1 2 9 0

UNIVERSIDADE D
COIMBRA

Christopher Neil Gadzinski

From Coordinate Algebras to the Pontryagin Maximum Principle

Dissertação no âmbito do Mestrado em Matemática Pura, orientada pela Professora Doutora Margarida Camarinha e apresentada ao Departamento de Matemática da Faculdade de Ciências e Tecnologia.

Novembro de 2021

From Coordinate Algebras to the Pontryagin Maximum Principle

Christopher Neil Gadzinski

1 2 9 0

UNIVERSIDADE D COIMBRA

Master in Mathematics

Mestrado em Matemática

MSc Dissertation | Dissertação de Mestrado

November 2021

Acknowledgements

I'd like to thank my thesis advisor, Margarida Camarinha, for helping me stay pointed in a good mathematical direction over this past year.

I'd also like to thank my mom for helping me stay healthy, and my dad for helping me start my education in mathematics.

Abstract

In the preface to *Introduction to Topology and Modern Analysis* [18], George Simmons writes:

> It seems to me that a worthwhile distinction can be drawn between two types of pure mathematics. The first—which unfortunately is somewhat out of style at present—centers attention on particular functions and theorems which are rich in meaning and history, like the gamma function and the prime number theorem, or on juicy individual facts, like Euler's wonderful formula
>
> $$1+\frac{1}{4}+\frac{1}{9}+\ldots=\frac{\pi^2}{6}.$$
>
> The second is concerned primarily with form and structure.

"Pure mathematics of the second type" will be our occupation in the first two chapters of this thesis. In Chapter 1, we will explain how smooth manifolds can be understood in terms of their *coordinate algebras*. This perspective is especially helpful to study the relationship between vector fields and their flows. In Chapter 2, we will borrow some motivation from classical mechanics and describe an additional algebraic operation (the *Poisson bracket*) equipped on the coordinate algebra of a cotangent bundle.

Our third and final chapter, however, is dedicated to an instance of "pure mathematics of the first type." Relying on the theoretical framework developed in the first two chapters, we will explain the *Pontryagin maximum principle* (PMP) for time-optimal trajectories of a control system.

The Pontryagin maximum principle, first published in 1956 [15], is today a fundamental result in optimal control theory. (Some references applying the PMP in recent research are [6], [3], and [7].) Our main reference in studying the PMP was the book [2], but effort has been devoted in this thesis to present the subject from an original perspective whenever possible.

Resumo

No prefácio do livro *Introduction to Topology and Modern Analysis* [18], George Simmons afirma:

> Parece-me que se pode fazer uma distinção relevante entre dois tipos de matemática pura. A primeira—que infelizmente parece estar um pouco fora de moda—interessa-se por funções e teoremas que são ricos em significado e interesse histórico, como a função gama e o teorema dos números primos, ou em factos isolados de interesse substancial, como a fórmula de Euler:
> $$1+\frac{1}{4}+\frac{1}{9}+\ldots=\frac{\pi^2}{6}.$$
> A segunda está interessada primariamente em forma e estrutura.

Será na "matemática pura do segundo tipo" que se centrará o nosso esforço nos primeiros 2 capítulos desta tese. No 1.º capítulo explicaremos como compreender variedades diferenciáveis em termos das suas álgebras de coordenadas. Esta perspetiva harmoniza-se com o estudo da relação entre campos de vetores e os seus fluxos. Encontramos motivação para o 2.º capítulo na mecânica clássica e aí descreveremos uma outra operação algébrica (os parênteses de Poisson) associada à álgebra de coordenadas de um fibrado cotangente. Contudo, o nosso terceiro e último capítulo é dedicado a um exemplo de "matemática pura do 1.º tipo". Usando a teoria desenvolvida nos 1.º e 2.º capítulos, explicaremos o *princípio do máximo de Pontryagin* (PMP) para trajetórias ótimas de um sistema de controlo.

O PMP, publicado originalmente em 1956 [15], é presentemente um resultado fundamental em teoria do controlo ótimo. (Algumas referências de aplicação do PMP em estudos recentes são [6], [3] e [7].) A nossa referência fundamental para o estudo do PMP foi o livro [2]. Contudo, sempre que possível, dedicámos um esforço a apresentar este tema de uma perspetiva original.

Table of Contents

Glossary of Notation

Diff	the category of smooth manifolds
Diff^*	the category of diffeomorphisms between smooth manifolds
Alg	the category of unital, commutative real algebras
PolySpc	the category of polynomial maps between Euclidean spaces
CartSpc	the category of smooth maps between Euclidean spaces
NGrp	the category of diffeomorphisms between real domains
$\mathrm{Hom}_{\mathscr{C}}(A,B)$	the set of morphisms from A to B in $\mathscr{C}$
$\mathrm{End}_{\mathscr{C}}(A)$	the set of morphisms from A to A in $\mathscr{C}$
$\mathrm{Aut}_{\mathscr{C}}(A)$	the set of isomorphisms from A to A in $\mathscr{C}$
M,N	smooth manifolds
Q,P	diffeomorphisms
X,Y,Z	vector fields
p,q	points
T	the tangent bundle functor
T^*	the cotangent bundle functor
C^∞	the coordinate algebra functor
$\mathrm{Aut}(M)$	the group of diffeomorphisms of M
$D(M)$	the Lie algebra of vector fields on M
$D^*(M)$	the Lie algebra of complete vector fields on M
$\{-,-\}$	the Poisson bracket on T^*M
$\overrightarrow{H}$	the Hamiltonian vector field $\{H,-\}$ on T^*M
H_X	the Hamiltonian function of a vector field $X \in D(M)$
$\overline{X}$	the Hamiltonian lift of a vector field, equal to $\overrightarrow{H}_X$
e^-	the exponential function for Lie groups, matrices, and scalars
$\exp(-)$	the exponential function for vector fields
$\overrightarrow{\exp}$, $\overleftarrow{\exp}$	the right and left chronological exponentials for vector fields
$\mathrm{Ad}\,Q$	the adjoint action of a diffeomorphism
$\mathrm{ad}\,X$	the adjoint action of a vector field

Chapter 1

Coordinate Algebras

When M is a smooth manifold, let $C^\infty M$ denote the space of smooth functions[1] from M to $\mathbb{R}$. We equip $C^\infty M$ with the natural operations of addition and multiplication and call this unital, commutative, real algebra the **coordinate algebra** of M.

It is standard to define a vector field on a manifold as a linear operator $X : C^\infty M \to C^\infty M$ verifying the “Leibniz rule"

$$X(fg) = X(f)g + fX(g) \tag{1.1}$$

for all $f, g \in C^\infty M$. In general, when A is a (unital, commutative, real) algebra, we define the **derivations** $D(A)$ of A as the vector space endomorphisms of A that respect Equation (1.1) for all $f, g \in A$. Let us take a moment to examine the significance of $D(A)$ when A is *finite-dimensional*.

In this case, we may think of $\mathrm{GL}(A)$, the group of vector space automorphisms of A, as a finite-dimensional Lie group. In the usual way, we think of its Lie algebra $\mathfrak{gl}(A)$ as the space of vector space endomorphisms of A. We also know that the set $\mathrm{Aut}(A)$ of automorphisms of A is a Lie subgroup of $\mathrm{GL}(A)$. Now, let $X \in \mathfrak{gl}(A)$, and suppose that e^{tX} is an automorphism of A for all t. Then, differentiating the homomorphism equation $e^{tX}(fg) = (e^{tX}f)(e^{tX}g)$ in the real parameter t gives

$$Xe^{tX}(fg) = (Xe^{tX}f)(e^{tX}g) + (e^{tX}f)(Xe^{tX}g).$$

With $t = 0$, this proves that X is a derivation of A. In fact, the converse is true: X is a derivation if and only if e^{tX} is an automorphism for all t. This means the following.

Theorem 1. *The Lie algebra of* $\mathrm{Aut}(A) \subseteq \mathrm{GL}(A)$ *is* $D(A)$*, the set of derivations of* A.

Proof. Recall that, when $G \subseteq H$ is an inclusion of Lie groups and $\mathfrak{g} \subseteq \mathfrak{h}$ is the associated inclusion of Lie algebras, an element $X \in \mathfrak{h}$ belongs to $\mathfrak{g}$ if and only if e^{tX} is in G for all $t \in \mathbb{R}$. We have already proven that X is a derivation if e^{tX} is an automorphism of A for all t, so it only remains to prove the converse.

Suppose X is a derivation of A, and let $f, g \in A$ be arbitrary. Our claim is that the equation

$$e^{tX}(fg) - (e^{tX}f)(e^{tX}g) = C_t = 0$$

[1]Throughout this thesis, “smooth” means infinitely differentiable.

holds for all $t \in \mathbb{R}$. Clearly, $C_0 = 0$. Now, we differentiate:

$$\begin{aligned}\dot{C}_t &= Xe^{tX}(fg) - (Xe^{tX}f)(e^{tX}g) - (e^{tX}f)(Xe^{tX}g) \\ &= X(e^{tX}(fg)) - X((e^{tX}f)(e^{tX}g)) = XC_t.\end{aligned}$$

We conclude that C_t solves the Cauchy problem

$$\begin{cases} C_0 = 0 \\ \dot{C}_t = XC_t, \end{cases}$$

and so must be constantly 0. □

At least in the case of finite-dimensional algebras, this shows that derivations are the *infinitesimal generators of automorphisms*. Perhaps vector fields—derivations of the coordinate algebra $C^\infty M$—have a similar relationship with automorphisms of $C^\infty M$. But, what is an automorphism of $C^\infty M$?

The answer, as we will show momentarily, is that automorphisms of the algebra $C^\infty M$ will encode *self-diffeomorphisms* of the manifold M. Indeed, besides tangent vectors and vector fields, also points, smooth maps, and manifolds themselves can be considered from an "algebraic point of view." This point of view is especially useful to clarify the relationship between vector fields and their flows; for instance, it will let us define a series formula for the flow of a non-autonomous (time-dependent) vector field. Such applications of the coordinate algebra formalism, called *chronological calculus*, were introduced by Agrachev et. al. in [1] and detailed in their book [2] on control theory. The relationship between homomorphisms and derivations of an algebra that we have hinted at here will also be helpful to keep in mind when, in the next chapter, we equip the cotangent bundle T^*M with an additional operation (the *Poisson bracket*) and study the derivations and automorphisms associated with this additional structure.

1.1 Why Coordinate Algebras?

Moments ago, we promised to show that algebra automorphisms of $C^\infty M$ are equivalent to diffeomorphisms of M. Actually, more is true: algebra homomorphisms from $C^\infty N$ to $C^\infty M$ are equivalent to smooth maps from M to N. This remarkable correspondence takes the form of a contravariant functor. (We will find a few occasions in this thesis to be inspired by category theory. For a reference on the subject, see [13].)

In the following, Diff is the category of smooth manifolds and Alg is the category of commutative, unital algebras over $\mathbb{R}$.

Definition 1. *The* ***coordinate algebra functor*** *is the contravariant functor*

$$\text{Diff} \xrightarrow{C^\infty} \text{Alg}$$

which takes a manifold M to its coordinate algebra $C^\infty M$ and a smooth map $\varphi\colon M \to N$ to the algebra homomorphism

$$\hat{\varphi}\colon C^\infty N \to C^\infty M$$
$$\hat{\varphi}(f) = f \circ \varphi.$$

Note that the map $\hat{\varphi}$ defined in this way clearly defines an algebra homomorphism; for example,

$$\hat{\varphi}(fg) = (fg) \circ \varphi = (f \circ \varphi)(f \circ \varphi) = \hat{\varphi}(f)\hat{\varphi}(g).$$

Functoriality of this assignment is also easy to prove. Our promised fact can now be stated in the following way.

Proposition 1. *The coordinate algebra functor is fully faithful.*

The core of this proposition is contained in the following weaker statement. (In [12], it is referred to as *Milnor's exercise*.)

Lemma 1 (Milnor's exercise)**.** *Let M be a manifold and let P be a manifold with one point. Then the coordinate algebra functor gives a bijection*

$$\mathrm{Hom}_{\mathrm{Diff}}(P, M) \cong \mathrm{Hom}_{\mathrm{Alg}}(C^\infty M, C^\infty P).$$

Of course, maps (smooth or otherwise) from P to M are in bijection with points of M. In what follows, we will regularly interpret points as inclusions of points and vice versa. Also, note that $C^\infty P$ is isomorphic to $\mathbb{R}$ with its usual algebra structure.

Our solution to Milnor's exercise is valid under the assumption that M is compact. For the more general case, we will apply a technical result whose proof we do not discuss. (See Chapter VIII of [12] for more details.)

Proof. Let $\varphi\colon C^\infty M \to C^\infty P$ be a homomorphism. For each element $f \in \ker\varphi$, let $Z_f \subseteq M$ be the zero set $f^{-1}(0)$. We will argue that there is at least one point x_0 common to all these zero sets. Then, it will follow that $\varphi(f) = f(x_0)$ for all $f \in C^\infty M$. Indeed, where 1 is the unit, $f - \varphi(f)1$ always belongs to the kernel of φ, so $x_0 \in Z_{f-\varphi(f)1}$, which forces $\varphi(f) = f(x_0)$. When x_0 is interpreted as an inclusion of P into M, this means that $\varphi = \hat{x}_0$, so we conclude that the assignment made by the coordinate algebra functor from $\mathrm{Hom}_{\mathrm{Diff}}(P, M)$ to $\mathrm{Hom}_{\mathrm{Alg}}(C^\infty M, C^\infty P)$ is surjective. Of course it is also injective; when p and q are different points, $f(p) \neq f(q)$ for some adequately chosen function $f \in C^\infty M$, so $\hat{p} \neq \hat{q}$.

We proceed to prove the existence of x_0. Our first observation is that, because the homomorphism φ preserves the unit element, it must also preserve inverses when they exist. Therefore, a non-vanishing element f of $C^\infty M$ cannot be in the kernel of φ, and so no zero set Z_f is empty. Our second observation is that the family of zero sets is closed under finite intersections. Indeed, given any $f, g \in \ker\varphi$, we also have $f^2 + g^2 \in \ker\varphi$, and

$$Z_f \cap Z_g = Z_{f^2+g^2}.$$

From these two observations, it follows that no finite family of zero sets can have empty intersection.

Suppose, nevertheless, that the whole family of zero sets had finite intersection. To arrive at a contradiction, we appeal to compactness. Suppose there is some $f_0 \in \ker\varphi$ for which Z_{f_0} is compact. Then Z_{f_0} is covered by the family of complements of zero sets, $\{Z_f^C : f \in \ker\varphi\}$. By compactness, there exists a finite set of functions $\{f_1,\dots,f_n\}$ for which the sets $Z_{f_i}^C$ cover Z_{f_0}. This means that the intersection $Z_{f_0} \cap Z_{f_1} \cap \dots \cap Z_{f_n}$ is empty, which is a contradiction.

The existence of a compact zero set Z_{f_0} is obviously true when M is compact. In general, it is possible to construct a function $g \in C^\infty M$ that is unbounded on every non-compact closed set. Such a function clearly has a compact zero set. To find an element in the kernel of φ with a compact zero set, it is enough to take $f_0 = g - \varphi(g)1$. □

The more general statement of Proposition 1 now follows.

Proof of Proposition 1. Let $\varphi, \psi \colon M \to N$ be smooth maps between manifolds and let $p \in M$ be a point. Assume $\hat{\varphi} = \hat{\psi}$. Then $\hat{p}\hat{\varphi} = \hat{p}\hat{\psi}$, and $\widehat{\varphi(p)} = \widehat{\psi(p)}$. By Lemma 1, we conclude that $\varphi(p) = \psi(p)$. This proves that the coordinate algebra functor is faithful.

Next, we show it is full. For any given homomorphism $h\colon C^\infty N \to C^\infty M$ and point $p \in M$, the composition $\hat{p}h$ is an element of $\mathrm{Hom}(C^\infty N, C^\infty P)$, and so Lemma 1 guarantees the existence of a unique point q for which $\hat{q} = \hat{p}h$. We define a (not necessarily smooth!) map $\varphi\colon M \to N$ by the equation

$$\widehat{\varphi(p)} = \hat{p}h.$$

To finish the proof, we will show that φ is smooth in a neighborhood of any point $p_0 \in M$. Of course, it will follow that $\hat{\varphi} = h$.

Let U be an open neighborhood of $\varphi(p_0) \in N$ equipped with a coordinate chart

$$\mathbf{x} = (x_1, \dots, x_n)\colon U \to \mathbb{R}^n.$$

By standard "bump function" techniques, we can find a compact neighborhood $K \subseteq U$ of $\varphi(p_0)$ and a bump function $b\colon N \to \mathbb{R}$ so that $b \equiv 1$ over K but $b \equiv 0$ outside of U. We extend the chart $\mathbf{x}$ to a smooth map $\mathbf{x}' = (x'_1, \dots, x'_n)\colon N \to \mathbb{R}$ by setting

$$x'_i(p) = \begin{cases} b(p)x_i(p) & : x \in U \\ 0 & : x \notin U. \end{cases}$$

By construction, $\mathbf{x}' = \mathbf{x}$ over K.

First of all, note that the compositions $x'_i \circ \varphi$ are smooth. Indeed, for a given point $p \in M$,

$$(x'_i \circ \varphi)(p) = \widehat{\varphi(p)}(x'_i) = (\hat{p}h)(x'_i) = (h(x'_i))(p),$$

and by hypothesis h takes values in $C^\infty M$. This means that $\mathbf{x}' \circ \varphi$ is smooth. It remains only to prove that, within some neighborhood V of p_0, φ takes values in K; then, we will know that $\mathbf{x} \circ \varphi|_V$ is a well-defined smooth map, meaning that φ is a smooth map near p_0.

This last fact follows from a similar technique. Consider another bump function $c\colon N \to \mathbb{R}$ that takes $c(\varphi(p_0)) = 1$ but satisfies $c \equiv 0$ outside of K. By the above argument, $c \circ \varphi$ is smooth, and so the set

$$V = \{p \in M : (c \circ \varphi)(p) > 0\} \subseteq M$$

is an open neighborhood of p_0. By construction of c, every element of V is sent to within K by φ, so this concludes the proof. □

A fully faithful functor is essentially injective on objects, so $C^\infty M$ determines M up to isomorphism.[2] Indeed, given an algebra A which we are told is the coordinate algebra of a manifold, we can define a set $M = \mathrm{Hom}(A, \mathbb{R})$ and say a map $\ell\colon M \to \mathbb{R}$ is smooth when it can be written as $\ell(\varphi) = \varphi(a)$ for some element $a \in A$. If A is the coordinate algebra of a manifold, this construction recovers the underlying manifold up to its smooth structure! Now that this equivalence is proven, we will frequently reinterpret smooth maps as algebra homomorphisms and vice versa without needing the hat notation.

At first glance, it may be quite surprising that manifolds can be encoded by algebraic structures in this way. However, category theory gives us an intuitive reason to expect Proposition 1, which we will now explain.

When $\mathscr{C}$ is an arbitrary category, the Yoneda lemma tells us that the contravariant functor

$$\begin{aligned} h^- &: \mathscr{C} \to \mathscr{C}^{\mathrm{Set}} \\ M &\mapsto h^M = \mathrm{Hom}(M, -) \end{aligned}$$

is fully faithful. Putting $\mathscr{C} = \mathrm{Diff}$, this tells us that M is determined up to isomorphism by the covariant functor

$$\mathrm{Hom}(M, -)\colon \mathrm{Diff} \to \mathrm{Set}$$

that records the smooth maps out of M. Now, let PolySpc be the subcategory of Diff whose objects are the spaces $\mathbb{R}^n$ for natural numbers $n \in \{0, 1, \ldots\}$ and whose morphisms are the polynomial functions between these spaces. We define the PolySpc-**algebra of a manifold** as follows.

Definition 2. *When M is a manifold, the* PolySpc-***algebra of*** *M is*

$$C^\infty(M) = \mathrm{Hom}_{\mathrm{PolySpc}}(M, -),$$

the restriction of $\mathrm{Hom}_{\mathrm{Diff}}(M, -)$ *to a functor from* PolySpc *to* Set*.*

As a Hom-functor, the PolySpc-algebra of a manifold preserves limits, and, in particular, products. It turns out to be natural to define the category of "abstract" PolySpc-algebras in the following way.

Definition 3. *A* PolySpc-*algebra is a product-preserving functor*

$$\mathrm{PolySpc} \longrightarrow \mathrm{Set},$$

and a morphism between PolySpc-*algebras is a natural transformation.*

[2] This is meant in the categorical sense: if $C^\infty M \cong C^\infty N$ for some other manifold N, then there is an isomorphism $M \cong N$.

Actually, PolySpc-algebras have a familiar interpretation.

Proposition 2. *The category of* PolySpc*-algebras is equivalent to* Alg*, and, under this equivalence, the coordinate algebra functor is naturally isomorphic to* $\mathrm{Hom}_{\mathrm{PolySpc}}(M,-)$.

A proof in full detail is not productive to include in this thesis, but once the basic idea is grasped it is straightforward. We will merely explain the way that a PolySpc-algebra originates an algebraic structure, and why it happens, in this algebraic structure, that the product distributes over the sum.

Proof hint for Proposition 2. Let $F\colon$ PolySpc $\to$ Set be a PolySpc-algebra. The set of elements for our commutative algebra will be $A = F(\mathbb{R})$. Since F is product-preserving, $F(\mathbb{R}^0)$ and $F(\mathbb{R}\times\mathbb{R})$ can be viewed as the one-point set 1 and the product $A\times A$ respectively. Let $p\colon \mathbb{R}^2 \to \mathbb{R}$ take $p(x,y) = xy$, let $s\colon \mathbb{R}^2 \to \mathbb{R}$ take $s(x,y) = x+y$, and, for every $\alpha \in \mathbb{R}$, let $u_\alpha\colon \mathbb{R}^0 \to \mathbb{R}$ take $u_\alpha(-) = \alpha$. We claim that the maps

$$\begin{cases} Fp\colon A\times A \to A \\ Fs\colon A\times A \to A \\ Fu_\alpha\colon 1 \to A \end{cases}$$

endow A with, respectively, a product, sum, and inclusion of scalars.

As an example, let us prove that the product Fp distributes over the sum Fs. Consider the following diagram over PolySpc.

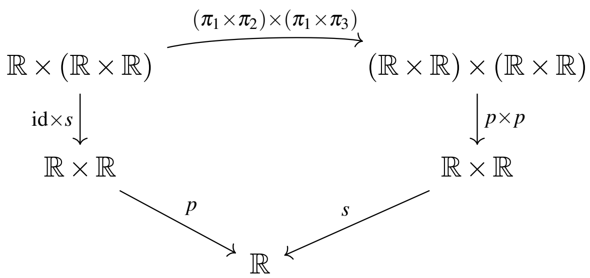

(The maps π_1, π_2 and π_3 are the three projections of the product $A\times A\times A$, and parenthesizations of products in the top row annotate the construction of the maps $\mathrm{id}\times s$ and $p\times p$, wherein $(-\times-)$ is applied as a bifunctor.) It is straightforward to check that this diagram commutes: an element $(a,b,c)\in\mathbb{R}^3$ is carried to $a(b+c)$ along the lower path, and carried to $ab+ac$ along the upper path. Because F is product-preserving, applying F takes the diagram above to an analogous commutative diagram in Set whose commutativity condition reads

$$(Fp)(a,(Fs)(b,c)) = (Fs)((Fp)(a,b),(Fp)(a,c)).$$

□

In summary, the coordinate algebra of a manifold M can be viewed as the restriction of the usual Yoneda embedding of M to PolySpc. The conclusion of Proposition 1 is that this restriction is not so severe as to break the normal conclusion of the Yoneda lemma.

It should be noted that the coordinate algebra $C^\infty M$ admits natural operations besides addition and multiplication. Indeed, every smooth map $g\colon \mathbb{R}^n \to \mathbb{R}$ gives us an operation on $C^\infty M$, namely

$$\begin{aligned} &C^\infty_g\colon (C^\infty M)^n \to C^\infty M \\ &C^\infty_g(f_1,\ldots,f_n) = (p \mapsto g(f_1(p),\ldots,f_n(p))). \end{aligned}$$

An algebra A equipped with an operation A_g for every smooth map $g\colon \mathbb{R}^n \to \mathbb{R}$, subject to certain natural axioms, is called a C^∞**-algebra**. The easiest way to describe a C^∞-algebra in rigor is as a product-preserving functor from CartSpc to Set, where CartSpc is the *full* subcategory of Diff generated by the spaces $\{\mathbb{R}^0,\mathbb{R}^1,\ldots\}$. The natural C^∞-algebra structure on $C^\infty M$ is equivalent to the restriction of the Yoneda embedding of M to CartSpc.

At this point, one may ask whether the coordinate algebra functor is essentially surjective on objects. This turns out to not be the case. For a counter-example, consider the quotient algebra

$$\mathbb{D} = \mathbb{R}[x]/\langle x^2\rangle.$$

A homomorphism from $C^\infty M$ to $\mathbb{D}$ is the same as a *tangent vector* of M. Indeed, where $\varphi\colon C^\infty M \to \mathbb{D}$ is given by $\varphi(f) = \varphi_0(f) + \varphi_1(f)x$ for real-valued functions φ_0 and φ_1, writing the homomorphism law $\varphi(fg) = \varphi(f)\varphi(g)$ in coordinates gives

$$\begin{cases} \varphi_0(fg) = \varphi_0(f)\varphi_0(g) \\ \varphi_1(fg) = \varphi_0(f)\varphi_1(g) + \varphi_1(f)\varphi_0(g). \end{cases}$$

This means that φ_1 satisfies the classical definition of a tangent vector at the point φ_0. We get the idea that the geometric structure associated with $\mathbb{D}$ is something like an interval of infinitesimal length![3] Although we will not prove this here, it can be shown that $\mathbb{D}$ is not the coordinate algebra of any manifold. Pursuing these questions further would lead us to consider certain *smooth toposes* which generalize the category Diff. For more details on these topics, see [14]. In this work we will deal only with classical manifolds; the coordinate algebra perspective will merely be a way to organize our calculations about them. We will also not find any occasion to invoke the C^∞-algebra structure of a coordinate algebra. However, knowing about C^∞-algebras makes the term "coordinate algebra" more natural; $C^\infty M$ is generated as a C^∞-algebra by some sort of "coordinate functions" of the sort that we have constructed in the proof of Proposition 1.[4]

Now, let us begin to explain some applications of the coordinate algebra framework.

Before a student of differential geometry learns the "classical definition" of tangent vectors as functionals on algebras of germs, they probably have the naive expectation that a tangent vector should be defined as a "velocity vector" taken by a trajectory of points. Unfortunately, the classical limit-based definition for "velocity vectors" only applies to curves in a topological vector space, and there appears to be no natural way to embed an arbitrary manifold into a TVS. Actually, this is not

[3] If we substitute the ideal $\langle x^2\rangle$ with $\langle x^k\rangle$ for some $k > 1$, a morphism from $C^\infty M$ to $\mathbb{D}$ becomes an object known as a "$(k-1)$-jet," which is something like an equivalence class of curves modulo their higher derivatives up to order $k-1$ at a given point.

[4] The term "coordinate ring" is quite standard in algebraic geometry, where this denomination is a little easier to defend.

true! We will see now that the representation of points as maps from $C^\infty M$ to $\mathbb{R}$ can be regarded as an embedding of M into a certain Fréchet space. More generally, we will be able to encode any smooth map φ as an element in a Fréchet space, and "tangent vectors to φ" will turn out to encode sections of the tangent bundle along φ.

Let us recall some important regularity properties for trajectories in a topological vector space.

Definition 4. *Let V be a TVS, let $I \subseteq \mathbb{R}$ be a compact interval, and let*

$$\begin{aligned} v\colon I &\to V \\ t &\mapsto v_t \end{aligned}$$

be a trajectory. In the following, let U quantify over all open neighborhoods of the origin in V. We say that v is:

1. ***Measurable*** *when, for all $w \in V$, $v^{-1}(w+U)$ is a measurable set.*
2. ***Locally bounded*** *when, for each $t_0 \in I$, there exists some real non-negative constants ε and C so that $v_s \in CU$ whenever $|s-t_0| < \varepsilon$.*
3. ***Lipschitz*** *when there exists a constant C so that, for every $t,s \in I$, there holds $v_t - v_s \in C|t-s|U$.*
4. ***Differentiable*** *at t when the limit*
$$\lim_{\varepsilon\to 0} \frac{v_{t+\varepsilon}-v_t}{\varepsilon}$$
converges in V.

In the following, all families are implicitly required to be measurable and locally bounded.

To perform computations with differentiable families, we'll frequently use the following generalized "product rule."

Proposition 3 (Product rule)**.** *Let X, Y and Z be topological vector spaces, let x and y be trajectories over some interval I into the spaces X and Y respectively, and let $P\colon X \times Y \to Z$ be a continuous bilinear map. If x and y are differentiable at $t_0 \in I$, then $z_t = P(x_t, y_t)$ is also differentiable at t_0, with derivative*

$$\dot{x}_{t_0} = P(\dot{x}_{t_0}, y_{t_0}) + P(x_{t_0}, \dot{y}_{t_0}).$$

Proof. The proof proceeds as it does in single-variable calculus:

$$\begin{aligned} \lim_{\varepsilon\to 0} \frac{z_{t_0+\varepsilon}-z_{t_0}}{\varepsilon} &= \lim_{\varepsilon\to 0} \frac{P(x_{t_0+\varepsilon}-x_{t_0}, y_{t_0+\varepsilon}) + P(x_{t_0}, y_{t_0+\varepsilon}-y_{t_0})}{\varepsilon} \\ &= \lim_{\varepsilon\to 0} P\left(\frac{x_{t_0+\varepsilon}-x_{t_0}}{\varepsilon}, y_{t_0+\varepsilon}\right) + \lim_{\varepsilon\to 0} P\left(x_{t_0}, \frac{y_{t_0+\varepsilon}-y_{t_0}}{\varepsilon}\right) \\ &= P(\dot{x}_{t_0}, y_{t_0}) + P(x_{t_0}, \dot{y}_{t_0}), \end{aligned}$$

where for the last equality we have used continuity of P. □

Now, we define a topology on the coordinate algebra of a manifold.

Definition 5. *Let K be a compact subset of M endowed with coordinates*

$$\mathbf{x} = (x_1, \ldots, x_n) \colon K \to \mathbb{R}^n,$$

and let $s \in \mathbb{N} = \{0, 1, \ldots\}$. The ***$C^\infty$-seminorm*** *$\|-\|_{s,K,\mathbf{x}} \colon C^\infty M \to \mathbb{R}$ is defined by*

$$\|a\|_{s,K,\mathbf{x}} = \sup \left\{ \left| \left(\frac{d^s}{dx^S} a \right) (p) \right| : p \in K, S \subseteq \{1, \ldots, n\} \right\}$$

where S runs over multisets of cardinality s.

Note that, where $s = 0$ and $K = M$, the C^∞-seminorm is just the supremum norm. For $s > 0$, the C^∞-seminorm is not a norm even with $K = M$, since it sends e.g. all constant functions to 0. However, it is straightforward to check that $\|-\|_{s,K,\mathbf{x}}$ is at least a seminorm.

Definition 6. *The* ***C^∞-topology*** *is the topology generated by all seminorms of the form $\|-\|_{s,K,\mathbf{x}}$.*

Convergence under the C^∞-topology corresponds to uniform convergence of partial derivatives on every compact set. Since it was induced by a family of seminorms, this topology makes $C^\infty M$ into a TVS. It is not hard to see that $C^\infty M$ is also complete under this topology. Furthermore, it can be shown, under the standard assumption that M is second-countable, that only a countable collection of C^∞-seminorms suffice to induce the topology on $C^\infty M$. Thus, the C^∞-topology makes $C^\infty M$ into a *Fréchet space*.

Finally, we need to topologize our spaces of maps between coordinate algebras. This is more straightforward: the vector spaces $\mathrm{Hom}_{\mathrm{TVS}}(C^\infty N, C^\infty M)$ of continuous linear maps between C^∞-topologies will take the topology of pointwise convergence.

We now affirm that various important maps between our newly defined topological spaces are continuous. Unfortunately, the proof of these facts had to be omitted.

Proposition 4. *The product $(- \cdot -) \colon C^\infty M \times C^\infty M \to C^\infty M$ is continuous.*

Proposition 5. *The composition map*

$$\begin{aligned}(- \circ -) \colon\ &\mathrm{Hom}_{\mathrm{TVS}}(C^\infty O, C^\infty N) \times \mathrm{Hom}_{\mathrm{TVS}}(C^\infty N, C^\infty M) \\ &\to \mathrm{Hom}_{\mathrm{TVS}}(C^\infty O, C^\infty M)\end{aligned}$$

is continuous.

Proposition 6. *When $\varphi \colon M \to N$ is a smooth map between manifolds, $\hat{\varphi} \colon C^\infty N \to C^\infty M$ is a continuous linear map between C^∞-topologies.*

Proposition 7. *The action of a vector field over M on elements of $C^\infty M$ is continuous.*

Proposition 4 means that $C^\infty M$ is, besides a topological vector space, a *topological algebra*. Proposition 5 justifies the use of the product rule to differentiate compositions of time-dependent families of smooth maps. Finally, Propositions 6 and 7 mean that smooth maps and vector fields are embedded as elements in our Fréchet spaces of operators.

We are now prepared to investigate the "tangent spaces" to Hom-sets of smooth maps.

Definition 7. *When $\mathscr{R}$ is a subset of a topological vector space V and $p \in \mathscr{R}$, the* ***tangent space*** $T_p\mathscr{R}$ *is the subset of elements $v \in V$ so that*

$$\begin{cases} q_0 = p \\ \dot{q}_0 = v \end{cases}$$

for some curve $q\colon \mathbb{R} \to \mathscr{R}$.

Definition 8. *Let $Q\colon A \to B$ be a homomorphism of unital algebras. A Q-derivation is a linear map $X\colon A \to B$ which, for every $f, g \in A$, satisfies*

$$X(fg) = X(f)Q(g) + Q(f)X(g).$$

Proposition 8. *Let Q be a smooth map from M to N. Every element of the tangent space*

$$T_Q \operatorname{Hom}_{\mathrm{TVS}}(C^\infty N, C^\infty M)$$

is a Q-derivation.

Proof. Let $P_t \in \operatorname{Hom}_{\mathrm{Alg}}(C^\infty N, C^\infty M)$ be such that $P_0 = Q$ and $\dot{P}_0 = X$. Since we are working in the pointwise topology, $\dot{P}_0 = X$ means that, for any $f \in C^\infty M$, $P_t(f)$ is a differentiable trajectory in $C^\infty N$ with derivative

$$\left(\frac{d}{dt}\right)_0 P_t(f) = X(f).$$

Now, let $g \in C^\infty M$ also be arbitrary, and develop $X(fg)$ with the homomorphism law and the product rule to obtain

$$\begin{aligned} X(fg) &= \left(\frac{d}{dt}\right)_0 P_t(f)P_t(g) = \dot{P}_0(f)Q(g) + Q(f)\dot{P}_t(g) \\ &= X(f)Q(g) + Q(f)X(g). \end{aligned}$$

□

If M is the one-point manifold and Q encodes the inclusion of a point, a Q-derivation is a tangent vector at Q. If $M = N$ and $Q = \mathrm{id}$, a Q-derivation is a vector field on M. More generally, a Q-derivation encodes what, in the classical perspective, we would call a section of TN along Q.

The objects we will most interested in, however, are vector fields and self-diffeomorphisms. The next proposition puts a tool linking these two kinds of objects—the exponential map—under the lens of coordinate algebras.

Proposition 9. *Let X be a complete vector field on M. Then the Cauchy problem*

$$\begin{cases} Q_0 = \mathrm{id} \\ \dot{Q}_t = Q_t X \end{cases} \tag{1.2}$$

has a unique solution, namely $Q_t = \exp(tX)$, among the class of Lipschitz trajectories

$$Q\colon \mathbb{R} \to \operatorname{End}_{\mathrm{TVS}}(C^\infty M).$$

In particular, $T_{\mathrm{id}}\,\mathrm{End}_{\mathrm{TVS}}(C^\infty M)$ *contains the set of complete vector fields on* M.

Notice how, under the coordinate algebra perspective, it becomes evident how the exponential map for vector fields is analogous to the exponential maps for matrices and scalars!

Partial proof of Proposition 9. Suppose Q_t is a Lipschitz family satisfying (1.2). Then, for any given point $p \in M$, $p_t = pQ_t$ must also be a Lipschitz trajectory satisfying the Cauchy problem

$$\begin{cases} p_0 = p \\ \dot{p}_t = pQ_tX = p_tX. \end{cases}$$

The second equation is simply the integral curve equation—in classical language, $\dot{p}_t = X(p_t)$. It follows from the theory of ODEs that p_t is uniquely determined as $p\exp(tX)$, and so Q_t must be the flow $\exp(tX)$. We must omit a proof that $\exp(tX)$ is in fact a Lipschitz family. □

Remark 1. Unfortunately, many technical details in this section regarding the topologization of operator spaces had to be overlooked. Our focus will be on "formal" questions more than on "technical" ones; we are interested more in the fact that the exponential map verifies an equation like

$$\frac{d}{dt}\exp(tX) = \exp(tX)X$$

than in the technical aspect of what it means for a family of smooth maps to be Lipschitz-continuous in a real parameter. However, before ignoring them in the sequel, let us take a moment to discuss why we are not using a simpler topologization.

Suppose that, rather than the C^∞-topology, we equipped $C^\infty M$ with the (weaker) topology of pointwise convergence of functions. A smooth map still gives a continuous map between coordinate algebras under this topology, and a family Q_t of smooth maps between manifolds would be Lipschitz or differentiable exactly when, for every point p, the path pQ_t is respectively Lipschitz or differentiable. Under these definitions, the previous proposition follows directly from what we know about the regularity of integral curves.

However, there are some serious drawbacks to this approach. Most significantly, this weaker topology makes any non-zero tangent vector M into an unbounded linear functional, so the families of points or maps can no longer be differentiated meaningfully.

Remark 2. The author conjectures that converse of Proposition 8 holds under the hypothesis that M is compact. We also conjecture that the inclusion proven in Proposition 9 is strict.

Now, we consider a generalization of the Cauchy problem (1.2) that dictates the flow of a *non-autonomous* vector field—a vector field that depends on time. The flows of non-autonomous vector fields will become important when we discuss control systems in Chapter 3. We will need several more technical results, stated without proof. (For more details on the missing proofs in this chapter, see [2].)

In the following, let $\mathrm{Aut}(M)$ be the group of diffeomorphisms of $C^\infty M$. All intervals I are assumed to be compact, and all ODEs are understood to hold almost everywhere over the domain of interest.

Theorem 2. *Let $Q\colon I \to \mathrm{Aut}(M)$ be a Lipschitz family. Then Q is almost everywhere differentiable. Furthermore, for arbitrary $t_0 \in I$, it is characterized as the unique Lipschitz family $P\colon I \to \mathrm{Aut}(M)$ solving the Cauchy problem*

$$\begin{cases} P_{t_0} = Q_{t_0} \\ \dot{P}_t = P_t X_t \end{cases} \tag{1.3}$$

where X_t is the vector field $Q_t^{-1}\dot{Q}_t$.

Theorem 3. *Let $X : I \to XM$ be a locally bounded vector field supported on some compact set $K \subseteq M$. Then, for arbitrary $t_0 \in I$, there is a unique Lipschitz family Q_t of self-diffeomorphisms solving the Cauchy problem*

$$\begin{cases} Q_{t_0} = \mathrm{id} \\ \dot{Q}_t = Q_t X_t \end{cases} \tag{1.4}$$

for almost all $t \in I$.

Theorem 4. *Let Q_t and P_t be families of smooth maps. If Q_t is Lipschitz, then so is Q_t^{-1}, and Q_t^{-1} is differentiable at any differentiability points of Q_t. Furthermore, if Q_t and P_t are both Lipschitz, then so is the composition Q_tP_t, and Q_tP_t is differentiable at any joint differentiability points of Q_t and P_t.*

Combining Theorem 4 with the product rule, we can compute an expression for the derivative of $P_t = Q_t^{-1}$. It's enough to differentiate the equation $\mathrm{id} = P_tQ_t$:

$$\begin{aligned} 0 &= \frac{d}{dt}P_tQ_t = \dot{P}_tQ_t + P_t\dot{Q}_t \\ \Rightarrow \dot{P}_t &= -P_t\dot{Q}_tP_t. \end{aligned}$$

In the conditions of Theorem 3, we will refer to the solution Q_t of (1.4) as a function of X_t with the notation

$$Q_t = \overrightarrow{\exp}\int_{t_0}^{t} X_\tau\, d\tau.$$

The family Q_t is called the **flow** of the non-autonomous vector field X_t, and the function $\overrightarrow{\exp}$ is called the **(right) chronological exponential**. (Note that $\overrightarrow{\exp}$ is really a function of the family X_t rather than just of some integral $\int_0^t X_\tau\, d\tau$!) By "stitching together" integral curves, we obtain the relation

$$\overrightarrow{\exp}\int_{t_0}^{t_1} X_\tau\, d\tau \circ \overrightarrow{\exp}\int_{t_1}^{t_2} X_\tau\, d\tau = \overrightarrow{\exp}\int_{t_0}^{t_2} X_\tau\, d\tau$$

for any $t_0, t_1, t_2 \in \mathbb{R}$. Along the same lines, we can write

$$Q_t^{-1} = \left(\overrightarrow{\exp}\int_0^t X_\tau\, d\tau\right)^{-1} = \overrightarrow{\exp}\int_t^0 X_\tau\, d\tau$$

for the value which a family P must take at 0 if it solves the Cauchy problem

$$\begin{cases} P_t = \mathrm{id} \\ \dot{P}_s = P_sX_s. \end{cases}$$

How can we describe the evolution of the inverse $P_t = Q_t^{-1}$ of the flow of X_t? By our computation above,

$$\dot{P}_t = -P_t \dot{Q}_t P_t = -P_t Q_t X_t P_t = -X_t P_t.$$

Thus, Q_t solves the Cauchy problem (1.4) if and only if P_t solves

$$\begin{cases} P_0 = \mathrm{id} \\ \dot{P}_t = -X_t P_t. \end{cases} \tag{1.5}$$

We conclude that a "left-handed" Cauchy problem of the form (1.5) also has a unique Lipschitz solution. We write the "left-handed flow" P_t as a function of X_t with the **left chronological exponential**,

$$P_t = \overleftarrow{\exp} \int_0^t -X_\tau \, d\tau.$$

In summary, the right and left chronological exponentials verify

$$\frac{d}{dt}\left(\overrightarrow{\exp} \int_0^t X_\tau \, d\tau\right) = \left(\overrightarrow{\exp} \int_0^t X_\tau \, d\tau\right) \circ X_t$$
$$\frac{d}{dt}\left(\overleftarrow{\exp} \int_0^t X_\tau \, d\tau\right) = X_t \circ \left(\overleftarrow{\exp} \int_0^t X_\tau \, d\tau\right)$$

and are related by the formula

$$\overleftarrow{\exp} \int_0^t X_\tau \, d\tau = \overrightarrow{\exp} \int_t^0 -X_\tau \, d\tau.$$

It is worthwhile to define left- and right-handed exponentials to deal with the integration of non-autonomous vector fields precisely because a non-autonomous vector field V_t might not commute with its flow Q_t—it might not be true that $V_t Q_t = Q_t V_t$ for every t. On the other hand, it is not necessary to define a handedness for the autonomous exponential because $\exp(tX)$ always commutes with X.

What does it mean for a vector field to commute with a diffeomorphism? In the next section, we will characterize the conditions in which a vector field is invariant under the *adjoint action* of a diffeomorphism or another vector field.

1.2 The Adjoint Actions

Let M be a manifold, let Q be a self-diffeomorphism of M, and let X be a vector field on M. The compositions QX and XQ are Q-derivations; in classical language, pQX is $X(Q(p))$, while pXQ is $Q_*(X(p))$. The conjugation QXQ^{-1}, on the other hand, is a derivation—that is, an id-derivation—which is also known as the pushforward of X by Q^{-1}. We will call this the *adjoint action* of Q on X.

In the following, let $D(M)$ be the Lie algebra of vector fields on M. When necessary, we will write $D^*(M)$ for the Lie subalgebra of complete vector fields.

Definition 9. *Let Q be an self-diffeomorphism of M. The* ***adjoint action*** $\operatorname{Ad} X$ *of Q is map from $D(M)$ to $D(M)$ given by*

$$(\operatorname{Ad} Q)X = QXQ^{-1}.$$

The following properties are fairly clear.

Proposition 10. *The adjoint action* Ad *is a representation of* $\operatorname{Aut}(M)$ *on $D(M)$ by Lie algebra automorphisms. This means the following:*

1. *For all diffeomorphisms P and Q,* $\operatorname{Ad} PQ = (\operatorname{Ad} P)(\operatorname{Ad} Q)$. *This means that* Ad *is an action of* $\operatorname{Aut}(M)$ *on $D(M)$.*

2. *For all diffeomorphisms Q,* $\operatorname{Ad} Q$ *is linear, and for all vector fields X and Y, we have* $(\operatorname{Ad} Q)[X,Y] = [(\operatorname{Ad} Q)X, (\operatorname{Ad} Q)Y]$. *This means that* $\operatorname{Ad} Q$ *is a Lie algebra homomorphism for each Q.*

Furthermore, if M is not zero-dimensional, then Ad *is a faithful representation:* $\operatorname{Ad} Q = \mathrm{id} \Leftrightarrow Q = \mathrm{id}$.

In classical language, we say that Q is a *Lie point symmetry* of X when the following diagram commutes.

$$\begin{array}{ccc} M & \xrightarrow{\;Q\;} & M \\ \big\downarrow X & & \big\downarrow X \\ TM & \xrightarrow{\;dQ\;} & TM \end{array}$$

In the coordinate algebra perspective, this has the following interpretation.

Definition 10. *A diffeomorphism Q is a* ***Lie point symmetry*** *of X when*

$$(\operatorname{Ad} Q)X = X.$$

Lie point symmetries of a vector field can also be characterized as symmetries of its family of integral curves or as transformations that commute with its flow.

Proposition 11. *Let Q be a diffeomorphism and X a complete vector field. The following conditions are equivalent:*

1. *Q is a Lie point symmetry of X.*

2. *Whenever c_t is an integral curve of X, $c_t Q$ is also an integral curve of X.*

3. *Q commutes with $\exp(tX)$ for every t.*

Proof. We will prove the three implications $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 1$.

- ($1 \Rightarrow 2$) Let c_t be an integral curve of X, so that $\dot{c}_t = c_t X$. By hypothesis, $QX = XQ$, so we have $\frac{d}{dt}(c_t Q) = c_t XQ = c_t QX$. This proves that $c_t Q$ is also an integral curve of X.

- ($2 \Rightarrow 3$) We will prove that $Q\exp(tX) = \exp(tX)Q$ by showing that, for any given point $p \in M$,

$$pQ\exp(tX) = p\exp(tX)Q.$$

Let c_t be the left-hand side of this equation, and d_t the right-hand side. It is obvious that c_t is an integral curve of X. Furthermore, since $p\exp(tX)$ is an integral curve of X, (2) implies that d_t is too. Since $c_0 = d_0 = pQ$, we conclude that $c_t = d_t$ for all t by uniqueness of integral curves.

- ($3 \Rightarrow 1$) Differentiating the equation $\exp(tX)Q = Q\exp(tX)$ at $t = 0$ gives $XQ = QX$.

□

Note that the identity $X\exp(X) = \exp(X)X$ is a consequence of the previous proposition, since $\exp(tX)\exp(X) = \exp((t+1)X) = \exp(X)\exp(tX)$. Indeed, every complete vector field X has a natural family of Lie point symmetries—its flow, $\exp(tX)$!

The adjoint action of diffeomorphisms also has a simple relationship with *non-autonomous* flows.

Proposition 12. *Let Q be a diffeomorphism and let $X\colon [0,t] \to D(M)$ be a vector field in the conditions of Theorem 3. Then*

$$Q\left(\overrightarrow{\exp}\int_0^t X_\tau\, d\tau\right)Q^{-1} = \left(\overrightarrow{\exp}\int_0^t (\operatorname{Ad} Q)X_\tau\, d\tau\right).$$

Proof. Let $P_t = \overrightarrow{\exp}\int_0^t X_\tau\, d\tau$. Using the fact that P_t is Lipschitz, it can be shown that QP_tQ^{-1} is also Lipschitz. Applying the product rule gives

$$\frac{d}{dt}QP_tQ^{-1} = QP_tX_tQ^{-1} = QP_tQ^{-1}(\operatorname{Ad} Q)X_t.$$

We have proven that $R_t = QP_tQ^{-1}$ is a Lipschitizian family solving the Cauchy problem

$$\begin{cases} R_0 = \mathrm{id} \\ \dot{R}_t = R_t(\operatorname{Ad} Q)X_t. \end{cases}$$

□

In particular, if Q is a Lie point symmetry of X_t for all t, then Q commutes with the flow of X_t.

Remark 3. Previously, we had characterized the exponential $Q_t = \exp(tX)$ by the differential equation $\dot{Q}_t = Q_tX$, but, as we have now proven, we could have used the equation $\dot{Q}_t = XQ_t$. This new form has an interesting interpretation from the coordinate algebra perspective: it gives a transport equation! Indeed, choosing coordinates $\mathbf{x} = (x_1, \ldots, x_n)$ where $X = \partial/\partial x_1$ and writing $g(t, \mathbf{x})$ for the value of $Q_t(f)$ at the point with coordinates $\mathbf{x}$, the equation $\dot{Q}_t(f) = XQ_t(f)$ reads

$$\frac{\partial}{\partial t}g(t,\mathbf{x}) = \frac{\partial}{\partial x_1}g(t,\mathbf{x}).$$

So far, we have explored the adjoint action $\operatorname{Ad} Q$ of a diffeomorphism Q. Next, we will define the adjoint representation of a *vector field* X, which we will denote by $\operatorname{ad} X$. In some sense, this can be viewed as the "differential" of the adjoint representation for smooth maps; that is, we define

$$(\operatorname{ad} X)Y = \left(\frac{d}{dt}\right)_0 (\operatorname{Ad}\exp(tX))Y = XY - YX. \tag{1.6}$$

Strictly speaking, $\exp(tX)$ only exists when X is complete, so the following definition is a formal extension of this formula.

Definition 11. *Let X be a vector field on M. The **adjoint action** $\operatorname{ad}X\colon D(M) \to D(M)$ of X is the Lie algebra automorphism given by*

$$(\operatorname{ad}X)Y = [X,Y] = XY - YX.$$

Just like Ad, the adjoint action ad can be viewed as a representation.

Proposition 13. *The adjoint action* ad *is a representation of $D(M)$ on $D(M)$ by derivations.*[5] *This means the following:*

1. $\operatorname{ad}[X,Y] = [\operatorname{ad}X, \operatorname{ad}Y]$, *where* $[\operatorname{ad}X, \operatorname{ad}Y] = \operatorname{ad}X \operatorname{ad}Y - \operatorname{ad}Y \operatorname{ad}X$. *This means that* ad *is a linear representation of $D(M)$ over $D(M)$.*
2. $\operatorname{ad}X$ *is linear and verifies* $(\operatorname{ad}X)[Y,Z] = [(\operatorname{ad}X)Y,Z] + [Y,(\operatorname{ad}X),Z]$, *meaning that* $\operatorname{ad}X$ *is a derivation of the Lie algebra $D(M)$.*

Note that both the homomorphism law for the adjoint action ad, namely

$$[[X,Y],Z] = (\operatorname{ad}[X,Y])(Z) = [\operatorname{ad}X, \operatorname{ad}Y](Z) = [X,[Y,Z]] - [Y,[X,Z]],$$

and derivation law for $\operatorname{ad}X$, namely

$$[X,[Y,Z]] = (\operatorname{ad}X)[Y,Z] = [(\operatorname{ad}X)Y,(\operatorname{ad}X),Z] = [[X,Y],Z] + [Y,[X,Z]],$$

are rearrangements of the *Jacobi identity*, $[X,[Y,Z]] + [Y,[Z,X]] + [Z,[X,Y]] = 0$. Furthermore, when X is complete, the derivation law for $\operatorname{ad}X$ results from differentiating the homomorphism law for $\operatorname{Ad}\exp(tX)$:

$$\begin{aligned}(\operatorname{ad}X)[Y,Z] &= \left(\frac{d}{dt}\right)_0 (\operatorname{Ad}\exp(tX))[Y,Z] = \left(\frac{d}{dt}\right)_0 [(\operatorname{Ad}\exp(tX))Y,(\operatorname{Ad}\exp(tX))Z] \\ &= \left[\left(\frac{d}{dt}\right)_0 (\operatorname{Ad}tX)Y,Z\right] + \left[Y,\left(\frac{d}{dt}\right)_0 (\operatorname{Ad}\exp(tX))Z\right] \\ &= [(\operatorname{ad}X)Y,Z] + [Y,(\operatorname{ad}X)Z].\end{aligned}$$

(Here we are using the fact that the Lie bracket a continuous bilinear operator under our topologization.) This is a memorable way to prove the Jacobi identity for vector fields.

We say that two vector fields X and Y *commute* when they commute as differential operators, meaning that $[X,Y] = 0$. Commutativity of vector fields is related to the idea of Lie point symmetry in the following way.

Proposition 14. *Let X and Y be complete vector fields. Then X and Y commute if and only if, for every t, $\exp(tX)$ is a Lie point symmetry of Y.*

[5] Actually, such an adjoint representation can be defined for any Lie algebra. We are using no special properties of $D(M)$ in this proposition.

Proof. Define the family $Z_t = (\operatorname{Ad}\exp(tX))Y$. It can be shown that this is a Lipschitz family of vector fields. We claim that Z_t is the unique Lipschitz solution to the Cauchy problem

$$\begin{cases} Z_0 = Y \\ \dot{Z}_t = (\operatorname{ad} X)Z_t. \end{cases} \tag{1.7}$$

This will prove the proposition, since the flow $\exp(tX)$ is a Lie point symmetry of Y exactly when $Z_t = Y$ for all t, and $Z_t = Y$ is a solution to the above problem exactly when $(\operatorname{ad} X)Y = Y$.

Suppose Z_t is a Lipschitz family verifying (1.7). Then the family $Z_t^* = (\operatorname{Ad}\exp(-tX))Z_t$ is also Lipschitz, and differentiating with the product rule gives

$$\begin{aligned}
\frac{d}{dt}Z_t^* &= \frac{d}{dt}\exp(-tX)Z_t\exp(tX) \\
&= -X\exp(-tX)Z_t\exp(tX) + \exp(-tX)Z_t\exp(tX)X \\
&+ \exp(-tX)[X,Z_t]\exp(tX) \\
&= -X\exp(-tX)Z_t\exp(tX) + \exp(-tX)Z_t\exp(tX)X \\
&+ X\exp(-tX)Z_t\exp(tX) - \exp(-tX)Z_t\exp(tX)X \\
&= 0.
\end{aligned}$$

Furthermore, $Z_0^* = Y$. We conclude that $Z_t^* = Y$, and so $Z_t = (\operatorname{Ad}\exp(tX))Y$. □

As happened with Proposition 11, we have the following natural extension of Proposition 14 to the chronological exponential. We omit the proof, which follows very analogously to the proof above.

Proposition 15. *Let X be a vector field and $Y : [0,t] \to D(M)$ a family of vector fields in the conditions of Theorem 3. If X commutes with Y_τ for almost all $\tau \in [0,t]$, then*

$$\overrightarrow{\exp}\int_0^t Y_\tau\, d\tau$$

is a Lie point symmetry of X.

Remark 4. The Cauchy problem (1.7) for the vector field Z_t above can also understood as a transport equation. If we choose a coordinate system $\mathbf{x} = (x_1,\ldots,x_n)$ for which $X = \partial/\partial x_1$ and write

$$Z_t = a_i(t,\mathbf{x})\frac{\partial}{\partial x_i}, \quad Y = b_i(\mathbf{x})\frac{\partial}{\partial x_i},$$

it simply dictates that

$$\begin{cases} a_i(0,\mathbf{x}) = b_i(\mathbf{x}) \\ \dfrac{\partial a_i}{\partial t}(t,\mathbf{x}) = \dfrac{\partial a_i}{\partial x_1}(t,\mathbf{x}), \end{cases}$$

which is a transport equation with velocity $-\partial/\partial x_1$. The vector fields X and Y will commute exactly when Y is unaffected by this transport process—meaning, in this coordinate system, that the coefficients b_i are constant in x_1. With this remark in mind, it is also easy to prove that the adjoint representation for vector fields is faithful.

Combining the previous proposition with one of our equivalent characterizations of Lie point symmetry proves the fundamental result that **complete vector fields commute iff their flows commute**.[6] Compared to the usual way this is done, the author thinks the approach we have taken here is particularly natural and memorable.

It is also interesting to note that everything stated above, except for the faithfulness of the adjoint representations, would apply equally well if we substituted $C^\infty M$ for $\mathrm{GL}(\mathbb{R}^n)$, $D^*(M)$ and $D(M)$ for $\mathfrak{gl}(\mathbb{R}^n)$, and the vector field exponential for the matrix exponential. Thinking in terms of coordinate algebras lets us do differential geometry as if we were doing matrix-valued calculus! However, there is a crucial topic from matrix-valued calculus that we have not yet discussed in our study of vector fields and flows: we have not yet formulated an *exponential series* for the flow of a vector field. In the next section, we will introduce a Taylor series-like formula that applies to the chronological exponential. We also tackle the problem of approximating the chronological exponential near a non-zero vector field by means of integral formulas called the *variations formulas*.

1.3 Approximating the Chronological Exponential

To develop a series formula for the differential equation

$$\begin{cases} Q_0 = \mathrm{id} \\ \dot{Q}_t = Q_t X_t, \end{cases}$$

we will rewrite it as an integral equation

$$Q_t = B_t(Q_-),$$

where B assigns the family Q_t of operators to the new family

$$B_t(Q_-) = \mathrm{id} + \int_0^t Q_\tau X_\tau \, d\tau.$$

Here, integration of a locally bounded vector-valued function over a real domain is defined as in [20]; in particular, the action of $B_t(Q_-)$ on an element $f \in C^\infty M$ is defined by

$$B_t(Q_-)(f) = f + \int_0^t Q_\tau X_\tau(f) \, d\tau,$$

where this integral is interpreted in the pointwise sense. Although we will not argue this here, it can be shown that the equation $Q_t = B_t(Q_-)$ (which, like before, must hold for all t) is equivalent, among the class of Lipschitz families Q_t, to the Cauchy problem above.

The advantage of rewriting our Cauchy problem in a fixed-point form is that it can now be iterated; any solution to $Q_t = B_t(Q_-)$ also verifies $Q_t = B_t^n(Q_-)$ for any natural n.

[6] Another classic result, which can be proven using the tools developed here, is that $\exp(X+Y) = \exp(X)\exp(Y)$ whenever $[X,Y] = 0$.

Let us compute $B_t^2(Q_-)$. We find

$$\begin{aligned} B_t^2(Q_-) &= \mathrm{id} + \int_0^t B_{\tau_1}(Q_-) X_{\tau_1}\, d\tau_1 \\ &= \mathrm{id} + \int_0^t \left(\mathrm{id} + \int_0^{\tau_1} Q_{\tau_2} X_{\tau_2}\, d\tau_2 \right) X_{\tau_1}\, d\tau_1. \end{aligned}$$

Since integration is being performed in the pointwise sense, it distributes with composition with operators on the right. Thus, we can rearrange this last expression as

$$\begin{aligned} B_t^2(Q_-) &= \mathrm{id} + \int_0^t \left(X_{\tau_1} + \int_0^{\tau_1} Q_{\tau_2} X_{\tau_2} X_{\tau_1}\, d\tau_2 \right) d\tau_1 \\ &= \mathrm{id} + \int_0^t X_{\tau_1}\, d\tau_1 + \int_0^t \int_0^{\tau_1} Q_{\tau_2} X_{\tau_2} X_{\tau_2}\, d\tau_2\, d\tau_1, \end{aligned}$$

where for the second equation we have also applied linearity. Continuing this process inductively leads to the following series expression, called the *Volterra series*.

Theorem 5 (Volterra series)**.** *Let X_t be a vector field in the conditions of Theorem 3. Let $\Delta_n(t)$ denote the domain*

$$\{(\tau_1, \ldots, \tau_n) : 0 \le \tau_n \le \ldots \le \tau_1 \le t\} \subseteq \mathbb{R}^n.$$

Then the flow Q_t of X_t can be expressed as

$$Q_t = \overrightarrow{\exp} \int_0^t X_\tau\, d\tau = \mathrm{id} + \sum_{k=1}^{n} \int_{\Delta_k(t)} X_{\tau_k} \ldots X_{\tau_1}\, d\tau_k \ldots d\tau_1 + R_n^t(X)$$

where the $R_n^t(X)$ has the explicit formula

$$R_n^t(X) = \int_{\Delta_{n+1}(t)} Q_{\tau_{n+1}} X_{\tau_n} \ldots X_{\tau_1}\, d\tau_{n+1} \ldots d\tau_1.$$

In [2] it is proven that, for any $a \in C^\infty M$,

$$\left\| R_n^t(X)(a) \right\|_{s,K} = \mathscr{O}(t^{n+1})$$

where the vector field X_t is held constant, and

$$\left\| R_n^t(\varepsilon X)(a) \right\|_{s,K} = \mathscr{O}(\varepsilon^{n+1})$$

where t is held constant.

The first-order truncation of the Volterra series,

$$Q_t \approx \mathrm{id} + \int_0^t X_\tau\, d\tau,$$

can be interpreted as approximating

$$\frac{pQ_t - pQ_0}{t}$$

by the average value of pX_t in the interval $[0,t]$. Now, consider the second-order truncation,

$$Q_t \approx \mathrm{id} + \int_0^t X_\tau \, d\tau + \int_0^t \int_0^{\tau_1} X_{\tau_2} X_{\tau_1} \, d\tau_2 d\tau_1. \tag{1.8}$$

To help us understand the behavior the second-order term appearing here, suppose X_t can be written as a linear combination

$$X = a_t V + b_t W$$

of two compactly-supported vector fields V and W. Let us write $A_t = \int_0^t a_\tau \, d\tau$ and $B_t = \int_0^t b_\tau \, d\tau$. Then, applying (1.8) gives

$$\begin{aligned}
\overrightarrow{\exp} \int_0^1 X_\tau \, d\tau &\approx \mathrm{id} + \int_0^1 (a_\tau V + b_\tau W) \, d\tau \\
&\quad + \int_0^1 \int_0^{\tau_1} (a_{\tau_2} V + b_{\tau_2} W)(a_{\tau_1} V + b_{\tau_1} W) \, d\tau_2 d\tau_1 \\
&= \mathrm{id} + A_1 V + B_1 W + \left(\int_0^1 A_\tau a_\tau \, d\tau \right) V^2 + \left(\int_0^1 A_\tau b_\tau \right) VW \\
&\quad + \left(\int_0^1 B_\tau a_\tau \, d\tau \right) WV + \left(\int_0^1 B_\tau b_\tau \, d\tau \right) W^2.
\end{aligned}$$

Now, suppose (A_t, B_t) draws a closed curve on the plane, meaning $A_1 = B_1 = 0$, so that the first-order terms in this expression will vanish. Suppose furthermore that the curve (A_t, B_t) is simple and runs counter-clockwise around a region of area η. Applying Stoke's theorem, we find that

$$\begin{aligned}
\int_0^1 A_\tau a_\tau \, d\tau &= 0 & \int_0^1 A_\tau b_\tau \, d\tau &= \eta \\
\int_0^1 B_\tau a_\tau \, d\tau &= -\eta & \int_0^1 B_\tau b_\tau \, d\tau &= 0.
\end{aligned}$$

Applying our characterization of the Volterra series' remainder, this proves altogether that

$$\overrightarrow{\exp} \int_0^1 \varepsilon X_\tau \, d\tau = \mathrm{id} + \varepsilon^2 \eta [V, W] + \mathscr{O}(\varepsilon^3).$$

In the case that (a_t, b_t) are defined piecewise by

$$\begin{cases}
(4,0) & : 0 \le t < 1/4 \\
(0,4) & : 1/4 \le 4 < 1/2 \\
(-4,0) & : 1/2 \le 4 < 3/4 \\
(0,-4) & : 3/4 \le 4 \le 1,
\end{cases}$$

this reduces to the formula

$$\exp(\varepsilon V) \exp(\varepsilon W) \exp(-\varepsilon V) \exp(-\varepsilon W) = \mathrm{id} + \varepsilon^2 [V, W] + \mathscr{O}(\varepsilon^3)$$

which can be checked with the exponential series alone.

Finally, suppose we want to develop an approximation of the chronological exponential $\overrightarrow{\exp}\int_0^t \tilde{V}_\tau\,d\tau$ for vector fields close to a given vector field V_t. For example, we might want to *differentiate* the chronological exponential at V_t; that is, get a first-order approximation for

$$\overrightarrow{\exp}\int_0^t V_\tau + \varepsilon W_\tau\,d\tau,$$

for small ε, where W_t is an arbitrary "perturbation". This will be made possible by the *variations formulas*.

Our analysis will proceed in two steps. First, we use an informal argument in the style of "variations of parameters" arguments to guess the form of the variations formulas. Finally, we will check rigorously that they hold.

Let V_t be a given vector field with flow $P_{t_0}^{t_1} = \overrightarrow{\exp}\int_{t_0}^{t_1} V_\tau\,d\tau$, and suppose δ_t is an "impulse" concentrated at a moment $s > 0$, so that

$$\overrightarrow{\exp}\int_0^t \delta_\tau\,d\tau = \begin{cases} \mathrm{id} & : t < s \\ \Delta & : t \geq s \end{cases}$$

for some element $\Delta \in \mathrm{Aut}(M)$. In this formalism, we would expect the flow of $V_t + \delta_t$ to be

$$\overrightarrow{\exp}\int_0^t V_\tau + \delta_\tau\,d\tau = \begin{cases} P_0^t & : t < s \\ P_0^s \Delta P_s^t & : t \geq s. \end{cases}$$

For $t > s$, we can rewrite this flow as

$$\overrightarrow{\exp}\int_0^t V_\tau + \delta_\tau\,d\tau = P_0^s \Delta P_s^0 P_0^t = P_0^t P_t^s \Delta P_s^t. \tag{1.9}$$

Extending the formalism of our "impulse" further to respect the commutativity of conjugation with the exponential (Proposition 12), we could propose that, for a self-diffeomorphism Q,

$$Q\Delta Q^{-1} = Q\left(\overrightarrow{\exp}\int_0^t \delta_\tau\,d\tau\right)Q^{-1} = \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\,Q)\delta_\tau\,d\tau.$$

Furthermore, because δ_t is only supported at $t = s$, we would expect that, e.g.,

$$P_0^s \Delta P_s^0 = \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\,P_0^s)\delta_\tau\,d\tau = \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\,P_0^\tau)\delta_\tau\,d\tau.$$

So, Equation (1.9) can be written as

$$\begin{aligned}\overrightarrow{\exp}\int_0^t V_\tau + \delta_\tau\,d\tau &= \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\,P_0^\tau)\delta_\tau\,d\tau \circ P_0^t \\ &= P_0^t \circ \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\,P_t^\tau)\delta_\tau\,d\tau.\end{aligned}$$

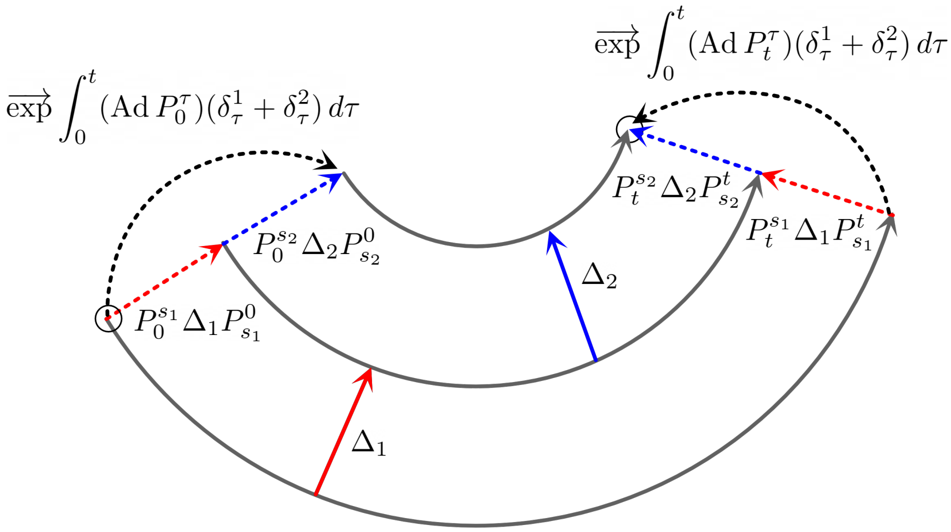

Fig. 1.1 The diagram underlying the formal variations formulas for the flow of a vector field perturbed by two impulses. Gray arrows indicate the flow of the unperturbed vector field V_t.

Now, suppose we have a set of n impulse terms, $\delta_t^1, \ldots, \delta_t^n$, which apply impulses $\Delta_1, \ldots, \Delta_n$ at times $0 < s_1 < \ldots < s_n < t$. Then we can apply the formulas above inductively to conclude

$$\begin{aligned}\overrightarrow{\exp}\int_0^t \left(V_\tau + \sum_i \delta_\tau^i\right) d\tau &= \prod_{i=1}^n \left(\overrightarrow{\exp}\int_0^t (\mathrm{Ad}\, P_0^\tau)\delta_\tau^i\, d\tau\right) \circ P_0^t \\ &= P_0^t \circ \prod_{i=1}^n \left(\overrightarrow{\exp}\int_0^t (\mathrm{Ad}\, P_t^\tau)\delta_\tau^i\, d\tau\right).\end{aligned}$$

Since the right chronological exponential of a sum of impulses is the same as the composition of their integrals, applying impulses that “happen earlier” first, we finally conclude the following variations formula for arbitrary sums of impulses.

$$\begin{aligned}\overrightarrow{\exp}\int_0^t \left(V_\tau + \sum_i \delta_\tau^i\right) d\tau &= \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\, P_0^\tau)\left(\sum_i \delta_\tau^i\right) d\tau \circ P_0^t \\ &= P_0^t \circ \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\, P_t^\tau)\left(\sum_i \delta_\tau^i\right) d\tau.\end{aligned}$$

Now, we posit that the same formula is valid when $\sum_i \delta_t^i$ is replaced by a legitimate vector field. With our formula in hand, it is not hard to prove that it is the right one.

Theorem 6. *Let V_t and W_t be vector fields in the conditions of Theorem 3, and let $P_{t_0}^{t_1} = \overrightarrow{\exp}\int_{t_0}^{t_1} V_\tau\, d\tau$. Then we have the equations*

$$\begin{aligned}\overrightarrow{\exp}\int_0^t V_\tau + W_\tau\, d\tau &= \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\, P_0^\tau) W_\tau\, d\tau \circ P_0^t \\ &= P_0^t \circ \overrightarrow{\exp}\int_0^t (\mathrm{Ad}\, P_t^\tau) W_\tau\, d\tau\end{aligned}$$

The two expressions above are called the **variations formulas**.

Proof. Let $Q_t = \overrightarrow{\exp}\int_0^t V_\tau + W_\tau\, d\tau$. We know that $P_t^0 Q_t$ and $Q_t P_t^0$ are Lipschitz families of self-diffeomorphisms, taking value id at $t = 0$. We first prove the equality

$$Q_t P_t^0 = \overrightarrow{\exp}\int_0^t (\operatorname{Ad} P_0^\tau) W_\tau\, d\tau \tag{1.10}$$

by proving that

$$\frac{d}{d\tau} Q_\tau P_\tau^0 = Q_\tau P_\tau^0 (\operatorname{Ad} P_0^\tau) W_\tau.$$

This follows directly from the identity $(d/d\tau)P_\tau^0 = -V_\tau P_\tau^0$ and the product rule:

$$\begin{aligned}\frac{d}{d\tau} Q_\tau P_\tau^0 &= Q_\tau (V_\tau + W_\tau) P_\tau^0 - Q_\tau V_\tau P_\tau^0 = Q_\tau W_\tau P_\tau^0 \\ &= Q_\tau P_\tau^0 P_0^\tau W_\tau P_\tau^0 = Q_\tau P_\tau^0 (\operatorname{Ad} P_0^\tau) W_\tau.\end{aligned}$$

This gives the first variations formula. To prove the second, just conjugate Equation (1.10):

$$\begin{aligned}P_t^0 Q_t = P_t^0 Q_t P_t^0 P_0^t &= P_t^0 \left(\overrightarrow{\exp}\int_0^t (\operatorname{Ad} P_0^\tau) W_\tau\, d\tau\right) P_0^t \\ &= \overrightarrow{\exp}\int_0^t (\operatorname{Ad} P_t^0 P_0^\tau) W_\tau\, d\tau = \overrightarrow{\exp}\int_0^t (\operatorname{Ad} P_t^\tau) W_\tau\, d\tau.\end{aligned}$$

□

Remark 5. Combining a variations formula with the first-order truncation of the Volterra series gives an integral representation for the differential of the autonomous exponential:

$$\begin{aligned}\exp(V + \varepsilon W) &= \overrightarrow{\exp}\int_0^1 V + \varepsilon W\, d\tau \\ &= \left(\overrightarrow{\exp}\int_0^1 (\operatorname{Ad}\exp(\tau V))\varepsilon W\, d\tau\right)\exp(V) \\ &= \left(\mathrm{id} + \varepsilon\int_0^1 (\operatorname{Ad}\exp(\tau V)) W\, d\tau + \mathscr{O}(\varepsilon^2)\right)\exp(V) \\ &= \exp(V) + \varepsilon\int_0^1 \exp(\tau V) W \exp((1-\tau)V)\, d\tau + \mathscr{O}(\varepsilon^2).\end{aligned}$$

The same formula holds for the matrix exponential.

1.4 Lie Groups from the Coordinate Algebra Perspective

Over the course of this chapter, we have been able to treat the group $\operatorname{Aut}(M)$ of diffeomorphisms of a manifold as a sort of Lie group. Most significantly, we have hinted that $D(M)$ (or perhaps $D^*(M)$) ought to be its Lie algebra, and defined adjoint actions like the ones usually defined for a Lie group. Now, as a point of curiosity, we review some basic facts about finite-dimensional Lie groups in the language of coordinate algebras. We hope to convince the reader that coordinate algebras can yield an elegant way to do differential geometry!

Let G be a Lie group. For each element $g \in G$, let L_g and R_g be the left and right translation maps,

$$hL_g = gh, \quad hR_g = hg.$$

A vector field $X \in D(G)$ is left-invariant when every left-translation is a Lie point symmetry of X. Right-invariant vector fields are defined analogously. Left translations and right translations themselves are also characterized by commutativity relations; a homomorphism $Q\colon C^\infty G \to C^\infty G$ is a left translation iff it commutes with every right translation. Indeed, if $QR_g = R_g Q$ for all $g \in G$, then

$$gQ = eR_g Q = eQR_g = gL_{eQ}.$$

The spaces of left- and right-invariant vector fields are closed under the Lie bracket. (Indeed, *any* polynomial in left- or right-invariant vector fields will commute with left- or right-translations.) It is also easy to check that the adjoint action of the inversion map $i\colon G \to G$ puts these two spaces of vector fields in bijection; for example, if X is left-invariant, then $(\mathrm{Ad}\, i)X$ is right-invariant:

$$iXiR_g = iXL_{g^{-1}}i = iL_{g^{-1}}Xi = R_g iXi.$$

It is customary to distinguish the space of left-invariant vector fields. We call this the Lie algebra, $\mathfrak{g}$, of the group. For every $X \in \mathfrak{g}$ and $g \in G$, commutativity with left-translations implies

$$gX = gL_{g^{-1}}L_g X = eXL_g.$$

In classical language, this equation reads $X(g) = (L_g)_* X(e)$, meaning that a left-invariant vector field X is determined by its value at e. Conversely, for any vector $v \in T_e G$, we can define a smooth vector field X that takes $gX = vL_g$. This gives a standard isomorphism $T_e G \cong \mathfrak{g}$ of vector spaces.

As we know, a left-invariant vector field X is always complete. Furthermore, its exponential $\exp X$ will be a *right-translation* by an element of the group, because

$$L_g(\exp X)L_{g^{-1}} = \exp(L_g X L_{g^{-1}}) = \exp X.$$

Identifying right translations with elements of G, this lets us see the exponential as a smooth map from $\mathfrak{g}$ to G. Finally, let us see how we can construct the functor LieGrp $\to$ LieAlg.

Lemma 2. *Suppose $\varphi\colon G \to H$ is a homomorphism of Lie groups. For every $X \in \mathfrak{g}$, there is a unique element $Y \in \mathfrak{h}$ so that $X\varphi = \varphi Y$.*

Proof. Note that the condition $X\varphi = \varphi Y$ forces

$$eY = e\varphi Y = eX\varphi.$$

So, let Y be the unique left-invariant vector field that takes $eY = eX\varphi$. This vector field is defined on elements $h \in H$ by the equation $hY = eX\varphi L_h$. It only remains to check that $X\varphi = \varphi Y$.

First, note that $L_g\varphi = \varphi L_{\varphi(g)}$ follows from the homomorphism law for φ; for each $h \in G$,

$$hL_g\varphi = \varphi(gh) = \varphi(g)\varphi(h) = h\varphi L_{\varphi(g)}.$$

Then, applying the formula for hY, we find

$$gX\varphi = eL_gX\varphi = eXL_g\varphi = eX\varphi L_{\varphi(g)} = g\varphi Y.$$

□

The map $X \mapsto Y$ established by this lemma defines the usual Lie algebra homomorphism associated with the Lie group homomorphism φ. To check that the map we have defined respects the Lie bracket, notice that, where $X\varphi = \varphi X'$ and $Y\varphi = \varphi Y'$,

$$\begin{aligned}[X,Y]\varphi &= XY\varphi - YX\varphi = X\varphi Y' - Y\varphi X' \\ &= \varphi X'Y' - \varphi Y'X' = \varphi[X',Y'].\end{aligned}$$

Chapter 2

Hamiltonian Systems

Consider a fundamental problem in classical mechanics: the movement of a system of particles under the influence of forces. In one formulation, we speak of a pair of generalized position and momentum vectors $(q, p) \in \mathbb{R}^n \times \mathbb{R}^n$ whose time-evolution is governed by the differential equation

$$\begin{cases} \dot{q} = p \\ \dot{p} = F \end{cases}$$

for a force $F \in \mathbb{R}^n$ that is a function of time, p, and q. (We are supposing that the mass of our system is normalized, so that the velocity $\dot{q}$ and momentum p can simply be identified.)

Now, suppose that F is only a function of q—that is, it can be viewed as a vector field on our position space—and furthermore that it is a conservative vector field, with $F = \nabla\Phi$ for some function $\Phi\colon \mathbb{R}^n \to \mathbb{R}$. Then our mechanical system has a conserved quantity, called the **Hamiltonian energy**:

$$\begin{aligned} &H\colon \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R} \\ &H(q,p) = \frac{1}{2}\|p\|^2 - \Phi(q). \end{aligned}$$

Indeed, the partial derivatives of H with respect to p and q are

$$\nabla_p H = p = \dot{q}, \quad \nabla_q H = -F = -\dot{p}, \tag{2.1}$$

so, differentiating in time along a trajectory of the system, we find

$$\begin{aligned} &\frac{d}{dt}H = \langle \dot{q}, \nabla_q H\rangle + \langle \dot{p}, \nabla_p H\rangle \\ &= \langle \nabla_p H, \nabla_q H\rangle + \langle -\nabla_q H, \nabla_p H\rangle = 0. \end{aligned}$$

The simple fact of having a conserved quantity is already very useful; having fixed initial conditions, a conserved quantity allows us to reduce the dimensionality of our phase space by one. When $n = 1$, this reduction is enough to let us integrate our differential equation by quadrature. Indeed, let H_0 be the value of H along a given trajectory of our system. Then, we have a differential equation for the evolution of the variable q,

$$\frac{1}{2}\dot{q}^2 = \Phi(q) + H_0.$$

In an interval of time when, say, $\dot{q} > 0$, it is equivalent that

$$\dot{q} = \sqrt{2(\Phi(q) + H_0)},$$

which, by separation of variables, has the implicit solution

$$t = \int_0^q \frac{1}{\sqrt{2(\Phi(y) + H_0)}}\,dy.$$

However, there is more to investigate when $n > 1$!

For reasons which will soon become apparent, we will think of $\mathbb{R}^{2n}$ as the cotangent bundle $T^*\mathbb{R}^n$. Where $\pi\colon T^*\mathbb{R}^n \to \mathbb{R}^n$ is the bundle projection, the identification $T^*\mathbb{R}^n \cong \mathbb{R}^{2n}$ is given by

$$\begin{aligned}
&\psi\colon T^*\mathbb{R}^n \to \mathbb{R}^{2n}\\
&\psi(\lambda) = (q_1, \ldots, q_n, p_1, \ldots, p_n)(\lambda)\\
&= \left(\pi_1(\lambda), \ldots, \pi_n(\lambda), \left\langle \lambda, \frac{\partial}{\partial x_1}\right\rangle, \ldots, \left\langle \lambda, \frac{\partial}{\partial x_n}\right\rangle\right).
\end{aligned}$$

In general, the Hamiltonian energy has a significance besides its status as a conserved quantity: it actually characterizes the equations of movement. Indeed, by the two partials computed in (2.1), the equations of movement for our problem can be written in the *Hamiltonian form*,

$$\begin{cases} \dot{q} = \nabla_p H \\ \dot{p} = -\nabla_q H. \end{cases} \tag{2.2}$$

We are lead to define the **Hamiltonian vector field** of any scalar function $H \in C^\infty T^*\mathbb{R}^n$ as

$$\overrightarrow{H} = \frac{\partial H}{\partial p_i}\frac{\partial}{\partial q_i} - \frac{\partial H}{\partial q_i}\frac{\partial}{\partial p_i}.$$

When used in this way, H is called a **Hamiltonian**. A differential equation written in terms of a Hamiltonian in the manner of Equation (2.2) is called a **Hamiltonian system**. Note that $\overrightarrow{H}$ determines H up to a constant.

Interpreting $\overrightarrow{H}(G)$ as a bilinear operator n the functions H and G leads us to the following important definition.

Definition 12. *The* ***Poisson bracket*** *on* $T^*\mathbb{R}^n$ *is the bilinear operator*

$$\{-,-\}\colon C^\infty T^*\mathbb{R}^n \times C^\infty T^*\mathbb{R}^n \to C^\infty T^*\mathbb{R}^n$$

given by

$$\{H, G\} = \frac{\partial H}{\partial p_i}\frac{\partial G}{\partial q_i} - \frac{\partial H}{\partial q_i}\frac{\partial G}{\partial p_i}.$$

In terms of the Poisson bracket, a function G is a conserved quantity of the Hamiltonian vector field $\overrightarrow{H}$ exactly when $\{H,G\}=0$. Above, we checked that H is a constant of motion of our system, which means that $\{H,H\}=0$. Actually, this fact does not depend on the special form we chose for H; the Poisson bracket is skew-symmetric, since

$$\{H,G\} = \frac{\partial H}{\partial p_i}\frac{\partial G}{\partial q_i} - \frac{\partial H}{\partial q_i}\frac{\partial G}{\partial p_i} = -\frac{\partial G}{\partial p_i}\frac{\partial H}{\partial q_i} + \frac{\partial G}{\partial q_i}\frac{\partial H}{\partial p_i} = -\{G,H\}.$$

An interesting consequence is that a function G is a conserved quantity of the vector field $\overrightarrow{H}$ exactly when H itself is a conserved quantity of the vector field $\overrightarrow{G}$. This shows a significant relationship between conserved quantities and *symmetries* of our system, which we will return to at the end of this chapter.

Not every vector field on $T^*\mathbb{R}^n$ is realized as a Hamiltonian vector field. If $X = a_i\partial/\partial q_i + b_i\partial/\partial p_i$ is a vector field on a simply connected domain, the Poincaré lemma tells us that we can find a function H so that $\partial H/\partial p_i = a_i$ and $-\partial H/\partial q_i = b_i$ iff the following $\binom{2n}{2}$ integrability conditions hold.

$$\begin{aligned} &\frac{\partial a_i}{\partial p_j} = \frac{\partial a_j}{\partial p_i}, \quad \frac{\partial b_i}{\partial q_j} = \frac{\partial b_j}{\partial q_i} && 1 \le i < j \le n \\ &\frac{\partial a_i}{\partial q_j} = -\frac{\partial b_j}{\partial p_i} && 1 \le i,j \le n \end{aligned}$$

Generally, a vector field verifying these equations will be called *locally Hamiltonian*. Since $T^*\mathbb{R}^n$ is contractible, locally Hamiltonian vector fields coincide with Hamiltonian vector fields at the moment. Curiously, the integrability conditions above turn out to be exactly the right conditions for X to be a derivation with respect to the Poisson bracket![1]

Proposition 16. *The vector field X is locally Hamiltonian iff*

$$X\{G,H\} = \{XG,H\} + \{G,XH\}$$

for all $G,H \in C^\infty T^*\mathbb{R}^n$.

Proof. Set

$$X = a_i\frac{\partial}{\partial q_i} + b_i\frac{\partial}{\partial p_i},$$

so that we can develop $\{XG,H\}$ as

$$\begin{aligned} \{XG,H\} &= \left\{a_i\frac{\partial G}{\partial q_i} + b_i\frac{\partial G}{\partial p_i}, H\right\} \\ &= \frac{\partial}{\partial p_j}\left(a_i\frac{\partial G}{\partial q_i} + b_i\frac{\partial G}{\partial p_i}\right)\frac{\partial H}{\partial q_j} - \frac{\partial}{\partial q_j}\left(a_i\frac{\partial G}{\partial q_i} + b_i\frac{\partial G}{\partial p_i}\right)\frac{\partial H}{\partial p_j}. \end{aligned}$$

[1] For an explanation of this mysterious fact, the author found Part II of the book [4] very helpful. The answer lies in the construction of the *symplectic form* and in the methods of Cartan calculus. Unfortunately, none of this material ended up making its way into the thesis. Besides, doing calculations in coordinates is a good exercise!

This expands to a sum of 8 terms, 4 of which involve partial derivatives of the coefficients a_i and b_i. These are

$$\frac{\partial a_i}{\partial p_j}\frac{\partial G}{\partial q_i}\frac{\partial H}{\partial q_j}+\frac{\partial b_i}{\partial p_j}\frac{\partial G}{\partial p_i}\frac{\partial H}{\partial q_j}-\frac{\partial a_i}{\partial q_j}\frac{\partial G}{\partial q_i}\frac{\partial H}{\partial p_j}-\frac{\partial b_i}{\partial q_j}\frac{\partial G}{\partial p_i}\frac{\partial H}{\partial p_j}.$$

To arrive at the corresponding terms in the expansion of $\{G,XH\}$, we simply multiply by negative one and swap G with H. Collecting common factors in the sum of all 8 of these terms gives

$$\begin{aligned}&\left(\frac{\partial a_i}{\partial p_j}-\frac{\partial a_j}{\partial p_i}\right)\frac{\partial G}{\partial q_i}\frac{\partial H}{\partial q_j}+\left(\frac{\partial b_i}{\partial p_j}+\frac{\partial a_j}{\partial q_i}\right)\frac{\partial G}{\partial p_i}\frac{\partial H}{\partial q_j}\\ &-\left(\frac{\partial a_i}{\partial q_j}+\frac{\partial b_j}{\partial p_i}\right)\frac{\partial G}{\partial q_i}\frac{\partial H}{\partial p_j}-\left(\frac{\partial b_i}{\partial q_j}-\frac{\partial b_j}{\partial q_i}\right)\frac{\partial G}{\partial p_i}\frac{\partial H}{\partial p_j}.\end{aligned}$$

Running G and H over linear functionals on $T^*\mathbb{R}^n$ shows that this sum vanishes iff X is locally Hamiltonian.

The 8 terms we have so far neglected from the expansion of $\{XG,H\}+\{G,XH\}$ are

$$\begin{aligned}&\left(a_i\frac{\partial^2 G}{\partial q_i\partial p_j}+b_i\frac{\partial^2 G}{\partial p_i p_j}\right)\frac{\partial H}{\partial q_j}-\left(a_i\frac{\partial^2 G}{\partial q_i\partial q_j}+b_i\frac{\partial^2 G}{\partial p_i\partial q_j}\right)\frac{\partial H}{\partial p_j}\\ &+\frac{\partial G}{\partial p_i}\left(a_j\frac{\partial^2 H}{\partial q_j\partial p_i}+b_j\frac{\partial^2 H}{\partial p_j\partial p_i}\right)-\frac{\partial G}{\partial q_i}\left(a_j\frac{\partial H}{\partial q_j\partial q_i}+b_j\frac{\partial H}{\partial p_j\partial q_i}\right),\end{aligned}$$

which can be factored as

$$\left(a_i\frac{\partial}{\partial q_i}+b_i\frac{\partial}{\partial p_i}\right)\left(\frac{\partial G}{\partial p_i}\frac{\partial H}{\partial q_i}-\frac{\partial G}{\partial q_i}\frac{\partial H}{\partial p_i}\right)=X\{G,H\}.$$

□

In particular, this also proves the Jacobi identity for the Poisson bracket; for all $F,G,H\in C^\infty T^*M$,

$$\{F,\{G,H\}\}=\{\{F,G\},H\}+\{G,\{F,H\}\}.$$

Since we have already checked bilinearity and antisymmetry, we conclude that the Poisson bracket is a Lie bracket. By a general property of adjoint representations of Lie algebras, we find that the Poisson bracket is related to the usual Lie bracket on $D(T^*\mathbb{R}^n)$ by the formula

$$[\overrightarrow{G},\overrightarrow{H}]=\overrightarrow{\{G,H\}}. \tag{2.3}$$

Indeed, we check explicitly that

$$\begin{aligned}&[\overrightarrow{G},\overrightarrow{H}]F=\{G,\{H,F\}\}-\{H,\{G,F\}\}\\ &=-\{G,\{F,H\}\}-\{\{F,G\},H\}=-\{F,\{G,H\}\}=\overrightarrow{\{G,H\}}F.\end{aligned}$$

Another basic consequence of the Jacobi identity for the Poisson bracket is *Poisson's theorem*, which tells us that the conserved quantities of a Hamiltonian system are closed under the Poisson bracket.

Theorem 7 (Poisson's theorem)**.** *Let $H \in C^\infty T^*\mathbb{R}^n$. Then the set*

$$C_H = \{F \in C^\infty T^*\mathbb{R}^n : \{H,F\} = 0\}$$

is a Lie sub-algebra of $C^\infty T^*\mathbb{R}^n$ with respect to the Poisson bracket.

Proof. Let $F, G \in C_H$. Clearly, any linear combination of F and G remains in C_H. Furthermore,

$$\{H,\{F,G\}\} = \{\{H,F\},G\} + \{F,\{H,G\}\} = \{0,G\} + \{F,0\} = 0$$

so we find that $\{F,G\}$ belongs to C_H as well. □

So far, we have defined the Poisson bracket on the cotangent bundle of $\mathbb{R}^n$ and performed calculations in coordinates. In our next section, we will show that the Poisson bracket is *coordinate invariant*, and can be defined naturally on any cotangent bundle.

2.1 Coordinate-Independence of the Poisson Bracket

In differential geometry, we think of a manifold as having no particular distinguished coordinate system. Instead, we say that it has a family of local coordinate systems—an atlas—which, together, describe its smooth structure. An operation associated with the manifold is characterized by its representation in each coordinate chart, but, for such coordinate representations to define a coherent operation on the manifold, we must check that they respect coordinate reparameterization. We need *coordinate-independence*.

To illustrate this idea, let us consider an operation that is *not* coordinate-independent: the gradient, ∇, which assigns a function $f\colon \mathbb{R}^n \to \mathbb{R}$ to a vector field

$$\nabla f = \frac{\partial f}{\partial x_i}\frac{\partial}{\partial x_i}.$$

Suppose M is a manifold admitting a global coordinate system, $\alpha\colon M \to U \subseteq \mathbb{R}^n$. The tangent bundle TM is then identified with TU by the differential of α, and we can define the gradient of a function $f\colon N \to \mathbb{R}$ with respect to our coordinate system[2] as

$$\nabla_\alpha f = (d\alpha)^{-1} \circ \nabla(f \circ \alpha^{-1}) \circ \alpha.$$

However, if $\beta\colon M \to V \subseteq \mathbb{R}^n$ is another coordinate system, the corresponding gradient operator ∇_β might not be the same! Let us check this.

Write $g = f \circ \beta^{-1}$ and $\varphi = \beta \circ \alpha^{-1}$, as in the following diagram.

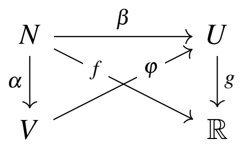

[2]Throughout this discussion, we temporarily revert to the classical perspective on vector fields and smooth maps.

We will compare $\nabla_\alpha f$ and $\nabla_\beta f$ as vector fields over V. We compute that

$$\begin{cases} d\beta \circ \nabla_\beta f \circ \beta^{-1} = \nabla g, \text{ and} \\ d\beta \circ \nabla_\alpha f \circ \beta^{-1} = d\varphi \circ \nabla(g \circ \varphi) \circ \varphi^{-1}. \end{cases}$$

In coordinates,

$$\begin{cases} \nabla g = \dfrac{\partial g}{\partial x_i} \dfrac{\partial}{\partial x_i}, \text{ and} \\ d\varphi \circ \nabla(g \circ \varphi) \circ \varphi^{-1} = d\varphi \left(\dfrac{\partial g}{\partial x_j} \dfrac{\partial \varphi_j}{\partial x_i} \dfrac{\partial}{\partial x_i} \right) = \dfrac{\partial \varphi_i}{\partial x_j} \dfrac{\partial g}{\partial x_k} \dfrac{\partial \varphi_k}{\partial x_j} \dfrac{\partial}{\partial x_i}. \end{cases}$$

Running g over the coordinate functions $x_1, \dots, x_n$ shows that ∇_α will only coincide with ∇_β if

$$\frac{\partial \varphi_i}{\partial x_j} \frac{\partial \varphi_k}{\partial x_j} = \delta_{i,k},$$

meaning that φ is a local isometry.

In the language of category theory, coordinate-independence can be viewed as a naturality property. Let NGrp be the category of diffeomorphisms between n-dimensional real domains. We will define two contravariant functors from NGrp to Set. The first is given by the coordinate algebra functor defined in the previous section, restricted from Diff to the subcategory NGrp, and post-composed by the inclusion Alg $\to$ Set. (We still denote this functor with C^∞.) The second is a functor $\mathscr{X}$, which will define temporarily in the following way.

Definition 13. *The vector field functor* $\mathscr{X}\colon$ NGrp $\to$ Set *takes a real domain* U *to its set* $\mathscr{X}U$ *of vector fields and a diffeomorphism* $\varphi\colon V \to U$ *to the pushforward operation*

$$\begin{gathered} \mathscr{X}\varphi\colon \mathscr{X}U \to \mathscr{X}V \\ (\mathscr{X}\varphi)(X) = (d\varphi)^{-1} \circ X \circ \varphi. \end{gathered}$$

From the point of view of these functors, what have shown above is that the maps $\nabla_\alpha\colon C^\infty U \to \mathscr{X}U$ defined by the gradient *do not* constitute a natural transformation from C^∞ to $\mathscr{X}$. Indeed, for a morphism $\varphi\colon V \to U$ in NGrp, the commutativity condition for the diagram

$$\begin{array}{ccc} C^\infty U & \xrightarrow{\;C^\infty\varphi\;} & C^\infty V \\ \big\downarrow{\scriptstyle\nabla_\alpha} & & \big\downarrow{\scriptstyle\nabla_\beta} \\ \mathscr{X}U & \xrightarrow{\;\mathscr{X}\varphi\;} & \mathscr{X}V \end{array}$$

can be written in classical notation, for an arbitrary element $g \in C^\infty U$, as

$$\begin{aligned} & (\mathscr{X}\varphi \circ \nabla_\alpha)(g) = (\nabla_\beta \circ C^\infty\varphi)(g) \\ \iff \quad & (d\varphi)^{-1} \circ \nabla g \circ \varphi = \nabla(g \circ \varphi) \\ \iff \quad & \nabla g = d\varphi \circ \nabla(g \circ \varphi) \circ \varphi^{-1}. \end{aligned}$$

This is exactly the equation that we investigated above and proved to be false in general. In summary, the gradient is not a *natural* operation on a manifold without a metric!

Although the Hamiltonian vector field construction on $T^*\mathbb{R}^n$ also maps functions to vector fields, it turns out to be coordinate-invariant.[3] Following the discussion above, we just need to check that the assignment of Hamiltonian vector fields is natural with respect to reparameterization of the underlying domain U. This will show that the coordinate algebra of any cotangent bundle is equipped with a natural Poisson bracket.

Let Diff* be the lluf[4] subcategory of Diff whose morphisms are diffeomorphisms. We now a introduce third functor, $T^*\colon \text{Diff}^* \to \text{Diff}^*$, which defines a canonical way for a diffeomorphism of manifolds to lift to a diffeomorphism of cotangent bundles.

Definition 14. *The* ***cotangent bundle functor*** *takes a manifold M to the cotangent bundle T^*M and a diffeomorphism $\varphi\colon M \to N$ to the* ***point transformation*** *defined by*

$$T^*\varphi\colon T^*M \to T^*N$$
$$T^*\varphi(\lambda) = (d\varphi^{-1}_{\varphi(\pi(\lambda))})^*(\lambda).$$

It is straightforward to check the functorial properties $T^*\mathrm{id}_M = \mathrm{id}_{T^*M}$ and $T^*(f \circ g) = T^*f \circ T^*g$. Note that, unlike what happens for the tangent bundle, there is no reasonable way to turn extend this functor to smooth maps in general. Now, consider T^* restricted to an endofunctor on NGrp. We state and prove our naturality property.

Proposition 17. *The assignment of Hamiltonian vector fields*

$$\eta_U\colon C^\infty T^*U \to \mathscr{X}T^*U$$
$$\eta_U(H) = \overrightarrow{H}$$

behaves naturally with respect to reparameterization of the domain U. That is, the maps η_U defined in this way form the sections of a natural transformation

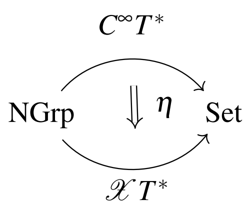

[3] This is quite surprising, given only what we have explained up to now! Again, the "mystery" could be explained by the tools of symplectic geometry—in this case, by the coordinate-invariant construction of the Liouville one-form and the symplectic form, and the definition of the Poisson bracket in terms of the latter.

[4] A lluf subcategory retains all objects from its supercategory. "Lluf" is "full" spelled backward.

Proof. Let $\varphi\colon U \to V$ be a reparameterization and let $\psi\colon T^*U \to T^*V$ be the corresponding point transformation. Our claim is that the diagram

$$\begin{array}{ccc} C^\infty T^*U & \xrightarrow{C^\infty \psi^{-1}} & C^\infty T^*V \\ \downarrow \eta_U & & \downarrow \eta_V \\ \mathscr{X} T^*U & \xrightarrow{\mathscr{X} \psi^{-1}} & \mathscr{X} T^*V \end{array}$$

commutes.

With respect to the standard coordinate system for $T^*\mathbb{R}^n$, the Jacobian matrix for ψ can be expressed in terms of the Jacobian J of φ as

$$L = \begin{bmatrix} J & 0 \\ 0 & (J^{-1})^T \end{bmatrix}.$$

Meanwhile, define the $2n \times 2n$ matrix

$$S = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix},$$

so that a Hamiltonian vector field $\overrightarrow{H}$ can be expressed as $\overrightarrow{H} = S\nabla f$. We can use these matrices to express the images of an element $H \in C^\infty T^*\mathbb{R}^n$ under the upper and lower paths of our diagram:

$$\begin{cases} (\mathscr{X}\psi^{-1} \circ \eta_U) f = d\psi \circ \overrightarrow{f} \circ \psi^{-1} = LS\nabla f, \text{ while} \\ (\eta_V \circ C^\infty \psi^{-1}) f = \overrightarrow{f \circ \psi^{-1}} = S(L^{-1})^T \nabla f. \end{cases}$$

However,

$$\begin{aligned} LS &= \begin{bmatrix} J & 0 \\ 0 & (J^{-1})^T \end{bmatrix} \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix} = \begin{bmatrix} 0 & J \\ -(J^{-1})^T & 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix} \begin{bmatrix} (J^{-1})^T & 0 \\ 0 & J \end{bmatrix} = S(L^{-1})^T, \end{aligned}$$

so the diagram indeed commutes. □

This justifies that the cotangent bundle T^*M of an arbitrary manifold (without a distinguished coordinate system) is equipped with a Poisson bracket in a natural way.

Definition 15. *The **Poisson bracket** on T^*M is the bilinear operator*

$$\{-,-\}\colon C^\infty T^*M \times C^\infty T^*M \to C^\infty T^*M$$

given in coordinates by Definition 12.

2.2 The Hamiltonian Lift

The properties of the Poisson bracket on T^*M we have just defined follow directly from the properties we checked earlier when working in coordinates. Altogether, they mean that the Poisson bracket on T^*M equips $C^\infty T^*M$ with the structure of a *Poisson algebra*.

Definition 16. *A **Poisson algebra** is a (unital, commutative) algebra A equipped with a **Poisson bracket**, denoted* $\{-,-\}$*, which verifies the following conditions for any* $F,G,H \in A$*:*

1. *Linearity:* $\{H,\alpha G+\beta F\} = \alpha\{H,G\}+\beta\{H,F\}$.
2. *Leibniz identity:* $\{H,FG\} = \{H,F\}G+F\{H,G\}$.
3. *Antisymmetry:* $\{H,G\} = -\{G,H\}$.
4. *Jacobi identity:* $\{H,\{G,F\}\} = \{\{H,G\},F\}+\{G,\{H,F\}\}$.

Properties 1 and 2 together mean that $\{H,-\}$ is a derivation of A. Properties 1, 3, and 4 together mean that $\{-,-\}$ is a Lie bracket over the vector space structure of A. It can also be checked that the Poisson bracket is compatible with our topologization of $C^\infty T^*M$.

Proposition 18. *The Poisson bracket is a continuous bilinear operator on* $C^\infty T^*M$*.*

Now, a new algebraic structure comes with a new class of maps preserving that structure. Homomorphisms and derivations of the Poisson algebras $C^\infty T^*M$ correspond to smooth maps and vector fields over T^*M with special geometric properties.

First, consider the case of vector fields.

Definition 17. *Let M be a manifold. A vector field X on* T^*M *is a **Poisson vector field** when it yields a derivation of the Poisson algebra* $C^\infty T^*M$*. That is, X must respect the equation*

$$X\{F,G\} = \{XF,G\}+\{F,XG\}$$

for all $F,G \in C^\infty T^*M$*. We denote the space of Poisson vector fields on a manifold* T^*M *by* $D_{\mathrm{Poiss}}(T^*M)$*. Meanwhile, X is a **Hamiltonian vector field** when there is some element* $H \in C^\infty T^*M$ *for which* $X = \{H,-\}$*.*

Poisson vector fields are the generalization of what, in the previous section, we called "locally Hamiltonian vector fields." Note that, by the properties of a Poisson algebra, every Hamiltonian vector field is a Poisson vector field, and both $D_{\mathrm{Poiss}}(T^*M)$ and the space of Hamiltonian vector fields will be closed under the Lie bracket. By Proposition 16, we also have a first-order partial differential equation that characterizes Poisson vector fields in coordinates. This will come in handy in a moment.

Next, we consider the special "structure-preserving" smooth maps between cotangent bundles.

Definition 18. *A smooth map* $Q\colon T^*M \to T^*N$ *is a **Poisson map** when it yields a homomorphism of Poisson algebras. That is, Q must respect the equation*

$$Q(\{F,G\}) = \{Q(F),Q(G)\}$$

for all $F,G \in C^\infty T^*M$*. The subgroup of* $\mathrm{Aut}(T^*M)$ *comprised by Poisson maps is denoted* $\mathrm{Aut}_{\mathrm{Poiss}}(T^*M)$*.*

A careful look at Proposition 17 shows that every point transformation is a Poisson map.

Proposition 19. *When $Q\colon M \to N$ is a diffeomorphism, the point transformation $T^*Q\colon T^*M \to T^*N$ is a Poisson map.*

Proof. It is enough to prove this in the special case that M and N are real domains. So, let Q be a morphism in NGrp. In Proposition 17, we proved that the diagram

$$\begin{array}{ccc} C^\infty T^*N & \xrightarrow{C^\infty T^*Q} & C^\infty T^*M \\ \downarrow \eta_U & & \downarrow \eta_V \\ \mathscr{X} T^*N & \xrightarrow{\mathscr{X} T^*Q} & \mathscr{X} T^*M \end{array}$$

commutes, where η_U and η_V send smooth maps F to their Hamiltonian vector fields $\{F,-\}$. Now, let $F \in C^\infty T^*N$, and write $P = C^\infty T^*Q$. Applying the definition of the functors $\mathscr{X}$ and C^∞ gives

$$(\mathscr{X} T^*Q \circ \eta_U)(F) = P\overrightarrow{F}P^{-1}, \text{ and}$$
$$(\eta_V \circ C^\infty T^*Q)(F) = \overrightarrow{P(F)}.$$

Since $P\overrightarrow{F}P^{-1} = \overrightarrow{P(F)}$, we conclude that, for all $G' \in C^\infty T^*M$,

$$P(\{F, P^{-1}(G')\}) = (P\overrightarrow{F}P^{-1})(G') = (\overrightarrow{P(F)})(G') = \{P(F), G'\}.$$

Putting $G' = P(G)$ proves that $P(\{F,G\}) = \{P(F), P(G)\}$, as desired. □

In the previous chapter, we embedded the group $\mathrm{Aut}(M)$ of self-diffeomorphisms of M into a Fréchet space of continuous linear operators on $C^\infty M$ and showed that the tangent space $T_{\mathrm{id}}\,\mathrm{Aut}(M)$ is a subset of $D(M)$. Now we ask: what is the tangent space $T_{\mathrm{id}}\,\mathrm{Aut}_{\mathrm{Poiss}}(M)$? The relationship between automorphisms and derivations we introduced at the beginning of our first chapter arises again in the following proposition.

Proposition 20. *Where $\mathrm{Aut}_{\mathrm{Poiss}}(T^*M)$ is regarded as a subset of $\mathrm{End}_{\mathrm{TVS}}\, C^\infty T^*M$,*

$$T_{\mathrm{id}}\,\mathrm{Aut}_{\mathrm{Poiss}}(T^*M) \subseteq D_{\mathrm{Poiss}}(T^*M).$$

Proof. Let $X \in T_{\mathrm{id}}\,\mathrm{Aut}_{\mathrm{Poiss}}(T^*M)$. Then there exists a curve Q_t in $\mathrm{Aut}_{\mathrm{Poiss}}(T^*M)$ with $Q_0 = \mathrm{id}$ and $\dot{Q}_0 = X$. Let $F, G \in C^\infty T^*M$. By hypothesis that Q_t is a Poisson map,

$$Q_t(\{F,G\}) = \{Q_t(F), Q_t(G)\}.$$

By Proposition 18, the generalized product rule can be used to differentiate the right hand side of this equality at 0. We conclude that

$$X\{F,G\} = \dot{Q}_0(\{F,G\}) = \{\dot{Q}_0(F), G\} + \{F, \dot{Q}_0(G)\} = \{XF, G\} + \{F, XG\}.$$

□

It can also be shown that the exponential of a complete Poisson vector field is a Poisson map.

Now, the point transformation functor gives a group homomorphism from $\mathrm{Aut}(M)$ to $\mathrm{Aut}_{\mathrm{Poiss}}(T^*M)$. From the theory of Lie groups, we know that a smooth group homomorphism $\varphi\colon G \to H$ gives rise to a homomorphism $h\colon \mathfrak{g} \to \mathfrak{h}$ of Lie algebras and that the exponential map verifies a naturality property with this functor; specifically, the following diagram commutes.

$$\begin{array}{ccc} G & \xrightarrow{\varphi} & H \\ {\scriptstyle \exp_{\mathfrak{g}}}\uparrow & & \uparrow{\scriptstyle \exp_{\mathfrak{h}}} \\ \mathfrak{g} & \xrightarrow{h} & \mathfrak{h} \end{array}$$

Does the homomorphism T^* give rise to a similar map from vector fields on M to Poisson vector fields on T^*M? Such a map turns out to exist and is called the *Hamiltonian lift*.

Definition 19. *Let $X \in D(M)$. We define the* ***Hamiltonian of*** *X as the smooth map*

$$H_X : T^*M \to \mathbb{R}$$
$$H_X(\lambda) = \langle \lambda, X \rangle.$$

The ***Hamiltonian lift*** *$\overline{X}$ of X is the Hamiltonian vector field $\overrightarrow{H}_X$.*

The following theorem is reminiscent of the naturality property that we described above for the exponential maps of Lie groups.

Theorem 8. *Suppose X_t is a vector field in the conditions of Theorem 3. Then*

$$T^*\left(\overrightarrow{\exp}\int_0^t X_\tau \, d\tau\right) = \overrightarrow{\exp}\int_0^t \overline{X}_\tau \, d\tau.$$

Our proof of this fact will be somewhat indirect. First of all, we state the following technical result (without proof) which tells that there exists *some* vector field Y_t for which

$$T^*\left(\overrightarrow{\exp}\int_0^t X_\tau \, d\tau\right) = \overrightarrow{\exp}\int_0^t Y_\tau \, d\tau.$$

Proposition 21. *Let Q be a family of diffeomorphisms. If Q is differentiable at t, then T^*Q is also differentiable at t. If Q is Lipschitz, then T^*Q is Lipschitz.*

Indeed, where $Q_t = \overrightarrow{\exp}\int_0^t X_\tau \, d\tau$ and $P_t = T^*Q_t$, it is enough to put $Y_t = P_t^{-1}\dot{P}_t$. To prove Theorem 8, it only remains to check that $Y_t = \overline{X}_t$ almost everywhere. This can be reduced to the following lemma, which, in a sense, describes the "differential of the cotangent bundle functor."

Lemma 3. *Suppose Q is a family of diffeomorphisms with $Q_0 = \mathrm{id}$ and $\dot{Q}_0 = X$. Then*

$$\left(\frac{d}{dt}\right)_0 T^*Q_t = \overline{X}.$$

To conclude that $Y_{t_0} = \overline{X}_{t_0}$ at a differentiability point $t_0 \neq 0$ of Q, note that

$$\begin{cases} X_{t_0} = Q_{t_0}^{-1}\dot{Q}_{t_0} = \left(\dfrac{d}{dt}\right)_0 U_t, \\ Y_{t_0} = P_{t_0}^{-1}\dot{P}_{t_0} = \left(\dfrac{d}{dt}\right)_0 T^*U_t \end{cases}$$

if we define $U_t = Q_{t_0}^{-1}Q_{t_0+t}$. We now proceed to the proof of Lemma 3, which is the substantial part of our argument.

Proof. Define the family $P_t = T^*Q_t$, and let $Y = \dot{P}_0$. To prove that $Y = \overline{X}$, we will deduce the expression of Y in coordinates. To simplify, we will assume that M admits a global coordinate chart $(x_1,\dots,x_n)$. An analogous proof is possible without this simplification.

Let $(q_1,\dots,q_n,p_1,\dots,p_n)$ be the global coordinate chart for T^*U associated with $(x_1,\dots,x_n)$. We will express the vector fields X and Y in coordinates as

$$Y = a_i\frac{\partial}{\partial q_i} + b_i\frac{\partial}{\partial p_i}, \quad X = c_i\frac{\partial}{\partial x_i}$$

for certain functions $a_i, b_i \in C^\infty T^*M$ and $c_i \in C^\infty M$. Our first step is to determine the coefficients a_i. By the properties of point transformations, we know that $P_t\pi = \pi Q_t$ holds for all t. Differentiating this equation gives us the relation

$$Y\pi = \dot{P}_0\pi = \pi\dot{Q}_0 = \pi X,$$

from which it follows that

$$a_i = Y(q_i) = (Y\pi)(x_i) = (\pi X)(x_i) = \pi(c_i).$$

Next, we consider the coefficients b_i. We know that

$$b_i = Y(p_i) = \left(\frac{d}{dt}\right)_0 P_t(p_i).$$

Furthermore, since P_t is a bundle homomorphism for each t, $P_t(p_i)$ is a linear function of the coefficients p_i. We conclude that b_i is also linear, taking the form

$$b_i = d_{i,j}p_j$$

for certain functions $d_{i,j}$ which only depend on the coordinates q_i.

Since P_t is a Poisson map for every t, Proposition 20 tells us that Y will be a Poisson vector field. To compute $d_{i,j}$, we recall one of the relations characterizing a Poisson vector field in coordinates:

$$\frac{\partial a_i}{\partial q_j} = -\frac{\partial b_j}{\partial p_i}.$$

Using this relation, we find that

$$d_{i,j} = \frac{\partial b_i}{\partial p_j} = -\frac{\partial a_j}{\partial q_i} = -\frac{\partial \pi(c_j)}{\partial q_i}.$$

Altogether, we have determined that

$$Y = \pi(c_i)\frac{\partial}{\partial q_i} - p_j \frac{\partial \pi(c_j)}{\partial q_i}\frac{\partial}{\partial p_i}.$$

In coordinates, the Hamiltonian H_X has the expression $H_X = \pi(c_i)p_i$, so the expression we have found for Y is exactly the Hamiltonian lift of X. □

2.3 Symmetry and Conservation Laws

Expressing a smooth dynamical system as a Hamiltonian vector field $\overrightarrow{H}$ on the cotangent bundle T^*M provides us with a variety of interesting results about its behavior. For example, *Liouville's theorem* tells us that $\overrightarrow{H}$ preserves a certain volume form on T^*M. Such a flow is called *incompressible*. This is a highly significant property; for example, it has been argued that the ergodic nature of incompressible flows justifies a version of the second law of thermodynamics for Hamiltonian systems [9]. The Hamiltonian formalism can also lead to the explicit integration of dynamical systems; if a Hamiltonian H on an $2n$-dimensional cotangent bundle admits $n-1$ constants of motion $C_1, \dots, C_{n-1}$ which also satisfy $\{C_i, C_j\} = 0$ for all $1 \leq i, j \leq n-1$, then the *angle-action variables* give us an explicit way to integrate the flow of $\overrightarrow{H}$ by quadrature.

In this last section, we return to the physical perspective with which we began the chapter and show one simple application of the Hamiltonian formalism: a derivation of the relationship between the *spatial symmetry* of a mechanical system and its *conservation laws*. As a concrete example, we will derive the conservation of linear and angular momentum in a system of interacting point masses.

Consider a system of n point-mass bodies in $\mathbb{R}^k$, of masses $(m_1, \dots, m_n)$, and suppose each pair of masses interacts according to a central, radially symmetric pair of equal and opposing forces. We can describe the dynamics of this system with $\binom{n}{2}$ radial potential functions $F_{i,j} \colon \mathbb{R}_{>0} \to \mathbb{R}$, one for each unordered pair $\{i, j\}$ of bodies. Where (q_i, p_i) are the position and momentum vectors of each point mass, the Hamiltonian $H \colon T^*\mathbb{R}^{nk} \to \mathbb{R}$ has the expression

$$H = \sum_{i=1}^{n} \frac{1}{2m_i}\|p_i\|^2 + \sum_{\substack{1 \leq i,j \leq n \\ i \neq j}} m_i m_j F_{i,j}(\|q_i - q_j\|).$$

The corresponding equations of movement for our system are

$$\begin{cases} \dot{q}_i &= \dfrac{\partial H}{\partial p_i} = \dfrac{p_i}{m_i}, \\ \dot{p}_i &= \dfrac{\partial H}{\partial q_i} = \displaystyle\sum_{j \neq i} m_i m_j \frac{\partial F_{i,j}(\|q_i - q_j\|)}{\partial q_i} = \sum_{j \neq i} m_i m_j \frac{F'_{i,j}(\|q_i - q_j\|)(q_i - q_j)}{\|q_i - q_j\|}. \end{cases}$$

Our Hamiltonian H is evidently invariant under isometries of $\mathbb{R}^k$. The next theorem shows us a simple and general way to use this symmetry to generate conserved quantities of our mechanical system.

Theorem 9 (Noether's theorem for spatial symmetries)**.** *Let $H \in C^\infty T^*\mathbb{R}^n$ be a Hamiltonian and let $X \in D(\mathbb{R}^n)$ be a vector field so that $(T^*\exp(tX))(H) = H$ for all $t \in \mathbb{R}$. Then H_X is a conserved quantity of $\overrightarrow{H}$.*

Proof. Since $(T^*\exp(tX))(H) = H$ constantly,

$$\frac{d}{dt}(T^*\exp(tX))H = 0.$$

On the other hand, using Theorem 8 to differentiate $T^*\exp(tX)$ gives

$$\frac{d}{dt}(T^*\exp(tX))H = \overline{X}(H) = \{H_X, H\}.$$

Combining these two equations and applying antisymmetry of the Poisson bracket lets us conclude that $\overrightarrow{H}(H_X) = \{H, H_X\} = 0$, as desired. □

Let us first apply Theorem 9 to translational symmetries. For an arbitrary vector $v \in \mathbb{R}^k$, let

$$X(v) = \left\langle v, \sum_{i=1}^{n} \frac{\partial}{\partial q_i} \right\rangle.$$

The flow $T^*\exp(tX(v))$ gives a translation

$$\begin{cases} q_i \mapsto q_i + tv \\ p_i \mapsto p_i \end{cases}$$

of the phase space $T^*\mathbb{R}^{nk}$, which patently preserves H. We conclude that our system has the conserved quantity

$$H_{X(v)} = \left\langle v, \sum_{i=1}^{n} p_i \right\rangle.$$

Since v is arbitrary, this proves that the sum $\sum_i p_i$ is conserved. Physically, we conclude that our mechanical system respects the conservation of linear momentum.

Conservation of angular momentum follows, analogously, from the invariance of our Hamiltonian under rotation. Let V be a $k \times k$ antisymmetric matrix, and define another vector field

$$Y(V) = \sum_{i=1}^{n} \left\langle Vq_i, \frac{\partial}{\partial q_i} \right\rangle.$$

The flow $T^*\exp(tY(V))$ now gives a rotation

$$\begin{cases} q_i \mapsto e^{tV} q_i \\ p_i \mapsto e^{-tV} p_i. \end{cases}$$

Since e^{tV} is an isometry for all t, it is easy to see that a transformation of this sort also leaves our system's Hamiltonian invariant. The conserved quantity derived from $Y(V)$ is

$$H_{Y(V)} = \sum_{i=1}^{n} p_i^T V q_i.$$

Running V over the space of anti-symmetric matrices gives us the tensor-valued conserved quantity $\sum_i p_i \wedge q_i$—an element of $\Lambda^2(\mathbb{R}^k)$ representing the angular momentum of our system. In $\mathbb{R}^3$ this quantity can be expressed with the cross product $p_i \times q_i$ and gives a total of 3 conserved scalar quantities, but in general the angular momentum will have $\binom{n}{2}$ scalar components.

Chapter 3

Some Optimal Control Theory

Let M be a manifold and consider a dynamical system occupying the state $p(t) \in M$ at time t.[1] The condition that p should be an integral curve of some vector field $X \in D(M)$ gives us a smooth, deterministic law for our system's evolution over time. Now, suppose we have some ability to *control* our system. We can formalize this in a very general way by replacing the integral curve equation

$$\dot{p}(t) = p(t)X$$

with an *inclusion*

$$\dot{p}(t) \in V(p(t))$$

for a certain set-valued function V. The trajectory p might no longer be uniquely determined from an initial condition $p(0) = P_0$; instead, we may have the ability to *choose* a trajectory out of our starting point from a set of admissible trajectories.

Definition 20. *A **control system** (M,V) is a manifold M paired with a function V which takes each point $p \in M$ to a subset*

$$V(p) \subseteq T_pM,$$

*called a **velocity set**.*

Definition 21. *Let I be a compact interval. A Lipschitz curve $p\colon I \to M$ is a **trajectory** of the control system (M,V) when, for almost every t, $\dot{p}(t) \in V(p(t))$.*

Although we have defined a control system here in terms of its velocity sets, it is also customary and useful to define a control system as a "smooth dynamical system with controllable parameters." For us, this is a *parameterization* of a control system.

Definition 22. *A **parameterization** (U,X) of a control system (M,V) is a set $U \subseteq \mathbb{R}^n$ and an assignment of controls $u \in U$ to vector fields $X_u \in D(M)$ so that*

$$V(p) = \{pX_u \colon u \in U\}.$$

[1] In this section, it will be clearer to use expressions like $p(t)$ rather than p_t to denote time-dependence.

The elements of U are called **control parameters**, and the vector field X_u is the **control-dependent vector field**. Of course, a control system can be defined by a parameterization, so we will speak in the following of **parameterized control systems** (M,U,X).

Remark 6. Note that not every "control system," as we have defined them, can be parameterized. For a trivial counterexample, observe that V can be chosen so that there is no smooth vector field X verifying

$$pX \in V(p)$$

for every $p \in M$. Furthermore, there be more than one way to parameterize a control system. Indeed, when (U,X) is a parameterization of a control system (M,V) and when

$$\varphi \colon M \times U \to U$$

is a map so that $\varphi(p,-)$ is a bijection of U for each $p \in M$, the control-dependent vector field

$$Y_u(p) = X_{\varphi(p,u)}(p)$$

is also a parameterization of (M,V), assuming Y_u is a smooth vector field for each $u \in U$. (In this case, the two parameterizations are called **feedback equivalent**.) In what follows, we will assume that our control system is parameterizable, and it will turn out to be helpful to fix an arbitrary parameterization.

Let P_0 and P_1 be points on M. When there is at least one trajectory of the control system (M,V) traveling from P_0 to P_1, there will typically be an infinite set of such trajectories. One may pose the problem of selecting a trajectory from this set that minimizes some "cost function." For example, we may want to find a trajectory that minimizes the time spent in transit. This is the problem of *time-optimal control*.

Let us begin with a control system where time-optimal control is easy to understand.

Example 1 (Movement on the line)**.** *Let $M = \mathbb{R}$ be the real line, let $U = [-1,1]$, and let X be the control-dependent vector field $X_u = u\partial/\partial x$.*

A trajectory of this parameterized control system is simply a Lipschitz path on $\mathbb{R}$ with velocity bounded by 1. In this case, it is easy to see how every pair of points on $\mathbb{R}$ is connected by a unique time-optimal trajectory. Our next example is a generalization of Example 1 where time-optimal control is no longer so straightforward.

Example 2 (Movement on a Riemannian manifold)**.** *Let (M,g) be a Riemannian manifold. Define velocity sets by $V(p) = \{v \in T_pM \colon g(v,v) \le 1\}$.*

A time-optimal trajectory of this system is the same as a geodesic on M. We recall from Riemannian geometry that geodesics are smooth curves and can be characterized, locally, as the solutions to a certain second-order differential equation.

Finally, we present a standard example from the control theory literature.

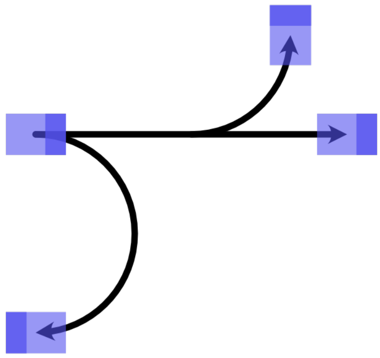

Fig. 3.1 Some Dubins paths of length π.

Example 3 (Movement of a Dubins car). *Let $M = S^1 \times \mathbb{R}^2$ be parameterized by coordinates (x_1, x_2, θ), let $U = [-1,1]$, and let*

$$X_u = \cos(\theta)\frac{\partial}{\partial x_1} + \sin(\theta)\frac{\partial}{\partial x_2} + u\frac{\partial}{\partial \theta}.$$

A Dubins car is a car, driving on a plane, with a steering apparatus controlled by the parameter u. The trajectory of a Dubins car projected onto (x_1, x_2)-space, called a Dubins path, is simply a smooth path on the plane with curvature between -1 and 1. As we will show later, a Dubins car taking a time-optimal trajectory will, at any given moment, either be going straight or else steering in one direction or the other as quickly as possible.

The geodesic problem and the time-optimal control problem for the Dubins car are quite different problems with qualitatively different solutions. However, as we will see, their solutions can be derived from a single result called the *Pontryagin maximum principle*.[2] In the last few pages of this thesis, we will use tools from the previous two chapters to sketch the proof of this result and see how it applies to our three toy control systems.

3.1 The Pontryagin Maximum Principle

Whereas the flow of a vector field mapped points to points, the "flow" of a control system will transform points into *sets* of points, called **attainable sets**. We will denote $\mathscr{A}_{P_0}(T)$ for the set of points attainable from P_0 by a trajectory in time T, and $\mathscr{A}_{P_0}$ for the union $\bigcup_{T\geq 0} \mathscr{A}_{P_0}(T)$ of points attainable from P_0 at all.

Now, suppose $P_1 \in \mathscr{A}_{P_0}$. We formalize the problem of *time-optimal control* as follows.

[2]The Pontryagin maximum principle is named after Lev Pontryagin, a Soviet mathematician. Pontryagin worked on algebraic topology and differential topology early in his career but later pivoted into applied mathematics. In 1955, his research group came into contact with members of the Russian airforce who were interested in minimum-time problems in the control of military aircraft [15]. Their first results, including a first form of the maximum principle, were published in 1956. A few years later, the work of Pontryagn and his team were made available in an English translation [16]. Notably, Pontryagin was blind since the age of 14.

Problem 1 (Time-optimal control)**.** *Find a pair (T, p), minimizing T over the positive real numbers, subject to the constraint that $p\colon [0,T] \to M$ is a trajectory of (M,V) with*

$$p(0) = P_0, \quad p(T) = P_1.$$

A trajectory solving this problem is called time-optimal.

Unfortunately, the existence theory for time-optimal control problems is out of the scope of this thesis. It suffices to say that, as a consequence of *Fillipov's theorem*, we have the following result.

Theorem 10. *Suppose the following assumptions are met:*

1. *The control space U is compact.*
2. *The velocity sets are convex, and $V(p) = \{0\}$ for all p outside of a certain compact set $K \subseteq M$.*
3. *Our control system admits a parameterization for which the control-dependent vector field X_u is Lipschitz in its parameter u.*

Then, for any P_0 and P_1 with $P_1 \in \mathscr{A}_{P_0}$, there is a time-optimal trajectory from P_0 to P_1.

Remark 7. Boundedness of velocity sets is an important condition. To see this, consider a modification of Example 1 where the control space $U = [-1,1]$ is replaced with the whole line. Now, any Lipschitz path is a trajectory of our control system, so we can travel between any pair of points by a trajectory of arbitrarily small duration ε. In this case, time-optimal control will have no solution except when $P_0 = P_1$. On the other hand, closedness and convexity are not essentially restrictive. If the velocity sets $V(p)$ of a control system (M,V) are bounded, defining new velocity sets $V^*(p)$ as the convex closures of the sets $V(p)$ gives a relaxed control system (M,V^*) whose trajectories can be "approximated arbitrary well" by trajectories from (M,V). (This is proven in [2].)

Now, we turn to the problem of characterizing time-optimal trajectories. Our argument, which culminates in the Pontryagin maximum principle, has four steps.

1. First, we show that time-optimal trajectories must be *boundary tracing*.
2. Next, we use techniques from Chapter 1 to understand the first-order variations that the end-point $p(T)$ of a trajectory undergoes when we modify its control function by so-called *needle-like variations*. This produces a set of tangent vectors at $p(T)$, which we call the *needle set*.
3. We argue that the needle set of a boundary-tracing trajectory must admit a *supporting covector*. This can already be interpreted as a sort of "maximum principle."
4. Finally, using the theory of Hamiltonian systems developed in Chapter 2, we reformulate our first-order optimality condition in terms of *extremal trajectories* on the cotangent bundle T^*M. This gives the *Pontryagin maximum principle* (PMP). Under certain conditions, the PMP reduces to the statement that boundary-tracing trajectories are projections of integral curves of a certain Hamiltonian system.

This analysis is not entirely straightforward, and in steps 2 and 3 we will need to omit some details. However, we hope the reader will be convinced of the basic argument which underpins the PMP.

We begin with the notion of a boundary-tracing trajectory.

Definition 23. *A trajectory* $p\colon [0,T] \to M$ *of a control system is boundary-tracing when*

$$p(T) \in \partial \mathscr{A}_{p(0)}(T).$$

It is easy to check that a boundary-tracing trajectory must also verify $p(t) \in \partial \mathscr{A}_{p(0)}(t)$ for every $t \in [0,T]$. The boundary-tracing property can be related to time-optimality with the following straightforward argument.

Proposition 22. *If* $p\colon [0,T] \to M$ *is time-optimal, then* p *is boundary-tracing on every sub-interval* $[0,t]$ *for* $0 < t < T$.

Proof. Suppose that $p(t)$ is in the interior of $\mathscr{A}_{p(0)}(t)$ for some $t \in [0,T)$. Then, for some $\varepsilon > 0$, $p(t+\varepsilon)$ also belongs to $\mathscr{A}_{p(0)}(t)$. This means that there is a trajectory $q\colon [0,t] \to M$ taking $p(0)$ to $p(t+\varepsilon)$ in time t. So, can construct a path $p'\colon [0,T-\varepsilon] \to M$ by taking

$$p'(s) \begin{cases} q(s) : s < t \\ p(s+\varepsilon) : s \geq t. \end{cases}$$

This is a trajectory of our control system taking $p(0)$ to $p(T)$ in time $T-\varepsilon$, contradicting the hypothesis that p was time-optimal. □

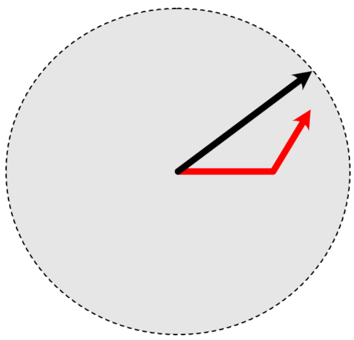

Fig. 3.2 Some trajectories of Example 2, taking M to be the Euclidean plane with its normal metric. Both trajectories take 1 unit of time and begin at P_0. The set $\mathscr{A}_{P_0}(1)$ is highlighted. The red trajectory is not boundary-tracing in an interval $[0,1-\varepsilon]$ and so cannot be time-optimal, while the black trajectory is boundary-tracing and is in fact time-optimal.

Remark 8. The author strongly suspects that a time-optimal trajectory $p\colon [0,T] \to M$ must be boundary-tracing on the whole interval $[0,T]$. However, the argument above does not extend to the case $t = T$ in an obvious way. A review of the literature has so far only uncovered a proof of this assertion in the special case of a *linear control system*. (See page 302 of [11].)

For our second step, we will need to assume that we are working with a parameterized control system. A technical result from [2] shows that, for every trajectory $p\colon [0,T] \to M$ of a parameterized control system, there exists a measurable and locally bounded **control function**, $u\colon [0,T] \to U$, so that p is an interval curve of the non-autonomous vector field $X_{u(t)}$. We will say that p is **driven by**

u. (Note that two control functions, distinct up to equality almost everywhere, may drive the same trajectory.)

We will also need the idea of a Lebesgue point. When D is a real interval, recall that $t \in D$ is a *Lebesgue point* of a function $f\colon D \to \mathbb{R}^n$ when

$$\lim_{\varepsilon \to 0} \frac{1}{2\varepsilon} \int_{t-\varepsilon}^{t+\varepsilon} f(s)\, ds = f(t).$$

By the Lebesgue differentiation theorem, we know that, since u is measurable, its Lebesgue points are a full-measure subset of the interval $[0,T]$.

Proposition 23. *Let $p\colon [0,T] \to M$ be a trajectory of a control system driven by the control function $u\colon [0,T] \to U$. Suppose $s \in [0,T]$ is a Lebesgue point of u, and denote $v = p(s)X_{u(s)}$. Then, for any other vector $w \in V(p(s))$, there exists a curve $\gamma\colon [0,\varepsilon] \to \mathscr{A}_{p(0)}(T)$ for which*

$$\dot\gamma(0) = (w - v)\, \overrightarrow{\exp} \int_s^T X_{u(t)}\, dt.$$

Proof. Since w belongs to the velocity set $V(p(s))$, there exists a control parameter $\tilde u \in U$ for which $p(s)X_{\tilde u} = w$. Consider the family of control functions $u_\varepsilon\colon [0,T] \to U$ defined, for $\varepsilon > 0$, by

$$u_\varepsilon(t) = \begin{cases} \tilde u & : |t-s| < \varepsilon/2 \\ u(t) & : \text{otherwise.} \end{cases}$$

The function u_ε is called a *needle-like variation* of u. Let us also denote D_ε for the difference

$$D_\varepsilon(t) = X_{u_\varepsilon(t)} - X_{u(t)} = \begin{cases} X_{\tilde u} - X_{u(t)} & : |t-s| < \varepsilon/2 \\ 0 & : \text{otherwise.} \end{cases}$$

By substituting u with u_ε, the trajectory p is deformed into a new trajectory p_ε, namely

$$p_\varepsilon(t) = p(0)\, \overrightarrow{\exp} \int_0^t \left(X_{u(t)} + D_\varepsilon(t)\right) dt.$$

Let us write $\gamma(\varepsilon) = p_\varepsilon(T)$. Note that $\gamma(0) = p(T)$. We will also write

$$P_{t_0}^{t_1} = \overrightarrow{\exp} \int_{t_0}^{t_1} X_{u(t)}\, dt.$$

Our goal is to show that $\dot\gamma(0) = (w-v)P_s^T$, which will complete the proof.

First, we use a variations formula (from Theorem 6) to express $\gamma(\varepsilon)$ in terms the flow of $(\operatorname{Ad} P_T^t)D_\varepsilon(t)$.

$$\begin{aligned} \gamma(\varepsilon) &= p(0)\, \overrightarrow{\exp} \int_0^T X_{u(t)}\, dt\ \overrightarrow{\exp} \int_0^T (\operatorname{Ad} P_T^t) D_\varepsilon(t)\, dt \\ &= \gamma(0)\, \overrightarrow{\exp} \int_0^T (\operatorname{Ad} P_T^t) D_\varepsilon(t)\, dt. \end{aligned} \tag{3.1}$$

Next, we use the first-order truncation of the Volterra series (Theorem 5) to estimate this flow. We find that

$$\begin{aligned}\gamma(\varepsilon) &= \gamma(0)\left(\mathrm{id} + \int_0^T (\mathrm{Ad}\,P_T^t) D_\varepsilon(t)\,dt + R(\varepsilon)\right)\\ &= \gamma(0) + \gamma(0)\int_0^T (\mathrm{Ad}\,P_T^t) D_\varepsilon(t)\,dt + \gamma(0)R(\varepsilon)\end{aligned}$$

for a certain remainder $R(\varepsilon)$. Unfortunately, we have not spent enough time on operator-valued calculus in this thesis to finish the proof entirely rigorously. We proceed informally.

Because $D_\varepsilon(t)$ is only non-zero when t is in the interval $[s-\varepsilon/2, s+\varepsilon/2]$, D_ε behaves like a term of order $\mathscr{O}(\varepsilon)$. Therefore, the remainder $R(\varepsilon)$ will be of order $\mathscr{O}(\varepsilon^2)$, and the differential of $\gamma(0)R(\varepsilon)$ vanishes. It remains only to differentiate the term

$$\begin{aligned}\gamma(0)\int_0^T (\mathrm{Ad}\,P_T^t) D_\varepsilon(t)\,dt &= \left(\int_0^T p(t) D_\varepsilon(t) P_t^T\,dt\right)\\ &= \left(\int_{s-\varepsilon/2}^{s+\varepsilon/2} p(t)(X_{\tilde{u}} - X_{u(t)}) P_t^T\,dt\right).\end{aligned}$$

Applying various regularity properties, including the hypothesis that s is a Lebesgue point of u, it is possible to conclude that

$$\begin{cases}\lim\limits_{\varepsilon\to 0}\dfrac{1}{\varepsilon}\displaystyle\int_{s-\varepsilon/2}^{s+\varepsilon/2} p(t) X_{\tilde{u}} P_t^T\,dt = wP_s^T, \text{ and}\\ \lim\limits_{\varepsilon\to 0}\dfrac{1}{\varepsilon}\displaystyle\int_{s-\varepsilon/2}^{s+\varepsilon/2} p(t) X_{u(t)} P_t^T\,d\tau = p(s) X_{u(s)} P_s^T = vP_s^T.\end{cases}$$

This shows that, as desired, $\dot{\gamma}(0) = (w-v)P_s^T$. □

Let us consider the set of all such vectors $(w-v)P_s^T$ guaranteed by the previous proposition.

Definition 24. *Let $p\colon [0,T] \to M$ be a trajectory of the parameterized control system (M,U,X) driven by a control function u. The* ***needle set*** *$\mathscr{N}$ of the pair (u,p) is the set of all vectors of the form*

$$(w - p(s)X_{u(s)})P_s^T$$

where $s \in [0,T]$ is a Lebesgue point of u and w is an element of $V(p(s))$.

The third step of our analysis is to give a geometric property of the needle set necessary for p to be a boundary-tracing trajectory.

Lemma 4. *Suppose $p\colon [0,T] \to M$ is a boundary-tracing trajectory driven by u. Then there is a covector $\Lambda \in T^*_{p(T)}M$ so that, for all vectors $\eta \in \mathscr{N}$, $\langle \Lambda, \eta\rangle \le 0$.*

Now, the inequality $\langle \Lambda, (w - p(s)X_{u(s)})P_s^T\rangle \le 0$ for can be rearranged as

$$\langle \Lambda, wX_{u(s)}P_s^T\rangle \le \langle \Lambda, p(s)X_{u(s)}P_s^T\rangle,$$

so the conclusion of this lemma can be reinterpreted in the following way.

Proposition 24. *Suppose $p\colon [0,T] \to M$ is a boundary-tracing trajectory driven by u. Then there is a covector $\Lambda \in T^*_{p(T)}M$ for which*

$$\langle \Lambda, p(s)X_{u(s)}P_s^T \rangle = \max_{v \in V(p(s))} \langle \Lambda, vP_s^T \rangle$$

for almost every $s \in [0,T]$.

If we make the simplifying assumption that $\mathscr{A}_{p(0)}(T)$ is locally diffeomorphic to a convex set near $p(T)$, Lemma 4 becomes a simple consequence of the supporting hyperplane theorem for convex sets.

Restricted argument for Lemma 4. Suppose that there is a neighborhood $V \subseteq M$ of $p(T)$ and a diffeomorphism $\varphi\colon V \to D$ for a certain real domain $D \subseteq \mathbb{R}^n$ so that $A = \varphi(\mathscr{A}_{p(0)}(T) \cap V)$ is a convex set. By the hypothesis that $p(T)$ is on the boundary of $\mathscr{A}_{p(0)}(T)$, $\varphi(p(T))$ must also be on the boundary of A. By the supporting hyperplane theorem, there exists a covector $L\colon \mathbb{R}^n \to \mathbb{R}$ that attains a maximum over A at $\varphi(p(T))$. The covector $\Lambda(v) = L(\varphi_*(v))$ is now easily checked to have the desired property. □

However, a totally rigorous proof is somewhat technical and therefore out of the scope of this thesis. One approach, explained in [17], is to show that the conic closure of $\mathscr{N}$ is a *Boltyansky approximating cone* for $\mathscr{A}_{p(0)}(T)$ at $p(T)$. It can be proven that, under no further assumptions, a Boltyansky approximating cone for a set S around a point p must be proper if p is not in its interior.

Our final step is to restate Proposition 24 in a form that will be much more amenable to applications. We begin with some definitions.

Definition 25. *When X_u is the parameterized vector field of a control system, let h_u be*

$$h_u\colon T^*M \to \mathbb{R}$$
$$h_u(\lambda) = H_{X_u} = \langle \lambda, X_u \rangle.$$

*The **maximized Hamiltonian** of a parameterized control system is the supremum*

$$H = \sup_{u \in U} h_u.$$

Note that this only depends on the underlying control system (M,V), since

$$H(\lambda) = \sup_{v \in V(\lambda\pi)} \langle \lambda, v \rangle.$$

It can be proven that, under the assumptions of Theorem 10, H is a Lipschitz function. Furthermore, since it is defined on each cotangent space as a supremum of linear functionals, H is guaranteed to be homogeneous; $H(\alpha\lambda) = \alpha H(\lambda)$ for any non-zero scalar α.

Definition 26. *The **Hamiltonian control system** associated with a parameterized control system (M,U,X) given by the vector fields X_u is the parameterized control system (T^*M,U,Y), where Y_u is the Hamiltonian lift $\overline{X}_u$.*

An ***extremal trajectory*** *of* (M,U,X) *is a trajectory* $\lambda\colon [0,T]\to T^*M$ *of the associated Hamiltonian control system which verifies the* ***maximum condition***

$$\lambda(t)\neq 0 \wedge \langle \lambda(t), \dot{\lambda}(t)\pi\rangle = H(\lambda(t)) \tag{3.2}$$

for almost all $t\in[0,T]$.

When u is a control function that drives λ, note that the maximum condition can be rewritten as

$$\lambda(t)\neq 0 \wedge \langle \lambda(t), \lambda(t)\pi X_{u(t)}\rangle = H(\lambda(t)).$$

Theorem 11 (Pontryagin maximum principle)**.** *Every boundary-tracing trajectory of a parameterized control system is the image, under the bundle projection* $\pi\colon T^*M\to M$*, of an extremal trajectory.*

Proof. Let $p\colon [0,T]\to M$ be boundary-tracing, driven by a control function $u\colon [0,T]\to U$. As before, we denote

$$P_{t_0}^{t_1} = \overrightarrow{\exp}\int_{t_0}^{t_1} X_{u(t)}\,dt.$$

Let us also write $Q_{t_0}^{t_1} = T^*P_{t_0}^{t_1}$, and let $\Lambda\in T^*_{p(T)}M$ be the non-zero covector guaranteed by Proposition 24.

Define $\lambda\colon [0,T]\to T^*M$ by $\lambda(t)=\Lambda Q_T^t$. By the properties of point transformations, we know that λ projects down to p: $\lambda(t)\pi = p(t)$. By Theorem 8, we also know that λ is a trajectory of the Hamiltonian control system, since

$$\lambda(\tau) = \lambda(0)\overrightarrow{\exp}\int_t^0 \overline{X}_{u(s)}\,ds.$$

Finally, by the definition of the point transformation $T^*P_T^t$,

$$\langle\lambda(t),v\rangle = \langle\Lambda Q_T^t, v\rangle = \langle\Lambda, vP_t^T\rangle$$

for any tangent vector $v\in T_{p(t)}$. By the construction of Λ, we conclude that λ is an extremal trajectory. □

Now, the behavior of a control system—and in particular, the nature of time-optimal and boundary-tracing trajectories—is captured by its velocity sets alone. However, because the Hamiltonian control system of (M,U,X) depends on the parameterization chosen, the definition of an extremal trajectory would also appear to depend on a parameterization. This is troubling; do different parameterizations of the same control system give different extremal trajectories? We will now show that, whenever the maximized Hamiltonian of our control system is differentiable, there is a convenient parameterization-invariant description of extremal curves.

First, note that for a smooth function $F\colon T^*M\to\mathbb{R}$, the value of the Hamiltonian vector field $\overrightarrow{F}(p)$ at a given point $p\in T^*M$ is a linear function of the differential of $(dF)(p)$. We will write $-^\sharp$ for each linear map $T_p^*T^*M\to T_pT^*M$ which makes

$$\overrightarrow{F}(p) = (dF)(p)^\sharp.$$

Proposition 25. *Let λ be an extremal trajectory of a control system driven by u, and let t be a moment where $\dot{\lambda}(t) = \lambda(t)\overline{X}_{u(t)}$. If H is differentiable at $\lambda(t)$, then*

$$\dot{\lambda}(t) = (dH)(\lambda(t))^\sharp.$$

Proof. Since λ is extremal, $h_{u(t)}(\lambda(t)) = H(\lambda(t))$. In general $H \geq h_{u(t)}$, so we conclude the function $H - h_{u(t)}$ attains a local minimum of 0 at $\lambda(t)$. Thus, if H is differentiable at $\lambda(t)$, we must have $(dH)(\lambda(t)) = (dh_{u(t)})(\lambda(t))$. Consequently,

$$\dot{\lambda}(t) = \lambda(t)\overline{X}_{u(t)} = \lambda(t)\overrightarrow{h}_{u(t)} = (dh_{u(t)})(\lambda(t))^\sharp = (dH)(\lambda(t))^\sharp.$$

□

When $H\colon T^*M \to R$ is everywhere differentiable, we may define a vector field by $\overrightarrow{H}(p) = (dH)(p)^\sharp$. When λ is an extremal trajectory driven by a control function u, almost every $t \in [0,T]$ safisfies the condition of Proposition 25, so the next theorem follows as a corollary.

Theorem 12. *If the maximized Hamiltonian is everywhere differentiable, then every extremal trajectory $\lambda\colon [0,T] \to T^*M$ is an integral curve of $\overrightarrow{H}$.*

Although we will not prove this here, it can also be shown that all such integral curves are extremal trajectories when H is at least C^2.

Let us take a moment to recognize the impressiveness of this result. When the velocity sets of our control system are convex—which, as we have mentioned earlier, we can assume w.l.o.g.—the maximized Hamiltonian is a very natural way to encode V as a function of the cotangent bundle. If H turns out to be differentiable, then the Pontryagin maximum principle says that time-optimal trajectories must be projections of integral curves of $\overrightarrow{H}$! The machinery used in the proof of the PMP falls away, and we are left with a somewhat magical and unexpected result. The author would classify this as an example of the "pure mathematics of the first type" defined by George Simmons.

When H is not differentiable, extremal trajectories cannot be characterized as integral curves. In fact, in the example of the Dubins car, we will find that the evolution of an extremal trajectory need not be determined by its initial condition. However, even in the general case, the class of extremal curves will still respect a certain family of "Lie point symmetries" (scaling of covectors) and a conserved quantity (the maximized Hamiltonian H).

Proposition 26. *Suppose $\lambda\colon [0,T] \to T^*M$ is an extremal trajectory of the parameterized control system (M,U,X). Then the following is true.*

1. *For any $\alpha > 0$, the scalar multiple $\alpha\lambda$ is also an extremal trajectory.*
2. *The maximized Hamiltonian H is a conserved quantity of λ.*

Proof. (1) Let u be a control function that drives λ. We will write $S_\alpha\colon T^*M \to T^*M$ for the scalar multiplication $S_\alpha(\lambda) = \alpha\lambda$, and $\lambda'(t) = \lambda(t)S_\alpha$.[3] If $\lambda(s) \neq 0$ and

$$\langle \lambda(s), \dot{\lambda}(s)\pi \rangle = H(\lambda(s)),$$

[3]We avoid the notation $\alpha\lambda(t)$ here because it may be understood to mean the scalar multiple of $\lambda(t)$ as a map from $C^\infty T^*M$ to $\mathbb{R}$.

then the same is true for λ', since $\dot{\lambda}'(s)\pi = \dot{\lambda}(s)\pi$ and H is homogeneous. To prove that λ' is an extremal trajectory, it only remains to prove that is a trajectory of the Hamiltonian control system.

We claim that λ' is also driven by the control function u. By Theorem 8, we know that

$$\lambda(0)S_\alpha \overrightarrow{\exp} \int_0^t \overline{X}_{u(s)}\, ds = \lambda(0)S_\alpha T^* \left(\overrightarrow{\exp} \int_0^t X_{u(s)}\, ds \right)$$

Since point transformations are linear maps, they commute with S_α, so

$$\lambda(0)S_\alpha T^* \left(\overrightarrow{\exp} \int_0^t X_{u(s)}\, ds \right) = \lambda(0) \left(\overrightarrow{\exp} \int_0^t X_{u(s)}\, ds \right) S_\alpha = \lambda(t)S_\alpha.$$

We conclude that

$$\lambda(t)S_\alpha = \lambda(0)S_\alpha \overrightarrow{\exp} \int_0^t \overline{X}_{u(s)}\, ds,$$

as desired.

(2) Let λ be an extremal trajectory with associated control u. It can be shown (using our assumptions above) that H is Lipschitz. It follows that $H \circ \lambda$ is almost everywhere differentiable, so it is enough to prove that $d/dt\, H(\lambda(t)) = 0$ at any almost differentiability point t.

For notational convenience, let us treat H as an element of $C^\infty T^*M$ so that we can write $H(\lambda(t))$ as $\lambda(t)(H)$. Since λ is extremal, $\lambda(s)(H - h_{u(t)})$ attains a minimum in the variable s when $s = t$, so

$$\frac{d}{dt}\lambda(t)(H) = \left(\frac{d}{ds}\right)_t \lambda(s)(h_{u(t)}).$$

However, since in addition $\dot{\lambda}(t) = \lambda(t)\overrightarrow{h}_{u(t)}$, for almost every t, we find that

$$\left(\frac{d}{ds}\right)_t \lambda(s)(h_{u(t)}) = (\lambda(t)\overrightarrow{h}_{u(t)})(h_{u(\tau)}) = \lambda(t)\{h_{u(t)}, h_{u(t)}\} = 0.$$

□

Finally, let us return to the case where H is smooth. Theorem 12 above shows that the problem of finding a time-optimal trajectory from P_0 to P_1 can be approached by solving the non-linear equation

$$\lambda_0 \exp(t\overrightarrow{H})\pi = P_1 \tag{3.3}$$

in the variables $\lambda_0 \in T^*_{P_0}$ and $t \in \mathbb{R}_{\geq 0}$. Actually, by homogeneity, λ_0 can be restricted to lie on the boundary of some open ball $B \subseteq T^*_{P_0}$. With this restriction, our system has n variables altogether, which is the number that would be necessary to parameterize the family of time-optimal paths when $\mathscr{A}_{P_0} = M$. When $\partial\mathscr{A}_{P_0}(T)$ is $(n-1)$-dimensional in some sense, we expect the map

$$\varphi \colon \partial B \subseteq T^*_{q_0}M \to M$$
$$\varphi(\lambda) = \lambda \exp(T\overrightarrow{H})\pi$$

to be some sort of parameterization of this boundary region by the sphere ∂B. This is exactly what happens in Example 2 when M is the Euclidean plane.

3.2 Applications of the PMP

In the final section of this chapter, we will investigate the nature of extremal trajectories in the three example control systems given earlier: movement on the line, movement on a Riemannian manifold, and movement of a Dubins car.

Example 4 (The PMP applied to movement on the line)**.** *Let $M = \mathbb{R}$, $U = [-1,1]$, and $X_u = u\partial/\partial x$. We easily compute that $\overline{X}_u = u\partial/\partial x$, so a trajectory λ of the associated Hamiltonian control system is just a pairing of a trajectory of the original control system with a non-zero constant, c, giving the coefficient of the covector λ in standard coordinates. A trajectory is extremal when the expression*

$$\langle \lambda, \lambda \pi X_u \rangle = cu$$

is always maximized in u, meaning that $u \equiv \operatorname{sgn}(c)$.

In this case, time-optimal trajectories coincide with the projections of extremal trajectories. However, we cannot expect this to happen in general without extra assumptions. One serious problem is the attainable set $\mathscr{A}_{P_0}$ may have empty interior, in which case *every* trajectory will be boundary-tracing and hence will be the projection of an extremal curve. For example, if we let $M = \mathbb{R}^2$ be parameterized by coordinates (x,y) and take $U = [-1,1]$ and $X_u = u\partial/\partial x$ as in the previous example, we find that the trajectory

$$\lambda : I \to T^*M \cong \mathbb{R}^2 \times (\mathbb{R}^2)^*$$
$$\lambda(t) = (p(t),(0,1))$$

is extremal for *any* trajectory p of the control system (M,U,X), including ones that are not time-optimal.

However, in practice, the Pontryagin maximum principle is frequently an effective tool to search for time-optimal trajectories. Our next example is another successful application of the PMP.

Example 5 (The PMP applied to movment on a Riemannian manifold)**.** *When (M,g) is a Riemannian manifold and (M,V) is the corresponding control system constructed in Example 2, the maximized Hamiltonian is given by the norm on T^*M dual to g. When the metric is expressed in coordinates as*

$$g = g_{i,j}\, dx_i \otimes dx_j,$$

we know that the dual metric is expressed as

$$g^* = g^{i,j} \frac{d}{dx_i} \otimes \frac{d}{dx_j}$$

where $g^{i,j}$ gives the inverse to the metric coefficients, in the sense that $g_{i,j}g^{j,k} = \delta_i^k$. Where λ_i are coordinates for the cotangent space, the maximized Hamiltonian for our problem is thus

$$H = \sqrt{\lambda_i \lambda_j g^{i,j}}.$$

This is a smooth function at a point where $\lambda \neq 0$, so extremal trajectories will be integral curves of the Hamiltonian vector field

$$\begin{aligned}\vec{H} &= \frac{\partial H}{\partial \lambda_i}\frac{d}{dx_i} - \frac{\partial H}{\partial x_i}\frac{d}{d\lambda_i}\\ &= \frac{\lambda_j g^{i,j}}{H(\lambda)}\frac{d}{dx_i} - \frac{\lambda_i \lambda_j}{2H(\lambda)}\frac{dg^{i,j}}{dx_k}\frac{d}{d\lambda_k}.\end{aligned}$$

Because the quantity H is conserved along integral curves and H is homogeneous, we can normalize λ so that $H = 1$ constantly. With this simplification, we conclude that an extremal trajectory satisfies the differential equations

$$\begin{cases}\dot{x}_i = \lambda_j g^{i,j}\\ \dot{\lambda}_i = -\dfrac{1}{2}\lambda_j \lambda_k \dfrac{dg^{j,k}}{dx_i}.\end{cases}$$

Although we will not discuss it here, the PMP can also be applied in an analogous way to solve the historically significant *brachistochrone problem* [19]. (In fact, the brachistochrone problem can be reduced to a special case of the geodesic problem.) However, the geodesic problem and the Brachistochrone problem are amenable to other types of analysis, like the calculus of variations. (See [10] for an overview of these techniques, and a derivation of the geodesic equation from the Euler-Lagrange equation.) The special interest of the Pontryagin maximum principle is that it generalizes, in principle, to *any* optimal control problem. The classical approaches to the geodesic problem and Brachistochrone problem, meanwhile, only succeed when we can state our problem as the minimization of a Lagrangian functional over an unconstrained space of functions. For example, calculus of variations could not be used to study our next example.

Example 6 (The PMP applied to control of a Dubins car)**.** *Recall the example of the Dubins car, where $M = \mathbb{R}^2 \times S^1$ is parameterized by coordinates (x_1, x_2, θ), $U = [-1, 1]$, and our control-dependent vector field is*

$$X_u = \cos(\theta)\frac{\partial}{\partial x_1} + \sin(\theta)\frac{\partial}{\partial x_2} + u\frac{\partial}{\partial \theta}.$$

Using canonical coordinates coordinates $(x_1, x_2, \theta, \xi_1, \xi_2, \eta)$ for the cotangent bundle, we compute that the Hamiltonian lift of X_u is

$$\overline{X}_u = \cos(\theta)\frac{\partial}{\partial x_1} + \sin(\theta)\frac{\partial}{\partial x_2} + u\frac{\partial}{\partial \theta} + (\xi_1 \sin(\theta) - \xi_2 \cos(\theta))\frac{\partial}{\partial \eta}.$$

We find that the coordinates (ξ_1, ξ_2) of a trajectory λ in the Hamiltonian control system are always constant. The evolution of η is also quite simple; since

$$\dot{\eta} = \xi_1 \sin(\theta) - \xi_2 \cos(\theta) = \xi_1 \dot{x}_2 - \xi_2 \dot{x}_1,$$

we conclude that $\eta = \eta_0 + \xi_1 x_2 - \xi_2 x_1$ for some constant η_0. The condition that

$$\langle \lambda, X_u \rangle = \xi_1 \cos(\theta) + \xi_2 \sin(\theta) + \eta u$$

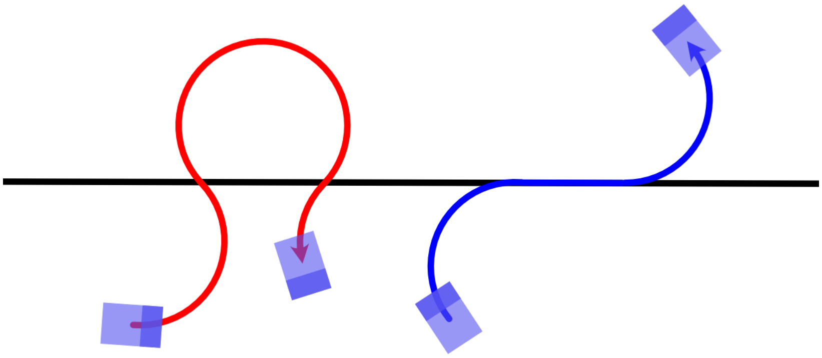

Fig. 3.3 Projections of two types of extremal curves for the Dubins car control system. We assume that (ξ_1, ξ_2, η_0) are such that $\eta > 0$ when the car is below the black line and $\eta < 0$ when the car is above.

is maximized in u means that $u = \operatorname{sgn} \eta$ at almost every moment where η is non-zero. Let us now put this information together to understand the behavior of extremal trajectories of this problem.

First of all, suppose $\xi_1 = \xi_2 = 0$. Then, since λ must be non-zero, the initial value of η must be non-zero. Of course, if $\xi_1 = \xi_2 = 0$, then η is constant over our trajectory. We conclude that u must be either constantly -1 or 1, depending on the sign of η.

Next, consider the more interesting case where (ξ_1, ξ_2) is non-zero. In these conditions, a control function u will produce an extremal curve exactly if u takes the value 1 almost everywhere when (x, y) is on the half-plane $\eta_0 + \xi_1 x_2 - \xi_2 x_1 > 0$ and takes the value -1 almost everywhere when (x, y) is on oppsite half-plane.

The reader can now see for themself what possibilities there are. If we begin at a point where $\eta \neq 0$, the car will initially turn to the right or the left depending on the sign of η. At some point, the position (x, y) of the car may reach the dividing line where $\eta = 0$. If $\dot{\eta} \neq 0$ at this point, the position of the car will move into the opposite half-plane, forcing the car to start turning in the opposite direction. If $\dot{\eta} = 0$—meaning that the circle driven by the car is tangent to the line $\eta = 0$—then we have a choice. On one hand, the control can be set to 0, making the car drive straight along the dividing line. However, at any future moment, we may decide to turn off the dividing line in either direction, initiating a new period of turning motion within one of the half-planes.

The case where $\eta = 0$ at the initial moment is handled similarly. It only remains to consider the case where our car is initially "pointing along the dividing line." In this situation, there is an extra subtlety: depending on the sign of (ξ_1, ξ_2), it may be impossible to begin turning in either direction.

From here, it is not much more difficult to conclude what the *time-optimal* trajectories of the Dubins car are. (This problem was originally solved, without the PMP, in [8]. For a more modern perspective on these results, see [5]. Also, see [3] for a more general problem of optimal control related to the Dubins car.)

A very curious fact, which the author has not yet had time to explore, is that the conserved quantities of an extremal curve for the Dubins car, namely,

$$H_1 = \xi_1,\, H_2 = \xi_2, \text{ and } H_3 = x_1\xi_2 - x_2\xi_1 + \eta,$$

are associated with the obvious infinitesimal symmetries of the underlying control system, namely

$$Y_1 = \frac{\partial}{\partial x}, \quad Y_2 = \frac{\partial}{\partial y},$$
$$Y_3 = x_1\frac{\partial}{\partial x_2} - x_2\frac{\partial}{\partial x_1} + \frac{\partial}{\partial \theta}.$$

(Actually, this symmetry-based analysis is what gave us the idea to write $\eta = \eta_0 + \xi_1 x_2 - \xi_2 x_1$. This simplification of the problem was not initially clear to us and allowed a significant simplification over the analysis of the Dubins car done in [2].) We hypothesize that this phenomenon generalizes, which would be a "Noether's theorem for the PMP!"

Conjecture 1. *Let (M, U, X) be a parameterized control system, and let Y be an autonomous vector field so that*

$$V(p)\exp(tY) = V(p\exp(tY))$$

for all $t \in \mathbb{R}$ and $p \in M$, where $V(p)\exp(tY)$ is the image of $V(p)$ under the pushforward of $\exp(tY)$*. Then H_Y is a conserved quantity of any extremal curve.*

Bibliography

[1] Andrei A. Agrachev and Revaz V. Gamkrelidze. The exponential representation of flow and the chronological calculus. *Mathematics of the USSR-Sbornik*, 35(6):727–785, June 1979. URL: http://stacks.iop.org/0025-5734/35/i=6/a=A01?key=crossref.769941c8cf18c8b3d9501ad1ed2decca, doi:10.1070/SM1979v035n06ABEH001623.

[2] Andrei A. Agrachev and Yuri L. Sachkov. Control theory from the geometric viewpoint. ISBN: 9783662064047, OCLC: 1157341037.

[3] Andrei A. Agrachev and Mario Sigalotti. On the Local Structure of Optimal Trajectories in R^3. *SIAM Journal on Control and Optimization*, 42(2):513–531, January 2003. URL: http://epubs.siam.org/doi/10.1137/S0363012902409246, doi:10.1137/S0363012902409246.

[4] Vladimir I. Arnold. *Mathematical Methods of Classical Mechanics*, volume 60 of *Graduate Texts in Mathematics*. Springer New York, New York, NY, 1978. URL: http://link.springer.com/10.1007/978-1-4757-1693-1, doi:10.1007/978-1-4757-1693-1.

[5] José Ayala. Length minimising bounded curvature paths in homotopy classes. *Topology and its Applications*, 193:140–151, 2015. URL: https://www.sciencedirect.com/science/article/pii/S0166864115002771, doi:https://doi.org/10.1016/j.topol.2015.06.008.

[6] Olivier Cots Bernard Bonnard and Boris Wembe. Zermelo navigation problems on surfaces of revolution and hamiltonian dynamics. *SIAM Journal on Control and Optimization*, 2021. URL: https://hal.archives-ouvertes.fr/hal-03209491/.

[7] Francesco Boarotto and Mario Sigalotti. Time-optimal trajectories of generic control-affine systems have at worst iterated fuller singularities, 2018. arXiv:1705.10055.

[8] L. E. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):497–516, 1957. URL: http://www.jstor.org/stable/2372560.

[9] Andreas Henriksson. On the second law of thermodynamics. URL: https://philarchive.org/rec/HENOTE-7.

[10] R. Herman. Euler equation and geodesics. 2018. URL: https://www.semanticscholar.org/paper/Euler-Equation-and-Geodesics-Herman/eab52fe62338775f01553c7aaaff666f5c3389a3.

[11] Velimir Jurdjevic. *Geometric control theory*. Number 52 in Cambridge studies in advanced mathematics. Cambridge University Press, Cambridge ; New York, NY, USA, 1997.

[12] Ivan Kolář, Jan Slovák, and Peter W. Michor. *Natural Operations in Differential Geometry*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1993. URL: http://link.springer.com/10.1007/978-3-662-02950-3, doi:10.1007/978-3-662-02950-3.

[13] Saunders Mac Lane. *Categories for the working mathematician*. Number 5 in Graduate texts in mathematics. Springer, New York, NY, 2nd. ed., softcover version of original hardcover edition 1998 edition, 2010.

[14] Ieke Moerdijk and Gonzalo E. Reyes. *Models for Smooth Infinitesimal Analysis*. Springer New York, New York, NY, 1991. OCLC: 851770169. URL: https://doi.org/10.1007/978-1-4757-4143-8.

[15] Hans Josef Pesch and Michael Plail. The maximum principle of optimal control: A history of ingenious ideas and missed opportunities. *Control and Cybernetics*, 38, 01 2009.

[16] Lev S. Pontryagin and Lucien W. Neustadt. *The mathematical theory of optimal processes*. Interscience Publ., New York, 1965. OCLC: 1069856997.

[17] Heinz Schättler and Urszula Ledzewicz. *Geometric Optimal Control*, volume 38 of *Interdisciplinary Applied Mathematics*. Springer New York, New York, NY, 2012. URL: http://link.springer.com/10.1007/978-1-4614-3834-2, doi:10.1007/978-1-4614-3834-2.

[18] George Finlay Simmons. *Introduction to topology and modern analysis*. R.E. Krieger Pub. Co, Malabar, Fla, 1983.

[19] Héctor J. Sussmann and Jan C. Willems. *The Brachistochrone Problem and Modern Control Theory*, pages 113–166. URL: https://www.worldscientific.com/doi/abs/10.1142/9789812778079_0004, arXiv:https://www.worldscientific.com/doi/pdf/10.1142/9789812778079_0004, doi:10.1142/9789812778079_0004.

[20] G. Thomas and Erik F. Integration of functions with values in locally convex Suslin spaces. *Transactions of the American Mathematical Society*, 212:61–61, 1975. URL: http://www.ams.org/jourcgi/jour-getitem?pii=S0002-9947-1975-0385067-1, doi:10.1090/S0002-9947-1975-0385067-1.