

RANDOM GROUPS CONTAIN SURFACE SUBGROUPS

DANNY CALEGARI AND ALDEN WALKER

1. INTRODUCTION

Gromov famously asked the following:

Surface Subgroup Question. *Let G be a one-ended hyperbolic group. Does G contain a subgroup isomorphic to the fundamental group of a closed surface with $\chi < 0$?*

Beyond its intrinsic appeal, and its obvious connections to the Virtual Haken Conjecture in 3-manifold topology (now a theorem of Agol [1]), one reason Gromov was interested in this question was the hope that such surface subgroups could be used as essential structural components of hyperbolic groups [9]. Our interest in this question is stimulated by a belief that surface groups (not necessarily closed) can act as a sort of “bridge” between hyperbolic geometry and symplectic geometry (through their connection to causal structures, quasimorphisms, stable commutator length, etc.).

Despite receiving considerable attention the Surface Subgroup Question is wide open in general, although in the specific case of hyperbolic 3-manifold groups it was positively resolved by Kahn and Markovic [10]. The main results of our paper may be summarized by saying that we show that Gromov’s question has a positive answer for *most* (hyperbolic) groups. In fact, the “executive summary” says that

- (1) *most groups* contain (many) surface subgroups;
- (2) these surface subgroups are *quasiconvex*—i.e., their intrinsic and extrinsic geometry is uniformly comparable on large scales; and
- (3) these surface subgroups can be *constructed*, and their properties *certified* quickly and easily.

Here “most groups” is a proxy for *random groups* in Gromov’s few relators or density models, to be defined presently.

In [8], Section 9 (also see [14]), Gromov introduced the notion of a *random group*. In fact, he introduced two such models: the *few relators* model and the *density* model. In either model one first fixes a free group F_k of rank $k \geq 2$ and a free generating set x_1, \dots, x_k , and adds ℓ random relators of some fixed length n . In one model ℓ is a constant, independent of n . In the other model $\ell = (2k - 1)^{Dn}$

Received by the editors April 5, 2013 and, in revised form, November 19, 2013, December 17, 2013, and January 27, 2014.

2010 *Mathematics Subject Classification*. Primary 20P05, 20F65, 57M07; Secondary 57M20.

©2014 American Mathematical Society
Reverts to public domain 28 years from publication

where now D is constant, independent of n . Explicitly:

Definition 1.0.1 (Few relators model). A random k -generator ℓ -relator group at length n is a group defined by a presentation

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

where the r_i are chosen randomly (with the uniform distribution) and independently from the set of all cyclically reduced cyclic words of length n in the x_i^\pm .

Definition 1.0.2 (Density model). A random k -generator group at density D (for some $0 < D < 1$) and at length n is a group defined by a presentation

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

where $\ell = (2k - 1)^{Dn}$, and where the r_i are chosen randomly (with the uniform distribution) and independently from the set of all cyclically reduced cyclic words of length n in the x_i^\pm .

Thus properly speaking, either model defines a *probability distribution* on finitely presented groups (in fact, on finite presentations) depending on constants k, ℓ, n in the few relators model, or on k, D, n in the density model.

If one is interested in a particular property of finitely presented groups, then one can compute for each n the probability that a random group as above has the desired property. If this probability goes to 1 as n goes to infinity, then one says that a random k -generator group (with ℓ relators; or at density D) has the given property *with overwhelming probability*.

Gromov showed that at any fixed density $D > 1/2$ random groups are trivial or isomorphic to $\mathbb{Z}/2\mathbb{Z}$, whereas at density $D < 1/2$ they are infinite, hyperbolic, and two-dimensional (with overwhelming probability), and in fact the “random presentation” determined as above is aspherical. Later, Dahmani, Guirardel, and Przytycki [7] showed that random groups at any density $D < 1/2$ are one-ended and do not split, and therefore (by the classification of boundaries of hyperbolic two-dimensional groups), have a Menger sponge as a boundary.

Random groups at density $D < 1/6$ are known to be cubulated (i.e., are equal to the fundamental groups of nonpositively curved compact cube complexes), and at density $D < 1/5$ to act cocompactly (but not necessarily properly) on a CAT(0) cube complex, by Ollivier and Wise [15]. On the other hand, groups at density $1/3 < D < 1/2$ have property (T), by Zuk [17] (further clarified by Kotowski and Kotowski [11]), and therefore cannot act on a CAT(0) cube complex without a global fixed point. A one-ended hyperbolic cubulated group contains a one-ended graph of free groups (see [6], Appendix A; this depends on the work of Agol [1]), and Calegari and Wilton [6] show that a *random* graph of free groups (i.e., a graph of free groups with random homomorphisms from edge groups to vertex groups) contains a surface subgroup. Thus one might hope that a random group at density $D < 1/6$ should contain a graph of free groups that is “random enough” so that the main theorem of [6] can be applied, and one can conclude that there is a surface subgroup.

Though suggestive, there does not appear to be an easy strategy to flesh out this idea. Nevertheless in this paper we are able to show directly that at *any* density $D < 1/2$ a random group contains a surface subgroup (in fact, many surface subgroups).

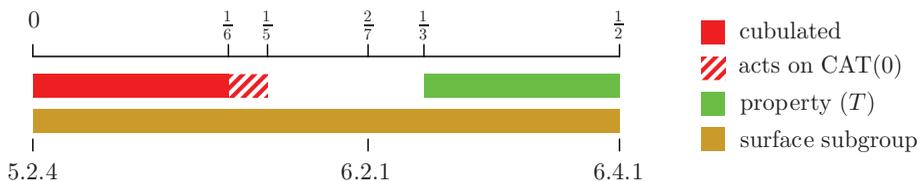


FIGURE 1. Random groups at different densities.

We give three proofs of this theorem, valid at different densities, with the final proof giving any density $D < 1/2$. Theorem 5.2.4 is valid for one-relator groups (informally $D = 0$), Lemma 6.2.1 gives $D < 2/7$, while our main Theorem 6.4.1 gives $D < 1/2$. Explicitly, we show the following.

Surfaces in Random Groups 6.4.1. *A random k -generator group at any density $D < 1/2$ and length n contains a surface subgroup with probability $1 - O(e^{-n^c})$. In fact, it contains $O(e^{n^c})$ surfaces of genus $O(n)$. Moreover, these surfaces are quasiconvex.*

This state of affairs is summarized in Figure 1. A modification of the construction (see Remark 6.4.3) shows that the surface subgroups can be taken to be homologically essential.

Along the way we prove some results of independent interest. The first of these (and the most technically involved part of the paper) is the *Thin Fatgraph Theorem*, which says that a “sufficiently random” homologically trivial collection of cyclic words Γ in a free group satisfies a strong combinatorial property: it can be realized as the oriented boundary of a trivalent fatgraph in which every edge is longer than some prescribed constant. This theorem is actually proved in a *relative* version, where after having realized a collection of subwords $\Gamma' \subset \Gamma$ as the oriented boundary of a *partial trivalent fatgraph* (i.e., a fatgraph with 3-valent interior vertices and 1-valent “boundary” vertices), the remainder $\Gamma'' := \Gamma - \Gamma'$ can be thought of as a collection of *tagged* cyclic words, where the tags indicate the boundary data (i.e., the way in which Γ'' lies inside Γ). Precise definitions of these terms are given in Section 3.1.

Thin Fatgraph Theorem 3.3.1. *For all $L > 0$, for any $T \gg L$ and any $0 < \epsilon \ll 1/T$, there is an N depending only on L so that if Γ is a homologically trivial collection of tagged loops such that for each loop γ in Γ :*

- (1) *no two tags in γ are closer than $4L$;*
- (2) *the density of the tags in γ is of order $o(\epsilon)$;*
- (3) *γ is (T, ϵ) pseudorandom;*

then there exists a trivalent fatgraph Y with every edge of length at least L so that $\partial S(Y)$ is equal to N disjoint copies of Γ .

If the rank of the group is 2, we can take $N = 1$ above; otherwise, we can take $N = 20L$.

The Thin Fatgraph Theorem strengthens one of the main technical theorems underpinning [5] and [6], and can be thought of as a kind of L^∞ theorem whose L^1 version (with optimal constants) is the main theorem of [3]. If r is a long random relator, the Thin Fatgraph Theorem lets us build a surface whose boundary consists

of a small number of copies of r and r^{-1} . By plugging in a disk along each boundary component, we obtain a closed surface in the one-relator group $\langle F_k \mid r \rangle$. If the surface is built correctly, it can be shown to be π_1 injective, with high probability. This is one of the most subtle parts of the construction, and ensuring that the surfaces we build are π_1 injective at this step depends on the existence of a so-called *Bead Decomposition* for r ; see Lemma 5.2.2. Thus we obtain the Random One Relator Theorem, whose statement is as follows.

Random One Relator Theorem 5.2.4. *Fix a free group F_k and let r be a random cyclically reduced word of length n . Then $G := \langle F_k \mid r \rangle$ contains a surface subgroup with probability $1 - O(e^{-n^c})$.*

The surfaces stay injective as more and more relators are added (in fact, these are the surfaces referred to in the main theorem) so this shows that random groups in the few relators model also contain surface subgroups for any fixed $\ell > 0$, with high probability.

There is an interesting tension here: the fewer relators, the harder it is to build a surface group, but the easier it is to show that it is injective. This suggests looking for surface subgroups in an arbitrary one-ended hyperbolic group at a very specific “intermediate” scale, perhaps at the scale $O(\delta)$ where δ is the constant of hyperbolicity with respect to an “efficient” (e.g., Dehn) presentation.

We conclude this introduction with three remarks.

First, it is worth spelling out some similarities and differences between our work and the breakthrough work of Kahn and Markovic [10]. The Kahn-Markovic argument depends crucially on the structure of hyperbolic 3-manifold groups as lattices in the semisimple Lie group $\mathrm{PSL}(2, \mathbb{C})$. By contrast, in this paper we are concerned with much more combinatorial classes of hyperbolic groups. Nevertheless, one common point of contact is the use of probability theory to *construct* surfaces, and the use of (hyperbolic) geometry to *certify* them as injective. In particular, because our surfaces are certified as injective by local methods, they end up being quasiconvex. It is an interesting question to identify the class of hyperbolic groups which contain non-quasiconvex (yet injective) surface subgroups (hyperbolic 3-manifold groups are now known to contain such groups since they are virtually fibered, again by Agol [1]).

Second, a large part of the difficulty in the proof of the Thin Fatgraph Theorem arises because we insist on building *oriented* surfaces. The advantage of this is that when our random groups G have nontrivial H_2 (which happens whenever $D > 0$ in the density model) the injective surfaces we construct can be chosen to be *homologically essential* in G . On the other hand, for the reader who is interested only in the existence of closed surface subgroups in G , the proof of the Thin Fatgraph Theorem can be considerably simplified. We explain this at the end of Section 4.

Third, the reader who is not already invested in the theory of random groups might complain that the few relators and density models seem rather special, insofar as the random relators are sampled from an especially simple probability distribution (i.e., the uniform distribution). One may consider a variation on the construction of a random group by fixing F_k and a stationary Markov process of entropy $\log(\lambda) > 0$ which successively generates the letters of reduced words in F_k , and define a random group at density D and length n to be obtained by adding λ^{nD} words of length n as relators, each generated independently by the Markov process.

Providing the Markov process is ergodic and has *full support*—i.e., providing that every finite reduced word σ has a positive probability of being generated—a random group in this model will contain surface groups with overwhelming probability for any $D < 1/2$. If we further assume that for a long random string generated by the Markov process and for any σ as above the expected number of copies of σ and of σ^{-1} are equal, then the surface subgroups can be chosen to be homologically essential.

2. BACKGROUND

In this section we describe some of the standard combinatorial language that we use in the remainder of the paper. Most important is the notion of *foldedness* for a map between graphs, as developed by Stallings [16]. We also recall some standard elements of the theory of small cancellation, which it is convenient to cite at certain points in our argument, though ultimately we depend on a more flexible version of small cancellation theory developed by Ollivier [13] specifically for application to random groups (his results are summarized in Section 6.1).

2.1. Fatgraphs and foldedness.

Definition 2.1.1. Let X and Y be graphs. A map $f : Y \rightarrow X$ is *simplicial* if it takes edges (linearly) to edges. It is *folded* if it is locally injective.

A folded map between graphs is injective on π_1 . The terminology of foldedness, and its first effective use as a tool in group theory, is due to Stallings [16].

Definition 2.1.2. A *fatgraph* is a graph Y together with a choice of cyclic order on the edges incident to each vertex. A fatgraph admits a canonical *fattening* to a surface $S(Y)$ in which it sits as a spine [so that $S(Y)$ deformation retracts to Y] in such a way that the cyclic order of edges coming from the fatgraph structure agrees with the cyclic order in which the edges appear in $S(Y)$. A *folded fatgraph over X* is a fatgraph Y together with a folded map $f : Y \rightarrow X$.

The case of most interest to us will be that X is a rose associated with a free generating set for a (finitely generated) free group F .

A folded fatgraph $f : Y \rightarrow X$ induces a π_1 injective map $S(Y) \rightarrow X$. The deformation retraction $S(Y) \rightarrow Y$ induces an immersion $\partial S(Y) \rightarrow Y$, and we may therefore think of $\partial S(Y)$ as a union of simplicial loops. Under f these loops map to immersed loops in X , corresponding to conjugacy classes in $\pi_1(X)$.

Conversely, given a homologically trivial collection of conjugacy classes Γ in $\pi_1(X)$ represented (uniquely) by immersed oriented loops in X , we may ask whether there is a folded fatgraph Y over X so that $\partial S(Y)$ represents Γ [by abuse of notation, we write $\partial S(Y) = \Gamma$]. Informally, we say that such a Γ *bounds a folded fatgraph*.

2.2. Small cancellation.

Definition 2.2.1. Let G have a presentation

$$G := \langle x_1, \dots, x_n \mid r_1, \dots, r_s \rangle$$

where the r_j are cyclically reduced words in the generators x_i^\pm . A *piece* is a subword that appears in two different ways in the relations or their inverses. A presentation satisfies the condition $C'(\lambda)$ for some λ if every piece σ in some r_i satisfies $|\sigma|/|r_i| < \lambda$.

Remark 2.2.2. Some authors use the notation $C'(\lambda)$ to indicate the weaker inequality $|\sigma|/|r_i| \leq \lambda$. This distinction will be irrelevant for us.

Associated with a presentation there is a connected 2-complex K with one vertex, one edge for each generator, and one disk for each relation. The 1-skeleton X for K is a rose for the free group on the generators. As is well known, a group satisfying $C'(1/6)$ is hyperbolic, and (if no relator is a proper power) the 2-complex K is aspherical (so that the group is of cohomological dimension at most 2).

Definition 2.2.3. Fix a group G with a presentation complex K and 1-skeleton X as above. A *surface over the presentation* is an oriented surface S with the structure of a cell complex together with a cellular map to K which is an isomorphism on each cell. The 1-skeleton Y of the CW complex structure on S inherits the structure of a fatgraph from S and its orientation, and this fatgraph comes together with a map to X . We say S has a *folded spine* if $Y \rightarrow X$ is a folded fatgraph.

If G is a small cancellation group, a surface with a folded spine can be certified as π_1 injective by the following combinatorial condition.

Definition 2.2.4. Let G be a group with a fixed presentation, and let S be an oriented surface over the presentation with a folded spine Y . We say S is α *convex* (for some $\alpha > 0$) with respect to the presentation if for every immersed path γ in Y which is a subword in some relation r_i^\pm with $|\gamma|/|r_i| \geq \alpha$, we actually have that γ is contained in $\partial S(Y)$ (i.e., it is in the boundary of a disk of S).

Lemma 2.2.5 (Injective surface). *Let G be a group with a presentation satisfying $C'(1/6)$ and such that no relator is a proper power; and let S be an oriented surface over the presentation with a folded spine Y . If S is $1/2$ convex, then it is π_1 injective.*

Proof. First we prove injectivity under the assumption that S is $1/2$ convex. Suppose not, so that there is some essential loop in $\pi_1(S)$ which is trivial in G . After a homotopy, we can assume this loop γ is *immersed* in Y . Since Y is folded, the image of γ in X is also immersed; i.e., it is represented by a cyclically reduced word in the generators. Since by hypothesis γ is trivial in G , there is a van Kampen diagram with γ as the boundary. We may choose γ and a diagram for which the number of faces is minimal.

The $C'(1/6)$ condition implies that there is a face D in the diagram which has at least $1/2$ of its boundary as a connected segment on γ ; this is sometimes called *Greedlinger's Lemma*. Then the hypothesis implies that this segment is actually contained in the boundary of a disk D' of S . Since G is $C'(1/6)$, it follows that D' and D bound the same relator in the same way, and we can therefore push γ across D' to obtain a van Kampen diagram with fewer faces and with the boundary an essential loop in S (homotopic to γ). But this contradicts the choice of the van Kampen diagram, and this contradiction shows that no such essential loop exists, i.e., that $\pi_1(S) \rightarrow G$ is injective. \square

Remark 2.2.6. If S is α convex for any fixed $\alpha < 1/2$, a similar argument shows that S is quasiconvex; since we shall prove quasiconvexity under more general geometric hypotheses in Theorem 6.4.1, and since this fact is not actually used in the paper, we do not justify this remark here.

In the sequel we usually say that a map from S to K is *injective* to mean that it is π_1 injective.

3. TRIVALENT FATGRAPHS

The purpose of this section is to prove the Thin Fatgraph Theorem 3.3.1, which implies that a (homologically trivial) collection of random cyclically reduced words bounds a trivalent fatgraph with long edges (i.e., in which *every* edge is as long as desired).

For concreteness the theorem is stated not for random words but for (sufficiently) pseudorandom words, and does not therefore really involve any probability theory. However, the (obvious) application in this paper is to random words and words obtained from them by simple operations.

3.1. Partial fatgraphs and tags. We are going to build folded fatgraphs with the prescribed boundary [i.e., given Γ we will build Y with $\Gamma = \partial S(Y)$]. In the process of building these fatgraphs we deal with intermediate objects that we call *partial fatgraphs* bounding part of Γ , and the part of Γ that is not yet bounded by a partial fatgraph is a collection of *cyclic words with tags*. This language is introduced in [5].

The (partial) fatgraphs will be built by taking disjoint pairs of segments in Γ with inverse labels (in X) and *pairing them*—i.e., associating them with opposite sides of an edge of the fatgraph. Once all of Γ is decomposed into such paired segments the fatgraph Y will be implicitly defined.

A *partial fatgraph* is, abstractly, the data of a pairing of some collection of disjoint pairs of segments in Γ . We imagine that this partial fatgraph Z has boundary $\partial S(Z) =: \Gamma'$ which is a subset of Γ . The difference $\Gamma - \Gamma'$ is a collection of paths whose end points are paired according to how they are paired in Γ' . The result is therefore a collection of cyclic words Γ'' , together with the data of the “germ” of the partial fatgraph Z at finitely many points. These extra data we refer to as *tags*, and we call this collection Γ'' a collection of *cyclic words with tags*.

Example 3.1.1. An example is illustrated in Figure 2. Starting with two reduced cyclic words $vzXwyZ$ and uxY we pair the subwords zX , xY and yZ along the edges of a tripod (as indicated in the figure) leaving “tagged” cyclic words $u \cdot w \cdot$ and $v \cdot$ as a remainder (in formulas the tags can be indicated by the punctuation character \cdot).

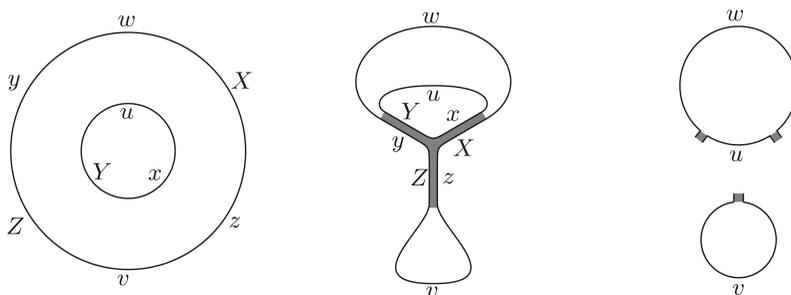


FIGURE 2. Two cyclic words are partially paired along a partial fatgraph (the grey tripod); what is left is two cyclic words with three tags.

3.2. Pseudorandomness. Random (cyclically reduced) words enjoy many strong equidistribution properties, at a large range of scales. For our purposes it is sufficient to have “enough” equidistribution at a sufficiently large *fixed* scale. To quantify this we describe the condition of *pseudorandomness*, and observe that random words are pseudorandom with high probability.

Definition 3.2.1. Let Γ be a cyclically reduced cyclic word in a free group F_k with $k \geq 2$ generators. We say Γ is (T, ϵ) *pseudorandom* if the following is true: if we pick any cyclic conjugate of Γ and write it as a product of reduced words w_i of length T (and at most one word v of length $< T$)

$$\Gamma := w_1 w_2 w_3 \cdots w_N v$$

then for every reduced word σ of length T in F_k , there is an estimate

$$1 - \epsilon \leq \frac{\#\{i \text{ such that } w_i = \sigma\}}{N} \cdot (2k)(2k - 1)^{T-1} \leq 1 + \epsilon$$

Similarly, we say that a collection of reduced words w_i of length T is ϵ *pseudorandom* if for every reduced word σ of length T in F_k the estimate above holds.

Lemma 3.2.2 (Random is pseudorandom). *Fix $T, \epsilon > 0$. Let Γ be a random cyclically reduced word of length n . Then Γ is (T, ϵ) pseudorandom with probability $1 - O(e^{-Cn})$.*

Proof. This is immediate from the Chernoff inequality for finite Markov chains (see Section 5.1 for a precise statement of the form of the Chernoff inequality we use, and for references). □

3.3. Thin Fatgraph Theorem. We now come to the main result in this section, the Thin Fatgraph Theorem. This says that any (sufficiently) pseudorandom homologically trivial collection of tagged loops, with sufficiently few and well-spaced tags, bounds a trivalent fatgraph with every edge as long as desired. Note that every trivalent graph (with reduced boundary) is automatically folded.

This theorem can be compared with [5], Theorem 8.9 which says that random homologically trivial words bound 4-valent folded fatgraphs, with high probability, and with [3], Theorem 4.1 which says that random homologically trivial words of length n bound (not necessarily folded) fatgraphs whose *average* valence is arbitrarily close to 3, and whose *average* edge length is as close to $\log(n)/2 \log(2k - 1)$ as desired (and moreover this quantity is sharp). It would be very interesting to prove (or disprove) that random homologically trivial words bound (with high probability) trivalent fatgraphs in which *every* edge has length $O(\log(n))$, but this seems to require new ideas.

Theorem 3.3.1 (Thin Fatgraph). *For all $L > 0$, for any $T \gg L$ and any $0 < \epsilon \ll 1/T$, there is an N depending only on L so that if Γ is a homologically trivial collection of tagged loops such that for each loop γ in Γ :*

- (1) *no two tags in γ are closer than $4L$;*
- (2) *the density of the tags in γ is of order $o(\epsilon)$;*
- (3) *γ is (T, ϵ) pseudorandom;*

then there exists a trivalent fatgraph Y with every edge of length at least L so that $\partial S(Y)$ is equal to N disjoint copies of Γ .

The notation $T \gg L$ means “for all T sufficiently large depending on L ,” and similarly $0 < \epsilon \ll 1/T$ means “for all ϵ sufficiently small depending on T .” The *density* of tags is just the number of tags divided by the length of γ , and the notation $o(\epsilon)$ just means something of negligible size compared to ϵ . The role of N will become apparent at the last step, where some combinatorial condition can be solved more easily over the rationals than over the integers (so that one needs to take a multiple of the original chain in order to clear denominators). In fact, in rank 2 we can actually take $N = 1$, and in a higher rank we can take $N = 20L$ (it is probably true that one can take $N = 1$ always, but this is superfluous for our purposes).

Except for the last step (which it must be admitted is quite substantial and takes up almost half the paper), the argument is very close to that in [5]. For the sake of completeness we reproduce that argument here, explaining how to modify it to control the edge lengths and valence of the fatgraph.

3.4. Experimental results. Theorem 3.3.1 asserts that long random words bound trivalent fatgraphs (up to taking sufficiently many disjoint copies). However, in order for the pseudorandomness to hold at scales required by the argument, it is necessary to consider random words of enormous length; i.e., on the order of a googol or more. On the other hand, experiments show that even words of modest length bound trivalent fatgraphs with high probability. To keep our experiment simple, we considered only the condition of bounding a trivalent graph, ignoring the question of whether the edges can all be chosen to be long.

In a free group of rank 3, we looked at between 100000 and 400000 cyclically reduced homologically trivial words of each even length from 10 to 120. The proportion of such words that bound trivalent fatgraphs is plotted in Figure 3. The vertical axis has a log scale to show some interesting features of the data. As one can see, bounding a trivalent fatgraph happens in practice for n far below the purview of Theorem 3.3.1. The curious local minimum at length ~ 50 is presumably a combinatorial artifact.

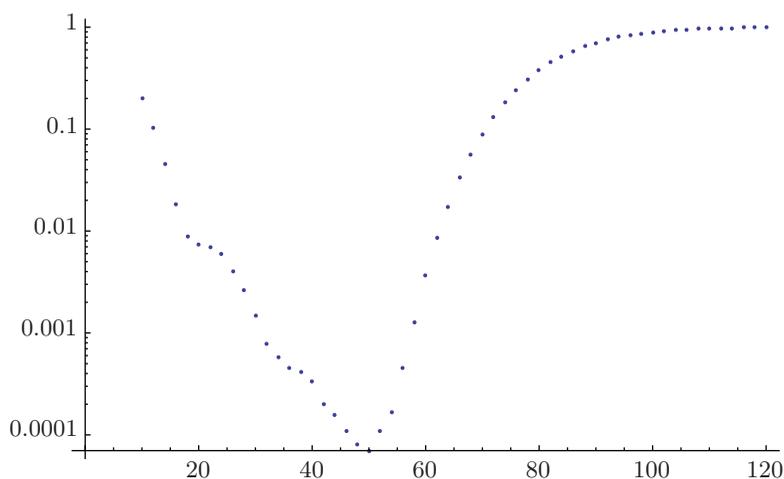


FIGURE 3. The experimental log fraction of random words of each length which bound a trivalent fatgraph.

3.5. Proof of the Thin Fatgraph Theorem. We now give the proof of Theorem 3.3.1. The proof proceeds in several steps. The first few steps are more probabilistic in nature. The last step is more combinatorial and quite intricate, and is deferred to Section 4.

Pick a γ in Γ . Now, γ is a cyclic word; starting at any letter we can express it in the form

$$\gamma = w_1 w_2 \cdots w_N v$$

where each $|w_i| = T$ and $N = \lfloor |\gamma|/T \rfloor$. Since γ is (T, ϵ) pseudorandom, the w_i are very well equidistributed among the reduced words of length T . Moreover, since by hypothesis (2) the density of tags is of order $o(\epsilon)$, the proportion of the w_i that contain a tag is also of order $o(\epsilon)$. In the next step we restrict attention to the w_i that do not contain a tag.

3.5.1. Tall poppies. Throughout the remainder of the proof we fix some T' which is an odd multiple of $10L$ with $1000L < T - T' \leq 2000L$ (in fact, something like $T - T' > 4L$ is sufficient, but there is no point in trying to optimize constants here). Note that we still have $T' \gg L$. For each w_i we let v_i be the initial subword of length T' . Note that the map which takes a reduced word of length T to its prefix of length T' takes the uniform measure to a multiple of the uniform measure, and therefore the v_i are also ϵ pseudorandom.

The first step is to create a collection of *tall poppies*. We fix some $v := v_i$ and read the letters one by one. As we read along, we look for a pair of inverse subwords x, X each of length $10L$ and separated by a subsegment y of length $40L$. Further we require that the copy of xyX should have the property that the x and X are maximal inverse subwords at their given locations, so that the result of pairing creates reduced tagged cyclic words. If the copy of xyX is not too close to a tag of γ (say, there is no tag within a $10L$ neighborhood), we create some partial fatgraph by identifying x to X ; this creates a tall poppy whose *stem* is x , and whose *flower* is y . Once we find and create a tall poppy, we look for each subsequent tall poppy at successive locations along v subject to the constraint that adjacent tall poppies are separated by subwords whose length is an even multiple of $10L$. Furthermore, we insist that the first tall poppy occurs at a distance an even multiple of $10L$ from the start of v . See Figure 4 for an example; the “dots” in the figure indicate units of $10L$.

For each v_i we fold off tall poppies as above. The result of this step is to create a partial fatgraph for each γ consisting of some tagged loop γ' (which is obtained from γ by cutting out all the xyX subwords and identifying end points) and a *reservoir* of flowers. Observe that every tagged cyclic word of length $40L$ occurs as a flower, and the set of tagged flowers is ϵ pseudorandom (conditioned on any compatible label on the tag). Note that as remarked above, we are only restricting attention to w_i that do not contain a tag of γ , so the operation of creating a tall poppy will never produce two tags that are too close together.

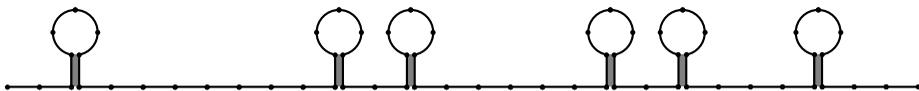


FIGURE 4. A word v of length $630L$ with 6 tall poppies folded off.

Informally, we say that the reservoir contains an *almost equidistributed collection* of tagged cyclic words of length $40L$. We can estimate the total number of flowers of each kind: at each location that a flower might occur, we require two subwords of length $10L$ to be inverse, which will happen with probability $(2k-1)^{-10L}$. The number of locations is roughly of size $O(|\gamma|/10L)$. So the number of copies of each tagged loop in the reservoir is of size $\delta \cdot |\gamma|$ (up to multiplicative error $1 \pm \epsilon$) for some specific positive $\delta > 0$ depending only on L .

3.5.2. Random cancellation. After cutting off tall poppies, the v_i become tagged words v'_i . Observe that the v'_i have variable lengths (differing from T' by an even multiple of $10L$) and have tags occurring at some subset of the points an even multiple of $10L$ from the start. The main observation to make is that the ϵ pseudorandomness of the v_i propagates to the ϵ pseudorandomness of the v'_i . That is, if σ is a reduced word of length $T' - m10L > 0$ for some even m , then among the v'_i of length $T' - m10L$, the proportion that is equal to σ is equal to $1/(2k)(2k-1)^{T'-m10L-1}$ up to a multiplicative error of size $1 \pm \epsilon$. This is immediate from the construction.

Recall that we chose T' to be an *odd* multiple of $10L$. This means that when we pair a segment v'_i labeled σ with a v'_j labeled σ^{-1} the tags of v'_i and v'_j do not match up, and in fact any two tags are no closer than distance $10L$. In fact, it is important that after pairing up inverse segments, the tagged loops that remain are reduced, so we write each v'_i in the form $l_i v''_i r_i$ where each of l_i, r_i has length $5L$, and pair v''_i with v''_j for some v''_j of the form $l_j v''_j r_j$ where $v''_j = (v''_i)^{-1}$, and the words $l_i r_j$ and $r_i l_j$ are reduced. By ϵ pseudorandomness, we can find such pairings of all but $O(\epsilon)$ of the v'_i in this way. Here, as in the previous subsection, we do *not* pair v'_i that contain one of the original tags of γ ; since the fraction of such v'_i is $o(\epsilon)$ [again by hypothesis (2)], the error term can be absorbed into the $O(\epsilon)$ term.

Thus the result of this pairing is to produce a trivalent partial fatgraph with all edges of length at least $5L$. Removing this from the v'_i produces a collection of tagged loops γ'' with $|\gamma''| = O(\epsilon \cdot |\gamma|)$.

3.5.3. Cancelling Γ'' from the reservoir. Let Γ'' be the union of all the γ'' , and pool the reservoirs from each γ into a single reservoir.

Notice that by construction, and by hypothesis (1) of the theorem, no tagged loop in Γ'' has two tags closer than distance $4L$ (in fact, it is only the original tags of γ which might be as close to each other as $4L$; the tags arising from tall poppies or by identifying the various v'_i in pairs will all be distance at least $20L$ apart).

For each tagged loop ν in Γ'' we can build a copy of ν^{-1} out of finitely many flowers in the reservoir in such a way that the result of pairing ν to this ν^{-1} is a trivalent partial fatgraph with all edges of length at least L , and the number of flowers that we need is proportional to $|\nu|/40L$. There is a slight subtlety here, in that the length of each flower is $40L$, and the result of partially gluing up a collection of cyclic words of even length always leaves an even number of letters unglued. Fortunately, the assumption that Γ is homologically trivial implies that $|\Gamma|$ itself is even, and since each flower also has an even number of letters, it follows that $|\Gamma''|$ is even. A flower with the cyclic word $xyzY$ can be partially glued to produce two tagged loops x and z , and if x and z are odd, each can be used to contribute to a copy of some ν^{-1} of odd total length. Since the number of odd $|\nu|$ is even, all of Γ'' can be cancelled in this way.

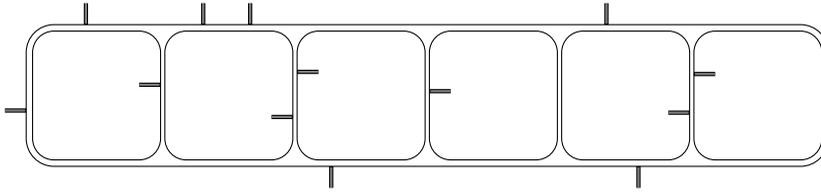


FIGURE 5. Cancelling a tagged loop ν of Γ'' by using flowers plus at most one loop of odd total length.

The construction of ν^{-1} from flowers plus at most one loop of odd total length, cancelling a tagged loop ν in Γ'' , is indicated in Figure 5. Each of the small loops in the figure has a length of order $40L$, and they are matched along segments roughly of order $10L$. Adjusting the length of the segments along which adjacent flowers are paired gives sufficient flexibility to build ν^{-1} (modulo the parity issue, which is addressed above). Notice that if ν contains a long string of tags, each a distance $\sim 4L$ from the next, we might need to attach two flowers near the midpoint between two adjacent tags, so that there might be some edges of length $2L$ in the trivalent partial fatgraph produced at this step. This is good enough to satisfy the conclusion of the theorem (with some room to spare).

Since $|\Gamma''| = O(\epsilon \cdot |\Gamma|)$, whereas the number of flowers of each kind in the reservoir is of order $\delta \cdot |\Gamma|$, if we take $\epsilon \ll \delta$ we can glue up all of Γ'' this way, at the cost of slightly adjusting the proportion of each kind of tagged loop in the reservoir.

3.5.4. *Gluing up the reservoir.* We are now left with an almost equidistributed collection of tagged loops of length $40L$ in the reservoir. Adding to the reservoir the contribution from each γ in Γ , and using the fact that Γ was homologically trivial, we see that the content of the reservoir is also homologically trivial. It remains to show that any such collection can be glued up to build a trivalent partial fatgraph with all edges of length at least L .

In fact, we only need two kinds of gluings to achieve this: gluings that result in partial fatgraphs that fatten to *annuli* and to *pants*. The argument is purely combinatorial, but quite intricate and involved, and makes up the content of Section 4.

Remark 3.5.1. At this point it is worth spelling out the modifications that need to be made to generalize the Thin Fatgraph Theorem to random chains generated by an ergodic stationary Markov process of full support, as discussed in the Introduction.

First, the definition of pseudorandom must be modified. Let's suppose that in our Markov model, the expected number of copies of a word σ in any sufficiently long string τ is $E(\sigma)|\tau|$ for some positive $E(\sigma)$. The correct definition of (T, ϵ) pseudorandomness of some word Γ in this context is that for any cyclic conjugate expressed in the form

$$\Gamma := w_1 w_2 \cdots w_N v$$

with each w_i of length T and at most one word v of length $< T$, for every reduced word σ of length T in F_k there is an estimate

$$1 - \epsilon \leq \frac{\#\{i \text{ such that } w_i = \sigma\}}{N} \cdot E(\sigma)^{-1} \leq 1 + \epsilon$$

Such pseudorandomness holds (with very high probability) for sufficiently long random words produced by the Markov process.

If one further assumes that $E(\sigma) = E(\sigma^{-1})$ for every σ , all steps of the argument above go through [the equality $E(\sigma) = E(\sigma^{-1})$ is used to ensure that after the random cancellation step the mass of the remainder is small compared to that of the reservoir] and we are left with a reservoir of loops, where the relative proportion of loops of kind σ and σ' is very close to $E(\sigma)/E(\sigma')$. Tagged loops with inverse labels σ and σ^{-1} for which the tags are not “too close” (under the orientation-reversing identification of σ with σ^{-1}) can be paired, and therefore we can reduce to the case of an almost equidistributed collection of tagged loops, at the cost of adjusting the constants.

If one does not assume that $E(\sigma) = E(\sigma^{-1})$ for every σ , the analogue of the Thin Fatgraph Theorem is *not* true on the nose. But for applications to the construction of surface subgroups by the method of Section 5 it is sufficient to apply the theorem to (subchains of) chains of the form $r \cup r^{-1}$ where r is a random relator; now the distribution of σ subwords in long segments of r very closely matches the distribution of σ^{-1} subwords in long segments of r^{-1} , and the construction goes through.

4. ANNULUS MOVES AND PANTS MOVES

In this section we show that an almost equidistributed collection of tagged loops of length $40L$ can be glued up to a trivalent partial fatgraph with all edges of length at least L . Together with the content of Section 3.5, this will conclude the proof of Theorem 3.3.1. The technical detail in this section is only necessary because we insist that our fatgraphs (surfaces) be orientable. There is a shortcut if we are willing to accept a nonorientable surface, explained in Section 4.4.

Remark 4.0.2. For the entirety of this section, we will rescale $40L$ to L . That is, we prove that an almost equidistributed collection of tagged loops of length L , where L is divisible by 4, can be glued up to a trivalent partial fatgraph with all edges of length at least $L/4$. This rescaling is intended to remove meaningless factors of 40 throughout the argument.

4.1. Pants and annuli. Let $S(L)$ be the set of tagged loops of length L , where L is divisible by 4. Let $W(L)$ be the vector space over \mathbb{Q} spanned by $S(L)$; that is, $W(L) = \mathbb{Q}[S(L)]$. We define $h : W(L) \rightarrow \mathbb{Z}^k$ to be the linear map so that $h(v)$ is the homology class of v . Finally, $V(L) = \ker h \subseteq W(L)$ is the vector space of homological trivial vectors. We are interested only in $V(L)$, not $W(L)$, so by “full dimensional,” we mean a full dimensional subset of $V(L)$. When we say that a vector *projectively bounds* a fatgraph, we mean that there is some multiple of the vector which has integer coordinates, and the collection of loops represented by the integral vector bounds a fatgraph. A (necessarily integral) vector *bounds* a fatgraph if the collection of loops that it represents bounds a fatgraph. The uniform vector of all 1’s will be of particular interest, and we denote it by $\mathbf{1}$.

We say that a fatgraph Y with the boundary a collection of loops in $S(L)$ is *thin* if Y is trivalent and the trivalent vertices of Y are pairwise distance at least $L/4$ apart, where the tags are counted as trivalent vertices. Let $C(L)$ be the subset of $V(L)$ of positive vectors which projectively bound a thin fatgraph. If $v, w \in C(L)$, then the disjoint union of the thin fatgraphs for v and w gives a thin fatgraph for

$v+w$. Also, the definition of $C(L)$ shows it to be closed under scalar multiplication. Hence, $C(L)$ is a cone. A variant of the `scallop` [4] algorithm gives an explicit hyperplane description of $C(L)$ and shows that it is a finite sided polyhedral cone, but we won't need this fact in the sequel.

We will build thin fatgraphs out of two kinds of pieces: (good pairs of) *pants* and (good) *annuli* (the terminology is supposed to suggest an affinity with the Kahn-Markovic proof of the Ehrenpreis conjecture, but one should not make too much of this). A good pair of pants is one whose edge lengths are all exactly $L/2$ and whose tags are each on different edges and exactly distance $L/4$ from the real trivalent vertices. Note the boundary of each such pair of pants lies in $V(L)$. A good annulus is a fatgraph annulus with boundary in $S(L)$ whose tags are distance at least $L/4$ apart. Hereafter, all pants and annuli are good.

Define an involution $\iota : S(L) \rightarrow S(L)$ which takes each loop to its inverse with the tag moved to the diametrically opposite position. There are several options for the tag at each position—for the definition of ι , we arbitrarily choose any pairing of the options to obtain an involution. There is a special class of annuli, which we call ι annuli, which have a boundary of the form $s + \iota(s)$. Notice that the collection of all ι annuli is a thin fatgraph which bounds the uniform vector $\mathbf{1}$.

The bulk of our upcoming work lies in manipulating *untagged* loops, and our result here is independently interesting, so we will need some complementary definitions. Let $S'(L)$ be the set of untagged loops of length L , let $W'(L) = \mathbb{Q}[S'(L)]$, and let $V'(L)$ be the vector space of homologically trivial vectors in $W'(L)$. We define a thin fatgraph and the uniform vector $\mathbf{1}' \in V'(L)$ as before. The set $C'(L) \subseteq V'(L)$ is the cone of vectors in $V'(L)$ which projectively bound thin fatgraphs. An untagged good pair of pants is a trivalent pair of pants whose edge lengths are exactly $L/2$, and an untagged annulus is simply an annulus whose boundary is two loops of length L . For untagged loops, $\iota : S' \rightarrow S'$ is simply inversion, and all annuli are ι annuli, although we may refer to them explicitly as ι annuli to emphasize their purpose.

For many applications, the property of a collection of loops that it *projectively* bounds a thin fatgraph is good enough (see, e.g., [5]), and this is in many ways a more pleasant property to work with, since the set of vectors (representing collections of loops) which projectively bound a thin fatgraph is a cone, whereas the set of vectors that *bound* (i.e., without resorting to taking multiples) is the intersection of this cone with an integer lattice. However, in this paper it is important to distinguish between “bounding” and “projectively bounding,” and therefore in the following propositions, we give both the stronger, technical “integral” statement and the weaker, cleaner “rational” one.

Proposition 4.1.1. *For any integral vector $v \in V'(L)$, there is $n \in \mathbb{N}$ so that $(L/2)v + n\mathbf{1}'$ bounds a collection of good pants and annuli. Consequently, $C'(L)$ is full dimensional and contains an open projective neighborhood of $\mathbf{1}'$.*

There is a stronger version without the $L/2$ factor if the free group has rank 2.

Proposition 4.1.2. *If the free group has rank 2, then for any integral vector $v \in V'(L)$, there is $n \in \mathbb{N}$ so that $v + n\mathbf{1}'$ bounds a collection of good pants and annuli.*

We believe that Proposition 4.1.2 is probably true for a higher rank, but the proof would be more complicated than we wish for a detail that we do not need.

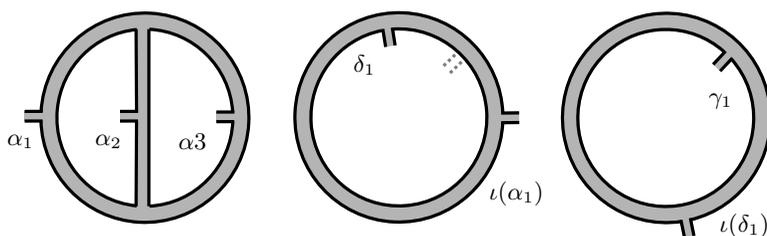


FIGURE 6. Adding annuli to fix an incorrectly tagged boundary. We want the loop γ_1 ; we have the loop α_1 , with the tag in the wrong place. The dotted gray tag indicates the proximity of the tags on $\iota(\alpha_1)$ and γ_1 . If it were farther away, we'd be done; however, it's too close, so we introduce δ_1 so that the distance between the tags on $\iota(\alpha_1)$ and δ_1 is at least $L/4$, and the distance between the tags on $\iota(\delta_1)$ and γ_1 is also at least $L/4$. The result is the desired boundary γ_1 , plus the ι pairs $\alpha_1 + \iota(\alpha_1) + \delta_1 + \iota(\delta_1)$. This procedure is repeated for α_2 and α_3 .

We delay the rather tedious proof of Proposition 4.1.1 in favor of stating the tagged version, which is a corollary and is the version we need.

Proposition 4.1.3. *For any integral vector $v \in V(L)$, there is $n \in \mathbb{N}$ so that $(L/2)v + n\mathbf{1}$ bounds a collection of good pants and annuli. Consequently, $C(L)$ is full dimensional and contains an open projective neighborhood of $\mathbf{1}$.*

Proof. Let us be given an integral vector $v \in V(L)$. Define $f : V(L) \rightarrow V'(L)$ to be the map $S(L) \rightarrow S'(L)$ which forgets the tag, extended by linearity. By Proposition 4.1.1, we can find a collection of pants and annuli which has boundary $(L/2)f(v) + m\mathbf{1}'$. Call this fatgraph Y' . Now place arbitrary tags in the forced positions on the pants (in the middle of the edges) and in allowed positions on the annuli (at least $L/4$ apart) to obtain Y . Clearly, Y is a thin fatgraph, and Y almost has boundary $(L/2)v + m\mathbf{1}$, as desired, but the tags are in the wrong places. We will fix this by simply adding annuli which “twist” the tags into the right positions.

Given some pair of pants in Y with boundary $\alpha_1 + \alpha_2 + \alpha_3$, let us focus on “twisting” the tag on α_1 . The loop α_1 corresponds to a loop γ_1 in $(L/2)v$, and the only difference is that the tag on α_1 is in a different position from the tag on γ_1 . There are two cases. If the tags on $\iota(\alpha_1)$ and γ_1 are at least $L/4$ apart, then we simply add an annulus with boundary $\iota(\alpha_1) + \gamma_1$. If the tags are closer than $L/4$, then we need two annuli: one with boundary $\iota(\alpha_1) + \delta_1$, and another with boundary $\iota(\delta_1) + \gamma_1$, where here δ_1 is the same loop as α_1 and γ_1 with the tag shifted so that the tags on δ_1 and $\iota(\alpha_1)$ are at least $L/4$ apart, and similarly for $\iota(\delta_1)$ and γ_1 . The result is that we have added annuli to resolve the boundary from α_1 to either $\alpha_1 + \iota(\alpha_1) + \gamma_1$ or $\alpha_1 + \iota(\alpha_1) + \delta_1 + \iota(\delta_1) + \gamma_1$; that is, ι pairs plus the desired tagged loop γ_1 . Figure 6 shows this operation.

After twisting all the tags in this manner, we are left with a collection of pants and annuli with boundary $(L/2)v + \Delta + \iota(\Delta) + A + \iota(A)$, where Δ and A are the α_i

and δ_i loops used to twist the tags. By adding ι annuli, we can make the boundary of Y be $(L/2)v + n\mathbf{1}$ for some $n \geq m$, as desired. \square

It remains to prove Proposition 4.1.1, which we now do.

4.2. Proof of Proposition 4.1.1. Let us be given an integral vector $v' \in V'(L)$. The vector v' represents a collection of loops $s' \in S'(L)$, for which we must find a thin fatgraph Y of the desired form. Our goal is to build a fatgraph which bounds $s' + t + \iota(t)$ for some $t \in S'(L)$. Adding ι annuli will then immediately finish the construction.

If we have a collection of annuli and pants which has boundary $s' + t$, then the problem reduces to finding a collection of annuli and pants which has a boundary of the form $\iota(t) + u + \iota(u)$ for some $u \in S'(L)$. We'll repeatedly apply this idea to simplify the problem by *attaching pants*. If we want to have a boundary which contains a loop γ , and we find a pair of pants with boundary $\gamma + \alpha + \alpha'$, then now we need only find a boundary containing $\iota(\alpha) + \iota(\alpha')$. In this case, we'll say that γ and $\iota(\alpha) + \iota(\alpha')$ are *pants equivalent*.

For this entire section, we will assume that our free group has rank 2 and is generated by a and b . In Section 4.3 we explain the extra details required to deal with free groups of higher rank.

Our strategy will be to start with s' and attach (many) pairs of pants which put all the loops in s' into a nice form. Then we attach more pants to further simplify the loops, and so on, eventually reducing to a case that is simple enough to handle by hand. A *run* in a loop is a maximal subword of the form a^p or b^p for some integer power $p \neq 0$. Note that any loop contains an even number of runs. We first reduce to the case that every loop has at most 4 runs, then to the case that every loop has 4 runs in a nice arrangement, then 2 runs, which we address directly.

For clarity, we separate the simplification into lemmas.

Lemma 4.2.1. *Any loop is pants equivalent to a collection of loops with at most 4 runs.*

Proof. To begin, we show how to attach a pair of pants to a loop which produces two loops, each of which has fewer runs than the initial loop. This method works whenever the number of runs is more than 4, so it reduces the loops in s' to a collection of loops with at most 4 runs. Let us be given a loop γ . The easiest way to visualize attaching a pair of pants is to simply draw a diameter d on γ between two antipodal vertices. Labeling the diameter d produces a pair of pants attached to γ . Note that we must be careful to label d compatibly with the labels adjacent to the vertices to which we attach d , so that the vertices do not fold. See Figure 7.

For concreteness, let us number the vertices of γ by $0, \dots, L - 1$, and we let d_i be the (oriented) diameter with initial vertex i [and thus terminal vertex $(i + L/2) \bmod L$]. Each diameter d_i divides γ into two pieces. Let x_i be the number of runs in the noncyclic subword of γ starting at index i and of length $L/2$; that is, the number of runs in the word to the “right” of d_i . Similarly, let y_i be the number of runs to the left. See Figure 8. Let r be the number of runs in γ . Note that $x_i + y_i$ may be greater than r . Specifically, $r \leq x_i + y_i \leq r + 2$. The important feature of these numbers is that $|x_i - x_{i+1}| \leq 1$ and $|y_i - y_{i+1}| \leq 1$. This is easily seen by considering the combinatorial possibilities that occur as we rotate the starting point i around the loop γ . We are particularly interested in *matched runs*, which are runs separated in either direction by the same number of other runs. That is,

matched runs are “directly across” from one another in the list of runs (we use scare quotes to emphasize that matched runs are not antipodal in the same sense that “antipodal vertices” are).

The “functions” x_i and y_i can be interpolated to piecewise linear functions on the circle; by applying the intermediate value theorem to these interpolations, we deduce that there is some point at which the interpolated graphs intersect. This can happen either at some value of i , in which case $x_i = y_i$, or between two values i and $i+1$, but by the discussion above, in this latter case $|x_i - y_i| \leq 1$ and $|x_{i+1} - y_{i+1}| \leq 1$. In either case, d_i must intersect two matched runs R and R' , perhaps on the boundaries of the runs. Now decrease i until one of the ends of d_i lies on the boundary of R or R' . This puts d_i into one of two combinatorial configurations, up to rotation and symmetry. See Figure 9. Note that the configuration on the right cannot occur at the intersection point, since $|x_i - y_i| = 2$.

We handle the two cases separately. First, we show the more generic case illustrated in Figure 9 on the left. Here we label the diameter entirely with the generator which is *not* the one labeling the bottom run, and in such a way as to minimize the number of runs in the resulting two loops. Figure 10, left, illustrates this. The sign of the labels on the diameter depends on the signs and orders of the generators around the end points, but the picture is equivalent. Notice that the number of runs in each of the resulting loops is at most $r/2 + 2$.

In the nongeneric case illustrated in Figure 9 in the middle, we label half of the diameter with one generator and the other half with the other, in a way which is compatible with the top and bottom labels. See Figure 10, right. In certain cases, it is possible to label the entire diameter with a single generator, and this reduces the number of runs still further, but we have illustrated the worst situation. We therefore compute again that the number of runs in each of the resulting loops is $r/2 + 2$.

As long as $r/2 + 2 < r$, or $r > 4$, this will produce two loops of strictly smaller length. Repeatedly attaching pants resolves our collection s' into a new collection of loops, all of which have at most 4 runs. \square

We have shown that an arbitrary collection of loops is pants equivalent to a collection of loops with at most 4 runs. To further reduce this to 2 runs, we first need to make the 4-run loops balanced. A 4-run loop is *balanced* if there is a pair of diameters d_a and d_b at right angles (with all end points spaced exactly $L/4$ apart) such that d_a starts touching one a run and ends touching the other, and similarly

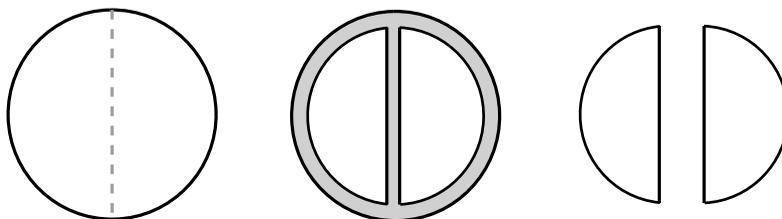


FIGURE 7. Attaching a diameter to a loop to form a pair of pants. The pair of loops on the right is pants equivalent to the loop on the left.

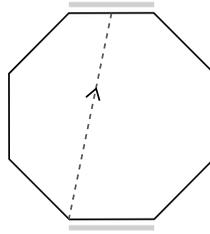


FIGURE 8. Moving the diameter one position changes x_i and y_i by at most 1. For the marked diameter, we have $x_i = 5$ and $y_i = 4$. A pair of matched runs is marked in grey.

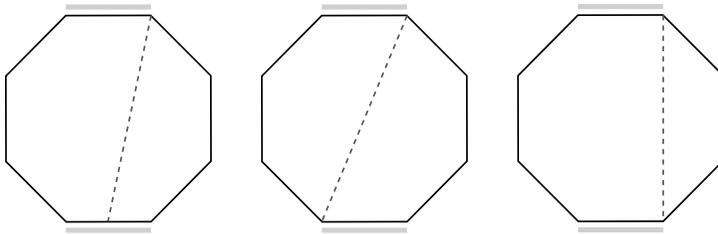


FIGURE 9. Possible configurations of d_i with respect to the matched runs R and R' . Up to rotation and symmetry, there are two. Note the configuration on the right cannot occur.

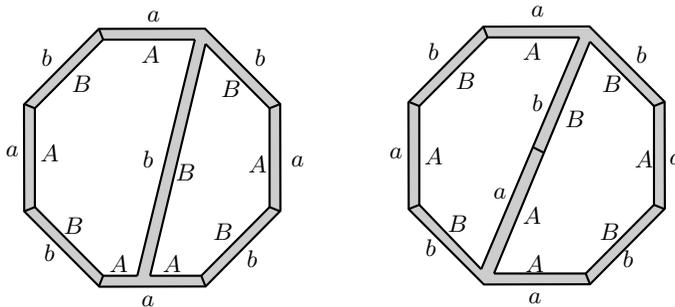


FIGURE 10. Labeling the diameter to reduce the number of runs. Each label represents a potentially long run of that generator.

for d_b with the b runs. Here *touching* a run means that the vertex on which the diameter starts or ends lies between two letters, at least one of which lies within the run. See Figure 11.

Lemma 4.2.2. *Any 4-run loop is pants equivalent to a collection of balanced loops and 2-run loops.*

Proof. Suppose we are given a 4-run loop. Without loss of generality, let us suppose there are at least as many a 's as b 's, and let $G = \#a - \#b$ be the *generator inequity*, recording how many more a 's than b 's there are. Let x and x' be the number of

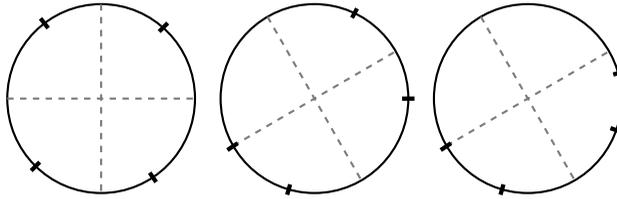


FIGURE 11. Examples of two balanced loops (left) and an unbalanced loop (right).

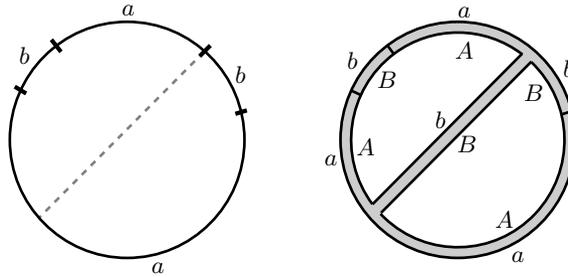


FIGURE 12. If the b runs are short, we can cut to reduce the generator inequity.

b 's in the longer and shorter b runs, respectively. Abusing notation, we'll also refer to the runs themselves as x and x' . Note that the a or b runs may have negative exponents, so they are actually runs of A or B . For simplicity, we'll use the "positive" notation.

First, let us eliminate the case that there are very few b 's. Suppose that $x + 2x' < L/2$. Consider the two diameters starting at the ends of x . There are two cases: if $x = x'$ and the runs are exactly antipodal, then drop a diameter between the middles of x and x' ; the diameter at right angles will touch both a runs, and the loop will be balanced, as desired. Otherwise, one of the diameters misses x' entirely, and we can cut to produce a 2-run loop and a loop whose generator inequity is strictly smaller [the roles of a and b are reversed, and the inequity becomes $2x' < 2(L/2 - x - x') = G$, i.e., smaller than the current inequity]. See Figure 12.

After repeatedly reducing the generator inequity, we may assume that $x + 2x' \geq L/2$. We conclude that $x \geq L/6$ and $x' \geq L/4 - x/2$, so $L/3 \leq x + x'$, and $G \leq L/3$. Our loop is now less degenerate, but it still might not be balanced. We'd like for x to be positioned opposite x' , as shown in Figure 11, left, so that we can simply draw two diameters and be done. However, the loop might look like Figure 11, right, in which x and x' are too close. To remedy this, we introduce the *triangle move at x* technique.

Algebraically, a triangle move at x takes in a word $a^{e_1} b^x a^{e_2} b^{x'}$ (for clarity, assume all the exponents are positive) such that $x + e_2 > L/2$ and $x < L/2$ and builds a pair of pants with boundary

$$a^{e_1} b^x a^{e_2} b^{x'} + A^{e_1 + L/2 - x} B^x A^{e_2 - (L/2 - x)} B^{x'} + a^{L/2 - x} b^x A^{L/2 - x} B^x$$

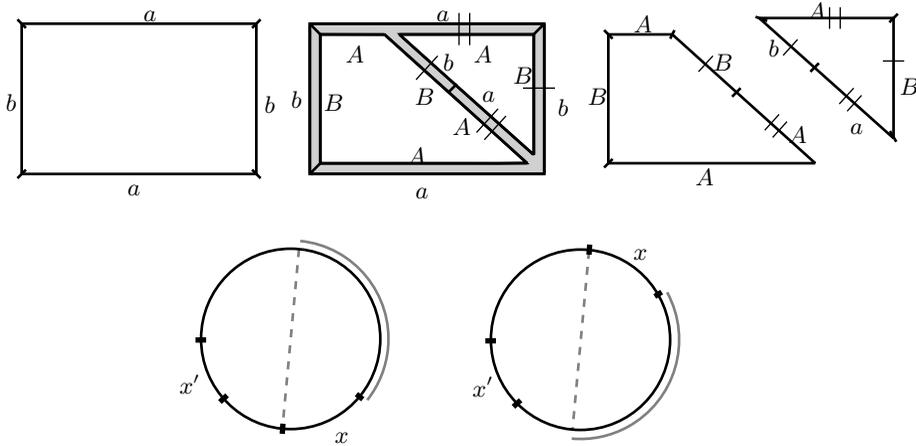


FIGURE 13. A triangle move shown explicitly, top, and schematically, bottom. The hash marks in the top picture indicate segments of equal length. Note the output is a balanced loop (opposite runs have exactly the same length), plus a loop in which $L/2 - x$ of the a 's on the top run have been shifted down. The schematic picture shows the effective result: a triangle move at x on the left loop produces a balanced loop (not shown), plus the loop on the right; note the runs have been shifted so that the right loop is now balanced.

The notation obscures the function of a triangle move, which is shown in Figure 13 and is as follows: it produces one balanced loop (opposite runs have the same length), and another loop with the same run sizes as the original one, except that $L/2 - x$ of the a 's in the top run have been shifted down to the bottom run. The signs of the generators may change as they shift, but we are only concerned with the lengths at this point. This is the critical feature of the triangle moves—shifting generators from one run to the other without disturbing anything else.

In the Figure 13 schematic, we are able to perform a triangle move at x in such a way that the loop becomes balanced. We will show that this can always be done. We need two things to happen simultaneously: x and x' must be opposite enough so that there is a diameter between them, and the diameter at right angles must also touch the a runs.

At this point, we split the argument into two cases. First, assume that $x \leq L/4$. In this case, x and x' are short enough that if we can find a diameter which touches both x and x' , then the diameter at right angles automatically touches the a runs, so the loop will be balanced. Consider the run x' : it casts a “shadow” directly opposite it so that if any part of x touches the shadow, we succeed in placing a diameter between the runs. Looking at the initial end point of x , we must place this end point in the *target region*, which we define to be the segment of size $x + x'$ ending directly antipodal to the end point of x' . See Figure 14.

Let t be the size of the target region, so $t = x + x'$. Doing a triangle move at x moves the initial end point of x by the *shift size*, which we denote s . Recall that $s = L/2 - x$. If we can show that $s \leq t$, then obviously we can shift x until it lies within the target region. Also recall that we reduced the generator inequity, so we

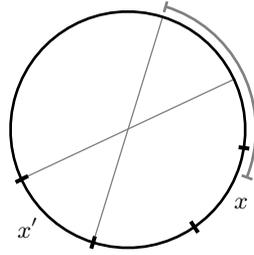


FIGURE 14. The run x' casts a shadow directly opposite itself. In the case that $x \leq L/4$, if we succeed in shifting x so it touches this shadow, the loop will be balanced. In other words, we need to place the initial end point of x inside the target region, shown in gray on the outside of the loop.

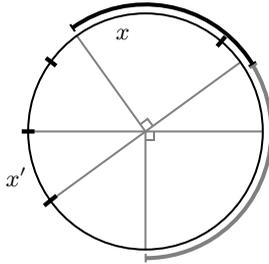


FIGURE 15. In the case that $x \geq L/4$, we need to be careful about the size of the target region (shown in gray). The initial end point of x can be placed anywhere in the target region, which starts exactly $L/4$ after the initial end point of x' and ends exactly $L/4+x$ before the final end point of x' . A segment of length x is shown adjacent to the target region for illustrative purposes.

have $x \geq L/6$ and $x + x' \geq L/3$. Putting these together, we have

$$s = \frac{L}{2} - x \leq \frac{L}{3} \leq x + x' = t$$

Therefore we do indeed have $s \leq t$, and we can balance the loop. This finishes the case that $x \leq L/4$.

For the other case, assume $x \geq L/4$. We will use the same technique, shifting x until the loop is balanced. Here we must be careful: if $x \geq L/4$, it is no longer obvious that there exists a diameter at right angles which exhibits the loop as balanced, so we must take this into account when setting the size of the target region. In this case, the target region starts exactly $L/4$ after the initial end point of x' and ends exactly $L/4 + x$ before the final end point of x' . Figure 15 shows the target region.

Computing the size, we find in this case that $t = L/2 + x' - x$. Again, the shift size is $s = L/2 - x$, so we immediately get

$$s = \frac{L}{2} - x < \frac{L}{2} - x + x' = t$$

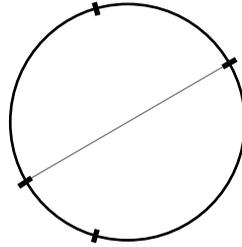


FIGURE 16. This loop has a diameter between corners.

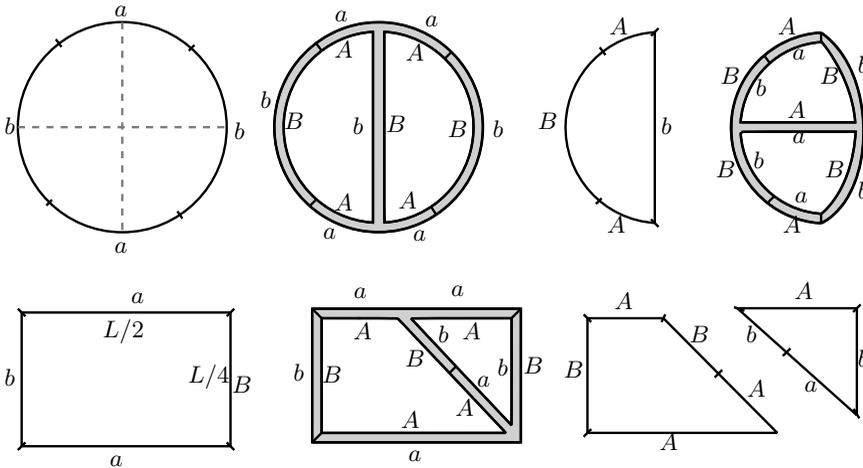


FIGURE 17. Applying the moves described in Lemma 4.2.3 to reduce a balanced loop to a loop with a diameter between corners. Lengths are not to scale, and some lengths are labeled. The operation is read left to right, top to bottom.

This completes the proof in the case that $x \geq L/4$ and thus the proof of the lemma. \square

At this point, we are left with a collection of balanced 4-run loops. For reducing the loop, even nicer than balanced is a loop with a *diameter between corners*. A diameter between corners is a diameter which starts between an a and a b run and ends between the other a and other b runs. See Figure 16.

Lemma 4.2.3. *Any balanced loop is pants equivalent to a collection of 2-run loops and loops with a diameter between corners.*

Proof. The proof of this lemma is essentially contained in the moves shown in Figure 17. Given a balanced loop, there exist diameters at right angles with ends in opposite runs. First, attach a pair of pants along the diameter between the a runs by labeling the diameter entirely with b 's. If one of the ends of this diameter touches a b run, it is necessary to be careful about the labels to ensure there is no folding. If both ends touch b runs, the loop already has a diameter between corners and we are done.

The result of attaching this pair of pants is two new loops. They have the same pattern, so we'll focus on one of them. It must be of the form (assuming positive

exponents) $b^{L/2}a^{e_1}b^{e_2}a^{e_3}$, and furthermore, because the original loop was balanced, the diameter starting exactly in the middle of the $b^{L/2}$ run must touch the other b run. Attach a pair of pants by labeling this diameter entirely with a . If the diameter touches one of the a runs, as always, we must label it to avoid folding. This results in two new loops, which either have two runs, or have the form $b^{L/4}a^{L/2}b^{e_1}a^{e_2}$ (again assuming positive exponents and new exponent variables e_1 and e_2).

Again these two loops have the same pattern, so we focus on one of them. The final step is to do a triangle move at the $b^{L/4}$ run. This shifts exactly $L/4$ of the a 's between runs, and results in a loop of the form $b^{L/4}a^{L/4}b^{e_1}a^{e_2+L/4}$. Observe this has a diameter between corners. The triangle move also produces a by-product loop; we did not stress this earlier because we did not need it, but the by-product loop has opposite runs of exactly the same length, so it has a diameter between corners. This completes the proof. \square

Lemma 4.2.4. *Any loop with a diameter between corners is pants equivalent to a collection of 2-run loops.*

Proof. This step requires at most two triangle moves, and again, is essentially described by a picture, which is shown in Figure 18.

Let us assume without loss of generality that our loop is of the form $b^{e_1}a^{e_2}b^{e_3}a^{e_4}$, where $|e_1| + |e_2| = |e_3| + |e_4| = L/2$. At this stage, we need to differentiate between positive and negative exponents. Suppose that e_2 and e_4 (the a runs) have the same sign. Then doing a triangle move at the run b^{e_1} produces a 2-run loop and the loop $a^{e_2}b^{\pm e_1}a^{-e_2}b^{-e_1}$. When following Figure 18, one must remember that the loop is on the inside of the pants, so the orientation is backward.

Therefore, $b^{e_1}a^{e_2}b^{e_3}a^{e_4}$ is pants equivalent to (a 2-run loop and) the inverse of this new loop, i.e., $b^{e_1}a^{e_2}b^{\mp e_1}a^{-e_2}$. Observe that this loop still has a diameter between corners, but now the signs on the a runs are different.

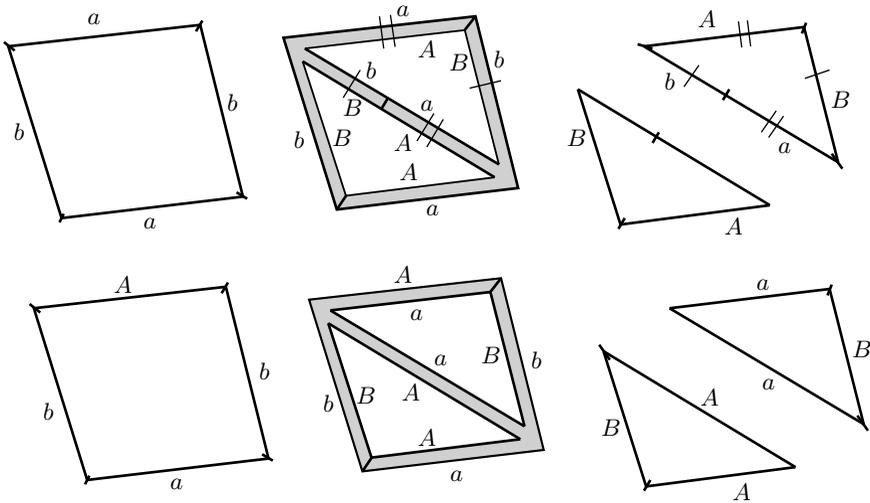


FIGURE 18. Performing triangle moves on loops with diameters between corners in order to produce 2-run loops. There are two possibilities, depending on whether an inverse pair appears.

We have reduced to a loop of the form $b^{e_1} a^{e_2} b^{e_3} a^{e_4}$, where $|e_1| + |e_2| = |e_3| + |e_4| = L/2$, and where the signs of e_2 and e_4 are different. In this case, we may attach a pair of pants by labelling the diameter entirely by a 's. This produces two 2-run loops, and completes the proof. \square

We are now left with a collection of only 2-run loops. For the next step, we will modify our collection of loops with 2 runs to put them in a standard form. A *uniform* loop has a single run, so just one generator appears. An *even* loop has two runs of the same length (so length $L/2$). The *type* of a loop with two runs is a pair that records which generators appear, so, for example, (a, B) .

Lemma 4.2.5. *Any collection of loops with 2 runs is pants equivalent to a collection of uniform loops, even loops, and at most one loop of each type.*

Proof. This step involves shifting and combining loops into even and uniform loops, which will leave a finite remainder. All of these operations are performed on loops of the same type. The first step is to arbitrarily select a generator to be the *small* generator on each loop. We'll choose b . Any loop whose b run has a length over $L/2$ can be cut with a diameter labeled with just a to produce an even loop and a loop with b run length less than $L/2$, as shown in Figure 19. An important feature of cutting to reduce the size of the b run is that it doesn't change the type of the loop. We have been somewhat casual about this thus far, because it makes the pictures easier to understand, but recall from the Introduction to this section that if we have a loop γ , and we find a pair of pants with boundary $\gamma + \alpha + \alpha'$, then the remaining problem is to find a collection of pants with boundary $\iota(\alpha) + \iota(\alpha')$; that is, recall that γ is not pants equivalent to $\alpha + \alpha'$, but rather is pants equivalent to $\iota(\alpha) + \iota(\alpha')$. Therefore, as shown in Figure 19, the input is the loop of type (a, b) , and the output is an even loop and another loop of type (a, b) . The same holds true for the other operations we describe here.

We aren't concerned with the even loops, so we turn our attention to the loops with b run length strictly less than $L/2$. There are two necessary operations here; the *trade*, in which we swap pieces of the b run between loops, and the *combine*, in which we combine two small b loops into a single one (and produce a uniform loop and two even loops as by-products). The combine operation works on any two loops whose total number of b 's is strictly less than $L/2$. These operations are shown in Figures 20 and 21. Algebraically, the trade operation takes in $a^{p_1} b^{t_1+t_2}$

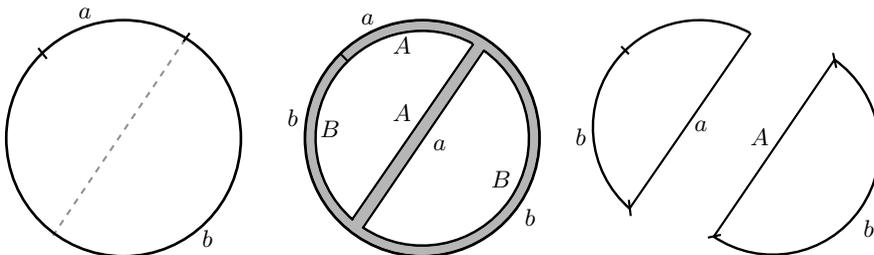


FIGURE 19. If a loop has a b run length larger than $L/2$, it can be cut to produce an even loop and a loop with a b run shorter than $L/2$. See the text for a discussion of why the inner boundary components become inverse on the right.

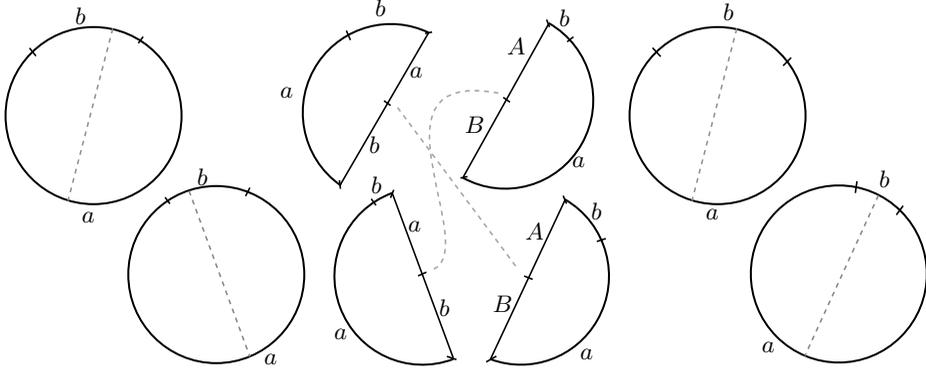


FIGURE 20. Trading pieces of b runs between loops.

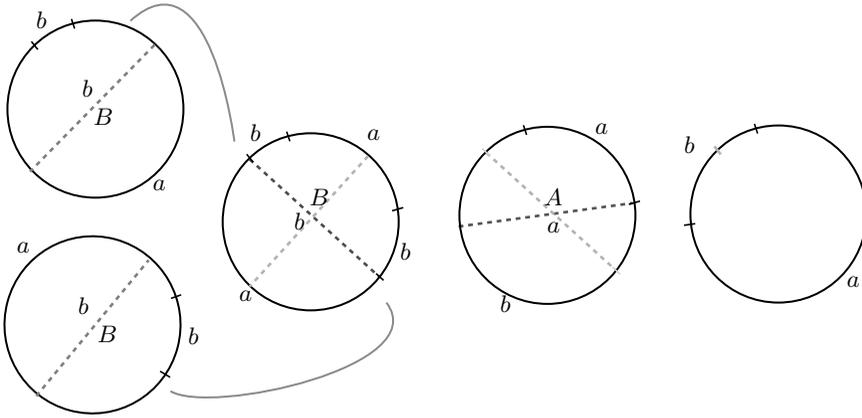


FIGURE 21. Combining b runs onto a single loop. The first step produces a by-product uniform a loop, the second an even loop, and the third another even loop.

and $a^{p_2}b^{w_1+w_2}$, where $t_1, t_2, w_2 > 0$ and $w_1 \geq 0$, and produces $a^{p_1}b^{t_1+w_2}$ and $a^{p_2}b^{w_1+t_2}$. The combine operator takes in $a^{p_1}b^{r_1}$ and $a^{p_2}b^{r_2}$, where $r_1 + r_2 < L/2$, and produces $a^{p_3}b^{r_1+r_2}$.

Notice that the trade operation requires only that *one* of the diameters be interior in the b run, or as written above, $w_2 > 0$ but $w_1 \geq 0$. Using the trade operation, we can take every b run and, for each one, move all but a single b onto a single chosen loop. Whenever this chosen loop contains a b run longer than $L/2$, we cut it off and start trading again. After this, we are left with a single loop with an unknown length b run, and possibly many loops with a single b . Then, we use the combine operation to combine these loops into a smaller number of loops with longer b runs. Then we trade the b mass to our chosen loop again, and so on.

After trading and combining as much as possible, we are left with either no loop (if the only remaining loop is even), a single loop with a b run of length less than $L/2$, or two loops whose combined run length is exactly $L/2$. This last case arises because we cannot use the combine operation on these loops. There is yet another sequence of moves, however, to resolve this: we use the trade operation to obtain

two loops with b runs of length exactly $L/4$. Then we cut and join them to produce a perfectly balanced loop with 4 runs. A single triangle move results in an even loop plus a commutator; we then apply Figure 18 to the commutator to get a pair of inverse loops.

Doing this to each loop type proves the lemma. □

The final step in the proof of Proposition 4.1.1 is to show that we can attach pants and annuli to the output of Lemma 4.2.5, that is, uniform loops, even loops, and a single remainder loop of each type, so that we have nothing left. Let $x_{a,b}$, $x_{a,B}$, $x_{A,b}$, and $x_{A,B}$ denote the run length of the single b run in the remainder loop of each type. Rescaling and considering arbitrary L , we can think of each variable as a real number in the interval $[0, 1/2)$. The fact that the entire collection of loops must be homologically trivial gives us two linear equations counting the homology contributions to a and b , respectively,

$$\begin{aligned} x_{a,b} + x_{A,b} - x_{a,B} - x_{A,B} &= k_1 \\ (1 - x_{a,b}) + (1 - x_{a,B}) - (1 - x_{A,b}) - (1 - x_{A,B}) &= k_2 \end{aligned}$$

Since there are even and uniform loops to consider, it is not *a priori* the case that $k_1 = k_2 = 0$. However, it *is* the case that $k_1, k_2 \in \frac{1}{2}\mathbb{Z}$ and $k_1 \pm k_2 \in \mathbb{Z}$, since the uniform and even loops change homology discretely by 1 and $1/2$, respectively. Now consider $k_1 + k_2 = 2(x_{A,b} - x_{a,B})$. Since $0 \leq x_{A,b}, x_{a,B} < 1$, we have $k_1 + k_2 \in (-1, 1)$, so $k_1 + k_2 = 0$. A similar argument shows that $k_1 - k_2 = 0$, so $k_1 = k_2 = 0$. Therefore, $x_{a,b} = x_{A,B}$ and $x_{a,B} = x_{A,b}$. These equalities show that actually, the remainder loops we have must be inverse pairs, so they are the boundary of two annuli. The homologically trivial collection of uniform and even loops can now be glued along entire runs, so is obviously pants equivalent to the empty collection.

That is, we are originally given a collection s' of loops of length L , and after applying Lemmas 4.2.1, 4.2.2, 4.2.3, 4.2.4, 4.2.5, and the argument above, we have shown that our original collection was pants equivalent to the empty collection, meaning we have successfully produced a collection of pants and annuli with boundary $s' + t + \iota(t)$, where t is the many intermediate boundaries we used to reduce s' . Adding in ι annuli, then, gives us a collection of pants and annuli which has boundary $s' + n\mathbf{1}'$, for some sufficiently large n .

Observe that we never need to duplicate our collection s' , or, equivalently, multiply v' by any factor. We have therefore proved the stronger Proposition 4.1.2 for rank 2 free groups.

4.3. Higher rank. We have completed the proof of Proposition 4.1.1 in the case that the free group has rank 2. We now describe the necessary modifications to the argument for higher rank free groups. Given a collection of loops $s' \in S'(L)$, the same technique of cutting with diameters works to show that s' is pants equivalent to a collection of 4-run loops. However, triangle moves no longer apply, since each of the 4 runs might be a run of a different generator.

The first step is to attach pants in such a way that we are left with 4-run loops, each of which only involves two generators. This is actually quite straightforward, since we have more freedom with the labels on the diameters that we attach. Figure 22 shows how to attach diameters to produce loops of the desired form which are pants equivalent to the original loop. Technically, the figures represent simply unions of pants, not the (nontrivalent) fatgraphs shown. They are drawn as shown

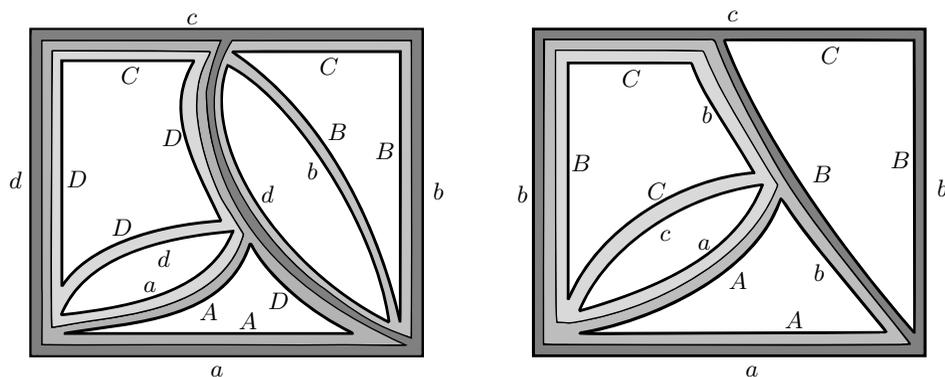


FIGURE 22. Reducing each loop to lie in a rank 2 subgroup. The fatgraph on the left is built out of folded pants only when b and d are distinct generators. If $b = d^{\pm 1}$, then we use the picture on the right. A similar picture holds when the vertical diameter has end points on runs of the same generator. The point is that all of the noninverse-matched pants boundaries lie in rank 2 subgroups.

to emphasize the point that we produce several other by-product loops, but they come in cancelling inverse pairs. All the interior diameters shown have length $L/2$, even though they are not drawn to scale.

We remark that the pictures in Figure 22 are general, up to rotation and reflection. The double diameter from top to bottom exists by the argument in the proof of Lemma 4.2.1, and the double diameter from the lower left corner to the middle exists because the first diameter has length $L/2$. We also remark that in a higher rank, it is possible that we have some 3-run loops. It is simple to cut these to 4-run loops and then apply the above argument.

After applying Lemma 4.2.5, we may assume that we are left entirely with uniform loops, even loops, and one 2-run loop of each type. Now, though, there are $4\binom{r}{2}$ loop types, which is too many to duplicate the linear-algebraic argument from the previous section. The simple solution is to take $L/2$ copies of our collection. Now each loop type is repeated exactly $L/2$ times, so when we re-collect the remainder, we are left with no remainder, so we have only uniform loops and even loops, which can be paired arbitrarily. Therefore, we have found a collection of pants and annuli which has boundary $(L/2)v' + n\mathbf{1}'$, which completes the proof of Proposition 4.1.1.

Remark 4.3.1. The statement of Proposition 4.1.3 doesn't quantify how n depends on the size of v , but this is technically necessary to deduce Theorem 3.3.1 from Proposition 4.1.3. Following the steps of the argument shows directly that $n = O(\|v\|_1)$; however, one can deduce the existence of such a linear bound on general grounds, as we now indicate. If we fix L , then the cone of homologically trivial collections of tagged loops of length L is a finite sided rational cone, so it has a finite Hilbert basis B . Applying Proposition 4.1.3 to each basis vector b gives a constant $n(b)$ so that $(L/2)b + n(b)\mathbf{1}$ bounds a collection of good pants and annuli. Since every integral v can be expressed as a disjoint union of copies of basis vectors, we obtain a uniform linear estimate for n . To deduce Theorem 3.3.1, we observe

that by making $\epsilon \ll \delta$, we can make the distribution of tagged loops as close to identical as desired.

Note that for random (rather than pseudorandom) words, if the mass of the reservoir is N , the central limit theorem says that the deviation from equidistribution will be of order \sqrt{N} .

4.4. Nonorientable surfaces. Proposition 4.1.3 shows that any sufficiently uniform collection s of tagged loops can be glued up into a thin fatgraph (in fact, can be glued up into just annuli and pants). We apply Proposition 4.1.3 as the last step in the proof of Theorem 3.3.1 to find many closed surface subgroups of random groups. We claim that Proposition 4.1.3 is the right way to do this, for reasons discussed at the end of this section. However, if we were only interested in the existence of surface subgroups, we can replace Proposition 4.1.3 with a trick which avoids most of the technical difficulty in this section.

Consider a collection s of tagged loops. Duplicate the collection once, so we assume that every loop in s appears an even number of times. Now take an untagged loop γ . In the collection s , there are many tagged copies of γ , and the tags are almost equidistributed around γ . Pair up these tagged loops such that in every pair, the tags are almost antipodal, and certainly distance $L/4$ apart. Let's consider a single pair γ_1 and γ_2 . They are both tagged copies of the loop γ . There's no annulus with *oriented* boundary $\gamma_1 + \gamma_2$, but there *is* an annulus bounding $\gamma_1 + \gamma_2$ in such a way that the orientation of one of the γ_i *disagrees* with that it inherits from the annulus. Performing this pairing for all pairs gives a collection of annuli bounding all the tagged γ loops.

This construction will give rise to *nonorientable* surfaces in random groups, which will be certified to be π_1 injective in the sequel. Taking an index 2 subgroup gives an oriented surface subgroup.

There are at least two good reasons to justify the hard work that went into the proof of Proposition 4.1.3. The first is that this proposition is of independent interest and can be used as a combinatorial tool in many contexts where it is important to build orientable surfaces or fatgraphs. The second is that the nonorientable surface subgroups built using the trick above will never be essential in H_2 , whereas the surfaces built using the full power of the Thin Fatgraph Theorem *can* be taken to be homologically essential in random groups whenever the density D is positive.

We remark that the original Kahn-Markovic construction of surface subgroups in hyperbolic 3-manifolds necessarily produced nonorientable (and therefore homologically inessential) quasifuchsian surfaces. Very recently, Liu and Markovic [12] have shown how to modify the construction—using substantial ingredients from the Kahn-Markovic proof of the Ehrenpreis conjecture—to build orientable quasifuchsian surfaces (projectively) realizing any homology class.

5. RANDOM ONE-RELATOR GROUPS

Throughout this section we fix a free group F_k with $k \geq 2$ generators, and we fix a free generating set. For some big (unspecified) constant n , we let r be a random cyclically reduced word of length n , and we consider the one-relator group $G := \langle F_k \mid r \rangle$.

In this section we will show that with probability going to 1 as $n \rightarrow \infty$, the group G contains a surface subgroup $\pi_1(S)$. The surface S in question can be built from N disks bounded by r and N disks bounded by r^{-1} , glued up along their boundary

in such a way that the 1-skeleton is a trivalent fatgraph with every edge of length $\geq L$, where L is some (arbitrarily big) constant fixed in advance, and $N \leq 20L$ is the constant in the Thin Fatgraph Theorem 3.3.1.

The group G is evidently $C'(\lambda)$ for any $\lambda > 0$ with probability going to 1 as $n \rightarrow \infty$, and therefore the injectivity of S can be verified by showing that the 1-skeleton of S does not contain a long path in common with r or r^{-1} , except for a path contained in the boundary of one of the $2N$ disks.

A trivalent graph Y in which every edge has length at least L has at most $2|Y| \cdot 2^{m/L}$ subpaths of length m , where $|Y|$ is the length of Y . In our context the trivalent graph Y will arise as the 1-skeleton of a surface S constructed as above, so $|Y| = O(n)$. In a free group of rank k there are (approximately) $(2k - 1)^m$ reduced words of length m , and the relator r contains at most $2n$ of them (allowing inverses). So if we fix any positive ϵ' , and take $m = \epsilon' \cdot n$, then providing we choose L so that $1/L \log_{2k-1} 2 < \epsilon'$ a simple counting argument shows that Y does not have any path of length m in common with an *independent* random word of length n , with probability $1 - O(e^{-n^c})$. However, Y and r are utterly dependent, and we must work harder to show that S is injective.

The key idea is the observation that disjoint subwords of a long random relator r are (almost) *independent of each other*. Informally, we fix some small positive δ , and break up $Nr \cup Nr^{-1}$ into pieces (called *beads*) of size $n^{1-\delta}$ which each bound their own trivalent fatgraph (by the Thin Fatgraph Theorem). Then subpaths in the fatgraph associated with one bead *will* be independent of the subpaths of $Nr \cup Nr^{-1}$ associated with another bead, and this argument can be made to work.

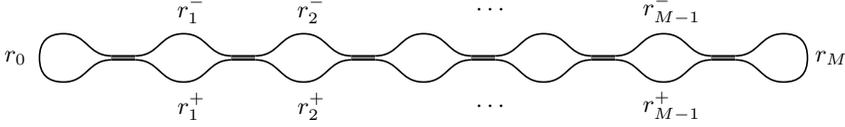
Remark 5.0.1. The “simple counting argument” alluded to above is a special case of Gromov’s intersection formula ([8] Section 9.A) which implies that two sets of independent random words of a fixed length whose (multiplicative) densities sum to less than 1 are typically disjoint. We apply this observation in a more substantial way in Section 6, especially in the proof of Theorem 6.4.1.

5.1. Independence and correlation. Since this is the first point in the argument where we are using the genuine randomness of the relators (rather than just pseudorandomness), some remarks are in order.

A random word in a finite alphabet (in the uniform distribution) has the property that any two disjoint subwords are independent. A random (cyclically) *reduced* word in the free group fails to have this property, since (for example) if uv are adjacent subwords, the last letter of u must not cancel the first letter of v (so the words are not really independent). However, such words have a slightly weaker property which is just as useful as independence in most circumstances; this property can be summarized by saying that *correlations decay exponentially*.

To explain the meaning of this, let’s fix reduced words u and v and a distance T . Suppose r is a random (reduced) word, and let’s write r as $abcd$ where $|a| = |u|$, where $|b| = T$, where $|c| = |v|$, and where $|d| = n - |a| - |b| - |c|$. The probability that $c = v$ is $1/(2k)(2k - 1)^{v-1}$. Saying that correlations decay exponentially means that the probability that $c = v$ *conditioned on* $a = u$ satisfies

$$1 - (2k - 2)^{-T} < \frac{\Pr(c = v \mid a = u)}{\Pr(c = v)} < 1 + (2k - 2)^{-T}$$

FIGURE 23. Bead decomposition of r .

This estimate is elementary; see, e.g., [3], Lemma 2.4 for a careful proof. In the more general context of stationary Markov chains, such exponential bounds on the deviation from independence in the tail are called *Chernoff bounds*.

In the sequel we are typically interested in estimating the probability of finding a pair of matching subwords which are separated by a considerable distance. Only order-of-magnitude estimates of the probability are important for our arguments. So in practice we can treat the subwords as though they were independent.

5.2. Beads. Let r be a random cyclically reduced word of length n . We fix some (small) positive constants C and δ ; later we will say how small they should be.

Definition 5.2.1. A *bead decomposition* of r is a decomposition of r into segments labeled (in cyclic order) $r_0, r_1^+, r_2^+, \dots, r_{M-1}^+, r_M, r_{M-1}^-, \dots, r_1^-$ where each r_i^\pm has length $n^{1-\delta} \cdot (1/2 + o(1))$ and r_0, r_M have length $n^{1-\delta} \cdot (1 + o(1))$ [so that $M = n^\delta \cdot (1 + o(1))$] so that the prefix of r_i^+ of length $C \log(n)$ is inverse to the suffix of r_i^- , and the prefix of r_M of length $C \log(n)$ is inverse to the suffix of r_0 .

See Figure 23.

Given a bead decomposition of r , we glue the mutually inverse subwords described above, thereby decomposing r into a sequence of loops (“beads”) of length $n^{1-\delta} \cdot (1 + o(1))$ with one or two tags. Denote these beads B_0, B_1, \dots, B_M . The edges of length $C \log(n)$ obtained by gluing the inverse prefixes/suffixes we refer to as the *lips* of the beads.

Lemma 5.2.2 (Bead decomposition). *There exists a bead decomposition with probability $1 - O(e^{-n^c})$.*

Proof. We start at an arbitrary location in r , and take this to be the approximate midpoint of r_0 . Extend this region $n^{1-\delta} \cdot (1/2 - n^{-\delta})$ in either direction, and then read the next pair of segments of length $n^{1-2\delta}$ synchronously until the first time we read off a pair of mutually inverse segments of length $C \log(n)$. If $C < (1 - 2\delta)/\log(2k - 1)$, then such a pair of mutually inverse segments will occur with probability $1 - O(e^{-n^c})$ by Chernoff’s inequality, for some $c > 0$ depending on δ . We then build successive segments r_i^\pm in order by the same procedure. There are n^δ such segments, and this polynomial term can be absorbed into the exponential estimate of probability at the cost of adjusting constants. \square

To think about the (lack of) correlation between subwords of the different B_i , the following mental picture is useful: imagine that we generate the word r by a Markov process letter by letter as we go, starting at the center of r_0 and building outward. In this model, the letters making up each successive B_i are only generated after we have already constructed B_j with $j < i$.

There is no reason to expect that the beads B_i are homologically trivial, but there is a trick to adjust them so that they are. We build a bead decomposition of r and of r^{-1} simultaneously, so that the beads of r^{-1} have inverse labels to the

beads of r . We denote the beads of r^{-1} by B_i^{-1} , and note that they are inverse (as tagged cyclic words) to the B_i . Then for each i the union $B_i \cup B_i^{-1}$ is homologically trivial.

By Theorem 3.3.1 for each i , and for some $N \leq 20L$, the collection $NB_i \cup NB_i^{-1}$ bounds a trivalent fatgraph Y_i with all edges of length at least L , with probability $1 - O(e^{-n^c})$. Since we can first build the fatgraph Y_i in a way which depends only on the substrings r_i^\pm , the Chernoff bound says that for any positive α there is $c(\alpha)$ so that with probability $1 - O(e^{-n^c})$ there are no paths in Y_i of length n^α in common with any segment in $r - r_i^\pm$. Summing over the n^δ different indices i , and absorbing this polynomial factor into the probability estimate (at the cost of adjusting constants), we see that with probability $1 - O(e^{-n^c})$ there is no index i and no path in any Y_i of length n^α in common with any segment of $r - r_i^\pm$.

Lemma 5.2.3 (No long path). *Let $\beta > 0$ be fixed. Let Y be the fatgraph obtained from the union of the Y_i associated with a bead decomposition as above. Then there is some positive $c(\beta)$ so that with probability $1 - O(e^{-n^c})$ every path in Y of length βn which appears in r or r^{-1} is in $\partial S(Y)$.*

Proof. Let γ be a path in Y of length βn and let $\gamma' \subset r$ (without loss of generality) have the same labels as γ . Then for any fixed $\alpha > 0$ there is a $c > 0$ so that for each i and each subsegment σ' of γ' of length n^α contained in B_i the corresponding subsegment σ of γ must have at least $(1 - o(1))$ of its length contained in Y_i , with probability $1 - O(e^{-n^c})$. By the definition of the bead decomposition, successive subsegments of γ' in adjacent B_i are joined by paths of length $C \log(n)$ running over the lip. The corresponding subsegments in γ that transition from Y_i to Y_{i+1} must *also* run over the lip, so there is another copy of the word on the lip contained in r within distance n^α of the lip. If the two copies are not distinct, so that γ and γ' overlap on a common path, then since Y is folded we must simply have $\gamma = \gamma'$ and γ is in $\partial S(Y)$ as claimed. Otherwise, there are two distinct copies of the lip contained in a segment of length n^α in r . See Figure 24.

If α is sufficiently small compared to C , the probability that two identical subwords of length $C \log(n)$ will occur in a specific segment of length n^α is arbitrarily small [in fact, of size $O(n^{-\alpha'})$ for some α' depending on α and C]. Explicitly, in a segment of length n^α , the expected length of the longest pair of identical subwords is $2\alpha \log_{2k-1}(n)$, so if we choose $C > 2\alpha / \log(2k-1)$ we obtain the desired estimate [with α' depending on the difference between C and $2\alpha / \log(2k-1)$]. Remember that $C < (1 - 2\delta) / \log(2k-1)$ and our only constraint on δ is that it is positive. So we can achieve $C > 2\alpha / \log(2k-1)$ subject to this constraint if α is sufficiently small.

If γ' has length βn it must therefore run over βn^δ successive lips in this way, and each time it runs over a lip it must contain two identical (or inverse) subwords of length $C \log(n)$ within distance n^α of each other; call each pair of identical (or inverse) subwords a *coincidence*. Each coincidence occurs with probability $O(n^{-\alpha'})$. The probability of successive coincidences at adjacent lips is not independent, but the Chernoff bound says that the probability of $O(n^\delta)$ such coincidences in succession is of order $O((n^{-\alpha'})^{n^\delta}) = O(e^{-n^c})$ for some c , and the lemma is proved. \square

We deduce the following theorem as a corollary.

Theorem 5.2.4 (Random One-Relator). *Fix a free group F_k and let r be a random cyclically reduced word of length n . Then $G := \langle F_k \mid r \rangle$ contains a surface subgroup with probability $1 - O(e^{-n^c})$.*

Proof. This follows from Lemma 5.2.3 and Lemma 2.2.5. □

There are $O(e^{n^c})$ many choices of bead decomposition, and most of these give rise to quasiconvex surface subgroups.

Definition 5.2.5. We call the surfaces constructed as above *beaded surfaces*.

Note that a beaded surface has genus $O(n)$. If $k = 2$, the least genus of a beaded surface is $o(n)$, since we can take $N = 1$ in the application of the Thin Fatgraph Theorem, and then as $n \rightarrow \infty$ we can take $L \rightarrow \infty$. It seems very likely that the least genus of a beaded surface is $o(n)$ for any fixed $k \geq 2$.

Remark 5.2.6. It will turn out (after the proof of Theorem 6.4.1) that beaded surfaces are quasiconvex (in fact, they stay quasiconvex even after adding many more random relators), but it is more efficient to give the proof of this in the next section.

6. RANDOM GROUPS

In this section we prove our main theorem, that a random group at density $D < 1/2$ contains a surface subgroup with probability $1 - O(e^{-n^c})$. In fact, our argument shows that it contains *many* subgroups [of genus $O(n)$]. Our argument depends on some elements of the theory of small cancellation developed for random groups by Ollivier [13], and we refer to that paper several times.

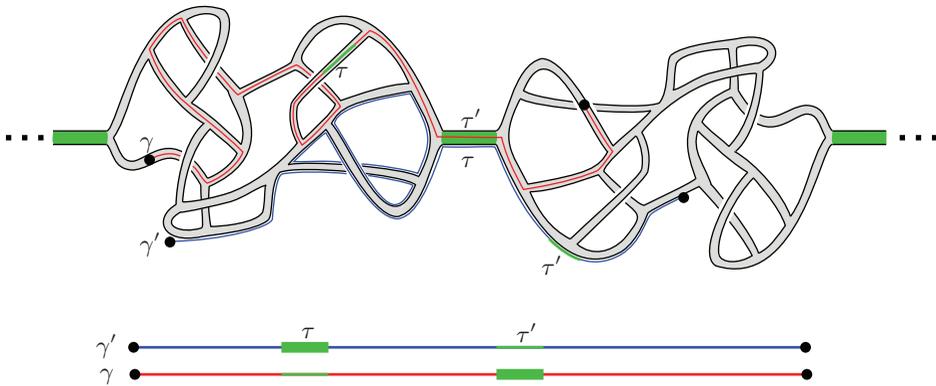


FIGURE 24. The figure shows part of a surface obtained by fattening the spine Y , with boundary contained in part of a single copy of r (for simplicity). Every time $\gamma' \subset r$ (blue) runs over a (thick green) lip τ [which has length $C \log(n)$], the path γ (red) must also run over τ at almost the same time. Since γ and γ' have the same labels, the copy of τ in γ gives rise to a *coincidence*; i.e., another copy τ' (thin green) of τ in γ' within distance n^α of the first copy.

6.1. Small cancellation in random groups. For later convenience, we here state three results from Ollivier [13] that we use in the sequel.

Theorem 6.1.1 (Ollivier [13], Theorem 2). *Let G be a random group at density D . Then for any positive ϵ , and any reduced van Kampen diagram \mathcal{D} containing m disks, we have*

$$|\partial\mathcal{D}| \geq (1 - 2D - \epsilon) \cdot nm$$

with probability $1 - O(e^{-n^c})$

Here the hardest part is to show that the same ϵ works for van Kampen diagrams of arbitrary size.

Theorem 6.1.2 (Ollivier [13], Corollary 3). *Let G be a random group at density D . Then the hyperbolicity constant δ of the presentation satisfies $\delta \leq 4n/(1 - 2D)$ with probability $1 - O(e^{-n^c})$.*

Theorem 6.1.3 (Ollivier [13], Theorem 6). *Let G be a random group at density D . Then for any positive ϵ , and for any reduced van Kampen diagram \mathcal{D} with at least two faces, there are at least two faces which have a (connected) piece on $\partial\mathcal{D}$ of length at least $n(1 - 5D/2 - \epsilon)$, with probability $1 - O(e^{-n^c})$.*

The statements of theorems in Ollivier's paper do not make the estimate of probability (as a function of n) explicit; however, these estimates are straightforward to derive from his methods (and in any case, we do not use them in the sequel).

6.2. Convexity. We now indicate how to use small cancellation arguments to find a surface subgroup at any $D < 2/7$. This is proved by a counting argument (Lemma 6.2.1), which is a model for the general case $D < 1/2$ (proved in Theorem 6.4.1).

Pick one relator r , and build a beaded surface S as in Section 5 whose spine is trivalent and with every edge of length $\geq L$ for some large (fixed) L .

Lemma 6.2.1. *Fix D . Then for any $\alpha > D$, a beaded surface S constructed by the method of Section 5 is α convex, with probability $1 - O(e^{-n^c})$.*

Proof. By Lemma 5.2.3 for any positive β , the spine Y of S has no subword of length βn in common with the relator r except for subwords occurring in $\partial S(Y)$. For any $\alpha > 0$ there are $O(2^{\alpha n/L})$ paths in Y of length αn . Define $\beta' = \log(2)\alpha/L$ so that $e^{\beta'n} = 2^{\alpha n/L}$, and note that by taking L sufficiently large, we can make β' as small as we want.

There are $(2k - 1)^{\alpha n}$ reduced words of length αn , and a random relator r' contains $2n$ subwords of this length counting inverses (which is polynomial in n , and therefore is absorbed into the exponential terms in our estimates), so a random relator r' has probability $O(e^{\beta'n - \log(2k-1)\alpha n})$ of having a subword of length αn in common with Y . If $\alpha > D$ and $\beta' < \log(2k - 1)(\alpha - D)$, then no relator $r' \neq r$ has a subword of length αn in common with Y , with probability $O(e^{-n^c})$. \square

We deduce by Lemma 2.2.5 that a random group contains a surface subgroup at any $D < 1/12$. However, Theorem 6.1.3 already improves this to $D < 2/7$.

6.3. van Kampen disks. Our strategy will be to show that the existence of a certain kind of van Kampen disk \mathcal{D} with boundary a cyclically reduced word in Y essential in S gives rise to a contradiction. Suppose that $\gamma \subset Y$ is an essential loop in S whose image is trivial in G , so that there is some van Kampen diagram \mathcal{D} with boundary γ . If some face in \mathcal{D} has boundary r or r^{-1} , and if this face has a segment of length more than βn in common with γ , then this face agrees with a disk of S , and we can find a smaller van Kampen diagram \mathcal{D}' by pushing across this disk. So in the sequel we will only consider loops $\gamma \subset Y$ essential in S bounding van Kampen disks \mathcal{D} which cannot be simplified by such a move. We call such a van Kampen disk *efficient*.

The following Lemma is standard.

Lemma 6.3.1 (Short shortcut). *Let G be a hyperbolic group with a presentation with respect to which it is δ hyperbolic. Let Γ be a cyclically reduced word in the generators which is trivial in G . Then there is a van Kampen disk \mathcal{D} with $|\partial\mathcal{D}| \leq 18\delta$ and a connected subpath $\gamma \subset \partial\mathcal{D}$ with $\gamma \subset \Gamma$ and $|\gamma| > |\partial\mathcal{D}|/2$.*

Note that if $\gamma' = \partial\mathcal{D} - \gamma$, then $|\gamma'| < |\gamma|$. In other words, γ' is a *shortcut*; hence the terminology.

Proof. In any δ -hyperbolic path metric space, for any $k > 8\delta$, a k -local geodesic (i.e., a 1-manifold for which every subpath of length at most k is a geodesic) is a (global) $(\frac{k+4\delta}{k-4\delta}, 2\delta)$ quasigeodesic; see [2], Ch. III. H, 1.13 p. 405. The loop Γ starts and ends at the same point, and is therefore not a k -local geodesic for $k \geq 9\delta$. Therefore some segment of length at most 9δ is not geodesic, and it cobounds a van Kampen disk \mathcal{D} with an honest geodesic. □

We deduce the following corollary.

Lemma 6.3.2. *Suppose that S is a beaded surface which is not π_1 injective. Then there are constants C and C' depending only on $D < 1/2$, a geodesic path γ in the spine Y of length at most Cn , and a van Kampen diagram \mathcal{D} containing at most C' faces so that $\gamma \subset \partial\mathcal{D}$ and $|\gamma| > |\partial\mathcal{D}|/2$.*

Proof. Theorem 6.1.2 says that $\delta \leq 4n/(1 - 2D)$, so by Lemma 6.3.1 it follows that there is such a disk \mathcal{D} with boundary of length at most $72n/(1 - 2D)$. On the other hand, by Theorem 6.1.1 we know $72n/(1 - 2D) \geq |\partial\mathcal{D}| \geq (1 - 2D - \epsilon) \cdot nC'$ where C' is the number of faces; in particular, C' is *bounded* in terms of D (and *independent* of n). □

The fact that C and C' can be chosen *independent* of n (but depending on $D < 1/2$, of course) is crucial for our purposes.

6.4. Surfaces in random groups. We can now prove the main theorem of the paper.

Theorem 6.4.1 (Surfaces in random groups). *A random group of length n and density $D < 1/2$ contains a surface subgroup with probability $1 - O(e^{-n^c})$. In fact, it contains $O(e^{n^c})$ surfaces of genus $O(n)$. Moreover, these surfaces are quasiconvex.*

Proof. Pick one relation r and build a beaded surface S by the method of Section 5. We have already shown that the 1-skeleton Y of S does not contain a path of length βn for any fixed positive β in common with r or r^{-1} except for paths in the boundary of a disk, with the desired probability.

Suppose S is not π_1 injective. Then by Lemma 6.3.2 there is an efficient van Kampen disk with boundary an essential loop in Y , containing a subdisk \mathcal{D} with at most C' faces, and at least half of its boundary equal to some path γ in the spine Y . We want to show that the existence of such a van Kampen disk is very unlikely, for fixed $D < 1/2$, and for n sufficiently big.

Fix a combinatorial type for the diagram. Then there are at most polynomial in n choices of edge lengths for the edges in the diagram. Choose a collection of edge lengths. Let $m \leq C'$ be the number of faces.

We estimate the probability that there is a way to label each face with a relator or its inverse compatible with some γ . We express the count in terms of *degrees of freedom*, measured multiplicatively, as powers of $(2k - 1)$ (Gromov uses the terms *density* and *codensity*; see the discussion in [8] pp. 269–272 expanded at length in [14]).

First, the choice of γ itself gives $n\beta'$ degrees of freedom, where $\beta' = \log(2)\alpha/L$, and $|\gamma| = \alpha n$, as in the proof of Lemma 6.2.1. Since $\alpha \leq C'/2$ is an absolute constant depending only on D , by choosing L big enough we can make β' as small as desired (we only really need $\beta' < 1/2 - D$), and therefore we may effectively neglect it in what follows.

Next we consider the disks with boundary label r or r^{-1} . Since the original disk was efficient, no face of \mathcal{D} with boundary label r or r^{-1} has a segment of more than βn in common with γ . Furthermore, a fixed random relation r of length n will have no piece in common with itself or its inverse of length $\epsilon'' n$ for any positive ϵ'' , with probability $1 - O(e^{-Cn})$; if we take $\epsilon'' = \beta$ for simplicity, we deduce that in a reduced diagram, no two disks with boundary label r or r^{-1} share a segment of their boundary in common of length more than βn .

In a van Kampen disk with at most m faces, the boundary of each face is decomposed into at most m segments, each of which is shared with another face or with the boundary. Thus each face with boundary label r or r^{-1} has at most $m\beta n$ of its boundary in common with γ or with other faces with label r or r^{-1} . Let \mathcal{D}' be the subdiagram obtained by cutting out the faces with boundary label r or r^{-1} , and let $\gamma' \subset \partial\mathcal{D}'$ be the union of $\gamma \cap \partial\mathcal{D}'$ with $\partial\mathcal{D}' - \partial\mathcal{D}$. Finally, let m' be the number of faces in \mathcal{D}' . Taking β sufficiently small, we can assume that $m\beta < 1/2$, and therefore $|\gamma'| \geq |\gamma|$ so that $|\gamma'| \geq |\partial\mathcal{D}'|/2$ and $m' \leq m$, with equality if and only if $\mathcal{D}' = \mathcal{D}$.

Each remaining choice of face gives nD degrees of freedom, and each segment in the interior of length ℓ imposes ℓ degrees of constraint. Similarly, γ' imposes $|\gamma'|$ degrees of constraint. Let \mathcal{J} denote the union of interior edges. Then $|\partial\mathcal{D}'| + 2|\mathcal{J}| = nm'$ so $|\gamma'| + |\mathcal{J}| \geq nm'/2$ because $|\gamma'| \geq |\partial\mathcal{D}'|/2$. On the other hand, the total degrees of freedom is $nm'D + n\beta' < nm'/2$, so no assignment is possible, with probability $1 - O(e^{-n^c})$. Summing the exceptional cases over the polynomial in n assignments of lengths, and the *finite* number of combinatorial diagrams, we see that S is injective, with probability $1 - O(e^{-n^c})$ for some c depending on D (and going to 0 as $D \rightarrow 1/2$).

In fact, the same argument implies that every geodesic path in the 1-skeleton of Y is actually quasigeodesic in K , by the proof of Lemma 6.3.1. Explicitly, for any fixed $k > 9\delta$ we can repeat the argument above with C' replaced by $C'k/9\delta$, and deduce that a geodesic in the 1-skeleton of (the universal cover of) S is mapped to a k -local geodesic in (the universal cover of) K . The theorem follows. \square

Remark 6.4.2. Since we may take $k > 9\delta$ arbitrarily large, the estimate in the proof of Lemma 6.3.1 actually shows that beaded surfaces can be taken to be $(1 + \epsilon)$ quasiconvex for any fixed $\epsilon > 0$.

Remark 6.4.3. The surfaces we build are homologically trivial, since they map nontrivially over only one disk bounded by a relator r , and with total degree 0. As remarked in the Introduction, because the surfaces produced by the Thin Fatgraph Theorem are *orientable*, a modification of our construction produces homologically essential surfaces.

If n is even, a random reduced word of length n in a free group of rank k is homologically trivial with probability $O(n^{-k/2})$. Since there are $(2k-1)^{nD}$ relators, there are an enormous number of such homologically trivial relators, and we can try to build a surface mapping over the associated disk with degree 1 (and therefore being homologically essential in G). Evidently, the only obstruction to finding such surfaces is to build a bead decomposition as in Lemma 5.2.2 where all the B_i are homologically trivial, while still preserving the property that correlations between distinct B_i decay exponentially fast. The probability that the naive construction of a bead decomposition (as in the lemma) applied to a random word will have this property is $(n^{-k/2})^{n^\delta}$, which is subexponential in n , so many of the relators will have this property, and we can build many homologically essential surfaces of genus $O(n)$. If n is odd, we can build a similar (homologically essential) beaded surface from two (judiciously chosen) relators.

Remark 6.4.4. The surfaces we build have genus $O(n)$ [or $o(n)$ for rank 2], and it is natural to wonder if this is the best possible. We conjecture not; in fact, we conjecture that the smallest genus injective surfaces in random groups are of genus $O(n/\log n)$ (at any density $D < 1/2$).

In fact, Theorem 4.16 of [3] gives a precise estimate of the geometry of the Gromov norm on $H_2(G; \mathbb{R})$. Let V be the vector space with the relators r_i as the basis, and let W be the kernel of the natural map $V \rightarrow H_1(F_k; \mathbb{R})$. Then we can identify W with $H_2(G; \mathbb{R})$, by Mayer-Vietoris. The vector space W inherits an L^1 norm from V with respect to its given basis. The Gromov norm on W (on random subspaces of fixed dimension) is (with overwhelming probability) proportional to this L^1 norm, with constant of proportionality $2 \log(2k-1)n/3 \log(n)$, up to a multiplicative error of size $1 + o(1)$ (there is a factor of 4 relative to the statement of Theorem 4.16 of [3]; this factor of 4 reflects the difference between the Gromov norm and the so-called *scl norm*). It seems plausible that classes in $H_2(G; \mathbb{Q})$ should be projectively represented by norm-minimizing surfaces; such surfaces will necessarily be injective. Again, it seems likely that one should not need to pass to a very big multiple of a class to find an extremal surface (at least for some classes); so there should be injective surfaces of genus $O(n/\log(n))$. We strongly suspect this order of magnitude is sharp.

ACKNOWLEDGMENTS

The authors would like to thank Misha Gromov, John Mackay, Yann Ollivier, Piotr Przytycki, Henry Wilton and the anonymous referee. The authors also would like to acknowledge the use of Colin Rourke's `pinlabel` program, and Nathan Dunfield's `labelpin` program to help add the (numerous) labels to the figures. The

first author was supported by NSF grant DMS 1005246, and the second author was supported by NSF grant DMS 1203888.

REFERENCES

- [1] Agol I., *The virtual Haken conjecture*, available at [arXiv:1204.2810](https://arxiv.org/abs/1204.2810). with an appendix with D. Groves and J. Manning.
- [2] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR1744486 (2000k:53038)
- [3] Danny Calegari and Alden Walker, *Random rigidity in the free group*, *Geom. Topol.* **17** (2013), no. 3, 1707–1744, DOI 10.2140/gt.2013.17.1707. MR3073933
- [4] Calegari D. and Walker A. *scallop*, Computer program available from the authors' webpages.
- [5] Calegari D. and Walker A., *Surface subgroups from linear programming*, available at [arXiv:1212.2618](https://arxiv.org/abs/1212.2618).
- [6] Calegari D. and Wilton H., *Random graphs of free groups contain surface subgroups*, available at [arXiv:1303.2700](https://arxiv.org/abs/1303.2700).
- [7] François Dahmani, Vincent Guirardel, and Piotr Przytycki, *Random groups do not split*, *Math. Ann.* **349** (2011), no. 3, 657–673, DOI 10.1007/s00208-010-0532-4. MR2755002 (2012b:20101)
- [8] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric group theory, Vol. 2 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, Cambridge, 1993, pp. 1–295. MR1253544 (95m:20041)
- [9] Gromov M., personal communication.
- [10] Jeremy Kahn and Vladimir Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, *Ann. of Math. (2)* **175** (2012), no. 3, 1127–1190, DOI 10.4007/annals.2012.175.3.4. MR2912704
- [11] Marcin Kotowski and Michał Kotowski, *Random groups and property (T): Żuk's theorem revisited*, *J. Lond. Math. Soc. (2)* **88** (2013), no. 2, 396–416, DOI 10.1112/jlms/jdt024. MR3106728
- [12] Liu Y. and Markovic V., *Homology of curves and surfaces in closed hyperbolic 3-manifolds*, available at [arXiv:1309.7418](https://arxiv.org/abs/1309.7418).
- [13] Yann Ollivier, *Some small cancellation properties of random groups*, *Internat. J. Algebra Comput.* **17** (2007), no. 1, 37–51, DOI 10.1142/S021819670700338X. MR2300404 (2008g:20096)
- [14] Yann Ollivier, *A January 2005 invitation to random groups*, *Ensaos Matemáticos [Mathematical Surveys]*, vol. 10, Sociedade Brasileira de Matemática, Rio de Janeiro, 2005. MR2205306 (2007e:20088)
- [15] Yann Ollivier and Daniel T. Wise, *Cubulating random groups at density less than 1/6*, *Trans. Amer. Math. Soc.* **363** (2011), no. 9, 4701–4733, DOI 10.1090/S0002-9947-2011-05197-4. MR2806688 (2012e:20095)
- [16] John R. Stallings, *Topology of finite graphs*, *Invent. Math.* **71** (1983), no. 3, 551–565, DOI 10.1007/BF02095993. MR695906 (85m:05037a)
- [17] A. Żuk, *Property (T) and Kazhdan constants for discrete groups*, *Geom. Funct. Anal.* **13** (2003), no. 3, 643–670, DOI 10.1007/s00039-003-0425-8. MR1995802 (2004m:20079)

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

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