Geom Dedicata (2010) 148:71–101
DOI 10.1007/s10711-009-9405-y

ORIGINAL PAPER

Epsilon–delta surgery over $\mathbb{Z}$

Steven C. Ferry

Received: 4 December 2008 / Accepted: 28 July 2009 / Published online: 19 August 2009

Abstract The purpose of this paper is to prove a controlled surgery exact sequence, including a stability theorem, as used in the construction of exotic homology manifolds. The approach is to show that this result is a formal consequence of the Chapman–Ferry Alpha-approximation theorem.

Keywords Controlled surgery · Homology manifold

Mathematics Subject Classification (2000) 57N65 · 57R65 · 57R67

1 Introduction

Let $p{:}X \to B$ be a map. We will say that a map $f{:}Y \to X$ is an ϵ-equivalence over B if there exist a map $g{:}X \to Y$ and homotopies $h_t{:}f \circ g \sim \text{id}$, $k_t{:}g \circ f \sim \text{id}$ so that the tracks $p \circ h_t(x)$ and $p \circ f \circ k_t(y)$ have diameter $<\epsilon$ for all $x \in X$ and $y \in Y$. Let $(M, \partial M)$ be a manifold. If $p{:}M \to B$ is a map, an ϵ-structure on $(M, \partial M)$ over B will mean an equivalence class of pairs (N, f), where $f{:}(N, \partial N) \to (M, \partial M)$ is an ϵ-equivalence over B which restricts to a homeomorphism of the boundaries. Pairs (N, f) and (N', f') are said to be ϵ*-related* if there is an homeomorphism $\phi{:}N \to N'$ so that $f' \circ \phi$ is ϵ-homotopic to f over B rel boundary. Our notion of equivalence for ϵ-structures is the equivalence relation generated by this relation. We will use the symbol $\mathcal{S}'_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$ to denote the collection of equivalence classes of ϵ-structures on M. The purpose of this paper is to prove an $\epsilon - \delta$ surgery exact sequence.

S. C. Ferry (✉)
Rutgers University, New Brunswick, NJ, USA
e-mail: sferry@math.rutgers.edu

Theorem 1.1 *If M^n is a compact topological manifold, $n \geq 6$, or $n \geq 5$ when $\partial M = \emptyset$, B is a finite polyhedron with the standard metric,*[1] *and $p : M \to B$ is a UV^1 map*[2]*, then there exist an $\epsilon_0 > 0$ and a $T > 0$ depending only on n and B so that for every $\epsilon \leq \epsilon_0$ there is a surgery exact sequence*

$$\dots H_{n+1}(B;\mathbb{L}) \dashrightarrow \mathcal{S}_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix} \longrightarrow [M,\ \partial M;\ \mathrm{G/TOP}] \longrightarrow H_n(B;\mathbb{L})$$

where $\mathbb{L}$ is the periodic L-spectrum of the trivial group and

$$\mathcal{S}_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix} = \mathrm{im}\left(\mathcal{S}'_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix} \to \mathcal{S}'_{T\epsilon}\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}\right).$$

Moreover, for $\epsilon \leq \epsilon_0$, $\mathcal{S}_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix} \cong \mathcal{S}_{\epsilon_0}\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$.

Definition 1.2 A connected locally finite polyhedron B with a metric bilipschitz equivalent to the standard metric is said to have bounded geometry if there is a (finite) bound on the number of vertices of B adjacent to a given vertex. We will call this bound the complexity of B. We will say that a topological manifold M has bounded geometry if there is an $L > 0$ so that M has a handle decomposition so that each handle H is L-bilipschitz equivalent to the standard handle of that dimension and index. Classically, a smooth manifold M is said to have bounded geometry if it has bounded sectional curvature and if there is a lower bound on the injectivity radius. Work of Cheeger et al. [7] shows that the classical definition implies the definition that we have given in the topological category.

Remark 1.3

(i) If a polyhedron of bounded geometry is given the standard metric, then up to isometries, there are only finitely many possibilities for vertex stars.
(ii) We can go further and declare a finite-dimensional ANR to have bounded geometry if it has a manifold mapping cylinder neighborhood $M(r) \to M$ so that $M(r)$ has bounded geometry as a topological manifold and so that the projection map $M(r) \to M$ is uniformly continuous. This level of generality is not needed in the present paper.

Addendum 1.4

(i) This surgery sequence is also valid in the smooth and PL categories. The point is that the PL and smooth surgery groups are the same as the TOP surgery groups, so the argument in Theorem 10.2 gives "squeezing" and a surgery exact sequence in those categories, as well. In the smooth and PL categories, $[M, \mathrm{G/TOP}]$ should be replaced by $[M, \mathrm{G/CAT}]$.
(ii) Theorem 1.1 is true as stated, i.e. with a linear relation between ϵ and $T\epsilon$, for B a polyhedron of bounded geometry with the standard metric provided that we use locally finite homology $H_n^{\ell f}(B;\mathbb{L})$ at the appropriate spots in the surgery sequence. Of course, M will be a noncompact manifold in this case. We will state our results for finite polyhedra and use addenda to discuss the extension to the bounded geometry case. The special case $B = P \times \mathbb{R}$, where P is a finite polyhedron, is used in the proof of Theorem 1.1 for finite polyhedra. The use of the standard metric is for definiteness. Theorem 1.1 remains true as stated in any metric bilipschitz equivalent to the standard metric.

[1] By the standard k-simplex, we mean the convex hull of unit vectors $e_0, \dots e_k$ in $\mathbb{R}^{k+1}$. By the "standard metric" on a simplicial complex, we mean the path metric in which each k simplex is isometric to the standard k-simplex and simplices are glued together via isometries on faces. See Sect. 1 part 7 of [2] for details.

[2] See Definition 4.1 below.

(iii) Theorem 1.1 is true for B a compact ANR, except that in that case we lose the linear dependence in the definition of $\mathcal{S}_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$. For B a compact ANR, the theorem should say that there is an $\epsilon_0 > 0$ so that for every $\epsilon \le \epsilon_0$ there is a $\delta > 0$ so that the surgery exact sequence is true with $\mathcal{S}_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix} = \mathrm{im}\left(\mathcal{S}'_\delta\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix} \to \mathcal{S}'_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}\right)$ and that for $\epsilon \le \epsilon_0$, $\mathcal{S}_\epsilon\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix} \cong \mathcal{S}_{\epsilon_0}\begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$.

The proof of Theorem 1.1 is quite easy in principle. The first basic slogan is that a π–π theorem lets you set up a surgery theory. It turns out to be straightforward to adapt the proof of the bounded π–π theorem of [10] to give an "epsilon–delta" π–π theorem. The result is a "pro" surgery theory. A second well-known slogan is that "once you understand one manifold with a given fundamental group well, you know a lot about all manifolds with that fundamental group." The "Alexander trick" of Sect. 9 of this paper shows that the rel boundary structure set of the projection map $N(B) \to B$ sending a regular neighborhood of a polyhedron B to B has trivial structure set. This uses the alpha approximation theorem of [6] and it allows us to compute the controlled surgery groups over B with $\mathbb{Z}$ coefficients. This is analogous to Sullivan's use of the generalized Poincaré conjecture to compute the homotopy groups of G/TOP, but in this case we reversed the process to compute the controlled surgery groups. Once the surgery groups are known to be stable, the stability for the structure set follows from a form of the five lemma. The rest of Theorem 1.1 follows immediately.

The most confusing part of this program is the "pro surgery theory," but an excellent model for this construction can be found in Chapman's development of a very general theory of controlled Whitehead torsion in [5]. Recapitulating, our basic plan for proving Theorem 1.1 was to combine the approaches of [5] and [10] to prove an ϵ–δ π–π theorem and use it to give a formal "chapter 9" development of a pro-surgery theory. We then computed the high-dimensional surgery groups by plugging the Chapman–Ferry alpha-approximation theorem into this theory. The lower-dimensional groups were then computed using a 4-periodic algebraic description of the groups. The stability of the surgery groups was then used to deduce the stability of the structure set in all dimensions. This version of the proof was presented in a series of five lectures at Notre Dame University in May 2002. The author would like to thank the topologists of Notre Dame for their hospitality during a pleasant visit.

This plan worked, but it turns out to have been overly elaborate. It is considerably simpler to work directly with the algebraically defined surgery groups. Applying the alpha approximation theorem shows the stability of the algebraic system in high dimensions and periodicity extends this immediately to all dimensions. One proceeds as before to prove the stability of the structure sets—again, this is basically a "pro" form of the five lemma. The surgery sequence of Theorem 1.1 then follows as usual.

The author apologizes for being slow in writing up this paper.[3] It does involve quite a lot of writing for a relatively straightforward result and the author hoped for some time to find a quick way of proving this theorem as a consequence of [10]. Sadly, this included an attempt called "Squeezing structures" which turned out to contain an error.

Here is the statement of the alpha-approximation theorem:

Theorem 1.5 (alpha-approximation theorem [6])

(i) *Let M^n be a closed topological manifold of dimension ≥ 5. For every $\epsilon > 0$ there is a $\delta > 0$ so that if $f : N \to M$ is a δ-homotopy equivalence from another manifold of the same dimension to M, then f is ϵ-homotopic to a homeomorphism.*

[3] A competing presentation of this material has appeared in [16]. The author disagrees with the historical statements in the introduction to that paper.

(ii) *If M^n is a compact topological manifold, then for every $\epsilon > 0$ there is a $\delta > 0$ so that if $f : N \to M$ is a δ-homotopy equivalence from a compact manifold of the same dimension to M such that $f | \partial N$ is a homeomorphism from ∂N to ∂M, then f is ϵ-homotopic rel boundary to a homeomorphism.*

Remark 1.6

(i) Since the manifolds in the theorem above are topological, the metric in which ϵ and δ are measured is merely a topological metric. The proof of the alpha-approximation theorem is a handle induction, so the size of δ is some fraction of the size of the smallest handle in a handle decomposition of M. As noted by Farrell and Jones, this means that the relation between ϵ and δ is linear in any metric which allows handle decompositions to be subdivided linearly. In particular, the relation is linear for PL manifolds with the standard metric. We will always assume that polyhedra in this paper have been given metrics which are Lipschitz equivalent to standard metrics.

(ii) The full statement of the theorem in [6] is valid for noncompact manifolds and uses open covers α and β in place of ϵ and δ, hence the name. The proof of the alpha approximation theorem is a handle induction, so the theorem is true in its original form—with ϵ's and δ's for manifolds of bounded geometry. One only has to deal with finitely many isomorphism type of handles in a PL manifold of bounded geometry.

(iii) The alpha-approximation theorem is valid in dimension 4 as an easy consequence of work of Freedman and Quinn. See [11].

(iv) The alpha-approximation theorem is a purely topological theorem. It is false in the PL and smooth categories.

2 Geometric algebra

One of the main ideas in proving the thin h-cobordism theorem and related results is to do ordinary algebraic topology while keeping track of the sizes of various homotopies and chain homotopy equivalences. To facilitate this, we follow [8,17,22], and introduce the language of geometric modules. One reason that geometric chain modules turn out to be useful is that in certain situations they allow us to use homological data to construct homotopy equivalences of non simply connected CW complexes without passing to the universal cover and/or dealing with modules over the group ring. The general strategy is that if we can keep our cell manipulations localized, then the loops that arise in our constructions will all bound small disks which can be found without invoking the universal cover. See Proposition 5.3 below.

Definition 2.1

(i) A geometric $\mathbb{Z}$-module on a space E is a free module $\mathbb{Z}[S]$ on a set S together with a map $f : S \to E$. In this paper, S will always be locally finite over E. We will often suppress the function f and pretend that the elements of S are points of E. Geometric Λ-modules can be defined similarly for any ring Λ.

(ii) A geometric morphism $h : \mathbb{Z}[S] \to \mathbb{Z}[T]$ of geometric $\mathbb{Z}$-modules with $f : S \to E$ and $g : T \to E$ is a homomorphism $\mathbb{Z}[S] \to \mathbb{Z}[T]$. If we write $h = (h_{st})$ with respect to the bases S and T, then the radius of h is $\sup_{s,t}\{d(f(s), g(t)) | h_{st} \neq 0\}$. This is less general than the definition from [22], but it will suffice for the purposes of this paper.

(iii) A geometric morphism $h : \mathbb{Z}[S] \to \mathbb{Z}[T]$ of radius ϵ is an ϵ-isomorphism if there is a geometric morphism k of radius ϵ, $k : \mathbb{Z}[T] \to \mathbb{Z}[S]$, so that $h \circ k = \text{id}$ and $k \circ h = \text{id}$.

Definition 2.2

(i) A geometric $\mathbb{Z}$-module chain complex C on E is a sequence of morphisms of geometric $\mathbb{Z}$-modules on E

$$C: \cdots \longrightarrow C_i \xrightarrow{d_i} C_{i-1} \xrightarrow{d_{i-1}} \cdots$$

such that $d \circ d = 0$. C will be called an ϵ-chain complex on E if each morphism d_i has radius $<\epsilon$.

(ii) A *chain map* f between geometric $\mathbb{Z}$-module chain complexes C and D is a sequence of geometric morphisms $f_i : C_i \to D_i$ so that $d_{i-1} \circ f_i = f_{i-1} \circ d_i$. The map f has *radius* ϵ if each f_i has radius ϵ.

(iii) A chain homotopy between two geometric chain maps f and g is a collection $\{H_i\}$ of geometric morphisms $H_i : C_i \to D_{i+1}$ so that $d_{i+1} \circ H_i + H_{i-1} \circ d_i = f_i - g_i$. The *radius* of H is its radius as a geometric morphism.

(iv) An ϵ-chain contraction $s : C_* \to C_*$ is an ϵ-chain homotopy between id and 0.

Example 2.3 A good example of a geometric $\mathbb{Z}$-module chain complex to keep in mind arises when X is a finite polyhedron. The simplicial k-chains on X form a geometric module $C_k(X)$ where each k-simplex σ is associated to its barycenter $\widehat{\sigma} \in X$. If the simplices in the subdivision have diameter $<\epsilon$, then the boundary map $\partial : C_k(X) \to C_{k-1}(X)$ has radius $<\epsilon$. It is often useful to extend this example in the following manner: Let $p : K \to B$ be a map from a finite polyhedron to a compact metric space. The simplicial chains of K give rise to a chain complex of geometric modules over B by associating each simplex with the image of its barycenter in B.

The biggest change in moving to this "controlled" or "epsilon–delta" world from ordinary algebraic topology is that we no longer have kernels or quotients, the problem being that except in very restricted circumstances it is difficult to assign a position in the underlying space to an element of a kernel or a quotient. Happily, other standard constructions of algebraic topology carry over to this situation without difficulty. In particular, we have the notions of the algebraic mapping cone[4] and algebraic mapping cylinder of a morphism of geometric chain complexes.

Definition 2.4 If E and F are geometric chain complexes over a space B, and $f : E \to F$ is a chain map, then

(i) The *algebraic mapping cone* of f is the chain complex $C(f)_k = E_{k-1} \oplus F_k$ with boundary map given by $\begin{pmatrix} \partial_E & 0 \\ (-1)^{\deg f} f & \partial_F \end{pmatrix}$.

(ii) The *algebraic mapping cylinder* of f is the chain complex $M(f)_k = E_k \oplus E_{k-1} \oplus F_k$ with boundary map given by $\begin{pmatrix} -\partial_E & (-1)^k \,\mathrm{id}_E & 0 \\ 0 & \partial_E & 0 \\ 0 & (-1)^k f & \partial_F \end{pmatrix}$.

If f is a δ-morphism, then $C(f)_*$ is a δ chain complex.

Since we do not have kernels and cokernels available in this setting, we do not have homology groups, so we must find a new proof that showing the contractibility of the mapping cone of a map f is equivalent to showing that f is a chain homotopy equivalence.

[4] It's a good thing that the algebraic mapping cone has no cone point.

Proposition 2.5 *Let E and F be geometric chain complexes over a space B, and let $f\colon E \to F$ be a δ-chain map.*

(i) *If $C(f)_*$ is δ-chain contractible, then f is a δ-chain homotopy equivalence.*
(ii) *If f is a δ-chain homotopy equivalence, then $C(f)_*$ is $k_0\delta$-chain contractible, where k_0 is a fixed integer.*

Proof The first part is a straightforward computation using the definitions. The second part amounts to showing that if f is a δ-chain homotopy equivalence, then there is a retraction $r\colon M(f)_* \to E_*$ which is $k_0\delta$-chain homotopic to the identity rel E_*. See Proposition 1.1 of [20], which gives an explicit formula for a controlled chain contraction. The analogous statement for spaces is proven in [9] using "mapping cylinder calculus". □

Here is an algebraic version of handle sliding or handle addition, as it is sometimes called. We will use this operation frequently to modify chain complexes.

Lemma 2.6 (handle sliding) *Given a δ-chain complex*

$$\cdots \to C_{n+1} \xrightarrow{\binom{a}{b}} C'_n \oplus C''_n \xrightarrow{(c\ d)} C_{n-1}$$

and a δ-morphism $s\colon C'_n \to C''_n$, there is a 2δ-isomorphism to the 2δ-chain complex

$$\cdots \to C_{n+1} \xrightarrow{\binom{a}{b-sa}} C'_n \oplus C''_n \xrightarrow{(c+ds\ d)} C_{n-1}$$

Proof The isomorphism is given by identity maps on the ends and $\begin{pmatrix} \mathrm{id} & 0 \\ -s & \mathrm{id} \end{pmatrix}$ in the middle. Notice that the slide performed a block row operation on the first boundary map and a compensating block column operation on the second one. If $E = \begin{pmatrix} \mathrm{id} & 0 \\ s & \mathrm{id} \end{pmatrix}$ in the discussion above, then the new boundary map on the left is $E^{-1}\partial$ and the new boundary map on the right is ∂E. □

Another very useful construction is cancellation of cells.

Lemma 2.7 *There is an integer k so that if a portion of a δ-chain complex looks like*

$$\cdots \to A \to B \oplus C \to D \oplus C' \to \cdots$$

with the composite

$$C \to B \oplus C \to D \oplus C' \to C'$$

a δ-isomorphism then the chain complex is $k\delta$-chain homotopy equivalent to a $k\delta$-chain complex

$$\cdots \to A \to B \to D \to \cdots$$

Proof Let the boundary map $B \oplus C \to D \oplus C'$ be given by $\begin{pmatrix} \alpha & \beta \\ \eta & \gamma \end{pmatrix}$. The map $\gamma\colon C \to C'$ is a δ-isomorphism, so an elementary column operation followed by an elementary row operation reduces the boundary map to $\begin{pmatrix} \alpha' & 0 \\ 0 & \gamma \end{pmatrix}$. Performing appropriate handle slides in $B \oplus C$ and

$D \oplus C'$ produces an isomorphic chain complex with boundary map of this diagonal form. It is now a simple matter to show that this chain complex is $k\delta$-chain homotopy equivalent to one of the form $\cdots \to A \to B \to D \to \cdots$.

If the original chain complex has the form

$$\cdots \longrightarrow A \xrightarrow{\binom{\sigma}{\tau}} B \oplus C \xrightarrow{\begin{pmatrix} \alpha & \beta \\ \eta & \gamma \end{pmatrix}} D \oplus C' \xrightarrow{(\lambda\ \mu)} E \longrightarrow \cdots$$

then the new complex has the form

$$\cdots \longrightarrow A \xrightarrow{\ \sigma\ } B \xrightarrow{\alpha - \beta\gamma^{-1}\eta} D \xrightarrow{\ \lambda\ } E \longrightarrow \cdots$$

In particular, the collapsed complex is what it "should" be if either $\beta{:}C \to D$ or $\eta{:}B \to C'$ is zero, i.e., if the cells in C are attached only to cells in C' or if no cells in B are attached to cells in C'. The condition on β allows us, for instance, to collapse a cone to a subcone, while the condition on η is to familiar "free face" condition from PL topology. □

Remark 2.8 The sophisticated reader may wonder where the Whitehead group has gone in this discussion. The K-theory vanishing theorem stated in the next section will show that the hypothesis that the isomorphism $C \to C'$ be simple is unnecessary when working with geometric $\mathbb{Z}$-modules.

Lemma 2.9 *If A is a geometric chain complex and $C \subset A$ is a geometric chain complex which is δ-contractible, then A is $k\delta$-chain homotopy equivalent to $(A - C)$, the chain complex obtained by deleting generators of C. The new complex $A - C$ is a δ-chain complex. Here, $k = k(n)$ is a fixed integer, where n is the dimension of C.*

Proof C is δ-chain homotopy equivalent to $\mathrm{Cone}(C)$ rel C, so A is δ-chain homotopy equivalent to the union of A with $\mathrm{Cone}(C)$ along C. Cancelling the cells of $\mathrm{Cone}(C)$ starting from the top dimension of $\mathrm{Cone}(C)$ gives the desired equivalence. Since the higher-dimensional cells in each collapse attach only into C, the boundary maps on $A - C$ are unchanged by this process. □

Here is an algebraic cell-trading lemma. It involves introducing cells, adding cells, and cancelling cells, the final result being that n-cells are "traded for" $(n+2)$-cells.

Lemma 2.10 *Suppose given an ϵ-chain complex decomposed as modules as $B_\# \oplus A_\#$ for which the boundary map has the form*

$$\begin{array}{ccc} B_\# & \oplus & A_\# \\ \downarrow & \swarrow & \downarrow \\ B_{\#-1} & \oplus & A_{\#-1} \end{array}$$

If there is an ϵ-chain homotopy s, with $(s|B_\#) = 0$, from the identity to a morphism which is 0 *on $A_\#$ for $\# < \ell$, then $B_\# \oplus A_\#$ is chain-homotopy equivalent to $B_\# \oplus A'_\#$ where $A'_\# = 0$ for $\# < \ell$ and $A'_\# = A_\#$ for $\# \geq \ell + 2$. The new chain complex is a $3^\ell\epsilon$-chain complex with a $3^\ell\epsilon$-contraction.*

Proof First introduce cancelling 1- and 2-cells corresponding to A_0 to obtain

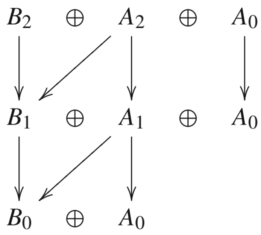

Now perform a handle slide using $s|{:}A_0 \to A_1$ to obtain

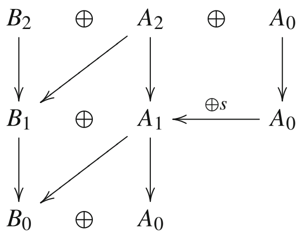

The lower map from A_0 to A_0 is the identity, so the lower copies of A_0 may be canceled to obtain

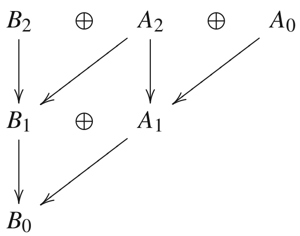

Repeat this process, and define $A'_\#$ so that $B_\# \oplus A'_\#$ is the resulting chain complex. □

Proposition 2.11 *If we have a commuting diagram:*

$$\begin{array}{ccc} A_* & \xrightarrow{\alpha} & B_* \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} \\ A'_* & \xrightarrow{\beta} & B'_* \end{array}$$

of δ-chain complexes, then there is a commuting diagram

$$\begin{array}{ccccc} A_* & \xrightarrow{\alpha} & B_* & \longrightarrow & C(\alpha) \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle f+g} \\ A'_* & \xrightarrow{\beta} & B'_* & \longrightarrow & C(\beta) \\ \downarrow & & \downarrow & & \downarrow \\ C(f) & \xrightarrow{\alpha+\beta} & C(g) & \longrightarrow & C(\alpha+\beta) \overset{iso}{=} C(f+g) \end{array}$$

Proof Chase the definitions. □

The next proposition gives us a controlled replacement for the five lemma.

Proposition 2.12 *If we have a commuting diagram*

$$\begin{array}{ccccc} A_* & \xrightarrow{\alpha} & B_* & \xrightarrow{i_B} & C(\alpha) \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle f+g} \\ A'_* & \xrightarrow{\beta} & B'_* & \xrightarrow{i_{B'}} & C(\beta) \end{array}$$

where $A_*,\ B_*,\ A'_*,\ B'_*$ *are n-dimensional* δ*-chain complexes and* g *and* $f+g$ *are* δ*-chain homotopy equivalences, then there is a* $k = k(n)$ *so that* f *is a* $k\delta$*-chain homotopy equivalence.*

Proof We extend the diagram to the right and down

$$\begin{array}{ccccccc} A_* & \xrightarrow{\alpha} & B_* & \xrightarrow{i_B} & C(\alpha) & \longrightarrow & C(i_B) \simeq \Sigma(A) \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle f+g} & & \downarrow{\scriptstyle \Sigma(f)} \\ A'_* & \xrightarrow{\beta} & B'_* & \xrightarrow{i_{B'}} & C(\beta) & \longrightarrow & C(i_{B'}) \simeq \Sigma(A') \\ & & \downarrow & & \downarrow & & \downarrow \\ & & C(g) & \longrightarrow & C(f+g) & \longrightarrow & C(\Sigma(f)) \end{array}$$

This is an algebraic version of the Puppe sequence. The first two mapping cones on the bottom are controlled chain contractible, so the last one is, as well. Here Σ just shifts the dimensions in the chain complex by one. This shows that $\Sigma(f)$ is a $k\delta$-equivalence for some k and that f is, as well. □

Remark 2.13 By continuing to the right, we can prove the analogous results whenever two of the three vertical maps are δ-chain homotopy equivalences.

Lemma 2.14 *Let* $A_* \xrightarrow{i} B_*$ *be a* δ*-chain map of* δ*-chain complexes so that there is a* δ*-chain map* $j{:}B_* \to A_*$ *with* $j \circ i\delta$*-chain homotopic to the identity. Then* B_* *is* 2δ*-chain homotopy equivalent to* $A_* \oplus C_*(i)$*.*

Proof B_* is δ-chain homotopy equivalent to $M_*(i)$. We will show that $M_*(i)$ splits. Writing $M_k(i) = A_k \oplus A_{k-1} \oplus B_k$, define $r{:}M_*(i) \to A_*$ by $r(a, a', b) = a + (-1)^k s(a') - j(b)$, where s is the chain homotopy from $j \circ i$ to the identity. This retracts $M_*(i)$ onto A_*, splitting $M_*(i)$ as $A_* \oplus C_*(i)$. □

3 Epsilon–delta K-theory

In this section, we will review controlled K-theory as described in [5] and [17].

Definition 3.1 Let h be an ϵ-automorphism of a geometric module $\mathcal{A}$ over a space B. We will say that h is ϵ*-elementary* if $\mathcal{A}$ can be written as a based direct sum $\mathcal{E} \oplus \mathcal{F}$ in such a way that h has matrix $\begin{pmatrix} I & * \\ 0 & I \end{pmatrix}$.

Definition 3.2 We will identify $\alpha{:}\mathcal{A} \to \mathcal{A}$ with $\alpha \oplus \mathrm{id} : \mathcal{A} \oplus \mathcal{F} \to \mathcal{A} \oplus \mathcal{F}$ for any geometric $\mathbb{Z}$-module $\mathcal{F}$ over B. If α and β are ϵ-automorphisms of $\mathcal{A}$, we write $\alpha \overset{\epsilon}{\sim} \beta$ if $\alpha \circ \beta^{-1}$ is ϵ-elementary. The relation $\overset{\epsilon}{\sim}$ generates an equivalence relation and we denote the set of equivalence classes of ϵ-automorphisms by $K_{1,\epsilon}(B)$. Direct sum makes this set into an additive semigroup. The Whitehead identities

$$\begin{pmatrix}\alpha & 0\\ 0 & \alpha^{-1}\end{pmatrix} = \begin{pmatrix}1&1\\0&1\end{pmatrix}\begin{pmatrix}1&0\\-1&1\end{pmatrix}\begin{pmatrix}1&1\\0&1\end{pmatrix}\begin{pmatrix}1&0\\\alpha&1\end{pmatrix}\begin{pmatrix}1&-\alpha^{-1}\\0&1\end{pmatrix}\begin{pmatrix}1&0\\\alpha&1\end{pmatrix}$$
$$\begin{pmatrix}\alpha & 0\\ 0 & \beta\end{pmatrix} = \begin{pmatrix}1&1\\0&1\end{pmatrix}\begin{pmatrix}1&0\\-1&1\end{pmatrix}\begin{pmatrix}1&1\\0&1\end{pmatrix}\begin{pmatrix}1&0\\\alpha&1\end{pmatrix}\begin{pmatrix}1&-\alpha^{-1}\\0&1\end{pmatrix}\begin{pmatrix}1&0\\\alpha&1\end{pmatrix}\begin{pmatrix}1&0\\0&\alpha\beta\end{pmatrix}.$$

show that the image of $K_{1,\epsilon}(B)$ in $K_{1,2\epsilon}(B)$ is an abelian group. We will define $Wh_\epsilon(B)$ to be $K_{1,\epsilon}(B)/\{\pm 1\}$.

The most unsatisfactory feature of this definition is the phrase "equivalence relation generated by $\overset{\epsilon}{\sim}$." The next lemma makes this relation more palatable.

Lemma 3.3 *If α and β are equivalent in $K_{1,\epsilon}(B)$, then $\alpha \overset{13\epsilon}{\sim} \beta$.*

This follows immediately from the next lemma.

Lemma 3.4 (Chapman's swindle) *If $\alpha, \beta{:}A \to A$ in $K_{1,\epsilon}(B)$ are automorphisms and $e_{\eta_k} \dots e_{\eta_1}\alpha = \beta$ with* radius$(\alpha_i) < \epsilon$ *for all $i, \alpha_i = e_{\eta_i} \dots e_{\eta_1}\alpha$, then there exist e_{ξ_j}: $\oplus_{i=1}^{2k+1} A \to \oplus_{i=1}^{2k+1} A$, $j = 1 \dots 13$, with* radius$(e_{\xi_j}) < \epsilon$ *for all j such that* $\left(e_{\xi_{13}} \dots e_{\xi_1}\right)$ $(\alpha \oplus \mathrm{id}) = \beta \oplus \mathrm{id}$.

Proof We have:

$$\begin{aligned}
&\alpha \oplus (\mathrm{id} \oplus \mathrm{id}) \oplus (\mathrm{id} \oplus \mathrm{id}) \oplus \cdots \oplus (\mathrm{id} \oplus \mathrm{id})\\
&\quad \overset{\epsilon}{\sim} \alpha \oplus \left(\alpha_1^{-1} \oplus \alpha_1\right) \oplus \left(\alpha_2^{-1} \oplus \alpha_2\right) \oplus \cdots \oplus \left(\alpha_k^{-1} \oplus \alpha_k\right)\\
&\quad = \left(\alpha \oplus \alpha_1^{-1}\right) \oplus \left(\alpha_1 \oplus \alpha_2^{-1}\right) \oplus \cdots \oplus \left(\alpha_{k-1} \oplus \alpha_k^{-1}\right) \oplus \alpha_k\\
&\quad \overset{\epsilon}{\sim} \left(\alpha\alpha_1^{-1} \oplus \mathrm{id}\right) \oplus \left(\alpha_1\alpha_2^{-1} \oplus \mathrm{id}\right) \oplus \cdots \oplus \left(\alpha_{k-1}\alpha_k^{-1} \oplus \mathrm{id}\right) \oplus \beta\\
&\quad = \left(e_{\eta_1}^{-1} \oplus \mathrm{id}\right) \oplus \left(e_{\eta_2}^{-1} \oplus \mathrm{id}\right) \oplus \cdots \oplus \left(e_{\eta_k}^{-1} \oplus \mathrm{id}\right) \oplus \beta\\
&\quad \overset{\epsilon}{\sim} (\mathrm{id} \oplus \mathrm{id}) \oplus (\mathrm{id} \oplus \mathrm{id}) \oplus \cdots \oplus (\mathrm{id} \oplus \mathrm{id}) \oplus \beta
\end{aligned}$$

The first "$\sim$" uses the first 6 term identity disjointly k times and the next line reassociates parentheses. The third line uses the second 6 term identity disjointly k more times and also uses the fact that $\alpha_k = \beta$. The last line combines k disjoint elementary operations. □

Theorem 3.5 (controlled $\widetilde{K}_1$ vanishing [17]) *For any finite polyhedron B there exist an $\epsilon_0 > 0$ and a k so that for any $\epsilon < \epsilon_0$ the map $Wh_\epsilon(B) \to Wh_{k\epsilon}(B)$ is zero.*

Remark 3.6

(i) The lemma and theorems above have a remarkable consequence: given a compact metric B, for every $\epsilon > 0$ there is a $\delta > 0$ so that every δ-automorphism with $\mathbb{Z}$-coefficients can be written (stably) as a product of at most 13 ϵ-elementary automorphisms.

(ii) The controlled vanishing theorem is also true as stated for polyhedra of bounded geometry. Since there are only finitely many isomorphism types of vertex stars, the inductive technique used to prove vanishing in the case where B is a finite polyhedron applies without alteration. The linearity is not stated in the original argument of Quinn, but follows easily from a subdivision argument due to Farrell–Jones.

The derivation of the surgery exact sequence of Theorem 1 will make extensive use of Quinn's Thin h-cobordism theorem, which we state here.

Theorem 3.7 (thin h-cobordism Theorem [17]) *Let B be a finite polyhedron. Then for every $\epsilon > 0$ there is a $\delta > 0$ so that if $n \geq 4$ and $p: M_0^n \to B$ is a $UV^1(\delta)$-map*[5] *and $(W; M_0, M_1)$ is a cobordism with strong deformation retractions r_t and s_t retracting W to M_0 and M_1 such that the lengths of the paths $p(r_1(r_t(x)))$ and $p(r_1(s_1(s_t(x))))$ are $<\delta$ for each $x \in W$, then there is a product structure on W so that the composition of p with the projection to M_0 is ϵ-homotopic to $p \circ r_1$.*

Addendum 3.8 The thin h-cobordism theorem remains true as stated if B is a compact ANR or a polyhedron with bounded geometry. For the latter, one simply notes that most existing proofs work "in parallel" to prove the bounded geometry case along with the finite case.[6] More specifically, note that in dimensions $n \geq 5$ the topological thin h-cobordism theorem parameterized by the identity map $M_0 \to M_0$ is an immediate corollary of the alpha approximation theorem. Using the given data, one constructs a homotopy equivalence from W to $M_0 \times [0, 1]$, which turns out to be controlled. One then applies the alpha approximation theorem twice, first to $M_1 \to M_0$ and then to $W \to M \times [0, 1]$ rel boundary.[7] The smooth and PL versions of the thin h-cobordism theorem can then be recovered using concordance smoothing results of Kirby–Siebenmann (see [14], p. 25). In particular, since any controlled torsion over M can be realized on a thin h-cobordism, this gives the vanishing of the epsilon Whitehead group controlled over a manifold of bounded geometry. Applying the "Alexander trick" of Proposition 9.2 extends this to show vanishing of the controlled Whitehead group parameterized over any polyhedron of bounded geometry. This general procedure of realizing and algebraic obstruction and then using geometric methods to prove that it vanishes will play a major role in this paper.

Definition 3.9 Let A be a geometric module on a metric space X.

(i) An ϵ-deformation $h: A \to A$ is an ϵ-morphism which is a product $\prod_{i=1}^{\ell} E_i$ of elementary ϵ-morphisms so that each product $\prod_{i=1}^{k} E_i$ is an ϵ-morphism for $k \leq \ell$.
(ii) An ϵ-projection p on A is an ϵ morphism $p: A \to A$ such that $p \circ p = p$. We say that p is geometric if A can be written as a direct sum $A_1 \oplus A_2$ of geometric submodules with $p|A_1 = \text{id}$ and $p|A_2 = 0$, that is, if p is the standard projection of A onto the summand A_1.

Theorem 3.10 (controlled $\widetilde{K}_0$ vanishing, Thm 8.4 [17]) *Let B be a finite polyhedron. There exist $\epsilon_0 > 0$ so that for every $\epsilon > 0$ with $\epsilon < \epsilon_0$ there is a $\delta > 0$ so that if A is a geometric $\mathbb{Z}$-module on B and $p: A \to A$ is an δ-projection, then there exist a geometric $\mathbb{Z}$-module C on B and a geometric projection $q: C \to C$ and ϵ-deformations H_1 and H_2 on $A \oplus C$ so that $H_1 \circ (p \oplus q) \circ H_2$ is geometric.*

5 See Definition 4.1 below.

6 The proofs of the finite case are already parallel processes, so no modification is necessary to extend them to the bounded geometry case.

7 Squeezing the "active" area of the homotopy into $[0, \epsilon]$ establishes control in the [0,1]-coordinate.

Remark 3.11 If $p{:}A \to A$ is an ϵ-projection of a geometric $\mathbb{Z}$-module over X, then $p + t(1-p)$ is a 2ϵ-isomorphism of the associated geometric $\mathbb{Z}[t, t^{-1}]$-module $A \otimes \mathbb{Z}[t, t^{-1}]$ with inverse $p + t^{-1}(1-p)$. One can show ([8,17]) that this automorphism is trivial in a controlled K_1 if and only if the projection p is trivial in controlled K_0 , i.e., that

(i) There exists k_1 so that if $H_1 \circ (p \oplus q) \circ H_2$ is geometric, as in Definition 3.9, then some stabilization of $p + t(1-p)$ is a $k_1\epsilon$-deformation.
(ii) There exists k_2 so that if $p + t(1-p)$ is an ϵ-deformation, then there exist $k_2\epsilon$-deformations H_1, H_2 and a geometric projection q as above so that $H_1 \circ (p \oplus q) \circ H_2$ is geometric.

Controlled $\widetilde{K}_0$ vanishing is the key algebraic ingredient in the proof of the following controlled version of Browder's $M \times \mathbb{R}$ theorem which is a special case of Quinn's approximate end theorem. See p. 283 of [17] for $n \geq 6$ and p. 505 of [18] for $n = 5$.

Theorem 3.12 (controlled $\mathrm{M} \times \mathbb{R}$) *Suppose that $n \geq 5$, and B is a finite polyhedron with the standard metric. Then there exist a $\delta_0 > 0$ and a $k > 0$ so that if $\delta < \delta_0$ and $K \to B$ is a $UV^1(\delta)$-map from a finite polyhedron to B and $W^n \to K \times \mathbb{R}^1$ is a proper δ-equivalence from an n-manifold W without boundary to $K \times \mathbb{R}^1$ over $B \times \mathbb{R}^1$, then there is a closed codimension one submanifold M of W which is a $k\delta$-strong deformation retract of W over B. The k and δ_0 depend on the dimension of W and the dimension of B.*

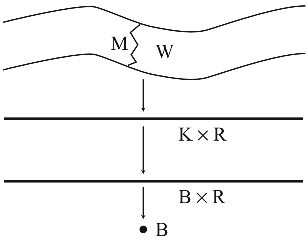

If W^n is a manifold with boundary and ∂W has a controlled splitting, then the splitting extends to the interior, provided that $n \geq 5$. The theorem is also valid if B is a polyhedron with bounded geometry. In this case, M will not be closed and the k and δ_0 depend on the dimension of W and the complexity of B.

Addendum 3.13 Most recent proofs of topological invariance of torsion also prove the vanishing of the controlled Whitehead group of geometric $\mathbb{Z}[t,\ t^{-1}]$-modules. This yields the vanishing of the controlled projective class group with $\mathbb{Z}$ coefficients, since that group embeds in the controlled Whitehead group of geometric $\mathbb{Z}[t,\ t^{-1}]$-modules. This is the algebraic key to proving the controlled $M \times \mathbb{R}$ theorem. The argument in [17] produces the desired splitting when it is given this algebraic data. As before, the proof extends without alteration to the case of polyhedral control spaces with bounded geometry. For readers who prefer to work with manifolds, a similar analysis to that of Addendum 3.8 establishes the vanishing of controlled projective class groups over these same spaces. After embedding the epsilon projective class group into the epsilon Whitehead group of $M \times S^1$ and realizing an obstruction by a controlled h-cobordism, one solves the manifold problem using Chapman's generalization of the alpha approximation theorem. See [4].

As before, realizing the algebraic problems as geometric problems, solving them, and using the well-definedness of the algebraic obstructions as in [8] or [17] completes the proof. This and the argument in Addendum 3.8 extend the thin h-cobordism theorem, end theorem and approximate end theorem to the situation of δ-control $k\delta$-vanishing over manifolds and polyhedra of bounded geometry.

4 Surgery below the middle dimension

We begin with some definitions.

Definition 4.1

(i) We will say that a map $p{:}K \to B$ between finite polyhedra is $UV^1(\delta)$ if for every map $\alpha{:}P^2 \to B$ of a 2-complex into B with lift $\alpha_0{:}P_0 \to K$ defined on a subcomplex P_0, there is a map $\bar{\alpha} : P \to K$ with $\bar{\alpha}|P_0 = \alpha_0$ so that $p \circ \bar{\alpha}$ is δ-homotopic to α.

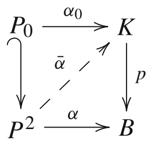

This can be thought of as saying that p has Čech δ-simply connected point-inverses. If ℓ is a loop near $p^{-1}(b)$ for some b, then the image of ℓ in B is contractible and the contraction can be δ-lifted to K, giving a contraction of ℓ close to $p^{-1}(x)$. The map p is said to be UV^1 if it is $UV^1(\delta)$ for every $\delta > 0$. More generally, we will say that a map $p{:}K \to B$ with B a not necessarily finite polyhedron is $UV^1(\delta)$ if it is proper[8] and satisfies the conditions above. See [15] for details.

(ii) If $p{:}P \to B$ is a $UV^1(\delta)$ control map, we will say that $f{:}M \to P$ is $(\delta,\ k)$-connected over B if whenever (L, L_0) is a CW pair with $\dim(L) \le k$ and $\alpha{:}L_0 \to M$ is a map such that there is a map $\beta{:}L \to P$ with $f \circ \alpha = \beta|L_0$, then there exist a map $\gamma{:}L \to M$ with $\gamma|L_0 = \alpha$ and a homotopy $h_t{:}L \to P \operatorname{rel} L_0$ with $h_0 = f \circ \gamma, h_1 = \beta$, and $\operatorname{diam}(p \circ h(\{x\} \times I)) < \delta$ for each $x \in L$.

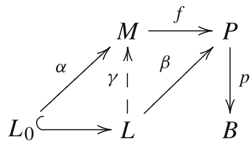

Definition 4.2 If P is a finite polyhedron and B is compact metric, we say that P is an unrestricted ϵ-Poincaré complex of formal dimension n over B if there exist a subdivision of P so that images of simplices have diameter $<\epsilon$ in B and so that there is a cycle y in the simplicial chains $C_n(P)$ so that $y \cap \_ : C^{\#}(P) \to C_{n-\#}(P)$ is an ϵ-chain homotopy equivalence. The definition of a restricted ϵ-Poincaré complex of formal dimension n is similar except that we require in addition that the control map $p{:}P \to B$ be $UV^1(\epsilon)$.

Addendum 4.3 If B is merely locally compact and P is a finite-dimensional locally finite complex, we will require that y be a locally finite cycle in the definitions of our ϵ-Poincaré complexes.

For simplicity, we will restrict our discussion below to the oriented case. The unoriented case can be handled as usual by using the orientation double cover. In a similar vein, we will omit mention of the orientation character in our definition of the ϵ-Wall groups below.

Definition 4.4 Let P be an unrestricted δ-Poincaré duality space of formal dimension n over a metric space B and let ν be a (TOP, PL or O) bundle over P. A δ-surgery problem or degree

[8] i.e., inverse images of compact sets are compact.

one normal map is a triple (M^n, ϕ, F) where $\phi : M \to P$ is a map from a closed topological n-manifold M to P such that $\phi_*([M]) = [P]$ and F is a stable trivialization of $\tau_M \oplus \phi^*\nu$. Two problems (M, ϕ, F) and $(\overline{M}, \bar{\phi}, \bar{F})$ are equivalent if there exist an $(n+1)$-dimensional manifold W with $\partial W = M \coprod \overline{M}$, a proper map $\Phi: W \to P$ extending ϕ and $\bar{\phi}$, and a stable trivialization of $\tau_W \oplus \Phi^*\nu$ extending F and $\bar{F}$. Such an equivalence is called a normal bordism. See p. 9 of [21] for further details.

We will use the notation $\begin{array}{ccc} M & \xrightarrow{\phi} & P \\ & & \downarrow \\ & & B \end{array}$ to denote a δ-surgery problem. When B is understood, we will shorten the notation to $\phi: M \to P$ or even to ϕ. We will follow tradition in pretending that our topological manifolds are PL in order to simplify details of the proofs. A good reference for the appropriate TOP material is [14].[9] In all cases, the bundle information is included as part of the data. Our theorem on surgery below the middle dimension and its proof are parallel to Theorem 1.2 on p. 11 of [21]. As usual, surgery below the middle dimension is unobstructed.

Theorem 4.5 *Let $(P^n, \partial P)$ be an unrestricted ϵ-Poincaré duality pair over a finite polyhedron $B, n \geq 6$, or $n \geq 5$ if ∂P is empty. Consider an ϵ-surgery problem $\phi: (M, \partial M) \to (P, \partial P)$. Then $\phi: (M, \partial M) \to (P, \partial P)$ is normally bordant to an ϵ-surgery problem $\bar{\phi}: (\overline{M}, \partial \overline{M}) \to (P, \partial P)$ such that $\bar{\phi}$ is $(\epsilon, \left[\frac{n}{2}\right])$-connected over B and $\bar{\phi}|: \partial \overline{M} \to \partial P$ is $(\epsilon, \left[\frac{n-1}{2}\right])$-connected.*

Proof We start by considering the case in which $\partial P = \emptyset$. Triangulate M so that ϕ is simplicial and the diameters $p \circ \phi(\tau), \tau \in M$ and $p(\sigma), \sigma \in P$, are $<\epsilon$. Replacing P by the simplicial mapping cylinder of ϕ, we can assume that $M \subset P$. We inductively define a bordism $U^{(i)}, -1 \leq i \leq \left[\frac{n+1}{2}\right]$ and maps $\Phi^{(i)}: U^{(i)} \to M \cup P^{(i)}$, so that $\partial U^{(i)} = M \coprod \overline{M}^{(i)}$ and so that $\Phi^{(i)}$ is an ϵ-homotopy equivalence. We begin by setting $U^{(-1)} = M \times I$, and letting $\Phi^{(-1)} \to P$ be $\phi \circ$ proj. Let $U^{(0)}$ be obtained from $U^{(-1)}$ by adding a disjoint $(n+1)$-ball corresponding to each 0-cell of $P - M$. The map $\Phi^{(0)}$ is constructed by collapsing each new ball to a point and sending the point to the corresponding 0-cell of $P - M$. Assume that $\Phi^{(i)}: U^{(i)} \to P$ has been constructed in such a way that $U^{(i)}$ is an abstract regular neighborhood of a complex consisting of M together with cells in dimensions $\leq i$ corresponding to the cells of $P - M$ in those dimensions. Assume further that $\Phi^{(i)}$ is the composition of the regular neighborhood collapse with a map which takes cells to corresponding cells. Each $(i+1)$-cell of $P - M$ induces an attaching map $S^i \to U^{(i)}$. If $2i + 1 \leq n$, general position allows us to move this map off of the underlying complex and approximate the attaching map by an embedding $S^i \to \overline{M}^{(i)}$. The bundle information tells us how to thicken this embedding to an embedding of of $S^i \times D^{n-i}$ and attach $(i+1)$-handles to $U^{(i)}$, forming $U^{(i+1)}$. We extend $\Phi^{(i)}$ to $\Phi^{(i+1)}$ in the obvious manner. This process terminates with the construction of $U^{\left[\frac{n+1}{2}\right]}$. Turning $U^{\left[\frac{n+1}{2}\right]}$ upside down, we see that $U^{\left[\frac{n+1}{2}\right]}$ is obtained from $\overline{M}^{\left[\frac{n+1}{2}\right]}$ by attaching handles of index $>\left[\frac{n+1}{2}\right]$. Thus, the composite map $\overline{M}^{\left[\frac{n+1}{2}\right]} \to P$ is $(\epsilon, \left[\frac{n}{2}\right])$-connected over P.

[9] In the present instance, we can embed M in a high-dimensional euclidean space, where it has a normal closed disk bundle by [14]. This bundle has a PL structure. If T is a k-skeleton for this normal disk bundle with $2k + 1 \leq \dim M$, then the restriction of the bundle projection to T can be approximated by an embedding. The image of T under this embedding can be used in the construction above in place of a skeleton of a triangulation of M. Alternatively, we could appeal directly to the theorem in [14] that shows that high-dimensional TOP manifolds have handlebody decompositions.

In case $\partial P \neq \emptyset$, the argument is similar. We first construct U over the ∂ (and, therefore, over a collar neighborhood of the boundary) and then construct U over the interior. □

Remark 4.6

(i) The Poincaré duality of P was never used. This result is true for arbitrary P and arbitrary maps $p: P \to B$. The space B can be an arbitrary metric space.

(ii) Notice that direct manipulation of cells and handles has replaced the usual appeals to homotopy theory and the Hurewicz–Namioka theorem. This is a general technique for adapting arguments from ordinary algebraic topology to ϵ-controlled topology.

(iii) The construction in the proof yields somewhat more—we wind up with $(\overline{M}, \partial\overline{M}) \subset (P, \partial P)$. When $n = 2k+1$, $\overline{M}$ and P are equal through the k-skeleton. When $n = 2k$, $\partial\overline{M}$ is equal to ∂P through the $(k-1)$-skeleton and $\overline{M}$ contains every k-cell of $P - \partial P$. Since $\overline{M} \to P$ is k-connected, every k-cell in ∂P is homotopic rel boundary to a map into M. By attaching a $(k+2)$-cell to this homotopy along a face, we can guarantee that for every k-cell in $\partial P - M$ there is a $(k+1)$-cell in P so that half of the boundary of the $(k+1)$-cell maps homeomorphically onto the k-cell and the other half maps into M. Alternatively, the same effect can be obtained by adding a collar to ∂P and giving it the product CW structure.

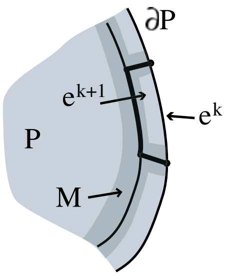

5 Controlled cell-trading

In this section, we prove a controlled version of Whitehead's cell-trading lemma and apply it to prove a useful controlled Hurewicz–Whitehead theorem. The operations we describe apply equally well to cells in a finely subdivided CW complex and to handles in a finely subdivided handle decomposition. We will use cell terminology throughout, except for the terms "handle addition" or "handle slide." These refer to the same mathematical operation.

We will write the operation of sliding handles in B over handles in C corresponding to $m: B \to C$ schematically as

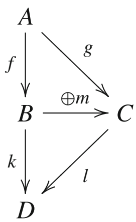

and call it adding the B-cells to C via m. When the sequence $A \to B \oplus C \to D$ is a part of a cellular chain complex, this operation is realized geometrically by handle-addition by taking each generator x in B and sliding it across $m(x)$. Changing the attaching maps of the cells this way has the effect described above on the cellular chains. If m is an ϵ-morphism, then the new chain complex is a 2ϵ-chain complex 2ϵ-isomorphic to the old one.

Lemma 5.1 *Let $\ell > 0$ be given and let B be a polyhedron with the standard metric. Then there exist $\delta_0 = \delta_0(\ell)$ and $k = k(\ell) > 0$ so that if $\delta < \delta_0$ and*

(i) *(X, Y) is a CW pair such that the image of each cell has diameter $<\delta$ in B. We call this a δ-CW pair.*
(ii) *$p:(X, Y) \to B$ is a map so that p and $p|Y$ are $UV^1(\delta)$-maps.*
(iii) *The cellular chain complex $C_\#(X)$ is decomposed as (based) modules $C_\#(Y) \oplus C_\#(X - Y)$ for which the boundary map has the form*

$$\begin{array}{ccc} C(Y)_\# & \oplus & C_\#(X-Y) \\ \downarrow & \swarrow & \downarrow \\ C_{\#-1}(Y) & \oplus & C_{\#-1}(X-Y) \end{array}$$

(iv) *There is an δ-chain homotopy s with $s|C(Y)_\# = 0$, from the identity to a morphism which is 0 on $C_\#(X - Y)$ for $\# < \ell$.*

then X may be changed by a simple homotopy equivalence of size $k\delta$ to a complex X', so that the cellular chains $C_\#(X')$ have the form $C(Y)_\# \oplus C_\#(X' - Y)$ where $C_\#(X' - Y) = 0$ for $\# < \ell$ and $C_\#(X' - Y) = C_\#(X - Y)$ for $\# \geq \ell + 2$.

Proof Using the $UV^1(\delta)$ condition we can trade away 0- and 1-cells of $X - Y$. Now perform the same operations as in the algebraic cell-trading lemma, but do them geometrically, using handle additions and cell cancellations, rather than algebraically. The constant δ_0 is included to guarantee that all intersections take place in simply connected regions of the space, so that geometric intersections can be manipulated to agree with algebraic intersection numbers without taking the global fundamental group into account. There are discussions of controlled cell-trading in Sect. 6 of [17] and on page 84 of [5]. □

Addendum 5.2 The argument above works for B a uniformly locally simply connected space, except that the relation between δ and $k\delta$ is no longer linear in the absence of a linear relationship between the diameter of a small loop and the diameter of a disk it bounds.

Controlled cell-trading is a very useful tool in epsilon–delta topology. Here's a controlled Whitehead theorem whose proof relies on cell-trading. Let B be a finite polyhedron endowed with the standard metric.

Proposition 5.3 (controlled Hurewicz–Whitehead) *Let an integer $n > 0$ be given. There exist a $k > 0$ and a $\delta_0 > 0$ depending on n so that if*

(i) *$\delta < \delta_0$*
(ii) *(X, Y) is an n-dimensional polyhedral pair with cells of size δ over B and $p:X \to B$ is a $UV^1(\delta)$-map such that $p|Y$ is also $UV^1(\delta)$.*
(iii) *$C_*(Y) \to C_*(X)$ is a δ-chain homotopy equivalence.*

Then $Y \to X$ is a $k\delta$-homotopy equivalence.

Proof By cell-trading, there is a k_1 depending on n so that (X, Y) is $k_1(n)\delta$-homotopy equivalent rel Y to a CW pair (X', Y) so that all cells of $X' - Y$ are of dimension $>n$. Let

$\phi{:}X \to X'$ and $\psi{:}X' \to X$ be the $k_1(n)\delta$-homotopy equivalence and homotopy inverse. By cellular approximation (or general position), we can take ϕ to be a map from X into $Y \subset X'$. This approximation loses as much as $3^{\dim X}\delta$ more in control, since cell trading may make the cells larger. The controlled homotopy from $\psi \circ \phi$ to the identity gives a controlled strong deformation retraction from X to Y, establishing the desired controlled homotopy equivalence. □

6 The epsilon–delta π–π theorem

Definition 6.1 If $f_\#{:}A_\# \to B_\#$ is a chain homomorphism, we define $K^\#(f)$ to be the algebraic mapping cone of $f^\#{:}B^\# \to A^\#$. We define $K_\#(f)$ to be the dual of $K^\#(f)$. Unraveling this, we see that $K^k(f) = A^k \oplus B^{k+1}$, $K_k(f) = A_k \oplus B_{k+1}$, and that up to signs in the boundary map, $K_\#(f)$ is the algebraic mapping cone of $f_\#$ with dimensions shifted by one.

Now, suppose that we are given a degree one map $\phi{:}M \to P$ from a manifold to a δ-Poincaré space over a metric space X. We need to show that the complex $K_\#(\phi)$ described above has $k\delta$-Poincaré duality for some $k = k(n)$. Following [19], Proposition 2.2, we have a controlled chain homotopy commuting diagram

$$\begin{array}{ccccc} C^k(P) & \xrightarrow{\phi^\#} & C^k(M) & \longrightarrow & K^k(\phi) \\ \cong\Big\downarrow [P]\cap\_ & & \cong\Big\downarrow [M]\cap\_ & & \\ C_{n-k}(P) & \xleftarrow{\phi_\#} & C_{n-k}(M) & & \end{array}$$

which splits the top row of the diagram up to chain homotopy. By Lemma 2.14, this gives us a chain-homotopy equivalence

$$C^*(M) \cong K^*(\phi) \oplus C^*(P).$$

Dualizing, we have

$$C_*(M) \cong K_*(\phi) \oplus C_*(P).$$

As in the classical case, these splittings preserve cap product with the fundamental class, so by Lemma 2.12 there is a controlled chain homotopy equivalence $[M] \cap \_ : K^*(\phi) \to K_{n-*}(\phi)$.

At this point, the reader whose primary interest is in getting to the proof of Theorem 1.1 as directly as possible should move forward to the next section.

The next lemma is standard, as in Browder et al. [3] or Siebenmann's thesis.

Lemma 6.2 *If M is a topological manifold and $H \subset M$ is a handle attached to ∂M, then the effect of excising the interior of H from $(M, \partial M)$ is to kill* $\mathrm{im}(H_*(H, H \cap \partial M) \to H_*(M, \partial M))$ *and to create homology in the next dimension corresponding to* $\ker(H_*(H, H \cap \partial M) \to H_*(M, \partial M))$.

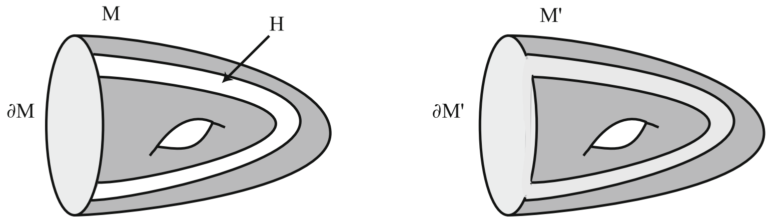

Theorem 6.3 (simply connected controlled π–π theorem) *If B is a finite polyhedron with the standard metric, then there exist $k > 0$ and $\epsilon_0 > 0$ so that if $(P^n, \partial P), n \geq 6$, is an ϵ-Poincaré duality space over B, $\epsilon \leq \epsilon_0$, and*

$$\begin{array}{ccc}(M, \partial M) & \xrightarrow{\phi} & (P, \partial P) \\ & & \downarrow p \\ & & B\end{array}$$

is an ϵ-surgery problem with bundle information assumed as part of the notation so that both $p: P \to B$ and $p|: \partial P \to B$ are $UV^1(\epsilon)$, then we may do surgery to obtain a normal bordism from $(M, \partial M) \to (P, \partial P)$ to $(M', \partial M') \to (P, \partial P)$, where the second map is a $k\epsilon$-homotopy equivalence of pairs. Here, k and ϵ_0 will depend on n.

Proof The argument is a translation into ϵ-terms of the bounded π–π Theorem of [10]. We first focus on the case $n = 2\ell$. By Theorem 4.5 we may do surgery below the middle dimension. We obtain a surgery problem $M' \xrightarrow{\phi'} P$ so that ϕ' is an inclusion which is the identity through dimension ℓ.

This means that cancelling cells in the algebraic kernel $K_\#(P, \partial P; M', \partial M')$ yields a complex which is 0 through dimension $\ell - 1$. Abusing the notation, we will assume that the chain complex $K_\#(P, \partial P; M', \partial M')$ is 0 for $\# \leq \ell - 1$. The generators of $K_{\ell-1}(P, M)$ correspond to ℓ-cells in $\partial P - M$. Cancelling these against the $(\ell + 1)$-cells described in Remark 4.6ii and leaving out the primes for notational convenience, we have

$$\begin{aligned} K_\#(P, \partial P; M, \partial M) &= 0 \qquad \# \leq \ell - 1 \\ K_\#(P, M) &= 0 \qquad \# \leq \ell - 1. \end{aligned}$$

Since $K^{n-\#}(P, \partial P; M, \partial M)$ is ϵ-chain homotopic to $K_\#(P, M)$, there is an algebraic ϵ-homotopy σ on $K^\#(P, \partial P; M, \partial M)$ satisfying $\sigma\delta + \delta\sigma = 1$ for $\# \geq \ell + 1$. Taking duals, there is an algebraic homotopy s on $K_\#(P, \partial P; M, \partial M)$ such that $s\partial + \partial s = 1$ for $\# \geq \ell + 1$. Since $K_\# = K_\#(P, \partial P; M, \partial M)$ is finite-dimensional, the "cell trading" procedure may be applied upside down, so that the $K_\#$ is changed to

$$0 \longrightarrow K'_{\ell+2} \xrightarrow{\partial} K'_{\ell+1} \xrightarrow{\partial} K_\ell \longrightarrow 0$$

together with a homotopy s so that $s\partial + \partial s = 1$ except at degree ℓ. Again, we leave out the primes for notational convenience. Corresponding to each generator of $K_{\ell+1}$ (and at a point near where the generator sits in the control space) we introduce a pair of cancelling $(\ell - 1)$- and ℓ-handles and excise the interior of the $(\ell - 1)$-handle from $(M, \partial M)$, modifying the map so that the new boundary maps to ∂P. The chain complex for this modified M is:

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & K_{\ell+3} & \longrightarrow & K_{\ell+2} & \longrightarrow & K_{\ell+1} & \longrightarrow & 0 \\
 & & & & & & \oplus & & \\
 & & & & & & K_{\ell+2} & &
\end{array}
$$

All generators of $K_\ell \oplus K_{\ell+1}$ are represented by discs. We may represent any linear combination of these discs by an embedded disc, and these embedded discs may be assumed to be disjoint by the usual piping argument. See p. [21], p. 39. The UV^1 condition on the interior is used here. We do surgery on the following elements: For each generator x of K_ℓ, we do surgery on $(x - \partial sx, sx)$ and for each generator y of $K_{\ell+2}$, we do surgery on $(0, \partial y)$. We can think of the process as introducing pairs of cancelling ℓ- and $(\ell+1)$-handles, performing handle additions with the ℓ-handles, and then excising the ℓ-handles from $(M, \partial M)$. The resulting chain complex is:

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & K_{\ell+2} & \xrightarrow{\partial} & K_{\ell+1} & \longrightarrow & K_\ell & & \\
 & & & & \oplus & \overset{1-\partial s}{\nearrow} & \oplus & & \\
 & & & & K_\ell & \xrightarrow{s} & K_{\ell+1} & \longrightarrow & 0 \\
 & & & & \oplus & \overset{\partial}{\nearrow} & & & \\
 & & & & K_{\ell+2} & & & &
\end{array}
$$

which is easily seen to be contractible, the contraction being

$$
\begin{array}{ccccccccc}
0 & \longleftarrow & K_{\ell+2} & \xleftarrow{s} & K_{\ell+1} & \longleftarrow & K_\ell & & \\
 & & & & \oplus & \overset{1-\partial s}{\swarrow} & \oplus & & \\
 & & & & K_\ell & \xleftarrow{\partial} & K_{\ell+1} & \longleftarrow & 0 \\
 & & & & \oplus & \overset{s}{\swarrow} & & & \\
 & & & & K_{\ell+2} & & & &
\end{array}
$$

Dualizing, we see that after surgery, $K_\#(P, \partial P; M, \partial M)$ is ϵ-chain contractible. Poincaré duality shows that $K_\#(P, M)$ is $k'\epsilon$-chain contractible. Together, these imply the $k''\epsilon$-chain contractibility of $K_\#(\partial P, \partial M)$. Using the controlled Hurewicz–Whitehead theorem now shows that $\partial P \to \partial M$ and $P \to M$ are $k\epsilon$-homotopy equivalences for some k depending on n. This application of the controlled Hurewicz–Whitehead theorem, Proposition 5.3, uses both UV^1 conditions. An easy argument composing deformations in the mapping cylinder of $(M, \partial M) \to (P, \partial P)$ completes the proof that $(M, \partial M) \to (P, \partial P)$ is a controlled homotopy equivalence.

To obtain the π–π-theorem in the odd dimensional case we resort to a trick.

(1) Cross with S^1 to get back to an even dimension and do the surgery.
(2) Go to the cyclic cover and split using the controlled $M \times \mathbb{R}$ theorem to obtain a controlled homotopy equivalence of the ends.

This completes the proof. □

Remark 6.4 There are a number of useful extensions of the π–π theorem. The theorem remains true in the presence of multiple boundary components, provided that the restrictions to the extra components are δ-equivalences. The theorem also remains true if the "active" boundary component is divided into two submanifolds with boundary provided that the original normal map is a δ-equivalence over one piece and satisfies the π–π condition over the other. In fact, the theorem remains true if the "inert" boundary components are unrestricted objects and the restrictions to these objects induce δ-duality at the chain level.

7 Remarks on pro theory

Throughout this paper, we will be working with systems of sets and/or groups. The purpose of this section is to establish some definitions and notation. The setup we're describing is pro-theory, but only for the comparatively uncomplicated case of systems indexed by the natural numbers. For the reader unfamiliar with these things, a good example to keep in mind is the system of homology groups near infinity in an open manifold. If we have a basis $\{U_i\}$ of neighborhoods of infinity, it's actually the system $\{H_k(U_i)\}$ as a whole that we're interested in, not the individual groups. A proper map from one manifold to another induces a homomorphism of homology systems at infinity, but a certain amount of passing to subsequences and reindexing is necessary in order to write down a pleasant commutative diagram representing the map of homology systems.

Definition 7.1

(i) In this paper, a system will be an inverse system of sets and maps indexed by the integers, most often positive, but sometimes including 0. We write such a system as $\{A_i, \alpha_i\}$ where $\alpha_i : A_i \to A_{i-1}$. The maps α_i are called bonding maps.

(ii) The relation of equivalence on systems is the equivalence relation generated by passing to subsequences. Of course we must also allow passing to "supersequences" in order to maintain symmetry. When we pass to subsequences we will automatically compose the bonding maps and reindex the remaining spaces.

(iii) Defining a map $\{A_i, \alpha_i\} \to \{A'_i, \alpha'_i\}$ of systems in full generality is messy. The official definition is $\lim_k \operatorname{colim}_j \operatorname{Maps}(A_j, A'_k)$. This allows allows a bit more flexibility in defining maps than we'll need in this paper. Suffice to say that after passing to subsequences and reindexing, maps can be represented by a commuting diagram of level-preserving maps

$$\begin{array}{ccccccccc} A_1 & \xleftarrow{\alpha_2} & A_2 & \xleftarrow{\alpha_3} & A_3 & \xleftarrow{\alpha_4} & A_4 & \longleftarrow & \cdots \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ A'_1 & \xleftarrow{\alpha'_2} & A'_2 & \xleftarrow{\alpha'_3} & A'_3 & \xleftarrow{\alpha'_4} & A_4 & \longleftarrow & \cdots \end{array}$$

(iv) If we have maps in both directions, it may not be possible to represent both by level-preserving maps in the same diagram and the best we can get by passing to subsequences and reindexing is a diagram like

$$\begin{array}{ccccccccc} A_1 & \xleftarrow{\alpha_2} & A_2 & \xleftarrow{\alpha_3} & A_3 & \xleftarrow{\alpha_4} & A_4 & \longleftarrow & \cdots \\ \downarrow{\scriptstyle\beta_1} & \nwarrow{\scriptstyle\gamma_2} & \downarrow{\scriptstyle\beta_2} & \nwarrow{\scriptstyle\gamma_3} & \downarrow{\scriptstyle\beta_3} & \nwarrow{\scriptstyle\gamma_4} & \downarrow{\scriptstyle\beta_4} & & \\ A'_1 & \xleftarrow{\alpha'_2} & A'_2 & \xleftarrow{\alpha'_3} & A'_3 & \xleftarrow{\alpha'_4} & A_4 & \longleftarrow & \cdots \end{array}$$

with commuting squares and parallelograms. Note that if all of the triangles commute, such a diagram implies equivalence of the systems $\{A_i, \alpha_i\}$ and $\{A'_i, \alpha'_i\}$, since it is easy to build a larger system which contains both as subsystems. The systems would still be equivalent if we only had $\alpha_{i-1} \circ \gamma_i \circ \beta_i = \gamma_{i-1} \circ \alpha'_i \circ \beta_i$ for each i, rather than strict commutativity in the diagram above. In general, composing with bonding maps to get commutativity is allowed as long as the commuting subdiagrams are cofinal in the original system. See [12] for more information, including a translation of $\lim_k \operatorname{colim}_j \operatorname{Maps}(A_j, A'_k)$ into something more readable.

Definition 7.2

(i) A system is Mittag–Leffler if it is equivalent to a system of epimorphisms.
(ii) A system is stable if it is equivalent to a system of isomorphisms.

After passing to subsequences and reindexing, stability of a system $\{A_i, \alpha_i\}$ leads to a commuting diagram like the one below, from which one can see that each α_i maps the image of A_{i+1} in A_i bijectively onto the image of A_i in A_{i-1}.

$$\begin{array}{ccccccccc} A_1 & \xleftarrow{\alpha_2} & A_2 & \xleftarrow{\alpha_3} & A_3 & \xleftarrow{\alpha_4} & A_4 & \longleftarrow & \cdots \\ \downarrow{\scriptstyle\beta_1} & \nwarrow{\scriptstyle\gamma_2} & \downarrow{\scriptstyle\beta_2} & \nwarrow{\scriptstyle\gamma_3} & \downarrow{\scriptstyle\beta_3} & \nwarrow{\scriptstyle\gamma_4} & \downarrow{\scriptstyle\beta_4} & & \\ A & \xleftarrow{\mathrm{id}} & A & \xleftarrow{\mathrm{id}} & A & \xleftarrow{\mathrm{id}} & A & \longleftarrow & \cdots \end{array}$$

8 Algebraic δ-surgery groups

In this section, we will define our algebraic δ-surgery groups over a polyhedron B. In even dimensions, we will define these groups to be systems of δ-Witt groups over B. Our $(2k+1)$-dimensional groups over B will be the $(2k+2)$-dimensional groups over $B \times \mathbb{R}$. Thus, in this section we will be looking at locally finite geometric chain complexes over finite polyhedra and noncompact polyhedra of bounded geometry. Restricting ourselves to locally finite polyhedra of this special form allows us to phrase our work in terms of epsilons and deltas, rather than working with collections of open covers.

In this section, we will define a sequence of abelian semigroups $\{L_{n,B,\delta}(e)\}$. We will show that this sequence is equivalent to a sequence of groups and homomorphisms and that this system serves as an appropriate system of Wall groups for epsilon–delta surgery. To begin with, we will restrict our attention to even-dimensional manifolds.

Definition 8.1 Let $\eta = \pm 1$. Let $I_\eta = \{0\}$ for $\eta = 1$ and $2\mathbb{Z}$ for $\eta = -1$. By a special geometric ($\mathbb{Z}-$)quadratic η-form over B, we will mean a triple (A, λ, μ) where A is a geometric $\mathbb{Z}$-module over B, $\lambda{:}A \times A \to \mathbb{Z}$ is $\mathbb{Z}$-bilinear, $\mu{:}A \to \mathbb{Z}/I_\eta$

(i) $\lambda(x, y) = \eta\lambda(y, x)$ $\qquad x, y \in A$
(ii) $\lambda(x, x) = \mu(x) + \eta\mu(x)$ $\qquad x \in A$
(iii) $\mu(x + y) - \mu(x) - \mu(y) = \lambda(x, y) \mod I_\eta$ $\qquad x, y \in A$
(iv) $\mu(xa) = a^2\mu(x)$ $\qquad x \in A,\ a \in \mathbb{Z}$
(v) $A\lambda{:}A \to A^*$ defined by $A\lambda(x)(y) = \lambda(x, y)$ for $x, y \in A$ is an isomorphism.[10]

[10] The usual simplicity condition is not needed because of the vanishing of controlled $\mathbb{Z}$-Whitehead groups. We are getting away with a certain amount in this section because the controlled Whitehead group vanishes and we do not have to prove simplicity at each stage of our argument.

The radius of this form is $<\epsilon$ if $\lambda(x, y) = 0$ when $d(x, y) \geq \epsilon$.

Definition 8.2

(i) Let $\eta = \pm 1$. If A is a geometric module over B, the nonsingular η-hyperbolic quadratic form on $A \oplus A$ is the form $\eta H(A)$ which has matrix $\begin{pmatrix} 0 & 1 \\ \eta & 0 \end{pmatrix}$ corresponding to each basis element of A with $\mu(1, 0) = \mu(0, 1) = 0$. This simply means that $\eta H(A)$ has a standard $2 \times 2\eta$-hyperbolic form (the intersection form of $S^\ell \times S^\ell$ for ℓ even and odd) corresponding to each basis element of A.

(ii) Two special geometric quadratic η-forms (A, λ, μ) and (A', λ', μ') of radius $\leq \delta$ are δ-isomorphic over B if there is a δ-isomorphism $h{:}A \to A'$ over B so that $\lambda'(h(x), h(y)) = \lambda(x, y)$ for all $x, y \in A$ and $\mu'(h(x)) = \mu(x)$ for all $x \in A$.

(iii) If (A, λ, μ) and (A', λ', μ') are geometric quadratic η-forms of radius $\leq \delta$, we will write $(A, \lambda, \mu) \overset{\delta}{\sim} (A', \lambda', \mu')$ if there are geometric modules F and G over B such that $(A', \lambda', \mu') \oplus \eta H(F)$ is δ-isomorphic to $(A, \lambda, \mu) \oplus \eta H(G)$.

(iv) We define $L_{\eta,B,\delta}(\mathbb{Z})$ to be the abelian semigroup of geometric special quadratic η-forms of radius $\leq \delta$, modulo the equivalence relation generated by $\overset{\delta}{\sim}$.

Here is the statement of the theorem which is our first main goal. It says that our surgery groups tell us, at least in the pro-sense, when we can do even-dimensional surgery.

Theorem 8.3 *Let $\ell \geq 3$ be given. Then $\exists\, \delta_0 > 0$ so that for $\delta < \delta_0$ the following holds: Let $\phi{:}(W^{2\ell}, \partial W) \to (P^{2\ell}, \partial P)$ be a degree one normal map with $p{:}P \to B$ and $p|\partial P$ both $UV^1(\delta)$-maps. Here, B is either a finite polyhedron or a polyhedron of bounded geometry. Let $\eta = (-1)^\ell$. Suppose, in addition, that $\phi|\,\partial W$ is a δ-equivalence. There is a number $k = k(n)$ so that that:*

(i) *There is a surgery obstruction $\sigma(\phi) \in L_{\eta,B,k\delta}(\mathbb{Z})$.*
(ii) *The image of $L_{\eta,B,\delta}(\mathbb{Z})$ in $L_{\eta,B,k\delta}(\mathbb{Z})$ is an abelian group.*
(iii) *$\sigma(\phi)$ is well-defined on normal bordism classes rel boundary in $L_{\eta,B,k^2\delta}(\mathbb{Z})$. In particular, if ϕ can be surgered to a δ-equivalence, then $\sigma(\phi) = 0$ in $L_{\eta,B,k^2\delta}(\mathbb{Z})$.*
(iv) *If $\sigma(\phi) = 0$ in $L_{\eta,B,k^2\delta}(\mathbb{Z})$, then ϕ can be surgered to a $k^3\delta$-equivalence.*
(v) *Every element of $L_{\eta,B,\delta}(\mathbb{Z})$ is realized on a manifold with two boundary components such that the restriction of ϕ to the first boundary component is the identity and the restriction of ϕ to the second boundary component is a $k\delta$-equivalence.*

Let $(P,\ \partial P)$ be a 2ℓ-dimensional δ-Poincaré duality pair with $UV^1(\delta)$-map $p{:}P \to B$ such that $p|\,\partial P$ is also $UV^1(\delta)$. Let $\phi{:}(N,\ \partial N) \to (P,\ \partial P)$ be a degree one normal map such that $\phi|\,\partial N$ is a δ-equivalence. We will assume that surgery has been done below the middle dimension as in Theorem 4.5.

The kernel complex is a $k'\delta$-Poincaré duality chain complex for some $k' = k'(n)$, with homology concentrated in dimension ℓ. Trading cells from the bottom and then flipping the complex over (algebraically) and trading down from the top shows that the kernel complex is $k''\delta$-chain homotopy equivalent to a complex of geometric modules with cells in only two dimensions, $\ell - 1$ and ℓ or, if we choose, ℓ and $\ell + 1$. This means that we have the diagram below, where φ and ψ are controlled chain-homotopy inverses.

$$
\begin{array}{ccccccccc}
0 & \longleftarrow & C_{\ell-1} & \overset{\partial}{\longleftarrow} & C_\ell & \longleftarrow & 0 & \longleftarrow & 0 \\
 & & \downarrow \psi & & \varphi \uparrow \ \downarrow \psi & & \varphi \uparrow & & \\
0 & \longleftarrow & 0 & \longleftarrow & C'_\ell & \overset{\partial}{\longleftarrow} & C'_{\ell+1} & \longleftarrow & 0
\end{array}
$$

Since the composition $\varphi \circ \psi$ is controlled chain homotopic to the identity, we have a diagram:

$$\begin{array}{ccccccc} 0 & \longleftarrow & C_{\ell-1} & \xleftarrow{\partial} & C_\ell & \longleftarrow & 0 \\ & & \downarrow \text{id} & \searrow^{s} & \downarrow \text{id} - \phi\circ\psi & & \\ 0 & \longleftarrow & C_{\ell-1} & \xleftarrow{\partial} & C_\ell & \longleftarrow & 0 \end{array}$$

with $\partial \circ s = \text{id}$.

Unfortunately, we are in a land without kernels, so this is not enough by itself to split the sequence

$$0 \longleftarrow C_{\ell-1} \underset{s}{\overset{\partial}{\rightleftarrows}} C_\ell \longleftarrow 0$$

However, we do have $(s \circ \partial) \circ (s \circ \partial) = (s \circ \partial)$, so $(s \circ \partial){:}C_\ell \to C_\ell$ is a controlled projection. The $\widetilde{K}_0$-vanishing result of Theorem 3.10 says that there exist a geometric projection g and ϵ deformations H_1 and H_2 so that $H_1((s \circ \partial) \oplus g)H_2$ is a geometric projection.

If we stabilize by doing trivial surgeries and adding cancelling pairs of ℓ- and $(\ell - 1)$-handles, we can assume that $H_1 \circ (s \circ \partial) \circ H_2$ is geometric. This means, in particular, that the image of $H_1 \circ s$ is geometric. By Chapman's swindle, we can assume, after stabilization, that H_1 is a product of no more than 13 elementary matrices.

Proposition 2.6 says that

$$0 \longleftarrow C_{\ell-1} \xleftarrow{\partial} C_\ell \longleftarrow 0$$

is isomorphic to the chain complex

$$0 \longleftarrow C_{\ell-1} \xleftarrow{\partial H_1^{-1}} C_\ell \longleftarrow 0$$

which is split by $H_1 s$. Since the image of $H_1 s$ is geometric, C_ℓ splits as geometric modules $C_\ell^1 \oplus C_\ell^2$ with $\partial H_1 | C_\ell^1{:}C_\ell^1 \to C_{\ell-1}$ a controlled isomorphism and $\partial H_1 | C_\ell^2 = 0$. Since the controlled Whitehaead group vanishes, another controlled handle slide arranges that $\partial H_1 | C_\ell^1$ takes generators to generators. At this point, we can cancel the cells in C_ℓ^1 against those of $C_{\ell-1}$.

The result is that the chain complex representing our surgery kernel has generators only in dimension ℓ. It follows from Poincaré duality and self-intersection, as in the classical case, that our surgery kernel has the structure of a special geometric $\mathbb{Z}$-quadratic η-form. This establishes part (i) of Theorem 8.3.

Next, suppose that the surgeries of part (i) have been performed and that our normal map $\phi{:}(N, \partial N) \to (P, \partial P)$ is normally cobordant rel boundary to another such degree one normal map, call it $\phi'{:}(N', \partial N) \to (P, \partial P)$. We can controlled surger the normal bordism rel ends and boundary to make the map from the bordism to $P \times I$ into an ℓ-connected map. By handle trading, first up from P and then down from the other end, it follows that ϕ is normally bordant to ϕ' via a bordism (with small handles, since any bordism can be subdivided) in dimensions ℓ and $\ell + 1$ and no handles outside of those dimensions.

We now look at the effect on the surgery kernel of passing through these layers of cells. Starting from the left in the diagram below, the first set of cells is trivially attached, so the algebraic effect on the surgery kernel is to add a geometric hyperbolic form to the algebraic kernel from P. This is the surgery kernel at the level of P'' below. On the other hand, we can begin at the right end of the bordism, where the surgery kernel is trivial, and see that adding the $(\ell + 1)$-cells, which become ℓ-cells when viewed from that side, exhibits P'''s surgery kernel as a hyperbolic form. Combined, this shows the forms representing the kernels on the two ends are stably equivalent. In particular, if ϕ' is a controlled homotopy equivalence, then the surgery kernel of ϕ must be controlled stably hyperbolic, i.e., it must be $\overset{k\delta}{\sim} 0$. Again, we have $k = k(n)$. Moreover, if the surgery kernel of a degree one normal map is stably controlled hyperbolic, we can proceed exactly as in [10], which is modelled on Chap. 5 of [21] to surger to a $k\delta$-equivalence for some $k = k(n)$. Since this step passes through geometry, we need to assume that the control space is a bounded geometry polyhedron endowed with the standard metric. This establishes parts (iii) and (iv) of Theorem 8.3.

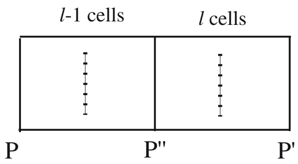

Finally, we note that the algebraic proof of Wall realization given in Theorem 5.8 of [21] works without modification to show that every geometric special quadratic $(-1)^{\ell}$-form can be realized on a manifold with boundary. The statement is given below.

Theorem 8.4 (Wall Realization) *Let B be a polyhedron with bounded geometry and let $n \geq 6$ be given. Then there exist $k = k(n)$ and δ_0 such that given a $UV^1(\delta)$-map $p: V^{n-1} \to B$ and $\alpha \in L_{n,B,\delta}(\mathbb{Z})$, $\delta < \delta_0$, we can represent the image of α in $L_{n,B,k\delta}(\mathbb{Z})$ by a map with $V \times I$ as target.*

We now return to salgebra. The next proposition is needed to prove part (ii) of Theorem 8.3.

Proposition 8.5 *If (A, λ, μ) is a geometric special quadratic η-form of radius $\leq \delta$, then $(A, \lambda, \mu) \oplus (A, -\lambda, -\mu) \overset{k\delta}{\sim} \eta H(A)$ for some $k = k(n)$.*

Proof We give a geometric argument, so our proof is only valid over polyhedra of bounded geometry. This is a deficiency of our "quick-and-dirty" geometric approach. It would be better to imitate Wall's algebraic proof in Lemma 5.4 of [21] and recover the result for arbitrary control spaces. Given a geometric special quadratic $(-1)^{\ell}$-form, we Wall realize it on a manifold $(W^{2\ell}; N, N')$ with boundary. Consider the $2\ell + 1$-dimensional surgery problem obtained by crossing our problem with $J = [0, 1]$.

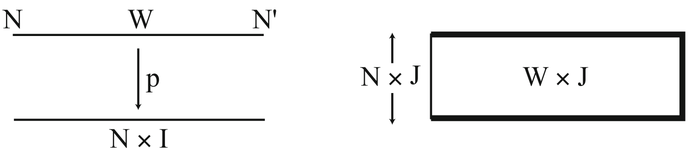

If (A, λ, μ) is the surgery kernel of $W \to N \times I$, then $W \times J$ gives a normal bordism from the dark region in the figure, which is $W \cup N' \times J \cup (-W)$, to $N \times J$, where we have a homeomorphism. The surgery kernel for the dark region is $(A, \lambda, \mu) \oplus (A, -\lambda, -\mu)$. This shows that a controlled surgery problem with kernel $(A, \lambda, \mu) \oplus (A, -\lambda, -\mu)$ can be solved and that the sum $(A, \lambda, \mu) \oplus (A, -\lambda, -\mu)$ is therefore stably hyperbolic. This completes the proof of Theorem 8.3 the k in the statement of the theorem is the maximum of the k's appearing in the proofs of parts (i)–(v). □

By part (ii) of Theorem 8.3, the system $\{L_{\eta,B,\delta_i}(\mathbb{Z})\}$ of abelian semigroups is equivalent to a system of groups. By abuse of notation, we will refer to $\{L_{\eta,B,\delta_i}(\mathbb{Z})\}$ as a system of groups.

We can now set up our system $\{L_{\eta,B,\delta_i}(\mathbb{Z})\}$ of even-dimensional algebraic surgery groups for B a polyhedron of bounded geometry. We can add elements of $L_{\eta,B,\delta_i}(\mathbb{Z})$ immediately this time, because direct sum does not increase the radius. Each element of $L_{\eta,B,\delta_i}(\mathbb{Z})$ has a negative in $L_{\eta,B,\delta_i}(\mathbb{Z})$, but the sum is only $\stackrel{k\delta}{\sim} 0$, so elements don't have inverses until we increase δ by a factor of k. Chapman's swindle then shows that the image of $L_{\eta,B,\delta_i}(\mathbb{Z})$ in $L_{\eta,B,\delta_{i-1}}(\mathbb{Z})$ is an abelian group, provided that δ_{i-1} is bigger than δ_i by a factor $k = k(n)$ for each i. Theorem 8.3 shows that in even dimensions and in a "pro" sense these groups do determine when surgery to a controlled homotopy equivalence is possible.

The next proposition shows that we can solve a controlled surgery problem over B if and only if we can solve (problem)$\times\mathbb{R}$ over $B \times \mathbb{R}$. Thus, we can always choose to work with even-dimensional surgery problems. Note that this is the basic philosophy of [10], as expressed in the introduction to that paper—the major difference being that here we use a product metric rather than a conelike metric.

Proposition 8.6 *Let $n \geq 6$ and let $p:(P^n, \partial P) \to B$ be $UV^1(\delta)$ with B a polyhedron of bounded geometry, and let $\phi:(N, \partial N) \to (P, \partial P)$ be a degree one normal map with $\phi|\,\partial N$ a δ-homotopy equivalence.*

- (i) *There exist a $k = k(n)$ and a $\delta_0 = \delta_0(n)$ so that if $\delta < \delta_0$ and $\phi \times \mathrm{id}:N \times \mathbb{R} \to (P, \partial P) \times R$ is normally cobordant rel ∂ to a δ-equivalence controlled over $B \times \mathbb{R}$, then ϕ is normally cobordant to a $k\delta$-equivalence controlled over B.*
- (ii) *If ϕ is normally cobordant to a δ-equivalence, then $\phi \times \mathrm{id} : N \times \mathbb{R} \to (P, \partial P) \times R$ is normally cobordant to a δ-equivalence.*

Proof Part (i) is a direct application of the controlled $M \times \mathbb{R}$ theorem stated in Sect. 2 plus the thin h-cobordism theorem. Part (ii) is clear. □

Definition 8.7 Let $\mathcal{N} = \mathcal{N}(M)$ be the set of normal bordism classes of degree one normal maps $(M, \partial M) \to (N, \partial N)$ which restrict to a homeomorphism on the boundary. The bordisms here should be through maps which are fixed on the boundary. By work of Sullivan, this collection is in one-to-one correspondence with $[N, \partial N;\ \mathrm{G/TOP}]$, making $\mathcal{N}$ into a group.

Remark 8.8 There also is a 1-1 correspondence between normal bordism classes of of degree one normal maps $W \to M \times \mathbb{R}$ and normal bordism classes of degree one normal maps $N \to M$. This means that the groups $L_{(-1)^{\ell+1},B\times\mathbb{R},\delta}(\mathbb{Z})$ give the obstructions for $(2\ell + 1) - \delta$-controlled dimensional surgery in the same "pro-" sense as the even-dimensional groups $L_{\eta,B,\delta}(\mathbb{Z})$ which we have already discussed.

The map in one direction is the product and the map in the other direction is given by transversality. To see that the composition $\mathcal{N}(M\times\mathbb{R}) \to \mathcal{N}(M) \to \mathcal{N}(M\times\mathbb{R})$ is the identity on bordism, note that we can arrange for the composition to be the identity on $B \times \{t_0\}$ for some t_0 and then enlarge a collar to make the composition equal to the identity everywhere. Again, this is the basic approach of [10].

Notation 8.9 By $L_{n,B,\delta}(\mathbb{Z})$, we will mean $L_{(-1)^{\ell},B,\delta}(\mathbb{Z})$ for $n = 2\ell$ even and $L_{(-1)^{\ell+1},B\times\mathbb{R},\delta}(\mathbb{Z})$ for $n = 2\ell + 1$ odd.

9 Stability of controlled wall systems for *n* large

Definition 9.1 A system consisting of groups $A_1 \longleftarrow A_2 \longleftarrow A_3 \longleftarrow \cdots$ is stable if it is equivalent to a sequence of groups and isomorphisms.

Our goal in this section is to prove the stability of the sequence

$$L_{q,B,\delta_1}(\mathbb{Z}) \longleftarrow L_{q,B,\delta_2}(\mathbb{Z}) \longleftarrow L_{q,B,\delta_3}(\mathbb{Z}) \longleftarrow \cdots$$

for all q whenever $\{\delta_i\}$ is a sequence of positive real numbers converging monotonically to zero. We will accomplish this by first proving the result for $q \geq 2\dim B + 2$ and then noting that the result for arbitrary q follows from the periodicity of the Wall groups. Let $N = N(B)$ be a regular neighborhood of B in $\mathbb{R}^q$, $q \geq 2\dim B + 2$, and let $\mathcal{N}$ be the set of normal bordism classes of degree one normal maps $(M, \partial M) \to (N, \partial N)$ which restrict to a homeomorphism on the boundary. By work of Sullivan, this collection is in one-to-one correspondence with $[N,\ \partial N;\ \text{G/TOP}]$, making $\mathcal{N}$ into a group.

Taking rel boundary δ-surgery obstructions gives us the commuting diagram below. In this special case, $N = N(B)$ has the form $N' \times [0,\ 1]$, where N' is a regular neighborhood of B in $\mathbb{R}^{q-1}$. Addition in $\mathcal{N}$ can be defined by gluing elements together along pieces of the boundary. The vertical maps in this diagram are homomorphisms, see [21], p. 111 for details.

$$\begin{array}{ccccccc}
\mathcal{N} & \xleftarrow{\text{id}} & \mathcal{N} & \xleftarrow{\text{id}} & \mathcal{N} & \longleftarrow & \cdots \\
\downarrow \sigma & & \downarrow \sigma & & \downarrow \sigma & & \\
L_{q,B,\delta_1}(\mathbb{Z}) & \longleftarrow & L_{q,B,\delta_2}(\mathbb{Z}) & \longleftarrow & L_{q,B,\delta_3}(\mathbb{Z}) & \longleftarrow & \cdots
\end{array}$$

We will show that the vertical maps are eventual isomorphisms. We begin with the argument for surjectivity. Consider an element α of $L_{q,B,\delta_2}(\mathbb{Z})$. By a rel boundary adaptation of Wall realization, this element is realized by a degree one normal map $\phi{:}(W, \partial W, N', N'') \to (N' \times I, \partial(N' \times I), N' \times 0, N' \times 1)$ where N' is a regular neighborhood of B in $\mathbb{R}^{q-1}$ as before and ϕ restricts to homeomorphisms over $N' \times 0 \cup \partial N' \times I$ and a $k\delta_2$-equivalence over B on N'', where $k = k(q)$. In order to show that ϕ is in the image of $\mathcal{N}$, we need to show that ϕ is normally bordant to ϕ', where ϕ' is a homeomorphism on the entire boundary.

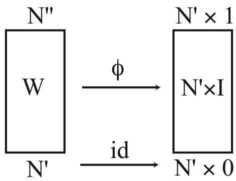

The restriction of ϕ to N'' is a homeomorphism on $\partial N''$. In order to show that ϕ is in the image of $\mathcal{N}$, we apply the alpha approximation theorem (Theorem 1.5 of the introduction) together with a variation on the classical Alexander trick. The figure below represents the restriction of ϕ to the top boundary of the previous picture.

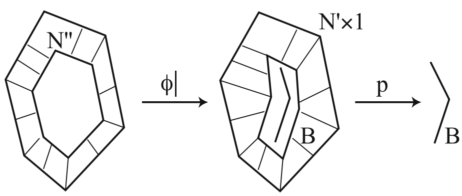

Here is our Alexander trick. Let p be the mapping cylinder projection $p{:}N'' \to B$. The homotopy equivalence ϕ is controlled over this map. Now, repeat the boundary homeomorphism over a boundary collar in N'', rescaling so that the image of this homeomorphism contains all but a very small regular neighborhood of B in $N' \times 1$. After this modification, $\phi|N''$ becomes a controlled homotopy equivalence over $N' \times 1$, not just over B, and we can use the alpha approximation theorem to find a small homotopy from $\phi|N''$ to a homeomorphism.[11] This means that if $k\delta_2$ is chosen to be so small that alpha approximation works on N' to produce a homeomorphism δ_1-homotopic to the original ϕ, then the image of α in $L_{q,B,\delta_1}(\mathbb{Z})$ is in the image of $\mathcal{N}$. This proves eventual surjectivity and we pass to a subsequence so that "surjectivity" takes place in the $(i-1)^{st}$ place for each i.

Now for the proof of injectivity. Consider an element $\phi \in \mathcal{N}$ whose image in, say, $L_{q,B,\delta_3}(\mathbb{Z})$ is trivial. This means that $\phi{:}M \to N$ with ϕ a homeomorphism over ∂N and that ϕ is bordant as a restricted δ_3-object to a δ_3-equivalence. We denote this bordism by $\Phi{:}P \to N \times I$. Our goal is to show that ϕ is equivalent in $\mathcal{N}$ to a homeomorphism.

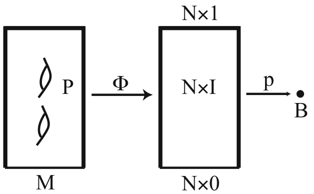

The restriction of Φ to $\partial P - \overset{\circ}{M}$ gives a δ_3-equivalence to $\partial(N \times I) - \overset{\circ}{N \times 1}$ which is a homeomorphism on the boundary. Reparameterizing, this gives a bordism from ϕ to a δ_3-equivalence over the shaded area. Using the same Alexander trick, the restriction of Φ to the shaded area is ϵ-homotopic rel boundary to a homeomorphism, where ϵ is related to δ_3 as in the alpha approximation theorem. This shows that the element $[\phi]$ is trivial in $\mathcal{N}$ and that the sequence $\{L_{q,B,\delta_i}(\mathbb{Z})\}$ is equivalent to the sequence $\mathcal{N} \leftarrow \mathcal{N} \leftarrow \cdots$.

The group $\mathcal{N}$ is isomorphic to $[N, \partial N; \mathrm{G/TOP}]$, which is in turn isomorphic to $H_q(N, \mathrm{G/TOP})$, $H_q(B, \mathrm{G/TOP})$, and $H_q(B, \mathbb{L}(e))$. Thus, we have shown:

Proposition 9.2 *Given B and $q \geq 2\dim B + 2$, we can choose a sequence δ_i of positive real numbers monotonically approaching zero so that the image of $L_{q,B,\delta_i}(\mathbb{Z})$ in $L_{q,B,\delta_{i-1}}(\mathbb{Z})$ is isomorphic to $H_q(B, \mathbb{L}(\mathbb{Z}))$ for all i. By periodicity of $L_{q,B,\delta_i}(\mathbb{Z})$ and $H_q(N, \mathrm{G/TOP})$, this establishes the same result for all q*

Remark 9.3 The proof above is closely related to the proof of the surgery exact sequence. The point of our Alexander trick is that it allows us to use the alpha approximation theorem to show that the rel boundary structure set $\mathcal{S}_\delta\begin{pmatrix} N \\ \downarrow \\ B \end{pmatrix}$ is equivalent to a trivial system and then t is equivalent to a trivial system and then to use this fact to prove that the normal maps are isomorphic to the surgery groups as systems. Rather than set up a structure set which is going to be zero (in this specific instance, anyway) we have chosen to delay setting up the official surgery sequence and make our argument on the level of individual representatives of elements of the structure-set-to-be.

[11] The key point here is that the alpha-approximation theorem only requires that homotopies have small diameter after projecting to the target. By stretching out a collar where $\phi|$ is a homeomorphism, we have replaced control over B by control over $N' \times 1$.

10 Stability of controlled structure sets

Next, we want to show how to use stability of the controlled L-groups to prove a similar stability result for manifold structures. We can always replace a $UV^1(\delta)$-map $p:M^n \to B$, M a manifold and $n \geq 5$, by a UV^1-map. Here is a theorem from [13] which is a modified version of a theorem of Bestvina [1].

Theorem 10.1

(i) *Let a finite polyhedron B and $n \geq 5$ be given. Then given $\epsilon > 0$ there is a $\delta > 0$ so that if $p':N^n \to B$ is a $UV^1(\delta)$-map from a compact manifold to B, then p' is ϵ-homotopic to a UV^1-map $p:N \to B$.*

(ii) *Let N be a compact manifold and suppose that a UV^1-map $q:N \to B$ onto a polyhedron B is given. Then for each $\epsilon > 0$ there is a $\delta = \delta(\epsilon, q) > 0$ such that for each map $f:M \to N$ of any compact PL manifold M with* $\dim M \geq 5$ *which is* 1*-connected with δ-control with respect to q, there is a UV^1-map $g:M \to N$ which is ϵ-close to f as measured in B.*

Here is the long-awaited stabilty theorem for structures.

Theorem 10.2 *Let $n \geq 5$. If M^n is a closed n-manifold and $p:M \to B$ is a UV^1-map, then for every $\epsilon > 0$ there is a $\delta > 0$ so that for any $\mu > 0$, if $\phi:N \to M$ is a δ-homotopy equivalence over B then there is an ϵ-homotopy over B from ϕ to a μ-homotopy equivalence over $B\psi:N \to M$. A similar result holds for rel boundary structures if M is a manifold with boundary.*

In other words, if we are willing to allow a homotopy of fixed size, then if we start with a sufficiently well-controlled homotopy equivalence we can improve the control of that homotopy equivalence by an arbitrary amount. This is the "squeezing theorem" for structures.[12]

Proof As above, the proof consists of working our way through the not-yet-existent surgery exact sequence. We start with the systems $\{L_{q,B,\delta_i}(\mathbb{Z})\}$ reindexed so that we have commutative diagrams:

$$
\begin{array}{ccccccc}
\mathcal{N}(M) & \xleftarrow{\text{id}} & \mathcal{N}(M) & \xleftarrow{\text{id}} & \mathcal{N}(M) & \longleftarrow & \cdots \\
\downarrow{\scriptstyle\sigma} & & \downarrow{\scriptstyle\sigma} & & \downarrow{\scriptstyle\sigma} & & \\
L_{n,B,\delta_1}(\mathbb{Z}) & \longleftarrow & L_{n,B,\delta_2}(\mathbb{Z}) & \longleftarrow & L_{n,B,\delta_3}(\mathbb{Z}) & \longleftarrow & \cdots \\
\uparrow & \swarrow{\scriptstyle\gamma_2} & \uparrow & \swarrow{\scriptstyle\gamma_3} & \uparrow & & \\
H_n(B;\mathbb{L}(\mathbb{Z})) & \longleftarrow & H_n(B;\mathbb{L}(\mathbb{Z})) & \longleftarrow & H_n(B;\mathbb{L}(\mathbb{Z})) & \longleftarrow & \cdots
\end{array}
$$

and

$$
\begin{array}{ccccccc}
\mathcal{N}(M\times I) & \xleftarrow{\text{id}} & \mathcal{N}(M\times I) & \xleftarrow{\text{id}} & \mathcal{N}(M\times I) & \longleftarrow & \cdots \\
\downarrow{\scriptstyle\sigma} & & \downarrow{\scriptstyle\sigma} & & \downarrow{\scriptstyle\sigma} & & \\
L_{n+1,B,\delta_1}(\mathbb{Z}) & \longleftarrow & L_{n+1,B,\delta_2}(\mathbb{Z}) & \longleftarrow & L_{n+1,B,\delta_3}(\mathbb{Z}) & \longleftarrow & \cdots \\
\uparrow & \swarrow{\scriptstyle\gamma_2} & \uparrow & \swarrow{\scriptstyle\gamma_3} & \uparrow & & \\
H_{n+1}(B;\mathbb{L}(\mathbb{Z})) & \longleftarrow & H_{n+1}(B;\mathbb{L}(\mathbb{Z})) & \longleftarrow & H_{n+1}(B;\mathbb{L}(\mathbb{Z})) & \longleftarrow & \cdots
\end{array}
$$

[12] Note that we do not claim to have constructed a fixed map which is a μ-equivalence for all μ.

Assume, further, that the δ_i's have been chosen so that for $q = n,\ n+1$

(i) Degree one normal maps $\phi: N \to M$ with vanishing surgery obstruction in $L_{q,B,\delta_i}(\mathbb{Z})$ are normally bordant to δ_{i-1}-equivalences.
(ii) Each $\alpha \in L_{q,B,\delta_i}(\mathbb{Z})$ can be Wall-realized by a normal bordism to an δ_{i-1}-equivalence over B.
(iii) Each δ_i-thin h-cobordism over B has a δ_{i-1}-product structure over B.
(iv) $\delta_0 = \epsilon$ and $\delta_4 = \mu$.

We take $\delta = \delta_3$. Suppose that we are given a δ_3-controlled homotopy equivalence $\phi: N \to M$. The surgery obstruction of $[\phi]$ vanishes in $L_{n,B,\delta_3}(\mathbb{Z})$, so an easy diagram chase shows that the surgery obstruction of $[\phi]$ vanishes in $L_{n,B,\delta_i}(\mathbb{Z})$ for all $i \geq 3$, so $[\phi]$ is normally bordant to some δ_i-equivalence for each i. Choose a δ_5-equivalence $\phi_5: N_5 \to M$ normally bordant to ϕ. The normal bordism $\Phi_5: W_5 \to M \times I$ has a surgery obstruction α in $L_{n,B,\delta_3}(\mathbb{Z})$. Wall realize an element $\bar{\alpha} \in L_{n,B,\delta_5}(\mathbb{Z})$ whose image in $L_{n,B,\delta_2}(\mathbb{Z})$ is the same as the image of $-\alpha$ starting with $\phi_5: N_5 \to M$ to get a normal bordism $\Phi: W \to M \times I$ from ϕ_5 to a δ_4-homotopy equivalence $\phi': N' \to M$.

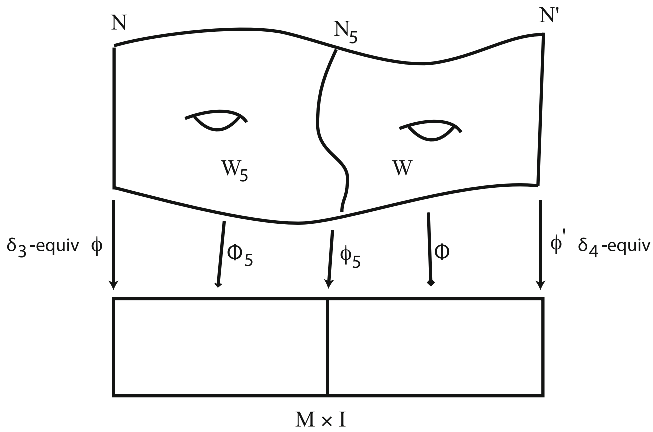

The obstruction to surgering $W_5 \cup W$ to a controlled h-cobordism dies in $L_{n,B,\delta_2}(\mathbb{Z})$, so we can surger the bordism rel boundary to a δ_1-h-cobordism which has a δ_0-product structure. This provides the desired ϵ-homotopy from ϕ to a δ_4-equivalence $\phi' \circ h$, where $h: N \to N'$ is the homeomorphism coming from the product structure. □

We now repeat our definition of δ-controlled structure sets.

Definition 10.3 Let M be a closed manifold and let $p: M \to B$ be a UV^1-map. We define $\mathcal{S}'_\delta \begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$ to consist of equivalence classes of δ-homotopy equivalences $\phi: N \to M$ over B modulo the relation that $\phi: N \to M$ and $\phi': N' \to M$ are equivalent if there is a homeomorphism $h: N \to N'$ so that ϕ is δ-homotopic to $\phi' \circ h$ over B. We declare $\mathcal{S}_\delta \begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$ to be the image of $\mathcal{S}'_{\delta_i} \begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$ in $\mathcal{S}'_{\delta_{i-1}} \begin{pmatrix} M \\ \downarrow \\ B \end{pmatrix}$.

Given stability of the system of surgery groups, Proposition 10.2 shows that the system $\left\{ \mathcal{S}_{\delta_i} \left(\begin{array}{c} M \\ \downarrow \\ B \end{array} \right) \right\}$ is stable as a sequence of sets, that is, it is equivalent to a sequence of bijections. The proposition shows immediately that the sequence is equivalent to a sequence of surjections and the relative version of the proposition shows that the sequence is equivalent to a system of injections. This completes our proof of stability in dimensions ≥ 5.

Remark 10.4 As usual, all of the above extends to the bounded geometry case. The most interesting extra step in this case is that we need to find a bounded geometry thickening of a polyhedron of bounded geometry. For this, we note that a polyhedron B of complexity ℓ immerses in a simplex of dimension $\ell + 1$. We obtain the required thickening of B by pulling back a regular neighborhood of $\Delta^{\ell+1}$ in $\mathbb{R}^{\ell+1}$. The rest of the argument goes through as above.

11 The surgery sequence, functoriality, and everything else

If any neophytes have made it this far, here is the proof of the surgery exact sequence, which we recall below:

$$\ldots H_{n+1}(B; \mathbb{L}) \dashrightarrow \mathcal{S}_\epsilon \left(\begin{array}{c} M \\ \downarrow \\ B \end{array} \right) \longrightarrow \mathcal{N}(M) \longrightarrow H_n(B; \mathbb{L})$$

If we start with a normal map, we can take its controlled surgery obstruction. If that dies, then it can be surgered to an ϵ-equivalence for any $\epsilon > 0$. Thus, it comes from a controlled structure. Conversely, if the normal map comes from a controlled structure, then its controlled surgery obstruction dies. Starting with a controlled structure, we can consider it as a normal map and ask if it is normally bordant to a homeomorphism. If it is, there is a controlled relative surgery obstruction in $H_{n+1}(B; \mathbb{L})$ to surgering the normal bordism to a controlled product. If this obstruction dies, then we can surger to a controlled product, which shows that the original structure woas controlled homotopic to a homeomorphism.

Wall realization gives an action of $H_{n+1}(B; \mathbb{L})$ on the controlled structures: start with a controlled homotopy equivalence $f: N \to M$ and Wall realize an element $\alpha \in H_{n+1}(B; \mathbb{L})$, obtaining a bordism (W_α, N, N'). The controlled equivalence $N' \to M$ is the result of acting on f by α. If a controlled structure goes to 0 in the normal maps, then it is obtained by acting on $\mathrm{id}: M \to M$ by some α.

If we have a UV^1-map $p: B \to B'$, then the entire surgery sequence controlled over B maps to the surgery sequence controlled over B'. This also works if the map is $UV^1(\delta)$ for $\delta \leq \delta_0$ with δ_0 depending on B' and p. If p happens to be Lipschitz with constant 1, then δ_0 only depends on B'.

Acknowledgment The author would like to thank the University of Chicago for hospitality while this paper was being written.

References

1. Bestvina, M.: Characterizing k-dimensional universal Menger compacta. Mem. Am. Math. Soc. **71**(380), vi+110 MR 89g:54083 (1988)

2. Bridson, M.R., Haefliger, A.: Metric spaces of non-positive curvature, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol 319. Springer, Berlin (1999) MRMR1744486 (2000k:53038)
3. Browder, W., Levine, J., Livesay, G.R.: Finding a boundary for an open manifold. Am. J. Math. **87**, 1017–1028 MR32#6473 (1965)
4. Chapman, T.A.: Locally homotopically unknotted embeddings of manifolds in codimension two are locally flat. Topology **18**(4), 339–348 MR 81a:57016 (1979)
5. Chapman, T.A. : Controlled Simple Homotopy Theory and Applications. Springer, Berlin MR 85d:57076 (1983) MR85d:57016
6. Chapman, T.A., Ferry, S.: Approximating homotopy equivalences by homeomorphisms. A. J. Math. **101**(3), 583–607 MR 81f:57007b (1979)
7. Cheeger, J., Müller, W., Schrader, R.: On the curvature of piecewise flat spaces. Commun. Math. Phys. **92**(3), 405–454 MR 85m:53037 (1984)
8. Connell, E.H., Hollingsworth, J.: Geometric groups and whitehead torsion. Trans. Am. Math. Soc. **140**, 161–181 MR 39#3485 (1969)
9. Ferry, S.: Homotopy, simple homotopy and compacta. Topology **19**(2), 101–110 MR 81j:57010 (1980)
10. Ferry, S.C., Pedersen, E.K.: Epsilon Surgery Theory. Novikov Conjectures, Index Theorems and Rigidity, vol. 2, pp.167–226 (Oberwolfach, 1993). Cambridge University Press, Cambridge MR 97g:57044 (1995)
11. Ferry, S.C., Weinberger, S.: Curvature, tangentiality, and controlled topology. Invent. Math. **105**(2), 401–414 MR 94c:57043 (1991)
12. Hastings, H.M.: Addendum to: Cech and Steenrod Homotopy Theories with Applications to Geometric Topology [Lecture Notes in Math., 542, Springer, Berlin, 1976; MR **55**#1347] by Hastings and David A. Edwards, Colloq. Math. **41**(1), 161–163 MR 80m:55020 (1979)
13. Kawamura, K.: An inverse system approach to Menger manifolds. Topol. Appl. **61**(3), 281–292 MR96a:54050 (1995)
14. Kirby, R.C., Laurence, C.: Siebenmann, Foundational Essays on Topological Manifolds, Smoothings, and Triangulations. Princeton University Press, Princeton, NJ. With notes by John Milnor and Michael Atiyah. Ann. Math. Stud. (88) MR 58#31082 (1977)
15. Lacher, R.C.: Cell-like mappings and their generalizations. Bull. Am. Math. Soc. **83**(4), 495–552 MR 58#31095 (1977)
16. Pedersen, E.K., Quinn, F., Ranicki, A.: Controlled surgery with trivial local fundamental groups, high-dimensional manifold topology, pp. 421–426. World Science Publication, River Edge, NJ, MR MR2048731 (2005e:57077) (2003)
17. Quinn, F.: Ends of maps. I. Ann. Math. **110**(2), 275–331 MR 82k:57009 (1979)
18. Quinn, F.: Ends of maps. III. Dimensions 4 and 5. J. Diff. Geom. **17**(3), 503–521 MR 84j:57012 (1982)
19. Ranicki, A.: The algebraic theory of surgery. II. Applications to topology. Proc. Lond. Math. Soc. **40**(2), 193–283 MR 82f:57024b (1980)
20. Ranicki, A.: The algebraic theory of finiteness obstruction. Math. Scand. **57**(1), 105–126 MR 87d:18014 (1985)
21. Wall, C.T.C.: Surgery on Compact Manifolds. Academic Press, London (1970), Lond. Math. Soc. Monogr. (1) MR 55 #4217
22. Yamasaki, M.: L-groups of crystallographic groups. Invent. Math. **88**(3), 571–602 MR 88c:57017 (1987)