

Fakultät für Mathematik und Informatik
Ruprecht-Karls-Universität Heidelberg

Master Thesis

**Compact Lorentz 3-folds
with non-compact isometry
groups**

Supervisor:

Prof. Dr. B. Pozzetti

Oskar Riedler

September 1, 2019

Eigenständigkeitserklärung

Hiermit bestätige ich, dass ich die vorliegende Masterarbeit mit dem Titel “Compact Lorentz 3-folds with non-compact isometry groups” selbstständig verfasst und keine anderen als die angegebenen Hilfsmittel benutzt habe.

Heidelberg, den September 1, 2019

Oskar Riedler

Abstract

It is a well known phenomenon that compact Lorentz-manifolds may admit a non-compact isometry group. The most readily accessible example is that of a flat Lorentz torus $\mathbb{T}^n = \mathbb{M}^n/\mathbb{Z}^n$, the quotient of Minkowski space by integer translations, here the isometry group is $O(n-1,1)_{\mathbb{Z}} \ltimes \mathbb{T}^n$. In [Fr 18] C. Frances provides a complete classification of all compact Lorentz 3-folds with non-compact isometry groups. Throughout [Fr 18] the language of Cartan geometries is used, and many of the initial steps, which show that such a manifold must admit many local Killing fields, are valid in this general setting. This thesis reviews this proof and provides an introduction to the theory of Cartan connections.

Zusammenfassung

Die Existenz von kompakten Lorentzmannigfaltigkeiten mit nicht-kompakter Isometriegruppe ist weit bekannt. Als Beispiel lässt sich ein flacher Lorentztorus $\mathbb{T}^n = \mathbb{M}^n/\mathbb{Z}^n$ angeben, welcher also Quotient vom Minkowski-Raum durch die ganzzahligen Translationen zu verstehen ist. In diesem Fall ist die Isometriegruppe $O(n-1,1)_{\mathbb{Z}} \ltimes \mathbb{T}^n$, welche nicht kompakt ist, da $O(n-1,1)_{\mathbb{Z}}$ ein Gitter im nicht-kompakten $O(n-1,1)$ ist. Die Klassifikation aller 3-dimensionalen Lorentzmannigfaltigkeiten mit nicht-kompakter Isometriegruppe wurde von C. Frances in [Fr 18] durchgeführt, der Hauptzweck dieser Arbeit ist es, diesen Beweis nach zu erarbeiten. Ein Merkmal von [Fr 18], ist dass die Sprache der Cartan Geometrie durchgehend verwendet wird und insbesondere die ersten Resultate, bei denen einer Zerlegung der Mannigfaltigkeiten in "Integrabilitätskomponenten" erfolgt, die eine große Menge von Killingfeldern besitzen, sind in diesem allgemeinen Kontext gültig. Der zweite Zweck dieser Arbeit ist dann eine Einführung in den relevanten Aspekten der Cartan Geometrie zu bieten.

Contents

1	Introduction	1
1.1	Compact spaces with non-compact automorphism groups	1
1.2	What is $C_3(\text{Lorentz})$?	3
1.3	Geometric structures?	4
2	Introduction to Cartan geometry	5
2.1	Model spaces	6
2.1.1	Definition of a Model space / Klein geometry	6
2.1.2	Examples	9
2.2	The Maurer-Cartan form	10
2.2.1	Local condition for isomorphism	10
2.2.2	The fundamental theorem of calculus	12
2.3	Building spaces out of model spaces	16
2.3.1	Glueing flat model space pieces together	16
2.3.2	Rolling the model space on an inhomogeneous manifold	20
2.4	Cartan connection	23
2.5	Local automorphisms and Killing fields	26
2.6	The curvature of a Cartan connection	30
3	Examples of Cartan geometries	33
3.1	Parallelisms	33
3.1.1	Geometry of a parallelism	35
3.1.2	Every Cartan geometry as a parallelism	36
3.2	Pseudo-Riemannian geometry	38
3.2.1	Conventions	40
3.2.2	Ehresmann connection	42

3.2.3	Cartan connection	46
3.2.4	Cartan-isomorphisms and isometries	48
3.2.5	Recovering a pseudo-Riemannian geometry from an $O(p, q) \times \mathbb{R}^n/O(p, q)$ Cartan geometry	50
3.2.6	Equivalence	55
4	Tool: Generalised curvature and the integrability locus	59
4.1	Covariant derivative	60
4.2	Generalised curvature	61
4.3	Integrability theorem	62
4.4	$M^{\text{int}} \subseteq \mathfrak{M}$	63
5	Classification Theorem: Closed Lorentz 3-folds with non-compact isometry group	65
5.1	The classification	65
5.2	Recursion leads to Killing fields	67
5.2.1	Isotropy algebra	68
5.2.2	Recurrence argument	69
5.2.3	$\text{Aut}(M)$ acts properly on \widehat{M}	76
5.3	Classification of components of \widehat{M}^{int} by their local Killing algebras	79
5.4	Overview of the proof of the main theorem	82
5.4.1	The locally homogenous case	82
5.4.2	The hyperbolic case	84
6	Example: Hyperbolic torus bundles	89
6.1	Calculating the relevant quantities	89
6.1.1	Principal bundle and Cartan connection on local charts	90
6.1.2	Local Killing fields	94
6.2	Hyperbolic torus bundles and the torus	96
6.2.1	Isometry group	97
6.2.2	Orthonormal basis and derived quantities	97
6.2.3	Cartan connection	98
6.2.4	Local Killing fields	98
	References	102

Chapter 1

Introduction

1.1 Compact spaces with non-compact automorphism groups

A standard result from functional analysis states that if (X, d) is a compact metric space, that the group of isometries $\text{Isom}(X, d)$ is also compact if endowed with the topology of uniform convergence. This follows from the Arzelá-Ascoli theorem. As a consequence it is impossible for a compact Riemannian manifold to have a non-compact isometry group.

This statement however is very particular to Riemannian geometry. For instance if one slightly generalises the word *isometry* to *conformal diffeomorphism* or changes the metric from a *Riemannian* metric to a *Lorentzian* metric, then there are counter examples:

1. Consider S^2 with the standard Riemannian metric. Then the conformal diffeomorphisms of S^2 are the Möbius transformations $PSL_2(\mathbb{C}) \cong O(3, 1)^0$. This group is not compact.
2. Consider \mathbb{T}^2 with the Lorentz-metric induced by $du dv$ where u, v are the eigenvectors of a hyperbolic element $A \in SL_2(\mathbb{Z})$. Then up to finite index the isometry group of \mathbb{T}^2 is $\mathbb{Z} \times_A \mathbb{T}^2$, where $(n, (a, b))$ acts a point (x, y) as $A^n(x, y) + (a, b)$. This group is not compact.

We remark however that such spaces are very special. For a generic choice of Riemannian / Lorentzian metric g on a given manifold M , the group of isometries or conformal diffeomorphisms is trivial. Spaces with “large” symmetry groups are the exception rather than the norm. One of the most basic ways to quantify the word “large” is the property of non-compactness. If the original manifold was already compact then non-compactness of the symmetry group is even more special, now the manifold itself is “small” and the symmetries cannot be non-compact simply

by pushing a point away to infinity. For example the question whether or not a given compact manifold admits a metric with a non-compact group of conformal symmetries is not trivial.

This question exists in a more general setting. The choice of looking at isometries or conformal diffeomorphisms of Riemannian or Lorentzian metrics is, while not entirely unmotivated, rather arbitrary. If one would put the question in a more general framework, then one could expect it to look like this:

Question 1. *Given a type of geometric structure K and a compact manifold M . Does M admit a geometric structure φ of type K so that the group of symmetries $\text{Aut}(M, \varphi) \subset \text{Diffeo}(M)$ is non-compact?*

At this point the words “geometric structure” still remain unmotivated and undefined. For the time being we will simply use this word as crutch in order for the groups $\text{Aut}(M, \varphi)$ to have meaning in a more general context. We already know the answer to the question in the case of Riemannian geometry, here the answer is: “No, never”.

In the case of Lorentz geometry the question admits some trivial answers. For example we know that S^2 admits no Lorentz metrics, in particular no Lorentz metrics with non-compact isometry groups.¹ For other manifolds it is more difficult.

For conformal geometry the answer is available in the literature. Here a theorem classifies all compact Riemannian manifolds with a non-compact group of conformal diffeomorphisms. The classification is somewhat surprising:

Theorem 1.1.1 (Obata & Lelong-Ferrand, [Ob 70] & [Le 71]). *Let (M, g) be a connected compact Riemannian manifold of dimension n , if $\text{Conf}(M, g)$ is non-compact then M is conformally equivalent to the euclidean sphere S^n .*

Thus the only connected manifolds for which the answer to Question 1 is positive is S^n . If one formulates it as a classification problem, then the condition of connectedness makes sense, as if $(M, [g])$ has non-compact automorphism group surely so will $M \sqcup X$ for any other manifold X . We reformulate Question 1 as a classification problem:

Question 2. *Given a type of geometric structure K , what connected and compact manifolds M of dimension n admit a geometric structure φ of type K so that the group of symmetries $\text{Aut}(M, \varphi)$ is non-compact?*

Tabulating our knowledge so far, let $C_n(K)$ denote, up to isomorphism, all connected compact manifolds of dimension n and a structure of type K for which the automorphism group is

¹The relevant statement is that a compact manifold admits a Lorentz metric if and only if its Euler-characteristic vanishes.

non-compact:

1. In the case of Riemannian geometry $C_n(\text{Riem. geom.}) = \emptyset$.
2. In the case of conformal geometry over Riemannian spaces $C_n(\text{Conf. Riem.}) = \{(S^n, [g_{\text{euc}}])\}$.

So in the two cases where we know the answer, there are very few such manifolds, underlining again the special nature of such examples. Due to the fact that such manifolds seem to be so exceptional, one finds in the literature a vague conjecture:

Vague general conjecture (D’Ambra & Gromov, [DG 90]). All triples $(M, \varphi, \text{Aut}(M, \varphi))$ where M is compact and $\text{Aut}(M, \varphi)$ is “sufficiently large” (e.g. non-compact) are almost classifiable.

In the case of Lorentz geometry we state here, without proof or motivation, that $C_2(\text{Lorentz}) = \{(\mathbb{T}^2, \lambda dudv) \mid \lambda \in \mathbb{R}, u, v \text{ eigenvectors of a hyperbolic } A \in SL_2(\mathbb{R})\}$. In particular one has that topologically \mathbb{T}^2 is the only 2-dimensional connected compact manifold admitting Lorentz-geometries with non-compact symmetries. The next step would be calculating $C_3(\text{Lorentz})$. This was carried out by C. Frances in [Fr 18], the main purpose of this thesis is to review this proof.

1.2 What is $C_3(\text{Lorentz})$?

Here we briefly summarise the results, beginning with the topological classification.

Theorem 1.2.1 (Topological classification). *Let (M, g) be a smooth, oriented and time-oriented closed 3-dimensional Lorentz manifold. If $\text{Isom}(M, g)$ is not compact, then M is diffeomorphic to:*

1. A quotient $\Gamma \backslash \widetilde{SL}_2(\mathbb{R})$, where $\Gamma \subset \widetilde{SL}_2(\mathbb{R})$ is any uniform lattice.
2. The 3-torus \mathbb{T}^3 or a torus bundle \mathbb{T}_A^3 where $A \in SL_2(\mathbb{Z})$ can be any hyperbolic or parabolic element.

The assumptions of oriented and time-oriented are of convenience in the formulation. Any Lorentz-manifold admits a cover of order at most 4 that is time-oriented and oriented, hence if the manifold has non-compact isometry group it is also covered by a manifold with non-compact isometry group. One may also describe the form of the metric in more detail:

Theorem 1.2.2. *Let (M, g) be a smooth, oriented and time-oriented closed 3-dimensional Lorentz manifold and suppose $\text{Isom}(M, g)$ is not compact. Let $(\widetilde{M}, \widetilde{g})$ denote the universal cover of (M, g) , then:*

1. If $M \cong \Gamma \backslash \widetilde{SL}_2(\mathbb{R})$, then \widetilde{g} is a Lorentzian, non-Riemannian, left-invariant metric on $\widetilde{SL}_2(\mathbb{R})$.

2. If $M \cong \mathbb{T}^3$ or \mathbb{T}_A^3 , then there exists a 1-periodic function $a : \mathbb{R} \rightarrow (0, \infty)$ so that \tilde{g} is isometric to one of the following metrics on \mathbb{R}^3 :

$$dt^2 + 2a(t)dxdy \quad \text{or} \quad a(x)(dt^2 + 2dxdy).$$

If M is a hyperbolic torus bundle, only the first case can occur, if M is a parabolic torus bundle, only the second case can occur. If \tilde{g} is locally homogenous (meaning a is constant), then it is flat (which can occur in all cases) or modelled on Lorentz-Heisenberg geometry (which can only occur in the case of a parabolic torus bundle).

Chapter 5 will describe the proof of these results. The initial step of this classification result makes use of a special decomposition of M into open sub-manifolds \mathcal{M}_i with $\bigcup_i \mathcal{M}_i$ being dense in M . One is able to show, in the context of a general “geometry”, that non-compactness of $\text{Aut}(M)$ and compactness of M imply that these components must admit a large number of local Killing fields (cf. Lemma 5.2.11). By analysing what kinds of local Killing algebras may appear on these components one achieves an initial classification of such manifolds, which may be leveraged, using now very concrete aspects of Lorentz geometry, into the complete classification.

1.3 Geometric structures?

In the formulation of Questions 1 and 2 (which are of course the same question) reference is made to “geometric structures”. Perhaps the most canonical way to understand this is in terms of Gromov’s rigid geometric structures, as introduced in [Gr 88]. Riemannian geometry, linear connections, conformal geometry may all be regarded as types of geometric structures in this definition. However the proper formulation of this requires the language of higher order jet bundles, which is often described as “unpleasant” (see e.g. [DG 90] Remark 0.7).

Another, less general, notion that can describe a large amount of geometric theories is that of Cartan geometry. Since our main reference, [Fr 18], makes liberal use of general results from this theory, we will choose it as the context in which Questions 1 and 2 are formulated. The secondary purpose of this thesis is to provide an introduction to this theory, mainly since this theory was not known to the author of the thesis prior to reading [Fr 18]. Chapter 2 motivates the form of this theory and describes some basic results. Chapter 3 then describes how this theory contains pseudo-Riemannian geometry.

In this context a moral extension of the classification result for Riemannian geometry is readily achieved: If M is a compact manifold then any Cartan geometry modelled on G/P where the “isotropy group” P is compact has a compact automorphism group (cf. Corollary 5.2.17).

Chapter 2

Introduction to Cartan geometry

The aim of this chapter is to provide an introduction to Cartan geometry. Cartan geometry is a setting in which one can “do geometry”. As such this chapter starts by trying to develop a story in which the resulting definition of a Cartan connection in Section 2.4 becomes natural. After these motivating sections a discussion of the structure that local automorphisms have in this theory follows.

A common description of Cartan geometry is that it provides a way of encoding what it means for a space to “infinitesimally” look like a certain homogenous model space.

As an analogy we refer to the relation between a Riemannian geometry and Euclidean space or the relation between Lorentz geometry and Minkowski space. A Riemannian manifold is a space that at each point looks infinitesimally like Euclidean space, concretely the metric is a structure making the individual tangent spaces isometric to Euclidean space. A Lorentz manifold is in this perspective exactly the same thing as a Riemannian manifold, except that instead of every point being infinitesimally isomorphic to Euclidean space, the points are isomorphic to Minkowski space.

In a sense one may view Cartan geometry as a generalisation of this setting, replacing the model Euclidean or Minkowski space with a more general class of homogenous model spaces. So while Riemannian geometry provides a way to describe inhomogenous or “lumpy” versions of Euclidean space, the Cartan geometry of a model space X provides a way of describing inhomogenous or lumpy versions of X .

The following diagram (stolen from [Sh 97]) illustrates these ideas:

$$\begin{array}{ccc}
\text{Euclidean space} & \xrightarrow{\text{inhomogenous}} & \text{Riemannian geometry} \\
\downarrow \text{generalisation} & & \downarrow \text{generalisation} \\
\text{Klein model spaces} & \xrightarrow{\text{inhomogenous}} & \text{Cartan geometry}
\end{array}$$

The arrow “generalisation” between Riemannian geometry and Cartan geometry is dashed, because while Cartan geometry modelled Euclidean space does describe Riemannian geometry, it is not a priori clear that it does this. The details showing the equivalence between Cartan geometries modelled on Euclidean space and Riemannian geometry are elaborated in Chapter 3.

The organisation of this chapter is as follows: In Section 2.1 the notion of a model space is defined and some classical examples are given, then in Section 2.2 a set of tools are developed in order to work with these spaces. Section 2.3 then examines how one can build spaces that locally look like a certain model space. Subsection 2.3.1 does this by glueing pieces of the original model space together along automorphism and finding an infinitesimal characterisation of such glued spaces. By dropping the integrability condition of this infinitesimal characterisation we obtain a different kind of space, here Subsection 2.3.2 motivates that this space is still infinitesimally the model space, but that it might now be lumpy. Section 2.4 then takes this as the definition of a Cartan connection and briefly remarks on the definition of automorphisms of such a structure, while Section 2.5 investigates the automorphism group and local Killing algebras in more detail. Section 2.6 defines the curvature of a Cartan connection and describes some of its most elementary properties.

2.1 Model spaces

2.1.1 Definition of a Model space / Klein geometry

Definition 2.1.1 (Model space). A pair $(M, \text{Aut}(M))$, where M is a connected manifold and $\text{Aut}(M) \subset \text{Diffeo}(M)$ is a subgroup of the diffeomorphism group of M , is called a model space if:

1. $\text{Aut}(M)$ is a Lie group.
2. $\text{Aut}(M)$ acts transitively on M .
3. $\text{Aut}(M)$ is finite dimensional as a Lie group.

$\text{Diffeo}(M)$ is given the open-compact topology and $\text{Aut}(M)$ the subspace topology of $\text{Diffeo}(M)$ in order for the first and third conditions to make sense. The first condition is one of necessity. $\text{Aut}(M)$ needs some kind of structure if we are to work with it in a manageable way. The second

is responsible for making the model space homogenous with respect to the symmetry group $\text{Aut}(M)$, while the third ensures that the group of symmetries are that of a “geometric” structure. This last comment bears some elaboration:

A common feature of many geometric theories is that the conditions for a map to be a symmetry are “rigid”, as opposed to “soft”. What this means is that given a symmetry, one has very little freedom to deform it without leaving the space of symmetries. For example a symmetry in Riemannian geometry, that is an isometry, is uniquely determined by its derivative at a point (provided the domain is connected), thus one has no freedom to deform an isometry if one wishes to keep its derivative at a point fixed.

Often the conditions for a map to be a symmetry can be formulated as a complete PDE, meaning that there exists an N so that the partial derivatives of order $N + 1$ of a solution may be expressed by the partial derivatives of order $\leq N$. In such a setting, knowledge of the first N derivatives at any point uniquely determine the symmetry (if the domain is connected), this follows for example by considering a curve γ and denoting with $J_f^N(t)$ all partial derivatives of $\leq N$ of f along γ . Then $\frac{d}{dt}J_f^N(t) = F(J_f^N(t), \dot{\gamma}(t))$ for a function F , if F is regular enough then $J_f^N(t)$ is uniquely determined by $J_f^N(0)$. If a symmetry is uniquely determined by the first N of derivatives at any point, then $\text{Aut}(M)$ will embed into a space that locally looks like $M \times \bigoplus_{k=0}^N (\mathbb{R}^n \otimes (\mathbb{R}^n)^*)^{\otimes k}$, and must thus be finite dimensional. Doing this formally would work by constructing the jet bundle of order N of M , on which $\text{Aut}(M)$ will act freely.

From this perspective the third condition is a way of ensuring that $\text{Aut}(M)$ is the symmetry group of a geometric structure on M . A slogan expressing this point of view could be: “If the symmetry group is infinite dimensional, then it’s topology and not geometry.”

Given a model space $(M, \text{Aut}(M))$ we consider, for any point $x \in M$, the stabiliser subgroup $\text{Stab}_{\text{Aut}(M)}(x) = \{ f \in \text{Aut}(M) \mid f(x) = x \}$. From the transitivity of the action of $\text{Aut}(M)$ on M , it is easy to see that for any $x, x' \in M$ there exists a $f \in \text{Aut}(M)$ so that

$$f \cdot \text{Stab}_{\text{Aut}(M)}(x) \cdot f^{-1} = \text{Stab}_{\text{Aut}(M)}(x'),$$

i.e. all stabiliser subgroups are conjugate to each other. This follows since if $f(x) = x'$ and $g(x) = x$ then $fg(f^{-1}x') = fg(x) = f(x) = x'$.

It is also elementary that $\text{Stab}_{\text{Aut}(M)}(x)$ is a closed subgroup of $\text{Aut}(M)$, as the topology of pointwise convergence is coarser than the compact-open topology (which for manifolds is the topology of uniform convergence on compact sets), and thus if $f_n \rightarrow f$ with $f_n(x) = x$ then $f(x) = x$. As such $\text{Stab}_{\text{Aut}(M)}(x)$ is a Lie subgroup of $\text{Aut}(M)$.

Proposition 2.1.2. *The coset space $\text{Aut}(M)/\text{Stab}_{\text{Aut}(M)}(x)$ is a manifold and is diffeomorphic to M .*

Proof. $\text{Aut}(M)/\text{Stab}_{\text{Aut}(M)}(x)$ is a quotient of a Lie group by a closed subgroup and thus a manifold by a general theorem of Lie groups (often called the quotient manifold theorem, see for example [Lee 13] Theorem 21.10). In fact $\text{Aut}(M)$ is a (right) $\text{Stab}_{\text{Aut}(M)}(x)$ -principal bundle over $\text{Aut}(M)/\text{Stab}_{\text{Aut}(M)}(x)$: The action $\text{Aut}(M) \times \text{Stab}_{\text{Aut}(M)}(x) \rightarrow \text{Aut}(M), (f, g) \mapsto f \cdot g$ is clearly a bundle map as well as smooth and free. Further it is transitive on each fibre, as these are always of the form $f \cdot \text{Stab}_{\text{Aut}(M)}(x)$ for some f . That form also makes it clear that the action of $\text{Stab}_{\text{Aut}(M)}(x)$ on any fibre is a diffeomorphism.

We consider the evaluation map $ev_x : \text{Aut}(M) \rightarrow M, f \mapsto f(x)$ which is smooth, as it is the restriction of the evaluation map on $\text{Diffeo}(M)$ which is already smooth. It is easily verified that this map only depends on the class of f in $\text{Aut}(M)/\text{Stab}_{\text{Aut}(M)}(x)$, for if $g \in \text{Stab}_{\text{Aut}(M)}(x)$ then $fg(x) = f(g(x)) = f(x)$. As such this map is actually a lift of the map $\text{Stab}_{\text{Aut}(M)}(x) \rightarrow M, [f] \mapsto f(x)$, in particular this second map is smooth.

Further it is surjective by transitivity of $\text{Aut}(M)$ and it is injective, for if $f(x) = g(x)$ then $g^{-1}f(x) = x$ and $g^{-1}f \in \text{Stab}_{\text{Aut}(M)}(x)$ so $[g] = [gg^{-1}f] = [f]$.

Lastly the differential is has no kernel. If γ_t were a path in $\text{Aut}(M)/\text{Stab}_{\text{Aut}(M)}(x)$ with $\dot{\gamma}_0 \neq 0$ but $D_{\gamma_0}ev_x(\dot{\gamma}_0) = 0$ then for any $f \in \text{Aut}(M)$ one must have that

$$\frac{d}{dt}ev_x(f \circ \gamma_t)|_{t=0} = \frac{d}{dt}f(\gamma_t(x))|_{t=0} = D_{\gamma_0(x)}f(D_{\gamma_0}ev_x(\dot{\gamma}_0)) = D_{\gamma_0}f(0) = 0.$$

Hence the map Dev_x must have non-zero kernel at every point. It is impossible for a bijective smooth map to have non-zero kernel at every point however, as can be seen by choosing a smooth vector field lying inside the kernel. The map ev_x must be invariant under the flow of that field, contradicting injectivity. So ev_x is a bijective immersion, which is a diffeomorphism. \square

This proposition instantly provides us with an extremely useful characterisation of the model spaces: they are coset spaces.

Corollary 2.1.3. *Given a Lie group G and a closed subgroup P denote with N the largest subgroup of P that is normal in G . Then $(G/N, G/P)$ is a model space.*

Proof. Note the action $G \times G/P, (f, [g]) \mapsto [fg]$ is well defined and transitive. However it is not effective, for if $n \in N$ and $g \in G$ there exists an $n' \in N$ so that $ng = gn'$, whence $[ng] = [g]$. This implies that the action of G/N on G/P is well defined. Further it is clear that this action is effective. \square

For the rest of the thesis we will use the following definition of a model space, which by the proposition and the corollary before is the same as the original one:

Definition 2.1.4 (Model space). A model space is a pair (G, P) where G is a finite-dimensional Lie group and $P \subseteq G$ is a closed subgroup so that G/P is connected and P admits no non-trivial subgroups that are normal in G . We will often denote the model space with G/P .

One could relax this definition by allowing a non-trivial normal (in G) subgroup N of P . In this case any pair (G, P) , $P \subseteq G$ closed would satisfy the definition. In the literature such a pair is called a Klein geometry, in the event $N = 1$ the Klein geometry is called effective and we recover our definition of a model space.

Asking for the model space to be effective will in our case be responsible for the fact that the automorphism group of a space M modelled on G/P acts effectively on M , i.e. an automorphism of M is uniquely determined by its action of M . In certain contexts, for example that of spin structures, it is useful to relax this condition.

2.1.2 Examples

Any pair (G, P) with $P \subseteq G$ provides an example of a Klein geometry, even with the restriction that the geometry be effective one can without difficulty write down any number of such pairs without there being any clue as to what these pairs have to do with geometry. For this reason we list here some classical model spaces, whose geometric character is clear.

1. **Euclidean space.** For $n \in \mathbb{N}$ consider Euclidean space \mathbb{R}^n , its automorphism group consists of all isometries of \mathbb{R}^n , namely $O(n) \ltimes \mathbb{R}^n$. This describes the model space (G, P) with $G = O(n) \ltimes \mathbb{R}^n$, $P = O(n)$.
2. **Minkowski space.** In the same way taking n -dimensional Minkowski space is described by the model space $(G, P) = (O(n-1, 1) \ltimes \mathbb{R}^n, O(n-1, 1))$. This may be further generalised to the flat model space of pseudo-Riemannian geometry of signature (p, q) , where $(G, P) = (O(p, q) \ltimes \mathbb{R}^n, O(p, q))$.
3. **Hyperbolic space.** A common description of hyperbolic space is the hyperboloid

$$\mathcal{H}_n = \{x \in \mathbb{R}^{n+1} \mid x_0^2 - \sum_{i=1}^n x_i^2 = -1, x_0 > 0\},$$

which is a subspace of $n + 1$ -dimensional Minkowski space. The symmetries of this set are $SO^0(1, n + 1)$ whereas the stabiliser of the point $(1, 0, \dots, 0)$ is $SO(n)$ as lying in

$SO(1) \times SO(n) \subset SO^0(1, n+1)$. So $\mathcal{H}_n = SO^0(1, n+1)/O(n)$. Other descriptions like the half plane model are described by the same (G, P) pair.

4. **The conformal sphere.** Consider $S^n \subset \mathbb{R}^{n+1}$ with the standard Euclidean metric. The group of conformal diffeomorphisms of S^n is isomorphic to $PO(n+1, 1)$, the stabiliser P of a point is then isomorphic to a subgroup of $PO(n+1, 1)$ preserving a given null ray, under this identification. $(PO(n+1, 1), P)$ then acts as the model space for the conformal sphere.
5. **Projective space.** The transformations of $\mathbb{R}P^n$ preserving projective lines and cross ratios is $PGL(n+1)$, the subgroup P preserving a point is then given by those transformations preserving a point, which we may view as the image of $\left\{ \begin{pmatrix} 1 & v^T \\ 0 & A \end{pmatrix} \middle| v \in \mathbb{R}^n, A \in GL(n) \right\}$ in $PGL(n+1)$.

2.2 The Maurer-Cartan form

We briefly pause to develop some technical notions. The Maurer-Cartan form defined in Definition 2.2.2 is an important notion and will inform later sections as well. Aside from that we find a condition that expresses, locally, when map is an automorphism of G/P (Lemma 2.2.4) as well as derive a fundamental theorem of calculus (Theorem 2.2.10) formulated for Lie group valued functions. These two statements are used in Theorem 2.3.6, which explains why a flat Cartan connection is the same as a local G/P structure.

2.2.1 Local condition for isomorphism

We already know what the automorphisms of the model space G/P look like, by definition they are G acting by left translations on this quotient. However before continuing we conduct a short investigation into these automorphisms. Our goal is to find a condition that tells us when a map $G/P \rightarrow G/P$ is locally an automorphism - meaning for any point in G/P there is a neighbourhood on which the map equals the restriction of an automorphism (that is an element of $G = \text{Aut}(G/P)$) acting via left-multiplication on G/P). To this end we first make the following remark:

Proposition 2.2.1. *Let G/P be a model space. Then any automorphism $[L_g] : G/P \rightarrow G/P$, $[p] \mapsto [gp]$ has a lift as a left multiplication $L_g : G \rightarrow G$, this lift is a bundle-automorphism of the P -principal bundle $G \rightarrow G/P$. This is the only left multiplication lifting $[L_g]$.*

Proof. If $\pi : G \rightarrow G/P$ is the quotient map and we denote with $L_g : G \rightarrow G$ the map $f \mapsto gf$ we

clearly have $[L_g] \circ \pi = \pi \circ L_g$, whence L_g is a bundle map over $[L_g]$. Since L_g has an inverse, namely $L_{g^{-1}}$, this map is a bundle-automorphism.

By effectiveness of the model G/P the action of G on G/P is effective, and then for any other $g' \in G$ the induced map $[L_{g'}]$ differs from $[L_g]$, so $L_{g'}$ cannot be a lift of $[L_g]$. \square

Remark. There are other bundle-automorphisms lifting $[L_g]$. For example if the centre $Z(P)$ of P is not trivial and $\gamma : G \rightarrow Z(P)$ is smooth path, then the map $G \rightarrow G$, $h \mapsto g \cdot h \cdot \gamma(h)$ is a bundle-automorphism lifting $[L_g]$. The uniqueness part of the proposition asserts only that there are no other left multiplications lifting $[L_g]$.

Proposition 2.2.1 shows that any automorphism of G/P must lift to an automorphism of the bundle $G \rightarrow G/P$. But there are many bundle-automorphisms, not all of them can be lifts of automorphisms of G/P . For example in the case where P is the trivial group $\{0\}$ the bundle-automorphisms of G/P are all of $\text{Diffeo}(G)$. The condition that a map lifts to a bundle-automorphism is then not really a definite tool in classifying automorphisms. In order to find the correct local condition we introduce the Maurer Cartan form:

Definition 2.2.2. Let G be a Lie group and for $g \in G$ denote with $L_g : G \rightarrow G$ the left-multiplication $f \mapsto gf$. The Maurer-Cartan form is defined as the $\text{Lie}(G) = T_1G$ valued 1-form:

$$\omega_{MC} : TG \rightarrow T_1G, \quad v \in T_gG \mapsto D_g L_{g^{-1}} [v].$$

Remark. The Maurer-Cartan form is smooth and provides the usual way in which the tangent spaces of a Lie group are identified, namely just by left translating these to the tangent space of identity. That makes sense because $\omega_{MC}|_{T_gG} : T_gG \rightarrow T_1G$ is an isomorphism of vector spaces for any $g \in G$. The details of showing smoothness and being an isomorphism on each tangent space are carried out in Proposition 3.1.2.

Proposition 2.2.3. Let $[L_g] : G/P \rightarrow G/P$, $[f] \mapsto [gf]$ be an automorphism and $L_g : G \rightarrow G$, $f \mapsto gf$ its lift. Then $L_g^*(\omega_{MC}) = \omega_{MC}$.

Proof. The equation $L_g^*(\omega_{MC}) = \omega_{MC}$ means $(\omega_{MC})_{L_g(h)} \circ D_h L_g = (\omega_{MC})_h$ for every $h \in G$. But by unpacking definitions

$$(\omega_{MC})_{L_g(h)} \circ D_h L_g = D_{gh} L_{(gh)^{-1}} \circ D_h L_g = D_h L_{(gh)^{-1}g} = D_h L_{h^{-1}} = (\omega_{MC})_h,$$

which is the desired relation. \square

The Maurer-Cartan form of G is then always preserved by the action of $\text{Aut}(G/P)$ on G . This is now a real condition, there are “far less” diffeomorphisms preserving ω_{MC} than there are

bundle-automorphisms (in fact the diffeomorphisms preserving ω_{MC} form a finite-dimensional subgroup of $\text{Diffeo}(G)$).

This condition is the condition we were looking for, if a map admits a lift to a bundle-isomorphism that preserves ω_{MC} then this map is locally an automorphism:

Lemma 2.2.4. *Let $U, V \subset G/P$ be connected and $F : \pi^{-1}(U) \rightarrow \pi^{-1}(V)$ a bundle-isomorphism so that $F^*(\omega_{MC}|_{\pi^{-1}(V)}) = \omega_{MC}|_{\pi^{-1}(U)}$. Then there exists a $g \in G$ so that $F = L_g|_U$.*

Proof. Let X be a constant vector field on G , that is there is $v \in \mathfrak{g}$ so that $X_x = (\omega_{MC})_x^{-1}(v)$ for all $x \in G$. Then if $x \in U$ we see $D_x F((\omega_{MC})_x^{-1}(v)) = (\omega_{MC})_{F(x)}^{-1}(v)$ and F preserves the constant fields. In particular, if $\phi_X^t(x)$ is the flow of X at x then $F(\phi_X^t(x)) = \phi_X^t(F(x))$, provided both $\phi_X^t(x)$ and x lie in U .

But $\phi_X^t(x) = x \cdot e^{tX_1}$, since the ω_{MC} -constant fields are the left-invariant fields. Then we have just seen $F(x \cdot e^{tX_1}) = F(x)e^{tX_1}$, provided both x and e^{tX_1} lie in U . In particular $F(x \cdot e^{tX_1}) = F(x)x^{-1}(xe^{tX_1})$ for all t in a small neighbourhood of 0 (chosen so that xe^{tX_1} remains in U), and thus $F(y) = F(x)x^{-1} \cdot y$ for all y in a small neighbourhood of x .

Since x was arbitrary we find that every point in U admits a neighbourhood $U_x \subseteq U$ on which F corresponds to a left-multiplication with an element g_x of G ($g_x = F(x)x^{-1}$). If $x, y \in U$ and $U_x \cap U_y \neq \emptyset$ then there must be a point $p \in G$ for which $g_x p = F(p) = g_y p$, from which $g_x = g_y$. By connectedness of U it follows that there exists a $g \in G$ with $F(x) = gx$ for all $x \in G$. \square

2.2.2 The fundamental theorem of calculus

In this section we describe the role the Maurer-Cartan form plays in formulating a ‘‘Fundamental Theorem of Calculus’’ for smooth Lie group valued functions. For the formulation of this theorem a new notion of derivative is introduced, which is more in line with notion of derivative as used in analysis.

In analysis the differential of a smooth map $f : \mathbb{R} \rightarrow \mathbb{R}$ is again a smooth map $f' : \mathbb{R} \rightarrow \mathbb{R}$. In the context of differential geometry however the most elementary way to view the differential of a map between manifolds $M \rightarrow N$ is by considering the induced map $TM \rightarrow TN$ between the tangent bundles. In the case $M = \mathbb{R} = N$ this is the map $Df : T\mathbb{R} \rightarrow T\mathbb{R}$. The difference between these f' and Df is easily described:

In \mathbb{R} all tangent spaces are canonically isomorphic (indeed, in most analysis courses one even considers all of them to be the same space), so there is an identification $T\mathbb{R} \cong \mathbb{R} \times \mathbb{R}$, under this

identification one has $Df(x, v) = (f(x), f'(x) \cdot v)$. Thus looking at f' corresponds to looking at Df and forgetting the base point $f(x)$ given by the original map.

This can be generalised to maps $M \rightarrow G$ valued in a Lie group, in this setting the Maurer-Cartan form provides a canonical way of identifying all the tangent spaces of G . Post-composing Df with the Maurer-Cartan form ω_{MC} will then retrieve a map $TM \rightarrow T_1G$ that “forgets” the underlying map f and keeps only the tangential information. We call this notion of derivative the Darboux derivative:

Definition 2.2.5 (Darboux derivative). Let G be a Lie group, $\mathfrak{g} = T_1G$ its Lie algebra and M a manifold. For a smooth function $f : M \rightarrow G$ define the Darboux derivative as the \mathfrak{g} -valued 1-form $\omega_f := f^*(\omega_{MC})$.

Note that for $x \in M$ one has $(\omega_f)_x = (\omega_{MC})_{f(x)} \circ D_x f = D_{f(x)} L_{f(x)^{-1}} \circ D_x f$, i.e. ω_f is precisely the map that forgets the basepoint of Df and retains only the tangency data.

Proposition 2.2.6. *Let M be a connected manifold and G a Lie group. If $f, g : M \rightarrow G$ are two functions so that $\omega_f = \omega_g$, then there exists an $h \in G$ so that $f = L_h \circ g$.*

Proof. Consider the map $fg^{-1} : M \rightarrow G, x \mapsto f(x)g^{-1}(x)$. Then

$$[(fg^{-1})^*(\omega_{MC})]_x = (\omega_{MC})_{fg^{-1}(x)} \circ D_x(fg^{-1}),$$

however $fg^{-1} = \mu \circ (1, \iota) \circ (f, g) \circ \Delta$, with $\mu : G \times G \rightarrow G$ the multiplication, $\iota : G \rightarrow G$ the inversion and $\Delta : G \rightarrow G \times G$ the diagonal map. It follows for any $v \in T_x M$ that

$$\begin{aligned} D_x(fg^{-1})[v] &= (D_{f(x)} R_{g^{-1}(x)} \circ D_x f)[v] + (D_{g(x)} L_{f(x)} \circ D_{g(x)} \iota \circ D_x g)[v] \\ &= (D_{f(x)} R_{g^{-1}(x)} \circ D_x f)[v] + (D_{g(x)} L_{f(x)} \circ -(D_1 L_{g(x)^{-1}} D_{g(x)} R_{g(x)^{-1}}) \circ D_x g)[v]. \end{aligned}$$

Post composing with $(\omega_{MC})_{fg^{-1}(x)}$ and using $(\omega_{MC})_{kh} \circ D_h L_k = (\omega_{MC})_h$ as well as $(\omega_{MC})_{hk} \circ D_h R_k = \text{Ad}(h^{-1})(\omega_{MC})_h$ (where $\text{Ad}(h)$ is the derivative of $k \mapsto hkh^{-1}$ at 1) one gets:

$$(\omega_{fg^{-1}})_x[v] = (\text{Ad}(g(x))\omega_f)_x[v] - (\text{Ad}(g(x))\omega_g)_x[v] = 0.$$

It follows that $D_x fg^{-1} = 0$ for any point x , whence by connectedness of M the map fg^{-1} must be constant. There must then be an $h \in G$ so that $f(x) = hg(x)$ for all $x \in M$. \square

Thus the form of a derivative of a map $M \rightarrow G$ is here always a \mathfrak{g} -valued 1-form on M , and the derivative uniquely determines the function up to a constant. A natural question is now, given such a form $\omega : TM \rightarrow \mathfrak{g}$, when does there exist an anti-derivative of ω , that is a function f with $\omega_f = \omega$?

To begin we show an additional property of ω_{MC} , this property will later express the fact that a model space is flat, or has “zero curvature”. At this point however it will provide us with a necessary condition for ω to have an anti-derivative (compare with the upcoming Corollary 2.2.9). We make a short definition in order to be able to formulate the property:

Definition 2.2.7. Let \mathfrak{g} be a Lie algebra with bracket $[\cdot, \cdot]_{\mathfrak{g}}$, M a manifold and $\eta, \sigma : TM \rightarrow \mathfrak{g}$ two \mathfrak{g} -valued 1-forms. Then the bracket $[\eta, \sigma]_{\mathfrak{g}}$ is defined as the \mathfrak{g} -valued 2-form given by

$$([\eta, \sigma]_{\mathfrak{g}})_x(v, w) = [\eta_x(v), \sigma_x(w)]_{\mathfrak{g}} + [\sigma_x(v), \eta_x(w)]_{\mathfrak{g}}$$

for any $x \in M$ and $v, w \in T_x M$.

Proposition 2.2.8 (Structural equation). *Let ω_{MC} denote the Maurer-Cartan form on a Lie group G and $\mathfrak{g} = \text{Lie}(G)$, then $d\omega_{MC} + \frac{1}{2}[\omega_{MC}, \omega_{MC}]_{\mathfrak{g}} = 0$.*

Proof. Let X, Y be two left-invariant vector fields on G , meaning $X_g = D_1 L_g [X_1]$ for all $g \in G$ from which $\omega_{MC}(X_g) = X_1$ follows and also $\omega_{MC}([X, Y]) = [X, Y]_1 = [X_1, Y_1]_{\mathfrak{g}}$, where the bracket on the left is the vector field commutator. Then

$$d\omega_{MC}(X, Y) = X(\omega_{MC}(Y)) - Y(\omega_{MC}(X)) - \omega_{MC}([X, Y]).$$

Since $\omega_{MC}(X)$ and $\omega_{MC}(Y)$ are constant, the first two terms drop out and we are left with

$$\begin{aligned} d\omega_{MC}(X, Y) &= -[X_1, Y_1]_{\mathfrak{g}} = -\frac{1}{2}([\omega_{MC}(X), \omega_{MC}(Y)]_{\mathfrak{g}} + [\omega_{MC}(X), \omega_{MC}(Y)]_{\mathfrak{g}}) \\ &= -\frac{1}{2}[\omega_{MC}, \omega_{MC}]_{\mathfrak{g}}(X, Y) \end{aligned}$$

from which $(d\omega_{MC} + \frac{1}{2}[\omega, \omega])(X, Y) = 0$ for such left-invariant X, Y . But these span the tangent space at any point and thus $d\omega_{MC} + \frac{1}{2}[\omega_{MC}, \omega_{MC}]_{\mathfrak{g}} = 0$. \square

Returning to the question of when a 1-form $\omega : TM \rightarrow \mathfrak{g}$ admits an anti-derivative, it immediately follows that one necessary conditions is that ω satisfies the structural equation:

Corollary 2.2.9. *Let G be a Lie group with Lie algebra \mathfrak{g} , M a manifold and $f : M \rightarrow G$ a smooth function. Then the Darbou derivative ω_f of f satisfies the equation $d\omega_f + \frac{1}{2}[\omega_f, \omega_f]_{\mathfrak{g}} = 0$.*

Proof. One has $\omega_f = f^*(\omega_{MC})$, that is $(\omega_f)_x(v) = (\omega_{MC})_{f(x)} D_x f [v]$ for any $x \in M$ and $v \in T_x M$. Now for $x \in M$ and $v, w \in T_x M$ it follows that

$$(d\omega_f + \frac{1}{2}[\omega_f, \omega_f]_{\mathfrak{g}})_x(v, w) = d\omega_{MC}(D_x f [v], D_x f [w]) + \frac{1}{2}[\omega_{MC}, \omega_{MC}]_{\mathfrak{g}}(D_x f [v], D_x f [w]) = 0$$

by Proposition 2.2.8. \square

Indeed this condition is the only thing that prevents us, at least locally, from finding an anti-derivative. For that reason the following theorem may be called a (local) fundamental theorem of calculus.

Theorem 2.2.10 (Fundamental theorem of calculus, [Sh 97] Theorem 6.1). *Let G be a Lie group with Lie algebra \mathfrak{g} and M a smooth manifold. Suppose $\omega : TM \rightarrow \mathfrak{g}$ is a \mathfrak{g} -valued 1-form on M so that $d\omega + \frac{1}{2}[\omega, \omega] = 0$. Then for each $x \in M$ there is a neighbourhood U of x and a function $f : U \rightarrow G$ so that $\omega|_U = \omega_f$.*

Proof ([Sh 97]). We will construct a distribution \mathcal{D} on $M \times G$ given by $\mathcal{D} = \ker(\pi_M^*(\omega) - \pi_G^*(\omega_{MC}))$. Further we will show that this distribution is integrable, and that π_M induces a local diffeomorphism from the leaves of the associated foliation to M . For $x \in M$ we then choose a neighbourhood U and a leaf L so that $\pi_M|_{\pi_M^{-1}(U) \cap L}$ is a diffeomorphism. The inverse must necessarily be of the form $\iota : U \rightarrow M \times G$, $y \mapsto (y, f(y))$ for a function f . By construction $\iota^*(\pi_M^*(\omega) - \pi_G^*(\omega_{MC})) = 0$, but this is equal to:

$$0 = \iota^*(\pi_M^*(\omega)) - \iota^*(\pi_G^*(\omega_{MC})) = (\pi_M \circ \iota)^*(\omega) - (\pi_G \circ \iota)^*(\omega_{MC}) = \omega - f^*(\omega_{MC}),$$

which shows that f is the desired anti-derivative.

We begin by showing that \mathcal{D} is a distribution, this follows from seeing that the map $D_{(y,g)}\pi_M|_{\mathcal{D}} : \mathcal{D}_{(y,g)} \rightarrow T_yM$, $(v, w) \mapsto v$ is a vector space isomorphism for any $y \in M, g \in G$. The map is clearly injective, for if $D_{(y,g)}\pi_M(v, w) = 0$ then $v = 0$, from which $\omega_{MC}(w) = \omega(v) = 0$ follows if $(v, w) \in \mathcal{D}_{(y,g)}$, this means that $(v, w) = 0$ implying injectivity. For surjectivity consider a $v \in T_yM$, then $\omega(v) \in T_1G$ and $w := (\omega_{MC})_g^{-1}(\omega(v)) \in T_gG$ with $\omega(v) - \omega_{MC}(w) = 0$, so $(v, w) \in \mathcal{D}_{(y,g)}$ is a pre-image of v under $D\pi_M$.

To see that $\mathcal{D} = \ker(\pi_M^*(\omega) - \pi_G^*(\omega_{MC}))$ is integrable, we note that if a distribution is of the form $\mathcal{D} = \ker(\Omega)$ for a vector-valued 1-form Ω , then it is integrable if and only if $d\Omega|_{\mathcal{D}} = 0$. This follows from considering a local basis $\{X_1, \dots, X_k\}$ of \mathcal{D} on an open set V . Then \mathcal{D} is integrable on V if and only if $[X_i, X_j] \in \mathcal{D}$ for all i, j , but this is equivalent to $\Omega([X_i, X_j]) = 0$. Now

$$d\Omega(X_i, X_j) = X_i(\Omega(X_j)) - X_j(\Omega(X_i)) - \Omega([X_i, X_j]) = -\Omega([X_i, X_j]),$$

and $\Omega([X_i, X_j]) = 0$ for all i, j if and only if $d\Omega(X_i, X_j) = 0$, which just means $d\Omega|_{\mathcal{D}} = 0$.

In our case $\Omega = \pi_M^*(\omega) - \pi_G^*(\omega_{MC})$, from which

$$d\Omega = \pi_M^*(d\omega) - \pi_G^*(d\omega_{MC}) = -\frac{1}{2}[\pi_M^*\omega, \pi_M^*\omega]_{\mathfrak{g}} + \frac{1}{2}[\pi_G^*\omega_{MC}, \pi_G^*\omega_{MC}]_{\mathfrak{g}}$$

follows by ω and ω_{MC} satisfying the property $d\omega + \frac{1}{2}[\omega, \omega]_{\mathfrak{g}} = 0$. If we plug $\pi_M^*\omega = \Omega + \pi_G^*\omega_{MC}$ into that equation and remember Definition 2.2.7 we find $d\Omega = -[\pi_G^*\omega_{MC}, \Omega]_{\mathfrak{g}} - \frac{1}{2}[\Omega, \Omega]_{\mathfrak{g}}$, whence $\Omega(X) = 0 = \Omega(Y)$ implies $d\Omega(X, Y) = 0$, so $d\Omega|_{\mathcal{D}} = 0$ and \mathcal{D} is integrable.

That for any leaf of the associated foliation the map $\pi_M : L \rightarrow M$ is a local diffeomorphism then follows from $D\pi$ inducing a vector space isomorphism $\mathcal{D}_{(y,g)} \rightarrow T_y M$ for any $(y, g) \in M \times G$, in particular if $(y, g) \in L$ then

$$D_{(y,g)}\pi_M|_{TL}: T_{(y,g)}L = \mathcal{D}_{(y,g)} \rightarrow T_y M$$

is a vectorspace isomorphism and $\pi_M|_L$ a local diffeomorphism. \square

Remark. Let M be a manifold and $\omega : TM \rightarrow \mathbb{R}$ a 1-form on M . It is well known that there exists a function $f : M \rightarrow \mathbb{R}$ with $df = \omega$ if and only if $d\omega = 0$ and the cohomology class of ω in $H^1(M)$ vanishes. The condition $d\omega = 0$, which locally guarantees the existence of such an f , is however the same condition as in Theorem 2.2.10, since $[\cdot, \cdot]_{\mathbb{R}} = 0$.

2.3 Building spaces out of model spaces

2.3.1 Glueing flat model space pieces together

Definition 2.3.1 (*G/P atlas*). Let M be a manifold and G/P a model space. A G/P chart in M is a triple (U, φ, V) where $U \subseteq G/P$ and $V \subseteq M$ are open and $\varphi : U \rightarrow V$ is a diffeomorphism. A G/P atlas of M is a collection \mathcal{A} of G/P charts covering M so that for any $(U_1, \varphi_1, V_1), (U_2, \varphi_2, V_2)$ in \mathcal{A} one has that

$$\varphi_1^{-1}\varphi_2 : U_2 \cap \varphi_2^{-1}(V_1) \rightarrow U_1 \cap \varphi_1^{-1}(V_2)$$

is on each connected component the restriction of an element of G acting by left-multiplication.

As is usual for the definition of an atlas we say two atlases \mathcal{A}_1 and \mathcal{A}_2 are equivalent if their union is again an atlas. We may note that equivalence in this sense is indeed an equivalence relation as our condition on the chart switching maps is a local condition, which ensures transitivity of the relation. To each atlas we can then associate then a maximal atlas, which is the union of all other atlases which are equivalent to the original one.

It is convenient to work with the maximal atlas, because we can then assume that the open sets of the charts are as fine as we please.

Definition 2.3.2 (*Glued space*). We call a manifold M together with an equivalence class $[\mathcal{A}]$ of G/P -atlases a locally G/P space or a G/P -glued space.

We briefly remind ourselves that G is a P -principal bundle over G/P and that the automorphisms of G/P (which are left multiplications by elements of G) lift to bundle-automorphisms of G (to left-multiplication with the same element of G). So if we construct a G/P -glued space by

literally glueing a collection of open sets U_i of G/P to one-another by maps that are locally automorphisms, we might as well glue the associated P -principal bundles $\pi^{-1}(U_i)$ together by the lifts of these maps.

Being more specific a G/P atlas of M induces an atlas on M of the bundle $G \rightarrow G/P$.

Definition 2.3.3. Let B, M be manifolds and $E \xrightarrow{\pi_E} B$ be a fibre-bundle over B with fibre F . An atlas on M of the bundle $E \rightarrow B$ is a collection of charts $\{(U_i, f_i, V_i) \mid i \in I\}$ where $f_i : U_i \rightarrow V_i$ is a diffeomorphism between $U_i \subset B$ is open and $V_i \subset M$ is open with V_i covering M together with a system of transition functions $\{g_{ij} : \pi_B^{-1}(f_j^{-1}(V_i \cap V_j)) \rightarrow \pi_B^{-1}(f_i^{-1}(V_i \cap V_j)) \mid i, j \in I\}$ where:

1. g_{ij} is a bundle-automorphism over $f_i^{-1}f_j$.
2. $g_{ij}g_{jk} = g_{ik}$ for all $i, j, k \in I$, where the domains are restricted to $\pi_B^{-1}(f_k^{-1}(V_k \cap V_j \cap V_i))$.
3. $g_{ii} = \text{id}$ for all $i \in I$.

This definition of an atlas is more general than what is usually considered useful. It is captured by, but less general than, the definition of a fibre bundle by coordinate transforms in [Gr 55]. The relevant statement is that such a data glues together to give the structure of a fibre bundle on M that is locally isomorphic to the bundle E over B via the chart maps f_i . We cite this as a standard result about fibre bundles:

Lemma 2.3.4. Let B, M be manifolds and $E \rightarrow B$ a fibre bundle over B with fibre F . An atlas on M of the bundle $E \rightarrow B$ induces a bundle $\widehat{M} \rightarrow M$ over M with fibre F so $\pi_M^{-1}(V_i)$ is isomorphic to the pullback bundle $(f_i^{-1})^*(\pi_B^{-1}(U_i))$ for all i .

In the case of a G/P atlas $\{U_i, \varphi_i, V_i \mid i \in I\}$ on M , the construction gives a P -principal bundle \widehat{M} . The transition functions in this case are $g_{ij} = \widehat{\varphi_i^{-1}\varphi_j}$, which is supposed to be the unique lift of $\varphi_i^{-1}\varphi_j$ by a (locally) left-multiplication as in Proposition 2.2.1. Specifically \widehat{M} is

$$\widehat{M} := \left(\prod_{i \in I} \pi^{-1}(U_i) \right) / \sim$$

with relation $(x, U_i) \sim (y, U_j)$ if $\pi(x) \in U_i \cap \varphi_i^{-1}\varphi_j(U_j)$ and $y = (\widehat{\varphi_j^{-1}\varphi_i})(x)$. We denote the quotient map $\widehat{M} \rightarrow M$, $[(x, U_i)] \mapsto \varphi_i(x)$ with π as well.

The transition functions $g_{ij} = \widehat{\varphi_i^{-1}\varphi_j}$ are in this case locally left-multiplications, meaning equal to a left-multiplication on every connected component of their domain. Such maps preserve the Maurer-Cartan form ω_{MC} by Proposition 2.2.3, meaning $g_{ij}^*(\omega_{MC}) = \omega_{MC}$. Since the bundle \widehat{M} is defined by gluing pieces of G together via g_{ij} , this means that locally defining

$$\omega : T\widehat{M} \rightarrow \mathfrak{g}, \quad \omega_{[x, U_i]}[v] = (\omega_{MC})_x[v]$$

for $v \in T_{\pi^{-1}(x)}G$ results in a defined form ω .

Proposition 2.3.5. *Let $\omega : T\widehat{M} \rightarrow \mathfrak{g}$ be the glued form as above, then:*

1. *For each $\widehat{x} \in \widehat{M}$ the map $\omega_{\widehat{x}} : T_{\widehat{x}}\widehat{M} \rightarrow \mathfrak{g}$ is an isomorphism of vector spaces.*
2. *For each $h \in P$ we have $\omega_{\widehat{x}\cdot h} \circ D_{\widehat{x}}R_h = \text{Ad}(h^{-1})\omega_{\widehat{x}}$ where R_h is the right-multiplication with h on the bundle \widehat{M} .*
3. *For each $\xi \in \mathfrak{p}$ one has $\omega_{\widehat{x}}(\frac{d}{dt}\widehat{x} \cdot \exp(t\xi)|_{t=0}) = \xi$, i.e. ω sends the fundamental fields of the P -principal bundle to their generators in \mathfrak{p} .*
4. $d\omega + \frac{1}{2}[\omega, \omega] = 0$.

Proof. These are all local properties, so if they hold for $\omega_{MC} : TG \rightarrow \mathfrak{g}$ then they hold for ω . But these properties were already verified for ω_{MC} (for example 4. is the structural equation Proposition 2.2.8) or are obvious. \square

So any G/P -glued space M admits the structure of a P -principal bundle $\widehat{M} \rightarrow M$ and a \mathfrak{g} -valued 1-form $\omega : T\widehat{M} \rightarrow \mathfrak{g}$ satisfying properties 1.-4. from Proposition 2.3.5. The data of a P -principal bundle over M with such a form ω is however equivalent to that of G/P atlas on M .

Theorem 2.3.6 (Flat connections have a G/P atlas). *If $\widehat{M} \rightarrow M$ is a P -principal bundle together with an \mathfrak{g} -valued 1-form ω satisfying conditions 1. to 4. from Proposition 2.3.5, then there is a G/P atlas on M inducing this bundle and 1-form.*

Proof. This is an application of the fundamental theorem of calculus, Theorem 2.2.10, and the local characterisation of left-multiplications Lemma 2.2.4. The rough idea is that the Properties 1.-4. and the fundamental theorem guarantee the existence of local bundle-isomorphisms $F : \pi^{-1}(V) \rightarrow \pi^{-1}(U)$ for $V \subset M$ a small open neighbourhood of any point and $U \subset G/P$ open. These bundle-isomorphisms will satisfy $F^*(\omega_{MC}) = \omega$. If we have another such isomorphism F' we find $(F^{-1} \circ F')^*(\omega_{MC}) = \omega_{MC}$ implying by Lemma 2.2.4 that $F^{-1} \circ F'$ is a left-multiplication. Thus if we take $\{(U, [F^{-1}], V)\}$ as a system of charts of M into G/P , the chart switching maps will be (locally at least) left-multiplications and we recover a G/P atlas on M .

If we let \widehat{M}' be the P -principal bundle induced by this atlas we define a map $\Psi : \widehat{M}' \rightarrow \widehat{M}$, $(x, U) \mapsto F^{-1}(x)$, which is well defined by construction of the atlas. Further it is locally a bundle-isomorphism since F is an isomorphism $\pi_M^{-1}(V) \rightarrow \pi_{G/P}^{-1}(U)$ and it is surjective since every point in M has a neighbourhood V with such an F . Lastly the map is injective since if $\Psi((x, U)) = \Psi((x', U'))$ then there must be $(U, [F], V), (U', [F'], V')$ with $\Psi((x, U)) \in \pi_M^{-1}(V) \cap \pi_M^{-1}(V')$ and

$F(\Psi((x, U))) = x$ and $F'(\Psi((x, U))) = x'$. In particular $(F'F^{-1})(x) = x'$, so (x, U) and (x', U') are identified in the glueing quotient.

This means that \widehat{M} and the just constructed \widehat{M}' are isomorphic as P -bundles, which is what we wanted to show. We now show that such an F does actually exist as well as the property $(F^{-1}F')^*(\omega_{MC}) = \omega_{MC}$:

Let $x \in M$ and $\widehat{x} \in \pi^{-1}(x)$. Then by property 4. and Theorem 2.2.10 there exists a neighbourhood $U_1 \subset \widehat{M}$ of \widehat{x} and a map $F : U_1 \rightarrow G$ with $F^*(\omega_{MC}) = \omega$.

In particular DF must map the ω constant fields to the ω_{MC} constant fields, which means that $F(\varphi_t^{\omega^{-1}(X)}(\widehat{y})) = \varphi_t^{\omega_{MC}^{-1}(X)}(F(\widehat{y}))$ for $\widehat{y} \in U_1$, $X \in \mathfrak{g}$ and t small enough. By point 3. the vertical ω constant fields are the fundamental fields of \widehat{M} . It follows for $\widehat{y} \in U_1$ and $\xi \in \mathfrak{p}$ that $F(\widehat{y} \cdot e^{t\xi}) = F(\widehat{y}) \cdot e^{t\xi}$ for t small enough.

We now make U_1 very small. Let $U_2 \subset U_1$ be open containing \widehat{x} and small enough that there exists a section $s : \pi(U_2) \rightarrow U_2$ and $U_3 \subset U_1$ even smaller so that there exists a $V \subset \mathfrak{p}$ open with $U_3 = \{s(y) \cdot \exp(V) \mid y \in \pi(U_2)\}$. This means that $U_3 \cong \pi(U_2) \times \exp(V)$ in the local trivialisation $U_2 \cong \pi(U_2) \times P$ given by the section s .

On U_3 we have by construction $F(\widehat{y} \cdot p) = F(\widehat{y}) \cdot p$ for all $\widehat{y} \in U_3$, $p \in P$ with $\widehat{y} \cdot p \in U_3$. We now forget U_1, U_2 and let U_3 be the domain of F . We fill up the vertical directions of U_3 by letting $U_4 = U_3 \cdot P = \pi^{-1}(U_3)$ and extending F to U_4 by:

$$F(\widehat{y} \cdot p) = F(\widehat{y}) \cdot p,$$

this makes F a P -bundle map from U_4 to its image in G and implies $D_{\widehat{y}p}F \circ D_{\widehat{y}}R_p = D_{F(\widehat{y})}R_p \circ D_{\widehat{y}}F$.

We still have $(F|_{U_3})^*(\omega_{MC}) = \omega|_{U_3}$ and want to see this for all of U_4 . We find for $\widehat{y} \in U_3$, $p \in P$ that

$$\begin{aligned} F^*(\omega_{MC})_{\widehat{y}p} &= (\omega_{MC})_{F(\widehat{y}p)} \circ D_{\widehat{y}p}F = (\omega_{MC})_{F(\widehat{y}p)} \circ D_{\widehat{y}p}F \circ D_{\widehat{y}}R_p \circ D_{\widehat{y}p}R_{p^{-1}} \\ &= (\omega_{MC})_{F(\widehat{y})p} \circ D_{F(\widehat{y})}R_p \circ D_{\widehat{y}}F \circ D_{\widehat{y}p}R_{p^{-1}} = \text{Ad}(p^{-1})(\omega_{MC})_{F(\widehat{y})} \circ D_{\widehat{y}}F \circ D_{\widehat{y}p}R_{p^{-1}} \\ &= \text{Ad}(p^{-1}) \cdot F^*(\omega_{MC})_{\widehat{y}} \circ D_{\widehat{y}p}R_{p^{-1}} = \text{Ad}(p^{-1}) \cdot \omega_{\widehat{y}} \circ D_{\widehat{y}p}R_{p^{-1}} = \omega_{\widehat{y}p}. \end{aligned}$$

where the second line uses the equality $D_{\widehat{y}p}F \circ D_{\widehat{y}}R_p = D_{F(\widehat{y})}R_p \circ D_{\widehat{y}}F$ just established as well as Property 2. from 2.3.5 holding for ω_{MC} . The third line then follows from $F^*(\omega_{MC})|_{U_3} = \omega|_{U_3}$ and Property 2. holding for ω .

So now we have that $F : U_4 \rightarrow G$ is a P -bundle map with $F^*(\omega_{MC}) = \omega$ on U_4 . Since ω is an

isomorphism on each tangent space, $D_{\hat{y}}F$ must be invertible for all $y \in U_4$ and F is actually a bundle-automorphism between U_4 and $F(U_4)$.

If we have another bundle-automorphism $F' : U_5 \rightarrow F'(U_5)$ with $(F')^*(\omega_{MC}) = \omega$, here U_5 is some other open set in \widehat{M} , then for each $\hat{y} \in U_5 \cap U_4$ we have

$$(F' \circ F^{-1})^*(\omega_{MC}) = (F^{-1})^*((F')^*(\omega_{MC})) = (F^{-1})^*(\omega) = \omega_{MC},$$

implying by Lemma 2.2.4 that $F'F^{-1}$ is locally a left multiplication with an element of G . \square

2.3.2 Rolling the model space on an inhomogeneous manifold

The previous section described a scenario in which a manifold M could be given by glueing together pieces of a model space G/P with glueing maps being automorphisms. We noted that such a datum was equivalent to the structure of a P -principal bundle over M together with a \mathfrak{g} -valued 1-form on this principal bundle satisfying the compatibility conditions listed in Proposition 2.3.5.

In this section we will drop condition 4., which ensured that the principal bundle \widehat{M} was locally isomorphic to G (which is viewed as a P -principal bundle over G/P). We will then try to establish that the data of such a bundle and form provides us a way of “rolling” the homogenous space G/P along paths in the manifold M . By noting that the end point of this rolling depends not just on the homotopy class of the path we draw the conclusion that M is inhomogeneous or “lumpy”.

In the following let G/P be a model space and M a manifold, we will consider a P -principal bundle $\widehat{M} \xrightarrow{\pi} M$ together with a \mathfrak{g} -valued 1-form $\omega : T\widehat{M} \rightarrow \mathfrak{g}$ satisfying conditons 1. - 3. from Proposition 2.3.5. This is called a Cartan connection, but we will not formally make that definition until Section 2.4.

Proposition 2.3.7. *Let $\sigma : [a, b] \rightarrow \widehat{M}$ a path. Then for any $g \in G$ there exists a unique map $\tilde{\sigma}_g : [a, b] \rightarrow G$, called the development of σ with base point g , so that $\tilde{\sigma}(a) = g$ and for which Darboux derivative $\omega_{\tilde{\sigma}_g}$ is $\omega \circ D\sigma : T\mathbb{R} \rightarrow \mathfrak{g}$, i.e. $\tilde{\sigma}_g$ is a primitive of $\omega \circ D\sigma$.*

Proof. Note that since there are no non-zero 2-forms on \mathbb{R} , that then

$$d(\omega \circ D\sigma) + \frac{1}{2}[\omega \circ D\sigma, \omega \circ D\sigma]_{\mathfrak{g}} = 0$$

and an anti-derivative of $\omega \circ D\sigma$ exists around each point in $[a, b]$ by the fundamental theorem. We assume these neighbourhoods are connected and then choose a finite covering U_1, \dots, U_N of $[a, b]$ by such neighbourhoods, further restricting $U_i \cap U_j$ to be empty for $|j - i| > 1$ and an interval else.

Let $\sigma'_1 : U_1 \rightarrow G$ be the anti-derivative on U_1 , we assume $\sigma'_1(a) = g$ as can be achieved by post-composing with $L_{g \cdot \sigma_1(a)^{-1}}$, which doesn't change the Darboux-derivative. Next for $i > 1$ we chose an $a_i \in U_i \cap U_{i-1}$ and assume the anti-derivative $\sigma'_i : U_i \rightarrow G$ to satisfy $\sigma'_i(a_i) = \sigma'_{i-1}(a_i)$ in the same manner. Since $U_i \cap U_{i-1}$ is connected and anti-derivatives are unique up to multiplication with a constant by Proposition 2.2.6 we find $\sigma'_i|_{U_i \cap U_{i-1}} = \sigma'_{i-1}|_{U_i \cap U_{i-1}}$.

Since each U_i only intersects U_{i+1} and U_{i-1} the procedure of glueing the individual σ'_i together to get a global $\tilde{\sigma}_g : [a, b] \rightarrow G$ is well defined. $\tilde{\sigma}_g$ remains an anti-derivative since the Darboux-derivatives carry only local information. Since $\tilde{\sigma}_g(a)$ is fixed to be g and anti-derivatives on connected sets are unique up to multiplication by a constant (as by Proposition 2.2.6) $\tilde{\sigma}_g$ is uniquely defined. \square

Corollary 2.3.8. *If $\sigma : [a, b] \rightarrow \widehat{M}$ is a path, then for any $g, h \in G$ one has $\tilde{\sigma}_g(t) = gh^{-1} \cdot \tilde{\sigma}_h(t)$.*

Proof. $\tilde{\sigma}_g$ and $gh^{-1}\tilde{\sigma}_h$ have the same Darboux-derivatives with $\tilde{\sigma}_g(a) = g = gh^{-1}h = gh^{-1}\tilde{\sigma}_h(a)$ so the two maps are equal by connectedness of $[a, b]$ and Proposition 2.2.6. \square

So while we see that we can develop paths $[a, b] \rightarrow \widehat{M}$ to paths in G , what we are really interested in is developing paths in M to paths in the model space G/P . In order for this to work two different lifts of the same path $[a, b] \rightarrow M$ to paths $[a, b] \rightarrow \widehat{M}$ must have similar developements in G . The following proposition is the technical step making this clear.

Proposition 2.3.9. *Let $h : [a, b] \rightarrow P$, $\sigma : [a, b] \rightarrow \widehat{M}$ be paths. Then the development $\widetilde{\sigma h_{h(a)}}$ of $\sigma h : [a, b] \rightarrow \widehat{M}$, $t \mapsto \sigma(t) \cdot h(t)$ with basepoint $h(a)$ is equal to $t \mapsto \tilde{\sigma}_1(t) \cdot h(t)$.*

Proof. Note that $\widetilde{\sigma h_{h(a)}}(a) = h(a) = 1 \cdot h(a) = \tilde{\sigma}_1(a) \cdot h(a)$. Next we check that the Darboux-derivatives of the two maps agree and then the maps must agree. First we look at $\tilde{\sigma}_1 \cdot h$:

$$(\tilde{\sigma}_1 \cdot h)^*(\omega_{MC})_t = \omega_{MC} \circ D_t(\tilde{\sigma}_1 \cdot h) = \omega_{MC} \circ D_{(\tilde{\sigma}_1(t), h(1))}\mu \circ (D_t\tilde{\sigma}_1, D_t h),$$

where $\mu : G \times G$ is the multiplication. We know that $D_{(a,b)}\mu[(v, w)] = D_a R_b[v] + D_b L_a[w]$ for $v \in T_a G, w \in T_b G$. So what we get is:

$$\begin{aligned} (\tilde{\sigma}_1 \cdot h)^*(\omega_{MC})_t &= \omega_{MC} \circ D_{\tilde{\sigma}_1(t)} R_{h(t)} \circ D_t \tilde{\sigma}_1 + \omega_{MC} \circ D_{h(t)} L_{\tilde{\sigma}_1(t)} \circ D_t h \\ &= \text{Ad}(h(t)^{-1}) \omega_{MC} \circ D_t \tilde{\sigma}_1 + \omega_{MC} \circ D_t h = \text{Ad}(h(t)^{-1}) \cdot (\sigma^*(\omega))_t + \omega_{MC} \circ D_t h, \end{aligned}$$

where the first equality of the second line follows from $\omega_{MC} \circ D_a R_b = \text{Ad}(b^{-1})\omega_{MC}$ and $\omega_{MC} \circ D_a L_b = \omega_{MC}$. Now we look at $(\widetilde{\sigma h_{h(a)}})^*(\omega_{MC})$:

$$(\widetilde{\sigma h_{h(a)}})^*(\omega_{MC}) = (\sigma \cdot h)^*(\omega) = \omega \circ D_t(\sigma \cdot h) = \omega \circ D_{(\sigma(t), h(t))}\mu_2 \circ (D_t \sigma, D_t h),$$

where $\mu_2 : \widehat{M} \times P \rightarrow \widehat{M}$ is the multiplication. One has $D_{(\widehat{x}, p)}\mu_2[(v, w)] = D_{\widehat{x}}R_p[v] + D_p ev_{\widehat{x}}[w]$ for $(\widehat{x}, p) \in \widehat{M} \times P$ and $v \in T_{\widehat{x}}\widehat{M}$, $w \in T_pP$. This implies:

$$\begin{aligned} (\widetilde{\sigma}h_{h(a)})^*(\omega_{MC}) &= \omega \circ D_{\sigma(t)}R_{h(t)} \circ D_t\sigma + \omega \circ D_{h(t)}ev_{\sigma(t)} \circ D_th \\ &= \text{Ad}(h(t)^{-1})\omega \circ D_t\sigma + \omega \circ D_{h(t)}ev_{\sigma(t)} \circ D_th \\ &= \text{Ad}(h(t)^{-1})(\sigma^*(\omega))_t + \omega \circ D_{h(t)}ev_{\sigma(t)} \circ D_th, \end{aligned}$$

where property 2. of ω from Proposition 2.3.5 was used in the second line. In order understand the second summand, we note that $\pi \circ ev_{\widehat{x}} : P \rightarrow M$ is constant, whence $D_{h(t)}ev_{\sigma(t)} \circ D_th$ is vertical, it is the differential of the map $\frac{d}{ds}\sigma(t)h(t+s)|_{s=0} = \frac{d}{ds}\sigma(t)h(t)h(t)^{-1}h(t+s)|_{s=0}$. As s varies $h(t)^{-1}h(t+s)$ is connected to the identity and we thus express it as $\exp(\xi(s))$ with $\xi(0) = 0$. This implies that $\frac{d}{ds}\sigma(t)h(t+s)|_{t=0} = \frac{d}{ds}\sigma(t)h(t) \cdot \exp(\xi(s))|_{s=0}$ is the evaluation of the fundamental field given by $\dot{\xi}(0)$ at $\sigma(t)h(t)$. Then $\omega(\frac{d}{ds}\sigma(t)h(t) \cdot \exp(\xi(s))|_{s=0}) = \dot{\xi}(0)$ by property 3. of ω as in Proposition 2.3.5. On the other hand

$$\dot{\xi}(0) = \frac{d}{ds}\exp(\xi(s))|_{s=0} = \frac{d}{ds}h(t)^{-1}h(t+s)|_{s=0} = D_{h(t)}L_{h(t)^{-1}} \circ D_th = \omega_{MC} \circ D_th.$$

This was the last step needed to see $(\widetilde{\sigma}_1 \cdot h)^*(\omega_{MC})_t = (\widetilde{\sigma}h_{h(a)})^*(\omega_{MC})$, and the proposition follows. \square

Corollary 2.3.10. *Let $\sigma : [a, b] \rightarrow M$ be a path and $\sigma^1, \sigma^2 : [a, b] \rightarrow \widehat{M}$ two lifts of σ with $\sigma^1(a) = \sigma^2(a)$ and $g \in G$. Then the maps $[\widetilde{\sigma}_g^1], [\widetilde{\sigma}_g^2] : [a, b] \rightarrow G \rightarrow G/P$ are the same.*

Proof. Note that there must be a function $h : [a, b] \rightarrow P$ with $h(a) = 1$ and $\sigma^1(t) = \sigma^2(t) \cdot h(t)$. Applying both Proposition 2.3.9 and Corollary 2.3.8 gives

$$\widetilde{\sigma}_g^1(t) = g \cdot \widetilde{\sigma}_1^1(t) = g\widetilde{\sigma}^2h_1(t) = g\widetilde{\sigma}_1^2(t)h(t) = \widetilde{\sigma}_g^2 \cdot h(t).$$

Composing both sides with the quotient $G \rightarrow G/P$ then results in the same path. \square

Definition 2.3.11. Let $\sigma : [a, b] \rightarrow M$ be a path, for a $\widehat{x} \in \pi^{-1}(\sigma(a))$, we denote with $r^{\widehat{x}}(\sigma) : [a, b] \rightarrow G/P$ the projection to G/P of the development at 1 of any lift at \widehat{x} of σ in \widehat{M} . We call $r^{\widehat{x}}(\sigma)$ the rolling map.

The previous corollary shows that this definition is well defined.

Remark. The choice of a point in the fibre $\widehat{x} \in \pi^{-1}(\sigma(a))$ amounts to choosing a specific way of glueing $T_1(G/P) = \mathfrak{g}/\mathfrak{p}$ to $T_{\sigma(a)}M$. This is because the kernel of the map $\pi_{\mathfrak{g} \rightarrow \mathfrak{g}/\mathfrak{p}} \circ \omega_{\widehat{x}} : T_{\widehat{x}}\widehat{M} \rightarrow \mathfrak{g} \rightarrow \mathfrak{g}/\mathfrak{p}$ is just $\ker(D_{\widehat{x}}\pi_{\widehat{M} \rightarrow M})$. Thus $\omega_{\widehat{x}}$ induces a map $[\omega_{\widehat{x}}] : T_{\pi(\widehat{x})}M \rightarrow \mathfrak{g}/\mathfrak{p}$, which remains an

isomorphism. The following diagram may make it clear:

$$\begin{array}{ccc} T_{\widehat{x}}\widehat{M} & \xrightarrow{\omega_{\widehat{x}}} & \mathfrak{g} \\ \downarrow D_{\widehat{x}}\pi & & \downarrow \pi' \\ T_{\pi(x)}M & \xrightarrow{[\omega_{\widehat{x}}]} & \mathfrak{g}/\mathfrak{p} \end{array}$$

here $\pi : \widehat{M} \rightarrow M$ and $\pi' : \mathfrak{g} \rightarrow \mathfrak{g}/\mathfrak{p}$ are the projections. The rolling map then amounts to choosing such an identification and then carrying out the infinitesimal translation of $\dot{\sigma}(0)$ in G/P . From this perspective choosing a different point \widehat{x}' in the fibre $\pi^{-1}(\sigma(a))$ should result in the rolling maps $r^{\widehat{x}'}(\sigma)$ and $r^{\widehat{x}}(\sigma)$ differing only by an initial rotation. The precise way in which this is true is the following proposition.

Proposition 2.3.12. *Let $\sigma : [a, b] \rightarrow M$ and $\widehat{x} \in \pi^{-1}(\sigma(a))$, $h \in P$. Then $r^{\widehat{x} \cdot h}(\sigma) = \text{Ad}(h^{-1})r^{\widehat{x}}(\sigma)$.*

Proof. Let σ^1 be a lift of σ to \widehat{M} with $\sigma^1(a) = \widehat{x}$. Then $\sigma^1(t) \cdot h$ is a lift of σ at $\widehat{x} \cdot h$, we will call it σ^2 . By Proposition 2.3.9 and Corollary 2.3.8 we have

$$\widetilde{\sigma^2}_1(t) = (\widetilde{\sigma^1 h})_1 = h^{-1}(\widetilde{\sigma^1 h})_h(t) = h^{-1}(\widetilde{\sigma^1})_1(t) = h^{-1}\widetilde{\sigma^1}_1(t) \cdot h = \text{Ad}(h^{-1}) \cdot \widetilde{\sigma^1}_1(t),$$

and it follows that $r^{\widehat{x} \cdot h}(\sigma) = \text{Ad}(h^{-1})r^{\widehat{x}}(\sigma)$. □

2.4 Cartan connection

We now formally define what a Cartan connection is. In this section we will also briefly derive some results about automorphisms and isomorphisms of Cartan connections. The relevant statement being that these maps, which are principal bundle maps, are uniquely determined by their base maps. Lemma 2.4.3 is the step proving this.

Definition 2.4.1 (Cartan connection). Let G be a Lie group, $P \subset G$ a closed subgroup and $\mathfrak{g} = \text{Lie}(G)$ and $\mathfrak{p} = \text{Lie}(P)$. A Cartan connection on a manifold M modelled on G/P is a right P -principal bundle $P \rightarrow \widehat{M} \xrightarrow{\pi} M$ together with a \mathfrak{g} -valued 1-form $\omega : T\widehat{M} \rightarrow \mathfrak{g}$ so that:

1. For each $\widehat{x} \in \widehat{M}$ the map $\omega_{\widehat{x}} : T_{\widehat{x}}\widehat{M} \rightarrow \mathfrak{g}$ is a linear isomorphism.
2. $\omega_{\widehat{x} \cdot h} \circ D_{\widehat{x}}R_h = \text{Ad}(h^{-1}) \cdot \omega_{\widehat{x}}$ for all $h \in P$, where R_h is the right multiplication $\widehat{x} \mapsto \widehat{x} \cdot h$.
3. For $\xi \in \mathfrak{p}$ one has $\omega_{\widehat{x}}(\frac{d}{dt}\widehat{x} \cdot \exp(t\xi)|_{t=0}) = \xi$, i.e. ω sends the fundamental fields of the P -principal bundle to their generators in \mathfrak{p} .

Definition 2.4.2 (Automorphism). An automorphism of a Cartan connection $(\widehat{M} \rightarrow M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ on a manifold M modelled on G/P is a bundle automorphism $F : \widehat{M} \rightarrow \widehat{M}$ (not necessarily over the identity) so that $F^*(\omega) = \omega$.

The following lemma makes clear, that while the definition of an automorphism/symmetry of a Cartan connection involves the auxiliary space \widehat{M} , the data that determines an automorphism is entirely contained in the base map $M \rightarrow M$.

Lemma 2.4.3. *Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan connection on a manifold M modelled on G/P . Let F be a bundle-automorphism $\widehat{M} \rightarrow \widehat{M}$ over the identity so that $F^*(\omega) = \omega$, then F is the identity map.*

Proof. We stress that G acts effectively on G/P and that G/P is connected, which are points we have included in Definition 2.1.4 of a model space.

Since F is a bundle-automorphism over the identity it is of the form $F(\widehat{x}) = \widehat{x} \cdot f(\widehat{x})$ for a function $f : \widehat{M} \rightarrow P$. We will denote the closed subgroup generated by the image of f in P by N and its Lie-algebra by \mathfrak{n} . We will show that N is normal in G , whence it must be $\{1\}$ since we are requiring model spaces to be effective as per Definition 2.1.4, proving the proposition.

First we show that N is normal in P . Note that $F(\widehat{x}p) = \widehat{x}pf(\widehat{x}p) \stackrel{!}{=} F(\widehat{x})p = \widehat{x}f(\widehat{x})p$, implying $p^{-1}f(\widehat{x})p = f(\widehat{x}p)$ for all $p \in P$, giving that N is normal in P .

Next we remark that F is the composition $\widehat{M} \rightarrow \widehat{M} \times P \rightarrow \widehat{M}$, $\widehat{x} \mapsto (\widehat{x}, f(\widehat{x})) \mapsto \widehat{x}f(\widehat{x})$. Its derivative is then equal to $D_{\widehat{x}}F = D_{\widehat{x}}R_{f(\widehat{x})} + D_{f(\widehat{x})}ev_{\widehat{x}} \circ D_{\widehat{x}}f$ where $ev_{\widehat{x}}$ is the map $P \rightarrow \widehat{M}$, $p \mapsto \widehat{x} \cdot p$. In particular $D_1ev_{\widehat{x}}$ maps \mathfrak{p} to the fundamental fields on \widehat{x} , so $\omega_{\widehat{x}} \circ D_1ev_{\widehat{x}}$ is the identity on \mathfrak{p} . Further we note $D_p ev_{\widehat{x}} = D_1ev_{\widehat{x}p} \circ D_pL_{p^{-1}}$, whence $\omega_{\widehat{x}} \circ D_p ev_{\widehat{x}} = (\omega_{MC})_p$. Combining this with $F^*(\omega) = \omega$ we find

$$\omega_{\widehat{x}} = \omega_{\widehat{x}} \circ D_{\widehat{x}}F = \text{Ad}(f(\widehat{x})^{-1})\omega_{\widehat{x}} + (\omega_{MC})_{f(\widehat{x})} \circ D_{\widehat{x}}f$$

since f is valued in N its Darboux derivative $\omega_{MC} \circ Df$ is valued in \mathfrak{n} , and we find for any \widehat{x} , $\xi \in \mathfrak{g}$:

$$\text{Ad}(f(\widehat{x})^{-1})\xi - \xi = \text{Ad}(f(\widehat{x})^{-1})\omega_{\widehat{x}}(\omega_{\widehat{x}}^{-1}\xi) - \omega_{\widehat{x}}(\omega_{\widehat{x}}^{-1}\xi) = \omega_{MC}D_{\widehat{x}}f(\omega_{\widehat{x}}^{-1}\xi),$$

i.e. $\text{Ad}(f(\widehat{x})^{-1})\xi - \xi \in \mathfrak{n}$. As this holds for the generating set $f(\widehat{M})$ we actually get for any $n \in N$ and $\xi \in \mathfrak{g}$ that $\text{Ad}(n)\xi - \xi \in \mathfrak{n}$. Taking for example $n = \exp(t\eta)$ for $\eta \in \mathfrak{n}$ and passing to the derivative gives $[\eta, \xi] \in \mathfrak{n}$ for any $\eta \in \mathfrak{n}, \xi \in \mathfrak{g}$. In other words \mathfrak{n} is an ideal in \mathfrak{g} .

Now let $\xi \in \mathfrak{g}$ arbitrary, $n \in N$, we have just seen that $\text{Ad}(n)\xi = \xi + \eta$ for some $\eta \in \mathfrak{n}$. By the Zassenhaus formula $\exp(\xi + \eta) = \exp(\xi)\exp(\eta) \cdot \exp(\frac{-1}{2}[\xi, \eta]) \cdot \exp(\frac{1}{6}(2[\eta, [\xi, \eta]] + [\xi, [\xi, \eta]])) \cdot \dots$

where the successive factors are exponentials of some sum of commutators with ξ and η . Since \mathfrak{n} is an ideal in \mathfrak{g} these exponentials always lie in N , which is a closed subgroup so the (possibly infinite but converging) product of the exponentials lies in N . But this gives an element n' of N for which:

$$n \exp(\xi) n^{-1} = \exp(\text{Ad}(n)\xi) = \exp(\xi) n'$$

holds, implying $\exp(-\xi) n \exp(\xi) \in N$ for all $\xi \in \mathfrak{g}$. Thus N is normal in the connected component G^0 of G . Now if $g \in G$, then since G/P is connected there is a $g_0 \in G^0$ and a $p \in P$ with $g = g_0 p$. Then $gN = g_0 p N = g_0 N p = N g_0 p = Ng$ since N is normal in P and in G^0 . Then N must be normal in G . \square

Carrying out the comment before the lemma explicitly:

Corollary 2.4.4. *If $f : M \rightarrow M$ admits a lift $F : \widehat{M} \rightarrow \widehat{M}$ so that $F^*(\omega) = \omega$, then this lift is unique.*

Proof. If $F' : \widehat{M} \rightarrow \widehat{M}$ were another such lift, then $F \circ F'^{-1}$ would be a bundle isomorphism over the identity. Further $(F \circ F'^{-1})^*(\omega) = (F'^{-1})^*(F^*(\omega)) = (F'^{-1})^*(\omega) = \omega$. The lemma then implies that $F \circ F'^{-1}$ is the identity map, ie $F = F'$. \square

Often we may abuse notation by understanding automorphisms to be maps $M \rightarrow M$, the above corollary justifies this. Additionally we can define the notion of isomorphism for two Cartan geometries over the same model space:

Definition 2.4.5 (Isomorphism). An isomorphism of two Cartan connections $(\widehat{M}_i \rightarrow M_i, \omega_i : T\widehat{M}_i \rightarrow \mathfrak{g})$, $i \in \{1, 2\}$ modelled on G/P is a bundle automorphism $F : \widehat{M}_1 \rightarrow \widehat{M}_2$ so that $F^*(\omega_2) = \omega_1$.

In this scenario Lemma 2.4.3 can be applied again to see that the isomorphism only depends on the base map. As before we may abuse notation and understand isomorphisms as maps between the base spaces.

Along the same vein, if $U \subset M$ is an open subset of a manifold M equipped with a Cartan geometry $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ modelled on G/P then $(\pi^{-1}(U) \xrightarrow{\pi} U, \omega|_{\pi^{-1}(U)} : \pi^{-1}(U) \rightarrow \mathfrak{g})$ is also a Cartan geometry modelled on G/P . This leads to the following natural definition:

Definition 2.4.6 (Local isomorphism). Let $i \in \{1, 2\}$ and $(\widehat{M}_i \rightarrow M_i, \omega_i : T\widehat{M}_i \rightarrow \mathfrak{g})$ two Cartan geometries modelled on G/P . A map $f : U_1 \rightarrow U_2$ between open subset $U_i \subset M_i$ is called a local isomorphism if it admits a lift $F : \pi_1^{-1}(U_1) \rightarrow \pi_2^{-1}(U_2)$ that is an isomorphism on the Cartan geometries induced on the U_i .

2.5 Local automorphisms and Killing fields

In this section we investigate the structure of local automorphisms and Killing fields. Results of particular interest may be that the automorphism group of a connected space is always finite dimensional with dimension bounded by $\dim(G)$ and that isomorphisms between connected spaces as well as Killing fields are uniquely determined by their lifts at a point. The section ends with the construction of the “analyticity locus”¹ \mathfrak{M} , which is an open and dense subset of the manifold on which the sheaf of Killing fields is locally constant.

For what follows $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ will denote a Cartan connection on a manifold M modelled on G/P .

Definition 2.5.1 (Constant fields). The vector fields on \widehat{M} of the form $X_{\widehat{x}} = \omega_{\widehat{x}}^{-1}(v)$ for some $v \in \mathfrak{g}$ and all $\widehat{x} \in \widehat{M}$ are called the constant fields.

Since ω is an isomorphism on each fibre, the constant fields span each tangent space. It follows that their flows at a point $\widehat{x} \in \widehat{M}$ fully explore some neighbourhood of this point.

Proposition 2.5.2. *Let $f : U \rightarrow V$ be a local isomorphism between two open and connected subspace of M , then f is uniquely determined by the value of its lift at an arbitrary point \widehat{x} in $\pi^{-1}(U)$.*

Proof. Let F denote the lift of f , and suppose $F(\widehat{x})$ is known. Since F is a bundle map $F(\widehat{x} \cdot h)$ is thus uniquely determined by $F(\widehat{x})$. Now note that F preserves the constant vector fields due to $F^*(\omega) = \omega$ and as such it also preserves their flow lines, which explore a neighbourhood of any point. If we know $F(\widehat{x})$ we know that F must map the flow lines through \widehat{x} to those at $F(\widehat{x})$, which uniquely determines F in a neighbourhood of \widehat{x} , and then also on a neighbourhood of $\pi^{-1}(\pi(\widehat{x}))$. By connectedness of U the proposition follows. \square

Definition 2.5.3 (Killing fields). A (locally defined) vector field X on M is a (local) Killing field if its flow is by (local) automorphisms of M . The sheaf of local Killing fields on M is denoted with \mathfrak{kll} .

We remark that \mathfrak{kll} is indeed a sheaf, since it is a sub-pre-sheaf of the sheaf of vector fields (thus a section is uniquely determined by its stalks) and since the property “flows by local automorphisms” is a local property (thus piecing together local Killing fields results again in a local Killing field).

Since the the flow of (local) automorphisms can be uniquely lifted to a flow on the bundle, Killing fields also admit unique lifts to vector fields on \widehat{M} :

¹The set doesn’t appear to have a name in the literature.

Proposition 2.5.4. *Let X be a Killing field on an open and connected subset $U \subset M$ and \widehat{X} its lift. X is uniquely determined by the value of \widehat{X} at an arbitrary point $\widehat{x} \in \pi^{-1}(U)$.*

Proof. It is enough to show that if \widehat{X} vanishes at a point then X must be the zero vector field. If \widehat{X} vanishes at a point, then the flow line at this point must be stationary. It follows that the associated one-parameter group of automorphisms always map this point into itself, by the preceding proposition these automorphisms must be the identity automorphisms and thus \widehat{X} is the zero vector field. \square

It is of interest that $\mathfrak{kill}(U)$ is a Lie-algebra, with bracket given by the vector field commutator. In order to get this result we first characterise how a field on \widehat{M} flowing by automorphisms must look like.

Lemma 2.5.5. *A vector field X defined on a neighbourhood of the form $\pi^{-1}(U) \subset \widehat{M}$ for $U \subset M$ open flows by local automorphisms if and only if it commutes with the constant fields and $D_{\widehat{x}}R_g(X_{\widehat{x}}) = X_{\widehat{x}g}$ for all $g \in P$.*

Proof. The condition $D_{\widehat{x}}R_g(X_{\widehat{x}}) = X_{\widehat{x}g}$ implies that if the flow $\varphi_X^t(\widehat{x})$ of X at \widehat{x} for time t is defined, then the flow $\varphi_X^t(\widehat{x}g)$ at $\widehat{x}g$ is also defined at time t for all g , so one can actually talk about the flow being by bundle-automorphisms. Further $\frac{d}{dt}\varphi_X^t(\widehat{x})g|_{t=0} = D_{\widehat{x}}R_g(X_{\widehat{x}})$ and $\frac{d}{dt}\varphi_X^t(\widehat{x}g)|_{t=0} = X_{\widehat{x}g}$, so the flow is by (local) bundle-automorphisms if and only if $D_{\widehat{x}}R_g(X_{\widehat{x}}) = X_{\widehat{x}g}$ for all $g \in P$.

Now let $C^i = \omega^{-1}(e^i)$ for e^i a basis for \mathfrak{g} , then we may write $\omega = \sum_i e^i \otimes (C^i)^*$ where $((C_{\widehat{x}}^1)^*, \dots, (C_{\widehat{x}}^n)^*)$ is the basis of $T_{\widehat{x}}^*\widehat{M}$ dual to $(C_{\widehat{x}}^1, \dots, C_{\widehat{x}}^n)$. One can easily check $\mathcal{L}_X(\omega) = \sum_i e^i \otimes (\mathcal{L}_X(C^i))^*$, making clear that X flows automorphisms if and only if $[X, C^i] = 0$ for all i , hence by linearity of the Lie bracket if and only if X commutes with every constant field. \square

Corollary 2.5.6. *For any $U \subset M$ the set of vector fields on $\pi^{-1}(U)$ flowing by local automorphisms forms a Lie-algebra with bracket given by the vector field commutator.*

Proof. We only need to check that the two conditions from Lemma 2.5.5 remain true upon taking the commutator. But the first condition follows from $DR_g([X, Y]) = [DR_g(X), DR_g(Y)]$, since R_g is a diffeomorphism. The second follows from the Jacobi identity: $[[X, Y], C] = [[X, C], Y] + [[C, Y], X]$. \square

This allows us to see that \mathfrak{kill} is a Lie-algebra:

Proposition 2.5.7. *If $U \subset M$ is open and $X, Y \in \mathfrak{kill}(U)$, then $[X, Y] \in \mathfrak{kill}(U)$.*

Proof. If \widehat{X}, \widehat{Y} denote the lifts of X, Y to $\pi^{-1}(U)$ then $D\pi(\widehat{X}) = X, D\pi(\widehat{Y}) = Y$ by definition. In particular $D\pi(\widehat{X}), D\pi(\widehat{Y})$ are actually well defined vector fields on U . By the previous corollary $\widehat{Z} := [\widehat{X}, \widehat{Y}]$ flows by automorphisms. One consequence is that it is right-invariant and then $Z := D\pi(\widehat{Z})$ is well defined. The flow of \widehat{Z} is a (bundle-)automorphism over the flow of Z , hence the flow of Z is by automorphisms. In particular $Z \in \mathfrak{kil}(U)$. But:

$$Z = D\pi(\widehat{Z}) = D\pi([\widehat{X}, \widehat{Y}]) = [D\pi(\widehat{X}), D\pi(\widehat{Y})] = [X, Y],$$

giving the proposition. □

We now investigate the dimension of $\mathfrak{kil}(U)$ as well as properties of the stalks \mathfrak{kil}_x .

Proposition 2.5.8. *For any connected open U one has $\dim(\mathfrak{kil}(U)) \leq \dim(\mathfrak{g})$. Also the group of automorphisms of U is finite dimensional and has dimension $\leq \dim(\mathfrak{g})$.*

Proof. The first point is clear since Killing fields are uniquely determined by a vector in a tangent space $T_x \widehat{M} \cong \mathfrak{g}$. The second follows since the action of the automorphism group on \widehat{M} is smooth, as $\text{Diffeo}(\widehat{M})$ acts smoothly on \widehat{M} and $\text{Aut}(M) \subset \text{Diffeo}(\widehat{M})$ is a closed subgroup. Then the one-parameter subgroups of $\text{Aut}(M)$ must be generated by Killing fields. □

Proposition 2.5.9. *For any $x \in M$ there exists a connected neighbourhood U_x so that $\mathfrak{kil}(U_x)$ is equal to (via the restriction map) to the stalk \mathfrak{kil}_x .*

Proof. Let U_α be a filtration of connected neighbourhoods of x converging to x . By virtue of Proposition 2.5.4 (or just general considerations with Killing fields) the restriction of local Killing fields to a smaller open set (both connected) is injective, i.e. $\mathfrak{kil}(U_\alpha)$ gets bigger as U_α gets smaller. By virtue of Proposition 2.5.8 these algebras cannot grow arbitrarily in dimension, and thus eventually the process stops. Thus there are neighbourhoods U_x of x so that $\mathfrak{kil}_x = \mathfrak{kil}(U_x)$. □

Proposition 2.5.10. *The dimension of \mathfrak{kil}_x can only increase locally, that is for every $x \in M$ there is an open neighbourhood U with $\dim(\mathfrak{kil}_y) \geq \dim(\mathfrak{kil}_x)$ for all $y \in U$.*

Proof. Let U_x be one of the sets from Proposition 2.5.9, then the composition

$$\mathfrak{kil}_x \xrightarrow{\cong} \mathfrak{kil}(U_x) \hookrightarrow \mathfrak{kil}_y$$

gives an injection from \mathfrak{kil}_x to \mathfrak{kil}_y . □

We will now describe a decomposition of the manifold M into open sub-manifolds \mathfrak{M}_i (with $\bigcup_i \mathfrak{M}_i$ being dense in M). This decomposition is special, because when restricted to \mathfrak{M}_i the sheaf

$\mathfrak{kil}|_{\mathfrak{M}_i}$ is locally constant. After 2.5.13 we remark on some of the consequences and uses of this decomposition.

Definition 2.5.11 (\mathfrak{M} and \mathfrak{M}_i). Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan connection modelled on G/P . Define

$$\mathfrak{M} := \{x \in M \mid \exists V_x \subseteq M \text{ open s.t. } r : \mathfrak{kil}(V_x) \rightarrow \mathfrak{kil}_y \text{ is an isomorphism } \forall y \in V_x\}$$

and denote with \mathfrak{M}_i the connected components of \mathfrak{M} .

By definition \mathfrak{M} is open, the following proposition shows that every open set contains points from \mathfrak{M} , hence \mathfrak{M} is also dense:

Proposition 2.5.12. *Let U be a non-empty open set. There exists a non-empty open set $V \subset U$ so that for any $x \in V$ the restriction map $\mathfrak{kil}(V) \rightarrow \mathfrak{kil}_x$ is an isomorphism.*

Proof. Let $z \in U$ and consider an open set U_z as given by Proposition 2.5.9, we may assume $U_z \subset U$ by intersecting the two and choosing a new U'_z to be the component containing z . Since the dimension of \mathfrak{kil}_y for $y \in U_z$ is bounded by Proposition 2.5.9, we choose a point $x \in U_z$ so that \mathfrak{kil}_x has maximal dimension. For this x we choose an open set U_x contained in U_z and having $\mathfrak{kil}(U_x) = \mathfrak{kil}_x$ as in Proposition 2.5.9. If y is another point in U_x then $\dim(\mathfrak{kil}_y) \leq \dim(\mathfrak{kil}_x) = \dim(\mathfrak{kil}(U_x))$ by construction, however the restriction $\mathfrak{kil}(U_x) \rightarrow \mathfrak{kil}_y$ must be injective, whence it must be an isomorphism. U_x then serves as the set V in the proposition. \square

Corollary 2.5.13. $\mathfrak{M} \subseteq M$ is open and dense.

It is immediate that the restriction $\mathfrak{kil}|_{\mathfrak{M}_i}$ of \mathfrak{kil} to \mathfrak{M}_i is a locally constant sheaf. A consequence of this is that any Killing field may be uniquely extended along a path in \mathfrak{M}_i , although some monodromy may appear (similar to how the complex logarithm may be uniquely extended along a path in $\mathbb{C} - \{0\}$, possibly with monodromy). In this way the components \mathfrak{M}_i form an arena in which the local Killing fields act as if everything were analytic², and thus act as useful places in which to begin constructions, by denseness of \mathfrak{M} one can also often return to such a component in a construction.

²This sentence can be understood in two ways. Firstly the Killing fields themselves behave like holomorphic functions in the way just described, secondly in [No 60] it is shown that on an analytic Riemannian manifold the sheaf of local Killing fields is locally constant (albeit with different words, cf. Theorem 2 of [No 60]). As remarked before the individual components \mathfrak{M}_i also enjoy this property, hence one could say \mathfrak{M}_i looks like an analytic manifold to \mathfrak{kil} . At this point is appropriate to remark that Proposition 2.5.12 is also contained in [No 60], which seems to be the origin of Corollary 2.5.13.

Another consequence of $\mathfrak{kill}|_{M_i}$ being locally constant is that the isomorphism class \mathfrak{kill}_x is an invariant of the component \mathfrak{M}_i . What kind of components \mathfrak{M}_i may appear, characterised by their local Killing algebras, provides an additional lever in a classification effort.

Finally we make some remarks comparing this decomposition with that given by the integrability locus (which is defined and discussed in Chapter 4)³. In comparing the consequences of the integrability theorem one notes that $M^{\text{int}} \subseteq \mathfrak{M}$ (and that M^{int} is also open and dense), but the reverse inclusion is not true in general⁴. The main difference between \mathfrak{M} and M^{int} is on the one hand the explicit characterisation of $M^{\text{int}} = CR(\mathcal{D}\kappa)$ (cf. Definition 4.2.3 and Theorem 4.3.2), which makes working with it easier, and the very useful property that every infinitesimal symmetry of $\mathcal{D}\kappa$ integrates to a Killing field on M^{int} , providing a canonical isomorphism of vector spaces $\ker(D_{\hat{x}}\mathcal{D}\kappa) \cong \mathfrak{kill}_{\pi(\hat{x})}$ inside of M^{int} .

2.6 The curvature of a Cartan connection

In Section 2.3 we had noticed, albeit without having introduced the notion of a Cartan connection, that a Cartan geometry can be “glued” by pieces of the model space if and only if $d\omega + \frac{1}{2}[\omega, \omega]_{\mathfrak{g}} = 0$. On the other hand if this expression is not zero then the method of rolling the model space on the manifold described in Section 2.3.2 will depend on more than just the homotopy class of the path along which we roll, motivating that such spaces have non-zero curvature. Indeed $d\omega + \frac{1}{2}[\omega, \omega]_{\mathfrak{g}}$ is the usual definition of the curvature of the Cartan connection ω .

Definition 2.6.1 (Curvature). If $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ is a Cartan connection on a manifold M modelled on G/P , then the curvature of ω is defined to be the \mathfrak{g} -valued 2-form

$$\kappa = d\omega + \frac{1}{2}[\omega, \omega]_{\mathfrak{g}}.$$

In this section we briefly remark that the curvature is P -equivariant and zero in vertical directions.

Proposition 2.6.2. *The curvature is equivariant under the action of P , meaning that*

$$R_g^*(\kappa) = \text{Ad}(g^{-1}) \cdot \kappa.$$

Proof. The compatibility of ω states $R_g^*(d\omega) = d(R_g^*\omega) = d(\text{Ad}(g^{-1}) \cdot \omega) = \text{Ad}(g^{-1})d\omega$. Together with $[\text{Ad}(g^{-1})v, \text{Ad}(g^{-1})w]_{\mathfrak{g}} = \text{Ad}(g)^{-1}[v, w]_{\mathfrak{g}}$ for all $v, w \in \mathfrak{g}$ we conclude $R_g^*(\kappa) = R_g^*d\omega + R_g^*\frac{1}{2}[\omega, \omega]_{\mathfrak{g}} = \text{Ad}(g^{-1})\kappa$. \square

³In order for this paragraph to make sense one must have already read Chapter 4, however these comments seem more appropriate here than there.

⁴Without providing a calculation we remark that $dx_1^2 + (1 + x_1^6)dx_2^2$ is a Riemannian metric on \mathbb{R}^2 for which the inclusion is strict. For this case $\mathfrak{M} = \mathbb{R}^2$ but $M^{\text{int}} = \mathbb{R}^2 - (\{0\} \times \mathbb{R})$.

We strengthen this result:

Proposition 2.6.3. *Let $f : \widehat{M} \rightarrow P$ be smooth and let $F : \widehat{M} \rightarrow \widehat{M}, \widehat{x} \mapsto \widehat{x} \cdot f(\widehat{x})$. Then $F^*(\kappa) = \kappa$.*

For the proof we refer to Lemma 5.3.9 in [Sh 97]. We can use this to find the following result:

Proposition 2.6.4. *The curvature is 0 in vertical directions, meaning if $\widehat{x} \in \widehat{M}, v, w \in T_{\widehat{x}}M$ and $D_{\widehat{x}}\pi(v) = 0$ then $\kappa_{\widehat{x}}(v, w) = 0 = \kappa_{\widehat{x}}(w, v)$.*

Proof. We assume that v is tangent to the fibre and choose a map $f : \widehat{M} \rightarrow P$ with $f(\widehat{x}) = 1$ and $\omega_{\widehat{x}}(v) = -D_{\widehat{x}}f(v)$. Then defining $F : \widehat{M} \rightarrow \widehat{M}$ by $\widehat{y} \mapsto \widehat{y} \cdot f(\widehat{y})$ we find, using the same calculation as in Lemma 2.4.3, that

$$F^*(\omega)_{F(\widehat{y})} = \text{Ad}(f(\widehat{x})^{-1})\omega_{\widehat{y}} + f^*(\omega_{MC})_{f(\widehat{y})},$$

implying $(\omega_{\widehat{x}}(D_{\widehat{x}}F(v)) = \omega_{\widehat{x}}(v) + D_{\widehat{x}}f(v) = 0$, whence $D_{\widehat{x}}F(v) = 0$.

However by Proposition 2.6.3 we know $F^*(\kappa) = \kappa$, so we find:

$$\kappa_{\widehat{x}}(v, w) = \kappa_{\widehat{x}}(D_{\widehat{x}}F(v), D_{\widehat{x}}F(w)) = \kappa_{\widehat{x}}(0, D_{\widehat{x}}F(w)) = 0,$$

implying the proposition. □

Chapter 3

Examples of Cartan geometries

3.1 Parallelisms

The simplest example of a Cartan geometry is that of a parallelism. Parallelisms are encoded by those geometries in which the stabiliser group P in the model geometry G/P is the trivial group. For many geometric considerations almost all of the details of the group G are superfluous in a way that will be made more precise in this section (cf. Lemma 3.1.4 and Proposition 3.1.5).

Additionally studying parallelisms is of interest, as every Cartan connection $(\widehat{M} \rightarrow M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ induces a parallelism on \widehat{M} . The geometry of this parallelism is connected to the geometry of the original Cartan connection. In particular the local Killing algebras of the original Cartan connection and the induced parallelism are canonically isomorphic and the two curvature functions are the same (cf. Corollary 3.1.8 and Proposition 3.1.9).

Definition 3.1.1 (V -parallelism). Let M be an n -dimensional manifold and V an n -dimensional vector space. A V -parallelism of M is a smooth vector-bundle isomorphism from TM into the trivial bundle $M \times V$.

Such a vector-bundle isomorphism is equivalent to the data of a smooth V -valued 1-form, which is an isomorphism at each tangent space, one often uses the name “parallelism” interchangeably for these two notions.

An important example of a parallelism is the Maurer-Cartan form:

Proposition 3.1.2. *Let G be a Lie group, then the Maurer-Cartan form $\omega_G : TG \rightarrow \mathfrak{g}$, which maps a vector $v \in T_g G$ to its left-translate $D_g L_{g^{-1}}[v]$ in $T_1 G = \mathfrak{g}$, is a \mathfrak{g} -parallelism of G .*

Proof. We carry out the details, verifying that the map is smooth and an isomorphism at each tangent space, then describe the induced trivialisation $TG \cong G \times \mathfrak{g}$.

To see that ω_G is smooth, note that the map $\mu : G \times G \rightarrow G, (g, h) \mapsto gh$ is smooth and then so is its differential

$$T(G \times G) \cong TG \times TG \rightarrow TG, \quad (w, v)_{(g, h)} \mapsto D_g R_h [w] + D_h L_g [v].$$

In particular, pre-composing the above differential with a map $TG \rightarrow T(G \times TG)$, where an element $v \in T_g G$ is mapped to $(0, v)_{(g^{-1}, g)} \in T_{(g^{-1}, g)}(G \times G)$ is smooth. The resulting map is however is nothing other than $TG \rightarrow TG, v \in T_g G \mapsto D_g L_{g^{-1}} [v]$, which is the Maurer-Cartan form ω_G .

ω_G is an isomorphism on each tangent space because for each $g \in G$ the map $G \rightarrow G, h \mapsto L_{g^{-1}}(h) = g^{-1}h$ is a diffeomorphism, and $\omega_G|_{T_g G} = D_g L_{g^{-1}}$ is the differential of that diffeomorphism at g .

Then \mathfrak{g} -parallelism is then nothing other than the map $TG \rightarrow G \times \mathfrak{g}, v \in T_g G \mapsto (g, \omega_G(v))$, its inverse is $G \times \mathfrak{g} \rightarrow TG, (g, A) \mapsto (\omega_G)_g^{-1}(A)$. \square

Proposition 3.1.3. *If $P = \{1\}$ is the trivial group, G an n -dimensional Lie group with Lie algebra \mathfrak{g} and M an n -dimensional manifold, then the data of Cartan connection on M modelled on the space G/P is the same as a \mathfrak{g} -parallelism of TM .*

Proof. A Cartan connection on M modelled on G/P is given by principal bundle \widehat{M} over M with structure group P and a \mathfrak{g} -valued 1-form ω so that:

1. $\omega_{\widehat{x}} : T_{\widehat{x}} \widehat{M} \rightarrow \mathfrak{g}$ is a linear isomorphism for any $\widehat{x} \in \widehat{M}$.
2. The fundamental fields of the bundle \widehat{M} are mapped to their generators in $\text{Lie}(P)$.
3. For any $h \in P$: $\omega_{\widehat{x} \cdot h} \circ D_{\widehat{x}} R_h = \text{Ad}(h^{-1}) \omega_{\widehat{x}}$, where R_h is the right-multiplication with h .

However since $P = \{1\}$ the only P -principal bundle over M is M itself with the identity map as projection. Further the only fundamental field is 0, which is conveniently mapped to $0 \in \text{Lie}(P)$ by any \mathfrak{g} -valued 1-form and thus the second condition is always true. The third condition must only be checked for $h = 1$, where it becomes $\omega_{\widehat{x}} = \omega_{\widehat{x}}$ for all \widehat{x} , which is a tautology.

Thus a Cartan connection on M modelled on G/P is the same as a \mathfrak{g} -valued 1-form $\omega : TM \rightarrow \mathfrak{g}$ which is a linear isomorphism on each tangent space, in other words it is the same as \mathfrak{g} -parallelism. \square

Example (\mathbb{R}^n -parallelism). An \mathbb{R}^n -parallelism of a manifold M is the same as a tuple of smooth vector fields (b_1, \dots, b_n) on TM that is a basis when restricted to every tangent space, in other words a global frame. For if (b_1, \dots, b_n) is such a tuple and denoting with (e^1, \dots, e^n) the canonical basis of \mathbb{R}^n , then the map $TM \rightarrow \mathbb{R}^n$, $v = \sum_i v_i b_{i,x} \in T_x M \mapsto \sum_i v_i e^i$ defines a smooth \mathbb{R}^n valued 1-form that is an isomorphism on each tangent space, and for each such form $\eta : TM \rightarrow \mathbb{R}^n$ the tuple $(\eta_x^{-1}(e^1), \dots, \eta_x^{-1}(e^n))$ defines a smoothly varying basis of each tangent space $T_x M$.

3.1.1 Geometry of a parallelism

Let M be a manifold, G a Lie group with Lie algebra \mathfrak{g} and $\omega : TM \rightarrow \mathfrak{g}$ a \mathfrak{g} -parallelism. For any real vector space V that has the same dimension as \mathfrak{g} there exist linear isomorphisms between \mathfrak{g} and V , and if one pre-composes such an isomorphism with ω one retrieves a V -parallelism of M .

In particular if we have another Lie group G' of the same dimension as G , then the \mathfrak{g} -parallelism, which is the same as a $G/\{1\}$ Cartan connection on M , induces a $\mathfrak{g}' = \text{Lie}(G')$ -parallelism on M , i.e. a $G'/\{1\}$ Cartan connection. The following lemma makes clear that the geometries of these two parallelisms are very similar:

Lemma 3.1.4. *Let $\omega : TM \rightarrow \mathfrak{g}$ be a Cartan connection modelled on $G/\{1\}$ and $\omega' : TM \rightarrow \mathfrak{g}'$ the Cartan connection modelled on $G'/\{1\}$ that is induced by ω and a linear isomorphism $\mathfrak{g} \rightarrow \mathfrak{g}'$. Then the local automorphisms of ω are the same as the local automorphisms of ω' .*

Proof. As remarked in Definition 2.4.6 a smooth map $f : U \rightarrow V$ with $U, V \subset M$ open, is a local automorphism of a Cartan geometry $\omega : \widehat{TM} \rightarrow \mathfrak{g}$ if and only if f lifts to a bundle isomorphism $F : \pi^{-1}(U) \rightarrow \pi^{-1}(V)$ so that $F^*(\omega|_{\pi^{-1}(U)}) = \omega|_{\pi^{-1}(V)}$. Here there is no bundle, so the only condition is $f^*(\omega) = \omega$. Then clearly

$$f^*(A \circ \omega) = A \circ \omega \circ Df = A \circ f^*(\omega)$$

and we have $f^*(\omega) = \omega \iff f^*(A \circ \omega) = A \circ \omega$ for an invertible A . □

Continuing in this direction we compare the curvature of two parallelisms $\omega : TM \rightarrow \mathfrak{g}$ and $\omega' : TM \rightarrow \mathfrak{g}'$ related in this way, that is there is a linear isomorphism $A : \mathfrak{g} \rightarrow \mathfrak{g}'$ so that $\omega'_x = A \circ \omega_x$ for all $x \in M$. The result is that these two curvatures are the same “up to a constant”, where the meaning of this needs to be made precise.

The curvature of a Cartan connection ω was defined in Section 2.6 to be the \mathfrak{g} -valued 2-form $\kappa = d\omega + \frac{1}{2}[\omega, \omega]_{\mathfrak{g}}$. If we use the parallelism $\omega : TM \rightarrow \mathfrak{g}$ we can interpret κ as a function on M ,

namely:

$$\kappa : M \mapsto \text{Hom}(\Lambda^2 \mathfrak{g}, \mathfrak{g}), \quad x \mapsto \left[(v, w) \mapsto d\omega_x(\omega_x^{-1}(v), \omega_x^{-1}(w)) + [v, w]_{\mathfrak{g}} \right].$$

Here Hom refers to vector space homomorphism, i.e. linear maps. Doing the same thing for the curvature for ω' gives:

$$\kappa' : M \rightarrow \text{Hom}(\Lambda^2 \mathfrak{g}', \mathfrak{g}'), \quad x \mapsto \left[(v, w) \mapsto A \cdot d\omega_x(\omega_x^{-1}(A^{-1}v), \omega_x^{-1}(A^{-1}w)) + [v, w]_{\mathfrak{g}'} \right].$$

Now since A is an isomorphism of vector spaces it induces an isomorphism of $A^* : \text{Hom}(\Lambda^2 \mathfrak{g}', \mathfrak{g}') \rightarrow \text{Hom}(\Lambda^2 \mathfrak{g}, \mathfrak{g})$, where for $\gamma' \in \text{Hom}(\Lambda^2 \mathfrak{g}', \mathfrak{g}')$ and $v, w \in \mathfrak{g}$ one defines

$$A^*(\gamma')(v, w) = A^{-1} \cdot \gamma'(Av, Aw).$$

With this isomorphism we can directly compare κ and κ' , for example by looking at the difference $\kappa(x) - A_*(\kappa'(x))$:

$$\begin{aligned} (\kappa(x) - A^*(\kappa'(x)))(v, w) &= d\omega_x(\omega_x^{-1}(v), \omega_x^{-1}(w)) + [v, w]_{\mathfrak{g}} - d\omega_x(\omega_x^{-1}(A^{-1}v), \omega_x^{-1}(A^{-1}w)) - [Av, Aw]_{\mathfrak{g}'} \\ &= [v, w]_{\mathfrak{g}} - [Av, Aw]_{\mathfrak{g}'}. \end{aligned}$$

We note that this difference is constant in x , meaning that the function $\kappa - A^*(\kappa') : M \rightarrow \text{Hom}(\Lambda^2 \mathfrak{g}, \mathfrak{g})$ does not depend on the point. The previous discussion yields the following proposition:

Proposition 3.1.5. *The expression $\kappa(x) - A^*(\kappa'(x))$ does not depend on the point x . That is: Changing the target space of the parallelism by composing with a linear isomorphism A only alters the curvature by a constant element of $\text{Hom}(\Lambda^2 \mathfrak{g}, \mathfrak{g})$ upon identifying $\text{Hom}(\Lambda^2 \mathfrak{g}, \mathfrak{g})$ and $\text{Hom}(\Lambda^2 \mathfrak{g}', \mathfrak{g}')$ via the linear isomorphism A^* .*

A large amount of geometric quantities that depend on the curvature are not changed if the curvature is modified by a constant function. As an example the integrability locus, defined in Chapter 4, can be characterised as the points where the “generalised curvature” $\mathcal{D}\kappa$ admits a neighbourhood on which it has constant rank (cf. Theorem 4.3.2). For two Cartan connections whose curvature functions differ by a constant the integrability locus is the same.

3.1.2 Every Cartan geometry as a parallelism

Let $(\widehat{M} \rightarrow M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan connection modelled on G/P . By definition $\omega : T\widehat{M} \rightarrow \mathfrak{g}$ is a \mathfrak{g} -parallelism of \widehat{M} . So the original Cartan connection ω_{orig} on M modelled on G/P induces a Cartan connection ω_{par} on \widehat{M} modelled on $G/\{1\}$. We will still denote the actual form $T\widehat{M} \rightarrow \mathfrak{g}$ with ω in both cases.

In this section we will investigate the relationship between the geometrical properties of ω_{orig} and ω_{par} .

Remark. Let $F : \widehat{M} \rightarrow \widehat{M}$ be a diffeomorphism, F is an automorphism of ω_{par} if and only if $F^*(\omega) = \omega$ and F induces an automorphism $M \rightarrow M$ of ω_{orig} if and only if F is a bundle-automorphism and $F^*(\omega) = \omega$. Every automorphism of ω_{orig} is then an automorphism of ω_{par} .

The first question would be: under what circumstances does $F^*(\omega) = \omega$ already imply that F is a bundle-automorphism? If P were connected it would follow:

Proposition 3.1.6. *Let $F : \widehat{M} \rightarrow \widehat{M}$ be a diffeomorphism with $F^*(\omega) = \omega$. Denote with P^0 the connected component containing identity of P , so that \widehat{M} is a P^0 bundle over $\tilde{M} = \widehat{M}/P^0$. Then F is a P^0 -bundle-automorphism.*

Proof. We have to check that $F(\widehat{x}p) = F(\widehat{x})p$ for any $p \in P^0$, $\widehat{x} \in \widehat{M}$. But since ω is a Cartan connection modelled on G/P we must have by property 3. of Definition 2.4.1 that the vertical ω -constant fields are the fundamental fields of the P -principal bundle \widehat{M} . Since $F^*(\omega) = \omega$ the map F must preserve the fundamental fields and thus their flows. But these flows are $t \mapsto \widehat{x} \cdot \exp(t\xi)$, so we find:

$$F(\widehat{x} \cdot \exp(t\xi)) = F(\widehat{x}) \cdot \exp(t\xi)$$

for all $\widehat{x} \in \widehat{M}$, $t \in \mathbb{R}$ and $\xi \in \mathfrak{p}$. This implies that F is a morphism of P^0 -bundles. \square

This is enough to show that any local automorphism of ω_{par} induces a local automorphism of ω_{orig} when restricted to a small enough open set:

Proposition 3.1.7. *Let $\widehat{x} \in \widehat{M}$, there exists a connected and arbitrarily small neighbourhood U of \widehat{x} so that every local automorphism $F : U \rightarrow V \subset \widehat{M}$ of ω_{par} induces a well defined ω_{orig} -automorphism $[F] : \pi(U) \rightarrow \pi(V)$.*

Proof. We let U' be small enough so that $\pi(U')$ has a section $s : \pi(U') \rightarrow U'$. We then let $U = s(U') \cdot P^0$ (which is open since P^0 is open in P). Any local automorphism F of ω_{par} defined on U must be a P^0 -bundle-automorphism as noted in the preceding proposition. In particular it follows that $[F] : \pi(U) \rightarrow \pi(V)$, $\pi(\widehat{x}) \mapsto \pi(F(\widehat{x}))$ is well defined. Further we can extend F to $U \cdot P = \pi^{-1}(\pi(U))$ simply by setting $F(\widehat{x} \cdot p) = F(\widehat{x})$, this preserves $F^*(\omega) = \omega$ (the calculation behind this step will not be carried out here but is identical to the calculation in Theorem 2.3.6). This means that $[F] : \pi(U) \rightarrow \pi(V)$ lifts to a P -bundle-automorphism and thus that $[F]$ is an ω_{orig} -automorphism. \square

In particular the ω_{par} -Killing fields on such sets U all flow by ω_{orig} -automorphisms. This gives the following useful result:

Corollary 3.1.8. *Let $\hat{x} \in \widehat{M}$, then $\mathfrak{kil}_{\hat{x}}^{\omega_{\text{par}}} \cong \mathfrak{kil}_{\pi(\hat{x})}^{\omega_{\text{orig}}}$, where the isomorphism is given by $X \mapsto D\pi[X]$ for X a Killing field in a small enough neighbourhood of \hat{x} .*

This will be used in the proof of the Integrability Theorem, as it implies it is no restriction to assume that the Cartan-connection ω is that of a parallelism when looking at germs of Killing fields.

We make one more comment:

Proposition 3.1.9. *Let κ^{par} be the curvature of ω_{par} and κ^{orig} the curvature of ω_{orig} . Then $\kappa^{\text{par}} = \kappa^{\text{orig}}$.*

Proof. This follows immediately since the definition of the curvature:

$$\kappa = d\omega + \frac{1}{2}[\omega, \omega]_{\mathfrak{g}},$$

which only depends on ω and not on properties of the bundle $\pi : \widehat{M} \rightarrow M$. □

3.2 Pseudo-Riemannian geometry

Pseudo-Riemannian geometry is a more complicated example of a Cartan geometry than that of parallelisms. This kind of geometry can be encoded in a Cartan connection modelled on the space G/P with $G = O(p, q) \times \mathbb{R}^n$ and $P = O(p, q)$. Being more precise a pseudo-Riemannian geometry of signature (p, q) on a manifold M may be understood as a pseudo-Riemannian metric g on M together with a metric connection, that is a covariant derivative ∇ so that $\nabla_X(g) = 0$ for all X .

Such a structure always induces a Cartan connection modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$ on M and is itself induced by such a Cartan connection. In this section we will describe a bijection:

$$\left\{ (g, \nabla), \text{ with } g \text{ a pseudo-Riemannian metric on } M \text{ and } \nabla \text{ a metric connection.} \right\} \leftrightarrow \left\{ \begin{array}{l} (\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n) \text{ a} \\ \text{Cartan connection on } M \text{ modelled on} \\ O(p, q) \times \mathbb{R}^n / O(p, q) \text{ (up to isomorphism)} \\ \text{over id : } M \rightarrow M. \end{array} \right\}$$

(Here we must identify any two Cartan connections if there is an isomorphism between them over the identity, this is necessary since if $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ is a Cartan connection and X is a set with a bijection $b : X \rightarrow \widehat{M}$, then we can pull all smooth, bundle and connection structures back onto X via b , giving a different Cartan connection on M describing exactly the same thing. Both g and ∇ do not have this “renaming freedom”.)

One should remark that in most contexts when one talks about a pseudo-Riemannian geometry one has as data only a metric g of signature (p, q) on M , and not an explicit metric connection. In this setting the Levi-Civita connection ∇^{LC} is usually canonically associated to g in order to give a covariant derivative on M . So to make the above bijection more compatible with practical considerations, we will describe the image of (g, ∇^{LC}) under the bijection:

$$\left\{ \begin{array}{l} (g, \nabla^{LC}), g \text{ a pseudo-Riemannian metric} \\ \text{on } M \text{ and } \nabla^{LC} \text{ the Levi-Civita connec-} \\ \text{tion of } g. \end{array} \right\} \leftrightarrow \left\{ \begin{array}{l} (\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n) \\ \text{a torsionless Cartan connection on } M \\ \text{modelled on } O(p, q) \times \mathbb{R}^n / O(p, q) \text{ (up} \\ \text{to isomorphism over } \text{id} : M \rightarrow M). \end{array} \right\}$$

The definition of when a Cartan connection modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$ is torsionless will be given in Section 3.2.6.

While these statements establish a correspondence between pseudo-Riemannian geometries and Cartan connections, this does not necessarily imply that the geometric content of the two notions is the same. To that end we additionally show that the isomorphisms of a Cartan connection are the same as isometries preserving the metric connection of the associated pseudo-Riemannian geometry. (In the case of the Levi-Civita connection any isometry preserves ∇^{LC} , thus here we show that isomorphisms between torsionless Cartan connections are described by isometries of the associated metrics.) This correspondence is in fact functorial, and will give an equivalence of categories.¹

The organisation is as follows: Section 3.2.1 briefly reviews the needed conventions and definitions of pseudo-Riemannian geometry, while Section 3.2.2 reviews the notion of an Ehresmann connection on a principal P -bundle and describes how a metric connection induces an Ehresmann connection on the orthonormal frame bundle. Section 3.2.3 puts these two sections together to construct the Cartan connection induced by a pseudo-Riemannian metric g and a metric connection ∇ . Section 3.2.4 checks that every isometry preserving the metric connection induces an isomorphism of the associated Cartan connections and vice versa, as well as checking the functoriality of this correspondence. Section 3.2.5 turns the process around and explains how an $O(p, q) \times \mathbb{R}^n / O(p, q)$ Cartan connection on a manifold M induces a pseudo-Riemannian metric and a metric connection on M . Section 3.2.6 then checks that these two constructions are inverse to each other, as well as checking that the case $\nabla = \nabla^{LC}$ corresponds to the torsionless Cartan connections.

¹Between the category [pseudo-Riemannian manifolds of signature (p, q) and metric connections with morphisms being isometries preserving the covariant derivative] and the category [Cartan connections modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$ with morphisms being isomorphisms of the connection]. Alternatively between the category [pseudo-Riemannian manifolds of signature (p, q) with morphisms being isometries] and the category [torsionless Cartan connections modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$ with morphisms being isomorphisms of Cartan connections].

3.2.1 Conventions

In this section we remark on notational conventions, for example on the ordering of an ONB of a symmetric bilinear form of signature (p, q) or the notation we adopt for the parallel transport, and describe the orthonormal frame bundle of a pseudo-Riemannian manifold. It is not the purpose of this section to provide an introduction to pseudo-Riemannian geometry; properly defining all relevant notions from this theory would exceed its intended scope.

Metric and metric connections

If g is a symmetric bilinear form on a vector space V let p be the maximal dimension of all subspaces of V on which g is positive definite and q the maximal dimension of subspaces on which g is negative definite. Then the signature of g is (p, q) , the important point being that the positive dimensions come first. We will denote with $O(p, q)$ the group of linear isometries of the standard bilinear form of signature (p, q) on \mathbb{R}^n , where $n = p + q$.

Definition 3.2.1 (Pseudo-Riemannian metric). Let M be a manifold of dimension $p + q$, a pseudo-Riemannian metric of signature (p, q) on M together is a smooth non-degenerate bilinear form $TM \otimes TM \rightarrow \mathbb{R}$ of signature (p, q) .

If g is a pseudo-Riemannian metric on a manifold M , then a covariant derivative ∇ is called a metric connection if $\nabla_X(g) = 0$ for all vector fields X . This is equivalent to the notion of parallel transport induced by ∇ to be by isometries. We briefly recall how parallel transport is defined.

Definition 3.2.2 (Parallel transport by ∇). If $\gamma : (-\epsilon, \epsilon) \rightarrow M$ is a path, ∇ a linear connection on M and $v \in T_{\gamma(a)}M$ a vector, then we define $P_a^b(t \mapsto \gamma(t)) [v_a] \in T_{\gamma(b)}M$ to be $v(b)$, where $v(t)$ is the solution of the ODE $\nabla_{\dot{\gamma}(t)}v(t) = 0$ with boundary condition $v(a) = v_a$.

This makes parallel transport a linear isomorphism $P_a^b(t \mapsto \gamma(t)) : T_{\gamma(a)}M \rightarrow T_{\gamma(b)}M$. By use of the formula $X(g(Y, Z)) = (\nabla_X g)(Y, Z) + g(\nabla_X Y, Z) + g(Y, \nabla_X Z)$ one notes that ∇ is a metric connection if and only if $P_a^b(t \mapsto \gamma(t))$ is an isometry from $T_{\gamma(a)}M \rightarrow T_{\gamma(b)}M$ for all smooth paths γ and a, b in the domain of γ . In fact the parallel transport uniquely determines ∇ , for if v^1, \dots, v^n is a basis of $T_{\gamma(a)}M$ and we parallel transport this along γ to a basis $v^1(t), \dots, v^n(t)$ then for any vector field $Y = \sum_i a_i(t)v^i(t)$ along γ we have:

$$\nabla_{\dot{\gamma}(t)}Y = \nabla_{\dot{\gamma}(t)} \sum_i a_i(t)v^i(t) = \sum_i \dot{a}_i(t)v^i(t) + \sum_i a_i(t)\nabla_{\dot{\gamma}(t)}v^i(t) = \sum_i \dot{a}_i(t)v^i(t). \quad (3.1)$$

We will have cause to refer to this equation again in the coming sections.

The Levi-Civita connection is a special metric connection. It is defined as the unique torsion free covariant derivative preserving the metric. Torsion free means in this case that $\nabla_X Y - \nabla_Y X -$

$[X, Y] = 0$ for all vector fields X, Y . This condition together with symmetry of g and $\nabla g = 0$ imply:

$$g(\nabla_X Y, Z) = \frac{1}{2} [X(g(Y, Z)) + Y(g(Z, X)) - Z(g(X, Y)) + g([X, Y], Z) - g([Y, Z], X) - g([X, Z], Y)],$$

for all vector fields X, Y, Z which determines $\nabla_X Y$. In particular there can be at most one torsion free metric connection. Since the formula above does indeed define a covariant derivative, any pseudo-Riemannian metric g has a unique torsion free metric connection, called the Levi-Civita connection. We remark that the above formula has a name, it is called the Koszul formula.

Frame bundle and orthonormal basis

For a vector space V with a non-degenerate symmetric bilinear form g of signature (p, q) an orthonormal basis is an ordered collection (b_1, \dots, b_{p+q}) of vectors so that $g(b_i, b_j) = \pm \delta_{ij}$. The sign $+$ is chosen when $i \leq p$, and $-$ if $i > p$, i.e. the “positive vectors come first”. If $U \subseteq M$ is open and (M, g) a pseudo-Riemannian manifold, then an orthonormal frame on U is an ordered collection of smooth vector fields (b_1, \dots, b_n) on U so that $(b_{1,x}, \dots, b_{n,x})$ is an orthonormal basis of $(T_x M, g_x)$ for all $x \in U$.

In the following sections we will denote with \widehat{M} the frame bundle of M , we may understand this to be the set

$$\widehat{M} = \{(x, b_1, \dots, b_n) \mid x \in M, (b_1, \dots, b_n) \text{ is an orthonormal basis of } T_x M\}.$$

Given the correct smooth structure this is a $O(p, q)$ -principal bundle, where the $O(p, q)$ action (on the right) is defined by

$$(x, b_1, \dots, b_n) \cdot g = \left(x, \sum_{i=1}^n g_{i1} b_i, \dots, \sum_{i=1}^n g_{ni} b_i \right).$$

We denote with $\pi : \widehat{M} \rightarrow M$ the projection. We briefly remark on how the structure of an $O(p, q)$ -principal bundle on \widehat{M} is constructed:

Let x be a point in M , then there is an open neighbourhood U of x admitting an orthonormal frame (b_1, \dots, b_n) on U . We then define $\varphi_U : U \times O(p, q) \rightarrow \pi^{-1}(U)$, $(x, g) \mapsto (x, \sum_{i=1}^n g_{i1} b_{i,x}, \dots, \sum_{i=1}^n g_{ni} b_{i,x})$. We then define a smooth structure on $\pi^{-1}(U)$ so that this is a diffeomorphism. If V is another open set with a smooth frame e_1, \dots, e_n , then $b_{i,x} = \sum_j f_{ij}(x) e_{j,x}$ where $f_{ij} : U \cap V \rightarrow \mathbb{R}$ is smooth. Since both e_1, \dots, e_n and b_1, \dots, b_n are orthonormal frames $(f_{ij}(x))_{1 \leq i, j \leq n}$ is a matrix in $O(p, q)$.

It is elementary to check that if d_1, \dots, d_n is a third basis and $e_{i,x} = \sum_{ij} f'_{ij}(x)d_{j,x}$, that then $b_{i,x} = \sum_{ijk} f'_{ij}(x)f_{jk}(x)e_{k,x}$. This means that such charts φ_U induce a system of $O(p, q)$ valued transition functions, defining the structure of an $O(p, q)$ -principal bundle on \widehat{M} .

3.2.2 Ehresmann connection

We briefly review how a pseudo-Riemannian metric g and a metric connection ∇ on M give rise to an Ehresmann connection on the induced frame bundle \widehat{M} . To this end we first define this terminology. Definition 3.2.3 and Proposition 3.2.5 provide two equivalent definitions of an Ehresmann connection. We then define the Ehresmann connection induced by (g, ∇) in the sense of Definition 3.2.3, but for working with it the relevant form will be as in Proposition 3.2.5. We refer to [KN 63] as a reference to this topic, specifically Chapters II and III.

Definition 3.2.3 (Ehresmann connection (horizontal bundle)). Let P be a Lie group, M a manifold and $E \xrightarrow{\pi} M$ be a P -principal bundle over M . Denote with $V = \ker(D\pi)$, which is a smooth sub-bundle of TE . An Ehresmann connection on E is a smooth sub-bundle H of TE so that:

1. H is a complement of V , that is $H \cap V = \{0\} \times E$ is the zero section and $H + V = TE$.
2. For any $x \in E$ and $p \in P$ one has $D_x R_p(H_x) = H_{xp}$, where R_p is the right-multiplication map $R_p : E \rightarrow E, x \mapsto x \cdot p$.

The bundle V is called the vertical bundle and the bundle H the horizontal bundle of the connection. An Ehresmann connection in this sense is then simply a right-invariant complementary bundle to the vertical bundle $V = \ker(D\pi)$. There is another equivalent definition of an Ehresmann connection we now remark upon.

For any principal bundle of a Lie group P with Lie algebra $\mathfrak{p} = \text{Lie}(P)$ the fundamental fields of the bundle are always defined. To be precise, for any $\xi \in \mathfrak{p}$ the fundamental field $X(\xi)$ is the vector field on E defined by $X(\xi)_x = \frac{d}{dt}x \cdot \exp(t\xi)|_{t=0}$ for $x \in E$. These fields are, by definition, always vertical. In fact at any point they span the vertical bundle.

We can always turn this around to get a \mathfrak{p} -valued 1-form $\alpha' : V \rightarrow \mathfrak{p}$. If we have a complementary bundle H so that $H + V = TE$ then we can extend α' by zero on H . On the other hand, any \mathfrak{p} -valued 1-form $\alpha : TE \rightarrow \mathfrak{p}$ agreeing with α' when restricted to V induces such a complement, namely $\ker(\alpha)$.

Definition 3.2.4 (Ehresmann connection (connection form)). Let P be a Lie group, M a manifold and $E \xrightarrow{\pi} M$ be a P -principal bundle over M . The connection form of an Ehresmann

connection H is defined to be the induced map \mathfrak{p} -valued 1-form $\alpha : TE \rightarrow \mathfrak{p}$, where for $x \in E$ and $h_x \in H_x, v_x \in V_x$ we define $\alpha_x(h_x + v_x) = \alpha'_x(v_x)$.

The connection form α uniquely determines the bundle H . In fact we may recast conditions 1. and 2. from Definition 3.2.3 into conditions on the connection form.

Proposition 3.2.5. *Let P be a Lie group, M a manifold and $E \xrightarrow{\pi} M$ be a P -principal bundle over M . A \mathfrak{p} -valued 1-form $\alpha : TE \rightarrow \mathfrak{p}$ is the connection form of an Ehresmann connection if and only if:*

1. *If $X(\xi)$ is a fundamental field then $\alpha_x(X(\xi)_x) = \xi$.*
2. *If $p \in P$ then $\alpha_{xp} \circ D_x R_p = \text{Ad}(p^{-1}) \cdot \alpha_x$.*

Proof. From the previous discussion it is obvious that point 1. of the proposition is equivalent to $H := \ker(\alpha)$ being a complementary bundle of V , i.e. equivalent to point 1. of Definition 3.2.3.

On the other hand one has $x \exp(t\xi) \cdot p = xp \cdot p^{-1} \exp(t\xi)p$, whence $D_x R_p(X_x(\xi)) = X_{xp}(\text{Ad}(p^{-1})\xi)$, implying $\alpha_{xp}(D_x R_p(v)) = \text{Ad}(p^{-1})\alpha_x(v)$ for all $v \in V_x$. Thus $\alpha \circ D_x R_p = \text{Ad}(p^{-1}) \cdot \alpha_x$ is only a condition when restricted to H_x , here it becomes:

$$\alpha_{xp}(D_x R_p(h)) = \text{Ad}(p^{-1})\alpha_x(h) \stackrel{!}{=} 0$$

for all $h \in H_x$. This condition is nothing other than $h \in H_x \implies D_x R_p(h) \in \ker(\alpha_{xp}) = H_{xp}$, in other words $H_{xp} \supseteq D_x R_p(H_x)$ for all $x \in E$. By comparing dimensions one notes the statement is actually $H_{xp} = D_x R_p(H_x)$. \square

We now specialise to pseudo-Riemannian geometry. We let $\widehat{M} \xrightarrow{\pi} M$ be the orthonormal frame-bundle of a manifold M equipped with a metric g of signature (p, q) .

Fixing a metric connection ∇ we will define the horizontal lift $\widehat{\gamma} : (-\epsilon, \epsilon) \rightarrow \widehat{M}$ of a path $\gamma : (-\epsilon, \epsilon) \rightarrow M$. All directions $\frac{d}{dt}\widehat{\gamma}$ of such horizontal lifts together will give a smooth sub-bundle of $T\widehat{M}$ satisfying conditions 1. and 2. of Definition 3.2.3, thus giving us an Ehresmann connection on \widehat{M} .

Definition 3.2.6. Let $\gamma : (-\epsilon, \epsilon) \rightarrow M$ be a path and $\widehat{x} \in \pi^{-1}(\gamma(0))$ an ONB of $T_{\gamma(0)}M$. We define the lift $\widehat{\gamma}(t)$ of γ at \widehat{x} as the parallel transport of \widehat{x} along γ , that is:

$$\widehat{\gamma}(t) = P_0^t(s \mapsto \gamma(s))[\widehat{x}] = (P_0^t(s \mapsto \gamma(s))[(\widehat{x})_1], \dots, P_0^t(s \mapsto \gamma(s))[(\widehat{x})_n]).$$

Here parallel transport is the parallel transport induced by the covariant derivative ∇ .

The relevant statements are that $\widehat{\gamma}$ is a well-defined (meaning here $P_0^t(s \mapsto \gamma(s)) [\widehat{x}]$ is an ONB for all t) smooth path $(-\epsilon, \epsilon) \rightarrow \widehat{M}$ and that the collection of possible vectors $\frac{d}{dt}\widehat{\gamma}(t)$ span a horizontal bundle H that is an Ehresmann connection on \widehat{M} (in the sense of 3.2.3).

Lemma 3.2.7. *Let $\gamma : (-\epsilon, \epsilon) \rightarrow M$ be a smooth path, then for any $\widehat{x} \in \pi^{-1}(\gamma(0))$ the lift $\widehat{\gamma}$ at \widehat{x} is a smooth path $(-\epsilon, \epsilon) \rightarrow \widehat{M}$. Further defining*

$$H_{\widehat{x}} := \left\{ \frac{d}{dt}\widehat{\gamma}(t)|_{t=0} \mid \widehat{\gamma} \text{ a lift at } \widehat{x} \text{ of a path through } \pi(\widehat{x}) \right\}$$

gives an Ehresmann connection H on \widehat{M} .

Proof. That the map is well-defined is clear, as parallel transport is a linear isometry and thus the parallel transport of an ONB is an ONB. We take care of smoothness and the bundle conditions by looking at special charts on which the bundle \widehat{M} is trivial.

Specifically let $x \in M$ and U a small coordinate neighbourhood of x diffeomorphic to the ball $B_1(0) \subset \mathbb{R}^n$ with $x \equiv 0$, we now identify U with $B_1(0)$ for local considerations. We choose an ONB (b_1, \dots, b_n) on $T_0B_1(0)$ and for $\vec{x} \in B_1(0)$ we parallel transport this basis along the path $t \mapsto t\vec{x}$ to get an ONB $(b_{1,\vec{x}}, \dots, b_{n,\vec{x}})$. Specifically $b_{i,\vec{x}}$ is determined by the solution of $\frac{d}{dt}v(t)^\mu + \Gamma_{\alpha\beta}^\mu(t\vec{x})\vec{x}^\alpha v(t)^\beta = 0$ (for convenience we used the Einstein sum convention, greek superscripts denote components) with $v(0) = b_i$ and $\Gamma_{\alpha\beta}^\mu$ the Christoffel symbols of ∇ on the coordinate chart. Since the coefficients of this ODE vary smoothly in \vec{x} the solution also varies smoothly in \vec{x} and $b_{i,\vec{x}}$ is a smooth vector field, giving us a smooth ONB on $B_1(0)$. This gives us a local trivialisation $B_1(0) \times O(p, q)$ of $\pi^{-1}(B_1(0)) \subseteq \widehat{M}$.

We now check that the lifts $\widehat{\gamma}$ are smooth in this trivialisation. Let $\gamma : (-\epsilon, \epsilon) \rightarrow B_1(0)$ be a smooth path through 0, then the parallel transport of b_i along γ results in a smooth vector field along γ . In particular expressing this vector field in basis $b_{i,\gamma(t)}$ results in the coefficients varying smoothly in t . In an equation:

$$P_0^t(s \mapsto \gamma(s)) [b_i] = \sum_j c_{ij}(t) b_{j,\gamma(t)}$$

where $c_{ij}(t)$ is smooth, note that $c_{ij}(t)$ are the components the base-change matrix from $(b_{1,\gamma(t)}, \dots, b_{n,\gamma(t)})$ to $P_0^t(s \mapsto \gamma(s)) [(b_1, \dots, b_n)]$. This means that the lift $\widehat{\gamma}$ of γ at (b_1, \dots, b_n) is given (in the chart $B_1(0) \times O(p, q)$) by

$$t \mapsto (\gamma(t), (c_{ij}(t))_{1 \leq i, j \leq n})$$

which is smooth.

Next we determine $H_{(0,(b_1,\dots,b_n))}$. First note that the ODE $\nabla_{\dot{\gamma}}v = 0$ is in coordinates

$$\frac{d}{dt}v(t)^\mu + \Gamma_{\alpha\beta}^\mu(\gamma(t))\dot{\gamma}^\alpha(t)v^\beta(t) = 0, \quad (3.2)$$

which implies that $\frac{d}{dt}v(0)$ is uniquely determined by $\dot{\gamma}(0)$ and $v(0)$. This is important here as it means that $\frac{d}{dt}\widehat{\gamma}(0)$ is uniquely determined by $\dot{\gamma}(0)$ and $\widehat{\gamma}(0)$. Then (in our special coordinate neighbourhood):

$$H_{(0,(b_1,\dots,b_n))} = \left\{ \frac{d}{dt}(\widehat{t \mapsto t\vec{x}})|_{t=0} \mid \vec{x} \in B_1(0) \right\} = T_{(0,1)}B_1(0) \subset T_{(0,1)}(B_1(0) \times O(p,q)),$$

where we remind ourselves that $t \mapsto t\vec{x}$ were the paths we used to define $(b_{1,\vec{x}}, \dots, b_{n,\vec{x}})$. So we have seen that $H_{(0,(b_1,\dots,b_n))}$ is a complement of $V_{(0,(b_1,\dots,b_n))} = T_{(0,(b_1,\dots,b_n))}O(p,q)$ verifying point 1. of Definition 3.2.3.

Next we note that parallel transport is linear, so if $A \in O(p,q)$ and we lift γ at $(\sum_i A_{i1}b_i, \dots, \sum_i A_{in}b_i) = (b_1, \dots, b_n) \cdot A$ then the resulting path will be $\widehat{\gamma}(t) \cdot A$. This means that the derivative of this alternative lift differs from the derivative of $\widehat{\gamma}$ by right-multiplication with A . This implies then

$$H_{(0,(b_1,\dots,b_n)) \cdot A} = D_{(0,(b_1,\dots,b_n))}R_A(H_{(0,(b_1,\dots,b_n))}),$$

verifying condition 2. of Definition 3.2.3.

The only thing left to check is that the distribution H is smooth. To see this we use the criterium that a distribution is smooth if it is (locally) spanned by smooth vector fields. Let $\vec{x} \in B_1(0)$, e^1, \dots, e^n be a basis of \mathbb{R}^n and $\gamma_i(t, \vec{x}) = \vec{x} + te^i$. We then define for $\vec{x} \in B_1(0)$, $A \in O(p,q)$

$$X_i(\vec{x}, A) = D_{(\vec{x},(b_{1,\vec{x}}, \dots, b_{n,\vec{x}}))}R_A \left[\frac{d}{dt}\widehat{\gamma_i(t, \vec{x})}|_{t=0} \right],$$

here the lift of $\gamma_i(t, \vec{x})$ is taken at $(b_{1,\vec{x}}, \dots, b_{n,\vec{x}})$, so $X_i(\vec{x}, A)$ is a vector in $T_{(\vec{x},(b_{1,\vec{x}}, \dots, b_{n,\vec{x}}))}B_1(0) \times O(p,q)$. By definition we have $H_{(\vec{x},(b_{1,\vec{x}}, \dots, b_{n,\vec{x}})) \cdot A} = \text{span}\{X_i \mid i = 1, \dots, n\}$, we will now check that $X_i(\vec{x}, A)$ varies smoothly over $B_1(0) \times O(p,q)$.

Since $\widehat{\gamma_i(t, \vec{x})}$ is a tuple of solutions to Equation 3.2 with boundary conditions $v(0) = b_{j,\vec{x}}$, $j = 1, \dots, n$ and $\gamma(t)$ replaced with $\gamma_i(t, \vec{x})$ and everything in the equation varies smoothly with the parameter \vec{x} , we find that $\frac{d}{dt}\widehat{\gamma_i(t, \vec{x})}|_{t=0}$ varies smoothly with \vec{x} (and is independent of A). Further the derivative of right multiplication is smooth (in all arguments), so $X_i(\vec{x}, A)$ is the application of a smooth function to a smooth function, hence itself smooth. \square

So a metric g and a metric connection ∇ induce an Ehresmann connection H on $\widehat{M} \xrightarrow{\pi} M$. We will be working with the connection form of this Ehresmann connection in the next section, Section 3.2.3.

An Ehresmann connection on \widehat{M} is equivalent to a metric connection. We sketch how the injectivity and surjectivity $\nabla \mapsto H$ can be seen. For a smooth path $\gamma : (-\epsilon, \epsilon) \rightarrow M$ and a basis (b_1, \dots, b_n) of $T_{\gamma(0)}M$ there is a unique lift $\widehat{\gamma} : (-\epsilon, \epsilon) \rightarrow \widehat{M}$ of γ with $\widehat{\gamma}(0) = (b_1, \dots, b_n)$ and for which $\widehat{\gamma}(t)$ is horizontal for all t (see [KN 63] Proposition 3.1 of Chapter II). This is then nothing other than a smoothly varying ONB $(b_1(t), \dots, b_n(t))$ along $\gamma(t)$, defining a parallel transport by $P_0^t(s \mapsto \gamma(s)) [\sum_i v_i b_i] = \sum_i v_i b_i(t)$. As demonstrated by Equation (3.1), a covariant derivative is uniquely determined by its associated notion of parallel transport, thus if this Ehresmann connection comes from a metric connection, then this connection is unique and the correspondence is injective. For construction of a covariant derivative on $\widehat{M} \times_{O(p,q)} \mathbb{R}^n = TM$ from this notion of parallel transport is we refer to [KN 63] Chapter III Paragraphs §1 and §2. By the fact the parallel transport is by isometries of the tangent spaces this covariant derivative is actually a metric connection.

3.2.3 Cartan connection

Let (M, g) be a pseudo-Riemannian manifold of signature (p, q) and $\widehat{M} \xrightarrow{\pi} M$ its frame-bundle and H the Ehresmann connection associated to a metric connection ∇ , denote with $\alpha : TM \rightarrow \mathfrak{so}(p, q)$ the connection form of H . We define:

Definition 3.2.8. For $\widehat{x} = (x, b_1, \dots, b_n) \in \widehat{M}$ define $\theta_{\widehat{x}} : T_{\widehat{x}}\widehat{M} \rightarrow \mathbb{R}^n$, $v \mapsto \mathcal{E}_{\widehat{x}}(D_{\widehat{x}}\pi[v])$, where

$$\mathcal{E}_{\widehat{x}} \left[\sum_i v_i b_i \right] = \sum_i v_i e^i,$$

where e^i denotes the canonical basis of \mathbb{R}^n .

The map $\theta_{\widehat{x}}$ thus projects a vector down to $T_{\pi(\widehat{x})}M$ and then expresses it in coordinates \widehat{x} . This map is often called the soldering form, although we will not use that name here, reserving it for a slightly different context.

Proposition 3.2.9. *Let $\widehat{x} \in \widehat{M}$, $p \in O(p, q)$, then $\theta_{\widehat{x} \cdot p} \circ D_{\widehat{x}}R_p = p^{-1} \cdot \theta_{\widehat{x}}$ (where $O(p, q)$ acts on \mathbb{R}^n in the usual way) and $\theta_{\widehat{x}}$ varies smoothly in \widehat{x} , that is $\theta : T\widehat{M} \rightarrow \mathbb{R}^n$ is a smooth 1-form.*

Proof. We begin with the equivariance relation $\theta_{\widehat{x} \cdot p} \circ D_{\widehat{x}}R_g = p^{-1} \cdot \theta_{\widehat{x}}$. First note that $D_{\widehat{x} \cdot p}\pi \circ D_{\widehat{x}}R_g = D_{\widehat{x}}\pi$, so the only thing to check is that $\mathcal{E}_{\widehat{x} \cdot p} = p^{-1} \cdot \mathcal{E}_{\widehat{x}}$. To this end:

$$b_i = \sum_{jk} (p^{-1})_{ji} p_{kj} b_k,$$

implying

$$\mathcal{E}_{\widehat{x} \cdot p} \left[\sum_i v_i b_i \right] = \mathcal{E}_{\widehat{x} \cdot p} \left[\sum_j \left(\sum_i (p^{-1})_{ji} v_i \right) \sum_k p_{kj} b_k \right] = \sum_{ij} (p^{-1})_{ji} v_i e^j = p^{-1} \cdot \mathcal{E}_{\widehat{x}} \left[\sum_i v_i b_i \right],$$

which implies the desired equivariance relation.

To see that the map varies smoothly we may, around any point $x \in M$, choose a neighbourhood U and an orthonormal frame (b_1, \dots, b_n) trivialising $\pi^{-1}(U) \cong U \times O(p, q)$. On this frame the map $\mathcal{E}_{\hat{x}} : T_{\pi(\hat{x})}U \rightarrow \mathbb{R}^n$ varies smoothly in both the horizontal and the vertical directions. Since the map $D_{\hat{x}}\pi$ is also smooth one gets that $\theta_{\hat{x}}$ varies smoothly in \hat{x} . \square

We now combine the forms α and θ together to get a $\mathfrak{so}(p, q) \times \mathbb{R}^n$ valued 1-form on $T\widehat{M}$:

Definition 3.2.10 (Cartan connection). We define $\omega : T\widehat{M} \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n$, $\omega = \alpha + \theta$.

Depending on taste this definition might be abusing notation a bit, since α and θ are not valued in the same vector space and thus cannot be summed. To be more concrete let $\iota_1 : \mathfrak{so}(p, q) \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n$ and $\iota_2 : \mathbb{R}^n \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n$ be the inclusions $\iota_1(p) = (p, 0)$, $\iota_2(v) = (1, v)$. Then $\omega := \iota_1 \circ \alpha + \iota_2 \circ \theta$.

Theorem 3.2.11. ($\widehat{M} \xrightarrow{\pi} M, \omega : \widehat{M} \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n$) is a Cartan connection on M modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$, that is:

1. $\omega_{\hat{x}}$ is an isomorphism of vector spaces.
2. For all $\hat{x} \in \widehat{M}$ and $p \in O(p, q)$ we have $\omega_{\hat{x}p} \circ D_{\hat{x}}R_p = \text{Ad}(p^{-1}) \cdot \omega_{\hat{x}}$.
3. For any $\xi \in \mathfrak{so}(p, q)$ and $\hat{x} \in \widehat{M}$ one has $\omega(\frac{d}{dt}\hat{x} \exp(t\xi)|_{t=0}) = \xi$, i.e. ω sends the fundamental fields to their generators.

Proof. For point 1. we show that θ restricted to $H_{\hat{x}}$ is an injective map (into \mathbb{R}^n) for any $\hat{x} \in \widehat{M}$. Since α is $\mathfrak{so}(p, q)$ valued and injective when restricted to V , and $T_{\hat{x}}\widehat{M} = H + V$ as well as $\theta(V) = 0 = \alpha(H)$, we find that $\alpha + \theta : T_{\hat{x}}\widehat{M} \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n$ is an injective linear map. By comparing dimensions it must be an isomorphism. The injectivity of $\theta_{\hat{x}}$ follows from $\ker(D_{\hat{x}}\pi) = V_{\hat{x}}$ and $\mathcal{E}_{\hat{x}} : T_x M \rightarrow \mathbb{R}^n$ being an isomorphism, whence $\theta_{\hat{x}} = \mathcal{E}_{\hat{x}} \circ D_{\hat{x}}\pi$ is injective on $H_{\hat{x}}$.

For point 2. we note that

$$(\omega_{\hat{x}p} \circ D_{\hat{x}}R_p)(v) = \iota_1((\alpha_{\hat{x}p} \circ D_{\hat{x}}R_p)(v)) + \iota_2((\theta_{\hat{x}p} \circ D_{\hat{x}}R_p)(v)) = \iota_1(\text{Ad}(p^{-1})\omega_{\hat{x}}(v)) + \iota_2(p^{-1}\theta_{\hat{x}}(v)).$$

Verifying $\iota_1(\text{Ad}(p^{-1})\xi) = \text{Ad}(p^{-1})\iota_1(\xi)$ and $\iota_2(p^{-1}v) = \text{Ad}(p^{-1})\iota_2(v)$ for $\xi \in \mathfrak{so}(p, q)$ and $v \in \mathbb{R}^n$ then gives point 2.

Point 3 follows by noting that the fundamental fields are vertical, so they are in $\ker(D\pi)$ and

$$\omega(\frac{d}{dt}\hat{x} \exp(t\xi)|_{t=0}) = \iota_1(\alpha(\frac{d}{dt}\hat{x} \exp(t\xi)|_{t=0})) = \iota_1(\xi) = \xi$$

follows for any $\xi \in \mathfrak{so}(p, q)$. \square

3.2.4 Cartan-isomorphisms and isometries

In this section we prove that the isometries of pseudo-Riemannian metrics sending a metric connection ∇ on U to a metric connection ∇' on V are the same as isomorphisms of the induced Cartan connections. This important step expresses that a pseudo-Riemannian metric g together with a metric connection ∇ has the same geometric content as the induced Cartan connection. Additionally we check the functoriality of this correspondence.

Remark. We say that a diffeomorphism $f : U \rightarrow V$ sends ∇ to ∇' if $Df(\nabla_X Y) = \nabla'_{Df(X)}(Df(Y))$ for any vector fields X, Y (note that since f is a diffeomorphism Df actually does map vector fields on U to vector fields on V). This implies for example if $\nabla_{\dot{\gamma}(t)} Y = 0$ that then $\nabla'_{\frac{d}{dt} f(\gamma(t))} Df(Y) = 0$, meaning that Df preserves the notion of parallel transport. As expressed by Equation (3.1) the property that Df sends fields that are parallel along γ to fields that are parallel along $f \circ \gamma$ is then equivalent to f sending ∇ to ∇' .

Proposition 3.2.12 (Isometries lift to isomorphisms). *Let U, V be two n -dimensional manifolds with pseudo-Riemannian metrics g, g' of signature (p, q) and metric connections ∇, ∇' . Then any isometry $f : U \rightarrow V$ sending ∇ to ∇' lifts to an isomorphism $F : \widehat{U} \rightarrow \widehat{V}$ of the induced Cartan connections ω_U, ω_V on \widehat{U}, \widehat{V} .*

Proof. We first remark that the condition that a map is an isomorphism is a local condition, i.e. we assume \widehat{U} is trivialised by a global frame (b_1, \dots, b_n) , i.e. $\widehat{U} \cong U \times O(p, q)$. As $f : U \rightarrow V$ is an isometry $(Df(b_1), \dots, Df(b_n))$ is a global frame of V , trivialising that bundle as well. We then define $F : U \times O(p, q) \rightarrow V \times O(p, q)$, $(x, A) \mapsto (f(x), A)$ for $x \in U, A \in O(p, q)$. By definition this is a bundle-isomorphism and smooth. It remains to check that $F^*(\omega_V) = \omega_U$.

We first check that $F^*(\alpha_V) = \alpha_U$. Since F is a bundle-isomorphism it maps the fundamental fields to the fundamental fields, in particular $F^*(\alpha_V)$ restricted to the vertical bundle is the same as α_U restricted to the vertical bundle. What is left is to check that DF maps the horizontal bundle H_U to the horizontal bundle H_V . Let $\gamma : (-\epsilon, \epsilon) \rightarrow U$ be a smooth path, define $\xi_{ij}(t)$ via

$$P_0^t(s \mapsto \gamma(s))(b_{i, \gamma(0)}) = \sum_j \xi_{ij}(t) b_{j, \gamma(t)}, \quad (3.3)$$

i.e. ξ_{ij} are the coefficients of the parallel transport of b_i along γ by ∇ . Since f sends ∇ to ∇' it must preserve parallel transport and we have:

$$P_0^t(s \mapsto f(\gamma))(Df(b_i)_{f(\gamma(0))}) = \sum_j \xi_{ij}(t) Df(b_j)_{f(\gamma(t))}. \quad (3.4)$$

So let $(\gamma(0), A)$ be a lift of γ at 0, we remind ourselves that this corresponds to the basis $(\sum_i A_{i1} b_{i, \gamma(0)}, \dots, \sum_i A_{in} b_{i, \gamma(0)})$. By Equation (3.3) the parallel transport of this basis along γ

by ∇ is $(\sum_i A_{i1}\xi_{ij}(t)b_{j,\gamma(t)}, \dots, \sum_i A_{in}\xi_{ij}(t)b_{j,\gamma(t)})$. This just means that the lift $\widehat{\gamma}$ of γ is the map $t \mapsto (\gamma(t), \xi(t)^T \cdot A)$. Equation (3.4) now implies in the same way that the lift of $f \circ \gamma$ at $(f(\gamma(0)), A)$ is $(f(\gamma(t)), \xi(t)^T \cdot A) = F(\gamma(t), \xi(t)^T \cdot A)$. This implies that the lift of a horizontal path is again horizontal, meaning that DF preserves the horizontal bundle.

Finally we check $F^*(\theta_V) = \theta_U$. Note that since F is a bundle-isomorphism we have $D\pi_V \circ DF = D\pi_U$. Whence $F^*(\theta_V) = \theta_U$ turns into whether for $x \in T_x M$ the expansion of $D_x f(v)$ in basis $(D_x f(b_{1,x}), \dots, D_x f(b_{n,x}))$ is the same as the expansion of v in basis $(b_{1,x}, \dots, b_{n,x})$. But that is trivially true. \square

Proposition 3.2.13 (Isomorphisms are over isometries). *Let U, V be two n -dimensional manifolds with pseudo-Riemannian metrics g, g' of signature (p, q) and metric connections ∇, ∇' . Then any isomorphism $F : \widehat{U} \rightarrow \widehat{V}$ of the induced Cartan connections ω_U, ω_V is a bundle-map over an isometry $f : U \rightarrow V$ that maps ∇ to ∇' .*

Proof. As the conditions that a map is an isometry and maps ∇ to ∇' are purely local conditions, we make U, V both small enough to admit global orthonormal frames (b_1, \dots, b_n) (on U) and (e_1, \dots, e_n) (on V). We then trivialise $\widehat{U} = U \times O(p, q)$ and $\widehat{V} = V \times O(p, q)$. Since F is a bundle-isomorphism it must necessarily be of the form:

$$F(x, A) = (f(x), g(x) \cdot A)$$

where $g(x) \in O(p, q)$ for all $x \in U$ and $f : U \rightarrow V$ is the map over which F is a bundle-map. Now let $\pi_{\mathbb{R}^n} : \mathfrak{so}(p, q) \times \mathbb{R}^n \rightarrow \mathbb{R}^n$. We have $\theta_U = \pi_{\mathbb{R}^n} \circ \omega_U$ and $\theta_V = \pi_{\mathbb{R}^n} \circ \omega_V$, whence $F^*(\theta_V) = F^*(\pi_{\mathbb{R}^n} \circ \omega_V) = \pi_{\mathbb{R}^n} \circ \omega_U = \theta_U$.

If we evaluate θ_U on a vector $(v, 0) \in T_{(x,A)}(U \times O(p, q))$ we by definition get the expansion of v in the orthonormal basis $(\sum_i A_{i1}b_{i,x}, \dots, \sum_i A_{in}b_{i,x})$. Evaluating this vector with $F^*(\theta_V)$ we get the expansion of $D_x f(v)$ in the orthonormal basis $(\sum_i (g(x)A)_{i1}e_{i,x}, \dots, \sum_i (g(x)A)_{in}e_{i,x})$. These two expansions must be equal, implying that $D_x f : T_x U \rightarrow T_{f(x)} V$ is an isometry. This holds for every point and then f itself must be an isometry.

The previous paragraph actually implies that $D_x f(\sum_i A_{i1}b_{i,x}) = \sum_i (g(x)A)_{i1}e_{i,x}$, meaning that $F((\sum_i A_{i1}b_{i,x}, \dots, \sum_{in} A_{in}b_{i,x})) = (D_x f(\sum_i A_{i1}b_{i,x}), \dots, D_x f(\sum_i A_{in}b_{i,x}))$, that is the image of F on a basis is the transformation of the basis by Df , or $F(\widehat{x}) = D_{\pi(\widehat{x})} f(\widehat{x})$. This is the key to checking that f maps ∇ to ∇' .

Let $\pi_{\mathfrak{so}(p,q)} : \mathfrak{so}(p, q) \times \mathbb{R}^n \rightarrow \mathfrak{so}(p, q)$, then $\pi_{\mathfrak{so}(p,q)} \circ \omega_U$ is the Ehresmann connection on U and $\pi_{\mathfrak{so}(p,q)} \circ \omega_V$ the Ehresmann connection on V . In particular $F^*(\pi_{\mathfrak{so}(p,q)} \circ \omega_U) = \pi_{\mathfrak{so}(p,q)} \circ \omega_V$ and F must map horizontal paths to horizontal paths. So if $\widehat{\gamma}(t)$ is a horizontal path in \widehat{U} , then

$F(\widehat{\gamma}(t)) = Df(\widehat{\gamma}(t))$ is horizontal in \widehat{V} . But a horizontal path $\widehat{\gamma}(t)$ is the parallel transport of the basis $\widehat{\gamma}(0)$ along $\pi_U(\widehat{\gamma}(t))$, and thus Df maps a parallel basis to a parallel basis, implying that f preserves parallel transport and thus sends ∇ to ∇' . \square

Proposition 3.2.14 (Functoriality). *Let U, V, W be pseudo-Riemannian manifolds of signature (p, q) , each endowed with a metric connection. Suppose $f : U \rightarrow V$ and $g : V \rightarrow W$ are isometries mapping the metric connections to each other and denote with $F : \widehat{U} \rightarrow \widehat{V}$ and $G : \widehat{V} \rightarrow \widehat{W}$ the induced isomorphisms on the associated Cartan connections. Then the isomorphism induced by $g \circ f$ is $G \circ F$.*

Proof. Note that $G \circ F$ is a an isomorphism of Cartan connections over the base map $g \circ f$. By Lemma 2.4.3 this is the only isomorphism over $g \circ f$ (note that $O(p, q) \times \mathbb{R}^n$ acts effectively on \mathbb{R}^n and that \mathbb{R}^n is connected, so this Klein geometry is effective and Lemma 2.4.3 applies). As such it must be equal to the lift of $g \circ f$. \square

3.2.5 Recovering a pseudo-Riemannian geometry from an $O(p, q) \times \mathbb{R}^n / O(p, q)$ Cartan geometry

Here we turn the previous process around. Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n)$ be a Cartan connection on an n -dimensional manifold M modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$, we wish to construct a pseudo-Riemannian metric g of signature (p, q) and a metric connection ∇ . For convenience we will denote $\mathfrak{so}(p, q) \times \mathbb{R}^n$ with \mathfrak{g} and $\mathfrak{so}(p, q)$ with \mathfrak{p} .

The rough idea for the construction of g is that for any Cartan connection $\omega : T\widehat{M} \rightarrow \mathfrak{g}$ on a manifold M modelled on G/P one considers the quotient of $\widehat{M} \times \mathfrak{g}/\mathfrak{p}$ by the P -action $(\widehat{x}, v) \mapsto (\widehat{x} \cdot p, \text{Ad}(p^{-1})v)$. This will result in a vector bundle $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$ over M , called the isotropy bundle. In fact this bundle is canonically isomorphic to the tangent bundle TM . One can then pull back any P -invariant structure defined on $\mathfrak{g}/\mathfrak{p}$ to TM , in our case $\mathfrak{g}/\mathfrak{p} = \mathbb{R}^n$ which admits a canonical $O(p, q)$ -invariant symmetric bilinear form of signature (p, q) . Pulling this form back to TM gives a Riemannian metric on M .

To get the metric connection one first constructs a bundle isomorphism from \widehat{M} to the frame bundle OM of the just defined Riemannian metric. One then makes use of the fact that \mathfrak{p} admits an $\text{Ad}(P)$ invariant complement \mathfrak{h} in \mathfrak{g} , that is \mathfrak{g} is a reducible P -module and $\mathfrak{p} \oplus \mathfrak{h}$ is a decomposition of \mathfrak{g} into two P -modules. Pulling \mathfrak{h} back to $\omega^{-1}(\mathfrak{h})$ gives a sub-bundle H of $T\widehat{M}$ complementing \mathfrak{p} , this will be the horizontal bundle of an Ehresmann connection.

Recovering a metric

Most of the steps in this process work for arbitrary Cartan connections. Keeping the context to this generality has the advantage of making everything more readable (things would be a lot uglier if every \mathfrak{g} were replaced by $\mathfrak{so}(p, q) \times \mathbb{R}^n$). For the following let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan geometry modelled on G/P .

Proposition 3.2.15. *The P -action*

$$(\widehat{M} \times \mathfrak{g}/\mathfrak{p}) \times P \rightarrow \widehat{M} \times \mathfrak{g}/\mathfrak{p}, \quad ((\widehat{x}, v), p) \mapsto (\widehat{x}p, \text{Ad}(p^{-1})v)$$

is smooth, free and proper.

Proof. Smoothness and freeness of the action are obvious. Properness is however also obvious since the action $\widehat{M} \times P \rightarrow \widehat{M}, (\widehat{x}, p) \mapsto \widehat{x}p$ is already proper. \square

By the quotient manifold theorem $(\widehat{M} \times \mathfrak{g}/\mathfrak{p})/P$ is then a manifold of dimension $\dim(\widehat{M}) + \dim(\mathfrak{g}/\mathfrak{p}) - \dim(P)$. We will denote it with $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$.

Definition 3.2.16 (Isotropy bundle). We call $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$ the isotropy bundle and define the projection map $\tilde{\pi} : \widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \rightarrow M, [\widehat{x}, v] \mapsto \pi(\widehat{x})$.

We remark that $\tilde{\pi}$ is well defined and smooth, smoothness follows in this case from the commutativity of the following diagram:

$$\begin{array}{ccc} \widehat{M} \times \mathfrak{g}/\mathfrak{p} & \xrightarrow{\quad} & \widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \\ & \searrow \pi & \swarrow \tilde{\pi} \\ & M & \end{array}$$

Proposition 3.2.17. $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \xrightarrow{\tilde{\pi}} M$ is a smooth vector bundle over M with fibre $\mathfrak{g}/\mathfrak{p}$.

Proof. The linear structure on the fibres of $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$ is inherited from that of $\widehat{M} \times \mathfrak{g}/\mathfrak{p}$ (the P -action is by vector bundle isomorphisms). We now construct a local trivialisation around any point.

Let $x \in M$ and $U \subset M$ a neighbourhood of x admitting a section $s : U \rightarrow \widehat{M}$ of $\widehat{M} \xrightarrow{\pi} M$. Then $\psi_s : U \times \mathfrak{g}/\mathfrak{p} \rightarrow \tilde{\pi}^{-1}(U), (x, v) \mapsto [s(x), v]$ is linear in v , additionally it factors via smooth maps $U \times \mathfrak{g}/\mathfrak{p} \rightarrow \widehat{M} \times \mathfrak{g}/\mathfrak{p} \rightarrow \widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$ and is thus smooth. It is injective since $[s(x), v] = [s(x'), v']$ implies $s(x) = s(x') \cdot p$ and then (since s is a section) $x = x'$ and then finally $v = v'$. Surjectivity is also clear since any $[\widehat{x}, v]$ can be written as $[s(\pi(\widehat{x}))p, v]$ for an appropriate $p \in P$. But then

$$[\widehat{x}, v] = [s(\pi(\widehat{x}))p, v] = [s(\pi(\widehat{x})), \text{Ad}(p)v] = \psi_s(\pi(\widehat{x}), \text{Ad}(p)v). \quad \square$$

Consider the map $\widehat{M} \times \mathfrak{g} \xrightarrow{\omega^{-1}} T\widehat{M} \xrightarrow{D\pi} TM$. Note that since the kernel of $D\pi$ is the vertical bundle, that is $\omega^{-1}(\widehat{M} \times \mathfrak{p})$, this induces a well defined map $\widehat{M} \times \mathfrak{g}/\mathfrak{p} \rightarrow TM$. Further by the compatibility of ω with right-multiplication and invariance of π we have for any $(\widehat{x}, v) \in \widehat{M} \times \mathfrak{g}$:

$$(D_{\widehat{x}}\pi \circ \omega_{\widehat{x}}^{-1})(\widehat{x}, v) = (D_{\widehat{x}}\pi \circ D_{\widehat{x}p}R_p \circ \omega_{\widehat{x}p})(\widehat{x}p, \text{Ad}(p)v) = (D_{\widehat{x}p}\pi \circ \omega_{\widehat{x}p}^{-1})(\widehat{x}p, \text{Ad}(p)v),$$

implying that $D\pi \circ \omega^{-1}$ is P -invariant, meaning that it actually induces a map $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \rightarrow TM$. We call this map Θ and briefly recall its construction in a diagram:

$$\begin{array}{ccccc} \widehat{M} \times \mathfrak{g} & \longrightarrow & \widehat{M} \times \mathfrak{g}/\mathfrak{p} & \longrightarrow & \widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \\ \omega^{-1} \downarrow & & \parallel & & \searrow \\ T\widehat{M} & & \widehat{M} \times \mathfrak{g}/\ker(D\phi \circ \omega^{-1}) & & \\ D\pi \downarrow & \swarrow & & \swarrow & \\ TM & \longleftarrow & & \longleftarrow & \Theta \end{array}$$

Definition 3.2.18 (Soldering form). The map $\Theta : \widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \rightarrow TM$, $[\widehat{x}, v] \mapsto D_{\widehat{x}}\pi(\omega_{\widehat{x}}^{-1}(v))$ defined above is called the soldering form.

Proposition 3.2.19. *The soldering form is a vector bundle isomorphism $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \rightarrow TM$.*

Proof. It is obvious that it is a bundle-map over the identity and linear on the fibres. We check injectivity and surjectivity. For surjectivity let $v \in T_x M$, then choose a $\widehat{x} \in \pi^{-1}(x)$ and a $w \in \mathfrak{g}$ with $D_{\widehat{x}}\pi(\omega_{\widehat{x}}^{-1}w) = v$. Then by definition $\Theta([\widehat{x}, w]) = v$. On the other hand let $\widehat{x}, \widehat{x}' \in \widehat{M}$ and $w, w' \in \mathfrak{g}$ and suppose $\Theta([\widehat{x}, w]) = \Theta([\widehat{x}', w'])$. First we see that there must be a p with $\widehat{x} = \widehat{x}'p$, as otherwise Θ sends these to different tangent spaces. Noting that $D_{\widehat{x}}\pi = D_{\widehat{x}p}\pi \circ D_{\widehat{x}}R_p$ and using the compatibility of ω with right-multiplication we get:

$$D_{\widehat{x}}\pi(\omega_{\widehat{x}}^{-1}w) = D_{\widehat{x}p}\pi(D_{\widehat{x}}R_p(\omega_{\widehat{x}}^{-1}w)) = D_{\widehat{x}p}\pi(\omega_{\widehat{x}p}^{-1}\text{Ad}(p^{-1})w) \stackrel{!}{=} D_{\widehat{x}p}\pi(\omega_{\widehat{x}p}^{-1}w'),$$

implying that $\omega_{\widehat{x}p}^{-1}(\text{Ad}(p^{-1})w - w') \in \ker(D_{\widehat{x}p}\pi)$, or $\text{Ad}(p^{-1})w - w' \in \mathfrak{p}$. This then means $[\widehat{x}', w'] = [\widehat{x}', \text{Ad}(p^{-1})w] = [\widehat{x}p, \text{Ad}(p^{-1})w] = [\widehat{x}, w]$ and Θ is injective. \square

Any P -invariant structure on $\mathfrak{g}/\mathfrak{p}$ can be pulled back to a structure on $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$, the preceding proposition then tells us that this is nothing other than a structure on TM . In our specific case $\mathfrak{g}/\mathfrak{p} = \mathfrak{so}(p, q) \times \mathbb{R}^n / \mathfrak{so}(p, q) = \mathbb{R}^n$, which admits a canonical $O(p, q)$ -invariant symmetric bilinear form of signature (p, q) . Pulling this back gives a symmetric bilinear form of signature (p, q) on each tangent space, that is a pseudo-Riemannian metric (the metric is smooth since the bilinear form on $\widehat{M} \times \mathfrak{g}/\mathfrak{p}$ is smooth, so the induced form on $\widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$ is also smooth).

Definition 3.2.20 (Metric induced by a Cartan connection). Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan connection modelled on $O(p, q) \ltimes \mathbb{R}^n / O(p, q)$. The pseudo-Riemannian metric g induced by ω is defined by $g_x(v, w) = \langle \Theta^{-1}(v), \Theta^{-1}(w) \rangle$ for $v, w \in T_x M$, where $\langle \cdot, \cdot \rangle$ is the standard scalar product on $\mathfrak{g}/\mathfrak{p} = \mathbb{R}^n$.

Remark. In [Sh 97] it is remarked that this procedure recovers a (pseudo-)Riemannian metric on M that is unique up to scalar multiple. This differs from our situation where the metric is unique. This difference comes from the fact that [Sh 97] is looking at $G = \text{Euc}(n)$, which is the symmetry group of the affine euclidean space. $\text{Euc}(n)$ is isomorphic to $O(n) \ltimes \mathbb{R}^n$, but not canonically so. Then $\text{Lie}(\text{Euc}(n))/\mathfrak{so}(n)$ is isomorphic to \mathbb{R}^n with the standard $O(n)$ action, but this isomorphism is no longer canonical. Thus one doesn't have a canonical invariant scalar product. What one does have is that the $O(n)$ invariant scalar products on this space form a one-dimensional ray, hence are unique up to a scalar multiple.

Recovering an Ehresmann connection

Let g denote the metric constructed above. Denote the orthonormal frame bundle of (M, g) with OM . Let $\widehat{x} \in \widehat{M}$ and denote with (e^1, \dots, e^n) the canonical basis of \mathbb{R}^n . Note that $(\Theta([\widehat{x}, e^1]), \dots, \Theta([\widehat{x}, e^n]))$ is an ONB of $T_{\pi(\widehat{x})}M$ and thus an element of OM .

Definition 3.2.21. We will denote the map $\widehat{M} \rightarrow OM$, $\widehat{x} \mapsto (\Theta([\widehat{x}, e^1]), \dots, \Theta([\widehat{x}, e^n]))$ with F .

Proposition 3.2.22. $F : \widehat{M} \rightarrow OM$ is an $O(p, q)$ -bundle isomorphism over the identity.

Proof. The most involved part is checking that this map is smooth. To that end we look at special charts of OM . Let $U \subset M$ be open admitting a smooth orthonormal frame (b_1, \dots, b_n) . For $\widehat{x} \in \widehat{M}$ denote with $A(\widehat{x})_{ij}$ the coefficients defined by

$$\Theta([\widehat{x}, e^i]) = \sum_j A(\widehat{x})_{ij} b_{j, \pi(\widehat{x})}.$$

These coefficients are smooth in \widehat{x} , since the composition $\widehat{M} \rightarrow \widehat{M} \times \mathfrak{g}/\mathfrak{p} \rightarrow \widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \rightarrow TM$, $\widehat{x} \mapsto (\widehat{x}, e^i) \mapsto [\widehat{x}, e^i] \mapsto \Theta([\widehat{x}, e^i])$ is smooth and $A(\widehat{x})_{ij}$ are the coefficients of expanding smoothly varying vector fields in the smoothly varying basis $b_{j, \pi(\widehat{x})}$. However in this trivialisation the map $F : \pi^{-1}(U) \rightarrow U \times O(p, q)$ is equal to $\widehat{x} \mapsto (\pi(\widehat{x}), A(\widehat{x}))$ and thus smooth.

Next we check that it is $O(p, q)$ -equivariant. If $A \in O(p, q)$ then

$$\begin{aligned} F(\widehat{x}A) &= (\Theta([\widehat{x}A, e^1]), \dots, \Theta([\widehat{x}A, e^n])) = (\Theta([\widehat{x}, A \cdot e^1]), \dots, \Theta([\widehat{x}, A \cdot e^n])) \\ &= \left(\Theta([\widehat{x}, \sum_i A_{i1} e^i]), \dots, \Theta([\widehat{x}, \sum_i A_{in} e^i]) \right) = \left(\sum_i A_{i1} \Theta([\widehat{x}, e^i]), \dots, \sum_i A_{in} \Theta([\widehat{x}, e^i]) \right) \\ &= (\Theta([\widehat{x}, e^1]), \dots, \Theta([\widehat{x}, e^n])) \cdot A = F(\widehat{x}) \cdot A. \end{aligned}$$

Thus F is a morphism of principal-bundles $\widehat{M} \rightarrow OM$. It obviously is a bundle-map over the identity and thus it is an automorphism. \square

Let $H = \omega^{-1}(\mathbb{R}^n) \subset \omega^{-1}(\mathfrak{so}(p, q) \times \mathbb{R}^n) = T\widehat{M}$. Note that H is a smooth sub-bundle of $T\widehat{M}$ since ω is a smooth 1-form and an isomorphism on every fibre. H complements $\ker(D\pi)$ as $\ker(D\pi) = \omega^{-1}(\mathfrak{so}(p, q))$ and \mathbb{R}^n complements $\mathfrak{so}(p, q)$ in \mathfrak{g} . We also remark that H is P -invariant, since

$$D_{\widehat{x}}R_g(H_{\widehat{x}}) = D_{\widehat{x}}R_g(\omega_{\widehat{x}}^{-1}(\mathbb{R}^n)) = \omega_{\widehat{x}g}^{-1}(\text{Ad}(g^{-1})\mathbb{R}^n) = \omega_{\widehat{x}g}^{-1}(\mathbb{R}^n) = H_{\widehat{x}g}$$

by compatibility of ω with the right-multiplication and the fact that \mathbb{R}^n is an $O(p, q)$ -invariant complement of $\mathfrak{so}(p, q)$. H thus fulfils the conditions in Definition 3.2.3 and is an Ehresmann connection on \widehat{M} .

Since we have a bundle-automorphism $F : \widehat{M} \rightarrow OM$, $DF(H)$ will be an Ehresmann connection on OM , thus defining a metric connection ∇ on M by the discussion in Section 3.2.2.

Definition 3.2.23 (Metric connection induced by a Cartan connection). Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan connection modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$, H the Ehresmann connection on \widehat{M} , g the pseudo-Riemannian metric induced by ω and $F : \widehat{M} \rightarrow OM$ the bundle-isomorphism to the frame bundle of g . Then the metric connection ∇ induced by ω is defined to be the metric connection induced by the Ehresmann connection $DF(H)$.

Remark. Let $(\widehat{M}_i \xrightarrow{\pi} M_i, \omega_i : T\widehat{M}_i \rightarrow \mathfrak{g})$, $i = 1, 2$ be two Cartan connections modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$ over M_1, M_2 and let $\Phi : \widehat{M}_1 \rightarrow \widehat{M}_2$ be an isomorphism over the base map ϕ . Denote with $F_i : \widehat{M}_i \rightarrow OM_i$ the isomorphisms into the orthonormal frame bundles of the induced metrics g_i on M_i . Then $F_2 \circ \Phi \circ F_1^{-1}$ is an isomorphism of Cartan connections $OM_1 \rightarrow OM_2$ over the base map ϕ . By Proposition 3.2.13 ϕ must be an isometry preserving the metric connections.

This means that the correspondence

$$(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g}) \mapsto (M, (g, \nabla)), \quad (\Phi : \widehat{M}_1 \rightarrow \widehat{M}_2 \text{ iso. over base } \phi) \mapsto (\phi : M_1 \rightarrow M_2)$$

(where (g, ∇) are the metric and metric connections induced by ω) is then well defined (and clearly functorial).

3.2.6 Equivalence

In Section 3.2.3 we constructed for a given pseudo-Riemannian metric g of signature (p, q) and a metric connection ∇ on a manifold M a Cartan connection $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$. In Section 3.2.5 we turned this around, constructing for such a Cartan connection a pseudo-Riemannian metric and a metric connection on the base manifold.

It remains to show that these constructions are inverse to each other. That is, show that $(g, \nabla) \rightarrow \omega \rightarrow (g', \nabla')$ recovers the original metric and connection, and that $\omega \rightarrow (g, \nabla) \rightarrow \omega'$ gives the original Cartan connection, up to an isomorphism over the identity. Theorem 3.2.24 takes care of the direction $(g, \nabla) \rightarrow \omega \rightarrow (g', \nabla')$, while Theorem 3.2.27 takes care of $\omega \rightarrow (g, \nabla) \rightarrow \omega'$.

We start by considering the path $(g, \nabla) \rightarrow \omega \rightarrow (g', \nabla')$.

Theorem 3.2.24. *Let M be a manifold g a pseudo-Riemannian metric of signature (p, q) and ∇ a metric connection. Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be the Cartan connection induced by (g, ∇) and (g', ∇') the metric and metric connection induced by ω , then $g = g'$ and $\nabla = \nabla'$.*

The first step is to understand the map $\Theta : \widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \rightarrow TM$:

Lemma 3.2.25. *Let $[\widehat{x}, v] \in \widehat{M} \times_P \mathfrak{g}/\mathfrak{p}$, where we identify $\mathfrak{g}/\mathfrak{p} = \mathbb{R}^n$. Then $\Theta([\widehat{x}, v]) = \mathcal{E}_{\widehat{x}}^{-1}(v)$.*

Proof. The map $\mathcal{E}_{\widehat{x}} : T_{\pi(\widehat{x})}M \rightarrow \mathbb{R}^n$ was defined in Definition 3.2.8, it describes the expansion of a vector in $T_{\pi(\widehat{x})}M$ in the basis \widehat{x} . We note that on the horizontal bundle H of \widehat{M} the restriction $D_{\widehat{x}}\pi : H_{\widehat{x}} \rightarrow T_{\pi(\widehat{x})}M$ is an isomorphism for each \widehat{x} . Then for $v \in \mathfrak{g}/\mathfrak{p} = \mathbb{R}^n$ we lift v to $\mathfrak{h} = \mathbb{R}^n$ (the complement of \mathfrak{p}), $\omega_{\widehat{x}}^{-1}(v)$ will then be equal to $\theta_{\widehat{x}}^{-1}(v)$, where we view the composition $\theta_{\widehat{x}} = \mathcal{E}_{\widehat{x}} \circ D_{\widehat{x}}\pi$ as a map $H_{\widehat{x}} \rightarrow \mathbb{R}^n$, where it is invertible. However per definition we have $\Theta([\widehat{x}, v]) = D_{\widehat{x}}\pi(\omega_{\widehat{x}}^{-1}v)$, which then just becomes $\Theta([\widehat{x}, v]) = \mathcal{E}_{\widehat{x}}^{-1}(v)$. \square

By construction the image $(\Theta([\widehat{x}, e^1]), \dots, \Theta([\widehat{x}, e^n]))$ is then an ONB of g' at $\pi(\widehat{x})$ (cf. the construction of g' following Definition 3.2.20). But since it is equal to $(\mathcal{E}_{\widehat{x}}^{-1}(e^1), \dots, \mathcal{E}_{\widehat{x}}^{-1}(e^n))$ by the previous lemma it is also an ONB of g at $\pi(\widehat{x})$. As a ONB uniquely determines the metric we find:

Corollary 3.2.26. *Let M be a manifold g a pseudo-Riemannian metric of signature (p, q) and ∇ a metric connection. If (g', ∇') is the metric and metric connection induced by the Cartan connection ω , which is in turn induced by a (g, ∇) , then $g = g'$.*

In particular the frame bundle of g' is equal to \widehat{M} , which was the frame bundle of g . The map F between the two frame bundles is then the identity. Now ∇' is defined to be the Ehresmann connection associated to $DF(H)$, where H is the Ehresmann connection associated to ∇ . Since F

is the identity one has $DF(H) = H$ and then $\nabla' = \nabla$. This step completes the proof of Theorem 3.2.24.

Theorem 3.2.27. *Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan connection modelled on $O(p, q) \times \mathbb{R}^n/O(p, q)$ and (g, ∇) the metric and metric connection induced by ω . If ω' denotes the Cartan connection induced by (g, ∇) , then ω and ω' are naturally isomorphic over the identity map $M \rightarrow M$.*

Proof. The relevant isomorphism is the map $F : \widehat{M} \rightarrow OM$ introduced in Definition 3.2.21, which is a bundle-isomorphism over the identity by Proposition 3.2.22. As the Ehresmann connection on OM is defined by $DF(H)$ where $H = \ker(\pi_{\mathfrak{so}(p,q)} \circ \omega)$, the connection forms of $DF(H)$ and H must be isomorphic by F , meaning $\pi_{\mathfrak{so}(p,q)} \circ \omega = F^*(\pi_{\mathfrak{so}(p,q)} \circ \omega')$. The only thing left to check is that $\pi_{\mathbb{R}^n} \circ \omega = F^*(\pi_{\mathbb{R}^n} \circ \omega')$.

This step is best summed up in the following diagram, we remind ourselves that $\mathfrak{h} = \mathbb{R}^n \subset \mathfrak{so}(p, q) \times \mathbb{R}^n$ is identified with $\mathfrak{g}/\mathfrak{p} = \mathfrak{so}(p, q) \times \mathbb{R}^n/\mathfrak{so}(p, q) = \mathbb{R}^n$:

$$\begin{array}{ccccc}
T_{\widehat{x}}\widehat{M} & \xrightarrow{D_{\widehat{x}}F} & T_{F(\widehat{x})}OM & & \\
\downarrow \pi_{\mathbb{R}^n} \circ \omega_{\widehat{x}} & \searrow D_{\widehat{x}}\pi & \swarrow D_{F(\widehat{x})}\pi' & & \downarrow \pi_{\mathbb{R}^n} \circ \omega'_{F(\widehat{x})} \\
& & T_{\pi(\widehat{x})}M & & \\
& & \uparrow \Theta([\widehat{x}, \cdot]) & \searrow \mathcal{E}_{F(\widehat{x})} & \\
\mathfrak{h} & \xlongequal{\quad} & \mathfrak{g}/\mathfrak{p} & \xlongequal{\quad} & \mathfrak{h}
\end{array}$$

We show that every block is commutative. The upper triangle is commutative because F is a bundle map over the identity, the right triangle is commutative as the route $D_{F(\widehat{x})}\pi' \circ \mathcal{E}_{F(\widehat{x})}$ is how the \mathbb{R}^n valued part of $\omega'_{F(\widehat{x})}$ is defined (see Definition 3.2.8). The lower left square is commutative because $\Theta([\widehat{x}, v])$ was defined as $D_{\widehat{x}}(\omega_{\widehat{x}}^{-1}v)$. It remains to check that the triangle on the lower right is commutative. Here we note that $F(\widehat{x})$ is the basis of $\pi(\widehat{x})$ given by $(\Theta([\widehat{x}, e^1]), \dots, \Theta([\widehat{x}, e^n]))$ as per the definition of F in Definition 3.2.21. Then $\mathcal{E}_{F(\widehat{x})}(\Theta([\widehat{x}, e^i])) = e^i$ by the definition of $\mathcal{E}_{F(\widehat{x})}$ in Definition 3.2.8, implying that these two maps are inverses of each other and that the lower right triangle commutes.

Thus the entire diagram commutes and $\pi_{\mathbb{R}^n} \circ \omega = F^*(\pi_{\mathbb{R}^n} \circ \omega')$, completing the step that the bundle-isomorphism (over id) F satisfies $F^*(\omega') = \omega$, i.e. that F is an isomorphism of Cartan geometries (over id).

What is left is to prove that this isomorphism is natural (from the identity functor to the functor $\omega \rightarrow (g, \nabla) \rightarrow \omega'$). Let $\Phi : \widehat{M} \rightarrow \widehat{N}$ be an isomorphism between two Cartan geometries over the

base map $\phi : M \rightarrow N$, and $O(\phi) : OM \rightarrow ON$ the isomorphism between OM and ON induced by ϕ . Then $F_N \circ O(\phi) \circ F_M^{-1}$ is an isomorphism between OM and ON over the base map ϕ , by Lemma 2.4.3 it must then be equal to $O(\phi)$, meaning the following diagram commutes:

$$\begin{array}{ccc} \widehat{M} & \xrightarrow{\Phi} & \widehat{N} \\ \downarrow F_M & & \downarrow F_N \\ OM & \xrightarrow{O(\phi)} & ON \end{array}$$

In other words the isomorphism F is natural. \square

It now remains to investigate which Cartan geometries correspond to the case where the metric connection is the Levi-Civita connection. The relevant statement is that for a metric g the Levi-Civita connection ∇^{LC} is the unique torsion free metric connection.

We will briefly describe how torsion may be described on the associated Cartan connection. For $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ a Cartan geometry modelled on $O(p, q) \times \mathbb{R}^n / O(p, q)$ denote with $\alpha : T\widehat{M} \rightarrow \mathfrak{so}(p, q)$ the map $\pi_{\mathfrak{so}(p, q)} \circ \omega$ and with $\theta : T\widehat{M} \rightarrow \mathbb{R}^n$ the map $\pi_{\mathbb{R}^n} \circ \omega$. Denote with $h : T\widehat{M} \rightarrow T\widehat{M}$ the projection onto the horizontal bundle $H = \ker(\alpha)$ given by the decomposition $T\widehat{M} = H \oplus \ker(D\pi)$. Then the torsion form $\tilde{\Theta} \in \Lambda^2 T\widehat{M}$ is defined by:

$$\tilde{\Theta}(v, w) = (d\theta)(hv, hw).$$

We remark upon the following Theorem:

Theorem 3.2.28. *If $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ is the Cartan connection associated to a Riemannian metric g and metric connection ∇ on M , then ω is torsion free (read $\tilde{\Theta} = 0$) if and only if ∇ is torsion free.*

This theorem may be formulated in the more general context of linear connections and their associated Ehresmann connections and fundamental forms. We will not prove it here, referring instead to Theorem 5.1 of Chapter III in [KN 63]. The theorem allows us to complete the formulation of the second equivalence initiated at the beginning of this entire section:

Definition 3.2.29. A Cartan connection $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ is called torsionless if the torsion form $\tilde{\Theta}$ vanishes.

Chapter 4

Tool: Generalised curvature and the integrability locus

In this section we consider a manifold M together with a Cartan connection $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ modelled on G/P . We will construct a way to view the derivatives of a map $\widehat{M} \rightarrow V$ for V a vector space as a function $\widehat{M} \rightarrow \text{Hom}(\mathfrak{g}, V)$, as opposed to a V -valued 1-form $T\widehat{M} \rightarrow V$. With this one can define a notion keeping track of the first $\dim(\mathfrak{g})$ derivatives of the curvature that is easy to manage, it will be called the generalised curvature.

The generalised curvature will be used to formulate a very useful theorem, the integrability theorem, which describes an open dense subset M^{int} of M on which any “infinitesimal” symmetry of the generalised curvature integrates to a Killing field. A consequence of the integrability theorem is that the sheaf of local Killing fields \mathfrak{kil} is locally constant on the components of M^{int} , which is a useful property as it implies Killing fields can always be uniquely developed in any direction inside of M^{int} (although this development may have a monodromy). We stress that the main point of the theorem is however *not* the *existence* of a dense open set on which \mathfrak{kil} is locally constant, we have already seen by elementary means that the “analyticity locus” \mathfrak{M} (cf. Definition 2.5.11) has this property.

The main use of the theorem is the explicit characterisation (inside of M^{int}) of $\mathfrak{kil}_{\pi(\widehat{x})}$ by $\ker(D_{\widehat{x}}\mathcal{D}\kappa)$, giving a simple and definite criterium of when a vector in $T_{\widehat{x}}\widehat{M}$ can be extended to the lift of a local Killing field inside of M^{int} . The fact that $M^{\text{int}} \subseteq \mathfrak{M}$ (the statement that \mathfrak{kil} is locally constant on M^{int}) is also very useful, as it serves to make the set M^{int} “even nicer”.

The integrability theorem appears in different forms in the literature. It was originally shown in [Gr 88] in the more general context of rigid geometrical structures (cf. Theorem 1.6.F in [Gr 88]). In the context of Cartan geometries a version of the integrability theorem was proven in [Me 11] for analytic Cartan geometries, [Pe 16] extended this result to smooth Cartan geometries. [Fr 16] provides a proof of a result that is slightly stronger than that of [Pe 16], to be precise the result of [Fr 16] is (in the notation of this chapter): $\widehat{M}^{\text{int}} \supseteq CR(\mathcal{D}\kappa)$. The other inclusion $\widehat{M}^{\text{int}} \subseteq CR(\mathcal{D}\kappa)$ is however elementary, and we carry out in Lemma 4.3.3, giving the formulation of the integrability theorem as in Theorem 4.3.2.

4.1 Covariant derivative

Definition 4.1.1 (Derivatives). For a smooth map $F : \widehat{M} \rightarrow V$ into a vector space we denote with $\mathcal{D}F$ the map $\widehat{M} \rightarrow \text{Hom}(\mathfrak{g}, V)$ induced from DF via the parallelism ω , that is

$$\mathcal{D}F : \widehat{M} \rightarrow \text{Hom}(\mathfrak{g}, V), \quad \widehat{x} \mapsto D_{\widehat{x}}F \circ \omega_{\widehat{x}}^{-1}.$$

We use $\mathcal{D}F$ as notation to prevent confusion with the usual derivative DF . In fact this procedure of parallelising DF generalises, allowing us to view any multilinear form on \widehat{M} as a function.

Remark. Let $\eta : T\widehat{M}^{\otimes k} \rightarrow V$ be a k -multilinear form on the tangent bundle valued in some vector space V . Using the parallelism ω we can view this as a map $\widehat{M} \rightarrow \text{Hom}(\mathfrak{g}^{\otimes k}, V)$ defined via

$$\widehat{x} \mapsto [(\mathbf{v}_1, \dots, \mathbf{v}_k) \mapsto \eta_{\widehat{x}}(\omega_{\widehat{x}}^{-1}\mathbf{v}_1, \dots, \omega_{\widehat{x}}^{-1}\mathbf{v}_k)].$$

We may further simplify the form of higher derivatives by making use of the fact that $\text{Hom}(V, \text{Hom}(W, U)) \cong V^* \otimes (W^* \otimes U) \cong (V \otimes W)^* \otimes U \cong \text{Hom}(V \otimes W, U)$ if V, W, U are finite dimensional vector spaces, here all identifications are canonical.

Definition 4.1.2 (Higher derivatives). Let V be a finite dimensional real vector space, and $F : \widehat{M} \rightarrow V$ smooth. For $k \in \mathbb{N}$ we inductively define $\mathcal{D}^k F : \widehat{M} \rightarrow \text{Hom}(\mathfrak{g}^{\otimes k}, V)$ via $\mathcal{D}^1 F = \mathcal{D}F$ and $\mathcal{D}^k F$ to be $\mathcal{D}\mathcal{D}^{k-1}F$ composed with the isomorphism taking $\text{Hom}(\mathfrak{g}, \text{Hom}(\mathfrak{g}^{\otimes k-1}, V))$ to $\text{Hom}(\mathfrak{g}^{\otimes k}, V)$.

We now introduce notation that keeps track of a function and all its derivatives in a simple manner.

Definition 4.1.3 (Covariant derivative). Let V be a finite dimensional real vector space and $F : \widehat{M} \rightarrow V$ smooth. For $k \in \mathbb{N}$ we define:

$$\mathcal{D}^k F : \widehat{M} \rightarrow V \oplus \bigoplus_{i=1}^k \text{Hom}(\mathfrak{g}^{\otimes i}, V), \quad \widehat{x} \mapsto (F(\widehat{x}), \mathcal{D}F(\widehat{x}), \dots, \mathcal{D}^k F(\widehat{x})).$$

$\mathcal{D}^k F$ is called the k -th covariant derivative of F .

Remark. If V, W are finite dimensional real vector spaces equipped with a representations ρ_V, ρ_W of P , then we equip $\text{Hom}(W, V)$ with the representation $\rho(g)A = \rho_W(g) \cdot A \cdot \rho_V(g^{-1})$.

Definition 4.1.4. Let V admit a representation ρ of P , we say a function $F : \widehat{M} \rightarrow V$ is P -equivariant if $F(\widehat{x} \cdot g) = \rho(g)F(\widehat{x})$ for all $\widehat{x} \in \widehat{M}$ and all $g \in P$.

We equip \mathfrak{g} with the adjoint representation of P , which allows us to formulate the following proposition:

Proposition 4.1.5. *Let V be a real vector space equipped with a representation ρ of P . Suppose $F : \widehat{M} \rightarrow V$ is P -equivariant, then $\mathcal{D}F : \widehat{M} \rightarrow \text{Hom}(\mathfrak{g}, V)$ is also P -equivariant.*

Proof. For $\mathbf{v} \in \mathfrak{g}$ let $\phi_t^{\mathbf{v}}$ be the flow of the constant field $\omega^{-1}\mathbf{v}$. First we check that

$$\phi_t^{\mathbf{v}}(\widehat{x}g) = \phi_t^{\text{Ad}(g)\mathbf{v}}(\widehat{x})g,$$

this follows from checking the derivatives of the both flows agree:

$$\frac{d}{dt}\phi_t^{\mathbf{v}}(\widehat{x}g)|_{t=0} = \omega_{\widehat{x}g}^{-1}(\mathbf{v}) = D_{\widehat{x}}R_g \omega_{\widehat{x}}^{-1}(\text{Ad}(g)\mathbf{v}) = \frac{d}{dt}\phi_t^{\text{Ad}(g)\mathbf{v}}(\widehat{x})g|_{t=0}.$$

(Note that $\omega_{\widehat{x}g} D_{\widehat{x}}R_g = \text{Ad}(g^{-1})\omega_{\widehat{x}}$, from whence $\omega_{\widehat{x}g}^{-1} = D_{\widehat{x}}R_g \omega_{\widehat{x}}^{-1} \text{Ad}(g)$.) Now

$$\begin{aligned} \mathcal{D}F(\widehat{x}g)[\mathbf{v}] &= D_{\widehat{x}g}F(\omega_{\widehat{x}g}^{-1}\mathbf{v}) = \frac{d}{dt}F(\phi_t^{\mathbf{v}}(\widehat{x}g)) = \frac{d}{dt}F(\phi_t^{\text{Ad}(g)\mathbf{v}}(\widehat{x})g) = \frac{d}{dt}[\rho_V(g^{-1})F(\phi_t^{\text{Ad}(g)\mathbf{v}}(\widehat{x}))] \\ &= \rho_V(g^{-1})(\mathcal{D}F(\widehat{x})[\text{Ad}(g)\mathbf{v}]) = (\rho(g^{-1})\mathcal{D}F(\widehat{x}))[\mathbf{v}], \end{aligned}$$

which is the desired statement. □

Corollary 4.1.6. *If $F : \widehat{M} \rightarrow V$ is P -equivariant then for any k both $\mathcal{D}^k F$ and $\mathcal{D}^k F$ are also P -equivariant.*

4.2 Generalised curvature

Reminder. The curvature of ω is defined as the \mathfrak{g} valued 2-form

$$\kappa(v, w) = d\omega(v, w) + [\omega(v), \omega(w)]_{\mathfrak{g}}.$$

In the following it will always be regarded as a map $\widehat{M} \rightarrow \text{Hom}(\Lambda^2\mathfrak{g}, \mathfrak{g})$.

Proposition 4.2.1. *The curvature is equivariant under the action of P , meaning that*

$$\kappa(\widehat{x} \cdot g) = \rho(g^{-1}) \cdot \kappa(\widehat{x}).$$

Proof. The compatibility of ω states

$$R_g^*(d\omega) = d(R_g^*\omega) = d(\text{Ad}(g^{-1}) \cdot \omega) = \text{ad}(g^{-1})d\omega,$$

from which $R_g^*(\kappa) = R_g^*d\omega + R_g^*\frac{1}{2}[\omega, \omega]_{\mathfrak{g}} = \text{Ad}(g^{-1})\kappa$ follows. \square

Definition 4.2.2 (Generalised curvature). We define $\mathcal{D}\kappa$ as the map $\mathcal{D}^{\dim(\mathfrak{g})}\kappa$ and call it the generalised curvature. Further $\text{Hom}(\mathfrak{g}^{\otimes k}, \text{Hom}(\Lambda^2\mathfrak{g}, \mathfrak{g}))$ will be denoted by \mathcal{W}_k and $\mathcal{W} := \bigoplus_{k=0}^n \mathcal{W}_k$, so that $\mathcal{D}\kappa$ is a map $\widehat{M} \rightarrow \mathcal{W}$.

Definition 4.2.3. We denote with $CR(\mathcal{D}\kappa)$ the set of points \widehat{x} in \widehat{M} for which the differential $D_{\widehat{y}}\mathcal{D}\kappa$ has constant rank for all points \widehat{y} in a neighbourhood of \widehat{x} .

Remark. By definition $CR(\mathcal{D}\kappa)$ is open and by basic analysis it is dense in \widehat{M} .

4.3 Integrability theorem

Definition 4.3.1 (Integrability locus). We denote with \widehat{M}^{int} those points \widehat{x} for which $\xi \in \ker(D_{\widehat{x}}\mathcal{D}\kappa)$ implies that there exists a local Killing field X around $\pi(\widehat{x})$ with lift $\widehat{X}_{\widehat{x}} = \xi$. \widehat{M}^{int} is called the integrability locus. We define $M^{\text{int}} := \pi(\widehat{M}^{\text{int}})$, which is also called the integrability locus.

Remark. $\widehat{M}^{\text{int}} = \pi^{-1}(M^{\text{int}})$, which follows from $\mathcal{D}\kappa$ being P -equivariant and Killing fields being right-invariant.

Theorem 4.3.2 (Integrability theorem). *Let $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ be a Cartan connection modelled on G/P . We have*

$$\boxed{\widehat{M}^{\text{int}} = CR(\mathcal{D}\kappa).}$$

In particular since $CR(\mathcal{D}\kappa)$ is open and dense, so is \widehat{M}^{int} . We briefly remark that the inclusion $\widehat{M}^{\text{int}} \subseteq CR(\mathcal{D}\kappa)$ is elementary (and not part of the formulation of the theorem in [Fr 16]). We carry out this step in a Lemma:

Lemma 4.3.3. $\widehat{M}^{\text{int}} \subseteq CR(\mathcal{D}\kappa)$.

Proof. We first remark that in general one has

$$\dim(\mathfrak{kill}_{\pi(\widehat{x})}) \leq \dim(\ker(D_{\widehat{x}}\mathcal{D}\kappa))$$

for any $\widehat{x} \in \widehat{M}$, which follows from $\mathcal{D}\kappa$ being invariant under automorphisms, hence the lift of any Killing field to \widehat{M} is in the kernel of $D\mathcal{D}\kappa$, and from Killing fields being uniquely determined by their lifts. Further any point $\widehat{x} \in \widehat{M}$ admits a neighbourhood $U \subseteq \widehat{M}$ so that

$$\dim(\mathfrak{kill}_{\pi(\widehat{x})}) \leq \dim(\mathfrak{kill}_{\pi(\widehat{y})}), \quad \dim(\ker(D_{\widehat{y}}\mathcal{D}\kappa)) \leq \dim(\ker(D_{\widehat{x}}\mathcal{D}\kappa))$$

for all $\widehat{y} \in U$. This follows from the fact that the dimension of the stalks of \mathfrak{kil} can only increase locally (cf. Proposition 2.5.10) and the fact that the rank of a smoothly parametrised linear map can only increase locally (hence the dimension of its kernel can only decrease locally). Finally if $\widehat{x} \in \widehat{M}^{\text{int}}$ then $\dim(\mathfrak{kil}_{\pi(x)}) = \dim(\ker(D_{\widehat{x}}\mathcal{D}\kappa))$ and we can combine the previous inequalities on the special set U to get for any $y \in U$:

$$\dim(\mathfrak{kil}_{\pi(\widehat{y})}) \leq \dim(\ker(D_{\widehat{y}}\mathcal{D}\kappa)) \leq \dim(\ker(D_{\widehat{x}}\mathcal{D}\kappa)) = \dim(\mathfrak{kil}_{\pi(\widehat{x})}) \leq \dim(\mathfrak{kil}_{\pi(\widehat{y})}),$$

so every \leq must be an $=$, and the dimension of the kernel of $D\mathcal{D}\kappa$ must be constant in a neighbourhood of \widehat{x} , giving $\widehat{x} \in CR(\mathcal{D}\kappa)$. \square

The proof of the other inclusion $CR(\mathcal{D}\kappa) \subseteq \widehat{M}^{\text{int}}$ is more involved. We refer to Annex A of [Fr 16], and to Theorem 2 of [Pe 16] (proven in Section 4), which has a more detailed discussion of a closely related formulation.

4.4 $M^{\text{int}} \subseteq \mathfrak{M}$

In the beginning of this chapter as well as in the discussion following the definition of \mathfrak{M} (cf. Definition 2.5.11) it was remarked that the sheaf of local Killing fields is locally constant on the components of M^{int} . While the conclusion is elementary, it is a good idea to carry out the details of this useful property.

Proposition 4.4.1. *The function $M \rightarrow \mathbb{N}, x \mapsto \dim(\mathfrak{kil}_x)$ is locally constant inside of M^{int} .*

Proof. Let $x \in M^{\text{int}}$, from the definition of M^{int} it follows that $\dim(\mathfrak{kil}_x) = \dim(\ker(D_{\widehat{x}}\mathcal{D}\kappa))$ for a point $\widehat{x} \in \pi^{-1}(x)$. However $\dim(\ker(D_{\widehat{x}}\mathcal{D}\kappa)) = \dim(\mathfrak{g}) - \text{Rank}(D_{\widehat{x}}\mathcal{D}\kappa)$, which is locally constant inside of \widehat{M}^{int} by the integrability theorem. \square

Proposition 4.4.2. *Let \mathcal{M} be a connected component of M^{int} , then the sheaf $\mathfrak{kil}|_{\mathcal{M}}$ is a locally constant sheaf.*

Proof. Let $x \in \mathcal{M}$ and $U_x \subset \mathcal{M}$ be a neighbourhood of x as in Proposition 2.5.9, that is a neighbourhood so that the restriction $\mathfrak{kil}(U_x) \rightarrow \mathfrak{kil}_x$ is an isomorphism. Since the restriction $\mathfrak{kil}(U_x) \rightarrow \mathfrak{kil}_y$ for any $y \in U_x$ is injective and by the previous proposition $\dim(\mathfrak{kil}(U_x)) = \dim(\mathfrak{kil}_x) = \dim(\mathfrak{kil}_y)$ one has that the restriction $\mathfrak{kil}(U_x) \rightarrow \mathfrak{kil}_y$ is an isomorphism. Since \mathcal{M} is connected this implies that $\mathfrak{kil}|_{\mathcal{M}}$ is a locally constant sheaf. \square

Indeed the statement that $\mathfrak{kil}(U_x) \rightarrow \mathfrak{kil}_y$ is an isomorphism for every $y \in U_x$ further implies:

Corollary 4.4.3. $M^{\text{int}} \subseteq \mathfrak{M}$.

Chapter 5

Classification Theorem: Closed Lorentz 3-folds with non-compact isometry group

5.1 The classification

[Fr 18] establishes a complete classification of all compact Lorentz 3-manifolds with non-compact isometry group. Strictly speaking, the orientable and time-orientable compact Lorentz 3-manifolds with non-compact isometry group are classified, but any other Lorentz 3-manifold with non-compact isometry group will admit a cover of order 4 by such a manifold.

[Fr 18] finds that, topologically, any such manifold must be diffeomorphic to:

Theorem 5.1.1 (Topological classification). *Let (M, g) be a smooth, oriented and time-oriented closed 3-dimensional Lorentz manifold. If $\text{Isom}(M, g)$ is not compact, then M is diffeomorphic to:*

1. *A quotient $\Gamma \backslash \widetilde{SL}_2(\mathbb{R})$, where $\Gamma \subset \widetilde{SL}_2(\mathbb{R})$ is any uniform lattice.*
2. *The 3-torus \mathbb{T}^3 or a torus bundle \mathbb{T}_A^3 where $A \in SL_2(\mathbb{Z})$ can be any hyperbolic or parabolic element.*

Here $\widetilde{SL}_2(\mathbb{R})$ denotes the universal cover of $SL_2(\mathbb{R})$, which (as a smooth manifold) is diffeomorphic

to \mathbb{R}^3 . The torus bundles \mathbb{T}_A^3 also have \mathbb{R}^3 as universal cover, we briefly define them:

Definition 5.1.2 (Torus bundle). Let $A \in SL_2(\mathbb{Z})$. The torus bundle \mathbb{T}_A^3 is defined to be the quotient $\mathbb{R}^3/(\mathbb{Z} \times_A \mathbb{Z}^2)$, where $\mathbb{Z} \times_A \mathbb{Z}^2$ the group structure on \mathbb{Z}^3 with multiplication defined by:

$$\begin{bmatrix} n_1, \\ \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} \end{bmatrix} \cdot \begin{bmatrix} n_2, \\ \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} \end{bmatrix} = \begin{bmatrix} n_1 + n_2, \\ \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + A^{n_1} \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} \end{bmatrix},$$

which acts on \mathbb{R}^3 by:

$$\begin{bmatrix} n, \\ \begin{pmatrix} a \\ b \end{pmatrix} \end{bmatrix} (t, x, y) := (t + n, a + (A^n)_{11}x + (A^n)_{12}y, b + (A^n)_{21}x + (A^n)_{22}y).$$

It is most convenient to describe what metrics may appear by giving their form them on the universal cover. The geometric form the classification theorem is here:

Theorem 5.1.3. *Let (M, g) be a smooth, oriented and time-oriented closed 3-dimensional Lorentz manifold and suppose $\text{Isom}(M, g)$ is not compact. Let $(\widetilde{M}, \widetilde{g})$ denote the universal cover of (M, g) , then:*

1. *If $M \cong \Gamma \backslash \widetilde{SL}_2(\mathbb{R})$, then \widetilde{g} is a Lorentzian, non-Riemannian, left-invariant metric on $\widetilde{SL}_2(\mathbb{R})$.*
2. *If $M \cong \mathbb{T}^3$ or \mathbb{T}_A^3 , then there exists a 1-periodic function $a : \mathbb{R} \rightarrow (0, \infty)$ so that \widetilde{g} is isometric to one of the following metrics on \mathbb{R}^3 :*

$$dt^2 + 2a(t)dxdy \quad \text{or} \quad a(x)(dt^2 + 2dxdy).$$

If M is a hyperbolic torus bundle, only the first case can occur, if M is a parabolic torus bundle, only the second case can occur. If \widetilde{g} is locally homogenous (meaning a is constant), then it is flat (which can occur in all cases) or modelled on Lorentz-Heisenberg geometry (which can only occur in the case of a parabolic torus bundle).

We remark briefly what it means for a metric to be non-Riemannian:

Definition 5.1.4. A Lorentz-metric on a connected manifold is called non-Riemannian at a point x if the flows of those Killing fields vanishing at x generated a non-compact subgroup of $O(T_x M)$. The metric is called non-Riemannian if it is non-Riemannian at every point.

A further statement is that this classification does not contain any superfluous spaces, meaning that every case described by the classification is actually realised.

The theorem relies heavily on the fact that compact Lorentz-manifolds with non-compact isometry groups must always have a large amount of continuous local symmetries, even if the isometry

group itself is discrete. In fact we will find, using a recurrence argument, that the metric must be non-Riemannian at almost every point of the manifold (if the manifold is connected). Additional arguments will yield that almost every point will admit a connected open neighbourhood on which the Killing algebra is at least 3 dimensional.

With these Killing fields one considers the partition of the manifold into the connected components of the integrability locus, which is described in Theorem 4.3.2. As remarked in Proposition 4.4.2 the sheaf of local Killing fields is locally constant on these components, and one can introduce a classification of the different components by their Killing algebras. The classification is described in Section 5.3.

The proof of the main theorems then is divided into 3 cases, depending on what kind of components (classified by their Killing algebras) appear in the manifold. In the first case one restricts all components to be locally homogenous and makes use of general structure theorems about locally homogenous 3-dimensional Lorentz geometries. For the second case one considers at least one component to be of “hyperbolic” type and for not all components to be locally homogenous. In this case one is able to find an embedded Lorentz-torus in the manifold and one can push it along its normal flow, which will recover all of M . In the last case one assumes that no hyperbolic components exist and the manifold is not locally homogenous. Here one must show that the manifold is then conformally flat, from which one can gain the classification result.

5.2 Recursion leads to Killing fields

For most of this section we remain in the setting of a general Cartan geometry (modelled on G/P). For the very last result we specify further to pseudo-Riemannian geometry of dimension > 2 . This section establishes that a compact space with a non-compact automorphism group and a measure that is automorphism-invariant possesses at least two locally defined Killing fields at every point¹, one of which is vertical at that point. An argument specific to pseudo-Riemannian geometry then generates a third locally defined Killing field out of the preceding two.

For the coming section M shall be a manifold together with a Cartan connection $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ modelled on G/P and we define $\mathfrak{g} = \text{Lie}(G)$, $\mathfrak{p} = \text{Lie}(P)$.

A central tool for this section is the integrability locus and the generalised curvature of a Cartan

¹More correctly: Every point so that every open neighbourhood of this point has non-zero measure has two independent locally defined Killing fields. This technicality will be overlooked by assuming that every open set has non-zero measure

connection. These were introduced in Chapter 4.

5.2.1 Isotropy algebra

Reminder. We denote with \mathfrak{kil} the sheaf of local Killing fields on M . When restricted to the components of M^{int} this is a locally constant sheaf, as by Proposition 4.4.2.

Definition 5.2.1 (Isotropy algebra). Let $x \in M$. We define the Isotropy algebra $\mathfrak{Is}(x)$ as:

$$\mathfrak{Is}(x) = \{X \in \mathfrak{kil}_x \mid X_x = 0\}.$$

That is $\mathfrak{Is}(x)$ consists of those local Killing fields that vanish at x .

Remark. The isotropy algebra $\mathfrak{Is}(x)$ is a Lie sub-algebra of \mathfrak{kil}_x , which follows from $[X, Y]_x = 0$ for two vector fields X, Y if both $X_x = 0$ and $Y_x = 0$.

Every local Killing field X on M lifts to a unique locally defined vector field on \widehat{M} , denoted by \widehat{X} , so that the flow is by local automorphisms, this was Proposition 2.5.4. If the Killing field \widehat{X} is zero at a point x , then the lift $\widehat{X}_{\widehat{x}}$ must be vertical for every $\widehat{x} \in \pi^{-1}(x)$. In particular $\omega_{\widehat{x}}(\widehat{X}_{\widehat{x}})$ lies in \mathfrak{p} .

Definition 5.2.2. Let $\widehat{x} \in \widehat{M}$ and $x = \pi(\widehat{x})$. We denote with $\phi_{\widehat{x}}$ the map

$$\phi_{\widehat{x}} : \mathfrak{Is}(x) \rightarrow \mathfrak{p}, \quad X \mapsto \omega_{\widehat{x}}(\widehat{X}_{\widehat{x}}).$$

Proposition 5.2.3. For every $\widehat{x} \in \widehat{M}$ the map $\phi_{\widehat{x}}$ is an injective Lie-algebra anti-homomorphism, that is $\phi_{\widehat{x}}([X, Y]_{\mathfrak{Is}(x)}) = [\phi_{\widehat{x}}(Y), \phi_{\widehat{x}}(X)]_{\mathfrak{p}}$ for any $X, Y \in \mathfrak{Is}(x)$.

Proof. Injectivity is immediate, since $\mathfrak{Is}(x) \rightarrow T_{\widehat{x}}\widehat{M}$, $X \mapsto \widehat{X}_{\widehat{x}}$ is injective (local Killing fields are uniquely determined by their lifts to a point, see Proposition 2.5.4) and $\omega_{\widehat{x}} : T_{\widehat{x}}\widehat{M} \rightarrow \mathfrak{g}$ is injective.

To see that $\phi_{\widehat{x}}$ is an anti-homomorphism of Lie-algebras let $X, Y \in \mathfrak{Is}(x)$. Then using Cartan's formula $L_X = \iota_X \circ d + d \circ \iota_X$ we note

$$\begin{aligned} \widehat{X}(\omega(\widehat{Y})) - \widehat{Y}(\omega(\widehat{X})) &= (\iota_{\widehat{X}} \circ d \circ \iota_{\widehat{Y}})(\omega) - (\iota_{\widehat{Y}} \circ d \circ \iota_{\widehat{X}})(\omega) \\ &= (-\iota_{\widehat{X}} \circ \iota_{\widehat{Y}})(d\omega) - (-\iota_{\widehat{Y}} \circ \iota_{\widehat{X}})(d\omega) + \iota_{\widehat{X}}(L_{\widehat{Y}}(\omega)) + \iota_{\widehat{Y}}(L_{\widehat{X}}(\omega)) \\ &= -d\omega(\widehat{Y}, \widehat{X}) + d\omega(\widehat{X}, \widehat{Y}) = 2d\omega(\widehat{X}, \widehat{Y}), \end{aligned}$$

where $L_{\widehat{X}}\omega = 0 = L_{\widehat{Y}}\omega$ since \widehat{X}, \widehat{Y} flow by automorphisms. The definition of the exterior derivative then gives:

$$d\omega(\widehat{X}, \widehat{Y}) = \widehat{X}(\omega(\widehat{Y})) - \widehat{Y}(\omega(\widehat{X})) - \omega([\widehat{X}, \widehat{Y}]) = 2d\omega(\widehat{X}, \widehat{Y}) - \omega([\widehat{X}, \widehat{Y}])$$

in other words $\omega([\widehat{X}, \widehat{Y}]) = d\omega(\widehat{X}, \widehat{Y})$. Now the curvature was defined as the two-form $\kappa(X, Y) = d\omega(X, Y) + [\omega(X), \omega(Y)]_{\mathfrak{g}}$. Using this and the fact that the curvature is zero on vertical vectors (cf. Proposition 2.6.4), as well as $\widehat{X}_{\widehat{x}}$ and $\widehat{Y}_{\widehat{x}}$ being vertical, we find:

$$[\phi_{\widehat{x}}(X), \phi_{\widehat{x}}(Y)]_{\mathfrak{p}} = [\omega_{\widehat{x}}(\widehat{X}_{\widehat{x}}), \omega_{\widehat{x}}(\widehat{Y}_{\widehat{x}})]_{\mathfrak{g}} = -(d\omega)_{\widehat{x}}(\widehat{X}_{\widehat{x}}, \widehat{Y}_{\widehat{x}}) = -\omega_{\widehat{x}}([\widehat{X}, \widehat{Y}]_{\widehat{x}}).$$

The only thing left to see is that $\widehat{[X, Y]} = [\widehat{X}, \widehat{Y}]$, meaning that the lift of the commutator is the commutator of the lifts. This was already shown in Proposition 2.5.7, where it followed from the relation $D\pi([\widehat{X}, \widehat{Y}]) = [D\pi(\widehat{X}), D\pi(\widehat{Y})]$ for right-invariant \widehat{X}, \widehat{Y} . \square

Proposition 5.2.4. *Let $x \in M^{\text{int}}$ and $\widehat{x} \in \pi^{-1}(x)$, denote the stabiliser of $\mathcal{D}\kappa(x) \in \mathcal{W}$ under the action of P with $P_{\mathcal{D}\kappa(\widehat{x})}$ (recall that \mathcal{W} is the codomain of $\mathcal{D}\kappa : \widehat{M} \rightarrow \mathcal{W}$ and is equipped with a representation of P under which $\mathcal{D}\kappa$ is P -equivariant). Then $\phi_{\widehat{x}}(\mathfrak{Is}(x)) = \text{Lie}(P_{\mathcal{D}\kappa(\widehat{x})})$.*

Proof. If $X \in \mathfrak{Is}(x)$ then the lift \widehat{X} is vertical at \widehat{x} . Since \widehat{X} must also be right-invariant by Lemma 2.5.5, \widehat{X} restricted to $\pi^{-1}(x)$ must be equal to a fundamental field. In particular its flow must be equal to $\varphi_{\widehat{X}}^t(\widehat{x}) = \widehat{x} \cdot \exp(t\omega_{\widehat{x}}(\widehat{X}_{\widehat{x}})) = \widehat{x} \cdot \exp(t\phi_{\widehat{x}}(X))$. Then this flow must be by isometries, and (for ρ the representation on \mathcal{W})

$$\rho(e^{t\phi_{\widehat{x}}(X)})\mathcal{D}\kappa(\widehat{x}) = \mathcal{D}\kappa(\widehat{x}e^{-t\phi_{\widehat{x}}(X)}) = \mathcal{D}\kappa(\widehat{x}),$$

implying $\phi_{\widehat{x}}(X) \in \text{Lie}(P_{\mathcal{D}\kappa(\widehat{x})})$.

On the other hand if $\xi \in \text{Lie}(P_{\mathcal{D}\kappa(\widehat{x})})$, then $\mathcal{D}\kappa(\widehat{x}e^{-t\xi}) = 0$, implying, by the integrability theorem, that $\omega_{\widehat{x}}^{-1}$ can be extended to the lift of a Killing field around \widehat{x} . But since $\omega_{\widehat{x}}^{-1}(\xi)$ is vertical, this Killing field must be in $\mathfrak{Is}(x)$ and $\phi_{\widehat{x}}$ applied to this field must be ξ , implying $\phi_{\widehat{x}}(\mathfrak{Is}(x)) \supseteq \text{Lie}(P_{\mathcal{D}\kappa(\widehat{x})})$. \square

These two propositions show the following statement:

Corollary 5.2.5. *For any $x \in M^{\text{int}}$ and $\widehat{x} \in \pi^{-1}(x)$ the map $\phi_{\widehat{x}}$ is an anti-isomorphism of Lie-algebras from $\mathfrak{Is}(x)$ to $\text{Lie}(P_{\mathcal{D}\kappa(\widehat{x})})$.*

5.2.2 Recurrence argument

In this section we specialise to the case of M having a non-compact automorphism group as well as being equipped with a finite automorphism invariant Borel measure, denoted with μ . A recurrence argument will provide us with a non-zero isotropy algebra $\mathfrak{Is}(x)$ at every point x for which all open neighbourhoods have non-zero measure. This last technicality will however be

overridden in the statements and proofs by further restricting the measure μ to be non-zero when evaluated on any open set, so that every point $x \in M$ has non-zero isotropy algebra $\mathfrak{Is}(x)$.

First we remark that for the situation we are interested, that is compact 3-dimensional Lorentz manifolds or $G/P = O(2, 1) \times \mathbb{R}^n / O(2, 1)$, such a measure μ exists and may be chosen canonically.

Remark. Let (M, g) be a pseudo-Riemannian manifold with metric g . M can be canonically equipped with an isometry-invariant measure, given in coordinates by the density $\sqrt{|\det(g)|}$ (for any open set U parametrised by coordinates x^i one has $\mu(U) = \int_U \sqrt{|\det(g)|} dx^1 \dots dx^n$, for larger sets one works with a partition of unity). As every point admits a coordinate neighbourhood for which $\sqrt{|\det(g)|}$ is bounded, it follows that for a compact pseudo-Riemannian manifold this measure is finite and by construction invariant under isometries. Further since the measure is locally defined by a density it is a Borel measure.

The above situation may be generalised to the case in which the adjoint action of P on $\mathfrak{g}/\mathfrak{p}$ is by linear maps of determinant ± 1 . Then any positive density on $\mathfrak{g}/\mathfrak{p}$ is automatically invariant under the action of P : $\sigma(\text{Ad}(p)v_1, \dots, \text{Ad}(p)v_n) = |\det(\text{Ad}(p))| \sigma(v_1, \dots, v_n) = \sigma(v_1, \dots, v_n)$. If we choose one such density and pull it back to TM by use of the soldering form $\Theta : \widehat{M} \times_P \mathfrak{g}/\mathfrak{p} \xrightarrow{\cong} TM$ (cf. Definition 3.2.18) we retrieve a smooth density on TM and thus a Borel measure on M . Further, a calculation shows that if $F : \widehat{M} \rightarrow \widehat{M}$ is an automorphism of ω over base-map $f : M \rightarrow M$, that then $\Theta([F(\widehat{x}), v]) = D_{\pi(\widehat{x})} f(\Theta([\widehat{x}, v]))$ holds², which implies that the density is invariant under automorphisms of the Cartan connection.

Poincaré recurrence

This subsection is a short detour about Poincaré recurrence, ending with Proposition 5.2.9 which is directly applicable to our situation. The definitions 5.2.6 and 5.2.7 as well as Theorem 5.2.8 are from [FK 02].

Let G be a locally compact second countable group and (X, \mathcal{B}, μ) a measure space with measure μ and σ -algebra \mathcal{B} . In applying the results of this section X will be a topological space and \mathcal{B} the Borel σ -algebra, further G will be equipped with the Haar measure, which is also a Borel measure. So in the interest of preventing notational clutter any reference to σ -algebras will be dropped from now on; they are always the Borel σ -algebras.

²Make use of $DF \circ \omega = \omega$ and $\pi \circ F = f \circ \pi$ in:

$$\begin{aligned} \Theta([F(\widehat{x}), v]) &= (D_{F(\widehat{x})} \pi \circ \omega_{F(\widehat{x})}^{-1})(v) = (D_{F(\widehat{x})} \pi \circ D_{\widehat{x}} F \circ \omega_{\widehat{x}}^{-1})(v) \\ &= (D_{\pi(\widehat{x})} f \circ D_{\widehat{x}} \pi \circ \omega_{\widehat{x}}^{-1})(v) = D_{\pi(\widehat{x})} f(\Theta([\widehat{x}, v])). \end{aligned}$$

Definition 5.2.6. Let $\Phi : X \times G \rightarrow X, (x, g) \mapsto \Phi_g(x)$ be a (right) G -action on X .

1. The action Φ is called measurable if the map $\Phi : X \times G \rightarrow X$ is measurable.
2. The measure μ is called G -invariant if Φ is measurable and $(\Phi_g)_*(\mu) = \mu$ for all $g \in G$.
3. The measure μ is called quasi-invariant with respect to G if Φ is measurable and $(\Phi_g)_*(\mu)$ and μ have the same zero sets for all $g \in G$.

As a remark, if $f : X \rightarrow Y$ is a measurable map then $f_*(\mu)$ defines a measure on Y via $f_*(\mu)(A) = \mu(f^{-1}(A))$ for any measurable set $A \subset Y$.

Definition 5.2.7 (Recurrent action). Let G act measurably on X and let μ be quasi-invariant with respect to G . The action of G is called recurrent if for any measurable $A \subset X$ the set $R_G(x, A) = \{g \in G \mid \Phi_g(x) \in A\}$ does not have compact closure for almost all $x \in A$.

Remark. If $G = \mathbb{R}$ or \mathbb{Z} this definition is equivalent to the usual notion of recurrence. For \mathbb{R} and \mathbb{Z} the only sets that do not have compact closure are the unbounded sets. It follows that a set of the form $\{g \in G \mid x \cdot g \in A\}$ does not have compact closure if and only if for every $t \in \mathbb{R}$ (or $n \in \mathbb{N}$) there is a $g \in G$ with $|g| \geq |t|$ (or $\geq n$) so that $x \cdot g \in A$. This also motivates saying that x returns to A if $R_G(x, A)$ does not have compact closure.

An action of a group G is then recurrent if and only if for every measurable set A the set of points $x \in A$ returning to A has full measure in A .

Theorem 5.2.8 (Poincaré recurrence, [FK 02], Theorem 2.2.6.). *Suppose G is a non-compact second countable locally compact group equipped with a measurable action on X so that μ is finite and G -invariant. Then the action of G on X is recurrent.*

Proof. Let $A \subseteq X$ be measurable, the goal is to prove that the set

$$B := \{x \in A \mid \exists K \subset G \text{ compact s.t. } x \cdot g \notin A \text{ for all } g \in K^C\}$$

is contained in a measure zero set. Any point $x \in A$ for which $R_G(x, A)$ has compact closure must by definition lie in B , and so the set of points in A returning to A must have full measure in A (provided it is measurable of course).

We begin the proof by getting a general topological argument out of the way. We show: There exists a countable family of compact sets $K_i \subset G$ so that $K_i \subseteq K_{i+1}$ and if $K \subset G$ is any compact set then there is an i with $K \subseteq K_i$.

To show this we make use of G being second countable and locally compact. We choose a countable basis \mathfrak{B} of X . By local compactness every point $x \in G$ admits a neighbourhood U_x with compact closure with compact closure, by definition there must be a $b_x \in \mathfrak{B}$ with $x \in b_x \subset U_x$,

then $\overline{b_x}$ is a compact neighbourhood containing x . This defines a set $\{\overline{b_x}\}_{\overline{b_x} \text{ compact}, b_x \in \mathfrak{B}}$, which is indexed by a subset of \mathfrak{B} and thus countable. We give it a numbering and define $K_i = \bigcup_{k=1}^i \overline{b_{x_i}}$. We first remark that $\bigcup_i K_i = G$, which follows from any $x \in G$ being contained in one of the b_x . Secondly if $K \subset G$ is compact, then $C = \{b_x \mid x \in K\}$ forms a covering of K by open sets, choosing a finite sub-covering C' and $j := \max\{i \in \mathbb{N} \mid b_{x_i} \in C'\}$ results in $K_j \supseteq K$.

The first consequence of this is that

$$B = \{x \in A \mid \exists i \in \mathbb{N} : x \cdot g \notin A \text{ for all } g \in K_i^C\},$$

for if there is a K with $x \cdot g \notin A$ for all $g \in K^C$, then there is an i with $K_i \supseteq K$, implying $K_i^C \subseteq K^C$ and then $xg \notin A$ for all $g \in K_i^C$, implying “ \subseteq ”. The relation “ \supseteq ” is however clear by the definition of B .

Now we show that B is contained in a measure zero set. Choose a countable dense subset L of G , since G is locally compact and second countable G is separable and such a subset exists. Inflating L by including all inverses and finite products changes neither the property that it is countable nor the property that it is dense, so we assume that L is a subgroup of G .

For a given $i \in \mathbb{N}$ we construct a sequence of elements $g_n \in L$ so that $g_n g_m^{-1} \in L \cap K_i^C$ for all n, m , $n \neq m$. This can be done inductively by taking g_1 any element of L and then g_n any element of $L \cap \bigcap_{m < n} K_i^C g_m$ (this intersection is not empty, since $K_i^C g_m$ is the complement of a compactum in a non-compact space and thus intersects every other complement of a compactum, further the finite intersection of all $K_i^C g_m$ is open and thus intersects L), the condition $g_n g_m^{-1} \in L \cap K^C$ then follows from $g_n g_m^{-1} \in L \cap (K^C g_m g_m^{-1}) = L \cap K^C$.

Now we define $B_i := \{x \in A \mid \text{for all } g \in L \cap K_i^C : xg \notin A\}$, which is equal to $\bigcap_{g \in L \cap K^C} (A \cap \Phi_g^{-1}(A^C))$ and thus measurable. For our sequence g_n and $n \neq m$ we note that for $x \in B_i$ we have $\Phi_{g_n g_m^{-1}}(x) \notin A$. Since $A \supseteq B_i$ this implies $\Phi_{g_n g_m^{-1}}(B_i) \cap B_i = \emptyset$, implying $\Phi_{g_n}(B_i) \cap \Phi_{g_m}(B_i) = \emptyset$. This means

$$\mu \left(\bigcup_{n=1}^{\infty} \Phi_{g_n}(B_i) \right) = \sum_{n=1}^{\infty} \mu(\Phi_{g_n}(B_i)) = \sum_{i=1}^{\infty} \mu(B_i),$$

implying $\mu(B_i) = 0$ by finiteness of μ .

Now

$$\bigcup_{i \in \mathbb{N}} B_i = \{x \in A \mid \exists i \in \mathbb{N} \text{ so that for all } g \in L \cap K_i^C : xg \notin A\}$$

necessarily has measure zero. But this set clearly contains B , since the condition $g \in L \cap K_i^C$ is a weaker condition than $g \in K_i^C$. This completes the proof up to one additional technicality,

namely the question why the return set itself is measurable. Here we need to make use of our family K_i once more:

The return set $\{x \in A \mid x \text{ returns to } A\}$ is equal to $\bigcap_{n=1}^{\infty} (A \cdot K_n^C) \cap A$: Considering “ \subseteq ”: if $x \in A$ so that $R_G(x, A)$ doesn't have compact closure, then $R_G(x, A) \cap K_i^C \neq \emptyset$ for every i (as otherwise $R_G(x, A)$ would be contained in some K_i and admit compact closure), whence $x \cdot K_i^C \in A$, implying points in the return set are in this intersection. For “ \supseteq ”: if $R_G(x, A)$ does have compact closure, then it must be contained in some K_i and $x \cdot K_i^C \notin A$, implying points not in the return set are not in this intersection. With this characterisation of the return set we remark that $A \cdot K_i^C$ is measurable, being equal to $\pi_X[\mu^{-1}(A) \cap (X \times (K_i^C)^{-1})]$, implying that the return set, which is a countable intersection of these sets, is measurable. \square

We specify this result to one that is more useful for our setting:

Proposition 5.2.9. *Let G be a non-compact Lie group acting on a metrisable space X with a finite G -invariant Borel measure. Then for almost every $x \in X$ there is a sequence $g_n \in G$, $g_n \rightarrow \infty$ with $xg_n \rightarrow x$.*

Proof. As a remark, $g_n \rightarrow \infty$ means that for every compact $K \subset G$ only finitely many terms of the sequence g_n lie in K . In other words $g_n \rightarrow \infty$ in the one-point compactification of G . It is worth noting that the one-point compactification of a finite dimensional Lie group is first countable, in particular the neighbourhood filter of ∞ admits countable basis (the K_i^C constructed in the previous proof could serve as such a basis). In this proof we will use the terminology “ x returns to A ” to denote that $R_G(x, A)$ has non-compact closure, here A is a measurable subset of X and $x \in A$.

Note that by Poincaré recurrence the action of G on X is recurrent, so if we chose a metric on X then for any $x \in X$ and $\epsilon > 0$ the points in the ball $B_\epsilon(x)$ returning to $B_\epsilon(x)$ have full measure in $B_\epsilon(x)$. If $\tilde{x} \in B_\epsilon(x)$ is such a point then $B_\epsilon(x) \subseteq B_{2\epsilon}(\tilde{x})$ from which $R_G(\tilde{x}, B_\epsilon(x)) \subseteq R_G(\tilde{x}, B_{2\epsilon}(\tilde{x}))$ follows. In particular \tilde{x} must return to $B_{2\epsilon}(\tilde{x})$, for if it didn't that would mean that $R_G(\tilde{x}, B_{2\epsilon}(\tilde{x}))$ would have compact closure, however if $A \subset B$ in a Hausdorff space and \bar{A} is not compact then \bar{B} cannot be compact (for it contains a closed non-compact subspace). Thus the points \tilde{x} in $B_\epsilon(x)$ that return to $B_{2\epsilon}(\tilde{x})$ have full measure in $B_\epsilon(x)$. So $A_{2\epsilon} = \{\tilde{x} \in X \mid \tilde{x} \text{ returns to } B_{2\epsilon}(\tilde{x})\}$ has full measure in $\bigcup_{x \in X} B_\epsilon(x) = X$.

The intersection $A = \bigcap_{n=1}^{\infty} A_{1/n}$ then also must have full measure in X . This intersection consists of points $x \in X$ that return to $B_\epsilon(x)$ for any $\epsilon > 0$, meaning for any $x \in A$ the sets $G_n(x) = \{g \in G \mid x \cdot g \in B_{1/n}(x)\}$ have non-compact closure. Since they have non-compact closure they must intersect every neighbourhood of ∞ in G . In particular if $\{N_n\}_{n \in \mathbb{N}}$ is a countable

neighbourhood basis of ∞ there exists a sequence g_n consisting of elements $g_n \in G_n(x) \cap N_n$. But $g_n \in N_n$ means $g_n \rightarrow \infty$ and $g_n \in G_n(x)$ means $d(x, xg_n) < \frac{1}{n}$, so $g_n x \rightarrow x$. \square

Recurrence provides Killing fields

We now return to the setting of Cartan geometries. First remark on a technical detail:

Proposition 5.2.10. *Let M be a connected manifold and $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ a Cartan connection on M modelled on G/P . Then $\text{Aut}(M)$ acts properly on \widehat{M} .*

This proposition will be proven in the next sub-section (Sub-section 5.2.3).

Lemma 5.2.11. *Let M be a connected compact manifold equipped with a Cartan connection $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ modelled on G/P so that $\text{Aut}(M)$ is not compact. Suppose that P is a linear algebraic group and M has a finite $\text{Aut}(M)$ invariant Borel measure so that every open set has non-zero volume.*

Then there is a dense subset M' of M so that for every $x \in M'$ and $\widehat{x} \in \pi^{-1}(x)$ one has that $\phi_{\widehat{x}}(\mathfrak{Is}(x)) \subset \mathfrak{p}$ is the Lie algebra of a non-compact subgroup of P .

Proof. As by Proposition 5.2.9 the set of points x for which there exists a sequence $f_n \in \text{Aut}(M)$, $f_n \rightarrow \infty$, $f_n(x) \rightarrow x$ has full measure. Since no open set has zero measure, these points are dense in M , in particular they are also dense in the integrability locus (which is open). Now choose a such a recurrent point x in the integrability locus M^{int} and $f_n \in \text{Aut}(M)$ such a sequence, fix an element \widehat{x} in the fibre $\pi^{-1}(x)$ as well. Denote with \widehat{f}_n the lift of f_n to \widehat{M} .

Since $f_n(x) \rightarrow x$ there exists a sequence $p_n \in P$ with $\widehat{f}_n(\widehat{x} \cdot p_n) = \widehat{f}_n(\widehat{x}) \cdot p_n \rightarrow \widehat{x}$. Now $\text{Aut}(M)$ acts properly on \widehat{M} , so p_n cannot admit a converging subsequence (as then $(\widehat{x} \cdot p_n, \widehat{f}_n(\widehat{x} \cdot p_n))$ would admit a converging subsequence, contradicting properness).

Further since $\mathcal{D}\kappa$ is invariant under automorphisms and P -equivariant one has $\rho(p_n^{-1}) \cdot \mathcal{D}\kappa(\widehat{x}) = \mathcal{D}\kappa(\widehat{f}_n(\widehat{x}) \cdot p_n) \rightarrow \mathcal{D}\kappa(\widehat{x})$ and p_n is ‘‘asymptotically in the stabiliser’’. Making this more rigorous, let \mathcal{O} denote the P -orbit of $\mathcal{D}\kappa(\widehat{x})$ and $P_{\mathcal{D}\kappa(\widehat{x})}$ the stabiliser of $\mathcal{D}\kappa(\widehat{x})$ in P , then \mathcal{O} is diffeomorphic to $P/P_{\mathcal{D}\kappa(\widehat{x})}$. In particular since $\mathcal{D}\kappa(\widehat{x}p_n) \rightarrow \mathcal{D}\kappa(\widehat{x})$ we find $[p_n] \rightarrow 1$ in $P/P_{\mathcal{D}\kappa(\widehat{x})}$, whence there exists a sequence $\delta_n \in P_{\mathcal{D}\kappa(\widehat{x})}$ with $p_n\delta_n \rightarrow 1$. Since $p_n \rightarrow \infty$ so too $\delta_n \rightarrow \infty$. It follows that the stabiliser $P_{\mathcal{D}\kappa(\widehat{x})}$ is not compact.

$P_{\mathcal{D}\kappa(\widehat{x})}$ is the stabiliser of an algebraic group action of an algebraic group, and thus also algebraic. As such it can only have finitely many connected components, and if $P_{\mathcal{D}\kappa(\widehat{x})}$ is not compact so too must be the identity component $P_{\mathcal{D}\kappa(\widehat{x})}^0$. But this identity component is the subgroup generated

by $\phi_{\bar{x}}(\mathfrak{I}\mathfrak{s}(x))$ as per Proposition 5.2.4. This then holds for every recurrent point in M^{int} , but these are dense since M^{int} is open and dense. \square

As a consequence of the previous lemma, there is a dense set of points $x \in M$ for which $\mathfrak{I}\mathfrak{s}(x) \neq 0$. This set can be upgraded to be at least all of M^{int} :

Proposition 5.2.12. *Let M be as in Lemma 5.2.11, and $x \in M^{\text{int}}$, then $\mathfrak{I}\mathfrak{s}(x) \neq 0$.*

Proof. Let $x \in M^{\text{int}}$, then every point y in the same component of M^{int} as x has an isomorphic $\mathfrak{k}\mathfrak{ill}_y$ by Proposition 4.4.2. Let U_x be a neighbourhood of x as in Proposition 2.5.9, that is $\mathfrak{k}\mathfrak{ill}(U_x) = \mathfrak{k}\mathfrak{ill}_x$, then the restriction of $\mathfrak{k}\mathfrak{ill}(U_x)$ to $\mathfrak{k}\mathfrak{ill}_y$ must be an isomorphism, since it is injective and $\dim(\mathfrak{k}\mathfrak{ill}(U_x)) = \dim(\mathfrak{k}\mathfrak{ill}_x) = \dim(\mathfrak{k}\mathfrak{ill}_y)$.

The map $ev : U_x \times \mathfrak{k}\mathfrak{ill}(U_x) \rightarrow T(U_x)$, $(y, X) \mapsto X_y$ is a smooth map, indeed it is a vector bundle morphism. For any $y \in U_x$ we have $\ker(ev_y) = \mathfrak{I}\mathfrak{s}(y)$, which is non-zero on a dense subset of U_x . By smoothness of the map ev this kernel must then be non-zero on all of U_x , in particular $\mathfrak{I}\mathfrak{s}(x) \neq 0$. \square

The conclusion can actually be extended to all of M , as one can show that the set of points for which $\mathfrak{I}\mathfrak{s}(x) \neq 0$ is closed, but we will not need this here. We now use $\mathfrak{I}\mathfrak{s}(x) \neq 0$ to find another Killing field:

Proposition 5.2.13. *Let M be as in Lemma 5.2.11 and $x \in M^{\text{int}}$. Then $\dim(\mathfrak{k}\mathfrak{ill}_x) \geq 2$.*

Proof. Let $x \in M^{\text{int}}$ and $X \in \mathfrak{k}\mathfrak{ill}_x$ a non-zero Killing field (though it may vanish at x). Let X be defined on some connected open neighbourhood U_x of x . Since X is a non-zero Killing field, its zero locus must be closed and nowhere dense. In particular there exists a point y in the same component of M^{int} with $X_y \neq 0$. It follows that $\mathfrak{k}\mathfrak{ill}_y$ is at least two dimensional, containing both X and an element of $\mathfrak{I}\mathfrak{s}(y)$, which are necessarily linearly independent in $\mathfrak{k}\mathfrak{ill}_y$.

As before the sheaf $\mathfrak{k}\mathfrak{ill}$ is locally constant on the components of M^{int} , whence $\dim(\mathfrak{k}\mathfrak{ill}_x) \geq 2$ also follows. \square

Specialising to pseudo-Riemannian geometry

As noted before, pseudo-Riemannian manifolds admit isometry-invariant measures. These are necessarily of finite volume for compact manifolds and no open set has vanishing volume. Further $P = O(p, q)$ in this setting, which is an algebraic group. Thus Lemma 5.2.11 and Propositions 5.2.12 and 5.2.13 hold here. Summarising: For every point $x \in M^{\text{int}}$ we have $\dim(\mathfrak{I}\mathfrak{s}(x)) \geq 1$, $\dim(\mathfrak{k}\mathfrak{ill}_x) \geq 2$ and for a dense subset of M^{int} $\mathfrak{I}\mathfrak{s}(x)$ generates a non-compact subgroup of $O(p, q)$.

Proposition 5.2.14. *Let M be as in Lemma 5.2.11 and specify to a torsion free connection modelled on $G/P = O(p, q) \ltimes \mathbb{R}^n / O(p, q)$, then $\dim(\mathfrak{k}\mathfrak{il}_x) \geq 3$ for every point $x \in M$.*

Proof. Denote the pseudo-metric on M with g . Note that since ω is torsion free the associated connection is the Levi-Civita connection. If X is a Killing field the equation $L_X g = 0$ implies, together with the Koszul formula, that $g(\nabla_V X, W) + g(\nabla_W X, V) = 0$ for any locally defined vector fields V, W . In [Fr 18] this equation is referred to as Clairault's equation.

Now consider a point $x \in \widehat{M}^{\text{int}}$ admitting a non-vanishing local Killing field X , let Y be a local Killing field so that $Y_x = 0$, that is $Y \in \mathfrak{I}\mathfrak{s}(x)$ (and $Y \neq 0$). Since the zero locus of Y is closed and nowhere dense, there will be a vector $v \in T_x M$ that is transverse to this zero locus so that $g(X_x, v) \neq 0$. Let $\gamma_v(t) = \exp_x(t, v)$ be a parametrised geodesic with derivative v . Clairault's equation (with $V = W = \dot{\gamma}(t)$) then implies

$$\frac{d}{dt}g(X_{\gamma(t)}, \dot{\gamma}(t)) = g(\nabla_{\dot{\gamma}(t)} X, \dot{\gamma}(t)) = 0$$

and the same with X replaced by Y . Thus $g(X_{\gamma(t)}, \dot{\gamma}(t))$ is constant $\neq 0$ and $g(Y_{\gamma(t)}, \dot{\gamma}(t))$ is also constant $= 0$. But since u is transverse to the zero locus of Y , $Y_{\gamma(t)}$ will be non-zero for small enough $t > 0$, we know fix such a t small enough that $\gamma(t)$ is still in M^{int} . These two equations imply that $X_{\gamma(t)}$ and $Y_{\gamma(t)}$ are not zero and linearly independent. But since every point in admits a local Killing field vanishing at that point, there must be a third local Killing field Z at $\gamma(t)$ with $Z_{\gamma(t)} = 0$. By construction Z is not a linear combination of X and Y , and then $\dim(\mathfrak{k}\mathfrak{il}_{\gamma(t)}) \geq 3$. Since $\gamma(t) \in M^{\text{int}}$, the result holds for all points in the same component of M^{int} as x . But x was arbitrary, thus it holds for all points in M^{int} . \square

5.2.3 $\text{Aut}(M)$ acts properly on \widehat{M}

In this sub-section we show that $\text{Aut}(M)$ acts properly on \widehat{M} . We briefly remind ourselves what it means for a group action to be proper.

Definition 5.2.15. Let G be a topological group and X a topological space with a (right) G -action. Then the G -action is proper iff the map $M \times G \rightarrow M \times M$, $(x, g) \mapsto (x, xg)$ is proper.

It is clear that a group action G on X is proper if and only if for any net (x_α, g_α) so that x_α and $x_\alpha g_\alpha$ both converge that then g_α must admit a convergent sub-net. In the event that G and X are second countable it is enough to only consider sequences in this formulation³.

³In second countable spaces the notions of compact and sequentially compact are equivalent.

Proposition 5.2.16. *Let M be a connected manifold and $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ a Cartan connection on M modelled on G/P . Then $\text{Aut}(M)$ acts properly on \widehat{M} .*

Proof. First we show the following related result: If $F_n \in \text{Aut}(M)$ and $\widehat{x} \in \widehat{M}$ so that $F_n(\widehat{x})$ converges, then there is an $F \in \text{Aut}(M)$ with $F_n \rightarrow F$ uniformly on compacta.

The main point is that the flows of the constant fields $\omega^{-1}(v)$, $v \in \mathfrak{g}$ explore a full neighbourhood of any point. By denoting the semi-group of local diffeomorphism given by composing such local flows with Γ , we note that Γ must then act transitively on M since M is connected. Since the automorphisms F_n preserve the constant fields (i.e. $DF_n(\omega^{-1}(v)) = \omega^{-1}(v)$) they commute with their flows, and hence with the elements of Γ .

Let $\widehat{y} \in \widehat{M}$ and $\varphi = \varphi_{z_1}^t \circ \dots \circ \varphi_{z_n}^t$ a local diffeomorphism from Γ with $\varphi(\widehat{x}) = \widehat{y}$. Then $F_n(\widehat{y}) = F_n(\varphi(\widehat{x})) = \varphi(F_n(\widehat{x}))$, in particular $F_n(\widehat{y})$ must converge for every $\widehat{y} \in \widehat{M}$. We denote with $F : \widehat{M} \rightarrow \widehat{M}$ the pointwise limit of F_n .

We can see that F is smooth by choosing an arbitrary point $\widehat{y} \in \widehat{M}$ and a neighbourhood of 0 in \mathfrak{g} for which the flow starting at y along constant fields is a diffeomorphism to a neighbourhood U of \widehat{y} . This gives a coordinate chart U . We may do the same for $F(\widehat{y})$ getting a neighbourhood of V , which we may assume (by making both U, V smaller) to be parametrised by the same open subset of \mathfrak{g} as U . Then since F commutes with the flow of constant fields the map F is the identity map in these coordinate charts, thus smooth on a neighbourhood of \widehat{y} (and hence everywhere since \widehat{y} was arbitrary).

So we have a smooth map $F : \widehat{M} \rightarrow \widehat{M}$ that commutes with the flows of the constant fields (hence preserve these and $F^*(\omega) = \omega$) and is the pointwise limit of bundle-automorphisms, hence is itself a bundle-automorphism. As such $F \in \text{Aut}(M)$.

Now we check that $F_n \rightarrow F$ uniformly on compacta. To do this step we will introduce a Riemannian metric on \widehat{M} for which $\text{Aut}(M)$ acts by isometries. Then we will see that any sequence of isometries converging pointwise to an isometry must converge uniformly on compacta to that isometry.

Let $\mathfrak{a}_1, \dots, \mathfrak{a}_n$ be a basis of \mathfrak{g} . Denote with g' the scalar product on \mathfrak{g} defined by $g'(\mathfrak{a}_i, \mathfrak{a}_j) = \delta_{ij}$ and denote with g the Riemannian metric on \widehat{M} defined by:

$$g_{\widehat{x}}(v, w) = g'(\omega_{\widehat{x}}^{-1}v, \omega_{\widehat{x}}^{-1}w) \quad \text{for } \widehat{x} \in \widehat{M} \text{ and } v, w \in T_{\widehat{x}}\widehat{M}.$$

Then any automorphism sends the ONB $\omega^{-1}(\mathfrak{a}_1), \dots, \omega^{-1}(\mathfrak{a}_n)$ to itself and as such is an isometry.

Now if f_n is a sequence of isometries of a metric space converging pointwise to an isometry f but not uniformly on compacta, there must be some compactum K with $\sup_{x \in K} d(f_n(x), f(x)) > \epsilon$. We choose a sequence $x_n \in K$ so that $d(f_n(x_n), f(x_n)) > \epsilon$ and by compactness of K we assume x_n converges to some element x . Then

$$d(f_n(x_n), f(x_n)) \leq \underbrace{d(f_n(x_n), f_n(x))}_{=d(x_n, x)} + d(f_n(x), f(x)) + \underbrace{d(f(x), f(x_n))}_{=d(x, x_n)},$$

by assumption, however every term on the right-hand side converges to 0, contradicting that the left-hand side is $> \epsilon$.

This proves the first step we wished to perform, namely that if $F_n \in \text{Aut}(M)$ and $\hat{x} \in \widehat{M}$ so that $F_n(\hat{x})$ converges, then there is an $F \in \text{Aut}(M)$ with $F_n \rightarrow F$ uniformly on compact. This is close to, but not directly the same as, the statement that $\text{Aut}(M)$ acts properly on \widehat{M} . In order to derive properness of the action from this point we continue viewing the automorphisms as isometries of a metric space and make the following remark:

Let X be a metric space and $(x_n, f_n) \in X \times \text{Isom}(X)$ so that $f_n(x_n)$ converges to a $y \in X$ and $x_n \rightarrow x \in X$. Then note that

$$d(f_n^{-1}(y), x) \leq d(f_n^{-1}(y), x_n) + d(x_n, x) = d(y, f_n(x_n)) + d(x_n, x),$$

where the right-hand side converges to zero.

Thus if $(\hat{x}_n, F_n) \in \widehat{M} \times \text{Aut}(M)$ so that $F_n(\hat{x}_n)$ and \hat{x}_n both converge, then there must be a point $\hat{y} \in \widehat{M}$ so that $F_n^{-1}(\hat{y})$ converges. Hence F_n^{-1} (and thus F_n) must converge in the topology of $\text{Aut}(M)$ and the action of $\text{Aut}(M)$ is proper on \widehat{M} . \square

This proposition can then be used to gain the following result:

Corollary 5.2.17. *If M is compact and $(\widehat{M} \xrightarrow{\pi} M, \omega : T\widehat{M} \rightarrow \mathfrak{g})$ is a Cartan connection on M modelled on G/P with P compact, then $\text{Aut}(M)$ is compact.*

Proof. If M and P are compact then \widehat{M} is a fibre bundle over a compact space with a compact fibre and hence itself compact. Then $\text{Aut}(M) \times \widehat{M}$ is the pre-image of the compact $\widehat{M} \times \widehat{M}$ under a proper mapping and must also be compact. \square

5.3 Classification of components of \widehat{M}^{int} by their local Killing algebras

The previous section has shown that if M is a compact pseudo-Riemannian manifold with non-compact isometry group then every point of the integrability locus admits an at least 3-dimensional algebra of local Killing fields and the isotropy algebra $\mathfrak{Is}(x)$ of every point is at least 1 dimensional. The construction of “the two additional Killing fields” (Propositions 5.2.13 and 5.2.14) beyond the one given by $\mathfrak{Is}(x)$ make clear that on a dense set in M we have at least “two horizontal local Killing fields”, meaning $\dim(\mathfrak{kill}_x) - \dim(\mathfrak{Is}(x)) \geq 2$, and at least “one vertical local Killing field”, meaning $\dim(\mathfrak{Is}(x)) \geq 1$.

Definition 5.3.1. Let $x \in M$. We denote with $h_x = \dim(\mathfrak{kill}_x) - \dim(\mathfrak{Is}(x))$ and call this the number of horizontal local Killing fields at x . We denote with $v_x = \dim(\mathfrak{Is}(x))$ and call this the number of vertical local Killing fields at x . In this section we will refer to the tuple (v_x, h_x) as the type of the point x .

Remark. For any point x the number h_x corresponds to the dimension of the \mathfrak{kill} orbit of x , that is the manifold explored by starting at x and then successively flowing along locally defined Killing fields (since \mathfrak{kill} is closed under action of the vector field commutator, this is indeed a manifold by the Frobenius integrability condition). As v_x corresponds to the dimension of $\mathfrak{Is}(x)$, which corresponds to those Killing fields admitting vertical lifts, one notes $h_x \leq \dim(M)$ and $v_x \leq \dim(P)$ for all x .

If one considers a connected component of M^{int} , denoted here with \mathcal{M} , then $\dim(\mathfrak{kill}_x)$ is constant as x varies over \mathcal{M} . The numbers (v_x, h_x) are however not necessarily constant and thus do not a priori act as invariants of the component \mathcal{M} or help distinguish it from other components. This remark however loses validity in the case of a 3-dimensional Lorentz manifold, where they are invariants of the components (although there is some wiggle room). We will see this now.

For 3-dimensional Lorentz manifolds

Now M is a compact 3-dimensional Lorentz manifold with a non-compact isometry group. The following fact is helpful in the classification:

Fact 5.3.2. *For any finite dimensional representation of $O(2, 1) \rightarrow GL(V)$ no vector can have a stabiliser of dimension 2.*

Reminder. By Corollary 5.2.5 for a point $x \in M^{\text{int}}$ the isotropy algebra $\mathfrak{Is}(x)$ generates the stabiliser of a point $\mathcal{D}\kappa(\widehat{x})$ under the representation of $O(2, 1)$ in \mathcal{W} . Thus $v_x = \dim(\mathfrak{Is}(x))$ is

not allowed to be 2.

Then $1 \leq v_x \leq 3 = \dim(O(2,1))$ and $v_x \neq 2$, so $v_x \in \{1, 3\}$. For h_x we have the restrictions $h_x \leq 3$ and $v_x + h_x \geq 3$. If $v_x = 3$ and $h_x \neq 0$, then h_x must be 3. This follows from the flows of the vertical fields acting on $T_x M$ as $O(2,1)^0$ does on \mathbb{R}^3 . If there is a single non-vertical Killing field here then by irreducibility of that representation then there must be a Killing field pointing in any direction of $T_x M$.

We make the arguments in the case $v_x = 3, h_x \neq 0$ more specific. Proposition 3.2.19 established that $TM \cong \widehat{M} \times_{O(2,1)} \mathfrak{so}(2,1) \times \mathbb{R}^3 / \mathfrak{so}(2,1) = \widehat{M} \times_{O(2,1)} \mathbb{R}^3$, it follows $T_x M \cong (\widehat{x} \cdot O(2,1)) \times_{O(2,1)} \mathbb{R}^3$ for a point $\widehat{x} \in \pi^{-1}(x)$, we choose a point \widehat{x} and use it to identify $T_x M$ with \mathbb{R}^3 . If X is a Killing field that is vertical at x , then as remarked in the proof of Proposition 5.2.4 the flow of \widehat{X} on $\widehat{x} \cdot O(2,1)$ is by multiplication with $\exp(t\phi_{\widehat{x}}(X))$. On $T_x M \cong \mathbb{R}^3$ this flow then must become $[\widehat{x}, v] \cdot \exp(t\phi_{\widehat{x}}(X)) = [\widehat{x} \cdot \exp(t\phi_{\widehat{x}}(X)), v] = [\widehat{x}, \exp(-t\phi_{\widehat{x}}(X))v]$. Since $\mathfrak{J}\mathfrak{s}(x)$ is maximal dimensional one finds that under this identification the flow of the vertical fields acts as $O(2,1)^0$ does on \mathbb{R}^3 .

One last remark: points of type (3,0) are nowhere dense in M^{int} . This follows from $v_x = \dim(\ker(ev_x))$ for a certain function ev_x (if $x \in M^{\text{int}}$), whence v_x cannot increase locally and the points y with $v_y \leq 2$ are open. But every neighbourhood of a (3,0) point must admit points y with $h_y \geq 1$, as otherwise the local Killing fields would be zero everywhere in this neighbourhood. Since inside M^{int} the number $h_y + v_y$ is locally constant and $v_y \in \{1, 3\}$ everywhere, this leaves only (1,2) type points if the neighbourhood is small enough. Thus: inside a component of M^{int} containing a (3,0) point the points of type (1,2) are open and dense.

This means that only points of type (3,3), (1,3), (1,2) or (3,0) can exist, while points of type (3,0) are nowhere dense.

Lemma 5.3.3 (Initial classification of components of M^{int} by their Killing algebras). *Let \mathcal{M} be a connected component of M^{int} . Then either:*

1. *For every point $x \in \mathcal{M}$ x is of type (3,3) and the generalised curvature $\mathcal{D}\kappa$ is constant on \mathcal{M} .*
2. *For every point $x \in \mathcal{M}$ x is of type (1,3) and \mathcal{M} is locally homogenous.*
3. *For a dense set of points $x \in \mathcal{M}$ x is of type (1,2), with there possibly existing a nowhere dense set of points of type (3,0).*

Proof. The proof follows from remembering that $\dim(\mathfrak{kill}_x)$ is locally constant on M^{int} and thus constant on \mathcal{M} . We now consider the different possible values of $\dim(\mathfrak{kill}_x)$.

In case $\dim(\mathfrak{kil}_x) = 6$ it is necessary that $(v_x, h_x) = (3, 3)$. Further we know $\text{Rank}(D_{\hat{x}}\mathcal{D}\kappa) = \dim(\mathfrak{g}) - \dim(\mathfrak{kil}_x) = 6 - 6$ inside such a component and then $\mathcal{D}\kappa$ must be constant on \mathcal{M} .

The case $\dim(\mathfrak{kil}_x) = 5$ cannot be realised, as no type (v_x, h_x) in $(3,3)$, $(1,3)$, $(1,2)$ or $(3,0)$ has $v_x + h_x = 5$.

The case $\dim(\mathfrak{kil}_x) = 4$ is only described by $(1,3)$ points. Since h_x corresponds to the dimension of points explored by starting at x and successively flowing along local Killing fields, we find that an open neighbourhood of any point is explored by flowing along Killing fields. From the connectedness of \mathcal{M} local homogeneity follows.

In the case of $\dim(\mathfrak{kil}_x) = 3$ we have already seen that the $(1,2)$ points form a dense open subset of \mathcal{M} , the complement consists of $(3,0)$ points but is nowhere dense.

No further cases are possible, since $\dim(\mathfrak{kil}_x) \geq 3$ must hold for every point. □

This classification is essentially a classification by $\dim(\mathfrak{kil}_x)$. We will now further this by splitting the locally homogenous components and the $(1,2)$ (on dense open set) components into two sub-classes:

Definition 5.3.4. A component of type $(1,3)$ or of type $(1,2)$ is called hyperbolic if there is a point $x \in M$ with $v_x = 1$ so that $\phi_{\hat{x}}(\mathfrak{J}\mathfrak{s}(x))$ generates a hyperbolic subgroup of $O(2,1)$. Else a component of this type is called parabolic.

This gives the final classification of a component \mathcal{M} to be used in the proof of the main theorem. It is convenient to define abbreviations for these components⁴:

A component \mathcal{M} is either

1. Of constant curvature **(cc)**.
2. Locally homogenous and hyperbolic **(hh)**.
3. Locally homogenous and parabolic **(hp)**.
4. Not locally homogenous and hyperbolic **(nh)**.
5. Not locally homogenous and parabolic **(np)**.

⁴As an exercise the reader may try to estimate how many times they can read variations of “ M has a non locally homogenous component and a hyperbolic component, but the non locally homogenous component need not be hyperbolic” without going mad.

5.4 Overview of the proof of the main theorem

The proof of the main theorem, described in section 5.1, now proceeds by separately considering 3 cases. What case we are in is determined by what kind of components M^{int} admits. Specifically the cases are:

1. The homogenous case: All components are either locally homogenous or of constant curvature, that is M has only **(cc)**, **(hh)** or **(hp)** components.
2. The hyperbolic case: Not all components are locally homogenous or of constant curvature and at least one component is hyperbolic, that is M has at least one **(nh)** or **(np)** component and at least one **(hh)** or **(nh)** component.
3. The parabolic case: Not all components are locally homogenous or of constant curvature and there are no hyperbolic components, that is M has at least one **(np)** component and no **(nh)** or **(hh)** components.

The following sections briefly sketch the steps of the proof of these cases (as developed in [Fr 18]), excluding the parabolic case (which is by far the longest). Due to the already excessive length of this thesis and time constraints proofs will not be carried out.

5.4.1 The locally homogenous case

Here we roughly describe the trajectory of the proof in the case that all components are locally homogenous or constant curvature, meaning of types **(cc)**, **(hh)** or **(hp)**. The cases covered by the theorem are:

Theorem 5.4.1. *If (M, g) is a connected, closed, oriented and time-orientable 3-dimensional Lorentz manifold, so that $\text{Isom}(M, g)$ is non-compact and the local Killing algebra of M is at least 4 dimensional on a dense open set, then either:*

1. M is diffeomorphic to \mathbb{T}^3 or \mathbb{T}_A^3 for a hyperbolic $A \in SL_2(\mathbb{Z})$ and M is flat.
2. M is diffeomorphic to \mathbb{T}_A^3 for a parabolic $A \in SL_2(\mathbb{Z})$ and M is either flat or locally isometric to Lorentz-Heisenberg geometry.
3. M is diffeomorphic to a quotient $\Gamma \backslash \widetilde{SL}_2(\mathbb{R})$, the metric is then induced by a left-invariant non-Riemannian metric on $\widetilde{SL}_2(\mathbb{R})$.

The proof of this theorem relies on a fine understanding of locally homogenous Lorentz 3-manifolds. The first step is to understand that restricting the manifold to only admit **(cc)**, **(hh)** or **(hp)**

components means that it is locally homogenous. One can then make use of Bieberbach theorems describing what the discrete subgroups $\Gamma \subset \text{Isom}(X)$ of the simply connected homogenous model spaces X admitting compact quotients $\Gamma \backslash X$ look like, thus describing the original manifold.

Theorem 5.4.2 ([Fr 16]). *Let (M, g) be a connected 3-dimensional Lorentz manifold admitting an open set on which the germs of Killing fields is at least 4 dimensional, then M is locally homogenous.*

Note that **(cc)** components have $\dim(\mathfrak{kil}_x) = 6$ for all x in the component and **(hh)** and **(hp)** components have $\dim(\mathfrak{kil}_x) = 4$ for all x in the component. Thus if every component of M^{int} is of one of these types then M is locally homogenous. In particular $M^{\text{int}} = M$ and there is only one component. Additionally since we have that $\mathfrak{Is}(x)$ generates a non-compact subgroup of $O(2, 1)$ on a dense subset of M by Lemma 5.2.11 it does so for every point (as M is locally homogenous, and the property that $\mathfrak{Is}(x)$ generates a non-compact subgroup is preserved by local isometries).

A complete classification of locally homogenous Lorentz geometries on compact manifolds is available in the literature:

Theorem 5.4.3 ([DZ 10]). *Suppose (M, g) is a closed 3-dimensional Lorentz manifold so that g is locally homogenous and non-Riemannian. Then g is locally isometric to one of the following:*

1. *A flat metric on \mathbb{R}^3 .*
2. *A Lorentzian, non-Riemannian, left-invariant metric on $\widetilde{SL}_2(\mathbb{R})$.*
3. *A Lorentz-Heisenberg metric on the group Heis.*
4. *The Lorentz-Sol metric on the group SOL.*

[Fr 18] remarks that in the flat case a Bieberbach rigidity theorem on \mathbb{R}^3 with Minkowski space is well known due to the works [Ca 89], [FG 83] and [GK 84]. In the case of Lorentz-Heisenberg and Lorentz-Sol geometries Bieberbach type theorems were developed in [DZ 10]. We summarise them here.

Theorem 5.4.4 (Bieberbach rigidity theorem for flat, Lorentz-Heisenberg and Lorentz-Sol manifolds). *Let (M, g) be a closed, 3-dimensional Lorentz manifold.*

1. *If (M, g) is flat there exists a discrete subgroup $\Gamma \subset \text{Isom}(\mathbb{R}^3)$ (here \mathbb{R}^3 has the Minkowski metric) such that (M, g) is isometric to the quotient $\Gamma \backslash \mathbb{R}^3$. Further there exists a connected 3 dimensional Lie group $G \subset \text{Isom}(\mathbb{R}^3)$ which is isometric to \mathbb{R}^3 , Heis or SOL and which acts simply transitively on \mathbb{R}^3 satisfying that $\Gamma_0 = G \cap \Gamma$ has finite index in Γ and is a uniform lattice in G .*
2. *If (M, g) is locally modelled on the Lorentz-Heisenberg geometry then M is isometric to*

a quotient $\Gamma \backslash \text{Heis}$ for a discrete subgroup $\Gamma \subset \text{Isom}(\text{Heis})$ and there exists a finite order subgroup $\Gamma_0 \subset \Gamma$ that is a lattice $\Gamma_0 \subset \text{Heis}$ acting by left translations.

3. If (M, g) is a locally modelled on the Lorentz-Sol geometry then M is isometric to a quotient $\Gamma \backslash \text{SOL}$ and the intersection $\Gamma_0 = \Gamma \cap \text{Isom}^0(\text{SOL})$ is a lattice $\Gamma_0 \subset \text{SOL}$ acting by left-translations.

From this one is able to prove:

Proposition 5.4.5 ([Fr 18], Propositions 4.3 and 4.6). *Let (M, g) be a closed, orientable and time-orientable 3-dimensional Lorentz manifold.*

1. *If (M, g) is flat and $\text{Isom}(M, g)$ is non-compact, then (M, g) is diffeomorphic to a torus or a torus-bundle \mathbb{T}_A^3 with A hyperbolic or parabolic in $SL_2(\mathbb{Z})$.*
2. *If (M, g) is modelled on Lorentz-Heisenberg metric and $\text{Isom}(M, g)$ is non-compact, then (M, g) is diffeomorphic to a parabolic torus bundle.*
3. *If (M, g) is modelled on Lorentz-Sol metric then $\text{Isom}(M, g)$ is compact.*

Leaving the case that (M, g) is locally isometric to a left-invariant non-Riemannian metric on $\widetilde{SL}_2(\mathbb{R})$. Here one is able to show that there is an anti-de Sitter metric on M preserved by a finite index subgroup of $\text{Isom}(M, g)$. The following Proposition then recovers the topological type of M :

Proposition 5.4.6 ([Fr 18], Proposition 4.7). *If (M, g) is a closed, oriented and time-oriented 3-dimensional anti-de Sitter manifold and $\text{Isom}(M, g)$ is non-compact, then there is a uniform lattice $\Gamma \subset \widetilde{SL}_2(\mathbb{R})$ so that M is diffeomorphic to $\Gamma \backslash \widetilde{SL}_2(\mathbb{R})$.*

5.4.2 The hyperbolic case

The theorem classifying the the hyperbolic case, that is the case when not all components are locally homogenous or constant curvature and there is at least one hyperbolic component, is the following:

Theorem 5.4.7. *If (M, g) is a connected, closed, orientable and time-orientable 3-dimensional Lorentz manifold, so that $\text{Isom}(M, g)$ is non-compact, M is not locally homogenous and M^{int} admits a hyperbolic component, then:*

1. *M is diffeomorphic to \mathbb{T}^3 or a hyperbolic torus bundle \mathbb{T}_A^3 ($A \in SL_2(\mathbb{R})$ hyperbolic).*
2. *The universal cover (M, g) is isometric to \mathbb{R}^3 with metric $dt^2 + a(t)dudv$ for $a(t)$ a periodic positive function.*

3. *There is an isometric action of the group SOL on (M, g) .*

A result of the locally homogenous case was that M is locally homogenous if and only if all of its components are locally homogenous. Thus in the theorem we could replace “is not locally homogenous” with “there exists a non-locally homogenous component” and it is clear why this theorem provides the classification result for the hyperbolic case.

A rough slogan of the proof of this theorem would be “one can fully explore M by pushing a 2-dimensional Anosov torus along its normal flow”. The proof then begins by finding this torus. The first step is to show the existence of a **(nh)** component. For the rest of this section M is now as in Theorem 5.4.7, meaning it has at least one **(nh)** or **(np)** component and at least one **(nh)** or **(hh)** component.

Proposition 5.4.8. *Let \mathcal{M} be a **(hh)** component of M . Then every neighbouring component \mathcal{M}' is also hyperbolic, that is either **(hh)** or **(nh)**. (Neighbouring means here $\partial\mathcal{M}' \cap \partial\mathcal{M} \neq \emptyset$.)*

In particular this implies the existence of a **(nh)** component, for if all hyperbolic components were **(hh)** then every component would be **(hh)** by connectedness of M and the above proposition. **(nh)** components are of type $(1, 2)$ on a dense open set, so for such points the orbit of local Killing fields is a two-dimensional submanifold of the component. Hyperbolicity of the component will ensure that we can choose a point such that this submanifold is a Lorentz-submanifold, in fact it will be a closed Lorentz-submanifold:

Proposition 5.4.9. *Let \mathcal{M} be a **(nh)** component, then there is a point \hat{x} in \mathcal{M} so that the orbit of local Killing fields of $\pi(\hat{x})$ is a closed Lorentz-surface in M .*

This surface will be denoted by Σ_0 from now on. The isometry group $\text{Isom}(M)$ must send Σ_0 to the Isom^{loc} orbit of Σ_0 , which by compactness of M must be a finite number of copies of the orbit of Σ_0 obtained by flowing along local Killing fields. But the orbit of Killing flows of Σ_0 is Σ_0 by definition. Thus the subgroup $H \subset \text{Isom}(M)$ leaving Σ_0 invariant is of finite index, in particular it is non-compact.

By a result in [Ze 96] (to be precise, Proposition 3.6 of that reference) if a sequence of isometries $f_n : N \rightarrow N$ of an n -dimensional Lorentz-manifold acts equicontinuously on a Lorentz-hypersurface of $T_x N$, then the sequence f_n already was equicontinuous. A more careful formulation of this result would imply:

Proposition 5.4.10. *The action of H on Σ_0 is proper. In particular the image of H in $\text{Isom}(\Sigma_0)$ generates a non-compact subgroup.*

This means that Σ_0 must be a closed Lorentz-surface with a non-compact isometry group. The classification of such surfaces is easily achieved.

Fact 5.4.11. *The only closed Lorentz-surface with non-compact isometry group is the flat 2-torus $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$ with metric $du dv$, where u, v are eigenvectors of a hyperbolic element $A \in SL_2(\mathbb{Z})$. Any $h \in \text{Isom}(\mathbb{T}^2)$ generating a non-compact subgroup must be conjugate to a power $A^n \circ \varphi$, where φ generates a compact subgroup (i.e. is a reflection or translation).*

So Σ_0 must be such a torus, and if h generates a non-compact subgroup of $\text{Isom}(M)$ leaving Σ_0 invariant then h must act as a hyperbolic linear transformation on $T_x \Sigma_0$ for some point $x \in \Sigma_0$.

Establishing the existence of this flat Anosov torus is the first part of the proof of the hyperbolic case. In the second part one pushes this torus along its geodesic normal flow, and notices that this flow eventually reconnects to the torus. In order to define this flow one chooses a normal field ν on Σ_0 . By setting $N = \mathbb{R} \times \Sigma_0$ one defines the normal flow f via

$$(t, z) \mapsto \exp_z(t\nu(z)) =: f(t, z).$$

While f is not necessarily well defined on all of N , it is true that f is well defined and smooth in some neighbourhood of $\{0\} \times \Sigma_0$. In fact by the inverse function theorem f will be invertible in a small enough neighbourhood, and one defines

$$T := \sup\{t \in \mathbb{R} \mid f : (0, t) \times \Sigma_0 \rightarrow M \text{ is an injective immersion}\}.$$

For practical reasons one sets the domain of f to be $(0, T) \times \Sigma_0$.

It turns out there exists a function $a : (0, T) \rightarrow \mathbb{R}_{>0}$ with $f^*(g)_{(t,z)} = dt^2 + a(t)g_0$, where g_0 is the metric on Σ_0 and g the metric on M . Further a can be extended smoothly to $[0, T]$ with $a(T) \neq 0 \neq a(0)$, which means the metric $f^*(g)$ can be extended to some neighbourhood $(-\delta, T + \delta) \times \Sigma_0$ in N . The following lemma then ensures that f can be extended to an isometric immersion $[0, T] \times \Sigma_0 \rightarrow M$:

Lemma 5.4.12 ([Fr 18], Proposition 5.8). *Let (L, \tilde{g}) be a Lorentz-manifold, $\Omega \subset L$ open so that $\bar{\Omega}$ is a manifold with boundary and $\partial\Omega$ is a Lorentz-hypersurface in L . If M is a compact Lorentz-manifold of the same dimension as L and $f : (\Omega, \tilde{g}) \rightarrow (M, g)$ is an injective isometric immersion, then f extends to a smooth isometric immersion $\bar{f} : \bar{\Omega} \rightarrow M$.*

So one can extend $\bar{f} : [0, T] \times \Sigma_0 \rightarrow M$. Of particular interest is what $\bar{f}|_{\{T\} \times \Sigma_0}$ does.

Proposition 5.4.13. *\bar{f} maps $\{T\} \times \Sigma_0$ diffeomorphically and isometrically to Σ_0 .*

Elementary considerations also show that then $\dot{\gamma}_z(T) = \nu(\gamma_z(T))$, that is the flow enters Σ_0 orthogonally and from “below”, so the normal flow can then be extended beyond this point, at least until time $T + T$ when it again lands in Σ_0 and one can again extend it. The consequence is then:

Corollary 5.4.14. *f can be extended to all of $\mathbb{R} \times \Sigma_0$ with $f(t+T, z) = f(t, f(t, z))$ and f being a local diffeomorphism. Equipping $\mathbb{R} \times \Sigma_0$ with metric $\tilde{g} = dt^2 + a(t)g_0$ makes f into an isometric immersion.*

The function $a : \mathbb{R} \rightarrow \mathbb{R}_{>0}$ in this corollary is periodic with period T . If we denote with B the map $\Sigma_0 \rightarrow \Sigma_0, z \mapsto f(T, z)$ then the quotient

$$N_B := \mathbb{R} \times \Sigma_0 / (t+T, z) \sim (t, Bz)$$

is a torus if B is conjugate to a translation, else a torus bundle (it being a 3-dimensional Klein bottle, i.e. B being conjugate to $-A^n$, is ruled out by the assumption that M is time-orientable and orientable). Since N is a covering space of N_B with decktransformations given by isometries \tilde{g} induces a metric on N_B . By construction of N_B , f induces a map $[f] : N_B \rightarrow M$ that is an injective isometric immersion. Since both spaces are compact, connected and of the same dimension this must be an isometry.

Once one is at this point verifying whats left Theorem 5.4.7 amounts to checking details.

Chapter 6

Example: Hyperbolic torus bundles

Theorem 5.1.1 provides the classification of oriented and time-oriented compact Lorentz 3-folds with non-compact isometry group. Geometrically these may be divided into 3 groups:

1. Hyperbolic torus bundles and the torus.
2. Parabolic torus bundles and the torus.
3. Locally homogenous $\widetilde{SL}_2(\mathbb{R})$ quotients.

This section investigates the first case. Concretely we find the isometry groups for these spaces (up to finite index), describe the Cartan connection explicitly and find the local Killing fields. We begin with an elaboration of the general computational procedure that finds these quantities and then calculate them for the specific example.

6.1 Calculating the relevant quantities

This section sketches a general procedure used to calculate the quantities of interest for the examples.

In this section we will be looking at pseudo-Riemannian manifold (M, g) of dimension n and g having signature (p, q) . The total space of the induced orthonormal frame bundle will be denoted with \widehat{M} and the projection $\widehat{M} \rightarrow M$ with π . When doing local operations on an open set U the pre-image $\pi^{-1}(U)$ will be denoted with \widehat{U} . Like \widehat{M} , \widehat{U} also is a principal $O(p, q)$ -bundle.

6.1.1 Principal bundle and Cartan connection on local charts

In this section we locally describe the Cartan connection and the principal bundle. The description chosen here works by using locally defined orthonormal frame-fields. Such fields provide a trivialisation of the fibre bundle on their domain of definition, and as such the global form of the fibre bundle must in general be obtained by gluing these trivialisations together.

Orthonormal frame field

In subsequent sections many constructions will depend explicitly on choosing a (local) orthonormal frame-field around any point in the manifold. As a reminder a frame-field defined on an open set $U \subset M$ will be given by n smooth vectorfields (b_1, \dots, b_n) defined on U so that $g_x(b_i, b_j) = \pm 1$ for all i, j and all points $x \in U$. This is equivalent to a smooth local section $U \rightarrow \pi^{-1}(U)$ of the frame bundle. Here we briefly describe some methods of how such frame-fields can be found.

In the examples considered it turns out that (M, g) admits a universal cover $(\widetilde{M}, \widetilde{g})$ where \widetilde{M} so that the covering map $\pi' : \widetilde{M} \rightarrow M$ is a local isometry and the cover has a globally defined orthonormal frame $(\widetilde{b}_1, \dots, \widetilde{b}_n)$. Since any point $x \in M$ admits a neighbourhood U for which $\pi'^{-1}(U)$ decomposes into a union of disjoint leaves on which π' is an isometry, choosing such a leaf induces an orthonormal frame $(D\pi'[\widetilde{b}_1], \dots, D\pi'[\widetilde{b}_n])$ on U .

Another, computationally more obscure, way of finding local frames around a point x would be to choose an orthogonal basis $(b_{1,x}, \dots, b_{n,x})$ of $T_x M$ and a neighbourhood V of 0 in $T_x M$ so that the exponential mapping is well-defined and a diffeomorphism on its image. On $U = \exp_x(V)$ we then define a frame-field by parallel transporting the basis $b_{i,x}$ radially:

$$b_{i, \exp_x(v)} = P_0^1(t \mapsto \exp_x(tv)) [b_{i,x}],$$

where $P_s^t(\gamma)$ denotes the parallel transport $T_{\gamma(s)} M \rightarrow T_{\gamma(t)} M$ along the parametrised curve γ .

A relevant quantity for calculations are the structure ‘‘constants’’ of the orthonormal frame. These will be used in obtaining the differential equations determining the local Killing fields as well as in expressing the curvature. We define them as follows:

Definition 6.1.1 (Structure constants). Let (b_1, \dots, b_n) be a frame on U . The structure constants or structure functions are the functions $\gamma_{ij}^k : U \rightarrow \mathbb{R}$, $i, j, k \in \{1, \dots, n\}$ determined by:

$$[b_i, b_j]_x = \sum_{ij} \gamma_{ij}^k(x) b_{k,x}$$

Remark. We remark that the structure constants are anti-symmetric in the indices i, j . Choosing the basis so that these structure constants are as simple as possible is helpful in simplifying calculations.

Locally trivialising the principal bundle

Having an orthonormal frame (b_1, \dots, b_n) on an open set U provides us with a local section $s : U \rightarrow \pi^{-1}(U)$, $x \mapsto (b_{1,x}, \dots, b_{n,x})$ of the principal bundle \widehat{M} . This local section provides us with a local trivialisation of the bundle \widehat{M} on U , to be specific the map

$$\Psi : U \times O(p, q) \rightarrow \pi^{-1}(U), \quad (x, g) \mapsto (b_{1,x}, \dots, b_{n,x}) \cdot g = \left(\sum_{i=1}^n g_{i1} b_{i,x}, \dots, \sum_{i=1}^n g_{in} b_{i,x} \right) \quad (6.1)$$

provides a bundle-isomorphism (here $U \times O(p, q)$ is given $O(p, q)$ action $(x, g) \cdot h = (x, gh)$). In other words a local orthonormal frame allows us to assume the bundle is trivial when we are doing purely local computations.

The Cartan connection on a trivialisation

Let (b_1, \dots, b_n) be an orthonormal frame on an open set U and $\Psi : U \times O(p, q) \rightarrow \pi^{-1}(U)$ the local trivialisation of the frame-bundle induced by this frame. We define $\omega' := \Psi^*(\omega|_{\pi^{-1}(U)})$, where ω is the Cartan connection on \widehat{M} . ω' is then a Cartan connection on $U \times O(p, q)$ and Ψ is by definition an isomorphism of $U \times O(p, q)$ into $\pi^{-1}(U)$. We will now sketch the form of ω' and understand this to be a (local) computation of ω .

Reminder. The Cartan connection $\omega : T\widehat{M} \rightarrow \mathfrak{so}(p, q) \times \mathbb{R}^n$ is the sum of the fundamental form $\theta : T\widehat{M} \rightarrow \mathbb{R}^n$, which assigns to each $v \in T_{\widehat{x}}\widehat{M}$ the “expression in coordinates \widehat{x} ” of $D_{\widehat{x}}\pi(v)$ (and from which one can recover the metric g), and the form $\alpha : T\widehat{M} \rightarrow \mathfrak{so}(p, q)$, which classically is the Ehresmann connection associated to the metric connection ∇ . In other words

$$\omega_{\widehat{x}}(v) = \alpha_{\widehat{x}}(v) + \theta_{\widehat{x}}(v)$$

for any $v \in T_{\widehat{x}}\widehat{M}$. We denote with α' and θ' their pullbacks via Ψ to $U \times O(p, q)$.

The following lemma gives the form of ω' on vertical vectors. The result is as expected, ω' is equal to the Maurer-Cartan form ω_{MC} when restricted to vertical vectors:

Lemma 6.1.2. *Let $(0, v)$ be a vertical vector in $T_{(x,g)}(U \times O(p, q)) \cong T_x U \times T_g O(p, q)$. Then $\alpha'_{(x,g)}(0, v) = D_g L_{g^{-1}}(v)$ (where L_g denotes the left multiplication with g) and $\theta'_{(x,g)}(0, v) = 0$.*

Proof. Note that θ' factors over $D\pi$, hence θ' is zero on vertical vectors. Further α' must map the fundamental fields to their generators, so

$$\xi = \alpha'_{(x,g)} \left(\frac{d}{dt}(x, g) \cdot \exp(t\xi)|_{t=0} \right) = \alpha'_{(x,g)}(D_{(x,1)}L_g(\xi)),$$

hence in the decomposition $T_{(x,g)}(U \times O(p, q)) = T_x U \times T_g O(p, q)$ one gets $\alpha'_{(x,g)}|_{T_g O(p,q)} = (D_1 L_g)^{-1} = D_g L_{g^{-1}}$. \square

The way ω' acts on horizontal vectors is more complicated (here we mean horizontal in the decomposition $U \times O(p, q)$), but this complication comes entirely from the Ehresmann connection; the action of θ' on horizontal vectors is clear:

Lemma 6.1.3. *Let $h = \sum_i h_i b_{i,x}$ be a vector in $T_x U$. Then, using the decomposition $T_{(x,g)}(U \times O(p, q)) \cong T_x U \times T_g O(p, q)$, one has $\theta'_{(x,g)}(h, 0) = \sum_{ij} (g^{-1})_{ij} h_j \mathbf{e}^i$ where \mathbf{e}^i is the i -th basis vector of \mathbb{R}^n .*

Proof. Reminding ourselves of the construction of the trivialisation $U \times O(p, q) \rightarrow \widehat{U}$, the point (x, g) corresponds to the basis $(\sum_i g_{i1} b_{i,x}, \dots, \sum_i g_{in} b_{i,x})$. Then $\theta'_{(x,g)}(0, h)$ is the expansion of the vector $h \in T_x U$ in the basis $(\sum_i g_{i1} b_{i,x}, \dots, \sum_i g_{in} b_{i,x})$ (compare to Definition 3.2.8). Note:

$$b_{j,x} = \sum_{ki} b_{k,x} g_{ki} (g^{-1})_{ij} = \sum_i (g^{-1})_{ij} \sum_j g_{ki} b_{k,x},$$

whence

$$\theta'_{(x,g)}(0, h) = \sum_j h_j \sum_i (g^{-1})_{ij} \theta'_{(x,g)} \left(\sum_k g_{ki} b_{k,x} \right) = \sum_{ij} (g^{-1})_{ij} h_j \mathbf{e}^i$$

follows. \square

Together the two lemmas tell us that for $h = \sum_i h_i b_{i,x} \in T_x U$ and $v \in T_g O(p, q)$ we have:

$$\omega'_{(x,g)}(h, v) = \left(D_g L_{g^{-1}}(v) + \alpha'_{(x,g)}(h), \sum_{ij} (g^{-1})_{ij} h_j \mathbf{e}^i \right).$$

Here $\alpha'_{(x,g)}(h)$ is the only term that still needs to be determined. In order to approach this we first define a vertical analog of the structure constants:

Definition 6.1.4 (Vertical structure constants). Let $U \subset M$ and (b_1, \dots, b_n) a smooth ONB on U . For $x \in U$ let:

1. $\varphi_i^t(x)$ denote the flow of the vector field b_i at x .
2. $A_i(t, x) := P_0^t(s \mapsto \varphi_i^s(x)) [b_{1,x}, \dots, b_{n,x}]$ be the parallel transport of the ONB $b_{1,x}, \dots, b_{n,x}$ along $\varphi_i^t(x)$.

3. Let $\mathcal{P}_i(t, x)$ denote the element of $O(p, q)$ so that $A_i(t, x) = (b_{1, \varphi_i^t(x)}, \dots, b_{n, \varphi_i^t(x)}) \cdot \mathcal{P}_i(t, x)$.
4. Let $\mathfrak{P}_i(x)$ denote $\frac{d}{dt} \mathcal{P}_i(t, x)|_{t=0}$, since $\mathcal{P}_i(0, x) = \mathbb{1}$ and so $\mathfrak{P}_i(x) \in \mathfrak{so}(p, q)$.

We call \mathfrak{P}_i the “vertical structure constants” of the trivialisation induced by the basis (b_1, \dots, b_n) .

Note that if $\mathfrak{P}_i(x) = 0$ for all $x \in U$ that then the basis (b_1, \dots, b_n) is parallel along the flows φ_i^t . In such a situation many computations will simplify, but this scenario is not to be expected. As a remark, which will not be proven, if the metric connection is torsion free then $\mathfrak{P}_i = 0$ implies that the metric is flat.

Lemma 6.1.5. *Let $h = \sum_i h_i b_{i,x}$ be vector in $T_x U$. Then, using the decomposition $T_{(x,g)}(U \times O(p, q)) \cong T_x U \times T_g O(p, q)$ one has: $\alpha'_{(x, \mathbb{1})}(h, 0) = -\sum_i h_i \mathfrak{P}_i(x)$.*

Proof. We show $\alpha'_{(x, \mathbb{1})}(b_{i,x}, 0) = -\mathfrak{P}_i(x)$ for any i . The idea of the proof is to consider a path $\gamma_i(t)$ with $\gamma(0) = x$ and $\dot{\gamma}(0) = b_{i,x}$. Denoting with $\hat{\gamma}$ the horizontal lift of γ at the basis $(b_{1,x}, \dots, b_{n,x})$ (cf. Definition 3.2.6) we remark that by definition $\alpha_{\hat{\gamma}(0)}(\frac{d}{dt} \hat{\gamma}(0)) = 0$. Hence pulling $\hat{\gamma}$ back to the trivialisation $U \times O(p, q)$ will give us

$$\alpha'_{(x, \mathbb{1})} \left(\frac{d}{dt} \gamma(0) - \frac{d}{dt} \hat{\gamma}(0) \right) = \alpha'_{(x, \mathbb{1})} \left(\frac{d}{dt} \gamma(0) \right).$$

However $\frac{d}{dt} \gamma(0) - \frac{d}{dt} \hat{\gamma}(0)$ is clearly vertical, hence $\alpha'_{(x, \mathbb{1})} \left(\frac{d}{dt} \gamma(0) \right) = \frac{d}{dt} \gamma(0) - \frac{d}{dt} \hat{\gamma}(0)$. Thus once we understand this difference we have understood $\alpha'_{(x, \mathbb{1})}(b_{i,x}, 0)$.

Now to be more concrete: For γ we choose $\varphi_i^t(x)$. By definition of the horizontal lift we have that $\widehat{\varphi_i^t(x)} = A_i(t, x)$. In the specific trivialisation given by the basis (b_1, \dots, b_n) the section $y \mapsto (b_{1,y}, \dots, b_{n,y})$ corresponds to $(y, \mathbb{1})$, hence the points $A_i(t, x) = (b_{1, \varphi_i^t(x)}, \dots, b_{n, \varphi_i^t(x)}) \cdot \mathcal{P}_i(t, x)$ correspond to $(\varphi_i^t(x), \mathbb{1} \cdot \mathcal{P}_i(t, x)) = (\varphi_i^t(x), \mathcal{P}_i(t, x))$. The differential of this at 0 is then equal to:

$$\left(\frac{d}{dt} \varphi_i^t(x)|_{t=0}, \frac{d}{dt} \mathcal{P}_i(t, x)|_{t=0} \right) = (b_{i,x}, \mathfrak{P}_i(x)).$$

By the preceding discussion it then follows that $\alpha'_{(x, \mathbb{1})}(b_{i,x}) = -\mathfrak{P}_i(x)$. □

We remind ourselves that the connection form α fulfils the condition: $\alpha_{\widehat{x} \cdot g}(D_{\widehat{x}} R_g(\xi)) = \text{Ad}(g^{-1}) \alpha_{\widehat{x}}(\xi)$ for any $\xi \in T_{\widehat{x}} \widehat{M}$. Hence one may read off:

$$\alpha'_{(x,g)}(h, 0) = \text{Ad}(g^{-1}) \alpha'_{(x, \mathbb{1})}(h, 0),$$

where we used that in the trivialisation $U \times O(p, q)$ the right-multiplication acts trivially on U . With this we consolidate Lemmas 6.1.2, 6.1.3, and 6.1.5 to get:

Corollary 6.1.6. *Let $U \subseteq M$ and (b_1, \dots, b_n) a smooth ONB on U , denote with ω' the pull-back of the Cartan connection on $U \times O(p, q)$. Then for $x \in U$ and $(h, v) = (\sum_i h_i b_{i,x}, v) \in T_{(x,g)}(U \times O(p, q))$ one has:*

$$\omega'_{(x,g)}(h, v) = \left(D_g L_{g^{-1}}(v) - \text{Ad}(g^{-1}) \sum_i h_i \mathfrak{P}_i(x), \sum_{ij} (g^{-1})_{ij} h_j e^i \right).$$

6.1.2 Local Killing fields

While a Killing field is defined as a vector field on M that flows by automorphisms, we have noted in Section 2.5 and Proposition 2.5.4 that any such field uniquely determines a vector field on \widehat{M} flowing by lifts of automorphisms. In determining the local Killing algebras we will describe when a vector field on \widehat{M} is such a lift. Thus we will in the rest of the chapter, for reasons of convenience, use the term ‘‘Killing field’’ to describe this *lift* on \widehat{M} of a Killing field on M .

We recall Lemma 2.5.5, which states that a vector field X on \widehat{M} is (the lift of) a Killing field if and only if it is right-invariant and commutes with all constant fields. This is useful, as in local coordinates the condition that X commutes with a family of vector fields can be formulated as a family of differential equations in the components of X . What is left to do in this section is to determine the form of the constant fields and from this the system differential equations determining when a vector field is a Killing field.

So we continue with the setting of the previous section and assume we have an orthonormal frame (b_1, \dots, b_n) on some open set U . This gives a trivialisation $U \times P$ of $\pi^{-1}(U)$ and the form of the Cartan connection on $U \times P$ is described in Corollary 6.1.6.

Lemma 6.1.7. *Adopting the notation of Corollary 6.1.6, the ω' -constant fields on $U \times O(p, q)$ are of the form:*

$$C(v, h)_{(x,g)} = \left(D_{\mathbb{1}} L_g(v) + \sum_{kj} g_{kj} h_j D_{\mathbb{1}} R_g(\mathfrak{P}_k(x)), \sum_{kj} g_{kj} h_j b_k \right)$$

and $\omega'(C(v, h)) = (v, h)$, for $h = \sum_i h_i e^i \in \mathbb{R}^n$ and $v \in \mathfrak{so}(p, q)$.

Proof. This is a direct calculation, applying Corollary 6.1.6. □

One immediately notices that this expression is rather unfriendly, in particular the fact that \mathfrak{P}_k depends on x and depends ‘‘quadratically’’ on g makes calculating commutators unwieldy. Whenever possible we will use tricks in order to bypass calculating with \mathfrak{P}_k .

In what follows we $U \subseteq M$ will be an open subset of a pseudo-Riemannian manifold M admitting an ONB (b_1, \dots, b_n) . The pull back of the Cartan connection ω to the trivialisation $U \times O(p, q)$ induced by the basis will be denoted with ω' . In order to simplify calculations we will now embed $O(p, q)$ in GL_n ($n = p + q$) and $\mathfrak{so}(p, q)$ in $M_{n \times n}$. The effect of this is that we may write identify $T_g O(p, q)$ with $g \cdot \mathfrak{so}(p, q) = \mathfrak{so}(p, q) \cdot g$, specifically the element $\mathfrak{a} \cdot g$ corresponds to $\sum_{ijk} (\mathfrak{a}_{ij} g_{jk}) \frac{\partial}{\partial g_{ik}}$. This will make calculating commutators easier.

Lemma 6.1.8. *If X is a right-invariant vector field on $U \times O(p, q)$, then X is necessarily of the form $X_{(x,g)} = H_x + D_1 R_g(\mathfrak{a}(x)) = H_x + \mathfrak{a}(x) \cdot g$ where H_x is a field on U and $\mathfrak{a}(x) \in \mathfrak{so}(p, q)$.*

Proof. Write $X_{(x,g)} = H_{(x,g)} + V_{(x,g)}$ with H horizontal and V vertical. By right-invariance we have $H_{(x,g)} = D_1 R_g(H_{(x,1)}) = H_{(x,1)}$ and $V_{(x,g)} = D_g R_g(V_{(x,1)}) = V_{(x,1)} \cdot g$. With $H_x := H_{(x,1)}$ and $\mathfrak{a}_x := V_{(x,1)}$ the proposition follows. \square

Since the fields b_1, \dots, b_n form a linear frame on TU , we may write any expand horizontal field H in the form:

$$H_x = \sum_l f_l(x) b_{l,x}.$$

In particular if we do this with the horizontal part of a Killing field we will recover a system of differential equations on the functions f_l :

Lemma 6.1.9. *Suppose $X = \mathfrak{a}(x) \cdot g + \sum_{l=1}^n f_l(x) b_{l,x}$ is a Killing field. Then $\mathfrak{a}(x)$ and $f_l(x)$ satisfy the following differential equation:*

$$\sum_l \gamma_{li}^k(x) f_l(x) - b_i(f_k(x)) + \mathfrak{a}_{ki}(x) = 0$$

for $i, k \in \{1, \dots, n\}$, where $\gamma_{ij}^k(x)$ are the structure functions as in Definition 6.1.1.

Proof. Let $C(0, h)_{(x,g)}$ be a constant field as in Lemma 6.1.7. We will abbreviate the form of $C(0, h)_{(x,g)}$ by writing

$$C_h = \sum_i (g \cdot h)_i (\mathfrak{P}_i \cdot g) + \sum_i (g \cdot h)_i b_i = C_h^{\text{vert}} + C_h^{\text{hor}}.$$

The first thing to notice is that when we take the commutator of C_h with X , that $[X, C_h^{\text{vert}}]$ is vertical, since the horizontal component of X does not depend on g . This is a relief, as then only $[X, C_h^{\text{hor}}]$ may have horizontal components and this term doesn't involve the vertical structure constants \mathfrak{P}_i . If we expand it we get:

$$[X, C_h^{\text{hor}}] = \sum_{il} (g \cdot h)_i [f_l b_l, b_i] + \sum_{mni} (\mathfrak{a}(x) \cdot g)_{mn} \partial_{g_{mn}} (g \cdot h)_i b_i - \sum_{imln} (g \cdot h)_i b_i (\mathfrak{a}_{ml}(x)) g_{ln} \partial_{g_{mn}}.$$

Dropping the last summand (which is vertical), making use of the structure constants, expanding $\partial_{g_{mn}}(g \cdot h)_i = \delta_{mi}h_n$, and re-indexing gives:

$$\sum_{ijkl} g_{ij} h_j f_l \gamma_{li}^k(x) b_k + \sum_{ijl} g_{ij} h_j b_i(f_l) b_l + \sum_{ij} (\mathbf{a}(x) \cdot g)_{ij} h_j b_i.$$

This may be further simplified to:

$$\sum_{ij} g_{ij} h_j \sum_k \left(\sum_l f_l(x) \gamma_{li}^k(x) + b_i(f_k) + \mathbf{a}(x)_{ki} \right) b_k.$$

This is the horizontal component of $[X, C_h]$, and as such must be 0 if X is a Killing field. We take the k -component in the basis $b_{k,x}$ of this equation and note that no linear relation exists on all of $O(p, q)$ and that h_j is arbitrary. These equations then imply (and are implied by):

$$\sum_l \gamma_{li}^k(x) f_l(x) - b_i(f_k(x)) + \mathbf{a}_{ki}(x) = 0$$

for all $i, k \in \{1, \dots, n\}$. □

6.2 Hyperbolic torus bundles and the torus

Here we give the isometry group (up to finite index), the local form of the Cartan connection, the local Killing algebras, the curvature and the integrability locus of the manifolds described by the hyperbolic case, that is those in Theorem 5.4.7.

We remind ourselves of their definition. Let $A \in SL_2(\mathbb{Z})$ be a hyperbolic matrix and u, v a choice of two eigenvectors of A (and denote the eigenvalues with λ, λ^{-1}). For a 1-periodic function $a : \mathbb{R} \rightarrow \mathbb{R}_{>0}$ we define a Riemannian metric \tilde{g} on \mathbb{R}^3 by $\tilde{g} = dt^2 + a(t)dudv$, in order to give meaning to dt, du and dv we give \mathbb{R}^3 coordinates (t, x, y) and understand A to act on the x, y coordinates in the usual way, whence du, dv are linear combinations of dx, dy . First we note that $\begin{pmatrix} t \\ x \\ y \end{pmatrix} \mapsto \begin{pmatrix} t \\ A \begin{pmatrix} x \\ y \end{pmatrix} \end{pmatrix}$ is an isometry, since $A^*(dudv) = d(Au)d(Av) = \lambda du \lambda^{-1} dv = dudv$. Now we define two discrete subgroups of $\text{Isom}(M, \tilde{g})$:

$$\Gamma_1 = \left\{ \begin{pmatrix} t \\ x \\ y \end{pmatrix} \mapsto \begin{pmatrix} t \\ x \\ y \end{pmatrix} + \begin{pmatrix} n_1 \\ n_2 \\ n_3 \end{pmatrix} \mid n_1, n_2, n_3 \in \mathbb{Z} \right\} \cong \mathbb{Z}^3$$

$$\Gamma_2 = \left\{ \begin{pmatrix} t \\ x \\ y \end{pmatrix} \mapsto \begin{pmatrix} t \\ A^{n_1} \begin{pmatrix} x \\ y \end{pmatrix} \end{pmatrix} + \begin{pmatrix} n_1 \\ n_2 \\ n_3 \end{pmatrix} \mid n_1, n_2, n_3 \in \mathbb{Z} \right\} \cong \mathbb{Z} \times_A \mathbb{Z}^2,$$

taking the quotient $\Gamma_1 \backslash \mathbb{R}^3$ we recover a torus \mathbb{T}^3 with a metric described in Theorem 5.4.7, the quotient $\Gamma_2 \backslash \mathbb{R}^3$ gives the hyperbolic mapping torus \mathbb{T}_A^3 with the same kind of metric. We

denote the metric on \mathbb{T}^3 and \mathbb{T}_A^3 with g , in the local coordinates given by the covering we write $g = dt^2 + a(t)dudv$.

In what follows we will be doing all calculations on $(\mathbb{R}^3, \tilde{g})$. Since the covering map is a local isometry any local calculation on \mathbb{R}^3 will work just as well on the quotient $\Gamma_i \backslash \mathbb{R}$.

6.2.1 Isometry group

Any isometry of $(\mathbb{T}^3, \tilde{g})$ or $(\mathbb{T}_A^3, \tilde{g})$ must lift to an isometry of $(\mathbb{R}^3, \tilde{g})$ that normalises Γ_1 or Γ_2 respectively, meaning if F is such a lift one must have $F\Gamma_i F^{-1} = \Gamma_i$. On the other hand any isometry of $(\mathbb{R}^3, \tilde{g})$ normalising Γ_1 or Γ_2 induces an isometry on \mathbb{T}^3 or \mathbb{T}_A^3 as well.

For generic (but 1-periodic) a the isometry group of (\mathbb{R}^3, g) is isomorphic to $\mathbb{Z} \times (O(1, 1) \ltimes \mathbb{R}^2)$, where $(n, (M, \xi))$ acts as

$$\begin{pmatrix} t \\ x \\ y \end{pmatrix} \mapsto \begin{pmatrix} t + n \\ B^{-1}MB \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} \xi_1 \\ \xi_2 \end{pmatrix} \end{pmatrix}.$$

here B base change map sending $u + v$ to e^1 and $u - v$ to e^2 .

These maps descend to isometries $\mathbb{Z} \ltimes_A \mathbb{T}^2$ on the tori.

6.2.2 Orthonormal basis and derived quantities

We note that \mathbb{R}^3 has a global ONB given by

$$(b_1, b_2, b_3)_{\mathbf{x}} = \left(\partial_t, \frac{1}{\sqrt{a(t)}}(\partial_v + \partial_u), \frac{1}{\sqrt{a(t)}}(\partial_u - \partial_v) \right). \quad (6.2)$$

We will now calculate the trivialisation of $\widehat{\mathbb{R}}^3$ induced by this basis as well as the structure constants γ_{ij}^k , the flows $\varphi_{b_i}^\tau$ of the basis and the vertical structure constants $\mathfrak{P}_i(x) \in \mathfrak{so}(2, 1)$ induced by the basis and the flows.

Using Equation (6.1) the trivialisation of $\widehat{\mathbb{R}}^3 \cong \mathbb{R}^3 \times O(2, 1)$ induced by this basis is:

$$((t, x, y), g) \mapsto \left(\sum_{i=1}^3 g_{i1} b_{i,(t,x,y)}, \sum_{i=1}^3 g_{i2} b_{i,(t,x,y)}, \sum_{i=1}^3 g_{i3} b_{i,(t,x,y)} \right).$$

The structure constants γ_{ij}^k from Definition 6.1.1 are here:

$$\gamma_{21}^2(t, x, y) = \frac{1}{2} \frac{a'(t)}{a(t)} = \gamma_{31}^3(t, x, y),$$

with all other values being determined by anti-symmetry or 0 and these calculations being elementary. Further one can check that:

$$\begin{aligned}\varphi_{b_1}^\tau(t, x, y) &= \begin{pmatrix} t + \tau \\ x \\ y \end{pmatrix} \\ \varphi_{b_2}^\tau(t, x, y) &= \begin{pmatrix} t \\ x \\ y \end{pmatrix} + \frac{\tau}{\sqrt{a(t)}}(u + v) \\ \varphi_{b_3}^\tau(t, x, y) &= \begin{pmatrix} t \\ x \\ y \end{pmatrix} + \frac{\tau}{\sqrt{a(t)}}(u - v)\end{aligned}$$

describe the flows of the basis (6.2). We may calculate:

$$\mathfrak{P}_1 = 0, \quad \mathfrak{P}_2(t, x, y) = \frac{a'(t)}{4a(t)} \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad \mathfrak{P}_3(t, x, y) = \frac{a'(t)}{4a(t)} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

but these matrices will not be used in the computations, since the differential equations given by Lemma 6.1.9 are enough here.

6.2.3 Cartan connection

By Corollary 6.1.6 the concrete form of the Cartan connection is:

$$\omega : T(\mathbb{R}^3 \times O(2, 1)) \rightarrow \mathfrak{so}(2, 1) \ltimes \mathbb{R}^3, \quad \left(\sum_i h_i b_{i, (t, x, y)}, \mathbf{a} \right)_{((t, x, y), g)} \mapsto \left(g^{-1} \cdot \mathbf{a}, \sum_{ij} (g^{-1})_{ij} h_j \mathbf{e}^i \right).$$

Here \mathbf{a} is viewed as an element of $T_g O(2, 1) = g \cdot \mathfrak{so}(2, 1)$.

6.2.4 Local Killing fields

We will now consider a connected open neighbourhood $U \subset \mathbb{R}^3$ of a point (t_0, x_0, y_0) and determine the local Killing fields on U . If we have need of it we will make U as small as we wish, so that this calculation computes the Killing algebra $\mathfrak{kill}_{(t_0, x_0, y_0)}$. There are 3 cases, either $a(t)$ is constant, $\frac{a'(t)}{a(t)}$ is constant or neither of the two is constant, corresponding to flat, locally homogenous and non-homogenous cases respectively. Lemmas 6.2.1, 6.2.2, 6.2.3 describe the form the Killing fields take for these 3 cases.

Since the covering map is a local isometry these (local) Killing fields are then the same as the fields on $\Gamma_i \backslash \mathbb{R}^3$, provided U is small enough.

We remind ourselves of Lemma 6.1.9, any local Killing field X is of the form $X_{((t,x,y),g)} = \sum_l f_l(t, x, y) b_{l,((t,x,y))} + \mathfrak{a} \cdot g$ where the functions f_l and the $\mathfrak{so}(2,1)$ element \mathfrak{a} must satisfy the following system of equations:

$$\sum_l \gamma_{li}^k(t) f_l(x) - b_i(f_k(t, x, y)) + \mathfrak{a}_{ik} = 0$$

for $k, i \in \{1, 2, 3\}$. These are 9 equations and we will go through them all. First we remark that for \mathfrak{a} to be in $\mathfrak{so}(2,1)$ we must have $\mathfrak{a}_{ii} = 0$, $\mathfrak{a}_{12} = -\mathfrak{a}_{21}$ and $\mathfrak{a}_{13} = \mathfrak{a}_{31}$, $\mathfrak{a}_{23} = \mathfrak{a}_{32}$. We start by investigating $k = 1$, $i \in \{1, 2, 3\}$:

$$\begin{aligned} 0 &= -b_1(f_1(t, x, y)) = -\partial_t(f_1(t, x, y)) \\ 0 &= -b_2(f_1(t, x, y)) + \mathfrak{a}_{21} \\ 0 &= -b_3(f_1(t, x, y)) + \mathfrak{a}_{31} \end{aligned}$$

implying that f_1 doesn't depend on t and $b_2(f_1), b_3(f_1)$ are constant. Specifically this implies:

$$f_1(t, x, y) = C_1 + \sqrt{a(t)} \cdot (C_2 + \mathfrak{a}_{21}(u(x, y) + v(x, y)) + \mathfrak{a}_{31}(u(x, y) - v(x, y))), \quad (6.3)$$

here $u(x, y), v(x, y)$ denote the u, v components of the vector (x, y) . However if $a(t)$ is not constant around t_0 we find that by $\partial_t f_1 = 0$ we must have $\mathfrak{a}_{21} = -\mathfrak{a}_{12} = 0 = \mathfrak{a}_{31} = \mathfrak{a}_{13}$, as well as $C_2 = 0$ so that $f_1(t, x, y)$ is constant. For the time being we then assume that $a(t)$ is not constant around t_0 . Checking the equations for $i = 1, k \in \{2, 3\}$ then gives:

$$\begin{aligned} 0 &= \gamma_{21}^2 f_2(t, x, y) - b_1(f_2(t, x, y)) = \frac{1}{2} \frac{a'(t)}{a(t)} f_2(t, x, y) - \partial_t f_2(t, x, y) \\ 0 &= \gamma_{31}^3 f_3(t, x, y) - b_1(f_3(t, x, y)) = \frac{1}{2} \frac{a'(t)}{a(t)} f_3(t, x, y) - \partial_t f_3(t, x, y) \end{aligned}$$

resulting in simple differential equations determining the t -dependence of f_2, f_3 . Solving these returns:

$$f_2(t, x, y) = \sqrt{a(t)} \cdot F_2(x, y), \quad f_3(t, x, y) = \sqrt{a(t)} \cdot F_3(x, y)$$

for two functions F_2, F_3 depending only on x and y . For convenience we will define for $l \in \{2, 3\}$:

$$\tilde{F}_l(u(x, y) + v(x, y), u(x, y) - v(x, y)) := F_l(x, y).$$

If we plug this into the equation for $i = 2 = k$ we get:

$$0 = \gamma_{12}^2 f_1 - b_2(f_2(t, x, y)) = -\frac{1}{2} \frac{a'(t)}{a(t)} f_1 - \sqrt{a(t)} \frac{1}{\sqrt{a(t)}} (\partial_u + \partial_v) F_2(x, y) = -\frac{1}{2} \frac{a'(t)}{a(t)} f_1 - \partial_1 \tilde{F}_2$$

since f_1 is constant and the second summand does not depend on t , we find that unless $\frac{a'(t)}{a(t)}$ is constant in t that f_1 must be zero. The case $\frac{a'(t)}{a(t)}$ being constant corresponds to $a(t) = C e^{\lambda t}$

in some neighbourhood of t_0 , which does not correspond to the generic case and will be treated separately. Thus from now on $f_1 = 0$, whence $\partial_1 \tilde{F}_2 = 0$. The very same steps imply for $i = 3 = k$ that $\partial_2 \tilde{F}_3 = 0$. Thus F_2 is a function depending only on $u(x, y) - v(x, y)$ and F_3 depends only on $u(x, y) + v(x, y)$. The only two cases which are left, $k = 2, i = 3$ and $k = 3, i = 2$, become:

$$\begin{aligned} 0 &= -\partial_2 \tilde{F}_2(u - v) + \mathfrak{a}_{32} \\ 0 &= -\partial_1 \tilde{F}_3(u + v) + \mathfrak{a}_{23}. \end{aligned}$$

The solutions of which are $\tilde{F}_2 = \alpha + \mathfrak{a}_{32}(u - v)$ and $\tilde{F}_3 = \beta + \mathfrak{a}_{23}(u + v)$ for constants α, β . We may recap everything in the following lemma:

Lemma 6.2.1. *Suppose $a(t)$ is not constant in any neighbourhood of t_0 and $\frac{a'(t)}{a(t)}$ is not constant in any neighbourhood of t_0 , then the Killing fields on a connected neighbourhood $U \subset \mathbb{R}^3$ of (t_0, x_0, y_0) are:*

$$\begin{aligned} X_{((t,x,y),g)} &= (\alpha + \mathfrak{a}_{23}(u(x, y) - v(x, y))) (\partial_u + \partial_v) \\ &\quad + (\beta + \mathfrak{a}_{23}(u(x, y) + v(x, y))) (\partial_u - \partial_v) + \sum_{i=1}^3 \mathfrak{a}_{ij} g_{jk} \frac{\partial}{\partial g_{ik}} \end{aligned}$$

where α, β are arbitrary constants in \mathbb{R} and \mathfrak{a} an arbitrary element of $\mathfrak{so}(2, 1)$ with only \mathfrak{a}_{23} and \mathfrak{a}_{32} non-zero (and necessarily $\mathfrak{a}_{23} = \mathfrak{a}_{32}$).

We now consider the case that $\frac{a'(t)}{a(t)}$ is constant in some neighbourhood U of t_0 . We will call this constant λ and $\lambda \neq 0$ as we are still assuming that $a(t)$ is not constant. We may keep going from equation (6.3). It becomes:

$$0 = -\frac{\lambda}{2} f_1 - \partial_1 \tilde{F}_2,$$

and we get $\tilde{F}_2 = -\frac{\lambda}{2} f_1 \cdot (u + v) + K_2(u - v)$, where K is a function. Similarly the $i = 3 = k$ equation implies $\tilde{F}_3 = -\frac{\lambda}{2} f_1 \cdot (u - v) + K_3(u + v)$. Finally the $i = 2, k = 3$ and $i = 3, k = 2$ equations are, as before:

$$-\partial_2 \tilde{F}_2 + \mathfrak{a}_{32} = 0 = -\partial_1 \tilde{F}_3 + \mathfrak{a}_{23}$$

implying $K_2(u - v) = \alpha + \mathfrak{a}_{32} \cdot (u - v)$ and $K_3(u + v) = \alpha + \mathfrak{a}_{23} \cdot (u + v)$. In the case $f_1 = 0$ the Killing field is then of the form as before. For $f_1 \neq 0$ we get a new field however:

Lemma 6.2.2. *Suppose $a(t)$ is not constant in any neighbourhood of t_0 . Then for any connected neighbourhood U of (t_0, x_0, y_0) so that $\frac{a'(t)}{a(t)} =: \lambda$ is constant on $U \subset \mathbb{R}^3$, the Killing fields are of the form:*

$$X'_{((t,x,y),g)} = X_{((t,x,y),g)} + f_1 \partial_t + -\lambda f_1 (u(x, y) \partial_u + v(x, y) \partial_v)$$

where X is a Killing field of the form given in Lemma 6.2.1.

Proof. Most of the work has been done, we just unpack (in the case $\alpha = \beta = \mathbf{a}_{23} = 0$) the term $f_2(x, y) b_2 = -\frac{\lambda}{2} f_1 \cdot (u - v)(\partial_u + \partial_v)$ and $f_3(x, y) b_3 = -\frac{\lambda}{2} f_1 \cdot (u + v)(\partial_u - \partial_v)$. Adding them together results in $-\lambda f_1 u \partial_u - \lambda f_1 v \partial_v$, which is the term in the parenthesis. \square

If $a(t)$ is constant the situation is different. In this case all γ_{ij}^k are zero. The equations a Killing field must obey thus become:

$$b_i(f_k(t, x, y)) = \mathbf{a}_{ik}.$$

We may solve this equation to get:

$$f_k(t, x, y) = C_k + \mathbf{a}_{1k}t + \sqrt{a} \mathbf{a}_{2k}(u + v) + \sqrt{a} \mathbf{a}_{3k}(u - v),$$

where C_k is an arbitrary number and a is the (constant) value of $a(t)$.

Lemma 6.2.3. *Suppose $a(t)$ is constant on some connected neighbourhood $U \subset \mathbb{R}^3$ of (t_0, x_0, y_0) . Then the Killing fields on U are of the form:*

$$X_{(t,x,y)} = \alpha \partial_t + \beta \partial_x + \gamma \partial_y + \sum_{k=1}^3 (\mathbf{a}_{1k}t + \sqrt{a} \mathbf{a}_{2k}(u + v) + \sqrt{a} \mathbf{a}_{3k}(u - v)) \mathbf{b}_k + \sum_{i=1}^3 \mathbf{a}_{ij} g_{jk} \frac{\partial}{\partial g_{ik}}$$

for some $\alpha, \beta, \gamma \in \mathbb{R}$ and $\mathbf{a} \in \mathfrak{so}(2, 1)$.

References

- [Ca 89] Y. Carrière, *Autour de la conjecture de L. Markus sur les variétés affines*. Invent. Math., **95**, (1989), 615–628.
- [DG 90] G. D’Ambra, M. Gromov, *Lectures on transformations groups: geometry and dynamics*, Surveys in Differential Geometry (Cambridge), (1990), 19–111.
- [DZ 10] S. Dumitrescu, A. Zeghib, *Géométries Lorentziennes de dimension 3: classification et complétude*, Geom. Dedicata **149** (2010), 243–273.
- [FG 83] D. Fried, W. Goldman, *Three-dimensional affine crystallographic groups*. Adv. Math., **47** (1), (1983), 1–49.
- [FK 02] R. Feres and A. Katok *Ergodic Theory and Dynamics of G-spaces (with special emphasis on rigidity phenomena)*, Handbook of dynamical systems **1A** (2002).
- [Fr 16] C. Frances *Variations on Gromov’s open-dense orbit theorem*. arXiv: <https://arxiv.org/abs/1605.05755>.
- [Fr 18] C. Frances *Lorentz dynamics on closed 3-manifolds*. arXiv: <https://arxiv.org/abs/1804.08695>.
- [GK 84] W. Goldman, Y. Kamishima, *The fundamental group of a compact flat Lorentz space form is virtually polycyclic*. J. Differential Geom., **19** (1), (1984), 233–240.
- [Gr 88] M. Gromov, *Rigid transformation groups*, Geometrie Differentielle, Travaux en cours, Hermann, Paris, **33**, (1988), 65–141.
- [Gr 99] M. Gromov, *Spaces and questions*. 1999, weblink: <https://www.ihes.fr/~gromov/wp-content/uploads/2018/08/SpacesandQuestions.pdf>.
- [Gr 55] A. Grothendieck, *A General Theory of Fibre Spaces with Structure Sheaf*, University of Kansas (1955).
- [Le 71] J. Lelong-Ferrand, *Transformations conformes et quasi-conformes de varietes riemanniens compactes*, Acad. Roy. Belg. Cl. Sci. Mem. Coll. in-8°, 2^e Ser., (1971) 1-44.
- [KN 63] S. Kobayashi and K. Nomizu *Foundations of Differential Geometry, Volume 1* Wiley (1963).
- [Lee 13] J. Lee *Introduction to Smooth Manifolds*, Graduate Texts in Mathematics (Springer), 2013.

- [Me 11] K. Melnick, *A Frobenius theorem for Cartan geometries, with applications*, L'Enseignement Mathématique (Série II) **57** (2011), no. 1-2, 57–89.
- [No 60] K. Nomizu, *On local and global existence of Killing fields*, Ann. of Math. **72** (1960), no. 2, 105–112.
- [Ob 70] M. Obata *Conformal transformations of Riemannian manifolds*, J. Differential Geometry **4** (1970) 311-333.
- [Pe 16] V. Pécastaing, *On two theorems about local automorphisms of geometric structures*, Ann. Inst. Fourier **66** (2016), no. 1, 175–208.
- [Sh 97] R.W. Sharpe *Differential Geometry: Cartan's generalisation of Klein's Erlangen program*, Springer Graduate Texts in Mathematics (1997)
- [Wi 07] D.K. Wise *Topological Gauge Theory, Cartan Geometry, and Gravity*, PhD thesis (2007)
- [Ze 96] A. Zeghib *Killing fields in compact Lorentz 3-manifolds*, J. Differential Geom., **43**, (1996), 859–894.