Is it worth learning differential geometric methods for modelling and control of mechanical systems?*

Andrew D. Lewis†

2007/01/11
Last updated: 2007/07/13

Abstract

Evidence is presented to indicate that the answer is, "Yes, sometimes."

Keywords. Riemannian geometry, affine differential geometry, control of mechanical systems.

AMS Subject Classifications (2020). 70E15, 70F25, 70Q05, 53B20, 93B05, 93B27.

1. Introduction

Before I get to the actual writing of this paper, let me say a few words about the manner and spirit in which it is written. Normally when writing technical articles I try to be rigorous and precise, giving clear definitions, stating clear theorems, etc. However, I believe that this should not be the tone of this article, given the intent of the special issue for which it is being written. Therefore, I make no claims of the article being precise, rigorous, or even remotely self-contained. Readers looking for a summary of differential geometric methods in mechanics will not find it here. Many such summaries have been written, and many of them are good. Moreover, many books have been written on the topic of differential geometric modelling of mechanical systems, and many of these are good. A reader who has an inkling that they want to learn the differential geometric method will already have ventured into this part of the literature, so it seems pointless to simply reproduce it, or a small part of it, here.

Instead I aim to give the reader some insights into *why* I think differential geometric modelling is important. In doing so I will (1) make free use of intuitive "definitions" of concepts and (2) make free use of terminology and notation that I have not defined, but is hopefully either clear or for which the absence of a precise definition is inobtrusive. Thus the paper should not be read as if it is a source of technical information, but rather as an exposition about technical information that can be found elsewhere.

One of my hopes for this article is that it will show that there are real advantages to be had from the differential geometric techniques of modelling and control of mechanical

---

systems. Consistent with this, I will also make the point that there are *dis*advantages to not using differential geometric techniques. Moreover, I will illustrate this latter point through particular instances. Therefore, the article will be, in a very few places I hope, somewhat provocative. In keeping with this, I have decided to write the article in a less formal and more personal style.

That all out of the way, let me get to the actual substance of the paper. The most common feedback I get from applied researchers not familiar with the mathematics I use is (1) "I do not understand your notation" and (2) "I am not sure you are getting any mileage from all of the mathematics." In this paper I will attempt to address these concerns.

First let me address the notation matter. I am sympathetic to this criticism in that I am certain that some parts of differential geometric notation seem almost intentionally cryptic if one is not used to them. I understand from my own experience that these notational issues can be a serious obstruction to understanding something new. However, my own experience also suggests that sometimes notation can be a placeholder for ideas, and it is the ideas behind the notation that should be focussed upon. I would, therefore, caution against disregarding something as being "mere notation" without taking the time to see whether the notation is actually a manifestation of something real. The affine connection framework for mechanics that I will review in this paper certainly makes use of notation that will be unfamiliar to those with a more classical mechanics background. However, I think that the framework brings with it ideas that can be useful; the notation is merely a consequence of the existence of these ideas.

The other matter—that of whether differential geometric methods actually contribute anything in practice—has real substance, and it is to this matter that this paper is primarily addressed. There is wonderful elegance and clarity in the differential geometric formulation of mechanics. Much of the work done in the area of overlap between differential geometry and mechanics—indeed, almost all of it until, maybe, the past fifteen years—has been in understanding the structure of the mathematical models of the mechanics. While such research is certainly justified on intellectual grounds by its outcome, to an engineer looking for solutions to specific problems, it is not clear that the differential geometric formulation contributes to these ends. Thus any skepticism that might be felt towards differential geometric techniques is justified, at least until it is clearly shown by people like me that these techniques contribute towards the solution of problems that are of interest. What I will try to do in this paper is indicate some of the kinds of problems where differential geometric ideas have made a clear and decisive contribution. I will argue that this contribution might arise in two ways.

1. It is sometimes the case that, merely by formulating a problem in differential geometric language, the right questions reveal themselves. This is an often overlooked contribution of precise mathematical formulations in general. In solving a problem it is often the question that is the real problem. Sometimes, once one has the right question, the problem either (a) solves itself or (b) reveals itself to be, in a precise way, impossible. For either eventuality, knowing the right question is much more than half the battle.

2. Sometimes the question one is asking is perfectly clear, and it is the solution that is the real problem. In some such cases (but definitely not all!) differential geometric methods are just the thing, and can lead one to unexpectedly simple solutions.

I should make some effort to position what I say in this paper relative to the growing

volume of literature on differential geometric methods for control of mechanical systems. It is impossible to give an accurate overview of the entirety of this research area, and I will not attempt to do so. Instead I will identify research directions that have achieved a somewhat polished state. Probably the first serious progress in uniting geometric mechanics and geometric control theory was made along Hamiltonian lines. The port-controlled Hamiltonian approach is representative of this, and a huge amount of work has been done here. I refer to the papers of Dalsmo and van der Schaft [1998] and van der Schaft [2000] as representative of this work. A book that allies this Hamiltonian approach with topics of interest to readers of this journal is that of Stramigioli [2001].

Another significant body of research, this connected firmly with the origins of geometric mechanics, is that undertaken by researchers including John Baillieul, Anthony Bloch, Peter Crouch, P. S. Krishnaprasad, Naomi Leonard, Jerry Marsden, and Dmitry Zenkov. An account of some of this work can be found in the recent book of Bloch [2003].

The approach to mechanics and control taken in this paper is the one developed in detail in my book [Bullo and Lewis 2004] coauthored with Francesco Bullo from the University of California, Santa Barbara. In this book we develop from "first principles" an approach to the modelling of mechanical systems using Riemannian and affine differential geometry. After this we outline some of the control problems that can be addressed using the differential geometric modelling techniques. I will not attempt to summarise the content of this book. One of the aims, instead, is to try to give the reader an idea of whether reading this book is worth their effort. For some readers I believe it *will* be worth the effort, but for others it may not be.

2. Modelling mechanical systems using affine differential geometry

In this section I present a quick overview of Chapter 4 of [Bullo and Lewis 2004]. The intention is to introduce the major mathematical players in the affine differential geometric formulation of mechanics. A novice reader expecting to read this section and walk away understanding affine differential geometry and its use in mechanics will be disappointed. The fact of the matter is that, if you are starting with only scant knowledge of differential geometry, there is simply no way one can appreciate a geometric formulation of mechanics in an afternoon. It will take a significant investment of effort and, more importantly, a probable abandoning of certain preconceptions learned during the course of a more classical mechanics education. Thus, what I attempt to do in this section is alert the reader to the "brute facts" of a geometric formulation of mechanics. These are the things that they will have to come to grips with in order to merely get off the ground.

2.1. The configuration manifold. One of the main reasons why differential geometry is an effective tool for formulating mechanics is that mechanical systems evolve naturally on a configuration space that is a manifold. That is to say, the set of configurations of a mechanical system is in natural 1–1 correspondence with a set that has the structure of a differentiable manifold; this manifold is called the configuration space and typically denoted by Q. This idea is stated clearly by Bullo and Lewis [2004].

At this point I could define what a manifold is, but this seems pointless since this definition can be found in many places, including in [Bullo and Lewis 2004]. Moreover, the reader will probably have a good idea, or at least some intuition, about what a manifold is.

Instead, therefore, let me indicate why the notion of a manifold is important in robotics. The reader is surely familiar with the fact that the configurations of a rigid body fixed at a point in some inertial frame are in 1–1 correspondence with the set $\mathsf{SO}(3)$ of 3×3 orthogonal matrices with determinant 1. What is the best way to view $\mathsf{SO}(3)$? Some choose to parameterise $\mathsf{SO}(3)$ with Euler angles. This makes it appear as though $\mathsf{SO}(3)$ is some subset of 3-dimensional Euclidean space, although clearly it is not. Other representations of $\mathsf{SO}(3)$ are used (e.g., quaternion representations) to attempt to get a handle on $\mathsf{SO}(3)$. But from my point of view, $\mathsf{SO}(3)$ is a manifold and that is that. (Actually, $\mathsf{SO}(3)$ is a special manifold, since it is also a Lie group, and this Lie group structure is extremely important in robotics. However, I elect not to get into this aspect of rigid body modelling here, but refer to [Bullo and Lewis 2004, Murray, Li, and Sastry 1994].) A manifold is, essentially, a set that can locally be parameterised by an open subset Euclidean space (e.g., by Euler angles), and different parameterisations are required to be compatible with one another. A specific choice of parameterisation is sometimes helpful in working on a concrete example. And the existence of parameterisations in general is sometimes useful in that it allows you to believe, at least temporarily, that you are working with familiar calculus in Euclidean space (typically by saying, "Let $(q^1, \ldots, q^n)$ be coordinates for Q"). However, the important thing about manifolds is that, *when used properly*, they force a coordinate independent restriction on the problem formulation. Thus one is forced to deal with the real structure of the problem and not one tacked on, *ad hoc*, by a choice of coordinates.

Let me try to explain this in a concrete context. In Figure 1 I have depicted two simple systems, systems that are frequently used as benchmark-type systems to implement algorithms and so on. In the top left is a simple mobile robot. Let me suppose the three wheels are actuated, so this is really just a fully actuated planar rigid body. The objective is to steer the body to the $\times$ with a prescribed orientation. In the bottom right in Figure 1 I depict the classical pendulum/cart system where the objective is to balance the pendulum in the upright configuration by application of a force to the cart. One approach to modelling the mobile robot problem is to declare the coordinates describing the configurations of the system to be (x, y, θ), and then proceed as if the configuration space is $\mathbb{R}^3$. Similarly, for the pendulum/cart, one could take coordinates (x, θ) and proceed as if the configuration space is $\mathbb{R}^2$. In practice this is often accentuated by writing the coordinates as $[x \ y \ \theta]^T$ and $[x \ \theta]^T$ to drive home the fact the these are "coordinate vectors." But the configuration spaces are not $\mathbb{R}^3$ or $\mathbb{R}^2$ since, in each case, θ and $\theta + 2k\pi$ describe the same configuration for every integer k. The configuration spaces are, in fact, $\mathsf{SE}(2)$ (which is naturally diffeomorphic to $\mathbb{R}^2 \times \mathbb{S}^1$) and $\mathbb{R} \times \mathbb{S}^1$, where $\mathbb{S}^1$ is the unit circle. (Other representations of the configuration spaces are possible, but they will be equivalent to what I have given in the precise way that they are diffeomorphic.)

A good question is, "Is it important to make this distinction?" I believe that it is, and let me give one reason why. It is possible for a mechanical control system with a configuration space of $\mathbb{R}^n$ to possess a globally stable equilibrium point using continuous feedback. However, it is *not* possible for a mechanical control system with a configuration space having rotational degrees of freedom to have a globally stable equilibrium point using continuous feedback. In particular, there cannot be equilibria for the systems depicted in Figure 1 that are globally stable using continuous feedback. Despite this, the literature is filled with papers where continuous globally stabilising feedbacks are designed for the systems in Figure 1 or systems like them. That is to say, merely by believing that the

Figure 1. These systems are not globally stabilisable using continuous feedback!

configuration manifold is Euclidean space many researchers are led to claim something impossible. This may seem as if it is not an important point on which to be fussy, but Bhat and Bernstein [2000] illustrate clearly why this matter of rotational degrees of freedom are, in fact, something the control designer ought to be aware of.

2.2. The tangent bundle to the configuration manifold. The configuration manifold describes the set of configurations of a mechanical system. The tangent bundle of the configuration manifold describes the set of configurations and velocities of a mechanical system. Intuitively, the tangent space at a point $q \in \mathsf{Q}$, denoted by $\mathsf{T}_q\mathsf{Q}$, is the set of all velocities at this configuration. From this intuitive description I hope it is evident that $\mathsf{T}_q\mathsf{Q}$ is a vector space of the same dimension as Q. The tangent bundle, denoted by TQ, is the union of all tangent spaces. The tangent bundle itself is a differentiable manifold, and inherits from Q a differentiable structure in a natural way. Coordinates for TQ are often denoted by $((q^1, \ldots, q^n), (v^1, \ldots, v^n))$. Thus TQ has twice the dimension of Q. The natural differentiable structure for TQ works out nicely since it demands that, if one has two compatible sets of coordinates $(q^1, \ldots, q^n)$ and $(\tilde{q}^1, \ldots, \tilde{q}^n)$ for Q, then the corresponding

velocities obey the rather Chain Rule-like formula

$$\tilde{v}^j = \frac{\partial \tilde{q}^j}{\partial q^k} v^k, \qquad j \in \{1, \ldots, n\}.$$

(Here I use the summation convention commonly employed in differential geometry where summation is implied over indices that are repeated.) The way one should think of the tangent bundle is as the state space for the system. That is to say, the dynamics of the system are uniquely determined by initial conditions in the tangent bundle, this by virtue of the equations of motion being second-order (as we shall soon see, in case you have forgotten).

One piece of useful notation concerns the natural basis for a tangent space induced by a set of coordinates. Thus, if $(q^1, \ldots, q^n)$ are coordinates for Q, there is a natural basis for each tangent space which is denoted by $\{\frac{\partial}{\partial q^1}, \ldots, \frac{\partial}{\partial q^n}\}$. This basis has the property that a tangent vector with components $(v^1, \ldots, v^n)$ has the representation $v^j \frac{\partial}{\partial q^j}$. The peculiar partial derivative notation for a tangent vector is motivated by the fact that tangent vectors are, in an appropriate sense, differential operators. But this interpretation need not concern us here.

There is one caveat I would like to make concerning the tangent bundle. The representation of tangent bundle coordinates as $((q^1, \ldots, q^n), (v^1, \ldots, v^n))$ would seem to imply that there is some sort of "product structure" to the tangent bundle. That is to say, it appears that the tangent bundle is the Cartesian product of the configuration manifold and the set of velocities. *This is not the case.* It is true that TQ is *locally* a product, but even this local product representation is not natural; it arises from a specific choice of coordinates and different coordinates give different local product representations. Thus, while the coordinates for TQ seem to validate the notation of representing points in TQ as (q, v), this representation is, in a precise way, incorrect. It certainly is the case that one would like to be able to specify the configuration at which a velocity lives. A good way to do this is by writing a point in TQ as v_q, thus declaring by the notation that this is a tangent vector in the tangent space at $q \in Q$.

2.3. External forces. In a geometric setting for mechanics, forces take values in the cotangent bundle. Let me quickly describe the cotangent bundle. The cotangent space at $q \in Q$ is the dual space of $T_q Q$ (i.e., the set of $\mathbb{R}$-valued linear maps on the vector space $T_q Q$) and is denoted by $T_q^* Q$. The cotangent bundle is the union of all the cotangent spaces and is denoted by $T^* Q$. Often authors feel compelled to attempt to give some intuitive idea of what cotangent vectors are, typically because it is felt that cotangent vectors are somehow more mysterious than tangent vectors. My own view is that the definition is perfectly adequate, and one just has to accept it and learn that various physical quantities, for example forces, naturally are represented as elements of the cotangent bundle.

But why should forces take values in the cotangent bundle? Well, what does a force do? This is actually a fairly involved metaphysical question, but let me sidestep the metaphysics and just answer the question: A force does work on a system as it moves. Precisely, the situation is this. Suppose that a mechanical system undergoes a motion described by a curve $t \mapsto \gamma(t)$ in Q. As it undergoes this motion, a force $t \mapsto F(t)$ is applied to the system (being a little vague about just where F lives for the moment). The work done is then

$$W = \int_I F(t) \cdot \gamma'(t) \, dt,$$

where I is the time interval on which the motion is defined, i.e., the work is the integral of the product of force with velocity. I have just written the product $F(t) \cdot \gamma'(t)$ as if it means something, but let me consider what F must be in order for this expression to have meaning. The expression $F(t) \cdot \gamma'(t)$ should live in $\mathbb{R}$ (work is a scalar quantity) and should be linear in $\gamma'(t)$. Thus $F(t)$ can live no place other than $\mathsf{T}^*_{\gamma(t)}\mathsf{Q}$. This is the explanation I give about why forces are cotangent bundle-valued, and my undergraduate students seem to buy it. Note that the "$\cdot$" in the expression $F(t) \cdot \gamma'(t)$ is not the "dot product." It represents an element of the dual space acting on a vector.

Just as there is a natural basis for tangent vectors induced by a choice of coordinates, there is also a basis for cotangent vectors associated with coordinates $(q^1, \ldots, q^n)$. This basis is denoted by $\{\mathrm{d}q^1, \ldots, \mathrm{d}q^n\}$ so that a force has the coordinate representation $F_j \mathrm{d}q^j$. It is not entirely obvious how one goes from the representation of forces in Newtonian mechanics to the geometric representation. This is discussed in [Bullo and Lewis 2004, Murray, Li, and Sastry 1994].

Forces in mechanics, at least when one is considering the control theory of mechanical systems, come in two flavours: external forces coming from the environment (potential forces, friction and dissipative forces, etc.) and control forces. Let me denote the total uncontrolled external force by F, noting that F may be a function of position, velocity, and/or time. The total control force I write as a linear combination of forces $F^1, \ldots, F^m$. Thus the control force is

$$\sum_{a=1}^{m} u^a F^a,$$

where I ask that the forces F^a be functions only of position, but allow the controls u^a to be functions of position, velocity, and/or time, as necessitated by the control scheme used.

2.4. The kinetic energy metric. The key to the affine differential geometric approach to mechanics is the fact that the kinetic energy of the system defines a Riemannian metric on the configuration space. Researchers in robotics will know this Riemannian metric as the "inertia matrix," the "mass matrix," or something similar. In the classical way of looking at things, the kinetic energy metric is a symmetric positive-definite matrix-valued function of configuration. In the geometric way of looking at things, the kinetic energy metric is a smooth assignment of an inner product to each tangent space of the configuration space. In differential geometry such an assignment is called a Riemannian metric. In the particular case where the assignment comes from the kinetic energy for a mechanical system I shall call the corresponding Riemannian metric the kinetic energy metric.

The notation used in [Bullo and Lewis 2004] for the kinetic energy metric is $\mathbb{G}$. Thus the inner product on the tangent space at $q \in \mathsf{Q}$ is denoted by $\mathbb{G}(q)$ (in the geometric formulation, the kinetic energy metric is a genuine map, it being a section of a certain tensor bundle), and the inner product of two tangent vectors $u_q, v_q \in \mathsf{T}_q\mathsf{Q}$ is denoted by $\mathbb{G}(q)(u_q, v_q)$, or simply by $\mathbb{G}(u_q, v_q)$ since the point $q \in \mathsf{Q}$ is explicit in the notation already. Therefore, the kinetic energy is the function on TQ given by

$$\mathrm{KE}(v_q) = \frac{1}{2}\mathbb{G}(v_q, v_q).$$

At a point $q \in \mathsf{Q}$, $\mathbb{G}(q)$ is an inner product on the tangent space $\mathsf{T}_q\mathsf{Q}$. Therefore, it is

defined by the n^2 components $\mathbb{G}_{jk} = \mathbb{G}(\frac{\partial}{\partial q^j}, \frac{\partial}{\partial q^k})$, $j, k \in \{1, \ldots, n\}$. The way in which one represents the kinetic energy metric in coordinates is then $\mathbb{G} = \mathbb{G}_{jk} \mathrm{d}q^j \otimes \mathrm{d}q^k$.

Once one has at hand a physical model for the system, including useful representations of the configuration space and its tangent bundle, the actual task of writing the kinetic energy metric is actually quite simple, at least in principle. Of course, in practice there are genuine issues involved with determining the inertial properties of one's system, etc., but there is not much one can say about these matters, within or without a geometric formulation of mechanics.

There is a useful bit of notation associated with a Riemannian metric $\mathbb{G}$. Since I have found that this notation can be off-putting to many readers, let me devote a few moments to an apology for it. Given a Riemannian metric $\mathbb{G}$ on Q there are naturally defined maps $\mathbb{G}^\flat \colon \mathsf{T}\mathsf{Q} \to \mathsf{T}^*\mathsf{Q}$ and $\mathbb{G}^\sharp \colon \mathsf{T}^*\mathsf{Q} \to \mathsf{T}\mathsf{Q}$ defined by

$$\mathbb{G}^\flat(v_q) \cdot u_q = \mathbb{G}(u_q, v_q), \qquad u_q, v_q \in \mathsf{T}_q\mathsf{Q},$$
$$\alpha_q \cdot \mathbb{G}^\sharp(\beta_q) = \mathbb{G}^{-1}(\alpha_q, \beta_q), \qquad \alpha_q, \beta_q \in \mathsf{T}^*_q\mathsf{Q},$$

where by $\mathbb{G}^{-1}$ I denote the inner product on $\mathsf{T}^*_q\mathsf{Q}$ induced by that on $\mathsf{T}_q\mathsf{Q}$. If $\mathbb{G}_{jk}$, $j, k \in \{1, \ldots, n\}$, are the components of $\mathbb{G}$ in a set of coordinates, then $\mathbb{G}^\flat$ is simply multiplication by this matrix of components. If $\mathbb{G}^{jk}$, $j, k \in \{1, \ldots, n\}$, are the components of the inverse of the matrix with components $\mathbb{G}_{jk}$, then $\mathbb{G}^\sharp$ is multiplication by this inverse matrix. The reason for the funny notation $\flat$ and $\sharp$ is that, in terms of indices, $\mathbb{G}^\flat$ takes an object whose indices are superscripts (a tangent vector) and converts them to an object whose indices are subscripts (a cotangent vector). Thus $\mathbb{G}^\flat$ "flattens" the indices. In a like manner, $\mathbb{G}^\sharp$ raises or "sharpens" the indices.

2.5. The Levi-Civita connection. The components of the description of a mechanical system I have described thus far—the configuration manifold, its tangent bundle, forces, and the kinetic energy metric–will be well known to anyone in robotics, although not always necessarily in just the way that I think of these things. However, the Levi-Civita connection that I will describe now is, if you are not familiar with it, really something new. While it can be related to some things that may be familiar, it *is* precisely none of these things. The Levi-Civita affine connection is what it is, and, for those unfamiliar with it, this will take a little hard swallowing.

Since I am making the affine connection formulation of mechanics the centrepiece of this article, I ought to say what an affine connection is. I suppose the reader to be familiar with the basic concept of a vector field on a manifold and some of the notation surrounding the differential geometry of dealing with vector fields. Let me work with a manifold Q (which may or not be the configuration manifold for a mechanical system). An affine connection[1] on Q is a mapping that assigns to two vector fields X and Y a third vector field denoted $\nabla_X Y$, and this assignment has the following properties:

1. the map $(X, Y) \mapsto \nabla_X Y$ is $\mathbb{R}$-bilinear;

2. for vector fields X and Y and a function f we have $\nabla_{fX} Y = f \nabla_X Y$;

---

[1] What I call an "affine connection" is called a "linear connection" by some authors, and these authors have something else (something that is related, however) in mind when they say "affine connection." However, the use of "affine connection" as I use it is fairly widespread.

3. for vector fields X and Y and a function f we have $\nabla_X(fY) = f\nabla_X Y + (\mathscr{L}_X f)Y$, where $\mathscr{L}_X$ is the Lie derivative with respect to X.

The vector field $\nabla_X Y$ is called the covariant derivative of Y with respect to X. So this is the definition, and how very uninsightful it is! However, the notion of an affine connection comes up naturally in mechanics, and I will explain this as we go along. Let me first say a few general things about affine connections.

1. An affine connection is extra structure on a manifold. That is to say, manifolds do not come equipped with a nice natural affine connection that you can go ahead and start using. You have to introduce your own. Thankfully, in mechanics, there are affine connections that come from the mechanical problem data. We shall see how these arise.

2. An affine connection ∇ is described in coordinates $(q^1, \ldots, q^n)$ by writing

$$\nabla_{\frac{\partial}{\partial q^j}} \frac{\partial}{\partial q^k} = \Gamma^l_{jk} \frac{\partial}{\partial q^l}, \qquad j, k \in \{1, \ldots, n\},$$

which is always possible for n^3 functions Γ^l_{jk}, $j, k, l \in \{1, \ldots, n\}$. These are called the Christoffel symbols for the affine connection and they clearly depend on a choice of coordinates. Using these, the covariant derivative of two general vectors fields is easily written:

$$\nabla_X Y = \left(\frac{\partial Y^l}{\partial q^j} X^j + \Gamma^l_{jk} X^j Y^k \right) \frac{\partial}{\partial q^l}.$$

3. Associated with every Riemannian metric $\mathbb{G}$ is a unique affine connection, called the Levi-Civita connection and denoted (by me) as $\overset{\mathrm{G}}{\nabla}$. There is a nice intrinsic characterisation of $\overset{\mathrm{G}}{\nabla}$ and for this I refer the reader to [Bullo and Lewis 2004]. Here let me give a hands-on, but rather grimy, characterisation by giving the formula for its Christoffel symbols:

$$\overset{\mathrm{G}}{\Gamma}^l_{jk} = \frac{1}{2} \mathbb{G}^{lr} \left(\frac{\partial \mathbb{G}_{rj}}{\partial q^k} + \frac{\partial \mathbb{G}_{rk}}{\partial q^j} - \frac{\partial \mathbb{G}_{jk}}{\partial q^r} \right), \qquad j, k, l, \in \{1, \ldots, n\}.$$

We shall see in Section 2.7 where this affine connection arises from.

4. Associated with an affine connection ∇ are special curves on $\mathbb{Q}$ called the geodesics of ∇. These admit a nice intrinsic description too, but let me simply give the differential equation in coordinates that these curves must satisfy:

$$\ddot{q}^l(t) + \Gamma^l_{jk} \dot{q}^j(t) \dot{q}^k(t) = 0, \qquad l \in \{1, \ldots, n\}. \tag{2.1}$$

In coordinate-free language this equation can be written as $\nabla_{\gamma'(t)} \gamma'(t) = 0$. Thus $\nabla_{\gamma'(t)} \gamma'(t)$ is the slick way of writing what is on the left-hand side of (2.1), where $t \mapsto \gamma(t)$ is the curve that in coordinates has the form $t \mapsto (q^1(t), \ldots, q^n(t))$.

It is through geodesics that one can most easily begin to get at what an affine connection "means." So let me say a few words about this.

(a) If $\nabla = \overset{\mathrm{G}}{\nabla}$ is the Levi-Civita connection, then geodesics are those curves which locally minimise length.

(b) Note that the coordinate expression in (2.1) for $\nabla_{\gamma'(t)} \gamma'(t)$ is a sum of "$\ddot{q}^l$" and "$\Gamma^l_{jk} \dot{q}^j(t) \dot{q}^k(t)$." *It is an act of unspeakable violence to break these terms apart.*

These terms always belong together, and together they mean "$\nabla_{\gamma'(t)}\gamma'(t)$." Separately they mean absolutely, positively nothing.[2] One consequence of this is that the expression "$\nabla_{\gamma'(t)}\gamma'(t)$" gives a way of understanding the acceleration along the curve γ. But note that this expression of acceleration is not intrinsic as it depends on a choice of affine connection. Some lengths are taken in Section 4.3.5 of [Bullo and Lewis 2004] to describe this.

The primary message of the preceding discussion is this:

Associated in a unique way with the kinetic energy metric $\mathbb{G}$ for a mechanical system is an affine connection called the Levi-Civita connection. The importance of this will be made clear in Section 2.7.

2.6. Nonholonomic constraints. The final ingredient in my description of mechanics before I write the equations of motion is nonholonomic constraints. Sometimes it seems to me as if there are as many techniques for deriving the equations of motion in the presence of nonholonomic constraints as there are people who have thought about doing this. These methods are, at least the correct ones, all distinguished by one simple fact: *they are all the same!* That is to say, Method X for writing equations of motion for a system subject to nonholonomic constraints is correct if and only if it agrees with Newton's and Euler's equations in the intersection where Method X and Newton's and Euler's equations apply. Many authors claim mystical advantages of their correct technique over other (necessarily equivalent) correct techniques. The mystical advantages too often are described in some metaphysical form, and Newton and Euler do not even get mentioned. My take on this is: If your technique is equivalent to Newton and Euler, then you are correct and so is everyone else whose method is equivalent to Newton and Euler. There can be real issues in practice, such as coming up with a technique that is easy to implement. By no means do I wish to assert that these practical considerations do not merit consideration. However, there are no "in principle" advantages of any one correct method over any other correct method.

That all being said, let me just say here what I mean when I say "nonholonomic constraint," and how these constraints get accounted for in the equations of motion. A nonholonomic constraint on a configuration manifold Q is a distribution $\mathcal{D}$ on Q. That is, we select a subspace $\mathcal{D}_q$ of each tangent space $\mathsf{T}_q\mathsf{Q}$ that describes the set of velocities admissible at that configuration. Generalisations are possible, for example, to time-dependent or nonlinear constraints. But this complication adds little to the conceptual picture.

The way in which nonholonomic constraints get incorporated into the equations of motion is by adding a force, called the constraint force, to the equations of motion. A constraint force is, by definition, one that does no work on admissible velocities. That is, at a configuration $q \in \mathsf{Q}$, a force $\lambda_q \in \mathsf{T}_q^*\mathsf{Q}$ is a constraint force if $\lambda_q \cdot v_q = 0$ for every $v_q \in \mathcal{D}_q$. The idea, as we shall see clearly in the next section, is that a constraint force is introduced into the problem as an unknown to be solved for, a little like, but not exactly like, a Lagrange multiplier.

---

[2]Okay, I am telling a fib here for emphasis. The expression $\ddot{q}^i$ makes sense as a local coordinate for the second jet bundle. But I do not wish to go in that direction, so let me stick to my expression of the unspeakable violence of rendering these terms asunder, since you cannot be led wrong in accepting it.

2.7. The equations of motion. As I suggested in the preceding section, the equations of motion for a mechanical system are given by Newton's and Euler's equations that are based on simple laws of force and moment balance. The precise variant of Newton's and Euler's equations one wishes to use—for example, the standard Euler–Lagrange equations, the Gibbs–Appell equations, Gauss's Principle of Least Action, the Principle of Virtual Work—is subject to taste and possibly other considerations. For my approach, the key is the following result.

2.1 Theorem: (General equations of motion) *Consider a mechanical system with configuration manifold* Q*, kinetic energy metric* $\mathbb{G}$*, external force* F*, control force* $\sum_{a=1}^{m} u^a F^a$*, and nonholonomic constraint* $\mathcal{D}$*. For a curve* $t \mapsto \gamma(t)$ *the following are equivalent:*

  (i) the force/moment balance equations of Newton/Euler hold along γ*;*

  (ii) γ *satisfies the equations*

$$\mathbb{G}^{\flat}(\overset{\mathrm{G}}{\nabla}_{\gamma'(t)}\gamma'(t)) = F(t, \gamma'(t)) + \sum_{a=1}^{m} u^a(t) F^a(\gamma(t)) + \lambda(t),$$

$$\gamma'(t) \in \mathcal{D}_{\gamma(t)},$$

    where $t \mapsto \lambda(t)$ *is a constraint force along* γ*.*

For unconstrained systems, i.e., where the nonholonomic constraint is $\mathcal{D} = \mathsf{TQ}$, this immediately gives the following corollary.

2.2 Corollary: (Equations of motion for unconstrained systems) *Consider an unconstrained mechanical system with configuration manifold* Q*, kinetic energy metric* $\mathbb{G}$*, external force* F*, and control force* $\sum_{a=1}^{m} u^a F^a$*. For a curve* $t \mapsto \gamma(t)$ *the following are equivalent:*

  (i) the force/moment balance equations of Newton/Euler hold along γ*;*

  (ii) γ *satisfies the equations*

$$\mathbb{G}^{\flat}(\overset{\mathrm{G}}{\nabla}_{\gamma'(t)}\gamma'(t)) = F(t, \gamma'(t)) + \sum_{a=1}^{m} u^a(t) F^a(\gamma(t)).$$

It is perhaps worth comparing these equations with the equations of motion in the usual form seen in the robotics and mechanics literature. The usual form of the equations is something like

$$\boldsymbol{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{C}(\boldsymbol{q}, \dot{\boldsymbol{q}})\dot{\boldsymbol{q}} + \boldsymbol{N}(t, \boldsymbol{q}, \dot{\boldsymbol{q}}) = \boldsymbol{Bu}.$$

In components the equations of Corollary 2.2 are

$$\mathbb{G}_{rl}\ddot{q}^l + \mathbb{G}_{rl}\overset{\mathrm{G}}{\Gamma}^l_{jk}\dot{q}^j\dot{q}^k = F_r(t, q, \dot{q}) + \sum_{a=1}^{m} u^a F^a_r(q), \qquad r \in \{1, \ldots, n\}.$$

Thus we can directly compare the terms in the two equations as per Table 1. Note that I *do not* separate $\boldsymbol{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{C}(\boldsymbol{q}, \dot{\boldsymbol{q}})\dot{\boldsymbol{q}}$, and neither should you, given my comments in item b in Section 2.5.

For systems with nonholonomic constraints the usual way to proceed is to eliminate the constraint force in some way. In terms of the affine connection picture I am painting, the magical thing is that, after elimination of the constraint force, one still gets equations

Table 1. Comparison of terms in "classical," "geometric component," and "intrinsic" formulation of equations of motion

Classical form	Geometric coordinate form	Intrinsic form
$\boldsymbol{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{C}(\boldsymbol{q},\dot{\boldsymbol{q}})\dot{\boldsymbol{q}}$	$\mathbb{G}_{rl}\ddot{q}^l + \mathbb{G}_{rl}\overset{\mathbb{G}}{\Gamma}{}^l_{jk}\dot{q}^j\dot{q}^k$	$\overset{\mathbb{G}}{\nabla}_{\gamma'(t)}\gamma'(t)$
$\boldsymbol{N}(t,\boldsymbol{q},\dot{\boldsymbol{q}})$	$-F_r(t,q,\dot{q})$	$-F(t,\gamma'(t))$
$\boldsymbol{B}\boldsymbol{u}$	$\sum_{a=1}^m u^a F_r^a(q)$	$\sum_{a=1}^m u^a F^a(\gamma(t))$

involving an affine connection, although it is not generally the Levi-Civita connection for some Riemannian metric. I will not explain the way this works here since the most direct explanation involves manipulations with affine connections that are one step beyond elementary, and I do not want to get into too much intricate stuff. The elimination of the constraint force also results in a modification of the external and control forces. Let me just summarise descriptively the final tagline and refer to [Bullo and Lewis 2004] for a thorough discussion of how this works.

For a mechanical system with configuration manifold Q, *kinetic energy metric* $\mathbb{G}$, *external force F, control force* $\sum_{a=1}^m u^a F^a$, *and nonholonomic constraint* $\mathcal{D}$, *there exists (1) an affine connection* $\overset{\mathcal{D}}{\nabla}$ *depending on* $\mathbb{G}$ *and* $\mathcal{D}$, *(2) a* $\mathsf{T}\mathsf{Q}$-*valued map Y depending on F,* $\mathbb{G}$, *and* $\mathcal{D}$, *and (3) vector fields $Y_1,\dots,Y_m$, depending on* $F^1,\dots,F^m$, $\mathbb{G}$, *and* $\mathcal{D}$, *such that, for a curve $t \mapsto \gamma(t)$ the following are equivalent:*

(i) the force/moment balance equations of Newton/Euler hold along γ;

(ii) γ satisfies the equations

$$\overset{\mathcal{D}}{\nabla}_{\gamma'(t)}\gamma'(t) = Y(t,\gamma'(t)) + \sum_{a=1}^m u^a(t)Y_a(\gamma(t)), \qquad (2.2)$$

$$\gamma'(0) \in \mathcal{D}_{\gamma(0)}.$$

Therefore, the point is that the equations of motion are given by the equations (2.2), and a fundamental rôle in the equations is played by an affine connection.

2.8. An example. In this section I will go quickly through an example to illustrate how the geometric concepts arise in a concrete context. The example I use is the so-called "snakeboard" system which is a model for a commercially available skateboard-like product.[3] The mathematical model I give originated in the paper of Lewis, Ostrowski, Murray, and Burdick [1994]. Here I shall present an affine differential geometric formulation of the equations of motion; further details on this model can be found in [Bullo and Lewis 2003, Bullo and Lewis 2004].

The model for the snakeboard is shown in Figure 2. The idea of the system is that there is a central coupler, at each end of which are wheels whose angle relative to the coupler is

---

[3]http://www.snakeboard.com/

Figure 2. The snakeboard

controlled. Atop the coupler, between the wheels, is an inertial rotor that is also controlled. Thus the system has five degrees of freedom and two controls. This is an example of a system with nonholonomic constraints since one places on the velocities of the system the constraint that the wheels roll without slipping.

Configuration manifold

In the model I consider, the angle of the front and back wheels relative to the coupler are fixed as being equal. This means that a configuration of the system is specified by (1) the position of the coupler, (2) the angle of the wheels, and (3) the angle of the rotor. The coupler is, essentially, a planar rigid body and it can take any configuration of a planar rigid body, i.e., a point in $\mathsf{SE}(2) \simeq \mathbb{R}^2 \times \mathbb{S}^1$. The wheels are then specified by an angle, i.e., by a point in $\mathbb{S}^1$. Similarly, the rotor angle is determined by a point in $\mathbb{S}^1$. The configuration space is then $\mathsf{Q} = \mathsf{SE}(2) \times \mathbb{S}^1 \times \mathbb{S}^1$. I use coordinates $(x, y, \theta, \psi, \phi)$ for Q, where the physical meaning of these coordinates is as indicated in Figure 2 (and with (x, y) being the position of the centre of mass of the coupler).

Note that in simple planar systems like the snakeboard, it is often easy to "by hand" determine the configuration manifold and useful coordinates for it. However, for complicated interconnected systems, this process can often be a challenge. Some guidance for doing this in a systematic way can be found in [Bullo and Lewis 2004, Murray, Li, and Sastry 1994]. Moreover, the modelling techniques outlined in [Bullo and Lewis 2004] make this step of determining the configuration manifold the step of paramount importance. That is to say, once this step is done properly, the other facets of the model will often fall into place fairly easily.

The tangent bundle to the configuration manifold

The coordinates $(x, y, \theta, \psi, \phi)$ induce natural coordinates $((x, y, \theta, \psi, \phi), (v_x, v_y, v_\theta, v_\psi, v_\phi))$ for the tangent bundle TQ. A point in the tangent bundle can be written as

$$v_x \frac{\partial}{\partial x} + v_y \frac{\partial}{\partial y} + v_\theta \frac{\partial}{\partial \theta} + v_\psi \frac{\partial}{\partial \psi} + v_\phi \frac{\partial}{\partial \phi},$$

relative to the basis of vector fields $\{\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial \theta}, \frac{\partial}{\partial \psi}, \frac{\partial}{\partial \phi}\}$ induced by the coordinate chart.

External forces

In the simple model I consider here, the only external forces are the control forces. These consist of torques applied to the wheels and the rotor. A general force F is then written in coordinates as

$$F = F_x \mathrm{d}x + F_y \mathrm{d}y + F_\theta \mathrm{d}\theta + F_\psi \mathrm{d}\psi + F_\phi \mathrm{d}\phi.$$

In the particular case of our control forces, a torque applied to the rotor will have the form $F^1 = u_\psi \mathrm{d}\psi$ and a torque applied to the wheels will have the form $F^2 = u_\phi \mathrm{d}\phi$. Here u_ψ and u_ϕ are to be thought of as the controls, and so are functions of time, position, and/or velocity, depending on what control task one is carrying out.

Kinetic energy metric

I use the following physical parameters:

m_c mass of coupler;
m_r mass of rotor;
m_w mass of each wheel assembly;
J_c inertia of coupler about center of mass;
J_r inertia of rotor about center of mass;
J_w inertia of wheel assembly about center of mass;
ℓ distance from coupler center of mass to wheel assembly.

The kinetic energy is easily determined by "elementary methods" to be

$$\mathrm{KE} = \frac{1}{2}(m_c + m_r + 2m_w)(v_x^2 + v_y^2) + \frac{1}{2}(J_c + J_r + 2(J_w + m_w \ell^2))v_\theta$$
$$+ \frac{1}{2}J_r v_\psi^2 + J_w v_\phi^2 + J_r v_\theta v_\psi.$$

In differential geometric notation the kinetic energy metric is then given by

$$\mathbb{G} = (m_c + m_r + 2m_w)(\mathrm{d}x \otimes \mathrm{d}x + \mathrm{d}y \otimes \mathrm{d}y) + (J_c + J_r + 2(J_w + m_w \ell^2))\mathrm{d}\theta \otimes \mathrm{d}\theta$$
$$+ J_r \mathrm{d}\psi \otimes \mathrm{d}\psi + 2J_w \mathrm{d}\phi \otimes \mathrm{d}\phi + J_r(\mathrm{d}\theta \otimes \mathrm{d}\psi + \mathrm{d}\psi \otimes \mathrm{d}\theta).$$

The Levi-Civita connection

Since the components of the kinetic energy metric are constants, independent of the coordinates we are using, the Christoffel symbols of the Levi-Civita connection associated to $\mathbb{G}$ are all identically zero. However, since this is a system with nonholonomic constraints, the Levi-Civita connection, *per se*, is not of much interest.

Nonholonomic constraints

The nonholonomic constraints for the snakeboard are prescribed by no slip conditions on the wheels. These conditions are given by the two constraints

$$-\sin(\theta + \phi)v_x + \cos(\theta + \phi)v_y - \ell\cos\phi v_\theta = 0,$$
$$-\sin(\theta - \phi)v_x + \cos(\theta - \phi)v_y + \ell\cos\phi v_\theta = 0.$$

To write a basis of vector fields that satisfy the constraints let me introduce the vector field

$$V = \cos\theta\frac{\partial}{\partial x} + \sin\theta\frac{\partial}{\partial y}.$$

One can then check by direct computation that the three vector fields

$$X_1 = \ell\cos\phi V - \sin\phi\frac{\partial}{\partial\theta}, \quad X_2' = \frac{\partial}{\partial\psi}, \quad X_3 = \frac{\partial}{\partial\phi}$$

form a basis for the set of velocities satisfying the constraints (after noting that one must exclude points in Q where $\phi = \pm\frac{\pi}{2}$ since, at these points, the dimension of the subspace of admissible velocities changes from three to four). Below we will see that it is useful to use a basis for the admissible velocities that are orthogonal with respect to G. One can check that if we define

$$X_2 = \frac{J_r\ell\cos\phi\sin\phi}{c_1(\phi)}V - \frac{J_r\sin^2\phi}{c_1(\phi)}\frac{\partial}{\partial\theta} + \frac{\partial}{\partial\psi},$$

where

$$c_1(\phi) = (m_c + m_r + 2m_w)\ell^2\cos^2\phi + (J_c + J_r + 2(J_w + m_w\ell^2))\sin^2\phi,$$

then the vector fields $\{X_1, X_2, X_3\}$ form a G-orthogonal basis for the set of velocities satisfying the constraints.

Constraint forces are those which annihilate the velocities satisfying the constraints. Thus a typical constraint force has the form

$$\lambda_1(-\sin(\theta + \phi)\mathrm{d}x + \cos(\theta + \phi)\mathrm{d}y - \ell\cos\phi\mathrm{d}\theta)$$
$$+ \lambda_2(-\sin(\theta - \phi)\mathrm{d}x + \cos(\theta - \phi)\mathrm{d}y + \ell\cos\phi\mathrm{d}\theta)$$

for some coefficients λ_1 and λ_2.

The equations of motion

As I mentioned above, there are many ways to write the equations of motion for a mechanical system subject to nonholonomic constraints. To represent the equations of motion using an affine connection as indicated in the equations (2.2), I follow the approach of Bullo and Žefran [2002] which is explained in detail in [Bullo and Lewis 2004]. The idea of this approach is that one uses coordinates for TQ that are *not* the natural coordinates, but rather coordinates adapted to the constraints. Thus one chooses a basis $\{X_1, \ldots, X_n\}$ of vector fields, the first r of which are a basis for the admissible velocities allowed by the constraints. In the particular version of this approach presented by Bullo and Žefran [2002], one uses a G-orthogonal basis of vector fields. For the snakeboard we thus use the orthogonal basis

$\{X_1, X_2, X_3\}$ of vector fields for the admissible velocities and write a typical admissible velocity as

$$\dot{q}^j(t)\frac{\partial}{\partial q^j} = v^\alpha(t)X_\alpha(\gamma(t)),$$

so defining v^1, v^2, and v^3. The functions v^1, v^2, and v^3 are often known in the literature as "pseudo-velocities." One can verify that

$$v^1 = \cos\theta\cos\phi\dot{x} + \sin\theta\cos\phi\dot{y} - l\sin\phi\dot{\theta} - \frac{J_r\ell\sin\phi}{c_1(\phi)}\dot{\psi},$$
$$v^2 = \dot{\psi},$$
$$v^3 = \dot{\phi},$$

provided that $(\dot{x}, \dot{y}, \dot{\theta}, \dot{\psi}, \dot{\psi})$ satisfy the constraints. Moreover, using the methodology outlined in [Bullo and Lewis 2004] one can compute the equations of motion to be

$$\dot{x} = \ell\cos\phi\cos\theta v^1 + \frac{J_r\ell\cos\phi\sin\phi\cos\theta}{c_1(\phi)}\dot{\psi},$$
$$\dot{y} = \ell\cos\phi\sin\theta v^1 + \frac{J_r\ell\cos\phi\sin\phi\sin\theta}{c_1(\phi)}\dot{\psi},$$
$$\dot{\theta} = -\sin\phi v^1 - \frac{J_r\sin^2\phi}{c_1(\phi)}\dot{\psi},$$
$$\dot{v}^1 = -\frac{(J_c + J_r + 2J_w - (m_c + m_r)\ell^2)\cos\phi\sin\phi}{c_1(\phi)}v^1\dot{\phi}$$
$$- \frac{J_r(m_c + m_r + 2m_w)\ell^2\cos\phi}{c_1(\phi)^2}\dot{\psi}\dot{\phi},$$
$$\ddot{\psi} = \frac{(m_c + m_r + 2m_w)\ell^2\cos\phi}{c_2(\phi)}v^1\dot{\phi}$$
$$+ \frac{J_r(m_c + m_r + 2m_w)\ell^2\cos\phi\sin\phi}{c_1(\phi)c_2(\phi)}\dot{\psi}\dot{\phi} + \frac{c_1(\phi)}{J_rc_2(\phi)}u_\psi,$$
$$\ddot{\phi} = \frac{1}{2J_w}u_\phi,$$

where

$$c_2(\phi) = (m_c + m_r + 2m_w)\ell^2\cos^2\phi + (J_c + 2(J_w + m_w\ell^2))\sin^2\phi.$$

These equations may not appear to be derived from the geodesic equations (2.1) since the first three equations are first-order in configuration. This, however, is simply a manifestation of the technique of using pseudo-velocities.

3. Some problems whose solution is facilitated by using differential geometry

The bottom line of the preceding section is equations (2.2), which include the equations of motion for a large class of mechanical systems. These equations are nice since they are so compact, general, and revealing of the geometric structure of the problem. However, the fact of the matter is that, if this is the end of what one can do with the affine connection

formalism, then serious questions can be raised about whether there is a payoff in learning the mathematics necessary to understand the formulae. Thankfully, there are problems beyond the mere formulation of the equations of motion that are effectively addressed within the differential geometric framework. In this section I investigate some of these. Since everything I will talk about here has already appeared in the literature, and since a detailed presentation of the results would necessitate reproducing huge chunks of the papers in question, I instead *describe* the results in the papers in question rather than reproduce them. I think that this approach is more in keeping with the objectives of the paper.

3.1. Controllability. One of the fundamental problems in the control of mechanical systems is, "Can a system be steered from Point A to Point B?" This is a controllability question. For systems that are fully actuated, i.e., when the control forces $F^1, \ldots, F^m$ span $\mathsf{T}^*_q\mathsf{Q}$ for each $q \in \mathsf{Q}$, these controllability questions are trivial since one can always apply a force to make the system go wherever you want. However, for underactuated systems, controllability questions are generally extremely difficult. Moreover, if the system is underactuated *and* if there are no uncontrolled external forces (e.g., no potential forces, so ruling out the multitude of pendulum systems that are often considered), then controllability becomes a difficult problem. The approach I take here is to give a brief general discussion of the sort of controllability problems I will consider, and then look at an example where I outline the controllability that can be deduced based on some more or less recent work that is heavily steeped in techniques of differential geometry.

The controllability problem I consider is purely a local one. Namely, I consider a mechanical system with configuration manifold Q and whose equations of motion are

$$\nabla_{\gamma'(t)}\gamma'(t) = \sum_{a=1}^{m} u^a(t)Y_a(\gamma(t)), \tag{3.1}$$

recalling from Section 2.7 that this equation models mechanical systems, possibly with nonholonomic constraints, that are subject to no uncontrolled external forces. I let $q_0 \in \mathsf{Q}$ and denote by $\mathcal{R}_\mathsf{Q}(q_0, \leq T)$ the set of configurations reachable from q_0 in time at most T. The system is small-time locally configuration controllable (STLCC) from q_0 if q_0 is an interior point of $\mathcal{R}_\mathsf{Q}(q_0, \leq T)$. This definition is a bit pared down from the most general definition, and I refer to [Bullo and Lewis 2004] for the full details. The results I will refer to in this section are from the work of Bullo and Lewis [2005], Lewis [2000], Lewis and Murray [1997], and Lewis and Tyner [2004]. All of the results in these papers rely heavily on the affine connection structure of the system equations (3.1).

As an example of the sort of system whose controllability can be described by the results I refer to, I consider the planar rigid body system depicted in Figure 3. This is a system moving on a planar surface orthogonal to the direction of the gravitational field, and I assume that there are no friction forces. One might want to think of this system as an idealised model of a hovercraft whose motive force is supplied by a single thrust fan for which the direction and magnitude of the force can be varied. One problem of interest for this system is to steer the system from rest in a configuration q_A to rest in a configuration q_B. In order to do this, one would require that q_B be in the set of configurations reachable from q_A. Thus the controllability problem is fundamental to the motion planning problem. In [Lewis and Murray 1997] it is shown that this system is, indeed, STLCC from every configuration.

Figure 3. A controlled planar body

The results in [Lewis and Murray 1997] are an affine differential geometric formulation of general accessibility results of Sussmann and Jurdjevic [1972] and local controllability results of Sussmann [1987].

Let me now consider a slight alteration of the system by supposing that the direction of the force is now fixed, say at an angle of $\frac{\pi}{2}$ as shown in Figure 4. Effectively, this


Figure 4. A controlled planar body with a single control input

system now has a single-input. It turns out that the controllability of this system cannot be determined using the results of Lewis and Murray [1997]. However, these results were refined in a certain direction in [Bullo and Lewis 2005] where it is shown that this system, and indeed any single-input system of the form (3.1), is controllable from at most a "thin" set[4] of points.

Finally let me consider another alteration of the model as depicted in Figure 5. For this model, the fan is assumed to have inertia, so this adds an extra degree of freedom to the model. Once again, this system was one whose controllability I was unable to analyse when I first encountered it. However, after some work by Lewis and Tyner [2004] it was shown that this system was only STLCC from the configurations shown in Figure 6 and $\mathsf{SE}(2)$-translations of these configurations.

The point of the presentation in this section might be said to be the following.

---

[4]Precisely, a set that is locally the set of zeros of a finite collection of analytic functions.

Figure 5. A controlled planar body with the fan dynamics mod-
elled

Figure 6. The configurations from which the planar body with fan
dynamics is controllable

1. Underactuated mechanical systems often have subtle local controllability properties.
   This is clearly evidenced in this section by the difficulty of determining the controllability
   of a simple planar body system with some simple, natural variations in the model.
2. By understanding the geometric structure of the model of the system, particularly the
   rôle played by the affine connection, it is possible to develop general results that allow
   the analysis of some nontrivial examples.

   Moreover, the controllability results referred to in this section lead naturally to some

techniques in motion planning that I now describe.

3.2. Trajectory planning using decoupling vector fields. In this section I shall describe what can be achieved in terms of converting the controllability results used in the preceding section to motion planning results. The idea I describe here was motivated by explorations of Arai, Tanie, and Shiroma [1998] and Lynch, Shiroma, Arai, and Tanie [2000]. The initial geometric presentation is due to Bullo and Lynch [2001] and the most polished form of the results appear in [Bullo and Lewis 2005]. The idea is fairly simple. For a given mechanical system whose equations of motion are given by

$$\nabla_{\gamma'(t)}\gamma'(t) = \sum_{a=1}^{m} u^a(t)Y_a(\gamma(t)),\tag{3.2}$$

one seeks vector fields $X_1, \ldots, X_k$, called decoupling vector fields, that have the property that their integral curves, and any reparameterisations of these integral curves, can be followed by a trajectory of the system (3.2). One then hopes to find sufficiently many such vector fields that any motion planning problem can be solved by concatenating their integral curves.

Again, rather than restating the results which are given in the language of affine differential geometry, let me give a problem for which the ideas lead to a solution of the motion planning problem. The system I consider is the snakeboard system discussed in Section 2.8. The motion planning strategy I present here is described by Bullo and Lewis [2003].

The objective is to steer the snakeboard from an initial configuration at rest to a desired configuration at rest. Bullo and Lewis [2003] show that the snakeboard has two decoupling vector fields and indicate how to use these to explicitly solve the motion planning problem. In Figure 7 I show the trace of a solution of the motion planning problem where the snakeboard is steered from the position at the origin to a specified final position. The motion planning controller that generates Figure 7 is very simple, but is directly derived from the controllability results of Bullo and Lewis [2005].

Other systems where the motion planning problem can be solved are given in Chapter 13 of Bullo and Lewis [2004].

The point of what I say in this section includes the following.

1. The idea of motion planning using decoupling vector fields is an excellent example of one where the right question is suggested by the differential geometric formulation.
2. Controllability is often seen as being a somewhat esoteric undertaking. Very often systems are controllable in a simple way, and in such cases sophisticated controllability techniques are more than what is needed. However, mechanical systems provide a rich class of systems where controllability is often not easy to characterise. Moreover, by understanding the controllability properties of a system, it is possible that one might be able to derive explicit solutions to the motion planning problem.
3. While the affine connection formulation of nonholonomic systems is fairly complicated, the application of this formulation to the solution of the motion planning problem for the snakeboard justifies the complication.

3.3. Energy shaping. The last idea I discuss in this section is something that is actually extremely complicated, and is still the topic of ongoing research. The idea is this. Suppose

Figure 7. A solution of the snakeboard motion planning problem

one is given a mechanical system, let me call it the "open-loop system," with kinetic energy metric $\mathbb{G}_{\mathrm{ol}}$ and with the uncontrolled external forces being those coming from a potential function V_{ol}. Using feedback, one wishes to transform the system into a mechanical system, let me call it the "closed-loop system," with kinetic energy metric $\mathbb{G}_{\mathrm{cl}}$ and external forces being those coming from a potential function V_{cl}. The value of this idea is that much can be said about the behaviour of a mechanical system by knowing properties of the potential function. In particular, equilibria and their stability are easily determined from the properties of the potential function. Thus this technique of "energy shaping," i.e., changing the energy of the system using feedback, is a useful tool for stabilisation. The original idea of potential shaping for fully actuated systems dates to [Takegaki and Arimoto 1981]. Potential shaping in the underactuated case was worked out by van der Schaft [1986]. The possibility of shaping the kinetic energy was considered first in a Lagrangian setting by Bloch, Chang, Leonard, and Marsden [2001] and Bloch, Leonard, and Marsden [2000] and in a Hamiltonian setting by Ortega, Spong, Gómez-Estern, and Blankenstein [2002].

In our affine differential geometric formulation energy shaping amounts to finding a state feedback u_{shp}^a, $a \in \{1, \ldots, m\}$, that translates the open-loop system to the closed-loop system. In equations this reads

$$\overset{\mathbb{G}_{\mathrm{ol}}}{\nabla}_{\gamma'(t)}\gamma'(t) + \mathbb{G}_{\mathrm{ol}}^\sharp \circ \boldsymbol{d}V_{\mathrm{ol}}(\gamma(t)) - \sum_{a=1}^m u_{\mathrm{shp}}^a(v_q)\mathbb{G}_{\mathrm{ol}}^\sharp \circ F^a(q) = \overset{\mathbb{G}_{\mathrm{cl}}}{\nabla}_{\gamma'(t)}\gamma'(t) + \mathbb{G}_{\mathrm{cl}}^\sharp \circ \boldsymbol{d}V_{\mathrm{cl}}(\gamma(t)).$$

The process of energy shaping can be broken into two steps, kinetic energy shaping and po-

tential energy shaping. For kinetic energy shaping one seeks feedback u_{kin}^a, $a \in \{1, \ldots, m\}$, such that

$$\sum_{a=1}^{m} u_{\text{kin}}^a(v_q)\mathbb{G}_{\text{ol}}^{\sharp} \circ F^a(q) = \underbrace{(\overset{\mathbb{G}_{\text{cl}}}{\nabla} - \overset{\mathbb{G}_{\text{ol}}}{\nabla})}_{D_{\text{ol,cl}}}(v_q, v_q), \tag{3.3}$$

where we use the fact that the difference of affine connections ($D_{\text{ol,cl}}$ in the above equation) is a tensor field. For potential energy shaping one seeks feedback u_{pot}^a, $a \in \{1, \ldots, m\}$, such that

$$\sum_{a=1}^{m} u_{\text{pot}}^a(q)\mathbb{G}_{\text{ol}}^{\sharp} \circ F^a(q) = \mathbb{G}_{\text{ol}}^{\flat} \circ \mathbb{G}_{\text{cl}}^{\sharp} \circ \boldsymbol{d}V_{\text{cl}}(q) - \boldsymbol{d}V_{\text{ol}}(q). \tag{3.4}$$

Equation (3.3) is a quasilinear first-order partial differential equation for the closed-loop metric $\mathbb{G}_{\text{cl}}$. Equation (3.4) is a linear first-order partial differential equation for the closed-loop potential function V_{cl}.

Thus energy shaping leads to partial differential equations and, particularly when kinetic energy shaping is considered, these partial differential equations are extremely complicated. It is difficult to even give conditions ensuring that these partial differential equations have interesting solutions (there always exists the trivial solution $\mathbb{G}_{\text{cl}} = \mathbb{G}_{\text{ol}}$ and $V_{\text{cl}} = V_{\text{ol}}$). A rather coarse general result is given by Auckly and Kapitanski [2002], and they make use of some Riemannian geometry in their development. A fairly refined result concerning part of the problem is given by Lewis [2006]. In both cases, the results are too complicated to reproduce here.

From a practical point of view, the jury is still out on the value of energy shaping as a means of stabilisation. Because of the complexity of the partial differential equations involved, only in simple examples has the method actually been implemented (see the references above). Myself, I am dubious about the merits of the technique being widely applicable. However, the mathematical questions arising from energy shaping are deep and interesting, and may possibly shed some light on more practical matters. However, at this point one cannot really be certain about any of this.

But the point is this.

1. Sometimes natural engineering questions can lead to extremely challenging problems that require sophisticated mathematics for their solution.

4. In conclusion

In this paper I have attempted to illustrate that there is value in the differential geometric approach to modelling and control of mechanical systems. In particular, in Section 3 I have attempted to make it clear that this value extends beyond a mere elegant formulation of the equations of motion. What I hope is that an uninitiated reader interested in the differential geometric approach might now attack the daunting literature with something of a sense of purpose, knowing that there are possible rewards awaiting them at the end of their journey.

References

Arai, H., Tanie, K., and Shiroma, N. [1998] *Nonholonomic control of a three-DOF planar underactuated manipulator*, Institute of Electrical and Electronics Engineers. Transactions on Robotics and Automation, **14**(5), pages 681–695, ISSN: 1042-296X, DOI: 10.1109/70.720345.

Auckly, D. R. and Kapitanski, L. V. [2002] *On the λ-equations for matching control laws*, SIAM Journal on Control and Optimization, **41**(5), pages 1372–1388, ISSN: 0363-0129, DOI: 10.1137/S0363012901393304.

Bhat, S. P. and Bernstein, D. S. [2000] *A topological obstruction to continuous global stabilization of rotational motion and the unwinding phenomenon*, Systems & Control Letters, **39**(1), pages 63–70, ISSN: 0167-6911, DOI: 10.1016/S0167-6911(99)00090-0.

Bloch, A. M. [2003] *Nonholonomic Mechanics and Control*, number 24 in Interdisciplinary Applied Mathematics, Springer-Verlag: New York/Heidelberg/Berlin, ISBN: 978-0-387-95535-3.

Bloch, A. M., Chang, D. E., Leonard, N. E., and Marsden, J. E. [2001] *Controlled Lagrangians and the stabilization of mechanical systems. II. Potential shaping*, Institute of Electrical and Electronics Engineers. Transactions on Automatic Control, **46**(10), pages 1556–1571, ISSN: 0018-9286, DOI: 10.1109/9.956051.

Bloch, A. M., Leonard, N. E., and Marsden, J. E. [2000] *Controlled Lagrangians and the stabilization of mechanical systems. I. The first matching theorem*, Institute of Electrical and Electronics Engineers. Transactions on Automatic Control, **45**(12), pages 2253–2270, ISSN: 0018-9286, DOI: 10.1109/9.895562.

Bullo, F. and Lewis, A. D. [2003] *Kinematic controllability and motion planning for the snakeboard*, Institute of Electrical and Electronics Engineers. Transactions on Robotics and Automation, **19**(3), pages 494–498, ISSN: 1042-296X, DOI: 10.1109/TRA.2003.810236.

— [2004] *Geometric Control of Mechanical Systems, Modeling, Analysis, and Design for Simple Mechanical Systems*, number 49 in Texts in Applied Mathematics, Springer-Verlag: New York/Heidelberg/Berlin, ISBN: 978-0-387-22195-3.

— [2005] *Low-order controllability and kinematic reductions for affine connection control systems*, SIAM Journal on Control and Optimization, **44**(3), pages 885–908, ISSN: 0363-0129, DOI: 10.1137/S0363012903421182.

Bullo, F. and Lynch, K. M. [2001] *Kinematic controllability and decoupled trajectory planning for underactuated mechanical systems*, Institute of Electrical and Electronics Engineers. Transactions on Robotics and Automation, **17**(4), pages 402–412, ISSN: 1042-296X, DOI: 10.1109/70.954753.

Bullo, F. and Žefran, M. [2002] *On mechanical systems with nonholonomic constraints and symmetries*, Systems & Control Letters, **45**(2), pages 133–143, ISSN: 0167-6911, DOI: 10.1016/S0167-6911(01)00173-6.

Dalsmo, M. and van der Schaft, A. J. [1998] *On representations and integrability of mathematical structures in energy-conserving physical systems*, SIAM Journal on Control and Optimization, **37**(1), pages 54–91, ISSN: 0363-0129, DOI: 10.1137/S0363012996312039.

Lewis, A. D. [2000] *Simple mechanical control systems with constraints*, Institute of Electrical and Electronics Engineers. Transactions on Automatic Control, **45**(8), pages 1420–1436, ISSN: 0018-9286, DOI: 10.1109/9.871752.

Lewis, A. D. [2006] *Potential energy shaping after kinetic energy shaping*, in *Proceedings of the 45th IEEE Conference on Decision and Control*, IEEE Conference on Decision and Control, (San Diego, CA, Dec. 2006), Institute of Electrical and Electronics Engineers, pages 3339–3344, DOI: 10.1109/CDC.2006.376885.

Lewis, A. D. and Murray, R. M. [1997] *Configuration controllability of simple mechanical control systems*, SIAM Journal on Control and Optimization, **35**(3), pages 766–790, ISSN: 0363-0129, DOI: 10.1137/S0363012995287155.

Lewis, A. D., Ostrowski, J. P., Murray, R. M., and Burdick, J. W. [1994] *Nonholonomic mechanics and locomotion: The Snakeboard example*, in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, IEEE International Conference on Robotics and Automation, (San Diego, CA, May 1994), Institute of Electrical and Electronics Engineers, pages 2391–2400, DOI: 10.1109/ROBOT.1994.351153.

Lewis, A. D. and Tyner, D. R. [2004] *Controllability of a hovercraft model (and two general results)*, in *Proceedings of the 43rd IEEE Conference on Decision and Control*, IEEE Conference on Decision and Control, (Paradise Island, Bahamas, Dec. 2004), Institute of Electrical and Electronics Engineers, pages 1204–1209, DOI: 10.1109/CDC.2004.1430205.

Lynch, K. M., Shiroma, N., Arai, H., and Tanie, K. [2000] *Collision-free trajectory planning for a 3-DOF robot with a passive joint*, The International Journal of Robotics Research, **19**(12), pages 1171–1184, ISSN: 0278-3649, DOI: 10.1177/02783640022068011.

Murray, R. M., Li, Z. X., and Sastry, S. S. [1994] *A Mathematical Introduction to Robotic Manipulation*, CRC Press: Boca Raton, FL, ISBN: 978-0-8493-7981-9.

Ortega, R., Spong, M. W., Gómez-Estern, F., and Blankenstein, G. [2002] *Stabilization of a class of underactuated mechanical systems via interconnection and damping assignment*, Institute of Electrical and Electronics Engineers. Transactions on Automatic Control, **47**(8), pages 1218–1233, ISSN: 0018-9286, DOI: 10.1109/TAC.2002.800770.

Stramigioli, S. [2001] *Modeling and IPC Control of Interactive Mechanical Systems—A Coordinate-Free Approach*, number 266 in Lecture Notes in Control and Information Sciences, Springer-Verlag: New York/Heidelberg/Berlin, ISBN: 978-1-85233-395-9.

Sussmann, H. J. [1987] *A general theorem on local controllability*, SIAM Journal on Control and Optimization, **25**(1), pages 158–194, ISSN: 0363-0129, DOI: 10.1137/0325011.

Sussmann, H. J. and Jurdjevic, V. [1972] *Controllability of nonlinear systems*, Journal of Differential Equations, **12**(1), pages 95–116, ISSN: 0022-0396, DOI: 10.1016/0022-0396(72)90007-1.

Takegaki, M. and Arimoto, S. [1981] *A new feedback method for dynamic control of manipulators*, Transactions of the ASME. Series G. Journal of Dynamic Systems, Measurement, and Control, **103**(2), pages 119–125, ISSN: 0022-0434, DOI: 10.1115/1.3139651.

van der Schaft, A. J. [1986] *Stabilization of Hamiltonian systems*, Nonlinear Analysis. Theory, Methods, and Applications, **10**(10), pages 1021–1035, ISSN: 0362-546X, DOI: 10.1016/0362-546X(86)90086-6.

— [2000] *Port-controlled Hamiltonian systems: Towards a theory for control and design of nonlinear physical systems*, Journal of the Society of Instrument & Control Engineers of Japan, **39**(2), pages 91–98, URL: http://www.lar.deis.unibo.it/news/files/pchs.pdf (visited on 07/14/2014).