 = \theta_1$. In particular, there is no restriction on $\tau(\bar{x}_0, \bar{z}_0)$ a priori and instead we set out to obtain one (we have not constructed a comparison triangle for $\Delta = \Delta(x_0, y_0, z_0)$ here, for example).

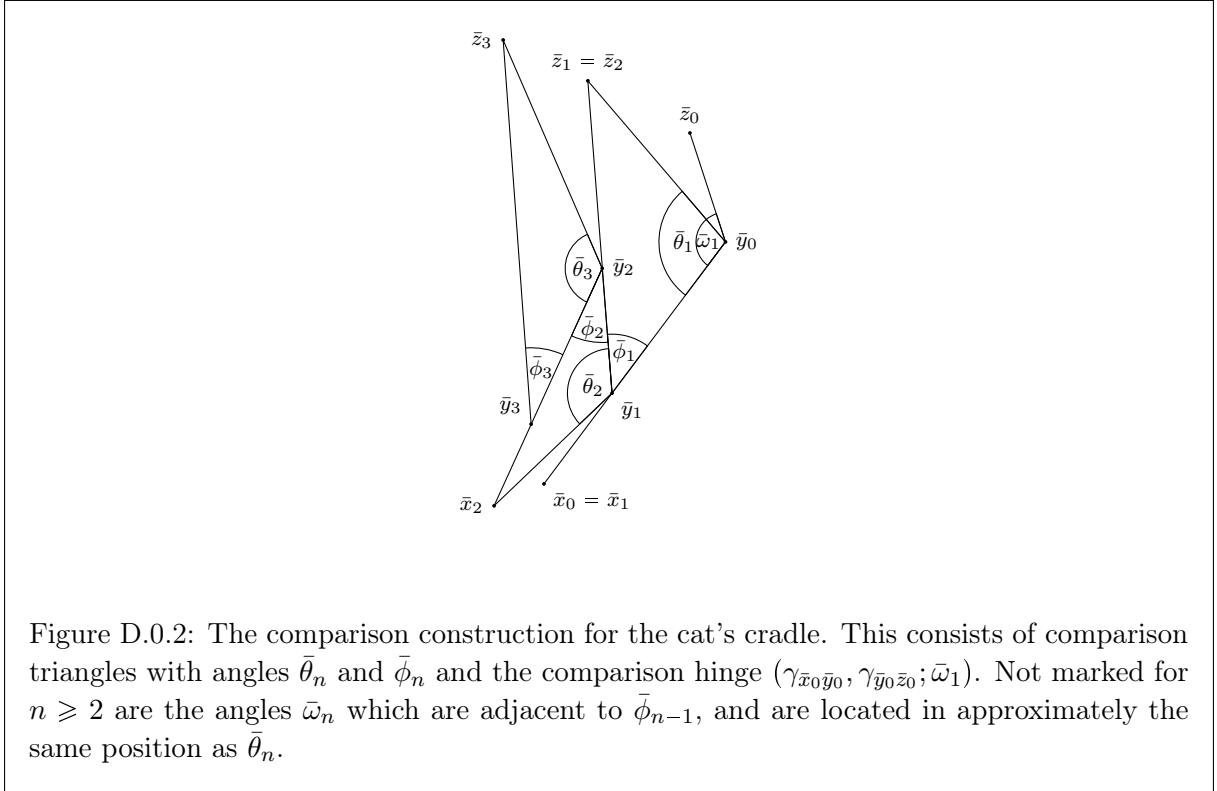


Figure D.0.2: The comparison construction for the cat's cradle. This consists of comparison triangles with angles $\bar{\theta}_n$ and $\bar{\phi}_n$ and the comparison hinge $(\gamma_{\bar{x}_0\bar{y}_0}, \gamma_{\bar{y}_0\bar{z}_0}; \bar{\omega}_1)$. Not marked for $n \geq 2$ are the angles $\bar{\omega}_n$ which are adjacent to $\bar{\phi}_{n-1}$, and are located in approximately the same position as $\bar{\theta}_n$.

Using this hinge, we now recursively construct the comparison triangles $\bar{\Delta}_n$ for $n \geq 1$. For odd n , fix \bar{y}_n on the distance realiser $\gamma_{\bar{x}_{n-1}, \bar{y}_{n-1}}$ such that $\tau(\bar{x}_{n-1}, \bar{y}_n) = \tau(x_{n-1}, y_n)$ and choose \bar{z}_n such that the timelike triangle $\bar{\Delta}_n = \Delta(\bar{y}_n, \bar{y}_{n-1}, \bar{z}_n)$ has the same side lengths as Δ_n with respect to τ , i.e. is a comparison triangle for Δ_n . Finally, set $\bar{x}_n = \bar{x}_{n-1}$. For even n , similarly fix \bar{y}_n on the distance realiser $\gamma_{y_{n-1}z_{n-1}}$ such that $\tau(\bar{y}_n, \bar{z}_{n-1}) = \tau(y_n, z_{n-1})$, construct a comparison triangle $\bar{\Delta}_n = \Delta(\bar{x}_n, \bar{y}_{n-1}, \bar{y}_n)$, and set $\bar{z}_n = \bar{z}_{n-1}$.

The choice of the two new points at each stage defines new angles. Denote by $\bar{\theta}_n$ the angle in $\bar{\Delta}_n$ at \bar{y}_{n-1} ,¹ by $\bar{\phi}_n$ the angle in $\bar{\Delta}_n$ at \bar{y}_n , and by $\bar{\omega}_{n+1}$ the angle of the remaining open (not

¹Note that the angle $\bar{\theta}_n = \tilde{\chi}_{y_{n-1}}(y_n, z_n)$ for odd n and $\bar{\theta}_n = \tilde{\chi}_{y_{n-1}}(y_n, x_n)$ for even n .

necessarily comparison) hinge $(\gamma_{\bar{y}_n \bar{x}_n}, \gamma_{y_n z_n}; \bar{\omega}_{n+1})$ adjacent to $\bar{\phi}_n$, see Figure D.0.2. Note that $\bar{\phi}_n = \bar{\omega}_{n+1}$, but with opposite sign, again by Proposition 8.2.6. As the angle condition holds at y_{n-1} and y_n in Δ_n by construction, we have $\theta_n \leq \bar{\theta}_n$ at q_{n-1} and $\phi_n \geq \bar{\phi}_n$ for the angle with sign $\sigma = -1$ at y_n . Furthermore, by construction, $\bar{\omega}_1 = \theta_1 \leq \bar{\theta}_1$. More generally, using the inequalities $\bar{\phi}_n = \bar{\omega}_{n+1}$, $\theta_n \leq \bar{\theta}_n$, $\phi_n \geq \bar{\phi}_n$, and Proposition 8.2.6 we obtain

$$\bar{\omega}_{n+1} = \bar{\phi}_n \leq \phi_n = \theta_{n+1} \leq \bar{\theta}_{n+1}, \tag{D.0.6}$$

for all $n \geq 1$. Consequently, $\bar{\omega}_n \leq \bar{\theta}_n$ for all $n \geq 1$ and the relative sizes of the angles are as depicted in Figure D.0.2. Hence, by monotonicity of the hyperbolic law of cosines, see Corollary 8.2.2, we have $\tau(\bar{x}_{n-1}, \bar{z}_{n-1}) \leq \tau(\bar{x}_n, \bar{z}_n)$ and the following sequence of inequalities holds:

$$\tau(\bar{x}_0, \bar{z}_0) \leq \tau(\bar{x}_1, \bar{z}_1) \leq \dots \tag{D.0.7}$$

At this stage, our construction is complete and we may move on to proving that the angle condition holds at y in $\Delta = \Delta(x, y, z)$. We do so by first illustrating that the hinge condition (8.2.8) holds at y in Δ . So assume for contradiction that $\tau(x, z) < \tau(\bar{x}_0, \bar{z}_0)$, that is, the hinge condition fails at y in Δ , and set $C := \tau(\bar{x}_0, \bar{z}_0) - \tau(x, z) > 0$. Since $\tau(\bar{x}_n, \bar{z}_n) \geq \tau(\bar{x}_0, \bar{z}_0)$ and $\tau(x, z) \geq l_n$ by (D.0.7) and (D.0.3) respectively, we find

$$\tau(\bar{x}_n, \bar{z}_n) - l_n \geq \tau(\bar{x}_0, \bar{z}_0) - \tau(x, z) = C > 0, \tag{D.0.8}$$

for all $n \geq 0$. Given a subsequence n_i such that $\tau(\bar{x}_{n_i}, \bar{z}_{n_i}) - l_{n_i} \rightarrow 0$, we would have a contradiction, implying that the hinge condition holds at y in $\Delta(x, y, z)$. Now let $\bar{\Delta} = \Delta(\bar{x}, \bar{y}, \bar{z})$ be the comparison triangle for $\Delta = \Delta(x, y, z)$ in $\mathbb{L}^2(K)$ and denote the angle at \bar{y} in $\bar{\Delta}$ by $\bar{\theta}$. As the hinge condition holds at y , we know that

$$\tau(\bar{x}, \bar{z}) = \tau(x, z) \geq \tau(\bar{x}_0, \bar{z}_0). \tag{D.0.9}$$

Since $\tau(\bar{x}, \bar{y}) = \tau(\bar{x}_0, \bar{y}_0)$ and $\tau(\bar{y}, \bar{z}) = \tau(\bar{y}_0, \bar{z}_0)$, law of (hyperbolic) cosines monotonicity (see Corollary 8.2.2) implies that $\bar{\theta} \geq \bar{\omega}_1 = \theta_1$, where θ_1 is the angle at y in $\Delta(x, y, z)$. That is, the angle condition holds at y in $\Delta(x, y, z)$, cf. Definition 10.1.10.

It remains to show that there exists a subsequence n_i such that $\tau(\bar{x}_{n_i}, \bar{z}_{n_i}) - l_{n_i} \rightarrow 0$. We do so by finding a subsequence n_i , such that the time separation between the vertices y_{n_i-1} and y_{n_i} of the triangle Δ_{n_i} (in their causal order) is uniformly bounded away from zero.

Assume for a contradiction that no subsequences n_i are such that the time separation between y_{n_i-1} and y_{n_i} is bounded away from zero. Then we have $\lim_{n \rightarrow \infty} \tau(y_{2n-1}, y_{2n}) =$




Figure D.0.3: The limiting configuration of the cat's cradle, when we assume that the time separation between y_{n_i-1} and y_{n_i} is not bounded away from zero for any subsequence n_i . The configuration consists of two “orthogonal” null segments — if y_a and y_c were to coincide, x and z would lie on a null segment with $\tau(x, z) = 0$. Since $\tau(x, z) > 0$, this is not possible.

$\lim_{n \rightarrow \infty} \tau(y_{2n+1}, y_{2n}) = 0$. Now consider the sequence of triples $\{(y_{2n-1}, y_{2n}, y_{2n+1})\}_{n \geq 1}$, which, by global hyperbolicity of X , lies in the compact set $J(x, z) \times J(x, z) \times J(x, z)$. Consequently, after passing to some subsequence n_i , the triples converge to a limit point (y_a, y_b, y_c) . Inspecting the time function, we see that

$$T(q_{2n-1}) - T(p) = d_T(p, q_{2n-1}) = \varepsilon L = d_T(p, q_{2n+1}) = T(q_{2n+1}) - T(p). \quad (\text{D.0.10})$$

Furthermore, since $q_{2n-1} \leq q_{2n}$, (10.1.3b) yields

$$T(q_{2n}) - T(q_{2n \pm 1}) = d_T(q_{2n \pm 1}, q_{2n}) = (1 - 2\varepsilon)L > 0. \quad (\text{D.0.11})$$

Hence, $T(q_{2n-1}) = T(q_{2n+1}) \neq T(q_{2n})$, which in the limit $n \rightarrow \infty$ implies that $q_a \neq q_b \neq q_c$. Furthermore, applying continuity of τ to $\lim_{n \rightarrow \infty} \tau(y_{2n-1}, y_{2n}) = \lim_{n \rightarrow \infty} \tau(y_{2n+1}, y_{2n}) = 0$ yields

$$\tau(y_a, y_b) = \tau(y_c, y_b) = 0. \quad (\text{D.0.12})$$

Consequently, y_b is either null related or spacelike related to both y_a and y_c , but not equal to either of them.

Recall that x, y_{2n+1} , and y_{2n} lie on a distance realiser, hence $\tau(x, y_{2n+1}) + \tau(y_{2n+1}, y_{2n}) = \tau(x, y_{2n})$. Again applying continuity of τ , we find

$$\tau(x, y_c) + \tau(y_c, y_b) = \tau(x, y_b). \quad (\text{D.0.13})$$

By Proposition 8.1.19, X is causally closed and hence $x \leq y_{2n+1} \leq y_{2n}$ implies $x \leq y_c \leq y_b$. In

particular, x, y_c , and y_b lie along a distance realiser with a non-constant null piece $\gamma_{y_c y_b}$. Thus by regularity of X , the whole distance realiser is null and therefore $\tau(x, y_b) = 0$.

Similarly, working from $\tau(y_{2n-1}, y_{2n}) + \tau(y_{2n}, z) = \tau(y_{2n-1}, z)$ and $y_{2n-1} \preceq y_{2n} \preceq z$, in the limit we find

$$\tau(y_a, y_b) + \tau(y_b, z) = \tau(y_a, z) \quad \text{and} \quad y_a \preceq y_b \preceq z, \quad (\text{D.0.14})$$

from which it follows that y_a, y_b , and z lie on a null distance realiser and $\tau(y_b, z) = 0$. See Figure D.0.3 for the limiting configuration. Therefore,

$$\lim_{i \rightarrow \infty} l_{2n_i} = \lim_{i \rightarrow \infty} (\tau(x, y_{2n_i}) + \tau(y_{2n_i}, z)) = \tau(x, y_b) + \tau(y_b, z) = 0 \quad (\text{D.0.15})$$

However, (D.0.3) states that l_n is a non-decreasing, positive sequence, hence we have a contradiction.

We now have a subsequence n_i such that the time separation between the vertices y_{n_i-1} and y_{n_i} is uniformly bounded away from zero, that is, the τ -length of the short side $\gamma_{y_{n_i-1} y_{n_i}}$ (or $\gamma_{y_{n_i} y_{n_i-1}}$ if the vertices have the opposite time orientation) in Δ_{n_i} is uniformly bounded away from zero. This implies that the τ -length of the longest side in Δ_{n_i} is also uniformly bounded away from zero by reverse triangle inequality. Furthermore, in the limit $n_i \rightarrow \infty$ we have triangle equality for Δ_{n_i} , since the excess of Δ_{n_i} is equal to $l_{n_i} - l_{n_i-1}$, which tends to zero by the discussion prior to (D.0.4). Consequently, in the limit, Δ_{n_i} either tends to a point or a degenerate admissible causal triangle, cf. Definition 8.1.7. Since two of the sides of Δ_{n_i} have τ -length bounded away from zero, at least two of the vertices remain timelike related in the limit, so the sequence of triangles cannot degenerate to a point and instead tends to a degenerate admissible causal triangle. The realisability lemma [116, Lemma 4.6] (see also our discussion following Definition 8.1.11) implies that the unique comparison triangle for a degenerate admissible triangle is a timelike line in $\mathbb{L}^2(K)$. It follows that the sequence of comparison triangles $\bar{\Delta}_{n_i}$ tends towards a timelike line.

By construction, the angle $\bar{\phi}_{n_i}$ lies between the two timelike sides of the triangle $\bar{\Delta}_{n_i}$ whose τ -lengths are uniformly bounded away from zero.² Since these timelike sides tend to segments of the same timelike line, the definition of an angle (8.2.1) in $\mathbb{L}^2(K)$ implies that $\bar{\phi}_{n_i} \rightarrow 0$ as $n_i \rightarrow \infty$. Furthermore, the identity $\bar{\omega}_{n+1} = \bar{\phi}_n$ implies that $\bar{\omega}_{n_i+1} \rightarrow 0$. As $\bar{\omega}_{n_i+1} = \sphericalangle_{\bar{y}_{n_i}}(\bar{x}_{n_i}, \bar{z}_{n_i})$ is the angle in the (not necessarily comparison) hinge $(\gamma_{\bar{x}_{n_i} \bar{y}_{n_i}}, \gamma_{\bar{y}_{n_i} \bar{z}_{n_i}}; \bar{\omega}_{n_i+1})$, we conclude from Corollary 8.2.3 that the excess in our sequence of hinges (interpreted as triangles) goes to zero, that is, $\tau(\bar{x}_{n_i}, \bar{z}_{n_i}) - \tau(\bar{x}_{n_i}, \bar{y}_{n_i}) - \tau(\bar{y}_{n_i}, \bar{z}_{n_i}) = \tau(\bar{x}_{n_i}, \bar{z}_{n_i}) - l_{n_i} \rightarrow 0$ along our subsequence, as required. \square

²Recall that comparison triangles have the same side-lengths as their associated ‘‘real’’ triangle and that $\bar{\phi}_{n_i}$ is adjacent to both the longest side of $\bar{\Delta}_{n_i}$ and either $\gamma_{\bar{y}_{n_i-1} \bar{y}_{n_i}}$ or $\gamma_{\bar{y}_{n_i} \bar{y}_{n_i-1}}$.

References

- [1] L. Napper, I. Roulstone, V. Rubtsov, and M. Wolf, *Monge–Ampère geometry and vortices*, *Non-linearity* **37** (2024) 045012 [2302.11604v2].
- [2] T. Beran, L. Napper, and F. Rott, *Alexandrov’s Patchwork and the Bonnet–Myers Theorem for Lorentzian length spaces*, *Trans. AMS (To Appear)* (2023) 34 [2302.11615 [math.DG]].
- [3] T. Beran, J. Harvey, L. Napper, and F. Rott, *A Toponogov globalisation result for Lorentzian length spaces*, *Math. Ann. (To Appear)* (2024) 37 [2309.12733 [math.DG]].
- [4] A. Kushner, V. V. Lychagin, and V. N. Rubtsov, *Contact geometry and non-linear differential equations*, Cambridge University Press, 2007 [doi].
- [5] A. Kolmogorov, I. Petrovsky, and N. Piskunov, *Investigation of the Equation of Diffusion combined with increasing of the substance and its application to a biology problem*, *Bulletin of Moscow State University Series A* **1** (1937) 1.
- [6] O. V. Rudenko, V. A. Khokhlova, and M. F. Hamilton, *Historical aspects of the Khokhlov–Zabolotskaya equation and its generalisations*, *J. Acoust. Soc. Am.* **126** (2009).
- [7] S. Chynoweth and M. Sewell, *Dual variables in semigeostrophic theory*, *Proc. R. Soc. Lond. A* **424** (1989) 155.
- [8] R. D’Onofrio, R. Ortenzi, and I. Roulstone, *Monge–Ampère geometry and the Eady problem*, 2405.09218 [math-ph].
- [9] J. D. E. Grant, *On Self-Dual Gravity*, *Phys. Rev. D* **48** (1993) 2606 [9301014v2 [gr-qc]].
- [10] B. Banos, *Complex solutions of Monge–Ampère equations*, *J. Geom. Phys.* **61** (2011) [1104.0362v2].
- [11] V. V. Lychagin, *Contact geometry and non-linear second-order differential equations*, *Russ. Math. Surv.* **34** (1979) 149.
- [12] A. M. Vinogradov, *Multivalued solutions and a principle of classification of non-linear differential equations*, *Dokl. Akad. Nauk* **210** (1973) 11.
- [13] A. M. Vinogradov and I. S. Krasil’shchik, *What is the Hamiltonian formalism?*, *Uspekhi Mat. Nauk* **30** (1975) 173.
- [14] A. M. Vinogradov and B. A. Kupershmidt, *The structures of Hamiltonian mechanics*, *Uspekhi Mat. Nauk* **32** (1977) 175.
- [15] V. V. Lychagin, V. N. Rubtsov, and I. V. Chekalov, *A classification of Monge–Ampère equations*, *Ann. Sci. Ec. Norm. Sup.* **26** (1993) 281.
- [16] V. N. Rubtsov and I. Roulstone, *Holomorphic structures in hydrodynamical models of nearly geostrophic flow*, *Proc. R. Soc. A* **457** (2001) 1519.
- [17] I. Roulstone, B. Banos, J. D. Gibbon, and V. N. Rubtsov, *Kähler geometry and Burgers’ vortices*, *Proc. Ukrain. Nat. Acad. Sci.* **16** (2009) 303 [nlin/0509023].
- [18] I. Roulstone, B. Banos, J. D. Gibbon, and V. N. Rubtsov, *A geometric interpretation of coherent structures in Navier–Stokes flows*, *Proc. R. Soc. A* **465** (2009) 2015.
- [19] B. Banos, V. N. Rubtsov, and I. Roulstone, *Monge–Ampère structures and the geometry of incompressible flows*, *J. Phys. A* **49** (2016) 244003 [1510.02327 [math-ph]].
- [20] J. D. Gibbon, *The three-dimensional Euler equations: where do we stand?*, *Physica D* **237** (2008)

- 1894.
- [21] S. Douady, Y. Couder, and M. E. Brachet, *Direct observation of the intermittency of intense vorticity filaments in turbulence*, *Phys. Rev. Lett.* **67** (1991) 983.
- [22] H. K. Moffatt, S. Kida, and K. Ohkitani, *Stretched vortices – the sinews of turbulence; large-Reynolds-number asymptotics*, *J. Fluid Mech.* **259** (1994) 241.
- [23] M. Larchevêque, *Equation de Monge–Ampère et écoulements incompressibles bidimensionnels*, *Comptes Rendus Acad. Sci. Paris Ser. II* **311** (1990) 33.
- [24] M. Larchevêque, *Pressure field, vorticity field, and coherent structures in two-dimensional incompressible turbulent flows*, *Theor. Comp. Fluid Dyn.* **5** (1993) 215.
- [25] J. Weiss, *The dynamics of enstrophy transfer in two-dimensional hydrodynamics*, *Physica D* **48** (1991) 273.
- [26] N. J. Hitchin, *The geometry of three-forms in six dimensions*, *J. Diff. Geom.* **55** (2000) 547 [[math/0010054](#)].
- [27] F. Cantrijn, A. Ibort, and M. De León, *On the geometry of multisymplectic manifolds*, *J. Austr. Math. Soc.* **66** (1999) 303.
- [28] J. C. Baez, A. E. Hoffnung, and C. L. Rogers, *Categorified symplectic geometry and the classical string*, *Commun. Math. Phys.* **293** (2010) 701 [[0808.0246](#) [[math-ph](#)]].
- [29] C. L. Rogers, *Higher symplectic geometry*, PhD thesis, University of California (2011) [[1106.4068](#) [[math-ph](#)]].
- [30] H. K. Moffatt, *The degree of knottedness of tangled vortex lines*, *J. Fluid Mech.* **35** (1969) 117.
- [31] H. K. Moffatt and R. L. Ricca, *Helicity and the Čaluşăreanu invariant*, *Proc. R. Soc. A* **439** (1992) 1906.
- [32] H. K. Moffatt and A. Tsinober, *Helicity in laminar and turbulent flow*, *Ann. Rev. Fluid Mech.* **24** (1992) 281.
- [33] J. Marsden and A. Weinstein, *Reduction of symplectic manifolds with symmetry*, *Rep. Math. Phys.* **5** (1974) 121.
- [34] C. Blacker, *Reduction of multisymplectic manifolds*, *Lett. Math. Phys.* **111** (2021) 31 [[2002.10062](#) [[math.SG](#)]].
- [35] B. Banos, *Non-degenerate Monge–Ampère structures in dimension 6*, *Lett. Math. Phys.* **62** (2002) 1 [[math/0211185](#)].
- [36] M. Wolf, I. Roulstone, V. Rubtsov, and J. McOrist, *On the Monge–Ampère geometry of incompressible fluid flows*, (unpublished) (2019).
- [37] V. Rubtsov. *Nonlinear PDEs, Their Geometry, and Applications* chapter Chapter 3: Geometry of Monge–Ampère structures, 95. Springer Birkhäuser 2019. Edited by R. A. Kycia, M. Ulan, and E. Schneider [[doi](#)].
- [38] M. Kossowski, *Local existence of multivalued solutions to analytic symplectic Monge–Ampère equations (the nondegenerate and type changing cases)*, *Indiana Univ. Math. J.* **40** (1991) 123.
- [39] A. Cannas da Silva. *Handbook of Differential Geometry* volume II chapter 3 - Symplectic Geometry, 81. Elsevier B.V. 2006. Edited by F. Dillen and L. Verstraelen.
- [40] X. Gràcia, J. de Lucas, X. Rivas, and N. Román-Roy, *On Darboux theorems for geometric structures induced by closed forms*, *Rev. Real Acad. Cienc. Exactas Fis. Nat. Ser. A-Mat* **118** (2024)

- [2306.08556v1].
- [41] J. Lee, *Introduction to Smooth Manifolds*, Springer New York, NY, 2012 [doi].
- [42] G. Ishikawa and Y. Machida, *Extra singularities of geometric solutions to Monge–Ampère equation of three variables*, *Kyoto Univ. Res. Inf. Repos.* **1502** (2006) 41.
- [43] G. Ishikawa and Y. Machida, *Monge–Ampère systems with Lagrangian pairs*, *SIGMA* **11** (2015) 32 [1503.01529 [math.DG]].
- [44] R. D’Onofrio, G. Ortenzi, I. Roulstone, and V. N. Rubtsov, *Solutions and singularities of the semigeostrophic equations via the geometry of Lagrangian submanifolds*, *Proc. R. Soc. A* **479** (2023) 20220682 [2209.13337 [math-ph]].
- [45] M. Nakahara, *Geometry, Topology, and Physics*, IoP Publishing, London, 2003 [doi].
- [46] V. Cruceanu, P. Fortuny, and P. Gadea, *A Survey on Paracomplex Geometry*, *Rocky Mtn. J. Math.* **26** (1996) 83.
- [47] A. Newlander and L. Nirenberg, *Complex Analytic Coordinates in Almost Complex Manifolds*, *Annals of Mathematics* **65** (1957) 391.
- [48] S. M. Webster, *A new proof of the Newlander–Nirenberg theorem*, *Math. Z.* **201** (1989) 303.
- [49] G. Nagaraj, *Note on the locally product and almost locally product structures*, *Proc. Indian Acad. Sci.* **65** (1967) 270.
- [50] V. Cortés, C. Mayer, T. Mohaupt, and F. Saueressig, *Special geometry of Euclidean supersymmetry 1. Vector multiplets*, *J. HEP.* **2004** (2004) [0312001v1 [hep-th]].
- [51] I. Vaisman, *Towards a double field theory on para-Hermitian manifolds*, *J. Math. Phys.* **54** (2013) 123507 [1209.0152 [math.DG]].
- [52] D. Cioranescu, P. Donato, and M. P. Poque, *An Introduction to Second Order Partial Differential Equations: Classical and Variational Solutions*, World Scientific, 2018 [doi].
- [53] B. Kruglikov, *Classification of Monge–Ampère equations with two variables*, *Banach Center Publ.* **50** (1999) 179.
- [54] R. D’Onofrio, *A note on optimal transport and Monge–Ampère geometry*, *J. Geom. Phys.* **186** (2023) 104771 [2302.05722 [math-ph]].
- [55] A. Enciso and D. Paralta-Salas, *Knotted vortex lines and vortex tubes in stationary fluid flows*, *Newslett. Eur. Math. Soc.* (2015) 26.
- [56] J. Jeong and F. Hussian, *On the identification of a vortex*, *J. Fluid Mech.* **285** (1995) 69.
- [57] G. Haller, *An objective definition of a vortex*, *J. Fluid. Mech.* **525** (2005) 1.
- [58] M. Chong, A. Perry, and B. Cantwell, *A general classification of three-dimensional flow fields*, *Phys. Fluids* **2** (1990) 765.
- [59] I. Roulstone, A. A. White, and S. A. Clough, *Geometric invariants of the horizontal velocity gradient tensor and their dynamics in shallow water flow*, *Q. J. R. Meteorol. Soc.* **140** (2014) 2527.
- [60] M. Ghomi and C. Xiong, *Nonnegatively curved hypersurfaces with free boundary on a sphere*, *Calc. Var. Partial Differ. Equ.* **58** (2019) [1707.03323v2].
- [61] S. Nakajima, *Über konvexe Kurven und Flächen*, *Tôhoku Math. J.* **29** (1928) 227.
- [62] H. Tietze, *Über Konvexität im kleinen und im großen und über gewisse den Punkter einer Menge zugeordnete Dimensionzahlen*, *Math. Z.* **28** (1928) 697.
- [63] C. Bjorndahl and Y. Karshon, *Revisiting Tietze–Nakajima: Local and global convexity for maps*,

- Canad. J. Math. **62** (2010) 975 [0701745].
- [64] J. Hunt, A. Wray, and P. Moin, *Eddies, streams, and convergence zones in turbulent flows*, Proceedings of the Summer Program 1988: Center for Turbulence Research (1988) 193.
- [65] Y. Dubief and F. Delcayre, *On coherent-vortex identification in turbulence*, J. Turbulence **1** (2000).
- [66] D. G. Ebin and J. Marsden, *Groups of diffeomorphisms and the motion of an incompressible fluid*, Ann. Math. **92** (1970) 102.
- [67] C. Cao, M. A. Rammaha, and E. S. Titi, *The Navier–Stokes equations on the rotating 2-D sphere: Gevrey regularity and asymptotic degrees of freedom*, Z. Angew. Math. Phys. **50** (1999) 341.
- [68] M. H. Kobayashi, *On the Navier–Stokes equations on manifolds with curvature*, J. Eng. Math **60** (2008) 55.
- [69] V. Pierfelice, *The incompressible Navier–Stokes equations on non-compact manifolds*, J. Geom. Anal. **27** (2017) 577 [1406.1644 [math.AP]].
- [70] C. H. Chan, M. Czubak, and M. M. Disconzi, *The formulation of the Navier–Stokes equations on Riemannian manifolds*, J. Geom. Phys. **121** (2017) 335 [1608.05114 [math.AP]].
- [71] M. Samavaki and J. Tuomela, *Navier–Stokes equations on Riemannian manifolds*, J. Geom. Phys. **148** (2020) 103543 [1812.09015 [math.NA]].
- [72] G. Rotskoff, *The Gauss–Bonnet theorem*, Masters’ thesis, University of Chicago (2010).
- [73] G. S. Birman and K. Nomizu, *The Gauss–Bonnet theorem for 2-dimensional space-times*, Michigan Math. J. **31** (1984) 77.
- [74] M. Steller, *A Gauss–Bonnet formula for metrics with varying signature*, Z. Anal. Anwend. **45** (2006) 143.
- [75] G. I. Taylor and A. E. Green, *Mechanism of the production of small eddies from large ones*, Proc. R. Soc. A **158** (1937) 499.
- [76] F. Xu, *SU(3)-structures and special Lagrangian geometries*, math/0610532.
- [77] F. Xu, *Geometry of SU(3) manifolds*, PhD thesis, Duke University (2008).
- [78] J. Wang and Y. Wang, *On the geometry of tangent bundles with rescaled metric*, 1104.5584 [math.DG].
- [79] A. Gezer and M. Altunbas, *Notes on the rescaled Sasaki type metric on the cotangent bundle*, Acta Math. Sci. **34** (2014) 162.
- [80] L. Woltjer, *A theorem on force-free magnetic fields*, Proc. Nat. Ac. Sci. **44** (1958) 489.
- [81] J. H. C. Whitehead, *An expression of Hopf’s invariant as an integral*, Proc. Nat. Ac. Sci. **33** (1947) 117.
- [82] R. L. Ricca and H. K. Moffatt, *The helicity of a knotted vortex filament*, Top. Asp. Dyn. Fluid. Plas. **218** (1992) 225.
- [83] G. Călugăreanu, *L’intégral de Gauss et l’analyse des noeuds tridimensionnels*, Rev. Math. Pures Appl. **4** (1959) 5.
- [84] G. Călugăreanu, *Sur les classes d’isotopie des noeuds tridimensionnels et leurs invariants*, Czech. Math. J. **11** (1961) 588.
- [85] X. Liu and R. L. Ricca, *The Jones polynomial for fluid knots from helicity*, J. Phys. A **45** (2012).
- [86] K. Ohkitani and J. D. Gibbon, *Numerical study of singularity formulation in a class of Euler and Navier–Stokes flows*, Phys. Fluids **12** (2000) 3181.

- [87] J. D. Gibbon, A. S. Fokas, and C. R. Doering, *Dynamically stretched vortices as solutions of the 3d Navier–Stokes equations*, *Physica D* **132** (1999) 497.
- [88] J. M. Burgers, *A mathematical model illustrating the theory of turbulence*, *Adv. Appl. Mech.* **1** (1948) 171.
- [89] L. W. Tu, *An Introduction to Manifolds*, Springer New York, 2010 [doi].
- [90] K. R. Meyer, *Symmetries and integrals in mechanics*, *Dyn. Sys.* **5** (1973) 259.
- [91] T. S. Lundgren, *Strained spiral vortex model for turbulent fine structure*, *Phys. Fluids* **25** (1982) 2193.
- [92] T. Dombre, U. Frisch, J. M. Greene, M. Hénon, A. Mehr, and A. Soward, *Chaotic streamlines in the ABC flows*, *J. Fluid Mech.* **31** (1986) 353.
- [93] W. M. Hicks, *Researches in vortex motion – part III: on spiral or gyrostatic vortex aggregates*, *Phil. Trans. R. Soc. A* **192** (1899) 33.
- [94] K. Abe, *Existence of vortex rings in Beltrami flows*, *Commun. Math. Phys.* **391** (2022) 873 [2008.09345 [math.AP]].
- [95] K. H. Prendergast, *The equilibrium of a self-gravitating incompressible fluid sphere with a magnetic field*, *Astrophysical J.* **123** (1956) 498.
- [96] M. J. M. Hill, *On a spherical vortex*, *Phil. Trans. R. Soc. A* **185** (1894).
- [97] R. L. Bryant, *Remarks on the geometry of almost complex 6-manifolds*, *Asian J. Math.* **10** (2006) 561 [0508428v2 [math.DG]].
- [98] B. Banos, *On symplectic classification of effective 3-forms and Monge–Ampère equations*, *Diff. Geom. Appl.* **19** (2003) 147 [math-ph/0003026].
- [99] B. Riemann, *Über die hypothesen welche der Geometrie zu grunde liegen*, Springer Spektrum, 2013 [doi].
- [100] S. Alexander, V. Kapovitch, and A. Petrunin, *Alexandrov Geometry: Foundations*, American Mathematical Society (Graduate Studies in Mathematics), 2024 [doi].
- [101] A. D. Alexandrov, *Über eine Verallgemeinerung der Riemannschen Geometrie*, *Schr. Forschungsinst. Math.* **1** (1957) 33.
- [102] V. A. Toponogov, *Riemann spaces with curvature bounded below*, *Uspekhi Mat. Nauk* **14** (1959) 87.
- [103] Y. Burago, M. Gromov, and G. Perelman, *Alexandrov spaces with curvature bounded below*, *Russ. Math. Surv.* **47** (1992) 1.
- [104] M. Gromov, *Synthetic geometry in Riemannian manifolds*, *Proc. Int. Congress Math, Helsinki* **1** (1978) 415.
- [105] M. Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Birkhauser, 1999 [doi].
- [106] V. Kapovitch, *Regularity of limits of non-collapsing sequences of manifolds*, *Geom. Funct. Anal.* **12** (2002) 121 [math/0109075].
- [107] K. Grove and P. Petersen, *Manifolds near the boundary of existence*, *J. Diff. Geom.* **33** (1991) 379.
- [108] N. Lebedeva, *Alexandrov spaces with maximal number of extremal points*, *Geom. Topol.* **19** (2015) 1493 [1111.7253v3].
- [109] D. Burago, Y. Burago, and S. Ivanov, *A Course in Metric Geometry*, American Mathematical Society (Graduate Studies in Mathematics), 2001.

- [110] G. Savaré, *Gradient flows and diffusion semigroups in metric spaces under lower curvature bounds*, C. R. Acad. Sci. Paris I (2007).
- [111] N. Gigli and F. Nobili, *A Differential Perspective on Gradient Flows on $CAT(\kappa)$ -spaces and Applications*, J. Geom. Anal. **31** (2021) 11780 [2012.12952 [math.MG]].
- [112] C. Villani, *Optimal Transport: Old and New*, Springer Verlag, 2009 [doi].
- [113] A. D. Rendall, *Theorems on existence and global dynamics for the Einstein equations*, Living Rev. Relativ. **8** (2005) [gr-qc/0505133].
- [114] J. A. G. Vickers, *Quasi-regular singularities and cosmic strings*, Class. Quantum Grav. **7** (1990) 731.
- [115] J. Podolský, C. Sämann, R. Steinbauer, and R. Švarc, *The global uniqueness and C^1 -regularity of geodesics in expanding impulsive gravitational waves*, Class. Quantum Grav. **33** (2016) 23 [1602.05020 [gr-qc]].
- [116] M. Kunzinger and C. Sämann, *Lorentzian Length Spaces*, Ann. Glob. Anal. Geom. **54** (2018) 399 [1711.08990v4].
- [117] E. H. Kronheimer and R. Penrose, *On the structure of causal spaces*, Proc. Camb. Philos. Soc. **67** (1967) 481.
- [118] S. G. Harris, *A triangle comparison theorem for Lorentz manifolds*, Indiana Univ. Math. J. **31** (1982) 289.
- [119] S. B. Alexander and R. L. Bishop, *Lorentz and semi-Riemannian Spaces with Alexandrov curvature bounds*, Comm. Anal. Geom. **16** (2008) 251 [0804.2493 [math.DG]].
- [120] J. D. E. Grant, M. Kunzinger, and C. Sämann, *Inextendibility of spacetimes and Lorentzian length spaces*, Ann. Global Anal. Geom. **55** (2019) 133 [1804.10423v5].
- [121] S. Alexander, M. Graf, M. Kunzinger, and C. Sämann, *Generalized cones as Lorentzian length spaces: Causality, curvature and singularity theorems*, Comm. Anal. Geom. **31** (2023) 1469 [1909.09575v2].
- [122] B. O'Neill, *Semi-Riemannian geometry with applications to relativity*, Pure and Applied Mathematics, Academic Press, Inc, 1983.
- [123] T. Beran, M. Kunzinger, and F. Rott, *On curvature bounds in Lorentzian length spaces*, J. LMS **110** (2024) 49 [2309.12062 [math.DG]].
- [124] W. Barrera, L. Mondes de Oca, and D. Solis, *Comparison theorems for Lorentzian length spaces with lower timelike curvature bounds*, Gen. Relativ. Gravit. **54** (2022) [2204.09612 [math.DG]].
- [125] T. Beran and C. Sämann, *Hyperbolic angles in Lorentzian length spaces and timelike curvature bounds*, J. LMS **107** (2023) 1823 [2204.09491v3].
- [126] T. Beran and F. Rott, *Gluing constructions for Lorentzian length spaces*, manuscripta math. **173** (2023) 667 [2201.09695 [math.DG]].
- [127] F. Rott, *Gluing of Lorentzian length spaces and the causal ladder*, Class. Quantum Grav. **40** (2023) [2209.06894 [math.DG]].
- [128] T. Beran, A. Ohanyan, F. Rott, and D. A. Solis, *The splitting theorem for globally hyperbolic Lorentzian length spaces with non-negative timelike curvature*, Lett. Math. Phys. **113** (2023) 43 [2209.14724 [math.DG]].
- [129] R. J. McCann and C. Sämann, *A Lorentzian analogue for Hausdorff dimension and measure*, Pure

- and Applied Analysis **4** (2022) 367 [2110.04386v2].
- [130] A. Burtscher and L. García-Heveling, *Time functions on Lorentzian length spaces*, Ann. Henri Poincaré (2024) 39 [2108.02693v2].
- [131] C. Sormani and C. Vega, *Null distance on a spacetime*, Class. Quantum Grav. **33** (2016) 29 [1508.00531v2].
- [132] M. Kunzinger and R. Steinbauer, *Null distance and convergence of Lorentzian length spaces*, Ann. Henri Poincaré **23** (2022) 4319 [2106.05393 [math.DG]].
- [133] U. Lang and V. Schroeder, *On Toponogov's comparison theorem for Alexandrov spaces*, Enseign. Math. **59** (2013) 325 [1207.3668 [math.MG]].
- [134] E. Minguzzi and S. Suhr, *Lorentzian metric spaces and their Gromov–Hausdorff convergence*, Lett. Math. Phys. **114** (2024) 64 [2209.14384v3].
- [135] O. Müller, *Gromov–Hausdorff metrics and dimensions of Lorentzian length spaces*, (To Appear) (2022) 28 [2209.12736v7].
- [136] M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature, Fundamental Principles of Mathematical Sciences*, Springer-Verlag Berlin Heidelberg, 1999.
- [137] U. Lang, *Length Spaces, Lecture Course on Metric Geometry – ETH Zurich* (2013).
- [138] W. Ballmann, *Lectures on Spaces of Nonpositive curvature (Series: OWS)*, Birkhäuser Basel, 1995 [doi].
- [139] P. Pizzetti, *Paragone fra due triangoli a lati uguali*, Rendiconti Lincei. Mat. App. **5** (1907) 6.
- [140] V. Pambuccian and T. Zamfirescu, *Paolo Pizzetti: The forgotten originator of triangle comparison geometry*, Hist. Math. **38** (2011) 415.
- [141] A. D. Alexandrov, *A theorem on triangles in metric space and some of its applications (Russian)*, Trudy Mat. Inst. Steklov Acad Sci USSR **38** (1951) 5.
- [142] C. Plaut, *Almost Riemannian Spaces*, J. Differential Geom. **34** (1991) 515.
- [143] C. Plaut, *Spaces of Wald-Berestovskii curvature bounded below*, J. Geom. Anal. **6** (1996) 113.
- [144] A. Petrunin, *A globalization for non-complete but geodesic spaces*, Math. Ann. **366** (2016) 387 [1208.3155v3].
- [145] S. B. Myers, *Riemannian manifolds in the large*, Duke Math. J. **1** (1935) 39.
- [146] S. B. Myers, *Riemannian manifolds with positive mean curvature*, Duke Math. J. **8** (1941) 401.
- [147] L. Aké Hau, A. J. Cabrera Pacheco, and D. A. Solis, *On the causal hierarchy of Lorentzian length spaces*, Class. Quantum Grav. **37** (2020) 21 [2003.03451 [gr-qc]].
- [148] E. Ling, *Aspects of \mathcal{C}^0 causal theory*, Gen. Relativ. Gravit. **52** (2020) [1911.04438v4].
- [149] T. Beran, M. Kunzinger, A. Ohanyan, and F. Rott, *The equivalence of smooth and synthetic notions of timelike sectional curvature bounds*, (To Appear) (2024) [2403.18077 [math.DG]].
- [150] S. Holst, *Gott time machines in the Anti-de Sitter space*, Gen. Relativ. Gravit. **28** (1996) 387 [gr-qc/9501010].
- [151] E. Minguzzi, *Lorentzian Causality Theory*, Living Rev. Relativ. **22** (2019) 203.
- [152] T. Beran, *Lorentzian length spaces*, Masters' Thesis, University of Vienna (2020).
- [153] U. Moschella. *Einstein, 1905–2005: Poincaré Seminar 2005* chapter The de Sitter and anti-de Sitter Sightseeing Tour, 120. Birkhäuser Basel 2006. [doi].
- [154] D. Christodoulou, *Mathematical Problems of General Relativity I*, European Mathematical Society,

- 2008 [doi].
- [155] J. Cheeger and D. Ebin, *Comparison Theorems in Riemannian geometry*, American Elsevier, 1975.
- [156] P. T. Chruściel and J. D. E. Grant, *On Lorentzian Causality with continuous metrics*, *Class. Quantum Grav.* **29** (2012) 145001 [1111.0400v3].
- [157] J. R. Munkres, *Topology: A first course*, Prentice–Hall, 1974.
- [158] J. K. Beem and P. E. Ehrlich, *Cut points, conjugate points, and Lorentzian comparison theorems*, *Math. Proc. Camb. Phil. Soc.* **86** (1979) 365.
- [159] F. Cavalletti and A. Mondino, *Optimal transport in Lorentzian synthetic spaces, synthetic time-like Ricci curvature lower bounds and applications*, *Camb. J. Math.* **12** (2024) 417 [2004.08934 [math.MG]].
- [160] M. Braun and M. Calisti, *Timelike Ricci bounds for low regularity spacetimes by optimal transport*, *Commun. Contemp. Math.* (2023) [2209.03802v3].
- [161] A. Petrunin, *Alexandrov meets Lott–Villani–Sturm*, *Münst. J. Math.* **4** (2011) 53 [1003.5948v3].
- [162] T. Beran, *Bonnet–Myers rigidity theorem for globally hyperbolic Lorentzian length spaces*, (To Appear) (2024) [2401.17017 [math.DG]].
- [163] T. Yamaguchi, *Homotopy type finiteness theorems for certain precompact families of Riemannian manifolds*, *Proc. Amer. Math. Soc.* **102** (1988) 660.
- [164] K. Grove, P. Petersen, and J. Wu, *Geometric finiteness theorems via controlled topology*, *Invent. Math.* **99** (1990) 205.
- [165] G. Perelman, *Alexandrov spaces with curvatures bounded from below II*, (Unpublished) (1991).
- [166] J. Harvey, *Equivariant Alexandrov geometry and orbifold finiteness*, *J. Geom. Anal.* **26** (2016) 1925 [1401.0531v3].
- [167] D. Jansen, *Notes on Pointed Gromov–Hausdorff convergence*, 1703.09595 [math.MG].
- [168] J. Noldus, *A Lorentzian Gromov–Hausdorff notion of distance*, *Class. Quant. Grav.* **21** (2004) 839 [0308074].
- [169] L. Bombelli, J. Lee, D. Meyer, and R. D. Sorkin, *Space-time as a causal set*, *Phys. Rev. Lett.* **59** (1987) 521.
- [170] F. Dowker, J. Henson, and R. D. Sorkin, *Quantum gravity phenomenology, Lorentz invariance and discreteness*, *Mod. Phys. Lett. A* **19** (2004) 1829 [gr-qc/0311055v3].
- [171] S. Surya, *The causal set approach to quantum gravity*, *Living Rev. Relativ.* **22** (2019) 75 [1903.11544v2].
- [172] S. Delahaies and I. Roulstone, *Hyper-Kähler geometry and semi-geostrophic theory*, *Proc. R. Soc. A* **466** (2010) 2113.
- [173] A. L. Besse, *Einstein manifolds*, Springer Verlag, 1987 [doi].