

On the nonexistence of
bilipschitz parameterizations
and geometric problems
about A_∞ -weights

Stephen Semmes

Abstract. How can one recognize when a metric space is bilipschitz equivalent to a Euclidean space? One should not take the abstraction of metric spaces too seriously here; subsets of \mathbb{R}^n are already quite interesting. It is easy to generate geometric conditions which are necessary for bilipschitz equivalence, but it is not clear that such conditions should ever be sufficient. The main point of this paper is that the optimistic conjectures about the existence of bilipschitz parameterizations are wrong. In other words, there are spaces whose geometry is very similar to but still distinct from Euclidean geometry. Related questions of bilipschitz equivalence and embeddings are addressed for metric spaces obtained by deforming the Euclidean metric on \mathbb{R}^n using an A_∞ weight.

1. Introduction.

How can one recognize when a metric space is bilipschitz equivalent to a Euclidean space? Recall that a map $f : M_1 \rightarrow M_2$ between two metric spaces M_1, M_2 is *bilipschitz* if there is a constant K such that $K^{-1} d_1(x, y) \leq d_2(f(x), f(y)) \leq K d_1(x, y)$ for all $x, y \in M_1$, where d_1

and d_2 are the metrics on M_1 and M_2 . We shall sometimes say *K-bilipschitz* to make the constant explicit. Two metric spaces are said to be bilipschitz equivalent if there is a bilipschitz mapping from one onto the other. If two metrics spaces are bilipschitz equivalent then they should have approximately the same behavior in terms of lengths, Hausdorff measure, topology, etc., and the question is whether some nice combination of such conditions can detect bilipschitz equivalence with Euclidean spaces.

Let $(M, d(\cdot, \cdot))$ be a metric space, and suppose that it is bilipschitz equivalent to \mathbb{R}^d with the Euclidean metric. What are some of the conditions that M must satisfy? For each pair of points $x, y \in M$ there must be a curve that joins them whose length is at most $Cd(x, y)$ for some constant C (which depends only on M). There must be a measure μ on M with the property that the μ -mass of a ball of radius $r > 0$ is approximately r^d , i.e., the μ -mass must be bounded from above and below by constants times r^d . In fact this must be true with μ equal to d -dimensional Hausdorff measure on M . There must be another constant C such that any metric ball B in M of radius r is contained in a d -dimensional topological ball U which is itself contained in a metric ball of radius Cr , and one could impose further restrictions on U . Suitable formulations of the Sobolev and Poincaré inequalities on M must also hold.

On the other hand, plenty of the familiar properties of \mathbb{R}^d do not have to be satisfied, even approximately, by a bilipschitz-equivalent metric space M . Bilipschitz mappings need not be smooth, or even C^1 , and so a bilipschitz-equivalent space could have a lot of corners. For instance, if $A : \mathbb{R}^d \rightarrow \mathbb{R}$ is any Lipschitz function (so that $|A(x) - A(y)| \leq C|x - y|$ for some C and all $x, y \in \mathbb{R}^d$), then the graph of A in \mathbb{R}^{d+1} equipped with the ambient Euclidean metric is bilipschitz equivalent to \mathbb{R}^d .

The main purpose of this paper is to provide examples which show that there is no hope for finding simple general conditions of the type just described which ensure the existence of a bilipschitz parameterization. The examples below will all be subsets of some Euclidean space (with the inherited metric) or “conformal” deformations of \mathbb{R}^d . Let us begin with the former, starting with a definition.

Definition 1.1. *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 $C_0 > 0$ such that*

$$(1.2) \quad C_0^{-1} r^d \leq H^d(E \cap B(x, r)) \leq C_0 r^d,$$

for all $x \in E$ and $r > 0$. Here (and forevermore) H^d denotes d -dimensional Hausdorff measure and $B(x, r)$ denotes the open ball with center x and radius r .

This condition is equivalent to the apparently more general version in which one merely asks that there exist a measure μ supported on E which satisfies (1.2). In other words, if such a measure exists, it has to be comparable in size to the restriction of d -dimensional Hausdorff measure to E .

Roughly speaking, a set is regular if it behaves measure-theoretically like \mathbb{R}^d , even though it may be very different geometrically. There are examples of regular sets which are self-similar Cantor sets, or snowflake curves, or tree-like objects. Regular sets can have non-integer dimension. Of course any set which is bilipschitz equivalent to \mathbb{R}^d is regular.

Theorem 1.3. *There is a 3-dimensional regular set E in \mathbb{R}^4 which is the image of a hyperplane under a global quasiconformal mapping from \mathbb{R}^4 onto itself but which is not bilipschitz equivalent to \mathbb{R}^3 . This quasiconformal mapping can also be taken to be Lipschitz continuous. The set E enjoys the additional property that there is a constant $L_0 > 0$ so that every pair of distinct points $x, y \in E$ is contained in a closed subset W of E which is L_0 -bilipschitz equivalent to a closed Euclidean 3-ball. In particular, x and y can be connected by a curve in E of length at most $L_0^2|x - y|$.*

A *quasiconformal mapping* is one which does not distort *relative* distances by more than a bounded factor. There are many equivalent characterizations, one of which is that $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is quasiconformal if there is a constant C so that for each $x \in \mathbb{R}^n$ and $r > 0$ there is an $s > 0$ so that $B(\Phi(x), s) \subseteq \Phi(B(x, r)) \subseteq B(\Phi(x), Cs)$. In other words, Φ maps a ball to a set which is trapped between two balls of comparable radii. It is important that the image radius s is allowed to be very different from r -otherwise this condition would reduce to bilipschitzness- but that the constant C not be allowed to depend on x or r . Recall that Φ is *Lipschitz continuous* if there is a constant C so that $|\Phi(x) - \Phi(y)| \leq C|x - y|$ for all x, y .

Although the image of a 3-plane under a global quasiconformal mapping (or under a quasisymmetric mapping) can be fractal, an image like this which is also regular of dimension 3 automatically has many of the properties that it would have if it were bilipschitz equivalent to \mathbb{R}^3 . Such a set must be uniformly rectifiable (in the sense of [DS4]),

every pair of points in the set can be connected by a path which is not too long compared to the distance between them, and suitable versions of Sobolev and Poincaré inequalities must hold on the set. This uses a method of Gehring [Ge] which allows one to control well the extent to which the quasiconformal mapping distorts distances in this situation. (See also [DS1] and [Se4].) Theorem 1.3 resolves (negatively) the obvious question of whether a quasiconformal image of a plane which is regular of the same dimension must actually be bilipschitz equivalent to a plane.

The set E in Theorem 1.3 is far from being bizarre or pathological. It will be smooth off a self-similar Cantor set, around which it does some spiralling. It can be approximated nicely by smooth submanifolds, and it agrees with a hyperplane outside of a compact set. The obstruction to the existence of a bilipschitz parameterization will be present locally in the sense that even small neighborhoods of the singular points will not admit bilipschitz parameterizations. By adjusting the parameters we can make this set E especially nice, so that the inverse of the quasiconformal parameterization is Hölder continuous of order as close to 1 as we like, while the mapping itself remains Lipschitz. Analogous statements for Sobolev spaces are also true. On the other hand, by choosing the parameters differently we can build E in such a way that there is no homeomorphism from E onto \mathbb{R}^3 which is locally Hölder continuous of any order $\alpha > 0$ given in advance. See Theorem 5.27 below.

The construction of these examples will be based on “Antoine’s Necklaces” [Mo] and the following lemma.

Lemma 1.4. *If U is an open set in \mathbb{R}^d and K is a closed subset of \mathbb{R}^d , and if K has Hausdorff dimension less than $d - 2$, then any loop in $U \setminus K$ can be contracted in $U \setminus K$ if it can be contracted in U . In particular, $\mathbb{R}^d \setminus K$ is simply connected.*

In other words, a set is invisible in terms of the properties of π_1 of its complement if it is thin enough. This is given in [MRV, Lemma 3.3, p. 9]. Actually, that result dealt only with the simple-connectivity of the relevant domains, but the same proof applies to this formulation. See also [LV] and [SS, p. 506].

This lemma provides us with a necessary condition for a set to be bilipschitz equivalent to \mathbb{R}^d , since any such metric space would have to enjoy a similar property, and it is this necessary condition that the examples in Theorem 1.3 will violate. That is, E will be constructed

in such a way that it contains a compact set with Hausdorff dimension less than 1 whose complement is not simply connected.

Let us consider briefly the conceptual ramifications of the existence of sets like E in Theorem 1.3. One could say that the above list of geometric necessary conditions for bilipschitz equivalence with a Euclidean space is simply too small, and that it should be enlarged to include the property in Lemma 1.4 (and others). I am inclined toward a different view. I am not optimistic that there are nice geometric characterizations of sets which are bilipschitz equivalent to a Euclidean space (although I would not be surprised if there were nice results in more narrowly focussed situations, in which π_1 conditions as in Lemma 1.4 would appear naturally). I think that the existence of sets like E in Theorem 1.3 means that one should view bilipschitz parameterizations as luxuries which are desirable but not reasonable to expect in general and also not crucial. In other words, these examples are sufficiently nice that they should be accommodated rather than excluded. Instead of looking for more stringent criteria to ensure the existence of a bilipschitz parameterization one can look for simpler conditions which imply a lot of good structure (of the type that these examples enjoy) if not an actual bilipschitz parameterization. In this connection the notion of uniform rectifiability (as in [DS4]) is very natural, because it incorporates many (but by no means all) of the nice features of bilipschitz equivalence with a Euclidean space while being much more flexible and easier to detect. See also [DS2], [DS3], [DS5], and [Se5].

Let us now consider analogous issues for some general “conformal” deformations of Euclidean geometry.

Definition 1.5. *A continuous weight ω on \mathbb{R}^d is a nonnegative continuous function whose zero set has Lebesgue measure zero. If A is a measurable subset of \mathbb{R}^d then $\omega(A)$ will be used to denote $\int_A \omega$. A continuous weight ω on \mathbb{R}^d is doubling (or satisfies a doubling condition) if there is a constant C so that $\omega(2B) \leq C\omega(B)$ for all balls B in \mathbb{R}^d , where $2B$ denotes the ball with the same center as B but twice the radius. We shall view ω as defining a measure on \mathbb{R}^d , and also a conformal deformation of Euclidean geometry. To this end we associate to ω the (possibly degenerate) distance function $D_\omega(x, y)$, which is the infimum of the ω -length of all rectifiable paths in \mathbb{R}^d which join x to y . (The ω -length of a path γ is defined to be $\int_\gamma \omega^{1/d} ds$, where ds denotes arclength measure.) We say that ω is a strong A_∞ continuous weight if it is doubling and if there is a $C > 0$ so that $C^{-1} \omega(B_{x,y})^{1/d} \leq D_\omega(x, y) \leq C \omega(B_{x,y})^{1/d}$*

for all $x, y \in \mathbb{R}^d$, where $B_{x,y}$ is the smallest closed Euclidean ball which contains x and y .

Constant functions provide trivial examples of strong A_∞ continuous weights. Less trivial examples are given by $\omega(x) = |x|^a$, $a \geq 0$. On the other hand the strong- A_∞ condition prevents ω from vanishing on a nontrivial line segment, or on any rectifiable curve for that matter, since $D_\omega(x, y)$ would then vanish for some $x \neq y$. However, strong A_∞ weights can vanish on Cantor sets of large Hausdorff dimension, as in Proposition 4.4.

The “strong- A_∞ ” condition was originally defined in [DS1]. The strange-looking name was motivated by an observation that is recalled below. The main point is that the strong- A_∞ condition by itself implies that the metric D_ω has many nice properties, without any smoothness assumptions on ω or anything like that. To understand this condition better it is helpful to consider $\delta_\omega(x, y) = \omega(B_{x,y})^{1/d}$ as some kind of distance function in its own right. Specifically, it is a quasimetric, which means that it satisfies all the conditions normally required of a metric except that the triangle inequality should be weakened to $\delta_\omega(x, z) \leq C(\delta_\omega(x, y) + \delta_\omega(y, z))$ for some C and all x, y, z . This condition is easy to verify using the doubling property of ω . Moreover, δ_ω is “quasiconformally” equivalent to the Euclidean metric, in the sense that its balls are approximately the same as Euclidean balls. More precisely, if $x \in \mathbb{R}^d$ and $r > 0$ are given, and if $R > 0$ is chosen so that $\omega(B(x, R)) = R^d$, then

$$(1.6) \quad B(x, C^{-1}R) \subseteq \{y \in \mathbb{R}^d : \delta_\omega(x, y) < r\} \subseteq B(x, CR).$$

Here C depends on the doubling constant of ω but not on x, r or R . On the other hand, R and r can be wildly different from each other, and either can be larger than the other. This fact (1.6) is easy to derive from the definition of δ_ω and the doubling condition on ω , and one can verify also that the ω -diameters of $B(x, C^{-1}R)$ and $B(x, CR)$ are both approximately equal to r , *i.e.*, they are bounded from above by $C'r$ and from below by $C'^{-1}r$ for some constant C' . The strong- A_∞ condition says that the geodesic distance D_ω is comparable in size to δ_ω , so that D_ω has these features too.

If ω is a strong A_∞ continuous weight, then the metric space (\mathbb{R}^d, D_ω) enjoys many of the same properties as ordinary Euclidean space. For instance, if β is a D_ω -ball, then there is a Euclidean (and hence topological) ball B containing β whose D_ω -diameter is bounded

by a constant times that of β . This can be derived from (1.6). Also, there is a constant $C > 0$ so that

$$(1.7) \quad C^{-1} r^d \leq \omega(\{y \in \mathbb{R}^d : D_\omega(x, y) < r\}) \leq C r^d,$$

because of (1.6) and the doubling condition. This is analogous to the regularity condition 1.2.

It was observed in [DS1] that the strong- A_∞ condition implies the (much older) A_∞ condition that there exist constants $p > 1$ and $C > 0$ such that

$$(1.8) \quad \left(\frac{1}{|B|} \int_B \omega^p\right)^{1/p} \leq C \frac{1}{|B|} \int_B \omega,$$

for all balls B in \mathbb{R}^d , where $|B|$ denotes the Lebesgue measure of B . In other words, the strong- A_∞ condition implies that $\omega \in L^p_{loc}$ with $p > 1$, and with a uniform and scale-invariant bound. This is basically a reformulation of a result of Gehring [Ge], and there is also a bound on averages of (small) negative powers of ω . (See [Ga] and [Jr] for more information about A_∞ weights.) These bounds on ω imply that D_ω -geometry is closer to Euclidean geometry than one might think. For instance, they imply the “uniform rectifiability” condition that every D_ω -ball in \mathbb{R}^d has a definite proportion (with respect to the ω -volume) which is uniformly bilipschitz equivalent to a subset of \mathbb{R}^d with the Euclidean metric. Alternatively one can use (1.8) and its cousins to obtain Sobolev space estimates on the identity mapping as a map from $(\mathbb{R}^d, |x - y|)$ to (\mathbb{R}^d, D_ω) and vice-versa. (Hölder continuity can be derived directly from the doubling condition, as in Proposition 4.22.)

The main result of [DS1] states that the analogues of the usual Sobolev and Poincaré inequalities on Euclidean spaces are also true for (\mathbb{R}^d, D_ω) when ω is a strong A_∞ weight. In view of all these common features between D_ω and the Euclidean metric when ω satisfies the strong- A_∞ condition it is natural to ask the following.

QUESTION 1.9. ([DS1]) *If ω is a strong A_∞ continuous weight on \mathbb{R}^d , then is \mathbb{R}^d equipped with the metric D_ω bilipschitz equivalent to \mathbb{R}^d equipped with the Euclidean metric?*

This is equivalent to asking whether there is a quasiconformal mapping $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$ whose jacobian J satisfies $C^{-1}\omega \leq J \leq C\omega$ for some constant C . It is not hard to show that the doubling and strong- A_∞

conditions on ω are necessary in order for such a quasiconformal mapping to exist. Continuity of ω is not necessary, though, but it is also not required in the “true” definition of a strong A_∞ weight in [DS1]. The restriction to continuous weights here allows us to avoid some technical complications which are not necessary for the purposes of this paper, but discontinuous weights are indispensable for the complete story. The general case of discontinuous weights is also discussed in [Se4].

The answer to Question 1.9 is no.

Theorem 1.10. *There is a strong A_∞ continuous weight ω on \mathbb{R}^3 for such that (\mathbb{R}^3, D_ω) is not bilipschitz equivalent to \mathbb{R}^3 equipped with the Euclidean metric. We can also take ω to be bounded. Moreover, there is a constant L_1 so that for every pair of distinct points $x, y \in \mathbb{R}^3$ there is a closed subset W of \mathbb{R}^3 which contains x and y and such that (W, D_ω) is L_1 -bilipschitz equivalent to a closed Euclidean 3-ball equipped with the standard Euclidean metric.*

The counterexamples for Theorem 1.10 will be based on the violation of the property in Lemma 1.4, and they can be taken to be powers of the distance to a set in \mathbb{R}^3 like an Antoine’s necklace. The resulting spaces (\mathbb{R}^3, D_ω) will then be embedded into \mathbb{R}^4 to produce the counterexamples for Theorem 1.3. In fact Theorem 1.3 has to include Theorem 1.10 (modulo continuity of the weight ω), because one can show that any set E as in Theorem 1.3 must be bilipschitz equivalent to (\mathbb{R}^3, D_ω) for some (not necessarily continuous) strong A_∞ weight ω . In our case this fact will come from the construction. (See Lemma 5.25.)

An amusing consequence of Theorem 1.10 is that the Sobolev-Poincaré inequalities obtained in [DS1] are nontrivial in the sense that they cannot be reduced to the Euclidean case by a change of variables. (However, for these particular examples there are much more direct arguments than the ones in [DS1] for the general case.)

By adjusting the parameters of the construction one can obtain various improvements of Theorem 1.10, just as for Theorem 1.3. See Theorem 4.20 and the remarks that follow it.

Theorem 1.10 raises the same kind of conceptual issues as Theorem 1.3 does. If one really wants to classify the weights that arise as jacobians of quasiconformal mappings (or even give nice criteria for this to happen), then one has to strengthen the strong- A_∞ condition.

Alternatively, one can take the view that a Euclidean space with its geometry deformed by a strong A_∞ weight as above is a rather nice object, even if it is not quite as nice as the standard Euclidean space, and that it should be accommodated rather than excluded. Again, I am inclined toward this “big tent” philosophy.

Let us consider now the following weaker version of Question 1.9.

QUESTION 1.11. *If ω is a strong A_∞ continuous weight on \mathbb{R}^d , then is (\mathbb{R}^d, D_ω) bilipschitz equivalent to a subset of some \mathbb{R}^n equipped with the Euclidean metric?*

It is known that the answer to Question 1.11 is yes for many strong A_∞ weights. The precise statement is complicated, but the main result in [Se4] says that if a strong A_∞ weight ω has the property that it can be made smaller to a nontrivial extent and remain a strong A_∞ weight, then the answer to Question 1.11 is yes for ω . This criterion is not satisfied by all strong A_∞ weights, but it is satisfied by the important subclass of A_1 weights (see Definition 2.8), and any strong A_∞ weight can be approximated by weights which satisfy this stronger condition. Roughly speaking, the strong A_∞ weights which do not satisfy this stronger condition are sitting at the boundary of the space of strong A_∞ weights.

Nonetheless, the answer to Question 1.11 is no.

Theorem 1.12. *There is a strong A_∞ weight on some \mathbb{R}^d such that (\mathbb{R}^d, D_ω) is not bilipschitz equivalent to any subset of any \mathbb{R}^n (equipped with the Euclidean metric).*

Of course the statement of Theorem 1.12 is much stronger than that of Theorem 1.10, except for the knowledge of the dimension d . However, the example used to prove Theorem 1.10 will have the property that one can embed (\mathbb{R}^d, D_ω) bilipschitzly in \mathbb{R}^4 , so that we actually know that there are counterexamples to Question 1.9 which are not pathological for Question 1.11. In fact the examples for Theorem 1.10 satisfy the stronger version of the strong- A_∞ condition in [Se4] mentioned above, and they are nicer, simpler, and more flexible than the examples for Theorem 1.12. The examples for Theorem 1.12 are not based on Lemma 1.4 or Antoine’s necklaces or anything like that, and unfortunately they are not very explicit.

Strong A_∞ weights are related to an abstract version of Question

1.11 in an interesting way. To explain this we need another definition.

Definition 1.13. *A metric space $(M, d(\cdot, \cdot))$ satisfies a doubling condition if there is a constant C so that any ball in M can be covered by at most C balls of half the radius.*

This condition shows up in [A1], [A2], [A3], [CW1], [CW2], and [Gr], but with different names. It provides a kind of boundedness on the geometry of M . It is satisfied by Euclidean spaces, and hence their subspaces, and it is not hard to show that (\mathbb{R}^d, D_ω) satisfies this doubling condition when ω is a strong A_∞ continuous weight, because of the doubling condition on ω in Definition 1.5.

The abstract version of Question 1.11 asks whether any metric space which satisfies a doubling condition must be bilipschitz equivalent to a subset of some \mathbb{R}^n . The answer is known to be no, and this point will be discussed more in Section 7. There is very nice positive result due to Assouad ([A3, Proposition 2.6, p. 436], see also [A1], [A2]), however.

Theorem 1.14 (Assouad). *If $(M, d(\cdot, \cdot))$ is a metric space which satisfies a doubling condition, then for each $\alpha \in (0, 1)$ the metric space $(M, d(\cdot, \cdot)^\alpha)$ is bilipschitz equivalent to a subset of some \mathbb{R}^n .*

Thus, given a metric space which satisfies a doubling condition, one can perturb the metric in order to get a space which embeds into a Euclidean space bilipschitzly. Although this perturbation is small in some ways, it does have the unfortunate feature that it enlarges the class of Lipschitz functions on the space enormously, in such a way as to destroy any sort of differentiability almost everywhere theorem as one has on Euclidean space. In fact the counterexamples to the $\alpha = 1$ case of Theorem 1.14 show that such destruction is necessary. (See Section 7.)

Note that the doubling condition is necessary for the embedding in Theorem 1.14 to exist (for any $\alpha > 0$). Also, Theorem 1.14 implies that every metric space which satisfies a doubling condition admits a quasisymmetric embedding into some Euclidean space. (Basically a quasisymmetric embedding is one that distorts relative distances by a bounded amount, like quasiconformal mappings on \mathbb{R}^d , but we shall not need the precise definition here.)

There is a converse to the fact that the deformations of Euclidean

spaces associated to strong A_∞ continuous weights satisfy the doubling condition in Definition 1.13.

Theorem 1.15. *If $(M, d(\cdot, \cdot))$ is a metric space which satisfies a doubling condition, then there is a positive integer n and a strong A_∞ continuous weight ω on \mathbb{R}^n such that $(M, d(\cdot, \cdot))$ is bilipschitz equivalent to a subset of $(\mathbb{R}^n, D_\omega(\cdot, \cdot))$.*

Thus Question 1.11 is equivalent to its abstract version for metric spaces which satisfy doubling conditions, and so the negative answer to the abstract version (which was known to Assouad) implies the negative answer to Question 1.11 itself (Theorem 1.12).

The next section provides more information about related results and open problems but is not essential to the rest of the paper. Antoine's necklaces are reviewed in Section 3, and they are used to prove Theorems 1.10 and 1.3 in Sections 4 and 5, respectively. Theorem 1.15 is proved in Section 6, and Theorem 1.12 is derived from it in Section 7. Section 8 is devoted to regular mappings, which are more flexible cousins of bilipschitz mappings.

2. Some related results and problems.

One of the most interesting examples of a set which is not bilipschitz equivalent to the expected standard model is provided by the following.

Theorem 2.1 (Edwards). *There is a finite polyhedron P (in a Euclidean space of modest dimension) which is homeomorphic to the 5-sphere S^5 and which contains a polygonal arc Γ such that $P \setminus \Gamma$ is not simply connected. In fact, for no open set $U \subseteq P$ which intersects Γ is it true that $U \setminus \Gamma$ is simply connected.*

Corollary 2.2. *If P and Γ are as above, then P is not bilipschitz equivalent to S^5 , and no open subset U of P which intersects Γ is bilipschitz equivalent to an open subset of \mathbb{R}^5 .*

The corollary follows from Edwards' theorem and Lemma 1.4 (as was explained in [SS, Remark (b), p. 504]). Lemma 1.4 implies that a closed set in S^5 with Hausdorff dimension less than 3 must always

have simply-connected complement, so that any homeomorphism from P onto S^5 takes Γ to a set with Hausdorff dimension at least 3. Thus no such homeomorphism can be Lipschitz or even Hölder continuous of order larger than $1/3$. Similar arguments apply to open sets U as in the corollary.

Edwards' example is of the following form. Let H be a compact smooth 3-dimension manifold which is not simply connected but which does have the same (integer) homology as S^3 . Many examples of such manifolds are known to exist, but note that if H were also simply connected then the (as yet unproven) Poincaré conjecture would imply that H is diffeomorphic to S^3 . Also, recall that a homology sphere of dimension 1 or 2 has to be a standard sphere. Now take P to be the "join" of H with a copy of S^1 . That is, P contains a copy of H and of S^1 , and for each point in H and each point in S^1 , P contains a line segment which joins the two points, these line segments are disjoint except possibly for necessary intersections at their endpoints in H and S^1 , and P does not contain any other points. We can do this in such a way that P is a polyhedron, by applying this to polyhedral copies of H and S^1 instead of smooth copies. The curve Γ mentioned above is simply the copy of S^1 inside P that we are using. It is easy to see that $P \setminus \Gamma$ is homotopy-equivalent to the polyhedral copy of H inside P , and hence is not simply connected. A little more thought gives the local version of this lack of simple-connectivity mentioned above.

The deep part of Edwards' theorem is that H can be chosen so that P is homeomorphic to S^5 . To understand this better we need to recall the notion of a suspension of a space. The *suspension* of H is just like the join of H with S^1 , except that we use S^0 instead. That is, we take two points off of H , and then we build a new space by joining each of these points with a line segment to each point in H . For example, if we take the suspension of a sphere, then we get a sphere of one higher dimension. One can check that P is the same as the suspension of the suspension (the double suspension) of H . Thus if we used S^3 instead of H to build P , then it would be immediately clear that P would be homeomorphic to S^5 , even piecewise-linearly. In the case of a non-simply connected homology sphere H this is less clear. The suspension of H is not a manifold, because the two cones points are not manifold points. The idea is that it is almost a manifold, and in fact almost a sphere, so that the second suspension makes it into a topological manifold. This is not at all obvious. Edwards proved that this worked for some homology spheres, and Cannon then proved that

all homology spheres work. Edwards then proved a very general theorem for recognizing when a topological space is a topological manifold (which is the key point here). See [C1], [C2], [Da], and [E].

Let us think a little about the geometric properties of the space P in Theorem 2.1. It is a finite polyhedron, *i.e.*, a finite union of (5-dimensional) simplices, it is homeomorphic to the 5-sphere, but it is not bilipschitz equivalent to the 5-sphere, let alone piecewise-linearly equivalent to S^5 . Nonetheless P has many of the same geometric properties as the 5-sphere. It is a 5-dimensional regular set in the sense of Definition 1.1, modulo the necessary restriction to $r < 1$ (say) in (1.2), since P is compact. It is easy to see that the usual Sobolev and Poincaré inequalities hold on P , and that other aspects of analysis work as well. It is not hard to prove that P also enjoys the following property: there is a $C > 0$ so that if $x \in P$ and $0 < r < C^{-1}$, then there is an open subset V of P such that $P \cap B(x, r) \subseteq V \subseteq P \cap B(x, Cr)$ and V is homeomorphic to a 5-ball. (This is true for any finite polyhedron which is a topological manifold -see [Se6, Section 11]- but in the case of P it is a little easier to see using its special form.) One can strengthen this statement to include uniform scale-invariant bounds on the moduli of continuity of the homeomorphisms between these sets V and the unit ball in \mathbb{R}^5 and their inverses. (Roughly speaking, there are really only two different V 's that one has to worry about, modulo some simple operations like dilating.)

Thus P has many of the same nice metric properties as a smooth manifold, even though it is not bilipschitz equivalent to one. It is in very much the same spirit as the examples mentioned in the introduction that will be constructed in the next sections, but it is much more impressive, since it is even a finite polyhedron. On the other hand, the examples constructed below have the advantages that they are easier to verify, they work in lower dimensions (3 instead of 5), and their parameterizability properties are better controlled (in terms of the existence or nonexistence of Hölder continuous coordinates, for instance; see Theorems 4.20 and 5.27).

Let us now consider more fully the question of the dimensions in which these various types of examples exist.

PROBLEM 2.3. *Are there analogues of Theorems 1.3 and 1.10 in dimension 2? For which dimensions d does Theorem 1.12 hold?*

The Edwards' examples work in all dimensions greater or equal

than 5, but the construction does not make sense in lower dimensions (because homology spheres are true spheres in dimensions 1 and 2). The examples given here for Theorems 1.3 and 1.10 could be adapted to higher dimensions, but the method breaks down completely in dimension 2. Of course many things are better in two dimensions, and there could even be positive results in that case. (For the record, these questions all degenerate into triviality when $d = 1$, because one can use arclength parameterizations to build the required mappings.)

As for Theorem 1.12, the proof will give a value of d which is computable in principle but whose smallness is not clear. See Section 7, and see Section 8 for some related questions.

There are some positive results which are special to dimension 2 but which address a slightly different question. That is, there are some reasonable geometric criteria for a 2-dimensional metric space to admit a quasisymmetric parameterization (which need not be bilipschitz). See [Se2, Section 5], [DS3, Section 6], and [HK]. The reason that dimension 2 is special here is that one has the uniformization theorem which can provide a conformal parameterization right from the start. (This is analogous to the special role of the arclength parameterization in dimension 1.) One still has the problem of passing from the infinitesimal conformality condition to distortion estimates at large scales, but this can be managed, and in fact the results of Heinonen and Koskela [HK] deal effectively with this problem under very general circumstances. Nonetheless, the difficulty posed by the absence of the uniformization theorem in higher dimensions remains, and in fact there are examples [S6] which show that the analogues of the 2-dimensional results fail in higher dimensions.

There is another special case of these questions which is not addressed by the known examples.

Definition 2.4. *Let M be a hypersurface in \mathbb{R}^{d+1} , and assume a priori that M is smooth and nice at ∞ . Given $\varepsilon \geq 0$, we say that M is ε -flat if*

$$(2.5) \quad D_M(x, y) \leq (1 + \varepsilon)|x - y|, \quad \text{for all } x, y \in M,$$

where $D_M(x, y)$ denotes the geodesic distance on M (the infimum of the lengths of all paths on M which join x to y), and if

$$(2.6) \quad (1 - \varepsilon)\nu_d r^d \leq H^d(M \cap B(x, r)) \leq (1 + \varepsilon)\nu_d r^d,$$

for all $x \in M$ and $r > 0$, where ν_d denotes the volume of the unit ball in \mathbb{R}^d .

It is easy to see that M must be a hyperplane if it is ε -flat with $\varepsilon = 0$.

PROBLEM 2.7. *If $M \subseteq \mathbb{R}^{d+1}$ is ε -flat, and if ε is small enough, then must it be true that M is bilipschitz equivalent to \mathbb{R}^d with a bilipschitz constant that depends only on d ? Is it bilipschitz equivalent to \mathbb{R}^d with a constant which tends to 1 uniformly as $\varepsilon \rightarrow 0$?*

If M admits a K -bilipschitz parameterization by \mathbb{R}^d , then (2.5) and (2.6) must hold with $\varepsilon \rightarrow 0$ as $K \rightarrow 1$, as one can easily check. The point of Problem 2.7 is to know whether the converse holds, with estimates which do not depend quantitatively on the a priori assumptions.

It turns out that the ε -flatness condition implies many other flatness conditions that are necessary for the existence of such a bilipschitz parameterization. For instance, it implies the existence of local coordinates with good estimates in terms of Hölder spaces and Sobolev spaces, as long as one stays away from the Lipschitz class (which is the question of Problem 2.7), and when $d = 2$ there is a positive result for quasisymmetric parameterizations, based on the uniformization theorem. (See [Se2].) Also, there are some sufficient conditions for the existence of bilipschitz parameterizations (with constant close to 1) in terms of conditions which are stronger than ε -flatness but which are roughly of the same order of magnitude [T2]. When $d = 2$ there is a nice sufficient condition in terms of the L^2 norm of the curvature being small [T1]. This condition is stronger than ε -flatness, and has approximately the same relationship with ε -flatness that the Sobolev space $W^{1,2}(\mathbb{R}^2)$ has with $\text{BMO}(\mathbb{R}^2)$ via the Sobolev embedding. When $d > 2$ it is not known whether the L^d norm of the principal curvatures being small is sufficient to ensure the existence of a bilipschitz parameterization. This curvature condition is the natural one, in that it scales correctly and implies ε -flatness.

There are also some nice equivalent characterizations of ε -flatness with small ε , in terms of the Gauss map having small BMO norm (and hence small oscillation in a certain sense) [Se3] and in terms of singular integral operators and Clifford analysis [Se1].

Despite all these good properties of ε -flat surfaces, I am pessimistic about Problem 2.7. However, I do not know any counterexamples, and

the type of examples given in this paper will not work for this.

One could also consider small constant versions of Question 1.9. When $d = 2$ one can reduce Problem 2.7 to such a question using [Se2].

There is another variant of Question 1.9 which has some hope. For this it is much less reasonable to deal only with continuous weights, and so we use the correct general definition.

Definition 2.8. *A nonnegative measurable function ω on \mathbb{R}^d which is positive almost everywhere is called an A_1 weight if there is a constant C such that*

$$(2.9) \quad \frac{1}{|B|} \int_B \omega(y) dy \leq C \operatorname{ess\,inf}_{y \in B} \omega(y),$$

for all balls B in \mathbb{R}^d .

The A_1 condition implies the strong- A_∞ condition, although the definition of $D_\omega(\cdot, \cdot)$ needs to be modified since we are not assuming continuity. (This issue is treated thoroughly in [Se4].) The A_1 condition is much stronger; it forbids any vanishing, while the strong- A_∞ condition forbids only certain kinds of vanishing. A simple example of an A_1 weight on \mathbb{R}^d is $\omega(x) = |x|^{-s}$ for $0 \leq s < d$. Similarly, an A_1 weight can blow up along a submanifold, but not as rapidly.

PROBLEM 2.10. *Does Question 1.9 have an affirmative answer for A_1 weights?*

Question 1.11 does have an affirmative answer in the case of A_1 weights [Se4]. The situation for Question 1.9 is unclear, but the method for producing an example as in Theorem 1.10 definitely does not work in the case of A_1 weights. (The whole point will be to make the weight vanish on a certain set, which an A_1 weight cannot do.)

See [Ga] and [Jr] for more information about A_1 weights.

There is a notion of “regular mappings” which is weaker (and more flexible) than bilipschitzness and for which there are some interesting results and problems related to Questions 1.9 and 1.11 and Problems 2.3 and 2.10. See Section 8.

For related questions and examples pertaining to strong- A_∞ weights, see [Se4], especially Sections 4 and 5.

3. Antoine's necklaces.

This section will be devoted to the necklaces of Antoine, which are Cantor sets in \mathbb{R}^3 whose complements are not simply connected. These sets will be used heavily in the next two sections, to construct the examples promised in the introduction. The basic reference for this section is [Mo, Chapter 18]. (See also [B] for higher-dimensional versions of Antoine's necklaces.)

Let k be a reasonably large positive integer. This is a parameter which is at our disposal; it needs to be reasonably large for the construction to work nicely ($k \geq 10^7$ is fine), but there is no reason for us to try to choose k as small as possible anyway. It will be important later for us to have the option of taking k to be arbitrarily large. The construction is a little nicer when k is even.

Fix a circle Γ_0 in \mathbb{R}^3 with radius 1. Let P_0 be a collection of k equally spaced points on Γ_0 . For each $p \in P_0$ choose a circle $\gamma_0(p)$ in \mathbb{R}^3 centered at p in such a way that all the $\gamma_0(p)$'s have the same radius $\rho(k)$ and the following properties are satisfied:

- (3.1) $k^{-1} \leq \rho(k) \leq 2\pi k^{-1}$,
- (3.2) $\text{dist}(\{p\} \cup \gamma_0(p), \{q\} \cup \gamma_0(q)) \geq (100k)^{-1}$, when $p \neq q$,
- $\gamma_0(p)$ and $\gamma_0(q)$ are linked (as circles in \mathbb{R}^3)
- (3.3) if and only if p and q are adjacent to each other
(as elements of P_0).

It is not hard to check that these circles actually exist. (It is helpful to remember that by taking k large we have that Γ_0 is very flat at the scale of k^{-1} . If we did not want adjacent circles to link, but instead to touch at a single point, then we would want to take $\rho(k) \approx \pi k^{-1}$. As it is, we need to take them a little larger, but (3.1) gives us enough room. We could also impose additional symmetry requirements on the $\gamma_0(p)$'s, but we shall not bother.) The union of the $\gamma_0(p)$'s, $p \in P_0$, link together in a necklace near Γ_0 , and we denote their union by $N(\Gamma_0)$. See [Mo, Figures 18.1 and 18.2, p. 127-8] for excellent pictures.

To build Antoine's necklace we shall replace each $\gamma_0(p)$ with a smaller copy of $N(\Gamma_0)$, and then replace each component of the resulting set with a smaller copy of $N(\Gamma_0)$, and so forth. To do this carefully we need to "mark" our circles.

Recall that a similarity on \mathbb{R}^n is an affine mapping which is a combination of a translation, (nonzero) dilation, and orthogonal transformation.

Definition 3.4. *A marked circle in \mathbb{R}^3 is a circle Γ together with an orientation-preserving similarity ϕ on \mathbb{R}^3 (called the marking) such that $\phi(\Gamma_0) = \Gamma$. We shall often let Γ denote both a marked circle (with ϕ not mentioned explicitly) and the circle as a set.*

Of course we shall always take Γ_0 to be marked in the obvious way (with $\phi =$ the identity).

Choose now markings ϕ_p for the circles $\gamma_0(p)$'s, $p \in P_0$. The choice of the markings is insignificant, but they need to be fixed once and for all. With the selection of these markings we can view $N(\Gamma_0)$ as a union of marked circles.

If Γ is a marked circle, with marking ϕ , then we set

$$(3.5) \quad N(\Gamma) = \phi(N(\Gamma_0)).$$

We shall consider $N(\Gamma)$ to be a union of marked circles, with the markings induced by ϕ in the obvious way. These circles are also naturally labelled by P_0 .

If Γ and Γ' are marked circles and if $\psi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is an orientation-preserving similarity which takes Γ to Γ' in a way which is compatible with the markings, then

$$(3.6) \quad \psi(N(\Gamma)) = N(\Gamma'),$$

and the markings of the constituent circles of $N(\Gamma_1)$ and $N(\Gamma_2)$ correspond under ψ in the obvious way.

If $E \subseteq \mathbb{R}^3$ is a finite union of marked circles, then we take $N(E)$ to be the union of $N(\alpha)$, where α runs through the constituent circles of E . Thus if \mathcal{C} denotes the set of subsets of \mathbb{R}^3 which are finite unions of marked circles, then N defines a mapping from \mathcal{C} to itself. Note that if $E, E' \in \mathcal{C}$, then $E \cup E' \in \mathcal{C}$ and $N(E \cup E') = N(E) \cup N(E')$. Also, orientation-preserving similarities act on \mathcal{C} in the obvious way, and this action commutes with N , by (3.6).

Set $A_0 = \Gamma_0$, $A_1 = N(\Gamma_0)$, and define A_l in general by $A_l = N^l(\Gamma_0)$, where N^l denotes the l^{th} power of N , viewed as a mapping on \mathcal{C} . That is, we view Γ_0 and the A_l 's as elements of \mathcal{C} , so that we can iterate N in this way. Each A_l is a union of k^l marked circles of radius

$\rho(k)^l$, and A_{l+1} is the union of the $N(\alpha)$'s, where α runs through the circles which make up A_l . The Hausdorff limit of the A_l 's will give a necklace of Antoine. Before we deal with the Hausdorff limit we should record some simple preliminary facts.

Lemma 3.7. *For each j, l with $0 \leq j \leq l$ we have that $A_l = \cup_{\alpha} N^{l-j}(\alpha)$, where the union is taken over all the constituent (marked) circles α in A_j , and the constituent circles of these two collections are marked in the same way.*

In other words, A_l and $\cup_{\alpha} N^{l-j}(\alpha)$ are the same as elements of \mathcal{C} . This follows easily from the definitions.

Lemma 3.8. *If $k \geq 10^7$, then every marked circle Γ in \mathbb{R}^3 satisfies*

$$\sup_{x \in N(\Gamma)} \text{dist}(x, \Gamma) \leq \rho(k) \text{radius}(\Gamma) < 10^{-6} \text{radius}(\Gamma).$$

To prove this one reduces to the case where $\Gamma = \Gamma_0$ and uses (3.1) and the definitions.

From now on we assume that k is at least 10^7 . Given a circle Γ in \mathbb{R}^3 , let $\tau(\Gamma)$ be the small solid torus containing Γ given by

$$(3.9) \quad \tau(\Gamma) = \{x \in \mathbb{R}^3 : \text{dist}(x, \Gamma) \leq 10^{-5} \text{radius} \Gamma\}.$$

If Γ and Γ' are two circles and ψ is a similarity that maps one onto the other, then $\psi(\tau(\Gamma)) = \tau(\Gamma')$. If E is a finite union of circles in \mathbb{R}^3 , then we define $\tau(E)$ to be the union of the sets obtained by applying τ to the constituent circles of E .

Lemma 3.10. $\tau(N^l(\Gamma)) \subseteq \tau(\Gamma)$ for all marked circles Γ and all $l \geq 0$.

This follows from Lemma 3.8 and the definitions by estimating the relevant geometric series.

Lemma 3.11. $\tau(A_j) \supseteq \tau(A_l)$ when $0 \leq j \leq l$.

This is an immediate consequence of Lemmas 3.10 and 3.7.

We now define our Antoine's necklace A by

$$(3.12) \quad A = \bigcap \tau(A_l).$$

This is a compact subset of \mathbb{R}^3 . It is the same as the Hausdorff limit of the A_l 's, but this definition is a little easier to use.

This set A is totally disconnected. To quantify this disconnectedness define $\tau_a(\Gamma)$ for $a > 0$ and Γ a circle in \mathbb{R}^3 to be the solid torus containing Γ given by

$$(3.13) \quad \tau_a(\Gamma) = \{x \in \mathbb{R}^3 : \text{dist}(x, \Gamma) \leq a \cdot 10^{-5} \text{ radius } \Gamma\}.$$

Thus if $a \geq 1$ then this is somewhat larger than $\tau(\Gamma)$.

Lemma 3.14. *Let α and α' be two distinct circles among those which make up A_l , $l \geq 1$. Then $\tau_{50}(\alpha)$ and $\tau_{50}(\alpha')$ are disjoint.*

To see this notice first that

$$(3.15) \quad \text{dist}(\gamma_0(p), \gamma_0(q)) > 700^{-1} \text{ radius } \gamma_0(p)$$

when $p, q \in P_0$, $p \neq q$, because of (3.2) and (3.1). This implies that $\tau_{50}(\gamma_0(p))$ and $\tau_{50}(\gamma_0(q))$ are disjoint when $p \neq q$, which is the same as the $l = 1$ case of Lemma 3.14. The case where $l > 1$ and α and α' have the same "parent" in A_{l-1} reduces to the $l = 1$ case by a similarity transformation. Suppose now that $l > 1$ and α and α' have different parents in A_{l-1} . Let δ be the parent of α , so that α is one of the circles in $N(\delta)$. Then

$$(3.16) \quad \tau_{50}(\alpha) \subseteq \tau_{50}(\delta)$$

because of Lemma 3.8. Thus the disjointness of $\tau_{50}(\alpha)$ and $\tau_{50}(\alpha')$ would follow if we could establish the disjointness of $\tau_{50}(\delta)$ and $\tau_{50}(\delta')$, where δ' is the parent of α' . By iterating this procedure we can reduce to the previous case where the two circles have the same parent. This proves the lemma.

For the record:

Theorem 3.17. $\mathbb{R}^3 \setminus A$ is not simply connected.

This is Theorem 4 of [Mo, p. 141].

Corollary 3.18. *No homeomorphism from \mathbb{R}^3 to itself can map A to a set with Hausdorff dimension less than 1.*

This follows from Theorem 3.17 and Lemma 1.4. It is very amusing, since A itself is homeomorphic to a standard Cantor set. Thus we cannot find a global homeomorphism on \mathbb{R}^3 which sends A to a standard Cantor set. See also [Mo, Theorem 5, p. 131].

Notice that when $k \rightarrow \infty$ A tends to the circle Γ_0 in the Hausdorff topology. The Hausdorff dimension A is always larger than 1, and it tends to 1 as $k \rightarrow \infty$.

There is also a “local” version of Corollary 3.18.

Corollary 3.19. *If U is an open set in \mathbb{R}^3 which intersects A , then any homeomorphism from U onto another open set in \mathbb{R}^3 must take $U \cap A$ to a set with Hausdorff dimension at least 1.*

Indeed, because of the obvious self-similarity of A , there must be a small copy \tilde{A} of A inside U . Theorem 3.17 implies that there is a small loop in $U \setminus \tilde{A}$ which is contractible in U but not in $U \setminus \tilde{A}$, and so Lemma 1.4 implies the corollary.

For more information about this kind of wildness phenomena see [R].

4. Strong- A_∞ weights.

This section will be devoted to Theorem 1.10 and some variants of it. We shall consider continuous weights of the form

$$(4.1) \quad \omega(x) = \min(1, \text{dist}(x, A)^{3s})$$

on \mathbb{R}^3 , where A is a compact set and $s > 0$ is at our disposal. In fact we shall take A to be the set constructed in the preceding section, with the parameter k also at our disposal. We shall see that this weight satisfies the *strong- A_∞ condition* and has other interesting properties. The typographically simpler weights of the form $\text{dist}(x, A)^{3s}$ would be practically as good, but it is nice to force ω to be constant off a compact set. There is nothing magical about the 3 in the exponent, it merely reflects the fact that we shall use ω as a density on \mathbb{R}^3 , and it simplifies some of the later formulae.

Our first task is to show that ω is a strong A_∞ weight. This point is clearer in a more general setting.

Definition 4.2. A closed set E in \mathbb{R}^n is said to be uniformly disconnected if there is a constant $C_0 > 0$ so that for each $x \in E$ and $r > 0$ there is a set F such that $E \cap B(x, r) \subseteq F \subseteq E \cap B(x, C_0 r)$ and $\text{dist}(F, E) \geq C_0^{-1} r$. (The latter will be considered to hold vacuously when $E = F$.)

In other words, $E \cap B(x, r)$ should be contained in a little island in E which is not too much larger and which is at a definite distance from the rest of E . The standard Cantor set has this property, as well as practically anything constructed in a similar manner, like the set A in the preceding section.

Proposition 4.3. The set A constructed in Section 3 is uniformly disconnected (for all k , which we require to be $\geq 10^7$, as in Section 3).

Proposition 4.4. If $E \subseteq \mathbb{R}^n$ is closed and uniformly disconnected, then the weight $\Omega(x) = \min(1, \text{dist}(x, E)^s)$ is a strong A_∞ continuous weight for all $s > 0$.

Corollary 4.5. If ω is defined on \mathbb{R}^3 as in (4.1), then ω is a strong A_∞ continuous weight for all $s > 0$ and $k \geq 10^7$.

Of course the corollary is an immediate consequence of the propositions.

Let us prove Proposition 4.3. We shall use freely the notations and results from the previous section. All constants C which appear in the argument below will be permitted to depend on k . Indeed, since A approximates a circle in the Hausdorff topology when k gets big, the uniform disconnectedness constant for A has to blow up as $k \rightarrow \infty$.

Let $x \in A$ and $r > 0$ be given, as in Definition 4.2. We may as well assume that $r < 10^{-10} \rho(k)$; if $r \geq 10^{-10} \rho(k)$, then we can simply take $F = A$, and the requirements of Definition 4.2 will be satisfied.

Choose $l \geq 1$ as large as possible so that $r < 10^{-10} \rho(k)^l$, and let α be the constituent circle in A_l such that $x \in \tau(\alpha)$. Let δ be the parent of α in A_{l-1} , and set $F = A \cap \tau(\delta)$. Then $B(x, r) \subseteq \tau(\delta)$, because $\text{dist}(x, \delta) < 10^{-6} \text{radius}(\delta) < 10^{-5} \text{radius}(\delta) - r$ by Lemma 3.8, our choice of l , and the fact that the radius of δ is $\rho(k)^{l-1}$. Thus $F \supseteq B(x, r) \cap A$. On the other hand, $\text{dist}(F, A \setminus F) \geq 10^{-5} \text{radius}(\delta)$ by Lemma 3.14 (applied to δ and its cousins, and with l replaced by $l-1$) when $l-1 \geq 1$. When $l-1 = 0$ we have that $\delta = \Gamma_0$ and $F = A$.

Since $r \geq 10^{-10} \rho(k)^{l+1} \geq C^{-1} \text{radius}(\delta)$ we have that $F \subseteq B(x, Cr)$. Altogether we conclude that F has the properties required in Definition 4.2, and Proposition 4.3 follows.

Now let us prove Proposition 4.4.

Lemma 4.6. *If $E \subseteq \mathbb{R}^n$ is closed and uniformly disconnected, then there is a $C > 0$ so that for each ball B in \mathbb{R}^n there is a subball $B' \subseteq B$ such that $B' \subseteq \mathbb{R}^n \setminus E$ and the radius of B' is at least C^{-1} times the radius of B .*

To prove this we may as well assume that B is centered on E . Indeed, either $(1/2)B$ is disjoint from E , in which case the conclusion of the lemma holds, or it is not, in which case we can find a ball contained in B with exactly half the radius and whose center is an element of E .

With B centered on E let us apply Definition 4.2 with x, r chosen so that $B(x, 2C_0r) = B$. This gives us a set $F \subseteq (1/2)B$ such that $\text{dist}(F, E \setminus F) \geq C_0^{-1}r$. Using this it is easy to check that B has a subball with the required properties.

One consequence of Lemma 4.6 is that E has Lebesgue measure zero. The vulgar reason for this is that the conclusion of Lemma 4.6 implies that E cannot have any points of density (in the sense of Lebesgue). A better reason is that E must even have Minkowski dimension less than n . At any rate we deduce that Ω as in Proposition 4.4 is at least a continuous weight.

Lemma 4.7. *If $E \subseteq \mathbb{R}^n$ is closed and uniformly disconnected, and if Ω is as in Proposition 4.4, then there is a $C > 0$ so that*

$$(4.8) \quad \sup_{2B} \Omega \leq C \frac{1}{|B|} \int_B \Omega,$$

for all balls B in \mathbb{R}^n .

By Lemma 4.6, every ball with radius 1 contains a ball disjoint from E whose radius is larger than some fixed positive number. Since Ω is bounded from below on at least one-half this ball we get that the right side of (4.8) is bounded from below when B has radius at least 1 (by an easy argument). This implies that (4.8) holds (with a suitable constant) when the radius of B is at least 1, since $\Omega \leq 1$ by definition.

Suppose that B has radius less than 1. If $3B$ is disjoint from E , then $\sup_{2B} \Omega \leq C \inf_B \Omega$ for a suitable constant C by the definition

of Ω and simple geometric considerations, and so (4.8) is satisfied. If $3B$ intersects E , then $\sup_{2B} \Omega \leq (5 \text{radius}(B))^s$, while the right side of (4.8) is bounded from below by $C^{-1} \text{radius}(B)^s$ for some constant C , because of Lemma 4.6 and the definition of Ω . In this case also (4.8) holds, and Lemma 4.7 follows.

Lemma 4.7 implies that Ω is doubling and that

$$D_\Omega(x, y) \leq C \Omega(B_{x,y})^{1/n}.$$

(See Definition 1.5.)

Lemma 4.9. *Suppose that $E \subseteq \mathbb{R}^n$ is closed and uniformly disconnected, and let $x, y \in \mathbb{R}^n$ and a curve γ which connects them be given. Then there is a point p on γ such that $|p - x| \leq |y - x|$ and $\text{dist}(p, E) \geq C^{-1}|y - x|$ for some C which does not depend on x, y , or p .*

This is easy to verify using the uniform disconnectedness condition.

To finish the proof of Proposition 4.4 it remains to show that if E and Ω are as above and $x, y \in \mathbb{R}^n$ and a curve γ which connects them are given, then $\int_\gamma \Omega^{1/n} ds \geq C^{-1} \Omega(B_{x,y})^{1/n}$, where $B_{x,y}$ and $\Omega(B_{x,y})$ are as in Definition 1.5. Apply Lemma 4.9 to get a point p as above. This means that there is a ball B centered at p and with radius equal to $C_1^{-1}|y - x|$ such that $2B$ is disjoint from E for some constant C_1 . This condition implies in turn that $\sup_B \Omega \leq C_2 \inf_B \Omega$ for some constant C_2 . Hence

$$(4.10) \quad \begin{aligned} \int_\gamma \Omega^{1/n} ds &\geq \int_{\gamma \cap B} \Omega^{1/n} ds \\ &\geq C_1^{-1} (\inf_B \Omega^{1/n}) |y - x| \geq C^{-1} \Omega(B)^{1/n}. \end{aligned}$$

On the other hand, $\Omega(B)^{1/n} \geq C^{-1} \Omega(B_{x,y})^{1/n}$, since Ω is doubling, $|p - x| \leq |y - x|$, and the radius of B is not too small compared to $|y - x|$, and hence $\int_\gamma \Omega^{1/n} ds \geq C^{-1} \Omega(B_{x,y})^{1/n}$, as desired. Thus Ω satisfies the strong- A_∞ condition, and the proof of Proposition 4.4 is complete.

Notice that for the weight Ω as in Proposition 4.4 the condition (1.8) is much more obvious than for generic strong A_∞ weights, because of Lemma 4.7.

Next, let A be as in Section 3 and ω be as in (4.1), and let us estimate the Hausdorff dimension of A with respect to D_ω in terms of the parameters s and k (from (4.1) and Section 3). For simplicity we shall call this the ω -Hausdorff dimension of A . The main point is to know when this dimension is less than 1 (or even small) so that we can apply Corollary 3.18 to get restrictions on Lipschitz or Hölder continuous maps from (\mathbb{R}^3, D_ω) to \mathbb{R}^3 equipped with the Euclidean metric.

To estimate the ω -Hausdorff dimension of A we need to cover A with little blobs, estimate the ω -diameter of the little blobs, and then bound the usual series. The computation is simplified by the homogeneity and self-similarity properties of A and ω . For each l let \mathcal{A}_l denote the collection of constituent circles in A_l . Then $\tau(A_l) = \cup_{\alpha \in \mathcal{A}_l} \tau(\alpha)$, by definition, and so

$$(4.11) \quad A \subseteq \bigcup_{\alpha \in \mathcal{A}_l} \tau(\alpha)$$

by (3.12). Thus we can use the $\tau(\alpha)$, $\alpha \in \mathcal{A}_l$, as the little blobs.

We need to estimate the ω -diameters of the $\tau(\alpha)$'s, where " ω -diameter" means the diameter with respect to D_ω . Observe first that $A \cap \tau(\alpha) \neq \emptyset$. (Lemma 3.10 is helpful in this regard.) Thus

$$(4.12) \quad \sup_{\tau(\alpha)} \omega \leq (2 \text{ radius}(\alpha))^{3s}$$

by definition of ω . Hence

$$(4.13) \quad \omega\text{-diameter}(\tau(\alpha)) \leq C \text{ radius}(\alpha)^{1+s}$$

by the definition of D_ω . By construction all the $\alpha \in \mathcal{A}_l$ have radius $\rho(k)^l$, and so we get

$$(4.14) \quad \omega\text{-diameter}(\tau(\alpha)) \leq C \rho(k)^{l(1+s)}.$$

Remember also that there are k^l elements of \mathcal{A}_l . Using this fact, (4.14), and the definition of Hausdorff measure we get the following.

Lemma 4.15. *The ω -Hausdorff dimension of A is less or equal than a if $\limsup_{l \rightarrow \infty} k^l \rho(k)^{l(1+s)a} < \infty$.*

It is not hard to see that the estimate on the ω -Hausdorff dimension provided by this lemma is sharp, but we shall not need that fact. (The main point is to use the natural probability measure on A to get lower bounds on ω -Hausdorff measure.)

Although we could go back now and analyze $\rho(k)$ carefully to get as much out of Lemma 4.15 as possible, a much cruder analysis will be sufficient for our purposes.

Lemma 4.16. *The ω -Hausdorff dimension of A is at most $3(1+s)^{-1}$.*

To see this we observe that

$$(4.17) \quad \sup_l k^l \rho(k)^{3l} < \infty.$$

This bound follows from the observation that $\rho(k)^{3l}$ is a constant multiple of the Lebesgue measure of each $\tau(\alpha)$, $\alpha \in \mathcal{A}_l$, so that the left side of (4.17) is dominated by the Lebesgue measure of a compact set in \mathbb{R}^3 . Lemma 4.16 follows immediately from (4.17) and Lemma 4.15. Of course Lemma 4.16 is not at all sharp, but it enjoys the simplicity of providing a bound which does not depend on k .

Lemma 4.18. *For each fixed $s > 0$ the ω -Hausdorff dimension of A is less or equal than $(1+s/2)^{-1}$ for k sufficiently large.*

Recall from (3.1) that $\rho(k) \leq 2\pi k^{-1}$. Thus Lemma 4.15 implies that the ω -Hausdorff dimension of A is less or equal than a if $\limsup_{l \rightarrow \infty} k^{l-l(1+s)a} (2\pi)^{l(1+s)a} < \infty$. If we take $a = (1+s/2)^{-1}$, then $1 - (1+s)a < 0$, and $k^{l-l(1+s)a} (2\pi)^{l(1+s)a} \rightarrow 0$ as $l \rightarrow \infty$ when k is sufficiently large. Lemma 4.18 follows.

We are now ready to Prove Theorem 1.10 and some variants of it. Let us first set some terminology. Let h be a map from an open subset U of \mathbb{R}^3 into \mathbb{R}^3 , viewed as a map from U equipped with the metric D_ω into \mathbb{R}^3 equipped with the Euclidean metric. We shall write this more succinctly as $h : (U, D_\omega(x, y)) \rightarrow (\mathbb{R}^3, |x - y|)$. We say that h is *locally Hölder continuous of order δ* if for each compact set $K \subseteq U$ there is a constant $C = C(K)$ such that

$$(4.19) \quad |h(x) - h(y)| \leq C D_\omega(x, y)^\delta,$$

for all $x, y \in K$. When this holds with $\delta = 1$ we say that h is *locally Lipschitz*.

Theorem 4.20. *Let A and k be as in Section 3, and let ω and $s > 0$ be as in (4.1). Let U be any open subset of \mathbb{R}^3 which intersects A , and let V be any other open subset of \mathbb{R}^3 . (For instance take $U = V = \mathbb{R}^3$.)*

a) *If $s > 2$ then there does not exist a homeomorphism $h : (U, D_\omega(x, y)) \rightarrow (V, |x - y|)$ which is locally Lipschitz.*

b) *For any $s > 0$ there does not exist a homeomorphism $h : (U, D_\omega(x, y)) \rightarrow (V, |x - y|)$ which is locally Lipschitz if k is large enough.*

c) *For any $s > 0$ there does not exist a homeomorphism $h : (U, D_\omega(x, y)) \rightarrow (V, |x - y|)$ which is locally Hölder continuous of order greater than $3(1 + s)^{-1}$.*

Like the lemmas that came before it, the bounds in Theorem 4.20 are not sharp.

For the proof of Theorem 4.20 we shall need the following.

Lemma 4.21. *If $U \subseteq \mathbb{R}^3$ is open, $h : (U, D_\omega(x, y)) \rightarrow (\mathbb{R}^3, |x - y|)$ is locally Hölder continuous of order δ , and $K \subseteq U$ is a compact set with ω -Hausdorff dimension less or equal than a , then the Euclidean Hausdorff dimension of $h(K)$ is less or equal than δa .*

This is a well-known and straightforward consequence of the definitions.

Recall now that Corollary 3.19 says that no homeomorphism $h : U \rightarrow V$ can send $U \cap A$ to a set of Euclidean Hausdorff dimension less than 1. Using this and Lemma 4.21, parts a) and c) of Theorem 4.20 follow from Lemma 4.16, and part b) follows from Lemma 4.18.

Notice that part c) contains a) as a special case, but the point of c) is more the fact that the Hölder exponent of any homeomorphism $h : (\mathbb{R}^3, D_\omega(x, y)) \rightarrow (\mathbb{R}^3, |x - y|)$ has to go to 0 as s gets large. By contrast we have the following simple and well-known fact.

Proposition 4.22. *If ω is any strong A_∞ weight on \mathbb{R}^3 (or \mathbb{R}^n , for that matter), then the identity mapping on \mathbb{R}^3 is locally Hölder continuous as a map from $(\mathbb{R}^3, D_\omega(x, y))$ to $(\mathbb{R}^3, |x - y|)$ of some positive order.*

Because of Definition 1.5, Proposition 4.22 reduces itself to the following.

Lemma 4.23. *If Ω is a continuous weight on \mathbb{R}^n which is doubling, then for each compact set $K \subseteq \mathbb{R}^n$ there exist positive constants C and N such that $\Omega(B(x, r)) \geq C^{-1}r^N$ for all $x \in K$ and $0 < r < 1$. Here C depends on K but N does not.*

Indeed, the doubling property implies that there is a constant C_0 such that $\Omega(B(y, t)) \geq C_0^{-1}\Omega(B(y, 2t))$ for all $y \in \mathbb{R}^n$ and $t > 0$. Let K , x , and r be given as in the lemma. If j is a positive integer which is sufficiently large so that $B(x, 2^j r) \supseteq K$, then $\Omega(B(x, r)) \geq C_0^{-j}\Omega(B(x, 2^j r)) \geq C_0^{-j}\Omega(K)$. The lemma follows from this estimate, since we can choose j to be $-\log_2 r + C$.

Proposition 4.22 can be improved in various ways. A similar argument can be used to show that the identity is locally Hölder continuous as a map in the other direction, i.e., from $(\mathbb{R}^3, |x - y|)$ to $(\mathbb{R}^3, D_\omega(x, y))$. The identity mapping is in truth quasisymmetric, essentially by definition of D_ω and the strong- A_∞ condition. This implies a scale-invariant version of Proposition 4.22 in particular. Estimates like the reverse Hölder inequality (1.8) can be used to show that the identity mapping as a map from $(\mathbb{R}^3, D_\omega(x, y))$ to $(\mathbb{R}^3, |x - y|)$ (or the other way around) satisfies local Sobolev space estimates. (These Sobolev spaces have to be formulated carefully, since we are dealing with metric spaces.)

One can think of part c) of Theorem 4.20 as saying that $(\mathbb{R}^3, D_\omega(x, y))$ can be as far as possible from being bilipschitz equivalent to $(\mathbb{R}^3, |x - y|)$ even though ω is a strong A_∞ weight. Part b) is a complement to this. It says that we can make the singularity of the weight ω as small as we want while still having $(\mathbb{R}^3, D_\omega(x, y))$ and $(\mathbb{R}^3, |x - y|)$ be bilipschitz inequivalent. In terms of the metrics this means that the identity mapping as a map from $(\mathbb{R}^3, D_\omega(x, y))$ to $(\mathbb{R}^3, |x - y|)$ is locally Hölder continuous of order δ where $\delta \rightarrow 1$ as $s \rightarrow 0$. On the other hand, the identity is trivially Lipschitz as a map from $(\mathbb{R}^3, |x - y|)$ to $(\mathbb{R}^3, D_\omega(x, y))$, since $\omega \leq 1$ by definition. As before, there are various improvements of the statement about Hölder continuity, in terms of Sobolev space estimates, for instance. The bottom line is that we can choose ω so that it is a strong A_∞ weight and so that $(\mathbb{R}^3, D_\omega(x, y))$ is as close as we want to being bilipschitz equivalent to $(\mathbb{R}^3, |x - y|)$ without actually being bilipschitz equivalent.

Note however that the strong- A_∞ constant of ω blows up as s gets small; thus part b) does not really provide counterexamples to suitable small-constant versions of Question 1.9 (as mentioned briefly in Section 2, after the discussion of Problem 2.7).

REMARK 4.24. If one took a non-wild version of the Cantor set A from the preceding section (e.g., by replacing (3.3) by the requirement that the $\gamma_0(p)$'s be unlinked), and defined ω as in (4.1), then $(\mathbb{R}^3, D_\omega(x, y))$ would be bilipschitz equivalent to $(\mathbb{R}^3, |x - y|)$, no matter how large s is. (There is also nothing special about dimension 3 here.) This is not very hard to prove. One could use the same construction as in the next section. (See Remark 5.28.)

To finish the proof of Theorem 1.10 it remains to establish the last part, about connecting an arbitrary pair of points by a set which is bilipschitz equivalent to a standard Euclidean ball.

Proposition 4.25. *Let A be as in Section 3 (with some choice of k) and let ω be as in (4.1) (with some choice of s). There is a constant L_2 (depending on k and s) with the property that for every pair of distinct points $p, q \in \mathbb{R}^3$ there is a closed subset W of \mathbb{R}^3 containing p and q such that (W, D_ω) is L_2 -bilipschitz equivalent to a closed Euclidean 3-ball with the standard Euclidean metric.*

The rest of this section will be devoted to the proof of Proposition 4.25. The proof is basically trivial, but one should be a little careful. The basic idea is to connect p and q by a curve which stays away from A as much as possible, and which is as smooth as it can be subject to this constraint (and has no self-intersections), and then to take W to be a fattened-up version of this curve. On this set W the "strangeness" (deviation from Euclidean geometry) of the metric D_ω will be essentially like the strangeness of D_ω on a curve, and we shall be able to get rid of it easily.

Let $p, q \in \mathbb{R}^3$ be given, $p \neq q$. Given a pair of points y, z in \mathbb{R}^3 and an ε in $(0, 1)$, let $S(y, z)$ denote the segment which joins y to z , and let $S_\varepsilon(y, z)$ be the set of points x in \mathbb{R}^3 such that $\text{dist}(x, S(y, z)) \leq \varepsilon \text{dist}(x, \{y, z\})$. Thus $S_\varepsilon(y, z)$ is the union of two truncated cones, one with vertex y , the other with vertex z . It is also bilipschitz equivalent to a Euclidean 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 $S_\varepsilon(y, z)$ rather than a round ball. Typically W will look like a twisted version of $S_\varepsilon(y, z)$, with some spiralling at the ends. The segment $S(y, z)$ will correspond to the curve mentioned in the previous paragraph.

Let us begin with a preliminary fact.

Lemma 4.26. *Suppose that X is a nonempty subset of $\mathbb{R}^3 \setminus A$ which satisfies $\text{dist}(X, A) \geq \mu \min\{\text{diam } X, 1\}$ for some $\mu > 0$. Then there is a constant $C = C(\mu)$ such that*

$$(4.27) \quad \sup_X \omega \leq C \inf_X \omega$$

and

$$(4.28) \quad C^{-1} \omega(x)^{1/3} |x - y| \leq D_\omega(x, y) \leq C \omega(x)^{1/3} |x - y|,$$

for all $x, y \in X$.

The first part (4.27) follows immediately from the definition (4.1) of ω . To establish the second part let us observe that

$$(4.29) \quad C^{-1} \omega(z) |B| \leq \omega(B) \leq C \omega(z) |B|,$$

whenever B is a Euclidean ball which contains some $z \in X$ and which has diameter at most the diameter of X . Here $|B|$ denotes the Euclidean volume of B . These inequalities follow from our assumptions on X and the definition of ω . (It is helpful to distinguish between the cases where $\text{dist}(X, A) \leq 1$ and $\text{dist}(X, A) \geq 1$.) To prove (4.28) we use the fact that ω is a strong- A_∞ weight (Corollary 4.5 and Definition 1.5) to reduce to (4.29). This proves the lemma.

We shall use heavily the notation and definitions from Section 3 in the rest of the proof of Proposition 4.25.

Lemma 4.30. *The conclusion of Proposition 4.25 holds when there is a circle α in some A_l such that $p, q \in \tau(\alpha) \setminus \tau(N^2(\alpha))$. The same is true when $p, q \in \mathbb{R}^3 \setminus \tau(N(\Gamma_0))$.*

To prove this let us first observe that we can find a mapping g from a closed Euclidean ball into $X = \tau(\alpha) \setminus \tau(N^2(\alpha))$ or $X = \mathbb{R}^3 \setminus \tau(N(\Gamma_0))$ (as appropriate) such that the image of g contains p, q and g

is bilipschitz with respect to the Euclidean metrics (and not D_ω) with a uniformly bounded constant. In the first case, where $X = \tau(\alpha) \setminus \tau(N^2(\alpha))$, this is true because X is a nice smooth (connected) domain, a torus with a few tori removed. These domains can all be realized as images of each other under similarities, which makes transparent the existence of estimates which do not depend on α . The second case requires a tiny bit of extra care but is standard. (Think about $\mathbb{R}^3 \setminus B(0, 1)$ first.) Once we have such a mapping g we can precompose with a dilation if necessary to get a uniformly bilipschitz map from a Euclidean ball into (\mathbb{R}^3, D_ω) whose image contains p and q (because of Lemma 4.26). This proves Lemma 4.30.

To deal with the remaining cases of Proposition 4.25 we cannot simply “localize” in this manner, but instead we have to connect p and q with chains of sets which each satisfy separately the hypotheses of Lemma 4.26. The next lemma covers the most interesting case, and its proof will take a while.

Lemma 4.31. *If $p, q \in A$, then the conclusions of Proposition 4.25 hold.*

Choose δ in some A_m such that $p, q \in \tau(\delta)$ and m is as large as possible. Note that $m < \infty$. Let α_l, β_l be the (unique) circles in A_l such that $p \in \tau(\alpha_l), q \in \tau(\beta_l)$, respectively, so that $\alpha_{l+1} \in N(\alpha_l)$ for each l , etc. Choose (arbitrarily) points p_l and q_l in the boundaries of $\tau(\alpha_l)$ and $\tau(\beta_l)$ for each $l > m$. Of course $p_l \rightarrow p$ and $q_l \rightarrow q$ as $l \rightarrow \infty$. Our bilipschitz ball W will be obtained by combining a family of smooth tubes which connect the successive p_l 's and q_l 's.

We should record some bounds on distances and diameters. Let us write $X \approx Y$ when the two quantities X and Y are each bounded by a constant times the other, where the constant is allowed to depend on our parameters k and s but nothing else. Thus

$$(4.32) \quad \begin{aligned} |p_l - p_{l+1}| &\approx \text{diam } \alpha_l, \\ D_\omega(p_l, p_{l+1}) &\approx D_\omega\text{-diam } \alpha_l \approx (\text{diam } \alpha_l)^{1+s}, \end{aligned}$$

and similarly for the q_l 's. Here “diam” refers to the diameter with respect to the Euclidean metric, while “ D_ω -diam” refers to the diameter with respect to D_ω . Remember that

$$(4.33) \quad \text{diam } \gamma = \rho(k)^l \text{diam } \Gamma_0 \quad \text{when } \gamma \in A_l,$$

because of the construction in Section 3. This implies that

$$(4.34) \quad D_\omega(p_l, p_{l+1}) \approx D_\omega(p_{l+1}, p_{l+2}),$$

for all l , and similarly for the q_l 's, and also

$$(4.35) \quad D_\omega(p_{m+1}, p_{m+2}) \approx D_\omega(p_{m+1}, q_{m+1}) \approx D_\omega(q_{m+1}, q_{m+2}).$$

Notice that $\alpha_l \neq \beta_l$ when $l > m$, because of the maximality of m . This implies that

$$(4.36) \quad |p_{m+1} - q_{m+1}| \geq C^{-1} \text{diam } \delta$$

for some constant C .

Fix any line J in \mathbb{R}^3 , and choose points z_l and w_l in J for $l > m$ in the following manner. These points are supposed to correspond to the p_l 's and q_l 's, and this will be made more precise soon. Choose z_{m+1} and w_{m+1} first, in such a way that $|z_{m+1} - w_{m+1}| = D_\omega(p_{m+1}, q_{m+1})$. Except for this constraint the specific choices do not matter. When $l > m + 1$ let z_l and w_l be the (unique) points such that $|z_l - z_{l-1}| = D_\omega(p_l, p_{l-1})$, $|w_l - w_{l-1}| = D_\omega(q_l, q_{l-1})$ for all $l > m + 1$ and such that the z_l 's and w_l 's are ordered correctly. This means that z_l is always on the opposite side of z_{l-1} from z_{l-2} , and similarly for the w_l 's, and that z_{m+2} lies on the opposite side of z_{m+1} from w_{m+1} , and that w_{m+2} lies on the opposite side of w_{m+1} from z_{m+1} . Let $z \in J$ be the limit of the z_l 's, and let $w \in J$ be the limit of the w_l 's. These points will correspond to our original p and q . Note that all the z_l 's and w_l 's lie on $S(z, w)$, because of our ordering, and that the sequences $\{|z_l - z_{l-1}|\}_l$ and $\{|w_l - w_{l-1}|\}_l$ are approximately geometric sequences, because of (4.32) and (4.33).

In order to prove Lemma 4.31 it suffices to find $\varepsilon > 0$ and a bilipschitz mapping f from $(S_\varepsilon(z, w), |x - y|)$ into (\mathbb{R}^3, D_ω) (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 little tubes, 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 A . 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. Before we do all these things let us define f initially on the z_l 's and w_l 's in the obvious way.

Sublemma 4.37. *Let \mathcal{E} be the set consisting of z, w , and the z_l 's and w_l 's for $l > m$, and define $f : \mathcal{E} \rightarrow \mathbb{R}^3$ by $f(z) = p$, $f(w) = q$, $f(z_l) = p_l$, and $f(w_l) = q_l$. Then f is bilipschitz as a map from $(\mathcal{E}, |x - y|)$ into (\mathbb{R}^3, D_ω) with uniformly bounded bilipschitz constant.*

We defined f so that it satisfies the bilipschitz condition for consecutive points in \mathcal{E} . In order to check the bilipschitz condition for pairs of points which are further apart it is helpful to make some additional observations. If $j > l + 1$, then $\tau(\alpha_j) \subseteq \tau(N^2(\alpha_l))$, because of Lemma 3.10. This implies that $\text{dist}(\tau(\alpha_l) \setminus \tau(N(\alpha_l)), \tau(\alpha_j)) \approx \text{diam } \alpha_l$, and similarly for the β 's. The constants implicit in this statement do not depend on l or j , as one can verify most easily using the self-similarity of the construction in Section 3. We also get that $D_\omega\text{-dist}(\tau(\alpha_l) \setminus \tau(N(\alpha_l)), \tau(\alpha_j)) \approx D_\omega\text{-diam } \alpha_l$, and similarly for the β 's, with a uniform choice of the constant. This is a variant of (4.28) which can also be derived from (4.29). (Note that $X = \tau(\alpha_l) \setminus \tau(N(\alpha_l))$ satisfies the hypotheses of Lemma 4.26.) Thus we can control the interaction between the various α 's, and between the β 's. We also have that the α 's and the β 's do not interact with each other. Specifically, $\tau(\alpha_j) \subseteq \tau(\alpha_{m+1})$ when $j > m$, because of Lemma 3.10, and similarly for the β 's. The maximality of m implies that $\alpha_{m+1} \neq \beta_{m+1}$, and so Lemma 3.14 yields $\text{dist}(\tau(\alpha_{m+1}), \tau(\beta_{m+1})) \approx \text{diam } \delta$. We can convert this into $D_\omega\text{-dist}(\tau(\alpha_{m+1}), \tau(\beta_{m+1})) \approx D_\omega\text{-diam } \delta$ using the definition of ω and D_ω . In other words we can control the interaction between all the α 's and all the β 's. It is easy to verify Sublemma 4.37 using these estimates.

Next we want to define f on $S(z, w)$. Let us first record a simple observation about curves which will provide the building blocks for this extension of f .

Sublemma 4.38. *Given any (marked) circle γ in \mathbb{R}^3 and any pair of points a, b in different components of the boundary of $\tau(\gamma) \setminus \tau(N(\gamma))$, we can find an arc σ in the closure of $\tau(\gamma) \setminus \tau(N(\gamma))$ which connects a to b and has the following properties: if u and v are two points on σ , then the length of the arc in σ which connects u to v is bounded by $C|u - v|$; inside $B(a, C^{-1}\text{diam } \gamma)$ the curve σ agrees with the line segment emanating from a which is orthogonal to the boundary of $\tau(\gamma) \setminus \tau(N(\gamma))$ at a and goes inside $\tau(\gamma) \setminus \tau(N(\gamma))$, and similarly for b ; if $u \in \sigma$, then $\text{dist}(u, \mathbb{R}^3 \setminus \{\tau(\gamma) \setminus \tau(N(\gamma))\}) \geq C^{-1}\text{dist}(u, \{a, b\})$ (so that*

σ 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 σ is bounded by $C(i)(\text{diam } \gamma)^{1-i}$. (This is the “scale-invariant” estimate on the higher derivatives.) Here the constants C and $C(i)$ depend only on the parameter k from Section 3.

This is an easy exercise. Note that $\tau(\gamma) \setminus \tau(N(\gamma))$ is connected, and that we can reduce to the case where $\gamma = \Gamma_0$ by using a similarity.

Sublemma 4.39. *There is a map $f : S(z, w) \rightarrow \mathbb{R}^3$ which is smooth away from the endpoints and satisfies the following properties: f is defined on \mathcal{E} as in Sublemma 4.37; f is bilipschitz as a map from $(S(z, w), |x - y|)$ into (\mathbb{R}^3, D_ω) with uniformly bounded bilipschitz constant; $f(S(z_{m+1}, w_{m+1}))$ is contained in the closure of $\tau(\delta) \setminus \tau(N(\delta))$, $f(S(z_l, z_{l+1}))$ is contained in the closure of $\tau(\alpha_l) \setminus \tau(N(\alpha_l))$ when $l > m$, and $f(S(w_l, w_{l+1}))$ is contained in the closure of $\tau(\beta_l) \setminus \tau(N(\beta_l))$ when $l > m$. In particular,*

$$(4.40) \quad \text{dist}(f(t), A) \approx \text{diam } \alpha_l, \quad \text{when } t \in S(z_l, z_{l+1}), l > m,$$

and similarly for $S(z_{m+1}, w_{m+1})$ and the $S(w_l, w_{l+1})$'s, and

$$(4.41) \quad \text{dist}(f(t), A)^{1+s} \approx \text{dist}(t, \{z, w\}), \quad \text{for all } t \in S(z, w).$$

Moreover, if $f^{(i)}$ denotes the i^{th} order derivative of f on $S(z, w)$, then

$$(4.42) \quad |f^{(i)}(t)| \leq C(i) \text{dist}(f(t), A) \text{dist}(t, \{z, w\})^{-i},$$

for all $t \in S(z, w) \setminus \{z, w\}$ and $i \geq 1$, where $C(i)$ depends on i and the parameters k and s but nothing else and $|\cdot|$ denotes the ordinary Euclidean norm.

Before explaining how to build f -which comes down to connecting arcs as in Sublemma 4.38 and parameterizing them at the correct “speed”- let us consider the slightly odd-looking bound for $|f^{(i)}|$ in (4.42). The first point to notice is that we could write (4.42) in many different ways using (4.41). When $i = 1$, for instance, (4.42) reduces to saying that the first derivative of f at t is bounded by $C \text{dist}(f(t), A)^{-s}$, which is compatible with the bilipschitz condition on f . (In fact the bilipschitzness requires that

$$(4.43) \quad |f'(t)| \approx \text{dist}(f(t), A)^{-s} \approx \text{dist}(f(t), A) \text{dist}(t, \{z, w\})^{-1}$$

and this estimate will also be clear from the proof of Sublemma 4.39.) The bounds on the higher derivatives of f in (4.42) basically mean that f is chosen to be as smooth as it can be on each of the segments $S(z_{m+1}, w_{m+1})$, $S(z_l, z_{l+1})$, and $S(w_l, w_{l+1})$, subject to the speed limit (4.43) and the fact that f will probably have to do some nontrivial turning on each of these segments.

To prove Sublemma 4.39 we begin by defining f on \mathcal{E} as in Sublemma 4.37. To define f on the segments $S(z_{m+1}, w_{m+1})$, $S(z_l, z_{l+1})$, and $S(w_l, w_{l+1})$ we connect the p_l 's and q_l 's by arcs as in Sublemma 4.38 and we take f to be certain parameterizations of these arcs. We cannot use arclength parameterizations, but instead we parameterize these arcs at roughly constant speed. "Roughly constant" means that the maximum speed is bounded by a constant times the minimal speed. On an $S(z_l, z_{l+1})$, for instance, this approximate speed is comparable to $|p_l - p_{l+1}|/|z_l - z_{l+1}|$, and this is in turn comparable to $(\text{diam } \alpha_l)^{-s}$, because of our choices of the p_l 's and z_l 's. Note that this average speed on $S(z_l, z_{l+1})$ is comparable to that of the adjacent intervals, as is also the case for $S(z_{m+1}, w_{m+1})$ and the $S(w_l, w_{l+1})$'s. We cannot use parameterizations which have truly constant speed, since that would lead to discontinuities of the derivative of f at the z_l 's and w_l 's. Instead we require that the speeds be roughly constant on these intervals while also making a gentle transition from one interval to the next. If we take some care to make the "gentle transitions" of the parameterizations of the adjacent arcs approximately as gentle as they can be, then the estimates for the higher derivatives of f in (4.42) will follow from the corresponding estimates in Sublemma 4.38 and the normalized behavior of the arcs in Sublemma 4.38 at their endpoints. It is not hard to see that the mapping f that we produce in this way is bilipschitz on the union of any two adjacent intervals among $S(z_{m+1}, w_{m+1})$, the $S(z_l, z_{l+1})$'s, and the $S(w_l, w_{l+1})$'s. This uses the properties of the curves in Sublemma 4.38 (especially the chord-arc property), the fact that D_Ω is approximately a constant multiple of the Euclidean metric on regions like $\tau(\alpha_l) \setminus \tau(N^2(\alpha_l))$ (by Lemma 4.26), and the fact that we chose the z_l 's and the w_l 's so that the constant multiples work out correctly. The bilipschitzness on all of $S(z, w)$ follows from an argument like the one used to prove Sublemma 4.37, using also the fact that $f(S(z_l, z_{l+1}))$ is contained in the closure of $\tau(\alpha_l) \setminus \tau(N(\alpha_l))$ when $l > m$, etc. The estimate (4.40) is an immediate consequence of the definition of f (and the constructions in Section 3), and (4.41) follows

from (4.40) and our choices of the z_l 's and w_l 's. This proves Sublemma 4.39.

From now on we assume that f is defined on $S(z, w)$ as in Sublemma 4.39. We want to extend f to some $S_\varepsilon(z, w)$. In the following we set $v_0 = (w - z)/|w - z|$ and we let $f'(t)$ denote the derivative of f in the direction v_0 . This is defined for all $t \in S(z, w) \setminus \{z, w\}$, and it is a vector in \mathbb{R}^3 .

Sublemma 4.44. *There is a smooth map F from $S(z, w) \setminus \{z, w\}$ into linear mappings on \mathbb{R}^3 such that each $F(t)$ is an orientation-preserving similarity (a combination of a rotation and a dilation) which satisfies $F(t)v_0 = f'(t)$ and*

$$(4.45) \quad |F^{(i)}(t)| \leq C(i) \operatorname{dist}(f(t), A) \operatorname{dist}(t, \{z, w\})^{-i-1},$$

for all $t \in S(z, w) \setminus \{z, w\}$ and $i \geq 0$, where $C(i)$ depends on i and the parameters k and s but nothing else. (In other words, F satisfies the same bounds on its derivatives as f' does.)

Notice first that the dilation factor for $F(t)$ must simply be $|f'(t)|$, which is in turn controlled by (4.43). In particular $|f'(t)|$ never vanishes.

Sublemma 4.44 is basically trivial but we should be a little careful. Let F' denote the derivative of F in the direction of v_0 . Instead of trying to choose F directly let us write $F'(t) = \Phi(t)F(t)$ and choose $\Phi(t)$ instead. We should choose Φ so that $f''(t) = \Phi(t)f'(t)$ and $\Phi(t) = (\log |f'(t)|)' I + \phi(t)$, where $\phi(t)$ is antisymmetric (with respect to the usual Euclidean inner product). If we set $v(t) = |f'(t)|^{-1}f'(t)$, then we can reformulate the constraint that $f''(t) = \Phi(t)f'(t)$ in terms of ϕ as $v'(t) = \phi(t)v(t)$. We can produce such a ϕ algebraically; we take $\phi(t)$ to be the of the rank 1 map which sends $v(t)$ to $v'(t)$ minus its transpose. This choice of $\phi(t)$ satisfies $v'(t) = \phi(t)v(t)$ because $v'(t)$ is orthogonal to $v(t)$ (since $|v(t)| = 1$ for all t).

We are now ready to define F by solving the differential equation. Choose $F(z_{m+1})$ to be any orientation-preserving similarity on \mathbb{R}^3 which satisfies $F(z_{m+1})v_0 = f'(z_{m+1})$, and extend F to all of $S(z, w) \setminus \{z, w\}$ by solving $F'(t) = \Phi(t)F(t)$. Our selection of $\Phi(t)$ ensures that each $F(t)$ is an orientation-preserving similarity on \mathbb{R}^3 which satisfies $F(t)v_0 = f'(t)$ for all t . It is easy to see that the derivatives of F satisfy the same estimates as the derivatives of f' do, because $\Phi(t)$ was chosen so that it satisfies the same estimates as $(\log |f'|)'$ does, i.e.,

$|\Phi^{(i)}(t)| \leq C(i) \operatorname{dist}(t, \{z, w\})^{-i-1}$ for all $t \in S(z, w) \setminus \{z, w\}$ and $i \geq 0$, and because F itself has the same size as f' . This proves Sublemma 4.44.

Let us now use F to extend f in the directions orthogonal to J . Let π denote the orthogonal projection of \mathbb{R}^3 onto J . Given $x \in \mathbb{R}^3$ such that $\pi(x) \in S(z, w) \setminus \{z, w\}$, set

$$(4.46) \quad f(x) = f(\pi(x)) + F(\pi(x))(x - \pi(x)).$$

Before we analyze this extension, let us record some simple observations about $S_\varepsilon(z, w)$. If $x \in S_\varepsilon(z, w)$, then

$$(4.47) \quad |x - \pi(x)| \leq \varepsilon \operatorname{dist}(x, \{z, w\}),$$

$$(4.48) \quad \begin{aligned} \operatorname{dist}(\pi(x), \{z, w\}) &\leq \operatorname{dist}(x, \{z, w\}) \\ &\leq 2 \operatorname{dist}(\pi(x), \{z, w\}), \end{aligned}$$

and

$$(4.49) \quad \pi(x) \in S(z, w)$$

when $0 < \varepsilon < 1/2$. These are all easy consequences of the definition of $S_\varepsilon(z, w)$. Actually, (4.49) holds as soon as $\varepsilon < 1$, and it implies that $S_\varepsilon(z, w)$ is contained in the domain of our extension of f .

Sublemma 4.50. *We can choose $\varepsilon > 0$ small enough so that f defines a bilipschitz map from $(S_\varepsilon(z, w), |x - y|)$ into (\mathbb{R}^3, D_ω) , with ε and the bilipschitz constant depending only on the parameters k and s .*

Using the definition of f and (4.45) (with $i = 0$) we get that

$$(4.51) \quad \begin{aligned} |f(x) - f(\pi(x))| &\leq C |x - \pi(x)| \operatorname{dist}(f(\pi(x)), A) \\ &\quad \cdot \operatorname{dist}(\pi(x), \{z, w\})^{-1}. \end{aligned}$$

If $x \in S_\varepsilon(z, w)$ and $\varepsilon < 1/2$ then

$$(4.52) \quad |f(x) - f(\pi(x))| \leq C \varepsilon \operatorname{dist}(f(\pi(x)), A),$$

because of (4.47) and (4.48). If ε is small enough then we get that

$$(4.53) \quad \frac{1}{2} \operatorname{dist}(f(\pi(x)), A) \leq \operatorname{dist}(f(x), A) \leq 2 \operatorname{dist}(f(\pi(x)), A).$$

Thus $f(x)$ will not stupidly fall into A too soon. From now on we assume that $\varepsilon < 1/2$ and that ε is small enough so that (4.53) holds.

Next we observe that

$$(4.54) \quad |\nabla^i f(x)| \leq C(i) \operatorname{dist}(f(x), A) \operatorname{dist}(x, \{z, w\})^{-i},$$

for all $x \in S_\varepsilon(z, w) \setminus \{z, w\}$ and $i \geq 1$, where $C(i)$ depends on i and the parameters k and s but nothing else. This is easy to check, using (4.42) and (4.45) to control the derivatives of this extension (4.46) of f , and then (4.48), and (4.53) to get the estimates in the form of (4.54) (i.e., to replace $\pi(x)$ with x when necessary). Let us check that

$$(4.55) \quad |\nabla f(x)| \leq C(i) \operatorname{dist}(f(x), A)^{-s},$$

for all $x \in S_\varepsilon(z, w) \setminus \{z, w\}$. Using (4.41), (4.53), and (4.48) we get that

$$(4.56) \quad \begin{aligned} \operatorname{dist}(f(x), A)^{1+s} &\approx \operatorname{dist}(f(\pi(x)), A)^{1+s} \\ &\approx \operatorname{dist}(\pi(x), \{z, w\}) \\ &\approx \operatorname{dist}(x, \{z, w\}). \end{aligned}$$

This and (4.54) imply (4.55). From (4.55) we obtain that f is Lipschitz as a map from $(S_\varepsilon(z, w), |x - y|)$ into (\mathbb{R}^3, D_ω) (modulo some small additional attention at the points z and w).

To show that f is bilipschitz let us consider first a special case. Let $\eta > 0$ be small, to be chosen soon, and let $t \in S(z, w)$ be given. Set $r = \operatorname{dist}(t, \{z, w\})$, and consider the ball $B = B(t) = B(t, \eta r)$. Let us show that if η is small enough, then the restriction of f to B is bilipschitz (as a map into (\mathbb{R}^3, D_ω)), with a uniformly bounded bilipschitz constant. Notice first that

$$(4.57.a) \quad \operatorname{dist}(x, \{z, w\}) \approx \operatorname{dist}(t, \{z, w\})$$

and

$$(4.57.b) \quad \operatorname{dist}(f(x), A) \approx \operatorname{dist}(f(t), A), \quad \text{when } x \in B,$$

by the triangle inequality, (4.51), (4.41), and the requirement that η be small. Also,

$$(4.58) \quad \operatorname{diam} f(B) \leq C \operatorname{dist}(f(t), A)$$

by (4.54) (with $i = 1$) and (4.57). Thus the hypotheses of Lemma 4.26 are satisfied with $X = f(B)$, and we conclude that

$$(4.59) \quad D_\omega(\xi, \zeta) \approx \text{dist}(f(t), A)^s |\xi - \zeta|, \quad \text{when } \xi, \zeta \in f(B).$$

To prove that f is bilipschitz on B as a map into (\mathbb{R}^3, D_ω) we should show that

$$(4.60) \quad \text{dist}(f(t), A)^s |f(x) - f(y)| \approx |x - y|, \quad \text{when } x, y \in B.$$

To do this we shall approximate f on B using Taylor's theorem.

Define $\psi(x)$ by $\psi(x) = f(t) + F(t)(x - t)$. This is the linear Taylor approximation to f at t . (See (4.46) and Sublemma 4.44.) Because $F(t)$ is a similarity with dilation factor $|f'(t)|$ we have that

$$(4.61) \quad |\psi(x) - \psi(y)| = |f'(t)| |x - y| \approx \text{dist}(f(t), A)^{-s} |x - y|,$$

because of (4.43). We can control $f - \psi$ using (4.54) and Taylor's theorem. In fact we are really interested in $\nabla(f - \psi)$, and we have that

$$(4.62) \quad \begin{aligned} \sup_B |\nabla f - \nabla \psi| &\leq C (\sup_B |\nabla^2 f|) \text{diam } B \\ &\leq C \text{dist}(f(t), A) \text{dist}(t, \{z, w\})^{-2} \eta r, \end{aligned}$$

because of (4.54), (4.57), and the definition of B . We can simplify this further using (4.41) and the definition of r to get

$$(4.63) \quad \sup_B |\nabla f - \nabla \psi| \leq C \eta \text{dist}(f(t), A)^{-s}.$$

This means that

$$(4.64) \quad |(f - \psi)(x) - (f - \psi)(y)| \leq C \eta \text{dist}(f(t), A)^{-s} |x - y|$$

when $x, y \in B$. If η is small enough then we get (4.60) from (4.61) and (4.64). Choose η small enough so that this is true (*i.e.*, the bilipschitzness of f on the ball $B = B(t)$), and let it be fixed from now on.

Suppose now that we are given any $x, y \in S_\varepsilon(z, w)$, and let us estimate $D_\omega(f(x), f(y))$ from below. If x and y both lie in a ball of the form $B(t)$ as above then we already have the estimate that we need from (4.60), and so we may assume that this is not the case. Note that $u \in B(\pi(u))$ for all $u \in S_\varepsilon(z, w)$ if ε is small enough (compared to η). If

$|x - y|$ is small compared to $\eta \operatorname{dist}(x, \{z, w\})$ then $y \in B(\pi(x))$ (if also ε is small enough), and we are assuming that this is not true. We are making the same assumption with the roles of x and y reversed, and so we obtain that

$$(4.65) \quad |x - y| \geq C^{-1} \eta \max\{\operatorname{dist}(x, \{z, w\}), \operatorname{dist}(y, \{z, w\})\},$$

for some constant C . (The η in (4.65) does not matter, since it is fixed now anyway.)

We need to estimate $D_\omega(f(x), f(y))$ from below, and we shall do this in terms of $D_\omega(f(\pi(x)), f(\pi(y)))$. We already know that f is bilipschitz on $S(z, w)$, and so

$$(4.66) \quad D_\omega(f(\pi(x)), f(\pi(y))) \geq C^{-1} |\pi(x) - \pi(y)|.$$

This implies that

$$(4.67) \quad D_\omega(f(\pi(x)), f(\pi(y))) \geq C^{-1} |x - y|$$

when ε is small enough, because (4.65) and (4.47) yield $|x - \pi(x)| + |y - \pi(y)| \leq C \eta^{-1} \varepsilon |x - y|$. To get from here to $D_\omega(f(x), f(y))$ we have to control some error terms. Since f is Lipschitz as a map from $(S_\varepsilon(z, w), |x - y|)$ into (\mathbb{R}^3, D_ω) we have that

$$(4.68) \quad D_\omega(f(x), f(\pi(x))) \leq C |x - \pi(x)| \leq C \varepsilon \operatorname{dist}(x, \{z, w\}).$$

The same estimate is true with x replaced by y . These estimates combined with (4.65) and (4.67) give the desired

$$(4.69) \quad D_\omega(f(x), f(y)) \geq C^{-1} |x - y|$$

when ε is small enough.

This completes the proof of Sublemma 4.50, and Lemma 4.31 follows. The remaining cases involve similar constructions of connecting curves and their tubular neighborhoods, and we shall treat them in less detail.

Lemma 4.70. *If $p, q \in \tau(\Gamma_0)$, then the conclusion of Proposition 4.25 holds.*

Again choose δ in some A_m so that $p, q \in \tau(\delta)$ and m is as large as possible. We may as well assume that one of p and q lies in $\tau(N^2(\delta))$,

since otherwise we can apply Lemma 4.30. This implies that $|p - q| \geq C^{-1} \text{diam } \delta$ for some constant C (which depends only on the parameter k from Section 3): if $|p - q|$ were small compared to $\text{diam } \delta$, then we could use the fact that one of p and q lies in $\tau(N^2(\delta))$ to conclude that $p, q \in \tau(\gamma)$ for some child $\gamma \in A_{m+1}$ 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 4.31. The difference is that now one or both of p and q may not lie in A , so that the sequences of α_l 's and β_l 's might stop in a finite number of steps. In fact, we could have that one of p or q lies in $\tau(\delta) \setminus \tau(N(\delta))$, so that there would be no α_l 's, or no β_l 's. Thus it may be necessary to make some adjustments to the construction at one or both of the "ends", but the estimates and underlying principles remain the same. (We use Sublemma 4.38 to find nice curves, we connect them as in Sublemma 4.39, we extend out to little neighborhoods of the curves as in (4.46), and we conclude as in Sublemma 4.50.) The details are left to the reader.

Lemmas 4.30, 4.31, and 4.70 cover all the possible locations of p and q except for $p \in \tau(N(\Gamma_0))$ and $q \in \mathbb{R}^3 \setminus \tau(\Gamma_0)$ (or the other way around). In this case we have that $|p - q|$ is bounded below by some fixed constant. This situation also lends itself to the same basic construction in Lemma 4.31. That is, we set $m = 0$, $\delta = \Gamma_0$, and we define α_l 's and β_l 's as before, except that these sequences will stop after finitely many steps if $p \notin A$. We can then connect p to the boundary of $\tau(\Gamma_0)$ by a sequence of smooth curves in the various $\tau(\alpha_l) \setminus \tau(N(\alpha_l))$'s. Since q now lies in $\mathbb{R}^3 \setminus \tau(\Gamma_0)$, we do not have to go through any contortions to connect it to the boundary of $\tau(\Gamma_0)$ in a nice way, we can simply do it. We can then combine these two curves and fatten them up as before. More precisely, this means that we can build a map f from a set S_ε like $S_\varepsilon(z, w)$ in the proof of Lemma 4.31 into \mathbb{R}^3 such that the restriction of f to one end of S_ε provides a connection from p to the boundary of $\tau(\Gamma_0)$ and the restriction of f to the rest of S_ε provides a connection from there to q . Note that the proportion of S_ε which corresponds to p will be much smaller than half of S_ε when $|p - q|$ is very large, in which case the part of S_ε that goes from the boundary of $\tau(\Gamma_0)$ to q will have to have a big bulge in the middle. This does not cause a problem, but one should be careful to map the bulge away from A . The remaining details are much like those in the proof of Lemma 4.31, and we leave them to the reader.

This completes the proof of Proposition 4.25.

5. The proof of Theorem 1.3.

In the previous section we saw how to build strong A_∞ continuous weights on \mathbb{R}^3 such that the associated metric on \mathbb{R}^3 is not bilipschitz equivalent to the Euclidean metric. In this section we want to build subsets of \mathbb{R}^n which have many of the same nice properties as \mathbb{R}^3 (say) without being bilipschitz equivalent to \mathbb{R}^3 (with the Euclidean metric). One way to do this would be to show that we can embed $(\mathbb{R}^3, D_\omega(x, y))$ bilipschitzly into \mathbb{R}^n with the Euclidean metric. There is a general result in [Se4] which could be applied to give such a bilipschitz embedding for the weights constructed in the previous section. In this special case, however, such embeddings can be produced with much less fuss than the general construction in [Se4]. Alternatively, one can go to some trouble and build embeddings with especially nice properties. That is what we shall do here.

The construction that we shall make in this section will proceed along the following lines. The first step will be to build a set F which is analogous to the set A from Section 3 except that at each stage we use slightly smaller circles. These circles will not be linked, unlike their predecessors. Then we shall build a sequence of diffeomorphisms on \mathbb{R}^4 which send the various approximations to A to the corresponding approximations to F . This would be impossible in \mathbb{R}^3 , because of the linking properties, but none of the circles are linked in \mathbb{R}^4 , and it will be easy to build these mappings. In the limit we shall obtain a quasi-conformal map on \mathbb{R}^4 which sends A to F and maps \mathbb{R}^3 to a reasonably well-behaved surface. This surface will however be bilipschitz equivalent to $(\mathbb{R}^3, D_\omega(x, y))$ with ω as in (4.1) for a suitable choice of s , and so we shall be able to choose the parameters in such a way that it is not bilipschitz equivalent to \mathbb{R}^3 with the Euclidean metric.

Thus we need to begin by extending the construction in Section 3, and in the following we use the notation, assumptions, and results of Section 3 freely.

Let $\mu \in (0, 1)$ be fixed but arbitrary. It will correspond eventually to the parameter s in (4.1). The parameter k from Section 3 should also be treated as fixed and will play the same role as before. However, for this section it will be convenient to assume that k is a little larger, and so we require that $k \geq 10^{10}$.

Our first task is to choose some circles $\beta_0(p)$, $p \in P_0$, which will be cousins to the γ_0 's in Section 3. Each $\beta_0(p)$ should be a circle in \mathbb{R}^3

centered at p with radius $\mu\rho(k)$, and we require also that

$$(5.1) \quad \text{dist}(\{p\} \cup \beta_0(p), \{q\} \cup \beta_0(q)) \geq (100k)^{-1} \quad \text{when } p \neq q.$$

Fix any circles with these properties. When μ is small we can simply take $\beta_0(p)$ to be the circle centered at p which is obtained from $\gamma_0(p)$ by the obvious dilation, but when μ is closer to 1 this will not work (because the circles could touch) and so the circles should be tilted a little (or a lot, for that matter). To simplify the discussion let us simply require that the $\beta_0(p)$'s be unlinked, which in fact they must be when μ is small enough.

In short, the $\beta_0(p)$'s are a bunch of little circles centered on Γ_0 and placed at regular intervals. Think of [Mo, Figure 18.1, p. 127], but with the circles being smaller and unlinked. If we wanted we could impose some symmetry conditions, but we shall not bother.

Next, fix markings ψ_p for the $\beta_0(p)$'s. It does not matter how the markings are selected, but they need to be fixed forever.

Take $M(\Gamma_0)$ to be the union of the $\beta_0(p)$'s, $p \in P_0$, but viewed as a union of marked circles. We can define $M(\Gamma)$ for any marked circle Γ in \mathbb{R}^3 just as in (3.5), so that $M(\Gamma)$ is again a union of marked circles, and these circles are labelled by P_0 in the obvious way. As in Section 3 we define $M(E)$ when E is a finite union of marked circles, so that M defines a mapping on the space \mathcal{C} of finite unions of marked circles in \mathbb{R}^3 with the same kinds of operational properties as N has (with respect to unions and the action of orientation-preserving similarities).

Define $F_l \in \mathcal{C}$, $l \geq 0$, in the same way that the A_l 's were before, but using M now instead of N . That is, we set $F_0 = \Gamma_0$, $F_1 = M(\Gamma_0)$, and $F_l = M^l(\Gamma_0)$, where M^l denotes the l^{th} power of M , viewed as a mapping on \mathcal{C} . Each F_l is a union of k^l marked circles of radius $(\mu\rho(k))^l$, and F_{l+1} is the union of the $M(\alpha)$'s, where α runs through the circles which make up F_l . The Hausdorff limit of the F_l 's will give a Cantor set which is not wild.

Most of the lemmas in Section 3 apply to F and M as well, and we summarize them in the following.

Scholium 5.2. a) For each j, l with $0 \leq j \leq l$ we have that $F_l = \cup_{\alpha} M^{l-j}(\alpha)$, where the union is taken over all the constituent (marked) circles α in F_j , and the constituent circles of these two collections are marked in the same way.

b) For any marked circle Γ in \mathbb{R}^3 we have that $\sup_{x \in M(\Gamma)} \text{dist}(x, \Gamma) \leq \mu\rho(k) \text{radius}(\Gamma) < 10^{-6} \text{radius}(\Gamma)$.

- c) $\tau(M_l(\Gamma)) \subseteq \tau(\Gamma)$ for all marked circles Γ and all $l \geq 0$.
- d) $\tau(F_j) \supseteq \tau(F_l)$ when $0 \leq j \leq l$.
- e) Let α and α' be two distinct circles among those which make up F_l , $l \geq 1$. Then $\tau_{50}(\alpha)$ and $\tau_{50}(\alpha')$ are disjoint.

This is proved in exactly the same way as Lemmas 3.7, 3.8, 3.10, 3.11, and 3.14. The point is that the $\beta_0(p)$'s have the same properties as the $\gamma_0(p)$'s, except that they are smaller and unlinked, and these changes do not matter for these statements. For instance, the combinatorics in a) are exactly the same as in Lemma 3.7, while b) and c) are no more than simple applications of the triangle inequality.

Let us now proceed to \mathbb{R}^4 . From now on we shall identify \mathbb{R}^3 with the $x_4 = 0$ hyperplane in \mathbb{R}^4 , so that all of our constructions (A , F , etc.) can be viewed as living also in \mathbb{R}^4 . Notice that every orientation-preserving similarity on \mathbb{R}^3 has a unique extension to an orientation-preserving similarity on \mathbb{R}^4 , and so we can view all such transformations as acting on \mathbb{R}^4 . In particular the similarities which provide markings for our circles will be viewed as acting on all of \mathbb{R}^4 .

Given a circle Γ in \mathbb{R}^4 and $a > 0$ set

$$(5.3) \quad T(\Gamma) = \{x \in \mathbb{R}^4 : \text{dist}(x, \Gamma) \leq 10^{-5} \text{ radius } \Gamma\},$$

$$(5.4) \quad T_a(\Gamma) = \{x \in \mathbb{R}^4 : \text{dist}(x, \Gamma) \leq a 10^{-5} \text{ radius } \Gamma\}.$$

These are the 4-dimensional versions of (3.9) and (3.13) in \mathbb{R}^4 , and they enjoy properties analogous to those for $\tau(\Gamma)$ and $\tau_a(\Gamma)$. If E is a finite union of circles, then we define $T(E)$ and $T_a(E)$ to be the union of the sets obtained by applying T or T_a to the constituent circles.

Lemma 5.5. a) If Γ is a marked circle in \mathbb{R}^3 , then $T(N_l(\Gamma)) \subseteq T(\Gamma)$ and $T(M_l(\Gamma)) \subseteq T(\Gamma)$ for all $l \geq 0$.

- b) $T(A_j) \supseteq T(A_l)$ and $T(F_j) \supseteq T(F_l)$ when $0 \leq j \leq l$.

c) Let α and α' be two distinct circles among those which make up A_l (or F_l), $l \geq 1$. Then $T_{50}(\alpha)$ and $T_{50}(\alpha')$ are disjoint.

This is proved in the same way as for the analogous results for τ . (Nothing special about \mathbb{R}^3 was used; it all came down to the triangle inequality.)

Our next task is to build a homeomorphism from \mathbb{R}^4 to itself which sends A to F and which is otherwise as nice as possible. This mapping will have to shrink distances with some severity near A , but away from A it will be a diffeomorphism. We shall produce this mapping through an iterative process, and we first need to construct some building blocks.

Lemma 5.6. *There is a smooth diffeomorphism $\Phi_0 : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ such that $\Phi_0(T(\Gamma_0)) = T(\Gamma_0)$, Φ_0 equals the identity on the complement of $T(\Gamma_0)$ and on a neighborhood of $\partial T(\Gamma_0)$, and the restriction of Φ_0 to a neighborhood of $T(\gamma_0(p))$ is, for each $p \in P_0$, the orientation-preserving similarity from $T(\gamma_0(p))$ onto $T(\beta_0(p))$ which is determined by their markings.*

For this lemma it is crucial that we are working in \mathbb{R}^4 instead of \mathbb{R}^3 . In \mathbb{R}^3 the $\gamma_0(p)$'s are linked, while the $\beta_0(p)$'s are not; in \mathbb{R}^4 , none of them are linked.

We shall obtain Φ_0 by composing a finite number of simpler pieces. It will be convenient to use also some auxiliary circles. For each $p \in P_0$ choose a circle $\delta_0(p)$ in \mathbb{R}^3 which is centered at p and which has radius $(10^4 k)^{-1}$. These circles should also be given markings. The specific choices do not matter.

Sublemma 5.7. *For each $p \in P_0$ there is a diffeomorphism $f : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ such that f equals the identity on $\mathbb{R}^4 \setminus T(\Gamma_0)$, on a neighborhood of $\partial T(\Gamma_0)$, and on neighborhoods of $T(\gamma_0(q))$ and $T(\delta_0(q))$ for each $q \in P_0 \setminus \{p\}$, and such that the restriction of f to a neighborhood of $T(\gamma_0(p))$ agrees with the orientation-preserving similarity which takes $\gamma_0(p)$ to $\delta_0(p)$ and which is determined by their markings. The analogous result holds for the β_0 's instead of the γ_0 's.*

Lemma 5.6 will follow once we have proved Sublemma 5.7. Indeed, if Sublemma 5.7 is true, then we can build a nice map on \mathbb{R}^4 which is the identity outside $T(\Gamma_0)$ and which sends $T(\gamma_0(p))$ to $T(\delta_0(p))$ for each $p \in P_0$ in the right way (in accordance with the markings), simply by composing the various pieces from Sublemma 5.7. We can then produce a similar map for the β_0 's instead of the γ_0 's, and Φ_0 is obtained by composing the previous map with the inverse of the second one.

Let us prove Sublemma 5.7. Let $p \in P_0$ be given. We want to pick up $T(\gamma_0(p))$ off the "floor" \mathbb{R}^3 , shrink it and turn it around as

necessary, and lay it down on $T(\delta_0(p))$ without disturbing any of the other $T(\gamma_0(q))$'s or $T(\delta_0(q))$'s or the complement of $T(\Gamma_0)$. This is very simple geometrically (or physically), and the following claim makes it precise.

Claim 5.8. *There is a smooth 1-parameter family of orientation-preserving similarities σ_t , $0 \leq t \leq 1$, on \mathbb{R}^4 such that σ_0 is the identity, σ_1 is the similarity which sends $\gamma_0(p)$ to $\delta_0(p)$ that is determined by the markings, and the compact set $\cup_{0 \leq t \leq 1} \sigma_t(T(\gamma_0(p)))$ lies in the interior of $T(\Gamma_0) \setminus \cup_{q \in P_0 \setminus \{p\}} (T(\gamma_0(q)) \cup T(\delta_0(q)))$.*

To see this, let H denote the $x_4 = 10^{-7}$ hyperplane. In the beginning (near $t = 0$) σ_t should simply make a translation in the x_4 direction by 10^{-7} , so that \mathbb{R}^3 is sent onto H . For the next period of time σ_t should preserve H while deforming the translate of $\gamma_0(p)$ in H into the translate of $\delta_0(p)$ in H . At the end σ_t should simply make a translation by -10^{-7} in the x_4 direction. If we do the deformation in H in the middle stage correctly then at the end we shall obtain the correct choice of σ_1 . Because we are requiring that $k \geq 10^{10}$ in this section (so that $\text{radius}(\beta_0(p)) \leq \text{radius}(\gamma_0(p)) \leq 10^{-9}$, by (3.1)), we obtain easily that $\cup_{0 \leq t \leq 1} \sigma_t(T(\gamma_0(p)))$ lies inside $T(\Gamma_0)$. It also remains disjoint from $\cup_{q \in P_0 \setminus \{p\}} (T(\gamma_0(q)) \cup T(\beta_0(q)))$. Indeed, in the first and final stages of the motion σ_t (when we are simply translating in the x_4 direction) this follows easily from (3.2) and (5.1), while in the middle stage we are moving around in H , and H is too far from $\mathbb{R}^3 = \{x_4 = 0\}$ to cause any problems, since $k \geq 10^{10}$. (Keep in mind that $\gamma_0(q)$'s, $\beta_0(q)$'s, and $\delta_0(q)$'s lie in \mathbb{R}^3 .) This proves the claim.

Next we want to extend the motion in Claim 5.8 to all of \mathbb{R}^4 with suitable properties. We use a standard argument for extending smooth isotopies, i.e., we convert the problem to one of extending vector fields.

Claim 5.9. *There is a smooth \mathbb{R}^4 -valued function $V(x, t)$ on $[0, 1] \times \mathbb{R}^4$ such that $V(x, t) = 0$ for all t when $x \in \mathbb{R}^4 \setminus T(\Gamma_0)$ or x lies in a neighborhood of $\partial T(\Gamma_0)$ or of $\cup_{q \in P_0 \setminus \{p\}} (T(\gamma_0(q)) \cup T(\delta_0(q)))$, and $V(\sigma_t(x), t) = (d\sigma_t/dt)(x)$ for all t when x lies in a neighborhood of $T(\gamma_0(p))$.*

This is an immediate consequence of Claim 5.8 and standard facts about smooth functions.

We are now almost finished with the proof of Sublemma 5.7. Let f_t be the flow associated to $V(x, t)$, so that f_0 is the identity and $(df_t/dt)(x) = V(f_t(x), t)$ for all x and t . Then $f_t(x) = x$ for all t when x lies in $\mathbb{R}^4 \setminus T(\Gamma_0)$ or a neighborhood of $\partial T(\Gamma_0)$ or a neighborhood of $\cup_{q \in P_0 \setminus \{p\}} (T(\gamma_0(q)) \cup T(\delta_0(q)))$, and $f_t(x) = \sigma_t(x)$ for all t when x lies in a neighborhood of $T(\gamma_0(p))$, by the uniqueness theorem for ordinary differential equations. Of course the f_t 's are diffeomorphisms also, and so $f = f_1$ has all the required properties. This proves Sublemma 5.7.

Now that Lemma 5.6 is established we want to build a more complicated homeomorphism by piecing together many copies of Φ_0 . This mapping will take A to F and we shall try to make it as nice as possible off of A . To do this we need some more notation and definitions.

Let \mathcal{A}_l and \mathcal{F}_l denote the collections of k^l marked circles which make up A_l and F_l , respectively. Let \mathcal{S}_l denote the Cartesian product of l copies of P_0 , so that elements of \mathcal{S}_l are finite sequences of length l with entries in P_0 . There are natural bijections from \mathcal{S}_l onto each of \mathcal{A}_l and \mathcal{F}_l which code the history of the circles. If $\alpha \in \mathcal{A}_l$ and α' is its parent in \mathcal{A}_{l-1} , so that $\alpha \in N(\alpha')$, then the sequence in \mathcal{S}_l associated to α is just the sequence in \mathcal{S}_{l-1} associated to α' together with one additional element of P_0 at the end to specify the position of α relative to α' . (See (3.5).) These codings of \mathcal{A}_l and \mathcal{F}_l induce a natural (bijective) correspondence $\theta_l : \mathcal{A}_l \rightarrow \mathcal{F}_l$.

Let \mathcal{S} denote the Cartesian product of countably many copies of P_0 , so that the elements of \mathcal{S} are sequences $\{q_j\}_{j \geq 1}$ which take values in P_0 . There are natural bijections from \mathcal{S} onto A and F which take a sequence of elements of P_0 and assign to it the point in A or F with that history. Alternatively, these points could be described as the limits of the circles in \mathcal{A}_l and \mathcal{F}_l corresponding to the initial l terms in the sequence in \mathcal{S} . These bijections induce a bijection $\theta : A \rightarrow F$. The homeomorphism on \mathbb{R}^4 that we are going to construct will be an extension of θ , and it will (in a certain sense) respect the codings provided by the θ_l 's on the complement of A .

Given a marked circle Γ in \mathbb{R}^3 , let us associate to it two compact sets $X(\Gamma)$ and $Y(\Gamma)$ by taking $X(\Gamma)$ to be the closure of $T(\Gamma) \setminus T(N(\Gamma))$ and $Y(\Gamma)$ to be the closure of $T(\Gamma) \setminus T(M(\Gamma))$. These sets are solid tori with k smaller solid tori removed. We are going to break up \mathbb{R}^4 into pieces using the $X(\Gamma)$'s and $Y(\Gamma)$'s and define our eventual homeomorphism initially on these various pieces before gluing them together. Notice that $\Phi_0(X(\Gamma_0)) = Y(\Gamma_0)$.

If β is a marked circle in \mathbb{R}^3 , let ϕ_β denote the similarity which provides the marking. (This will not be confused with ϕ_p , which provides the marking of $\gamma_0(p)$, because p is not a circle.) Define Φ_β to be the composition $\phi_\beta \circ \Phi_0 \circ (\phi_\beta)^{-1}$. In other words this is a copy of Φ_0 which lives near β instead of Γ_0 . Thus Φ_β has the following properties: it is a diffeomorphism on \mathbb{R}^4 which is the identity on the complement of $T(\beta)$ and on a neighborhood of $\partial T(\beta)$ and which takes $T(\beta)$ to itself; it sends $N(\beta)$ to $M(\beta)$, and it preserves the labellings of the circles in $N(\beta)$ to $M(\beta)$ by P_0 ; its restriction to a neighborhood of $T(\gamma)$ is an orientation-preserving similarity for each circle γ in $N(\beta)$, and this similarity is the one determined by the markings (and implicitly the labellings of the circles by P_0 also); and Φ_β sends $X(\beta)$ onto $Y(\beta)$. These properties are all easy consequences of the analogous statements for Φ_0 and Γ_0 .

Given $\alpha \in \mathcal{A}_l$ let ξ_α denote the orientation-preserving similarity which takes α to $\beta = \theta_l(\alpha) \in \mathcal{F}_l$ and which is the one determined by the markings. Set $\Psi_\alpha = \Phi_\beta \circ \xi_\alpha$. The Ψ_α 's will provide the building blocks for the homeomorphism that we want to build. Let us summarize some of their important properties in a lemma.

Lemma 5.10. *Suppose that $\alpha \in \mathcal{A}_l$ and $\beta = \theta_l(\alpha) \in \mathcal{F}_l$. Ψ_α has the following properties: it is a diffeomorphism on \mathbb{R}^4 which sends $T(\alpha)$ to $T(\beta)$; it agrees with ξ_α outside $T(\alpha)$ and on a neighborhood of $\partial T(\alpha)$; it maps $N(\alpha)$ to $M(\beta)$; if γ is one of the circles in $N(\alpha)$, then $\Psi_\alpha(\gamma)$ is the same as the circle $\delta = \theta_{l+1}(\gamma) \in \mathcal{F}_l$, and the restriction of Ψ_α to a neighborhood of $T(\gamma)$ agrees with the orientation-preserving similarity that takes γ to δ and is determined by the markings (i.e., ξ_γ); and $\Psi_\alpha(X(\alpha)) = Y(\beta)$.*

These properties are all easy to verify from the definitions and Lemma 5.6.

Let \mathcal{A} and \mathcal{F} denote the union of all the \mathcal{A}_l 's and \mathcal{F}_l 's for $l = 0, 1, \dots$. The homeomorphism H that we really want is defined as follows:

$$(5.11) \quad H = \begin{cases} \text{the identity,} & \text{on } \mathbb{R}^4 \setminus T(\Gamma_0), \\ \Psi_\alpha, & \text{on } X(\alpha) \text{ for } \alpha \in \mathcal{A}_l, \\ \theta, & \text{on } A. \end{cases}$$

We need to check that this is well defined, etc.

Lemma 5.12. \mathbb{R}^4 is the union of $\mathbb{R}^4 \setminus T(\Gamma_0)$, the sets $X(\alpha)$ for $\alpha \in \mathcal{A}$, and A . $\mathbb{R}^4 \setminus T(\Gamma_0)$ is disjoint from all the $X(\alpha)$'s except $X(\Gamma_0)$, and $X(\alpha), X(\alpha')$ intersect, $\alpha, \alpha' \in \mathcal{A}$, if and only if one of α and α' is the parent of the other, in which case $X(\alpha)$ and $X(\alpha')$ intersect only in a component of the boundary. Neither $\mathbb{R}^4 \setminus T(\Gamma_0)$ nor any of the $X(\alpha)$'s contain any elements of A . The analogous statements for $Y(\beta), \beta \in \mathcal{F}$, and F are also true.

This follows easily from Lemma 5.5, the definitions of A and F , etc. The main point is that if $\alpha, \alpha' \in \mathcal{A}$, then either one of α and α' is an ancestor of the other, say α' is an ancestor of α , in which case $T(\alpha) \subseteq T(\alpha')$, or $T(\alpha)$ and $T(\alpha')$ are disjoint (because they have distinct common ancestors in some \mathcal{A}_l).

Lemma 5.13. H is well defined and smooth off A .

For instance, we have taken H to be the identity on the complement of $T(\Gamma_0)$ and to be Ψ_{Γ_0} on $X(\Gamma_0)$, and the two share the torus $\partial T(\Gamma_0)$ as part of their boundaries. However, $\Psi_{\Gamma_0} = \Phi_0$ by definitions, and so we really have $H = \Phi_0$ on the union of $\mathbb{R}^4 \setminus T(\Gamma_0)$ and $X(\Gamma_0)$.

Similarly, if $\alpha \in \mathcal{A}, \alpha \neq \Gamma_0$, and $\alpha' \in \mathcal{A}$ is its parent, then $X(\alpha')$ and $X(\alpha)$ have a torus as their common boundary (namely, $\partial T(\alpha)$). However, on a neighborhood of $\partial T(\alpha)$ both Ψ_α and $\Psi_{\alpha'}$ agree with ξ_α , and so H is smooth across $\partial T(\alpha)$.

Lemma 5.14. For any α in any \mathcal{A}_l we have that $H(X(\alpha)) = Y(\theta_l(\alpha))$ and $H(T(\alpha)) = T(\theta_l(\alpha))$.

The first part about $X(\alpha)$ is an immediate consequence of the definition (5.11). For the second part we begin by observing that $T(\alpha)$ is the union of $X(\gamma)$ over all $\gamma \in \mathcal{A}$ descended from α (including α itself) together with $A \cap T(\alpha)$, and that the analogous statement holds for a $\beta \in \mathcal{F}$, but with $Y(\cdot)$ instead of $X(\cdot)$. Next we observe that $H(A \cap T(\alpha)) = F \cap T(\theta_l(\alpha))$. Indeed, Lemma 5.5 implies that $A \cap T(\alpha)$ consists precisely of the points in A which are “descended” from α , and similarly $F \cap T(\theta_l(\alpha))$ consists of the points in F which are descended from $\theta_l(\alpha)$, and the two correspond under H because H is defined to be the same as θ on A . The second part of the lemma follows from the first part and these observations.

Lemma 5.15. *H is continuous on all of \mathbb{R}^4 .*

The continuity of H away from A follows from Lemma 5.13, while the continuity on A follows from Lemma 5.14.

Lemma 5.16. *H is a homeomorphism on \mathbb{R}^4 , and it is a diffeomorphism from $\mathbb{R}^4 \setminus A$ onto $\mathbb{R}^4 \setminus F$.*

This follows from Lemmas 5.14 and 5.12, the fact that the individual Ψ_α 's are diffeomorphisms, etc.

Next let us estimate the differential of H , which we denote by dH , and which is well defined off A . More precisely, if $x \in \mathbb{R}^4 \setminus A$, then dH_x will be used to denote the differential of H at x as a linear transformation.

Lemma 5.17. *Choose $s > 0$ so that $\mu = \rho(k)^s$, set $t = s(1+s)^{-1}$, and set $\lambda(x) = \min\{1, \text{dist}(x, A)^s\}$ and $\nu(y) = \max\{1, \text{dist}(y, F)^{-t}\}$. Then there is a constant C so that*

$$C^{-1} \lambda(x) |v| \leq |dH_x(v)| \leq C \lambda(x) |v|$$

and

$$C^{-1} \nu(y) |v| \leq |dH_y^{-1}(v)| \leq C \nu(y) |v|,$$

for all $x \in \mathbb{R}^4 \setminus A$, $y \in \mathbb{R}^4 \setminus F$, and $v \in \mathbb{R}^4$, where $|v|$ denotes the Euclidean norm of the vector v .

To see this we need to first reexpress λ and ν in more useful forms.

Sublemma 5.18. *$\lambda(x) \approx 1$ on $\mathbb{R}^4 \setminus T(\Gamma_0)$, and $\lambda(x) \approx \mu^l$ on $X(\alpha)$ for any $\alpha \in \mathcal{A}_l$, $l \geq 0$. Here $a \approx b$ means that each of a, b is bounded by a constant times the other. Similarly, $\nu(y) \approx 1$ on $\mathbb{R}^4 \setminus T(\Gamma_0)$ and $\nu(y) \approx \mu^{-l}$ on $Y(\beta)$ for any $\beta \in \mathcal{F}_l$, $l \geq 0$.*

That $\lambda(x) \approx 1$ and $\nu(y) \approx 1$ on $\mathbb{R}^4 \setminus T(\Gamma_0)$ simply reflects the fact that A and F lie in the interior of $T(\Gamma_0)$ (which follows from $A \subseteq \tau(N(\Gamma_0))$ and $F \subseteq \tau(M(\Gamma_0))$, for instance). Using the definitions of s and t the remaining parts come down to

$$(5.19) \quad \begin{aligned} \text{dist}(x, A) &\approx \rho(k)^l, & \text{on } X(\alpha), \\ \text{dist}(y, F) &\approx (\mu \rho(k))^l, & \text{on } Y(\beta), \end{aligned}$$

for $\alpha \in \mathcal{A}_l$ and $\beta \in \mathcal{F}_l$. These estimates are easy to check, using $\emptyset \neq A \cap T(\alpha) \subseteq T(N^2(\alpha))$, $\text{radius}(\alpha) = \rho(k)^l$, $\emptyset \neq F \cap T(\beta) \subseteq T(M^2(\beta))$, and $\text{radius}(\beta) = (\mu \rho(k))^l$. (Do not forget Lemma 3.10 and Scholium 5.2.c.) This proves the sublemma.

As for Lemma 5.17, notice first that it is true on $\mathbb{R}^4 \setminus T(\Gamma_0)$, since H equals the identity there. Fix now an α in some \mathcal{A}_l , so that H agrees with Ψ_α on $X(\alpha)$ by the definition (5.11). Let us check that

$$(5.20) \quad |d(\Psi_\alpha)_x(v)| \approx \mu^l |v|,$$

for all x and v . Recall that, by definition, $\Psi_\alpha = \Phi_\beta \circ \xi_\alpha = \phi_\beta \circ \Phi_0 \circ (\phi_\beta)^{-1} \circ \xi_\alpha$, where $\beta = \theta_l(\alpha)$. Φ_0 is a single diffeomorphism which equals the identity outside $T(\Gamma_0)$, and so its differential distorts the Euclidean norm only by a bounded factor. Since ϕ_β is a similarity, it distorts distances by the same constant factor everywhere, and so the presence of ϕ_β and its inverse cancel each other out. Thus we are left with ξ_α , which is a similarity which maps α to β . This means that the dilation factor of ξ_α is simply the ratio of the diameters of β and α . By construction the diameter of the former is $(\mu \rho(k))^l$ while the diameter of the latter is $\rho(k)^l$, and so we get (5.20). The required estimates on dH and dH^{-1} follow easily. (Do not forget Lemma 5.14.)

Lemma 5.21. *H is Lipschitz continuous and continuously differentiable on all of \mathbb{R}^4 .*

We already know that H is smooth off A , but it is not hard to see that the differential of H exists and vanishes at points in A . This can be derived from Lemma 5.14, for instance, and the fact that the radius of $\theta_l(\alpha)$ is μ^l times the radius of α for all $\alpha \in \mathcal{A}_l$. Lemma 5.21 follows from Lemma 5.17 and the boundedness of λ .

REMARK 5.22. Although H is C^1 everywhere, it is not a C^1 diffeomorphism across A , because its differential vanishes there. However, there is a simple way to approximate H by C^1 diffeomorphisms. Define H_m by $H_m =$ the identity on $\mathbb{R}^4 \setminus T(\Gamma_0)$, $H_m = \Psi_\alpha$ on $X(\alpha)$ for $\alpha \in \mathcal{A}_l$, $l < m$, and $H_m = \xi_\alpha$ on $T(\alpha)$ when $\alpha \in \mathcal{A}_m$. This is approximately the same as the definition (5.11) of H , except at levels m and below, where it is flattened out. These mappings satisfy suitable versions of the preceding lemmas, and in particular the differentials of the H_m^{-1} 's satisfy the same sort of estimates as in Lemma 5.17, uniformly in m .

This observation can be useful in making it easier to derive properties of H from the bounds on dH^{-1} (e.g., when verifying the precise Sobolev space properties of H^{-1}).

Lemma 5.23. *H is quasiconformal.*

According to the usual definition this follows from Lemmas 5.16 and 5.17. However it is not hard to verify directly the a priori stronger condition that there is a $C > 0$ so that for all $x \in \mathbb{R}^4$ and $r > 0$ there is an $R > 0$ (which may depend on x and r) such that $B(H(x), R) \subseteq H(B(x, r)) \subseteq B(H(x), CR)$. To do this one considers separately the cases where $r \geq 1$, $r < 1$ and r is small compared to $\text{dist}(x, A)$, and $r < 1$ but r is not small compared to $\text{dist}(x, A)$. These cases can be treated by reducing to facts about the Ψ_α 's (which are clearly uniformly quasiconformal) and Lemma 5.14.

It is not too difficult to describe completely the manner in which H distorts distances, but this is slightly gory and best left as an exercise. We should at least make the connection with the preceding section by formulating Lemma 5.17 (and some consequences of it) in terms of the metrics associated to strong A_∞ weights.

Define $\Omega(x)$ on \mathbb{R}^4 by $\Omega(x) = \lambda(x)^4$, and define $\omega(x)$ on \mathbb{R}^3 by $\omega = \lambda(x)^3$. Thus ω is the same as in (4.1), with s chosen as in Lemma 5.17. These are both strong A_∞ weights on their respective domains, by Propositions 4.3 and 4.4. This uses also the simple fact that A is uniformly disconnected as subset of \mathbb{R}^4 , and not just \mathbb{R}^3 . (In fact uniform disconnectedness is an intrinsic property of a metric space.) Let $D_\Omega(\cdot, \cdot)$ and $D_\omega(\cdot, \cdot)$ be the associated metrics on \mathbb{R}^4 and \mathbb{R}^3 , respectively.

Lemma 5.24. $D_\omega(x; y) = D_\Omega(x, y)$ for all $x, y \in \mathbb{R}^3$.

This is not an accident, but it is sort of pleasant that we get an actual equality and not just an equivalence in size. Let $x, y \in \mathbb{R}^3$ be given, and let γ be a rectifiable curve in \mathbb{R}^4 which connects x to y . Let γ' denote the projection of γ to \mathbb{R}^3 . One can compute that the Ω -length of γ is at least as big as the ω -length of γ' , and that the two are equal if $\gamma \subseteq \mathbb{R}^3$ to begin with. (The first part comes down to the fact that $\text{dist}(z, A)$ is decreased by projecting z onto \mathbb{R}^3 , while the second is just a question of unwinding definitions.) This implies the lemma.

Lemma 5.25. *H is bilipschitz as a map from $(\mathbb{R}^4, D_\Omega(x, y))$ to $(\mathbb{R}^4, |x - y|)$, and as a map from $(\mathbb{R}^3, D_\omega(x, y))$ to $(E, |x - y|)$, where $E = H(\mathbb{R}^3)$.*

The first part is basically a reformulation of Lemma 5.17. Strictly speaking, the fact that H is Lipschitz as a map from $(\mathbb{R}^4, D_\Omega(x, y))$ to $(\mathbb{R}^4, |x - y|)$ is an immediate consequence of Lemmas 5.17 and 5.21, but one should be a little more careful about the Lipschitzness in the reverse direction. (Note however that Lemma 5.14 can be very useful in providing control near the singular points, so that one can concentrate on the smooth parts, which are more amenable to calculus. Alternatively one can approximate H by diffeomorphisms as in Remark 5.22 in order to reduce the problem to calculus.) The second part follows from the first and Lemma 5.24.

Note that Lemma 5.25 implies that for any pair of points $x, y \in E$ there is a closed subset of E containing x and y which is bilipschitz equivalent to a closed Euclidean 3-ball, with a uniformly bounded bilipschitz constant, because of the corresponding property for $(\mathbb{R}^3, D_\omega(x, y))$ (Proposition 4.25).

Lemma 5.26. *E is a regular set of dimension 3 (as in Definition 1.1).*

This can be derived from the second part of Lemma 5.25 and the corresponding general fact for strong A_∞ weights (1.7). Alternatively, one could go back to the definitions of E, H, \dots , and simply compute directly, using the fact that H is basically a uniform contraction by a known quantity on each $X(\alpha)$, and Lemma 5.14, etc.

At this stage we can read off Theorem 1.3 and some variants of it from Lemma 5.25, Proposition 4.25, and Theorem 4.20.

Theorem 5.27. *With the notation as above, $H : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ is a quasi-conformal mapping which is also Lipschitz continuous and $E = H(\mathbb{R}^3)$ is a 3-dimensional regular set when $k \geq 10^{10}$ and $\mu \in (0, 1)$. Every pair of points in E is contained in a closed subset of E which is bilipschitz equivalent to a closed Euclidean 3-ball, with a uniformly bounded bilipschitz constant (which depends on k and μ .) Define $s > 0$ by $\mu = \rho(k)^s$, and let U be a (relatively) open subset of E which intersects $F = H(A)$, and let V be any open subset of \mathbb{R}^3 . (For instance take $U = E, V = \mathbb{R}^3$.)*

a) *If $s > 2$ then there does not exist a homeomorphism $h : (U, |x -$*

$y|) \rightarrow (V, |x - y|)$ which is locally Lipschitz.

b) For any $s > 0$ there does not exist a homeomorphism $h : (U, |x - y|) \rightarrow (V, |x - y|)$ which is locally Lipschitz if k is large enough.

c) For any $s > 0$ there does not exist a homeomorphism $h : (U, |x - y|) \rightarrow (V, |x - y|)$ which is locally Hölder continuous of order greater than $3(1 + s)^{-1}$.

The Euclidean metric is made crudely explicit in a)-c) above to avoid confusion with the other metrics which have been used.

Notice that any value of $s > 0$ can occur by choosing $\mu \in (0, 1)$ correctly.

The remarks that follow Theorem 4.20 have counterparts in this setting too. Quasiconformal mappings and their inverses are always Hölder continuous, but part c) of Theorem 5.27 shows that any homeomorphism from E to \mathbb{R}^3 must have small Hölder exponent if s is large enough. Thus E can be chosen so that it is very far from being bilipschitz equivalent to \mathbb{R}^3 . On the other hand, we can take s to be as small as we like as long as we take k to be large enough, in such a way that we can make E be extremely close to being bilipschitz equivalent to \mathbb{R}^3 without actually being bilipschitz equivalent. For instance, H^{-1} will be locally Hölder continuous of order as close to 1 as we like if we take s small enough. There are similar statements in terms of Sobolev spaces (using Lemma 5.17).

REMARK 5.28. Suppose that we defined A in the same way as in Section 3 except that we replaced (3.3) with the requirement that the $\gamma_0(p)$'s be unlinked. Then $(\mathbb{R}^3, D_\omega(x, y))$ would be bilipschitz equivalent to $(\mathbb{R}^3, |x - y|)$, no matter how large s is. This follows from the same construction as above. The only place in the preceding construction where we really needed to be in \mathbb{R}^4 was in the proof of Lemma 5.6, but if the $\gamma_0(p)$'s are unlinked then Lemma 5.6 works in \mathbb{R}^3 .

6. Another interesting class of strong- A_∞ weights.

In this section we shall give another construction of strong A_∞ weights which can be viewed as a more refined version of Proposition 4.4 and from which Theorem 1.15 will be derived. We begin with a definition.

Definition 6.1. A closed set $E \subseteq \mathbb{R}^n$ will be called a snowflake of order s , $s > 1$, if there is a metric $\rho(x, y)$ on E and a constant $C > 0$ such that

$$(6.2) \quad C^{-1} |x - y|^s \leq \rho(x, y) \leq C |x - y|^s,$$

for all $x, y \in E$.

In this definition a metric simply means a nonnegative symmetric function which satisfies the triangle inequality, etc.

The usual Von Koch snowflake curve in \mathbb{R}^2 is a snowflake of order s for a computable value of s . In general we can use Assouad's Theorem 1.14 to generate plenty of snowflakes in the sense of Definition 6.1, with the order s given by α^{-1} , where α is as in Theorem 1.14. One can show that uniformly disconnected sets (in the sense of Definition 4.2) are snowflakes of order s for all $s > 1$. On the other hand, no snowflake of order $s > 1$ can contain a nonconstant rectifiable curve, for if $\gamma(t)$, $0 \leq t \leq 1$, were such a curve, and if $\rho(x, y)$ is as in Definition 6.1, then $\rho(\gamma(0), \gamma(1))$ would have to be zero, because of the assumptions on ρ . (Think about $\sum \rho(\gamma(t_i), \gamma(t_{i+1}))$ for partitions $\{t_i\}$ of $[0, 1]$ of small mesh.)

Theorem 6.3. Let $E \subseteq \mathbb{R}^n$ be a snowflake of order s . Then $\omega(x) = \text{dist}(x, E)^{n(s-1)}$ is a strong A_∞ continuous weight on \mathbb{R}^n such that

$$(6.4) \quad C^{-1} D_\omega(x, y) \leq |x - y|^s \leq C D_\omega(x, y),$$

for some C and all $x, y \in E$, where D_ω is as in Definition 1.5. (Actually, (6.4) holds as soon as one of x or y lies in E .)

In other words, if ρ is as in Definition 6.1, then we can build our strong A_∞ weight ω in such a way that D_ω is comparable to ρ on E . This kind of extension principle works in much greater generality, i.e., we take a distance function on a set E with certain properties, and we build a strong A_∞ weight on all of \mathbb{R}^n whose associated distance function essentially gives this distance function back again on E . For the present purposes however this more modest result is adequate and it has the nice feature of being much simpler.

Theorem 1.15 will follow once we prove Theorem 6.3. Indeed, if $(M, d(x, y))$ is a metric space which satisfies a doubling condition, then Assouad's Theorem 1.14 implies that $(M, d(x, y))^\alpha$ is bilipschitz equivalent to some subset of some \mathbb{R}^n for any given value of $\alpha \in (0, 1)$. The

closure of this subset is an s -snowflake with $s = \alpha^{-1}$ (with ρ coming from $d(\cdot, \cdot)$), and the continuous weight ω provided by Theorem 6.3 satisfies the requirements of Theorem 1.15.

The rest of this section will be devoted to the proof of Theorem 6.3. From now on let E and ω be as in Theorem 6.3, let $\rho(\cdot, \cdot)$ be as in Definition 6.1, and let $D_\omega(\cdot, \cdot)$, the ω -length, and $B_{x,y}$ be defined as in Definition 1.5. Set $\delta_\omega(x, y) = \omega(B_{x,y})^{1/n}$, as in Section 1. The following will be the main step.

Lemma 6.5. *There is a constant C so that $D_\omega(x, y) \geq C^{-1}|x - y|^s$ for all $x, y \in \mathbb{R}^n$.*

Let $x, y \in \mathbb{R}^n$ be given, and let Γ be a curve in \mathbb{R}^n that joins them. More precisely, Γ should be a continuous map from $[0, L]$ into \mathbb{R}^n for some L with $\Gamma(0) = x$ and $\Gamma(1) = y$, and we shall assume that Γ is parameterized by arclength, so that Γ is Lipschitz continuous with norm 1 and $|x - y| \leq L$ in particular. We want to show that the ω -length of Γ is bounded from below by $C^{-1}|x - y|^s$. The idea is that this is easy when Γ stays far away from E , where ω is large, and that we can use our snowflake condition to get estimates when Γ gets close to E . To make this precise we need to break up Γ into simpler pieces that stay away from E and then recombine them.

Let Δ denote the collection of closed subintervals of $[0, L]$ which are dyadic with respect to $[0, L]$. That is, if $L = 1$ these intervals are dyadic in the usual sense, but in general they are not quite the usual dyadic intervals because we take $[0, L]$ itself, its two halves, the two halves of those, etc. These dyadic intervals are all of the form $[k2^{-j}L, (k+1)2^{-j}L]$ for some nonnegative integers j and k .

Set $F = \{t \in [0, L] : \Gamma(t) \in E\}$ and $U = [0, L] \setminus F$. Let \mathcal{M} denote the collection of maximal dyadic subintervals I of $[0, L]$ such that

$$(6.6) \quad 10|I| \leq \inf_{t \in I} \text{dist}(\Gamma(t), E),$$

where $|I|$ denotes the length of the interval I . By standard reasoning U is the union of the elements of \mathcal{M} , and two distinct intervals in \mathcal{M} are either disjoint or intersect in one of their common endpoints.

Sublemma 6.7. $|x - y|^s \leq C \omega\text{-length}(\Gamma)$ when $[0, L] \in \mathcal{M}$.

Indeed, in this case we have that $\text{dist}(\Gamma(t), E) \geq 10L$ for all $t \in [0, L]$, and since $L \geq |x - y|$ (because we are using an arclength

parameterization) we get the desired estimate from the definitions of ω and the ω -length.

From now on we assume that $[0, L] \notin \mathcal{M}$.

Sublemma 6.8. *If $I \in \mathcal{M}$ then $\inf_{t \in I} \text{dist}(\Gamma(t), E) \leq 30 |I|$.*

Indeed, otherwise the (dyadic) parent of I would satisfy (6.6), in contradiction to the maximality of I .

Let Γ_J denote the restriction of Γ to any closed subinterval J of $[0, L]$. A finite sequence I_1, I_2, \dots, I_k of closed intervals will be called a chain if the the right endpoint of I_j is the same as the left endpoint of I_{j+1} for each $j < k$.

Sublemma 6.9. *There is a constant $C > 0$ so that if $J = [a, b]$ is the union of a chain of intervals I_1, I_2, \dots, I_k in \mathcal{M} , then $|a - b|^s \leq C \omega\text{-length}(\Gamma_J)$.*

Pick $v_j \in E$ such that $\text{dist}(\Gamma_{I_j}, v_j) = \text{dist}(\Gamma_{I_j}, E)$ for $j = 1, 2, \dots$. Thus $\text{dist}(\Gamma_{I_j}, v_j) \leq 30 |I_j|$ by Sublemma 6.8. Using $\rho(\cdot, \cdot)$ we get that

$$\begin{aligned} |v_1 - v_k|^s &\leq C \rho(v_1, v_k) \\ (6.10) \quad &\leq C \sum_{j=1}^{k-1} \rho(v_j, v_{j+1}) \leq C \sum_{j=1}^{k-1} |v_j - v_{j+1}|^s. \end{aligned}$$

On the other hand

$$\begin{aligned} |v_j - v_{j+1}| &\leq \text{dist}(\Gamma_{I_j}, v_j) + \text{dist}(\Gamma_{I_{j+1}}, v_{j+1}) \\ (6.11) \quad &\quad + \text{diam}(\Gamma_{I_j} \cup \Gamma_{I_{j+1}}) \\ &\leq C(|I_j| + |I_{j+1}|) \end{aligned}$$

by our choice of the v_i 's, Sublemma 6.8, and the fact that we chose Γ to be parameterized by arclength. Using (6.6) we get that the ω -length of each Γ_{I_j} is at least $C^{-1}|I_j|^s$, and so we conclude from (6.11) that $|v_j - v_{j+1}|^s \leq C \omega\text{-length}(\Gamma_{I_j} \cup \Gamma_{I_{j+1}})$. Putting this back into (6.10) we obtain

$$(6.12) \quad |v_1 - v_k|^s \leq C \sum_{j=1}^{k-1} \omega\text{-length}(\Gamma_{I_j} \cup \Gamma_{I_{j+1}}) \leq C \omega\text{-length}(\Gamma_J).$$

In the same manner we can obtain that $|a - v_1| \leq \text{dist}(\Gamma_{I_1}, v_1) + \text{diam}(\Gamma_{I_1}) \leq C |I_1|$ and that $|I_1|^s \leq C \omega\text{-length}(\Gamma_{I_1})$, whence $|a - v_1|^s \leq$

$C\omega$ -length(Γ_{I_1}). Similarly $|b - v_k|^s \leq C\omega$ -length(Γ_{I_k}). Combining these estimates with (6.12) we get the desired conclusion.

Sublemma 6.13. *Let J be a maximal subinterval of U , and let a, b be its endpoints. Then $|a - b|^s \leq C\omega$ -length(Γ_J).*

This follows from the previous lemma and a limiting argument.

Sublemma 6.14. *Let $\varepsilon > 0$ be given. There is a finite chain of points $0 = t_0 < t_1 < \dots < t_m = L$, $t_j \in F \cup \{0, L\}$, such that for each j either $|t_j - t_{j+1}| < \varepsilon$ or t_j, t_{j+1} are the endpoints of an interval J as in Sublemma 6.13.*

This is an easy exercise.

We are now ready to finish the proof of Lemma 6.5. Let $\varepsilon > 0$ be given, and let $\{t_j\}$ be as in Sublemma 6.14. Then $\Gamma(t_j) \in E$ when $1 \leq j < m$, and so using $\rho(\cdot, \cdot)$ we get that

$$\begin{aligned}
 (6.15) \quad |\Gamma(t_1) - \Gamma(t_{m-1})|^s &\leq C \rho(\Gamma(t_1), \Gamma(t_{m-1})) \\
 &\leq C \sum_{j=1}^{m-2} \rho(\Gamma(t_j), \Gamma(t_{j+1})) \\
 &\leq C \sum_{j=1}^{m-2} |\Gamma(t_j) - \Gamma(t_{j+1})|^s.
 \end{aligned}$$

Hence $|\Gamma(t_0) - \Gamma(t_m)|^s \leq C \sum_{j=0}^{m-1} |\Gamma(t_j) - \Gamma(t_{j+1})|^s$. The terms in this sum fall into two categories. The first are the terms for which t_j, t_{j+1} are the endpoints of an interval J as in Sublemma 6.13. The sum of these terms is at most $C\omega$ -length(Γ), by Sublemma 6.13. For the remaining terms we have $|t_j - t_{j+1}| < \varepsilon$. Since we are assuming that Γ is parameterized by arclength, for such a j we have that $|\Gamma(t_j) - \Gamma(t_{j+1})|^s \leq \varepsilon^{s-1} |\Gamma(t_j) - \Gamma(t_{j+1})|$, and so the sum of these terms is dominated by ε^{s-1} times the Euclidean length of Γ . Altogether we get that

$$(6.16) \quad |x - y|^s = |\Gamma(t_0) - \Gamma(t_m)|^s \leq C\omega\text{-length}(\Gamma) + C\varepsilon^{s-1}\text{length}(\Gamma).$$

Sending ε to 0 we get that $|x - y|^s \leq C\omega$ -length(Γ), which proves Lemma 6.5.

In order to derive Theorem 6.3 from Lemma 6.5 we just need to make some simple observations.

Lemma 6.17. *If B is a ball in \mathbb{R}^n such that $3B$ is disjoint from E , then $\sup_{2B} \omega \leq C \inf_{2B} \omega$ for some constant C which does not depend on B .*

This follows from the definition of ω and simple geometric considerations.

Lemma 6.18. *There is a constant C so that $D_\omega(x, y) \geq C^{-1} \delta_\omega(x, y)$ for all $x, y \in \mathbb{R}^n$.*

If $3B_{x,y}$ touches E , then $\delta_\omega(x, y) \leq C|x - y|^s$ by definition of δ_ω and ω , and the desired inequality follows from Lemma 6.5. If $3B_{x,y}$ is disjoint from E , then it is easy to derive the required inequality from Lemma 6.17 and the definition of $D_\omega(x, y)$.

Next we need the following estimate on the thinness (or porosity) of E .

Lemma 6.19. *There is a constant C so that for each ball $B(x, r)$ in \mathbb{R}^n we can find a point $z \in B(x, r/2)$ such that $\text{dist}(z, E) \geq C^{-1}r$.*

Let y be any point in $\partial B(x, r/2)$, and let γ denote the segment which joins x to y . If $2B_{x,y}$ is disjoint from E then there is nothing to prove. If not, then we can apply Lemma 6.5 to conclude that the ω -length of γ is at least $C^{-1}|x - y|^s$. This implies that $\sup_{z \in \gamma} \text{dist}(z, E) \geq C^{-1}|x - y|$, because of the definitions of ω and the ω -length, and the lemma follows.

Lemma 6.20. *E has Lebesgue measure zero, and there is a constant $C > 0$ so that*

$$(6.21) \quad \sup_{2B} \omega \leq C \frac{1}{|B|} \int_B \omega,$$

for all balls B in \mathbb{R}^n .

The first part is a consequence of Lemma 6.19, which implies that E can have no points of density. (See also the remarks after Lemma

4.6.) The second part follows from Lemma 6.17 when $3B$ is disjoint from E . If $3B$ intersects E , then $\sup_{P_{2B}} \omega \leq (5 \text{ radius}(B))^{n(s-1)}$, while the right side of (6.21) is bounded from below by $C^{-1} \text{radius}(B)^{n(s-1)}$ for some constant C , because of Lemma 6.19 and the definition of ω . This proves Lemma 6.20.

Lemma 6.20 implies that ω is a continuous weight which is also doubling. Another easy consequence of (the second part of) Lemma 6.20 is that $D_\omega(x, y) \leq C \delta_\omega(x, y)$. Combining this with Lemma 6.18 we obtain that ω is a strong A_∞ continuous weight.

To prove the first inequality in (6.4), notice that if at least one of x, y lies in E , then $\text{dist}(z, E) \leq |x - y|$ for all z on the line segment that joins x and y , and the ω -length of this line segment is less or equal than $C|x - y|^s$. The second inequality in (6.4) comes from Lemma 6.5 and is true for all x, y , and the proof of Theorem 6.3 is now complete.

7. Metric spaces which do not admit bilipschitz embeddings into Euclidean spaces.

Theorem 7.1. *There is a metric space M which satisfies a doubling condition (as in Definition 1.13) but which is not bilipschitz equivalent to a subset of any Euclidean space.*

This theorem was known to Assouad, and it is an easy consequence of [P], but it does not seem to have been stated explicitly anywhere.

To prove Theorem 7.1 we use the 3-dimensional Heisenberg group (or any other Carnot group, as in [P, Definition 1.2]) equipped with its Carnot metric. For the sake of simplicity we leave the precise definitions to [P] (see especially paragraph 1.1 on p. 3), but basically the Carnot metric on the Heisenberg group is a distance function that is defined by minimizing the length of the “horizontal” curves which connect a given pair of points, where a curve is said to be horizontal if at each point it is tangent to a certain (completely nonintegrable) distribution of planes. This distance function is invariant under group translations and scales a certain way under a natural family of dilations. The built-in degeneracy of the metric leads to a certain fractal quality of the resulting metric space. For instance, the 3-dimensional Heisenberg group has Hausdorff dimension 4 with respect to its Carnot metric.

The Heisenberg group with its Carnot metric certainly satisfies a

doubling condition. The group of translations and dilations can be employed to reduce this property to the (true) statement that the unit ball centered at the origin can be covered by a finite number of balls of radius $1/2$.

One of the reasons that the Heisenberg group with its Carnot metric (and other Carnot groups) are so interesting geometrically is that real-valued Lipschitz functions on them are differentiable almost everywhere. This is a special case of [P, Theorem 2, p. 4]. The precise notion of “differentiability” is given in [P, Paragraph 1.3, p. 4], but it comes down to the usual idea: one “blows up” the given function at a given point and asks that a limiting object exist, and one uses the group of translations and dilations to realize the tangent map as a map on the original Heisenberg group (or Carnot group). The theorem in [P] states not only the existence of the differential almost everywhere, but also its realizability as a group homomorphism which is compatible with the respective groups of dilations.

Let us call the 3-dimensional Heisenberg group with its Carnot metric M . If M had a bilipschitz embedding f into some Euclidean space \mathbb{R}^n , then the aforementioned result would imply that f is differentiable almost everywhere in the sense of [P]. The blowing-up procedure used to define the differential scales in the natural way, so that the differential is bilipschitz since f itself is. This gives a contradiction, because any homomorphism from the 3-dimensional Heisenberg group into \mathbb{R}^n must have a kernel which is at least 1-dimensional (all commutators in the Heisenberg group must be mapped to 0 by the homomorphism) and hence cannot be bilipschitz.

Theorem 2 in [P] on the differentiability almost everywhere of Lipschitz functions on the Heisenberg group (or other Carnot groups) actually allows the mapping to take values in another Carnot group and not just the real line. The special (linear) case of real-valued (and hence \mathbb{R}^n -valued) Lipschitz functions should be much older than [P], although I did not find a reference. It is certainly within the realm of the usual subelliptic analysis on Carnot groups, and I doubt that it would be very difficult to adapt the standard methods in Harmonic Analysis for proving the differentiability a.e. of Lipschitz functions on Euclidean spaces (using maximal functions, etc.) to the case of the Heisenberg group and other Carnot groups using the standard tools for doing analysis on these groups (as in [Fo], [FS1], [FS2], [Je], [St2]).

Theorem 7.1 together with Theorem 1.15 imply Theorem 1.12, but unfortunately it is not at all clear how small we can take the dimension d

to be. By this method one would first have to know the best dimension n in Assouad's Theorem 1.14 when M is taken to be the Heisenberg group with the Carnot metric and α is allowed to be any positive number. In any case n is at least 5.

It would be very interesting to have some alternative construction for Theorem 7.1 which is more direct. Aside from understanding how small the dimension d in Theorem 1.12 can be, it would be good to have a simpler and more direct understanding of why examples as in Theorem 7.1 exist.

A related problem is to find other examples of metric spaces for which there is some kind of rigidity theorem along the lines of "real-valued Lipschitz functions are differentiable almost everywhere". I do not know of any examples which are not somehow based on Euclidean geometry or the geometry of Carnot groups, nor do I know whether any such examples should exist. This is related to the WALA and GWALA in [DS4]. (See [DS4, p. 45-6 and Chapter III.4]. The issue there is to decide whether certain uniform rigidity properties of Lipschitz functions on a set $E \subseteq \mathbb{R}^n$ should force E to be "uniformly rectifiable". This problem was not resolved satisfactorily in [DS4].) Notice that no such results are true for self-similar Cantor sets or snowflakes; Lipschitz functions on these types of sets are as flabby as Hölder continuous functions on \mathbb{R}^n of order less than 1.

8. Regular mappings.

Definition 8.1. *Let M and N be metric spaces. A mapping $f : M \rightarrow N$ is said to be regular if it is Lipschitz continuous and if there is a constant $C > 0$ so that if B is a ball in N then $f^{-1}(B)$ can be covered by at most C balls in M of the same radius as B .*

Note that no requirements are being imposed on the position of these balls in M which cover $f^{-1}(B)$.

To my knowledge this kind of condition was first considered in the context of Euclidean spaces (with the metric perhaps deformed by a weight) in [D1], [D2]. The definition given in [D1], [D2] is slightly different but equivalent to this one in the case of Euclidean spaces.

In practice we shall only be considering metric spaces which satisfy a doubling condition, and so bilipschitz embeddings will automatically be regular. Notice that the bilipschitz condition can be reformulated

as meaning that f is Lipschitz and $f^{-1}(B)$ is contained in a single ball whose radius is allowed to be larger than the radius of B by only a bounded factor. Roughly speaking, bilipschitzness is a uniform and scale-invariant version of injectivity, while regularity is a uniform and scale-invariant version of the requirement that a map have bounded multiplicity. (Note that regular maps have bounded multiplicity.)

A simple example of a regular mapping is $f : \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = |x|$. It is easy to draw curves in the plane whose arclength parameterization is regular. They are allowed to cross themselves many times, but they are not allowed to have too much mass concentrate in a disk.

Regular mappings have many of the same nice properties as Lipschitz and bilipschitz mappings. For instance, the regularity condition scales the same way as the Lipschitz and bilipschitz conditions, and the composition of two regular mappings is regular. Regular mappings also share some of the important features of bilipschitz mappings, *e.g.*, they can increase Hausdorff measure of any dimension by only a bounded factor.

There is an analogue of Theorems 7.1 and 1.12 for regular mappings.

Theorem 8.2. *There is a metric space M which satisfies a doubling condition but which does not admit a regular mapping into any Euclidean space.*

Corollary 8.3. *There is a strong- A_∞ weight on some \mathbb{R}^d such that (\mathbb{R}^d, D_ω) does not admit a regular mapping into any Euclidean space.*

Theorem 8.2 is proved in exactly the same way as Theorem 7.1. Take M to be the 3-dimensional Heisenberg group with its Carnot metric, and suppose that $f : M \rightarrow \mathbb{R}^n$ is regular. In particular it is Lipschitz, and so it is differentiable almost everywhere. Because of the natural scale-invariance of the regularity condition we have that the differential of f is also regular whenever it exists. Since the differential is almost always a group homomorphism, we get a contradiction as before, because the kernel of such a homomorphism has dimension at least one.

Corollary 8.3 is an immediate consequence of Theorem 8.2 and Theorem 1.15.

The questions posed at the end of the preceding section (concerning

the possible small values of d in Corollary 8.3 and alternative constructions for Theorem 8.2 and Corollary 8.3) are also open in the case of regular mappings. Here is another one.

PROBLEM 8.4. *Let M be a metric which satisfies a doubling condition. If M admits a regular mapping into some Euclidean space, must M admit a bilipschitz embedding into one also?*

Of course the examples above were the same for Theorems 7.1 and 8.2.

Proposition 8.5. *The answer to Problem 8.4 is affirmative if and only if it is affirmative for the special case of metric spaces of the form $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$, where ω is a strong A_∞ continuous weight.*

This is proved using Theorem 1.15, and the argument is given at the end of the section.

Regular mappings from Euclidean spaces into other (larger) Euclidean spaces are quite interesting. It turns out that such mappings have a considerable amount of bilipschitz behavior, and in particular that they have “large bilipschitz pieces”. See [D1], [D2], [D3], [D4], and [Js], and see [DS3] for a related notion of “weakly bilipschitz”. This good behavior is sufficient to ensure the L^p boundedness of singular integral operators on the image of the mapping, as in [D1], [D2]. Regular mappings are also flexible enough so that there are some general existence results. For instance, suppose that E is a d -dimensional regular subset of \mathbb{R}^n which is “uniformly rectifiable” in the sense of [DS4]. This means that inside each ball centered on E there should be a substantial fraction of E which is bilipschitz equivalent to a subset of \mathbb{R}^d , with uniform bounds. Then part of the main result of [DS2] is that there is an A_1 weight ω on \mathbb{R}^d and a regular mapping ϕ from the metric space $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$ into \mathbb{R}^{n+1} whose image contains E . (See Definition 2.8 for the definition of an A_1 weight, and note that for this result we need to allow discontinuous weights.) It is not known whether this last result is true with the weight ω simply taken to be constant. Because a regular image of $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$ is necessarily uniformly rectifiable for any A_1 weight ω , this question comes down to the following.

PROBLEM 8.6. *Suppose that ϕ is a regular map from $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$*

into some \mathbb{R}^n , where ω is an A_1 weight. Is there then a regular map $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^m$ for some $m \geq n$ such that the image of ψ contains the image of ϕ (with \mathbb{R}^n viewed as a subspace of \mathbb{R}^m)?

Remember that Problem 2.10 asks whether $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$ is bilipschitz equivalent to \mathbb{R}^d with the Euclidean metric when ω is an A_1 weight. If this is true then Problem 8.6 also has an affirmative answer, with ψ simply a reparameterization of ϕ . For that matter Problem 8.6 would have an affirmative answer if there is a regular mapping from \mathbb{R}^d with the Euclidean metric onto $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$.

The main result of [DS2] contains a result of the same type as Problem 8.6. Specifically, the image of a regular mapping from $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$ into some \mathbb{R}^n for ω a strong A_∞ weight is uniformly rectifiable and hence is contained in the image of a regular mapping associated to an A_1 weight. Thus A_1 weights are natural for Problem 8.6. (See also the discussion of open problems in [DS2, Section 21].)

A special case of the main result in [Se4] is that $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$ is bilipschitz equivalent to a subset of some \mathbb{R}^n (with the Euclidean metric) when ω is an A_1 weight. If, for a particular ω , this subset could be put inside the image of a regular mapping $\psi_0 : \mathbb{R}^d \rightarrow \mathbb{R}^m$ as in Problem 8.6, then the answer to Problem 8.6 would be affirmative in general for ω , i.e., all other regular mappings from $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$ into Euclidean space would also be good. To see this we need an auxiliary fact (which is useful to know anyway).

Proposition 8.7. *Suppose that $E \subseteq \mathbb{R}^n$ is closed and that $f : E \rightarrow \mathbb{R}^k$ is regular. Then there exists a regular mapping $F : \mathbb{R}^n \rightarrow \mathbb{R}^{k+n+1}$ such that F agrees with f on E , modulo the identification of \mathbb{R}^k with a subspace of \mathbb{R}^{k+n+1} in the obvious way.*

Assuming Proposition 8.7 for the moment let us apply it to the assertion in the preceding paragraph. Let ω be an A_1 weight on \mathbb{R}^d , and let $E \subseteq \mathbb{R}^n$ be chosen so that E (with the induced Euclidean metric) is bilipschitz equivalent to $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$. Suppose also that E is contained in the image of a regular mapping $\psi_0 : \mathbb{R}^d \rightarrow \mathbb{R}^m$. We may as well assume that $m = n$, because each of m and n can be increased without difficulty. Let $\phi : (\mathbb{R}^d, D_\omega(\cdot, \cdot)) \rightarrow \mathbb{R}^k$ be some other regular map, so that we want to find another regular map ψ as in Problem 8.6. Since E is bilipschitz equivalent to $(\mathbb{R}^d, D_\omega(\cdot, \cdot))$ we get a regular mapping $f : E \rightarrow \mathbb{R}^k$ which is just a reparameterization of ϕ .

Proposition 8.7 provides us with an extension F of f to all of \mathbb{R}^n , and the composition $\psi = F \circ \psi_0$ is a regular mapping from \mathbb{R}^d into \mathbb{R}^{k+n+1} whose image contains the image of ϕ , as desired.

Proposition 8.7 is a simpler and cruder version of the result in [D2, Section 4]. As such it is very similar to [DS2, Proposition 17.4], and the proof below is essentially the same as the argument in [DS2, Section 17], modulo some additional simplifications which are possible in this case.

Lemma 8.8. *Given any closed set $E \subseteq \mathbb{R}^n$ we can find a Lipschitz function $\rho : \mathbb{R}^n \rightarrow \mathbb{R}^n$ which vanishes on E and which has the property that if B is a ball such that $2B \subseteq \mathbb{R}^n \setminus E$, then the restriction of ρ to B is regular, with a uniformly bounded constant.*

In fact we can do this in such a way that $\rho|_B$ is bilipschitz, with a uniform bound, at least if we allow ourselves to increase the dimension of the target space. This is basically what happens in [DS2, Section 17], but it is a little simpler to just get regularity.

Note that the mapping ρ in Lemma 8.8 cannot be orientation preserving, because (as pointed out to me by Juha Heinonen) quasiregular mapping theory would then force E to be discrete.

To prove Lemma 8.8 let $\{Q_i\}$ be a Whitney decomposition of $\mathbb{R}^n \setminus E$ into dyadic cubes, as described on [St1, p. 167ff]. We are going to build ρ by gluing together a bunch of little maps on the Whitney cubes Q_i . We shall do this via an induction on skeleta, with the next observation supplying the induction step.

Sublemma 8.9. *Given a j -dimensional cube Q in \mathbb{R}^n and a regular mapping $\sigma : \partial Q \rightarrow \mathbb{R}^l$ with constant C_0 (as in Definition 8.1) we can find an extension $\Sigma : Q \rightarrow \mathbb{R}^{l+1}$ of σ (modulo the obvious identification of \mathbb{R}^l with $\mathbb{R}^l \times \{0\}$) which is regular with constant $C C_0$, where C depends only on the dimensions. (Here ∂Q refers to the “polyhedral” boundary of Q , since Q will be its own topological boundary when $j < n$. Also, $j = 1$ and $l = 0$ are allowed here, with \mathbb{R}^l interpreted as the trivial vector space with only the zero element when $l = 0$.)*

This is easy to prove. Assume for simplicity that Q is centered at the origin and that $0 \in \sigma(\partial Q)$. Let tQ denote the image of Q under the mapping $x \rightarrow tx$. If $x \in \partial(tQ)$, define $\Sigma(x)$ by taking the \mathbb{R}^l part of $\Sigma(x)$ to be $t\sigma(t^{-1}x)$ (0 when $x = 0$) and the last coordinate

to be $(1 - t) \text{diam } Q$. It is easy to check that this defines a regular mapping with the correct estimate. Note that in the $j = 1$ case this gives a piecewise-linear mapping on a segment which vanishes at the endpoints, is positive in the middle, and does not preserve orientations.

To define ρ , we begin by setting $\rho = 0$ on E and also on all the vertices of the Whitney cubes Q_i . At this stage we can view ρ as a map into \mathbb{R}^0 . We extend ρ to the various edges of the Whitney cubes using Sublemma 8.9. Actually, we have to be a little careful; let us call an edge of a Whitney cube *minimal* if it does not properly contain an edge of another Whitney cube, and let us call the collection of all these minimal edges the “minimal edges of the Whitney decomposition”. Thus the edge of any Whitney cube is the union of (a bounded number of) minimal edges, disjoint except at the vertices, and the minimal edges contained in any Whitney cube Q_i cannot be smaller than a fixed constant times the sidelength of Q_i . This follows from the fact that if two Whitney cubes intersect, then the intersection must be a face of one of the two cubes (of some dimension j , $0 \leq j < n$), and the two cubes must have approximately the same size. (In fact one can take the “fixed constant” mentioned above to be $1/4$. See [St1, Proposition 1, p. 169].) We extend ρ to the edges of the Whitney cubes by applying Sublemma 8.9 to the minimal edges, which then takes care of all the others. This gives rise to a map into \mathbb{R}^1 which is regular on each minimal edge, and hence on any edge of any Whitney cube, with a uniformly bounded regularity constant. We then extend ρ to the various squares using Sublemma 8.9, so that this part of ρ takes values in \mathbb{R}^2 . Note that a mapping on the boundary of a square is regular if it is continuous at the vertices and regular on each of the four sides of the square, so that we can apply Sublemma 8.9. Also, as before, we should really work with the “minimal squares of the Whitney decomposition”, etc. Repeating this argument for the various dimensions up to n we get a map ρ which is defined and regular with a bounded constant on each of the Whitney cubes, and which is also continuous as one passes from a Whitney cube to its neighbor. Because ρ vanishes at the vertices of the Whitney cubes we also obtain that $|\rho(x)| \leq C \text{dist}(x, E)$ for some constant C and all x . This implies that ρ is continuous across E , and it is not hard to check that ρ is Lipschitz on all of \mathbb{R}^n , and not just on the various Whitney cubes. The regularity property on balls stated in Lemma 8.8 follows easily from the fact that any such ball is covered by a bounded number of Whitney cubes. This proves Lemma 8.8.

Now let us prove Proposition 8.7. Let $f : E \rightarrow \mathbb{R}^k$ be given (and regular), and extend it to a Lipschitz map (also denoted by f) from \mathbb{R}^n into \mathbb{R}^k . Identify \mathbb{R}^{k+n+1} with $\mathbb{R}^k \times \mathbb{R}^n \times \mathbb{R}$, and identify \mathbb{R}^k with $\mathbb{R}^k \times \{0\} \times \{0\}$. Let $\rho : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be as in Lemma 8.8, and define $h : \mathbb{R}^n \rightarrow \mathbb{R}$ by $h(x) = \text{dist}(x, E)$. Take F to be (f, ρ, h) .

Let B be a ball in \mathbb{R}^{k+n+1} with radius r , and let t denote the last coordinate of the center of B . We need to show that $F^{-1}(B)$ can be covered by a bounded number of balls of radius r .

Let β denote the ball in \mathbb{R}^k which is the projection of B . By assumption we can cover $f^{-1}(M\beta) \cap E$ with at most $C(M)$ balls of radius r for any given $M > 0$. Let E_s denote the set of points in \mathbb{R}^n at distance at most s from E .

It will be convenient to distinguish some cases concerning the relative values of t and r . If $t < -r$ then every point in B has negative last coordinate. In this case $F^{-1}(B) = \emptyset$, since h is nonnegative. Now suppose that $-r < t \leq 2r$. Then $F^{-1}(B) \subseteq E_{3r}$, by definition of h . For each point $z \in F^{-1}(B)$ there is a point $y \in E$ such that $|y - z| \leq 3r$, and so $|f(y) - f(z)| \leq 3Lr$, where L denotes the Lipschitz norm of f . That is to say, z lies within $3r$ of $f^{-1}((3L+1)\beta) \cap E$. From this it follows easily that $F^{-1}(B)$ can be covered by a bounded number of balls of radius r , since this is true for $f^{-1}((3L+1)\beta) \cap E$, by assumption.

It remains to deal with the case where $t > 2r$. In this case we have that $F^{-1}(B) \subseteq E_{2t} \setminus E_{t/2}$, again because of the definition of h . Let B' denote the ball in \mathbb{R}^{k+n+1} with the same center as B but with radius t , so that $B' \supseteq B$. Then B' is a ball of the type considered in the previous case, and so $F^{-1}(B')$ is covered by a bounded number of balls of radius t . Since $F^{-1}(B) \subseteq \mathbb{R}^n \setminus E_{t/2}$ we obtain that $F^{-1}(B)$ is covered by a bounded number of balls whose doubles do not touch E . Because ρ is regular on each of these balls, with bounded constant, we conclude that $F^{-1}(B)$ is covered by a bounded number of balls with radius r . This proves Proposition 8.7.

Let us now prove Proposition 8.5. The "only if" part is trivial, since the metric spaces coming from strong A_∞ weights satisfy doubling conditions. Conversely, let $(M, d(\cdot, \cdot))$ be a metric space which satisfies a doubling condition and which admits a regular mapping into some Euclidean space. Fix $\alpha \in (0, 1)$, e.g., take $\alpha = 1/2$. By Assouad's Theorem 1.14 we can find a set E in some \mathbb{R}^n such that $(M, d(\cdot, \cdot)^\alpha)$ is bilipschitz equivalent to E (with the Euclidean metric). Thus E is a snowflake of order $1/\alpha$, in the sense of Definition 6.1, and so we can

apply Theorem 6.3 with $s = 1/\alpha$ to get a strong A_∞ weight ω such that $(E, D_\omega(\cdot, \cdot))$ is bilipschitz equivalent to $(M, d(\cdot, \cdot))$. Our assumption that M admits a regular mapping into some Euclidean space translates into the condition that there is a regular mapping $g : (E, D_\omega(\cdot, \cdot)) \rightarrow \mathbb{R}^k$ for some k . In order to prove Proposition 8.5 it suffices to show that there is a regular mapping $G : (\mathbb{R}^n, D_\omega(\cdot, \cdot)) \rightarrow \mathbb{R}^{k+n+1}$ which agrees with g on E , modulo the usual identification of \mathbb{R}^n with a subspace of \mathbb{R}^{k+n+1} . (The sufficiency of this statement comes down to the fact that $(M, d(\cdot, \cdot))$ is bilipschitz equivalent to a subset of some Euclidean space if $(\mathbb{R}^n, D_\omega(\cdot, \cdot))$ is, since $(M, d(\cdot, \cdot))$ is bilipschitz equivalent to $(E, D_\omega(\cdot, \cdot))$.)

The construction of the mapping G is completely analogous to the proof of Proposition 8.7; we simply have make some adjustments for the weight.

Lemma 8.10. *There is a Lipschitz mapping $\tau : (\mathbb{R}^n, D_\omega(x, y)) \rightarrow (\mathbb{R}^n, |x - y|)$ which vanishes on E and which has the property that if B is a ball such that $2B \subseteq \mathbb{R}^n \setminus E$, then $\tau : (B, D_\omega(x, y)) \rightarrow (\mathbb{R}^n, |x - y|)$ is regular, with a uniformly bounded constant.*

This is proved in much the same way as Lemma 8.8 was. Let $\{Q_i\}$ be a Whitney decomposition of $\mathbb{R}^n \setminus E$, and define τ initially on E and on the vertices of the Whitney cubes Q_i as a map into $\mathbb{R}^0 = \{0\}$ which vanishes identically. We want to extend τ to the edges of the Whitney cubes, the higher-dimensional faces, and eventually to the Whitney cubes themselves as before, except that on these various faces we should build τ so that it is regular as a map with the D_ω -metric on the domain and the Euclidean metric on the range. In order to do this we need to have a version of Sublemma 8.9 where σ and Σ are regular as maps with the D_ω -metric on the domain and with the Euclidean metric on the range. This version of Sublemma 8.9 is true, at least when the cube Q in Sublemma 8.9 is contained in a Whitney cube (as it always is for our application). Indeed, $D_\omega(x, y)$ is comparable in size to $(\text{diam } Q_i)^\delta |x - y|$ when x, y lie in a Whitney cube Q_i , where $\delta = 1/\alpha - 1$, and hence a mapping on a subset A of Q_i is regular with respect to $D_\omega(x, y)$ with a bounded constant if and only if it equals $(\text{diam } Q_i)^\delta$ times a mapping which is regular with respect to the Euclidean metric with a bounded constant. This allows the aforementioned D_ω version of Sublemma 8.9 to be derived from the original Euclidean statement. Once we have this D_ω version of Sublemma 8.9 we can construct τ in

the same manner as before, and we get a map into $(\mathbb{R}^n, |x - y|)$ which is regular with respect to $D_\omega(x, y)$ on each Whitney cube Q_i , and with a bounded constant. The regularity property on balls required in Lemma 8.10 again follows from the observation that any such ball is covered by a bounded number of Whitney cubes. On the other hand we have that $|\tau(x)| \leq C(\text{diam } Q_i)^{1/\alpha}$ for x in a Whitney cube Q_i , because τ vanishes on the vertices of Q_i (by construction) and is uniformly Lipschitz with respect to D_ω on Q_i , and this is the same as saying that $|\tau(x)|$ is bounded by a constant times the D_ω distance from x to E for any $x \in \mathbb{R}^n$ (because of Theorem 6.3, for instance). Using this bound it is not hard to show that τ has the correct Lipschitz property, and Lemma 8.10 follows.

Next define $h(x)$ for $x \in \mathbb{R}^n$ to be the D_ω -distance from x to E . This is about the same as the Euclidean distance to E raised to the power $1/\alpha$, by Theorem 6.3. Notice that h is Lipschitz as a map from $(\mathbb{R}^n, D_\omega(x, y))$ into \mathbb{R} equipped with the Euclidean metric.

Extend g to a Lipschitz map from $(\mathbb{R}^n, D_\omega(\cdot, \cdot))$ to \mathbb{R}^k equipped with the Euclidean metric. The existence of such an extension follows from the fact that $(\mathbb{R}^n, D_\omega(\cdot, \cdot))$ is a metric space, but in this case one could also use the methods of the Whitney extension theorem (as in [St1, Chapter VI]). Take G to be (g, τ, h) , with \mathbb{R}^{k+n+1} identified with $\mathbb{R}^k \times \mathbb{R}^n \times \mathbb{R}$, and identify \mathbb{R}^k with $\mathbb{R}^k \times \{0\} \times \{0\}$.

The proof that G is regular is practically the same as in the proof of Proposition 8.7 above. Let B be a (Euclidean) ball in \mathbb{R}^{k+n+1} with radius r , and let t denote the last coordinate of the center of B . We need to show that $G^{-1}(B)$ can be covered by a bounded number of ω -balls of ω -radius r . As before, there is nothing to do when $t < -r$.

For this proof let E_s denote the set of points $z \in \mathbb{R}^n$ such that the ω -distance from z to E is $\leq s$, i.e., $h(z) \leq s$. If we assume now that $-r < t < 2r$, then $G^{-1}(B) \subseteq E_{3r}$. Thus for each $z \in G^{-1}(B)$ there is a $y \in E$ with $|g(y) - g(z)| \leq L D_\omega(y, z) \leq 3Lr$, where L is the Lipschitz norm of G . This implies that the ω -distance from z to $g^{-1}((3L+1)\beta) \cap E$ is $\leq 3r$, where β is again the projection of B in \mathbb{R}^k . This implies that $G^{-1}(B)$ is contained in the union of a bounded number of ω -balls of radius $(3L+4)r$, since $g|_E$ is assumed to be regular, and the doubling property for the metric space $(\mathbb{R}^n, D_\omega(x, y))$ allows us to conclude that $G^{-1}(B)$ is covered by a bounded number of ω -balls of radius r , which is what we wanted.

Now suppose that $t > 2r$. In this case we have that $G^{-1}(B) \subseteq$

$E_{2t} \setminus E_{t/2}$, because of the definition of h . Let B' denote the (Euclidean) ball in \mathbb{R}^{k+n+1} with the same center as B but with radius t , so that $B' \supseteq B$. Then B' is a ball of the type considered in the previous case, and so $G^{-1}(B)$ is covered by a bounded number of ω -balls of radius t . Since $G^{-1}(B) \subseteq \mathbb{R}^n \setminus E_{t/2}$ we obtain that $G^{-1}(B)$ is covered by a bounded number of Euclidean balls whose doubles do not touch E . This is not too hard to check, but one must be a little careful. (One way to do this uses the observation that for each $\varepsilon > 0$ we can cover $G^{-1}(B)$ by at most $C(\varepsilon)$ ω -balls of radius εt , because of the doubling property for (\mathbb{R}^n, D_ω) . If ε is small enough, then these smaller ω -balls will be contained in Euclidean balls whose doubles do not touch E . Alternatively, all the relevant ω -balls of radius t which cover $G^{-1}(B)$ lie in E_{4t} , and we can use the fact that that D_ω is comparable to a constant multiple of the Euclidean metric on $E_{10t} \setminus E_{t/5}$ in this special situation.) Therefore τ is regular on each of these balls, with bounded constant, and we conclude that $G^{-1}(B)$ is covered by a bounded number of ω -balls with radius r .

Thus we obtain that G is regular, which is what we wanted. This completes the proof of Proposition 8.5.

Acknowledgements. I am grateful to some topologists for telling me some things that I needed to know, but I am embarrassed for taking so long to realize their importance. I would also like to thank Juha Heinonen for pointing out the ridiculousness of one of my sublemmas before it was too late, and for other comments as well.

References.

- [A1] Assouad, P., *Espaces Métriques, Plongements, Facteurs*. Thèse de Doctorat, Université de Paris XI, France, 1977.
- [A2] Assouad, P., Étude d'une dimension métrique liée à la possibilité de plongement dans R^n . *C. R. Acad. Sci. Paris* **288** (1979), 731-734.
- [A3] Assouad, P., Plongements Lipschitziens dans R^n . *Bull. Soc. Math. France* **111** (1983), 429-448.
- [B] Blankinship, W. A., Generalization of a construction of Antoine. *Ann. of Math.* **53** (1951), 276-297.
- [C1] Cannon, J., The characterization of topological manifolds of dimension $n \geq 5$. *Proceedings I.C.M.* (Helsinki, 1978), 449-454.

- [C2] 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** (1971), Springer-Verlag.
- [CW2] Coifman, R., Weiss, G., Extensions of Hardy spaces and their use in analysis. *Bull. Amer. Math. Soc.* **83** (1977), 569-645.
- [Da] Daverman, R., *Decompositions of Manifolds*, Academic Press, 1986.
- [D1] David, G., Opérateurs intégraux singuliers sur certaines courbes du plan complexe. *Ann. Sci. École. Norm. Sup.* **17** (1984), 157-189.
- [D2] David, G., Opérateurs d'intégrale singulière sur les surfaces régulières. *Ann. Sci. École. Norm. Sup.* **21** (1988), 225-258.
- [D3] David, G., Morceaux de graphes lipschitziens et intégrales singulières sur un surface. *Revista Mat. Iberoamericana* **4** (1988), 73-114.
- [D4] David, G., *Wavelets and Singular Integrals on Curves and Surfaces*. Lecture Notes in Math. **1465** (1991), Springer-Verlag.
- [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 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.
- [E] Edwards, R., The topology of manifolds and cell-like maps. *Proceedings I.C.M.* (Helsinki, 1978), 111-127.
- [Fo] Folland, G., Subelliptic estimates and function spaces on nilpotent Lie groups. *Arkiv Math.* **13** (1975), 161-207.
- [FS1] Folland, G., Stein, E., Estimates for the $\bar{\partial}_b$ -complex and analysis on the Heisenberg group. *Comm. Pure Appl. Math.* **27** (1974), 429-522.
- [FS2] Folland, G., Stein, E., Hardy Spaces on Homogeneous Groups. *Math. Notes* **28** (1982), Princeton University Press.
- [Ga] Garnett, J., *Bounded Analytic Functions*, Academic Press, 1981.
- [Ge] Gehring, F. W., The L^p integrability of the partial derivatives of a quasiconformal mapping. *Acta Math.* **130** (1973), 265-277.

- [Gr] Gromov, M., *Structures Métriques pour les Variétés Riemanniennes*. (J. Lafontaine, P. Pansu, ed.), Cedric/Fernand Nathan, 1981.
- [HK1] Heinonen, J., Koskela, P., Definitions of quasiconformality. *Invent. Math.* **120** (1995), 61-79.
- [Je] Jerison, D., The Poincaré inequality for vector fields satisfying Hörmander's condition. *Duke Math. J.* **53** (1986), 503-523.
- [Jr] Journé, J. L., *Calderón-Zygmund Operators, Pseudodifferential Operators, and the Cauchy Integral of Calderón*. Lecture Notes in Math. **994**, Springer-Verlag, 1983.
- [Js] Jones, P., Lipschitz and bi-Lipschitz functions. *Revista Mat. Iberoamericana* **4** (1988), 115-122.
- [LZ] Luukkainen, J., Väisälä, J., Elements of Lipschitz topology. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **3** (1977), 85-122.
- [Mo] Moise, E. E., *Geometric Topology in Dimensions 2 and 3*. Graduate Texts in Math. **47**, Springer-Verlag, 1977.
- [MRV] Martio, O., Rickman, S., Väisälä, J., Topological and metric properties of quasiregular mappings. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **488** (1971), 1-31.
- [P] Pansu, P., Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un. *Ann. of Math.* **129** (1989), 1-60.
- [R] Rushing, T. B., Hausdorff dimension of wild fractals. *Trans. Amer. Math. Soc.* **334** (1992), 597-613.
- [Se1] Semmes, S., Chord-arc surfaces with small constant I. *Advances in Math.* **85** (1991), 198-223.
- [Se2] Semmes, S., Chord-arc surfaces with small constant II: Good parameterizations. *Advances in Math.* **88** (1991), 170-199.
- [Se3] Semmes, S., Hypersurfaces in \mathbb{R}^n whose unit normal has small BMO norm. *Proc. Amer. Math. Soc.* **112** (1991), 403-412.
- [Se4] Semmes, S., Bilipschitz mappings and strong A_∞ weights. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **18** (1993), 211-248.
- [Se5] Semmes, S., Finding structure in sets with little smoothness. *Proc. I.C.M. (Zürich, 1994)*, 875-885. Birkhäuser, 1995.
- [Se6] Semmes, S., Good metric spaces without good parameterizations. *Revista Mat. Iberoamericana* **12** (1996), 187-275.
- [SS] Siebenmann, L., Sullivan, D., On complexes that are Lipschitz manifolds. In *Geometric Topology*, J. Cantrell, ed., Academic Press, 1979, 503-525.
- [St1] Stein, E. M., *Singular Integrals and Differentiability Properties of Functions*, Princeton University Press, 1970.

- [St2] Stein, E. M., *Harmonic Analysis: Real-Variable Methods, Orthogonality, and Oscillatory Integrals*, Princeton University Press, 1993.
- [T1] Toro, T., Surfaces with generalized second fundamental form in L^2 are Lipschitz manifolds. *J. Differential Geom.* **39** (1994), 65-102.
- [T2] Toro, T., Geometric conditions and existence of bilipschitz parameterizations. *Duke Math. J.* **77** (1995), 193-227.

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
`semmes@ihes.fr`

* The author is supported by the U.S. National Science Foundation.