Geometric Algebra

Frank Quinn
Department of Mathematics
Virginia Polytechnic Institute and State University
Blacksburg, Virginia USA 24061

Introduction

In this paper we develop a hybrid sort of algebra, whose morphisms involve paths in a space. The primary purpose is to elucidate and extend the algebra with ϵ estimates developed in [3, 4-6, 1]. The setting is also fruitful for investigating relationships between the topology of a space X, and the algebra of $R[\pi_1 X]$ modules.

The first section presents the definitions of geometric R-modules on a space, and their morphisms. We show that by allowing appropriate "homotopies" of morphisms we can recover either ordinary $R[\pi_1 X]$ homomorphisms of free modules, or ϵ homomorphisms of geometric modules. Then we show that if $K \longrightarrow X$ is a map from a CW complex to a space, the cellular chains of K can be seen in a very natural way as a geometric chain complex on X.

Section two gives decompositions of $R[\pi]$ chain complexes, corresponding to amalgamated free product decompositions of π. The approach is to geometrically realize the free product structure as the structure induced on the fundamental group of a space X by a codimension 1 subspace Y. Then we use the equivalence of 1.1 to represent chain complexes by geometric ones on X. Intersections of the geometric structure of the complex with Y then show how to decompose the complex. The main purpose of this is to illustrate the technique, which we anticipate will apply to algebraic K and L theory.

Finally in section three, geometric versions of the Whitehead group are defined. These are shown to be the obstruction groups for the thin h-cobordism theorem, a controlled version of the usual result.

Section 1: Geometric modules and morphisms

Suppose that X is a topological space and R a ring. A *Geometric* R-*module* on X is defined to be a free module R[S] and a map of the basis $f:S \longrightarrow X$. We require that geometric modules be *locally finite* in the sense that every point in X has a neighborhood whose preimage in S is finite. So for example a geometric module on a compact space has a finite basis.

A *geometric morphism* of geometric modules is defined to be a locally finite

Partially supported by the National Science Foundation

algebraic sum of paths between generators. More specifically suppose $f_i:S_i \longrightarrow X$ are bases for geometric modules, i = 0,1. A morphism $h:R[S_0] \longrightarrow R[S_1]$ is a sum $\Sigma m_j\lambda_j$ where $m_j \epsilon R$ and λ_j is a path. The data for a path consists of elements $x_i \epsilon S_i$ and a map $\lambda:[0,t] \longrightarrow X$ with $\lambda(0)=f_0(x_0)$ and $\lambda(t)=f_1(x_1)$. Here t is a real number $t \geq 0$. Finally we require that for each $y \epsilon S_0$ there are only finitely many paths λ_j starting at y which have nonzero coefficient.

In a morphism we allow deletion of a path with coefficient 0 (or conversely insertion of such a path). We also identify $(m+n)\lambda$ with $(m\lambda)+(n\lambda)$.

We describe how to compose two morphisms. If $f = \Sigma m_j\lambda_j:R[S_1] \longrightarrow R[S_2]$, and $g = \Sigma n_k\alpha_k:R[S_0] \longrightarrow R[S_1]$, then $fg = \Sigma(m_jn_k)(\lambda_j\alpha_k)$. The sum is taken over all pairs (j,k) such that the end of α_j is equal to the starting point of λ_k in S_1. The Moore composition of paths is used: given $\lambda:[0,t] \longrightarrow X$ and $\alpha:[0,u] \longrightarrow X$ then $\lambda\alpha:[0,t+u] \longrightarrow X$ is defined by $\lambda\alpha(s)=\alpha(s)$ for $s \leq u$ and $\lambda\alpha(s)=\lambda(s-t)$ for $s \geq u$. Notice we are writing compositions of paths from right to left, so that it will agree with the notation for composition of functions.

The composition is associative, and there is a unit (the unit in the ring times the constant paths defined on [0,0]). Geometric modules and morphisms therefore form a category. This is not a directly useful category because there are too many paths.

There is a forgetful functor from geometric morphisms to ordinary R-module homomorphisms, defined by forgetting the paths. Explicitly, if $h = \Sigma m_j\lambda_j:R[S_0] \longrightarrow R[S_1]$ then we can define an R-homomorphism $h':R[S_0] \longrightarrow R[S_1]$ by $h'(s) = \Sigma_i(\Sigma_j m_j)t_i$. Here the outer summation is over $t_i \epsilon S_1$, and the inner summation is over j such that the path λ_j goes from s to t_i.

We will define several notions of "homotopy" of geometric morphisms. The goal is to obtain useful intermediate stages between the rigidity of geometric morphisms and the laxity of ordinary algebra over R.

1.1 Unrestricted homotopy of morphisms A homotopy of a morphism is obtained by changing all the paths in the morphism by homotopy holding the endpoints fixed. Since we are using Moore paths, a "homotopy" is allowed to change the interval on which the path is defined. Form the category whose morphisms are homotopy classes of morphisms of geometric R-modules on X. We claim that if X is connected and locally 1-connected then this category is naturally equivalent to the category of free $R[\pi_1 X]$ modules, with a restriction on rank. (If X is compact, the modules are finitely generated. If X is noncompact and separable the modules are countably generated, etc.) To simplify the discussion assume that X is compact.

Let $\tilde{X}$ denote the universal cover of X, which exists since X is locally 1-connected. Given a geometric module R[S], with map $S \longrightarrow X$, form the pullback

$$\begin{array}{ccc} \tilde{S} & \longrightarrow & \tilde{X} \\ \downarrow & & \downarrow \\ S & \longrightarrow & X \end{array}$$

then the action of $\pi_1 X$ on $\tilde{S}$ gives $R[\tilde{S}]$ the structure of a finitely generated free $R[\pi_1 X]$ module.

Next suppose that $\Sigma m_j \lambda_j : R[S_0] \longrightarrow R[S_1]$ is a morphism of geometric R-modules on X. The paths lift into the universal cover to give a $\pi_1 X$ equivariant family of paths from $\tilde{S}_0$ to $\tilde{S}_1$. This defines a lift of the morphism itself to an equivariant morphism $R[\tilde{S}_0] \longrightarrow R[\tilde{S}_1]$. Forget the paths to get a $R[\pi_1 X]$ homomorphism.

Notice that there is a unique homotopy class of paths between any two points in $\tilde{X}$, so no information is lost in forgetting the paths in the lifted morphism.

Now we go the other way, from $\pi_1 X$ modules to geometric modules. To a free module $R[\pi_1 X][S]$ we associate the geometric module $R[S]$, with $S \longrightarrow X$ the map to the basepoint. To a $\pi_1 X$ homomorphism $\Sigma m_i p_i$ with $p_i \in \pi_1 X$, choose representative loops α_i for p_i and form the geometric morphism $\Sigma m_i \alpha_i$.

It is straightforward to see that the constructions are inverses, and define an equivalence of categories. The benefits of thinking of $\pi_1 X$ modules this way are explored in section 2.

1.2 ϵ homotopy Suppose X is a metric space, and $\epsilon > 0$. We say a *homotopy* $h: Y \times I \longrightarrow X$ *has radius less than* ϵ if for each $y \in Y$ the arc $h(y \times I)$ lies in the ball of radius ϵ about $h(y,0)$. In particular this gives a notion of ϵ homotopy of morphisms of geometric modules. This notion is most useful when the morphisms themselves are small. We say a *morphism has radius less than* ϵ if each path λ_i in the morphism lies in the ball of radius ϵ about its starting point $\lambda_i(0)$.

Notice that ϵ homotopy is not an equivalence relation: the composition of ϵ homotopies has radius at best 2ϵ. In fact the situation is often worse than this. If X is not compact then it is necessary to use control *functions* $\epsilon: X \longrightarrow (0, \infty)$, (in which case the ball of radius ϵ at x means the ball of radius $\epsilon(x)$). When ϵ is a function the composition of two ϵ homotopies may be much larger than 2ϵ. We describe how to deal with this in section 3.1.

If the paths in an ϵ morphism are discarded, we get a homomorphism $f': R[S_0] \longrightarrow R[S_1]$ with the property that if we write $f'(s_i)$ as $\Sigma m_{i,j} t_j$ then the coefficient $m_{i,j}$ is zero if t_j is not in the ϵ ball about s_i. This is an ϵ homomorphism in the sense of Connell and Hollingsworth [3], and Quinn [4, 5]. Morphisms which are ϵ homotopic determine the same ϵ homomorphism. As with $\pi_1 X$ homomorphisms if X is locally 1-connected there is a converse to this construction, at least in the appropriate ϵ sense.

Suppose X is locally 1-connected. Then given $\epsilon>0$ (a function if X is not compact) there is $\delta>0$ such that any two points within δ can be joined by a path of radius ϵ. This means paths can be chosen to represent a homomorphism of radius less than δ as an ϵ geometric morphism. Similarly there is γ so that loops of radius less than γ are nullhomotopic by homotopies of radius less than ϵ. This means that any two representations of a homomorphism by geometric morphisms of radius less than γ are ϵ homotopic. Together these observations imply that for sufficiently small γ, γ homomorphisms determine geometric morphisms well defined up to ϵ homotopy.

If X is not locally 1-connected then this correspondence between metric and geometric ϵ theories breaks down. For many purposes it is the geometric theory which is more fundamental. A precursor of geometric morphisms was developed by Chapman [1] to allow non-locally 1-connected control spaces X in certain controlled manifold theorems, as in section 3.

1.3 Controlled homotopy This is a combination of 1.1 and 1.2: suppose $f:E \longrightarrow X$ is a map, X is a metric space, and $\epsilon>0$. Consider homotopies of morphisms in E whose compositions with f have radius less than ϵ (in X).

Suppose f is a projection of a product $X \times Y \longrightarrow X$. If Y is locally 1-connected we can use the universal cover as in 1.1 to obtain $R[\pi_1 Y]$ homomorphisms of geometric modules over X. If X is also locally 1-connected, then we can proceed as in 1.2 to see that the geometric theory of R-modules on $X \times Y$ with ϵ control in X is essentially equivalent to the ϵ metric theory of $R[\pi_1 Y]$-modules on X.

In some more general situations we can generalize from the product situation and think of geometric algebra on E with ϵ control in X as being like $R[\pi_1 f^{-1}(x)]$ metric algebra. In other words let the coefficient ring vary from point to point in X. In some cases (eg. if E is a "stratified system of fibrations over X", Quinn [5]) this can be made precise. In general, however, it seems best to stick with the geometric description.

The controlled version will be applied in section 3.

1.4 Geometric cellular chains Suppose that K is a CW complex, and $f:K \longrightarrow X$ is a map. We interpret the cellular chain complex of K as a geometric $\mathbb{Z}$ complex over X.

The cellular chain group $C_k(K)$ is the free abelian group generated by the k-cells of K. To give this the structure of a geometric module we introduce notation for the cells in K. Let K^k denote the k-skeleton. Let $\theta_s:D^k \longrightarrow K^k$ denote inclusions of k-cells, where s is in an index set S_k. C_k is then by definition $\mathbb{Z}[S_k]$. Define functions $S_k \longrightarrow X$ by mapping s to $0 \in D^k$, applying θ_s to get a point in K, and then applying f to get a point in X. To ensure that this is locally finite we should assume something like: each point in X has a neighborhood U such that $f^{-1}(U)$ is contained in a finite subcomplex. Assuming this the C_k become geometric $\mathbb{Z}$-modules.

The next step is to define geometric boundary homomorphisms $\partial:C_k \longrightarrow C_{k-1}$. In the definition of a CW complex the maps of the k-cells carry the boundary into the k-1 skeleton: $\partial\theta_s:S^{k-1} \longrightarrow K^{k-1}$. These are the attaching maps for the k-cells, so in fact K^k is defined to be K^{k-1} with cells attached by these $\partial\theta_s$. If k>1 then the boundary homomorphism is defined by $\partial s_i = \Sigma d_{i,j}t_j$, where $s_i \in S_k$, $t_j \in S_{k-1}$, and $d_{i,j}$ is the degree of the map $\partial\theta_{si}$ on the cell $\theta_{tj}:D^{k-1} \longrightarrow K^{k-1}$.

Assume that the attaching maps for the k-cells are transverse to the center points of the k-1 cells. The inverse image $(\partial\theta_s)^{-1}(0_t)$ is then a finite set of points, and at each point there is a sign +1 or -1 depending on whether $\partial\theta$ preserves or reverses orientation at that point. The degree of $\partial\theta_s$ on the cell θ_t is the sum of these signs. Define paths in X by taking the radial line in D^k from 0 to the points $(\partial\theta_s)^{-1}(0_t)$ and composing with θ_s and f. The geometric boundary morphism is defined by adding up these paths times the sign of $\partial\theta_s$ on the endpoint. It is clear from the construction that forgetting the paths yields the ordinary boundary homomorphism.

The boundary $\partial:C_1 \longrightarrow C_0$ is defined slightly differently, since degrees are not defined for 0-cells. The 1-cells are arcs, and the ordinary boundary of a 1-cell is defined to be the beginning point minus the endpoint. The geometric boundary is defined to be the arc from the center to the beginning, minus the arc from the center to the end.

It is suggested that the reader draw a picture of a 2-simplex, and draw in the geometric chain groups and boundary morphisms.

Next consider the composition $\partial\partial$. In traditional complexes this is equal to zero. In the geometric context there is a homotopy to 0. To see this note that the center points in the k-2 cells are codimension k-2 in K^{k-2}, in the sense that they have neighborhoods which are products with D^{k-2}. Form a 1-complex in K^{k-1} by adding to these points the rays to the centers of the k-1 cells. Since the attaching maps are transverse to the centers, this 1-complex is also codimension k-2, except at the centers of the k-1 cells.

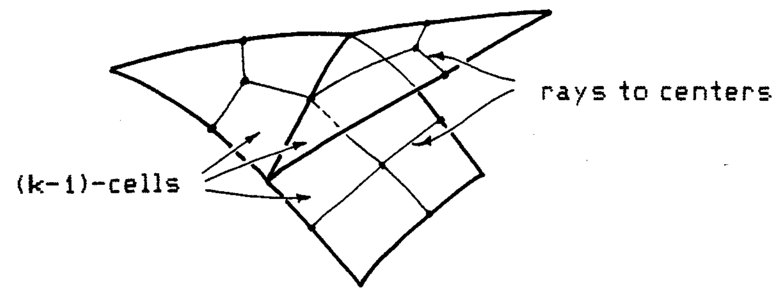

Assume the attaching maps of the k-cells are transverse to this 1-complex. The inverse images in S^{k-1} are then 1-complexes. The vertices are the inverses of the centers of the k-1 cells. There are arcs between these, and disjoint circles. The circles are not useful to us. The cone on the arcs (union of radial lines from the centers of the k-cells) define maps of 2-disks into K.

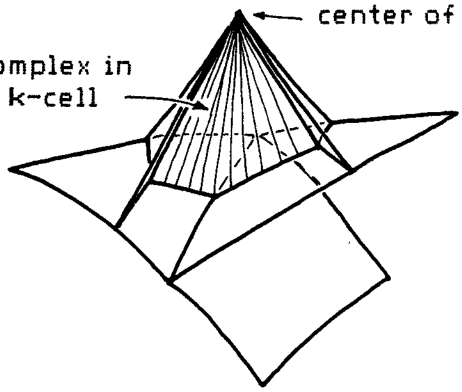

These define a homotopy of $\partial\partial$ to 0. In more detail, each of these 2-disks can be deformed to a map of a square into K. One vertex goes to the center of a k-cell, the adjacent edges to radial lines to centers of (k-1)-cells, and the remaining edges go to radial lines in (k-1)-cells to the center of a (k-2)-cell (see the illustration above). This is a homotopy between two of the paths in the composition $\partial\partial$. Consideration of orientations shows that these paths have opposite sign, so the pair of signed paths are homotopic to a single path with coefficient 0. Therefore up to homotopy they cancel. Finally it is not hard to see that each path in $\partial\partial$ occurs in exactly one of these squares, so the entire composition is homotopic to 0.

This entire collection of data, geometric modules C_k, geometric morphisms ∂, and the homotopy of $\partial\partial$ to 0, forms a geometric chain complex.

Note that if each cell in K has image of sufficiently small diameter in X then the morphisms and homotopies in the geometric complex have radius less than ϵ. Forgetting paths then gives the ϵ chain complexes constructed in Quinn [6, p.271]. Passing to unrestricted homotopy classes gives the classically defined $\mathbb{Z}[\pi_1 X]$ chain complex.

Section 2: Splitting of chain complexes

In this section we suppose that π is a group which is a generalized free product, and construct corresponding splittings of chain complexes over $R[\pi]$. The result itself is not particularly striking. Rather the proof is supposed to suggest benefits of the geometric point of view even in purely algebraic situations. For example it may be that the decomposition theorems for algebraic K-theory (cf. Waldhausen [7, 8]) could be obtained this way. I am told that early preprints of Waldhausen's work have constructions similar to ones used here.

Proposition *Suppose* π *is a homotopy pushout of morphisms* $\alpha,\beta:A \longrightarrow B$ *of groupoids, and the composition* $A \longrightarrow \pi$ *is an injection on each component of* A. *Then any finitely generated free* $R[\pi]$ *chain complex is chain equivalent to a pushout* $a,b:E\otimes_A R[\pi] \longrightarrow F\otimes_B R[\pi]$, *where* E, F *are complexes over* $\mathbb{Z}[A]$, $\mathbb{Z}[B]$, *and* a, b *are chain maps over* α, β.

Proof The homotopy pushout hypothesis on π means that there is a space (CW complex) X with a subspace Y with a neighborhood homeomorphic to $Y\times\mathbb{R}$. The fundamental group of X is π, the fundamental groupoids of Y, X−Y are A, B respectively, and α, β are induced by the inclusions $Y \longrightarrow Y\times(0,\pm\infty) \longrightarrow X-Y$. (Groupoids are disjoint unions of groups. Here they occur as the union of fundamental groups of components of disconnected spaces, see [7].)

Suppose that C_* is a finitely generated free chain complex over $R[\pi_1 X]$. Represent C_* as a geometric R-complex on X. Tha data for this are bases $S_i \longrightarrow X$ for the chain groups C_i, geometric morphisms $\partial:R[S_i] \longrightarrow R[S_{i-1}]$, and homotopies $\partial^2\sim 0$. For simplicity we will assume that all paths in the morphisms are defined on the unit interval I. The homotopies then consist of maps of squares $I^2 \longrightarrow X$; the (0,0) corner goes to an element of S_{i+1}, the adjacent sides to paths in ∂_{i+1}, and the remaining sides to paths in ∂_i.

We will say that a geometric complex is "special" if the paths $\lambda:I \longrightarrow X$ have $\lambda^{-1}(Y)$ either I, {1}, or ϕ, and the homotopies $h:I^2 \longrightarrow X$ have $h^{-1}(Y)$ either I^2, or properly contained in $\partial(I^2)$. We claim that such a complex splits in the desired way.

Suppose C_* has this special form. Define E_* to be the submodule of C_* generated by basis elements which map into Y. E_* and the restriction of the boundary morphisms in C_* define a geometric complex on Y; by hypothesis if a path in ∂ starts in Y it stays in Y. The composition ∂^2 is homotopic to 0 in X, but the homotopies are squares with entire boundary mapping to Y. Such squares are required to map to Y, so ∂^2 is nullhomotopic in Y.

Next define F_* by "doubling" E inside C: replace each basis element of C with image in Y by two elements, with images $y\times\{-1\}$ and $y\times\{+1\}$ in $Y\times\mathbb{R} \subset X$. The boundary morphism is that of F in each copy of F, and unchanged in the rest of C except for paths which terminate at a point in Y. Such paths by hypothesis intersect Y in only the final endpoint. Just before the path hits Y, it is either on the + or − side of Y in $Y\times\mathbb{R}$. Push the path off Y, to terminate at the appropriate $y\times\{\pm1\}$. The homotopies of ∂^2 to 0 also can be pushed off Y. F is therefore a geometric complex over X−Y.

There are chain maps $a,b:E_* \longrightarrow F_2$ defined by inclusions of the copies of E over $Y\times\{\pm1\}$. C_* is the quotient of F_* by the image $(a-b)E_*$, hence chain equivalent to the pushout. Passing to homotopy classes of morphisms gives complexes over $R[\pi_1 Y]$, $R[\pi_1(X-Y)]$. This gives the decomposition required for the proposition.

The proposition therefore will follow if we show that an arbitrary geometric chain complex on X is equivalent to a special one.

The *underlying 2-complex* of a geometric chain complex C_* is formed from the geometric data. The vertices are the union of the generators for the chain groups. The edges are the paths with nonzero coefficient in the boundary homomorphisms. The 2-cells are the squares in the homotopy $\partial^2 \sim 0$. Denote this underlying complex by UC. The maps of the pieces fit together to give a map UC $\longrightarrow$ X. The underlying complex is filtered by dimension in the chain complex; define U_iC to be $U(C_*, * \leq i)$. Finally note that the 1-cells are oriented, in that one end has lower filtration than the other.

We say a filtered 2-complex mapping to X is "special" if $f^{-1}(Y)$ is a subcomplex, and if the vertex of highest filtration of a cell is in $f^{-1}(Y)$ then the entire cell is also. In these terms the proof of the proposition is reduced to: show that every geometric chain complex is equivalent to one whose underlying 2-complex is special.

Assume, as an induction hypothesis, that the i-1 filtration $U_{i-1}C$ is special. By small homotopy holding U_{i-1} fixed we may assume that there is a neighborhood N of U_{i-1} in U_i such that $N-U_{i-1} \subset X-Y$. Then we may assume that U_i-U_{i-1} is transverse to Y. The inverse image will therefore be a 1-complex. Squares in U_i-U_{i-1} intersect this 1-complex in arcs with ends on the upper edges, and circles.

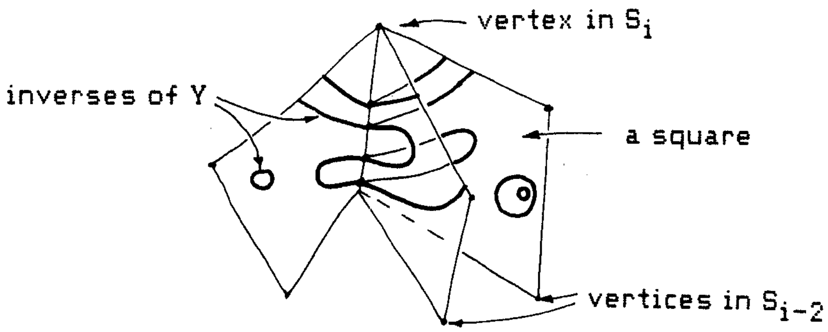

The first step in simplifying the intersection is to note that if there is an arc with both ends on one side of a square, then it encloses a disk in the square. We can push the edge across this disk (draging along any other squares which share that edge). This operation may generate new arcs and circles, but it reduces the number of intersections with the edges.

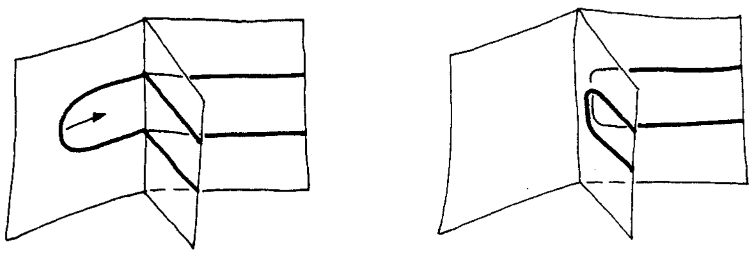

push across

By induction on the intersections with edges, we may assume there are no such arcs.

Next we eliminate the circles interior to each square. These circles map to Y, and the map on the disk the circle bounds in the square gives a nullhomotopy in X. By the hypothesis of injectivity of the fundamental group of Y, these circles are also nullhomotopic in Y. Using the nullhomotopy in Y we can redefine the map to take the disk the circle bounds to Y. This disk can then be pushed off Y.

These changes in the CW complex define changes in the chain complex. Note that the edges, and therefore the boundary morphisms, are changed by homotopy. The homotopies $\partial^2 \sim 0$ are changed by more than homotopy, but that is acceptible; only their existence is part of the data.

Now consider one edge with vertex v in S_i, and consider the point intersection nearest to v of the edge with the inverse of Y. Let $L \subset U_i$ be the component of inverse of Y containing this intersection point. We claim that the region in U_i between v and L is isomorphic with the cone v*L (see the illustration below). For this it is sufficient to show that every intersection of L with an edge is the first intersection of the edge with the inverse of Y. To see this, suppose there is one which is not the first. Choose a path (=sequence of 1-cells) in L from the intersection which is a first, to one which is not. Somewhere in the path there is a single arc so that one end is a first intersection and the other is not. This implies the existence of an arc with both ends on one edge, which contradicts the earlier improvement.

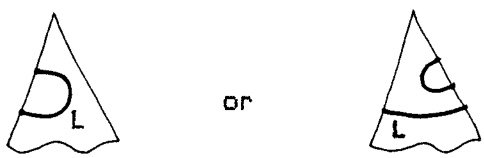

Now construct a new complex U'_i by inserting an arc between v and the cone on L. There is a map $U_i \longrightarrow U'_i$ defined by mapping the cone v*L to the arc. We modify the map $U \longrightarrow X$ so that $U_i \longrightarrow X$ factors through this map.

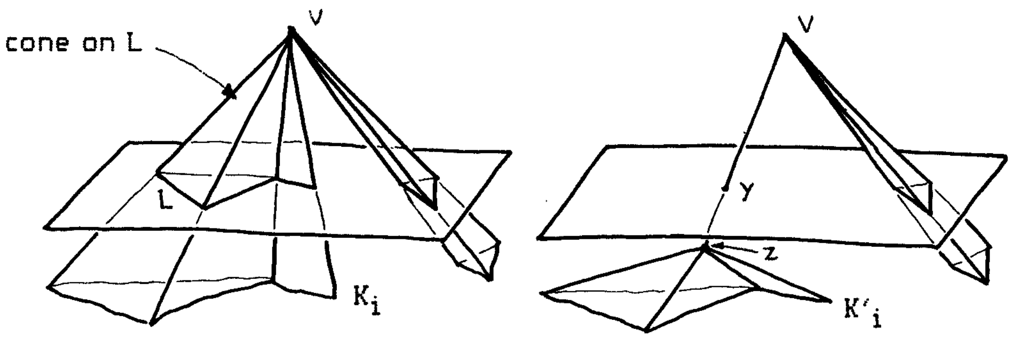

Choose a maximal tree in L, and a collapse of it to a vertex w. Let vw denote the edge between v and w. Then the collapse defines a homotopy of the cone on the tree into L∪(vw). The remaining edges in L define loops in Y which are nullhomotopic in X. The injectivity hypothesis implies that they are also nullhomotopic in Y. A nullhomotopy gives a new map of the cone on this edge, into Y union the image of vw. This factors the map through U′. Push U′ off Y as in the picture above, to leave an intersection point with the arc.

The next step is to define a new complex C′ equivalent to C, which has U'_i as filtration i in the underlying 2-complex. For this we introduce some notation in U′. Let z denote the new cone point, so the arc has been inserted between v and z. Let y denote the intersection point of the arc vz with Y, and let vy, zy denote the paths in vz from the endpoints to y. Let D denote the complex with R[y] in dimension i-1, R[z] in dimension i, and boundary 1(zy). C′ will be defined by modifying the boundary homomorphisms in C⊕D.

Write $\partial: C_i \longrightarrow C_{i-1}$ as $\partial = a+b$, where a consists of the pieces of ∂ whose paths pass through L in U_i, and b is all the rest. Note that the homotopy $\partial(a+b) \sim 0$ breaks into homotopies $\partial a \sim 0$ and $\partial b \sim 0$: since L is an entire component of the intersection with Y, any square in the homotopy with one edge on a path in a must have the other edge on a path in a as well. Note a is defined on the module R[v]. Define $a': R[z] \longrightarrow C_{i-1}$ so that a is homotopic to $a'(vz)$. We use this to construct isomorphisms of the chain groups of C⊕D, and define C′ to have boundary morphisms obtained by conjugation by these isomorphisms. Explicitly C′ is the bottom line in the diagram:

$$\begin{array}{ccccccccc}
\longrightarrow C_{i+1} & \xrightarrow{\begin{bmatrix}\partial \\ 0\end{bmatrix}} & C_i \oplus R[z] & \xrightarrow{\begin{bmatrix} a+b & 0 \\ 0 & zy\end{bmatrix}} & C_{i-1}\oplus R[y] & \xrightarrow{[\partial\ 0]} & C_{i-2} \longrightarrow \\
\downarrow 1 & & \downarrow \begin{bmatrix} 1 & 0 \\ vz & 1\end{bmatrix} & & \downarrow \begin{bmatrix} 1 & a'(yz) \\ 0 & 1\end{bmatrix} & & \downarrow 1 \\
\longrightarrow C_{i+1} & \xrightarrow[\begin{bmatrix}\partial \\ (vz)\partial\end{bmatrix}]{} & C_i \oplus R[z] & \xrightarrow[\begin{bmatrix} b & a' \\ -vy & zy\end{bmatrix}]{} & C_{i-1}\oplus R[y] & \xrightarrow[{[\partial\ 0]}]{} & C_{i-2} \longrightarrow
\end{array}$$

The vertex y is special in C′, and there are fewer components of the inverse image of Y in U_iC'. Therefore by iterating the construction we can get a complex equivalent to C with filtration i special.

This completes the induction step, and shows that we can find an equivalent complex whose entire underlying 2-complex is special. As indicated above, this implies the proposition.

Section 3: The controlled h-cobordism theorem

The classical h-cobordism theorem states that an h-cobordism with vanishing Whitehead torsion is isomorphic to a product. In the contolled version of this there is a map to a metric space X, $\epsilon>0$ is given, and we want a product structure such that the image in X of each product arc lies in the ball of radius ϵ about its beginning point. Another way to say this is that the product structure has radius less than ϵ as a homotopy. We will see that the obstruction to this lies in a controlled version of the Whitehead group.

The section begins with some generalities on control functions, necessary on noncompact control spaces. Then ϵ Whitehead groups are defined, and the theorem is proved.

3.1 Control functions Suppose X is a metric space, and $\epsilon:X \longrightarrow (0,\infty)$ is a map. If $f,g:Y \longrightarrow X$ are functions then we say g is *within* ϵ *of* f, and write $d(f,g)<\epsilon$, if $d(f(y),g(y)) < \epsilon(f(y))$ for each $y\in Y$. Notice that this usually does not imply that $d(g,f)<\epsilon$, and the triangle inequality does not hold. To deal with this we introduce some notation.

Suppose α,β are maps $X \longrightarrow (0,\infty)$. Define $\alpha\#\beta(x)$ to be $\max\{\alpha(x)+\beta(y),\ d(x,y)<\alpha(x)\}$. If α, β are constant then $\alpha\#\beta = \alpha+\beta$.

It follows easily that if $d(f,g)<\alpha$ and $d(g,h)<\beta$ then $d(f,h) < \alpha\#\beta$. Also for functions α,β,δ, $\alpha\#(\beta\#\delta) \le (\alpha\#\beta)\#\delta$ (both expressions are maxima of $\alpha(x)+\beta(y)+\delta(z)$, and more values of z are allowed in the second expression). Denote by $n*\beta$ the n-fold iteration of this operation, with parentheses arranged to give the largest value. For example $4*\beta$ means $((\beta\#\beta)\#\beta)\#\beta$, and $m*(n*\beta)) \le (mn)*\beta$. If β is constant then $n*\beta$ is the ordinary product $n\beta$.

These expressions are usually used as upper bounds. Note that if an expression with any arrangement of parentheses is an upper bound, then the largest arrangement is an upper bound as well. This is why the notation $n*\beta$ is useful.

In the geometric algebra context note that if f,g are geometric morphisms with radius less than α,β respectively, and gf is defined, then gf has radius less than $\alpha*\beta$. Compositions of homotopies behave similarly.

3.2 Whitehead groups Suppose E is a space. The Whitehead group $Wh(R[\pi_1E])$ is defined to be the set of equivalence classes of isomorphisms of free based modules over $R[\pi_1 E]$. The equivalence relation is generated by direct sums with identity isomorphisms, and by composition with triangular automorphisms. In this context an automorphism is triangular if there is an ordering of the basis of the module so that the matrix expression is zero below the diagonal, and the diagonal entries are units in R times elements of $\pi_1 E$.

Usually a triangular matrix is required to have diagonal entries all equal to 1. The group obtained with this definition is the reduced K-group $\tilde{K}_1(R[\pi_1 E])$. The Whitehead group is obtained from this by dividing by the subgroup generated by the automorphisms of $R[\pi_1 E]$ given by products (unit in R)(element of $\pi_1 E$). But this is equivalent to allowing such products on the diagonal of triangular matrices.

Using the equivalence of section 1.1 we can describe $Wh(R[\pi_1 E])$ as equivalence classes of geometric isomorphisms of finitely generated geometric R-modules on E. We obtain a version with ϵ control by adding ϵ to this description, as in 1.3.

Suppose $p:E \longrightarrow X$ is a map, and X is a metric space. A *geometric ϵ isomorphism* of geometric R-modules on Eis a morphism of radius $<\epsilon$ (measured in X) with an "inverse" also of radius $<\epsilon$, such that the compositions are ϵ homotopic to the identity morphisms. Unfortunately it is possible to have a morphism of radius $<\epsilon$ which is an isomorphism, but whose inverse has very large radius. Therefore the estimate on the radius of the inverse must be included in the definition.

Suppose $M = R[S]$ is a geometric R-module on E. A geometric morphism $A:M \longrightarrow M$ is *(upper) triangular* provided there is an ordering of the basis of M such that A has no paths from t to s unless $t \leq s$, and if $t = s$ there is exactly one path, whose coefficient is a unit in R.

We observe that a triangular morphism is an isomorphism: it can be written as $D(I+B)$ where D is diagonal with entries a unit times a loop, and B has entries strictly above the diagonal. D^{-1} is obtained by inverting the units and reversing the loops, and $(I+B)^{-1} = I+\Sigma_1^n(-B)^i$, where n is large enough so that $B^{n+1} = 0$. Note that the radius of the inverse depends on the radius of the original, and this n.

We define a *deformation* of a geometric module to be a sequence $A_n A_{n-1} \cdots A_1$ of triangular morphisms. We write it as a product, and think of it that way, but actually need to keep track of a little more information than is retained in the product. What is needed is a refined version of the radius.

The *underlying 1-complex* of a morphism is the collection of paths which occur in the morphism. We think of these as trees eminating from the basis of the module, by identifying the beginning points of all paths coming from a given basis element. The underlying 1-complex of a sequence $A_2 A_1$ again consists of a tree for each basis element: start with the tree for A_1 beginning at the element, and at the end of each branch add a copy of the tree of A_2 which begins there. Trees for a sequence with n terms are built up similarly. We now say that *a sequence* $A_n A_{n-1} \ldots A_1$ *has radius less than* ϵ if for each basis element s the tree in the underlying 1-complex starting at s lies inside the ball of radius $\epsilon(s)$ about s. As with the radius of isomorphisms, we must require that the trees for the inverse sequence $A^{-1}{}_1 A^{-1}{}_2 \ldots A^{-1}{}_n$ lie inside the ϵ balls about their starting points as well.

Notice that the paths which occur in the composition of the sequence are just paths in these trees. The radius of the composition is therefore less than or equal to the radius of the sequence. The difference is that in the composition we allow deletion of paths when the coefficients cancel, whereas no cancellations are allowed in the underlying 1-complex. Therefore the radius of the composition may be strictly smaller.

Finally if A is a geometric morphism, an ϵ deformation of A is a composition $D_0 A D_1$, where D_0, D_1 are ϵ deformations of the range and domain modules of A.

Definition *Suppose* $p:E \longrightarrow X$ *is a map to a metric space, and* $\epsilon: X \longrightarrow (0,\infty)$ *is given. Then* $Wh(X,p,\epsilon)$ *is defined to be equivalence classes of geometric isomorphisms on* E *with radius* $<\epsilon$ *in* X, *with equivalence relation generated by direct sum with identity morphisms, and homotopies and deformations of radius* $< 3*\epsilon$.

The next lemma shows this to be a convenient place to work. However, see the comments following the proof.

Lemma *Direct sum induces an abelian group structure on* $Wh(X,p,\epsilon)$. *Further if* ϵ *isomorphisms are equivalent in this sense, then there is a* $9*\epsilon$ *deformation between isomorphisms* $9*\epsilon$ *homotopic to appropriate stabilizations of the originals.*

Proof of the lemma Since direct sum clearly induces an abelian monoid structure, the point of the first statement is that there are inverses. If A is an isomorphism there is a matrix identity

$$\begin{bmatrix} I & A \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ -A^{-1} & I \end{bmatrix} \begin{bmatrix} I & A \\ 0 & I \end{bmatrix} \begin{bmatrix} A & 0 \\ 0 & A^{-1} \end{bmatrix} \begin{bmatrix} I & -I \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ I & I \end{bmatrix} \begin{bmatrix} I & -I \\ 0 & I \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix}$$

If A is a geometric morphism of radius $<\epsilon$, the left side of the equation is a $3*\epsilon$ deformation of $A \oplus A^{-1}$. (The left three, and right three, terms are triangular.) When the left side is multiplied out there are terms like $A - AA^{-1}A$, so the composition is actually $3*\epsilon$ homotopic to $I \oplus I$ rather than equal to it. This shows that A^{-1} is a $3*\epsilon$ additive inverse for A.

Next suppose that there is a sequence A_i of ϵ isomorphisms, such that there is a $3*\epsilon$ deformation from A_i to A_{i+1}, and i goes from 1 to n+1. Consider the sequence of deformations

$$A_1 \sim A_1 \oplus \Sigma_1^n (I \oplus I) \sim A_1 \oplus \Sigma_1^n (A_i^{-1} \oplus A_i) \sim A_1 \oplus \Sigma_1^n (A_i^{-1} \oplus A_{i+1}) =$$
$$[\Sigma_1^n (A_i \oplus A_i^{-1})] \oplus A_{n+1} \sim [\Sigma_1^n (I \oplus I)] \oplus A_{n+1} \sim A_{n+1}.$$

The first and last are stabilizations, the second and fifth are the $3*\epsilon$ deformations which cancel inverses, and the third is the sum of the deformations $A_i \sim A_{i+1}$. The composition of these give a $9*\epsilon$ deformation from a stabilization of A_1 to a stabilization of A_{n+1}.

Usually there will be stabilizations in the deformations $A_i \sim A_{i+1}$. These are easily incorperated in the above argument.

Remarks The point of the second statement in the lemma is that geometric control is not lost by allowing arbitrarily many deformations. If we simply compose the sequence of deformations $A_i \sim A_{i+1}$ we get a deformation of radius $3n*\epsilon$, which can be arbitrarily large. The method for getting a short deformation comes from Quinn [5, lemma 4.4]; its use in this context is due to Chapman [2, theorem 3.5].

In general the groups $Wh(X,p,\epsilon)$ are quite mysterious. If ϵ is much larger than δ then the image of $Wh(X,p,\delta)$ in $Wh(X,p,\epsilon)$ is sometimes more accessible. Chapman [2] gives criteria for this image to be trivial, in terms of vanishing of ordinary Whitehead groups of $\pi_1(p^{-1}(U))$, for open sets U in X. In a more rigid setting (p a stratified system of fibrations) but no condition on $\pi_1(p^{-1}(U))$, Quinn [5] shows the image to be a generalized homology group of X.

3.3 Controlled h-cobordisms Suppose $\delta:X \longrightarrow (0,\infty)$ is given. Then a manifold triad $(W,\partial_0 W,\partial_1 W)$ with a map $f:W \longrightarrow X$ is a *(δ,h)-cobordism* if f is a proper map, and there are deformation retractions of W to $\partial_i W$ which have radius $<\delta$ in X.

The question we consider is: when does a (δ,h)-cobordism have a product structure $W \simeq (\partial_i W)\times I$ of radius $<\epsilon$ in X? This has been considered at length in the literature; the objective here is just to see the obstructions in geometric algebraic terms.

Some local control of the fundamental group is necessary. For this fix a map $p:E \longrightarrow X$. A map $f:W \longrightarrow E$ is *relatively δ,1 connected* (over X) if for every relative 2-complex (K,L) and commutative diagram

$$\begin{array}{ccc} L & \longrightarrow & W \\ \downarrow & & \downarrow f \\ K & \longrightarrow & E \end{array}$$

there is a map $K \longrightarrow W$ which agrees with the given map on L, and whose composition with f is within ϵ of the given map into E, measured in X.

Theorem *Suppose* $f:W \longrightarrow E$, $p:E \longrightarrow X$, *and* $\delta:X \longrightarrow (0,\infty)$ *are given, so that* pf *is a (δ,h)-cobordism over* X. *Then there is a well-defined invariant* $q_1(W,\partial_0 W) \in Wh(X,p,9*\delta)$ *which vanishes if* W *has a* δ *product structure. Conversely there is a function* $k(n)$ *such that if* $n\geq 6$ *and* $q_1(W,\partial_0 W) = 0$ *then* W *has a product structure of radius* $< k(n)*\delta$.

Proof This is theorem 3.1 of Quinn [5], with some minor refinements in estimates and the use of geometric instead of metric algebra. We outline that proof. A similar statement, with a more geometric definition of the Whitehead group, is given by Chapman [2, 14.2].

Choose a handlebody structure on $(W,\partial_0 W)$ with handles whose images in X have diameter less than δ. By *diameter less than* δ we mean here that if x and y are in the set, then $d(x,y) < \delta(x)$. If W is smooth or PL such a handle structure can be defined from a fine triangulation.

A handlebody structure has a spine, which is a CW complex structure on $(W,\partial_0 W)$. For example the spine of a handlebody structure obtained from a triangulation is just the triangulation itself. We require that the CW structure also have cells of diameter less than δ. This is automatic if the handles are small enough.

We also require the CW structure to be *saturated*. This means that the attaching map $\theta:S^n \longrightarrow K^n$ for an n+1 cell has image a union of cells. In other words, if a cell intersects the image of θ, it is contained in the image. If the CW structure is a triangulation it is automatically saturated. In the topological case it is not hard to arrange saturation.

The saturation condition is used to push things rapidly into skeleta. Suppose that $K \longrightarrow X$ is a saturated complex of dimension n whose cells have diameter less than δ in X. Suppose $h:L \longrightarrow K$ is a map, L a j-complex with $j<n$. L can be pushed off the n-cells of K, to obtain a map h_{n-1} of L into the n-1 skeleton with $d(h,h_{n-1}) < \delta$ (measured in X, as always). Similarly we can push L off the n-1 cells, and repeat until we have h_j mapping L into the j-skeleton of K. Since L has been moved n-j times, in general we only know that $d(h,h_j) < (n-j)*\delta$. However if K is saturated, then a point which is moved out of the interior of a cell in one push stays in the image of the boundary of that cell during later pushes. Therefore $d(h,h_j) < \delta$.

Choose a δ deformation retraction of W to $\partial_0 W$, $h:W\times I \longrightarrow W$. Put the product CW structure on $(W\times I,(\partial_0 W)\times I)$. Use the fact that the structure on W is saturated to get a deformation retraction h' with $d(h,h') < \delta$, $d(h',h) < \delta$, and which preserves skeleta. In particular note that images of cells under h' have diameter $< 3*\delta$.

Apply the cellular chain construction of 1.4 to obtain $C_* = C_*(W,\partial_0 W)$, a geometric chain complex of radius $<\delta$. The deformation retraction h' defines a chain homotopy $s:C_* \longrightarrow C_{*+1}$ of radius $<3*\delta$, such that $s\partial+\partial s$ is $4*\delta$ homotopic to the identity map of C_*.

Consider the morphism $(s\partial s+\partial s\partial):\Sigma_{(j\ \mathrm{even})}C_j \longrightarrow \Sigma_{(j\ \mathrm{odd})}C_j$. This is a $9*\delta$ isomorphism; it has radius $<7*\delta$ and the same formula from odd j back to the even ones is a $9*\delta$ inverse. We define $q_1(W,\partial_0 W)$ to be the equivalence class $[s\partial s+\partial s\partial]$ in $Wh(X,p,9*\delta)$.

The first step in proving the theorem is to show the invariance. If s' is another $4*\delta$ chain contraction then s' is $4*\delta$ homotopic to $s+(s's\partial-\partial s's)$. Using this we can write

$(s'\partial s'+\partial s'\partial)$ as $(I+D)(s\partial s+\partial s\partial)$, where for example $D = \partial ss'\partial ss'\partial s\partial s$. D raises dimension in $\Sigma_{(j\ odd)}C_j$ so I+D is triangular. It has radius less than $3(9)*\delta$ so is an equivalence in $Wh(X,p,9*\delta)$. This shows that q_1 is independent of the deformation retraction.

Now suppose there is another handle decomposition of $(W,\partial_0 W)$ satisfying the conditions above. Then there is a 1-parameter family of handlebody structures, all of diameter less than δ, joining them. In this family the handles can change by isotopy, cancelling pairs can be introduced or be cancelled, and handle additions can occur. Isotopy changes the chain complex only by homotopy. The other changes occur at isolated points in the parameter arc. Choose points between these and apply the construction. This gives a sequence of $9*\delta$ isomorphisms, with adjacent ones related by homotopy and a single modification. Cancelling pairs change the complex by addition of an identity morphism. Handle additions change the boundary morphisms by product with a triangular morphism, so change the isomorphism in the same way. We conclude that adjacent isomorphisms in the sequence are equivalent, hence the ends are equivalent in $Wh(X,p,9*\delta)$.

This shows that q_1 is well defined. If W has a δ product structure then the product structure is a handle structure with no handles. The chain complex is therefore trivial, and the invariant equal to 0.

The converse is essentially proved in section 6 of Quinn [4], which is independent of the rest of that paper. The estimate there is $k(n) = (54+3n)^n$, but this can be improved quite a lot using the "saturation" idea above.

The idea of the proof is to first show that $(W,\partial_0 W)$ has a handlebody structure with handles only in two adjacent dimensions. In this case the boundary morphism between the geometric chain groups is an isomorphism, which when mapped to E represents $\pm q_1$. The hypothesis that $q_1 = 0$ implies, by the lemma in 3.2, that there is an $81*\delta$ deformation of the image of the boundary morphism to the empty morphism. This deformation takes place in E. Since $W \longrightarrow E$ is relatively $(\delta,1)$-connected, the paths and homotopies of the deformation lift back to give a deformation in W. The proof of [4, section 6] then shows how to use such a deformation to cancel the remaining handles.

This completes the sketch of the controlled h-cobordism theorem. We remark that recent work of M. Freedman and the author shows that this theorem also applies to 5-dimensional topological h-cobordisms, provided the local fundamental groups are "poly-(finite or cyclic)".

References

[1] T. A. Chapman, *Controlled boundary and h-cobordism theorems*, Trans. AMS 280 (1983) 73-95.

[2] T. A. Chapman, *Controlled simple homotopy theory and applications*, Lecture Notes in Mathematics 1009 (1983) Springer-Verlag.

[3] E.H. Connell and J. Hollingsworth, *Geometric groups and Whitehead torsion*, Transactions AMS 140 (1969) 161-181.

[4] F. Quinn, *Ends of Maps I*, Annals of Math 110 (1979) 275-331.

[5] F. Quinn, *Ends of Maps II*, Invent. Math 68 (1982) 353-424.

[6] F. Quinn, *Resolutions of homology manifolds, and the topological characterization of manifolds*, Invent. Math. 72 (1983) 267-284.

[7] F. Waldhausen, *Whitehead groups of generalized free products* Springer Lecture Notes 342 (1973) 155-179.

[8] F. Waldhausen, *Algebraic K-theory of generalized free products*, Annals of Math. 108 (1978) 135-256.