

HOMOLOGICAL STABILITY FOR MODULI SPACES OF HIGH DIMENSIONAL MANIFOLDS. II

SØREN GALATIUS AND OSCAR RANDAL-WILLIAMS

ABSTRACT. We prove a homological stability theorem for moduli spaces of manifolds of dimension $2n$, for attaching handles of index at least n , after these manifolds have been stabilised by countably many copies of $S^n \times S^n$.

Combined with previous work of the authors, we obtain an analogue of the Madsen–Weiss theorem for any simply-connected manifold of dimension $2n \geq 6$.

1. INTRODUCTION AND STATEMENT OF RESULTS

Let W be a smooth compact connected manifold of dimension $2n \geq 2$ and let $\text{Diff}_\partial(W)$ denote the topological group of diffeomorphisms of W relative to its boundary. This paper concludes our study ([GRW14b, GRW14a]) of the cohomology of the classifying space $B\text{Diff}_\partial(W \sharp g(S^n \times S^n))$ as well as some variants, resulting in a formula for its cohomology in the limit $g \rightarrow \infty$ in terms of a purely homotopy theoretic construction.

The main result is a new *stable homological stability* theorem, about the dependency of (a variant of) $B\text{Diff}_\partial(W \sharp g(S^n \times S^n))$ on W in the limit $g \rightarrow \infty$. Roughly speaking, it asserts that this limit depends only on the n -skeleton of W together with the restriction of the tangent bundle, up to tangential homotopy equivalence.

When combined with our earlier work on the subject, we deduce a homotopy theoretic formula for the cohomology of $B\text{Diff}_\partial(W \sharp g(S^n \times S^n))$ for finite g in a range of degrees increasing with g . The formula applies when $2n \geq 6$ and W is simply-connected. It is well-suited for calculations and has applications to positive scalar curvature ([BERW14]), topological Pontryagin classes ([Wei15]), and homotopy equivalences of manifolds ([BM14]).

1.1. Statement of the main theorem. In order to state our homological stability results more precisely, we must first introduce a variant of the space $B\text{Diff}_\partial(W)$ which incorporates certain tangent bundle information. As in our earlier paper [GRW14b] the clearest statements can be made by not fixing a manifold W , but rather considering all manifolds with a specified boundary at once.

To define this modified classifying space, we fix a map $\theta : B \rightarrow BO(2n)$ classifying a vector bundle $\theta^*\gamma_{2n}$ over B , and we shall suppose throughout that B is path connected. We shall call a bundle map $TW \rightarrow \theta^*\gamma_{2n}$ a θ -structure on W .

Definition 1.1. Let W be a $2n$ -dimensional manifold with boundary P , and $\hat{\ell}_P : \varepsilon^1 \oplus TP \rightarrow \theta^*\gamma_{2n}$ be a bundle map. Let $\text{Bun}_{n,\partial}^\theta(TW; \hat{\ell}_P)$ denote the space of all bundle maps $\hat{\ell} : TW \rightarrow \theta^*\gamma_{2n}$ which are equal to $\hat{\ell}_P$ when restricted to P and whose underlying map $\ell : W \rightarrow B$ is n -connected. Let

$$\mathcal{N}_n^\theta(P, \hat{\ell}_P) = \coprod_{[W]} \text{Bun}_{n,\partial}^\theta(TW; \hat{\ell}_P) // \text{Diff}_\partial(W),$$

Date: November 11, 2016.

2010 Mathematics Subject Classification. 57R90, 57R15, 57R56, 55P47.

where the disjoint union is taken over compact manifolds W , equipped with a diffeomorphism $\partial W \cong P$, one in each diffeomorphism class relative to P . The topological group $\text{Diff}_\partial(W)$ is the group of diffeomorphisms of W restricting to the identity on a neighbourhood of ∂W , and “//” denotes the homotopy orbit space.

If M is a cobordism between P and Q such that (M, Q) is $(n-1)$ -connected, and is equipped with a θ -structure $\hat{\ell}_M$ restricting to $\hat{\ell}_P$ over P and $\hat{\ell}_Q$ over Q , then there is an induced map

$$(1.1) \quad - \cup_Q (M, \hat{\ell}_M) : \mathcal{N}_n^\theta(Q, \hat{\ell}_Q) \longrightarrow \mathcal{N}_n^\theta(P, \hat{\ell}_P),$$

which on the path components corresponding to W is induced by the $\text{Diff}_\partial(W)$ -equivariant map $\text{Bun}_{n,\partial}^\theta(TW; \hat{\ell}_Q) \rightarrow \text{Bun}_{n,\partial}^\theta(T(W \cup_Q M); \hat{\ell}_P)$ defined by extending bundle maps by $\hat{\ell}_M$, and extending diffeomorphisms of W to $W \cup_Q M$ by the identity map of M .

Our main result will concern the homological effect of (1.1) once all manifolds in sight have been stabilised by connect-sum with countably many copies of $S^n \times S^n$. We shall give a point-set model for forming this stabilisation in Section 2, but for now the following explanation of diagram (1.3) below will suffice.

Let $p : [0, 1] \times D^{2n-1} \hookrightarrow M$ be an embedding sending $\{0\} \times D^{2n-1}$ into Q and $\{1\} \times D^{2n-1}$ into P . (In particular, P and Q must be non-empty.) We may then form the connect-sum

$${}_Q H = ([0, 1] \times Q) \# (S^n \times S^n)$$

inside the image of the embedding $[0, 1] \times D^{2n-1} \hookrightarrow [0, 1] \times Q$ in a standard way; we may analogously form ${}_P H$. Unless P and Q are path connected, the diffeomorphism types of ${}_P H$ and ${}_Q H$ may depend on p and not just P and Q , but we suppress this dependence from the notation. Sliding along the thickened path p induces a diffeomorphism

$$(1.2) \quad M \cup_P ({}_P H) \cong ({}_Q H) \cup_Q M$$

relative to the common boundaries P and Q .

Definition 1.2. Let $W_{1,1} = S^n \times S^n \setminus \text{int}(D^{2n})$ and let us say that a θ -structure $\hat{\ell} : TW_{1,1} \rightarrow \theta^* \gamma$ is *admissible* if there is a pair of orientation-preserving embeddings $e, f : S^n \times D^n \hookrightarrow W_{1,1} \subset S^n \times S^n$ whose cores $e(S^n \times \{0\})$ and $f(S^n \times \{0\})$ intersect transversely in a single point, such that each of the θ -structures $e^* \hat{\ell}$ and $f^* \hat{\ell}$ on $S^n \times D^n$ extend to \mathbb{R}^{2n} for some orientation-preserving embeddings $S^n \times D^n \hookrightarrow \mathbb{R}^{2n}$.

The manifolds ${}_P H$ and ${}_Q H$ each contain an embedded copy of $W_{1,1}$ and we shall say that θ -structures on ${}_P H$ and ${}_Q H$ are admissible if they restrict to admissible θ -structures on $W_{1,1}$. Choosing an admissible θ -structure $\hat{\ell}_{{}_P H}$ on ${}_P H$ which agrees with $\hat{\ell}_P$ on $\{0\} \times P$, we obtain a θ -structure $\hat{\ell}_M \cup \hat{\ell}_{{}_P H}$ on $M \cup_P ({}_P H)$, and hence via the diffeomorphism (1.2) a θ -structure on $({}_Q H) \cup_Q M$. Restricted to ${}_Q H$ this is an admissible θ -structure $\hat{\ell}_{{}_Q H}$ which agrees with $\hat{\ell}_Q$ on $\{1\} \times Q$. Restricted to M it gives a new θ -structure $\hat{\ell}'_M$, which restricts to new θ -structures $\hat{\ell}'_P$ over P and $\hat{\ell}'_Q$ over Q . These manifolds with θ -structure induce maps in the diagram

$$(1.3) \quad \begin{array}{ccc} \mathcal{N}_n^\theta(Q, \hat{\ell}_Q) & \xrightarrow{-\cup_Q ({}_Q H, \hat{\ell}_{{}_Q H})} & \mathcal{N}_n^\theta(Q, \hat{\ell}'_Q) \\ \downarrow -\cup_Q (M, \hat{\ell}_M) & & \downarrow -\cup_Q (M, \hat{\ell}'_M) \\ \mathcal{N}_n^\theta(P, \hat{\ell}_P) & \xrightarrow{-\cup_P ({}_P H, \hat{\ell}_{{}_P H})} & \mathcal{N}_n^\theta(P, \hat{\ell}'_P), \end{array}$$

and the diffeomorphism (1.2) induces a homotopy between the two compositions. Iterating this construction, we obtain an induced map from

$$\mathrm{hocolim}(\mathcal{N}_n^\theta(Q, \hat{\ell}_Q) \rightarrow \mathcal{N}_n^\theta(Q, \hat{\ell}'_Q) \rightarrow \mathcal{N}_n^\theta(Q, \hat{\ell}''_Q) \rightarrow \cdots)$$

to the analogous homotopy colimit for $(P, \hat{\ell}_P)$. This limiting map is denoted $-\cup_Q(M, \hat{\ell}_M^{(\infty)})$ in the following theorem, which is the main new result of this paper.

Theorem 1.3. *The map*

$$-\cup_Q(M, \hat{\ell}_M^{(\infty)}) : \mathrm{hocolim}_{g \rightarrow \infty} \mathcal{N}_n^\theta(Q, \hat{\ell}_Q^{(g)}) \longrightarrow \mathrm{hocolim}_{g \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(g)})$$

induces an isomorphism on homology with all abelian systems of local coefficients.

Remark 1.4. A system of local coefficients \mathcal{L} on a space X is called *abelian* if for each point $x \in X$ the action of $\pi_1(X, x)$ on the fibre $\mathcal{L}(x)$ factors through an abelian group. In particular, all constant coefficient systems are abelian. If X is path connected and based then an abelian system of local coefficients is equivalent to the data of a module over the group ring $\mathbb{Z}[H_1(X; \mathbb{Z})]$.

1.2. Stable homology. The homology of the limiting spaces

$$\mathrm{hocolim}_{g \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(g)})$$

may in many cases be deduced from the results of [GRW14b]. In fact, in Section 7 we shall revisit the results of [GRW14b] and see that the two sections of that paper concerning “surgery on objects” may be replaced by the homological stability results of the present paper, by following a line of argument inspired by [Til97]. We shall see that in all cases the homology of the limiting space agrees with the homology of the infinite loop space of a certain Thom spectrum. This approach using homological stability does not give substantially simpler or easier proofs of the results of [GRW14b], but it does work in more generality. For example, we shall be able to remove the restriction $2n \geq 6$ from [GRW14b].

Recall that we write $\mathrm{MT}\theta$ for the Thom spectrum of the virtual vector bundle $-\theta^* \gamma_{2n}$ over B , and $\Omega^\infty \mathrm{MT}\theta$ for its infinite loop space.

Theorem 1.5. *If $\mathcal{N}_n^\theta(P, \hat{\ell}_P) \neq \emptyset$ then there is a map*

$$\mathrm{hocolim}_{g \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(g)}) \longrightarrow \Omega^\infty \mathrm{MT}\theta$$

which induces an isomorphism on homology, and is in fact acyclic.

In fact, we will construct a specific acyclic map, using a cobordism category \mathcal{C}_θ . The space $\mathcal{N}_n^\theta(P, \hat{\ell}_P)$ will be modelled as a subspace of the space of morphisms $\mathcal{C}_\theta(\emptyset, (P, \hat{\ell}_P))$ in \mathcal{C}_θ , and hence comes with a canonical map

$$\mathcal{N}_n^\theta(P, \hat{\ell}_P) \subset \mathcal{C}_\theta(\emptyset, (P, \hat{\ell}_P)) \longrightarrow \Omega_{[\emptyset, (P, \hat{\ell}_P)]} \mathrm{BC}_\theta,$$

where BC_θ denotes the classifying space (geometric realisation of the nerve) and $\Omega_{[\emptyset, (P, \hat{\ell}_P)]}$ denotes the space of paths between the points corresponding to the objects \emptyset and $(P, \hat{\ell}_P)$. The main theorem of [GMTW09] provides a weak equivalence from this space of paths, when non-empty, to $\Omega^\infty \mathrm{MT}\theta$. This is the map which the theorem is about.

Remark 1.6. A map $f : X \rightarrow Y$ is *acyclic* if for any system of local coefficients \mathcal{L} on Y the induced map $f_* : H_*(X; f^* \mathcal{L}) \rightarrow H_*(Y; \mathcal{L})$ is an isomorphism. This is equivalent to the homotopy fibre of f over any point of Y being acyclic, i.e. having the integral homology of a point.

In the statement of Theorem 1.5 the target has abelian fundamental group. Thus every system of local coefficients is abelian, and so the map being acyclic is

equivalent to inducing an isomorphism on homology with all abelian systems of local coefficients.

Theorem 1.5, and the slightly more general Theorem 7.3 below, improves on [GRW14b, Theorem 1.8] in a few ways. Firstly, it applies to manifolds of all even dimensions, and not just dimensions $2n \geq 6$; secondly, it does not require that θ be “spherical” (a technical condition introduced in [GRW14b]); thirdly, it requires a simpler form of stabilisation (only copies of $W_{1,1}$ have to be glued on, rather than a “universal θ -end”); fourthly, the conclusion is that the map is acyclic rather than merely a homology equivalence. In the special case $2n = 2$ we recover the main result of [MW07]. In that case the acyclicity of the map seems to be new (when $\pi_1 MT\theta \neq 0$), as does the extension to completely general $\theta : B \rightarrow BO(2)$.

1.3. Finite genus and closed manifolds. We can combine Theorem 1.3 with our earlier work [GRW14a] to obtain results about the homology of the spaces $\mathcal{N}_n^\theta(P; \hat{\ell}_P)$ before stabilising by $W_{1,1}$. If we write $\mathcal{M}_n^\theta(W, \hat{\ell}_W) \subset \mathcal{N}_n^\theta(Q, \hat{\ell}_Q)$ for the path component containing $(W, \hat{\ell}_W)$, then the map (1.1) restricts to a map

$$(1.4) \quad - \cup_Q (M, \hat{\ell}_M) : \mathcal{M}_n^\theta(W, \hat{\ell}_W) \longrightarrow \mathcal{M}_n^\theta(W', \hat{\ell}_{W'}),$$

where $W' = W \cup_Q M$ and $\hat{\ell}_{W'}$ is the bundle map obtained by gluing $\hat{\ell}_W$ and $\hat{\ell}_M$. Recall from [GRW14a] that the *genus* of a θ -manifold is defined as

$$g^\theta(W, \hat{\ell}_W) = \max \left\{ g \in \mathbb{N} \mid \begin{array}{l} \text{there are } g \text{ disjoint copies of } W_{1,1} \text{ in } W, \\ \text{each with admissible } \theta\text{-structure} \end{array} \right\}$$

and the *stable genus* is defined as

$$\bar{g}^\theta(W, \hat{\ell}_W) = \max \{ g^\theta(W \sharp W_{k,1}, \hat{\ell}_W^{(k)}) - k \mid k \in \mathbb{N} \},$$

where $W \sharp W_{k,1}$ is the manifold obtained from W by cutting out a disc and forming the boundary connected sum with k copies of $W_{1,1}$. This glued manifold is equipped with a θ -structure $\hat{\ell}_W^{(k)}$ obtained by extending the restriction of $\hat{\ell}_W$ by any admissible structures on the $W_{1,1}$. In Section 8 we shall explain how to deduce the following two finite-genus results from the results in this paper and [GRW14a]. (In fact we shall deduce slightly more general results which imply the corollaries but which also imply some homological stability for $H_0(\mathcal{N}_n^\theta(-); \mathbb{Z})$.)

Corollary 1.7. *Assume $2n \geq 6$, B is simply-connected, and $(M, \hat{\ell}_M)$ is a θ -cobordism between P and Q such that (M, Q) is $(n-1)$ -connected, and denote by $\mathcal{M}_n^\theta(W, \hat{\ell}_W) \subset \mathcal{N}_n^\theta(Q, \hat{\ell}_Q)$ the path component containing $(W, \hat{\ell}_W)$. Write $g = \bar{g}^\theta(W, \hat{\ell}_W)$, and let \mathcal{L} be an abelian coefficient system on the target of the map (1.4). Then the induced map*

$$(- \cup_Q (M, \hat{\ell}_M))_* : H_i(\mathcal{M}_n^\theta(W, \hat{\ell}_W); (- \cup_Q M)^* \mathcal{L}) \longrightarrow H_i(\mathcal{M}_n^\theta(W', \hat{\ell}_{W'}); \mathcal{L})$$

is an isomorphism, provided one of the following three assumptions hold:

- (i) $2i \leq g - 3$, \mathcal{L} is constant, and θ is spherical;
- (ii) $3i \leq g - 4$ and \mathcal{L} is constant;
- (iii) $3i \leq g - 4$, $g \geq 5$, and either $g \geq 7$ or θ is spherical.

The situation is formally quite similar to both [HV04] and [Wah08], from which we learnt the following trick. Suppose one wishes to prove homological stability for a stabilisation map α which increases genus and leaves the boundary unchanged and also for another stabilisation map β which changes the boundary. If α and β commute up to homotopy and stability has already been shown for α , then it suffices to prove stability for β after taking the colimit over stabilising infinitely many times by α .

Corollary 1.8. *There is a map*

$$\mathcal{M}_n^\theta(W, \hat{\ell}_W) \longrightarrow \Omega_0^\infty \text{MT}\theta$$

which, under the same conditions on n and B as in Corollary 1.7, is acyclic in degrees satisfying $3* \leq \bar{g}^\theta(W, \hat{\ell}_W) - 4$, and induces an isomorphism on integral homology in degrees satisfying $2* \leq \bar{g}^\theta(W, \hat{\ell}_W) - 3$ if in addition θ is spherical.

As in Theorem 1.5 and the discussion following it, Corollary 1.8 concerns a specific map into a component of a path space of BC_θ , which by [GMTW09] may be identified with the basepoint component $\Omega_0^\infty \text{MT}\theta$ of the infinite loop space $\Omega^\infty \text{MT}\theta$.

1.4. Other tangential structures. The definition of $\mathcal{N}_n^\theta(P, \hat{\ell}_P)$ used a space $\text{Bun}_{n, \partial}^\theta(TW; \hat{\ell}_P)$, consisting of all n -connected bundle maps with fixed boundary condition. If we replace this by the space of all (i.e. not necessarily n -connected) bundle maps with fixed boundary condition, then the results of this paper may still be used to deduce the homology of the resulting moduli space in a stable range. Again we state the result one path component at a time, and for simplicity we shall restrict attention to closed manifolds in this introduction.

Let $\theta' : B' \rightarrow BO(2n)$ be a map, and let W be a closed $2n$ -dimensional manifold and $\hat{\ell}'_W : TW \rightarrow (\theta')^* \gamma$ a θ' -structure. We shall write $\text{Bun}^{\theta'}(TW)$ for the space of all (not necessarily n -connected) θ' -structures and

$$\mathcal{M}^{\theta'}(W, \hat{\ell}'_W) \subset \text{Bun}^{\theta'}(TW) // \text{Diff}(W)$$

for the path component containing the image of $\hat{\ell}'_W$. Let $\ell'_W : W \rightarrow B'$ be the map underlying $\hat{\ell}'_W$ and let $W \rightarrow B \rightarrow B'$ be a Moore–Postnikov n -stage factorisation of ℓ'_W , i.e. a factorisation into an n -connected cofibration $\ell_W : W \rightarrow B$ and an n -co-connected fibration $u : B \rightarrow B'$. Let us write $\theta = \theta' \circ u : B \rightarrow BO(2n)$. In Section 9 below, we explain how to deduce the following result from Corollary 1.8.

Corollary 1.9. *In this situation the topological monoid $\text{hAut}(u)$, consisting of weak equivalences $B \rightarrow B$ commuting with the map $u : B \rightarrow B'$, acts on the space $\Omega^\infty \text{MT}\theta$ and there is a map*

$$\mathcal{M}^{\theta'}(W, \hat{\ell}'_W) \longrightarrow (\Omega^\infty \text{MT}\theta) // \text{hAut}(u).$$

Considered as a map to the path component which it hits, if $2n \geq 6$ and W is simply-connected then this map is acyclic in degrees satisfying $3* \leq \bar{g}^{\theta'}(W, \hat{\ell}'_W) - 4$, and induces an isomorphism in integral homology in degrees satisfying $2* \leq \bar{g}^{\theta'}(W, \hat{\ell}'_W) - 3$ if in addition θ' is spherical.

The proof given in Section 9 again concerns a specific map, and we will also treat the case $\partial W \neq \emptyset$. Let us briefly spell out two important special cases.

Taking $\theta' = \text{Id} : BO(2n) \rightarrow BO(2n)$ and letting $W \xrightarrow{\ell_W} B \xrightarrow{u} BO(2n)$ be a Moore–Postnikov n -stage factorisation of a Gauss map $\ell'_W : W \rightarrow BO(2n)$, the space $\mathcal{M}^{\theta'}(W, \hat{\ell}'_W)$ is a model for $B\text{Diff}(W)$ and we obtain a map

$$(1.5) \quad B\text{Diff}(W) \longrightarrow (\Omega^\infty \text{MT}\theta) // \text{hAut}(u)$$

which, as long as $2n \geq 6$ and W is simply-connected, is a homology isomorphism (respectively acyclic) in a range of degrees depending on $g(W)$, onto the path component which it hits.

Similarly we can take $\theta' : BSO(2n) \rightarrow BO(2n)$ to be the orientation cover and let $W \xrightarrow{\ell_W} B \xrightarrow{u^+} BSO(2n)$ be a Moore–Postnikov n -stage factorisation of an *oriented* Gauss map $\ell'_W : W \rightarrow BSO(2n)$. Then the space $\mathcal{M}^{\theta'}(W, \hat{\ell}'_W)$ is a model for $B\text{Diff}^+(W)$ and we obtain a map

$$B\text{Diff}^+(W) \longrightarrow (\Omega^\infty \text{MT}\theta) // \text{hAut}(u^+)$$

which, as long as $2n \geq 6$ and W is simply-connected, is a homology isomorphism (respectively acyclic) in a range of degrees depending on $g(W)$, onto the path component which it hits.

1.5. Acknowledgements. S. Galatius was partially supported by NSF grants DMS-1105058 and DMS-1405001, as well as the Danish National Research Foundation through the Centre for Symmetry and Deformation (DNRF92) and ERC-682992. O. Randal-Williams was partially supported by the Herchel Smith Fund and EPSRC grant EP/M027783/1.

CONTENTS

1. Introduction and statement of results	1
1.1. Statement of the main theorem	1
1.2. Stable homology	3
1.3. Finite genus and closed manifolds	4
1.4. Other tangential structures	5
1.5. Acknowledgements	6
2. Definitions	7
2.1. Cobordism categories	7
2.2. Standard θ -structures	9
2.3. Stable homological stability	10
2.4. Models for moduli spaces of manifolds	10
2.5. Elementary simplifications of Theorem 2.15	13
3. Proof of Theorem 2.15: stability for $W_{1,1}$	15
4. Proof of Theorem 2.15: stability for handles of index n	16
4.1. Support and interchange of support	16
4.2. Constructing auxiliary cobordisms	17
4.3. A semi-simplicial resolution	20
4.4. Resolving composition with $M \circ r(M)$	21
5. Contractibility of the higher-dimensional arc complex	25
6. Proof of Theorem 2.15: stability for handles of index $n < k < 2n$	34
6.1. Constructing auxiliary cobordisms	34
6.2. A semi-simplicial resolution	35
6.3. Resolving composition with M	37
7. Stable homology and group-completion	38
7.1. The group-completion argument	39
7.2. Proof of Theorem 7.3	43
8. Finite genus and stability for closed manifolds	45
8.1. Finite genus and homological stability	45
8.2. Proof of Corollary 1.7 for $P \neq \emptyset$	46
8.3. Proof of Corollary 1.7 in general	47
8.4. Stable homology and proof of Corollary 1.8	48
9. General tangential structures	49
Appendix A. Homology equivalences and local coefficients	52
A.1. Local coefficient systems	52
A.2. Invariance properties for abelian coefficients	54
A.3. Ladder diagrams and abelian equivalences	56
A.4. The generalised group completion theorem with local coefficients	57
References	59

2. DEFINITIONS

In the following we shall suppose that $n > 0$. We emphasise that we allow $n = 1$ and $n = 2$, and we will point out when these cases need special treatment.

2.1. Cobordism categories. It is convenient to state the precise version of Theorem 1.3 using the language of cobordism categories; these are also useful for the proof of Theorem 1.3, and will be essential for the proof of Theorem 1.5. See [GMTW09, GRW10, GRW14b] for further background on these categories. Throughout the paper we shall work with manifolds inside $\mathbb{R} \times \mathbb{R}^\infty$, and will write e_0 for the basis vector of the first copy of \mathbb{R} . Apart from this, we will write e_1, e_2, \dots, e_k for the basis vectors of $\mathbb{R}^k \subset \mathbb{R}^\infty$.

Definition 2.1. Let \mathcal{C}_θ be the (non-unital) category with objects given by

- (i) a $(2n - 1)$ -dimensional closed submanifold $P \subset \mathbb{R}^\infty$, and
- (ii) a bundle map $\hat{\ell}_P : \varepsilon^1 \oplus TP \rightarrow \theta^* \gamma$.

A morphism from $(P, \hat{\ell}_P)$ to $(Q, \hat{\ell}_Q)$ is given by

- (i) an $s \in (0, \infty)$,
- (ii) a $2n$ -dimensional submanifold $W \subset [0, s] \times \mathbb{R}^\infty$ which for some $\varepsilon > 0$ intersects

$$([0, \varepsilon] \times \mathbb{R}^\infty) \cup ((s - \varepsilon, s] \times \mathbb{R}^\infty)$$

in $([0, \varepsilon] \times P) \cup ((s - \varepsilon, s] \times Q)$,

- (iii) a bundle map $\hat{\ell}_W : TW \rightarrow \theta^* \gamma$ which restricts to $\hat{\ell}_P$ on $\{0\} \times P$ and to $\hat{\ell}_Q$ on $\{s\} \times Q$.

The composition of $(s, W, \hat{\ell}_W) : (P, \hat{\ell}_P) \rightsquigarrow (Q, \hat{\ell}_Q)$ and $(s', W', \hat{\ell}_{W'}) : (Q, \hat{\ell}_Q) \rightsquigarrow (R, \hat{\ell}_R)$ is denoted $(s', W', \hat{\ell}_{W'}) \circ (s, W, \hat{\ell}_W)$ and defined as

$$(s + s', W \cup (W' + s \cdot e_0), \hat{\ell}_W \cup \hat{\ell}_{W'}) : (P, \hat{\ell}_P) \rightsquigarrow (R, \hat{\ell}_R).$$

We make \mathcal{C}_θ into a topological (non-unital) category using the evident isomorphism with the topological (non-unital) category described in [GRW14b, Definition 2.6]. In particular, the topology on the morphism space $\mathcal{C}_\theta((P, \hat{\ell}_P), (Q, \hat{\ell}_Q))$ is that of a subquotient of

$$(0, \infty) \times \left(\coprod_{[W]} \text{Emb}(W, [0, 1] \times \mathbb{R}^\infty) \times \text{Bun}^\theta(TW) \right),$$

where W ranges over the set of all cobordisms from P to Q , one in each diffeomorphism class.

Definition 2.2. Let $L \subset (-\infty, 0] \times \mathbb{R}^{\infty-1}$ be a compact $(2n - 1)$ -dimensional submanifold which near $\{0\} \times \mathbb{R}^{\infty-1}$ coincides with $(-\infty, 0] \times \partial L$, and let $\hat{\ell}_L : \varepsilon^1 \oplus TL \rightarrow \theta^* \gamma_{2n}$ be a θ -structure.

Let $\mathcal{C}_{\theta, L} \subset \mathcal{C}_\theta$ be the subcategory with

- (i) objects those $(P, \hat{\ell}_P)$ such that $P \cap ((-\infty, 0] \times \mathbb{R}^{\infty-1}) = L$ and $\hat{\ell}_P|_L = \hat{\ell}_L$,
- (ii) morphisms those $(s, W, \hat{\ell}_W)$ such that

$$W \cap ([0, s] \times ((-\infty, 0] \times \mathbb{R}^{\infty-1})) = [0, s] \times L$$

and $\hat{\ell}_W$ restricts to the bundle map $T([0, s] \times L) \rightarrow \varepsilon^1 \oplus TL \xrightarrow{\hat{\ell}_L} \theta^* \gamma_{2n}$.

We shall generally abbreviate θ -manifolds by the name of the underlying smooth manifold, say X , and write $\hat{\ell}_X$ for its θ -structure. Using this convention, for any

object $P \in \mathcal{C}_\theta$ and any $s \in (0, \infty)$ there is an morphism $(s, [0, s] \times P)$, having underlying manifold $[0, s] \times P \subset [0, s] \times \mathbb{R}^\infty$ and θ -structure given by

$$\hat{\ell}_{[0, s] \times P} : T([0, s] \times P) \longrightarrow \varepsilon^1 \oplus TP \xrightarrow{\hat{\ell}_P} \theta^* \gamma_{2n}.$$

We shall call such morphisms ‘‘cylindrical’’. We can make the analogous construction for objects in $\mathcal{C}_{\theta, L}$.

Remark 2.3. In this paper it suffices to work with the non-unital category \mathcal{C}_θ defined above. If identities are desired, they could be added by replacing the morphism space with the pushout $[0, \infty) \times \text{Ob}(\mathcal{C}_\theta) \leftarrow (0, \infty) \times \text{Ob}(\mathcal{C}_\theta) \rightarrow \text{Mor}(\mathcal{C}_\theta)$, along the map $(0, \infty) \times \text{Ob}(\mathcal{C}_\theta) \rightarrow \text{Mor}(\mathcal{C}_\theta)$ defined by sending $(t, P, \hat{\ell}_P)$ to $(t, [0, t] \times P, \hat{\ell}_P)$.

This motivates calling a morphism $M : P \rightsquigarrow Q$ in \mathcal{C}_θ an *isomorphism* if there exists a morphism $N : Q \rightsquigarrow P$ such that $M \circ N$ and $N \circ M$ are in the same path components of $\mathcal{C}_\theta(Q, Q)$ and $\mathcal{C}_\theta(P, P)$ as the cylindrical morphisms $[0, s] \times Q$ and $[0, s] \times P$. Pre- and postcomposing with an isomorphism $M : P \rightsquigarrow Q$ then define weak equivalences $\mathcal{C}_\theta(Q, R) \rightarrow \mathcal{C}_\theta(P, R)$ and $\mathcal{C}_\theta(R, P) \rightarrow \mathcal{C}_\theta(R, Q)$. Similarly for $\mathcal{C}_{\theta, L}$. The most important examples of isomorphisms will be cobordisms whose underlying manifold is diffeomorphic to a cylinder, but with non-cylindrical θ -structure or embedding.

Definition 2.4 ([GRW14b]). Let $\mathcal{C}_{\theta, L}^{n-1} \subset \mathcal{C}_{\theta, L}$ be the subcategory having the same objects, and having those morphisms $(s, M) : P \rightsquigarrow Q$ for which the pair (M, Q) is $(n-1)$ -connected.

The categories $\mathcal{C}_{\theta, L}$ do not really depend on L , but rather only on ∂L : if L' is another collared θ -manifold which is equal to L near its boundary, then we obtain an isomorphism of topological categories $\mathcal{C}_{\theta, L} \xrightarrow{\sim} \mathcal{C}_{\theta, L'}$ by cutting out L and gluing in L' (and similarly on morphisms). In order to make use of this, we make the following definition.

Definition 2.5. If P is an object of $\mathcal{C}_{\theta, L}$, let $P^\circ \subset [0, \infty) \times \mathbb{R}^{\infty-1}$ be the θ -manifold obtained by cutting out $\text{int}(L)$. Similarly, if (s, W) is a morphism in $\mathcal{C}_{\theta, L}$, let $W^\circ \subset [0, s] \times [0, \infty) \times \mathbb{R}^{\infty-1}$ be the θ -cobordism obtained by cutting out $[0, s] \times \text{int}(L)$.

Let $\mathcal{C}_{\theta, \partial L}$ be the (non-unital) category defined by

$$\text{Ob}(\mathcal{C}_{\theta, \partial L}) = \{(P^\circ, \hat{\ell}_{P^\circ}) \mid (P, \hat{\ell}_P) \in \text{Ob}(\mathcal{C}_{\theta, L})\}$$

$$\text{Mor}(\mathcal{C}_{\theta, \partial L}) = \{(s, W^\circ, \hat{\ell}_{W^\circ}) \mid (s, W, \hat{\ell}_W) \in \text{Mor}(\mathcal{C}_{\theta, L})\}$$

and made into a topological category by insisting that the functor $\mathcal{C}_{\theta, L} \rightarrow \mathcal{C}_{\theta, \partial L}$ given by $X \mapsto X^\circ$ is an isomorphism of topological categories.

Definition 2.6. Let $\mathcal{C}_{\theta, \partial L}^{n-1} \subset \mathcal{C}_{\theta, \partial L}$ be the subcategory having the same objects, and having those morphisms $(s, M^\circ) : P^\circ \rightsquigarrow Q^\circ$ for which the pair (M°, Q°) is $(n-1)$ -connected.

This topological category shall play an especially important role and for brevity we shall often write $\mathcal{D} = \mathcal{C}_{\theta, \partial L}^{n-1}$ when θ and ∂L are fixed and understood. When using the notation \mathcal{D} , we shall implicitly assume that $\partial L \neq \emptyset$.

Remark 2.7. If $(s, M^\circ) : P^\circ \rightsquigarrow Q^\circ \in \mathcal{C}_{\theta, \partial L}^{n-1}$ then the cobordism $M = ([0, s] \times L) \cup M^\circ$ is $(n-1)$ -connected relative to $Q = L \cup Q^\circ$, which defines a functor $\mathcal{C}_{\theta, \partial L}^{n-1} \rightarrow \mathcal{C}_{\theta, L}^{n-1}$. This is not in general an isomorphism of categories, but if $(L, \partial L)$ is $(n-1)$ -connected then it is. (See [GRW14b, Lemma 7.2] for a related discussion.)

We shall generally write P, Q, \dots for objects and W, M, \dots for morphisms of $\mathcal{C}_{\theta, \partial L}$ rather than $\mathcal{C}_{\theta, L}$, and reserve the $(-)^\circ$ notation for when we are directly comparing elements in $\mathcal{C}_{\theta, L}$ with their corresponding elements in $\mathcal{C}_{\theta, \partial L}$.

For use in Section 7, let us recall the following result from [GRW14b]. Together with the isomorphism $\mathcal{D} \cong \mathcal{C}_{\theta,L}^{n-1}$ from Remark 2.7 and the equivalence $BC_\theta \simeq \Omega^{\infty-1}MT\theta$ from [GMTW09], it determines the weak homotopy type of $B\mathcal{D}$ for $(n-1)$ -connected $(L, \partial L)$.

Theorem 2.8 ([GRW14b]). *The maps*

$$(2.1) \quad BC_{\theta,L}^{n-1} \longrightarrow BC_{\theta,L} \longrightarrow BC_\theta,$$

induced by the inclusion functors, are weak equivalences.

Proof. This first equivalence is [GRW14b, Theorem 3.1], one of the major technical results of that paper (“surgery on morphisms”). The second is Corollary 2.17 of op. cit. \square

2.2. Standard θ -structures. Let $\theta : B \rightarrow BO(2n)$ be a tangential structure, and recall that we insist that B is path connected. Choose once and for all a “basepoint” θ -structure on the vector space \mathbb{R}^{2n} , i.e. a bundle map $\tau : \mathbb{R}^{2n} \rightarrow \theta^*\gamma_{2n}$ (depending on the orientability of $\theta^*\gamma_{2n}$, there are one or two possible such choices, up to homotopy). This induces a canonical θ -structure $\hat{\ell}_X^\tau$ on any framed $2n$ - or $(2n-1)$ -manifold X .

Definition 2.9. The *standard framing* $\xi_{S^n \times D^n}$ on $S^n \times D^n$ is the framing induced by the embedding

$$\begin{aligned} S^n \times D^n &\longrightarrow \mathbb{R}^{n+1} \times \mathbb{R}^{n-1} = \mathbb{R}^{2n} \\ (x; y_1, y_2, \dots, y_n) &\longmapsto (2e^{-\frac{y_1}{2}} x; y_2, \dots, y_n). \end{aligned}$$

This framing induces a θ -structure $\hat{\ell}_{S^n \times D^n}^\tau$, and we say that a θ -structure on $S^n \times D^n$ is *standard* if it is homotopic to $\hat{\ell}_{S^n \times D^n}^\tau$.

Definition 2.10. Let $W_{1,1} = S^n \times S^n \setminus \text{int}(D^{2n})$ be obtained by removing an open disc inside $D_+^n \times D_+^n$, the product of the two upper hemispheres. A θ -structure $\hat{\ell}$ on $W_{1,1}$ is *standard* if both θ -structures $\bar{e}^*\hat{\ell}$ and $\bar{f}^*\hat{\ell}$ on $S^n \times D^n$ are standard, where \bar{e} and \bar{f} are the embeddings defined by

$$\begin{aligned} \bar{e} : S^n \times D^n &\longrightarrow W_{1,1} \subset S^n \times S^n \subset \mathbb{R}^{n+1} \times \mathbb{R}^{n+1} \\ (x, y) &\longmapsto \left(x; \frac{y}{2}, -\sqrt{1 - \left|\frac{y}{2}\right|^2} \right) \end{aligned}$$

and

$$\begin{aligned} \bar{f} : S^n \times D^n &\longrightarrow W_{1,1} \subset S^n \times S^n \subset \mathbb{R}^{n+1} \times \mathbb{R}^{n+1} \\ (x, y) &\longmapsto \left(-\frac{y}{2}, -\sqrt{1 - \left|\frac{y}{2}\right|^2}; x \right), \end{aligned}$$

Remark 2.11. In Definition 1.2 we said that a θ -structure $\hat{\ell}$ on $W_{1,1}$ was *admissible* if there are orientation-preserving embeddings $e, f : S^n \times D^n \rightarrow W_{1,1}$ with cores intersecting transversely in one point, such that each of the θ -structures $e^*\hat{\ell}$ and $f^*\hat{\ell}$ on $S^n \times D^n$ extend to \mathbb{R}^{2n} for some orientation-preserving embeddings $S^n \times D^n \hookrightarrow \mathbb{R}^{2n}$. In particular, as $\bar{e}(S^n \times D^n) \cup \bar{f}(S^n \times D^n) \subset W_{1,1}$ is an isotopy equivalence it follows that if $\hat{\ell}$ is admissible then there is an embedding $\phi : W_{1,1} \hookrightarrow W_{1,1}$ such that $\phi^*\hat{\ell}$ is standard.

The embeddings \bar{e} and \bar{f} are orientation-preserving and the union of their images is isotopy equivalent to $W_{1,1}$. Therefore there exists a framing $\xi_{W_{1,1}}$ of $W_{1,1}$ such that $\bar{e}^*\xi_{W_{1,1}}$ and $\bar{f}^*\xi_{W_{1,1}}$ are each homotopic to $\xi_{S^n \times D^n}$, so the associated θ -structure $\hat{\ell}_{W_{1,1}}^\tau$ is standard.

Lemma 2.12. *The space of standard θ -structures on $W_{1,1}$ fixed on a disc in the boundary is path-connected.*

Proof. This is [GRW14a, Lemma 7.7], under a choice of identification of $W_{1,1}$ with the manifold H of that paper. \square

2.3. Stable homological stability. In this section we shall formulate our stable homological stability theorem, which is a strengthening of Theorem 1.3.

Definition 2.13. We say that a composable sequence of cobordisms

$$K|_0 \xrightarrow{K|_{[0,1]}} K|_1 \xrightarrow{K|_{[1,2]}} K|_2 \xrightarrow{K|_{[2,3]}} K|_3 \rightsquigarrow \dots$$

in $\mathcal{C}_{\theta, \partial L}$ is a θ -end in $\mathcal{C}_{\theta, \partial L}$ if each cobordism $K|_{[i, i+1]}$ satisfies

- (i) it is $(n-1)$ -connected relative to both $K|_i$ and $K|_{i+1}$, and
- (ii) it contains an embedded copy of $W_{1,1}$ with standard θ -structure, in a path component intersecting ∂L .

For a θ -end K in $\mathcal{C}_{\theta, \partial L}$, let

$$\mathcal{F}_i(P) \subset \mathcal{C}_{\theta, \partial L}(P, K|_i)$$

be the subspace consisting of those θ -cobordisms (s, M) such that $\ell_M : M \rightarrow B$ is n -connected.

Lemma 2.14. *The map $\mathcal{C}_{\theta, \partial L}(Q, K|_i) \rightarrow \mathcal{C}_{\theta, \partial L}(P, K|_i)$ induced by precomposing with any $M \in \mathcal{D}(P, Q)$ sends $\mathcal{F}_i(Q)$ into $\mathcal{F}_i(P)$. Similarly, postcomposing with $K|_{[i, i+1]}$ induces compatible maps $\mathcal{F}_i(P) \rightarrow \mathcal{F}_{i+1}(P)$.*

In other words, \mathcal{F}_i defines a subfunctor of $\mathcal{C}_{\theta, \partial L}(-, K|_i) : \mathcal{D}^{op} \rightarrow \text{Top}$ and postcomposition with $K|_{[i, i+1]}$ defines a natural transformation $\mathcal{F}_i \Rightarrow \mathcal{F}_{i+1}$.

Proof. If $(s, W) \in \mathcal{F}_i(Q)$ and $(t, M) : P \rightsquigarrow Q \in \mathcal{D}$, then we must show that the map $\ell_M \cup \ell_W : M \cup_Q W \rightarrow B$ is n -connected. As (M, Q) is $(n-1)$ -connected, so is $(M \cup_Q W, W)$, and hence we have a factorisation

$$\ell_W : W \longrightarrow M \cup_Q W \xrightarrow{\ell_M \cup \ell_W} B$$

of an n -connected map where the first map is $(n-1)$ -connected: it follows that the second map is n -connected too.

For the second claim, given $(s, W) \in \mathcal{F}_i(P)$ we must show that $(1+s, K|_{[i, i+1]} \circ W) \in \mathcal{F}_{i+1}(P)$, i.e. that $\ell_{K|_{[i, i+1]} \circ W}$ is n -connected. This follows from the same argument as above, using that $(K|_{[i, i+1]}, K|_i)$ is $(n-1)$ -connected. \square

Theorem 1.3 can now be stated economically as follows.

Theorem 2.15 (Stable homological stability). *For any fixed θ -end in $\mathcal{C}_{\theta, \partial L}$ as in Definition 2.13, write $\mathcal{F}(P) = \text{hocolim}_{i \rightarrow \infty} \mathcal{F}_i(P)$. Then the map $\mathcal{F}(Q) \rightarrow \mathcal{F}(P)$ induced by any $M \in \mathcal{D}(P, Q)$ is an abelian homology equivalence.*

2.4. Models for moduli spaces of manifolds. In order to explain how to deduce Theorem 1.3 from Theorem 2.15, we will make a particular choice of model for the Borel constructions $\text{Bun}_{n, \partial}^\theta(TW; \hat{\ell}_P) // \text{Diff}_\partial(W)$ which form path components of $\mathcal{N}_n^\theta(P; \hat{\ell}_P)$. This involves a comparison between two models for spaces of manifolds which will be of further use in Section 7, so we give it in full generality.

In [GRW14b, Proposition 2.16] we have explained how to start with a $(2n-1)$ -manifold $L \subset \{0\} \times (-\infty, 0] \times \mathbb{R}^{\infty-1}$ with collared boundary $\partial L \subset \{0\} \times \{0\} \times \mathbb{R}^{\infty-1}$, and rotate it in $(-\infty, 0] \times \mathbb{R} \times \mathbb{R}^{\infty-1}$ around the origin in the first two coordinate directions to obtain a new manifold $\bar{L} \subset \{0\} \times [0, \infty) \times \mathbb{R}^{\infty-1}$. The union $D(L) = L \cup \bar{L}$ is the double of L , and the subset $V_L \subset (-\infty, 0] \times \mathbb{R} \times \mathbb{R}^{\infty-1}$ swept out by

the rotation gives a nullbordism $V_L : \emptyset \rightsquigarrow D(L)$. The manifold V_L is diffeomorphic to $[0, 1] \times L$, after unbending corners, and in particular the inclusion $L \hookrightarrow V_L$ is a homotopy equivalence. Therefore a θ -structure $\hat{\ell}_L$ on $\varepsilon^1 \oplus TL$ may be extended to one, $\hat{\ell}_{V_L}$, on V_L . We then give \bar{L} the θ -structure induced by restriction from V_L .

We may equally well begin with a $(2n - 1)$ -manifold $P \subset \{0\} \times [0, \infty) \times \mathbb{R}^{\infty-1}$ with collared boundary $\partial P \subset \{0\} \times \{0\} \times \mathbb{R}^{\infty-1}$ and rotate it in $(-\infty, 0] \times \mathbb{R} \times \mathbb{R}^{\infty-1}$. We write $\bar{P} \subset \{0\} \times (-\infty, 0] \times \mathbb{R}^{\infty-1}$ for the manifold so obtained, and induce a θ -structure on \bar{P} from one on P in the same way.

In particular, if $P \in \mathcal{C}_{\theta, \partial L}$ is an object then $\bar{P} \subset \{0\} \times (-\infty, 0] \times \mathbb{R}^{\infty-1}$ is equal (as a θ -manifold) to L near their boundary, so there are isomorphisms

$$\mathcal{C}_{\theta, L} \cong \mathcal{C}_{\theta, \partial L} = \mathcal{C}_{\theta, \partial \bar{P}} \cong \mathcal{C}_{\theta, \bar{P}}.$$

A similar procedure associates to a $2n$ -manifold $N \subset [0, s] \times [0, \infty) \times \mathbb{R}^{\infty-1}$ with collared boundary $(\{0\} \times P) \cup ([0, s] \times \partial L) \cup (\{s\} \times P)$ a rotated submanifold $\bar{N} \subset [0, s] \times (-\infty, 0] \times \mathbb{R}^{\infty-1}$ together with a rotated θ -structure $\hat{\ell}_{\bar{N}}$.

Definition 2.16. For objects P and Q of $\mathcal{C}_{\theta, \partial L}$, let $\langle P, Q \rangle = \bar{P} \cup Q \in \mathcal{C}_{\theta, \bar{P}}$. For a morphism $(s, M) : Q \rightsquigarrow Q'$ let $\langle P, M \rangle = ([0, s] \times \bar{P}) \cup M : \langle P, Q \rangle \rightsquigarrow \langle P, Q' \rangle$, and for a morphism $(s, N) : P' \rightsquigarrow P$ let $\langle N, Q \rangle = \bar{N} \cup ([0, s] \times Q) : \langle P, Q \rangle \rightsquigarrow \langle P', Q \rangle$. Both are morphisms in \mathcal{C}_{θ} .

Lemma 2.17. For any $P, Q \in \mathcal{C}_{\theta, \partial L}$ the composition

$$(2.2) \quad \mathcal{C}_{\theta, \partial L}(P, Q) = \mathcal{C}_{\theta, \partial \bar{P}}(P, Q) \hookrightarrow \mathcal{C}_{\theta}(D(P), \langle P, Q \rangle) \xrightarrow{-\circ V_{\bar{P}}} \mathcal{C}_{\theta}(\emptyset, \langle P, Q \rangle)$$

is a weak homotopy equivalence.

If $M : Q \rightsquigarrow Q'$ and $N : P' \rightsquigarrow P$ are morphisms in $\mathcal{C}_{\theta, \partial L}$ then the squares

$$\begin{array}{ccc} \mathcal{C}_{\theta, \partial L}(P, Q) & \xrightarrow{M \circ -} & \mathcal{C}_{\theta, \partial L}(P, Q') & & \mathcal{C}_{\theta, \partial L}(P, Q) & \xrightarrow{-\circ N} & \mathcal{C}_{\theta, \partial L}(P', Q) \\ \downarrow \simeq & & \downarrow \simeq & & \downarrow \simeq & & \downarrow \simeq \\ \mathcal{C}_{\theta}(\emptyset, \langle P, Q \rangle) & \xrightarrow{\langle P, M \rangle \circ -} & \mathcal{C}_{\theta}(\emptyset, \langle P, Q' \rangle) & & \mathcal{C}_{\theta}(\emptyset, \langle P, Q \rangle) & \xrightarrow{\langle N, Q \rangle \circ -} & \mathcal{C}_{\theta}(\emptyset, \langle P', Q \rangle) \end{array}$$

commute up to homotopy.

Proof sketch. Up to smoothing corners, (2.2) is given by gluing on an invertible cobordism. The squares each commute because the two compositions are given by gluing on cobordisms which are diffeomorphic, by a diffeomorphism preserving θ -structures up to homotopy. For more details, see the proof of [GRW14b, Proposition 7.5]. \square

For an object $P \in \mathcal{C}_{\theta}$ we have a homeomorphism

$$\mathcal{C}_{\theta}(\emptyset, P) \cong \coprod_{[W]} \left((0, \infty) \times \text{Emb}_{\partial}(W, (-1, 0] \times \mathbb{R}^{\infty}) \times \text{Bun}_{\partial}^{\theta}(TW; \hat{\ell}_P) \right) / \text{Diff}_{\partial}(W),$$

where the disjoint union is taken over manifolds W with boundary P , one in each diffeomorphism class. As the action of $\text{Diff}_{\partial}(W)$ on $\text{Emb}_{\partial}(W, (-1, 0] \times \mathbb{R}^{\infty})$ is free and has slices [BF81], and $(0, \infty)$ is contractible, this quotient is a model for the Borel construction so there is a weak equivalence

$$\mathcal{C}_{\theta}(\emptyset, P) \simeq \coprod_{[W]} \text{Bun}_{\partial}^{\theta}(TW; \hat{\ell}_P) // \text{Diff}_{\partial}(W),$$

where the disjoint union is taken over manifolds W with boundary P , one in each diffeomorphism class.

The spaces $\mathcal{N}_n^{\theta}(P)$ and $\mathcal{M}_n^{\theta}(W, \hat{\ell}_W)$ in the following definition are point-set models for the spaces $\mathcal{N}_n^{\theta}(P)$ and $\mathcal{M}_n^{\theta}(W, \hat{\ell}_W)$ in the introduction.

Definition 2.18.

- (i) For $P \in \mathcal{C}_\theta$ let $\mathcal{N}_n^\theta(P) \subset \mathcal{C}_\theta(\emptyset, P)$ denote the subspace of those nullbordisms (s, W) of P such that $\ell_W : W \rightarrow B$ is n -connected.
- (ii) For $(W, \hat{\ell}_W) \in \mathcal{N}_n^\theta(P)$, let $\mathcal{M}_n^\theta(W, \hat{\ell}_W)$ denote the path component of $\mathcal{N}_n^\theta(P)$ containing $(W, \hat{\ell}_W)$.
- (iii) For a manifold W with boundary P , let $\mathcal{M}^\theta(W; \hat{\ell}_P) \subset \mathcal{C}_\theta(\emptyset, P)$ be the subspace of those $(X, \hat{\ell}_X)$ such that X is diffeomorphic to W relative to P .

Proof of Theorem 1.3, using Theorem 2.15. We may embed M in $[0, 1] \times \mathbb{R}^\infty$ as a cobordism $Q \rightsquigarrow P$, and after changing this embedding by an isotopy we may suppose that M intersects $[0, 1] \times [0, \infty) \times \mathbb{R}^{\infty-1}$ precisely in the image of $p : [0, 1] \times D^{2n-1} \hookrightarrow M$ and that $[0, 1] \times D^{2n-1} \xrightarrow{p} M \subset [0, 1] \times \mathbb{R}^\infty$ is given by $(t, x) \mapsto (t, e(x))$ for an embedding $e : D^{2n-1} \rightarrow [0, \infty) \times \mathbb{R}^{\infty-1}$.

Let $L_Q = Q \cap ((-\infty, 0] \times \mathbb{R}^{\infty-1})$, $L_P = P \cap ((-\infty, 0] \times \mathbb{R}^{\infty-1})$, and $D = P \cap ([0, \infty) \times \mathbb{R}^{\infty-1})$. Then L_P , L_Q , and D all have equal boundaries, which we call ∂L . By rotating these and the manifold $N = M \cap ([0, 1] \times (-\infty, 0] \times \mathbb{R}^{\infty-1})$ in the first two coordinate directions, we obtain a cobordism

$$\overline{N} : \overline{L_P} \rightsquigarrow \overline{L_Q} \in \mathcal{C}_{\theta, \partial L}.$$

Now let

$$D = K|_0 \xrightarrow{K|_{[0,1]}} K|_1 \xrightarrow{K|_{[1,2]}} K|_2 \rightsquigarrow \dots$$

be a θ -end in the category $\mathcal{C}_{\theta, \partial L}$, where each $K|_{[i, i+1]}$ is obtained from $[0, 1] \times K|_i$ by the boundary connect-sum at $\{1\} \times K|_i$ with $W_{1,1}$ having a standard θ -structure. There is then a commutative diagram

$$(2.3) \quad \begin{array}{ccccc} \mathcal{C}_{\theta, \partial L}(\overline{L_Q}, K|_0) & \xrightarrow{K|_{[0,1]}^{\circ-}} & \mathcal{C}_{\theta, \partial L}(\overline{L_Q}, K|_1) & \xrightarrow{K|_{[1,2]}^{\circ-}} & \mathcal{C}_{\theta, \partial L}(\overline{L_Q}, K|_2) \cdots \\ \downarrow -\circ \overline{N} & & \downarrow -\circ \overline{N} & & \downarrow -\circ \overline{N} \\ \mathcal{C}_{\theta, \partial L}(\overline{L_P}, K|_0) & \xrightarrow{K|_{[0,1]}^{\circ-}} & \mathcal{C}_{\theta, \partial L}(\overline{L_P}, K|_1) & \xrightarrow{K|_{[1,2]}^{\circ-}} & \mathcal{C}_{\theta, \partial L}(\overline{L_P}, K|_2) \cdots \end{array}$$

which Lemma 2.17 shows is equivalent to the homotopy commutative diagram

$$(2.4) \quad \begin{array}{ccccc} \mathcal{C}_\theta(\emptyset, Q) & \xrightarrow{Q^H \circ -} & \mathcal{C}_\theta(\emptyset, Q') & \xrightarrow{Q'^H \circ -} & \mathcal{C}_\theta(\emptyset, Q'') \cdots \\ \downarrow M \circ - & & \downarrow M' \circ - & & \downarrow M'' \circ - \\ \mathcal{C}_\theta(\emptyset, P) & \xrightarrow{P^H \circ -} & \mathcal{C}_\theta(\emptyset, P') & \xrightarrow{P'^H \circ -} & \mathcal{C}_\theta(\emptyset, P'') \cdots \end{array}$$

used to form the map of homotopy colimits in the statement of Theorem 1.3. Now Theorem 2.15 implies that, after passing to the subdiagram of (2.3) consisting of those components represented by cobordisms (s, W) such that $\ell_W : W \rightarrow B$ is n -connected, the map on homotopy colimits is an abelian homology equivalence.

The corresponding subdiagram of (2.4) consists of the maps $(M, \hat{\ell}_M^{(g)}) \circ - : \mathcal{N}_n^\theta(Q, \hat{\ell}_Q^{(g)}) \rightarrow \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(g)})$, and as a diagram in the homotopy category it is isomorphic to the diagram made out of the squares (1.3). In principle the homotopies implied in the diagram (2.4) could be different from the ones described informally in the introduction and since different choices of homotopies can lead to non-homotopic maps of homotopy colimits, we may not immediately conclude that the stabilised maps in Theorems 1.3 and 2.15 are isomorphic as arrows in the homotopy category. However, as we shall explain in Lemma A.10, if the induced map of homotopy colimits is an abelian homology equivalence for one choice of homotopies then it is so for all choices. \square

2.5. Elementary simplifications of Theorem 2.15. There are several easy simplifications which can be made before embarking on the proof of Theorem 2.15. In order to phrase these simplifications, it is convenient to introduce the following notation: let $\mathcal{W} \subset \mathcal{D}$ be the subcategory of those morphisms $M : P \rightsquigarrow Q$ in \mathcal{D} such that the induced map $\mathcal{F}(Q) \rightarrow \mathcal{F}(P)$ is an abelian homology equivalence. Theorem 2.15 is then equivalent to the assertion that $\mathcal{W} = \mathcal{D}$, but this has the advantage that intermediate results can be stated: we shall show that \mathcal{W} contains larger and larger classes of morphisms, until it is clear that it coincides with the entire category \mathcal{D} .

The first simplification is a purely formal saturation property of abelian homology equivalences.

Lemma 2.19. *The subcategory $\mathcal{W} \subset \mathcal{D}$ has the 2-out-of-3 property.*

Proof. This will follow from the fact that for maps of spaces $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, if any two of f , g , and $g \circ f$ are abelian homology equivalences, so is the third.

Ordinary homology equivalences satisfy 2-out-of-3, so if any two of f , g , and $g \circ f$ are abelian homology equivalences then all three spaces have the same first homology, and so have the same collection of abelian local coefficient systems available. In particular, we may assume that all abelian local coefficient systems are pulled back from Z . The claim then follows by functoriality of homology with local coefficients. \square

Lemma 2.20. *If \mathcal{W} contains every morphism $M : P \rightsquigarrow Q$ in \mathcal{D} whose underlying smooth cobordism has a handle structure relative to Q consisting of a single handle of index k where $n \leq k \leq 2n$, then $\mathcal{W} = \mathcal{D}$.*

Proof. If $2n \neq 4$ then this is easy: a cobordism $M : P \rightsquigarrow Q$ in \mathcal{D} is $(n-1)$ -connected relative to Q by definition, so admits a handle structure relative to Q which only has handles of index at least n . For $2n = 2$ this is clear, by cancelling any handles of index 0, and for $2n \geq 6$ this may be seen by handle-trading, as in the proof of the s -cobordism theorem (see e.g. [Ker65]). Cutting M into the corresponding elementary cobordisms, each of which lie in \mathcal{W} by assumption, it follows that M does too.

If $2n = 4$ such handle-trading is not immediately available, but becomes available after connect-sum with many copies of $S^2 \times S^2$. More precisely, let us write $H_0 : P_{-1} \rightsquigarrow P_0 = P$ for any morphism with underlying smooth bordism $([-1, 0] \times P) \sharp (S^n \times S^n)$ and θ structure extending $\hat{\ell}_P$ on $\{0\} \times P$, and repeat to get composable morphisms $H_{-k} : P_{-k-1} \rightsquigarrow P_{-k}$ with underlying smooth bordism $([-k-1, -k] \times P) \sharp (S^n \times S^n)$ for all $k \geq 0$. Then it follows from [Qui83, Theorem 1.2] that for sufficiently large k , the cobordism

$$M' : P_{-k} \xrightarrow{H_{-(k-1)}} \dots \xrightarrow{H_{-2}} P_{-2} \xrightarrow{H_{-1}} P_{-1} \xrightarrow{H_0} P_0 = P \xrightarrow{M} Q$$

admits a handle structure relative to Q which only has relative handles of index at least 2, and so by assumption lies in \mathcal{W} . The cobordism $H_0 \circ \dots \circ H_{-(k-1)}$ also admits a handle structure relative to P which only has relative handles of index at least 2, so also lies in \mathcal{W} ; by Lemma 2.19, it follows that M also lies in \mathcal{W} . \square

The manifold ∂L has finitely many path-components, and each $K|_{[i, i+1]}$ contains an embedded $W_{1,1}$ with standard θ -structure in a path component intersecting ∂L . Therefore there is a path component $\partial_0 L \subset \partial L$ such that infinitely-many $K|_{[i, i+1]}$ contain an embedded $W_{1,1}$ with standard θ -structure in the path component of $\partial_0 L$. We call the path-component of $\partial_0 L$ the *basepoint component*. By composing some of the cobordisms $K|_{[i, i+1]}$, and rescaling, we may assume that each $K|_{[i, i+1]}$ contains an embedded $W_{1,1}$ with standard θ -structure in the basepoint component.

For a cobordism $M : P \rightsquigarrow Q$ which has a handle structure relative to Q consisting of a single handle of index $k \geq n$, it makes sense to ask whether the handle is attached to the basepoint component i.e. whether the attaching map has image which intersects the basepoint component. When $k > 1$ the image of such an attaching map is always connected, so is required to lie inside the basepoint component; when $k = n = 1$ we require that at least one of the two components lie inside the basepoint component.

Lemma 2.21. *If \mathcal{W} contains every morphism $M : P \rightsquigarrow Q$ in \mathcal{D} whose underlying smooth cobordism has a handle structure relative to Q consisting of a single handle of index $n \leq k < 2n$ attached to the basepoint component, then $\mathcal{W} = \mathcal{D}$.*

Proof. Let $M : P \rightsquigarrow Q \in \mathcal{D}$ be a morphism having a handle structure relative to Q consisting of a single handle of index $n \leq k \leq 2n$. If we can show that $M \in \mathcal{W}$ then the previous lemma applies and gives the desired conclusion.

First suppose that $k = 2n$: then M is obtained from Q by attaching a single $2n$ -handle to Q , along an entire component $\phi : S^{2n-1} \hookrightarrow Q$. Choose surgery data $\varphi : S^0 \times D^{2n-1} \hookrightarrow Q$ sending one disc to the basepoint component and the other into the image of ϕ , and let $U : Q \rightsquigarrow R$ be the trace of the surgery along φ . Then the composition $U \circ M$ consists of a $2n$ -handle and a $(2n - 1)$ -handle relative to R , and moreover these are cancelling handles by construction: thus $U \circ M$ is a cylinder and so is invertible in \mathcal{D} and hence lies in \mathcal{W} . On the other hand U consists of a single $(2n - 1)$ -handle relative to R attached to the basepoint component, so $U \in \mathcal{W}$ by hypothesis: thus $M \in \mathcal{W}$ by Lemma 2.19.

Now suppose that $n \leq k < 2n$. We will construct morphisms

$$\begin{array}{ccc} P & \xrightarrow{M} & Q \\ \downarrow T' & & \downarrow T \\ R & \xrightarrow{M'} & S \end{array}$$

in \mathcal{D} such that $T \circ M$ and $M' \circ T'$ lie in the same path component of $\mathcal{D}(P, S)$, and $T, M', T' \in \mathcal{W}$. It will then follow from Lemma 2.19 that $M \in \mathcal{W}$, as required.

Suppose that M consists of a single handle attached via a map $\phi : \partial D^k \times D^{2n-k} \hookrightarrow Q$, which does not land in the basepoint component. Choose surgery data $\varphi : S^0 \times D^{2n-1} \hookrightarrow Q$, disjoint from L and from ϕ , such that surgery along it connects one of the path components intersecting the image of ϕ to the basepoint component, and let $T : Q \rightsquigarrow S$ be the trace of the surgery along φ . Considered relative to S , it consists of a single $(2n - 1)$ -handle attached to the basepoint component of S , so $T \in \mathcal{W}$.

The cobordism $T \circ M : P \rightsquigarrow S$ has a k -handle and a $(2n - 1)$ -handle relative to S , and these are attached along disjoint embeddings. Thus it is isotopic to a composition $M' \circ T' : P \rightsquigarrow S$ where $M' : R \rightsquigarrow S$ has a k -handle relative to S and $T' : P \rightsquigarrow R$ has a $(2n - 1)$ -handle relative to R . The k -handle of M' is attached along ϕ , but as S is the result of surgery on Q along φ the map ϕ now lands in the basepoint component, and so $M' \in \mathcal{W}$. The $(2n - 1)$ -handle of T' is attached to the basepoint component of R , because by construction it consists of a 1-handle relative to P with one end attached in the basepoint component of P , so $T' \in \mathcal{W}$. This provides the required data for the argument given above. \square

Thus to prove Theorem 2.15 it suffices to show that every cobordism $M : P \rightsquigarrow Q$ in \mathcal{D} which has a handle structure relative to Q consisting of a single handle of index $n \leq k < 2n$ attached to the basepoint component induces an abelian homology equivalence. Our proof occupies the following four sections, and proceeds by induction on k .

3. PROOF OF THEOREM 2.15: STABILITY FOR $W_{1,1}$

For any $P \in \mathcal{D}$, let us write

$$H_P : 'P \rightsquigarrow P$$

for any morphism in \mathcal{D} obtained from $[0, 1] \times P$ by forming the boundary connect-sum with $W_{1,1}$ at $\{0\} \times P$ inside the basepoint component of P (that is, the path component of $\partial_0 L$). Similarly, let

$${}_P H : P \rightsquigarrow P'$$

be any morphism in \mathcal{D} obtained from $[0, 1] \times P$ by forming the boundary connect-sum with $W_{1,1}$ at $\{1\} \times P$ inside the basepoint component of P . Both ${}_P H$ and H_P are equipped with any θ -structure which is standard when restricted to the embedded $W_{1,1}$, and is equal to $\hat{\ell}_P$ when restricted to P . We shall establish the following case of Theorem 2.15 for any choice of such morphisms H_P and ${}_P H$.

Theorem 3.1. *For each $P \in \mathcal{D}$, the morphism H_P lies in \mathcal{W} .*

Proof. Recall that the functor \mathcal{F} is defined as the objectwise homotopy colimit of the \mathcal{F}_i , so the map $- \circ H_P : \mathcal{F}(P) \rightarrow \mathcal{F}(P')$ which we must show is an abelian homology equivalence is the induced map on horizontal homotopy colimits of the commutative diagram

$$\begin{array}{ccccccc} \mathcal{F}_0(P) & \xrightarrow{K|_{[0,1]^\circ-}} & \mathcal{F}_1(P) & \xrightarrow{K|_{[1,2]^\circ-}} & \mathcal{F}_2(P) & \xrightarrow{K|_{[2,3]^\circ-}} & \cdots \\ \downarrow -\circ H_P & & \downarrow -\circ H_P & & \downarrow -\circ H_P & & \\ \mathcal{F}_0(P') & \xrightarrow{K|_{[0,1]^\circ-}} & \mathcal{F}_1(P') & \xrightarrow{K|_{[1,2]^\circ-}} & \mathcal{F}_2(P') & \xrightarrow{K|_{[2,3]^\circ-}} & \cdots \end{array}$$

We shall show that for each square in this diagram there are maps

$$\Delta_i^{\text{top}}, \Delta_i^{\text{bottom}} : \mathcal{F}_i(P') \longrightarrow \mathcal{F}_{i+1}(P)$$

which respectively make the top and bottom triangles of the squares in which they lie commute up to homotopy; it will then follow from Proposition A.13 that the induced maps on horizontal homotopy colimits are abelian homology equivalences. In order to do this, we will pass to a homotopy equivalent model of this diagram. For the reader familiar with [GRW14b], this is completely analogous to Lemma 7.15 of that paper.

Lemma 2.17 shows that $\mathcal{F}_i(P) \simeq \mathcal{N}_n^\theta(\langle P, K|_i \rangle) \subset \mathcal{C}_\theta(\emptyset, \langle P, K|_i \rangle)$, the subspace consisting of those nullbordisms $(s, W) : \emptyset \rightsquigarrow \langle P, K|_i \rangle$ for which $\ell_W : W \rightarrow B$ is n -connected. Similarly for P' and the morphism $K|_{[i, i+1]}$. Hence the square in which we are trying to find diagonal maps (up to homotopy) may be replaced by

$$\begin{array}{ccc} \mathcal{N}_n^\theta(\langle P, K|_i \rangle) & \xrightarrow{\langle P, K|_{[i, i+1]}^\circ-} & \mathcal{N}_n^\theta(\langle P, K|_{i+1} \rangle) \\ \langle H_P, K|_i \rangle^\circ- \downarrow & & \downarrow \langle H_P, K|_{i+1} \rangle^\circ- \\ \mathcal{N}_n^\theta(\langle P', K|_i \rangle) & \xrightarrow{\langle P', K|_{[i, i+1]}^\circ-} & \mathcal{N}_n^\theta(\langle P', K|_{i+1} \rangle). \end{array}$$

We claim that $X = \langle H_P, K|_i \rangle$ may be θ -embedded into $Y = \langle P, K|_{[i, i+1]} \rangle$ relative to $\langle P, K|_i \rangle$. If this is the case, the complement of such an embedding gives a cobordism $Z : \langle P', K|_i \rangle \rightsquigarrow \langle P, K|_{i+1} \rangle$ which is $(n-1)$ -connected relative to $\langle P, K|_{i+1} \rangle$, and so $Z \circ -$ defines a diagonal map in the square making the top triangle commute up to homotopy, as required.

The cobordism $X : \langle P, K|_i \rangle \rightsquigarrow \langle P', K|_i \rangle$ is—by definition of H_P —obtained from $\langle P, K|_i \rangle$ by forming the boundary connect-sum with $W_{1,1}$ with a standard θ -structure inside the basepoint component of $\{1\} \times \langle P, K|_i \rangle$. On the other hand,

the cobordism $Y : \langle P, K|_i \rangle \rightsquigarrow \langle P, K|_{i+1} \rangle$ contains an embedded $W_{1,1}$ with standard θ -structure, in the basepoint component (as we re-indexed the sequence of cobordisms $K|_{[i,i+1]}$ in order to have this property). There is therefore an embedding e of underlying manifolds from X to Y sending the $W_{1,1}$ in X to that in Y . By Lemma 2.12 the θ -structure $e^*\hat{\ell}_Y$ is homotopic to $\hat{\ell}_X$ relative to $\langle P, K|_i \rangle$, and extending this homotopy to $\hat{\ell}_Y$ it follows that e is an embedding of θ -manifolds.

Similarly, $\langle H_P, [0, 1] \times K|_{i+1} \rangle$ may be θ -embedded into $\langle [0, 1] \times P', K|_{[i,i+1]} \rangle$ relative to $\langle P', K|_{i+1} \rangle$, which provides a diagonal map making the bottom triangle commute up to homotopy. \square

4. PROOF OF THEOREM 2.15: STABILITY FOR HANDLES OF INDEX n

In this section and the next we shall prove Theorem 2.15 in the case of a cobordism which admits a handle structure relative to its incoming boundary consisting of a single n -handle attached to the basepoint component. In terms of the subcategory $\mathcal{W} \subset \mathcal{D}$ from Section 2.5, the precise statement we prove is the following.

Theorem 4.1. *If $M : P \rightsquigarrow Q$ is a morphism in \mathcal{D} whose underlying smooth cobordism admits a handle structure relative to Q consisting of a single n -handle attached to the basepoint component of Q , then $M \in \mathcal{W}$.*

The detailed proof is cumbersome, but the strategy can be explained informally as follows. It suffices to prove that the composition of M and the morphism $r(M)$, which attaches the n -handle “dual” to that of M , is in \mathcal{W} , since this then applies to $r(M)$, and we may deduce that all three morphisms in the composition $(M \circ r(M)) \circ r(r(M)) = M \circ (r(M) \circ r(r(M)))$ induce isomorphisms on $H_1(-; \mathbb{Z})$ and homology with abelian coefficients. If the handle of M is attached trivially, the composition of $M \circ r(M)$ is essentially a model for the cobordism H_P considered in the previous section. We shall then use a simplicial resolution to reduce the general case to the trivially attached case.

4.1. Support and interchange of support. We first introduce the following notion of support of a cobordism, which will be used not only in this section but also in Section 6.

Definition 4.2. The *support* $\text{supp}(W)$ of a cobordism $(s, W) : P \rightsquigarrow Q$ in \mathcal{C}_θ is the smallest closed set $A \subset \mathbb{R}^\infty$ such that

$$W \cap ([0, s] \times (\mathbb{R}^\infty \setminus A)) = [0, s] \times (P \setminus A)$$

as θ -manifolds. Or, equivalently, such that

$$W \cap ([0, s] \times (\mathbb{R}^\infty \setminus A)) = [0, s] \times (Q \setminus A)$$

as θ -manifolds.

This notion of support is specific to our particular model of cobordism categories using manifolds embedded in euclidean space; it does not have meaning for “abstract” cobordisms. We used a similar notion in [GRW14b, p. 268].

Definition 4.3. If $(s, W) : P \rightsquigarrow Q$ and $(s', W') : Q \rightsquigarrow R$ are cobordisms in \mathcal{C}_θ such that $\text{supp}(W) \cap \text{supp}(W') = \emptyset$, then we let the *interchange of support* be the θ -cobordisms

$$\mathcal{R}_{W'}(W) = (W \setminus ([0, s] \times \text{supp}(W'))) \cup ([0, s] \times (R \setminus \text{supp}(W)))$$

and

$$\mathcal{L}_W(W') = (W' \setminus ([0, s'] \times \text{supp}(W))) \cup ([0, s'] \times (P \setminus \text{supp}(W'))).$$

The composition $(s, \mathcal{R}_{W'}(W)) \circ (s', \mathcal{L}_W(W'))$ may be formed in \mathcal{C}_θ , and is a morphism from P to R . Furthermore, there is a path

$$t \mapsto \tau_t(W, W') : [0, 1] \longrightarrow \mathcal{C}_\theta(P, R)$$

from $(s', W') \circ (s, W)$ to $(s, \mathcal{R}_{W'}(W)) \circ (s', \mathcal{L}_W(W'))$ given by sliding $W \cap ([0, s] \times \text{supp}(W))$ in the positive direction and $W' \cap ([0, s'] \times \text{supp}(W'))$ in the negative direction.

This manoeuvre may be described graphically as in Figure 1.




FIGURE 1. The interchange of support path. The grey regions indicate where the cobordisms are not cylindrical.

4.2. Constructing auxiliary cobordisms. We begin by showing that the morphism $M : P \rightsquigarrow Q$ for which we shall prove Theorem 4.1 may be assumed to be of a rather standard form. As we shall need the analogous result later, when dealing with cobordisms having a single k -handle for $n \leq k \leq 2n$, we begin by working in this generality.

Construction 4.4. The (reverse) trace of the surgery along $\partial D^k \times D^{2n-k} \subset \partial D^k \times \mathbb{R}^{2n-k}$ gives a smooth manifold $T \subset [0, 1] \times \mathbb{R}^k \times \mathbb{R}^{2n-k}$ such that

- (i) T agrees with $[0, 1] \times \partial D^k \times \mathbb{R}^{2n-k}$ outside of $[0, 1] \times D^k \times D^{2n-k}$,
- (ii) $T \cap ((1 - \varepsilon, 1] \times \mathbb{R}^{2n}) = (1 - \varepsilon, 1] \times \partial D^k \times \mathbb{R}^{2n-k}$ for some $\varepsilon > 0$, and
- (iii) $T \cap ([0, \varepsilon] \times \mathbb{R}^{2n}) = [0, \varepsilon] \times P$ for some $\varepsilon > 0$, with P diffeomorphic to $\mathbb{R}^k \times \partial D^{2n-k}$,

and comes equipped with embeddings of the k -handle and its dual $(2n - k)$ -handle which we denote

$$\begin{aligned} \phi_T : D^k \times \mathbb{R}^{2n-k} &\hookrightarrow T \\ \varphi_T : \mathbb{R}^k \times D^{2n-k} &\hookrightarrow T. \end{aligned}$$

We have $\phi_T(\partial D^k \times \mathbb{R}^{2n-k}) \subset T \cap (\{1\} \times \mathbb{R}^{2n})$ and $\varphi_T(\mathbb{R}^k \times \partial D^{2n-k}) \subset T \cap (\{0\} \times \mathbb{R}^{2n})$, and we may also arrange that $\phi_T(x, y) = (1, x, y) \in T \cap (\{1\} \times \mathbb{R}^{2n})$ for all $(x, y) \in \partial D^k \times \mathbb{R}^{2n-k}$.

Suppose given a $(2n - 1)$ -manifold $Q \subset \mathbb{R}^\infty$ and an embedding $\sigma : \mathbb{R}^k \times \mathbb{R}^{2n-k} \rightarrow \mathbb{R}^\infty$ such that $\sigma^{-1}(Q) = \partial D^k \times \mathbb{R}^{2n-k}$. Then $\sigma|_{\partial D^k \times \mathbb{R}^{2n-k}}$ is a diffeomorphism from $\partial D^k \times \mathbb{R}^{2n-k}$ onto an open subset of Q and we may define an elementary cobordism $M_\sigma \subset [0, 1] \times \mathbb{R}^\infty$ by

$$M_\sigma = ([0, 1] \times (Q \setminus \sigma(\partial D^k \times D^{2n-k}))) \cup ((\text{Id}_{[0,1]} \times \sigma)(T)).$$

We then have

$$\partial M_\sigma = (\{0\} \times P_\sigma) \cup (\{1\} \times Q)$$

for a closed $(2n - 1)$ -manifold $P_\sigma \subset \mathbb{R}^\infty$ obtained from Q by surgery along $\sigma|_{\partial D^k \times D^k}$. (Even though the manifold P_σ depends on Q and σ and not P we shall use the notation P_σ to emphasise that it is a model for P in Lemma 4.5 below.) The

embedding ϕ_T of a handle in T then induces an embedding of a handle into M , for which we shall write

$$\phi = (\text{Id}_{[0,1]} \times \sigma) \circ \phi_T : D^k \times \mathbb{R}^{2n-k} \rightarrow M_\sigma.$$

This embedding satisfies $\phi(\partial D^k \times \mathbb{R}^{2n-k}) \subset \{1\} \times Q$ and in fact $\phi(x, y) = (1, \sigma(x, y))$ for $(x, y) \in \partial D^k \times \mathbb{R}^{2n-k}$.

If Q is endowed with a θ -structure $\hat{\ell}_Q$, then pulling back along σ endows $\partial D^k \times \mathbb{R}^{2n-k}$ with a θ -structure, and hence the subspace

$$T_0 = (\{1\} \times \partial D^k \times \mathbb{R}^{2n-k}) \cup ([0, 1] \times \partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k})) \subset T$$

inherits a θ -structure. The maps of pairs

$$(D^k, \partial D^k) \times D^{2n-k} \xrightarrow{\phi_T|_{D^k}} (T, \{1\} \times \partial D^k \times \mathbb{R}^{2n-k}) \xrightarrow{\text{incl}} (T, T_0)$$

are relative homotopy equivalences, and hence to extend the θ -structure from T_0 to T is the same, up to homotopy, as extending the bundle map

$$(4.1) \quad T(\partial D^k \times D^{2n-k}) \oplus \varepsilon^1 \xrightarrow{D\sigma \oplus \varepsilon^1} TQ \oplus \varepsilon^1 \xrightarrow{\hat{\ell}_Q} \theta^* \gamma_{2n}$$

to $T(D^k \times D^{2n-k})$. In what follows we shall, whenever given an extension of (4.1), tacitly pick an extension of θ -structures from T_0 to T in the corresponding homotopy class. Given such an extension to T , we obtain a θ -structure on M_σ , and hence by restriction a θ -structure on P_σ . The θ -cobordism $(1, M_\sigma)$ so obtained has support in $\sigma(D^k \times D^{2n-k})$.

Lemma 4.5. *Let $M : P \rightsquigarrow Q$ be a morphism in \mathcal{D} whose underlying smooth cobordism admits a handle structure relative to Q consisting of a single k -handle. Then*

- (i) *there exists an embedding $\sigma : \mathbb{R}^k \times \mathbb{R}^{2n-k} \rightarrow [0, \infty) \times \mathbb{R}^{\infty-1}$ such that $\sigma^{-1}(Q) = \partial D^k \times \mathbb{R}^{2n-k}$, and an extension of $(\sigma|_{\partial D^k \times D^{2n-k}})^* \hat{\ell}_Q$ to $T(D^k \times D^{2n-k})$, so that we may form the θ -cobordism*

$$M_\sigma : P_\sigma \rightsquigarrow Q \in \mathcal{D},$$

- (ii) *for a suitable choice of embedding and extension in (i), there exists an isomorphism $U : P \rightsquigarrow P_\sigma$ in \mathcal{D} such that $M_\sigma \circ U$ and M lie in the same component of $\mathcal{D}(P, Q)$.*

Proof. Suppose for simplicity that the cobordism M has length 2. The handle structure of M gives an embedding of smooth manifolds $M_\sigma \hookrightarrow M$ relative to Q , whose complement is diffeomorphic to $[0, 1] \times P$ relative to $\{0\} \times P$, and we give M_σ the θ -structure from M . After changing the embedding of $M \subset [0, 2] \times \mathbb{R}^\infty$ relative to $(\{0\} \times P) \cup (\{2\} \times Q)$ we may assume that it is *equal* to the composition of M_σ and an isomorphism U . \square

We now return to the particular case where $k = n$. If Theorem 4.1 holds for the cobordism M_σ produced by this lemma, then it also holds for M , since U is an isomorphism. Hence without loss of generality we may assume that M is of the form M_σ with respect to some embedding $\sigma : \mathbb{R}^n \times \mathbb{R}^n \rightarrow [0, \infty) \times \mathbb{R}^{\infty-1}$ with $\sigma^{-1}(Q) = \partial D^n \times \mathbb{R}^n$ and some extension of $(\sigma|_{\partial D^n \times D^n})^* \hat{\ell}_Q$ to $T(D^n \times D^n)$. We shall henceforth make this assumption, and then also $P_\sigma = P$. The proof of Theorem 4.1 shall use certain compositions

$$R \xrightarrow{r(M)} P \xrightarrow{M} Q \xrightarrow{V_p} S_p$$

of M and cobordisms $r(M)$ and V_p for $p \geq 0$, which we first define. We shall make use of the reflection diffeomorphism $r(t, x) = (1 - t, x) : [0, 1] \times \mathbb{R}^\infty \rightarrow [0, 1] \times \mathbb{R}^\infty$.

Construction 4.6. Construct a cobordism $r(M) : R \rightsquigarrow P$ having underlying manifold the reflection $r(M)$ of $M = M_\sigma$, so the underlying manifold of R is the same as that of Q . It remains to describe a θ -structure on $r(M)$. The manifold M has a single n -handle relative to P , namely

$$\varphi = (\text{Id}_{[0,1]} \times \sigma) \circ \varphi_T|_{\mathbb{R}^n \times D^n} : \mathbb{R}^n \times D^n \longrightarrow M$$

which has $\varphi(\mathbb{R}^n \times \partial D^n) \subset \{0\} \times P$. The reflection of φ gives an n -handle $r(\varphi)$ in $(r(M), P)$, and any θ -structure on this handle extending $(r(\varphi)|_{\partial D^n \times D^n})^* \hat{\ell}_P$ induces a θ -structure on $r(M)$ supported in $\sigma(D^n \times D^n)$. We choose the θ -structure on this handle by insisting that the induced θ -structure on

$$S^n \times D^n \cong r(\varphi)(D^n \times D^n) \cup_P \varphi(D^n \times D^n) \subset r(M) \cup_P M$$

is diffeomorphic to the standard one described in Section 2.2.

Construction 4.7. We construct a cobordism $V_0 : Q \rightsquigarrow S_0$ having support in $\sigma(D^n \times (3e_1 + D^n))$ by letting

$$V_0 = ([0, 1] \times (Q \setminus \sigma(\partial D^n \times (3e_1 + D^n)))) \cup ((\text{Id}_{[0,1]} \times \sigma)(3e_1 + r(T))),$$

which contains the handle

$$\psi = (\text{Id}_{[0,1]} \times \sigma) \circ (3e_1 + r(\phi_T|_{D^n \times D^n})) : D^n \times D^n \longrightarrow V_0$$

which has $\psi(\partial D^n \times D^n) \subset \{0\} \times Q$. We fix the θ -structure on V_0 by insisting that the θ -structure on

$$S^n \times D^n \cong \psi(D^n \times (3e_1 + D^n)) \cup_Q \psi(D^n \times D^n) \subset M \cup_Q V_0$$

is the standard one described in Section 2.2.

More generally, for each $p \geq 0$ we construct a cobordism $V_p : Q \rightsquigarrow S_p$ having support in $\cup_{i=0}^p \sigma(D^n \times (3(i+1)e_1 + D^n))$ by performing the above construction simultaneously inside each of $\sigma(D^n \times (3(i+1)e_1 + D^n))$ for $i = 0, 1, \dots, p$. We write $\psi_i : D^n \times D^n \rightarrow V_p$ for the n -handles relative to Q , $i = 0, \dots, p$.

The composable cobordisms $V_p : Q \rightsquigarrow S_p$ and $M \circ r(M) : R \rightsquigarrow Q$ arising from these constructions have disjoint support, so may be subjected to interchange of support. Recall from Section 3 that we write $H_P : P' \rightsquigarrow P$ for any morphism in \mathcal{D} obtained from $[0, 1] \times P$ by forming the boundary connect-sum with $W_{1,1}$ at $\{0\} \times P$.

Lemma 4.8. *For each $p \geq 0$, the cobordism $\mathcal{R}_{V_p}(M \circ r(M))$ obtained by interchange of support is H_{S_p} composed with an isomorphism.*

The cobordism $M \circ r(M)$ is obtained from Q by attaching an n -handle and its dual, and as the support of $M \circ r(M)$ is disjoint from that of V_p the cobordism $\mathcal{R}_{V_p}(M \circ r(M))$ is also obtained by attaching an n -handle and its dual. The proof of this lemma consists of showing that in this case the n -handle is trivially attached.

Proof. In order that a cobordism $W : S \rightsquigarrow T$ be isotopic to H_T composed with an isomorphism, it is enough that it contain an embedded copy of $W_{1,1}$ on which the induced θ -structure is the standard one described in Section 2.2, and so that the smooth cobordism (without θ -structure) $(W \setminus \text{int}(W_{1,1})) \cup_{\partial W_{1,1}} D^{2n} : S \rightsquigarrow T$ is diffeomorphic to $[0, 1] \times T$ relative to T .

By construction of the θ -structure on $r(M)$,

$$S^n \times D^n \cong r(\varphi)(D^n \times D^n) \cup \varphi(D^n \times D^n) \subset r(M) \cup_P M$$

has a standard θ -structure. This is disjoint from $\text{supp}(V_p)$ so can also be found inside the cobordism $\mathcal{R}_{V_p}(M \circ r(M))$.

The n -handle ϕ of M relative to Q has support disjoint from $\text{supp}(V_p)$, so can also be considered as an n -handle of $\mathcal{R}_{V_p}(M \circ r(M))$ relative to S_p . The manifold S_p is, tautologically, the result of doing surgery on the embeddings

$$\sigma|_{\partial D^n \times (3(i+1)e_1 + D^n)} : \partial D^n \times D^n \hookrightarrow Q \quad i = 0, 1, \dots, p,$$

and so, using only that surgery on $\sigma|_{\partial D^n \times (3e_1 + D^n)}$ has been performed, the embedding $\sigma|_{\partial D^n \times D^n} : \partial D^n \times D^n \hookrightarrow Q \setminus \text{supp}(V_p) \subset S_p$ can be isotoped into a collar neighbourhood of $S_p \subset \mathcal{R}_{V_p}(M \circ r(M))$ and extended there to an embedding of $D^n \times D^n$ such that the result of joining it with the n -handle ϕ of $\mathcal{R}_{V_p}(M \circ r(M))$ gives an embedded $S^n \times D^n$. Furthermore, by our choice of θ -structure on V_p this may be arranged to have standard θ -structure. By construction, the core of this $S^n \times D^n$ intersects that of the previous paragraph transversely in a single point, so their union gives an embedded $W_{1,1}$, on which the θ -structure is standard. \square

4.3. A semi-simplicial resolution. In this section we explain how to construct certain semi-simplicial spaces $\mathcal{Y}_j(P)_\bullet$ augmented over the spaces $\mathcal{F}_j(P)$ of cobordisms introduced in Definition 2.13, when we have fixed a θ -end

$$K|_0 \xrightarrow{K|_{[0,1]}} K|_1 \xrightarrow{K|_{[1,2]}} K|_2 \xrightarrow{K|_{[2,3]}} K|_3 \longrightarrow \dots$$

and P is an object of the category \mathcal{D} equipped with a little extra structure. For clarity we work in slightly more generality than we will eventually need.

Definition 4.9. Fix an object $P \in \mathcal{D}$, an element $W = (s, W) \in \mathcal{F}_j(P)$ for some $j \geq 0$, an embedding $\chi : \partial D^n \times (1, \infty) \times \mathbb{R}^{n-1} \hookrightarrow P$, and a 1-parameter family $t \mapsto \hat{\ell}_t^{\text{std}}$, $t \in (2, \infty)$, of θ -structures on $D^n \times D^n$ such that $\hat{\ell}_t^{\text{std}}|_{\partial D^n \times D^n} = \chi^* \hat{\ell}_P|_{\partial D^n \times (t \cdot e_1 + D^n)}$.

Let $Y(W)_0 = Y(W, \chi, \hat{\ell}_t^{\text{std}})_0$ be the set of tuples (t, c, \hat{L}) consisting of a $t \in (2, \infty)$, an embedding $c : (D^n \times D^n, \partial D^n \times D^n) \hookrightarrow (W, P)$, and a path of θ -structures $\tau \mapsto \hat{L}(\tau)$ on $D^n \times D^n$, $\tau \in [0, 1]$, such that

- (i) there is a $\delta > 0$ such that $c(x, v) = \chi(\frac{x}{|x|}, v + t \cdot e_1) + (1 - |x|) \cdot e_0$ for $1 - |x| < \delta$,
- (ii) the image $C = c(D^n \times D^n)$ is disjoint from $([0, s] \times L) \cup (\{s\} \times K|_j)$, and $c^{-1}(P) = \partial D^n \times D^n$,
- (iii) the restriction $\ell_W|_{W \setminus C} : W \setminus C \rightarrow B$ is n -connected,
- (iv) $\hat{L}(0) = c^* \hat{\ell}_W$, $\hat{L}(1) = \hat{\ell}_t^{\text{std}}$, and $\hat{L}(\tau)|_{\partial D^n \times D^n}$ is independent of $\tau \in [0, 1]$.

We topologise $Y(W)_0$ as a subspace of

$$\mathbb{R} \times \text{Emb}(D^n \times D^n, [0, \infty) \times \mathbb{R}^\infty) \times \text{Bun}^\theta(D^n \times D^n)^I.$$

Let $Y(W)_p = Y(W, \chi, \hat{\ell}_t^{\text{std}})_p \subset (Y(W)_0)^{p+1}$ be the subset consisting of tuples $(t_0, c_0, \hat{L}_0, t_1, c_1, \hat{L}_1, \dots, t_p, c_p, \hat{L}_p)$ such that

- (i) the images C_i of the c_i are disjoint,
- (ii) $t_0 < t_1 < \dots < t_p$,
- (iii) the restriction $\ell_W|_{W \setminus (\cup_i C_i)} : W \setminus (\cup_i C_i) \rightarrow B$ is n -connected.

We topologise $Y(W)_p$ as a subspace of the $(p+1)$ -fold product of $Y(W)_0$. The collection $Y(W)_\bullet$ has the structure of a semi-simplicial space, where the i th face map forgets (t_i, c_i, \hat{L}_i) .

We now wish to combine all of the $Y(W)_\bullet$ for all $W = (s, W) \in \mathcal{F}_j(P)$ into a single semi-simplicial space.

Definition 4.10. Fix a $P \in \mathcal{D}$, an embedding $\chi : \partial D^n \times (1, \infty) \times \mathbb{R}^{n-1} \hookrightarrow P$, and a 1-parameter family $t \mapsto \hat{\ell}_t^{\text{std}}$, $t \in (2, \infty)$, of θ -structures on $D^n \times D^n$ such that $\hat{\ell}_t^{\text{std}}|_{\partial D^n \times D^n} = \chi^* \hat{\ell}_P|_{\partial D^n \times (t \cdot e_1 + D^n)}$.

Let $\mathcal{Y}_j(P)_p = \mathcal{Y}_j(P, \chi, \hat{\ell}_t^{\text{std}})_p$ be the set of tuples $(s, W; x)$ with $(s, W) \in \mathcal{F}_j(P)$ and $x \in Y(W, \chi, \hat{\ell}_t^{\text{std}})_p$. Topologise this set as a subspace of

$$\mathcal{F}_j(P) \times (\mathbb{R} \times \text{Emb}(D^n \times D^n, [0, \infty) \times \mathbb{R}^\infty) \times \text{Bun}^\theta(D^n \times D^n)^{p+1}).$$

The collection $\mathcal{Y}_j(P)_\bullet$ has the structure of a semi-simplicial space augmented over $\mathcal{F}_j(P)$, where the i th face maps forgets (t_i, c_i, \hat{L}_i) , and the augmentation map just remembers the underlying θ -cobordism (s, W) .

The maps $K|_{[j, j+1]} \circ - : \mathcal{F}_j(P) \rightarrow \mathcal{F}_{j+1}(P)$ given by the θ -end lift to semi-simplicial maps

$$\begin{aligned} (K|_{[j, j+1]} \circ -)_\bullet : \mathcal{Y}_j(P)_\bullet &\longrightarrow \mathcal{Y}_{j+1}(P)_\bullet \\ (s, W; x) &\longmapsto (s + 1; K|_{[j, j+1]} \circ W; x), \end{aligned}$$

and we let $\mathcal{Y}(P)_\bullet \rightarrow \mathcal{F}(P)$ be the augmented semi-simplicial space obtained as the levelwise homotopy colimit.

Lemma 4.11. *The map $|\mathcal{Y}_j(P)_\bullet| \rightarrow \mathcal{F}_j(P)$ is a quasifibration, with fibre $|Y(W)_\bullet|$ over $(s, W) \in \mathcal{F}_j(P)$.*

Proof. This may be proved in the same way as the analogue of Lemma 6.8 in [GRW14a, Section 7]. \square

The homotopy fibre of a map between mapping telescopes is weakly homotopy equivalent to the mapping telescope of the homotopy fibres. Hence the homotopy fibre of $|\mathcal{Y}(P)_\bullet| \rightarrow \mathcal{F}(P)$ over a point $(s, W) \in \mathcal{F}_j(P) \subset \mathcal{F}(P)$ is weakly homotopy equivalent to

$$(4.2) \quad \text{hocolim}_{g \rightarrow \infty} |Y(K|_{[j, j+g]} \circ W, \chi, \hat{\ell}_t^{\text{std}})_\bullet|.$$

In Section 5 we will prove the following theorem.

Theorem 4.12. *The space (4.2) is weakly contractible, and hence the forgetful map $|\mathcal{Y}(P)_\bullet| \rightarrow \mathcal{F}(P)$ is a weak equivalence.*

In the rest of this section we shall explain how to deduce Theorem 4.1 from Theorem 4.12.

4.4. Resolving composition with $M \circ r(M)$. We now continue with the notation of Section 4.2, in particular we have the cobordisms $M \circ r(M) : R \rightsquigarrow Q$ and $V_0 : Q \rightsquigarrow S_0$ and the embedding $\sigma : \partial D^n \times \mathbb{R}^n \hookrightarrow Q$, and we wish to resolve the map

$$- \circ M \circ r(M) : \mathcal{F}(Q) \longrightarrow \mathcal{F}(R).$$

We choose, once and for all, a 1-parameter family of embeddings

$$h_t : D^n \times D^n \hookrightarrow V_0, \quad t \in (2, \infty)$$

so that h_3 is the map ψ from Construction 4.7, and $h_t|_{\partial D^n \times D^n}(x, v) = \sigma(x, t \cdot e_1 + v)$. We then let $\hat{\ell}_t^{\text{std}} = h_t^* \hat{\ell}_{V_0}$.

The θ -structure on V_0 has been chosen (in Construction 4.7) so that the induced θ -structure on the sphere

$$S^n \times D^n \cong \phi(D^n \times (3e_1 + D^n)) \cup h_3(D^n \times D^n) \subset M \cup_Q V_0$$

is homotopic to the standard one; it follows that the same is true for the spheres $\phi(D^n \times (te_1 + D^n)) \cup h_t(D^n \times D^n)$ for all t .

The previous section then gives an augmented semi-simplicial space

$$(4.3) \quad \mathcal{Y}(Q)_\bullet = \mathcal{Y}(Q, \sigma|_{\partial D^n \times (1, \infty) \times \mathbb{R}^{n-1}}, \hat{\ell}_t^{\text{std}})_\bullet \longrightarrow \mathcal{F}(Q),$$

and Theorem 4.12 shows that this augmentation map becomes a weak homotopy equivalence after geometric realisation.

The underlying manifold of R is equal to the underlying manifold of Q , so the embedding σ can also be considered as an embedding $\sigma : \partial D^n \times \mathbb{R}^n \hookrightarrow R$. The θ -structures $\sigma^* \hat{\ell}_Q$ and $\sigma^* \hat{\ell}_R$ are not necessarily equal, but they become equal when restricted to $\partial D^n \times (1, \infty) \times \mathbb{R}^{n-1}$ (as $M \circ r(M)$ is supported in $\sigma(\mathbb{R}^n \times D^n)$). Hence we have defined an augmented semi-simplicial space

$$(4.4) \quad \mathcal{Y}(R)_\bullet = \mathcal{Y}(R, \sigma|_{\partial D^n \times (1, \infty) \times \mathbb{R}^{n-1}}, \hat{\ell}_t^{\text{std}})_\bullet \longrightarrow \mathcal{F}(R).$$

This also becomes a weak homotopy equivalence after geometric realisation, by Theorem 4.12.

We wish to cover the map $- \circ M \circ r(M) : \mathcal{F}(Q) \rightarrow \mathcal{F}(R)$ by a map of semi-simplicial resolutions $\mathcal{Y}(Q)_\bullet \rightarrow \mathcal{Y}(R)_\bullet$. The idea is that for any embedding $c : D^n \times D^n \hookrightarrow W$ of a thickened n -handle, as in Definition 4.9 but with Q replacing P , there is an induced thickened n -handle embedded in $W \circ M \circ r(M)$ —provided that the boundary of the handle is disjoint from the support of $M \circ r(M)$ —defined by gluing a cylinder $([0, 1] \times \partial D^n) \times D^n$ to the handle and embedding it into $M \circ r(M)$ in a cylindrical way. In order to define a map of semi-simplicial spaces in this way, we need to reparametrise the source

$$(([0, 1] \times \partial D^n) \times D^n) \cup_{\partial D^n \times D^n} (D^n \times D^n)$$

of the resulting embedding by a diffeomorphism to $\partial D^n \times D^n$. Of course we also need to explain what to do with the other data (θ -structures and so on) in the simplices of $Y(W)_\bullet$. Let us first explain the reparametrisation. We shall have to do the same later, so the following definition is given in greater generality than we need here: we will only use the case $k = n$ in this section.

Definition 4.13. Let $c : D^{2n-k} \times D^k \hookrightarrow [0, \infty) \times \mathbb{R}^\infty$ be an embedding which is collared in the sense that there is a $\delta > 0$ such that

$$c(x, v) = c\left(\frac{x}{|x|}, v\right) + (1 - |x|) \cdot e_0 \quad \text{if } 1 - |x| < \delta.$$

We may then define subsets

$$\begin{aligned} A_t &= [0, t] \times c(\partial D^{2n-k} \times D^k) \\ B_t &= c(D^{2n-k} \times D^k) + t \cdot e_0 \end{aligned}$$

whose union is a smooth submanifold of $[0, \infty) \times \mathbb{R}^\infty$. Define the *extrusion* of c of length $t \in [0, \infty)$, $\varepsilon_t(c) : D^{2n-k} \times D^k \hookrightarrow [0, \infty) \times \mathbb{R}^\infty$, by

$$(x, v) \mapsto \begin{cases} c((1+t) \cdot x, v) + t \cdot e_0 & |x| \leq \frac{1}{1+t} \\ c\left(\frac{x}{|x|}, v\right) + (1+t)(1 - |x|) \cdot e_0 & |x| \geq \frac{1}{1+t}. \end{cases}$$

In particular $\varepsilon_0(c) = c$. We have

$$\begin{aligned} A_t &= \varepsilon_t(c)\left(\left(D^{2n-k} \setminus \frac{1}{1+t} \text{int}(D^{2n-k})\right) \times D^k\right) \\ B_t &= \varepsilon_t(c)\left(\frac{1}{1+t} D^{2n-k} \times D^k\right), \end{aligned}$$

so that $A_t \cup B_t = \varepsilon_t(c)(D^{2n-k} \times D^k)$.

Let $\hat{\ell}$ be a θ -structure on $D^{2n-k} \times D^k$. The diffeomorphism

$$(x, v) \mapsto c(x, v) + t \cdot e_0 : D^{2n-k} \times D^k \longrightarrow B_t$$

endows B_t with a θ -structure, and we may extend this uniquely to a θ -structure on $A_t \cup B_t$ which is cylindrical over A_t . We define the *extrusion* of $\hat{\ell}$ of length $t \in [0, \infty)$, $\varepsilon_t(\hat{\ell})$, to be the θ -structure on $D^{2n-k} \times D^k$ obtained by pulling back the θ -structure just constructed along the diffeomorphism $\varepsilon_t(c) : D^{2n-k} \times D^k \rightarrow A_t \cup B_t$.

Composition with $M \circ r(M)$ gives a map $\mathcal{F}_j(Q) \rightarrow \mathcal{F}_j(R)$, and we may lift it to a semi-simplicial map $(-\circ M \circ r(M))_\bullet : \mathcal{Y}_j(Q)_\bullet \rightarrow \mathcal{Y}_j(R)_\bullet$ by the formula

$$(W; t, c, \hat{L}) \mapsto (W \circ M \circ r(M); t, \varepsilon_2(c), \varepsilon_2(\hat{L}))$$

on 0-simplices, and the analogous formula on the higher simplices. This construction commutes (strictly) with the semi-simplicial maps $(K|_{[j, j+1]} \circ -)_\bullet : \mathcal{Y}_j(Q)_\bullet \rightarrow \mathcal{Y}_{j+1}(Q)_\bullet$ and $(K|_{[j, j+1]} \circ -)_\bullet : \mathcal{Y}_j(R)_\bullet \rightarrow \mathcal{Y}_{j+1}(R)_\bullet$, and so induces a semi-simplicial map $(-\circ M \circ r(M))_\bullet : \mathcal{Y}(Q)_\bullet \rightarrow \mathcal{Y}(R)_\bullet$ on the stabilisations. The diagrams

$$\begin{array}{ccc} \mathcal{Y}(Q)_p & \xrightarrow{-\circ M \circ r(M)} & \mathcal{Y}(R)_p \\ \downarrow & & \downarrow \\ \mathcal{F}(Q) & \xrightarrow{-\circ M \circ r(M)} & \mathcal{F}(R) \end{array}$$

commute for all $p \geq 0$, and by Theorem 4.12 we may then view $\mathcal{Y}(Q)_\bullet \rightarrow \mathcal{Y}(R)_\bullet$ as a simplicial resolution of $\mathcal{F}(Q) \rightarrow \mathcal{F}(R)$.

Proposition 4.14. *For each $p \geq 0$, the map*

$$(-\circ M \circ r(M))_p : \mathcal{Y}(Q)_p \rightarrow \mathcal{Y}(R)_p$$

is an abelian homology equivalence.

Proof. The idea, cf. diagram (4.6) below, is to identify this map up to homotopy with a map gluing on a cobordism H_S as studied in Section 3.

Recall from Construction 4.7 that the cobordism $V_p : Q \rightsquigarrow S_p$ contains canonical handles

$$\psi_i : D^n \times D^n \rightarrow V_p \quad 0 \leq i \leq p,$$

and the θ -structure on V_p is determined (up to homotopy relative to Q) by its restriction to these handles. We may therefore suppose that the θ -structure on V_p is such that $\psi_i^* \hat{\ell}_{V_p} = \hat{\ell}_{3(i+1)}^{\text{std}} = h_{3(i+1)}^* \hat{\ell}_{V_0}$ for each i , because the θ -structure

$$h_i^* \hat{\ell}_{V_0} \cup (\phi|_{D^n \times (t e_1 + D^n)})^* \hat{\ell}_M : T(S^n \times D^n) \rightarrow \theta^* \gamma_{2n}$$

is standard for $t = 3$ (by definition of $\hat{\ell}_{V_0}$), and hence for all t (as being standard is homotopy invariant).

Let \hat{L}_i be the constant path $\hat{\ell}_{3(i+1)}^{\text{std}}$. This determines an injection

$$(4.5) \quad \begin{array}{l} \mathcal{F}_j(S_p) \rightarrow \mathcal{Y}_j(Q)_p \\ (s, X) \mapsto (s+1, X \circ V_p; 3, \psi_0, \hat{L}_0, 6, \psi_1, \hat{L}_1, \dots, 3(p+1), \psi_p, \hat{L}_p). \end{array}$$

We claim that this injection is a weak homotopy equivalence. In order to see this, let E denote the space of triples $((s, W), e, \hat{L})$ consisting of an $(s, W) \in \mathcal{F}_j(Q)$, an embedding $e : V_p \hookrightarrow W$ relative to Q , and a path of θ -structures \hat{L} from $e^* \hat{\ell}_W$ to $\hat{\ell}_{V_p}$ which is constant over $Q \subset V_p$. We topologise E as a subspace of $\mathcal{F}_j(Q) \times \text{Emb}(V_p, [0, \infty) \times \mathbb{R}^\infty; Q) \times \text{Bun}^\theta(TV_p)^I$. The map

$$\begin{array}{l} E \rightarrow \mathcal{Y}_j(Q)_p \\ (s, W, e) \mapsto (s, W; 3, e \circ \psi_0, \psi_0^* \hat{L}, 6, e \circ \psi_1, \psi_1^* \hat{L}, \dots, 3(p+1), e \circ \psi_p, \psi_p^* \hat{L}) \end{array}$$

is then a weak homotopy equivalence, because the embedding

$$([0, \varepsilon] \times Q) \cup \left(\prod_{i=0}^p \psi_i(D^n \times D^n) \right) \rightarrow V_p$$

is an isotopy equivalence for sufficiently small ε . On the other hand, the map

$$\begin{aligned} E &\longrightarrow \text{Emb}(V_p, [0, \infty) \times \mathbb{R}^\infty; Q) \\ (s, W, e) &\longmapsto e \end{aligned}$$

is a fibration (by the parameterised isotopy extension theorem) and the base space is contractible (by the parameterised form of Whitney's embedding theorem). The fibre of this map over the canonical embedding $V_p \subset [0, 1] \times \mathbb{R}^\infty$ is the space of those $N \in \mathcal{F}_j(Q)$ which contain the cobordism V_p , and this is weakly homotopy equivalent to $\mathcal{F}_j(S_p)$. This proves the claim that (4.5) is a weak homotopy equivalence.

The cobordisms $V_p : Q \rightsquigarrow S_p$ and $M \circ r(M) : R \rightsquigarrow Q$ are composable and have disjoint support. We may perform interchange of support on these cobordisms, giving composable cobordisms

$$\mathcal{L}_{M \circ r(M)}(V_p) : R \rightsquigarrow S'_p \quad \mathcal{R}_{V_p}(M \circ r(M)) : S'_p \rightsquigarrow S_p$$

and the embeddings ψ_i have image inside $\mathcal{L}_{M \circ r(M)}(V_p)$ and satisfy $\psi_i^* \hat{\ell}_{\mathcal{L}_{M \circ r(M)}(V_p)} = \hat{\ell}_{3(i+1)}^{\text{std}}$. Thus $\mathcal{L}_{M \circ r(M)}(V_p)$ has the same relationship to R as V_p does to Q , and in particular the inclusion

$$\begin{aligned} \mathcal{F}_j(S'_p) &\longrightarrow \mathcal{Y}_j(R)_p \\ (s, X) &\longmapsto (s+1, X \circ \mathcal{L}_{M \circ r(M)}(V_p); 3, \psi_0, \hat{L}_0, 6, \psi_1, \hat{L}_1, \dots, 3(p+1), \psi_p, \hat{L}_p). \end{aligned}$$

is a weak homotopy equivalence for the same reason as (4.5) is.

Consider the diagram

$$(4.6) \quad \begin{array}{ccc} \mathcal{F}(S_p) & \xrightarrow{-\circ \mathcal{R}_{V_p}(M \circ r(M))} & \mathcal{F}(S'_p) \\ \downarrow \simeq & & \downarrow \simeq \\ \mathcal{Y}(Q)_p & \xrightarrow{(-\circ M \circ r(M))_p} & \mathcal{Y}(R)_p, \end{array}$$

in which the vertical maps are weak homotopy equivalences by taking the limit $j \rightarrow \infty$ of the weak homotopy equivalences established above. By Lemma 4.8 the cobordism $\mathcal{R}_{V_p}(M \circ r(M))$ is isotopic to H_{S_p} composed with an isomorphism in \mathcal{D} , and so by Theorem 3.1 the map $-\circ \mathcal{R}_{V_p}(M \circ r(M))$ is an abelian homology equivalence. Hence if the square commutes up to homotopy then we have proved the proposition.

The two directions around the square are given by

$$X \mapsto (X \circ \mathcal{R}_{V_p}(M \circ r(M)) \circ \mathcal{L}_{M \circ r(M)}(V_p); 3, \psi_0, \hat{L}_0, 6, \psi_1, \hat{L}_1, \dots, 3(p+1), \psi_p, \hat{L}_p)$$

for the upper composition and

$$X \mapsto (X \circ V_p \circ M \circ r(M); 3, \varepsilon_2(\psi_0), \varepsilon_2(\hat{L}_0), \dots, 3(p+1), \varepsilon_2(\psi_p), \varepsilon_2(\hat{L}_p))$$

for the lower composition. The path $t \mapsto \tau_t$ in $\mathcal{D}(R, S_p)$ from $V_p \circ M \circ r(M)$ to $\mathcal{R}_{V_p}(M \circ r(M)) \circ \mathcal{L}_{M \circ r(M)}(V_p)$ given by interchange of support (cf. Section 4.1) may be described as follows. For a cobordism $(s, W) : A \rightsquigarrow B$ we write

$$\widehat{W} = ((-\infty, 0] \times A) \cup W \cup ([s, \infty) \times B)$$

for its extension, then we let τ_t denote the θ -manifold

$$(t \cdot e_0 + \widehat{M \circ r(M)} \setminus \text{supp}(V_p)) \cup ((2-2t) \cdot e_0 + \widehat{V_p} \setminus \text{supp}(M \circ r(M)))$$

intersected with $[0, 3] \times \mathbb{R}^\infty$, which is schematically shown in Figure 2.

The extruded embedding $\varepsilon_{2t}(\psi_i)$ has image inside τ_{1-t} , by inspection of the formulae for the two objects. Furthermore the bundle map $\varepsilon_{2t}(\psi_i)^* \hat{\ell}_{\tau_{1-t}} : T(D^n \times$




FIGURE 2. Interchanging the supports of V_p and $M \circ r(M)$. The grey regions indicate where the cobordisms are not cylindrical.

$D^n) \rightarrow \theta^* \gamma_{2n}$ agrees with the extruded bundle map $\varepsilon_{2t}(\hat{L}_i) = \varepsilon_{2t}(\hat{\ell}_{3(i+1)}^{\text{std}})$, by construction. Hence sending X to

$$(X \circ \tau_{1-t}; 3, \varepsilon_{2t}(\psi_1), \varepsilon_{2t}(\hat{L}_1), 6, \varepsilon_{2t}(\psi_2), \varepsilon_{2t}(\hat{L}_2), \dots, 3(1+p), \varepsilon_{2t}(\psi_p), \varepsilon_{2t}(\hat{L}_p))$$

gives a homotopy in $\mathcal{Y}(R)_p$ between the two compositions in (4.6), as required. \square

Now, for any system \mathcal{L} of abelian local coefficients on $\mathcal{F}(R)$ we have a map of spectral sequences

$$\begin{array}{ccccc} E_{p,q}^1(P) = H_q(\mathcal{Y}(Q)_p; \mathcal{L}) & \Longrightarrow & H_{p+q}(|\mathcal{Y}(Q)_\bullet|; \mathcal{L}) & \xrightarrow{\sim} & H_{p+q}(\mathcal{F}(Q); \mathcal{L}) \\ \downarrow & & \downarrow & & \downarrow \\ E_{p,q}^1(R) = H_q(\mathcal{Y}(R)_p; \mathcal{L}) & \Longrightarrow & H_{p+q}(|\mathcal{Y}(R)_\bullet|; \mathcal{L}) & \xrightarrow{\sim} & H_{p+q}(\mathcal{F}(R); \mathcal{L}) \end{array}$$

where we have continued to write \mathcal{L} for this system of local coefficients pulled back along any of the maps occurring in this set-up. By Proposition 4.14 the map of E^1 -pages is an isomorphism, as \mathcal{L} is an *abelian* system of local coefficients. Thus the middle vertical map is also an isomorphism. The two rightmost horizontal maps are isomorphisms because the augmentation maps (4.3) and (4.4) are weak homotopy equivalences, and hence the rightmost vertical map is too. This shows that

$$- \circ M \circ r(M) : \mathcal{F}(Q) \longrightarrow \mathcal{F}(R)$$

induces an isomorphism on homology with abelian coefficients, so $M \circ r(M) \in \mathcal{W}$.

Proof of Theorem 4.1, assuming Theorem 4.12. We must show that $M \in \mathcal{W}$. However, the discussion above holds for any cobordism which admits a handle structure relative to its outgoing boundary consisting of a single n -handle attached to the basepoint component. In particular, it applies to $r(M)$ too, so $r(M) \circ r(r(M)) \in \mathcal{W}$. By the argument of Lemma 2.19, all three morphisms in the composition $M \circ r(M) \circ r(r(M))$ are in \mathcal{W} . This establishes Theorem 4.1. \square

5. CONTRACTIBILITY OF THE HIGHER-DIMENSIONAL ARC COMPLEX

In this section we will give the proof of Theorem 4.12. It is convenient to work with two more flexible approximations to the semi-simplicial space of Definition 4.9. In the first we relax the condition that the handles be embedded: we allow them to be immersed, as long as their cores $c(D^n \times \{0\})$ are still embedded.

Definition 5.1. Fix data $(W, \chi, \hat{\ell}_t^{\text{std}})$ as in Definition 4.9. Let us write $\bar{Y}(W)_0 = \bar{Y}(W, \chi, \hat{\ell}_t^{\text{std}})_0$ for the set of triples (t, c, \hat{L}) where $t \in (2, \infty)$, $c : D^n \times D^n \looparrowright W$ is an immersion, and $\hat{L}(\tau)$ is a path of θ -structures on $D^n \times D^n$, $\tau \in [0, 1]$ such that

- (i) $c|_{D^n \times \{0\}}$ is an embedding, and there is a $\delta > 0$ such that $c(x, v) = \chi(\frac{x}{|x|}, v + t \cdot e_1) + (1 - |x|) \cdot e_0$ for $1 - |x| < \delta$,
- (ii) $c(D^n \times D^n)$ lies outside of $([0, s] \times L) \cup (\{s\} \times K|_j)$, and $c^{-1}(P) = \partial D^n \times D^n$,
- (iii) the restriction $\ell_W|_{W \setminus c(D^n \times \{0\})} : W \setminus c(D^n \times \{0\}) \rightarrow B$ is n -connected,
- (iv) $\hat{L}(0) = c^* \hat{\ell}_W$, $\hat{L}(1) = \hat{\ell}_t^{\text{std}}$, and $\hat{L}(\tau)|_{\partial D^n \times D^n}$ is independent of $\tau \in [0, 1]$.

We topologise $\bar{Y}(W)_0$ as in Definition 4.9, but using the C^∞ topology on the space of immersions rather than embeddings. Let $\bar{Y}(W)_p \subset (Y(W)_0)^{p+1}$ be the subspace consisting of ordered tuples $((t_0, c_0, \hat{L}_0), \dots, (t_p, c_p, \hat{L}_p))$ such that

- (i) the sets $c_i(D^n \times \{0\})$ are disjoint,
- (ii) $t_0 < t_1 < \dots < t_p$,
- (iii) ℓ_W restricts to an n -connected map $W \setminus (\cup_i c_i(D^n \times \{0\})) \rightarrow B$.

We shall also write $\bar{Y}^\delta(W)_p$ for the set $\bar{Y}(W)_p$ equipped with the discrete topology.

In the most flexible approximation, we further relax the condition that the cores be embedded, as long as they are immersed and in general position. We also forget condition (iii), and topologise it discretely.

Definition 5.2. Fix data $(W, \chi, \hat{\ell}_t^{\text{std}})$ as in Definition 4.9. Let us write $\hat{Y}^\delta(W)_0 = \hat{Y}^\delta(W, \chi, \hat{\ell}_t^{\text{std}})_0$ for the set of triples (t, c, \hat{L}) where $t \in (2, \infty)$, $c : D^n \times D^n \looparrowright W$ is an immersion, and \hat{L} is a path of θ -structures on $D^n \times D^n$, such that

- (i') the immersion $c|_{D^n \times \{0\}}$ is self-transverse and has no triple points, and there is a $\delta > 0$ such that $c(x, v) = \chi(\frac{x}{|x|}, v + t \cdot e_1) + (1 - |x|) \cdot e_0$ for $1 - |x| < \delta$.

as well as (ii) and (iv) of Definition 5.1. Note that the data is *not* required to satisfy (iii). Let $\hat{Y}^\delta(W)_p \subset (\hat{Y}^\delta(W)_0)^{p+1}$ be the subset consisting of ordered tuples $((t_0, c_0, \hat{L}_0), \dots, (t_p, c_p, \hat{L}_p))$ such that

- (i) the immersions $c_i|_{D^n \times \{0\}}$ are in general position (i.e. pairwise transverse and without triple intersections),
- (ii) $t_0 < t_1 < \dots < t_p$.

As usual, this data defines a semi-simplicial set $\hat{Y}^\delta(W)_\bullet$.

Lemma 5.3. *The map $|Y(W)_\bullet| \rightarrow |\bar{Y}(W)_\bullet|$ is a weak homotopy equivalence.*

Proof. The inclusion $Y(W)_\bullet \subset \bar{Y}(W)_\bullet$ is a levelwise weak homotopy equivalence, by precomposing an immersion $h : D^n \times D^n \looparrowright W$ which is an embedding of the core, and hence of a neighbourhood of the core, with an isotopy from the identity map of $D^n \times D^n$ to an embedding into a small neighbourhood of $(\partial D^n \times D^n) \cup (D^n \cup \{0\})$. \square

By this lemma, in order to prove Theorem 4.12 it is enough to show that

$$\text{hocolim}_{g \rightarrow \infty} |\bar{Y}(K|_{[j, j+g]} \circ W, \chi, \hat{\ell}_t^{\text{std}})_\bullet|$$

is weakly contractible. Since homotopy colimit commutes with geometric realisation, we may equivalently prove weak contractibility of the realisation of the semi-simplicial space

$$[p] \longmapsto \text{hocolim}_{g \rightarrow \infty} \bar{Y}(K|_{[j, j+g]} \circ W, \chi, \hat{\ell}_t^{\text{std}})_p.$$

Since each map $(K|_{[i, i+1]} \circ -)_p$ forming the colimit diagram is the inclusion of a subspace, we may replace the homotopy colimit by the actual colimit. We shall therefore prove weak contractibility of the semi-simplicial space whose space of p -simplices is

$$\bar{Y}(K|_{[j, \infty)} \circ W)_p = \text{colim}_{g \rightarrow \infty} \bar{Y}(K|_{[j, j+g]} \circ W, \chi, \hat{\ell}_t^{\text{std}})_p.$$

(The proof would work for the homotopy colimit, but it is notationally convenient to work with the actual colimit.) We shall first prove that the underlying semisimplicial set $\bar{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet$ has contractible realisation, and then use the techniques of [GRW14a] to deduce weak contractibility in the topologised case.

Lemma 5.4. *The realisation $|\widehat{Y}^\delta(W)_\bullet|$ is contractible.*

Proof. We first show that $\widehat{Y}^\delta(W)_0 \neq \emptyset$. Pick some $t \geq 2$ and consider the commutative square

$$\begin{array}{ccc} \partial D^n \times D^n & \xrightarrow{\chi(-, -+t \cdot e_1)} & PC \xrightarrow{\cong} W \\ \downarrow & \searrow \scriptstyle g & \downarrow \scriptstyle \ell_W \\ D^n \times D^n & \xrightarrow{\ell_t^{\text{std}}} & B. \end{array}$$

As ℓ_W is n -connected, and the pair $(D^n \times D^n, \partial D^n \times D^n)$ is homotopy equivalent to $(D^n, \partial D^n)$, there exists a dashed diagonal map g making the top triangle commute and the bottom triangle commute up to homotopy. As ℓ_W and ℓ_t^{std} are covered by bundle maps $\hat{\ell}_W$ and $\hat{\ell}_t^{\text{std}}$, this provides a bundle map $\hat{g} : T(D^n \times D^n) \rightarrow TW$ such that $\hat{\ell}_W \circ \hat{g}$ is homotopic to $\hat{\ell}_t^{\text{std}}$ through bundle maps, via a homotopy which is constant over $\partial D^n \times D^n$. By Smale–Hirsch theory, the pair (g, \hat{g}) may be homotoped relative to $\partial D^n \times D^n$ to a pair of the form (c, Dc) , for c an immersion and Dc its differential. Then $c^* \hat{\ell}_W = \hat{\ell}_W \circ Dc$ is still homotopic to $\hat{\ell}_t^{\text{std}}$, and if we choose such a homotopy, \hat{L} , then we have constructed an element $(t, c, \hat{L}) \in \widehat{Y}^\delta(W)_0$.

To prove contractibility assuming non-emptiness, we must prove that for each $k \geq 1$, any map $f : \partial I^k \rightarrow |\widehat{Y}^\delta(W)_\bullet|$ extends to a map from I^k . By the simplicial approximation theorem, we can assume that f is simplicial with respect to some PL triangulation of ∂I^k . For each vertex $v_i \in \partial I^k$ in the triangulation, there is then given an element $f(v_i) = (t_i, c_i, \hat{L}_i) \in \widehat{Y}^\delta(W)_0$. By a suitable application of Thom’s transversality theorem, we may let $(t, c, \hat{L}) \in \widehat{Y}^\delta(W)_0$ be a slight perturbation of one of the (t_i, c_i, \hat{L}_i) such that $t \neq t_j$ and $c(D^n \times \{0\}) \pitchfork c_j(D^n \times \{0\})$ for all j . A further perturbation will remove triple intersections between the cores. Then $f(\partial I^k)$ is contained in the star of (t, c, \hat{L}) , and therefore f extends to the cone $C(\partial I^k) \cong I^k$. \square

The (rather lengthy) proof of the following proposition makes essential use of the passage to the limit $g \rightarrow \infty$.

Proposition 5.5. *The natural maps $\bar{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet \rightarrow \widehat{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet$ induce a weak equivalence on geometric realisation.*

Proof. To ease notation throughout the proof, let us write $M = K|_{[j,\infty)} \circ W$. Consider a map

$$f : (I^k, \partial I^k) \longrightarrow (|\widehat{Y}^\delta(M)_\bullet|, |\bar{Y}^\delta(M)_\bullet|)$$

which we may assume simplicial with respect to some PL triangulation of I^k . We shall explain how to homotope f to a map with image in $|\bar{Y}^\delta(M)_\bullet|$. If the image of some simplex $\sigma < I^k$ is a p -simplex $f(\sigma) = ((t_0, c_0, \hat{L}_0), \dots, (t_p, c_p, \hat{L}_p))$, then there can be three reasons why $f(\sigma)$ is not in this subcomplex: firstly, the cores $D_i = c_i(D^n \times \{0\})$ may not be embedded; secondly, they may not be pairwise disjoint; thirdly, the restriction $\ell|_{M \setminus \cup D_i} : M \setminus \cup D_i \rightarrow B$ may not be n -connected. All three problems will be fixed using the infinite supply of embedded copies of $W_{1,1} = S^n \times S^n \setminus \text{int}(D^{2n})$ given by the θ -end, using a well-known local move. Let us first explain it.

Suppose that $h_0, h_1 : D^n \times D^n \looparrowright M$ are two immersed handles, such that each of the cores $h_i|_{D^n \times \{0\}}$ is self-transverse and has no triple points, and they also meet

transversely. Around a point of intersection x_0 between the two cores, we may find a coordinate chart $\varphi : D^n \times D^n \hookrightarrow M$ inside which the cores are $\varphi(D^n \times \{0\})$ and $\varphi(\{0\} \times D^n)$ respectively. By scaling the handles in the meridian direction near the discs $h_0^{-1}\varphi(D^n \times \{0\})$ and $h_1^{-1}\varphi(\{0\} \times D^n)$, and making a change of coordinates, we may obtain new immersed handles h'_i which agree with the old h_i outside a small neighbourhood of $\varphi(D^n \times D^n)$, and which intersect $\varphi(D^n \times D^n)$ in $\varphi(D^n \times \frac{1}{2}D^n)$ and $\varphi(\frac{1}{2}D^n \times D^n)$ respectively. This preliminary move is illustrated in Figure 3.




FIGURE 3. Shrinking and straightening handles near an intersection point so that they intersect a coordinate chart cleanly.

We may then find an embedded path in M from the boundary of $\varphi(D^n \times D^n)$ to an embedded copy of $W_{1,1}$ having standard θ -structure, ensure that it is disjoint from the cores of all the h'_i , and thicken it up to obtain an embedding $\varphi' : (D^n \times D^n) \natural W_{1,1} \hookrightarrow M$, so that the handles intersect the image of φ' in precisely $\varphi'(D^n \times \frac{1}{2}D^n)$ and $\varphi'(\frac{1}{2}D^n \times D^n)$.




FIGURE 4. Extending the coordinate chart φ to an embedding of $(D^n \times D^n) \natural W_{1,1}$, and using it to remove a point of intersection.

Now, as illustrated in Figure 4, there are *disjoint* embeddings

$$\begin{aligned} j_0 : D^n \times \frac{1}{2}D^n &\longrightarrow (D^n \times D^n) \natural W_{1,1}, \\ j_1 : \frac{1}{2}D^n \times D^n &\longrightarrow (D^n \times D^n) \natural W_{1,1}, \end{aligned}$$

which near $\partial(D^n \times D^n)$ are given by $j_i(x, v) = (x, v)$. (These may be constructed as follows: $(D^n \times D^n) \natural W_{1,1}$ is diffeomorphic to $S^n \times S^n \setminus \text{int}(D^n_- \times D^n_-)$, the manifold obtained by removing points having both coordinates in the lower hemisphere. This diffeomorphism can be taken to be the identity on that part of the boundary of $(D^n \times D^n) \natural W_{1,1}$ away from where the connect-sum was formed. Inside $S^n \times S^n \setminus \text{int}(D^n_- \times D^n_-)$ the discs $D^n_+ \times \frac{1}{2}D^n_-$ and $\frac{1}{2}D^n_- \times D^n_+$ are disjoint, and satisfy the required boundary condition.) Replacing the map h'_0 on $(h'_0)^{-1}(\varphi(D^n \times D^n))$ by

$$(h'_0)^{-1}(\varphi(D^n \times D^n)) \xrightarrow[\sim]{\varphi^{-1} \circ h'_0} D^n \times \frac{1}{2}D^n \xrightarrow{j_0} (D^n \times D^n) \natural W_{1,1} \xrightarrow{\varphi'} M,$$

and the map h'_1 on $(h'_1)^{-1}(\varphi(D^n \times D^n))$ by

$$(h'_1)^{-1}(\varphi(D^n \times D^n)) \xrightarrow[\cong]{\varphi^{-1} \circ h'_1} \frac{1}{2}D^n \times D^n \xrightarrow{j_1} (D^n \times D^n) \natural W_{1,1} \xrightarrow{\varphi'} M,$$

we obtain new handles h''_i which no longer intersect at x_0 , and whose cores outside of $\varphi'(D^n \times D^n) \natural W_{1,1}$ are unchanged.

Finally, we claim that the θ -structures $(h'_0)^*\ell_M$ and $(h''_0)^*\ell_M$ on $D^n \times D^n$, which are already equal outside of $(h'_0)^{-1}(\varphi(D^n \times D^n))$, are homotopic relative to the complement of this subset. This will make use of the embedded copy of $W_{1,1}$ having *standard* θ -structure. If we write $\ell = (\varphi')^*\ell_M$ it suffices to see that the two embeddings

$$D^n \times \frac{1}{2}D^n \longrightarrow (D^n \times D^n) \natural W_{1,1},$$

given by the standard embedding and by j_0 , pull back ℓ to θ -structures which are homotopic relative to $\partial D^n \times \frac{1}{2}D^n$. The two embeddings differ by forming the ambient connected-sum (inside $\varphi'(D^n \times D^n) \natural W_{1,1}$) with the sphere

$$S^n \times \{0\} \subset S^n \times D_+^n \subset S^n \times S^n \setminus \text{int}(D_-^n \times D_-^n) \approx W_{1,1}.$$

As the θ -structure on this sphere is standard, it extends over the contractible space

$$S^n \times D_+^n \approx (D^{n+1} \setminus \frac{1}{2}\mathring{D}^{n+1}) \times D^{n-1} \subset D^{n+1} \times D^{n-1},$$

and so the θ -structure obtained after forming the ambient connected-sum is homotopic to the original one, as required. The analogous claim holds for $(h'_1)^*\ell_M$ and $(h''_1)^*\ell_M$.

Note that as the argument above took place locally near the intersection point, it works equally well for a transverse self-intersection of a core of an immersed handle.

We now explain how to implement these moves. We first explain the argument in the case $n > 1$, where the cores (having codimension at least 2) cannot separate. In the case $n = 1$ the argument must be reorganised slightly, and we shall explain that at the end.

Step 1. We first explain how the map f may be changed by a homotopy so that for each vertex $v \in I^k$ with $f(v) = (t, c, \hat{L})$, the manifold $D = c(D^n \times \{0\})$ is embedded. If $v \in \text{int}(I^k)$ with $f(v) = (t, c, \hat{L})$ is a vertex which does not satisfy this condition, then the map $c|_{D^n \times \{0\}}$ is an immersion with a finite number of transverse self-intersections.

Firstly, let c' be obtained by perturbing c a small amount ε in the normal direction $e_1 \in \mathbb{R}^n$. Precisely, we have a field of vectors tangent to M defined on $\chi(\partial D^n \times (1, \infty) \times D^{n-1})$ given by the unit vector in the $(1, \infty)$ direction. We may extend this to a compactly supported smooth vector field on M which is nowhere tangent to $c(D^n \times \{0\})$, and (choosing a Riemannian metric) obtain a 1-parameter family φ_t of compactly-supported diffeomorphisms of M . Then c' is the map obtained by applying φ_ε to c for some small ε ; if ε is small enough, the core of the immersion c' has as many points of self-intersection as that of c , and is transverse to the core of c .

Secondly, let c'' be obtained from c' by applying the move using $W_{1,1}$ described above to remove a point of self-intersection. We make sure that the copy of $W_{1,1}$, and the path used to form the connect-sum, are disjoint from the core $c'(D^n \times \{0\})$ so that no new self-intersections are created when performing this move. We may also ensure that no non-transverse intersections with other cores are created.

To obtain a zero-simplex $(t + \varepsilon, c'', \hat{L}'')$, it remains to find the path \hat{L}'' . The θ -structures $(c')^*\hat{\ell}_M$, $(c'')^*\hat{\ell}_M$, and $\hat{\ell}_{t+\varepsilon}^{\text{std}}$ are all equal on $\partial D^n \times D^n$. The isotopy φ_s , $s \in [0, \varepsilon]$, gives paths $(c')^*\hat{\ell}_M \rightsquigarrow c^*\hat{\ell}_M$ and $\hat{\ell}_{t+\varepsilon}^{\text{std}} \rightsquigarrow \hat{\ell}_t^{\text{std}}$ which become equal when restricted to $\partial D^n \times D^n$. Conjugating the path $\hat{L} : c^*\hat{\ell}_M \rightsquigarrow \hat{\ell}_t^{\text{std}}$ by these

determines a path $(c')^*\hat{\ell}_M \rightsquigarrow \hat{\ell}_{t+\varepsilon}^{\text{std}}$ which restricts to a nullhomotopic loop over $\partial D^n \times D^n$. We may then use homotopy extension to find a path $\hat{L}' : (c')^*\hat{\ell}_M \rightsquigarrow \hat{\ell}_{t+\varepsilon}^{\text{std}}$ which is constant over $\partial D^n \times D^n$. Finally, by our discussion of the move above, $(c')^*\hat{\ell}_M$ and $(c'')^*\hat{\ell}_M$ are homotopic relative to $\partial D^n \times D^n$, so we may find a path $\hat{L}'' : (c'')^*\hat{\ell}_M \rightsquigarrow \hat{\ell}_{t+\varepsilon}^{\text{std}}$ which is constant over $\partial D^n \times D^n$.

By construction the vertices (t, c, \hat{L}) and $(t + \varepsilon, c'', \hat{L}'')$ span a 1-simplex, and if $f(w) = (t_i, c_i, \hat{L}_i)$ for w a vertex adjacent to v , then $(t + \varepsilon, c'', \hat{L}'')$ and (t_i, c_i, \hat{L}_i) are adjacent too (when ε is small enough, the relative order of $t + \varepsilon$ and t_i is the same as that of t and t_i , and if c and c_i are in general position so are c'' and c_i). Hence we may define a simplicial map

$$F : [0, 1] \times I^k \longrightarrow |\widehat{Y}^\delta(M)_\bullet|$$

by $F(0, v) = f(v) = (t, c, \hat{L})$, $F(1, v) = (t + \varepsilon, c'', \hat{L}'')$, and $F(-, w) = f(w)$ for $w \neq v$. This is a homotopy which is constant on ∂I^k , and $F(1, -)$ sends v to an immersion whose core has one fewer self-intersection. After finitely many applications of this technique, we may change f by a homotopy so that each disc is embedded.

Step 2. We now explain how to modify f so that it sends adjacent vertices in I^k to embeddings with disjoint cores. For each pair of adjacent vertices (v, u) in I^k mapping to a 1-simplex $((t_0, c_0, \hat{L}_0), (t_1, c_1, \hat{L}_1))$ for which $D_0 = c_0(D^n \times \{0\})$ intersects $D_1 = c_1(D^n \times \{0\})$, we create embeddings c''_0 and c''_1 by the process described in Step 1: we perturb each c_i a little in the e_1 -direction, and then use the move to eliminate a point of intersection (using a path from the intersection point to a $W_{1,1}$ which is disjoint from the D_i and from any adjacent cores). We obtain $(t_0 + \varepsilon, c''_0, \hat{L}''_0)$ disjoint from (t_0, c_0, \hat{L}_0) and $(t_1 + \varepsilon, c''_1, \hat{L}''_1)$ disjoint from (t_1, c_1, \hat{L}_1) , with c''_0 and c''_1 having one fewer point of intersection than c_0 and c_1 did. Hence we may define a simplicial map

$$F : [0, 1] \times I^k \longrightarrow |\widehat{Y}^\delta(M)_\bullet|$$

by $F(0, u) = f(u) = (t_0, c_0, \hat{L}_0)$, $F(1, u) = (t_0 + \varepsilon, c''_0, \hat{L}''_0)$, similarly for v , and $F(-, w) = f(w)$ for $w \neq u$ or v . Since c''_i is disjoint from c_i this gives a homotopy of pairs $(I^k, \partial I^k) \longrightarrow (|\widehat{Y}^\delta(M)_\bullet|, |\overline{Y}^\delta(M)_\bullet|)$ after which the total number of intersection points has been reduced by one. After finitely many applications of this technique, we obtain a map f sending adjacent vertices in I^k to embeddings with disjoint cores.

Step 3. Finally, we fix simplices $\sigma < I^k$ with $f(\sigma) = ((t_0, c_0, \hat{L}_0), \dots, (t_p, c_p, \hat{L}_p))$ for which the restriction $\ell_{M \setminus \cup D_i} : M \setminus \cup D_i \rightarrow B$ is not n -connected. The discs D_i that we have cut out have codimension n , so the inclusion $M \setminus \cup D_i \hookrightarrow M$ is $(n-1)$ -connected. Thus $\ell_{M \setminus \cup D_i}$ can fail to be n -connected either because it is not surjective on π_n or because it is not injective on π_{n-1} . Thus we set

$$(5.1) \quad \begin{aligned} K_\sigma &= \text{Ker}(\pi_{n-1}(M \setminus \cup D_i) \rightarrow \pi_{n-1}(B)) \\ I_\sigma &= \text{Im}(\pi_n(M \setminus \cup D_i) \rightarrow \pi_n(B)) \end{aligned}$$

and aim to kill the groups K_σ and $\pi_n(B)/I_\sigma$. As the map $\ell_M : M \rightarrow B$ is n -connected, it follows from the long exact sequence of a triple that the map

$$\pi_n(M, M \setminus \cup D_i) \longrightarrow \pi_n(B, M \setminus \cup D_i)$$

is surjective. If $n > 2$, so M , $M \setminus \cup D_i$, and B have a common fundamental group, π , then by the Hurewicz theorem $\pi_n(M, M \setminus \cup D_i)$ is generated as a $\mathbb{Z}[\pi]$ -module by the $(p+1)$ meridian spheres of the handles which have been cut out, so it follows that $\pi_n(B, M \setminus \cup D_i)$ is a finitely generated $\mathbb{Z}[\pi]$ -module, and hence that K_σ is also a finitely generated $\mathbb{Z}[\pi]$ -module; in fact, it is generated by the meridian

spheres of the handles that have been removed. If $n = 2$ then the Seifert–van Kampen theorem shows that analogous claim is true: K_σ is normally generated by the meridian circles of the handles that have been removed. Furthermore, if K_σ is trivial, then (if $n = 2$ the map $\pi_1(M \setminus \cup D_i) \rightarrow \pi_1(M) = \pi$ is an isomorphism and) there is an exact sequence

$$\pi_n(M \setminus \cup D_i) \longrightarrow \pi_n(B) \longrightarrow \pi_n(B, M \setminus \cup D_i) \longrightarrow 0$$

with rightmost term a finitely generated $\mathbb{Z}[\pi]$ -module. Thus $\pi_n(B)$ is generated as a $\mathbb{Z}[\pi]$ -module by $I_\sigma \subset \pi_n(B)$ along with finitely many elements (this makes sense, and holds, for $n \geq 2$). We will explain how to kill these finitely many extra elements.

Let us describe a general construction, and some of its properties. Let $\{v\} < I^k$ be an interior vertex with $f(v) = (t, c, \hat{L})$. Suppose that we modify c by choosing a standard $W_{1,1}$, and a path from a point in $D = c(D^n \times \{0\})$ to the $W_{1,1}$, both disjoint from all the cores of all vertices in $\text{Lk}(t, c, \hat{L})$ (which is possible as $n \geq 2$), and then form a new embedding c' by perturbing c a small amount ε in the e_1 -direction, and then forming the connect-sum of its core with the core of $\bar{e}(S^n \times D^n) \subset W_{1,1}$ along the path. The θ -structure $(c')^* \hat{\ell}_M$ is homotopic to $c^* \hat{\ell}_M$ extending the standard homotopy on the boundary, because the θ -structure on $\bar{e}(S^n \times D^n) \subset W_{1,1}$ is standard, and hence by the same argument as in Step 1 there is a path $\hat{L}' : (c')^* \hat{\ell}_M \rightsquigarrow \hat{\ell}_{t+\varepsilon}^{\text{std}}$. The following claim describes how to use this construction to modify $f : I^k \rightarrow |\hat{Y}^\delta(M)_\bullet|$, and the effect of the modification on the groups I_σ and K_σ defined in (5.1) above.

Claim 5.6. In the situation described above, let $f' : I^k \rightarrow |\hat{Y}^\delta(M)_\bullet|$ be the simplicial map obtained from f by changing its value at the single interior vertex $v \in I^k$: $f'(w) = f(w)$ for vertices $w \neq v$, but $f'(v) = (t + \varepsilon, c', \hat{L}')$. Then the maps f and f' are homotopic relative to ∂I^k , and furthermore

- (i) For any simplex $\sigma = (v, w_1, \dots, w_p)$, if $\sigma' = (v', w_1, \dots, w_p)$ then there is a surjection $K_\sigma \rightarrow K_{\sigma'}$, sending the class of the meridian of the core of w_i to the class of the meridian of the core of w_i , and whose kernel contains the class of the meridian of the core of v .
- (ii) For any simplex $\sigma = (v, w_1, \dots, w_p)$, if $\sigma' = (v', w_1, \dots, w_p)$, then $I_\sigma \subset I_{\sigma'}$.

Proof. The statement about the relative homotopy is immediate, using a homotopy defined as in the previous steps.

For (i), consider the region X formed by the union of $c(D^n \times D^n)$, the standard $W_{1,1}$, and a thickening of the chosen path between them, inside which the (perturbed) core $D = c(D^n \times \{0\})$ is connect-summed with $\bar{e}(S^n \times \{0\}) \subset W_{1,1}$ to form the modified core $D' = c'(D^n \times \{0\})$. If $\sigma = (v, w_1, \dots, w_p)$ is a simplex then all the cores $D_i = w_i(D^n \times \{0\})$ are disjoint from X . Thus there are maps

$$K_{\sigma'} \xleftarrow{l} \text{Ker}(\pi_{n-1}(M \setminus (X \cup_i D_i)) \rightarrow \pi_{n-1}(B)) \xrightarrow{r} K_\sigma,$$

which are both surjective by transversality, as X and the D_i have n -dimensional cores. We claim that r is also injective. If $g : S^{n-1} \rightarrow M \setminus (X \cup_i D_i)$ becomes trivial in K_σ , then there is a nullhomotopy $\bar{g} : D^n \rightarrow M \setminus (D \cup_i D_i)$. After possibly perturbing it, if this nullhomotopy intersects the core of X then it must do so by intersecting $W_{1,1} \subset X$. But $\partial W_{1,1} \cong S^{2n-1}$ is n -connected, so \bar{g} can be rechosen to miss $W_{1,1}$, and hence to lie in $M \setminus (X \cup_i D_i)$. Thus r is an isomorphism. On the other hand, under $l \circ r^{-1}$ the meridian of D in K_σ maps to the meridian of D' in $K_{\sigma'}$, which is nullhomotopic as the meridian of $\bar{e}(S^n \times \{0\}) \subset W_{1,1}$ is.

For (ii), let $\psi : S^n \rightarrow M \setminus (D \cup_i D_i)$. After perturbing it, if this map intersects the core of X then it must do so by intersecting $W_{1,1} \subset X$. But if ψ' is obtained as

above by rechoosing the part of the image of ψ in $W_{1,1}$, the homotopy class of ψ' in M is obtained from that of ψ by the addition of a class in the $\mathbb{Z}[\pi]$ -submodule generated by $\bar{e}, \bar{f} \in \pi_n(W_{1,1})$. As the θ -structure on $W_{1,1}$ is standard, $\ell_M \circ \bar{e}$ and $\ell_M \circ \bar{f}$ are nullhomotopic, and so $[\ell_M \circ \psi] = [\ell_M \circ \psi'] \in \pi_n(B)$. Thus $I_{\sigma'} \supset I_\sigma$. \square

Let us now explain how to use the Claim above to inductively fix simplices $\sigma < I^k$ with $f(\sigma) = ((t_0, c_0, \hat{L}_0), \dots, (t_p, c_p, \hat{L}_p))$ for which the restriction $\ell_{M \setminus \cup D_i} : M \setminus \cup D_i \rightarrow B$ is not n -connected. Firstly, apply the construction described above once for each interior vertex $v \in I^k$. This changes the map f by a homotopy, and for every simplex σ we now have that K_σ is generated by the meridians of the cores of its vertices, but also by Claim 5.6 (i) that the classes of all these meridians are trivial. Thus $K_\sigma = 0$ for every simplex $\sigma < I^k$.

It remains to explain how to achieve $I_\sigma = \pi_n(B)$ for all simplices $\sigma < I^k$. By the discussion before Claim 5.6, when $K_\sigma = 0$ it follows that the fundamental group of $M \setminus \cup D_i$ is also π and that $\pi_n(B)$ is generated as a $\mathbb{Z}[\pi]$ -module by I_σ along with finitely-many additional elements. We may thus choose finitely-many maps $\{g_\alpha : S^n \rightarrow M\}_{\alpha \in J}$ such that $[\ell_M \circ g_\alpha] \in \pi_n(B)$ are the additional generators, and we may suppose that the g_α are transverse to the discs D_i . Using the same move to eliminate intersection points as in all the previous steps, with embedded paths and copies of $W_{1,1}$ chosen disjoint from all the c_i , the g_α , and any cores in the link of any vertex of σ , we iteratedly remove points of intersection between the c_i and the g_α until they are disjoint. As in Steps 1 and 2, this changes the map f by a simplicial homotopy. The effect of this move on each c_i is to connect-sum its core with the core of $\bar{e}(S^n \times D^n) \subset W_{1,1}$, and so is described by Claim 5.6. In particular, we must still have $K_\tau = 0$ for every simplex $\tau < I^k$ by Claim 5.6 (i), and must have $I_\sigma \subset I_{\sigma'}$ by Claim 5.6 (ii). But in addition, the map obtained from g_α is now disjoint from the cores of the vertices of σ , so gives a class $[g'_\alpha] \in \pi_n(M \setminus \cup_i D_i)$. Furthermore, although $[g_\alpha] \neq [g'_\alpha] \in \pi_n(M)$, we have $[\ell_M \circ g_\alpha] = [\ell_M \circ g'_\alpha]$, because the spheres with which we took the connected sum of g_α were inside copies of $W_{1,1}$ with standard θ -structure, and so had nullhomotopic maps to B . Thus the classes $[\ell_M \circ g_\alpha] \in \pi_n(B)$ also lie in $I_{\sigma'}$, but this means that $I_{\sigma'} = \pi_n(B)$. We have therefore arranged that $\ell_{M \setminus \cup D_i} : M \setminus \cup D_i \rightarrow B$ is n -connected. By a further application of Claim 5.6 (i) and (ii), this cannot be undone as we go on to fix other simplices.

The case $n = 1$. In this case we use the above techniques, but with some reorganisation. We sketch the necessary changes.

We first show that $\bar{Y}^\delta(M)_0$ is non-empty. To do this, choose a vertex in the non-empty set $\hat{Y}^\delta(M)_0$, represented by a tuple (t, c, \hat{L}) with $c|_{D^1 \times \{0\}}$ a self-transverse immersed arc having no triple points. This immersed arc may separate M , but its complement will always have a (unique) non-compact path component, which we shall call the “non-compact region”. Step 1 can now be performed as follows: choose a self-intersection point which is in the closure of the non-compact region, and choose a path in the non-compact region to a copy of $W_{1,1}$, then use the move described above to remove this intersection point. Continuing in this way, we may remove all self-intersection points, and hence suppose that $c|_{D^1 \times \{0\}}$ is an embedding. We then perform the move in Step 3, forming the connect-sum of $D = c(D^1 \times \{0\})$ with a copy of $S^1 \times \{x\} \subset (S^1 \times S^1) \setminus \text{int}(D^2) = W_{1,1}$ along a path in the non-compact region: this ensures that the arc D is non-separating (which in this dimension is the analogue of K_σ being trivial). We then proceed as in Step 3, to ensure that $\ell_M|_{M \setminus D} : M \setminus D \rightarrow B$ is 1-connected (as D is non-separating, we may choose all the necessary paths to be disjoint from it). We have produced the data of a vertex of $\bar{Y}^\delta(M)$.

We now show that $|\overline{Y}^\delta(M)_\bullet|$ is contractible. Let $f : S^{k-1} \rightarrow |\overline{Y}^\delta(M)_\bullet|$ be a map, which we may assume is simplicial with respect to some PL triangulation $|K| \approx S^{k-1}$, and let $(t_0, c_0, \hat{L}_0) \in \overline{Y}^\delta(M)_0$ be a vertex (for example, one constructed as above). As K has finitely many vertices, we may perturb the data (t_0, c_0, \hat{L}_0) so that its core $D = c_0(D^1 \times \{0\})$ is transverse to the core $D_v = c_v(D^1 \times \{0\})$ of $f(v) = (t_v, c_v, \hat{L}_v)$ for every $v \in K$. Furthermore, by stretching we may ensure that D is not disjoint from the non-compact region of $M \setminus \cup_{v \in K} D_v$. Now we may proceed as in Step 2, using paths in the non-compact region of $(M \setminus \cup_{v \in K} D_v) \setminus D$ to remove intersections between D and the D_v . This can be done so that it changes f by a homotopy, by ensuring that when forming a parallel copy of D_v it remains disjoint from any $D_{v'}$ that it was already disjoint from. This move also changes (t_0, c_0, \hat{L}_0) and hence D , but does not change the fact that it intersects the boundary of the non-compact region of $(M \setminus \cup_{v \in K} D_v) \setminus D$. After finitely-many applications we may therefore suppose that all D_v are disjoint from D .

Finally, we choose a $W_{1,1}$ and a path from D to it in the non-compact region, and change c_0 by forming the connect-sum of its core with that of $\bar{e}(S^1 \times D^1) \subset W_{1,1}$. For any simplex $\sigma < S^{k-1}$ the map $M \setminus \cup_{v \in \sigma} D_v \rightarrow B$ is 1-connected. The arc D is non-separating in $M \setminus \cup_{v \in \sigma} D_v$, as $\bar{e}(S^1 \times \{0\})$ is non-separating in $W_{1,1}$, and furthermore D has a dual circle C in $W_{1,1}$. It follows that $(M \setminus \cup_{v \in K} D_v) \setminus D$ is path connected, and that $\pi_1(M \setminus \cup_{v \in K} D_v)$ is generated by $\pi_1((M \setminus \cup_{v \in K} D_v) \setminus D)$ and the conjugacy class of C . But C is nullhomotopic in B , as the $W_{1,1}$ had standard θ -structure, and so the map $(M \setminus \cup_{v \in K} D_v) \setminus D \rightarrow B$ is 1-connected and hence $f(\sigma)$ spans a simplex with (t_0, c_0, \hat{L}_0) . This means that f has image in the star of (t_0, c_0, \hat{L}_0) , so is nullhomotopic. \square

Lemma 5.7. *The space $|\overline{Y}(K|_{[j,\infty)} \circ W)_\bullet|$ is weakly contractible.*

Proof. We have proved that $|\overline{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet|$ is contractible, so it remains to show that the map

$$|\overline{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet| \longrightarrow |\overline{Y}(K|_{[j,\infty)} \circ W)_\bullet|$$

induced by the identity is a weak homotopy equivalence. To this end, we proceed as in the proof of [GRW14a, Theorem 5.6]. In more detail, we first observe that for any $\sigma = ((t_0, c_0, \hat{L}_0), \dots, (t_p, c_p, \hat{L}_p)) \in \overline{Y}(K|_{[j,\infty)} \circ W)_p$, there is a subcomplex $F(\sigma)_\bullet \subset \overline{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet$ whose q -simplices are those $\sigma' = ((t'_0, c'_0, \hat{L}'_0), \dots, (t'_q, c'_q, \hat{L}'_q))$ satisfying

$$(5.2) \quad ((t_0, c_0, \hat{L}_0), \dots, (t_p, c_p, \hat{L}_p), (t'_0, c'_0, \hat{L}'_0), \dots, (t'_q, c'_q, \hat{L}'_q)) \in \overline{Y}^\delta(K|_{[j,\infty)} \circ W)_{p+q+1}.$$

The same argument that we used to prove contractibility of $|\overline{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet|$ applies to $F(\sigma)_\bullet$ and proves it has contractible realisation. (That space is a slight enlargement of $\overline{Y}^\delta((K|_{[j,\infty)} \circ W) \setminus \cup_i c_i(D^n \times \{0\}))_\bullet$, and the restriction of $\ell_{K|_{[j,\infty)} \circ W} |_{K|_{[j,\infty)} \circ W \setminus \cup_i c_i(D^n \times \{0\})}$ is still n -connected.) Then define the subspace $D_{p,q} \subset \overline{Y}(K|_{[j,\infty)} \circ W)_p \times \overline{Y}^\delta(K|_{[j,\infty)} \circ W)_q$ consisting of those (σ, σ') satisfying (5.2). The obvious forgetful maps make $D_{\bullet,\bullet}$ into a bi-semi-simplicial space, augmented in both directions. The augmentation $D_{p,q} \rightarrow \overline{Y}(K|_{[j,\infty)} \circ W)_p$ then induces a map $|D_{p,\bullet}| \rightarrow \overline{Y}(K|_{[j,\infty)} \circ W)_p$ with fibre $|F(\sigma)_\bullet|$ over σ , which is contractible. The argument of [GRW14a, Theorem 5.6] shows that this augmentation map is a weak equivalence (and in fact a Serre fibration). By the argument of [GRW14a, Lemma 5.8], the resulting weak equivalence $|D_{\bullet,\bullet}| \rightarrow |\overline{Y}(K|_{[j,\infty)} \circ W)_\bullet|$ factors up to homotopy through the contractible space $|\overline{Y}^\delta(K|_{[j,\infty)} \circ W)_\bullet|$, proving that all three spaces are weakly contractible. \square

6. PROOF OF THEOREM 2.15: STABILITY FOR HANDLES OF INDEX $n < k < 2n$

In this section we shall prove the following instance of Theorem 2.15.

Theorem 6.1. *If $M : P \rightsquigarrow Q$ is a morphism in \mathcal{D} whose underlying smooth cobordism admits a handle structure relative to Q consisting of a single k -handle, with $n \leq k < 2n$, attached to the basepoint component of Q , then $M \in \mathcal{W}$.*

The proof shall be by induction on k , where the case $k = n$ has already been established by Theorem 4.1. Thus we suppose that $k > n$ and that Theorem 6.1 holds for elementary cobordisms of index $k - 1$. As in the case $k = n$ the detailed proof is again cumbersome, but the strategy of the induction step can be explained informally as follows. Suppose first that there exists an elementary cobordism V of index $k - 1$ such that the composition $V \circ M = V \cup_Q M$ is defined and is diffeomorphic to a trivial cobordism. Then the composition $V \circ M$ is in \mathcal{W} and by induction V is in \mathcal{W} so by the 2-out-of-3 property we deduce that M is in \mathcal{W} . This proves the theorem for elementary bordisms M admitting a “left inverse” elementary cobordism V in this sense. We shall then use a simplicial resolution to reduce the general case to the case where M admits a left inverse.

6.1. Constructing auxiliary cobordisms. Recall that to the data of an object $Q \in \mathcal{D}$, an embedding $\sigma : \mathbb{R}^k \times \mathbb{R}^{2n-k} \rightarrow [0, \infty) \times \mathbb{R}^{\infty-1}$ with $\sigma^{-1}(Q) = \partial D^k \times \mathbb{R}^{2n-k}$, and an extension of a bundle map $(\sigma|_{\partial D^k \times D^{2n-k}})^* \hat{\ell}_Q : T(D^k \times D^{2n-k})|_{\partial D^k \times D^{2n-k}} \rightarrow \theta^* \gamma$ to $T(D^k \times D^{2n-k})$ we have associated an object P_σ and a morphism $M_\sigma : P_\sigma \rightsquigarrow Q$ in \mathcal{D} . This was explained in detail in Construction 4.4, but we remind the reader that P_σ was obtained by surgery along $\sigma|_{\partial D^k \times \mathbb{R}^{2n-k}}$ and M_σ is the associated trace of the surgery. By Lemma 4.5 it suffices to prove Theorem 6.1 for elementary bordisms $M = M_\sigma$ of this special type. In particular, M has support in $\sigma(D^k \times D^{2n-k})$.

Let us consider D^k as a subspace of $D^{k-1} \times D^1$, writing coordinates as $(y, z) = (y_1, y_2, \dots, y_{k-1}, z)$, and write $\iota : D^{k-1} \rightarrow \partial D^k$ for the diffeomorphism onto the lower hemisphere which is inverse to the stereographic projection $(y_1, \dots, y_k) \mapsto \frac{1}{1-y_1}(y_2, \dots, y_k)$. For each $t \in (2, \infty)$ we have an embedding

$$\begin{aligned} \mu_t : D^{2n-k} \times D^k &\longrightarrow [0, 2] \times \partial D^k \times \mathbb{R}^{2n-k} \\ (x; y, z) &\longmapsto \left(\frac{3}{2}(1 - |x|^2)(1 + \frac{1}{3}z), \iota(y), t(1 + \frac{1}{3}z)x\right). \end{aligned}$$

Composing μ_t with $[0, 2] \times \sigma$ gives an embedding into $[0, 2] \times Q$, and this embedding when $t = 3 + 3i$, $i \in \mathbb{N}$, will be important for us. We write

$$\hat{\ell}_i^{\text{triv}} = (([0, 2] \times \sigma) \circ \mu_{3+3i})^* \hat{\ell}_Q$$

for the θ -structure which $\hat{\ell}_Q$ induces on $D^{2n-k} \times D^k$ via these embeddings, and

$$\partial\psi_i : \partial D^{2n-k} \times D^k \longrightarrow Q$$

for the restriction of $([0, 2] \times \sigma) \circ \mu_{3+3i}$ to $\partial D^{2n-k} \times D^k$.

The following is analogous to Construction 4.7, except that we now do surgery on the spheres $(\partial\psi_i)(\partial D^{2n-k} \times \{0\})$, which are *meridian* to $\sigma(\partial D^k \times \{0\})$, instead of the spheres $\sigma(\partial D^n \times \{3(i+1)\})$ which were *parallel* to $\sigma(\partial D^n \times \{0\})$.

Construction 6.2. For each $p \geq 0$ we will construct a pair of composable cobordisms $(1, U_p) : Q \rightsquigarrow R_p$ and $(1, V_p) : R_p \rightsquigarrow Q$ such that U_p has support outside of $\text{supp}(M)$, and $V_p \circ U_p$ is an isomorphism.

Let $(1, U_p) : Q \rightsquigarrow R_p$ be the cobordism obtained as the simultaneous (forward) trace of the surgeries along the disjoint embeddings

$$\partial\psi_i : \partial D^{2n-k} \times D^k \longrightarrow Q \quad i = 0, 1, \dots, p,$$

with θ -structure constructed using the extensions $\hat{\ell}_i^{\text{triv}}$ of $(\partial\psi_i)^*\hat{\ell}_Q$. This is analogous to the cobordism $V_p : Q \rightsquigarrow S_p$ in Construction 4.7, except that we are now doing surgery on several spheres meridian to the core of σ instead of several spheres parallel to the core of σ . It is possible to arrange that the support of U_p is disjoint from that of M , and we do so. We write $\psi_i : D^{2n-k} \times D^k \rightarrow U_p$ for the embedding of the i th handle.

By construction, there is an embedding $i : U_p \hookrightarrow [0, 2] \times Q$ relative to $\{0\} \times Q$, and a homotopy of θ -structures $i^*\hat{\ell}_Q \simeq \hat{\ell}_{U_p}$ relative to Q . Using the isotopy extension theorem, and the homotopy extension property for bundle maps, this gives a θ -cobordism $(1, V_p) : R_p \rightsquigarrow Q$ and a path from $V_p \circ U_p$ to $[0, 2] \times Q$ in $\mathcal{D}(Q, Q)$.

The composable cobordisms $M : P \rightsquigarrow Q$ and $U_p : Q \rightsquigarrow R_p$ have disjoint support by construction, so may be subjected to interchange of support, giving composable cobordisms $\mathcal{L}_M(U_p) : P \rightsquigarrow P'_p$ and $\mathcal{R}_{U_p}(M) : P'_p \rightsquigarrow R_p$.

Lemma 6.3. *The cobordism $V_p \circ \mathcal{R}_{U_p}(M) : P'_p \rightsquigarrow Q$ has a handle structure relative to Q having handles of index $(k-1)$ only.*

Proof. The cobordism V_p consists of $(p+1)$ handles of index $(k-1)$ relative to Q , and $\mathcal{R}_{U_p}(M)$ consists of a single handle of index k relative to R_p . We claim that the handle of $\mathcal{R}_{U_p}(M)$ may be cancelled against one of the handles of V_p , leaving a cobordism with p handles of index $(k-1)$. We shall explain the case $p=0$; the argument is the same in general, working only with the innermost handle.

As an abstract manifold R_0 is the result of doing surgery on Q along

$$\partial\psi_0 : \partial D^{2n-k} \times D^k \longrightarrow Q.$$

The image of $\partial\psi_0$ is disjoint from the set $\sigma(\partial D^k \times D^{2n-k}) \subset Q$ onto which the handle of M is attached, so we may consider $\sigma|_{\partial D^k \times D^{2n-k}}$ as the attaching map for the unique handle of $\mathcal{R}_{U_p}(M)$ as well. In Q , $\partial\psi_0$ is the embedding of a (thickened) meridian sphere of $\sigma(\partial D^k \times \{0\})$. The cobordism V_0 is the trace of a surgery along a $(2n-k)$ -sphere, which in terms of the surgered manifold

$$R_0 \approx (Q \setminus \partial\psi_0(\partial D^{2n-k} \times \text{int} D^k)) \cup_{\partial D^{2n-k} \times \partial D^k} (D^{2n-k} \times \partial D^k)$$

can be described as the union of the disc $D^{2n-k} \times \{*\}$ with the meridian disc $\sigma(\{*\} \times 2D^{2n-k})$ for some $*$ $\in \partial D^k$. In particular, it intersects $\sigma(\partial D^k \times \{0\})$ transversely in a single point, which shows that the handle of $\mathcal{R}_{U_p}(M)$ cancels the handle of V_0 , as required. \square

6.2. A semi-simplicial resolution. In this section we shall construct augmented semi-simplicial spaces $\mathcal{Z}_j(P)_\bullet \rightarrow \mathcal{F}_j(P)$, for any $P \in \mathcal{D}$ equipped with some auxiliary data. These semi-simplicial spaces will play roles to those in Section 4.3, although their definition is different.

Definition 6.4. Fix a $P \in \mathcal{D}$, a $W = (s, W) \in \mathcal{F}_j(P)$ for some $j \geq 0$, an embedding $\chi : \partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k}) \hookrightarrow P$ and a θ -structure $\hat{\ell}^{\text{std}}$ on $[0, 2] \times \partial D^k \times \mathbb{R}^{2n-k}$ which restricts to $\chi^*\hat{\ell}_P$ on $\{0\} \times \partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k})$.

Let $Z(W)_0 = Z(W, \chi, \hat{\ell}^{\text{std}})_0$ be the set of tuples (t, c, \hat{L}) consisting of a $t \in (2, \infty)$, an embedding $c : D^{2n-k} \times D^k \hookrightarrow W$, and a path of θ -structures $\hat{L} \in \text{Bun}^\theta(T(D^{2n-k} \times D^k))^I$ such that

- (i) there is a $\delta > 0$ such that $c(x, v) = \chi \circ \mu_t(\frac{x}{|x|}, v) + (1 - |x|) \cdot e_0$ for $1 - |x| < \delta$, where we have used that $\mu_t(\partial D^{2n-k} \times D^k) \subset \{0\} \times \partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k})$ to form $\chi \circ \mu_t$,
- (ii) the image of c is disjoint from $[0, s] \times L$, and $c^{-1}(P) = \partial D^{2n-k} \times D^k$,
- (iii) \hat{L} is a path from $c^*\hat{\ell}_W$ to $\mu_t^*\hat{\ell}^{\text{std}}$ which is constant over $\partial D^{2n-k} \times D^k$.

We topologise $Z(W)_0$ as a subspace of

$$\mathbb{R} \times \text{Emb}(D^{2n-k} \times D^k, [0, \infty) \times \mathbb{R}^\infty) \times \text{Bun}^\theta(T(D^{2n-k} \times D^k))^I.$$

Let $Z(W)_p = Z(W, \chi, \hat{\ell}^{\text{std}})_p \subset (Z(W, \chi, \hat{\ell}^{\text{std}})_0)^{p+1}$ be the subset consisting of tuples $(t_0, c_0, \hat{L}_0, t_1, c_1, \hat{L}_1, \dots, t_p, c_p, \hat{L}_p)$ such that

- (i) each (t_i, c_i, \hat{L}_i) lies in $Z(W)_0$,
- (ii) the c_i are disjoint,
- (iii) $t_0 < t_1 < \dots < t_p$.

We topologise $Z(W)_p$ as a subspace of $(Z(W, \chi, \hat{\ell}^{\text{std}})_0)^{p+1}$. The collection $Z(W)_\bullet$ has the structure of a semi-simplicial space, where the i th face map is given by forgetting (t_i, c_i, \hat{L}_i) .

We now combine all of the $Z(W)_\bullet$ into a single augmented semi-simplicial space.

Definition 6.5. Fix a $P \in \mathcal{D}$, a $j \geq 0$, an embedding $\chi : \partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k}) \hookrightarrow P$ and a θ -structure $\hat{\ell}^{\text{std}}$ on $[0, 2] \times \partial D^k \times \mathbb{R}^{2n-k}$ which restricts to $\chi^* \hat{\ell}_P$ on $\{0\} \times \partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k})$.

Let $\mathcal{Z}_j(P)_p = \mathcal{Z}_j(P, \chi, \hat{\ell}^{\text{std}})_p$ be the set of tuples $(s, W; x)$ with $(s, W) \in \mathcal{F}_j(P)$ and $x \in Z(W, \chi, \hat{\ell}^{\text{std}})_p$. Topologise this set as a subspace of

$$\mathcal{F}_j(P) \times (\mathbb{R} \times \text{Emb}(D^{2n-k} \times D^k, [0, \infty) \times \mathbb{R}^\infty) \times \text{Bun}^\theta(T(D^{2n-k} \times D^k))^I)^{p+1}.$$

The collection $\mathcal{Z}_j(P)_\bullet$ has the structure of a semi-simplicial space augmented over $\mathcal{F}_j(P)$, where the i th face maps forgets (t_i, c_i, \hat{L}_i) , and the augmentation map just remembers the underlying θ -manifold (s, W) .

The main result concerning these semi-simplicial spaces is the following, which is analogous to Theorem 4.12.

Theorem 6.6. *If $k > n$, then for any data $(\chi, \hat{\ell}^{\text{std}})$ as in Definition 6.5 the augmentation map $|\mathcal{Z}_j(P)_\bullet| \rightarrow \mathcal{F}_j(P)$ is a weak homotopy equivalence for all j .*

The proof of this theorem is rather easier than the corresponding Theorem 4.12, which is suggested by the fact that while that theorem only holds in the limit $j \rightarrow \infty$, this theorem holds for finite j . The reason is that Theorem 4.12 concerns n -dimensional submanifolds of a $2n$ -manifold, which were made disjoint using the infinite supply of $W_{1,1}$'s available in the limit, whereas Theorem 6.6 concerns $(2n - k)$ -dimensional submanifolds of a $2n$ -manifold, which can be made disjoint merely by general position since $2n - k < n$.

Proof of Theorem 6.6. Just as in Lemma 4.11, the map $|\mathcal{Z}_j(P)_\bullet| \rightarrow \mathcal{F}_j(P)$ is a quasifibration with fibre $|Z(W)_\bullet|$ over $(s, W) \in \mathcal{F}_j(P)$. Hence it will be enough to show that $|Z(W)_\bullet|$ is weakly contractible for each W .

Let $\overline{Z}(W)_\bullet$ be defined analogously to $Z(W)_\bullet$ with the exception that for a tuple $(t_0, c_0, \hat{L}_0, t_1, c_1, \hat{L}_1, \dots, t_p, c_p, \hat{L}_p)$ to span a p -simplex we require only the weaker condition

- (ii') the embeddings $c_i|_{D^{2n-k} \times \{0\}}$ are disjoint.

The inclusion $Z(W)_\bullet \rightarrow \overline{Z}(W)_\bullet$ is a levelwise weak homotopy equivalence by the argument of Lemma 5.3, so it is enough to show that $|\overline{Z}(W)_\bullet|$ is weakly contractible for each W . It is easy to verify that $\overline{Z}(W)_\bullet$ is a topological flag complex in the sense of [GRW14b, Definition 6.1], and we claim that it satisfies the conditions of [GRW14b, Theorem 6.2]. (We recall these conditions below.) This immediately implies the result.

Condition (i), which says that the augmentation map has local sections, is vacuous because $\overline{Z}(W)_\bullet$ is augmented only over a point.

Next we establish condition (ii): that the augmentation map is surjective, or in other words that $\overline{Z}(W)_0$ is not empty. Let $t \in (2, \infty)$ and consider the commutative diagram of bundle maps

$$\begin{array}{ccc} T(D^{2n-k} \times D^k)|_{\partial D^{2n-k} \times D^k} & \xrightarrow{D\chi \circ \mu_t} & TW|_P \xrightarrow{\cong} TW \\ \downarrow & \dashrightarrow^{\hat{c}} & \downarrow \hat{\ell}_W \\ T(D^{2n-k} \times D^k) & \xrightarrow{\mu_t^* \hat{\ell}^{\text{std}}} & \theta^* \gamma. \end{array}$$

As $\ell_W : W \rightarrow B$ is n -connected, and $(D^{2n-k}, \partial D^{2n-k}) \times D^k$ only has relative cells of dimension strictly less than n , there is a dashed bundle map \hat{c} making the top triangle commute, and the bottom triangle commute up to a homotopy of bundle maps which is constant over $\partial D^{2n-k} \times D^k$. By Smale–Hirsch theory, we may find an immersion $c : D^{2n-k} \times D^k \looparrowright W$ extending $\chi \circ \mu_t|_{\partial D^{2n-k} \times D^k} : \partial D^{2n-k} \times D^k \rightarrow P$ and with differential homotopic to \hat{c} relative to $\partial D^{2n-k} \times D^k$. By general position of the core $D^{2n-k} \times \{0\}$ and shrinking, this may be supposed to be an embedding, which then satisfies $c^* \hat{\ell}_W = \hat{\ell}_W \circ Dc \simeq \hat{\ell}_W \circ \hat{c} \simeq \mu_t^* \hat{\ell}^{\text{std}}$.

Finally we establish condition (iii): that any finite collection of vertices of $\overline{Z}(W)_\bullet$ each span a 1-simplex with some other common vertex. Let $\{(t_i, c_i, \hat{L}_i)\}_{j \in J}$ be a finite collection of 0-simplices, and choose (as in part (ii)) another (t, c, \hat{L}) having $t \ll t_i$. The embedding $c|_{D^{2n-k} \times \{0\}}$ may be perturbed relative to the boundary to make it disjoint from every $c_i|_{D^{2n-k} \times \{0\}}$, as these cores are $(2n-k)$ -dimensional and $(2n-k) + (2n-k) < 2n$ as we have supposed that $k > n$. After changing c in this way (using isotopy extension), the 0-simplex (t, c, \hat{L}) spans 1-simplex with each (t_i, c_i, \hat{L}_i) , as required. \square

6.3. Resolving composition with M . We now come to the proof of Theorem 6.1 proper. We have a morphism $(1, M) : P \rightsquigarrow Q \in \mathcal{D}$ which we have supposed is of the form M_σ for some $\sigma : \mathbb{R}^k \times \mathbb{R}^{2n-k} \hookrightarrow [0, \infty) \times \mathbb{R}^{\infty-1}$ and some extension of $(\sigma|_{\partial D^k \times D^{2n-k}})^* \hat{\ell}_Q$ to $T(D^k \times D^{2n-k})$. Gluing on M defines a map $- \circ M : \mathcal{F}(Q) \rightarrow \mathcal{F}(P)$. The restricted embedding $\sigma|_{\partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k})}$ has image outside of the support of M , so may be considered as an embedding into either Q or P . Let $\hat{\ell}^{\text{std}}$ be the θ -structure on $[0, 2] \times \partial D^k \times \mathbb{R}^{2n-k}$ given by

$$T([0, 2] \times \partial D^k \times \mathbb{R}^{2n-k}) = \varepsilon^1 \oplus T(\partial D^k \times \mathbb{R}^{2n-k}) \xrightarrow{\varepsilon^1 \oplus D\sigma} \varepsilon^1 \oplus TQ \xrightarrow{\hat{\ell}_Q} \theta^* \gamma_{2n}.$$

Taking the limit $j \rightarrow \infty$ in Definition 6.5 gives augmented semi-simplicial spaces

$$\begin{aligned} \mathcal{Z}(Q)_\bullet &= \mathcal{Z}(Q, \sigma|_{\partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k})}, \hat{\ell}^{\text{std}})_\bullet \longrightarrow \mathcal{F}(Q) \\ \mathcal{Z}(P)_\bullet &= \mathcal{Z}(P, \sigma|_{\partial D^k \times (\mathbb{R}^{2n-k} \setminus D^{2n-k})}, \hat{\ell}^{\text{std}})_\bullet \longrightarrow \mathcal{F}(P), \end{aligned}$$

both of which become homotopy equivalences after geometric realisation. We wish to cover $- \circ M : \mathcal{F}(Q) \rightarrow \mathcal{F}(P)$ by a semi-simplicial map $(- \circ M)_\bullet : \mathcal{Z}(Q)_\bullet \rightarrow \mathcal{Z}(P)_\bullet$, and can do this using the extrusion construction of Definition 4.13, via the formula

$$(W; t, c, \hat{L}) \longmapsto (W \circ M; t, \varepsilon_1(c), \varepsilon_1(\hat{L}))$$

on 0-simplices, and the analogous formula on higher simplices. This commutes with face maps, and defines a semi-simplicial map $(- \circ M)_\bullet : \mathcal{Z}(Q)_\bullet \rightarrow \mathcal{Z}(P)_\bullet$.

Proposition 6.7. *For each $p \geq 0$ the map $(- \circ M)_p : \mathcal{Z}(Q)_p \rightarrow \mathcal{Z}(P)_p$ is an abelian homology equivalence.*

Proof. The cobordism $U_p : Q \rightsquigarrow R_p$ provided by Construction 6.2 has embeddings $\psi_i : D^{2n-k} \times D^k \rightarrow U_p$ for $i = 0, 1, \dots, p$ extending $\partial\psi_i$, and $\psi_i^* \hat{\ell}_{U_p}$ is equal to

$\hat{\ell}_i^{\text{triv}} = \mu_{3+3i}^* \hat{\ell}^{\text{std}}$. Hence, letting $\hat{L}_i = \hat{\ell}_i^{\text{triv}}$ be the constant homotopy, we have a map

$$\begin{aligned} \mathcal{F}(R_p) &\longrightarrow \mathcal{Z}(Q)_p \\ X &\longmapsto (X \circ U_p; 3, \psi_0, \hat{L}_0, 6, \psi_1, \hat{L}_1, \dots, 3 + 3p, \psi_p, \hat{L}_p). \end{aligned}$$

This map is a weak homotopy equivalence, as may be proved in the same way as the corresponding step of Proposition 4.14.

If $\mathcal{L}_M(U_p) : P \rightsquigarrow P'_p$ and $\mathcal{R}_{U_p}(M) : P'_p \rightsquigarrow R_p$ are the cobordisms produced by the interchange of support, then $\mathcal{L}_M(U_p)$ also contains the handles ψ_i carrying the tangential structures $\hat{\ell}_i^{\text{triv}}$. Thus we may form the analogous map $\mathcal{F}(P'_p) \rightarrow \mathcal{Z}(P)_p$, which is a weak homotopy equivalence by the same argument. Now consider the diagram

$$(6.1) \quad \begin{array}{ccccc} \mathcal{F}(Q) & \xrightarrow{-\circ V_p} & \mathcal{F}(R_p) & \xrightarrow{\simeq} & \mathcal{Z}(Q)_p \\ & \searrow & \downarrow & & \downarrow \\ & & \mathcal{F}(P'_p) & \xrightarrow{\simeq} & \mathcal{Z}(P)_p \end{array} \begin{array}{l} \\ \\ \downarrow \\ \downarrow \end{array} \begin{array}{l} \\ \\ -\circ \mathcal{R}_{U_p}(M) \\ (-\circ M)_p \end{array}$$

The left triangle commutes by definition of the diagonal map. The cobordism V_p is obtained from Q by attaching $(k-1)$ -handles, so lies in \mathcal{W} by inductive hypothesis. By Lemma 6.3 the cobordism $V_p \circ \mathcal{R}_{U_p}(M)$ is obtained from Q by attaching $(k-1)$ -handles, so also lies in \mathcal{W} by inductive hypothesis. Hence by Lemma 2.19 the cobordism $\mathcal{R}_{U_p}(M)$ also lies in \mathcal{W} , so induces an isomorphism on homology with all abelian coefficient systems.

Hence if the bottom square commutes up to homotopy then we have proved this proposition. But this follows by the argument for the analogous step of Proposition 4.14. \square

This proposition shows that the semi-simplicial map $(-\circ M)_\bullet : \mathcal{Z}(Q)_\bullet \rightarrow \mathcal{Z}(P)_\bullet$ is a levelwise abelian homology equivalence, and as the augmentation maps for these semi-simplicial spaces both become weak homotopy equivalences after geometric realisation, the spectral sequence argument from the end of Section 4 shows that $-\circ M : \mathcal{F}(Q) \rightarrow \mathcal{F}(P)$ is an abelian homology equivalence too. Thus $M \in \mathcal{W}$, which finishes the proof of Theorem 6.1.

By the discussion in Section 2.5, this finishes the proof of Theorem 2.15.

7. STABLE HOMOLOGY AND GROUP-COMPLETION

In this section we shall explain how Theorem 1.3 may be combined with the ‘‘surgery on morphisms’’ part of [GRW14b] to re-prove and strengthen the main result of that paper.

For a pair of objects $A, B \in \mathcal{C}_\theta$, there is a map $\mathcal{C}_\theta(A, B) \rightarrow \Omega_{[A, B]} B\mathcal{C}_\theta$ to the space of paths in $B\mathcal{C}_\theta$ from the point represented by A to that represented by B . We shall consider the subspaces $\mathcal{N}_n^\theta(P) \subset \mathcal{C}_\theta(\emptyset, P)$ of Definition 2.18, and we shall establish a theorem which describes the effect of the map

$$\mathcal{N}_n^\theta(P) \hookrightarrow \mathcal{C}_\theta(\emptyset, P) \longrightarrow \Omega_{[\emptyset, P]} B\mathcal{C}_\theta$$

on homology, after suitable stabilisation. Our proof shall make use of the group-completion theorem applied to the category $\mathcal{D} = \mathcal{C}_{\theta, \partial L}^{n-1}$ and the weak equivalence $BD \simeq B\mathcal{C}_\theta$ from Theorem 2.8.

Definition 7.1. Analogously to Definition 2.13, we say that a composable sequence of cobordisms

$$K|_0 \xrightarrow{K|_{[0,1]}} K|_1 \xrightarrow{K|_{[1,2]}} K|_2 \xrightarrow{K|_{[2,3]}} K|_3 \rightsquigarrow \dots$$

in \mathcal{C}_θ is a θ -end in \mathcal{C}_θ if each $K|_{[i,i+1]}$ satisfies

- (i) it is $(n-1)$ -connected relative to both $K|_i$ and $K|_{i+1}$, and
- (ii) it contains an embedded copy of $W_{1,1}$ with standard θ -structure.

We shall often refer to such a θ -end in \mathcal{C}_θ by the non-compact θ -manifold $K \subset [0, \infty) \times (-1, 1)^\infty$ obtained by composing all of these cobordisms.

Remark 7.2. This definition of θ -end generalises the notion of a universal θ -end in [GRW14b, Definition 1.7]; in terms of the characterisation in [GRW14b, Addendum 1.9], it omits conditions (i) and (ii).

If K is a θ -end in \mathcal{C}_θ then there are induced maps

$$K|_{[i,i+1]} \circ - : \mathcal{N}_n^\theta(K|_i) \longrightarrow \mathcal{N}_n^\theta(K|_{i+1}),$$

as $W \cup_{K|_i} K|_{[i,j]}$ has n -connected structure map to B , by assumption (i) above. The following should be considered as a strengthened version of [GRW14b, Theorem 1.8].

Theorem 7.3. *Let K be a θ -end in \mathcal{C}_θ such that $\mathcal{N}_n^\theta(K|_0) \neq \emptyset$. Then the map*

$$\operatorname{hocolim}_{i \rightarrow \infty} \mathcal{N}_n^\theta(K|_i) \longrightarrow \operatorname{hocolim}_{i \rightarrow \infty} \Omega_{[\emptyset, K|_i]} BC_\theta$$

is acyclic.

Similarly to [GRW14b, Section 7.4], the proof of Theorem 7.3 will use the group-completion theorem for categories. In fact, the group-completion theorem will be used to prove a weaker preliminary result, which we state in Proposition 7.5 below; we shall deduce Theorem 7.3 from it by a further application of Theorem 2.15.

We first make the following crucial observation: if the tangential structure θ is such that $\mathcal{N}_n^\theta(K|_0)$ is non-empty for some $K|_0$, then there is an n -connected map $\ell_W : W \rightarrow B$ from a compact manifold. Cells of dimension at least $(n+1)$ can be attached to W in order to make this map a weak homotopy equivalence, and hence, by [Wal65, Theorem A], B satisfies Wall's finiteness condition (F_n) . We shall therefore assume this property throughout this section.

7.1. The group-completion argument. To state the preliminary result, let $L \subset (-1, 0] \times \mathbb{R}^{\infty-1}$ be such that $(L, \partial L)$ is $(n-1)$ -connected, and as in Definition 2.6 consider the category $\mathcal{D} = \mathcal{C}_{\theta, \partial L}^{n-1} \subset \mathcal{C}_{\theta, \partial L}$ consisting of those cobordisms which are $(n-1)$ -connected relative to their outgoing boundaries. As $(L, \partial L)$ is $(n-1)$ -connected, Remark 2.7 shows that the natural map $\mathcal{C}_{\theta, \partial L}^{n-1} \rightarrow \mathcal{C}_{\theta, L}^{n-1}$ is an isomorphism of categories.

Definition 7.4. For $A, B \in \mathcal{D}$, let $\mathcal{D}_n(A, B) \subset \mathcal{D}(A, B)$ be those path-components represented by cobordisms $W : A \rightsquigarrow B$ such that the structure map $\ell_W : W \rightarrow B$ is n -connected (recall that $W \in \mathcal{D}(A, B)$ has had $[0, 1] \times \operatorname{int}(L)$ cut out).

The notation $\mathcal{D}_n(A, B)$ should not be taken to imply that this defines a subcategory \mathcal{D}_n of \mathcal{D} : it does not; but $\mathcal{D}_n(A, -)$ is a subfunctor of $\mathcal{D}(A, -)$. The following is our preliminary version of Theorem 7.3.

Proposition 7.5. *Suppose that B satisfies Wall's finiteness condition (F_n) , and let K' be a θ -end in $\mathcal{C}_{\theta, \partial L}$ such that $\ell_{K'|_0}$ is $(n-1)$ -connected. Then the map*

$$\operatorname{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(\overline{L}, K'|_i) \longrightarrow \operatorname{hocolim}_{i \rightarrow \infty} \Omega_{[L, K'|_i]}^\infty B\mathcal{D}$$

is acyclic.

Lemma 7.6. *Suppose that B satisfies Wall's finiteness condition (F_n) , and let K' be a θ -end in $\mathcal{C}_{\theta, \partial L}$ such that $\ell_{K'|_0}$ is $(n-1)$ -connected. Then there is another θ -end K'' in $\mathcal{C}_{\theta, \partial L}$ such that $K''|_0 = K'|_0$ and for all i the structure map $\ell_{K''|_{[i, \infty)}} : K''|_{[i, \infty)} \rightarrow B$ is n -connected.*

Proof. It is enough to show that for any $P \in \mathcal{C}_{\theta, \partial L}$ such that ℓ_P is $(n-1)$ -connected there is a morphism $W_P : P \rightsquigarrow P'$ in $\mathcal{C}_{\theta, \partial L}$ such that

- (i) $\ell_{P'}$ is $(n-1)$ -connected,
- (ii) W_P is $(n-1)$ -connected relative to either end,
- (iii) W_P contains an embedded copy of $W_{1,1}$ with standard θ -structure,
- (iv) $\ell_{W_P} : W_P \rightarrow B$ is n -connected.

Then we can take K'' to be given by $W_{K'|_0}, W_{(K'|_0)', W_{(K'|_0)''}, \dots$. Let us first suppose that $n \geq 3$, and we shall explain the necessary changes for small n later.

If we write $\pi = \pi_1(B)$ and $P^{(n-1)}$ for an $(n-1)$ -skeleton of P , then it follows from the definition of Wall's condition (F_n) that $\pi_n(B, P^{(n-1)})$ is a finitely-generated $\mathbb{Z}[\pi]$ -module, and so from the long exact sequence of the triple $(B, P, P^{(n-1)})$ that $\pi_n(B, P)$ is finitely-generated too. Thus by the long exact sequence for the pair (B, P) the $\mathbb{Z}[\pi]$ -module

$$\text{Ker}((\ell_P)_* : \pi_{n-1}(P) \rightarrow \pi_{n-1}(B))$$

is also finitely generated. We can represent $\mathbb{Z}[\pi]$ -module generators of this kernel by finitely many maps

$$\alpha_1, \alpha_2, \dots, \alpha_k : S^{n-1} \rightarrow P,$$

and as P has dimension $(2n-1)$ we may suppose that these maps are disjoint embeddings. Because they become nullhomotopic in B , and the tangent bundle of P is pulled back from B , the embeddings α_i have stably trivial normal bundles, but as their normal bundles are of dimension $n > \dim(S^{n-1})$ they must in fact be unstably trivial. Thus we may upgrade the α_i to disjoint embeddings

$$\hat{\alpha}_1, \hat{\alpha}_2, \dots, \hat{\alpha}_k : S^{n-1} \times D^n \hookrightarrow \text{int}(P).$$

The trace of the (simultaneous) surgeries along these maps gives a cobordism $W' : P \rightsquigarrow P''$ in \mathcal{D} . Consider the diagram

$$\begin{array}{ccc} \pi_{n-1}(P) & \longrightarrow & \pi_{n-1}(W') \\ & \searrow & \downarrow \\ & & \pi_{n-1}(B). \end{array}$$

As the diagonal map is surjective, the kernel of the diagonal map is contained in the kernel of the horizontal map, and the horizontal map is surjective, it follows that the vertical map is an isomorphism.

Now $\pi_n(B, W')$ is a quotient of $\pi_n(B, P)$ (by the long exact sequence of the triple (B, W', P) and the fact that $\pi_{n-1}(W', P) = 0$) and so a finitely-generated $\mathbb{Z}[\pi]$ -module. The exact sequence

$$\dots \rightarrow \pi_n(W') \rightarrow \pi_n(B) \rightarrow \pi_n(B, W') \rightarrow 0$$

shows that $\pi_n(B)$ is generated as a $\mathbb{Z}[\pi]$ -module by the image of $\pi_n(W')$ along with finitely many additional elements, $\beta_1, \beta_2, \dots, \beta_l : S^n \rightarrow B$. For each β_i , the map $S^n \xrightarrow{\beta_i} B \xrightarrow{\theta} BO(2n)$ may be lifted to $\hat{\beta}_i : S^n \rightarrow BO(n)$, and the disc bundle $D^n \rightarrow E_i \xrightarrow{\pi} S^n$ classified by $\hat{\beta}_i$ has tangent bundle classified by $E_i \xrightarrow{\pi} S^n \xrightarrow{\beta_i} B \xrightarrow{\theta} BO(2n)$, and so is endowed with a θ -structure by the lift $\beta_i \circ \pi$. The map $\ell_{E_i} : E_i \xrightarrow{\pi} S^n \xrightarrow{\beta_i} B$ hits $[\beta_i] \in \pi_n(B)$. Let $W_P : P \rightsquigarrow P'$ be the manifold obtained from $W' : P \rightsquigarrow P''$ by forming the boundary connect-sum at P'' with the θ -manifolds E_1, E_2, \dots, E_l , as well as with a copy of $W_{1,1}$.

By construction $\pi_n(\ell_{W_P})$ is surjective, as its image contains the image of $\pi_n(\ell_{W'})$ and the $[\beta_i]$ and these generate $\pi_n(B)$. Furthermore, $\pi_{n-1}(W') \rightarrow \pi_{n-1}(W_P)$ is

an isomorphism, as homotopically W_P is obtained by wedging n -spheres on to W' . Thus properties (ii)–(iv) hold. For property (i) note that the composition

$$P' \hookrightarrow W_P \xrightarrow{\ell_{W_P}} B$$

is $(n-1)$ -connected as the first map is $(n-1)$ -connected and the second is n -connected.

Let us now explain the necessary changes for small n . If $n = 2$ the first step should be interpreted as saying that $\text{Ker}((\ell_P)_* : \pi_1(P) \rightarrow \pi_1(B))$ is normally finitely generated as a subgroup of $\pi_1(P)$. This no longer follows from Wall's condition (F_2) , but is instead a standard exercise in group theory: a normal subgroup $N \triangleleft G$ of a finitely presented group is normally finitely generated if and only if G/N is finitely presented. The technique described can then be used to kill the finitely many normal generators of this kernel, giving $W' : P \rightsquigarrow P''$, and the same argument shows that $\pi_1(\ell_{W'})$ is an isomorphism. Now we apply Wall's (F_2) to the map $(W')^{(2)} \rightarrow W' \rightarrow B$, and hence deduce that $\pi_2(B, W')$ is a finitely generated $\mathbb{Z}[\pi]$ -module. The argument is then concluded as above.

If $n = 1$ the first step should be interpreted as saying that $\pi_0(P)$ is a finite set, which is true as P is compact. We can then perform finitely many 0-surgeries on it to obtain a connected cobordism $W' : P \rightsquigarrow P''$. As B satisfies Wall's (F_1) the group π is finitely generated, so in particular is generated by the image of $\pi_1(W')$ and finitely many additional elements, $\beta_1, \beta_2, \dots, \beta_l : S^1 \rightarrow B$. Performing the construction described on these gives a cobordism having the required properties. \square

Lemma 7.7. *Suppose that B satisfies Wall's finiteness condition (F_n) , and let K'' be a θ -end in $\mathcal{C}_{\theta, \partial L}$ such that the structure map $\ell_{K''|_{[i, \infty)}} : K''|_{[i, \infty)} \rightarrow B$ is n -connected for all i . Then for each $X \in \mathcal{D}$, the inclusion*

$$\text{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(X, K''|_i) \hookrightarrow \text{hocolim}_{i \rightarrow \infty} \mathcal{D}(X, K''|_i)$$

is a weak equivalence.

Proof. Before taking homotopy colimits, the map of direct systems is a levelwise inclusion of a collection of path components, so it certainly induces an injection on π_0 on homotopy colimits. We shall show that it also induces a surjection on π_0 on homotopy colimits, from which it follows that it is a weak homotopy equivalence by considering the induced map on homotopy groups with all basepoints.

To show that it is surjective on π_0 , let $W : X \rightsquigarrow K''|_i \in \mathcal{D}$. The commutative square

$$\begin{array}{ccc} K''|_i & \xrightarrow{(n-1)\text{-connected}} & K''|_{[i, \infty)} \\ \downarrow (n-1)\text{-connected} & & \downarrow n\text{-connected} \\ W & \xrightarrow{\ell_W} & B \end{array}$$

shows that ℓ_W is $(n-1)$ -connected, and we must glue on some $K''|_{[i, k]}$ to W in order to make it n -connected. The proof of this is similar to that of Lemma 7.6, and uses all of the same techniques. We give it here for $n \geq 3$, leaving the necessary changes for $n \leq 2$ to the reader.

Write $\pi = \pi_1(B)$. As B is assumed to satisfy the finiteness condition (F_n) and W is a finite CW-complex, as in the proof of Lemma 7.6 we deduce that

$$\text{Ker}((\ell_W)_* : \pi_{n-1}(W) \rightarrow \pi_{n-1}(B))$$

is also a finitely-generated $\mathbb{Z}[\pi]$ -module. As $\pi_{n-1}(K''|_i) \rightarrow \pi_{n-1}(W)$ is onto, we can represent $\mathbb{Z}[\pi]$ -module generators of this kernel by finitely many maps $\alpha_1, \alpha_2, \dots, \alpha_k : S^{n-1} \rightarrow K''|_i$, and because these become nullhomotopic in B , they must also be nullhomotopic in $K''|_{[i, \infty)}$, and hence must in fact be nullhomotopic

in $K''|_{[i,j]}$ for some $i \ll j$. Let $W' = W \cup_{K''|_i} K''|_{[i,j]}$. As in the proof of Lemma 7.6 it follows that $\ell_{W'} : W' \rightarrow B$ is an isomorphism on π_{n-1} .

Wall's condition (F_n) still implies that $\pi_n(B, W')$ is a finitely-generated $\mathbb{Z}[\pi]$ -module. The exact sequence

$$\cdots \rightarrow \pi_n(W') \rightarrow \pi_n(B) \rightarrow \pi_n(B, W') \rightarrow 0$$

shows that $\pi_n(B)$ is generated as a $\mathbb{Z}[\pi]$ -module by the image of $\pi_n(W')$ along with finitely many additional elements, $\beta_1, \beta_2, \dots, \beta_l : S^n \rightarrow B$. As the map $K''|_{[j,\infty)} \rightarrow B$ is n -connected the β 's can be lifted to $K''|_{[j,\infty)}$, and hence to $K''|_{[j,k]}$ for some $j \ll k$.

Letting $W'' = W' \cup_{K''|_j} K''|_{[j,k]}$, the map $(\ell_{W''})_* : \pi_n(W'') \rightarrow \pi_n(B)$ is surjective, as its image contains both the image of $(\ell_{W'})_*$ and the elements $[\beta_i]$. The map $\ell_{W''}$ still induces an isomorphism on lower homotopy groups, so is n -connected: hence $W'' \in \mathcal{D}_n(X, K''|_k)$, as required. \square

Proof of Proposition 7.5. We shall use a version of the group-completion theorem for categories, specifically the version given as Theorem A.14 in the appendix. For assumption (i) of that theorem we require that the combined source/target map is a Serre fibration, which follows from the isotopy extension theorem (for deforming the underlying manifolds) and the homotopy extension property for the restriction map $\text{Bun}^\theta(TW) \rightarrow \text{Bun}^\theta(TW|_{\partial W})$ (for deforming θ -structures). Let K'' be a θ -end in $\mathcal{C}_{\theta, \partial L}$ provided by Lemma 7.6. In order to apply Theorem A.14 we must show that any morphism in \mathcal{D} is sent to an abelian homology equivalence by applying

$$\text{hocolim}_{i \rightarrow \infty} \mathcal{D}(-, K''|_i) : \mathcal{D}^{\text{op}} \rightarrow \text{Top}.$$

But by Lemma 7.7 we may replace this by $\text{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(-, K''|_i)$, and it is the content of Theorem 2.15 that this sends every morphism to an abelian homology equivalence. Theorem A.14 then allows us to conclude that for any object $X \in \mathcal{D}$,

$$(7.1) \quad \text{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(X, K''|_i) \rightarrow \text{hocolim}_{i \rightarrow \infty} \Omega_{[X, K''|_i]} B\mathcal{D}$$

is an acyclic map.

We must now pass from this result to the corresponding statement for K' . Inductively choose a sequence of composable cobordisms

$$\cdots \rightsquigarrow \bar{L}|_{-2} \xrightarrow{\bar{L}|_{[-2,-1]}} \bar{L}|_{-1} \xrightarrow{\bar{L}|_{[-1,0]}} \bar{L}|_0 = \bar{L}$$

by letting $\bar{L}|_{[-i-1,-i]}$ be obtained from $[0, 1] \times \bar{L}|_{-i}$ by forming the boundary connect-sum with $W_{1,1}$ at $\{0\} \times \bar{L}|_{-i}$, and consider the diagram

$$\begin{array}{ccc} \text{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(\bar{L}, K''|_i) & \longrightarrow & \text{hocolim}_{j \rightarrow \infty} \text{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(\bar{L}_{-j}, K''|_i) \\ & & \uparrow \\ & & \text{hocolim}_{j \rightarrow \infty} \mathcal{D}_n(\bar{L}_{-j}, K'|_0) \\ & & \downarrow \\ \text{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(\bar{L}, K'|_i) & \longrightarrow & \text{hocolim}_{j \rightarrow \infty} \text{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(\bar{L}_{-j}, K'|_i). \end{array}$$

Each of these four maps is an abelian homology equivalence: the horizontal ones by Theorem 2.15, and the vertical ones by the analogue of Theorem 2.15 in which

stabilisation is formed on the left. Comparing this diagram with the analogous one of homotopy colimits of path spaces of $B\mathcal{D}$, (7.1) implies that

$$(7.2) \quad \operatorname{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(\bar{L}, K'|_i) \longrightarrow \operatorname{hocolim}_{i \rightarrow \infty} \Omega_{[\bar{L}, K'|_i]} B\mathcal{D}$$

is an abelian homology equivalence, and hence acyclic. \square

7.2. Proof of Theorem 7.3. Let K be a θ -end in \mathcal{C}_θ , pick a self-indexing Morse function $f : K|_0 \rightarrow [0, 2n - 1]$, and let $L = f^{-1}([0, n - \frac{1}{2}])$ with induced θ -structure $\hat{\ell}_L$. By construction, L has a handle structure with handles of index at most $(n - 1)$; equivalently, it may be obtained from ∂L by attaching handles of index at least n . After moving $K|_0 \subset \mathbb{R}^\infty$ by an isotopy, we may suppose that $L = K|_0 \cap ((-\infty, 0] \times \mathbb{R}^{\infty-1})$ as θ -manifolds: $K|_0$ is then an object of $\mathcal{C}_{\theta, L}$, giving an object $K|_0^\circ$ of $\mathcal{C}_{\theta, \partial L}$.

Lemma 7.8. *The proper embedding $K \hookrightarrow [0, \infty) \times (-1, 1)^\infty$ may be changed by an isotopy, and the bundle map $\hat{\ell}_K$ changed by a homotopy, relative to $K|_0$, so that $(K, \hat{\ell}_K)$ remains a θ -end in \mathcal{C}_θ and so that*

$$K \cap ([0, \infty) \times (-\infty, 0] \times \mathbb{R}^{\infty-1}) = [0, \infty) \times L$$

as θ -manifolds.

Proof. Supposing that $K|_i \cap ((-\infty, 0] \times \mathbb{R}^{\infty-1}) = L$ as θ -manifolds, we shall show that $[i, i + 1] \times L$ may be embedded into $K|_{[i, i+1]}$ relative to $\{i\} \times L$ such that $\{i + 1\} \times L$ lies in $K|_{i+1}$. The claim will then follow, by the isotopy extension theorem and induction.

As L is a $(2n - 1)$ -manifold with a handle structure only having $(n - 1)$ -handles and smaller, any embedding of a k -sphere into $K|_i$ can be isotoped off of L as long as $k \leq n - 1$.

Let us first suppose that $2n \geq 6$. As $(K|_{[i, i+1]}, K|_{i+1})$ is $(n - 1)$ -connected, by handle trading as in the proof of the s -cobordism theorem ([Ker65]) we can find a handle structure on $K|_{[i, i+1]}$ relative to $K|_i$ having only n -handles and lower: these are attached along spheres of dimension at most $n - 1$, so their attaching maps can be isotoped off of L , and hence $K|_{[i, i+1]}$ contains an embedded $[i, i + 1] \times L$. The same argument applies for $2n = 2$.

If $2n = 4$ we must modify the argument slightly. The above goes through after perhaps changing $K|_{[i, i+1]}$ by connect-sum with finitely many copies of $S^2 \times S^2$ so that handle-trading becomes available (just as in the proof of Lemma 2.20). As K contains countably many $S^2 \times S^2$ -summands by property (ii) of θ -ends in \mathcal{C}_θ , we may realise this by sliding in enough such summands down from $K|_{[i+1, \infty)}$. \square

Once K has been prepared using this lemma, we may form $K|_{[i, i+1]}^\circ$, giving a θ -end in $\mathcal{C}_{\theta, \partial L}$. Let us write

$$\mathcal{C}_{\theta, \partial L, n}(A, B) \subset \mathcal{C}_{\theta, \partial L}(A, B)$$

for the subspace of those cobordisms $(W, \hat{\ell}_W)$ such that $\ell_W : W \rightarrow B$ is n -connected. The weak homotopy equivalence

$$\mathcal{C}_{\theta, \partial L}(\bar{L}, K|_i^\circ) \xrightarrow{-\cup L} \mathcal{C}_\theta(D(L), K|_i) \xrightarrow{-\circ V_L} \mathcal{C}_\theta(\emptyset, K|_i)$$

of Lemma 2.17 thus restricts to a map

$$(7.3) \quad \mathcal{C}_{\theta, \partial L, n}(\bar{L}, K|_i^\circ) \longrightarrow \mathcal{N}_n^\theta(K|_i)$$

which is a weak homotopy equivalence.

By assumption $\mathcal{N}_n^\theta(K|_0) \simeq \mathcal{C}_{\theta, \partial L, n}(\bar{L}, K|_0^\circ)$ is non-empty, so let $W^\circ : \bar{L} \rightsquigarrow K|_0^\circ$ lie in $\mathcal{C}_{\theta, \partial L, n}(\bar{L}, K|_0^\circ)$, choose a self-indexing Morse function $f : W \rightarrow [0, 2n]$, and let $V_0^\circ = f^{-1}([n - \frac{1}{2}, 2n]) : P_0^\circ \rightsquigarrow K|_0^\circ$. By construction, V_0° has a handle structure

relative to P_0° having all handles of index at least n . The cobordism $K|_{[0,1]}^\circ \circ V_0^\circ : P_0^\circ \rightsquigarrow K|_1^\circ$ is thus $(n-1)$ -connected relative to P_0° and contains an embedded $W_{1,1}$ with standard θ -structure, so after changing it by a path in $\mathcal{C}_{\theta,\partial L}(P_0^\circ, K|_1^\circ)$ we may factorise it as

$$P_0^\circ \xrightarrow{P_0^\circ H} P_1^\circ \xrightarrow{V_1^\circ} K|_1^\circ$$

for some cobordism V_1° , which will be $(n-1)$ -connected relative to P_1° . Continuing in this way, we obtain a homotopy commutative diagram

$$(7.4) \quad \begin{array}{ccccccc} P_0^\circ & \xrightarrow{P_0^\circ H} & P_1^\circ & \xrightarrow{P_1^\circ H} & P_2^\circ & \xrightarrow{P_2^\circ H} & \dots \\ \downarrow V_0^\circ & & \downarrow V_1^\circ & & \downarrow V_2^\circ & & \\ K|_0^\circ & \xrightarrow{K|_{[0,1]}^\circ} & K|_1^\circ & \xrightarrow{K|_{[1,2]}^\circ} & K|_2^\circ & \xrightarrow{K|_{[2,3]}^\circ} & \dots \end{array}$$

in the category $\mathcal{C}_{\theta,\partial L}$.

The composition $\ell_{P_0^\circ} : P_0^\circ \hookrightarrow W^\circ \xrightarrow{\ell_W} B$ has the second map n -connected by assumption, and W° is obtained from P_0° by attaching cells of dimension n and higher so $P_0^\circ \hookrightarrow W^\circ$ is $(n-1)$ -connected, and hence $\ell_{P_0^\circ}$ is $(n-1)$ -connected. The inclusions $P_i^\circ \hookrightarrow P_i^\circ H \hookrightarrow P_{i+1}^\circ$ are both $(n-1)$ -connected, so it follows that all $\ell_{P_i^\circ}$ are $(n-1)$ -connected.

Applying $\mathcal{C}_{\theta,\partial L,n}(\overline{L}, -)$ to (7.4), we obtain a homotopy commutative diagram of spaces, and choosing homotopies making each square commute gives a map

$$(7.5) \quad \operatorname{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta,\partial L,n}(\overline{L}, P_i^\circ) \longrightarrow \operatorname{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta,\partial L,n}(\overline{L}, K|_i^\circ).$$

It follows from Lemma A.10 that the property of this map being an abelian homology equivalence is independent of this choice of homotopies.

Lemma 7.9. *The map (7.5) is an abelian homology equivalence.*

Proof. Similarly to the proof of Proposition 7.5 we consider the commutative square

$$\begin{array}{ccc} \operatorname{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta,\partial L,n}(\overline{L}, P_i^\circ) & \longrightarrow & \operatorname{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta,\partial L,n}(\overline{L}, K|_i^\circ) \\ \downarrow & & \downarrow \\ \operatorname{hocolim}_{j \rightarrow \infty} \operatorname{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta,\partial L,n}(\overline{L}_{-j}, P_i^\circ) & \longrightarrow & \operatorname{hocolim}_{j \rightarrow \infty} \operatorname{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta,\partial L,n}(\overline{L}_{-j}, K|_i^\circ) \end{array}$$

where $\dots \rightsquigarrow \overline{L}_{-2} \xrightarrow{\overline{L}|_{[-2,-1]}} \overline{L}_{-1} \xrightarrow{\overline{L}|_{[-1,0]}} \overline{L}_0 = \overline{L}$ is a sequence of cobordisms formed by letting $\overline{L}|_{[-i-1,-i]}$ be obtained from $[0,1] \times \overline{L}_{-i}$ by boundary connect-sum with $W_{1,1}$ at $\{0\} \times \overline{L}_{-i}$.

The claim then follows as the vertical maps are abelian homology equivalences by Theorem 2.15, and the lower map is an abelian homology equivalence by the analogue of Theorem 2.15 in which stabilisation is formed on the left. \square

If $W^\circ \in \mathcal{C}_{\theta,\partial L,n}(\overline{L}, P_i^\circ)$ then the map ℓ_{W° is n -connected and $\ell_{P_i^\circ}$ is $(n-1)$ -connected. Thus (W°, P_i°) is $(n-1)$ -connected, and so $W^\circ \in \mathcal{D}_n(\overline{L}, P_i^\circ)$, and hence $\mathcal{C}_{\theta,\partial L,n}(\overline{L}, P_i^\circ) = \mathcal{D}_n(\overline{L}, P_i^\circ)$. The argument is now completed using the

commutative diagram

$$\begin{array}{ccccc}
 \mathrm{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta, \partial L, n}(\bar{L}, P_i^\circ) & \longrightarrow & \mathrm{hocolim}_{i \rightarrow \infty} \mathcal{C}_{\theta, \partial L, n}(\bar{L}, K|_i^\circ) & \xrightarrow{\simeq} & \mathrm{hocolim}_{i \rightarrow \infty} \mathcal{N}_n^\theta(K|_i) \\
 \parallel & & \downarrow & & \downarrow \\
 \mathrm{hocolim}_{i \rightarrow \infty} \mathcal{D}_n(\bar{L}, P_i^\circ) & & & & \\
 \downarrow & & \downarrow & & \downarrow \\
 \mathrm{hocolim}_{i \rightarrow \infty} \Omega_{[\bar{L}, P_i^\circ]} BC_{\theta, \partial L} & \xrightarrow{\simeq} & \mathrm{hocolim}_{i \rightarrow \infty} \Omega_{[\bar{L}, K|_i^\circ]} BC_{\theta, \partial L} & \xrightarrow{\simeq} & \mathrm{hocolim}_{i \rightarrow \infty} \Omega_{[\emptyset, K|_i]} BC_\theta.
 \end{array}$$

Consider the horizontal maps: the left hand top map is the abelian homology equivalence (7.5); the right hand top map is given levelwise by the weak homotopy equivalences (7.3); the left hand bottom map is given levelwise by concatenation of paths, so is a weak homotopy equivalence; the right hand bottom map is given by the weak homotopy equivalence $BC_{\theta, \partial L} \cong BC_{\theta, L} \rightarrow BC_\theta$ and concatenation of paths. Finally, the left hand vertical map is an abelian homology equivalence by Proposition 7.5 and the weak homotopy equivalence $BD \rightarrow BC_{\theta, \partial L}$ from Theorem 2.8. Thus, it is at this point that we rely on a major technical result of [GRW14b]. It follows that the right hand vertical map is also an abelian homology equivalence, which finishes the proof of Theorem 7.3.

Finally, we remind the reader of the theorem of Galatius–Madsen–Tillmann–Weiss [GMTW09] concerning the homotopy type of BC_θ : there is a weak homotopy equivalence

$$BC_\theta \simeq \Omega^{\infty-1} \mathrm{MT}\theta,$$

defined using a parametrised Pontryagin–Thom construction. This allows us to translate Theorem 7.3 into the form given in Theorem 1.5.

8. FINITE GENUS AND STABILITY FOR CLOSED MANIFOLDS

In this section we shall combine Theorems 1.3 and 1.5 with our results from [GRW14a] to prove Corollaries 1.7 and 1.8.

8.1. Finite genus and homological stability. By Remark 2.11, the definition of θ -genus given in Section 1.3 is equivalent to

$$g^\theta(W) = \max \left\{ g \in \mathbb{N} \mid \begin{array}{l} \text{there are } g \text{ disjoint copies of } W_{1,1} \text{ in } W, \\ \text{each with standard } \theta\text{-structure} \end{array} \right\},$$

and if W has non-empty boundary then we have defined the *stable θ -genus* to be

$$\bar{g}^\theta(W) = \max \{ g^\theta(W \natural k(W_{1,1})) - k \mid k \in \mathbb{N} \}.$$

If W is a closed θ -manifold, we now define $\bar{g}^\theta(W)$ to be $\bar{g}^\theta(W \setminus D^{2n})$.

Definition 8.1. A *graded space* is a pair (X, h_X) of a space X and a continuous map $h_X : X \rightarrow \mathbb{Z}$. We write $X_{h_X=n} = h_X^{-1}(n)$ for the subspace of degree n . A (degree zero) map $f : (X, h_X) \rightarrow (Y, h_Y)$ of graded spaces is a continuous map $f : X \rightarrow Y$ such that $h_Y \circ f = h_X$. Similarly, a degree k map is a continuous map $f : X \rightarrow Y$ such that $h_Y \circ f = h_X + k$. The homology of a graded space (X, h_X) with any system of local coefficients \mathcal{L} on X acquires an extra grading

$$H_i(X; \mathcal{L}) = \bigoplus_{n \geq 0} H_i(X; \mathcal{L})_{h_X=n},$$

where $H_i(X; \mathcal{L})_{h_X=n} = H_i(X_{h_X=n}; \mathcal{L})$. Maps of graded spaces respect this additional grading (a degree k map induces a map with a shift of k).

We may use the function $\bar{g}^\theta : \mathcal{N}_n^\theta(P) \rightarrow \mathbb{Z}$ to grade the space $\mathcal{N}_n^\theta(P)$. Recall that in Section 3 we have defined for each θ -manifold P a bordism ${}_P H : P \rightarrow P'$ whose outgoing boundary P' is diffeomorphic to P but possibly with a different θ -structure. The map $- \circ {}_P H : \mathcal{N}_n^\theta(P) \rightarrow \mathcal{N}_n^\theta(P')$ then has degree 1 with respect to this new grading, and by [GRW14a, Theorem 1.4] we have the following.

Theorem 8.2. *If $2n \geq 6$, B is simply-connected, and \mathcal{L} is an abelian local coefficient system on $\mathcal{N}_n^\theta(P')$, then the map*

$$(- \circ {}_P H)_* : H_i(\mathcal{N}_n^\theta(P); (- \circ {}_P H)^* \mathcal{L})_{\bar{g}^\theta = g} \longrightarrow H_i(\mathcal{N}_n^\theta(P'); \mathcal{L})_{\bar{g}^\theta = g+1}$$

is an epimorphism if $3i \leq g - 1$, and an isomorphism if $3i \leq g - 4$ (if \mathcal{L} is constant and θ is spherical, it is an epimorphism if $2i \leq g - 1$, and an isomorphism if $2i \leq g - 3$).

8.2. Proof of Corollary 1.7 for $P \neq \emptyset$. We wish to combine Theorem 8.2 with Theorem 1.3 in order to deduce that for $M : P \rightsquigarrow Q$ a cobordism which is $(n - 1)$ -connected relative to Q and has $P \neq \emptyset$, the map

$$- \circ M : \mathcal{N}_n^\theta(Q) \longrightarrow \mathcal{N}_n^\theta(P)$$

induces an isomorphism in homology in a particular range of homological degrees. However, it is easy to see that $\bar{g}^\theta(- \circ M) - \bar{g}^\theta(-)$ need not be constant, so we shall need another grading of $\mathcal{N}_n^\theta(P)$ for this to be a graded map.

Lemma 8.3. *Let $M : P \rightsquigarrow Q$ be a θ -cobordism which is $(n - 1)$ -connected relative to Q . Then there is a function*

$$\bar{g}_M^\theta : \pi_0(\mathcal{N}_n^\theta(P)) \longrightarrow \mathbb{Z}$$

such that $\bar{g}_M^\theta(- \circ M) = \bar{g}^\theta(-)$, $\bar{g}_M^\theta(- \circ {}_P H) = \bar{g}_M^\theta(-) + 1$, and $\bar{g}_M^\theta(-) \leq \bar{g}^\theta(-)$.

Proof. Let us first suppose that $P \neq \emptyset$. Then Theorem 1.3 applies, and the map

$$- \circ M : \operatorname{hocolim}_{g \rightarrow \infty} \mathcal{N}_n^\theta(Q, \hat{\ell}_Q^{(g)}) \longrightarrow \operatorname{hocolim}_{g \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(g)})$$

is a homology equivalence so in particular a bijection on π_0 . Furthermore, as $\bar{g}^\theta(W \natural W_{1,1}) = \bar{g}^\theta(W) + 1$, there is an induced function

$$\bar{g}^\theta : \operatorname{colim}_{g \rightarrow \infty} \pi_0(\mathcal{N}_n^\theta(Q, \hat{\ell}_Q^{(g)})) \longrightarrow \operatorname{colim}_{+1} \mathbb{Z} \cong \mathbb{Z}.$$

We may therefore define \bar{g}_M^θ by

$$\bar{g}_M^\theta : \pi_0(\mathcal{N}_n^\theta(P)) \longrightarrow \operatorname{colim}_{g \rightarrow \infty} \pi_0(\mathcal{N}_n^\theta(P, \hat{\ell}_P^{(g)})) \xleftarrow{\sim} \operatorname{colim}_{g \rightarrow \infty} \pi_0(\mathcal{N}_n^\theta(Q, \hat{\ell}_Q^{(g)})) \xrightarrow{\bar{g}^\theta} \mathbb{Z},$$

which satisfies $\bar{g}_M^\theta(W \circ M) = \bar{g}^\theta(W)$ by construction. That $\bar{g}_M^\theta(- \circ {}_P H) = \bar{g}_M^\theta(-) + 1$ follows immediately from the same property of $\bar{g}^\theta(-)$. By the definition of \bar{g}_M^θ , when $W \circ k({}_P H) = W' \circ M$ for some W' then we have

$$\bar{g}_M^\theta(W) = \bar{g}^\theta(W') - k \leq \bar{g}^\theta(W' \circ M) - k = \bar{g}^\theta(W \circ k({}_P H)) - k = \bar{g}^\theta(W).$$

If $P = \emptyset$ then we factor M as $M : \emptyset \xrightarrow{D^{2n}} S^{2n-1} \xrightarrow{M'} Q$, and define

$$\bar{g}_M^\theta : \pi_0(\mathcal{N}_n^\theta(\emptyset)) \xrightarrow{\sim} \pi_0(\mathcal{N}_n^\theta(S^{2n-1})) \xrightarrow{\bar{g}_{M'}^\theta} \mathbb{Z},$$

which also satisfies $\bar{g}_M^\theta(W \circ M) = \bar{g}^\theta(W)$, by the way we have defined the stable θ -genus of a closed manifold. \square

If we use the function $\bar{g}_M^\theta : \mathcal{N}_n^\theta(P) \rightarrow \mathbb{Z}$ to grade $\mathcal{N}_n^\theta(P)$, then by Lemma 8.3 we have an induced map

$$(8.1) \quad (- \circ M)_* : H_i(\mathcal{N}_n^\theta(Q); (- \circ M)^* \mathcal{L})_{\bar{g}^\theta = g} \longrightarrow H_i(\mathcal{N}_n^\theta(P); \mathcal{L})_{\bar{g}_M^\theta = g}$$

for each local coefficient system \mathcal{L} on M . The following proves Corollary 1.7 in the case $P \neq \emptyset$.

Proposition 8.4. *If $2n \geq 6$ and B is simply-connected, then for any system of abelian coefficients \mathcal{L} on $\mathcal{N}_n^\theta(P)$ and any cobordism $M : P \rightsquigarrow Q$ which is $(n-1)$ -connected relative to Q and which has $P \neq \emptyset$, the map (8.1) is an isomorphism*

- (i) for $2i \leq g-3$ if \mathcal{L} is constant and θ is spherical, or
- (ii) for $3i \leq g-4$ if \mathcal{L} is constant, or
- (iii) for $3i \leq g-4$ if \mathcal{L} extends to $\text{hocolim}_{h \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(h)})$.

The extension condition in (iii) always holds for $g \geq 7$, and for $g \geq 5$ if θ is spherical.

Proof. Consider the commutative diagram

$$\begin{array}{ccc} \mathcal{N}_n^\theta(Q)_{\bar{g}^\theta = g} & \longrightarrow & \text{hocolim}_{h \rightarrow \infty} \mathcal{N}_n^\theta(Q, \hat{\ell}_Q^{(h)})_{\bar{g}^\theta = g+h} \\ \downarrow - \circ M & & \downarrow - \circ M \\ \mathcal{N}_n^\theta(P)_{\bar{g}_M^\theta = g} & \longrightarrow & \text{hocolim}_{h \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(h)})_{\bar{g}_M^\theta = g+h}, \end{array}$$

where the right-hand map is an abelian homology equivalence, by Theorem 1.3.

By assumption \mathcal{L} is pulled back from $\text{hocolim}_{h \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(h)})$, so defines a coefficient system on every space in this commutative square. By Theorem 8.2 the top horizontal map induces an isomorphism on $H_i(-; \mathcal{L})$ as long as $3i \leq g-4$ (or $2i \leq g-3$ if θ is spherical). As $\bar{g}_M^\theta = g$ implies that $\bar{g}^\theta \geq g$, Theorem 8.2 also implies that the lower horizontal map induces an isomorphism on $H_i(-; \mathcal{L})$ in this range of degrees, and so the left-hand vertical map induces an isomorphism on homology in the same range.

If $g \geq 7$, or if θ is spherical and $g \geq 5$, then by Theorem 8.2 the lower horizontal map induces a bijection on path components and an isomorphism on $H_1(-; \mathbb{Z})$, and so the two spaces have the same collections of abelian local coefficient systems. In particular \mathcal{L} is pulled back from $\text{hocolim}_{h \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(h)})$. \square

8.3. Proof of Corollary 1.7 in general. We will now finish the proof of Corollary 1.7 by explaining how to remove the hypothesis $P \neq \emptyset$ from Proposition 8.4. In order to do so, it is enough to prove the analogue of that proposition for the cobordism $(D^{2n}, \hat{\ell}_{D^{2n}}) : \emptyset \rightsquigarrow S^{2n-1}$, as any $M : \emptyset \rightsquigarrow Q$ can be factorised as $M : \emptyset \xrightarrow{D^{2n}} S^{2n-1} \xrightarrow{M'} Q$, and Proposition 8.4 applies to M' .

We will construct a resolution of the degree zero map of graded spaces

$$(8.2) \quad - \circ D^{2n} : \mathcal{N}_n^\theta(S^{2n-1}) \longrightarrow \mathcal{N}_n^\theta(\emptyset)$$

such that the induced map on the space of p -simplices in this resolution is of the form covered by Proposition 8.4. This is entirely analogous to the argument given for surfaces in [RW16, Section 11]. The resolution is conceptually very similar to that used in Section 7, but is a little different in detail.

Definition 8.5. Let $\mathcal{R}(P)_0$ be the set of tuples $(s, W; c, \hat{L})$ where $(s, W) \in \mathcal{N}_n^\theta(P)$, $c : D^{2n} \hookrightarrow (-\infty, 0) \times \mathbb{R}^\infty$, and $\hat{L} \in \text{Bun}^\theta(TD^{2n})^I$ are such that

- (i) c has image in W ,

(ii) \hat{L} is a path from $c^*\hat{\ell}_W$ to $\hat{\ell}_{D^{2n}}$.

We topologise $\mathcal{R}(P)_0$ as a subspace of

$$\mathcal{N}_n^\theta(P) \times \text{Emb}(D^{2n}, (-\infty, 0] \times \mathbb{R}^\infty) \times \text{Bun}^\theta(TD^{2n})^I.$$

Let $\mathcal{R}(P)_p$ consist of tuples $(s, W; c_0, \hat{L}_0, \dots, c_p, \hat{L}_p)$ such that

- (i) each $(s, W; c_i, \hat{L}_i)$ lies in $\mathcal{R}(P)_0$,
- (ii) the c_i are disjoint.

We topologise $\mathcal{R}(P)_p$ as a subspace of the $(p+1)$ -fold fibre product of the projection map $\mathcal{R}(P)_0 \rightarrow \mathcal{N}_n^\theta(P)$. The collection $\mathcal{R}(P)_\bullet$ has the structure of a semi-simplicial space augmented over $\mathcal{N}_n^\theta(P)$, where the i th face maps forgets (c_i, \hat{L}_i) , and the augmentation map just remembers (s, W) . Composition with $\bar{g}^\theta : \mathcal{N}_n^\theta(P) \rightarrow \mathbb{Z}$ gives it the structure of a semi-simplicial graded space augmented over $\mathcal{N}_n^\theta(P)$.

This semi-simplicial space replaces $\mathcal{L}_j(P)_\bullet$ in Section 6, and the following theorem may be proved in precisely the same way as Theorem 6.6.

Theorem 8.6. *The augmentation map $|\mathcal{R}(P)_\bullet| \rightarrow \mathcal{N}_n^\theta(P)$ is a weak homotopy equivalence (onto the components consisting of non-empty manifolds).*

The map (8.2) is covered by a semi-simplicial map $(-\circ D^{2n})_\bullet : \mathcal{R}(S^{2n-1})_\bullet \rightarrow \mathcal{R}(\emptyset)_\bullet$ given on 0-simplices by $(s, W; c, \hat{L}) \mapsto ((s, W) \circ (1, D^{2n}); c + e_0, \hat{L})$, and by the analogous formula on higher simplices.

Proposition 8.7. *If $2n \geq 6$ and B is simply-connected, the map of graded spaces*

$$(-\circ D^{2n})_p : \mathcal{R}(S^{2n-1})_p \longrightarrow \mathcal{R}(\emptyset)_p$$

induces an isomorphism on homology in the range described in Proposition 8.4.

Proof. An argument completely analogous to that of Propositions 4.14 and 6.7 identifies the map in question with the map of graded spaces

$$(-\circ D^{2n}) : \mathcal{N}_n^\theta(\sqcup^{p+2} S^{2n-1}) \longrightarrow \mathcal{N}_n^\theta(\sqcup^{p+1} S^{2n-1}).$$

As $\sqcup^{p+1} S^{2n-1}$ is non-empty for all $p \geq 0$, this induces isomorphisms in the range claimed, by Proposition 8.4. \square

The full statement of Corollary 1.7 now follows from the usual spectral sequence argument, namely that for any system \mathcal{L} of abelian local coefficients on $\mathcal{N}_n^\theta(\emptyset)$ the map of spectral sequences induced by $(-\circ D^{2n})_\bullet$ has the form

$$\begin{array}{ccc} E_{p,q}^1 = H_q(\mathcal{R}(S^{2n-1})_p; \mathcal{L})_{\bar{g}^\theta=g} & \Longrightarrow & H_{p+q}(|\mathcal{R}(S^{2n-1})_\bullet|; \mathcal{L})_{\bar{g}^\theta=g} \\ \downarrow & & \downarrow \\ \tilde{E}_{p,q}^1 = H_q(\mathcal{R}(\emptyset)_p; \mathcal{L})_{\bar{g}^\theta=g} & \Longrightarrow & H_{p+q}(|\mathcal{R}(\emptyset)_\bullet|; \mathcal{L})_{\bar{g}^\theta=g}, \end{array}$$

and by Theorem 8.6 its target can also be identified with the map induced by (8.2) on homology with \mathcal{L} -coefficients. Proposition 8.7 shows that the map of E^1 -pages is an isomorphism in a certain range of degrees, and Corollary 1.7 follows.

8.4. Stable homology and proof of Corollary 1.8. If $(W, \hat{\ell}_W)$ is a θ -manifold with non-empty boundary $(P, \hat{\ell}_P)$ and such that $\ell_W : W \rightarrow B$ is n -connected, then we can consider the commutative diagram

$$\begin{array}{ccccc} \mathcal{M}_n^\theta(W, \hat{\ell}_W) & \longrightarrow & \mathcal{N}_n^\theta(P, \hat{\ell}_P) & \longrightarrow & \text{hocolim}_{g \rightarrow \infty} \mathcal{N}_n^\theta(P, \hat{\ell}_P^{(g)}) \\ & \searrow & \downarrow & & \downarrow \\ & & \Omega_{[0, (P, \hat{\ell}_P)]} BC_\theta & \xrightarrow{\cong} & \text{hocolim}_{g \rightarrow \infty} \Omega_{[\emptyset, (P, \hat{\ell}_P^{(g)})]} BC_\theta \end{array}$$

By Corollary 1.7 the composition along the top, restricted to the path component which it hits, is an abelian homology isomorphism in degrees $3* \leq \bar{g}^\theta(W, \hat{\ell}_W) - 4$, and is a homology isomorphism in degrees $2* \leq \bar{g}^\theta(W, \hat{\ell}_W) - 3$ if θ is spherical. By Theorem 7.3 the right hand vertical map is acyclic, and the target has the homotopy type of a loop space so all coefficient systems on it are abelian. Therefore the diagonal map is a homology equivalence (or an acyclic map) onto the path component which it hits in the same range of degrees. This establishes Corollary 1.8 in this case, after replacing BC_θ by the homotopy equivalent $\Omega^{\infty-1}MT\theta$.

If $(W, \hat{\ell}_W)$ is a θ -manifold with empty boundary and such that $\ell_W : W \rightarrow B$ is n -connected, let $(W', \hat{\ell}_{W'}) : \emptyset \rightsquigarrow (P, \hat{\ell}_P)$ be a θ -cobordism obtained by subtracting a $2n$ -disc $D^{2n} = (D^{2n}, \hat{\ell}_{D^{2n}})$ from W . We obtain a homotopy commutative diagram

$$\begin{array}{ccccc} \mathcal{M}_n^\theta(W', \hat{\ell}_{W'}) & \longrightarrow & \mathcal{N}_n^\theta(P, \hat{\ell}_P) & \longrightarrow & \Omega_{[\emptyset, P]}BC_\theta \\ \downarrow -\circ D^{2n} & & \downarrow -\circ D^{2n} & & \simeq \downarrow -\circ D^{2n} \\ \mathcal{M}_n^\theta(W, \hat{\ell}_W) & \longrightarrow & \mathcal{N}_n^\theta(\emptyset) & \longrightarrow & \Omega BC_\theta \end{array}$$

where, by Corollary 1.7 and the case treated already, the left hand vertical map and the top composition are both abelian homology isomorphisms in degrees $3* \leq \bar{g}^\theta(W, \hat{\ell}_W) - 4$, and homology isomorphisms in degrees $2* \leq \bar{g}^\theta(W, \hat{\ell}_W) - 3$ if θ is spherical. Therefore the bottom composition has the same property, which establishes Corollary 1.8.

9. GENERAL TANGENTIAL STRUCTURES

So far we have only considered $2n$ -manifolds W equipped with θ -structures $\hat{\ell}_W$ such that the underlying map $\ell_W : W \rightarrow B$ is n -connected. In this section we shall formulate and prove a generalisation of Corollary 1.8 which allows for manifolds equipped with general tangential structures. We shall in particular prove Corollary 1.9, but also extend it to manifolds with non-empty boundary.

Definition 9.1. Let $\pi : E \rightarrow B$ be a fibration and $i : P \rightarrow E$ be a cofibration, in the usual model structure on Top (compactly generated weakly Hausdorff topological spaces, weak homotopy equivalences, Serre fibrations, induced cofibrations). Let $\text{hAut}(\pi, i)$ be the space of maps $f : E \rightarrow E$ which are weak equivalences and such that $\pi \circ f = \pi$ and $f \circ i = i$.

Composition of maps makes $\text{hAut}(\pi, i)$ into a grouplike topological monoid. If γ is a vector bundle over B , and we let $\pi^*\gamma$ be the pulled back vector bundle over E , then each f induces an endomorphism $\hat{f} : \pi^*\gamma \rightarrow \pi^*\gamma$ by sending $(e, v) \in \pi^*\gamma \subset E \times \gamma$ to $(f(e), v) \in E \times \gamma$, which lies in $\pi^*\gamma$ as f commutes with π . Furthermore, \hat{f} is the identity over the subspace $i(P) \subset E$.

To apply this to the situation at hand, let $\theta : B \xrightarrow{u} B' \xrightarrow{\theta'} BO(2n)$ be a pair of tangential structures with a map u between them, which is a fibration, and write $\hat{u} : \theta^*\gamma_{2n} \rightarrow (\theta')^*\gamma_{2n}$ for the induced bundle map. Let W be a $2n$ -dimensional manifold and $\hat{\ell}_{\partial W} : TW|_{\partial W} \rightarrow \theta^*\gamma_{2n}$ be a boundary condition with underlying map $\ell_{\partial W} : \partial W \rightarrow B$ a cofibration. Then $\text{hAut}(u, \ell_{\partial W})$ acts on $\text{Bun}_n^\theta(TW, \hat{\ell}_{\partial W})$ via $(f, \hat{\ell}) \mapsto \hat{f} \circ \hat{\ell}$, and this action commutes with the map

$$\begin{array}{ccc} \text{Bun}_n^\theta(TW, \hat{\ell}_{\partial W}) & \longrightarrow & \text{Bun}^{\theta'}(TW, \hat{u} \circ \hat{\ell}_{\partial W}) \\ \hat{\ell} & \longmapsto & \hat{u} \circ \hat{\ell} \end{array}$$

because $\hat{u} \circ \hat{f} = \hat{u}$.

Lemma 9.2. *Assume as above that $\ell_{\partial W} : \partial W \rightarrow B$ is a cofibration and that $u : B \rightarrow B'$ is a fibration. Assume in addition that the map u is n -co-connected. Then the induced map*

$$\mathrm{Bun}_n^\theta(TW, \hat{\ell}_{\partial W}) // \mathrm{hAut}(u, \ell_{\partial W}) \longrightarrow \mathrm{Bun}^{\theta'}(TW, \hat{u} \circ \hat{\ell}_{\partial W})$$

is a homotopy equivalence onto the path components which it hits.

Proof. Since u is a fibration, so is the induced map

$$(9.1) \quad \mathrm{Bun}_n^\theta(TW, \hat{\ell}_{\partial W}) \longrightarrow \mathrm{Bun}^{\theta'}(TW, \hat{u} \circ \hat{\ell}_{\partial W}),$$

and the fibre over $\hat{\ell}'_W$ may be identified with the space of relative lifts

$$\begin{array}{ccc} \partial W & \xrightarrow{\ell_{\partial W}} & B \\ \downarrow & \nearrow & \downarrow u \\ W & \xrightarrow{\ell'_W} & B' \end{array}$$

which are also n -connected maps. If such a lift exists, choose one and call it $f : W \rightarrow B$. We shall consider this f as a morphism in the category $\partial W/\mathrm{Top}/B'$ of spaces over B' and under ∂W ; this category has the structure of a simplicial model category and W is cofibrant and B is both fibrant and cofibrant, cf. e.g. [GS07, Example 1.7 and Definition 4.11]. Composition with f defines a map

$$(9.2) \quad \mathrm{Map}_{\partial W/\mathrm{Top}/B'}(B, B) \longrightarrow \mathrm{Map}_{\partial W/\mathrm{Top}/B'}(W, B)$$

between mapping simplicial sets, which can be seen to be a weak equivalence by induction on cells in an approximation to (B, W) by a relative CW complex. To give more details, let us write $F : (\partial W/\mathrm{Top}/B')^{\mathrm{op}} \rightarrow \mathrm{sSet}$ for the functor represented by the object B . Any representable functor takes pushout diagrams to pullback diagrams and by general properties of (simplicial) model categories, F sends cofibrations to fibrations and sends weak equivalences between cofibrant objects to weak equivalences. In particular, if X is a cofibrant object and $Y = X \cup D^k$ is obtained by attaching a cell $D^k \rightarrow B'$ to X along some map $\partial D^k \rightarrow X$ over B' , then there is the (strict) pullback diagram

$$\begin{array}{ccc} F(Y) & \longrightarrow & \mathrm{Map}_{\mathrm{Top}/B'}(D^k, B) \\ \downarrow & & \downarrow \\ F(X) & \longrightarrow & \mathrm{Map}_{\mathrm{Top}/B'}(\partial D^k, B), \end{array}$$

in which the vertical maps are actually fibrations. Up to homotopy, a point in $\mathrm{Map}_{\mathrm{Top}/B'}(\partial D^k, B)$ is described by a map from ∂D^k to a homotopy fibre of u and the space of extensions to D^k is a model for the homotopy fibre of the horizontal maps in the square. In particular the fibres are weakly contractible when $k \geq n+1$ if the homotopy fibres of u are $(n-1)$ -types, as we have assumed. Hence by induction on cells $F(Y) \rightarrow F(X)$ is a weak equivalence whenever X and Y are cofibrant and Y is obtained from X by attaching cells of dimension at least $n+1$. Up to weak equivalence the morphism $W \rightarrow B$ is of this form, and hence (9.2) is a weak equivalence.

The simplicial sets in (9.2) are the singular simplicial sets of the associated mapping spaces, and composition with f takes the subspace of the mapping space $\mathrm{map}_{\partial W/\mathrm{Top}/B'}(B, B)$ consisting of weak equivalences onto the subspace of the mapping space $\mathrm{map}_{\partial W/\mathrm{Top}/B'}(W, B)$ consisting of n -connected maps. Hence we have shown that the homotopy fibres of (9.1) are either empty or weakly equivalent to $\mathrm{hAut}(u, \ell_{\partial W})$, in which case the equivalence is given by acting on an element. Thus

after forming the Borel construction by $\mathrm{hAut}(u, \ell_{\partial W})$, the fibres of the map (9.1) become empty or contractible. \square

We obtain a map $\mathcal{M}_n^\theta(W, \hat{\ell}_{\partial W}) // \mathrm{hAut}(u, \ell_{\partial W}) \rightarrow \mathcal{M}^{\theta'}(W, \hat{u} \circ \hat{\ell}_{\partial W})$ by taking the further Borel construction by $\mathrm{Diff}_\partial(W)$, and this is also a weak homotopy equivalence onto those path components which it hits. Therefore, in terms of the point-set models introduced in Definition 2.18, the map

$$(9.3) \quad \mathcal{M}_n^\theta(W, \hat{\ell}_{\partial W}) // \mathrm{hAut}(u, \ell_{\partial W}) \longrightarrow \mathcal{M}^{\theta'}(W, \hat{u} \circ \hat{\ell}_{\partial W})$$

is also a weak homotopy equivalence onto those path components which it hits.

Each $f \in \mathrm{hAut}(u)$ gives a bundle map $\hat{f} : \theta^* \gamma_{2n} \rightarrow \theta'^* \gamma_{2n}$ which in turn defines a continuous endofunctor of \mathcal{C}_θ (by postcomposition on all θ -structures), and hence a continuous endomorphism of BC_θ . This defines an action of $\mathrm{hAut}(u)$ on BC_θ . Moreover, if $(\partial W, \hat{\ell}_{\partial W}) \in \mathcal{C}_\theta$ is an object satisfying the assumption above that $\ell_{\partial W} : \partial W \rightarrow B$ is a cofibration, then $\mathrm{hAut}(u, \ell_{\partial W})$ preserves $(\partial W, \hat{\ell}_{\partial W}) \in \mathcal{C}_\theta$. In this case $\mathrm{hAut}(u, \ell_{\partial W})$ acts on both the morphism space $\mathcal{C}_\theta(\emptyset, (\partial W, \hat{\ell}_{\partial W}))$ and on the path space $\Omega_{[\emptyset, (\partial W, \hat{\ell}_{\partial W})]} BC_\theta$, and the map

$$\mathcal{C}_\theta(\emptyset, (\partial W, \hat{\ell}_{\partial W})) \longrightarrow \Omega_{[\emptyset, (\partial W, \hat{\ell}_{\partial W})]} BC_\theta$$

is equivariant for this action.

Definition 9.3. If $(W, \hat{\ell}_W)$ is a θ -manifold, we write $\Omega_{(W, \hat{\ell}_W)} BC_\theta$ for the path component of $\Omega_{[\emptyset, (\partial W, \hat{\ell}_{\partial W})]} BC_\theta$ represented by the θ -cobordism $(W, \hat{\ell}_W) : \emptyset \rightsquigarrow (\partial W, \hat{\ell}_{\partial W})$.

If $(W, \hat{\ell}'_W)$ is a θ' -manifold and $\ell'_W : W \xrightarrow{\ell'_W} B \xrightarrow{u} B'$ is a Moore–Postnikov n -stage of ℓ'_W , then we write $\Omega_{(W, \hat{\ell}'_W)} BC_\theta$ for the union of those path components of $\Omega_{[\emptyset, (\partial W, \hat{\ell}_{\partial W})]} BC_\theta$ represented by the $\pi_0 \mathrm{hAut}(u, \ell_{\partial W})$ -orbit of $(W, \hat{\ell}_W)$.

Lemma 9.4. *If $(W, \hat{\ell}'_W)$ is a θ' -manifold and $\ell'_W : W \xrightarrow{\ell'_W} B \xrightarrow{u} B'$ is a Moore–Postnikov n -stage of ℓ'_W , then $\bar{g}^{\theta'}(W, \hat{\ell}'_W) = \bar{g}^\theta(W, \hat{\ell}_W)$.*

Proof. By the definition of the stable genus, it is enough to show that $g^{\theta'}(W, \hat{\ell}'_W) = g^\theta(W, \hat{\ell}_W)$. By definition of standard θ - or θ' -structure on $S^n \times D^n$, and hence on $W_{1,1}$, composing with u sends standard θ -structures to standard θ' -structures, and so certainly $g^{\theta'}(W, \hat{\ell}'_W) \geq g^\theta(W, \hat{\ell}_W)$. Conversely, if $\hat{\ell} : T(S^n \times D^n) \rightarrow \theta^* \gamma_{2n}$ is a θ -structure such that $\hat{u} \circ \hat{\ell}$ is a standard θ' -structure, then $\hat{u} \circ \hat{\ell}$ extends over the embedding $S^n \times D^n \hookrightarrow \mathbb{R}^{2n}$ described in Section 2.2, so we have a commutative square

$$\begin{array}{ccc} S^n \times D^n & \xrightarrow{\ell} & B \\ \downarrow & \nearrow & \downarrow u \\ \mathbb{R}^{2n} & \longrightarrow & B'. \end{array}$$

The map u is n -co-connected, as it arises from the Moore–Postnikov factorisation of ℓ'_W . Therefore the square admits a diagonal map, showing that $\hat{\ell}$ is a standard θ -structure. \square

The following theorem implies Corollary 1.9 from the introduction, but also allows manifolds with non-empty boundary.

Theorem 9.5. *Let $\theta' : B' \rightarrow BO(2n)$ be a tangential structure, and $(W, \hat{\ell}'_W)$ be a θ' -manifold. Let $\ell'_W : W \xrightarrow{\ell'_W} B \xrightarrow{u} B'$ be a Moore–Postnikov n -stage of ℓ'_W , i.e. a*

factorisation into an n -connected cofibration ℓ_W and an n -co-connected fibration u , and $\theta = \theta' \circ u$. Then there is a map

$$\alpha_W^{\theta'} : \mathcal{M}^{\theta'}(W, \hat{\ell}'_W) \longrightarrow \left(\Omega_{[W, \hat{\ell}'_W]} BC_\theta \right) // \text{hAut}(u, \ell_{\partial W})$$

which, if $2n \geq 6$ and W is simply-connected, is acyclic in degrees $3* \leq \bar{g}^{\theta'}(W, \hat{\ell}'_W) - 4$, and is a homology isomorphism in degrees $2* \leq \bar{g}^{\theta'}(W, \hat{\ell}'_W) - 3$ if θ' is spherical.

Proof. We have $\text{hAut}(u, \ell_{\partial W})$ -equivariant maps

$$\alpha_W^\theta : \mathcal{M}_n^\theta(W; \hat{\ell}_{\partial W}) \longrightarrow \mathcal{C}_\theta(\emptyset, (\partial W, \hat{\ell}_{\partial W})) \longrightarrow \Omega_{[\emptyset, (\partial W, \hat{\ell}_{\partial W})]} BC_\theta$$

and so on taking Borel constructions and using (9.3) we obtain a diagram

$$\begin{array}{ccc} \mathcal{M}^{\theta'}(W; \hat{u} \circ \hat{\ell}_{\partial W}) & \longleftarrow & \mathcal{M}_n^\theta(W; \hat{\ell}_{\partial W}) // \text{hAut}(u, \ell_{\partial W}) \\ & & \downarrow \alpha_W^\theta // \text{hAut}(u, \ell_{\partial W}) \\ & & (\Omega_{[\emptyset, (\partial W, \ell_{\partial W})]} BC_\theta) // \text{hAut}(u, \ell_{\partial W}). \end{array}$$

The horizontal map is a weak homotopy equivalence onto the path components which it hits, and the $\text{hAut}(u, \ell_{\partial W})$ -orbit of $\hat{\ell}_W$ hits the path component containing $\hat{\ell}'_W$. Restricting to this component and inverting the homotopy equivalence, we obtain a map we define to be $\alpha_W^{\theta'}$.

Any θ -structure in the $\text{hAut}(u, \ell_{\partial W})$ -orbit of $\hat{\ell}_W$ becomes homotopic to $\hat{\ell}'_W$ after applying u , so has stable θ -genus equal to $\bar{g}^{\theta'}(W, \hat{\ell}'_W)$ by Lemma 9.4. On each such path component, by Corollary 1.8 the map α_W^θ is acyclic in degrees $3* \leq \bar{g}^{\theta'}(W, \hat{\ell}'_W) - 4$, and a homology isomorphism in degrees $2* \leq \bar{g}^{\theta'}(W, \hat{\ell}'_W) - 3$ if θ' is spherical (θ is spherical if and only if θ' is, as $u : B \rightarrow B'$ is n -co-connected), onto the path component which it hits. It therefore remains so after taking the Borel construction, and so the right-hand map has the same property. \square

Finally, we observe that the maps used in establishing the main theorem of [GMTW09], $BC_\theta \simeq \Omega^{\infty-1} \text{MT}\theta$, are all $\text{hAut}(\theta)$ -equivariant, which allows us to translate Theorem 9.5 into the form given in the introduction.

APPENDIX A. HOMOLOGY EQUIVALENCES AND LOCAL COEFFICIENTS

Parts of this appendix are loosely based on one written by Johannes Ebert for an early (uncirculated) draft of [BERW14]. The authors are grateful to him for allowing us to make use of it.

A.1. Local coefficient systems. We shall take the point of view that a local coefficient system on a space X is a functor $\mathcal{L} : \pi_1(X) \rightarrow \text{Ab}$ from the fundamental groupoid of X to the category of abelian groups. It is important that local systems form a *category*, whose morphisms $T : \mathcal{L} \rightarrow \mathcal{L}'$ are the natural transformations. In particular each coefficient system \mathcal{L} has an endomorphism ring $R = \text{End}(\mathcal{L})$, and in fact the functor $\mathcal{L} : \pi_1(X) \rightarrow \text{Ab}$ naturally factors through the category of R -modules. In particular the values $\mathcal{L}(x)$ are naturally $\mathbb{Z}[\pi_1(X, x)]$ - R -bimodules.

Singular homology assigns to such a local coefficient system \mathcal{L} a chain complex $C_*(X; \mathcal{L})$ and homology groups $H_*(X; \mathcal{L})$, and a natural transformation $T : \mathcal{L} \rightarrow \mathcal{L}'$ induces maps of chains and of homology $T : H_*(X; \mathcal{L}) \rightarrow H_*(X; \mathcal{L}')$. If $f : X \rightarrow Y$ is a continuous map and $\mathcal{L} : \pi_1(Y) \rightarrow \text{Ab}$ is a coefficient system, then there is a pulled-back coefficient system $f^*\mathcal{L} : \pi_1(X) \rightarrow \text{Ab}$ and an induced map

$$f_* : H_*(X; f^*\mathcal{L}) \longrightarrow H_*(Y; \mathcal{L}).$$

For a fixed \mathcal{L} the singular chains and the homology groups $H_*(Y; \mathcal{L})$ are naturally modules over the ring $R = \text{End}(\mathcal{L})$, and the functoriality with respect to $f : X \rightarrow Y$ is compatible with this structure.

For any space X with coefficient system \mathcal{L} there exists for any $x \in H_k(X; \mathcal{L})$ a finite CW complex K and a map $f : K \rightarrow X$, such that x is in the image of $f_* : H_*(K; f^*\mathcal{L}) \rightarrow H_*(X; \mathcal{L})$.

A.1.1. Homology equivalences.

Definition A.1. A local coefficient system $\mathcal{L} : \pi_1(X) \rightarrow \text{Ab}$ is *constant* if the action of $\pi_1(X, x)$ on $\mathcal{L}(x)$ is trivial for all $x \in X$. The system is *abelian* if the commutator subgroup of $\pi_1(X, x)$ acts trivially for all $x \in X$.

Definition A.2. Let $f : X \rightarrow Y$ be a continuous map, and for each coefficient system \mathcal{L} on Y let

$$(A.1) \quad f_* : H_n(X; f^*\mathcal{L}) \longrightarrow H_n(Y; \mathcal{L})$$

be the map induced by f .

- (i) If (A.1) is an isomorphism for all n and all \mathcal{L} then f is called an *acyclic map*.
- (ii) If (A.1) is an isomorphism for all n and all abelian \mathcal{L} then f is called an *abelian equivalence*.
- (iii) If (A.1) is an isomorphism for all n and all constant \mathcal{L} then f is called a *homology equivalence*.

We remark that the class of coefficient systems $\mathcal{L} : \pi_1(Y) \rightarrow \text{Ab}$ for which (A.1) is an equivalence is closed under filtered colimits and extensions. Hence, if $f : X \rightarrow Y$ is an abelian equivalence then (A.1) will be an isomorphism for any coefficient system \mathcal{L} obtained from abelian ones by extensions and filtered colimits. This class of local systems can alternatively be described as those which for all $y \in Y$ satisfy that $\mathcal{L}(y)$ admits no non-trivial homomorphism $\mathcal{L}(y) \rightarrow W$ for any $\mathbb{Z}[\pi_1(Y, y)]$ -module W with $W^{[\pi_1(Y, y), \pi_1(Y, y)]} = 0$.

A.1.2. Special abelian coefficient systems.

Definition A.3. A *special abelian coefficient system* on a space X consists of a commutative ring R and a coefficient system $\mathcal{L} : \pi_1(X) \rightarrow \text{Mod}_R$ such that $\{x \in X \mid \mathcal{L}(x) \neq 0\}$ is a path component of X , and for x in this path component $\mathcal{L}(x)$ is a free R -module of rank 1.

Note that a special abelian coefficient system is indeed an abelian coefficient system, as $\pi_1(X, x)$ acts on $\mathcal{L}(x)$ via R^\times , an abelian group.

Lemma A.4. *If $f : X \rightarrow Y$ is a continuous map such that $f_* : H_*(X; f^*\mathcal{L}) \rightarrow H_*(Y; \mathcal{L})$ is an isomorphism for all special abelian coefficient systems \mathcal{L} on Y , then f is an abelian equivalence.*

Proof. If $\mathcal{L} : \pi_1(Y) \rightarrow \text{Mod}_{\mathbb{F}_2}$ is the unique special abelian coefficient system supported on the path component given by $y \in \pi_0(Y)$, then the dimension of $H_0(X; f^*\mathcal{L})$ as an \mathbb{F}_2 vector space is the cardinality of the inverse image of y under $\pi_0(X) \rightarrow \pi_0(Y)$. Hence f induces a bijection $\pi_0(X) \rightarrow \pi_0(Y)$.

We can then restrict to the case where both X and Y are path connected and deduce that $f_* : H_1(X; \mathbb{Z}) \rightarrow H_1(Y; \mathbb{Z})$ is an equivalence by using the constant coefficient system $\mathbb{Z} : \pi_1(Y) \rightarrow \text{Mod}_{\mathbb{Z}}$, which is special abelian. Then we pick a point $y \in Y$ and use the Hurewicz homomorphism $\pi_1(Y, y) \rightarrow A = H_1(Y)$ to define the special abelian coefficient system $\mathcal{L}_y : \pi_1(Y) \rightarrow \text{Mod}_{\mathbb{Z}[A]}$ by

$$\mathcal{L}_y(z) = \mathbb{Z}[A] \otimes_{\mathbb{Z}[\pi_1(Y, y)]} \mathbb{Z}[\pi_1(Y)(y, z)].$$

Then $H_*(Y; \mathcal{L}_y)$ and $H_*(X; f^* \mathcal{L}_y)$ calculate the homology of the universal abelian covers and if f_* induces an isomorphism between these, then it induces an isomorphism for all abelian coefficient systems. \square

The following lemma is the main advantage of special abelian systems over all systems. In general if $f : X \rightarrow Y$ and \mathcal{L} is an arbitrary coefficient system on Y , then the natural map $\text{End}(\mathcal{L}) \rightarrow \text{End}(f^* \mathcal{L})$ can be very far from surjective if the system is not abelian, even when X and Y are path connected.

Lemma A.5. *Let X and Y be path connected. If $f, g : X \rightarrow Y$ are two maps inducing equal homomorphisms $H_1(X; \mathbb{Z}) \rightarrow H_1(Y; \mathbb{Z})$, then $f^* \mathcal{L}$ and $g^* \mathcal{L} : \pi_1(X) \rightarrow \text{Mod}_R$ are isomorphic for any special abelian $\mathcal{L} : \pi_1(Y) \rightarrow \text{Mod}_R$.*

Let X be path connected and $\mathcal{L}, \mathcal{L}' : \pi_1(X) \rightarrow \text{Mod}_R$ be isomorphic special abelian coefficient systems. Then the R -module $\text{Hom}_R(\mathcal{L}, \mathcal{L}')$, consisting of natural transformations of functors $\pi_1(X) \rightarrow \text{Mod}_R$, is free of rank one.

Let $f : X \rightarrow Y$ be a map between path connected spaces and let $\mathcal{L}, \mathcal{L}' : \pi_1(Y) \rightarrow \text{Mod}_R$ be isomorphic special abelian coefficient systems. Then the natural map

$$\text{Hom}_R(\mathcal{L}, \mathcal{L}') \longrightarrow \text{Hom}_R(f^* \mathcal{L}, f^* \mathcal{L}')$$

is an isomorphism.

Proof. To verify the first claim we pick a basepoint $x \in X$. Then \mathcal{L} gives homomorphisms $\pi_1(Y, f(x)) \rightarrow R^\times$ and $\pi_1(Y, g(x)) \rightarrow R^\times$ given by equal elements of $H^1(Y; R^\times)$, and the isomorphism classes of $f^* \mathcal{L}$ and $g^* \mathcal{L}$ are determined by the image of this element under the equal maps $f^* = g^* : H^1(Y; R^\times) \rightarrow H^1(X; R^\times)$.

To verify the second claim we also use a basepoint $x \in X$. It suffices to prove that the restriction map $\text{Hom}_R(\mathcal{L}, \mathcal{L}') \rightarrow \text{Hom}_R(\mathcal{L}(x), \mathcal{L}'(x))$ is an isomorphism. It is injective as X is path connected, and any choice of an isomorphism $\mathcal{L} \cong \mathcal{L}'$ gives an element of $\text{Hom}_R(\mathcal{L}, \mathcal{L}')$ which restricts to a generator of the R -module $\text{Hom}_R(\mathcal{L}(x), \mathcal{L}'(x)) \cong R$.

The third claim now follows because the homomorphism is between free R -modules of rank one and sends a generator to a generator. \square

A.1.3. Homotopies. Finally, let us discuss homotopies. The two inclusion $i_0, i_1 : X \rightarrow I \times X$ induce two functors $\pi_1(i_0), \pi_1(i_1) : \pi_1(X) \rightarrow \pi_1(I \times X)$, and there is a canonical natural isomorphism $\pi_1(i_0) \Rightarrow \pi_1(i_1)$ given on the object x by the path from $i_0(x) = (0, x)$ to $i_1(x) = (1, x)$ along the interval. If R is an associative ring and $\mathcal{L} : \pi_1(I \times X) \rightarrow \text{Mod}_R$ is a coefficient system, we may compose \mathcal{L} with this natural transformation to get an induced isomorphism $T : i_0^* \mathcal{L} \rightarrow i_1^* \mathcal{L}$ of functors into Mod_R . For any homotopy $H : I \times X \rightarrow Y$ from f to g and coefficient system $\mathcal{L} : \pi_1(Y) \rightarrow \text{Mod}_R$ we therefore have an induced isomorphism

$$T_H : f^* \mathcal{L} \longrightarrow g^* \mathcal{L}$$

of functors $\pi_1(X) \rightarrow \text{Mod}_R$, fitting into a commutative diagram

$$(A.2) \quad \begin{array}{ccc} H_*(X; f^* \mathcal{L}) & \xrightarrow{f_*} & H_*(Y; \mathcal{L}) \\ T_H \downarrow & & \parallel \\ H_*(X; g^* \mathcal{L}) & \xrightarrow{g_*} & H_*(Y; \mathcal{L}). \end{array}$$

A.2. Invariance properties for abelian coefficients. The following observation is immediate from diagram (A.2).

Proposition A.6. *Let $f, g : X \rightarrow Y$ be homotopic maps and let $\mathcal{L} : \pi_1(Y) \rightarrow \text{Mod}_R$ be any coefficient system.*

- (i) For any n , $f_* : H_n(X; f^* \mathcal{L}) \rightarrow H_n(Y; \mathcal{L})$ is surjective if and only if $g_* : H_n(X; g^* \mathcal{L}) \rightarrow H_n(Y; \mathcal{L})$ is surjective.
- (ii) For any n , $f_* : H_n(X; f^* \mathcal{L}) \rightarrow H_n(Y; \mathcal{L})$ is injective if and only if $g_* : H_n(X; g^* \mathcal{L}) \rightarrow H_n(Y; \mathcal{L})$ is injective.

In particular, f is a homology equivalence, abelian equivalence, or acyclic map, if and only if g has that property. \square

It turns out that in the case of abelian equivalences, the assumption in the above proposition that f and g be homotopic can be replaced by the following weaker assumption, called “weakly homotopic” by [Ada71].

Definition A.7. Two parallel maps

$$X \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} Y$$

are *almost homotopic* if for all finite CW complexes K and maps $i : K \rightarrow X$ we have $i \circ f \simeq i \circ g$.

Proposition A.8. Let X and Y be path connected, let $f, g : X \rightarrow Y$ be almost homotopic and let $\mathcal{L} : \pi_1(Y) \rightarrow \text{Mod}_R$ be a special abelian coefficient system.

- (i) If $f_* : H_k(X; f^* \mathcal{L}) \rightarrow H_k(Y; \mathcal{L})$ is surjective, then $g_* : H_k(X; g^* \mathcal{L}) \rightarrow H_k(Y; \mathcal{L})$ is also surjective.
- (ii) If $f_* : H_k(X; f^* \mathcal{L}) \rightarrow H_k(Y; \mathcal{L})$ is injective, then $g_* : H_k(X; g^* \mathcal{L}) \rightarrow H_k(Y; \mathcal{L})$ is also injective.

Proof. Without loss of generality we may assume X is a CW complex. The surjectivity part in fact holds true for any coefficient system, with essentially the same proof as in Proposition A.6: Any homology class in $H_n(Y; \mathcal{L})$ is by assumption the image of some $x \in H_n(X; f^* \mathcal{L})$ under f_* , and any such x is supported on some finite subcomplex $K \subset X$, but the map

$$(g|_K)_* : H_n(K; g|_K^* \mathcal{L}) \rightarrow H_n(Y; \mathcal{L})$$

has the same image as $(f|_K)_*$, by diagram (A.2).

Injectivity is not true for general coefficient systems, but when $\mathcal{L} : \pi_1(Y) \rightarrow \text{Mod}_R$ is special abelian we may argue as follows. Suppose $f_* : H_n(X; f^* \mathcal{L}) \rightarrow H_n(Y; \mathcal{L})$ is injective and suppose $x \in \text{Ker}(g_*) \subset H_n(X; g^* \mathcal{L})$. Again x is supported on some finite subcomplex $K \subset X$ and by assumption we may choose a homotopy $H_K : f|_K \simeq g|_K$, inducing an isomorphism of coefficient systems $T_{H_K} : f|_K^* \mathcal{L} \rightarrow g|_K^* \mathcal{L}$ and a commutative diagram

$$\begin{array}{ccc} H_*(K; f|_K^* \mathcal{L}) & \xrightarrow{(f|_K)_*} & H_*(Y; \mathcal{L}) \\ T_{H_K} \downarrow & & \parallel \\ H_*(K; g|_K^* \mathcal{L}) & \xrightarrow{(g|_K)_*} & H_*(Y; \mathcal{L}). \end{array}$$

Since any class in $H_1(X; \mathbb{Z})$ is supported on a finite subcomplex of X , we see that $f_* = g_* : H_1(X; \mathbb{Z}) \rightarrow H_1(Y; \mathbb{Z})$. If we assume, as we may, that K is path connected, the special abelianness of \mathcal{L} implies by Lemma A.5 that

$$\text{Hom}_R(f^* \mathcal{L}, g^* \mathcal{L}) \longrightarrow \text{Hom}_R(f|_K^* \mathcal{L}, g|_K^* \mathcal{L})$$

is an isomorphism of R -modules, and in particular the isomorphism T_{H_K} is induced by an isomorphism $T \in \text{Hom}_R(f^*\mathcal{L}, g^*\mathcal{L})$. We may therefore form a diagram

$$\begin{array}{ccccc} H_*(K; f|_K^*\mathcal{L}) & \longrightarrow & H_*(X; f^*\mathcal{L}) & \xrightarrow{f_*} & H_*(Y; \mathcal{L}) \\ T_{H_K} \downarrow & & \downarrow T & & \parallel \\ H_*(K; g|_K^*\mathcal{L}) & \longrightarrow & H_*(X; g^*\mathcal{L}) & \xrightarrow{g_*} & H_*(Y; \mathcal{L}), \end{array}$$

whose horizontal maps are induced by the maps of underlying spaces and whose vertical maps are induced by isomorphisms of coefficient systems, in which the left-hand square and the outer rectangle commute. Now the element $x \in H_n(X; g^*\mathcal{L})$ lifts to an $x' \in H_n(K; g|_K^*\mathcal{L})$ by assumption, and by the commutativity of the outer rectangle we have $(f|_K)_* T_{H_K}^{-1}(x') = (g|_K)_*(x') = g_*(x) = 0$. As f_* is injective and the left-hand square commutes, it follows that $x = 0$, as required. \square

Corollary A.9. *Let $f, g : X \rightarrow Y$ be almost homotopic maps. Then f is an abelian equivalence if and only if g is.*

Proof. It is easy to reduce to the case of path connected spaces, and by Lemma A.4 it suffices to consider special abelian coefficient systems, where the claim follows from Proposition A.8. \square

A.3. Ladder diagrams and abelian equivalences. A *ladder diagram* is a diagram of spaces of the shape

$$(A.3) \quad \begin{array}{ccccccc} X_0 & \xrightarrow{g_1} & X_1 & \xrightarrow{g_2} & X_2 & \longrightarrow & \cdots \\ \downarrow f_0 & & \downarrow f_1 & & \downarrow f_2 & & \\ Y_0 & \xrightarrow{h_1} & Y_1 & \xrightarrow{h_2} & Y_2 & \longrightarrow & \cdots \end{array}$$

A commutative ladder diagram is one in which each square commutes, and a homotopy commutative ladder diagram is one in which each square commutes up to homotopy.

For a ladder diagram as above we shall write X_∞ for the telescope (homotopy colimit) of the top row and Y_∞ for the bottom row. If the diagram commutes, there is an induced map $X_\infty \rightarrow Y_\infty$. If the diagram homotopy commutes, then a choice of homotopies $H_i : h_i \circ f_{i-1} \simeq f_i \circ g_i$ induces a map $X_\infty \rightarrow Y_\infty$. The homotopy class of the map of telescopes in general depends on the choice of homotopies H_i , but we have the following result.

Lemma A.10. *For a homotopy commutative ladder diagram (A.3), any two choices of homotopies $H_i : h_i \circ f_{i-1} \simeq f_i \circ g_i$ induce almost homotopic maps $X_\infty \rightarrow Y_\infty$. In particular, the property of whether the induced map $X_\infty \rightarrow Y_\infty$ is an abelian equivalence depends only on the underlying commutative ladder diagram in the homotopy category, not on the choices of homotopies H_i .*

Proof. The induced maps become equal when restricted to each $X_i \subset X_\infty$ and hence they become homotopic when restricted to the finite telescope of $(X_0 \rightarrow \cdots \rightarrow X_i)$ for all i . Any map from a finite CW complex will factor through one of these finite telescopes. \square

Definition A.11. A map $f : X \rightarrow Y$ is an *almost homotopy equivalence* if there exists a map $g : Y \rightarrow X$ such that $f \circ g$ and $g \circ f$ are each almost homotopic to the identity.

The following is immediate from Corollary A.9 and this definition.

Corollary A.12. *An almost homotopy equivalence is an abelian equivalence. \square*

The main technical result of this subsection is as follows.

Proposition A.13. *A homotopy commutative ladder diagram as in (A.3) induces an almost homotopy equivalence $X_\infty \rightarrow Y_\infty$, and hence an abelian equivalence, provided that the following condition holds: for each n , there exists a $k > n$, a map $i : Y_n \rightarrow X_k$ such that $f_k \circ i \simeq h_k \circ \cdots \circ h_{n+1}$, as well as a map $j : Y_n \rightarrow X_k$ such that $j \circ f_n \simeq g_k \circ \cdots \circ g_{n+1}$.*

Proof. After restricting to cofinal subsequences and reindexing, we may assume that we have maps $i_n, j_n : Y_{n-1} \rightarrow X_n$, making each triangle in the diagram

$$\begin{array}{ccc} X_{n-1} & \xrightarrow{g_n} & X_n \\ f_{n-1} \downarrow & \nearrow j_n & \downarrow f_n \\ Y_{n-1} & \xrightarrow{h_n} & Y_n, \\ & \nearrow i_n & \end{array}$$

homotopy commute, but possibly i_n and j_n are not homotopic. If we write $k_n = j_n \circ f_{n-1} \circ i_{n-1} : Y_{n-2} \rightarrow X_n$, we may calculate

$$\begin{aligned} f_n \circ k_n &= f_n \circ j_n \circ f_{n-1} \circ i_{n-1} \simeq f_n \circ g_n \circ i_{n-1} \simeq h_n \circ f_{n-1} \circ i_{n-1} \simeq h_n \circ h_{n-1}, \\ k_n \circ f_{n-2} &= j_n \circ f_{n-1} \circ i_{n-1} \circ f_{n-2} \simeq j_n \circ h_{n-1} \circ f_{n-2} \simeq j_n \circ f_{n-1} \circ g_{n-1} \\ &= g_n \circ g_{n-1}. \end{aligned}$$

In other words, after passing to a further subsequence and reindexing we have a single diagonal map

$$\begin{array}{ccc} X_{n-1} & \xrightarrow{g_n} & X_n \\ f_{n-1} \downarrow & \nearrow k_n & \downarrow f_n \\ Y_{n-1} & \xrightarrow{h_n} & Y_n, \end{array}$$

such that both triangles homotopy commute. The resulting diagram may be reinterpreted as a homotopy commutative diagram

$$\begin{array}{ccccccc} X_0 & \xrightarrow{g_1} & X_1 & \xrightarrow{g_2} & X_2 & \longrightarrow & \cdots \\ \downarrow f_0 & & \downarrow f_1 & & \downarrow f_2 & & \\ Y_0 & \xrightarrow{h_1} & Y_1 & \xrightarrow{h_2} & Y_2 & \longrightarrow & \cdots \\ \downarrow k_1 & & \downarrow k_2 & & \downarrow k_3 & & \\ X_1 & \xrightarrow{g_2} & X_2 & \xrightarrow{g_3} & X_3 & \longrightarrow & \cdots \end{array}$$

such that the vertical maps from the top row to the bottom row are homotopic to $g_n : X_{n-1} \rightarrow X_n$. Any choice of homotopies in the squares now induce maps $k \circ f : X_\infty \rightarrow Y_\infty \rightarrow X_\infty$, whose composition $X_\infty \rightarrow X_\infty$ is almost homotopic to the “shift map” of the telescope, and hence almost homotopic to the identity map of X_∞ . By a similar argument we see that $f \circ k : Y_\infty \rightarrow X_\infty \rightarrow Y_\infty$ is almost homotopic to the identity, and hence $f : X_\infty \rightarrow Y_\infty$ is an almost homotopy equivalence, as claimed. \square

A.4. The generalised group completion theorem with local coefficients.

Let \mathcal{C} be a (possibly non-unital) topological category, and write $N_\bullet \mathcal{C}$ for the nerve, so that $N_0 \mathcal{C}$ is the space of objects and $N_1 \mathcal{C}$ is the space of morphisms. Let BC be the geometric realisation of the nerve as a *semi-simplicial* space. The space of morphisms from a to b is denoted $\mathcal{C}(a, b)$.

Fix a sequence of objects and morphisms

$$c_0 \xrightarrow{f_1} c_1 \xrightarrow{f_2} c_2 \xrightarrow{f_3} \cdots$$

in \mathcal{C} . Then, for objects a and b of \mathcal{C} and a morphism $g : a \rightarrow b$, we obtain a commutative ladder diagram

$$\begin{array}{ccccccc} \mathcal{C}(b, c_0) & \xrightarrow{f_1 \circ -} & \mathcal{C}(b, c_1) & \xrightarrow{f_2 \circ -} & \mathcal{C}(b, c_2) & \longrightarrow & \cdots \\ \downarrow - \circ g & & \downarrow - \circ g & & \downarrow - \circ g & & \\ \mathcal{C}(a, c_0) & \xrightarrow{f_1 \circ -} & \mathcal{C}(a, c_1) & \xrightarrow{f_2 \circ -} & \mathcal{C}(a, c_2) & \longrightarrow & \cdots \end{array}$$

The generalised version of the group completion theorem is then as follows.

Theorem A.14. *Let \mathcal{C} be a (perhaps non-unital) topological category and let c_n and f_n ($n \geq 0$) be as above. Suppose in addition that*

- (i) *The map $(d_0, d_1) : N_1\mathcal{C} \rightarrow (N_0\mathcal{C})^2$ is a fibration.*
- (ii) *For any morphism $g \in \mathcal{C}(a, b)$, the induced map*

$$(A.4) \quad \text{hocolim}_{i \rightarrow \infty} \mathcal{C}(b, c_i) \rightarrow \text{hocolim}_{i \rightarrow \infty} \mathcal{C}(a, c_i)$$

is an abelian homology equivalence.

Then the natural map

$$(A.5) \quad \text{hocolim}_{i \rightarrow \infty} \mathcal{C}(a, c_i) \longrightarrow \text{hocolim}_{i \rightarrow \infty} \Omega_{[a, c_i]} BC$$

is acyclic.

Proof sketch. If we replace “abelian homology equivalence” and “acyclic” in the statement by “homology equivalence”, this is a standard argument based on the notion of homology fibration from [MS76], cf. [Til97, Theorem 3.2], [GMTW09, §7], or [GRW14b, §7.4], with a small modification to deal with the possible lack of units. Let us outline the argument. For $i \geq 0$ let E_\bullet^i be the semi-simplicial space with $E_p^i \subset N_{p+1}\mathcal{C}$ the subspace consisting of sequences of $(p+1)$ composable morphisms ending in c_i , i.e. the inverse image of c_i under $d_0^{p+1} : N_{p+1}\mathcal{C} \rightarrow N_0\mathcal{C}$. Postcomposing the last morphism with $f_{i+1} : c_i \rightarrow c_{i+1}$ defines a semi-simplicial map $F_\bullet^i : E_\bullet^i \rightarrow E_\bullet^{i+1}$, and instead appending f_{i+1} to the sequence of composable morphisms defines maps $E_p^i \rightarrow E_{p+1}^{i+1}$ which supply a homotopy between $|F_\bullet^i| : |E_\bullet^i| \rightarrow |E_\bullet^{i+1}|$ and the constant map to the vertex $f_{i+1} \in E_0^{i+1}$. Thus the mapping telescope $\text{hocolim}_{i \rightarrow \infty} |E_\bullet^i|$ is contractible.

The map $d_{p+1} : N_{p+1}\mathcal{C} \rightarrow N_p\mathcal{C}$ restricts to a fibration $E_p^i \rightarrow N_p\mathcal{C}$, by assumption (i), which induces a quasi-fibration

$$\text{hocolim}_{i \rightarrow \infty} E_p^i \longrightarrow N_p\mathcal{C}.$$

By [MS76, Propositions 4 and 5] this becomes a homology fibration after taking geometric realisation, using assumption (ii). Hence the inclusion of the fiber over $a \in BC$ into the homotopy fiber is a homology equivalence, and since the domain is contractible this inclusion may be identified with (A.5).

When the assumption is strengthened to (A.4) being an abelian equivalence, the conclusion may be strengthened to (A.5) being an abelian equivalence as well. (See [MP15] for a treatment of the main results of [MS76] in this setting.) Since the codomain has abelian fundamental group, this implies the map is acyclic. \square

REFERENCES

- [Ada71] J. F. Adams, *A variant of E. H. Brown's representability theorem*, *Topology* **10** (1971), 185–198. MR 0283788
- [BERW14] Boris Botvinnik, Johannes Ebert, and Oscar Randal-Williams, *Infinite loop spaces and positive scalar curvature*, arXiv:1411.7408, 2014.
- [BF81] E. Binz and H. R. Fischer, *The manifold of embeddings of a closed manifold*, *Differential geometric methods in mathematical physics* (Proc. Internat. Conf., Tech. Univ. Clausthal, Clausthal-Zellerfeld, 1978), *Lecture Notes in Phys.*, vol. 139, Springer, Berlin, 1981, With an appendix by P. Michor, pp. 310–329.
- [BM14] Alexander Berglund and Ib Madsen, *Rational homotopy theory of automorphisms of highly connected manifolds*, arXiv:1401.4096, 2014.
- [GMTW09] Søren Galatius, Ib Madsen, Ulrike Tillmann, and Michael Weiss, *The homotopy type of the cobordism category*, *Acta Math.* **202** (2009), no. 2, 195–239.
- [GRW10] Søren Galatius and Oscar Randal-Williams, *Monoids of moduli spaces of manifolds*, *Geom. Topol.* **14** (2010), no. 3, 1243–1302.
- [GRW14a] ———, *Homological stability for moduli spaces of high dimensional manifolds. I*, arXiv:1403.2334v2, 2014.
- [GRW14b] ———, *Stable moduli spaces of high-dimensional manifolds*, *Acta Math.* **212** (2014), no. 2, 257–377.
- [GS07] Paul Goerss and Kristen Schemmerhorn, *Model categories and simplicial methods*, *Interactions between homotopy theory and algebra*, *Contemp. Math.*, vol. 436, Amer. Math. Soc., Providence, RI, 2007, pp. 3–49.
- [HV04] Allen Hatcher and Karen Vogtmann, *Homology stability for outer automorphism groups of free groups*, *Algebr. Geom. Topol.* **4** (2004), 1253–1272. MR 2113904
- [Ker65] Michel A. Kervaire, *Le théorème de Barden–Mazur–Stallings*, *Comment. Math. Helv.* **40** (1965), 31–42.
- [MP15] Jeremy Miller and Martin Palmer, *A twisted homology fibration criterion and the twisted group-completion theorem*, *Q. J. Math.* **66** (2015), no. 1, 265–284.
- [MS76] Dusa McDuff and Graeme Segal, *Homology fibrations and the “group-completion” theorem*, *Invent. Math.* **31** (1975/76), no. 3, 279–284.
- [MW07] Ib Madsen and Michael Weiss, *The stable moduli space of Riemann surfaces: Mumford's conjecture*, *Ann. of Math. (2)* **165** (2007), no. 3, 843–941.
- [Qui83] Frank Quinn, *The stable topology of 4-manifolds*, *Topology Appl.* **15** (1983), no. 1, 71–77.
- [RW16] Oscar Randal-Williams, *Resolutions of moduli spaces and homological stability*, *J. Eur. Math. Soc.* **18** (2016), no. 1, 1–81.
- [Til97] Ulrike Tillmann, *On the homotopy of the stable mapping class group*, *Invent. Math.* **130** (1997), no. 2, 257–275.
- [Wah08] Nathalie Wahl, *Homological stability for the mapping class groups of non-orientable surfaces*, *Invent. Math.* **171** (2008), 389–424.
- [Wal65] C. T. C. Wall, *Finiteness conditions for CW-complexes*, *Ann. of Math. (2)* **81** (1965), 56–69.
- [Wei15] Michael Weiss, *Dalian notes on Pontryagin classes*, arXiv:1507.00153, 2015.

E-mail address: galatius@stanford.edu

DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY, STANFORD CA, 94305

E-mail address: o.randal-williams@dpmms.cam.ac.uk

CENTRE FOR MATHEMATICAL SCIENCES, WILBERFORCE ROAD, CAMBRIDGE CB3 0WB, UK