Acta Math., 212 (2014), 257–377
DOI: 10.1007/s11511-014-0112-7

Stable moduli spaces of high-dimensional manifolds

by

Søren Galatius

Stanford University
Stanford, CA, U.S.A.

Oscar Randal-Williams

University of Cambridge
Cambridge, U.K.

Dedicated to Ib Madsen on the occasion of his 70th birthday.

Contents

Galatius was partially supported by NSF grants DMS-0805843 and DMS-1105058 and the Clay Mathematics Institute. Randal-Williams was supported by the Herchel Smith Fund. Both authors were supported by ERC Advanced Grant No. 228082, and the Danish National Research Foundation through the Centre for Symmetry and Deformation.

1. Introduction and statement of results

For any smooth compact manifold W, the diffeomorphism group $\mathrm{Diff}(W)$ has a classifying space $B\mathrm{Diff}(W)$. This classifies smooth fibre bundles with fibre W, in the sense that for a smooth manifold B, there is a natural bijection between the set of isomorphism classes of smooth fibre bundles $\pi: E \to B$ with fibre W and the set $[B, B\mathrm{Diff}(W)]$ of homotopy classes of maps. The cohomology groups $H^k(B\mathrm{Diff}(W))$ therefore give characteristic classes of such bundles, and it is desirable to understand as much as possible about these cohomology groups. The difficulty of this question depends highly on W: it is essentially completely understood when the dimension of W is 0 or 1, and much effort has been devoted to understanding the case where the dimension of W is 2. Mumford [Mu] formulated a conjecture about the case where $W = \Sigma_g$ is an oriented surface of genus g, in the limit $g\to\infty$. If we let $\mathrm{Diff}(\Sigma_g, D^2)$ denote the diffeomorphism group which fixes some chosen disc $D^2\subset\Sigma_g$, Mumford's conjecture predicted an isomorphism

$$\varprojlim H^*(B\mathrm{Diff}(\Sigma_g, D^2);\mathbb{Q})\cong\mathbb{Q}[\kappa_1,\kappa_2,\kappa_3,\ldots]$$

for certain classes $\kappa_i\in H^{2i}(B\mathrm{Diff}(\Sigma_g, D^2))$. Mumford's conjecture was finally proved by Madsen and Weiss [MW] in a strengthened form.

The goal of the present paper is to prove analogues of the Madsen–Weiss theorem and Mumford's conjecture for manifolds of higher dimension. We have results for manifolds of any even dimension greater than 4. As an interesting special case of our results, we completely determine the stable rational cohomology ring

$$\varprojlim H^*(B\mathrm{Diff}(W_g, D^{2n});\mathbb{Q}),$$

where $W_g=\#^g(S^n\times S^n)=g(S^n\times S^n)$ denotes the connected sum of g copies of $S^n\times S^n$. To state our result, we recall that for each characteristic class of oriented $2n$-dimensional vector bundles $c\in H^{2n+k}(B\mathrm{SO}(2n))$, we can define the associated *generalised Mumford–Morita–Miller* class of a smooth fibre bundle $\pi: E\to B$ with oriented $2n$-dimensional fibres as

$$\kappa_c(E)=\pi_!(c(T_\pi E))\in H^k(B),$$

where $T_\pi E$ is the fibrewise tangent bundle of π (when π is a submersion of smooth manifolds, this is simply the kernel of $D\pi: TE\to\pi^*TB$). When the fibre is taken to be W_g, there is a corresponding universal class $\kappa_c\in H^k(B\mathrm{Diff}(W_g, D^{2n}))$ which for $k>0$ is compatible with increasing g.

THEOREM 1.1. *Let $2n>4$ and let $\mathcal{B}\subset H^*(B\mathrm{SO}(2n);\mathbb{Q})$ be the set of monomials in the classes $e, p_{n-1}, p_{n-2}, \ldots, p_{\lceil (n+1)/4\rceil}$ of total degree greater than $2n$. Then the natural map*

$$\mathbb{Q}[\kappa_c: c\in\mathcal{B}]\longrightarrow\varprojlim H^*(B\mathrm{Diff}(W_g, D^{2n});\mathbb{Q})$$

is an isomorphism.

The strengthened form of Mumford's conjecture proved by Madsen and Weiss states that a certain map

$$\operatorname*{hocolim}_{g\to\infty} B\mathrm{Diff}(\Sigma_g, D^2)\longrightarrow\Omega_0^\infty\,\mathrm{MTSO}(2)$$

induces an isomorphism in integral homology. We will prove a similar homotopy-theoretic strengthening of Theorem 1.1, which also applies to more general manifolds.

1.1. Definitions and recollections

To state the main results in their general form, we recall the following definitions.

1.1.1. Classifying spaces

We shall use the following model for the classifying space $B\mathrm{Diff}(W,\partial W)$ of the topological group of diffeomorphisms of a compact manifold W, restricting to the identity on a neighbourhood of ∂W. We first pick an embedding $\partial W \hookrightarrow \{0\}\times\mathbb{R}^\infty$ and let $\mathrm{Emb}^\partial(W,(-\infty,0]\times\mathbb{R}^\infty)$ denote the space of all extensions to an embedding of W (required to be standard on a collar neighbourhood of ∂W). We then let

$$B\mathrm{Diff}(W,\partial W) = \mathrm{Emb}^\partial(W,(-\infty,0]\times\mathbb{R}^\infty)/\mathrm{Diff}(W,\partial W)$$

be the orbit space. If W is closed and $A\subset W$ is a compact codimension-0 submanifold, we write $B\mathrm{Diff}(W,A)=B\mathrm{Diff}(W\setminus\mathrm{int}(A),\partial A)$. The construction of $B\mathrm{Diff}(W,\partial W)$ has the following naturality property: any inclusion $W\subset W'$ of a codimension-0 submanifold induces a continuous map $B\mathrm{Diff}(W,\partial W)\to B\mathrm{Diff}(W',\partial W')$, well defined up to homotopy. (On the point-set level it depends on a choice of embedding of the cobordism $W'\setminus\mathrm{int}(W)$ into $[0,1]\times\mathbb{R}^\infty$.) For example, a choice of inclusion $W_g\setminus\mathrm{int}(D^{2n})\to W_{g+1}$ induces a map $B\mathrm{Diff}(W_g,D^{2n})\to B\mathrm{Diff}(W_{g+1},D^{2n})$; these define the inverse system in Theorem 1.1.

1.1.2. Thom spectra

For any space B and any map $\theta\colon B\to B\mathrm{O}(d)$, where $B\mathrm{O}(d)=\mathrm{Gr}_d(\mathbb{R}^\infty)$, there is a Thom spectrum $\mathrm{MT}\theta=B^{-\theta}$ constructed in the following way: First, we let

$$B(\mathbb{R}^n) = \theta^{-1}(\mathrm{Gr}_d(\mathbb{R}^n)).$$

The Grassmannian $\mathrm{Gr}_d(\mathbb{R}^n)$ carries an $(n-d)$-dimensional vector bundle $\gamma_n^\perp$, the orthogonal complement of the tautological bundle. Then the nth space of the spectrum $\mathrm{MT}\theta$ is the Thom space $B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp}$. The associated infinite loop space is the direct limit

$$\Omega^\infty\,\mathrm{MT}\theta = \operatorname*{colim}_{n\to\infty}\Omega^n\big(B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp}\big),$$

and we shall write $\Omega^\infty_0\,\mathrm{MT}\theta$ for the basepoint component. The rational cohomology of this space is easy to describe; in the case where the bundle classified by θ is oriented, it is as follows: for each $c\in H^{d+k}(B)$, there is a corresponding "generalised Mumford–Morita–Miller class" $\kappa_c\in H^k(\Omega^\infty\,\mathrm{MT}\theta)$, and $H^*(\Omega^\infty_0\,\mathrm{MT}\theta;\mathbb{Q})$ is the free graded-commutative algebra on the classes κ_c, where c runs through a basis for the vector space $H^{>d}(B;\mathbb{Q})$. We describe the general case in §2.5.

1.1.3. Moore–Postnikov towers

Let us first recall that a topological space X is said to be *n-connected*, for an integer $n\geqslant -1$, if any map $S^k\to X$ admits an extension to $D^{k+1}\to X$ for $-1\leqslant k\leqslant n$. For non-empty X this is equivalent to the vanishing of $\pi_k(X,x)$ for $0\leqslant k\leqslant n$ and all $x\in X$. A map $f\colon X\to Y$ is *n-connected* if all its homotopy fibres are $(n-1)$-connected. If $n\geqslant 0$ and $\pi_0(X)\to\pi_0(Y)$ is surjective, this is equivalent to the vanishing of the relative homotopy groups $\pi_k(M_f,X,x)$ for all $k\leqslant n$ and all $x\in X$. Here M_f denotes the mapping cylinder of f; henceforth we shall write $\pi_k(Y,X,x)$ for $\pi_k(M_f,X,x)$ when the map f is clear. If $A\subset X$ is a subspace, we shall say that the pair (X,A) is *n-connected* if the inclusion map $A\to X$ is n-connected.

Similarly, a space X is said to be *n-co-connected* (or an *$(n-1)$-type*) if any map $S^k\to X$ admits an extension to D^{k+1} for $k\geqslant n$, and a map $X\to Y$ is *n-co-connected* if all homotopy fibres are n-co-connected.

It is well known (cf. [H, Theorem 4.71]) that for any map $f\colon A\to X$ of spaces and any $n\geqslant 0$, there is a factorisation $f\colon A\xrightarrow{g}B\xrightarrow{h}X$ with the property that g is n-connected and h is n-co-connected, and moreover such a factorisation is unique up to weak homotopy equivalence. This is the nth stage of the Moore–Postnikov tower for the map $A\to X$. It will be important for us that this factorisation can be made strictly functorial in the map $f\colon A\to X$, and we briefly recall one way of achieving this. Define factorisations $f=h_k\circ g_k\colon A\to B^k\to X$ for $k\geqslant n$ by setting $B^n=A$ and inductively letting B^{k+1} be the relative CW complex obtained from B^k by attaching one $(k+1)$-cell for each commutative square of the form

$$\begin{array}{ccc} S^k & \longrightarrow & D^{k+1} \\ \big\downarrow & & \big\downarrow \\ B^k & \xrightarrow{h_k} & X. \end{array}$$

Then $h_k\colon B^k\to X$ extends canonically to $h_{k+1}\colon B^{k+1}\to X$, and we may let $B=\bigcup_{k\geqslant n}B^k$.

In the case where $A=\{x\}$ is a point, $B=X\langle n\rangle\to X$ is the n-connective cover of the based space (X,x), characterised by the property that $\pi_i(X\langle n\rangle,x)=0$ for $0\leqslant i\leqslant n$, and that $\pi_i(X\langle n\rangle,x)\to\pi_i(X,x)$ is an isomorphism for $i>n$. (Some authors write $X\langle n+1\rangle$ for what we denote $X\langle n\rangle$.)

1.2. Connected sums of copies of $S^n\times S^n$

We can now state our homotopy-theoretic version of Theorem 1.1, generalising Madsen–Weiss' theorem to dimension $2n$ (recall that we assume $2n>4$ throughout). As before, we write $W_g=\#^g(S^n\times S^n)=g(S^n\times S^n)$ for the connected sum of g copies of $S^n\times S^n$.

If we pick a disc $D^{2n} \subset W_g$, there is a classifying space $B\mathrm{Diff}(W_g, D^{2n})$ and there are maps $B\mathrm{Diff}(W_g, D^{2n}) \to B\mathrm{Diff}(W_{g+1}, D^{2n})$ induced by taking connected sum with one more copy of $S^n \times S^n$. Let $\theta^n: B\mathrm{O}(2n)\langle n\rangle \to B\mathrm{O}(2n)$ be the n-connective cover, and $\mathrm{MT}\theta^n$ be the associated Thom spectrum. Let us also say that a continuous map is a *homology equivalence* if it induces an isomorphism in integral homology (and hence for homology or cohomology with coefficients in any spectrum, at least after replacing both spaces by weakly equivalent CW complexes).

THEOREM 1.2. *Let $2n>4$. There is a homology equivalence*

$$\operatorname*{hocolim}_{g\to\infty} B\mathrm{Diff}(W_g, D^{2n}) \longrightarrow \Omega_0^\infty \mathrm{MT}\theta^n.$$

More generally, if W is any $(n-1)$-connected closed $2n$-manifold which is parallelisable in the complement of a point, there is a homology equivalence

$$\operatorname*{hocolim}_{g\to\infty} B\mathrm{Diff}(W\# W_g, D^{2n}) \longrightarrow \Omega_0^\infty \mathrm{MT}\theta^n.$$

It is easy to deduce Theorem 1.1 from Theorem 1.2. In [GRW2] we proved that the maps $B\mathrm{Diff}(W_g, D^{2n}) \to B\mathrm{Diff}(W_{g+1}, D^{2n})$ induce isomorphisms in integral homology up to degree $\lfloor \frac{1}{2}(g-4)\rfloor$ (cf. also [BM]). Thus, Theorem 1.2 also determines the homology and cohomology of $B\mathrm{Diff}(W_g, D^{2n})$ in this range.

1.3. The moduli space of highly connected null-bordisms

The determination, in Theorems 1.1 and 1.2, of the stable homology and cohomology of the space $B\mathrm{Diff}(W\# g(S^n\times S^n), D^{2n})$ is a special case of Theorem 1.8 below, in which we determine the stable homology of $B\mathrm{Diff}(W)$ for more general manifolds W. We also consider manifolds equipped with an additional *tangential structure*, defined as follows.

Definition 1.3. Let $\theta: B\to B\mathrm{O}(2n)$ be a map. A *θ-structure* on a $2n$-dimensional manifold W is a bundle map $\ell: TW\to\theta^*\gamma$, i.e. a fibrewise linear isomorphism. Such a pair (W,ℓ) will be called a *θ-manifold*. A *θ-structure* on a $(2n-1)$-dimensional manifold M is a bundle map $\varepsilon^1\oplus TM\to\theta^*\gamma$. If ℓ is a θ-structure on W, the *induced* structure on ∂W is obtained by composing with a certain isomorphism $\varepsilon^1\oplus T(\partial W)\to TW|_{\partial W}$. In fact, there are two such isomorphisms: One comes from a collar $[0,1)\times\partial W\to W$ of ∂W. Differentiating this gives an isomorphism $\varepsilon^1\oplus T(\partial W)\to TW|_{\partial W}$, and the resulting θ-structure on ∂W will be called the *incoming* restriction. Another comes from a collar $(-1,0]\times\partial W\to W$; this is the *outgoing* restriction. When W is a cobordism, i.e. a compact manifold together with a partition $\partial W=\partial_{\mathrm{in}}W\amalg\partial_{\mathrm{out}}W$ of its boundary, we will generally use the incoming restriction to induce a θ-structure on the incoming boundary of W and the outgoing restriction on the outgoing boundary.

Let $\ell_0: TW|_{\partial W} \to \theta^*\gamma$ be a θ-structure on ∂W, and $\mathrm{Bun}^\partial(TW, \theta^*\gamma; \ell_0)$ denote the space of all bundle maps $\ell: TW \to \theta^*\gamma$ which restrict to ℓ_0 over ∂W, equipped with the compact-open topology. The group $\mathrm{Diff}(W, \partial W)$ of diffeomorphisms of W which restrict to the identity near ∂W acts on $\mathrm{Bun}^\partial(TW, \theta^*\gamma; \ell_0)$ by precomposing a bundle map with the differential of a diffeomorphism.

The most general case of our theorem concerns the *moduli space of highly connected null-bordisms*, defined as follows.

Definition 1.4. Let $P \subset \mathbb{R}^\infty$ be a closed $(2n-1)$-dimensional manifold with θ-structure $\ell_P: \varepsilon^1 \oplus TP \to \theta^*\gamma$. A *null-bordism* is a pair (W, ℓ_W), where $W \subset (-\infty, 0] \times \mathbb{R}^\infty$ is a compact manifold with $\partial W = \{0\} \times P$ and $(-\varepsilon, 0] \times P \subset W$ for some $\varepsilon > 0$, and $\ell_W: TW \to \theta^*\gamma$ is a θ-structure satisfying $\ell_W|_{\partial W} = \ell_P$. A null-bordism (W, ℓ_W) is *highly connected* if (W, P) is $(n-1)$-connected, and the *moduli space of highly connected null-bordisms* is the set $\mathcal{N}^\theta(P, \ell_P)$ of all highly connected null-bordisms of (P, ℓ_P). It is topologised as

$$\coprod_W (\mathrm{Emb}^\partial(W, (-\infty, 0] \times \mathbb{R}^\infty) \times \mathrm{Bun}^\partial(TW, \theta^*\gamma; \ell_P))/\mathrm{Diff}(W, \partial W), \tag{1.1}$$

where the disjoint union is over compact manifolds W with $\partial W = P$ for which (W, P) is $(n-1)$-connected, one of each diffeomorphism class.

If $K \subset [0, 1] \times \mathbb{R}^\infty$ is a cobordism with collared boundary $\partial K = (\{0\} \times P_0) \cup (\{1\} \times P_1)$ we say that K is *highly connected* if each pair $(K, \{i\} \times P_i)$ is $(n-1)$-connected. If K is equipped with a θ-structure ℓ_K restricting to ℓ_0 and ℓ_1 on the boundaries, then there is an induced map $\mathcal{N}^\theta(P_0, \ell_0) \to \mathcal{N}^\theta(P_1, \ell_1)$ defined by taking union with K and subtracting 1 from the first coordinate. (If (W, P_0) is $(n-1)$-connected then $(W \cup K, K)$ is too, and it follows from the long exact sequence of the triple $(W \cup K, K, P_1)$ that $(W \cup K, P_1)$ is also $(n-1)$-connected.)

This moduli space classifies smooth families of null-bordisms of P, in the sense that if B is a smooth manifold without boundary, there is a natural bijection between the set of homotopy classes $[B, \mathcal{N}^\theta(P, \ell_P)]$ and the set of equivalence classes of triples (π, φ, ℓ), where $\pi: E \to B$ is a proper submersion (i.e. smooth fibre bundle), φ is a diffeomorphism $\partial E \cong B \times P$ over B, such that $(E, \partial E)$ is $(n-1)$-connected, and ℓ is a θ-structure on the fibrewise tangent bundle $T_\pi E = \mathrm{Ker}(D\pi)$ restricting to ℓ_P on the boundary of each fibre.

Let us also introduce notation for each of the disjoint summands in (1.1).

Definition 1.5. Let W be a compact $2n$-dimensional manifold, and $\ell_0: TW|_{\partial W} \to \theta^*\gamma$ be a θ-structure on ∂W. We shall write

$$B\mathrm{Diff}^\theta(W; \ell_0) = (E\mathrm{Diff}(W, \partial W) \times \mathrm{Bun}^\partial(TW, \theta^*\gamma; \ell_0))/\mathrm{Diff}(W, \partial W)$$

for the homotopy orbit space of the action of $\mathrm{Diff}(W,\partial W)$ on $\mathrm{Bun}^\partial(TW,\theta^*\gamma;\ell_0)$. If $\ell\colon TW\to\theta^*\gamma$ is a particular extension, we shall write $B\mathrm{Diff}^\theta(W;\ell_0)_\ell\subset B\mathrm{Diff}^\theta(W;\ell_0)$ for the path component containing ℓ.

Using the model $E\mathrm{Diff}(W,\partial W)=\mathrm{Emb}^\partial(W,(-\infty,0]\times\mathbb{R}^\infty)$, we have the homeomorphism

$$\mathcal{N}^\theta(P,\ell_P)=\coprod_W B\mathrm{Diff}^\theta(W;\ell_P).$$

Definition 1.6. A tangential structure $\theta\colon B\to B\mathrm{O}(2n)$ is called *spherical* if any θ-structure on the lower hemisphere $\partial_-D^{2n+1}\subset\partial D^{2n+1}$ extends to some θ-structure on the whole sphere. (If B is path connected, this is equivalent to the sphere S^{2n} admitting a θ-structure.)

Most of the usual structures, for example SO, Spin, Spin^c, etc. are spherical, but some are not, e.g. framings. Theorem 1.8 below determines the homology of $\mathcal{N}^\theta(P,\ell_P)$ after stabilising with cobordisms in the (P,ℓ_P)-variable. The following definition makes the stabilisation procedure precise.

Definition 1.7. Let $\theta\colon B\to B\mathrm{O}(2n)$ be spherical, and $K\subset[0,\infty)\times\mathbb{R}^\infty$ be a submanifold with θ-structure ℓ_K. For $A\subset[0,\infty)$, we let $(K|_A,\ell_K|_A)$ denote the pair

$$(K\cap x_1^{-1}(A),\ell_K|_{K\cap x_1^{-1}(A)}),$$

which will again be a θ-manifold when A is an interval whose endpoints are regular values of $x_1\colon K\to[0,\infty)$. We shall assume that each natural number n is a regular value of x_1. If $M\subset\mathbb{R}^\infty$ is the manifold such that $K|_n=\{n\}\times M$ we also impose the existence of a cylindrical collar, i.e. an open neighbourhood $U\subset[0,\infty)$ of n such that $K|_U=U\times M$.

(i) Let $W\subset[0,1]\times\mathbb{R}^\infty$ be a cobordism with θ-structure ℓ_W, and suppose that $(W|_0,\ell_W|_0)=(K|_0,\ell_K|_0)$. We say that (K,ℓ_K) *absorbs* (W,ℓ_W) if there exists an embedding $j\colon W\to K$ which is the identity on $W|_0=K|_0$, such that $\ell_K\circ Dj\colon TW\to\theta^*\gamma$ is homotopic to ℓ_W relative to $W|_0$. That $K|_{[i,\infty)}$ absorbs a θ-bordism $W\subset[i,i+1]\times\mathbb{R}^\infty$ is defined similarly.

(ii) We say that (K,ℓ_K) is a *universal* θ*-end* if for each integer $i\geqslant 0$, $K|_{[i,i+1]}$ is a highly connected cobordism and $K|_{[i,\infty)}$ absorbs W for any highly connected cobordism $W\subset[i,i+1]\times\mathbb{R}^\infty$ with θ-structure ℓ_W such that $(W|_i,\ell_W|_i)=(K|_i,\ell_K|_i)$.

For example, in dimension 2 with $\theta=\mathrm{Id}_{BO(2)}$, we can construct a universal θ-end by letting each $K|_{[i,i+1]}$ be diffeomorphic to $\mathbb{R}P^2$ with two discs removed. For $\theta=\theta^n\colon B\mathrm{O}(2n)\langle n\rangle\to B\mathrm{O}(2n)$, a universal θ-end can be constructed by letting each $K|_{[i,i+1]}$ be diffeomorphic to $S^n\times S^n$ with two discs removed. In many other cases, a universal

θ-end K can be constructed as the infinite iteration of a single self-bordism $K|_{[0,1]}$. In particular, this will be the case in the examples in §1.5 below.

As we shall see, universal θ-ends are unique up to isomorphism in the following sense. If (K, ℓ_K) and (K', ℓ'_K) are two universal θ-ends with $K|_0 = K'|_0$, then there exists a diffeomorphism $K \to K'$ preserving θ-structure up to homotopy, relative to $K|_0$. More generally, given a highly connected cobordism (W, ℓ_W) from $K|_0$ to $K'|_0$, there exists a similar diffeomorphism from K to $W \cup K'$.

THEOREM 1.8. *Let $2n > 4$ and let $\theta: B \to BO(2n)$ be spherical. Let (K, ℓ_K) be a universal θ-end with $\mathcal{N}^\theta(K|_0, \ell_K|_0) \neq \varnothing$. Then there is a homology equivalence*

$$\operatorname*{hocolim}_{i\to\infty} \mathcal{N}^\theta(K|_i, \ell_K|_i) \longrightarrow \Omega^\infty \mathrm{MT}\theta',$$

where

$$\theta': B' \longrightarrow B \xrightarrow{\theta} BO(2n)$$

is the n-th stage of the Moore–Postnikov tower for $\ell_K: K \to B$.

The property of being a universal θ-end can often be checked in practice, using the following addendum, as it is essentially a homotopical property.

ADDENDUM 1.9. *Let $\theta: B \to BO(2n)$ be spherical, let $K \subset [0, \infty) \times \mathbb{R}^\infty$ be a submanifold such that $K|_{[i,i+1]}$ is a highly connected cobordism for each integer i, and let ℓ_K be a θ-structure on K. Then (K, ℓ_K) is a universal θ-end if and only if the following conditions hold*:

(i) *For each integer i, the map $\pi_n(K|_{[i,\infty)}) \to \pi_n(B)$ is surjective, for all basepoints in K.*

(ii) *For each integer i, the map $\pi_{n-1}(K|_{[i,\infty)}) \to \pi_{n-1}(B)$ is injective, for all basepoints in K.*

(iii) *For each integer i, each path component of $K|_{[i,\infty)}$ contains a submanifold diffeomorphic to $(S^n \times S^n) \setminus \operatorname{int}(D^{2n})$, which in addition has null-homotopic structure map to B.*

Remark 1.10. The maps in all the theorems above are induced by the Pontryagin–Thom construction. We shall briefly explain this in the setting of Theorem 1.8, after replacing $\mathcal{N}^\theta(P, \ell_P)$ by a weakly equivalent space, and refer the reader to [MT, §2.3] for further details. First we say that a submanifold $W \subset (-\infty, 0] \times \mathbb{R}^{q-1}$ with collared boundary is *fatly embedded* if the canonical map from the normal bundle νW to $\mathbb{R}^q$ restricts to an embedding of the disc bundle into $(-\infty, 0] \times \mathbb{R}^{q-1}$. In that case the Pontryagin–Thom collapse construction gives a continuous map from $[-\infty, 0] \wedge S^{q-1}$ to the Thom space of νW. Secondly we replace $\theta'\colon B'\to B\mathrm{O}(2n)$ by a fibration, and redefine $\mathcal{N}^\theta(P,\ell_P)$ as a space of pairs (W,ℓ'_W), where $W\subset(-\infty,0]\times\mathbb{R}^\infty$ is a fatly embedded submanifold, collared near $\partial W=\{0\}\times P$, and $\ell'_W\colon W\to B'$ is a continuous map such that $\theta'\circ\ell'_W\colon W\to B\mathrm{O}(2n)=\mathrm{Gr}_{2n}(\mathbb{R}^\infty)$ is *equal* to the Gauss map and whose restriction to ∂W is equal to a specified map $\ell'_P\colon P\to B'$ lifting ℓ_P. There is a forgetful map from the space of such pairs to the space in Definition 1.4, and standard homotopy-theoretic methods imply that it is a weak equivalence (see Corollary 7.17). If $P\subset\mathbb{R}^{q-1}\subset\mathbb{R}^\infty$, the Pontryagin–Thom construction (composed with ℓ'_P) gives a point

$$\alpha(P,\ell'_P)\in\Omega^{q-1}(B'(\mathbb{R}^q)^{(\theta')^*\gamma_q^\perp})\subset\Omega^{\infty-1}\,\mathrm{MT}\theta',$$

and if $(W,\ell'_W)\in\mathcal{N}^\theta(P,\ell_P)$ has $W\subset(-\infty,0]\times\mathbb{R}^{q-1}$, it gives a path

$$\alpha(W,\ell'_W)\colon[-\infty,0]\longrightarrow\Omega^{q-1}(B'(\mathbb{R}^q)^{(\theta')^*\gamma_q^\perp})\subset\Omega^{\infty-1}\,\mathrm{MT}\theta',$$

starting at the basepoint and ending at $\alpha(P,\ell_P)$. If we write $\Omega_{\varnothing,P}\Omega^{\infty-1}\,\mathrm{MT}\theta'$ for the space of all such paths, then this construction determines a map

$$\alpha\colon\mathcal{N}^\theta(P,\ell_P)\longrightarrow\Omega_{\varnothing,P}\Omega^{\infty-1}\,\mathrm{MT}\theta'.$$

The non-compact manifold $K\subset[0,\infty)\times\mathbb{R}^\infty$ admits a homotopically unique θ'-structure ℓ'_K lifting its θ-structure and extending the canonical θ'-structure on $P=K|_0$. The Pontryagin–Thom construction applied to each cobordism $K|_{[i,i+1]}$ then gives a path $\alpha(K|_{[i,i+1]},\ell'_K|_{[i,i+1]})\colon[0,1]\to\Omega^{\infty-1}\,\mathrm{MT}\theta'$ from $\alpha(K|_i,\ell'_K|_i)$ to $\alpha(K|_{i+1},\ell'_K|_{i+1})$. The maps α then give a map of direct systems, which on direct limits is

$$\operatorname*{hocolim}_{i\to\infty}\mathcal{N}^\theta(K|_i,\ell_K|_i)\longrightarrow\operatorname*{hocolim}_{i\to\infty}\Omega_{\varnothing,K|_i}\Omega^{\infty-1}\,\mathrm{MT}\theta'.$$

Finally, the maps in the direct system on the right-hand side are all homotopy equivalences, so the direct limit is equivalent to its zeroth term, and a *choice* of path from $\alpha(K|_0,\ell'_K|_0)$ to $\varnothing$ identifies the zeroth term with $\Omega^\infty\,\mathrm{MT}\theta'$.

Remark 1.11. It is often useful to consider the homology equivalence in Theorem 1.8 one path component at a time, so we spell out the resulting statement using the notation of Definition 1.5. Any path component of the infinite loop space $\Omega^\infty\,\mathrm{MT}\theta'$ is homotopy equivalent to the basepoint component $\Omega_0^\infty\,\mathrm{MT}\theta'$. On the left-hand side of the homology equivalence, the path component of an element $(W,\ell_W)\in\mathcal{N}^\theta(K|_0,\ell_K|_0)$ is the homotopy colimit of the spaces

$$B\mathrm{Diff}^\theta(W\cup K|_{[0,i]};\ell_i)_{\ell_W\cup\ell_K|_{[0,i]}}.$$

Conversely, given any triple (W, K, ℓ), where $K \subset [0,\infty)\times\mathbb{R}^\infty$ is a non-compact submanifold such that the subset $K|_{[i,i+1]}\subset K$ is a highly connected cobordism for each integer $i\geqslant 0$, $W\subset(-\infty,0]\times\mathbb{R}^\infty$ is a compact manifold with collared boundary $\partial W = K|_0$ such that $(W,\partial W)$ is $(n-1)$-connected, and $\ell\colon T(W\cup K)\to\theta^*\gamma$ is a bundle map, the Pontryagin–Thom construction described in Remark 1.10 provides a map

$$\operatorname*{hocolim}_{i\to\infty} B\mathrm{Diff}^\theta(W\cup K|_{[0,i]};\ell_i)_\ell \longrightarrow \Omega_0^\infty \mathrm{MT}\theta', \tag{1.2}$$

where $\theta'\colon B'\to B\xrightarrow{\theta} B\mathrm{O}(2n)$ is obtained from the nth Moore–Postnikov stage of the underlying map $W\cup K\to B$. If it is the case that K is a universal θ'-end then $\ell'_K\colon K\to B'$ is n-connected (cf. Addendum 1.9), so θ' is also obtained from the nth Moore–Postnikov stage of $\ell_K\colon K\to B$, and so by Theorem 1.8 the map (1.2) is a homology isomorphism. In particular, Theorem 1.2 can be deduced this way: If we let each $K|_{[i,i+1]}$ be diffeomorphic to $([0,1]\times S^{2n-1})\#(S^n\times S^n)$ and let $\theta=\mathrm{Id}_{B\mathrm{O}(2n)}$, then $\theta'=\theta^n\colon B\mathrm{O}(2n)\langle n\rangle\to B\mathrm{O}(2n)$, and K is a universal θ^n-end. Similarly, all examples in §1.5 below arise in this way.

Let us also remark that the homotopy colimit (1.2) may be replaced by the strict colimit $B\mathrm{Diff}_c^\theta(W\cup K;\ell)$, defined by

$$B\mathrm{Diff}_c^\theta(W\cup K;\ell) = (E\mathrm{Diff}_c(W\cup K)\times\mathrm{Bun}_c(T(W\cup K),\theta^*\gamma;\ell))/\mathrm{Diff}_c(W\cup K),$$

where $\mathrm{Diff}_c(W\cup K)$ is the topological group of compactly supported diffeomorphisms of the non-compact manifold $W\cup K$, and $\mathrm{Bun}_c(T(W\cup K),\theta^*\gamma;\ell)$ is the space of bundle maps which agree outside of a compact subset of $W\cup K$ with ℓ.

1.4. Algebraic localisation

There is one final algebraic version of our main theorem. Fix P, a closed $(2n-1)$-manifold with θ-structure $\ell_P\colon\varepsilon^1\oplus TP\to\theta^*\gamma$. As explained in Definition 1.4, a cobordism (K,ℓ_K) from (P,ℓ_P) to itself with $K\subset[0,1]\times\mathbb{R}^\infty$, which is $(n-1)$-connected with respect to both boundaries, gives a self-map of $\mathcal{N}^\theta(P,\ell_P)$ defined by $W\mapsto(W\cup_P K)-e_1$. We shall write $\mathcal{K}_0$ for the set of isomorphism classes of such (K,ℓ_K), where we identify (K,ℓ_K) with $(K',\ell_{K'})$ if there is a diffeomorphism $\varphi\colon K\to K'$ which is the identity near ∂K such that $\varphi^*\ell_{K'}$ is homotopic to ℓ_K relative to ∂K. It is clear that the homotopy class of the self-map of $\mathcal{N}^\theta(P,\ell_P)$ induced by (K,ℓ_K) depends only on the isomorphism class of (K,ℓ_K), and we get an action of the non-commutative monoid $\mathcal{K}_0$ on $H_*(\mathcal{N}^\theta(P,\ell_P))$. Our theorem determines the algebraic localisation

$$H_*(\mathcal{N}^\theta(P,\ell_P))[\mathcal{K}^{-1}]$$

at a certain commutative submonoid $\mathcal{K}\subset\mathcal{K}_0$ which we now describe.

We say that a θ-cobordism $K: P\rightsquigarrow P$ has *support* in a closed subset $A\subset P$ if it contains $[0,1]\times(P\setminus A): P\setminus A\rightsquigarrow P\setminus A$ as a sub-cobordism with the product θ-structure. We let $\mathcal{K}\subset\mathcal{K}_0$ consist of those elements which admit a representative with support in a regular neighbourhood of a simplicial complex of dimension at most $n-1$ inside P, and prove the following lemma.

LEMMA 1.12. *The subset $\mathcal{K}\subset\mathcal{K}_0$ is a commutative submonoid.*

We may localise the $\mathbb{Z}[\mathcal{K}]$-module $H_*(\mathcal{N}^\theta(P,\ell_P))$ at any submonoid $\mathcal{L}\subset\mathcal{K}$. The content of Theorem 1.13 below is an isomorphism

$$H_*(\mathcal{N}^\theta(P,\ell_P))[\mathcal{L}^{-1}]\cong H_*(\Omega^\infty \mathrm{MT}\theta')$$

under certain conditions, where $\theta': B'\to B\xrightarrow{\theta} B\mathrm{O}(2n)$ is the $(n-1)$-st stage of the Moore–Postnikov tower for $\ell_P: P\to B$. (Note that in Theorem 1.8 we used the nth stage instead.) To describe the isomorphism explicitly, recall that in Remark 1.10 we described a map

$$\mathcal{N}^\theta(P,\ell_P)\longrightarrow \Omega^\infty \mathrm{MT}\theta',$$

compatible with gluing highly connected cobordisms of (P,ℓ_P) equipped with θ'-structures, and hence the induced map

$$H_*(\mathcal{N}^\theta(P,\ell_P))\longrightarrow H_*(\Omega^\infty \mathrm{MT}\theta') \tag{1.3}$$

is a map of $\mathbb{Z}[\mathcal{K}']$-modules, where the monoid $\mathcal{K}'$ is defined like $\mathcal{K}$ but using θ' instead of θ. An obstruction-theoretic argument, which we explain in more detail in §7.6, shows that the natural map $\mathcal{K}'\to\mathcal{K}$ is a bijection, so (1.3) is naturally a homomorphism of $\mathbb{Z}[\mathcal{K}]$-modules.

THEOREM 1.13. *Let $2n>4$ and let $\theta: B\to B\mathrm{O}(2n)$ be spherical. Let P be a closed $(2n-1)$-manifold with θ-structure $\ell_P: \varepsilon^1\oplus TP\to\theta^*\gamma$, such that $\mathcal{N}^\theta(P,\ell_P)$ is non-empty. Then the morphism (1.3) induces an isomorphism*

$$H_*(\mathcal{N}^\theta(P,\ell_P))[\mathcal{K}^{-1}]\longrightarrow H_*(\Omega^\infty \mathrm{MT}\theta').$$

Furthermore, localisation at a submonoid $\mathcal{L}\subset\mathcal{K}$ agrees with localisation at $\mathcal{K}$, provided $\mathcal{L}$ satisfies the following conditions:

(i) *The group $\pi_n(B)$ is generated by the subgroups $\mathrm{Im}(\pi_n(K)\to\pi_n(B))$, $K\in\mathcal{L}$.*

(ii) *The subgroup of $\pi_{n-1}(P)$ generated by $\mathrm{Ker}(\pi_{n-1}(P)\to\pi_{n-1}(K))$, $K\in\mathcal{L}$, contains $\mathrm{Ker}(\pi_{n-1}(P)\to\pi_{n-1}(B))$.*

(iii) *There is an element of $\mathcal{L}$ containing a submanifold diffeomorphic to*

$$(S^n\times S^n)\setminus\mathrm{int}(D^{2n}).$$

(*There is a bijection* $\pi_0(P)=\pi_0(K)$, *and if* P *is not connected, conditions* (i)–(iii) *are required to hold for each path component of* P.)

Applying the functor $\mathrm{Hom}_{\mathbb{Z}[\mathcal{K}]}(-,\mathbb{Q})$ to both sides of the isomorphism in the theorem identifies the subring of $H^*(\mathcal{N}^\theta(P,\ell_P);\mathbb{Q})$ consisting of $\mathcal{K}$-invariants with

$$H^*(\Omega_0^\infty \mathrm{MT}\theta';\mathbb{Q}).$$

Observing that these classes are also invariant under the larger monoid $\mathcal{K}_0$, we deduce the isomorphism

$$H^*(\mathcal{N}^\theta(P,\ell_P);\mathbb{Q})^{\mathcal{K}_0}\cong H^*(\Omega_0^\infty \mathrm{MT}\theta';\mathbb{Q}).$$

The left-hand side can be interpreted as characteristic classes of certain bundles, invariant under fibrewise gluing of trivial bundles.

1.5. Examples and applications

Recall that the connective cover $B\mathrm{O}(d)\langle k\rangle$ is $B\mathrm{SO}(d)$ if $k=1$, $B\mathrm{Spin}(d)$ if $k=2,3$, and is often called $B\mathrm{String}(d)$ if $k=4,5,6,7$. We will write $\mathrm{MTSO}(d)$, $\mathrm{MTSpin}(d)$ and $\mathrm{MTString}(d)$ for the corresponding Thom spectra. As special cases of Theorem 1.8 we have the following maps, which become homology equivalences in the limit $g\to\infty$. All are deduced from Theorem 1.8 and Addendum 1.9 as in Remark 1.11, with $\theta=\mathrm{Id}_{B\mathrm{O}(2n)}$:

$$\begin{aligned}
B\mathrm{Diff}(g(S^3\times S^3),D^6)&\longrightarrow \Omega_0^\infty \mathrm{MTSpin}(6),\\
B\mathrm{Diff}(g(\mathbb{H}P^2\#\overline{\mathbb{H}P}^2),D^8)&\longrightarrow \Omega_0^\infty \mathrm{MTSpin}(8),\\
B\mathrm{Diff}(g(S^4\times S^4),D^8)&\longrightarrow \Omega_0^\infty \mathrm{MTString}(8),\\
B\mathrm{Diff}(g(S^5\times S^5),D^{10})&\longrightarrow \Omega_0^\infty \mathrm{MTString}(10),\\
B\mathrm{Diff}(g(S^6\times S^6),D^{12})&\longrightarrow \Omega_0^\infty \mathrm{MTString}(12),\\
B\mathrm{Diff}(g(S^7\times S^7),D^{14})&\longrightarrow \Omega_0^\infty \mathrm{MTString}(14),\\
B\mathrm{Diff}(g(\mathbb{O}P^2\#\overline{\mathbb{O}P}^2),D^{16})&\longrightarrow \Omega_0^\infty \mathrm{MTString}(16).
\end{aligned}$$

A slightly different type of example is given by $B\mathrm{Diff}(\mathbb{C}P^3\#g(S^3\times S^3),U)$, where $U\subset\mathbb{C}P^3$ is a tubular neighbourhood of $\mathbb{C}P^1$. In this case the stable homology is that of $\Omega_0^\infty \mathrm{MTSpin}^c(6)$, where $B\mathrm{Spin}^c(6)$ is the homotopy fibre of the map

$$\beta w_2: B\mathrm{SO}(6)\longrightarrow K(\mathbb{Z},3).$$

An example where we need a more complicated stabilisation (not induced by connected sum) comes from $\mathbb{R}P^6$. The map $\mathbb{R}P^6\to B\mathrm{O}(6)$ lifts canonically to a 3-connected map $\mathbb{R}P^6\to B\mathrm{Pin}^-(6)$, where $\theta: B\mathrm{Pin}^-(6)\to B\mathrm{O}(6)$ is the homotopy fibre of

$$w_2+w_1^2: B\mathrm{O}(6)\longrightarrow K(\mathbb{Z}/2\mathbb{Z},2).$$

The standard self-indexing Morse function $f\colon \mathbb{R}P^6 \to [0,6]$ given by

$$f(x_0;\ldots;x_6) = \sum_{i=0}^{6} i x_i^2$$

has one critical point of each index, and we let $W = f^{-1}([0,2.5]) \cong \mathbb{R}P^2 \times D^4$. Cutting out a parallel copy of W gives a θ-bordism $\widetilde{K} \cong f^{-1}([2.5,3.5])$ from $\partial W = \mathbb{R}P^2 \times S^3$ to $-\partial W$ (i.e. $\mathbb{R}P^2 \times S^3$ equipped with the opposite θ-structure). Hence $K_0 = \widetilde{K} \circ (-\widetilde{K})$ is a cobordism from ∂W to itself, and we let K be the infinite iteration. In this situation we get a stable homology equivalence

$$B\mathrm{Diff}((\mathbb{R}P^2 \times D^4) \cup_\partial g K_0, \partial) \longrightarrow \Omega_0^\infty \mathrm{MTPin}^-(6).$$

Another interesting special case concerning the manifolds $W_g = \#^g(S^n \times S^n)$ is the following. Let (Y,y) be a pointed space, and consider the homotopy orbit space

$$\mathcal{S}_g^n(Y,y) = (E\mathrm{Diff}(W_g, D^{2n}) \times \mathrm{Map}((W_g, D^{2n}), (Y,y)))/\mathrm{Diff}(W_g, D^{2n}).$$

We can determine the stable homology of these spaces using a Pontryagin–Thom map

$$\coprod_{g \geqslant 0} \mathcal{S}_g^n(Y,y) \longrightarrow \Omega^\infty(Y\langle n-1\rangle_+ \wedge \mathrm{MT}\theta^n), \tag{1.4}$$

defined as in Remark 1.10. Any map $f\colon (S^n, D^n) \to (Y,y)$ may be composed with the projection to the first coordinate $S^n \times S^n \to S^n$ to give a map $(W_1, D^{2n} \cup D^{2n}) \to (Y,y)$, representing an element $[W_f] \in \mathcal{K}$, and we let $\mathcal{L} \subset \mathcal{K}$ be the submonoid generated by the $[W_f]$. It is easy to check that $\mathcal{L}$ satisfies the conditions of Theorem 1.13 for the tangential structure $Y\langle n-1\rangle \times B\mathrm{O}(2n)\langle n\rangle \to B\mathrm{O}(2n)$. This shows that (1.4) induces a map from the stabilised homology

$$\bigg(\bigoplus_{g \geqslant 0} H_*(\mathcal{S}_g^n(Y,y))\bigg)[\mathcal{L}^{-1}] \longrightarrow H_*(\Omega^\infty(Y\langle n-1\rangle_+ \wedge \mathrm{MT}\theta^n)) \tag{1.5}$$

which is an isomorphism, after restricting to appropriate path components. This result is a generalisation of the result of Cohen and Madsen [CM], who proved the special case where $2n=2$ and Y is simply connected. (The case $2n=2$ was generalised to non-simply connected Y in [GRW1].)

As a final application, in [GRW3] we deduce a generalisation of the detection result of Ebert [E1]. We will prove that for any abelian group k and any non-zero cohomology class $c \in H^*(\Omega_0^\infty \mathrm{MTSO}(2n); k)$, there exists a bundle $\pi\colon E \to B$ of smooth oriented manifolds, such that the characteristic class associated with c is non-vanishing in $H^*(B;k)$. (The case $k = \mathbb{Q}$ was proved by Ebert.)

1.6. Cobordism categories and outline of proof

Finally, let us say a few words about our method of proof, which follows the strategy in [GRW1] and [GMTW]. A central object is the *cobordism category* $\mathcal{C}_\theta(\mathbb{R}^N)$, whose morphisms are cobordisms $W\subset[0,t]\times\mathbb{R}^N$ of dimension $2n$ and whose objects are closed $(2n-1)$-dimensional manifolds $M\subset\mathbb{R}^N$, both equipped with θ-structures.

Remark 1.14. The applications described above use only the case where morphisms are even-dimensional. Many of our results about cobordism categories are valid for odd-dimensional cobordisms as well, but we do not know an interpretation in terms of stable homology in that case. In fact, Ebert [E2] has shown that there are non-trivial classes in $H^*(\Omega_0^\infty \operatorname{MTSO}(2n+1);\mathbb{Q})$ which are trivial when restricted to any $B\mathrm{Diff}^+(M,\partial M)$. Thus there can be no analogue of e.g. Theorem 1.8, expressing $H_*(\Omega_0^\infty \operatorname{MTSO}(2n+1))$ as a direct limit of $H_*(B\mathrm{Diff}(W\cup K|_{[0,i]},K|_i))$'s. It is an interesting question to find an odd-dimensional analogue of our results.

In the limit $N\to\infty$, the main result of [GMTW] gives a weak equivalence

$$\Omega B\mathcal{C}_\theta \simeq \Omega^\infty \operatorname{MT}\theta. \tag{1.6}$$

As in [GRW1], our strategy will be to find subcategories $\mathcal{C}\subset\mathcal{C}_\theta$, as small as possible, such that the inclusion induces a weak equivalence $\Omega B\mathcal{C}\to\Omega B\mathcal{C}_\theta$. The proof of Theorem 1.8 will consist of applying a version of the "group-completion" theorem to a very small subcategory of $\mathcal{C}_\theta$. Let us briefly describe the conditions we impose on objects and morphisms of this subcategory.

Let L be a $(2n-1)$-manifold with boundary which admits a handle structure with no handles of index n or larger, and let ℓ_L be a θ-structure whose underlying map $L\to B$ is $(n-1)$-connected. Then we pick a (collared) embedding $L\to(-\infty,0]\times\mathbb{R}^\infty$, and consider the subcategory $\mathcal{C}_{\theta,L}\subset\mathcal{C}_\theta$ where objects $M\subset\mathbb{R}\times\mathbb{R}^\infty$ satisfy $M\cap((-\infty,0]\times\mathbb{R}^\infty)=L$ and morphisms $W\subset[0,t]\times\mathbb{R}\times\mathbb{R}^\infty$ satisfy $W\cap([0,t]\times(-\infty,0]\times\mathbb{R}^\infty)=[0,t]\times L$. For both objects and morphisms, these identities are required to hold as θ-manifolds. In §2 we prove that the inclusion map induces a weak equivalence

$$B\mathcal{C}_{\theta,L}\longrightarrow B\mathcal{C}_\theta. \tag{1.7}$$

Secondly, we filter $\mathcal{C}_{\theta,L}$ by *connectivity of morphisms*: for $-1\leqslant\kappa\leqslant n-1$, the subcategory $\mathcal{C}_{\theta,L}^\kappa$ has the same objects as $\mathcal{C}_{\theta,L}$, but a morphism W from M_0 to M_1 is required to satisfy that the inclusion $M_1\to W$ is κ-connected. In §3 we prove that the inclusion map induces a weak equivalence

$$B\mathcal{C}_{\theta,L}^\kappa\longrightarrow B\mathcal{C}_{\theta,L}. \tag{1.8}$$

(In the case where $\kappa=0$, this is the "positive boundary subcategory", and this case was proved in [GMTW].)

Thirdly, we filter $\mathcal{C}_{\theta,L}^{\kappa}$ by *connectivity of objects*: for $-1\leqslant l\leqslant\min\{n-2,\kappa\}$, the subcategory $\mathcal{C}_{\theta,L}^{\kappa,l}\subset\mathcal{C}_{\theta,L}^{\kappa}$ is the full subcategory on those objects where the structure map $M\to B$ induces an injection $\pi_i(M)\to\pi_i(B)$ for all $i\leqslant l$ and all basepoints, or equivalently the inclusion $L\to M$ is l-connected. In §4 we prove that the inclusion map induces a weak equivalence

$$B\mathcal{C}_{\theta,L}^{\kappa,l}\longrightarrow B\mathcal{C}_{\theta,L}^{\kappa}. \tag{1.9}$$

(In the case where $l=0$ and B is connected, this is the full subcategory on objects which are path connected, and this case was proved in [GRW1].)

We have now reduced to $\mathcal{C}_{\theta,L}^{n-1,n-2}$, the full subcategory on those objects for which the inclusion $L\to M$ is $(n-2)$-connected. In the fourth and final step of the filtration we let $\mathcal{C}$ denote the full subcategory on those objects M which can be obtained from L by attaching handles of index at least n. (This is equivalent to the condition that $M\setminus\operatorname{int}(L)$ is diffeomorphic to a handlebody with handles of index at most $n-1$, which if $n>3$ is in turn equivalent to the inclusion $L\to M$ being $(n-1)$-connected.) In §5 we prove that the inclusion map induces a weak equivalence

$$\Omega B\mathcal{C}\longrightarrow\Omega B\mathcal{C}_{\theta,L}^{n-1,n-2} \tag{1.10}$$

provided that θ is spherical.

Now, given a closed $(2n-1)$-manifold P with θ-structure ℓ_P, we will show how to obtain a θ-manifold L as described above, such that the space $\mathcal{N}^{\theta}(P,\ell_P)$ occurs as a space of morphisms in the category $\mathcal{C}$. The weak equivalences (1.6)–(1.10) establish the homotopy equivalence

$$\Omega B\mathcal{C}\simeq\Omega^{\infty}\,\mathrm{MT}\theta,$$

and the proof of Theorem 1.8 in this case will be completed by applying a suitable version of the "group-completion" theorem to the map $\mathcal{N}^{\theta}(P,\ell_P)\to\Omega B\mathcal{C}$.

The weak equivalences (1.8)–(1.10) are established using a parameterised surgery procedure, and the proof depends on the contractibility of certain spaces of surgery data. Contractibility is proved in a similar way in all three cases, and we defer this to §6. Finally, in §7 we explain how to use a version of the group-completion theorem to prove Theorem 1.8 and tie things together.

§3–§6 contain the main technical steps, but on a first reading it is possible to skip to §7 after reading §2, to see the overall structure of the argument. The reader mainly interested in Theorems 1.1 and 1.2 can take

$$\theta=\theta^n\colon B\mathrm{O}(2n)\langle n\rangle\longrightarrow B\mathrm{O}(2n)$$

and $L \cong D^{2n-1}$ in the above outline and throughout the paper. Considering only this special case would not significantly simplify the main technical steps in §3–§6, but the group-completion arguments in §7 do simplify, and we incorporate a separate discussion of this case in §7.1.

2. Definitions and recollections

2.1. Tangential structures

Throughout this paper, an important role will be played by the notion of a tangential structure on manifolds. This will be important even for the proofs of theorems which do not explicitly mention tangential structures on manifolds. However, for the proofs of Theorems 1.1 and 1.2, the structure $\theta = \theta^n: B\mathrm{O}(2n)\langle n\rangle \to B\mathrm{O}(2n)$ suffices.

Definition 2.1. A *tangential structure* is a map $\theta: B \to B\mathrm{O}(d)$. A *$\theta$-structure* on a d-manifold W is a bundle map (i.e. fibrewise linear isomorphism) $\ell: TW \to \theta^*\gamma$. A *$\theta$-manifold* is a pair (W, ℓ). More generally, a θ-structure on a k-manifold M (with $k \leqslant d$) is a bundle map $\ell: \varepsilon^{d-k} \oplus TM \to \theta^*\gamma$.

Given vector bundles U and V of the same dimension, but not necessarily over the same space, we write $\mathrm{Bun}(U, V)$ for the subspace of $\mathrm{map}(U, V)$ (with the compact-open topology) consisting of the bundle maps. Thus, $\mathrm{Bun}(TW, \theta^*\gamma)$ is the space of θ-structures on W.

2.2. Spaces of manifolds

We recall the definition and main properties of spaces of submanifolds, from [GRW1]. Fix a tangential structure $\theta: B \to B\mathrm{O}(d)$.

Definition 2.2. For an open subset $U \subset \mathbb{R}^n$, we denote by $\Psi_\theta(U)$ the set of pairs (M^d, ℓ) where $M^d \subset U$ is a smooth d-dimensional submanifold that is closed as a topological subspace, and ℓ is a θ-structure on M.

We denote by $\Psi_{\theta_{d-m}}(U)$ the set of pairs (M, ℓ) where $M \subset U$ is a smooth $(d-m)$-dimensional submanifold that is closed as a topological subspace, and ℓ is a θ-structure on M, i.e. a bundle map $\varepsilon^m \oplus TM \to \theta^*\gamma$.

In [GRW1, §2] we have defined a topology on these sets so that $U \mapsto \Psi_{\theta_{d-m}}(U)$ defines a continuous sheaf of topological spaces on the site of open subsets of $\mathbb{R}^n$. We will not give full details of the topology again here, but remind the reader that the topology is "compact-open" in flavour: disregarding tangential structures, points nearby to M are those which near some large compact subset $K\subset U$ look like small normal deformations of M. In particular, a typical neighbourhood of the empty manifold $\varnothing\in\Psi_\theta(U)$ consists of all those manifolds in U disjoint from some compact K.

Definition 2.3. We define $\psi_\theta(n,k)\subset\Psi_\theta(\mathbb{R}^n)$ to be the subspace consisting of those θ-manifolds (M,ℓ) such that $M\subset\mathbb{R}^k\times(-1,1)^{n-k}$. We make the analogous definition of $\psi_{\theta_{d-m}}(n,k)$.

2.3. Semi-simplicial spaces and non-unital categories

Let Δ denote the category of finite non-empty totally ordered sets and monotone maps, the simplicial indexing category. Let $\Delta_{\mathrm{inj}}\subset\Delta$ denote the subcategory with the same objects but only injective monotone maps as morphisms. For a category $\mathcal{C}$, a *simplicial object in* $\mathcal{C}$ is a contravariant functor $X:\Delta\to\mathcal{C}$, and a *semi-simplicial object in* $\mathcal{C}$ is a contravariant functor $X:\Delta_{\mathrm{inj}}\to\mathcal{C}$. A map of (semi-)simplicial objects is a natural transformation of functors.

We call a semi-simplicial object in the category of topological spaces a *semi-simplicial space*. More concretely, it consists of a space $X_n=X(0<1<...<n)$ for each $n\geqslant 0$, and face maps $d_i: X_n\to X_{n-1}$ defined for $i=0,...,n$ satisfying the simplicial identities $d_id_j=d_{j-1}d_i$ for $i<j$. We often denote a semi-simplicial space by $X_\bullet$, where we treat $\bullet$ as a place-holder for the simplicial degree.

The *geometric realisation* of a semi-simplicial space $X_\bullet$ is defined to be

$$|X_\bullet|=\coprod_{n\geqslant 0}(X_n\times\Delta^n)/\sim,$$

where Δ^n denotes the standard topological n-simplex and the equivalence relation is generated by $(d_i(x),y)\sim(x,d^i(y))$, where $d^i:\Delta^n\to\Delta^{n+1}$ is the inclusion of the ith face. This space is given the quotient topology. We shall need to make reference to specific points in geometric realisations: a point $y\in|X_\bullet|$ is uniquely written as $y=(x,t)$ with $x\in X_p$ and $t\in\operatorname{int}(\Delta^p)$.

The *k-skeleton* of $|X_\bullet|$ is

$$|X_\bullet|^{(k)}=\coprod_{n=0}^{k}(X_n\times\Delta^n)/\sim$$

with the quotient topology, and one easily checks that $|X_\bullet|=\bigcup_{k\geqslant 0}|X_\bullet|^{(k)}$ with the direct limit topology. A useful consequence of this is the following: a map from a compact space to $|X_\bullet|$ factors through a finite skeleton (since we do not have degeneracies, it suffices to prove this for the terminal semi-simplicial space, where it is clear). We recall the following result (cf. e.g. [S2, Proposition A.1 (ii)]).

LEMMA 2.4. *If $X_\bullet \to Y_\bullet$ is a map of semi-simplicial spaces such that each $X_n \to Y_n$ is a weak homotopy equivalence, then $|X_\bullet| \to |Y_\bullet|$ is too.*

Remark 2.5. The term *semi-simplicial object* we have defined above is not quite standard (though is gaining popularity) and deserves some justification. Our justification is that it agrees with Eilenberg and Zilber's original usage of "semi-simplicial complex" [EZ]. Another is that the alternative used in the literature is Δ-*space*, but as Δ is the indexing category for full simplicial objects this seems counterintuitive.

A *non-unital topological category* $\mathcal{C}$ consists of a pair of spaces $(\mathcal{O}, \mathcal{M})$ of objects and morphisms, equipped with source and target maps $s, t: \mathcal{M} \to \mathcal{O}$. We let $\mathcal{M} \times_{t\mathcal{O}s} \mathcal{M}$ denote the fibre product made with the maps t and s, and require in addition a composition map $\mu: \mathcal{M} \times_{t\mathcal{O}s} \mathcal{M} \to \mathcal{M}$ which satisfies the evident associativity requirement.

A non-unital topological category $\mathcal{C}$ has a semi-simplicial nerve, generalising the simplicial nerve of a topological category [S1]. Define $N_\bullet \mathcal{C}$ by $N_0 \mathcal{C} = \mathcal{O}$ and

$$N_k \mathcal{C} = \mathcal{M} \times_{t\mathcal{O}s} \mathcal{M} \times_{t\mathcal{O}s} \ldots \times_{t\mathcal{O}s} \mathcal{M}, \quad k > 0,$$

being the space of k-tuples of composable morphisms, and let the face maps be given by composing and forgetting morphisms, as in the simplicial nerve of a topological category. We define the *classifying space* of a non-unital topological category by

$$B\mathcal{C} = |N_\bullet \mathcal{C}|.$$

2.4. Definition of the cobordism categories

For convenience in the rest of the paper, we introduce the following notation. All of our constructions will take place inside $\mathbb{R} \times \mathbb{R}^N$, and we write $x_1: \mathbb{R} \times \mathbb{R}^N \to \mathbb{R}$ for the projection to the first coordinate. Given a manifold $W \subset \mathbb{R} \times \mathbb{R}^N$ and a set $A \subset \mathbb{R}$, we write

$$W|_A = W \cap x_1^{-1}(A),$$

if it is a manifold, and we also write $\ell|_A$ for the restriction of a θ-structure ℓ on W to this manifold.

Our definition of the cobordism category of θ-manifolds is similar to that of [GRW1] (the only difference is that here will we only define a non-unital category); it follows that of [GMTW] in spirit, but is different in some technical points. (These slight variations all have equivalent classifying spaces.) We use the spaces of manifolds of §2.2 in order to describe the point-set topology of these categories.

Definition 2.6. For each $\varepsilon>0$ we let the non-unital topological category $\mathcal{C}_\theta(\mathbb{R}^N)_\varepsilon$ have space of objects $\psi_{\theta_{d-1}}(N,0)$. The space of morphisms from (M_0,ℓ_0) to (M_1,ℓ_1) is the subspace of those $(t,(W,\ell))\in\mathbb{R}\times\psi_\theta(N+1,1)$ such that $t>0$ and

$$W|_{(-\infty,\varepsilon)}=(\mathbb{R}\times M_0)|_{(-\infty,\varepsilon)}\in\Psi_\theta((-\infty,\varepsilon)\times\mathbb{R}^N)$$

and

$$W|_{(t-\varepsilon,\infty)}=(\mathbb{R}\times M_1)|_{(t-\varepsilon,\infty)}\in\Psi_\theta((t-\varepsilon,\infty)\times\mathbb{R}^N).$$

Here $\mathbb{R}\times M_i$ denotes the θ-manifold with underlying manifold $\mathbb{R}\times M_i\subset\mathbb{R}\times\mathbb{R}^N$ and θ-structure

$$T(\mathbb{R}\times M_i)\longrightarrow\varepsilon^1\oplus TM_i\xrightarrow{\ell_i}\theta^*\gamma.$$

Composition in this category is defined by

$$(t',W')\circ(t,W)=(t+t',W|_{(-\infty,t]}\cup(W'+te_1)|_{[t,\infty)}),$$

where $W'+te_1$ denotes the manifold W' translated by t in the first coordinate. We topologise the total space of morphisms as a subspace of $(0,\infty)\times\psi_\theta(N+1,1)$.

If $\varepsilon<\varepsilon'$ there is an inclusion $\mathcal{C}_\theta(\mathbb{R}^N)_{\varepsilon'}\subset\mathcal{C}_\theta(\mathbb{R}^N)_\varepsilon$, and we define $\mathcal{C}_\theta(\mathbb{R}^N)$ to be the colimit over all $\varepsilon>0$.

Note that a morphism $(t,(W,\ell))$ in this category is uniquely determined by the restriction $(t,(W|_{[0,t]},\ell|_{[0,t]}))$. We often think of morphisms in this category as being given by such restricted manifolds, but the topology on the space of morphisms is best described as we did above.

As explained in the introduction, we will also require a version of this category where the objects and morphisms contain a fixed codimension-zero submanifold. In order to define this, we let

$$L\subset\left(-\tfrac{1}{2},0\right]\times(-1,1)^{N-1}$$

be a compact $(d-1)$-manifold which near $\{0\}\times\mathbb{R}^{N-1}$ agrees with $(-1,0]\times\partial L$. Furthermore, we let $\ell|_L:\varepsilon^1\oplus TL\to\theta^*\gamma$ be a θ-structure on L. Near ∂L we require that the structure is a product (i.e. that translation in the collar direction preserves the structure). Such an $\ell|_L$ makes $\mathbb{R}\times L$ into a θ-manifold with boundary, and we make the following definition.

Definition 2.7. The topological subcategory $\mathcal{C}_{\theta,L}(\mathbb{R}^N)\subset\mathcal{C}_\theta(\mathbb{R}^N)$ has space of objects those (M,ℓ) such that

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

as θ-manifolds. It has space of morphisms from (M_0, ℓ_0) to (M_1, ℓ_1) given by those $(t, (W, \ell))$ such that

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

as θ-manifolds.

Remark 2.8. The category $\mathcal{C}_{\theta,L}(\mathbb{R}^N)$ does not really depend on L, but only on ∂L. It is sometimes convenient to think of the interior of L as being cut out, so that objects in the category are manifolds with boundary ∂L and morphisms are cobordisms between manifolds with boundary which are trivial along the boundary.

If we take $L = D^{d-1}$, then the category $\mathcal{C}_{\theta,L}(\mathbb{R}^N)$ is equivalent to the category of "manifolds with basepoint" defined in [GRW1, Definition 4.2]. That case is sufficient for the proofs of Theorems 1.1 and 1.2.

The subject of our main technical theorem, from which we shall show how to obtain results on diffeomorphism groups in §7, is certain subcategories of $\mathcal{C}_{\theta,L}(\mathbb{R}^N)$ where we require the morphisms to have certain connectivities relative to their outgoing boundaries, and objects to be those (M, ℓ_M) whose Gauss map $M \to B$ (i.e. the map underlying $\ell_M : \varepsilon^1 \oplus TM \to \theta^* \gamma$) has a certain injectivity range on homotopy groups.

Definition 2.9. For an integer $\kappa \geqslant -1$, the topological subcategory

$$\mathcal{C}^{\kappa}_{\theta,L}(\mathbb{R}^N) \subset \mathcal{C}_{\theta,L}(\mathbb{R}^N)$$

has the same space of objects. It has space of morphisms from (M_0, ℓ_0) to (M_1, ℓ_1) given by those $(t, (W, \ell))$ such that the pair $(W|_{[0,t]}, W|_t)$ is κ-connected. Thus this is the subcategory on those morphisms which are κ-connected relative to their outgoing boundary.

The category $\mathcal{C}^0_\theta$ is the "positive boundary category" as in [GMTW], where each path component of a cobordism is required to have non-empty outgoing boundary.

Definition 2.10. For an integer $l \geqslant -1$, the topological subcategory

$$\mathcal{C}^{\kappa,l}_{\theta,L}(\mathbb{R}^N) \subset \mathcal{C}^{\kappa}_{\theta,L}(\mathbb{R}^N)$$

is the full subcategory on those objects (M, ℓ) such that the map

$$\ell_* : \pi_i(M) \longrightarrow \pi_i(B)$$

is injective for all $i \leqslant l$ and all basepoints. (In our main application in §7, the map $L \to B$ will be $(l+1)$-connected. In that case the requirement is equivalent to (M, L) being l-connected.)

For our final definition we specialise to even dimensions.

Definition 2.11. Let $d=2n$ and $\mathcal{A}\subset\pi_0(\mathrm{Ob}(\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)))$ be a collection of path components of the space of objects. The topological subcategory

$$\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)\subset\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$$

is the full subcategory on the subspace of those objects in $\mathcal{A}$.

For $N=\infty$, we shall often denote $\mathcal{C}_\theta(\mathbb{R}^\infty)=\mathrm{colim}_N\,\mathcal{C}_\theta(\mathbb{R}^N)$ by $\mathcal{C}_\theta$, and similarly with any decorations.

2.5. The homotopy type of the cobordism category

The main theorem of [GMTW] identifies the homotopy type $\Omega B\mathcal{C}_\theta$ in terms of the infinite loop space of a certain Thom spectrum $\mathrm{MT}\theta$.

Recall from the introduction that given a map $\theta\colon B\to B\mathrm{O}(d)=\mathrm{Gr}_d(\mathbb{R}^\infty)$ we let

$$B(\mathbb{R}^n)=\theta^{-1}(\mathrm{Gr}_d(\mathbb{R}^n))$$

and define $\gamma_n^\perp\to\mathrm{Gr}_d(\mathbb{R}^n)$ to be the orthogonal complement of the tautological bundle. The canonical map $B(\mathbb{R}^n)\to B(\mathbb{R}^{n+1})$ pulls back $\theta^*\gamma_{n+1}^\perp$ to $\theta^*\gamma_n^\perp\oplus\varepsilon^1$ and hence we obtain pointed maps

$$B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp}\wedge S^1\longrightarrow B(\mathbb{R}^{n+1})^{\theta^*\gamma_{n+1}^\perp}$$

of Thom spaces, which form a spectrum $\mathrm{MT}\theta$. Its associated infinite loop space is

$$\Omega^\infty\,\mathrm{MT}\theta=\operatorname*{colim}_{n\to\infty}\Omega^n(B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp}).$$

THEOREM 2.12. (Galatius–Madsen–Tillmann–Weiss [GMTW]) *There is a canonical map*

$$\Omega B\mathcal{C}_\theta\longrightarrow\Omega^\infty\,\mathrm{MT}\theta$$

which is a weak homotopy equivalence.

We write $\Omega_0^\infty\,\mathrm{MT}\theta$ for the basepoint component of $\Omega^\infty\,\mathrm{MT}\theta$, and now describe the rational cohomology of this space. The map $B\xrightarrow{\theta}B\mathrm{O}(d)\xrightarrow{\det}B\mathrm{O}(1)$ on fundamental groups defines a character $w_1\colon\pi_1(B)\to\mathbb{Z}^\times$, and we write $H^*(B;\mathbb{Q}^{w_1})$ for the rational cohomology of B with local coefficients given by this character. For each n there are evaluation maps

$$\mathrm{ev}\colon\Sigma^n\Omega^n(B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp})\longrightarrow B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp},$$

and so we can define the dashed map in the diagram

$$\begin{tikzcd} H^{*+d}(B(\mathbb{R}^n);\mathbb{Q}^{w_1}) \arrow[r, dashed] \arrow[d, equal, "\text{Thom iso.}"'] & H^*(\Omega^n(B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp});\mathbb{Q}) \arrow[d, equal, "\text{suspension iso.}"] \\ \widetilde{H}^{*+n}(B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp};\mathbb{Q}) \arrow[r, "\mathrm{ev}^*"] & \widetilde{H}^{*+n}(\Sigma^n\Omega^n(B(\mathbb{R}^n)^{\theta^*\gamma_n^\perp});\mathbb{Q}) \end{tikzcd}$$

by commutativity. Taking limits and restricting to the basepoint component, we obtain a map

$$\sigma\colon H^{*+d}(B;\mathbb{Q}^{w_1}) \longrightarrow H^*(\Omega_0^\infty \mathrm{MT}\theta;\mathbb{Q})$$

(there is no $\lim^1$ contribution as we are working over a field). The right-hand side is a graded-commutative algebra, so σ extends to the free graded-commutative algebra on the part of $H^{*+d}(B;\mathbb{Q}^{w_1})$ of degree >0,

$$\Lambda(H^{*+d>0}(B;\mathbb{Q}^{w_1})) \longrightarrow H^*(\Omega_0^\infty \mathrm{MT}\theta;\mathbb{Q}).$$

This is an isomorphism of graded-commutative algebras.

2.6. Poset models

A key step in the proofs of [GMTW] and [GRW1] identifying the infinite loop space $B\mathcal{C}_\theta$ is to first identify this classifying space with the classifying space of a certain topological poset. The result holds for all variations of the cobordism category mentioned above; we prove the general result in Proposition 2.14 below.

Definition 2.13. Let

$$D_\theta \subset \mathbb{R}\times\mathbb{R}_{>0}\times\psi_\theta(N+1,1)$$

denote the subspace of tuples $(t,\varepsilon,(W,\ell))$ such that $[t-\varepsilon,t+\varepsilon]$ consists of regular values for $x_1\colon W\to\mathbb{R}$. Define a partial order on D_θ by

$$(t,\varepsilon,(W,\ell))<(t',\varepsilon',(W',\ell'))$$

if and only if $(W,\ell)=(W',\ell')$ and $t+\varepsilon<t'-\varepsilon$.

Define the full subposet $D_{\theta,L}\subset D_\theta$ to consist of those tuples $(t,\varepsilon,(W,\ell))$ such that $W\cap(\mathbb{R}\times(-\infty,0]\times\mathbb{R}^{N-1})=\mathbb{R}\times L$ as θ-manifolds.

If $\mathcal{C}\subset\mathcal{C}_{\theta,L}(\mathbb{R}^N)$ is a subcategory which consists of entire path components of the object and morphism spaces of $\mathcal{C}_{\theta,L}(\mathbb{R}^N)$, let $D^{\mathcal{C}}_{\theta,L}\subset D_{\theta,L}$ be the smallest subposet consisting of entire path components of the object and morphism spaces of $D_{\theta,L}$ which contains those tuples $(t,\varepsilon,(W,\ell))$ such that $W|_t\in\mathrm{Ob}(\mathcal{C})$, and those morphisms $(t,\varepsilon,(W,\ell))<(t',\varepsilon',(W',\ell'))$ with $W|_{[t,t']}\in\mathrm{Mor}(\mathcal{C})$.

Proposition 2.14. *Let $\mathcal{C}\subset\mathcal{C}_{\theta,L}(\mathbb{R}^N)$ be a subcategory which consists of entire path components of the object and morphism spaces of $\mathcal{C}_{\theta,L}(\mathbb{R}^N)$. Then there is a weak homotopy equivalence*

$$B\mathcal{C}\simeq BD^{\mathcal{C}}_{\theta,L}.$$

Proof. We introduce an auxiliary topological poset $D^{\mathcal{C},\perp}_{\theta,L}$ which maps to both $D^{\mathcal{C}}_{\theta,L}$ and $\mathcal{C}$. It is the subposet of $D^{\mathcal{C}}_{\theta,L}$ consisting of $(t,\varepsilon,(W,\ell))$ such that (W,ℓ) is a product over $(t-\varepsilon,t+\varepsilon)$. This condition means that if we write $W|_t=\{t\}\times M$ and give M the inherited θ-structure, then

$$W|_{(t-\varepsilon,t+\varepsilon)}=(t-\varepsilon,t+\varepsilon)\times M$$

as θ-manifolds. Then there is a zig-zag of functors

$$D^{\mathcal{C}}_{\theta,L}\longleftarrow D^{\mathcal{C},\perp}_{\theta,L}\longrightarrow\mathcal{C},$$

where the first arrow is the inclusion of the subposet and the second is the functor that sends a morphism $(a<b,W,\ell)$ to the manifold $(W|_{[a,b]}-ae_1)$ extended cylindrically in $(-\infty,0]\times\mathbb{R}^N$ and $[b-a,\infty)\times\mathbb{R}^N$. This induces a zig-zag diagram

$$N_kD^{\mathcal{C}}_{\theta,L}\longleftarrow N_kD^{\mathcal{C},\perp}_{\theta,L}\longrightarrow N_k\mathcal{C},$$

and we prove that both maps are weak equivalences for all k in the same way as in [GRW1, Theorem 3.9]. □

Applying the above construction to the categories $\mathcal{C}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)$ we obtain topological posets $D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)$ and weak homotopy equivalences

$$B\mathcal{C}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)\simeq BD^{\kappa,l}_{\theta,L}(\mathbb{R}^N). \tag{2.1}$$

Similarly, when we specialise to the case $d=2n$ and let $\mathcal{A}\subset\pi_0(\mathrm{Ob}(\mathcal{C}^{n-1,n-2}_{\theta,L}(\mathbb{R}^N)))$ be a collection of path components of objects, we obtain weak homotopy equivalences

$$B\mathcal{C}^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)\simeq BD^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N). \tag{2.2}$$

2.7. The homotopy type of $\mathcal{C}_{\theta,L}(\mathbb{R}^N)$

In [GRW1, Theorems 3.9 and 3.10] we proved that there is a weak homotopy equivalence $BD_\theta(\mathbb{R}^N)\simeq\psi_\theta(N+1,1)$, which combined with Proposition 2.14 gives

$$B\mathcal{C}_\theta(\mathbb{R}^N)\simeq BD_\theta(\mathbb{R}^N)\simeq\psi_\theta(N+1,1). \tag{2.3}$$

(Strictly speaking, in that paper we worked with a version of $D_\theta(\mathbb{R}^N)$ where $\varepsilon=0$, but the obvious map induces a levelwise weak equivalence of nerves.) For the purposes of this paper we require a slightly stronger version of this result, taking into account the submanifold L.

Proposition 2.15. *There are weak homotopy equivalences*

$$B\mathcal{C}_{\theta,L}(\mathbb{R}^N)\simeq BD_{\theta,L}(\mathbb{R}^N)\simeq\psi_{\theta,L}(N+1,1),$$

where $\psi_{\theta,L}(N+1,1)\subset\psi_\theta(N+1,1)$ *is the subspace consisting of those* (W,ℓ) *such that* $W\cap(\mathbb{R}\times(-\infty,0]\times\mathbb{R}^{N-1})=\mathbb{R}\times L$ *as* θ*-manifolds.*

Proof. The proof of [GRW1, Theorem 3.10] applies verbatim. □

Proposition 2.16. *The inclusion*

$$i\colon\psi_{\theta,L}(N+1,1)\longrightarrow\psi_\theta(N+1,1)$$

is a weak homotopy equivalence.

Proof. This is similar to [GRW1, Lemma 4.6], which is essentially the case $L=D^{d-1}$. It requires careful analysis of θ-structures, so let us, for this proof only, denote the θ_{d-1}-structure on L by $\ell_L\colon\varepsilon^1\oplus TL\to\theta^*\gamma$. We first want to construct the *double* $D(L)$ of L as a θ_{d-1}-manifold, and a canonical θ-null-bordism of it. Recall that L is a submanifold of $\left(-\frac{1}{2},0\right]\times(-1,1)^{N-1}$ which we identify with

$$\{0\}\times\left(-\tfrac{1}{2},0\right]\times(-1,1)^{N-1}\subset(-1,0]\times\left(-\tfrac{1}{2},0\right]\times(-1,1)^{N-1}.$$

Let $V\subset(-1,0]\times\left(-\frac{1}{2},\frac{1}{2}\right)\times(-1,1)^{N-1}$ denote the subset swept out by rotating L around $(0,0)$ in the half-plane $(-1,0]\times(-1,1)$. Since L was collared, this subset is a d-dimensional submanifold with boundary, and L lies in its boundary. We define $D(L)=\partial V$, and $\bar{L}=D(L)\setminus\operatorname{int}(L)$. The inclusion $L\hookrightarrow V$ is a homotopy equivalence, so there is a unique extension up to homotopy

$$\begin{tikzcd}
\varepsilon^1\oplus TL \arrow[r, "\ell_L"] \arrow[d] & \theta^*\gamma \\
TV, \arrow[ur, dashed] &
\end{tikzcd}$$

where the vertical map sends ε^1 to the outwards pointing vector. This restricts to a θ-structure on $D(L)$, and hence on $\bar{L}$, and V gives a θ-cobordism $V\colon\varnothing\rightsquigarrow D(L)$.

Similarly, we can rotate L in the half-plane $[0,1)\times(-1,1)$ around the point $\left(0,-\frac{1}{2}\right)$ to obtain a submanifold of $[0,1]\times[-1,0]\times(-1,1)^{N-1}$, extending to a θ-cobordism $U\subset[0,1]\times[-1,0]\times(-1,1)^{N-1}$, ending at $\{1\}\times[-1,0]\times\partial L$ and starting at $\{0\}\times(L\cup(\bar{L}-e_1))$, where $\bar{L}-e_1\subset\left[-1,\frac{1}{2}\right)\times(-1,1)^{N-1}$ denotes the parallel translate of $\bar{L}$.

The θ-manifolds U and V give us the tools we need. $D(L)$ is a submanifold of $\left(-\frac{1}{2},\frac{1}{2}\right)\times(-1,1)^{N-1}$, so we have a θ-manifold $\mathbb{R}\times D(L)\subset\mathbb{R}\times\left(-\frac{1}{2},\frac{1}{2}\right)\times(-1,1)^{N-1}$. We define a map

$$r\colon\psi_\theta(N+1,1)\longrightarrow\psi_{\theta,L}(N+1,1),$$

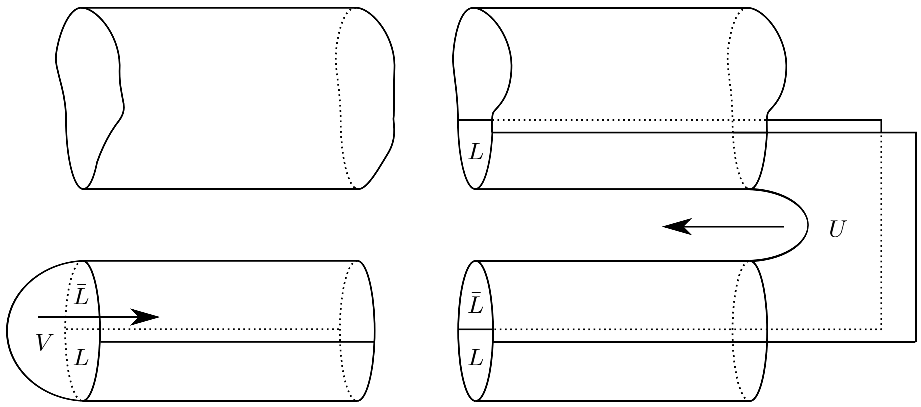

Figure 1. Adding and removing L.

which given $(W, \ell) \subset \mathbb{R} \times (-1,1) \times (-1,1)^{N-1}$ applies the unique increasing affine diffeomorphism $(-1,1) \cong \left(\frac{1}{2}, 1\right)$ to its second coordinate, and then takes the (disjoint) union with $\mathbb{R} \times D(L)$.

The composition $i \circ r$ is homotopic to the identity as the θ-null-bordism V of $D(L)$ may be used to push the cylinder $\mathbb{R} \times D(L)$ off to the right. A similar argument, pushing U to the left, proves that the composition $r \circ i$ is homotopic to the identity. Figure 1 shows how. $\square$

Combining this proposition with Proposition 2.15 and the homotopy equivalence (2.3) gives the following corollary.

Corollary 2.17. *For any pair (L, ℓ_L) as in Definition 2.7, the inclusion*

$$B\mathcal{C}_{\theta,L}(\mathbb{R}^N) \longrightarrow B\mathcal{C}_\theta(\mathbb{R}^N)$$

is a weak homotopy equivalence.

2.8. A more flexible model

From the poset models of §2.6 we construct the semi-simplicial spaces

$$D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_\bullet = N_\bullet D^{\kappa,l}_{\theta,L}(\mathbb{R}^N).$$

The remarks of §2.6 and Proposition 2.15 show that the geometric realisations of these semi-simplicial spaces are models for the classifying spaces of the categories $\mathcal{C}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)$ in which we are interested. The benefit of working with these semi-simplicial spaces instead of the cobordism categories is that we can often make constructions which are *not* functorial, yet give well-defined maps between geometric realisations of the semi-simplicial spaces involved.

To make this technique easier to apply, we will define an auxiliary semi-simplicial space $X_\bullet^{\kappa,l}$. (The space also depends on θ, N and L, but we suppress these from the notation.) We will prove that its geometric realisation is weakly equivalent to $B\mathcal{C}_{\theta,L}^{\kappa,l}(\mathbb{R}^N)$, but it will be easier to construct a simplicial map into $X_\bullet^{\kappa,l}$ than into $N_\bullet\mathcal{C}_{\theta,L}^{\kappa,l}(\mathbb{R}^N)$ or $D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet$. We will also write $X_\bullet^{\kappa}$ for $X_\bullet^{\kappa,-1}$.

Definition 2.18. Let $\theta\colon B\to B\mathrm{O}(d)$, N and L be as before, and let $\kappa, l\geqslant -1$ be integers. Define $X_\bullet^{\kappa,l}$ to be the semi-simplicial space with p-simplices consisting of certain tuples $(a,\varepsilon,(W,\ell))$ such that $a=(a_0,...,a_p)\in\mathbb{R}^{p+1}$, $\varepsilon=(\varepsilon_0,...,\varepsilon_p)\in(\mathbb{R}_{>0})^{p+1}$, and $(W,\ell)\in\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N)$, satisfying

(i) $W\subset(a_0-\varepsilon_0,a_p+\varepsilon_p)\times(-1,1)^N$;

(ii) W and $(a_0-\varepsilon_0,a_p+\varepsilon_p)\times L$ agree as θ-manifolds on the subspace $x_2^{-1}(-\infty,0]$;

(iii) $a_{i-1}+\varepsilon_{i-1}<a_i-\varepsilon_i$ for all $i=1,...,p$;

(iv) for each pair of regular values $t_0<t_1\in\bigcup_{i=1}^p(a_i-\varepsilon_i,a_i+\varepsilon_i)$, the cobordism $W|_{[t_0,t_1]}$ is κ-connected relative to its outgoing boundary;

(v) for each regular value $t\in(a_i-\varepsilon_i,a_i+\varepsilon_i)$, the map

$$\pi_j(W|_t)\longrightarrow\pi_j(B),$$

induced by $\ell|_t$, is injective for all basepoints and all $j\leqslant l$.

We topologise this set as a subspace of $\mathbb{R}^{p+1}\times(\mathbb{R}_{>0})^{p+1}\times\Psi_\theta((-1,1)\times\mathbb{R}^N)$, where we use the standard affine diffeomorphism $(-1,1)\cong(a_0-\varepsilon_0,a_p+\varepsilon_p)$ to identify the sets $\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N)$ and $\Psi_\theta((-1,1)\times\mathbb{R}^N)$. The jth face map is given by forgetting a_j and ε_j, and if $j=0$, composing with the restriction map

$$\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N)\longrightarrow\Psi_\theta((a_1-\varepsilon_1,a_p+\varepsilon_p)\times\mathbb{R}^N),$$

and similarly if $j=p$.

There are semi-simplicial maps $D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet\to X_\bullet^{\kappa,l}$, which on p-simplices are given by sending $(a,\varepsilon,(W,\ell))$ with $(W,\ell)\in\Psi_\theta(\mathbb{R}\times\mathbb{R}^N)$ to the same thing restricted down to $\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N)$.

The semi-simplicial space $X_\bullet^{\kappa,l}$ is easier to map into (by a semi-simplicial map) than $D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet$ for two reasons. Firstly, we do not require that the intervals $(a_i-\varepsilon_i,a_i+\varepsilon_i)$ consist entirely of regular values: instead we allow critical values, and conditions (iv) and (v) ensure that the critical values do not affect the essential properties of the space. Secondly, we discard those parts of the manifold outside of $(a_0-\varepsilon_0,a_p+\varepsilon_p)$, and so do not need to worry about controlling parts of the manifold outside of the region.

Definition 2.19. In the case $d=2n$, with $\mathcal{A}\subset\pi_0(\mathrm{Ob}(\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)))$ being a collection of path components of objects, we make the entirely analogous definition of $X_\bullet^{n-1,\mathcal{A}}$. Precisely, in Definition 2.18 we replace condition (v) by

(v′) for each regular value $t\in(a_i-\varepsilon_i,a_i+\varepsilon_i)$, the θ_{d-1}-manifold $(W|_t,\ell|_t)$ lies in $\mathcal{A}$.

The following is our main result concerning these models, and together with (2.1) and (2.2) provides weak homotopy equivalences $B\mathcal{C}_{\theta,L}^{\kappa,l}(\mathbb{R}^N)\simeq|X_\bullet^{\kappa,l}|$ and, in the case $d=2n$, $B\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)\simeq|X_\bullet^{n-1,\mathcal{A}}|$.

PROPOSITION 2.20. *Let κ and l satisfy the inequalities in Definition 2.18. The semi-simplicial map $D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet\to X_\bullet^{\kappa,l}$, and when $d=2n$ also $D_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)_\bullet\to X_\bullet^{n-1,\mathcal{A}}$, induce weak homotopy equivalences after geometric realisation.*

Proof. For the proof we introduce an auxiliary semi-simplicial space $\overline{X}_\bullet^{\kappa,l}$. Its p-simplices are those tuples

$$(a,\varepsilon,(W,\ell))\in\mathbb{R}^{p+1}\times(\mathbb{R}_{>0})^{p+1}\times\psi_\theta(N+1,1)$$

satisfying the conditions of Definition 2.18, except that the interval $(a_0-\varepsilon_0,a_p+\varepsilon_p)$ is replaced by $\mathbb{R}$ in (i) and (ii). We can regard $D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet$ as a subspace of $\overline{X}_\bullet^{\kappa,l}$ and write r for the inclusion, and we have a factorisation

$$D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet\xrightarrow{r}\overline{X}_\bullet^{\kappa,l}\longrightarrow X_\bullet^{\kappa,l}.$$

The map $\overline{X}_\bullet^{\kappa,l}\to X_\bullet^{\kappa,l}$ is a weak homotopy equivalence in each simplicial degree, by methods similar to [GRW1, Theorem 3.9]. Briefly, in simplicial degree p choose—continuously in the data $(a_0,a_p,\varepsilon_0,\varepsilon_p)$—diffeomorphisms $(a_0-\varepsilon_0,a_p+\varepsilon_p)\cong\mathbb{R}$ which are the identity on $[a_0,a_p]$. Using this family of diffeomorphisms to stretch gives a map $X_p^{\kappa,l}\to\overline{X}_p^{\kappa,l}$, which is homotopy inverse to the restriction map $\overline{X}_p^{\kappa,l}\to X_p^{\kappa,l}$.

To show that the first map r induces a weak homotopy equivalence on geometric realisation, we use a technique which we shall use many times in this paper. That is, we consider a commutative diagram

$$\begin{tikzcd}
\partial D^n \arrow[r, "\hat{f}"] \arrow[d] & {|D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet|} \arrow[d, hook, "|r|"] \\
D^n \arrow[r, "f"] \arrow[ur, dashed, "F"] & {|\overline{X}_\bullet^{\kappa,l}|}
\end{tikzcd}$$

and show that the pair of maps $(f,\hat{f})$ may be changed by a homotopy of such maps to a new pair which admits a dashed diagonal map making both triangles commute. This shows that $|r|$ is a weak homotopy equivalence, as claimed. Let us take the opportunity to point out that $|r|$ is a continuous injection because it is the geometric realisation of a levelwise continuous injection, but it does not follow that $|r|$ is a homeomorphism onto its image, even though r has this property levelwise.

For each $x \in D^n$ the point $f(x)$ is a tuple $(t, a, \varepsilon, (W(x), \ell))$, where $t \in \operatorname{int}(\Delta^p)$ and $(a, \varepsilon, (W(x), \ell)) \in \overline{X}_p^{\kappa,l}$, and we may choose a pair (a^x, ε^x) such that

$$[a^x - \varepsilon^x, a^x + \varepsilon^x] \subset \bigcup_{i=0}^{p} (a_i - \varepsilon_i, a_i + \varepsilon_i) \setminus \{a_i\}$$

and that $[a^x - \varepsilon^x, a^x + \varepsilon^x]$ consists of regular values of $x_1: W(x) \to \mathbb{R}$. By properness of $x_1: W(x) \to \mathbb{R}$, there is a neighbourhood $U_x \ni x$ for which $[a^x - \varepsilon^x, a^x + \varepsilon^x]$ still consists of regular values. The U_x's cover D^n and we let $\{U_j\}_{j \in J}$ be a finite subcover. We may suppose that $a^j \neq a^k$, as otherwise we may change the cover by letting $U'_j = U_j \cup U_k$ with $(a^j)' = a^j = a^k$ and $(\varepsilon^j)' = \min\{\varepsilon^j, \varepsilon^k\}$. Once the a^j are distinct, we may shrink the ε^j so that the intervals $[a^j + \varepsilon^j, a^j - \varepsilon^j]$ are pairwise disjoint, and so that no a_i lies in such an interval.

As the intervals $[a^j + \varepsilon^j, a^j - \varepsilon^j]$ are chosen to consist of regular values, the data $\{(U_j, a^j, \varepsilon^j)\}_{j \in J}$, together with a choice of partition of unity subordinate to the cover by the U_j's, determine a map $F': D^n \to |D_{\theta,L}^{-1,-1}(\mathbb{R}^N)_\bullet|$ with the same underlying family of θ-manifolds. As $[a^j - \varepsilon^j, a^j + \varepsilon^j] \subset \bigcup_{i=0}^{p} (a_i - \varepsilon_i, a_i + \varepsilon_i)$, this new family satisfies conditions (iv) and (v) of Definition 2.18 (as the old family did) so F' actually lifts further to a map $F: D^n \to |D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet|$. There is a homotopy H from $|r| \circ F$ to f as follows: on underlying θ-manifolds it is constant, but on the interval data we first use the straight-line homotopy from the data $\{(a^j, \varepsilon^j)\}_{j \in J}$ to the data $\{(a_i, \varepsilon)\}_{i=0}^{p}$, where we choose $\varepsilon \leqslant \min_i \varepsilon_i$ small enough so that $[a_i - \varepsilon, a_i + \varepsilon]$ is disjoint from the $[a^j - \varepsilon^j, a^j + \varepsilon^j]$. This straight-line homotopy is in the barycentric coordinates: as the intervals are all disjoint, the join of the simplices they describe also lies in $|D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet|$, and so there is a canonical straight line between them. Then we use the obvious homotopy from the data $\{(a_i, \varepsilon)\}_{i=0}^{p}$ to the data $\{(a_i, \varepsilon_i)\}_{i=0}^{p}$ that stretches the ε's. The same construction gives a homotopy $\widehat{H}$ from $F|_{\partial D^n}$ to $\hat{f}$ such that $|r| \circ \widehat{H} = H|_{\partial D^n}$, which is the data we required.

The case when $d = 2n$ and $\mathcal{A}$ is chosen is identical. □

3. Surgery on morphisms

In this section we wish to study the filtration

$$\mathcal{C}_{\theta,L}^{\kappa}(\mathbb{R}^N) \subset \ldots \subset \mathcal{C}_{\theta,L}^{1}(\mathbb{R}^N) \subset \mathcal{C}_{\theta,L}^{0}(\mathbb{R}^N) \subset \mathcal{C}_{\theta,L}^{-1}(\mathbb{R}^N) = \mathcal{C}_{\theta,L}(\mathbb{R}^N)$$

and in particular establish the following theorem. The reader mainly interested in Theorems 1.1 and 1.2 can take $d=2n$, $\theta=\theta^n\colon B\mathrm{O}(2n)\langle n\rangle\to B\mathrm{O}(2n)$, $L\cong D^{2n-1}$ and $N=\infty$ (but the proof does not simplify much in this special case).

THEOREM 3.1. *Suppose that the following conditions are satisfied*:
(i) $2\kappa\leqslant d-2$;
(ii) $\kappa+1+d<N+1$;
(iii) *L admits a handle decomposition only using handles of index at most $d-\kappa-2$. Then the map*

$$B\mathcal{C}^{\kappa}_{\theta,L}(\mathbb{R}^N)\longrightarrow B\mathcal{C}^{\kappa-1}_{\theta,L}(\mathbb{R}^N)$$

is a weak homotopy equivalence.

The proof of Theorem 3.1 consists of performing surgery on morphisms, in order to make them more highly connected relative to their outgoing boundary. Making this idea into a proof has two main ingredients. Firstly, we construct for each morphism in $\mathcal{C}^{\kappa-1}_{\theta,L}$ a contractible space of surgery data. The space is defined in Definition 3.2, and the precise statement is Theorem 3.4. Secondly, we implement the surgery described by the surgery data, using a standard one-parameter family of manifolds defined in §3.2.

In order to motivate some of the more technical constructions, let us first give an informal account of this technique. For simplicity, we suppose that $N=\infty$, $L=\varnothing$ and $\kappa=0$. We first discuss a technique which works when there are no tangential structures to keep track of, and then explain a small modification which makes it work for any tangential structure.

We first apply the equivalence (2.1) to reduce the problem to studying the map

$$BD^0\longrightarrow BD^{-1}$$

of classifying spaces of posets. Let

$$\sigma=(t_0,t_1;a_0,a_1;\varepsilon_0,\varepsilon_1;W)\in BD^{-1}$$

be a point on a 1-simplex (for example), where $(t_0,t_1)\in\Delta^1$ are the barycentric coordinates. We will describe a way of producing a path from its image in $|X_\bullet^{-1}|$ into the subset $|X_\bullet^0|$. The proof of Theorem 3.1 will be a systematic, parameterised version of this construction. If the cobordism $W|_{[a_0,a_1]}$ is already 0-connected relative to its outgoing boundary, then the image of σ in $|X_\bullet^{-1}|$ already lies in the subset $|X_\bullet^0|$, and we are done. If not, we may choose a finite set of distinct points

$$\{f_\alpha\colon *\to W|_{[a_0,a_1]}\}_{\alpha\in\Lambda}$$

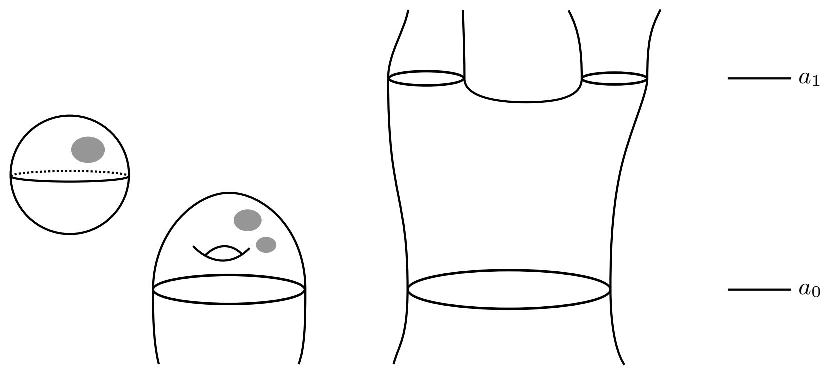

Figure 2. An element of N_1D^{-1} which is not in N_1D^0, together with surgery data.

such that the pair $\big(W|_{[a_0,a_1]}, (W|_{a_1}) \cup \bigcup_{\alpha\in\Lambda} f_\alpha(*)\big)$ is 0-connected. We then choose tubular neighbourhoods of these points to obtain codimension-0 embeddings

$$\hat{f}_\alpha\colon \{-1\}\times\mathbb{R}^d \longrightarrow W|_{[a_0,a_1]},$$

which we can extend to embeddings

$$e_\alpha\colon [-2,0]\times\mathbb{R}^d \longrightarrow \mathbb{R}\times\mathbb{R}^\infty$$

sending $\{-2\}\times\mathbb{R}^d$ into $(-\infty, a_0-\varepsilon_0)\times\mathbb{R}^\infty$ and $\{0\}\times\mathbb{R}^d$ into $(a_1+\varepsilon_1,\infty)\times\mathbb{R}^\infty$. As the original points $f_\alpha(*)$ were distinct, we may suppose that the embeddings e_α are disjoint from each other, and only intersect W in $\{-1\}\times\mathbb{R}^d$. In Figure 2 we have shown a typical example of the case $d{=}2$: The original bordism is not 0-connected relative to its outgoing boundary, but we have chosen the e_α's and depicted the images $e_\alpha(\{-1\}\times\mathbb{R}^d)$ as the shaded discs. (One of the discs in the figure is redundant; it will be important that we allow such redundant surgery data.)

Now on each $e_\alpha([-2,0]\times\mathbb{R}^d)$ we do the surgery move shown in Figure 3, a move similar in spirit, though much simpler, than that described in [GMTW, §6.2]. More precisely, Figure 3 describes a continuous 1-parameter family of d-manifolds

$$\mathcal{P}_t \subset [-2,1]\times\mathbb{R}^d, \quad t\in[0,1],$$

depicted (for $d{=}2$) by its values at times $t=0,\frac{1}{4},\frac{2}{4},\frac{3}{4},1$. The family comes equipped with functions to $\mathbb{R}$, depicted in the figure as the height function (projection onto the vertical axis), such that under the embedding e_α the height $a_0-\varepsilon_0$ corresponds to the bottom of the pictures in Figure 3 and the height $a_1+\varepsilon_1$ corresponds to the dashed line in the figure, and anything above the dashed line will actually end up being forgotten in a moment.

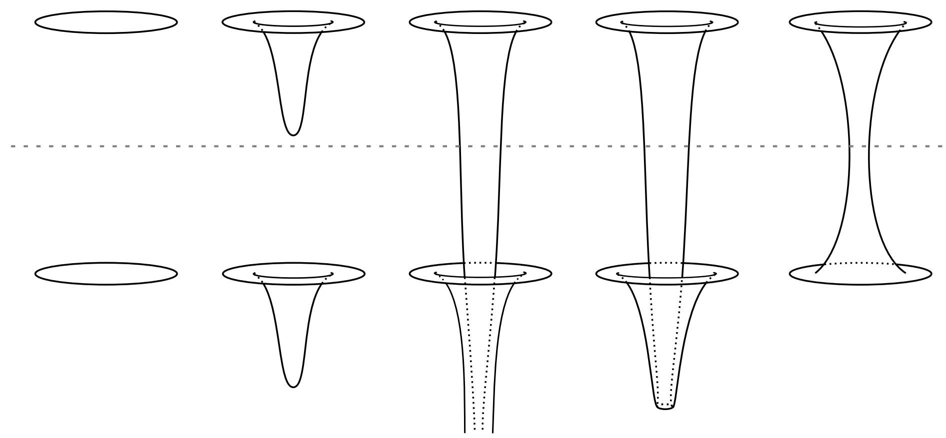

Figure 3. The simple move for surgery on morphisms.

(We find it useful to depict it anyway, in order to better explain the surgery move: The manifold at time 1 is obtained from the manifold at time 0 by performing connected sum, alias zero-surgery.) The family starts at the manifold $\mathcal{P}_0{=}S^0\times\mathbb{R}^d$, and we may cut out each $e_\alpha([-2,0]\times\mathbb{R}^d)$ from W and glue in the part of $\mathcal{P}_t$ below the dashed line, to obtain a one-parameter family of manifolds W_t, each equipped with a height function $W_t\to\mathbb{R}$, with $W_0{=}W$. The values $\{a_0,a_1\}$ do not remain regular throughout this move, so this does not describe a path in the space BD^{-1}. However, it does describe a path in the space $|X_\bullet^{-1}|$. Furthermore, at the end of the move we obtain a manifold $W_1{=}\overline{W}$ such that $(\overline{W}|_{[a_0,a_1]},\overline{W}|_{a_1})$ is 0-connected, and hence a point in $|X_\bullet^0|$. By Proposition 2.20, this proves that $\pi_0(BD^0)\to\pi_0(BD^{-1})$ is surjective, as required.

This surgery move generalises easily to the case when N is finite (but sufficiently large), $L\neq\varnothing$ and $\kappa>0$ (the analogue of the surgery move will start with $S^\kappa\times\mathbb{R}^{d-\kappa}$ and end with $\mathbb{R}^{\kappa+1}\times S^{d-\kappa-1}$, equipped with appropriate height functions). However, it does not generalise well to the case of arbitrary tangential structures (to understand how it can fail, we suggest that the reader attempt to impose a family of framings to the family of 2-manifolds in Figure 3). One way to fix this would be to use the surgery move described in [GMTW, §6.2], but that does not seem to generalise to $\kappa>0$. Instead we modify the surgery move in Figure 3 as shown in Figure 4.

The refined surgery move still begins with $S^0\times\mathbb{R}^d$, but it ends with $(\mathbb{R}^1\times S^{d-1})\setminus\{p_1\}$ for a point $p_1\in\mathbb{R}^1\times S^{d-1}$. The height function is modified so that it goes to $-\infty$ (or at least below $a_0-\varepsilon_0$) at p_1. As we shall see (in the proof of Proposition 3.6, where we also explain the analogous process for $\kappa>0$) there is a canonical way of extending any tangential structure on $\{-1\}\times\mathbb{R}^d$ to the resulting one-parameter family of manifolds.

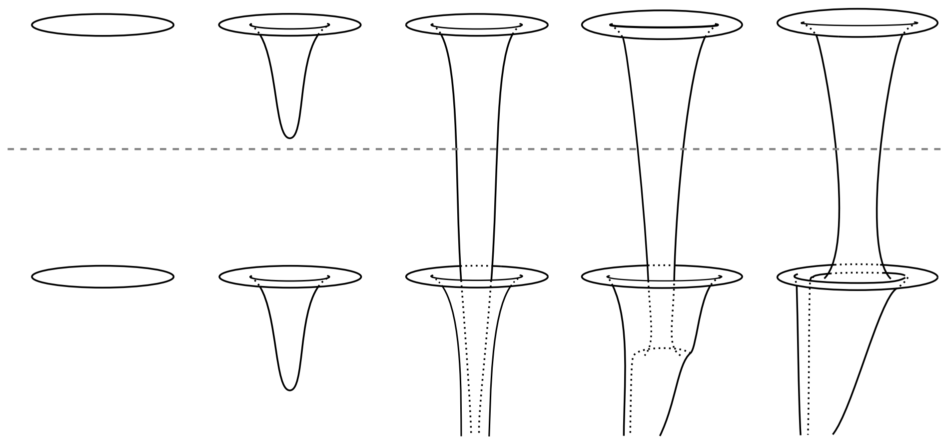

Figure 4. The refined move for surgery on morphisms.

3.1. Surgery data

In order to implement the ideas discussed above, we will fatten the semi-simplicial space $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_\bullet$ up to a bi-semi-simplicial space $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$ which includes suitable surgery data. The space $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$ is described in Definition 3.3 below, using the following notation. Let $V\subset\overline{V}\subset\mathbb{R}^{\kappa+1}\times\mathbb{R}^{d-\kappa}$ be the subspaces

$$V=(-2,0)\times\mathbb{R}^d \quad \text{and} \quad \overline{V}=[-2,0]\times\mathbb{R}^d,$$

and let $h\colon\overline{V}\to[-2,0]\subset\mathbb{R}$ denote the projection to the first coordinate, which we call the *height function*. Let $\partial_-D^{\kappa+1}\subset\partial D^{\kappa+1}$ denote the lower hemisphere (explicitly, it is given by $\partial_-D^{\kappa+1}=\partial D^{\kappa+1}\cap h^{-1}([-1,0])$). We shall also use the notation $[p]^\vee=\Delta([p],[1])$ when $[p]\in\Delta_{\mathrm{inj}}$. The elements of $[p]^\vee$ are in bijection with $\{0,...,p+1\}$, using the convention that $\varphi\colon[p]\to[1]$ corresponds to the number i with $\varphi^{-1}(1)=\{i,i+1,...,p\}$. Finally, we fix once and for all an infinite set Ω.

Definition 3.2. Let $x=(a,\varepsilon,(W,\ell_W))\in D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$ and define $Z_0(x)$ to be the set of triples (Λ,δ,e), where $\Lambda\subset\Omega$ is a finite set, $\delta\colon\Lambda\to[p]^\vee$ is a function, and

$$e\colon\Lambda\times\overline{V}\hookrightarrow\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$$

is an embedding, satisfying the following conditions:

(i) on every subset $(x_1\circ e|_{\{\lambda\}\times\overline{V}})^{-1}(a_k-\varepsilon_k,a_k+\varepsilon_k)\subset\{\lambda\}\times\overline{V}$, the height function $x_1\circ e$ coincides with the height function h up to an affine transformation;

(ii) e sends $\Lambda\times h^{-1}(0)$ into $x_1^{-1}(a_p+\varepsilon_p,\infty)$;

(iii) for $i>0$, e sends $\delta^{-1}(i)\times h^{-1}\left(-\frac{3}{2}\right)$ into $x_1^{-1}(a_{i-1}+\varepsilon_{i-1},\infty)$;

(iv) e sends $\Lambda\times h^{-1}(-2)$ into $x_1^{-1}(-\infty, a_0-\varepsilon_0)$;
(v) $e^{-1}(W)=\Lambda\times\partial_- D^{\kappa+1}\times\mathbb{R}^{d-\kappa}$;
(vi) writing $D_i=e(\delta^{-1}(i)\times\partial_- D^{\kappa+1}\times\{0\})$ for $i\in[p]^\vee$, the pair

$$(W|_{[a_{i-1},a_i]}, W|_{a_i}\cup D_i|_{[a_{i-1},a_i]})$$

is κ-connected for each $i\in\{1, ..., p\}$.

A typical example of a surgery datum is (partially) depicted in Figure 2. In that figure $d=2$ and $\kappa=0$, and only the image $e(\Lambda\times\partial_- D^{\kappa+1}\times\mathbb{R}^{d-\kappa})\subset W$ is shown. Let us explain the role that elements of $Z_0(x)$ will play in the proof of Theorem 3.1. In §3.3 we shall describe a surgery process which to an element of $Z_0(x)$ associates a path from the image of x under the forgetful map $D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p\to X^{\kappa-1}_p$ to a point in the subspace X^κ_p, formalising the one-parameter family depicted in Figure 4 (the image of e is automatically disjoint from L). Conflating the three different simplicial spaces modelling $B\mathcal{C}^\kappa_{\theta,L}(\mathbb{R}^N)$ from §2.4, §2.6 and §2.8, we can thus think of an element of $Z_0(x)$ as a surgery datum for making morphisms κ-connected relative to their outgoing boundary (they start out being only $(\kappa-1)$-connected). To explain this in more detail, it is helpful to write $\Lambda_i=\delta^{-1}(i)$ for $i\in[p]^\vee\cong\{0, ..., p+1\}$. For $0<i<p+1$, the restriction of e to $\Lambda_i\times\overline{V}$ is then the surgery data which will be used to make the bordism $W|_{[a_{i-1},a_i]}$ κ-connected relative to its outgoing boundary. (The embeddings associated with the outer values $i=0$ and $i=p+1$ play a more technical role; they will make the surgery construction compatible with face maps in the p direction.)

In order to prove Theorem 3.1 we would like to have a contractible space of surgery data, but $Z_0(x)$ is usually far from contractible (we only defined $Z_0(x)$ as a set, but it would be disconnected in any reasonable topology). To fix that, we extend the definition to a semi-simplicial set $Z_\bullet(x)$ whose set of q-simplices is the subset $Z_q(x)\subset Z_0(x)^{q+1}$ consisting of $(q+1)$-tuples which are disjoint (i.e. the subsets $\Lambda\subset\Omega$ are disjoint and the maps e have disjoint images). Allowing also x to vary gives rise to a bi-semi-simplicial space $D^\kappa_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$, whose (p,q)-simplices are p-chains in the poset $D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)$ equipped with $(q+1)$-tuply redundant surgery data. To fix notation we spell this out in the following definition.

Definition 3.3. Let $x=(a,\varepsilon,(W,\ell_W))\in D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$ and $q\geqslant 0$, define $Z_q(x)$ to be the set of triples (Λ,δ,e), where $\Lambda\subset\Omega$ is a finite set, $\delta\colon\Lambda\to[p]^\vee\times[q]$ is a function, and

$$e\colon\Lambda\times\overline{V}\longhookrightarrow\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$$

is an embedding, subject to the requirement that for each $j\in[q]$, the restriction of e to $\delta^{-1}([p]^\vee\times\{j\})\times\overline{V}$ defines an element of $Z_0(x)$. We shall write $\Lambda_{i,j}=\delta^{-1}(i,j)$ and $e_{i,j}=e|_{\Lambda_{i,j}\times\overline{V}}$ for $i\in[p]^\vee$ and $j\in[q]$.

We then define a bi-semi-simplicial space $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$ as a set by

$$D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}=\{(x,y):x\in D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p \text{ and } y\in Z_q(x)\}$$

topologised as a subspace of

$$D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p\times\bigg(\coprod_{\Lambda\subset\Omega}C^{\infty}(\Lambda\times\overline{V},\mathbb{R}^{N+1})\bigg)^{(p+2)(q+1)}.$$

The space $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$ is functorial in $[p]\in\Delta_{\mathrm{inj}}$ by composing each $\delta_j:\Lambda_j\to[p]^\vee$ with the induced map $[p']^\vee\to[p]^\vee$ and functorial in $[q]\in\Delta_{\mathrm{inj}}$ in the same way as $Z_\bullet(x)$. Explicitly, the face map d_i in the q direction forgets the embeddings $e_{*,i}$ and in the p direction takes the union of $e_{i,*}$ and $e_{i+1,*}$. We shall write $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,-1}=D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$, and there is an augmentation map $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}\to D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,-1}$ which forgets all surgery data.

The main result concerning this bi-semi-simplicial space is the following, whose proof we defer until §6.

Theorem 3.4. *Under the assumptions of Theorem 3.1, the augmentation map*

$$D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}\longrightarrow D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_{\bullet}$$

induces a weak homotopy equivalence after geometric realisation.

In fact, we shall prove this theorem with condition (i) of Theorem 3.1 replaced by the weaker condition $2\kappa\leqslant d-1$. The stronger assumption $2\kappa\leqslant d-2$ will be used in Lemma 3.7.

3.2. The standard family

We will now construct a one-parameter family $\mathcal{P}_t$, $t\in[0,1]$, of submanifolds of the space $V=(-2,0)\times\mathbb{R}^d$ which formalises the family of manifolds depicted in Figure 4 (more precisely, it corresponds to the part of Figure 4 which is below the dashed line). In the process, we will also define a family $\mathcal{P}'_t$ formalising the family depicted in Figure 3. (This simpler family would suffice for proving Theorem 3.1 in the case without tangential structures; we shall not actually use the simpler family, but it is perhaps helpful to keep in mind.) The manifolds $\mathcal{P}_0$ and $\mathcal{P}_1$ (and $\mathcal{P}'_1$) shall be defined by intersecting $V\subset\mathbb{R}^{d+1}$ with submanifolds of the larger space $\mathbb{R}^{d+1}$, denoted $\widetilde{\mathcal{P}}_0$, $\widetilde{\mathcal{P}}_1$, and $\widetilde{\mathcal{P}}'_1$. (These larger manifolds include the parts of Figures 3 and 4 above the dashed line.) For following the construction, it might also be useful to have the case $d=1$ and $\kappa=0$ in mind, which is depicted in Figure 5 (although these dimensions do not satisfy the inequality $2\kappa\leqslant d-2$).

First define the element $\widetilde{\mathcal{P}}_0 \in \Psi_d(\mathbb{R}\times\mathbb{R}^\kappa\times\mathbb{R}^{d-\kappa})$ as

$$\widetilde{\mathcal{P}}_0 = \partial D^{\kappa+1}\times\mathbb{R}^{d-\kappa}.$$

This manifold is depicted in Figure 5a (for $d=1$ and $\kappa=0$) and the first frame in Figures 3 and 4 (for $d=2$ and $\kappa=0$). We then choose a function $\varphi: [0,\infty)\to[0,\infty)$ which is the identity function on a neighbourhood of $\left[\frac{1}{2},\infty\right)$, takes value $\frac{1}{4}$ near 0, and has $\varphi''\geqslant 0$. The map defined by

$$\begin{aligned} g': \mathbb{R}^{\kappa+1}\times\partial D^{d-\kappa} &\longrightarrow D^{\kappa+1}\times\mathbb{R}^{d-\kappa}, \\ (x,y) &\longmapsto \left(\frac{x}{\varphi(|x|)}, \varphi(|x|)y\right), \end{aligned}$$

is an embedding with inverse $(u,v)\mapsto(|v|u, v/|v|)$, and we shall write $\widetilde{\mathcal{P}}_1'$ for its image, which is depicted in Figure 5b (for $d=1$ and $\kappa=0$) and in the last frame of Figure 3 (for $d=2$ and $\kappa=0$).

Finally, we define $\widetilde{\mathcal{P}}_1$ by modifying the embedding g' in the following way: The subset $\widetilde{\mathcal{P}}_1'$ agrees with the subset $\partial D^{\kappa+1}\times\mathbb{R}^{d-\kappa}$ in a neighbourhood of the region defined by $|v|\geqslant\frac{1}{2}$, and in particular it contains the $(d-\kappa-1)$-sphere defined by $u=-e_1$ and $|v|=\frac{1}{2}$, which we shall temporarily denote by S. This sphere is "level" in the sense that the height function $u_1: \widetilde{\mathcal{P}}_1'\to[-1,1]$ takes the constant value -1 on S. We shall modify the embedding g' by "tilting" a small neighbourhood of S so that the height function on the tilted sphere is instead the Morse function $v_1-\frac{3}{2}$. The resulting subset

$$\widetilde{\mathcal{P}}_1 \subset [-2,1]\times D^\kappa\times\mathbb{R}^{d-\kappa}$$

is depicted in the last frame of Figure 4 (in the case $d=2$ and $\kappa=0$) and in Figure 5c (in the case $d=1$ and $\kappa=0$). To define $\widetilde{\mathcal{P}}_1$ more precisely, we first pick a bump function $\lambda: \mathbb{R}\to[0,1]$ supported in a small neighbourhood of 0, increasing on $(-\infty,0]$, decreasing on $[0,\infty)$, and having $\lambda^{-1}(1)=\{0\}$, and define a bump function $\tau: \mathbb{R}^{\kappa+1}\times\mathbb{R}^{d-\kappa}\to[0,1]$ supported in a small neighbourhood of S as the product

$$\tau(u,v) = \lambda(u_1+1)\lambda(u_2)\,...\,\lambda(u_{\kappa+1})\lambda(2|v|-1).$$

We choose λ with support in $[-1,1]$, and small enough that the support of τ will be contained in the region where $\widetilde{\mathcal{P}}_1'$ agrees with $\partial D^{\kappa+1}\times\mathbb{R}^{d-\kappa}$, and we shall verify presently that the smooth function defined by

$$\begin{aligned} j: D^{\kappa+1}\times\mathbb{R}^{d-\kappa} &\longrightarrow [-2,1]\times D^\kappa\times\mathbb{R}^{d-\kappa}, \\ (u,v) &\longmapsto (u,v)+(v_1-1)\tau(u,v)e_1, \end{aligned}$$

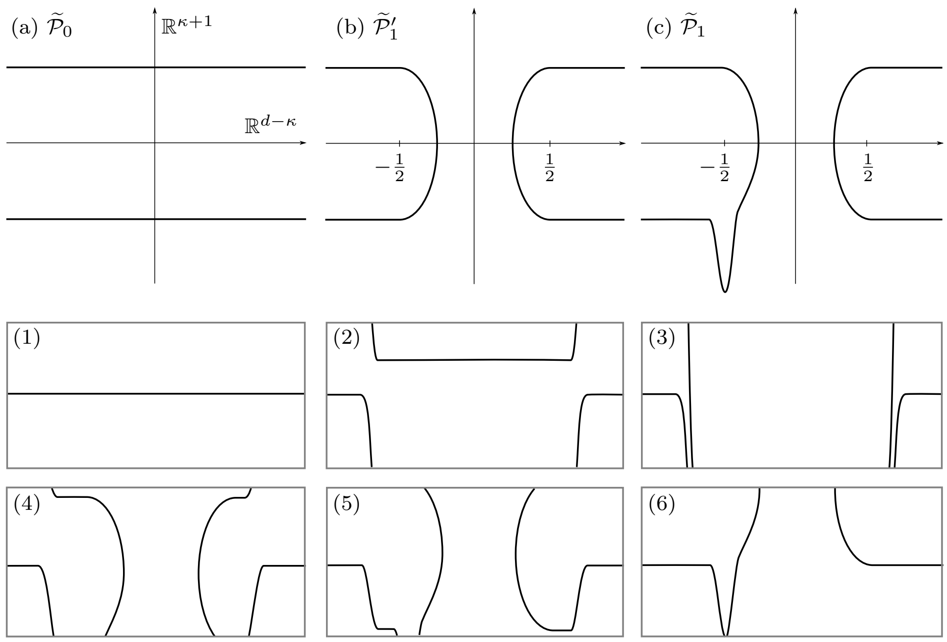

Figure 5. (a)–(c) The submanifolds $\widetilde{\mathcal{P}}_0$, $\widetilde{\mathcal{P}}_1'$ and $\widetilde{\mathcal{P}}_1$ of $\mathbb{R}^{\kappa+1}\times\mathbb{R}^{d-\kappa}$ in the case $d=1$ and $\kappa=0$. (1)–(6) The path $\mathcal{P}_t$ of submanifolds of V.

is an embedding. We then let $\widetilde{\mathcal{P}}_1$ be the image of the composition

$$g=j\circ g'\colon\mathbb{R}^{\kappa+1}\times\partial D^{d-\kappa}\longrightarrow[-2,1]\times D^{\kappa}\times\mathbb{R}^{d-\kappa}$$

and define

$$\mathcal{P}_0,\mathcal{P}_1\in\Psi_d(V)$$

by intersecting the manifolds $\widetilde{\mathcal{P}}_0$ and $\widetilde{\mathcal{P}}_1$ with the open set $V=(-2,0)\times\mathbb{R}^d$.

To see that the function j is indeed an embedding, we first note that it is the restriction of a function $[-1,1]\times\mathbb{R}^d\to\mathbb{R}\times\mathbb{R}^d$ defined by the same formula. Since the extended function commutes with the projection to the $\mathbb{R}^d$ coordinates, it suffices to prove that

$$u_1\longmapsto j_1(u_1,u_2,...,u_{\kappa+1},v)$$

defines an embedding $[-1,1]\to\mathbb{R}$ for any $(u_2,...,u_{\kappa+1},v)\in\mathbb{R}^d$, where j_1 denotes the first coordinate of j. We calculate

$$\frac{\partial j_1}{\partial u_1}(u,v)=1+\lambda'(u_1+1)\big((v_1-1)\lambda(2|v|-1)\big)\lambda(u_2)\,...\,\lambda(u_{\kappa+1}).$$

By the assumption on λ, we will have $\lambda'(u_1+1)\leqslant 0$ as $u_1\geqslant -1$. We can also conclude that $(v_1-1)\lambda(2|v|-1)\leqslant 0$, as $\lambda(2|v|-1)=0$ unless $|v|\leqslant 1$, in which case $v_1-1\leqslant 0$. Since $\lambda(u_2)\,...\,\lambda(u_{\kappa+1})\geqslant 0$, we conclude that $\partial j_1/\partial u_1\geqslant 1$.

To construct $\mathcal{P}_t \in \Psi_d(V)$ for intermediate values of $t \in [0,1]$, we first observe that $\widetilde{\mathcal{P}}_0$ and $\widetilde{\mathcal{P}}_1$ agree near the subset $|v| \geqslant 1$. (And $\widetilde{\mathcal{P}}'_1$ agrees with $\widetilde{P}_0$ near the larger subset $|v| \geqslant \frac{1}{2}$.) Starting with the two submanifolds $\widetilde{\mathcal{P}}_0, \widetilde{\mathcal{P}}_1 \subset \mathbb{R} \times \mathbb{R}^\kappa \times \mathbb{R}^{d-\kappa}$, we then pull the entire region $\{(u,v) : |v| < 1\}$ downwards, much in the same fashion as we tilted the sphere S, i.e. we compose with an ambient diffeomorphism which subtracts a non-negative amount from the first coordinate. We pull far enough so that the region where the submanifolds may disagree is moved completely outside of V. This will give two one-parameter families of submanifolds which, upon restricting to V, give two paths in $\Psi_d(V)$ starting at $\mathcal{P}_0$ and $\mathcal{P}_1$ and ending at the same point in $\Psi_d(V)$. Concatenating one path with the reverse of the other, we get the desired path from $\mathcal{P}_0$ to $\mathcal{P}_1$.

Spelling this pulling-down process out in a little more detail, we first choose a function $\varrho: [0,\infty) \to [0,\infty)$ taking the value 1 near $[0,1]$, the value 0 near $[2,\infty)$, and which is strictly decreasing on $\varrho^{-1}((0,1))$. We then define embeddings

$$
\begin{aligned}
H_t: \mathbb{R} \times \mathbb{R}^\kappa \times \mathbb{R}^{d-\kappa} &\longrightarrow \mathbb{R} \times \mathbb{R}^\kappa \times \mathbb{R}^{d-\kappa}, \\
(u,v) &\longmapsto (u,v) - t\varrho(|v|)e_1,
\end{aligned}
$$

which for all t restrict to the identity near the region defined by $|v| \geqslant 2$. Define one-parameter families of manifolds by

$$
\begin{aligned}
\mathcal{P}_t^0 &= V \cap H_t(\widetilde{\mathcal{P}}_0) = (H_{-t}|_V)^{-1}(\widetilde{\mathcal{P}}_0), \\
\mathcal{P}_t^1 &= V \cap H_t(\widetilde{\mathcal{P}}_1) = (H_{-t}|_V)^{-1}(\widetilde{\mathcal{P}}_1).
\end{aligned}
$$

The second description shows that these are closed subsets of V and describe continuous functions $\mathbb{R} \to \Psi_d(V)$. It is easy to see that we have $\mathcal{P}_t^0 = \mathcal{P}_t^1 \in \Psi_d(V)$ for $t \geqslant 3$, and we then define the path $\mathcal{P}_t$ as the concatenation

$$
\mathcal{P}_0 = \mathcal{P}_0^0 \rightsquigarrow \mathcal{P}_3^0 = \mathcal{P}_3^1 \rightsquigarrow \mathcal{P}_0^1 = \mathcal{P}_1
$$

in $\Psi_d(V)$, reparameterised so that the path has length 1. We collect the most important properties of this family in Proposition 3.6 below.

Remark 3.5. The one-parameter family of d-manifolds described above is in fact strongly related to the more usual description of performing κ-surgery on a d-manifold, which we briefly recall (see e.g. [Mi1, pp. 12 ff.]). Let $Q(u,v) = -|u|^2 + |v|^2$, where as usual $(u,v) \in \mathbb{R}^{\kappa+1} \times \mathbb{R}^{d-\kappa}$. Then the inverse image $Q^{-1}(s)$ is a smooth d-manifold for $s \neq 0$, diffeomorphic to $\partial D^{\kappa+1} \times \mathbb{R}^{d-\kappa}$ for $s < 0$ and to $\mathbb{R}^{\kappa+1} \times \partial D^{d-\kappa}$ for $s > 0$. The classical description of a κ-surgery consists of cutting out a copy of $Q^{-1}(-3)$ from a d-manifold and replacing it with $Q^{-1}(3)$. The *trace* of the surgery is a cobordism equipped with a

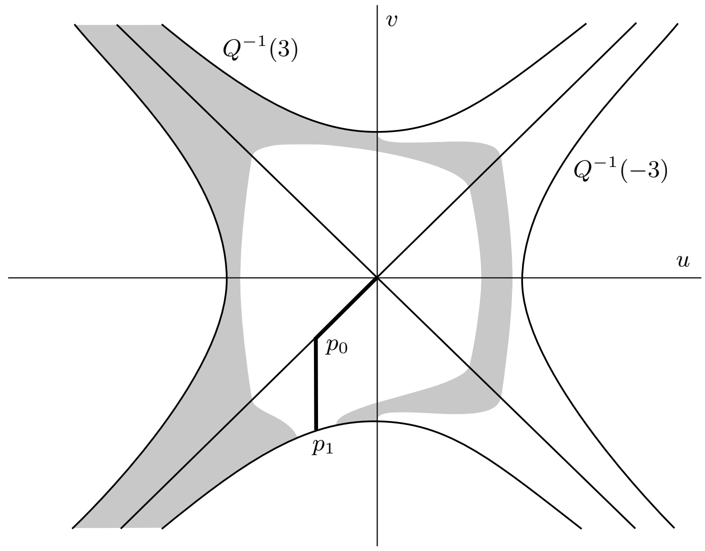

Figure 6. The image of the embedding $\mathcal{P}\to Q^{-1}([-3,3])$.

Morse function with a single critical point, interpolating between the original manifold and the surgered one.

To explain the relation between this classical picture and the one-parameter family we have defined above, we shall exhibit a continuous family of open embeddings $\mathcal{P}_t^0\to Q^{-1}(t-3)$ and $\mathcal{P}_t^1\to Q^{-1}(3-t)$ gluing to a continuous family of embeddings $\mathcal{P}_t\to Q^{-1}(6t-3)$. Indeed, for $t\in[0,3]$ the function $(u,v)\mapsto H_t(u/|u|,v)$ defines a diffeomorphism from $Q^{-1}(t-3)$ to $H_t(\widetilde{\mathcal{P}}_0)$ which restricts to a diffeomorphism from an open subset of $Q^{-1}(t-3)$ to $\mathcal{P}_t^0$. Similarly the function $(u,v)\mapsto H_t\circ g(u,v/|v|)$ defines a diffeomorphism from $Q^{-1}(3-t)$ to $H_t(\widetilde{\mathcal{P}}_1)$ restricting to a diffeomorphism from an open subset of $Q^{-1}(3-t)$ to $\mathcal{P}_t^1$. The inverses of these diffeomorphisms give the desired smooth embeddings $\mathcal{P}_t^0\to Q^{-1}(t-3)$ and $\mathcal{P}_t^1\to Q^{-1}(3-t)$, which fit together at $t=3$.

The one-parameter family $t\mapsto\mathcal{P}_t$ has "total space" given by

$$\mathcal{P}=\{(t,x)\in[0,1]\times V : x\in\mathcal{P}_t\},$$

and the above remarks give a continuous embedding $\mathcal{P}\to Q^{-1}([-3,3])$. Let us briefly discuss the image of this embedding, which is depicted as the shaded area in Figure 6 in the case $d=1$ and $\kappa=0$.

The point $p_1=\left(-\frac{1}{2},0,-\frac{1}{2}\sqrt{13},0\right)$ has $Q(p_1)=3$, and corresponds to the point in $\widetilde{\mathcal{P}}_1$ at the bottom of the tilted sphere, i.e. the global minimum of the height function $\widetilde{\mathcal{P}}_1\to[-2,1]$. Since this bottom point is outside V, the point p_1 is not in the image of the embedding of $\mathcal{P}$. Likewise, the pulled-down global minimum of the height function

$H_3(\widetilde{\mathcal{P}}_1)\to[-5,1]$ corresponds to the point $p_0=\left(-\frac{1}{2},0,-\frac{1}{2},0\right)\in Q^{-1}(0)$, and the entire straight line from p_0 to p_1 is disjoint from the image of the embedding $\mathcal{P}\to Q^{-1}([-3,3])$. Finally, the straight line from 0 to p_0, which takes place inside $Q^{-1}(0)$, is also disjoint from the embedding $\mathcal{P}$. Thus we have obtained a homeomorphism from $\mathcal{P}$ to an open subset of the set

$$\mathcal{Q}=Q^{-1}([-3,3])\setminus([0,p_0]\cup[p_0,p_1]), \tag{3.1}$$

which is contractible.

PROPOSITION 3.6. *For $d\geqslant 2$ and $2\kappa\leqslant d-1$, the one-parameter family $\mathcal{P}_t\in\Psi_d(V)$, defined for $t\in[0,1]$, has the following properties*:

(i) *The height function, i.e. the restriction of $h\colon V\to(-2,0)$ to $\mathcal{P}_t\subset V$, has isolated critical values.*

(ii) *$\mathcal{P}_0=\operatorname{int}(\partial_-D^{\kappa+1})\times\mathbb{R}^{d-\kappa}$, where $\partial_-D^{\kappa+1}=\partial D^{\kappa+1}\cap([-1,0]\times\mathbb{R}^{\kappa})$.*

(iii) *Independently of $t\in[0,1]$ we have*

$$\mathcal{P}_t\setminus(\mathbb{R}^{\kappa+1}\times B_3^{d-\kappa}(0))=\operatorname{int}(\partial_-D^{\kappa+1})\times(\mathbb{R}^{d-\kappa}\setminus B_3^{d-\kappa}(0)).$$

For ease of notation we write $\mathcal{P}_t^{\partial}$ for this closed subset of $\mathcal{P}_t$.

(iv) *For all t and each pair of regular values $-2<a<b<0$ of the height function, the pair*

$$(\mathcal{P}_t|_{[a,b]},\mathcal{P}_t|_b\cup\mathcal{P}_t^{\partial}|_{[a,b]}) \tag{3.2}$$

is κ-connected.

(v) *For each pair of regular values $-2<a<b<0$ of the height function, the pair*

$$(\mathcal{P}_1|_{[a,b]},\mathcal{P}_1|_b)$$

is κ-connected.

Furthermore, if $\mathcal{P}_0$ is equipped with a θ-structure ℓ we can upgrade this, continuously in ℓ, to a one-parameter family $\mathcal{P}_t(\ell)\in\Psi_\theta(V)$ starting from $(\mathcal{P}_0,\ell)$ such that

(iii′) *The path $\mathcal{P}_t(\ell)$ is constant as θ-manifolds near $\mathcal{P}_t^{\partial}$.*

Proof. We have seen properties (i)–(iii) during the construction (the statement in (iii) would still be true with 3 replaced by 2, but we wish to emphasise the smaller set). For property (iv) we consider two cases depending on the value of a. In the case $a>-1$, the pair (3.2) is homotopy equivalent to the pair

$$(\mathcal{P}_t|_{[a,b]},\mathcal{P}_t|_b),$$

using e.g. the gradient flow trajectories of h to deform $\mathcal{P}_t^{\partial}|_{[a,b]}$ back to $\mathcal{P}_t^{\partial}|_b$. In the case $a<-1$ we consider the *modified height function*, defined using the coordinates $(u,v)\in$

$\mathbb{R}^{\kappa+1}\times\mathbb{R}^{d-\kappa}$ as $\bar{h}(u,v)=h(u,v)+\lambda(|v|)$, where $\lambda:[0,\infty)\to[0,\infty)$ is a smooth function which is 0 on $[0,4]$ and restricts to a diffeomorphism $(4,\infty)\to(0,\infty)$. This modification ensures that $\bar{h}$ is a proper function on $\mathcal{P}_t$. With this definition $\mathcal{P}_t\cap\bar{h}^{-1}([a,b])$ is contained in $\mathcal{P}_t\cap h^{-1}([a,b])$, and $\mathcal{P}_t\cap\bar{h}^{-1}(b)$ is contained in $\mathcal{P}_t|_b\cup\mathcal{P}_t^\partial|_{[a,b]}$, which may be seen as follows. Let $(u,v)\in P_t\cap\bar{h}^{-1}([a,b])$: if $|v|\leqslant 4$ then $h(u,v)=\bar{h}(u,v)\in[a,b]$ and we are done; if $|v|>2$ then (u,v) lies in the part of $\mathcal{P}_t$ that does not vary with t, i.e. $\partial D^{\kappa+1}\times\mathbb{R}^{d-\kappa}$, and so $h(u,v)=u_1>-1>a$ by assumption, and $b\geqslant\bar{h}(u,v)>h(u,v)$ too. The same reasoning treats $\mathcal{P}_t\cap\bar{h}^{-1}(b)$.

We claim that the inclusion of pairs

$$(\mathcal{P}_t\cap\bar{h}^{-1}([a,b]),\mathcal{P}_t\cap\bar{h}^{-1}(b))\longrightarrow(\mathcal{P}_t|_{[a,b]},\mathcal{P}_t|_b\cup\mathcal{P}_t^\partial|_{[a,b]}) \tag{3.3}$$

is a homotopy equivalence. To define a homotopy inverse, we first consider the continuous, piecewise smooth function $\varrho_t:[0,\infty)\to(0,\infty)$ defined for $t\leqslant b$ by

$$\begin{aligned}\varrho_t(s)&=1 &&\text{for } s\in[0,2],\\ \varrho_t(s)&=\frac{\lambda^{-1}(b-t)}{s} &&\text{for } s\in[3,\infty),\end{aligned}$$

and by linear interpolation for $s\in[2,3]$. Then the function $(u,v)\mapsto(u,v\varrho_{u_1}(|v|))$ restricts to a homotopy inverse of (3.3), where both homotopies are given by straight lines in $\mathbb{R}^{d+1}$.

In either case, the connectivity question is reduced to studying the inverse image of an interval relative to its outgoing boundary and can be studied as in ordinary Morse theory one critical level at a time. The proof of (iv) will be finished once we establish that for each critical value of $\bar{h}:\mathcal{P}_t\to\mathbb{R}$ in the interval (a,b), the function can be perturbed in a neighbourhood of the critical set contained in $\bar{h}^{-1}((a,b))$ to a Morse function with no more than one critical point, and of index at most $d-\kappa-1$. (In the case $a>-1$ we have $h=\bar{h}$ near any critical point of h, so it suffices to consider $\bar{h}$.) It is easy to verify that $\bar{h}:\mathcal{P}_t^0\to\mathbb{R}$ has at most two critical values in $(-2,0)$. One critical value moves with t and is homotopically Morse of index 0 for $0\leqslant t<1$ and index κ for $1<t<3$ (meaning that the function can be perturbed to a Morse function with one critical point of that index). The other is at -1 and can be cancelled (meaning that the function can be perturbed to a non-singular function there). Since $2\kappa\leqslant d-1$ and hence $\kappa\leqslant d-\kappa-1$, the index is at most $d-\kappa-1$ as claimed. Similarly, one verifies that $\bar{h}:\mathcal{P}_t^1\to\mathbb{R}$ has at most two critical values in $(-2,0)$, one of which is -1 and can be cancelled, the other of which moves with t and is homotopically Morse of index $d-\kappa-1$.

Property (v) can be proved in a similar way. In the case $a<-1<b$ the pair is a relative $(d-1)$-cell, so it is $(d-2)$-connected and hence κ-connected (since $d\geqslant 2$ and $2\kappa\leqslant d-1$). In all other cases the inclusion $\mathcal{P}_t|_b\to\mathcal{P}_t|_{[a,b]}$ is a homotopy equivalence.

To establish the extra properties which can be obtained given a θ-structure ℓ on $\mathcal{P}_0=\mathrm{int}(\partial_- D^{\kappa+1})\times\mathbb{R}^{d-\kappa}$, we again use the "total space" $\mathcal{P}=\{(t,x)\in[0,1]\times\mathbb{R}^{d+1}:x\in\mathcal{P}_t\}$ and its identification with an open subset of the manifold $\mathcal{Q}$ of (3.1). The tangent bundles $T\mathcal{P}_t$ assemble to a d-dimensional vector bundle $T_v\mathcal{P}\to\mathcal{P}$ which then becomes identified with the restriction of the vector bundle $T_v\mathcal{Q}=\mathrm{Ker}(DQ: T\mathcal{Q}\to T[-3,3])$ and since both $\mathcal{P}_0$ and $\mathcal{Q}$ are contractible, there is no obstruction to picking a vector bundle map $r: T_v\mathcal{Q}\to T\mathcal{P}_0$ which is the identity (with respect to the identifications) over $\mathcal{P}_0$ and each $\mathcal{P}_t^\partial=\mathcal{P}_0^\partial\subset\mathcal{P}_0$. We can then restrict r to $r_t: T\mathcal{P}_t\to T\mathcal{P}_0$ and let $\mathcal{P}_t(\ell)$ have the θ-structure $\ell\circ r_t$. $\square$

Let $(a,\varepsilon,(W,\ell_W),e)\in D^\kappa_{\theta,L}(\mathbb{R}^N)_{p,0}$, with $e=\{e_{i,0}\}_{i=0}^{p+1}$ (where we omit the set Λ and the function $\delta:\Lambda\to[p]^\vee$ from the notation). We construct a one-parameter family of θ-manifolds

$$\mathcal{K}_e^t(W,\ell_W)\in\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N),\quad t\in[0,1],$$

by letting it be equal to $W|_{(a_0-\varepsilon_0,a_p+\varepsilon_p)}$ outside of the images of the $e_{i,0}|_{\Lambda_{i,j}\times V}$, and on each $e_{i,0}(\{\lambda\}\times V)$ we let it be given by $e_{i,0}(\{\lambda\}\times\mathcal{P}_t(\ell_W\circ De_{i,0}))$. This gives a θ-manifold, as, by the properties established above, $\mathcal{P}_t(\ell_W\circ De_{i,0})$ and $\mathcal{P}_0(\ell_W\circ De_{i,0})$ agree as θ-manifolds near the set $(-2,0)\times\mathbb{R}^\kappa\times(\mathbb{R}^{d-\kappa}\setminus B_3^{d-\kappa}(0))$.

Lemma 3.7. *Let $2\kappa\leqslant d-2$. The tuple $(a,\varepsilon,\mathcal{K}_e^t(W,\ell_W))$ is an element of $X_p^{\kappa-1}$. If $t=1$ or $(W,\ell_W)\in D^\kappa_{\theta,L}(\mathbb{R}^N)_p$, then $(a,\varepsilon,\mathcal{K}_e^t(W,\ell_W))$ lies in the subspace $X_p^\kappa\subset X_p^{\kappa-1}$.*

Proof. We must verify conditions (i)–(v) of Definition 2.18 (with $l=-1$). Condition (i) is true by definition, and certainly (ii) is satisfied as the embeddings $e_{i,0}$ are disjoint from $\mathbb{R}\times L$. For (iii) and (v) there is nothing to say.

For (iv), consider regular values $a<b\in\bigcup_{i=0}^p(a_i-\varepsilon_i,a_i+\varepsilon_i)$ of the height function

$$x_1: W_t=\mathcal{K}_e^t(W,\ell_W)\longrightarrow\mathbb{R}.$$

The cobordism $W_t|_{[a,b]}$ is obtained from $W|_{[a,b]}$ by cutting out embedded images of cobordisms $\mathcal{P}_0|_{[a_\lambda,b_\lambda]}$ indexed by $\lambda\in\Lambda=\coprod_{i=0}^{p+1}\Lambda_{i,0}$ and gluing in $\mathcal{P}_t|_{[a_\lambda,b_\lambda]}$, where $a_\lambda<b_\lambda$ are regular values of the height function on $\mathcal{P}_0$ and $\mathcal{P}_t$. If we denote by X the complement of the embedded $e_{i,0}(\mathrm{int}(\partial_- D^{\kappa+1})\times B_3^{d-\kappa}(0))$ in the manifold $W|_{[a,b]}$, there are homotopy pushout squares

$$\begin{array}{ccc} X|_b & \longrightarrow & W_t|_b \\ \downarrow & & \downarrow \\ X & \longrightarrow & W_t|_b\cup X \end{array}$$

and

$$\begin{tikzcd} \coprod_{\lambda\in\Lambda}(\mathcal{P}_t|_{b_\lambda}\cup\mathcal{P}_t^\partial|_{[a_\lambda,b_\lambda]}) \arrow[r] \arrow[d] & W_t|_b\cup X \arrow[d] \\ \coprod_{\lambda\in\Lambda}\mathcal{P}_t|_{[a_\lambda,b_\lambda]} \arrow[r] & W_t|_{[a,b]}. \end{tikzcd}$$

The left-hand map of the second square is a disjoint union of the maps discussed in property (iv) of Proposition 3.6, and so is κ-connected. As this square is a homotopy pushout, the right-hand map is also κ-connected.

The pair $(X, X|_b)$ is obtained from the manifold pair $(W|_{[a,b]}, W|_b)$ by cutting out embedded copies of $(D^\kappa, \partial D^\kappa)$. By transversality we see that this does not change relative homotopy groups in dimensions $*\leqslant d-\kappa-2$, which includes $*\leqslant\kappa$ by our assumption that $2\kappa\leqslant d-2$. In particular, suppose the pair $(W|_{[a,b]}, W|_b)$ is k-connected, with $k\leqslant\kappa$, then the pair $(X, X|_b)$ is k-connected too. As the first square above is a homotopy pushout square, the inclusion $W_t|_b\to W_t|_b\cup X$ also has this connectivity. Hence the composition $W_t|_b\to W_t|_b\cup X\to W_t|_{[a,b]}$ has the same connectivity as $W|_b\to W|_{[a,b]}$, up to a maximum of κ. This establishes that the tuple $(a,\varepsilon,\mathcal{K}_e^t(W,\ell_W))$ is an element of $X_p^{\kappa-1}$, and also that it lies in X_p^κ if (W,ℓ_W) lies in $D_{\theta,L}^\kappa(\mathbb{R}^N)$. When $t=1$, there is a little more to say.

Step 1. Suppose $a<b\in(a_i-\varepsilon_i, a_i+\varepsilon_i)$. Then $(W|_{[a,b]}, W|_b)$ is ∞-connected and so $(W_1|_{[a,b]}, W_1|_b)$ is κ-connected, by the discussion above.

Step 2. Suppose $a\in(a_{i-1}-\varepsilon_{i-1}, a_{i-1}+\varepsilon_{i-1})$ and $b\in(a_i-\varepsilon_i, a_i+\varepsilon_i)$. We now do the surgeries for $\Lambda_{i,0}\subset\Lambda$ first, giving a family of manifolds $\widetilde{W}_t$. We claim that the pair $(\widetilde{W}_1|_{[a,b]}, \widetilde{W}_1|_b)$ is κ-connected. Once this is established, doing the remaining surgeries to obtain W_1 does not change this property, as we have seen above.

Recall from Definition 3.2 (vi) that the pair $(W_0|_{[a,b]}, W_0|_b\cup D_{i,0}|_{[a,b]})$ is κ-connected, where

$$D_{i,0}=e_{i,0}(\Lambda_{i,0}\times\partial_-D^{\kappa+1}\times\{0\})\subset W=W_0.$$

If we write

$$\widetilde{D}_{i,0}=e_{i,0}(\Lambda_{i,0}\times\partial_-D^{\kappa+1}\times\{v\})\subset W=W_0$$

for some $v\in\mathbb{R}^{d-\kappa}\setminus B_4^{d-\kappa}(0)$, then the pair $(W_0|_{[a,b]}, W_0|_b\cup\widetilde{D}_{i,0}|_{[a,b]})$ is also κ-connected. Now the subset $\widetilde{D}_{i,0}\subset W$ is contained in $e_{i,0}(\Lambda_{i,0}\times\mathcal{P}_0^\partial)$, so we can regard $\widetilde{D}_{i,0}$ as a subset of $\widetilde{W}_t$ for all $t\in[0,1]$. The same transversality argument as before now shows that $(X, X|_b\cup\widetilde{D}_{i,0}|_{[a,b]})$ is also κ-connected, and the same gluing argument shows that $(\widetilde{W}_t|_{[a,b]}, \widetilde{W}_t|_b\cup\widetilde{D}_{i,0}|_{[a,b]})$ is κ-connected for all $t\in[0,1]$. When $t=1$, Proposition 3.6 (v) shows that the inclusion $\widetilde{D}_{i,0}|_{[a,b]}\to\widetilde{W}_1|_{[a,b]}$ is homotopic relative to $\widetilde{D}_{i,0}|_b$ to a map into $\widetilde{W}_1|_b$, and hence $(\widetilde{W}_1|_{[a,b]}, \widetilde{W}_1|_b)$ is κ-connected.

Step 3. For general $a<b\in\bigcup_{i=0}^{p}(a_i-\varepsilon_i, a_i+\varepsilon_i)$, we may choose regular values in each intermediate interval $(a_j-\varepsilon_j, a_j+\varepsilon_j)$. By the previous case, this expresses $W_1|_{[a,b]}$ as a composition of cobordisms which are all κ-connected relative to their outgoing boundaries, and hence the composition also has that property. □

3.3. Proof of Theorem 3.1

We begin with the composition

$$|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}| \longrightarrow |D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_{\bullet}| \longrightarrow |X^{\kappa-1}_{\bullet}|,$$

where the first map (induced by the augmentation) is a homotopy equivalence by Theorem 3.4 and the second is a homotopy equivalence by Proposition 2.20. We will define a homotopy

$$\mathscr{S}\colon [0,1]\times|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}| \longrightarrow |X^{\kappa-1}_{\bullet}|$$

starting from this map so that $\mathscr{S}(1,-)$ factors through $|X^{\kappa}_{\bullet}|\to|X^{\kappa-1}_{\bullet}|$, which is a continuous injection. Furthermore, there is an injection

$$|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet}| \hookrightarrow |D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}| \hookrightarrow |D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}|$$

using the empty collection of surgery data, and $\mathscr{S}$ will be constant on the image of this injection. The existence of a homotopy with these properties establishes Theorem 3.1 as follows: there is a diagram

$$\begin{tikzcd}
|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet}| \arrow[r] \arrow[d] & |X^{\kappa}_{\bullet}| \arrow[d] \\
|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}| \arrow[r, "\mathscr{S}(0,-)"] \arrow[ur, "\mathscr{S}(1,-)"] & |X^{\kappa-1}_{\bullet}|,
\end{tikzcd}$$

where the square commutes, the horizontal maps are weak homotopy equivalences, the top triangle commutes exactly and the bottom triangle commutes up to the homotopy $\mathscr{S}$. Taking homotopy groups we see that the vertical maps are also weak equivalences. Under the equivalence $B\mathcal{C}^{\kappa}_{\theta,L}(\mathbb{R}^N)\simeq|X^{\kappa}_{\bullet}|$, and similarly for $\kappa-1$, we obtain Theorem 3.1.

To define the surgery map $\mathscr{S}$ we will give a collection of maps

$$\mathscr{S}_{p,q}\colon [0,1]\times D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}\times\Delta^q \longrightarrow X^{\kappa-1}_p$$

compatible on their faces. The construction of the last section gives a one-parameter family

$$\begin{aligned}\mathcal{K}^r\colon D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,0} &\longrightarrow X^{\kappa-1}_p,\\ (a,\varepsilon,W,e) &\longmapsto (a,\varepsilon,\mathcal{K}^r_e(W)),\end{aligned}$$

for $r\in[0,1]$, such that $\mathcal{K}^1$ lands in X^{κ}_p. When $q=0$, we set

$$\mathscr{S}_{p,0}(r,(a,\varepsilon,W,e)) = (a,\varepsilon,\mathcal{K}^r_e(W))\in X^{\kappa-1}_p.$$

More generally, for $q\geqslant 0$ we have $e=\{e_{i,j}\}$, and for each j we get an element

$$(a,\varepsilon,W,e_{*,j})\in D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,0}.$$

We then set

$$\mathscr{S}_{p,q}(r,(a,\varepsilon,W,e),s) = (a,\varepsilon,\mathcal{K}^{\bar{s}_q r}_{e_{*,q}}\circ\ldots\circ\mathcal{K}^{\bar{s}_0 r}_{e_{*,0}}(W)),$$

where $\bar{s}_j=s_j/\max_k s_k$. Note that some $\bar{s}_j$ is always equal to 1, so when $r=1$, some $\mathcal{K}^1_{e_{*,j}}$ is applied to W making each morphism κ-connected relative to its outgoing boundary. The remaining $\mathcal{K}^{\bar{s}_k}_{e_{*,k}}$ do not change this property, by Lemma 3.7, and so the map $\mathscr{S}_{p,q}(1,-)$ factors through the subspace X^{κ}_p.

The resulting map from $\coprod_{q\geqslant 0}([0,1]\times D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}\times\Delta^q)$ factors through a map

$$\mathscr{S}_p\colon [0,1]\times|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,\bullet}|\longrightarrow X^{\kappa-1}_p,$$

which together form a map of semi-simplicial spaces with geometric realisation

$$\mathscr{S}\colon [0,1]\times|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}|\longrightarrow |X^{\kappa-1}_{\bullet}|.$$

On the image of $|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet}|$, the homotopy is constant as there is no surgery data. At $r=1$ it factors through $|X^{\kappa}_{\bullet}|$. This finishes the proof of Theorem 3.1.

4. Surgery on objects below the middle dimension

In this section we wish to study the filtration

$$\mathcal{C}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)\subset\ldots\subset\mathcal{C}^{\kappa,1}_{\theta,L}(\mathbb{R}^N)\subset\mathcal{C}^{\kappa,0}_{\theta,L}(\mathbb{R}^N)\subset\mathcal{C}^{\kappa,-1}_{\theta,L}(\mathbb{R}^N)=\mathcal{C}^{\kappa}_{\theta,L}(\mathbb{R}^N)$$

and in particular establish the following theorem. The reader mainly interested in Theorems 1.1 and 1.2 can take $d=2n$, $\kappa=n-1$, $l\leqslant n-2$, $\theta=\theta^n\colon B\mathrm{O}(2n)\langle n\rangle\to B\mathrm{O}(2n)$, $L\cong D^{2n-1}$, and $N=\infty$ (but the proof does not simplify much in this special case).

THEOREM 4.1. *Suppose that the following conditions are satisfied*:
(i) $2(l+1)<d$;
(ii) $l\leqslant\kappa$;
(iii) $l\leqslant d-\kappa-2$;
(iv) $l+2+d<N+1$;
(v) *L admits a handle decomposition only using handles of index at most $d-l-2$*;
(vi) *the map $\ell_L\colon L\to B$ is $(l+1)$-connected.*
Then the map

$$B\mathcal{C}_{\theta,L}^{\kappa,l}(\mathbb{R}^N)\longrightarrow B\mathcal{C}_{\theta,L}^{\kappa,l-1}(\mathbb{R}^N)$$

is a weak homotopy equivalence.

We remark that under the assumptions of Theorem 3.1, (iii) and (v) in the theorem above are implied by (ii).

The proof will be similar in spirit to that of the last section, in so far as we will define a contractible space of surgery data and describe a surgery move which compresses $B\mathcal{C}_{\theta,L}^{\kappa,l-1}(\mathbb{R}^N)$ into $B\mathcal{C}_{\theta,L}^{\kappa,l}(\mathbb{R}^N)$. In the same way that the surgery move of the last section was a refinement of that of [GMTW], the surgery move we use in this and the next section is a refinement of that of [GRW1]. Let us first give an informal account of this move, and for simplicity suppose that $N=\infty$, that we have no tangential structure (i.e. we consider $\theta=\mathrm{Id}\colon B\mathrm{O}(d)\to B\mathrm{O}(d)$), that $L=\varnothing$, and that $d>2$, $l=0$ and $\kappa=0$. We first apply the equivalence (2.1) to reduce the problem to studying the map

$$BD^{0,0}\longrightarrow BD^{0,-1}$$

of classifying spaces of posets. Let

$$\sigma=(t_0,t_1;a_0,a_1;\varepsilon_0,\varepsilon_1;W)\in BD^{0,-1}$$

be a point on a 1-simplex (for example), and let us suppose that $W|_{a_1}$ is already connected (so $\pi_0(W|_{a_1})$ injects into $\pi_0(B\mathrm{O}(d))$). We will describe a way of producing a path from the image of this point in $|X_\bullet^{0,-1}|$ into the subset $|X_\bullet^{0,0}|$.

If $W|_{a_0}$ is already connected, then the point σ already lies in $|X_\bullet^{0,0}|$ and there is nothing to prove. Otherwise, let us choose disjoint embeddings

$$\{f_\alpha\colon S^0\hookrightarrow W|_{a_0}\}_{\alpha\in\Lambda}$$

such that if we perform 0-surgery along (thickenings of) all of these embeddings, the resulting $(d-1)$-manifold is connected. As $\kappa=0$, the cobordism $W|_{[a_0,a_1]}$ is path connected relative to its top, and so we can extend the f_α to smooth maps

$$\hat{f}_\alpha\colon(a_0-\varepsilon_0,a_1+\varepsilon_1)\times S^0\longrightarrow W$$

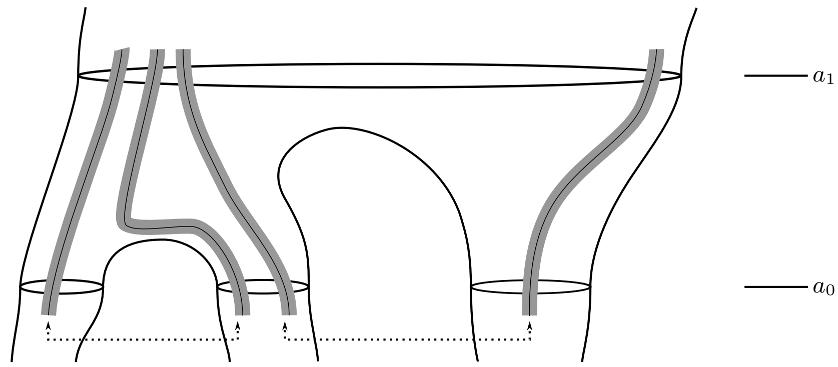

Figure 7. An element of $N_1D^{0,-1}$ which is not in $N_1D^{0,0}$, together with two pieces of surgery data for the level a_0.

such that the standard height function (i.e. the projection to $(a_0-\varepsilon_0, a_1+\varepsilon_1)$) and $x_1\circ\hat{f}_\alpha$ agree inside $(x_1\circ\hat{f}_\alpha)^{-1}\big(\bigcup_{i=0}^{p}(a_i-\varepsilon_i, a_i+\varepsilon_i)\big)$. As we have supposed that $d>2$, we may assume that these $\hat{f}_\alpha$ are mutually disjoint embeddings. By taking a tubular neighbourhood, we extend the $\hat{f}_\alpha$ to embeddings

$$\hat{e}_\alpha\colon(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}^{d-1}\times S^0\hookrightarrow W$$

which are still mutually disjoint, and extend these further to disjoint embeddings

$$e_\alpha\colon(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}^{d-1}\times D^1\hookrightarrow\mathbb{R}\times\mathbb{R}^\infty$$

such that $e_\alpha^{-1}(W)=(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}^{d-1}\times S^0$. It is clear that we can arrange the same relationship between the standard height function on $(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}^{d-1}\times D^1$ and the function $x_1\circ e_\alpha$ as we have over $(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}^{d-1}\times S^0$. In Figure 7 we have shown a typical example (the picture has $d=2$, but the reader should imagine a slightly larger d): the original manifold does not have path-connected level set at the level a_0, but we have chosen two e_α's and depicted the images $e_\alpha((a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}^{d-1}\times S^0)$ as the shaded parts.

The surgery move is then given by gluing in the one-parameter family shown in Figure 8 along each e_α. The family depicted there (with $d=2$, but again the reader should imagine a larger d) starts at the manifold $(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}^{d-1}\times S^0$, replaces it with the trace of a 0-surgery on $\mathbb{R}^{d-1}\times S^0$, sliding the critical value down from $a_1+\varepsilon_1$ to $a_0-\frac{1}{2}\varepsilon_0$. This does not define a path in $BD^{0,-1}$, as $(a_1-\varepsilon_1, a_1+\varepsilon_1)$ will contain a critical value at some points during the path. However, it does define a path in $|X_\bullet^{0,-1}|$. Furthermore, if we let $\overline{W}$ be the manifold obtained at the end of the path, then $\overline{W}|_{a_0}$ is obtained from $W|_{a_0}$ by doing 0-surgery along the data $\{\hat{e}_\alpha|_{\{a_0\}\times\mathbb{R}^{d-1}\times S^0}\}_{\alpha\in\Lambda}$ and so is connected. Also, the manifold $\overline{W}|_{a_1}$ is obtained from $W|_{a_1}$ by doing 0-surgery along the

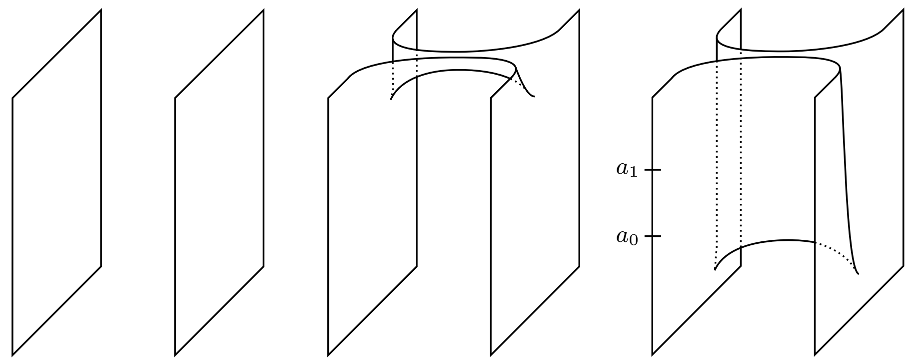

Figure 8. The surgery move for surgery on objects below the middle dimension.

data $\{\hat{e}_\alpha|_{\{a_1\}\times\mathbb{R}^{d-1}\times S^0}\}_{\alpha\in\Lambda}$, and as it was connected to start with (and $d>2$), it remains connected. Hence $(t_0, t_1; a_0, a_1; \varepsilon_0, \varepsilon_1; \overline{W})\in|X_\bullet^{0,0}|$, as required.

This surgery move generalises well to $l>0$, to finite (but sufficiently large) N, and to non-empty L, but to make it work with general tangential structures θ we must equip the surgery data $\{e_\alpha\}_{\alpha\in\Lambda}$ with extra data describing how to induce a θ-structure on the surgered manifold. We will first give a definition of θ-surgery, then describe the standard family, and finally go on to describe the semi-simplicial space of surgery data analogous to that of §3.1.

4.1. θ-surgery

Consider a $(d-1)$-dimensional θ-manifold (M, ℓ_M) and an embedding

$$e\colon \mathbb{R}^{d-l-1}\times S^l \hookrightarrow M,$$

and let C be the d-dimensional cobordism obtained as the trace of the surgery along e. Thus $\partial_{\mathrm{in}}C=M$ and $\partial_{\mathrm{out}}C=\overline{M}$ is the result of the surgery. The data of a *θ-surgery* on M is an embedding e as above along with a θ-structure ℓ on C which agrees with ℓ_M on M. This induces a θ-structure on $\overline{M}$. We will typically give the data of a θ-surgery extending an embedding e by giving an extension of the θ-structure $\ell_M\circ De$ on $\mathbb{R}^{d-l-1}\times S^l$ to $\mathbb{R}^{d-l-1}\times D^{l+1}$.

Given just an embedding $e_0\colon S^l\to M$, the simultaneous choice of an extension to $\mathbb{R}^{d-l-1}\times S^l\hookrightarrow M$ and a θ-structure on the trace of the resulting surgery has the following alternative description. The embedding e_0 induces an embedding

$$e_1 = I\times e_0\colon I\times S^l \longrightarrow I\times M,$$

where $I=[0,1]$. We shall write A for the annulus $A=D^{l+1}\setminus\operatorname{int}\big(\frac{1}{2}D^{l+1}\big)$ which we identify with $I\times S^l$ using a diffeomorphism which identifies $\partial D^{l+1}\subset A$ with $\{0\}\times S^l$ (for example $x\mapsto(2(1-|x|),x/|x|)$). We have the following diagram of bundle maps and bundle injections:

$$\begin{tikzcd}
TD^{l+1}|_A \arrow[r] \arrow[d] & T(I\times S^l) \arrow[r, hook, "De_1"] \arrow[d] & T(I\times M) \arrow[r, "\ell_M"] \arrow[d] & \theta^*\gamma \arrow[d] \\
A \arrow[r, "\cong"] & I\times S^l \arrow[r, "e_1"] & I\times M \arrow[r] & B.
\end{tikzcd}$$

The composition is a bundle injection $TD^{l+1}|_A\to\theta^*\gamma$ covering a map $A\to B$, and we claim the necessary data is an extension to a bundle injection $TD^{l+1}\to\theta^*\gamma$. Indeed, such a bundle injection determines a $(d-l-1)$-dimensional complement $V\to D^{l+1}$ and a bundle map $TD^{l+1}\oplus V\to\theta^*\gamma$, and $V|_A$ is canonically identified with the normal bundle of e_1. The disc bundle $D(V)$ is diffeomorphic to $D^{l+1}\times D^{d-l-1}$ and hence the pushout of the diagram

$$I\times M \longleftarrow D(V|_A)\longrightarrow D(V)$$

is the trace of a surgery on M along a thickening of e_0, but in this description it comes with a canonical θ-structure.

An obstruction-theoretic argument shows that any extension of the underlying map $A\to B$ to a map $f\colon D^{l+1}\to B$ may be lifted to an extension of the bundle injection if $2(l+1)\leqslant d$, and that such a lift is unique up to homotopy when $2(l+1)<d$. Indeed, the problem is equivalent to injecting TD^{l+1} into $f^*(\theta^*\gamma)$ with given injection over ∂D^{l+1}, and since both bundles are trivial because D^{l+1} is contractible, this is in turn equivalent to extending a map $\partial D^{l+1}\to O(d)/O(d-l-1)$ to a map from D^{l+1}; since the Stiefel manifold $O(d)/O(d-l-1)$ is $(d-l-2)$-connected, this is possible for $l\leqslant d-l-2$ with homotopically unique extension for $l<d-l-2$. Thus any map

$$M\cup_{e_0}D^{l+1}\longrightarrow B$$

extending the map $M\to B$ underlying ℓ_M occurs as the trace of some θ-surgery on M, restricted to the core of the attached $(l+1)$-handle, provided $l+1\leqslant\frac{1}{2}d$.

4.2. The standard family

Let us construct the one-parameter family of manifolds depicted in Figure 8. Choose a smooth function $\varrho\colon\mathbb{R}\to\mathbb{R}$ which is the identity on $\big(-\infty,\frac{1}{2}\big)$, has nowhere negative derivative, and has $\varrho(t)=1$ for all $t\geqslant 1$. We define

$$K=\{(x,y)\in\mathbb{R}^{d-l}\times\mathbb{R}^{l+1}: |y|^2=\varrho(|x|^2-1)\},$$

a smooth d-dimensional submanifold, contained in $\mathbb{R}^{d-l}\times D^{l+1}$, which outside of the set $B^{d-l}_{\sqrt{2}}(0)\times D^{l+1}$ is identically equal to $\mathbb{R}^{d-l}\times S^l$. Let

$$h=x_1\colon K\longrightarrow \mathbb{R}$$

denote the first of the x coordinates, which is the height function we will consider on K. This function has precisely two critical points, both non-degenerate: $(-1,0,...,0)$ of index $l+1$ and $(1,0,...,0)$ of index $d-l-1$.

We now define a one-parameter family of d-dimensional submanifolds $\mathcal{P}_t$ inside $(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1}$ in the following way. Pick a smooth one-parameter family of embeddings

$$\lambda_s\colon (-6,-2)\longrightarrow(-6,0),$$

such that $\lambda_0=\mathrm{Id}$, that $\lambda_s|_{(-6,-5)}=\mathrm{Id}$ for all s, and that $\lambda_1(-4)=-1$. Then we get an embedding $\lambda_t\times\mathrm{Id}_{\mathbb{R}^d}\colon(-6,-2)\times\mathbb{R}^d\to(-6,0)\times\mathbb{R}^d$ and define

$$\mathcal{P}_t=(\lambda_t\times\mathrm{Id}_{\mathbb{R}^d})^{-1}(K)\in\Psi_d((-6,-2)\times\mathbb{R}^{d-l-1}\times\mathbb{R}^{l+1}).$$

It is easy to verify that $\mathcal{P}_t$ agrees with $(-6,-2)\times\mathbb{R}^{d-l-1}\times S^l$ outside of the region $(-5,-2)\times B^{d-l-1}_{\sqrt{2}}(0)\times D^{l+1}$, independently of t: if (s,x,y) is outside of this region, we either have $s\leqslant -5$, so $|\lambda_t(s)|^2=|s|^2\geqslant 25$, or $|x|^2\geqslant 2$, so in either case $|\lambda_t(s)|^2+|x|^2-1\geqslant 1$ and hence

$$\varrho(|\lambda_t(s)|^2+|x|^2-1)=1.$$

Therefore $(\lambda_t(s),x,y)\in K$ if and only if $|y|=1$.

We shall also need a tangentially structured version of this construction, given a structure $\ell\colon TK|_{(-6,0)}\to\theta^*\gamma$. For this purpose, let $\omega\colon[0,\infty)\to[0,1]$ be a smooth function such that $\omega(r)=0$ for $r\geqslant 2$ and $\omega(r)=1$ for $r\leqslant\sqrt{2}$. We define a one-parameter family of embeddings by

$$\begin{aligned}\psi_t\colon(-6,-2)\times\mathbb{R}^{d-l-1}\times\mathbb{R}^{l+1}&\longrightarrow(-6,0)\times\mathbb{R}^{d-l-1}\times\mathbb{R}^{l+1},\\ (s,x,y)&\longmapsto(\lambda_{t\omega(|x|)}(s),x,y).\end{aligned}$$

It is easy to see that we also have $\psi_t^{-1}(K)=(\lambda_t\times\mathrm{Id}_{\mathbb{R}^d})^{-1}(K)=\mathcal{P}_t$ for all t: the two functions $\lambda_t\times\mathrm{Id}_{\mathbb{R}^d}$ and ψ_t agree on those (s,x,y) with $|x|\leqslant 2$ since $\omega(|x|)=1$ there, and outside this region both inverse images agree with $(-6,-2)\times\mathbb{R}^{d-l-1}\times S^l$. We define a θ-structure on $\mathcal{P}_t$ by pullback along ψ_t, which as ψ_t is independent of t outside of $(-6,-2)\times B^{d-l-1}_{\sqrt{2}}(0)\times D^{l+1}$ gives a constant family of θ-structures outside of this region. This gives a family $\mathcal{P}_t(\ell)\in\Psi_\theta((-6,-2)\times\mathbb{R}^{d-l-1}\times\mathbb{R}^{l+1})$, and we record some important properties in the following proposition. We will omit ℓ from the notation when it is unimportant.

PROPOSITION 4.2. *The elements*

$$\mathcal{P}_t(\ell)\in\Psi_\theta((-6,-2)\times\mathbb{R}^{d-l-1}\times\mathbb{R}^{l+1})$$

are θ*-submanifolds of* $(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1}$ *satisfying the following properties*:

(i) $\mathcal{P}_0(\ell)=K|_{(-6,-2)}=(-6,-2)\times\mathbb{R}^{d-l-1}\times S^l$ *as* θ*-manifolds.*

(ii) *For all* t, $\mathcal{P}_t(\ell)$ *agrees with* $K|_{(-6,-2)}$ *as* θ*-manifolds outside of the region* $(-5,-2)\times B_2^{d-l-1}(0)\times D^{l+1}$.

(iii) *For all* t *and each pair of regular values* $-6<a<b<-2$ *of the height function* $h:\mathcal{P}_t\to\mathbb{R}$, *the pair*

$$(\mathcal{P}_t|_{[a,b]},\mathcal{P}_t|_b)$$

is $(d-l-2)$*-connected.*

(iv) *For each regular value* a *of* $h:\mathcal{P}_t\to(-6,-2)$, *the manifold* $\mathcal{P}_t|_a$ *is either isomorphic to* $\mathcal{P}_0|_a$ *or is obtained from it by* l*-surgery.*

(v) *The only critical value of* $h:\mathcal{P}_1\to(-6,-2)$ *is* -4, *and for* $a\in(-4,-2)$, $\mathcal{P}_1|_a$ *is obtained by* l*-surgery from* $\mathcal{P}_0|_a=\mathbb{R}^{d-l-1}\times S^l$ *along the standard embedding.*

In (iv) *and* (v), *the* θ*-structure on the surgered manifold is determined* (*up to homotopy, cf.* §4.1) *by the* θ*-structure on* $K|_{(-6,0)}$.

The precise meaning of the word isomorphic in (iv) above is the following: By (ii) we know that the manifolds are *equal* outside $(-5,-2)\times B_2^{d-l-1}(0)\times D^{l+1}$. Being isomorphic means that the identity extends to a diffeomorphism which preserves θ-structures up to a homotopy of bundle maps which is constant outside $(-5,-2)\times B_2^{d-l-1}(0)\times D^{l+1}$.

Proof. (i) and (ii) follow easily from the properties of λ_t and ψ_t, and the fact that K agrees with $\mathbb{R}^{d-l}\times S^l$ outside $B_{\sqrt{2}}^{d-l}\times\mathbb{R}^{l+1}$. It follows from the properties of ω that the θ-structures agree outside $B_2^{d-l}\times\mathbb{R}^{l+1}$. For (iii), the height function $\mathcal{P}_t\to(-6,-2)$ has at most one critical point, which is non-degenerate of index $l+1$. If the critical value is in (a,b), then the pair is $(d-l-2)$-connected, otherwise $\mathcal{P}|_{[a,b]}$ deformation retracts to $\mathcal{P}|_b$. The fact that the height function has at most one critical point, of index $l+1$, also implies (iv) by definition of surgery (and θ-surgery, cf. §4.1). Finally, the property that $\lambda_1(-4)=-1$ and $\lambda_1(-5)=-5$ implies that $h:\mathcal{P}_1\to(-6,-2)$ does have a critical point of index $l+1$, with critical value -4, which proves (v). □

4.3. Surgery data

We can now describe the semi-simplicial space of surgery data out of which we will construct a "perform surgery" map. In the next section we will describe how to construct this map.

Before doing so, we choose once and for all, smoothly in the data $(a_i, \varepsilon_i, a_p, \varepsilon_p)$, increasing diffeomorphisms

$$\varphi = \varphi(a_i, \varepsilon_i, a_p, \varepsilon_p): (-6, -2) \cong (a_i - \varepsilon_i, a_p + \varepsilon_p) \tag{4.1}$$

sending -4 to $a_i - \frac{1}{2}\varepsilon_i$ and -5 to $a_i - \frac{3}{4}\varepsilon_i$.

Definition 4.3. Let $x = (a, \varepsilon, (W, \ell_W)) \in D^{\kappa, l-1}_{\theta, L}(\mathbb{R}^N)_p$, and write $M_i = W|_{a_i}$. Define the set $Y_q(x)$ to consist of tuples $(\Lambda, \delta, e, \ell)$, where $\Lambda \subset \Omega$ is a finite subset of the fixed infinite set Ω, $\delta: \Lambda \to [p] \times [q]$ is a function,

$$e: \Lambda \times (-6, -2) \times \mathbb{R}^{d-l-1} \times D^{l+1} \hookrightarrow \mathbb{R} \times (0, 1) \times (-1, 1)^{N-1}$$

is an embedding, and $\ell: T(\Lambda \times K|_{(-6,0)}) \to \theta^* \gamma$ is a bundle map, satisfying the conditions below. We shall write $\Lambda_{i,j} = \delta^{-1}(i, j)$ and

$$e_{i,j}: \Lambda_{i,j} \times (a_i - \varepsilon_i, a_p + \varepsilon_p) \times \mathbb{R}^{d-l-1} \times D^{l+1} \longrightarrow \mathbb{R} \times (0, 1) \times (-1, 1)^{N-1}$$

for the embedding obtained by restricting e and reparameterising using (4.1).

(i) $e^{-1}(W) = \Lambda \times (-6, -2) \times \mathbb{R}^{d-l-1} \times S^l$. We let

$$\partial e: \Lambda \times (-6, -2) \times \mathbb{R}^{d-l-1} \times S^l \hookrightarrow W$$

denote the embedding restricted to the boundary.

(ii) For $t \in \bigcup_{k=0}^{p} (a_k - \varepsilon_k, a_k + \varepsilon_k)$, we have $(x_1 \circ e_{i,j})^{-1}(t) = \Lambda_{i,j} \times \{t\} \times \mathbb{R}^{d-l-1} \times D^{l+1}$.

(iii) The composition $\ell_W \circ D\partial e: T(\Lambda \times K|_{(-6,-2)}) \to \theta^* \gamma$ agrees with the restriction of ℓ.

If we let $\ell_{i,j}$ denote the restriction of ℓ to $T(\Lambda_{i,j} \times K|_{(-6,0)})$, the data $(e_{i,j}, \ell_{i,j})$ is enough to perform θ-surgery on M_i (as $K|_{(-6,0)}$ contains the trace of an l-surgery on the $(d-1)$-manifold $K|_{-2}$), and we further insist that

(iv) For each $j = 0, ..., q$ and $i = 0, ..., p$, the resulting θ_{d-1}-manifold $\overline{M}_i$ has the property that $\pi_k(\overline{M}_i) \to \pi_k(B)$ is injective for $k \leqslant l$.

For each x, $Y_\bullet(x)$ is a semi-simplicial set in the same way as in Definition 3.2.

Note that the set $Y_q(x)$ consists of those $(q+1)$-tuples of elements of $Y_0(x)$ which are disjoint. A typical example of a surgery datum is (partially) depicted in Figure 7. The figure has $p = 1$, $q = 0$, $d = 2$ and $\kappa = l = 0$ (although this case does not satisfy the requirement $2(l+1) < d$, so the reader should imagine a larger value of d). Only the image $e(\Lambda \times (-6, -2) \times \mathbb{R}^{d-l-1} \times S^l) \subset W$ is shown.

Definition 4.4. We define a bi-semi-simplicial space $D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$ (augmented in the second semi-simplicial direction) as a set by

$$D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,q} = \{(x,y) : x \in D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_p \text{ and } y \in Y_q(x)\},$$

and topologise it as a subspace of

$$D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_p \times \left(\coprod_{\Lambda\subset\Omega} (C^\infty(\Lambda\times V, \mathbb{R}^{N+1}) \times \mathrm{Bun}(T(\Lambda\times K|_{(-6,0)}), \theta^*\gamma))\right)^{(p+1)(q+1)},$$

where V denotes the manifold $(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1}$. Explicitly, the face map d_k in the q direction forgets the surgery data $(e_{i,j},\ell_{i,j})$ with $j=k$, and the face map d_k in the p direction forgets both the surgery data $(e_{i,j},\ell_{i,j})$ with $i=k$ and the kth regular value.

The main result about this bi-semi-simplicial space of manifolds equipped with surgery data is the following, whose proof we defer until §6.

Theorem 4.5. *Under the assumptions of Theorem* 4.1, *the maps*

$$|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}| \longrightarrow |D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}| \longrightarrow |D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_{\bullet}|$$

are weak homotopy equivalences, where the first map is the inclusion of 0*-simplices and the second is the augmentation, in the second simplicial direction.*

In fact, we shall prove this theorem assuming the conditions of Theorem 4.1 except (iii). That condition will be used in the proof of Lemma 4.6.

4.4. Proof of Theorem 4.1

We now go on to prove Theorem 4.1, so suppose that the conditions in the statement of that theorem are satisfied: $2(l+1)<d$, $l\leqslant\kappa$, $l\leqslant d-\kappa-2$, $l+2+d<N$, L admits a handle decomposition using only handles of index at most $d-l-2$, and the map $\ell_L\colon L\to B$ is $(l+1)$-connected.

Let $(a,\varepsilon,(W,\ell_W),e,\ell)\in D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,0}$. For each $i=0,...,p$, we have an embedding $e_i=e_{i,0}$ and a bundle map $\ell_i=\ell_{i,0}$, from which we shall construct a one-parameter family of elements $\mathcal{K}^t_{e_i,\ell_i}(W,\ell_W)\in\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N)$, $t\in[0,1]$, as follows. Changing the first coordinate of the manifolds $\mathcal{P}_t(\ell_i)$ by composing with the reparametrisation functions of (4.1), we get a family of manifolds

$$\overline{\mathcal{P}}_t(\ell_i)\in\Psi_\theta((a_i-\varepsilon_i,a_p+\varepsilon_p)\times\mathbb{R}^{d-l-1}\times\mathbb{R}^{l+1})$$

having all the properties of Proposition 4.2, where property (v) now holds for all values in the interval $\left(a_i-\frac{1}{2}\varepsilon_i, a_p+\varepsilon_p\right)$, as this is the image of the interval $(-4,-2)$ under the reparametrisation (4.1). Then for $t\in[0,1]$, let

$$\mathcal{K}^t_{e_i,\ell_i}(W,\ell_W)\in\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N)$$

be equal to $W|_{(a_0-\varepsilon_0,a_p+\varepsilon_p)}$ outside the image of e_i, and on

$$e_i(\Lambda_i\times(a_i-\varepsilon_i,a_p+\varepsilon_p)\times\mathbb{R}^{d-l-1}\times D^{l+1})$$

be given by $e_i(\Lambda_i\times\overline{\mathcal{P}}_t(\ell_i))$. This gives a θ-manifold, because $\Lambda_i\times\overline{\mathcal{P}}_t(\ell_i)$ and $\Lambda_i\times\overline{\mathcal{P}}_0(\ell_i)$ agree as θ-manifolds outside of $\left(a_i-\frac{3}{4}\varepsilon_i,a_p+\varepsilon_p\right)\times B^{d-l-1}_2(0)\times D^{l+1}$.

As the embeddings e_i are all disjoint, this procedure can be iterated, and for a tuple $t=(t_0,...,t_p)\in[0,1]^{p+1}$ we let

$$\mathcal{K}^t_{e,\ell}(W,\ell_W)=\mathcal{K}^{t_p}_{e_p,\ell_p}\circ...\circ\mathcal{K}^{t_0}_{e_0,\ell_0}(W,\ell_W)\in\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N).$$

LEMMA 4.6. *Firstly, the tuple* $\left(a,\frac{1}{2}\varepsilon,\mathcal{K}^t_{e,\ell}(W,\ell_W)\right)$ *is an element of* $X^{\kappa,l-1}_p$. *Secondly, if* $t_i=1$*—so the surgery for the regular value* a_i *is fully done—then for any regular value* b *of* $x_1\colon\mathcal{K}^t_{e,\ell}(W,\ell_W)\to\mathbb{R}$ *in the interval* $\left(a_i-\frac{1}{2}\varepsilon_i,a_i+\frac{1}{2}\varepsilon_i\right)$ *we have that*

$$\pi_j(\mathcal{K}^t_{e,\ell}(W,\ell_W)|_b)\longrightarrow\pi_j(B)$$

is injective for $j\leqslant l$.

Proof. For the first part we must verify the conditions of Definition 2.18. Conditions (i)–(iii) are immediate from the properties of (a,ε) that we start with and the disjointness of the surgery data from $\mathbb{R}\times L$.

For condition (iv) we proceed exactly as in the proof of Lemma 3.7. For regular values $a<b\in\bigcup_{i=0}^p(a_i-\varepsilon_i,a_i+\varepsilon_i)$ of the height function $x_1\colon\mathcal{K}^t_{e,\ell}(W,\ell_W)\to\mathbb{R}$, it follows from Definition 4.3 (ii) that the manifold $\mathcal{K}^t_{e,\ell}(W,\ell_W)|_{[a,b]}$ is obtained from $W|_{[a,b]}$ by cutting out embedded images of cobordisms $\mathcal{P}_0|_{[a_\lambda,b_\lambda]}$ and gluing in $\mathcal{P}_t|_{[a_\lambda,b_\lambda]}$, where $a_\lambda<b_\lambda$ are regular values of the height function on $\mathcal{P}_0$ and $\mathcal{P}_t$. By property (iii) of the standard family, the pair $(\mathcal{P}_t|_{[a_\lambda,b_\lambda]},\mathcal{P}_t|_{b_\lambda})$, and hence the homotopy equivalent pair

$$(\mathcal{P}_t|_{[a_\lambda,b_\lambda]},\mathcal{P}_t|_{b_\lambda}\cup([a_\lambda,b_\lambda]\times(\mathbb{R}^{d-l-1}\setminus B^{d-l-1}_2(0))\times S^l)),$$

is $(d-l-2)$-connected, and so as we have supposed that $l\leqslant d-\kappa-2$ it is in particular κ-connected. We now continue as in the proof of Lemma 3.7.

For condition (v), let $b\in\left(a_i-\frac{1}{2}\varepsilon_i, a_i+\frac{1}{2}\varepsilon_i\right)$ be a regular value of the height function on $\mathcal{K}^t_{e,\ell}(W,\ell_W)$, and define θ_{d-1}-manifolds

$$\overline{M}=\mathcal{K}^t_{e,\ell}(W,\ell_W)|_b \quad \text{and} \quad M=W|_b.$$

By Proposition 4.2 (iv), the θ_{d-1}-manifold $\overline{M}$ is obtained from M by performing θ-l-surgeries. Let $C\colon M\rightsquigarrow\overline{M}$ be the θ-cobordism given by the trace of these surgeries. We have the commutative diagram

$$\begin{tikzcd} \pi_j(M) \arrow[r, "i"] \arrow[dr] & \pi_j(C) \arrow[d] & \pi_j(\overline{M}) \arrow[l, "\bar\imath"'] \arrow[dl] \\ & \pi_j(B) & \end{tikzcd} \tag{4.2}$$

and C is obtained by attaching $(l+1)$-cells to M or by attaching $(d-l-1)$-cells to $\overline{M}$. Hence i is surjective for $j\leqslant l$ and $\bar\imath$ is bijective for $j\leqslant d-l-3$. The condition $2(l+1)\leqslant d$ implies that $l-1\leqslant d-l-3$, and the left-hand diagonal map is injective for $j\leqslant l-1$, so the right-hand diagonal map is injective for $j\leqslant l-1$ too.

We now prove the second part, so suppose that $t_i=1$. We construct the manifold $\mathcal{K}^t_{e,\ell}(W,\ell_W)$ by first taking $\mathcal{K}^1_{e_i,\ell_i}(W,\ell_W)$ and then performing the remaining surgeries to it. Let $\widetilde{M}=\mathcal{K}^1_{e_i,\ell_i}(W,\ell_W)|_b$, so that $\overline{M}$ is obtained from $\widetilde{M}$ by l-surgery.

We first show that $\pi_j(\widetilde{M})\to\pi_j(B)$ is injective for $j\leqslant l$. By property (iv) of the complex of surgery data, (e_i,ℓ_i) is enough surgery data on $M=W|_b$ to make the map on π_l injective after performing it. By property (v) of the standard family, as $b>a_i-\frac{1}{2}\varepsilon_i$ the manifold $\widetilde{M}$ has all of this surgery done, and so $\pi_j(\widetilde{M})\to\pi_j(B)$ is injective for $j\leqslant l$.

By the previous argument, with M replaced by $\widetilde{M}$ in (4.2), the remaining surgeries do not change this injectivity property. □

In the composition

$$|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}|\longrightarrow|D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_\bullet|\longrightarrow|X^{\kappa,l-1}_\bullet|$$

both maps are homotopy equivalences by Theorem 4.5 and Proposition 2.20 respectively. There is also an injection

$$|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_\bullet|\longrightarrow|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|\longrightarrow|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}|$$

using the empty collection of surgery data, and the second map is a weak homotopy equivalence by Theorem 4.5.

Similarly to the last section, we now construct maps which implement the surgery move. However, because a $(p,0)$-simplex now contains a $(p+1)$-tuple of surgery data, rather than a $(p+2)$-tuple as in the last section, the formal details are slightly different. We define a map

$$\begin{aligned}\mathscr{S}_p\colon [0,1]^{p+1}\times D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,0} &\longrightarrow X^{\kappa,l-1}_p,\\ (t,(a,\varepsilon,(W,\ell_W),e,\ell)) &\longmapsto \big(a,\tfrac{1}{2}\varepsilon,\mathcal{K}^t_{e,\ell}(W,\ell_W)\big),\end{aligned}$$

which has the desired range by the first part of Lemma 4.6, and furthermore sends $(1,...,1)\times D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,0}$ into $X^{\kappa,l}_p$. On the boundary of the cube, this map has further distinguished properties: one is given by the second part of Lemma 4.6. The second is that, by Proposition 4.2 (i), we have an equality $\mathcal{K}^0_{e_i,\ell_i}(W')=W'$ of θ-submanifolds of $(a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N$. Thus we obtain the formula

$$d_i\mathscr{S}_p(d^it,x)=\mathscr{S}_{p-1}(t,d_ix), \tag{4.3}$$

where $d^i\colon[0,1]^p\to[0,1]^{p+1}$ adds a zero in the ith position, and the d_i are the face maps of the semi-simplicial spaces $D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}$ and $X^{\kappa,l-1}_\bullet$.

We wish to assemble the maps $\mathscr{S}_p$ to a homotopy $\mathscr{S}\colon[0,1]\times|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|\to|X^{\kappa,l-1}_\bullet|$. Hence we define $\lambda,\psi\colon\Delta^p\to[0,1]^{p+1}$ by the formulæ

$$\lambda_i(t)=\min\{1,2\bar{t}_i\}\quad\text{and}\quad\psi_i(t)=\max\{0,2\bar{t}_i-1\}$$

for $0\leqslant i\leqslant p$, where again $\bar{t}_i=t_i/\max_j t_j$, and a map $H\colon[0,1]\times\Delta^p\to[0,1]^{p+1}\times\Delta^p$ by

$$H(s,t)=\left(s\lambda(t),\frac{\psi(t)}{\sum_{j=0}^p\psi_j(t)}\right).$$

These may be used to form the composition

$$F_p\colon[0,1]\times D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,0}\times\Delta^p\xrightarrow{H} D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,0}\times[0,1]^{p+1}\times\Delta^p\xrightarrow{\mathscr{S}_p\times\Delta^p} X^{\kappa,l-1}_p\times\Delta^p.$$

LEMMA 4.7. *These maps glue to a homotopy* $\mathscr{S}\colon[0,1]\times|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|\to|X^{\kappa,l-1}_\bullet|$.

Proof. The points $F_p(s,x,d^it)$ and $F_{p-1}(s,d_ix,t)$ are identified under the usual face maps among the $X^{\kappa,l-1}_p\times\Delta^p$. This follows immediately from the formula (4.3) and the observation that $\lambda(d^it)=d^i(\lambda(t))$, $\psi(d^it)=d^i(\psi(t))$ and $\sum_{j=0}^p\psi_j(d^it)=\sum_{j=0}^p\psi_j(t)$. □

Proof of Theorem 4.1. We claim that the map $\mathscr{S}(1,-)\colon|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|\to|X^{\kappa,l-1}_\bullet|$ factors through the continuous injection $|X^{\kappa,l}_\bullet|\to|X^{\kappa,l-1}_\bullet|$. Informally, this is because the functions λ_i and ψ_i have the property that either $\lambda_i(t)=1$ or $\psi_i(t)=0$. In the first case, the surgery on the ith critical level is performed completely (and the other surgeries do not affect the connectivity at level a_i). In the second case, the surgery is performed incompletely, but as $\psi_i(t)=0$, the behaviour at a_i is discarded. This pointwise argument does not prove continuity of the factorisation, but that can be seen at the level of the maps F_p, since the domain of F_p is covered by the 2^{p+1} closed sets obtained by requiring for each i either $\lambda_i(t)=1$ or $\psi_i(t)=0$, on each of which the map $F_p(1,-)$ composed with $X_p^{\kappa,l-1}\times\Delta^p\to|X_\bullet^{\kappa,l-1}|$ factors continuously through $\coprod_{p'\leqslant p}(X_{p'}^{\kappa,l}\times\Delta^{p'})\to|X_\bullet^{\kappa,l}|$.

The homotopy $\mathscr{S}$ is constant when precomposed with $|D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_\bullet|\to|D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_{\bullet,0}|$, and by the argument in §3.3 we deduce the weak equivalence in Theorem 4.1. $\square$

5. Surgery on objects in the middle dimension

We now restrict our attention to even dimensions, and write $d=2n$. Given a collection of path components $\mathcal{A}\subset\pi_0(\mathrm{Ob}(\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)))$, in Definition 2.11 we defined

$$\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)\subset\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$$

to be the full subcategory on this collection of objects. To state our main theorem concerning these subcategories, we first need two definitions.

Definition 5.1. The reflection of a morphism $C=(t,(W,\ell))\in\mathcal{C}_{\theta,L}(\mathbb{R}^N)$ is the morphism $\overleftarrow{C}=(t,(\overleftarrow{W},\overleftarrow{\ell}))$ defined as follows. The underlying manifold $\overleftarrow{W}\subset[0,t]\times\mathbb{R}^N$ is given by the reflection of W in the hyperplane $\left\{\frac{1}{2}t\right\}\times\mathbb{R}^N$ and the structure $\overleftarrow{\ell}\colon T\overleftarrow{W}\to\theta^*(\gamma)$ is defined by precomposing ℓ with the induced bundle map $T\overleftarrow{W}\to TW$. This construction defines an isomorphism of topological categories

$$\mathcal{C}_{\theta,L}(\mathbb{R}^N)^{\mathrm{op}}\longrightarrow\mathcal{C}_{\theta,\bar{L}}(\mathbb{R}^N).$$

In particular the reflection sends an object M to an object $\overleftarrow{M}$ with the same underlying manifold, but the tangential structure is precomposed with the bundle endomorphism $(-1)\oplus\mathrm{Id}\colon\varepsilon^1\oplus TM\to\varepsilon^1\oplus TM$. This new object is *not equal* to the original M and the reflection $\overleftarrow{C}$ of a morphism $M\rightsquigarrow N$ is not a morphism $N\rightsquigarrow M$, but sometimes it will be after some change of tangential structure.

Definition 5.2. We say a tangential structure θ is *reversible* if whenever there is a morphism $C\colon M\rightsquigarrow N$ in $\mathcal{C}_{\theta,L}$, there also exists a morphism $C'\colon N\rightsquigarrow M$ in this category, whose underlying manifold is equal to that of the reflection $\overleftarrow{C}\colon\overleftarrow{N}\rightsquigarrow\overleftarrow{M}$.

Although it seems that reversibility is a property of the pair (θ,L) rather than just θ, in Proposition 5.7 we prove that it is equivalent to θ being *spherical*, as defined in §1 (i.e. the $2n$-sphere admits a θ-structure extending any given structure on the lower hemisphere). In particular, the definition of reversibility does not depend on L.

We can now state our main theorem concerning these subcategories, analogous to Theorem 4.1 but in the middle dimension. The reader mainly interested in Theorems 1.1 and 1.2 can take $\theta=\theta^n: B\mathrm{O}(2n)\langle n\rangle \to B\mathrm{O}(2n)$, $L\cong D^{2n-1}$, $N=\infty$, and $\mathcal{A}$ being the class of objects which are either diffeomorphic to S^{2n-1} with its standard smooth structure and θ-structure or are *not* θ-bordant to S^{2n-1}. (Again, the proof does not simplify much in this special case.)

THEOREM 5.3. *Suppose that*

(i) $2n\geqslant 6$;

(ii) $3n<N$;

(iii) θ *is reversible*;

(iv) L *admits a handle decomposition only using handles of index at most* $n-1$;

(v) *the map* $\ell_L: L\to B$ *is* $(n-1)$*-connected*;

(vi) *the natural map* $\mathcal{A}\to\pi_0(B\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N))$ *is surjective.*

Then

$$B\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)\longrightarrow B\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$$

is a weak homotopy equivalence.

The surgery move that we will employ is similar to that of the last section, but has a crucial difference. In the last section, when we performed the surgery move to make a_i be a good regular value, we glued a family of manifolds having the effect of performing l-surgery on the level sets $W|_{a_i}$, but at the same time performing l-surgery on all higher level sets. In §4, $l<\frac{1}{2}(d-2)=n-1$, and therefore performing l-surgery on a $(2n-1)$-manifold which is l-connected preserves its l-connectedness. In this section, we will need to change level sets by doing $(n-1)$-surgery on $(2n-1)$-manifolds, and this is much more delicate. To explain why the approach of §4 needs modification, we refer to Figure 8 in that section, depicting the case $n=1$: The level set at a_1 is already connected and we have picked surgery data for making the level set at a_0 connected, but gluing in the surgery move shown in Figure 8 would disconnect the level set at a_1.

Instead we use a modified surgery move, which will let us perform $(n-1)$-surgery on a level set $W|_a$ and leave all other level sets $W|_b$ unchanged, except when b is very close to a. For $n=1$, this was done in [GRW1], and the construction there generalises to higher n. Let us briefly recall and depict the case $n=1$.

Suppose that we have a point $\sigma=(t_0,t_1;a_0,a_1;\varepsilon_0,\varepsilon_1;W)\in BD^{0,-1}$ on a 1-simplex, and that the level set $W|_{a_1}$ is already path connected. We will describe a way to produce a path from the image of this point in $|X_\bullet^{0,-1}|$ to the subset $|X_\bullet^{0,0}|$. We start with the

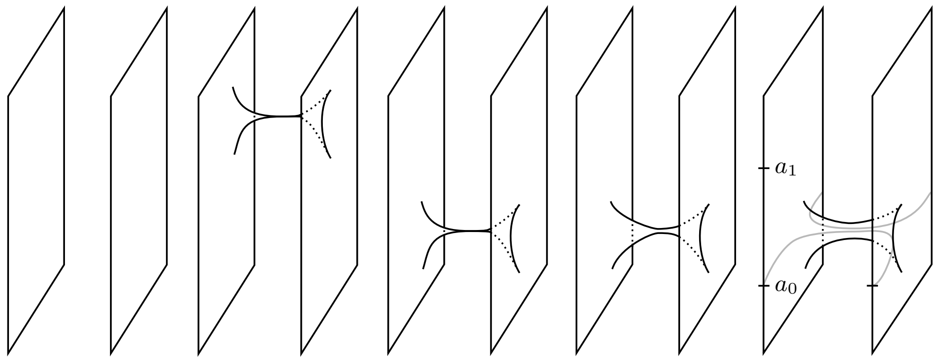

Figure 9. The surgery move on objects in the middle dimension in the case $n=1$ and $l=0$. The pictures show a one-parameter family of 2-manifolds embedded in $\mathbb{R}^4$, depicted by their projection to $\mathbb{R}^3$. The manifolds are embedded throughout the move: drawing only three coordinates causes the apparent singularity. In the last frame we have indicated the level set at a_0 in light grey, to emphasise that it *is* modified by the surgery.

same surgery data as in §4, a collection of embeddings

$$\{e_\alpha\colon (a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}\times D^1 \hookrightarrow \mathbb{R}\times\mathbb{R}^\infty\}_{\alpha\in\Lambda},$$

such that if surgery is performed along each $\{a_0\}\times\mathbb{R}\times\partial D^1 \hookrightarrow W|_{a_0}$ then the level set at a_0 becomes connected. Then we thicken the e_α to embeddings

$$\{\bar{e}_\alpha\colon \mathbb{R}\times(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}\times D^1 \hookrightarrow \mathbb{R}\times\mathbb{R}^\infty\}_{\alpha\in\Lambda},$$

which intersect W only in $\{0\}\times(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}\times S^0$. A typical example with $n=1$ of that part of the data lying in W is shown in Figure 7 in §4. Figure 9 depicts a one-parameter family of manifolds starting with $\{0\}\times(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}\times S^0$, shown in five representative instants.

We now glue this family into the image of each $\bar{e}_\alpha$. This defines a path in the space $|X_\bullet^{0,-1}|$, and if the handle in Figure 9 at time 1 is "thin" enough (with respect to the height function) so that it does not affect the level-set at a_1, then the manifold $\overline{W}$ obtained at the end of the path has $\overline{W}|_{a_0}$ and $\overline{W}|_{a_1}$ both path connected and so lies in $|X_\bullet^{0,0}|$.

Let us also briefly explain our motivation for the additional coordinate direction we added to $\bar{e}_\alpha$, and the singular-looking form of the one-parameter family in Figure 9. We could have constructed a similar one-parameter family of 2-manifolds inside the space $(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}\times D^1$ which slides down a tube of positive width (with respect to the height function) until it arrives at height a_0. However, during this family the centre of the tube will at some point lie at height a_1, so a_1 will be a regular value inside $(a_1-\varepsilon_1, a_1+\varepsilon_1)$ whose level-set may well not be path connected. Thus, even if we started with $\sigma\in|X_\bullet^{0,0}|$

we may have left this subset. This does not affect the argument above, which showed that $|X_\bullet^{0,0}|\to|X_\bullet^{0,-1}|$ is 0-connected, but becomes a problem when trying to show higher connectivity: it is important that we always construct *relative* homotopies. To avoid this problem, we use the family of Figure 9, where the tube has width zero while it is moving through the level-set at a_1: thus when the centre of the tube is at height $t\in(a_1-\varepsilon_1, a_1+\varepsilon_1)$, t is not regular and we impose no condition on its level set, whereas the level set at $t'\neq t$ is unchanged up to diffeomorphism. The extra coordinate direction added to $\bar{e}_\alpha$ is necessary as we cannot make the tube have width zero and be embedded in $(a_0-\varepsilon_0, a_1+\varepsilon_1)\times\mathbb{R}\times D^1$: we require an extra dimension.

In order to make sense of this surgery move in the presence of θ-structures, we must again equip the one-parameter family of manifolds shown in Figure 9 with θ-structures which start at a given structure, are constant near the vertical boundaries, and at the end of the path the level sets above and below the handle should be isomorphic as θ-manifolds to the level sets before the handle was added. This last property does not hold in general: for example, if we equip the original manifold in Figure 9 with a framing, one may easily see (using the Poincaré–Hopf theorem) that there is no framing on the final manifold consistent with these requirements. As we will see, this problem goes away when θ is assumed to be reversible. Let us first discuss the reversibility condition in more detail.

5.1. Reversibility

Recall that a tangential structure $\theta\colon B\to B\mathrm{O}(d)$ is called *spherical* if any θ-structure on a disc $D\subset S^d$ extends to one on S^d. (When B is path connected, this is equivalent to the d-sphere admitting any θ-structure at all.) Let us first discuss some related conditions on tangential structures $\theta\colon B\to B\mathrm{O}(d)$.

Definition 5.4. A tangential structure $\theta\colon B\to B\mathrm{O}(d)$ is *once-stable* if there exists a map $\bar{\theta}\colon \overline{B}\to B\mathrm{O}(d+1)$ and a commutative diagram

$$\begin{tikzcd} B \arrow[r] \arrow[d, "\theta"'] & \overline{B} \arrow[d, "\bar{\theta}"] \\ B\mathrm{O}(d) \arrow[r] & B\mathrm{O}(d+1) \end{tikzcd} \tag{5.1}$$

which is homotopy cartesian, i.e. the induced map from B to the homotopy pullback is a weak equivalence.

A tangential structure θ is *weakly once-stable* if there exists such a diagram which is d-cartesian, i.e. the induced map from B to the homotopy pullback is d-connected.

From the commutative diagram (5.1), there is a bundle map $\varepsilon^1\oplus\theta^*\gamma\to\bar\theta^*\gamma$. Hence a θ-structure $TW\to\theta^*\gamma$ on a d-manifold W induces a bundle map $\varepsilon^1\oplus TW\to\bar\theta^*\gamma$. If θ is weakly once-stable we may deduce the converse, that a bundle map $\varepsilon^1\oplus TW\to\bar\theta^*\gamma$ is homotopic to one that arises from a θ-structure. More precisely, we have the following useful lemma.

LEMMA 5.5. *Let $\theta: B\to B\mathrm{O}(d)$ be weakly once-stable. Let W be a d-manifold and let $\ell: TW|_A\to\theta^*\gamma$ be a θ-structure defined on a closed submanifold $A\subset W$. Then ℓ extends to a θ-structure $TW\to\theta^*\gamma$ if and only if the bundle map $\varepsilon^1\oplus TW|_A\to\varepsilon^1\oplus\theta^*\gamma\to\bar\theta^*\gamma$ extends to a bundle map over all of W.*

Proof. Without loss of generality, we may assume that θ and $\bar\theta$ are Serre fibrations. Let us write $s: B\mathrm{O}(d)\to B\mathrm{O}(d+1)$ for the stabilisation map, and let us pick a classifying map $t: W\to B\mathrm{O}(d)$ for the tangent bundle. Tangential structures on TW then correspond to lifts of t along some fibration, and tangential structures on $\varepsilon^1\oplus TW$ correspond to lifts of $s\circ t$ along some fibration.

We write $\tilde\theta: \widetilde{B}\to B\mathrm{O}(d)$ for the pullback of $\bar\theta$, so the commutative diagram (5.1) gives a map $i: B\to\widetilde{B}$ over $B\mathrm{O}(d)$. A $\bar\theta$-structure on $\varepsilon^1\oplus TW$ is then nothing but a $\tilde\theta$-structure on TW. By definition of being weakly once-stable, the map i is d-connected. Now, the situation described in the statement is a lifting problem

$$\begin{tikzcd} A \arrow[r] \arrow[d] & B \arrow[d, "i"] \\ W \arrow[r] \arrow[ur, dashed] & \widetilde{B}, \end{tikzcd}$$

which has a solution as (W,A) has cells of dimension at most d and i is d-connected. □

LEMMA 5.6. *The tangential structure $\theta: B\to B\mathrm{O}(d)$ is weakly once-stable if and only if it is spherical.*

Proof. Given any bundle map $\ell: TS^d|_D\to\theta^*\gamma$ we can of course extend the stabilised map to $\varepsilon^1\oplus TS^d\to\varepsilon^1\oplus\theta^*\gamma$, and if θ is weakly once-stable, the above lemma implies that the θ-structure extends to all of S^d.

Conversely, given a spherical structure $\theta: B\to B\mathrm{O}(d)$ we define $\overline{B}$ in diagram (5.1) as the dth Moore–Postnikov factorisation of the composite map $s\circ\theta: B\to B\mathrm{O}(d+1)$, where we again write $s: B\mathrm{O}(d)\to B\mathrm{O}(d+1)$ for the stabilisation map. For each choice of base-point in B, we may write F for the homotopy fibre of $B\to\overline{B}$. The diagram then induces a map from F to the homotopy fibre of s, which is weakly equivalent to S^d, and both spaces are $(d-1)$-connected. The long exact sequence in homotopy groups gives rise to the diagram

$$\begin{tikzcd}
\pi_{d+1}(\overline{B}) \arrow[r] \arrow[d, equal] & \pi_d(F) \arrow[r] \arrow[d] & \pi_d(B) \arrow[r] \arrow[d] & \pi_d(\overline{B}) \arrow[d, hook] \\
\pi_{d+1}(B\mathrm{O}(d+1)) \arrow[r] & \pi_d(S^d) \arrow[r] \arrow[ur, dashed] & \pi_d(B\mathrm{O}(d)) \arrow[r] & \pi_d(B\mathrm{O}(d+1)).
\end{tikzcd}$$

The map $S^d \to B\mathrm{O}(d)$ classifies the tangent bundle of S^d, and the assumption that θ be spherical guarantees the existence of a dashed arrow in the diagram making the bottom triangle commute. Now an easy diagram chase proves that $\pi_d(F) \to \pi_d(S^d)$ is surjective, so $F \to S^d$ is d-connected. Since this holds for any basepoint in B (and the map from B to the homotopy pullback obviously is 0-connected), the diagram (5.1) is d-cartesian as claimed. □

We now show that these conditions on θ are also equivalent to reversibility.

Proposition 5.7. *The tangential structure θ is reversible if and only if it is spherical.*

Proof. If θ is reversible and a θ-structure on D^d is given, we think of D^d as a morphism from the empty set to S^{d-1}. By assumption, a compatible θ-structure exists on the disc, thought of as a morphism from S^{d-1} to the empty set.

For the reverse direction we use Lemmas 5.5 and 5.6. Given a bordism

$$C = (t, (W, \ell)) \colon M \rightsquigarrow N$$

we obtain a reflected bordism

$$\overleftarrow{C} = (t, (\overleftarrow{W}, \overleftarrow{\ell})) \colon \overleftarrow{N} \rightsquigarrow \overleftarrow{M}$$

as in Definition 5.1, and we wish to find a new tangential structure $\ell' \colon T\overleftarrow{W} \to \theta^*\gamma$ such that $C' = (t, (\overleftarrow{W}, \ell'))$ is a morphism $N \rightsquigarrow M$ in $\mathcal{C}_{\theta,L}$. Near the subset

$$A = (\{0\} \times N) \cup ([0,t] \times L) \cup (\{t\} \times M) \subset \overleftarrow{W},$$

we must define ℓ' as $\overleftarrow{\ell} \circ R$, where the map $R \colon T\overleftarrow{W}|_A \to T\overleftarrow{W}|_A$ is the bundle automorphism which changes sign in the first coordinate, and it suffices to prove that this bundle map $T\overleftarrow{W}|_A \to \theta^*\gamma$ extends to all of $T\overleftarrow{W}$. By Lemmas 5.5 and 5.6 it is enough to extend the stabilised map $\varepsilon^1 \oplus T\overleftarrow{W}|_A \to \varepsilon^1 \oplus \theta^*\gamma$, but this is now easy. Indeed, the first basis vector gives a non-zero section of $T\overleftarrow{W}|_A \subset \varepsilon^1 \oplus T\overleftarrow{W}|_A$, which may be extended to a non-vanishing section V of $\varepsilon^1 \oplus T\overleftarrow{W}$. The bundle automorphism $R' \colon \varepsilon^1 \oplus T\overleftarrow{W} \to \varepsilon^1 \oplus T\overleftarrow{W}$ which multiplies by -1 in the line field spanned by V and by $+1$ in its orthogonal complement then extends $\varepsilon^1 \oplus R \colon \varepsilon^1 \oplus T\overleftarrow{W}|_A \to \varepsilon^1 \oplus T\overleftarrow{W}|_A$, and the composition $(\varepsilon^1 \oplus \overleftarrow{\ell}) \circ R'$ gives the required extension. □

One key property of reversible tangential structures is that they allow us to form connected sums of θ-manifolds, which of course is not possible in general: the connected sum of framed manifolds is not typically framable. In fact, more is true.

PROPOSITION 5.8. *Let (M,ℓ_M) be a d-dimensional θ-manifold, and suppose that*

$$e_0: S^l \hookrightarrow M$$

is an embedded sphere such that the map $S^l\to B$ induced by $\ell_M\circ e_0$ is null-homotopic. Then, if θ is reversible and $2(l+1)\leqslant d+1$, there is an extension of e_0 to an embedding

$$e: S^l\times D^{d-l} \hookrightarrow M$$

such that the surgered manifold

$$\overline{M}=(M\setminus \operatorname{int}(e(S^l\times D^{d-l})))\cup_{S^l\times S^{d-l-1}}(D^{l+1}\times S^{d-l-1})$$

admits a θ-structure which agrees with ℓ_M on $M\setminus\operatorname{int}(e(S^l\times D^{d-l}))$.

Proof. As θ is reversible it is also weakly once-stable, so let $\bar\theta:\overline{B}\to B\mathrm{O}(d+1)$ be a tangential structure exhibiting it as such, and $\bar\ell_M: M\to B\to\overline{B}$ be the induced $\bar\theta$-structure. The composition $\bar\ell_M\circ e_0: S^l\to B\to\overline{B}$ is null-homotopic, so we apply the discussion in §4.1 to the $\bar\theta$-manifold M and the embedding e_0. As $2(l+1)\leqslant d+1$, that discussion shows that there is a framing of the normal bundle of e_0, giving the embedding e, such that the trace $V: M\rightsquigarrow\overline{M}$ of the surgery along e has a $\bar\theta$-structure $\bar\ell_V$ such that $\bar\ell_V|_{\overline{M}}$ agrees with $\bar\ell_M$ on $M\setminus\operatorname{int}(S^l\times D^{d-l})$.

Thus the θ-structure ℓ_M restricted to $M\setminus\operatorname{int}(S^l\times D^{d-l})\subset\overline{M}$ extends to a $\bar\theta$-structure $\bar\ell_V|_{\overline{M}}$ on $\overline{M}$, so by the proof of Lemma 5.5 it also extends to a θ-structure, as claimed. □

For tangential structures that are once-stable (not just weakly), we can say that for a d-manifold W with a fixed θ-structure $\ell_0: TW|_{\partial W}\to\theta^*\gamma$, the stabilisation map

$$\mathrm{Bun}^\partial(TW,\theta^*\gamma;\ell_0)\longrightarrow \mathrm{Bun}^\partial(\varepsilon^1\oplus TW,\bar\theta^*\gamma;\varepsilon^1\oplus\ell_0)$$

is a weak homotopy equivalence. (Weakly once-stable only implies that this map is 0-connected.) We shall not make explicit use of this stronger condition in this paper, but point out that most of the naturally occurring tangential structures *are* once-stable. In particular, the following construction will be our main source of once-stable tangential structures. Let W be a connected d-dimensional manifold and $\tau: W\to B\mathrm{O}(d)$ be its Gauss map. For each k we have the kth Moore–Postnikov factorisations of τ,

$$W\xrightarrow{j_k} B_W(k)\xrightarrow{p_k} B\mathrm{O}(d).$$

Then $\theta_W(k)=p_k$ is a tangential structure and j_k gives a canonical $\theta_W(k)$-structure on W.

Lemma 5.9. *The tangential structure $\theta_W(k)\colon B_W(k)\to B\mathrm{O}(d)$ is once-stable for any $k<d$.*

Proof. We let $\overline{B}_W(k)$ denote the same Moore–Postnikov construction applied to the composition $W\to B\mathrm{O}(d)\to B\mathrm{O}(d+1)$. The claim then follows as $B\mathrm{O}(d)\to B\mathrm{O}(d+1)$ is d-connected. □

Remark 5.10. There do exist tangential structures which are reversible but not once-stable, which justifies our emphasis on reversibility. An interesting example is the map $B\mathrm{U}(3)\to B\mathrm{O}(6)$, which is reversible as S^6 admits an almost-complex structure, but is not once-stable: if it were pulled back from a fibration $f\colon\overline{B}\to B\mathrm{O}(7)$, one can easily use the Serre spectral sequence to check that the kernel of the map f^* on $\mathbb{F}_2$-cohomology would be the ideal $I=(w_1,w_3,w_5)\subset H^*(B\mathrm{O}(7);\mathbb{F}_2)$, but this is not closed under the action of the Steenrod algebra as $\mathrm{Sq}^4(w_5)=w_4w_5+w_3w_6+w_2w_7\notin I$.

5.2. The standard family

We will prove Theorem 5.3 by performing $(n-1)$-surgery on objects until we reach an object in $\mathcal{A}$, just as in §4 we performed l-surgery on objects to make them l-connected (relative to L). As in that section, the surgery shall be performed by gluing in a suitable family of manifolds along certain families of embeddings, whose existence we shall prove in §6. The standard family to be glued in is very similar to that in §4, and is depicted for $n=1$ in Figure 9. The reader may compare that figure with Figure 8 to get a feeling for the similarities and differences between the two constructions. In §4 we started with a certain submanifold $K\subset\mathbb{R}^{d-l}\times\mathbb{R}^{l+1}$. In this section, we shall use the same manifold, with $d=2n$ and $l=n-1$. Recall that we first chose a smooth function $\varrho\colon\mathbb{R}\to\mathbb{R}$ which is the identity on $\left(-\infty,\frac{1}{2}\right)$, has nowhere negative derivative, and $\varrho(t)=1$ for all $t\geqslant1$, and we let

$$K=\{(x,y)\in\mathbb{R}^{n+1}\times\mathbb{R}^n:|y|^2=\varrho(|x|^2-1)\}.$$

The first coordinate restricts to a Morse function $h=x_1\colon K\to\mathbb{R}$ with exactly two critical points: $(-1,0,...,0;0)$ and $(+1,0,...,0;0)$ both of index n.

In §4, we constructed from K a one-parameter family of manifolds

$$\mathcal{P}_t\subset(-6,-2)\times\mathbb{R}^{d-l-1}\times\mathbb{R}^{l+1},$$

obtained from K by moving the lowest critical point downwards as $t\in[0,1]$ increases, as in Figure 8. In this section we shall need a one-parameter family

$$\mathcal{P}_t\subset\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times\mathbb{R}^n$$

which we will construct from $\{0\}\times K$ by moving *both* critical points down as $t\in[0,1]$ increases, and varying the distance between the corresponding critical values as shown in Figure 9. As an intermediate step we construct a two-parameter family $\mathcal{P}_{t,w}$ where $w\geqslant 0$ controls the distance between the two critical values. In order for the manifold to stay embedded in the limit $w=0$, we need the extra ambient dimension.

Let us first construct a one-parameter family of submanifolds $K_w\subset\mathbb{R}\times\mathbb{R}^{n+1}\times D^n$ such that $K_1=\{0\}\times K$. Let $\mu:\mathbb{R}\to[0,1]$ be a smooth function such that $\mu^{-1}(0)=[2,\infty)$, $\mu^{-1}(1)=(-\infty,\sqrt{2}\,]$ and $\mu'<0$ on $(\sqrt{2},2)$, and define a one-parameter family of embeddings

$$\begin{aligned}\varphi_w:\mathbb{R}^{n+1}\times D^n &\longrightarrow \mathbb{R}\times\mathbb{R}^{n+1}\times D^n,\\ (x,y)&\longmapsto (x_1(1-w)\mu(|x|), x_1(1-(1-w)\mu(|x|)), x_2, ..., x_{n+1}, y).\end{aligned}$$

We now let

$$K_w=\varphi_w(K)\subset\mathbb{R}\times\mathbb{R}^{n+1}\times D^n$$

for $w\in[0,1]$, and we remark that these manifolds are all diffeomorphic (to each other and to K), as it is only the embedding that we are varying. The critical points of the height function $h: K_w\to\mathbb{R}$ correspond to the critical points of the function

$$\begin{aligned}h_w: K&\longrightarrow\mathbb{R},\\ (x,y)&\longmapsto x_1(1-(1-w)\mu(|x|)).\end{aligned}$$

In the region $K\cap\{(x,y):|x|^2\leqslant 2\}$ the function μ is constantly 1, and $h_w=wx_1$. As long as $w>0$, this function is Morse and has critical points $(\pm1,0,0,...,0)$, with critical values $\pm w$. The manifold $K\cap\{(x,y):|x|^2>2\}$ is equal to $\{(x,y)\in\mathbb{R}^{n+1}\times\mathbb{R}^n:|y|=1 \text{ and } |x|^2>2\}$, so in this region it makes sense to take the partial derivative of h_w with respect to the coordinate x_1 and we calculate

$$\frac{\partial h_w}{\partial x_1}=1-(1-w)\left[\mu(|x|)+\mu'(|x|)\frac{x_1^2}{|x|}\right].$$

When $|x|^2>2$ we have $\mu(|x|)<1$ and $\mu'(|x|)\leqslant 0$, so the square bracket is strictly smaller than 1. Since $0\leqslant(1-w)\leqslant 1$ we see that $\partial h_w/\partial x_1>0$, so h_w has no critical points in $K\cap\{(x,y):|x|^2>2\}$. To summarise, we have shown that for $w>0$ the critical points of the height function $h: K_w\to\mathbb{R}$ are $\varphi_w(\pm1,0,...,0)$, are Morse of index n, and lie at heights $\pm w$. When $w=0$ the function $h: K_0\to\mathbb{R}$ is constantly 0 on the entire subset

$$\varphi_0(K\cap\{(x,y):|x|^2\leqslant 2\})\approx S^n\times D^n,$$

but has no other critical values.

We now define a two-parameter family of d-dimensional submanifolds $\mathcal{P}_{t,w}$ inside $\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times D^n$ in much the same way as $\mathcal{P}_t$ was constructed from K in §4.1. Apart from the extra width parameter, the main difference is that in this section we will use a larger part of K, including *both* critical points. Pick a smooth one-parameter family of embeddings $\lambda_s\colon(-6,-2)\to(-6,2)$, such that $\lambda_0=\mathrm{Id}$, $\lambda_s|_{(-6,-5)}=\mathrm{Id}$ for all s, $\lambda_1(-4)=-1$ and $\lambda_1(-3)=1$. Then we get embeddings

$$\mathrm{Id}_{\mathbb{R}}\times\lambda_t\times\mathrm{Id}_{\mathbb{R}^{2n}}\colon\mathbb{R}\times(-6,-2)\times\mathbb{R}^{2n}\longrightarrow\mathbb{R}\times(-6,2)\times\mathbb{R}^{2n}$$

and define

$$\mathcal{P}_{t,w}=(\mathrm{Id}_{\mathbb{R}}\times\lambda_t\times\mathrm{Id}_{\mathbb{R}^{2n}})^{-1}(K_w)\in\Psi_d(\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times\mathbb{R}^n).$$

As in §4.2, it is easy to verify that $\mathcal{P}_{t,w}$ agrees with $\{0\}\times(-6,-2)\times\mathbb{R}^n\times S^{n-1}$ outside of $(-2,2)\times(-5,-2)\times B_2^n(0)\times D^n$, independently of t and w.

We shall also need a tangentially structured version of this construction, given a structure $\ell\colon TK|_{(-6,2)}\to\theta^*\gamma$. For this purpose, let $\omega=\mu\colon\mathbb{R}\to[0,1]$ be the function defined above and define a one-parameter family of embeddings by

$$\begin{aligned}\psi_t\colon\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times\mathbb{R}^n&\longrightarrow\mathbb{R}\times(-6,2)\times\mathbb{R}^n\times\mathbb{R}^n,\\ (s;x_1,...,x_{n+1};y)&\longmapsto(s;\lambda_{t\omega(|x|)}(x_1),x_2,...,x_{n+1};y).\end{aligned}$$

As in §4.2, it is easy to see that we also have $\psi_t^{-1}(K_w)=(\mathrm{Id}_{\mathbb{R}}\times\lambda_t\times\mathrm{Id}_{\mathbb{R}^{2n}})^{-1}(K_w)=\mathcal{P}_{t,w}$, and we define a θ-structure on $\mathcal{P}_{t,w}$ by pullback along ψ_t. This gives a continuous two-parameter family

$$\mathcal{P}_{t,w}(\ell)\in\Psi_\theta(\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times\mathbb{R}^n).$$

Let $P\colon[0,1]\to[0,1]^2$ be the piecewise linear path with $P(0)=(0,0)$, $P\left(\frac{1}{2}\right)=(1,0)$ and $P(1)=(1,1)$, and define a continuous one-parameter family

$$\mathcal{P}_t(\ell)=\mathcal{P}_{P(t)}(\ell)\in\Psi_\theta(\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times\mathbb{R}^n).$$

We will omit ℓ from the notation when it is unimportant. We record some important properties of this family in Proposition 5.12 below, using the following definition.

Definition 5.11. Let $\ell\colon TK\to\theta^*\gamma$ be a θ-structure on K. Recall that outside of the region $\mathbb{R}\times B_2^n(0)\times D^n$ the manifold K agrees with $\mathbb{R}\times\mathbb{R}^n\times S^{n-1}$. We say that ℓ is *extendible* if the θ-structure $\ell|_{\mathbb{R}\times(\mathbb{R}^n\setminus B_2^n(0))\times S^{n-1}}$ extends to a θ-structure on the whole of $\mathbb{R}\times\mathbb{R}^n\times S^{n-1}$.

PROPOSITION 5.12. *Suppose that ℓ is extendible. The elements*

$$\mathcal{P}_t(\ell) \in \Psi_\theta(\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times\mathbb{R}^n)$$

are θ-submanifolds of $\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times D^n$ satisfying the following properties:

(i) $\mathcal{P}_0(\ell)=K_1|_{(-6,-2)}=\{0\}\times(-6,-2)\times\mathbb{R}^n\times S^{n-1}$ *as θ-manifolds.*

(ii) *For all t, $\mathcal{P}_t(\ell)$ agrees with $K_1|_{(-6,-2)}$ as a θ-manifold, outside of the region $(-2,2)\times(-5,-2)\times B^n_2(0)\times D^n$.*

(iii) *For all t and each pair of regular values $-6<a<b<-2$ of the height function $h\colon\mathcal{P}_t\to\mathbb{R}$, the pair*

$$(\mathcal{P}_t|_{[a,b]},\mathcal{P}_t|_b)$$

is $(n-1)$-connected.

(iv) *If a is outside of $(-4,-3)$ and is a regular value of $h\colon\mathcal{P}_t(\ell)\to(-6,-2)$ then the manifold $\mathcal{P}_t(\ell)|_a$ is isomorphic to $\mathcal{P}_0(\ell)|_a=\{0\}\times\{a\}\times\mathbb{R}^n\times S^{n-1}$ as a θ-manifold. If a is inside of $(-4,-3)$ and is a regular value of h then the manifold $\mathcal{P}_t(\ell)|_a$ is either isomorphic to $\mathcal{P}_0(\ell)|_a$ as a θ-manifold, or is obtained from it by $(n-1)$-surgery along the standard embedding.*

(v) *The critical values of $h\colon\mathcal{P}_1(\ell)\to(-6,-2)$ are -4 and -3. For $a\in(-4,-3)$, $\mathcal{P}_1(\ell)|_a$ is obtained by $(n-1)$-surgery from $\mathcal{P}_0(\ell)|_a=\{0\}\times\mathbb{R}^n\times S^{n-1}$ along the standard embedding.*

In (iv) *and* (v), *the θ-structure on the surgered manifold is determined* (*up to homotopy*) *by the θ-structure ℓ on $K|_{(-6,2)}$.*

The precise meaning of the word isomorphic in (iv) above is the following: By (ii) we know that the manifolds are *equal* outside $(-2,2)\times(-5,-2)\times B^n_2(0)\times D^n$. Being isomorphic means that the identity extends to a diffeomorphism which preserves θ-structures up to a homotopy of bundle maps which is constant outside $(-2,2)\times(-5,-2)\times B^n_2(0)\times D^n$.

Proof. Statements (i) and (ii) follow easily from the properties of λ_t and ψ_t, and the fact that K agrees with $\mathbb{R}^{n+1}\times S^{n-1}$ outside $B^{n+1}_2(0)\times\mathbb{R}^n$. For (iii), the Morse function $\mathcal{P}_{t,w}\to(-6,-2)$ has at most two critical points, both of index n. If a critical value is in (a,b), then the pair is $(n-1)$-connected, otherwise it is ∞-connected as $\mathcal{P}_{t,w}|_{[a,b]}$ deformation retracts to $\mathcal{P}_{t,w}|_b$.

To see (iv) and (v), first suppose that $t\in\left[0,\frac{1}{2}\right]$. Then $P(t)$ has second coordinate 0, so $\mathcal{P}_t(\ell)$ has width zero. The function $K_0\to\mathbb{R}$ has exactly one critical value, 0, and therefore $K_0|_a$ is diffeomorphic to $\{a\}\times\mathbb{R}^n\times S^{n-1}$ for all regular values a, and extendibility implies that they are also isomorphic as θ-manifolds. If instead $t\in\left[\frac{1}{2},1\right]$, then $P(t)$ has first coordinate 1, and so $\mathcal{P}_t(\ell)$ is obtained from some K_w using the embedding $\mathrm{Id}_\mathbb{R}\times\lambda_1\times\mathrm{Id}_{\mathbb{R}^{2n}}$. The fact that the function $K_w\to\mathbb{R}$ has exactly two critical points with value $\pm w$ implies that $K_w|_a$ is diffeomorphic to $\{a\}\times\mathbb{R}^n\times S^{n-1}$ for regular values $a\in\mathbb{R}\setminus(-1,1)\subset\mathbb{R}\setminus(-w,w)$ and extendibility implies that they are also isomorphic as θ-manifolds. When $w=1$, for regular values $a\in(-1,1)$ we have that $K_1|_a$ is obtained from $\{a\}\times\mathbb{R}^n\times S^{n-1}$ by $(n-1)$-surgery along the standard embedding. □

5.3. Surgery data

We can now describe the semi-simplicial space of surgery data in the middle dimension. It is similar to the space of surgery data below the middle dimension, but because the standard family constructed in the previous section uses more of the manifold K, each surgery datum will include a choice of θ-structure on more of K.

Before doing so, we choose once and for all, smoothly in the data $(a_i,\varepsilon_i,a_p,\varepsilon_p)$, increasing diffeomorphisms

$$\psi=\psi(a_i,\varepsilon_i,a_p,\varepsilon_p)\colon(-6,-2)\cong(a_i-\varepsilon_i,a_p+\varepsilon_p) \tag{5.2}$$

sending $[-4,-3]$ linearly onto $\left[a_i-\tfrac{1}{2}\varepsilon_i,a_i+\tfrac{1}{2}\varepsilon_i\right]$.

Definition 5.13. Let $x=(a,\varepsilon,(W,\ell_W))\in D_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)_p$, and write $M_i=W|_{a_i}$. Define the set $Y_q(x)$ to consist of tuples (Λ,δ,e,ℓ), where $\Lambda\subset\Omega$ is a finite subset of the fixed infinite set Ω, $\delta\colon\Lambda\to[p]\times[q]$ is a function,

$$e\colon\Lambda\times\mathbb{R}\times(-6,-2)\times\mathbb{R}^n\times D^n\hookrightarrow\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$$

is an embedding, and ℓ is a bundle map $T(\Lambda\times K)\to\theta^*\gamma$. (In Definition 4.3, it was only defined on $T(\Lambda\times K|_{(-6,0)})$.) As in Definition 4.3, we write $\Lambda_{i,j}=\delta^{-1}(i,j)$,

$$e_{i,j}\colon\Lambda_{i,j}\times(a_i-\varepsilon_i,a_i+\varepsilon_i)\times\mathbb{R}^n\times D^n\longrightarrow\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$$

for the embedding obtained by restricting e and reparameterising using (5.2), and $\ell_{i,j}$ for the restriction of ℓ to $T(\Lambda_{i,j}\times K|_{(-6,0)})$. This data is required to satisfy the following conditions:

(i) $e^{-1}(W)=\Lambda\times\{0\}\times(-6,-2)\times\mathbb{R}^n\times\partial D^n$. We let

$$\partial e\colon\Lambda\times\{0\}\times(-6,-2)\times\mathbb{R}^n\times\partial D^n\hookrightarrow W$$

denote the embedding restricted to the boundary.

(ii) For $t\in\bigcup_{k=0}^{p}(a_k-\varepsilon_k,a_k+\varepsilon_k)$, we have $(x_1\circ e_{i,j})^{-1}(t)=\Lambda_{i,j}\times\mathbb{R}\times\{t\}\times\mathbb{R}^n\times D^n$.

(iii) The composition $\ell_W\circ D(\partial e)\colon T(\Lambda\times K|_{(-6,-2)})\to\theta^*\gamma$ agrees with the restriction of ℓ.

(iv) For each $\lambda\in\Lambda$, the restriction of ℓ to $T(\{\lambda\}\times K)$ is extendible.

For each j, the data $(e_{i,j}, \ell_{i,j})$ is enough to perform θ-surgery on M_i (as $K|_{(-6,0)}$ contains the trace of an $(n-1)$-surgery on the $(d-1)$-manifold $K|_{-2}$), and we further insist that

(v) The resulting θ-manifold $\overline{M}_i$ lies in $\mathcal{A}$.

For each x, $Y_\bullet(x)$ is a semi-simplicial set.

A typical example of a surgery datum is (partially) depicted in Figure 7 in §4. The figure has $p=1$, $q=0$ and $n=1$. Only the image $e(\Lambda\times\{0\}\times(-6,-2)\times\mathbb{R}^n\times\partial D^n)\subset W$ is shown.

Define a bi-semi-simplicial space $D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$ (augmented in the second semi-simplicial direction) from this, as in Definition 4.4. The main result about this bi-semi-simplicial space of manifolds equipped with surgery data is the following, whose proof we defer until §6.

Theorem 5.14. *Under the assumptions of Theorem* 5.3, *the maps*

$$|D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}| \longrightarrow |D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}| \longrightarrow |D^{n-1,n-2}_{\theta,L}(\mathbb{R}^N)_{\bullet}|$$

are weak homotopy equivalences, where the first map is the inclusion of 0*-simplices and the second is the augmentation, in the second simplicial direction.*

5.4. Proof of Theorem 5.3

The proof of this theorem will be almost identical with that of Theorem 4.1. Thus, suppose that the conditions in the statement of Theorem 5.3 are satisfied, and let $(a,\varepsilon,(W,\ell_W),e,\ell)\in D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{p,0}$. For each $i=0,...,p$, we have an embedding $e_i=e_{i,0}$ and a bundle map $\ell_i=\ell_{i,0}$, and precisely as in §4.4 we may construct a one-parameter family of elements $\mathcal{K}^t_{e_i,\ell_i}(W,\ell_W)\in\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N)$ for $t\in[0,1]$. From this, for each tuple $t=(t_0,...,t_p)\in[0,1]^{p+1}$ we may form the element

$$\mathcal{K}^t_{e,\ell}(W,\ell_W)=\mathcal{K}^{t_p}_{e_p,\ell_p}\circ...\circ\mathcal{K}^{t_0}_{e_0,\ell_0}(W,\ell_W)\in\Psi_\theta((a_0-\varepsilon_0,a_p+\varepsilon_p)\times\mathbb{R}^N).$$

To apply the same proof as that of Theorem 4.1, we need an analogue of Lemma 4.6 to tell us how the manifold improves when we apply the various surgery operations.

Lemma 5.15. *Firstly, the tuple* $\big(a,\tfrac{1}{2}\varepsilon,\mathcal{K}^t_{e,\ell}(W,\ell_W)\big)$ *is an element of* $X^{n-1,n-2}_p$. *Secondly, if* t_i *is* 1*—so the surgery for the regular value* a_i *is fully done—then for each regular value* $b\in\big(a_i-\tfrac{1}{2}\varepsilon_i,a_i+\tfrac{1}{2}\varepsilon_i\big)$ *of* $x_1\colon\mathcal{K}^t_{e,\ell}(W,\ell_W)\to\mathbb{R}$, *the* θ*-manifold* $\mathcal{K}^t_{e,\ell}(W,\ell_W)|_b$ *lies in* $\mathcal{A}$.

Proof. For the first part we must verify the conditions of Definition 2.18. This part of the argument of Lemma 4.6 applies equally well when $\kappa=n-1$ and $l=n-1$.

For the second part, we suppose $t_i=1$. Let $b\in\left(a_i-\frac{1}{2}\varepsilon_i, a_i+\frac{1}{2}\varepsilon_i\right)$ be a regular value of the height function on $\mathcal{K}^t_{e,\ell}(W,\ell_W)$ and define θ-manifolds

$$\overline{M}=\mathcal{K}^t_{e,\ell}(W,\ell_W)|_b, \quad \widetilde{M}=\mathcal{K}^{t_i}_{e_i,\ell_i}(W,\ell_W)|_b \quad \text{and} \quad M=W|_b.$$

By Definition 5.13 (v), performing surgery on M using the data (e_i,ℓ_i) gives a θ-manifold in $\mathcal{A}$. By Proposition 5.12 (v), $\mathcal{K}^{t_i}_{e_i,\ell_i}(W,\ell_W)$ has this surgery done, so $\widetilde{M}$ lies in $\mathcal{A}$. Now $\overline{M}$ is obtained from $\widetilde{M}$ by applying the remaining operations $\mathcal{K}^{t_j}_{e_j,\ell_j}$ for $j\neq i$, but by Proposition 5.12 (iv), applying each of these only changes $\widetilde{M}$ up to isomorphism (because $b\in\left(a_i-\frac{1}{2}\varepsilon_i, a_i+\frac{1}{2}\varepsilon_i\right)$, so it is not in $\left(a_j-\frac{1}{2}\varepsilon_j, a_j+\frac{1}{2}\varepsilon_j\right)$), so $\overline{M}$ lies in $\mathcal{A}$. $\square$

As in §4.4 we define a map

$$\begin{aligned}\mathscr{S}_p\colon [0,1]^{p+1}\times D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{p,0} &\longrightarrow X^{n-1,n-2}_p,\\ (t,(a,\varepsilon,(W,\ell_W),e,\ell)) &\longmapsto \left(a,\tfrac{1}{2}\varepsilon,\mathcal{K}^t_{e,\ell}(W,\ell_W)\right),\end{aligned}$$

which has the desired range by the first part of Lemma 5.15. The argument of §4.4 gives maps

$$F_p\colon [0,1]\times D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{p,0}\times\Delta^p\longrightarrow X^{n-1,n-2}_p\times\Delta^p$$

gluing to a homotopy $\mathscr{S}\colon[0,1]\times|D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|\to|X^{n-1,n-2}_\bullet|$ which is constant on the image of $|D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_\bullet|\hookrightarrow|D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|$. It also provides a factorisation of the map $\mathscr{S}(1,-)$ through the continuous injection $|X^{n-1,\mathcal{A}}_\bullet|\to|X^{n-1,n-2}_\bullet|$. The argument in §3.3 then gives the weak equivalence in Theorem 5.3.

6. Contractibility of spaces of surgery data

In order to finish the proofs of the results of the last three sections, we must supply proofs of Theorems 3.4, 4.5 and 5.14 concerning the bi-semi-simplicial spaces of manifolds equipped with surgery data.

6.1. The first part of Theorems 4.5 and 5.14

Theorems 4.5 and 5.14 both assert that two maps are weak equivalences. In either theorem, the proof for the first map will use that the second map is a weak equivalence, but is otherwise simpler, so we first consider the maps

$$|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|\longrightarrow|D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}| \quad\text{and}\quad |D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{\bullet,0}|\longrightarrow|D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}|.$$

Proof (*assuming the second part*). The proof in both cases is the same, so let us write $D_{\bullet,\bullet}$ for either $D_{\theta,L}^{\kappa,l}(\mathbb{R}^N)_{\bullet,\bullet}$ or $D_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)_{\bullet,\bullet}$. We define, for this proof only, a bi-semi-simplicial space $D'_{\bullet,\bullet}$ in the same way as $D_{\bullet,\bullet}$ except that the usual inequalities $a_i+\varepsilon_i<a_{i+1}-\varepsilon_{i+1}$ and $\varepsilon_i>0$ are replaced by $a_i\leqslant a_{i+1}$ and $\varepsilon_i\geqslant 0$ (so the intervals $[a_i-\varepsilon_i, a_i+\varepsilon_i]$ are allowed to overlap).

The inclusion $D_{\bullet,\bullet}\hookrightarrow D'_{\bullet,\bullet}$ is easily seen to be a levelwise weak homotopy equivalence, by spreading the a_i out and making the ε_i positive but small, so it is enough to work with $D'_{\bullet,\bullet}$ throughout and show that $|D'_{\bullet,0}|\to|D'_{\bullet,\bullet}|$ is a weak homotopy equivalence.

To do so, we describe a retraction $r\colon|D'_{\bullet,\bullet}|\to|D'_{\bullet,0}|$ which will be a weak homotopy inverse to the inclusion. The map r only modifies the a_i and barycentric coordinates, and does not change the underlying manifold $W\in\psi_\theta(N+1,1)$. There is a map

$$D'_{p,q}\longrightarrow D'_{(p+1)(q+1)-1,0}$$

given by considering $p+1$ regular values, each equipped with $q+1$ pieces of surgery data, as $(p+1)(q+1)$ not-necessarily distinct regular values, each with a *single* piece of surgery data. There is also a map

$$\Delta^p\times\Delta^q\longrightarrow\Delta^{(p+1)(q+1)-1}\subset\mathbb{R}^{(p+1)(q+1)}$$

with $(j+(q+1)i)$-th coordinate given by $(t,s)\mapsto s_it_j$. Taking the product of these maps gives

$$r_{p,q}\colon D'_{p,q}\times\Delta^p\times\Delta^q\longrightarrow D'_{(p+1)(q+1)-1,0}\times\Delta^{(p+1)(q+1)-1}$$

which glue together to give the map $r\colon|D'_{\bullet,\bullet}|\to|D'_{\bullet,0}|$. It is clear that r is a retraction (i.e. left inverse to the inclusion), so the induced map on homotopy groups is surjective. To see that it is injective, we use the map $|D'_{\bullet,\bullet}|\to|D'_{\bullet}|$ induced by the augmentation in the second bi-semi-simplicial direction (by forgetting all surgery data). This is a weak equivalence by the second part of Theorem 4.5 or 5.14 respectively, but it clearly factors as

$$|D'_{\bullet,\bullet}|\xrightarrow{r}|D'_{\bullet,0}|\longrightarrow|D'_{\bullet}|,$$

where the second map is again induced by the augmentation in the second bi-semi-simplicial direction. Therefore r is also injective on homotopy groups, and hence a weak homotopy equivalence. □

6.2. A simplicial technique

In order to give the proofs of Theorems 3.4, 4.5 and 5.14, we need a technique for showing that for certain augmented semi-simplicial spaces $X_\bullet\to X_{-1}$, the map $|X_\bullet|\to X_{-1}$ is a weak homotopy equivalence. The semi-simplicial spaces occurring in those theorems are all of the following special type.

Definition 6.1. Let $X_\bullet \to X_{-1}$ be an augmented semi-simplicial space. We say it is an *augmented topological flag complex* if

(i) the map $X_n \to X_0 \times_{X_{-1}} \ldots \times_{X_{-1}} X_0$ to the $(n+1)$-fold fibre product—which takes an n-simplex to its $n+1$ vertices—is a homeomorphism onto its image, which is an open subset;

(ii) a tuple $(v_0, \ldots, v_n) \in X_0 \times_{X_{-1}} \ldots \times_{X_{-1}} X_0$ lies in X_n if and only if $(v_i, v_j) \in X_1$ for all $i<j$.

If elements $v, w \in X_0$ lie in the same fibre over X_{-1} and $(v, w) \in X_1$, we say that w is *orthogonal* to v. (We do not require the relation to be symmetric, although in our applications it will be.) If $X_{-1}=*$ we omit the adjective augmented.

The semi-simplicial space $Z_\bullet(a, \varepsilon, (W, \ell_W)) \to *$ from Definition 3.2 and the semi-simplicial spaces $Y_\bullet(a, \varepsilon, (W, \ell_W)) \to *$ from Definitions 4.3 and 5.13 are topological flag complexes. Furthermore we have that $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,\bullet} \to D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$ from Definition 3.3, $D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,\bullet} \to D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_p$ from Definition 4.4, and $D^{n-1,\mathcal{A}}_{\theta,L}(\mathbb{R}^N)_{p,\bullet} \to D^{n-1,n-2}_{\theta,L}(\mathbb{R}^N)_p$ from §5.3 are all augmented topological flag complexes. In all cases this is immediate from the definition: firstly, a p-simplex of these semi-simplicial spaces consists of $(p+1)$-tuples of surgery data, which are all 0-simplices; secondly, the pieces of surgery data are subject to the requirement that they are all disjoint, but disjointness is a property that can be verified pairwise.

Theorem 6.2. *Let $X_\bullet \to X_{-1}$ be an augmented topological flag complex. Suppose that the following are true*:

(i) *The map $\varepsilon: X_0 \to X_{-1}$ has local lifts of any map from a disc, i.e. given a map $f: D^n \to X_{-1}$, a point $x \in D^n$ and a point $p \in \varepsilon^{-1}(f(x))$, there is an open neighbourhood $U \subset D^n$ of x and a map $F: U \to X_0$ such that $\varepsilon \circ F = f|_U$ and $F(x)=p$.*

(ii) *$\varepsilon: X_0 \to X_{-1}$ is surjective.*

(iii) *For any $p \in X_{-1}$ and any (non-empty) finite set $\{v_1, \ldots, v_n\} \subset \varepsilon^{-1}(p)$ there exists a $v \in \varepsilon^{-1}(p)$ with $(v_i, v) \in X_1$ for all i.*

Then $|X_\bullet| \to X_{-1}$ is a weak homotopy equivalence.

Condition (ii) can be viewed as the $n=0$ analogue of condition (iii), but we prefer to keep the cases $n=0$ and $n>0$ separate.

Remark 6.3. To motivate the proof of this theorem, let us first consider the case where $X_{-1}=*$ and each X_i is discrete, so $|X_\bullet|$ has the structure of a Δ-complex. Then any map $f: S^n \to |X_\bullet|$ may be homotoped to be simplicial, for some triangulation of S^n, and so hits finitely many vertices $v_1, ..., v_k$. By (iii) there exists a $v \in X_0$ such that (v_i, v) is a 1-simplex for all i. But then the map f extends to the join

$$f * \{v\}: S^n * \{v\} \longrightarrow |X_\bullet|$$

and so f is null-homotopic.

The proof we give below follows this in spirit, although is necessarily more complicated when the X_i carry a topology. To deal with the topology, we require the following technical result.

PROPOSITION 6.4. *Let $Y_\bullet$ be a semi-simplicial set, and X be a Hausdorff space. Let $Z_\bullet \subset Y_\bullet \times X$ be a sub-semi-simplicial space which in each degree is an open subset. For $x \in X$, let $Z_\bullet(x) \subset Y_\bullet$ be the sub-semi-simplicial set defined by $Z_\bullet \cap (Y_\bullet \times \{x\}) = Z_\bullet(x) \times \{x\}$ and suppose that $|Z_\bullet(x)|$ is contractible for all $x \in X$. Then the map $\pi: |Z_\bullet| \to X$ is a Serre fibration with contractible fibres.*

Proof. This follows from [GRW2, Proposition 2.7] and [We, Lemma 2.2]. □

COROLLARY 6.5. *Let Ω be a set, X be a Hausdorff space, and let*

$$P \subset \mathbb{N} \times \Omega \times X$$

be a subset which is open (when $\mathbb{N}$ and Ω are given the discrete topology) and such that the projection $P \to \mathbb{N} \times X$ is surjective. We give $\mathbb{N} \times \Omega \times X$ the partial order defined by

$$(n, \alpha, x) < (m, \beta, y) \quad \text{if and only if} \quad n < m \text{ and } x = y,$$

and give the subspace P the induced order. Then the natural map $\pi: |N_\bullet P| \to X$ is a Serre fibration with contractible fibres.

Proof. We apply Proposition 6.4 with $Y_\bullet = N_\bullet(\mathbb{N} \times \Omega)$ and $Z_\bullet = N_\bullet P$. For $x \in X$, the semi-simplicial subset $Z_\bullet(x) \subset N_\bullet(\mathbb{N} \times \Omega)$ is contractible by the argument in Remark 6.3. □

Since the topology on the realisation of a semi-simplicial space is a quotient topology, there is no description of maps *into* it by a universal property. The following technical lemma concerns this question, and is the final technical preliminary for the proof of Theorem 6.2.

LEMMA 6.6. *Let $X_\bullet$ be a semi-simplicial space and write points in the realisation $|X_\bullet|$ as $(x; t)$ where $x \in X_p$ and $t = (t_0, ..., t_p) \in \Delta^p$. Write $\bar{t} = (\bar{t}_0, ..., \bar{t}_p)$ with $\bar{t}_i = t_i / \max_j t_j$, and let $r: |X_\bullet| \to |X_\bullet|$ be the map which in the $\bar{t}$-coordinates replaces all $\bar{t}_i$ by $\max\{0, 2\bar{t}_i - 1\}$.*

(i) *The map r is continuous and homotopic to the identity. If $X_\bullet$ is augmented, the homotopy is fibrewise over $|X_\bullet| \to X_{-1}$.*

For $a \in (0,1)$, let $U_a \subset |X_\bullet|$ be the subset where no $\bar{t}_i$ is equal to a.

(ii) *The subsets $U_a \subset |X_\bullet|$ are open and any infinite set of U_a's form a cover of $|X_\bullet|$.*

Finally, let $r_a: U_a \to \coprod_{q\geqslant 0}(X_q \times \Delta^q)$ be the function defined for $x \in X_p$ and $t \in \Delta^p$ by $r_a(x;t) = (\theta^*(x); \theta^*(t))$, where $\theta = \theta_{t,a}: [q] \to [p]$ is the unique order-preserving monomorphism whose image is $\{i \in [p] : \bar{t}_i > a\}$, and $\theta^*: \Delta^p \to \Delta^q$ is the function given in the $\bar{t}$-coordinates as $(\bar{t}_0, ..., \bar{t}_p) \mapsto (\bar{t}_{\theta(0)}, ..., \bar{t}_{\theta(q)})$.

(iii) *The function $r_a: U_a \to \coprod_{q\geqslant 0}(X_q \times \Delta^q)$ is continuous (where $U_a \subset |X_\bullet|$ is given the subspace topology), and for $a \in \left(0, \frac{1}{2}\right)$ the diagram*

$$\begin{tikzcd}
U_a \arrow[r, "r_a"] \arrow[d, "\mathrm{incl}"'] & \coprod_{q\geqslant 0}(X_q \times \Delta^q) \arrow[r, "\mathrm{quot}"] & {|X_\bullet|} \arrow[d, "r"] \\
{|X_\bullet|} \arrow[rr, "r"'] & & {|X_\bullet|}
\end{tikzcd}$$

is commutative.

Proof. The map r is continuous because it is induced from a continuous map

$$\coprod_{p\geqslant 0}(X_p \times \Delta^p) \longrightarrow \coprod_{p\geqslant 0}(X_p \times \Delta^p).$$

The obvious straight-line homotopy is continuous and fibrewise.

The subset $U_a \subset |X_\bullet|$ is open in the quotient topology because its inverse image in $X_p \times \Delta^p$ is the intersection of the $p+1$ open sets defined by $\bar{t}_i \neq a$. The point $(x;t) \in |X_\bullet|$ will be in U_a except if $a = \bar{t}_i$ for some i, so the set $\{a : (x;t) \notin U_a\}$ is finite.

The commutativity of the diagram is easily verified on the set level, but the continuity of r_a requires an argument. Let us write $m: \coprod_{q\geqslant 0}(X_q \times \Delta^q) \to |X_\bullet|$ for the quotient map. Since $|X_\bullet|$ has the quotient topology from m and $U_a \subset |X_\bullet|$ is open, the subspace topology on $U_a \subset |X_\bullet|$ agrees with the quotient topology from $m^{-1}(U_a) \to U_a$ and therefore it suffices to see that the composition $r_a \circ m: m^{-1}(U_a) \to \coprod_{q\geqslant 0}(X_q \times \Delta^q)$ is continuous. This composition is in fact an idempotent self-map of $m^{-1}(U_a)$ locally given by face maps $X_q \to X_p$ and projection maps in the $\bar{t}$-coordinates, and therefore it is continuous. □

Corollary 6.7. *Let $X_\bullet$ be an augmented topological flag complex with augmentation ε, and let $f_0, f_1: K \to |X_\bullet|$ be two continuous maps with $|\varepsilon| \circ f_0 = |\varepsilon| \circ f_1$. Suppose that for each $x \in K$, the vertices of $f_1(x)$ are orthogonal to the vertices of $f_0(x)$, i.e. that $f_0(x) = (x;s)$ and $f_1(x) = (y;t)$ with $(x,y) \in X_{p+q+1} \subset X_p \times X_q$, $s \in \mathrm{int}(\Delta^p)$ and $t \in \mathrm{int}(\Delta^q)$. Then f_0 and f_1 are homotopic, fibrewise over X_{-1}.*

Proof. For each $x \in K$, there is an obvious straight line in $X_{p+q+1} \times \Delta^{p+q+1}$ which maps to a path in $|X_\bullet|$ from $f_0(x)$ to $f_1(x)$. These assemble to a canonical well-defined function $[0,1] \times K \to |X_\bullet|$, but it is not completely obvious whether this is continuous.

In the case where f_0 and f_1 admit factorisations as

$$K \longrightarrow X_p \times \Delta^p \longrightarrow |X_\bullet| \quad \text{and} \quad K \longrightarrow X_q \times \Delta^q \longrightarrow |X_\bullet|,$$

it is clear that this construction does give a continuous homotopy

$$[0,1] \times K \longrightarrow X_{p+q+1} \times \Delta^{p+q+1} \longrightarrow |X_\bullet|$$

using that $X_{p+q+1} \subset X_p \times X_q$ has the subspace topology. Since continuity may be checked locally, we also get a canonical continuous homotopy when there are factorisations of f_0 and f_1 locally near each point in K. In the general case, Lemma 6.6 implies that $r \circ f_0$ and $r \circ f_1$ admit such local factorisations, so we get homotopies $f_0 \simeq r \circ f_0 \simeq r \circ f_1 \simeq f_1$, fibrewise over X_{-1}. $\square$

Proof of Theorem 6.2. We begin with an element of the relative homotopy group of the pair of spaces $(X_{-1}, |X_\bullet|)$,

$$\begin{tikzcd} \partial D^k \arrow[r, "\hat{f}"] \arrow[d] & {|X_\bullet|} \arrow[d, "|\varepsilon|"] \\ D^k \arrow[r, "f"] \arrow[ur, dashed] & X_{-1}. \end{tikzcd}$$

We will show that there exists a diagonal map making the lower triangle commute and the upper triangle commute up to fibrewise homotopy.

We first explain how to construct a continuous map $D^k \to |X_\bullet|$ making the lower triangle commute, ignoring the upper triangle for the moment. To do this we first pick an infinite set Ω (topologised discretely) and note that it suffices to find open sets $P_n \subset \Omega \times D^k$ together with maps $g_n: P_n \to X_0$ with the properties that the projection $\pi_n: P_n \to D^k$ is surjective, that $\varepsilon \circ g_n = f \circ \pi_n$, and that for all $x \in D^k$ and $n < m$, any $p \in \pi_n^{-1}(x)$ and $q \in \pi_m^{-1}(x)$ have $(g_n(p), g_m(q)) \in X_1$. Namely, given such (P_n, g_n) we can let

$$P = \bigcup_{n \geqslant 0} (\{n\} \times P_n) \subset \mathbb{N} \times \Omega \times D^k$$

and assemble the g_n to a simplicial map $g: N_\bullet P \to X_\bullet$, fitting into the diagram

$$\begin{tikzcd} \partial D^k \arrow[rr, bend left, "\hat{f}"] \arrow[d] & {|N_\bullet P|} \arrow[r, "|g|"'] \arrow[d, "\pi"'] & {|X_\bullet|} \arrow[d, "|\varepsilon|"] \\ D^k \arrow[r, equal] & D^k \arrow[r, "f"] & X_{-1}. \end{tikzcd}$$

By Corollary 6.5, the map $\pi\colon |N_\bullet P| \to D^k$ is a Serre fibration with contractible fibres, so we may pick a section $s\colon D^k \to |N_\bullet P|$. Then the composition $|g| \circ s\colon D^k \to |X_\bullet|$ gives the required map. Afterwards, we shall explain how to improve the construction in order to have $(|g|\circ s)|_{\partial D^k}$ be fibrewise homotopic to $\hat{f}$.

The (P_n, g_n) will be constructed by an inductive procedure, for which it is useful to construct a slightly stricter structure. If we write $\overline{P}_n \subset \Omega \times D^k$ for the closure, we will demand an extension $\bar{g}_n\colon \overline{P}_n \to X_0$ satisfying

(i) the projection $\bar{\pi}_n\colon \overline{P}_n \to D^k$ is proper, and the restriction $\pi_n\colon P_n \to D^k$ is surjective;

(ii) $\varepsilon \circ \bar{g}_n = f \circ \bar{\pi}_n$;

(iii) for all $x \in D^k$ and $n < m$, any $p \in \bar{\pi}_n^{-1}(x)$ and $q \in \bar{\pi}_m^{-1}(x)$ have $(\bar{g}_n(p), \bar{g}_m(q)) \in X_1$.

The properness of $\bar{\pi}_n$ is equivalent to the compactness of $\overline{P}_n$, which in turn is equivalent to the image of $\overline{P}_n$ in Ω being finite. For the construction, we first pick for each $x \in D^k$ an element $g_x(x) \in \varepsilon^{-1}(f(x))$ which is orthogonal to each element of the finite set $\bigcup_{i<n} \bar{g}_i(\bar{\pi}_i^{-1}(x))$, as is possible by assumption. Then, since ε has local lifts of any map from a disc, we can extend to a map $g_x\colon V_x \to X_0$ which is a lift of $D^k \to X_{-1}$, defined on a neighbourhood V_x of x. The maps

$$\bar{g}_i \times g_x\colon \overline{P}_i \times_{D^k} V_x \longrightarrow X_0 \times_{X_{-1}} X_0$$

for $i < n$ all send $\overline{P}_i \times_{D^k} \{x\}$ into the open subset X_1, so by properness of $\bar{\pi}_i$ we can ensure that all these maps have image in X_1, after perhaps shrinking the open set V_x. If we let $U_x \subset V_x$ be a smaller neighbourhood of x with $\overline{U}_x \subset V_x$, then g_x restricts to a continuous map $\overline{U}_x \to X_0$. The sets U_x give an open cover of D^k, and we let $U_{x_1}, ..., U_{x_m}$ be a finite subcover. Finally, we pick distinct $\omega_1, ..., \omega_m \in \Omega$, disjoint from the image of $\bigcup_{i<n} P_i \to \Omega$ and let

$$P_n = \bigcup_{i=1}^{m} (\{\omega_i\} \times U_{x_i}) \subset \Omega \times D^k$$

and define the map $\bar{g}_n\colon \overline{P}_n \to X_0$ by $\bar{g}_n(\omega_i, y) = g_{x_i}(y)$. The sequence of $(P_n, \bar{g}_n)$ thus constructed will satisfy the properties (i)–(iii) above.

The construction of the lift $|g| \circ s\colon D^k \to |X_\bullet|$ so far has not used the given $\hat{f}$ in any way, so we cannot expect $(|g|\circ s)|_{\partial D^k}$ and $\hat{f}$ to be equal or even fibrewise homotopic. We shall add an extra step preceding the above inductive construction of $(\overline{P}_n, \bar{g}_n)$, in order to fix this. The idea is quite simple: For each $x \in \partial D^k$ the point $\hat{f}(x)$ involves only finitely many vertices. If we can arrange that the vertices of $(|g|\circ s)(x)$ are orthogonal to those of $(r \circ \hat{f})(x)$, then Corollary 6.7 provides a fibrewise homotopy $\hat{f} \simeq r \circ \hat{f} \simeq |g| \circ s$, where $r\colon |X_\bullet| \to |X_\bullet|$ is the function from Lemma 6.6. In the notation of that lemma we may cover $|X_\bullet|$ by the open sets U_a, $a \in \left(0, \frac{1}{2}\right)$. Writing $U_a = \coprod_{p \geqslant 0} U_{a,p}$ with $U_{a,p} = r_a^{-1}(X_p \times \Delta^p)$, we may cover ∂D^k by the open sets $\hat{f}^{-1}(U_{a,p})$, each of which has a function $r_a \circ \hat{f}$ to $X_p \times \Delta^p$. Projecting to $X_p \subset X_0^{p+1}$ and taking adjoints give functions

$$g_{a,p}\colon \hat{f}^{-1}(U_{a,p}) \times [p] \longrightarrow X_0,$$

with the property that if $x \in \hat{f}^{-1}(U_{a,p})$, then all the vertices of $(r \circ \hat{f})(x)$ are contained in $g_{a,p}(\{x\} \times [p])$. By compactness we may find a finite open cover $\partial D^k = \bigcup_{i \in I} V_i$ such that for all $i \in I$ there exist a_i and p_i with $\overline{V}_i \subset \hat{f}^{-1}(U_{a_i,p_i})$. We may then let

$$\overline{P}_{-1} = \coprod_{i \in I} (\overline{V}_i \times [p_i]),$$

let $\bar{g}_{-1}\colon \overline{P}_{-1} \to X_0$ be given by the restrictions of the $g_{a,p}$, and let $\bar{\pi}_{-1}\colon \overline{P}_{-1} \to \partial D^k$ be given by the inclusions $\overline{V}_i \subset \partial D^k$. (The resulting $\overline{P}_{-1}$ will not be a subset of $\Omega \times D^k$, but that is unimportant, as long as we remember the maps $\bar{g}_{-1}$ and $\bar{\pi}_{-1}$.) We then construct the $(\overline{P}_n, \bar{g}_n)$ as above, demanding that (iii) hold also for $n=-1$. Proceeding as above with

$$P = \bigcup_{n \geqslant 0} (\{n\} \times P_n) \subset \mathbb{N} \times \Omega \times D^k,$$

the resulting map $|g| \circ s|_{\partial D^k}$ and the map $r \circ \hat{f}$ will satisfy the assumptions of Corollary 6.7: for $x \in \partial D^k$, all vertices of $(r \circ \hat{f})(x)$ will be orthogonal to all vertices of $(s \circ |g|)(x)$. Therefore they are homotopic, fibrewise over X_{-1}. $\square$

6.3. Proof of Theorem 3.4

Recall that this theorem states that the augmentation

$$D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet} \longrightarrow D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_{\bullet}$$

induces a weak homotopy equivalence after geometric realisation, as long as the conditions of Theorem 3.1 are satisfied. In fact, we only require the following weaker set of conditions:

(i) $2\kappa \leqslant d-1$;

(ii) $\kappa+1+d < N+1$;

(iii) L admits a handle decomposition only using handles of index at most $d-\kappa-2$.

We will use Theorem 6.2 to prove that for each p the augmentation map induces a weak equivalence

$$|D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,\bullet}| \longrightarrow D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p.$$

Theorem 6.2 does not apply directly to the augmentation $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,\bullet}\to D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$, but we will show that it does apply after replacing with weakly equivalent spaces.

Recall that an element of $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$ consists of an element

$$(a,\varepsilon,(W,\ell_W))\in D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p,$$

together with an element

$$(\Lambda,\delta,e)\in Z_q(a,\varepsilon,(W,\ell_W)),$$

where $\Lambda\subset\Omega$ is a finite set equipped with a map $\delta\colon\Lambda\to[p]^\vee\times[q]=\{0,...,p+1\}\times\{0,...,q\}$ and e is an embedding $e\colon\Lambda\times\overline{V}\hookrightarrow\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$.

Definition 6.8. The *core* of $\overline{V}$ is the submanifold

$$C=[-2,0]\times D^\kappa\times\{0\}\subset\overline{V}=[-2,0]\times\mathbb{R}^\kappa\times\mathbb{R}^{d-\kappa}.$$

Let $\widetilde{Z}_\bullet(a,\varepsilon,(W,\ell_W))$ be the semi-simplicial space defined as in Definition 3.2 except that instead of demanding that $e\colon\Lambda\times\overline{V}\to\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$ be an embedding, we demand only that it be a smooth map which restricts to an embedding of a neighbourhood of $\Lambda\times C$. We still require that e satisfy the numbered conditions listed in Definition 3.2. Let $\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}\to D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_\bullet$ be the augmented bi-semi-simplicial space defined as in Definition 3.3, but using $\widetilde{Z}_\bullet(x)$ instead of $Z_\bullet(x)$.

PROPOSITION 6.9. *The inclusion* $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}\hookrightarrow\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$ *induces a weak homotopy equivalence in each bidegree, and so on geometric realisation.*

Proof. Choose an isotopy h_t, $t\in[0,\infty)$, from the identity of $\mathbb{R}^{d-\kappa}$ which fixes $\{0\}$ throughout, and is such that $h_t(\mathbb{R}^{d-\kappa})\subset B_{1/t}(0)$ for large t. Similarly, choose an isotopy i_t, $t\in[0,\infty)$, from the identity of $\mathbb{R}^\kappa$ which fixes D^κ throughout, and is such that $i_t(\mathbb{R}^\kappa)\subset B_{1+1/t}(0)$ for large t. Combining these on the last two factors of $\overline{V}=[-2,0]\times\mathbb{R}^\kappa\times\mathbb{R}^{d-\kappa}$, we obtain an isotopy of embeddings $j_t\colon\overline{V}\to\overline{V}$, $t\in[0,\infty)$, such that $j_0=\mathrm{Id}$, $j_t|_C=\mathrm{Id}_C$ for all t and $j_t(\overline{V})$ is contained in the $(1/t)$-neighbourhood of C for large t. It also has the property that every j_t preserves the submanifold $\mathrm{int}(\partial_-D^{\kappa+1})\times\mathbb{R}^{d-\kappa}$ and fixes the height function $h\colon\overline{V}\to[-2,0]$.

Precomposing the embedding $e\colon\Lambda\times\overline{V}\to\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$ with the maps $\mathrm{Id}_\Lambda\times j_t$ induces a deformation

$$[0,\infty)\times\widetilde{Z}_q(a,\varepsilon,(W,\ell_W))\longrightarrow\widetilde{Z}_q(a,\varepsilon,(W,\ell_W))$$

and in turn $[0,\infty)\times\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}\to\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$. Elements of $\widetilde{Z}_q(a,\varepsilon,(W,\ell_W))$ have disjoint cores, so in a compact family $K\to\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$, there exists an $\varepsilon>0$ such that the ε-neighbourhoods of all cores are also disjoint. Composing with the deformation of $\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$, the map from K will eventually deform into $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$. The deformation of $\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$ we have constructed preserves the subspace $D^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,q}$, and hence the relative homotopy groups vanish. □

In order to prove Theorem 3.4, we will show that for each p the map

$$|\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,\bullet}| \longrightarrow D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$$

is a weak homotopy equivalence, by applying Theorem 6.2. Hence we must verify the conditions of that theorem. First we establish condition (i).

Proposition 6.10. *The map $\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,0}\to D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$ has local lifts of any map from a disc.*

Proof. Let $f: D^k\to D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$ be a continuous map, let $x\in D^k$ be a point such that $f(x)=(a,\varepsilon,(W,\ell_W))$, and let $(\Lambda,\delta,e)\in\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$ be the data describing a lift of $f(x)$. Choose $t_0<a_0-\varepsilon_0$ and $t_1>a_p+\varepsilon_p$ which are regular values for the height function $x_1: W\to\mathbb{R}$, and such that $(x_1\circ e)(\Lambda\times\overline{V})\subset(t_0,t_1)$. There is an open neighbourhood $U\subset D^k$ of x such that the t_i remain regular values of the height function on each manifold underlying $f(u)$ for $u\in U$.

The map $U\hookrightarrow D^k\xrightarrow{f} D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$ has graph $\Gamma\subset U\times\mathbb{R}^N$. All fibres of the projection $\pi: \Gamma|_{[t_0,t_1]}\to U$ are diffeomorphic to the same manifold $M=W|_{[t_0,t_1]}$. Sending a point in U to its fibre defines a function

$$\begin{aligned} E: U &\longrightarrow \mathrm{Emb}_\partial(M,[t_0,t_1]\times(-1,1)^N)/\mathrm{Diff}(M),\\ u &\longmapsto \pi^{-1}(u), \end{aligned}$$

where Emb_∂ denotes embeddings which send the boundary to the boundary, and the definition of the topology on $\Psi_\theta(\mathbb{R}\times\mathbb{R}^N)$ makes this continuous (manifolds near to a point $W\in\psi_\theta(N+1,1)\subset\Psi_\theta(\mathbb{R}\times\mathbb{R}^N)$ look like a section of the normal bundle of W inside a compact set, e.g. inside $[t_0,t_1]\times[-1,1]^N$).

We now require two results on spaces of embeddings. Firstly, the map

$$\mathrm{Emb}_\partial(M,[t_0,t_1]\times(-1,1)^N)\longrightarrow \mathrm{Emb}_\partial(M,[t_0,t_1]\times(-1,1)^N)/\mathrm{Diff}(M)$$

is well known to be a principal $\mathrm{Diff}(M)$-bundle, and has local sections (see e.g. [BiFi]). Thus, after perhaps passing to a smaller open neighbourhood, which we will still call U, E has a lift $\widetilde{E}: U\to\mathrm{Emb}_\partial(M,[t_0,t_1]\times(-1,1)^N)$, and we will write $h=\widetilde{E}(x)$.

Secondly, we need the following modification of a technical theorem of Cerf [C, p. 293] (the "first isotopy and extension theorem"), an especially elementary proof of which was given by Lima [Li]. We follow Lima's proof.

Lemma 6.11. *Let* $C \subset [t_0, t_1]$ *be a closed subset and* $S \subset \mathrm{Emb}_\partial(M, [t_0, t_1] \times (-1, 1)^N)$ *be the open subset of those embeddings* e *for which* $\pi_1 \circ e: M \to [t_0, t_1]$ *has no critical values inside* C.

Given an $h \in S$, *there is a neighbourhood* U' *of* h *in* S *and a continuous map*

$$\varphi: U' \longrightarrow \mathrm{Diff}([t_0, t_1] \times (-1, 1)^N)$$

such that $\varphi(g) \circ h$ *and* g *have the same image, and* $\varphi(g)$ *is height-preserving over* C, *for all* $g \in U'$.

Proof. Consider M to be a submanifold of $[t_0, t_1] \times (-1, 1)^N$ via h. We choose a tubular neighbourhood $\pi: T \to M$ of radius ε which over the boundary and $x_1^{-1}(C)$ has fibres contained in level sets of x_1 (this is possible as C is closed and consists of regular values). If $g \in S$ is sufficiently close to h, it will have image in T and we may define an element $\bar{\varphi}(g) \in C^\infty(M, M)$ by

$$\bar{\varphi}(g)(x) = \pi(g(x)).$$

This is a diffeomorphism for $g = h$, and so there is a neighbourhood U'' of h in S where this remains true. We get a function $\bar{\varphi}: U'' \to \mathrm{Diff}(M)$ and for each $g \in U''$ we define a new embedding $G = G(g): M \to [t_0, t_1] \times (-1, 1)^N$ by $G = g \circ (\bar{\varphi}(g)^{-1})$. It has the same image as g and has $\pi(G(x)) = x$. Therefore x and $G(x)$ have the same height when $x \in x_1^{-1}(C)$.

Let λ be a bump function which is 1 on $\left[0, \frac{1}{4}\varepsilon\right)$ and 0 on $\left[\frac{1}{2}\varepsilon, \infty\right)$. Now let

$$\varphi(g)(x) = x + \lambda(|x - \pi(x)|)(G(\pi(x)) - \pi(x))$$

define a smooth self-map $\varphi(g)$ of $[t_0, t_1] \times (-1, 1)^N$, which is the identity outside a compact subset. For $g = h$ it is the identity, and so there is a smaller neighbourhood U' of h in S where it remains a diffeomorphism, since these form an open subset of the smooth maps. We obtain a function $\varphi: U' \to \mathrm{Diff}_c([t_0, t_1] \times (-1, 1)^N)$.

By construction $\varphi(g) \circ h(x) = \varphi(g)(x) = x + (G(x) - x) = G(x)$, so $\varphi(g) \circ h$ has the same image as g. Also, if $x \in x_1^{-1}(C)$ then the vector $G(\pi(x)) - \pi(x)$ has no component in the x_1 direction, so $x_1(\varphi(g)(x)) = x_1(x)$ and $\varphi(g)$ is height-function preserving over C. □

We now continue with the proof of Proposition 6.10. Applying the above lemma with $C = \bigcup_{i=0}^p [a_i - \varepsilon_i, a_i + \varepsilon_i]$, we find that after possibly shrinking U there is a map

$$\varphi: U \longrightarrow \mathrm{Diff}([t_0, t_1] \times (-1, 1)^N)$$

taking values in diffeomorphisms which are height-preserving over C, such that the graph $\Gamma|_{[t_0, t_1]} \subset U \times [t_0, t_1] \times (-1, 1)^N$ is obtained from $W|_{[t_0, t_1]}$ by applying the family of diffeomorphisms φ.

The element $(\Lambda,\delta,e)\in\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$ has surgery data

$$e:\Lambda\times\overline{V}\hookrightarrow[t_0,t_1]\times(0,1)\times(-1,1)^{N-1},$$

so we attempt to define a section $F:U\to\widetilde{D}^{\kappa}_{\theta,L}(\mathbb{R}^N)_{p,0}$ by $u\mapsto(f(u),\Lambda,\delta,\varphi(u)\circ e)$. We must verify that $(\Lambda,\delta,\varphi(u)\circ e)$ is indeed an element of $\widetilde{Z}_0(f(u))$ by checking the conditions of Definition 3.2. Conditions (i)–(iv) hold as φ is height preserving over each $[a_i-\varepsilon_i,a_i+\varepsilon_i]$. Condition (v) holds by construction, as $\varphi(u)$ is a diffeomorphism which carries M into $\pi^{-1}(u)$. The truth or falsity of condition (vi) is locally constant in U, but it holds at the point x so by replacing U with a smaller neighbourhood of x we may ensure that it holds everywhere. Thus we have produced a lift, as required. □

Next, we establish condition (iii) in Theorem 6.2.

PROPOSITION 6.12. *Fix a point* $(a,\varepsilon,(W,\ell_W))\in D^{\kappa-1}_{\theta,L}(\mathbb{R}^N)_p$, *and consider a non-empty collection* $v_1,...,v_k\in\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$ *of pieces of surgery data* (*not necessarily forming a* $(k-1)$-*simplex*). *Then, if* $2\kappa<d$ *and* $\kappa+1+d<N+1$, *there exists a piece of surgery data* $v\in\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$ *such that each* (v_i,v) *is a* 1-*simplex.*

Proof. Each v_j is given by a set Λ^j (which is a subset of the infinite set Ω), a function $\delta^j:\Lambda^j\to[p]^\vee$ and a map $e^j:\Lambda^j\times\overline{V}\to\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$, satisfying certain properties. We first pick a set Λ which is disjoint from all Λ^j and a bijection $\varphi:\Lambda\to\Lambda^1$, let $\delta=\delta^1\circ\varphi:\Lambda\to[p]^\vee$, and then set

$$\tilde{e}=e^1\circ(\varphi\times\mathrm{Id}_{\overline{V}}):\Lambda\times\overline{V}\longrightarrow\mathbb{R}\times(0,1)\times(-1,1)^{N-1}.$$

This gives a new element of $\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$, but it is of course not orthogonal to v_1 (and not necessarily orthogonal to the other v_j). We then perturb $\tilde{e}$ inside the class of functions satisfying the requirements of Definition 3.2, to a new function

$$e:\Lambda\times\overline{V}\longrightarrow\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$$

whose core is in general position with respect to the cores of the v_j. More explicitly, $\tilde{e}$ restricts to a map

$$\Lambda\times\partial_-D^{\kappa+1}\times\mathbb{R}^{d-\kappa}\longrightarrow W,$$

and we first perturb this so that $\Lambda\times\partial_-D^{\kappa+1}\times\{0\}$ is transverse in W to the corresponding part of the other embeddings, and remains disjoint from L, then we extend this perturbation to a map $e:\Lambda\times\overline{V}\to\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$ whose restriction to the interior of C is transverse to the corresponding part of the other embeddings. In the first step we make κ-dimensional manifolds transverse in a d-dimensional manifold, and in the second we make $(\kappa+1)$-dimensional manifolds disjoint in an $(N+1)$-dimensional manifold. As $2\kappa<d$ and $2(\kappa+1)\leqslant\kappa+d+1<N+1$, the new core will actually be disjoint from all other cores, producing the required element $v\in\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$. □

Finally, we establish condition (ii) of Theorem 6.2.

PROPOSITION 6.13. *The set* $\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$ *of surgery data is non-empty as long as* $2\kappa<d$, $\kappa+1+d<N+1$, *and* L *admits a handle decomposition only using handles of index at most* $d-\kappa-2$.

Proof. For each $i=1,...,p$ we consider the pair $(W|_{[a_{i-1},a_i]},W|_{a_i})$. The bordism $W|_{[a_{i-1},a_i]}$ is homotopy equivalent to a relative CW complex $(X',W|_{a_i})$ with finitely many relative cells (one for each critical point of a Morse function, for example). Since the pair is $(\kappa-1)$-connected, we may use the cell trading lemma (cf. [G, Proposition 4.2.1]) to replace any relative cell of dimension $k<\kappa$ by a $(k+2)$-cell, inductively obtaining a homotopy-equivalent relative CW complex $(X,W|_{a_i})$ with the same (finite) total number of cells, each of which now has dimension at least κ. Since $2\kappa<d$, the homotopy equivalence $(X,W|_{a_i})\to(W|_{[a_{i-1},a_i]},W|_{a_i})$ may be assumed to restrict to a smooth embedding of the relative κ-cells. If we pick a subset $\Lambda_{i,0}\subset\Omega$ with one element for each relative κ-cell (choosing disjoint sets for each i), we may therefore pick an embedding

$$\hat{e}_{i,0}\colon\Lambda_{i,0}\times(D^\kappa,\partial D^\kappa)\longrightarrow(W|_{[a_{i-1}+\varepsilon_{i-1},a_i+\varepsilon_i]},W|_{a_i+\varepsilon_i}),$$

which we may assume to be collared on $[a_i-\varepsilon_i,a_i+\varepsilon_i]$, such that the pair

$$(W|_{[a_{i-1},a_i]},W|_{a_i}\cup\mathrm{Im}(\hat{e}_{i,0})|_{[a_{i-1},a_i]})$$

is κ-connected. Furthermore, $\mathbb{R}\times L\subset W$ has a core of dimension at most $d-\kappa-1$, by our assumption on the indices of handles of L, and so we may suppose that the embedding $\hat{e}_{i,0}$ is disjoint from $\mathbb{R}\times L$. As $2\kappa<d$, we may also suppose that the images of the $\hat{e}_{i,0}$ are mutually disjoint.

The embedding

$$\hat{e}_{i,0}|_{\Lambda_{i,0}\times\partial D^\kappa}\colon\Lambda_{i,0}\times\partial D^\kappa\times\{0\}\longrightarrow W|_{a_i+\varepsilon_i}\subset W|_{[a_i+\varepsilon_i,a_{i+1}+\varepsilon_{i+1}]}$$

extends to an embedding of $\Lambda_{i,0}\times\partial D^\kappa\times[0,1]$, where the set $\Lambda_{i,0}\times\partial D^\kappa\times\{1\}$ is sent into $W|_{a_{i+1}+\varepsilon_{i+1}}$ and is collared on the ε-neighbourhoods of both boundaries. This may be seen as follows: to extend $\hat{e}_{i,0}|_{\Lambda_{i,0}\times\partial D^\kappa}$ to a continuous map having this property is possible as $\pi_{\kappa-1}(W|_{[a_i,a_{i+1}]},W|_{a_{i+1}})=0$, but this may then be perturbed to be an embedding as $2\kappa<d$. As above, this may be made disjoint from $\mathbb{R}\times L$, and they can be made mutually disjoint.

We may glue the two embeddings together. Using a suitable diffeomorphism

$$D^\kappa\approx D^\kappa\cup(\partial D^\kappa\times[0,1]),$$

this gives a new embedding of $\Lambda_{i,0}\times D^\kappa$. Continuing in this way, we obtain an extension of $\hat{e}_{i,0}$ to an embedding

$$\tilde{e}_{i,0}\colon \Lambda_{i,0}\times(D^\kappa,\partial D^\kappa)\longrightarrow (W|_{[a_{i-1}+\varepsilon_{i-1},a_p+\varepsilon_p]}, W|_{a_p+\varepsilon_p})$$

which is disjoint from $\mathbb{R}\times L$, and which are mutually disjoint. Identifying D^κ with the disc $\partial_- D^{\kappa+1}\subset[-1,0]\times\mathbb{R}^{\kappa+1}$ gives a height function $D^\kappa\to[-1,0]$ and if we pick the diffeomorphisms $D^\kappa\approx D^\kappa\cup(\partial D^\kappa\times[0,1])$ carefully, we can arrange it so that on each $\tilde{e}_{i,0}^{-1}(W|_{(a_k-\varepsilon_k,a_k+\varepsilon_k)})$, the embedding $\tilde{e}_{i,0}$ is height-function preserving up to an affine transformation.

We now want to extend the $\tilde{e}_{i,0}$ from $\Lambda_{i,0}\times(\partial_- D^{\kappa+1}\times\{0\})\subset\Lambda_{i,0}\times\overline{V}$ to the whole of $\Lambda_{i,0}\times\overline{V}$ so that it satisfies the conditions of Definition 3.2. Since $\kappa+1+d<N+1$, there is no trouble with extending the maps $\tilde{e}_{i,0}$ to disjoint maps $e_{i,0}$ from $\Lambda_{i,0}\times\overline{V}$ to $\mathbb{R}\times(0,1)\times(-1,1)^{N-1}$ satisfying conditions (i)–(v) of Definition 3.2: we first extend each $\tilde{e}_{i,0}$ to an embedding of $[-2,0]\times\mathbb{R}^\kappa\times\{0\}$ (which is possible as $2(\kappa+1)<d+\kappa+1<N+1$), then make this intersect W only in $\partial_- D^{\kappa+1}$ (which is possible as $\kappa+1+d<N+1$), and finally thicken it up by $\mathbb{R}^{d-\kappa}$ (which is possible as $\partial_- D^{\kappa+1}$ and $[-2,0]\times\mathbb{R}^\kappa\times\{0\}$ are both contractible). Property (vi) is ensured by the way we chose $\hat{e}_{i,0}$.

Then, we let $\Lambda=\coprod_{i=0}^{p+1}\Lambda_{i,0}$, $\delta\colon\Lambda\to[p]^\vee$ be given by $\delta(\Lambda_{i,0})=i\in[p]^\vee$, and $e=\coprod_{i=0}^{p+1}e_{i,0}$. The data (Λ,δ,e) thus lies in $\widetilde{Z}_0(a,\varepsilon,(W,\ell_W))$. □

6.4. Proof of Theorem 4.5

We have already proved the first part of this theorem in §6.1. Recall that the second part states that the augmentation map

$$D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}\longrightarrow D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_\bullet$$

induces a weak homotopy equivalence after geometric realisation, as long as the conditions of Theorem 4.1 are satisfied. In fact, we only require the following weaker set of conditions:

(i) $2(l+1)<d$;

(ii) $l\leqslant\kappa$;

(iii) $l+2+d<N+1$;

(iv) L admits a handle decomposition only using handles of index at most $d-l-2$;

(v) the map $\ell_L\colon L\to B$ is $(l+1)$-connected.

We will proceed as in the previous section. Recall that each point of $D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,0}$ lying over $(a,\varepsilon,(W,\ell_W))\in D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_p$ is a tuple (Λ,δ,e,ℓ), where $\Lambda\subset\Omega$ is a subset,

$\delta\colon \Lambda \to [p]\times[0]$ is a function,

$$e\colon \Lambda\times(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1} \hookrightarrow \mathbb{R}\times(0,1)\times(-1,1)^{N-1}$$

is an embedding, and $\ell\colon T(\Lambda\times K|_{(-6,0)})\to\theta^*\gamma$ is a bundle map (where K is defined in §4.2). Let us define

$$C=(-6,-2)\times\{0\}\times D^{l+1}\subset(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1}$$

and call it the *core*. Shrinking in the $\mathbb{R}^{d-l-1}$ direction gives an isotopy from the identity map of $(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1}$ into any neighbourhood of its core.

Definition 6.14. Let $\widetilde{Y}_\bullet(a,\varepsilon,(W,\ell_W))$ be the semi-simplicial space defined as in Definition 4.3, except that we only ask for e to be a smooth map which restricts to an embedding on a neighbourhood of $\Lambda\times C\subset\Lambda\times(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1}$. Note that condition (iv) still makes sense: although the surgery data is no longer disjoint, it is still disjoint when restricted to a small enough neighbourhood of each core.

Let $\widetilde{D}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}\to D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_\bullet$ be the augmented bi-semi-simplicial space defined as in Definition 4.4, but using $\widetilde{Y}_\bullet(a,\varepsilon,(W,\ell_W))$ instead of $Y_\bullet(a,\varepsilon,(W,\ell_W))$.

We have the following analogue of Proposition 6.9, although the proof is slightly more complicated in this case, due to the tangential structures on the surgery data.

PROPOSITION 6.15. *The inclusion* $D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}\hookrightarrow\widetilde{D}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$ *induces a weak homotopy equivalence in each bidegree*, *and so on geometric realisation.*

Proof. This is very similar to Proposition 6.9. We pick an isotopy of maps

$$\varrho_t\colon\mathbb{R}^{d-l-1}\longrightarrow\mathbb{R}^{d-l-1},\quad t\in[0,\infty),$$

which starts at the identity, has $\varrho_t(0)=0$ for all t, and has image in the ball of radius $1/t$ for all t. Applying ϱ_t in the $\mathbb{R}^{d-l-1}$ direction gives an isotopy of self-embeddings of $\Lambda\times(-6,-2)\times\mathbb{R}^{d-l-1}\times D^{l+1}$. Similarly, we can get an isotopy of self-embeddings of the manifold $K|_{(-6,0)}$ from §4.2, which applies ϱ_t in the $\mathbb{R}^{d-l-1}$ direction on $h^{-1}((-6,-2])$, is the identity on $h^{-1}((-\sqrt{2},0))$, and interpolates inbetween. Precomposing with these isotopies gives a homotopy of self-maps of $\widetilde{D}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$, which eventually deforms any compact space into $D^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{\bullet,\bullet}$. $\square$

Therefore it is enough to show that for each p, the augmentation map

$$\widetilde{D}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,\bullet}\longrightarrow D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_p$$

which forgets all surgery data induces a weak homotopy equivalence after geometric realisation, which we do by establishing the conditions of Theorem 6.2. The proofs that conditions (i) and (iii) hold are very similar to the analogous case in §6.3, so we consider those first.

PROPOSITION 6.16. *The map* $\widetilde{D}^{\kappa,l}_{\theta,L}(\mathbb{R}^N)_{p,0}\to D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_p$ *has local lifts of any map from a disc.*

Proof. This is exactly as in the proof of Proposition 6.10. □

PROPOSITION 6.17. *Fix a point* $(a,\varepsilon,(W,\ell_W))\in D^{\kappa,l-1}_{\theta,L}(\mathbb{R}^N)_p$, *and let* $v_1,...,v_k\in\widetilde{Y}_0(a,\varepsilon,(W,\ell_W))$ *be a non-empty collection of pieces of surgery data. Then, if* $2(l+1)<d$ *and* $l+2+d<N+1$, *there exists a piece of surgery data* $v\in\widetilde{Y}_0(a,\varepsilon,(W,\ell_W))$ *such that each* (v_i,v) *is a* 1*-simplex.*

Proof. This is essentially the same as Proposition 6.12: first we let $v=v_1$, then we perturb it to have its cores transverse to the cores of all the v_j. We first do the perturbation on the part of the cores inside W. On the boundary the cores are $(l+1)$-dimensional, so they are disjoint when they are transverse as $2(l+1)<d$. We now make sure the cores intersect W only on their boundary, which is possible as $l+2+d<N+1$. We finally make sure that the cores are also disjoint on their interiors, which is possible as

$$(l+2)+(l+2)\leqslant(l+2)+d<N+1.$$

All these perturbations need to preserve the condition that the embedding be level-preserving near the critical levels (condition (ii) of Definition 4.3), but that is not a problem. □

Finally, we establish condition (ii).

PROPOSITION 6.18. *The set* $\widetilde{Y}_0(a,\varepsilon,(W,\ell_W))$ *is non-empty as long as* $2(l+1)<d$, $l\leqslant\kappa$, $l+2+d<N+1$, L *admits a handle decomposition only using handles of index at most* $d-l-2$, *and the map* $\ell_L\colon L\to B$ *is* $(l+1)$*-connected.*

Proof. Writing $M_i=W|_{a_i}$, we consider the map $\ell\colon M_i\to B$, which by assumption is l-connected. We first argue that ℓ factors as $M_i\subset X\to B$, where (X,M_i) is a relative CW complex with no relative cells of dimension less than $l+1$, finitely many relative cells of dimension $l+1$, as well as (arbitrarily many) higher cells, and $X\to B$ is a weak equivalence. This is a purely homotopy-theoretic question, and we may replace L and M_i by finite CW complexes, and B by a CW complex obtained from L by attaching cells of dimension at least $l+2$. The mapping cylinder of a cellular approximation to ℓ is then a CW complex with finite $(l+1)$-skeleton. M_i is a subcomplex, and as in the proof of Proposition 6.13 the cell-trading lemma implies that the mapping cylinder may be replaced by a relative CW complex (X,M_i) as claimed.

Picking a bijection between the set of relative $(l+1)$-cells of (X,M_i) and a subset $\Lambda_i\subset\Omega$, the attaching maps of the cells assemble to a map

$$g_i\colon\Lambda_i\times S^l\longrightarrow M_i,$$

which may be assumed to be an embedding since $2l<d-1$. We are now in the situation of §4.1 and hence the relative $(l+1)$-skeleton

$$X^{l+1} = M_i \cup_{g_i} (\Lambda_i \times D^{l+1})$$

occurs as the core of the $(l+1)$-handles in the trace of a (multiple) θ-surgery along an embedding

$$f_i|_{a_i}: \Lambda_i \times \{a_i\} \times \mathbb{R}^{d-l-1} \times S^l \hookrightarrow M_i$$

extending g_i. The trace of this surgery is a θ-cobordism C_i from M_i to a θ-manifold $\overline{M}_i$. The relative cells of (X^{l+1}, M_i) are the cores of the handles in C_i, and the structure map $\ell_{C_i}: C_i \to B$ restricts to the $(l+1)$-connected map $X^{l+1} \to B$, so ℓ_{C_i} is also $(l+1)$-connected. The pair $(C_i, \overline{M}_i)$ is $(d-l-2)$-connected and hence $(l+1)$-connected, since we assume $2(l+1)<d$. It follows that the structure map $\overline{M}_i \to B$ is also $(l+1)$-connected and in particular induces an injection in homotopy groups up to degree l. The sets $\Lambda_i \subset \Omega$ may be assumed to be disjoint and the maps $f_i|_{a_i}$ may be assumed to have image disjoint from L, as L only has handles of index less than $d-l-1$. To produce an element of $\widetilde{Y}_0(a, \varepsilon, (W, \ell_W))$ it remains to extend the surgery data $f_i|_{a_i}$ to the rest of the manifold $\Lambda_i \times (a_i - \varepsilon_i, a_p + \varepsilon_p) \times \mathbb{R}^{d-l-1} \times D^{l+1}$.

Since $(l+1)+(d-1)<N$, the map $f_i|_{a_i}$ extends to an embedding

$$e_i|_{a_i}: \Lambda_i \times \{a_i\} \times \mathbb{R}^{d-l-1} \times D^{l+1} \hookrightarrow \{a_i\} \times (0,1) \times (-1,1)^{N-1}$$

which intersects $W|_{a_i}$ precisely on the boundary. This embedding may be further extended to an embedding

$$\Lambda_i \times (a_i - \varepsilon_i, a_i + \varepsilon_i) \times \mathbb{R}^{d-l-1} \times D^{l+1} \hookrightarrow (a_i - \varepsilon_i, a_i + \varepsilon_i) \times (0,1) \times (-1,1)^{N-1}$$

using just the cylindrical structure of W over $(a_i - \varepsilon_i, a_i + \varepsilon_i)$, but we wish to extend it to an embedding of $\Lambda_i \times (a_i - \varepsilon_i, a_p + \varepsilon_p) \times \mathbb{R}^{d-l-1} \times D^{l+1}$, which is cylindrical over each $(a_j - \varepsilon_j, a_j + \varepsilon_j)$ and intersects W precisely on the boundary. We will do this by extending it step-by-step over each interval $[a_j, a_{j+1}]$: if it is defined up to a_j we have an embedding

$$e_i|_{a_j}: \Lambda_i \times \{a_j\} \times \mathbb{R}^{d-l-1} \times D^{l+1} \hookrightarrow \{a_j\} \times (0,1) \times (-1,1)^{N-1},$$

and as the pair $(W|_{[a_j, a_{j+1}]}, W|_{a_{j+1}})$ is κ-connected and $l \leqslant \kappa$, on the boundary this extends to a smooth map

$$f_i|_{[a_j, a_{j+1}]}: \Lambda_i \times [a_j, a_{j+1}] \times \mathbb{R}^{d-l-1} \times S^l \longrightarrow W|_{[a_j, a_{j+1}]}$$

sending $\Lambda_i \times \{a_{j+1}\} \times \mathbb{R}^{d-l-1} \times S^l$ to $W|_{a_{j+1}}$. The core $\Lambda_i \times [a_j, a_{j+1}] \times \{0\} \times S^l$ has dimension $l+1 < \frac{1}{2}d$, so by a general-position argument $f_i|_{[a_j, a_{j+1}]}$ can be isotoped to an embedding. Similarly, we may arrange that it is cylindrical near the necessary ε-neighbourhoods of the ends and has image disjoint from $[a_j, a_{j+1}] \times L$. Finally, as $l+2+d < N+1$ we may extend $f_i|_{[a_j, a_{j+1}]}$ to an embedding

$$e_i|_{[a_j,a_{j+1}]}\colon \Lambda_i \times [a_j, a_{j+1}] \times \mathbb{R}^{d-l-1} \times D^{l+1} \hookrightarrow [a_j, a_{j+1}] \times (0,1) \times (-1,1)^{N-1}$$

which is cylindrical over each $(a_j - \varepsilon_j, a_j + \varepsilon_j)$ and intersects W precisely on the boundary. In total we obtain an embedding

$$e_i\colon \Lambda_i \times (a_i - \varepsilon_i, a_p + \varepsilon_p) \times \mathbb{R}^{d-l-1} \times D^{l+1} \hookrightarrow (a_i - \varepsilon_i, a_p + \varepsilon_p) \times (0,1) \times (-1,1)^{N-1}$$

which is cylindrical over each $(a_j - \varepsilon_j, a_j + \varepsilon_j)$ and intersects W precisely on the boundary. Furthermore, by doing the above in increasing order of i, we can ensure that the different e_i have disjoint cores: while constructing e_i we make sure that its core stays disjoint from those of the e_j for all $j < i$, which is possible as $2(l+1) < d$ and $2(l+2) < N+1$.

We let $\Lambda = \coprod_{i=0}^p \Lambda_i$, $\delta\colon \Lambda \to [p] \times [0]$ be given by $\delta(\Lambda_i) = (i, 0)$, and e be given by $\coprod_{i=0}^p e_i$, reparameterised using the $\varphi(a_i, \varepsilon_i, a_p, \varepsilon_p)$. Then the data (Λ, δ, e) gives the embedding part of the data of an element of $\widetilde{Y}_0(a, \varepsilon, (W, \ell_W))$, and the bundle part comes from the θ-structures ℓ_{C_i} above. $\square$

6.5. Proof of Theorem 5.14

Recall that the statement of the theorem is as follows. We work in dimension $2n$, and fix a tangential structure θ which is reversible (cf. Definition 5.2), a $(2n-1)$-manifold with boundary L equipped with θ-structure, and a collection $\mathcal{A} \subset \pi_0(\mathrm{Ob}(\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)))$ of objects. This allows us to define the augmented bi-semi-simplicial space

$$D_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)_{\bullet,\bullet} \longrightarrow D_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)_{\bullet}$$

of surgery data, and the second part of Theorem 5.14 states that if the conditions of Theorem 5.3 are satisfied, then the induced map on geometric realisation is a weak homotopy equivalence. (We have already proved the first part of Theorem 5.14 in §6.1.) We recall that these conditions are

(i) $2n \geqslant 6$;
(ii) $3n < N$;
(iii) θ is reversible;
(iv) L admits a handle decomposition only using handles of index at most $n-1$;
(v) $\ell_L\colon L \to B$ is $(n-1)$-connected;
(vi) the natural map $\mathcal{A} \to \pi_0(B\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N))$ is surjective.

Note that the penultimate condition implies that for any object M the map $M \to B$ induced by the tangential structure induces a surjection on homotopy groups in degrees less than n.

In many respects the proof of this theorem is very similar to what we did in §6.4, but in that section we often used the inequality $2(l+1)<d$ so that pairs of transverse $(l+1)$-dimensional submanifolds of a d-manifold are automatically disjoint. In Theorem 5.14, $d=2n$ and the analogue of l is $n-1$ so this observation fails. Instead, we will use a version of the Whitney trick to separate n-dimensional submanifolds of our $2n$-manifolds; this accounts for the restriction $2n\geqslant 6$ in the statement of the theorem.

We proceed precisely as in Definition 6.14 by for $(a,\varepsilon,(W,\ell_W))\in D_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)_p$ letting $\widetilde{Y}_\bullet(a,\varepsilon,(W,\ell_W))$ be the analogue of $Y_\bullet(a,\varepsilon,(W,\ell_W))$ from Definition 5.13, where instead of asking that e be an embedding, we only ask for it to be a smooth map which restricts to an embedding on a neighbourhood of $\Lambda\times C$. We use this to define the bi-semi-simplicial space $\widetilde{D}_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)_{\bullet,\bullet}$, and by the same argument as in the proof of Proposition 6.15, the inclusion

$$D_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)_{\bullet,\bullet} \hookrightarrow \widetilde{D}_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)_{\bullet,\bullet}$$

is a weak homotopy equivalence in each bidegree. We are now left to verify the conditions of Theorem 6.2 for the augmented semi-simplicial spaces

$$\widetilde{D}_{\theta,L}^{n-1,\mathcal{A}}(\mathbb{R}^N)_{p,\bullet} \longrightarrow D_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)_p.$$

That the map on 0-simplices has local lifts of any map from a disc is proved as in the previous two sections.

Proposition 6.19. *Fix a point $(a,\varepsilon,(W,\ell_W))\in D_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)_p$, and consider a non-empty collection $v_1,...,v_k\in\widetilde{Y}_0(a,\varepsilon,(W,\ell_W))$ of pieces of surgery data. Then if $2n\geqslant 6$ and $3n<N$ there exists a piece of surgery data $v\in\widetilde{Y}_0(a,\varepsilon,(W,\ell_W))$ such that each (v_i,v) is a 1-simplex.*

Proof. Let us write $v_j=(\Lambda^j,\delta^j,e^j,\ell^j)$. First we let $v=v_1$, then we perturb it to a nearby embedding which has its cores transverse to the cores of all the v_j. We first do the perturbation on the part of the cores inside W. On the boundary the cores are n-dimensional, so when they are transverse they intersect in a finite set of points. We now make sure the cores intersect W only on their boundary, which is possible as $(n+1)+2n<N+1$. We finally make sure that the cores are also disjoint on their interiors, which is possible as $2(n+1)<N+1$.

We are left with surgery data v whose core is disjoint from the cores of v_j away from W, and on W intersects the other cores transversely. It has a finite number of transverse intersections with all the other cores in W, so it is enough to give a procedure which reduces the number of intersections by 1. We will use the following technique similar to the Whitney trick, which for later use we call the *half Whitney trick* (similar to the *piping* of [Wa, p. 40]). Let x be such an intersection point, between v and some v_j. More precisely, suppose it is a point of intersection of the cylinders

$$e_i(\Lambda_i \times (-6,-2) \times \{0\} \times S^{n-1}) \quad \text{and} \quad e_k^j(\Lambda_k^j \times (-6,-2) \times \{0\} \times S^{n-1}).$$

Claim 6.20. Let $T \subset \mathbb{R}^2$ denote the triangle $\{(x,y) : |x|-1 \leqslant y \leqslant 0\}$ and U be a small open neighbourhood of it, e.g. defined by $|x|-1-\varepsilon < y < \varepsilon$. There is a *half Whitney disc* $w: U \hookrightarrow W$ such that

(i) w is disjoint from $\mathbb{R} \times L$;

(ii) $w|_{[-1,1]\times\{0\}}$ is a path in $W|_{a_p}$ which on its interior is disjoint from all the cores;

(iii) the inverse image of the first cylinder is the line on ∂T from $(0,-1)$ to $(-1,0)$; the inverse image of the second cylinder is the line from $(0,-1)$ to $(1,0)$;

(iv) the height functions $x_1 \circ w$ and $y: T \to \mathbb{R}$ agree up to an affine transformation inside each $(x_1 \circ w)^{-1}(a_j - \varepsilon_j, a_j + \varepsilon_j)$.

Given such a disc, we can extend it to a standard neighbourhood

$$w(U) \times \mathbb{R}^{n-1} \times \mathbb{R}^{n-1} \subset W$$

as in the proof of [Mi2, Theorem 6.6]. Note that the argument is easier in this case as we are cancelling intersection points against the boundary instead of against each other, and so no framing problems arise. We can further extend this to a neighbourhood

$$w(U) \times \mathbb{R}^{n-1} \times \mathbb{R}^{n-1} \times \mathbb{R}^{N+1-2n} \subset \mathbb{R} \times (0,1) \times \mathbb{R}^{N-1}.$$

There is a compactly supported vector field on U which is $\partial/\partial x$ on T, and we extend it using bump functions in the euclidean directions to this open subset of $\mathbb{R} \times (0,1) \times \mathbb{R}^{N-1}$. The flow associated with this vector field gives a one-parameter family of diffeomorphisms φ_t, and flowing e_i using φ_t will eventually lead to a new e_i whose core has one fewer intersection point with the other cores (at least inside $W|_{(a_0-\varepsilon_0, a_p]}$, but we can then use the cylindrical structure of $W|_{(a_p-\varepsilon_p, a_p+\varepsilon_p)}$ to remove any intersections above a_p). It will still satisfy condition (ii) of Definition 5.13 by property (iv) above. If the vector field is chosen carefully (i.e. in the kernel of dx_1 near inverse images of a_i), it will also satisfy (ii), and the other conditions are clear.

It remains to prove the claim. If it were not for property (iv) the argument is clear: choose an embedded path from x in each cylinder up to $W|_{a_p}$. Together these give an element of $\pi_1(W|_{[a_i,a_p]}, W|_{a_p})$ which is 0 as $\kappa=n-1\geqslant 2$, and so this extends to a continuous map $w|_T: T\to W$ which gives these two paths along the lower part of its boundary and lies in $W|_{a_p}$ in the top part of its boundary. As $2\cdot 2<2n$, this map may be perturbed to be an embedding into W, still enjoying these two properties. Finally, as $2+n<2n$, $w|_T$ can be made disjoint from the other cores on its interior. This may now be extended to a map on U, enjoying properties (ii) and (iii).

To obtain property (iv) as well, we instead build up the embedding $w|_T$ in pieces inside each $W|_{[a_j,a_{j+1}]}$, which is possible as each $\pi_1(W|_{[a_j,a_{j+1}]}, W|_{a_{j+1}})$ is 0. □

In the proof of Proposition 6.18, it was easy to see that for an object $M\in\mathcal{C}_{\theta,L}^{\kappa,l-1}(\mathbb{R}^N)$ there exists a piece of θ-surgery data $e:\Lambda\times\mathbb{R}^{d-l-1}\times S^l\hookrightarrow M$ such that the resulting manifold $\overline{M}$ has $\pi_l(\overline{M})\to\pi_l(B)$ injective, and thus satisfies condition (iv) of Definition 4.3. In the present situation we have $M\in\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$ and require surgery data so that $\overline{M}\in\mathcal{A}$, to satisfy condition (v) of Definition 5.13. This is rather more difficult, and we first describe how to accomplish this step.

Lemma 6.21. *Let $M\in\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$ be an object, and suppose that θ is reversible, $2n\geqslant 6$, $3n<N$, L has a handle structure with only handles of index at most $n-1$, the map $\ell_L: L\to B$ is $(n-1)$-connected, and $\mathcal{A}$ contains an object in the same path component of $B\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$ as M. Then there is a piece of θ-surgery data, given by an embedding $e:\Lambda\times\mathbb{R}^{d-l-1}\times S^l\hookrightarrow M$ disjoint from L and a compatible bundle map*

$$T(\Lambda\times\mathbb{R}^{d-l-1}\times D^{l+1})\longrightarrow\theta^*\gamma,$$

such that the resulting surgered manifold $\overline{M}$ lies in $\mathcal{A}$.

Proof. Part of this proof is very similar to [Kr, pp. 722–724]. We shall assume in this proof that B is path connected, the general case then follows by considering one path component at a time. Then all objects and morphisms of $\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$ are path-connected manifolds, so relative connectivity may be verified using relative homotopy groups with any basepoint.

We first claim that if there is a morphism $W: M_0\rightsquigarrow M_1\in\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$, then there is another, W' say, with the property that (W', M_0) is also $(n-1)$-connected. (By definition, (W', M_1) is $(n-1)$-connected.) In fact, we claim that it is possible to do surgery along a finite set of embeddings of $S^{n-1}\times D^{n+1}$ into the interior of W (and disjoint from L), such that the resulting cobordism W' is $(n-1)$-connected with respect to both boundaries. Let us first point out that doing any such $(n-1)$-surgery does not change the property that $\pi_k(W, M_1)=0$ for $k\leqslant n-1$: up to homotopy it amounts to cutting out a manifold of codimension $n+1$ and then attaching cells of dimension n and $2n$. We have assumed that $L\to B$ is $(n-1)$-connected so $\pi_k(M_0)\to\pi_k(B)$ is surjective for

$k \leqslant n-1$. Since $M_0 \in \mathcal{C}_{\theta,L}^{n-1,n-2}$, it is an isomorphism for $k \leqslant n-2$ and similarly for M_1. As $\pi_k(M_1) \to \pi_k(W)$ is an isomorphism for $k \leqslant n-2$, we conclude that $\pi_k(M_0) \to \pi_k(W)$ is an isomorphism for $k \leqslant n-2$, but it need not be surjective for $k=n-1$. In fact, the long exact sequence in homotopy groups identifies the cokernel with $\pi_{n-1}(W, M_0) \cong H_{n-1}(\widetilde{W}, \widetilde{M}_0)$. The fact that $\pi_{n-1}(M_0) \to \pi_{n-1}(B)$ is surjective implies that the composition

$$\operatorname{Ker}(\pi_{n-1}(W) \to \pi_{n-1}(B)) \longrightarrow \pi_{n-1}(W) \longrightarrow \pi_{n-1}(W, M_0)$$

is still surjective. By the Hurewicz theorem, $\pi_{n-1}(W, M_0) \cong H_{n-1}(\widetilde{W}, \widetilde{M}_0)$ is finitely generated as a module over π_1, and we have proved that there exist finitely many elements

$$\alpha_i \in \operatorname{Ker}(\pi_{n-1}(W) \to \pi_{n-1}(B))$$

which generate the cokernel of $\pi_{n-1}(M_0) \to \pi_{n-1}(W)$. These elements may be represented by disjoint embedded spheres in the interior of W, and as L has a handle structure with only handles of index at most $n-1$ they can be made disjoint from L. By Proposition 5.8 these embedded spheres can be framed so that the result W' of performing surgery on them has a θ-structure which agrees with the old one on the boundary and on L. Both pairs (W', M_0) and (W', M_1) are now $(n-1)$-connected, so this gives the required cobordism.

We now return to the proof of the lemma. There is a zig-zag of morphisms in the category $\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N)$ from M to an object of $\mathcal{A}$, as $\mathcal{A}$ was assumed to hit the path component of M. By the above discussion we can suppose that it is a zig-zag of θ-cobordisms which are $(n-1)$-connected relative to both ends. Then, by reversibility, we can reverse the backwards-pointing arrows and obtain a single morphism

$$(C, \ell_C)\colon (M, \ell_M) \rightsquigarrow (A, \ell_A) \in \mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N),$$

which is $(n-1)$-connected relative to both ends, so $A \in \mathcal{A}$ and $\pi_*(C, A) = \pi_*(C, M) = 0$ for $* \leqslant n-1$.

If such a cobordism C admits a Morse function with only critical points of index n, then the descending manifolds of the critical points, and ℓ_C restricted to them, gives the required θ-surgery data. It remains to produce such a Morse function.

If $\pi_1(L)=0$ then all of the manifolds appearing above are also simply connected, and we deduce by Poincaré duality and the universal coefficient theorem that $H_*(C, M)$ is concentrated in degree n and is free abelian. We can choose a self-indexing Morse function on C and as in the proof of the h-cobordism theorem we can first modify it to have no critical points of index 0 or 1 [Mi2, Theorem 8.1], do the same to the negative of the Morse function to remove critical points of index $2n$ and $2n-1$, and finally by the basis theorem [Mi2, Theorem 7.6] we can diagonalise the differentials in the Morse homology complex, and so modify the Morse function to only have critical points of index n.

When $\pi_1(L)\neq 0$ we must go to a little more trouble, and use techniques from the proof of the s-cobordism theorem. As these are less well known, we go into more detail, but recommend [Lü] and [Ke] for details of that argument. As above, pick a self-indexing Morse function on C and let us write

$$\pi=\pi_1(L)=\pi_1(M)=\pi_1(C)=\pi_1(A)$$

for the common fundamental group, and $\mathbb{Z}[\pi]$ for its integral group ring.

When $M\hookrightarrow C$ is 1-connected, [Mi2, Theorem 8.1] is still true: we may modify the Morse function to have no critical points of index 0 or 1, and as above do the same on the opposite Morse function to eliminate critical points of index $2n$ and $2n-1$. The cores of the handles given by this Morse function on the universal cover give a cell complex with cellular chain complex $C_*(\widetilde{C},\widetilde{M})$, and $C_*(\widetilde{C},\tilde{A})$ for the opposite Morse function. These are chain complexes of based free $\mathbb{Z}[\pi]$-modules, and geometric Poincaré duality gives an isomorphism

$$C_*(\widetilde{C},\widetilde{M})\cong \mathrm{Hom}_{\mathbb{Z}[\pi]}(C_{2n-*}(\widetilde{C},\tilde{A}),\mathbb{Z}[\pi])$$

of chain complexes, by sending basis elements to their "dual" basis elements (we use the convention of [Wa, Chapter 2] to interchange right and left $\mathbb{Z}[\pi]$-module structures; when C is not orientable, this involves the orientation character).

We calculate $0=\pi_*(C,A)=\pi_*(\widetilde{C},\tilde{A})=H_*(\widetilde{C},\tilde{A};\mathbb{Z})$ for $*\leqslant n-1$. As the chain complex $C_*(\widetilde{C},\tilde{A})$ is one of free $\mathbb{Z}[\pi]$-modules we have a universal coefficient spectral sequence [Le, Theorem 2.3],

$$\mathrm{Ext}^q_{\mathbb{Z}[\pi]}(H_p(\widetilde{C},\tilde{A}),\mathbb{Z}[\pi])\Rightarrow H^{p+q}(\mathrm{Hom}_{\mathbb{Z}[\pi]}(C_*(\widetilde{C},\tilde{A}),\mathbb{Z}[\pi]))\cong H_{2n-p-q}(\widetilde{C},\widetilde{M};\mathbb{Z}),$$

and so the chain complex $C_*(\widetilde{C},\widetilde{M})$ is acyclic in degrees $*\geqslant n+1$. Furthermore, we also have $0=\pi_*(C,M)=\pi_*(\widetilde{C},\widetilde{M})=H_*(\widetilde{C},\widetilde{M};\mathbb{Z})$ for $*\leqslant n-1$, so the homology of $C_*(\widetilde{C},\widetilde{M})$ is concentrated in degree n. By the usual modification technique, we can use handle exchanges to modify the Morse function to only have critical points of index n and $n-1$. We are left with a short exact sequence of $\mathbb{Z}[\pi]$-modules

$$0\longrightarrow H_n(\widetilde{C},\widetilde{M};\mathbb{Z})\longrightarrow C_n(\widetilde{C},\widetilde{M})\xrightarrow{\partial_n} C_{n-1}(\widetilde{C},\widetilde{M})\longrightarrow 0.$$

The rightmost term is a free $\mathbb{Z}[\pi]$-module and so this sequence is split: in particular, $H_n(\widetilde{C},\widetilde{M};\mathbb{Z})$ is stably free as a $\mathbb{Z}[\pi]$-module. If $H_n(\widetilde{C},\widetilde{M};\mathbb{Z})$ is not actually free as a $\mathbb{Z}[\pi]$-module, there cannot exist a Morse function on C with only critical points of index n. In this case we replace C by $C\#g(S^n\times S^n)$ for g sufficiently large (this manifold admits a θ-structure: $S^n\times D^n$ is parallelisable, so admits a θ-structure, and hence by reversibility its double $S^n\times S^n$ does too; the connected sum is then formed by applying Proposition 5.8 to the embedding $S^0\to C\amalg(S^n\times S^n)$ sending one point to C and one to $S^n\times S^n$, which satisfies the hypotheses as B is path connected). This has the effect of adding on a large free $\mathbb{Z}[\pi]$-module to $H_n(\widetilde{C},\widetilde{M};\mathbb{Z})$, so we may assume that this homology group is now free, and pick a basis for it.

Choosing a splitting of the short exact sequence above, we obtain an isomorphism

$$C_n(\widetilde{C},\widetilde{M})\cong H_n(\widetilde{C},\widetilde{M};\mathbb{Z})\oplus C_{n-1}(\widetilde{C},\widetilde{M}) \tag{6.1}$$

of based free $\mathbb{Z}[\pi]$-modules, and so an element of $K_1(\mathbb{Z}[\pi])$. However, the basis we chose for $H_n(\widetilde{C},\widetilde{M};\mathbb{Z})$ was not geometrically meaningful and we are free to change it. After possibly stabilising C further, it is possible to choose a basis for which (6.1) represents the zero class in $K_1(\mathbb{Z}[\pi])$, and hence in the Whitehead group $\mathrm{Wh}(\pi)$ too (note that further stabilisation may even be required if C is an h-cobordism). We may then use the modification lemma to rearrange the index-n critical points of the Morse function so that $\partial_n: C_n(\widetilde{C},\widetilde{M})\to C_{n-1}(\widetilde{C},\widetilde{M})$ is simply projection onto the first few basis elements: this allows us to cancel all the critical points of index $n-1$. ☐

Proposition 6.22. *$\widetilde{Y}_0(a,\varepsilon,(W,\ell_W))$ is non-empty as long as $3n<N$, $2n\geqslant 6$, θ is reversible, L admits a handle structure with only handles of index at most $n-1$, $\ell_L: L\to B$ is $(n-1)$-connected, and the natural map $\mathcal{A}\to\pi_0(B\mathcal{C}_{\theta,L}^{n-1,n-2}(\mathbb{R}^N))$ is surjective.*

Proof. Let $d=2n$ and $l=n-1$. We follow the proof of Proposition 6.18, with a few changes. As in that proof, the first step is to produce for each $W|_{a_i}$ the θ-surgery data $f_i|_{a_i}$. The method described in that proposition no longer works, and we use Lemma 6.21 to produce the necessary data instead. From this point up to constructing the maps $e_i|_{(a_i-\varepsilon_i,a_i+\varepsilon_i)}$ there is no difference, and the argument given in the proof of Proposition 6.18 goes through.

It remains to explain how given an embedding $e_i|_{a_j}$ we can extend it to $e_i|_{[a_j,a_{j+1}]}$. We proceed in the same way: we have the embedding

$$f_i|_{a_j}:\Lambda_i\times\{a_j\}\times\mathbb{R}^n\times S^{n-1}\hookrightarrow W|_{a_j}$$

disjoint from L, which extends to a continuous map

$$f_i|_{[a_j,a_{j+1}]}:\Lambda_i\times[a_j,a_{j+1}]\times\mathbb{R}^n\times S^{n-1}\longrightarrow W|_{[a_j,a_{j+1}]}$$

as $(W|_{[a_j,a_{j+1}]},W|_{a_{j+1}})$ is $(n-1)$-connected by assumption. We can again make this be a self-transverse immersion of the core, but this no longer implies that the core is embedded: it will have isolated points of self-intersection. As $2n\geqslant 6$ we can remove these using the half Whitney trick, as in the proof of Proposition 6.19. The core may still intersect the core of $[a_j,a_{j+1}]\times L$, as they are both of dimension n inside a $2n$-manifold, but we can again use the half Whitney trick to separate them. Given $f_i|_{[a_j,a_{j+1}]}$ which is an embedding of the core and whose core is disjoint from that of $[a_j,a_{j+1}]\times L$, we can shrink in the $\mathbb{R}^n$ direction and isotope it to get an embedding disjoint from $[a_j,a_{j+1}]\times L$, and then extend this to $e_i|_{[a_j,a_{j+1}]}$ as in the proof of Proposition 6.18.

This gives the required embeddings e_i, which are then combined as in the proof of Proposition 6.18 to get (Λ,δ,e), the embedding part of the data of an element of $\widetilde{Y}_0(a,\varepsilon,(W,\ell_W))$. The remaining bundle part of the data consists of an extendible (cf. Definition 5.11) θ-structure ℓ on $\Lambda\times K$ which agrees with $\ell_W\circ D(\partial e)$ on $\Lambda\times K|_{(-6,-2)}$, and is such that the effect of the θ-surgery described by this data (i.e. the restriction of ℓ to $K|_{(-6,0]}$) lies in $\mathcal{A}$. We will describe a construction which for each $\lambda\in\Lambda$ produces a θ-structure ℓ_λ on $K\subset\mathbb{R}^{n+1}\times\mathbb{R}^n$; these are then combined in the obvious way.

Firstly, there is a unique θ-structure on the subspace

$$K|_{(-6,-2)}=((-6,-2)\times\mathbb{R}^n)\times S^{n-1}, \tag{6.2}$$

such that the embedding ∂e preserves θ-structures (i.e. satisfies requirement (iii) of Definition 5.13). Secondly, the manifold $K|_{(-6,0]}\subset\mathbb{R}^{n+1}\times\mathbb{R}^n$ is obtained from $K|_{(-6,-2)}$ by attaching an n-handle. To extend the θ-structure over this n-handle requires a null-homotopy of the map $S^{n-1}\to B$ induced by the θ-structure on (6.2), and this is provided as part of the θ-surgery data in Lemma 6.21. Finally, we need to prove that this structure extends to a θ-structure over all of K, which is furthermore extendible. To see this, let us write $X\subset\mathbb{R}^{n+1}\times\mathbb{R}^n$ for the union

$$X=K\cup(\mathbb{R}^{n+1}\times S^{n-1}).$$

This subset is not a manifold, but the union (taken inside $T(\mathbb{R}^{n+1}\times\mathbb{R}^n)$) of the vector bundles TK and $T(\mathbb{R}^{n+1}\times S^{n-1})$ is a vector bundle over X which we shall denote TX, and it suffices to extend the bundle map $T(K|_{(-6,0]})\to\theta^*\gamma$ to a bundle map $TX\to\theta^*\gamma$: the restriction to TK will then be extendible. But TK is naturally identified with a subbundle of the trivial bundle $T(\mathbb{R}^{n+1}\times\mathbb{R}^n)|_X$ with trivial one-dimensional complement and therefore we get a trivialisation $\varepsilon^1\oplus TX\cong\varepsilon^{2n+1}$. Since $K|_{(-6,0]}$ is contractible, there is then no problem with extending the stabilised bundle map $\varepsilon^1\oplus T(K|_{(-6,0]})\to\varepsilon^1\oplus\theta^*\gamma$ to a bundle map $\varepsilon^1\oplus TX\to\varepsilon^1\oplus\theta^*\gamma$. By Lemma 5.5, we may therefore also extend the unstabilised $T(K|_{(-6,0]})\to\theta^*\gamma$ to a bundle map $TX\to\theta^*\gamma$ as desired. (The argument in Lemma 5.5 still applies even though X is not a manifold; it is still a CW complex of dimension $d=2n$ which is all we used in the proof.) $\square$

7. Proofs of the main theorems

In this section, we use the results of §§3–6 to prove the theorems stated in §1. As explained in Remark 1.11, Theorem 1.2 follows from Theorem 1.8, which we prove in full detail in §§7.2–7.5 below. Nevertheless, we shall first outline in some detail how to deduce Theorem 1.2 directly from the results in §§3–6, in the hope of putting the general case in a useful perspective. As in [GRW1], the parameterised surgery results of the previous sections allow us to prove a general theorem about the direct limit of moduli spaces of manifolds, independent of homological stability results.

In the following we shall work entirely in even dimension $d=2n>4$ and always set $N=\infty$. (Suitably interpreted, all results hold for sufficiently large finite N, but we shall not pursue this here.)

7.1. Outline of the proof of Theorem 1.2

To apply the theorems in §3 and §4, we must specify a structure $\theta\colon B\to B\mathrm{O}(2n)$ and a $(2n-1)$-dimensional manifold L with θ-structure $\ell_L\colon\varepsilon^1\oplus TL\to\theta^*\gamma$. For the purpose of deducing Theorem 1.2, we let $\theta=\theta^n\colon B\mathrm{O}(2n)\langle n\rangle\to B\mathrm{O}(2n)$ be the n-connective cover, and let $L\subset(-1,0]\times\mathbb{R}^N$ be a $(2n-1)$-manifold with collared boundary, diffeomorphic to D^{2n-1}. Now, the inclusion functors induce weak equivalences

$$B\mathcal{C}_{\theta^n,L}^{n-1,n-2}\simeq B\mathcal{C}_{\theta^n,L}^{n-1}\simeq B\mathcal{C}_{\theta^n,L}\simeq\psi_{\theta^n,L}(\infty,1)\simeq\psi_{\theta^n}(\infty,1)\simeq\Omega^{\infty-1}\,\mathrm{MT}\theta^n,$$

obtained by applying Theorem 4.1 $n-1$ times, Theorem 3.1 n times, Proposition 2.15 and Proposition 2.16, respectively, composed with the weak homotopy equivalence

$$\psi_{\theta^n}(\infty,1)\simeq\Omega^{\infty-1}\,\mathrm{MT}\theta^n$$

from [GRW1, Theorem 3.12].

To apply the result of §5, we must specify a subset $\mathcal{A}\subset\pi_0(\mathrm{Ob}(\mathcal{C}_{\theta^n,L}^{n-1,n-2}))$. There is a unique path component of $\mathrm{Ob}(\mathcal{C}_{\theta^n,L}^{n-1,n-2})$ consisting of manifolds diffeomorphic to S^{2n-1} (with its standard smooth structure). Letting $\mathcal{A}$ consist of this path component, $\mathcal{C}_{\theta^n,L}^{n-1,\mathcal{A}}$ is the full subcategory of $\mathcal{C}_{\theta^n,L}^{n-1,n-2}$ on the objects in $\mathcal{A}$. It is clear that θ^n is spherical and hence reversible (cf. Proposition 5.7), that $L\cong D^{2n-1}$ admits a handle decomposition using only handles of index less than n (since a single 0-handle suffices), and that the map $\ell_L\colon D^{2n-1}\to B\mathrm{O}(2n)\langle n\rangle$ is $(n-1)$-connected (it is even n-connected). Theorem 5.3 would give the weak equivalence $B\mathcal{C}_{\theta^n,L}^{n-1,\mathcal{A}}\simeq B\mathcal{C}_{\theta^n,L}^{n-1,n-2}$, except that that theorem requires $\mathcal{A}$ to contain at least one object from each path component of $B\mathcal{C}_{\theta^n,L}^{n-1,n-2}$, which may not hold here. Therefore we let $\bar{\mathcal{A}}\subset\pi_0(\mathrm{Ob}(\mathcal{C}_{\theta^n,L}^{n-1,n-2}))$ be the union of $\mathcal{A}$ and the set of path components of objects which map to a path component of $B\mathcal{C}^{n-1,n-2}_{\theta^n,L}$ disjoint from that of $\mathcal{A}$. Theorem 5.3 does apply to $\bar{\mathcal{A}}$ and gives the weak equivalences

$$B\mathcal{C}^{n-1,\bar{\mathcal{A}}}_{\theta^n,L} \simeq B\mathcal{C}^{n-1,n-2}_{\theta^n,L} \simeq \Omega^{\infty-1}\,\mathrm{MT}\theta^n.$$

By definition, the inclusion $B\mathcal{C}^{n-1,\mathcal{A}}_{\theta^n,L} \subset B\mathcal{C}^{n-1,\bar{\mathcal{A}}}_{\theta^n,L}$ is just the inclusion of a path component, and hence becomes a homeomorphism after taking the based loop space, so we get the weak equivalence

$$\Omega B\mathcal{C}^{n-1,\mathcal{A}}_{\theta^n,L} \simeq \Omega^{\infty}\,\mathrm{MT}\theta^n.$$

The category $\mathcal{C}^{n-1,\mathcal{A}}_{\theta^n,L}$ is not quite a monoid, since it contains multiple objects (namely all those manifolds diffeomorphic to S^{2n-1}), but the space of objects is path connected, and we let $\mathcal{M}$ be the endomorphism monoid of some chosen object. By an argument similar to the proof of Lemma 6.11, the map $N_p\mathcal{C}^{n-1,\mathcal{A}}_{\theta^n,L} \to (N_0\mathcal{C}^{n-1,\mathcal{A}}_{\theta^n,L})^{p+1}$ is a fibre bundle and hence a Serre fibration with fibre $N_p\mathcal{M}$. Then the Bousfield–Friedlander theorem ([BoFr, Theorem B.4] or the earlier special case [Ma, Theorem 12.7]) implies that the inclusion $B\mathcal{M} \to B\mathcal{C}^{n-1,\mathcal{A}}_{\theta^n,L}$ is a weak equivalence (this can also be seen more geometrically as in [GRW1, Proposition 4.26]). Altogether, we obtain a weak equivalence $\Omega B\mathcal{M} \simeq \Omega^\infty\,\mathrm{MT}\theta^n$.

The monoid $\mathcal{M}$ is described up to homotopy as

$$\mathcal{M} \simeq \coprod_W B\mathrm{Diff}(W, D),$$

where W ranges over $(n-1)$-connected closed $2n$-manifolds admitting a θ^n-structure, and $D \subset W$ is a submanifold equipped with a diffeomorphism $D \cong D^{2n}$. (Admitting a θ^n-structure is equivalent to being parallelisable over the n-skeleton. Since the pair (W, D) is $(n-1)$-connected, the space of θ^n-structures is contractible when it is non-empty, cf. Lemma 7.16.) In this description, the monoid structure corresponds to connected sum and therefore $\mathcal{M}$ is homotopy commutative. The classical "group-completion" theorem (cf. [McDS]) then gives an isomorphism in homology

$$H_*(\mathcal{M})[\pi_0\mathcal{M}^{-1}] \xrightarrow{\cong} H_*(\Omega^\infty\,\mathrm{MT}\theta^n),$$

where the left-hand side denotes the ring $H_*(\mathcal{M})$ localised by inverting the multiplicative subset $\pi_0\mathcal{M}$. Finally, we claim that the localisation on the left-hand side may be calculated by inverting only the element of $\pi_0\mathcal{M}$ corresponding to $T = S^n \times S^n$. To see this, we use that if W is an element of $\mathcal{M}$, then there is another element $\overline{W}$ with the same underlying manifold, but where the identification $D^{2n} \cong D \subset W$ is changed by an orientation-reversing diffeomorphism. Then the connected sum $W \# \overline{W}$ may be identified with $\partial((W \setminus \mathrm{int}(D)) \times [0, 1])$, which we claim is diffeomorphic to the connected sum of

b copies of T, where b=rank($H_n(W)$). This can be seen by picking a minimal Morse function on the bounding manifold $(W\setminus\mathrm{int}(D))\times[0,1]$. (It has homology $\mathbb{Z}$ in degree 0 and $\mathbb{Z}^b$ in degree n and is parallelisable; cancelling critical points in a Morse function as in [Mi2, §7 and §8] proves that $(W\setminus\mathrm{int}(D))\times[0,1]$ is diffeomorphic to the boundary connected sum of b copies of $S^n\times D^{n+1}$.) Therefore the element $[W]\in\pi_0\mathcal{M}$ is invertible in the ring $H_*(\mathcal{M})[T^{-1}]$, with inverse $[T]^{-b}[\overline{W}]$. The localisation by inverting the element $[T]$ may be calculated as a direct limit, and hence we have the homology equivalence

$$\mathrm{hocolim}(\mathcal{M}\xrightarrow{\cdot T}\mathcal{M}\xrightarrow{\cdot T}\dots)\longrightarrow\Omega^\infty\,\mathrm{MT}\theta^n,$$

which upon restricting to the appropriate path component gives Theorem 1.2.

We now embark on the detailed proof of Theorem 1.8, which will occupy §7.2–§7.5, as follows. Suppose given a spherical tangential structure $\theta\colon B\to B\mathrm{O}(2n)$, a $(2n-1)$-manifold L which admits a handle structure with handles of index less than n, and a θ-structure $\ell_L\colon\varepsilon^1\oplus TL\to\theta^*\gamma$ such that the underlying map $L\to B$ is $(n-1)$-connected. In this situation the results of §3–6 apply, and will be summarised in §7.2 below as a weak equivalence between $\Omega^\infty\,\mathrm{MT}\theta$ and the loop space of the classifying space of a category $\mathcal{C}$. Then in §7.4 we apply a version of the "group-completion" theorem to relate the homology of $\Omega B\mathcal{C}$ to the homology of morphism spaces of $\mathcal{C}$, suitably localised using the theory of universal θ-ends developed in §7.3. In §7.5 we explain how to apply these results to prove Theorem 1.8. Finally, in §7.6 we explain how to deduce the results about algebraic localisation from Theorem 1.8.

7.2. The category $\mathcal{C}$

Suppose that $2n>4$, let $\theta\colon B\to B\mathrm{O}(2n)$ be a spherical tangential structure, and L be a $(2n-1)$-dimensional manifold with boundary which admits a handle structure using handles of index at most $n-1$. Let ℓ_L be a θ-structure on L, *and suppose that the underlying map* $L\to B$ *is* $(n-1)$*-connected*.

Picking a collared embedding $L\hookrightarrow\left(-\frac{1}{2},0\right]\times(-1,1)^{\infty-1}$, we have defined a category $\mathcal{C}_{\theta,L}^{n-1,n-2}$. Finally, let $\mathcal{A}\subset\pi_0(\mathrm{Ob}(\mathcal{C}_{\theta,L}^{n-1,n-2}))$ be the set of objects (M,ℓ) for which $M\setminus\mathrm{int}(L)$ is diffeomorphic to a handlebody with handles of index at most $n-1$. In Definition 2.11 we have defined

$$\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}\subset\mathcal{C}_{\theta,L}^{n-1,n-2}$$

as the full subcategory on those objects contained in $\mathcal{A}$.

Definition 7.1. A morphism in $\mathcal{C}_{\theta,L}$ is a manifold $W\subset[0,t]\times(-1,1)^\infty$ with

$$W\cap x_2^{-1}((-\infty,0])=[0,t]\times L$$

(equality as θ-manifolds). Write

$$W^\circ=W\setminus([0,t]\times\operatorname{int}(L))$$

for morphisms and similarly

$$M^\circ=M\setminus\operatorname{int}(L)$$

for objects. Morphisms or objects $X\in\mathcal{C}_{\theta,L}$ are then completely determined by X°. We shall consider the category $\mathcal{C}$ defined by

$$\begin{aligned}\operatorname{Ob}(\mathcal{C})&=\{M^\circ:M\in\operatorname{Ob}(\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}})\},\\ \operatorname{Mor}(\mathcal{C})&=\{W^\circ:W\in\operatorname{Mor}(\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}})\},\end{aligned}$$

made into a topological category by insisting that the functor $\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}\to\mathcal{C}$ given by $X\mapsto X^\circ$ is an isomorphism of topological categories.

Whether or not the θ-manifolds M° and W° define an object and a morphism of $\mathcal{C}$ seems to depend on the manifolds $M=M^\circ\cup L$ and $W=W^\circ\cup([0,t]\times L)$. However, the following lemma gives an intrinsic characterisation in terms of M° and W° alone. Let us first point out that the blanket assumption that $L\to B$ be $(n-1)$-connected is equivalent to $\partial L\to B$ being $(n-1)$-connected, as follows from the long exact sequence in homotopy groups for the triple $\partial L\to L\to B$ and the assumption that $(L,\partial L)$ is $(n-1)$-connected.

LEMMA 7.2. (i) *An object $M\in\mathcal{C}_{\theta,L}$ is in $\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$ if and only if the manifold M° can be obtained from its boundary by attaching handles of index at least n.*

(ii) *A morphism $W\in\mathcal{C}_{\theta,L}$ whose source M and target N are objects of $\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$ is in $\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$ if and only if the pair (W°,N°) is $(n-1)$-connected. This in turn happens if and only if $(W^\circ,\partial W^\circ)$ is $(n-1)$-connected.*

(iii) *For a morphism $W\in\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$ from M to N, the pairs (W,M) and (W°,M°) are also $(n-1)$-connected.*

Proof. Part (i) is true by definition of $\mathcal{A}$.

For the first part of (ii) we first note that if (W°,N°) is $(n-1)$-connected then $(W^\circ\cup_{N^\circ}N,N)$ is also $(n-1)$-connected, and W deformation retracts to $W^\circ\cup_{N^\circ}N$. For the other direction, assume that (W,N) is $(n-1)$-connected and write $\pi=\pi_1(W)$ with respect to some basepoint. We then have isomorphisms

$$H_*(W,N;\mathbb{Z}[\pi])\cong H_*(W^\circ\cup N,N;\mathbb{Z}[\pi])\cong H_*(W^\circ,N^\circ;\mathbb{Z}[\pi]),$$

where the first isomorphism follows from the deformation retraction and the second from excision. It follows that $H_*(W^\circ, N^\circ; \mathbb{Z}[\pi])$ vanishes in degrees $n-1$ and below. Since $(L, \partial L)$ is $(n-1)$-connected, the inclusions $N^\circ \to N$ and $W^\circ \to W^\circ \cup N \simeq W$ are also $(n-1)$-connected. As $n \geqslant 3$, the spaces N°, N, W° and $W^\circ \cup N$ have isomorphic π_0 and π_1 for any basepoint in N°. Therefore the inclusion $N^\circ \to W^\circ$ induces an isomorphism on π_0 and on π_1 with any basepoint, and the inclusion of universal covers at any basepoint is $(n-1)$-connected.

For the second part of (ii), we use that the inclusion $\partial L \to M^\circ$ is $(n-1)$-connected and therefore $N^\circ \to \partial W^\circ \simeq N^\circ \cup_{\partial L} M^\circ$ is too. Then the long exact sequence in homotopy groups for the triple $N^\circ \to \partial W^\circ \to W^\circ$ implies that (W°, N°) is $(n-1)$-connected if and only if $(W^\circ, \partial W^\circ)$ is.

For (iii), both inclusions $\partial L \to N^\circ$ and $N^\circ \to W^\circ$ are $(n-1)$-connected. The composition is homotopic to $\partial L \to M^\circ \to W^\circ$, so it follows from the long exact sequence in homotopy groups of this triple that (W°, M°) is also $(n-1)$-connected. The connectivity of (W, M) follows as in (ii). □

Our work in §3–§6 determines the homotopy type of the space $\Omega B\mathcal{C}$, as follows. (We emphasise again that in this section $L \to B$ is assumed to be $(n-1)$-connected and $\theta\colon B \to B\mathrm{O}(2n)$ is assumed to be spherical.)

Theorem 7.3. *There is a weak equivalence*

$$\Omega B\mathcal{C} \simeq \Omega^\infty \mathrm{MT}\theta,$$

where loops are based at any object $P^\circ \in \mathrm{Ob}(\mathcal{C})$, and $\mathrm{MT}\theta$ is the Thom spectrum associated with $\theta\colon B \to B\mathrm{O}(2n)$.

Proof. This is identical with the argument given in §7.1. Briefly, we define the set $\bar{\mathcal{A}}$ to be the union of $\mathcal{A}$ and all objects not in a path component of $B\mathcal{C}_{\theta,L}^{n-1,n-2}$ containing an element of $\mathcal{A}$, and use the string of weak equivalences

$$B\mathcal{C}_{\theta,L}^{n-1,\bar{\mathcal{A}}} \simeq B\mathcal{C}_{\theta,L}^{n-1,n-2} \simeq B\mathcal{C}_{\theta,L}^{n-1} \simeq B\mathcal{C}_{\theta,L} \simeq \psi_{\theta,L}(\infty,1) \simeq \psi_\theta(\infty,1) \simeq \Omega^{\infty-1} \mathrm{MT}\theta$$

as well as the homeomorphism $B\mathcal{C} \cong B\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$, and the fact that the inclusion

$$B\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}} \longrightarrow B\mathcal{C}_{\theta,L}^{n-1,\bar{\mathcal{A}}}$$

is a homeomorphism onto the path components it hits. □

The relevance of the category $\mathcal{C}$ to Theorem 1.8 is evident from Proposition 7.5 below.

Definition 7.4. Recall from the proof of Proposition 2.16 that we constructed a θ-manifold $D(L)$ which is diffeomorphic to the double of L. This contains $L \subset D(L)$ with its standard θ-structure, and we write $\bar{L}$ for the θ-manifold $D(L) \setminus \mathrm{int}(L)$. As L has a handle structure with handles of index at most $n-1$, $D(L)$ can be obtained from L by attaching handles of index at least n. We extend the embedding of L to an embedding $D(L) \to (-1,1)^\infty$ to get objects $D(L) \in \mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$ and $D(L)^\circ = \bar{L} \in \mathcal{C}$.

PROPOSITION 7.5. *For any object $P \in \mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$, there is a weak equivalence*

$$\varphi_P : \mathcal{C}(\bar{L}, P^\circ) \longrightarrow \mathcal{N}^\theta(P, \ell_P)$$

such that if $K: P \rightsquigarrow P'$ is a morphism in $\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$, then the diagram

$$\begin{tikzcd}
\mathcal{C}(\bar{L}, P^\circ) \arrow[r, "K^\circ \circ -"] \arrow[d, "\varphi_P"'] & \mathcal{C}(\bar{L}, (P')^\circ) \arrow[d, "\varphi_{P'}"] \\
\mathcal{N}^\theta(P, \ell_P) \arrow[r, "K \circ -"] & \mathcal{N}^\theta(P', \ell_{P'})
\end{tikzcd}$$

commutes, i.e. φ_P is a natural transformation of functors $\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}} \to \mathbf{Top}$.

Proof. We define the map φ_P as the composition

$$\varphi_P : \mathcal{C}(\bar{L}, P^\circ) \cong \mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}(D(L), P) \xrightarrow{-\circ V} \mathcal{C}_\theta^{n-1}(\varnothing, P) \xrightarrow{\simeq} \mathcal{N}^\theta(P, \ell_P),$$

where $V: \varnothing \rightsquigarrow D(L)$ is the θ-cobordism constructed in the proof of Proposition 2.16 and the last map is $(t, W) \mapsto W - te_1$. It is clear that the square commutes, so it remains to show that φ_P is a homotopy equivalence. Both its source $\mathcal{C}(\bar{L}, P^\circ)$ and its target $\mathcal{N}^\theta(P, \ell_P)$ can be described as

$$\coprod_W (\mathrm{Emb}^\partial(W, \mathbb{R}^\infty) \times \mathrm{Bun}^\partial(TW, \theta^*\gamma))/\mathrm{Diff}^\partial(W),$$

where W runs over $2n$-dimensional manifolds with boundary $\bar{L} \cup P^\circ$ which are $(n-1)$-connected relative to their boundary (in the source, this uses Lemma 7.2 (ii)), but the boundary conditions imposed on the embeddings and the bundle maps are different. Up to homotopy, the map φ_P glues an invertible bordism between the two boundary conditions. (More precisely, on the homotopy equivalent subspace of $\mathcal{C}(\bar{L}, P^\circ)$ consisting of morphisms (t, W) with $t=1$, we glue $V \cup_{\{0\} \times L} ([0,1] \times L)$, which after smoothing corners is equivalent to gluing a trivial cobordism.) $\square$

7.3. Universal θ-ends and the proof of Addendum 1.9

Let $\theta: B \to BO(2n)$ be spherical. Recall from Definition 1.7 that a *universal θ-end* is a submanifold $K \subset [0,\infty)\times\mathbb{R}^\infty$ with θ-structure ℓ_K such that $x_1: K \to [0,\infty)$ has the natural numbers as regular values. We insist that

(i) Each $K|_{[i,i+1]}$ is a highly connected cobordism, i.e. is $(n-1)$-connected relative to either end.

(ii) For each highly connected θ-cobordism $W: K|_i \rightsquigarrow P$, there is an embedding $j: W \hookrightarrow K|_{[i,\infty)}$, and a homotopy $\ell_K \circ Dj \simeq \ell_W$, both relative to $K|_i$.

We wish to have the notion of universal θ-end available to us in the cobordism category $\mathcal{C}$. Let $K|_0, K|_1, ...$ be a sequence of objects in $\mathcal{C}$, and $K|_{[i-1,i]}: K|_{i-1} \to K|_i$ be a sequence of morphisms in $\mathcal{C}$. For integers $0 \leqslant a < b$, let us write

$$K|_{[a,b]} = K|_{[b-1,b]} \circ K|_{[b-2,b-1]} \circ ... \circ K|_{[a,a+1]}$$

for the composition of the morphisms from $K|_a$ to $K|_b$. There are natural inclusions $K|_{[0,a]} \subset K|_{[0,a+1]} \subset ...$ and we let K denote the union: a non-compact smooth manifold with θ-structure. The symbol $K|_{[a,b]}$ is not ambiguous, and we can also make sense of $K|_{[a,\infty)} = \bigcup_{b>a} K|_{[a,b]}$.

Definition 7.6. Say that a non-compact manifold K of the above form is a *universal θ-end in $\mathcal{C}$* if, in the notation just introduced, properties (i) and (ii) above hold, where in (ii) we require W to be a morphism in $\mathcal{C}$.

Let us remark that it would be natural to impose a slightly stronger condition in (ii), namely that the embedding and the homotopy be relative to the slightly larger set $K|_i \cup ([0,t]\times\partial L)$, when $W \subset [0,t]\times[0,1)\times(-1,1)^\infty$. In fact the two conditions are equivalent, as the inclusion $K|_i \to K|_i \cup ([0,t]\times\partial L)$ is isotopic to a diffeomorphism (after unbending the corner of $K|_i \cup ([0,t]\times\partial L)$).

Proposition 7.8 below proves a version of Addendum 1.9 for universal θ-ends in $\mathcal{C}$. Before giving the proof, we make some preparations.

LEMMA 7.7. *Let $W: N \rightsquigarrow M$ be a highly connected cobordism between closed manifolds. There exist cobordisms $F: M \rightsquigarrow M$ and $G: N \rightsquigarrow N$ such that $F \circ W$ and $W \circ G$ both admit handle structures using only handles of index n. Similarly, if W is a morphism in the category $\mathcal{C}$, then F and G can be taken to be morphisms in this category, with the same conclusion (in this case, attaching handles along embeddings $S^{n-1}\times D^n \hookrightarrow \mathrm{int}(N)$).*

Proof. The pairs (W, M) and (W, N) are both $(n-1)$-connected, either by assumption or by Lemma 7.2. If we let F and G be sufficiently large multiples of

$$([0,1]\times M)\#(S^n\times S^n) \quad \text{and} \quad ([0,1]\times N)\#(S^n\times S^n),$$

respectively, then, by the method used in the proof of Lemma 6.21, both $F\circ W$ and $W\circ G$ admit the required handle decompositions. $\square$

Proposition 7.8. *Let $K|_{[i,i+1]}$ be a sequence of composable morphisms in $\mathcal{C}$ and let $K=\bigcup_{i\geqslant 1} K|_{[0,i]}$ be the infinite composition. Then (K,ℓ_K) is a universal θ-end in $\mathcal{C}$ if and only if the following conditions hold*:

(i) *For each integer i, the map $\pi_n(K|_{[i,\infty)})\to\pi_n(B)$ is surjective, for all basepoints in K.*

(ii) *For each integer i, the map $\pi_{n-1}(K|_{[i,\infty)})\to\pi_{n-1}(B)$ is injective, for all basepoints in K.*

(iii) *For each integer i, each path component of $K|_{[i,\infty)}$ contains a submanifold diffeomorphic to $(S^n\times S^n)\setminus\mathrm{int}(D^{2n})$, which in addition has null-homotopic structure map to B.*

Proof. To prove the "if" direction, we must show that for each integer i and each highly connected cobordism $W: K|_i\rightsquigarrow P$ with θ-structure ℓ_W, there is an embedding $j: W\hookrightarrow K|_{[i,\infty)}$ and a homotopy $\ell_K\circ Dj\simeq\ell_W$, all relative to $K|_i$.

By Lemma 7.7, for any such W there is a cobordism $F: P\rightsquigarrow P$ so that $W\circ F$ admits a handle structure with handles of index n only, so it suffices to consider the case where W consists of a single n-handle relative to $K|_i$, attached along some embedding of $S^{n-1}\times D^n$ into $K|_i$. We need to find an extension of this embedding into $K|_{[i,\infty)}$ (with the correct homotopy class of θ-structure). The map $S^{n-1}\times D^n\to K|_i\to K|_{[i,\infty)}$ is null-homotopic by assumption (ii): it is certainly null-homotopic when composed with $K|_{[i,\infty)}\to B$, because that composition is equal to the composition $S^{n-1}\times D^n\to K_i\to W\to B$. Thus there is a continuous map $f: W\to K|_{[i,\infty)}$ relative to $K|_i$. Furthermore, as $\pi_n(K|_{[i,\infty)})\to\pi_n(B)$ is surjective by assumption (i), we can change f by adding on elements of $\pi_n(K|_{[i,\infty)})$ so that

$$W \xrightarrow{f} K|_{[i,\infty)} \xrightarrow{\ell_K} B$$

is homotopic relative to $K|_i$ to ℓ_W. The θ-structures on W and K now give bundle isomorphisms

$$TW\cong\ell_W^*\theta^*\gamma\cong f^*\ell_K^*\theta^*\gamma \quad\text{and}\quad TK|_{[i,\infty)}\cong\ell_K^*\theta^*\gamma,$$

and hence an isomorphism $TW\cong f^*TK|_{[i,\infty)}$ relative to $K|_i$, i.e. $f: W\to K$ is covered by a bundle map $TW\to TK$, which near $K|_i$ is the derivative of the embedding. By Smale–Hirsch theory, we may therefore homotope $f: W\to K|_{[i,\infty)}$ to an immersion, without changing it near $K|_i$.

Finally, we explain how to replace the immersion $f: W\to K|_{[i,\infty)}$ by an embedding. It suffices to make f an embedding near a core $(D^n,\partial D^n)\subset(W,K|_i)$ of the n-handle, and we shall write $\hat{f}: D^n\to K|_{[i,\infty)}$ for the restriction of f. After changing f by a small

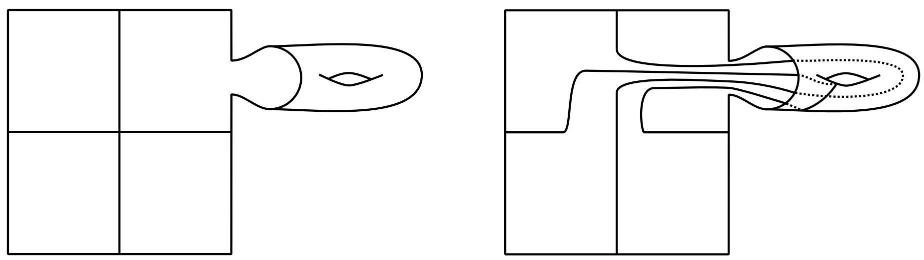

Figure 10. Disjoint discs inside $(D^n\times D^n)\natural((S^n\times S^n)\setminus\mathrm{int}(D^{2n}))$.

isotopy, we may assume that all self-intersections of $\hat{f}$ are transverse. We shall explain how to remove one self-intersection point of $\hat{f}$, changing the homotopy class of f in the process. Around a self-intersection point, choose a coordinate $\mathbb{R}^n\times\mathbb{R}^n\hookrightarrow K|_{[i,\infty)}$ so that $\mathbb{R}^n\times\{0\}$ and $\{0\}\times\mathbb{R}^n$ give local coordinates around the two preimages of the double point. By assumption (iii) we can find an embedded $(S^n\times S^n)\setminus\mathrm{int}(D^{2n})\subset K|_{[i,\infty)}$ with null-homotopic map to B. We can also assume it is disjoint from the image of $\hat{f}$, since D^n is compact. Then we choose an embedded path from this $(S^n\times S^n)\setminus\mathrm{int}(D^{2n})$ to the patch $\mathbb{R}^n\times\mathbb{R}^n$ disjoint from the image of $\hat{f}$, and thicken it up: inside this we have a subset diffeomorphic to the boundary connected sum

$$(D^n\times D^n)\natural((S^n\times S^n)\setminus\mathrm{int}(D^{2n})), \tag{7.1}$$

which the image of $\hat{f}$ intersects in $D^n\vee D^n=(D^n\times\{0\})\cup(\{0\}\times D^n)$. This situation is depicted (for $n=1$) in the first picture in Figure 10. Inside this subset there is a pair of disjointly embedded discs which are equal to the standard pair near the boundary, as shown in the second picture of Figure 10. (The construction in the figure works for any n. More formally a manifold diffeomorphic, but not equal, to (7.1) may be obtained by performing the connected sum of $D^n\times D^n$ and $S^n\times S^n$ at the origin $(0,0)\in D^n\times D^n$ in a way that identifies neighbourhoods of the wedge points in

$$D^n\vee D^n\subset D^n\times D^n \quad\text{and}\quad S^n\vee S^n\subset S^n\times S^n.$$

The resulting manifold has the disjoint discs

$$D_1=(D^n\times\{0\})\#(S^n\times\{*\}) \quad\text{and}\quad D_2=(\{0\}\times D^n)\#(\{*\}\times S^n)$$

and is diffeomorphic (relative to $(\partial D^n\times\{0\})\amalg(\{0\}\times\partial D^n)$) to (7.1) where the connected sum is performed away from $D^n\vee D^n\subset D^n\times D^n$. The second picture in Figure 10 shows the image of D_1 and D_2 under such a diffeomorphism.)

We can modify $\hat{f}$ by redefining it to have these discs as image instead. This reduces by 1 the number of geometric self-intersections of $\hat{f}$, and up to homotopy we have added an element of $\pi_n((S^n \times S^n)\setminus\mathrm{int}(D^{2n}))$ to the homotopy class of $\hat{f}$. As

$$(S^n \times S^n)\setminus\mathrm{int}(D^{2n}) \longrightarrow K|_{[i,\infty)} \longrightarrow B$$

was null-homotopic, we have not changed the homotopy class of $\hat{f}$ in B.

After finitely many steps, we have changed $\hat{f}$ to an embedding. The corresponding embedding $f: W \to K|_{[i,\infty)}$ (obtained by thickening $\hat{f}$ up again) is homotopic to the original one after composing with $\ell_K: K|_{[i,\infty)} \to B$, so $\ell_K \circ f \simeq \ell_W$ relative to $K|_i$. Hence the induced θ-structure on W is homotopic to the given one relative to $K|_i$.

To prove the "only if" direction, we must prove that any universal θ-end (K, ℓ_K) satisfies the three conditions. It is clear that (iii) is necessary: For any i we can let W be the boundary connected sum of the cylinder $[i, i+1] \times K|_i$ and the (parallelisable) manifold $(S^n \times S^n)\setminus\mathrm{int}(D^{2n})$ equipped with a trivial θ-structure. Universality implies that this admits an embedding into $K|_{[i,\infty)}$, and hence $(S^n \times S^n)\setminus\mathrm{int}(D^{2n})$ does too.

For property (i), it suffices to prove that for any i and any $\alpha \in \pi_n(B)$, there exists a morphism $W_\alpha \in \mathcal{C}(K|_i, P)$ for some P, with $\alpha \in \mathrm{Im}(\pi_n(W_\alpha) \to \pi_n(B))$. To construct such a manifold, we may represent α by a map $S^n \to B$ and lift the composition

$$\theta \circ \alpha: S^n \longrightarrow B \longrightarrow B\mathrm{O}(2n)$$

to a map $f: S^n \to B\mathrm{O}(n)$. If we let $D \to S^n$ be the disc bundle of the vector bundle classified by f, the tangent bundle of D is classified by $\theta \circ \alpha$, and therefore admits a θ-structure whose underlying map $S^n \simeq D \to B$ represents α. We can then let W_α be the boundary connected sum of $[i, i+1] \times K|_i$ and D.

Finally, for property (ii), we use that each $K|_{[j,j+1]}$ is a highly connected cobordism to see that $\pi_{n-1}(K|_i) \to \pi_{n-1}(K|_{[i,\infty)})$ is surjective. It therefore suffices to prove that for any $\alpha \in \mathrm{Ker}(\pi_{n-1}(K|_i) \to \pi_{n-1}(B))$, there exists a morphism $W_\alpha \in \mathcal{C}(K|_i, P)$ for some P, with $\alpha \in \mathrm{Ker}(\pi_{n-1}(K|_i) \to \pi_{n-1}(W_\alpha))$. We may represent α by an embedding $S^{n-1} \to K|_i$. Since the composition $S^{n-1} \to K|_i \to B \to B\mathrm{O}(2n)$ is trivial, the normal bundle of the embedding is stably trivial and hence trivial, so we may extend to an embedding

$$f: S^{n-1} \times D^n \longrightarrow K|_i.$$

The underlying manifold of the morphism W_α is then defined as the trace of surgery along f, and a θ-structure is constructed from a choice of null-homotopy of $S^{n-1} \to K|_i \to B$. □

Proof of Addendum 1.9. The assumptions of the addendum are exactly as in Proposition 7.8, and the same argument applies. □

The following three propositions establish further useful properties of universal θ-ends. The first proposition gives a refinement of property (ii), which lets us exert more control on the behaviour of the embedding j which is provided by (ii). Propositions 7.10 and 7.11 give strong existence and uniqueness properties for universal θ-ends (and universal θ-ends in $\mathcal{C}$), which essentially say that a universal θ-end (K, ℓ_K) is determined up to diffeomorphism (respecting θ-structures) by $(K|_0, \ell_K|_0)$.

PROPOSITION 7.9. *If* (K, ℓ_K) *is a universal* θ*-end* (*or a universal* θ*-end in* $\mathcal{C}$) *then it also satisfies*

(ii$'$) *For each highly connected* θ*-cobordism* $W: K|_i \rightsquigarrow P$*, there is a* $k \gg i$*, an embedding* $j: W \hookrightarrow K|_{[i,k]}$*, and a homotopy* $\ell_K \circ Dj \simeq \ell_W$*, both relative to* $K|_i$*, such that the complement of* $j(W)$ *is a cobordism* $Z: P \rightsquigarrow K|_k$ *which is highly connected.*

Proof. Let us treat the case of a universal θ-end; working in $\mathcal{C}$ can be done in the same way. As W is $(n-1)$-connected relative to either end, Lemma 7.7 applies, and for sufficiently large g, the manifold $W' = W \# g(S^n \times S^n) = (([0,1] \times P) \# g(S^n \times S^n)) \circ W$ admits a handle structure relative to $K|_i$ using handles of index n only.

By universality, there is an embedding of θ-manifolds $j': W' \hookrightarrow K|_{[i,k']}$ relative to $K|_i$. We wish to modify this embedding, and increase k', so that if

$$\{e_\alpha: (D^n \times D^n, D^n \times S^{n-1}) \hookrightarrow (W', K|_i)\}_{\alpha \in I}$$

denotes the collection of relative n-handles of W', there exist embedded spheres

$$\{f_\beta: S^n \hookrightarrow K|_{[i,k']}\}_{\beta \in I}$$

so that

$$e_\alpha(\{0\} \times D^n) \pitchfork f_\beta(S^n) = \begin{cases} \varnothing, & \text{if } \alpha \neq \beta, \\ \{*\}, & \text{if } \alpha = \beta. \end{cases}$$

This will be done by inductively modifying j' on the image of e_α, one α at a time. For each α, the composition $j' \circ e_\alpha: D^n \times D^n \to W' \to K|_{[i,k']}$ may be extended to an embedding $\mathbb{R}^n \times (D^n, \partial D^n) \to (K|_{[i,k']}, K|_i)$ and by property (iii) of Proposition 7.8 this may be further extended to an embedding

$$c: (\mathbb{R}^n \times D^n) \natural ((S^n \times S^n) \setminus \mathrm{int}(D^{2n})) \longrightarrow K|_{[i,\infty)},$$

disjoint from the other handles and any f_β already constructed, such that the structure map to B is null-homotopic on the embedded $(S^n \times S^n) \setminus \mathrm{int}(D^{2n})$. Then the image $j'(e_\alpha(D^n \times D^n))$ is contained in the image of c, and it is depicted in the first part of Figure 11.

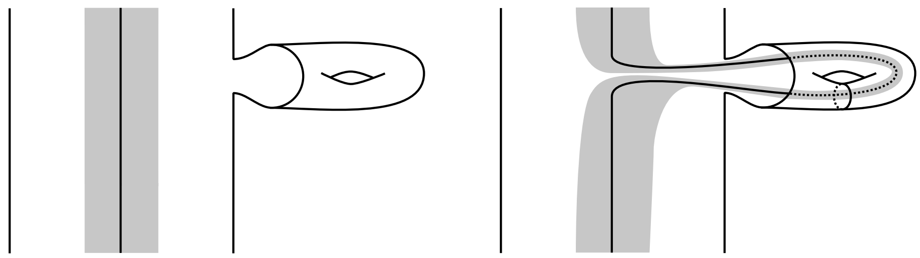

Figure 11. Modifying j' on a handle of W'.

Just as in the proof of Proposition 7.8 and Figure 10, we may then change the map j' on the image of e_α as indicated in the second part of Figure 11: The image of the handle under the modified j' is the shaded strip in the figure. Its homotopy class is changed by adding one of the summands in $S^n \vee S^n \subset (S^n \times S^n) \setminus \mathrm{int}(D^{2n}) \subset K|_{[i,\infty)}$. Because the structure map $K \to B$ is null-homotopic on this $S^n \vee S^n$, we have not changed the homotopy class of the composition $W' \to K|_{[i,\infty)} \to B$. We can then use the inclusion of the other summand $S^n \to S^n \vee S^n \to K|_{[i,\infty)}$ as f_β, shown as the small transverse sphere in Figure 11.

We denote by $j' \colon W' \hookrightarrow K|_{[i,k]}$ this improved embedding, and by Z' the complement of the interior of the image of j'. We now prove that the cobordism Z' is highly connected. This can be verified one path component at a time, so we may assume that Z' is path connected and pick a basepoint. The cobordism W' is itself highly connected, and since $n \geqslant 3$, composing with W' from either side will preserve fundamental groups. Therefore $K|_i$, W', P, Z', $K|_{[i,k]}$ and $K|_k$ all have the same fundamental group, which we denote by π. The $\mathbb{Z}[\pi]$-module $H_n(W', K|_i; \mathbb{Z}[\pi])$ is free with basis $[e_\alpha]$, so $\mathrm{Hom}(H_n(W', K|_i; \mathbb{Z}[\pi]), \mathbb{Z}[\pi])$ is free on the dual basis elements $[e_\alpha]^\vee$. The intersection pairing gives a map

$$H_n(K|_{[i,k]}; \mathbb{Z}[\pi]) \longrightarrow \mathrm{Hom}(H_n(W', K|_i; \mathbb{Z}[\pi]), \mathbb{Z}[\pi]),$$

which by construction sends $[f_\alpha]$ to $[e_\alpha]^\vee$, and so is surjective. This map can be factored as

$$\begin{tikzcd} H_n(K|_{[i,k]}) \arrow[r] & H_n(K|_{[i,k]}, K|_k) \arrow[r, "\varphi"] & H_n(K|_{[i,k]}, Z') \arrow[dll, "\cong"] \\ H_n(W', P) \arrow[r, "\cong"] & H^n(W', K|_i) \arrow[r, "\cong"] & \mathrm{Hom}(H_n(W', K|_i), \mathbb{Z}[\pi]), \end{tikzcd}$$

where all homology and cohomology is with coefficients in $\mathbb{Z}[\pi]$. The third map in the composition is an isomorphism by excision, the fourth by Poincaré duality (where we again use the conventions of [Wa, Chapter 2] to interchange left and right $\mathbb{Z}[\pi]$-modules using the orientation character). The fifth is an isomorphism because $(W', K|_i)$ is $(n-1)$-connected so the higher Ext terms in the universal coefficients spectral sequence (cf. [Le, Theorem 2.3]) vanish. Therefore the map φ is surjective and it follows from the long exact sequence in $\mathbb{Z}[\pi]$-homology for the triple $(K|_{[i,k]}, Z', K|_k)$ that $H_{n-1}(Z', K|_k; \mathbb{Z}[\pi]) = 0$. The same long exact sequence shows the vanishing of $H_*(Z', K|_k; \mathbb{Z}[\pi])$ in lower degrees, and hence $(Z', K|_k)$ is $(n-1)$-connected.

That the pair (Z', P) is $(n-1)$-connected follows by the long exact sequence for $\mathbb{Z}[\pi]$-homology of the triple $(K|_{[i,k]}, W', K|_i)$ and excision $(Z', P) \sim (K|_{[i,k]}, W')$. Finally, we note that if Z denotes the complement of the image of $j = j'|_W$, then we have

$$Z = Z' \circ (([0,1] \times P) \# g(S^n \times S^n))$$

so it is also a highly connected cobordism. □

Proposition 7.10. *Let (K, ℓ_K) and $(K', \ell_{K'})$ be universal θ-ends, and suppose we are given a highly connected θ-cobordism $W: K|_0 \rightsquigarrow K'|_0$. Then there is a diffeomorphism $\varphi: W \cup_{K'|_0} K' \cong K$, and a homotopy $\ell_K \circ D\varphi \simeq \ell_{W \cup K'}$, both relative to $K|_0$. Furthermore, there is a weak homotopy equivalence*

$$\operatorname*{hocolim}_{i\to\infty} \mathcal{N}^\theta(K|_i, \ell_K|_i) \simeq \operatorname*{hocolim}_{i\to\infty} \mathcal{N}^\theta(K'|_i, \ell_{K'}|_i).$$

Proof. By replacing K' with $W \cup_{K'|_0} K'$, we may as well assume that $K|_0 = K'|_0$ as θ-manifolds, and that W is the trivial cobordism. As K' is a universal θ-end, we may find an embedding of θ-manifolds $j_1: K|_{[0,1]} \hookrightarrow K'|_{[0,k_1']}$ relative to $K|_0$, and by Proposition 7.9 we may suppose its complement $Z_1: K|_1 \rightsquigarrow K'|_{k_1'}$ is highly connected. Now, as K is a universal θ-end, we may find an embedding of θ-manifolds $j_1': Z_1 \hookrightarrow K|_{[1,k_1]}$ relative to $K|_1$, again with highly connected complement $Z_2: K'|_{k_1'} \rightsquigarrow K|_{k_1}$. Together, j_1^{-1} and j_1' give an embedding of θ-manifolds $K'|_{[0,k_1']} \hookrightarrow K|_{[0,k_1]}$. Continuing in this way, we produce the required diffeomorphism φ and homotopy.

For the second part, note that we have constructed a direct system

$$\mathcal{N}^\theta(K|_0) \xrightarrow{K|_{[0,1]}\circ -} \mathcal{N}^\theta(K|_1) \xrightarrow{Z_1 \circ -} \mathcal{N}^\theta(K'|_{k_1'}) \xrightarrow{Z_2 \circ -} \mathcal{N}^\theta(K|_{k_1}) \longrightarrow \dots$$

which contains cofinal subsystems which are also cofinal in either of the direct systems used to form the homotopy colimits in the statement. □

Proposition 7.11. *Let $\pi_n(B)$ be countable. Then for any object $(M, \ell_M) \in \mathcal{C}$ there is a universal θ-end (K, ℓ_K) in $\mathcal{C}$ with $(K|_0, \ell_K|_0) = (M, \ell_M)$. Moreover, $K \cup ([0,\infty) \times L)$ is then a universal θ-end.*

Proof. In the proof of Proposition 7.8 we saw that for each

$$\alpha \in \operatorname{Ker}(\pi_{n-1}(M) \to \pi_{n-1}(B)),$$

there exists a morphism $W_\alpha \in \mathcal{C}(M,P)$ with $\alpha \in \operatorname{Ker}(\pi_{n-1}(M) \to \pi_{n-1}(W_\alpha))$, and for each element $\alpha \in \pi_n(B)$, there exists a morphism $W_\alpha \in \mathcal{C}(M,P)$ with $\alpha \in \operatorname{Im}(\pi_n(W_\alpha) \to \pi_n(B))$. A priori, the target P depends on α, but as θ has been assumed to be spherical, it is reversible (by Proposition 5.7), and we may find another morphism $P \rightsquigarrow M$; after composing, we may assume that $M=P$ so we have endomorphisms $W_\alpha \in \mathcal{C}(M,M)$. We then construct a universal θ-end in $\mathcal{C}$ by letting $K|_i = M$ for each integer $i \geqslant 0$ and letting each $K|_{[i,i+1]}$ be of the form $W_{\alpha_i} \# (S^n \times S^n)$, where the α_i form a sequence of elements of $\pi_n(B) \cup \operatorname{Ker}(\pi_{n-1}(M) \to \pi_{n-1}(B))$ in which each element occurs infinitely often. (This is possible because $\pi_n(B)$ is assumed to be countable and $\pi_{n-1}(M)$ is automatically countable.) It then follows from Proposition 7.8 that K is a universal θ-end in $\mathcal{C}$.

It is obvious that gluing $[0,\infty) \times L$ to a universal θ-end in $\mathcal{C}$ gives a universal θ-end, since the homotopical properties in Proposition 7.8 are clearly preserved. □

Corollary 7.12. *Let (K, ℓ_K) be a universal θ-end for which $P = (K|_0, \ell_K|_0)$ is an object of $\mathcal{C}_{\theta,L}^{n-1,\mathcal{A}}$. Then we may isotope the proper embedding $K \to [0,\infty) \times (-1,1)^\infty$ and homotope the bundle map $\ell_K \colon TK \to \theta^*\gamma$, both relative to $K|_0$, after which K is of the form $K^\circ \cup ([0,\infty) \times L)$ where K° is a universal θ-end in $\mathcal{C}$.*

Proof. By Proposition 7.8, the structure map $\ell_K \colon K \to B$ induces a surjection on π_n. Since K is a manifold, $\pi_n(K)$, and hence $\pi_n(B)$, is countable, so there exists a universal θ-end in $\mathcal{C}$, by Proposition 7.11. Denoting this by K°, the θ-manifold $K^\circ \cup ([0,\infty) \times L)$ is a universal θ-end, and hence by Proposition 7.10 is isomorphic to the original K. □

7.4. Group completion

Let us return to the category $\mathcal{C}$ of §7.2. Assigning to a morphism $W \in \mathcal{C}(P_0, P_1)$ the corresponding 1-simplex in the nerve of $\mathcal{C}$ gives a continuous map $\mathcal{C}(P_0,P_1) \to \Omega_{P_0,P_1} B\mathcal{C}$, analogous to the map $\mathcal{M} \to \Omega B\mathcal{M}$ in the outline in §7.1. As in that section, the effect in homology can be studied by a version of the "group-completion" theorem. The classical group-completion theorem concerns a topological *monoid* M, and says that the map $H_*(M) \to H_*(\Omega BM)$ is an algebraic localisation at the multiplicative subset $\pi_0(M) \subset H_*(M)$. The group-completion theorem holds under the assumption that this localisation *admits a calculus of right fractions*, cf. [McDS]. A similar result holds for topological *categories*, and here implies that $H_*(\Omega B\mathcal{C})$ is a suitable direct limit of $H_*(\mathcal{C}(P_0,P_1))$, generalising the localisation in the monoid case. As in the monoid case, some assumption is needed in order to apply the group-completion theorem: Lemma 7.15 below can be seen as a multi-object version of admitting a calculus of right fractions.

THEOREM 7.13. *Let*

$$K|_0 \xrightarrow{K|_{[0,1]}} K|_1 \xrightarrow{K|_{[1,2]}} K|_2 \xrightarrow{K|_{[2,3]}} K|_3 \xrightarrow{K|_{[3,4]}} \ldots$$

be a sequence of composable morphisms in $\mathcal{C}$ such that K is a universal θ-end in the category $\mathcal{C}$ and $\mathcal{C}(\bar{L}, K|_0) \neq \varnothing$. Then there is a map

$$\operatorname*{hocolim}_{i\to\infty} \mathcal{C}(\bar{L}, K|_i) \longrightarrow \Omega B\mathcal{C}$$

which is a homology equivalence.

The proof will be based on Proposition 7.14 below. Let $F_i: \mathcal{C}^{\mathrm{op}} \to \mathbf{Top}$ denote the representable functor $\mathcal{C}(-, K|_i)$, in the sense of [GMTW, Proposition 7.1] and the discussion preceding it. Let $F_\infty: \mathcal{C}^{\mathrm{op}} \to \mathbf{Top}$ denote the (objectwise) homotopy colimit of the natural transformations $F_i \to F_{i+1}$ given by right composition with $K|_{[i,i+1]}$. The following proposition is the key ingredient for proving Theorem 7.13.

PROPOSITION 7.14. *The functor F_∞ sends each morphism in $\mathcal{C}$ to a homology equivalence.*

Proof. By Lemma 7.7, it suffices to prove that F_∞ sends any cobordism admitting a handle structure with a single n-handle to a homology isomorphism: indeed, in the notation of that lemma, for any cobordism W, the functor F_∞ sends both $W \circ F$ and $G \circ W$ to homology isomorphisms, but then it must send W to one as well. We therefore consider a cobordism $W \in \mathcal{C}(N, M)$ admitting a handle structure with a single n-handle. The cobordism W gives a map of direct systems

$$\begin{tikzcd}
\mathcal{C}(M, K|_0) \arrow[r] \arrow[d, "- \circ W"'] & \mathcal{C}(M, K|_1) \arrow[r] \arrow[d, "- \circ W"'] & \mathcal{C}(M, K|_2) \arrow[r] \arrow[d, "- \circ W"'] & \mathcal{C}(M, K|_3) \arrow[r] \arrow[d, "- \circ W"'] & \cdots \\
\mathcal{C}(N, K|_0) \arrow[r] & \mathcal{C}(N, K|_1) \arrow[r] & \mathcal{C}(N, K|_2) \arrow[r] & \mathcal{C}(N, K|_3) \arrow[r] & \cdots .
\end{tikzcd}$$

Taking homotopy colimits of the rows gives a map $F_\infty(M) \to F_\infty(N)$, and Lemma 7.15 below implies that the induced map on homology is a bijection, finishing the proof of Proposition 7.14. □

Proof of Theorem 7.13. We will apply [GMTW, Proposition 7.1] to the functor F_∞. Each of the maps $N_0(\mathcal{C} \downarrow K|_i) = N_0(\mathcal{C} \wr F_i) \to N_0(\mathcal{C})$ is a fibre bundle by an argument similar to the proof of Lemma 6.11 and hence a Serre fibration. Using that $F_\infty = \operatorname{hocolim}_{i\to\infty} F_i$

sends all morphisms to homology equivalences, it follows as in [GMTW, Proposition 7.1] that the square

$$\begin{tikzcd} F_\infty(\bar{L}) \arrow[r] \arrow[d] & B(\mathcal{C}\wr F_\infty) \arrow[d] \\ \{\bar{L}\} \arrow[r] & B\mathcal{C} \end{tikzcd}$$

is homology cartesian and so the induced map from $F_\infty(\bar{L})$ to the homotopy fibre of $B(\mathcal{C}\wr F_\infty)\to B\mathcal{C}$ is a homology equivalence. Since $B(\mathcal{C}\wr F_\infty)\simeq\operatorname{hocolim}_{i\to\infty} B(\mathcal{C}\wr F_i)$ is contractible, this homotopy fibre is equivalent to $\Omega B\mathcal{C}$ establishing Theorem 7.13.

(As stated, [GMTW, Proposition 7.1] requires $\operatorname{hocolim}_{i\to\infty} N_0(\mathcal{C}\wr F_i)\to N_0\mathcal{C}$ to be a Serre fibration, and it is not clear that the telescope of a direct system of Serre fibrations over the same base is again a Serre fibration. Such a telescope is however easily seen to have the weak covering homotopy property, i.e. the WCHP of [D, Definition 5.1], for CW complexes. This property is closed under pullback and implies quasifibration, which is sufficient here.) □

Lemma 7.15. *Let $W\colon N\rightsquigarrow M$ be a cobordism which is obtained by attaching a single n-handle to N. For each i there is a $k\geqslant i$ such that the commutative square*

$$\begin{tikzcd} \mathcal{C}(M,K|_i) \arrow[r, "K|_{[i,k]}\circ -"] \arrow[d, "-\circ W"'] & \mathcal{C}(M,K|_k) \arrow[d, "-\circ W"] \\ \mathcal{C}(N,K|_i) \arrow[r, "K|_{[i,k]}\circ -"'] \arrow[ur, dashed] & \mathcal{C}(N,K|_k) \end{tikzcd} \tag{7.2}$$

admits a dashed map making the top triangle commute up to homotopy, and a (possibly different) dashed map making the bottom triangle commute up to homotopy.

Proof. The objects M and N in $\mathcal{C}$ are submanifolds of $[0,1)\times\mathbb{R}^\infty$ (with θ-structure), and the morphism $W\in\mathcal{C}(N,M)$ is a submanifold of $[0,t]\times[0,1)\times\mathbb{R}^\infty$. Rotating W in the first two coordinate directions gives a submanifold

$$\overline{W}\subset[0,t]\times(-1,0]\times\mathbb{R}^\infty$$

with incoming boundary $\{0\}\times\overline{M}$ and outgoing boundary $\{t\}\times\overline{N}$. As in the proof of Proposition 2.16, the θ-structure on M extends to a θ-structure on the closed manifold $\overline{M}\cup M\subset(-1,1)\times\mathbb{R}^\infty$, giving an object $\langle M,M\rangle\in\mathcal{C}_\theta^{n-1}$ with a canonical null-bordism

$$V\in\mathcal{C}_\theta^{n-1}(\varnothing,\langle M,M\rangle).$$

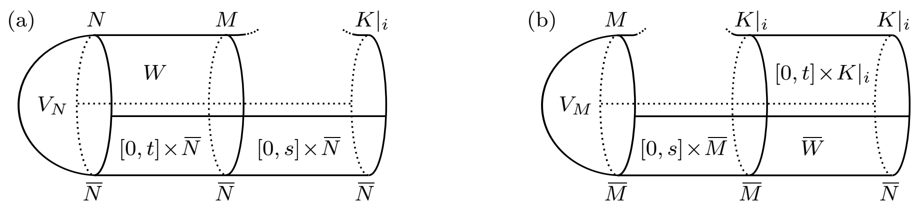

Figure 12. Schematic description of the two compositions in the diagram, along the top (a) and along the bottom (b). The unlabelled cobordism represents a point in $\mathcal{C}(M, K|_i)$, of length s, and V_X denotes the canonical null-bordism of $\overline{X} \cup X$.

Similarly, we have objects $\langle M, K|_i\rangle = \overline{M} \cup K|_i$ and $\langle N, K|_i\rangle = \overline{N} \cup K|_i$, and the submanifold $\overline{W} \cup ([0,t] \times K|_i) \subset [0,t] \times (-1,1) \times \mathbb{R}^\infty$ inherits a θ-structure from W, giving an element of $\mathcal{C}_\theta^{n-1}(\langle M, K|_i\rangle, \langle N, K|_i\rangle)$ which we shall denote $\langle W, K|_i\rangle$. We now consider a diagram

$$\begin{tikzcd}
\mathcal{C}(M, K|_i) \arrow[r, "- \circ W"] \arrow[d, "\simeq"'] & \mathcal{C}(N, K|_i) \arrow[d, "\simeq"] \\
\mathcal{N}^\theta(\langle M, K|_i\rangle) \arrow[r, "\langle W, K|_i\rangle \circ -"] & \mathcal{N}^\theta(\langle N, K|_i\rangle)
\end{tikzcd}$$

similar to the diagram in Proposition 7.5. The vertical maps are defined as in that proposition, but with L and $\bar{L}$ replaced by $\overline{M}$ and M (or $\overline{N}$ and N for the right column), and are therefore weak equivalences by the same argument. The two compositions from the top left corner to the bottom right corner of the diagram are induced by gluing isomorphic θ-manifolds with equal boundary, and are therefore homotopic; Figure 12 gives a schematic view of the manifolds involved.

The diagram of solid arrows in (7.2) may now be replaced by

$$\begin{tikzcd}
\mathcal{N}^\theta(\langle M, K|_i\rangle) \arrow[r, "\langle M, K|_{[i,k]}\rangle \circ -"] \arrow[d, "\langle W, K|_i\rangle \circ -"'] & \mathcal{N}^\theta(\langle M, K|_k\rangle) \arrow[d, "\langle W, K|_k\rangle \circ -"] \\
\mathcal{N}^\theta(\langle N, K|_i\rangle) \arrow[r, "\langle N, K|_{[i,k]}\rangle \circ -"] & \mathcal{N}^\theta(\langle N, K|_k\rangle),
\end{tikzcd}$$

where $\langle M, K|_{[i,k]}\rangle = ([i,k] \times \overline{M}) \cup K|_{[i,k]} \subset [i,k] \times (-1,1) \times \mathbb{R}^\infty$, and similarly for $\langle N, K|_{[i,k]}\rangle$.

Let us first show that there is a dashed map making the top triangle commute up to homotopy, for some $k \gg i$. We wish to find an embedding (of θ-manifolds) of $\langle W, K|_i\rangle$ into $\langle M, K|_{[i,k]}\rangle$ relative to $\langle M, K|_i\rangle$, with complement being a θ-cobordism

$Z\colon \langle N, K|_i\rangle \rightsquigarrow \langle M, K|_k\rangle$. If we can ensure that $(Z, \langle M, K|_k\rangle)$ is $(n-1)$-connected, then gluing on Z gives a map

$$Z\circ -\colon \mathcal{N}^\theta(\langle N, K|_i\rangle) \longrightarrow \mathcal{N}^\theta(\langle M, K|_k\rangle)$$

making the top triangle commute (as $\langle W, K|_i\rangle \circ Z \cong \langle M, K|_{[i,k]}\rangle$ as θ-manifolds), as required.

By definition of the category $\mathcal{C}$, $\overline{M}$ is obtained from its boundary, ∂L, by attaching handles of index at least n. Thus, by transversality, the attaching map for the n-handle of $\overline{W}$ relative to $\overline{M}$ may be assumed to have image in a collar neighbourhood $[-\varepsilon, 0]\times\partial L \subset \overline{M}$. Thus $\langle W, K|_i\rangle$ may be obtained from $\langle M, K|_i\rangle$ by attaching a single n-handle along

$$f\colon S^{n-1}\times D^n \hookrightarrow [0,\varepsilon]\times\partial L \subset K|_i,$$

so up to diffeomorphism (relative to its incoming boundary) the cobordism $\langle W, K|_i\rangle$ is of the form $\langle M, W'\rangle$ for some cobordism $W'\colon K|_i \rightsquigarrow X$ in $\mathcal{C}$. As $K|_{[i,\infty)}$ is a universal θ-end in the category $\mathcal{C}$, there exists an embedding of θ-manifolds $j'\colon W' \hookrightarrow K|_{[i,k]}$ relative to $K|_i$, for some $k \gg i$, and by Proposition 7.9 we may assume that its complement Z' is highly connected. Gluing M back in, we obtain an embedding $j\colon \langle W, K|_i\rangle \hookrightarrow \langle M, K|_{[i,k]}\rangle$ relative to $\langle M, K|_i\rangle$ whose complement $Z \cong \langle M, Z'\rangle$ is highly connected, as required.

To produce the dashed map making the bottom triangle commute up to homotopy, we must produce an embedding relative to $\langle N, K|_k\rangle$ of $\langle W, K|_k\rangle$ into $\langle N, K|_{[i,k]}\rangle$, for some suitably large k, with an appropriate connectivity condition on its complement. As we shall explain, this reduces to the same embedding problem as for the upper triangle. We have collar neighbourhoods $[-\varepsilon, 0]\times\partial L \subset \overline{N}$ and $[0,\varepsilon]\times\partial L \subset K|_i$, and as above, we can suppose $\overline{W}$ is obtained from $\overline{N}$ by attaching a single n-handle along a map

$$f_0\colon S^{n-1}\times D^n \hookrightarrow [-\varepsilon, 0]\times\partial L \subset \overline{N}.$$

We now consider $\overline{N}$ to lie inside

$$([i-\varepsilon, i]\times(\overline{N}\cup K|_i))\cup K|_{[i,\infty)},$$

where we may extend f_0 inside $[i-\varepsilon, i]\times[-\varepsilon,\varepsilon]\times\partial L$ to an embedding $f_{[0,1]}$ of $[0,1]\times S^{n-1}\times D^n$ so that $f_1 = f_{[0,1]}|_{\{1\}\times S^{n-1}\times D^n}$ is an embedding into $\{i\}\times[0,\varepsilon]\times\partial L \subset K|_i$. Since $K|_{[i,\infty)}$ is a universal θ-end in $\mathcal{C}$, we may extend the embedding f_1 to an embedding of the handle $\{1\}\times D^n\times D^n$ into $K|_{[i,k]}$ for some $k \gg i$, having θ-structure homotopic to the one given on the n-handle of W, and as in the first part of the proof of Proposition 7.9 we can ensure that there is an embedded n-sphere in $K|_{[i,k]}$ intersecting the core of this handle transversely in precisely one point. Choosing a further extension $f_{[-1,0]}$ of f_0 to an embedding $[-1,0]\times S^{n-1}\times D^n \hookrightarrow [i,k]\times \overline{N}$ which sends $\{-1\}\times S^{n-1}\times D^n$ to $\{k\}\times \overline{N}$, we altogether obtain an embedding

$$(([-1,1]\times S^{n-1})\cup(\{1\}\times D^n))\times D^n \longrightarrow ([i-\varepsilon,i]\times(\overline{N}\cup K|_i))\cup K|_{[i,k]}\cup([i,k]\times\overline{N}).$$

The source of this map is diffeomorphic to a tubular neighbourhood of the n-handle in W, and the target is diffeomorphic relative to $K|_k\cup\overline{N}=\langle N,K|_k\rangle$ to $K|_{[i,k]}\cup([i,k]\times\overline{N})=\langle N,K|_{[i,k]}\rangle$, so this induces an embedding

$$e\colon\langle W,K|_k\rangle \hookrightarrow \langle N,K|_{[i,k]}\rangle$$

relative to $\langle N,K|_k\rangle$, with complement Z being a cobordism from $\langle N,K|_i\rangle$ to $\langle M,K|_k\rangle$.

It remains to show that Z is a highly connected cobordism. As usual, all manifolds in sight have the same fundamental group, π, and we proceed as in the second part of the proof of Proposition 7.9. We have ensured that the map

$$H_n(\langle N,K|_{[i,k]}\rangle,\langle N,K|_i\rangle;\mathbb{Z}[\pi]) \longrightarrow H_n(\langle N,K|_{[i,k]}\rangle,Z';\mathbb{Z}[\pi])\cong\mathbb{Z}[\pi]$$

in the long exact sequence for the triple $(\langle N,K|_{[i,k]}\rangle,Z',\langle N,K|_i\rangle)$ is surjective, as under excision, Poincaré duality, and universal coefficients it corresponds to the map

$$H_n(\langle N,K|_{[i,k]}\rangle,\langle N,K|_i\rangle;\mathbb{Z}[\pi]) \longrightarrow H_n(W,N;\mathbb{Z}[\pi])^\vee\cong\mathbb{Z}[\pi]$$

sending a cycle to the functional represented by intersection with that cycle, and we made sure that there was a sphere in $\langle N,K|_{[i,k]}\rangle$ intersecting the core of the embedded n-handle of (W,N) transversely in one point. It then follows from the long exact sequence for the triple $(\langle N,K|_{[i,k]}\rangle,Z',\langle N,K|_i\rangle)$ that $(Z',\langle N,K|_i\rangle)$ is $(n-1)$-connected. The long exact sequence for the triple $(\langle N,K|_{[i,k]}\rangle,\langle W,K|_k\rangle,\langle N,K|_k\rangle)$ and excision

$$(\langle N,K|_{[i,k]}\rangle,\langle W,K|_k\rangle)\sim(Z',\langle M,K|_k\rangle)$$

shows that $(Z',\langle M,K|_k\rangle)$ is also $(n-1)$-connected. □

The argument above can *not* be improved to show that F_∞ sends each morphism in $\mathcal{C}$ to a weak homotopy equivalence, since the dashed maps we constructed in no sense preserve basepoints. The case $n=0$ gives rise to the following example from [McDS]: we have $F_\infty(\varnothing)\simeq\mathbb{Z}\times B\Sigma_\infty$ and the morphism $1\colon\varnothing\rightsquigarrow\varnothing$ given by a single point induces the shift map on Σ_∞, that is, the map induced by the self-embedding given by $\{1,2,\dots\}\cong\{2,3,\dots\}\hookrightarrow\{1,2,\dots\}$. This is not surjective, so the map is not a homotopy equivalence; it is however a homology equivalence, by the argument we have presented.

7.5. Proof of Theorem 1.8

We revert to the situation of Theorem 1.8: We are given a spherical $\theta: B \to BO(2n)$ and a universal θ-end (K, ℓ_K) with $\mathcal{N}^\theta(K|_0, \ell_K|_0) \neq \varnothing$. In order to apply the results of the previous sections, we shall produce a new structure $\theta': B' \to BO(2n)$ and a θ'-manifold L whose structure map $L \to B'$ is $(n-1)$-connected.

Let $\theta': B' \to B \xrightarrow{\theta} BO(2n)$ be obtained as the nth stage of the Moore–Postnikov tower of $\ell_K: K \to B$, and ℓ'_K be the θ'-structure on K given by the Moore–Postnikov factorisation.

LEMMA 7.16. *Let W be a manifold with boundary ∂W, and suppose that $(W, \partial_0 W)$ is $(n-1)$-connected, for a collection of boundary components $\partial_0 W \subset \partial W$. Let $\theta: B \to BO(2n)$ and $\theta': B' \to BO(2n)$ be two tangential structures and $f: B' \to B$ be a fibrewise map whose homotopy fibres are $(n-2)$-types. If $\ell'_{\partial_0 W}$ is a θ'-structure on $\partial_0 W$ with underlying θ-structure $\ell_{\partial_0 W}$, then*

$$\mathrm{Bun}_\partial(TW, (\theta')^*\gamma; \ell'_{\partial_0 W}) \longrightarrow \mathrm{Bun}_\partial(TW, \theta^*\gamma; \ell_{\partial_0 W}) \tag{7.3}$$

is a weak homotopy equivalence.

Proof. As the homotopy fibres of $f: B' \to B$ are $(n-2)$-types and $(W, \partial_0 W)$ is $(n-1)$-connected, the space of lifts

$$\begin{tikzcd} \partial_0 W \arrow[r, "\ell'_{\partial_0 W}"] \arrow[d, hook] & B' \arrow[d, "f"] \\ W \arrow[r, "\ell_W"'] \arrow[ur, dashed] & B \end{tikzcd}$$

is contractible, for each θ-structure ℓ_W on W restricting to $\ell_{\partial_0 W}$ on the boundary. But this space of lifts is easily identified with the homotopy fibre of the map (7.3) over the point ℓ_W. $\square$

By Proposition 7.8, the map $K \to B$ induces an injection in π_{n-1} and a surjection in π_n, so the nth and $(n-1)$-st stages of the Moore–Postnikov tower actually agree, and in particular the homotopy fibres of $B' \to B$ are $(n-2)$-types. The following results allow us to work with θ'-manifolds instead of θ-manifolds for many purposes.

COROLLARY 7.17. *The natural map induces a weak equivalence*

$$\mathcal{N}^{\theta'}(K|_i, \ell'_K|_i) \xrightarrow{\simeq} \mathcal{N}^\theta(K|_i, \ell_K|_i).$$

Proof. This follows from Lemma 7.16 in the case $\partial_0 W = \partial W$ by forming the homotopy orbit space by the action of $\mathrm{Diff}(W, \partial W)$ and taking disjoint union over all W with $\partial W = K|_i$ for which $(W, \partial W)$ is $(n-1)$-connected. $\square$

LEMMA 7.18. *The θ'-manifold (K,ℓ'_K) is a universal θ'-end.*

Proof. We verify the conditions of Definition 1.7. The cobordisms $K|_{[i-1,i]}$ are highly connected, as we have assumed that K is a universal θ-end. If $(W: K|_i \rightsquigarrow P, \ell'_W)$ is a highly connected θ'-cobordism, with underlying θ-structure ℓ_W, then by assumption there is an embedding $j: W \to K|_{[i,\infty)}$ and a homotopy $\ell_K \circ Dj \simeq \ell_W$, all relative to $K|_i$, but then by Lemma 7.16 there is also a homotopy $\ell'_K \circ Dj \simeq \ell'_W$ relative to $K|_i$. $\square$

Proof of Theorem 1.8. Recall that the theorem asserts a homology equivalence between the homotopy colimit of the direct system

$$\mathcal{N}^\theta(K|_0,\ell_K|_0) \xrightarrow{K|_{[0,1]}} \mathcal{N}^\theta(K|_1,\ell_K|_1) \xrightarrow{K|_{[1,2]}} \mathcal{N}^\theta(K|_2,\ell_K|_2) \xrightarrow{K|_{[2,3]}} \dots \tag{7.4}$$

and the infinite loop space $\Omega^\infty \mathrm{MT}\theta'$. By Corollary 7.17 and Lemma 7.18, it suffices to prove the theorem in the case $\theta=\theta'$, i.e. when $\ell_K: K \to B$ is n-connected. In order to apply Theorem 7.13, we first need to define a θ-manifold L (in order to have the category $\mathcal{C}$ defined). To do so, we pick a self-indexing Morse function $f: K|_0 \to [0,2n-1]$ and let $L=f^{-1}\big([0,n-\frac{1}{2}]\big)$. Then the inclusions $L \to K|_0$ and $K|_0 \to K$ are both $(n-1)$-connected, so the structure map $L \to B$ is $(n-1)$-connected and we have defined the category $\mathcal{C}$, satisfying Theorem 7.13. By Proposition 7.10 we may replace (K,ℓ_K) with any other universal θ-end without changing the homotopy type of the homotopy colimit (7.4), as long as $K|_0$ is unchanged, and by Corollary 7.12 there exists a universal θ-end of the form $K^\circ \cup ([0,\infty)\times L)$, where K° is a universal θ-end in $\mathcal{C}$. Now, by Proposition 7.5 the direct system (7.4) is homotopy equivalent to

$$\mathcal{C}(\bar{L}, K|_0^\circ) \xrightarrow{K|^\circ_{[0,1]}} \mathcal{C}(\bar{L}, K|_1^\circ) \xrightarrow{K|^\circ_{[1,2]}} \mathcal{C}(L, K|_2^\circ) \xrightarrow{K|^\circ_{[2,3]}} \dots .$$

By Theorem 7.13, the homotopy colimit is homology equivalent to $\Omega B\mathcal{C}$, which in turn is weakly equivalent to $\Omega^\infty \mathrm{MT}\theta = \Omega^\infty \mathrm{MT}\theta'$, by Theorem 7.3. $\square$

7.6. Proofs of Lemma 1.12 and Theorem 1.13

Let us first show that $\mathcal{K} \subset \mathcal{K}_0$ is a submonoid, and that it is commutative. Recall that $\mathcal{K}_0$ was the set of isomorphism classes of highly connected cobordisms $K \subset [0,1]\times\mathbb{R}^\infty$ with θ-structure, starting and ending at (P,ℓ_P), and that $\mathcal{K}$ is the subset admitting representatives containing $[0,1]\times(P\setminus A)$ with product θ-structure, where $A \subset P$ is a closed regular neighbourhood of a simplicial complex of dimension at most $n-1$ inside P. Let $K_0, K_1: P \rightsquigarrow P$ be two such cobordisms and let K_i have support in A_i, a regular neighbourhood of a simplicial complex X_i of dimension at most $n-1$. As P is $(2n-1)$-dimensional, we can perturb the X_i to be disjoint and then shrink the A_i so they are disjoint. But if K_0 and K_1 have support in the disjoint sets A_0 and A_1, then $K_0 \circ K_1$ has support in $A_0 \amalg A_1$ which is a regular neighbourhood of $X_0 \amalg X_1$ which is again a simplicial complex of dimension at most $n-1$. Furthermore $K_0 \circ K_1$ is isomorphic to the θ-bordism K_{01} which is supported in $A_0 \amalg A_1$ and agrees with K_i on $[0,1] \times A_i$, and this in turn is isomorphic to $K_1 \circ K_0$, so $\mathcal{K}$ is commutative.

Recall that we have a monoid map $\mathcal{K}' \to \mathcal{K}$, where $\mathcal{K}'$ is defined like $\mathcal{K}$, but using θ' instead of θ, where $\theta' : B' \to B \xrightarrow{\theta} B\mathrm{O}(2n)$ was defined by letting $\ell_P : P \to B' \to B$ be the $(n-1)$-st stage of the Moore–Postnikov tower. We saw in Corollary 7.17 that the map $\mathcal{N}^{\theta'}(P, \ell'_P) \to \mathcal{N}^{\theta}(P, \ell_P)$ is a weak equivalence, and we claim that an obstruction-theoretic argument similar to that of Lemma 7.16 shows that $\mathcal{K}' \to \mathcal{K}$ is an isomorphism. Explicitly, $[0,1] \times P$ has a canonical lift of its θ-structure to a θ'-structure. If an element of $\mathcal{K}$ is represented by a cobordism K supported in $A \subset P$, it contains the subset

$$(\{0\} \times P) \cup ([0,1] \times (P \setminus A)) \cup (\{1\} \times P)$$

which has a canonical θ'-structure. Because A is a regular neighbourhood of a simplicial complex of dimension at most $n-1$, the manifold K is obtained up to homotopy from this subset by attaching cells of dimension at least n, so up to homotopy there is a unique extension of the lift. This shows that $\mathcal{K}' \to \mathcal{K}$ is a bijection.

Before embarking on the proof of Theorem 1.13, we establish the following useful strengthening of assumption (iii) of that theorem.

LEMMA 7.19. *Let $[W] \in \mathcal{K}$ be such that each path component of W contains a submanifold diffeomorphic to $(S^n \times S^n) \setminus \mathrm{int}(D^{2n})$. Then each path component of $3W$, that is, the composition $W \circ W \circ W$, contains such a submanifold which in addition has nullhomotopic structure map to B.*

Proof. Let us suppose that W is path connected: otherwise we repeat the argument below for each path component. Finding an embedded $(S^n \times S^n) \setminus \mathrm{int}(D^{2n})$ is equivalent to finding two embedded n-spheres with trivial normal bundles, which intersect transversely at a single point. By assumption, this holds for W so we have

$$(S^n \times S^n) \setminus \mathrm{int}(D^{2n}) \hookrightarrow W \xrightarrow{\ell_W} B,$$

which in π_n induces a homomorphism $\mathbb{Z} \oplus \mathbb{Z} = \pi_n((S^n \times S^n) \setminus \mathrm{int}(D^{2n})) \to \pi_n(B)$, sending the basis elements to $x, y \in \pi_n(B)$.

In a separate copy of W we have a framed embedding

$$S^n \times \{*\} \overset{\text{reflection}}{\cong} S^n \times \{*\} \hookrightarrow (S^n \times S^n) \setminus \mathrm{int}(D^{2n}) \hookrightarrow W$$

which in $\pi_n(B)$ gives the element $-x$. Thus in $2W$, the connected sum of this embedded framed sphere and the original one gives an embedded framed sphere with null-homotopic map to B. Using the third copy of W we can fix the remaining sphere, without changing the property that the two spheres intersect transversely in one point. □

We shall first prove Theorem 1.13 under an additional countability hypothesis, namely we prove the following result.

PROPOSITION 7.20. *Let θ, (P, ℓ_P) and $\mathcal{K}$ be as in Theorem 1.13, and let $\mathcal{L} \subset \mathcal{K}$ be a submonoid satisfying conditions* (i)–(iii) *of that theorem. Assume in addition that $\mathcal{L}$ is countable. Then the induced morphism*

$$H_*(\mathcal{N}^\theta(P, \ell_P))[\mathcal{L}^{-1}] \longrightarrow H_*(\Omega^\infty \mathrm{MT}\theta')$$

is an isomorphism.

Proof. By countability of $\mathcal{L}$, we may pick a sequence of θ-manifolds $(K|_{[i,i+1]}, \ell_i)$ which are self-bordisms of (P, ℓ_P) representing elements of $\mathcal{L}$, in a way that each element of $\mathcal{L}$ is represented infinitely often. We then let K be the infinite composition of the $K|_{[i,i+1]}$, and deduce from Addendum 1.9 that (K, ℓ_K) is a universal θ-end. (That property (iii) of the addendum is satisfied follows from assumption (iii) and Lemma 7.19.) Then Theorem 1.8 gives a homology equivalence

$$\operatorname{hocolim} \mathcal{N}^\theta(P, \ell_P) \longrightarrow \Omega^\infty \mathrm{MT}\theta'',$$

where the homotopy colimit is over composition with the $K|_{[i,i+1]}$, and the map

$$\theta'' \colon B'' \longrightarrow B \xrightarrow{\theta} B\mathrm{O}(2n)$$

is constructed from the Moore–Postnikov n-stage of ℓ_K. Obstruction theory provides a map $B' \to B''$ over B and under P, and combining Addendum 1.9 and the $(n-1)$-connectedness of (K, P) shows that this map is a weak homotopy equivalence, so

$$\mathrm{MT}\theta'' \simeq \mathrm{MT}\theta'.$$

Taking homology turns the homotopy colimit into a colimit of the $\mathbb{Z}[\mathcal{L}]$-module $H_*(\mathcal{N}^\theta(P, \ell_P))$ over multiplying with elements of $\mathcal{L}$, each element occurring infinitely many times. But that precisely calculates the localisation at $\mathcal{L}$, as $\mathcal{L}$ is commutative so the localisation may be computed by right fractions. □

The proposition above proves Theorem 1.13 in the case where $\mathcal{K}$ is countable. (To apply Proposition 7.20 with $\mathcal{L} = \mathcal{K}$, we need to check that conditions (i)–(iii) hold. This is proved using the manifolds W_α from the proof of Proposition 7.8.) We will deduce the general case by a colimit argument, based on the following result.

Corollary 7.21. *Let θ, (P,ℓ_P) and $\mathcal{K}$ be as in Theorem 1.13, and let $\mathcal{L}\subset\mathcal{K}$ be a submonoid satisfying conditions* (ii) *and* (iii) *of that theorem, but not necessarily* (i). *Assume in addition that $\mathcal{L}$ is countable. Then the induced morphism*

$$H_*(\mathcal{N}^{\theta_\mathcal{L}}(P,\ell_P^\mathcal{L}))[\mathcal{L}^{-1}] \longrightarrow H_*(\Omega^\infty \mathrm{MT}\theta_\mathcal{L}) \tag{7.5}$$

is an isomorphism, where $\theta_\mathcal{L}\colon B_\mathcal{L}\to B\mathrm{O}(2n)$ is obtained as the n-th Moore–Postnikov factorisation of a certain map $\ell\colon X_\mathcal{L}\to B$, defined as follows. Each self-bordism (L,ℓ_L) representing an element of $\mathcal{L}$ has incoming boundary $P\subset K$, and we let $X_\mathcal{L}$ be obtained by gluing every such L along their common incoming boundary; *the structure maps ℓ_L then glue to the map $\ell\colon X_\mathcal{L}\to B$.*

Proof. We will explain how this is an instance of Proposition 7.20, with θ replaced by $\theta_\mathcal{L}$. The map $\ell_P\colon P\to B$ lifts canonically to $\ell_P^\mathcal{L}\colon P\to B_\mathcal{L}$, by restricting the first map in the factorisation $X_\mathcal{L}\to B_\mathcal{L}\to B$, but for the statement to make sense we should explain how to regard $\mathcal{L}$ as a submonoid of $\mathcal{K}^\mathcal{L}$, where $\mathcal{K}^\mathcal{L}$ is defined like $\mathcal{K}$, but with $\theta_\mathcal{L}$ in place of θ. In fact, the same obstruction-theoretic argument as we used in the beginning of §7.6 to show bijectivity of $\mathcal{K}'\to\mathcal{K}$ will show that the natural map $\mathcal{K}^\mathcal{L}\to\mathcal{K}$ is injective and that the image contains $\mathcal{L}$: by the assumption on support of representatives of elements of $\mathcal{K}$ and $\mathcal{K}^\mathcal{L}$, we will only ever need to lift structure maps over cells of dimension n or higher; the space of relative lifts over an n-cell is either contractible or empty, and over higher cells it is always contractible.

As each element of $\mathcal{L}$ may canonically be given the structure of a self θ'-cobordism of (P,ℓ'_P), the two relevant Moore–Postnikov factorisations fit into the following diagram

$$\begin{tikzcd}
X_\mathcal{L} \arrow[r] & B_\mathcal{L} \arrow[r] \arrow[d] & B \arrow[d, equal] \\
P \arrow[u] \arrow[r] & B' \arrow[r] & B,
\end{tikzcd}$$

where the top row is an nth stage factorisation and the bottom row is an $(n-1)$-st stage. The inclusion $P\to X_\mathcal{L}$ is $(n-1)$-connected, and it follows formally that the map $B_\mathcal{L}\to B'$ induces an injection in homotopy groups except possibly π_{n-1} and a surjection in homotopy groups except possibly π_n. But the assumed property (ii) ensures that in fact $\pi_{n-1}(B_\mathcal{L})\to\pi_{n-1}(B')$ is an isomorphism, and thus $\pi_*(B_\mathcal{L})\to\pi_*(B')$ is injective in degree n and an isomorphism in all other degrees. In π_n, we claim that

$$\mathrm{Im}(\pi_n(X_\mathcal{L})\to\pi_n(B'))=\mathrm{Im}\biggl(\bigoplus_{[K]\in\mathcal{L}}\pi_n(K)\to\pi_n(B')\biggr). \tag{7.6}$$

Indeed, both sides are subgroups of $\pi_n(B')$ containing the image of $\pi_n(P)\to\pi_n(B')$, so it suffices to prove that the subgroups have the same image in $\pi_n(B',P)$, but this follows from the homotopy excision isomorphism

$$\bigoplus_{[K]\in\mathcal{L}} \pi_n(K,P) \xrightarrow{\sim} \pi_n(X_{\mathcal{L}},P),$$

which may be proved using the Blakers–Massey theorem.

The equality (7.6) shows that the monoid $\mathcal{L}\subset\mathcal{K}^{\mathcal{L}}$ satisfies the assumption (i) of Proposition 7.20 with respect to $\theta_{\mathcal{L}}$, and by construction it also satisfies (ii) and (iii). In the conclusion from the proposition we should put the $(n-1)$-st Moore–Postnikov stage of $P\to B_{\mathcal{L}}$, but this map is already $(n-1)$-connected so the factorisation is equivalent to $P\to B_{\mathcal{L}}=B_{\mathcal{L}}$. □

Proof of Theorem 1.13. Let $\mathcal{L}\subset\mathcal{K}$ satisfy the conditions of Theorem 1.13. Then for each countable submonoid $\mathcal{L}'\subset\mathcal{L}$, there are maps $B_{\mathcal{L}'}\to B$ from Corollary 7.21, and if $\mathcal{L}''\subset\mathcal{L}'$ is a submonoid we also have a factorisation $B_{\mathcal{L}''}\to B_{\mathcal{L}'}\to B$, as our description of $B_{\mathcal{L}}$ is strictly functorial in the monoid $\mathcal{L}$ (using a functorial model for Moore–Postnikov factorisation, cf. §1.1.3). Therefore we may form the colimit of the isomorphisms (7.5) over the poset of countable submonoids $\mathcal{L}'\subset\mathcal{L}$.

Using the manifolds W_α constructed in the proof of Proposition 7.8 and the observation that the poset of countable submonoids is filtered, it is easy to see that the natural map from the homotopy colimit of the $B_{\mathcal{L}'}$ to B' is a weak equivalence, and hence the homotopy colimit of the $\Omega^\infty \mathrm{MT}\theta_{\mathcal{L}'}$ is $\Omega^\infty \mathrm{MT}\theta'$. If we can prove that the homotopy colimit of the spaces $\mathcal{N}^{\theta_{\mathcal{L}'}}(P,\ell_P^{\mathcal{L}'})$ is $\mathcal{N}^{\theta'}(P,\ell'_P)\simeq\mathcal{N}^{\theta}(P,\ell_P)$, Theorem 1.13 will therefore follow as the direct limit of the isomorphism (7.5).

We saw in Lemma 7.16 that the map $\mathcal{N}^{\theta'}(P,\ell'_P)\to\mathcal{N}^{\theta}(P,\ell_P)$ is a weak equivalence. A similar obstruction-theoretic argument as in that lemma shows that the map $\mathcal{N}^{\theta_{\mathcal{L}'}}(P,\ell_P^{\mathcal{L}'})\to\mathcal{N}^{\theta}(P,\ell_P)$ induces an injection on π_0 and a weak equivalence of each path component. Namely, $\mathcal{N}^{\theta_{\mathcal{L}'}}(P,\ell_P^{\mathcal{L}'})$ is up to homotopy a disjoint union of path components of $\mathcal{N}^{\theta}(P,\ell_P)$; the component containing (W,ℓ_W) is included precisely when ℓ_W admits a lift to a $\theta_{\mathcal{L}'}$-structure. Up to homotopy, the system of spaces $\mathcal{N}^{\theta_{\mathcal{L}'}}(P,\ell_P^{\mathcal{L}'})$ therefore just consists of including more and more components of $\mathcal{N}^{\theta}(P,\ell_P)$, including all of them in the colimit. □

References

[BM] Berglund, A. & Madsen, I., Homological stability of diffeomorphism groups. *Pure Appl. Math. Q.*, 9 (2013), 1–48.

[BiFi] Binz, E. & Fischer, H. R., The manifold of embeddings of a closed manifold, in *Differential Geometric Methods in Mathematical Physics* (Clausthal-Zellerfeld, 1978), Lecture Notes in Phys., 139, pp. 310–329. Springer, Berlin–New York, 1981.

[BoFr] Bousfield, A. K. & Friedlander, E. M., Homotopy theory of Γ-spaces, spectra, and bisimplicial sets, in *Geometric Applications of Homotopy Theory* (Evanston, IL, 1977), Lecture Notes in Math., 658, Part II, pp. 80–130. Springer, Berlin–Heidelberg, 1978.

[C] Cerf, J., Topologie de certains espaces de plongements. *Bull. Soc. Math. France*, 89 (1961), 227–380.

[CM] Cohen, R. L. & Madsen, I., Surfaces in a background space and the homology of mapping class groups, in *Algebraic Geometry* (Seattle, 2005), Proc. Sympos. Pure Math., 80, Part 1, pp. 43–76. Amer. Math. Soc., Providence, RI, 2009.

[D] Dold, A., Partitions of unity in the theory of fibrations. *Ann. of Math.*, 78 (1963), 223–255.

[E1] Ebert, J., Algebraic independence of generalized MMM-classes. *Algebr. Geom. Topol.*, 11 (2011), 69–105.

[E2] — A vanishing theorem for characteristic classes of odd-dimensional manifold bundles. *J. Reine Angew. Math.*, 684 (2013), 1–29.

[EZ] Eilenberg, S. & Zilber, J. A., Semi-simplicial complexes and singular homology. *Ann. of Math.*, 51 (1950), 499–513.

[GMTW] Galatius, S., Madsen, I., Tillmann, U. & Weiss, M., The homotopy type of the cobordism category. *Acta Math.*, 202 (2009), 195–239.

[GRW1] Galatius, S. & Randal-Williams, O., Monoids of moduli spaces of manifolds. *Geom. Topol.*, 14 (2010), 1243–1302.

[GRW2] — Homological stability for moduli spaces of high dimensional manifolds. Preprint, 2012. arXiv:1203.6830 [AT].

[GRW3] — Detecting and realising characteristic classes of manifold bundles. To appear in *Algebraic Topology: Applications and New Directions* (Stanford, 2012), Contemp. Math., Amer. Math. Soc., Providence, RI.

[G] Geoghegan, R., *Topological Methods in Group Theory*. Graduate Texts in Mathematics, 243. Springer, New York, 2008.

[H] Hatcher, A., *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.

[Ke] Kervaire, M. A., Le théorème de Barden–Mazur–Stallings. *Comment. Math. Helv.*, 40 (1965), 31–42.

[Kr] Kreck, M., Surgery and duality. *Ann. of Math.*, 149 (1999), 707–754.

[Le] Levine, J., Knot modules. I. *Trans. Amer. Math. Soc.*, 229 (1977), 1–50.

[Li] Lima, E. L., On the local triviality of the restriction map for embeddings. *Comment. Math. Helv.*, 38 (1964), 163–164.

[Lü] Lück, W., A basic introduction to surgery theory, in *Topology of High-Dimensional Manifolds* (Trieste, 2001), ICTP Lecture Notes, 9, Part 1, pp. 1–224. Abdus Salam Int. Cent. Theoret. Phys., Trieste, 2002.

[MT] Madsen, I. & Tillmann, U., The stable mapping class group and $Q(\mathbb{C}\mathrm{P}^\infty_+)$. *Invent. Math.*, 145 (2001), 509–544.

[MW] Madsen, I. & Weiss, M., The stable moduli space of Riemann surfaces: Mumford's conjecture. *Ann. of Math.*, 165 (2007), 843–941.

[Ma] May, J. P., *The Geometry of Iterated Loop Spaces*. Lecture Notes in Math., 271. Springer, Berlin–Heidelberg, 1972.

[McDS] McDuff, D. & Segal, G., Homology fibrations and the "group-completion" theorem. *Invent. Math.*, 31 (1975/76), 279–284.

[Mi1] Milnor, J., *Morse Theory*. Annals of Mathematics Studies, 51. Princeton University Press, Princeton, NJ, 1963.

[Mi2] — *Lectures on the h-Cobordism Theorem*. Princeton University Press, Princeton, NJ, 1965.

[Mu] Mumford, D., Towards an enumerative geometry of the moduli space of curves, in *Arithmetic and Geometry*, Progr. Math., 36, Vol. II, pp. 271–328. Birkhäuser, Boston, MA, 1983.

[S1] Segal, G., Classifying spaces and spectral sequences. *Inst. Hautes Études Sci. Publ. Math.*, 34 (1968), 105–112.

[S2] — Categories and cohomology theories. *Topology*, 13 (1974), 293–312.

[Wa] Wall, C. T. C., *Surgery on Compact Manifolds*. London Mathematical Society Monographs, 1. Academic Press, London–New York, 1970.

[We] Weiss, M., What does the classifying space of a category classify? *Homology Homotopy Appl.*, 7 (2005), 185–195.

Søren Galatius
Department of Mathematics
Stanford University
Stanford, CA 94305
U.S.A.
galatius@stanford.edu

Oscar Randal-Williams
Department of Pure Mathematics
and Mathematical Statistics
University of Cambridge
Wilberforce Road
Cambridge CB3 0WB
U.K.
o.randal-williams@dpmms.cam.ac.uk

Received September 19, 2012
Received in revised form November 12, 2013