Lecture 3: Interlude on sheaves with values in $D(\mathbb{Z})$

Dustin Clausen

May 6, 2021

Let's finish up our discussion of abstract matters. We want to work in the ∞-category of sheaves with values in $D(\mathbb{Z})$ instead of always having to choose injective resolutions in our proofs. The definition of a sheaf with values in $D(\mathbb{Z})$ is exactly the same as that of a sheaf with values in Ab, at least provided you used the correct definition of a sheaf with values in Ab, namely the one based on sieves.

Definition 1. *Let X be a topological space. A* sieve *on X is a set $\mathfrak{U}$ of open subsets of X such that if $V \in \mathfrak{U}$ and $V' \subset V$, then $V' \in \mathfrak{U}$. If*

$$U = \cup_{V \in \mathfrak{U}} V,$$

we say that the sieve $\mathfrak{U}$ covers U.

If you have an open cover of U, you can make a sieve $\mathfrak{U}$ which covers U by saying that $V \in \mathfrak{U}$ means that V lies in some element of the cover. This is called the sieve generated by the cover. Obviously, all sieves are of this form. The picture is that the elements of the cover are like the holes our other open sets have to fall through, as flour through a sieve.

We will always view a sieve as a poset under inclusion.

Definition 2. *Let X be a topological space, and let $\mathcal{F} \in PSh(X, D(\mathbb{Z}))$ be a presheaf with values in $D(\mathbb{Z})$, meaning an element of the functor ∞-category $Fun(Open(X)^{op}, D(\mathbb{Z}))$. We say that $\mathcal{F}$ is a* sheaf *if for all sieves $\mathfrak{U}$ on X covering $U \in Open(X)$,*

$$\mathcal{F}(U) \xrightarrow{\sim} \varprojlim_{V \in \mathfrak{U}^{op}} \mathcal{F}(V).$$

Informally, the sheaf condition says that giving a section over U is the same as giving compatible sections over all elements of the covering sieve $\mathfrak{U}$. If $\mathfrak{U}$ is generated by a cover $\{U_i\}_{i \in I}$ and our sheaves are taking values in an ordinary category instead, say Ab, then it is elementary to see that $\varprojlim_{V \in \mathfrak{U}^{op}} \mathcal{F}(V)$ identifies with the subset of $\prod_{i \in I} \mathcal{F}(U_i)$ consisting of those local elements which agree on intersections. What is the ∞-analog?

To answer this, and for many other purposes, it's helpful to introduce the presheaf $h_U \in PSh(X, \mathcal{S})$ for $U \in Open(X)$, defined by $h_U(V) = *$ for $V \subset U$ and $\varnothing$ otherwise. In other words, this is the Yoneda image of U. The Yoneda lemma shows that $\mathrm{Map}(h_U, \mathcal{F}) = \mathcal{F}(U)$ for all $\mathcal{F} \in PSh(X, \mathcal{S})$.

A sieve $\mathfrak{U}$ can similarly be encoded in terms of its associated presheaf $h_{\mathfrak{U}}$ given by $h_{\mathfrak{U}}(V) = *$ for $V \in \mathfrak{U}$ and $\varnothing$ otherwise. We have a natural map $h_{\mathfrak{U}} \to h_U$, in fact an inclusion. As with any presheaf, there are many different ways of "presenting" $h_{\mathfrak{U}}$ in terms of the h_V.

Lemma 3. *In $PSh(X,\mathcal{S})$, we have*

$$h_{\mathfrak{U}} = \varinjlim_{V\in\mathfrak{U}} h_V.$$

If $\mathfrak{U}$ is the sieve generated by a cover $\{U_i\}_{i\in I}$ of U, we have

$$h_{\mathfrak{U}} = \varinjlim_{[n]\in\Delta^{op}} \coprod_{f:[n]\to I} h_{U_f},$$

where $U_f = U_{f_0}\cap\ldots\cap U_{f_n}$, and also

$$h_{\mathfrak{U}} = \varinjlim_{V\in\mathfrak{U}'} h_V$$

where $\mathfrak{U}'\subset\mathfrak{U}$ is the subposet consisting of intersections $\cap_{i\in J}U_i$ for all non-empty finite subsets $J\subset I$.

Proof. In all cases the idea is the same: colimits in a functor category are computed objectwise. So it suffices to evaluate on an arbitrary $W\in Open(X)$ and check a colimit in $\mathcal{S}$. When $W\notin\mathfrak{U}$, all terms involved are $\varnothing$ and we just note that an arbitrary colimit of $\varnothing$ is $\varnothing$. When $W\in\mathfrak{U}$ we have $h_{\mathfrak{U}}(W) = *$, and to match this up with what we see on the right hand side we'll use the criterion that a colimit of copies of $*$ is still $*$ provided the indexing diagram is contractible.

In the first case we thus need to see that the poset of elements of $\mathfrak{U}$ containing W is contractible; this is true because there is a terminal object, namely W. In the third case we need to see that the poset of elements of $\mathfrak{U}'$ containing W is $*$; this is true because that poset has all non-empty finite intersections. Any such poset is contractible, because we can write it as a filtered colimit of finite subposets, and any finite subposet includes into another finite poset by a null-homotopic map because we can just add the intersection to the poset. Finally, in the second case, let $J\subset I$ be the set of $i\in I$ with $W\subset U_i$. The simplicial set we're taking a colimit over identifies with the nerve of the category with object set J and a unique morphism between any two objects. This category is equivalent to $*$, so the simplicial set is simplicially contractible so the colimit in $\mathcal{S}$ is $*$ as desired. □

By the first claim, the sheaf condition on a presheaf of spaces is equivalent to requiring $Map(h_U,\mathcal{F})\overset{\sim}{\to} Map(h_{\mathfrak{U}},\mathcal{F})$ whenever $\mathfrak{U}$ is a covering sieve of U. But then by the other conditions this can also be rephrased more "concretely" if we choose an open cover generating $\mathfrak{U}$. For example if $\{V,V'\}$ generates, then by part 3 the sheaf condition simply says

$$\mathcal{F}(U)\overset{\sim}{\to}\mathcal{F}(V)\times_{\mathcal{F}(V\cap V')}\mathcal{F}(V').$$

Actually, the general sheaf condition can pretty much be reduced to this very simple one.

Proposition 4. *Let $\mathcal{F}$ be a presheaf with values in $\mathcal{S}$ (or $D(\mathbb{Z})$ or any other ∞-category, by Yoneda). Then $\mathcal{F}$ is a sheaf if and only if the following conditions are satisfied:*

1. *$\mathcal{F}(\varnothing) = *$.*

2. *For any two open subsets V,V',*

$$\mathcal{F}(V\cup V')\overset{\sim}{\to}\mathcal{F}(V)\times_{\mathcal{F}(V\cap V')}\mathcal{F}(V').$$

3. *If an open subset U is the filtered union of a set $\mathcal{U}$ of open subsets, then viewing $\mathcal{U}$ as a poset under inclusion,*

$$\mathcal{F}(U)\overset{\sim}{\to}\varprojlim_{V\in\mathcal{U}^{op}}\mathcal{F}(V).$$

Proof. First suppose $\mathcal{F}$ is a sheaf. Point 1 follows from the fact that the empty sieve covers the empty set. Point 2 follows from point 3 of the previous lemma and point . To verify point 3, consider the sieve $\mathfrak{U}$ generated by $\mathcal{U}$. Then $\varinjlim_{V \in \mathcal{U}} h_V = h_{\mathfrak{U}}$ as we see by evaluating on an arbitrary open W. (In fact, this is a filtered union of subpresheaves of sets, and filtered unions of sets are also filtered colimits in $\mathcal{S}$ as homotopy groups commute with filtered colimits.) Mapping out to $\mathcal{F}$ and using the sheaf condition gives the claim.

Now suppose 1, 2, and 3 are satisfied, and let $\mathfrak{U}$ be a sieve covering U. Suppose the cover $\{U_i\}_{i \in I}$ generates $\mathfrak{U}$. We can write $\mathfrak{U}$ as the filtered union of the sieves $\mathfrak{U}_J$ generated by the finite subsets $J \subset I$. Thus $\varinjlim_J h_{\mathfrak{U}_J} \tilde{\to} h_{\mathfrak{U}}$. This iso extends to a commutative diagram with $\varinjlim_J h_{\cup_{i \in J} U_i} \to h_U$, which is an iso on mapping out to $\mathcal{F}$ by part 1. Thus we reduce to checking the sheaf condition in the case of a sieve generated by a finite cover. Then similarly, working by induction on the number of elements in the cover, we reduce to either a cover with two elements or a cover with no elements where it amounts to conditions 1 and 2 by part 3 of the previous lemma. □

Now let's get back to $D(\mathbb{Z})$-valued sheaves. We can also define a presheaf with values in $D(\mathbb{Z})$ by

$$\mathbb{Z}[h_U](V) = \mathbb{Z}[h_U(V)],$$

so $\mathbb{Z}$ for $V \subset U$ and 0 otherwise, and similarly for $\mathbb{Z}[h_{\mathfrak{U}}]$. It follows that

$$Map(\mathbb{Z}[h_U], \mathcal{F}) = Map(\mathbb{Z}, \mathcal{F}(U))$$

and

$$Map(\mathbb{Z}[h_{\mathfrak{U}}], \mathcal{F}) = \varprojlim_{V \in \mathfrak{U}^{op}} Map(\mathbb{Z}, \mathcal{F}(V)) = Map(\mathbb{Z}, \varprojlim_{V \in \mathfrak{U}^{op}} \mathcal{F}(V)),$$

so the sheaf condition can be rephrased as saying that $Map(\mathcal{A}, \mathcal{F}) = 0$ for any $\mathcal{A}$ in the stable cocomplete full subcategory generated by $cofib(\mathbb{Z}[h_{\mathfrak{U}}] \to \mathbb{Z}[h_U])$. In fact we can make a much more refined claim.

Lemma 5. *The inclusion $Sh(X, D(\mathbb{Z})) \to PSh(X, D(\mathbb{Z}))$ admits a left adjoint $\mathcal{F} \mapsto \mathcal{F}^{sh}$. For a presheaf $\mathcal{F}$, we have $\mathcal{F}^{sh} = 0$ if and only if $\mathcal{F}$ lies in the stable co-complete subcategory generated by $cofib(\mathbb{Z}[h_{\mathfrak{U}}] \to \mathbb{Z}[h_U])$ for all sieves $\mathfrak{U}$ covering U.*

Proof. This follows from the theory of *presentable* ∞-categories as you can read about in Lurie's book "Higher topos theory". The ∞-category $D(\mathbb{Z})$ is presentable because it is compactly generated. Therefore so is $PSh(X, D(\mathbb{Z}))$ as functor categories to presentable ∞-categories are also presentable. For a presentable ∞-category $\mathcal{C}$ and any set S of arrows we can consider the full subcategory $\mathcal{C}[S^{-1}] \subset \mathcal{C}$ of those objects X with $Map(f, X)$ an iso for all $f \in S$, and the general theory says that $\mathcal{C}[S^{-1}]$ is presentable, the inclusion has a left adjoint, and the class of arrows inverted by this left adjoint identifies with the "strongly saturated class" generated by S. (Also, the left adjoint $\mathcal{C} \to \mathcal{C}[S^{-1}]$ does make $\mathcal{C}[S^{-1}]$ into the initial ∞-category with a functor from $\mathcal{C}$ inverting the arrows in S, but that's not relevant for us.)

In the stable case, by passing to cofibers, we can talk about killing a set of objects instead of inverting a set of arrows, and the strongly saturated class generated by the set of arrows corresponds simply to the smallest cocomplete stable subcategory generated by the class of objects. Putting things together the claim follows. □

Now we'd like to identify in more concrete terms this cocomplete stable subcategory of those $\mathcal{A} \in PSh(X; D(\mathbb{Z}))$ such that $\mathcal{A}^{sh} = 0$, at least under boundedness hypotheses. First, for an arbitrary $\mathcal{F} \in PSh(X; D(\mathbb{Z}))$ and $n \in \mathbb{Z}$, define

$$H_n(\mathcal{F}) \in PSh(X; Ab)$$

to be the presheaf of n^{th} homology groups. By sheafifying this presheaf of abelian groups, we can also consider the sheaf of abelian groups $H_n(\mathcal{F})^{sh} \in Sh(X; Ab)$.

Proposition 6. *Let $\mathcal{A} \in PSh(X; D(\mathbb{Z}))$.*

1. *If $\mathcal{A}^{sh} = 0$, then $H_n(\mathcal{F})^{sh} = 0$ for all $n \in \mathbb{Z}$.*
2. *If $\mathcal{A}$ is bounded above and $H_n(\mathcal{A})^{sh} = 0$ for all $n \in \mathbb{Z}$, then $\mathcal{A}^{sh} = 0$.*

Proof. From the previous proposition, the class of $\mathcal{A}$ as in 1 is generated under shifts and colimits by $cofib(\mathbb{Z}[h_{\mathfrak{U}}] \to \mathbb{Z}[h_U])$. Because homology groups have long exact sequences and commute with filtered colimits, and sheafifcation is an exact colimit preserving functor, to prove 1 it therefore suffices to check it for the generating acyclic $\mathcal{A} = cofib(\mathbb{Z}[h_{\mathfrak{U}}] \to \mathbb{Z}[h_U])$. But indeed if we restrict to any element of $\mathfrak{U}$ then $\mathbb{Z}[h_{\mathfrak{U}}] \to \mathbb{Z}[h_U]$ is the isomorphism $\mathbb{Z} = \mathbb{Z}$, whereas if we restrict to any element not contained in U it is the isomorphism $0 = 0$. Thus the map is a local isomorphism, so the map on sheafified homology groups is an isomorphism, as desired.

For property 2, by the lemma which follows we know that $\mathcal{A}^{sh}$ is also bounded above. Now we note that the top nonzero homology presheaf of a sheaf, is itself a sheaf (of abelian groups). This is because the sheaf condition is a limit and hence is preserved by the truncation $\tau_{\geq d}$. By part 1, it follows that the top nonzero homology presheaf of $\mathcal{A}^{sh}$ is a sheaf whose sheafification vanishes, hence is 0. Thus $\mathcal{A}^{sh} = 0$. □

Lemma 7. *If $\mathcal{A}$ is bounded above (there is a $d \in \mathbb{Z}$ with $H_n(\mathcal{A})) = 0$ for all $n > d$), then so is $\mathcal{A}^{sh}$, and the same d works.*

Proof. This is actually fairly subtle. To prove it we need to know a bit better how to actually produce the sheafification functor. It arises as a transfinite composition of so-called *Cech constructions $\mathcal{F} \mapsto \mathcal{F}^{\dagger}$*, where the sections over U

$$\mathcal{F}^{\dagger}(U) = \varinjlim_{\mathfrak{U}} \varprojlim_{V \in \mathfrak{U}} \mathcal{F}(V)$$

is the filtered colimit over all covering sieves of U of the limit over that sieve. This should be plausible if you've seen the analogous formula for sheafification with values in ordinary categories, where you just need to apply this Cech construction twice instead of transfinitely often. More generally, for a sheaf of n-truncated spaces you'd need to apply it $n+1$ times. Naively you might think that for a sheaf of arbitrary spaces you'd just need to iterated it countably often, but as usual in the theory of presentable ∞-categories, based on the small object argument, you a priori need a transfinite composition depending on the size of the sieves (though perhaps here $\aleph_1$ always suffices). Anyway proving this formula for sheafification in the ∞-world is one of the main technical results in Lurie's treatment of sheaves in "Higher topos theory", and let us just take it for granted.

Since arbitrary limits and filtered colimits preserve $D(\mathbb{Z})_{\leq d}$, we see that the lemma is a consequence of this description. □

Corollary 8. *Suppose $\mathcal{F} \to \mathcal{G}$ is a map of presheaves. If the induced map on sheafifications is an isomorphism, then the induced map on sheafified homology groups is an isomorphism. If both $\mathcal{F}$ and $\mathcal{G}$ are bounded above, the converse holds.*

Corollary 9. *A map $\mathcal{F} \to \mathcal{G}$ of bounded above presheaves identifies $\mathcal{G}$ with the sheafification of $\mathcal{F}$ if and only if it is an iso on sheafified homology groups and $\mathcal{G}$ is a sheaf.*

Now let's get to injective resolutions.

Proposition 10. *Suppose $\mathcal{I}_\bullet$ is a homologically bounded above chain complex of injective sheaves of abelian groups on X. Then the presheaf $|\mathcal{I}_\bullet| \in PSh(X, D(\mathbb{Z}))$ is a sheaf, where*

$$|\mathcal{I}_\bullet|(U) := |\mathcal{I}(U)_\bullet|.$$

Proof. We have $|\mathcal{I}_\bullet| \tilde{\to} \varprojlim |F^{\leq p}\mathcal{I}_\bullet|$, the inverse limit of the brutal truncations, as this is true sectionwise by checking on homology. Thus we can reduce to the case where $\mathcal{I}_\bullet$ only lives in finitely many degrees. But again by filtering using the brutal truncation and inductively using fiber-cofiber sequences we can reduce to where it's concentrated in a single degree, which may as well be degree 0. Thus it suffices to show that if $\mathcal{I}$ is an injective sheaf of abelian groups, then

$$\mathcal{I}(U) = \varprojlim_{V \in \mathfrak{U}^{op}} \mathcal{I}(V),$$

where $\mathfrak{U}$ is a sieve covering U and the limit is taken in $D(\mathbb{Z})$.

To prove this, consider the cofiber $C := cofib(\varinjlim_{V \in \mathfrak{U}} \mathbb{Z}[h_V] \to \mathbb{Z}[h_U])$ in the ∞-category of presheaves. This is the realization of a complex of presheaves $C_\bullet$ where each term is a direct sum of free presheaves $\mathbb{Z}[h_W]$ on open subsets W. (One can see this explicitly just by staring at C, but that's also not necessary: it's a general fact that any presheaf concentrated in non-negative degrees is realized by such a complex. This can be proved analogously to how we proved that every element of $D(\mathbb{Z})$ is the realization of a chain complex of free $\mathbb{Z}$-modules in the previous lecture, using that every presheaf of abelian groups admits a surjection from such a direct sum.)

Recall, by the correspondence between complexes and certain filtered objects, that this means that C has an increasing filtration where the n^{th} associated graded is concentrated in degree n and has n^{th} homology given by a direct sum of copies of $\mathbb{Z}[h_W]$'s. It follows that if $\mathcal{F} \in PSh(X; D(\mathbb{Z}))$, then $Map(C, \mathcal{F})$, or rather the mapping object in $D(\mathbb{Z})$, admits a dual filtration where the n^{th} associated graded is given by

$$Map(\Sigma^n \oplus_I \mathbb{Z}[h_{W_i}], \mathcal{F}) = \Sigma^{-n} \prod_{i \in I} \mathcal{F}(W_i).$$

If $\mathcal{F}$ actually lives in degere zero and hence corresponds to a presheaf of abelian gropus, then this is concentrated in degree $-n$, and it follows again by the correspondence between filtered objects and chain complexes that $Map(C, \mathcal{F})$ is realized by the cochain complex $Hom(C_\bullet, \mathcal{F})$. If furthermore $\mathcal{F}$ is a sheaf of abelian groups, this is the same as $Hom(C_\bullet^{sh}, \mathcal{F})$. Now, we need to show that if $\mathcal{F} = \mathcal{I}$ is an injective sheaf, then this cochain complex of abelian groups has vanishing cohomology. (For then its realization will too, hence $Map(C, \mathcal{I}) = 0$ as desired.) But because mapping out to an injective object is an exact functor, for that it suffices to see that the sheafified homology groups of $C_\bullet$ vanish. However $|C_\bullet| = C$ dies on sheafification by construction, hence has vanishing sheafified homology groups. □

Corollary 11. *Suppose that $\mathcal{F}_\bullet$ is an arbitrary bounded above complex of presheaves of abelian groups, and choose any map $\mathcal{F}_\bullet \to \mathcal{I}_\bullet$ to a bounded above complex of injective sheaves, such that the induced map on sheafified homology groups is an iso. Then the map in $PSh(X, D(\mathbb{Z}))$ given by*

$$|\mathcal{F}_\bullet| \to |\mathcal{I}_\bullet|$$

realizes the target as the sheafification of the source.

Proof. Since the two presheaves are bounded above and the map is an isomorphism on sheafified homology, the map is an iso on sheafification. But the target is a sheaf. □

In particular, given a sheaf of abelian groups $\mathcal{F}$, we have the formula

$$H^n(X; \mathcal{F}) = H_{-n}\Gamma(X; \mathcal{F}[0]^{sh}),$$

explaining what sheaf cohomology "is". The notation makes the right hand side look weirder than it really is. The point is this: if we have a presheaf of abelian groups, we can either sheafify it with values in Ab or with values in $D(\mathbb{Z})$. The latter contains much more information. The degree 0 homology will recover the sheafification with values in Ab, and the lower homology will give you the sheaf cohomology of that sheafification.

Here's an example of something which is straightforward to prove with this perspective, but quite tricky to prove with the standard tools such as we used in the first lecture; see Sella's rather recent paper "Comparison of sheaf cohomology and singular cohomology" which was the first to prove this result in this generality. (It was classically known with a paracompactness assumption. See also Dan Petersen's paper arXiv:2102.06927 for a very similar argument to the one we'll be giving here, but in model category language.)

Theorem 12. *Let X be a topological space which is locally contractible. Then for any abelian group A, we have $H^*_{sing}(X; A) \simeq H^*(X, A)$, i.e. singular cohomology agrees with sheaf cohomology.*

Proof. Pulling back from the point gives a comparison map of complexes of presheaves on X

$$A \to C^*_{sing}(-, A).$$

We need to see that the induced map

$$|A[0]| \to |C^*_{sing}(-, A)|$$

is the sheafification in $D(\mathbb{Z})$. We check that it's an iso on sheafification and that the target is a sheaf. For the first, as the complexes are bounded above we can check on sheafified homology groups. Then this is true by local contractibility; actually all that's needed is that each stalk of the singular cohomology graded presheaf is 0 in negative degrees and A in degree 0.

What remains is to see that the target is a sheaf. By the correspondence between chain complexes and filtered objects, and using that the terms in the singular chain complex are free modules, we find that $|C^*_{sing}(U, A)|$ is the internal hom from $|C^{sing}_*(U, \mathbb{Z})|$ to $|A[0]|$. Thus it suffices to show $|C^{sing}_*(-, \mathbb{Z})|$ is a sheaf with values in $D(\mathbb{Z})^{op}$. We use the criterion of Proposition 4, or rather its formal dual. As simplices are compact, mapping out from them commutes with filtered unions of open subsets, hence the singular chain complex commutes with such colimits, hence so does its realization. This shows the third condition. As for the second, it is exactly Mayer-Vietoris for singular homology, and the first is trivial. □

Remark 13. *Combining with the de Rham theorem, we get isomorphisms*

$$H^*_{dR}(M;\mathbb{R}) \simeq H^*(M;\mathbb{R}) \simeq H^*_{sing}(M;\mathbb{R})$$

between de Rham cohomology and singular cohomology on a manifold. It might look like this resulting composite isomorphism is very abstract. How do we know it is induced by the "usual" map gotten by integrating over simplices?

Actually, that's easy to check. First of all, to define the integration over simplices we can't use arbitrary continuous simplices as appear in the singular chain complex. The most convenient choice is smooth simplices, so let's instead consider the C^∞-analog of the singular complex, where the degree n term is the free abelian group on the set of smooth maps from Δ^n to U. One can similarly establish Mayer-Vietoris and copy the above proof to see that this C^∞-analog also computes sheaf cohomology with $\mathbb{R}$-coefficients.

On the other hand there is a natural comparison map of cochain complexes

$$i : C^*_{dR}(-;\mathbb{R}) \to C^*_{C^\infty-sing}(-;\mathbb{R})$$

given by integrating over simplices, which commutes with the differential by Stokes' theorem for simplices. Again pulling back from the point shows that the resolutions $\mathbb{R} \to C^*_{dR}(-;\mathbb{R})$ and $\mathbb{R} \to C^*_{C^\infty-sing}(-;\mathbb{R})$ are compatible with this comparison map. Using to $|-|$ pass to presheaves with values in $D(\mathbb{Z})$, the two resolution maps become isomorphisms on sheafification by the theorem; but as $|C^*_{dR}(-;\mathbb{R})|$ and $|C^*_{C^\infty-sing}(-;\mathbb{R})|$ are already sheaves it follows that $|i|$ is an isomorphism, and indeed is the unique isomorphism making the triangle commute (by the universal property of sheafification). Thus we learn that the abstract isomorphism referred to above is induced on cohomology by the concrete integration map i, and in particular i is a quasi-isomoprhism, which is perhaps the more classical statement of de Rham's theorem.

Note that for this argument it was important to work in $D(\mathbb{Z})$, not in the two other options of chain complexes (too rigid) or graded abelian groups (just remembering the cohomology — not enough structure). It's only in $D(\mathbb{Z})$ that everything is described by a simple universal property which lets you prove uniqueness claims.

In fact, the abstract perspective gives you even more, namely it can be use to produce *the integration map i (and prove Stokes' theorem). Indeed, because the abstractly-produced comparison isomorphism between $|C^*_{dR}|$ and $|C^*_{C^\infty-sing}|$ is functorial in smooth maps, we can also use it to define a relative version, say for a manifold with corners relative to its boundary, simply by taking the fiber of the restriction map. For an n-simplex there is a canonical fundamental class in relative homology in degree n, so this means if we have a class in degree n relative de Rham cohomology of an n-simplex we get to write down a real number by pairing with the fundamental class. But any degree n differential form defines a class in relative de Rham cohomology, because even on the level of chain complexes it maps to zero on the boundary simply because the boundary has dimension $< n$. This gives a perfectly acceptable definition of what it means to integrate an n-form over a smooth n-simplex in a smooth manifold M, and it makes Stokes' theorem, plus the de Rham theorem, more or less a tautology.*

Exercise 14. *A topological space X is said to have covering dimension $\leq d$ if every open cover of X has a refinement $\{U_i\}_{i\in I}$ such that $U_{i_1} \cap \ldots \cap U_{i_{n+1}} = \varnothing$ whenever $i_1, \ldots, i_{n+1}$ are distinct elements of I. You will prove that if a topological space has the property that every open subset has covering dimension $\leq d < \infty$, then if $\mathcal{F} \in Sh(X; D(\mathbb{Z}))$ is an* arbitrary *sheaf, not necessarily bounded above, it holds that $H_n(\mathcal{F})^{sh} = 0 \forall n \in \mathbb{Z}$ implies $\mathcal{F} = 0$. (In other words, every sheaf is* hypercomplete.*) Thus* *for such topological spaces, of which every manifold is an example, one can remove the word "bounded above" from every result in this lecture. It's enough to do just* two *of the three parts of this exercise, and you can use a part even if you didn't prove it.*

Let X be a topological space satisfying this dimension hypothesis: every open subset has dimension $\leq d$.

1. *Suppose P is a poset with Krull dimension $\leq d$, which means that for any chain*

$$p_0 \leq p_1 \leq \ldots \leq p_{d+1}$$

   we have $p_i = p_{i+1}$ for some i. Show that for any functor $F : P \to D(\mathbb{Z})_{\geq 0}$, we have $\varprojlim F \in D(\mathbb{Z})_{\geq -d}$.

2. *Using the formula for sheafification in terms of Cech constructions, deduce that if $\mathcal{F} \in PSh(X; D(\mathbb{Z})_{\geq 0})$, then $\mathcal{F}^{sh} \in PSh(X; D(\mathbb{Z})_{\geq -d})$.*

3. *Deduce that if $\mathcal{F} \in Sh(X; D(\mathbb{Z}))$, then $\mathcal{F} \tilde{\to} \varprojlim_n (\tau_{\leq n}\mathcal{F})^{sh}$. Use this to deduce the desired fact that if the sheafified homology groups of $\mathcal{F}$ vanish, then $\mathcal{F} = 0$. Conclude also from part 2 that for any sheaf of abelian groups on X, its cohomology vanishes in degrees $> d$.*

Exercise 15. *The point of this exercise is to show that the various derived functors one normally uses when discussing sheaves can also be seen from a "purely derived" point of view using ∞-categories, without any resolutions. Just like in the lecture where we saw that the derived functors of the global sections functor are just the plain old global sections if you work with derived sheaves. It's enough to just do* two *of the three parts of this exercise.*

First, recall that if $f : X \to Y$ is a map of topological spaces, then the pushforward functor $f_* : PSh(X) \to PSh(Y)$ is simply given by composition with $f^{-1} : Open(Y) \to Open(X)$. This works for any target ∞-category, and it preserves the sheaf condition and hence restricts to a functor $f_* : Sh(X) \to Sh(Y)$.

1. *If $\mathcal{F} \in \mathrm{Sh}(X; Ab)$, write $|\mathcal{F}| \in \mathrm{Sh}(X; D(\mathbb{Z}))$ for the sheafification of $\mathcal{F}$ viewed as a presheaf with values in $D(\mathbb{Z})$. For $i \geq 0$, show that $(H_{-i} f_* |\mathcal{F}|)^{sh}$ identifies with $R^i f_*(\mathcal{F}) \in \mathrm{Sh}(X; Ab)$, the i^{th} right derived functor of the pushforward functor on sheaves of abelian groups.*

2. *Recall the pullback functor on presheaves, defined as the left adjoint to pushforward, is a left Kan extension along f^{-1} and can be pointwise calculated as the (filtered) colimit*

$$(f^* \mathcal{G})(U) = \varinjlim_{V \in Open(Y), f^{-1}(V) \supset U} \mathcal{G}(V).$$

   This works for any target ∞-category with all colimits. If the target ∞-category is presentable, it follows that we get the pullback on sheaves by sheafifying the presheaf pullback.

   Now, for sheaves of abelian groups, pullback is exact so there's no need to derive it. This motivates the following.

   Show that if $\mathcal{G} \in Sh(Y, Ab)$, then $|f^* \mathcal{G}| = f^* |\mathcal{G}|$ where again we use $|\cdot|$ to mean pass to the presheaf with values in $D(\mathbb{Z})$ then sheafify, and we use f^* for the pullback functor both on sheaves with values in Ab and with values in $D(\mathbb{Z})$.

3. *Suppose X is a Hausdorff space, and define the functor $\Gamma_c(X;-) : Sh(X, Ab) \to Ab$ by setting*

$$\Gamma_c(X;\mathcal{F}) = \{s \in \mathcal{F}(X) | s\,|_{X \smallsetminus K} = 0 \textit{ for some compact } K \subset X\}.$$

This is called the group of sections with compact support. Prove that this functor is left exact and so right derived functors $R^i\Gamma_c(X;-)$ can be defined. On the other hand we can directly define the derived analog $\Gamma_c(X;-) : Sh(X, D(\mathbb{Z})) \to D(\mathbb{Z})$ by

$$\Gamma_c(X;\mathcal{F}) = \varinjlim_{K \subset X \textit{ compact}} Fib\left(\mathcal{F}(X) \to \mathcal{F}(X \smallsetminus K)\right).$$

Show that if $\mathcal{F} \in Sh(X, Ab)$, then for $i \geq 0$ we have

$$R^i\Gamma_c(X;\mathcal{F}) = H_{-i}\Gamma_c(X;|\mathcal{F}|).$$