ANNALES DE L'INSTITUT FOURIER

Vincent Guirardel *&* Gilbert Levitt
Random subgroups, automorphisms, splittings

RANDOM SUBGROUPS, AUTOMORPHISMS, SPLITTINGS

by Vincent GUIRARDEL & Gilbert LEVITT

————

ABSTRACT. — We show that, if H is a random subgroup of a finitely generated free group $\mathbb{F}_k$, only inner automorphisms of $\mathbb{F}_k$ may leave H invariant. A similar result holds for random subgroups of toral relatively hyperbolic groups, and more generally of groups which are hyperbolic relative to slender subgroups. These results follow from non-existence of splittings over slender groups which are relative to a random group element. Random subgroups are defined using random walks or balls in a Cayley tree of $\mathbb{F}_k$.

In the free group $\mathbb{F}_k$, we also prove the following deterministic result: if a cyclically reduced word $h \in \mathbb{F}_k$ contains all reduced words of length L, then $\mathbb{F}_k$ has no splitting relative to h over a subgroup of rank $\leqslant (k-1)(L-2)$.

RÉSUMÉ. — Nous montrons que si H est un sous-groupe aléatoire d'un groupe libre de type fini $\mathbb{F}_k$, tout automorphisme de $\mathbb{F}_k$ préservant H est intérieur. Nous prouvons un résultat similaire pour les sous-groupes aléatoires de groupes hyperboliques toriques, et plus généralement de groupes hyperboliques relativement à des sous-groupes sveltes. Ces résultats découlent de la non-existence de scindements au-dessus de sous-groupes sveltes qui sont relatifs à un élément aléatoire. Les sous-groupes aléatoires peuvent être définis en termes de marches aléatoires ou de boules dans le graphe de Cayley de $\mathbb{F}_k$.

Dans le cas du groupe libre $\mathbb{F}_k$, nous démontrons aussi le résultat déterministe suivant : si un mot cycliquement réduit $h \in \mathbb{F}_k$ contient tous les mots réduits de longueur L, alors $\mathbb{F}_k$ n'a pas de scindement relatif à h au-dessus d'un sous-groupe de rang $\leqslant (k-1)(L-2)$.

1. Introduction

When studying an automorphism of a group G, it is often useful to consider invariant subgroups. For instance, irreducible automorphisms of free groups are defined by considering invariant free factors [4].

One may also fix a subgroup H and consider the group $\mathrm{Aut}(G, H)$ of automorphisms of G leaving H invariant, or the group of automorphisms of H extending to automorphisms of G. There may be many of those, for instance if H is a free factor or a direct factor. The authors have proved that, conversely, if H is a non-cyclic subgroup of a finitely generated free group G, and every automorphism of H extends to G, then H is a free factor [10].

At the other extreme, Schupp proved in [23] that any H may be embedded into a group G so that only inner automorphisms of H extend. The proof uses small cancellation, so G is defined in an explicit, but ad hoc, way. A starting point of the present paper is the idea that this non-extension phenomenon should be fairly common, and we express this by the following principle:

> If H is a very complicated subgroup of G, then very few automorphisms of G leave H invariant.

But this principle is not valid in full generality. For instance, if $G = \mathbb{Z}^n$, every subgroup is left invariant by many automorphisms.

In any case, $\mathrm{Aut}(G, H)$ will always contain the group $\mathrm{Inn}_H(G) < \mathrm{Inn}(G)$ defined as the set of all conjugations by elements of H. In this paper we prove:

THEOREM 1.1 (Theorem 7.2). — *Assume that G is hyperbolic relative to a finite family $\mathcal{P}$ of slender subgroups. If H is a random subgroup of G, then $\mathrm{Aut}(G, H)/\mathrm{Inn}_H(G)$ is finite.*

Recall that a group A is slender if A and all its subgroups are finitely generated. To define a random subgroup of G, we fix $p \geqslant 1$ and we let H be generated by p independent random walks of length n (see Definition 5.1 for details). The conclusion of the theorem then holds with probability going to 1 as $n \to \infty$. We rely on results of Maher–Sisto [18] about random walks, and our assumptions are the same as in their paper (Theorem 1.1 would apply to subgroups generated by elements chosen randomly independently in balls, as in Theorem 1.4 below, if the results of [18] were known to hold in that context).

We believe that $\mathrm{Inn}_H(G)$ is actually equal to $\mathrm{Aut}(G, H)$ when G is torsion-free, but our methods do not allow us to prove it unless G is a free group $\mathbb{F}_k$.

Let X be a free basis of $\mathbb{F}_k$. For the standard simple random walk on $\mathbb{F}_k$, associated to the uniform measure on $X^{\pm 1} = X \cup X^{-1}$, we show that

generically $\mathrm{Inn}_H(G)$ is precisely equal to $\mathrm{Aut}(G, H)$. More precisely, we show that a random subgroup H is malcharacteristic in the following sense:

DEFINITION 1.2 ([16]). — A subgroup $H < G$ is malcharacteristic if any $\alpha \in \mathrm{Aut}(G)$ such that $\alpha(H) \cap H \neq \{1\}$ belongs to $\mathrm{Inn}_H(G)$.

THEOREM 1.3 (Theorem 7.6). — Fix $k \geqslant 2$ and $p \geqslant 1$. Let $H = \langle w_1, \ldots, w_p \rangle \subset \mathbb{F}_k$ be the subgroup generated by p independent simple random walks $w_1, \ldots, w_p$ of length n with respect to a free basis.

With probability going to 1 exponentially fast as $n \to +\infty$, the subgroup H is malcharacteristic. In particular, $\mathrm{Aut}(G, H) = \mathrm{Inn}_H(G)$.

There is a similar result if one chooses the elements $w_1, \ldots, w_p$ independently randomly in the ball of radius n (for the word metric associated to X). This answers Questions 4 and 6 of [16] positively.

THEOREM 1.4 (Theorem 7.5). — Fix $k \geqslant 2$ and $p \geqslant 1$. With probability going to 1 exponentially fast as $n \to +\infty$, the subgroup $H \subset \mathbb{F}_k$ generated by p elements w_i chosen randomly independently in the ball of radius n is malcharacteristic (and therefore $\mathrm{Aut}(G, H) = \mathrm{Inn}_H(G)$).

The proof of these results uses Whitehead's peak reduction and equidistribution of subwords, so is specific to free groups.

The proof of Theorem 1.1 uses the connection between automorphisms and splittings (i.e. decompositions of G as the fundamental group of a graph of groups). This is well-known in the context of (relatively) hyperbolic groups since Paulin's paper [22] constructing an action of G on an $\mathbb{R}$-tree for G a hyperbolic group with $\mathrm{Out}(G)$ infinite (one then applies Rips's theory of groups acting on $\mathbb{R}$-trees [3] to get a splitting of G over a virtually cyclic group).

Another key idea of the present paper is a non-splitting principle. Recall that a splitting of G is relative to an element h or a subgroup H if h (or H) is contained in a conjugate of a vertex group (in other words, h or H fixes a point in the Bass–Serre tree).

NON-SPLITTING PRINCIPLE. — If h is a very complicated element of a group G, it is universally hyperbolic: there is no splitting of G relative to h (in other words, if G acts on a tree with no global fixed point, then h does not fix a point).

Unfortunately this is false, even in free groups: given any $h \in \mathbb{F}_k$, there is an epimorphism $\mathbb{F}_k \twoheadrightarrow \mathbb{Z}$ which kills h, hence a splitting of $\mathbb{F}_k$ relative to h (this splitting is over an infinitely generated group, but this may be

remedied using standard approximation techniques, see [15] for instance).
By imposing conditions on edge groups, however, one can get the following
valid version of the non-splitting principle:

THEOREM 1.5 (Corollary 6.6). — *Let G be a non-slender group which
is hyperbolic relative to a finite family of slender subgroups.*

- *Let w_n be given by a random walk on G. With probability going
  to 1 as $n \to \infty$, there is no splitting of G over a slender subgroup
  relative to w_n.*
- *If H is a random subgroup, then with probability going to 1 as
  $n \to \infty$ the group H acts freely in every non-trivial G-tree with
  slender edge stabilizers (as in Theorem 1.1, H is generated by p
  independent random walks as in [18]).*

JSJ decompositions of relatively hyperbolic groups are acylindrical
(see [12]), and this allows us to apply the results of [18]. See Corollary 6.7
for a similar result about torsion-free CSA groups, which also have acylin-
drical JSJ decompositions.

In the case of free groups, it was proved by Cashen–Manning [6] that $\mathbb{F}_k$
has no cyclic splitting relative to an element h represented by a cyclically
reduced word containing all reduced words of length 3 as subwords. The key
technical result used to prove Theorem 1.5 is the following generalization
of this fact.

THEOREM 1.6 (Theorem 3.1). — *Let S be a tree with an action of a
finitely generated group G with slender edge stabilizers.*

*There exist finitely many segments $I_i \subset S$ of length at most 4 with the
following property: if $h \in G$ is hyperbolic in S and its axis contains a
translate of each I_i, then h remains hyperbolic in every non-trivial tree T
with slender edge stabilizers such that edge stabilizers of S are elliptic in T.*

This result applies in particular with S a JSJ tree of G over slender
groups and shows that there is no splitting over a slender group relative to
a sufficiently complicated element.

Our proof of Theorem 1.6 is self-contained and only uses basic Bass–Serre
theory. It is inspired by ideas of Otal [20] and Cashen–Manning.

We also generalize Cashen–Manning's result in the following way.

THEOREM 1.7 (Theorem 4.1). — *Let $k \geqslant 2$ and $L \geqslant 2$. Let $h \in \mathbb{F}_k$
be a cyclically reduced word containing all reduced words of length L as
subwords. If $\mathbb{F}_k$ splits over a subgroup isomorphic to $\mathbb{F}_r$ relative to h, then
$r > (k-1)(L-2)$.*

In other words: if h is complicated, all splittings relative to h are over groups of large rank. The bound is sharp (see Proposition 4.3).

2. Notations and conventions

We will always denote by G a finitely generated group. We consider actions of G on simplicial trees T which are *minimal* (there is no proper invariant subtree). We allow the *trivial* action (T is a point). We write G_v for the stabilizer of a vertex v, and G_e for the stabilizer of an edge e.

We assume that G acts *without inversion* (if $g \in G$ leaves an edge invariant, it fixes its endpoints), and there is no *redundant vertex* (if v is a vertex of valence 2, there is $g \in G$ having v as its unique fixed point).

We equip T with the simplicial metric (every edge has length 1). A *segment I* is the geodesic joining two vertices. The *translates* of I are the segments gI, for $g \in G$.

A *splitting* of G is an isomorphism of G with the fundamental group of a graph of groups Γ, or equivalently an action on a tree T (the Bass–Serre tree of the splitting). Splittings are always assumed to be non-trivial: vertex groups are proper subgroups of G (so T is not a point). A splitting Γ is *over* a subgroup H if H is an edge group. When H is cyclic, Γ is a *cyclic splitting*.

An element $g \in G$, or a subgroup $H \subset G$, is *elliptic* in T if it fixes a point in T. We then say that T (or the corresponding splitting) is *relative* to g or H. If g is not elliptic, it is *hyperbolic* and has an *axis*, a line on which it acts as a translation.

A tree $\widehat{T}$ is a *refinement* of T if one obtains T from $\widehat{T}$ by collapsing each edge belonging to some G-invariant set to a point.

If T, T' are two trees with an action of G, one says that T is *elliptic with respect to T'* if every edge stabilizer of T is elliptic in T'. This implies (see [12, Proposition 2.2]) that T has a refinement $\widehat{T}$ which *dominates* T', in the sense that there exists a G-equivariant map from $\widehat{T}$ to T'.

A group G is *slender* if G and all its subgroups are finitely generated. Equivalently, whenever G acts on a tree, there is a fixed point or an invariant line [9, Lemma 1.1].

A subgroup $H \subset G$ is *almost malnormal* if there exists C such that $gHg^{-1} \cap H$ has cardinality at most C for all $g \notin H$.

We denote by $\mathbb{F}_k$ the free group of rank k. Given a free basis X, a word $w = a_1 \ldots a_q$ with $a_i \in X^{\pm 1}$ is (freely) *reduced* if $a_{i+1} \neq a_i^{-1}$ for $i = 1, \ldots, q-1$, *cyclically reduced* if in addition $a_q \neq a_1^{-1}$. A *subword* of w is a word $a_i \ldots a_j$ with $1 \leqslant i \leqslant j \leqslant q$. The *m-prefix* of w is the word $a_1 \ldots a_m$.

If $w = a_1 \ldots a_q$ is reduced, its *length* $|w|$ is q. In general, we identify a reduced word and the corresponding element of $\mathbb{F}_k$.

Any finitely generated subgroup $H \subset \mathbb{F}_k$ has a *Stallings graph* Θ. It has a base vertex 1, its edges are oriented and labelled by elements of X. The elements of H are precisely the words represented by immersed paths with both endpoints 1. One may construct Θ by letting H act on the Cayley graph $\mathrm{Cay}(\mathbb{F}_k, X)$, restricting to the convex hull of the H-orbit of the base vertex, and taking the quotient by the action of H.

3. A non-splitting theorem

This section is devoted to the proof of Theorem 3.1, which restricts the ways in which a group G may split relative to a complicated enough element h. Theorem 1.5 will be proved in Section 6 by combining Theorem 3.1 with results by Maher–Sisto [18].

THEOREM 3.1. — *Let S be a tree with an action of a finitely generated group G. Assume that S is locally finite or has slender edge stabilizers.*

There exists a finite set $\mathcal{I}$ of segments $I_i \subset S$ of length at most 4 with the following property: if $h \in G$ is hyperbolic in S and its axis contains a translate of each I_i, then h remains hyperbolic in every non-trivial tree T with slender edge stabilizers such that S is elliptic with respect to T.

Remark 3.2.

- Our implicit assumption that S has no redundant vertex is important to bound the length of the I_i's.
- If S has no vertex of valence 2, the I_i's may be taken to be of length at most 3. Applying the theorem to the action of $\mathbb{F}_k$ on its Cayley tree yields Cashen–Manning's theorem [6]: $\mathbb{F}_k$ has no cyclic splitting

relative to a cyclically reduced word h containing all reduced words of length $\leqslant 3$ as subwords.

- The assumption that edge stabilizers of T are slender may be weakened to saying that some edge stabilizer of T is *slender in S*: it fixes a point or leaves a line invariant in S.

Proof. — We may assume that S is not a point or a line: the theorem is trivial if S is a point, easy if S is a line (in this case $S = T$). We may also assume that there is only one orbit of edges in T: if a group E contains a hyperbolic element g and e is an edge in the axis of g, then collapsing all edges except those in the orbit of e yields a tree with one orbit of edges in which g is hyperbolic; one may argue similarly, using a ray going to an end fixed by E, if E contains no hyperbolic element but fixes no point in T (this cannot occur if E is finitely generated).

We start the proof by performing several constructions, starting with a tree T as in the theorem. The assumption that S is elliptic with respect to T implies that there exists a refinement R of S together with an equivariant map $f : R \to T$ (see [12, Proposition 2.2] for instance). We may assume that f sends each vertex to a vertex, and each edge to a point or an edge-path.

We fix an edge $e \subset T$, with midpoint m. We declare one component of $T \setminus \{m\}$ to be positive, the other negative. We consider the set $M = f^{-1}(m)$. It is G_e-invariant, contains no vertex, and we claim that M/G_e is finite. Indeed, up to subdividing R, we may assume that the image of any edge is an edge or a vertex. If $x_1, x_2 \in M$ are in two edges e_1, e_2 belonging to the same G-orbit, say $e_2 = ge_1$, then $x_2 = gx_1$. In particular, g fixes $m = f(x_2) = f(x_1)$, so $g \in G_e$. This proves that the cardinality of M/G_e is bounded by the number of G-orbits of edges of R.

Let $\ell \subset R$ be any proper G_e-invariant subtree. There is one because G_e is slender, hence fixes a point or leaves a line invariant, and S is not a point or a line. For later use (in the proof of Theorem 4.1), we do not assume yet that ℓ is a point or a line.

We fix an integer C such that M is contained in the C-neighborhood ℓ_C of ℓ. Each component of $R \setminus \{\ell_C\}$ is mapped into a single component of $T \setminus \{m\}$, and we label it positive or negative accordingly. Any ray $\rho \subset R$ having compact intersection with ℓ thus inherits a sign (a ray is an isometric image of $[0, +\infty)$).

LEMMA 3.3. — *Let $w \in G$ be hyperbolic in R. Assume that its axis A_w has compact intersection with ℓ, and its ends have different signs. Then w is hyperbolic in T.*

Proof. — Assume that w fixes a vertex x in T, say in the positive component of $T \setminus \{m\}$. Let y be any point of A_w, and $z = f(y)$. Replacing w by w^{-1} if needed, we may assume that $w^n y$ goes to the negative end of A_w as $n \to +\infty$. Then $w^n z = f(w^n y)$ and $w^{n+1} z = w(w^n z)$ are in the negative component of $T \setminus \{m\}$ for n large. Since x is in the positive component and w maps $[x, w^n z]$ to $[x, w^{n+1} z]$, the element w fixes the segment $[x, m]$, in particular it fixes m. We deduce that $w^n z$ is in the negative component for every n (positive or negative), and both ends of the axis of w in R are negative, a contradiction. $\square$

We say that a vertex $p \in \ell$ is a *boundary vertex* if there is at least one edge incident on p which is not contained in ℓ (if ℓ is a line, this means that p is a vertex of ℓ having valence $\geqslant 3$ in R). Given a boundary vertex $p \in \ell$, we consider rays ρ with origin p such that $\rho \cap \ell = \{p\}$. Such rays have a sign (positive or negative). We say that p is *positive* (resp. *negative*) if all rays ρ with origin p such that $\rho \cap \ell = \{p\}$ are positive (resp. negative), and *mixed* otherwise (see Figure 3.1).
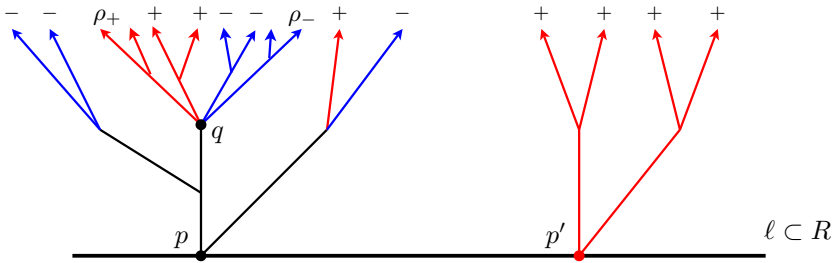

Figure 3.1. p is a mixed vertex, p' is positive.

We shall distinguish two cases.

Case 1: ℓ contains a mixed vertex p. — In this case, we can find a vertex q in ℓ_C, with projection to ℓ equal to p, such that both a positive ray ρ_+ and a negative ray ρ_- with origin p pass through q, and the sign of a ray ρ passing through q only depends on the edge through which ρ exits q (one can take for q a point projecting to p, and furthest from p with the property that there are rays of both signs with origin p passing through q). Note that we allow $p = q$.

Case 2: ℓ has no mixed vertex (this cannot happen if ℓ is a point). — In this case, each boundary vertex $p \in \ell$ inherits a sign, and both signs occur in ℓ.

After these preliminary constructions, we recall that S is locally finite or has slender edge stabilizers. We first suppose that it is locally finite. Vertex stabilizers of S then contain an edge stabilizer with finite index, hence are elliptic in any T as in the theorem. This implies that S dominates T, so we may take $R = S$ and view ℓ as a subtree of S.

We define a finite set $\mathcal{I}$ by choosing a representative for each G-orbit of segments of length 4 in S. We consider T as in the theorem, h whose axis in S contains a translate of each segment of length 4, and we show that h is hyperbolic in T. We have distinguished two cases (depending on T and ℓ).

In Case 1, some translate of the axis of h passes through q and contains the exit edges of ρ_+ and ρ_-. Lemma 3.3 implies that some conjugate of h, hence also h itself, is hyperbolic in T.

In Case 2, we recall that we may take ℓ to be a line, so ℓ contains a positive p_+ and a negative p_- which are at distance 1 or 2 (at distance 1 if all vertices of ℓ are boundary vertices); indeed, since S has no redundant vertex and is not a line, there are no adjacent vertices of valence 2 in S. The intersection of some translate of the axis of h with ℓ is precisely the segment p_+p_-, and hyperbolicity of h follows from Lemma 3.3.

The argument when S has slender edge stabilizers but is not locally finite is more complicated because there may be infinitely many G-orbits of segments of length $\leqslant 4$. Also, we may have to take $R \neq S$ (and R depends on T), but this issue is easily dealt with.

In order to construct a suitable finite family $\mathcal{I}$ (independent of T) we use the case $k = 3$ of the following lemma, whose proof we defer.

LEMMA 3.4. — Let $k \geqslant 1$. Let H be a finitely generated group acting on an infinite set $\mathcal{X}$ with finitely many orbits. Assume that all point stabilizers H_x are slender. The action of H on $\mathcal{X}$ extends to an action on a graph Δ with vertex set $\mathcal{X}$ such that:

- there are finitely many H-orbits of edges in Δ;
- Δ is k-connected: it cannot be disconnected by removing $k - 1$ vertices (we use terminology from graph theory: 1-connected means connected, 3-connected means that there is no separating pair).

Let v be a vertex of S. We consider the action of its stabilizer G_v (which is finitely generated because G and edge stabilizers of S are [7, Lemma 8.32]) on the link L_v of v in S (the set of incident edges). Point stabilizers for this action are edge stabilizers of S, hence slender. We apply the lemma with $k = 3$. We get a graph Δ_v with vertex set L_v and no separating pair (if L_v is finite, we let Δ_v be the complete graph with vertex set L_v). Since

G_v acts on Δ_v, we may perform this construction G-equivariantly for all vertices v of S.

Edges of Δ_v join two elements of the link of v, we view them as segments of length 2 centered at v in S. Considering these segments for every v, we obtain a family of segments of length 2 consisting of finitely many G-orbits (because Δ_v/G_v has finitely many edges and there are finitely many G-orbits of vertices in S), and we include a representative of each orbit in $\mathcal{I}$.

We also consider representatives for G-orbits of edges of S bounded by two vertices having valence at least 3, and for orbits of segments of length 2 whose midpoint has valence 2 (this is a finite set of orbits). For each such ε we choose two extensions $\varepsilon_1\varepsilon\varepsilon_2$ and $\varepsilon_1'\varepsilon\varepsilon_2'$ of ε to segments of length 3 or 4 respectively, with edges $\varepsilon_i \neq \varepsilon_i'$. We then add to $\mathcal{I}$ the four segments $\varepsilon_1\varepsilon\varepsilon_2$, $\varepsilon_1\varepsilon\varepsilon_2'$, $\varepsilon_1'\varepsilon\varepsilon_2$, $\varepsilon_1'\varepsilon\varepsilon_2'$.

Having constructed $\mathcal{I}$, we now consider T as in the theorem and h whose axis in S contains a translate of each I_i in $\mathcal{I}$, and we show that h is hyperbolic in T. We first assume that $R = S$. Since G_e is slender we may assume that ℓ is a point or a line, and we consider the two cases introduced above.

In Case 1, we fix a mixed vertex $p \in \ell$. Recall that we have defined a vertex $q \in \ell_C$ projecting to p. Let L_q be the link of q in S (the set of incident edges). We first define at most two special incident edges at q. If $q \notin \ell$ (i.e. if $q \neq p$), the edge pointing towards p is the only special edge. If $q \in \ell$ and ℓ is a line, both edges contained in ℓ are special. There is no special edge if $q \in \ell$ and ℓ is a point.

Because of the way we defined q, non-special incident edges ζ at q may be given a sign: they are positive or negative, depending on whether rays with origin p exiting q through ζ are positive or negative, and both signs occur.

Using Lemma 3.4, we have constructed a graph Δ_q with vertex set L_q having no separating pair. This graph remains connected when we remove the vertices corresponding to the (at most two) special edges. The remaining vertices correspond to incident edges ζ at q which are positive or negative. Since both signs occur, we may find a positive edge ζ_+ and a negative edge ζ_- which are adjacent in Δ_q. Because of the way we constructed $\mathcal{I}$, some translate of the axis of h in S contains $\zeta_+ \cup \zeta_-$ and Lemma 3.3 implies that h is hyperbolic in T, as required.

In Case 2, as in the locally finite case, ℓ contains a positive p_+ and a negative p_- which are either adjacent or at distance 2 (separated by a

vertex of valence 2). We included four extensions of $\varepsilon = p_+ p_-$ in the G-orbit of $\mathcal{I}$, and one of them at least intersects ℓ only along $p_+ p_-$. Some translate of the axis of h contains this extension, and h is hyperbolic in T by Lemma 3.3.

To complete the proof of Theorem 3.1, we need to consider the case when $R \neq S$. Let $\pi : R \to S$ be a collapse map. Note that, if ε is any open edge of R, both components of $R \setminus \varepsilon$ have unbounded image in S. We define $\bar{\ell} = \pi(\ell)$, a point or a line, and $\bar{\ell}_C$ its C-neighbourhood in S. The sign assignment of components of $R \setminus \ell_C$ induces one for components of $S \setminus \bar{\ell}_C$, with both signs appearing. Lemma 3.3 applies in S because if the axis of w in S has compact intersection with $\bar{\ell}$ and its ends have two different signs, then the same holds for the axis of w in R, so the rest of the proof is the same as when $R = S$. □

Proof of Lemma 3.4. — The proof is by induction on k. If $k = 1$, we just need Δ to be connected. This is easy to achieve, using finite generation of H and finiteness of $\mathcal{X}/H$.

In the general case, we construct $\Delta_1 \subset \Delta_2 \subset \Delta_3 = \Delta$ by successively adding H-orbits of edges (each Δ_i is a graph with vertex set $\mathcal{X}$ on which H acts with finite quotient). At each step we specify a finite set of edges, and we obtain Δ_{i+1} from Δ_i by adding the H-orbits of these edges.

As explained above, we may find a connected graph Δ_1. Given an element $x \in \mathcal{X}$ (which we view as a vertex of Δ_1), we view its link in Δ_1 as the set of vertices adjacent to x. It is itself a graph L_x (possibly with no edge): there is an edge between y and y' in L_x if and only if there is one in Δ_1. The stabilizer H_x acts naturally on this graph L_x.

It is easy to check that L_x/H_x is finite, so by induction we may add finitely many H_x-orbits of edges to L_x in order to make it $(k-1)$-connected (if L_x is finite, we make it a complete graph). We view these added edges as edges between elements of $\mathcal{X}$, and since $\mathcal{X}/H$ is finite we obtain a connected Δ_2 with the property that all links of vertices are $(k-1)$-connected (or complete finite graphs).

We now enlarge Δ_2 in order to obtain Δ_3 with the additional property that each edge is contained in a $(k-1)$-simplex (a complete subgraph with k vertices). We claim that $\Delta = \Delta_3$ is then k-connected.

Fix a subset $\mathcal{X}_0$ of cardinality $k-1$ in $\mathcal{X}$. We must be able to join any two vertices x, y in $\mathcal{X} \setminus \mathcal{X}_0$ by a path in Δ avoiding $\mathcal{X}_0$. Since Δ is connected, we may find a path from x to y. It suffices to consider the case when this path is of the form $xz_1 \ldots z_p y$ with the z_i's distinct elements of $\mathcal{X}_0$.

First suppose $p = 1$, so that x and y belong to the link of z_1, which is $(k-1)$-connected. The intersection of this link with $\mathcal{X}_0$ has cardinality at most $k-2$, so we may join x to y in the complement of $\mathcal{X}_0$. If $p \geqslant 2$, we consider the edge $z_1 z_2$. It is contained in a $(k-1)$-simplex, which has k vertices so contains a vertex $z \notin \mathcal{X}_0$. We then replace the path $x z_1 \ldots z_p y$ by the concatenation of $x z_1 z$ and $z z_2 \ldots z_p y$ and use induction on p. $\square$

4. Splittings of free groups

We view $\mathbb{F}_k$ as the set of reduced words on a set X of cardinality k.

THEOREM 4.1. — *Let $k \geqslant 2$ and $L \geqslant 2$. Let $h \in \mathbb{F}_k$ be a cyclically reduced word containing all reduced words of length L as subwords. If h is elliptic in a splitting of $\mathbb{F}_k$ over a subgroup isomorphic to $\mathbb{F}_r$, then $r > (k-1)(L-2)$.*

In other words: if h is complicated, all splittings relative to h are over groups of large rank.

When $L = 2$, the theorem says that h is not contained in a proper free factor, a result due to Whitehead. When $L = 3$ there is no splitting of $\mathbb{F}_k$ relative to h over $\mathbb{F}_r$ if $r \leqslant k-1$ (the case $r = 1$ is due to Cashen–Manning [6]).

Proof. — We argue as in the proof of Theorem 3.1, with S the Cayley graph $\mathrm{Cay}(\mathbb{F}_k, X)$ (a locally finite tree), T the Bass–Serre tree of a splitting over $\mathbb{F}_r$, and e an edge of T (note that $G_e \simeq \mathbb{F}_r$ is not slender if $r \geqslant 2$). We let $\ell \subset S$ be any point in S if G_e is trivial, the minimal G_e-invariant subtree otherwise (it is a proper subtree because G_e has infinite index: otherwise G would fix a point in T and the splitting would be trivial).

As in the proof of Theorem 3.1, we distinguish Case 1 and Case 2. In Case 1 (there exists a mixed vertex in ℓ), no new argument is needed since we assume $L \geqslant 2$. In Case 2, all boundary vertices of ℓ are positive or negative, but we can no longer find boundary vertices p_+, p_- with distance at most 2 (this required ℓ to be a line). In fact, if h as in the theorem is elliptic in T, any boundary vertices p_+, p_- of opposite signs must be at least $(L-1)$-apart: otherwise the axis of a conjugate of h intersects ℓ precisely in the segment $p_+ p_-$, so h is hyperbolic in T by Lemma 3.3.

Choose a pair of boundary vertices of opposite signs $p_+, p_- \in \ell$ whose distance D is minimal. We have seen $D \geqslant L-1$. Every vertex of ℓ which is not a boundary vertex has valence $2k$ in ℓ, so all vertices between p_+ and p_- have valence $2k$ in ℓ. The quotient map from ℓ to ℓ/G_e (a regular

covering with group G_e) is injective on the segment $p_+ p_-$: otherwise the images of p_+, p_- in ℓ/G_e have distance less than D, so some pair $p_+, g p_-$ with $g \in G_e$ contradicts our choice of D because the sign assignment is G_e-invariant. The quotient graph ℓ/G_e therefore has at least $D-1$ vertices with valence $2k$. Since it has no vertex of valence 1 and its fundamental group has rank r, we get $r > (k-1)(D-1) \geqslant (k-1)(L-2)$. $\qquad\square$

The theorem may be generalized, for instance to the following statement.

THEOREM 4.2. — *Let G be a finitely generated group acting on a tree S with finite stabilizers (so G is virtually free). Let L be an integer, and $\mathcal{H}$ a family of elements of G such that every segment of length $\leqslant L$ in S is contained in a translate of the axis of an element of $\mathcal{H}$.*

If G splits relative to $\mathcal{H}$ over a subgroup which is virtually $\mathbb{F}_r$, then $r \geqslant L/4$ (and $r \geqslant L/2$ if S has no vertices of valence 2).

We leave details to the reader.

We now show that the bound in Theorem 4.1 is optimal (at least for L even).

PROPOSITION 4.3. — *For each $k \geqslant 2$ and each even $L = 2i \geqslant 2$, there is a splitting of $\mathbb{F}_k$ over a group of rank $r = (k-1)(L-2)+1$ relative to a cyclically reduced h containing all reduced words of length L.*

Remark 4.4. — We are not sure of optimality for L odd. For instance, there seems to be no splitting of $\mathbb{F}_2$ over $\mathbb{F}_2$ relative to an h containing all reduced words of length 3.

LEMMA 4.5 (see Figure 4.1). — *Fix $k \geqslant 2$. For $i = 1, 2, \ldots$, there are subgroups A_i, C_i, and splittings of $\mathbb{F}_k$ as $A_i *_{C_i} A_{i+1}$, such that:*

- A_i *has rank $i(k-1)$;*
- C_i *has index 2 in A_i, hence has rank $2i(k-1)-1$;*
- *all reduced words of length i may be read as labels of paths in the Stallings graph of A_{i+1} starting at the base vertex.*

The last item ensures that A_{i+1} contains cyclically reduced elements containing all reduced words of length $L = 2i$, so the proposition follows from the lemma.

The Stallings graphs of the groups A_i and C_i are pictured on Figure 4.1 in the case of $\mathbb{F}_2 = \langle a, b \rangle$. One easily checks that $C_i \subset C_{i+1}$, and that $A_{i+1} = \langle A_{i-1}, C_i \rangle$. The initial splitting is $A_1 *_{C_1} A_2 = \langle b \rangle *_{\langle b^2 \rangle} \langle a, b^2 \rangle$. For all i one obtains the splitting $A_i *_{C_i} A_{i+1}$ from $A_i *_{C_{i-1}} A_{i-1}$ by folding $C_i < A_i$ along the edge, thus replacing C_{i-1} by C_i and A_{i-1} by $\langle A_{i-1}, C_i \rangle = A_{i+1}$.
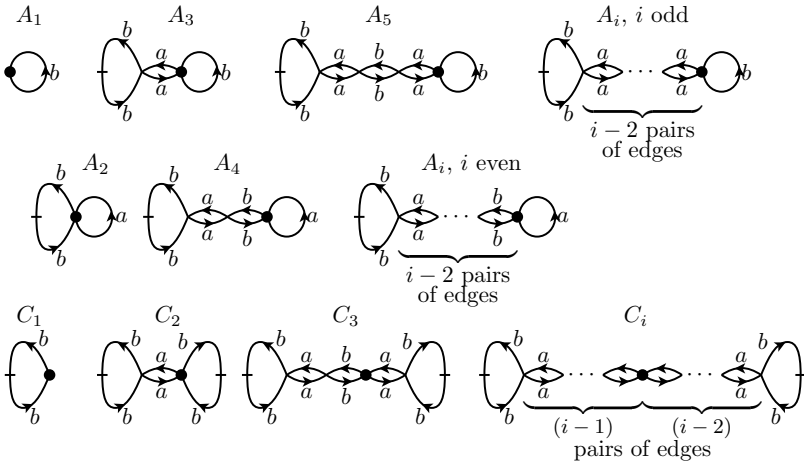
Figure 4.1. The Stallings graphs of the subgroups A_i, C_i of $\mathbb{F}_2$ (the bullet represents the base vertex).

The reader may check that these splittings have the required properties. For $k > 2$, one adds $k - 2$ loops labelled by the extra generators at each vertex of each Stallings graph.

5. Random walks (after Maher–Sisto [18])

DEFINITION 5.1 (Random subgroup, random element). — *Let G be a finitely generated group. Let μ be a probability measure on G whose support is finite and generates G as a semigroup. We fix $p \geqslant 1$, and we consider a subgroup $H \subset G$ generated by p elements $w_{1,n}, \ldots, w_{p,n}$ arising from independent random walks of length n generated by μ. We call H a random subgroup of G. When $p = 1$, we call $w_n = w_{1,n}$ a random element.*

Remark 5.2. — The assumptions on μ and H may be weakened to those of [18].

Recall that G acts *acylindrically* on a tree S if there exist numbers K and C such that stabilizers of segments of length K have cardinality at most C (this is sometimes called almost acylindrical, and agrees with the general definition of acylindricity given in [18]).

The following theorem will ensure that Theorem 3.1 applies to non-trivial elements of random subgroups if S is acylindrical.

THEOREM 5.3. — *Assume that G is not virtually cyclic and acts acylindrically on a non-trivial tree S. Let $\mathcal{I}$ be a finite family of segments $I_i \subset S$. Let $H = \langle w_{1,n}, \dots, w_{p,n} \rangle$ be a random subgroup as in Definition 5.1.*

With probability going to 1 as $n \to \infty$, the group H is freely generated by $w_{1,n}, \dots, w_{p,n}$, the action of H on S is free, and the axis of any non-trivial $h \in H$ contains a translate of each I_i.

We explain how to derive this theorem from [18]. Since the action on S is acylindrical and G is not virtually cyclic, S is irreducible (there is no fixed point, no fixed end, no invariant line), so the action is non-elementary in the sense of [18]. By the main theorem of [18], the $w_{i,n}$'s freely generate H with probability going to 1, and $HE(G)$ is hyperbolically embedded in G (with $E(G)$ the maximal finite normal subgroup of G).

Choose a basepoint $x_0 \in S$, and fix a hyperbolic element $g \in G$ such that some fundamental domain for the action of g on its axis contains a translate of each I_i (one finds such a g by applying Lemma 4.3 of [21] inductively).

We first consider the case $p = 1$ and we let γ_n be the segment between x_0 and $w_n x_0$. Applying Proposition 3.2.4 of [18] with $\varepsilon = 1/4$ and L large with respect to the constant K_0 and the translation length of g, we deduce that the middle half of γ_n contains a translate of each I_i with probability going to 1 (we note that Proposition 3.2.4 of [18] actually assumes that g is in the support of μ, but taking n_0 such that g is in the support of μ^{*n_0}, the result holds for n a multiple of n_0, and the general case follows easily). By Proposition 3.2.5 of [18], this also holds for the axis of w_n.

For $p > 1$, we consider the smallest H-invariant subtree $S_H \subset S$ containing x_0. It follows from Propositions 6.3 and 6.5 of [18] that, with probability going to 1, the action of H on S_H is free, and the quotient looks like a rose: it is the union of a central tree C with diameter $< \varepsilon n$ (for some arbitrarily small $\varepsilon > 0$) and p arcs $\theta_1, \dots, \theta_p$ of length $> (L - \varepsilon)n$ attached to C (with $L > 0$ the drift of the random walk), and moreover the image of the axis of $w_{i,n}$ in S_H/H is the union of θ_i with an arc contained in C (compare the central tree property, see e.g. [2, Section 3.1] and Subsection 7.2.1). The image of the axis of any non-trivial $h \in H$ in S_H/H contains one of the θ_i's, and the result follows.

6. Non-splitting relative to random elements

One basic theme of this paper is that a group has no non-trivial splitting relative to a random element (or a random subgroup). As explained in

the introduction, one must impose restrictions on the edge groups of the splitting. In the case of $\mathbb{F}_k$, combining Theorems 4.1 and 5.3 yields:

THEOREM 6.1. — *Fix $r \geqslant 1$. Let $w_n \in \mathbb{F}_k$ be a random element as in Definition 5.1. With probability going to 1 as $n \to \infty$, there is no splitting of $\mathbb{F}_k$ relative to w_n over a group of rank at most r.*

Proof. — Choose L such that $r \leqslant (k-1)(L-2)$. Apply Theorem 5.3 to the action of $\mathbb{F}_k$ on its Cayley tree. With probability going to 1, the axis of w_n contains (translates of) all segments of length L, so w_n is hyperbolic in every splitting over a group of rank $\leqslant r$ by Theorem 4.1. $\qquad\square$

The following related result was suggested by the referee.

THEOREM 6.2. — *Consider a free group $\mathbb{F}_k$, equipped with a free basis. There exists $\varepsilon > 0$ with the following property: with probability going to 1 as $n \to \infty$, there is no splitting of $\mathbb{F}_k$ over a subgroup of rank at most $\varepsilon \log n$ relative to an element w chosen randomly in the ball of radius n, or given by a simple random walk of length n with respect to a free basis.*

Sketch of proof. — Arguing as in Subsection 7.2, one shows that, for ε small enough, a random word w in the sphere of radius n contains all subwords of length $\varepsilon \log n$ with probability going to 1. It easily follows that this also holds for a random word in the ball of radius n or for a word sampled from a simple random walk of length n. One then applies Theorem 4.1. $\qquad\square$

In more general finitely generated groups, we get:

THEOREM 6.3. — *Assume that G is not virtually cyclic and acts acylindrically on a non-trivial tree S which is locally finite or has slender edge stabilizers. Let H be a random subgroup as in Definition 5.1. With probability going to 1 as $n \to \infty$, the group H acts freely in every non-trivial tree T with slender edge stabilizers such that S is elliptic with respect to T.*

Proof. — Use Theorem 5.3 and apply Theorem 3.1 to all non-trivial elements of H. $\qquad\square$

We refer the reader to [12] for details about the JSJ decompositions used in the next results.

COROLLARY 6.4. — *Assume that G is finitely presented, not virtually $\mathbb{Z}^2$, and the JSJ decomposition of G over virtually cyclic subgroups is acylindrical. Let w_n be a random element of G. With probability going to 1 as $n \to \infty$, there is no splitting of G over a virtually cyclic subgroup relative to w_n.*

Proof. — Let S be an acylindrical JSJ tree over virtually cyclic subgroups. By definition of the JSJ decomposition, S is elliptic in every tree T with virtually cyclic edge stabilizers, so the result follows by applying Theorem 6.3 to S, provided that S is not a trivial tree (a point).

If S is a point, there are two cases: rigid or flexible (see [12, Definition 2.14]). In the rigid case, G does not split over a virtually cyclic group so the theorem is empty. In the flexible case, it follows from Theorem 6.2 and Proposition 6.38 of [12] that either G is virtually $\mathbb{Z}^2$, contrary to our hypothesis, or G maps onto the fundamental group of a closed hyperbolic 2-orbifold Σ with finite kernel (it is QH with finite fiber, see [12, Theorem 6.2]).

For simplicity we assume that Σ is a surface rather than an orbifold. We apply [18] to the action of G on the hyperbolic plane $\mathbb{H}^2$ (viewed as the universal cover of Σ). We fix a closed geodesic γ which fills Σ. By Propositions 3.2.4 and 3.2.5 of [18], there exists a constant K such that, for any compact segment $A \subset \gamma$, the axis of w_n is K-close to some translate of A with high probability. If A is chosen long enough, this implies that the closed geodesic representing (the conjugacy class of the image of) w_n meets every simple closed geodesic δ, so w_n is hyperbolic in the splitting of G dual to δ. The result follows since every splitting of G over a slender group is dual to a simple closed geodesic of Σ (see e.g. [12, Sections 5.1.2 and 5.2]). □

Remark 6.5. — The same argument shows that, if G is finitely presented, not slender, and its slender JSJ decomposition is acylindrical, then, with probability going to 1 as $n \to \infty$, there is no splitting of G over a slender subgroup relative to w_n.

COROLLARY 6.6. — *Let G be a non-slender group which is hyperbolic relative to a finite family of slender subgroups P_i.*

- *Let w_n be a random element of G. With probability going to 1 as $n \to \infty$, there is no splitting of G over a slender subgroup relative to w_n.*
- *If H is a random subgroup, then with probability going to 1 as $n \to \infty$ the group H acts freely in every non-trivial tree with slender edge stabilizers.*

Proof. — We let S be a JSJ tree over slender groups relative to the parabolic subgroups P_i (which we may assume not to be virtually cyclic), see [12, in particular Theorem 9.18 and Corollary 4.16]. It is 2-acylindrical (stabilizers of segments of length 3 are finite with bounded cardinality), and

its edge stabilizers are elliptic in every tree T with slender edge stabilizers. As in the previous proof, we apply Theorem 6.3 to S. In the flexible case, G is QH with finite fiber by Theorem 9.18 of [12]. $\qquad\square$

COROLLARY 6.7. — *Let G be a torsion-free CSA group. Let w_n be a random element of G. With probability going to 1 as $n \to \infty$, there is no splitting of G over a finitely generated abelian subgroup relative to w_n.*

Recall that a group is CSA if its maximal abelian subgroups are malnormal.

Proof. — If all abelian subgroups of G are finitely generated (hence slender), the proof is the same as that of the previous corollary, using a JSJ decomposition over abelian groups relative to all non-cyclic abelian subgroups (see [12, Theorem 9.5]).

In general, we apply Corollary 9.1 of [12] with $\mathcal{A}$ the family of all finitely generated abelian subgroups, $\mathcal{S}$ the family of all abelian subgroups (note that conditions (4b) and (4c) of the corollary are satisfied), and $\mathcal{H} = \emptyset$. We obtain a tree $S = (T_a)_c^*$ which is a JSJ tree over $\mathcal{A}$ (hence has finitely generated edge stabilizers) relative to all non-cyclic abelian subgroups. It is compatible with every tree T with edge stabilizers in $\mathcal{A}$, in particular it is elliptic with respect to T, and we can argue as before. $\qquad\square$

Remark 6.8. — If G is finitely presented, there is no splitting T of G relative to w_n over *any* abelian subgroup. To see this, we apply Theorem 6.36 of [12], with $\mathcal{A}$ the family of *all* abelian subgroups and $\mathcal{H} = \emptyset$. By Theorem 2.20 of [12], there is a JSJ tree S over $\mathcal{A}$ with finitely generated (hence slender) edge stabilizers. The edge stabilizers of T do not have to be slender, but they are slender in S and we use Remark 3.2.

COROLLARY 6.9. — *Assume that G has infinitely many ends, and let w_n be a random element of G. With probability going to 1 as $n \to \infty$, there is no splitting of G over a slender subgroup relative to w_n.*

Proof. — Apply Theorem 6.3 with S a tree with finite edge stabilizers (such an S is elliptic with respect to any T). $\qquad\square$

Remark 6.10. — The result remains true if the edge group of the splitting is only assumed not to split over a finite group.

COROLLARY 6.11. — *Assume that G splits over a slender almost malnormal subgroup H, and let w_n be a random element of G. With probability going to 1 as $n \to \infty$, there is no splitting of G over a slender subgroup relative to w_n and H.*

Recall that H is almost malnormal if there exists C such that $gHg^{-1} \cap H$ has cardinality at most C for all $g \notin H$.

Proof. — Apply Theorem 6.3 with S the given splitting of G over H. Almost malnormality of H implies that it is acylindrical. Edge stabilizers of S are conjugate to H, hence elliptic in any tree T relative to H.   □

7. Automorphisms

We now prove several results saying that few automorphisms of a given group G leave a random subgroup invariant. We shall consider relatively hyperbolic groups, before focusing on the specific case of free groups.

Before doing that, we note the following consequence of Theorem 7.6 or 7.14 of [11].

THEOREM 7.1. — *Let G be a hyperbolic group, and $g \in G$ an element of infinite order. If G does not split relative to g over a virtually cyclic group with infinite center, the stabilizer of g in $\mathrm{Aut}(G)$ is virtually cyclic (it is virtually generated by the conjugation by g).*

In particular, using [6], we see that, if $w \in \mathbb{F}_k$ is represented by a cyclically reduced word containing every reduced word of length 3, then the stabilizer of w in $\mathrm{Aut}(\mathbb{F}_k)$ is virtually cyclic.

7.1. Relatively hyperbolic groups

We refer to [5, 19] for basic notions concerning relatively hyperbolic groups.

THEOREM 7.2. — *Assume that G is hyperbolic relative to a finite family $\mathcal{P}$ of slender subgroups. Let $H = \langle w_{1,n}, \ldots, w_{p,n} \rangle$ be a random subgroup as in Definition 5.1. With probability going to 1 as $n \to \infty$, the subgroup $\mathrm{Inn}_H(G) \in \mathrm{Aut}(G)$ generated by conjugations by elements of H has finite index in the group $\mathrm{Aut}(G, H)$ of automorphisms of G leaving H invariant.*

The proof requires a lemma. A group P is *small* if it does not contain $\mathbb{F}_2$.

LEMMA 7.3. — *Let $\mathcal{P}_0$ be a finite family of small finitely generated subgroups P_i which are not virtually cyclic. Assume that G is hyperbolic relative to $\mathcal{P}_0$, and also relative to $\mathcal{P}_0 \cup \{H\}$ with H infinite and finitely generated. Also assume that non-small subgroups of H have finite centralizer (in H hence also in G).*

If $\mathrm{Inn}_H(G)$ has infinite index in $\mathrm{Aut}(G, H)$, then G has a splitting over a small group, and this splitting is relative to the P_i's and to some $H_0 \subset H$ which is equal to H or contains $\mathbb{F}_2$.

If $\mathcal{P}_0 = \emptyset$ and $H = \{1\}$, the lemma reduces to the standard statement that a hyperbolic group G with $\mathrm{Out}(G)$ infinite splits over a small group.

Proof. — Our assumptions imply that the group $\mathrm{Out}(G, \mathcal{P}_0 \cup \{H\}) \subset \mathrm{Out}(G)$ of outer automorphisms sending the P_i's and H to conjugates is infinite. This is because H is infinite and almost malnormal (so inner automorphisms of G leaving H invariant are conjugations by elements of H), and every automorphism of G maps P_i to a conjugate of some P_j (so $\mathrm{Out}(G, \mathcal{P}_0)$ has finite index in $\mathrm{Out}(G)$).

We view G as hyperbolic relative to $\mathcal{P}_0 \cup \{H\}$, and we apply Corollary 7.13 of [11] with $\mathcal{P} = \mathcal{P}_0 \cup \{H\}$ and $\mathcal{H}$ empty. We get a graph of groups decomposition Γ of G relative to the P_i's and H, with edge groups small or contained in H (up to conjugacy). The lemma is proved if some edge group of Γ is small, so we assume that all edge groups are conjugate to subgroups of H.

The assumption about centralizers implies that the group of twists of Γ is finite (see [11] for definitions not given here). By Corollary 7.13 of [11], infiniteness of $\mathrm{Out}(G, \mathcal{P}_0 \cup \{H\})$ implies that Γ has a vertex group G_v with $\mathrm{Out}(G_v; \mathrm{Inc}_v^{(t)})$ infinite: G_v has infinitely many outer automorphisms acting on incident edge groups as conjugations by elements of G_v.

Corollary 7.13 of [11] also implies that G_v is a maximal parabolic subgroup. It has to be conjugate to H: otherwise it would be conjugate to some P_i, and incident edge groups would be small.

The automorphisms in $\mathrm{Out}(G_v; \mathrm{Inc}_v^{(t)})$ extend to G and we get that $\mathrm{Out}(G, \mathcal{P}_0) \cap \mathrm{Out}(G, \{H_0\}^{(t)})$ is infinite for some non-small $H_0 \subset H$ (an incident edge group at v). We now view G as hyperbolic relative to $\mathcal{P}_0$ only, and we apply Corollary 7.13 of [11] with $\mathcal{P} = \mathcal{P}_0$ and $\mathcal{H} = \{H_0\}$. We get a splitting of G which is relative to the P_i's and H_0, over a group which is virtually cyclic or contained in some P_i (up to conjugacy), hence small. $\square$

Proof of Theorem 7.2. — First assume that G is torsion-free. By [18], with probability going to 1, the group H is free and malnormal. In order to apply Lemma 7.3, we just need to check that G is hyperbolic relative to $\mathcal{P} \cup \{H\}$ (we may assume with no loss of generality that no $P_i \in \mathcal{P}$ is virtually cyclic).

Recall [8, Proposition 4.28] that G being hyperbolic relative to $\mathcal{P}$ is equivalent to $\mathcal{P}$ being hyperbolically embedded in (G, X) (with X a finite generating set of G relative to $\mathcal{P}$). By Theorem 2.5 of [18], the group H is

quasi-isometrically embedded and geometrically separated in $\mathrm{Cay}(G, X \cup \mathcal{P})$ with probability going to 1, hence ([1, Theorem 3.9]) $\mathcal{P} \cup \{H\}$ is hyperbolically embedded in (G, X), i.e. G is indeed hyperbolic relative to $\mathcal{P} \cup \{H\}$ (since X is finite).

If Theorem 7.2 is false, Lemma 7.3 provides a splitting of G over a slender subgroup which contradicts Corollary 6.6.

We now allow torsion. Let $E(G)$ be the maximal finite normal subgroup of G, and $\overline{H} = HE(G)$. With probability going to 1 it is virtually free (hence satisfies the condition on centralizers in Lemma 7.3) and G is hyperbolic relative to $\mathcal{P} \cup \{\overline{H}\}$ as above. Since $\mathrm{Aut}(G, H) \subset \mathrm{Aut}(G, \overline{H})$ (because $E(G)$ is characteristic) and H has finite index in $\overline{H}$, it suffices to prove that $\mathrm{Inn}_{\overline{H}}$ has finite index in $\mathrm{Aut}(G, \overline{H})$. If this does not hold, Lemma 7.3 provides a splitting relative to some infinite subgroup $\overline{H}_0 \subset \overline{H}$. Since $\overline{H}_0 \cap H$ is non-trivial and fixes a point in this splitting, this contradicts Corollary 6.6. $\qquad \square$

7.2. Free groups

Let $\mathbb{F}_k$ be a free group of rank $k \geqslant 2$. We denote by ad_h the inner automorphism $g \mapsto hgh^{-1}$.

DEFINITION 7.4 ([16]). — *A subgroup $H < \mathbb{F}_k$ is malcharacteristic if, for any $\alpha \in \mathrm{Aut}(F_k)$ such that $\alpha(H) \cap H \neq \{1\}$, there exists $h \in H$ such that $\alpha = \mathrm{ad}_h$.*

Clearly, if H is malcharacteristic and non-trivial, then the only automorphisms of G preserving H are conjugations by elements of H. With the notation of Theorem 7.2, this says that $\mathrm{Aut}(G, H) = \mathrm{Inn}_H(G)$ (exactly, not up to finite index).

In this section, we prove two results saying that random subgroups of the free group are malcharacteristic, one for groups generated by elements chosen randomly independently in a ball of large radius, and one for groups generated by elements coming from independent simple random walks.

We fix a free basis X of $\mathbb{F}_k$. We view elements $g \in \mathbb{F}_k$ as reduced words in $X^{\pm 1}$, and we write $|g|$ for the length of g. Balls are defined using the generating set $X^{\pm 1}$, and we consider the simple random walk where $w_n = s_1 \cdots s_n$ with $s_1, \ldots, s_n$ chosen randomly and independently in $X^{\pm 1}$ (equipped with the uniform measure).

We say that an event occurs *with probability going to 1 exponentially fast* as $n \to +\infty$ if the probability that it does not occur is bounded by $C\kappa^{-n}$ for some constants $C, \kappa > 0$.

THEOREM 7.5. — *Fix $k \geqslant 2$ and $p \geqslant 1$. With probability going to 1 exponentially fast as $n \to +\infty$, the subgroup $H \subset \mathbb{F}_k$ generated by p elements w_i chosen randomly independently in the ball of radius n is malcharacteristic.*

See Theorem 8.5 and Proposition 8.7 of [14] for the case $p = 1$.

THEOREM 7.6. — *Fix $k \geqslant 2$ and $p \geqslant 1$. Let $H = \langle w_{1,n}, \ldots, w_{p,n} \rangle$ be the subgroup generated by p elements $w_{1,n}, \ldots, w_{p,n}$ arising from independent simple random walks of length n in $\mathbb{F}_k$.*

With probability going to 1 exponentially fast as $n \to +\infty$, the subgroup H is malcharacteristic.

COROLLARY 7.7. — *In the setting of Theorems 7.5 and 7.6, the only automorphisms of $\mathbb{F}_k$ preserving H are conjugations by elements of H.*

Both theorems are special cases of the following general statement.

PROPOSITION 7.8. — *Fix $k \geqslant 2$ and $p \geqslant 1$. Let $w_{1,n}, \ldots, w_{p,n}$ be independent random variables in $\mathbb{F}_k$ satisfying the following conditions:*

- *(Radial symmetry) Given n and i, the probability that $w_{i,n} = g$ only depends on the length of the element $g \in \mathbb{F}_k$.*
- *(Positive drift) There exists $L > 0$ such that, for each i, the probability that $|w_{i,n}| > Ln$ goes to 1 exponentially fast as $n \to +\infty$.*
- *(Subexponential growth) For any $\theta > 0$, the probability that $|w_{i,n}| \leqslant e^{\theta n}$ goes to 1 exponentially fast as $n \to +\infty$.*

Then, with probability going to 1 exponentially fast as $n \to +\infty$, the subgroup $H = \langle w_{1,n}, \ldots, w_{p,n} \rangle$ is malcharacteristic.

This proposition clearly implies Theorems 7.5 and 7.6 (it is well-known that the simple random walk on $\mathbb{F}_k$ has positive drift $1 - \frac{1}{k}$).

We shall now prove the proposition. For simplicity, we sometimes write *generically* to mean with probability going to 1 exponentially fast as $n \to +\infty$.

Many arguments already appear in [2] or [14], but we do not have information about the distribution of the lengths $|w_{i,n}|$, so we will have to use a Fubini-type argument, working with spheres rather than balls (this is made possible by radial symmetry). More precisely:

LEMMA 7.9. — *Let C, θ, κ be positive numbers. Suppose that, given n and numbers $A_1, \ldots, A_p$ with $Ln \leqslant A_i \leqslant Ce^{\theta n}$, the probability that words $w_1, \ldots, w_p$ with $|w_i| = A_i$ (chosen uniformly independently on spheres of radius A_i) satisfy a given property is at least $1 - C\kappa^{-n}$ (independently of*

the A_i's). Then the words $w_{1,n}, \ldots, w_{p,n}$ satisfy the property with probability going to 1 exponentially fast as $n \to +\infty$.

Proof. — This follows from radial symmetry, since $Ln \leqslant |w_{i,n}| \leqslant Ce^{\theta n}$ holds generically by positive drift and subexponential growth. □

7.2.1. The central tree property (see for instance [2])

Let $w_{1,n}, \ldots, w_{p,n}$ be as in the proposition. We fix n, and we let Θ be the Stallings graph of H. The elements $w_{i,n}$, indeed all elements of H, are represented by immersed paths with both endpoints the base vertex 1.

The central tree property says that, generically, the graph Θ looks like a rose. The *m-prefix* of a word is its initial subword of length m.

LEMMA 7.10. — *Fix $\lambda < \frac{L}{2}$. With probability going to 1 exponentially fast as $n \to +\infty$, the $2k$ elements $w_{i,n}^{\pm 1}$ have length $\geqslant \lambda n$ and have distinct λn-prefixes.*

We shall consistently neglect the fact that numbers such that λn are not necessarily integers (so that we should write $[\lambda n]$ instead).

DEFINITION 7.11 (Central tree, outer loops). — *Viewing the words $w_{i,n}^{\pm 1}$ as loops based at 1 in Θ, their initial segments of length λn are all distinct, so form a central tree $\mathcal{C} \subset \Theta$ with $2p$ or $2p + 1$ leaves (1 may be a leaf). The complement of $\mathcal{C}$ in Θ consists of p arcs of length $> (L - 2\lambda)n$ called the outer loops.*

Proof of Lemma 7.10. — Fix n and $A \geqslant \lambda n$ (we do not use subexponential growth in this proof). The number of reduced words of length A is $\gamma_A = 2k(2k - 1)^{A-1}$. Among those, the number of words w such that w and w^{-1} have the same λn-prefix is at most $\widetilde{\gamma}_{A;n} = 2k(2k - 1)^{A-1-\lambda n}$, since w is completely determined by its $(A - \lambda n)$-prefix. The probability that an element w chosen at random among elements of length A has the same λn-prefix as w^{-1} is therefore bounded by $\widetilde{\gamma}_{A;n}/\gamma_A = (2k - 1)^{-\lambda n}$.

Since $(2k - 1)^{-\lambda n}$ goes to 0 exponentially fast as $n \to \infty$, Lemma 7.9 implies that, for each i, the λn-prefixes of $w_{i,n}$ and $w_{i,n}^{-1}$ are generically different.

The argument for $w_{i,n}$ and $w_{j,n}^{\pm 1}$ is similar. Fix $A_1, \ldots, A_p$ bigger than λn. The number $\gamma_{A_1, \ldots, A_p; n}$ of p-tuples $(w_1, \ldots, w_p)$ with $|w_i| = A_i$ is $(2k)^p \prod_{\alpha=1}^{p} (2k - 1)^{A_\alpha - 1}$. The number $\widetilde{\gamma}_{A_1, \ldots, A_p; i, j, n}$ of those for which w_i and $w_j^{\pm 1}$ have the same λn-prefix is bounded by twice the same product, but with the term $(2k - 1)^{A_j - 1}$ replaced by $(2k - 1)^{A_j - 1 - \lambda n}$, so the ratio $\widetilde{\gamma}_{A_1, \ldots, A_p; i, j, n}/\gamma_{A_1, \ldots, A_p; n}$ is bounded by $2(2k - 1)^{-\lambda n}$. □

7.2.2. Whitehead minimality

As in [14], we use Whitehead's peak reduction. We refer to [17, Section I.4] for the basic definitions and results of this theory (the reader unfamiliar with it may skip the definitions and simply combine Lemma 7.13 and Proposition 7.17).

It is now more convenient to work with cyclically reduced elements. If $g = s_1 \ldots s_l$ is a cyclically reduced word with $s_i \in X^{\pm 1}$, its cyclic permutations are the words $s_i \ldots s_l s_1 \ldots s_{i-1}$. As in [17], the set of all cyclic permutations of g is called a cyclic word, it corresponds to a conjugacy class in $\mathbb{F}_k$.

A *relabeling automorphism* of $\mathbb{F}_k$ is an automorphism preserving $X^{\pm 1}$.

DEFINITION 7.12 (Strictly Whitehead minimal, [14, Definition 1.3]). — *A cyclically reduced element $g \in \mathbb{F}_k$ is strictly Whitehead minimal if $|\varphi(g)| > |g|$ for any Whitehead automorphism φ which is not inner and is not a relabeling automorphism.*

A cyclically reduced word $g \in \mathbb{F}_k$ is strictly Whitehead minimal if and only if all its cyclic conjugates are. We thus say that the corresponding cyclic word is strictly Whitehead minimal.

We use peak reduction in the following form.

LEMMA 7.13 (see [14, Proposition 4.3]). — *Strictly Whitehead minimal elements have minimal length in their $\mathrm{Aut}(\mathbb{F}_k)$-orbit. If two cyclically reduced elements $g, h \in \mathbb{F}_k$ are strictly Whitehead minimal and $\alpha(g) = h$ for some $\alpha \in \mathrm{Aut}(\mathbb{F}_k)$, then α is the composition of an inner automorphism and a relabeling automorphism; one passes from g to h by a cyclic permutation and a relabeling.*

7.2.3. Equidistribution

As observed in [14], one deduces from Proposition I.4.16 of [17] that a word g is strictly Whitehead minimal if all words of length 2 in $X^{\pm 1}$ appear with approximately the same frequency in g.

Given a freely reduced word $g = s_1 \ldots s_\ell$ with $s_i \in X^{\pm 1}$, and a letter $u \in X^{\pm 1}$, let
$$P_u(g) = \frac{1}{\ell} \#\{i \leqslant \ell \mid s_i = u\}$$
be the proportion of u's among the letters of g.

Given a couple of letters $(u, v) \in X^{\pm 1} \times X^{\pm 1}$ with $u \neq v^{-1}$, we also define
$$P_{uv}(g) = \frac{1}{\ell - 1} \#\{i \leqslant \ell - 1 \mid s_i = u, s_{i+1} = v\},$$

the frequency of uv in g.

We define $P_u(g)$ and $P_{uv}(g)$ similarly if g is a cyclic word, except that we agree that $s_{\ell+1} = s_1$ and we define $P_{uv}(g) = \frac{1}{\ell}\#\{i \leqslant \ell \mid s_i = u, s_{i+1} = v\}$.

DEFINITION 7.14 (ε-equidistributed). — *Given $\varepsilon > 0$, say that a reduced word $g \in \mathbb{F}_k$ (or a cyclic word representing a conjugacy class in $\mathbb{F}_k$) is ε-equidistributed if :*

(1) $|P_u(g) - \frac{1}{2k}| \leqslant \varepsilon$ *for every $u \in X^{\pm 1}$;*
(2) $|P_{uv}(g) - \frac{1}{2k(2k-1)}| \leqslant \varepsilon$ *for every couple of letters $(u, v) \in X^{\pm 1} \times X^{\pm 1}$ with $u \neq v^{-1}$.*

(note that (2) implies (1), with a different ε).

LEMMA 7.15 ([14, Lemma 4.8]). — *Given k, there exists ε_0 such that, if a cyclic word g is ε_0-equidistributed, then g is strictly Whitehead minimal.*

LEMMA 7.16 ([14, Proposition 5.3]). — *Let γ_n be the number of reduced words of length n in $\mathbb{F}_k$, and let $\gamma_n(\varepsilon)$ be the number of ε-equidistributed reduced words of length n. For any $\varepsilon > 0$, the ratio $\gamma_n(\varepsilon)/\gamma_n$ goes to 1 exponentially fast as $n \to \infty$.*

We can now state:

PROPOSITION 7.17. — *Let H be as in Proposition 7.8. The following holds with probability going to 1 exponentially fast as $n \to +\infty$: for every non-trivial element $g \in H$, the cyclic reduction $\bar{g}$ of g is strictly Whitehead minimal.*

Proof. — We deduce this from the preceding lemmas and the central tree property (Lemma 7.10). We show that $\bar{g}$ is ε_0-equidistributed, with ε_0 provided by Lemma 7.15. Fix $\varepsilon_2 < \varepsilon_1 < \varepsilon_0$.

By Lemma 7.16 and radial symmetry, the words $w_{i,n}$ are ε_2-equidistributed generically. One obtains their cyclic reduction $\bar{w}_{i,n}$ by removing initial and terminal subwords, whose length is bounded by the central tree property; applying Lemma 7.10 with λ small enough (depending on ε_2 and ε_1), we deduce that, generically, the cyclic words $\bar{w}_{i,n}$ are ε_1-equidistributed.

We now consider the cyclic reduction $\bar{g}$ of a non-trivial $g \in H$. It is represented by an immersed loop in the Stallings graph Θ. Generically, this loop consists of arcs of length $< 2\lambda n$ contained in the central tree $\mathcal{C}$ (see Subsection 7.2.1) and outer loops of length $> (L - 2\lambda)n$. Frequencies are controlled in outer loops, and $\bar{g}$ is ε_0-equidistributed if λ is small enough (depending now on ε_1 and ε_0). $\qquad\square$

7.2.4. Matching subwords

The following lemma is a variation on a standard fact (see [2, Lemmas 4.5 and 4.6] and the references given there).

LEMMA 7.18. — *Let $0 < \beta < L/2$. With probability going to 1 exponentially fast as $n \to \infty$, the $2p$ words $w_{i,n}^{\pm 1}$ have length at least βn, and all their subwords of length βn are distinct.*

Proof. — This is similar to the proof of Lemma 7.10, but we have to use subexponential growth: it implies that, generically, $|w_{i,n}| \leqslant (2k-1)^{\beta n/3}$ for $i = 1, \ldots, p$.

We fix n, and numbers $A_1, \ldots, A_p$ with $\beta n \leqslant A_i \leqslant (2k-1)^{\beta n/3}$. We consider words w_i with $|w_i| = A_i$ chosen independently uniformly on the respective spheres. Thanks to Lemma 7.9, it suffices to bound the probability that some word of length βn appears twice in the words $w_i^{\pm 1}$ by some $C\kappa^{-n}$, with C, κ independent of $n, A_1, \ldots, A_p$.

First suppose that some word u of length βn appears in both w_i and $w_j^{\pm 1}$ for some fixed i, j with $i \neq j$. There are $A_j - \beta n$ possibilities for the location of u within $w_j^{\pm 1}$. Once this is fixed, w_j is determined by the letters outside of u, hence by two reduced words whose lengths add up to $A_j - \beta n$. It follows that the number of possibilities for w_j is bounded by $\widetilde{\gamma}_{A_j;n} = (A_j - \beta n)(2k)^2(2k-1)^{A_j - \beta n - 2}$. The probability that the words w_i and $w_j^{\pm 1}$ have a common subword of length βn is thus bounded by $\widetilde{\gamma}_{A_j;i,j,n}/2k(2k-1)^{A_j} = 2k(A_j - \beta n)(2k-1)^{-\beta n - 2}$, hence by $2k(2k-1)^{-2\beta n/3}$ because $A_j \leqslant (2k-1)^{\beta n/3}$.

Now suppose that u appears twice in $\{w_i, w_i^{-1}\}$. We then have two subwords u_1 and u_2 in w_i, each equal to $u^{\pm 1}$. There are $4(A_i - \beta n)^2$ possibilities for the location and sign of u_1, u_2. Fix one.

The key remark is the following. If we consider the set Z of cardinality A_i whose elements are the letters of w_i, and an equivalence relation on Z identifying each letter of u_1 with the corresponding letter of u_2, there is a subset $Y \subset Z$ of cardinality $A_i - \beta n$, consisting of one or two intervals and meeting each equivalence class. In particular, w_i is determined by $A_i - \beta n$ letters. We conclude by checking that the ratio between $4(A_i - \beta n)^2(2k)^2(2k-1)^{A_i - \beta n - 2}$ and $2k(2k-1)^{A_i - 1}$ is bounded by some $C(2k-1)^{-\beta n/3}$ provided that $A_i \leqslant (2k-1)^{\beta n/3}$. $\quad\square$

We generalize Lemma 7.18 as follows (compare Lemma 8.3 of [14]).

LEMMA 7.19. — *Let φ be a relabeling automorphism other than the identity. Let $0 < \beta < L/2$. With probability going to 1 exponentially fast*

as $n \to \infty$, the words $w_{i,n}^{\pm 1}$ cannot contain both a word u of length βn and its image by φ.

Proof. — This is proved as the previous lemma if $\varphi(u) \neq u$. If $\varphi(u) = u$, some $w_{i,n}$ has a subword of length βn all of whose letters are fixed by φ. Since φ is not the identity, these letters belong to a set of cardinality at most $2k - 2$. It is easily checked that this happens with probability going to 0 exponentially fast. □

7.2.5. Malcharacteristic subgroups

We can now prove Proposition 7.8. With probability going to 1 exponentially fast, H satisfies the conclusion of Lemmas 7.10, 7.18, 7.19 and Proposition 7.17 (with numbers λ and β which we choose so that $3\lambda < \beta < L/9$).

Consider $\alpha \in \mathrm{Aut}(\mathbb{F}_k)$ such that $\alpha(H) \cap H \neq \{1\}$. Fix a non-trivial element $h_1 \in H$ such that $h_2 = \alpha(h_1) \in H$. Denote by $\bar{h}_i = a_i h_i a_i^{-1}$ the cyclic reduction of h_i, and consider the automorphism $\theta = \mathrm{ad}_{a_2} \circ \alpha \circ \mathrm{ad}_{a_1}^{-1}$ sending $\bar{h}_1$ to $\bar{h}_2$.

The elements $\bar{h}_i$ are strictly Whitehead minimal by Proposition 7.17. By Lemma 7.13, they differ by a cyclic permutation and a relabeling φ, and $\theta = ad_g \circ \varphi$ for some $g \in G$. We claim that φ has to be the identity, so that θ and α are inner.

We view each $\bar{h}_i$ as an immersed loop in the Cayley graph of H. By the central tree property (Lemma 7.10), they consist of short arcs contained in the central tree and outer loops. Choose a subword u of $\bar{h}_1$ of length $3\beta n$ contained in an outer loop, hence in some $w_{i,n}$ (this is possible because $\beta < L/9$ and $\lambda < \beta/3$). The word $\varphi(u)$ appears as a subword of $\bar{h}_2$. Since $\lambda < \beta/3$, some subword of length βn of $\varphi(u)$ is contained in an outer loop. Lemma 7.19 now implies that φ is trivial.

We have proved that any α such that $\alpha(H) \cap H \neq \{1\}$ is inner. We deduce $\alpha \in \mathrm{Inn}_H(G)$ from malnormality of H. Malnormality of subgroups generated by elements chosen randomly in balls is known ([13], [2, Theorem 4.3]). Malnormality for H as in Proposition 7.8 follows from Lemma 7.18 as in the proof of Theorem 4.3 of [2].

BIBLIOGRAPHY

[1] Y. Antolín, A. Minasyan & A. Sisto, "Commensurating endomorphisms of acylindrically hyperbolic groups and applications", *Groups Geom. Dyn.* **10** (2016), no. 4, p. 1149-1210.

[2] F. Bassino, A. Martino, C. Nicaud, E. Ventura & P. Weil, "Statistical properties of subgroups of free groups", *Random Struct. Algorithms* **42** (2013), no. 3, p. 349-373.

[3] M. Bestvina & M. Feighn, "Stable actions of groups on real trees", *Invent. Math.* **121** (1995), no. 2, p. 287-321.

[4] M. Bestvina & M. Handel, "Train tracks and automorphisms of free groups", *Ann. Math.* **135** (1992), no. 1, p. 1-51.

[5] B. H. Bowditch, "Relatively hyperbolic groups", *Int. J. Algebra Comput.* **22** (2012), no. 3, article no. 1250016 (66 pages).

[6] C. H. Cashen & J. F. Manning, "Virtual geometricity is rare", *LMS J. Comput. Math.* **18** (2015), no. 1, p. 444-455.

[7] D. E. Cohen, *Combinatorial group theory: a topological approach*, London Mathematical Society Student Texts, vol. 14, Cambridge University Press, 1989, x+310 pages.

[8] F. Dahmani, V. Guirardel & D. V. Osin, "Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces", *Mem. Am. Math. Soc.* **245** (2017), no. 1156, p. v+152.

[9] M. J. Dunwoody & M. E. Sageev, "JSJ-splittings for finitely presented groups over slender groups", *Invent. Math.* **135** (1999), no. 1, p. 25-44.

[10] V. Guirardel & G. Levitt, in preparation.

[11] ———, "Splittings and automorphisms of relatively hyperbolic groups", *Groups Geom. Dyn.* **9** (2015), no. 2, p. 599-663.

[12] ———, *JSJ decompositions of groups*, Astérisque, vol. 395, Société Mathématique de France, 2017, vii+165 pages.

[13] T. Jitsukawa, "Malnormal subgroups of free groups", in *Computational and statistical group theory (Las Vegas, NV/Hoboken, NJ, 2001)*, Contemporary Mathematics, vol. 298, American Mathematical Society, 2002, p. 83-95.

[14] I. Kapovich, P. E. Schupp & V. Shpilrain, "Generic properties of Whitehead's algorithm and isomorphism rigidity of random one-relator groups", *Pac. J. Math.* **223** (2006), no. 1, p. 113-140.

[15] G. Levitt & F. Paulin, "Geometric group actions on trees", *Am. J. Math.* **119** (1997), no. 1, p. 83-102.

[16] A. D. Logan, "Every group is the outer automorphism group of an HNN-extension of a fixed triangle group", *Adv. Math.* **353** (2019), p. 116-152.

[17] R. C. Lyndon & P. E. Schupp, *Combinatorial group theory*, Classics in Mathematics, Springer, 2001, reprint of the 1977 edition, xiv+339 pages.

[18] J. Maher & A. Sisto, "Random subgroups of acylindrically hyperbolic groups and hyperbolic embeddings", *Int. Math. Res. Not.* (2019), no. 13, p. 3941-3980.

[19] D. V. Osin, "Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems", *Mem. Am. Math. Soc.* **179** (2006), no. 843, p. vi+100.

[20] J.-P. Otal, "Certaines relations d'équivalence sur l'ensemble des bouts d'un groupe libre", *J. Lond. Math. Soc.* **46** (1992), no. 1, p. 123-139.

[21] F. Paulin, "The Gromov topology on ℝ-trees", *Topology Appl.* **32** (1989), no. 3, p. 197-221.

[22] ———, "Sur les automorphismes extérieurs des groupes hyperboliques", *Ann. Sci. Éc. Norm. Supér.* **30** (1997), no. 2, p. 147-167.

[23] P. E. Schupp, "A characterization of inner automorphisms", *Proc. Am. Math. Soc.* **101** (1987), no. 2, p. 226-228.

Vincent GUIRARDEL
Univ Rennes, CNRS, IRMAR - UMR 6625
35000 Rennes (France)
vincent.guirardel@univ-rennes1.fr

Gilbert LEVITT
Laboratoire de Mathématiques Nicolas Oresme
(LMNO)
Université de Caen et CNRS (UMR 6139)
14000 Caen (France)
(Pour Shanghai : Normandie Univ, UNICAEN,
CNRS, LMNO, 14000 Caen, France)
levitt@unicaen.fr