 29 January 1999

THE HODGE DUAL OPERATOR

Tevian Dray
Department of Mathematics, Oregon State University, Corvallis, OR 97331, USA
tevian@math.oregonstate.edu

The Hodge dual operator $*$ is one of the 3 basic operations on differential forms. (The other 2 are wedge product $\wedge$ and exterior differentiation d.) However most treatments consider only positive-definite inner products, and there are at least 2 standard ways of generalizing this to inner products of arbitrary signature. We outline here a construction of the Hodge dual operator which works for any signature, resulting in a particular choice of signs. [1]

1. NOTATION

We assume throughout that V is an n-dimensional vector space with a non-degenerate inner product g, also called the (inverse) metric. We further assume that an ordered orthonormal basis $\{\sigma^i \in V, i = 1, ..., n\}$ is given, so that

$$g^{ij} := g(\sigma^i, \sigma^j) = \pm\delta^{ij} \tag{1}$$

where δ^{ij} is the Kronecker delta symbol, which is 1 if $i = j$ and 0 otherwise. We define the *signature* s of g to be the number of $-$ signs occurring on the right-hand-side of this equation. [2]

The spaces $\bigwedge^p V$ of p-vectors are constructed as usual using the wedge product $\wedge$. In particular, we have

$$\bigwedge^0 V = \mathbb{R} \qquad \bigwedge^1 V = V \qquad \bigwedge^p V = \{0\} \quad (p > n) \tag{2}$$

and

$$\dim\left(\bigwedge^p V\right) = \binom{n}{p} = \frac{n!}{p!(n-p)!} \tag{3}$$

The ordered basis determines an *orientation* by selecting a preferred n-vector, namely

$$\omega = \sigma^1 \wedge \sigma^2 \wedge ... \wedge \sigma^n \in \bigwedge^n V \tag{4}$$

[1] Our treatment parallels that in Flanders, which however uses the opposite convention. Bishop & Goldberg only considers the positive-definite case, for which all 3 treatments agree.

[2] This is nonstandard, as it is customary to define the signature as the *difference* between the number of + and $-$ signs.

A basis of $\bigwedge^p V$ is given by $\{\sigma^I\}$, where

$$\sigma^I = \sigma^{i_1} \wedge ... \wedge \sigma^{i_p} \tag{5}$$

where the indices are assumed to satisfy $1 \le i_1 < ... < i_p \le n$. The label $I = \{i_k, k = 1, ..., p\}$ is called the index set; the basis of $\bigwedge^p V$ consists of all σ^I as I ranges over the allowed index sets (of length p).

A p-vector $\alpha \in \bigwedge^p V$ is called *decomposable* if and only if there exist vectors $\alpha^i \in V$ with

$$\alpha = \alpha^1 \wedge ... \wedge \alpha^p \tag{6}$$

The inner product can be extended to $\bigwedge^0 V$ via $g(1,1) = 1$, and to $\bigwedge^p V$ by requiring that

$$g(\alpha, \beta) = \det\Big(g(\alpha^i, \beta^j)\Big) \tag{7}$$

for any decomposable p-vectors α and β. Equivalently,

$$g(\sigma^I, \sigma^J) = \left(\prod_{k=1}^{p} g(\sigma^{i_k}, \sigma^{i_k})\right) \delta^{IJ} = \pm\delta^{IJ} \tag{8}$$

so that the σ^I are in fact an orthonormal basis for $\bigwedge^p V$. (Note that the same symbol g is used for the different inner products in the different vector spaces $\bigwedge^p V$). In particular, we have

$$g(\omega, \omega) = \prod_{k=1}^{n} g(\sigma^i, \sigma^i) = (-1)^s \tag{9}$$

so that the norm of the preferred n-vector is determined by the signature.

2. ABSTRACT DEFINITION

We start with an important lemma.

Lemma: *Given any real-valued linear function f on a vector space W with (nondegenerate) inner product g, there exists a unique element $\beta \in W$ such that*

$$f(\alpha) = g(\alpha, \beta) \quad \forall \alpha \in W$$

Proof: To prove this, simply expand α and β in terms of the basis. From linearity, we must have $f(\sigma^i) = g(\sigma^i, \beta)$, and since the basis is orthonormal we obtain

$$\beta = \sum_i g(\sigma^i, \sigma^i)\, f(\sigma^i)\, \sigma^i$$

It is important to realize that this lemma applies to *any* vector space; we will now apply it to $\bigwedge^{n-p} V$.

Fix a p-vector $\lambda \in \bigwedge^p V$. For *any* $\theta \in \bigwedge^{n-p} V$, $\lambda \wedge \theta$ is an n-vector. But all n-vectors are multiples of ω, and we can define a function f_λ via

$$\lambda \wedge \theta = f_\lambda(\theta)\, \omega \tag{10}$$

With this definition, f_λ is clearly a linear function on $\bigwedge^{n-p}V$. Therefore, the lemma tells us that there is a unique element $\phi \in \bigwedge^{n-p}V$ such that

$$f_\lambda(\theta) = g(\theta, \phi) \quad \forall \theta \in \bigwedge^{n-p} V \tag{11}$$

We finally define the *Hodge dual* $*\lambda$ of λ to be [3]

$$*\lambda = (-1)^s \phi \in \bigwedge^{n-p} V \tag{12}$$

Equivalently, the Hodge dual of a p-vector λ is the $(n{-}p)$-vector defined by

$$\boxed{\lambda \wedge \theta = (-1)^s g(\theta, *\lambda)\, \omega} \qquad \forall \theta \in \bigwedge^{n-p} V \tag{13}$$

(This definition makes sense because of the lemma.)

3. PRACTICAL DEFINITION

Since the $*$ operator is linear, it is enough to compute it in a basis. We will use our orthonormal bases $\{\sigma^I\}$ for each of the vector spaces $\bigwedge^p V$. Furthermore, we can always permute the basis without changing the orientation ω, so long as we restrict ourselves to *even* permutations. But any $\{\sigma^I\}$ can be brought to the form $\sigma^1 \wedge ... \wedge \sigma^p$ by an even permutation. [4] It is therefore sufficient to determine $*\lambda$ for the basis p-vector

$$\lambda = \sigma^1 \wedge ... \wedge \sigma^p \tag{14}$$

Using the boxed definition above, we obtain for any basis $(n{-}p)$-vector $\sigma^J \in \bigwedge^{n-p}V$

$$\lambda \wedge \sigma^J = (-1)^s g(\sigma^J, *\lambda)\, \omega \tag{15}$$

But the LHS of this equation is clearly 0 unless $J = \{p+1, ..., n\}$. Since the σ^J's are themselves orthonormal, the RHS then tells us that $*\lambda$ only has a σ^J component for J as given above, or in other words that

$$*\lambda = c\, \sigma^{p+1} \wedge ... \wedge \sigma^n \tag{16}$$

for some real number c. But for this value of J

$$\lambda \wedge \sigma^J = \omega \tag{17}$$

which leads to

$$1 = (-1)^s g(\sigma^J, c\, \sigma^J) \tag{18}$$

or equivalently

$$c = \frac{(-1)^s}{g(\sigma^J, \sigma^J)} = \frac{g(\omega, \omega)}{g(\sigma^J,\, \sigma^J)} = g(\lambda, \lambda) \tag{19}$$

Putting this all together, we obtain

$$\boxed{*(\sigma^1 \wedge ... \wedge \sigma^p) = g(\sigma^1, \sigma^1) ... g(\sigma^p, \sigma^p)\, \sigma^{p+1} \wedge ... \wedge \sigma^n} \tag{20}$$

---

[3] The definition in Flanders is equivalent to omitting the factor of $(-1)^s$. One suspects that Bishop & Goldberg would have omitted this factor as well, but they only considers positive-definite inner products, for which $s = 0$.

[4] If $p = n - 1$, only one of $\pm\sigma^I$ can be brought to this form by an even permutation. But $*\sigma^I$ can trivially be computed from $*(-\sigma^I)$ if necessary.

4. EXAMPLES

To see how this works in practice, consider first the case of ordinary Euclidean 2-space, that is $V = \mathbb{R}^2$ with a positive-definite inner product. For compatibility with later work, we will write the basis as $\{dx, dy\}$. The inner product can be conveniently expressed in terms of the *metric* or line element

$$ds^2 = dx^2 + dy^2 \tag{21}$$

which also encodes the fact that our basis is orthonormal. Since we are in the case of Euclidean (also called Riemannian) signature, namely $s = 0$, all basis vectors have norm $+1$, so that the g's can be omitted from the formula at the end of the last section. Choosing our preferred 2-vector to be [5]

$$\omega = dx \wedge dy = dy \wedge (-dx) \tag{22}$$

we obtain

$$\begin{aligned} *1 &= dx \wedge dy \\ *dx &= dy \\ *dy &= -dx \\ *(dx \wedge dy) &= 1 \end{aligned} \tag{23}$$

Another example is the case $V = \mathbb{R}^3$ with Lorentzian signature, namely $s = 1$, which is called Minkowski 3-space. The metric is

$$ds^2 = dx^2 + dy^2 - dt^2 \tag{24}$$

corresponding to the orthonormal basis $\{dx, dy, dt\}$. Note that

$$g(dt, dt) = -1 \tag{25}$$

We take the preferred 3-form, or volume element, to be

$$\omega = dx \wedge dy \wedge dt = dy \wedge dt \wedge dx = dt \wedge dx \wedge dy \tag{26}$$

which leads to

$$\begin{aligned} *1 &= dx \wedge dy \wedge dt \\ *dx &= dy \wedge dt \\ *dy &= dt \wedge dx \\ *dt &= -dx \wedge dy \\ *(dx \wedge dy) &= dt \\ *(dy \wedge dt) &= -dx \\ *(dt \wedge dx) &= -dy \\ *(dx \wedge dy \wedge dt) &= -1 \end{aligned} \tag{27}$$

Unlike the previous example, the minus signs here arise exclusively due to the presence of dt, whose "squared norm" is negative. (This would not be true had we failed to use a "cyclic" basis for 2-vectors.)

---

[5] It is only for 1-vectors in 2 dimensions that the minus sign in ω can not be eliminated by an appropriate ordering of the basis. An alternative in this case is to write $\omega = (-dy) \wedge dx$ and $*(-dy) = dx$.

5. PROPERTIES

We now derive some useful properties of the $*$ operator.

First of all, with $I = \{1, ..., p\}$ and $J = \{p{+}1...n\}$ as above, what is $*\sigma^J$? We have

$$\begin{aligned}\sigma^I &= \sigma^1 \wedge ... \wedge \sigma^p \\ \sigma^J &= \sigma^{p+1} \wedge ... \wedge \sigma^n\end{aligned} \tag{28}$$

The practical definition (20) of the Hodge dual is just

$$*\sigma^I = g(\sigma^I, \sigma^I)\, \sigma^J \tag{29}$$

which follows from the abstract definition (13) since $\sigma^I \wedge \sigma^J = \omega$. Two special cases deserve mention, namely 0-vectors and n-vectors, for which we have

$$\boxed{*1 = \omega} \qquad \boxed{*\omega = (-1)^s} \tag{30}$$

in all cases.

In words, the dual of an (orthonormal) basis p-vector is the $n{-}p$-vector obtained by "wedging" together (in a particular order) all the basis 1-vectors *not* appearing in the given p-vector, then multiplying by the norm of that p-vector. The "particular order" is such that the product of the p-vector with the remaining 1-vectors is just the preferred n-vector ω; if the order is wrong, it can be corrected either by interchanging any two 1-vectors, or by multiplying by -1.

We use this description in order to determine $*\sigma^J$. This will clearly be some multiple of σ^I, but what multiple? But

$$\sigma^J \wedge \sigma^I = (-1)^{p(n-p)} \sigma^I \wedge \sigma^J = (-1)^{p(n-p)} \omega \tag{31}$$

or in other words

$$\sigma^J \wedge \left((-1)^{p(n-p)} \sigma^I\right) = \omega \tag{32}$$

from which it finally follows that

$$*\sigma^J = g(\sigma^J, \sigma^J) \left((-1)^{p(n-p)} \sigma^I\right) \tag{33}$$

This in turn allows us to work out $**\sigma^I$:

$$\begin{aligned}**\sigma^I &= * \left(g(\sigma^I, \sigma^I)\, \sigma^J\right) = g(\sigma^I, \sigma^I)\; *\sigma^J \\ &= g(\sigma^I, \sigma^I) g(\sigma^J, \sigma^J) (-1)^{p(n-p)} \sigma^I = (-1)^{p(n-p)+s} \sigma^I\end{aligned} \tag{34}$$

since

$$g(\sigma^I, \sigma^I)\, g(\sigma^J, \sigma^J) = g(\omega, \omega) = (-1)^s \tag{35}$$

The formula for the double dual is most easily remembered as

$$\boxed{** = (-1)^{p(n-p)+s}} \tag{36}$$

The double dual can in turn be used to compute one of the most important properties of the $*$ operator. Applying the abstract definition (13) to the $n-p$-vector $*\beta$ leads to

$$*\beta \wedge \alpha = (-1)^s g(\alpha, **\beta)\,\omega \tag{37}$$

for any p-vector α. Evaluating the double dual on the right, and commuting the vectors on the left causes all the minus signs to go away, and we are left with

$$\boxed{\alpha \wedge *\beta = g(\alpha, \beta)\,\omega} \tag{38}$$

which holds for any two p-vectors α and β. It was to ensure the validity of this identity that we inserted the factor of $(-1)^s$ in the original definition. In fact, many authors take (38) as the *definition* of the $*$ operator, although this requires some checking to make sure it is well-defined.

Finally, we can turn (38) around and express the inner product g entirely in terms of $*$, thus showing that these two concepts are equivalent. Explicitly, we have

$$\boxed{g(\alpha, \beta) = (-1)^s *(\alpha \wedge *\beta)} \tag{39}$$

where we have used (30). In Euclidean 3-space, this yields a formula for the ordinary "dot" product. (Try it!)