

DIV, GRAD, CURL ARE DEAD

William L. Burke
©1993, 1994, 1995 William L. Burke

Contents

Preface

Introduction

CHAPTER ZERO. PROLEGOMENA

1. Patterns
2. Linear Algebra
3. Affine Algebra
4. Computer Algebra Systems

CHAPTER ONE. CALCULUS

5. Vectors
6. Calculus of Curves
7. What is Calculus?
8. Linear Transformations
9. The Differential in Optics

PRELIMINARY
Draft II

CHAPTER TWO. GRADIENTS AND FLOW

10. The Differential of a Function
11. The algebra of 1-forms
12. 1-forms in mechanics
13. Flow in 2D
14. Density
15. Flow in 3D
16. Drawing Vectors: The Eight Icons

CHAPTER THREE. DIFFERENTIAL FORMS

17. Ordinary Differential Forms
18. Twisted Forms
19. Maps and Pullback

CHAPTER FOUR. CALCULUS OF FORMS

20. Exterior Derivatives
21. Integrals
22. Stokes' theorem
23. Symmetries and conservation laws
24. Inverse of Stokes' theorem
25. Delta Functions

CHAPTER FIVE. GEOMETRY WITH FORMS

- 26. Tensor algebra
- 27. Star Operator
- 28. Symmetry and Lie Derivatives
- 29. Old Vector Calculus
- 30. Simplicial Calculus

CHAPTER SIX. ELECTRODYNAMICS

- 31. Maxwell's equations
- 32. Electrostatics
- 33. Energy in Electrostatics
- 34. Magnetism
- 35. Maxwell's equations in spacetime

CHAPTER SEVEN. MECHANICS

- 36. Hamiltonian Mechanics
- 37. Lagrangian Mechanics
- 38. Fluid Mechanics

CHAPTER EIGHT. WAVES

- 39. Dispersive Waves
- 40. Group Velocity
- 41. Water Waves

APPENDICES

- 42. Forms Calculator

Preface

In 1978 Walter Thirring wrote in his multivolume series on theoretical physics:

The best and latest mathematical methods to appear on the market have been used whenever possible. In doing this many an old and trusted favorite of the older generation as been forsaken, as I deemed it best not to hand dull and worn-out tools down to the next generation.

It is now 1993, and we are still teaching and using the old clumsy tools. As much as I would like to, I am not trying to bury vector calculus. Vector calculus will be longer lived than the typewriter keyboard or Fortran. My goal is to provide support for those students who want to learn the modern methods, but whose textbooks and teachers can provide no help.

In teaching graduate-level mechanics or electrodynamics, I have found that spending two weeks on differential forms speeds up the treatment enough to easily pay back the two weeks of investment. This book started out as a short set of notes for that purpose, and chapters three, four, and five read alone should still serve that purpose. I have struggled to keep this short. In *Mathematical TeX by Example*, Arvin Borde writes about his "original plan for a slim mathematical companion." His book ended up at 352 pages. My original goal was to keep this below 80 pages, keeping the treatment of the tools for manipulating differential forms complete, but including only tempting glimpses of the applications. These glimpses are intended to counter any impression that this is math for maths sake.

The immediate motivation for writing this came from reading *Icons and Symmetries* by Simon Altmann. There he shows how using a representation with the wrong symmetry seriously impeded the discovery of the laws of magnetism. I realized that differential forms carried his program even further. The electromagnetic theory given here has no right-hand rules in it at all. Electromagnetic theory is, after all, a right-left symmetric theory.

A proper regard for the symmetry of objects leads to a stratification of the various three-degree geometric objects into vectors, 1-forms, 2-forms, twisted 1-forms, and twisted 2-forms. Relations

between objects in different strata require operators with additional geometric structure. To relate E and D in electrodynamics requires the metric structure of either Euclidean space or a crystal lattice. A relation between E and H would also require an operator that distinguished between left and right handedness.

The two goals of this book are thus: provide the generalization of vector calculus that works in any dimension and with any metric, and secondly, develop representations of geometric objects with the correct symmetries.

Lets be frank. This is a polemic, arguing for the inclusion of the mathematics of the last three decades into traditional vector calculus. There is, quite correctly, considerable opposition to this, based on the idea: "If it ain't broke, don't fix it." Were this a mere change in notation, it would make no sense to change things. It is not a mere change in notation, however, but a basic change in the fundamental concepts. The new concepts are better for unarguable reasons: they correctly represent a larger symmetry group, and therefore correctly represent more features of the real world. We have an intuitive appreciation of symmetry, even before we have formalized the concept. A representation that violates the symmetry leaves a bad taste in your mouth. I would argue that mechanics, for example, is unnecessarily difficult precisely because the basic symmetries are not properly represented. This appears in the notoriously unintuitive distinction between centrifugal and centripetal force.

The present work is written for students who have had traditional calculus and vector calculus courses. Most of my examples are from physics, and a basic knowledge of mechanics, electrodynamics, and optics is assumed. Most junior science majors would meet these requirements if they are taking classical mechanics and electrodynamics at the same time. Originally I thought of this as the "Differential Forms Samizdat." With the publication of the excellent books by Edwards, Bressoud, and Bamberg and Sternberg, it is clear that differential forms have become mainstream. This should be a painless way to jump into that stream.

Version 2.* is still only a rough draft, but the organization has now been pretty much fixed. My original goal of keeping this at around 80 pages has been abandoned, both because of the realities of the material, and the realities of the publishing world.

Introduction

There are two main ideas that we will cover. One is the evolution in calculus away from the concept of a derivative toward the concept of the differential. This makes single-variable and multi-variable calculus fit smoothly together, and also paves the way for the ideas of catastrophe theory.

The second idea is to use both vectors and their duals, the differential forms. This allows calculus to be applied in spaces with no natural Euclidean metric, such as thermodynamics or relativity, and without the need to be in three dimensions. Remarkably, such a generalization turns out to be, even in three dimensions, computationally convenient.

CHAPTER ZERO. PROLEGOMENA

Here is some preliminary material that will be of interest to some readers.

I believe that some of the advantages of the differential forms approach come from the alignment of the notation, the concepts, and in particular the rendering with the natural symmetries of the objects. This natural alignment I feel makes them easy to think about. As an experiment, I include a section at the start in which purely abstract pattern puzzles are presented, somewhat in the spirit of those "apple:orange::five:x" problems on old I.Q. tests. Most people will be able to solve the puzzles without any understanding of what the puzzles are about. This, I claim, makes my point that the renderings are natural.

Note that I am going to continually distinguish between the concepts, their representation, and their rendering. For example, ~~their~~ ^{are} abstract vectors, their representation with triplets of numbers, and their rendering as arrows in space.

After the right-brain exercise, I am going to include some basic facts on linear algebra, multilinear algebra, affine algebra, and multi-affine algebra. Actually I would rather call these linear geometry, etc. but I follow the historical usage here. You may have taken a course on linear algebra. This is to repair the omissions of such a course, which now is typically only a course on matrix manipulation. The necessity for this has only slowly dawned on me, as the result of email with local mathematicians along the lines of:

When do you guys treat dual spaces in linear algebra?

We don't.

What! How can that be?

there

1. Patterns

One of the reasons why the use of differential forms is so easy is that the patterns involved are natural, and accord with the basic symmetry of the problem involved. Just for fun I have collected up some of the geometric patterns here, without any explanation. Just by looking at the examples you should be able to scope out the answers to many of these puzzles. After you have read the material, you might want to return to these puzzles.

In all of these puzzles you are looking for a rule that will be invariant under general linear transformations. You can use parallel lines in your constructions, and you can subdivide a line or stretch it by some factor. One thing you cannot do is to make any use of perpendicularity. Also, the signs of the results must follow from some natural rule from the signs given; you are not allowed to use any right-hand rule.

[This section is incomplete. I have not yet decided if it is a good idea.]

Addition

The most basic feature of all of these differential forms and their related vector objects is that they have a linear structure: you can scale them and add them together.

Consider the addition of vectors. Instead of looking at

$$c = a + b,$$


let us look for c such that

$$a + b + c = 0,$$

and expect the resulting pattern to be symmetric in the three vectors. The usual picture of this is, I claim, the only picture that you can draw that has the correct symmetries.

One of our geometric objects will have an icon in two dimensions consisting of a pair of parallel lines. These lines have no definite length, which makes them a bit unusual to think about, and makes it rather tricky to program a computer to make a reasonably aesthetic rendering.

Figure 1-1. The pattern for the addition of vectors.



One of the lines is special, and changing the sign is the same as changing the special line. Again I will claim that the addition pattern is uniquely determined by the symmetries of the above icons, and has to be that shown in the second figure.

Puzzles

You should be able to use the symmetries to guess the answers to these puzzles.

The addition rules that have been used above of course fix the scaling behavior of the objects.

These 1-forms and twisted 1-forms act on vectors and a twisted version of vectors. The basic pattern (again, this is all you can do given the invariances) is shown in the next figure.

Figure 1-2. The pattern for the addition of 1-forms. Note that the second line in each icon is rather short, to avoid cluttering up the pattern.




Figure 1-3. Add these two 1-forms.




Figure 1-4. Figure out an addition rule for these things which are like 1-forms but with a different parity. The parallel lines are all of indefinite length. I call these twisted 1-forms.




Figure 1-5. Puzzles requiring you to scale 1-forms and twisted 1-forms.




Figure 1-6. A 1-form acting on a vector, with one for you below it.





Figure 1-7. You should be able to guess the action of a twisted 1-form on the equivalent of a twisted vector.



2. Linear Algebra

Any set of objects that can be multiplied by numbers and added to each other is said to form a linear space. Real numbers themselves are an example as are the usual vectors. Less familiar examples would be apples and oranges (with a sharp knife) and the continuous functions on the unit interval. Row vectors, column vectors, and matrices are also all linear spaces.

Associated with any linear space is another linear space called its dual. Elements of the dual space are operators which map elements of the given linear space into the real numbers in a fashion that preserves the linear structure. That is, the map of the sum of two vectors is the sum of the maps. The most familiar duality is that between row vectors and column vectors, with the map provided by matrix multiplication. If your linear space is a shopping cart full of groceries, then the check-out clerk is a linear operator on that space. For the continuous functions on the unit interval, the map

$$f(x) \mapsto \int_0^1 x f(x) dx,$$

would be a linear operator on the f 's.

Surprisingly, the duals to the ordinary vectors are rarely pictured. If we render the vectors by the usual little arrows, line segments with arrows on one end, these correspond to displacements if we line up the tails of the vectors all at a single point which we call the origin. Linear operators on the vectors are the linear functions on the space. The value of the linear operator is just the value of the linear function at the head of the arrow.

We can make a drawing (which we will call a rendering) of these linear maps of vectors in two dimensions by drawing the contour lines corresponding to unit function value. If we wish to ignore the particular location of the origin, then we will also have to draw the contour line associated with the value zero. This then is the icon for these duals to vectors: pairs of parallel lines, with one of them singled out as the unit line, as opposed to the zero line. We do that by putting an arrowhead transversely on the unit line. In the preceding section on patterns there were many uses of this icon.

In three dimensions we can still render the dual space to vectors without undue mental strain. We just require a pair of parallel planes, with one of them singled out with a tick mark of some kind.

The most important fact about linear spaces is that they have a number called the dimension. The dual to a linear space has the same dimension.

3. Affine Algebra

Given a linear space, think of it as three dimensional vector space for concreteness, then the planes through the origin are subspaces that themselves have a linear structure. The planes that do not pass through the origin do not have a linear structure. For example, they do not have an origin. The weaker structure that they have is called affine structure.

What is defined on an affine space, as these are called, are sums of elements such that the sum of the coefficients is unity. If a and b are in an affine space, then

$$c = \frac{a + b}{2}$$

is defined. We would call it the midpoint between a and b . You can easily see that the planes in a linear space that do not pass through the origin are closed under the above operation.

Many of the linear objects that one talks about are really affine objects. We routinely call a function linear even though $f(0)$ doesn't vanish.

Computer Graphics

Affine structure plays a very important role in computer graphics. How, you might wonder, do we represent curves so that computers can deal with them. One approach would be to refer to the familiar catalog of lines, circles, ellipses, and so on. There is a better scheme, though. It was invented initially to describe automobile fenders; fortunately some enlightened auto companies had some mathematicians on their staffs. Here is the scheme to describe non-linear curves using nothing but affine structure. The PostScript type faces used to print these notes are all defined in this manner.

First we need to describe two different approaches to curves. Both have merits and faults. You can describe a circle in the plane by an equation

$$x^2 + y^2 = 1,$$

or by a map

$$\theta \mapsto (\cos \theta, \sin \theta).$$

The first form is called implicit, the second parametrized. Over the early decades of this century it was finally realized that the behavior of parametrized curves was more orderly and formed a better basis for an extension of calculus. This extension is called calculus on manifolds.

Affine structure allows us to describe straight lines parametrically using just two points, called control points. Call them a and b . The rule is

$$u \mapsto (1 - u)a + ub.$$

Because the coefficients in the sum add up to one, this is well defined in an affine space.

To describe a quadratic curve, use three control points, call them aa , ab , and bb . The rule is

$$u \mapsto (1 - u)((1 - u)aa + uab) + u((1 - u)ab + ubb).$$

This uses three repeats of the basic affine operation.

All of this is easily generalized to higher powers, more dimensions, and from curves to surfaces. Behind it all is the idea of a symmetric, multiaffine function, a function which takes several arguments and is affine in each one. For a quadratic curve we need a biaffine map, Φ . In terms of two numbers, a and b , the control points are

$$aa = \Phi(a, a)$$

$$ab = \Phi(a, b)$$

$$bb = \Phi(b, b)$$

and the curve is the map

$$u \mapsto \Phi((1 - u)a + ub, (1 - u)a + ub).$$

4. Computer Algebra Systems

I consider the aid of a computer algebra system essential to the work of any modern scientist. To be without one is like a carpenter without a Skill saw: underpowered.

The most obvious need for computer algebraic manipulation is the presence of antisymmetric multiplication in differential forms calculations. This is really the last blow to one's hopes of getting the signs right. Couple this with the expansion of intermediate steps in calculations to include sums of many terms, and you see that this is much more than just the need to correct the sign of the final answer.

There is a much deeper use of a CAS, however. The best way to learn a subject is to explain the subject to a naive listener. There is no possible listener more patient, more thorough, and more naive, than a computer algebra system. I continually am amazed at what more advanced pattern recognizing brain passes over without notice, but which stops a computer algebra system cold. When setting up the system that I use, I finally noticed the independence of much of the formalism over any specification of what the independent variables were. This is a creative sloppiness, I think, although I have not yet fully digested the idea. I just mention it here to show that you can get some high-level ideas from a CAS.

I personally find Mathematica the most congenial system to use. Any CAS will be the most complicated computer program that you run on your system. The learning curve is steep, and you will certainly have to put in your thousand hours on the road to wizard. Shortchanging your system to save a few dollars is not a good idea. The differential forms system that I use is available on my homepage.

<http://www.ucolick.org/~burke/home.html>

It is discussed in an appendix here.

CHAPTER ONE. CALCULUS

The calculus of differential forms is a replacement and improvement of vector calculus. We start by looking over differential calculus, trying to get the Big Picture. This requires us to take the modern viewpoint that concentrates on the differential, the local linear approximation, rather than the derivative, which is how most of you learned calculus. The derivative is a useful shorthand only for the differential in the case of maps $\mathbb{R} \rightarrow \mathbb{R}$. In general we will find conceptual simplification in the shift from an emphasis on the representation of objects to an emphasis on the intrinsic properties of the objects themselves. The differential of a function of several variables is the entire array of partial derivatives. While each partial derivative is easy to compute, their meaning is a collective property of the ensemble of them.

5. Vectors

One object that we can all agree on is the displacement vector. To represent operations like: go north three feet and then west for five feet, the vector and its icon, the line segment with an arrowhead are useful and incorporate the correct symmetries.

An example of a symmetry properly represented here is the reflection of a vector in a line containing the vector. The arrow icon is invariant under this operation. So too the most primitive notion of a displacement. Refinements of the displacement concept will break this symmetry. You might be driving from point A to point B, and in which case you had better drive on the correct side of the road. Driving is not invariant under such reflections.

Now I will argue that our thinking and iconography should properly reflect the symmetries of the situation. But there are significant contrary examples to keep in mind. We use a pencil line both to represent a line like a path across a field, or a crack in pottery, which has the above reflection symmetry, and also to represent an edge, like the boundary of a plank or the edge of a cliff. There are famous visual puns that exploit this ambiguity, and every mechanical drawing student has to laboriously learn how to disambiguate representations into proper concrete objects. We will say more about this question of orienting boundaries when we discuss the divergence theorem.

The representation of positions and displacements on a map has more symmetry than the real world. We are all familiar with the distortions inevitable on large-scale maps of the earth. Right angles on the map need not represent right-angles on the ground. The maps are distorted. In a small region we can describe this distortion as a shear: what happens when you push a deck of cards sideways. Our representation of displacement vectors works perfectly well on a sheared map. Just as on a Euclidean map, if displacements A and B add up to C, then they also add that way on any sheared version of the map.

Your first reaction might be that this is a bad thing. A representation with too much invariance could be misleading. In fact, we will do just the opposite, and ensure that all of our representations have this more general symmetry: invariance under general linear transformations. Why? It comes down to calculus, and the central role of

linear transformations in calculus.

This invariance under general linear transformations will show up in our icons through the elimination of the concept of perpendicularity. Parallelism is ok, but not perpendicularity. Of course one never follows a principle too closely. The arrowheads on our vector icons involve perpendicularity in their construction. To be general-linear-transformation correct, we should use a more general arrowhead, subjecting the Euclidean special case to a different general linear transformation each time we draw one.



In the context of vectors there is an important relation known as duality. Given a set of vectors, that is, objects which can be scaled and added, one naturally comes to think about linear operators. The simplest linear operators are functions which take in a vector and produce a number.

Example: Let us think of shopping carts full of groceries as vectors. The addition operation is to dump everything from two carts into one. To scale by a factor of two, just double the quantities of everything. A linear operator on these vectors would be, for example, the check out clerk. Every vector (shopping cart) is assigned a price.

[Certain features of real markets are not compatible with linear structure. For example, you are not usually allowed to purchase half a tube of toothpaste, for example.]

For reasons that escape me duality is considered too difficult to introduce, and linear algebra classes instead start with the idea of matrices. These are linear operators whose values are not numbers, but other vector spaces. Elementary linear algebra in the United States usually means beginning matrix algebra.

6. Calculus of Curves

The calculus of curves is the process by which we assign an instantaneous velocity vector to a particle moving nonuniformly through space. Geometrically, the velocity vector describes the path the particle would have taken if it had continued on with uniform motion.

The tangent vector is defined by the following limit process. This limit process should do two things: isolate the behavior of the particle in the immediate neighborhood of the point in question, and find a finite representation of that behavior. If we are given the motion of the particle in the (x, y) plane by two functions $x(t)$ and $y(t)$, and if we want the velocity vector at $t = t_0$, then we isolate the behavior at the point t_0 by successively considering the piece of the curve between t_0 and $t_0 + 1$, then between t_0 and $t_0 + 1/2$, then between t_0 and $t_0 + 1/4$, and so on. To get a finite representation, at each stage we expand our map of the (x, y) plane by a factor of two, and draw the straight line connecting the endpoints of the segment of the curve.




Figure 6-1. A few steps in the construction of the tangent vector to a curve.

This process can be done at any point on the path of the particle.

The result of doing it everywhere is too messy to really draw, but the next figure shows the idea.





Figure 6-2. The collection of local approximations to a curve.

A practical example of this comes up in mechanics, where a particle free from external forces moves in uniform motion. I had just built a buffing wheel, a cloth disc covered with polishing compound rotating at high speed. As I was finishing building this I thought, "Oh no, the polishing compound is going to fly off and hit me in the face."


Of course that doesn't happen. Compound that frees itself from the wheel flies off in whatever direction it was going. This is exactly the velocity vector construction given above.

Figure 6-3. Unphysical but intuitive behaviour of a buffing wheel.



Y

Figure 6-4. Correct physical behaviour of a buffing wheel.



Y

With a PhD in physics, I was a bit embarrassed by this wrong intuition.




Figure 7-1. The differential of the function $y = x^2$ at $y = 1$.

Only for single-variable calculus is it possible to package up all of the differential behavior in an object of the same type, here a function of one variable. This makes single-variable calculus misleading, and why I started with vector calculus.

Once we have the notion of the differential of a function, we can use it, usually in the form of the Taylor's series

$$f(x) = f(x_0) + (x - x_0) f'(x_0) + \dots$$

to find representations of other differentials.

The differential of the parametrized curve

$$s \mapsto (x(s), y(s))$$

is the vector with components

$$x'(s), \quad y'(s),$$

which we write

$$x'(s) \frac{\partial}{\partial x} + y'(s) \frac{\partial}{\partial y}.$$

Let me carefully parse the above mathematical sentence. The expression $x'(s)$ is a function

$$x'(s) : \mathbf{R} \rightarrow \mathbf{R}; s \mapsto \frac{dx}{ds}(s),$$

and for any value of s this is a number. The symbol $(\partial/\partial x)$ denotes the x -basis vector. At the point (x_0, y_0) this is the tangent to the parametrized curve

$$s \mapsto (s + x_0, y_0),$$

and since this is already a linear map, the tangent to the curve is the curve itself.

What does it mean to multiply a linear curve by some number? We have to answer this question if we are to claim that the linear curves have linear structure. The answer: you take the same curve and just travel faster along it. Thus

$$r \frac{\partial}{\partial x} = s \mapsto (rs + x_0, y_0).$$

And what does it mean to add two curves? Again this is straightforward:

$$\begin{aligned} s \mapsto (\alpha s + x_0, \beta s + y_0) + s \mapsto (\gamma s + x_0, \delta s + y_0) = \\ s \mapsto ((\alpha + \gamma)s + x_0, (\beta + \delta)s + y_0), \end{aligned}$$

and in pictures




Figure 7-2. Addition of two linear parametrized curves.

8. Linear Transformations

The structures discussed in this chapter, two dual linear spaces, have an important invariance. The operations of addition, scaling, and evaluation all commute with linear transformations. In this section, I will discuss linear transformations more carefully.

Linear transformations can be represented by matrices acting on a vector space by the usual multiplication law. Closely related transformations are the affine transformations, which also allow you to translate the spaces, and projective transformations, which preserve the straight lines but not the linear structure. Projective transformations are familiar from photography. They do not preserve parallel lines, for example.

In the plane there are three basic types of linear transformations. Dilations expand the space uniformly in both directions. This is the only rotationally symmetric linear transformation.




Figure 8-1. Dilations and rotations of the plane.

Next in complexity are the rotations. A pure rotation would be represented by an antisymmetric matrix. Finally there are the shears. A pure shear expands one direction and contracts an orthogonal di-

rection so as to preserve the area. There are two independent shears separated by a 45° rotation. See figure 2.




Figure 8-2. Two pure shears in the plane.

I will often refer to a shear which is a mixture of these pure types, in which a vector is transformed according to


$$x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} \mapsto (x + \alpha y) \frac{\partial}{\partial x} + y \frac{\partial}{\partial y}.$$

The x -axis is invariant and all other vectors are slid horizontally an amount depending upon their vertical position. This is what happens when a deck of playing cards is pushed sideways. It is a combination of a pure shear and a rotation.

One thing that does not commute with linear transformations is the concept of orthogonality. Since I used orthogonality above in my description of the types of shears, that description is in fact not invariant under linear transformations. It presupposes a given Euclidean geometry framework. Even the concept of a symmetric matrix is not invariant.

We are going to separate the operations that are not invariant under linear transformations, like Euclidean geometry, and put them in with explicit operators. Thus, while we will not be able to talk of just perpendicular, we will be able to say that using the metric \mathcal{E} , two lines are perpendicular. It is by no means obvious at this point that it is worth all the bother of doing this. Trust me.

Figure 8-3. A purely horizontal shear.



9. The Differential in Optics

Sometimes an idea can be appreciated better when it is seen in a more realistic, complex situation. You might find the application of the idea of a differential to optics useful. If not, then this section can be skipped without loss.




Figure 9-1. A simple optical situation.

In the above figure I show an optical set up consisting of light rays coming from a point source, passing through a lens with spherical surfaces, bouncing off of a spherical mirror, going through the lens again, and finally forming an imperfect image of the source. The image is imperfect because a simple lens is not able to form a perfect image for light rays making finite angles with the axis. We can define a limit process for this situation that will isolate from this nonlinear situation, the special behaviour of light rays that are close to the axis and nearly parallel to it. The limit process will replace the original rays with rays from a source half as far away from the axis, and reduce all of the slope of the rays by the same factor. Finally, we rescale the

drawing by magnifying the transverse direction by a factor of two. If this process goes to a well defined limit, then we have the analog of the differential in calculus. Here it does, and the limit is called Gaussian optics, or the paraxial approximation.




Figure 9-2. Two steps in the paraxial limit.

In the above figure I show the original rays and two steps in the paraxial limit all drawn on the same figure. In the next figure I show the second stage in the process appropriately rescaled. Next I show the limit itself. In the limit, the rays do form a perfect image of the source. Note how the lens surfaces appear flat in the limit. You might think about why the reflection does not make equal angles with the surface normal.




Figure 9-3. The second stage in the paraxial limit, and the limit itself.

CHAPTER TWO. GRADIENTS AND FLOW

Here we cover some examples of differential forms. The presentation is primarily graphical. I want you to get the ideas down before we get swamped in a lot of unfamiliar, but necessary, notation. The basis for all of the calculus of differential forms is the idea of the gradient of a function. The point to be made here is that this is not, as you were taught in vector calculus, a vector. That assumption restricts you to working with Euclidean geometry. Here we remove that restriction, and find in the process that things become not just more general, but simpler. We end up with examples of more complicated differential forms that come up in the discussion of the flow of something. This will be essential when we come to discuss conservation laws.

10. The Differential of a Function

Suppose we have a function of several variables. I use two variables here just so I can draw the results. We can represent its differential as follows. We want to keep expanding the view in order to zoom in on the local behavior of the function. And we need to find an icon to represent the differential that maintains its size and shape as we expand. This differential and its icon will represent the gradient of the function, and it is called a 1-form.

If we represent functions by their contour maps, then we can represent the differential by drawing contour lines, not with unit spacing, but with a spacing of ϵ when the map has been blown up by a factor of $1/\epsilon$.




Figure 10-1. Contour maps of the function $y^2 + (x-1)^2$, and its affine approximation at $0, 0, 1 - 2x$.

If the function is smooth, then in this process the contour lines straighten out and their spacing becomes uniform. This is the contour map of a plane. Just as the differential of a function of one variable was the best fitting tangent line on its graph, the differential of a function of several variables is a plane tangent to the graph of the function.

To represent the gradient it is sufficient to just draw two of the contour lines, and to indicate somehow which line is uphill. The icon in common use is shown in the next figure.



Fig 10-2. The icon for the above gradient.

From a Taylor's series expansion we can write a function locally as

$$f(x, y) = f(x_0, y_0) + (x - x_0) \frac{\partial f}{\partial x}(x_0, y_0) + (y - y_0) \frac{\partial f}{\partial y}(x_0, y_0) + \dots$$

and the partial derivatives represent the information in the differential. The differential is just the collection of all of partial derivatives assembled into a coherent geometric whole.

The coordinates themselves are functions, and we can take their gradients as a basis. If we call the differential of the x coordinate dx , then the differential of a function looks just like the chain rule

$$\frac{\partial f}{\partial x} dx + \frac{\partial f}{\partial y} dy.$$


Note here that dx and dy are not infinitesimals, but perfectly finite objects, the gradients of the functions x and y .

An example of the physical measurement of a gradient came up when I wanted to survey the surface of my pool table (the home team advantage). I borrowed an accurate level from the machine shop and using a shim of standard thickness, I could construct the local gradient at any point.

Figure 10-3. A mechanical equivalent of a 1-form.



Figure 10-4. Drawing the gradient 1-form.



I just moved the shim out along the level until the bubble was centered. If I did this in two different directions, then I would determine the position and direction of the contour line, and thus the gradient.

A mental picture of this field of 1-forms easily let me visualize the force that would disturb the straight line motion of slow shots. Straight pool anyone?

11. The Algebra of 1-forms

The 1-forms that we developed in the last section represent the gradient of a function. They have the algebraic properties of a linear vector space.

A 1-form that is doubled, that is, multiplied by two, represents a steeper gradient and has its contour lines closer together. 1-forms shrink as they increase in magnitude. The zero 1-form has its contour lines infinitely far apart, and with uncertain direction.




Figure 11-1. Doubling a 1-form.

The addition of 1-forms is just like the addition of functions. In terms of our icons addition is given by the rule in the next figures.

I prefer the symmetrical view of it given in the second figure.

Taking the negative of a 1-form just means putting the tick on the other contour line.

The linear vector space of 1-forms is dual to the linear vector space of vectors. Given a vector and a 1-form, this means that we




Figure 11-2. Addition of two 1-forms to give a third: $a + b = c$.




Figure 11-3. Addition in the symmetrical form $a + b + c = 0$.

can find a number, their inner product. For our icons we just count the net number of contour lines crossed by the vector.

Our basis for vectors and our basis for 1-forms are related so that




Figure 11-4. Duality between vectors and 1-forms: $v \cdot a = 2, w \cdot b = -3$.

we have

$$\begin{aligned} \frac{\partial}{\partial x} \cdot dx &= 1, \\ \frac{\partial}{\partial x} \cdot dy &= 0, \\ \frac{\partial}{\partial y} \cdot dx &= 0, \\ \frac{\partial}{\partial y} \cdot dy &= 1. \end{aligned}$$

So far we have been talking only about local linear approximations, and everything so far is invariant under general linear transformations. If you shear, for example, the law of addition, you get another valid picture of the addition law.

Similarly, the basis vectors and 1-forms in oblique coordinates also satisfy the above duality relations.

Figure 11-5. Basis vectors and 1-forms.




Figure 11-6. A transformation of the addition law.





Figure 11-7. Basis vectors and 1-forms in oblique coordinates.



12. 1-forms in Mechanics

We now come to the first surprise. Force is not a vector, but a 1-form. The most direct way to see this is to think of the work done by a force. Force is the operator that takes in a displacement, a vector, and tells you how much work was done. This makes forces dual to vectors, i.e. 1-forms.

There is a temptation here to bring in Euclidean geometry. Our gradient 1-form could be turned into a gradient vector. In rectangular coordinates just replace dx by $\partial/\partial x$, and so on. When dealing with the motion of a particle, the Euclidean geometry comes in with the isotropic behavior of mass. Only experience shows that it is more convenient and useful to resist this and to stay with the less familiar but more general representation.

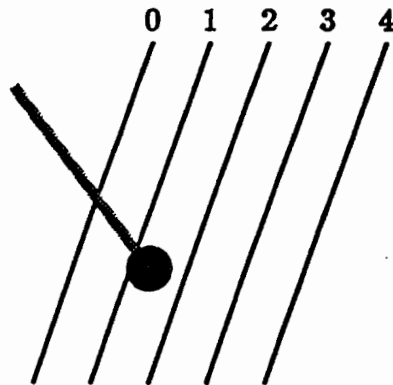


Figure 12-1. Measuring a generalized coordinate.

The notion that is free of Euclidean geometry comes into its own in situations involving generalized coordinates that are not necessarily Euclidean. Let me describe one such application, the reciprocity relations that apply when a situation of static equilibrium is disturbed by small forces. Consider any Rube Goldberg arrangement of springs

and rods. Focus your attention on two parts of the system that are free to move. The position of any part can be described by a generalized coordinate, and you should imagine that a measurement of this generalized coordinate is done by locating a point on the part and see where it falls on a grid of parallel lines.

The part may have several degrees of freedom, and need several such coordinates to fully determine its position. Here we may only be measuring one such coordinate.

Now you can also apply forces to these parts, and when you do the system will come to a new equilibrium. The reciprocity theorem is nothing more than the expansion of the potential energy surface about its minimum, and the equality of the two mixed partial derivatives, yet it is still a surprising and little used theorem. The reciprocity theorem states that if a force Q is applied to one degree of freedom, and it produces a displacement d in some other degree of freedom, then that same force Q applied to the second degree of freedom, produces the same displacement d in the first. This despite the fact that the rest of the system might be disposed quite differently in the two cases, as will show up in the example below.

What does it mean to apply a force to a degree of freedom? Why, just align the force 1-form with the coordinate lines. And how to measure the strength of the force? Again, just use the coordinate lines. This is precisely the definition of generalized force that is used in Lagrange's equations.

Look at a double pendulum with equal weights and equal length rods. The same horizontal force is applied first at the top joint, and then at the bottom joint. The deflection of the top joint when the force is applied to the bottom mass will equal the deflection of the bottom mass when the force is applied to the top joint. To see that this follows from the reciprocity theorem, use generalized coordinates as sketched in the next figure.

I have given them equal linear spacing. When the force is applied to the top joint, both joints have the same displacement. When the force is applied to the bottom joint, it deflects through twice the angle and three times the displacement, but the top joint deflects the same as before. It does this since, by the reciprocity theorem, its displacement when the force is applied below is the same as the deflection of the lower point when the force is applied above, and in that case the deflections were equal.

Figure 12-2. A double pendulum as an example.





Figure 12-3. Generalized coordinates for the double pendulum.



After all these years, I still find this theorem remarkable. Here, however, note how natural its statement and application is when the force is treated as a 1-form.

Just as we can have a function, a rule which assigns a different number to every point, so too we can have a 1-form field, a rule which tells us how to construct a different 1-form at every point. The gravitational force field on my pool table is an example of this.

Given such a force field, one can ask how much work is done moving along a particular path. We want to calculate a line integral. Now 1-forms were set up precisely for this, and sometimes they are even defined as those geometric objects that can be integrated along oriented paths.

To define the value of a line integral, imagine breaking it up into little pieces. Treat the pieces as vectors, and use the 1-form field to turn the vectors into numbers. Then add up all the numbers. Do this with finer and finer pieces until the answer doesn't change. In practice, you do the usual integral calculus manipulations, which amount to the same thing.

From now on you should view the integrand of a line integral as a 1-form. No surprise to you, I bet, that 2-forms will be integrated over surfaces, and so on.

Not every 1-form field is the gradient of a function. Look at

$$\omega = (1 + x) dy.$$




Figure 12-4. Typical line integral.

If we integrate around the edge of the unit square, we find a net value of +1, not the zero that we would get integrating the gradient of a function. We will later find out how to decide whether a given 1-form is the gradient of a function or not. This will give us the replacement for the usual curl operator.

13. Flow in Two Dimensions

Now turn to the construction of a proper geometric representation for flow. In this section I will work only in the plane. Suppose we have something measurable flowing in the plane. To describe this flow locally we want an operator that will answer the question: given a little piece of a line σ , how much of the measurable stuff flows across it? In the calculus limit, of course; thus we expect this to be a linear operator.

Your first guess for the little piece of a line is to represent it as a vector. But wait. A vector is invariant under reflection, and this little piece of a boundary has an inside and an outside, and is not invariant under reflection. An icon with the proper reflection symmetry is shown in the next figure.

The icon has a definite length, but the arrowhead is across the line rather than along it.

Can we construct a vector space of such objects? Scaling is obvious. What about addition? We can model a definition on the addition of vectors, but with the natural sign convention of the next figure.

What about dual objects, the operators on such icons? You can guess that they will be like 1-forms, but the tick marks will be placed in some other manner. This is shown in the next figure.

Such an object represents the flux of measurable stuff according to the following story. The lines of the icon are aligned so that none of the stuff flows across the lines. The spacing of the lines is such that a unit amount of the stuff is flowing between the lines. The arrowheads are located along the direction of flow of the stuff.

Contrast this with the picture of a gradient in Figure 9-5.

The only problem, really, is what to call these variations on the 1-form theme. There is no agreement among mathematicians. Some call them W -forms, after Hermann Weyl, whose keen appreciation of symmetry led him to these objects. Some call them forms of odd type,




Fig 13-1. Icon for edge line.




Fig 13-3. The representation of flux.

Figure 13-2. The picture of $c = a + b$.






Figure 13-4. The picture of the spacing condition.



but an odd 2-form is apt to confuse. The most commonly used name is twisted 1-form. This is derogatory in exactly the same way that saying Mrs. John Smith relegates the poor woman to second class status. The twisted forms are just as fundamental as the untwisted ones. But "twisted" is as good as we can do, in my opinion.

Figure 13-5. The representation of a gradient.



14. Density

Related to the idea of flow, and apparently simpler, is the idea of density. We pursue the now familiar theme: find a pair of linear spaces dual to each other to represent the idea. We need to represent both the idea of density, and also tiny pieces of surface. The density operator will take a tiny piece of surface and tell you how much measurable stuff is on that piece of surface.

A reasonable guess for the representation of a tiny piece of surface would be to use a pair of vectors, one along each edge of the surface. You might slip into Euclidean thinking at this point, and say orthogonal vectors, but remember that is a no-no. The density will be a linear operator on the pair of vectors.

To use this representation we need an important property of area. This property is what makes the area idea really area. If you shear a parallelogram, the area does not change. We need not know how to measure area. That requires metric notions. We can compare two areas, however, just using linearity and the shear invariance.




Figure 14-1. Equal areas under shear.




Figure 14-2. Scaling of areas.

To compare two areas, use shears to line up the sides of the two parallelograms. Then just compare the relative lengths of the side and multiply the numbers. This uses the linearity that comes from the calculus limit.

The same manipulations involved in the comparison of two areas allow us to add two areas of different shapes.

To picture density, one can draw a parallelogram of a size to include a unit amount of material. The usual calculus limit is going on here, of course. The shape of the parallelogram doesn't matter, only the area that it encloses.

Now that we know how to compare areas we can define the action of a density on a piece of surface. Merely compare the given area with that enclosing unit amount of stuff.

The density parallelograms must include a sign to tell whether it contains positive or negative stuff. This allows us to find the laws for scaling and adding densities. Just as with 1-forms, these behave in a manner opposite to the scaling and addition of areas.

It is also possible to have densities with a different handedness. These would change sign under a reflection. These densities operate on areas which have either a clockwise or counterclockwise orientation

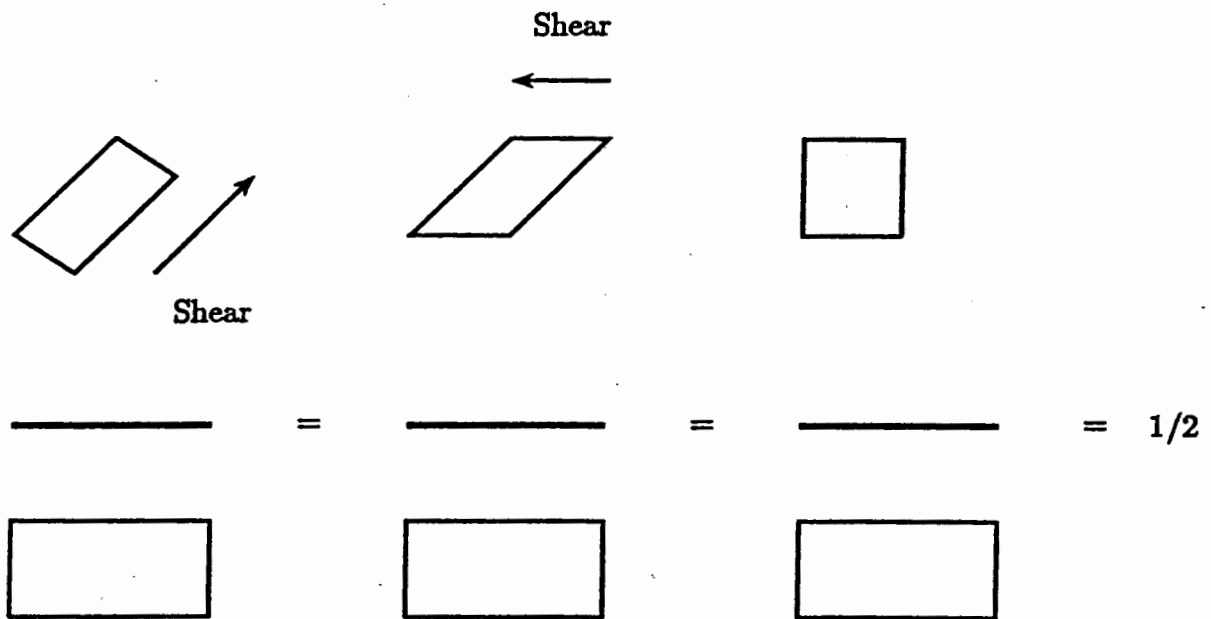


Figure 14-3. Comparing two areas. The ratio here is $1/2$.

specified. These might seem to be weird and unnecessary objects, but that is not so. In fact these densities are the natural extensions of the 1-forms, and are called 2-forms. The other densities are twisted 2-forms. This flies in the face of nearly everybody's intuitive sense of twisted, which is one reason that a better name for them would be desirable.

Even though it will be easier to define twisted 2-forms, and they relate more naturally to 1-forms, in practice most of the densities that you will meet will be twisted. Not all, however, say two thirds of them.

If you think of the density as an operator on pairs of vectors, then scaling and linearity demand that it be a linear operator on each vector. The invariance under shears gives us a further requirement.




Figure 14-4. Addition of two areas.

We must have the area spanned by two vectors a and b to be the same as that spanned by $a + \lambda b$ and b for all values of the number λ . Using linearity, this says that the area of b and b vanishes, which makes sense. Apply this to the area spanned by $a + b$ with itself, and you conclude that the area spanned by a and b is the negative of that spanned by b and a . The invariance under shears makes the bilinear operator antisymmetric in its two arguments. Thus the operator needs to know not just the two vectors, but also which one comes first and which second. Thus the area must have a handedness specified, a little arrow if you like, going from the first vector to the second. This handed area is operated on by 2-forms. Most area is not handed. That is why we will need to define twisted 2-forms to operate on these.

We capture this algebraic structure by defining an antisymmetric multiplication on 1-forms. This is similar to the cross product in




Figure 14-5. A density and an area composing to give a number, the amount of stuff in the area.

ordinary vector analysis. Unlike the vector cross product, it does not depend on any metric, and works as well in spacetime as in Euclidean geometry. Everyone uses the wedge \wedge for this multiplication, when they bother to include a specific symbol for it,

$$\alpha \wedge \beta = -\beta \wedge \alpha.$$

This object is defined to be the operator which acts on pairs of vectors according to

$$\alpha \wedge \beta \cdot (a, b) = (\alpha \cdot a)(\beta \cdot b) - (\alpha \cdot b)(\beta \cdot a).$$

Since $(\alpha \cdot a)$ and $(\beta \cdot b)$ are numbers, the multiplication on the right hand side is well defined. Because α and β are required to be 1-forms it has the required linearity on a and b . Note that we have explicitly made it antisymmetric,


$$\alpha \wedge \beta = -\beta \wedge \alpha.$$

When the 1-forms are basis forms it is customary to leave out the wedge and write $dx dy$ rather than $dx \wedge dy$. This makes integrals of 2-forms look familiar.

$$\frac{1}{2} \times \square = \text{rectangle}$$

$$\text{rectangle} + \square = \text{rectangle}$$

Figure 14-6. Scaling and addition of densities.



v

Figure 14-7. Addition of densities with clockwise signs.

15. Flow in Three Dimensions

The descriptions using differential forms are very similar in different dimensions. Often you can describe the two dimensional case, which is easy to draw, and then just reinterpret the symbols to cover three or higher dimensions. This smooth march through the dimensions is a valuable feature of differential forms. You do have to get past some peculiarities of low dimensions, such as the accident that $n - 1 = 1$ in two dimensions.

Gradient

The gradient in three dimensions is easy to render. A function has not contour lines but contour surfaces. A linear function has parallel, equidistant contour planes. We can pick two of these to render the gradient. We can represent a gradient by using the gradients of the three coordinate functions as a basis.




Figure 15-1. A gradient in three dimensions.

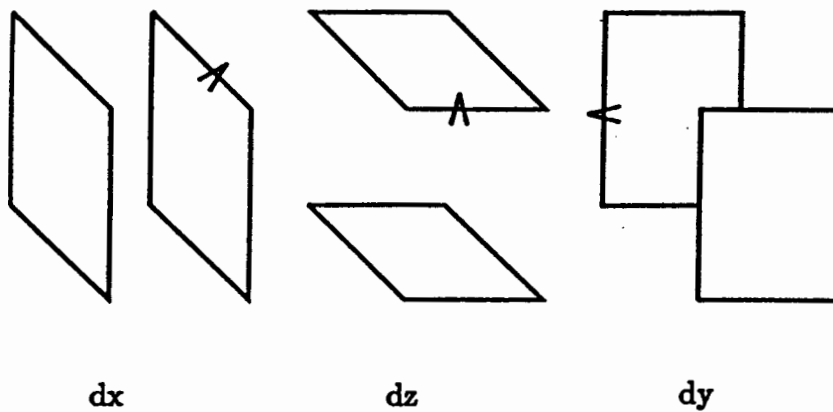


Figure 15-2. A coordinate-derived set of basis forms in three dimensions.

Flux

To study flow in three dimensions one needs the idea of a shapeless area covered in the last section. Renderings of forms are difficult to interpret at first because they involve lines and planes of indefinite extent, second because the figures often have no definite shape. The geometric figure for rendering flow is a prism with a parallelogram for a cross section. The length of the prism is indefinite, as is the shape of the cross section. Three equivalent renderings are sketched in the next figure.

The prism is chosen so that nothing flows through the walls, and a unit amount of stuff flows through the prism. Given a definite piece of surface, the flux prism answers the question: how much stuff flows through the surface. Think of the prism as a cookie cutter. You want




Figure 15-3. Three equivalent flux renderings.

to know how many cookies you can cut out of a given piece of dough. Of course here we are thinking of reshaping the cookie cutter rather than the dough as you finish up the fragments.

The prism needs an arrow directed along it to indicate the direction of flow, and the surface that the prism acts on needs an inside and an outside indicated. We want the signed flux across the surface.

This flux prism is one type of differential form, called a 2-form because it acts on two dimensional surfaces. In general, in n dimensions flux will be represented by an $(n - 1)$ form.

16. Drawing Vectors: The Eight Icons

Drawing vectors? How can there be anything to say about this? What is there beyond drawing a line with an arrowhead on one end?

Well, vectors come in eight different geometric types, and each one has a different rendering, a different icon if you will. Only one of these is the line with an arrowhead on it. All of these come up in electrodynamics.




Figure 16-1. The eight icons for vectors. All of these would be conventionally denoted by \hat{x} . The electrodynamic fields \mathbf{B} , \mathbf{H} , \mathbf{E} , and \mathbf{D} would be represented by the icons in the bottom row, in that order.

The objects in all eight classes have three degrees of freedom, and these three components transform in the same way under proper rotations. This justifies calling them all vectors. These objects are stratified into classes by their behavior under other transformations, inversions and dilations. You are probably already familiar with the stratification of vectors and pseudovectors under inversion. By ignoring dilations that stratification missed much of the geometry of the situation, but it was a start.

A reflection about the center line in figure one exchanges an object with one of opposite behavior under inversions. A reflection about

a horizontal line connects objects with opposite behavior under dilations. These objects are connected by the usual vector space duality. For example, a field intensity can be integrated along a line. Small pieces of lines look like displacement vectors. Flux densities can be integrated over surfaces, and the dual object looks like a little piece of surface, with a direction through the surface specified. Objects within the right-hand and left-hand squares are connected by the operation called Hodge-duality. Unlike the other two operations, this one involves the metric. In electrodynamics, this is signalled by the presence of the operators ϵ and μ .

Two Dimensional Flow

Let us start by studying the geometrically simpler case of the flow of some conserved stuff (you might think of water) in two dimensions. How would you measure the flow rate? You could find the flux by lowering a container into the flow, and collecting the stuff for some standard time interval.

If the flow were smoothly varying, then over a small enough region, the amount collected would be linear in the length of the region over which the stuff was collected. If you measure the flux for two regions which have different directions, then you can figure out the flux across any other region as well. Just construct a triangular region, with one side the region in question, and the other two parallel to the two measured lines. Using linearity we can find out how much flows in each of these two sides, and use conservation to find out how much flows out the other side. In practice we would take the two regions of measurement parallel to the coordinate lines.

The usual vector representation of this flux comes from a representation of the line across which the flow is to be measured by a vector perpendicular to the line, pointed across the surface in the direction of measurement, and with a length proportional to the length of the line. Call this vector S . The flux vector is the vector J such that the flux through a surface is given by the dot product $J \cdot S$.

Now I show how to draw a picture of this flux vector that we will see is more geometrically correct. The icon will be a pair of parallel lines forming a rectangle, with a direction specified along the lines. Such a picture can be specified by two perpendicular vectors, a and b . These vectors must be related to J by the following equations. The





Figure 16-2. Calculating flux across a line in terms of two measurements.


Figure 16-3. Drawing the flux vector corresponding to two orthogonal flow measurements across unit intervals.



vector \mathbf{b} does not have a unique length.

$$\begin{aligned} \mathbf{a} \cdot \mathbf{b} &= 0, \\ \mathbf{a} \cdot \mathbf{J} &= 0, \\ \mathbf{b} \cdot \mathbf{J} &> 0, \\ |\mathbf{a}| |\mathbf{J}| &= 1. \end{aligned}$$

Figure 16-4. The geometrically correct icon for the flux vector, and its representation by two ordinary vectors.



Why do I consider this icon more geometrically natural? Look at the next figure, which shows two representations of a flow question. In the first case, you can immediately see that there are two units flowing across the given line. In the second, you can't do anything until I also tell you what a unit length vector is. Now we are going to do the same geometrical constructions for three dimensions.

Three Dimensional Vector Icons

The most familiar objects are in the class containing displacement vectors and velocities. Under an inversion, taking

$$x \mapsto -x,$$

$$y \mapsto -y,$$

$$z \mapsto -z.$$

these displacement vectors change sign. Under a dilation

$$x \mapsto \alpha x,$$

$$y \mapsto \alpha y,$$

$$z \mapsto \alpha z.$$




Figure 16-5. Two representations of the flux vector used to compute the flux across a line segment.

the magnitude of the vector scales by the same factor of α . This is the same scaling behavior as the arrow icon.

The next class of vectors are what I call field intensity vectors, such as the electric field in electrodynamics, and gradients in general. Under the above dilation the magnitude of a gradient scales like $1/\alpha$. Stretching a function decreases its gradient. The icon for a field intensity vector that has this scaling behavior is to take a pair of parallel planes. The shape and extent of the planes is irrelevant. Only the spacing matters. Furthermore, there is a mark on one of the planes to denote the "uphill" direction.

If we describe the icon by three ordinary vectors, \mathbf{a} , \mathbf{b} , \mathbf{c} , then the icon is related to the usual gradient vector \mathbf{w} by the relations


$$\begin{aligned}\mathbf{w} \cdot \mathbf{c} &= 1, \\ \mathbf{w} \times (\mathbf{a} \times \mathbf{b}) &= 0.\end{aligned}$$

The lengths of \mathbf{a} and \mathbf{b} are irrelevant, as are changes of shape such as

$$\mathbf{a} \mapsto \mathbf{a} + \lambda \mathbf{b}.$$

The next class of vectors are what I call flux density vectors. The prototypes in electrodynamics are the \mathbf{D} field and the current density. These measure how much of a given stuff passes through some particular surface. If we dilate by a factor of α , the flux changes

Figure 16-6. The field intensity icon.



by $1/\alpha^2$. The icon with this invariance is a rectangular prism of indefinite length and shape, but with a specific direction and cross-sectional area. If we represent the icon by three vectors, the relation between the usual vector w and the icon is given by

$$\begin{aligned} w \cdot (a \times b) &= \text{sgn}(a \cdot b \times c), \\ w \times c &= 0. \end{aligned}$$

The term involving the sign of the triple product is there to make the direction of c agree with the direction of w , and to make it independent of an interchange of a and b .


If we leave out the factor involving the sign of the triple product, then we have a similar icon, but with a direction of circulation specified around the prism, rather than a direction along the prism. This is the class of vector fields that the magnetic field B belongs to.

Finally, if we add a sign term to the field intensity vector, we find the icon for the H field in electrodynamics. I call this a twisted field intensity, and the B field a twisted flux density. These are shown in figure one.

Discussion

There are five reasons to use the calculus of differential forms rather than the usual vector calculus: ease of computations, need to go to

Figure 16-7. The flux density icon.



dimension other than three, need to use a metric other than Euclidean, need to use oblique coordinates, and desire to draw geometrically correct pictures. By geometrically correct, I mean that the pictures have the same invariances as the concepts themselves. This work here has addressed only the last point.

Mathematicians call field intensity vectors 1-forms, and twisted flux density vectors 2-forms. The flux density vectors are called twisted 2-forms, and the twisted field intensities are called twisted 1-forms. If you want to work in a manner that is not specific to three dimensions, it makes more sense to call something twisted when it is the sign of the triple product in its definition.

CHAPTER THREE. DIFFERENTIAL FORMS

The preceding two chapters were motivational and heuristic. Here we get down to the actual rules for manipulating differential forms. The following three chapters can be read on their own for those with less interest in the pictorial rendering of differential forms.

In this chapter we cover the basic algebraic operations. Here we will find the replacement for the cross product that is superior on several counts. It will work in any number of dimensions. It is not dependent upon Euclidean geometry, and so can be used also in special relativity. Finally, it is associative, and so we can dispense with a lot of parentheses and special rules. In the next chapter we cover the analysis of forms, how we replace the operations div , grad , and curl with a single more general operation.

17. Ordinary Differential Forms

Differential forms are polynomials in the algebra that one gets by defining an antisymmetric multiplication between 1-forms. This multiplication is denoted by the \wedge operator, read "wedge". Often \wedge is omitted between basis forms, and we will write

$$dx dy \text{ to mean } dx \wedge dy.$$

This multiplication is associative

$$dx(dy dz) = (dx dy)dz = dx dy dz,$$

and distributes over addition

$$dx(dy + dz) = dx dy + dx dz.$$

We will define our basic operations on monomials of basis forms and use this linearity to pass to the general case.

In what follows the following shorthand will be convenient. I will write the general basis 1-form as dq , and use greek, $d\alpha$, for the monomial which is the product of any number of basis 1-forms.

Contraction

Differential forms are operators which act on ordered sets of vectors. That is, the operator on (a,b,c) may not be the same as on (b,a,c) . For differential forms it will be antisymmetric in its arguments.

If we have any function of n arguments, then we can supply it with one argument, and consider the result as a function still waiting for now $(n-1)$ arguments. In computer science this is called a curried function. Here we call this operation "contraction." In old tensor it would be called "summing over the repeated index." We write the operation with the symbol \lrcorner , read "angle":

$$\omega \cdot (a, \dots) = a \lrcorner \omega \cdot (\dots).$$

Mathematicians often write the contraction $a \lrcorner \omega$ as $i_a(\omega)$.

Everywhere in this book that I use the centered dot it means evaluation, and will only be used for linear operators. Thus we write $f \cdot x$ to mean $f(x)$, and save on parentheses.

In terms of monomials the rule for contraction is, in our shorthand,

$$\frac{\partial}{\partial q} \lrcorner dq d\alpha = d\alpha.$$

If there is no dq anywhere in the differential form then the contraction is zero. If the dq factor is not in front, then you must use the anti-symmetry to move it to the front, keeping track of the sign changes.

Example:

$$\begin{aligned} \frac{\partial}{\partial x} \lrcorner dx dy &= dy, \\ \frac{\partial}{\partial y} \lrcorner dx dy &= -dx, \\ \frac{\partial}{\partial z} \lrcorner dx dy &= 0. \end{aligned}$$

Example: A 2-form in three dimensions acts on pairs of vectors as follows. Suppose we want to compute

$$dx dy \cdot \left(\frac{\partial}{\partial x} + \frac{\partial}{\partial y}, \frac{\partial}{\partial x} - \frac{\partial}{\partial y} \right).$$

This is equivalent to

$$\left(\frac{\partial}{\partial x} - \frac{\partial}{\partial y} \right) \lrcorner \left(\frac{\partial}{\partial x} + \frac{\partial}{\partial y} \right) \lrcorner dx dy,$$

with the missing parentheses implied so that \lrcorner is right associative:

$$\left(\frac{\partial}{\partial x} - \frac{\partial}{\partial y} \right) \lrcorner \left(\left(\frac{\partial}{\partial x} + \frac{\partial}{\partial y} \right) \lrcorner dx dy \right).$$

Contracting with the 2-form gives us

$$\begin{aligned} \left(\frac{\partial}{\partial x} + \frac{\partial}{\partial y} \right) \lrcorner dx dy &= \frac{\partial}{\partial x} \lrcorner dx dy + \frac{\partial}{\partial y} \lrcorner dx dy, \\ &= dy - dx. \end{aligned}$$

Contracting further gives us

$$\begin{aligned} \left(\frac{\partial}{\partial x} - \frac{\partial}{\partial y} \right) \lrcorner (dy - dx) &= \frac{\partial}{\partial x} \lrcorner (dy - dx) - \frac{\partial}{\partial y} \lrcorner (dy - dx) \\ &= -2. \end{aligned}$$

The contraction of a 1-form and a vector uses the usual duality relationship discussed before.

Rendering Forms

We know what differential forms do as operators: they act on ordered sets of vectors in an antisymmetric fashion yielding numbers. We also know how to represent them: as homogenous polynomials of basis 1-forms. Now I discuss how to draw pictures of them. One nice feature of differential forms is that it is easy to draw honest pictures of them in one, two, and three dimensions.

[In old tensor they would be called antisymmetric covariant tensors.]

To render a 1-form ω we will draw the pair of hyperplanes, the sets of all vectors v such that

[A hyperplane in n dimensions is a surface of dimension $(n-1)$.]

$$\omega \cdot v = 0,$$

$$\omega \cdot v = 1,$$

and place some some kind of distinguishing mark on the $\omega \cdot v = 1$ hyperplane.

To draw those r-forms which can be written as a single monomial, not necessarily in a given basis,

$$\alpha \wedge \beta \wedge \dots,$$


draw the "eggcrate" picture you get by superimposing the drawings for α and β, \dots separately, and add some indication of the order. In one, two, and three dimensions, all differential forms can be written as monomials in some basis. This is not true for four or higher dimensions.

Example:

$$\begin{aligned} dx dy + dy dz + dz dx &= (dx - dz) dy - (dx - dz) dx \\ &= (dx - dz) \wedge (dy - dx). \end{aligned}$$

Example:

$$dx dy + dz dw.$$



Example:

We can render $dx dy$ by the unit square with a counterclockwise arrow.

Example: The linear operations for the contraction


$$\left(2\frac{\partial}{\partial x} + \frac{\partial}{\partial y}\right) \lrcorner dx dy$$

are as follows.

(i) shear the figure, preserving its area, so that one side lines up with the vector.

(ii) rescale the figure so that the side parallel to the vector is the same length as the vector.

(iii) use the two sides parallel to the vector as a rendering of the result of the contraction, which is a 1-form. Use the counterclockwise arrow to decide which line to put the tick on.



What we have done here is to use linear transformations which do not change the value of $dx dy$ to reduce the problem to a linear transformation of the canonical problem

$$\frac{\partial}{\partial x} \lrcorner dx dy = dy.$$

18. Twisted Differential Forms

Twisted differential forms differ from ordinary differential forms in their behavior under reflection. Here is a notation that I use to denote the symmetry behavior of an object. For a given object T , I denote by $\{T\}$ the equivalence class of objects having the same behavior under reflections.

Example:

$$\{dx\} = \left\{ \frac{\partial}{\partial x} \right\},$$

since they both change sign under a reflection along the x axis, and don't change sign under y or z reflections.

I will use differential forms to represent the simplest elements of these equivalence classes, with the rule

$$\{k d\alpha\} = \{d\alpha\} \quad \text{if } k > 0$$

and

$$\{-d\alpha\} = -\{d\alpha\}.$$

We will represent twisted objects by multiplying or pairing them with one of these orientations for the whole space, that is, a $\{d\alpha\}$ with $d\alpha$ having n factors in n dimensions.

Example: The density 2-form in the plane is a twisted 2-form. It should be invariant under x or y reflections. The ordinary 2-form $dx dy$ does not have that property, but it will if we multiply it by an orientation:

$$\{dx dy\} dx dy.$$

This has the correct behavior under scaling and addition, since it is a 2-form, and the orientation gives it the correct behavior under reflections.

Example: The flow 1-forms in the plane are

$$\{dy dx\} dx \quad \text{and} \quad \{dx dy\} dy.$$

[Very few of the books on differential forms bother to deal with twisted forms, and this clutters their work up with spurious right hand rules.]

Both have the correct symmetry.

We need a convention to connect this representation of twisted forms to the constructions and renderings given earlier. A consistent convention is the following. If $\{\beta\}$ is the orientation of the twisted form

$$\{\beta\} = \{\{\Omega\} \alpha\},$$

then

$$\{\beta \wedge \alpha\} = \{\Omega\}.$$

Ω is an n -form. Be careful with the double brackets. The outer set of curly brackets means that we are discussing the orientation of what is inside, which, because of the second set of curly brackets, is a twisted form.

Example: Applied to the flux in the x direction, we let

$$\{\beta\} = \{\{dx dy\} dy\},$$

then

$$\{\beta \wedge dy\} = \{dx dy\}$$

and so

$$\{\beta\} = \{dx\}.$$

This gives the flow in the direction of dx , which we would call positive flux, and put the arrow on it pointing in the direction of the x -axis.

Example: Flow in the y direction follows the symmetrical pattern:

$$\{dy dx\} dx.$$

Other books on differential forms get around the issue of twisted forms by picking a standard orientation and using it automatically with all twisted forms. This leads to some rather ugly results: the x flux above is given by dy , while the y -flux is given by $-dx$.

I simplify the above notation by adopting the following shorthand:

(1) denote all of the missing basis forms by writing them in with a hat over them;

(2) arrange the factors, changing the sign if necessary, so that they are in the same order as the factors in the orientation;

(3) since the orientation information is now redundant, leave it out;

(4) use the hat notation only for twisted forms;

(5) the twisted n-form has nothing left out, write it with a hat over the space in front of it.

Example: The x and y fluxes in two dimensions are:

$$\widehat{dx} dy \quad \text{and} \quad \widehat{dy} dx.$$

In three dimensions the x, y, and z fluxes are:

$$\widehat{dx} dy dz = \widehat{dx} dz dy,$$

$$\widehat{dy} dx dz = \widehat{dy} dz dx,$$

$$\widehat{dz} dx dy = \widehat{dz} dy dx.$$

Note how this notation is symmetric on the explicit factors, and antisymmetric on the missing factors. This leads to the final rule in our shorthand notation for twisted tensors:

(6) leave out the explicit factors. Their order doesn't matter.

[I am tempted to call this the ghost or Cheshire cat notation.]

Example: In two dimensions the flux in the x-direction is

$$\begin{aligned} \widehat{dx} &= \\ \widehat{dx} dy &= \\ &= \{dx dy\} \widehat{dx} dy = \{dx dy\} dy. \end{aligned}$$

Example: In three dimensions the flux in the x-direction is

$$\widehat{dx} = \{dx dy dz\} dy dz.$$

Example: In three dimensions

$$\widehat{1} = \widehat{dx dy dz} = \{dx dy dz\} dx dy dz.$$

Example: Look at a constant magnetic field. This can be produced by currents circulating in the $z = 0$ plane. The symmetry of this is the following:

- (1) unchanged by a reflection $z \mapsto -z$;
- (2) reversed by reflections $x \mapsto -x$ and $y \mapsto -y$;
- (3) reversed by the interchange $x \mapsto y$ and $y \mapsto x$.

These are satisfied by the orientation $\{B\} = \{dx dy\}$. This is the symmetry of an untwisted 2-form, B, or a twisted 1-form, H. See Figure 13-1.




Figure 18-1. The sources that generate a magnetic field in the z direction, and a 2-form and a twisted 1-form with the same symmetry.

Wedge Products

The rules for the wedge product with twisted forms are:

$$\begin{aligned}\alpha \wedge \{\Omega\}\beta &= \{\Omega\}\alpha \wedge \beta, \\ &= \{\Omega\}(\alpha \wedge \beta), \\ \{\Omega\}\alpha \wedge \{\Omega\}\beta &= \alpha \wedge \beta.\end{aligned}$$

where $\{\Omega\}$ is an orientation. In terms of the hat notation:

$$d\alpha \wedge (\widehat{d\gamma d\alpha} d\beta) = \widehat{d\gamma} d\alpha d\beta,$$

or in maximally compressed notation

$$d\alpha \wedge (\widehat{d\gamma d\alpha}) = \widehat{d\gamma}.$$

Wedging two twisted forms leads to an untwisted result

$$(\widehat{d\gamma} \widehat{d\beta} d\alpha) \wedge (\widehat{d\gamma} \widehat{d\alpha} d\beta) = d\beta d\alpha,$$

(note the reversal) or in compressed notation

$$(\widehat{d\gamma} \widehat{d\beta}) \wedge (\widehat{d\gamma} \widehat{d\alpha}) = d\beta d\alpha.$$

19. Maps and Pullback

One can generalize the idea of a function so that the source space can have any dimension and the target space any other dimension. These more generalized relations are called maps.

Example: A parametrized curve in three dimensions is a map from the reals to three dimensional space.

We denote maps by any symbol we wish, and display the action of the map with a statement: $M : S \rightarrow T$, read "M maps the S to T." This may be followed by a statement of the specific rules for the map.

Example: The map

$$\gamma : \mathbb{R} \rightarrow \mathbb{R}^3; s \mapsto (\cos s, \sin s, s),$$

describes a helix in three space as a parametrized curve.

If we have a set of points in the source space, that set can be mapped to the target set, point by point. If we have a parametrized curve in the source space, then this generates a parametrized curve in the target space.

Example: If we have

$$M : \mathbb{R}^2 \rightarrow \mathbb{R}^3; (u, v) \mapsto (X(u, v), Y(u, v), Z(u, v)),$$

then given a parametrized curve

$$\gamma : \mathbb{R} \rightarrow \mathbb{R}^2; s \mapsto (U(s), V(s))$$

we are lead to a parametrized curve in \mathbb{R}^3

$$M_*\gamma : \mathbb{R} \rightarrow \mathbb{R}^3; s \mapsto (X(U(s), V(s)), Y(U(s), V(s)), Z(U(s), V(s))).$$

The curve $M_*\gamma$ is called the pushforward of the curve γ .

Since the tangents to curves are themselves linear curves, these too can be pushed forward. The rule looks remarkably like the chain

rule, which is indeed used in the derivation. This motivated our strange looking notation for tangent vectors. The basis vectors are pushed forward according to the rules:

$$\frac{\partial}{\partial x} \mapsto \frac{\partial Y}{\partial x} \frac{\partial}{\partial y}$$

where there is a summation if there are several y coordinates, and several equations if there are several x coordinates. Derive this using a Taylor's Series expansion and the chain rule.

Example: The basis vectors in the above example push forward according to

$$\begin{aligned} \frac{\partial}{\partial u} &\mapsto \frac{\partial X}{\partial u} \frac{\partial}{\partial x} + \frac{\partial Y}{\partial u} \frac{\partial}{\partial y} + \frac{\partial Z}{\partial u} \frac{\partial}{\partial z}, \\ \frac{\partial}{\partial v} &\mapsto \frac{\partial X}{\partial v} \frac{\partial}{\partial x} + \frac{\partial Y}{\partial v} \frac{\partial}{\partial y} + \frac{\partial Z}{\partial v} \frac{\partial}{\partial z}. \end{aligned}$$

A dual situation exists with functions. If we have a function on the target space

$$f : T \rightarrow \mathbf{R},$$

then this leads to a function on the source space. To compute the function at a point in the source space, follow the map to a point in the target space and then evaluate the function there. This related function is called the pullback of the original function.


Likewise the gradients of functions, being linear functions, can also be pulled back. See figure 19-1. for the situation where the map is being used as in the above examples to define a two dimensional subspace of three dimensional space.

There you can see why pullback is sometimes called sectioning. The subspace cuts out a section of the 1-form, and defines a 1-form on the subspace.

Example: In the above example the basis 1-forms pullback according to the rules

$$\begin{aligned} dX &\mapsto \frac{\partial X}{\partial u} du + \frac{\partial X}{\partial v} dv, \\ dY &\mapsto \frac{\partial Y}{\partial u} du + \frac{\partial Y}{\partial v} dv, \\ dZ &\mapsto \frac{\partial Z}{\partial u} du + \frac{\partial Z}{\partial v} dv, \end{aligned}$$

Figure 19-2. A one form being pulled back onto a subspace.



which indeed looks just like the chain rule.

Be warned that it makes no sense to try to pushforward a 1-form or to pull back a tangent vector.

To pullback a twisted form the subspace must be given a transverse orientation. Look at a subspace defined such that each term of a monomial $d\alpha$ pulls back to zero. It is said to have a transverse orientation if we pick out either $\{d\beta\}$ or $-\{d\beta\}$, where $d\beta$ is a monomial containing all of the remaining forms.

Example: If a hypersurface is specified by a coordinate $q = \text{constant}$, and a basis for the whole space is given by $dq d\alpha$, then the usual (inner) orientation for the hypersurface would be a choice of either $\{d\alpha\}$ or $-\{d\alpha\}$. A transverse orientation is a choice of either $\{dq\}$ or $-\{dq\}$.

If the transverse orientation of a subspace defined by the pullbacks of all the 1-forms in a monomial $d\alpha$ vanishing is given by $\{d\beta\}$, then the pullback rule is

$$\widehat{d\alpha} \widehat{d\beta} d\gamma \mapsto \widehat{d\beta} d\gamma.$$

Example: In three-space, let us pull back the twisted 1-form $(-dy, \Omega)$ onto the $z = 0$ surface from above. Treating

[It is good practice to use an upper case letter for x when it denotes a dependent variable, and lower case for x when it is an independent variable.]


$z = 0$ as the boundary of the upper volume, we have z decreasing to the outside, so the orientation $-\{dz\}$ corresponds to $\{d\alpha\}$ above.

The pullback is

$$-dy \widehat{dz} \widehat{dx} = \widehat{dz} dy \widehat{dx} \mapsto -dy \widehat{dx} = \widehat{dx} dy.$$

The x -flux, recall is $\widehat{dx} dy$, so this is a flux in the negative x -direction.

Figure 19-3. Pull back $-dy \widehat{dz} \widehat{dx}$ from above.



CHAPTER FOUR. CALCULUS OF FORMS

This chapter continues the basic rules for the manipulation of differential forms. It covers the rules for differentiation and integration.

22. Stokes' Theorem

Integrating a differential form around the boundary of a region gives the same result as integrating its exterior derivative over the region itself. This statement includes both the fundamental theorem of calculus and Stokes' Theorem in vector calculus. It also covers the analogs of these theorems in higher dimensions. There is a version that applies to twisted forms and that one includes the divergence theorem and its analogs in other dimensions.

I will not prove Stokes' Theorem here. The most reasonable course of action is to use Stokes' Theorem to define the exterior derivative operator d . Then one finds the rule for calculating it given earlier. On a lattice there is no other way to calculate d except as the operator that makes Stokes' theorem work.

Example: Let us integrate the 1-form

$$\omega = x dy$$

around the unit square in the right upper quadrant. Go around the square in the counter-clockwise direction. Only the right edge contributes to the integral, along the line $x = 1$. There we have

$$\int_{\partial\Gamma} \omega = \int_0^1 dy = 1.$$

Now the exterior derivative is

$$d\omega = dx dy.$$

This is the unit density 2-form, and so

$$\int_{\Gamma} d\omega = \int_0^1 \left(\int_0^1 dy \right) dx = 1.$$

These agree, as I claimed.

We will use the partial derivative sign to denote the boundary of a region. Stokes' Theorem says that always

$$\int_{\Gamma} d\omega = \int_{\partial\Gamma} \omega.$$

To complete this we need to know how to orient the boundary of a region. I will only give the rule for the simple case where everything is lined up with the coordinates. If the boundary of the region Γ is given by $q = \text{constant}$, with q increasing to the outside of the region, then we must orient the boundary $\partial\Gamma$ so that

$$\{dq \wedge \{\partial\Gamma\}\} = \{\Gamma\}.$$

Example: For our preceding example, the boundary on the right is given by $x = 1$, and $x > 1$ lies outside the region. Hence we must satisfy

$$dx \wedge \{\partial\Gamma\} = \{\Gamma\}.$$

If we give Γ the usual orientation $\{dx dy\}$, then this is satisfied with

$$\{\partial\Gamma\} = \{dy\}.$$

Since we are integrating upward on this leg, the answer is indeed positive.

Stokes' Theorem requires that both the space and its boundary be oriented consistently. Don't apply Stokes' Theorem to a space that cannot be oriented.

Example: Take a Mobius strip. Its boundary is a closed curve. You cannot use Stokes' Theorem to relate a line integral along the edge to the integral of its exterior derivative over the Mobius strip. If you run an electric current through the edge, it looks just like a coil with two turns. The magnetic field will hardly penetrate the Mobius strip, however. To apply Maxwell's equations you need to use the less smooth and certainly less obvious orientable surface that spans the curve. This is a tricky surface to visualize. I sketch it below. You may want to cut a model out of cardboard and verify that it is orientable and that the Mobius strip is not.

Divergence Theorem

Stokes's Theorem also applies to twisted tensors integrated over regions with transverse orientations. I call this case the divergence




Figure 22-1. Simplicial picture of a two-turn current loop and the oriented surface that it bounds in addition to bounding the Mobius strip. Here line CF is vertical. The four simplicial faces are ABC, CDE, EFC, and FAC.

theorem. It states that

$$\int_{(\Gamma, \Omega)} d(\omega, \Omega) = \int_{(\partial\Gamma, \Omega)} (\omega, \Omega).$$

Here Ω is an n -form giving the orientation. For the case of transverse orientations, the orientation of the boundary must satisfy

$$\{(\partial\Gamma, \Omega)\} = \{(\Gamma, \Omega)\} \wedge dq.$$

Example: Suppose that a flux 2-form in three dimensions is given by the twisted 2-form

$$j = u \widehat{dx} dy dz + v \widehat{dy} dz dx + w \widehat{dz} dx dy,$$

Then

$$dj = \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} \right) \widehat{dx dy dz}.$$

The flux through the faces of a unit cube would be computed like this: for the $x = 1$ face the transverse orientation is $\{dx\}$. To pull j back onto this surface we use

$$\begin{aligned} \widehat{d\alpha d\beta d\gamma} &\mapsto \widehat{d\beta d\gamma}, \\ u \widehat{dx dy dz} &\mapsto u \widehat{dy dz}, \\ v \widehat{dy dz dx} &\mapsto 0, \\ w \widehat{dz dx dy} &\mapsto 0. \end{aligned}$$

23. Conservation Laws

Conservation laws can be expressed very naturally using differential forms.

Example: Look at cars moving along a one-lane road. Let $\rho(x, t)$ be the local density of cars, and $k(x, t)$ the local flux of cars. The number of cars between $x = a$ and $x = b$ will be

$$\int_a^b \rho(x, t) dx.$$

The number of cars that cross $x = a$ from time $t = c$ to $t = d$ will be

$$\int_c^d k(a, t) dt.$$

The net loss in the number of cars in the region $a < x < b$ over that time will be

$$\int_a^b \rho(x, c) dx - \int_a^b \rho(x, d) dx.$$

If cars are not created or destroyed, then this decline must be balanced by the cars coming in and out of the region:

$$\int_c^d k(a, t) dt - \int_c^d k(b, t) dt.$$

If we bring all the terms to one side, and properly straighten out the signs, we have

$$\int_a^b \rho(x, c) dx + \int_c^d k(b, t) dt + \int_b^a \rho(x, d) dx + \int_d^c k(a, t) dt = 0.$$

Let us define a 1-form on this two dimensional spacetime

$$\begin{aligned} j &\equiv \rho \widehat{dt} + k \widehat{dx}, \\ j &= \rho \widehat{dt} dx + k \widehat{dx} dt, \\ j &= -\rho dx dt dx + k dx dt dt, \\ j &= -\rho dx + k dt. \end{aligned}$$

The first form is our preferred shorthand for twisted forms. In the second line I add in the explicit terms, but still leave out the orientation. In the third line I add in the orientations specifically. In the fourth line I use the notation that you will find in most books on differential forms, where the standard orientation is omitted as a shorthand. Note the peculiar minus sign that arises from this convention.

Now let Γ be the rectangular region of spacetime, then we have found that

$$\int_{\partial\Gamma} j = 0.$$

To assert this for all possible regions Γ is the general statement of the conservation of cars. Using the divergence theorem we can write this

$$\int_{\Gamma} dj = 0.$$

If this is to be true for all regions, then

$$dj = 0,$$

is our conservation law in differential form. This would normally be written

$$\frac{\partial k}{\partial x} - \frac{\partial \rho}{\partial t} = 0.$$

All conservation laws will fit the above pattern: an n dimensional space with a twisted $(n - 1)$ -form with vanishing exterior derivative.

Example: Here is a twisted 3-form that describes the conservation of electric charge in spacetime.

$$j \equiv J_x \widehat{dx} + J_y \widehat{dy} + J_z \widehat{dz},$$

$$j = J_x \widehat{dx} dy dz dt + J_y \widehat{dy} dx dz dt + J_z \widehat{dz} dx dy dt + \rho \widehat{dt} dx dy dz.$$

The exterior derivative is

$$dj = \frac{\partial J_x}{\partial x} + \frac{\partial J_y}{\partial y} + \frac{\partial J_z}{\partial z} + \frac{\partial \rho}{\partial t} = 0.$$

Note that it doesn't matter in what order you add in the missing three terms in the first definition of the current since

$$\begin{aligned} \widehat{dx} dy dz dt &= \{dx dy dz dt\} dy dz dt, \\ &= \{dx dz dy dt\} dz dy dt, \\ &= \widehat{dx} dz dy dt. \end{aligned}$$

24. Inverse of Stokes' Theorem

From the basic theorem

$$dd = 0,$$

we can find solutions of the differential form equation

$$d\omega = 0,$$

by taking

$$\omega = d\phi,$$

for any form ϕ one rank lower than ω .

Example: In electrodynamics Maxwell's equation for the B field reads

$$dB = 0,$$

and we want to introduce a 1-form potential A such that

$$B = dA.$$

Can this always be done? That is the question for this section. A similar question comes up with the scalar potential: Does $dE = 0$ mean we can always find a potential such that $E = -dV$?

A form which can be written in terms of a potential is called exact, as in the sense of an exact differential in thermodynamics. A form with zero exterior derivative is called closed, as in a surface with no boundary.

[Recall that the exterior derivative and boundary operators are adjoint.]

The answer to this question is involved, and one of the advantages of differential forms is the order that it brings to this situation. Again I will give no proofs. This is quite a complicated subject, in fact. We ask if there are closed forms that are not exact. This question is equivalent to its adjoint: are there surfaces that have no boundary that are not themselves boundaries? Such a surface is called a cycle. For simple spaces this adjoint question is easy.

Example: On the surface of the sphere there are no closed curves that are not boundaries. The equator, for example, bounds two discs, the northern and southern hemispheres.

On the surface of a torus there are two intrinsically different curves that are not boundaries. They go around the torus in the two possible ways. These cycles cannot be shrunk to zero while staying in the surface.

Example: Look at an annular space in two dimensions, the plane with a central hole removed. Any curve encircling the hole in the middle is a cycle.

Look at the 1-form $d\theta$. This is a closed 1-form, but it is not exact. The notation $d\theta$ might lead you to believe that it is the d of θ , but θ is not a true function; it is not single-valued over the region. Normally you don't care about this, but in topology these things matter.

Now we can state the solution to our problem: you can write

$$\omega = d\alpha$$

provided that

$$d\omega = 0$$

and

$$\int_{\Gamma} \omega = 0$$

for all cycles Γ .

Example: Continuing the above example: $d\theta$ is not exact because the integral around any circle enclosing the origin has a non-zero value

$$\int_0^{2\pi} d\theta = 2\pi.$$

If you have a closed form with a non-zero value integrated over a cycle you can subtract out that degree of freedom.

Example: In the above case, if the form had

$$\int_{\Gamma} \omega = A,$$

then

$$\omega - \frac{1}{2\pi A} d\theta$$

will have zero integral, and will therefore be exact. Thus we can write

$$\omega = \frac{1}{2\pi A} d\theta + d\alpha,$$

for some α , and this gives a representation of all the possible solutions.

25. Delta Functions

With a notation for integration that does not depend upon dimension, there is little or no need for the conventional notation for functions confined to subspaces or even points. The connections between the dimensions are provided by the pullback formulae.

I intend to look through some of the books on advanced delta-functionology and verify that they have been made totally redundant by the differential forms notation. I will start with *Mathematical Tools for Changing Spatial Scales in the Analysis of Physical Systems* by William G. Gray *et al.*

CHAPTER FIVE. GEOMETRY WITH FORMS

Not all the geometrical objects that we deal with can be represented by differential forms. In particular the metric properties of a space are represented by symmetric tensors. There are tricks that we can use to represent them in our language, however, and these are the subject of this chapter.

The techniques here are very powerful, and when they were introduced in general relativity back in the 1970's, they made calculations easier by a factor to more than ten or twenty.

26. Tensor Algebra

There are some situations where a tensor multiplication that is not antisymmetric is more natural. The general ordered multiplication of tensors is indicated by the tensor product symbol \otimes . If u and v are tensors, and a and b are in spaces dual to the spaces containing u and v , then we define

$$u \otimes v \cdot (a, b) = (u \cdot a)(v \cdot b).$$

The multiplication on the left is tensor multiplication, the multiplication on the right hand side is just the ordinary multiplication of numbers. Here u and v need not be tensors of the same type, all that is necessary is that they be linear operators.

[Remember the centered dot for evaluation of a linear operator.]

If k is a real number, then we require that

$$k(u \otimes v) = (ku) \otimes v = u \otimes (kv).$$

Example: The wedge product of two 1-forms can be written

$$\alpha \wedge \beta = \alpha \otimes \beta - \beta \otimes \alpha.$$

The metric properties of a space are expressed by a symmetric tensor. Euclidean geometry uses the tensor \mathcal{E}

$$\mathcal{E} = dx \otimes dx + dy \otimes dy.$$

The usual dot product between two vectors a and b is

$$a \bullet b = \mathcal{E} \cdot (a, b).$$

[I use a bullet for the vector dot product to distinguish it from the centred dot.]

The spacetime geometry of special relativity is described by a tensor \mathcal{M}

$$\mathcal{M} = dx \otimes dx - dt \otimes dt$$

All of special relativity is contained in the above equation and in its four dimensional extension:

$$\mathcal{M} = dx \otimes dx + dy \otimes dy + dz \otimes dz - dt \otimes dt.$$

When you see geometry metric tensors written

$$\mathcal{E} = dx^2 + dy^2,$$

you are seeing a shorthand notation where the tensor product symbol has been left out.

27. The Star Operator

The calculus of differential forms is an efficient tool for calculating with antisymmetric tensors. To include symmetric tensors is important, and we resort to a trick for this. You are already familiar with this trick in the realm of vectors. One commonly represents little areas (antisymmetric pairs of vectors) with their surface normal vectors. In the same way, the star operator represents antisymmetric pairs of 1-forms with a twisted 1-form.

To dot two forms together one needs to use the metric structure, just as one would to dot two vectors together. We can think of the metric as map

$$\mathcal{E} : V \times V \rightarrow \mathbf{R}$$

or as the partially evaluated operator

$$\mathcal{E} : V \rightarrow (V \rightarrow \mathbf{R}).$$

That is, we can think of \mathcal{E} as an operator which maps vectors to 1-forms, and then the natural duality leads to the dot product. To stay in the realm of forms, we can mimic this action.

Introduce an operator that maps r -forms to twisted $(n-r)$ -forms:

$$\alpha \mapsto *\alpha.$$

Now given another form of the same type as α , we can form

$$\beta \wedge (*\alpha)$$

and this is an n -form. These are all similar except for magnitude, and we can compare the magnitude with that of the unit volume element, which we will write $*1$:

$$\begin{aligned} \beta \wedge (*\alpha) &= *1 \cdot (\beta \bullet \alpha), \\ &= *(\beta \bullet \alpha). \end{aligned}$$

[The centered dot for metric inner product will be rarely used in this book. To distinguish it from mere evaluation of a linear operator, I am using a very bold dot.]

This definition of the dot product is bilinear in its arguments. We will have to arrange the definition of $*$ so that it is symmetric, and then we will have arranged a representation of our metric that stays within the domain of differential forms.

The tool for this will be the orthonormal frame. Look at a basis of 1-forms that are all of unit size according to the metric, and mutually orthogonal. We will denote these by $\omega^x, \omega^y, \dots$, although ω^x, \dots would be more in parallel with dx, \dots

Because they are orthonormal we can write the metric in terms of them

$$\mathcal{E} = \omega^x \otimes \omega^x + \omega^y \otimes \omega^y,$$

using the tensor product of the last section. Following the conventions used with the coordinate basis, we will write the arbitrary such basis form ω^i and the wedge product of an arbitrary number of them by ω^A .

The rule for the metric operator $*$ when the metric is positive definite is

$$*\omega^A = \widehat{\omega^A}, \quad \text{Positive Definite Metric}$$

Example: In three dimensional Euclidean space dx, dy, dz are already an orthonormal basis. Thus

$$\begin{aligned} *dx &= \widehat{dx}, \\ *dy &= \widehat{dy}, \\ *dz &= \widehat{dz}, \\ *dx dy &= \widehat{dx} \widehat{dy}, \\ *1 &= \widehat{1}. \end{aligned}$$

Example: In two dimension polar coordinates we have an orthonormal basis

$$\begin{aligned} \omega^r &= dr, \\ \omega^\theta &= r d\theta. \end{aligned}$$

and

$$\begin{aligned} *_2 dr &= *\omega^r = \widehat{\omega^r} \omega^\theta \\ &= \widehat{dr} r d\theta = r \widehat{dr} d\theta \\ *_2 d\theta &= \frac{1}{r} \widehat{d\theta} dr, \\ *_2 1 &= \widehat{1} \omega^r \omega^\theta = \widehat{1} r dr d\theta. \end{aligned}$$

When there are two sets of basis forms being used, it is not safe to leave out basis forms when describing twisted forms. Leaving them in makes only a small complication.

[The subscript here indicates that this is a different star operator from the one in three dimensions.]

The metric inner product of two r -forms is given by

$$*(\alpha \circ \beta) = (\alpha \circ \beta) * 1 = \alpha \wedge * \beta.$$

To see that this is symmetric, let

$$\alpha = \omega^A.$$

This entails no loss of generality because this is an algebraic argument that only involves the value at a single point. Now β must involve exactly the same factors, or the wedge product will vanish. But if we have

$$\beta = k d\omega^A$$

then it is clear that

$$\alpha \wedge * \beta = \beta \wedge * \alpha.$$

To represent Lorentz geometry we need a more general rule, adding one minus sign for each timelike basis form being starred:

$$*\omega^A = (-)^p \widehat{\omega^A} \omega^B,$$

where p is the number of timelike factors in ω^A . In Lorentz geometry the number of timelike factors will be either zero or one.

Example: For two dimensional spacetime we can pick an orthonormal basis

$$\omega^t = dt,$$

$$\omega^x = dx.$$

Then we have

$$*_L \omega^t = -\widehat{\omega^t} \omega^x$$

$$= -\widehat{dt} dx,$$

$$*_L \omega^x = \widehat{\omega^x} \omega^t$$

$$= \widehat{dx} dt,$$

and we see that

$$\begin{aligned} \omega^t \wedge *_L \omega^t &= dt \wedge (-\widehat{dt}) dx, \\ &= -\widehat{1} dt dx = -\widehat{1} dx dt \end{aligned}$$

Also we have

$$*_L 1 = -\widehat{1} dx dt,$$

so we would say that

$$\omega^t \cdot \omega^t = -1.$$

Likewise you can show that

$$\omega^x \cdot \omega^x = 1.$$

28. Symmetry and the Lie Derivative

The notion of symmetry is closely connected with conservation laws. For some systems every continuous symmetry leads to a conservation law.

Under a continuous transformation, every point moves along a curve, and the tangent vector to that curve is a representation of an infinitesimal transformation. We will restrict our attention to infinitesimal transformations.

Example: An infinitesimal translation along the x -axis is represented by the vector field

$$\frac{\partial}{\partial x}.$$

A rotation about the origin by

$$y \frac{\partial}{\partial x} - x \frac{\partial}{\partial y},$$

and a Lorentz transformation by

$$t \frac{\partial}{\partial x} + x \frac{\partial}{\partial t}.$$

Given a transformation, a good question to ask is what objects are not changed by the transformation. We call such objects invariant under the transformation. We could also say they have such-and-such a symmetry.

Example: The function r has rotational symmetry, but not translational symmetry.

Example: A function $f(x, y)$ is invariant under an infinitesimal transformation given by the vector field k if

$$k \lrcorner df = 0.$$

We want to generalize the differentiation shown in the last example so that we can decide if a given differential form is invariant or not.

Example: The 1-form dy is invariant under the transformation

$$y \frac{\partial}{\partial x},$$

which represents a shear. To see this pictorially, transform each point in the rendering of dy by a pair of parallel lines. The picture of dy slides into itself.

The differentiation which measures the change under infinitesimal transformations is called Lie differentiation. It is very easy for differential forms. The Lie derivative \mathcal{L} is given by

$$\mathcal{L}_Q \frac{\partial}{\partial q} f dq d\alpha = Q \frac{\partial f}{\partial q} dq d\alpha + f dQ \wedge \alpha.$$

It satisfies the important identity

$$\mathcal{L}_k \omega = k \lrcorner d\omega + d(k \lrcorner \omega).$$

Example:

$$\mathcal{L}_{y \frac{\partial}{\partial x}} dy = y \frac{\partial}{\partial x} (ddy) + d(y \frac{\partial}{\partial x} \lrcorner dy) = 0.$$

Isometries

One can also study the symmetries of a metric. Here the focus is not on which metric tensors have a given symmetry, but rather, given a metric, what are its symmetries. If the metric is given in terms of its components, with summation over repeated indices

$$g_{\mu\nu} dx^\mu \otimes dx^\nu,$$

then a vector

$$k = k^\mu \frac{\partial}{\partial x^\mu}$$

is a symmetry provided that

$$g_{\mu\nu,\sigma} k^\sigma + g_{\mu\sigma} k^\sigma_{,\nu} + g_{\nu\sigma} k^\sigma_{,\mu} = 0.$$

Here the comma indicates a partial derivative. This is by no means obvious; I am just quoting it without proof, since the proof really

needs more tensor analysis tools than you have. Such a symmetry vector is called a Killing vector.

Example: It is not obvious, but the symmetries of the two dimensional Euclidean metric are the three independent Killing vectors:

$$\begin{aligned} & \frac{\partial}{\partial x}, \\ & \frac{\partial}{\partial y}, \\ & y \frac{\partial}{\partial x} - x \frac{\partial}{\partial y}. \end{aligned}$$

The only result you will need for Killing vectors is the relation

$$\mathcal{L}_k * \alpha = * \mathcal{L}_k \alpha.$$

29. Old Vector Calculus

You might be expecting here a translation guide along the lines of: the dot product of two 1-forms is given by

$$\alpha \cdot \beta \mapsto *^{-1}(\alpha \wedge * \beta),$$

and the cross product by

$$\alpha \times \beta \mapsto *(\alpha \wedge \beta).$$

This is not the case. The hidden stratification of three-degree-of-freedom objects shows up here in a multiplicity of translations for dot and cross products. I will give examples here drawn mainly from electrodynamics.

First we note that the stratification must be respected in equations. Consider Ohm's Law. It makes no sense to write this in forms

$$J = \sigma E,$$

with σ a scalar parameter, because J is a twisted 2-form, (integrate it over a two-sided surface) while E is a 1-form, (integrate along a directed line). The geometrically correct translation must be something like

$$J = \sigma * E.$$

Even this can be criticized, and the relativistically correct form requires that the material velocity be specified, to

$$J = \sigma * (E - v \lrcorner B).$$

There are a variety of dot products in electrodynamics. The electrostatic energy density will be a twisted 3-form

$$E \wedge *E,$$

and similarly magnetic energy

$$B \wedge *B.$$

Relativistically these cannot be separated, and only the combination

$$E \wedge *E + B \wedge *B$$

makes sense if you are moving between different states of motion.

Another important invariant is the dot product $E \cdot B$. This too can be written without reference to a metric as

$$E \wedge B.$$

Likewise, the local density for ohmic heating becomes the twisted 3-form

$$J \wedge E.$$

The most frequent cross product is $v \times B$, which translates into

$$-v \lrcorner B,$$

again independent of any choice of metric.

What about the Poynting vector $E \times B$? This will be a twisted 2-form, giving the flux of energy across a two-sided surface. It translates into

$$E \wedge *B.$$

More difficult are cross products like $J \times B$. This gives the local density of force on a conducting fluid in magnetohydrodynamics. This is really three different conservation laws, one for x-momentum, and so on. To translate this into forms we need to explicitly introduce the conserved quantity by giving the Killing vector field that generates it. In terms of this vector k we have

$$J \wedge k \lrcorner B.$$

This is a collection of twisted 3-forms, one for each Killing vector. In Euclidean space there are three Killing vectors representing translational symmetries. We have something like a "vector" of twisted 3-forms, but this is not a "vector" in the sense of vector field, like the k . It is not attached to any point, but has components that require an integration over all space. In a similar manner we can also include the angular momentum conservation laws. Really we have not three but six conservation laws.

For a loop of current in a magnetic field, the k-momentum balance involves

$$I \int_{\Lambda} k \lrcorner B.$$

Note that here we are integrating an untwisted 1-form over a loop with inner orientation specified by the direction of current flow.

30. Simplicial Calculus

I want to just sketch here the start of a geometrically natural lattice calculus that would provide the underpinnings for finite-difference, finite-element, and boundary-element calculations. The use of differential forms and their analogs is particularly well suited to a discussion of conservation laws. The approximations involved in the discretization appear only at the algebraic level. Unlike finite-difference methods, we do not replace differential equations with infinite-order equations.

The main utility for this section for us will be the nice pictures it provides for the boundary operator and exterior differentiation.

[A relaxed and extensive discussion of this can be found in Bamberg and Sternberg.]

Elements

We will consider cellular lattices embedding in Euclidean space. The embedding may be time dependent or not, as the problem demands. The elements of the lattice are vertices, edges, faces, and cells, in three dimensions. The lattice element of dimension r will be referred to generally as an r -cell. These cells need not have the minimum number of vertices. That is, the faces can be squares, not triangles. This makes things like the orientation and the boundary operator more complicated, but allows us to consider the dual lattice as well. The dual of a simplicial lattice is not usually simplicial.

We need to orient these lattice elements. For our case of lattices embedded in Euclidean space, we specify the orientation by giving any r -form which pulls back to a non-zero r -form on the lattice element. A twisted orientation for a lattice element in n dimensions can be specified by any $(n - r)$ -form which pulls back to zero on the r -cell. An ordinary orientation for the r -cells of a lattice naturally defines a twisted orientation on the dual lattice.

Chains and Cochains

To analog of a vector will be the formal sum of any number of 1-cells, with coefficients that can be reals although they will usually be integers. The formal sums of r -cells are called r -chains. They are the lattice versions of curves, surfaces, and volumes. A picture of a

1-chain is a picture of the lattice, with a weight written next to all of the non-zero 1-cells.

The linear operators on r-chains are called r-cochains. These are the analogs of differential forms on the lattice. The picture of a 1-cochain looks just like that of a 1-chain. Are they really the same thing? To see the difference, you need to consider a refinement of the lattice. If you subdivide an edge into two pieces, one after the other, then they go to a new chain with weight 1 on each piece. A particular 1-cochain that have a value of 1 on the original edge, will now have a value of 1/2 on each of the new edges, assuming that it was divided symmetrically. This representation of both chains and cochains in terms of chains is similar to our decision to represent both vectors and forms in the tangent space.

Example: The electrostatic field in two dimensions will be a 1-cochain giving the potential difference across any edge. This has an ordinary orientation. The electric flux is a twisted 1-cochain, and this will give the flux through any edge. The electric charge will be a twisted 2-cochain, which tells you how much charge is contained on each face.

The equations of electrostatics will relate the flux and the field strength, and express the balance between the net flux out of a face and the charge it contains.

Boundary Operator

For a simplex the boundary operator ∂ is defined by

$$\partial(p_0 \dots p_r) = \sum_i (-)^i (p_0 \dots \widehat{p}_i \dots p_r),$$

where $(p_0 \dots p_r)$ denotes the r-simplex with vertices p_0 and so on, and $(p_0 \dots \widehat{p}_i \dots p_r)$ is the (r-1)-simplex which is missing the vertex p_i .

The boundary of any cell can be found by taking a simplicial decomposition and adding up the boundaries of all the simplices. The boundary of a chain is the sum of the boundaries of its cells, with the appropriate weights and signs.

This is consistent with the orientation rule for the boundaries of regions

$$d\phi \wedge \{\partial\Gamma\} = \{\Gamma\},$$

where ϕ is a function which increases to the outside of Γ . Both sides of the above equation are equivalence classes of r-forms.




Figure 30-1. The boundaries of two 1-chains and a twisted 1-chain in two dimensions.

∇




Figure 30-2. The boundaries of a 2-chain and a twisted 2-chain in two dimensions.

∇

Coboundary Operator

The adjoint to the boundary operator is d , the coboundary operator

defined by

$$d\omega \cdot g = \omega \cdot \partial g,$$

for any r -chain g and $(r-1)$ -cochain ω . The operator d is the lattice version of the exterior derivative d .




Figure 30-3. The coboundaries of a 0-chain and a 1-chain in two dimensions.




Figure 30-4. The coboundaries of a twisted 0-chain and a twisted 1-chain in two dimensions.



v

CHAPTER SIX. ELECTRODYNAMICS

Electrodynamics is the ideal field theory to show off the advantages of differential forms. You should be able to follow the developments in this chapter after or during your junior year electrodynamics class. The section on estimating capacitance and inductance is provided to show that forms are not just useful in the formulation of the theory, but also in doing hard calculations.

31. Electromagnetic Fields and Equations

The geometric picture of the four electromagnetic fields follows from their integral properties. The electric field E integrated along a path tells us how much work is done by a charge carried along that path. Clearly E is a 1-form.

The flux D integrated over a sphere tells us how much charge is inside. D is a twisted 2-form.

The time derivative of the magnetic flux through a 2-surface gives us the voltage induced around its boundary. Using Stokes' Theorem, this means that B will be an ordinary 2-form. Finally, H will be a twisted 1-form. The density of electrical charge will be given by a twisted 3-form, and the current density by a twisted 2-form. The current J through a surface is related to the integral of H around its boundary. This shows that H is indeed a twisted 1-form.

These are all given in a form suitable for a curvilinear coordinate system, but one which is not moving. Time is treated as a separate variable here, and the operator d does not include time differentiation.

Maxwell's equations in natural units, cgs Gaussian units with the speed of light set equal to unity, are

$$\begin{aligned}\frac{\partial B}{\partial t} &= -dE, \\ \frac{\partial D}{\partial t} &= dH - 4\pi J, \\ dB &= 0, \\ dD &= 4\pi\rho.\end{aligned}$$


In addition we need relations between E and D and between B and H . In vacuum we have

$$\begin{aligned}D &= *E, \\ H &= *B.\end{aligned}$$

The force law for a charged particle is

$$F = q(E - v \lrcorner B).$$

As I mentioned before, force is a 1-form. Integrated along a line it gives you work.



32. Electrostatics

To go deeper into the differential forms calculus, let us look at the physically simpler situation of electrostatics. The vacuum equations simplify to

$$\begin{aligned}dE &= 0, \\dD &= 4\pi\rho, \\D &= *E.\end{aligned}$$

Over a region like all of space, where there are not any holes, we can infer from the fact that E is closed that we can introduce a potential (the sign is a historical convention)

$$E = -dV.$$

In a region with no charge this leads to Laplace's equation

$$d * dV = 0.$$

Note that you cannot bring a derivative through a star, otherwise this would be a trivial equation.

Let us write out Laplace's equation in spherical polar coordinates. The d operator is simple in a coordinate basis

$$dr, d\theta, d\phi,$$

while the star operator is simple in an orthonormal basis

$$\begin{aligned}\omega^r &= dr, \\ \omega^\theta &= r d\theta, \\ \omega^\phi &= r \sin \theta d\phi.\end{aligned}$$

First we take the gradient of the potential:

$$dV = \frac{\partial V}{\partial r} dr + \frac{\partial V}{\partial \theta} d\theta + \frac{\partial V}{\partial \phi} d\phi.$$

Next convert to an orthonormal basis:

$$dV = \frac{\partial V}{\partial r} \omega^r + \frac{1}{r} \frac{\partial V}{\partial \theta} \omega^\theta + \frac{1}{r \sin \theta} \frac{\partial V}{\partial \phi} \omega^\phi.$$

Now we can star this

$$*dV = \frac{\partial V}{\partial r} \widehat{\omega^r} \omega^\theta \omega^\phi + \frac{1}{r} \frac{\partial V}{\partial \theta} \widehat{\omega^\theta} \omega^r \omega^\phi + \frac{1}{r \sin \theta} \frac{\partial V}{\partial \phi} \widehat{\omega^\phi} \omega^r \omega^\theta.$$

And now convert back to a coordinate basis, getting ready for the final exterior derivative:

$$*dV = r^2 \sin \theta \frac{\partial V}{\partial r} \widehat{dr} d\theta d\phi + \sin \theta \frac{\partial V}{\partial \theta} \widehat{d\theta} dr d\phi + \frac{1}{\sin \theta} \frac{\partial V}{\partial \phi} \widehat{d\phi} dr d\theta.$$

Now the final exterior derivative. Because of the antisymmetry, this involve only one differentiation per term. Collecting up the results we have

$$d * dV = \left(\frac{\partial}{\partial r} (r^2 \frac{\partial V}{\partial r}) + \frac{\partial}{\partial \theta} (\sin \theta \frac{\partial V}{\partial \theta}) + \frac{\partial}{\partial \phi} \frac{\partial^2 V}{\partial \theta^2} \right) \widehat{dr} d\theta d\phi.$$

Conformal Mapping

In two dimensions the equations of electrostatics have a particularly simple geometric picture. We know that we can draw E by drawing the contour lines of V . In charge-free regions, we can draw flux tubes of D . How are these flux tubes related to E ?

The relation between E and D is given by the algebraic equations

$$D = *D.$$

No derivatives are involved. Thus we can pursue this argument by picking special coordinates for E without any loss of generality. Take

$$E = E_0 dx,$$

at some point. Then

$$D = E_0 \widehat{dx} dy.$$

The lines representing E are perpendicular to the lines of D , and they have the same spacing. Thus the lines of E , called equipotentials, and the lines of D , called streamlines, meet everywhere at right angles, and form what are called "curvilinear squares." The squares are only true for lines infinitesimally close together, but is still visible even in the finite case. You can often sketch a crude but servicable approximation to a solution of Laplace's equation this way.

These squares do not depend on the actual field strength. Doubling the strength of the E field doubles the strength of D , in all dimensions. But only in two dimensions is the D field also a 1-form. In three dimensions you do not have curvilinear cubes because when you double the E field, this factor of two is spread over two other dimensions, $\sqrt{2}$ for each.

This field strength invariant behavior of two dimensional electrostatics is called conformal symmetry. Every analytic function in complex variable theory leads to a picture with exactly these properties, and this is a powerful technique for solving Laplace's equation in two dimensions.

The same argument used above gives us the energy in the configuration

$$W = \frac{1}{2}CV^2 = \frac{Q^2}{2C}.$$

Estimating Capacitance

We can combine the two expressions for energy into a powerful tool for estimating the capacity coefficients for system too difficult to solve exactly.

Suppose we guess the voltage V everywhere in space, call this V_e , making sure that we have the correct voltage on the conductors. This leads to an estimate of the electric field, E_e , which will be closed but not satisfy Maxwell's equations. We now calculate the energy in this estimate of the field, and discover that first-order mistakes in the guess of the field leads only to a second-order mistake in the energy. This is pretty much like getting some information for nothing.

Write the estimate in terms of a difference field e :

$$E_e = E + e,$$

then

$$\int E_e \wedge *E_e = \int (E + e) \wedge (*E + *e).$$

The linear error terms in here are identical, because

$$e \wedge *E = E \wedge *e.$$

Now introduce the potential going with e . This is possible because e is closed over the entire space, even inside the conductors. The linear error term is thus

$$2 \int e \wedge *E = -2 \int dv \wedge *E.$$

Now integrate by parts, using

$$d(v * E) = dv \wedge *E + v d * E.$$

The surface integral is taken on the surfaces of the conductors. Since the estimate has the correct voltage there by assumption, we have $v = 0$ on the surfaces and this surface term vanishes. It is only a surface term after using Stokes' Theorem, which is just moving the d down to a boundary operator ∂ .

[Remember, wedge star acts like a dot product.]

[The trick in doing this is to guess where you want to get a d into the expression.]

The volume integral involves $d * E$, the divergence of the true solution. We do not allow free charge inbetween the conductors, and so this vanishes as well, since

$$d * E = 4\pi\rho = 0.$$

Now we notice that the quadratic term is just like our energy term. Thus it is everywhere positive. This leads to a true inequality if we drop it:

$$C \leq \frac{1}{4\pi V^2} \int E_e \wedge *E_e,$$

with a quadratic error.

You can pursue a similar calculation for the case where you guess not E but D , picking D so that the correct amount of charge is on the conductors. This leads to a bound going the other way:

$$C \geq \frac{4\pi Q^2}{\int D_e \wedge *D_e}.$$

Example: Look at a square conductor of width 2 inside a larger grounded square of width 4. We will make two very crude guesses just so you get the idea. In two dimensions the conformal symmetry means that the capacity is a dimensionless number, and we needn't specify units for the lengths.

We guess equipotentials to be squares. We will do the integrals over just one of the eight symmetrical octants, the one from 1:30 to 3:00 o'clock. The E_e field is

$$E = dx.$$


We don't have to pick any particular size for the electric field because of linearity. The scaling of E and V cancel out.

$$*E = \widehat{dx} dy.$$

The energy density is

$$dx \wedge \widehat{dx} dy = \widetilde{dx} dy,$$

and so the energy integral is just the area of the octant, which is $\frac{3}{2}$. The voltage between the conductor and the



[Scandalous degree of crudeness, eh what?]

ground plane is $V = -1$. Thus the estimate of the capacity is

$$C \leq \frac{8}{4\pi} \times \frac{3}{2} = \frac{3}{\pi}.$$


For the flux, let us pick

$$*E = \widehat{dx} dy, \quad y < 1.$$

That is, the flux is just confined to the square. This leads to an estimate

$$C \geq \frac{2}{\pi}.$$

Even these crude guesses at fields have narrowed down the capacitance to 50%. It also shows that the capacitance is finite even with the sharp corner. You could pursue a similar argument to show that even an infinitely thin rectangle would have a finite capacitance.



34. Magnetism

Magnetostatics is an all together more complex case than electrostatics. The equations look the same, except they have the 1-form and 2-form interchanged:

$$dH = 4\pi J,$$

$$dB = 0.$$

What is missing is the nice boundary condition corresponding to conductors being equipotential surfaces. The corresponding potential for magnetostatics is not a scalar potential but a 1-form potential:

$$B = dA.$$

Now first of all surfaces like this on which B pulls back to zero are nowhere near as common as conductors, surfaces on which E pulls back to zero. Either we talk about superconductors, or we talk about moderate frequency alternating magnetic fields, low enough in frequency so that radiation effects can be ignored, and high enough in frequency so that the skin depth is small and the fields do not penetrate the conductors. Cases where the currents are distributed throughout the conductor are also tractable, but even more complicated than the case where all the current is confined to the surface.

Even once we accept the idea of a perfect conductor, the extra complications involved in having surface currents rather than surface charge will be nontrivial. In fact, there are important topological consequences. The number of degrees of freedom in a configuration will not be, as in the electrostatics case, the number of separate conductors, but rather the total number of holes in the conductors. Spherical conductors surprisingly will have, in fact, no magnetic degrees of freedom, and a figure eight conductor will have two.

Energy

To compute the assembly energy of a system of currents distributed on a finite number of isolated perfect conductors we need to depart from the approximate equations above. The work done will depend on the small E field generated by

$$dE = -\frac{\partial B}{\partial t}.$$

This electric field is called the “back electromotive force,” and usually gets either neglected or treated as a mystery in undergraduate courses. A rather similar argument to the electrostatics case leads to the pretty expression for energy in the configuration in terms of field quantities

$$W = \frac{1}{8\pi} \int_{\Gamma} B \wedge H.$$

Since this is not an electrodynamics book, take this as a given.

[An easy derivation, in fact. There is an ok argument in Griffiths.]

To find the degrees of freedom in the system, let us convert this to surface integrals over all of the perfect conductors. First we bring in the vector potential A . We can do this because we have

$$dB = 0,$$

over all of space, even inside the conductors. The motivation for this should be clear. We need to get a d into the equation, then use integration by parts and Stokes’ Theorem to get surface integrals. Thus

$$W = \frac{1}{8\pi} \int_{\Gamma} dA \wedge H.$$

There are by assumption no currents in the region between the perfect conductors, so there we have $dH = 0$, and so we get the surface integral

$$W = \frac{1}{8\pi} \int_{\partial\Gamma} A \wedge H.$$

On the surface we have the pullback of the B field vanishing, so we have on the surface

$$dA = 0.$$

There will be a discontinuity in the H field at the surface, and this will require that there be a surface current on the conductor. The equation


$$dH = 4\pi J,$$

integrated around a little loop transverse to the surface relates the surface current K to the pullback of the H field:

$$K = \frac{1}{4\pi} \text{Pullback}(H).$$

We have not bothered to introduce any special notation for the pullback. The picture of this operation shows how natural it is, including the need to know on what side of the surface the H field comes from.

Figure 34-1. Surface current caused by the pullback of H .



This surface current is conserved. To see this, look at the integral of K around a closed loop in the surface. Break this up into a lot of the little loops used to define K in terms of H . This relates the integral of K to the integral of H over two loops, one just outside and one just inside the surface. The one inside is clearly zero. The one outside is also zero:

$$\int_{\partial\Gamma} H = \int_{\Gamma} dH = \int_{\Gamma} 4\pi J = 0.$$

where Γ is any surface spanning the inside of the loop which remains outside of the surface. Since the integral of K around any closed loop is zero, we have

$$dK = 0.$$

This is important. On the surface of the conductor, which is topologically non trivial, both K and A are closed.


The energy integral can be written completely in terms of surface quantities.

$$W = \frac{1}{2} \int_{\partial\Sigma} A \wedge K.$$

This is to be compared with the electrostatic case, which was just a simple sum of voltage time charge for each conductor. That simplification came about because the voltage was constant, and could be factored out of the integral over the surface, leaving the total charge as the result. Remarkably, the same thing happens here.

Let us deal only with a simple torus. There are two ways to cut the torus that will remove the hole. These are sketched in the next figure.

Figure 34-2. Two cuts (cycles) that remove the hole in the torus without cutting it apart.



Neither of these has a boundary in the surface. The first of them has a boundary inside of the conductor. That implies that the integral of B over that surface is zero, and hence that the integral of A around that cycle vanishes. For the other cut, we want to see that the net current across that cut is zero. A net current across it would be a net integral of H just outside the conductor. But the integral just outside is the boundary of a region of space, albeit not on the surface. On this surface there is no current, hence that integral must vanish.

We are left with the possibility of net current K across the first cut, and a net A around the second. These are the two conjugate variables that correspond to voltage and charge. They are called current and magnetic flux. Now we show that we can factor the energy integral into the form

$$W = \frac{1}{2} \int A \int K.$$

In the surface integral we can move the current around as long as we keep the same net current around the ring and we will have the same

energy. To see that look at the difference of the two currents $K_1 - K_2$. Since this has no net integral around the ring, it can be written in terms of a potential; integrate by parts and this D either will hit the A term or the boundary of the surface. Both vanish. Exactly the same argument can be made for the vector potential.

Thus we can confine the current to a small strip around the ring, and the vector potential to a small strip through the center. They can both be made constant in the small rectangle where they cross, and hence the integral factors. See the figure for the modified current and vector potential.




Figure 34-3. Deformed current and vector potential shown on flattened torus.

Thus we can write

$$W = \frac{1}{2} I \Phi,$$

where I is the total flux and Φ the net flux.

One can now follow exactly the same path as with capacitance to find upper and lower bounds on the inductance.

35. Maxwell's Equations in Spacetime

We can also find a fully four dimensional formulation of electrodynamics. For this we use a four dimensional version of the star operator. I will denote it by $*_4$ to keep things straight.

Example:

$$*_4 dx = \widehat{dx} = \widehat{dx} dy dz dt = (dy dz dt, \{dx dy dz dt\}),$$

$$*_4 dy = \widehat{dy},$$

$$*_4 dz = \widehat{dz},$$

$$*_4 dt = -\widehat{dt}.$$

A minus sign in the last equation because dt is a timelike 1-form.

$$*_4(dx dy) = \widehat{dx} \widehat{dy},$$

$$*_4(dz dt) = -\widehat{dz} \widehat{dt},$$

and so on.

We can assemble a 2-form for half of the components of the electromagnetic field by taking

$$F = B + E \wedge dt.$$

We are using dx both as a basis 1-form in 3D and in 4D. If we now take the four dimensional exterior derivative of this, we find

$$d_4 F = dB + dt \wedge \frac{\partial B}{\partial t} + dE \wedge dt.$$

Note that when you d the new forms appear at the front, that is why the dt is first in its term.

$$d_4 F = -dt dE + dE dt = 0.$$

[Be careful in switching the two factors. Moving a 1-form, dt , through a 2-form dE involve two sign changes.]

We have to be more careful with the twisted forms. We start by taking the star of our field 2-form:

$$\begin{aligned} *_4 F &= *_4 B - *_4(dt E), \\ &= *_4 B + \widehat{dt} *_4 E. \end{aligned}$$

This leads us to define the twisted 2-form

$$G = H + \widehat{dt} D.$$

[Minus sign because there is a timelike 1-form being starred.]

Example: Suppose that we have in 3D

$$\begin{aligned} D &= t \widehat{dx}, \\ H &= x \widehat{dx} \widehat{dx}. \end{aligned}$$

Then

$$G = x \widehat{dx} \widehat{dx} + t \widehat{dt} \widehat{dx}.$$

We take the exterior derivative of this. Remember that we work from right to left with twisted forms.

$$d_4 G = \widehat{dx} - \widehat{dx} = 0.$$

We take the 4D exterior derivative, being very careful with the signs

$$d_4 G = dH + \widehat{dt} dD - \frac{\partial D}{\partial t}.$$

[You may have to work this out several times until you have made an even number of sign mistakes.]

Using Maxwell's equations

$$d_4 G = \widehat{dt} 4\pi\rho + 4\pi J.$$

The right hand side is the twisted current 3-form

$$j = \widehat{dt}\rho + J.$$

Example: For a unit charge density we have

$$\begin{aligned} \rho &= \widehat{1}, \\ j &= \widehat{dt}, \end{aligned}$$

and for unit current in the x direction

$$j = \widehat{dx}.$$

Thus we have Maxwell's equations in 4D:

$$\begin{aligned} d_4 F &= 0, \\ d_4 G &= 4\pi j, \\ *_4 F &= G. \end{aligned}$$

In a material with dielectric or magnetic properties we would have to change the third equation.

CHAPTER SEVEN. CLASSICAL MECHANICS

Here I just sketch without any derivation the geometric structure of classical mechanics. The idea is to show how the ideas adapt to the geometry of mechanics without getting into the details.

The most geometrical formulation of mechanics is the Hamiltonian form when time does not appear explicitly. The geometry here is the symplectic geometry of a 2-form on phase space. Even the less geometric Lagrangian form fits well with our language, however.

I give just a brief sketch of continuum mechanics, showing how even a theory with considerable dependence on a symmetric tensor can be geometrized with differential forms, although the extreme simplicity of the Navier Stokes equations seems like a miracle to me.

36. Hamiltonian Mechanics

We saw that the structure of a metric space was given by a symmetric second-rank tensor. Classical mechanics has its structure determined in part by an antisymmetric second-rank tensor, a 2-form. The space is phase space, the space of generalized coordinates q_i and generalized momenta p^i . The 2-form is given by

$$\Omega = \Sigma_i dp_i \wedge dq^i.$$

It is called the canonical 2-form. Note that it is a closed 2-form. The equations of motion of a given dynamical system follow from the Hamiltonian function, which is a function on phase space, according to

$$u \lrcorner \Omega = -dH.$$

Here u is the vector velocity of the system as it moves in phase space. The parametrization need not be time; often time is just one of the coordinates in phase space.

If we write the velocity vector in a two dimensional phase space

$$u = \dot{q} \frac{\partial}{\partial q} + \dot{p} \frac{\partial}{\partial p},$$

where the symbols \dot{q} and \dot{p} are just convenient mnemonic names, not time derivatives. This follows our convention that lower case letters are independent variables. Thus

$$\begin{aligned} u \lrcorner dp \wedge dq &= \dot{p} dq - \dot{q} dp, \\ &= -\frac{\partial H}{\partial q} dq - \frac{\partial H}{\partial p} dp. \end{aligned}$$

The system will move along integral curves of this vector field, given by maps

$$t \mapsto (q, p) = (Q(t), P(t))$$

such that

$$\begin{aligned} \frac{dQ}{dt} &= \dot{q} = \frac{\partial H}{\partial p}, \\ \frac{dP}{dt} &= \dot{p} = -\frac{\partial H}{\partial q}. \end{aligned}$$

[We switch to captial letters for dependent variables.]

It would be easy to go back and add summation signs to these equations and produce a calculation for any dimension.

Integral Invariants

The canonical 2-form can be written in terms of a canonical 1-form

$$\Omega = d\theta = d(\sum_i p_i dq^i).$$

The action is the integral of this canonical 1-form around any closed loop in phase space

$$S = \int_{\partial\Sigma} p dq = \int_{\Sigma} dp \wedge dq.$$

The change in this integral as the curve is moved by the equations of motion is given by the Lie Derivative

$$\begin{aligned} \dot{S} &= \int_{\Sigma} \mathcal{L}_u dp dq, \\ &= \int_{\Sigma} d(u \lrcorner (dp dq)), \\ &= \int_{\Sigma} d(dH) = 0. \end{aligned}$$

The value of this integral is invariant.

One can derive other integral invariants. The most important is found from

$$\int \Omega \wedge \Omega \wedge \dots$$

with enough Ω s to make this a volume integral in phase space. The invariance of this integral is usually called Liouville's Theorem. It lies at the heart of the second law of thermodynamics.

37. Lagrangian Mechanics

There is also a nice geometric picture of Lagrange's equations in mechanics. These are

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}} \right) - \frac{\partial L}{\partial q} = F,$$

where $L = T - V$ is the Lagrangian function, a scalar function on the space of positions, velocities, and time. I have left off an index on all of these quantities for clarity. It should read:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}^i} \right) - \frac{\partial L}{\partial q^i} = F_i.$$

The F_i are generalized forces in the sense of the principle of virtual work. The work done is

$$dW = \sum_i F_i dq^i$$

in the sense of p.v.w. Of course we interpret the right hand side as a 1-form, and recognize the F_i as the components of a 1-form.

To geometrize these equations we need to identify the intrinsic structure of the space (q, \dot{q}, t) . If we look at a possible motion of a particle in (q) space:

$$t \mapsto (Q(t)),$$

then this leads naturally to a curve in (q, \dot{q}, t) space:

$$t \mapsto \left(Q(t), \frac{dQ}{dt}(t), t \right)$$

called the "lift" of the original curve. Not all curves are lifts. Those curves which are lifted from position space have the 1-forms

$$\alpha = dq - \dot{q} dt$$

all pull back to zero on the curve.

The geometric structure of Lagrange's equations is that the system moves in (q, \dot{q}, t) space along curves for which the above forms pull back to zero as well as the forms

$$d \left(\frac{\partial L}{\partial \dot{q}} \right) - \frac{\partial L}{\partial q} dt.$$

[You can see how this enforces the requirement that the q-dot coordinate actually be the time derivative by just dividing through by dt.]

This provides us with $2n$ 1-forms in a $(2n+1)$ -dimensional space. This will determine unparametrized curves through the space. Time is just a coordinate here, except insofar as it occurs in a simple manner in $dq - \dot{q} dt$.

Example: For a free particle in one dimension we have the three dimensional space (q, \dot{q}, t) , with a Lagrangian function

$$L = \frac{1}{2}\dot{q}^2,$$

and two 1-forms

$$dq - \dot{q} dt,$$

$$d\dot{q},$$

determine the path. These describe tubes along which the system point must move. See the figure.




Figure 37-1. Free particle motion.

In this formulation it is not apparent that the terms involving the Lagrangian function are really independent of the coordinates. The easiest way to see that is to recognize that these are the Euler-Lagrange equations for a variational problem, extremizing the integral

$$\int L dt.$$

You could have motivated the choice of velocity space by asking the question: on what space is $L dt$ truly a 1-form? It certainly isn't in (q, t) space, for example.

38. Fluid Mechanics

While any partial differential equation can be cast into the language of differential forms by going to a phase space that is large enough, it requires a special structure to the equations for them to be described by differential forms on 3-space itself. Here we see how the Navier-Stokes equations can be so written. This is a remarkable simplification. It allows these equations to be written in curvilinear coordinates without using covariant derivatives.

Two natural geometric objects come to mind when you think about the state of a fluid: the tangent vector describing the local velocity of a fluid particle, and the 2-form representing the mass flux. Our representation scheme will use this mass flux 2-form and its associated mass density 3-form. We are considering using curvilinear coordinates for the spatial dimensions, and treat time as a parameter, just as in the usual Maxwell's theory. In this way the thermodynamic state variables are all scalars. I will ignore them for this brief discussion. Also, fluids with intrinsic spin angular momentum will be ignored, and so we do not need separate equations for the angular momentum budget.

To keep as close to familiar representations as possible, I will write the mass density $*\rho$ so that ρ will be an ordinary scalar field. The mass flux will be a 2-form that I will write as $*\rho\gamma$. In rectangular coordinates the 1-form γ is given by

$$\gamma = v_x dx + v_y dy + v_z dz,$$

where v_x is the x -component of the fluid velocity, and so on. The conservation of matter is written

$$*\frac{\partial\rho}{\partial t} + d(*\rho\gamma) = 0.$$

Similar balance laws can be written for the other scalar variables such as entropy.

The interesting balance laws involve momentum, since it is a vector quantity. It would seem to require covariant derivatives. To discuss the balance of momentum we need expressions for its density, and for its transport by convention, pressure, and viscous stresses.

Since the language of forms includes only antisymmetric quantities, the treatment of the symmetric stress-energy tensor will take some circumlocution. We must write separate balance laws for each of the momentum components, x -momentum, and so on. While this might seem to be geometrically barbaric, note that these conservation laws depend upon the symmetries of the metric, and such symmetries are represented by vector fields (called Killing vector fields) such as $\frac{\partial}{\partial x}$.

The density of k -momentum is just the mass density 3-form $*\rho$ times the k -component of the velocity,

$$(k \lrcorner \gamma).$$

Pressure transfers x -momentum in the negative x -direction, and so on. This flux of momentum will be given by

$$k \lrcorner *p.$$

The momentum convected with the fluid will be represented by a 2-form aligned in the $*\gamma$ direction. The correct flux will be quadratic in velocity, and is

$$*(\rho(k \lrcorner \gamma)\gamma).$$

At this point we can write down the Euler equations, which describe the momentum balance for the case of no viscosity. The conservation of k -momentum requires

$$*\frac{\partial}{\partial t}(\rho(k \lrcorner \gamma)) + d[\rho(k \lrcorner \gamma) \wedge (*\gamma) + k \lrcorner *p] = 0.$$

Now one can manipulate these equations and make a marvelous simplification. Here one can write these scalar equations in the form of the vector k contracted into a 1-form. The calculation uses an unfamiliar identity that I will give later. Using this identity along with some standard forms manipulations leads to the Euler equations in the form

$$k \lrcorner \left[\rho \frac{\partial \gamma}{\partial t} + \frac{\rho}{2} d|\gamma|^2 + dp + *\rho(*d\gamma \wedge \text{gamma}) \right] = 0$$

Now there are three Killing vectors for the three translations, and these are linearly independent at every point. Thus the 1-form in brackets must vanish. Thus we have the Euler equations

$$\rho \frac{\partial \gamma}{\partial t} + \frac{\rho}{2} d|\gamma|^2 + dp + *\rho(*d\gamma \wedge \text{gamma}) = 0.$$

There you see the Euler equations written in curvilinear coordinates without the use of covariant derivatives. In fact, the $*$ operator is even outside of all the differentiations.

We now add the viscous stresses to this. The easiest approach is to write out all possible 2-forms involving the velocity 1-form γ , the Killing vector k , one differentiation d , and one $*$ (to match the parity of the other terms). Then one tries to construct from them the flux of k -momentum known to describe viscosity.

The possible terms are

$$*d(k \lrcorner \gamma), *k \lrcorner d\gamma, \text{ and } k \lrcorner d* \gamma.$$

We can now write these out in rectangular components and compare with the usual expression for viscous momentum flux. Using the usual viscosity coefficients, we find that the flux of k -momentum caused by viscous stress is

$$\eta(*k \lrcorner d\gamma + 2*d(k \lrcorner \gamma)) + (\zeta - \frac{2}{3}\eta)k \lrcorner d* \gamma.$$

These give the correct expressions for the flux of linear momentum, but not for angular momentum. Some of the differentiations hit the Killing vector, which is only constant for translations. The above is not the stress tensor, since it is not function-linear in k .

To find the Navier-Stokes equations we have to add the divergence of this flux to our system. This flux is fairly easy to find using the identities given below. We are forced to assume that the 1-form κ which is the covariant form of k satisfies $d\kappa = 0$. This is true for the translations. The resulting equation is

$$\frac{\partial \gamma}{\partial t} + \frac{1}{2}d|\gamma|^2 + *(d\gamma \wedge \gamma) + \frac{dp}{\rho} + \frac{\eta}{\rho} *d*d\gamma - \frac{3\zeta + 4\eta}{3\rho}d*d* \gamma = 0.$$

[Warning! Here I have not properly distinguished between star and its inverse.]

For incompressible fluids we have the usual Navier-Stokes equations

$$\frac{\partial \gamma}{\partial t} + \frac{1}{2}d|\gamma|^2 + *(d\gamma \wedge \gamma) + \frac{dp}{\rho} + \frac{\eta}{\rho} *d*d\gamma = 0.$$

For some uses it is more convenient to write this as a Lie derivative, using the velocity vector v

$$v = \lrcorner \gamma.$$

We have

$$*(d\gamma \wedge \gamma) = v \lrcorner d\gamma,$$

and so

$$\frac{\partial \gamma}{\partial t} + \mathcal{L}_v \gamma - \frac{1}{2} d|\gamma|^2 + \frac{dp}{\rho} + \frac{\eta}{\rho} * d * d\gamma = 0.$$

In terms of the vorticity 2-form

$$\omega = d\gamma$$

we have

$$\frac{\partial \omega}{\partial t} + \mathcal{L}_v \omega + \frac{\eta}{\rho} d * d * \omega = 0,$$

from which Kelvin's circulation theorem follows for $\eta = 0$

$$\mathcal{L}_{\left(\frac{\partial}{\partial t} + v\right)} \omega = 0.$$

For some applications it is convenient to use time-dependent curvilinear coordinates, for example when considering rotating fluids, or Lagrangian coordinates. In this case we must introduce an additional spatial vector field λ giving the velocity of the coordinates relative to inertial space. The time differentiation must now be changed according to

$$\frac{\partial \gamma}{\partial t} \mapsto \mathcal{L}_{\left(\frac{\partial}{\partial t} + \lambda\right)} \gamma = \frac{\partial \gamma}{\partial t} + \mathcal{L}_\lambda \gamma.$$

[I am not totally convinced by the above paragraph, myself.]

Special Identity

Until I tried to derive the Euler equations I was unacquainted with the following identity; nor can it be found in the common books.

$$\frac{1}{2} k \lrcorner d|\gamma|^2 = *(* \gamma \wedge d(k \lrcorner \gamma)) - k \lrcorner (* (* d\gamma \wedge \gamma)).$$

There are some peculiar signs that arise in higher dimensions. This is only true for three dimensional Euclidean space. Here k must be a Killing vector.

CHAPTER EIGHT. DISPERSIVE WAVES

The phenomena of group velocity and phase velocity are nicely geometrized once you realize that group velocity is a tangent vector and that phase velocity is a 1-form.

In addition, the picture of a wave packet of dispersive waves provides a nice geometric picture, much needed, of the momenta in Hamilton's equations.

39. Dispersive Waves

[There is a liesurely treatment of this, with special attention to waves on deep water, in my book Spacetime, Geometry, Cosmology.]

We want to develop a geometric description of waves which are perfectly coherent. By perfect coherence, we mean that you can define a global phase function. Thus you can describe the wave with functions of the form

$$\psi = A(t, x) \cos(2\pi\theta(t, x)).$$

Both the amplitude and the phase are functions of position. We will only consider the case where the amplitude and the phase gradients change slowly with respect to the phase itself.

[These are the tamer regions of the WKB approximation.]

Example: If you look at the pattern of light that you see on the bottom of a swimming pool, you see a wave field with only partial coherence. If you follow a minimum in the light, it will not be continuous, but disappear in an ambiguous manner. A coherent wave lets you follow the crests and troughs forever, and they are conserved.

Phase velocity

The most important geometric property of a coherent wave field is the alignment of the wavecrests and wave troughs in spacetime. We will be building a theory that nowhere depends on the metric, and we will be indifferent as to whether we are in several spatial dimensions or a spacetime situation. For clarity I will always speak of it as if it were spacetime. When we discuss optics in a later section, we will be in a purely spatial situation. This disposition of the wave throughs in spacetime is specified by the gradient of the phase θ :

$$d\theta = \nu dt + k dx.$$

I use ν instead of ω to go with the 2π that I have inserted into the definition of the phase. I do this to follow the Fourier transform conventions of Bracewell. Fourier theory is needed if you are to consider situations of partial coherence.

A major difference between what we do here and what you will see in all of the books is that I am treating all coordinates equally.

Most of the time you will see plane waves written

$$\psi = \cos(2\pi(\nu t - kx)),$$

in which time and space are treated differently. Since we want to proceed smoothly between spaces of different dimension and signature, we should not do this. This will bring in a minus sign that is avoided by the above barbarity, and our waves with positive frequency will have a negative wave number when they are right-going.


You need to exercise a little care with this representation. The wave with both ν and k reversed in sign represents the same wave.

The geometric interpretation of ν and k is that they are densities. ν is the density of wavecrests along the time axis. k is the density of waves along the space axis. The phase gradient is then the flux of wave-crests 1-form, and the identity

$$dd\theta = 0,$$

can be interpreted as the conservation law for wave crests.

Figure 39-1. The basic geometry of the wave crests in spacetime.



The geometrically unaware could try to describe the geometry of wavecrests by defining a vector that follows a wave crest in spacetime. The folly of this is apparent if you try to do it with even one more dimension. This is what makes the conventional treatment of phase

velocity a dead end, and why we need to recognize that phase velocity is a 1-form.

Dispersion Relation

The phase gradients are constrained by the wave equation that describes the waves. This relation is called a dispersion relation, vaguely reflecting its origin in optics. If the wave equation has constant coefficients, then you can find the dispersion relation by just inserting a plane wave solution into the equation.

Example: The wave equation for an elastic beam with compression is, up to irrelevant dimensioned constants

$$\psi_{tt} + \psi_{xx} + \psi_{xxxx} = 0.$$

This leads to a dispersion relation

$$-4\pi^2\nu^2 - 4\pi^2k^2 + 16\pi^4k^4 = 0.$$

Interactions

If two waves are to interact via some linear relation, then the fact that different frequencies are linearly independent gives us the requirement that the interacting waves must have their phase gradients equal up to a sign when they are pulled back to the set on which they interact. You can use this to find the relations across an interface which leads to Snell's Law in optics, or the relation which determines the antenna pattern for a linear antenna in radio waves, or the discrete Laue conditions that apply across the discrete lattice in x-ray diffraction.

40. Group Velocity

For dispersive waves all of the information about the frequency of the wave, its wavenumber, and its amplitude is propagated along curves called characteristics, whose tangents are the group velocity vectors. This situation has a very nice geometric picture provided we use both the space of tangent vectors and the less familiar space of 1-forms.

[This is closely related to what is called the reciprocal lattice in solid state physics.]

Geometric Derivation

We argue at first in two dimensions. Look at two plane wavetrains with nearly the same phase gradient. The two sets of wave crests form a Moire pattern in spacetime. Physically, there are beats between the two different waves.




Figure 40-1. The Moire pattern between two wavetrains defines the group velocity direction.

The line of beats is the line along which amplitude information propagates. Since time is just another coordinate, there is no particular length for the tangent vector, and we are free to choose any


convenient normalization. We choose for the group velocity u

$$d\theta \cdot u = 1.$$

In the space of vectors this construction looks like the next figure.

[The technical name for this space is the tangent space, and the space of 1-forms is called the cotangent space.]

Figure 40-2. The relation between phase and group velocity shown in the space of tangent vectors.



An similar construction holds in the space of 1-forms. Note the change in the labels, however.

Note that in the space of 1-forms, a 1-form is represented by a point. To make the point obvious we can draw an arrow from the origin to that point, just as we do with vectors. In that space vectors are represented by the set of 1-forms α such that $\alpha \cdot u = 1$. The duality between the spaces is nicely seen in the duality between these two pictures.

Let me remove extraneous lines from Figure 40-1 and redraw just the essence, the two phase gradients and the group velocity vector.

[Actually, what I have done is to draw the picture not in spacetime but in the tangent space to spacetime.]

If we now draw our two beating wavetrains in 1-form space, we get the following figure.

Figure 40-3. The same relation shown in the space of 1-forms. This is a numerically correct translation of the preceding figure. The arrows on the axes are all of unit length.






Figure 40-4. Two beating wavetrains drawn in the space of vectors.



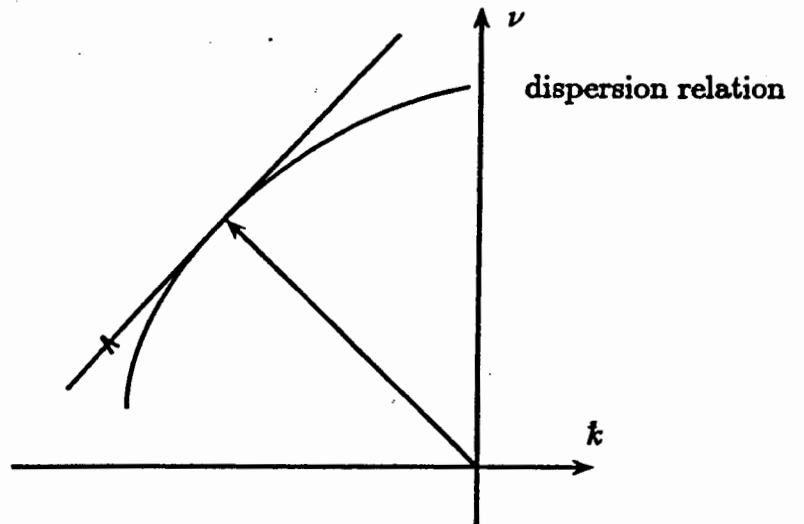
Now we want to take the limit as the two wavetrains get closer together in the space of 1-forms. That is, their frequencies and wavenumbers approach each other. Since both wavetrains are to be solutions of whatever wave equation we are considering, both lie on

Figure 40-5. Two beating wave-trains in the space of 1-forms.



the dispersion relation. Clearly the group velocity will be tangent to the dispersion relation in the limit. The general pattern is shown in the next figure.

Figure 40-6. The limit as the wave-trains merge.



In two dimensions this is the usual result. When the dispersion relation is given by the specifying the frequency as a function of

wavenumber the speed is dv/dk .

In more dimensions you just need to draw a tangent plane to a hypersurface and properly arrange the normalization. In three dimensions we have

$$u = \left(\frac{\partial W}{\partial k} \frac{\partial}{\partial x} + \frac{\partial W}{\partial l} \frac{\partial}{\partial y} + \frac{\partial W}{\partial \nu} \frac{\partial}{\partial t} \right) / \left(\frac{\partial W}{\partial k} k + \frac{\partial W}{\partial l} l + \frac{\partial W}{\partial \nu} \nu \right).$$

Algebraic Derivation

The above geometric discussion is very good for showing you how to geometrize the idea of group velocity. It is not really convincing unless you already know that group velocity exists. Here is a straightforward argument, unfortunately lacking in geometric intuition.

We start with a dispersion relation

$$W(k, \nu) = 0.$$

We are not letting the dispersion relation depend on x or t . We take k and ν to refer to a specific wavetrain solution, then they will be functions of x and t , and we can differentiate the above relation

$$\begin{aligned} \frac{\partial W}{\partial k} \frac{\partial k}{\partial x} + \frac{\partial W}{\partial \nu} \frac{\partial \nu}{\partial x} &= 0, \\ \frac{\partial W}{\partial k} \frac{\partial k}{\partial t} + \frac{\partial W}{\partial \nu} \frac{\partial \nu}{\partial t} &= 0, \end{aligned}$$

Now recall that k and ν are already partial differentials, the gradient of the phase, so we have

$$\frac{\partial k}{\partial t} = \frac{\partial \nu}{\partial x}.$$

This lets us write the above equations

$$\begin{aligned} \frac{\partial W}{\partial k} \frac{\partial k}{\partial x} + \frac{\partial W}{\partial \nu} \frac{\partial k}{\partial t} &= 0, \\ \frac{\partial W}{\partial k} \frac{\partial \nu}{\partial x} + \frac{\partial W}{\partial \nu} \frac{\partial \nu}{\partial t} &= 0, \end{aligned}$$

These equations state that both k and ν are constant in the direction of the vector

$$v = \frac{\partial W}{\partial k} \frac{\partial}{\partial x} + \frac{\partial W}{\partial \nu} \frac{\partial}{\partial t}.$$

[This discussion comes from my book *Spacetime, Geometry, Cosmology*.]

This is the same group velocity vector, except for the overall length, that we had before.

Now we need to consider dispersion relations that depend on position in spacetime. Then the wavenumber and frequency will not be constant along the characteristics along the group velocity directions. There will now be further terms in the differentiations

$$\begin{aligned}\frac{\partial W}{\partial k} \frac{\partial k}{\partial x} + \frac{\partial W}{\partial \nu} \frac{\partial \nu}{\partial x} + \frac{\partial W}{\partial x} &= 0, \\ \frac{\partial W}{\partial k} \frac{\partial k}{\partial t} + \frac{\partial W}{\partial \nu} \frac{\partial \nu}{\partial t} + \frac{\partial W}{\partial t} &= 0.\end{aligned}$$

From these we can read off the rate of change along the group velocity direction. These equations are

$$\begin{aligned}\dot{k} &= -\frac{\partial W}{\partial x}, \\ \dot{\nu} &= -\frac{\partial W}{\partial t}.\end{aligned}$$

The dot is differentiation with respect to a parameter. This is not time, since we are treating all of the coordinates equally, but rather the wave action. We also have equations for the characteristic curves

$$\begin{aligned}\dot{x} &= \frac{\partial W}{\partial k}, \\ \dot{t} &= \frac{\partial W}{\partial \nu}.\end{aligned}$$

From these we see the remarkable result that wavepackets follow the rules of Hamiltonian dynamics.

Wave Diagrams

Return to the geometry. Every wavetrain corresponds to a point on the dispersion relation, and the tangent plane at that point determines a vector, the group velocity vector. You can use this to map the dispersion relation surface point by point into the tangent space. I call the resulting surface a wave diagram.

[This is my own invention, and you will not find this general construction or the name elsewhere.]

The wave diagram is interpreted in a dual manner to the dispersion relation. Points on the wave diagram correspond to wavetrains with specified group velocity. The tangent to the wave diagram shows you the phase gradient of that wavetrain.

Example: For the elastic beam with

$$W = \nu^2 + k^2 - k^4.$$

The phase gradient is

$$d\theta = \nu dt + k dx,$$

and the group velocity, properly normalized, is

$$u = \frac{\nu}{2k^2} ddt + \frac{2k^2 - 1}{k} ddx.$$

That is

$$\dot{t} = \frac{\sqrt{k^2 - 1}}{2k}$$
$$\dot{x} = \frac{2k^2 - 1}{k}$$

See the next figure for the wave diagram. It is shown as the envelope of the phase gradients.




Figure 40-7. The wave diagram for the elastic beam. This is only schematic.

41. Water Waves

A nice, visible, example of the foregoing ideas about dispersive waves can be seen in the waves on the ocean: deep water waves in the linear approximation with surface tension neglected. You should probably read this section and the last one in parallel. New material on symmetries and conservation laws is introduced here.

Dispersion Relation

We can write the usual dispersion relation for water waves

$$\omega^2 = gk,$$

as

$$v^4 = k^2 + l^2.$$

We have picked new units for space to get rid of the physical constant, we have squared the dispersion relation to get waves going in both directions, and we have made the obvious extension to two spatial directions. Restore physical units by replacing lengths according to

$$L \mapsto \frac{L}{2\pi g}.$$

The phase gradient is

$$d\theta = v dt + k dx + l dy,$$

and this statement really serves to define k and l , including the sign that we will use.

Had we included surface tension we would have used

$$\omega^2 = k\left(g + \frac{T}{\rho}k^2\right).$$

[In Spacetime, Geometry, Cosmology I use this to make a physical model of special relativity.]

Wave Diagram

In (x, t) spacetime we can draw a wave diagram as follows. From the dispersion relation in the form

$$W = v^2 + k = 0$$

we find that the group velocity vector u

$$u = \dot{x} \frac{\partial}{\partial x} + \dot{t} \frac{\partial}{\partial t}$$

is given by

$$\dot{x} = -\frac{1}{k},$$
$$\dot{t} = \frac{2}{\sqrt{-k}}.$$

We can eliminate the parameter k

$$\dot{t}^2 - 4\dot{x} = 0.$$

The wave diagram is a parabola. The parameter k is negative, and these waves are all going to the right. The left-going waves would be described by the parabola

$$\dot{t}^2 + 4\dot{x} = 0.$$

The wave diagram in three dimensional spacetime (x, y, t) is just the surface of revolution formed by rotating the above parabola around the time axis.




Figure 41-1. The geometry for right-going water waves.

You can see here the famous result that the group velocity is half the speed of the wavecrests. Note also that longer waves travel faster than shorter ones. One direct result of this is that long sail boats sail faster than short ones.

Chirp

In the winter the Pacific coast is treated to large swells that come from storms in the South Pacific, beyond Hawaii. Over a few days the period of the waves will decrease, from around 17 or 18 seconds to 8 seconds or so. Then it will jump up again as waves from a new storm dominate.

You can easily find the distance to the storm, and verify the surprising result that this wave energy comes from thousands of miles away, just by measuring the rate of change of the frequency, called the chirp, of the wavetrain. Waves with wavenumber k travel a distance L in a time T

$$T = 2L\sqrt{-k} = 2L\nu.$$

We use the group velocity since it is the energy that we want to track, not the wavecrests. Thus waves with frequency ν are received at a time t

$$\nu(t) = \frac{t}{2L}$$

and the chirp is

$$\frac{d\nu}{dt} = \frac{1}{2L}$$

and so

$$L = \frac{1}{2\dot{\nu}} = \frac{g}{4\pi\dot{\nu}}.$$

I have restored the physical units in the last equation.

Symmetry

Dispersive waves follow the laws of Hamiltonian dynamics. Thus we have a conservation law associated with every symmetry. We only need this in its simplest form: if a coordinate does not appear in the dispersion relation, then its associated wavenumber is constant along the group velocity ray.

You can use this to estimate the direction that the ocean waves discussed above really came from. The direction near the beach is severely affected by refraction as the depth becomes comparable to the wavelength. In fact, the refractive effects make all the waves

appear to come in perpendicular to the beach. This makes beaches effective traps for wave energy. Wave which reflect off of rocks and local irregularities are bound by refraction to bend back to the shore again and not escape to "infinity".

If the topography is symmetric, that is, we are treating the case of a straight beach, the the wavenumber along the beach is conserved. In addition, the frequency of the wave is conserved. From this frequency you can find the deep-water wavelength. Thus, in deep water, beyond where you can see, we have

$$k^2 + l^2 = \nu^4.$$

Where this ν can be measured on the beach since it is invariant. The wavenumber k can also be measured on the beach, assumed to run along the x -axis. Thus the deep water angle ϕ is given by

$$\sin^2 \phi = \frac{k^2}{k^2 + l^2},$$

that is

$$\sin \phi = \frac{k}{\nu^2} = \frac{4\pi g k}{\nu^2}.$$

Again I have restored dimensioned variables in the last equation.

Partial Coherence

One striking observation is that the waves coming in to the beach are grouped into sets with a statistical regularity. A few larger waves will be followed by a few tiny ones, and so on. The surfers say roughly that every fifth wave is the biggest. This is only true statistically.

The physical effect is cause by there being a mixture of frequencies present at the beach at any one time: short, slow waves from early on in the distant storm combined with long, fast waves from later in the storm. If the storm duration is T , and its distance L , then the spread of frequencies is

$$\Delta\nu = \frac{T}{2L}.$$

You need the previous result on the chirp to estimate both T and L . The spectral purity is

$$\frac{\Delta\nu}{\nu} = \frac{T}{2L\nu} = \frac{\pi g T P}{L}.$$

where P is the period. From this you can estimate the number of waves per packet.

APPENDICES.

APPENDICES

42. Forms Calculator

Here is a calculator implemented in Mathematica to simplify expressions involving differential forms. As the name implies, the intent was to make a minimal implementation, with little attention paid to efficient input or pretty output. It will get the signs right for you, and the simplicity of the implementation will demonstrate the efficiency of a rule-based system for this kind of programming. Setting up such a system is a good way to debug your understanding of forms.

The key idea behind all of this is that Mathematica manipulates expressions as trees. It looks for patterns, and when it finds a match, substitutes something else for the pattern. When it finds no further matches, or none at all, it returns the expression. I will use the "fallen tree" representation for drawing trees in ascii text streams. This is the style used in the trn newsreader. The trees are written on their sides, hence my name for it. The vector expression

$$\frac{\partial}{\partial x} + 2\frac{\partial}{\partial y} + \frac{\partial}{\partial z}$$

would be written as an expression tree in Mathematica which I will write in this notation

```
Plus--Vector--x
  |-Times--2
  |   \-Vector--y
  \-Vector--z
```

Coordinates

You might want to start working with forms in three dimensional Euclidean space. There are special definitions to simplify the input and output for this case in the file `euclid3.m`. The coordinates there are x , y , and z . It is surprising how many of the rules for forms do not depend on any knowledge of the full list of coordinates. I didn't realize this until I developed this system.

The shorthand names for the basis forms defined in `euclid3.m` are `dx`, `dx dy`, `dx dy dz` and so on. No spaces here, Mathematica treats these as a single symbol. The basis vectors are given by `px`, and so on. The "absent" symbol \hat{dx} is denoted by `tx`.

There is also an output function `Efilter` that you can use to beautify the output, setting

```
In[] := $Post = Efilter
```

Vectors

A vector will be a linear combination of monomial terms. Each of these monomials will be a tree with one branch. The head of the tree is a node labeled "Vector". The symbol on the branch indicates the direction of the vector relative to the coordinate system. No rules are needed for these monomials to form a vector space.

Forms

Differential forms will be linear combinations of monomial terms. This time there will be a number of branches; an r -form will have r branches below a node labeled "Form". The 2-form $dx dy$ would be represented by the tree

```
Form--x
      \-y
```

and written `Form[x,y]`. The shorthand in `euclid3.m` would let you write this `dx dy`, no space.

We need to explicitly accomodate the fact that a 0-form is a number with a rule

```
Form[] := 1
```

The antisymmetry allows us to simplify expressions involving forms. This is not done continually, in the interests of efficiency, but the following rule is invoked as needed:

```
FSimp := Form[ a__ ] :=>
      Signature[ Form[a] ] Sort[Form[a]]
```

Duality

The operation which contracts a vector with a form is denoted by the operator "angle".. The expression

$$a \frac{\partial}{\partial x} b dx dy dz$$

is represented by

```
angle--Times--a
      |          \-Vector--x
```

```

\Times--b
  \Form--x
    |-y
    \-z

```

There are rules needed to enforce the bilinearity of this operation, such as

```

angle[ a_, b_ + c_] := angle[a,b] + angle[a,c]

```

which is the replacement

```

angle--a
  \-Plus--b      ->  Plus--angle--a
    \-c          |      \-b
                \-angle--a
                  \-c

```

The essential rule replaces

```

angle--Vector--u
  \-Form--u      ->  \-Form--a
    |-a          |      \-b
    \-b

```

with due attention to the signs.

There are also rules needed to take care of the fact that we represent the zero vector and the zero forms by the number zero.

Wedge Product

The wedge product is denoted by the tree

```

wedge--Form
  |-Form
  ...
  \-Form

```

I have put in some rules that let you denote the wedge product with `NonCommutativeMultiply`, `**`, but these are not fool proof at present. Intuitively you expect `**` to bind more weakly than `*`, but that can't be done in `Mma` as I understand it. So it gets confused on

```

dx ** 3 dy

```

Exterior Derivative

The exterior derivative operator is denoted `d[.]`. Functions need to have their arguments explicitly indicated, since the differentiations are done by the `Mathematica` operator `Dt`. Thus

```

In := d[ f[x,y] ]
      (1,0)          (0,1)

```

Out = f [x,y] Form[x] + f [x,y] Form[y]

Lie Bracket

The Lie bracket (Jacobi bracket) is represented by

Lie--Vector

\-Vector

and linear combinations.

Lie Derivative of Forms

This is represented by

Lie--Vector

\-Form

and linear combinations.

Twisted Forms

These are put in using the "absent symbol" notation. The twisted form

$$\hat{d}x\hat{d}y$$

will be represented by the expression tree

TForm--x

\-y

and written TForm[x,y]. The shorthand in euclid3.m would allow you to write this txt. The unit pseudoscalar is called tmax, and in Euclidean 3-space we have

$$tmax = TForm[x,y,z]$$

The Length of the variable tmax is used whenever a result depends upon the dimension of the space.

The wedge of an ordinary form with a twisted form in that order does not depend upon the dimension of the space, but the other two products do. Thus the exterior derivative of a twisted form can be computed without knowing the dimension of the space or the names of any coordinates beyond those given.

Pullback

Defined so far only for ordinary forms:

Pullback--Form

\-Rule

where the second argument is the set of rules that describe the map. For the x -axis as a parametrized curve, the rules would be

$$pc = \{x \rightarrow s, y \rightarrow 0, z \rightarrow 0\}$$

The points on the curve are found from the replacement

$$\{x, y, z\} /. pc$$

and the pullback of a form α onto the x -axis would be found from

$$\text{Pullback}[pc, \alpha]$$

Hodge Star

The present implementation is rather simple minded. It assumes that the coordinates are orthonormal, and that the coordinate t , if present, is timelike. Also both star and its inverse are denoted by star :

$$\text{star--Form} \quad \text{and} \quad \text{star--TForm}$$

Testing

There is a file of test expressions that has been used to validate the system, and that should be used if you modify the rules to ensure that you have not broken anything. The file is called `trip.m`, and will run when you load it after `forms.m`.

1-form 10-1
Divergence Theorem 22-2
Gaussian optics 9-2
Killing vector 28-3
Laplace's equation 32-1
Lie derivative 28-2
Liouville's theorem 36-2
Maxwell's equations 31-1
Stokes' Theorem 22-1, 24-1
Taylor's series 10-2, 7-2
basis 11-2
buffing wheel 6-2
calculus 7-1
canonical 2-form 36-1
capacitance 33-2
chain rule 19-2
closed 24-1
conformal mapping 32-2
contour map 10-1
contraction 17-1
cycle 24-1
density 14-1
derivative 7-1
dilation 8-1
displacement vector 5-1
duality 11-3
electrostatics 32-1
exact 24-1
flow 13-1
generalized coordinates 12-1, 36-1
gradient 10-1
hypersurface 19-3
inductance 34-5
invariance 5-1
lift 37-1
magnetostatics 34-1
orientations 18-1
paraxial approximation 9-2
phase space 36-1
pool 10-2
pullback 19-2, 19-2, 20-2
pushforward 19-1
reciprocity 12-1
rotation 8-1
shear 8-1
shorthand 17-1
star operator 27-1
symmetry 18-1, 5-1
tangent vector 6-1
tensor product 26-1
transverse orientation 19-3
uniqueness theorem 33-2
virtual work 37-1
wedge 14-5, 17-1
Möbius strip 22-2