ON GERBS

David Saumitra Chatterjee

Trinity College, Cambridge

A dissertation submitted towards the degree of Doctor of Philosophy at the University of Cambridge
June 1998

The line that can be drawn is not the true line.

—Lao Zi.

Contents

Chapter 1

Introduction

We define a new approach to the concept of a gerbe and illustrate it with applications to differential and algebraic geometry.

Our aim is to obtain a *geometric realisation of abelian cohomology.* Just as nowhere-zero functions in $\mathrm{H}^0(X;\underline{\mathbb{C}}^*)$ (or $\mathcal{O}^*$) on a smooth manifold X have a winding number realising classes in $\mathrm{H}^1(X;\mathbb{Z})$, and smooth line-bundles classified by $\mathrm{H}^1(X;\underline{\mathbb{C}}^*)$ realise classes in $\mathrm{H}^2(X;\mathbb{Z})$, by a "gerbe" we mean some more-or-less concrete realisation of $\mathrm{H}^2(X;\underline{\mathbb{C}}^*)$ and $\mathrm{H}^3(X;\mathbb{Z})$.

Then we would expect such a thing to have links to "curvature" 3-forms, codimension-3 submanifolds, holonomy over surfaces and the like. Physicists have long been comfortable with the idea [11] that when they find a field which looks like a closed integer-valued differential 2-form which is well-defined up to the exterior derivative of a 1-form, what they have actually uncovered is a line bundle with connection (or a vector bundle, given more exotic coefficients). In recent years, all sorts of obscure fields have been cropping up in investigations into string theory, amongst which are closed integral $(n+1)$-forms, well-defined up to $\mathrm{d}(n\text{-form})$. The string community seems to have become quite quite blasé about such putative "n-gerbs". There is clearly a need to understand their underlying character.

Gerbes were originally defined by Giraud [18] in the early 1970s in an attempt to understand *non-abelian* second cohomology, invoking the most rarefied aspects of algebraic geometry such as stacks, torsors and toposes. The 1993 book of Brylinski [5], taking inspiration from Deligne, tries to use abelian gerbes to attack $\mathrm{H}^3(\mathbb{Z})$ on a smooth manifold. This retains a degree of the impenetrability of the original, with torsors and sheaves of categories omnipresent, but is the starting point for current research. We can compare this with knowing only the abstract coherent-sheaf definition of a line bundle: whilst the algebraic approach has of course proven its worth in modern geometry, it would be nice to have a more hands-on description.

In reaction to this, Murray [37] defines "bundle-gerbes". Just as with a line bundle, it is possible to describe a gerbe in terms of a concrete total space with certain properties sitting over the base manifold X, but inevitably the fibres (which are spaces of paths) are infinite-dimensional; and this approach also has problems with defining the appropriate equivalence classes. Before turning to gerbes, Brylinski too considers infinite-dimensional bundles, starting from the projective Hilbert space bundles of the 1960s work of Dixmier and Douady [12] concerning $\mathrm{H}^3(\mathbb{Z})$. Atiyah has remarked that, after 2-forms, the next concept of "curvature" seems to require an infinite-dimensional setting rather than stepping up by some finite degree: this was in the context of path integrals and quantum field theory, but it is also in tune with these various attempts at definitions of gerbes.

In contrast, our work advocates a new approach in seeking to understand what is involved in a "gerbe"—we follow a third route to line bundles, that of local trivialisations and transition functions. The underlying idea is nothing more than to take Čech theory quite literally, guided by the principle that gerbes are to line bundles what line bundles are to functions. We believe this offers the simplest and most explicit way to get to grips with what a gerbe entails. We can show that our structures are (particular expressions of) gerbes, and our equivalence classes correspond with equivalence classes of gerbes. To emphasise their stripped-down nature, we call our structures not gerbes but *gerbs*.

It seems that the tendencies—at least implicit in both Brylinski's and Murray's work—towards seeing gerbes as essentially infinite dimensional (or categorical) and principal (rather than vector) structures are in fact not at all vital. Whilst comparison of our method with those of [5] and [37] is clearly of interest, we have avoided devoting space to line-by-line translations, feeling it more worthwhile to develop new examples in which the efficacy of the gerb definition can be tested.

Since our thesis is that it is possible to do many things on manifolds without invoking stacks or torsors, our technology is deliberately low-powered throughout, requiring little more than Čech cohomology and a basic familiarity with line bundles and their connections. Our use of algebraic geometry is gentlemanly and transcendental rather than professional and algebraic. (To lay bare our religious affiliation [15], we declare that the word "torsor" has no further role in this work.) Our ambition is to be simple but not simplistic; and thereby perhaps to offer a useful alternative approach to those of [5] and [37].

A guide to the chapters

Essentially we define in the smooth and the holomorphic cases three structures—gerbs, their connections, and their sections—and the natural equivalences between them. This theory takes up about a quarter of the dissertation. (The fact that it can be covered so rapidly is we hope a virtue.)

Chapter 2 lays out the basics of smooth abelian gerbs and connections, works through the appropriate notions of gauge transformation or equivalence, and offers a few minor examples to consolidate the definitions. This includes a brief comment on "principal" gerbs (2.3.3).

Chapter 3 defines smooth gerb sections, which in deference to [5] we call "objects" since these are the most natural link to category-theoretic gerbes. (Compare this with moving between a vector bundle and a sheaf, its sheaf of sections.) We outline this step, but not in depth since that would amount to a regurgitation of most of chapter 5 of Brylinski.

The case of holomorphic gerbs is in broad terms the same as in these first two chapters, but unsurprisingly there are more subtle questions of existence to consider, and these are taken up in chapter 5. Before that, however, we wish to pursue a key example in depth. Chapter 4 is concerned with holomorphic gerbs, but on deliberately simple spaces ($\mathbb{P}^3$ and $\mathbb{C}^3$) so that the technicalities of chapter 5 are not a concern. The example is a twistor-theoretic one, in which we transform between a purely holomorphic setting and one with infinitesimal structure, or connection. Twistor theory has proved so natural in the context of vector bundles in a wide range of examples [3, 49] that we view this as a significant test: given a class in $\mathrm{H}^2(\mathcal{O}^*)$, to what does this correspond on the other side? If we believe $\mathrm{H}^2(\mathcal{O}^*)$ to be related to holomorphic gerbs (as we do, in chapters 2 and 5) then its transform is showing us the essential content of a gerb with connection—and so our definitions had better agree with it. The outcome is highlighted in theorems (4.4.1) and (4.5.1): not only are gerbs quite appropriate fields on $\mathbb{P}^3$, but "n-gerbs" are similarly efficacious on $\mathbb{P}^{n+1}$.

Chapter 6 begins from the idea that, given that $\mathrm{H}^3(\mathbb{Z})$ is identified by Poincaré with $\mathrm{H}_{n-3}(\mathbb{Z})$, we can hope for gerbs to have interesting links with real codimension-3 submanifolds. The smooth category is too loose to offer much, but there are discoveries to be made once we fix a metric and consider harmonic theory. It turns out that we can move between a certain natural type of gerb with flat connection; a submanifold which is the singularity of an abelian monopole; and a torus which is a direct analogue of the Jacobian variety of line-bundles. Whilst theorem (6.4.2) is perhaps not quite as compelling as our two twistor-theoretic chapters, we can certainly claim that it is a meaningful extension of its starting place, which is Brylinski's elementary example (chapter 7 of [5]) of a point in S^3. (We should also declare a debt to Kodaira's paper [29], an accidentally-uncovered gem whose influence on chapter 6 is by now well-buried since at the requisite times—fifty years apart—neither the author nor the reader was aware of de Rham's concept of a current.)

Chapter 7 is concerned with holonomy. Given the notion of a "error 2-form" from chapter 3, a definition of the holonomy of a surface is immediate. Brylinski on the other hand has a rather ornate approach to the holonomy of a loop: on the space of loops, he defines a line bundle instead of a function. We explore this idea further, but since we end up at a limited and well-known location (the correspondence between monopoles on S^3 and line bundles on the quadric surface) we cannot view this chapter as fully satisfactory. Our constructions seem interesting in their own right, however, and we offer them in the hope that a more compelling application awaits.

We return to twistor theory for chapter 8, looking at the Ward correspondence between even-dimensional quadrics and their collections of linear subspaces. This is a more complicated setting than chapter 4, and we cannot handle the general n-gerb case in full. We can certainly deal with gerbs themselves though, finding a natural correspondence between holomorphic gerbs on one side and anti-self-dual connections on the other. We also lay out two approaches to higher-order structures, which seem to work convincingly enough to conjecture that there is indeed a general Ward correspondence between gerbs of arbitrary degree and anti-self-dual connections.

We conclude in chapter 9 with outlines of further directions. There is unsurprisingly a notion of divisor for holomorphic gerbs—indeed, there are two notions. We have considered no sufficiently compelling example to let us choose between the two, and so we merely lay out some ideas for consideration. Secondly, and in reaction to this, we consider an alternative and currently-fashionable type of "divisor": a special Lagrangian submanifold of a Calabi-Yau 3-fold. Such spaces are sufficiently rigid that there ought to be some obvious gerb associated with them, but we can offer only the vaguest of hints as to how to construct it. No doubt the string theorists can think of something.

Citations in square brackets [1] are listed in the bibliography which follows the final chapter. Round brackets (2.1.1) refer to chapters in the text and sections within them.

Acknowledgements

As is proper and traditional, my first thanks are due to my supervisors: initially Michael Atiyah, who opened a store of geometric treasures in my first year, beginning with the Klein quadric with which this thesis ends; and subsequently and especially Nigel Hitchin, who guided me through all of the work contained here and more, and was never anything less than committed and positive even when I was not. I do not expect to discover any more beautiful ideas in my life than revealed to me by these two people, and I am most grateful to them.

Equally inspiring has been the presence of Graeme Segal, who always seems to have the best way of understanding anything when I have the wit to follow. Nick Shepherd-Barron, Dennis Barden and Pelham Wilson—who first taught me what pure maths involves in Michaelmas 1989—have all coped with good humour with a series of foolish questions.

The really foolish questions were inflicted on others. I thank Richard Hinchliffe, for long and valuable time-series on the effects of vodka and peach schnapps; Chris Reynolds and Phil Boalch, for the physics and dynamics of discs (with and without X-rays); Jochen Theis, for keeping half the department library in our office; Balázs Szendroi, for practice with accents; John Allen, for extreme violence; Richard Hind, for a most illuminating seminar on Gromow theory; Roger Moore, for research on European taxation policies (and with best wishes to Kirsten); Birte Feix, for professor-monitoring services; Ginge Laurie, for inducting me into to the Porphyry Pillar; in prospect, Mark Wigoder-Daniels, for coming back from wherever he may be; Valerio Toledano-Laredo, for the other half of the department library; Dave Morris, for postcards of dinosaurs; Charlie Wartnaby, for looking after a very good friend of mine (and with best wishes to Clare and especially to Georgie); Gareth McAleese, for making everyone else seem perfectly normal; and my brother, for not finishing before me.

I also wish to record my gratitude to Nick Denyer and Diana Paterson of Side C, for looking after me for so many years; the Master and Fellows of Trinity College, for a research scholarship and an internal graduate studentship; my parents, without whom I expect I'd be somewhere else completely; and the people of Northern Ireland, for their courage on Friday 22 May 1998.

And as is proper and traditional, my final thanks are due with greatest love and devotion to the one who makes it all worthwhile. Yee Yen, I haven't forgotten …

May all beings find happiness and freedom from suffering.

Chapter 2

The definition of a gerb

All structures are oriented unless otherwise declared. "Smooth" is used throughout to mean C^∞.

2.1 Locally-trivialised gerbs

We work over a smooth real finite-dimensional manifold X.

Definition 2.1.1 (Gerb) *A (smooth, abelian, locally-trivialised) gerb*

$$\mathcal{G}(I, \Lambda, \theta)$$

is defined by the following data—

- *An open cover of X*

$$\{U_i : i \in I\} \quad \textit{with} \quad \bigcup_I U_i = X$$

  (and we shall write for instance $U_{i,j}$ to mean $U_i \cap U_j$);

- *A smooth complex line-bundle Λ_i^j existing over $U_{i,j}$ for each ordered pair $(i, j), i \neq j$, such that Λ_i^j and Λ_j^i are dual to each other;*

- *For each ordered triple of distinct indices (i, j, k), a smooth nowhere-zero section*

$$\theta_{i,j,k} \in \Gamma(U_{i,j,k}; \Lambda_i^j \otimes \Lambda_j^k \otimes \Lambda_k^i)$$

  such that the sections $\theta_{i,j,k}$ of reorderings of a triple (i, j, k) are related in the natural way.

One further constraint is to be obeyed: on four-fold intersections we require that

$$\delta\theta = 1.$$

To clarify: over any four-fold intersection $U_{i,j,k,l}$ we can tensor together the four sections $\theta_{i,j,k}, \ldots,$ to give a trivialisation $\delta\theta$ of the line-bundle

$$\begin{aligned} \delta^2\Lambda \quad = \quad & \left(\Lambda_i^j \otimes \Lambda_j^k \otimes \Lambda_k^i\right) \bigotimes \left(\Lambda_i^j \otimes \Lambda_j^l \otimes \Lambda_l^i\right)^{-1} \\ & \bigotimes \left(\Lambda_i^k \otimes \Lambda_k^l \otimes \Lambda_l^i\right) \bigotimes \left(\Lambda_j^k \otimes \Lambda_k^l \otimes \Lambda_l^j\right)^{-1}. \end{aligned}$$

But thanks to the duality condition $\Lambda_a^b = (\Lambda_b^a)^{-1}$, this untidy bundle is *canonically* trivial ($\delta^2 = 0$). We are insisting that $\delta\theta$ is in fact this canonical section.

Now we must begin to free ourselves from the self-imposed prison of our choice of cover.

Definition 2.1.2 (Refinement) *A refinement of a local trivialisation is—*

- *A refinement of the cover, ie. a new cover and a map*

$$\{V_a : a \in A\} \quad \textit{with} \quad \bigcup_A V_a = X,$$
$$r : A \to I \quad : \quad V_a \subset U_{r(a)} \ (\forall a);$$

- *Line-bundles*

$$\Lambda_a^b := \begin{cases} \Lambda_{r(a)}^{r(b)}|_{V_{a,b}} & \textit{if } r(a){\neq}r(b), \\ \textit{trivial, with trivialising section } 1_a^b & \textit{if } r(a){=}r(b); \end{cases}$$

- *Sections*

$$\theta_{a,b,c} := \begin{cases} \theta_{r(a),r(b),r(c)} & \textit{if all different,} \\ 1_a^b \otimes 1_b^c \otimes 1_c^a & \textit{if } r(a){=}r(b){=}r(c), \\ \textit{eg. } 1_a^b \otimes \textit{can} & \textit{if eg. } r(a){=}r(b){\neq}r(c) \end{cases}$$

  over each $V_{a,b,c}$.

In the last line *can* is to mean the canonical section of $\Lambda_{r(a)}^{r(c)} \otimes \Lambda_{r(c)}^{r(a)}$.

Proposition 2.1.3 *A refinement of a gerb is a gerb.*

Proof—We must check that our new sections obey $\delta\theta = 1$. This is a straightforward matter of testing the various cases for which $r(a), r(b), r(c), r(d)$ are equal or distinct. □

Proposition 2.1.4 *A locally trivialised gerb naturally gives rise to a class in*

$$\mathrm{H}^2(X; \underline{\mathbb{C}}^*).$$

Here $\underline{\mathbb{C}}^*$ is the sheaf of smooth nowhere-zero complex functions on X. Recall that, given the exponential sequence

$$0 \to \mathbb{Z} \xrightarrow{2\pi i} \underline{\mathbb{C}} \xrightarrow{\exp} \underline{\mathbb{C}}^* \to 0$$

and given the fact that, being soft [21], $\underline{\mathbb{C}}$ has no cohomology, this group is equal to

$$\mathrm{H}^2(X; \underline{\mathbb{C}}^*) \quad \cong \quad \mathrm{H}^3(X; 2\pi i\mathbb{Z}).$$

Proof—Refine the cover such that now all the Λ_i^j are trivialisable, and take trivialisations. Using these as coordinates, θ becomes a Čech cocyle

$$t \in \mathrm{Z}^2(X; \underline{\mathbb{C}}^*)$$

thus inducing the cohomology class.

Different trivialisations shift the cocyle by

$$\delta \mathrm{C}^1(X; \underline{\mathbb{C}}^*).$$

A different refinement gives the same class in standard Čech fashion [21], by going to a common refinement. So the class of $\mathcal{G}(I, \Lambda, \theta)$ is independent of choices. □

Gerbs come with definitions of restriction and pull-back and the like, but it seems unnecessary to write these down since they derive directly from the well-known behaviour of bundles, sections and Čech cocycles. We shall simply use such properties as needed.

Why do we not insist that the bundles Λ be trivialised? (In many proofs, we shall refine so as to ensure this.) The reason is that many natural examples of gerbs use non-trivial bundles: for instance, in (2.3.6) we consider a degree-1 line bundle on a punctured ball, and to break this up into subsets would be unnecessarily artificial.

0-Gauge transformations

There is an obvious concept of equivalence between trivialisations.

Definition 2.1.5 (Equivalence of locally-trivialised gerbs) *Two gerbs*

$$\begin{cases} \mathcal{G}(I, \Lambda, \theta) \\ \mathcal{H}(J, M, \eta) \end{cases}$$

are equivalent, if there exists a common refinement A of their covers I, J, and isomorphisms

$$\Lambda_a^b \xrightarrow{\cong} K_a^b \quad over \quad V_{a,b}$$

on the pairwise intersections of this common refinement, such that the isomorphisms induce

$$\theta_{a,b,c} \mapsto \eta_{a,b,c}.$$

This is an equivalence relation. If we are given two local trivialisations over the same cover, and given particular isomorphisms inducing an equivalence, we might reasonably say the trivialisations describe one and the same gerb. Given a single local trivialisation of a gerb, a *0-gauge transformation* of the gerb is an automorphism of each of the Λ_i^j, ie. a family of line-bundle gauge-transformations. This gives a new local trivialisation that is equivalent to the original.

Theorem 2.1.6 *The collection of equivalence classes of gerbs is canonically identified with*

$$\mathrm{H}^2(X; \underline{\mathbb{C}}^*).$$

Proof—It remains to show that equivalent trivialisations induce the same class; and in the other direction, that Čech cocycles induce gerbs; and that equivalent cocycles give rise to equivalent gerbs.

For the first: go to a mutual refinement in which we can take bases for the line-bundles on each $V_{a,b}$. Then we clearly have a Čech equivalence of the cocycles representing θ, η, and as before altering the refinement or the bases does not alter the class.

For the second: an open cover $\{U_i\}$ and a cocycle

$$t_{i,j,k} \in \mathrm{Z}^2(X; \underline{\mathbb{C}}^*)$$

define an equivalence class of locally-trivialised gerbs, as follows. Choose trivial bundles Λ_i^j with bases λ_i^j. Then we define the section

$$\theta_{i,j,k} \ := \ t_{i,j,k} \lambda_i^j \otimes \lambda_j^k \otimes \lambda_k^i$$

which obeys $\delta\theta = 1$ since $\delta t = 1$. Different choices for λ_i^j give equivalent trivialisations of the same gerb.

Finally: that a different cover and cocycle lying in the same Čech cohomology class gives an equivalent gerb is yet another repetition of the same argument. □

Note that the proof also demonstrates that gerbs exist, on any X.

Definition 2.1.7 (Global trivialisation of a gerb) *A gerb $\mathcal{G}(I, \Lambda, \theta)$ is globally trivialised by displaying bases λ_i^j for each line-bundle Λ_i^j such that $\delta\lambda = \theta$, ie. such that on each $U_{i,j,k}$*

$$\lambda_i^j \otimes \lambda_j^k \otimes \lambda_k^i = \theta_{i,j,k}.$$

Definition 2.1.8 (Trivial gerb) *A gerb is trivial, if equivalent to a globally-trivialised gerb.*

Proposition 2.1.9 *A gerb is trivial, iff its class in $\mathrm{H}^2(X; \underline{\mathbb{C}}^*)$ is zero.*

Proof—A gerb is trivial, iff it has a globally-trivialisable refinement. Such a global trivialisation can be represented by the trivial cocycle $1_{i,j,k}$ on its cover I. □

We remark that, given two gerbs over X and an equivalence between them, by going to a suitable common refinement we can define their difference in the obvious manner: we obtain the ratio of the line-bundles on each V_a^b as a canonically-trivial line-bundle; and these trivialisations have coboundary section

$$\theta^{-1} \otimes \eta.$$

Proposition 2.1.10 *The difference of two equivalent gerbs is a trivial gerb; and if given a particular equivalence between the two, the difference comes with a particular global trivialisation.* □

This is in exact analogy with isomorphic line-bundles. Further, we shall later discuss the analogue of a section of a line-bundle (3.1.1). Then a trivial gerb will be seen to have global "sections"; and global sections of a trivial gerb with given global trivialisation are naturally global line-bundles (3.1.3)—just as global sections of a trivialised line-bundle are global functions.

The module structure of $\mathrm{H}^2(X; \underline{\mathbb{C}}^*) \cong \mathrm{H}^3(X; \mathbb{Z})$ over $\mathbb{Z}$ is expressed by a simple operation.

Definition 2.1.11 (Tensor product) *The product of two gerbs is obtained by refining to a common local trivialisation, and tensoring the various pairs of line-bundles and sections in the obvious way.*

Proposition 2.1.12 *The product of two gerbs is a gerb. Tensoring induces the $\mathbb{Z}$-module addition of equivalence classes in* $\mathrm{H}^2(X;\underline{\mathbb{C}}^*)$.

Proof—That products are gerbs is immediate. It is also immediate that tensoring $\mathcal{G}$ by either of two equivalent gerbs gives isomorphisms between the bundles forming the two products, making them equivalent. So the construction passes to equivalence classes.

On the module structure: on a common refinement adequate to trivialise all Λ_i^j, tensoring amounts to multiplication of Čech representatives $s_{i,j,k}$ and $t_{i,j,k}$, thus inducing $\mathbb{Z}$-addition. □

2.2 Differential structures

We must break up the differential geometry of gerbs into two steps. The first has a simple description in terms of standard geometry; the second remains in general abstract, but we can hope to make it more concrete in particular situations.

The idea can be described in local-coordinate terms. Consider a locally-trivialised *line-bundle* with transitions f_i^j. Then a connection is given by smooth complex local 1-forms A_i such that

$$\delta A - \mathrm{d}\log f = 0$$

and its curvature is then $F = \mathrm{d}A$.

We should be thinking in terms of the Čech-de Rham double complex

$$\mathrm{C}^p(X;\mathcal{A}^q)$$

writing $\mathcal{A}^q$ for the sheaf of smooth complex q-forms on X. Equivalence classes of bundles with connection become hypercohomology classes in

$$\mathbb{H}^1(X;\ \underline{\mathbb{C}}^* \xrightarrow{\mathrm{d}\log} \mathcal{A}^1).$$

What we shall do for gerbs amounts to nothing more than this: given a cocycle $t_{i,j,k}$ representing θ, choose 1-forms α_i^j such that

$$\delta\alpha + \mathrm{d}\log t = 0.$$

(The sign is to be consistent with the standard total differential $\delta + (-1)^p\mathrm{d}$ on the Čech-de Rham complex.) Then we have the second step of choosing local 2-forms β_i such that

$$\delta\beta - \mathrm{d}\alpha = 0$$

which leaves us with a global 3-form curvature $\Omega = \mathrm{d}\beta$.

0-Connections

Definition 2.2.1 (0-Connection) *On a locally-trivialised gerb, a 0-connection is a family of line-bundle connections*

$$\nabla_i^j \text{ on } \Lambda_i^j$$

such that the induced connection on each threefold intersection obeys

$$\nabla\theta = 0.$$

Note that a 0-transform of the gerb induces line-bundle gauge transforms of the ∇_i^j. We view this transformed 0-connection as *equivalent* to the original.

What is the difference between two 0-connections on a given gerb? Fix some local trivialisation over X. Then the difference is a collection of 1-forms

$$\nabla' - \nabla = \alpha_i^j \in \mathrm{C}^1(X; \mathcal{A}^1).$$

Since $(\nabla - \nabla')(\theta) = 0$ we know α is in fact a cocycle. In the smooth case we can go further

$$\mathrm{Z}^1(X; \mathcal{A}^1) = \delta \mathrm{C}^0(X; \mathcal{A}^1)$$

which is pertinent to (2.2.5).

Proposition 2.2.2 *Every gerb (expressed over some open cover I) has a 0-connection (expressed on that same cover).*

Proof—It suffices to observe that the sheaf $\mathcal{A}^1$ is soft [21].

To expand: choose an arbitrary connection ∇_i^j for each bundle Λ_i^j. Then on $U_{i,j,k}$ we find

$$\nabla\theta = \eta \otimes \theta$$

for some $\eta \in \mathrm{C}^2(X; \mathcal{A}^1)$ with respect to the cover I.

Further, $\delta\eta = 0$. This is because on 4-fold intersections $U_{i,j,k,l}$ by definition $\delta\theta$ is equal to the canonical section *can*; but the induced connection $\delta^2\nabla$ is necessarily zero on *can*, whether or not $(\delta\nabla)\theta = 0$ on 3-fold intersections. (We see this by taking local bases, for instance.)

Now choose a smooth real partition of unity ρ subordinate to I. In standard fashion this gives us η as a coboundary

$$\zeta_i^j := \sum_k \rho_k \cdot \eta_{i,j,k} \qquad \Rightarrow \qquad \delta\zeta = \eta$$

thus offering us

$$\nabla' := \nabla - \zeta \qquad \Rightarrow \qquad \nabla'\theta = 0$$

as a 0-connection on the given cover. □

If we were discussing the holomorphic category, there would be an obstruction here to the existence of a holomorphic 0-connection. We shall consider this in chapter 5.

1-Connections and 1-gauge-transformations

Since we know that gerbs are related to $\mathrm{H}^2(X; \underline{\mathbb{C}}^*)$, we expect some sort of 3-form "curvature" in $\mathrm{H}^3(X; \mathbb{Z})$. The general way to reach this stage is disappointingly short on geometry—

Definition 2.2.3 (1-Connection) *Given a locally-trivialised gerb with 0-connection, a 1-connection is any choice of local 2-forms*

$$\beta_i \in \mathrm{C}^0(X; \mathcal{A}^2)$$

satisfying

$$\delta\beta = F.$$

Here F is the 1-cocycle of curvature 2-forms of the line-bundles

$$F_i^j := \text{Curvature of } \nabla_i^j.$$

Proposition 2.2.4 *1-connections exist for any 0-connection, on the given cover.*

Proof—Just as for 0-connections, a subordinate partition of unity suffices. □

For a given 0-connection, two 1-connections clearly differ by a global 2-form.

Definition 2.2.5 (1-Gauge-transformation) *Given a locally-trivialised gerb, a 1-transform is a cochain*

$$\gamma \in \mathrm{C}^0(X; \mathcal{A}^1)$$

which acts as follows—

- *The bare gerb is left unscathed.*
- *A 0-connection, if given, transforms as*

$$\nabla \mapsto \nabla + \delta\gamma.$$

- *A 1-connection, if given over the 0-connection, transforms as*

$$\beta \mapsto \beta + \mathrm{d}\gamma.$$

We already know that any two 0-connections are related by a 1-transform (as remarked before (2.2.2)). The purpose of the definition is to give the appropriate concept of equivalence of 1-connections. (Note that if we say "gerb with 1-connection", we mean to say "gerb with 0-connection and compatible 1-connection".)

Definition 2.2.6 (Equivalence of gerbs with 1-connection) *Suppose we are given two locally-trivialised gerbs $\mathcal{G}, \mathcal{H}$ over X with two 0-connections, and 1-connections on the 0-connections. Then they are equivalent, given three conditions—*

- *The gerbs are equivalent.*

View this as a 0-transform letting us identify the gerbs. We now have two 0- and 1-connections on one gerb. We insist that there exist a 1-transform γ obeying—

- $$(\nabla')_i^j = \nabla_i^j + \delta\gamma_i.$$

- $$(\beta')_i = \beta_i + \mathrm{d}\gamma_i.$$

Consider the case of a fixed gerb and a fixed 0-connection. Then the set of 1-connections equivalent to any given one is clearly

$$(\beta')_i = \beta_i + \mathrm{d}\Gamma$$

for any *global* 1-form Γ.

Theorem 2.2.7 *The set of equivalence classes of gerbs with 1-connection is*

$$\mathbb{H}^2(X;\ \underline{\mathbb{C}}^* \xrightarrow{\mathrm{d}\log} \mathcal{A}^1 \xrightarrow{\mathrm{d}} \mathcal{A}^2).$$

Proof—This is no deeper than previous remarks. The definitions deliberately slot into the definition of hypercohomology of a complex [21]. □

Curvature 3-forms and Chern classes

Definition 2.2.8 (Curvature and Chern form) *The curvature of a locally-trivialised gerb with 1-connection is the globally-defined closed smooth 3-form*

$$\Omega := \mathrm{d}\beta_i.$$

The Chern form of this 1-connection is

$$\omega := \frac{i}{2\pi}\Omega.$$

Proposition 2.2.9 *Two 1-connections on the same gerb define the same curvature 3-form if and only if they are equivalent.*

Proof—If equivalent, note $\mathrm{d}^2\gamma = 0$. If $\mathrm{d}(\beta - \beta')_i = 0$, then refine and use the Poincaré lemma to find suitable γ_i. □

Theorem 2.2.10 *The de Rham class of ω in $\mathrm{H}^3_{\mathrm{dR}}(X)$ is independent of choice of trivialisation, of 0-connection and of 1-connection, and equals the image of*

$$[\mathcal{G}] \in \mathrm{H}^2(X;\underline{\mathbb{C}}^*)$$

in the long exact sequence of

$$0 \longrightarrow \mathbb{Z} \xrightarrow{2\pi i} \underline{\mathbb{C}} \xrightarrow{\exp} \underline{\mathbb{C}}^* \longrightarrow 0$$

(followed by the inclusion $\mathbb{Z} \hookrightarrow \mathbb{C}$).

Proof—Standard. □

Perhaps it will be convenient to write $\mathrm{H}^p_{\mathrm{dR}}(X;\mathbb{Z})$ for the image in de Rham cohomology of $\mathrm{H}^p(X;\mathbb{Z})$. We shall try to be consistent in writing $[\mathcal{G}]$ for the class in $\mathrm{H}^2(X;\underline{\mathbb{C}}^*)$, and describe the class

$$[\omega] \in \mathrm{H}^3_{\mathrm{dR}}(X;\mathbb{Z}) \subset \mathrm{H}^3(X;\mathbb{C})$$

as the *Chern class* $c(\mathcal{G})$ of the gerb. In this smooth case the exponential long exact sequence degenerates, and moving to the Chern class loses nothing but torsion.

2.3 The simplest examples

We start with gerbs of the trivial class in $\mathrm{H}^2(X;\underline{\mathbb{C}}^*)$.

2.3.1 The vacuous gerb

Note that we can create gerbs using covers of less than three sets: we simply have no data to specify on intersections that do not exist.

For instance, given a single set $U = X$ there is a single vacuous gerb trivialisable on this cover, with nothing to specify; and it comes with a unique vacuous 0-connection. The collection of 1-connections on this trivialisation is the set of all 2-forms β on U, and two 1-connections are 1-gauge equivalent if they differ by $\mathrm{d}\gamma$ for some global 1-form γ over U. Of course, all choices yield the same class of curvature

$$[\mathrm{d}\beta] \quad = \quad 0$$

and two choices yield identical Ω iff they are equivalent.

This example will feature in chapters 4 and 8, in which we study such local gerbs under various twistor correspondences.

2.3.2 Coboundary gerbs

Now over a general cover U_i we consider some representative

$$t_{i,j,k} \quad = \quad \delta s_{i,j}$$

for $[\mathcal{G}] = 0 \in \mathrm{H}^2(X;\underline{\mathbb{C}}^*)$. Given the gerb represented by $t_{i,j,k}$ as in the proof of (2.1.4), that t is a coboundary is equivalent to saying that there exist bases λ_i^j such that $\theta = \delta\lambda$, ie. we have a globally-trivialisable gerb as in (2.1.7).

Given such a trivialising basis for the Λ_i^j, any 0-connection is a coboundary of 1-forms

$$\alpha_i^j = \delta\zeta_k \quad : \quad \zeta_k \in \mathrm{C}^0(\mathcal{A}^1)$$

and fixing the 0-connection specifies ζ up to a global 1-form.

Choosing such a ζ gives a possible 1-connection for this 0-connection

$$\beta_i \quad := \quad \mathrm{d}\zeta_i$$

which has zero curvature Ω. Such choices (for fixed 0-connection) give 1-connections varying by d(global 1-form) rather than by a general global 2-form.

Now consider some simple non-trivial gerbs $[\mathcal{G}] \neq 0$.

2.3.3 Spin^c-bundles and principal gerbs

Consider the oriented frame bundle of an oriented Riemannian manifold; this is a principal SO(n)-bundle. Whether or not we can lift this to the universal cover Spin(n) of SO(n) is a matter for the second Stiefel-Whitney class [35]

$$w_2 \quad \in \quad \mathrm{H}^2(X; \mathbb{Z}_2).$$

Consider also lifts to $\mathrm{Spin}^c(n) = \mathrm{Spin}(n) \times_{\pm 1} S^1$, which sits over the special orthogonal group

$$0 \to S^1 \to \mathrm{Spin}^c(n) \to \mathrm{SO}(n) \to 0.$$

Define $W_3 \in \mathrm{H}^3(X; \mathbb{Z})$ to be the image of w_2 in the long exact sequence of

$$0 \to \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \to \mathbb{Z}_2 \to 0.$$

Then the oriented frame bundle lifts to a Spin^c-bundle iff W_3 vanishes [30]. (This is always true for instance on a simply-connected 4-manifold, as featured in Seiberg-Witten theory.)

It is straighforward to put a gerb interpretation on this well-known fact. Take a cover for X such that we can lift over each U_i. Then put line-bundles Λ_i^j on the intersections with some trivialisations λ_i^j. We also have S^1-valued functions on the $U_{i,j}$ given by the failure of the local lifts to extend; and on triple intersections the coboundary of these functions is some

$$\eta \quad \in \quad \mathrm{Z}^2(X; \mathbb{Z}_2)$$

which represents the second Stiefel-Whitney class. We define the line-bundle section

$$\theta \quad := \quad \eta \cdot \delta\lambda$$

on $U_{i,j,k}$, thus fixing a smooth gerb.

Proposition *The equivalence class of this gerb is*

$$[\mathcal{G}] \quad \cong \quad W_3.$$

(So $\mathcal{G}$ is (smooth-) trivial iff there is a global lift to Spin^c*.)*

Proof—Viewing $\mathbb{Z}_2 = \{\pm 1\}$ as a subsheaf of $\underline{\mathbb{C}}^*$, we re-interpret the short exact sequence $\mathbb{Z} \to \mathbb{Z} \to \mathbb{Z}_2$ as one of (locally constant) sheaves with inclusions

$$\left\{ \begin{array}{ccccccccc} 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{2\pi i} & \pi i\mathbb{Z} & \xrightarrow{\exp} & \mathbb{Z}_2 & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{2\pi i} & \underline{\mathbb{C}} & \xrightarrow{\exp} & \underline{\mathbb{C}}^* & \longrightarrow & 0. \end{array} \right.$$

The cocycle η represents both the Stiefel-Whitney class $w_2 \in \mathrm{H}^2(X; \mathbb{Z}_2)$ (when viewed as a cocycle of $\{\pm 1\}$) and also the gerb class $[\mathcal{G}] \in \mathrm{H}^2(X; \underline{\mathbb{C}}^*)$ (as a cocycle of $\underline{\mathbb{C}}^*$). The images W_3 and the coboundary of $[\mathcal{G}]$ in the two long exact sequences are then both represented by the cocycle

$$\frac{1}{2\pi i}\delta \log \eta \quad \in \quad \mathrm{Z}^3(X;\mathbb{Z})$$

which suffices. □

Note that the Chern class $c(\mathcal{G})$ is in fact zero, since the image of $[\mathcal{G}]$ in $\mathrm{H}^3(X;\mathbb{Z})$ has order 2 and vanishes in de Rham cohomology.

Other cases of central extensions can be handled in the same way. One could argue that this example would be better handled by a "principal" gerb, in that we should be thinking of the Λ_i^j as principal $\mathbb{Z}_2$-bundles. Then our construction amounts to taking a principal $\mathbb{Z}_2$-gerb (with class $[\mathcal{G}] = w_2$) and viewing it as a principal $\underline{\mathbb{C}}^*$-gerb (with class $[\mathcal{G}] \cong W_3$). The second is trivial iff there is a lift to Spin^c, the first iff there is a lift to Spin.

We shall not pursue this since our other examples are quite reasonably described by "vector" gerbs. We merely remark that our initial definition of a gerb (2.1.1) should perhaps have described itself not simply as "abelian" but as a "rank-1 vector $\mathbb{C}^*$-gerb". There is an immediate generalisation of the definition to principal gerbs of any abelian group, and to their representations. How to handle the non-abelian case is altogether another matter.

2.3.4 Partitions of unity

As we have seen, partitions of unity can be a useful auxiliary tool in constructing connections. Suppose given a locally-trivialised gerb with 0-connection, and choose also a smooth partition of unity $\{\rho_i : i \in I\}$ subordinate to the given cover.

Then the 0-connection gives a cocycle of 2-forms F_i^j (the curvatures of the line-bundle connections), and we can define a 1-connection by

$$\beta_i \quad := \quad \sum_{k\neq i} \rho_k F_i^k$$

which obeys

$$\begin{aligned}\beta_i - \beta_j \quad &= \quad \rho_j F_i^j - \rho_i F_j^i + \sum_{k\neq i,j} \rho_k (F_i^k - F_j^k) \\ &= \quad (\rho_j + \rho_i + \sum_{k\neq i,j} \rho_k) F_i^j\end{aligned}$$

(since $\delta F = 0$)

$$= \quad F_i^j$$

as required. The curvature of this 1-connection is not zero.

(This example is just an expansion of the proof of (2.2.4).)

2.3.5 Holomorphic gerbs

On a complex manifold there is an uncomplicated translation from smooth gerbs to holomorphic. Write $\mathcal{O}$ for the sheaf of holomorphic functions, and $\mathcal{O}^*$ for the nowhere-zero functions. Most of the results of this chapter survive unaltered.

Definition (Holomorphic gerb) *A locally-trivialised holomorphic gerb is a smooth gerb (2.1.1) on a complex manifold X, whose bundles Λ_i^j and sections $\theta_{i,j,k}$ are holomorphic.*

Under the obvious notion of holomorphic equivalence (or 0-gauge transformation), the collection of equivalence classes over X is naturally identified with

$$\mathrm{H}^2(X;\mathcal{O}^*).$$

The class of such a gerb will be written as $[\mathcal{G}]$. Note that this is now quite distinct from the Chern class $c(\mathcal{G})$, which is defined by the holomorphic exponential sequence

$$0 \longrightarrow \mathbb{Z} \xrightarrow{2\pi i} \mathcal{O} \xrightarrow{\exp} \mathcal{O}^* \longrightarrow 0$$

as the image (under inclusion $\mathbb{Z} \hookrightarrow \mathbb{C}$) of

$$[\mathcal{G}] \in \mathrm{H}^2(\mathcal{O}^*) \to \mathrm{H}^3(\mathbb{Z}) \to \mathrm{H}^3(\mathbb{C}).$$

Accordingly, we expect analogues of the Atiyah class [1] to obstruct the existence of holomorphic gerb connections. This matter and others will be studied in chapter 5.

2.3.6 Points in 3-manifolds

Here we begin to consider the relation between gerbs and codimension 3. Pick a finite collection of disjoint points

$$p_i \in X^3 \quad : \quad i \in I$$

in some real 3-manifold X^3 (oriented as ever). Choose a cover of the form

$$\begin{cases} U_0 & := \; X \setminus \bigcup_{i \in I}\{p_i\} \\ U_i & := \; \text{open-ball neighbourhood of } p_i \end{cases}$$

such that all U_i, U_j are disjoint: any pairwise intersection looks like a punctured ball around one of the points, and there are no triple intersections.

So we define a gerb completely by installing line-bundles on the $U_{0,i}$. It is reasonable to choose these to be (isomorphic to) the standard degree-1 monopole line-bundle over the punctured ball. Without a means to pick particular bundles within this isomorphism class, we have at least fixed an equivalence class of gerbs (2.1.5).

Given particular choices of bundles, any connection on each of these suffices to fix a 0-connection on the gerb.

This obviously extends to arbitrary integral cycles

$$R \; := \; n_i p_i \; \in \; \mathrm{C}_0(X;\mathbb{Z})$$

in that, about a weighted point, we would wish to put the line-bundle Λ_i^0 of appropriate degree.

Proposition *The Chern class* $c(\mathcal{G}) \in \mathrm{H}^3(X;\mathbb{C})$ *is Poincaré-dual to*

$$[R] \quad \in \quad \mathrm{H}_0(X;\mathbb{C}).$$

Proof—Pick a 0-connection. This has curvature 2-forms F on each punctured ball, of degree equal to the multiplicity n_i. Pick a partition of unity: this gives a curvature 3-form supported only on the punctured balls

$$\Omega \quad = \quad \mathrm{d}\rho_i \wedge F_i^0$$

on each U_i for $i \neq 0$. But now

$$\begin{aligned} \int_X \omega &= \sum_{i\neq 0} \Big(\int_{S^2} \frac{i}{2\pi} F_i^0 \cdot \int_{r=0}^{r_i} \mathrm{d}\rho_i \Big) \\ &= \sum_{i\neq 0} n_i \end{aligned}$$

as desired. (The integrals are hopefully self-explanatory.) □

In chapter 6 we shall extend this example to a codimension-3 submanifold in X^n.

Chapter 3

Objects and errors

There is a clear analogue of the concept for line-bundles of nowhere-zero section, which we discuss in this chapter. Rather than compound our repetition of bundle terminology, we shall speak not of *sections* of gerbs, but of local (or global) *objects*. This is also in recognition of the fact that these things will demonstrate the link between our gerbs and the original category-theoretic *gerbes* (3.4).

3.1 Objects and trivial gerbs

Definition 3.1.1 (Object) *Given a locally-trivialised gerb $\mathcal{G}(I, \Lambda, \theta)$, an object is specified by*

- *line-bundles L_i defined over each U_i;*
- *bundle isomorphisms on each $U_{i,j}$*

$$m_i^j : L_i \xrightarrow{\cong} \Lambda_i^j \otimes L_j;$$

such that the composition on each 3-fold intersection

$$L_i \longrightarrow \Lambda_i^j \Lambda_j^k \Lambda_k^i \otimes L_i$$

is exactly

$$m_i^j \circ m_j^k \circ m_k^i \equiv \theta_{i,j,k} \otimes \textit{identity}.$$

Proposition 3.1.2 *An object exists, if and only if the gerb is trivial*

$$[\mathcal{G}] = 0.$$

Proof—Refine sufficiently; this does not affect the class of $\mathcal{G}$. Take bases for all bundles. View θ as a cocycle $t \in \mathrm{Z}^2(X; \underline{\mathbb{C}}^*)$.

Suppose there exists an object. With respect to the selected bases, we can view its isomorphisms m_i^j as a cochain

$$m \in \mathrm{C}^1(X; \underline{\mathbb{C}}^*) \;:\; \delta m = t$$

so that $0 = [t] = [\theta] = [\mathcal{G}]$.

In reverse, if $\mathcal{G}$ is trivial it has a globally-trivial refinement (2.1.7) with bases λ_i^j. Impose trivial bundles L_i with bases l_i and define isomorphisms

$$m_i^j : l_i \mapsto \lambda_i^j \otimes l_j$$

to give an object. □

Recall that (2.1.9) described triviality in a "coordinate-dependent" way, in terms of a coboundary of bases. This result in terms of a set of (non-trivial) line-bundles is more geometric. We should think of the difference between the two just like descriptions of a line-bundle either with transitions which are a coboundary or with a more abstract nowhere-zero section.

An object on an already globally-trivialised gerb gives a global line-bundle: the trivialisation singles out choices for bases λ, thus permitting us to use

$$m : l_i \mapsto \lambda \otimes l_j$$

as glue for a global bundle

$$m' : l_i \mapsto l_j \quad : \quad \delta m' = 0.$$

Just as an explicitly-trivial line-bundle turns (nowhere-zero) sections into global (nowhere-zero) functions (and any such function is possible), the global bundle l_i is very dependent on the trivialisation: varying the trivialisation by $\mathrm{Z}^1(\underline{\mathbb{C}}^*)$ turns a given object into any global line-bundle.

Example: a local object

In general, objects always exist locally, but not necessarily globally. As an example (for a general gerb) of a local object around a point in some set U_a of our cover, we can use the local trivialisation over U_a as follows. Restrict the gerb to U_a, so the new cover is now

$$V_i := \begin{cases} U_a & \text{if } i = a \\ U_i \cap U_a & \text{otherwise} \end{cases}$$

with the same bundles Λ_b^c restricted to their new domains.

Then a local object is given by an arbitrary bundle $\mathcal{L}$ over U_a, inducing

$$L_i := \begin{cases} \mathcal{L} & \text{if } i = a \\ \mathcal{L} \otimes \Lambda_i^a & \text{otherwise} \end{cases}$$

with the canonical isomorphisms for m_i^j.

Proposition 3.1.3 *In general, the difference between two objects is a global line-bundle.*

("Global" here means global over the support of the objects.)
Proof—Go to a common trivialisation. Then we have two objects L_i and K_i on U_i whose maps induce

$$m^{-1} \otimes n : L_i^{-1} \otimes K_i \to L_j^{-1} \otimes K_j$$

with composition equal to the identity on triple intersections. So we use $m^{-1}\otimes n$ as glue to create a global line-bundle "$L^{-1}\otimes K$". (Its two factors do not exist globally in general.) □

Compare (2.1.10). This result implies that, locally, any object has the form of the above example for some general $\mathcal{L}$.

Definition 3.1.4 (Equivalence of objects) *Given an gerb with object, we extend the idea of equivalence (2.1.5) by asking for isomorphisms on U_i of the bundles L_i and K_i, that commute with the maps m and n.*

A *0-transform* on an object is an equivalence with itself.

Proposition 3.1.5 *Two objects on the same gerb are equivalent iff they differ by a trivial bundle.*

Proof—An equivalence

$$\chi_i : L_i \to K_i$$

defines local nowhere-zero sections of $L_i^{-1}\otimes K_i$. Commuting with the maps $(m^{-1}\otimes n)_i^j$ means that these join up to a global nowhere-zero section of "$L^{-1}\otimes K$", trivialising it.

Conversely, if this difference has a global basis, any such choice induces local maps from L_i to K_i that commute with the object maps m and n. □

3.2 0-Connections on objects

Definition 3.2.1 (Object 0-connection) *Given a trivial gerb with 0-connection, and an object, we can define a 0-connection* on the object *as bundle connections for each L_i such that the isomorphisms m_i^j preserve the structure*

$$\nabla|_{L_i} \overset{m}{\cong} \nabla|_{\Lambda_i^j\otimes L_j}$$

(so the 2-curvatures obey

$$\delta F_i = F_i^j).$$

Proposition 3.2.2 *Object 0-connections exist.*

Proof—Since $[\mathcal{G}]=0$, we can refine and take a global trivialisation (2.1.7)

$$\delta\lambda = \theta$$

and bases l_i for the line-bundles of the object.

Now the isomorphisms m can be viewed in terms of a cocycle

$$l_i \mapsto m_i^j \lambda_i^j \otimes l_j$$
$$m_i^j \in \mathrm{Z}^1(X;\underline{\mathbb{C}}^*).$$

Say the 0-connection is given by 1-forms α_i^j. We want a 0-connection on the object, represented by 1-forms A_i. The constraint

$$\nabla l = \nabla(m\cdot\lambda\otimes l)$$

says that

$$A_i - A_j = \alpha_i^j + \mathrm{d}\log m_i^j.$$

But the right-hand side is a cocycle

$$\delta(\alpha + \mathrm{d}\log m) = -\mathrm{d}\log t + \mathrm{d}\log t = 0$$

and the obstruction to the existence of A_i is just

$$[\alpha + \mathrm{d}\log m] \in \mathrm{H}^1(X; \mathcal{A}^1) = 0.$$

□

Proposition 3.2.3 *Given a gerb with 0-connection and object, two object 0-connections differ by a global 1-form.*

Proof—Given α and m, any object connection A_i obeys

$$\delta A = \alpha + \mathrm{d}\log m$$

as in the proof of (3.2.2). So any alternative object connection A'_i differs by local 1-forms with vanishing coboundary. □

(An object with "connection" means one with 0-connection.)

Proposition 3.2.4 *Given a gerb with 0-connection, any two objects (with object connections) differ by a global line-bundle (with line-bundle connection).*

Proof—Without the connections, this is just (3.1.3). On each U_i the object connections induce a connection on the restricted global bundle $L_i^{-1} \otimes {L'}_i$. The transitions $m^{-1} \otimes m'$ preserve this connection. □

A converse clearly holds: given an object (with connection) and any global line-bundle (with any connection), we can create a new object (with connection).

A 0-gauge transformation of an object alters its 0-connection in the usual way. Then we can relate equivalence to the above result.

Proposition 3.2.5 *Given a gerb with 0-connection, two objects with connection are equivalent iff their difference is a trivial line-bundle with flat connection.*

Proof—As in (3.1.5) and (3.2.4). □

Example: a local object 0-connection

Suppose given a gerb with 0-connection. Expanding on the example preceding (3.1.3) of a local object defined on a set U_a of the cover, choose any bundle $\mathcal{L}$ with connection ∇ over this space.

Combined with the connections on the Λ_i^j given by the gerb 0-connection, this induces a connection on each L_i of the object (ie. $\mathcal{L}$ or $\mathcal{L} \otimes \Lambda_i^a$). These are certainly compatible with the canonical maps m, and so we have a 0-connection for the object.

Proposition 3.2.6 *Every object with connection has this local form.*

Proof—See (3.2.4). □

3.3 Errors and 1-connections

Now that we have defined objects, we can give another example of a 1-connection, that despite its simplicity will feature widely in later chapters.

Definition 3.3.1 (Objective 1-connection) *Given a 0-connection, and an object with 0-connection, the 1-connection defined by*

$$\beta_i := F_i$$

is said to be objective.

That is, we just take the curvature of each L_i. This gives zero gerb curvature

$$\Omega = \mathrm{d}\beta_i = 0.$$

(Inevitably the *class* of Ω vanishes, due to triviality.) Given triviality, the most general form of 1-connection is an objective 1-connection plus any *global* 2-form.

There is an alternative description of the same construction.

Definition 3.3.2 (Error 2-form) *Given a trivial gerb with a general 1-connection, and an object with connection, the error form is the global 2-form*

$$\epsilon := \beta_i - F_i \quad (\forall U_i).$$

Note that the error obeys

$$\Omega = \mathrm{d}\epsilon$$

and that the gerb 1-connection is objective for the given object connection iff $\epsilon = 0$.

Slightly more generally (3.2.4), given a gerb with 1-connection and an object with connection, the 1-connection is in fact objective (for some possibly *different* object connection) iff $\mathrm{d}\epsilon = 0$ (ie. $\Omega = 0$) and ϵ is integral

$$\left[\frac{i}{2\pi}\epsilon\right] \in \mathrm{H}^2_{\mathrm{dR}}(X; \mathbb{Z}).$$

We might try to use the error to give a different emphasis to the definition (2.2.3) of 1-connection.

Definition 3.3.3 (1-connection (alternative)) *Fix a gerb with 0-connection. Then define a 1-connection to be a rule that, to every locally-defined object with 0-connection, assigns a smooth local 2-form ϵ. The rule must behave well under transformations, as below.*

The transformation properties required are simple. First, any other object with 0-connection differs by a bundle with connection, of curvature 2-form F, say. The 1-connection is to follow such a shift by

$$\epsilon \mapsto \epsilon + F.$$

(The same holds if the object connection is changed but not the underlying object—this merely forces F to be exact.) Second, a 1-gauge transform γ is to alter the 1-connection by

$$\epsilon \mapsto \epsilon + \mathrm{d}\gamma.$$

Proposition 3.3.4 *This definition is equivalent to (2.2.3).*

Proof—Given β, our ϵ are just the error 2-forms for any local object with connection, and they transform appropriately.

Conversely, given the new version of ϵ, recall the local object with connection described after (3.2.5). Given a locally-trivialised gerb with 0-connection, and a rule providing the alternative 1-connection, take on each set U_i the object with connection generated by $(\mathcal{L}, \nabla)$ being *trivial and flat.* This gives 2-forms ϵ_i, and we define β_i to equal these. Compare (3.3.2), and note that now $F_i = 0$: tracing through the example shows that these β do form a 1-connection under the original definition. □

3.4 Gerbs and gerbes

To write out Giraud's full definition of a gerbe [18] would be long and painful, and indeed the point of this thesis is that a simple-minded differential geometer can make do without the full algebraic-geometric technology. Following Brylinski [5], we restrict to abelian gerbes over a smooth manifold X. Even this is quite indigestible enough, but we summarise the exercise.

Sheaves of categories

A gerbe with band $\underline{\mathbb{C}}^*$ over X is a sheaf of categories satisfying certain properties.

To have a sheaf of categories means that for every open set in X, or indeed for every smooth $f : Y \to X$ which is locally a diffeomorphism, we have a category, possibly empty of objects. There are various niceness requirements under composition of local diffeomorphisms and the like.

Further, we insist that every morphism be invertible (the category is a groupoid); that the sheaf of automorphisms of any object is locally isomorphic to $\underline{\mathbb{C}}^*$; that, given two objects of any category, they are at least locally isomorphic; and that there exists some *surjective* local diffeomorphism whose category is non-empty.

There is a natural definition of equivalence between gerbes over X, and happily the equivalence classes correspond naturally with

$$\mathrm{H}^2(X; \underline{\mathbb{C}}^*).$$

Details are omitted due to the number of unilluminating commutative diagrams involved; but we now have a natural bijection between classes of gerbs and of gerbes.

Theorem 3.4.1 *Gerbs are gerbes.*

Proof—(Sketch)—Take a local trivialisation of $\mathcal{G}$ over X. Given an inclusion or indeed a local diffeomorphism

$$f : Y \to X$$

this defines by pull-back a local trivialisation of a gerb $f^*\mathcal{G}$ over Y. Then to Y we assign the collection of “objects” over Y in the sense of (3.1.1). These are then objects of a category, with morphisms given (when possible) by equivalence of objects (3.1.4).

To show the existence of a non-empty category for a surjective map: we clearly cannot take just X itself, unless $\mathcal{G}$ is trivial. In general we can take the disjoint union

$$Y := \coprod_{a \in A} U_a$$

with the obvious map to X. The pullback gerb on Y does contain objects, since on each U_a we can put the object described before (3.1.3) with, say, all $\mathcal{L}$ being restrictions of $X \times \mathbb{C}$. □

Apart from involving merely bundles rather than categories, a further advantage of our approach is that we do not need this extra layer of objects and local diffeomorphisms. This cuts out a level of bureaucracy, and lets us ignore objects until they are explicitly useful.

Connective structure and curvings

Brylinski [5], following Deligne, has gone to some trouble to define what he calls a "connective structure" on a gerbe. We make no attempt to reproduce this, except to claim without proof that a 0-connection gives a connective structure via the family of possible 0-connections on objects. The Čech version advocated here seems rather more transparent.

We cannot claim much superiority in our definition of 1-connection over that of a "curving". That they agree is uncomplicated to prove using our alternative definition (3.3.3). Our equivalence classes of gerbs with 1-connection (2.2.7) canonically correspond with the appropriate classes of gerbes.

3.5 Gerbs and bundles

Perhaps it would help to clarify the relationships between our various constructions if we make explicit some parallels with line-bundles.

Nowhere-zero function	Line-bundle
Line-bundle	Gerb (with 0-connection?)
Connection	1-connection
Nowhere-zero section	Object (with connection?)
Connection 1-form	Error 2-form
Curvature 2-form	Curvature 3-form
Chern class	Chern class

This work offers no definition merely because of such a supposed analogy. The table should be viewed as hindsight rather than as justification. Parts that seem obscure as yet will hopefully make more sense later—eg. compare the "connection 1-form" entry with the definition of gerb holonomy in (7.1). That the analogue of a bundle should perhaps be seen as a gerb *with 0-connection* is implied for instance by (2.3.1)—in which it costs no more to fix a 0-connection than it does to give the underlying gerb—but perhaps not by the 3-gerb construction (4.5). A comparison of a basis of a bundle with an object with connection, rather than with just an object, is again supported by the holonomy idea of (7.1) but not by (4.5). (All three are supported by the alternative 1-connection definition (3.3.3).)

The moral is clearly that such fancies as feature in this table should take a back seat to concrete geometry. Barring some more theory in chapter 5, which considers the holomorphic case $\mathcal{O}^*$ in the light of the smooth $\underline{\mathbb{C}}^*$, the pursuit of such examples takes up the rest of this thesis.

Chapter 4

A twistor correspondence on projective space

In this chapter we take advantage of the admirable capacity of twistor theory for converting between holomorphic and differential constructs in a mathematically meaningful way [3, 49]. We consider a very simple geometry: complex projective space $\mathbb{P}^{n+1}$ and its dual (actually slightly modified). Ignoring all the intricacies of chapter 2, we start on one side merely with an element of $\mathrm{H}^n(X, \mathcal{O}^*)$ and ask what comes out on the other side. If we believe that $\mathrm{H}^n(\mathcal{O}^*)$ is to do with holomorphic "n-gerbs"—and given faith in the track record of twistor theory—we must expect whatever comes out to represent an "n-gerb with connection" even if we do not know what such a thing might be.

So the attitude of this chapter is to consider the geometry of $\mathbb{P}^{n+1}$ in its own right, and only at the end to compare what we find with our previous definitions. The general case is described in the first two sections. Then in (4.3) we interpret $n = 1$ as a transform between line-bundles; and we find in (4.4) that the $n = 2$ case is indeed talking of gerbs. In both of these examples, we show some explicit calculations that may help illuminate the abstract cohomological work of (4.1). Finally, we discuss what the general case tells us of higher-order constructions (4.5).

We should point out that in the line-bundle case $n = 1$, the result (4.3.1) has long been known in the twistor-theory industry, although it appears not to have been published. We discuss it in order to illuminate the gerb case $n = 2$.

All structures in this chapter are holomorphic, until declared otherwise in section (4.5). For the definitions of holomorphic gerb and connection, we rely on (2.3.5) and a little common sense. (For the details, see chapter 5—but the point of this chapter is to work in a situation in which the gerbs themselves are as simple as possible, and chapter 5 is hopefully superfluous.)

4.1 The geometry of punctured projective space

Throughout this chapter we are working within $\mathbb{P}^{n+1}$, its dual $\mathbb{P}^*$ and the correspondence space between them. We view this as a twistor transform, taking for instance points on one side to hyperplanes on the other, in the traditional fashion [21, 39].

We need something slightly more complicated than this, to get sufficient structure (eg. there are very few interesting line bundles on $\mathbb{P}^n$). The next simplest geometry is obtained by removing a point

$$\begin{cases} \mathcal{P} & := & \mathbb{P}^{n+1} - \text{one point} \\ \mathcal{Q} & := & \{\text{all hyperplanes } \mathbb{P}^n \subset \mathcal{P}\} = \text{affine } \mathbb{C}^{n+1} \subset \mathbb{P}^{n+1*} \end{cases}$$

with the correspondence space

$$\mathcal{F} := \{(p, \pi) \ : \ p \in \mathcal{P}, \ \pi \in \mathcal{Q}, \ p \in \pi\} \subset \mathbb{P}^{n+1} \times \mathbb{P}^{n+1*}$$

(if the notational abuse is pardoned) projecting to each via maps f, g respectively.

Over $\mathcal{P}$ we see that $\mathcal{F}$ is an affine $\mathbb{C}^n$-bundle, whilst its projection ($\mathcal{F} \xrightarrow{g} \mathcal{Q}$) is very simple: $\mathcal{F}$ is a trivial $\mathbb{P}^n$-bundle over $\mathcal{Q}$. In particular,

$$\mathcal{F} = \mathbb{P}(T\mathcal{Q})$$

for $T\mathcal{Q}$ the holomorphic tangent bundle of $\mathcal{Q}$.

We can best describe $\mathcal{P}$ by choosing a hyperplane $\pi_0 \subset \mathcal{P}$ (which corresponds to a point in $\mathcal{Q}$). Over $\pi_0 = \mathbb{P}^n$, the space $\mathcal{P}$ is naturally an $\mathcal{O}(1)$ line-bundle; and $\mathcal{Q}$ is a vector space

$$\mathcal{Q} = \mathrm{H}^0(\pi_0; \mathcal{P}).$$

Three properties of $\mathcal{P}$

First we note that $\mathcal{P}$ does carry enough cohomology to be of value.

Proposition 4.1.1 $\mathrm{H}^n(\mathcal{P}; \mathcal{O})$ *is not zero. In particular, given two distinct hyperplanes* π_0 *and* η *in* $\mathcal{P}$*, and viewing* η *as a section of* $\mathcal{P}$ *over the base-space* π_0*,*

$$\mathrm{H}^n(\mathcal{P}; \mathcal{O}) \cong \bigoplus_{i \geq n+1} \eta^i \ \mathrm{H}^n(\pi_0; \mathcal{O}(-i)).$$

Proof—We use Čech cohomology. We take the standard open cover $\{U_r\}_{r=0,1,\ldots,n}$ for $\pi_0 = \mathbb{P}^n$. Choose it such that U_0 is the complement of $\pi_0 \cap \eta$. This lifts to a cover for $\mathcal{P}$, which is acyclic for the sheaf $\mathcal{O}$ since in general any covering by $\mathbb{C}^a \times (\mathbb{C}^*)^b$ is acyclic [21]. (Indeed, this cover is acyclic for all other sheaves we shall require.) Any class in $\mathrm{H}^n(\mathcal{P}; \mathcal{O})$ thus has a representative (on the $(n{+}1)$-fold intersection) of the form

$$\sum_{\substack{i \geq 0 \\ b_r \in \mathbb{Z}}} \frac{t_{i,b_1,\ldots,b_n} \cdot e^i}{\zeta_0^{b_1} \cdots \zeta_n^{b_n}}$$

for constants t and affine coordinates ζ_r on $U_0 = \mathbb{C}^n \subset \mathbb{P}^n$; and for e the coordinate in the fibre direction with respect to the basis vector η.

A little cordinate manipulation (eliminating terms which are coboundaries of $\mathrm{C}^{n-1}(\mathcal{P}; \mathcal{O})$) soon shows that in fact there is a *unique* representative of the form

$$\sum_{\substack{0 < b_r \\ \sum b_r < i}} \frac{t \cdot e^i}{\zeta_0^{b_1} \cdots \zeta_n^{b_n}}$$

(and any such function certainly gives a class). But if we now fix i, representatives

$$\sum_{\substack{0 < b_r \\ \sum b_r < i}} \frac{t_{i,b_1,\ldots,b_n}}{\zeta_0^{b_1} \cdots \zeta_n^{b_n}}$$

correspond exactly to classes in $\mathrm{H}^n(\pi_0; \mathcal{O}(-i))$. □

Secondly, we remark on the relation between $\mathcal{O}$ and $\mathcal{O}^*$ over $\mathcal{P}$.

Proposition 4.1.2 *$\mathrm{H}^n(\mathcal{P}; \mathcal{O})$ is naturally isomorphic to the kernel of*

$$\mathrm{H}^n(\mathcal{P}; \mathcal{O}^*) \to \mathrm{H}^{n+1}(\mathcal{P}; \mathbb{Z})$$

in the exponential long exact sequence. And $\mathrm{H}^{n+1}(\mathcal{P}; \mathbb{Z})$ is torsion-free, and vanishes if n is even.

Proof—The point is that an adequate cover for any $\pi_0 \subset \mathcal{P}$ lifts to one for $\mathcal{P}$; so for instance we can use the form of representatives of $\mathrm{H}^n(\mathcal{P}; \mathcal{O})$ from the proof of (4.1.1).

Alternatively, there is an exact correspondence between cocycles in the two sets

$$\begin{cases} \mathrm{Z}^n(\pi_0; \mathbb{Z}), \\ \mathrm{Z}^n(\mathcal{P}; \mathbb{Z}). \end{cases}$$

But in the exponential sequence, the map

$$\mathrm{H}^{n-1}(\pi_0; \mathcal{O}^*) \to \mathrm{H}^n(\pi_0; \mathbb{Z})$$

surjects (since $\mathrm{H}^n(\pi_0; \mathcal{O}) = 0$). So, by lifting elements of $\mathrm{Z}^{n-1}(\pi_0; \mathcal{O}^*)$ to $\mathcal{P}$, the map

$$\mathrm{H}^{n-1}(\mathcal{P}; \mathcal{O}^*) \to \mathrm{H}^n(\mathcal{P}; \mathbb{Z})$$

also surjects, which suffices.

Given that $\mathcal{P}$ is diffeomorphic to a vector bundle over $\mathbb{P}^n$, whose $\mathbb{Z}$-cohomology is well-known [21], the second claim is clear. □

Thirdly, we consider the infinitesimal neighbourhoods of a general hyperplane $\pi \subset \mathcal{P}$. Given that $\mathcal{I}$ is the ideal of π in $\mathcal{O}_\mathcal{P}$—the sheaf of functions vanishing on π—it is a general fact [49] that

$$\mathcal{O}_\pi \equiv \mathcal{O}_\mathcal{P}/\mathcal{I}$$

restricted to π. More precisely, but still restricted to π,

$$\mathcal{O}_\mathcal{P} = \mathcal{O}_\pi \oplus \mathcal{I}/\mathcal{I}^2 \oplus \mathcal{I}^2/\mathcal{I}^3 \oplus \ldots$$

and that

$$\mathcal{I}^j/\mathcal{I}^{j+1} = S^q N^*$$

where $S^q N^* \cong \mathcal{O}(-q)$ is the q-th symmetric power of the conormal bundle. The q-th formal neighbourhood of π (for $q \geq 1$) is the scheme $(\pi; \mathcal{O}^q)$, where

$$\mathcal{O}^q := \mathcal{O}_\mathcal{P}/\mathcal{I}^{q+1}.$$

We view this as an extension of $(\pi; \mathcal{O})$. From the above, there is an exact sequence

$$0 \to S^q N^* \to \mathcal{O}^q \to \mathcal{O}^{q-1} \to 0$$

where $\mathcal{O}^0 := \mathcal{O}$.

The normal bundle of π is of degree one (by adjunction, or by intersection properties in $\mathbb{P}^{n+1}$). We are interested in extending from $\mathrm{H}^{n-1}(\pi; \mathcal{O})$ to the neighbourhoods of π, ie. in

$$\mathrm{H}^{n-1}(\mathbb{P}^n; \mathcal{O}(-q)) \to \mathrm{H}^{n-1}(\pi; \mathcal{O}^q) \to \mathrm{H}^{n-1}(\pi; \mathcal{O}^{q-1}) \to \mathrm{H}^n(\mathbb{P}^n; \mathcal{O}(-q)).$$

But in fact

$$\mathrm{h}^n(\mathbb{P}^n; \mathcal{O}(-q)) = \begin{cases} 0 & q \leq n \\ 1 & q = n+1 \end{cases}$$

by Serre duality; and furthermore

$$\mathrm{h}^{n-1}(\mathbb{P}^n; \mathcal{O}(-q)) = 0 \quad (\forall q).$$

These groups are given by the following, of which we shall later have more need.

Proposition 4.1.3 *Writing Ω^p for the sheaf of holomorphic p-forms,*

$$\mathrm{h}^q(\mathbb{P}^n; \Omega^p(k)) = \begin{cases} \binom{k+n-p}{k}\binom{k-1}{p} & q=0,\ k>p=0,\ldots,n; \\ \binom{-k-1}{n-p}\binom{p-k}{-k} & q=n,\ k<p-n=-n,\ldots,0; \\ 1 & k=0,\ p=q=0,\ldots,n; \\ 0 & \textit{otherwise.} \end{cases}$$

Proof—Bott [4]; or as summarised in [39]. □

The interest in this vanishing of cohomology of the normal bundle is that it shows a unique extension exists from $\mathrm{H}^{n-1}(\pi; \mathcal{O})$ to $\mathrm{H}^{n-1}(\pi; \mathcal{O}^n)$, with an obstruction to further extension lying in $\mathrm{H}^n(\pi; S^n N^*)$. In itself this is vacuous unless $n = 1$, since $\mathrm{H}^{n-1}(\pi; \mathcal{O}) = 0$ (from the above); we shall take a more productive line when we discuss the $n = 1$ case (4.3).

A twistor correspondence

Clearly the non-trivial content of our geometry is the bundle $(\mathcal{F} \xrightarrow{f} \mathcal{P})$, and understanding this is the task in this section. The key observation of this chapter is—

Theorem 4.1.4 *There is a natural equivalence*

$$\mathrm{H}^n(\mathcal{P}, \mathcal{O}) \cong \mathrm{H}^0(\mathcal{Q}; \Omega^n)/\mathrm{H}^0(\mathcal{Q}; \mathrm{d}\Omega^{n-1}).$$

The device by which we approach this is the tautological rank-n vector bundle

$$\mathcal{W} \to \mathcal{F}$$

of tangent vectors in $\mathcal{F}$ along the directions of the fibres of

$$\mathcal{F} \xrightarrow{f} \mathcal{P}.$$

This comes with a well-defined operator $\mathrm{d}^{\mathcal{W}}$ differentiating along the same directions, which for instance takes local functions on $\mathcal{F}$ to local sections of $\mathcal{W}^*$.

Proof—Start from the left-hand side: choose a Čech representative for a class

$$\theta \in [\theta] \in \mathrm{H}^n(\mathcal{P}, \mathcal{O})$$

taken with respect to a standard open cover for $\mathcal{P}$ as in (4.1.1). This representative lifts to

$$f^*\theta \in f^*[\theta] \in \mathrm{H}^n(\mathcal{F}, \mathcal{O})$$

(for which the lift of our cover is still acyclic).

But this group is zero by lemma (4.2.1); and also, by definition of $\mathcal{W}$ we see

$$\mathrm{d}^{\mathcal{W}} f^*\theta = 0.$$

So there exists some

$$\mu_{n-1} \in \mathrm{C}^{n-1}(\mathcal{F}; \mathcal{O}) \quad : \quad \delta\mu_{n-1} = f^*\theta$$

which obeys

$$\mathrm{d}^{\mathcal{W}} \mu_{n-1} \in \mathrm{Z}^{n-1}(\mathcal{F}; \mathcal{W}^*).$$

As we shall see in lemma (4.2.3),

$$\mathrm{H}^{n-r}(\mathcal{F}; \Lambda^r \mathcal{W}^*) = 0 \quad r = 0, \ldots, n-1$$

so that we can progress from

$$\mathrm{d}^{\mathcal{W}} \mu_{n-1} = \delta\mu_{n-2} \quad : \quad \mathrm{d}^{\mathcal{W}} \mu_{n-2} \in \mathrm{Z}^{n-2}(\mathcal{F}; \Lambda^2 \mathcal{W}^*)$$

as far as

$$\mathrm{d}^{\mathcal{W}} \mu_2 = \delta\mu_1 \quad : \quad \mathrm{d}^{\mathcal{W}} \mu_1 \in \mathrm{Z}^1(\mathcal{F}; \Lambda^{n-1} \mathcal{W}^*)$$

to give finally a choice of

$$\mu_0 \quad \Rightarrow \quad \mathrm{d}^{\mathcal{W}} \mu_0 \in \mathrm{H}^0(\mathcal{F}; \det \mathcal{W}^*)$$

which according to lemma (4.2.4) is canonically equal to

$$\mathrm{H}^0(\mathcal{Q}; \Omega^n)$$

and we have reached a representative for the right-hand side of the theorem.

It is clear—since $\mathrm{H}^{n-r-1}(\mathcal{F}; \Lambda^r \mathcal{W}^*) = 0$ for $r = 0, \ldots, n-2$ by lemma (4.2.3)—that different choices of representatives cannot affect the outcome except at the term μ_0. So the n-form $g_* \mathrm{d}^{\mathcal{W}} \mu_0$ is well-defined up to an element of

$$\mathrm{H}^0(\mathcal{F}; \mathrm{d}^{\mathcal{W}} \Lambda^{n-1} \mathcal{W}^*)$$

which, after an unproblematic differentiation, is canonically identified in lemma (4.2.4) with

$$\mathrm{H}^0(\mathcal{Q}; \mathrm{d}\Omega^{n-1})$$

and we have proved half of the theorem.

Now certainly there exists a Poincaré lemma for $\mathrm{d}^{\mathcal{W}}$

$$\mathrm{d}^{\mathcal{W}} \zeta = 0 \;\Rightarrow\; \zeta = \mathrm{d}^{\mathcal{W}} \eta$$

for some local η, and so the other direction of proof proceeds painlessly. The maps in the two directions are inverses (on equivalence classes). We leave this unwritten. □

4.2 Lemmas for theorem (4.1.4)

Recall that $(\mathcal{F} \xrightarrow{g} \mathcal{Q})$ is a trivial $\mathbb{P}^n$-bundle over $\mathcal{Q} \cong \mathbb{C}^{n+1}$. The Leray spectral sequence [32, 19] can easily be applied in this context: given a sheaf $\mathcal{S}$ over $\mathcal{F}$, its direct image sheaves $\mathrm{R}^q\mathcal{S}$ over $\mathcal{Q}$ (where $q \geq 0$) are such that their stalks at a point $y \in \mathcal{Q}$ are

$$H^q(g^{-1}(y); \mathcal{S}).$$

Then the Leray spectral sequence starts from

$$E_2^{p,q} := \mathrm{H}^p(\mathcal{Q}; \mathrm{R}^q\mathcal{S})$$

and abuts to

$$\mathrm{H}^*(\mathcal{F}; \mathcal{S}).$$

Since $g^{-1}(y) = \mathbb{P}^n$, which supports very little cohomology, the spectral sequences we shall need (taking $\mathcal{S} := \mathcal{O}$ or $\Lambda^r\mathcal{W}^*$) will degenerate rapidly.

Lemma 4.2.1 $\mathrm{H}^p(\mathcal{F}; \mathcal{O}) = 0$ *for all* $p > 0$.

Proof—Since

$$\mathrm{h}^q(\mathbb{P}^n; \mathcal{O}) = \begin{cases} 1 & q = 0 \\ 0 & \text{otherwise,} \end{cases}$$

we know that

$$\mathrm{R}^q(\mathcal{O}_\mathcal{F}) = \begin{cases} \mathcal{O}_\mathcal{Q} & q = 0 \\ 0 & q > 0 \end{cases}$$

so that in the Leray spectral sequence

$$E_2^{p,q} = 0$$

except for $p = q = 0$. □

The proof of theorem (4.1.4) requires the vanishing of a collection of cohomology groups of the form $\mathrm{H}^r(\mathcal{F}; \Lambda^s\mathcal{W}^*)$. To check this, we first identify $\mathcal{W}$ up to isomorphism over a single $\mathbb{P}^n \subset \mathcal{F}$.

Lemma 4.2.2 *Restricting to any choice of fibre* $g^{-1}(y) = \pi = \mathbb{P}^n$ *of* $(\mathcal{F} \xrightarrow{g} \mathcal{Q})$,

$$\mathcal{W} \cong T^* \otimes \mathcal{O}(1)$$

where T^* *is the holomorphic cotangent bundle over* π.

Proof—Whilst $\mathcal{W}$ is a rank-n sub-bundle of the tangent bundle of $\mathcal{F}$, it is also a sub-bundle of the rank-$(n{+}1)$ normal bundle N of $\pi \subset \mathcal{F}$. (N is trivial, though not canonically so; and is not naturally a sub-bundle of $T\mathcal{F}$.) So there is a sequence

$$0 \to \mathcal{W} \to N \to \mathcal{O}(a) \to 0$$

for some integer a.

In fact, $a = 1$. This is because $f_*\pi \subset \mathcal{P}$ is a $\mathbb{P}^n$ whose normal bundle in $\mathcal{P}$ is canonically the $\mathcal{O}(1)$ bundle with total space $\mathcal{P}$ itself—and whilst the lift of this line-bundle cannot canonically be identified with a sub-bundle of $(N \to \pi)$, it is certainly isomorphic to the quotient of $\mathcal{W} \hookrightarrow N$.

We have previously remarked that $\mathcal{F} = \mathbb{P}(T\mathcal{Q})$. Restricting to $y \in \mathcal{Q}$ says that $\pi = \mathbb{P}(V)$, where $V := T_y\mathcal{Q}$. Note that we can canonically identify N with the trivial pull-back bundle g^*V over $\mathbb{P}(V)$. But this bundle is part of the Euler sequence

$$0 \to T^*(1) \to g^*V \to \mathcal{O}(1) \to 0$$

over $\mathbb{P}(V)$ [21].

Comparing our two exact sequences, it suffices to show that up to isomorphism there is only one possible kernel for an exact sequence of bundles

$$\underbrace{\mathcal{O} \oplus \cdots \oplus \mathcal{O}}_{n+1} \to \mathcal{O}(1) \to 0$$

over $\mathbb{P}^n$. A general map

$$\bigoplus_0^n \mathcal{O} \to \mathcal{O}(1)$$

is an element of

$$\bigoplus_0^n \mathrm{H}^0(\mathbb{P}^n; \mathcal{O}(1)).$$

We need a *surjective* map, which amounts to choosing $n{+}1$ linearly-independent elements of the vector space $\mathrm{H}^0(\mathbb{P}^n; \mathcal{O}(1))$. Since this is the same as choosing an element of

$$GL(n+1, \mathbb{C}) = \mathrm{Aut}(\mathcal{O} \oplus \cdots \oplus \mathcal{O})$$

we find that an automorphism of the trivial bundle exists that takes the kernel sub-bundle of any choice of map to the kernel of any other. So the existence of two exact sequences

$$\begin{cases} 0 \to \mathcal{W} \to N \to \mathcal{O}(1) \to 0 \\ 0 \to T^*(1) \to N \to \mathcal{O}(1) \to 0 \end{cases}$$

over π demonstrates that

$$\mathcal{W} \cong T^*(1)$$

as claimed. □

This result, together with (4.1.3), enables us to calculate the cohomology of $\mathcal{W}^*$ over any π, and thus its cohomology over all of $\mathcal{F}$, by Leray. We can now confirm the vanishing of various groups required for the proof of the theorem (4.1.4).

Lemma 4.2.3 *For all $p > 0$, and for all r,*

$$\mathrm{H}^p(\mathcal{F}; \Lambda^r\mathcal{W}^*) = 0.$$

(This subsumes lemma (4.2.1), by taking $r = 0$.)
Proof—By (4.2.2) and (4.1.3), we know that

$$\begin{aligned} \mathrm{h}^q(\pi; \Lambda^r\mathcal{W}^*) &= \mathrm{h}^q(\mathbb{P}^n; \Lambda^r T \otimes \mathcal{O}(-r)) \\ &= \mathrm{h}^{n-q}(\mathbb{P}^n; \Omega^r(r-n-1)) \\ &= 0 \end{aligned}$$

given that $q > 0$, and independently of r. Thus the sheaves

$$\mathrm{R}^q(\Lambda^r \mathcal{W}^*) = 0$$

for $q > 0$; and the noughth direct image is necessarily a trivial bundle over $\mathcal{Q}$ (of rank equal to the rank of the bundle of r-forms on $\mathcal{Q}$, of which more later (4.2.4)). The Leray sequence of $\Lambda^r \mathcal{W}^*$ collapses at E_2, with no non-zero terms except

$$E_2^{0,0} = \mathrm{H}^0(\mathcal{Q}; \mathrm{R}^0)$$

and we are done. □

The final information we need is the link between holomorphic forms on $\mathcal{Q}$ and sections of $\mathcal{W}^*$ on $\mathcal{F}$.

Lemma 4.2.4 *There are natural isomorphisms*

$$\begin{aligned} \mathrm{H}^0(\pi; \Lambda^r \mathcal{W}^*) &\cong \Lambda^r T_y^* \mathcal{Q} \\ \mathrm{H}^0(\mathcal{F}; \Lambda^r \mathcal{W}^*) &\cong \mathrm{H}^0(\mathcal{Q}; \Omega^r) \end{aligned}$$

for all $r = 0, \ldots, n$ and for any $y \in \mathcal{Q}$ with $\pi := g^{-1}\{y\}$.

(The lemma is immediate in the case $r = 0$, which we ignore).
Proof—Recall from (4.2.2) that there is a canonical identification

$$N = g^* T\mathcal{Q}$$

of bundles over $\mathcal{F}$, where N restricts to the normal bundle over each π; and secondly that we have an Euler sequence

$$0 \to \mathcal{O}(-1) \to N^* \to \mathcal{W}^* \to 0$$

over each fibre or over all of $\mathcal{F}$, where $\mathcal{O}(-1)$ is a line-bundle over $\mathcal{F}$ that restricts to a standard tautological line over each fibre.

Considering the exterior powers of the first observation, we find canonical isomorphisms

$$\begin{aligned} \Lambda^r T_y^* \mathcal{Q} &\cong \mathrm{R}^0(\Lambda^r N^*)|_y \\ &\cong \mathrm{H}^0(\pi; \Lambda^r N^*) \\ \mathrm{H}^0(\mathcal{Q}; \Omega^r) &\cong \mathrm{H}^0(\mathcal{Q}; \mathrm{R}^0 \Lambda^r N^*) \\ &\cong \mathrm{H}^0(\mathcal{F}; \Lambda^r N^*). \end{aligned}$$

The exterior powers of the Euler sequence [24] are, for $r \geq 1$,

$$0 \to \Lambda^{r-1} T(-r) \to \Lambda^r N^* \to \Lambda^r \mathcal{W}^* \to 0$$

as bundles over π or $\mathcal{F}$. (Here T is the tangent bundle in the fibre directions of ($\mathcal{F} \xrightarrow{g} \mathcal{Q}$).) Its long exact sequence gives

$$\mathrm{H}^0(\Lambda^{r-1} T(-r)) \to \mathrm{H}^0(\Lambda^r N^*) \to \mathrm{H}^0(\Lambda^r \mathcal{W}^*) \to \mathrm{H}^1(\Lambda^{r-1} T(-r))$$

taken over either π or $\mathcal{F}$. It remains to show that the first and fourth terms vanish in both cases.

Over π we know that

$$\begin{aligned} \mathrm{h}^q(\pi; \Lambda^{r-1}T(-r)) &= \mathrm{h}^{n-q}(\mathbb{P}^n; \Omega^{r-1}(r-n-1)) \\ &= 0 \end{aligned}$$

for all q; and $q = 0$ and 1 are sufficient for the first claim of the lemma. To extend over $\mathcal{F}$, note that the vanishing for all q means that all direct images R^q vanish, and the second term of Leray for $\Lambda^{r-1}T(-r)$ is completely empty. Thus we find

$$\mathrm{H}^p(\mathcal{F}; \Lambda^{r-1}T(-r)) = \mathrm{H}^p(\pi; \Lambda^{r-1}T(-r)) = 0$$

for all $p \geq 0$ and for all $r \geq 1$. □

This completes the proof of theorem (4.1.4).

4.3 Line-bundles: $n = 1$

The result for $n = 1$ is well-known to those in the know, but seems to have remained unpublished—

Theorem 4.3.1 *In the case $n = 1$, so that*

$$\begin{cases} \mathcal{P} = \mathbb{P}^2 - \textit{one point} \\ \mathcal{Q} = \{\textit{all lines in } \mathcal{P}\} \subset \mathbb{P}^{2*}, \end{cases}$$

there is a natural identification between holomorphic isomorphism classes of

- *line-bundles on $\mathcal{P}$ with zero Chern class; and*
- *line-bundles on $\mathcal{Q}$ with connection.*

Proof—By the exponential exact sequence (4.1.2), we can view classes of the first type as elements of

$$\mathrm{H}^1(\mathcal{P}; \mathcal{O}).$$

Since $\mathcal{Q} \cong \mathbb{C}^2$, any holomorphic bundle over $\mathcal{Q}$ is trivial; and a class of the second type is a holomorphic 1-form defined up to $\mathrm{d}f$ for any holomorphic function f.

These are identified under the twistor correspondence (4.1.4). □

We can view this observation as confirmation that the geometry of (4.1) is indeed relevant to the concerns of chapter 2. This supports our claim that the interpretation we shall put on the $n = 2$ case in (4.4) is not a coincidence, and does provide validation of the definition of a gerb. To back this up, we use the rest of this section to show that our design of twistor correspondence is an entirely reasonable way to treat a line bundle over $\mathcal{P}$.

Infinitesimal extensions over $l \in \mathcal{P}$

Choose a particular line bundle $\mathcal{L}$ that is trivial on each projective line $l \in \mathcal{P}$: this constraint is equivalent to the vanishing of the Chern class, and amounts to selecting a representative

$$\theta \in [\theta] \in \mathrm{H}^1(\mathcal{P}; \mathcal{O})$$

The twistor correspondence first lifts $\mathcal{L}$ to $\mathcal{F}$ via f^*. Then consider the noughth direct image sheaf on $\mathcal{Q}$: the constraint on $\mathcal{L}$ says that

$$\mathrm{R}^0(f^*\mathcal{L}) \to \mathcal{Q}$$

is in fact a line bundle, since

$$\mathrm{H}^0(l;\mathcal{L}) \cong \mathbb{C}.$$

But we can do more if we note, not only that $\mathcal{L}$ is trivialisable on any $(l;\mathcal{O})$, but that such a trivialisation has a unique extension to the first-order infinitesimal neighbourhood $(l;\mathcal{O}^1)$ of l. Further, there exists an obstruction in general to extension to the second-order neighbourhood.

These observations come from tensoring by $\mathcal{L}$ the extension sequence (4.1)

$$\big(0 \to S^qN^* \to \mathcal{O}^q \to \mathcal{O}^{q-1} \to 0\big) \otimes \mathcal{L}.$$

Since $\mathcal{L}|_l \cong \mathcal{O}$, the vanishing or otherwise of groups in the long exact sequence follows exactly the results of (4.1) for $\mathrm{H}^r(l;S^qN^*)$, namely

$$\mathrm{H}^0(l;N^*\otimes\mathcal{L}) = \mathrm{H}^1(l;N^*\otimes\mathcal{L}) = 0$$
$$\mathrm{H}^1(l;S^2N^*\otimes\mathcal{L}) \cong \mathbb{C}.$$

This is hinting that the sheaf over $\mathcal{Q}$ is not merely a (trivial, holomorphic) line bundle, but comes with a particular induced holomorphic connection whose curvature may not be zero. We can use the language of extensions explicitly to write down the connection and its curvature obstruction in terms of the Čech cocycle θ.

Given the multiplicity of perspectives on line bundles and their connections, it is not surprising that the same result can be reached in a number of ways. For instance, we can take the approach of (4.1.4) and write down transition functions for $\mathcal{W}$, thereby permitting us to evaluate $\mathrm{d}^{\mathcal{W}}f^*\theta$. Alternatively, we might expect considerations of parallel transport to proffer a route of attack.

Holonomy

Given a trivialisation over some $l_0 \cong \mathbb{P}^1 \subset \mathcal{P}$—ie. a non-zero section of $(\mathcal{L} \to l_0)$, or equivalently a non-zero point in the fibre of the direct image $R^0(f^*\mathcal{L})$ over the point $l_0 \in \mathcal{Q}$, with the obvious abuse of notation—there is a natural parallel translation of this along affine complex lines in $\mathcal{Q}$ through l_0, which we now describe.

If we pick a second $l_1 \in \mathcal{Q}$, then in $\mathcal{P}$ these two projective lines intersect in a single point (corresponding in $\mathcal{Q}$ to the $\mathbb{C}^1$ containing l_0 and l_1). The section of $\mathcal{L}$ over l_0 extends uniquely to one over l_1, since it singles out a (non-zero) point in the fibre over the intersection $l_0 \cap l_1 \in \mathcal{P}$; and this extends because $\mathcal{L}|_{l_1}$ is trivial.

Thus we have a parallel translation of the fibre of $\mathrm{R}^0(f^*\mathcal{L})$ over any $l_0 \in \mathcal{Q}$, in any straight-line direction in $\mathcal{Q}$. Infinitesimally, this creates a lift to any point in $R^0(f^*\mathcal{L})|_{l_0}$ of the holomorphic tangent space $T_{l_0}\mathcal{Q}$ of $\mathcal{Q}$ at l_0, and we have a holomorphic connection for $\mathrm{R}^0(f^*\mathcal{L})$.

If we try to extend the section from l_0 to a third line l_2, then there is no reason why the sections over l_1 and l_2 should agree at their single point of intersection (unless they happen to have a common intersection with l_0). That is, we expect non-trivial holonomy around triangles in $\mathcal{Q}$.

(We might anticipate this result via classical geometry: two lines in $\mathcal{P}$ form a degenerate conic in $\mathbb{P}^2$. A non-degenerate conic would be a rational curve, on which any line bundle such that $c_1 = 0$ is trivial. But three lines form a degenerate cubic; and a non-degenerate cubic would be an elliptic curve, which could certainly support topologically-trivial but holomorphically-nontrivial line bundles.)

We can take coordinates for $\mathcal{P}$ and a cocycle θ, and calculate directly that all this is indeed the case. The holonomy vanishes if and only if $\mathcal{L}$ is trivial, ie. if and only if

$$[\theta] = 0 \in \mathrm{H}^1(\mathcal{P}; \mathcal{O}).$$

Coordinate calculation

Let us illustrate one of the three approaches discussed above: that of infinitesimal extensions. (A second, the twistor transform via $\mathcal{W}$, will be displayed in (4.4).)

Pick a base $l_0 \in \mathcal{P}$. This is covered by two affine-$\mathbb{C}^1$ patches, with coordinates z and $w = 1/z$. Since $\mathcal{P}$ is an $\mathcal{O}(1)$-bundle over l_0, we have two $\mathbb{C}^2$ -patches for $\mathcal{P}$ with coordinates

$$\begin{cases} (z,p) & w = 1/z \\ (w,q) & q = p/z. \end{cases}$$

We also have coordinates $(a,b) \in \mathbb{C}^2$ for $\mathcal{Q}$, corresponding to the projective line

$$\begin{cases} p &= a + bz \\ q &= aw + b \end{cases}$$

in $\mathcal{P}$.

Choose a cocycle

$$\theta = \theta(z,p) = \sum_{0<j<i} \theta_i^j \frac{p^i}{z^j}.$$

This corresponds to a line-bundle $\mathcal{L}$ with total space covered by

$$\begin{cases} (z,p,s) \\ (w,q,t) \\ t = e^{-\theta} s. \end{cases}$$

(In general we could take an isomorphic bundle by adding coboundary terms, but we shall just take this representative, as in (4.1.1).)

Pick some projective line $l \subset \mathcal{P}$, corresponding to $(a,b) \in \mathcal{Q}$. Restricted to l, $\mathcal{L}$ is trivial with sections

$$s = c \cdot e^{\theta^{0+}}, \quad t = c \cdot e^{-\theta^-}$$

where c is any constant and we view θ restricted to l as a function of (z,a,b); and then "θ^{0+}" means terms in non-negative powers of z. (To hold that $\theta = \theta(z,a,b)$ is to consider it as a cocycle on $\mathcal{F}$ in terms of $(\mathcal{F} \xrightarrow{g} \mathcal{Q})$; but we pass over this until section (4.4).)

We seek to extend this section to the first formal neighbourhood of l. That is, we require

$$\begin{cases} s = s_0 + \nu s_1 & \nu := p - a - bz \\ t = t_0 + \mu t_1 & \mu := q - aw - b = \nu/z \end{cases}$$

in which s_0 and s_1 are functions of (z, a, b), and t_r are functions of (w, a, b), obeying transition functions

$$t \equiv e^{-\theta} s \mod \mathcal{I}^2$$

ie. we set ν^2 to zero. Modulo this,

$$\begin{aligned} \exp(-\theta(z,p)) &= \exp(-\theta(z,a,b) - \nu\partial\theta) \\ &= \exp(-\theta(z,a,b)) \cdot \big(1 - \nu\partial\theta\big) \end{aligned}$$

where

$$\partial\theta := \frac{\partial\theta}{\partial p}(z,p)\big|_{p=a+bz}$$

so that

$$\begin{aligned} t_0 + \mu t_1 &= e^{-\theta}(s_0 + \nu s_1) \\ &= e^{-\theta(z,a,b)}\big(s_0 + \nu(s_1 - \partial\theta \cdot s_0)\big). \end{aligned}$$

The terms of degree zero in ν just describe a section of $\mathcal{L}$ over l, so we set s_0 and t_0 to equal our previous choice $c \neq 0$ of such a section. We need also to show that there are now unique non-singular choices of $O(\nu)$ terms. But the constraint is

$$\frac{t_1}{t_0} = \frac{1}{w}\big(\frac{s_1}{s_0} - \partial\theta\big)$$

and the presence of the $1/w$ means that indeed there is no choice: we must define

$$\begin{aligned} s_1 &:= s_0 \cdot (\partial\theta)^{0+} \\ t_1 &:= \frac{t_0}{w} \cdot (\partial\theta)^{-} \end{aligned}$$

to avoid singularities at $z = 0$ or $w = 0$.

Finally, consider a putative extension to the second infinitesimal neighbourhood: recall that we expect an obstruction in general. Suppose there exist some

$$\begin{cases} s = s_0 + \nu s_1 + \nu^2 s_2 & \nu := p - a - bz \\ t = t_0 + \mu t_1 + \mu^2 t_2 & \mu := q - aw - b = \nu/z \end{cases}$$

equated under the transition modulo ν^3. To this order

$$\begin{aligned} \exp(-\theta(z,p)) &= \exp(-\theta(z,a,b) - \nu\partial\theta - \frac{\nu^2}{2}\partial^2\theta) \\ &= \exp(-\theta(z,a,b)) \cdot \big(1 - \nu\partial\theta + \frac{\nu^2}{2}((\partial\theta)^2 - \partial^2\theta)\big) \end{aligned}$$

and the $O(1)$, $O(\nu)$ are set as before. But the second-order sections then require

$$\frac{t_2}{t_0} = \frac{1}{w^2}\big(\frac{s_2}{s_0} - \frac{X}{2}\big)$$

where

$$X = X(z,a,b) = \big((\partial\theta)^{0+}\big)^2 - \big((\partial\theta)^{-}\big)^2 + \partial^2\theta$$

and so, thanks to the factor $1/w^2$, a division of X (as a Laurent series in z) to give non-singular functions s_2, t_2 is obstructed by the terms of X of order $1/z$. Such terms exist only in $\partial^2\theta$, and equal

$$\sum_{0<i<j} \begin{pmatrix} i \\ j-1 \end{pmatrix} \theta_i^j a^{i-j-1} b^{j-1}$$

which vanishes if and only if $[\theta] = 0$.

If we were to calculate the connection on the induced bundle over $\mathcal{Q}$ by either of the two other techniques discussed above, we would in fact find a well-defined curvature 2-form

$$\sum_{0<i<j} \begin{pmatrix} i \\ j-1 \end{pmatrix} (\theta_i^j a^{i-j-1} b^{j-1}) \mathrm{d}a \wedge \mathrm{d}b.$$

Compare the expression for the gerb 3-form at the end of (4.4).

4.4 Gerbs: $n = 2$

Theorem 4.4.1 *Consider $\mathcal{P}$ and $\mathcal{Q}$ in the case $n = 2$*

$$\begin{cases} \mathcal{P} &= \mathbb{P}^3 - \textit{one point} \\ \mathcal{Q} &= \{\textit{all planes in } \mathcal{P}\} \subset \mathbb{P}^{3*}. \end{cases}$$

There is a canonical identification between holomorphic equivalence classes of

- *gerbs on $\mathcal{P}$; and*
- *gerbs with 1-connection on $\mathcal{Q}$.*

Proof—The first alternative corresponds to elements of

$$\mathrm{H}^2(\mathcal{P}; \mathcal{O}^*) \cong \mathrm{H}^2(\mathcal{P}; \mathcal{O})$$

by (4.1.2). The second corresponds as in (2.3.1) to elements of

$$\mathrm{H}^0(\mathcal{Q}; \Omega^2)/\mathrm{H}^0(\mathcal{Q}; \mathrm{d}\Omega^1).$$

Apply theorem (4.1.4). □

The 3-curvature of the gerb over $\mathcal{Q}$ is the exterior derivative of any 2-form representative of the class in $\mathrm{H}^0(\mathcal{Q}; \Omega^2)/\mathrm{H}^0(\mathcal{Q}; \mathrm{d}\Omega^1)$.

Strictly, the remainder of this section is superfluous, but as in (4.3) it seems apposite to illustrate the rather austere result above with some explicit demonstration. We might take a holonomy point of view, buoyed by the classical observation that two or three planes describe degenerate quadric or cubic surfaces, whose non-degenerate versions would be merely $\mathbb{P}^2$ (on which any gerb would be trivial). But if we had a non-degenerate *quartic*, a K3 surface, then since $\mathrm{H}^2(\mathcal{O})$ is now non-zero it can support non-trivial gerbs. So, given an object on a plane in $\mathcal{P}$, we expect to be able to extend it to an object over any three planes; but some obstruction to a fourth extension will carry the "holonomy" around a complex tetrahedron in $\mathcal{Q}$ (7.1).

More mundanely, we shall content ourselves with taking coordinates for $\mathcal{P}$ and $\mathcal{Q}$, and displaying transition functions for $\mathcal{W}^*$. Then we work through the twistor transform of a gerb over $\mathcal{P}$ to calculate the resultant 3-curvature on $\mathcal{Q}$. Even in this low-dimensional case $n = 2$, there are many coordinate patches to track and we shall not present every step.

The geometry of $\mathcal{F}$ and $\mathcal{W}^*$

Fix some base projective plane $\pi_0 \subset \mathcal{P}$, and cover π_0 with three affine sets (z, w), (u, v), (s, t). Hence coordinate patches for $\mathcal{P}$

$$\begin{cases} (z, w, p) \in U_2 & z = u/v = 1/t \\ (u, v, q) \in U_1 & w = 1/v = s/t \\ (s, t, r) \in U_0 & p = q/v = r/t, \text{ etc.} \end{cases}$$

The bundle $(\mathcal{F} \xrightarrow{f} \mathcal{P})$ has three lifted patches. A point in $\mathcal{F}_2$ for instance corresponds to a point $(z, w, p) \in U_2$ and a plane π containing it. The plane is specified by two further coordinates: take σ_2 to be the value of q at the point $u = v = 0$, and τ_2 to be r at $s = t = 0$. We have similarly

$$\begin{cases} (z, w, p, \sigma_2, \tau_2) \in \mathcal{F}_2 & \\ (u, v, q, \sigma_1, \tau_1) \in \mathcal{F}_1 & \sigma_1 = \tau_2,\ \tau_1 = p - \tau_2 z - \sigma_2 w \\ (s, t, r, \sigma_0, \tau_0) \in \mathcal{F}_0 & \sigma_0 = p - \tau_2 z - \sigma_2 w = \tau_1,\ \tau_0 = \sigma_2 \end{cases}$$

By definition the bundle $(\mathcal{W} \to \mathcal{F})$ has local bases of sections

$$\partial/\partial\sigma_i,\ \partial/\partial\tau_i$$

and, with respect to the dual bases, sections of $\mathcal{W}^*$ have coefficients

$$\begin{aligned} \begin{pmatrix} e_2 \\ f_2 \end{pmatrix} &= \begin{pmatrix} 0 & -w \\ 1 & -z \end{pmatrix} \begin{pmatrix} e_1 \\ f_1 \end{pmatrix} \\ &= \begin{pmatrix} -w & 1 \\ -z & 0 \end{pmatrix} \begin{pmatrix} e_0 \\ f_0 \end{pmatrix} \\ \begin{pmatrix} e_1 \\ f_1 \end{pmatrix} &= \begin{pmatrix} -u & 1 \\ -v & 0 \end{pmatrix} \begin{pmatrix} e_2 \\ f_2 \end{pmatrix} \\ \begin{pmatrix} e_0 \\ f_0 \end{pmatrix} &= \begin{pmatrix} 0 & -t \\ 1 & -s \end{pmatrix} \begin{pmatrix} e_2 \\ f_2 \end{pmatrix}. \end{aligned}$$

We are now in a position to study $\mathcal{W}^*$ directly, if we wish, rather than rely on the sheaf cohomology of (4.1). For instance, we can now observe the Chern class of $\mathcal{W}^*$ restricted to any plane $\pi \subset \mathcal{F}$ by taking two sections

$$\begin{pmatrix} e_2 \\ f_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ or } \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

to show that $c = 1 + h + h^2$ (where h is the class of a line in $\mathbb{P}^2$). It can be shown that $\mathcal{W}^*|_\pi$ has no sub-line-bundle, whilst its splitting type is uniformly $\mathcal{O} \oplus \mathcal{O}(1)$.

Since $\mathcal{W}^*$ is homogeneous by definition, and thus uniform, Van de Ven shows that it must be a sum of lines or a twisted tangent bundle [51, 39]. Then its Chern class forces it up to isomorphism to be

$$\mathcal{W}^* \cong T \otimes \mathcal{O}(-1)$$

over any π, agreeing with the general result (4.2.2). (A third and most immediate proof of this result arises from a direct comparison of transition functions for $\mathcal{W}$ and T^*.)

The projection $(\mathcal{F} \xrightarrow{g} \mathcal{Q})$ can be dealt with by the same three patches

$$(a, b, c, \zeta_i, \eta_i) \in \mathcal{F}_i$$

—projecting by g to coordinates $(a, b, c) \in \mathcal{Q}$—which correspond to the coordinates on $(\mathcal{F} \xrightarrow{f} \mathcal{P})$ by

$$\begin{aligned} a &= p - \tau_2 z - \sigma_2 w \\ b &= \sigma_2 \\ c &= \tau_2 \\ \zeta_2 &= z \\ \eta_2 &= w \end{aligned}$$

so that the point $(a, b, c) \in \mathcal{Q}$ has over it the $\mathbb{P}^2$-fibre (ζ_i, η_i), and corresponds in $\mathcal{P}$ with the plane

$$p = cz + bw + a.$$

Computation of the twistor correspondence

So much for the underlying geometry. Now on the triple intersection $U_{2,1,0} \subset \mathcal{P}$ we pick some gerb represented as

$$\theta = \theta(z, w, p) = \sum_{\substack{s,t>0 \\ r>s+t}} \theta_r^{s,t} \frac{p^r}{z^s w^t}$$

(plus any coboundary, which we ignore). This lifts to $f^*\theta$ in $\mathcal{F}$, which as a function of $(z, w, p, \sigma_2, \tau_2)$ looks the same. (Equivalently, $\mathrm{d}^{\mathcal{W}} f^*\theta$ vanishes.)

We convert to coordinates $(a, b, c, \zeta_2, \eta_2)$ of $(\mathcal{F} \xrightarrow{g} \mathcal{Q})$, and represent $f^*\theta$ as a coboundary in the form

$$\theta = \delta\mu$$

where we might choose for example

$$\begin{aligned} \mu_2^1 &:= \text{all terms in non-negative powers of } \zeta_2 \\ \mu_1^0 &:= \text{all terms in negative powers of } \zeta_2 \text{ and of } \eta_2 \\ \mu_0^2 &:= \text{all negative powers of } \zeta_2 \text{ and non-negative powers of } \eta_2. \end{aligned}$$

In general,

$$\begin{aligned} \mathrm{d}^{\mathcal{W}}\mu &= \frac{\partial\mu}{\partial\sigma}\mathrm{d}\sigma + \frac{\partial\mu}{\partial\tau}\mathrm{d}\tau \\ &= \left(\frac{\partial\mu}{\partial b} - \eta\frac{\partial\mu}{\partial a}\right)\mathrm{d}b + \left(\frac{\partial\mu}{\partial c} - \zeta\frac{\partial\mu}{\partial a}\right)\mathrm{d}c \end{aligned}$$

whilst

$$\begin{aligned}\theta &= \sum \theta^r_{s,t} \cdot (a + c\zeta + b\eta)^r / \zeta^s \eta^t \\ &= \sum \theta^r_{s,t} \left(\begin{array}{c} r \\ i,j \end{array} \right) (a^{r-i-j} c^i b^j) \cdot \zeta^{i-s} \eta^{j-t}\end{aligned}$$

summed over all $r \geq 3$, non-negative i and j such that $i+j \leq r$, and positive s and t such that $s+t<r$. The combinatorial symbol is to mean

$$\left(\begin{array}{c} r \\ i,j \end{array} \right) := \frac{r!}{i!j!(r-i-j)!}.$$

Our choice of coboundary μ is then

$$\begin{aligned}\mu^1_2 &= \text{terms with } s \leq i \\ \mu^0_1 &= \text{terms with } s > i \text{ and } t \leq j \\ \mu^2_0 &= \text{terms with } s > i \text{ and } t > j.\end{aligned}$$

Next is the differentiation $\mathrm{d}^{\mathcal{W}}\mu$ and a further choice of coboundary

$$\delta\nu = \mathrm{d}^{\mathcal{W}}\mu.$$

Then the objects $\mathrm{d}^{\mathcal{W}}\nu_i$ will naturally be holomorphic 2-forms on $\mathcal{Q}$, and their exterior derivative is the well-defined curvature form we seek to exhibit. This is a messy task: we eschew a full demonstration of consistency in favour of dealing with μ^0_1 only, and finally differentiating ν_0 to display the 3-curvature.

The first differentiation gives

$$\begin{aligned}\mathrm{d}^{\mathcal{W}}\mu^0_1 &= \sum \theta^r_{s,t} \cdot \frac{\mathrm{d}b}{\zeta^s_2 - i} \Big(\frac{-r!}{i!(t-1)!(r-i-t)!} a^{r-t-i} b^{t-1} c^i \Big) \\ &+ \sum \theta^r_{s,t} \cdot \frac{\mathrm{d}c}{\eta^t_2 - j} \Big(\frac{-r!}{(s-1)!j!(r-s-j)!} a^{r-s-j} b^j c^{s-1} \Big)\end{aligned}$$

where the first term is summed over (r, s, t) as ever and also $0 \leq i < s$; and the second is over (r, s, t) and $0 \leq j < t$. We want some

$$\begin{aligned}\nu_2 &= B_2 \mathrm{d}\sigma_2 + C_2 \mathrm{d}\tau_2 \\ &= B_2 \mathrm{d}b + C_2 \mathrm{d}c \\ \nu_1 &= B_1 \mathrm{d}\sigma_1 + C_1 \mathrm{d}\tau_1 \\ &= (-\eta_2 C_1)\mathrm{d}b + (B_1 - \zeta_2 C_1)\mathrm{d}c \\ \nu_0 &= B_0 \mathrm{d}\sigma_0 + C_0 \mathrm{d}\tau_0 \\ &= (C_0 - \eta_2 B_0)\mathrm{d}b + (-\zeta_2 B_0)\mathrm{d}c\end{aligned}$$

where the B and C are functions of $(a, b, c, \zeta_i, \eta_i)$ that are non-singular on f^*U_i. In terms of ζ_2 and η_2, this means that B_1 and C_1 can contain only non-positive powers of η_2, and terms in non-negative powers of ζ_2/η_2; and B_0 and C_0 have non-positive powers of ζ_2, and non-negative powers of η_2/ζ_2.

The requirement for ν is that

$$\begin{aligned}(\delta\nu)^0_1 &= (\eta_2(B_0 - C_1) - C_0)\mathrm{d}b + (\zeta_2(B_0 - C_1) + B_1)\mathrm{d}c \\ &= \mathrm{d}^{\mathcal{W}}\mu^0_1.\end{aligned}$$

Considering the cofficient of db in $\mathrm{d}^{\mathcal{W}}\mu_1^0$, we could for instance set $-C_0$ equal to this, and take $B_0 = C_1$ equal to any function purely in (a, b, c). (The only alternative is to add such a function to C_0 and subtract this $\times 1/\eta_2$ from C_1.) And so on: the simplest solution consistent with what is so far displayed is

$$\begin{aligned} B_0 &= 0 \\ C_0 &= -\text{the d}b\text{ term of d}^{\mathcal{W}}\mu_1^0 \\ B_1 &= +\text{the d}c\text{ term of d}^{\mathcal{W}}\mu_1^0 \\ C_1 &= 0 \end{aligned}$$

and without further justification we claim that these choices do indeed extend to a fully consistent ν. Given this, the three fields $\mathrm{d}^{\mathcal{W}}\nu_i$ should each equal the same 1-connection 2-form on $\mathcal{Q}$. If $\mathrm{d}^{\mathcal{W}}\nu_0$ is calculated, for example, a string of cancellations leaves

$$\begin{aligned} \mathrm{d}^{\mathcal{W}}\nu_0 &= \sum \theta^r_{s,t}\Big(\frac{r!}{(t-1)!(s-1)!(r-t-s)!}a^{r-t-s}b^{t-1}c^{s-1}\Big)\cdot \mathrm{d}b\wedge \mathrm{d}c \\ &= \frac{\partial^2}{\partial b\partial c}\Big(\theta(a,b,c,\zeta_2,\eta_2)|_{\text{terms of order zero in }\zeta_2,\eta_2}\Big)\cdot \mathrm{d}b\wedge\mathrm{d}c \end{aligned}$$

which is explicitly independent of fibre-coordinates in $(\mathcal{F} \xrightarrow{g} \mathcal{Q})$, and so descends to a 2-form over $\mathcal{Q}$.

The exterior derivative of this in $\mathcal{Q}$ is the 3-curvature of the gerb, which is independent of all choices made. After such involved calculations, its form is reassuringly straightforward

$$\begin{aligned} \Omega &= \mathrm{d}(g_*\mathrm{d}^{\mathcal{W}}\nu_0) \\ &= \frac{\partial^3}{\partial a\partial b\partial c}\Big(\theta(a,b,c,\zeta_2,\eta_2)|_{\text{order zero in }\zeta_2,\eta_2}\Big)\cdot \mathrm{d}a\wedge\mathrm{d}b\wedge\mathrm{d}c \\ &= \sum_{\substack{s,t>0 \\ r>s+t}} \theta^{s,t}_r \begin{pmatrix} r \\ s,t \end{pmatrix} \frac{\partial^3}{\partial a\partial b\partial c} a^{r-s-t}b^t c^s \cdot \mathrm{d}a\wedge\mathrm{d}b\wedge\mathrm{d}c \end{aligned}$$

and vanishes if and only if $\theta = 0$. Starting from $\theta(z, w, p)$, therefore, the net result is to substitute

$$p = cz + bw + a$$

retaining only those terms of order zero in z and in w, and to differentiate with respect to a, b and c. Compare the end of (4.3)

4.5 Higher-order gerbs

Given that theorem (4.1.4) extends beyond the case $n = 2$, this seems an appropriate juncture at which to consider higher extensions of the gerb concept.

0-, 1- and 2-gerbs

Let us summarise the pattern so far developed. Suppose given a manifold X with an open cover U_i. We shall describe the smooth case: the holomorphic definitions are similar, though we do not pause to worry about matters of existence and Atiyah obstructions.

The fundamental structure is a *0-gerb*, defined to be a global $\underline{\mathbb{C}}^*$-function on X. No equivalences, no connective structures, no objects. Such things are classified by $\mathrm{H}^0(X;\underline{\mathbb{C}}^*)$, and have a curvature 1-form just equal to $\mathrm{d}\log$ of the function.

A *1-gerb* is given by $\underline{\mathbb{C}}^*$-functions (0-gerbs) on each $U_{i,j}$ such that their coboundary is trivial—a global line bundle, in other words. A 0-connection is a collection of local 1-forms A_i whose coboundary is the curvature of each 0-gerb ($\delta A = \mathrm{d}\log f$, ie. a line-bundle connection). An object is defined by a 0-gerb on each U_i such that the coboundary is identified with the given 0-gerb on $U_{i,j}$, thus describing a global nowhere-zero section. A 0-equivalence (a gauge transformation) is given by arbitrary 0-gerbs on each U_i (an element of $\mathrm{C}^0(X;\underline{\mathbb{C}}^*)$), which act by their coboundary on the transition 0-gerbs, and by their curvatures on the 0-connection 1-forms A_i. Spaces of equivalence classes are $\mathrm{H}^1(\mathcal{O}^*)$ and $\mathbb{H}^1(\underline{\mathbb{C}}^* \xrightarrow{\mathrm{d}\log} \mathcal{A}^1)$.

A *2-gerb*—or a gerb—has bundles (1-gerbs) on each $U_{i,j}$ with trivialisations (objects of 1-gerbs) on $U_{i,j,k}$ whose coboundary 0-gerb vanishes. Bundle 0-connections making the trivialisations parallel describe a 0-connection. An object with connection is as per usual. 0- and 1-transforms give equivalence classes in second cohomology $\mathrm{H}^2(\underline{\mathbb{C}}^*)$ and in

$$\mathbb{H}^2(\underline{\mathbb{C}}^* \xrightarrow{\mathrm{d}\log} \mathcal{A}^1 \xrightarrow{\mathrm{d}} \mathcal{A}^2).$$

We refer to a gerb here as a 2-gerb merely because this matches the pattern $n = 2$ of theorem (4.1.4). It is also perhaps convenient to start at level zero with functions. A category theorist following the Giraud/Brylinski definition would no doubt prefer something different: a bundle (a sheaf of 0-categories) would be a 0-gerbe, and a gerbe (a sheaf of categories, or 1-categories) a 1-gerbe. Given useful definitions of higher-order categories, the sheaf description should offer an alternative approach to higher cohomology.

No doubt the categorists can claim to offer the most fundamental approach to geometry. But this does not seem likely to improve the unwieldly nature of the categorical approach to gerbes—as viewed by those unwilling to fraternise with torsors—and one can conceive that each stage could become increasingly obscure. There is also Murray's infinite-dimensional approach in defining higher bundle gerbes [7], which struggles in handling issues of equivalence. Meanwhile our rather scruffy Čech approach can easily step up the gears, as we shall now observe.

3-gerbs

We sketch the natural constructions. (Full expansions and proofs can be given, granted no more sophistication than that required by chapter 2.)

Given a manifold X with cover U_i, a *3-gerb* is made up of: gerbs $\mathcal{G}_i^j$ on pairwise intersections (such that $\mathcal{G}_j^i$ is the dual gerb); objects $\Gamma_{i,j,k}$ of the coboundary gerbs on $U_{i,j,k}$ (which we thus must insist be trivial); and trivialisations $\theta_{i,j,k,l}$ of the coboundaries of the objects (which are naturally line-bundles—they are required to be trivial) such that $\delta\theta = can$, where can is the canonical section (which exists automatically) of $\delta^2\Gamma$.

Some elaboration: given a locally-trivial gerb over some cover, its *dual* is described by the same cover with the dual bundles and the dual section on 3-

fold intersections. Then the tensor-product of the two (2.1.11) is again on the same cover, with trivialised bundles on pairwise intersections and the canonical trivialisation on 3-fold sets. (So in Čech terms, we have $t_{i,j,k} \equiv 1$.)

Back to the 3-gerb: whilst each object Γ is not itself a line-bundle, nevertheless we can view their coboundary

$$\Gamma \otimes \Gamma^{-1} \otimes \Gamma \otimes \Gamma^{-1}$$

on $U_{i,j,k,l}$ as a pair of ratios of objects. Each ratio is naturally a line-bundle (3.1.3). So to ask for a *trivialisation* θ of $\delta\Gamma$ is legitimate.

And on $U_{i,j,k,l}$ the double coboundary of Γ consists of a tensor product of ten terms of the form

$$(\Gamma \otimes \Gamma^{-1})_{i,j,k} \otimes (\text{others}).$$

But $\Gamma \otimes \Gamma^{-1}$ is a canonically-trivial line-bundle, thus giving *can*.

Equivalences

To find a Čech cocycle in all of this, refine such that the $\mathcal{G}_i^j$ are trivial and select an object for each. Then, relative to the coboundary of these objects, the object $\Gamma_{i,j,k}$ is a line-bundle: refine such that we can pick a trivialisation. Then $\theta_{i,j,k,l}$ becomes a $\underline{\mathbb{C}}^*$-function $t_{i,j,k,l}$ with respect to this basis, and necessarily has vanishing coboundary

$$t \in \mathrm{Z}^3(\underline{\mathbb{C}}^*).$$

(This also shows the existence of 3-gerbs.)

Changing the bases shifts t by $\delta\mathrm{C}^2(\underline{\mathbb{C}}^*)$, whilst changing the objects of $\mathcal{G}$ shifts by δ^2 of line-bundles, which is canonically trivial. So to no great surprise we find

$$[t] \in \mathrm{H}^3(\underline{\mathbb{C}}^*).$$

What is the natural definition of equivalence? Suppose given two 3-gerbs $(\mathcal{G}, \Gamma, \theta)$ and $(\mathcal{H}, \Delta, \tau)$ on a common refinement. Since $\delta(\mathcal{G}_i^j)$ and $\delta(\mathcal{H}_i^j)$ both carry objects, they are both trivial and so 0-equivalent. But it need not be that their objects are 0-equivalent (3.1.4). We say that these *3-gerbs are equivalent*, if there exists a gerb-equivalence of $\delta\mathcal{G}$ and $\delta\mathcal{H}$ that makes the objects Γ and Δ equivalent (ie. their difference is a *trivial* bundle) and sends θ to τ.

Given two equivalent 3-gerbs, and comparing with the construction of $[t]$ above, we take bases for the trivial difference bundles between Γ and Δ to find that equivalence shifts by

$$t' = t + \delta(\text{something}).$$

Thus the space of equivalence classes is $\mathrm{H}^3(\underline{\mathbb{C}}^*)$.

Connections

A *0-connection* on a 3-gerb $(\mathcal{G}, \Gamma, \theta)$ is determined by a gerb 0-connection on each $\mathcal{G}_i^j$ and object connections with respect to this on each Γ. The coboundary of the object connections is a bundle connection, for which we insist θ is parallel.

A *1-connection* is given by gerb 1-connections on the $\mathcal{G}_i^j$ with respect to their existing 0-connections. These induce a 1-connection on $\delta\mathcal{G}$—this 1-connection must be *objective* (3.3.1) with respect to the extant object connection.

Now there exists a cocycle of 3-form curvatures Ω_i^j. A *2-connection* is a family of local 3-forms

$$\Pi_i \in \mathrm{C}^0(X; \mathcal{A}^3)$$

such that

$$\delta\Pi = \Omega,$$

thereby leaving us with a *global 4-form curvature*

$$M := \mathrm{d}\Pi_i \quad (\forall i).$$

The various levels of gauge transformation are clear: a *0-transform* is a 3-gerb equivalence as above, which affects the 0-connection in the obvious way. A *1-transform* is a family of 1-transforms of the $\mathcal{G}_i^j$, and so affects the 1- and 0-connections but not the underlying 3-gerb. A *2-transform* is a cochain of 2-forms

$$\zeta_i \in \mathrm{C}^0(X; \mathcal{A}^2)$$

which leaves untouched the 3-gerb and its 0-connection, sends the 1-connection to

$$\beta_i^j \mapsto \beta_i^j + \delta\zeta_i$$

and shifts the 2-connection by

$$\Pi_i \mapsto \Pi_i + \mathrm{d}\zeta_i.$$

The full space of equivalence classes is then

$$\mathbb{H}^3(\underline{\mathbb{C}}^* \xrightarrow{\mathrm{d}\log} \mathcal{A}^1 \xrightarrow{\mathrm{d}} \mathcal{A}^2 \xrightarrow{\mathrm{d}} \mathcal{A}^3)$$

by arguments which are by now transparent.

n-gerbs

By now it should be clear that we can continue without needing any new technology—Čech cohomology is quite adequate to define gerbs of arbitrary order. We resist all further demonstrations bar one.

Theorem 4.5.1 *Given the spaces*

$$\begin{cases} \mathcal{P} & := & \mathbb{P}^{n+1} - \textit{one point} \\ \mathcal{Q} & := & \{\textit{all hyperplanes } \mathbb{P}^n \subset \mathcal{P}\} \end{cases}$$

for arbitrary $n \geq 0$, there is a canonical identification between holomorphic equivalence classes of

- *n-gerbs on $\mathcal{P}$ with vanishing Chern class; and*
- *n-gerbs with $(n-1)$-connection on $\mathcal{Q}$.*

(And the Chern class constraint is automatic in case n is even.)

Proof—This is a repetition of (4.1.2) and (4.1.4). Even without writing out explicit definitions of holomorphic n-gerbs, we can be confident that their equivalence classes on $\mathcal{P}$ are elements of $\mathrm{H}^n(\mathcal{P}; \mathcal{O}^*)$; and that n-gerbs with connection on $\mathbb{C}^{n+1}$ are classified by n-forms up to exact forms

$$\mathrm{H}^0(\mathcal{Q}; \Omega^n)/\mathrm{H}^0(\mathcal{Q}; \mathrm{d}\Omega^{n-1}).$$

(The result applies even when $n = 0$, for which $\mathcal{P}$ and $\mathcal{Q}$ are identified, and equal $\mathbb{C}^1$). □

Chapter 5

Holomorphic gerbs

Holomorphic gerbs on a complex manifold X have already been discussed briefly (2.3.5); and they arise naturally in chapter 4 in a simple enough form not to require further details. But of course there is more to be said, and in this chapter we pursue their theory more fully.

5.1 Equivalence classes and moduli

Recall that a holomorphic gerb has a cohomology class

$$[\mathcal{G}] \quad \in \quad \mathrm{H}^2(X;\mathcal{O}^*)$$

and a Chern class

$$c(\mathcal{G}) \quad \in \quad \mathrm{H}^3_{\mathrm{dR}}(X;\mathbb{Z}).$$

The Chern class is real; and we can restrict its type. By

$$\mathrm{H}^{(2,1)+(1,2)}_{\mathrm{dR}}(X)$$

on a general complex manifold X, we mean those de Rham classes with some d-closed representative containing types (2,1) and (1,2) only.

Proposition 5.1.1 *The possible Chern classes $c(\mathcal{G}) \in \mathrm{H}^3_{\mathrm{dR}}(X;\mathbb{Z})$ are exactly those classes lying in*

$$\mathrm{H}^{(2,1)+(1,2)}_{\mathrm{dR}}(X;\mathbb{Z}).$$

Proof—We seek the image in $\mathrm{H}^3(X;\mathbb{C})$ of the kernel (in the exponential sequence) of

$$\mathrm{H}^3(X;\mathbb{Z}) \to \mathrm{H}^3(X;\mathcal{O}).$$

It suffices to show that the real part of the kernel of

$$\mathrm{H}^3(X;\mathbb{C}) \xrightarrow{i^*} \mathrm{H}^3(X;\mathcal{O})$$

(induced by $\mathbb{C} \hookrightarrow \mathcal{O}$) is

$$\mathrm{H}^{(2,1)+(1,2)}(X;\mathbb{C}).$$

Construct these groups via resolutions of their sheaves [50]

$$\begin{cases} 0 & \rightarrow & \mathbb{C} & \rightarrow & \underline{\mathbb{C}} & \stackrel{\mathrm{d}}{\rightarrow} & \mathcal{A}^1 & \stackrel{\mathrm{d}}{\rightarrow} & \dots \\ 0 & \rightarrow & \mathcal{O} & \rightarrow & \underline{\mathbb{C}} & \stackrel{\bar\partial}{\rightarrow} & \mathcal{A}^{0,1} & \stackrel{\bar\partial}{\rightarrow} & \dots \end{cases}$$

so that i^* is induced by sending a d-closed 3-form to its (0,3)-component. The kernel is those classes with some representative devoid of this component. Its real part is those classes with some representative containing only (2,1)- and (1,2)-terms. □

Now the exponential sequence yields the moduli space of gerbs of fixed topological type as a complex torus

$$0 \rightarrow \frac{\mathrm{H}^2(\mathcal{O})}{\mathrm{H}^2(\mathbb{Z})} \rightarrow \mathrm{H}^2(\mathcal{O}^*) \rightarrow \mathrm{H}^3(\mathbb{Z}).$$

The first term involves only $\mathrm{H}_{\mathrm{dR}}^{(2,0)+(0,2)}(\mathbb{Z})$; and the image on the right is exactly $\mathrm{H}_{\mathrm{dR}}^{(3,0)+(0,3)}(\mathbb{Z})$ plus all torsion in $\mathrm{H}^3(\mathbb{Z})$.

The torus is not a Jacobian under the standard definition (intermediate or otherwise), since such things arise in odd degree, but the pattern remains that set by the theory of line-bundles.

5.2 Holomorphic connections

A *holomorphic 0-connection* is a 0-connection (2.2.1) on a holomorphic gerb, each of whose line-bundle connections is holomorphic. (A line-bundle connection is holomorphic if it behaves as

$$\nabla : \mathcal{L} \rightarrow \mathcal{L} \otimes \Omega^1$$

mapping between holomorphic sheaves rather than to the more general smooth sheaf $\mathcal{L} \otimes \mathcal{A}^1$.) Unlike the smooth case (2.2.2), these need not exist. The reason is similar to that for the line-bundle Atiyah class in $\mathrm{H}^1(\Omega^1)$.

Proposition 5.2.1 *There is an obstruction*

$$\alpha_0 \in \mathrm{H}^2(X; \Omega^1)$$

to the existence of a holomorphic 0-connection.

Proof—Given a cocycle

$$t \in \mathrm{Z}^2(X; \mathcal{O}^*)$$

representing θ in some local trivialisation, a choice of holomorphic 0-connection amounts to a choice of 1-forms

$$A \in \mathrm{C}^1(X; \Omega^1)$$

such that

$$\delta A = \mathrm{d} \log t \in \mathrm{Z}^2(X; \Omega^1).$$

But unlike in the smooth case, in which $\mathrm{H}^2(\mathcal{A}^1) \equiv 0$, this need not be possible

$$\alpha_0 := \left[\frac{1}{2\pi i}\mathrm{d}\log t\right] \in \mathrm{H}^2(X;\Omega^1) \neq 0$$

in general. □

Given a holomorphic 0-connection A_i^j, a *holomorphic 1-connection* over it is a 1-connection whose local 2-forms β are holomorphic. Again, unlike (2.2.4), these may not exist: this requires the vanishing of

$$\alpha_1 := \left[\frac{1}{2\pi i}\mathrm{d}A\right] \in \mathrm{H}^1(X;\Omega^2).$$

If the obstructions α_0 and α_1 vanish, then the gerb supports a curvature 3-form which is holomorphic, and so the Chern class $c(\mathcal{G})$ has a representative of type (3,0).

Under the obvious notions of holomorphic 0- and 1-gauge transformation, the space of holomorphic equivalence classes of gerbs with connection is

$$\mathbb{H}^2(X;\ \mathcal{O}^* \xrightarrow{\mathrm{d}\log} \Omega^1 \xrightarrow{\mathrm{d}} \Omega^2).$$

We remark merely that in this holomorphic case it is no longer true that any two 0-connections are related by a 1-transform, since in general

$$\mathrm{Z}^1(X;\Omega^1) \quad \neq \quad \delta\mathrm{C}^0(X;\Omega^1).$$

All this tightens up given access to the Hodge decomposition [21, 50].

Compact kähler manifolds

Theorem 5.2.2 *Let $\mathcal{G}$ be a holomorphic gerb on a compact kähler manifold X, with Chern class $c(\mathcal{G})$. Then*

$$\begin{aligned} c^{0,3} = \overline{c^{3,0}} &= 0 \\ c^{1,2} = \overline{c^{2,1}} &= \alpha_0. \end{aligned}$$

Further, α_1 is zero when it is defined. That is, the Chern class is essentially the obstruction to the existence of a holomorphic 0-connection; and given a holomorphic 0-connection, there always exists a holomorphic 1-connection.

Proof—Aside from the reality constraint

$$c^{p,q} \quad = \quad \overline{c^{q,p}}$$

the proof amounts to understanding how to move from Čech cohomology $\mathrm{H}^3(X;\mathbb{C})$ to its Hodge components $\mathrm{H}^{p,q}$. This can be discussed in the language of the spectral sequence of the double complex $\mathrm{C}^q(X;\Omega^p)$; or equivalently as follows.

The short exact sequences

$$\begin{cases} 0 \to \mathbb{C} \hookrightarrow \mathcal{O} \xrightarrow{\partial} \partial\mathcal{O} \to 0 \\ 0 \to \partial\Omega^{r-1} \hookrightarrow \Omega^r \xrightarrow{\partial} \partial\Omega^{r+1} \to 0 \end{cases}$$

on a complex manifold (for $r \geq 1$) give rise to long exact sequences

$$\begin{cases} * & \to & \mathrm{H}^2(\partial\mathcal{O}) & \to & \mathrm{H}^3(\mathbb{C}) & \to & \mathrm{H}^3(\mathcal{O}) & \to & * \\ * & \to & \mathrm{H}^1(\partial\Omega^1) & \to & \mathrm{H}^2(\partial\mathcal{O}) & \to & \mathrm{H}^2(\Omega^1) & \to & * \\ * & \to & \mathrm{H}^0(\partial\Omega^2) & \to & \mathrm{H}^1(\partial\Omega^1) & \to & \mathrm{H}^1(\Omega^2) & \to & * \\ & & 0 & \to & \mathrm{H}^0(\partial\Omega^2) & \to & \mathrm{H}^0(\Omega^3) & \to & * \end{cases}$$

for which restricting to a compact kähler manifold ensures that all maps to and from groups $*$ vanish. The reason for this is that such maps have the form

$$\mathrm{H}^q(\Omega^p) \xrightarrow{\partial} \mathrm{H}^q(\partial\Omega^p)$$

but if for instance we define these groups by resolutions

$$0 \to \Omega^p \hookrightarrow \mathcal{A}^{p,0} \xrightarrow{\overline{\partial}} \mathcal{A}^{p,1} \xrightarrow{\overline{\partial}} \dots$$

and take the harmonic representatives for classes in

$$\mathrm{H}^q(\Omega^p) := \mathrm{H}^{p,q}_{\overline{\partial}}$$

then the map to $\mathrm{H}^q(\partial\Omega^p)$ must vanish. (Alternatively, if we wish to define $\mathrm{H}^q(\Omega^p)$ by Čech cohomology, we can show that there exist (non-unique) Čech representatives which are ∂-closed on the intersections; and that taking ∂ of a representative induces the map on cohomology.)

Given these sequences, we have sufficient grip on the components of $\mathrm{H}^3(\mathbb{C})$ to demonstrate the theorem. Suppose given a class in $\mathrm{H}^3(\mathbb{C})$. Its (0,3)-component is its image under the above map to $\mathrm{H}^3(\mathcal{O})$, as is clear (given X compact kähler) from the proof of (5.1.1). Suppose this component vanishes: then the class pulls back to $\mathrm{H}^2(\partial\mathcal{O})$ and thereby projects to $\mathrm{H}^2(\Omega^1)$. This is its (1,2)-component, since (by converting to Dolbeault) it has a map to $\mathrm{H}^3(\mathbb{C})$ which is a (one-sided) inverse. And so on: if the class has no (0,3)- and (1,2)-terms, it lies in $\mathrm{H}^1(\partial\Omega^1)$, whose map to $\mathrm{H}^1(\Omega^2)$ is its (2,1)-component; and similarly for the final sequence.

Now to the theorem proper. It is already known that

$$\begin{aligned} c^{0,3} &= \overline{c^{3,0}} = 0 \\ c^{1,2} &= \overline{c^{2,1}} \end{aligned}$$

by (5.1.1) and reality. Given some Čech cocycle

$$t \in \mathrm{Z}^2(X; \mathcal{O}^*)$$

representing θ, the gerb $\mathcal{G}$ is described in $\mathrm{H}^3(\mathbb{C})$ by any choice of

$$\frac{1}{2\pi i}\delta \log t \in \mathrm{Z}^3(X; \mathbb{C})$$

(which also represents its explicitly-trivial $\mathrm{H}^3(\mathcal{O})$-component). Restricting to $\mathrm{H}^2(\partial\mathcal{O})$ means to consider

$$\frac{1}{2\pi i}\partial \log t \in \mathrm{Z}^2(X; \partial\mathcal{O})$$

which, viewed as an element of $\mathrm{Z}^2(\Omega^1)$, represents both $c^{1,2}$ and α_0 by (5.2.1).

It remains to consider α_1. But this is defined only when $\alpha_0 = 0$. In the same way that

$$\alpha_0 \quad = \quad c^{1,2}$$

it turns out that

$$\alpha_1 \quad = \quad c^{2,1}$$

which necessarily already vanishes. □

5.3 Compatible connections and hermitian gerbs

Compatibility with the holomorphic structure

Definition 5.3.1 (Compatible 0- and 1-connection) *Given a holomorphic gerb $\mathcal{G}$, a smooth 0-connection is compatible with the holomorphic structure if each line-bundle connection is compatible, ie. holomorphic sections have connection 1-form of type (1,0)*

$$\nabla_i^j : \Lambda_i^j \to \Lambda_i^j \otimes \mathcal{A}^{1,0}.$$

Given a compatible 0-connection, a smooth 1-connection is compatible if

$$\beta_i^{0,2} = 0 \quad (\forall i).$$

Proposition 5.3.2 *Such connections always exist.*

Proof—Suitably trivialised, a general smooth 0-connection requires

$$\alpha \in \mathrm{C}^1(\mathcal{A}^1)$$

such that

$$\delta\alpha = -\mathrm{d}\log t \in \mathrm{Z}^2(\mathcal{A}^{1,0}).$$

Given α, take its (1,0) component, which suffices.

Similarly, we then seek 2-forms β_i such that

$$\delta\beta = F_i^j$$

but the right-hand side vanishes in $\mathrm{H}^1(\mathcal{A}^2)$ and has no (0,2) part, and the same trick applies. □

(Thus for a holomorphic gerb there always exist connections for which the (0,3)-component of curvature vanishes.)

Since we are dealing with holomorphic gerbs, a 0-transform is described by holomorphic isomorphisms of the Λ_i^j, as in (2.3.5). We obtain the expected equivalence classes

$$\mathbb{H}^2(X;\ \mathcal{O}^* \xrightarrow{\mathrm{d}\log} \mathcal{A}^{1,0} \xrightarrow{\mathrm{d}} \mathcal{A}^{2,0+1,1})$$

by defining 1-transforms to be given by

$$\gamma_i \in \mathrm{C}^0(X; \mathcal{A}^{1,0})$$

which is different from both the smooth and the holomorphic cases.

Hermitian gerbs

Definition 5.3.3 (Hermitian gerb) *A hermitian structure on a smooth gerb is a hermitian metric on each line-bundle Λ such that $|\theta| = 1$.*

A 0-connection is compatible with the metric, if each ∇_i^j is compatible in the line-bundle sense (ie. if a norm-1 section has imaginary connection 1-form).

A compatible 1-connection is one such that each β_i is imaginary.

All these structures exist: given arbitrary line-bundle metrics, the failure of θ to be unitary is a 2-cocycle of the sheaf of smooth positive real functions, which is soft. Given any smooth 0-connection on a hermitian gerb, and unitary sections of the Λ, the failure to be compatible is a 1-cocycle of smooth real 1-forms, which again is necessarily a coboundary. Rectifying this, the F_i^j are imaginary, and then the imaginary parts of any smooth β_i form a compatible 1-connection.

Hermitian gerbs are classified by $\mathrm{H}^2(X; \underline{S}^1)$; and we must define 0-transforms to be collections of isometries. The natural extension of

$$\underline{S}^1 \xrightarrow{\mathrm{d}\log} \mathcal{I}^1$$

as an exact sequence (in which $\mathcal{I}^p$ is the sheaf of smooth purely imaginary p-forms) is

$$\underline{S}^1 \xrightarrow{\mathrm{d}\log} \mathcal{I}^1 \xrightarrow{\mathrm{d}} \mathcal{I}^2.$$

So the appropriate notion of 1-transform is a smooth 1-transform in which the local 1-forms γ_i are purely imaginary; and the set of equivalences of hermitian gerbs with connection is

$$\mathbb{H}^2(X;\ \underline{S}^1 \xrightarrow{\mathrm{d}\log} \mathcal{I}^1 \xrightarrow{\mathrm{d}} \mathcal{I}^2).$$

Proposition 5.3.4 *A holomorphic hermitian gerb has a unique 0-connection compatible with both structures.*

Proof—Each ∇_i^j is uniquely specified by the analogous line-bundle result [21]. □

Note that it is similarly clear that hermitian structures exist on holomorphic gerbs. Further, a 1-connection β compatible with both structures must be purely of type (1,1), thus proving (5.1.1) for a second time by cruder but more transparent means. Given one such 1-connection, any other differs by an arbitrary smooth global imaginary (1,1)-form.

(It is not clear whether the failure of such a 1-connection to be unique means our definition of compatibility is incomplete. Perhaps our original definition of a general 1-connection is inadequate, just as a definition of line-bundle connection involving merely throwing down some local 1-forms would miss the true geometric issue. But the categorical definition of a gerbe [5] fares no better on this score; and we can hardly expect every line-bundle theorem to translate unchanged into one on gerbs.)

5.4 Objects for a holomorphic gerb

A holomorphic object—the definition mimics (3.1.1)—exists iff the gerb is holomorphically trivial

$$[\mathcal{G}] = 0 \in \mathrm{H}^2(X; \mathcal{O}^*).$$

Holomorphic objects always exist locally; and the example preceding (3.1.3) applies as in the smooth case.

The existence of an object ensures the existence of a holomorphic gerb connection (since $[\mathcal{G}] = 0$ ensures $c(\mathcal{G}) = 0$). Locally any object carries a holomorphic object connection: this is because any object locally looks like the example referenced above, and so is dictated by a line-bundle $\mathcal{L}$. Whilst $\mathcal{L}$ need not be topologically trivial, we can restrict further to enforce this—and upon choosing a holomorphic connection $(\mathcal{L}, \nabla)$ we obtain a holomorphic object with connection as in the smooth example following (3.2.5).

As this implies, there need not exist a holomorphic object connection defined over the *whole* of the object—given holomorphic 1-forms α_i^j and functions m_i^j describing the line-bundle maps, there is an obstruction

$$[\alpha + \mathrm{d}\log m] \in \mathrm{H}^1(\Omega^1)$$

which does not vanish in general.

Example: the Riemann sphere

Consider an example, as an antidote to these layerings of definitions: the Riemann sphere $\mathbb{P}^1$. All structures will be assumed holomorphic.

Any gerb is trivial, and carries holomorphic connections, since by the exponential sequence $\mathrm{H}^2(\mathbb{P}^1; \mathcal{O}^*) = 0$. Consider first a single open set

$$U := \mathbb{P}^1$$

so that there is no data to specify for any gerb trivialisable on this cover, nor indeed for a 0-connection. A 1-connection is a holomorphic 2-form over U, which is necessarily zero. Any object for this gerb is a line-bundle $\mathcal{O}(j)$; and there is an object connection over all of $\mathbb{P}^1$ iff $j = 0$.

Example: same space, different cover

In the first example, at least there exists *some* object that can support a holomorphic object connection. This is not true in general: the reason is that there is nothing to force the obstruction to be integral in $\mathrm{H}^1(\Omega^1)$. If it were, we could take a second object differing from the first by a global bundle whose Chern class cancelled the obstruction (3.1.3).

As a second simple illustration, consider again $\mathbb{P}^1$ but now with gerbs trivialisable over a two-set cover

$$\begin{cases} U := & \mathbb{P}^1 - p_0 \\ V := & \mathbb{P}^1 - p_1 \\ W := & U \cap V = \mathbb{C}^* \end{cases}$$

so a gerb is specified by a line-bundle Λ over W. Since $\mathrm{H}^1(W; \mathcal{O}^*) = 0$, we can take a trivialisation of Λ. A 0-connection is then a 1-form α over W, which let us say is of type

$$\alpha = c \cdot \mathrm{d}\log z$$

for some complex constant c. An object is given by bundles on U and V, which we trivialise, and a map which (given the three bases) amounts to a general element

$$m \in \mathrm{H}^0(W; \mathcal{O}^*).$$

An object connection exists if and only if

$$[\alpha + \mathrm{d}\log m] = 0 \in \mathrm{H}^1(\mathbb{P}^1; \Omega^1)$$

where the right-hand side is spanned by $\mathrm{d}\log z$. But whilst $[\mathrm{d}\log m]$ is always integral in $\mathrm{H}^2(\mathbb{P}^1; \mathbb{C})$, $[\alpha]$ is integral only when $c \in \mathbb{Z}$. So in general there is no object that can carry an object connection.

Compatibility

Given a holomorphic gerb and object, and a 0-connection compatible with the holomorphic structure, a smooth object connection is *compatible* if each bundle connection ∇_i is compatible with L_i. Such object connections always exist, thanks to the obstruction being

$$[\alpha + \mathrm{d}\log m] \in \mathrm{H}^1(\mathcal{A}^1) = 0.$$

A *hermitian object* (for a hermitian gerb) is an object with metrics on each L_i, such that the maps $m : L \to \Lambda \otimes L$ are isometries. These are obstructed by $\mathrm{H}^1(\underline{\mathbb{R}})$, which vanishes.

If $\mathcal{G}$ is holomorphic and hermitian (and $[\mathcal{G}] = 0$ of course), then there exist objects supporting both structures for the same reason. Such an object is not unique: it can be tensored by any global holomorphic hermitian line-bundle.

Chapter 6

j-Equivalence of codimension-three submanifolds

In this chapter, we expand on the example (2.3.6) by considering not points in 3-manifolds but a real codimension-3 submanifold in X^n, with the extra structure of supposing that X comes with a fixed metric. In this context, a useful way to move between submanifolds R^{n-3} and differential forms (describing connections) is through the classical language of currents. We use nothing but the basics of this theory [10, 21].

We shall write down very general equations which describe a wide class of 1-connections on these gerbs. Then we find that considering the possibility that such a 1-connection is in fact *objective* (3.3.1) starts to lead somewhere interesting.

The underlying idea is this: given a point in a 3-manifold, it is natural in the context of gerbs to want to put a degree-1 line bundle on a punctured neighbourhood of the point (2.3.6). It is then perhaps unsurprising that such gerbs are related to the equation for a (singular, abelian) *monopole* on that neighbourhood (given a metric on the space). In more general codimension-3 contexts, such monopoles keep appearing: and objective 1-connections lead us there (6.4.2).

As ever, all structures are oriented unless specified otherwise.

(We shall always try to write *real* dimensions in this chapter. Note also that we shall now start using the notation of currents, with the convention that the degree of a submanifold as a current—written as a superscript—is its real codimension, so the submanifold R^{n-3} is also the current R^3. In compensation, a differential form α^r is also a current α^r. Forms are always held to be smooth, but currents may be smooth or singular.)

6.1 Motivation from line-bundles

The greatest interest in line-bundles and codimension-2 submanifolds is of course in the holomorphic version. In this chapter our gerbs will be merely Riemannian, yet we quickly show that this approach can carry a good amount of the data of a holomorphic line-bundle.

Holomorphic bundles on kähler manifolds

Suppose given a complex codimension-1 submanifold $Y^{n-2} \subset X^n$ on a compact kähler manifold.

Y determines a holomorphic line-bundle: if we have Y described by local holomorphic functions f_i, then we have an explicit description of the bundle by transitions. Take then any global meromorphic section s of the bundle, and any hermitian metric on the bundle.

Then it is standard [21] that, as a global current, s satisfies the Poincaré-Lelong equation

$$\frac{i}{2\pi}\bar{\partial}\partial S = c - Y$$

where

$$S := \log \|s\|^2$$

$$c := \frac{i}{2\pi}F$$

and F as ever is the bundle curvature.

The curvature is displayed explicitly as follows, with respect to the bases of the local trivialisation. Off Y we have the connection 1-form

$$A_0 := \ \partial S$$

and on U_i we take

$$A_i := \ \partial S - \partial \log |f_i|^2$$

all obeying $F = \mathrm{d}A_i$ $(\forall i)$.

Smooth version

Start again, on a compact Riemannian manifold X^n with Laplacian operator Δ, and fix a real codimension-2 submanifold Y^{n-2}. (The 2 is not significant and any proper codimension would work.) Pick a global smooth closed 2-form c in the same (current-) cohomology class as Y. Then there exists [10] a global 2-current u solving

$$\frac{i}{2\pi}\Delta u = c - Y.$$

(In particular, u is a smooth 2-form off Y and is unique up to any global harmonic 2-form.)

On small enough open sets U_i, we can also locally solve

$$\frac{i}{2\pi}\Delta v_i = Y$$

such that the v_i are closed currents (6.3.1). Do this for an open cover $\{U_i\}_{i\in I}$ of Y, and write U_0 for the complement of Y.

We can now write down a global 1-current

$$a_0 := \mathrm{d}^* u$$

which is smooth off Y, and smooth local 1-forms

$$a_i := \mathrm{d}^* u + \mathrm{d}^* v_i$$

so that around any point we have some 1-form obeying

$$\frac{i}{2\pi}\mathrm{d}a = c.$$

From smooth to holomorphic

If we now carry this out for our divisor on kähler X, and writing Λ for the adjoint operator of right-multiplication by the kähler form, we find that this purely real construction almost gives us the full line-bundle. Precisely, we obtain

$$\begin{aligned} S :=& 2i\Lambda u \quad \text{globally} \\ |f_i| :=& \exp(-i\Lambda v_i) \quad \text{on } U_i \end{aligned}$$

in the sense that these S and $|f_i|$ behave the same as those defined from the line-bundle, ie. they satisfy Poincaré-Lelong and define the appropriate connection 1-forms.

One thing the real approach lacks is that we need to know *in advance* that Y is holomorphic. (If we had started with local holomorphic equations $f_i = 0$, this would have been apparent.) We need the complex structure of X to get from $|f|$ to f. We shall keep playing with the real case even so in this chapter, since for gerbs we do not already know what the answer will be and it is useful to see how far the smooth theory goes.

6.2 Codimension-three submanifolds

Given smooth compact manifolds $R^{n-3} \subset X^n$, we choose an open cover of X of the form

$$\begin{aligned} U_0 &:= X \backslash R \\ R &\subset \bigcup_{i \neq 0} U_i \end{aligned}$$

on which for convenience we further impose that the U_i be suitable to trivialise the normal bundle N of R and that all $U_{i,j,\ldots k}$ for $i, j, \ldots \neq 0$ be contractible. (When we write $i, j, \ldots$, we shall now assume them all non-zero.) We implicitly view N embedded in X as a tubular neighbourhood of R [35].

Now double intersections are either contractible (the $U_{i,j}$) or homotopic to the 2-sphere ($U_{0,i}$). Over the former spaces, put trivial line-bundles; over the latter, put degree-1 Λ_i^0. We seek to extend the construction of (2.3.6) to obtain a gerb up to equivalence on X determined by R.

Proposition 6.2.1 *There exists a gerb $\mathcal{G}(R)$ on this (I, Λ).*

Proof—We need a cocycle θ of sections of the bundles on three-fold intersections $U_{i,j,k}$ or $U_{0,i,j}$.

Choose arbitrary bases, say e_i^j for Λ_i^j and v_i^j for $\Lambda_i^0 \otimes \Lambda_0^j$. Then any nowhere-zero section θ on $\bigotimes_1^3 \Lambda$ is given by scalars s and t, where

$$\begin{cases} \theta_{i,j,0} := & s_i^j \cdot (e_i^j \otimes v_j^i) \quad : \ s_i^j \in \mathrm{C}^1(N \backslash R; \underline{\mathbb{C}}^*) \\ \theta_{i,j,k} := & t_{i,j,k} \cdot (e_i^j \otimes e_j^k \otimes e_k^i) \quad : \ t_{i,j,k} \in \mathrm{C}^2(N; \underline{\mathbb{C}}^*) \end{cases}$$

(with respect to covers $\{U_i\}_{i \neq 0}$ for N and $\{U_i \cap U_0\}_{i \neq 0}$ for $N \backslash R$.) It remains to show that there exist s and t such that $\delta\theta = 1$, the canonical section of $\delta^2 \Lambda$.

On four-fold intersections of the form $U_{i,j,k,l}$ this says simply that t must be a cocycle, since $\delta^2 e$ is necessarily the canonical section.

Secondly and finally, on intersections $U_{i,j,k,0}$ we need

$$\begin{aligned} 1 &= \delta s \cdot \delta e \otimes \delta v^* \otimes \theta^*_{i,j,k} \\ \Rightarrow \quad t_{i,j,k} &= (\delta s \cdot \delta v^*)_{i,j,k} \end{aligned}$$

where we freely drop uninteresting canonical sections (such as $e \otimes e^*$) from our notation. To interpret this constraint, note that even though (so far) the sections v_i^j and the scalars s_i^j have been arbitrary, nevertheless all choices give the same cohomology class

$$\sigma := [\delta v^*] \equiv [\delta(s \cdot v^*)] \in \mathrm{H}^2(N \backslash R; \underline{\mathbb{C}}^*)$$

in which by δv^* we mean the coefficient of the section δv^* with respect to the canonical section of the bundle

$$(\Lambda_i^0 \Lambda_0^j) \otimes (\Lambda_j^0 \Lambda_0^k) \otimes (\Lambda_k^0 \Lambda_0^i).$$

The freedom we have is to vary the class $[t_{i,j,k}] \in \mathrm{H}^2(N; \underline{\mathbb{C}}^*)$, and our need is that there be a choice of t such that the inclusion map

$$\mathrm{H}^2(N; \underline{\mathbb{C}}^*) \xrightarrow{i^*} \mathrm{H}^2(N \backslash R; \underline{\mathbb{C}}^*)$$

induces

$$i^* : t \mapsto \sigma.$$

To demonstrate this, it is convenient to make two changes. First, take choices of logarithm for the representatives $t|_S$ and $\delta v^*|_S$ so that we view our classes as lying in $\mathrm{H}^3(\mathbb{Z})$ rather than in $\mathrm{H}^2(\underline{\mathbb{C}}^*)$. (The choices do not affect the result.)

Second, pick a metric for the fibre bundle N (for instance, one induced from a Riemannian metric on X); and thus write B for the unit-ball sub-bundle of N, and S for the unit-sphere sub-bundle. Perform excision of the set $N \backslash B$ from the pair $(N, N \backslash R)$.

Now the pair sequence $S \to B \to (B, S)$ induces a long exact sequence [36]

$$\longrightarrow \mathrm{H}^3(B; \mathbb{Z}) \xrightarrow{i^*} \mathrm{H}^3(S; \mathbb{Z}) \longrightarrow \mathrm{H}^4(B, S; \mathbb{Z}) \longrightarrow$$

with $[t|_S]$ lying in the left term and σ in the middle term. We need to show that σ lies in the image of $\mathrm{H}^3(B; \mathbb{Z})$; and that the representative t is adequate to express a class whose image is σ.

The second is true because the map on cohomology is just restriction, which is essentially all we have done to the representive $t_{i,j,k}$.

The first is true because we have representatives for σ, such as $\frac{1}{2\pi i}\delta \log \delta v^*$, which are explicitly constant on each fibre not just of S but indeed of B, and thus provide representatives of a preimage.

We have proved (after taking a series of choices) that a cocycle section θ exists, and thus that we have a gerb. □

It is clear by considering the choices made in the proof that our gerb is fixed only up to equivalence, just as in (2.3.6).

Theorem 6.2.2 *The Chern class* $c(\mathcal{G}) \in \mathrm{H}^3(X;\mathbb{C})$ *of this gerb is Poincaré-dual to the homology class of* R.

Proof—Pick a 0-connection for $\mathcal{G}$ over the given cover (2.2.2), with bundle curvature 2-forms f_i^j. Pick a subordinate partition of unity ρ, giving a 1-connection

$$b_i \quad := \quad \sum_{k \neq i} \rho_k \cdot f_i^k$$

with curvature 3-form

$$o = \mathrm{d}b_0 = \sum_{k \neq 0} \mathrm{d}\rho_k \wedge f_0^k$$

as viewed in U_0. Note that $o \equiv 0$ off U_0, and indeed off the unit-ball bundle B if we choose ρ appropriately; and that on any fibre S^2 of the sphere bundle intersecting U_k, $\frac{i}{2\pi} f_0^k$ has degree -1.

Now we turn to $[R]$. Its dual in $\mathrm{H}^3_{\mathrm{dR}}(X)$ is represented by a Thom form T on the normal bundle, extended by zero to X [35]. To describe such a form, take a fibre-volume 2-form V on the sphere bundle S and extend it by pull-back to all of the punctured normal bundle $N \backslash R$. Note that V has degree 1 on any fibre S^2. Now we can take

$$T \quad := \quad \mathrm{d}(1 - \rho_0) \wedge V$$

as a representative for the Thom class. (This supposes that $\rho_0 = 1$ on S, which can certainly be arranged.)

It remains to show that

$$\left[\frac{i}{2\pi} o\right] \quad = \quad [T]$$

but

$$\frac{i}{2\pi} o - T \quad = \quad \sum_{i \neq 0} \mathrm{d}\rho_i \wedge \left(V + \frac{i}{2\pi} f_0^1\right)$$

Since $V + \frac{i}{2\pi} f_0^1$ is trivial in de Rham cohomology on $U_{i,0}$ (because it is of degree zero), there exist 1-forms γ_i on $U_{i,0}$ such that

$$\begin{aligned} V + \frac{i}{2\pi} f_0^1 \quad &= \quad \mathrm{d}\gamma_i \\ \Rightarrow \quad \frac{i}{2\pi} o - T \quad &= \quad \mathrm{d} \sum_{i \neq 0} \mathrm{d}\rho_i \wedge \gamma_i \end{aligned}$$

where the right-hand side is the exterior derivative of a *global* smooth 2-form. □

6.3 Gerb connections via the Laplacian

1-connections in codimension 3

We construct a somewhat more geometric 1-connection for the gerb $\mathcal{G}(R)$ of (6.2). Fix a Riemannian metric on X. Write $\Delta := \mathrm{dd}^* + \mathrm{d}^*\mathrm{d}$ for the Laplace operator. Pick any target for the Chern 3-form

$$\omega \in c(R) \in \mathrm{H}^3_{\mathrm{dR}}(X;\mathbb{Z})$$

(meaning the image of integral cohomology in de Rham cohomology)—for instance, the unique harmonic representative.

We can certainly globally solve the equation

$$\frac{i}{2\pi}\Delta\gamma = \omega - R$$

for a 3-current γ, since ω and R share a cohomology class. Further, γ is automatically a smooth closed 3-form on U_0 (though singular on R) and is unique up to any global harmonic 3-form. (We could specify γ precisely by insisting it be globally exact.)

Take a cover U_0, $U_{i\neq 0}$ as in (6.2) for which we further insist the U_i be small enough that we can solve

$$\frac{i}{2\pi}\Delta\sigma_i = R$$

as a current equation for smooth 3-forms σ, but requiring that σ be closed. (Say $\sigma_0 := 0$ on U_0.)

Proposition 6.3.1 *This is always possible locally.*

Proof—We apply a series of standard results, taken for instance from de Rham [10], chapters 30 and 32.

Certainly local solutions exist. On their being closed: we need to know things about the square-summability of the currents involved. For this, it is more transparent (due to smoothness) to consider local solutions to

$$\frac{i}{2\pi}\Delta\tau_i = \omega$$

for which we seek to enforce that $\mathrm{d}\tau = 0$, and then to take $\sigma_i := \tau_i - \gamma$.

If U is small, then $\omega = \mathrm{d}\mu$ is exact, with μ smooth. Shrinking U if necessary, μ is square-summable. So it decomposes into a sum of three unique square-summable currents: closed and coclosed; exact but orthogonal to the first class; and coexact but orthogonal to the first class. (This is Theorem 24 of de Rham, due to Kodaira [29].) Since μ is smooth, its three components are smooth. Without loss of generality, μ is of the third type only.

So $\mu = \mathrm{d}^*\nu$, again with ν smooth. And we can again shrink U slightly away from its boundary to be sure that ν is also square-summable. It decomposes as above; take ν to be its (smooth) exact component.

Now $\mathrm{d}\nu = 0$, so $\Delta\nu = \omega$ and we can solve for τ. □

Given these choices we define local smooth 2-forms over U_i (including $i = 0$)

$$\beta_i := \mathrm{d}^*(\gamma + \sigma_i)$$

which all satisfy

$$d\beta_i = \Omega \quad (\forall i)$$

where by Ω we mean $-2\pi i\omega$ (2.2.8). (We would like a 1-connection with these β as our 2-forms). Note in particular that β_0 extends as a global current, smooth off R.

Choosing line-bundle curvatures

On each $U_{i,j}$ we see that

$$\begin{aligned} 0 &= \Delta(\sigma_i - \sigma_j) \\ &= dd^*(\sigma_i - \sigma_j) \end{aligned}$$

so we define closed (and coclosed) 2-forms

$$F_i^j := d^*(\sigma_i - \sigma_j)$$

which obey

$$\beta_i - \beta_j = F_i^j.$$

We want these closed 2-forms to be the curvatures of line-bundles Λ_i^j as described in (6.2). We must show that they are integral of the appropriate degree.

Proof—For $i, j \neq 0$ there is nothing to prove. Our cover is such that $U_{i,j}$ is contractible; so certainly F_i^j has trivial de Rham class as desired.

The $U_{i,0}$ are homotopic to 2-spheres, and in particular their second homologies are each generated by a fibre S^2 of the sphere bundle of the normal bundle of R^{n-3}. Say B^3 is the ball it contains. And whilst F_i^0 is a smooth closed 2-form off R, the definition of σ_i forces it to extend as a unique current from $U_{i,0}$ to all of U_i such that

$$\frac{i}{2\pi} dF_i^0 = R$$

which thus behaves as

$$\begin{aligned} \int_{S^2} \frac{i}{2\pi} F_i^0 &= \int_{B^3} \frac{i}{2\pi} dF_i^0 \\ &= B \wedge R \\ &= 1 \end{aligned}$$

just as needed for a degree-1 bundle. □

The existence of a 0-connection

Theorem 6.3.2 *Pick a topological gerb $\mathcal{G}(R)$ in the manner of (6.2). Then there exists a 0-connection such that the curvatures of its line-bundle connections ∇_i^j of Λ_i^j are the harmonic 2-forms F_i^j.*

Corollary 6.3.3 *The 2-forms β_i give a 1-connection for $\mathcal{G}(R)$, with curvature Ω.* □

Proof—of the theorem—The proof of (6.2.2) provides 0- and 1-connections ∇, f, b, o. We shall adapt this into the desired F, β, Ω.

We know by (6.2.2) and by the definition of Ω that there exists a global 2-form B such that

$$\Omega = o + \mathrm{d}B.$$

In particular, $b_i + B$ is still a 1-connection for the 0-connection ∇_i^j.

Now there exist local 1-forms ψ_i for $i \neq 0$ such that

$$\beta_i - b_i - B = \mathrm{d}\psi_i$$

since the left hand side is closed and U_i is contractible. So (viewing ψ_i as a 1-transform) β_i is a 1-connection for $\nabla + \delta\psi$. This means that

$$\delta\beta = \mathrm{Curvature}(\nabla + \delta\psi)$$

but by construction

$$\delta\beta_i = F_i^j$$

so the new 0-connection $\nabla + \delta\psi$ has bundle curvatures F_i^j as required.

It remains to consider the case $i = 0$, ie. to show that for some choice of B

$$[\beta_0 - b_0 - B] = 0 \in \mathrm{H}^2_{\mathrm{dR}}(X - R)$$

noting that we are still free to vary B by any closed 2-form defined over all of X. We shall demonstrate this separately for the two cases in which $R^{n-3} \subset X^n$ is or is not a boundary.

Suppose first that

$$[R] \neq 0 \in \mathrm{H}_{n-3}(X)$$

so that in the long exact homology sequence of the pair (X, R), the map

$$\mathrm{H}_{n-2}(X, R) \to \mathrm{H}_{n-3}(R)$$

vanishes (ie. R is not a boundary in X). This leaves

$$0 = \mathrm{H}_{n-2}(R) \to \mathrm{H}_{n-2}(X) \to \mathrm{H}_{n-2}(X, R) \to 0$$

which (with real or complex coefficients) is canonically isomorphic to

$$0 \to \mathrm{H}^2_{\mathrm{dR}}(X) \to \mathrm{H}^2_{\mathrm{dR}}(X - R) \to 0$$

so that for some closed globally-defined B'

$$[\beta_0 - b_0 - B] = [B'] \in \mathrm{H}^2_{\mathrm{dR}}(X - R)$$

and modifying $B \mapsto B + B'$ does the job.

If finally $[R] = 0$, we no longer have the isomorphism above. Instead there are global forms B_Ω and B_o such that

$$\begin{aligned} & \Omega = \mathrm{d}B_\Omega\ , \quad o = \mathrm{d}B_o \\ \Rightarrow \quad & [\beta_0 - B_\Omega] = [b_0 - B_o] \in \mathrm{H}^2_{\mathrm{dR}}(X - R) \end{aligned}$$

and taking $B := B_\Omega - B_o$ is enough. □

Variations

Given R and a trivialisation U_i for N, the freedom to vary choices of Λ, e and v amounts to (a restricted set of) 0-transforms, or gerb equivalences.

Given also a 0-connection with line-bundle curvatures F, these 0-transforms act as expected on the 0-connection. Fixing the gerb, the 0-connection can vary only by a closed 1-transform $\gamma \in \mathrm{C}^1(X; \mathrm{d}\mathcal{A}^0)$, which does not affect β or F.

Given ω, β_0 is unique whilst the other β_i are unique up to d(1-form). This same 1-form must then act as a 1-transform on the 0-connection in order that the new β_i remain a 1-connection.

Changing ω—necessarily by a global exact form—takes us to an inequivalent 1-connection.

So if we *restrict to harmonic* ω (which is uniquely specified by R) then our gerb with 1-connection on X is fixed by R up to equivalence. (It is fixed more firmly than that, since by construction it has certain degree-1 singularities along R, but this will do for now.)

6.4 j-Equivalence of different submanifolds

An obvious question is to consider variation in R itself: how does this affect the gerb? So consider the case $[R] = 0$ (or, if we prefer, $[R_1] = [R_2]$ and then $R := R_1 - R_2$ as a singular-homology cycle).

The underlying topological gerb is clearly not interesting since $[R] \equiv [\mathcal{G}]$. So consider connections on $\mathcal{G}$.

One approach that springs naturally from chapter 3 is to create an *object* for the gerb $\mathcal{G}(R)$ now that it is trivial. In particular, consider the possibility that the β_i in fact represent an *objective* 1-connection (3.3.1), ie. the 1-connection 2-forms are the curvatures of an object 0-connection—or equivalently, the global 2-form ϵ vanishes (3.3.2). This requires the β_i to be closed and integral (to $\frac{i}{2\pi}$).

Closure is easy: we must take $\omega = 0$ (which has already been imposed, having restricted ourselves to harmonic $\omega \in [R]$).

Since the U_i are contractible, we know the β_i are integral for $i \neq 0$. What about U_0? There are several ways to describe the constraint on the integrality of $\beta_0|_{U_0}$.

Since R is now a boundary, pick some singular homology chain $\partial C^{n-2} = R$. It gives a relative class

$$[C] \in \mathrm{H}_{n-2}(X, R; \mathbb{Z}) \cong \mathrm{H}^2(X - R; \mathbb{Z}).$$

Both C and β_0 are global 2-currents, and as such decompose uniquely into exact, harmonic and coexact parts. We write these as subscripts.

Definition 6.4.1 (j-map) *We have a natural map given by integration over any choice of C*

$$\mathrm{j} : \{\textit{exact } (n-3)\textit{-cycles in } \mathbb{Z}\textit{-singular homology}\} \longrightarrow \textit{Jac}$$

where Jac is the quotient by $\mathrm{H}_{n-2}(X; \mathbb{Z})$ of the dual of the vector space of real harmonic forms

$$\mathcal{H}^{n-2}(X; \mathbb{R})^*.$$

This real torus has a well-defined base-point 0.

The name is meant to imply a link with intermediate Jacobians [21]; but ours is not such a torus even if X is compact kähler, since then $n-2$ would be even and intermediate Jacobians are defined in odd degree. On the other hand, there are analogous definitions of $\mathfrak{j}$ in all real codimensions R^{n-c} for $c = 1, \ldots, n$, and in codimension 2 we obtain

$$\frac{\mathcal{H}^{n-1}(X;\mathbb{R})^*}{\mathrm{H}_{n-1}(X;\mathbb{Z})}$$

which (if $n = 2N$ and X is compact kähler) is equal as a real torus to the first intermediate Jacobian, the Picard variety

$$\frac{\mathrm{H}^{0,1}(X)}{\mathrm{H}_{2N-1}(X;\mathbb{Z})}$$

of which more in (6.5). In the various even real codimensions, we form the various such intermediate Jacobians.

Theorem 6.4.2 *The following are equivalent—*

1. *Some 1-connection β of type (6.3) is objective;*
2. *On U_0, $\frac{i}{2\pi}\beta_0$ is integral*

$$\left[\frac{i}{2\pi}\beta_0|_{U_0}\right] \in \mathrm{H}^2_{\mathrm{dR}}(X - R;\mathbb{Z});$$

3. *In X, R^{n-3} is a "Dirac monopole" in the sense that there is a line bundle with connection and a "Higgs field" $(n-3)$-form ϕ, both defined off R, such that*

$$F = *\mathrm{d}\phi$$

—the Bogomolnyi equation, if $n = 3$—subject to a certain boundary constraint, below;

4. *Any choice of C has integral coexact part on U_0*

$$[C_{coexact}|_{U_0}] \in \mathrm{H}^2_{\mathrm{dR}}(X - R;\mathbb{Z});$$

5. *Any choice of C has integral harmonic part globally*

$$C_{harmonic} \in \mathcal{H}^2(X;\mathbb{Z});$$

6. *The $\mathfrak{j}$-image of R is trivial*

$$\mathfrak{j} : R \mapsto 0 \in Jac.$$

We consider the interesting conditions to be (1.), (3.) and (6.)—the others are given mainly to clarify the proof, which follows. For the monopole (3.), the appropriate behaviour at the singularity R is that ϕ extends as a global singular current satisfying

$$*\Delta\phi = 2\pi i R$$

(and consequently F extends globally such that

$$\mathrm{d}F = 2\pi i R).$$

Recall that in this section $[R] = 0$, so such equations have global solutions.

For reference we note [10] that in general on p-forms

$$\begin{cases} ** & \equiv & (-1)^{np+p} \\ \mathrm{d}^* & \equiv & (-1)^{np+n+1} * \mathrm{d} * \end{cases}$$

Proof—(1.⇔2.) Already discussed.

(2.⇒3.) We know

$$\begin{aligned} F_0 =& \beta_0 \\ =& \mathrm{d}^* \gamma \\ =& (-1)^{n+1} * \mathrm{d} * \gamma \end{aligned}$$

so define

$$\phi^{n-3} := (-1)^{n+1} * \gamma.$$

Since $\mathrm{d}\beta_0 = 2\pi i R$ and γ is closed, this has the appropriate behaviour at R.

(3.⇒2.) Define $\gamma := *\phi$, which then obeys

$$\Delta\gamma = 2\pi i R$$

and so $[\frac{i}{2\pi}\beta_0]$ is the Chern class of a line-bundle, and is integral.

(2.⇔4.) In general, β_0 is the (unique) global coexact current such that

$$\frac{i}{2\pi}\mathrm{d}\beta_0 = -R.$$

But for any chain such that $\partial C = R$, C_{coexact} is the (unique) global coexact current such that

$$\mathrm{d}C_{\text{coexact}} = (-1)^n \partial C_{\text{coexact}} = (-1)^n R$$

and so each is integral as a current on U_0 if and only if the other is.

(4.⇔5.) Note that

$$[C] \in \mathrm{H}_{n-2}(X, R; \mathbb{Z}) \cong \mathrm{H}^2(X - R; \mathbb{Z})$$

is exactly equivalent—in its evaluation on forms in $\mathrm{H}_{\mathrm{dR}}^{n-2}(X, R; \mathbb{Z})$—to the global current

$$C_{\text{coclosed}} = C_{\text{harmonic}} + C_{\text{coexact}}.$$

So C_{coclosed} is integral on such forms. Then (4.) holds if and only if C_{harmonic} is similarly integral. C_{harmonic} extends globally as a smooth form. If in fact (5.) holds and C_{harmonic} is integral *globally*—ie. evaluated on $\mathrm{H}_{\mathrm{dR}}^{n-2}(X; \mathbb{Z})$—then then it is so in particular on forms compactly supported in U_0. It remains to show the converse.

The long exact homology sequence of (X, R) is dual to

$$\cdots \to \mathrm{H}_{\mathrm{dR}}^{n-3}(R) \to \mathrm{H}_{\mathrm{dR}}^{n-2}(X, R) \xrightarrow{(\text{inclusion})^*} \mathrm{H}_{\mathrm{dR}}^{n-2}(X) \to \mathrm{H}_{\mathrm{dR}}^{n-2}(R) = 0.$$

In particular, any class in $\mathrm{H}^{n-2}_{\mathrm{dR}}(X;\mathbb{Z})$ has a compactly-supported representative (off R). But as a global harmonic form, C_{harmonic} has the same pairing with all representatives in a given class. So if its pairing is integral on any compactly-supported representative, it is so on any representative.

(5.⇔6.) This is tautologous. □

(Currents have a dual pairing on differential forms. So when discussing for example C_{harmonic} evaluated on $\mathrm{H}^{n-2}_{\mathrm{dR}}(X;\mathbb{Z})$, we mean its pairing with representative forms of classes in the image of say singular cohomology in de Rham cohomology

$$\mathrm{H}^{n-2}(X;\mathbb{Z}) \to \mathrm{H}^{n-2}_{\mathrm{dR}}(X).$$

Implicit in this is independence from choice of representative.)

Definition 6.4.3 (j-Equivalence) *Two (closed) topologically-equivalent submanifolds are* j*-equivalent, if their difference maps to zero*

$$\mathfrak{j} : R_1 - R_2 \mapsto 0.$$

Considering the case $n = 3$ of points in a 3-manifold (2.3.6), we see that there is nothing to show if X is simply-connected (since there are no global harmonic 1-forms and j is trivial). Otherwise it is reasonable to ask whether or not two points are related by an abelian monopole interpolating between them. We can view this as a finer equivalence than the merely topological, as a smooth analogue of linear equivalence of divisors.

6.5 "j-Equivalence" in codimension 2

We make explicit the analogy between j-equivalence and linear equivalence of divisors, by expanding on section (6.1). To do this with purely real machinery requires that we restrict to the case for which the metric carries all we need to know about the complex structure: a Riemann surface, $n = 2$.

Riemann surfaces

Take a topologically-trivial collection of points $[Y] = 0 \in \mathrm{H}_0(X;\mathbb{Z})$. As before, we can solve globally

$$\frac{i}{2\pi}\Delta u = -Y$$

thus giving global currents

$$b := \mathrm{d}^* u$$
$$a := *b.$$

From these we have a complex global current

$$A := a - ib$$

which is a smooth 1-form off Y. Further, $\mathrm{d}A = 0$ off Y; and the real part a of A is an exact smooth form off Y

$$a = +\mathrm{d}(*u)$$

and so has *zero periods.* And in fact A is a (1,0)-form, since given $n = 2$ and $p = 1$ the complex structure obeys

$$\begin{aligned} \mathcal{I} &\equiv -* \\ \Rightarrow \quad \mathcal{I}A &= -*a + i*b \\ &= +iA. \end{aligned}$$

"j-Equivalence" on Riemann surfaces

Much of (6.4.2) can be repeated in arbitrary non-zero codimension. In particular, there is an obvious analogue of the j-map in codimension 2, and the imaginary part of A has *integral periods* (to $\frac{i}{2\pi}$) if and only if this $\mathrm{j}'(Y) = 0$. As we have seen, the torus in this case is just the Picard variety viewed as a real manifold (6.4.1).

Theorem 6.5.1 *In this situation,* $\mathrm{j}'(Y) = 0$ *if and only if* L *is linearly-equivalent to the trivial divisor.*

Proof—Assume $\mathrm{j}'(Y) = 0$. Define

$$\begin{aligned} f :&= \exp\left(i - \int *\mathrm{d}\right)(*u) \\ &= \exp(i*u + \int b) \end{aligned}$$

where the integration means to fix some base-point and integrate along any path (avoiding Y) to the point in question. (Note that, since u is imaginary, the modulus of f is $\exp i*u$ and its argument is $-i\int b$.) This is a well-defined complex function off Y, since b is closed off Y, and if a closed loop surrounds a point of Y then $\mathrm{d}b = 2\pi i Y$ ensures that the integral jumps by an integer. If two paths differ by some non-contractible loop, the integral jumps by an integer precisely because b has integral periods (ie.$\mathrm{j}'(Y) = 0$).

We claim that (off Y) this f is holomorphic. This would follow if $\mathrm{d}\log f$ had type (1,0). In fact we find

$$\mathrm{d}\log f = i(a - ib)$$

which we already know to be (1,0).

Finally, we must show that f extends over Y as a meromorphic function with divisor Y. We shall derive globally the Poincaré-Lelong formula; then it will be straightforward to show that the singularities of f must be of the appropriate order and the work is done. As before, we use Λ for the adjoint to right-multiplication by the kähler form.

$$\begin{aligned} \Lambda(\frac{i}{\pi}\partial\overline{\partial}\log|f|) &= -\frac{1}{\pi}\Lambda\partial\overline{\partial}(*u) \\ &= -\Delta(\frac{i}{2\pi}*u) \end{aligned}$$

(since $\Delta \equiv -2i\Lambda\partial\overline{\partial}$ for 0-currents on any kähler manifold [21])

$$= + * Y$$

(since $*$ commutes with Δ). Note also that acting on 2-forms in $n = 2$,

$$*\Lambda = 1$$

so that applying $*$ to both sides yields

$$\frac{i}{\pi}\partial\overline{\partial}\log|f| = +Y$$

which is the Poincaré-Lelong equation as desired.

We claim this shows f to be meromorphic. It suffices to work locally in a disc around $Y = n$ copies of the point $z = 0$, for some integer $n \in \mathbb{Z}$. Cauchy's formula [21] says that over this disc

$$\frac{1}{2\pi i}\overline{\partial}\partial\log z^n = Y.$$

But Poincaré-Lelong shows that globally

$$\frac{1}{2\pi i}\overline{\partial}\partial\log f = Y.$$

So we know the local distribution f/z^n is harmonic, and is in particular smooth over $z = 0$. Thus f does not have an essential singularity there, and is locally $f = z^m \cdot g$ for some holomorphic non-zero g and some integer $m \geq n$. By the same reasoning, z^n/f is harmonic; and so $m = n$.

So in fact f is a global meromorphic function cutting out the divisor Y.

The converse follows smoothly in reverse. Given f, define

$$u := -i * \log|f|$$

which then obeys $\frac{i}{2\pi}\Delta u = -Y$ and from which we define the currents a and b as before. It remains to show that $\frac{i}{2\pi}b$ is integral when evaluated on $\mathrm{H}_1(X - Y; \mathbb{Z})$.

But given some representative cycle γ, the fact that

$$b := \mathrm{d}^* u = i\mathrm{dArg}f$$

ensures that

$$\int_\gamma \frac{i}{2\pi}b$$

is integer-valued. Consequently, $\mathrm{j}'(Y) = 0$. □

Chapter 7

Holonomy of loops and of surfaces

Given a line-bundle with connection on X^n, parallel translation defines a function (the holonomy function) on the infinite-dimensional space of loops on X. There are two ways to extend this notion to gerbs. First, it would seem unsurprising if a gerb with 1-connection were to specify a function on the space of closed surfaces in X. This is in fact the case (7.1), and a demonstration makes clear that the error 2-form ϵ of (3.3.2) should be thought of as the equivalent of a bundle's connection 1-form.

Second, guided by the principle that gerbs are to line-bundles what bundles are to functions, we might hope that (since a bundle creates a function) a gerb creates a line-bundle on the loop space. Brylinski [5] describes such a construction, crediting Gawędzki [17] and Deligne for the initial ideas. We can offer a more geometric interpretation in (7.2) with the tools of chapters 2 and 3.

This alternative view of the holonomy of gerbs is in some sense viewing a surface as a family of loops, and so is more complicated than (7.1). In return, the complications suggest interesting new avenues. For instance, recall that in chapter 6 a real harmonic-theoretic viewpoint was only partly enlightening when applied to holomorphic line-bundles (6.1)—except for the special case $n = 2$, for which knowing the metric gave complete control of the complex structure (6.5). We might hope the analogous gerbs (in $n = 3$) to be similarly friendly. If we notice further that in $n = 3$ there is known to be a natural almost-complex structure on the space of loops [34], we might start wondering whether we have in some sense a *holomorphic* bundle on the loop space: and even whether this bundle might be holomorphically trivial when the gerb is $\jmath$-trivial (6.4.1).

It seems reasonable in this context to upgrade the definition of the line bundle (7.3). Then, however, the curvature of a typical loop makes it hard to be explicit on the infinite-dimensional loop space; and it becomes natural to want to restrict to the space of geodesics. Here a nice theory arises (7.4), barring the most inconvenient fact that the set of geodesics is a manifold for very few spaces. The obvious example is $X = S^3$, whose geodesics form a complex quadric surface $\mathbb{Q}^2 = \mathbb{P}^1 \times \mathbb{P}^1$.

7.1 The holonomy of a surface

Start with a completely general smooth gerb with 1-connection on a general manifold X^n. For any surface Σ^2, we define a number, its holonomy, which in some sense is the result of "parallel translation" of the gerb around Σ.

Definition 7.1.1 (Holonomy of a surface) *Given a gerb $\mathcal{G}$ with 1-connection on X, and a smooth compact surface $\Sigma \in X$ (without boundary), pick any object of $\mathcal{G}|_\Sigma$ with 0-connection. The holonomy of Σ is then defined to be*

$$\exp - \int_\Sigma \epsilon$$

where ϵ is the error 2-form (3.3.2) of the object connection.

Since $\mathrm{H}^3(\Sigma;\mathbb{Z}) = 0$, objects exist for $\mathcal{G}|_\Sigma$. If the surface is a boundary $\Sigma^2 = \partial M^3$, then Stokes lets us calculate its holonomy via the 3-curvature

$$\exp - \int_\Sigma \epsilon = \exp - \int_M \Omega.$$

This definition is offered because objects and errors are the natural structures to seek to apply, given chapter 3; but we then see that this is in close analogy with the line-bundle holonomy around a loop

$$\exp - \int_\gamma A$$

where we trivialise the bundle over γ, yielding a connection 1-form A. That this is the correct definition justifies the remarks made in table (3.5), that an object with connection is a "trivialisation" of $\mathcal{G}$ and the error form is the equivalent of the connection 1-form.

Theorem 7.1.2 *The holonomy of Σ is independent of the choice of connection on the object, and of the choice of object. The holonomy of any Σ is constant within its homology class iff the 1-connection on $\mathcal{G}$ is flat (ie. $\omega = 0$).*

Proof—Fix a choice of object for $\mathcal{G}|_\Sigma$. Two object 0-connections differ by a 1-form α over Σ, by (3.2.3). From the proof of that result, we see that the errors differ by $\mathrm{d}\alpha$. Since Σ is closed, the integral is unchanged thereby.

Two choices of object (with 0-connection) differ by a line-bundle with connection over Σ. By the proof of (3.2.4), the errors differ by the curvature of this line-bundle. But this is integral

$$\int_\Sigma F = -2\pi i c_1[\Sigma] \in -2\pi i\mathbb{Z}$$

and thus the exponential is unchanged. So the holonomy is well-defined.

If $\omega = 0$ over X, then $[\mathcal{G}] = 0$ and we can take a *global* object with connection. In particular such a global error 2-form is closed, since

$$0 = \omega = \frac{i}{2\pi}\mathrm{d}\epsilon.$$

So its integral over Σ is determined by the surface's homology class.

The converse: if the holonomy is independent of the surface within its class, take Σ to be the boundary of a contractible ball B^3. Pick an object over the whole of B. Since $[\Sigma] = 0$,

$$1 = \exp - \int_\Sigma \epsilon = \exp 2\pi i \int_B \omega$$

and, since we can take arbitrarily small 3-balls spanning any element of $\Lambda^3 T_p X$ around any point $p \in X$, ω must vanish. □

7.2 A line-bundle on loop space

Start with a general gerb with 0-connection on some smooth X^n. Write LX for the space of compact connected oriented 1-dimensional submanifolds of X, ie. images of embeddings $\gamma : S^1 \to X$ of an oriented unparametrised circle. (Which particular space of loops we take is not crucial.) Brylinski defines a principal $\mathbb{C}^*$-bundle for his "gerbe with connective structure". Whilst our version looks a little different, it is directly inspired by an attempt to understand Brylinski's bundle.

Construction of the line-bundle

An open set in X lifts to one in LX (ie. all loops contained within it). We choose an open cover U_a for X such that its lifts cover LX, and such that $\mathcal{G}$ restricted to each U_a is trivial. (Note that these sets need have nothing directly to do with any given local trivialisation of $\mathcal{G}$.) Such a cover does exist: for instance, for every point in LX, we slightly thicken the corresponding loop $\gamma \subset X$ into an open solid torus. Since this set U_γ has vanishing $\mathrm{H}^3(\mathbb{Z})$, the restricted gerb is trivial.

Accordingly, choose on each U_a an object with 0-connection: and on the lift of the set, we place a trivialised line bundle $\hat{U}_a \times \mathbb{C}^1$. We must declare transitions for this cover.

On an intersection $U_{a,b}$ the difference of the two objects fixes a line bundle with connection (3.2.4), and its holonomy is a $\underline{\mathbb{C}}^*$-function on $\hat{U}_{a,b}$. We take this to be the transition function between the two patches on LX.

Proposition 7.2.1 *This construction is a well-defined smooth bundle on loop space, given a gerb with 0-connection on X.*

Proof—Such functions are smooth, so the bundle (if it exists) is smooth. Going from U_b to U_a has, as the difference of the two objects, the inverse bundle and connection; and thus the transition from $\hat{U}_b$ to $\hat{U}_a$ is the reciprocal of that in the other direction. Finally, on a three-fold intersection the double coboundary of the objects is canonically the trivial bundle and connection, and so the transitions on LX have zero coboundary.

Thus we have a smooth bundle on loop space, given the initial choices of objects and connections on each U_a. But if we consider alternative choices, they amount to gauge transformations of the trivialisations on each $\hat{U}_a$, and the holonomy on $\hat{U}_{a,b}$ is unchanged. □

We remark that there is no need to work over the whole of LX: we are quite entitled to restrict to subsets if we wish. Also, it will later be convenient for us to choose an arbitrary basepoint γ_a in each $\hat{U}_a$, and to insist that $\hat{U}_a$ be contractible. These can be accomplished, for instance, by taking open sets which are thickenings into tori of given loops γ_a.

Proposition 7.2.2 *This line bundle from a gerb with 0-connection is naturally identifiable with the line bundle of the principal $\mathbb{C}^*$-bundle of [5] (6.2) constructed from the corresponding gerbe with connective structure.* □

Connection from 1-Connection

Suppose that $\mathcal{G}$ also comes with a 1-connection β_i. Given an object with 0-connection over U_a, we now have an error 2-form over U_a

$$\epsilon := F_i - 2\pi i \beta_i.$$

View a tangent vector to a point in LX

$$v \in T_\gamma LX$$

as a vector field on S^1

$$v \in \Gamma(S^1; \gamma^* TX)$$

which is transverse to TS^1. With respect to the local trivialisation on $\hat{U}_a$, define a connection by its local 1-form A

$$A : v \mapsto \int_\gamma \mathrm{i}(v)\epsilon$$

in which i is to mean contraction between vectors and forms in X.

Proposition 7.2.3 *This connection is well-defined, with curvature 2-form*

$$F : (u, v) \mapsto \int_\gamma i(u, v) \mathrm{d}\epsilon$$

so that

$$\frac{i}{2\pi} F(u, v) = \int_\gamma i(u, v)\omega.$$

Proof—A direct translation of Brylinski [5] (6.2.2). We merely flag the two main points. First, transgression (integration of forms around a loop γ) is a morphism from the de Rham complex on X to that on LX. In particular,

$$\mathrm{d} \int_\gamma (\text{form}) = \int_\gamma \mathrm{d}(\text{form}).$$

Second, if the error form is ϵ for some object with 0-connection, then on any alternative choice it differs by the curvature 2-form of the difference bundle. This makes the connection 1-form consistent with the transition function, the holonomy. □

Higher gerbs

It is established that (2-) gerbs with 1-connection on X generate functions (0-gerbs) on the space of 2-manifolds in X, and generate bundles (1-gerbs) with (0-) connection on the space of loops. Given the corresponding facts about bundles with connection on X, and indeed $\underline{\mathbb{C}}^*$-functions on X, the pattern might be expected to continue without any great struggle to higher orders. A 3-gerb with 2-connection (4.5) should define a function on the set of 3-manifolds; a bundle with connection on the set of 2-manifolds; and a gerb with 1-connection on the loop space.

Although no new ideas should be necessary, we shall not write this out. Working case-by-case seems inefficient: perhaps some suitably algebraic notation will make this a tautology in all degrees.

7.3 Loops in a 3-manifold

For the rest of this chapter we restrict to a Riemannian X^3. Our starting point is the observation of Marsden and Weinstein [34] that LX has a natural almost-complex structure: given a vector $v \in T_\gamma LX$ (viewed now as a vector field on the loop γ that is *normal*, not just transverse), and writing $\dot\gamma$ for the unit tangent vector along the loop, define

$$\mathcal{I}(v) \quad := \quad \dot\gamma \times v$$

to be the vector product in three dimensions. This almost-complex structure is not integrable [31], but its Nijenhuis tensor vanishes [5] (so that if the manifold were *finite-dimensional* we would have integrability, by Newlander and Nirenberg [38]).

We see that, given a real vector v in the loop space (a normal vector field to γ), a vector of type (1,0) has the form

$$v \quad - \quad i\dot\gamma \times v$$

whilst

$$v \quad + \quad i\dot\gamma \times v$$

is of type (0,1). Then, as Brylinski notes, the curvature

$$F : (a, b) \mapsto \int_\gamma \mathrm{i}(a, b)\mathrm{d}\epsilon$$

of the line-bundle of (7.2) is a (1,1)-form (eg. by taking $(u, v, \dot\gamma)$ to be a real orthonormal triple and testing F on combinations of these types of vectors).

In finite dimensions, Griffiths [20] pointed out that a (1,1) curvature would define a unique holomorphic structure on the line-bundle: take the $\overline{\partial}$-operator to be the (0,1) part of the connection. Given the Poincaré $\overline{\partial}$-lemma, there then exist local holomorphic sections of the bundle.

In the context of the infinite-dimensional LX, this formal application of the Poincaré lemma would then drag in questions of exactly "how holomorphic" the bundle really is. (Lempert [31] for instance considers various more or less strong notions of integrability of the almost-complex structure on LX.) We do not propose to enter into this discussion. Instead, we remark that the particular sections s described in (7.2) are certainly not holomorphic, since for instance exterior derivatives of the transition functions with respect to such sections are not of type (1,0). We seek a modified bundle which is explicitly holomorphic.

Harmonic 1-connections

Various restrictions will come into play as we proceed, but for the moment consider a general smooth gerb $\mathcal{G}$ with 0-connection on a Riemannian 3-manifold X.

Definition 7.3.1 (Holonomy line-bundle $\mathcal{L}$) *Construct a line-bundle $\mathcal{L}$ with connection on LX as in (7.2) but with two changes—*

- *we insist that the local object connections on each U_a be "harmonic", in the sense that their curvature 2-forms are* **coclosed** *as well as closed; and*
- *the transition function is now not merely the holonomy function on $U_{a,b}$, but is multiplied by the function*

$$exp\ i \int_{\gamma_0}^{\gamma} T\Big))$$

where T is the 1-form on LX

$$T(v) \quad := \quad \int_{\gamma'} \mathrm{d}x \ \ast F(v)$$

which is to be integrated from the base-point γ_a to γ by any path in $\hat{U}_a$. (Here $\mathrm{d}x$ is induced by the metric on X; and F is the curvature 2-form of the line bundle on $U_{a,b}$.)

To show that this gives a smooth line-bundle, we must add to (7.2) the existence of such special object connections and the good behaviour of $\int T$.

The definition of T will be clearer when we consider the almost-complex structure on LX (7.4.1): indeed, that calculation is how T was found in the first place. Essentially, we are thinking of the holonomy function as the "modulus" of a complex transition function (although it need not be real), and then T is the "argument" that makes the function almost-holomorphic (meaning that its derivative is of type (1,0)).

We do not insist on the name "harmonic" above: it merely seems less awkward than "coclosed". Originally we expected a condition of *coexactness* to be necessary (due to its links with the Bogomolnyi equation) but this seems not to be the case. Perhaps the failure to exploit that condition expains why this chapter eventually runs into the sands, as we shall see.

Proposition 7.3.2 *Suppose the 0-connection of $\mathcal{G}$ is "harmonic" (in that the 2-form curvatures F_i^j are coclosed in some—and thus any—local trivialisation of the gerb). Then such object connections exist as demanded above, in some neighbourhood of any loop $\gamma_0 \subset X$.*

From here on, we restrict to such harmonic 0-connections on $\mathcal{G}$. For example, the 0-connection constructed in (6.3) from a submanifold R using harmonic theory has the form

$$F_i^j = \mathrm{d}^*(\sigma_i - \sigma_j)$$

which is eminently suitable.

Proof—We know objects with some general smooth 0-connection exist on the open set U_a. We shall take such a 0-connection and modify it.

First, note that if the forms F_i^j are coclosed, then $\delta F_i = F_i^j$ means that $\mathrm{d}^* F_i$ has vanishing Čech coboundary, ie.

$$\xi := \mathrm{d}^* F_i$$

is a global coclosed 1-form over U_a. If we shrink U_a so that $\mathrm{H}^2_{\mathrm{dR}}(U_a) = 0$ (for instance, by taking our sets to be thickened loops as heretofore suggested) then any such ξ is in fact coexact

$$\xi = \mathrm{d}^* \chi.$$

Suppose that χ can be taken to be *exact*

$$\chi = \mathrm{d}\psi$$

for some 1-form (defined over U_a). Then modify the object 0-connection by

$$\nabla_i \mapsto \nabla_i - \psi$$

to find that now each local F_i is coclosed as desired.

It remains to prove the supposition. This follows from Kodaira's decomposition [29, 10]: shrinking U so that χ is square-integrable, it decomposes into three unique smooth components as in (6.3.1), two of which are coclosed. Being interested only in $\mathrm{d}^*\chi$, we may insist that χ is of the third type, which is exact, and we are done. □

To answer the second difficulty, given that $\hat{U}_a$ is assumed to be contractible we need only show that T is closed for its integral to be well-defined.

Calculation of $\mathrm{d}T$

Now that the object connections are coclosed, so is the curvature F of the difference bundle on $U_{a,b}$.

Rather than worry directly about how to differentiate $\mathrm{d}x$, pick a fixed parametrisation $\mathrm{d}s$ for the circle—ie. take $\gamma : S^1 \to X$ as a map so that now

$$T(v) = \int_{S^1} \mathrm{d}s \, \big(l * F(v)\big)$$
$$l := \left|\frac{\mathrm{d}x}{\mathrm{d}s}\right|$$

and view fields as sections over S^1 of bundles associated to $\gamma^* TX$. Choose the map such that l is constant along the curve, so it equals the length of the loop with respect to the Riemannian metric. (Now, whilst $\dot{\gamma}$ remains the unit-length tangent vector, we shall write $\tau = l\dot{\gamma}$ for the tangent $\mathrm{d}/\mathrm{d}s$ with respect to the given parametrisation.)

For vectors u, v, it is general that

$$\mathrm{d}T(u,v) = u(Tv) - v(Tu) - T([u,v]).$$

On $\hat{U}_a$, therefore,

$$\begin{aligned}\mathrm{d}T(u,v) &= \int_{S^1} \mathrm{d}s\Big(*F(v)u(l) - *F(u)v(l) \\ &\qquad + l\big(u(*F(v)) - v(*F(u)) - *F([u,v])\big)\Big) \\ &= \int_{S^1} \mathrm{d}s\Big(*F(v)u(l) - *F(u)v(l)\Big)\end{aligned}$$

since $*F$ is closed. Now we need $u(l)$.

With metric g and Levi-Civita connection ∇ on X,

$$\begin{aligned} l^2 &= g(\tau,\tau) \\ \Rightarrow \quad u(l) &= \frac{1}{l}\, g(\tau, \nabla_u \tau).\end{aligned}$$

Since the connection is torsion-free

$$\nabla_u \tau - \nabla_\tau u = [u,\tau]$$

which vanishes since u is a deformation of the loop. The vector fields are orthogonal: differentiating $0 = g(u,\tau)$ gives

$$\begin{aligned} 0 &= g(\nabla_\tau u, \tau) + g(u, \nabla_\tau \tau) \\ \Rightarrow \quad u(l) &= -\frac{1}{l}\, g(u, \nabla_\tau \tau)\end{aligned}$$

So, writing $g(\nabla_\tau \tau)$ for the 1-form which is metric-equivalent to $\nabla_\tau \tau$,

$$\begin{aligned}\mathrm{d}T(u,v) &= \int_{S^1} \mathrm{d}s \cdot \frac{1}{l}\Big(*F \wedge g(\nabla_\tau \tau)\Big)(u,v) \\ &= \int_\gamma \mathrm{d}x\Big(*F \wedge g(\nabla_{\dot{\gamma}} \dot{\gamma})\Big)(u,v)\end{aligned}$$

since $\dot{\gamma}(l) = 0$.

This is not zero in general.

Proposition 7.3.3 *Suppose X is such that the set of closed geodesics $G \subset LX$ is a submanifold of the loop space. Then T restricted to G is a closed 1-form.*

Proof—Restricting to geodesics kills $\nabla_{\dot{\gamma}}\dot{\gamma}$. □

From here on, we insist that X have this property. (Unfortunately, there are very few examples.)

Theorem 7.3.4 *With the restrictions made so far—*

- *G is a manifold, and*
- *the 0-connection of $\mathcal{G}$ is harmonic*

—the construction (7.3.1) gives a smooth line-bundle over G. □

7.4 The holonomy bundle is holomorphic

The virtue of all this becomes clearer when we consider the almost-complex structure $\mathcal{I}$ on LX. We would wish to declare the special sections of (7.3.1) to be holomorphic. This does not hold for the old version (7.2). With the new, we need at least that the transitions preserve holomorphicity: in particular, that $\mathrm{d}\log(\text{transition})$ is a (1,0)-form. This is $\mathrm{d}\log(\text{holonomy}) + iT$.

Proposition 7.4.1 *On $(LU_0, \mathcal{I})$, $\mathrm{d}\log(\textit{holonomy}) + iT$ is a (1,0)-form.*

Proof—

$$\begin{aligned}
\mathcal{I}(\mathrm{d}\log(\text{holonomy}) + iT)(v) &= (\mathrm{d}\log(\text{holonomy}) + iT)(\mathcal{I}v) \\
&= (\mathrm{d}\log(\text{holonomy}) + iT)(\dot\gamma \times v) \\
&= -\int_\gamma \mathrm{i}(\dot\gamma \times v)F \;+\; i\int_\gamma \mathrm{d}x \cdot *F(\dot\gamma \times v) \\
&= \int_\gamma \mathrm{d}x \Big(i *F(\dot\gamma \times v) - F(\dot\gamma \times v, \dot\gamma)\Big) \\
&= \int_\gamma \mathrm{d}x \Big(- iF(v, \dot\gamma) - *F(v)\Big) \\
&= +i\big(\mathrm{d}\log(\text{holonomy}) + iT\big)(v)
\end{aligned}$$

as is clear by viewing $(\dot\gamma, v, \dot\gamma \times v)$ as an ordered orthogonal triple. □

So if we wish to take our harmonic generating sections to be holomorphic, we know now that the transition functions respect this formally in the almost-complex manifold LX. But we also know that for the transitions to be well-defined we must restrict to the putative submanifold of geodesics. If G were a manifold, counting degrees of freedom shows that we could expect it to be real 4-dimensional. It is not yet clear that the almost-complex structure restricts to G, ie. that $\mathcal{I}(TG) = TG$. If that were true, the fact that the Nijenhuis tensor vanishes would mean that the almost-complex structure is in fact integrable (by finite-dimensionality) and we would have a complex surface.

Theorem 7.4.2 *Given all previous constraints, $\mathcal{L}|_G$ is explicitly a holomorphic line-bundle (in that the local sections due to harmonic objects are holomorphic) iff X is Einstein*

$$\mathit{Ricci} = c\, g$$

for a constant scalar c.

*Proof—*It remains to show that $\mathcal{I}$ preserves TG if and only if X is Einstein. Given a geodesic γ, take a vector

$$u \in T_\gamma G \subset T_\gamma LX$$

ie. a Jacobi field [43] along the loop in X

$$\nabla_{\dot\gamma}\nabla_{\dot\gamma} u = R_{\dot\gamma,u}\dot\gamma.$$

We are interested in whether $v := \dot\gamma \times u$ is also Jacobi.

In this proof we use $*'$ to mean the Hodge dual acting on skew-symmetric products of vectors, but acting as the identity on forms. So for instance on vectors $a \times b = *'(a \wedge b)$. Since $*'$ depends purely on the metric, it is parallel under the Levi-Civita connection. So

$$\begin{aligned} \nabla_{\dot\gamma} v &= g(\dot\gamma, \nabla *' (\dot\gamma \wedge u)) \\ &= *' g(\dot\gamma, \nabla(\dot\gamma \wedge u)) \\ &= *'(\dot\gamma \wedge \nabla_{\dot\gamma} u) \end{aligned}$$

since γ is a geodesic. Similarly,

$$\begin{aligned} \nabla_{\dot\gamma}\nabla_{\dot\gamma} v &= \dot\gamma \times \nabla_{\dot\gamma}\nabla_{\dot\gamma} u \\ &= \dot\gamma \times R_{\dot\gamma,u}\dot\gamma \\ &= \frac{1}{|u|^2}\,\dot\gamma \times \Big(g(R_{\dot\gamma,u}\dot\gamma, u)u + g(R_{\dot\gamma,u}\dot\gamma, v)v \Big) \\ &= \frac{1}{|u|^2}\Big(g(R_{\dot\gamma,u}\dot\gamma, u)v - g(R_{\dot\gamma,u}\dot\gamma, v)u \Big). \end{aligned}$$

The deformation v is Jacobi if and only if this equals

$$R_{\dot\gamma,v}\dot\gamma = \frac{1}{|u|^2}\Big(g(R_{\dot\gamma,v}\dot\gamma, v)v + g(R_{\dot\gamma,v}\dot\gamma, u)u \Big)$$

so a necessary and sufficient condition is

$$\begin{aligned} g(R_{\dot\gamma,u}\dot\gamma, u) &= g(R_{\dot\gamma,v}\dot\gamma, v) \\ g(R_{\dot\gamma,u}\dot\gamma, v) &= 0. \end{aligned}$$

In general, the Ricci tensor is a sum of sectional curvatures: given an orthonormal basis,

$$Ricci(e_i, e_j) = \sum_{k=1}^{n} g(R_{e_j,e_k} e_i, e_k).$$

In three dimensions, this is invertible: say the basis is e, f, g. Then the full Riemann tensor is given by

$$\begin{aligned} g(R_{e,f}e, f) &= \frac{1}{2}(Ricci_{e,e} + Ricci_{f,f} - Ricci_{g,g}) \\ g(R_{f,g}f, g) &= \frac{1}{2}(Ricci_{f,f} + Ricci_{g,g} - Ricci_{e,e}) \\ g(R_{g,e}g, e) &= \frac{1}{2}(Ricci_{g,g} + Ricci_{e,e} - Ricci_{f,f}) \\ g(R_{e,f}g, f) &= Ricci_{g,e} \\ g(R_{f,g}e, g) &= Ricci_{e,f} \\ g(R_{g,e}f, e) &= Ricci_{f,g}. \end{aligned}$$

Thus our constraint says that pointwise

$$\begin{aligned} Ricci_{u,u} &= Ricci_{v,v} \\ Ricci_{u,v} &= 0 \end{aligned}$$

at any point in X, with $\dot{\gamma}, u, v$ an orthogonal triple for any choice of unit vector $\dot{\gamma}$. This holds if and only if in any orthonormal basis at any point, the Ricci tensor has all cross-terms vanishing and all diagonal terms equal

$$Ricci \quad = \quad cg$$

for some function c. In general, Bianchi forces c to be a constant. □

It is immediate from the above formulae (and is standard [42]) that any 3-manifold is Einstein if and only if it has constant sectional curvature, which then equals c.

A new connection

We need to upgrade the connection of (7.2) to take account of the complex structure.

Definition 7.4.3 (Connection on $\mathcal{L}$) *Given a 1-connection on $\mathcal{G}$, "harmonic" in that the local 2-forms β_i are coclosed, and a local basis of $\mathcal{L}$ on $\hat{U}_a$ as in (7.3.1), we define a connection 1-form*

$$A(v) \quad := \quad \int_\gamma \mathrm{i}(v)\epsilon \; - \; i \int_\gamma \mathrm{d}x \cdot \mathrm{i}(v) * \epsilon.$$

Such 1-connections do exist: in (6.3) we have for instance

$$\beta_i := \mathrm{d}^*(\gamma + \sigma_i).$$

Just as in our previous existence remarks, this would still work if we felt a need to impose the tighter constraint of coexactness rather than coclosedness.

Note also that at this point our "harmonic" terminology looks a little bedraggled, since we might expect a harmonic 1-connection to be one such that the curvature 3-form Ω is harmonic. No matter.

Proposition 7.4.4 *This defines a connection compatible with the holomorphic structure. The curvature*

$$\mathrm{d}A(u, v) \quad = \quad \int_\gamma \mathrm{i}(u, v)\Omega$$

is of type (1,1).

Proof—As in (7.2), we check consistency with the transition. (That the 1-connection is coclosed implies that $*\epsilon$ is closed.) Also A is of type (1,0), giving compatibility. We have already remarked upon the type of F. □

Conclusion

In the normal run of affairs, we would at this point consider some enlightening example to demonstrate that our effort has not been in vain. Unfortunately, the only example of a 3-manifold all of whose geodesics are closed is the sphere S^3, for which G is a quadric surface. There are things to say here but we leave them unsaid since, in different forms, they have largely been said before. Jones [28] and Pedersen [41] discuss line bundles on $\mathbb{Q}^2$ and abelian monopoles on S^3: and, as we have seen in chapter 6, there is little difference between a gerb and a monopole on spaces with vanishing $\mathrm{H}^2_{\mathrm{dR}}$.

However, the development in this chapter is perhaps merely misapplied rather than without value, and we retain it in the hope that this be so. Rather than dwell on the details of S^3, therefore, we shall move on to a much more productive exercise—the natural generalisation of the Ward correspondence to abelian gerbs on $\mathbb{Q}^6$.

Chapter 8

Twistor theory of quadric 6-folds

We consider a new twistor correspondence, as a companion to chapter 4. This time, rather than work with an affine-linear $\mathcal{Q} = \mathbb{C}^{n+1}$, we take a "quadratic" $\mathbb{C}^{2n+2}$. That is, we begin with a non-degenerate complex quadric

$$\mathbb{Q}^{2n+2} \quad \subset \quad \mathbb{P}^{2n+3}$$

on which there are two families of linear $\mathbb{P}^{n+1}$. The *twistor space* of this quadric is one of these families

$$\mathbb{T}_{n+1} \quad := \quad \{\text{all } \alpha\text{-planes } \mathbb{P}^{n+1} \subset \mathbb{Q}^{2n+2}\}$$

and is a complex manifold of dimension $(n+1)(n+2)/2$ for all $n \geq 0$.

Just as in chapter 4, we wish to start with a purely holomorphic n-gerb on $\mathbb{T}_{n+1}$ and see what this induces on $\mathbb{Q}^{2n+2}$ via their correspondence space. As before, it is appropriate to remove a closed subset from each side. We shall find that the transformation leaves us with a (holomorphic) n-gerb with *anti-self-dual* connection, in the sense that the curvature $(n+1)$-form vanishes on each α^{n+1}. In the line-bundle case $n = 1$, this is just the abelian version of the Ward correspondence [48, 2, 49] for which $\mathbb{Q}^4$ is the Klein quadric of lines in

$$\mathbb{T}_1 \quad = \quad \mathbb{P}^3.$$

The Ward correspondence is usually considered to apply to vector bundles on the 4-sphere. A number of attempts have been made to broaden it (see for instance [49, 8] and the many references therein). Restricting to the abelian case for obvious reasons, and without worrying about reality conditions that would transfer us from the quadric to the sphere, we aim to show enough to claim that n-gerbs are the natural fields to which the correspondence generalises. (So for instance if one claimed to know what an SU(2)-gerb might be, here is an obvious place to test it.)

Initially we discuss a general even-dimensional quadric $\mathbb{Q}^{2n+2}$ [22, 21]. Unlike chapter 4 however, we cannot complete the correspondence for arbitrary n since the twistor spaces $\mathbb{T}_{n+1}$ are now not as simple as were $\mathbb{P}^{n+1}$. At least we can make clear what further information is need (8.2). All the gaps can be bridged in the cases $n = 0, 1, 2$ for which $\mathbb{T}_{n+1}$ is understood, and we elucidate these in detail in (8.3) and (8.4). Of these, only $n = 2$ is a new result.

All structures in this chapter are holomorphic unless stated otherwise.

8.1 The geometry of quadrics

We gather some classical results which describe our correspondence [22, 21, 49, 8].

Quadrics, spheres and linear subspaces

Consider for $n \geq 0$ the spaces

$$\begin{cases} \mathbb{Q}^{2n+2} & \subset & \mathbb{P}^{2n+3} \\ \mathbb{T}_{n+1} & := & \{\text{all } \alpha\text{-planes } \mathbb{P}^{n+1} \subset \mathbb{Q}^{2n+2}\} \end{cases}$$

and their correspondence space

$$\mathcal{C} \quad := \quad \{(p, a) \ : \ p \in \mathbb{Q}^{2n+2}, \ a \in \mathbb{T}_{n+1}, \ p \in \alpha(a)\}.$$

This is a $\mathbb{P}^{n+1}$-bundle over $\mathbb{T}_{n+1}$, with the fibre over a point $a \in \mathbb{T}_{n+1}$ corresponding naturally to the plane $\alpha(a) \subset \mathbb{Q}^{2n+2}$ represented by a.

It is also a $\mathbb{T}_n$-bundle over $\mathbb{Q}$, since any α through $p \in \mathbb{Q}$ lies in the degenerate quadric

$$T_p\mathbb{Q} \quad \cap \quad \mathbb{Q}$$

and thus corresponds to a unique $\mathbb{P}^n$ in the "celestial" quadric at p

$$\mathbb{P}(T_p\mathbb{Q} \cap \mathbb{Q}) \quad \subset \quad \mathbb{P}(T_p\mathbb{Q}).$$

As this is a smooth $2n$-fold quadric, it has two families of $\mathbb{P}^n$. Its α^n-planes, parametrised by $\mathbb{T}_n(p) \cong \mathbb{T}_n$, correspond to the α^{n+1}-planes through p.

Another way of describing this is to start with $\mathbb{T}_{n+1}$ and define its dual twistor space as the space of all $\mathbb{T}_n \subset \mathbb{T}_{n+1}$: this is then a quadric $\mathbb{Q}^{2n+2}$, and $\mathcal{C}$ is the correspondence space for moving in either direction.

If we view $\mathbb{Q}$ as the complexification of a real quadric (the sphere S^{2n+2} for mathematicians, or a compactification of Minkowski space $\mathbb{R}^{1,2n+1}$ for physicists, with standard metric), then no α can hit the real quadric in more than one point, and through each point runs exactly one α. So $\mathbb{T}_{n+1}$ is a smooth $\mathbb{T}_n$-bundle over S^{2n+2}. From this perspective, it is natural to consider the fibre $\mathbb{T}_n(p)$ to be the space of orthogonal complex structures on T_pS

$$\mathbb{T}_n(p) \quad = \quad \mathrm{SO}(2n+2)/\mathrm{U}(n+1)$$

and the whole twistor space to be the space of orthogonal almost-complex structures on the Riemannian manifold S^{2n+2}

$$\begin{aligned} \mathbb{T}_{n+1} \quad &= \quad \mathrm{SO}(2n+3)/\mathrm{U}(n+1) \\ &= \quad \mathrm{SO}(2n+4)/\mathrm{U}(n+2) \end{aligned}$$

(although we shall not dwell on this approach).

What are these spaces? In low dimensions one finds

$$\begin{cases} \mathbb{T}_0 & = & \mathbb{P}^0 \\ \mathbb{T}_1 & = & \mathbb{P}^1 \\ \mathbb{T}_2 & = & \mathbb{P}^3 \\ \mathbb{T}_3 & = & \mathbb{Q}^6 \end{cases}$$

which we briefly justify: $\mathbb{T}_0$ = one point is of no interest except notationally, as the trivial fibre of the correspondence space over $\mathbb{Q}^2$.

Since $\mathbb{Q}^2 = \mathbb{P}^1 \times \mathbb{P}^1$, one ruling gives the α-lines and the other the β-lines, each parametrised by the other. So $\mathbb{T}_1$ is any choice of β-$\mathbb{P}^1 \subset \mathbb{Q}^2$.

The space of all $\mathbb{P}^1$ (=$\mathbb{T}_1$) in $\mathbb{P}^3$ is the Grassmanian of 2-planes in $\mathbb{C}^4$, known to one and all as the Klein quadric

$$\mathbb{Q}^4 \subset \mathbb{P}(\Lambda^2\mathbb{C}^4).$$

So $\mathbb{P}^3 = \mathbb{T}_2$.

To justify $\mathbb{T}_3 = \mathbb{Q}^6$ would mean delving into the triality of SO(8) for which we have no use except this result, so we shall be brief. The 8 derives from the $\mathbb{R}^8$ which, complexified and projectivised, becomes the $\mathbb{P}^7$ in which sits the standard quadric $\mathbb{Q}_0$. SO(8,$\mathbb{C}$) has an order-3 automorphism interchanging its standard representation with its two spinor representations, which give rise to further quadric 6-folds $\mathbb{Q}_+$ and $\mathbb{Q}_-$. It turns out that a point in $\mathbb{Q}_+$ corresponds to an α in $\mathbb{Q}_0$ (and to a β in $\mathbb{Q}_-$); whilst a point in $\mathbb{Q}_0$ gives a β in $\mathbb{Q}_+$ (and an α in $\mathbb{Q}_-$). On triality in general, see [9, 16]; and for examples of attempts at twistor theory in this context, see [33, 26, 45, 8].

None of these references is relevant to us except as background: what we need is that $\mathbb{T}_3 = \mathbb{Q}_+$, and that given a point $p \in \mathbb{Q}_0$ the space $\mathbb{T}_2(p)$ is naturally a β^3-plane $\subset \mathbb{Q}_+$.

Removing points

The standard Ward correspondence can afford to work over the whole of S^4 since there exist anti-self-dual SU(2)-bundles (for instance) on this space. In the abelian situation, we do not have such an ability—a harmonic $(n+1)$-curvature form on S^{2n+2} must vanish—and must instead work locally. Removing a point from S^{2n+2} provides enough freedom: taking this as our guide but without wanting to work on the sphere itself, what does this mean for $\mathbb{T}$ and $\mathbb{Q}$?

We start by fixing for all time an arbitrary basepoint point $x \in \mathbb{Q}^{2n+2}$. Knowing that $\mathbb{T}_{n+1}$ is a $\mathbb{T}_n$-bundle over the sphere, we shall have to remove the whole of $\mathbb{T}_n(x)$ from $\mathbb{T}^{n+1}$. This corresponds to all α-planes through x—these cut out the whole of the null cone of x, the degenerate quadric

$$T_x\mathbb{Q} \cap \mathbb{Q},$$

and we thus remove all these points from $\mathbb{Q}$. It is standard that this leaves behind an affine $\mathbb{C}^{2n+2}$. Every α^{n+1} in $\mathbb{Q}$ that avoids x cuts out a $\mathbb{C}^{n+1}$ in this space, hitting the null cone of x in a $\mathbb{P}^n$. Each point p off the null cone corresponds to some $\mathbb{T}_n \subset \mathbb{T}_{n+1}$ that is disjoint from the special $\mathbb{T}_n(x)$.

Thus we have a satisfactory two-way conversion between

$$\begin{cases} \mathcal{S} & := & \mathbb{Q}^{2n+2} - T_x\mathbb{Q} \\ & = & \text{affine } \mathbb{C}^{2n+2} \\ & = & \{\text{all } \mathbb{T}_n \subset \mathcal{T}\} \\ \mathcal{T} & := & \mathbb{T}_{n+1} - \mathbb{T}_n(x) \\ & = & \{\text{all } \alpha\text{-planes} \subset \mathbb{Q}^{2n+2} \text{ not fully contained in } T_x\mathbb{Q}\} \end{cases}$$

over which we still write $\mathcal{C}$ for the correspondence space. This remains a $\mathbb{T}_n$-bundle over $\mathcal{S}$ (and is necessarily holomorphically trivial) with projection map g; and is now an affine $\mathbb{C}^{n+1}$-bundle over $\mathcal{T}$, with projection f.

Tangent and normal bundles

As in chapter 4, the tautological vector bundle along the fibres of $(\mathcal{C} \xrightarrow{f} \mathcal{T})$ will be called

$$\mathcal{W} \to \mathcal{C}$$

and comes with a differential operator $\mathrm{d}^{\mathcal{W}}$. Given a point $p \in \mathcal{S}$, we write

$$\mathcal{N} \to \mathbb{T}_n(p)$$

for the normal bundle of $\mathbb{T}_n(p) \subset \mathcal{T}$: this induces a bundle over all of $\mathcal{C}$ which we also call $\mathcal{N}$. Both $\mathcal{W}$ and the tangent bundle to the fibres of g are subbundles of the tangent bundle of $\mathcal{C}$; and $\mathcal{N}$ is naturally its quotient by their direct sum. As in chapter 4, we do not need the tautological g-fibre bundle because its quotient—the normal bundle of the fibres of g—is canonically the pullback of the tangent bundle T of $\mathcal{S}$

$$0 \to \mathcal{W} \to g^*T \to \mathcal{N} \to 0$$

and g^*T is trivial since $\mathcal{S} = \mathbb{C}^{2n+2}$. By construction, both $\mathcal{N}$ and $\mathcal{W}$ are homogeneous over each $\mathbb{T}_n(p)$=SO$(2n+2)/$U$(n+1)$ fibre.

8.2 A conjectural transform

A toy correspondence: $n = 0$

To motivate our programme, consider the simplest case. A point $x \in \mathbb{Q}^2 = \mathbb{P}^1 \times \mathbb{P}^1$ lies on a single α_x and a single β_x. Their union is the null cone: removing it leaves

$$\mathcal{S} = \mathbb{C}^2.$$

The twistor space $\mathbb{T}_1$ can then be identified with β_x, since through each point of β_x runs exactly one α. Removing the single α in the null cone (ie. $\mathbb{T}_0(x)$) leaves

$$\mathcal{T} = \beta_x - \{x\} = \mathbb{C}^1.$$

The correspondence space $\mathcal{C}$ is $\mathcal{S}$ itself.

Note also that having chosen any $p \in \mathcal{S}$, and thus its corresponding $\mathbb{T}_0(p) \subset \mathcal{T}$ (ie. a point in $\mathcal{T}$), the total space of the normal bundle $\mathcal{N}(\mathbb{T}_0(p))$ can be identified with the space $\mathcal{T}$.

Proposition 8.2.1 *In this situation there is a canonical correspondence between holomorphic equivalence classes of 0-gerbs on $\mathcal{T}$, and of anti-self-dual 0-gerbs on $\mathcal{S}$.*

Proof—An equivalence class of 0-gerbs on $\mathcal{T}$ is a nowhere-zero holomorphic function (4.5). This lifts to $\mathcal{C}$ as a function on $\mathbb{C}^2$ that is constant in α-directions, and is thus a 0-gerb on $\mathcal{S}$ with the same property. This means its 1-curvature ($\mathrm{d}\log$ of the function) vanishes in α-directions.

The converse is not quite so demanding. □

Surprisingly, this content-free model is—barring technicalities—a complete guide to higher-dimensional quadrics.

The general case: $n \geq 0$

Conjecture 8.2.2 *For arbitrary $n \geq 0$, there is a canonical correspondence between holomorphic equivalence classes of*

- *n-gerbs on*
$$\mathcal{T} \subset \mathbb{T}_{n+1} := \{\textit{all } \alpha\textit{-planes } \mathbb{P}^{n+1} \subset \mathbb{Q}^{2n+2}\}$$
with zero "Chern class", in the sense that the class of the gerb lies in the kernel of $\mathrm{H}^n(\mathcal{T};\mathcal{O}^*) \to \mathrm{H}^{n+1}(\mathcal{T};\mathbb{Z})$; and

- *n-gerbs on*
$$\mathcal{S} \subset \mathbb{Q}^{2n+2}$$
with anti-self-dual connection, in that the $(n+1)$-form curvature vanishes when restricted to any $\alpha^{n+1} \cap \mathcal{S}$.

Proof—(Sketch)—As in the proof of (4.1.4), we begin with a representative for an equivalence class on $\mathcal{T}$, held to be some Čech cocycle

$$\theta \in [\theta] \in \mathrm{H}^n(\mathcal{T};\mathcal{O}).$$

(We have not shown this cohomology group to be non-empty.) Lift to $\mathcal{C}$

$$[f^*\theta] \in \mathrm{H}^n(\mathcal{C};\mathcal{O})$$

and conjecture this group to be empty (8.2.3). Then $f^*\theta$ is a coboundary: choose some cochain

$$\nu_{n-1} \in \mathrm{C}^{n-1}(\mathcal{C};\mathcal{O})$$

such that $\delta\nu = f^*\theta$, and differentiate along the $\mathcal{W}$-fibre directions to find

$$\mathrm{d}^{\mathcal{W}}\nu_{n-1} \in \mathrm{Z}^{n-1}(\mathcal{C};\mathcal{W}^*).$$

Conjecture (8.2.3) supposes further that

$$\mathrm{H}^{n-i}(\mathcal{C};\Lambda^i\mathcal{W}^*) = 0 \qquad (\forall i = 0,\ldots,n-1)$$

so that we can keep choosing coboundaries

$$\nu_{n-1-i} \in \mathrm{C}^{n-1-i}(\mathcal{C};\Lambda^i\mathcal{W}^*) \quad : \quad \delta\nu_{n-1-i} = \mathrm{d}^{\mathcal{W}}\nu_{n-i}$$

and keep differentiating until we reach

$$d^{\mathcal{W}}\nu_0 \in \mathrm{H}^0(\mathcal{C}; \Lambda^n \mathcal{W}^*).$$

Given the vanishing of $\mathrm{H}^{n-i-1}(\mathcal{C}; \Lambda^i \mathcal{W}^*)$ for $i \leq n-2$ (8.2.3), we would find ν_0 to be unique up to a global $d^{\mathcal{W}}$-exact section.

Our second requirement (8.2.4) claims that this descends to give an n-form on $\mathcal{S}$. We suppose also that this n-form is unique up to a global exact form. Further, $d(g_* d^{\mathcal{W}} \nu_0)$ vanishes on α-planes by (8.2.5), and we thus have an anti-self-dual connection (fixed up to $(n-1)$-equivalence) on the standard vacuous n-gerb on $\mathcal{S}$.

To go in the opposite direction: up to equivalence an n-gerb on $\mathcal{S}$ can be held to be vacuous (2.3.1), with the vacuous $(n-2)$-connection. The only choice in $(n-1)$-connection is then that of a global n-form—well-defined up to a global exact form—whose derivative is anti-self-dual. Such a field lifts to a section of $\Lambda^n \mathcal{W}^*$ (see comment after (8.2.4)); and anti-self-duality means that this section is $d^{\mathcal{W}}$-closed (8.2.5). (This is not automatic, since the rank of $\mathcal{W}^*$ is $n+1$, not n.) The application of a Poincaré lemma for $d^{\mathcal{W}}$ then permits us to drop down the ladder to some element of

$$\mathrm{Z}^n(\mathcal{C}; \mathcal{O})$$

which, being $d^{\mathcal{W}}$-closed, thus descends to $\mathcal{T}$. The freedom of an exact form on $\mathcal{S}$—which selects a unique section of $d^{\mathcal{W}} \Lambda^{n-1} \mathcal{W}^*$—varies this cocycle by a $d^{\mathcal{W}}$-closed coboundary in $\mathcal{C}$, and we are done.

This second direction is rigorous. If the first were to proceed as our outline suggests, the two would be inverses up to equivalence and the correspondence would be complete. □

Open issues

We gather together the gaps in the programme for conjecture (8.2.2).

Conjecture 8.2.3 *We need certain cohomology groups on $\mathcal{C}$ to vanish*

$$\begin{aligned} \mathrm{H}^{n-i}(\mathcal{C}; \Lambda^i \mathcal{W}^*) &= 0 \qquad i = 0, \ldots, n-1; \\ \mathrm{H}^{n-i-1}(\mathcal{C}; \Lambda^i \mathcal{W}^*) &= 0 \qquad i = 0, \ldots, n-2. \end{aligned}$$

Note that $\mathcal{W}$ is a sub-bundle of $g^* T\mathcal{S}$, and so the operator $d^{\mathcal{W}}$ is well-defined. We write Π for the map (induced by projection) from $g^* \Omega^i$ to $\Lambda^i \mathcal{W}^*$.

Conjecture 8.2.4 *We also suppose the existence of natural isomorphisms induced by g^* and Π*

$$\begin{aligned} \mathrm{H}^0(\mathcal{S}; \Omega^n) &\cong \mathrm{H}^0(\mathcal{C}; \Lambda^n \mathcal{W}^*) \\ \mathrm{H}^0(\mathcal{S}; d\Omega^{n-1}) &\cong \mathrm{H}^0(\mathcal{C}; d^{\mathcal{W}} \Lambda^{n-1} \mathcal{W}^*). \end{aligned}$$

We note at least that it is immediate—since the $\mathbb{T}_n$ are compact—that g^* generates isomorphisms

$$\mathrm{H}^0(\mathcal{S}; \Omega^n) \cong \mathrm{H}^0(\mathcal{C}; g^* \Omega^n).$$

Indeed, the proof of (8.2.5) demonstrates that there exist *injections* of the left-hand groups of (8.2.4) into the right, which suffices to make the direction from $\mathcal{S}$ to $\mathcal{T}$ rigorous in (8.2.2).

Proposition 8.2.5 *An i-form η defined locally on $\mathcal{S}$ ($i \geq 0$) obeys*

$$\mathrm{d}^{\mathcal{W}}(g^*\eta) \quad = \quad 0$$

iff $\mathrm{d}\eta$ *vanishes when restricted to any* α^{n+1}.

Proof—First, work pointwise at some $p \in \mathcal{S}$. Pick also a point $a \in \mathbb{T}_n(p)$ in the fibre over p, ie. pick some α-plane through p. Then by definition of $\mathcal{W}$ there is a vector-space identity

$$\mathcal{W}|_{(p,a)} \quad \cong \quad \alpha(a)|_{\mathcal{S}}$$

where the right-hand space has base-point p. So, if given a form at p

$$\zeta \quad \in \quad \Omega^i|_p$$

of degree $1 \leq i \leq n+1$, we see by definition that

$$\zeta|_\alpha = 0 \ (\forall \alpha \ni p) \quad \Leftrightarrow \quad \Pi(g^*\zeta) = 0.$$

(Still working pointwise in $\mathcal{S}$, we remark that Π injects as a linear map

$$\Omega^i|_p \to \mathrm{H}^0(\mathbb{T}_n(p); g^*\Omega^i) \to \mathrm{H}^0(\mathbb{T}_n(p); \Lambda^i\mathcal{W}^*)$$

for all $0 \leq i \leq n$. To prove this, it remains to show that if η^i vanishes on each $\alpha \ni p$, then it vanishes at p. Since the null cone

$$T_p\mathcal{S} \cap \mathcal{S}$$

spans $T_p\mathcal{S}$, we need only show that η vanishes on any rank-i vector space through p lying in the null cone. But it is a general fact that such a space—for $i \leq n$—must lie on some α through $p \in \mathbb{Q}^{2n+2}$.)

Now suppose that η is defined locally in $\mathcal{S}$. Then

$$\begin{aligned} \mathrm{d}^{\mathcal{W}} g^*\eta &= \Pi(\mathrm{d}g^*\eta) \\ &= \Pi g^*(\mathrm{d}\eta) \end{aligned}$$

and vanishes iff $\mathrm{d}\eta$ is zero in all α-directions. □

All of the incomplete claims can be justified for $n \leq 2$, which suffices to prove the gerb case. To go further would require greater knowledge of the twistor spaces

$$\mathbb{T}_{n+1} \quad = \quad \mathrm{SO}(2n+3)/\mathrm{U}(n+1)$$

than we can yet bring to bear.

8.3 Line bundles on $\mathbb{P}^3$: $n = 1$ done rigorously

It remains only to make explicit the transform from $\mathcal{T}$ to $\mathcal{S}$. In this setting, $\mathcal{T} = \mathbb{P}^3 - \mathbb{P}^1 = \mathbb{T}_2 - \mathbb{T}_1(x)$. Choosing some $p \in \mathcal{S}$ singles out a second, disjoint, $\mathbb{P}^1 = \mathbb{T}_1(p)$ in $\mathcal{T}$. Its normal bundle $\mathcal{N}$ is easily seen to be isomorphic to $\mathcal{O}(1) \oplus \mathcal{O}(1)$, and indeed the total space of $\mathcal{N}$ can naturally be identified with

$\mathcal{T}$—any point in $\mathcal{T}$ defines a unique $\mathbb{P}^2$ (spanned by that point and the missing $\mathbb{P}^1$) that forms a fibre of $\mathcal{N}$.

(This same phenomenon occurs trivially in $n = 0$, and also occurs in $n = 2$. Presumably it holds in general. The vector space of sections of $\mathcal{N}$ is for general n canonically identified with $\mathcal{S}$ (with zero-point p), since each parametrises the family of all $\mathbb{T}_n \subset \mathcal{T}$, but why $\mathcal{T}$ should be the total space of $\mathcal{N}$ is not so clear.)

The following result is the abelian Ward correspondence, and as such is not original.

Theorem 8.3.1 *Conjecture (8.2.2) holds true for $n = 1$: our twistor transformation naturally identifies holomorphic classes of—*

- *line bundles on*

$$\mathcal{T} \quad = \quad \mathbb{P}^3 - \mathbb{P}^1$$

  with zero Chern class; and

- *line bundles on the affine Klein quadric*

$$\mathcal{S} \quad \subset \quad \mathbb{Q}^4$$

  with anti-self-dual connection.

Proof—Choose some local Čech trivialisation

$$\theta \quad \in \quad \mathrm{Z}^1(\mathcal{T}; \mathcal{O})$$

for a topologically-trivial line bundle. Then $f^*\theta$ is $\mathrm{d}^{\mathcal{W}}$-closed, and we need that it be a coboundary

$$\mathrm{H}^1(\mathcal{C}; \mathcal{O}) \quad = \quad 0.$$

(This is all that is required to prove conjecture (8.2.3) for $n = 1$.) We use Leray for the $\mathbb{P}^1$-bundle ($\mathcal{C} \xrightarrow{g} \mathcal{S}$): the direct image sheaves $\mathrm{R}^q(\mathcal{O})$ vanish except for $q = 0$ (which gives a trivial line bundle over $\mathcal{S}$), and so $\mathrm{H}^p(\mathcal{C}; \mathcal{O}) = 0$ for all $p > 0$, which more than suffices.

Choosing ν_0 as in (8.2.2) gives

$$\mathrm{d}^{\mathcal{W}}\nu_0 \quad \in \quad \mathrm{Z}^0(\mathcal{C}; \mathcal{W}^*)$$

and we require Π to induce an isomorphism

$$\mathrm{H}^0(\mathcal{C}; g^*\Omega^1) \quad \cong \quad \mathrm{H}^0(\mathcal{C}; \mathcal{W}^*)$$

after which we can certainly descend to $\mathcal{S}$.

To prove this, expand the sequence

$$0 \to \mathcal{N}^* \to g^*\Omega^1 \xrightarrow{\Pi} \mathcal{W}^* \to 0$$

into

$$\to \mathrm{H}^0(\mathcal{C}; \mathcal{N}^*) \to \mathrm{H}^0(\mathcal{C}; g^*\Omega^1) \xrightarrow{\Pi} \mathrm{H}^0(\mathcal{C}; \mathcal{W}^*) \to \mathrm{H}^1(\mathcal{C}; \mathcal{N}^*) \to .$$

But all direct images $\mathrm{R}^q(\mathcal{N}^*)$ of $\mathcal{N}$ vanish by the usual methods; and so by Leray the two middle groups are canonically isomorphic.

Since $d^{\mathcal{W}}\nu_0$ is now in fact the pullback of a 1-form on $\mathcal{S}$, we apply (8.2.5) to show that the curvature 2-form is anti-self-dual.

An alternative choice of ν_0 differs by a global function on $\mathcal{C}$, which—since necessarily constant on any $\mathbb{P}^1$-fibre—is merely a global function on $\mathcal{S}$.

It is clear that this transform is the inverse of the already-justified transform in the other direction, and our mini-Ward correspondence is complete. □

8.4 Gerbs on $\mathbb{Q}_+^6$: $n = 2$ done rigorously

Remove the degenerate 5-fold quadric through a point x in a smooth quadric 6-fold $\mathbb{Q}_0$, to give $\mathcal{S}$. The quadric of the title is not this one but the space

$$\mathbb{T}_3 \quad = \quad \mathbb{Q}_+$$

of α^3-planes in $\mathbb{Q}_0$. The point $x \in \mathbb{Q}_0$ corresponds to a β-plane $\mathbb{T}_2(x) \subset \mathbb{Q}_+$ and removing this leaves $\mathcal{T}$.

Choose any point $p \in \mathcal{S}$, thus defining a second β-plane $\mathbb{T}_2(p) \subset \mathbb{Q}_+$ disjoint from the first. We need to know the rank-3 normal bundle $\mathcal{N}$ of this submanifold.

Proposition 8.4.1 *The manifold $\mathcal{T}$ is naturally the total space of a rank-3 affine bundle over the dual $\mathbb{P}^3$ to $\mathbb{T}_2(x)$.*

Proof—$\mathbb{T}_2(x)$ is a $\beta^3 \subset \mathbb{Q}_+$. It is classical that through each point in $\mathbb{Q}_+\backslash\beta$ runs a unique α such that $\alpha \cap \beta = \mathbb{P}^2$. (The generic intersection between an α- and a β-plane is a point.) Conversely, each $\mathbb{P}^2 \subset \beta$ lies on exactly one such α. These α are thus parametrised by the dual of β; and their affine parts $\alpha - \mathbb{P}^2$ form the fibres of a holomorphic bundle whose total space is $\mathbb{Q}_+ - \beta = \mathcal{T}$. □

Proposition 8.4.2 *Choosing any second disjoint $\beta = \mathbb{T}_2(p) \subset \mathcal{T}$, the affine bundle (8.4.1) naturally becomes the normal vector bundle $\mathcal{N}$ of $\mathbb{T}_2(p)$.*

Proof—The fibres now become vector spaces, since their intersections with $\mathbb{T}_2(p)$ single out choices of zero vector. Since each fibre is an α that is transverse to $\mathbb{T}_2(p)$, we can identify it with the fibre of the normal bundle. □

So again we have the curious fact that, given any two disjoint $\mathbb{T}_n \subset \mathbb{T}_{n+1}$, removing one leaves canonically the total space of the normal bundle of the other.

Proposition 8.4.3 *There exists an isomorphism*

$$\mathcal{N}(\mathbb{T}_2(p)) \quad \cong \quad \Omega^1 \otimes \mathcal{O}(2)$$

of bundles on $\mathbb{P}^3$.

Proof—We resort to coordinates $(x_0 : x_1 : x_2 : x_3 : y_0 : y_1 : y_2 : y_3) \in \mathbb{P}^7$ and take $\mathbb{Q}_+$ to be the hypersurface

$$x_0y_0 + x_1y_1 + x_2y_2 + x_3y_3 \quad = \quad 0.$$

Without elucidating the full system of α and β in this quadric [21], we summarise the relevant calculations. We single out two disjoint 3-planes

$$\begin{aligned} \mathbb{T}_2(x) &: \quad x_i = 0\ (\forall i) \\ \mathbb{T}_2(p) &: \quad y_i = 0\ (\forall i) \end{aligned}$$

—since these are disjoint, they must be from the same family. View these as our two β-planes. Then $\mathcal{S}$ is the vector-$\mathbb{C}^6$ of skew 4×4 matrices $B_{i,j}$ which correspond to β-planes spanned by the four points P_j with coordinates

$$\begin{aligned} x_i &:= \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \\ y_i &:= B_{i,j} \end{aligned}$$

for $j = 0, \ldots, 3$. Take the four standard patches

$$U_i \quad : \quad (x_i \neq 0)$$

for $\mathbb{T}_2(p) = \mathbb{P}^3$. Over the first patch U_0, the unique α that hits $\mathbb{T}_2(p)$ in one point $(1 : a_0 : b_0 : c_0 : 0 : 0 : 0 : 0)$ and $\mathbb{T}_2(x)$ in a $\mathbb{P}^2$ is the space spanned by the four points

$$\left\{ \begin{array}{llllllll} (1 & : a_0 & : b_0 & : c_0 & : 0 & : 0 & : 0 & : 0) \\ (0 & : 0 & : 0 & : 0 & : -a_0 & : 1 & : 0 & : 0) \\ (0 & : 0 & : 0 & : 0 & : -b_0 & : 0 & : 1 & : 0) \\ (0 & : 0 & : 0 & : 0 & : -c_0 & : 0 & : 0 & : 1). \end{array} \right.$$

This gives an open $\mathbb{C}^3 \times \mathbb{C}^3 \subset \mathcal{T}$ with coordinates $(a_0, b_0, c_0) \times (u_0, v_0, w_0)$ parametrising points

$$(1 : a_0 : b_0 : c_0 : -(a_0 u_0 + b_0 v_0 + c_0 w_0) : u_0 : v_0 : w_0)$$

—a local trivialisation of $\mathcal{N}$—together with another three such patches lifted up from the other three sets $U_{i \geq 1}$ in the cover of $\mathbb{T}_2(p)$. The numbers (a_0, b_0, c_0) are coordinates along $\mathbb{T}_2(p)$; and (u_0, v_0, w_0) are coordinates in the fibre directions of $\mathcal{N}(\mathbb{T}_2(p))$. In the chart U_1, this same point has the alternative description

$$(a_1 : 1 : b_1 : c_1 : u_1 : -(a_1 u_1 + b_1 v_1 + c_1 w_1) : v_1 : w_1)$$

from which we can read off the transitions of $\mathcal{N}$.

We now have control over $\mathcal{N}$. For instance, its space of sections is 6-dimensional (and naturally equals $\mathcal{S}$). In terms of coordinates over U_0, a section is given by six constants (E, F, G, H, I, J) which cut out the $\mathbb{P}^3$

$$\left\{ \begin{aligned} u_0 &:= E + H b_0 - I c_0 \\ v_0 &:= F - H a_0 + J c_0 \\ w_0 &:= G + I a_0 - J b_0 \end{aligned} \right.$$

(which is a $\beta \subset \mathbb{Q}_+$, as required).

To show that in fact $\mathcal{N} \cong \Omega^1 \otimes \mathcal{O}(2)$, we could for instance calculate the transitions for Ω^1. Slightly more elegantly, we know that $\mathcal{N}$ is homogeneous—and hence that it must be isomorphic to a twisted (co-) tangent bundle or a sum of lines [14]. We can calulate its Chern class to be

$$c(\mathcal{N}) \quad = \quad 1 + 2h + 2h^2 + 0$$

(where h is the positive generator of $\mathrm{H}^*(\mathbb{P}^3;\mathbb{Z})$) by picking sections as above; or, since $\mathcal{N}$ is globally generated by its sections, we need merely read off that its determinant is $\mathcal{O}(2)$. Either constraint suffices to fix the holomorphic isomorphism class from amongst the limited homogeneous alternatives. □

For a coordinate-free method of proof, adjunction and the Chern class of $\mathbb{P}^7$ give the Chern class of $\mathbb{Q}^6$. This and the class of $\mathbb{P}^3$ give that of the normal to $\mathbb{P}^3$. Then the quoted homogeneity result fixes $\mathcal{N}$ as discussed.

Gerbs and twistors

It is time to plug the gaps in conjecture (8.2.2).

Theorem 8.4.4 *There is a canonical equivalence between holomorphic classes of*

- *gerbs on*

$$\mathcal{T} = \mathcal{N} = \{\textit{all } \alpha^3 \subset \mathbb{Q}_0 \textit{ not containing } x\} \subset \mathbb{T}_3 = \mathbb{Q}_+$$

  (these necessarily have zero Chern class, since $\mathcal{T}$ is homotopic to $\mathbb{P}^3$ and has vanishing $\mathrm{H}^3(\mathbb{Z})$); and

- *gerbs with 1-connection on*

$$\mathcal{S} = \mathrm{H}^0(\mathbb{T}_2(x)^*;\mathcal{N}) = \{\textit{all } \beta^3 \subset \mathbb{Q}_+ \textit{ disjoint from } \mathbb{T}_2(x)\} \subset \mathbb{Q}_0$$

  with anti-self-dual curvature 3-form.

Proof—The first problem in the sketch of (8.2.2) is that we need

$$\mathrm{H}^1(\mathcal{C};\mathcal{O}) = \mathrm{H}^2(\mathcal{C};\mathcal{O}) = 0.$$

But $\mathcal{C}$ is a $\mathbb{P}^3$-bundle over $\mathbb{C}^6$, and Leray shows that $\mathrm{H}^p(\mathcal{C};\mathcal{O}) = 0$ for all $p > 0$.

So we can move to $\mathrm{H}^1(\mathcal{C};\mathcal{W}^*)$. We need to know that this too vanishes. But $\mathcal{W}^*$ is homogeneous over each $\mathbb{P}^3$; and the sequence

$$0 \to \mathcal{W} \to g^*T \to \mathcal{N} \to 0$$

shows that

$$c(\mathcal{W}^*) \quad = \quad 1 + 2h + 2h^2$$

so that in fact we find that $\mathcal{W}^*$ is isomorphic to $\mathcal{N} \cong \Omega^1 \otimes \mathcal{O}(2)$. Its direct images follow from (4.1.3)—R^0 is a trivial rank-6 bundle over $\mathcal{S}$, but all others are empty so that

$$\mathrm{H}^p(\mathcal{C};\mathcal{W}^*) \quad = \quad 0$$

for all $p > 0$. We now have a global section of $\Lambda^2\mathcal{W}^*$ (unique up to one of $\mathrm{d}^{\mathcal{W}}\mathcal{W}^*$).

To show that such a section is in fact a pull-back of a 2-form on $\mathcal{S}$, we merely calculate that the ranks of

$$\Omega^2_{\mathcal{S}}|_p \quad = \quad \mathrm{H}^0(g^{-1}(p); g^*\Omega^2)$$

and of

$$\mathrm{H}^0(g^{-1}(p); \Lambda^2 \mathcal{W}^*)$$

both come to 15. Since we know there is a natural injection of global sections (8.2.5), this suffices.

As ever, we can merely reverse our steps to see that the direction from $\mathcal{S}$ to $\mathcal{T}$ takes us back to the same class from which we started. □

On reality

One attractive feature of the $n = 1$ case is the way it descends to the real quadric S^4. As is well-known [3], the transform can begin with a real anti-self-dual 2-form curvature on S^4, which lifts to a (1,1)-form in $\mathbb{T}_2 = \mathbb{P}^3$. This connection then provides an integrable holomorphic structure on the bundle [20].

We shall not attempt to continue this in $n = 2$, but restrict ourselves to some remarks. First, if n is even—so the dimension of $\mathbb{Q}^{2n+2}$ is not divisible by 4—then any (anti-) self-dual $(n+1)$-form cannot be purely real, since $*^2 = -1$.

Second, consider representations. In $n = 1$ (ie. real 4-dimensional space), the bundle of complex 2-forms has four irreducible sub-bundles under U(2) which combine into two of SO(4)

$$\begin{aligned} \text{self-dual} &: \quad \Omega^{2,0} \oplus \langle k \rangle \oplus \Omega^{0,2} \\ \text{anti-self-dual} &: \quad \Omega^{1,1}_p \end{aligned}$$

in which p means the primitive part and $\langle k \rangle$ is the ray spanned by the kähler form k. This is why the (real) anti-self-dual 2-form curvature lifts to a form of type (1,1).

In $n = 2$, or $\mathbb{R}^6$ and its complexification, things are less tidy. The six irreducible parts of Ω^3 under U(3) gather together under SO(6) as

$$\begin{aligned} \text{self-dual} &: \quad \Omega^{2,1}_p \oplus (\Omega^{0,1} \wedge k) \oplus \Omega^{0,3} \\ \text{anti-self-dual} &: \quad \Omega^{3,0} \oplus (\Omega^{1,0} \wedge k) \oplus \Omega^{1,2}_p \end{aligned}$$

(if by self-duality we mean $* = +i$ rather than $-i$). So suppose we fix a complex 3-form at a point $p \in S^3$. The fibre $\mathbb{T}_2(p)$ of $\mathbb{T}_3$ over p is the collection of compatible complex structures on the real tangent plane of p; and so lifting the 3-form amounts to viewing it under each of these complex structures in turn. Now the representation theory above shows that anti-self-duality imposes merely the condition that *the lift of the 3-form has vanishing (0,3)-component at every point of the fibre*; and conversely any pull-back that has this property is anti-self-dual at p. (Similarly, self-duality is equivalent to having vanishing (3,0)-component under all compatible complex structures.)

Perhaps this is hinting that we should be considering 1-connections that are merely *compatible* with the holomorphic structure of the gerb (5.3.1) (rather than being themselves holomorphic), since such connections necessarily have no (0,3)-curvature; but there is a slight subtlety. Such a 1-connection (for which by definition the β_i have zero (0,2)-part) cannot merely be a lift of the 1-connection of S^6. This is because $\Lambda^2 T^*(\mathbb{R}^6) \otimes \mathbb{C}$ is irreducible under SO(6), so that a 2-form which has vanishing (0,2)-component under all compatible complex structures is in fact identically zero. We do not pursue a solution.

8.5 The twistors of chapter 4

It is not a coincidence that the ideas of chapter 4 translate so easily into this quadric problem. We remark first on the case $n = 2$, and then outline an approach to the general case.

Punctured projective 3-space

Recall that the transform of (4.4) was between a gerb on $\mathbb{P}^3$ minus a point, and a gerb with 1-connection on the dual $\mathbb{P}^3$ minus a hyperplane. This is in fact a sub-transform of that of (8.4), as we now demonstrate.

We know that each $\beta^3 \subset \mathbb{Q}_+$ corresponds to a point in $\mathbb{Q}_0$. Consider some fixed $\alpha^3 \subset \mathbb{Q}_+$, which we call A, corresponding to some $\beta \subset \mathbb{Q}_0$ labelled B. Now A intersects with $\mathbb{T}_2(x)$ (which is a β^3) in either a $\mathbb{P}^2$ or a point—the former if $x \in B$ and the latter otherwise. The fibres of $\mathcal{N}$ for instance lie in the former class.

We choose A to be of the latter type. This means that restricting to $\mathcal{T}$ gives

$$\begin{aligned} A \cap \mathcal{T} &= \mathbb{P}^3 - \text{one point} \\ B \cap \mathcal{S} &= \mathbb{P}^3 - \mathbb{P}^2 \end{aligned}$$

(the $\mathbb{P}^2$ being the intersection with the degenerate quadric through x).

If we now restrict the correspondence $\mathcal{C}$ to these subspaces, it is unproblematic to show that this is nothing other than the correspondence ($\mathcal{P} \leftarrow \mathcal{F} \rightarrow \mathcal{Q}$) of (4.4), ie. that between a projective space and its dual. Further, the twistor transformation from $\mathrm{H}^2(A; \mathcal{O})$ to a gerb with 1-connection on B is the same in both cases. We can see already that the discovery of chapter 4 that we naturally get a 1-connection along (partial) β-planes B in $\mathcal{S}$ is a hint towards the fact uncovered in (8.4.4) that the 3-curvature vanishes on α-planes.

This observation clearly suggests an alternative method of proof, by joining up these mini-transformations over the whole of $\mathcal{T}$ and $\mathcal{S}$. We do not follow this through, but instead highlight that this idea is not restricted to $n = 2$—indeed, Ward's original arguments in $n = 1$ were along such lines.

The correspondence between $\mathbb{T}$ and $\widetilde{\mathbb{T}}$

In $\mathbb{Q}^{2n+2}$, consider the space $\mathcal{L}$ of all linear $\mathbb{P}^n \subset \mathbb{Q}$. It is classical that each such $\mathbb{P}^n$ lies on a unique α^{n+1} and a unique β^{n+1}. Also, a given α cuts out a $\mathbb{P}^{n+1}$ in $\mathcal{L}$, which is naturally identified with the dual projective space of α. Thus $\mathcal{L}$ is a $\mathbb{P}^{n+1}$-bundle over $\mathbb{T}_{n+1}$, for which each fibre is the dual of the α-space to whose modulus $a \in \mathbb{T}_{n+1}$ it projects.

Writing $\widetilde{\mathbb{T}}_{n+1}$ for the moduli of β in $\mathbb{Q}^{2n+2}$, $\mathcal{L}$ is similarly a β^*-bundle over $\widetilde{\mathbb{T}}_{n+1}$. Further, we claim that a given fibre β^* identifies naturally with a $\mathbb{P}^{n+1} \subset \mathbb{T}_{n+1}$ (and similarly from $\mathbb{T}_{n+1}$ to $\widetilde{\mathbb{T}}_{n+1}$).
Proof—Fix a point in $\widetilde{\mathbb{T}}_{n+1}$ ie. a fibre $\beta^* \subset \mathcal{L}$. Through each point of β^* runs exactly one fibre α^*, because any point in $\mathcal{L}$ gives a $\mathbb{P}^n \subset \mathbb{Q}$, through which lies one α and one β. And each α^* intersects β^* in not more than one point, since an α and a β cannot share more than one $\mathbb{P}^n$. Each point of β^* thus cuts exactly one fibre of ($\mathcal{L} \rightarrow \mathbb{T}$), and then projects as claimed. □

Proposition 8.5.1 *Such a dual of a β-plane $\beta^* \subset \mathbb{T}_{n+1}$ intersects the given $\mathbb{T}_n(x)$ in*

- *a hyperplane of β^*, if $x \in \beta$; or*
- *a single point, otherwise.*

Proof—The sets β^* and $\mathbb{T}_n(x)$ share a point a iff the corresponding $\alpha(a) \subset \mathbb{Q}$ shares a $\mathbb{P}^n$ with β. Now β either lies fully inside the degenerate null cone, or intersects it in a $\mathbb{P}^n$. The latter would mean that β^* and $\mathbb{T}_n(x)$ share exactly one point, determined by that one $\mathbb{P}^n$; whilst the former requires us to find all $\mathbb{P}^n \subset \beta$ that define an α also containing x. This constraint is equivalent to requiring that the subspace $\mathbb{P}^n \subset \beta$ itself contains x; and the set of such $\mathbb{P}^n$ is a hyperplane in β^*. □

We expect similar ideas to yield a proof that $\mathcal{T}$ is in general the total space of the bundle $\mathcal{N}$, as found in the special cases $n \leq 2$. The value of this would be that it helps us get to $\mathcal{W}$ (which we need for conjectures (8.2.3) and (8.2.4)) since we also expect a general isomorphism

$$\mathcal{N} \quad \cong \quad \mathcal{W}^*$$

This claim should be straighforward to prove by the method of (4.2.2).

Theorem 8.5.2 *Fix such a β^{n+1} (such that $\beta^* \cap \mathbb{T}_n(x)$ is one point). Defining*

$$\begin{cases} \mathcal{P} & := \quad \beta^* - \mathbb{T}_n(x) \\ \mathcal{Q} & := \quad \{\textit{all hyperplanes } \mathbb{P}^n \subset \mathcal{P}\} = \beta \setminus \textit{Null quadric of } x. \end{cases}$$

Then the correspondence space

$$\mathcal{F} := \{(a,p) \;:\; a \in \mathcal{P},\; p \in \mathcal{Q},\; a \in p\}$$

is a sub-bundle of the correspondence

$$\mathcal{C} := \{(a,p) \;:\; a \in \mathcal{P},\; p \in \mathcal{Q},\; p \in \alpha(a)\}$$

Proof—This is almost tautologous, given the above discussion. □

If conjecture (8.2.2) were valid as described, presumably the twistor transformation of (4.1.4) and (4.5.1) would be its restriction. Inverting that point of view, attempting to glue together those chapter-4 fields should open up a second possible method for making an honest theorem of our conjecture.

Chapter 9

Concluding remarks

This chapter is an appendix of incomplete ideas and suggestions for further research.

9.1 Divisors and rulers

We present some plausible concepts of divisor for a holomorphic gerb. First, recall the Čech approach to divisors on a complex manifold X [21]. With $\mathcal{M}^*$ defined to be the sheaf of local meromorphic functions, not identically zero, a divisor (or "1-divisor" for present purposes) can be seen as an element of

$$\mathrm{Z}^0(X;\mathcal{M}^*/\mathcal{O}^*)$$

being cut out by local meromorphic functions whose ratio is holomorphic. This fixes a line-bundle, since such ratios on $U_{i,j}$ provide local transitions; and if alternative meromorphic functions were chosen, the differences on each U_i would combine as gauge transformations. This fits into the long exact sequence of

$$0 \to \mathcal{O}^* \to \mathcal{M}^* \to \mathcal{M}^*/\mathcal{O}^* \to 0$$

alongside notions of equivalence of line bundles and linear equivalence of divisors under global meromorphic functions.

Definition 9.1.1 (2-Divisor) *Given a cover U_i of a complex manifold X, a 2-divisor (or gerb divisor) is a Čech cocycle*

$$D \quad \in \quad \mathrm{Z}^1(X;\mathcal{M}^*/\mathcal{O}^*).$$

That is, we take (ordinary) divisors D_i^j on each $U_{i,j}$ whose boundaries on triple intersections cancel out.

Proposition 9.1.2 *A 2-divisor D generates a holomorphic gerb $\mathcal{G}(D)$.*

Proof—Each D_i^j gives us a bundle Λ_i^j. To get a section θ, choose defining meromorphic functions f for each D_i^j locally in $U_{i,j,k}$ (not necessarily over all of $U_{i,j,k}$). This gives a local trivialisation of $\delta\Lambda$, and the coboundary δf is a local $\mathcal{O}^*$-function. An alternative f shifts Λ_i^j by a gauge transformation and shifts the function δf by a matching factor. So the local functions δf can thus be viewed as a section $\theta_{i,j,k}$ of $\delta\Lambda$ throughout each $U_{i,j,k}$.

That $\delta\theta = 0$ follows since locally this is represented by $\delta^2 f$. So we have a gerb $\mathcal{G}$.

Varying the functions cutting out D varies the local trivialisation by holomorphic 0-gauge transformations, and the gerb is well-defined. □

Proposition 9.1.3 *Any gerb defined from a 2-divisor in the image of* $\mathrm{Z}^1(X;\mathcal{M}^*)$ *is trivial.*

Proof—The Λ_i^j are trivial, because each D_i^j is cut out by a single meromorphic function f_i^j over all of $U_{i,j}$. Since $\delta f = 1$, the section θ is the coboundary of sections of the Λ. So in fact any choice of representative meromorphic functions gives an explicit global trivialisation of $\mathcal{G}(D)$. □

Proposition 9.1.4 *If a 2-divisor D is a coboundary*

$$D \quad \in \quad \delta\mathrm{C}^0(X;\mathcal{M}^*/\mathcal{O}^*)$$

then its gerb $\mathcal{G}(D)$ is trivial.

Proof—Given a 1-divisor D_i on each U_i, we get line bundles L_i. Since $D_i^j = \delta D_i$, we find that $\Lambda_i^j = L_i \otimes L_j^*$, which gives canonical maps

$$m_i^j : L_i \mapsto \Lambda_i^j \otimes L_j$$

and thus a global object for $\mathcal{G}(D)$. □

It thus seems that we would wish to define the group of *linear equivalence classes* of 2-divisors to be

$$\frac{\mathrm{H}^1(X;\mathcal{M}^*/\mathcal{O}^*)}{\mathrm{H}^1(X;\mathcal{M}^*)}$$

which injects to $\mathrm{H}^2(X;\mathcal{O}^*)$ in the long exact sequence of $(\mathcal{O}^* \to \mathcal{M}^*)$. This would then say that two 2-divisors generate equivalent gerbs if and only if they are linearly-equivalent.

Existence of 2-divisors

If we fix a holomorphic gerb $\mathcal{G}$, then can we say that there must be some 2-divisor that generates it? This would mean that in some local trivialisation, the bundles Λ_i^j each had meromorphic sections (global over $U_{i,j}$) whose coboundary was the holomorphic section θ.

One might hope that a result analogous to that for line-bundles holds: that such 2-divisors exist for gerbs on projective varieties $X \subset \mathbb{P}^N$. The proof for bundles $\mathcal{L}$ amounts to noting that $\mathbb{P}^N$ has a positive line-bundle $\mathcal{O}(1)$, so that $\mathcal{L}$ can be tensored with arbitrarily large positive powers of $\mathcal{O}(1)|_X$ eventually yielding a holomorphic section, and thus a meromorphic section of $\mathcal{L}$. But since no such "positive gerb" exists ($\mathrm{H}^3(\mathbb{P}^N;\mathbb{Z}) = 0$) some other method is required.

Rulers

Now we take chapter 6 rather than chapter 5 as our starting point. One conceptual problem with 2-divisors (and their higher analogues for n-gerbs) is that—if we view objects as the "sections" of a gerb—a 2-divisor does not link directly with what we might wish a "meromorphic section" to be. Instead, consider real codimension-3 submanifolds.

Definition 9.1.5 (Ruler) *Given a holomorphic gerb $\mathcal{G}$ on a complex manifold X, a ruler is a smooth codimension-3 submanifold R such that (for some open cover U_i of X, with $U_0 := X\backslash R$) there exist complex manifolds $D_i^0 \subset U_{i,0}$ with*

$$\partial D_i^0 \quad = \quad R|_{U_i}$$

in the sense of manifolds (or singular chains) with boundary on U_i.

Roughly, a ruler is the boundary of a 2-divisor. As the definition of 2-divisor stands, however, the boundary of each $D_i^j \subset U_{i,j} \subset X$ could be quite unpleasant, and also needs to be apportioned between U_i and U_j. It seems better to start with R as a given.

We discuss manifolds, but it may be appropriate to permit R to be a singular homology cycle.

Sufficient conditions for R are that it be *real-analytic* and *maximally complex.* This latter means that under the almost-complex structure $\mathcal{I}$ of X, at every point $x \in R$ the real tangent space of R has a subspace

$$T_x R \quad \cap \quad \mathcal{I}(T_x R)$$

of maximal dimension (in this case, real dimension $2n-4$). (No doubt a familiarity with the literature on CR manifolds [27] would have a bearing on alternative conditions on R.) Harvey and Lawson [23] show that, given these two constraints, R lies locally in a unique complex hypersurface (without local boundary). Then there are exactly two possibilities for D_i^0, namely either "half" of this 1-divisor.

Without loss of generality, we suppose our cover to be as in (6.2): the $U_{i\neq 0}$ suffice to trivialise the normal bundle of R; and all intersections $U_{i,j,\ldots k}$ (all labels non-zero) are contractible.

Proposition 9.1.6 *A ruler R generates a gerb $\mathcal{G}(R)$.*

Proof—A choice of D_i^0 is a 2-divisor: for i and j non-zero, we define D_i^j by

$$D_i^j + D_j^0 + D_0^i \quad = \quad 0$$

ie. D_i^j is empty if D_i^0 and D_j^0 are on the same "side", and their set-theoretic union (a complex hypersurface containing R) otherwise. Then by construction we have a cocycle of 1-divisors, or a 2-divisor, and hence a gerb $\mathcal{G}(D)$.

What if we had a different 2-divisor? Apart from a different choice of cover, for which we merely take a common refinement, the only change can be that, on a given U_i^0, D' is on the other side of R as D. One finds that such a change induces a 0-gauge transformation, and accordingly $\mathcal{G} = \mathcal{G}(R)$ is well-defined. □

We can see this construction in coordinates, after a fashion. Suppose in $U_i = \mathbb{C}^n$ we have real coordinates $(x_1, y_1, \ldots, x_n, y_n)$ or complex coordinates $(z_1, \ldots, z_n)$ in the usual manner; and that R is the submanifold $z_n = 0$, $y_{n-1} = 0$. (We do not consider what constraints enable such coordinates to exist.) Of the two choices for D, we take $z_n = 0$, $y_{n-1} > 0$. To see the bundle Λ_i^0, we break up $\mathbb{C}^n \backslash R$ into three sets

$$\begin{cases} U &= \{y_{n-1} > 0\} \\ V &= \{y_{n-1} < 0\} \\ W &= \{z_n \neq 0\} \end{cases}$$

on which D is cut out by defining equations

$$\begin{cases} f_U &= z_n \\ f_V &= 1 \\ f_W &= 1 \end{cases}$$

respectively. Then Λ is of degree 1 on the space $\mathbb{C}^n \backslash \mathbb{R}^{2n-3}$ (which is homotopic to S^3).

A ruler $R \subset X$ faces an obvious topological constraint. The choice of ± 1 on each $U_{i,j}$ according to whether D_i^0 and D_j^0 are equal or opposite is a representative for the Stiefel-Whitney class w_1. If this vanishes (ie. R is orientable) then necessarily the Euler class of the normal bundle vanishes. This is because we can then take all D_i^0 to be on the same side, and they provide a smooth global nowhere-zero normal section. (We can thus take just two sets U_0 and U_1 for $\mathcal{G}(R)$, with a degree-1 line-bundle running along the whole punctured normal bundle.)

Proposition 9.1.7 *If R is the boundary of a global complex hypersurface $D \subset X$ (with boundary), then the gerb $\mathcal{G}(R)$ is trivial.*

Proof—The gerb generated by D has a global object: the line-bundle of D extends to all of U_0, and we take a trivial bundle on U_i. $\square$

We can also construct a gerb from R without passing to 2-divisors explicitly, by following the approach of (6.2.1). It is almost obvious (and easy to show) that the two give the same gerb $\mathcal{G}(R)$.

Proposition 9.1.8 *The cohomology classes*

$$\begin{cases} [R] & \in & \mathrm{H}_{2n-3}(X;\mathbb{C}) \\ c(\mathcal{G}(R)) & \in & \mathrm{H}^3_{\mathrm{dR}}(X) \end{cases}$$

are Poincaré-dual.

Proof—Given the above comment, this is just (6.2.2). $\square$

Being optimistic, we might take this as evidence that 2-divisors and rulers (or something very similar) are quite reasonable constructs.

Rulers and subjects

A meromorphic section of a line bundle gives a submanifold D and a trivialisation of the bundle off D. For gerbs, we have concentrated only on non-zero sections, or objects. Here we outline a possible position on singular objects, or "subjects". Note first that a gerb $\mathcal{G}(R)$ generated by a ruler comes with an object on $X\backslash R$, by choosing a trivialised bundle on U_0 and taking Λ_0^i on each $U_{i,0}$ (seen as the restriction to $X\backslash R$ of U_i). Any two such objects are naturally equivalent.

Definition 9.1.9 (Subject) *Suppose given a holomorphic gerb $\mathcal{G}$ on a complex manifold X. A subject of the gerb is specified by—*

- *a ruler R of X, and*
- *an object of $\mathcal{G}|_{X\backslash R}$*

such that there exists a 0-equivalence of $\mathcal{G}$ with $\mathcal{G}(R)$ over all of X that, restricted to $X\backslash R$, induces a 0-equivalence of the object with the object of $\mathcal{G}(R)$ defined above.

That is, the two objects on $X\backslash R$ differ by a trivial holomorphic line-bundle.

Note that we work only indirectly with R, unlike the usual line-bundle case where a meromorphic section explicitly cuts out the divisor D. To improve this, perhaps we need to decide on an appropriate notion of *singular line-bundle*. What we want to say is that the line-bundles of the object have degree-1 poles around R, in the sense of bundles on S^3 (or have poles of whatever degree corresponds to the local multiplicity of R).

Proposition 9.1.10 *If a holomorphic gerb has a subject with ruler R, then the classes $[\mathcal{G}]$ and $[R]$ are Poincaré-dual.*

Proof—$\mathcal{G}$ is 0-equivalent to $\mathcal{G}(R)$. □

9.2 Special Lagrangian submanifolds

Perhaps the main initial motivation for this dissertation was the recent flurry of interest in special Lagrangian submanifolds [23] of Calabi-Yau 3-folds begun by Strominger, Yau and Zaslow [46] in the context of mirror symmetry. Since these are rather particular codimension-3 submanifolds, one might expect that they are linked to rather particular gerbs and that their gerbs might provide a new and useful way to understand them. To the extent that special Lagrangian submanifolds deserve a status on a par with complex submanifolds—a view expressed for instance by Hitchin [25]—we could hope that special Lagrangian submanifolds give a more interesting notion of ruler than the holomorphic version in (9.1) (for some non-holomorphic but similarly rigid type of gerb).

We still hope this to be the case, but can offer little by way of results. In this section we merely sketch some connections with the ideas of chapter 6 on j-equivalence, but without bringing gerbs themselves into play. This section is thus very much a collection of scattered remarks rather than of settled ideas.

We shall find at least that there are intriguing hints of a relationship with current ideas of Donaldson and Thomas [13]. One way of reading that paper is as a complexification of classical gauge theory in dimensions 2, 3 and 4: another is as a search for the links between connections and submanifolds. They naturally end up looking at calibrated submanifolds of Calabi-Yau 3-folds. (We should also mention a recent preprint of Tyurin [47] exploring similar territory.)

Kähler forms and holomorphic forms

Definition 9.2.1 (Calabi-Yau manifold) *A Calabi-Yau manifold X is a kähler manifold of complex dimension m with a covariant-constant holomorphic m-form; or equivalently a Riemannian manifold with holonomy contained in $SU(m)$.*

It has become traditional in this context to write the kähler form as ω and the holomorphic form as Ω. Since throughout this text these symbols have other uses, we must break with tradition and write k and K respectively, for which we apologise. We also split K

$$K \quad = \quad \mu + i\nu$$

into real and imaginary parts. As usual X is held to be compact.

Without loss of generality, normalise K such that

$$(-1)^{m(m-1)/2} \frac{i^m}{2^m} \cdot K \wedge \overline{K} \quad = \quad \frac{k^m}{m!}$$

—this fixes K up to a global constant $\exp(it)$ and means for instance that

$$\begin{aligned} *\mu &= \nu \\ *\nu &= -\mu \end{aligned}$$

when m is odd, and

$$\begin{aligned} *\mu &= \mu \\ *\nu &= \nu \end{aligned}$$

when m is even. There is a local unitary basis

$$\mathrm{d}s^2 \quad = \quad \sum_1^m \phi_i \otimes \overline{\phi}_i$$

such that

$$K \quad = \quad \phi_1 \wedge \cdots \wedge \phi_m.$$

(Note that $*$ is the conjugate-linear operator

$$* : \mathcal{A}^{p,q} \to \mathcal{A}^{m-p,m-q}$$

of the convention for instance of Griffiths and Harris [21], and not the complex-linear operator

$$* : \mathcal{A}^{p,q} \to \mathcal{A}^{m-q,m-p}$$

of Wells [50], who writes $\overline{*}$ for our $*$.)

Definition 9.2.2 (Special Lagrangian submanifold) *A smooth real submanifold $R \subset X$ of a symplectic manifold is Lagrangian, if the kähler form k restricts to zero on R and if the real dimension of R is half that of X.*

If X is Calabi-Yau and R is symplectic, then R is a special Lagrangian submanifold if also ν vanishes on R; or equivalently if the restriction of μ is the induced volume form on R.

This follows Harvey and Lawson [23], who would call μ the *calibration*. They work in $\mathbb{C}^m$ but point out that Calabi-Yau manifolds are the appropriate extension. A special Lagrangian submanifold is volume-minimising in its homology class, just as are complex submanifolds (for which a similar approach can be taken with calibration $k^d/d!$). Alternative special Lagrangian calibrations are obtained by rotating K by some $\exp(it)$, so that for instance in [25] one takes ν to be the volume and μ to vanish.

Local coordinates

If we hope to treat special Lagrangian submanifolds as analogues of divisors of holomorphic 1-gerbs (line bundles), we might try to use the local equations cutting out the submanifold to define transitions for a gerb (of some special structure).

Such equations are described in [23] in $\mathbb{C}^m$, and to translate to a local manifold is straightforward. Suppose R is locally the common zero of real smooth functions f_1,..., f_m which are independent (in that the $\mathrm{d}f_i$ are linearly independent at each point). Then define two sets of vector fields

$$\begin{aligned} u_i &:= \mathrm{i}(g^{-1})\mathrm{d}f_i \\ v_i &:= \mathrm{i}(k^{-1})\mathrm{d}f_i \end{aligned}$$

so that the u_i are the gradient vector fields and are normal to R. Then R is Lagrangian iff for all i and j

$$k(u_i, u_j) = 0$$

and then the v_i are tangent to R. Now R is special iff

$$\nu(v_1, \ldots, v_m) = 0.$$

To convert into complex coordinates, split the local unitary basis into real and imaginary parts

$$\phi_i = \alpha_i + i\beta_i$$

so that the Riemannian metric, the real part of the hermitian metric, is

$$g = \alpha_1 \otimes \alpha_1 + \beta_1 \otimes \beta_1 + \ldots \quad .$$

Then

$$v_i = \sum_j \left(\alpha_j^*(f_i)\beta_j^* - \beta_j^*(f_i)\alpha_j^* \right)$$

and the condition that ν vanish on these vectors becomes, if m is odd,

$$0 \quad = \quad \text{Re det } \bar{\phi}_i^*(f_j)$$

(and Im det if m is even). Unfortunately, such equations are more complicated than the equivalents in the case of a holomorphic divisor, and how to construct an explicit gerb from them remains elusive.

j-Equivalence and complex tori

The virtue of theorem (6.4.2) is that it shows a correspondence between gerbs (with objective 1-connection) and *two* other structures: monopoles off a submanifold, and tori of harmonic forms. This means that even without knowing what a special Lagrangian gerb might be, we can still consider the remaining two ideas. Recall that our general principle of j-equivalence is to consider two submanifolds (of codimension at least one) in the same homology class, and to choose an arbitrary chain whose boundary is their difference

$$\partial C \quad = \quad R_1 - R_2$$

so that

$$[C] \quad \in \quad \mathrm{H}_*(X, R; \mathbb{Z})$$

ie. $C_{\text{harmonic}} + C_{\text{coexact}}$ is integral in its evaluation on compactly-supported closed forms on $X \backslash (R_1 \cup R_2)$. Then we have established that C_{coexact} is itself integral, iff $\mathrm{j}(R_1) = \mathrm{j}(R_2)$ in the appropriate torus. But C_{coexact} is the unique global coexact current whose exterior derivative is (up to sign) $R_1 - R_2$, and so now it restricts to a closed smooth integral form on $X \backslash (R_1 \cup R_2)$. The question for us now is what sort of submanifolds give rise to interesting structure for such equivalences.

First, note that the particular formulation of chapter 6 is not the only possible. We chose to dwell on harmonic language therein to give some sort of rigidity or naturality to our connections, but we can in fact cope without a metric, as follows. In general we need a *trivial gerb with flat 1-connection* on a smooth manifold. Since $[\mathcal{G}] = 0$, choose any global object with 0-connection, thus giving a global error form ϵ which is closed since $\Omega = 0$. This singles out a class in $\mathrm{H}^2_{\text{dR}}(X)$. An alternative object with connection varies ϵ by some closed integral 2-form $[F] \in \mathrm{H}^2_{\text{dR}}(X; \mathbb{Z})$. So we have a well-defined class

$$\mathrm{j}'(\mathcal{G}) \quad \in \quad \frac{\mathrm{H}^2_{\text{dR}}(X)}{\mathrm{H}^2_{\text{dR}}(X; \mathbb{Z})}.$$

This is zero iff the 1-connection is objective and iff the original j is zero (on compact Riemannian X, and supposing $\mathcal{G}$ to be generated from some submanifold R), thus returning us to theorem (6.4.2). The proofs of these two claims are transparent.

A further space that has featured so far is the moduli of holomorphic gerbs

$$\frac{\mathrm{H}^2(\mathcal{O})}{\mathrm{H}^2(\mathbb{Z})}$$

with zero Chern class (5.1). Taking X to be compact and kähler, we see that this complex torus is a part of the real torus

$$Jac = \frac{\mathcal{H}^2(X;\mathbb{R})}{\mathrm{H}^2(\mathbb{Z})}$$

and we can try to tease apart the rest of Jac in terms of holomorphic connections on $\mathcal{G}$.

Instead, we choose another illustration. We repeat the harmonic approach of chapter 6, now insisting that $R_1 - R_2$ is the difference of two special Lagrangian submanifolds in a Calabi-Yau 3-fold X. We need a few facts about such spaces first.

The structure of a special Lagrangian submanifold

Proposition 9.2.3 *On a compact Calabi-Yau 3-manifold, the harmonic component $\omega := R_{\mathrm{harmonic}}$ of a special Lagrangian submanifold R^3 of volume v takes the form*

$$\begin{aligned} \omega &= \frac{iv}{8}K + \pi + \overline{\pi} - \frac{iv}{8}\overline{K} \\ &= -\frac{v}{4}\nu + \pi + \overline{\pi} \end{aligned}$$

in which π is a global primitive harmonic (2,1)-form.

Proof—First, note that ω is real and that each component is individually harmonic. Next we claim that at every point of X

$$\omega \wedge \nu = 0.$$

To show this, consider a real smooth bump function f of total integral $X[*f] = 1$. If the support of f has interior in a "small" neighbourhood of $p \in X$, then

$$\begin{aligned} (\omega \wedge \nu)|_p &= Vol|_p * (\omega \wedge \nu)|_p \\ &= Vol|_p \cdot \int_X *(\omega \wedge \nu)|_p \cdot f \; Vol \\ &\simeq Vol|_p \cdot \int_X f * (\omega \wedge \nu) \; Vol \\ &= Vol|_p \cdot \int_X \omega \wedge (f\nu) \\ &= Vol|_p \cdot \omega[f\nu] \\ &= Vol|_p \cdot \omega[(f\nu)_{\mathrm{harmonic}}] \end{aligned}$$

since ω is harmonic. But since ν is harmonic, and the harmonic component $f_{\mathrm{harmonic}} = 1$, this is

$$\begin{aligned} \omega[1 \cdot \nu] &= R[\nu] \\ &= 0 \end{aligned}$$

and the claim is proven by localising the bump function arbitrarily finely.

As a consequence, and since

$$\omega^{3,0} + \omega^{0,3} \quad = \quad cK + \overline{cK}$$

for some global harmonic function (ie. constant) c, we find that in fact c must be imaginary. The general form in complex dimension m for this calculation is that

$$\begin{aligned}
\omega^{3,0} + \omega^{0,3} &= (-1)^m \frac{v}{2^{m-1}} * \mu \\
&= (-1)^m \frac{iv}{2^m} \left(\overline{K} + (-1)^m K \right).
\end{aligned}$$

In fact we can calculate c directly, using the condition that μ restricts to the volume form on R

$$\begin{aligned}
v &= R[\mu] \\
&= \omega[\mu] \\
&= \omega^{3,0+0,3}[\mu] \\
&= \int_X \frac{c}{2} K \wedge \overline{K} + \int_X \frac{\overline{c}}{2} \overline{K} \wedge K \\
&= 4i(\overline{c} - c) \int_X \frac{i}{8} K \wedge \overline{K} \\
&= 4i(\overline{c} - c)
\end{aligned}$$

and so c equals $iv/8$.

It remains to show that the (2,1)-part of R_{harmonic} must be primitive [21], which means in this context that $\omega^{2,1} \wedge k = 0$. We shall demonstrate that this follows from the constraint that k vanishes restricted to R.

Note that ω is trivial in its evaluation on $\mathrm{H}^1_{\mathrm{dR}}(X) \wedge k$, since for any closed 1-form α

$$\begin{aligned}
\int_X \omega \wedge \alpha \wedge k &= \int_R \alpha \wedge k \\
&= 0
\end{aligned}$$

by the Lagrangian constraint. Thus the component of ω in the dual of $\mathrm{LH}^1_{\mathrm{dR}}(X)$ vanishes. (L is the operation of right exterior product with k, which again is the convention of Griffiths and Harris rather than of Wells.) But this dual is $\mathrm{LH}^1_{\mathrm{dR}}(X)$ itself: for instance, we have the general formula

$$*L^r \xi^{p,q} \quad = \quad (-1)^{(p+q)(p+q-1)/2} \frac{r!}{(n-p-q-r)!} i^{p-q} L^{n-p-q-r} \overline{\xi}^{q,p}$$

if ξ is primitive [50, 21]. This proves the claim. □

Two further remarks: first, if we want $[R] = 0$ we must consider the difference of two special Lagrangian submanifolds and not a single manifold. This is because a special Lagrangian submanifold is volume-minimising in its homology class. Second, we might ask what structure can be put on the real 4-manifold $\partial C = R$. The shortest answer is that C cannot be a complex surface, since the fact that the tangent 3-planes of R are Lagrangian means that their complex span $T_pR \oplus \mathcal{I}T_pR$ is the whole of T_pX, rather than a complex hyperspace as required.

We can now refine our torus in imitation of the way the Hodge decomposition breaks up $\mathcal{H}^2(X;\mathbb{C})$. Given $[R_1] = [R_2]$, consider as in (6.3) solutions γ to

$$\frac{i}{2\pi}\Delta\gamma = \omega - R$$

for which we impose that the Chern form is harmonic $\omega = R_{\text{harmonic}}$ (and thus zero). Then γ is a global closed current. The fact that R is special Lagrangian means that γ splits by type as

$$\gamma = \phi\nu + \eta^{2,1} + \overline{\eta}$$

where ϕ is real and η is primitive. Then since $*\nu = -\mu$

$$\begin{aligned} *\gamma &= -\frac{\phi}{2}K + i\eta + \text{conjugate} \\ \mathrm{d} * \gamma &= -\frac{1}{2}\overline{\partial}\phi \wedge K + i\partial\eta + i\overline{\partial}\eta + \text{conjugate} \end{aligned}$$

which we can simplify, since $\mathrm{d}\gamma = 0$ means that

$$\begin{aligned} 0 &= -\frac{i}{2}\overline{\partial}\phi \wedge K + \partial\eta \\ 0 &= \overline{\partial}\eta + \partial\overline{\eta} \\ \Rightarrow \quad \mathrm{d} * \gamma &= (-\overline{\partial}\phi \wedge K + \text{conjugate}) + 2i\overline{\partial}\eta \\ &= (-\overline{\partial}\phi \wedge K + \text{conjugate}) + \mathrm{L}\alpha \end{aligned}$$

for some real primitive (1,1)-form α. (Primitive, because η being primitive means $\mathrm{L}\eta$=0, so that $\mathrm{L}\overline{\partial}\eta$=0.) Then

$$\begin{aligned} \beta_0 &:= \mathrm{d}^*\gamma \\ &= (*_3\partial\phi + \text{conjugate}) + \alpha^{1,1} \end{aligned}$$

where $*_3$ is defined from the above formulae. The notation is taken from Donaldson and Thomas [13]. The three components of β_0 are each primitive and globally coexact; and their derivative vanishes off R

$$\mathrm{d}\beta_0 = 2\pi i R.$$

Given R, we know β_0 to be unique (as the unique global coexact current whose derivative is $2\pi i R$). Then the $\mathfrak{j}$-image of R vanishes iff there is a monopole on X with singularity R for which the general Bogomolnyi equation $F = *\mathrm{d}$(3-form) of chapter 6 now has the more exotic form

$$F = (*_3\partial\phi + \text{conjugate}) + \alpha^{1,1}$$

for real primitive coexact α.

Complex submanifolds

We can break up the Higgs field in similar ways for submanifolds of real codimension 2 and 4 as well as 3: again, a comparison with Donaldson and Thomas [13] is suggestive.

Consider first the case of two homologous complex surfaces $D = D_1 - D_2$ in the Calabi-Yau 3-fold X. Then we can say

$$\begin{aligned} \Delta\gamma &= 2\pi i D \\ \gamma &= ak + \zeta \end{aligned}$$

for a real function a and a real primitive (1,1)-form ζ. Now

$$\begin{aligned} 0 &= \mathrm{d}\gamma \\ \Rightarrow \quad \mathrm{d}\zeta &= -\mathrm{d}a \wedge k \end{aligned}$$

and so

$$\begin{aligned} \mathrm{d} * \gamma &= \mathrm{d}a \wedge k^2/2 - \mathrm{d}\zeta \wedge k \\ &= \frac{3}{2}\mathrm{d}a \wedge k^2 \\ \Rightarrow \quad \beta_0 &= -\frac{3}{2} * \mathrm{L}^2 \mathrm{d}a \\ &= 3i\partial a + \text{conjugate.} \end{aligned}$$

(The general form on a compact kähler m-fold X is

$$\begin{aligned} \beta_0 &= \Big(1 + (m-1)!\Big) \cdot i\partial a + \text{conjugate} \\ &= -4\pi\Big(1 + (m-1)!\Big)\mathrm{d}^c a \end{aligned}$$

where a is a 0-current as above, for a pair D of homologous divisors.)

This is globally coexact, and closed off D. If further it is integral (to $\frac{i}{2\pi}$), a Bogomolnyi-type equation

$$A^{1,0} = -\frac{3}{2} * \overline{\partial} a \wedge k^2$$

(and its conjugate) constrain a real closed 1-form $A + \overline{A}$, which under integrality leads us to define

$$\int i(A + \overline{A})$$

as the argument of some complex function on $X \backslash D$: this is expected, as in (6.5), to be a globally meromorphic function cutting out the divisor D, which is then linearly equivalent to zero.

In real codimension 4, suppose that C is the difference between two homologous complex curves, and consider

$$\begin{aligned} \Delta\gamma &= 2\pi i C \\ \gamma &= f \cdot k^2 + \lambda \wedge k \end{aligned}$$

for a real function f and real primitive (1,1)-form λ. Then we find

$$\begin{aligned} d^*\gamma &= -2 * (\partial f \wedge k) + *\partial\lambda + \text{conjugate} \\ &= 2i\partial f \wedge k + *\partial\lambda + \text{conjugate} \end{aligned}$$

and the Bogomolnyi equation in case of j-triviality becomes that this equal a closed integral 3-form on $X \backslash C$. It would be interesting to compare such "gerb monopoles" (and the corresponding 3-gerb on all of X) with standard notions of algebraic equivalence of curves lying in an algebraic family.

Hyperkähler manifolds

Finally, we remark that there is a similar array of structures if we consider submanifolds of hyperkähler manifolds: 4-manifolds are uninteresting to the extent that a special Lagrangian submanifold in one complex structure is merely a complex curve in another; but hyperkähler 8-manifolds are leading us into codimension four, and thus 3-gerbs. Again the idea is to use the representation-theoretic decomposition of the appropriate bundles of forms [44]. This section has enough poorly-understood formulae however, and we shall add no more.

9.3 Non-abelian gerbs

Whilst this dissertation has dwelt exclusively on abelian gerbs, we can be sure that the most interesting future developments and applications will be in non-abelian contexts. A naive attempt to apply the approach of definition (2.1.1) to vector bundles (or non-abelian principal bundles) runs aground very rapidly; and yet since the original purpose of *gerbes* [18] is to understand non-abelian cohomology, there is surely some (relatively) straightforward gerb construction also. Aside from trying to digest the ideas of Giraud, not a prospect we confess that fills the heart with joy, three lines of attack present themselves.

The Ward correspondence

The twistor correspondence of chapter 8 between even-dimensional quadrics and their spaces of linear submanifolds seems to give a quite convincing local transformation between holomorphic gerbs and anti-self-dual connections. Knowing that the case of $\mathbb{Q}^4$ and $\mathbb{P}^3$ is much more interesting when non-abelian bundles are brought into play, we can perhaps look with some confidence for the 6-fold quadric to be a natural playground for non-abelian gerbs.

Passing to, say, SU(2) might let us work with global fields rather than just local, as in the bundle case. But a cautious start would be to remain local with a Lie-algebra valued curvature 3-form on $\mathcal{S} = \mathbb{C}^6$, and to ask what structures this induces on $\mathcal{T} = \mathbb{Q}^6_+ \backslash \mathbb{P}^3$. From $\mathcal{S}$ to $\mathcal{T}$ is not the more interesting direction in which to move, unfortunately, since it will not be clear just how much of the resulting data is necessary for a holomorphic structure on the "gerb" over $\mathcal{T}$. It might at least be suggestive.

Monopoles

In chapter 6 we see that abelian gerbs are closely linked with abelian monopoles with singularities along a codimension-three submanifold. We might thus begin with a singular non-abelian monopole, to search out some of the aspects of a non-abelian gerb. Now the Higgs field is valued in the Lie algebra of the group. It seems fair to begin with a monopole on a compact 3-manifold, with singularities at a cycle of points R of class $[R] = 0$. The BPS boundary condition for monopoles on $\mathbb{R}^3$

$$\phi \quad = \quad \phi_\infty + \psi \frac{1}{r} + O\Big(\frac{1}{r}\Big)$$

might be suitable—here ϕ_∞ takes values in an adjoint orbit in the Lie algebra, and ψ is a map from the sphere at infinity to the algebra. Thus our ubiquitous current equation

$$\Delta\gamma \quad = \quad 2\pi i R$$

will no doubt be enhanced by choices of adjoint orbits along R.

In this context we mention work of Pauly [40] on singular monopoles on compact 3-manifolds. He discusses the idea that, locally around each pole, the monopole corresponds to an anti-self-dual connection on a punctured 4-ball (a Hopf bundle over the punctured 3-ball). The pole is controlled by insisting that the connection extend over the whole 4-ball. It is not possible for topological reasons to extend the lift globally over the 3-manifold: but arguably this is not a problem for us, since (unless j=0) we do not expect a global monopole either. Its local purpose is merely to formalise the singularity of the subject of the gerb.

Strings and branes

We end with another key original motivation for this work, from theoretical physics. String theorists are nowadays compactifying all sorts of extended objects and finding various obscure fields induced upon them. Quite a cottage industry has arisen in hunting "gerbes", by which seems to be meant little more than closed integral n-forms well-defined up to the exterior derivative of an $(n-1)$-form. The categorical definition of gerbes seems not to be widely assimilated, for understandable reasons; neither does that of bundle gerbes (although Murray and collaborators have considered quantum anomalies and the like, most recently in [6]).

This therefore offers a set of examples on which to test our approach to gerbs: partly for the sake of developing its geometric consequences, but also in the hope that it will be of direct use in physical problems. Indeed, to the extent that situations arise in which putative "non-abelian gerbes" seem to play a role, we may find valuable hints as to a mathematical non-abelian codification. Witten has spoken for instance, not entirely frivolously, of configurations of k branes each of which supports a U(1)-gerbe, coalescing to form a single brane with U(k)-gerb. To seek to understand such processes more precisely would clearly have value both for mathematics and for physics, a unifying ambition which would be entirely proper and traditional.

Bibliography

[1] MF Atiyah—*Complex analytic connections in fibre bundles*—Transactions of the American Mathematical Society **85**, 181-207 (1957).

[2] MF Atiyah—*The geometry of Yang-Mills fields*—Lezioni Fermiane, Accademia Nazionale dei Lincei e Scuola Normale Superiore (Pisa 1979).

[3] MF Atiyah, NJ Hitchin and IM Singer—*Self-duality in four-dimensional Riemannian geometry*—Proceedings of the Royal Society of London **A362**, 425-461 (1978).

[4] R Bott—*Homogeneous vector bundles*—Annals of Mathematics **66**, 203-248 (1957).

[5] JL Brylinski—*Loop spaces, characteristic classes and geometric quantization*—Progress in Mathematics 107, Birkhäuser (Boston 1993).

[6] AL Carey, J Mickelsson and MK Murray—*Bundle gerbes applied to quantum field theory*—preprint `hep-th/9711133` (1997).

[7] AL Carey, MK Murray and BL Wang—*Higher bundle gerbes and cohomology classes in gauge theories*—Journal of Geometry and Physics **21**, 183-197 (1997).

[8] DMJ Calderbank—*Geometric aspects of spinor and twistor analysis*—PhD thesis, University of Warwick (1995).

[9] E Cartan—*The theory of spinors*—Hermann (Paris 1966); or *Leçons sur la théorie des spineurs*—Hermann (Paris 1937).

[10] G de Rham—*Differentiable manifolds*—Grundlehren 266, Springer (Berlin 1984); or *Variétés différentiables*—Hermann (Paris 1955).

[11] PAM Dirac—*Quantised singularities in the electromagnetic field*—Proceedings of the Royal Society of London **A133**, 60-72 (1931).

[12] J Dixmier and A Douady—*Champs continus d'espaces hilbertiens et de C^*-algèbres*—Bulletin de la Société Mathématique de France **91**, 227-284 (1963).

[13] SK Donaldson and RP Thomas—*Gauge theory in higher dimensions*—preprint, Oxford University (1996).

[14] G Elencwajg, A Hirschowitz and M Schneider—*Les fibrés uniformes de rang au plus n sur* $\mathbb{P}_n(\mathbb{C})$ *sont ceux qu'on croit*—in "Vector bundles and differential equations" edited by A Hirschowitz, Progress in Mathematics 7, Birkhäuser (Boston 1980).

[15] DS Freed—*Bundle gerbes*—Mathematical Reviews `98a:55016` (1998).

[16] W Fulton and J Harris—*Representation theory*—Graduate Texts in Mathematics 129, Springer (New York 1991).

[17] K Gawędzki—*Topological actions in two-dimensional quantum field theories*—in "Non-perturbative quantum field theory" edited by G 'tHooft and others, NATO ASI Series B 185, Plenum (New York 1988).

[18] J Giraud—*Cohomologie non-abélienne*—Grundlehren 179, Springer (Berlin 1971).

[19] R Godement—*Topologie algébrique et theorie des faisceaux*—Hermann (Paris 1958).

[20] PA Griffiths—*The extension problem in complex analysis II*—American Journal of Mathematics **88**, 366-446 (1966).

[21] PA Griffiths and J Harris—*Principles of algebraic geometry*—Wiley (New York 1978).

[22] J Harris—*Algebraic geometry*—Graduate Texts in Mathematics 133, Springer (New York 1993).

[23] R Harvey and HB Lawson—*Calibrated geometries*—Acta Mathematica **148**, 47-157 (1982).

[24] F Hirzebruch—*Topological methods in algebraic geometry*—Grundlehren 131, Springer (Berlin 1966); or *Neue topologische Methoden in der algebraischen Geometrie*—Springer (Berlin 1956).

[25] NJ Hitchin—*The moduli space of special Lagrangian submanifolds*—preprint `dg-ga/9711002` (1997).

[26] LP Hughston—*Applications of* $SO(8)$ *spinors*—in "Gravitation and geometry" edited by W Rindler and A Trautman, Bibliopolis (Naples 1987).

[27] H Jacobowitz—*An introduction to CR structures*—Mathematical Surveys and Monographs 32, American Mathematical Society (Providence 1990).

[28] PE Jones—*Minitwistors*—DPhil thesis, University of Oxford (1984).

[29] K Kodaira—*Harmonic fields in Riemannian manifolds*—Annals of Mathematics **50**, 587-665 (1949).

[30] HB Lawson and M-L Michelsohn—*Spin geometry*—Princeton UP (Princeton 1989).

[31] L Lempert—*Loop spaces as complex manifolds*—Journal of Differential Geometry **38**, 519-543 (1993).

[32] J Leray—*L'homologie d'un espace fibré dont la fibre est connexe*—Journal de Mathématiques Pures et Appliquées **29**, 169-213 (1950).

[33] YI Manin and Hoàng Lê Minh—*The Radon-Penrose transformation for the group SO(8)*—Mathematics of the USSR—Izvestiya **28**, 189-200 (1987); or Izvestiya Akademii Nauk SSSR Seriya Matematicheskaya **50** 195-206 (1986).

[34] JE Marsden and AD Weinstein—*Coadjoint orbits, vortices and Clebsch variables for incompressible fluids*—Physica **D7**, 305-323 (1983).

[35] JW Milnor and JD Stasheff—*Characteristic classes*—Annals of Mathematics Studies 76, Princeton UP (Princeton 1974).

[36] JR Munkres—*Elements of algebraic topology*—Benjamin Cummings (Menlo Park 1984).

[37] MK Murray—*Bundle gerbes*—Journal of the London Mathematical Society, Second Series **54**, 403-416 (1996).

[38] A Newlander and L Nirenberg—*Complex analytic coordinates in almost complex manifolds*—Annals of Mathematics **65**, 391-404 (1957).

[39] C Okonek, M Schneider and H Spindler—*Vector bundles on complex projective spaces*—Progress in Mathematics 3, Birkhäuser (Boston 1980).

[40] M Pauly—*Monopole moduli spaces for compact 3-manifolds*—Mathematishe Annalen **311**, 125-146 (1998).

[41] H Pedersen—*Einstein metrics, spinning top motions and monopoles*—Mathematische Annalen **274**, 35-59 (1986).

[42] P Petersen—*Riemannian geometry*—Graduate Texts in Mathematics 171, Springer (New York 1998).

[43] T Sakai—*Riemannian geometry*—Translations of Mathematical Monographs 149, American Mathematical Society (Providence 1996).

[44] S Salamon—*Riemannian geometry and holonomy groups*—Pitman Research Notes 201, Longman (Harlow 1989).

[45] MJ Slupinski—*The twistor space of the conformal six sphere and vector bundles on quadrics*—Journal of Geometry and Physics **19**, 246-266 (1996).

[46] A Strominger, ST Yau and E Zaslow—*Mirror symmetry is T-duality*—Nuclear Physics **B479**, 243-259 (1996).

[47] A Tyurin—*Special Lagrangian geometry and slightly deformed algebraic geometry*—preprint `math/9806006` (1998).

[48] RS Ward—*On self-dual gauge fields*—Physics Letters **A61**, 81-82 (1977).

[49] RS Ward and RO Wells—*Twistor geometry and field theory*—Cambridge UP (Cambridge 1990).

[50] RO Wells—*Differential analysis on complex manifolds*—Graduate Texts in Mathematics 65, Springer (New York 1980).

[51] AJHM van de Ven—*On uniform vector bundles*—Mathematische Annalen **195**, 245-248 (1972).