–VERSION 7.1 –Mar 10, 1995–

TWISTED FORMS

Twisted Differential Forms as they Should Be

William L. Burke

Here are practical rules for calculating with differential forms, twisted and untwisted. It introduces what I have found to be the best and simplest notation for twisted forms. While differential forms are elegant, simple, and efficient, their use is hampered by numerous right-hand rules. To eliminate these requires the additional complexity of twisted forms, with further rules of operation. Here I have tried to make the simplifications involved in using twisted forms clearly outweigh the additional complexity.

I have found that the rules for calculations with differential forms can best be described by using expressions for the operators in terms of their actions on monomials of basis forms. The extension to general forms is immediate using the linearity of the operators. You could call these Cartesian differential forms were it not for the fact that they are independent of Euclidean geometry. We will also simplify matters by taking all surfaces of integration to lie along right-angled coordinate lines. This could be called the Manhatten theory of forms.

The ideas and notation used here are completely congruent with the method of moving frames used in General Relativity for efficient Riemann tensor calculations.

We will use two kinds of monomials here. One corresponds to a coordinate–derived basis.

> **Example:** The familiar three dimensional coordinates lead to the basis 1–forms dx, dy, and dz.

Such a basis is called *holonomic* because these basis forms are exact, that is, they are the gradients of functions.

We will write the generic basis 1–form as dq, and the exterior product of an indeterminant number of them as $d\alpha$ and $d\beta$.

The other monomials involve an orthonormal basis, and will be described a bit later.

Basic Operations

The operations of exterior multiplication, $\wedge$, exterior differentiation, d, and contraction, $\rfloor$, are all easily represented when acting on our monomials. In a monomial all the 1–forms are implicitly wedged together. Thus we have

$$d\alpha \wedge d\beta = d\alpha\, d\beta,$$

which is shorthand for relations such as

$$dx \wedge dy\, dz = dx\, dy\, dz.$$

The multiplication is antisymmetric under exchange of adjacent 1-forms.

Example: There is only one independent 3–form in three dimensions. We have

$$dx\, dy\, dz = -dy\, dx\, dz,$$

and so on.

For any function f, we have

$$d\,(\, f\, d\alpha\,) = df \wedge d\alpha,$$

wedged because df is not a monomial.

For contraction we must either put the relevant basis 1–form at the front

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

or take into account the sign changes that would occur. We use the shorthand

$$d\alpha\, dq\, d\beta = (-)^{\alpha}\, dq\, d\alpha\, d\beta,$$

where $(-)^{\alpha}$ for α an r–form is $(-1)^{r}$.

The Lie derivative $\pounds$ is given by

$$\pounds_{Q\frac{\partial}{\partial q}}\, f\, dq\, d\alpha = Q\frac{\partial f}{\partial q}\, dq\, d\alpha + f\, dQ \wedge \alpha.$$

Rendering Forms

One delight of using forms is that they have iconic representations that are faithful to the symmetries of the objects being represented. To the newcomer these icons have several unusual features: lines and surfaces of indefinite size, objects of undetermined shape, and peculiar methods for describing the negative of an object.

I don't know the full provenance of these "eggcrate" pictures, but they can all be found in Schouten's 1951 book. The simplest of these pictures is that of a 1-form. The 1-form can represent the gradient of a function (although not all 1-form fields are gradients). At a single point the 1-forms looks like the essence of a gradient: in two dimensions it is represented by two contour lines, with an arrow pointing in the uphill direction. In fact, only the arrowhead is needed here.

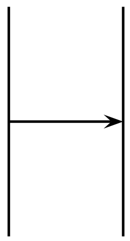

Figure 1. The representation of a 1-form in two dimensions, here the 1-form dx.

The picture of a 1-form has two lines of indefinite length. Another common icon is the one for a 2-form. In three dimensions, this involves a prism of indefinite length and indefinite shape. It does have to maintain, via the parallelogram rule, a definite area. Using our Manhatten rules, it looks unnecessarily right-angled.

The 2-form has its plus/minus-hood specified by a sense of rotation around the cylinder. We will specify this by listing the "snake" of vectors that describe, head-to-tail, this rotation. In the second figure, this sequence is $dydz$.

The picture of the rule for the contraction of a vector and a form is, in Manhattan monomial form, shown in Figure 3. The action of contraction is to chop the tail end of the snake.

In contrast, exterior differentiation, adds on to the tail of the snake.

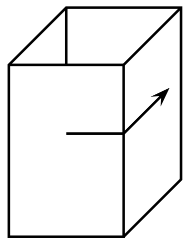

Figure 2. The representation of a 2-form in three dimensions. The 2-form here is $dydz$.

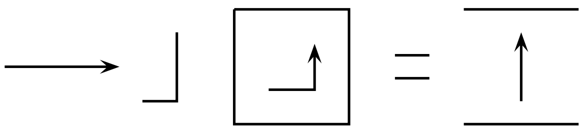

Figure 3. The representation of contraction.

Twisted Forms

Twisted r–forms have different symmetry properties from untwisted r–forms. We arrange this by defining them to be pairs, an ordinary r–form and an orientation of n–space. We represent the orientation of n–space by a ray of n–forms, which I write using curly brackets:

$$\{d\alpha\} \equiv \{k\, d\alpha\}, \qquad k > 0.$$

Curly brackets around any geometric object denote the symmetry equivalence class of the object. Abstractly, and to Mathematica, $\{dx\,dy\}$ is just a symbol with the property that

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

and we can think of this symbol as multiplying the ordinary differential form.

Example: The symmetry class of the vector $\dfrac{\partial}{\partial x}$ is $\{dx\}$. It is odd, that is, changes sign, under x reflection. It is even under y and z reflections.

The definition of this symmetry operation in general, beyond the Manhattan context here, requires some care.

Example: The two orientations of a two dimensional space are $\{dx\,dy\}$ and $\{dy\,dx\} = -\{dx\,dy\}$.

Example: An example of a twisted 1–form would be

$$(dx, \{dx\,dy\}) = -(dx, \{dy\,dx\}).$$

Unlike dx, this twisted 1–form is unchanged by reflections $x \mapsto -x$, and it does change sign if y is reversed.

I will sometimes use Ω to refer to an unspecified orientation. The conventional notation for twisted forms is to pick a reference orientation, $\{dx\,dy\,dz\}$ say, and refer all twisted forms to this orientation. Then one need not bother to write the orientation down explicitly. However, the definitions are clearer if we explicitly display the orientation rather than refer everything to a standard orientation. Thus I use what I call *William's Twisted Notation*. First let me introduce the hat notation for missing factors. This is fairly standard. An missing factor can be written with a hat over it, so that is absence is explicitly noted.

Example: Thus I could write the above twisted 1–form $(\widehat{dy}\,dx, \{dx\,dy\})$.

Notation: Now I extend this to the denotation of twisted forms using the conventions:

(1) all missing factors are to be written in appropriately hatted, before the actually present factors;

(2) the order of all factors, missing or not, matches the orientation;

(3) now that the orientation is redundant information, it is left out;

(4) the hat notation is only used for twisted forms. If it is hatted, you know that it is twisted.

(5) since the order of the present factors does not matter, leave them out as well.

Example: So, finally, I write the above form simply as $-\widehat{dy}$.

This works for everything except twisted n-forms. For these nothing is missing. You can either write these out explicitly, put in a factor of $\widehat{1}$, or use the star operator to be introduced in the next section to write it $*1$.

To connect this representation to our pictures we need to know whether, for example, $\widehat{dy}$ points up or down. In *Applied Differential Geometry* (ADG) the correspondence is given by

$$\{(\alpha, \Omega)\} \wedge \alpha = \Omega,$$

where $\{(\alpha, \Omega)\}$ stands for the orientation of the twisted form (α, Ω). This does not depend upon Ω since there is one Ω on each side of the equation. The curly brackets describe the symmetry properties of the object. In hat notation

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

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 1.

Ascii Shorthand

To denote these basis vectors and forms in plain text, I use the shorthands

$$px \equiv \frac{\partial}{\partial x}$$
$$tx \equiv \widehat{dx}$$

Pullback

To pull back a twisted form onto a subspace, the subspace must be given a transverse orientation. This will come either from the induced orientation of a boundary (see the section on integration to come), or from the location of the form on one side or the other of a surface, as when writing junction conditions. To pullback a twisted form onto a subspace lined up with the coordinate axes, let the subspace be defined by an equation

$$\eta = 0,$$

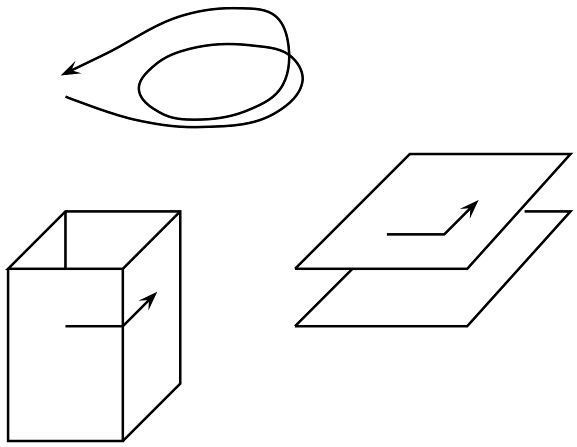

Figure 4. The sources that generate a magnetic field in the z direction, and a 2-form and a twisted 1-form with the same symmetry.

or possibly several such equations. The basis forms $d\eta$ all pull back to zero on the subspace. In our Manhattan context, these are all supposed to be coordinates. The transverse orientation of the subspace will be given by either $\{d\eta\}$ or $-\{d\eta\}$.

The pullback rule is

$$(\omega, \{d\eta\, d\alpha\}) \mapsto (\omega, \{d\alpha\}).$$

To get a twisted differential form on the subspace we need to select an orientation for the subspace. This requires us to discard part of the orientation chosen for the full space. The above choice, in which the discard is from the "tail of the snake" is necessary so that it will commute with d, which acts on the "head of the snake." In terms of the "hat" notation

$$\widehat{d\eta}\,\widehat{d\alpha} \mapsto \widehat{d\alpha}.$$

Example: In three–space, let us pull back the twisted 1–form $\widehat{dz}\,\widehat{dx}$ onto the $z = 0$ surface, taking the natural definition of the outside to be $z > 0$. The pullback is $\widehat{dx}$, sketched in Figure 5.

Integration

The rules are especially simple if the integration surfaces line up with the coordinates. One can integrate an r–form only over an oriented, r–dimensional space. A twisted form,

Figure 5. Pull back of $\widehat{dz}\,\widehat{dx}$.

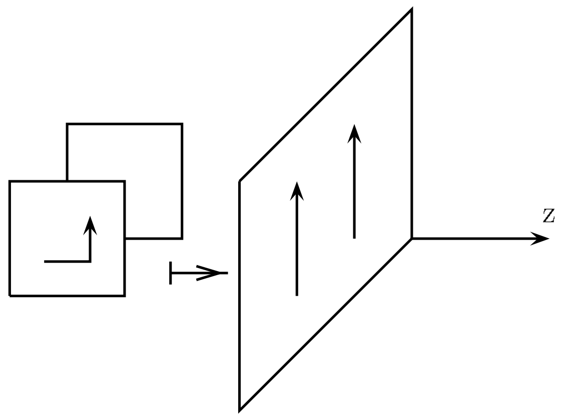

on the other hand, can only be integrated over an unoriented r–dimensional space. If α is an untwisted form, then

$$\int_{(\Gamma,\Omega)} \alpha$$

is evaluated by writing the basis forms in α in the order specified by the orientation; then converting any holonomic basis to a coordinate basis; then doing it as a conventional Riemann multiple integral.

For a twisted n-form α the integral that makes sense is

$$\int_{\Gamma} (\alpha, \Omega).$$

For a twisted p-form (α, Ω) we need a transverse orientation of Γ in order to be able to pullback (α, Ω). This can come from the equations of an implicit representation, as in the discussion of pullback.

Example: The electric potential difference is the integral $\int_\Gamma E$, and requires an oriented line integral.

The electric flux is trickier. We can integrate the twisted 2–form $*E$ over a 2–surface. But we need to pull the twisted 2–form back from our 3–space, and this pullback will require us to pick a transverse orientation. Here this transverse orientation points to the outside of the surface.

Boundaries of spaces acquire orientations naturally. The inner orientation of a boundary satisfies

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

where the coordinate q is increasing to the outside. The outer (transverse) orientation of a boundary must be consisted with Stokes' Theorem, and this requires

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

Examples

First we work out a Stokes' Theorem example, with untwisted forms.

Example: To evaluate a situation with a curl, and relate it to a line integral, go to three dimensions and take

$$E = x\, dy, \qquad dE = dx\, dy.$$

Integrate over the unit square in the positive quadrant. Orient the square using $\{dx\, dy\}$. Remember that we have to orient the surface in order to have the integral well defined. The only side of the square which contributes is the $x = 1$ side. On that side we have x increasing to the outside. Thus

$$dx \wedge \{\partial\Gamma\} = \{dx\, dy\}$$

and the orientation of that piece is upward,

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

Thus

$$\int_{\partial\Gamma} E = \int_{\Gamma} dE = 1.$$

Remember that the curl is not twisted.

Now we work out a Divergence Theorem example, Stokes' Theorem for twisted forms.

Example: We do an integral over a spherical surface. Let us work in spherical polar coordinates. To integrate the radial flux $D = a^2 \sin\theta\, \widehat{dr}$ over a sphere of radius a, defined by

$$r - a = 0,$$

the transverse orientation of the surface $r = a$ is given by $\{dr\}$. The rule for pullback gives us a pullback of $\widehat{1}$.

$$\int_{\Sigma} D = \int\int a^2 \sin\theta\, d\theta\, d\phi = 4\pi a^2.$$

The pullback rule is needed to describe the conditions associated with surface currents, surface charges, and the like. The easiest cases to consider are those in which a form may possibly be discontinuous across a surface.

For untwisted forms, the junction conditions are immediately expressed using pullback. The pullbacks of E or B from the two sides of the surface of discontinuity must be the same. You derive this by integrating over the classic "pillbox", or nowadays "tuna-fish can".

With twisted forms the situation is more complicated because there can be a surface charge and current, and so the conditions on D and H involve these surface distributions. There is a natural notion of transverse orientation coming from the gaussian pillbox, and it will be opposite on the two sides of the surface. We lack a good name for this differenced pullback.

Example: Find the junction conditions on H, coming from the equation

$$dH = 4\pi J.$$

Align coordinates so that the surface of discontinuity is $z = 0$. Integrate around a little loop in the $x = 0$ plane, and let its height in the z–direction going to zero. The loop is the boundary of a long thin rectangle Γ, which we give a transverse orientation in the plus x direction, $\{dx\}$. The boundary $\partial\Gamma$ is our loop, and it therefore has a transverse orientation defined by the above expression. On the upper line segment (where z increases to the outside of Γ) this orientation will be

$$\{dx\} \wedge dz = \{dx\, dz\}.$$

It is opposite on the lower segment.

The H field can be written in terms of components

$$H = H_x \widehat{dy}\,\widehat{dz} + H_y \widehat{dz}\,\widehat{dx} + H_z \widehat{dx}\,\widehat{dy},$$

and only the H_y term will contribute here. This pulls back to $-H_y$. The integration on the lower loop gives us a contribution that is the negative of this.

Now we want to describe this in terms of a twisted 1-form on the surface $z = 0$. That twisted 1-form will have a $\widehat{dx}$–component which will be the H_y field for $z < 0$ minus the H_y field for $z > 0$. In both cases, these contributions can be computed by computing the pullback of H to the surface with an orientation pointing in toward the surface.

Using Stokes' Theorem we can relate this to the surface current.

Figure 6. Geometry for computing the junction condition on H.

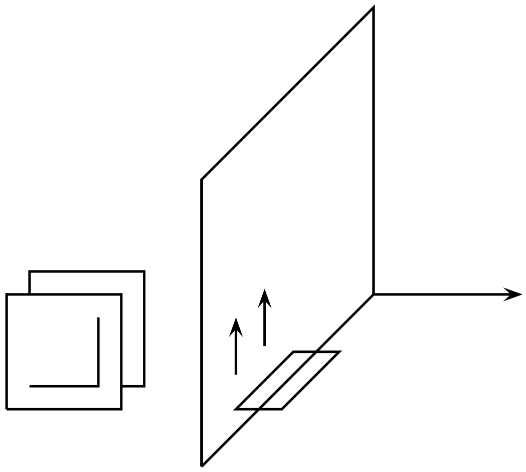

Unfortunately one cannot define pullback so that it both commutes with d and always uses the same orientation in the junction condition integrals. For the D field, one must use the orientation pointing away from the surface.

Star Operator

When we have a metric it is also useful to use a basis of 1–forms that has been normalized. This does not replace the usual coordinate basis. Operations using d are simpler with a coordinate basis. Operations involving $*$ are simpler in an orthonormal basis. An orthonormal basis need not be holonomic. A coordinate basis always is.

Example: The orthonormal basis for spherical polar coordinates is

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

We will write the generic orthonormal basis 1–form as ω^i, and an indeterminant exterior product of basis 1–forms as ω^A.

From ADG the basic rule for the star operator is given by

$$*(\beta\wedge\omega^i) = \sharp\omega^i \,\lrcorner\, *\beta,$$

where the $\sharp$ operator uses the metric to convert the 1–form ω^i to a vector. To start this recursive definition, we need to know $*1$. We define it as the unit twisted n–form,

$$*1 = (\omega^i \wedge \omega^A, \{\omega^i \wedge \omega^A\}).$$

As with a coordinate basis, I am going to leave out the wedge between these orthonormal basis forms, writing:

$$*1 = (\omega^i \, \omega^A, \{\omega^i \, \omega^A\}).$$

The notation $*1$ provides us with another notation for twisted n–forms.

For $\beta = 1$ we have

$$\begin{aligned} *(1 \wedge \omega^i) &= \frac{\partial}{\partial \omega^i} \lrcorner * 1 \\ &= \frac{\partial}{\partial \omega^i} \lrcorner (\omega^i \, \omega^A, \{\omega^i \, \omega^A\}) \\ &= (\frac{\partial}{\partial \omega^i} \lrcorner \omega^i) \widehat{\omega^i}. \end{aligned}$$

That is,

$$*\omega^i = \widehat{\omega^i} \, \omega^A.$$

and repeating that we can see that for a positive definite metric

$$*\omega^A = \widehat{\omega^A}.$$

If there are p timelike basis 1–forms in ω^A then we have an additional factor of $(-1)^p$.

The star operator provides a "metric" on each space of p-forms,

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

Definitions of $*$ which don't use twisted forms need to select an orientation to convert $*1$ to an untwisted form. They do this by "defining" $*1$ to be one of the two unit n–forms. Here we need no such choice.

There are two ways that come to mind for extending the $*$ operator to twisted forms. The most natural rule would be to use the inverse of the star already defined. The other rule would be to make it act similarly on twisted and untwisted forms, so that one can read and work without bothering with whether the forms are twisted or not. That rule would be

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

whereas the inverse rule would give

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

These differ by a sign $(-)^{AB}$. In three dimensions, this sign is always plus, and so the issue there is moot. Either convention works if you use it consistently. I recommend using the inverse, thus avoiding the ugly rule

$$**\,(\omega^A) = *(\widehat{\omega^A\,\omega^B}) = (-)^{AB}\omega^A, \qquad \text{NOT OUR STAR!}$$

although the ugly rule is used by default by all those who don't know about twisted forms.

The star operator follows the above rules in all dimensions, and with care you can write results in a form independent of dimension.

Example: The equations of electrostatics in differential forms are

$$\begin{aligned} dE &= 0, \\ d*E &= 4\pi\rho, \\ E &= -dV. \end{aligned}$$

The field energy is

$$W = \frac{1}{8\pi}\,E \wedge *E.$$

Example:

The wedge operator acts on twisted forms as you would expect:

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

but note that

$$dq \wedge (\widehat{dq\,d\alpha}) = *1 \neq dq\,d\alpha.$$

Twisted Units

The two unit pseudoscalars, twisted zero forms, in three dimensions

$$\widehat{dx}\,\widehat{dy}\,\widehat{dz} \quad \text{and} \quad \widehat{dz}\,\widehat{dy}\,\widehat{dx}$$

are used to represent the choice of a particular orientation, and the use of that choice to convert twisted forms to untwisted, and vice versa. Either one defines a parity on the space. Since these relations are not parity invariant, the generic notation for these, $\widehat{\omega^A}$, must be used carefully.

I will refer to either of the two twisted units as $\widehat{\Omega}$. The usual star operator used by those who do not use twisted forms is $\widehat{\Omega}*$. Multiplication by $\widehat{\Omega}$ converts a twisted $(n-r)$–form to an untwisted $(n-r)$–form. This notation makes explicit the need for them to choose an orientation, and lets you see which results are even in $\widehat{\Omega}$ and so independent of that choice.

We have

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

and this gives us an explicit formula for $*^{-1}$

$$*^{-1}\alpha = (-)^{\alpha(n-\alpha)}\, \widehat{\Omega} * \widehat{\Omega}\, \alpha.$$

Here α is twisted. Since there are two factors of $\widehat{\Omega}$ used, it doesn't matter whether you use the right–handed or the left–handed twisted unit.

> **Example:** Look at the star operator in two dimensional Euclidean space, $*_2$. If we take
>
> $$\widehat{\Omega} = \widehat{dx}\,\widehat{dy},$$
>
> then
>
> $$\widehat{\Omega} *_2\ dx = \widehat{\Omega}\,\widehat{dx}\,dy = dy,$$
>
> and
>
> $$\widehat{\Omega} *_2\ dy = \widehat{\Omega}\,\widehat{dy}\,dx = -dx,$$
>
> The culprit responsible for this ugly sign is clearly $\widehat{\Omega}$.

Codifferential

There is an operator somewhat like an adjoint to d for expressions of the form

$$\int_\Gamma \alpha \wedge *\beta,$$

which we will write δ. It is to satisfy

$$\int_\Gamma d\alpha \wedge *\beta = \int_\Gamma \alpha \wedge *\delta\beta \quad + \quad \text{boundary terms.}$$

We have

$$\int_\Gamma d\alpha \wedge *\beta = (-)^{\alpha-1} \int_\Gamma \alpha \wedge *d\beta \quad + \quad \text{boundary terms.}$$

and so

$$\delta\,\beta = (-)^{\beta}\ *^{-1}\, d * \beta.$$

We can write this in terms of our twisted unit:

$$\delta\,\beta = (-)^{\beta}\,(-)^{(\beta+1)(n-\beta-1)}\gamma * \gamma d * \beta.$$

If we pick γ such that $d\gamma = 0$, then it falls into the form

$$\delta\,\beta = (-)^{n(\beta+1)-1}\gamma * d\,\gamma * \,\beta.$$

Example: You can write Gauss' Law using the codifferential:

$$\begin{aligned} d * E &= 4\pi\rho \\ *^{-1} d * E &= 4\pi *^{-1}\rho, \\ \delta E &= -4\pi *^{-1}\rho. \end{aligned}$$

Here $*^{-1}\rho$ is the usual scalar charge density.

Since δ maps r–forms to $(r-1)$–forms, it is sometimes called the divergence operator, but that is misleading. On 1–forms it acts like the divergence, but on 2–forms it acts like the curl.

Laplacian

The generalization of the Laplacian to differential forms is

$$\Box = d\delta + \delta d.$$

On functions this gives us

$$\Box\phi = \delta d\phi = - *^{-1} d * d\phi.$$

This agrees with ∇^2 except for a sign. All mathematicians use this sign for $\Box$. Clearly we have

$$\begin{aligned} \Box d &= d\Box \\ \Box\delta &= \delta\Box. \end{aligned}$$

An exercise with nearly unlimited potential for sign errors will show that

$$\Box\,\gamma * = \gamma * \Box.$$

A hard result of elliptic operator theory shows that any form ω can be written

$$\omega = d\alpha + \delta\beta + \phi,$$

with

$$\Box\,\phi = 0.$$

The forms α, β, and ϕ are all orthogonal in the sense

$$\int_\Gamma \alpha \wedge *\beta = 0.$$

Summary

The pictorial representation of the twisted form $\widehat{d\alpha}$ has an orientation

$$\{\,\widehat{d\alpha}\,d\beta\} = \{d\alpha\}.$$

The star operator acts on an orthonormal basis according to

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

where there are p timelike basis forms in ω^A. Any factors ω^B can be used to complete the basis, in any order.

Pullback with orientation $\{d\eta\}$ follows

$$\widehat{d\eta}\,\widehat{d\alpha} \mapsto \widehat{d\alpha}.$$

The boundary orientation is given by

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

with q increasing to the outside.