FLUID MECHANICS AND MATHEMATICAL STRUCTURES

PHILIP BOYLAND
Department of Mathematics
University of Florida
Gainesville, FL 32611-8105, USA
E-mail: *boyland@math.ufl.edu*

To Ed Spiegel on his 70th birthday.

Abstract. This paper provides an informal survey of the various mathematical structures that appear in the most basic models of fluid motion.

1. Introduction

Fluid mechanics is the source of many of the ideas and concepts that are central to modern Mathematics. The vestiges of these origins remain in the names of mathematical structures: flows, currents, circulation, frames, ..., and in the names of the distinguished mathematician-scientists which pervade the terminology of both fields: Euler, Bernoulli, Stokes, Cauchy, Lagrange, ... Mathematicians have abstracted and vastly generalized basic fluid mechanical concepts and have created a deep and powerful body of knowledge that is unfortunately now mostly inaccessible to fluid mechanicians, while mathematicians themselves have lost all but a passing knowledge of the physical origins of many of their basic notions.

This paper provides an informal survey of the various mathematical structures that appear in the most basic models of fluid motion. As befits this volume it is strictly pedagogical. There are no original results, and most of what we say here has been known for at least a century in one guise or another. In style it attempts to reproduce the informal tutorials that took place during the Newton Institute program in Fall, 2000. Mathematical objects are described informally with an emphasis on their meaning in Fluid Mechanics. Many formulas are stated without proof. The paper is meant to be a friendly introduction to basic concepts and ideas and will hopefully provide an intuitive foundation for a more careful study of the literature. Readers are strongly encouraged to consult the references given in each section.

In describing the fluid model we give primacy to the actual evolution of the fluid, to the fluid flow, as opposed to the velocity fields. This agrees with our direct experience of seeing fluids move as well as being very natural in the mathematical progression we describe. While conceptually valuable, this point of view is computationally very impractical. The basic equations of Fluid Mechanics are framed in terms of the velocity fields. Solutions of the equations are difficult to obtain, but formulas for the fluid motions are virtually nonexistent.

One of the goals of the paper is to introduce students of fluid mechanics to the "world view" of modern Mathematics. A common view among mathematicians is that Mathematics is the study of sets with structures and of the transformations between them. From this point of view, the basic model of Fluid Mechanics is the set consisting of the fluid body and the various additional structures that allow one to discuss such properties as continuity, volume, velocity and deformation. The flow itself is a transformation of the fluid body to itself and a basic object of study is the transport of structures under the flow.

It is considered a mathematical virtue to use just those structures which are needed in a given situation and no more. The goal is not to be mindlessly abstract but rather to discover and illustrate what is essential and fundamental to the task at hand. This has the added advantage that conclusions have the widest applicability. Thus we develop the fluid model step by step, introducing structures to fit certain needs in the modelling process. When possible coordinate free terminology and notation is used. We invoke the usual argument in its favor: anything fundamental shouldn't depend on the choice of coordinates and being forced to treat computational objects as global entities provides a new and sometimes valuable perspective on familiar operations. We do, however, try to connect new objects with familiar formula, and we try to be clear when formulas are only valid in usual Euclidean space.

We have attempted to keep the paper accessible to a wide audience, but some fluid mechanical terminology is used without explanation. General readers may consult any standard text; the mathematically inclined may prefer [22], [7], [16] or [5].

2. Basic Kinematics and Mathematical Structures

We begin by describing the most fundamental mathematical structures used in modelling the simplest type of fluid behavior. Unless otherwise stated the fluid is three-dimensional; two-dimensional flow can be treated similarly.

2.1. THE FLUID REGION, FLUID MAPS AND CARDINALITY

The first step is to model the fluid region itself. In Fluid Mechanics it is usually said that the fluid is a continuum. In mathematical language this says that, at least locally, the fluid looks like usual two or three-dimensional space. We allow the possibility that on large scales the fluid is not flat, but rather can curve back on itself like the surface of the (almost) spherical earth. Objects with a local structure like the plane or three space, but perhaps nontrivial global behavior are called 2- or 3-dimensional manifolds. A fluid particle is a mathematical point in the manifold. The generic fluid region or fluid body is denoted B.

The modelling of the fluid as a manifold has many implications and some obvious problems. It is clearly false on very small scales as it completely ignores the molecular nature of fluids. In addition, a mathematical point obviously has no physical meaning. Nonetheless, the theories and technologies based on continuum models have been wildly successful and we can proceed with confidence using a manifold model.

As a fluid flows particles are transported. In terms of the model the fluid motion is defined as the collective motion of the particles. It is assumed that the individual particles do not split into pieces, and each has a well defined future and remains distinguished for all times. This means that the evolution is described by the transformation that takes the initial position of a particle as input and gives the position after a time T as the output. Thus for each time T the evolution of the fluid after time T is described by a function (or map), f, from the initial region occupied by the fluid, B_0, to the region occupied after time T, B_T. The usual notation for this is $f : B_0 \to B_T$, indicating that f is a function with domain B_0 and range B_T. The function f is called the *time-T fluid map*. The typical point in a fluid region is labelled $\mathbf{p}$, $\mathbf{q}$, etc. Note that these label the geometric position in the region and not the fluid particles, and so they do not move with the flow.

During their evolution it is further assumed that fluid particles do not overlap, coalesce or collide. Thus if two points are distinct, then their positions after time T are different, or $\mathbf{p} \neq \mathbf{q}$ implies $f(\mathbf{p}) \neq f(\mathbf{q})$, and so f is a one-to-one (or injective) function. Since by definition, B_T is the fluid region after time T, we have that $f : B_0 \to B_T$ is an onto (or surjective) function, which is just to say that for any $\widehat{\mathbf{p}}$ in B_T there is a $\mathbf{p}$ in B_0 with $f(\mathbf{p}) = \widehat{\mathbf{p}}$, i.e. the particle at $\widehat{\mathbf{p}}$ came from somewhere.

Summarizing, the most basic assumptions on a fluid map is that it is one to one and onto, i.e. it is a bijection, or what is sometimes termed a one-to-one correspondence. This means that f preserves the "number" of fluid points. This idea is formalized as the most fundamental property of a set, its *cardinality*. By definition, two sets have the same cardinality exactly when there is a one-to-one correspondence between them. While this may not seem like much information about the fluid map it is worth noting that there are sets with different cardinality.

Exercise 2.1 Show that there is a one-to-one correspondence between the set of points in the interval, $0 \leq x \leq 1$, and the set of points in the unit square in the plane, $0 \leq x \leq 1,\ 0 \leq y \leq 1$, but that these sets have a different cardinality from the set of integers $0, \pm 1, \pm 2, \ \ldots$.

A bijection f always has an inverse map, usually denoted f^{-1}, which by definition satisfies $f \circ f^{-1} = f^{-1} \circ f = \mathrm{id}$. While the inverse of a fluid map always exists, there is no assumption in general that it is a physically reasonable fluid map, though in certain cases it can be.

Elementary set theory is covered in the first sections of most textbooks in Algebra or Topology, [13] is a standard text, and [26] is recreational but informative.

2.2. STRUCTURES, TRANSPORT, AND CATEGORIES

At this point our fluid is described by the most basic of mathematical objects, sets and a transformation. We pause now to introduce some general notions about mathematical structures and how they are transported by transformations.

Mathematical structures take many forms. The most basic examples are an *order structure* which declares which elements are larger than others, a *topology* which is a distinguished collection of subsets, and an *algebraic* structure which is a binary operation on the set which satisfies certain rules. Other common structures are made out of various kinds of functions into or out of the set. For example, a velocity field is a map out of the set, it assigns a vector to each point in the set. A loop is a map from the circle into a set. The collection of all tangent vectors to a space are turned into the tangent bundle, which in turn is the basis for other structures like tensors, Riemannian metrics and differential forms. The general intuition is that a structure is something that "lives on top of " the set, representing additional information and properties.

If a transformation between underlying sets is sufficiently well behaved it induces a map on the various structures, transforming a structure on one set to the same kind of structure on the other set. For example, in ideal magnetohydrodynamics the fluid flow transports the magnetic vector field at time zero to the magnetic vector field after a time T. The general mathematical rubric is that transport in the direction of the map is called *push forward* while *pull back* is an induced map that goes in the direction opposite to the given map. If the map we are studying is a fluid map, then the push forward transports the structure in forward time while pull back refers to transport from the future backward, usually back to the initial state of the system. The standard notation puts a star on the given map to indicate the induced map on structures, with a subscript indicating a push forward and a superscript indicating a pull back. Thus if f is a fluid map and α is a structure (say a magnetic field) then $f_*(\alpha)$ is this vector field pushed forward by the fluid evolution and $f^*(\alpha)$ indicates its pull back.

Since a fluid map f is invertible, pulling back is the same as pushing forward by the inverse f^{-1}, and we will freely pass between the two. Depending on the structures it is usually most natural to either pull back or else push forward. As a general guideline when a structure is defined by a map into the set it is pushed forward. For example, a loop in the fluid region is a map $\Gamma : S^1 \to B_0$, and this naturally pushes forward to the same kind of structure on B_T since the composition $f \circ \Gamma$ is a map $S^1 \to B_T$ called $f_*\Gamma$. On the other hand, a structure defined by maps out of the set, say a scalar field α on B_T, is most naturally pulled back because $\alpha \circ f = f^*\alpha$ is a map out of B_0. Another basic pull back is that of a set. If N is a subset of the fluid region B_T, then its pull back is a subset of B_0 defined by $f^*(N) = \{\mathbf{p} \text{ in } B_0 : f(\mathbf{p}) \text{ is in } B_T\}$; this set is also often denoted $f^{-1}(N)$.

The notion of a mathematical category and its morphisms formalizes many of the notions of sets, structures and transformations. The most basic category is the collection of sets and other categories consist of sets with some specific additional structure such as topological spaces, vector spaces, smooth manifolds, or flows on manifolds. The *morphisms* associated with a category are the maps between sets that preserve the structure, i.e. the induced map on structures is defined and a transported structure is compatible with the pre-existing one. For example, linear transformation are the morphisms of vector spaces and continuous functions are the morphisms of topological spaces.

An *isomorphism* in a category is a morphism that is a bijection and its inverse is also a morphism. Most of the commonly studied categories have their own traditional name for their isomorphism such as bijection in the category of sets, homeomorphism in the category of topological spaces, and isometry in the case of geometries. Within a given category isomorphic objects are considered indistinguishable, and indeed a rather formal definition of a mathematical area of study is that it studies those properties that are preserved by the morphisms in its particular category.

The text [15] is a standard reference.

Exercise 2.1 clearly indicates that a bijection can scramble a fluid domain very badly, prompting the requirement that fluid motions preserve the continuum nature of the fluid, not tearing or distorting it too wildly. This requirement is the essence of Topology and its morphisms, the continuous functions, and is expressed in the assumption that fluid particles which are sufficiently close initially should still be close after evolving for time T. The technicalities arise from defining "sufficiently close" and "still be close".

A topology as a structure on a set is simply a collection of distinguished subsets called the *open sets* (there are conditions which are not important here). Open sets are the allowable neighborhoods of the point and *neighborhood* is used synonymously with open set. A property is true for points sufficiently close to $\mathbf{p}$ if it is true for all open sets containing $\mathbf{p}$. In the category of sets with topologies (topological spaces) the morphisms or structure preserving maps are the continuous functions. A function $g : X \to Y$ between two topological spaces is called *continuous* if the pull back of an open set is open. Thus for each open set U in Y it is required that $f^*(U)$ is an open set in X, or by acting on all open sets, it is required that the pull back of the topology on Y is compatible with the topology on X.

Sequences provide a simple way to connect this abstraction with the more physical condition that nearby fluid particles should stay near each other. In terms of a given topology, the sequence $(\mathbf{p}_n)$ converges to the point $\mathbf{p}_0$, written $\mathbf{p}_n \to \mathbf{p}_0$, when any neighborhood V that contains $\mathbf{p}$ also contains the tail of the sequence, i.e. there is an N so that $n > N$ implies that $\mathbf{p}_n \in V$. If f is the time T fluid map, the flow preserving the continuum structure clearly requires that $\mathbf{p}_n \to \mathbf{p}_0$ implies $f(\mathbf{p}_n) \to f(\mathbf{p}_0)$. This says that any neighborhood U of $f(\mathbf{p}_0)$ must contain the tail of the sequence $f(\mathbf{p}_n)$. But transporting this data back to the fluid at time zero we see that the convergence after time T requires that $\mathbf{p}_n \in f^*(U)$ for $n > N$. But since $\mathbf{p}_n \to \mathbf{p}_0$, this would be achieved by requiring that $f^*(U)$ is itself a neighborhood, which is the case if f is continuous in the sense defined above.

Exercise 2.2 In the standard topology on $\mathbb{R}^n$ a set U is open when for every $\mathbf{p} \in U$ there is an r, so that the ball $B_r(\mathbf{p}) = \{\mathbf{q} : |\mathbf{q} - \mathbf{p}| < r\}$ is wholly contained in the set U. Show that in this case the definition of continuity of a function using the pull back of open sets is the same as the usual one: f is continuous at $\mathbf{p}$ if for all ϵ there exists a δ so that

$$|\mathbf{q} - \mathbf{p}| < \delta \text{ implies } |f(\mathbf{q}) - f(\mathbf{p})| < \epsilon \ .$$

In the fluid model particles which are sufficiently close after time T should come from particles which were close at time zero. Thus we assume that the inverse of a fluid map f^{-1} is also continuous. Such a continuous bijection with a continuous inverse is called a *homeomorphism*. Homeomorphisms are the isomorphisms in the category of topological spaces. Under a homeomorphism the pull back and the push forward of each open set is an open set, and so the pull back of the topology on Y is exactly the topology on X, and vice versa. Thus the homeomorphism has altered only labels and has altered nothing related to the topology. Thus it makes sense to define Topology as the study of those properties that are invariant under homeomorphisms. These properties include the connectedness, the number of holes, and the dimension of the space. So in the model these topological properties of the fluid region do not change as the fluid evolves.

In requiring fluid maps to be homeomorphisms we are excluding some common fluid behaviors. When we sip our coffee, the fluid tears, and we see the droplets of rain coalescing into the continuum of the creek. It is also worth noting that what is commonly called "The Continuity Equation" in Fluid Mechanics is not a statement about the kind of continuity discussed in this section, but rather a statement about the preservation of the mass of the continuum during its evolution (see sections 2.5 and 4.6).

The books [20] and [8] are standard Topology texts.

2.4. FLUID MOTIONS AND MATHEMATICAL FLOWS

At this point a fluid map describes the evolution of the fluid up to fixed, but arbitrary time. The complete evolution is described by a family of maps, ϕ_t, one for each time t. The particle at position $\mathbf{p}$ at time 0 is at the position $\phi_t(\mathbf{p})$ after time t, and so what was denoted f above is now denoted ϕ_T. If we fix $\mathbf{p}$ and vary t, the positions $\phi_t(\mathbf{p})$ sweep out the *trajectory, flow line or path* of the particle. For simplicity of exposition we assume that the fluid motion can be defined for all forward and backwards time and so ϕ_t is defined for all $t \in \mathbb{R}$. In terms of fluid maps, *slip boundary conditions* say that the fluid region is preserved by the fluid motion, $B_t = B$ for all t. The *no slip* conditions require that on the boundary of B, $\phi_t = \mathrm{id}$ for all t. Henceforth we always assume that our fluid maps satisfy at least slip conditions.

The family of fluid maps ϕ_t is best described as a single map on a bigger space including both the fluid domain and time, $\phi : B \times \mathbb{R} \to B$, with $\phi(\mathbf{p}, t) = \phi_t(\mathbf{p})$. In the previous subsection we required that ϕ_t is a homeomorphism for each t. It is clear that we must also require continuity as the time is varied so that fluid particles don't suddenly jump during their evolution. This is accomplished by requiring that the full function ϕ is continuous, i.e. continuous in both variables.

Remark 2.3 It is traditional in Fluid Mechanics to distinguish the Eulerian and the Lagrangian perspective. In the first, the labels are pinned down in the fluid region, in mathematical language they are local coordinates in the manifold modelling the fluid. In the Lagrangian perspective, the fluid particle is labelled and that label remains on the particle as it moves. In addition, the Eulerian perspective is a field theory, focusing on the velocity vector field while the Lagrangian focuses on the motions or flow lines of particles. The point of view here is mixed: we always label using a fixed, pinned down coordinate system in the fluid region and never use Lagrangian or advected coordinates, but the primary focus is on the motion of particles. It is also traditional in Fluid Mechanics to use $\mathbf{x}$ for the Eulerian coordinate, and to denote a trajectory as $\mathbf{x}(t)$, indicating how the coordinate varies as a function of time. This use of the same symbol to indicate both a coordinate and a function of time makes many mathematicians uncomfortable; it represents an unacceptable confusion of different classes of mathematical objects. Thus we denote the time evolution by a distinct function, ϕ, which encapsulates all the time and space evolution. To avoid confusion with the training of fluid mechanicians, while maintaining the notation $\mathbf{x} = (x_1, x_2, x_3)$ for the local coordinates, $\mathbf{p}$ and $\mathbf{q}$ are used here to denote the typical point in the fluid region (labelled by the fixed coordinate system) rather than the more common mathematical x or $\mathbf{x}$.

There is a useful distinction among fluid motions based on how the future of particles depends on the starting time. Pick a fixed, but arbitrary location in the fluid region and monitor the future of the particle that is there at time 0 and also that of the particle that is there at some later time. If these futures are always the same the flow is *steady*, otherwise it is *unsteady*.

This distinction has a nice formulation in terms of the algebra of the time parameter, i.e. of the real line. Begin the evolution of the fluid particle at $\mathbf{p}$ and flow for time t, now at this point begin again and flow for time s, so we are at the point $\phi_s(\phi_t(\mathbf{p}))$. On the other hand, we could start at $\mathbf{p}$, flow for time t and then continue without restarting time for another time interval s, ending up at the point $\phi_{s+t}(\mathbf{p})$. The flow is steady exactly when these points are the same for all s, t and points $\mathbf{p}$, or in terms of the fluid maps, $\phi_s \circ \phi_t = \phi_{s+t}$. This group law coupled with the fact that ϕ_0 is the identity map says that a steady flow is an *action* of the real numbers on the region B. This implies, in particular, that the inverse maps be computed by reversing time, $\phi_t^{-1} = \phi_{-t}$

Within most of mathematics the word *flow* refers to this kind of action of the reals on a space. In particular, mathematical flows always correspond to steady fluid flows. For this reason the general, perhaps unsteady, evolution of a fluid is called here a *fluid motion* reserving the word flow for steady flows.

Two simple examples of one-dimensional flows are the linear flow, $\phi_t(p) = p + rt$, and the exponential flow, $\phi_t(p) = pe^{rt}$, where r is a real number. In these cases the group law of the action corresponds to the standard distributive and exponential laws, respectively. Flows correspond to solutions of time-independent differential equations, and it is remarkable that any such solution yields a flow with nice algebraic properties.

The simplest unsteady fluid motions are periodic. A fluid motion (or fluid map ϕ_t) is called *periodic* if starting at an initial time t at the point $\mathbf{p}$ yields the same future as starting at the same point at a time P later, where P is called the *period.* This implies that $\phi_{t+P} = \phi_t \circ \phi_P$, and letting $t = (n-1)P$ for an integer n, that $\phi_{nP} = \phi_P \circ \phi_P \circ \ldots \circ \phi_P$ (n times). Thus if we define the *Poincáre* map as $g = \phi_P$, its iterates (repeated self compositions) satisfy $g^n = \phi_{nP}$, and so they describe the evolution of the fluid. The Poincáre map is often called the stroboscope map, as it corresponds to viewing the fluid under a light that turns on once each period.

If we have a fluid motion ϕ_t on a region B and a homeomorphism h from B to another region $\hat{B}$, we can push forward the flow to one on $\hat{B}$ by the rule $h_*(\phi_t) = h \circ \phi_t \circ h^{-1}$. The motion for time T on $\hat{B}$ is obtained by first coming back to B via h^{-1}, then flowing from that point via ϕ_T, and then going back to B via h. If the homeomorphism is from B to itself, it represents a change of coordinates, and if $h_*\phi_t = \psi_t$ one says that two flows ϕ_t and ψ_t are the same fluid motion up to change of coordinates by h.

2.5. INCOMPRESSIBILITY, MASS CONSERVATION AND MEASURES

If there are no mechanisms for inflow or outflow as a fluid subregion evolves its mass stays the same. Similarly, if there are no mechanisms to compress or expand regions of the fluid, then the volume of fluid subregions is preserved under the evolution. Both of these properties can be included in the model by postulating measures that are invariant under the fluid motion.

A measure is a rule which assigns a non-negative number to various subsets of the main set (again, there are additional technicalities that are not crucial here). In a fluid model, the measure of a subset could represent the mass of the fluid in the subset, or perhaps the volume of the subset. Thus if $m(U)$ is the mass of the fluid contained in the region U, then mass conservation of the flow ϕ_t is expressed by $m(\phi_t(U)) = m(U)$ for all measurable sets U and all times t. Similarly, if $\nu(U)$ is the usual three-dimensional volume (Lebesgue measure) of a set U in $\mathbb{R}^3$, then a fluid flow is *incompressible* if $\nu(\phi_t(U)) = \nu(U)$ for all measurable sets U and all times t.

Measures are most commonly used for integration, and one writes $\int_U \alpha \, d\nu$ for the usual integral of the real valued function α over the region U, and thus $\nu(U) = \int_U \, d\nu$. A *density* ρ is a real-valued function which represents the "measure per unit volume", or more precisely, ρ is a density for the measure m, exactly when $m(U) = \int_U \rho \, d\nu$. One then writes $dm = \rho \, d\nu$.

If $g : X \to Y$ is a function, we can use it to transport measures. Given a measure m on X, its *push forward* to Y is $g_* m$. It is defined by the rule that the measure $g_* m$ assigns to a subset V of Y is the measure of $g^*(V)$, or in symbols, $g_* m(V) = m(g^*(V))$. Thus a fluid motion ϕ_t preserves the mass given by the measure m exactly when $\phi_{t*} m = m$, for all t (in the unsteady case m may be time dependent). Similarly, incompressibility is expressed by $\phi_{t*}\nu = \nu$, for all t.

¿From a strictly kinematic point of view there is little difference between preserving a mass m or preserving the volume ν. This is consequence of a theorem of Oxtoby and Ulam [21] (Moser [19] in the smooth category). If the mass density ρ is a reasonable function and for simplicity we assume that the total mass equals the total volume, $\int \rho \, d\nu = \int \, d\nu$, then the theorems give a homeomorphism (diffeomorphism) h of B to itself which pushes forward m to ν, $h_* m = \nu$. Thus if ϕ_t preserves m then by changing coordinates by h, the fluid motion $h_*(\phi_t)$ preserves ν.

Elementary measure theory and integration is covered in most textbooks in Analysis, [12] is a measure theory standard.

2.6. DIFFERENTIABILITY

The laws of mechanics are expressed using derivatives, and indeed, the development of mechanics went hand in hand with the invention of calculus. The continuity of the fluid motion arose from the need to preserve the continuum structure of the fluid body. But the assumption that the fluid region is locally Euclidean gives more than just a local topology, there is also the algebraic structure of Euclidean space as a vector space. The assumption of differentiability of fluid motion means that on small enough scales the fluid map looks like a linear map of a vector space to itself. This is made precise by Taylor's theorem and its converse which says that a map f is differentiable exactly when it is linear up to an error of specified type. More precisely, in Euclidean space the matrix A is the first derivative of the map f at the point $\mathbf{p}$ exactly when

$$f(\mathbf{p} + \mathbf{h}) = f(\mathbf{p}) + A\mathbf{h} + o(|\mathbf{h}|) \ , \tag{1}$$

with $|\mathbf{h}|$ the usual Euclidean norm and $o(|\mathbf{h}|)$ representing a quantity $r(\mathbf{h})$ that goes to zero sufficiently fast so that $\frac{r(\mathbf{h})}{|\mathbf{h}|} \to 0$ as $\mathbf{h} \to 0$.

The matrix A is, of course, the usual Jacobian matrix, for which there are numerous notations; its (i, j)-th entry is $\partial f_i / \partial x_j$. When f is real valued, i.e. has only one component, then the derivative matrix is a row vector representing the gradient, ∇f. It is common in Fluid Mechanics to also use ∇ in the multidimensional case, and so we shall denote the full derivative as ∇f, or $\nabla f(\mathbf{p})$. A function f is *continuously differentiable* if a matrix A satisfying (1) exists at each point $\mathbf{p}$ and the matrices vary continuously with the evaluation point. In this case one writes $f \in C^1$.

Higher derivatives can be defined using the analog of (1). A function is C^2 when it locally looks like a quadratic polynomial up to an error of $o(|\mathbf{h}|^2)$. The notation for the analog of (1) grows complicated, but in the simplest case of a real valued function α, C^2 means that

$$\alpha(\mathbf{p}+\mathbf{h}) = \alpha(\mathbf{p}) + \nabla\alpha \cdot \mathbf{h} + \mathbf{h}^T \cdot H(\alpha) \cdot \mathbf{h} + o(|\mathbf{h}|^2) \ ,$$

where $H(\alpha)$ is the Hessian matrix of α, $(\frac{\partial^2\alpha}{\partial x_i \partial x_j})$, and the Hessian must vary continuously with the evaluation point. Continuing, a function is C^k if it locally looks like a degree k polynomial up to order $o(|\mathbf{h}|^k)$. If a function is C^k for all k, it is called C^∞. A condition that is even stronger than C^∞ is the requirement that the Taylor's series converges in a neighborhood of each point. Such functions are called *real analytic* and their class is denoted C^ω.

Exercise 2.4 Let the real valued function α be defined by $\alpha(0) = 0$ and for $x \neq 0$, $\alpha(x) = \exp(-1/x^2)$. Show that all orders of derivatives at 0 are equal to 0 and thus the Taylor's series about 0 is identically zero. However, the function α is not identically zero in any neighborhood of the origin, and so α is an example of a C^∞ function that is not C^ω.

On a manifold, a function is C^k if it is C^k in local coordinates in the domain and the range, and we need to treat the derivative as a linear transformation rather than a matrix. The matrix represents the linear transformation in particular coordinates, but it is the linear transformation that has coordinate free meaning.

In formulating differentiability from the structural point of view, the objects in the category are C^∞-manifolds. These are manifolds in which C^∞-compatible local coordinates can be defined. The morphisms are C^∞-maps, and the isomorphisms are called *diffeomorphisms*, A diffeomorphism is a C^∞-bijection whose inverse is also C^∞. This category is the subject of Differential Topology. Henceforth we assume that fluid maps ϕ_t are C^∞-diffeomorphisms for each t, and further, that the family itself is C^∞ in t. It is worth noting that C^3 is enough differentiability to do virtually all of Fluid Mechanics, but we shall not deal with that level of technicalities here.

Introductions to Differential Topology are given in [17] and [24], and [11], [6], and [14] are standard texts.

3. Vector Fields and the Tangent Bundle

At this point we have developed all the basic kinematic assumptions for the simplest models of fluid flow. Forces enter into the models by their action on velocity fields, the prototypical example of a vector field. Vector fields as a mathematical structure are the basis of many additional structures and operations, and in this section we discuss some of those that are valuable in describing fluid motions. Principal among these are the tangent bundle and other structural bundles that live "above" the fluid body. The fluid motion and its derivatives induce an action on these bundles that transport the various structures. The nature of this transport is the key to describing geometric and topological aspects of the fluid motion.

3.1. VECTOR FIELDS AND THE VELOCITY FIELD

Given a fluid motion ϕ_t, its *velocity field* $\mathbf{u}$ at the point $\mathbf{p}$ at a time t is the instantaneous velocity of the fluid particle that occupies that point at that time. Note that this is the particle that started at $\phi_{-t}(\mathbf{p})$ at time zero, and so

$$\mathbf{u}(\mathbf{p}, t) = \frac{\partial \phi}{\partial t}(\phi_{-t}(\mathbf{p}), t) \ ,$$

or

$$\mathbf{u}(\phi_t(\mathbf{p}), t) = \frac{\partial \phi}{\partial t}(\mathbf{p}, t) \ . \tag{2}$$

This is sometimes called the *advection equation*. Given a velocity field $\mathbf{u}$, its fluid motion ϕ_t is obtained by solving the differential equation given by (2). It is easy to check that a fluid motion is steady as defined in Section 2.4 if and only if its velocity field is time independent, $\mathbf{u}(\mathbf{p}, t) = \mathbf{u}(\mathbf{p})$, and the motion is T-periodic exactly when the velocity field is T-periodic, $\mathbf{u}(\mathbf{p}, t + T) = \mathbf{u}(\mathbf{p}, t)$.

Exercise 3.1 The planar velocity field $\mathbf{u}(\mathbf{p}) = (-p_2, p_1)$ represents uniform counter-clockwise fluid motion. Show that its flow is $\phi_t(p_1, p_2) = (p_1 \cos(t) - p_2 \sin(t), p_1 \sin(t) + p_2 \cos(t))$, and that in this case the group property $\phi_t \phi_s = \phi_{s+t}$ is a reflection of standard trigonometric identities.

The velocity field at a fixed time is an example of a *vector field*, i.e. a rule that assigns a vector to each point in the fluid region. Other common examples of vector fields are the magnetic field and the vorticity field. We will adopt the fairly standard mathematical convention that the terminology "vector field" refers to a *time independent* field. Through equation (2) mathematical flows and vector fields go hand in hand and it is usual to pass freely from one to the other without comment.

In certain circumstances the time parameterization of a flow is not important, and one wants to think of the trajectories as just curved lines in space. In this case it is common in the sciences to call the curve constituting the trajectory a *field line*, as in a magnetic field line and a vortex (field) line. The mathematical name for a way to fill a manifold with a collection of curves is a *one-dimensional foliation*, and the individual curves are called the *leaves*. Note that the nonzero scalar multiple of a given nonvanishing vector field, $\alpha\,\mathbf{u}$, gives the same foliation as the original field, but gives rise to a different flow.

3.2. THE TANGENT BUNDLE

There is a common and quite understandable confusion that occurs early in a student's mathematical training; is the base of a vector at the origin or at the point along a curve to which it is tangent? This usually leads to the mathematical vector space called $\mathbb{R}^3$ being somehow conceptually distinct in the student's mind from the physical notion of velocity vectors.

The resolution of this unsatisfactory state is to attach a copy of the vector space $\mathbb{R}^3$ to each point of three-dimensional space and then one can accommodate all possible vectors being based at all possible points. The resulting object is the *tangent bundle* of three space, and is denoted $T\mathbb{R}^3$. It is equal to $\mathbb{R}^3 \times \mathbb{R}^3$, with the first factor holding base points and the second the tangent vectors. The collection of vectors attached to a single basepoint $\mathbf{p}$ is called the *tangent space* or *fiber* at that point and is written $T_{\mathbf{p}}\mathbb{R}^3$.

A vector field assigns a vector to each basepoint in $\mathbb{R}^3$. Thus properly speaking a vector field is a map $\mathbf{U} : \mathbb{R}^3 \to T\mathbb{R}^3$. Amongst all such maps vector fields have the distinguishing property that they map a basepoint to a vector based at that point. This is symbolically expressed by defining the projection $\pi : T\mathbb{R}^3 \to \mathbb{R}^3$ which takes a vector to its basepoint $\pi(\mathbf{p}, \vec{v}) = \mathbf{p}$. A vector field then satisfies $\pi \circ \mathbf{U} = \mathrm{id}$. Since it picks out just one vector "above" each point, a vector field is called a *cross section* of the tangent bundle.

The velocity field studied in Exercise 3.1 is formally written as $\mathbf{U}(p_1, p_2) = (p_1, p_2, -p_2, p_1)$. It is clear why in practise the first two components are usually suppressed, but it is equally clear that they need to be there in the proper understanding of a vector field. A compromise is to use the base point as a subscript, and so, for example, $\mathbf{v}_{\mathbf{p}}$ represents a vector based at $\mathbf{p}$.

For a fluid region B contained in $\mathbb{R}^3$, its tangent bundle is denoted TB and is also a Cartesian product $TB = B \times \mathbb{R}^3$. For a 3-dimensional manifold the collection of tangent vectors based at a single point is always the vector space $\mathbb{R}^3$, and locally any tangent bundle looks like a product, but in general the entire bundle is not a product. The simplest example is in one dimension lower and is provided by the two sphere, S^2. Now if in fact, $TS^2 = S^2 \times \mathbb{R}^2$, then for each point $\mathbf{p}$ on the sphere, we pick the vector $(\mathbf{p}, (1, 0))$, yielding a nonvanishing vector field on the sphere. This contradicts the "hairy ball" theorem (see, for example, [11]).

3.3. TRANSPORT OF VECTOR FIELDS

The next step is to understand how vectors and vector fields are moved by the fluid, or more generally transported by a diffeomorphism. Recall that Taylor's theorem

$$f(\mathbf{p} + \mathbf{h}) = f(\mathbf{p}) + \nabla f(\mathbf{p})\mathbf{h} + o(|\mathbf{h}|)$$

says that up to an error, on small scales the function f acts like the linear map $\nabla f(\mathbf{p})$. In the sciences it is common to use $\delta\mathbf{p}$ in place of $\mathbf{h}$, representing a differential or infinitesimal displacement in the actual space. However, in Differential Topology the linear action of the derivative is lifted from the manifold up to the attached tangent vectors. Thus the function f induces a map on the tangent bundle with the map itself transporting basepoints and the derivative transporting tangent vectors. This induced map is called the *tangent map* and is written as $Tf(\mathbf{p}, \mathbf{v_p}) = (f(\mathbf{p}), \nabla f(\mathbf{p})\mathbf{v_p})$.

The transport of vectors can also be viewed through the chain rule. Let us focus on a vector $\mathbf{v_p}$ and say it is a vector in a magnetic field and that $\gamma(s)$ is a parameterization of the field line through $\mathbf{p}$, with $\gamma(0) = \mathbf{p}$ and $\gamma'(0) = \mathbf{v_p}$. Now we transform the entire field by the diffeomorphism f. The field line through $\mathbf{p}$ is transformed to a curve parameterized as $f(\gamma(s))$ and so the vector $\mathbf{v_p}$ is sent to the corresponding tangent vector of the image,

$$\left.\frac{df(\gamma(t))}{dt}\right|_{t=0} = \nabla f(p)\gamma'(0) = \nabla f(p)\mathbf{v_p} \ .$$

Thus we see that the appropriate transport of the vector $\mathbf{v_p}$ is its image under the linear transformation ∇f. Note that under the tangent map a tangent vector is not rigidly transported, but is acted on by the linear map that best reflects the stretching, shearing, and rotation that occurs near the point $\mathbf{p}$ as it is transformed by f.

Since the tangent map can act on all tangent vectors at once, a diffeomorphism $f : M \to N$ induces a tangent map $Tf : TM \to TN$. It also acts on vector fields, pushing forward a vector field $\mathbf{v}$ defined on M to one defined on N, denoted $f_*(\mathbf{v})$. It is also common to use the pull back $f^*(\mathbf{v})$, defined by pushing forward by f^{-1}. A vector field $\mathbf{v}$ is said to be *invariant* under the action of a diffeomorphism f if $f_*(\mathbf{v}) = \mathbf{v}$. Given a fluid motion ϕ_t and a vector field $\mathbf{v}$ (say a magnetic field), then its *push forward* or *transport* under the flow is $\phi_{t*}(\mathbf{v})$. The vector field is *invariant under the flow* or *frozen into the fluid* if $\phi_{t*}(\mathbf{v}) = \mathbf{v}$ for all times t.

As noted above, vector fields and flows go hand in hand. Since a diffeomorphism can be viewed as a change of coordinates one would expect that pushing forward a vector field and then constructing its flow should give the same result as pushing forward the flow of the original vector field. More precisely, if $\mathbf{u}$ has flow ϕ_t and $\mathbf{v}$ has flow ψ_t, we have $f_*\mathbf{u} = \mathbf{v}$ if and only if $f_*\phi_t = \psi_t$, which by definition says that $f \circ \phi_t \circ f^{-1} = \psi_t$.

3.4. LIE DERIVATIVES

As a general notion, the Lie derivative of a structure with respect to the vector field $\mathbf{u}$ measures the rate of change of the structure as it is transported by the flow of $\mathbf{u}$. Assume now that $\mathbf{u}$ and its corresponding flow ϕ_t are steady and that $\mathbf{v}$ is also time independent, then the Lie derivative of $\mathbf{v}$ with respect to $\mathbf{u}$ is

$$\mathcal{L}_{\mathbf{u}}\mathbf{v} = \left.\frac{d(\phi_t^*\mathbf{v})}{dt}\right|_{t=0} .$$

The Lie derivative is sometimes called "the fisherman's derivative" since it corresponds to sitting at one point and measuring the rate of change as the transported vector field goes by. In Euclidean space a computation yields

$$\mathcal{L}_{\mathbf{u}}\mathbf{v} = (\mathbf{u} \cdot \nabla)\mathbf{v} - (\mathbf{v} \cdot \nabla)\mathbf{u} . \tag{3}$$

Although the Lie derivative by definition only measures what is happening at time $t = 0$, it also captures other times as well. This is expressed in the formula

$$\frac{d}{dt}(\phi_t^*\mathbf{v}) = \phi_t^*(\mathcal{L}_{\mathbf{u}}\mathbf{v}) , \tag{4}$$

which says that at any time the derivative of the pull back of $\mathbf{v}$ is the pull back of the Lie derivative. Equation (4) immediately yields that $\mathcal{L}_{\mathbf{u}}\mathbf{v} = 0$ if and only if $\phi_t^*\mathbf{v} = \mathbf{v}$, and so the vanishing of the Lie derivative is a differential condition that implies that $\mathbf{v}$ is frozen into the flow of $\mathbf{u}$.

The Lie derivative $\mathcal{L}_{\mathbf{u}}\mathbf{v}$ is sometimes written as the Lie bracket $[\mathbf{u}, \mathbf{v}]$ and it has many marvellous algebraic and analytic properties. We mention just two here. The first is that $[\mathbf{u}, \mathbf{v}] = -[\mathbf{v}, \mathbf{u}]$, and so $\mathcal{L}_{\mathbf{u}}\mathbf{v} = 0$ implies $\mathcal{L}_{\mathbf{v}}\mathbf{u} = 0$. In Fluid Mechanics one usually thinks of a fluid flow with velocity field $\mathbf{u}$ and a different kind of physical object, say a magnetic field $\mathbf{v}$, as being transported in the flow. But both are vector fields and can be used to generate flows. If $\mathbf{v}$ is frozen in the flow of $\mathbf{u}$, then we can turn $\mathbf{v}$ into a flow and $\mathbf{u}$ will be frozen into that. Another nice property is that when $[\mathbf{u}, \mathbf{v}] = 0$, the corresponding flows commute, i.e. $\phi_t \circ \psi_s = \psi_s \circ \phi_t$ for all t and s.

We can also define the Lie derivative of a time independent scalar field α as

$$\mathcal{L}_{\mathbf{u}}\alpha = \left.\frac{d(\phi_t^*\alpha)}{dt}\right|_{t=0} .$$

As with vector fields, $\mathcal{L}_{\mathbf{u}}\alpha = 0$ means that α is frozen in, i.e. constant on trajectories of the flow. Since by definition $\phi_t^*\alpha(\mathbf{p}) = \alpha(\phi_t(\mathbf{p}))$, we see that in Euclidean space, the Lie derivative of a steady scalar field is the same as its material derivative and is computed as $D\alpha/Dt = (\mathbf{u}\cdot\nabla)\alpha$.

Example 3.2 The prototypical fluid mechanical example of a frozen in field is the vorticity field $\boldsymbol{\omega} = \mathrm{curl}(\mathbf{u})$ for a steady, incompressible, constant density, Euler flow. In Euclidean space the vector field satisfies the equation

$$(\mathbf{u}\cdot\nabla)\mathbf{u} = -\nabla P \ , \tag{5}$$

where we use a capital P for the pressure and assume that the density is one. Letting $\Phi = \frac{1}{2}(\mathbf{u}\cdot\mathbf{u}) + P$ be the Bernoulli function, standard vector identities turn (5) into

$$\nabla(\Phi) = \mathbf{u}\times\boldsymbol{\omega} \ . \tag{6}$$

Dotting this by $\mathbf{u}$ gives $\mathcal{L}_{\mathbf{u}}\Phi = 0$ and so Φ is constant on flow lines. Taking the curl of (6) gives

$$0 = \mathrm{curl}(\mathbf{u}\times\boldsymbol{\omega}) \ , \tag{7}$$

and vector identities with (3) yield that $\mathcal{L}_{\mathbf{u}}\boldsymbol{\omega} = 0$. More can be obtained by assuming that $\mathbf{u}\times\boldsymbol{\omega} \neq 0$, except at perhaps a finite number of points. Via (6) this implies that except for a finite number of exceptions any level set S of Φ (called a Lamb surface) is a two-manifold. Further, (6) says that $\boldsymbol{\omega}$ as well as $\mathbf{u}$ are tangent to S. Thus restricted to S, since $\mathcal{L}_{\mathbf{u}}\boldsymbol{\omega} = 0$, the flow and the (artificial) flow made from $\boldsymbol{\omega}$ commute. We can now invoke a classical theorem of Liouville which says that S has to be a two-torus or a topological cylinder and the flow of $\mathbf{u}$ on it has constant direction and magnitude (perhaps after a change of coordinates). This is an outline of a basic piece of the Bernoulli-Lamb-Arnol'd theorem. See [4] and [10] for more details.

For an unsteady vector field $\mathbf{u}_t$, the Lie derivative $\mathcal{L}_{\mathbf{u}_t}$ is by definition a family of derivatives, one for each t. To compute a member of the family, one freezes a time and then computes the Lie derivative with respect to that vector field, thus using the time t streamlines and *not* the unsteady flow of $\mathbf{u}_t$. If $\mathbf{v}$ is steady, then (4) with ϕ_t the unsteady fluid motion of $\mathbf{u}_t$ still holds and $\mathbf{v}$ is frozen in if and only if $\mathcal{L}_{\mathbf{u}_t}\mathbf{v} = 0$ for all t. If $\mathbf{v}_t$ is also time dependent then

$$\frac{d\phi_t^*\mathbf{v}_t}{dt} = \phi_t^*(\frac{\partial\mathbf{v}_t}{\partial t}) + \phi_t^*(\mathcal{L}_{\mathbf{u}_t}\mathbf{v}_t) = \phi_t^*(\frac{\partial\mathbf{v}_t}{\partial t} + \mathcal{L}_{\mathbf{u}_t}\mathbf{v}_t) \ .$$

Thus since ϕ_t is a diffeomorphism, the condition for $\mathbf{v}_t$ to be frozen in, $\phi_{t*}v_t = v_0$, can be written

$$\frac{\partial \mathbf{v}_t}{\partial t} + \mathcal{L}_{\mathbf{u}_t}\mathbf{v}_t = 0 \tag{8}$$

for all t.

For more information on the Lie derivative see [25], [3], or [1].

4. Geometry, Metrics, and Connections

4.1. THE NEED FOR ADDITIONAL STRUCTURE

As a fluid flows, subregions of fluid are deformed by the surrounding fluid. The forces involved in these deformation are, in fact, what determine the equations that characterize fluid motions. Since the fluid maps are diffeomorphisms, all topological properties of the evolving subregions stay the same. The nature of the deformation lies in changing angles and lengths, and is therefore geometric. Thus we need a geometric structure on the flow region.

There are several other ways in which geometric considerations can be seen entering into mechanics. Most simply, the magnitude of a velocity vector is required for the kinetic energy. In addition, we have seen that a velocity vector lies in the tangent bundle, and so the acceleration (the velocity of the velocity) lies in the tangent bundle of the tangent bundle. Thus a force vector and the acceleration live in different mathematical objects, and there is no way to equate them as required by Newton's second law.

The acceleration of the fluid is the rate of change of the velocity field along a trajectory, and is thus a special case of what in Fluid Mechanics is called the material derivative. For simplicity, let $\mathbf{u}$ and $\mathbf{v}$ be steady. The derivative we require is the instantaneous rate of change of one vector field in the direction of another. In Mathematics this is called the directional derivative of $\mathbf{v}$ in the direction of $\mathbf{u}$ and is defined in Euclidean space by

$$\nabla_{\mathbf{u}}\mathbf{v}(\mathbf{p}) = \left.\frac{d(\mathbf{v}(\phi_t(\mathbf{p}))}{dt}\right|_{t=0} .$$

The chain rule and the advection equation (2) then yield that $\nabla_{\mathbf{u}}\mathbf{v} = (\nabla\mathbf{v})\cdot\mathbf{u}$, which is more commonly written in Fluid Mechanics as $(\mathbf{u}\cdot\nabla)\mathbf{v}$ where $\nabla\mathbf{v}$ is the derivative matrix of $\mathbf{v}$, sometimes called the velocity gradient 2-tensor, and has components $\frac{\partial v_i}{\partial x_j}$. To uncover the implicit assumptions in this calculation let us return to the definition of the derivative,

$$\nabla_{\mathbf{u}}\mathbf{v}(\mathbf{p}) = \lim_{t\to 0}\frac{\mathbf{v}(\phi_t(\mathbf{p})) - \mathbf{v}(\mathbf{p})}{t} . \tag{9}$$

Thus computing the derivative requires the subtraction of $\mathbf{v}(\phi_t(\mathbf{p}))$ and $\mathbf{v_p}$. But note that the vector $\mathbf{v}(\phi_t(\mathbf{p}))$ lives in the tangent space attached to point $\phi_t(\mathbf{p})$ while $\mathbf{v_p}$ is attached to the point $\mathbf{p}$, and thus they lie in different vector spaces and their difference has no meaning. When we compute this limit in Euclidean space we implicitly identify the two tangent spaces by rigidly transporting one onto the other, but on a manifold there is no natural way to do this without some additional structure.

So we see that there are, in fact, two geometric structures required to proceed with the fluid model. The first is a way to measure geometric quantities like angles and lengths, and the second is a way to compare the geometry at different points in the fluid body. These needs are fulfilled by a Riemannian metric and parallel transport, respectively. The infinitesimal version of parallel transport is the differentiation of one vector field with respect to another, and it is this process that is most fundamental and has been entitled a *linear connection*.

An introduction to Differential Geometry is given in [6], the chapter in [18] is short and sweet, and the redoubtable [23] is comprehensive and comprehensible.

4.2. CONNECTIONS AND COVARIANT DERIVATIVES

We now assume that our fluid region has a linear connection. A connection defines a way of taking derivatives, a process. It is rather algebraic in character, taking a pair of vectors fields $\mathbf{u}$ and $\mathbf{v}$ and returning a third denoted $\nabla_{\mathbf{u}}\mathbf{v}$. It is required to satisfy properties that make it behave like the familiar directional derivative in Euclidean space.

- It is linear in the $\mathbf{u}$ slot with respect to multiplication by scalar fields, $\nabla_{\alpha\mathbf{u}+\hat{\alpha}\hat{\mathbf{u}}}\mathbf{v} = \alpha\nabla_{\mathbf{u}}\mathbf{v} + \hat{\alpha}\nabla_{\hat{\mathbf{u}}}\mathbf{v}$.
- It is linear in the $\mathbf{v}$ slot with respect to multiplication by constants, $\nabla_{\mathbf{u}}(r\mathbf{v} + \hat{r}\hat{\mathbf{v}}) = r\nabla_{\mathbf{u}}\mathbf{v} + \hat{r}\nabla\hat{\mathbf{v}}$.
- When one multiplies by a scalar field in the $\mathbf{v}$ slot it must satisfy the product rule, $\nabla_{\mathbf{u}}(\alpha\mathbf{v}) = \alpha\nabla_{\mathbf{u}}\mathbf{v} + (\mathcal{L}_{\mathbf{u}}\alpha)\mathbf{v}$.

The connection itself is indicated by the symbol ∇, and $\nabla_{\mathbf{u}}\mathbf{v}$ is called the *covariant derivative* of $\mathbf{v}$ in the direction of $\mathbf{u}$. The standard connection on Euclidean space is the usual convective or directional derivative, $\nabla_{\mathbf{u}}\mathbf{v} = (\mathbf{u}\cdot\nabla)\mathbf{v}$.

In order to use the connection to compare different points we first connect them with a curve, which we parameterize as $\gamma(s)$, with s not necessarily the arc length. For simplicity we assume that γ is the integral curve of a vector field $\mathbf{u}$, i.e. the derivative of the curve $\gamma'(s)$ at the point $\gamma(s)$ is the value of the vector field at that point, $\gamma'(s) = \mathbf{u}(\gamma(s))$. Physically this curve could correspond to a fluid trajectory or perhaps a field line. We define the *covariant derivative* of the vector field $\mathbf{v}$ along γ as

$$\frac{D\mathbf{v}}{Ds}(s) = \nabla_{\mathbf{u}}\mathbf{v}(\gamma(s)) \ .$$

Note that $\frac{D\mathbf{v}}{Ds}$ is a vector field which is defined along γ. The vector field $\mathbf{v}$ is said to be parallel along γ if $\frac{D\mathbf{v}}{Ds} = 0$ on the whole curve. If γ starts at the point $\mathbf{p}$ and $\mathbf{w_p}$ is a tangent vector, there is a unique way to continue $\mathbf{w_p}$ to a vector field that is parallel along γ. In this way any tangent vector at $\mathbf{p}$ can be *parallel transported* to a tangent vector based at $\gamma(s)$. This process defines the parallel transport maps $\mathcal{P}_s$ from the tangent space at $\mathbf{p}$ to the tangent space at $\gamma(s)$.

Although we have assumed the existence of a connection and used it to define parallel transport note that our discussion is consistent in the sense that using $\mathcal{P}$ as the analog of implicit rigid translation in (9) does correctly compute the derivative because

$$\nabla_{\mathbf{u}}\mathbf{v}(\mathbf{p}) = \left.\frac{d\mathcal{P}_s^*\mathbf{v}}{ds}\right|_{s=0} = \lim_{s\to 0} \frac{\mathcal{P}_s^*\mathbf{v} - \mathbf{v}(\mathbf{p})}{s} \ . \tag{10}$$

It is important to note that the parallel transport of tangent vectors between two points usually depends on the curve we choose between them. In fact the only situation in which all parallel transport is independent of path is when there is no curvature. In spite of this, the limit in (10) is independent of the choice of curve, and in fact, it may be used to define $\nabla_{\mathbf{w_p}}\mathbf{v}$ for a given vector $\mathbf{w_p}$, since the limit will not depend on how $\mathbf{w_p}$ is extended to a vector field on the whole fluid region.

Given a connection, the acceleration of fluid particles in the steady case is $\nabla_{\mathbf{u}}\mathbf{u}$ which is, as required, a vector field in the tangent bundle and not in the tangent tangent bundle.

4.3. RIEMANNIAN METRICS

We also need to quantify such geometric notions as lengths, angles and volumes. In Euclidean space this is done via the usual inner product $\vec{u}\cdot\vec{v} = \sum u_i v_i$, which defines a length as $\|\vec{v}\| = (\vec{v}\cdot\vec{v})^{\frac{1}{2}}$ and an angle using $\frac{\vec{v}\cdot\vec{u}}{\|\vec{v}\|\|\vec{u}\|}$. In general, an *inner product* on a vector space is a rule that takes two vectors and returns a number. A usual notation is $\langle\vec{u},\vec{v}\rangle$ or $\iota(\vec{u},\vec{v})$. It is further required that an inner product be linear in each argument, symmetric $\iota(\vec{u},\vec{v}) = \iota(\vec{v},\vec{u})$, and positive definite $\vec{v}\neq 0$ implies $\iota(\vec{v},\vec{v}) > 0$. A linear isomorphism L (i.e. a linear bijection) induces a pull back on a inner product as $L^*(\iota)(\vec{u},\vec{v}) = \iota(L(\vec{u}),L(\vec{v}))$. The isomorphism is said to be an *isometry* between the inner products $\hat{\iota}$ and ι if it preserves the inner product, $L^*(\iota) = \hat{\iota}$.

On a curved manifold the way of measuring lengths and angles vary from point to point and so the geometry is specified by a *Riemannian metric* or *metric tensor* which is a family of inner products, with one defined on the tangent space to each point. The metric itself is denoted g, the particular inner product at the point $\mathbf{p}$ is denoted $g_\mathbf{p}$ (but the subscript is often suppressed), and its value on a pair of vectors $\mathbf{v}_\mathbf{p}$ and $\mathbf{u}_\mathbf{p}$ based at $\mathbf{p}$ is $g_\mathbf{p}(\mathbf{v}_\mathbf{p}, \mathbf{u}_\mathbf{p})$. A metric is often indicated by its components in a basis as g_{ij}, or as a line element $ds^2 = g_{11}\, dx_1^2 + g_{22}\, dx_2^2 + \ldots$. Recall from Section 3.3 that a diffeomorphism $h : \hat{M} \to M$ induces a linear isomorphism on tangent spaces via the tangent maps, $Th_\mathbf{p} : T\hat{M}_\mathbf{p} \to TM_{h(\mathbf{p})}$. If M has metric g, then each of these $Th_\mathbf{p}$ induce a pull back of $g_\mathbf{p}$, which in turns creates a pull back of the entire metric which is denoted h^*g. If $\hat{M}$ has a metric $\hat{g}$, then h is called an *isometry* if Th is an isometry on every tangent space, i.e. if $h^*g = \hat{g}$. Isometries are the isomorphisms in the category of Riemannian manifolds.

The Riemannian metric immediately allows to define the speed, or magnitude of a velocity vector as $\|\mathbf{u}_\mathbf{p}\|_g = g_\mathbf{p}(\mathbf{u}_\mathbf{p}, \mathbf{u}_\mathbf{p})^{1/2}$ and the length of the curve $\gamma : [a, b] \to B$ as

$$\int_a^b \|\gamma'(s)\|_g \, ds \ . \tag{11}$$

Parallel transport as defined by the connection was introduced as the analog of rigid translation in Euclidean space. In particular, it should preserve lengths and angles, and so it should preserve the inner products on tangent spaces. Thus we now require our connection to be *compatible with the metric* in the sense that the parallel transport maps, $\mathcal{P}_s$, must induce an isometry of the tangent spaces at points along the curve.

There are actually many connections that are compatible with a given metric. To get a unique connection, we also require that the connection be compatible with the Lie derivative in the same way it is in Euclidean space

$$\mathcal{L}_\mathbf{u}\mathbf{v} = \nabla_\mathbf{u}\mathbf{v} - \nabla_\mathbf{v}\mathbf{u} \ .$$

If this condition holds the connection is said to be *symmetric* or *torsion free*. The fundamental theorem of Riemannian geometry states that there is a unique symmetric connection (called the *Riemannian* or *Levi-Civita* connection) that is compatible with a given Riemannian metric. Henceforth we assume that our fluid region has a Riemannian metric g and ∇ is its connection.

4.4. VOLUME FORMS

In Euclidean space we get the volume of a box from the product of its side lengths. A Riemannian metric determines a way of computing length and so it also provides a way to compute a volume. As with the inner product we begin with the notion of a linear volume on $\mathbb{R}^3$. Given three vectors, the Euclidean volume of the parallelepiped they span is the determinant of the matrix whose columns are the three vectors. The appropriate notion of volume generalizes this situation. A volume element σ is a map that takes three vectors and returns a number, $\sigma(\vec{u}, \vec{v}, \vec{w})$. Since a volume element is supposed to act like the usual determinant, it is required to be linear in all three arguments and interchanging two vectors must change the sign. A linear isomorphism L induces a pull back on a volume element as $L^*(\sigma)(\vec{u}, \vec{v}, \vec{w}) = \sigma(L(\vec{u}), L(\vec{v}), \mathrm{L}(\vec{w}))$.

Exercise 4.1 If σ and $\widehat{\sigma}$ are volume elements on $\mathbb{R}^3$, show that there is a number r so that $\widehat{\sigma} = r\,\sigma$, i.e. there is a single r that works for all choices of triples of input vectors.

A family of volume elements on a manifold, one for each tangent space, is called a *volume form*. A diffeomorphism, f, induces a pull back, $f^*\chi$, of a volume form χ in a manner completely analogous to the pull back of an inner product. The diffeomorphism *preserves the volume form* if $f^*\chi = \chi$.

To construct a volume form that is compatible with the metric first note that as a consequence of the exercise, given an inner product on $\mathbb{R}^3$ and an orthonormal basis b there is a unique volume element which gives it volume 1. Switching exactly two elements in the chosen orthonormal basis yields a basis $\hat{b}$ with volume -1. Every other orthonormal basis will be assigned a volume of 1 or -1, and the two classes are distinguished by whether the basis in question can be continuously moved to b or to $\hat{b}$. A Riemannian manifold has an inner product on each tangent space, and so we may choose an orthonormal basis on each tangent space and thus obtain a compatible volume element. Since parallel transport is an isometry on tangent spaces, it takes orthonormal bases to orthonormal bases. If our volume elements are to fit together in a nice family, the volume of an orthonormal basis should not be changed by parallel transport. If the volume elements can be chosen so this is possible, the manifold is called *orientable*. The simplest example of a non-orientable manifold is the two-dimensional Möbious band with the metric induced from how it lives in $\mathbb{R}^3$. As one traverses the core circle, an orthonormal basis will come back with one vector flipped, resulting in a basis assigned the opposite area it had initially.

We now assume that the fluid region is orientable, and for definiteness we pick a orthonormal basis b on one tangent space and a volume element which assigns it a volume 1. Then we can find a family of volume elements, one on each tangent space, which give all parallel translates of b the volume 1. The resulting object is called a *Riemannian volume form*, and we denote it μ. Note that switching a pair of elements of the original basis b would result in the negative of the volume form we have chosen. Applying Exercise 4.1

on each tangent space shows that for any volume form χ there is a scalar field r called the *density* with $\chi = r\mu$.

In (11) we computed the length of a curve by integrating the norm of its tangent vectors as defined by the metric. In the same way we can compute the volume of a region by integrating the value of the volume form on its tangent vectors. This is done most succinctly by using the pull back under its parameterization. We begin with the case where U is a region in $\mathbb{R}^3$ with the Euclidean form μ_E. Since μ_E should give the usual volumes, if χ is a volume form on U with density r we define

$$\int_U \chi = \int_U r\mu_E = \int_U r(x,y,z)\; dx\, dy\, dz\ .$$

Now if V is a region in a manifold that fits into a single coordinate chart it has a parameterization in terms of a subset of Euclidean space. This is a diffeomorphism h from some region U in $\mathbb{R}^3$ onto V. If χ is a volume form on the manifold define

$$\int_V \chi = \int_U h^*\chi\ ,$$

where to be very careful we must insist that the choice of h was such that $\int_U h^*(\mu) > 0$ for the Riemannian volume form μ on V, or equivalently that $h^*(\mu) = r\,\mu_E$ for a positive function r . If the region V is large, chop it up into smaller pieces that fit into charts and add the integrals.

Volumes and measures are closely related and in many cases one can pass from one to the other. If μ a Riemannian volume form, the volume form χ is *positive* if its density r is a positive function. Given a positive volume form we can define a measure via integration $m(V) = \int_V \chi$ on nice sets V and on nastier sets using limits of nice sets. In a steady fluid motion if ρ is the mass density, then $\chi = \rho\mu$ is the *mass form* and the flow is mass preserving if that form is preserved, $\phi_t^*(\chi) = \chi$. This implies that the corresponding measure is preserved, $\phi_t^*(m) = m$. One advantage which volume forms have over measures is that they are linear objects and so can be added, subtracted, differentiated, etc.

Recall that in Euclidean space the Jacobian $J(f)$ of a diffeomorphism f is the determinant of the Jacobian matrix, i.e. $\det(\nabla f)$. Since the determinant of a linear transformation is equal to the (signed) volume of the image of the unit cube, and the derivative (the Jacobian matrix) is the best local linear approximation of the map, we see that the Jacobian computes the local change in volume under the map f. It follows fairly easily that for a diffeomorphism f, $f_*(\mu_E) = J(f)\mu_E$. Thus, in particular, a fluid motion is incompressible exactly when $J(\phi_t) = 1$ for all times and at all points. Similarly, a flow preserves a mass with density ρ if and only if $\rho\mu_E = \phi_t^*(\rho\mu_E) = \phi_t^*(\rho)\phi_t^*(\mu_E) = \phi_t^*(\rho)J(\phi_t)\mu_E$, which reduces to a scalar equation $\phi_t^*(\rho) = J(\phi_t)\rho$. The same derivation works if ϕ_t is unsteady and ρ is time dependent in which case the scalar equation is written in the familiar form $\rho = J\rho_0$.

On the general Riemannian manifold case one *defines* the Jacobian (with respect to μ) of the diffeomorphism h as the scalar field $J(h)$ that satisfies $h_*(\mu) = J(h)\mu$, which is to say that $J(h)$ is the density of $h_*(\mu)$. This is a common strategy in coordinate free definitions, a property of an object in Euclidean space is taken as the coordinate free definition.

4.5. THE LIE, COVARIANT AND MATERIAL DERIVATIVES

Both the Lie derivative and the covariant derivative measure a rate of change of one vector field with respect to another. In this subsection we compare them and remark on their generalizations. We put the definitions side by side

$$\nabla_{\mathbf{u}}\mathbf{v} = \left.\frac{d\mathcal{P}_t^*(\mathbf{v})}{dt}\right|_{t=0} \quad \text{and} \quad \mathcal{L}_{\mathbf{u}}\mathbf{v} = \left.\frac{d\phi_t^*(\mathbf{v})}{dt}\right|_{t=0}, \tag{12}$$

and note that the only difference is the way in which the vector field is pulled back. In the Lie case one pulls back the advected vector field using the derivative of the flow and so the Lie derivative measures the deformations of the advected vector field. In contrast, in the covariant derivative one pulls back using an isometry, and so this derivative does not measure the deformations during advection, but rather just how the vector field is changing with respect to the metric. If $\mathcal{L}_{\mathbf{u}}\mathbf{v} = 0$, it means that the vector field $\mathbf{v}$ is transported to itself, i.e. $\phi_{t*}\mathbf{v} = \mathbf{v}$, which is a strictly topological notion, and indeed the Lie derivative requires only Differential Topology. On the other hand, $\nabla_{\mathbf{u}}\mathbf{v} = 0$ along a trajectory means that the vector field is constant along the trajectory *as measured by the metric*, and indeed $\nabla_{\mathbf{u}}\mathbf{v} = 0$ is used to define the geometric notion of parallel transport.

In analogy with (12) we may compute the Lie and covariant derivative of any kind of structure that can be pulled back and subtracted. Such structures include Riemannian metrics and volume forms. These are examples of *contravariant k-tensors* which are families of multilinear maps, one for each tangent space, which take k tangent vectors as input and return a number at each point, so as a global object they return a scalar field. Note that what most mathematicians call a contravariant tensor is called a covariant tensor by most engineers and physicists, and vice versa. Only one kind of tensor arises here so we drop the potentially confusing adjective; a Riemannian metric is a 2-tensor and a volume form is a 3-tensor. In general, if τ

is a tensor, define

$$\nabla_{\mathbf{u}}\tau = \left.\frac{d\mathcal{P}_t^*(\tau)}{dt}\right|_{t=0} \quad \text{and} \quad \mathcal{L}_{\mathbf{u}}\tau = \left.\frac{d\phi_t^*(\tau)}{dt}\right|_{t=0} . \tag{13}$$

In this case the Lie and covariant derivatives have the same dynamical meaning as in the vector field case. The condition $\mathcal{L}_{\mathbf{u}}\tau = 0$ says that $\phi_t^*(\tau) = \tau$ and so τ is frozen into the flow. On the hand, $\nabla_{\mathbf{u}}\tau = 0$, means that from the point of view of the metric, τ is not changing along trajectories.

Example 4.2 The Riemannian connection associated with a metric g was required to have the property that each $\mathcal{P}_s$ is an isometry. This says that $\mathcal{P}_t^* g = g$, or that $\nabla_{\mathbf{u}} g = 0$ for any $\mathbf{u}$. In contrast, $\mathcal{L}_{\mathbf{u}} g = 0$ is a very rare situation. It says that the metric is invariant under the flow, and so each ϕ_t is an isometry. Such a $\mathbf{u}$ is called a *Killing field* for the metric and in Euclidean space with the usual metric any flow of a Killing field is either a rigid rotation or a translation, or composition of these.

A scalar field α is a 0-tensor and in this case the two derivatives are equal, $\mathcal{L}_{\mathbf{u}}\alpha = \nabla_{\mathbf{u}}\alpha$. This is because the pull backs are the same; an advected scalar field does not feel the local deformations during the evolution.

The Lie and covariant derivatives behave differently when transported by a diffeomorphism h. The Lie derivative satisfies $h^*(\mathcal{L}_{\mathbf{u}}\mathbf{v}) = \mathcal{L}_{h^*\mathbf{u}} h^*\mathbf{v}$. This is sometimes called the "naturalness" of the Lie derivative, and expresses the fact that the definition of the Lie derivative requires only Differential Topology and so is preserved by a diffeomorphism. Thus the same definition works before and after the action of h, or in symbols $h^*(\mathcal{L}) = \mathcal{L}$. On the other hand, the covariant derivative depends on the metric, and so $h^*(\nabla)$ is the Riemannian connection $\hat{\nabla}$ of the pulled back metric $h^*(g)$, and so $h^*(\nabla_{\mathbf{u}}\mathbf{v}) = \hat{\nabla}_{h^*\mathbf{u}} h^*\mathbf{v}$.

In Euclidean space the *material derivative* with respect to $\mathbf{u}$ is the operator $\frac{D}{Dt} = \frac{\partial}{\partial t} + (\mathbf{u} \cdot \nabla)$ and it computes the rate of change of say a time dependent vector field $\mathbf{v}$ along the flow as

$$\frac{D\mathbf{v}}{Dt}(\phi_t(\mathbf{p})) = \frac{\partial \mathbf{v}(\phi_t(\mathbf{p}), t)}{\partial t} . \tag{14}$$

The obvious generalization using the covariant derivative $\frac{D}{Dt} = \frac{\partial}{\partial t} + \nabla_{\mathbf{u}}$ maintains this meaning. Thus, for example, the acceleration of a fluid particle under the perhaps unsteady velocity field $\mathbf{u}$ is given as usual by $\frac{D\mathbf{u}}{Dt}$, and for a perhaps time dependent scalar field the familiar condition $\frac{D\alpha}{Dt} = 0$ says that the scalar field is passively transported.

4.6. DIV, GRAD, CURL AND THE LAPLACIAN

We may now give coordinate free definitions of the standard operations on velocity fields. In Euclidean space these can be built out of the velocity gradient tensor, $\nabla \mathbf{u} = (\partial u_i/\partial x_j)$. Its generalization is the *covariant derivative* of $\mathbf{u}$, also denoted $\nabla \mathbf{u}$. In the previous subsections we took covariant derivatives in the direction of $\mathbf{u}$. To define the derivative of $\mathbf{u}$ itself recall that a connection is required to be linear over scalar fields in the direction in which one is differentiating, i.e. $\nabla_{\mathbf{v}}\mathbf{u}$ is linear in the $\mathbf{v}$ slot. Thus $\nabla \mathbf{u}$ can be defined as the linear transformation on each tangent space which inputs a vector $\mathbf{v_p}$ and outputs the rate of change of $\mathbf{u}$ in that direction $\nabla_{\mathbf{v_p}}\mathbf{u}$. Thus $\nabla \mathbf{u}(\mathbf{v})$ means $\nabla_{\mathbf{v}}\mathbf{u}$.

In a similar vein, given a scalar field α its derivative is the linear functional that inputs a vector $\mathbf{v_p}$ and outputs the rate of change of α in that direction $\nabla_{\mathbf{v_p}}\alpha$. Since the notation $\nabla\alpha$ has a commonly accepted meaning as a vector field, the linear functional is called $d\alpha$, and so $d\alpha(\mathbf{v_p}) = \nabla_{\mathbf{v_p}}\alpha$. The *gradient vector field* $\nabla\alpha$ is defined using the metric as the unique vector field that satisfies $d\alpha(\mathbf{v_p}) = g(\nabla\alpha, \mathbf{v_p})$ for all vectors $\mathbf{v_p}$. This means that $\nabla\alpha$ is the direction of maximum increase of α as measured by the metric, in accord with the usual situation in Euclidean space where $\nabla_{\mathbf{v_p}}\alpha = \mathbf{v_p}\cdot\nabla\alpha$.

Since $\nabla \mathbf{u}$ gives a linear transformation on each tangent space, its trace is independent of the choice of basis and so we may define $\mathrm{div}(\mathbf{u}) = \mathrm{trace}(\nabla \mathbf{u})$. As in the Euclidean case the divergence measures the infinitesimal rate of change of volumes as they are transported which is expressed by $\mathcal{L}_{\mathbf{u}}\mu = \mathrm{div}(\mathbf{u})\mu$, where μ is the Riemannian volume form. Thus as usual, ϕ_t is incompressible exactly when $\mathrm{div}(\mathbf{u}) = 0$. If the perhaps time dependent density is ρ_t, then the mass form is $\rho_t\mu$, and conservation of mass says that $\phi_t^*(\rho_t\mu) = \rho_0\mu$. Thus using the analog of (8) for tensors we have that

$$0 = \frac{\partial(\rho_t\mu)}{\partial t} + \mathcal{L}_{\mathbf{u}}(\rho_t\mu)\ .$$

Now Leibnitz' rule for the Lie derivative says that $\mathcal{L}_{\mathbf{u}}(\rho_t\mu) = (\mathcal{L}_{\mathbf{u}}\rho_t)\mu + \rho_t\mathcal{L}_{\mathbf{u}}\mu = (\nabla_{\mathbf{u}}\rho_t)\mu + \rho_t\,\mathrm{div}(\mathbf{u})\mu$, using the definition of the divergence and the equality of the Lie and covariant derivatives for scalar fields. Since μ is time independent $\frac{\partial(\rho_t\mu)}{\partial t} = \frac{\partial\rho_t}{\partial t}\mu$. Equating the coefficients of μ and using the definition of the material derivative yields the *mass conservation equation* or *continuity equation*

$$\frac{D\rho_t}{Dt} + \rho_t\mathrm{div}(\mathbf{u}) = 0\ . \tag{15}$$

Note that the preceding derivation did not require $\mathbf{u}$ to be steady.

In Euclidean space the symmetric and skew symmetric parts of $\nabla \mathbf{u}$ yield the deformation and rotation tensor. To formulate the generalization we require a transpose. Recall that in $\mathbb{R}^3$, the transpose of a linear transformation A with respect to inner product ι is the unique linear transformation A^T that satisfies $\iota(A^T(\vec{v}), \vec{w}) = \iota(\vec{v}, A(\vec{w}))$. Thus working with the metric g we can define $(\nabla\mathbf{u})^T$ as satisfying $g((\nabla\mathbf{u})^T(\mathbf{v_p}), \mathbf{w_p}) = g(\mathbf{v_p}, \nabla\mathbf{u}(\mathbf{w_p}))$ on each tangent space.

We then define the symmetric part of $\nabla\mathbf{u}$ as $D(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + (\nabla\mathbf{u})^T)$ and the skew symmetric part as $\Omega(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} - (\nabla\mathbf{u})^T)$, and so $\nabla\mathbf{u} = D(\mathbf{u}) + \Omega(\mathbf{u})$. Since D is symmetric, it has three orthogonal principal directions and D diagonalizes in that basis. The diagonal elements represent the infinitesimal deformation rates in each direction and so D is called the *deformation tensor* (although at this point it is formally a linear transformation). In Euclidean space Ω is related to $\boldsymbol{\omega} = \mathrm{curl}(\mathbf{u})$ by the formula

$$\Omega(\vec{v}) = \frac{1}{2}\boldsymbol{\omega} \times \vec{v} \ , \tag{16}$$

for any vector $\vec{v}$, and so Ω is sometimes called the *rotation tensor*.

We can find our way to the covariant definition of the curl by dotting both sides of (16) by $\vec{w}$, yielding

$$\vec{w} \cdot \Omega(\vec{v}) = \frac{1}{2}\vec{w} \cdot (\boldsymbol{\omega} \times \vec{v}) = \frac{1}{2}\det(\vec{w}, \boldsymbol{\omega}, \vec{v}) \ , \tag{17}$$

where we have used the standard identity connecting the triple scalar product and the determinant of the matrix whose columns are the vectors. All of the terms in (17) have a covariant generalization, so in a now familiar move we define $\mathrm{curl}(\mathbf{u})$ as the unique vector field $\boldsymbol{\omega}$ satisfying

$$g(\vec{w}, \Omega(\vec{v})) = \frac{1}{2}\mu(\vec{w}, \boldsymbol{\omega}, \vec{v}) \ . \tag{18}$$

Since the metric quantifies deformation, one would expect a close relationship between the deformation tensor and the metric. Since the metric is a 2-tensor, we will change D from a linear transformation to a 2 tensor. There is a standard way to do this using the metric called lowering the indices or the $\flat$ operator. The $\widetilde{D(\mathbf{u})}$ which is associated with $D(\mathbf{u})$ is the unique 2-tensor $\widetilde{D(\mathbf{u})}$ with $\widetilde{D(\mathbf{u})}(\mathbf{w_p}, \mathbf{v_p}) = g(D(\mathbf{u})(\mathbf{w_p}), \mathbf{v_p})$. The relationship of the deformation tensor and the metric is expressed by $\mathcal{L}_{\mathbf{u}}g = 2\widetilde{D(\mathbf{u})}$, which says that $\widetilde{D(\mathbf{u})}$ exactly measures how the metric is deformed as it is advected by the flow. If we also turn Ω into a 2-tensor we find that the curl satisfies $\mu(\mathbf{w_p}, \mathrm{curl}(\mathbf{u}), \mathbf{v_p}) = 2\widetilde{\Omega(\mathbf{u})}(\mathbf{w_p}, \mathbf{v_p})$.

The Laplacian of a scalar field is covariantly defined as $\triangle(\alpha) = \mathrm{div}(\nabla\alpha)$. Since the first derivative of a vector field requires the use of a connection one might suspect that the Laplacian of a vector field would require yet more structure. This is fortunately not the case. We can use (13) to define the covariant derivative of $\nabla\mathbf{u}$ in a given direction. Treating the result as a function of the direction we get an object denoted $(\nabla\nabla)\mathbf{u}$ which takes two vectors as input and gives another as output. The Laplacian of $\mathbf{u}$ is the trace of this object, $\triangle\mathbf{u} = \mathrm{trace}((\nabla\nabla)\mathbf{u})$. To be clear on the type of trace we are taking, if we choose an orthonormal basis, e_i, with respect to the metric on each tangent space, then $\triangle\mathbf{u} = \sum_i (\nabla\nabla)\mathbf{u}(e_i, e_i)$.

The Laplacian we have just defined is sometimes called the *analyst's Laplacian.* The *topologist's Laplacian* is defined using the (negative of the) analog of the Euclidean space formula $\triangle\mathbf{u} = \nabla(\mathrm{div}(\mathbf{u})) - \mathrm{curl}\,\mathrm{curl}\,\mathbf{u}$. After a sign switch the two Laplacians differ by the Ricci curvature.

Differential forms are the other common way to give coordinate free definitions of the standard vector calculus notions. Both points of view are important and have their virtues: forms work well with integration and are directly connected to the underlying topology, but the covariant derivative is most naturally related to velocity fields. The two methods are intimately connected and we have, in fact, already encountered the 1-form $d\alpha$, the 2-form $\tilde{\Omega}$, and the three form μ. Forms were not explicitly discussed here only because space and time limitations demanded the most direct path to the goal. The reader is urged to consult [9], [25], [2], [3] or [1].

5. Equations

We now have the mathematical equipment to bring forces into the model and state the basic dynamical equations of Fluid Mechanics. This is familiar material for fluid mechanicians, but for completeness we give a brief summary. Assume that there are no external forces and so the only forces to consider are internal, the force that the fluid body exerts on a subregion across its boundary. The force per unit area on the bounding surface is the *stress* and its exact form is encapsulated in the existence and properties of the Cauchy stress tensor. The usual derivation of the basic dynamic equations in Euclidean space invokes Newton's laws to say that the rate of change of momentum of a patch of fluid is equal to the total surfaces forces on it. If it is assumed that the only stresses are normal to the bounding surface, one obtains Euler's equation

$$\rho(\frac{\partial\mathbf{u}}{\partial t} + \nabla_{\mathbf{u}}\mathbf{u}) = -\nabla P \ . \tag{E}$$

If tangential components of the stress are included one obtains the Navier-Stokes equation

$$\rho(\frac{\partial\mathbf{u}}{\partial t} + \nabla_{\mathbf{u}}\mathbf{u}) = -\nabla P + \bar{\mu}\triangle\mathbf{u} \ . \tag{NS}$$

In these equations P is the pressure, ρ is the mass density and $\bar{\mu}$ is the viscosity (not a volume form!). Note that all the operations involved in the equations have been covariantly defined. However, there are serious and subtle problems in trying to formulate the usual derivations of (NS) on a manifold. We refer the reader to [16] for a careful exposition and merely remark that the Laplacian in (NS) is the analyst's Laplacian.

Depending on the relative importance of viscosity in the fluid system under study either (E) or (NS) is adapted as the basic dynamical equation. The appropriate boundary conditions are slip and no slip, respectively. For a complete system in which all variables are determined additional equations must be added. In the most common situations it can be assumed that the viscosity is constant throughout the fluid, mass is conserved, and the fluid is incompressible. After inclusion of the mass conservation equation (15), incompressibility is equivalent to $\frac{D\rho}{Dt} = 0$, or that the density of a particle remains constant as it is transported. The simplest compressible systems use thermodynamic considerations to justify the assumption that the pressure and the density are functionally dependent.

There are three obligatory remarks to be made. The first is that all our mathematical modelling would be meaningless if it were not for the fact that the resulting models and equations give results that agree extremely well with experimental data. The second is that the existence-uniqueness theory of the Navier-Stokes and Euler equations is still far from being understood. The third is that what's in this paper just sets the stage; the real action is the understanding and prediction of fluid behavior.

One advantage of having defined all operations covariantly is that changing coordinates or regions with a diffeomorphism $h : B \to \hat{B}$ preserves the property of being a solution to a system of fluid equations. More precisely, if $\mathbf{u}$ satisfies a system on B with respect to the metric g, then $h_*\mathbf{u}$ satisfies the same system on $\hat{B}$. That's the good news. The bad news that the operations in the equations on $\hat{B}$ such as ∇ and $\triangle$ must be defined in terms of the metric h_*g which may not be the metric which you care about. This begs a question:

Question 5.1 Is there a physical meaning to doing Fluid Mechanics with a general Riemannian metric?

We only hazard a few remarks. If there is a general physical interpretation of a curved metric, it cannot involve an intrinsic property of the fluid because everything in the fluid is advected, and the metric (at least as developed here) stays fixed on the manifold. There are few cases where it is clear that fluid flows over a curved space. One is the surface of the earth. Another is in very large scale fluid mechanical models in Cosmology where one can need to take into account the curvature of space-time. Also note that changing from Euclidean into curvilinear, non-orthogonal coordinates forces one to work with the push forward of the Euclidean metric. This is a rather special metric, however, being by definition isometric to the Euclidean one and therefore lacking curvature. The last remark comes from the philosophy of Mathematics described in the introduction: by understanding fluid mechanics in the most general context in which it makes sense, one gains new insights into the particular cases of interest.

Acknowledgements

Thanks to David Groisser for several very useful conversations during the preparation of this paper.

References

1. Abraham, R., Marsden, J. & Ratiu, T. (1988) *Manifolds, Tensor Analysis, and Applications.* Springer Verlag.
2. Aris, R. (1962) *Vectors, Tensors, and the Basic Equations of Fluid Mechanics.* Dover Publications.
3. Arnol'd, V. (1989) *Mathematical Methods in Classical Mechanics.* Springer Verlag.
4. Arnol'd, V. & Khesin, B. (1998) *Topological Methods in Hydrodynamics.* Springer Verlag.
5. Batchelor, G.K. (1967) *An Introduction to Fluid Mechanics.* Cambridge University Press.
6. Boothby, W. (1986) *An Introduction to Differentiable Manifolds and Riemannian Geometry.* Academic Press.
7. Chorin, A. & Marsden, J. (1993) *A Mathematical Introduction to Fluid Mechanics.* Springer Verlag.
8. Christenson, C. & Voxman, W. (1998) *Aspects of Topology,* 2nd ed. BCS Associates.
9. Flanders, H. (1963) *Differential Forms with Applications to the Physical Sciences.* Academic Press.
10. Ghrist, R. & Komendarczyk, R. (2001) Topological features of inviscid flows. this volume.
11. Guillemin, V. & Pollack, A. (1974) *Differential Topology.* Prentice-Hall.
12. Halmos, P. (1950) *Measure theory.* Van Nostrand.
13. Halmos, P. (1960) *Naive set theory.* Van Nostrand.
14. Hirsch, M. (1976) *Differential Topology.* Springer Verlag.
15. MacLane, S. (1998) *Categories for the Working Mathematician,* 2nd ed. Springer Verlag.
16. Marsden, J. & Hughs, T. (1994) *Mathematical Foundations of Elasticity.* Dover Publications.
17. Milnor, J. (1965) *Topology from the Differentiable Viewpoint.* University Press of Virginia.
18. Milnor, J. (1969) *Morse Theory.* Princeton University Press.
19. Moser, J., (1965) On the volume elements on a manifold. *Trans. Amer. Math. Soc.* **120**, 286–294.
20. Munkres, J. (1975) *Topology; a First Course.* Prentice-Hall.
21. Oxtoby, J. & Ulam, S. (1941) Measure-preserving homeomorphisms and metrical transitivity. *Ann. of Math.* **(2) 42**, 874–920.
22. Serrin, J. (1959) Mathematical Principles of Classical Fluid Mechanics. *Handbuch der Physik* **VIII/1**.

23. Spivak, M. (1979) *A Comprehensive Introduction to Differential Geometry,* vol. I-V. Publish or Perish, Inc.
24. Spivak, M. (1965) *Calculus on Manifolds.* W. A. Benjamin.
25. Tur, A. & Yanovsky, V. (1993) Invariants in dissipationless hydrodynamic media. *J. Fluid Mech* **248**, 67–106.
26. Vilenkin, N. (1968) *Stories about Sets.* Academic Press.