ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 5 (2001) 399–429
Published: 24 April 2001

The compression theorem I

Colin Rourke
Brian Sanderson

Mathematics Institute, University of Warwick
Coventry, CV4 7AL, UK

Email: cpr@maths.warwick.ac.uk and bjs@maths.warwick.ac.uk

URL: http://www.maths.warwick.ac.uk/~cpr/ and ~bjs/

Abstract

This the first of a set of three papers about the *Compression Theorem*: if M^m is embedded in $Q^q \times \mathbb{R}$ with a normal vector field and if $q - m \geq 1$, then the given vector field can be *straightened* (ie, made parallel to the given $\mathbb{R}$ direction) by an isotopy of M and normal field in $Q \times \mathbb{R}$.

The theorem can be deduced from Gromov's theorem on directed embeddings [5; 2.4.5 (C′)] and is implicit in the preceeding discussion. Here we give a direct proof that leads to an explicit description of the finishing embedding.

In the second paper in the series we give a proof in the spirit of Gromov's proof and in the third part we give applications.

AMS Classification numbers Primary: 57R25

Secondary: 57R27, 57R40, 57R42, 57R52

Keywords: Compression, embedding, isotopy, immersion, straightening, vector field

Proposed: Robion Kirby
Seconded: Yasha Eliashberg, David Gabai

Received: 25 January 2001
Revised: 2 April 2001

1 Introduction

We work throughout in the smooth (C^∞) category. Embeddings, immersions, regular homotopies etc will be assumed either to take boundary to boundary or to meet the boundary in a codimension 0 submanifold of the boundary. Thus for example if $f\colon M \to Q$ is an immersion then we assume that either $f^{-1}\partial Q = \partial M$ or $f^{-1}\partial Q$ is a codimension 0 submanifold of ∂M. In the latter case we speak of M having *relative boundary*. The tangent bundle of a manifold W is denoted $T(W)$ and the tangent space at $x \in W$ is denoted $T_x(W)$. Throughout the paper, "normal" means independent (as in the usual meaning of "normal bundle") and not necessarily perpendicular.

This paper is about the following result:

Compression Theorem *Suppose that M^m is embedded in $Q^q \times \mathbb{R}$ with a normal vector field and suppose that $q - m \geq 1$. Then the vector field can be straightened (ie, made parallel to the given $\mathbb{R}$ direction) by an isotopy of M and normal field in $Q \times \mathbb{R}$.*

Thus the theorem moves M to a position where it projects by vertical projection (ie "compresses") to an immersion in Q.

The theorem can be deduced from Gromov's theorem on directed embeddings [5; 2.4.5 (C′)] and is implicit in the discussion which precedes Gromov's theorem. Here we present a proof that is different in character from Gromov's proof.

Gromov uses the technique of "convex integration" which although simple in essence leads to very complicated embeddings. In the second paper in this series [14] we give a proof of Gromov's theorem (and deduce the compression theorem). This proof is in the spirit of Gromov's and makes clear the complicated (and uncontrolled) nature of the resulting embeddings.

By contrast the proof that we present in this paper is completely constructive: given a particular embedding and vector field the resulting compressed embedding can be described explicitly. A third, and again different, proof of the Compression Theorem has been given by Eliashberg and Mishachev [3].

The method of proof allows for a number of natural addenda to be proved and in particular we can straighten a sequence of vector fields. More precisely suppose that M is embedded in $Q \times \mathbb{R}^n$ with n independent normal vector fields, then M is isotopic to an embedding in which each vector field is parallel to the corresponding copy of $\mathbb{R}$. This result solves an old problem (posed by Bruce Williams at the 1976 Stanford Conference [2; problem 6]) though it should be noted that this solution could also have been deduced from Gromov's theorem at any time since the publication of his book.

The third paper in this series [15] concerns applications of the Compression Theorem (and its addenda): we give new and constructive proofs for immersion theory [7, 17] and for the loops–suspension theorem of James, May, Milgram and Segal [8, 10, 12, 16]. We give a new approach to classifying embeddings of manifolds in codimension one or more, which leads to theoretical solutions, and we consider the general problem of simplifying (or specifying) the singularities of a smooth projection up to C^0–small isotopy and give a theoretical solution in the codimension ≥ 1 case.

Two examples of immediate application of the compression theorem are the following:

Corollary 1.1 *Let π be a group. There is a classifying space $BC(\pi)$ such that the set of homotopy classes $[Q, BC(\pi)]$ is in natural bijection with the set of cobordism classes of framed submanifolds L of $Q \times \mathbb{R}$ of codimension 2 equipped with a homomorphism $\pi_1(Q \times \mathbb{R} - L) \to \pi$.*

Corollary 1.2 *Let Q be a connected manifold with basepoint $*$ and let M be any collection of disjoint submanifolds of $Q - \{*\}$ each of which has codimension ≥ 2 and is equipped with a normal vector field. Define the* vertical *loop space of Q denoted $\Omega^{\text{vert}}(Q)$ to comprise loops which meet given tubular neighbourhoods of manifolds in M in straight line segments parallel to the given vector field. Then the natural inclusion $\Omega^{\text{vert}}(Q) \subset \Omega(Q)$ (where $\Omega(Q)$ is the usual loop space) is a weak homotopy equivalence.*

Proofs The first corollary is a special case of the classification theorem for links in codimension 2 given in [4; theorem 4.15], the space $BC(\pi)$ being the rack space of the conjugacy rack of π. To prove the second corollary suppose given a based map $f\colon S^n \to \Omega(Q)$ then the adjoint of f can be regarded as a map $g\colon S^n \times \mathbb{R} \to Q$ which takes the ends of $S^n \times \mathbb{R}$ and $\{*\} \times \mathbb{R}$ to the basepoint. Make g transverse to M to create a number of manifolds embedded in $S^n \times \mathbb{R}$ and equipped with normal vector fields. Apply the compression theorem to each of these (the local version proved in section 4 of this paper). The result is to deform g into the adjoint of a map $S^n \to \Omega^{\text{vert}}(Q)$. This shows that $\Omega^{\text{vert}}(Q) \subset \Omega(Q)$ induces a surjection on π_n. A similar argument applied to a homotopy, using the relative compression theorem, proves injectivity. (The vertical loop space is introduced in Wiest [18]; for applications and related results see [18, 19].) □

The result for one vector field (sufficient for the above applications) has a particularly simple global proof given in the next section (section 2). In section 3 we describe some of the geometry which results from this proof and in section 4 we localise the proof and show that the isotopy can be assumed to be arbitrarily small in the C^0 sense. This small version is needed for straightening multiple vector fields and leads at once to deformation and bundle versions, which lead to the connection with Gromov's results.

Finally section A is an appendix which contains proofs of the general position and transversality results that are needed in sections 2 and 4.

Acknowledgements We are grateful to Bert Wiest for observing corollary 1.2 and to David Mond for help with smooth general position. We are also extremely grateful to Chris French who has carefully read the main proof and found some technical errors which have now been corrected. We are also grateful to David Spring and Yasha Eliashberg for comments on the connection of our results with those of Gromov.

2 The global proof

We think of $\mathbb{R}$ as vertical and the positive $\mathbb{R}$ direction as upwards. We call M *compressible* if the vector field always points vertically up. Note that a compressible embedding covers an immersion in Q.

Assume that Q is equipped with a Riemannian metric and use the product metric on $Q\times\mathbb{R}$. Call a normal field *perpendicular* if it is everywhere orthogonal to M.

A perpendicular vector field α is said to be *grounded* if it never points vertically down. More generally α is said to be *ε–grounded* if it always makes an angle of at least ε with the downward vertical, where $\varepsilon > 0$.

Compression Theorem 2.1 *Let M^m be a compact manifold embedded in $Q^q \times \mathbb{R}$ and equipped with a normal vector field. Assume $q - m \geq 1$ then M is isotopic to a compressible embedding.*

The method of proof allows a number of extensions, but note that the addenda to the local proof (in section 4) give stronger statements for the first two of these addenda:

Addenda

(i) (Relative version) *Let C be a compact set in Q. If M is already compressible in a neighbourhood of $C \times \mathbb{R}$ then the isotopy can be assumed fixed on $C \times \mathbb{R}$.*

(ii) (Parametrised version) *Given a parametrised family $M_t^m \subset Q_t^q$ of embeddings with a normal vector field, $t \in K$, where K is a compact manifold of dimension k and $q-m-k \geq 1$, then there is a parametrised family of isotopies to compressible embeddings (and there is a relative version similar to (i)).*

(iii) *If in theorem 2.1 or in addendum (ii) the fields are all perpendicular and grounded, then the dimension condition can be relaxed to $q-m \geq 0$ and there is no dimension condition on K.*

We need the following lemma.

Lemma 2.2 *Under the hypotheses of theorem 2.1 the normal field may be assumed to be perpendicular and grounded.*

The lemma follows from general position. Call the vector field α and without loss assume that α has unit length everywhere. Note that the fact that α is normal (ie, independent of the tangent plane at each point of M) does not imply that it is perpendicular; however we can isotope α without further moving M to make it perpendicular. α now defines a section of $T(Q \times \mathbb{R}) \,|\, M$ and vertically down defines another section. The condition that $q - m > 0$ implies that these two sections are not expected to meet in general position. A formal proof can be found in the appendix, see corollary A.5.

Proof of theorem 2.1 We shall prove the main result first and then the addenda. The first move is to apply the lemma which results in the normal vector field α being perpendicular and grounded. By compactness of M, α is in fact ε–grounded for some $\varepsilon > 0$.

We now define an operation on α given by rotating it towards the upward vertical: Choose a real number μ with $0 < \mu < \varepsilon$. Consider a point $p \in M$ at which $\alpha(p)$ does not point vertically up and consider the plane $P(p)$ in $T_p(Q \times \mathbb{R})$ defined by the vector $\alpha(p)$ and the vertical. Define the vector $\beta(p)$ to be the vector in the plane $P(p)$ obtained by rotating $\alpha(p)$ through an angle $\frac{\pi}{2} - \mu$ in the direction towards vertically up, unless this rotation carries $\alpha(p)$ past vertically up, when we define $\beta(p)$ to *be* vertically up. If $\alpha(p)$ is already vertically up, then we again define $\beta(p)$ to be vertically up. The rotation of α to β is called *upwards rotation*. If we wish to be precise and refer to the chosen real number μ then we say *μ–upwards rotation*.

Figure 1 shows the extreme case when α is pointing as far down as possible:

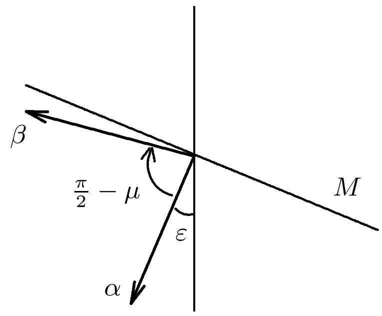

Figure 1: upwards rotation

The operation of upwards rotation, just defined, yields a *continuous* but not in general *smooth* vector field. However the operation may be altered to yield a smooth vector field by using a bump function to phase out the amount of rotation as the rotated vector approaches vertical. The properties of the resulting vector field for the proof (below) are not altered by this smoothing.

The second move is to perform μ–upwards rotation of α. The resulting vector field β has the property that it is still normal (though not now perpendicular) to M and has a positive vertical component of at least $\sin(\varepsilon - \mu)$. These facts can both be seen in figure 1; β makes an angle of at least $\varepsilon - \mu$ above the horizontal and at least μ with M. The line marked M in the figure is where the tangent plane to M at p might meet $P(p)$.

Now extend β to a global unit vector field γ on $Q \times \mathbb{R}$ by taking β on M and the vertically up field outside a tubular neighbourhood N of M in $Q \times \mathbb{R}$ and interpolating by rotating β to vertical along radial lines in the tubular neighbourhood.

Call the flow defined by γ the *global flow* on $Q \times \mathbb{R}$. Notice that this global flow is determined by α and the two choices of $\mu < \varepsilon$ and of the tubular neighbourhood N of M.

Since the vertical component of γ is positive and bounded away from zero (by $\sin(\varepsilon - \mu)$) any point will flow upwards in the global flow as far as we like in finite time.

Now let M flow in the global flow. In finite time we reach a region where γ is vertically up. Since $\gamma|M = \beta$ is normal to M at the start of the flow, $\gamma|M$ remains normal to M throughout the flow and we have isotoped M together with its normal vector field to a compressible embedding. $\square$

Proofs of the addenda

To prove addendum (i) we modify the global flow to be stationary on $C \times \mathbb{R}$ as follows.

Suppose that $x \in Q \times \mathbb{R}$ and $s \in \mathbb{R}$ then let $x + s$ denote the point obtained by moving x vertically by s. Let the global flow defined by γ be given by $x \mapsto f_t(x)$ then the *modified global flow* is given by $x \mapsto f_t(x) - t$. In words, this flow is obtained by flowing along γ for time t and then flowing back down the unit downward flow again for time t.

We apply this flow to M and the vector field γ. To avoid confusion use the notation γ_t, M_t for the effect of this modified global flow on γ, M at time t respectively. Note that $\gamma_t(x) = \gamma(x + t)$, where we have identified the tangent spaces at x and $x + t$ in the canonical way, and that $\gamma_t | M_t$ is normal to M_t for all t.

Now the modified flow is stationary whenever and wherever γ_t is vertically upwards and in particular it is stationary on $C \times \mathbb{R}$ for all t. Moreover it has similar properties to the unmodified global flow in that, after finite time, any compact set reaches a region where the vector field γ_t is vertically up.

For (ii) the dimension condition implies that the vector fields can all be assumed to be both perpendicular and grounded. This is a general position argument and the details are in the appendix, see corollary A.5. By compactness there is a global $\varepsilon > 0$ such that each vector field is ε–grounded. Now apply the main proof for each $t \in K$ and observe that the resulting flows vary smoothly with $t \in K$.

For (iii) observe that the dimension conditions were only used to prove that the vector fields were all grounded and only compactness was used thereafter. □

Remark 2.3

The modified global flow defined in the proof of addendum (i) can be regarded as given by a time-dependent vector field as follows. Let u denote the unit vertical vector field and γ the vector field which defines the global flow. Then the modified global flow is given by the vector field γ^* where

$$\gamma^*(x,t) = \gamma(x+t) - u.$$

Note also that the vector field carried along by γ^* (denoted γ_t above) is $\gamma^* + u$. Here we have again identified the tangent spaces at x and $x+t$ in the canonical way.

3 Pictures

The global proof given in the last section hides a wealth of geometry. In this section we reveal some of this geometry. We start by drawing a sequence of pictures, where M has codimension 2 in $Q \times \mathbb{R}$, which are the end result of the isotopy given by the proof in particular situations. These pictures contain all the critical information for constructing an isotopy of a general manifold of codimension 2 with normal vector field to a compressible embedding. We shall explain how this works in the local setting of the next section.

After the sequence of codimension 2 pictures, we describe some higher codimension situations and then we describe how the compression desingularises a map in a particular case (the removal of a Whitney umbrella). We finish with an explicit compression of an (immersed) projective plane in $\mathbb{R}^4$ which uses many of the earlier pictures. The image of the projection on $\mathbb{R}^3$ changes from a sphere with cross-cap to Boy's surface.

1 in 3

Consider the vector field on an angled line in $\mathbb{R}^3$ which rotates once around the line as illustrated in figure 2.

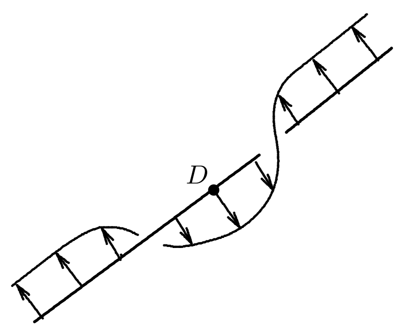

Figure 2: perpendicular field

Upwards rotation replaces this vector field by one which is vertically up outside an interval and rotates just under the line as illustrated in figure 3.

Seen from on top this vector field has the form illustrated in figure 4.

Now apply the global flow. The interval where the vector field is not vertically up flows upwards more slowly and at the same time it flows under and to both sides on the original line. The result is the twist illustrated in figure 5.

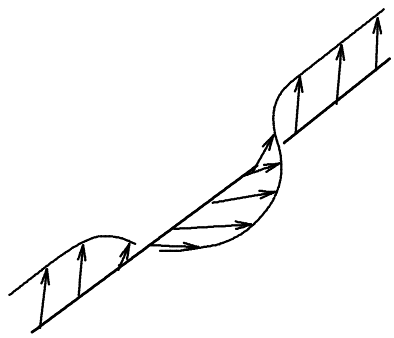

Figure 3: upwards rotated field

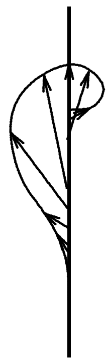

Figure 4: the upwards rotated field seen from on top

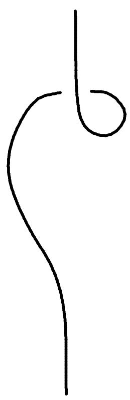

Figure 5: the effect of the global flow

2 in 4

Now observe that the twist constructed above has two possible forms depending on the slope of the original line. So now consider the surface in 4–space (with normal vector field) which is described by the moving line as it changes slope in 3–space from one side of horizontal to the other. The vector field so described is not grounded since in the middle of the movement the line is horizontal and one vector points vertically down. But a small general position shift moves this normal vector one side (ie, into the past or the future) and makes the field grounded. We can then draw the end result of the isotopy provided by the compression theorem as the sequence of pictures in figure 6 which describe an embedded 2–space in 4–space.

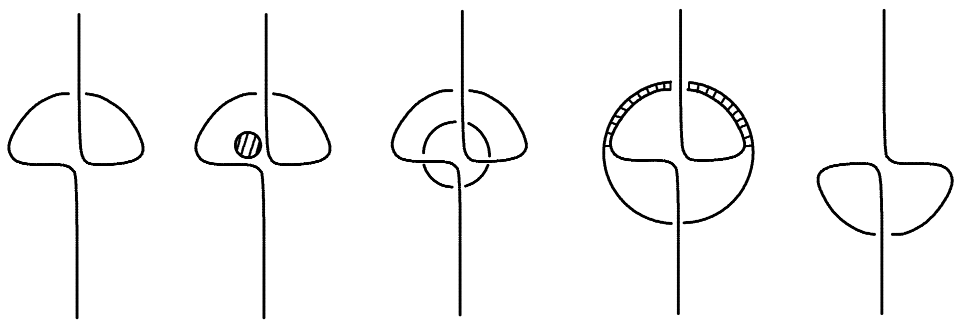

Figure 6: embedding of 2 in 4

Notice that the pictures are not accurate but diagrammatic; the end result of the actual flow produces a surface which, described as a moving picture starts like figure 5 and ends with a rotation of figure 5. However the combinatorial structure of the pictures is the same as that produced by the flow. Read the figure as a moving picture from left to right. The picture is static before time 0 (figure 6.0); at time 1 a small disc appears (a 0–handle) and the boundary circle of this disc grows until it surrounds the main twist (figure 6.3). At this point a 1–handle bridges across from the circle to the twist. The 1–handle has been drawn very wide, because then the effect of the bridge can be clearly seen as the replacement of the upwards twist by the downwards twist in figure 6.4. Note that the 1–handle in figure 6.3 cancels with the 0–handle in figure 6.1 so the topology of the surface is unaltered and the whole sequence describes an embedded 2–plane in 4–space. Indeed we can see the small disc, whose boundary grows from times 1 to 3, as a little finger pushed into and under the surface pointing into the past.

Notice that there is a triple point in the projected immersion between 6.2 and 6.3 as the circle grows past the double point.

3 in 5

The embedding of 2 in 4 just constructed again has an asymmetry (corresponding to the choice of past or future for the general position shift). The pictures drawn were for the choice of shift to the past. The choice of shift to the future produces a similar picture with the finger pointing to the right. We can move from one set of pictures to the other by a similar construction illustrated in figure 7, which should be thought of as a moving sequence of moving pictures of 1 in 3 and hence describes an embedding of 3 in 5.

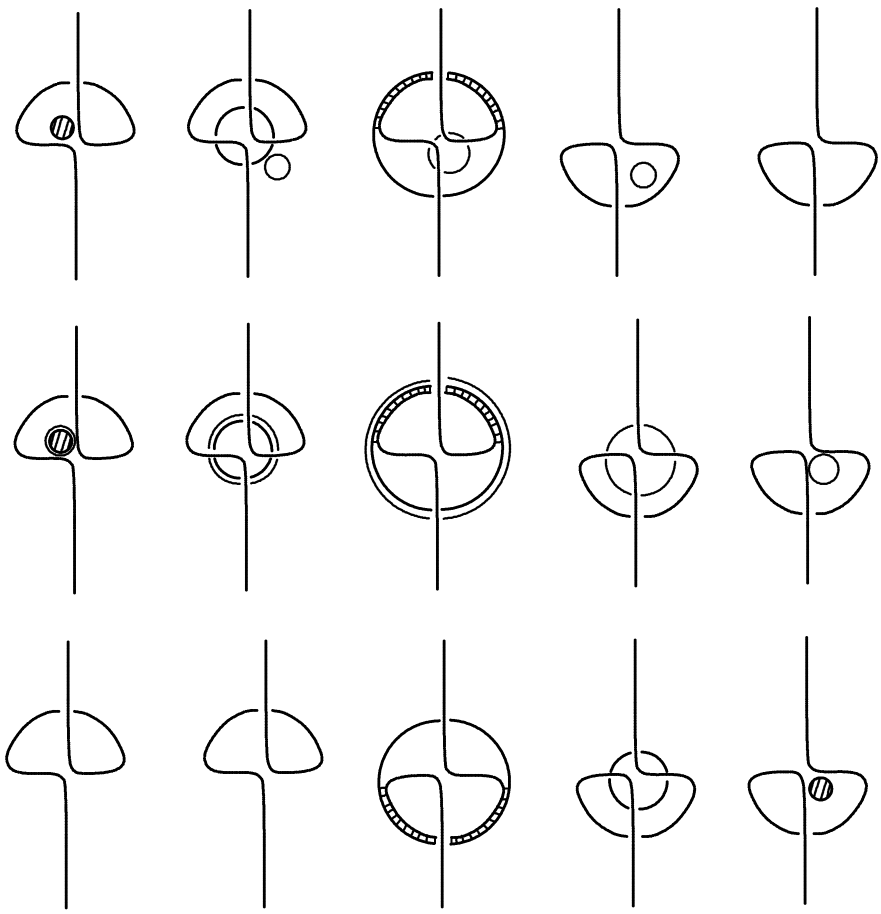

Figure 7: embedding of 3 in 5

This figure should be understood as follows: Time passes down the page. At times before 0 the picture is static and is the same as figure 6. Just before time 1 (the top row of the figure) a small 2–sphere has appeared (a 0–handle). At time 1 this sphere has grown a little and now appears as the three little circles in the middle three slices (think of the sphere as a circle which grows from a point and then shrinks back down). At time 2 (the middle row of the figure) the sphere has grown to the point where it encloses the finger. At this point a 1–handle bridges between the finger and the sphere and this has the effect of flipping the finger over from a left finger to a right finger as shown in the bottom row (time 3).

It should now be clear how to continue this sequence of constructions to construct embeddings of n–space in $(n+2)$–space for all n.

Higher codimensions

We shall now describe the end result of the straightening process for the simplest situation in higher codimensions, namely an inclined plane with the analogue of the twist field of figure 5. More precisely we start with an embedding of $\mathbb{R}^c$ in $\mathbb{R}^{2c+1}$ with a perpendicular field which points up outside a disc, down at the centre of the disc and at other points of the disc has direction given by identifying the disc (rel boundary) with the normal sphere to $\mathbb{R}^c$ in $\mathbb{R}^{2c+1}$. We shall describe the result of applying the compression theorem. The move which achieves this result is a similar twist to the codimension 2 twist of figures 2 to 5 above. We shall describe an immersion of $\mathbb{R}^c$ in $\mathbb{R}^{2c}$ with a single double point covered by an embedding in $\mathbb{R}^{2c+1}$.

The case $c = 2$ can be described as follows:

Take a straight line in $\mathbb{R}^3$. Put a twist in it. Pass the line through itself to get the opposite twist. Pull straight again. This moving picture of immersions of $\mathbb{R}^1$ in $\mathbb{R}^3$ defines an immersion of $\mathbb{R}^2$ in $\mathbb{R}^4$. Lying above it is a moving picture of $\mathbb{R}^1$ in $\mathbb{R}^4$, where at the critical stage we have the $\mathbb{R}^1$ embedded in a 3-dimensional subspace (and have a twist like figure 5 there).

We construct the general picture inductively on c. Suppose it is constructed for c. So we have an embedding of $\mathbb{R}^c$ in $\mathbb{R}^{2c+1}$ lying over an immersion in $\mathbb{R}^{2c}$ with a single double point. Now consider this immersion as being in $\mathbb{R}^{2c+1}$ then we can undo the immersion in two different ways: lift the double point off upwards and then pull flat and ditto downwards. Performing the reverse of one of these undoings followed by the other describes a moving picture of $\mathbb{R}^c$ in $\mathbb{R}^{2c+1}$, ie, an immersion of $\mathbb{R}^{c+1}$ in $\mathbb{R}^{2c+2}$ with a single double point. This is the projection of the next standard picture. To get the lift, just use the extra dimension to lift off the double point.

The move which produces the immersions just described is very similar to the move described by Koschorke and Sanderson in [9; Theorem 5.2], see in particular the picture on [9; page 216].

Removal of a Whitney umbrella

We now return to the initial example described in figures 2 to 5, but from a different point of view. The compression isotopy is a sequence of embeddings of a line in $\mathbb{R}^3$ with normal vector fields. The whole sequence defines an embedding of a plane in $\mathbb{R}^4$ with a normal vector field. The projection on $\mathbb{R}^3$ of this plane has a singularity—in fact a Whitney umbrella. To see this, we have reproduced this sequence in figure 8 below. We have changed the isotopy to an equivalent one which shows the singularity clearly.

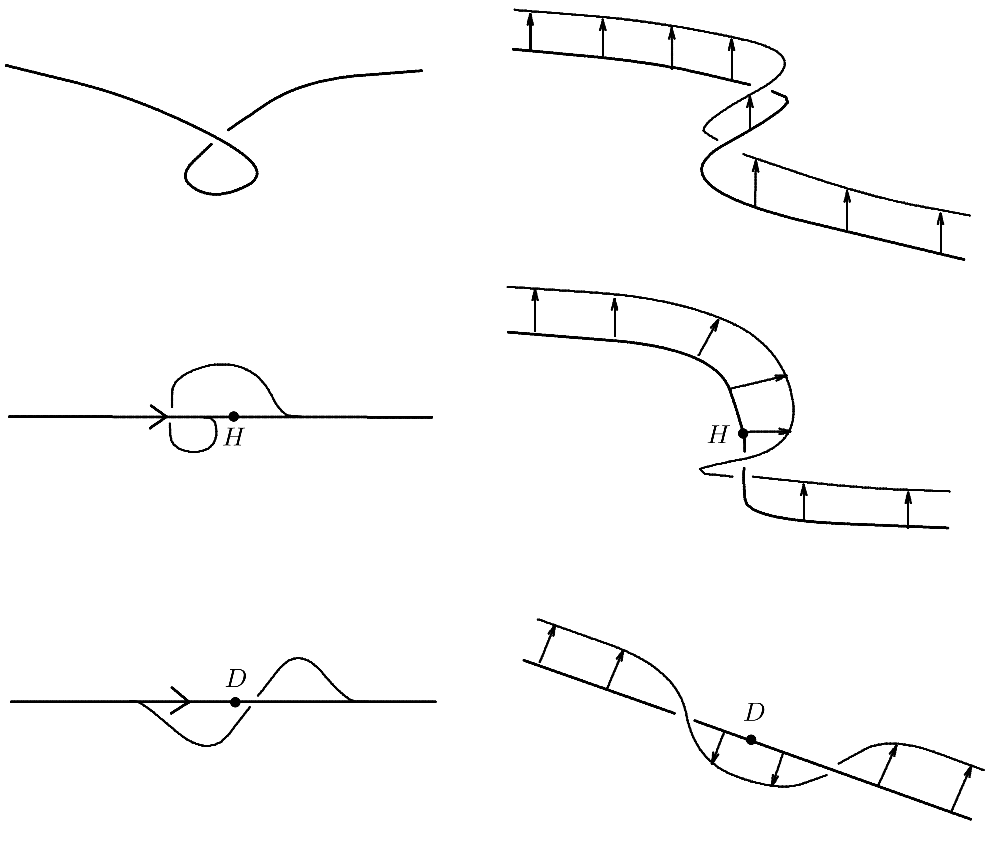

Figure 8: Surface above a Whitney umbrella

The left hand sequence in figure 8 is the view from the top, with the curve drawn out by the tip of the vector shown (in most places the vector is pointing up so this curve coincides with the submanifold), and the right hand sequence is the view from the side. The arrows in the bottom two pictures of the left sequence indicate slope. Arrows point *downhill.* Reading up the page gives an equivalent compression isotopy to that pictured in figures 2 to 5. The whole sequence defines an embedding of $\mathbb{R}^2$ in $\mathbb{R}^4$ with normal vector field which projects to a map of $\mathbb{R}^2$ to $\mathbb{R}^3$. The image has a line of double points at the top terminating in a singular point (the image of H) which can be seen to be a Whitney umbrella. The surface is flat at the bottom and has a 'ripple' at the top (ie, a region of the form (curve shaped like the letter α)$\times I$). The ripple shrinks to a point at the umbrella.

If we now apply the compression theorem to this embedding with vector field, then the image changes into a sheet with a continuous ripple (a twist is created in each slice from the H–slice down). Thus the Whitney umbrella is desingularised by having a new ripple spliced into it.

Cross-cap to Boy's surface

We finish by describing a complete compression. In this example we apply the compression theorem to an immersion (rather than an embedding). We shall see in the next section that the compression theorem can be applied to any immersion with a normal vector field by working locally. Here we anticipate this. Consider the well-known non-immersion of P^2 in $\mathbb{R}^3$ with a line of double points and two singularities (Whitney umbrellas) — a 2–sphere with cross-cap — illustrated in figure 9.

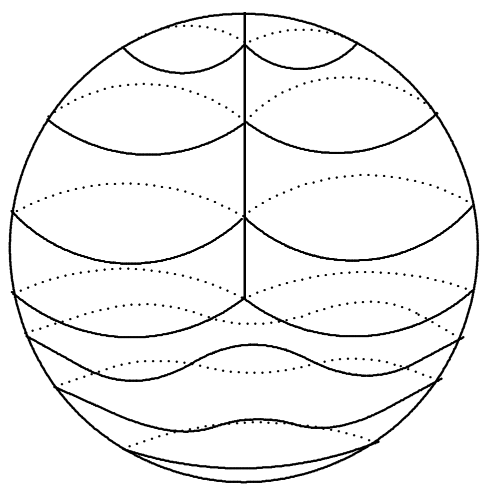

Figure 9: non-immersion of P^2 in $\mathbb{R}^3$

This is covered by an embedding in $\mathbb{R}^4$ which we can think of as a sequence of 3–dimensional slices, which start with a small circle in the shape of a 'figure 8' (corresponding to a level just below the top singular point), move to a genuine 'figure 8' at the bottom singular point and then resolve into a circle. Now it is an old result of Whitney that no embedding in $\mathbb{R}^4$ admits a normal vector field, so we modify this embedding to an immersion with a single double point in order to construct a normal vector field. The result, in a more general projection, together with the vector field is illustrated in figure 10 below.

In figure 10 we have sliced the sphere-with-cross-cap by planes which are roughly horizontal but tilted a little into general position and then lifted to an immersion in $\mathbb{R}^4$. The vector information is given using the conventions defined in the left-hand pictures in figure 8. Ignoring the vectors for the time being, the sequence starts with a 0–handle producing a small circle which passes through the top Whitney umbrella at H_1 (think of a plane tilted to the left slicing the sphere-with-cross-cap). The double point occurs at Q and then a 1–handle (a bridge) at B and the bottom Whitney umbrella at H_2 (think of the slicing plane titled backwards here). Then the resulting circle shrinks to a 2–handle at the bottom.

We now explain the vector field. At the top it is up (towards the eye) and then two opposite twist fields appear (think of a twist on a 'U' shaped curve). The right-hand twist field passes through H_1 to form a twist in the curve. The local sequence here is the same as figure 8 read upwards. The left hand twist field continues downwards, but notice that at X the slope (indicated by the arrows) changes. The field near X is the same as that used to create the figure 6 sequence described earlier. Another figure 8 sequence (reflected this time) changes the twist resulting from the bridge into an opposite twist field and the two twist fields cancel. We now draw the result of the compression.

Figure 11 comprises two views of the immersion obtained by projecting the resulting compressible immersion in $\mathbb{R}^4$ to an immersion in $\mathbb{R}^3$. The twist fields have all been replaced by ripples (as explained in connection with figure 8 above). This is indicated in the left view by a heavy black line near the top and the heavy dashed line near the middle. The heavy line represents a ripple on the *outside* of the top sheet and the dashed line represents a ripple on the *inside* of the top sheet. The square containing the point X should be imagined to be enlarged and contain a copy of the immersion in $\mathbb{R}^3$ obtained by projecting figure 6. Notice that figure 6 can be seen to start with a twist to the right of the line and ends with a twist to the left of the line. The projected immersion therefore starts with a ripple on one side of the sheet and ends with a ripple on the other side.

To the right in figure 11 is a sequence of cross-sections of the immersions, roughly on the same levels as the left-hand view. These cross-sections make

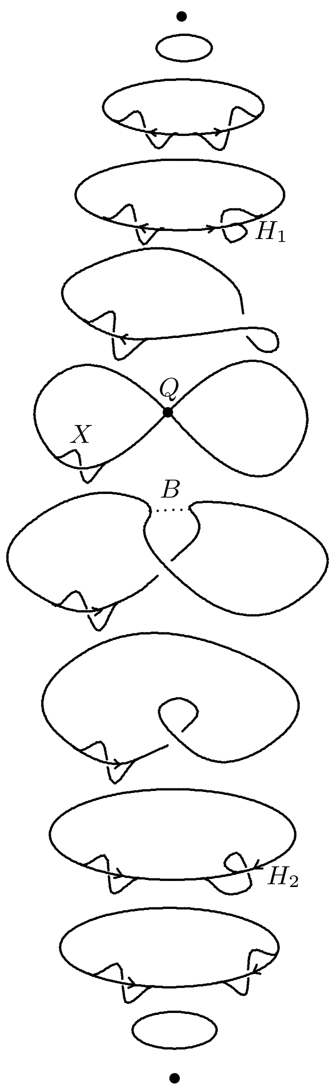

Figure 10: covering immersion in $\mathbb{R}^4$

clear what is happening near the singular points. Near the top the ripple turns around to join the 'big ripple' which is the right half of the cross-cap. Near the middle it joins into the back sheet of the Whitney umbrella. The small dotted line in the fourth section from the bottom at the right is a bridge (the 1–handle). The square marked X should again be imagined replaced by the sequence in figure 6 and notice that this sequence contains a triple point. The sequence labelled Y (to be inserted in the square marked Y) can be very simply

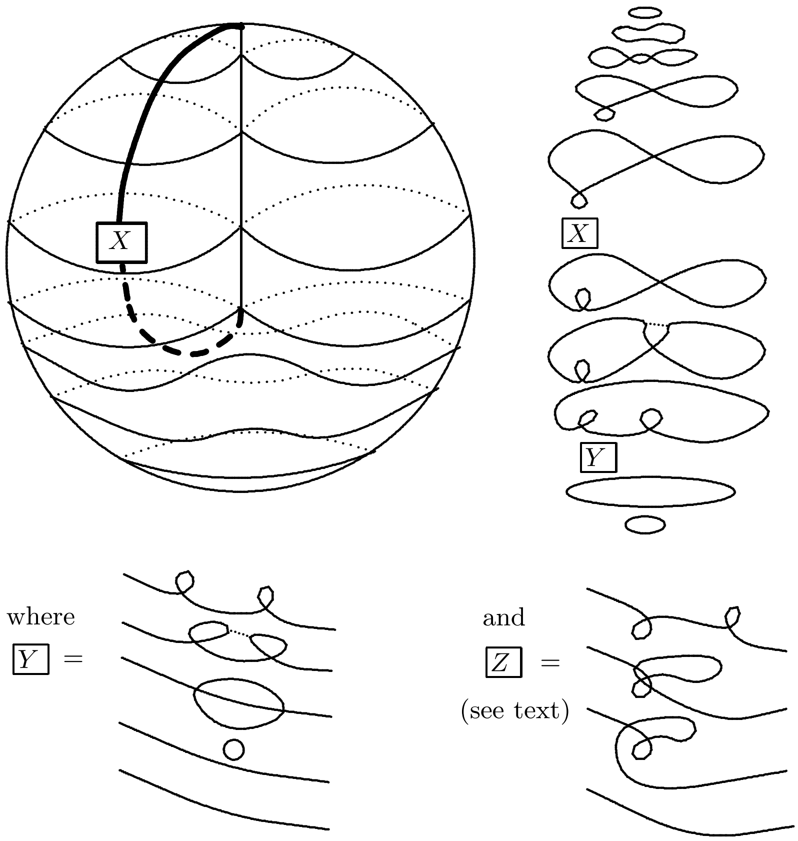

Figure 11: immersion of P^2 in $\mathbb{R}^3$

visualised as an immersion: it is just an immersion in which a ripple makes a 'U' shape.

The final immersion is isotopic to Boy's surface. This is best seen by sliding the X sequence down to meet the Y sequence. The two sequences can then be combined in a single sequence (with no critical levels) detailed as Z in the figure. The immersion now has just three critical levels and can be seen to coincide with Philips' projection of Boy's surface [13].

4 The local proof

We now prove that the isotopy constructed in the Compression Theorem can be assumed to be arbitrarily small in the C^0 sense. This implies that it can be assumed to take place within an arbitrary neighbourhood of the given embedding. We can then apply the result to an immersion by working in an induced regular neighbourhood and this opens the way to an inductive proof of the 'multi-compression theorem' which states that a number of vector fields can be straightened simultaneously.

What we shall do is to identify a certain submanifold of M which we call the *downset*. The closure of the downset is a manifold with boundary. We call the boundary the *horizontal set*. We shall see that nearly all the straightening process can be made to take place in an arbitrary neighbourhood of the downset and, by choosing this neighbourhood to be sufficiently small, the straightening becomes arbitrarily small. It will help in understanding the proof to refer to the pictures in the last section. In figure 2 the downset is one point, namely D and note that the initial vector field twists once around M in an interval centred on D. The straightening becomes smaller (and tighter) if the initial twist is made to take place in a smaller interval. In the surface which determines figure 6 the downset is a line and in figure 7 it is a plane. In each case it can be seen that the straightening takes place in a neighbourhood of the downset and can be made small by choosing this neighborhood sufficiently small.

In figure 8 the downset is a half-open interval (comprising points like D in the middle of twist fields) with boundary at H, which is the horizontal set in this example.

In the final example (figures 9 to 11) the downset again comprises a point in the middle of all the twist fields. It thus forms an interval in the neighbourhood of which the ripple in figure 11 is constructed. It has as boundary the two singular points H_1 and H_2, which form the horizontal set in this example.

We start by defining the horizontal set. This is independent of the normal vector field and defined for embeddings in $Q \times \mathbb{R}^r$ for any $r \geq 1$.

The horizontal set

Suppose given a submanifold M^m in $Q^q \times \mathbb{R}^r$ where $m \leq q$. Think of Q as *horizontal* and $\mathbb{R}^r$ as *vertical*. Define the *horizontal set* of M, denoted $H(M)$, by $H(M) = \{(x,y) \in M \mid T_{(x,y)}(\{x\} \times \mathbb{R}^r) \subset T_{(x,y)}(M)\}$. Now suppose Q has a metric and give $Q \times \mathbb{R}^r$ the product metric, then $H(M)$ is the set of points in M which have *horizontal* normal fibres, ie, (after the obvious identification) $H(M) = \{(x,y) \in M \mid \nu_{(x,y)} \subset T_x(Q)\}$ where $\nu_{(x,y)}$ denotes the fibre of the normal bundle at (x,y) defined by the metric.

Proposition 4.1 *After a small isotopy of M in $Q \times \mathbb{R}^r$ the horizontal set $H(M)$ can be assumed to be a submanifold of M.*

The proof (given in the appendix, see proposition A.2) shows more. $H(M)$ is given locally by transversality to the appropriate Grassmannian.

The downset

After application of the proposition with $r = 1$ the corresponding horizontal set H is a submanifold of M of codimension $c = q - m + 1$. Let ψ be the unit perpendicular field on $M - H$ defined by choosing the $\mathbb{R}$ coordinate to be maximal, so ψ points 'upmost' in its normal fibre. Suppose M is equipped with a unit perpendicular field α. Call $D = \{x \in M \mid \alpha(x) = -\psi(x)\}$ the *downset* of M; so D is where α points *downmost* in the normal fibre.

Proposition 4.2 *After a small isotopy of α we can assume that α is transverse to $-\psi$, which implies that D is a manifold, and further we can assume that the closure of D is a manifold with boundary H.*

The proof is given in the appendix, see proposition A.3.

Localisation

The following considerations explain why the downset is the key to the straightening process. Suppose that there is no downset (and no horizontal set). Then our perpendicular vector field α can be canonically isotoped to point to the upmost position in each normal fibre (ie, to the section ψ) by isotoping along great circles. Once in this upmost position, upwards rotation straightens the field.

More generally, let $\overline{D}$ denote the closure of the downset after proposition 4.2 (so that $\overline{D}$ is a manifold with boundary H) and let W be a tubular neighbourhood of $\overline{D}$ in M. Then α can be canonically isotoped to ψ on $\overline{(M - W)}$ and this isotopy extended to W via a collar of ∂W in W. We call this isotopy *localisation* because it localises the straightening problem in W and we call the resulting field the *localised field*.

Remark It is possible to give an explicit description of the localised field in terms of the transversality map which defines D. We sketch this description (which will not be used). The field is upmost outside a tubular neighbourhood V of D. A fibre of V comes by transversality from a neighbourhood of the "south pole" in the normal sphere, and by composing with a standard stretch of this neighbourhood over the sphere, can be identified with the normal sphere. The vector field at a point on the fibre points in the direction thus determined.

Now let φ be the gradient field on M determined by the projection $M \subset Q \times \mathbb{R} \to \mathbb{R}$. Thus for $p \in M$ the tangent vector φ_p, if non-zero, points *upmost* in the tangent space $T_p(M)$. Note that our normal field is grounded if and only if the zeros of φ are not on D.

The proof of the following lemma is again to be found in the appendix, see corollary A.6.

Lemma 4.3 *Suppose that α is perpendicular and grounded. Let U be a tubular neighbourhood of H. By a small isotopy of α we may assume that $\overline{D - U}$ is in general position with respect to φ in the following sense: Given $\delta > 0$ there is a neighbourhood V of $\overline{D - U}$ in M such that each component of intersection of an integral curve of φ with V has length $< \delta$.*

We are now ready to prove the main result of this section.

Local compression theorem 4.4 *Given $\varepsilon > 0$ and the hypotheses of the compression theorem, there is an isotopy of M to a compressible embedding which moves each point a distance $< \varepsilon$.*

Proof By lemma 2.2 we can assume that the given vector field α is perpendicular and grounded. We shall straighten α in two moves. The first move is a C^0–small move which takes place near the downset D, which we shall call the *local move* and which contains the meat of the straightening process. The second move is a C^∞–small global move. Both moves use the modified global flow defined in the proof of addendum (i) to 2.1, though in the local move, its effect is restricted to a neighbourhood of the downset. There is one important observation about the flow near the horizontal set which we need to make at the outset.

Observation Near the horizontal set the upward rotated vector field β is nearly vertical (since α is nearly horizontal there) and the same is true of the globalised field γ (see proof of 2.1). Hence the (time-dependent) generating field γ^* for the modified global flow (see remark 2.3) is initially small near H. More precisely, suppose that we restrict γ to a neighbourhood W of H in $Q \times \mathbb{R}$ and phase it out to be vertical outside a slightly larger neighbourhood W', then the isotopy defined by the corresponding modified global flow is C^∞–small. Indeed by choosing W, W' and μ (to define the upwards rotation) sufficiently small, we can make this isotopy arbitrarily small.

The local move will eliminate the downset except in a small neighbourhood of the horizontal set. Moreover the upwards rotated field β changes, if at all, to be more upright. Thus the global move, which is defined by the vector field left after the local move, has precisely the form of the flow described in the observation and will be arbitrarily C^∞–small.

For the local move we proceed as follows. Apply 4.1 and 4.2 so that D is a manifold with boundary H. Now choose a small neighbouhood U of H in M containing a smaller neighbourhood U' say. Apply 4.3 to place $\overline{D-U'}$ in general position with respect to the gradient flow φ and choose $\delta > 0$ small enough that U' has distance $> \delta$ from $\overline{M-U}$ and let V be the neighbourhood of $\overline{D-U'}$ corresponding to δ given by 4.3. We can assume that δ is also small compared to the scale of α, ie, compared with the distance over which the direction of α changes. Localise α using $U \cup V$.

We now make the constructions made in the proof of 2.1 and addendum (i). Choose a suitably small $\mu > 0$ and apply μ–upwards rotation to give the vector field β. Choose a ν–tubular neighbourhood N of M in $Q\times\mathbb{R}$ and globalise β to γ and let γ^* be the (time-dependent) vector field which gives the corresponding modified global flow. We shall make a final modification to this vector field. We need to observe that vertical projection near $\overline{D-U'}$ is an immersion in Q so we can choose a real number ω such that any two points, $x \in \overline{D-U'}$ and $y \in M$, which project to the same point of Q are a distance at least 3ω apart. (Note that ω depends on U'.)

The final modification to the flow is to multiply the time-dependent vector field $\gamma^*(x,t)$, which generates the flow, by $\rho(t)$ where ρ is a bump function with value 1 for $t \le \omega$ and 0 for $t \ge 2\omega$. (The purpose of this final modification is to terminate the effect of the flow once the straightening near $D - U'$ has occurred.) We shall call this final modification *phasing out* the flow. This modified flow determines the local move.

We now prove that, for any sufficiently small choice of U' (which then defines ω), provided the other parameters and neighbourhoods used to define the flow (namely μ, δ, V and ν) are chosen appropriately then the local move will move each point by less than $\frac{\varepsilon}{2}$. We have already observed that the flow is initially small inside $N|U$ (indeed smoothly small) and moreover the flow is initially stationary outside $N\,|\,V\cup U$ and as we shall see it is always C^∞–small in these places. We need to examine the flow in $N|V$.

In what follows we shall be near $M - H$ where vertical projection is an immersion, so we can use local coordinates of three types: coordinates in M, horizontal coordinates perpendicular to M, which we will call *sideways* coordinates and one vertical coordinate. Without loss we can assume that the fibres of N over V are compatible with these local coordinates. This implies that the foliation of N given by restricting N to flow lines of φ is locally invariant under sideways and vertical translations. We call this foliation Υ. We can also assume without loss that the rotation to vertical which defines the global field γ in terms of β is described in local coordinates of this type.

We now consider the effect of the local flow near D and for small time. The key observation is the following. On D the vector field α lies in the same vertical plane as φ. This is just because α is perpendicular to M. Since β is obtained from α by rotation in this plane, the same is true of β. Now off D, α has a component perpendicular to M in the vertical plane of φ plus a sideways component and this is still true after upwards rotation. By choice of N these facts imply that β is tangent to leaves of Υ and by definition of the global field, the same is true of γ and hence initially of γ^*. It follows that flowlines of the modified global flow lie in leaves of Υ or vertical translates.

Now consider a vector v of γ and suppose that it has a vertically upwards component of less than $\frac{\sqrt{2}}{2}$ (ie, it makes an angle $\frac{\pi}{4}$ or less above the horizontal). Then, since it is obtained by upwards rotation of $\pi/2-\mu$ or more from the corresponding vector of α, and this vector has a sideways component and a component under φ at most $\pi/2-\mu$ below horizontal, the combined sideways component and component under φ of v must have size at least $\frac{\sqrt{2}}{2}$. Rotating back upwards and subtracting the unit vertical vector to get the corresponding vector of γ^* we can see the following:

Note A vector of γ^* *either* has a vertically downwards component of at most $1-\frac{\sqrt{2}}{2}$ *or* has a combined sideways component and component under φ of at least $\frac{\sqrt{2}}{2}$.

We are now ready to estimate the distance that a point moves under the local flow. It will help to think of the flow generated by γ^* on a leaf of Υ in the following terms. Let η be a flowline of φ in V (therefore of length $<\delta$). Think of the leaf $L(\eta)$ determined by η and its vertical translates as a river with vertical sides. The flow is stationary on the banks (and at many times and places in the middle as well). The flow is generated by a disturbance (γ^* at time zero on $L(\eta)$) which moves with unit speed downwards. The total height h of this disturbance is controlled by the size of the leaf (determined by δ and N) and ν, the size of N. Furthermore the horizontal size of $L(\eta)$ is also controlled by δ and ν. Finally notice that the choice of δ small compared to the scale of α implies that the direction of any sideways movement is roughly constant on $L(\eta)$. It now follows from the note made above that each point of η either reaches the banks of the river or a point where the disturbance has passed (which then remains stationary) after a time at most $\sqrt{2}(h+\delta+\nu)$, which we can make as small as we please by choice of δ and N. Moreover since the speed a point moves is $<\sqrt{2}$ we may assume that each point is moved at most $\frac{\varepsilon}{2}$.

We can also understand the effect of phasing out in terms of the same river picture. Suppose η is outside U' then the river meets M again at a distance approximately 3ω from η; the approximation is because meetings with M are not necessarily horizontal—but by choosing the river to be narrow enough (ie δ and ν small enough) this approximation is as good as we please. Thus phasing out has the effect of killing the disturbance after it has done its work on η and before it can affect any other point of M. Finally notice that by choice of δ any point of $\overline{V - U}$ lies on such a flowline outside U'.

We have proved that, by choosing the parameters suitably, the local move straightens the vector field on $\overline{V - U}$ in a move which moves points at most $\frac{\varepsilon}{2}$ and moreover the field outside U does not move any points outside $\overline{V - U}$. It follows from the observation made near the outset that $M - \overline{V - U}$ moves only a C^∞–small amount.

This completes the local move. We observe that the field would at this point be straight if we had not phased out the flow. However the only place where the field might not be straight is now near H, where the phasing out might have stopped the global straightening before it was finished. Moreover we can see that the local move has not adversely affected the upwards rotated field β near H. The only substantial effect on β (before the flow) comes from localisation of α, which moves vectors in α generally upwards and has the same effect on β. (This point will be discussed in more detail in the slightly more general setting of the proof of the full immersion theorem in part III—see [15; figure 1] and adjacent text for more detail here.) The flow also moves vectors upwards. Thus we are now in the situation described in the observation near the start and the field can be finally straightened by a C^∞–small global move. By choosing U and μ small enough, this final move also moves points less than $\frac{\varepsilon}{2}$. □

Remark In the proof, we could have straightened α in one C^0–small move by not phasing out the local flow. We could then have estimated the total effect of this flow: each point is affected only a finite number of times by the flow near D. However this would have obscured the local nature of the proof. Indeed we can see a local picture generated by the proof which depends on the behaviour of φ near D. In particular it depends on the intersection multiplicity. For example the picture given in figure 7 corresponds to a point of threefold intersection multiplicity. It is possible to describe the whole proof combinatorially in terms of such local pictures. The first sequence of pictures described in section 2 is sufficient for the codimension 2 case. In this description we would stratify D according to the intersection multiplicity with φ (see [1, 11]) and then construct the straightening isotopy locally, starting at points of highest multiplicity, and working down the stratification, using the pictures constructed in section 2 as 'templates'.

Straightening multiple vector fields

Now let M be embedded in $Q \times \mathbb{R}^n$ and suppose that M is equipped with n linearly independent normal vector fields. We say that M is *parallel* if the n vector fields are parallel to the n coordinate directions in $\mathbb{R}^n$.

Corollary : Multi-compression Theorem 4.5 *Suppose that M^m in embedded in $Q^q \times \mathbb{R}^n$ with n independent normal vector fields and that $q - m \geq 1$. Then M is isotopic (by a C^0–small isotopy) to a parallel embedding.*

Proof Apply the compression theorem to the first normal vector field. By the theorem M can be isotoped to make this field parallel to the first coordinate axis. Then M lies over an immersion in $Q \times \mathbb{R}^{n-1}$ with $n-1$ independent normal fields. (This can be seen by thinking of the remaining $n-1$ fields as determining a embedding of $M \times D^{n-1}$ in $Q \times \mathbb{R}^n$ which is compressed into $Q \times \mathbb{R}^{n-1}$ by the straightening of the first field.) Now consider an induced neighbourhood of M pulled back by the immersion in $Q \times \mathbb{R}^{n-1}$. This neighbourhood is made of patches of $Q \times \mathbb{R}^{n-1}$ glued together and we can apply the local compression theorem to isotope M within this neighbourhood until the second vector field is parallel to the second coordinate axis. This isotopy determines a regular homotopy within $Q \times \mathbb{R}^{n-1}$ which lifts to an isotopy of M in $Q \times \mathbb{R}^n$ finishing with an embedding which has the first two normal fields parallel to the first two axes of $\mathbb{R}^n$. Continue in this way until all vector fields are parallel. □

Addenda

The local compression theorem admits a number of extensions. Notice that these addenda improve considerably on the addenda to theorem 2.1 with the exception of addendum (iii) to 2.1 for which the analogue can be seen to be false: If a grounded vector field on M in $Q \times \mathbb{R}$ could be compressed by a *small* isotopy when $q - m = 0$ then by working in an induced neighbourhood (as in the proof above) a grounded *immersion* could be compressed. But it is easy to construct an immersion of S^1 in $\mathbb{R}^2$ with grounded perpendicular vector field, see figure 12. This immersion cannot be compressed since S^1 does not immerse in $\mathbb{R}^1$.

There are similar addenda to the multi-compression theorem which we leave the reader to state and prove.

(i) (Non-compact relative version) *The local compression theorem is true without the hypothesis that M is compact, moreover there is a far stronger relative version. Let C be a closed set in M. If M is already compressible in a neighbourhood of C then the isotopy can be assumed fixed on*

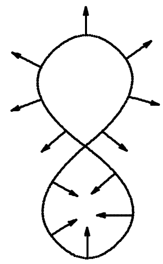

Figure 12: a grounded field on an immersed S^1

C. Moreover we can replace the real number ε in the statement of the theorem by a function ε from M to the positive reals. The conclusion is that each point x moves a distance less than $\varepsilon(x)$.

(ii) *If each component of M has relative boundary then the dimension condition can be relaxed to $q - m \geq 0$.* (See note 4.6 below.)

(iii) (Parametrised version) *Given a parametrised family $M_t^m \subset Q_t^q$ of embeddings with a normal vector field, $t \in K$, where K is a manifold of dimension k and $q - m - k \geq 1$, then there is a parametrised family of small isotopies to compressible embeddings (and there is a relative version similar to (i)).* (Note that the Normal Deformation Theorem 4.7 below admits a parametrised version with no additional dimensional condition.)

(iv) *If in (iii) the fields are all perpendicular and grounded, then the dimension condition can be relaxed to $q-m \geq 1$ and there is no dimension condition on K.*

(v) *If each component of M has relative boundary then the dimension conditions in (iii) and (iv) can each be relaxed by 1.*

(vi) *All the above statements apply to immersions of M in Q (or families of immersions). The result is a small regular homotopy (or family) to a compressible immersion (or family).*

Proofs The relative version in addendum (i), with M compact, follows from the method of proof: notice that the first (local) move in the proof of 4.4 can be assumed fixed near C and the effect of the second (C^∞–small) move near C is to move M a little through compressible embeddings. This movement can be cancelled by a small local isotopy. For the non-compact and variable ε versions, use the relative version and a patch by patch argument. To prove (ii) use a handle decomposition of M with no m–handles, and inductively apply the proof to the core of each handle, working relative to the union of the previous handles. Notice that once the core of a handle has been straightened then a small neighbourhood is also straight. But we can shrink the handle into such a neighbourhood. For (iii) we notice that the vector field can be assumed to be grounded by the dimension condition (see corollary A.5) and the two transversality results used (4.1 and 4.2) both have parametrised versions (proved in the appendix, as propositions A.2 and A.3) which state that the union of the horizontal sets and downsets in each fibre can be assumed to be manifolds. We can then define localisation near the downset in the same way as in the unparametrised version and move D in into general position with respect to the flow on $K \times M$ which is φ in each fibre, see corollary A.6. The proof then proceeds as before. Addendum (iv) follows since the dimension condition was used only to establish groundedness and for (v) we use a handle argument as in (ii). Finally for (vi) we work on an induced regular neighbourhood as in the proof of 4.5. □

Note 4.6 Strictly speaking, the isotopy in part (ii) in the case $q = m$ is not small. It is of the form: shrink to the neighbouhood of a spine and then perform a small isotopy. Thus it small in the sense that it takes place inside an arbitrarily small neighbourhood of the initial embedding. But notice that this is precisely what is needed for the proof of the multi-compression theorem.

Bundle and deformation versions

The relative version of the multi-compression theorem leads at once to a bundle version: suppose that instead of $M^m \subset Q \times \mathbb{R}^n$, we have M contained in the total space W of a bundle η^n/Q such that there is a subbundle ξ of TW defined at M and isomorphic to the pull back of η to M. Then there is an isotopy of M and ξ realising this isomorphism. The proof is to work locally where η is trivial and apply the multi-compression theorem to straighten ξ (ie realise the isomorphism with η). This can be further generalised by replacing η by a second subbundle of TW. It is convenient to state this result as a deformation theorem, which then admits a parametrised version with no extra conditions on dimension:

Normal Deformation Theorem 4.7 *Suppose that $M^m \subset W^w$ and that ξ^n is a subbundle of TW defined in a neighbourhood U of M such that $\xi|M$ is normal to M in W and that $m+n < w$. Suppose given $\varepsilon > 0$ and a homotopy of ξ through subbundles of TW defined on U finishing with the subbundle ξ'. Then there is an isotopy of M in W which moves points at most ε moving M to M' and covered by a bundle homotopy of $\xi|M$ to $\xi'|M'$ (and in particular $\xi'|M'$ is normal to M').*

Proof We work in small compact patches where we can assume that ξ and ξ' are trivial and we use the proof of the multi-compression theorem. Thinking of ξ as comprising n linearly independent vector fields, consider the first vector field α. By compactness the total angle that the homotopy of ξ moves α is bounded and we can choose r and a sequence of homotopies $\xi = \xi_0 \simeq \xi_1 \simeq \ldots \simeq \xi_r = \xi'$ so that for each $s = 1, \ldots, r-1$ the images of α in ξ_{s-1} move through an angle less than $\pi/2$ in the homotopy to ξ_s. Now apply the local proof of the compression theorem. Think of the image α_1 of α in ξ_1 as vertically up, then if α is made perpendicular it is grounded and the proof gives a small isotopy of M covered by a straightening of α which moves α by a homotopy which can be seen to be a deformation of the given homotopy (both homotopies take place in the contractible neighbourhood of α_1 comprising vectors which make an angle $< \pi$ with α_1). Repeating this r times we find a small isotopy of M, α to M', α' where α' is the first vector of ξ'. This moves ξ to ξ'' say and then if we consider the homotopy h of ξ'' to ξ' which is the reverse of the isotopy followed by the given homotopy of ξ to ξ' then h moves α' through a contractible loop and contracting this loop we obtain a homotopy of ξ'' to ξ' fixing α'. We now project onto the orthogonal complement of α' (as in the proof of the mulit-compression theorem) and proceed to straighten the next vector in the same way, again following the given homotopy. □

Addenda The covering bundle homotopy is a deformation of the given one. There is a relative version which follows directly from the proof and, since all fields are grounded, there is a parametrised version with no extra hypotheses, see proof of addenda (iii) and (iv) to the multi-compression theorem given above.

The proof can readily be modified to construct an isotopy which "follows" the given homotopy of ξ, in other words if ξ_i is the position of ξ at time i in the homotopy and M_i the position of M at time i in the isotopy, then $\xi_i | M_i$ is normal to M_i for each i. This is done by breaking the homotopy into very small steps. However this makes the final isotopy far less explicit and in that case there is a simpler proof given in Part II [14; section 3].

Finally there is a codimension 0 ($m + n = w$) version which it is worth spelling out in detail:

Suppose in the normal deformation theorem that $m + n = w$ and that M is open or has boundary and that X is a spine of M. Then there is an isotopy of M of the form: shrink into a neighbourhood of X followed by a small isotopy in W, carrying M to be normal to ξ'.

Remark The Normal Deformation Theorem is close to Gromov's theorem on directed embeddings [5; 2.4.5 C′] from which it can be readily deduced. For more detail here see the final remarks in [14].

References

[1] **J M Boardman**, *Singularities of differentiable maps*, Publ. IHES, 33 (1967) 21–57

[2] **W Browder**, **W C Hsiang**, *Some problems on homotopy theory, manifolds and transformation groups*, (Proceedings of the Stanford conference in algebraic and geometric topology, 1976), AMS Proc. Sympos. Pure Math. XXXII, 2 (1978) 251–267

[3] **Y M Eliashberg**, **N M Mishachev**, *Wrinkling of smooth maps II: wrinkling of embeddings and K Igusa's theorem*, Topology, 39 (2000) 711–732

[4] **R Fenn**, **C Rourke**, **B Sanderson**, *James bundles and applications*, `http://www.maths.warwick.ac.uk/~cpr/ftp/james.ps`

[5] **M Gromov**, *Partial differential relations*, Springer–Verlag (1986)

[6] **V Guillemin**, **A Pollack**, *Differential Topology*, Prentice–Hall (1974)

[7] **M Hirsch**, *Immersions of manifolds*, Trans. Amer. Math. Soc. 93 (1959) 242–276

[8] **I James**, *Reduced product spaces*, Annals of Math. 62 (1955) 170–197

[9] **U Koschorke**, **B Sanderson**, *Geometric interpretations of the generalised Hopf invariant*, Math. Scand. 41 (1977) 199–217

[10] **J P May**, *The geometry of iterated loop spaces*, Springer Lectures Notes Series, 271 Springer-Verlag (1972)

[11] **J N Mather**, *Generic projections*, Annals of Math. 98 (1973) 226–245

[12] **R J Milgram**, *Iterated loop spaces*, Annals of Math. 84 (1966) 386–403

[13] **A Philips**, *Turning a surface inside out*, Scientific American, 214 (1966) 112–120

[14] **Colin Rourke**, **Brian Sanderson**, *The compression theorem II: directed embeddings*, Geometry and Topology, 5 (2001) 431–440, `arxiv:math.GT/0003026`

[15] **Colin Rourke**, **Brian Sanderson**, *The compression theorem III: applications*, to appear, see `http://www.maths.warwick.ac.uk/~cpr/ftp/compIII.ps` for a preliminary version

[16] **G Segal**, *Configuration-spaces and iterated loop-spaces*, Invent. Math. 21 (1973) 213–221

[17] **S Smale**, *The classification of immersions of spheres in Euclidean spaces*, Annals of Math. 69 (1959) 327–344

[18] **B Wiest**, *Loop spaces and the compression theorem*, Proc. Amer. Math. Soc. 128 (2000) 3741–3747

[19] **B Wiest**, *Rack spaces and loop spaces*, J. Knot Theory Ramifications, 8 (1999) 99–114

A Appendix: transversality and general position

In this appendix we prove the transversality and general position results needed in the main proofs (in sections 2 and 4).

Let M^m and K^k be manifolds and let $F\colon K \times M \to \mathbb{R}^n$ be a K*–family of embeddings* of M in $\mathbb{R}^n$. Ie, F is a smooth map such that $F \mid \{t\} \times M$ is a smooth embedding for each $t \in K$. We denote $F(\{t\} \times M)$ by M_t. Let E_n be the group of isometries of $\mathbb{R}^n$. Define $J\colon E_n \times K \times M \to \mathbb{R}^n$ by $J(u,t,x) = uF(t,x)$. Let $\mathcal{G}_{n,m}$ be the Grassmannian of m-planes in $\mathbb{R}^n$. For $u \in E_n$ let $F_u\colon K \times M \to \mathcal{G}_{n,m}$ be given by $F_u(t,x) = T_{J(u,t,x)}J(\{u\} \times \{t\} \times M)$.

Proposition A.1 *Given a neighbourhood N of the identity of E_n and a submanifold W of $\mathcal{G}_{n,m}$ there is an element $u \in N$ such that $F_u\colon K \times M \to \mathcal{G}_{n,m}$ is transverse to W.*

Proof Let $G\colon E_n \times K \times M \to \mathcal{G}_{n,m}$ be given by $G(u,t,x) = F_u(t,x)$. It is sufficient to prove the derivative $TG_{(u,t,x)}$ is surjective for any (u,t,x), see [6, page 68]. In fact we prove a stronger result: the derivative on tangents to each $E_n \times \{t\} \times \{x\}$ at any (u,t,x) is surjective.

Represent a tangent vector to the Grassmannian by $\dot\alpha(0)$, where $\alpha\colon (-\epsilon,\epsilon) \to \mathcal{G}_{n,m}$ and $\alpha(0) = G(u,t,x)$. Let $q\colon O_n \to \mathcal{G}_{n,m}$ be the fibration given by $q(v) = v(\alpha(0))$ and choose $\beta\colon (-\epsilon,\epsilon) \to O_n$ such that $q\beta = \alpha$. Finally define $\gamma\colon (-\epsilon,\epsilon) \to E_n$ by $\gamma(s) = \mathrm{tr}_{J(u,t,x)} \circ \beta(s) \circ \mathrm{tr}_{-J(u,t,x)}$ where tr_z denotes translation by z. Then after identifying E_n with $E_n \times \{p\} \times \{x\}$ we have $G\gamma = \alpha$ and so $TG(\dot\gamma(0)) = \dot\alpha(0)$ as required. □

Suppose now Q^q is a manifold and we are given a family of embeddings $F\colon K \times M \to Q \times \mathbb{R}^r$ with $m \le q$. Define $H(K \times M)$, the *horizontal set*, to consist of those points $(t,x) \in K \times M$ such that the tangents to $\mathbb{R}^r$ at $F(t,x)$ are contained in the image of TF. Similarly define $C(K \times M)$, the *critical set*, to consist of points $(t,x) \in K \times M$ such that the image of tangents at (t,x) are tangent to Q. In case $Q = \mathbb{R}^q$, by considering $\mathcal{G}_{q,m-r} \subset \mathcal{G}_{q+r,m}$ and $\mathcal{G}_{q,m} \subset \mathcal{G}_{q+r,m}$, we can assume by, A.1, that after a small Euclidean motion $H(K \times M)$ and $C(K \times M)$ are submanifolds of codimension $r(q+r-m)$ and rm respectively.

The terminology has been chosen so that if M_t has a perpendicular normal field then the field must be horizontal at points $F(t,x)$ where $(t,x) \in H(K \times M)$, and if $(t,x) \in C(K \times M)$ and $r = 1$, then x is a critical point of the function $M \to \mathbb{R}$ given by $x \mapsto \pi_2 F(t,x)$, where $\pi_2\colon Q \times \mathbb{R} \to \mathbb{R}$ is projection.

Proposition A.2 *Given a family of embeddings $F\colon K\times M^m \to Q^q\times\mathbb{R}^r$ with $m \le q$, then there is an isotopy of F, ie a homotopy F_t through families of embeddings, such that $F = F_0$ and the horizontal set determined by F_1 is a submanifold of $K\times M$. Similarly for the critical set.*

Proof We prove the case of the horizontal set. The critical set is similar. The case $Q = \mathbb{R}^q$ follows from A.1. For the general case we use a standard patch by patch argument. Choose a locally finite cover of $K\times M$ by discs of the form $D = D_1 \times D_2$ such that the corresponding half discs $\frac{1}{2}D = \frac{1}{2}D_1 \times \frac{1}{2}D_2$ cover and such that each disc has image above a Euclidean patch in Q. Suppose then $F(D) \subset U \times \mathbb{R}^r$ where U is a Euclidean patch. By A.1 we can compose $F|D$ with a small Euclidean motion e of $U\times\mathbb{R}^r$ so that the horizontal set in a neighbourhood of $\frac{1}{2}D$ becomes a manifold. By choosing a path from e to the identity in E_{q+r} we can phase out the movement using a collar on $\frac{1}{2}D$ in D and then extend by the identity to $K\times M$.

We now argue by induction. Suppose the horizontal set in a neighbourhood of the union of the first k such $\frac{1}{2}$–discs is a manifold and consider the $(k+1)$-st. By choosing e as above with a sufficiently short path to the identity, the horizontal set in a neighbourhood $(k+1)$-st $\frac{1}{2}$–disc becomes a manifold without disturbing the property that the horizontal set is a manifold in (a possibly smaller) neighbourhood of the first k $\frac{1}{2}$–discs. Since the cover is locally finite, this inductively defined isotopy defines the required isotopy of F. □

We now restrict to the case of application in sections 2 and 4 namely $r = 1$ and suppose that $K\times M \to Q\times\mathbb{R}$ is a K–family of embeddings equipped with normal vector fields. Ie, we assume that the normal vector fields α vary continuously with $t \in K$. We identify $K\times M$ with its image in $K\times Q\times\mathbb{R}$ by the embedding $(t,x)\mapsto(t,F(t,x))$ and we think of α as a single vector field on $K\times M$ in $K\times Q\times\mathbb{R}$. We denote the vector field on M_t in Q corresponding to $t\in K$ by α_t.

Choose metrics and assume that α is a unit vector field such that α_t is perpendicular to M_t in $Q\times\mathbb{R}$ for each $t\in K$. We define the *downset* exactly as in section 4. Let ν be the normal bundle on $K\times M$ in $K\times Q\times\mathbb{R}$ given as the union of the normal bundles $\nu(M_t \subset Q\times\mathbb{R})$ for $t\in K$. Let ψ be the vector field on $K\times M - H(K\times M)$ in $K\times Q\times\mathbb{R}$ given by choosing $\psi_t(x)$ to point up the line of steepest ascent in $\nu_{(t,x)}$. Then the downset $D\subset K\times M - H(K\times M)$ comprises points (t,x) where $\alpha_t(x) = -\psi_t(x)$ ie, D comprises all points (t,x) such that α_t points down the line of steepest descent in $\nu_{(t,x)}$. Notice that the downset of α is precisely the union of the downsets of α_t over $t\in K$.

Proposition A.3 *After a small isotopy of α we can assume that α is transverse to $-\psi$, which implies that D is a manifold, and further we can assume that the closure of D is a manifold with boundary H.*

Proof By A.2 the horizontal set can be assumed to be a manifold H say. Since we are in the case $r = 1$, the codimension of H in $K\times M$ is $c = q+1-m$. Now assume for simplicity that $Q = \mathbb{R}^q$. Then H is given (by proposition A.1) as the transverse preimage of $\mathcal{G}_{q,m-1} \subset \mathcal{G}_{q+1,m}$, which (using the metrics) can also be seen as the transverse preimage of $\mathcal{G}_{q,c} \subset \mathcal{G}_{q+1,c}$.

Now we can identify the restriction of the canonical disc bundle $\bar{\gamma}_{q+1,c}$ to $\mathcal{G}_{q,c}$ with a closed tubular neighbourhood of $\mathcal{G}_{q,c}$ by $(p, P) \mapsto P_p \in G_{q+1,c}$, where P_p is the subspace spanned by the vector $(p, -|p|) \in \mathbb{R}^q \times \mathbb{R}^1 = \mathbb{R}^{q+1}$ and the subspace of P orthogonal to p.

Now consider α near H. Consider a fibre D_x of a tubular neighbourhood U of H in $K \times M$. Let $\pi\colon U \to H$ be the projection of U. We can identify the fibres of ν over D_x with ν_x by choosing an isomorphism of $\pi^*\nu$ with $\nu|U$. This done we can homotope α to be constant over D_x. Since ν is transverse to $\mathcal{G}_{q,c}$ we find that, if U is chosen small enough, the map $-\psi$ on ∂D_x followed by projection to the unit sphere is a diffeomorphism. So $-\psi$ and α on ∂D_x meet in one point. We now get a collar on H by considering variable x and variable radius for D_x. Consequently we have α transverse to $-\psi$ on $U - H$, and after a further homotopy of αwe have α transverse to $-\psi$ on $K \times M - H$.

The general case follows by a patch by patch argument as in A.2. □

Lemma A.4 *We can realise any small isotopy of D fixed near H by a small isotopy of α.*

Proof Consider a small patch in D. Locally D is given as the tranvserse preimage of $-\psi$ by α. A small local isotopy of α maintains transversality and has the effect of moving D inside its normal bundle in $K \times M$ to an arbitrary section near the zero section. By a combination of such moves we can realise any small isotopy of D fixed near H by a small isotopy of α. □

Corollary A.5 *Suppose that $q - m - k \geq 1$ then by a small isotopy of F and α we may assume that α_t is perpendicular and grounded for each $t \in K^k$.*

Proof By A.2 we may assume that the critical set of F is a manifold of dimension k. But recall that α_t is grounded if and only if the downset of α_t is disjoint from the critical set of M_t. The result now follows from the lemma. □

Now let φ_t be the gradient vector field on M_t for each $t \in K$ and let φ be the vector field on $K \times M$ such that $\varphi \,|\, \{t\} \times M = \varphi_t$.

Corollary A.6 *Suppose that α is perpendicular and grounded. Let U be a tubular neighbourhood of H in $K \times M$. By a small isotopy of α we may assume that $\overline{D - U}$ is in general position with respect to φ in the following sense: Given $\delta > 0$ there is a neighbourhood V of $\overline{D - U}$ in M such that each component of intersection of an integral curve of φ with V has length $< \delta$.*

Proof By the lemma and the theory of smooth maps of codimension 0 [1, 11], we may assume that each integral curve of φ meets $\overline{D - U}$ in a discrete set. The result now follows for compact K and M by a standard accumulation argument. For the non-compact case, we use a patch by patch argument as in A.2. □