

Good metric spaces
without
good parameterizations

Stephen Semmes

Abstract. A classical problem in geometric topology is to recognize when a topological space is a topological manifold. This paper addresses the question of when a metric space admits a quasisymmetric parameterization by providing counterexamples to the obvious optimistic conjectures, or, in other words, by providing examples of spaces with many Euclidean-like properties which are nonetheless substantially different from Euclidean geometry. These examples are geometrically self-similar versions of classical topologically self-similar examples from geometric topology, and they can be realized as codimension 1 subsets of Euclidean spaces. Unlike earlier examples going back to Rickman, these sets enjoy good bounds on their geodesic distance functions and good mass bounds (Ahlfors regularity). They are also smooth except for reasonably tame degenerations near small sets, they are uniformly rectifiable, and they have good properties in terms of analysis (like Sobolev and Poincaré inequalities). The construction also produces uniform domains which have many nice properties but which are not quasiconformally equivalent to balls.

1. Introduction.

How can one recognize when a metric space admits a quasisymmetric, bilipschitz, or homeomorphic parameterization by a Euclidean space?

For the purposes of this paper it will be sufficient to consider only subsets of Euclidean spaces instead of abstract metric spaces, and so we restrict the generality of our definitions accordingly. (Note however the results of Assouad [A1], [A2], [A3] on embedding metric spaces nicely into Euclidean spaces, and quasisymmetrically in particular.) A mapping $f : \mathbb{R}^d \rightarrow \mathbb{R}^n$ is said to be *quasisymmetric* (or a *quasisymmetric embedding*) if it is not constant and if there exists a homeomorphism $\eta : [0, \infty) \rightarrow [0, \infty)$ such that

$$(1.1) \quad |x - a| \leq t|x - b| \quad \text{implies} \quad |f(x) - f(a)| \leq \eta(t)|f(x) - f(b)|$$

for all $t > 0$ and $x, a, b \in \mathbb{R}^d$. This condition means that f distorts relative distances in a bounded way, *e.g.*, if a is much closer to x than b is, then the analogous statement is true for $f(a)$, $f(x)$, and $f(b)$. However, it does not prevent $|f(x) - f(a)|$ from being wildly different from $|x - a|$. The bilipschitz condition requires the stronger property that f distort absolute distances by only a bounded amount, *i.e.*,

$$(1.2) \quad C^{-1}|x - y| \leq |f(x) - f(y)| \leq C|x - y|$$

for some C and all $x, y \in \mathbb{R}^d$.

SIMPLE EXAMPLES. The map $x \mapsto x|x|^{a-1}$ on \mathbb{R}^d is quasisymmetric whenever $a > 0$, but it is bilipschitz only when $a = 1$. The mapping $x \mapsto x \exp(|x|)$ is a homeomorphism on \mathbb{R}^d but is not quasisymmetric, because quasisymmetric maps cannot grow faster than polynomially at infinity. The mapping on \mathbb{R} given by $x \mapsto |x|^a$ when $x \geq 0$ and $x \mapsto -|x|^b$ when $x \leq 0$ is a homeomorphism for all $a, b > 0$, but it is quasisymmetric only when $a = b$.

One can make more amusing examples by introducing some spiralling. Let $\{\theta_t\}$, $t \in \mathbb{R}$, be a one-parameter family of rotations on \mathbb{R}^d which is Lipschitz continuous in t . Then $x \mapsto \theta_{\log|x|}(x|x|^{a-1})$ is a quasisymmetric map on \mathbb{R}^d when $a > 0$, and it is bilipschitz when $a = 1$. For embeddings of \mathbb{R}^d into \mathbb{R}^n one can also introduce plenty of corners.

To put the above question about the existence of parameterizations into perspective let us recall a wonderful theorem of Edwards ([E], see

also [C1], [C2], [C3], [D]) to the effect that there exist finite polyhedra of dimension 5 (say) which are homeomorphic to the standard 5-sphere but not bilipschitz equivalent to it. (This formulation of the “double suspension theorem” uses also the observation made in [SS, p. 504, Remark (b)].) As this result indicates, there is some serious technology in topology for showing that a space admits a homeomorphic parameterization by a Euclidean space even when one might not expect that to be true. (See Section 2 for a little more information.)

Because of the existence of these strange polyhedral spheres and other examples in [Se5] I have come to the conclusion that bilipschitz parameterizations are too limited for understanding the structure of sets with little smoothness but reasonable behavior. In other words, the strange polyhedral spheres are just finite polyhedra, and the examples constructed in [Se5] are also very reasonable in their behavior, and so they should not receive all the blame if the bilipschitz condition is too stingy to include them in its parameterizations. This does not mean that there is no meaningful characterization of sets which admit bilipschitz parameterizations, nor that there are not plenty of interesting criteria for the existence of bilipschitz parameterizations, but simply that these criteria cannot include some otherwise very reasonable sets. So far the only general criteria known seem to be the ones in [SS] (for polyhedra) and [To1], [To2]. (See [Se5] for some open problems.)

What about quasisymmetric parameterizations? All of the examples in [Se5] were quasisymmetrically equivalent to a Euclidean space, and it is not known whether the strange polyhedral spheres of Edwards and Cannon are quasisymmetrically equivalent to standard spheres. The bottom line of this paper (Theorem 1.12) is that there exist spaces with many good properties but which do not admit quasisymmetric parameterizations. The most compelling of these is a variant of an example in [FS] of a discrete group of homeomorphisms which is not topologically conjugate to a uniformly quasiconformal group. Before getting to that let us start from scratch and consider the question of quasisymmetric parameterizations more thoroughly.

Suppose first that $d = 1$. A set $E \subseteq \mathbb{R}^n$ admits a *quasisymmetric parameterization* by \mathbb{R} if and only if it is a Jordan curve and it satisfies the Ahlfors “3-point condition”, *i.e.*, there is a constant C so that if $x, y \in E$ and A is the arc which connects them then $\text{diam } A \leq C|x - y|$. (See [TV, Section 4].)

For the $d = 2$ case there is a result [Tu2, Lemma 4] to the effect that the product of a nonrectifiable arc with a line segment cannot be

embedded quasisymmetrically into \mathbb{R}^2 . Assertions of this nature were established first for snowflake curves by Rickman, and then general results were obtained by Väisälä and Tukia. (See [Tu2], [V1].) Keep in mind that there are plenty of nonrectifiable Jordan curves which satisfy the Ahlfors 3-point condition, like snowflake curves. This result allows one to build counterexamples to many reasonable conjectures about the existence of quasisymmetric (and quasiconformal) mappings (see [Tu2]), and thereby makes it difficult if not impossible to find reasonable criteria for the existence of a quasisymmetric parameterization when $d = 2$ without imposing some conditions on the mass. However, it turns out that if we do impose such a condition, then there is a nice positive result (Theorem 1.6 below). Before stating it we need a couple of definitions.

Definition 1.3. *A subset E of \mathbb{R}^n is said to satisfy Condition (*) (with dimension d) if it is closed and if there is a constant C such that for each $x \in E$ and $r > 0$ there is a (relatively) open set $U \subseteq E$ such that $E \cap B(x, r) \subseteq U \subseteq E \cap B(x, Cr)$ and U is homeomorphic to a d -ball.*

Condition (*) is necessary for E to be quasisymmetrically equivalent to \mathbb{R}^d . (A quasisymmetric map takes a ball to a set with approximately the same shape as a ball.) It is not sufficient, however, even when $d = 2$, because of the examples described above. (Note that [Tu2] also covers products of unrectifiable curves with higher-dimensional (standard) cells. See [AV] for related results for products of topological cells when both are permitted to have dimension larger than 1.)

When $d = 1$ Condition (*) implies that E is a Jordan curve which satisfies the Ahlfors 3-point condition, and so it does actually imply quasisymmetric equivalence with \mathbb{R} . In general Condition (*) tries to capture some of the geometry implied by the existence of a quasisymmetric parameterization. For instance, Condition (*) forbids cusps and long thin tubes, not to mention crossings.

The next condition requires that the given set be well behaved in terms of Hausdorff measure, and it will be used to avoid the above counterexamples.

Definition 1.4. *A subset E of \mathbb{R}^n is said to be (Ahlfors) regular of dimension d if it is closed and if there is a constant $C > 0$ so that*

$$(1.5) \quad C^{-1} r^d \leq H^d(E \cap B(x, r)) \leq C r^d$$

for all $x \in E$ and $r > 0$. Here H^d denotes d -dimensional Hausdorff measure (and not cohomology).

A d -plane in \mathbb{R}^n is regular with dimension d , and the same is true of any set which is bilipschitz equivalent to a d -plane. In general (Ahlfors) regularity means that E behaves measure-theoretically like a d -plane, but it can still be very different from a d -plane geometrically. For instance, for each $0 < d < n$ there are Cantor sets in \mathbb{R}^n which are regular with dimension d . There are also snowflake curves and tree-like sets which are regular (with dimension larger than 1). Note that regularity is not necessary for quasisymmetric equivalence with \mathbb{R}^d . However, when $d = 2$ regularity and Condition (*) together imply the existence of a quasisymmetric parameterization, modulo some a priori smoothness assumptions.

Theorem 1.6. *Suppose that $E \subseteq \mathbb{R}^n$ is regular and satisfies Condition (*), both with dimension 2. Suppose also that E is smooth and well behaved at infinity. (We need to assume enough to ensure that E is conformally equivalent to the plane.) Then E is quasisymmetrically equivalent to \mathbb{R}^2 , with a choice of η as in (1.1) which depends on the constants from Definitions 1.3 and 1.4 but which does not depend on our a priori smoothness assumptions in a quantitative way.*

Theorem 1.6 was proved in [Se2, Section 5]. (See [DS3, Section 6] for a related result.) The argument in [Se2] went in two steps, as follows: the a priori assumptions on E together with the uniformization theorem were used to obtain the existence of a conformal parameterization of E , and then classical methods were used to show that the geometric assumptions imply that the conformal parameterization is quasisymmetric with uniform bounds. It turns out that the second step can be made to work in great generality; it is proved in [HK] that a quasiconformal parameterization of a metric space which satisfies certain simple geometric properties is actually quasisymmetric. (See [HK] for the precise statement.) The first step -the uniformization theorem- is of course very special to dimension 2, and in fact Theorem 1.6 fails in dimension 3, as we shall see.

Note that if a set E admits a quasisymmetric parameterization f by \mathbb{R}^d and is regular with dimension d then f satisfies the same kind of estimates as in [Ge]. (That is, the pull-back of $H^d|_E$ to \mathbb{R}^d via f is an A_∞ weight, by the same argument as in [Ge]. See [DS1] and [Se3].

This result implies strong restrictions on the way that f can distort distances, and it implies local Sobolev space conditions on f and its inverse.) Such a set E has many of the same properties as it would if it admitted a bilipschitz parameterization. (This would not be true if E were regular but with a different dimension than d .) When $d = 2$ it is not known whether E might actually be bilipschitz equivalent to \mathbb{R}^2 under these conditions, but there are counterexamples when $d = 3$, by [Se5].

Let us consider a couple of other geometric conditions on sets of roughly the same spirit as (*). The first is a more uniform version of (*), while the second is weaker and is given in terms of contractability properties.

Definition 1.7. *A subset E of \mathbb{R}^n is said to satisfy Condition (**) (with dimension d) if it is closed and if there is a constant C and a locally bounded function $\omega : [0, \infty) \rightarrow [0, \infty)$ with $\lim_{t \rightarrow 0} \omega(t) = 0$ such that for each $x \in E$ and $r > 0$ there is a (relatively) open set $U \subseteq E$, $E \cap B(x, r) \subseteq U \subseteq E \cap B(x, Cr)$, and a homeomorphism g from the unit ball B_d in \mathbb{R}^d onto U such that*

$$(1.8) \quad |g(y) - g(z)| \leq r \omega(|y - z|), \quad \text{for all } y, z \in B_d,$$

and

$$(1.9) \quad |g^{-1}(v) - g^{-1}(w)| \leq \omega(r^{-1}|v - w|), \quad \text{for all } v, w \in U.$$

Roughly speaking, Condition (**) does for sets what quasisymmetry does for mappings.

The difference between this and Condition (*) is that we require here a uniform bound on the moduli of continuity of the homeomorphic parameterizations of the topological d -balls U (and also on the moduli of continuity of their inverses). We do not require that this modulus of continuity be anything in particular - *e.g.*, we do not require Hölder continuity- and we have been careful in (1.8) and (1.9) to make the estimates scale-invariant.

Finite polyhedra which are also topological manifolds (without boundary) provide an amusing class of sets which satisfy (**), or rather the obvious counterpart of (**) for compact sets, in which we consider only small radii. (See Section 11.) The results of Edwards and Cannon imply that there are many strange examples of such polyhedra.

Notice that Condition (**) is satisfied if E admits a quasisymmetric parameterization by \mathbb{R}^d . The difference between the two properties is basically that the homeomorphisms g in Definition 1.7 are allowed to depend on x, r in a completely arbitrary fashion, while the existence of a quasisymmetric parameterization means that all the g 's can be obtained from a global parameterization of E in a certain way.

One can also look at the uniformity required in Condition (**) in terms of compactness. To understand this point it is helpful to recall the following compactness property of quasisymmetric mappings.

Lemma 1.10. *Suppose that $f_j : \mathbb{R}^3 \rightarrow \mathbb{R}^4$ is a sequence of quasisymmetric embeddings which satisfy (1.1) with a fixed choice of η . Suppose also that there are two points $a_0, a_1 \in \mathbb{R}^3$ and a positive constant C such that $|f_j(a_i)| \leq C$ and $|f_j(a_0) - f_j(a_1)| \geq C^{-1}$ for all i, j . Then there is a subsequence of $\{f_j\}$ which converges uniformly on compact subsets of \mathbb{R}^3 to an η -quasisymmetric embedding $f : \mathbb{R}^3 \rightarrow \mathbb{R}^4$.*

This is a well-known and simple consequence of the Arzela-Ascoli theorem. It implies compactness properties for the set of subsets of \mathbb{R}^4 which are quasisymmetrically equivalent to \mathbb{R}^3 with respect to the Hausdorff topology. Condition (**) has a similar compactness property built in to the definition, but without the benefit of a single parameterization which incorporates all the estimates. By contrast, we shall see that Condition (*) fails to enjoy such compactness.

We shall also consider a condition weaker than (*).

Definition 1.11. *A subset E of \mathbb{R}^n satisfies Condition (†) if it is a topological manifold and if there is a $C > 0$ such that $x \in E$ and $r > 0$ imply that $E \cap B(x, r)$ can be contracted to a point inside $E \cap B(x, Cr)$.*

This kind of uniform contractability condition has gained in prominence in recent years. (See [F1], [F2], [GP], [GPW], [P].) Condition (*) implies (†), and the two have similar features, e.g., they both prevent cusps and long thin tubes. It turns out that (†) implies (*) stably, in the sense that (†) for E implies (*) for $E \times \mathbb{R}$, by a theorem of Ferry [F2], at least when E has topological dimension greater or equal than 4. (Basically this comes from [F2, Theorem 4.1], but I am cheating slightly here, because Ferry works only with compact sets and the corresponding versions of (†) and (*).)

The following is the main result of this paper.

Theorem 1.12. a) *There exists a set in \mathbb{R}^4 which satisfies (†) (with dimension 3) but not (*).*

b) *There exists a set in \mathbb{R}^4 which satisfies (*) (with dimension 3) but not (**).*

c) *There exists a set in \mathbb{R}^4 which satisfies (**) (with dimension 3) but which does not admit a quasisymmetric parameterization.*

All of these sets can be chosen to be Ahlfors regular with dimension 3 and to have the property that there is a constant L_0 so that every pair of points p, q in the set is contained in a closed subset W of the set such that W is L_0 -bilipschitz equivalent to a closed Euclidean 3-ball. (In particular, p and q can be connected by a curve inside the set with length less or equal than $L_0^2 |p - q|$. Thus the Euclidean distance on these sets are comparable in size to the internal geodesic distances on them.) These sets can also be taken to agree with a 3-plane outside a large ball, and to be homeomorphic to \mathbb{R}^3 .

For the record, the statement that two sets are C -bilipschitz equivalent means that there is a bijection between them which satisfies (1.2) for all pairs of points in the domain of the mapping.

The sets in Theorem 1.12 are smooth away from small singular sets, and the degeneracies near the singularities are well controlled (and have a natural self similarity to them). The statement about bilipschitz balls implies also that these sets are uniformly rectifiable in the sense of [DS4]. Thus there is nothing fractal happening, not even asymptotically. These sets are also well behaved in terms of analysis, and in particular there are Sobolev and Poincaré inequalities for them. (See Sections 9 and 10.)

Note that examples as in c) existed before, because of [Tu2], but without Ahlfors regularity (of the correct dimension), or the property about pairs of points being contained in bilipschitz balls, or the good bounds on the geodesic distances.

Part a) of Theorem 1.12 answers a question of Ferry and shows that the stabilization in his theorem is necessary. (He gave a slightly different version of this necessity in [F2, Theorem 4].)

What does Theorem 1.12 mean? The glib answer is that it means that one should assume more than Ahlfors regularity and (**) if one wants to have geometric criteria for the existence of a quasisymmetric parameterization. I am inclined more to the view that the examples

of Theorem 1.12 mean that quasimetric mappings are too rigid to accommodate some reasonable geometric phenomena, and that there is an interesting middle zone where geometry is approximately Euclidean and good enough for a lot of analysis but still substantially different from Euclidean geometry.

Note that quasiconformal parameterizations of the sets promised in Theorem 1.12 must be quasimetric and hence cannot exist. See the comments after Corollary 3.104.

The examples promised in Theorem 1.12 are basically geometric reformulations of classical examples from geometric topology and were inspired by the pictures in [D, Section 9]. Geometric topology is unusual in mathematics for its wealth of concrete examples, including some very interesting (topological) quotients of Euclidean spaces with a lot of topological self-similarity. The main point of Theorem 1.12 is that one can construct sets which contain the same topological information as in these examples but for which the topological self-similarity is converted into actual geometric self-similarity. Specifically, the examples for a), b), and c) of Theorem 1.12 correspond to the construction of the Whitehead continuum, Bing's dogbone space, and Bing doubling, respectively. It is easy to construct similar sets corresponding to other examples of the same type, but the main points of this general procedure are illustrated well by these three examples. The topological features of these classical examples imply interesting properties of the sets constructed here, from which the requirements of a) and b) will follow immediately. In the case of c) there is an additional point which was covered in [FS], and in fact Theorem 1.12.c) turns out to be almost just a reformulation of an example in [FS].

Juha Heinonen pointed out to me that the complementary components of the sets used to prove Theorem 1.12 are "uniform domains" (see (7.13) and (7.14)), and hence cannot be quasiconformally equivalent to a ball. (If they were, then there would have to be a quasimetric extension to the boundary, which is impossible.) In particular, if these domains are equipped with their quasihyperbolic metrics, then they cannot be bilipschitzly equivalent to the standard hyperbolic space. This is amusing, because these complementary domains are otherwise so nice. They are topological balls, with topologically tame boundaries, and we can even build them so that there are bilipschitz reflections across their boundaries (*i.e.*, across the sets promised in Theorem 1.12). The complementary components of the set in c) are particularly nice, because they satisfy a version of (**) adapted to domains. (See Theorem 8.1

below.) One could also build examples with the same properties using [Tu2, Lemma 4] (as discussed on [Tu2, p. 518]), but not with the control on the mass, uniform rectifiability, etc., which is available here.

Note that there are general results in [V2] which permit one to conclude that the complementary components of a set are uniform domains under natural uniform conditions on the topology of the set itself. One can also derive higher-order versions of the uniform domain condition in this way. In our case we shall be able to check easily and directly that the domains are uniform, but one should keep the general results in mind.

To understand the details of this paper it could be very helpful to have a copy of [D] handy. In particular [D] has excellent pictures. I shall provide references which permit one to do without [D], but the wonderfully clear book [D] provides the advantages of one-stop shopping. Good general references for aspects of geometric topology related to the topics of this paper include [C1], [C2], [D], [E], [K], and many papers of R. H. Bing.

The paper [Se5] addresses similar questions about bilipschitz instead of quasimetric mappings and relies on different examples from geometric topology (like Antoine's necklaces).

Some background information about geometric topology will be given in the next section, and a general construction will begin after that. Specific examples corresponding to Theorem 1.12 are described in Sections 4, 5, and 6, and the complements of these sets are discussed in Sections 7 and 8. Section 9 deals with analysis on these sets, with suitable versions of Sobolev and Poincaré inequalities established in Section 10. The simple fact that finite polyhedra which are topological manifolds satisfy the analogue of (**) for compact sets is given in Section 11, and the last section is devoted to miscellaneous remarks.

2. Some geometric topology.

The examples for Theorem 1.12 come from classical geometric topology, but before getting to that let us review briefly some of the more modern activity concerning the problem of finding homeomorphic parameterizations of a space. A good reference is the expository paper [C1], which begins with the following:

Recognition Problem 2.1. *Find a short list of topological properties,*

reasonably easy to check, that characterize topological manifolds among topological spaces.

In dimensions 5 and higher there is now a reasonable characterization of topological manifolds coming from work of Edwards and Quinn [E], [Q1], [Q2], modulo a locally defined integer obstruction whose non-triviality was only recently established by Bryant, Ferry, Mio, and Weinberger [B^∞]. That is, they construct spaces which are “almost” manifolds but which have the wrong value of the aforementioned obstruction. Their construction is quite complicated, and the spaces they produced are not well understood. These spaces should be very interesting, because they are so well behaved but still distinct from Euclidean topology, and one can hope that they have interesting geometric realizations on which one can do some analysis.

To put the recognition problem into perspective it is helpful to consider the special case of finite polyhedra. It turns out that a finite polyhedron K is a topological manifold (without boundary) if and only if the link of every simplex in K is a homology sphere of the correct dimension, and if the links of vertices are simply connected when $\dim K > 2$. This result follows from the theorem of Edwards and Cannon [C2], [C3], [E] to the effect that double suspensions of homology spheres are topological spheres. When the dimension is 4 one must also use the Freedman theory [Fr]. As indicated in the introduction, it is known that the local homeomorphic parameterizations promised in this theorem cannot be bilipschitz in certain cases (see [SS] for details), and the existence of quasisymmetric homeomorphisms is an open problem.

The bottom line is that topologists have some serious technology for establishing the existence of homeomorphic parameterizations. They tend not to provide any estimates on the extent to which their parameterizations distort distances, but there are also examples which show that the homeomorphisms which they produce have to be complicated.

Of course the preceding discussion about the existence of homeomorphic parameterizations ignores local-to-global issues of the type addressed by the Schönflies theorem, the annulus and Poincaré conjectures, the H and S cobordism theorems, and surgery theory. The local and global questions are not truly separate -*e.g.*, the properties of a space at infinity can be reformulated in terms of local properties of the one-point compactification near the new point- and the distinction is particularly dubious in the context of parameterizations with

scale-invariant bounds, as in bilipschitz and quasimetric conditions, where an obstruction to doing something globally can give rise to obstruction to doing it locally with a uniform bound.

Now let us consider a much more restricted version of the recognition problem which will be more directly connected to the proof of Theorem 1.12. There is an old result of R. L. Moore to the effect that a Hausdorff topological space X is homeomorphic to S^2 if it can be realized as the image of a continuous mapping $f : S^2 \rightarrow X$ with the property that $f^{-1}(x)$ and $S^2 \setminus f^{-1}(x)$ are nonempty and connected for each $x \in X$. This is not to say that f is itself a homeomorphism; f could collapse an arc or a disk to a point, for instance. Observe that if we collapse a circle in S^2 to a point we get an X which consists of two 2-spheres touching at a point. In this case both the hypothesis and conclusion of Moore's theorem fail to hold.

What happens in 3 dimensions? It turns out that there are some subtle negative results. To explain this it is helpful to introduce some auxiliary notions. (See [D] for details.) A *decomposition* G of \mathbb{R}^3 is simply a partition of \mathbb{R}^3 , i.e., a collection of disjoint subsets of \mathbb{R}^3 whose union is all of \mathbb{R}^3 . (See [D, bottom of p. 7], or [K, p. 86].) Given a decomposition G of \mathbb{R}^3 we can form the usual quotient space \mathbb{R}^3/G (as a topological space). Under reasonable hypotheses (which include the requirement that each element of G be a closed subset of \mathbb{R}^3) one knows that \mathbb{R}^3/G is a Hausdorff space, and every decomposition that we shall consider will have a Hausdorff quotient. One would like to know when \mathbb{R}^3/G is a topological manifold, which might even be homeomorphic to \mathbb{R}^3 itself. Moore's theorem provides a nontrivial criterion for this in 2 dimensions, but this is a much harder problem in 3 dimensions.

If F is a closed subset of \mathbb{R}^3 , then we can define a decomposition associated to it by taking G to consist of F together with $\{x\}$ for all $x \in \mathbb{R}^3 \setminus F$. Thus \mathbb{R}^3/G is simply the space that one gets by shrinking F to a point and leaving the rest of \mathbb{R}^3 alone. For instance, if one takes F to be a line segment, then \mathbb{R}^3/G is homeomorphic to \mathbb{R}^3 again. This is also true if one takes F to be a (standard) closed 2-disk, or a (standard) closed 3-ball. If F is taken to be the unit sphere, then \mathbb{R}^3/G is homeomorphic to \mathbb{R}^3 with a 3-sphere attached to it at one point.

What happens if F is a (standard) circle? Set $X = \mathbb{R}^3/G$, and let p denote the point in X which corresponds to F . Then $X \setminus \{p\}$ is not simply connected, because it is homeomorphic to $\mathbb{R}^3 \setminus F$, and we can take a circle which links F to get a homotopically nontrivial loop in $X \setminus \{p\}$. This implies that X is not homeomorphic to \mathbb{R}^3 , and a

local version of the same argument shows that X is not a topological manifold at p .

This example shows that the most naive transcription of Moore's theorem does not work in 3-dimensions, since the complement of F in \mathbb{R}^3 is connected. There is a more interesting example based on the Whitehead continuum W in \mathbb{R}^3 . The main properties of this continuum are the following: (i) it is cell-like, which means that it can be contracted to a point inside of any open set which contains it; (ii) when viewed as a subset of \mathbb{S}^3 , its complement is contractible (unlike a circle, as in the previous example); (iii) there is an open set $U \subseteq \mathbb{R}^3$ which contains W such that there are loops in arbitrarily small neighborhoods of W which do not intersect W but which cannot be contracted to a point in $U \setminus W$. Equivalently, although $\mathbb{S}^3 \setminus W$ is contractible, it is not simply connected at infinity. Notice that a standard line segment satisfies (i) and (ii) above, but not (iii). If we let G be the decomposition associated to $F = W$ as before, then $X = \mathbb{R}^3/G$ is again not homeomorphic to \mathbb{R}^3 . Indeed, if $p \in X$ corresponds to W , then $X \setminus \{p\}$ is simply connected in this case, but there is an open set V (the image in X of U in (iii)) which contains p and which has the property that there are loops in $X \setminus \{p\}$ which are as close as we want to p but which are not homotopically trivial in $V \setminus \{p\}$. Thus X is not a topological manifold at p .

The Whitehead continuum can be constructed through an iterative procedure as follows. One starts with a solid torus T in \mathbb{R}^3 and another solid torus T_1 embedded inside T as in [D, Figure 9-7, p. 68]. (See also [K, p. 81ff].) T_1 should be embedded into T in such a way that it is homotopically trivial but clasped to make it isotopically nontrivial. That is, T_1 cannot be deformed to a standard torus inside T without crossing itself. One then iterates this construction by identifying T_1 with T to get a new torus T_2 inside T_1 , and then repeating the process indefinitely to get a sequence of nested solid tori T_j . The Whitehead continuum is obtained by taking the intersection of the T_j 's. The key property (iii) above can be reduced to the fact that the meridional circle in ∂T pictured in [D, Figure 9-7] cannot be contracted to a point in T without touching the intersection of the T_j 's. (See [D, Proposition 9, p. 76] and the remarks which precede it. See also [K, p. 82].)

We shall use the Whitehead continuum in Section 4 to prove Theorem 1.12.a), and we shall employ similar iterative constructions for b) and c). Before we consider specific examples in detail we should formulate this iterative procedure in more general terms, starting with a basic definition from [D, bottom of p. 61].

Definition 2.2. A defining sequence in \mathbb{R}^3 is a sequence $\{C_i\}$ of closed subsets of \mathbb{R}^3 such that each C_i is the closure of a bounded open set with smooth boundary and such that C_{i+1} is contained in the interior of C_i for each i . The C_i 's need not be connected.

Given a defining sequence $\{C_i\}$ in \mathbb{R}^3 , we can define a decomposition G of \mathbb{R}^3 by taking the elements of G to be the components of $\cap C_i$ together with the singletons from $\mathbb{R}^3 \setminus \cap C_i$. In the earlier discussion of the Whitehead continuum each $C_i = T_i$ had only one component, as did $\cap C_i$. In the other examples considered here the number of components of C_i grows exponentially.

The iterative procedures that we shall use will be represented by the following. For the record, a "domain" is a connected open set.

Definition 2.3. An initial package \mathcal{P} consists of a bounded smooth domain D in \mathbb{R}^3 , a finite collection D_1, \dots, D_n of smooth subdomains with disjoint closures contained in D , and mappings ϕ_j , $j = 1, \dots, n$, such that each ϕ_j is a diffeomorphism from a neighborhood of \overline{D} onto a neighborhood of \overline{D}_j which maps D onto D_j .

For example, in the construction of the Whitehead continuum we had an initial package (with $n = 1$) consisting of $D = T$, $D_1 = T_1$, and any reasonable choice of ϕ_1 .

To an initial package \mathcal{P} we can associate a defining sequence $\{C_i\}$ by taking C_0 to be \overline{D} , C_1 to be $\cup \overline{D}_j$, C_2 to be the union of the images of the \overline{D}_k 's under ϕ_j inside each \overline{D}_j , and so forth. Thus C_i is the union of n^i images of \overline{D} under various compositions of the ϕ_j 's.

As explained in [D, Section 9], there are some very interesting decompositions of \mathbb{R}^3 which can be obtained from an initial package in this way. One such decomposition G , due to Bing [B3], has a nonmanifold quotient \mathbb{R}^3/G (called "the dogbone space") even though G has the property of being *cellular*. A compact set K in \mathbb{R}^3 is said to be cellular if for each open set $U \subseteq \mathbb{R}^3$ with $U \supseteq K$ there is a topological 3-ball contained in U which contains K , and a decomposition G of \mathbb{R}^3 is said to be cellular if each of its elements is a cellular set. (See p. 35 and Corollary 2A on p. 36 of [D].) The Whitehead continuum is definitely not cellular [D, p. 76, Proposition 9], and at one point it was apparently hoped that a cellular quotient of \mathbb{S}^3 would be \mathbb{S}^3 again. Bing's example shows that this is not true, and we shall use it to prove Theorem 1.12.b).

There is an older example which Bing considered in [B1] for which the quotient is homeomorphic to \mathbb{R}^3 but in a nontrivial way. For instance, this decomposition has a symmetry about a 2-plane which gives rise to an involution on S^3 whose fixed point set is a wild sphere. This example was also used in [FS], and it will be used here to prove Theorem 1.12.c).

Our next task will be to take the construction of a decomposition and an associated quotient of \mathbb{R}^3 from an initial package and adapt it in such a way as to have better geometric properties. For a bare bones version of Theorem 1.12 all we really need to do is deform the Euclidean metric smoothly on D in such a way that the ϕ_j 's become similarities near ∂D . A straightforward iterative construction would then allow us to build a metric with respect to which the ϕ_j 's are similarities on all of \overline{D} . This metric would deteriorate near the elements of the decomposition, as it should. In this way we can build a metric space which is topologically equivalent to the decomposition space and which has nice geometric self-similarity properties. For the sake of concreteness and other benefits it is better to construct these spaces as subsets of Euclidean spaces, and the examples used in this paper even fit into \mathbb{R}^4 . In fact their embeddings into \mathbb{R}^4 and their complementary components have some especially nice properties for which we shall take extra care to make manifest. If one simply wants to build such sets without worrying about extra properties then some of the efforts and assumptions in the next section are unnecessary. (See the remarks after Definition 3.2.)

3. The general construction.

The next definition describes excellent packages, which consist of an initial package \mathcal{P} together with objects which will allow us to convert \mathcal{P} into a topologically equivalent object in \mathbb{R}^4 for which the analogue of the ϕ_j 's are similarities. Recall that a (Euclidean) similarity on \mathbb{R}^4 is an affine transformation which is a combination of a translation, dilation by a positive number (called the dilation factor of the similarity), and an orthogonal transformation.

Convention 3.1. *P denotes the $x_4 = 0$ hyperplane in \mathbb{R}^4 , and from now on we identify it with \mathbb{R}^3 , so that any objects living on \mathbb{R}^3 (like an initial package) will be viewed as living on P .*

Definition 3.2. *An excellent package \mathcal{E} consists of an initial package $\mathcal{P} = \{D, D_1, \dots, D_n, \phi_1, \dots, \phi_n\}$ as in Definition 2.3 together with a bounded smooth domain Ω in \mathbb{R}^4 , smooth subdomains $\omega_1, \dots, \omega_n$ of Ω with disjoint closures contained in Ω , another collection of smooth subdomains $\Omega_1, \dots, \Omega_n$ with disjoint closures contained in Ω , orientation-preserving similarities ψ_j , $1 \leq j \leq n$, on \mathbb{R}^4 which all preserve P and have the same dilation factor $\rho \in (0, 1/10)$, and a diffeomorphism θ on \mathbb{R}^4 which satisfy the following properties:*

$$(3.3) \quad \overline{\Omega} \cap P = \overline{D}, \quad \overline{\omega_j} \cap P = \overline{D_j}, \quad \text{and} \\ \partial\Omega, \partial\omega_j \text{ all intersect } P \text{ transversely};$$

$$(3.4) \quad \psi_j(\Omega) = \Omega_j \text{ for all } j;$$

$$(3.5) \quad \theta(\Omega) = \Omega \text{ and } \theta = \text{the identity on } \mathbb{R}^4 \setminus \Omega \\ \text{and on a neighborhood of } \partial\Omega;$$

$$(3.6) \quad \theta(\omega_j) = \Omega_j \text{ and } \theta = \psi_j \circ \phi_j^{-1} \\ \text{on a neighborhood of } \overline{D_j} \text{ in } P \text{ for each } j.$$

Thus Ω and the ω_j 's are fattened-up versions of D and the D_j 's in \mathbb{R}^4 , the Ω_j 's are "straightened" versions of the ω_j 's, and θ converts the slightly twisted D_j 's into the straighter $\Omega_j \cap P$'s at the cost of deforming P inside Ω .

These excellent packages have more structure than we actually need for Theorem 1.12. Instead of the requirement that θ exist as a diffeomorphism on all of \mathbb{R}^4 it would be enough to have θ as an embedding of P into \mathbb{R}^4 , say. If we were also willing to work in a larger space than \mathbb{R}^4 then this weaker version of an excellent package would exist for any initial package. (\mathbb{R}^7 is plenty large enough.) As it is, the above definition of an excellent package imposes some topological restrictions on an initial package. These restrictions will be satisfied in the examples that we shall consider, and the extra structure that we obtain will be pleasant to have. Note that the kind of bare-bones deformation of the metric for an initial package described at the end of Section 2 involves no topological obstructions whatsoever.

For the rest of this section we shall assume that we are given an excellent package \mathcal{E} as above and build some surfaces from it. This construction is similar to the one used in [Se5, Section 5] (with Lemma 5.6 there providing the excellent package).

Define $M^1 \subseteq \mathbb{R}^4$ by $M^1 = \theta(P)$. Thus M^1 is the same as P outside Ω , it agrees with $\psi_j(D) = \Omega_j \cap P$ inside Ω_j , and it is some smooth manifold in $\Omega \setminus \cup \Omega_j$. We want to define a sequence of submanifolds M^j with many twists like the ones in M^1 , but to do this we should first sort out the relevant codings.

Let \mathcal{S}_l denote the finite sequences $\alpha = \{a_i\}_{i=1}^l$ with l terms such that $a_i \in \{1, \dots, n\}$ for all i . Thus we can identify \mathcal{S}_1 with $\{1, \dots, n\}$, and it will be useful to consider \mathcal{S}_0 as a set with just one element, the empty sequence, sometimes denoted \emptyset . Define ψ_α and Ω_α for $\alpha \in \mathcal{S}_l$ recursively in the following manner. If $l = 1$, so that we can view α as an element of $\{1, \dots, n\}$, then we simply take ψ_α and Ω_α to be as in Definition 3.2 above. If $l > 1$ and $\alpha = \{a_i\}_{i=1}^l$, then let $\alpha' = \{a_i\}_{i=1}^{l-1}$ be the “parent” of α in \mathcal{S}_{l-1} and set $\Omega_\alpha = \psi_{\alpha'}(\Omega_{a_l})$ and $\psi_\alpha = \psi_{\alpha'} \circ \psi_{a_l}$. We view the empty sequence in \mathcal{S}_0 as being the parent of the elements of \mathcal{S}_1 , so that the preceding equations hold with $\Omega_\emptyset = \Omega$ and ψ_\emptyset taken to be the identity. We shall call two elements of \mathcal{S}_l “siblings” if they have the same parent, and extend this terminology to the Ω_α ’s as well. With these conventions we have the following properties for ψ_α and Ω_α for each $\alpha \in \mathcal{S}_l$, $l \geq 1$, with α' the parent of α :

$$(3.7) \quad \begin{aligned} \psi_\alpha \text{ is a similarity with dilation factor } \rho^l, \\ \text{and } \psi_\alpha(P) = P, \end{aligned}$$

$$(3.8) \quad \Omega_\alpha = \psi_\alpha(\Omega),$$

$$(3.9) \quad \begin{aligned} \overline{\Omega}_\alpha \subset \Omega_{\alpha'} \text{ and} \\ \overline{\Omega}_\alpha \text{ is disjoint from its siblings in } \Omega_{\alpha'}. \end{aligned}$$

The Ω_α ’s have additional nesting properties, which we state as a lemma.

Lemma 3.10. *Suppose that $\alpha \in \mathcal{S}_l$ and $\beta \in \mathcal{S}_k$. Then either Ω_α and Ω_β are disjoint (and have disjoint closures), or $\Omega_\alpha \subseteq \Omega_\beta$, in which case $l \geq k$ and α is a descendant of β , or $\Omega_\beta \subseteq \Omega_\alpha$, in which case $k \geq l$ and β is a descendant of α . In particular Ω_α and Ω_β are disjoint (and have disjoint closures) when $k = l$ and $\alpha \neq \beta$.*

To see this choose $\gamma \in \mathcal{S}_m$ to be the common ancestor of α and β with m as large as possible (but perhaps = 0). If γ is equal to either α or β , then one is an ancestor of the other, and we are in business. Otherwise, α and β are descended from distinct children of γ , and disjointness follows from (3.9).

We can define ϕ_α and D_α in the same way as ψ_α and Ω_α , so that ϕ_α and D_α are given to us as part of our initial package when $\alpha \in \mathcal{S}_1$, $\phi_\alpha = \phi_{\alpha'} \circ \phi_{\alpha_l}$ when $l > 1$, $\alpha = \{a_i\}_{i=1}^l$, and α' is the parent of α , and D_α is defined to be $\phi_\alpha(D)$. The D_α 's satisfy the same nesting properties as the Ω_α 's, in the sense that the analogue of Lemma 3.10 for the D_α 's is true, with the same proof. If we set $C_l = \cup_{\alpha \in \mathcal{S}_l} \overline{D}_\alpha$, then this is equivalent to the defining sequence mentioned after Definition 2.3.

Set

$$(3.11) \quad F = \bigcap_{l=1}^{\infty} \bigcup_{\alpha \in \mathcal{S}_l} \overline{\Omega}_\alpha,$$

and let \mathcal{S} denote the collection of all infinite sequences $s = \{s_i\}$ of elements of $\{1, \dots, n\}$, so that \mathcal{S} is the natural limit of the \mathcal{S}_l 's. Thus F is a Cantor set in \mathbb{R}^4 which actually lies in P , because of (3.7) and (3.8), and there is a natural bijection $f : \mathcal{S} \rightarrow F$ which is defined in the obvious manner. (Each element s of \mathcal{S} determines a nested sequence of Ω_α 's which converges to a point (by (3.7)-(3.9)), and $f(s)$ is defined to be this point. Conversely, every element of F must arise from such a nested sequence of Ω_α 's, by Lemma 3.10.) If we perform the same construction for the D'_α 's instead of the Ω_α 's, then we might not get a Cantor set but a more complicated set with nontrivial components. These components are the nontrivial elements of the decomposition associated to the defining sequence C_l above (as described just after Definition 2.2). (See Sublemma 3.40.)

Set $Y = \overline{\Omega} \setminus \cup_{j=1}^n \Omega_j$ and $Y_\alpha = \psi_\alpha(Y)$. These compact sets in \mathbb{R}^4 are the closures of smooth domains. As before let us set $Y_\emptyset = Y$, where \emptyset is the empty sequence in \mathcal{S}_0 . If $\alpha \in \mathcal{S}_l$ and $\beta \in \mathcal{S}_k$, and if Y_α intersects Y_β , then one of α, β is the parent of the other, and their intersection equals $\partial\Omega_\gamma$, where γ is whichever of α and β is the child. This is an easy consequence of Lemma 3.10, (3.9), and the definitions. Notice also that

$$(3.12) \quad \mathbb{R}^4 = (\mathbb{R}^4 \setminus \Omega) \cup \left(\bigcup_{l=0}^m \bigcup_{\alpha \in \mathcal{S}_l} Y_\alpha \right) \cup \left(\bigcup_{\alpha \in \mathcal{S}_{m+1}} \Omega_\alpha \right)$$

for each $m \geq 0$, and

$$(3.13) \quad \mathbb{R}^4 = (\mathbb{R}^4 \setminus \Omega) \cup \left(\bigcup_{l=0}^{\infty} \bigcup_{\alpha \in \mathcal{S}_l} Y_\alpha \right) \cup F.$$

The basic building block for the construction of the M^j 's is

$$(3.14) \quad \Sigma = \theta \left(\overline{D} \setminus \bigcup_{j=1}^n D_j \right).$$

This is a compact embedded 3-dimensional submanifold of \mathbb{R}^4 with boundary which satisfies $\Sigma \subset Y$ and $\Sigma = P$ on a neighborhood of ∂Y inside Y (by (3.5) and (3.6)). Setting $\Sigma_\alpha = \psi_\alpha(\Sigma)$ for $\alpha \in \mathcal{S}_l$, $l \geq 0$ (so that $\Sigma_\emptyset = \Sigma$ when $l = 0$, as usual), we have that

$$(3.15) \quad \Sigma_\alpha \subset Y_\alpha$$

and

$$(3.16) \quad \Sigma_\alpha = P \text{ on a neighborhood of } \partial Y_\alpha \text{ inside } Y_\alpha.$$

Define M^j for $j \geq 1$ by

$$(3.17) \quad M^j = (P \setminus D) \cup \left(\bigcup_{l=0}^{j-1} \bigcup_{\alpha \in \mathcal{S}_l} \Sigma_\alpha \right) \cup \left(\bigcup_{\alpha \in \mathcal{S}_j} \Omega_\alpha \cap P \right).$$

In accordance with our usual conventions for $l = 0$ we set $M^0 = P$. It is easy to see that this agrees with $M^1 = \theta(P)$ from before, because of the definitions and the properties of θ . All the M^j 's are embedded smooth submanifolds of \mathbb{R}^4 , because of (3.16) and the disjointness properties of the Y_α 's. Also,

$$(3.18) \quad M^j = M^k \text{ outside } \bigcup_{\alpha \in \mathcal{S}_j} \Omega_\alpha, \quad \text{if } j \leq k.$$

This implies in particular that the M^j 's are converging in the Hausdorff topology to

$$(3.19) \quad M = (P \setminus D) \cup \left(\bigcup_{l=0}^{\infty} \bigcup_{\alpha \in \mathcal{S}_l} \Sigma_\alpha \right) \cup F.$$

This is a smooth embedded submanifold away from the Cantor set $F \subseteq P$. It will sometimes be convenient to denote M by M^∞ to make the notation more uniform.

These submanifolds have the self-similarity property that

$$(3.20) \quad \Omega_\alpha \cap M^j = \psi_\alpha(\Omega \cap M^{j-l}) \quad \text{when } \alpha \in \mathcal{S}_l \text{ and } 0 \leq l \leq j.$$

Here we allow $j = \infty$, which gives the relevant property for M . This equality is not hard to derive from the definitions (which were chosen precisely so that this would be true).

These are basically the sets that we are interested in, but for technical reasons we shall need to define another set \widetilde{M} which is not quite as singular as M but which contains a (small) copy of each M^j . To do this let $\{B_k\}_{k=1}^\infty$ be a sequence of balls in \mathbb{R}^4 with disjoint doubles whose radii tend to 0 and whose centers lie on P and converge to the origin, and let $A_k : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ be affine mappings composed of translations and (nonzero) dilations such that $A_k(\widetilde{\Omega}_k) \subseteq B_k$ and $A_k(P) = P$. Set $\widetilde{M}^k = A_k(M^k)$, so that $\widetilde{M}^k = P$ outside B_k (and the interesting part of \widetilde{M}^k is contained in B_k). Let \widetilde{M} be the subset of \mathbb{R}^4 such that $\widetilde{M} = \widetilde{M}^k$ inside each B_k and $\widetilde{M} = P$ outside $\bigcup_k B_k$. (Compare with [Tu1, Example on p. 69] and [HY, Example 6.6].) For convenience we also require that there exist balls B'_k with radius ≤ 100 radius B_k such that $B_l \subseteq B'_k$ when $l \geq k$ but $B_l \cap B'_k = \emptyset$ when $l < k$. This is not hard to arrange, by taking the B_k 's to be $B(z_k, 2^{-2k-3})$ with $|z_k| = 2^{-2k}$ and setting $B'_k = B(0, 2^{-2k+1})$, for instance. For the clarity of future arguments it is best for us to simply require that the B_k 's be chosen in this manner, and with all the z_k 's lying on the same line through the origin, rather than worry about the level of generality in which we can do this construction.

It is clear that the M 's have a lot of self-similarity, but we also need to keep track of their topological properties. Recall the definition of the D_α 's and C_l 's (just before (3.11)), and let G be the decomposition associated to the defining sequence $\{C_l\}$, as discussed just after Definition 2.2. We want to show that M is homeomorphic to \mathbb{R}^3/G , where G is the decomposition associated to our initial package in the manner described after Definition 2.3. We also want to build some parameterizations of the M^j 's which approximate the aforementioned homeomorphism in a nice way.

Lemma 3.21. *There exist diffeomorphisms h_j from \mathbb{R}^4 onto itself, $0 \leq j < \infty$, and a continuous mapping h from \mathbb{R}^4 onto itself with the following properties: $h_j(\mathbb{R}^3) = M^j$ and $h(\mathbb{R}^3) = M$; $h = h_j = \theta$ on $\mathbb{R}^4 \setminus (\bigcup_{i=1}^n \theta^{-1}(\Omega_i))$; $h_j \rightarrow h$ uniformly on \mathbb{R}^4 ; $h(D_\alpha) = M \cap \Omega_\alpha$ for α in any \mathcal{S}_l and $h_j(D_\alpha) = M^j \cap \Omega_\alpha$ for α in any \mathcal{S}_l with $l \leq j$; h*

is constant on each element of the decomposition G , and it induces a homeomorphism from \mathbb{R}^3/G onto M . There is also a homeomorphism \tilde{h} from \mathbb{R}^4 onto itself which maps \mathbb{R}^3 onto \tilde{M} and which equals the identity off of each B_k (and off each $A_k(\overline{\Omega})$), which maps each $B_k \cap P$ onto $B_k \cap \tilde{M}$, and which is a diffeomorphism away from the origin.

Let us define first little copies of θ on the various Ω_α 's. Set

$$(3.22) \quad \theta_\alpha = \psi_\alpha \circ \theta \circ \psi_\alpha^{-1},$$

so that

$$(3.23) \quad \theta_\alpha(\Omega_\alpha) = \Omega_\alpha \text{ and } \theta_\alpha = \text{the identity on } \mathbb{R}^4 \setminus \Omega_\alpha \\ \text{and on a neighborhood of } \partial\Omega_\alpha.$$

If α is the empty sequence in S_0 , then $\theta_\alpha = \theta$. If $\beta = \{b_i\}_{i=1}^{l+1} \in S_{l+1}$ is the child of α with $b_{l+1} = p$, then

$$(3.24) \quad \theta_\alpha = \psi_\beta \circ \phi_p^{-1} \circ \psi_\alpha^{-1} \text{ on a neighborhood of } \psi_\alpha(\overline{D}_p) \text{ in } P$$

by (3.6).

Let g_l denote the composition of all θ_α for $\alpha \in S_l$. Because the Ω_α 's for $\alpha \in S_l$ are pairwise disjoint (by Lemma 3.10), these θ_α 's commute, and so we need not to worry about how we do the composition. Note that $g_0 = \theta$,

$$(3.25) \quad g_l(\Omega_\beta) = \Omega_\beta \quad \text{whenever } \beta \in S_k, k \leq l,$$

and

$$(3.26) \quad g_l = \text{the identity on } \mathbb{R}^4 \setminus \left(\bigcup_{\gamma \in S_k} \Omega_\gamma \right) \quad \text{when } k \leq l.$$

These observations follow easily from Lemma 3.10.

Define h_j for $j \geq 1$ by

$$(3.27) \quad h_j = g_{j-1} \circ g_{j-2} \circ \cdots \circ g_0.$$

These are obviously diffeomorphisms, and $h_1 = g_0 = \theta$. We can take h_0 to be the identity, for completeness. Note that

$$(3.28) \quad h_k \circ h_j^{-1} = \text{the identity on } \mathbb{R}^4 \setminus \left(\bigcup_{\alpha \in S_j} \Omega_\alpha \right)$$

and

$$(3.29) \quad h_k \circ h_j^{-1}(\Omega_\alpha) = \Omega_\alpha, \quad \text{for all } \alpha \in \mathcal{S}_j$$

when $j \leq k$. In particular

$$(3.30) \quad h_j = h_1 \text{ on } \mathbb{R}^4 \setminus \left(\bigcup_{i=1}^n \theta^{-1}(\Omega_i) \right) \text{ when } j > 1.$$

Let us check that

$$(3.31) \quad \begin{aligned} h_j &= \psi_\alpha \circ \phi_\alpha^{-1} \text{ on } \overline{D}_\alpha \text{ (and hence } h_j(D_\alpha) = \Omega_\alpha \cap P) \\ &\text{when } \alpha \in \mathcal{S}_j. \end{aligned}$$

We do this by induction. When $j = 1$ this reduces to (3.6) and the definitions of D_α and Ω_α . Suppose now that we know (3.31) for some value of j and that we want to verify it for $j + 1$. Let $\beta \in \mathcal{S}_{j+1}$ be given, and let $\alpha \in \mathcal{S}_j$ be its parent. Set $p = b_{j+1} \in \{1, \dots, n\}$, where $\beta = \{b_i\}_{i=1}^{j+1}$. From our induction hypothesis we get that $h_j(\overline{D}_\beta) \subseteq h_j(D_\alpha) \subseteq \Omega_\alpha$, and so $h_{j+1} = g_j \circ h_j = \theta_\alpha \circ h_j$ on \overline{D}_β . Our induction hypothesis also gives $h_j(D_\beta) = \psi_\alpha \circ \phi_\alpha^{-1}(D_\beta)$. By definitions we have that $D_\beta = \phi_\beta(D)$ and $\phi_\beta = \phi_\alpha \circ \phi_p$, and so $\phi_\alpha^{-1}(D_\beta) = \phi_\alpha^{-1}(\phi_\beta(D)) = \phi_p(D) = D_p$. Thus $h_j(\overline{D}_\beta) = \psi_\alpha(\overline{D}_p)$, and this permits us to use (3.24) to obtain $h_{j+1} = \psi_\beta \circ \phi_\beta^{-1}$ on \overline{D}_β from our induction hypothesis that (3.31) holds for j and α . This in turn implies that $h_{j+1}(D_\beta) = \psi_\beta(D)$, and this last is the same as $\Omega_\beta \cap P$ because of (3.3), (3.7), and (3.8). This proves (3.31).

Observe that

$$(3.32) \quad h_k(D_\alpha) \subseteq \Omega_\alpha, \quad \text{when } \alpha \in \mathcal{S}_j, j \leq k.$$

This follows from (3.31) and (3.29).

Set $E = \overline{D} \setminus \bigcup_{j=1}^n D_j$ and $E_\alpha = \phi_\alpha(E)$ for any α in any \mathcal{S}_l . Thus E_α is the same as \overline{D}_α with the children of D_α removed (i.e., the D_β 's with $\beta \in \mathcal{S}_{l+1}$ a child of α). As usual we have $E_\emptyset = E$ for the empty sequence in \mathcal{S}_0 . Let us check that

$$(3.33) \quad h_k(E_\alpha) = \Sigma_\alpha, \quad \text{when } \alpha \in \mathcal{S}_j, j < k.$$

It suffices to check this when $k = j + 1$, because $h_k \circ h_{j+1}^{-1}$ is the identity on Σ_α when $\alpha \in \mathcal{S}_j$ and $k \geq j + 1$, because of (3.29) and (3.15).

Because $h_j(E_\alpha) \subseteq h_j(\overline{D}_\alpha) \subseteq \overline{\Omega}_\alpha$ we get that $g_j = \theta_\alpha$ on $h_j(E_\alpha)$, while (3.31) implies that $h_j(E_\alpha) = \psi_\alpha(E)$. Hence $h_{j+1}(E_\alpha) = g_j(h_j(E_\alpha)) = \theta_\alpha(\psi_\alpha(E)) = \psi_\alpha(\theta(E))$ by (3.22). Using (3.14) we get $\psi_\alpha(\theta(E)) = \psi_\alpha(\Sigma) = \Sigma_\alpha$, which proves (3.33).

It is now easy to check that $h_j(D_\alpha) = M^j \cap \Omega_\alpha$ for α in any \mathcal{S}_l with $l \leq j$, using (3.33), (3.31), the definition (3.17) of M^j , and the nesting properties of the Ω_β 's (as in Lemma 3.10). We also have that

$$(3.34) \quad h_k(\mathbb{R}^3 \setminus \left(\bigcup_{\alpha \in \mathcal{S}_k} D_\alpha \right)) \subseteq \mathbb{R}^4 \setminus \left(\bigcup_{\alpha \in \mathcal{S}_k} \Omega_\alpha \right)$$

because of (3.33). This implies that

$$(3.35) \quad h_m = h_k \text{ on } \mathbb{R}^3 \setminus \left(\bigcup_{\alpha \in \mathcal{S}_k} D_\alpha \right), \quad \text{when } m \geq k,$$

because of (3.28).

From (3.28) and (3.29) we have that

$$(3.36) \quad \|h_j - h_k\|_\infty \leq \sup_{\alpha \in \mathcal{S}_j} \text{diam } \Omega_\alpha = \rho^j \text{diam } \Omega \quad \text{when } j \leq k,$$

where ρ is the dilation factor of the ψ_j 's (as in Definition 3.2). This implies that the h_j 's converge uniformly to a continuous mapping $h : \mathbb{R}^4 \rightarrow \mathbb{R}^4$. Notice that

$$(3.37) \quad h = h_j \text{ on } \mathbb{R}^3 \setminus \left(\bigcup_{\alpha \in \mathcal{S}_j} D_\alpha \right),$$

$$(3.38) \quad h = h_1 = \theta \text{ on } \mathbb{R}^4 \setminus \left(\bigcup_{i=1}^n \theta^{-1}(\Omega_i) \right),$$

and

$$(3.39) \quad h(\overline{D}_\alpha) \subseteq \overline{\Omega}_\alpha, \quad \text{for any } \alpha \in \mathcal{S}_j \text{ and any } j,$$

because of (3.35), (3.30), and (3.32).

Let $\{C_l\}$ be the defining sequence associated to the D_α 's, so that $C_l = \bigcup_{\alpha \in \mathcal{S}_l} \overline{D}_\alpha$, and set $C = \bigcap_{l=1}^\infty C_l$. By definition our decomposition G of \mathbb{R}^3 consists of the components of C and the singletons in $\mathbb{R}^3 \setminus C$. We understand h well on $\mathbb{R}^3 \setminus C$, because of (3.37), and we want to understand it on C . We must first analyze the components of C .

Let \mathcal{S} be as before (defined just after (3.11)). Given a sequence $s = \{s_i\}_{i=1}^\infty \in \mathcal{S}$, let A_s denote the intersection of the \overline{D}_α 's which come from ancestors α of s . That is, $\alpha \in \mathcal{S}_j$ is an ancestor of s if $\alpha = \{s_i\}_{i=1}^j$. It is easy to check that the A_s 's are pairwise disjoint, using the nesting properties of the D_α 's.

Sublemma 3.40. $C = \bigcup_{s \in \mathcal{S}} A_s$, and the A_s 's are the connected components of C .

C contains all the A_s 's by definition. Conversely, if $p \in C$, then for each j there is an $\alpha_j \in \mathcal{S}_j$ such that $p \in \overline{D}_{\alpha_j}$. The analogue of Lemma 3.10 for the D 's instead of the Ω 's implies that α_{j+1} must be a child of α_j for each j . This means that the α_j 's combine to form a sequence $s \in \mathcal{S}$, and it is clear that $p \in A_s$.

Each A_s is connected. For this we use the fact that each D_α is connected, since our original domain D is (by the definition of domain). Thus A_s is the decreasing intersection of compact connected sets, and it is an elementary general fact that A_s must itself be connected under these conditions. (If A_s were disconnected, it would be contained in the union of two disjoint open sets, each of which intersects it, and the same would then have to be true for some \overline{D}_α .)

Each A_s is a component of C . If not, there would be an A_s and an A_t which lie in the same component of C , with $s \neq t$. Let $\alpha \in \mathcal{S}_j$ be the common ancestor of s and t with j as large as possible (but perhaps equal to 0). Then C is contained in the union of the complement of D_α and the closures of the children of D_α , and A_s and A_t lie in different children of D_α . Since the closures of these children are disjoint and lie in D_α we see that A_s and A_t cannot both touch the same component of C . This proves Sublemma 3.40.

Recall the (bijective) mapping $f : \mathcal{S} \rightarrow F$ defined just after the definition of the Cantor set F in (3.11). We have that

$$(3.41) \quad h(p) = f(s), \quad \text{for all } p \in A_s \text{ and } s \in \mathcal{S},$$

because of (3.39) and the definitions of A_s and $f(s)$. It is easy to check that

$$(3.42) \quad h(D_\alpha) = M \cap \Omega_\alpha, \quad \text{for any } \alpha \text{ in any } \mathcal{S}_l,$$

because of (3.19), (3.37), (3.18), (3.41), and the corresponding statement for the h_j 's (just before (3.34)). In other words, h looks like the

h_j 's away from C , M looks like the M^j 's away from F , and C and F correspond under h the way that they should because of (3.41).

It is not hard to check that h induces a homeomorphism from \mathbb{R}^3/G onto M . Sublemma 3.40 provides us with a complete understanding of the decomposition G , we know how h behaves on and near the nontrivial elements of G because of (3.41) and (3.42), and we understand h away from C because of (3.37).

The last part of Lemma 3.21, about \tilde{h} , is an easy consequence of the earlier part. That is, we set $\tilde{h} =$ the identity off the $A_k(\Omega)$'s and $\tilde{h} = A_k \circ h_k \circ A_k^{-1}$ on each $A_k(\Omega)$. Because each h_k equals the identity off Ω and maps Ω onto itself it is easy to see that this is a homeomorphism, and it maps \mathbb{R}^3 onto \tilde{M} by definition. It is also clear that it is a diffeomorphism away from the origin.

This completes the proof of Lemma 3.21.

RECAPITULATION 3.43. We started with an excellent package, which contained an initial package. From the initial package we can generate a decomposition of \mathbb{R}^3 as in Section 2. The excellent package is, roughly speaking, a topologically equivalent version of the initial package which lives in \mathbb{R}^4 and has better geometric properties (like self-similarity). We have now used the excellent package to construct sets M , M^j , and \tilde{M} , and Lemma 3.21 tells us that they have the correct compatibility with the decomposition associated to the initial package. To prove Theorem 1.12 we shall choose specific excellent packages (in Sections 4, 5, and 6) and use the resulting sets M , \tilde{M} .

Before we consider specific examples we shall establish some general properties of the construction which will be relevant for all of the examples. For this we shall need a case-by-case analysis of the positions of balls centered on these sets whose basic structure will be used repeatedly in this paper, and so we establish it first. For the moment we restrict ourselves to M and M^j and forget about \tilde{M} .

Lemma 3.44. *Let $x \in M^j$ and $r > 0$ be given, where we allow $j = \infty$ in which case $M^j = M$. Let $a > 0$ be a small number that we get to choose but which should be small enough so that $2a \text{diam } \Omega < \text{dist}(\Omega_p, \Omega_q)$ for all $p, q = 1, \dots, n, p \neq q$. Then one of the following alternatives holds:*

- i) $B(x, r) \cap \Omega = \emptyset$,
- ii) $B(x, r) \cap \Omega \neq \emptyset$ and $r \geq a \text{diam } \Omega$,

iii) $B(x, r) \cap \Omega \neq \emptyset$, $r < a \operatorname{diam} \Omega$, and $B(x, r) \not\subseteq \Omega$,

iv) there is an α in some \mathcal{S}_l , $0 < l \leq j$, such that $B(x, r) \subseteq \Omega_\alpha$ and $r \geq a \operatorname{diam} \Omega_\alpha$,

v) there is an α in some \mathcal{S}_l , $0 \leq l \leq j$, such that $B(x, r) \subseteq \Omega_\alpha$, $r \leq a \operatorname{diam} \Omega_\alpha$, and either $l = j$ or $B(x, r) \cap Y_\alpha \neq \emptyset$.

Recall that Y_α was defined just before (3.12), and note that $l = j$ is not an option when $j = \infty$.

Lemma 3.44 is quite straightforward. If none of the first three cases hold then we have $r < a \operatorname{diam} \Omega$ and $B(x, r) \subset \Omega$. Choose $l \leq j$ as large as possible so that $B(x, r) \subset \Omega_\alpha$ for some $\alpha \in \mathcal{S}_l$, where $l = 0$ is allowed. If also $r \geq a \operatorname{diam} \Omega_\alpha$, then iv) obtains (unless $l = 0$, in which case ii) was already satisfied). If $r \leq a \operatorname{diam} \Omega_\alpha$, then v) has to hold because we took l to be as large as possible. That is, if $l < j$ and $B(x, r) \cap Y_\alpha = \emptyset$, then we could replace Ω_α by one of its children. For this last step we need to know that a is sufficiently small so that $B(x, r)$ cannot touch two different children of Ω_α , and the condition on a in the lemma ensures precisely this.

Let us now establish the Ahlfors regularity of these sets.

Lemma 3.45. *If $\rho^3 n < 1$, where ρ, n are as in Definition 3.2, then the sets M , M^j , and \widetilde{M} are all regular with dimension 3, and with a constant that is bounded independently of j .*

Notice first that M , M^j , and \widetilde{M} are all closed.

We should begin with some preliminary facts. The first is that

$$(3.46) \quad H^3(F) = 0.$$

This follows from the definition (3.11) of F , which implies that $F \subseteq \bigcup_{\alpha \in \mathcal{S}_l} \overline{\Omega}_\alpha$ for each l , so that

$$(3.47) \quad H^3(F) \leq \limsup_{l \rightarrow \infty} \sum_{\alpha \in \mathcal{S}_l} (\operatorname{diam} \overline{\Omega}_\alpha)^3 \leq \limsup_{l \rightarrow \infty} n^l \rho^{3l} (\operatorname{diam} \Omega)^3$$

by definition of Hausdorff measure. (Do not forget that \mathcal{S}_l has n^l elements.) This implies (3.46), since we are assuming that $\rho^3 n < 1$. Notice that a similar argument implies that

$$(3.48) \quad \text{Hausdorff dimension}(F) \leq d, \quad \text{if } \rho^d n < 1.$$

In our examples we shall have the freedom to choose ρ to be as small as we wish, and so we can make F have Hausdorff dimension as small as we wish.

Next we check that

$$(3.49) \quad H^3(\Omega \cap M^j) \leq C < \infty,$$

for some constant C which does not depend on j . In view of (3.46) and the definitions (3.17) and (3.19) of M^j and M we are reduced to estimating

$$(3.50) \quad \sum_{l=0}^{j-1} \sum_{\alpha \in \mathcal{S}_l} H^3(\Sigma_\alpha) + \sum_{\alpha \in \mathcal{S}_j} H^3(\Omega_\alpha \cap P).$$

Since Hausdorff measure behaves properly under similarities we have that $H^3(\Sigma_\alpha) = \rho^{3l} H^3(\Sigma)$ when $\alpha \in \mathcal{S}_l$ and $H^3(\Omega_\alpha \cap P) = \rho^{3j} H^3(\Omega \cap P)$ when $\alpha \in \mathcal{S}_j$, and so (3.50) reduces to

$$(3.51) \quad \sum_{l=0}^{j-1} n^l \rho^{3l} H^3(\Sigma) + n^j \rho^{3j} H^3(\Omega \cap P).$$

The desired bound follows from our assumption that $\rho^3 n < 1$, since $H^3(\Sigma)$ and $H^3(\Omega \cap P)$ are finite.

Notice also that each Σ_α satisfies the ‘‘compact’’ version of regularity, namely that

$$(3.52) \quad C_0^{-1} s^3 \leq H^3(\Sigma_\alpha \cap B(y, s)) \leq C_0 s^3,$$

for some constant C_0 (which does not depend on α) and all $y \in \Sigma_\alpha$ and $0 < s < \text{diam } \Sigma_\alpha$. For Σ itself (3.52) is a consequence of its smoothness, while for the general Σ_α (3.52) reduces to the case of Σ because everything behaves properly under similarities. The same reasoning implies that

$$(3.53) \quad C_0^{-1} s^3 \leq H^3(\Omega_\alpha \cap P \cap B(y, s)) \leq C_0 s^3,$$

for some constant C_0 (which does not depend on α) and all $y \in \Omega_\alpha$ and $0 < s < \text{diam } \Omega_\alpha$.

To prove Lemma 3.45 let us deal first with M and M^j , and let us take $M = M^\infty$ as before. Let $x \in M^j$ and $r > 0$ be given, so that we want to show that

$$(3.54) \quad C^{-1} r^3 \leq H^3(M^j \cap B(x, r)) \leq C r^3,$$

for some constant C which does not depend on x , r , or j . Fix an $a > 0$ which is small enough so that Lemma 3.44 can be applied, and also so that

$$(3.55) \quad 2a \operatorname{diam} \Omega < \rho \operatorname{dist}(\partial\Omega, \Omega_p) < \operatorname{dist}(\partial\Omega, \Omega_p),$$

for $p = 1, \dots, n$. Lemma 3.44 provides us with five alternatives to consider separately. In case i) we have that $M^j \cap B(x, r) = P \cap B(x, r)$, and in particular that $x \in P$, and (3.54) follows.

Now suppose that ii) holds. In this case the upper bound in (3.54) is automatic, because of the corresponding bound for P and (3.49). The lower bound is slightly a nuisance, but it is not deep. Let us first check that

$$(3.56) \quad H^3((P \setminus \Omega) \cap B(x, r)) \geq C^{-1} r^3,$$

for some constant C when $x \in P \setminus \Omega$. When $r \leq 10 \operatorname{diam} \Omega$ this follows from the smoothness of Ω , while for $r \geq 10 \operatorname{diam} \Omega$ it holds also for $x \in \Omega$, because $(P \setminus \Omega) \cap B(x, r)$ must then contain the intersection of P with a ball of radius $r/3$. These observations imply that in order to establish the lower bound in (3.54) we may as well assume that $x \in \Omega$ and $r \leq 10 \operatorname{diam} \Omega$. Since we are already assuming in ii) that $r \geq a \operatorname{diam} \Omega$, we have that r is bounded and bounded from below. Let m be the smallest positive integer such that $\rho^{m+1} < a/2$, and notice that m does not depend on x , r , or j . The main point in the rest of the argument is that the parts of M^j which correspond to levels above m do not really matter. If $j \leq m$, then either $x \in \Sigma_\alpha$ for an $\alpha \in \mathcal{S}_k$ with $k < j$ or $x \in \Omega_\alpha \cap P$ for some $\alpha \in \mathcal{S}_j$ (see (3.17)). In this case we can derive the desired lower bound in (3.54) from (3.52) or (3.53) (applied with $s =$ a small (but not too small) multiple of r). Thus we may assume that $j > m$, so that $\Sigma_\beta \subseteq M^j$ whenever $\beta \in \mathcal{S}_l$, $l \leq m$. If $x \in \Sigma_\alpha$ for an $\alpha \in \mathcal{S}_k$ with $k \leq m$, then we can again derive the lower bound in (3.54) from (3.52). We are left with the case where $x \in \Omega_\alpha$ for some $\alpha \in \mathcal{S}_{m+1}$, because of the nesting properties of the Ω_γ 's (as in Lemma 3.10). Our choice of m implies that $\operatorname{diam} \Omega_\alpha < r/2$, and hence there is a point y in $B(x, r/2)$ which lies on (the boundary of) Σ_β , where $\beta \in \mathcal{S}_m$ is the parent of α . This permits us to reduce to (3.52) again since $H^3(M^j \cap B(x, r)) \geq H^3(\Sigma_\beta \cap B(y, r/2))$. This establishes (3.54) when ii) holds.

When iii) holds, our assumption (3.55) on a implies that $B(x, r)$ cannot touch the Ω_p 's, $1 \leq p \leq n$, so that $M^j \cap B(x, r) = ((P \setminus \Omega) \cup$

$\Sigma) \cap B(x, r)$. In this case (3.54) is an immediate consequence of the smoothness of Σ and D .

If iv) holds, then we can reduce to $l = 0$ (and $\Omega_\alpha = \Omega$) using the self-similarity property (3.20), and this is just a special case of ii).

We are left with v). As before we can use (3.20) to reduce to $l = 0$ and $\Omega_\alpha = \Omega$. If $j = 0$ then $M^j = P$ and (3.54) is immediate. If $j = 1$ then $M^j = M^1$ is smooth and (3.54) is again clear. Thus we may assume that $j > 1$. The key observation now is that

$$(3.57) \quad B(x, r) \cap \Omega_\gamma = \emptyset, \quad \text{for all } \gamma \in \mathcal{S}_2 .$$

To see this first notice that (3.55) implies that

$$(3.58) \quad 2a \operatorname{diam} \Omega < \operatorname{dist}(\partial\Omega_\beta, \Omega_\gamma)$$

whenever $\beta \in \mathcal{S}_1$ and $\gamma \in \mathcal{S}_2$ is a child of β . (Do not forget that ρ is the dilation factor of the similarity ψ_β , and do not forget (3.8) either.) Our assumption v) implies that $B(x, r)$ intersects Y (defined just before (3.12)), and so if $B(x, r)$ intersected some Ω_γ , $\gamma \in \mathcal{S}_2$, then (3.58) would not be true, since we have also that $r < a \operatorname{diam} \Omega$ from v). Thus (3.57) is true, which implies that

$$(3.59) \quad B(x, r) \cap M^j = B(x, r) \cap \left(\Sigma \cup \left(\bigcup_{\beta \in \mathcal{S}_1} \Sigma_\beta \right) \right)$$

since we have also that $B(x, r) \subseteq \Omega$ by v). The regularity estimate (3.54) follows easily (from (3.52), for instance).

This proves that the M^j 's are regular with a uniformly bounded constant. It remains to deal with \widetilde{M} . Let $x \in \widetilde{M}$ and $r > 0$ be given, and recall the definition of \widetilde{M} and the related notation (just after (3.20)). If $B(x, r)$ is disjoint from all the B_k 's, then $B(x, r) \cap \widetilde{M} = B(x, r) \cap P$, and we are in business. If $B(x, r) \subseteq 2B_k$ for some k , then $x \in A_k(M^k)$, $B(x, r) \cap \widetilde{M} = B(x, r) \cap A_k(M^k)$, and the necessary estimates on $H^3(B(x, r) \cap \widetilde{M})$ follow from the regularity of M^k . From this case we get that

$$(3.60) \quad H^3(\widetilde{M} \cap B_k) \leq C (\operatorname{radius} B_k)^3 ,$$

for all k . Now suppose that $B(x, r)$ intersects some B_k 's but that it is not contained in any $2B_k$. Let K denote the set of k 's such that $B(x, r)$ intersects B_k . Since $B(x, r) \not\subseteq 2B_k$ for any k we have that

$2r \geq \text{radius } B_k$ when $k \in K$, and so $\cup_{k \in K} B_k \subseteq B(x, 5r)$. Because the B_k 's are all disjoint and centered on P we obtain from this that

$$(3.61) \quad \sum_{k \in K} (\text{radius } B_k)^3 \leq (5r)^3$$

since $H^3((\cup_{k \in K} B_k) \cap P) \leq H^3(B(x, 5r) \cap P)$. Because (3.60) implies that

$$(3.62) \quad \begin{aligned} H^3(B(x, r) \cap \widetilde{M}) &\leq H^3(B(x, r) \cap P) \\ &+ \sum_{k \in K} H^3(B(x, r) \cap \widetilde{M} \cap B_k) \leq C r^3, \end{aligned}$$

for some constant C , we get the upper bound that we need from (3.61). For the lower bound we observe that

$$(3.63) \quad H^3(B(x, r) \cap P \cap (2B_k \setminus B_k)) \geq C^{-1} (\text{radius } B_k)^3$$

when $k \in K$, because $B(x, r)$ intersects such a B_k but is not contained in $2B_k$. Hence

$$(3.64) \quad \begin{aligned} H^3(B(x, r) \cap \widetilde{M}) &\geq H^3\left((B(x, r) \cap P) \setminus \left(\bigcup_{k \in K} B_k\right)\right) \\ &\geq C^{-1} H^3(B(x, r) \cap P), \end{aligned}$$

for some constant C . This uses also the disjointness of the $2B_k$'s. Thus we get the lower bound on $H^3(B(x, r) \cap \widetilde{M})$ that we needed, so that \widetilde{M} is also regular.

This completes the proof of Lemma 3.45. The same basic structure of the argument will be used repeatedly in this paper. That is, we shall need to prove various properties about balls on M or on an M^j , and we shall use Lemma 3.44 to distinguish some cases. Case i) will always be trivial, and iii) and v) will typically be easy because of the smoothness of M and the M^j 's away from the singular set F . Cases ii) and iv) are generally about the same as each other and often require more specific information about the excellent package. To simplify these future arguments we collect first some information in the following lemma that will be common to many of them.

Lemma 3.65. *Consider M^j , $j \leq \infty$, with $j = \infty$ corresponding to M . There exist a small constant a and a large constant C_0 (depending on*

the excellent package but not on j) so that if $x \in M^j$ and $r > 0$ satisfy i), iii), or v) in Lemma 3.44, or if they satisfy ii) and also $j < \infty$, then there is an open set W in \mathbb{R}^4 with $B(x, r) \subseteq W \subseteq B(x, C_0 r)$ and a diffeomorphism λ from $B(0, r)$ onto W such that $\lambda(P \cap B(0, r)) = M^j \cap W$. In the cases i), iii), and v) (but not ii)) we can also take λ to be C_0 -bilipschitz, i.e.,

$$(3.66) \quad C_0^{-1} |y - z| \leq |\lambda(y) - \lambda(z)| \leq C_0 |y - z|,$$

for all $y, z \in B(0, r)$.

In other words, M^j can be flattened out in a nice way near $B(x, r)$. Note the scale-invariance of the bilipschitz condition (3.66). This is very important when we are working with small balls near the singular set, because we shall not have uniform bounds on higher derivatives of λ .

Let us prove Lemma 3.65. Let a be small, to be chosen soon, and let $x \in M^j$ and $r > 0$ be given. If (x, r) satisfies i) in Lemma 3.44, then $B(x, r) \cap M^j = B(x, r) \cap P$, and we can simply take $W = B(x, r)$ and λ to be a translation. If ii) holds, so that $B(x, r) \cap \Omega \neq \emptyset$ and $r \geq a \operatorname{diam} \Omega$, then $B(x, Cr) \supseteq \Omega$ for $C = 1 + a^{-1}$. In this case $B(x, Cr) \cap M^j = h_j(B(x, Cr) \cap P)$, where h_j is the diffeomorphism (for $j < \infty$) promised in Lemma 3.21, and so we can simply take $W = h_j(B(x, Cr))$ and λ to be a translation of h_j . If iii) holds, then $B(x, r)$ must intersect $\partial\Omega$. If a is small enough, then $B(x, r) \cap M^j = B(x, r) \cap ((P \setminus \Omega) \cup \Sigma)$. The right hand side is a smooth submanifold with boundary, and $B(x, r)$ stays away from the boundary when a is sufficiently small. In this case M^j is a small smooth perturbation of a 3-plane inside $B(x, r)$, and it is easy to get the desired W and λ (with $C_0 = 2$, for instance). Before we deal with v) let us introduce some auxiliary notation and definitions. We may as well assume that $j \geq 2$, because the $j = 0$ case is trivial and we can simply use the diffeomorphism h_1 provided by Lemma 3.21 to get Lemma 3.65 when $j = 1$.

Given $\beta \in \mathcal{S}_k$, $k \leq j$, set $\Sigma'_\beta = \Sigma_\beta$ when $k < j$ and $\Sigma'_\beta = \overline{\Omega}_\beta \cap P$ when $k = j$. This is just a convenient way to allow for the slightly exceptional case where $k = j < \infty$ without having make additional statements. Define N_α for α in some \mathcal{S}_l , $l \leq j$, in the following manner. If $0 < l \leq j - 1$ set

$$(3.67) \quad N_\alpha = \Sigma'_\delta \cup \Sigma'_\alpha \cup \left(\bigcup \Sigma'_\beta \right),$$

where δ is the parent of α in \mathcal{S}_{l-1} and the union is taken over the β 's in \mathcal{S}_{l+1} which are children of α . If $l = j$ define N_α in the same way

except that we drop the Σ'_β 's (since the children of α do not matter). If $l = 0$ then replace Σ'_δ in (3.67) with $P \setminus D$ but keep the rest. In each case N_α is a smooth embedded submanifold of \mathbb{R}^4 which contains Σ'_α , with some room to spare. Specifically, there is a constant $C > 0$ so that

$$(3.68) \quad \{z \in M^j : \text{dist}(z, \Sigma'_\alpha) < C^{-1} \text{diam } \Omega_\alpha\} \subseteq N_\alpha,$$

and in fact the left hand side does not get too close to the boundary of N_α . This constant C does not depend on α , j , or l ; this is easy to check, using the usual self-similarity argument based on (3.20) to reduce to the (finitely many) $l = 0, 1$ cases.

Now suppose that v) in Lemma 3.44 is true, and let α, l be as in v). Remember from v) that either $l = j$ or $B(x, r)$ intersects Y_α . If a is small enough then we have that $B(x, r) \cap M^j = B(x, r) \cap N_\alpha$, because of (3.68), and $B(x, r)$ stays away from the boundary of N_α . In this case we can get the desired W and λ as soon as a is small enough, for the same reasons of smoothness as in case iii). We can even get uniform estimates (which do not depend on α or l) because the self-similarity provided by (3.20) permits us to reduce the problem to a finite number of models for the N_α 's. This proves Lemma 3.65.

Next we deal with the property about bilipschitz balls in the conclusion of Theorem 1.12. This will occupy us for the remainder of the section, and the reader may wish to skip the long argument for the time being. Let us assume that

$$(3.69) \quad P \setminus D \text{ and } \overline{D} \setminus \bigcup D_i \text{ are connected.}$$

This assumption will hold in all of our examples.

Proposition 3.70. *Assuming that our initial package satisfies (3.69), there is a constant L so that if $E = M, \widetilde{M}$, or M^j for some $j < \infty$, then every pair of distinct points $p, q \in E$ is contained in a closed subset W of E such that W is L -bilipschitz equivalent to a closed Euclidean 3-ball. In particular there is a curve in E which joins them and which has length less or equal than $L^2 |p - q|$.*

This proposition (which is a variant of [Se5, Proposition 4.25]) is basically trivial, but it takes some space to do it with a moderate amount of care. Suppose for simplicity that we are working with M rather than M^j or \widetilde{M} . The first step is to connect p and q by a nice

curve. “Nice” means in particular that the curve should avoid the singularities of M as much as possible. It may be necessary for the curve to do a fair amount of looping near p and q , because of the twisting of M , but we can understand this in a clear and simple way because we understand the singularities of M so well (by construction). The second step is to “fatten up” this curve to get a subset of M which is bilipschitz equivalent to a 3-ball. The amount of fattening is allowed to degenerate linearly as we move toward the endpoints of the curve, because we only want a bilipschitz ball (as opposed to a smooth ball). By choosing the curve to move away from the singularities of M as fast as possible we shall have that M is very flat near points on the curve (at the appropriate scale), and this will allow us to fatten the curve sufficiently.

The proof will show that we can choose W so that $W \setminus \{p, q\}$ is contained in the smooth part of E (i.e., it is disjoint from F when $E = M$ and it does not contain the origin when $E = \widetilde{M}$), and that the bilipschitz equivalence between W and a ball can be taken to be smooth away from p and q . This observation will be useful in Section 10.

Much of the structure of an excellent package is unnecessary for Proposition 3.70, in the same way as discussed just after Definition 3.2. In particular the existence of these bilipschitz balls involves only “internal” properties of the M 's, and not their relationship with the ambient space.

Let us now begin the proof of Proposition 3.70. Suppose first that $E = M^j$, where $j = \infty$ is allowed, and let $p, q \in M^j$ be given, $p \neq q$. We may as well assume that $j \geq 2$, since $M^0 = P$ and M^1 is bilipschitz equivalent to P (via the diffeomorphism h_1 from Lemma 3.21).

Given u, v in P and $\varepsilon > 0$ let $S(u, v)$ denote the segment which connects u to v and let $S_\varepsilon(u, v)$ be the set of points x in P such that $\text{dist}(x, S(u, v)) \leq \varepsilon \text{dist}(x, \{u, v\})$. Thus $S_\varepsilon(u, v)$ is the union of two truncated cones, one with vertex u , the other with vertex v . It is also bilipschitz equivalent to a Euclidean 3-ball, with a bilipschitz constant which depends only on ε . In order to produce a set W as in the proposition it is better to think of W as being bilipschitz equivalent to some $S_\varepsilon(u, v)$ rather than a round ball. Typically W will look like a twisted version of $S_\varepsilon(u, v)$ which spirals around at the ends. (To get the smoothness we want one should smooth out the spherical “corner” in the middle of $S_\varepsilon(u, v)$, but that is easy to do.)

It will be more convenient in the proof to use a slightly different analysis of cases than the one in Lemma 3.44. We begin with an easy one.

Lemma 3.71. *If $p, q \in M^j$ satisfy $p, q \in N_\alpha$ for some $\alpha \in \mathcal{S}_l$, $l \leq j$, then the conclusions of Proposition 3.70 hold for M^j with this choice of p and q . Here N_α is as in (3.67).*

To see this, forget about all this specific notation for a moment, and let N be a compact connected smooth 3-dimensional embedded submanifold of \mathbb{R}^4 (with boundary). Then any pair of points in N are contained in a closed subset of N which is bilipschitz equivalent to a closed Euclidean 3-ball. This is not hard to prove, and we leave it as an exercise. Lemma 3.71 is a special case of this statement, at least when $l > 0$, modulo the issue of getting uniform bounds on the bilipschitz constants. These uniform bounds come from the self-similarity property (3.20), which ensures that the N_α 's are all similar to a finite collection of models. The argument for $l = 0$ is similar but modifications are needed because N_α is now unbounded (but equal to P outside Ω). In bounded regions this case behaves in the same way as the previous one, but in unbounded regions it behaves like the corresponding question for $P \setminus B$, where B is some ball. The point is simply that one must sometimes be careful to choose W so that it avoids the hole. Again the details are left to the reader.

The next lemma covers the most interesting case for Proposition 3.70.

Lemma 3.72. *The conclusions of Proposition 3.70 hold for M^j when $j = \infty$ and $p, q \in F$.*

Choose $\delta \in \mathcal{S}_m$ such that $p, q \in \Omega_\delta$ and m is as large as possible. Let α_l and β_l be the unique elements of \mathcal{S}_l such that $p \in \Omega_{\alpha_l}$ and $q \in \Omega_{\beta_l}$. Thus $\alpha_m = \beta_m = \delta$, and α_{m+1} and β_{m+1} are both children of δ , but they are distinct children, since m is maximal. Choose points $p_l \in \partial\Sigma_{\alpha_l}$ and $q_l \in \partial\Sigma_{\beta_l}$ in an arbitrary manner. There is a constant C which depends only on our excellent package so that

$$(3.73) \quad C^{-1} \text{diam } \Omega_{\alpha_l} \leq |p_l - p_{l+1}| \leq C \text{diam } \Omega_{\alpha_l} ,$$

$$(3.74) \quad C^{-1} \text{diam } \Omega_{\beta_l} \leq |q_l - q_{l+1}| \leq C \text{diam } \Omega_{\beta_l} ,$$

$$(3.75) \quad C^{-1} \operatorname{diam} \Omega_\delta \leq |p_{m+1} - q_{m+1}| \leq C \operatorname{diam} \Omega_\delta ,$$

$$(3.76) \quad C^{-1} \operatorname{diam} \Omega_\delta \leq |p - q| \leq C \operatorname{diam} \Omega_\delta .$$

Note that $p_l \rightarrow p$ and $q_l \rightarrow q$ as $l \rightarrow \infty$. We shall build our bilipschitz ball W by combining a family of smooth tubes which connect the successive p_l 's and q_l 's.

Remember that $\operatorname{diam} \Omega_\alpha = \rho^l \operatorname{diam} \Omega$ when $\alpha \in S_l$. Thus $\{|p_l - p_{l+1}|\}_l$ and $\{|q_l - q_{l+1}|\}_l$ are approximately geometric sequences.

In order to prove Lemma 3.72 it suffices to find $\varepsilon > 0$ and a bilipschitz mapping $f : S_\varepsilon(z, w) \rightarrow M$ (with uniform choices of ε and the bilipschitz constant) such that $f(z) = p$ and $f(w) = q$. We shall define f in stages. To understand how f is constructed it is helpful to visualize the region $f(S_\varepsilon(z, w))$ that we shall have to construct. It will be a union of 3-dimensional tubes in M , where the tubes connect the successive p_l 's and q_l 's. These tubes will be diffeomorphic to rectangles and they will be neither too thin nor too close to F . To build these tubes we shall first choose some smooth Jordan arcs which connect the successive p_l 's and q_l 's, and the tubes will be little tubular neighborhoods of these arcs. The next sublemma deals with the existence of these Jordan arcs.

Sublemma 3.77. *Given any α in any S_l and any pair of points a, b in different components of the boundary of Σ_α , we can find an arc $\gamma \subseteq \Sigma_\alpha$ which connects a to b and has the following properties: if u and v are two points on γ , then the length of the arc in γ which connects u to v is bounded by $C|u - v|$; inside $B(a, C^{-1} \operatorname{diam} \Omega_\alpha)$ the curve γ agrees with the line segment in P which emanates from a , is orthogonal to $\partial \Sigma_\alpha$ at a , and goes inside Σ_α , and similarly for b (remember that the part of Σ_α near its boundary lies in P); if $u \in \gamma$, then $\operatorname{dist}(u, \partial \Sigma_\alpha) \geq C^{-1} \operatorname{dist}(u, \{a, b\})$ (so that γ does not get close to the boundary except near the endpoints); for each positive integer i the Euclidean norm of the i^{th} derivative of the arclength parameterization of γ is bounded by $C(i) (\operatorname{diam} \Omega_\alpha)^{1-i}$. (Notice that this estimate is scale-invariant.) These constants C and $C(i)$ depend only on our excellent package.*

This is an easy exercise. The main points are that we can reduce to the case where $l = 0$ and $\Sigma_\alpha = \Sigma$ using the self-similarity principle (3.20), and that Σ is a smooth connected (by (3.69)) compact manifold with boundary which agrees with P near its boundary. Thus we can certainly connect any pair of points in Σ with a curve, but by being a little

bit careful we can choose the curve so that it avoids self-intersections and the boundary of Σ , we can make it smooth, etc.

Next we connect a sequence of curves as provided by Sublemma 3.77.

Sublemma 3.78. *There exist points z and w in P and a bilipschitz map $f : S(z, w) \rightarrow M$ with uniformly bounded bilipschitz constant which satisfy the following properties: $f(z) = p$ and $f(w) = q$;*

$$(3.79) \quad C^{-1} \operatorname{dist}(t, \{z, w\}) \leq \operatorname{dist}(f(t), F) \leq C \operatorname{dist}(t, \{z, w\}),$$

for all $t \in S(z, w)$ (so that the image of f avoids the singular set F as much as possible); f is smooth away from the endpoints z and w , and if $f^{(i)}$ denotes the i^{th} order derivative of f on $S(z, w)$, $f' = f^{(1)}$, then $|f'(t)| = 1$ and

$$(3.80) \quad |f^{(i)}(t)| \leq C(i) \operatorname{dist}(t, \{z, w\})^{-i+1},$$

for all $t \in S(z, w) \setminus \{z, w\}$ and $i \geq 1$. These constants depend only on our excellent package (and i in the case of $C(i)$).

Given $l > m$ let γ_l be the curve provided by Sublemma 3.77 for $\alpha = \alpha_l$, $a = p_l$, and $b = p_{l+1}$. Thus the length of γ_l is comparable to $|p_l - p_{l+1}|$, which is controlled by (3.73). Similarly, let $\tilde{\gamma}_l$ be the curve which corresponds to β_l, q_l , and q_{l+1} , and let γ be the curve which corresponds to δ, p_{m+1} , and q_{m+1} . We shall choose f so that its image is precisely the union of all these curves together with p and q .

Notice first that the sum of the lengths of all the γ 's is controlled by a convergent geometric series and hence is finite. In fact this total length is comparable to $|p - q|$, because of (3.76). Choose $z, w \in P$ so that $|z - w|$ equals the sum of the lengths of the γ 's. (Except for this z and w can be arbitrary.) Set $f(z) = p$ and $f(w) = q$, and define f on $S(z, w)$ in such a way that it is really just the concatenation of the arclength parameterizations of the γ 's, ordered in the obvious way. This mapping is smooth on $S(z, w) \setminus \{z, w\}$, because of the properties of the γ 's in Sublemma 3.77.

Each arc γ_l corresponds to a segment I_l in $S(z, w)$, and the length of I_l and its distance to z are both comparable to $|p_l - p_{l+1}|$. Similar statements apply to the other γ 's. Using these observations it is easy to

derive (3.80) from the corresponding property in Sublemma 3.77, while (3.79) uses also the fact that

$$(3.81) \quad C^{-1} \text{diam } \Omega_\alpha \leq \text{dist}(x, F) \leq C \text{diam } \Omega_\alpha$$

when $x \in \Sigma_\alpha$.

The Lipschitzness of f and the fact that $|f'| = 1$ on $S(z, w) \setminus \{z, w\}$ follow from the fact that we are using the arclength parameterizations of the γ 's. It remains to show that f is bilipschitz, i. e., that $|f(u) - f(v)| \geq C^{-1}|u - v|$ whenever $u, v \in S(z, w)$. If $f(u)$ and $f(v)$ lie in the same arc among the γ 's then this follows immediately from the chord-arc property in Sublemma 3.77. If $f(u)$ and $f(v)$ lie in adjacent arcs among the γ 's then this estimate can also be derived from the properties of these arcs given in Sublemma 3.77. If they do not lie in adjacent γ 's - in particular if one of u or v equals z or w - then we use the fact that we know where the arcs lie in terms of the Σ_α 's together with the various nesting and separation properties of the Y_α 's and Ω_α 's. For instance, we know that $\Sigma_{\alpha_l} \subseteq \Omega_{\alpha_{m+1}}$ and $\Sigma_{\beta_l} \subseteq \Omega_{\beta_{m+1}}$ when $l > m$, and we know that $\text{dist}(\Omega_{\alpha_{m+1}}, \Omega_{\beta_{m+1}}) \geq C^{-1} \text{diam } \Omega_\delta$ since $\alpha_{m+1} \neq \beta_{m+1}$ by the maximality of m . This implies that $|f(u) - f(v)| \geq C^{-1}|p - q| \geq C^{-1}|u - v|$ when $f(u)$ lies in one of the γ_l 's and $f(v)$ lies in one of the $\tilde{\gamma}_k$'s. Similarly, $\text{dist}(\Sigma_{\alpha_l}, \Omega_{\alpha_k}) \geq C^{-1} \text{diam } \Omega_{\alpha_l}$ when $k > l + 1$, as one can check by reducing to the case where $k = l + 2$ and $l = 0$, using the nesting properties in Lemma 3.10 and self-similarity (3.20). This implies that $|f(u) - f(v)| \geq C^{-1} \text{diam } \Omega_{\alpha_l}$ if $f(u)$ lies in γ_l and $f(v)$ lies in γ_k with $k > l + 1$. In this case we have that $|u - v| \leq C \text{diam } \Omega_{\alpha_l}$ and hence $|f(u) - f(v)| \geq C^{-1}|u - v|$. The other cases can be handled in the same way, and we obtain that f is indeed bilipschitz. This proves Sublemma 3.78.

From now on we assume that f is defined on $S(z, w)$ as in Sublemma 3.78. We want to extend f to $S_\varepsilon(z, w)$ for suitable $\varepsilon > 0$. To do this we analyze the unit normals to M along the image of f , then we shall determine a first-order (linear) approximation to this extension of f , and then we build a true extension of f .

Sublemma 3.82. *There is a smooth function $\nu : S(z, w) \setminus \{z, w\} \rightarrow \mathbb{R}^4$ such that $\nu(t)$ is normal to M at $f(t)$, $|\nu(t)| = 1$, and the derivatives of ν satisfy*

$$(3.83) \quad |\nu^{(i)}(t)| \leq C(i) \text{dist}(t, \{z, w\})^{-i},$$

for all $t \in S(z, w) \setminus \{z, w\}$ and $i \geq 1$. These constants depend only on i and our excellent package.

This is easy to see. The point is that there is a smooth unit normal on Σ simply because Σ is smooth, and there are smooth unit normals to the Σ_α 's by self-similarity. We can get ν by piecing together these unit normals. The estimates (3.83) come from the estimates on f in Sublemma 3.78 and the self-similarity of the Σ_α 's. This proves Sublemma 3.82.

From now on (in this proof of Lemma 3.72) we let ν be as in the Sublemma. Set $v_0 = (w-z)/|w-z|$ and $v_1 = (0, 0, 0, 1)$ (= the standard unit normal to P). When we write $f'(t)$ we shall mean the derivative of f at t in the direction v_0 .

Sublemma 3.84. *There is a smooth map ϕ from $S(z, w) \setminus \{z, w\}$ into rotations on \mathbb{R}^4 such that ϕ satisfies $\phi(t)v_0 = f'(t)$, $\phi(t)v_1 = \nu(t)$, and*

$$(3.85) \quad |\phi^{(i)}(t)| \leq C(i) \operatorname{dist}(t, \{z, w\})^{-i},$$

for all $t \in S(z, w) \setminus \{z, w\}$ and $i \geq 0$, where $C(i)$ depends only on i and our excellent package.

Let us first resolve the "differentiated" version of this problem, in which we construct a family of antisymmetric linear mappings which will turn out to be $\phi'(t)\phi(t)^{-1}$. Given $t \in S(z, w) \setminus \{z, w\}$ define linear mappings $\psi_0(t), \psi_1(t) : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ by

$$(3.86) \quad \psi_0(t)\zeta = f''(t)\langle f'(t), \zeta \rangle + \nu'(t)\langle \nu(t), \zeta \rangle,$$

$$(3.87) \quad \psi_1(t)\zeta = \nu(t)\langle \nu'(t), f'(t) \rangle \langle f'(t), \zeta \rangle \\ + f'(t)\langle f''(t), \nu(t) \rangle \langle \nu(t), \zeta \rangle.$$

Note that $f'(t)$ and $\nu(t)$ are orthogonal to each other for all t , so that $\psi_0(t)f'(t) = f''(t)$ and $\psi_0(t)\nu(t) = \nu'(t)$. We also have that

$$(3.88) \quad \langle f''(t), f'(t) \rangle = 0, \quad \langle \nu'(t), \nu(t) \rangle = 0, \\ \langle \nu'(t), f'(t) \rangle + \langle \nu(t), f''(t) \rangle = 0,$$

because $\langle f'(t), f'(t) \rangle$, $\langle \nu(t), \nu(t) \rangle$, and $\langle \nu(t), f'(t) \rangle$ are all constant and hence have vanishing derivative. (Remember from Sublemma 3.78 that $|f'(t)| \equiv 1$.) This implies that $\psi_1(t)$ is antisymmetric and that $\psi_0^t(t)f'(t) = \psi_1(t)f'(t)$ and $\psi_0^t(t)\nu(t) = \psi_1(t)\nu(t)$ for all t , where ψ_0^t denotes the transpose of ψ_0 . Set $\psi = \psi_0 - \psi_0^t + \psi_1$. Thus $\psi(t)$ is antisymmetric and $\psi(t)f'(t) = f''(t)$ and $\psi(t)\nu(t) = \nu'(t)$ for all t .

We want to define ϕ now by solving the differential equation $\phi' = \psi\phi$. Let u be the midpoint of $S(z, w)$, and take $\phi(u)$ to be any rotation which satisfies $\phi(u)v_0 = f'(u)$, $\phi(u)v_1 = \nu(u)$. With this choice made we can extend ϕ to all of $S(z, w) \setminus \{z, w\}$ by solving the aforementioned differential equation. Because ψ is always antisymmetric we get that every $\phi(t)$ is a rotation. Our choice of ψ also ensures that $\phi(t)v_0 = f'(t)$ and $\phi(t)v_1 = \nu(t)$ for all t . To get the bounds (3.85) we observe that

$$(3.89) \quad |\psi^{(i)}(t)| \leq C(i) \operatorname{dist}(t, \{z, w\})^{-i-1},$$

for all $t \in S(z, w) \setminus \{z, w\}$ and $i \geq 0$, where $C(i)$ depends only on i and our excellent package. This follows from the definition of ψ and straightforward computation. The bounds for ϕ follow easily from this (and the fact that every $\phi(t)$ is a rotation, and hence has norm one). This proves Sublemma 3.84.

Sublemma 3.90. *There is a small number $\eta > 0$ so that f admits an extension to a smooth mapping (also called f) from $S_\eta(z, w)$ into M such that the differential of this extension at $t \in S(z, w)$ equals the restriction of $\phi(t)$ to P , where ϕ is as in Sublemma 3.84, and such that*

$$(3.91) \quad |\nabla^i f(x)| \leq C(i) \operatorname{dist}(x, \{z, w\})^{-i+1},$$

$$(3.92) \quad C^{-1} \operatorname{dist}(x, \{z, w\}) \leq \operatorname{dist}(f(x), F) \leq C \operatorname{dist}(x, \{z, w\}),$$

for all $x \in S_\eta(z, w) \setminus \{z, w\}$ and $i \geq 1$. These constants η , C , and $C(i)$ depend only on our excellent package.

There are several ways to prove this, all of them boring. Here's one. Define an auxiliary extension g of f as a map from $S_\eta(z, w)$ into \mathbb{R}^4 by taking g to be affine in the directions perpendicular to $S(z, w)$, with the affine mapping chosen in the obvious way using ϕ from Sublemma 3.84. It is not hard to check that g satisfies the analogues of (3.91) and (3.92), using (3.80) and (3.85), at least if η is small enough (for the first inequality in (3.92)). In particular g is Lipschitz, with a uniform bound. This Lipschitz bound implies that the image of g stays very close to M compared to its distance to F when η is small. To make this precise let π denote the orthogonal projection of P onto the line through $S(z, w)$. Then

$$(3.93) \quad \begin{aligned} |g(x) - f(\pi(x))| &= |g(x) - g(\pi(x))| \\ &\leq C|x - \pi(x)| \\ &\leq C\eta \operatorname{dist}(x, \{z, w\}) \end{aligned}$$

when $x \in S_\eta(z, w)$. We are using here the fact that $\pi(x) \in S(z, w)$ when $x \in S_\eta(z, w)$ (and $\eta < 1$) and also the Lipschitzness of g .

We want to define $f(x)$ for $x \in S_\eta(z, w)$ by taking the point on M which is closest to $g(x)$. A priori this is dangerous, but here we need only deal with pieces of M which are far from the singular set compared to the length scale at which we are working, and there is no problem with the nearest-point-projections on such small pieces of M . More precisely, there is a small constant $a > 0$ depending only on our excellent package so that if $\xi \in M \setminus F$ and $B = B(\xi, a \operatorname{dist}(\xi, F))$, then the mapping Π on B which takes a point and sends it to the (unique) nearest point in M is well defined, smooth, and satisfies

$$(3.94) \quad \sup_B |\nabla^i \Pi| \leq C(i) \operatorname{dist}(\xi, F)^{-i+1},$$

for $i \geq 1$. This is not hard to prove, using the smoothness of M away from F and the self-similarity property (3.20) to get the uniform estimates. (To do this from scratch one must compute a little to reduce to the inverse function theorem.)

Once we have these nearest-point-projections on small balls like B we can get f as in Sublemma 3.90 by projecting g onto M . This will only work when η is small enough, which ensures that the image points of g lie in balls like the ones just described, because of (3.93) and (3.79). It is not hard to check that this definition of f satisfies (3.91) and (3.92). Also, we defined g so that it had the correct differential along $S(z, w)$, and the nearest-point-projections onto M don't change that (because the image of the differential lies in the tangent space to M). This proves Sublemma 3.90.

Sublemma 3.95. *Let $f : S_\eta(z, w) \rightarrow M$ be as in Sublemma 3.90. If $\varepsilon > 0$ is small enough, then the restriction of f to $S_\varepsilon(z, w)$ is bilipschitz, with ε and the bilipschitz constant depending only on our excellent package.*

The point here is that we chose ϕ carefully to make the extension spread out in the right way. We shall show first that f is bilipschitz on certain small balls using our choice of ϕ , and then we shall use the bilipschitzness of f on $S(z, w)$ to control the global behavior of f .

Given $t \in S(z, w)$ set $r = \operatorname{dist}(t, \{z, w\})$ and $B = B(t) = B(t, br) \cap P$, where $b > 0$ is small and to be chosen. Let us show that if b is small enough then the restriction of f to B is bilipschitz with a bounded

constant. This is an easy consequence of Sublemma 3.90. Let $A(x)$ be the affine function from P into \mathbb{R}^4 defined by $A(x) = f(t) + \phi(t)x$, where ϕ is as in Sublemma 3.84. This is the affine Taylor approximation to f at t , and it preserves distances, since $\phi(t)$ is a rotation. We can estimate $f - A$ (or rather its gradient) on B using (3.91), and we get

$$(3.96) \quad \begin{aligned} \sup_B |\nabla(f - A)| &\leq C b r \sup_B |\nabla^2 f| \\ &\leq C b r \sup_B \text{dist}(x, \{z, w\})^{-1} \leq C b. \end{aligned}$$

This uses also the fact that $\text{dist}(x, \{z, w\}) \approx r$ when $x \in B$ (assuming $b \leq 1/2$, say). If b is small enough then we conclude that

$$(3.97) \quad \frac{1}{2} |x - y| \leq |f(x) - f(y)| \leq 2 |x - y|, \quad \text{when } x, y \in B,$$

since A preserves distances. Choose such a b and let it be fixed from now on.

Now let x, y be any pair of points in $S_\varepsilon(z, w)$, and let us check the bilipschitz condition for them. We already know from (3.91) that f is Lipschitz on $S_\varepsilon(z, w)$, and so we need only concern ourselves with getting a lower bound for $|f(x) - f(y)|$. We may assume that x and y do not both belong to any ball $B(t), t \in S(z, w)$ as above, since we have (3.97) already. Thus $y \notin B(\pi(x))$, and this implies that $|x - y| \geq 10^{-2} b \text{dist}(x, \{z, w\})$ if ε is small enough. We get the same inequality with the roles of x and y reversed, and so

$$(3.98) \quad |x - y| \geq 10^{-2} b \max\{\text{dist}(x, \{z, w\}), \text{dist}(y, \{z, w\})\}.$$

Since f is bilipschitz on $S(z, w)$ (Sublemma 3.78) we get that

$$(3.99) \quad |f(\pi(x)) - f(\pi(y))| \geq C^{-1} |\pi(x) - \pi(y)|.$$

This implies that

$$(3.100) \quad |f(\pi(x)) - f(\pi(y))| \geq C^{-1} |x - y|,$$

because (3.98) yields $|x - \pi(x)| + |y - \pi(y)| \leq C b^{-1} \varepsilon |x - y| \leq 10^{-1} |x - y|$ when ε is small enough. To get back to $|f(x) - f(y)|$ we observe that

$$(3.101) \quad |f(\pi(x)) - f(x)| \leq C \varepsilon \text{dist}(x, \{z, w\}) \leq C b^{-1} \varepsilon |x - y|$$

because of the Lipschitzness of f , the fact that $x \in S_\varepsilon(z, w)$, and (3.98). We also have the same estimate for y instead of x , and we conclude from (3.100) that

$$(3.102) \quad |f(x) - f(y)| \geq C^{-1} |x - y|$$

if ε is small enough.

This completes the proof of Sublemma 3.95, and Lemma 3.72 follows.

Lemma 3.103. *The conclusions of Proposition 3.70 hold for M^j when $p, q \in \Omega$, whether or not $j = \infty$.*

Again choose δ in some \mathcal{S}_m so that $p, q \in \Omega_\delta$ and m is as large as possible. We may as well assume that $m < j - 1$ and that one of p and q lies in Ω_γ for some $\gamma \in \mathcal{S}_{m+2}$, since otherwise we can apply Lemma 3.71. This implies that $|p - q| \geq C^{-1} \text{diam } \Omega_\delta$ for some constant C (which depends only on the excellent package); if $|p - q|$ were small compared to $\text{diam } \delta$, then we could use the fact that one of p and q lies in an Ω_γ , $\gamma \in \mathcal{S}_{m+2}$, to conclude that $p, q \in \Omega_\zeta$ for some child $\zeta \in \mathcal{S}_{m+1}$ of δ , in contradiction to the maximality of m .

Under these conditions we can apply the same basic construction as in the proof of Lemma 3.72. It can happen now that now one or both of p and q does not lie in F , or that $j < \infty$, so that the sequences of α_l 's and β_l 's might have to stop in a finite number of steps. In fact, we could have that one of p or q lies in Σ_δ , so that there would be no α_l 's, or no β_l 's. Thus it may be necessary to modify the construction at one or both "ends", but the estimates and underlying principles remain the same. One chooses points like the p_l 's and the q_l 's, one connects these points with nice curves in M^j (using Sublemma 3.77, extended slightly to include Σ'_α 's when $j < \infty$), one combines the curves and parameterizes the union by a bilipschitz map as in Sublemma 3.78, one extends this mapping as in Sublemma 3.90 (using a good family of frames as in Sublemma 3.84), and then one checks bilipschitzness as in Sublemma 3.95. The details are left to the reader.

Lemmas 3.71, 3.72, and 3.103 cover all the possible locations of $p, q \in M^j$ except for $p \in \Omega_\alpha$ for some $\alpha \in \mathcal{S}_1$ and $q \in P \setminus \Omega$ (or the other way around). (See (3.17) and (3.19).) In this case we have that $|p - q| \geq C^{-1}$ for some constant C . Set $m = 0$ and let δ be the empty sequence in \mathcal{S}_0 , so that $\Omega_\delta = \Omega$. We can choose $\alpha_l \in \mathcal{S}_l$ and $p_l \in \partial\Omega_{\alpha_l}$ as in the beginning of the proof of Lemma 3.72, except that

these sequences may stop in a finite number of steps. We can use the same basic construction as in the proof of Lemma 3.72 to connect p to an auxiliary point on the boundary of $\partial\Omega$ through a nice sequence of curves, as in Sublemmas 3.77 and 3.78. We can connect q to this auxiliary point in a nice way from $P \setminus \Omega$, simply because $P \cap \Omega$ is a bounded smooth domain in P and $P \setminus \Omega$ is connected (by (3.69)). (This is similar to part of Lemma 3.71.) If we do these things in a non-stupid manner then we can fatten up this connection between p and q to get a bilipschitz 3-ball in M^j which contains them. There is a minor difference in this situation, however. If S_ε denotes the analogue of $S_\varepsilon(z, w)$ (from the proof of Lemma 3.72) adapted to this situation, then the proportion of S_ε devoted to the connection from p to $\partial\Omega$ will be comparable in size to the diameter of Ω (a positive constant). If $|p - q|$ is very large, then this will be a small proportion of S_ε , much less than half, and most of S_ε will be devoted to the connection from q to $\partial\Omega$. This does not pose a serious problem, but it does mean that the bulge in the middle of S_ε should be placed away from Ω , where everything is flat. The details are again left to the reader.

This proves Proposition 3.70 in the case where $E = M$ or M^j . Suppose now that $E = \widetilde{M}$. Let $p, q \in \widetilde{M}$ be given, $p \neq q$. Let $\{B_k\}$ be the sequence of balls used in the definition of \widetilde{M} (just after (3.20)), and let \widetilde{M}^k be the affine image of M^k with the interesting part squeezed into B_k , as in the definition of \widetilde{M} . If $p, q \in (3/2)B_k$ for some k , then we can use the previous result for M^k to obtain that p and q are contained in a set $W' \subseteq \widetilde{M}^k$ which is bilipschitz equivalent to a Euclidean 3-ball. If one is careful about the previous construction one can choose W' so that $W' \subseteq 2B_k$, but one can also simply force this to happen, in the following manner. We can choose W' so that $W' \subseteq CB_k$ for some uniformly bounded constant C ; if this is not true to begin with, it simply means that W' is unnecessarily large, and we can replace it with a smaller subset. Let Ψ be a bilipschitz map from CB_k into $2B_k$ which equals the identity on $(3/2)B_k$ and maps $(CB_k \setminus (3/2)B_k) \cap P$ into $2B_k \cap P$. If we set $W = \Psi(W')$, then $W \subseteq \widetilde{M}^k \cap 2B_k$ and hence $W \subseteq \widetilde{M}$. We also have that W is uniformly bilipschitz equivalent to a Euclidean 3-ball and contains p and q . Thus the case where $p, q \in (3/2)B_k$ for some k can be reduced to the previous results.

If neither p nor q lie in any B_k , then it is not hard to show directly that they are contained in a subset W of $P \setminus (\cup B_k) \subseteq \widetilde{M}$ which is bilipschitz equivalent to a Euclidean 3-ball with a uniform bound. It is easier to think of W as being bilipschitz equivalent to a set like $S_\varepsilon(u, v)$

as described at the beginning of the proof of Proposition 3.70 rather than a standard ball, so that it is easier to visualize the way the ends are placed into $P \setminus (\cup B_k)$. It is also convenient to choose the B_k 's in the rather specific manner described in the paragraph after (3.20).

We are left with the case where p lies in some B_k and q lies outside $(3/2)B_k$. Suppose first that q does not lie in any other B_l . We can use the same kind of argument as in Lemmas 3.72 and 3.103 to connect p to an auxiliary point in $\partial((3/2)B_k)$ in a nice way (as in Sublemma 3.78), using a family of smooth arcs. We can then connect from there to q inside $P \setminus (\cup B_l)$ by a more direct construction, since the structure of $P \setminus (\cup B_l)$ is so simple. These two connections can be combined and then filled out to get a bilipschitz 3-ball in \widetilde{M} which connects p and q . This combination and filling-out will be realized concretely as a bilipschitz map from a set of the form $S_\varepsilon(z, w)$ into \widetilde{M} , with the restriction of this map to one end of $S_\varepsilon(z, w)$ providing the connection from p to $\partial((3/2)B_k)$, and the rest corresponding to the connection from there to q . The diameter of $S_\varepsilon(z, w)$ should be comparable to $|p - q|$, while the diameter of the piece at the end corresponding to the connection from p to $\partial((3/2)B_k)$ should be comparable to the radius of B_k . These sizes are not inconsistent with each other, because $|p - q|$ is at least one-half the radius of B_k . It may well be that $|p - q|$ is much larger than the radius of B_k , in which case we should be slightly careful to put the middle of $S_\varepsilon(z, w)$ far away from the B_l 's. This type of detail is awkward but not at all deep.

If instead q does lie in some B_l , then we use the method of Lemmas 3.72 and 3.103 to connect p to an auxiliary point in $\partial((3/2)B_k)$ and q to an auxiliary point in $\partial((3/2)B_l)$ by well-behaved curves (as in Sublemma 3.78). We can connect these auxiliary points in a nice way inside $P \setminus (\cup B_m)$ by a direct construction, and these three connections can be combined and filled out in such a way as to get a bilipschitz 3-ball in \widetilde{M} which connects p and q . If we think of the combination of these connections as being represented by a bilipschitz map from a set of the form $S_\varepsilon(z, w)$, then the connections from p to $\partial((3/2)B_k)$ and from q to $\partial((3/2)B_l)$ will correspond to pieces of $S_\varepsilon(z, w)$ at the two ends of $S_\varepsilon(z, w)$. These two pieces will have sizes comparable to the radii of B_k and B_l , respectively, while the diameter of $S_\varepsilon(z, w)$ should be about the same as $|p - q|$. These sizes are consistent with each other, because $|p - q|$ is at least one-half the radius of each of B_k and B_l . As usual, if $|p - q|$ is much larger than the radii of B_k and B_l , then we

have to be careful to map the large middle of $S_\varepsilon(z, w)$ away from the various B_i 's.

This completes the proof of Proposition 3.70, modulo various packages of details left to the reader. In all cases the argument can be understood in terms of connecting sequences of smooth curves together and then parameterizing these curves and filling out these parameterizations to get a bilipschitz mapping defined on a set of the form $S_\varepsilon(z, w)$. There are some minor variations among the various cases - whether the sequence of constituent curves is finite or infinite, whether we work only in bounded regions like the Ω_α 's or we have to go wandering outside to the large flat regions of the M 's, or whether we have to worry about where the "bulge" in the middle of $S_\varepsilon(z, w)$ should be sent - but the actual constructions are simpler than their gory detailed descriptions.

Note that if we were only interested in the last part of Proposition 3.70 (about connecting p and q by a curve in E whose length is bounded by a constant times $|p - q|$) then the preceding proof would simplify substantially. For instance, under the assumptions of Lemma 3.72 we would need much less than Sublemma 3.78.

Corollary 3.104. *If our initial package satisfies (3.69) and $E = M, \widetilde{M}$, or M^j , $j < \infty$, then E is linearly locally connected, with constants that do not depend on j in the latter case. This means that there is a constant C so that for each $x \in E$ and $t > 0$ we have that any two points in $E \cap B(x, t)$ lie in the same component of $E \cap B(x, Ct)$, and any two points in $E \setminus B(x, t)$ lie in the same component of $E \setminus B(x, C^{-1}t)$.*

This is an easy consequence of Proposition 3.70. One could also prove it more directly, in the same way that it is much easier to connect pairs of points in the M 's by curves which are not too long than it is to get the bilipschitz balls in Proposition 3.70. If one were to try to give a more direct proof of the corollary, then Lemma 3.44 would be rather convenient for proving the second part of the linear local connectedness condition, but it is better to apply it to $B(x, (1 + a^{-1})^{-1}t)$, where a is as in Lemma 3.65, than to $B(x, t)$ itself.

The main result of [HK] provides a good reason to care about linear local connectedness. This result states that if a 3-dimensional regular set E is linearly locally connected, and if E admits a homeomorphic parameterization by \mathbb{R}^3 by a mapping which satisfies the pointwise definition of quasiconformality, then in fact this mapping must be quasymmetric. Thus the sets promised in Theorem 1.12 cannot even be

quasiconformally equivalent to \mathbb{R}^3 . In this particular situation it is a little ridiculous to cite the theorem in [HK], because the same result could be obtained more directly using standard results about quasiconformal mappings between domains in \mathbb{R}^n and the fact that the Σ_α 's are all smooth and similar to each other. (One could also use Proposition 3.70.)

In the next three sections we consider specific examples of excellent packages.

4. The Whitehead example.

We shall continue to use the definitions and notation of the preceding section here.

Let D be a smooth solid torus in \mathbb{R}^3 , and let D_1 be another smooth solid torus whose closure is contained in D and which is clasped inside of D in the usual manner for generating the Whitehead continuum, *i.e.*, in the manner shown in [D, p. 68, Figure 9-7] (see also [K, p. 81ff]). Note that this clasping prevents D_1 from being isotopic to a standard (small) solid torus in D , but D_1 is homotopically trivial in D . In other words, we can deform D_1 inside D into something small, but D_1 has to cross itself along the way. Let ϕ_1 be a smooth diffeomorphism which maps a neighborhood of \overline{D} onto $\overline{D_1}$ and sends D onto D_1 . (One can think of simply bending D around in \mathbb{R}^3 to get D_1 .) This gives us an initial package, and note that it satisfies (3.69).

To get an excellent package we take Ω and ω_1 to be (4-dimensional) solid tori in \mathbb{R}^4 which satisfy (3.3). The property of being a "4-dimensional solid torus" here means in particular that Ω is diffeomorphic to $D \times (-1, 1)$ and similarly for ω . It is helpful to take ω_1 to be much flatter than Ω in the x_4 direction, to provide plenty of room to move around. We take ρ to be small and positive but otherwise at our disposal, and we choose ψ_1 to be a combination of a translation (by an element of P) and dilation by ρ such that $\Omega_1 = \psi_1(\Omega)$ and its closure lie in a ball in Ω . To build θ we use the fact that (unlike D_1 and D) ω_1 is not really clasped inside Ω , because of the freedom of movement provided by the extra dimension. We can lift up one end of the clasping part of ω_1 (in the x_4 direction) while leaving the other clasping end alone, and then we can bring it around in Ω and shrink it until ω_1 is deformed into Ω_1 . One can do this process in such a way that $\omega_1 \cap P = D_1$ is deformed to $\Omega_1 \cap P$, and one can even extend this

motion to all of \mathbb{R}^4 without ever moving points outside of Ω . These observations are more easily made rigorous using the following well-known “isotopy extension” result.

Lemma 4.1. *Let O be a bounded open set in \mathbb{R}^n and let K be a compact subset of O . Suppose that $g(x, t)$ is a smooth \mathbb{R}^n -valued mapping defined for x in a neighborhood of K and t in $[0, 1]$ such that $g(x, 0) = x$ for all x and $g(\cdot, t)$ is a diffeomorphism onto its image for each t . Assume also that $g(K \times [0, 1]) \subseteq O$. Then there exists a smooth mapping $G : \mathbb{R}^n \times [0, 1] \rightarrow \mathbb{R}^n$ such that $g(x, t) = G(x, t)$ when $x \in K$, $G(\cdot, t)$ is a diffeomorphism on \mathbb{R}^n for each t , $G(x, t) = x$ when $x \in \mathbb{R}^n \setminus O$, and $G(x, 0) = x$ for all x .*

To see this we begin by observing that there is a smooth time-dependent vector field $V(y, t)$ on \mathbb{R}^n such that $V(g(x, t), t) = \partial g(x, t) / \partial t$ for all $x \in K$ and $t \in [0, 1]$ and such that $V(y, t)$ vanishes whenever x lies outside any prescribed neighborhood of $g(K \times [0, 1])$. Indeed, we start by defining V on pairs of the form $(g(x, t), t)$ for x near K and t in $[0, 1]$, and we can extend V to all $(x, t) \in \mathbb{R}^n \times [0, 1]$ using a cut-off function. Let $G(x, t)$ be the solution of the ordinary differential equation $\partial G(x, t) / \partial t = V(G(x, t), t)$ for all $(x, t) \in \mathbb{R}^n \times [0, 1]$, with the initial condition $G(x, 0) = x$ for all x . The uniqueness theorem for ordinary differential equations and standard facts imply that G has the required properties. This proves Lemma 4.1.

Using Lemma 4.1 and the deformation process described above one can build a diffeomorphism θ on \mathbb{R}^4 which satisfies (3.5), $\theta(\omega_1) = \Omega_1$, and also $\theta(D_1) = \Omega_1 \cap P$. If one is careful one can get that $\theta = \psi_1 \circ \phi_1^{-1}$ on a neighborhood of \overline{D}_1 in P , but it is simpler to just redefine ϕ_1 so that this is true.

The conclusion of all this is that we get an excellent package associated to our initial package, and we can take ρ to be as small as we want. The construction of Section 3 produces sets M^j and \widetilde{M} .

Theorem 4.2. *\widetilde{M} satisfies (†) but not (*). The M^j 's for $j < \infty$ satisfy (†) with constants which remain bounded, but they do not satisfy (*) with bounded constants.*

This result together with Lemma 3.45 and Proposition 3.70 imply Theorem 1.12.a).

Of course the M^j 's for $j < \infty$ and \widetilde{M} are homeomorphic to \mathbb{R}^3 (and hence topological manifolds) by Lemma 3.21.

Let us first prove the bounded contractability conditions for the M^j 's. Let a be chosen as in Lemma 3.65, and let $0 \leq j < \infty$, $x \in M^j$, and $r > 0$ be given. We want to check that we can contract $B(x, r) \cap M^j$ to a point inside $B(x, Cr) \cap M^j$ for some constant C which does not depend on x, r , or j . Lemma 3.44 permits us to consider separately the five cases listed there. All but case iv) are covered by Lemma 3.65, and so we suppose that iv) holds.

Let Ω_δ be the parent of Ω_α , and let us show that

$$(4.3) \quad \Omega_\alpha \cap M^j \text{ can be contracted to a point inside } \Omega_\delta \cap M^j.$$

To prove this we may as well assume that $l = 1$, so that $\Omega_\delta = \Omega$, since otherwise we can use ψ_δ to pull everything back to Ω . (This will change M^j to M^{j-l} , as in (3.20), but that is okay.) Using the homeomorphism h_j from Lemma 3.21 we can reduce further to the problem of contracting D_1 to a point in D . This we can do, because of the specific choice of the Whitehead initial package. (It is not true for arbitrary initial packages.)

This proves that the M^j 's satisfy (†) uniformly, and so we consider now \widetilde{M} . Recall its definition and related notation from the paragraph after (3.20). Part of this argument also works in general, and so we split it off as a separate lemma.

Lemma 4.4. *Suppose that we have started with any excellent package and constructed \widetilde{M} as in Section 3. Let $b > 0$, $x \in \widetilde{M}$ and $r > 0$ be given, and assume that for each k we have either $B(x, r) \cap B_k = \emptyset$ or $r > b$ radius B_k , where $\{B_k\}$ is the sequence of balls in the definition of \widetilde{M} . Then there is a relatively open set $U \subseteq \widetilde{M}$ which is homeomorphic to a 3-ball and satisfies $B(x, r) \cap \widetilde{M} \subseteq U \subseteq B(x, Cr) \cap \widetilde{M}$, where C depends on b but not on x or r .*

Let x, r be given as in the lemma. If $B(x, r) \cap B_k = \emptyset$ for all k then $B(x, r) \cap \widetilde{M} = B(x, r) \cap P$ and there is nothing to do. Suppose that $B(x, r) \cap B_k \neq \emptyset$ for some k and choose such a k which is as small as possible. Let B'_k be as in the paragraph defining \widetilde{M} (just after (3.20)), so that radius $B'_k \leq 100$ radius B_k , $B_l \subset B'_k$ when $l \geq k$, and $B_l \cap B'_k = \emptyset$ when $l < k$. Let \widetilde{h} be as in Lemma 3.21, and set $U = \widetilde{h}((B(x, r) \cup B'_k) \cap P)$. This is a topological 3-ball, because the

union of two intersecting (standard) 3-balls is a topological 3-ball. We also have that

$$(4.5) \quad U = \left((B(x, r) \cup B'_k) \setminus \left(\bigcup_{l \geq k} B_l \right) \right) \cap P \cup \left(\left(\bigcup_{l \geq k} B_l \right) \cap \widetilde{M} \right) \\ \subseteq B(x, r) \cup B'_k,$$

because of the properties of \tilde{h} in Lemma 3.21. The hypothesis of Lemma 4.4 guarantees that radius $B_k \leq b^{-1}r$, which implies in turn that $\text{diam } U \leq Cr$ for a suitable constant C . It is easy to check also that $B(x, r) \cap \widetilde{M} \subseteq U$ because of the properties of \tilde{h} in Lemma 3.21. This proves Lemma 4.4.

Let us now show that \widetilde{M} satisfies (\dagger) . We already know that it is a topological manifold, and so we need only check the contractability condition. Let $b > 0$ be small, which we get to choose. Let $x \in \widetilde{M}$ and $r > 0$ be given. Because of Lemma 4.4 we may as well assume that there is a k such that $B(x, r)$ intersects B_k and $r \leq b$ radius B_k . These conditions imply in particular that $B(x, r) \subseteq 2B_k$. Hence $B(x, r) \cap \widetilde{M} = B(x, r) \cap A_k(M^k)$, and since we already know that (\dagger) holds uniformly for the M^j 's, and since the similarity A_k does not affect the (\dagger) property (or its constant) we get that there is a constant C' such that $B(x, r) \cap A_k(M^k)$ can be contracted to a point inside $B(x, C'r) \cap A_k(M^k)$. If b is small enough (depending on C'), then $B(x, C'r) \subseteq 2B_k$ too, so that $B(x, C'r) \cap A_k(M^k) = B(x, C'r) \cap \widetilde{M}$. Thus $B(x, r) \cap \widetilde{M}$ can be contracted to a point inside $B(x, C'r) \cap \widetilde{M}$, which is what we wanted. This proves that \widetilde{M} also satisfies (\dagger) .

(Incidentally, the choice of b in the preceding argument is a little bit stupid, in the sense that if one looks carefully one sees that a far less small choice of b would work fine. However, this additional complication is not needed for the proof.)

It remains to show that $(*)$ is bad. This will be derived from a famous property of the Whitehead continuum. Since $n = 1$ we have that each S_l has only one element, and so if the defining sequence $\{C_l\}$ is constructed from our initial package as described after Definition 2.3, then each C_l has only one component, namely D_α for $\alpha =$ the unique element of S_l . Set $W = \bigcap_l C_l$. This is the Whitehead continuum, and it is the only nondegenerate element of the decomposition G associated to $\{C_l\}$ as discussed just after Definition 2.2. We shall sometimes find it convenient to view W as a subset of the 3-sphere S^3 .

Recall that a compact set $K \subseteq \mathbb{R}^3$ is said to be cellular if for each open set $V \supseteq K$ there is a topological 3-ball U such that $K \subseteq U \subseteq V$.

(See [D, p. 35].) A decomposition G of \mathbb{R}^3 is said to be cellular if each element of G is a cellular subset of \mathbb{R}^3 . (See [D, p. 36, Corollary 2A].) A topological space X is said to be simply connected at ∞ if for each compact set $K \subseteq X$ there is a compact set $L \subseteq X$ with $L \supseteq K$ such that every loop in $X \setminus L$ can be contracted to a point in $X \setminus K$. (See [K, p. 83].) Note that $S^3 \setminus K$ is simply-connected at ∞ if $K \subseteq \mathbb{R}^3$ is cellular.

Proposition 4.6. *W is not a cellular subset of \mathbb{R}^3 , and in fact $S^3 \setminus W$ is not simply connected at ∞ .*

The first statement is a reformulation of [D, p. 76, Proposition 9] (see also the top of p. 69 of [D]). The second statement is discussed on [K, p. 82-83].

Lemma 4.7. *If \widetilde{M} satisfies (*), or if the M^j 's satisfy (*) with a uniformly bounded constant, then W is cellular.*

Suppose first that the M^j 's satisfy (*) with a uniformly bounded constant C_0 . Let l be given, and let j be larger than l and at our disposal. Let α and β be the unique elements of \mathcal{S}_l and \mathcal{S}_j , respectively, and fix a point x in $\Omega_\beta \cap M^j$. Our assumption on the M^j 's implies that there is a topological 3-ball U such that

$$(4.8) \quad \Omega_\beta \cap M^j \subseteq U \subseteq B(x, C_0 \text{ diam } \Omega_\beta) \cap M^j.$$

If $j - l$ is sufficiently large, depending on C_0 , then $B(x, C_0 \text{ diam } \Omega_\beta) \subseteq \Omega_\alpha$, so that $U \subseteq \Omega_\alpha \cap M^j$. We can use the homeomorphism h_j in Lemma 3.21 to bring U back to \mathbb{R}^3 , and we get a topological 3-ball $V = h_j^{-1}(U)$ such that $D_\beta \subseteq V \subseteq D_\alpha$. Since l was arbitrary and we obtain that W is cellular.

The same argument works if we assume that \widetilde{M} satisfies (*). The point is that \widetilde{M} contains a copy of $\Omega \cap M^j$ for each j , and we never left Ω in the preceding argument. This proves Lemma 4.7.

Proposition 4.6 and Lemma 4.7 imply that \widetilde{M} does not satisfy (*), and that the M^j 's do not satisfy (*) with a uniformly bounded constant. This proves Theorem 4.2.

REMARK 4.9. Lemma 4.7 works for any excellent package. That is, if we start with an excellent package and construct M^j and \widetilde{M} as in

Section 3, and if either \widetilde{M} satisfies (*) or the M^j 's satisfy (*) with a uniformly bounded constant, then the decomposition G associated to the initial package (as in Section 2) is cellular, *i.e.*, each element of G is a cellular subset of \mathbb{R}^3 . This can be proved by exactly the same argument as in Lemma 4.7, with minor additional complications when the number n of D_i 's is greater than 1.

REMARK 4.10. The set M corresponding to this excellent package is homeomorphic to \mathbb{R}^3 with the Whitehead continuum contracted to a point. This famous non-manifold has the property that its product with \mathbb{R} is homeomorphic to \mathbb{R}^4 . (See [K, p. 87, Theorem 1] and [D, p. 83, Theorem 3].) This space $M \times \mathbb{R}$ arises in one description of the simplest Casson handles, as in p. 86 and the bottom half of p. 83 of [K]. These Casson handles have been shown in [Bi] to be exotic in a certain sense (with respect to their smooth structures). One can imagine that the quasiconformal version of this exoticness is also true, because of [DoS].

5. Bing's dogbone space.

We shall continue to use freely the notations and definitions of Section 3.

To prove Theorem 1.12.b) we shall use the construction of Bing's dogbone space, which is given in [B3] and as [D, p. 64, Example 4]. For this we define an initial package by taking D to be a smooth solid two-handled torus which is embedded in \mathbb{R}^3 in the standard way (no funny business with the two handles). We also take $n = 4$ and the D_i 's to be solid two-handled tori in D arranged in the manner pictured in [B3, p. 486, Figure 1] and in [D, p. 65, Figure 9-4] (and with diffeomorphisms ϕ_i as in Definition 2.3 chosen in a non-stupid manner). This arrangement satisfies the requirements of Definition 2.3 and (3.69), and it has the additional feature that each D_i is (individually) embedded in D in a topologically standard (unlinked) manner. In particular,

$$(5.1) \quad \begin{array}{l} \text{for each } i = 1, 2, 3, 4 \text{ there is a topological 3-ball } U_i \\ \text{such that } \overline{D}_i \subseteq U_i \subseteq D. \end{array}$$

This is obvious from the pictures; the point is that each D_i is, as an individual domain, not linking with itself inside D in any way. (As a group the D_i 's are definitely linked, not homologically, but in the sense that they cannot be disentangled by an isotopy on D . This is related to

the fact that the U_i that works for one i will have to intersect the other D_j 's.) Condition (5.1) implies that the decomposition G associated to this initial package (as discussed in Section 2) is cellular (in the sense described in the preceding section). However, the quotient is not a manifold.

Theorem 5.2. (Bing, [B3, p. 498, Theorem 13]) \mathbb{R}^3/G is not a manifold.

For the proof of Theorem 1.12.b) (5.1) and Theorem 5.2 are the only properties that we shall need. Note that there are other examples in Section 9 (beginning on p. 61) of [D] of cellular decompositions of \mathbb{R}^3 which are obtained from initial packages and whose quotients are not topological manifolds. Actually, [D] does not address directly the issue of manifold quotients, only the stronger property of "shrinkability" of the decomposition, but a theorem of Armentrout [Ar] implies that the two properties are equivalent for (cellular) decompositions (as remarked near the bottom of p. 22 of [D]). The bottom line is that there are other examples that we could use to get Theorem 1.12.b).

We can also build an excellent package for this initial package. Let Ω and ω_i , $1 \leq i \leq 4$, be solid versions of D and the D_i 's in \mathbb{R}^4 which satisfy (3.3). The phrase "solid version" means that Ω should be diffeomorphic to $D \times (-1, 1)$, and similarly for the ω_i 's. It is also a good idea to require that the ω_i 's lie in a thin slab $\{x \in \mathbb{R}^4 : |x_4| < \varepsilon\}$, while Ω should contain a much fatter slab around most of D (and near the ω_i 's in particular). This allows us to translate an ω_i "up" (in the positive x_4 direction) away from the other ω 's and to move it around up there without getting too close to the others. This ensures that the ω_i 's are not linked in \mathbb{R}^4 in any manner. Let ψ_i , $1 \leq i \leq 4$, be similarities on \mathbb{R}^4 with a common dilation factor ρ which map P to itself and which send Ω to domains Ω_i with disjoint closures in Ω . It is convenient to require also that the Ω_i 's stay away from all the ω_i 's. We do not care too much about the specific value of ρ but it should be reasonably small and we may take it to be as small as we want. The main point now is that we can build a θ as in Definition 3.2. To do this we take ω_1 , we lift it up in the positive x_4 direction away from the slab $\{x \in \mathbb{R}^4 : |x_4| < \varepsilon\}$, we shrink it and slide it around until it is the same as Ω_1 but translated up a bit, and then we set it down onto Ω_1 . In this whole process we take care not to touch the ω_i 's or the Ω_i 's for $i \neq 1$, and also to remain inside Ω the whole time. We then repeat the process

for $\omega_2, \omega_3, \omega_4$. Lemma 4.1 allows us to extend these deformations to all of \mathbb{R}^4 in such a way that points in or near the complement of Ω are not moved, and points in the ω_i 's and the Ω_i 's move only when they are supposed to. As a result we get a diffeomorphism θ as in Definition 3.2. (To be honest, it is simpler to first build θ and then choose the ϕ_i 's so that (3.6) holds, rather than adjusting θ to fit the ϕ_i 's.)

Thus we have an excellent package associated to our initial package, where ρ can be made as small as we want, and so we can use the construction of Section 3 to produce sets M, M^j , and \widetilde{M} .

Theorem 5.3. *\widetilde{M} satisfies (*) but not (**). The M^j 's, $0 \leq j < \infty$, satisfy (*) with a uniformly bounded constant, but they do not satisfy (**) with uniform choices of the constant and modulus of continuity.*

Let us first show that the M^j 's satisfy (*) with a uniformly bounded constant. Let a be as in Lemma 3.65, and let $j, x \in M^j$, and $r > 0$ be given. Lemma 3.44 allows us to consider separately the cases i)-v) listed there, but Lemma 3.65 implies that we need only consider iv). Let α be as in iv) in Lemma 3.44, and let $\delta \in \mathcal{S}_{l-1}$ be its parent. It suffices to show that there exists a relatively open set U_α in M^j which is homeomorphic to a 3-ball and satisfies

$$(5.4) \quad \Omega_\alpha \cap M^j \subseteq U_\alpha \subseteq \Omega_\delta \cap M^j.$$

For the usual self-similarity reasons (i.e., (3.20), but with α replaced by δ) we can reduce this to the case where $l = 1$ (and j is replaced by $j - l + 1$). This case reduces to (5.1) because of Lemma 3.21. (Take U_α to be the image under h_{j-l+1} of the appropriate U_i .) This proves (5.4) and the fact that the M^j 's satisfy (*) with a bounded constant.

To show that \widetilde{M} satisfies (*), one uses Lemma 4.4 to reduce to the previous fact for the M^j 's. The argument is practically identical to the corresponding step in Section 4 (just after (4.5)), and we do not repeat it.

Next we show that (**) is bad for this excellent package.

Lemma 5.5. *If \widetilde{M} satisfies (**), or if the M^j 's satisfy (**) with uniform choices of the constant and modulus of continuity, then M is a topological manifold. (This works for any excellent package, and not just the particular ones considered in this section.)*

This is pretty straightforward. We have very precise control on the convergence of the M^j 's to M which implies convergence in the Hausdorff topology in particular. A uniform version of (**) would force M to satisfy (**) also, and of course (**) certainly implies that M is a topological manifold. If \widetilde{M} satisfies (**), then we use the fact that \widetilde{M} contains a translation and dilation of the most interesting part of each M^j . Since (**) is preserved by similarities we can use the same argument to conclude that M is a topological manifold.

Theorem 5.3 now follows from Lemma 5.5, Lemma 3.21 (which states that M is homeomorphic to \mathbb{R}^3/G , where G is the decomposition associated to our initial package), and Theorem 5.2. Theorem 1.12.b) follows from Theorem 5.3, Lemma 3.45, and Proposition 3.70.

Incidentally, the fact that Bing's dogbone space \mathbb{R}^3/G could be embedded topologically in \mathbb{R}^4 was observed long ago [Cu].

REMARK 5.6. Note that we can make the singular set F of M have Hausdorff dimension as small as we want in this example, by taking the parameter ρ to be small. See (3.48).

6. Bing doubling.

We shall use the definitions and notations from Section 3 freely in this section.

To prove Theorem 1.12.c) we use another example studied by Bing [B1], [D, p. 62, Example 1]. We start with a smooth solid torus D in \mathbb{R}^3 , and we take $n = 2$ and D_1, D_2 to be two disjoint smooth solid tori in D which are folded over and linked as in [B1, p. 357, Figure 3] and [D, p. 63, Figure 9-1]. Each of these two tori are (separately) embedded in a topologically trivial manner in D , and we have that

$$(6.1) \quad \begin{array}{l} \text{there are open sets } U_1, U_2 \subseteq D \text{ with } \overline{D}_i \subseteq U_i \text{ such that} \\ \overline{U}_1, \overline{U}_2 \text{ are each diffeomorphic to the closed unit 3-ball.} \end{array}$$

However, as a pair, D_1 and D_2 are linked, in the sense that they cannot be pulled apart by an isotopy of D onto itself. Note that D , D_1 , and D_2 satisfy (3.69).

Let Ω and ω_1, ω_2 be solid versions of the D 's in \mathbb{R}^4 which satisfy (3.3). As before, "solid version" means that Ω should be diffeomorphic

to $D \times (-1, 1)$, with $\Omega \cap P$ corresponding to $D \times \{0\}$, and similarly for the ω_i 's. One should not go out of one's way to choose them stupidly, and we shall see in the next section that it is a good idea to choose them to be symmetric about P . As in the preceding section, it is better to choose ω_1, ω_2 to lie much closer to P than Ω , so that we can disentangle ω_1 from ω_2 in Ω with ease. For instance, it is convenient to require that there exist $\varepsilon > 0$ so that we have the following solid version of (6.1):

$$(6.1') \quad \omega_i \subseteq U_i \times (-\varepsilon, \varepsilon) \subseteq \bar{U}_i \times [-\varepsilon, \varepsilon] \subseteq \Omega.$$

Let us ask also that $\omega_i \supseteq D_i \times (-b, b)$ for some $b > 0$ (for a minor technical convenience in Section 8).

Let ψ_1 and ψ_2 be similarities on \mathbb{R}^4 with the same dilation factor ρ which map P onto itself and which send Ω onto domains Ω_1, Ω_2 in Ω with disjoint closures and which stay away from the ω_i 's. As usual it is good for ρ to be small, and we can take it to be as small as we want. For the same reason as in Section 5 we can build a mapping θ which satisfies the requirements of Definition 3.2. We lift up ω_1 (in the positive x_4 direction), deform it into a copy of Ω_1 sitting just over Ω_1 , and drop it onto Ω_1 , and then we repeat the process for Ω_2 , taking care that the deformations stay inside Ω and do not disturb the other players (*i.e.*, we do not touch ω_2 when deforming ω_1). We can use Lemma 4.1 to extend these deformations to all of \mathbb{R}^4 in such a way that points outside Ω never move and points in $\omega_1, \omega_2, \Omega_1, \Omega_2$ move only when they are supposed to. In the end we get a mapping θ with the right properties. As usual, we can simply define the ϕ_i 's as in Definition 2.3 from this construction of θ (or make unnecessary efforts to adjust θ to previous choices of the ϕ_i 's). The bottom line is that we have initial and excellent packages in this case, and so we get the associated decomposition G of \mathbb{R}^3 (as in Section 2) and the sets M, M^j constructed in Section 3. (We do not need \widetilde{M} for this example.)

It turns out that this decomposition is well behaved topologically, but for nontrivial reasons.

Theorem 6.2 (Bing [B1]) \mathbb{R}^3/G is homeomorphic to \mathbb{R}^3 , and in fact there exists a homeomorphism f from \mathbb{R}^3/G onto \mathbb{R}^3 which agrees with the "identity" on the complement of $D_1 \cup D_2$.

This result is given in [B1, Section 3, Paragraph III]. See also [B5].

Theorem 6.3. M satisfies (**).

Let a be chosen as in Lemma 3.65. Let $x \in M$ and $r > 0$ be given, so that we want to find a topological ball U and a parameterization of it which satisfy the conditions in Definition 1.7. If x, r satisfy i), iii), or v) in Lemma 3.44, then we are in business, because of Lemma 3.65. Suppose that we are in case iv) in Lemma 3.44. Because of (3.20) (with $j = \infty$, so that $M^{j-l} = M$) we can reduce to the case where $l = 1$. Note that there is nothing fishy going on here with the uniform estimates, because they are all chosen to behave properly under similarities. (Do not forget that the radius r changes also with the similarity.) Let h be as in Lemma 3.21, and choose $i = 1, 2$ so that $\Omega_\alpha \cap M = h(D_i)$, where Ω_α is as in iv) in Lemma 3.44 (and hence $\alpha \in \mathcal{S}_1$). Set $U = h(U_i)$, where U_i is as in (6.1). Lemma 3.21 tells us that h descends to a homeomorphism from \mathbb{R}^3/G onto M , and so \bar{U} is homeomorphic to \bar{U}_i/G . Our mapping f above (from Theorem 6.2) provides us with a homeomorphism from \bar{U}_i/G onto \bar{U}_i . The bottom line is that U is a topological ball whose closure is homeomorphic to a closed 3-ball. Since there are only two choices here there is no problem with getting the uniform estimates required in Definition 1.7, since continuous maps between compact sets are uniformly continuous. In other words, in case iv) we get our uniformity because the self-similarity (3.20) allows us to reduce to a finite number of mappings. The whole construction in Section 3 was designed to make this happen.

We are left with ii) in Lemma 3.44, which is slightly a nuisance but not deep. Set $r_1 = r_1(r) = 2r + \text{diam } \Omega$, so that

$$(6.4) \quad r_1 \leq C_0 r$$

(with $C_0 = 2 + a^{-1}$) by ii), and let h be as in Lemma 3.21 again. Notice that $B(x, r_1) \supseteq \Omega$, and that

$$(6.5) \quad h(B(x, r_1) \cap P) = ((B(x, r_1) \cap P) \setminus D) \cup (\Omega \cap M) \supseteq B(x, r) \cap M.$$

Setting $U = h(B(x, r_1) \cap P)$, we have that $B(x, r) \cap M \subseteq U = B(x, r_1) \cap M$.

Let F be the homeomorphism from \mathbb{R}^3 onto M obtained in the following manner. We know that h descends to a homeomorphism from \mathbb{R}^3/G onto M , and that f above is a homeomorphism from \mathbb{R}^3/G onto \mathbb{R}^3 , and we take F to be the composition of the former with the inverse of the latter. From the properties of h and f we get that F equals the identity off D and that F maps D onto $\Omega \cap M$. Thus $U = F(B(x, r_1))$, and U is a topological ball in particular. We need to check that we can parameterize U with the right kind of uniform estimates for the moduli of continuity.

Since F is a homeomorphism which equals the identity off the compact set D there is a locally bounded function $\zeta : [0, \infty) \rightarrow [0, \infty)$ such that $\lim_{t \rightarrow 0} \zeta(t) = 0$,

$$(6.6) \quad |F(y) - F(z)| \leq \zeta(|y - z|) \quad \text{when } y, z \in \mathbb{R}^3,$$

$$(6.7) \quad |F^{-1}(v) - F^{-1}(w)| \leq \zeta(|v - w|) \quad \text{when } v, w \in M,$$

and

$$(6.8) \quad \zeta(t) = 2t \quad \text{when } t > \text{diam } \Omega.$$

These estimates are not quite what we want, because they do not scale properly. The estimates that we need come down to the existence of a function $\xi : [0, \infty) \rightarrow [0, \infty)$ such that $\lim_{t \rightarrow 0} \xi(t) = 0$,

$$(6.9) \quad |F(y) - F(z)| \leq r \xi(r_1^{-1}|y - z|)$$

when $y, z \in B(x, r_1) \cap \mathbb{R}^3$, and

$$(6.10) \quad |F^{-1}(v) - F^{-1}(w)| \leq r_1 \xi(r^{-1}|v - w|)$$

when $v, w \in B(x, r) \cap M$.

It is important here that ξ not depend on r (or x); if not for this we could simply compute ξ from ζ . As it is, we are lead to try

$$(6.11) \quad \xi(t) = \sup\{r^{-1} \zeta(r_1 t) + r_1^{-1} \zeta(r t) : r \geq a \text{diam } \Omega\},$$

(where r_1 is still related to r as above). This is actually finite and locally bounded for all t , because of (6.8), but we need to check that $\lim_{t \rightarrow 0} \xi(t) = 0$. Let ε in $(0, 1)$ be given, and choose $\delta \in (0, 1)$ so that $\zeta(t) < \varepsilon$ when $t < \delta$. Suppose that $t < C_0^{-1} \varepsilon \delta$ (where C_0 is as in (6.4)), and let us show that $\xi(t)$ is small. Let ξ_1, ξ_2 , and ξ_3 be defined as in (6.11), but where the supremum is limited to the ranges $r > t^{-1} \text{diam } \Omega$, $\varepsilon^{-1} \leq r \leq t^{-1} \text{diam } \Omega$, and $a \text{diam } \Omega \leq r < \varepsilon^{-1}$. Thus ξ is the same as the maximum of the ξ_i 's. For ξ_1 we use (6.8) and (6.4) to get

$$(6.12) \quad \xi_1(t) \leq r^{-1} (2r_1 t) + r_1^{-1} (2r t) \leq C t \leq C \varepsilon.$$

For ξ_2 we use the fact that $\sup\{\zeta(s) : s \leq C_0 \text{diam } \Omega\} < \infty$ to obtain

$$(6.13) \quad \xi_2(t) \leq C \varepsilon.$$

As for ξ_3 we simply use (6.4) and the fact that r is bounded from below to get

$$(6.14) \quad \xi_3(t) \leq C \sup\{\zeta(s) : s \leq \delta\} \leq C \varepsilon.$$

Altogether we get that $\xi(t) \leq C\varepsilon$ when $t < C_0^{-1} \varepsilon \delta$, and so $\lim_{t \rightarrow 0} \xi(t) = 0$, as desired. This completes the proof of Theorem 6.3.

Theorem 6.15. *M does not admit a quasisymmetric parameterization by \mathbb{R}^3 .*

This is basically a small perturbation of [FS, Theorem 2.1]. This result in [FS] says that a certain discrete group of homeomorphisms on \mathbf{S}^3 is not homeomorphically conjugate to a group of quasiconformal mappings with uniformly bounded dilatation. The present story amounts to building a different metric space where a topologically equivalent form of this group acts uniformly quasiconformally, and we are concluding that this metric space cannot be quasiconformally equivalent to \mathbf{S}^3 . However it is easier to prove the theorem directly. The idea of the proof is that the sets $\Omega_\alpha \cap M$ all look alike and are reasonably well-shaped, while any homeomorphism from M to \mathbb{R}^3 has to twist at least some of these sets rather severely, and more so than a quasisymmetric map can.

To make this precise, let g be a homeomorphism from M onto \mathbb{R}^3 , let \mathcal{T}_l^0 denote the collection of subsets of M of the form $\overline{\Omega}_\alpha \cap M$ with $\alpha \in \mathcal{S}_l$, and let \mathcal{T}_l be the set of images of elements of \mathcal{T}_l^0 under g . We want to show that no matter how g is chosen the geometry of some of the elements of \mathcal{T}_l will have to degenerate as $l \rightarrow \infty$. This will come down to a lemma in [FS].

Let F be the homeomorphism described just after (6.5). Using F we obtain that the elements of \mathcal{T}_l^0 are homeomorphic to smooth solid tori in \mathbb{R}^3 when $l = 0, 1$, and they even have neighborhoods in M which are homeomorphic to solid tori. The same is true for all l because of the self-similarity property (3.20). Thus the elements of \mathcal{T}_l are locally flat topological solid tori in \mathbb{R}^3 . Let T denote the unique element of \mathcal{T}_0 . This is the solid torus in which all the action takes place, because it contains all the others.

If $\Gamma \in \mathcal{T}_l$, define $\text{length}(\Gamma)$ to be the infimum of the (parameterized) length of the loops in the interior of Γ which represent a generator in the fundamental group of Γ . The next lemma is a consequence of [FS, p. 81, Lemma 2.2].

Lemma 6.16. *There is a constant $C > 0$ so that*

$$(6.17) \quad 2^{-l} \sum_{\Gamma \in \mathcal{T}_l} \text{length}(\Gamma) \geq C^{-1},$$

for all l . In particular, for each l there is at least one $\Gamma \in \mathcal{T}_l$ such that $\text{length}(\Gamma) \geq C^{-1}$.

The second statement follows from the first, because \mathcal{T}_l has 2^l elements (since $n = 2$ in this example). The idea behind the first part is pretty simple. Fix an l , and choose for each $\Gamma \in \mathcal{T}_l$ a loop γ in the interior of Γ which represents a generator of the fundamental group of Γ . Consider the combination of all 2^l of these γ 's. The point is that every time we increase l by 1 we get to double the number of γ 's, but we also have to double the number of times that this system of curves “goes around” in T . (“Goes around” should be interpreted geometrically, and not homologically, because these curves do not go around in T homologically at all.) It is easy to believe this after staring at the pictures, and a proof is given in [FS].

Lemma 6.18. *Assume that $g : M \rightarrow \mathbb{R}^3$ is actually quasimetric. Then there is a constant $C > 0$ so that $\text{length} \Gamma \leq C \text{diam} \Gamma$ for all $\Gamma \in \mathcal{T}_l$ and any l .*

Let $\Gamma \in \mathcal{T}_l$ be given, and let $\alpha \in \mathcal{S}_l$ be chosen so that $\Gamma = g(\overline{\Omega}_\alpha \cap M)$. Let us first choose a loop γ_0 in $\Omega_\alpha \cap M$ which represents a generator of its fundamental group and which comes from some fixed smooth curve in $\Omega \cap M$. That is, we first make a nice choice of such a curve (call it τ) in $\Omega \cap M$; and then we use the self-similarity property (3.20) and take $\gamma_0 = \psi_\alpha(\tau)$. Thus γ_0 is smooth at the scale of $\varepsilon \text{diam} \Omega_\alpha$ and $\text{dist}(\gamma_0, M \setminus \Omega_\alpha) \geq \delta \text{diam} \Omega_\alpha$ for fixed $\varepsilon, \delta > 0$ which do not depend on α or l .

Set $\gamma = g(\gamma_0)$. This represents a generator in the fundamental group of Γ , but it may have infinite length. However, there is a fixed $\delta' > 0$, which does not depend on Γ , such that $\text{dist}(\gamma, \mathbb{R}^3 \setminus \Gamma) \geq \delta' \text{diam} \Gamma$. This follows from the corresponding property of γ_0 and the quasimetric condition. This gives us enough room to deform γ inside Γ to a loop with length less or equal than $C \text{diam} \Gamma$, as desired. (To be honest, to check this carefully one should notice that γ cannot oscillate too many times at the scale of $(\delta'/10) \text{diam} \Gamma$, say, because of the smoothness property of γ and the quasimetric of g . This

implies that we only need to make a bounded number of modifications to γ in a bounded number of little balls. Note that the condition $\text{dist}(\gamma, \mathbb{R}^3 \setminus \Gamma) \geq \delta' \text{diam } \Gamma$ does not prevent γ from looping around many, many times in a little neighborhood of itself, even though we know that this cannot happen here.)

Theorem 6.15 follows from Lemmas 6.16 and 6.18, because $\text{diam } \Gamma$ tends to 0 as $l \rightarrow \infty$ when $\Gamma \in \mathcal{T}_l$ (since the corresponding statement is true for \mathcal{T}_l^0 and g is continuous). Theorem 1.12 .c) now follows from Theorems 6.3 and 6.15 together with Lemma 3.45 and Proposition 3.70.

REMARK 6.19. I have a philosophical explanation for Theorem 6.15 which I cannot back up with a proof but which I would like to share with the reader. The self-similarity properties of M stem from the fact that we constructed Σ as in (3.14) so that its big boundary component is similar to each of its small boundary components. The corresponding part of our initial package is the set $D \setminus \cup_j \overline{D}_j$, and one reason that we cannot build a quasisymmetric parameterization of M is that $D \setminus \cup_j \overline{D}_j$ does not possess a version of this property. Specifically, if ∂D were conformally equivalent to each ∂D_j , then we could try to build a quasiconformal parameterization of M by building a suitable quasiconformal map from $D \setminus \cup_j \overline{D}_j$ onto Σ and putting copies of it on top of itself. In order to have the dilatation not build up and become unbounded in the limit we need this quasiconformal building block to be conformal at the ends, which is why we would need to have ∂D be conformally equivalent to each ∂D_j . However, it is not true that 2-dimensional tori are all conformally equivalent, and this is the source of the problem. I do not know how to turn this explanation into an alternative proof of Theorem 6.15, but I think that it would be interesting to do so, especially for the purpose of understanding other examples of this type. (See also Section 12.)

7. The complementary components, part 1.

Throughout this section we assume that we are given an excellent package as in Definition 3.2, and we shall consider the behavior of the complementary components of M , the M^j 's, and \widetilde{M} . We shall use freely the definitions and notation of Section 3.

Let U^+ and U^- denote the two components of $\mathbb{R}^1 \setminus P$, i.e., U^+ is

the set of points $x \in \mathbb{R}^4$ such that $x_4 > 0$, and similarly for U^- . Define X^+ and X^- by

$$(7.1) \quad X^\pm = \left(\overline{\Omega} \setminus \bigcup_{j=1}^n \omega_j \right) \cap U^\pm.$$

In this section we require that

$$(7.2) \quad X^\pm, \Omega \cap U^\pm, \text{ and } U^\pm \setminus \Omega \text{ are connected.}$$

This condition is not necessarily minimal, but it is valid in the examples in Sections 4, 5, and 6 and others like them, and it is nicely clear.

Lemma 7.3. *Each of M , the M^j 's, and \widetilde{M} has exactly two complementary components in \mathbb{R}^4 .*

For \widetilde{M} and the M^j 's this is an immediate consequence of Lemma 3.21. For M we have to be slightly more careful.

Let $g_l : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ be as in the proof of Lemma 3.21, $l = 0, 1, 2, \dots$ and define a mapping $e : \mathbb{R}^4 \setminus M \rightarrow \mathbb{R}^4$ by

$$(7.4) \quad e(x) = \lim_{l \rightarrow \infty} g_0^{-1} \circ g_1^{-1} \circ \dots \circ g_l^{-1}.$$

This limit exists because $g_k(x) = x$ for any $x \in \mathbb{R}^4 \setminus M$ and all sufficiently large k , by (3.26). Also, if h_l is as in (3.27),

$$(7.5) \quad h_l(e(x)) = x,$$

for any $x \in \mathbb{R}^4 \setminus M$ and all sufficiently large l . In particular we have that e maps $\mathbb{R}^4 \setminus M$ into $\mathbb{R}^4 \setminus P$. Since e equals the identity outside Ω we conclude that $\mathbb{R}^4 \setminus M$ has at least two components.

Let us analyze $\mathbb{R}^4 \setminus M$ a little more in order to show that it has at exactly two components. Recall the definition of Y (shortly before (3.12)) and set $Y^+ = \theta(X^+)$, $Y^- = \theta(X^-)$. Set $Y_\alpha^+ = \psi_\alpha(Y^+)$ and $Y_\alpha^- = \psi_\alpha(Y^-)$ for any α in any \mathcal{S}_l , and define V^+ and V^- by

$$(7.6) \quad V^\pm = (U^\pm \setminus \Omega) \cup \bigcup_{l=0}^{\infty} \bigcup_{\alpha \in \mathcal{S}_l} Y_\alpha^\pm.$$

It is easy to see that

$$(7.7) \quad \mathbb{R}^4 \setminus M = V^+ \cup V^-,$$

because of (3.13) and (3.19). Let us check that

$$(7.8) \quad V^+ \text{ and } V^- \text{ are connected.}$$

To see this note that if $\alpha \in \mathcal{S}_l$, $l \geq 1$, and δ is the parent of α , then

$$(7.9) \quad Y_\alpha^+ \text{ touches } Y_\delta^+$$

(at their respective boundaries), and similarly for Y_α^- . When $l = 1$ (7.9) reduces easily to the fact that the ψ_j 's from Definition 3.2 are required to be orientation-preserving. The general case can be reduced to this one by using the fact that $\psi_\alpha = \psi_\delta \circ \psi_j$ for some j (as in the definition of ψ_α just before (3.7)). We also have that

$$(7.10) \quad Y^+ \text{ touches } U^+ \setminus \Omega$$

(at their boundaries), and similarly for Y^- , and (7.8) follows easily from these observations and our assumption (7.2). This proves Lemma 7.3.

Note that the complementary domains of M^j also admit an expression like (7.7). That is, if we define V_j^+ and V_j^- by

$$(7.11) \quad V_j^\pm = (U^\pm \setminus \Omega) \cup \left(\bigcup_{l=0}^{j-1} \bigcup_{\alpha \in \mathcal{S}_l} Y_\alpha^\pm \right) \cup \left(\bigcup_{\alpha \in \mathcal{S}_j} \Omega_\alpha \cap U^\pm \right),$$

then

$$(7.12) \quad \mathbb{R}^4 \setminus M^j = V_j^+ \cup V_j^-,$$

because of (3.12) and (3.17), and V_j^\pm are connected for the same reasons as for (7.8). Thus V_j^\pm are the two complementary components of M^j .

Next we want to show that these various complementary domains are *uniform domains*. Recall that a domain O in \mathbb{R}^4 is a uniform domain if there exists a constant C so that for each pair of points $x, y \in O$ we can find a path Γ in O which connects x and y and satisfies

$$(7.13) \quad \text{diam } \Gamma \leq C |x - y|$$

and

$$(7.14) \quad \text{dist}(z, \partial O) \geq C^{-1} \text{dist}(z, \{x, y\}) \quad \text{when } z \in \Gamma.$$

(There are many equivalent characterizations of uniform domains, and this condition is a little weaker in appearance than some of the others.) Bounded smooth domains are uniform domains, but a domain with an outward-pointing cusp is not. Note that this condition is scale-invariant.

Lemma 7.15 *The complementary components of M , the M^j 's, and \widetilde{M} in \mathbb{R}^4 are uniform domains, with constants bounded independently of j (in the case of the M^j 's).*

This is a fairly straightforward exercise. We could derive this lemma from the general results in [V2], but let us give instead a direct argument. Consider M and the M^j 's and their complementary components V^+ and V_j^+ . Let us call sets of the form $U^+ \setminus \Omega$, Y_α^+ ($\alpha \in \mathcal{S}_l, l < j$), and $U^\pm \cap \Omega_\alpha$ ($\alpha \in \mathcal{S}_j$) *building blocks* for V^+ or V_j^+ , as appropriate. Notice that these building blocks are all uniform domains, with uniformly bounded constant. In the case of $U^+ \setminus \Omega$ this is true because it has smooth boundary and looks like a half-space at infinity. For the other building blocks this uniform estimate follows from the fact that they are all similar to one of a finite number of bounded smooth domains. The union of two building blocks which touch is also a uniform domain with bounded constant, for the same reasons. (Remember also that a Y_α and a Y_β can touch only when one of α and β is the parent of the other, as observed just before (3.12).)

Suppose now that we are given two points x, y in V^+ or V_j^+ . We want to connect them with a curve which satisfies (7.13) and (7.14) (with $O = V^+$ or V_j^+ , as appropriate). If x and y lie in the same building block then we are in business, by the preceding remarks, and also when they lie in different building blocks which touch. Thus we may assume that x and y lie in disjoint building blocks. Consider first the case where $x \in Y_\alpha^+$ and $y \in Y_\beta^+$ for some $\alpha \neq \beta$, $\alpha \in \mathcal{S}_l, \beta \in \mathcal{S}_k, k, l < j$. Our assumption of disjointness implies that neither of α or β is the parent of the other. Let δ be the "last" common ancestor of α and β , so that either $\alpha = \delta, \beta = \delta$, or α and β are descended from different children of δ . Notice that

$$(7.16) \quad \text{dist}(Y_\alpha^+, Y_\beta^+) \leq |x - y|.$$

$$(7.17) \quad \text{diam } Y_\alpha^+ + \text{diam } Y_\beta^+ \leq 2 \text{diam } \Omega_\delta \leq C \text{dist}(Y_\alpha^+, Y_\beta^+).$$

(The last inequality reduces via the similarity ψ_δ to the fact that the distance between the children of Ω is positive (when δ is different from

both α and β), and to the fact that the distance between Y and any of the Ω_η 's with $\eta \in \mathcal{S}_2$ is positive (when δ equals one of α and β). We get a uniform estimate because the similarity allows us to reduce to a finite number of cases.) We can build a curve which goes from x to ∂Y_α^+ , then ascends to ∂Y_γ^+ for the various successive ancestors γ of α until we get to δ , and then goes down to ∂Y_γ^+ for the successive descendants γ of δ until we get to ∂Y_β^+ , from which we can connect easily to β . This curve will satisfy (7.13) because it will be contained in Ω_δ and because $\text{diam } \Omega_\delta \leq C|x - y|$, by (7.16) and (7.17). If we are a little bit careful we can choose the curve so that (7.14) also holds. (The point is to stay as far away from $\partial Y_\gamma^+ \cap M^j$ as possible - *i.e.*, at distance greater or equal than $C^{-1} \text{diam } Y_\gamma^+$ - for all the intermediate γ 's. In Y_α^+ and Y_β^+ it may be necessary to let the curve get close to the boundary of V^+ or V_j^+ because of the positions of x and y .) The remaining possible situations where x and y lie in disjoint building blocks (*e.g.*, $x \in U^+ \setminus \Omega$, or $y \in U^+ \cap \Omega_\alpha$ for an $\alpha \in \mathcal{S}_j$) are handled in essentially the same way.

Thus one can show that V^+ and the V_j^+ 's are uniform domains, with bounded constants. The same argument works for the V^- 's. It is easy to show that the complementary components of \tilde{M} are uniform domains using the corresponding statement for the M^j 's and the construction of \tilde{M} (just after (3.20)). For this one should go through the usual analysis of cases, *i.e.*, the cases where the given pair of points both lie in the same $2B_k$, or they both lie outside all the B_k 's, or they lie in different $2B_k$'s, or one lies in some $2B_k$ and the other lies outside all the other B_l 's. The details are left to the reader. This completes the proof of Lemma 7.15.

In our examples we also have that there exist bilipschitz reflections across M , the M^j 's, and \tilde{M} . To see this we first define a suitable symmetry condition for an excellent package.

Definition 7.18. *An excellent package as in Definition 3.2 is said to be symmetric if Ω , the ω_j 's and the Ω_j 's are symmetric about P , and if the restriction of θ to a neighborhood of each $\bar{\omega}_j$ commutes with the obvious reflection about P .*

It is important here that we do not require that θ commute with the reflection about P everywhere, because that will not be true in the interesting examples.

Lemma 7.19. *We can choose the excellent packages in Sections 4, 5, and 6 to be symmetric.*

This is straightforward but slightly tedious. There is absolutely no problem about making the various Ω 's and ω_j 's be symmetric, but the symmetry condition on θ is slightly more complicated. The main point is that we always obtained θ by first building suitable isotopies on the ω 's and then extending them to all of Ω using Lemma 4.1. The first step is the only one that really matters here and is the one which is most easily controlled. It is easy to check that it can be carried out in such a manner as to get the desired symmetry condition. In the examples in Sections 5 and 6, for instance, these isotopies on the ω 's could be taken to be translations in the positive x_4 direction, followed by a certain "unwinding" operation centered on the relevant vertical translate P' of P , followed by a translation back down to P . One need only demand that this middle unwinding operation be symmetric with respect to P' , which is easily accomplished, because this unwinding operation really comes from a 3-dimensional process. The example in Section 4 is a bit different, because one translates part of ω up a little while leaving the other part alone, but it is also not difficult to handle. We leave the details as an exercise. (Keep in mind that the intermediate stages of the deformation do not have to be symmetric about P , only the end result.)

Proposition 7.20. *If our excellent package is symmetric, then there exists a bilipschitz reflection r on \mathbb{R}^4 across M , and there exist reflections r_j across each M^j which are uniformly bilipschitz. These reflections all agree with the standard reflection across P outside Ω . There is also a bilipschitz reflection across \tilde{M} .*

Let τ denote the standard reflection across P , and for any α in any S_l define σ_α by

$$(7.21) \quad \sigma_\alpha = \psi_\alpha \circ \theta \circ \tau \circ \theta^{-1} \circ \psi_\alpha^{-1}.$$

This is the same as taking $\theta \circ \tau \circ \theta^{-1}$, which agrees with τ outside Ω and on a neighborhood of its boundary (by (3.5)) but does something complicated inside, and then transporting it to Ω_α using ψ_α . Our symmetry assumptions in Definition 7.18 imply that $\theta \circ \tau \circ \theta^{-1}$ also agrees with τ on a neighborhood of each $\bar{\Omega}_i$, and hence

$$(7.22) \quad \sigma_\alpha = \tau \text{ on } \mathbb{R}^4 \setminus Y_\alpha \text{ and on a neighborhood of } \partial Y_\alpha$$

(where Y_α is as defined just before (3.12)). Also,

$$(7.23) \quad \sigma_\alpha(Y_\alpha^+) = Y_\alpha^- \quad \text{and} \quad \sigma_\alpha(Y_\alpha^-) = Y_\alpha^+,$$

by definition of Y_α^\pm , and

$$(7.24) \quad \sigma_\alpha(x) = x \quad \text{for all } x \in \Sigma_\alpha,$$

because of the definition of Σ_α (just before (3.15)) and the fact that τ fixes every element of P .

Define r by taking r to be the identity on M , $r = \tau$ on $\mathbb{R}^4 \setminus \Omega$, and $r = \sigma_\alpha$ on each Y_α . Define r_j to be the identity on M^j , $r_j = \tau$ on $\mathbb{R}^4 \setminus \Omega$, and $r_j = \sigma_\alpha$ on each Y_α with $\alpha \in S_l$, $l < j$, and $r_j = \tau$ on Ω_α when $\alpha \in S_j$. Clearly r and the r_j 's are involutions, and one can use the formulae (7.6) and (7.11) for the complementary components of M and M^j and (7.23) to show that r and the r_j 's exchange the complementary components of M and M^j . It is not hard to check that r and the r_j 's are continuous (and even smooth in the case of the r_j 's), because of the way that the Ω_α 's fit together, and because of (7.22). One should be a little bit careful about the continuity of r and the r_j 's across M and the M^j 's, but the only slightly tricky issue is the continuity of r at points in the singular set F of M . For this it is useful to observe that

$$(7.25) \quad r(\Omega_\alpha) = \Omega_\alpha,$$

for all α in any S_l . This observation is easy to derive from the definition of r and the simpler fact that $r(Y_\alpha) = Y_\alpha$ for all α . (Note that the analogue of (7.25) for the r_j 's holds as well.)

The uniform Lipschitz conditions are easy to check. The main point is that the σ_α 's are uniformly Lipschitz. For the r_j 's this is enough because the smoothness of the r_j 's permits their uniform Lipschitzness to be derived from a bound on their gradients, which we get from the corresponding bounds for the σ_α 's. In the case of r one must be a little more careful, since it is not smooth across M , but the continuity of r across M allows one to piece the local Lipschitz conditions together. Uniform bilipschitzness follows from uniform Lipschitzness and the fact that these mappings are involutions (and hence are their own inverses).

As usual, one can deal with \widetilde{M} by treating separately the little pieces which look like M^j 's (as in the definition of \widetilde{M} , just after (3.20)).

RECAPITULATION 7.26. Under mild additional conditions on our excellent package which are satisfied in the examples, we obtain that M , the M^j 's, and \tilde{M} have exactly two complementary components which are uniform domains, and there are uniformly bilipschitz reflections across the M 's.

8. The complementary components, part 2.

Throughout this section we use the excellent package described in Section 6, taking a modicum of care to ensure that the additional requirements of the preceding section ((7.2) and symmetry) are met. We shall use freely the definitions and notation of Section 3, and in particular we assume that M has been constructed as in Section 3. Let V^\pm be the complementary domains of M , as in (7.6) and (7.7).

Theorem 8.1. a) V^\pm are uniform domains, and there is a bilipschitz reflection on \mathbb{R}^4 which equals the identity outside Ω , fixes every point on M , and interchanges V^+ and V^- . Neither V^+ nor V^- is quasiconformally equivalent to a ball (or a half-space).

b) There is a homeomorphism ν from \mathbb{R}^4 onto itself such that $\nu =$ the identity outside Ω , $\nu = \theta$ outside $\omega_1 \cup \omega_2$, $\nu(\omega_i) = \Omega_i$ for $i = 1, 2$, and ν maps P onto M .

c) There exist a constant $C > 0$ and a locally bounded function $\eta : [0, \infty) \rightarrow [0, \infty)$ with $\lim_{t \rightarrow 0} \eta(t) = 0$ such that for each $x \in M$ and $r > 0$ there is an open set $W \subseteq \mathbb{R}^4$ with $B(x, r) \subseteq W \subseteq B(x, Cr)$ and a homeomorphism μ from $B(0, 1)$ onto W which satisfy $\mu(B(0, 1) \cap P) = M \cap W$,

$$(8.2) \quad |\mu(y) - \mu(z)| \leq r \eta(|y - z|), \quad \text{for all } y, z \in B(0, 1),$$

and

$$(8.3) \quad |\mu^{-1}(p) - \mu^{-1}(q)| \leq \eta(r^{-1}|p - q|), \quad \text{for all } p, q \in U.$$

Part c) is a stronger version of (**) for M which takes into account the ambient space as well. It implies that V^+ and V^- are *strongly uniform* domains in the sense of [HY]. This follows from the compactness result in [HY, Theorem 4.4].

One should keep in mind that the boundaries of these “Bing domains” are Ahlfors regular (at least if we choose ρ small enough, as in Lemma 3.45), uniformly rectifiable (in the sense of [DS4]), and have well-behaved geodesic distance functions (as in Theorem 1.12 and Lemma 3.69). One can construct much simpler examples which satisfy the properties in Theorem 8.1 using the methods of [Tu2], but the boundaries of these examples are not so well behaved.

The fact that V^\pm are uniform domains and the existence of the bilipschitz reflection were proved in the preceding section (Lemma 7.15 and Proposition 7.20). It follows that V^\pm cannot be quasiconformally equivalent to a ball, since, as pointed out in the introduction, well-known results about uniform domains then imply that the quasiconformal map would extend to a quasisymmetric map between the closures, which would contradict the fact that M is not quasisymmetrically equivalent to \mathbb{R}^3 . This gives a).

Let us assume b) for the moment and derive c). We use the same argument as in the proof of Theorem 6.3. We choose a as in Lemma 3.65, and we consider separately the cases i)-v) in Lemma 3.44. The cases i), iii), and v) are covered by Lemma 3.65. We shall use the next lemma to deal with case iv), and we shall consider ii) after that.

Lemma 8.4. *There exist open subsets O_1 and O_2 of Ω such that $\Omega_i \subseteq O_i$ for $i = 1, 2$ and each \overline{O}_i is homeomorphic to $\overline{B}(0, 1)$ via a mapping which sends $M \cap O_i$ onto $P \cap B(0, 1)$.*

Set $O'_i = U_i \times (-\varepsilon, \varepsilon)$, where U_i and ε are as in (6.1) and (6.1'). Clearly each \overline{O}'_i is homeomorphic to $\overline{B}(0, 1)$ via a mapping which takes P to itself, since the U_i 's are closed topological 3-balls. Thus if ν is as in b) then $O_i = \nu(O'_i)$, $i = 1, 2$, are open subsets of Ω such that $O_i \supseteq \Omega_i$ and each \overline{O}_i is homeomorphic to $\overline{B}(0, 1)$ via a mapping which sends M to P . This proves Lemma 8.4.

Now suppose that x, r satisfy iv) in Lemma 3.44. If the l from iv) equals 1, then we can take W to be one of the O_i 's from Lemma 8.4. When $l > 1$ we can reduce to the $l = 1$ case using the self-similarity property (3.20). These choices of W admit homeomorphic parameterizations with the right kind of equicontinuity conditions because they all reduce to the two (uniformly continuous) models in Lemma 8.4 by self-similarity.

The remaining case, where x, r satisfies ii), is handled in exactly

the same way as the corresponding step in the proof of Theorem 6.3 (with ν playing the role of F). That is, one must go through the song-and-dance of (6.11), but that part of the argument applies equally well to the current situation. This proves c) given b).

We are left with proving b). Recall that G is the decomposition of \mathbb{R}^3 which is associated to the present initial package as in Section 2 (just after Definition 2.2). (See also the discussion around Sublemma 3.40.) Let G' denote the decomposition of \mathbb{R}^4 which extends G trivially. That is, G' consists of the elements of G (viewed as subsets of \mathbb{R}^4) together with the singletons from $\mathbb{R}^4 \setminus P$. Our first task in proving b) is to come to grips with \mathbb{R}^4/G' .

Proposition 8.5. *There is a homeomorphism ξ from \mathbb{R}^4/G' onto \mathbb{R}^4 which sends P/G onto P and which equals "the identity" outside $\omega_1 \cup \omega_2$.*

In other words, there is a nice extension of the homeomorphism given by Theorem 6.2 to \mathbb{R}^4 . To prove this we use the following.

Lemma 8.6. *There is a continuous 1-parameter family f_t , $t \in [0, 1]$, of continuous mappings from \mathbb{R}^3 onto itself such that f_0 is the identity, each f_t agrees with the identity outside $D_1 \cup D_2$ (where D_1 and D_2 come from our initial package), each f_t for $t < 1$ is a homeomorphism, and f_1 induces a homeomorphism from \mathbb{R}^3/G onto \mathbb{R}^3 .*

Thus the homeomorphism in Theorem 6.2 can be deformed to the identity in a natural way. This follows from Bing's construction of a homeomorphism as in Theorem 6.2. Specifically, Bing produces an increasing sequence of integers $\{j_i\}$ with $j_i \geq 1$ and homeomorphisms $\{T_i\}$ on \mathbb{R}^3 such that T_i maps the \overline{D}_α 's with $\alpha \in \mathcal{S}_{j_i}$ to sets of diameter less than $\frac{1}{i}$ and $T_i = T_{i-1} \circ T'_i$, where T'_i takes each \overline{D}_β to itself when $\beta \in \mathcal{S}_{j_{i-1}}$ and T'_i equals the identity outside of all these D_β 's. Each of the T'_i 's can be deformed continuously to the identity through homeomorphisms which also take each of these D_β 's to themselves and equal the identity off of the D_β 's. This follows from the construction; Bing obtains T'_i by sliding the D_α 's around, $\alpha \in \mathcal{S}_{j_i}$, without ever doing anything outside the D_β 's. (The point is to slide the D_α 's around to make them have very small diameter.) This sequence $\{T_i\}$ converges to a mapping T which induces a homeomorphism from \mathbb{R}^3/G onto \mathbb{R}^3 . Indeed, by construction $T_l = T_i$ on the complement of the D_α 's for $\alpha \in \mathcal{S}_{j_i}$ when $l \geq i$, and $T_l(D_\alpha) = T_i(D_\alpha)$ has diameter less than $1/i$

for each $\alpha \in S_{j_i}$, and these properties imply that $\{T_i\}$ converges to a mapping which induces a homeomorphism from \mathbb{R}^3/G onto \mathbb{R}^3 .

To get the deformation described in Lemma 8.6 we take f_0 to be the identity, f_1 to be T , f_t to be T_i when $t = i/(i + 1)$, and on $(i-1)/i < t < i/(i+1)$ we deform T_{i-1} into T_i through homeomorphisms by deforming the identity to T'_i through homeomorphisms which map each D_α to itself for $\alpha \in S_{j_i}$ and which equal the identity off these D_α 's. (For $i = 1$ this makes sense with T_0 taken to be the identity.) It is not hard to check that this deformation has the properties described in Lemma 8.6.

Proposition 8.5 is a straightforward consequence of Lemma 8.6. Let $b > 0$ be such that $\overline{D}_i \times [-b, b] \subseteq \omega_i$ for $i = 1, 2$. We define ξ by setting it to be the identity when $|x_4| > b$ and by taking ξ to be the obvious copy of f_t on the $x_4 = \pm b(1 - t)$ 3-planes when $0 \leq t \leq 1$. That is, ξ should take these 3-planes to themselves, and be the same as f_t modulo the obvious vertical translation down to P and back. It is easy to check that this choice of ξ has the right properties. This proves Proposition 8.5.

To prove b) in Theorem 1.2 it is enough to produce a mapping $\zeta : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ which agrees with θ outside $\omega_1 \cup \omega_2$, maps P onto M , and induces a homeomorphism from \mathbb{R}^4/G' onto \mathbb{R}^4 . Indeed, if we can build such a mapping ζ , then the desired ν will result from Proposition 8.5. We would like to simply take $\zeta = h$, where h is as in Lemma 3.21, but it is not completely clear that h has the right properties on \mathbb{R}^4 , i.e., $h^{-1}(\Omega_\alpha)$ might leak out into $\mathbb{R}^4 \setminus P$ further than we want even when $\alpha \in S_l$ for l large. Rather than attempt some fine analysis we shall modify h a little bit brutally to avoid this problem. Note that h has some nice self-similarity properties that we do not need to replicate in ζ .

Given $\alpha \in S_l$ define ω_α to be $h^{-1}(\Omega_\alpha)$. This is compatible with the choice of ω_j 's in our initial package, because of Lemma 3.21 and (3.6), and we could have defined the ω_α 's through the same kind of recursive constructions as in Section 3, but this amounts to the same thing as the present definition. Set $C_l^* = \cup_{\alpha \in S_l} \overline{\omega}_\alpha$, in analogy with the definition of the defining sequence $\{C_l\}$ associated to the D_α 's (just before (3.11)). Set $C^* = \cap_l C_l^*$, and let G^* denote the decomposition of \mathbb{R}^4 associated to the defining sequence $\{C_l^*\}$ in the manner described just after Definition 2.2.

Lemma 8.7. *h is constant on each element of G^* , and it induces a homeomorphism from \mathbb{R}^4/G^* onto \mathbb{R}^4 .*

The proof of this is practically the same as for the corresponding statement for \mathbb{R}^3/G in Lemma 3.21. Let us briefly review the highlights. The ω_α 's (and their closures) enjoy the same nesting properties as do the Ω_α 's (as described in Lemma 3.10). Let \mathcal{S} be, as before (just after (3.11)), the set of infinite sequences which take values in $\{1, \dots, n\}$. To each element s of \mathcal{S} we associate a set A_s^* which is the intersection of the $\bar{\omega}_\alpha$'s which correspond to the ancestors $\alpha \in \mathcal{S}_l$ of the sequence s , just as for the D_α 's (before Sublemma 3.40).

Sublemma 8.8. *$C^* = \bigcup_{s \in \mathcal{S}} A_s^*$, and the A_s^* 's are the connected components of C^* .*

This is the analogue of Sublemma 3.40 for the ω_α 's, and the same proof works here, except that we should verify that the $\bar{\omega}_\alpha$'s are connected. A priori we have a problem, since h is not a homeomorphism, but in fact we have that

$$(8.9) \quad \omega_\alpha = h_l^{-1}(\Omega_\alpha), \quad \text{for all } \alpha \in \mathcal{S}_l,$$

where h_l is as in the proof of Lemma 3.21. This equality is an easy consequence of the definitions in the proof of Lemma 3.21, but let us write it out. Notice first that

$$(8.10) \quad h \circ h_j^{-1} = \text{the identity on } \mathbb{R}^4 \setminus \bigcup_{\delta \in \mathcal{S}_j} \Omega_\delta,$$

and

$$(8.11) \quad h \circ h_{l+1}^{-1}(\bar{\Omega}_\beta) = \bar{\Omega}_\beta, \quad \text{for all } \beta \in \mathcal{S}_{l+1}.$$

These follow immediately from (3.28), (3.29), and the definition of h as the limit of the h_k 's. (Actually, one should think a little about the inclusion " \supseteq " in (8.11). To derive this inclusion from its counterpart in (3.29) one can use the compactness of $\bar{\Omega}_\beta$ and the uniform convergence of the h_k 's to h .) Even though h is not a homeomorphism we can convert (8.11) into $(h \circ h_{l+1}^{-1})^{-1}(\bar{\Omega}_\beta) = \bar{\Omega}_\beta$ using (8.10) and the fact that the $\bar{\Omega}_\beta$'s, $\beta \in \mathcal{S}_{l+1}$, are disjoint (Lemma 3.10). Since $\bar{\Omega}_\beta \subseteq \Omega_\alpha$ when β is a child of α (by (3.9)) we conclude that

$$(8.12) \quad (h \circ h_{l+1}^{-1})^{-1}(\Omega_\alpha) = \Omega_\alpha \quad \text{when } \alpha \in \mathcal{S}_l$$

using also (8.10) again. From (3.27) we have that $h_{l+1} = g_l \circ h_l$, where g_l is as in the proof of Lemma 3.21, and from (3.25) we obtain $g_l(\Omega_\alpha) = \Omega_\alpha$ when $\alpha \in \mathcal{S}_l$. Thus (8.12) yields $(h \circ h_l^{-1})^{-1}(\Omega_\alpha) = (h \circ h_{l+1}^{-1} \circ g_l)^{-1}(\Omega_\alpha) = \Omega_\alpha$. This proves (8.9).

It follows that the $\bar{\omega}_\alpha$'s are connected (since each h_l is a homeomorphism), which is the fact we needed for Sublemma 8.8.

Recall the (bijective) mapping $f : \mathcal{S} \rightarrow F$ defined just after (3.11). We have that

$$(8.13) \quad h(p) = f(s), \quad \text{for all } p \in A_s^* \text{ and } s \in \mathcal{S}.$$

This is the analogue of (3.41) in this situation, and it follows from chasing definitions. Because of Sublemma 8.8 this says exactly that h is constant on each component of C^* , and hence on the nontrivial elements of G^* . Thus h induces a continuous map from \mathbb{R}^4/G^* into \mathbb{R}^4 . It is not hard to see that it is actually a bijection and even a homeomorphism, using (8.13), (8.10), (8.11), (3.13), and the fact that the h_j 's are all homeomorphisms. This proves Lemma 8.7.

Because of Lemma 8.7, the proof of Theorem 8.1.b) comes down to showing that the decompositions G^* and G' of \mathbb{R}^4 are equivalent in a way which does not disturb points in P . Let us first check that they "agree" on P . To do this we begin with the observation that

$$(8.14) \quad \omega_\alpha \cap P = D_\alpha, \quad \text{for all } \alpha \in \mathcal{S}_l$$

(and any l). By (8.9) the left side is the same as $h_l^{-1}(\Omega_\alpha) \cap P$, and since $h_l(P) = M^l$ this is the same as $h_l^{-1}(\Omega_\alpha \cap M^l)$. Since $h_l(D_\alpha) = \Omega_\alpha \cap M^l$, by Lemma 3.21, we get (8.14). Thus if C is as in Sublemma 3.40 and C^* is as in Sublemma 8.8, then $C = C^* \cap P$, because the corresponding statement is true for the defining sequences $\{C_l\}$ and $\{C_l^*\}$, by (8.14). Moreover, if A_s and A_s^* are as in Sublemmas 3.40 and 8.8, then

$$(8.15) \quad A_s = A_s^* \cap P, \quad \text{for all } s \in \mathcal{S},$$

by the definitions of A_s and A_s^* . This makes precise the sense in which G^* and G' "agree" on P , since the A_s 's and A_s^* 's are the only nontrivial elements of these decompositions.

Lemma 8.16. *There is a homeomorphism σ from \mathbb{R}^4/G^* onto \mathbb{R}^4/G' which “equals the identity” on P and on the complement of $\omega_1 \cup \omega_2$, and which sends ω_i/G^* to ω_i/G' for $i = 1, 2$.*

The statement of the lemma is an abuse of language whose meaning is hopefully clear to the reader. It is “justified” by (8.15) and the fact that both the decompositions G^* and G' are trivial (consisting only of singletons) on the complement of $\omega_1 \cup \omega_2$.

To prove Lemma 8.16 we shall construct a sequence of diffeomorphisms $\{\sigma_l\}_{l=1}^\infty$ on \mathbb{R}^4 with the following properties (for all l):

$$(8.17) \quad \sigma_l = \text{the identity on } P \cup (\mathbb{R}^4 \setminus (\omega_1 \cup \omega_2)), \text{ and} \\ \sigma_l(\omega_i) = \omega_i, \text{ for } i = 1, 2;$$

$$(8.18) \quad \sigma_{l+1} = \sigma_l \quad \text{on } \mathbb{R}^4 \setminus \bigcup_{\alpha \in \mathcal{S}_l} \omega_\alpha;$$

$$(8.19) \quad \sigma_{l+1}(\omega_\alpha) = \sigma_l(\omega_\alpha), \quad \text{for all } \alpha \in \mathcal{S}_l;$$

$$(8.20) \quad \text{every point in } \sigma_l(\overline{\omega}_\alpha) \text{ lies within } l^{-1} \text{ of } \overline{D}_\alpha, \\ \text{for all } \alpha \in \mathcal{S}_l \text{ when } l \geq 2.$$

We take σ_1 to be the identity, and to construct the later σ_l 's we use an iterative construction based on the following.

Sublemma 8.21. *For each $\varepsilon > 0$ and any $\alpha \in \mathcal{S}_l$ (with l arbitrary) we can find a diffeomorphism of \mathbb{R}^4 onto itself which equals the identity on P and on the complement of ω_α and which maps $\overline{\omega}_\beta$ to a set which lies within ε of \overline{D}_β for each of the two children β of α in \mathcal{S}_{l+1} .*

If $l = 0$, so that $\omega_\alpha = \Omega$ and the children of ω_α are simply ω_1 and ω_2 , then this sublemma follows from the way that Ω , ω_1 , and ω_2 were chosen (as smooth solid 4-dimensional tori which are cut in half by the 3-plane P in the standard way, etc.). In other words, we can just shrink the ω_i 's in Ω as close to the D_i 's as we want, without disturbing any points in P or outside of Ω . One can do this rather explicitly, but one could also use Lemma 4.1 (or rather a small variant of it).

For an arbitrary $\alpha \in \mathcal{S}_l$ we can reduce to the preceding case as follows. We know from (8.9) that $h_l(\omega_\alpha) = \Omega_\alpha = \psi_\alpha(\Omega)$, and we have that $h_l(\omega_\alpha \cap P) = h_l(D_\alpha) = \Omega_\alpha \cap P = \psi_\alpha(\Omega \cap P)$ by (8.14), (3.31), (3.8), and (3.7). We want to show that $h_l(\omega_\beta)$ is the same as the image

under ψ_α of ω_1 or ω_2 when β is a child of α . If β is a child of α , then $\omega_\beta = h_{l+1}^{-1}(\Omega_\beta)$ by (8.9), and so $h_l(\omega_\beta) = h_l \circ h_{l+1}^{-1}(\Omega_\beta) = g_l^{-1}(\Omega_\beta)$ by the definition (3.27) of h_{l+1} , where g_l is as in the proof of Lemma 3.21. We also have $g_l^{-1}(\Omega_\beta) = \theta_\alpha^{-1}(\Omega_\beta)$ because of the definition of g_l (just before (3.25)). The fact that this is the same as the image under ψ_α of ω_1 or ω_2 follows from (3.22) and (3.6), using also the definitions of Ω_β and ψ_α (just before (3.7)) to get that $\psi_\alpha^{-1}(\Omega_\beta)$ is one of Ω_1 or Ω_2 . Now that we understand $h_l(\omega_\beta)$ we can reduce Sublemma 8.21 for ω_α to the $l = 0$ case of the preceding paragraph using h_l . Note that this reduction of the general case to Ω involves mappings with large distortions, so that we are not getting good estimates (as a function of l and ε), but we do not care. This proves Sublemma 8.21.

It is now a simple matter to construct the σ_l 's recursively. Given σ_l , one builds σ_{l+1} by composing σ_l on the right with diffeomorphisms on each of the ω_α 's, $\alpha \in \mathcal{S}_l$, which are chosen as in Sublemma 8.21. Because these diffeomorphisms equal the identity on P and outside the corresponding ω_α 's the properties (8.17), (8.18), and (8.19) are maintained. We get (8.20) for σ_{l+1} simply by choosing ε in Sublemma 8.21 to be suitably small. This choice of ε will depend on the modulus of continuity of σ_l , but we do not mind.

Once we have the σ_l 's we simply take σ to be their limit as $l \rightarrow \infty$. This makes sense as an \mathbb{R}^4 -valued mapping only as long as we stay away from C^* , but in fact the limit exists as a map from \mathbb{R}^4/G^* onto \mathbb{R}^4/G' because of (8.20). It is not hard to check that this choice of σ has the required properties. This proves Lemma 8.16.

Because of Lemmas 8.16 and 8.7 we can form the map $\zeta = h \circ \sigma^{-1}$ which gives a homeomorphism from \mathbb{R}^4/G' onto \mathbb{R}^4 . Lemmas 8.16 and 3.21 imply that ζ agrees with θ outside $\omega_1 \cup \omega_2$ and that it maps P onto M . By composing ζ with the inverse of the map promised in Proposition 8.5 we get a homeomorphism ν as in Theorem 8.1.b). This completes the proof of Theorem 8.1.

9. Analysis on these sets.

To what extent can we do analysis on the sets M , M^j , and \widetilde{M} ? Because the geometry of these sets is approximately Euclidean we can hope that much of the usual analysis on Euclidean spaces also works on

these sets. For this we should remember to require that ρ in Definition 3.2 satisfy $\rho^3 n < 1$, so that (3.46) and the conclusions of Lemma 3.45 hold.

We can get a lot of mileage out of the fact that the M^j 's are all smooth, M is smooth away from a compact singular set with 3-dimensional measure zero, and \widetilde{M} is smooth away from a single point. For instance we can talk about smooth functions (away from the singular sets) and we can use Stokes' theorem (if we are a little bit careful about the singularities). We also have automatically the "local" results in real analysis, like the existence almost everywhere of Lebesgue points of locally integrable functions, points of density of measurable sets, and derivatives of Lipschitz functions. We can get more quantitative results from real analysis using the fact that our sets are Ahlfors regular and hence "spaces of homogeneous type" in the sense of [CW1], [CW2], with uniform bounds in the case of the M^j 's. The uniform rectifiability of these sets implies additional analytical information, as in [Da], [DS2], [DS3], [DS4]. Of course the M 's are much nicer than most regular or uniformly rectifiable sets, because they degenerate only near small sets, and in a moderate way.

What about Sobolev and Poincaré inequalities? To what extent can we control a function on M in terms of its gradient? If these sets were quasisymmetrically equivalent to \mathbb{R}^3 then they would have to be well behaved for Sobolev and Poincaré inequalities. Indeed, if a quasisymmetric parameterization were to exist, it would distort Hausdorff measure by only an " A_∞ weight", by the method of Gehring [Ge] (see also [DS1], [Se3]), and one could then obtain Sobolev and Poincaré inequalities from the results in [DS1] (since the A_∞ weight in question would have to be a strong A_∞ weight). This approach is not available here, but we can verify Sobolev and Poincaré inequalities directly using the bilipschitz balls provided by Proposition 3.70, as we shall do in the next section. (One could do with less than these bilipschitz balls, but there is not much point in that here. The method of [Se1] (see also [DS3, Section 6]) would work too, but it is unnecessarily indirect for the present situation.)

It would be interesting to find nice ways to see the strangeness of the M 's in terms of analysis on them or on their complementary domains, *e.g.*, in terms of harmonic functions, or Clifford holomorphic functions (see [BDS]), or nonlinear potential theory.

10. Sobolev and Poincaré inequalities.

The main result of this section is that the sets M , M^j , or \widetilde{M} from Section 3 satisfy Sobolev and Poincaré inequalities under mild assumptions. This will be a rather simple consequence of Proposition 3.70. Much of the structure of excellent packages will not really be needed, as in the comments which follow Definition 3.2, and there is nothing special about the dimension 3. For that matter, we could do with less than the bilipschitz balls provided by Proposition 3.70, but since we have them we may as well use them.

Let us begin with some general definitions. Let us call a subset of \mathbb{R}^4 a “singular submanifold” if it is a smooth embedded 3-dimensional submanifold away from a singular set which is closed and has zero 3-dimensional Hausdorff measure. This definition accommodates the sets M and \widetilde{M} , while the M^j 's are everywhere smooth already. All of our integrals, L^p norms, etc., on singular submanifolds will be taken with respect to 3-dimensional Hausdorff measure, which we shall denote by “ dx ” or something similar. If f is a locally integrable function on a singular submanifold then its differential df makes sense as a distribution (or more precisely a current, as in [Fe]) which is defined away from the singular set. As a practical matter one should not take this business about currents too seriously here, because we shall always work with the differential on the smooth part of our set, in such a way that all the computations can be reduced to similar problems on a piece of \mathbb{R}^3 using local coordinates. We shall always assume that df is locally integrable away from the singular set, so that we can use the Euclidean metric to define $|df|$ and hence L^p norms and other integrals of $|df|$. When we say that “ df is locally integrable” on our singular submanifold we shall mean that $|df|$ is actually locally integrable across the singular set, and hence has an unambiguous locally integrable extension to the whole singular submanifold, since the singular set has measure zero. This extension may differ from the natural distributional definitions of df near the singular set -i.e., we could be dropping Dirac masses or other singular contributions to the “true” df - but we do not care. We shall always do our distribution theory away from the singular set and then extend brutally across the singular set.

Proposition 10.1. *Suppose that $E = M, \widetilde{M}$, or M^j for some j , that ρ in Definition 3.2 satisfies $\rho^3 n < 1$, and that our excellent package satisfies (3.69). There exist constants L_1 and C_1 , depending only on our excellent package, so that if $p, q \in E$, f is a locally integrable function*

on $E \cap B(p, L_1|p - q|)$, p and q are Lebesgue points for f , and df is locally integrable on $E \cap B(p, L_1|p - q|)$ (in the sense just described when $E = M$ or \widetilde{M}), then

$$(10.2) \quad |f(p) - f(q)| \leq C_1 \int_{E \cap B(p, L_1|p - q|)} \left(\frac{1}{|p - z|^2} + \frac{1}{|q - z|^2} \right) |df(z)| dz.$$

The statement that p is a Lebesgue point for f means that

$$(10.3) \quad \lim_{r \rightarrow 0} r^{-3} \int_{E \cap B(p, r)} |f(x) - f(u)| du = 0.$$

The “ du ” here refers to 3-dimensional Hausdorff measure, as indicated above. Almost all points are Lebesgue points.

To prove Proposition 10.1 let us first review the situation for \mathbb{R}^3 . Suppose that h is a smooth function on a ball B in \mathbb{R}^3 with radius r . Then

$$(10.4) \quad r^{-6} \int_B \int_B |h(u) - h(v)| du dv \leq C r^{-3} \int_B r |dh(w)| dw,$$

where C does not depend on h, B , or r . The proof of this is quite easy. We can express $h(u) - h(v)$ in terms of an integral of dh over the line segment which connects u to v , and (10.4) is obtained by averaging this formula over all u and v and applying Fubini’s theorem. This Poincaré-type inequality (10.4) also holds when h is merely locally integrable and has locally integrable distributional first derivatives, because of standard approximation arguments (as in the proof of [St, p. 122, Proposition 1]).

Now suppose that x and y lie in \overline{B} , h is locally integrable and has locally integrable first derivatives on B , and that x and y are Lebesgue points for h in the sense of (10.3) with $E = B$. (If x or y lies in ∂B , so that they are not in the putative domain of h , then we simply assume that h is also defined at x and y in such a way that (10.3) holds.) Then

$$(10.5) \quad |h(x) - h(y)| \leq C \int_B \left(\frac{1}{|x - z|^2} + \frac{1}{|y - z|^2} \right) |dh(z)| dz.$$

This is also well-known, but let us quickly go through a proof. Let us assume for simplicity that B is centered at the origin and has radius

one. Set $B_k(x) = B((1 - 2^{-k})x, 2^{-k})$ for $k \geq 0$, and define $B_k(y)$ similarly. Thus $B_0(x) = B_0(y) = B$ and $B_{k+1} \subseteq B_k \subseteq B$ for all k (by the triangle inequality). If we let $\text{Av}(h, Z)$ denote the average of h over the subset Z of B (with positive measure), then our assumption that x and y are Lebesgue points implies that $\lim_{k \rightarrow \infty} \text{Av}(h, B_k(x)) = h(x)$ and similarly for y . Hence

$$(10.6) \quad |h(x) - h(y)| \leq \sum_{k=0}^{\infty} |\text{Av}(h, B_k(x)) - \text{Av}(h, B_{k+1}(x))| + \sum_{k=0}^{\infty} |\text{Av}(h, B_k(y)) - \text{Av}(h, B_{k+1}(y))|,$$

since $B_0(x) = B_0(y)$. We also have that

$$(10.7) \quad |\text{Av}(h, B_k(x)) - \text{Av}(h, B_{k+1}(x))| \leq C \int_{B_k(x)} 2^{-2k} |dh(z)| dz,$$

and similarly with x replaced by y , because of (10.4) applied to B_k . It is easy to derive (10.5) from these inequalities.

Proposition 10.1 follows immediately from (10.5) and Proposition 3.70. We are also using here the comment made shortly after the statement of Proposition 3.70 to the effect that the bilipschitz 3-ball W in Proposition 3.70 can be chosen to be smooth away from p and q , and the bilipschitz parameterization of W can be taken to be smooth away from these points. This permits us to avoid technical issues concerning the distribution theory. Other than that we are simply using the bilipschitz invariance of (10.2) and (10.3) in a brutal way.

Proposition 10.8. *Suppose that $E = M, \widetilde{M}$, or M^j for some j , that ρ in Definition 3.2 satisfies $\rho^3 n < 1$, and that our excellent package satisfies (3.69). Then there exist constants L_2 and C_2 , depending only on our excellent package, so that if B is a ball with radius r centered on E and f is a locally integrable function on $E \cap L_2 B$ such that df is also locally integrable there (in the sense described before Proposition 10.1 when $E = M$ or \widetilde{M}), then we have the Poincaré inequality*

$$(10.9) \quad r^{-6} \int_{E \cap B} \int_{E \cap B} |f(x) - f(y)| dx dy \leq C_2 r^{-3} \int_{E \cap L_2 B} r |df(z)| dz.$$

If f is a locally integrable function on E with compact support and if df is locally integrable on E , then

$$(10.10) \quad |f(x)| \leq C \int_E \frac{1}{|z-x|^2} |df(z)| dz,$$

for some C (depending only on our excellent package) and almost all $x \in E$.

To prove this let us first check that

$$(10.11) \quad \int_{B(z,t) \cap E} \frac{1}{|x-z|^2} dx \leq C t,$$

for all $z \in E$ and $t > 0$. Using the Ahlfors regularity of E we get that

$$(10.12) \quad \int_{(B(z,2s) \setminus B(z,s)) \cap E} \frac{1}{|z-x|^2} dx \leq C s,$$

for all $s > 0$, and (10.11) follows from this by summing the obvious geometric series.

Once we have (10.11) we can derive (10.9) from Proposition 10.1 by simply averaging (10.2) over p and q and using Fubini's theorem. The pointwise inequality (10.10) is an immediate consequence of (10.2). (Just take y to be far far away.)

Let us now check that Sobolev embeddings work for the sets E as in Proposition 10.8. Consider first the potential operator I_1 on functions on E defined by

$$(10.13) \quad I_1(g)(x) = \int_E \frac{1}{|x-z|^2} g(z) dz.$$

This operator has exactly the same $L^p \rightarrow L^q$ mapping properties on any regular set of dimension 3 as on \mathbb{R}^3 itself, *i.e.*, it maps $L^p(E)$ into $L^q(E)$ when $1 < p < 3$ and $1/q = 1/p - 1/3$. This is not hard to prove -the point is that the two situations are essentially the same at the level of this kind of measure theory- and one can simply mimic the proof of the corresponding result on \mathbb{R}^3 ([St, p. 119, Theorem 1]). This is really just a minor variation on [St, p. 121, Comment 1.4]. Alternatively, one could invoke theorems from the real method of interpolation of Banach spaces or other general results.

The usual Sobolev embedding theorems on E follow immediately from the $L^p \rightarrow L^q$ mapping properties for the potential operator I_1 and (10.10), at least for $p > 1$. This method breaks down for $p = 1$, but one can establish isoperimetric inequalities for these sets using the Poincaré inequality (10.9) and a covering lemma. (A very similar argument was used in [DS1].)

If df lies in L^p for some $p > 3$, then one can modify f on a set of measure zero to get a function which is Hölder continuous of order $1 - 3/p$. This is an easy consequence of (10.2) and Hölder's inequality. If $df \in L^3$ then f lies in $BMO(E)$, because of (10.9).

The bottom line is that Propositions 10.1 and 10.8 permit us to verify that many of the usual results about functions on \mathbb{R}^3 which satisfy Sobolev space conditions also work on M , \widetilde{M} , and the M^j 's. Of course the preceding list is not exhaustive.

11. A remark about polyhedra.

Proposition 11.1. *A finite d -dimensional polyhedron P (in some \mathbb{R}^n) which is a topological manifold (without boundary) satisfies the analogue of (**) for compact sets (i.e., the conditions of Definition 1.7 hold when r is sufficiently small).*

These polyhedra can be pretty strange, because of the results of Edwards and Cannon on the double suspensions of spheres (as mentioned in Sections 1 and 2).

Proposition 11.1 is no surprise, but it seems worthwhile to record it in view of the gap between (**) and the existence of a quasisymmetric parameterization established by Theorem 1.12, and since we know that there are interesting examples of these polyhedra.

Let us prove the proposition. Let P be as above, and note that P must have pure dimension d . Since P is a polyhedron we can give it the structure of a simplicial complex. That is, we can realize P as a (finite) union of d -dimensional simplices in such a way that the collection of all these simplices together with all their faces (of any dimension, including 0 (vertices)) have the property that when any two of them intersect, either one is a face of the other or the two intersect in a common face. We shall call a simplex in P "distinguished" if it is one of these basic simplices or one of their faces (as opposed to a random simplex floating around in \mathbb{R}^n). All simplices in this discussion are assumed to be closed,

and ∂A will be used to denote the geometric (or simplicial) boundary of A (as opposed to the uninteresting topological boundary of A as a subset of \mathbb{R}^n). If A is just a vertex, then we interpret ∂A to be the empty set.

Finite polyhedra obviously have a lot of homogeneity to them, and the proof of Proposition 11.1 merely requires a precise formulation of this homogeneity. We begin with a preliminary fact.

Let A, X be distinguished simplices in P , with $A \subseteq X$. Then A is a face of X , and we can order the vertices v_1, \dots, v_l of X in such a way that v_1, \dots, v_j are the vertices of A . Thus A is the convex hull of v_1, \dots, v_j and X is the convex hull of v_1, \dots, v_l . The $j - 1$ -plane $Q(A)$ determined by A can be described as the set of points of the form $\sum_{i=1}^j \lambda_i v_i$ where $\sum_i \lambda_i = 1$ and the λ_i 's are real numbers. (Points in A correspond to restricting ourselves to λ_i 's which are nonnegative.) Let $Q_0(A)$ denote the set of points of the form $\sum_{i=1}^j \lambda_i v_i$ where $\sum_i \lambda_i = 0$. This is just a translation of $Q(A)$ which contains the origin. Notice that $Q(A)$ is preserved by translations by elements of $Q_0(A)$.

Let $H(A, X)$ be the set of points of the form $\sum_i \lambda_i v_i$ with $\sum_i \lambda_i = 1$ and $\lambda_i \geq 0$ when $i > j$. Notice that $X \subseteq H(A, X)$. If A is just a vertex then $H(A, X)$ is the cone with vertex A generated by X , while in general it is a product of the plane determined by A with a cone. Note that $H(A, X)$ is preserved by translations of elements of $Q_0(A)$. One can also check that $H(A, X)$ is preserved by dilations by positive constants centered at elements of $Q(A)$, i.e., mappings of the form $z \mapsto a + t(z - a)$ for $a \in Q(A)$ and $t > 0$. We shall eventually use these symmetries of $H(A, X)$ to make precise the homogeneity of P , but first we establish the following.

Lemma 11.2. *Let A, X be as in the preceding paragraphs, and let $a \in A$ and $p \in H(A, X) \setminus X$ be given. Then there is a point x on the line segment that joins a and p with the property that x lies in a face of X which does not contain (all of) A . (It could be that $x = a \in \partial A$).*

By assumptions, $a = \sum_i \alpha_i v_i$ and $p = \sum_i \lambda_i v_i$, where $\sum_i \alpha_i = \sum_i \lambda_i = 1$, $\alpha_i \geq 0$ when $i \leq j$, $\alpha_i = 0$ when $i > j$, and $\lambda_i \geq 0$ when $i > j$. Choose t to be the first real number ≥ 0 such that $(1 - t)\alpha_i + t\lambda_i = 0$ for some $i \leq j$ ($t = 0$ is possible), and set $x = \sum_i ((1 - t)\alpha_i + t\lambda_i)v_i$. Our choice of t implies that $(1 - t)\alpha_i + t\lambda_i \geq 0$ for all $i \leq j$, and this is true when $i > j$ as well, by our assumptions. It is easy to check that x lies in a face of X which does not contain A . Also, $t < 1$, because

$p \notin X$ (and hence $\lambda_i < 0$ for some $i \leq j$). This proves the lemma.

Given a distinguished simplex A in P let $S(A)$ denote the star of A , i.e., the union of all the distinguished simplices in P which contain A as a face. Because P has pure dimension d one may simply take the union of the d -dimensional distinguished simplices which contain A as a face. (The d -dimensional ones will contain all the others.) If A is a vertex then $S(A)$ is a neighborhood of A (in P) with nice properties. If A has positive dimension, then every point in $A \setminus \partial A$ lies in the interior of $S(A)$ (relative to P), but this is not necessarily true for points in ∂A . Let $H(A)$ be the union of $H(A, X)$ for all distinguished simplices X in P which contain A , and note that $S(A) \subseteq H(A)$.

For a distinguished simplex A in P let $C(A)$ denote the union of all distinguished simplices X which do not contain A . Thus $C(A)$ contains ∂A in particular, and $C(A)$ is approximately the same as the complement of $S(A)$. In fact $C(A)$ is the union of the complement of $S(A)$ and the faces of the d -simplices in $S(A)$ which do not contain A .

There is a constant k such that

$$(11.3) \quad \text{dist}(a, C(A)) \geq k \text{dist}(a, \partial A),$$

for any distinguished simplex A in P and all $a \in A$. This comes down to the fact that if X is a distinguished simplex in P , and if X does not contain A , then either X is disjoint from A , X is contained in ∂A , or X meets A in a face of ∂A and makes a definite angle with A . Keep in mind also that there are only finitely many A 's and X 's around, so that we can easily choose k to be independent of them. When A is a vertex the correct version of (11.3) is that $\text{dist}(A, C(A))$ is bounded from below.

Let us call a ball $B(z, r)$ centered on P "good" with respect to a distinguished simplex A in P if $z \in A$ and $B(z, r) \cap C(A) = \emptyset$. The key property of a good ball is that

$$(11.4) \quad B(z, r) \cap P = B(z, r) \cap H(A)$$

when $B(z, r)$ is good relative to A . Clearly $B(z, r) \cap P = B(z, r) \cap S(A)$ when $B(z, r)$ is good, and so we get one inclusion in (11.4) from the fact that $S(A) \subseteq H(A)$. If we have a point p in $B(z, r) \cap H(A)$ which does not lie in $S(A)$, then Lemma 11.2 produces a point x on the line segment which joins p and z (so that $x \in B(z, r)$) with the property that x lies in a distinguished simplex in P which does not contain A . This

means that $x \in C(A)$, which contradicts the assumption that $B(z, r)$ is good. This proves (11.4).

If $B(y, s)$ is another ball which is good relative to A then the obvious translation and dilation which sends $B(z, r)$ onto $B(y, s)$ must send $B(z, r) \cap P$ onto $B(y, s) \cap P$. This follows from (11.4), because each $H(A, X)$ is preserved by this translation and dilation (see the paragraph just before Lemma 11.2), and so $H(A)$ is also preserved by these mappings. Thus the intersection of a good ball with P is equivalent (by a translation and a dilation) to one of finitely many models. To prove Proposition 11.1 we need to show that we can ignore the bad balls.

Lemma 11.5. *Given any constant $L > 0$, there exist positive constants K, ε (depending on L and P) with the property that if $x \in P$ and $0 < r < \varepsilon$ then there is a ball $B(\xi, t)$ such that $B(\xi, Lt)$ is good with respect to some simplex, $B(\xi, t) \supset B(x, r)$, and $t \leq Kr$.*

Let x and r be given, and let A be a distinguished simplex in P which contains x . The ball $B(x, Lr)$ is itself good if $Lr \leq k \operatorname{dist}(x, \partial A)$, where k is as in (11.3), and so we assume that this inequality is not true. This means that there is a point $y \in \partial A$ such that $B(x, r) \subset B(y, Cr)$ for a suitable constant $C = C(L)$. Thus y lies in a lower-dimensional simplex, and we can repeat the argument to conclude that either $B(y, LCr)$ is good or $B(y, Cr)$ is contained in $B(z, C^2r)$ for some z in a lower dimensional simplex. Repeating this as many times as necessary (but at most d times) we reduce to the case of vertices. A ball centered at a vertex is good as soon as its radius is small enough. Lemma 11.5 follows easily from this.

Let us analyze the the structure of the $H(A)$'s some more. For each distinguished simplex A in P we have that $H(A)$ is a topological manifold. Indeed, near a point in $A \setminus \partial A$ $H(A)$ looks like P , because of (11.4) and the existence of good points, and therefore $H(A)$ is a topological manifold near such points. This implies that all points of $H(A)$ are manifold points, because $H(A)$ is invariant under translations by elements of $Q_0(A)$ and dilations centered at points in $Q(A)$ (as discussed just before Lemma 11.2). In fact we get that there is a constant L so that if $a \in A$ and $r > 0$ then there is a relatively open set U of $H(A)$ with $B(x, r) \cap H(A) \subseteq U \subseteq B(x, Lr) \cap H(A)$ such that U is a topological d -ball and its closure is homeomorphic to the closed unit d -ball. This follows easily from the invariance properties of $H(A)$ just mentioned. (That is, we choose U once for some fixed ball, and

then we use the translations and dilations to extend this choice to other balls.) We can also choose these sets U so that they are all translates and dilates of each other for a fixed A . Hence there are only finitely many models total, since there are only finitely many A 's. We can also choose the constant L so that it does not depend on A , since there are only finitely many A 's.

We can now finish the proof of Proposition 11.1. Let $x \in P$ and $r > 0$ be given. Lemma 11.5 implies that if r is small enough, then there is a ball $B(\xi, t)$ such that $B(\xi, Lt)$ is good with respect to some simplex A , $B(\xi, t) \supset B(x, r)$, and $t \leq Kr$. From the observations of the preceding paragraph we obtain that there is relatively open set U in $H(A)$ which is a topological d -ball and satisfies $B(\xi, t) \cap H(A) \subseteq U \subseteq B(\xi, Lt) \cap H(A)$. Since $B(\xi, Lt)$ is good we actually have that $B(\xi, Lt) \cap H(A) = B(\xi, Lt) \cap P$. As mentioned above, we can also choose U so that its closure is homeomorphic to a closed ball and so that U can be reduced by translations and dilations to one of a finite set of models. This implies that P satisfies the compact version of (**), as promised. (Of course it is the finiteness of the set of the models for U which gives us the uniform estimates as in (**), as opposed to (*). Do not forget that continuous maps between compact sets are uniformly continuous.)

12. Concluding remarks.

Geometric topology provides a lot of technology for building homeomorphic parameterizations of a set, and it provides many interesting examples, but little is known about quantitative estimates on the geometric complexity of these parameterizations. The most basic example that I know is the type of mapping constructed by Bing in [B1], for which optimal estimates are (apparently) still unknown (but see [B5] and the remarks at the end of Section 2 of [FS]). The construction in Section 3 gives an interesting formulation to this problem, namely, what kind of estimates can be satisfied by a homeomorphic parameterization of the set M from Section 6 by \mathbb{R}^3 ? We know from [B1] that such a parameterization exists, and we know that it cannot be quasimetric, but it seems to be unknown whether M can be parameterized by a Hölder continuous map whose inverse is also Hölder continuous (for instance), let alone the possible range of orders of Hölder continuity (if

any). I am also unaware of any estimates (other than uniform continuity) for the homeomorphisms between the strange polyhedral spheres of Edwards and Cannon and standard Euclidean spheres.

One can wonder whether Condition (**), even in combination with Ahlfors regularity, implies the existence of local homeomorphic parameterizations with a universal choice (or family of choices) of modulus of continuity ω as in Definition 1.7. For instance, do Condition (**) and Ahlfors regularity imply the existence of homeomorphic local parameterizations which are Hölder continuous, of some universal order, or even of just some positive order? This leads back to questions about the type of mapping produced in [B1] for the set M from Section 6, or the kind of local coordinates that might exist for the strange polyhedral spheres of Edwards and Cannon. (I am not optimistic.) Note that there is a positive result when $d = 2$, because of Theorem 1.6.

One can also ask whether the combination of (**) and Ahlfors regularity (of the same dimension) imply the existence of local homeomorphic parameterizations which satisfy Sobolev space conditions, as well as their inverses. (One should be a little careful in the formulation of Sobolev space conditions for a map into a metric space. I prefer to use maximal functions, as in [Se2].) In the case of quasimetric parameterizations there are pretty strong results of this type, with the Sobolev exponent p larger than the dimension, because of a method of Gehring [Ge] (see also [DS1] and [Se3]). The general case would be more complicated (if true at all), because one would have to modify the parameterization. Note that the 2-dimensional case is again special, because of Theorem 1.6. In all dimensions one can get “uniform rectifiability” (in the sense of [DS4]) under even more general conditions, by [DS5]. Uniform rectifiability implies the existence of some non-homeomorphic parameterizations with otherwise very good estimates, *i.e.*, there are bilipschitz parameterizations of large pieces of the set, and the set can be put inside a bigger one which admits a controlled parameterization that allows a limited amount of crossing. (See also [DS2].)

Geometric topology is particularly successful at parameterizing a set if one is permitted to “stabilize” first by taking a Cartesian product with the real line or some \mathbb{R}^k . Perhaps there is a general theorem to the effect that (*), possibly in combination with Ahlfors regularity, implies (**) after stabilizing. This should be compared with the result of Ferry [F2] for (†) mentioned in the introduction. A positive conjecture is supported by the examples discussed in Sections 4 and 5. In each

of those situations the non-manifold M becomes homeomorphic to \mathbb{R}^4 after taking a Cartesian product with the real line. (See [K, p. 87, Theorem 1] for the set discussed in Section 4, [B4] for the set discussed in Section 5, and [D, Section 11] for both cases and others.) This fact together with self-similarity allows us to obtain (**) for the product by the same arguments as used in Section 6 to prove Theorem 6.3. Other examples can be produced as in of [D, Section 11].

One can also ask whether stabilization makes it easier to get quasisymmetric parameterizations. That is, if M is the non-manifold constructed as in Sections 4 or 5, then does $M \times \mathbb{R}$ admit a quasisymmetric parameterization by \mathbb{R}^4 ? I am pessimistic, for reasons like those in Remark 6.19.

See also [Se5] for some related topics.

Acknowledgements. I would like to thank Juha Heinonen for many valuable consultations. We have been discussing these kinds of questions for some time, and it was he who told me about [FS]. I am also grateful to a splintered cabal of topologists for revealing to me a fraction of their ancient knowledge.

References.

- [Ar] Armentrout, S., Cellular decompositions of 3-manifolds that yield 3-manifolds. *Memoirs Amer. Math. Soc.* **107** (1971).
- [A1] Assouad, P., Espaces Métriques, Plongements, Facteurs, Thèse de Doctorat, Université de Paris XI, 1977.
- [A2] Assouad, P., Étude d'une dimension métrique liée à la possibilité de plongement dans \mathbb{R}^n . *C. R. Acad. Sci. Paris* **288** (1979), 731-734.
- [A3] Assouad, P., Plongements Lipschitziens dans \mathbb{R}^n . *Bull. Soc. Math. France* **111** (1983), 429-448.
- [AV] Alestalo, P., Väisälä, J., Quasisymmetric embeddings of products of cells into the Euclidean space. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **19** (1994), 375-392.
- [B1] Bing, R. H., A homeomorphism between the 3-sphere and the sum of two solid horned spheres. *Ann. of Math.* **56** (1952), 354-362.
- [B2] Bing, R. H., Upper semicontinuous decompositions of E^3 . *Ann. of Math.* **65** (1957), 363-374.
- [B3] Bing, R. H., A decomposition of E^3 into points and tame arcs such that the decomposition space is topologically different from E^3 . *Ann. of*

- Math.* **65** (1957), 484-500.
- [B4] Bing, R. H., The Cartesian product of a certain nonmanifold and a line is E^4 . *Ann. of Math.* **70** (1959), 399-412.
- [B5] Bing, R. H., Shrinking without lengthening. *Topology* **27** (1988), 487-493.
- [Bi] Bizaca, Ž., An explicit family of exotic Casson handles. *Proc. Amer. Math. Soc.* **123** (1995), 1297-1302.
- [BDS] Brackx, F., Delanghe, R., Sommen, F., *Clifford Analysis*, Pitman, 1982.
- [B^∞] Bryant, J., Ferry, S., Mio, W., Weinberger, S., Topology of homology manifolds. *Bull. Amer. Math. Soc.* **28** (1993), 324-328.
- [C1] Cannon, J., The recognition problem: What is a topological manifold?. *Bull. Amer. Math. Soc.* **84** (1978), 832-866.
- [C2] Cannon, J., The characterization of topological manifolds of dimension $n \geq 5$. *Proceedings I.C.M., Helsinki* (1978), 449-454.
- [C3] Cannon, J., Shrinking cell-like decompositions of manifolds: codimension 3. *Ann. of Math.* **110** (1979), 83-112.
- [CW1] Coifman, R. Weiss, G., *Analyse Harmonique Non-commutative sur Certains Espaces Homogènes*. Lecture Notes in Math. **242**, Springer-Verlag, 1971.
- [CW2] Coifman, R., Weiss, G., Extensions of Hardy spaces and their use in analysis. *Bull. Amer. Math. Soc.* **83** (1977), 569-645.
- [Cu] Curtis, M., An imbedding theorem. *Duke Math. J.* **24** (1957), 349-352.
- [D] Daverman, R., *Decompositions of Manifolds*. Academic Press, 1986.
- [Da] David, G., *Wavelets and Singular Integrals on Curves and Surfaces*. Lecture Notes in Math. **1465**, Springer-Verlag, 1991.
- [DS1] David, G., Semmes, S., Strong A_∞ -weights, Sobolev inequalities, and quasiconformal mappings. In *Analysis and Partial Differential Equations*, edited by C. Sadosky, Lecture Notes in Pure and Appl. Math. **122**, Marcel Dekker, 1990.
- [DS2] David, G., Semmes, S., *Singular integrals and rectifiable sets in \mathbb{R}^n : au-delà des graphes lipschitziens*. Astérisque **193**, Soc. Math. France, 1991.
- [DS3] David, G., Semmes, S., Quantitative rectifiability and Lipschitz mappings. *Trans. Amer. Math. Soc.* **337** (1993), 855-889.
- [DS4] David, G., Semmes, S., *Analysis of and on Uniformly Rectifiable Sets*. Math. Surveys and Monographs **38**, Amer. Math. Soc., 1993.
- [DS5] David, G., Semmes, S., to exist.
- [DoS] Donaldson, S., Sullivan, D., Quasiconformal 4-manifolds. *Acta Math.* **163** (1989), 181-252.

- [E] Edwards, E., The topology of manifolds and cell-like maps. *Proceedings I.C.M., Helsinki* (1978), 111-127.
- [Fe] Federer, H., *Geometric Measure Theory*, Springer-Verlag, 1969.
- [F1] Ferry, S., Topological finiteness theorems for manifolds in Gromov-Hausdorff space. *Duke Math. J.* **74** (1994), 95-106.
- [F2] Ferry, S., Counting simple-homotopy types of topological manifolds. Preprint.
- [Fr] Freedman, M. H., The topology of four-dimensional manifolds. *F. Diff. Geometry* **17** (1982), 357-453.
- [FS] Freedman, M., Skora, R., Strange actions of groups on spheres. *J. Diff. Geom.* **25** (1987), 75-98.
- [Ge] Gehring, F. W., The L^p integrability of the partial derivatives of a quasiconformal mapping. *Acta Math.* **130** (1973), 265-277.
- [GP] Greene, R., Petersen V, P., Little topology, big volume. *Duke. Math. J.* **67** (1992), 273-290.
- [GPW] Grove, K., Petersen V, P., Wu, J.-Y., Geometric finiteness theorems via controlled topology. *Invent. Math.* **99** (1990), 205-213; erratum, **104** (1991), 221-222.
- [HK] Heinonen, J., Koskela, P., Definitions of quasiconformality. *Invent. Math.* **120** (1995), 61-79.
- [HY] Heinonen, J., Yang, S., Strongly uniform domains and periodic quasiconformal maps. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **20** (1995), 123-148.
- [K] Kirby, R., *The Topology of 4-manifolds*. Lecture Notes in Math. **1374**, Springer-Verlag, 1989.
- [P] Petersen V, P., Gromov-Hausdorff convergence of metric spaces. *Proc. Symp. Pure Math.* **54** (1993), 489-504.
- [Q1] Quinn, F., Resolutions of manifolds, and the topological characterization of manifolds. *Invent. Math.* **72** (1983), 267-284.
- [Q2] Quinn, F., An obstruction to the resolution of homology manifolds. *Michigan Math. J.* **301** (1987), 285-292.
- [Se1] Semmes, S., Differentiable function theory on hypersurfaces in \mathbb{R}^n (without bounds on their smoothness). *Indiana Math. J.* **39** (1990), 985-1004.
- [Se2] Semmes, S., Chord-arc surfaces with small constant II: Good parameterizations. *Advances in Math.* **88** (1991), 170 -199.
- [Se3] Semmes, S., Bilipschitz mappings and strong A_∞ weights. *Ann. Acad. Sci. Fenn. Serie A I Math.* **18** (1993), 211-248.
- [Se4] Semmes, S., Finding structure in sets with little smoothness. To appear in *Proc. I.C.M., Zürich* (1994), Birkhäuser.
- [Se5] Semmes, S., On the nonexistence of bilipschitz parameterizations and

- geometric problems about A_∞ weights. To appear in *Revista Mat. Iberoamericana*.
- [SS] Siebenmann, L., Sullivan, D., On complexes that are Lipschitz manifolds. In *Geometric Topology*, J. Cantrell, ed., Academic Press, 1979, 503-525.
- [St] Stein, E. M., *Singular Integrals and Differentiability Properties of Functions*. Princeton Univ. Press, 1970.
- [To1] Toro, T., Surfaces with generalized second fundamental form in L^2 are Lipschitz manifolds. *J. Diff. Geom.* **39** (1994), 65-102.
- [To2] Toro, T., Geometric conditions and existence of bilipschitz parameterizations. *Duke Math. J.* **77** (1995), 193-227.
- [Tu1] Tukia, P., The planar Schönflies theorem for Lipschitz maps. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **5** (1980), 49-72.
- [Tu2] Tukia, P., A quasiconformal group not isomorphic to a Möbius group. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **6** (1981), 149-160.
- [TV] Tukia, P., Väisälä, J., Quasisymmetric embeddings of Euclidean spaces. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **5** (1980), 97-114.
- [V1] Väisälä, J., Quasisymmetric maps of products of curves into the plane. *Rev. Roumaine Math. Pures Appl.* **33** (1988), 147-156.
- [V2] Väisälä, J., Metric duality in Euclidean spaces. Preprint.

Recibido: 15 de febrero de 1.995

Stephen Semmes*
 Department of Mathematics
 Rice University
 Houston, Texas 77251, U.S.A.
 semmes@math.rice.edu

and

I.H.E.S.
 35 Route de Chartres
 91440 Bures-sur-Yvette, FRANCE

* supported by the U.S. National Science Foundation.