SOME RESULTS FOR COMPUTING AUTOMORPHISM GROUPS OF CARTAN GEOMETRIES

JACOB W. ERICKSON

Abstract. Following a brief introduction to Cartan geometries, a few techniques related to the computation of the group of automorphisms of a Cartan geometry are discussed, including the use of covering maps and an expression for the Lie algebra of the automorphism group in terms of the curvature of the geometry at a point.

Contents

1. Introduction

When working with a Riemannian manifold, it is often useful to consider its group of isometries. Given the amount of information that might be gleaned from knowing the isometry group, ranging from bounds on sectional curvature to the dimension of certain homology groups, it is perhaps surprising, at first, to learn that this group is so well-behaved; by a well-known theorem of Myers and Steenrod, this group will always have the structure of a Lie group if the Riemannian manifold has only finitely many connected components.

Unfortunately, the problem of computing these isometry groups is, in general, highly nontrivial. Even to compute the Lie algebra of the isometry group, one must usually solve a system of differential equations, and even then the isometry group may have more than one connected component. Thus, short of computer-assisted methods, more elegant methods for computing these groups are incredibly convenient in practice.

More generally, we can look at the group of automorphisms of a Cartan geometry, a convenient and versatile generalization of Riemannian manifolds that defines structures in a way that appeals to the idea of symmetry much more clearly. This group of automorphisms is a Lie group as well, and many of the techniques that are useful for computing isometry groups remain useful in this broader context. This paper will focus on these techniques to compute the Lie groups of automorphisms of Cartan geometries.

Date: September 7, 2016.
The author composed this paper as part of the 2016 University of Chicago Summer REU.

2. Preliminaries

To motivate the definition of a Cartan geometry, consider an m-dimensional connected Riemannian manifold (M, g), which carries a structure based on that of m-dimensional Euclidean geometry. Imagine attempting to describe its structure in terms of m-dimensional Euclidean isometries. At each point $p \in M$, we have a local notion of rotations, reflections, and translations; we can "translate" points around p in some direction, and we can "reflect" points or "rotate" them about p locally, though these transformations might not extend to global symmetries of the manifold. The Lie group of Euclidean symmetries, which we will call $\mathfrak{I}(m)$, is also of a higher dimension than M, so to include these local Euclidean transformations as part of the space, we need to extend M. Since the rotations and reflections do not move their base points, we might as well extend M by $O(m)$, the group of rotations and reflections of Euclidean space fixing 0, to get a principal $O(m)$-bundle over M. Then, to include these local Euclidean transformations, we just need a map from each tangent space of this new principal $O(m)$-bundle to $\mathfrak{i}(m)$, the Lie algebra of $\mathfrak{I}(m)$, which acts in such a way as to make the Lie algebra $\mathfrak{o}(m) \leq \mathfrak{i}(m)$ of $O(m)$ consistent in some sense with the right action of $O(m)$ on this principal $O(m)$-bundle. With this, we get an object that we will call a Cartan geometry of type $(\mathfrak{i}(m), O(m))$ over M.[1]

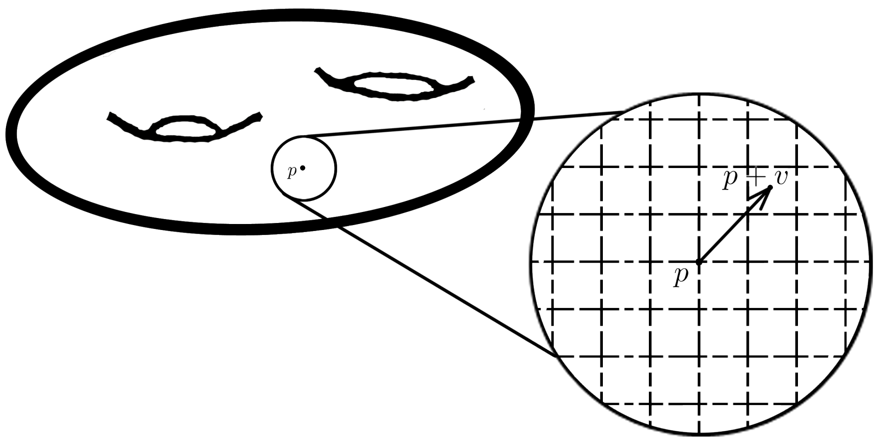

Figure 1. A local translation in normal coordinates

Definition 2.1. A *Cartan structure type* is a pair $(\mathfrak{g}, H)$, where H is a Lie group and $\mathfrak{g}$ is a Lie algebra with a specified Lie subalgebra $\mathfrak{h} \leq \mathfrak{g}$ isomorphic to the Lie algebra of H.

It is important to note that the Lie subalgebra $\mathfrak{h}$ in the above definition is, as it says, a specific Lie subalgebra of $\mathfrak{g}$. If we were to choose a different subalgebra that was also isomorphic to the Lie algebra of H, then we would get a different Cartan structure type. For most of this paper, we will simply refer to a Cartan structure type $(\mathfrak{g}, H)$ and implicitly assume a choice of $\mathfrak{h}$; this is often done in practice because the choice of $\mathfrak{h}$ is apparent from context. However, to insure clarity for those still

[1] It should be noted that this only defines the Riemannian manifold up to scale, for the same reason that the Erlangen Program approach to Euclidean geometry by writing $\mathbb{R}^m$ as $\mathfrak{I}(m)/O(m)$ only determines distance on $\mathbb{R}^m$ up to scale.

new to the subject, we shall explicitly specify the Lie subalgebra $\mathfrak{h}$ in the examples below.

Alternatively, we could use a pair (G, H), where G is a Lie group and H is a closed subgroup, to specify the type of structure instead of using what we call a Cartan structure type. This is the approach taken by [1]. However, for Lie groups G and G' with Lie algebra $\mathfrak{g}$ and containing H as a closed subgroup that determines the same Lie subalgebra of $\mathfrak{g}$, the pairs (G, H) and (G', H) essentially determine the same structures. Thus, unless we need to fix a choice of Lie group G for the Lie algebra $\mathfrak{g}$, it seems more sensible to use Cartan structure types as defined above.

For a Cartan structure type $(\mathfrak{g}, H)$, we can extend the adjoint representation of H on $\mathfrak{h}$ to all of $\mathfrak{g}$, since any choice of Lie group G with Lie algebra $\mathfrak{g}$ and with $H \leq G$ determining the appropriate subalgebra of $\mathfrak{g}$ will have the same adjoint representation on $\mathfrak{g}$ when restricted to H.

Suppose H is a Lie group and P is a principal H-bundle. In what follows, we will denote the right action of an element $h \in H$ on a point $p \in P$ by either ph or $R_h p$.

Definition 2.2. Suppose M is a connected smooth manifold and $(\mathfrak{g}, H)$ is a Cartan structure type. Then, a *Cartan geometry of type* $(\mathfrak{g}, H)$ over M is a pair $(\mathcal{G}, \omega)$, where $\mathcal{G}$ is a principal H-bundle over M and ω is a $\mathfrak{g}$-valued 1-form on $\mathcal{G}$ such that the following are satisfied:

- For every $p \in \mathcal{G}$, $\omega_p : T_p\mathcal{G} \to \mathfrak{g}$ is a linear isomorphism;
- For every $h \in H$, $R_h^* \omega = \mathrm{Ad}_{h^{-1}} \omega$;
- For every $Y \in \mathfrak{h}$ and $p \in \mathcal{G}$, $\frac{d}{dt}\big|_0 p \exp(tY) = \omega_p^{-1}(Y)$.

Remark. Some mathematicians attempt to describe this structure in terms of "rolling without slipping." The idea is to assign structure by, as the name implies, rolling a homogeneous space along the manifold without letting it slip. While this can be a pleasant visual for those outside the subject, it can be confusing for someone not quite used to working with these objects. To see why, try to visualize a plane rolling along another plane without letting it slip; without an understanding of how the idea works, the notion is objectively absurd. Thus, to prevent confusion, the author recommends ignoring this visual while first learning the basics of the subject.

In this way, we can look at projective manifolds, affine manifolds, and a plethora of other geometric structures as instances of the same class of objects. For more on how this works, see [6].

Suppose G is a Lie group acting on a smooth manifold M. We will denote the left action of an element $g \in G$ on a point $p \in M$ by either gp or $L_g p$.

Example 2.3. Denote each element of $\mathfrak{I}(n)$ by a pair (p, A), where $p \in \mathbb{R}^n$ and $A \in O(n)$. We can write the group operation of $\mathfrak{I}(n)$ as $(p, A)(q, B) = (Aq+p, AB)$. The Lie algebra $\mathfrak{i}(n)$ is then given by pairs (v, X), with $v \in \mathbb{R}^n$ and $X \in \mathfrak{o}(n)$, with bracket given by $[(v, X), (w, Y)] = (Xw - Yv, [X, Y])$.

The pair $(\mathfrak{I}(n), \omega_{\mathfrak{I}(n)})$, where $\omega_{\mathfrak{I}(n)} : (p, A) \mapsto L_{(p,A)^{-1}*}$ is the Maurer-Cartan form on $\mathfrak{I}(n)$, is a Cartan geometry of type $(\mathfrak{i}(n), O(n))$ over $\mathbb{R}^n$, where the Lie algebra of $O(n)$ is taken to be the subalgebra of elements of the form $(0, X) \in \mathfrak{i}(n)$. To see that this is a Cartan geometry, note that $\mathfrak{I}(n)$ does, in fact, form a principal $O(n)$-bundle over $\mathbb{R}^n$, with right action $R_B : (p, A) \mapsto (p, AB)$, and that $\omega_{\mathfrak{I}(n)}$ satisfies all of the appropriate conditions:

- For every $(p, A) \in \mathfrak{I}(n)$, left multiplication $L_{(p,A)^{-1}}$ by $(p, A)^{-1}$ is a diffeomorphism on $\mathfrak{I}(n)$, so the pushforward of $L_{(p,A)^{-1}}$ at (p, A) is a linear isomorphism from $T_{(p,A)}\mathfrak{I}(n)$ to $\mathfrak{i}(n)$;

- For every $B \in O(n)$ and every $(p, A) \in \mathfrak{I}(n)$,
$$\begin{aligned} R_B^* \omega_{\mathfrak{I}(n)}(X_{(p,A)}) &= L_{(p,AB)^{-1}*} \circ R_{B*}(X_{(p,A)}) \\ &= L_{(0,B)^{-1}*} \circ R_{B*} \circ L_{(p,A)^{-1}*}(X_{(p,A)}) \\ &= \mathrm{Ad}_{B^{-1}} \omega_{\mathfrak{I}(n)}(X_{(p,A)})\,; \end{aligned}$$
- For every $Y \in \mathfrak{h}$, $\omega_{\mathfrak{I}(n)}^{-1}(Y)$ is the left-invariant vector field corresponding to Y, so the flow of $\omega_{\mathfrak{I}(n)}^{-1}(Y)$ is just $R_{\exp tY}$.

More generally, given a Lie group G with Lie algebra $\mathfrak{g}$ and a closed subgroup H, (G, ω_G), where $\omega_G : g \mapsto L_{g^{-1}*}$ is the Maurer-Cartan form on G, is an example of a Cartan geometry of type $(\mathfrak{g}, H)$. These examples, called Klein geometries, are particularly nice to keep in mind, since they are the inspiration behind the definition of a Cartan geometry.

Definition 2.4. Suppose G is a Lie group with Lie algebra $\mathfrak{g}$, and H is a closed subgroup of G. Then, (G, ω_G) is a Cartan geometry of type $(\mathfrak{g}, H)$ called the *Klein geometry* modeled on (G, H).

One obstruction to a Cartan geometry being a Klein geometry is the curvature of the geometry.

Definition 2.5. Suppose $(\mathcal{G}, \omega)$ is a Cartan geometry of type $(\mathfrak{g}, H)$. Then, the *curvature*[2] Ω of $(\mathcal{G}, \omega)$ is given by $\Omega = \mathrm{d}\omega + \frac{1}{2}[\omega, \omega]$, where the bracket $[\cdot, \cdot]$ on the right-hand side is the bilinear map determined by $[X \otimes \alpha, Y \otimes \beta] = [X, Y] \otimes (\alpha \wedge \beta)$ if we identify the space of $\mathfrak{g}$-valued differential forms on $\mathcal{G}$ with $\mathfrak{g} \otimes \Gamma(T^\vee \mathcal{G})$. If $\Omega = 0$, then the geometry is *flat*.

Unfortunately, the above definition of curvature as a 2-form with values in a Lie algebra makes it seem like an esoteric algebraic object. This disservice is further compounded by what seems to be a ubiquitous lack of effort to demonstrate the geometric intuition behind the definition. In truth, there is a very satisfying geometric interpretation of curvature, which we present below.

Suppose $(\mathcal{G}, \omega)$ is a Cartan geometry of type $(\mathfrak{g}, H)$. Let G be a Lie group with Lie algebra $\mathfrak{g}$ and let $X, Y \in \mathfrak{g}$. In essence, the bracket of X and Y is the endpoint of an "infinitesimal path" in G starting at the identity, moving along X, then moving along Y, then going back along X, and finally going back along Y. Likewise, at a point $p \in \mathcal{G}$, the Lie bracket of $\omega^{-1}(X)$ with $\omega^{-1}(Y)$ is the endpoint of the "infinitesimal path" starting at p and moving around along $\omega^{-1}(X)$ and $\omega^{-1}(Y)$. Using our connection ω, we can put $[\omega^{-1}(X), \omega^{-1}(Y)]_p$ in $\mathfrak{g}$ as well. The difference between the endpoint of the standard "loop", given by the bracket of the Lie algebra $[X, Y]$, and the endpoint of the "loop" from the geometry, given as $\omega_p([\omega^{-1}(X), \omega^{-1}(Y)])$, is the curvature applied to the oriented tangent parallelogram $\omega^{-1}(X) \wedge \omega^{-1}(Y)$ at p. That is,
$$\Omega_p(\omega^{-1}(X) \wedge \omega^{-1}(Y)) = [X, Y] - \omega_p([\omega^{-1}(X), \omega^{-1}(Y)]).$$

Thus, a flat geometry is just one where moving along these loops in the geometry is infinitesimally the same as moving along loops in a Lie group with the appropriate Lie algebra; in other words, $(\mathcal{G}, \omega)$ is flat if and only if ω^{-1} is a Lie algebra isomorphism onto its image in $\Gamma(T\mathcal{G})$, the Lie algebra of vector fields under the Lie bracket.

As in Riemannian geometry, there is also a notion of completeness for Cartan geometries.[3]

---

[2]To see how this notion of curvature is related to curvature in Riemannian geometry, see Theorem 4.3 of [3]

[3]Indeed, these notions are equivalent in the Riemannian case; see Theorem 5.7 of [3].

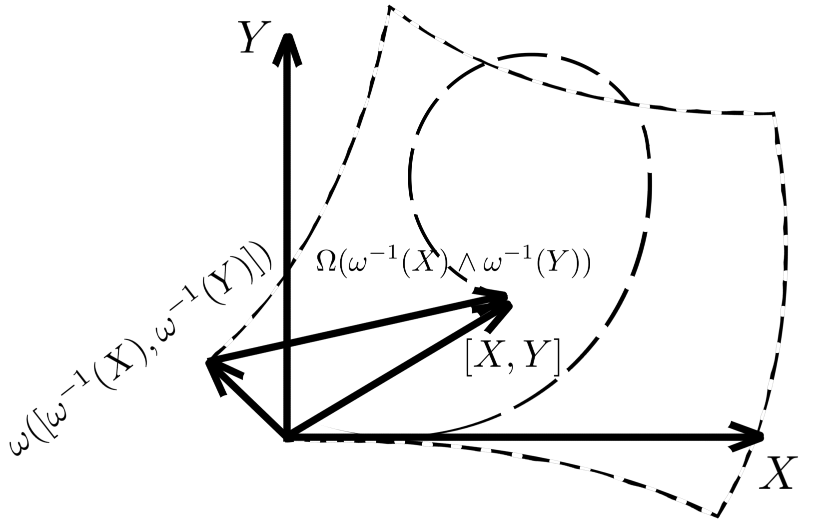

FIGURE 2. A depiction of curvature as a "difference of loops"

Definition 2.6. A Cartan geometry $(\mathcal{G}, \omega)$ of type $(\mathfrak{g}, H)$ is *complete* if and only if $\omega^{-1}(X)$ is a complete vector field[4] for every $X \in \mathfrak{g}$.

The intuition behind completeness, as it is for Riemannian manifolds, is that a geometry is complete when it is not "missing" any points. A good example of a geometry that is "missing" points is that of the open ball with the obvious Riemannian structure given to it as a subset of Euclidean space: in other words, $(B_r(0) \times O(n), \omega_{\mathbb{R}^n \rtimes O(n)}|_{B_r(0)\times O(n)})$ with $r > 0$. Here, the vector fields corresponding to infinitesimal translations are not complete because their flows cannot go beyond the boundary of the open ball; the part of the geometry that they are flowing toward is "missing."

In order to relate different geometries, we will also need to specify a notion of morphisms between them.

Definition 2.7. Suppose $(\mathcal{G}, \omega)$ and $(\mathcal{Q}, \eta)$ are Cartan geometries of type $(\mathfrak{g}, H)$. A *geometric map* is a (smooth) right H-equivariant map $\varphi : \mathcal{G} \to \mathcal{Q}$ such that $\varphi^*\eta = \omega$. A map is a *geometric isomorphism* if and only if it is a bijective geometric map.

We will be primarily interested in the groups $\operatorname{Aut}(\mathcal{G}, \omega)$ of geometric automorphisms of Cartan geometries $(\mathcal{G}, \omega)$. We will call $\operatorname{Aut}(\mathcal{G}, \omega)$ the *symmetry group* of $(\mathcal{G}, \omega)$.

Definition 2.8. Suppose $(\mathcal{G}, \omega)$ is a Cartan geometry of type $(\mathfrak{g}, H)$. The *symmetry group* of $(\mathcal{G}, \omega)$ is the (Lie) group $\operatorname{Aut}(\mathcal{G}, \omega)$ of geometric isomorphisms from $(\mathcal{G}, \omega)$ to itself under composition. The *symmetry algebra* of $(\mathcal{G}, \omega)$ is the Lie algebra $\mathfrak{aut}(\mathcal{G}, \omega)$ of complete vector fields ξ on $\mathcal{G}$ such that $\mathcal{L}_\xi\omega = 0$ and $R_{h*}\xi = \xi$ for all $h \in H$, where the bracket for $\mathfrak{aut}(\mathcal{G}, \omega)$ is given by the negative of the Lie bracket.

We state the following iconic result without proof.

Proposition 2.9 (Chapter IV, Theorem VII in [5])**.** *Let S be a group of diffeomorphisms of a smooth manifold M and let $\mathfrak{s}$ be the space of all complete vector*

[4] Recall that a complete vector field is a vector field X such that its flow $t \mapsto \exp(tX)$ is defined for all $t \in \mathbb{R}$.

fields whose flows are contained in S. If the Lie algebra generated by $\mathfrak{s}$ under the Lie bracket is finite-dimensional, then S is a (possibly not second countable) Lie group with Lie algebra $\mathfrak{s}$.

When we prove, for all Cartan geometries $(\mathcal{G},\omega)$, that $\mathfrak{aut}(\mathcal{G},\omega)$ is finite-dimensional (see Corollary 3.3), the above will show that $\mathrm{Aut}(\mathcal{G},\omega)$ is a (possibly not second countable) Lie group, and since we are choosing to work over a connected base manifold, it is guaranteed to be second countable.[5]

For a more cogent introduction to Cartan geometries, the author highly recommends [6], with the additional advice to carefully check each result thoroughly. For a more advanced treatment, see [1].

3. The Lie algebra $\mathfrak{aut}(\mathcal{G},\omega)$

Let $(\mathcal{G},\omega)$ be a Cartan geometry of type $(\mathfrak{g},H)$.

While we ultimately seek to compute $\mathrm{Aut}(\mathcal{G},\omega)$, a related and slightly more tractable problem is to find the Lie algebra $\mathfrak{aut}(\mathcal{G},\omega)$ of $\mathrm{Aut}(\mathcal{G},\omega)$. By knowing $\mathfrak{aut}(\mathcal{G},\omega)$, we can deduce the identity component of $\mathrm{Aut}(\mathcal{G},\omega)$. Instead of the standard method of computing the symmetry algebra, which involves solving a system of differential equations and then checking which solutions are complete, we will instead rewrite $\mathfrak{aut}(\mathcal{G},\omega)$ in terms of $\mathfrak{g}$ and curvature and develop a way to determine the existence of an element of $\mathfrak{aut}(\mathcal{G},\omega)$ with a given value at a point.

Consider $\xi\in\mathfrak{aut}(\mathcal{G},\omega)$. Since $\mathcal{L}_\xi\omega=0$, we have that

$$\begin{aligned}0=(\mathcal{L}_\xi\omega)(Y)&=\xi(\omega(Y))-\omega([\xi,Y])\\&=\mathrm{d}\omega(\xi\wedge Y)+Y(\omega(\xi))\\&=\Omega(\xi\wedge Y)-[\omega(\xi),\omega(Y)]+Y(\omega(\xi))\end{aligned}$$

for all $Y\in\Gamma(T\mathcal{G})$. Note that, if we set $Y=\omega^{-1}(X)$ for any $X\in\mathfrak{g}$, the first term on the far right side of the first line is zero, since it is just differentiating a constant function, and since ω is an isomorphism at each point, this tells us that $[\xi,\omega^{-1}(X)]=0$ for all $X\in\mathfrak{g}$. Because $\omega_p^{-1}(\mathfrak{g})=T_p\mathcal{G}$ for every $p\in\mathcal{G}$, we are led to the following result.

Lemma 3.1. $\mathcal{L}_\xi\omega=0$ *if and only if* $[\xi,\omega^{-1}(X)]=0$ *for all* $X\in\mathfrak{g}$.

The following theorem is a straightforward elaboration on Theorem 1.5.11 on page 97 of [1], and essentially follows the same proof.

Theorem 3.2. *Suppose $(\mathcal{G},\omega)$ is a Cartan geometry of type $(\mathfrak{g},H)$. Then, for each $p\in\mathcal{G}$, $\mathfrak{aut}(\mathcal{G},\omega)$ is isomorphic to the subspace $\omega_p(\mathfrak{aut}(\mathcal{G},\omega))\subseteq\mathfrak{g}$ with bracket given by $(X,Y)\mapsto[X,Y]-\Omega_p(\omega^{-1}(X)\wedge\omega^{-1}(Y))$.*

Proof. Suppose $\xi\in\mathfrak{aut}(\mathcal{G},\omega)$. Then, $\mathcal{L}_\xi\omega=0$, so by Lemma 3.1, $[\xi,\omega^{-1}(X)]=0$ for all $X\in\mathfrak{g}$. This implies that their flows commute, so that

$$\xi_{\exp(t\omega^{-1}(X))p}=\exp(t\omega^{-1}(X))_*\xi_p$$

for all $t\in\mathbb{R}$ such that this is well-defined. Thus, ξ_p determines ξ on the connected component of p, and because we must also have $R_{h*}\xi=\xi$ for all $h\in H$, ξ_p determines ξ on all of $\mathcal{G}$. Thus, for each $p\in\mathcal{G}$, $\omega_p:\mathfrak{aut}(\mathcal{G},\omega)\to\mathfrak{g}$ is a vector space isomorphism onto its image.

Suppose $\xi,\nu\in\mathfrak{aut}(\mathcal{G},\omega)$. Recall from the computation above that

$$0=(\mathcal{L}_\xi\omega)(\nu)=\Omega(\xi\wedge\nu)-[\omega(\xi),\omega(\nu)]+\nu(\omega(\xi)).$$

[5]On the other hand, the isometry group of the Riemannian manifold $(M\times\mathbb{Z},\mathrm{g})$, where g is independent of $\mathbb{Z}$, has uncountably many connected components because it contains the permutation group of the integers as a subgroup of its component group.

Adding $\omega([\xi,\nu])$ to each side, this gives us

$$\begin{aligned}\omega([\xi,\nu]) &= \Omega(\xi\wedge\nu) - [\omega(\xi),\omega(\nu)] + (\nu(\omega(\xi)) - \omega([\nu,\xi]))\\ &= \Omega(\xi\wedge\nu) - [\omega(\xi),\omega(\nu)] + (\mathcal{L}_\nu\omega)(\xi)\\ &= \Omega(\xi\wedge\nu) - [\omega(\xi),\omega(\nu)].\end{aligned}$$

Since the bracket on $\mathfrak{aut}(\mathcal{G},\omega)$ is given by $(X,Y)\mapsto -[X,Y]$, the negative of the Lie bracket, this proves the result. □

Corollary 3.3. *The dimension of the Lie algebra $\mathfrak{aut}(\mathcal{G},\omega)$ is less than or equal to the dimension of $\mathfrak{g}$. In particular, $\mathfrak{aut}(\mathcal{G},\omega)$ is finite-dimensional and, by Proposition 2.9, $\mathrm{Aut}(\mathcal{G},\omega)$ is a Lie group with Lie algebra $\mathfrak{aut}(\mathcal{G},\omega)$.*

Corollary 3.4. *If there exists a point $p\in\mathcal{G}$ such that $\Omega_p = 0$, then $\mathfrak{aut}(\mathcal{G},\omega)$ is isomorphic to a subalgebra of $\mathfrak{g}$.*

In the above proof, we use Lemma 3.1 to tell us that, if $p\in\mathcal{G}$ and $\xi\in\mathfrak{aut}(\mathcal{G},\omega)$, then $\xi_{\exp(t\omega^{-1}(X))p} = \exp(t\omega^{-1}(X))_*\xi_p$ for sufficiently small $t\in\mathbb{R}$. This allows us to push ξ_p around the connected component of $p\in\mathcal{G}$ to determine ξ on that component. Additionally, as we noted in the proof, $R_{h*}\xi = \xi$ for all $h\in H$, so we can push ξ_p to the other components using H. Together, these determine ξ on all of $\mathcal{G}$ by its value at p.

Now, suppose we wish to compute $\mathfrak{aut}(\mathcal{G},\omega)$. Instead of solving a system of differential equations, we can instead pick elements of a given tangent space and just determine whether the above process determines complete vector fields from their initial values in that tangent space. This allows us to determine $\mathfrak{aut}(\mathcal{G},\omega)$ simply by knowing the flows of $\omega^{-1}(X)$ for each $X\in\mathfrak{g}$.

Example 3.5. Consider an ellipsoid given as the locus of $x^2+y^2+\frac{z^2}{4}=1$ and give it the Riemannian structure it inherits as a submanifold of Euclidean 3-space. Let $(\mathcal{G},\omega)$ be the Cartan geometry of type $(\mathfrak{i}(2), O(2))$ associated to that structure, where the Lie algebra of $O(2)$ is associated with the subalgebra of elements of the form $(0,X)\in\mathfrak{i}(2)$. Pick a point $p\in\mathcal{G}$ over the point $(0,0,2)$ on the ellipsoid. Then, the horizontal vectors at p do not determine well-defined vector fields using the pushing process described above, since their value at some other point q not in the same fiber as p will depend on the flows used to get from p to q. Tangent vectors pointing along the fiber at p, on the other hand, do determine a vector field; this should be visually clear. Thus, the symmetry algebra is a 1-dimensional Lie algebra, which we can identify with $\mathfrak{o}(2)$ over the point p.

4. Using covering maps

Another convenient way to compute the symmetry group of a Cartan geometry would be to determine it algebraically from another geometry. Using regular geometric covering maps, we can accomplish this with minimal effort.

Definition 4.1. A *(regular) geometric covering map* is a (regular)[6] covering map that is also a geometric map.

When looking at the symmetry group of a Cartan geometry and the symmetry group of another Cartan geometry that regularly covers it, it is visually clear that the symmetries of the cover should descend to symmetries on the geometry being covered, unless these symmetries do not make sense on the covered geometry. For example, in the Riemannian case, consider the standard covering of the flat torus by Euclidean space. Translations on the flat torus still make sense, but rotations,

[6] Recall that a regular covering map is a covering map π whose deck transformations, for each p, are transitive over $\pi^{-1}(p)$.

in general, do not; thus, the translations descend to symmetries on the flat torus, but not most rotations. The following result simply restates this visual intuition rigorously.

It should be noted that the proof of the result below is almost exactly the same as the proof for a weaker claim given on page 211 of [2].

Proposition 4.2. *Suppose* $\pi : (\mathcal{G}, \omega) \to (\mathcal{Q}, \eta)$ *is a regular geometric covering map with deck transformation group* $\Gamma \leq \mathrm{Aut}(\mathcal{G}, \omega)$*. Then,* $\mathrm{Aut}(\mathcal{Q}, \eta)$ *is isomorphic to* $N_{\mathrm{Aut}(\mathcal{G},\omega)}(\Gamma)/\Gamma$*, where* $N_{\mathrm{Aut}(\mathcal{G},\omega)}(\Gamma)$ *is the normalizer of* Γ *in* $\mathrm{Aut}(\mathcal{G}, \omega)$*.*

Proof. Consider the map

$$\begin{aligned} \rho\colon N_{\mathrm{Aut}(\mathcal{G},\omega)}(\Gamma) &\to \mathrm{Aut}(\mathcal{Q}, \eta), \\ \phi &\mapsto \pi \circ \phi \circ \pi^{-1}. \end{aligned}$$

Note that ρ is well-defined, since (by definition) $\phi(\Gamma p) = \Gamma\phi(p)$ for all $p \in \mathcal{G}$ whenever $\phi \in N_{\mathrm{Aut}(\mathcal{G},\omega)}(\Gamma)$. We wish to show that ρ is a surjective Lie group homomorphism with kernel Γ.

Suppose $\phi, \psi \in N_{\mathrm{Aut}(\mathcal{G},\omega)}(\Gamma)$. Then,

$$\begin{aligned} \rho(\phi) \circ \rho(\psi) &= (\pi \circ \phi \circ \pi^{-1}) \circ (\pi \circ \psi \circ \pi^{-1}) \\ &= \pi \circ \phi \circ \psi \circ \pi^{-1} \\ &= \rho(\phi \circ \psi). \end{aligned}$$

Thus, ρ is a Lie group homomorphism.

Let $\alpha \in \mathrm{Aut}(\mathcal{Q}, \eta)$ and let $\widehat{\alpha} \in \rho^{-1}(\alpha)$. Then, for $p \in \mathcal{G}$ and $\gamma \in \Gamma$,

$$\pi(\widehat{\alpha}(\gamma(p))) = \alpha(\pi(\gamma(p))) = \alpha(\pi(p)) = \pi(\widehat{\alpha}(p)).$$

Thus, there exists a $\gamma' \in \Gamma$ such that $\widehat{\alpha}(\gamma(p)) = \gamma'(\widehat{\alpha}(p))$. Since $\widehat{\alpha} \circ \gamma$ and $\gamma' \circ \widehat{\alpha}$ are both in $\rho^{-1}(\alpha)$ and $\widehat{\alpha} \circ \gamma(p) = \gamma' \circ \widehat{\alpha}(p)$, it follows that $\widehat{\alpha} \circ \gamma = \gamma' \circ \widehat{\alpha}$, so $\widehat{\alpha} \in N_{\mathrm{Aut}(\mathcal{G},\omega)}(\Gamma)$. Because $\rho(\widehat{\alpha}) = \alpha$ and α was arbitrary, this shows ρ is surjective.

Finally, we have $\ker(\rho) = \rho^{-1}(\mathrm{id}_{\mathcal{Q}}) = \pi^{-1} \circ \pi = \Gamma$, which completes the proof. □

Thus, if we wish to find the symmetry group of a Cartan geometry, then we can instead find the symmetry group of a geometry regularly covering it.

Example 4.3. Let us look more closely at the above example of the flat torus. Starting with the Klein geometry $(\mathfrak{I}(n), \omega_{\mathfrak{I}(n)})$ modeled on $(\mathfrak{I}(n), O(n))$. It has symmetry group $\mathfrak{I}(n)$, and we will quotient by the subgroup Γ of integer translations, which consists of elements of the form $(c, \mathbb{1})$, where $c \in \mathbb{Z}^n$. Clearly, elements of the form $(p, \mathbb{1})$ will commute with Γ, so they will be in the normalizer of Γ. Rotations and reflections are a bit more difficult. Since $(0, A)(c, \mathbb{1})(0, A)^{-1} = (Ac, \mathbb{1})$ must be in Γ for $(0, A)$ to be in the normalizer, we must have A take integer translations to integer translations. In other words, as a matrix, A must only have integer entries. Thus, the symmetry group of the flat torus will be $N_{\mathfrak{I}(n)}(\Gamma)/\Gamma$, which is isomorphic to $\mathbb{T}^n \rtimes O_n(\mathbb{Z})$.

5. Complete, flat Cartan geometries

Cartan geometries that are both complete and flat are very well-behaved; in fact, as the following result states, they are all geometrically covered by a Klein geometry.

It should be noted that a weaker version of the following result appears on page 213 of [6]. Surprisingly, our proof of this more general result is much less elaborate.

Proposition 5.1. *Suppose* $(\mathcal{G}, \omega)$ *is a complete flat Cartan geometry of type* $(\mathfrak{g}, H)$*. Then, for some Lie group* G *with Lie algebra* $\mathfrak{g}$*, there exists a regular geometric covering map* $\pi : (G, \omega_G) \to (\mathcal{G}, \omega)$*.*

Proof. Because $(\mathcal{G},\omega)$ is flat, $\omega^{-1}:\mathfrak{g}\to\Gamma(T\mathcal{G})$ is a Lie algebra isomorphism onto its image, and because $(\mathcal{G},\omega)$ is complete, the one-parameter subgroups of diffeomorphisms generated by $\omega^{-1}(\mathfrak{g})\approx\mathfrak{g}$ generate a connected Lie group G_{e}, with Lie algebra $\mathfrak{g}$, that acts on $\mathcal{G}$ by its inclusion in the diffeomorphism group.

Since, by definition of a Cartan geometry, $t\mapsto R_{\exp(tY)}p$ is the integral curve of $\omega^{-1}(Y)$ through p whenever $p\in\mathcal{G}$ and $Y\in\mathfrak{h}$, we have $H_{\mathrm{e}}<G_{\mathrm{e}}$. With this, we extend G_{e} to a Lie group G (with Lie algebra $\mathfrak{g}$) containing H and acting on $\mathcal{G}$ in the obvious way, with $h\in H$ acting by $h\cdot p=R_{h^{-1}}(p)$.

Note that G acts transitively on $\mathcal{G}$: the action of G_{e} is transitive on each connected component, essentially by the same standard argument used to show that a neighborhood of the identity in a topological group generates the identity component of that group, and the action of H is transitive on the fiber.

Fix $o\in\mathcal{G}$ and let $\pi:G\to\mathcal{G}$ be $g\mapsto g^{-1}\cdot o$. Then, π is a regular covering map of $\mathcal{G}$, and its group of deck transformations Γ is the isotropy subgroup of G at o. Moreover, since the one-parameter subgroups of $\omega^{-1}(\mathfrak{g})$ generated G_{e}, we know that ω_G will just be $\pi^*\omega$, and $\pi(gh)=h^{-1}g^{-1}\cdot o=R_h(g^{-1}\cdot o)=R_h\circ\pi(g)$ for all $h\in H$, where R_h is the right action of $h\in H$ on $\mathcal{G}$ as a principal H-bundle, so π is a geometric map as well. □

The proof of the following lemma is from Theorem 1.2.4 and Proposition 1.5.2(2) in [1].

Lemma 5.2. $\mathrm{Aut}(G,\omega_G)=G$.

Proof. Let $\mu:(a,b)\mapsto ab$ be the group operation for G and let $(\cdot)^{-1}:g\mapsto g^{-1}$ be the inverse operation. For all $g,h\in G$, $X\in T_gG$, and $Y\in T_hG$, we have

$$\mu^*\omega_G(X,Y)=\mathrm{Ad}_{h^{-1}}\omega_G(X)+\omega_G(Y)$$

and

$$(\cdot)^{-1*}\omega_G(X)=-\mathrm{Ad}_g\omega_G(X).$$

Suppose $f_1,f_2:G\to G$ are such that $f_1^*\omega_G=f_2^*\omega_G$.

Define $h:g\mapsto f_1(g)f_2(g)^{-1}$. We wish to show that f_1 and f_2 differ only by a constant. In other words, we wish to show that h is constant. It suffices to show that $h^*\omega_G=0$. We can write h out as

$$h=\mu\circ(\mathrm{id}_G,(\cdot)^{-1})\circ(f_1,f_2)\circ(g\mapsto(g,g)).$$

Thus,

$$\begin{aligned}h^*\omega_G(X)&=\mu^*\omega_G(f_{1*}X,(\cdot)^{-1}_*\circ f_{2*}X)\\&=\mathrm{Ad}_{(\cdot)^{-1}\circ f_2}(f_1^*\omega_G(X)+f_2^*(\cdot)^{-1*}\omega_G(X))\\&=\mathrm{Ad}_{(\cdot)^{-1}\circ f_2}(f_1^*\omega_G(X)-f_2^*\omega_G(X))\\&=0.\end{aligned}$$

Thus, every symmetry is a left multiplication by an element of G.

Conversely, left multiplication commutes with right multiplication, so every left multiplication by an element of G is a symmetry. □

By combining Proposition 5.1 with Lemma 5.2 and Proposition 4.2, we arrive at the following result.

Corollary 5.3. *If $(\mathcal{G},\omega)$ is a complete flat Cartan geometry of type $(\mathfrak{g},H)$, then its symmetry group is of the form $N_G(\Gamma)/\Gamma$, where G is a Lie group covering $\mathcal{G}$ with Lie algebra $\mathfrak{g}$ and with $H\leq G$.*

6. Homogeneous Cartan geometries

Often, homogeneous Riemannian manifolds are the nicest spaces on which to work. Their role in Riemannian geometry can be extended to general Cartan geometries.

The propositions below are from [4], which goes much further in depth than we will here. For convenience, we will omit the proofs.

Definition 6.1. A Cartan geometry $(\mathcal{G}, \omega)$ of type $(\mathfrak{g}, H)$ over M is *homogeneous* if and only if there exists a Lie group S acting transitively on M such that, for each $s \in S$, the action of s on M lifts to a symmetry of $(\mathcal{G}, \omega)$.

Let K be a Lie subgroup of S, and let $\psi : K \to H$ be a Lie group homomorphism. Then, we will denote by $S \times_K H$ the principal H-bundle given by the quotient of $S \times H$ by $k(s, h) = (sk^{-1}, \psi(k)h)$ over S/K.

Proposition 6.2 (Theorem 2.2.6(i) in [4])**.** *Suppose $\mathcal{G}$ is a principal H-bundle over M and S is a Lie group that acts transitively on M such that, for each $s \in S$, the action of s on M lifts to a right H-equivariant map on $\mathcal{G}$. Further, suppose K is the isotropy subgroup of S for some point $o \in \mathcal{G}$. Then, there exists a unique Lie group homomorphism $\psi : K \to H$ such that $\mathcal{G}$ is isomorphic as a principal H-bundle to $S \times_K H$.*

The above result shows that we might as well write our principal H-bundle $\mathcal{G}$ as $S \times_K H$ for some $\psi : K \to H$, since all homogeneous Cartan geometries will take this form.

We can even get an explicit expression for the Cartan connection of a homogeneous Cartan geometry.

Proposition 6.3 (Theorem 4.2.1 in [4])**.** *Let $(S \times_K H, \omega)$ be a homogeneous Cartan geometry of type $(\mathfrak{g}, H)$. Then, there exists a unique linear map $\alpha : \mathfrak{s} \to \mathfrak{g}$ descending to an isomorphism from $\mathfrak{s}/\mathfrak{k}$ to $\mathfrak{g}/\mathfrak{h}$ such that $\alpha|_{\mathfrak{k}} = \psi_{*e}$ and $\alpha \circ \mathrm{Ad}_k = \mathrm{Ad}_{\psi(k)} \circ \alpha$ for every $k \in K$ such that ω lifts to*

$$(s, h) \mapsto ((X_S, X_H) \mapsto \mathrm{Ad}_{h^{-1}} \circ \alpha(\omega_S(X_S)) + \omega_H(X_H))$$

on $S \times H$.

Homogeneous Cartan geometries conveniently take care of the "horizontal" symmetries, since they are determined as part of the definition, so all that is left is to find the symmetries directed along the fibers. From the above information, we can determine a few details regarding the symmetry algebra.

Suppose S acts faithfully. Let $o = (\mathrm{e}_S, \mathrm{e}_H) \in S \times_K H$. By Theorem 3.2, $\mathfrak{aut}(\mathcal{G}, \omega)$ is given by its image in $\mathfrak{g}$ under the bracket $(X, Y) \mapsto [X, Y] - \Omega_o(\omega^{-1}(X) \wedge \omega^{-1}(Y))$. For $\alpha(X), \alpha(Y) \in \alpha(\mathfrak{s})$, this becomes $[\alpha(X), \alpha(Y)] - ([\alpha(X), \alpha(Y)] - \alpha([X, Y])) = \alpha([X, Y])$. Since Ω is horizontal, if either argument of the bracket is in $\mathfrak{h}$, then the bracket is just the bracket on $\mathfrak{g}$. Additionally, since S acts transitively on the base manifold, searching for other infinitesimal symmetries using the pushing method from above can be restricted to elements of $\omega_o^{-1}(\mathfrak{h})$.

It was originally the intention of the author to develop, in much more detail than the feebler results of the above paragraph, an easy way to completely determine the symmetry group of a homogeneous Cartan geometry. To see why there is hope for results in this direction, note that, for any Cartan geometry $(\mathcal{G}, \omega)$, $\mathrm{Aut}(\mathcal{G}, \omega)$ will always act freely on $\mathcal{G}$, so we can embed $\mathrm{Aut}(\mathcal{G}, \omega)$ into $\mathcal{G}$ using the map $\phi \mapsto \phi(p)$ for some fixed $p \in \mathcal{G}$. Here, we are concerned with a homogeneous Cartan geometry $(S \times_K H, \omega)$, and $\mathrm{Aut}(S \times_K H, \omega)$ embeds into $S \times_K H$ using $\phi \mapsto \phi(o)$. This embedding will contain the orbit of S on o, so finding any other symmetries restricts to looking along the fiber. Informally, this means that the symmetries that have no "S-component" embed into H. The author had wished to find a way to link the algebraic structure of these other symmetries to H, but this was, perhaps, too ambitious.

Acknowledgements. The author would like to thank Joel Specter and Subhadip Chowdhury for supervising this project. He would also like to thank Professor J. Peter May for his helpful commentary on the paper.

References

[1] A. Čap and J. Slovák, *Parabolic Geometries I: Background and General Theory*, Mathematical Surveys and Monographs, vol. 154, American Mathematical Society, 2009.

[2] L. S. Charlap, *Bieberbach Groups and Flat Manifolds*, Springer-Verlag, 1986.

[3] J. W. Erickson, *Some elementary results regarding reductive Cartan geometries* (2016), available at `arXiv:1607.01451`.

[4] M. Hammerl, *Homogeneous Cartan geometries*, diploma thesis, University of Vienna, 2006, `http://www.mat.univie.ac.at/~cap/theses.html`.

[5] R. S. Palais, *A global formulation of the Lie theory of Transformation Groups*, Memoirs of the American Mathematical Society **22** (1957).

[6] R. W. Sharpe, *Differential Geometry: Cartan's Generalization of Klein's Erlangen Program*, Graduate Texts in Mathematics, vol. 166, Springer-Verlag New York, 1997.