Introduction to Gerbes

Yiannis Loizides, Spring 2015

These notes introduce gerbes, which are "realizations" of degree 3 integral cohomology, analogous to the way line bundles "realize" degree 2 integral cohomology. We begin with the very concrete Hitchin-Chatterjee description in terms of transition line bundles, following Hitchin's notes ("Lectures on special Lagrangian submanifolds"). We then move on to Murray's bundle gerbes (following his notes "An introduction to bundle gerbes"). Although slightly more abstract, these lead to elegant constructions of some important examples. They are also a useful stepping stone to the still more abstract picture (Giraud, Brylinski) in terms of sheaves of groupoids, which we touch on only very briefly.

1 Transition line bundles

We will describe the Hitchin-Chatterjee picture in terms of hermitian line bundles (so all of our line bundles are equipped with a hermitian structure). As usual, this is equivalent to using $U(1)$ principal bundles. We could instead work with line bundles and principal $\mathbb{C}^\times$ principal bundles. It makes no difference (the structure group can always be reduced to $U(1)$), but as a differential geometer I like $U(1)$-bundles, so that is what I will use. For a line bundle L, I will write $Q = U(L)$ for the corresponding principal $U(1)$-bundle. By unitary section of $L|_U$, I will mean a section s such that $||s|| = 1$.

We begin by briefly reviewing the situation with line bundles. Let U_a be an open cover, and suppose we have $U(1)$-valued functions g_{ab} on double overlaps, $g_{ab} = g_{ba}^{-1}$. We require that these satisfy the cocycle relation on triple overlaps:

$$\delta g = g_{ab}g_{bc}g_{ca} = 1.$$

Then, we get a line bundle L in the usual way, with the g_{ab} acting as transition functions. The cocycle condition says that g_{ab} is a Cech 1-cocycle for the sheaf of $U(1)$-valued functions. Suppose g is a Cech coboundary, that is, $g_{ab} = h_a h_b^{-1} \Rightarrow h_a = g_{ab} h_b$. This says that the collection $\{h_a\}$ defines a global section of $U(L)$, so L is trivial. More generally representatives for the same Cech 1-cocycle give isomorphic line bundles. We've shown that line bundles are classified by $H^1(M, C^\infty(U(1)))$. By the exponential sequence, this is isomorphic to $H^2(M, \mathbb{Z})$. The class of the line bundle in H^2 is denoted $c_1(L)$; it is the first Chern class.

For gerbes, we want to try to do something analogous with $H^2(M, C^\infty(U(1)))$. Again the exponential sequence tells us that this is equivalent to $H^3(M, \mathbb{Z})$ (as advertised). So the starting point is an open cover U_a and a Cech 2-cocycle $g_{abc} : U_{abc} \to U(1)$. So on U_{abcd}, g satisfies

$$(\delta g)_{abcd} = g_{bcd} g_{acd}^{-1} g_{abd} g_{abc}^{-1} = 1.$$

We want in the end to have a definition independent of choices (cover and representative). This will come later; for now I'll follow Hitchin, who says "this is a working definition, and we are going to make gerbes work for us".

Definition 1.1. A *gerbe* is specified by an open cover U_a, and a Cech 2-cocycle with values in the sheaf of $U(1)$-valued smooth functions.

We will say that a gerbe is *trivial* if g is a coboundary. A *trivialization* of a gerbe is just a choice of cochain $f = f_{ab}$ such that $g = \delta f$. Now suppose f, f' are trivializations. Consider $h = f'/f$. Then

$$\delta h = \delta f' / \delta f = g/g = 1.$$

(Remember we are using multiplicative notation, so moving the δ inside the fraction is like linearity when we use additive notation!) This says that h is a 1-cocycle, hence it defines a line bundle! So we have shown that:

- Given a trivial gerbe, any two trivializations differ by a line bundle.

This is analogous to the fact that given a trivial line bundle, any two unitary global sections differ by a function $f : M \to U(1)$.

We obtain an example of a trivial gerbe by restricting the cocycle g to one of the open sets, say U_0. To see this, note that

$$1 = (\delta g)_{0bcd} = g_{bcd} g_{0cd}^{-1} g_{0bd} g_{0bc}^{-1} \Rightarrow g_{bcd} = f_{cd} f_{db} f_{bc},$$

where

$$f_{ab} := g_{0ab}.$$

(This makes sense as a cochain ONLY on U_0, since g_{0ab} is only defined on U_{0ab}.) So $g|_{U_0} = \delta f$, so we've found a trivialization of g over U_0.

On double overlaps, U_{ab} we would get two potentially different trivializations of $g|_{U_{ab}}$, coming from each of U_a and U_b as above. As we showed before, these differ by a line bundle L_{ab}. We introduce some notation, which has the goal of replacing Cech cocycles with line bundles. For a collection of line bundles $L_{a_1 \ldots a_n}$ defined on n-fold intersections, we define

$$(\delta L)_{a_1 \ldots a_{n+1}} = L_{a_2 \ldots a_{n+1}} L^{-1}_{a_1 a_3 \ldots} \ldots L^{\pm 1}_{a_1 \ldots a_n},$$

exactly analogous to Cech cocycles. Notice that $\delta\delta L$ is canonically trivial; for example, for L_{ab} we have

$$(\delta L)_{abc} = L_{ab} L_{bc} L_{ca},$$

and so $\delta\delta L$ is a product of 12 line bundles, which cancel in pairs (all (3 choose 2) times 2 possible ordered pairs of indices from the set abc appear). So $\delta\delta L_{abcd} \simeq \mathbb{C} \times U_{abcd}$ has a canonical trivializing section denoted 1. Similarly, if $s = s_{a_1 a_2 \ldots}$ is a collection of unitary sections of $L_{a_1 a_2}$ (alternating under permutations of indices), then we can define δs as a collection of sections of δL by putting (we do the case s_{abc} for simplicity):

$$\delta s_{abc} = s_{bc} s^{-1}_{ac} s_{ab} \in \Gamma(\delta L_{abc})$$

(tensor product section of a tensor product of line bundles). This is again analogous to Cech cocycles. We apply the same notation to the corresponding principal $U(1)$ bundles.

This leads to the second definition of gerbe:

Definition 1.2. A *gerbe* is specified by an open cover U_a, together with *transition line bundles* $L_{ab} = L^{-1}_{ba}$ on double overlaps such that the line bundle δL_{abc} over U_{abc} is trivial. We must also have a choice of unitary trivialization θ_{abc} of δL_{abc}, such that $\delta\theta = 1$ over quadruple overlaps.

It is equivalent to work with the corresponding unitary frame bundles $Q_{ab} = U(L_{ab})$. Then, an equivalent way of specifying the trivializing section is to have a choice of bundle isomorphism:

$$Q_{ab} \otimes Q_{bc} \to Q_{ac};$$

the condition $\delta\theta = 1$ then translates into the assumption that these bundle isomorphisms satisfy "associativity". We could also simply say that we have a unitary isomorphism $L_{ab} \otimes L_{bc} \to L_{ac}$, which we then just restrict to the unit circle bundles.

Here the trivialization θ is needed to get back the Cech cocycle, and is *essential* to the definition. Indeed, if we were to take a good cover, then all the transition line bundles will be trivial, and *all* the non-trivial topological data of the gerbe is contained in the choices of the θ's. In general, the non-trivial topological data is carried by a mixture of the θ's and the transition line bundles themselves. (An example is the Hopf bundle, where the trivializations carry no information (triple intersections are empty), and all the non-trivial topological information is contained in the transition line bundle.)

Given a gerbe as defined above, here is how we can recover a Cech 2-cocycle as in the first definition. Take a good cover, then all the L_{ab} are trivial. Choose trivializing unitary sections $\tau_{ab} = \tau_{ba}^{-1}$. Then

$$\theta_{abc} = g_{abc}\tau_{ab}\tau_{bc}\tau_{ca}$$

for some function $g_{abc} : U_{abc} \to U(1)$, and $\delta\theta = 1$ implies that this is a Cech 2-cocycle.

Another remark about the terminology: we call these line bundles "transition line bundles", but what do they transition between? This depends on the context they are being used in. But one important point that initially confused me is that you shouldn't think of them as transitioning between different line bundles—indeed if we have a good cover and $L_{ab} = L_a L_b^{-1}$ then δL_{abc} is canonically trivial, and suppose θ is this canonical section (this is what we would mean by saying that L_{ab} transition between line bundles L_a). Then choosing τ_a unitary trivializing sections of L_a, $\tau_{ab} := \tau_a\tau_b^{-1}$, we get that $\delta\tau = 1 \Rightarrow g_{abc} = 1$ by the above calculation. So the Cech 2-cocycle is trivial.

Example 1: the lifting gerbe. Let $P \to M$ be a principal G-bundle, and suppose $\tilde{G}$ is a central extension of G:

$$1 \to U(1) \to \tilde{G} \to G \to 1.$$

(An important example is the group $\mathrm{Spin}^c(n)$, which is a central extension of $SO(n)$.) Suppose we want to find a principal $\tilde{G}$ bundle $\tilde{P} \to M$ which "covers" P. In general this might not be possible, and the obstruction is encoded in a gerbe which we now describe.

Choose a trivializing cover U_a for the bundle P, and let $g_{ab} : U_{ab} \to G$ be the transition functions. What we seek are lifts of these to functions $\tilde{g}_{ab} : U_{ab} \to \tilde{G}$ such that

$$\tilde{g}_{ab}\tilde{g}_{bc}\tilde{g}_{ca} = 1.$$

Suppose we have two such lifts $\tilde{g}_{ab}, \tilde{g}'_{ab}$, defining two lifts $\tilde{P}, \tilde{P}'$. Since $\tilde{g}, \tilde{g}'$ both map to the same element in G, the above exact sequence implies that their quotient is in $U(1)$, and thus we get $U(1)$-valued functions $f_{ab} := \tilde{g}'_{ab}/\tilde{g}_{ab}$. These functions satisfy the cocycle relation, and so define a line bundle L (or equivalently a principal $U(1)$-bundle $Q = U(L)$). Also by construction,

$$\tilde{g}'_{ab} = \tilde{g}_{ab} f_{ab},$$

which just says that $\tilde{P}' = \tilde{P} \otimes Q$. [1] That is, any two choices of lifts (when they exist) differ by a principal $U(1)$-bundle.

Now, since P is trivial over each U_a, we can find a lift $\tilde{P}_a$ over U_a. However the restrictions of $\tilde{P}_a$ and $\tilde{P}_b$ to the double intersection U_{ab} need not agree: by the computation above they will differ by a principal $U(1)$ bundle in general, say $Q_{ab} = U(L_{ab})$ where L_{ab} is a hermitian line bundle over U_{ab}, and we have isomorphisms

$$\tilde{P}_a \otimes Q_{ab} = \tilde{P}_b.$$

(By writing "=" we really mean that we have *chosen* an isomorphism.) On triple intersections we have

$$\tilde{P}_a Q_{ab} Q_{bc} Q_{ca} = \tilde{P}_b Q_{bc} Q_{ca} = \tilde{P}_c Q_{ca} = \tilde{P}_a,$$

[1] The tensor product here means take the principle $\tilde{G}$ bundle whose transition functions are the product of the transition functions for the two bundles. This works since $U(1) \subset \tilde{G}$ is in the centre of $\tilde{G}$. More abstractly, we take the fibre product of $\tilde{P}$ and Q, and then quotient by the appropriate action of $U(1)$.

which shows that δL_{abc} is trivial: let θ_{abc} be the unique section of δQ such that $\tilde{g} \mapsto \tilde{g} \otimes \theta_{abc}$ is inverse to the isomorphism above.

Everything seems to depend strongly on all the choices—we'll get a much more elegant description of this gerbe when we pass to bundle gerbes. The above was a good example of why the L_{ab} can be thought of transition line bundles: here transitioning between different lifts to a central extension.

Example 2: the "Hopf gerbe"/magnetic monopole. The first example of a non-trivial line bundle is the Hopf bundle on S^2: we get this by a "clutching construction", using for transition function along the equator a map $S^1 \to U(1)$ with winding number 1. Its chern class is the generator for $H^2(S^2)$. We can do a similar construction for S^3, but now instead of a transition function we have a transition line bundle. The equator is now S^2, and we can use the Hopf line bundle as our transition line bundle. We get a gerbe, representing the generator of $H^3(S^3)$. (Don't need to specify a trivializing section here, as triple intersections are empty. Indeed this is an example of the opposite extreme, where all the nontrivial topological data is contained in the transition line bundles.)

Actually we can be more general: let M be any 3-manifold. Given a point $p \in M$ and a ball B around p, $\partial B \simeq S^2$ and we can do the same thing. This is analogous to the construction of the holomorphic line bundle L_p on a Riemann surface.

2 Connections and curvature

Definition 2.1. A *connection* (or connective structure) on a gerbe (L_{ab}, θ_{abc}) consists of a hermitian connection ∇_{ab} on L_{ab} and collection of 2-forms F_a on U_a such that

1. $\nabla_{abc}\theta_{abc} = 0$ (here ∇_{abc} is the induced connection on δL_{abc})
2. $F_b - F_a =: F_{ab}$ (on double overlaps) is the curvature of ∇_{abc}.

The forms F_a are sometimes called the *curving* of the connection, or sometimes the B-field (think: the B-fields we encountered last semester, for

Courant algebroids!) On double overlaps we have $d(F_b - F_a) = dF_{ab} = 0$, which shows that

$$G|_{U_a} := dF_a,$$

defines a global 3-form called the *3-curvature.*

Take a good cover, and let τ_{ab} denote choices of unitary trivializing sections for L_{ab}. Recall that the Cech 2-cocycle g of the gerbe is given by

$$\theta_{abc} = g_{abc}(\delta\tau)_{abc}.$$

Let A_{ab} be connection 1-forms relative to the chosen trivializations, that is

$$\nabla_{ab}\tau_{ab} = -iA_{ab}\tau_{ab}.$$

Then $A_{ab} + A_{bc} + A_{ca}$ is the connection 1-form on δL_{abc} relative to the trivializing section $\delta\tau_{abc}$. Thus

$$0 = \nabla_{abc}\theta_{abc} = (dg_{abc})\tau_{abc} - ig_{abc}(A_{ab} + A_{bc} + A_{ca})\tau_{abc},$$

which gives

$$iA_{ab} + iA_{bc} + iA_{ca} = g_{abc}^{-1}dg_{abc}.$$

Proposition 2.2. *The cohomology class of the 3-curvature is the image in* $H^3(M, \mathbb{R})$ *of the class representing the gerbe in* $H^3(M, \mathbb{Z}) \simeq H^2(M, C^\infty(U(1)))$.

Proof. (Sketch) We need to recall how to get a Cech cocycle beginning with a de Rham cocycle. Suppose G is a closed 3-form. Let U_a be a good cover. Then by Poincare $G|_{U_a} = dF_a$ for some 2-forms F_a. On double overlaps $d(F_b - F_a) = G - G = 0$, so we can find 1-forms A_{ab} defined on double overlaps such that $F_b - F_a = dA_{ab}$. Then $\delta A_{abc} = A_{ab} + A_{bc} + A_{ca}$ is a sum of six terms which cancel in pairs. So we can find functions $f_{abc} : U_{abc} \to \mathbb{R}$ such that $A_{ab} + A_{bc} + A_{ca} = df_{abc}$.

On quadruple overlaps, let $k_{abcd} = (\delta f)_{abcd} = f_{bcd} - f_{acd} + f_{abd} - f_{abc}$. Then $d(\delta f)_{abcd} = 0$, and so the k_{abcd} are *constants.* This is the desired Cech 3-cocycle (for the *constant* sheaf $\mathbb{R}$). If G was the image of an integral class, then we can make all choices such that these constants are *integers.* Then define

$$g_{abc} = \exp(2\pi i f_{abc});$$

this is the desired (multiplicative) Cech 2-cocycle—the fact that the constants are integers ensures that $\delta g = 1$ even though $\delta f \neq 0$ in general. (This formula for g is $\exp \circ \delta^{-1}(k)$, which is the inverse of the boundary map $\delta \circ \log$ in the exponential long exact sequence.) □

3 Holonomy

Definition 3.1. A connection on a gerbe is *flat* if its 3-curvature vanishes.

Note that this does not imply that the gerbe is trivial, but only that the Cech 3-cocycle is torsion. When a line bundle is flat, it can still be non-trivial; and this can be captured through the holonomy of the connection. Similarly we will next define the holonomy of a flat gerbe.

Suppose $(L_{ab}, \theta, \nabla_{ab}, F_a)$ is a flat gerbe. Choose a good cover, so

$$F_a = dB_a \Rightarrow F_b - F_a = dA_{ab} = d(B_b - B_a)$$

so there exist functions f_{ab} on double overlaps such that

$$A_{ab} - B_b + B_a = df_{ab}.$$

We get therefore

$$id(f_{ab} + f_{bc} + f_{ca}) = g_{abc}^{-1} dg_{abc}.$$

Since U_{abc} is contractible, we can define $\log g_{abc}$ by choosing a value at a point $p \in U_{abc}$ (i.e. a choice of branch) and then everything else is determined by analytic continuation. Thus choosing a branch we can re-write this as

$$d(if_{ab} + if_{bc} + if_{ca} - \log g_{abc}) = 0.$$

So we get a constants $c_{abc} \in \mathbb{R}$ (divided by i here). If we choose a different branch of log, then $\log g$ changes by $2\pi i$ times an integer, so the constants c_{abc} change by $2\pi\mathbb{Z}$. Thus although c_{abc} depends on the choice of branch of log, the equivalence classes $[c_{abc}/2\pi] \in \mathbb{R}/\mathbb{Z}$ do not.

Definition 3.2. The class of the Cech 2-cocycle $c_{abc}/2\pi$ in $H^2(M, \mathbb{R}/\mathbb{Z})$ is called the *holonomy* of the connection. (N.B. here $\mathbb{R}/\mathbb{Z}$ is the constant sheaf!)

Hitchin points out an elegant interpretation of this. Recall that a flat line bundle with connection has holonomy around a curve. A gerbe has holonomy around a *surface*. This is because for a closed connected orientable surface $S \subset M$ we have $H^2(S, \mathbb{R}/\mathbb{Z}) \simeq \mathbb{R}/\mathbb{Z}$ (constant sheaf, NOT S^1-valued functions). On the other hand, the pullback of G to S is zero, since S is 2-dimensional. So we get a holonomy $c_{abc}/2\pi \in \mathbb{R}/\mathbb{Z}$, an angle of holonomy around the surface!

For line bundles, if we have a flat connection AND the holonomy is trivial, then the line bundle is trivial (and we get a trivializing section by parallel transport). A similar result holds for gerbes. Suppose c_{abc} is a coboundary, so $c = \delta k$, where k_{ab} are constants in $2\pi\mathbb{R}/\mathbb{Z}$. Define

$$h_{ab} = \exp(if_{ab} - ik_{ab}),$$

(well-defined since exp handles the quotient by $\mathbb{Z}$). Then the equations above imply

$$h_{ab}h_{bc}h_{ca} = g_{abc}, \Rightarrow g = \delta h,$$

so g is a coboundary and the gerbe is trivial.

But more than that, we have a particularly nice trivialization given by h (it is the analog of a covariantly constant trivialization obtained by parallel transport in the line bundle case)—Hitchin calls the h that arrises in this way a *flat trivialization*. Now suppose we have two different flat trivializations h, h'. As before $t = h'/h$ defines a line bundle L. But it is not just any line bundle! Using the equations above, we have

$$\begin{aligned} iB_b - iB_a - iA_{ab} &= idf_{ab} = d\log h_{ab} \\ iB'_b - iB'_a - iA_{ab} &= idf'_{ab} = d\log h'_{ab} \end{aligned}$$

and subtracting these shows that if we define $C_a = B'_a - B_a$ then

$$iC_b - iC_a = g_{ab}^{-1}dg_{ab} \Rightarrow iC_b = iC_a + g_{ab}^{-1}dg_{ab}$$

which says the C_a are a collection of connection 1-forms on L. Moreover

$$dC_a = dB'_a - dB_a = F_a - F_a = 0,$$

so this connection is flat. In other words, the difference of two flat trivializations of a gerbe is a *flat line bundle*: this is strictly stronger than before (when we found that the difference between two ordinary trivializations is a line bundle), since e.g. the chern class of this line bundle in this case will be torsion. This shows the sense in which flat trivializations are particularly nice examples of trivializations.

4 Loop space picture

Suppose we pullback the gerbe $\mathcal{G}$ to a loop $\ell : S^1 \to M$ $(f \in LM)$. Since the loop is only 1-dimensional, the pullback gerbe $\ell^*\mathcal{G}$ is flat with trivial holonomy (Cech 2-cocycle c_{abc} will also become trivial!). This means that there will be flat trivializations of $\mathcal{G}|_\ell$. Moreover, we showed above that the space of flat trivializations is acted on transitively by the space of flat line bundles. A flat line bundle on S^1 is uniquely determined, up to gauge equivalence, by its holonomy, which is an element of $U(1)$.

Identify two flat trivializations if they differ by a flat line bundle with trivial holonomy (a flat line bundle on S^1 may still have a non-trivial connection; however in this case the only gauge invariant is the holonomy of the connection): let $\mathcal{Q}_\ell$ denote the resulting "moduli space of flat trivializations" of $\mathcal{G}|_\ell$. Summarizing, given a loop ℓ, we have a moduli space $\mathcal{Q}_\ell$ of flat trivializations, which carries a free and transitive action of $U(1)$ (the latter being viewed as the space of flat line bundles on ℓ). These glue together to give a principle $U(1)$ bundle $\mathcal{Q} \to LM$. So we've replaced the gerbe with a line bundle, by passing to the loop space!

And we get more structure. A path in the loop space $F : [0,1] \times S^1 \to M$ gives a cylinder in M. Since this is a 2-dimensional surface, the pullback of $\mathcal{G}$ under F is flat. Moreover, as $[0,1] \times S^1$ deformation retracts to S^1, the pullback gerbe also has trivial holonomy (for this latter fact, it is important to view this as a pullback to $[0,1] \times S^1$, since e.g. the image of F could well be a closed surface, e.g. S^2). So we can choose a flat trivialization for $F^*\mathcal{G}$. Restricting this flat trivialization to the two endpoints, we obtain flat trivializations on $\{0\} \times S^1$ and $\{1\} \times S^1$; this defines a mapping from the moduli space of flat trivializations over ℓ_0 to that over ℓ_1. In other words, we get a method of doing parallel transport in the bundle $\mathcal{Q}$ over loop space. The curvature of the resulting connection on $\mathcal{Q}$ is the transgression of the 3-curvature on M.

5 Bundle gerbes

There is a nice explanation of taking duals/products of $U(1)$ bundles on page 4 of Murray's notes.

Now is a good time to summarize some of the basic functorial properties and structures associated with gerbes. These are all highly analogous to line bundles.

1. We can take duals and products of gerbes. These operations make sense for Cech cocycles, and they also make sense in the transition line bundle picture.

2. Gerbes can be pulled back by smooth maps.

3. Associated to a gerbe $\mathcal{G}$ is a characteristic class $c(\mathcal{G})$, which pulls back like the gerbe, and satisfies

$$c(\mathcal{G}_1 \otimes \mathcal{G}_2) = c(\mathcal{G}_1) + c(\mathcal{G}_2), \qquad c(\mathcal{G}^*) = -c(\mathcal{G}).$$

4. There is a notion of connection (or "connective structure") on a gerbe. It has a curvature, which is a globally defined 3-form. This 3-form is a representative in de Rham cohomology for the image of $c(\mathcal{G})$ in cohomology with real coefficients.

5. A connection on a gerbe has holonomy around any 2-dimensional closed submanifold S.

The description of gerbes in terms of Cech cocycles for a good cover was highly local. We moved somewhat away from this in the Hitchin-Chatterjee picture, which allowed some of the non-trivial data of the gerbe to be packaged in not necessarily local things: the transition line bundles, which could be non-trivial bundles over non-contractible sets. The bundle gerbe picture moves further in this direction.

5.1 Surjective submersions

We will want to consider surjective submersions $\pi : Y \to M$. There are two main examples: fibre bundles and open covers. Given an open cover $\mathcal{U} = \{U_a\}$, write $Y_{\mathcal{U}}$ for the disjoint union. Write $Y^{(p)}$ (or simply Y^p for brevity) for the p-fold fibre product. There are maps $\pi_i : Y^p \to Y^{p-1}$ for $i = 1, ..., p$. In the case $Y_{\mathcal{U}}$, note that these spaces are disjoint unions of all *ordered* double intersections, triple intersections, etc. For example, for two sets, Y^2 is a disjoint union of 4 double intersections. Thus restricting some form G to each U_a becomes simply $\pi^* G$, and so on.

A basic property of surjective submersions is that they admit local sections. What this means is that we can always find an open cover $\mathcal{U} = \{U_a\}$ and local sections $s_a : U_a \to Y$ of the submersion π. These can be bundled into a map $s : Y_{\mathcal{U}} \to Y$, and this type of thing will always allow you to get back to the local description of the Hitchin-Chatterjee theory. Surjective submersions to M form a category, with morphisms being what you think they are. One nice feature is that "refinements of open covers" are special cases of morphisms in this category.

Define

$$\delta : \Omega(Y^p) \to \Omega(Y^{p+1}), \qquad \delta = \sum_{i=1}^{p} (-1)^{i-1} \pi_i^*.$$

Note that in the case $Y_{\mathcal{U}}$ this is the Cech differential on sheaves of forms. Also

$$0 \to \Omega(M) \to \Omega(Y) \to \Omega(Y^2) \to \cdots$$

is the Cech resolution of the sheaf. It is acyclic by partitions of unity. And this holds for general Y in fact!

More generally for a function $g : Y^{p-1} \to H$ where H is an abelian group, δg makes sense. And for a bundle $P \to Y^{p-1}$ we define a bundle $\delta(P) \to Y^p$ as the alternating tensor product of pullbacks. It is immediate that $\delta\delta P$ is canonically trivial.

5.2 Definition of bundle gerbes

We simply replace $Y_{\mathcal{U}}$ with a more general surjective submersion!

Definition 5.1. A *bundle gerbe* over M is a pair (L, Y) where $Y \to M$ is a surjective submersion and $L \to Y^2$ is a hermitian line bundle, satisfying:

1. There is a chosen unitary trivializing section θ of $\delta(L) \to Y^3$.
2. $\delta(\theta) = 1$.

Taking $Y = Y_{\mathcal{U}}$ gives back Hitchin-Chatterjee. This can be packaged slightly differently: let $Q \to Y^2$ be the associated unit circle bundle. Then (1) is equivalent to the existence of an isomorphism

$$m : \pi_3^*(Q) \otimes \pi_1^*(Q) \to \pi_2^*(Q).$$

This map is called the *bundle gerbe multiplication.* The reason for this, and the reason why this is nice, is that (2) then says that m is associative. Explicity, we take θ^*, a unitary section of $\delta(Q)^*$; by duality, it pairs naturally with a section of $\pi_3^*(Q) \otimes \pi_1^*(Q)$, with output in $\pi_2^*(Q)$. This is essentially the way we've been thinking about "products" in hLAB recently, as homomorphisms in $\mathrm{Hom}(A^{\otimes 2}, A)$. Viewed in these terms, I believe that $\delta\theta$ is the *associator.*

(Note from Stevenson, 2000.) Over a point (y_1, y_2, y_3) the section θ is defined by taking elements $u \in P_{(y_1,y_2)}$ and $v \in P_{(y_2,y_3)}$ and setting

$$\theta_{(y_1,y_2,y_3)} = u \otimes m(u \otimes v)^* \otimes v \in P_{(y_1,y_2)} \otimes P^*_{(y_1,y_3)} P_{(y_2,y_3)}.$$

From this we see that θ^* indeed implements the bundle gerbe multiplication. To see that $\delta\theta$ being the associator is plausible, note that over a point (y_1, y_2, y_3, y_4), $\delta\theta$ lives in the product of 12 copies $P_{(y,z)}$ where y, z run through all possible ordered pairs. Using $P_{(y,z)} = P^*_{(z,y)}$, these cancel in pairs and give the trivial bundle. Also, we can view $\delta\theta$ as performing the product $((ab)c)((c^{-1}b^{-1})a^{-1})$. This matches numerically: 4 dual bundles are used to eat the elements used in the first two products ab and $c^{-1}b^{-1}$, with the results stored in 2 bundles (6 so far). Then to form the next products (with c and with a^{-1} respectively), another 4 dual bundles are used to eat the elements, and then the results stored in the remaining 2 bundles (no up to 12 bundles used). The remaining two bundles where the results are stored are dual to each other (as follows because we used all inverses for one triple product), and they can be paired to give an element of $U(1)$. This element is the associator; the phase relating the two ways of doing the product. Saying that it is 1, says that the product is associative.

5.3 Sheaf of groupoids

For each point $m \in M$, we get a groupoid as follows. The objects are elements of the fibre Y_m, and the morphisms between $y_1, y_2 \in Y_m$ are $P_{(y_1,y_2)}$. The bundle gerbe multiplication gives composition. Note that this groupoid is transitive, and that the group of morphisms of a point is $U(1)$. So we get a bundle of $U(1)$-groupoids over M.

5.4 Reformulation of functorial properties

1. Pullbacks: $f : N \to M$. Take f^*Y, then there is a map $f : f^*Y \to Y$ covering f, and so a map $f^2 : (f^*Y)^2 \to Y^2$ which we use to pullback P.

2. Dual: just take P^*.

3. Products: given (P, Y), (Q, X) take the fibre product of Y, X, and the tensor product $P \otimes Q$.

4. Characteristic class. Choose a good cover $\mathcal{U}$, sufficiently small that we have sections $s_a : U_a \to Y$. Then we get sections $s_{ab} = (s_a, s_b) : U_{ab} \to Y^2$. Let $\tau_{ab} : U_{ab} \to P_{ab} = s_{ab}^* P$ be a section (exists as the cover is good). Likewise put $\theta_{abc} = s_{abc}^* \theta$. On triple overlaps we have

$$\tau_{ab}\tau_{bc}\tau_{ca} = g_{abc}\theta_{abc},$$

   for some $g_{abc} : U_{abc} \to U(1)$. This is the characteristic class. (Note: the left-hand-side here really means $(\pi_c^*\tau_{ab}) \otimes (\pi_a^*\tau_{bc}) \cdots$.)

5. Connections: simply a hermitian connection ∇ on $L \to Y^2$, such that θ is flat for $\delta(\nabla)$ (where here $\delta(\nabla)$ denotes the connection on $\delta(L)$ obtained by pullbacks and by the usual prescription for tensor products of bundles). Equivalently, m, ∇ satisfy a Leibniz rule. This immediately gives back Hitchin-Chatterjee for $Y = Y_{\mathcal{U}}$. As before there is a (non-unique) form F on Y such that $\delta(F)$ is the curvature of ∇ (this is the "curving" or "B-field"), and dF is basic, and so $dF = \pi^* G$ for some 3-form G on the base, the 3-curvature. Similar to above, we can use local sections s_a for a good cover, to obtain equations as before.

Definition 5.2. A bundle gerbe is *trivial* if there is a hermitian line bundle $R \to Y$ such that $(L, Y) \simeq (\delta(R), Y)$. The choice of R and of the isomorphism is called a trivialization. Note here that $\delta\delta(R)$ is canonically trivial, so we are taking $\theta = 1$ for the bundle gerbe $\delta(R)$. Say that two bundle gerbes are stably isomorphic if $\mathcal{G}_1^* \otimes \mathcal{G}_2$ is trivial. Then turns out that stable isomorphism classes are exactly classified by the characteristic class.

5.5 Lifting bundle gerbes

One nice feature of bundle gerbes is that they give a nice construction of the lifting gerbe considered before. We consider a central extension, as before, and a principal G-bundle $Y \to M$—this *will* be our surjective submersion! (Notice the way it is global now: we are far from a good cover!) There is a natural map $f : Y^2 \to G$ defined by $y_2 = y_1 f(y_1, y_2)$. The central extension $\tilde{G}$ is, in particular, a principle $U(1)$-bundle over G. We simply pull it back using the map f to get the $U(1)$ bundle $P \to Y^2$. (If you want a line bundle, then take the associated line bundle.) Given $g \in G$, let $\tilde{G}_g$ denote the fibre of $\tilde{G}$ over g. Then the group multiplication on G respects the fibres in the natural way (this is just because $\tilde{G} \to G$ is a group homomorphism). Now let (y_1, y_2, y_3) be in Y^3. We seek a bundle gerbe multiplication, that is, a map

$$m : Y_{(y_1,y_2)} \otimes Y_{(y_2,y_3)} \to Y_{(y_1,y_3)}$$

satisfying associativity. But, if we think about how pullback bundles are defined, $Y_{(y,z)} = \tilde{G}_{f(y,z)}$ (i.e. the fibre over (y, z) *is* the fibre of $\tilde{G}$ over the image point). So the multiplication above is already defined by the group multiplication on $\tilde{G}$, and associativity follows from associativity of multiplication on $\tilde{G}$!

We now give some examples of this beautiful construction. In addition to spin-c structures that we mentioned before, we also have:

1. $PU(H)$ bundles. Let H be a Hilbert space, and $Y \to M$ a principal $PU(H)$ bundle. Then we have a central extension $U(H) \to PU(H)$ and so we get a gerbe. The characteristic class is called the *Dixmier-Douady class*. This was, I believe, the earliest (?) "geometric realization" of degree-3 integral cohomology.

2. The basic gerbe over a compact, connected, simple Lie group G. Under these assumptions it is known that $H^3(G, \mathbb{Z}) = \mathbb{Z}$. (In de Rham, this is generated by the Cartan 3-form.) We want to get the gerbe representing the generator. Let PG denote the space of based paths, and $PG \to G$ the path fibration (given by the evaluation-at-1 map). The kernel of the evaluation map is ΩG, so $PG \to G$ is a principal ΩG bundle (check: the endpoint of the path doesn't change because loop in ΩG ends at the identity). There is a basic central extension,

$$1 \to U(1) \to \tilde{\Omega G} \to \Omega G \to 1.$$

The lifting bundle gerbe for this is what we seek!

5.6 One application: WZW term

Let Σ be a compact Riemann surface. There is a theory whose configuration space consists of maps $g : \Sigma \to G$ with G a compact simple Lie group. The action has a kinetic term and a more interesting term which was defined by Witten as follows: choose a 3-manifold X with $\partial X = \sigma$ and assume we can extend g to X. Then the interesting interaction term is

$$\int_X g^* \eta.$$

This is dependent on the choice of extension. However, choosing different extensions changes the result by an integer (using integrality of η). So in the quantum theory (after we take $e^{2\pi i S}$) it disappears. Using gerbes we can instead think of this as the holonomy of the basic gerbe around the surface Σ! This eliminates X from the discussion (and we no longer have to make the additional assumption on g).