The d-orbifold programme. Lecture 5 of 5: D-orbifold homology and cohomology, and virtual cycles

Dominic Joyce, Oxford University

May 2014

Work in progress, no papers yet.

However, you can find a previous version of this project using Kuranishi spaces instead of d-orbifolds at arXiv:0710.5634 (survey), and arXiv:0707.3572.

Warning: these papers should not be trusted, as the Kuranishi spaces material is dodgy, but they do show the basic idea.

Plan of talk:

1. Introduction

Several important areas of symplectic geometry involve 'counting' moduli spaces $\bar{\mathcal{M}}$ of J-holomorphic curves — Gromov–Witten invariants, Lagrangian Floer cohomology, Symplectic Field Theory, ... To do this, we first endow $\bar{\mathcal{M}}$ with an appropriate geometric structure (smooth manifold, Kuranishi space, polyfold, d-orbifold), and then associate a *virtual class* (or *virtual cycle*, or *virtual chain*) to $\bar{\mathcal{M}}$, in some (co)homology theory. That is, we need a bridge between moduli spaces and homological algebra.

If things go nicely (they may not), *geometric* relations between moduli spaces translate cleanly into *algebraic* relations between their virtual classes / virtual chains. The geometers can then go home for tea, and leave a specialist in homological algebra to define Lagrangian Floer cohomology, Fukaya categories, SFT, etc.

If things do not go nicely, it is probably the geometers' fault.

For example, moduli spaces $\bar{\boldsymbol{\mathcal{M}}}_k(\gamma)$ of J-holomorphic discs in homology class γ in a symplectic manifold with boundary in a Lagrangian L and k boundary marked points are Kuranishi spaces or d-orbifolds with corners satisfying the boundary equation

$$\partial\bar{\boldsymbol{\mathcal{M}}}_k(\gamma) \cong \textstyle\coprod_{i+j=k,\ \alpha+\beta=\gamma} \bar{\boldsymbol{\mathcal{M}}}_{i+1}(\alpha) \times_L \bar{\boldsymbol{\mathcal{M}}}_{j+1}(\beta). \qquad (1)$$

Ideally, one would like $\bar{\boldsymbol{\mathcal{M}}}_k(\gamma)$ to have a *virtual chain* $[\bar{\boldsymbol{\mathcal{M}}}_k(\gamma)]_{\rm virt}$ in the chains $C_*(L^k)$ of some homology theory of L^k, satisfying

$$\partial[\bar{\boldsymbol{\mathcal{M}}}_k(\gamma)]_{\rm virt} = \textstyle\sum_{i+j=k,\ \alpha+\beta=\gamma} [\bar{\boldsymbol{\mathcal{M}}}_{i+1}(\alpha)]_{\rm virt} \bullet_L [\bar{\mathcal{M}}_{j+1}(\beta)]_{\rm virt}, \quad (2)$$

where $\bullet_L$ is an associative, supercommutative intersection product defined *at the chain level*, not just on homology. (Compare with the wedge product of forms in de Rham cohomology.)

In [FOOO], this is done by perturbing the Kuranishi spaces using multisections to get a ($\mathbb{Q}$-weighted, non-Hausdorff) manifold, triangulating this by simplices to get a chain in singular homology. This process is acutely painful, because singular homology does not play at all nicely with Kuranishi spaces, and much of the algebraic complexity of [FOOO] is due to the problems this causes — especially, perturbing Kuranishi spaces to transverse.
Note too that singular chains $C_*^{\rm sing}(L)$ do not have a nice intersection product $\bullet_L$ defined at the chain level, so (2) does not make sense. Basically, singular homology is not a good choice.

As I understand it, the proposed virtual cycle construction for polyfolds (HWZ arXiv:0711.0781) is very similar to [FOOO] (curiously, this is not mentioned), except that after perturbing by multisections, they integrate differential forms over the ($\mathbb{Q}$-weighted, non-Hausdorff) manifold. So $[\bar{\mathcal{M}}_k(\gamma)]_{\rm virt}\in C^\infty(\Lambda^\bullet L^k)^*$ is a *current* on L^k. Currents also have no natural chain product $\bullet_L$.

Introduction
D-manifold bordism
D-orbifold homology
D-orbifold cohomology

D-orbifold homology and cohomology

I propose to define new (co)homology theories $dH_*(Y,R)$, $dH^*(Y,R)$ of a manifold or orbifold Y, with coefficients in a $\mathbb{Q}$-algebra R, isomorphic to ordinary homology and compactly-supported cohomology, but in which the (co)chains are 1-morphisms $\mathbf{f}:\mathbf{X}\to Y$ for $\mathbf{X}$ a compact, oriented d-orbifold with g-corners (or stratified d-orbifold), plus extra 'gauge-fixing data' $\mathbf{G}$.
Forming virtual classes for moduli spaces in d-orbifold (co)homology is almost trivial, there is no need to perturb, one just chooses gauge-fixing data $\mathbf{G}$, which is always possible.
There is an associative, supercommutative intersection product $\bullet_L$ defined on chains $dC_*(L)$. For the moduli spaces $\bar{\mathcal{M}}_k(\gamma)$, in a single induction one can choose virtual chains $[\bar{\mathcal{M}}_k(\gamma)]_{\rm virt}\in dC_*(L^k)$ satisfying (2). Because of this, the homological algebra in [FOOO] can be drastically simplified if we use d-orbifold homology instead of singular homology.

2. Ordinary bordism and d-manifold bordism

As a warm-up, we discuss bordism and d-manifold bordism.

Let Y be a manifold. Define the *bordism group* $B_k(Y)$ to have elements $\sim$-equivalence classes $[X,f]$ of pairs (X,f), where X is a compact oriented k-manifold without boundary and $f:X\to Y$ is smooth, and $(X,f)\sim(X',f')$ if there exists a $(k+1)$-manifold with boundary W and a smooth map $e:W\to Y$ with $\partial W\cong X\amalg -X'$ and $e|_{\partial W}\cong f\amalg f'$. It is an abelian group, with addition $[X,f]+[X',f']=[X\amalg X,f\amalg f']$.
If Y is oriented of dimension n, there is a supercommutative, associative intersection product $\bullet:B_k(Y)\times B_l(Y)\to B_{k+l-n}(Y)$ given by $[X,f]\bullet[X',f']=[X\times_{f,Y,f'}X',\pi_Y]$, choosing X,f, X',f' in their bordism classes with $f:X\to Y$, $f':X'\to Y$ transverse.
There is a natural morphism $\Pi_{\rm bo}^{\rm hom}:B_k(Y)\to H_k(Y,\mathbb{Z})$ given by $\Pi_{\rm bo}^{\rm hom}:[X,f]\mapsto f_*([X])$, for $[X]\in H_k(X,\mathbb{Z})$ the fundamental class.

Similarly, define the *derived bordism group* $dB_k(Y)$ to have elements $\approx$-equivalence classes $[\mathbf{X},\mathbf{f}]$ of pairs (X,f), where $\mathbf{X}$ is a compact oriented d-manifold with $\mathrm{vdim}\,\mathbf{X}=k$ and $\mathbf{f}:\mathbf{X}\to Y$ is a 1-morphism in **dMan**, and $(\mathbf{X},\mathbf{f})\approx(\mathbf{X}',\mathbf{f}')$ if there exists a d-manifold with boundary $\mathbf{W}$ with $\mathrm{vdim}\,\mathbf{W}=k+1$ and a 1-morphism $\mathbf{e}:\mathbf{W}\to\mathbf{Y}$ in $\mathbf{dMan^b}$ with $\partial\mathbf{W}\simeq\mathbf{X}\amalg-\mathbf{X}'$ and $\mathbf{e}|_{\partial\mathbf{W}}\cong\mathbf{f}\amalg\mathbf{f}'$. It is an abelian group, with $[\mathbf{X},\mathbf{f}]+[\mathbf{X}',\mathbf{f}']=[\mathbf{X}\amalg\mathbf{X},\mathbf{f}\amalg\mathbf{f}']$.
If Y is oriented of dimension n, there is a supercommutative, associative intersection product $\bullet:dB_k(Y)\times dB_l(Y)\to dB_{k+l-n}(Y)$ given by $[\mathbf{X},\mathbf{f}]\bullet[\mathbf{X}',\mathbf{f}']=[\mathbf{X}\times_{\mathbf{f},Y,\mathbf{f}'}\mathbf{X}',\boldsymbol{\pi}_Y]$, with no transversality condition on $\mathbf{X},\mathbf{f}$, $\mathbf{X}',\mathbf{f}'$.

There is a natural morphism $\Pi_{\rm bo}^{\rm dbo} : B_k(Y) \to dB_k(Y)$ mapping $[X, f] \mapsto [X, f]$.

Theorem

$\Pi_{\rm bo}^{\rm dbo} : B_k(Y) \to dB_k(Y)$ is an isomorphism for all k, with $dB_k(Y) = 0$ for $k < 0$.

This holds as every d-manifold can be perturbed to a manifold.

Composing $(\Pi_{\rm bo}^{\rm dbo})^{-1}$ with $\Pi_{\rm bo}^{\rm hom} : B_k(Y) \to H_k(Y, \mathbb{Z})$ gives a morphism $\Pi_{\rm dbo}^{\rm hom} : dB_k(Y) \to H_k(Y, \mathbb{Z})$. We can interpret this as a *virtual class map* for compact, oriented d-manifolds. In particular, this is an easy proof that *the geometric structure on d-manifolds is strong enough to define virtual classes.*

Virtual classes (in homology over $\mathbb{Q}$) also exist for compact oriented d-orbifolds, though the proof is harder.

3. D-orbifold homology

3.1. The basic idea

Let Y be a manifold or orbifold, and R a $\mathbb{Q}$-algebra. We define a complex of R-modules $\bigl(dC_*(Y, R), \partial\bigr)$, whose homology groups $dH_*(Y, R)$ are the *d-orbifold homology* of Y.

Similarly to the definition of d-manifold bordism, chains in $dC_k(Y, R)$ for $k \in \mathbb{Z}$ are R-linear combinations of equivalence classes $[\mathbf{X}, \mathbf{f}, \mathbf{G}]$ with relations, where $\mathbf{X}$ is a compact, oriented d-orbifold with g-corners (or stratified d-orbifold) with dimension k, $\mathbf{f} : \mathbf{X} \to Y$ is a 1-morphism in $\mathbf{dOrb^{gc}}$, and $\mathbf{G}$ is some extra 'gauge-fixing data' associated to $\mathbf{X}$, with many possible choices.

I won't give all the relations on the $[\mathbf{X}, \mathbf{f}, \mathbf{G}]$. Two examples are:

$$[\mathbf{X}_1 \amalg \mathbf{X}_2, \mathbf{f}_1 \amalg \mathbf{f}_2, \mathbf{G}_1 \amalg \mathbf{G}_2] = [\mathbf{X}_1, \mathbf{f}_1, \mathbf{G}_1] + [\mathbf{X}_2, \mathbf{f}_2, \mathbf{G}_2], \tag{3}$$

$$[-\mathbf{X}, \mathbf{f}, \mathbf{G}] = -[\mathbf{X}, \mathbf{f}, \mathbf{G}], \tag{4}$$

where $-\mathbf{X}$ is $\mathbf{X}$ with the opposite orientation.

Gauge-fixing data – first properties

Here 'gauge-fixing data' is the key to the whole story. I won't tell you what it is — in fact, this doesn't really matter — but I will give you lists of axioms we need it to satisfy. Here are the first:

(i) For any d-orbifold with g-corners $\mathbf{X}$ (not necessarily compact or oriented) we have a nonempty set $\mathrm{Gauge}(\mathbf{X})$ of choices of 'gauge-fixing data' $\mathbf{G}$ for $\mathbf{X}$.

(ii) If $\mathbf{f}:\mathbf{X}\rightarrow\mathbf{Y}$ is étale (e.g. an equivalence, or inclusion of an open) we have a *pullback map* $\mathbf{f}^*:\mathrm{Gauge}(\mathbf{Y})\rightarrow\mathrm{Gauge}(\mathbf{X})$. If $\mathbf{f},\mathbf{g}$ are 2-isomorphic then $\mathbf{f}^*=\mathbf{g}^*$. Pullbacks are functorial.

(iii) There is a *boundary map* $|_{\partial\mathbf{X}}:\mathrm{Gauge}(\mathbf{X})\rightarrow\mathrm{Gauge}(\partial\mathbf{X})$. We regard it as a pullback along the (non-étale) $\mathbf{i_X}:\partial\mathbf{X}\rightarrow\mathbf{X}$. The maps $|_{\partial\mathbf{X}}$ combine functorially with other pullbacks.

(iv) (Desirable.) The map $\mathbf{U}\rightarrow\mathrm{Gauge}(\mathbf{U})$ for $\mathbf{U}\subseteq\mathbf{X}$ open is a *sheaf* on $\mathbf{X}$, using pullbacks by inclusions $\mathbf{U}\hookrightarrow\mathbf{V}\subseteq\mathbf{X}$.

Boundary operators

Note that d-orbifolds $\mathbf{X}$ can have virtual dimension $\mathrm{vdim}\,\mathbf{X}<0$, so $dC_k(Y,R)\neq 0$ for all $k<0$, although $dH_k(Y,R)=0$ for $k<0$. The boundary operator $\partial:dC_k(Y,R)\rightarrow dC_{k-1}(Y,R)$ maps

$$\partial:\bigl[\mathbf{X},\mathbf{f},\mathbf{G}\bigr]\longmapsto\bigl[\partial\mathbf{X},\mathbf{f}\circ\mathbf{i_X},\mathbf{G}|_{\partial\mathbf{X}}\bigr].$$

We have a natural 1-morphism $\mathbf{i_X}:\partial\mathbf{X}\rightarrow\mathbf{X}$ and an equivalence $\partial^2\mathbf{X}\simeq\partial\mathbf{X}_{\mathbf{i_X},\mathbf{X},\mathbf{i_X}}\partial\mathbf{X}$. Thus there is an orientation-reversing involution $\boldsymbol{\sigma}:\partial^2\mathbf{X}\rightarrow\partial^2\mathbf{X}$ swapping the two factors of $\partial\mathbf{X}$. This satisfies $\mathbf{i_X}\circ\mathbf{i}_{\partial\mathbf{X}}\circ\boldsymbol{\sigma}\cong\mathbf{i_X}\circ\mathbf{i}_{\partial\mathbf{X}}$. Hence $\mathbf{G}|_{\partial^2\mathbf{X}}$ is $\boldsymbol{\sigma}$-invariant. Using this and (4) we show that $\partial^2=0$, so $dH_*(Y,R)$ is well-defined.

Here is a property of gauge-fixing data with prescribed boundary values. 'Only if' is necessary by (i)–(iii) as above.

(v) Suppose $\mathbf{H}\in\mathrm{Gauge}(\partial\mathbf{X})$. Then there exists $\mathbf{G}\in\mathrm{Gauge}(\mathbf{X})$ with $\mathbf{G}|_{\partial\mathbf{X}}=\mathbf{H}$ if and only if $\boldsymbol{\sigma}^*(\mathbf{H}|_{\partial^2\mathbf{X}})=\mathbf{H}|_{\partial^2\mathbf{X}}$.

Pushforwards, relation to singular homology

Suppose $g : Y \to Z$ is a smooth map of manifolds or orbifolds. Define an R-linear *pushforward* $g_* : dC_k(Y,R) \to dC_k(Z,R)$ by $g_* : [\mathbf{X}, \mathbf{f}, \mathbf{G}] \mapsto [\mathbf{X}, g \circ \mathbf{f}, \mathbf{G}]$. Then $g_* \circ \partial = \partial \circ g_*$, so this induces $g_* : dH_k(Y,R) \to dH_k(Z,R)$. Pushforwards are functorial.

Singular homology $H_*^{\rm sing}(Y,R)$ may be defined using $(C_*^{\rm sing}(Y,R), \partial)$, where $C_k^{\rm sing}(Y,R)$ is spanned by *smooth* maps $f : \Delta_k \to Y$, for Δ_k the standard k-simplex, thought of as a manifold with corners.

We define an R-linear map $F_{\rm sing}^{\rm dH} : C_k^{\rm sing}(Y,R) \to dC_k(Y,R)$ by

$$F_{\rm sing}^{\rm dH} : f \longmapsto \left[\Delta_k, f, \mathbf{G}_{\Delta_k}\right],$$

with $\mathbf{G}_{\Delta_k}$ some standard choice of gauge-fixing data for Δ_k. Then $F_{\rm sing}^{\rm dH} \circ \partial = \partial \circ F_{\rm sing}^{\rm dH}$, so that $F_{\rm sing}^{\rm dH}$ induces morphisms $F_{\rm sing}^{\rm dH} : H_k^{\rm sing}(Y,R) \to dH_k(Y,R)$.

3.2. The main result (hopefully)

The main result of the theory (I haven't proved it yet!) will be:

"Theorem"

$F_{\rm sing}^{\rm dH} : H_k^{\rm sing}(Y,R) \to dH_k(Y,R)$ *is an isomorphism for all* $k \in \mathbb{Z}$.

Some remarks:

- Whether or not the "Theorem" is actually true depends on the choice of relations on the $[\mathbf{X}, \mathbf{f}, \mathbf{G}]$ in the definition of $dC_k(Y)$, and on the definition (or at least the properties) of gauge-fixing data, neither of which I have explained. The aim is to design them to make the "Theorem" hold.
- Forming virtual classes/virtual chains is easy. Suppose $\bar{\mathcal{M}}$ is a moduli d-orbifold, with evaluation map $\mathbf{ev} : \bar{\mathcal{M}} \to Y$. Choose gauge-fixing data $\mathbf{G}$ for $\bar{\mathcal{M}}$, which is possible by (i). Then $[\bar{\mathcal{M}}, \mathbf{ev}, \mathbf{G}] \in dC_k(Y)$ is a virtual chain for $\bar{\mathcal{M}}$. If $\partial\bar{\mathcal{M}} = \emptyset$ then $[[\bar{\mathcal{M}}, \mathbf{ev}, \mathbf{G}]] \in dH_k(Y)$ is a virtual class for $\bar{\mathcal{M}}$.

- Obviously, d-orbifold homology is not a new invariant, it's just ordinary homology. The point is that it has special properties which make it more convenient than competing homology theories (e.g. singular homology) for some tasks:

(a) D-orbifold homology is very well adapted for forming *virtual cycles* and *virtual chains* in moduli problems. It is particularly powerful for moduli spaces 'with corners', as in Lagrangian Floer homology and Symplectic Field Theory.
Using d-orbifold homology instead of singular homology could simplify [FOOO] considerably.

(b) Issues to do with *transversality* — for instance, defining intersection products on (not necessarily transverse) chains — often disappear in d-orbifold homology.

- To prove the "Theorem", the main issue is to show that any cycle $\sum_{i=1}^{m} c_i[\mathbf{X}_i, \mathbf{f}_i, \mathbf{G}_i]$ in $dC_k(Y, R)$ is homologous to a singular cycle $\sum_{j=1}^{n} d_j[\Delta_k, g_j, \mathbf{G}_{\Delta_k}]$.
The proof of this involves using homologies and relations to first cut the $\mathbf{X}_i$ into small pieces which are global quotients $[\mathbf{X}'/G]$ for d-manifolds $\mathbf{X}'$, replacing $[\mathbf{X}'/G]$ by $\frac{1}{|G|}\mathbf{X}'$ to get a cycle $\sum_{i=1}^{m'} c_i'[\mathbf{X}_i', \mathbf{f}_i', \mathbf{G}_i']$ involving d-manifolds with g-corners $\mathbf{X}_i'$ rather than d-orbifolds, perturbing the $\mathbf{X}_i'$ to manifolds with g-corners, and triangulating these by simplices.
- Essentially, this uses all the messy bits of [FOOO] I was complaining about: perturbation by multisections, triangulation by simplices, arbitrary choices with boundary compatibilities. We do have to solve all the same issues. But now, these messy bits are repackaged in the proof of the "Theorem". I only have to do them once, and you won't have to read it; you don't meet them every time you use theory.

3.3. Why we need gauge-fixing data: an example

D-manifold bordism $dB_k(Y)$ in §2 worked fine without gauge-fixing data. So what would go wrong if we did d-orbifold homology without gauge-fixing data, or equivalently, set $\mathrm{Gauge}(\mathbf{X}) = \{\mathrm{pt}\}$ for all $\mathbf{X}$, so there is only ever one choice for $\mathbf{G}$?

It turns out that without gauge-fixing data, we would have $dH_k(Y, R) = 0$ for all k, Y. We explain this by an example; finding out what goes wrong shows more properties we need of gauge-fixing data.

Take $Y = \mathrm{pt}$, the point. Then $dH_*(\mathrm{pt}, R)$ is a ring, with identity $[[\mathrm{pt}, \mathrm{id}]] \in dH_0(Y, R)$. We will show that $[[\mathrm{pt}, \mathrm{id}]] = 0$, so $dH_*(\mathrm{pt}, R) = 0$ as the identity is zero.

For any orbifold Z, by writing $Z \cong Z \times \mathrm{pt}$, we can make $dH_*(Z, R)$ into a module over $dH_*(\mathrm{pt}, R) = 0$. Hence $dH_*(Z, R) = 0$.

For $k \in \mathbb{Z}$, define $\mathbf{X}_k$ to be the 'standard model' d-manifold of virtual dimension 0, with Kuranishi neighbourhood $(\mathbb{CP}^1, \mathcal{O}(k), 0)$. Then $[\mathbf{X}_k, \pi] \in dC_0(\mathrm{pt}, R)$, supposing we are defining d-orbifold homology without gauge-fixing data, and $\partial[\mathbf{X}_k, \pi] = 0$ as $\partial\mathbf{X}_k = 0$, so $[[\mathbf{X}_k, \pi]] \in dH_0(\mathrm{pt}, R)$.

By perturbing the section $s = 0$ to a generic section $\tilde{s}$, which has k transverse zeros counted with signs, we can show that

$$[[\mathbf{X}_k, \pi]] = k[[\mathrm{pt}, \mathrm{id}]]. \tag{5}$$

Divide $\mathbb{CP}^1$ into two hemispheres D^+, D^- with common boundary $\mathcal{S}^1$. Let $\mathbf{X}_k^\pm$ be the corresponding d-manifolds with boundary, so that $\mathbf{X}_k = \mathbf{X}_k^+ \cup_{\mathcal{S}^1} \mathbf{X}_k^-$. Then we can show that

$$[[\mathbf{X}_k, \pi]] = \big[[\mathbf{X}_k^+, \pi] + [\mathbf{X}_k^-, \pi]\big]. \tag{6}$$

However, $\mathbf{X}_k^\pm$ are independent of $k \in \mathbb{Z}$ up to equivalence, as $\mathcal{O}(k)|_{D^\pm}$ is trivial. So (5)–(6) show that $[[\mathrm{pt}, \mathrm{id}]] = 0$, as we want.

What went wrong in this example?

To prove that $[\mathbf{X}_k, \boldsymbol{\pi}] \sim [\mathbf{X}_l, \boldsymbol{\pi}]$ for $k \neq l \in \mathbb{Z}$ we cut $\mathbf{X}_k$ into two pieces $\mathbf{X}_k^{\pm}$, whose boundary $\mathbf{Y} = \partial \mathbf{X}_k^{\pm}$ is a d-manifold of virtual dimension -1, a circle $\mathcal{S}^1$ with obstruction bundle $\mathbb{R}^2 \times \mathcal{S}^1 \to \mathcal{S}^1$.

To make $\mathbf{X}_l$ we glue $\mathbf{X}_k^{+}, \mathbf{X}_k^{-}$ back together along their boundaries in a different way, applying an automorphism of $\mathbf{Y} = \partial \mathbf{X}_k^{\pm}$ which rotates the obstruction bundle $l - k$ times as one goes round $\mathcal{S}^1$.

Thus, this example involves an automorphism of $\mathbf{Y}$, of infinite order, which acts nontrivially upon the obstructions of $\mathbf{Y}$.

This suggests that problems arise when chains $[\mathbf{X}, \boldsymbol{f}, \mathbf{G}]$ have *infinite automorphism groups* $\mathrm{Aut}(\mathbf{X}, \boldsymbol{f}, \mathbf{G})$.

More properties of gauge-fixing data

One property of gauge-fixing data we could impose is:

(vi) (Desirable.) Suppose $\mathbf{X}$ is a compact d-orbifold with g-corners, and $\mathbf{G} \in \mathrm{Gauge}(\mathbf{X})$. Then $\mathrm{Aut}(\mathbf{X}, \mathbf{G})$ is a finite group.

Actually, (vi) contradicts the sheaf property (iv) in examples where $\mathbf{X}$ has infinitely many connected components. But (vi) is stronger than we need; probably 'every finitely generated subgroup of $\mathrm{Aut}(\mathbf{X}, \mathbf{G})$ is finite' is enough, and there are weaker versions of (vi) consistent with (iv), '$\mathrm{Aut}(\mathbf{X}, \mathbf{G})$ is locally finite'.

In the proof of the "Theorem", this arises because sometimes we need to average over an automorphism group $\Gamma \subseteq \mathrm{Aut}(\mathbf{X}, \boldsymbol{f}, \mathbf{G})$ of a chain $[\mathbf{X}, \boldsymbol{f}, \mathbf{G}]$ – for instance, if we choose a non Γ-invariant deformation $\tilde{\mathbf{X}}$ of $\mathbf{X}$, and average over images of $\tilde{\mathbf{X}}$ under Γ. If Γ is not finite, this does not make sense. This is also one reason why we take the coefficient ring R to be a $\mathbb{Q}$-algebra, so that $\frac{1}{|\Gamma|} \in R$.

4. D-orbifold cohomology
4.1. The basic idea

I will also define a cochain complex $\big(dC^*(Y,R),\mathrm{d}\big)$ with cohomology $dH^*(Y,R)$, called *d-orbifold cohomology*.
The basic idea is to use Poincaré duality:

- If Y is a compact oriented n-manifold, then $H^k(Y,R)\cong H_{n-k}(Y,R)$. So roughly we take $dC^k(Y,R)=dC_{n-k}(Y,R)$. This is still true for orbifolds Y, as R is a $\mathbb{Q}$-algebra.
- If Y is noncompact, then $H^k_{\rm cs}(Y,R)\cong H_{n-k}(Y,R)$, where $H^*_{\rm cs}(Y,R)$ is compactly-supported cohomology. So in the analogue of the "Theorem", we want $dH^k(Y,R)\cong H^k_{\rm cs}(Y,R)$.
- If Y is not oriented, then instead of taking $\mathbf{X}$ oriented in cochains $[\mathbf{X},\mathbf{f},\mathbf{G}]$, we take $\mathbf{f}:\mathbf{X}\to Y$ *cooriented*. That is, the line bundle $\Lambda^{\rm top}\mathbb{L}_{\mathbf{X}}\otimes\mathbf{f}^*(\Lambda^{\rm top}TY)$ on $\mathbf{X}$ is oriented.

So, in a similar way to d-manifold homology, let Y be a manifold or orbifold, and R a $\mathbb{Q}$-algebra. We define a complex of R-modules $\big(dC^*(Y,R),\partial\big)$, whose homology groups $dH^*(Y,R)$ are the *d-orbifold cohomology* of Y.
Cochains in $dC_k(Y,R)$ for $k\in\mathbb{Z}$ are R-linear combinations of equivalence classes $[\mathbf{X},\mathbf{f},\mathbf{C}]$ with relations, where $\mathbf{X}$ is a compact d-orbifold with g-corners (or stratified d-orbifold) with virtual dimension $\dim Y-k$, $\mathbf{f}:\mathbf{X}\to Y$ is a cooriented 1-morphism in $\mathbf{dOrb^{gc}}$, and $\mathbf{G}$ is some extra 'co-gauge-fixing data' associated to $\mathbf{f}:\mathbf{X}\to Y$, with many possible choices, which we explain shortly.

As a general principle, in chains $[\mathbf{X},\mathbf{f},\mathbf{G}]$ in d-orbifold homology we work with absolute data on $\mathbf{X}$, but in cochains $[\mathbf{X},\mathbf{f},\mathbf{C}]$ in d-orbifold cohomology we work with relative data for $\mathbf{f}:\mathbf{X}\to Y$.

4.2. Co-gauge-fixing data

We can't (or shouldn't) define d-orbifold cohomology using gauge-fixing data, as it won't have the right functoriality: gauge-fixing data on $\mathbf{X}$ is good for defining pushforwards, but (compactly-supported) cohomology should have (proper) pullbacks.

So instead, cochains $dC^k(Y)$ in d-orbifold cohomology will be R-linear combinations of equivalence classes $[\mathbf{X}, \mathbf{f}, \mathbf{C}]$, where $\mathbf{C}$ is *co-gauge-fixing data* associated to the 1-morphism $\mathbf{f} : \mathbf{X} \to Y$, rather than just to the d-orbifold with corners $\mathbf{X}$.

In fact the target Y doesn't need to be a manifold or orbifold: we can associate co-gauge-fixing data to any *w-submersion* $\mathbf{f} : \mathbf{X} \to \mathbf{Y}$ in $\mathbf{dOrb^{gc}}$, where w-submersions induce injective maps $H^{-1}(\mathrm{d}f|_x) : O_y\mathbf{Y} \to O_x\mathbf{X}$ on obstructions. If Y is an orbifold, any $\mathbf{f} : \mathbf{X} \to \mathbf{Y}$ is a w-submersion, as the obstructions $O_y Y$ are zero.

Co-gauge-fixing data – first properties

(i) Any w-submersion $\mathbf{f} : \mathbf{X} \to \mathbf{Y}$ in $\mathbf{dOrb^{gc}}$ has a nonempty set $\mathrm{CoGauge}(\mathbf{f} : \mathbf{X} \to \mathbf{Y})$ of *co-gauge-fixing data* $\mathbf{C}$, which depends on $\mathbf{f}$ only up to 2-isomorphism.

(ii) If $\mathbf{f} : \mathbf{X} \to \mathbf{Y}$ is étale, we are given *identity co-gauge-fixing data* $\mathbf{1_f} \in \mathrm{CoGauge}(\mathbf{f} : \mathbf{X} \to \mathbf{Y})$.

(iii) If $\mathbf{f} : \mathbf{X} \to \mathbf{Y}$, $\mathbf{g} : \mathbf{Y} \to \mathbf{Z}$ are w-submersions, we have an associative *composition* $\circ : \mathrm{CoGauge}(\mathbf{g} : \mathbf{Y} \to \mathbf{Z}) \times \mathrm{CoGauge}(\mathbf{f} : \mathbf{X} \to \mathbf{Y}) \to \mathrm{CoGauge}(\mathbf{g} \circ \mathbf{f} : \mathbf{X} \to \mathbf{Z})$.

(iv) (Desirable.) If $\mathbf{f} : \mathbf{X} \to \mathbf{Y}$ is a w-submersion, then $\mathbf{U} \mapsto \mathrm{CoGauge}(\mathbf{f}|_{\mathbf{U}} : \mathbf{U} \to \mathbf{Y})$ for open $\mathbf{U} \subseteq \mathbf{X}$ is a sheaf, where if $\mathbf{U} \subseteq \mathbf{V} \subseteq \mathbf{X}$ are open then the restriction map $\mathrm{CoGauge}(\mathbf{f}|_{\mathbf{V}} : \mathbf{V} \to \mathbf{Y}) \to \mathrm{CoGauge}(\mathbf{f}|_{\mathbf{U}} : \mathbf{U} \to \mathbf{Y})$ is $\mathbf{C} \mapsto \mathbf{C} \circ \mathbf{1}_{\mathbf{U} \hookrightarrow \mathbf{V}}$.

Co-gauge-fixing data – first properties

(v) Suppose we are given a 2-Cartesian diagram in $\mathbf{dOrb^{gc}}$

$$\begin{array}{ccc} \mathbf{W} & \xrightarrow{\quad \mathbf{f} \quad} & \mathbf{Y} \\ \big\downarrow{\mathbf{e}} & \eta \Uparrow \quad g & \mathbf{h}\big\downarrow \\ \mathbf{X} & \xrightarrow{\quad\quad\quad} & \mathbf{Z}, \end{array}$$

with $\mathbf{g}$ (and hence $\mathbf{f}$) a w-submersion. Then there is a *pullback map* $\mathbf{e}^* : \mathrm{CoGauge}(\mathbf{g} : \mathbf{X} \to \mathbf{Z}) \to \mathrm{CoGauge}(\mathbf{f} : \mathbf{W} \to \mathbf{Y})$. These are functorial, together with composition, under combining 2-Cartesian squares into larger rectangles.

Let $g : Y \to Z$ be a smooth map of manifolds or orbifolds. Define an R-linear *pullback map* $g^* : dC^k(Z, R) \to dC^k(Y, R)$ by $g^* : [\mathbf{X}, \mathbf{f}, \mathbf{C}] \mapsto [\mathbf{X} \times_{\mathbf{f},Z,g} Y, \pi_Y, \pi^*_{\mathbf{X}}(\mathbf{C})]$, where $\pi^*_{\mathbf{X}}(\mathbf{C})$ is as in (v). Then $g^* \circ \mathrm{d} = \mathrm{d} \circ g^*$, so this induces $g^* : dH^k(Z, R) \to dH^k(Y, R)$ in the usual way.

4.4. Relation of gauge-fixing and co-gauge-fixing data

We can identify gauge-fixing data for $\mathbf{X}$ with co-gauge-fixing data for $\boldsymbol{\pi} : \mathbf{X} \to \mathrm{pt}$. In fact, both d-orbifold homology and cohomology are part of a larger 'bivariant theory' $dBH_*(g : Y \to Z, R)$ of smooth maps $g : Y \to Z$, mixing homology and cohomology, with chains $[\mathbf{X}, \mathbf{f}, \mathbf{C}]$ for $\mathbf{f} : \mathbf{X} \to Y$ and $\mathbf{C}$ co-gauge-fixing data for $g \circ \mathbf{f} : \mathbf{X} \to Z$, with $dH_k(Y, R) \cong dBH_k(\pi : Y \to \mathrm{pt}, R)$ and $dH^k(Y, R) \cong dBH_{-k}(\mathrm{id}_Y : Y \to Y, R)$.

This may have symplectic applications: for moduli spaces of J-holomorphic discs in Lagrangian Floer theory, it may be more functorial to associate 'incoming' boundary marked points to the homology of L, and 'outgoing' ones to the cohomology of L, mixing homology and cohomology.

4.5. Isomorphism with compactly-supported cohomology

Let Y be oriented of dimension n. Choose co-gauge-fixing data $\mathbf{G}_\pi$ for $\pi : Y \to *$. Define $\Pi_{\rm dcoh}^{\rm dhom} : dC^k(Y,R) \to dC_{n-k}(Y,R)$ by $\Pi_{\rm dcoh}^{\rm dhom} : [\mathbf{X},\mathbf{f},\mathbf{C}] \mapsto [\mathbf{X},\mathbf{f},\mathbf{G}_\pi \circ \mathbf{C}]$. This induces $\Pi_{\rm dcoh}^{\rm dhom} : dH^k(Y,R) \to dH_{n-k}(Y,R)$. I would like to prove:

"Theorem"

These maps $\Pi_{\rm dcoh}^{\rm dhom} : dH^k(Y,R) \to dH_{n-k}(Y,R)$ are isomorphisms, and independent of the choice of $\mathbf{G}_\pi$.

If so, the composition

$$H^k_{\rm cs}(Y,R) \xrightarrow[\text{duality}]{\text{Poincaré}} H^{\rm sing}_{n-k}(Y,R) \xrightarrow{F_{\rm sing}^{\rm dH}} dH_{n-k}(Y,R) \xrightarrow{(\Pi_{\rm dcoh}^{\rm dhom})^{-1}} dH^k(Y,R)$$

gives isomorphisms $H^*_{\rm cs}(Y,R) \cong dH^*(Y,R)$, as we want.