

Chern–Weil Forms and Abstract Homotopy Theory

Dexter Chua

1	The statement	1
2	The proof	7

1 THE STATEMENT

The main theorem of the Freed–Hopkins paper *Chern–Weil forms and abstract homotopy theory* is that Chern–Weil forms are the only natural way to get a differential form from a principal G -bundle.

Theorems along these lines are of interest historically. It is an important ingredient in the heat kernel proof of the Atiyah–Singer index theorem. Essentially, the idea of the proof is to use the heat equation to show that there is *some* formula for the index of a vector bundle in terms of the derivatives of the metric, and then by invariant theory, this must be given by the Chern–Weil forms we know and love. One then computes this for sufficiently many examples to figure out exactly which characteristic class it is, as Hirzebruch originally did for his signature formula.

To state the theorem, we work in the category $\mathrm{Shv}(\mathrm{Man}, \mathcal{S})$. For the purposes of this theorem, it actually suffices to work with sheaves of groupoids, i.e. $\mathrm{Shv}(\mathrm{Man}, \tau_{\leq 1}\mathcal{S})$. This only requires 2-category theory instead of ∞ -category theory. However, working with ∞ -categories presents no additional difficulty, and is what we shall do.

We now introduce the main characters of the story.

Example 1. Any $M \in \mathrm{Man}$ defines a representable (discrete) sheaf, which we denote by M again.

Example 2. Any sheaf of sets on Man is in particular sheaf of (discrete) spaces. Thus, for $p \geq 0$, we have a discrete sheaf

$$\Omega^p \in \mathrm{Shv}(\mathrm{Man}, \mathcal{S}).$$

This is in fact a sheaf of vector spaces, and moreover, there are linear natural transformations $d : \Omega^p \rightarrow \Omega^{p+1}$. Thus, we get a sheaf of chain complexes Ω^\bullet , and

$$[M, \Omega^\bullet] = \Omega^\bullet(M).$$

In general, for any sheaf \mathcal{F} , we can think of $\Omega^\bullet(\mathcal{F}) \equiv [\mathcal{F}, \Omega^\bullet]$ as the de Rham complex of \mathcal{F} .

From now on, fix G a Lie group.

Example 3. For $M \in \mathbf{Man}$, define $B_{\nabla}G(M)$ to be the groupoid of principal G -bundles on M with connection and isomorphisms, which we think of as a 1-truncated space. This defines $B_{\nabla}G \in \mathbf{Shv}(\mathbf{Man}, \mathcal{S})$.

The main theorem is

Theorem 4. *The Chern–Weil homomorphism induces an isomorphism*

$$(\mathrm{Sym}^\bullet \mathfrak{g}^*)^G \xrightarrow{\sim} \Omega^\bullet(B_{\nabla}G).$$

To prove the theorem, we consider the universal principal G -bundle $E_{\nabla}G \rightarrow B_{\nabla}G$. The point is that $E_{\nabla}G$ admits a much more explicit description, and then we use $B_{\nabla}G = E_{\nabla}G//G$ to understand $B_{\nabla}G$ itself.

$E_{\nabla}G$ can be described explicitly as follows:

Example 5. Define $E_{\nabla}G(M)$ to be the groupoid of trivialized G -bundles on M with connection. Equivalently, this is the groupoid of connections on the trivial G -bundle $M \times G \rightarrow G$. So $E_{\nabla}G \cong \Omega^1 \otimes \mathfrak{g}$.

There is then a natural map $E_{\nabla}G(M) \rightarrow B_{\nabla}G(M)$, which one can easily check is the universal principal G -bundle. Our next claim is that $B_{\nabla}G(M) = E_{\nabla}G(M)//G$, which is clear once we know what the latter is.

Definition 6. Let $\mathcal{F} \in \mathbf{Shv}(\mathbf{Man}, \mathcal{S})$, and let $\alpha : G \times \mathcal{F} \rightarrow \mathcal{F}$ be an action by G . Explicitly, for each $M \in \mathbf{Man}$, there is a group action

$$\mathrm{Hom}_{\mathbf{Man}}(M, G) \times \mathcal{F}(M) \rightarrow \mathcal{F}(M)$$

where $\mathrm{Hom}_{\mathbf{Man}}(M, G)$ is given the pointwise group structure. We can then define the *action groupoid*

$$(\mathcal{F}//G)_\bullet = G^{\times \bullet} \times \mathcal{F} \in \mathbf{Shv}(\mathbf{Man}, \mathcal{S})^{\Delta^{\mathrm{op}}}.$$

The homotopy quotient of \mathcal{F} by G is then

$$\mathcal{F}//G = |(\mathcal{F}//G)_\bullet|.$$

Note that this geometric realization is taken in the category $\mathbf{Shv}(\mathbf{Man}, \mathcal{S})$. To compute this, one takes the geometric realization in the category of presheaves, then sheafify.

We then see that $B_{\nabla}G = E_{\nabla}G//G$. Explicitly, the action of the gauge group can be described as follows — given $g : M \rightarrow G$ and $\alpha \in E_{\nabla}G(M) = \Omega^1(M; \mathfrak{g})$, we have

$$g \cdot \alpha = g^* \theta + \mathrm{Ad}_{g^{-1}} \alpha.$$

Remark 7. Formally, to prove that $B_{\nabla}G = E_{\nabla}G//G$, we first form the quotient of $E_{\nabla}G$ by G in the category of presheaves. Since $E_{\nabla}G$ is discrete, this is given by (the nerve of) the action groupoid of the G -action on $E_{\nabla}G$. This gives the presheaf of trivial principal G -bundles. To show that the sheafification is $B_{\nabla}G$, observe that there is a natural map from this presheaf to $B_{\nabla}G$, and it is an equivalence on stalks since all principal G -bundles on contractible spaces are trivial. So it induces an isomorphism after sheafification.

Our proof then naturally breaks into two steps. First, we compute $\Omega^{\bullet}(E_{\nabla}G)$, and then we need to know how to compute $\Omega^{\bullet}(\mathcal{F}//G)$ from $\Omega^{\bullet}(\mathcal{F})$ for any discrete sheaf \mathcal{F} .

We first do the second part.

Lemma 8. *Let $\mathcal{F} \in \text{Shv}(\text{Man}, \mathcal{S})$ be a discrete sheaf with a G -action $\alpha : G \times \mathcal{F} \rightarrow \mathcal{F}$. Then $\Omega^{\bullet}(\mathcal{F}//G)$ is the subcomplex of $\Omega^{\bullet}(\mathcal{F})$ consisting of the ω such that*

- (1) $\alpha^*\omega|_{\{g\} \times \mathcal{F}} = \omega$ for all $g \in G$; and
- (2) $\iota_{\xi}\omega = 0$ for all $\xi \in \mathfrak{g}$.

The first condition says ω should be G -invariant, and the second condition says ω is suitably “horizontal”.

Remark 9. Let us explain what we mean by $\iota_{\xi}\omega$. In general, for M a manifold and X is a vector field on M , we can define $\iota_X : \Omega^p(M \times N) \rightarrow \Omega^{p-1}(M \times N)$ for all manifolds N . Then by left Kan extension, this induces a map $\iota_X : \Omega^p(M \times \mathcal{F}) \rightarrow \Omega^{p-1}(M \times \mathcal{F})$ for all $\mathcal{F} \in \text{Shv}(\text{Man}, \mathcal{S})$.

Now if \mathcal{F} has a G -action and $\xi \in \mathfrak{g}$, then ξ induces an invariant vector field on G , which we also call ξ . We then define $\iota_{\xi} : \Omega^p(\mathcal{F}) \rightarrow \Omega^{p-1}(\mathcal{F})$ by the following composition

$$\Omega^p(\mathcal{F}) \xrightarrow{\alpha^*} \Omega^p(G \times \mathcal{F}) \xrightarrow{\iota_{\xi}} \Omega^{p-1}(G \times \mathcal{F}) \longrightarrow \Omega^{p-1}(\{e\} \times \mathcal{F}) = \Omega^{p-1}(\mathcal{F}),$$

where the last map is induced by the inclusion.

This gives us a very explicit method to compute the natural transformation $\iota_{\xi}\omega$ for $\omega \in \Omega^p(\mathcal{F})$ and $\xi \in \mathfrak{g}$. Given a test manifold M and $\phi \in \mathcal{F}(M)$, which we think of as a natural transformation $\phi : M \rightarrow \mathcal{F}$, we form the composite

$$G \times M \xrightarrow{1 \times \phi} G \times \mathcal{F} \xrightarrow{\alpha} \mathcal{F} \xrightarrow{\omega} \Omega^p$$

This defines a differential form $\eta \in \Omega^p(G \times M)$. Then we have

$$(\iota_{\xi}\omega)_M(\phi) = \iota_{\xi}\eta|_{\{e\} \times M}.$$

Proof. We have

$$\Omega^p(\mathcal{F} // G) = \Omega^p(|(\mathcal{F} // G)_\bullet|) = \text{Tot}(\Omega^p((\mathcal{F} // G)_\bullet)).$$

Since $(\mathcal{F} // G)_\bullet$ is a simplicial discrete sheaf, its totalization can be computed by

$$\Omega^p(\mathcal{F} // G) = \ker \left(\Omega^p(\mathcal{F}) \xrightarrow{\text{pr}^* - \alpha^*} \Omega^p(G \times \mathcal{F}) \right),$$

where $\text{pr} : G \times \mathcal{F} \rightarrow \mathcal{F}$ is the projection.

To prove the lemma, we have to show that $\text{pr}^*\omega = \alpha^*\omega$ iff the conditions in the lemma are satisfied. This follows from the more general claim below with $\eta = \alpha^*\omega - \text{pr}^*\omega$.

Claim. *Let M be a manifold and \mathcal{F} a sheaf. Then $\eta \in \Omega^p(M \times \mathcal{F})$ is zero iff*

- (1) $\eta|_{\{x\} \times \mathcal{F}} = 0$ for all $x \in M$
- (2) $\iota_X \eta = 0$ for any vector field X on M .

The conditions (1) and (1') match up exactly. Unwrapping the definition of ι_ξ and noting that $\iota_X \text{pr}^*\omega = 0$ always, the only difference between (2) and (2') is that in (2), we only test on invariant vector fields on G , instead of all vector fields, and we only check the result is zero after restricting to a fiber $\{e\} \times \mathcal{F}$. The former is not an issue because the condition $C^\infty(G)$ -linear and the invariant vector fields span as a $C^\infty(G)$ -module. The latter also doesn't matter because we have assumed that $\alpha^*\omega$ is invariant.

To prove the claim, if \mathcal{F} were a manifold, this is automatic, since the first condition says η vanishes on vectors in the N direction while the second says it vanishes on vectors in the M direction.

If \mathcal{F} were an arbitrary sheaf, we know η is zero when pulled back along any map $(1 \times \phi) : M \times N \rightarrow M \times \mathcal{F}$ where N is a manifold, by naturality of the conditions. But since $M \times \mathcal{F}$ is a colimit of such maps, η must already be zero on $M \times \mathcal{F}$. \square

Now it remains to describe $\Omega^\bullet(E_{\nabla} G) = \Omega^\bullet(\Omega^1 \otimes \mathfrak{g})$. More generally, for any vector space V , we can calculate $\Omega^\bullet(\Omega^1 \otimes V)$. We first state the result in the special case where $V = \mathbb{R}$.

Theorem 10.

$$\Omega^p(\Omega^1) \cong \mathbb{R} \text{ for all } p \geq 0.$$

For $p = 2q$, it sends ω to $(d\omega)^q$. For $p = 2q + 1$, it sends ω to $\omega \wedge (d\omega)^q$.

The general case is no harder to prove, and the result is described in terms of the Koszul complex.

Definition 11. Let V be a vector space. The *Koszul complex* $\text{Kos}^\bullet V$ is a differential graded algebra whose underlying algebra is

$$\text{Kos}^\bullet V = \bigwedge^\bullet V \otimes \text{Sym}^\bullet V.$$

For $v \in V$, we write v for the corresponding element in $\bigwedge^1 V$, and \tilde{v} for the corresponding element in $\text{Sym}^1 V$. We set $|v| = 1$ and $|\tilde{v}| = 2$. The differential is then

$$d(v) = \tilde{v}, \quad d(\tilde{v}) = 0.$$

Theorem 12. *For any vector space V , we have an isomorphism of differential graded algebras*

$$\eta : \text{Kos}^\bullet V^* \xrightarrow{\sim} \Omega^\bullet(\Omega^1 \otimes V).$$

In particular,

$$\Omega^\bullet(E_{\nabla} G) = \text{Kos}^\bullet \mathfrak{g}^*.$$

Explicitly, for $\ell \in V^ = \bigwedge^1 V^*$, the element $\eta(\ell) \in \Omega^1(\Omega^1 \otimes V)$ is defined by*

$$\eta(\ell)(\alpha \otimes v) = \langle v, \ell \rangle \alpha$$

for $\alpha \in \Omega^1$ and $v \in V$. This is then extended to a map of differential graded algebras. In other words, the theorem says every natural transformation

$$\omega_M : \Omega^1(M; V) \rightarrow \Omega^p(M)$$

is (uniquely) a linear combination of transformations of the form

$$\sum \alpha_i \otimes v_i \mapsto \sum_{I,J} M_{I,J}(v_{i_1}, \dots, v_{i_k}, v_{j_1}, \dots, v_{j_\ell}) \alpha_{i_1} \wedge \dots \wedge \alpha_{i_k} \wedge d\alpha_{j_1} \wedge \dots \wedge d\alpha_{j_\ell}$$

where $M_{I,J}$ is anti-symmetric in the first k variables and symmetric in the last ℓ .

Using this, we conclude

Theorem 13. *The Chern–Weil homomorphism gives an isomorphism*

$$(\text{Sym}^\bullet \mathfrak{g}^*)^G \xrightarrow{\sim} \Omega^\bullet(B_{\nabla} G),$$

and the differential on $\Omega^\bullet(B_{\nabla} G)$ is zero.

Note that this $\text{Sym}^\bullet \mathfrak{g}^*$ is different from that appearing in the Koszul complex.

Proof. We apply the criteria in Lemma 8. The first condition is the G -invariance condition, and translates to the $(\dots)^G$ part of the statement. So we have to check that the forms satisfying the second condition are isomorphic to $\text{Sym}^\bullet \mathfrak{g}^*$.

To do so, we have to compute the action of ι_ξ on $E_\nabla G$ following the recipe in Remark 9. Fix $\omega \in \Omega^p(E_\nabla G)$ and $\xi \in \mathfrak{g}$.

Let $\phi : M \rightarrow E_\nabla G$ be a trivial principal G -bundle with connection $A \in \Omega^1(M; \mathfrak{g})$. The induced principal G -bundle on $G \times M$ under the action then has connection $\theta + \text{Ad}_{g^{-1}} A$. So by definition,

$$(\iota_\xi \omega)_M(A) = \iota_\xi (\omega(\theta + \text{Ad}_{g^{-1}} A)) \Big|_{\{e\} \times M}.$$

To compute the action on $\text{Kos}^\bullet \mathfrak{g}^*$, it suffices to compute it on $\bigwedge^1 \mathfrak{g}^*$ and $\text{Sym}^1 \mathfrak{g}^*$.

1. If $\lambda \in \mathfrak{g}^* = \bigwedge^1 \mathfrak{g}^*$, then $\lambda(A) = \langle A, \lambda \rangle$, and

$$\iota_\xi \langle \theta + \text{Ad}_{g^{-1}} A, \lambda \rangle = \langle \iota_\xi \theta + \iota_\xi \text{Ad}_{g^{-1}} A, \lambda \rangle.$$

We know $\iota_\xi \theta = \xi$, and $\iota_\xi \text{Ad}_{g^{-1}} A = 0$ since $\text{Ad}_{g^{-1}} A$ vanishes on all vectors in the G direction. So we know

$$\iota_\xi \lambda = \langle \xi, \lambda \rangle \in \bigwedge^0 \mathfrak{g}^*.$$

2. Next, $\tilde{\lambda}(A) = \langle dA, \lambda \rangle$. We compute

$$\begin{aligned} \iota_\xi \langle d(\theta + \text{Ad}_{g^{-1}} A), \lambda \rangle \Big|_{\{e\} \times M} \\ = \iota_\xi \left\langle -\frac{1}{2}[\theta, \theta] + \text{Ad}_{\text{d}g^{-1}} \wedge A + \text{Ad}_{g^{-1}} dA, \lambda \right\rangle \Big|_{\{e\} \times M} \\ = \langle -\text{Ad}_\xi A, \lambda \rangle = \langle A, -\text{Ad}_\xi^* \lambda \rangle. \end{aligned}$$

So

$$\iota_\xi \tilde{\lambda} = -\text{Ad}_\xi^* \lambda \in \bigwedge^1 \mathfrak{g}^*.$$

First observe that in $\bigwedge^\bullet \mathfrak{g}^*$, the only elements killed by ι_ξ are those in $\bigwedge^0 \mathfrak{g}^* \cong \mathbb{R}$. To take care of the Sym part, set

$$\Omega_\lambda = \tilde{\lambda} + \frac{1}{2}[\lambda, \lambda].$$

Since $\tilde{\lambda}(A) = \langle dA, \lambda \rangle$, we see that $\Omega_\lambda(A) = \langle \Omega_A, \lambda \rangle$, where Ω_A is the curvature, and one calculates $\iota_\xi \Omega_\lambda = 0$. By a change of basis, we can identify

$$\text{Kos}^\bullet \mathfrak{g}^* \cong \bigwedge^\bullet \mathfrak{g}^* \otimes \text{Sym}^\bullet \langle \Omega_\lambda : \lambda \in \mathfrak{g}^* \rangle,$$

and ι_ξ vanishes on the second factor entirely. So we are done. \square

More generally, the same proof shows that

Theorem 14. *If M is a smooth manifold, the de Rham complex of $M \times (\Omega^1 \otimes V)$ is $\Omega(M; \text{Kos } V^*)^\bullet$ (the total complex of $\Omega^\bullet(M; \text{Kos}^\bullet V^*)$).*

In particular, if M has a G -action, then $M \times E_{\nabla}G//G$ is exactly the Cartan model for equivariant de Rham cohomology.

This would follow immediately if we had a result that says $\Omega^\bullet(M \times \mathcal{F}) \cong \Omega^\bullet(M) \hat{\otimes} \Omega^\bullet(\mathcal{F})$, and since $\Omega^\bullet(E_{\nabla}G)$ is finite dimensional, the completed tensor product is the usual tensor product.

2 THE PROOF

We now prove of Theorem 12. The $p = 0$ case is trivial, so assume $p > 0$.

Recall that we have to show that any natural transformation

$$\omega_M : \Omega^1(M; V) \rightarrow \Omega^p(M)$$

is (uniquely) a linear combination of transformations of the form

$$\sum \alpha_i \otimes v_i \mapsto \sum_{I, J} M_{I, J}(v_{i_1}, \dots, v_{i_k}, v_{j_1}, \dots, v_{j_\ell}) \alpha_{i_1} \wedge \dots \wedge \alpha_{i_k} \wedge d\alpha_{j_1} \wedge \dots \wedge d\alpha_{j_\ell}.$$

The uniqueness part is easy to see since we can extract $M_{I, J}$ by evaluating $\omega_M(\alpha)$ for M of dimension large enough. So we have to show every ω_M is of this form.

The idea of the proof is to first use naturality to show that for $x \in M$, the form $\omega_M(\alpha)_x$ depends only on the N -jet of α at x for some large but finite number N (of course, *a posteriori*, $N = 1$ suffices). Once we know this, the problem is reduced to one of finite dimensional linear algebra and invariant theory.

Lemma 15. *For $\omega \in \Omega^p(\Omega^1 \otimes V)$ and $\alpha \in \Omega^1(M; V)$, the value of $\omega_M(\alpha)$ at $x \in M$ depends only on the N -jet of α at p for some N . In fact, $N = p$ suffices.*

We elect to introduce the constant N , despite it being equal to p , because the precise value does not matter.

Proof. Suppose α and α' have identical p -jets at x . Then there are functions f_0, f_1, \dots, f_p vanishing at p and $\beta \in \Omega^1(M; V)$ such that

$$\alpha' = \alpha + f_0 f_1 \cdots f_p \beta.$$

The first step is to replace the f_i with more easily understood coordinate functions. Consider the maps

$$M \xrightarrow{1_M \times (f_0, \dots, f_p)} M \times \mathbb{R}^{p+1} \xrightarrow{\text{pr}_1} M.$$

Let $\tilde{\alpha}, \tilde{\beta}$ be the pullbacks of the corresponding forms under pr_1 , and t_0, \dots, t_p the standard coordinates on \mathbb{R}^{p+1} . Then $\alpha, f_0 f_1 \cdots f_p \beta$ are the pullbacks of $\tilde{\alpha}, t_0 t_1 \cdots t_p \tilde{\beta}$ under the first map.

So it suffices to show that $\omega_{M \times \mathbb{R}^{p+1}}(\tilde{\alpha})$ and $\omega_{M \times \mathbb{R}^{p+1}}(\tilde{\alpha} + t_0 t_1 \cdots t_p \tilde{\beta})$ agree as p -forms at $(x, 0)$.

The point now is that by multilinearity of a p -form, it suffices to evaluate these p -forms on p -tuples of standard basis vectors (after choosing a chart for M), and there is at least one i for which the ∂_{t_i} is not in the list. So by naturality we can perform this evaluation in the submanifold defined by $t_i = 0$, in which these two p -forms agree. \square

By naturality, we may assume $M = W$ is a vector space and x is the origin. The value of $\omega_W(\alpha)$ at the origin is given by a map

$$\tilde{\omega}_W : J^N(W; W^* \otimes V) \rightarrow \bigwedge^p W^*,$$

where $J^N(W; W^* \otimes V)$ is the space of N -jets of elements of $\Omega^1(W; V)$. This is a finite dimensional vector space, given explicitly by

$$J^N(W; W^* \otimes V) = \bigoplus_{j=0}^N \text{Sym}^j W^* \otimes W^* \otimes V.$$

Under this decomposition, the j^{th} piece captures the j^{th} derivatives of α . Throughout the proof, we view $\text{Sym}^j W^*$ as a *quotient* of $(W^*)^{\otimes j}$, hence every function on $\text{Sym}^j W^*$ is in particular a function on $(W^*)^{\otimes j}$.

At this point, everything else follows from the fact that $\tilde{\omega}_W$ is functorial in W , and in particular $\text{GL}(W)$ -invariant.

Lemma 16. *$\tilde{\omega}_W$ is a polynomial function.*

This lemma is true in much greater generality — it holds for any set-theoretic natural transformation between “polynomial functors” $\text{Vec} \rightarrow \text{Vec}$. Here a set-theoretic natural transformation is a natural transformations of the underlying set-valued functors. This is a polynomial version of the fact that a natural transformation between additive functors is necessarily additive, because being additive is a *property* and not a structure.

Proof. Write

$$F(W) = \bigoplus_{j=0}^N \text{Sym}^j W^* \otimes W^* \otimes V, \quad G(W) = \bigwedge^p W.$$

We think of these as a functor $\text{Vec} \rightarrow \text{Vec}$ (with V fixed). The point is that for $f \in \text{Hom}_{\text{Vec}}(W, W')$, the functions $F(f), G(f)$ are polynomial in f . This together with naturality will force $\tilde{\omega}_W$ to be polynomial as well.

To show that $\tilde{\omega}_W$ is polynomial, we have to show that if $v_1, \dots, v_n \in F(W)$, then $\tilde{\omega}_W(\sum \lambda_i v_i)$ is a polynomial function in $\lambda_1, \dots, \lambda_n$. Without loss of generality, we may assume each v_i lives in the $(j_i - 1)$ th summand (so that the summand has j_i tensor powers of W^*).

Fix a number j such that $j_i \mid j$ for all i . We first show that $\tilde{\omega}_W(\sum \lambda_i^j v_i)$ is a polynomial function in the λ_i 's.

Let $f : W^{\oplus n} \rightarrow W^{\oplus n}$ be the map that multiplies by λ_i^{j/j_i} on the i th factor, and $\Sigma : W^{\oplus n} \rightarrow W$ be the sum map. Consider the commutative diagram

$$\begin{array}{ccccc} F(W^{\oplus n}) & \xrightarrow{F(f)} & F(W^{\oplus n}) & \xrightarrow{F(\Sigma)} & F(W) \\ \downarrow \tilde{\omega}_{W^{\oplus n}} & & \downarrow \tilde{\omega}_{W^{\oplus n}} & & \downarrow \tilde{\omega}_W \\ G(W^{\oplus n}) & \xrightarrow{G(f)} & G(W^{\oplus n}) & \xrightarrow{G(\Sigma)} & G(W) \end{array}$$

Let $\tilde{v}_i \in F(W^{\oplus n})$ be the image of v_i under the inclusion of the i th summand. Then $x = \sum \tilde{v}_i$ gets sent along the top row to $\sum \lambda_i^j v_i$. On the other hand, $\tilde{\omega}_{W^{\oplus n}}(x)$ is some element in $G(W^{\oplus n})$, and whatever it might be, the image along the bottom row gives a polynomial function in the λ_i^{j/j_i} , hence in the λ_i . So we are done.

We now know that for any finite set v_1, \dots, v_n , we can write

$$\tilde{\omega}_W(\lambda_1^j v_1 + \dots + \lambda_n^j v_n) = \sum_{r_1, \dots, r_m} a_R \lambda_1^{r_1} \dots \lambda_n^{r_n}.$$

We claim each r_i is a multiple of j (if the corresponding a_R is non-zero). Indeed, if we set $\lambda_i = (\mu_i^j - \nu_i^j)^{1/j}$, then the result must be a polynomial in the μ_i and ν_i as well, since it is of the form $\tilde{\omega}_W(\sum \mu_i^j v_i - \nu_i^j v_i)$. But $\sum a_R (\mu_1^j - \nu_1^j)^{r_1/j} \dots (\mu_n^j - \nu_n^j)^{r_n/j}$ is polynomial in μ_i, ν_i if and only if $j \mid r_i$.

Now by taking j th roots, we know $\tilde{\omega}_W(\sum \lambda_i v_i)$ is polynomial in the λ_i when $\lambda_i \geq 0$. That is, it is polynomial when restricted to the cone spanned by the v_i 's. But since the v_i 's are arbitrary, this implies it is polynomial everywhere. \square

Lemma 17. *Any non-zero $\mathrm{GL}(W)$ -invariant linear map $\otimes^M W^* \rightarrow \bigwedge^p W^*$ has $M = p$ and is a multiple of the anti-symmetrization map. In particular, any such map is anti-symmetric.*

Proof. For convenience of notation, replace W^* with W . Since the map is in particular invariant under $\mathbb{R}^\times \subseteq \mathrm{GL}(W)$, we must have $M = p$. By Schur's lemma, the second part of the lemma is equivalent to claiming that if we decompose $W^{\otimes p}$ as a direct sum of irreducible $\mathrm{GL}(W)$ representations, then $\bigwedge^p W$ appears exactly once. In fact, we know the complete decomposition of $W^{\otimes p}$ by Schur–Weyl duality.

Let $\{V_\lambda\}$ be the set of irreducible representations of S_p . Then as an $S_p \times \mathrm{GL}(W)$ -representation, we have

$$W^{\otimes p} = \bigoplus_{\lambda} V_\lambda \otimes W_\lambda,$$

where $W_\lambda = \text{Hom}_{S_p}(V_\lambda, W^{\otimes p})$ is either zero or irreducible, and are distinct for different λ . Under this decomposition, $\bigwedge^p W$ corresponds to the sign representation of S_p . \square

So we know $\tilde{\omega}_W$ is a polynomial in $\bigoplus_j \text{Sym}^j W^* \otimes W^* \otimes V$, and is anti-symmetric in the W^* . So the only terms that can contribute are when $j = 0$ or $j = 1$. In the $j = 1$ case, it has to factor through $\bigwedge^2 W^* \otimes V$. So $\tilde{\omega}_W$ is polynomial in $(W^* \otimes V) \oplus (\bigwedge^2 W^* \otimes V)$. This exactly says $\omega_W(\alpha)$ is given by wedging together α and $d\alpha$ (and pairing with elements of V^*).