Lecture 1: de Rham cohomology of real manifolds

Dustin Clausen

July 22, 2021

1 de Rham's theorem

This class is about algebraic de Rham cohomology, but for motivation and historical framing we'll start with de Rham cohomology of (smooth real) manifolds. First, recall the *de Rham complex* of a manifold M, denoted

$$\Omega^\bullet(M) = \left[\Omega^0(M) \xrightarrow{d} \Omega^1(M) \xrightarrow{d} \Omega^2(M) \to \ldots\right].$$

It has the following properties, which characterize it as a functor of M:

1. $\Omega^\bullet(M)$ is a commutative differential graded algebra, contravariantly functorial in the manifold M;
2. $\Omega^0(M)$ identifies with the ring $C^\infty(M)$ of smooth real-valued functions on M, functorially in M;
3. For arbitrary M and $k \geq 0$, the presheaf $U \mapsto \Omega^k(U)$ on open subsets of M is a sheaf;
4. If U is an open subset of $\mathbb{R}^n$ with coordinate functions $x_1, \ldots, x_n \in C^\infty(U)$, then:
   (a) As a $C^\infty(U)$-module, $\Omega^1(M)$ is free of rank n with basis $dx_1, \ldots, dx_n$;
   (b) The multiplication map $\Lambda^k_{C^\infty(U)}\Omega^1(U) \to \Omega^k(U)$ is an isomorphism; in particular the latter is free of rank $\binom{n}{k}$ with basis the $dx_{i_1}\ldots dx_{i_k}$ for $1 \leq i_1 < \ldots < i_k \leq n$ and vanishes for $k > n$;
   (c) For $f \in C^\infty(U)$ we have $df = (\partial_1 f)dx_1 + \ldots + (\partial_n f)dx_n$.

It's clear from these properties that the differentials in the complex are completely determined by the first differential $C^\infty(M) \to \Omega^1(M)$ by the Leibniz rule. The essential content in the assertion $d^2 = 0$ is the theorem on equality of mixed partial derivatives. The essential content in the fact that the contravariant functoriality of $C^\infty(M)$ extends to the CDGA $\Omega^\bullet(M)$ (in particular ensuring the independence of coordinates in the above description) is the chain rule for multivariable functions. This de Rham complex somehow packages these fundamental theorems of calculus in several variables into the compact algebraic structure of a CDGA.

Write $H^k_{dR}(M)$ for the k^{th} cohomology group of the complex $\Omega^\bullet(M)$; this is *de Rham cohomology*. The main theorem of this lecture states that de Rham cohomology identifies with the "usual" cohomology of the underlying topological space M.

Theorem 1. *Let M be a smooth manifold. Then there is a canonical and functorial isomorphism*

$$H^k_{dR}(M) \simeq H^k(M; \underline{\mathbb{R}}),$$

where the right hand side denotes the sheaf cohomology of M with coefficients in the constant sheaf with values the (discrete) abelian group $\mathbb{R}$.

To construct the isomorphism and make the statement more useful, note that the constant sheaf $\underline{\mathbb{R}}$, which by definition is the sheafification of the constant presheaf with values $\mathbb{R}$, is the same as the sheaf of locally constant functions with values in $\mathbb{R}$. As every locally constant function is smooth, this gives $\underline{\mathbb{R}} \subset \Omega^0$; in fact more precisely we have

$$\underline{\mathbb{R}} = \ker(d : \Omega^0 \to \Omega^1),$$

as a smooth function is locally constant if and only if all its partial derivatives vanish. In particular, this furnishes a comparison map of complexes of sheaves on M

$$\underline{\mathbb{R}} \to \Omega^\bullet,$$

where $\underline{\mathbb{R}}$ is viewed as a complex concentrated in degree zero and $\Omega^\bullet$ is the complex of sheaves defined by $U \mapsto \Omega^\bullet(U)$. Then the de Rham theorem breaks up into the following two claims:

Proposition 2. *Let M be a smooth manifold. Then:*

1. *The comparison map $\underline{\mathbb{R}} \to \Omega^\bullet$ is a quasi-isomoprhism* of complexes of sheaves *on M, i.e. the sheafification of the presheaf $U \mapsto H^k_{dR}(U)$ is 0 for $k > 0$ and (as already mentioned above) identifies with $\underline{\mathbb{R}}$ for $k = 0$.*
2. *For $q > 0$ and $p \geq 0$, the sheaf cohomology group $H^q(M; \Omega^p)$ is 0.*

Before proving the proposition, let's explain why it implies the de Rham theorem. This is some bit of homological algebra and it can be organized in different ways. For now we'll give the most basic one, but at the end we'll use this as a springboard for a more thorough discussion of what's going on.

Lemma 3. *The previous proposition implies the de Rham theorem. More generally, suppose $\mathcal{A}$ is a sheaf of abelian groups on a topological space X and $\mathcal{A} \simeq \left[\mathcal{A}^0 \to \mathcal{A}^1 \to \ldots\right]$ is a quasi-isomorphism to a complexes of sheaves. If for each $p \geq 0$ and $q > 0$ we have $H^q(X; \mathcal{A}^p) = 0$, then*

$$H^*(X; \mathcal{A}) \simeq H^* \left[\mathcal{A}^0(X) \to \mathcal{A}^1(X) \to \ldots\right].$$

Proof. Recall that $H^q(X; \mathcal{A})$ is defined as the q^{th} right derived functor of the left exact global sections functor from abelian sheaves on M to abelian groups. It is therefore computed by choosing an injective resolution of $\mathcal{A}$, passing to global sections, then taking cohomology of the resulting complex of abelian groups. However, it is a general result in homological algebra that, for computing right derived functors of left exact functors, one need not use anything so strong as an injective resolution; it suffices to have an acyclic resolution. □

Now we prove the proposition. The two halves are completely separate. For the second one, here is a more general claim.

Lemma 4. *Let $\mathcal{F}$ be any sheaf on M which admits the structure of a module sheaf over the sheaf of rings $U \mapsto C^\infty(U)$. Then $H^q(M;\mathcal{F}) = 0$ for $q > 0$.*

Proof. First we show that the global sections functor for sheaves of C^∞-modules is exact. Since left exactness is clear, this means we need to see that if $f : \mathcal{A} \to \mathcal{B}$ is a surjection of C^∞-module sheaves, then $\mathcal{A}(M) \to \mathcal{B}(M)$ is surjective. The key will be that C^∞-functions have partition of unity. Suppose given $b \in \mathcal{B}(M)$. Then by surjectivity of f, there is an open cover $\{U_i\}_{i\in I}$ of M and sections $a_i \in \mathcal{A}(U_i)$ with $f(a_i) = b|_{U_i}$. Since M is paracompact, by refining we can assume the cover locally finite, and then we can choose a partition of unity $(\varphi_i \in C^\infty(M))_{i\in I}$ subordinate to the cover. Because $\mathrm{supp}(\varphi_i) \subset U_i$, the section $\varphi_i \cdot a_i \in \mathcal{A}(U_i)$ extends uniquely to $\mathcal{A}(M)$ requiring it to be 0 outside the support of φ_i. Moreover the sum $\sum_i \varphi_i a_i$ is locally finite and glues and is therefore globally well-defined, and similarly we justify the manipulations

$$f(\sum_i \varphi_i a_i) = \sum_i \varphi_i f(a_i) = (\sum_i \varphi_i) b = b$$

again by working locally and using local finiteness of the cover.

This shows that if we define cohomology of C^∞-module sheaves as the derived functors of the global sections functor, then the higher cohomology vanishes. Thus we need to see that C^∞-module cohomology identifies with the usual cohomology of the underlying abelian group sheaf. However, this is a standard lemma valid for any sheaf of rings replacing C^∞: one checks that an injective C^∞-module sheaf is flasque, and that a flasque sheaf of abelian groups is acyclic. □

Now for part 1 of the proposition, since there is a basis for the topology of any manifold consisting of open subsets diffeomorphic to $\mathbb{R}^d$, it suffices to show the following *Poincaré Lemma*:

Lemma 5. *We have $H^k_{dR}(\mathbb{R}^d) = 0$ for $k > 0$ and $\mathbb{R}$ for $k = 0$.*

Proof. More to ease notation than anything else, let's use a bit of functional analytic machinery to reduce to the case $d = 1$. Recall that $C^\infty(\mathbb{R}^d)$, endowed with the topology of uniform convergence of all derivatives on all compact subsets, is a Frechet space, and

$$C^\infty(\mathbb{R}^{d+e}) = C^\infty(\mathbb{R}^d)\widehat{\otimes}_{\mathbb{R}} C^\infty(\mathbb{R}^e)$$

for the projective tensor product of Frechet spaces. (In fact, these are *nuclear* Frechet spaces, so the answer is the same for any reasonable completed tensor product of Frechet spaces, of which there are many.) Endowing each $\Omega^q(\mathbb{R}^d)$ with the induced Frechet structure as a finite direct sum of copies of $C^\infty(\mathbb{R}^d)$, then the de Rham complex is a complex of Frechet spaces, and it follows that the same Künneth statement holds for the de Rham complexes:

$$\Omega^\bullet(\mathbb{R}^{d+e}) = \Omega^\bullet(\mathbb{R}^d)\widehat{\otimes}_{\mathbb{R}} \Omega^\bullet(\mathbb{R}^e).$$

Considering Frechet spaces as an additive category, the tensor product is a bi-additive functor. Thus it suffices to show that $\Omega^\bullet(\mathbb{R})$ is the direct sum of $\mathbb{R}[0]$ and a complex which is chain null-homotopic. But indeed $\Omega^\bullet(\mathbb{R}) = \left[C^\infty(\mathbb{R}) \overset{d/dx}{\to} C^\infty(\mathbb{R})\right]$ is the direct sum of its subcomplexes $[\mathbb{R} \to 0]$ and $[C^\infty(\mathbb{R})_0 \to C^\infty(\mathbb{R})]$ where $C^\infty(\mathbb{R})_0 = \{f \in C^\infty(\mathbb{R}) : f(0) = 0\}$, and the latter subcomplex is null-homotopic because the differential is an isomoprhism with inverse $g \mapsto G(x) = \int_0^x g(t)dt$ by the fundamental theorem of calculus. (By the open mapping theorem, we don't even need to worry about the inverse map being continuous, though it's also easy to check.) □

Thus we've proved that the de Rham cohomology agrees with the sheaf cohomology with $\mathbb{R}$-coefficients. In the second half of this lecture, I'd like to revisit the homological algebra fact that sheaf cohomology can be computed by an arbitrary acyclic resolution, and view it from increasingly more and more modern perspectives. This will prepare us for when we move to complex manifolds.

2 Some homological algebra of sheaves

We used the fact that sheaf cohomology $H^*(X;\mathcal{A})$ can be calculated by means of any *acyclic* resolution $\mathcal{A} \simeq \mathcal{A}^\bullet$. However, the proof we gave of this fact doesn't tell us what would happen if we had an arbitrary resolution, where the acyclicity of the $\mathcal{A}^p$ failed. But one should get good information even without the acyclicity, and this will become important for us later. To give a very simple example, a short exact sequence of sheaves

$$0 \to \mathcal{A} \to \mathcal{B} \to \mathcal{C} \to 0$$

can be reinterpreted as a quasi-isomorphism of complexes $\mathcal{A} \simeq [\mathcal{B} \to \mathcal{C}]$ or in other words a two-term resolution of $\mathcal{A}$, and we know in this case, without any hypothesis on $\mathcal{B}$ and $\mathcal{C}$, that there is a long exact sequence on cohomology, which at least tells us something about the cohomology of $\mathcal{A}$ even if it doesn't completely determine it.

So what is the statement which generalizes the long exact sequence in cohomology to the case of a resolution of length longer than 2?

Proposition 6. *If $\mathcal{A}$ is a sheaf and $\mathcal{A} \simeq \left[\mathcal{A}^0 \to \mathcal{A}^1 \to \dots\right]$ is a resolution of $\mathcal{A}$, then there is a convergent spectral sequence*

$$E_1^{p,q} = H^q(M;\mathcal{A}^p) \Rightarrow H^{p+q}(M;\mathcal{A})$$

where the d_1 differential is induced by functoriality from the differential on our complex of sheaves.

If the complex has two terms there can be no differentials after d_1 for degree reasons and it follows that this spectral sequence exactly reproduces the long exact sequence (check and see why if you've never done this before — it's not immediately clear!). On the other hand if each $\mathcal{A}^p$ is acyclic then the E_1-page lives entirely on the $p = 0$ line and we get the conclusion

$$H^*(M;\mathcal{A}) = H^*\left[\mathcal{A}^\bullet(M)\right].$$

So those are two nice consequences. Now let's give the proof.

Proof. Probably the best way to organize this is via the theory of *hypercohomology*, which generalizes sheaf cohomology from the case of a single sheaf to the case of a complex of sheaves. The hypercohomology of a (cohomologically bounded below) complex of sheaves is calculated by choosing a quasi-isomorphism to a (cohomologically bounded below) complex of injective sheaves, passing to global sections, then taking cohomology of the resulting complex of abelian groups. It is well-defined, functorial, and invariant under quasi-isomorphism, for the same reasons as usual sheaf cohomology is: essentially, the well-definedness of injective resolutions up to chain homotopy.

In particular, the cohomology of $\mathcal{A}$ identifies with the hypercohomology of any resolution $\mathcal{A}^\bullet$. Now we can actually drop the condition that $\mathcal{A}^\bullet$ is quasi-isomorphic to a single sheaf and work with a general cochain complex of sheaves: we claim that there is a spectral sequence of the same form

$$E_1^{p,q} = H^q(M;\mathcal{A}^p) \Rightarrow H^{p+q}(M;\mathcal{A}^\bullet)$$

abutting to the hypercohomology.

The explantation for this is that any cochain complex carries a canonical decreasing filtration, the so-called *brutal filtration*, with

$$F^{\geq p}\mathcal{A}^\bullet = \left[0 \to \dots \to 0 \to \mathcal{A}^p \to \mathcal{A}^{p+1} \to \dots\right].$$

The p^{th} associated graded $F^{\geq p}/F^{\geq p+1}$ for this filtration is $\mathcal{A}^p[-p]$, which has hypercohomology

$$H^*(X; \mathcal{A}^p[-p]) = H^{*-p}(X; \mathcal{A}^p).$$

Moreover the hypercohomology of $F^{\geq p}\mathcal{A}^\bullet$ vanishes in cohomological degrees $< p$ because the complex itself only starts in degree p. If we lived in the **fantasy world** where there was a chain-complex valued exact functor $R\Gamma(X; -)$ calculating hypercohomology, then we would deduce the existence of a filtered complex $F^{\geq \cdot}R\Gamma(X; \mathcal{A}^\bullet)$ where the underlying complex calculates hypercohomology of $\mathcal{A}^\bullet$, the associated graded calculates the hypercohomology of $\mathcal{A}^p[-p]$, and the p^{th} filtered piece calculates the hypercohomology of $F^{\geq p}\mathcal{A}^\bullet$. Then the spectral sequence of a filtered complex will give us what we want; the vanishing in a range of the p^{th} filtered piece ensures convergence.

Now, there is no such functorial chain complex calculating the hypercohomology. However, we can argue around this by a construction known as a *Cartan-Eilenberg resolution* of the complex $\mathcal{A}^\bullet$. This is in particular a double complex of injective sheaves such that each row gives an injective resolution of $\mathcal{A}^p$. The total complex gets a filtration in the usual way filtering by rows, and it satisfies our needs after passing to global sections. □

Remark 7. *There are actually two useful filtrations on an arbitrary cochain complex $\mathcal{A}^\bullet$. Besides the brutal filtration used above, there is also the* canonical filtration, *an increasing filtration $\tau^{\leq \cdot}\mathcal{A}^\bullet$ with*

$$\tau^{\leq d}\mathcal{A}^\bullet = \left[\mathcal{A}^0 \to \mathcal{A}^1 \to \dots \to \mathcal{A}^{d-1} \to \ker(d) \to 0 \to 0 \to \dots\right].$$

This definition is arranged so that the natural map $\tau^{\leq d}\mathcal{A}^\bullet \to \mathcal{A}^\bullet$ induces an isomorphism on cohomology groups in degrees $\leq d$ but the source has vanishing cohomology in degrees $> d$.

There are two major differences between these two filtrations:

1. *The brutal filtration is a convergent* decreasing *filtration: the complex is recovered as the* inverse limit *of the $\mathcal{A}^\bullet/F^{\geq p}\mathcal{A}^\bullet$. The canonical filtration is a convergent* increasing *filtration: the complex is recovered as the* colimit *of the $\tau^{\leq d}\mathcal{A}^\bullet$.*

2. *Both filtrations are functorial in maps of chain complexes. But only the canonical filtration sends quasi-isomorphisms to quasi-isomoprhisms. If you have an object of the derived category but not an honest chain complex representing it, you have access to the canonical filtration but not the brutal filtration.*

If you're like me, you find all of the business with injective resolutions in the previous proofs very opaque and unsatisfying. It just doesn't feel right to hack your way out of functoriality issues such as arose in the construction of the spectral sequence associated to the brutal truncation. Also, important constructions such as sheaf cohomology should be prima facie canonical, and their intuitive meaning should be transparent. Neither aspect is very clear from the discussion in terms of injective resolutions.

Fortunately, the modern theory of ∞-categories fixes these problems and, for functoriality issues such as arose in the construction of the spectral sequence, gives an honest and true replacement for the false idea that hypercohomology is computed by an exact functor to chain complexes. We'll discuss this in the next lecture, but let's give a teaser now. There is an ∞-category $D(\mathbb{Z})$ called the derived ∞-category of abelian groups. It admits a canonical functor from the ordinary category of chain complexes:

$$Ch_{\mathbb{Z}} \to D(\mathbb{Z}),$$

and its homotopy category identifies with the usual derived category of abelian groups (the localization of $Ch_{\mathbb{Z}}$ at the quasi-isomorphisms). We denote this functor by

$$C \mapsto |C|.$$

In particular, if you have a cohomologically bounded below cochain complex of sheaves $\mathcal{A}^\bullet$ on a topological space X, you can view it a sheaf with values in $Ch_{\mathbb{Z}}$, whence by composition a *presheaf* $|\mathcal{A}^\bullet|$ with values in $D(\mathbb{Z})$, namely

$$|\mathcal{A}^\bullet|(U) = |\mathcal{A}^\bullet(U)|.$$

Now, the whole point is that the functor $Ch_{\mathbb{Z}} \to D(\mathbb{Z})$ does not preserve limits, so there's no reason for this to be a sheaf. Thus it makes sense to want to sheafify it. And that's exactly what injective resolutions accomplish, back up at the level of chain complexes:

Theorem 8. *Suppose $\mathcal{A}^\bullet \to \mathcal{I}^\bullet$ is an injective resolution of $\mathcal{A}^\bullet$: a quasi-isomoprhism to a complex of injective sheaves. Then the induced map of presheaves with values in $D(\mathbb{Z})$*

$$|\mathcal{A}^\bullet| \to |\mathcal{I}^\bullet|$$

identifies $|\mathcal{I}^\bullet|$ with the sheafification of $|\mathcal{A}^\bullet|$.

So this is what these injective resolutions are really doing: they're giving a non-canonical model for the canonical operation of sheafification of presheaves with values in $D(\mathbb{Z})$.

In particular, the cohomology of a sheaf is "just" the global sections of the sheafification of the induced presheaf with values in $D(\mathbb{Z})$. Since sheafification is defined by a universal property, it is perfectly functorial and we can produce spectral sequences without resorting to any clever resolutions by double complexes. We can in fact entirely work in the world of sheaves and presheaves with values in $D(\mathbb{Z})$ and avoid discussion of injective resolutions altogether.

Exercise. Let M be a manifold and $N \subset M$ a closed submanifold. Define a subcomplex $\Omega^\bullet(M;N) \subset \Omega^\bullet(M)$ of the de Rham complex of M, consisting of differential forms whose coefficients in local coordinates satisfy the property that they, and all their partial derivatives of all orders, vanish on N. Establish a long exact sequence of the form

$$\ldots \to H^k(\Omega^\bullet(M;N)) \to H^k(M;\underline{\mathbb{R}}) \to H^k(N;\underline{\mathbb{R}}) \to H^{k+1}(\Omega^\bullet(M;N)) \to \ldots$$

by identifying $H^k(\Omega^\bullet(M;N))$ with the sheaf cohomology of the extension by zero of the constant sheaf $\underline{\mathbb{R}}$ on the complement $M \setminus N$.