

COXETER GROUPS AND RANDOM GROUPS

DANNY CALEGARI

ABSTRACT. For every dimension d , there is an infinite family of convex cocompact reflection groups of isometries of hyperbolic d -space — the *superideal* (simplicial and cubical) reflection groups — with the property that a random group at any density less than a half (or in the few relators model) contains quasiconvex subgroups commensurable with some member of the family, with overwhelming probability.

CONTENTS

| | |
|------------------------------------|----|
| 1. Introduction | 1 |
| 2. Statements of the main theorems | 2 |
| 3. Spines | 4 |
| 4. Constructing spines | 6 |
| 5. Acknowledgments | 17 |
| References | 17 |

1. INTRODUCTION


The recent combined work of Wise [7], Ollivier-Wise [6] and Agol [1] shows that random groups at density $< 1/6$ virtually embed in right-angled Artin groups. Our main theorem is complementary to this (at least in flavor), and is concerned with virtually embedding certain classes of Coxeter groups in random groups (at any density).

A *superideal* hyperbolic polyhedron is obtained by taking a polyhedron P in real projective space (of some dimension) and intersecting it with the region H bounded by a quadric (which is identified with hyperbolic space in the Klein model), so that one obtains an infinite hyperbolic polyhedron $P \cap H$. If P is a regular polyhedron and H is centered at the center of P , one obtains a *regular superideal polyhedron* $P \cap H$, for which the symmetries of P are realized as hyperbolic isometries of $P \cap H$. For judiciously chosen P , the dihedral angles of $P \cap H$ might be of the form $2\pi/m$ for some integer m , and hyperbolic space can then be tiled by copies of $P \cap H$. The symmetry group of this tiling is a *superideal reflection group*.

In every dimension d , there are two interesting infinite families of convex cocompact hyperbolic reflection groups — the *superideal simplex* reflection group, and the *superideal cube* reflection group. We denote the simplex groups by $\Delta(m, d)$ and the cube groups by $\square(m, d)$, where the dihedral angles of the superideal regular polyhedra in each case are $2\pi/m$.

Date: November 6, 2014.

These reflection groups are Coxeter groups; for $\Delta(m, d)$ the Coxeter diagram is



whereas for $\square(m, d)$ the Coxeter diagram is



where in each case there are $d + 1$ nodes.

The main result we prove in this paper is that for any d , a random finitely presented group at any density less than a half (or in the few relators model) contains quasiconvex subgroups which are commensurable with some $\Delta(m, d)$ for some large m , and with some $\square(m, d)$ for some large (possibly different) m , with overwhelming probability.

The two cases are treated in a similar way, but the combinatorial details of the argument are different in each case. The precise statements of our main theorems are the Superideal Simplex Theorem and the Superideal Cube Theorem, stated in § 2. The proof of these theorems is carried out in § 3 and § 4.

These theorems generalize previous joint work of the author with Alden Walker [2] and with Henry Wilton [3], and the architecture of the proof greatly resembles the proofs of the main theorems in those papers. The new ideas in this paper are mainly combinatorial in nature.

2. STATEMENTS OF THE MAIN THEOREMS

We first discuss the superideal simplex groups $\Delta(m, d)$.

Example 2.1. Some low-dimensional examples are familiar:

- In any dimension d , taking $m = 3$ gives the finite Coxeter group A_{d+1} , which is just the symmetric group S_{d+2} . Similarly, in any dimension d , taking $m = 4$ gives the finite Coxeter group BC_{d+1} , the symmetry group of the $(d + 1)$ -dimensional cube (or equivalently of the $(d + 1)$ -dimensional cross-polytope), known as the *hyperoctahedral group*.
- In dimension $d = 2$ there is an (ordinary) hyperbolic simplex with angles $2\pi/m$ whenever $m \geq 7$. The groups $\Delta(m, 2)$ are all commensurable, and are commensurable with the fundamental groups of all closed surfaces of genus at least 2.
- In dimension $d = 3$ the ideal regular simplex has $m = 6$; the group $\Delta(6, 3)$ is commensurable with the fundamental group of the complement of the figure 8 knot. For $m \geq 7$ the superideal “simplices” have infinite volume, but the groups they generate are convex cocompact, with limit set a Sierpinski carpet. As $m \rightarrow \infty$ the limit sets converge to the Apollonian gasket, with Hausdorff dimension approximately 1.3057.
- In dimension $d \geq 4$ the simplices are genuinely superideal whenever $m \geq 5$.

Our first main theorem says that for every fixed dimension d , a “generic” finitely presented group will contain some subgroup isomorphic to a finite index subgroup of some element of the $\Delta(m, d)$ family. Explicitly, we show:

Superideal Simplex Theorem. *For any fixed d , a random group at any density $< 1/2$ or in the few relators model contains (with overwhelming probability) a subgroup*

commensurable with the Coxeter group $\Delta(m, d)$ for some $m \geq 7$, where $\Delta(m, d)$ is the superideal simplex group with Coxeter diagram



where there are $d + 1$ nodes.

Example 2.2. Some special cases of this theorem were already known:

- Taking $d = 2$, this theorem implies that a random group contains surface subgroups (with overwhelming probability). This is the main theorem of Calegari-Walker in the paper [2].
- Taking $d = 3$, this theorem implies that a random group contains a subgroup isomorphic to the fundamental group of a hyperbolic 3-manifold with totally geodesic boundary. This is the main theorem of Calegari-Wilton in the paper [3]. In fact, the Commensurability Theorem proved in that paper, is *exactly* the statement of our main theorem in the case $d = 3$.


We now discuss the superideal cube groups $\square(m, d)$.

Example 2.3. The low-dimensional examples are again familiar:

- (1) In any dimension d , taking $m = 3$ gives the finite coxeter group BC_{d+1} , the hyperoctahedral group. Taking $m = 4$ gives the *hypercubic honeycomb* \tilde{C}_{d+1} , which is virtually abelian, and acts cocompactly on Euclidean space.
- (2) In dimension $d = 2$ the group $\square(m, 2)$ is the symmetry group of the tiling by regular right-angled m -gons, whenever $m \geq 5$. These groups are all commensurable with each other, and with the fundamental groups of all closed surfaces of genus at least 2 (and with all $\Delta(m', 2)$).
- (3) In dimension $d = 3$ the group $\square(5, 3)$ is the symmetry group of the tiling by regular right-angled dodecahedra; this is commensurable with the (orbifold) fundamental group of the orbifold with underlying space the 3-sphere, and cone angle π singularities along the components of the Borromean rings, as made famous by Thurston. For $m \geq 6$ the groups $\square(m, 3)$ are convex cocompact but not cocompact. Similarly, in dimension $d = 4$ the group $\square(5, 4)$ is the symmetry group of the tiling by regular right-angled 120-cells, whereas $\square(m, 4)$ is convex cocompact but not cocompact for $m \geq 6$.
- (4) In dimension $d \geq 5$ the groups $\square(m, d)$ are never cocompact when $m \geq 5$.

The analog of the Superideal Simplex Theorem for such groups is the following:

Superideal Cube Theorem. *For any fixed d , a random group at any density $< 1/2$ or in the few relators model contains (with overwhelming probability) a subgroup commensurable with the Coxeter group $\square(m, d)$ for some $m \geq 5$, where $\square(m, d)$ is the superideal cube group with Coxeter diagram*



where there are $d + 1$ nodes.

The case $d = 2$ reduces to the existence of closed surface subgroups, proved by Calegari-Walker in [2], but all other cases are new.

We prove these theorems under the following two simplifying assumptions:

- (1) that the length n is divisible by a finite list of specific integers (implicitly depending on the density D); and

- (2) that the number k of free generators is sufficiently large depending on d in either case; for the Superideal Simplex Theorem we assume $2k - 1 \geq d + 1$, whereas for the the Superideal Cube Theorem we assume $2k - 1 \geq 2d + 1$.

These assumptions simplify the combinatorics, but they are superfluous and at the end of § 4.8 we indicate why they can be dispensed with, and the theorems are true in full generality.

3. SPINES

The proof of our main theorems is entirely direct and constructive; it depends on building certain 2-dimensional complexes subject to local and global combinatorial constraints from pieces which correspond to relators in the random group (in a precise way). In this section we describe these combinatorial constraints.

3.1. Sets as modules. We frequently need to discuss finite sets together with an action of a finite group. If a finite group H acts on a set A we say that A is an H -module. The cases of interest in this paper are:

- the symmetric group S_d acting on a d -element set by the standard permutation action; and
- the hyperoctahedral group BC_d acting on a $2d$ -element set of d pairs by the standard action, which permutes the pairs and acts on each pair as $\mathbb{Z}/2\mathbb{Z}$.

By abuse of notation, we refer to a set together with such structure as an S_d -module or BC_d -module respectively. If A is a $2d$ element set with a BC_d -module structure, and a is an element of A , the other element of the pair containing a is said to be *antipodal* to a .

3.2. Graphs and 2-complexes. Let G be a random group with k generators at density $D < 1/2$ and length n . That means G is the group defined by a presentation

$$G := \langle x_1, \dots, x_k \mid r_1, r_2, \dots, r_\ell \rangle$$

where $\ell = (2k - 1)^{nD}$, and where each r_i is chosen randomly and independently from the set of all (cyclically) reduced words in the free group $F_k := \langle x_1, \dots, x_k \rangle$ of length n , with the uniform distribution. For an introduction to random groups see [4] or [5].

We let $r = r_1$, and let G_r denote the random 1-relator group with presentation

$$G_r := \langle x_1, \dots, x_k \mid r \rangle$$

The group G is the fundamental group of a 2-complex K , with one vertex, with one edge for each generator, and with one 2-cell for each relator. We denote the 1-skeleton of K by X ; thus X is a k -fold rose. The group G_r is the fundamental group of a subcomplex K_r , whose 1-skeleton is equal to X , and which contains only one 2-cell, the cell attached by the relation $r = r_1$.

Definition 3.1. A *graph over X* is a graph Σ together with a simplicial immersion from Σ into X .

A *morphism over X* between oriented graphs $\Sigma \rightarrow X$, $\Sigma' \rightarrow X$ over X is a simplicial immersion $\Sigma \rightarrow \Sigma'$ such that the composition $\Sigma \rightarrow \Sigma' \rightarrow X$ is the given immersion from $\Sigma \rightarrow X$.

If Σ is a graph over X , the oriented edges of Σ can be labeled by the x_i and their inverses, by pulling back the labels from the edges of X . Conversely, a graph with oriented edges labeled by the x_i and their inverses, has the structure of a graph over X providing no adjacent oriented edges have inverse labels.

Let L be a finite union of simplicial circles (i.e. circles subdivided into edges) with oriented edges labeled by generators or their inverses, in such a way that the cyclic word on each component of L is r . Then L is a graph over X .

We will build a simplicial graph Σ over X together with a morphism $L \rightarrow \Sigma$ over X satisfying certain conditions. These conditions are different for the Superideal Simplex Theorem and for the Superideal Cube Theorem, but they have common features, which we now describe.

Definition 3.2. An (m, d) -regular simplicial spine or (m, d) -regular cubical spine is a graph Σ over X , together with a morphism $L \rightarrow \Sigma$ over X with the following properties:

- (1) The graph Σ is obtained by subdividing a regular graph. That is, all the vertices are either 2-valent (we call these *internal*) or of the same fixed valence > 2 (we call these *genuine*). In the simplex case, each genuine vertex is $d+1$ -valent; in the cube case, each genuine vertex is $2d$ -valent and further admits the structure of a BC_d -module. By abuse of notation, we omit the adjective “genuine” when discussing higher valence vertices unless the meaning would be ambiguous.
- (2) The map $L \rightarrow \Sigma$ is d to 1 in the simplex case, and $2(d-1)$ to 1 in the cube case; i.e. each edge of Σ is the preimage of exactly d (resp. $2(d-1)$) edges of L . Furthermore, the set of preimages of each edge has the structure of an S_d (resp. BC_{d-1})-module.
- (3) For each vertex v of Σ and each pair of distinct incident edges e, e' which are not antipodal in the cube case, there is exactly one segment of L of length 2 whose midpoint maps to v and whose adjacent edges map to e and e' . In the cubical case, if $L(e)$ is the set of $2(d-1)$ edges of L mapping to e with its BC_{d-1} -module structure, then the $2(d-1)$ edges in L adjacent to $L(e)$ map bijectively to the $2(d-1)$ edges of Σ not equal to e or its antipode; this latter set also has a natural BC_{d-1} -module structure by restriction, and we require that this map of BC_{d-1} -modules respect the module structure.
- (4) Each component of L has exactly m vertices which map to genuine vertices of Σ .
- (5) Σ satisfies the *cocycle condition* (to be defined shortly).

Let $L \rightarrow \Sigma$ be a spine satisfying the first three conditions to be (m, d) -regular. The mapping cylinder M is a 2-complex which admits a canonical local embedding into \mathbb{R}^d in such a way that the (combinatorial) symmetries of the link of each vertex or edge are realized by isometries of the ambient embedding (taking module structure into account). Taking a tubular neighborhood of this canonical thickening gives rise to a D^{d-2} -bundle over M with a flat connection with holonomy in the symmetric group S_{d-1} in the simplicial case, or BC_{d-2} in the cubical case, acting by the standard representation. The restriction of this bundle to each component of L therefore determines a conjugacy class in S_{d-1} or BC_{d-2} respectively. The *cocycle condition* is the condition that this conjugacy class is trivial, for each component of L .

Let $L \rightarrow \Sigma$ be an (m, d) -regular spine. Form the mapping cylinder M and then let \overline{M} be the 2-complex obtained from M by attaching a disk to each component of L . Since each component of L has the label r , the tautological immersion $L \rightarrow X$ extends to a cellular immersion $\overline{M} \rightarrow K_r$.

Furthermore, by construction, the fundamental group of \overline{M} is commensurable with the Coxeter group $\Delta(m, d)$ in the simplicial case, and $\square(m, d)$ in the cubical case. We will show in the next section that for r a sufficiently long random word, we can construct an (m, d) -regular spine (of either kind) in which all the maximal paths in Σ containing only internal vertices in their interior (informally, the “topological edges” of Σ) are bigger than any specified constant λ . This is the main ingredient necessary to show that the immersion $\overline{M} \rightarrow K_r \rightarrow K$ is π_1 -injective, and stays injective when a further $(2k - 1)^{nD} - 1$ random relators are added.

4. CONSTRUCTING SPINES

We start with L , a finite union of simplicial circles each labeled with the relator r . The spine Σ is obtained from L by identifying oriented edges of L in groups of d or $2(d - 1)$, respecting labels, in such a way as to satisfy the conditions of a regular spine. We thus obtain Σ as the limit of a (finite) sequence of quotients $L \rightarrow \Sigma_i$ limiting to Σ , where each $L \rightarrow \Sigma_i$ is a morphism over X . Thus, each intermediate Σ_i should immerse in X . We call an identification of subgraphs of Σ_{i-1} with the same labels, giving rise to a quotient $\Sigma_{i-1} \rightarrow \Sigma_i$ over X a *legal quotient* if Σ_i is immersed in X . We assume in the sequel that our quotients are always legal.

In our intermediate quotients Σ_i , some edges are in the image of d (resp. $2(d - 1)$) distinct edges of L , and some are in the image of a unique edge of L . We call edges of the first kind *glued*, and edges of the second kind *free*.

In order to simplify notation, in the sequel we let \mathfrak{d} denote either d in the simplicial case, or $2(d - 1)$ in the cubical case.

We choose a big number λ divisible by \mathfrak{d} , and an even number N with $N \gg \lambda$, and assume for the sake of convenience that $\lambda \cdot N$ divides n , the length of the relator r (the size of λ that we need will ultimately depend only on the density D , but not on the length n of the relators). Then we subdivide each component of L into consecutive *segments* of length λ .

4.1. Creating beachballs. A *block* in L is a sequence of N consecutive segments. Since we assume that $\lambda \cdot N$ divides n , each component of L can be subdivided into $n/(\lambda \cdot N)$ consecutive blocks. Each block is made up of alternating *odd segments* and *even segments*. We say that a \mathfrak{d} -tuple of blocks is *compatible* if their odd segments can all be identified in groups of \mathfrak{d} (in the order in which they appear in the blocks) in such a way that the resulting quotient is legal. This means that each \mathfrak{d} -tuple of even segments at the same location in the \mathfrak{d} blocks must start and end with *different* letters, and the same must be true for the last \mathfrak{d} -tuple of even segments in the blocks immediately preceding the given collection of blocks. The existence of a compatible collection of blocks requires $2k - 1 \geq \mathfrak{d}$.

When \mathfrak{d} blocks are glued together along their odd segments, each collection of even segments *except* for the last one is identified at their vertices, producing a *beachball* — a graph with two vertices and \mathfrak{d} edges, each joining one vertex to the other. See Figure 1.




FIGURE 1. A beachball for $\mathfrak{d} = 9$

We say that a $(\mathfrak{d} + 1)$ -tuple of blocks is *supercompatible* if each sub \mathfrak{d} -tuple is compatible. Evidently, the existence of supercompatible tuples requires $2k - 1 \geq \mathfrak{d} + 1$, which is our second simplifying assumption.

Recall that $\mathfrak{d} = d$ in the simplicial case, and $\mathfrak{d} = 2(d - 1)$ in the cubical case. Suppose r is a random reduced word of length n . Since r is random, if n is sufficiently long it is possible to partition a proportion $(1 - \epsilon)$ of the blocks into compatible \mathfrak{d} -tuples, and glue their odd segments, for ϵ depending only on n , and therefore $1/\epsilon \gg N, \lambda$, with probability $1 - O(e^{-n^c})$. This leaves a proportion ϵ of the blocks left unglued. For each of these unglued blocks, find a collection of \mathfrak{d} glued blocks so that the $\mathfrak{d} + 1$ tuple is supercompatible. If we unglue the collection of \mathfrak{d} glued blocks, and then take $\mathfrak{d} + 1$ copies of our graph, we can glue the resulting $\mathfrak{d}(\mathfrak{d} + 1)$ blocks in compatible groups of \mathfrak{d} .

In the cubical case, in order to give each collection of edges the natural structure of a BC_{d-1} -module, it is convenient to first take some finite number of disjoint copies of L and partition them into $2(d - 1)$ subsets of equal size. The group BC_{d-1} acts on each set of $2(d - 1)$ copies of L as a standard representation, giving these components the structure of a BC_{d-1} -module. When we partition blocks into \mathfrak{d} -tuples (where $\mathfrak{d} = 2(d - 1)$ in the cubical case), each tuple should contain exactly one block from each subset, so that the elements of the block inherit a natural BC_{d-1} action.

Thus at the end of this step, we have constructed Σ_1 which consists of collections of glued edges of length λ (each the image of \mathfrak{d} edges of L), a *reservoir* of beachballs with every edge of length λ and all possible edge labels in almost uniform distribution of possible edge labels, and a *remainder* consisting of a \mathfrak{d} -valent graph in which every edge has length λ . Note too that the mass of the remainder is of order $O(1/N)$, where by assumption we have chosen $N \gg \lambda$. Furthermore in the cubical case, for each component X of the reservoir or remainder, the edges of X incident to each vertex inherit a natural BC_{d-1} -module structure.

4.2. Covering move. The next move is called the *covering move*, and its effect is to *undo* some of the gluing of some previous step, and then to *reglue* in such a way that the net effect is to transform some collection of beachballs into more complicated graphs, each of which is a finite cover of a beachball, with prescribed topology and edge labels. This move does *not* necessarily preserve the total mass of beachballs, and can be used to adjust the total mass of beachballs of any specific kind, at the cost of transforming some other (prescribed) set of beachballs into covers.

To describe this move, we first take some consecutive strings of beachballs and arrange them in a “matrix” form, where each string of beachballs is a column of the matrix, and the different strings make up the rows of the matrix. See Figure 2 for a matrix consisting of 7 columns of 3 consecutive beachballs.




FIGURE 2. A matrix of 7 columns of 3 consecutive beachballs. In this and subsequent figures, glued edges are in red and free edges are in black.

The covering move takes as input a matrix of s columns of r consecutive beachballs with compatible labels on corresponding glued segments. This means the following. In each column there are $r + 1$ consecutive glued segments interspersed with the r beachballs. Thus between each pair of consecutive rows of beachballs there is a row of glued segments, each consisting of \mathfrak{d} glued edges with some label. We require that all the labels on glued edges in the same row are equal.

In the cubical case, the edges in each column have a natural structure of a BC_{d-1} -module. Thus it is possible to “trivialize” the module structure on the union of edges in each row, giving them the structure of a product of a standard BC_{d-1} -module with an s -element set.

The covering move pulls apart each interior row of glued segments (i.e. the rows between a pair of consecutive row of beachballs) into $\mathfrak{d}s$ edges, all with the same label, then permutes them into s new sets of \mathfrak{d} elements and reglues them. More specifically, if we identify the set of $\mathfrak{d}s$ edges with the product $\{1, \dots, s\} \times \{1, \dots, \mathfrak{d}\}$ where BC_{d-1} acts on the second factor, we permute each $\{1, \dots, s\} \times j$ factor by some element of S_s . Thus the move is described on each of these $r - 1$ rows by an element $p_j \in S_s^{\mathfrak{d}}$, the product of \mathfrak{d} copies of the symmetric group S_s .

For each $1 \leq j \leq r$ the s beachballs in the j th row are rearranged by this move into a graph which is a (possibly disconnected) covering space of a beachball of degree s ; which covering space is determined by the permutations p_{j-1} and p_j . If $p_{j-1} = p_j$ it is a trivial covering space, in the sense that it consists just of s disjoint beachballs. However, even if it is trivial as a covering space, if p_{j-1} and p_j are not the identity, the edges making up each beachball might be different from the edges making up the beachballs before the move. We allow two possibilities for the labels on the new beachball covers in each row:

- (1) the beachball cover in the given row can have arbitrary topology, and the labels on the edges are legal; or
- (2) the beachball cover in the given row is trivial (i.e. $p_{j-1} = p_j$) and on each beachball the edges all have the *same* label.

In the former case we do nothing else to the row. In the latter case, since all the edge labels are equal (and the gluing is not legal as it stands), we collapse the beachball to the interior of a glued segment of length 3λ . So the net effect of the covering move is to transform some set of beachballs into covering spaces of the same mass, while possibly eliminating some set of beachballs of comparable mass.

Two kinds of beachball covering spaces are especially important in what follows; these are

- a 2-fold cover of a beachball called a *barrel*; and
- a \mathfrak{d} -fold cover of a beachball called a *bipart*, whose underlying graph is the complete bipartite graph $K_{\mathfrak{d}, \mathfrak{d}}$.

See Figure 3.




FIGURE 3. A barrel and a bipart for $\mathfrak{d} = 4$

We now describe two especially useful examples of covering moves which will be used in the sequel.

Example 4.1 (Elimination). The elimination move trades $2\mathfrak{d}$ beachballs for 2 biparts, and eliminates a collection of \mathfrak{d} beachballs with the same labels. It is convenient to perform this move in such a way that the biparts created have labels of

covering type. This means that the edge labels on the bipart are pulled back from some legal edge labeling on a beachball under the covering projection. See Figure 4.




FIGURE 4. The elimination move for $\mathfrak{d} = 4$

Example 4.2 (Rolling barrels). The rolling move trades 4 beachballs for 2 barrels. This move is self-explanatory. See Figure 5.




FIGURE 5. The rolling move for $\mathfrak{d} = 4$

4.3. Tearing the remainder. We now describe a modification of the gluing which replaces the remainder with a new remainder of similar mass, but composed entirely of biparts. In the process, a set of beachballs of similar mass are transformed into beachballs.

The tear move is applied to two vertices v, v' in a component Y of the remainder. It uses up the two beachballs in the reservoir which are adjacent to v and v' , and a further $3(\mathfrak{d} - 1)$ beachballs drawn from the middle of the reservoir in consecutive groups of three. The result of the move “tears” Y apart at v and v' , producing \mathfrak{d}

copies of each vertex v_i, v'_i for $1 \leq i \leq \mathfrak{d}$, each v_i joined to v'_i by $\mathfrak{d} - 1$ new edges (these new $\mathfrak{d}(\mathfrak{d} - 1)$ edges coming from $(\mathfrak{d} - 1)$ of the beachballs), and transforms the remaining $2\mathfrak{d}$ beachballs into two biparts. In the process, $2\mathfrak{d}$ glued edges are unglued, and then reglued in a different configuration; thus it is necessary for the labels on each set of \mathfrak{d} edges to agree. The tear move is illustrated in Figure 6; as in § 4.2, the segments before and after the move have the same vertical coordinate, and the gluing respects $BC_{\mathfrak{d}-1}$ -module structure on edges in the cubical case.




FIGURE 6. The tear move for $\mathfrak{d} = 4$.

If e, e' are a pair of edges of Y with vertices v_1, v_2 and v'_1, v'_2 , applying the tear move to the pair (v_1, v'_1) and (v_2, v'_2) creates a barrel, and reduced the complexity of the rest of Y . So after applying $O(1/N)$ tear moves, we can transform the entire remainder to a mass of $O(1/N)$ barrels and biparts, while the reservoir still contains almost exactly the same number of beachballs in almost exactly the uniform distribution.

4.4. Hypercube gluing. At this stage of the construction we have a reservoir consisting of an almost equidistributed collection of beachballs, and a relatively small mass of pieces consisting of biparts and barrels. In this section and the next we explain how to glue up beachballs. There are two different (but related) constructions depending on whether we are in the simplicial or cubical case. In the simplicial case we use the *hypercube gluing*.

A finite cover of the hypercube gluing can be used to glue up all the barrels and biparts produced by the tear moves, together with some other complementary collection of pieces produced by the covering move. The existence of such a combinatorial cover in which the pieces from the remainder can be included follows by LERF for free groups, and it is straightforward to use the covering move to produce complementary pieces of the correct topology and labels.

After dealing with the remainder, almost all the mass of the reservoir can be completely glued up using hypercube gluing, and the error can be transformed away using the elimination move, while simultaneously transforming some prescribed collection of beachballs into biparts of covering type, which can themselves be glued up by a covering of the hypercube gluing.




FIGURE 7. Beachball immersed in the 1-skeleton of a cube in dimensions 2, 3 and 4.

It remains to describe the hypercube gluing (recall that we are in the simplicial case). Let C denote the 1-skeleton of the d -dimensional cube, and with each edge of length λ/d . Let e_1, \dots, e_d denote vectors of length λ/d aligned positively along the d coordinate axes in order. There is a path of length λ from the vertex $(0, 0, \dots, 0)$ to the vertex $(\lambda/d, \lambda/d, \dots, \lambda/d)$ of C , obtained by concatenating straight segments of length λ/d in the order e_1, e_2, \dots, e_d . Taking d cyclic permutations of this sequence of paths gives d paths of length λ between extremal vertices of C , whose union is a beachball. This is illustrated in Figure 7 in dimensions 2, 3 and 4.

There are 2^{d-1} pairs of extremal axes of C , and the union of 2^{d-1} suitably labeled beachballs can be arranged along the 1-skeleton of C as above. Gluing the beachballs in this pattern is legal, and the resulting graph locally satisfies the first two conditions to be a (m, d) -regular simplicial spine.

4.5. Lens gluing. The analog of the hypercube gluing for cubical spines is the *lens gluing*. This is the most complicated move in our construction, and we first describe this move in low dimensions before giving the definition in generality.

In dimension 2 a beachball has degree 2 (i.e. it has 2 edges); $BC_1 = \mathbb{Z}/2\mathbb{Z}$ acts as the full permutation group. Two beachballs with suitable labels can be identified along their boundaries, laid out along a circle; see Figure 8.




FIGURE 8. Lens gluing $d = 2$

In dimension 3 a beachball has degree 4; $BC_2 = D_4$ acts by the dihedral group, preserving or reversing a circular order on the 4 edges. Eight beachballs with

suitable labels can be glued up along a $K_{4,4}$ graph in S^3 , where each edge of each beachball runs along a segment of length 2 in the $K_{4,4}$ graph, as in Figure 9 (where one of the vertices is at “infinity” in the figure):




FIGURE 9. Lens gluing $d = 3$

The lens gluing in dimension d is closely related to the hypercube gluing in dimension $d - 1$. Recall the definition of the hypercube gluing move from § 4.4. In this move, 2^{d-2} beachballs of degree $d - 1$ are draped along the 1-skeleton of a $d - 1$ -dimensional cube; each edge of each beachball is subdivided into $d - 1$ segments which run between antipodal vertices of the cube. We let C denote the 1-skeleton of the cube, and denote the immersion of a beachball b into C by $b \rightarrow C$.

For any graph L , the *spherical graph* associated to L is the graph with one S^0 (i.e. two vertices) for each vertex in L , and one $S^0 * S^0$ for each edge in L . If v is a vertex of L , we denote the corresponding pair of vertices in $S(L)$ by v_0 and v_1 . Similarly, if e is an edge of L , we denote the corresponding edges in $S(L)$ by e_{00} , e_{01} , e_{10} and e_{11} (in fact this construction generalizes in an obvious way to simplicial complexes of arbitrary dimension). There is a canonical simplicial map from $S(L) \rightarrow L$ which forgets the 01 labels.

Now we apply this construction to C ; the lens gluing will map $4 \times 2^{d-2}$ beachballs of degree $2(d - 1)$ into $S(C)$ in such a way that the projection to C will map each 4 beachballs onto each degree $(d - 1)$ beachball mapping into C in the hypercube gluing, with each edge of the beachball downstairs in the image of two edges upstairs.

If $b \rightarrow C$ is a beachball, and $b' \rightarrow S(C)$ is one of the four beachballs upstairs mapping to it, each edge of b maps to a path in C of length $d - 1$, which is covered by two paths in $S(C)$ of length d . Each such path is determined (given its image in C) by a word of length d in the alphabet $0, 1$. Our beachballs upstairs have the property that the pair of lifts of each edge of b are described by the *same* pair of 01 words. Moreover, the collection of 4 pairs of 01 words describing the lifts of b are the same for all $b \rightarrow C$. So to describe the lens move we just need to give 4 sets of pairs of 01 words of length d . We call these *height pairs*. An explicit formula for $d > 2$ is given by

$$\text{height pairs are } \begin{cases} (00^{d-2}0, & 01^{d-2}0) \\ (0(01)^{\frac{d-2}{2}}1, & 0(10)^{\frac{d-2}{2}}1) \\ (1(1100)^{\frac{d-2}{4}}x, & 1(0011)^{\frac{d-2}{4}}x) \\ (1(1001)^{\frac{d-2}{4}}(1-x), & 1(0110)^{\frac{d-2}{4}}(1-x)) \end{cases}$$

where x is 0 or 1 depending on the parity of d , and where an expression like $w^{p/q}$ for w a word of length q means the initial string of length p of the word $w^\infty := www \dots$. The meaning of this formula is best explained by a picture; see Figure 10.




FIGURE 10. Height graphs for $d = 9$

The case $d = 2$ is degenerate; the case $d = 3$ gives rise to the height pairs

$$(000, 010), (001, 011), (111, 101), (110, 100)$$

which reproduces four of the beachballs in the lens gluing for $d = 3$ described above (the other four beachballs are in the preimage of the other beachball in the hypercube gluing).

Note that there is an ambiguity in our choice of labels of each pair of vertices of $S(C)$ over v by $v0$ and $v1$; thus the group $\mathbb{Z}/2\mathbb{Z}^{\text{vert}(C)}$ acts by automorphisms of $S(C)$ over C . Thus although this is not evident in the formulae, under this symmetry group all four height pairs are in the same orbit.

The case $d = 4$ is hard to draw without the diagram becoming cluttered; Figure 11 shows the 4 beachballs of degree 6 in the $d = 4$ lens gluing which project to one beachball of degree 3 in the $d = 3$ hypercube gluing.

4.6. Regularity. The third condition for (m, d) -regularity is that each component of L has exactly m vertices which map to high-valent vertices of Σ . The only move which adjusts the number of such vertices on each component of L is the elimination move. This move takes pieces which may be drawn from any component L ; so we simply take enough disjoint copies of L , and symmetrize the components from which the moves are drawn, so that we apply this move the same number of times to every component. This will ensure the third condition for (m, d) -regularity.




FIGURE 11. Four beachballs in the $d = 4$ lens gluing which project to one beachball in the $d = 3$ hypercube gluing.

4.7. Cocycle condition. The cocycle is sensitive to more combinatorial data than we have been using so far; it depends not only on the set of labels appearing as the edges in a beachball, but the way in which these labels come from consecutive even segments on each of \mathfrak{d} blocks in a compatible \mathfrak{d} -tuple. Gluing different compatible \mathfrak{d} -tuples gives rise to the same beachball collections but with different cocycles (relative to a trivialization along the odd segments of the blocks). If the original gluing was done randomly, there are pairs of compatible segments whose gluings give rise to the same set of beachballs, but with cocycles differing by any element in the symmetric group S_{d-1} in the simplicial case, or any element of BC_{d-1} in the cubical case; adjusting the gluing by interchanging elements of these pairs gives the same collection of beachballs and the same remainder (and therefore the same gluing problem) but with the cocycle adjusted in the desired way. Boundedly many moves of this kind for each component produces an (m, d) -regular simplicial or cubical spine with all topological edges of length at least λ .

4.8. Few generators. We now indicate how to modify our arguments and constructions so that they hold without our simplifying assumptions.

The first assumption — that the length n of the relators is divisible by some big fixed constant $\lambda \cdot N$ — is easy to dispense with. We let ρ be the remainder when dividing n by $\lambda \cdot N$, then find \mathfrak{d} copies of a subword of r of length $\lambda \cdot N + \rho$ so that it is legal to glue these \mathfrak{d} words in \mathfrak{d} distinct copies of r ; what is left after this step is a graph with edges whose lengths are divisible by $\lambda \cdot N$, and the rest of the construction can go through as before.

The second assumption — that the number k of free generators satisfies $2k - 1 \geq 2d + 1$ — can be finessed by looking more carefully at the construction. It can be verified that we do not really use the hypothesis that the word r is random with respect to the uniform measure on reduced words of length n ; instead we can make do with a much weaker hypothesis, namely that we generate random words by *any* stationary ergodic Markov process, whose subword distribution of any fixed length has enough symmetry, and such that there are at least $\mathfrak{d} + 1$ letters that may follow any legal substring.

The relevance of this is as follows: a random reduced word of length n in a free group of rank k can be thought of as a random reduced word of length n/s in a free group of rank $k(2k - 1)^{s-1}$ (whose symmetric generating set is the set of reduced words in F_k of length s) generated by a certain stationary ergodic Markov process. We can perform matching of subwords and build spines in this new “alphabet”; the result will not immerse in X , since there might be folding of subtrees of diameter

at most $2s$. But the spine will still have the crucial properties that the fundamental group of the 2-complex \overline{M} is commensurable with $\Delta(m, d)$, and that the number of immersed paths of length ν grows like $(2k - 1)^{s\nu/\lambda}$ where λ is as big as we like. Taking s big enough to “simulate” a group of rank at least $d + 1$, we can then take λ much bigger so that s/λ is arbitrarily small. This small exponential growth rate of subpaths is the key to proving that $\pi_1(\overline{M})$ maps injectively with quasiconvex image.

4.9. Conclusion of the theorem. The remainder of the argument is almost identical to the arguments in [2] and [3], and make use in the same way of the mesoscopic small cancellation theory for random groups developed by Ollivier [5]. It depends on two ingredients:

- (1) a *bead decomposition*, exactly analogous to that in [2], § 5 or [3], § 4; and
- (2) a *mesoscopic small cancellation argument*, which applies techniques of Ollivier as in [2], § 6 or [3], § 6.

We very briefly summarize the argument below just to indicate that no new ideas are needed, but refer the reader to the cited references for details.

A bead decomposition is an (m, d) -regular simplicial or cubical spine Σ which is made up of $O(n^\delta)$ pieces of size $O(n^{1-\delta})$ arranged in a circle, with each piece separated from the next by a *neck* (an unusually long glued segment) of length $C \log n$. We can construct a spine as in [3], § 4 as follows. Fix some small positive constant δ and write r as a product

$$r = r_1 s_1 r_2 s_3 \cdots r_m s_m$$

where each r_i has length approximately $n^{1-\delta}$ and each s_i has length approximately n^δ , and we choose lengths so that m is divisible by \mathfrak{d} . We then fix a small positive constant $C < \delta / \log(2k - 1)$ and look for a common subword x of length $C \log n$ in $s_i, s_{i+m/\mathfrak{d}}, \dots, s_{i+(\mathfrak{d}-1)m/\mathfrak{d}}$ with indices taken mod m , and then glue these distinct copies of x into unusually long segments (called *lips*) of what will become the spine Σ .

The lips partition the remainder of L into subsets b_i each with mass $O(n^{1-\delta})$, and we can extend the partial gluing along the lips into an (m, d) -regular simplicial or cubical spine by gluing intervals in each b_i independently. As in [2], § 5 and [3], § 4 the resulting spine Σ has the property that for any positive β , any immersed segment $\gamma \rightarrow \Sigma$ with length βn whose image in X lifts to r or r^{-1} , already lifts to L (i.e. it appears in the boundary of a disk of \overline{M}). From this and the fact that a random 1-relator group is $C'(\mu)$ for any positive μ it follows that $\overline{M} \rightarrow K_r$ is π_1 -injective and its image is quasiconvex.

Finally, the fact that the valence of Σ is uniformly bounded and the length of every segment is at least λ where we may choose λ as big as we like, implies that $\overline{M} \rightarrow K_r$ stays π_1 -injective and quasiconvex when we attach the disks corresponding to the remaining $(2k - 1)^{nD} - 1$ random relations. This depends on Ollivier’s mesoscopic small cancellation theory [5]; the argument is exactly the same as that in [2], § 6 and [3], § 6.

This completes the proof of the Superideal Simplex Theorem and Superideal Cube Theorem.

5. ACKNOWLEDGMENTS

Danny Calegari was supported by NSF grant DMS 1405466. The \LaTeX macro for the symbol \square was copied from `pmtmacros20.sty` by F. J. Ynduráin.

REFERENCES

- [1] I. Agol, *The virtual Haken conjecture*, With an appendix by Agol, Daniel Groves, and Jason Manning. Doc. Math. **18** (2013), 1045–1087
- [2] D. Calegari and A. Walker, *Random groups contain surface subgroups*, Jour. AMS to appear; available at arXiv:1304.2188
- [3] D. Calegari and H. Wilton, *3-manifolds everywhere*, preprint; arXiv:1404.7043
- [4] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric group theory, vol. 2; LMS Lecture Notes **182** (Niblo and Roller eds.) Cambridge University Press, Cambridge, 1993
- [5] Y. Ollivier, *Some small cancellation properties of random groups*, Internat. J. Algebra Comput. **17** (2007), no. 1, 37–51
- [6] Y. Ollivier and D. Wise, *Cubulating random groups at density less than 1/6*, Trans. AMS **363** (2011), no. 9, 4701–4733
- [7] D. Wise, *The structure of groups with a quasiconvex hierarchy*, preprint.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO, CHICAGO, ILLINOIS, 60637
E-mail address: dannyc@math.uchicago.edu