Zeitschrift: L'Enseignement Mathématique

Herausgeber: Commission Internationale de l'Enseignement Mathématique

Band: 66 (2020)

Heft: 3-4

Artikel: The Martin boundary of relatively hyperbolic groups with virtually abelian parabolic subgroups

Autor: Dussaule, Matthieu / Gekhtman, Ilya / Gerasimov, Victor / Potyagailo, Leonid

DOI: https://doi.org/10.5169/seals-919580

L'Enseignement Mathématique (2) 66 (2020), 341–382 DOI 10.4171/LEM/66-3/4-3

The Martin boundary of relatively hyperbolic groups with virtually abelian parabolic subgroups

Matthieu DUSSAULE, Ilya GEKHTMAN,
Victor GERASIMOV and Leonid POTYAGAILO

Abstract. Given a probability measure on a finitely generated group, its Martin boundary is a way to compactify the group using the Green's function of the corresponding random walk. We give a complete topological characterization of the Martin boundary of finitely supported random walks on relatively hyperbolic groups with virtually abelian parabolic subgroups. In particular, in the case of nonuniform lattices in the real hyperbolic space $\mathbb{H}^n$, we show that the Martin boundary coincides with the $CAT(0)$ boundary of the truncated space.

Mathematics Subject Classification (2020). Primary 20F65, 20F67, 31C35; Secondary 60B15, 57M07, 22D05.

Keywords. Random walks on groups, hyperbolic groups relatively to a system of virtually abelian subgroups, Martin and Bowditch compactifications.

1. Introduction and statement of results

1.1. Random walks on relatively hyperbolic groups. A probability measure μ on a countable group Γ determines a Γ-invariant Markov chain with transition probabilities $p(x,y)=\mu(x^{-1}y)$, called a random walk.

Connecting asymptotic properties of this random walk to the geometry of Cayley graphs of Γ has been a fruitful line of research. One way to do this is through relating the Green's function of μ to some natural metric on Γ, and the probabilistically defined Martin boundary of μ to some geometric boundary of Γ.

The Green's function G of (Γ,μ) is defined as

$$G(x,y)=\sum_{n=0}^{\infty}\mu^{*n}(x^{-1}y).$$

It describes the amount of time a random path starting at x is expected to spend at y. We now fix a basepoint o in Γ. For each $y \in \Gamma$ the function $K_y : \Gamma \to \mathbb{R}$ defined by $K_y(x) = G(x,y)/G(o,y)$ is called a Martin kernel. The map $y \to K_y$ defines an embedding of Γ in the space of functions $\Gamma \to \mathbb{R}^+$. The closure of Γ in this space is called the Martin compactification $\overline{\Gamma}_\mu$ and $\partial_\mu \Gamma = \overline{\Gamma}_\mu \setminus \Gamma$ is called the Martin boundary. One can equally define the Martin compactification as the horospherical (Busemann) compactification with respect to the Green distance $d_G(x,y) = -\ln \frac{G(x,y)}{G(o,y)}$ on Γ [BB]. These definitions also make sense for more general measures μ. One of the main interest of the Martin boundary is its connection with harmonic functions. We will give more details in Section 2.

Giving a geometric description of the Martin boundary is often a difficult problem. Margulis showed that for centered finitely supported random walks on nilpotent groups, the Martin boundary is trivial [Mar]. On the other hand, for noncentered random walks with exponential moment on an abelian group of rank k, Ney and Spitzer [NS] showed that the Martin boundary is homeomorphic to a sphere of dimension $k-1$. For a hyperbolic group equipped with a finitely supported measure, Ancona [Anc] proved that the Martin boundary coincides with the Gromov boundary of the group. His proof relies on his famous following inequality which also implies that Γ is hyperbolic with respect to the Green distance. There exists $c \geq 0$ such that for any $x, y, z \in \Gamma$ lying in this order on a word geodesic,

$$d_G(x,y) + d_G(y,z) \leq d_G(x,z) + c.$$

Recall that an action $\Gamma \curvearrowright T$ is *minimal* and *non-elementary* if T is a minimal compact space invariant under the action and it contains more than two points. An action $\Gamma \curvearrowright T$ is called *convergence* if the induced action on the space of distinct triples of T is properly discontinuous. A minimal, non-elementary convergence action on a compact metric space T such that every point of T is either conical or bounded parabolic (see definitions in Section 3) is called *geometrically finite*. A finitely generated group Γ is called hyperbolic relative to a system of subgroups $\mathcal{P}$ if Γ admits a geometrically finite action on a metrisable compactum T such that the elements of $\mathcal{P}$ are the stabilizers of the parabolic points. The space T is called the Bowditch boundary and we will denote it by $\partial_B \Gamma$. It is known to be the Gromov boundary of a proper geodesic Gromov hyperbolic space on which Γ acts properly discontinuously and isometrically [Bow].

Gekhtman, Gerasimov, Potyagailo and Yang proved in [GGPY] for any finitely generated group Γ that the following generalized Ancona inequality is satisfied. For every $x, y, z \in \Gamma$,

$$d_G(x,y) + d_G(y,z) \leq d_G(x,z) + A\big(\delta_y^f(x,z)\big). \tag{1}$$

The inequality (1) is similar to that of Ancona but there are no restrictions on the group to be hyperbolic and on the triple of points x, y, z to belong to the same word-geodesic in the Cayley graph. On the other hand, the universal constant of Ancona is replaced with a function $A : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ which is a non-increasing function such that $A(t) \to +\infty$ once $t \to 0$. Its argument in (1) is a visibility function $\delta_y^f(x, z)$ which is the Floyd distance between x and z from y. This distance is obtained by rescaling the word distance with a quickly decreasing scalar function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ called *Floyd function* (see Section 3.1 for more details).

One of the corollaries of (1) obtained in [GGPY] states that if the group Γ is relatively hyperbolic then the identity map on the group extends to a continuous equivariant surjection from the Martin boundary to the Bowditch boundary of Γ. Moreover the preimage of any conical point under this map is a singleton. Determining the Martin boundary is thus reduced in this case to describing the preimages of a countable set of parabolic points.

The goal of this paper is to show that if the maximal parabolic subgroups are virtually abelian, then the Martin boundary of a finitely supported random walk is obtained by a blow-up construction at parabolic limit points of the Bowditch boundary. More precisely, one can define the induced chain on any neighborhood of a maximal parabolic subgroups, using the first return kernel to this neighborhood. We replace every parabolic limit point ξ at Bowditch boundary, stabilized by the parabolic subgroup $P \in \mathcal{P}$ with the Martin boundary of this return kernel on some fixed-sized neighborhood of P, which is the sphere of dimension $k-1$, where k is the rank of P (see Section 3.2 for more details).

This result was already known in some partial cases. Woess determined in [Woe1] the homeomorphism type of the Martin boundary for finitely supported nearest neighbor random walks on free products of the form $\mathbb{Z}^{d_1} * \mathbb{Z}^{d_2}$ (which are hyperbolic relative to each abelian factor). For such random walks, he proved that the Martin boundary is obtained by adding spheres of dimension $d_1 - 1$ for each left coset $\gamma\mathbb{Z}^{d_1}$ and spheres of dimension $d_2 - 1$ for each left coset $\gamma\mathbb{Z}^{d_2}$ to the set of infinite words. We also note that in the Woess' theorem one does not need to assume that the free factors are abelian. Precisely, the Martin boundary of a nearest neighbor random walk on $\Gamma_1 * \Gamma_2$ is obtained by gluing copies of the Martin boundary of the induced chains on Γ_1 and Γ_2 to the set of infinite words.

Another previously known example is given by finitely supported random walks on a non-cocompact lattice in the hyperbolic upper half-plane $\mathbb{H}^2$. In this case the Bowditch boundary is the limit set which coincides with the unit circle, and the group is also hyperbolic. So by Ancona's theorem, the Martin boundary is the Gromov boundary of the group which is homeomorphic in this case to the set of ends of a free group obtained by blowing-up at each parabolic fixed point of rank 1 into two points. However, in general, for geometrically finite Kleinian groups (even for non-cocompact lattices in the hyperbolic 3-space $\mathbb{H}^3$) the homeomorphism type of the Martin boundary of finitely supported random walks was not known before. For sake of completeness, let us also mention the work of Ballman and Ledrappier [BL]. They study discretization of the Brownian motion on Riemannian manifolds. In particular, they identify the Martin boundary for some classes of measures on geometrically finite Kleinian groups. Those measures however are not in general finitely supported.

In view of the above discussion, we ask the following question.

Question 1.1. Let Γ be a relatively hyperbolic group and let μ be a probability measure on Γ whose finite support generates Γ as a semigroup. Let φ be the equivariant continuous surjective map from the Martin boundary $\partial_\mu\Gamma$ to the Bowditch boundary $\partial_B\Gamma$ constructed in [GGPY]. Is the φ-preimage of a parabolic point homeomorphic to the Martin boundary of the induced chain on a bounded neighborhood of the corresponding parabolic subgroup ?

In this paper, we give a positive answer to this question when the group Γ is hyperbolic relatively to virtually abelian subgroups. This class of groups includes: geometrically finite subgroups of the group of isometries of the real hyperbolic space $\mathbb{H}^n$, limit groups, and finitely generated groups acting freely on $\mathbb{R}^n$-trees (the latter groups are hyperbolic relatively to abelian subgroups by [Gui], [Dah2]).

Definition 1.2. We will call a boundary, obtained by replacing a fixed point of a parabolic subgroup of rank k by a sphere of dimension $k-1$, *parabolic blow-up* boundary (PBU-boundary for short).

We refer to Section 3.2 for more details on this definition. Dussaule [Dus] generalized Woess' result [Woe1] and gave a geometric description of the Martin boundary for any finitely supported random walk on free products of abelian groups, identifying it with the visual boundary of a $CAT(0)$ space on which the group acts cocompactly. In particular it is a PBU-boundary. The key technical result of [Dus] extends results of Ney and Spitzer [NS] to more general chains. It states that the Martin boundary of non-centered -or strictly sub-Markov- chains on $\mathbb{Z}^k\times\{1,\dots,N\}$, $N\in\mathbb{N}$, is a sphere of dimension $k-1$ (see Proposition 4.6 below). One of the main results of the paper is the following.

Theorem 1.3. *Let Γ be a non-elementary finitely generated relatively hyperbolic group with respect to a collection of infinite virtually abelian subgroups. Let μ be a measure on Γ whose finite support generates Γ as a semigroup. Then, the Martin boundary is a PBU-boundary.*

The following direct corollary answers the above question in the case when all parabolic subgroups are virtually abelian.

Corollary 1.4. *Suppose that all the assumptions of Theorem 1.3 are valid. Let $\varphi : \partial_\mu \Gamma \to \partial_B \Gamma$ be an equivariant continuous map from the Martin boundary to the Bowditch boundary. Then for every parabolic point $p \in \partial_B \Gamma$ the set $\varphi^{-1}(p)$ is homeomorphic to the Martin boundary of the stabilizer H of p with respect to the induced measure.*

Another central result of the paper is the following.

Theorem 1.5. *Let Γ be a non-elementary finitely generated relatively hyperbolic group with respect to a collection of infinite virtually abelian subgroups. Let μ be a measure on Γ whose finite support generates Γ as a semigroup. Then, every point of the Martin boundary corresponds to a minimal harmonic function.*

There is a particularly simple geometric construction of a PBU-boundary when Γ is a non-uniform lattice in the real hyperbolic n-space $\mathbb{H}^n$. By removing from $\mathbb{H}^n$ a Γ-equivariant collection of disjoint horoballs based at parabolic fixed points and considering the induced shortest-path metric on the complement, we obtain a $CAT(0)$ space on which Γ acts cocompactly. One can easily check that the visual boundary of this $CAT(0)$ space is a PBU-boundary. In particular, when $n = 3$, the PBU-boundary is homeomorphic to the sphere $\mathbb{S}^2$ with a countable and dense set of discs removed. It is then homeomorphic to the Sierpinski carpet (see [Rua, Theorem 4.1, Corollary 4.2] for a proof and see also [TW, Theorem 1] for a more general statement). Concluding this discussion we obtain from Theorem 1.3 the following.

Corollary 1.6. *Let Γ be a non-uniform lattice in the real hyperbolic space $\mathbb{H}^n$. Let μ be a probability measure on Γ whose finite support generates Γ as a semigroup. Then, the Martin boundary is equivariantly homeomorphic to the $CAT(0)$ boundary of the truncated space. In particular if $n = 3$, it is homeomorphic to the Sierpinski carpet.*

1.2. Brief description of difficulties and ideas of the proofs. To prove Theorem 1.3 we define the induced random walk on each parabolic subgroup P as the first return kernel on P. This chain happens to be strictly sub-Markov. The aim would be to show that the Martin boundary of the induced chain on P coincides with the limit set of P in the original Martin boundary of (Γ, μ). By [GGPY, Corollary 8.3] this would be true if we knew *a priori* that every point in the preimage of ξ in the Martin boundary of (Γ, μ) corresponds to a *minimal* harmonic function. However to prove the minimality, we need to have a precise description of the Martin boundary of the induced sub-Markov chain on parabolic subgroups. So we proceed in the opposite way: we first characterize the preimages of parabolic points and a posteriori we obtain that they correspond to minimal harmonic functions.

Our proofs of Theorems 1.3 and 1.5 use both the inequality (1) and the generalization of the theorem of Ney and Spitzer [NS] given in [Dus]. Roughly stated, we show that the preimage of a parabolic point on the Bowditch boundary, with stabilizer P, is homemorphic to the Martin boundary of a neighborhood of P with a finite (though not probability) measure induced by the first return times, which by the result of [Dus] is a sphere of the appropriate dimension.

To show that the Martin boundary $\partial_\mu \Gamma$ is a PBU-boundary, we have to deal with two types of trajectories, namely those converging to conical points in the Bowditch boundary and those whose projections converge in the geometric boundary of a parabolic subgroup. To treat the first type of trajectories, we use results of [GGPY] which imply that whenever a sequence g_n converges to a conical limit point then it converges to a unique point of the Martin boundary (see Proposition 5.2).

For the second type of trajectories, we study the induced random walk on a parabolic subgroup P. An additional difficulty, not mentioned above, is that this random walk is not finitely supported. Using results of Gerasimov and Potyagailo [GP3], we prove that the induced random walk on a sufficiently large neighborhood of P has large exponential moments. Applying then several results of [Dus], we show that if a sequence g_n converges in the boundary of P, then it converges in the Martin boundary. Furthermore two different points in the boundary of P correspond to two different points in the Martin boundary.

To finish this discussion we stress that the above methods crucially use several times that the parabolic subgroups are virtually abelian. So already in the case where the group Γ is hyperbolic relatively to virtually nilpotent subgroups, Question 1.1 remains open.

1.3. Organization of the paper. The paper is divided into six main parts, besides the introduction.

Section 2 is devoted to giving the necessary probabilistic background on random walks, Markov chains, and their Martin boundaries. In particular, we will explain the relationship between the Martin boundary and harmonic functions. This relationship was actually the reason for introducing the Martin boundary in the first place.

Section 3 is about relatively hyperbolic groups. In Sections 3.1, we give the definition of those groups and state results of Gerasimov and Potyagailo about the interplay between the Floyd distance and the geometry of the Cayley graph of such groups. We also define properly what is a PBU-boundary and what is the geometric compactification of a parabolic subgroup in Section 3.2.

In Section 4, we give the necessary geometric background on Martin boundaries for the proof of our main theorem. In Section 4.1, we restate the inequality (1) obtained in [GGPY]. This inequality will be used throughout the proofs, especially when we deal with trajectories converging to conical limit points. In Section 4.2, we state the results of Dussaule about Martin boundaries of chains on $\mathbb{Z}^k \times \{1, \ldots, N\}$. This part is a bit technical and we extend his results to deal later with trajectories converging in the geometric boundary of a parabolic subgroup.

In Section 5, we prove our main theorem, Theorem 1.3. We first deal with conical limit points in Section 5.1, using results of Section 4.1 and then with parabolic subgroups in Section 5.2, using results of Section 4.2.

In Section 6, we prove Theorem 1.5, that is, every point of the Martin boundary is a minimal harmonic function. Again, we will deal separately with trajectories converging to conical limit points and trajectories converging in the boundary of parabolic subgroups.

In Section 7, we give a geometric construction of a PBU-boundary, using a construction of Dahmani in [Dah1]. We also state some questions suggested by our results.

2. Martin boundaries of random walks

Let us give here a proper definition of the Martin boundary and the minimal Martin boundary. In this paper, we deal with random walks on groups, but during the proofs, we will restrict the random walk to thickenings of peripheral subgroups and we will not get actual random walks. Thus, we need to define Martin boundaries for more general transition kernels.

Consider a countable space E and equip E with the discrete topology. Fix some base point o in E. Consider a transition kernel p on E with finite total mass, that is $p : E \times E \to \mathbb{R}_+$ satisfies

$$\forall x \in E, \sum_{y \in E} p(x, y) < +\infty.$$

It is often required that the total mass is 1 and in that case, the transition kernel defines a Markov chain on E. In general, we will say that p defines a chain on E and we will sometimes assume that this chain is sub-Markov, that is the total mass is at most 1. If μ is a probability measure on a finitely generated group Γ, then $p(g,h) = \mu(g^{-1}h)$ is a probability transition kernel and the Markov chain is the random walk associated to μ.

Define in this context the Green's function G as

$$G(x,y) = \sum_{n\geq 0} p^{(n)}(x,y) \in [0,+\infty],$$

where $p^{(n)}$ is the nth convolution power of p, i.e.

$$p^{(n)}(x,y) = \sum_{x_1,\dots,x_{n-1}} p(x,x_1)p(x_1,x_2)\cdots p(x_{n-1},y).$$

Definition 2.1. Say that the chain defined by p is finitely supported if for every $x \in E$, the set of $y \in E$ such that $p(x,y) > 0$ is finite.

Definition 2.2. Say that the chain defined by p is irreducible if for every $x, y \in E$, there exists n such that $p^{(n)}(x,y) > 0$.

For a Markov chain, this means that one can go from any $x \in E$ to any $y \in E$ with positive probability. In this setting, the Green's function $G(x,y)$ is closely related to the probability that a μ-governed path starting at x ever reaches y. Indeed, the latter quantity is equal to $\frac{G(x,y)}{G(y,y)}$ (see [Woe2, Lemma 1.13.(b)].

Notice that in the case of a random walk on a group Γ, the Green's function is invariant under left multiplication, so that $G(x,x) = G(o,o)$ for every x. Thus, up to some multiplicative constant, $G(x,y)$ is the probability to go from x to y. We denote this probability by $\mathbb{P}(x \to y)$ in the following, so that $G(x,y) = G(o,o)\mathbb{P}(x \to y)$. Moreover, the irreducibility of the chain is equivalent to the condition that the support of the measure μ generates Γ as a semigroup.

In particular, in the context of Theorem 1.3, the transition kernel defined by the probability measure μ is irreducible.

We will also use the following definition during our proofs.

Definition 2.3. Say that the chain defined by p is strongly irreducible if for every $x, y \in E$, there exists n_0 such that $\forall n \geq n_0$, $p^{(n)}(x,y) > 0$.

We will also assume that the chain is transient, meaning that the Green's function is everywhere finite. For a Markov chain, this just means that almost surely, a path starting at x returns to x only a finite number of times.

Consider an irreducible transient chain p. For $y \in E$, define the Martin kernel based at y as

$$K_y(x) = \frac{G(x,y)}{G(o,y)}.$$

The Martin compactification of E with respect to p (and o) is a compact metrizable space containing E as an open and dense space, whose topology is described as follows. A sequence y_n in E converges to a point in the Martin compactification if and only if the sequence K_{y_n} converges pointwise. Up to isomorphism, it does not depend on the base point o and we denote it by $\overline{E}_p$. We also define the Martin boundary as $\partial_p E = \overline{E}_p \setminus E$. We refer to [Saw] for a complete construction of the Martin compactification.

Seeing the Martin kernel K as a function of two variables x and y, the Martin compactification is then the smallest compact space M in which E is open and dense and such that K can be extended to the space $E \times M$, continuously on the second variable. If $\tilde{y} \in \overline{E}_p$, denote by $K_{\tilde{y}}$ the extension of the Martin kernel.

In the particular case of a symmetric Markov chain, that is a Markov chain satisfying $p(x,y) = p(y,x)$, the Green's distance, which was defined by Brofferio and Blachère in [BB] as

$$d_G(x,g) = -\ln \mathbb{P}(x \to y) = -\ln \frac{G(x,y)}{G(y,y)},$$

is actually a metric and the Martin compactification of E with respect to the Markov chain p is the horofunction compactification of E for this metric.

Now, assume that $E = \Gamma$ is a finitely generated group and that the transition kernel p is defined by a probability measure μ. In that case, denote by $\overline{\Gamma}_\mu$ the Martin compactification and by $\partial_\mu \Gamma$ the Martin boundary. The action by (left) multiplication of Γ on itself extends to an action of Γ on $\overline{\Gamma}_\mu$.

One important aspect of the Martin boundary is its relation with harmonic functions. Recall that if p is a transition kernel on a countable space E, a harmonic function is a function $\phi : E \to \mathbb{R}$ such that $p\phi = \phi$, that is,

$$\forall x \in E, \phi(x) = \sum_{y \in E} p(x,y)\phi(y).$$

We have the following key property (see [Saw, Theorem 4.1]).

Proposition 2.4 (Martin Representation Theorem). *Let p be a irreducible transient transition kernel on a countable space E. For any non-negative harmonic function ϕ, there exists a measure ν on the Martin boundary $\partial_\mu E$ of E such that*

$$\forall x \in E, \phi(x) = \int_{\partial_\mu E} K_{\tilde{x}}(x) d\nu(\tilde{x}).$$

Let ϕ be a non-negative harmonic function. It is called minimal if any other non-negative harmonic function ψ such that $\psi(x) \leq \phi(x)$ for every $x \in E$ is proportional to ϕ by a constant. The minimal Martin boundary is the set

$$\partial_\mu^m E = \{\tilde{x} \in \partial_\mu E, K_{\tilde{x}}(\cdot) \text{ is minimal harmonic}\}.$$

It is thus a subset of the full Martin boundary $\partial_\mu E$. A classical representation theorem of Choquet shows that for any non-negative harmonic function ϕ, one can choose the support of the measure ν lying in $\partial_\mu^m E$. The measure ν is then unique (see [KSK, Chapter 10.7]). In other words, for any such function ϕ, there exists a unique measure μ_ϕ on $\partial_\mu^m E$ such that

$$\forall x \in E, \phi(x) = \int_{\partial_\mu^m E} K_{\tilde{x}}(x) \mathrm{d}\mu_\phi(\tilde{x}).$$

3. Relatively hyperbolic groups

3.1. Relative hyperbolicity and the Floyd metric. Let Γ be a finitely generated group. The action of Γ on a compact Hausdorff space T is called a convergence action if the induced action on triples of distinct points of T is properly discontinuous. If T is a metrizable compactum then the action $G \curvearrowright T$ is convergence if and only if every sequence of distinct elements g_n in Γ contains a subsequence g_{n_k} such that $g_{n_k} x \to a$ and for all $x \in X$ with at most perhaps one exceptional point.

Suppose $\Gamma \curvearrowright T$ is a convergence action. The set of accumulation points $\Lambda\Gamma$ of any orbit Γx $(x \in T)$ is called the *limit set* of the action. As long as $\Lambda\Gamma$ has more than two points, it is uncountable and the unique minimal closed Γ-invariant subset of T. The action is then said to be non-elementary. In this case, the orbit of every point in $\Lambda\Gamma$ is infinite. The action is *minimal* if $\Lambda\Gamma = T$.

A point $\zeta \in \Lambda\Gamma$ is called *conical* if there is a sequence g_n of Γ and distinct points $\alpha, \beta \in \Lambda\Gamma$ such that $g_n\zeta \to \alpha$ and $g_n\eta \to \beta$ for all $\eta \in T \setminus \{\zeta\}$. The point $\zeta \in \Lambda\Gamma$ is called bounded parabolic if it is the unique fixed point of its stabilizer in Γ, which is infinite and acts cocompactly on $\Lambda\Gamma \setminus \{\zeta\}$. The stabilizers of bounded parabolic points are called (maximal) parabolic subgroups. The convergence action $\Gamma \curvearrowright T$ is called geometrically finite if every point of $\Lambda\Gamma \subset T$ is either conical or bounded parabolic.

Definition 3.1. Let $\mathcal{P}$ be a collection of subgroups of Γ. We say that Γ is hyperbolic relative to $\mathcal{P}$ if there exists some compactum T on which Γ acts minimally and geometrically finitely and the maximal parabolic subgroups are the elements of $\mathcal{P}$.

Such a compactum is then unique up to Γ-equivariant homeomorphism [Bow] and is called the Bowditch boundary of $(\Gamma, \mathcal{P})$. The group Γ is said to be relatively hyperbolic if it is hyperbolic relative to some collection of subgroups, or equivalently if it admits a geometrically finite convergence action on some compactum. The group Γ is non-elementary relatively hyperbolic if it admits a non-elementary geometrically finite convergence action on some infinite compactum.

Since Γ is assumed to be finitely generated, every maximal parabolic subgroup is finitely generated too (see [Ger1, Main Theorem (d)]). Then, by Yaman's results [Yam], it follows that if $\Gamma \curvearrowright T$ is a minimal geometrically finite action, then there exists a proper geodesic Gromov hyperbolic space X on which Γ acts properly discontinuously by isometries and a Γ-equivariant homeomorphism $T \to \partial X$.

Let Γ be a group hyperbolic relative to a collection of parabolic subgroups $\mathcal{P}$. The set $\mathcal{P}$ is invariant under conjugacy, since the set of parabolic limit points is invariant under the group action. Furthermore, the set $\mathcal{P}$ contains at most finitely many conjugacy classes of maximal parabolic subgroups (see [Tuk, Theorem 1B]).

We now discuss the Floyd distance and the Floyd boundary. Let $f : \mathbb{R}^+ \to \mathbb{R}^+$ be a function satisfying two conditions: $\sum_{n\geq 0} f_n < \infty$; and there exists a $\lambda \in (0,1)$ such that $1 \geq f_{n+1}/f_n \geq \lambda$ for all $n \in \mathbb{N}$. The function f is then called the *rescaling function* or the *Floyd function.*

Let S be a finite system of generators of Γ, we denote by $C_S\Gamma$ the Cayley graph, and by $d(.,.)$ the word distance. Fix a basepoint $o \in \Gamma$ and rescale $C_S\Gamma$ by declaring the length of an edge σ to be $f(d(o,\sigma))$. The induced shortpath metric on $C_S\Gamma$ is called the *Floyd metric* with respect to the basepoint o and Floyd function f and denoted by $\delta_o^f(.,.)$. Its Cauchy completion (whose topology does not depend on the basepoint) is called the Floyd compactification $\overline{\Gamma}_f$ and $\partial_f\Gamma = \overline{\Gamma}_f \setminus \Gamma$ is called the Floyd boundary. Karlsson showed that the action of a group on its Floyd boundary is always convergence [Kar, Theorem 2]. On the other hand, if Γ is relatively hyperbolic and if the Floyd function f is not decreasing exponentially too fast, Gerasimov [Ger2, Map theorem] proved that there is continuous Γ-equivariant surjection (*Floyd map*) from the Floyd boundary to the Bowditch boundary. Furthermore, Gerasimov and Potyagailo [GP2, Theorem A] proved that the preimage of any conical point by this map is a singleton and the preimage of a parabolic fixed point p is the limit set for the action of its stabilizer Γ_p on $\partial_f\Gamma$. In particular if Γ_p is an amenable non-virtually cyclic group then its limit set on the Floyd boundary is a point. Consequently, when Γ is hyperbolic relative to a collection of infinite amenable subgroups which are not virtually cyclic, the Floyd boundary is homeomorphic to the Bowditch boundary.

Let $\alpha : I \to C_S\Gamma$ be a (finite or infinite) geodesic for the Cayley metric. We also fix a relatively hyperbolic structure on Γ and let $\mathcal{P}$ be a system of maximal parabolic subgroups of it. A point $v = \alpha(t_0) \in \alpha$ $(t_0 \in I)$ is said to be (ϵ, R)-*deep* if there exists $g \in \Gamma$ and $P \in \mathcal{P}$ such that the image $\alpha(]t_0 - R, t_0 + R[)$ is contained in the ϵ-neighborhood of gP. Otherwise, $p \in \alpha$ is called an (ϵ, R)-*transition point* (or simply (ϵ, R)-transitional) of α. Gerasimov and Potyagailo proved the following key property.

Proposition 3.2 ([GP3, Corollary 5.10]). *For each $\epsilon > 0$, $R > 0$ there exists $\delta > 0$ such that the following holds. For any 3 distinct points x, y, z in Γ, if y is situated in a (word) distance D from an (ϵ, R)-transition point of a word geodesic $[x, z]$ between x and z, then $\delta_y^f(x, z) \geq \delta$.*

Remark about the proof. The argument of [GP3, Corollary 5.10] guarantees that the *shortcut distance* $\overline{\delta}_y^f(x, z)$ on the Bowditch compactification of the Cayley graph is bounded below by some constant $\delta > 0$. This distance is obtained by transferring the Floyd metric by the Floyd map constructed by Gerasimov [Ger2]. It satisfies the following inequality (see [GP1, Section 3.1]): $\delta_y^f(x, z) \geq \overline{\delta}_y^f(x, z)$. This gives the above statement (see [Ger2] and [GP1] for more details). We note also that the shortcut distance provides a necessary and sufficient criterion in the above context: the point y is a (ϵ, R)-transition point if and only if $\overline{\delta}_y^f(x, z) \geq \delta > 0$ where (ϵ, R) and δ determine each other.

3.2. Geometric compactifications. We now give a precise definition of a PBU-boundary. We first define the geometric boundary of an infinite, virtually abelian, finitely generated group. Let P be such a group, so that there exists a subgroup of P isomorphic to $\mathbb{Z}^k$, for some $k \geq 1$, with finite index in P. Then, any section $P/\mathbb{Z}^k \to P$ provides an identification between P and $\mathbb{Z}^k \times \{1, \dots, N\}$ for some $N \geq 1$. Let g_n be a sequence in P and identify g_n with $(z_n, j_n) \in \mathbb{Z}^k \times \{1, \dots, N\}$. Say that the sequence g_n converges to a point in the boundary of P if z_n tends to infinity and $\frac{z_n}{\|z_n\|}$ converges to some point in the sphere $\mathbb{S}^{k-1}$. Here, $\|\cdot\|$ is the Euclidean norm. This defines what we call the geometric boundary ∂P of P.

Actually, since P is virtually abelian, it admits a proper and cocompact action on a $CAT(0)$ space, see [BH, Remark 7.3 (2)]. The geometric boundary of P we defined coincides with the $CAT(0)$ boundary of P, which is the visual boundary of a $CAT(0)$ space $\mathbb{R}^k$ on which P acts properly and cocompactly. Note that the $CAT(0)$ boundary of a finitely generated group is not well defined in general since $CAT(0)$ boundaries are not invariant under quasi-isometries. However, the $CAT(0)$ boundary of a virtually abelian group is well defined. More generally, if a relatively hyperbolic group Γ with respect to virtually abelian subgroups acts isometrically on a CAT(0) space then the boundary $\partial\Gamma$ is a well defined $CAT(0)$ boundary, see [HK, Theorem 1.2.1, Theorem 1.2.2 (3)]. In particular, the geometric boundary of P neither depends on the choice of the abelian subgroup $\mathbb{Z}^k$ of finite index in P nor on the choice of the section $P/\mathbb{Z}^k \to P$.

The definition of the geometric boundary immediately implies the following.

Lemma 3.3 (*Perspectivity property*). *Let $P \in \mathcal{P}$ and z_n, z'_n two sequences converging to ∂P such that the distance $||z - z'||$ is bounded. Then they both converge to the same point in ∂P.*

Proof. By definition of the geometric topology $||z_n|| \to \infty$ (and so is z'_n). Assume that the sequence $\frac{z_n}{||z_n||}$ converges to $\theta \in \mathbb{S}^{k-1}$ where k is the rang of P. Since $z'_n = z_n + u_n$ and $||u_n||$ is bounded it follows that $\lim_{n\to\infty} \frac{z'_n}{||z'_n||} = \theta$. The same argument works in the opposite sens. □

Remark. The property established in the lemma we call *perspectivity property*. It is true for virtually abelian parabolic subgroups but it is largely unknown in other cases (in particular for nilpotent parabolic subgroups, see also the concluding Section 7.3).

If F is a finite set, we define the geometric boundary of the product $P \times F$ as follows. First identify P with $\mathbb{Z}^k \times \{1, \dots, N\}$ as before. This provides an identification between $P \times F$ and $\mathbb{Z}^k \times \{1, \dots, N'\}$ for some other integer $N' \geq 1$. As above, a sequence g_n in $P \times F$ is said to converge in the geometric boundary if its projection onto $\mathbb{Z}^k$ under this identification converges in the geometric boundary of $\mathbb{Z}^k$. This slight generalization will be useful in the following. Indeed, for technical reasons, when studying sequences converging in ∂P, we will not restrict the random walk to parabolic subgroups but to bounded neighborhoods of them.

Suppose now that Γ is a finitely generated group hyperbolic relative to a collection $\mathcal{P}$ of infinite subgroups. We will assume through the paper that every parabolic subgroup is virtually abelian. Let $\mathcal{P}_0 \subset \mathcal{P}$ denote a full subset of representatives of conjugacy classes of parabolic subgroups. As mentioned, by [Tuk, Theorem 1B], $\mathcal{P}_0$ is finite. We define a parabolic blow-up (PBU) boundary for Γ relative to this choice of $\mathcal{P}_0$, although our definition will not depend on it up to equivariant homeomorphism, since we will prove that the Martin boundary is a PBU-boundary (and the Martin boundary does not depend on $\mathcal{P}_0$).

Fixing a word metric d on Γ, for $A \subset \Gamma$ and $g \in \Gamma$, let

$$\mathrm{proj}_A(g) = \{h \in A : d(h, g) = d(A, g)\}$$

be the set of closest point projections of g to A. For a subset $F \subset \Gamma$ let

$$\mathrm{proj}_A(F) = \bigcup_{g \in F} \mathrm{proj}_A(g).$$

Let $\pi_A : \Gamma \to A$ be any function with $\pi_A(g) \in \mathrm{proj}_A(g)$. The following result will be used several times further on.

Proposition 3.4 ([GP3, Proposition 4.3]). *The diameter of* $\mathrm{proj}_P(g)$ *is finite and bounded uniformly and independently on* $g \in \Gamma$ *and* $P \in \mathcal{P}$.

In particular when Γ is a hyperbolic group relatively to virtually abelian subgroups, then by Lemma 3.3 for a sequence $g_n \in \Gamma$ the convergence of $\pi_P(g_n)$ to the geometric boundary of P does not depend on the choice of π_P.

We will use boundaries throughout the paper. We fix the following terminology. A *compactification* $\overline{\Gamma}$ of Γ is a *metrizable* compact space, containing Γ as an open and dense space, endowed with a group action by homeomorphisms $\Gamma \curvearrowright \overline{\Gamma}$ that extends the action by left multiplication on Γ. Then, $\partial\Gamma := \overline{\Gamma} \setminus \Gamma$ is called a boundary of Γ.

Definition 3.5. Let Γ be a group hyperbolic relative to a collection $\mathcal{P}$ of virtually abelian subgroups. Fix a finite set of representatives of conjugacy classes of parabolic subgroups $\mathcal{P}_0 \subset \mathcal{P}$. A parabolic blow-up (PBU) boundary of $(\Gamma, \mathcal{P}_0)$ is a boundary $\partial\Gamma$ such that the following holds.

A sequence g_n in Γ converges to a point in $\partial\Gamma$ if and only if g_n tends to infinity and either g_n converges to a conical point in the Bowditch boundary or there exist $g \in \Gamma$ and a parabolic subgroup $P \in \mathcal{P}_0$ such that $g^{-1}\pi_{gP}(g_n)$ converges in the geometric boundary of P.

We start proving that the topology on the PBU-boundary of $(\Gamma, \mathcal{P}_0)$ is well defined.

Proposition 3.6. *Let* Γ *be a group hyperbolic relative to a collection* $\mathcal{P}$ *of virtually abelian subgroups. Let* $\partial_1\Gamma$ *and* $\partial_2\Gamma$ *be two PBU-boundaries. Then, the identity on* Γ *extends to an equivariant homeomorphism from* $\Gamma \cup \partial_1\Gamma$ *to* $\Gamma \cup \partial_2\Gamma$.

Proof. We first define a map Φ from $\partial_1\Gamma$ to $\partial_2\Gamma$. Let $\xi \in \partial_1\Gamma$. Since Γ is dense in $\Gamma \cup \partial_1\Gamma$, there exists a sequence g_n of points of Γ converging to ξ. By Definition 3.5 this implies that g_n also converges to some point $\tilde{\xi}$ in $\partial_2\Gamma$. Let us prove that $\tilde{\xi}$ only depends on ξ. Assume that g_n' is another sequence converging to ξ in $\partial_1\Gamma$, so that g_n' converges to some $\tilde{\xi}'$ in $\partial_2\Gamma$. Consider the sequence g_n'' defined by $g_{2n}'' = g_{2n}$ and $g_{2n+1}'' = g_{2n+1}'$. Then, g_n'' also converges to ξ in $\partial_1\Gamma$ so it converges to a point in $\partial_2\Gamma$. This proves that $\tilde{\xi} = \tilde{\xi}'$. We define a map $\Phi : \Gamma \cup \partial_1\Gamma \to \Gamma \cup \partial_2\Gamma$ as an extension of the identity map on Γ given by $\Phi(\xi) = \tilde{\xi}$.

By construction, whenever a sequence g_n in Γ converges to $\xi \in \partial_1\Gamma$, $\Phi(g_n)$ converges to $\Phi(\xi)$. Similarly, we define a map $\Psi : \Gamma \cup \partial_2\Gamma \to \Gamma \cup \partial_1\Gamma$. If g_n in Γ converges to $\xi \in \partial_1\Gamma$ and converges to $\zeta \in \partial_2\Gamma$, then $\Phi(\xi) = \zeta$ and $\Psi(\zeta) = \xi$. Thus, $\Phi \circ \Psi = Id$ on $\partial_2\Gamma$ and $\Psi \circ \Phi = Id$ on $\partial_1\Gamma$. Obviously, $\Phi \circ \Psi = \Psi \circ \Phi = Id$ on Γ, so that Φ and Ψ are inverse bijections. Hence, we only need to prove that Φ and Ψ are continuous and equivariant. By symmetry, we only need to prove that Φ is continuous and equivariant.

Let x_n be a sequence in $\Gamma \cup \partial_1\Gamma$ converging to some x. Assume first that $x \in \Gamma$. Since Γ is open in $\Gamma \cup \partial_1\Gamma$, $x_n \in \Gamma$ for large enough n, so that x_n converges to $x \in \Gamma$ and so $\Phi(x_n)$ converges to $\Phi(x)$. Consider now some $x \in \partial_1\Gamma$. Since both compactifications are assumed to be metrizable, we choose arbitrary distances d_i on $\Gamma \cup \partial_i\Gamma$, $i = 1, 2$.

Our goal is to construct a sequence g_n in Γ that also converges to x and satisfies that $d_1(g_n, x_n) \leq 1/n$ and $d_2(g_n, \Phi(x_n)) \leq 1/n$. Whenever $x_n \in \Gamma$, we set $g_n = x_n$. Let n be such that $x_n \in \partial_1\Gamma$. Since Γ is dense in $\Gamma \cup \partial_1\Gamma$, there exists a sequence $g_{n,m}$ converging to x_n when m tends to infinity. By construction of Φ, this implies that $g_{n,m}$ converges to $\Phi(x_n)$ when m tends to infinity. We can thus find large enough m that we denote by m_n such that, letting $g_n = g_{n,m_n}$, we have $d_1(g_n, x_n) \leq 1/n$ and $d_2(g_n, \Phi(x_n)) \leq 1/n$. Now, x_n converges to x and $d_1(g_n, x) \leq d_1(g_n, x_n) + d(x_n, x)$, so that g_n also converges to x. By construction of Φ, this implies that g_n converges to $\Phi(x)$ in $\partial_2\Gamma$. Finally, $d_2(\Phi(x_n), \Phi(x)) \leq d_2(\Phi(x_n), g_n) + d_2(g_n, \Phi(x))$, so that $\Phi(x_n)$ also converges to $\Phi(x)$. This shows that Φ is continuous.

Let us prove that Φ is equivariant to conclude. Since Φ is the identity on Γ, for any $g, g' \in \Gamma$, $\Phi(gg') = gg' = g \cdot \Phi(g')$. Assume that $\xi \in \partial_1\Gamma$ and let $g \in \Gamma$. Choose a sequence g_n converging to ξ. Then, gg_n converges to $g \cdot \xi$. By construction of Φ, gg_n converges to $\Phi(g \cdot \xi)$ in $\partial_2\Gamma$. Also, by construction, g_n converges to $\Phi(\xi)$ so that gg_n converges to $g \cdot \Phi(\xi)$ in $\partial_2\Gamma$. This shows that $\Phi(g \cdot \xi) = g \cdot \Phi(\xi)$, which concludes the proof. □

The following result implies that Definition 3.5 matches the rough definition given in the introduction (Definition 1.2).

Proposition 3.7. *Let Γ be a group hyperbolic relative to a collection $\mathcal{P}$ of virtually abelian subgroups. Let $\partial\Gamma$ be a PBU-boundary. Then, the identity on Γ extends to a continuous equivariant surjective map*

$$\Gamma \cup \partial\Gamma \to \Gamma \cup \partial_B\Gamma,$$

where $\partial_B\Gamma$ is the Bowditch boundary. Moreover, the preimage of a conical limit point is a singleton and the preimage of a parabolic point is an $n-1$-dimensional sphere where n is the rank of its stabilizer.

Proof. We first construct a map $\Phi_{PBU} : \partial\Gamma \to \partial_B\Gamma$. Let $\zeta \in \partial\Gamma$. Since Γ is dense in $\Gamma \cup \partial\Gamma$, there exists a sequence of points g_n in Γ converging to ζ. There are two cases. First, if g_n converges to a conical limit point α, then we define $\Phi_{PBU}(\zeta) = \alpha$. Otherwise, there exists $P \in \mathcal{P}_0$ and $g \in \Gamma$ such that $g^{-1}\pi_{gP}(g_n)$ converges in the geometric boundary of P. Let α be the parabolic limit point fixed by P and define $\Phi_{PBU}(\zeta) = g\alpha$. We extend this map to a map $\Phi_{PBU} : \Gamma \cup \partial\Gamma \to \Gamma \cup \partial_B\Gamma$ declaring Φ_{PBU} to be the identity on Γ, as in the proof of Proposition 3.6.

The group Γ is dense in both compactifications and both compactifications are metrizable. Indeed, by definition, a PBU-boundary is assumed to be metrizable and it follows from Bowditch's construction [Bow] that $\Gamma \cup \partial_B\Gamma$ is metrizable. Note also that $\Gamma \cup \partial_B\Gamma$ admits the shortcut metric mentioned in the Remark after Proposition 3.2. Therefore, it is sufficient to prove that whenever a sequence g_n in Γ converges to $\zeta \in \partial\Gamma$, g_n converges to $\Phi_{PBU}(\zeta) \in \partial_B\Gamma$. If $\Phi_{PBU}(\zeta)$ is conical, this is given by the construction of Φ_{PBU}. Otherwise, $\Phi_{PBU}(\zeta) = \alpha$ is parabolic. Letting $P \in \mathcal{P}_0$ and $g \in \Gamma$ be such that the stabilizer of $g^{-1}\alpha$ is P, $g^{-1}\pi_{gP}(g_n)$ converges to ζ. In particular, $\pi_{gP}(g_n)$ tends to infinity so that g_n converges to α. By construction, Φ_{PBU} is surjective and equivariant.

To conclude, we note that the preimage of a conical limit point is a single point by construction and that the preimage of a parabolic limit point α is homeomorphic to the geometric boundary of $P \in \mathcal{P}_0$, where P is the stabilizer of $g^{-1}\alpha$ for some $g \in \Gamma$. This geometric boundary of P is an $n-1$-dimensional sphere where n is the rank of P. This concludes the last part of the proof. □

Our main theorem states that the Martin boundary of a finitely supported random walk is a PBU-boundary, so that in particular, such a boundary always exists and does not depend on the choice of generators. In Section 7 we will also give a geometric construction of a PBU-boundary based on a compactifications of relatively hyperbolic groups introduced by F. Dahmani which in turns is based on a classical topological $\mathcal{Z}$-compactification.

4. Topology of Martin boundaries

4.1. Generalized Ancona's inequality. Suppose Γ is a finitely generated group. Let μ be a probability measure whose finite support generates Γ as a semigroup and let G be the associated Green's function.

Denote by $G(x,z;B_R^c(y))$ the Green's function from x to z conditioned by not visiting the ball of center y and radius R, that is

$$G(x,z;B_R^c(y)) = \sum_{k\geq 0} \mathbb{P}_x\big(X_k = z \mid X_l \notin B_R(y),\ l \in \{1,\dots,k-1\}\big).$$

For a hyperbolic group Ancona's inequality, mentioned in the Introduction, can be restated in the multiplicative form as follows. There exists a uniform constant C, depending only on the hyperbolicity constant of the group, such that for any three distinct points x, y, z lying along a word geodesic in this order in the Cayley graph, one has

$$\frac{1}{C}G(x,y)G(y,z) \leq G(x,z) \leq CG(x,y)G(y,z).$$

Ancona used this inequality to identify the Martin boundary of hyperbolic groups with their Gromov boundary. To apply this theory to relatively hyperbolic groups, we will need the following result of Gekhtman, Gerasimov, Potyagailo and Yang which implies the inequality (1). Fix a Floyd function f.

Theorem 4.1 ([GGPY, Theorem 1.3]). *For each $\epsilon > 0$ and $\delta > 0$ there exists $R > 0$ such that for all $x, y, w \in \Gamma$ satisfying that $\delta_w^f(x,y) \geq \delta > 0$, the probability that the random walk starting at x and conditioned to reach y avoids a ball in the Cayley graph centered at w of radius R is at most ϵ. In terms of the Green function, it can be stated as*

$$G\big(x,y;B_R^c(w)\big) \leq \epsilon G(x,y). \tag{2}$$

By Proposition 3.2, the points x, y, w satisfy $\delta_w^f(x,y) \geq \delta$ for some fixed $\delta > 0$ as soon as the point w is within a word distance D of a transition point on a word geodesic $[x,y]$. If the group Γ is word-hyperbolic then all points on a geodesic in the Cayley graph are transition points for a uniform constant (depending only on the hyperbolicity constant), hence Theorem 4.1 implies the Ancona's inequality in this case, see [GGPY, Corollary 1.4].

4.2. Martin boundaries of thickened abelian groups. To understand the behavior of $K_{g_n}(g)$, when g_n converges in the geometric boundary of a parabolic subgroup, we will introduce the transition kernel of the first return to the corresponding subgroup P. We will then get a sub-Markov chain on P and we will show that we can identify this first-return-chain with a $\mathbb{Z}^k$-invariant sub-Markov chain on $\mathbb{Z}^k \times \{1, \dots, N\}$ for the standard action $z \cdot (z', k) = (z + z', k)$ (see Lemma 5.10). We will then use results for such chains.

In [Dus], the author shows that under some technical assumptions, the Martin boundary of such a chain on $\mathbb{Z}^k \times \{1, \dots, N\}$ coincides with the geometric boundary. In this setting, the geometric boundary is defined as in Section 3.2. Namely, a sequence (z_n, j_n) in $\mathbb{Z}^k \times \{1, \dots, N\}$ converges to a point in the geometric boundary if z_n tends to infinity and $\frac{z_n}{\|z_n\|}$ converges in the unit sphere $\mathbb{S}^{k-1}$. We now introduce the assumptions of [Dus] and we will later show that they are satisfied in our setting.

Consider a $\mathbb{Z}^k$-invariant chain p on the product space $\mathbb{Z}^k \times \{1, \dots, N\}$. For every function defined on $\mathbb{Z}^k \times \{1, \dots, N\}$, the $\{1, \dots, N\}$ coordinate will be considered as an index. For example, the transition kernel will be written as $p_{j_1,j_2}(z_1, z_2)$, its nth power of convolution as $p^{(n)}_{j_1,j_2}(z_1, z_2)$, the Green's function as $G_{j_1,j_2}(z_1, z_2)$ and the Martin kernel as $K_{j_1,j_2}(z_1, z_2)$. We can thus see these functions as the entries of $N \times N$ matrices. Assume that the chain p is strongly irreducible, that is, for every $j_1, j_2 \in \{1, \dots, N\}$ and for every $z_1, z_2 \in \mathbb{Z}^k$, there exists n_0 such that for every $n \geq n_0$, $p^{(n)}_{j_1,j_2}(z_1, z_2) > 0$. As we will see later (see Lemma 5.1), strong irreducibility is not too much to ask and we will be able to reduce our study of irreducible chains to strongly irreducible ones.

In [NS], Ney and Spitzer show that the Martin boundary of a strongly irreducible, finitely supported, noncentered random walk on $\mathbb{Z}^k$ coincides with the $CAT(0)$ boundary. Their proof is based on the study of minimal harmonic functions which are of the form $z \in \mathbb{Z}^k \mapsto \mathrm{e}^{u \cdot z}$ for some $u \in \mathbb{R}^k$ satisfying the condition

$$\sum_{z \in \mathbb{Z}^k} p(0, z)\mathrm{e}^{u \cdot z} = 1. \tag{3}$$

In our setting, for $u \in \mathbb{R}^k$, we define the $N \times N$ matrix $F(u)$ whose entries are given by

$$F_{j_1,j_2}(u) = \sum_{z \in \mathbb{Z}^k} p_{j_1,j_2}(0, z)\mathrm{e}^{u \cdot z}.$$

The entries of this matrix may be infinite. We restrict our attention to the set where they are finite and denote this set by $\mathcal{F}_0$:

$$\mathcal{F}_0 = \left\{u \in \mathbb{R}^k, \forall j_1, j_2 \in \{1, \dots, N\}, F_{j_1,j_2}(u) < +\infty\right\}.$$

We also denote by $\mathcal{F}$ the interior of $\mathcal{F}_0$.

Let $M > 0$. Let p be a chain on $\mathbb{Z}^k \times \{1, \dots, N\}$. Say that p has exponential moments up to M if for every $j, j' \in \{1, \dots, N\}$,

$$\sum_{z \in \mathbb{Z}^k} p_{j,j'}(0, z)\mathrm{e}^{M\|z\|} < +\infty.$$

We will show in Proposition 5.11 that the chain has exponential moments. Hence every coefficient $F_{j_1,j_2}(u)$ is finite for small u. It follows that $\mathcal{F}_0$ contains a small ball, and so the set $\mathcal{F}$ is not empty.

Lemma 4.2. *For every $u \in \mathcal{F}_0$, the matrix $F(u)$ has non-negative entries. Furthermore, this matrix is strongly irreducible, meaning that there exists $n \geq 0$ such that $F(u)^n$ has positive entries.*

Proof. The calculation provided in [Dus, Lemma 3.2] shows that the entries of $F(u)^n$ are given by

$$F_{j_1,j_2}(u)^n = \sum_{z \in \mathbb{Z}^k} p^{(n)}_{j_1,j_2}(0,z) \mathrm{e}^{u \cdot z}.$$

Strong irreducibility of $F(u)$ is deduced from strong irreducibility of p. □

Since $F(u)$ is strongly irreducible, it follows from the Perron–Frobenius theorem (see [Sen, Theorem 1.1]) that $F(u)$ has a dominant positive eigenvalue, that is an eigenvalue $\lambda(u)$ which is positive and such that for every other eigenvalue $\lambda \in \mathbb{C}$, $|\lambda| < \lambda(u)$. Moreover, any eigenvector associated to $\lambda(u)$ has non-zero coordinates and we can assume that every coordinate is positive. The analog of Equation (3) will be

$$\lambda(u) = 1. \tag{4}$$

Denote by D the set where $\lambda(u)$ is at most 1: $D = \{u \in \mathcal{F}, \lambda(u) \leq 1\}$. The two technical assumptions of [Dus] on the chain p are the following.

Assumption 4.3. The set D is compact.

Assumption 4.4. The minimum of the function λ is strictly smaller than 1.

Since $\lambda(u)$ is a dominant eigenvalue, it is analytic in u (see Proposition 8.20 in [Woe2]). For $u \in \mathcal{F}$, denote by $\nabla\lambda(u)$ the gradient of λ with respect to u. We have the following.

Lemma 4.5 ([Dus, Lemma 3.22]). *Under Assumptions 4.3 and 4.4, the set*

$$\{u \in \mathbb{R}^k, \lambda(u) = 1\}$$

is homeomorphic to $\mathbb{S}^{k-1}$. An explicit homeomorphism is given by

$$u \in \{u \in \mathbb{R}^k, \lambda(u) = 1\} \mapsto \frac{\nabla\lambda(u)}{\|\nabla\lambda(u)\|}.$$

This provides a homeomorphism φ between $\mathbb{S}^{k-1}$ and the geometric boundary of $\mathbb{Z}^k \times \{1, \dots, N\}$ constructed as follows. Let (z_n, j_n) be a sequence in $\mathbb{Z}^k \times \{1, \dots, N\}$ converging to a point $\tilde{z}$ in the geometric boundary $\partial(\mathbb{Z}^d \times \{1, \dots, N\})$. Then z_n tends to infinity and $\frac{z_n}{\|z_n\|}$ converges to a point θ in the unit sphere $\mathbb{S}^{k-1}$. There exists a unique $u \in \{u \in \mathbb{R}^k, \lambda(u) = 1\}$ such that $\theta = \frac{\nabla\lambda(u)}{\|\nabla\lambda(u)\|}$. Then, define $\varphi(\tilde{z}) = u$.

The Martin boundary is defined up to the choice of a base point. Fix such a base point $(z_0, j_0) \in \mathbb{Z}^k \times \{1, \dots, N\}$. Now, we can state that the Martin boundary coincides with the geometric boundary.

Proposition 4.6 ([Dus, Proposition 3.29]). *Let p be a strongly irreducible transition kernel on $\mathbb{Z}^k\times\{1,\dots,N\}$ which is $\mathbb{Z}^k$-invariant and satisfies Assumptions 4.3 and 4.4. If $z_n\in\mathbb{Z}^k$ converges to $\tilde{z}\in\partial\mathbb{Z}^k$, let $u=\varphi(\tilde{z})$. Then, for every $z\in\mathbb{Z}^k$ and for every $j_1,j_2\in\{1,\dots,N\}$, there exists a constant $C_{j_1}>0$ which only depends on j_1 such that $K_{j_1,j_2}(z,z_n)$ converges to $C_{j_1}e^{u\cdot(z-z_0)}$.*

Consider now a chain p on $\mathbb{Z}^k\times\mathbb{N}$. If $N\geq 1$, define the induced chain p_N as the chain of the first return to $\mathbb{Z}^k\times\{1,\dots,N\}$, that is, if $(z,j),(z',j')\in\mathbb{Z}^k\times\{1,\dots,N\}$,

$$
\begin{aligned}
&p_N\big((z,j),(z',j')\big)=p\big((z,j),(z',j')\big)\\
&+\sum_{k\geq 1}\sum_{\substack{(z_1,j_1),\dots,(z_k,j_k)\\ j_1,\dots,j_k>N}} p\big((z,j),(z_1,j_1)\big)p\big((z_1,j_1),(z_2,j_2)\big)\dots p\big((z_k,j_k)(z,j')\big).
\end{aligned}
$$

Denote by G the Green's function associated to p and by G_N the Green's function associated to the induced chain p_N. Then, we have the following lemma.

Lemma 4.7. *The restriction to $\mathbb{Z}^k\times\{1,\dots,N\}$ of the Green's function G coincides with the Green's function G_N.*

Proof. Every trajectory from (z,j) to (z',j') for the initial chain p defines a trajectory from (z,j) to (z',j') for p_N, by conditioning the trajectory on the successive passages through $\mathbb{Z}^k\times\{1,\dots,N\}$. Every trajectory for p_N is uniquely obtained in such a way. Summing over all trajectories, the two Green's functions coincide. □

We also have the following proposition.

Proposition 4.8. *Let p be a $\mathbb{Z}^k$-invariant, finitely supported, strongly irreducible transition kernel on $\mathbb{Z}^k\times\mathbb{N}$. Then, there exist $N_0\geq 0$ and $M>0$ such that whenever $N\geq N_0$ and the chain p_N has exponential moments up to M, p_N satisfies Assumption 4.3.*

Proof. We will prove that there exist $\mu>0$ and $M>0$, such that for sufficiently large N, whenever the chain p_N has exponential moments up to M, then

$$
\{u\in\mathbb{R}^k,\lambda(u)\leq 2\}\subset\{u\in\mathbb{R}^k,\|u\|\leq\mu\}\subset\mathcal{F}. \tag{5}
$$

Denote by $(e_1,\dots,e_k)$ the canonical basis in $\mathbb{R}^k$. Since the chain p is strongly irreducible, there exists n_i such that for every $n\geq n_i$, $p^{(n)}((0,1),(e_i,1))>0$. Thus, there exists n_0 such that for every i, $p^{(n_0)}((0,1),(e_i,1))>0$, so that there is a path of length n_0 from $(0,1)$ to $(e_i,1)$. These paths stay in $\mathbb{Z}^k\times\{1,\dots,N_0\}$

for some N_0, since the chain p is finitely supported. Thus, for every $N \geq N_0$, for every i, the restricted chain p_N satisfies $p_N^{(n_0)}((0,1),(e_i,1)) \geq a$, for some $a > 0$.

Let $N \geq N_0$ and $u \in \mathbb{R}^k$, and let us fix $L \geq 0$. There exists $\mu > 0$ such that if $\|u\| \geq \mu$, then at least one of the $e^{u \cdot e_i}$ is larger than $\frac{L}{a}$, so that $F_{1,1}(u)^{n_0} \geq L$. Moreover, if p_N has exponential moments up to $\mu+1$, then $F(u)$ has finite entries for $\mu \leq \|u\| \leq \mu+1$ and so does $F(u)^{n_0}$. Let $v(u)$ be an eigenvector associated to $\lambda(u)$. Then, it is an eigenvector of $F(u)^{n_0}$ associated to $\lambda(u)^{n_0}$. Since $F(u)$ is strongly irreducible, $v(u)$ has non-zero coordinates and we can even choose $v(u)$ with strictly positive coordinates. Denote by $v(u)(1)$ its first coordinate. Then, $v(u)(1)\lambda(u)^{n_0} \geq F_{1,1}(u)^{n_0}v(u)(1)$ so that $\lambda(u)^{n_0} \geq F_{1,1}(u)^{n_0} \geq L$.

Consequently, $\lambda(u)^{n_0}$, hence $\lambda(u)$, can be made arbitrarily large, when enlarging $\|u\|$. Moreover, if p_N has sufficiently large exponential moments, then $\lambda(u)$ is well defined for arbitrarily large $\|u\|$. Inclusion (5) now follows from these two facts. This proves that the sub-level $\lambda(u) \leq 1$, is bounded, thus compact and contained in the open set $\lambda(u) < 2$, which is included in $\mathcal{F}$. Thus, Assumption 4.3 is satisfied. □

We will also use the following.

Lemma 4.9. *Let p be a $\mathbb{Z}^k$-invariant, strongly irreducible chain on $\mathbb{Z}^k \times \{1,\dots,N\}$. If p is (strictly) sub-Markov, then it satisfies Assumption 4.4.*

Proof. The fact that the chain is strictly sub-Markov means that the matrix $F(0)$ defined above is strictly sub-stochastic. In particular, its dominant eigenvalue $\lambda(0)$ satisfies $\lambda(0) < 1$ and the minimum of λ is strictly less than 1, so Assumption 4.4 is satisfied. □

Combining the explicit formula given in Proposition 4.6 together with Proposition 4.8 and Lemma 4.9, we obtain the following corollary which describes convergence in the Martin boundary for a chain on $\mathbb{Z}^k \times \mathbb{N}$.

Corollary 4.10. *Let p be a $\mathbb{Z}^k$-invariant, finitely supported, strongly irreducible transition kernel on $\mathbb{Z}^k \times \mathbb{N}$ such that:*

(a) *For large enough N, the induced chain p_N on $\mathbb{Z}^k \times \{1,\dots,N\}$ is strictly sub-Markov.*

(b) *For all M there exists an $N_0 > 0$ such that for $N > N_0$, the chain p_N has exponential moments up to M.*

Then, a sequence (z_n, j_n) in $\mathbb{Z}^k \times \mathbb{N}$, with $\sup(j_n) < +\infty$, converges to a point in the Martin boundary of p if and only if $\|z_n\|$ tends to infinity and $\frac{z_n}{\|z_n\|}$ converges to a point of $\mathbb{S}^{k-1}$.

Thus, the fact that (z_n, j_n) converges to a point in the Martin boundary does not depend on the sequence (j_n) as long as it remains bounded. In particular, (z_n, j_n) converges in the Martin boundary if and only if its projection $(z_n, 0)$ on $\mathbb{Z}^k \times \{0\}$ converges in the Martin boundary and the limits are the same.

5. Convergence of Martin kernels: Proof of Theorem 1.3

Let Γ be a hyperbolic group relative to a collection $\mathcal{P}$ of infinite virtually abelian subgroups. Let μ be a measure on Γ whose finite support generates Γ as a semigroup. In this section we prove that the Martin boundary is a PBU-boundary, proving Theorem 1.3. Recall that ∂P denotes the geometric boundary of a maximal parabolic subgroup P defined in Section 3.2. We fix a finite set $\mathcal{P}_0$ of representatives of conjugacy classes of $\mathcal{P}$. We will deal separately with sequences converging to conical limit points and sequences converging in $g\partial P$ for some coset gP of a parabolic subgroup $P \in \mathcal{P}_0$. For the second case, we will apply results of Section 4.2. It will be more convenient to deal with strongly irreducible chains. Thus, we first show that we can reduce to such chains.

In a very general context, consider a chain p on a countable space E. Define the modified chain $\tilde{p}$ on E by

$$\tilde{p}(x_1, x_2) = \frac{1}{2}\Delta(x_1, x_2) + \frac{1}{2}p(x_1, x_2),$$

where $\Delta(x_1, x_2) = 0$ if $x_1 \neq x_2$ and 1 otherwise. Denote by $\tilde{p}^{(n)}$ the nth convolution power of $\tilde{p}$. Also denote by $\tilde{G}$ the associated Green's function:

$$\tilde{G}(x_1, x_2) = \sum_{n \geq 0} \tilde{p}^{(n)}(x_1, x_2).$$

We have the following (see [Woe2, Lemma 9.2]).

Lemma 5.1. *With these notations, $\frac{1}{2}\tilde{G}(x_1, x_2) = G(x_1, x_2)$ and thus the Martin kernels are the same.*

In our context, this means we can assume that $\mu(e) > 0$, and so the random walk is strongly irreducible.

5.1. Convergence to conical limit points. We first study conical limit points. We prove the following.

Proposition 5.2. *Consider a sequence g_n of Γ that converges to a conical point α in the Bowditch boundary. Then, g_n converges to a point in the Martin boundary.*

This is a consequence of the following two results of [GGPY].

Proposition 5.3 ([GGPY, Theorem 7.4]). *The identity map $\Gamma \to \Gamma$ extends to a continuous equivariant surjection F from the Martin compactification to the Bowditch compactification.*

Proposition 5.4. [GGPY, Corollary 7.14] *The preimage $F^{-1}(\alpha)$ of a conical limit point α consists of a single point.*

Indeed, let g_n converge to a conical limit point α and assume that g_n does not converge in the Martin compactification. By compactness, g_n has two subsequences that converge to two distinct points in the Martin boundary, which are both mapped to α by F. This is a contradiction. $\square$

5.2. Convergence in parabolic subgroups. In this section, we prove the convergence of the Martin kernels $K_{g_n}(\cdot)$ when g_n converges to a point in the geometric boundary of a parabolic subgroup P. We fix a finite set $\mathcal{P}_0$ of conjugacy classes of parabolic subgroups. Let $P \in \mathcal{P}_0$. By assumption, P contains a subgroup isomorphic to $\mathbb{Z}^k$ with finite index. Any section $P/\mathbb{Z}^k \to P$ provides an identification between P and $\mathbb{Z}^k \times \{1, \dots, N\}$.

Let p_n be a sequence in P. Identify then p_n with $(z_n, j_n) \in \mathbb{Z}^k \times \{1, \dots, N\}$. By definition, the sequence p_n converges to a point in the boundary ∂P of P if and only if z_n tends to infinity and $\frac{z_n}{\|z_n\|}$ converges to some point in $\mathbb{S}^{k-1}$. Denote the corresponding point in $\mathbb{S}^{k-1}$ by θ and say that p_n converges to θ. Our goal is to prove the following.

Proposition 5.5. *Let g_n be a sequence in Γ. If $\pi_P(g_n)$ converges to a point of ∂P, then g_n and $\pi_P(g_n)$ both converge in the Martin boundary $\partial_\mu \Gamma$ to the same point.*

Corollary 5.6. *Let $g \in \Gamma$ and let $P \in \mathcal{P}_0$. If g_n is a sequence such that $g^{-1}\pi_{gP}(g_n)$ converges to a point of ∂P, then g_n converges to some point in the Martin boundary.*

Proof of the Corollary. By Proposition 3.4 the diameter of $\mathrm{proj}_P(g)$ is uniformly bounded. In our situation, $\pi_P(g^{-1}g_n) \in (\mathrm{proj}_P(g^{-1}g_n) = g^{-1}(\mathrm{proj}_{gP}g_n))$, so the element $\pi_P(g^{-1}g_n)$ is within a bounded distance of $g^{-1}\pi_{gP}(g_n)$. Thus, $\pi_P(g^{-1}g_n)$ also converges to some point in ∂P. Proposition 5.5 shows that $g^{-1}g_n$ converges to some point in the Martin boundary, hence so does g_n. $\square$

The key point in the proof of Proposition 5.5 is the following statement.

Proposition 5.7. *Suppose g_n is a sequence with $\sup_n d(g_n, P) < \infty$. Then, g_n converges to a point in the Martin boundary of Γ if and only if $\pi_P(g_n)$ converges to a point of ∂P. Moreover, in that case, g_n and $\pi_P(g_n)$ both converge in the Martin boundary of Γ to the same point.*

The rest of this section is devoted to prove Proposition 5.7. Proposition 5.5 will then be deduced from it.

Let $\eta \geq 0$ and let $N_\eta(P)$ be the η-neighborhood of P. We introduce the chain p corresponding to the first return to $N_\eta(P)$, defined as in Section 2. Namely, for $g, g' \in N_\eta(P)$ denote by $p(g, g')$ the probability that the random walk starting at g returns to $N_\eta(P)$, and first does so at g'. In other words $p(g, g') = G(g, g'; N_\eta^c(P))$. We will see that the probability that the random walk starting at g never goes back to $N_\eta(P)$ is positive (see Lemma 5.10). Thus, p is not a probability transition kernel and defines a sub-Markov chain on $N_\eta(P)$. Nevertheless, one can still define the Green's function associated to p as

$$G_p(g, g') = \sum_{n \geq 0} p^{(n)}(g, g'), \quad g, g' \in N_\eta(P),$$

where $p^{(n)}$ is the nth power of convolution of p. According to Lemma 4.7, we have the following.

Lemma 5.8. *The Green's function G_p coincides with the restriction to $N_\eta(P)$ of the Green's function G_μ associated to the initial random walk.*

We also have the following property.

Lemma 5.9. *The chain p is strongly irreducible.*

Proof. The proof is based on the same idea as the proof of Lemma 5.8. First, the initial random walk is irreducible. Now, every trajectory for p comes from a trajectory for the random walk on the whole group, after excluding points that do not stay in the neighborhood of P. Thus, there is a positive proportion (for p) of paths from any point $g \in N_\eta(P)$ to any other point $g' \in N_\eta(P)$. This proves that p is irreducible. Now, recall that we assumed that $\mu(e) > 0$ (see Lemma 5.1), so that $p(g, g) > 0$ and thus p is strongly irreducible. □

In light of Lemma 5.8, to prove Proposition 5.7, it suffices to show that a sequence satisfying its conditions converges to a point in the Martin boundary of $N_\eta(P)$ with the induced chain p.

We first notice that, as a set, Γ can be identified P-equivariantly with $P \times \mathbb{N}$. Indeed, P acts by left multiplication on Γ and the quotient is countable. We order elements in the quotient according to their distance to P. It follows that the η-neighborhood $N_\eta(P)$ can be P-equivariantly identified with $P \times \{1, \dots, N_\eta\}$. Moreover, if $\eta' \leq \eta$, the set $P \times \{1, \dots, N_{\eta'}\}$ identified with $N_{\eta'}(P)$ is a subset of $P \times \{1, \dots, N_\eta\}$ identified with $N_\eta(P)$.

Now, identifying P with $\mathbb{Z}^k \times F$, where F is finite, we identify the group Γ with $\mathbb{Z}^k \times \mathbb{N}$. Thus, the μ-random walk can be considered as a $\mathbb{Z}^k$-invariant Markov chain q on $\mathbb{Z}^k \times \mathbb{N}$ and the restriction of the random walk to $N_\eta(P)$ coincides with the restriction of the chain q to $\mathbb{Z}^k \times \{1, \dots, \tilde{N}_\eta\}$ for some $\tilde{N}_\eta$ To simplify the notations, we will write $N = \tilde{N}_\eta$.

Let g_n be a sequence in $N_\eta(P)$ and identify g_n with $(z_n, j_n) \in \mathbb{Z}^k \times \{1, \dots, N\}$. Notice that the projection of g_n to P converges in the geometric boundary ∂P of P if and only if (z_n, j_n) converges in the geometric boundary of $\mathbb{Z}^k \times \{1, \dots, N\}$, since in both cases, the sequence converges in the geometric boundary if and only if z_n tends to infinity and $\frac{z_n}{\|z_n\|}$ converges to a point in the sphere.

To prove Proposition 5.7, it suffices to show that the Markov chain q on $\mathbb{Z}^k \times \mathbb{N}$ and its induced chain p on $\mathbb{Z}^k \times \{1, \dots, N\}$ satisfy the conditions of Corollary 4.10. Thus, we just need to show that for large enough η, the induced chain on $N_\eta(P)$ has sufficiently large exponential moments and is strictly sub-Markov.

Lemma 5.10. *The induced chain p is strictly sub-Markov.*

Proof. It suffices to show that there exists $g \in N_\eta(P)$ such that

$$\sum_{g' \in N_\eta(P)} p(g, g') < 1.$$

This follows from the fact that the μ-random walk starting at g with $d(g, P) = \eta$ has a positive probability of never returning to $N_\eta(P)$. This, in turn, follows from the fact that the random walk almost surely converges to a conical point (see, for example, [GGPY, Theorem 9.8, Theorem 9.14]). □

For $M > 0$, recall that p is said to have exponential moments up to M if for every $j, j' \in \{1, \dots, N\}$,

$$\sum_{z \in \mathbb{Z}^k} p_{j,j'}(0, z) e^{M\|z\|} < +\infty.$$

Proposition 5.11. *Let $M > 0$. For large enough η, p has exponential moments up to M.*

The proof of Proposition 5.11 will be divided into several steps. We will use the following geometric preliminary results.

Proposition 5.12 ([GP3, Proposition 8.5]). *There are constants $a_0, D > 0$, independent of the parabolic subgroup P, such that if α is a geodesic with endpoints α_1, α_2 satisfying $d(\mathrm{proj}_P(\alpha_i), \alpha) > D$, then*

$$\mathrm{diam}\big(\mathrm{proj}_P(\alpha)\big) < a_0.$$

The following two results follow from Proposition 5.12. Note that the first one also follows from [Hru, Corollary 8.2].

Corollary 5.13. *There exists c_0 such that the following holds. If $g_1, g_2 \in \Gamma$ and if $g_i' \in \mathrm{proj}_P g_i$ for $i = 1, 2$ then, $d(g_1', g_2') \leq d(g_1, g_2) + c_0$.*

Proof. Consider a geodesic α from g_1 to g_2. If α stays outside the D-neighborhood of P, then $d(g_1', g_2') \leq a_0$.

Otherwise, denote by α_1, respectively α_2 the first, respectively last point of α within a distance D of P and let $\alpha_i' \in \mathrm{proj}_P(\alpha_i)$. Applying Proposition 5.12, we get $d(g_i', \alpha_i') \leq a_0$. Also, $d(\alpha_i', \alpha_i) \leq D$, so that the triangle inequality yields

$$d(g_1', g_2') \leq 2a_0 + 2D + d(\alpha_1, \alpha_2) \leq d(g_1, g_2) + c_0,$$

where $c_0 = 2D + 2a_0$. □

Corollary 5.14. *For large enough $a > 0$, the function $\rho : \mathbb{R}_+ \to \mathbb{R}_+$ defined by*

$$\rho(\eta) = \inf\Big\{d(g_1, g_2) : d\big(\pi_P(g_1), \pi_P(g_2)\big) \geq a, d(g_i, P) > \eta\Big\}$$

tends to infinity as $\eta \to \infty$.

Proof. Choose a and η larger than the constants a_0 and D from Proposition 5.12. Let $g_1, g_2 \in \Gamma$ be such that $d(g_1, P) > \eta$ and $d(g_1, g_2) \leq \eta - D$. Let α be a geodesic connecting g_1 and g_2. By the triangle inequality, $d(\alpha, P) \geq D$. Consequently, we have $diam(\mathrm{proj}_P(\alpha)) < a_0 \leq a$. In particular $d(\pi_P(g_1), \pi_P(g_1)) < a$. Thus, the conditions $d(g_1, P) > \eta$ and $d(\pi_P(g_1), \pi_P(g_1)) \geq a$ imply that $d(g_1, g_2) > \eta - D$, completing the proof. □

The following classical lemma is a consequence of the fact that the spectral radius of μ is less than 1. This follows in turn from non-amenability of Γ, according to a famous result of Kesten, see [Kes].

Lemma 5.15 (Kesten). *Denote by μ^{*n} the nth power of convolution of the measure μ. There exists $\alpha > 0$ such that for every $g \in \Gamma$,*

$$\mu^{*n}(g) \leq \mathrm{e}^{-\alpha n}.$$

We can now prove Proposition 5.11. Let $z \in \mathbb{Z}^k$ and $j, j' \in \{1, \dots, N\}$. If the first return to $N_\eta(P)$ starting at $(0, j)$ is at (z, j'), there is a path $Z_0, \dots, Z_{n+1}$ such that $Z_0 = (0, j)$, $Z_{n+1} = (z, j')$ and $Z_l \notin N_\eta(P)$ for $1 \leq l \leq n$. Note that $d(Z_l, Z_{l+1}) \leq r(\mu)$, where $r(\mu)$ only depends on the support of μ. Thus, if $\eta \geq 3r(\mu)$, then $Z_0, Z_{n+1} \notin N_{2\eta/3}(P)$ as soon as $n \geq 1$, which will hold if $\|z\|$ is large enough. Moreover by the triangle inequality, any geodesic from Z_l to Z_{l+1} stays outside of $N_{\eta/3}(P)$, for $0 \leq l \leq n$.

Define a path ϕ from Z_0 to Z_{n+1} by gluing together geodesics from Z_l to Z_{l+1}. Then, the length of ϕ is at most $nr(\mu)$. The parabolic subgroup P together with the induced word distance is quasi-isometric to its subgroup $\mathbb{Z}^k$ together with the Euclidean distance. In particular, the word distance between 0 and z in P is larger than $\Lambda\|z\|$, where $||z||$ stands for the Euclidean norm of $z \in \mathbb{Z}^k$ and Λ only depends on the quasi-isometry parameters.

Denote by v_i the vertices of the path ϕ. We claim that for the constants a and c_0 from Corollary 5.14 and Corollary 5.13, assuming that $\mathrm{diam}(\mathrm{proj})_P\phi \geq a$ (for sufficiently large $||z||$), we can choose points $y_k = v_{i_k}$ $(k = 1, \dots, l)$ such that their projections $\tilde{y}_k = \pi_P(y_k)$ satisfy

$$a \leq d(\tilde{y}_k, \tilde{y}_{k+1}) \leq a_1 = a + 1 + c_0.$$

Indeed, let $\tilde{v}_j = \pi_P(v_j)$ and $y_1 = v_1$. We argue by induction. Assuming that the index k_{i-1} was already chosen, let $k_i = \min\{k > k_{i-1}, d(\tilde{y}_{i-1}, \tilde{v}_k) \geq a\}$. If such an index exists, we define $y_i = v_{k_i}$. Otherwise, we set $l = i - 1$. By our choice, we have $d(\tilde{y}_{i-1}, \tilde{v}_{k_i-1}) < a$. Applying now Corollary 5.13, the triangle inequality yields $d(\tilde{y}_{i-1}, \tilde{y}_i) \leq d(\tilde{y}_{i-1}, \tilde{v}_{k_i-1}) + d(\tilde{v}_{k_i-1}, \tilde{v}_{k_i}) \leq a + 1 + c_0 = a_1$, as we claimed.

By gluing together a path from 0 to $\tilde{y}_1$, paths from $\tilde{y}_i$ to $\tilde{y}_{i+1}$ and a path from $\tilde{y}_l$ to z, we get a path from 0 to z inside P whose length is thus larger than $\Lambda\|z\|$ and is at most la_1, for $d(\tilde{y}_i, \tilde{y}_{i+1}) \leq a_1$. Hence, $l \geq \frac{\Lambda\|z\|}{a_1}$. By definition of the function ρ in Corollary 5.14, $d(y_j, y_{j+1}) \geq \rho(\eta/3)$. Thus, the length of ϕ is at least $\frac{\Lambda}{a_1}\rho(\eta/3)\|z\|$, where $\rho(\eta/3)$ tends to infinity, as η tends to infinity.

Fix $R_0 \geq 0$. Then, for large enough η, $\rho(\eta/3) \geq \frac{a_1}{\Lambda} R_0 r(\mu)$, so the length of ϕ is at least $R_0 r(\mu)\|z\|$. Recall that n is the number of steps of the trajectory from $(0, j)$ to (z, j') and that the length of ϕ is at most $nr(\mu)$. We thus have $n \geq R_0\|z\|$. Note that R_0 can be taken arbitrarily large, provided that η is chosen large enough. Summing over all trajectories of the random walk from $g = (0, j)$ to $g_z = (z, j')$ that stay outside $N_\eta(P)$, we have

$$p_{j,j'}(0, z) \leq \sum_{n \geq R_0\|z\|} \mu^{*n}(g^{-1}g_z).$$

Lemma 5.15 shows that

$$p_{j,j'}(0,z) \leq \sum_{n \geq R_0\|z\|} \mathrm{e}^{-\alpha n} \leq \mathrm{e}^{-\alpha R_0\|z\|} \sum_{n\geq 0} \mathrm{e}^{-\alpha n}.$$

To prove Proposition 5.11, it suffices to choose R_0 so that $R_0\alpha > M$. □

We can now complete the proof of Proposition 5.7, using Corollary 4.10. Indeed, Lemma 5.9 shows that the induced chain on (arbitrary) bounded neighborhoods of P is strongly irreducible, while Lemma 5.10 shows that it is strictly sub-Markov (Condition a) of Corollary 4.10) and Proposition 5.11 shows that it has sufficiently high exponential moment (Condition b) of Corollary 4.10).

Thus, Corollary 4.10 implies that the Martin compactification of the induced chain on bounded neighborhoods of P coincides with the geometric compactification of P, and together with Lemma 5.8 this implies Proposition 5.7. □

To prove Proposition 5.5, we now show that we can reduce to the case of a sequence that stays in a uniform neighborhood of the parabolic subgroup P. The proof is based on the following strategy. Assume that the sequence g_n leaves every bounded neighborhood of P, but its projections to P still converge to a point θ in ∂P. Proposition 5.7 applied to $\pi_P(g_n)$ guarantees that $\pi_P(g_n)$ converges to a point in the Martin boundary. We want to prove that the same is true for g_n. In other words, we want to prove that K_{g_n} converges pointwise. Generalized Ancona inequalities show that to go from the basepoint e (the neutral element of Γ) or from an arbitrary point g to g_n, the random walk visits $\pi_P(g_n)$ with high probability. Thus, $G(g,g_n)$ is close to $G(g,\pi_P(g_n))G(\pi_P(g_n),g_n)$ and $G(e,g_n)$ is close to $G(e,\pi_P(g_n))G(\pi_P(g_n),g_n)$, so that $K_{g_n}(g)$ is close to $K_{\pi_P(g_n)}(g)$. Convergence for $K_{g_n}(g)$ then follows from convergence for $K_{\pi_P(g_n)}(g)$.

We now give a formal proof. Let $p_n = \pi_P g_n$ be a projection point of g_n to P. By assumption we have $\lim_{n\to\infty} p_n = \xi \in \partial_\mu\Gamma$. By Lemma 8.2 of [GGPY], there is a uniform $\delta > 0$ with $\liminf_{n\to\infty} \delta^f_{p_n}(g,g_n) \geq \delta$ for all $g \in \Gamma$. Let $\epsilon > 0$. Consider any $g \in \Gamma$. By Theorem 4.1, there is an $\eta > 0$ such that for large enough n,

$$G\big(g,g_n;B^c_\eta(p_n)\big) \leq \epsilon G(g,g_n) \tag{6}$$

and

$$G\big(e,g_n;B^c_\eta(p_n)\big) \leq \epsilon G(e,g_n). \tag{7}$$

Assume η is also large enough to satisfy Proposition 5.11. Decomposing a path from e to g_n according to its last visit to $B_\eta(p_n)$, we can write

$$G(e,g_n) = \sum_{u_n\in B_\eta(p_n)} G(e,u_n)G\big(u_n,g_n;B^c_\eta(p_n)\big) + G\big(e,g_n;B^c_\eta(p_n)\big) \tag{8}$$

and similarly,

$$G(g, g_n) = \sum_{u_n \in B_\eta(p_n)} G(g, u_n) G\big(u_n, g_n; B^c_\eta(p_n)\big) + G\big(g, g_n; B^c_\eta(p_n)\big). \tag{9}$$

By Proposition 5.7 we know that for any $u_n \in B(p_n, \eta)$ also converges to ξ. Then $G(g, u_n)/G(e, u_n)$ converges to $K_\xi(g)$ and this is independent of the sequence u_n. Hence, for large enough n, we have

$$(1-\epsilon) K_\xi(g) \le \frac{G(g, u_n)}{G(e, u_n)} \le (1+\epsilon) K_\xi(g). \tag{10}$$

Combining (6), (9) and (10), we obtain for all large n that

$$G(g, g_n) \le \sum_{u_n \in B_\eta(p_n)} (1+\epsilon) K_\xi(g) G(e, u_n) G\big(u_n, g_n; B^c_\eta(p_n)\big) + \epsilon G(g, g_n),$$

so that

$$(1-\epsilon) G(g, g_n) \le (1+\epsilon) K_\xi(g) \sum_{u_n \in B_\eta(p_n)} G(e, u_n) G\big(u_n, g_n; B^c_\eta(p_n)\big).$$

Then, using (8), $(1-\epsilon) G(g, g_n) \le (1+\epsilon) K_\xi(g) G(e, g_n)$. Similarly, using (7), (8), (9) and (10), we get a lower bound, so that for large enough n,

$$\frac{1-\epsilon}{1+\epsilon} K_\xi(g) \le \frac{G(g, g_n)}{G(e, g_n)} \le \frac{1+\epsilon}{1-\epsilon} K_\xi(g).$$

Since $\epsilon > 0$ is arbitrary we get that $K(g, g_n)$ converges to $K_\xi(g)$. This holds for every $g \in \Gamma$, so that g_n converges to ξ in the Martin boundary, completing the proof of Proposition 5.5. □

It follows now from Corollary 5.6 and Proposition 5.2 that a sequence $g_n \in \Gamma$ converges in the Martin compactification of Γ if one of the following conditions is satisfied:

(1) either g_n converges to a conical point of the Bowditch boundary,

(2) or for some parabolic subgroup $P \in \mathcal{P}_0$ and some $g \in \Gamma$, $g^{-1}\pi_{gP}(g_n)$ converges to a point of ∂P,

To complete the proof that the Martin boundary is a PBU-boundary we need to show the converse: namely that if g_n converges to a point in the Martin boundary, then it satisfies either (1) or (2).

Suppose g_n converges to a point in the Martin boundary. By Proposition 5.3, g_n converges to a point α in the Bowditch boundary. If α is conical, then (1) holds. Suppose now that α is parabolic with stabilizer gPg^{-1} ($P \in \mathcal{P}_0$, $g \in \Gamma$) and assume that (2) is not satisfied. We have $g^{-1}\pi_{gP}(g_n) \in \mathrm{proj}_P(g^{-1}g_n)$ and by Proposition 3.4 the latter set has a uniformly bounded diameter. Thus the elements $\pi_P(g^{-1}g_n)$ and $g^{-1}\pi_{gP}(g_n)$ are within a bounded distance of each other, hence by Lemma 3.3 $\pi_P(g^{-1}g_n)$ cannot converge to a point in ∂P. Since g_n converges to the parabolic limit point α, the quantity $||\pi_{gP}(g_n)||$ goes to infinity. Thus, there are sub-sequences h_n and h'_n of g_n with $\pi_P(g^{-1}h_n)$ and $\pi_P(g^{-1}h'_n)$ converging to different points of ∂P. By Proposition 5.7, $\pi_P(g^{-1}h_n)$ and $\pi_P(g^{-1}h'_n)$ converge to different points ξ and ξ' in $\partial_\mu\Gamma$. Furthermore, by Proposition 5.5, $g^{-1}h_n$ converges to the same point in $\partial_\mu\Gamma$ as $\pi_P(g^{-1}h_n)$ and $g^{-1}h'_n$ converges to the same point in $\partial_\mu\Gamma$ as $\pi_P(g^{-1}h'_n)$. Thus h_n and h'_n converge to different points of $\partial_\mu\Gamma$, contradicting our assumption on g_n.

This proves that the Martin boundary is a PBU-boundary, ending the proof of Theorem 1.3. □

6. Minimality

In this section we prove Theorem 1.5, namely the minimality of the Martin boundary. We will use the following, which is a direct consequence of [Dus, Proposition 6.3].

Proposition 6.1. *There exists an $\eta_0 > 0$ such that for $\eta > \eta_0$ the following holds. For any distinct $\alpha_0, \alpha_1 \in \partial P$ there exists a neighborhood $\mathcal{U}$ of α_1 (not containing α_0) in ∂P and a sequence g_n of $N_\eta P$ such that*

(1) *either $K_\alpha(g_n)$ tends to infinity, uniformly over $\alpha \in \mathcal{U}$ and $K_{\alpha_0}(g_n)$ stays bounded away from infinity.*

(2) *or $K_\alpha(g_n)$ stays bounded away from 0, uniformly over $\alpha \in \mathcal{U}$ and $K_{\alpha_0}(g_n)$ converges to 0.*

We now prove the following.

Theorem 6.2. *Let Γ be hyperbolic relative to a collection of virtually abelian subgroups. Let μ be a probability measure on Γ whose finite support generates Γ as a semigroup. Then every point of the Martin boundary $\partial_\mu\Gamma$ corresponds to a minimal harmonic function.*

Proof. By Theorem 1.3, $\partial_\mu\Gamma$ is a PBU-boundary. This means that there is a Γ-equivariant surjective map $F : \partial_\mu\Gamma \to \partial_B\Gamma$ such that if $a \in \partial_B\Gamma$ is conical, $F^{-1}(a)$ is a single point and if $a \in \partial_B\Gamma$ is parabolic, $F^{-1}(a) = \partial P$ where P is the stabilizer of a and ∂P denotes its geometric boundary. Notice that $F : \partial_\mu\Gamma \to \partial_B\Gamma$ is the same map as the the map $\psi : \partial_{\mathcal{M}}\Gamma \to \partial_B\Gamma$ constructed in [GGPY, Corollary 1.7].

Let $\alpha_0 \in \partial_\mu \Gamma$. Then K_{α_0} is a positive harmonic function. By the Choquet representation theorem, there exists a finite Borel measure $\nu_0 = \nu^{\alpha_0}$ on the Martin boundary, with support contained in the minimal Martin boundary $\partial_\mu^m \Gamma$ such that for all $g \in \Gamma$

$$K_{\alpha_0}(g) = \int_{\partial_\mu^m \Gamma} K_\alpha(g) d\nu_0(\alpha).$$

To prove minimality of α_0 it suffices to show that the support of ν_0 consists of the single point α_0. In Corollary 7.9 of [GGPY] the authors deduce the following result from the inequality (1) (see Theorem 4.1).

Lemma 6.3. *The support of* ν_0 *is contained in* $F^{-1}\big(F(\alpha_0)\big)$.

If $F(\alpha_0)$ is conical, then $F^{-1}(F(\alpha_0))$ is a single point [GGPY, Corollary 7.14]. Lemma 6.3 then implies that the support of ν_0 is a single point so that α_0 is minimal.

On the other hand, if α_0 is a parabolic point of the Bowditch boundary with stabilizer P, Theorem 1.3 implies that $F^{-1}(F(\alpha_0)) = \partial P$. Thus we know that ν_0 is supported on $\partial P \cap \partial_\mu^m \Gamma$.

Now, suppose α_1 is a point of ∂P distinct from α_0. By Proposition 6.1, there exists a neighborhood $\mathcal{U}$ of α_1, not containing α_0, and a sequence g_n such that

(1) either $K_\alpha(g_n)$ tends to infinity, uniformly over $\alpha \in \mathcal{U}$ and $K_{\alpha_0}(g_n)$ stays bounded away from infinity.

(2) or $K_\alpha(g_n)$ stays bounded away from 0, uniformly over $\alpha \in \mathcal{U}$ and $K_{\alpha_0}(g_n)$ converges to 0.

Thus, in the first case, for large enough n and for all $\alpha \in \mathcal{U}$, we have $K_\alpha(g_n) \geq R_n$, where R_n tends to infinity. Then by definition,

$$K_{\alpha_0}(g_n) = \int_{\alpha \in \partial P} K_\alpha(g_n) d\nu_0(\alpha) \geq \int_{\alpha \in U} K_\alpha(g_n) d\nu_0(\alpha) \geq R_n \nu_0(\mathcal{U}).$$

As $K_{\alpha_0}(g_n)$ stays bounded away from infinity, it follows that $\nu_0(\mathcal{U}) = 0$.

In the second case, for large enough n and for all $\alpha \in \mathcal{U}$, we have $K_\alpha(g_n) \geq C$ for some constant C. Then,

$$K_{\alpha_0}(g_n) = \int_{\alpha \in \partial P} K_\alpha(g_n) d\nu_0(\alpha) \geq C \nu_0(\mathcal{U}).$$

As $K_{\alpha_0}(g_n) \to 0$ as $n \to \infty$ it follows again that $\nu_0(\mathcal{U}) = 0$. In both cases, the support of ν_0 does not contain α_1. We conclude that the support of ν_0 consists only of α_0, so K_{α_0} must be minimal. □

7. Another viewpoint on the PBU-boundary

The aim of this section is to prove that the PBU-compactification of a hyperbolic group relatively to a system of virtually abelian subgroups constructed in the paper is equivalent to a well-known compactification constructed for general relatively hyperbolic groups by F. Dahmani [Dah1].

7.1. Relative Cayley graph. In this section we summarize several facts which will be used further on. Let Γ be a group generated by a finite symmetric set S. We denote by $\mathcal{G}$ the Cayley graph $C_S\Gamma$ of Γ with respect to the system S.

Let us also fix a family $\mathcal{P}$ of subgroups of Γ satisfying two properties: it is invariant under conjugation in Γ and there is a finite subset $\mathcal{P}_0 = \{P_1, \dots, .P_k\} \subset \mathcal{P}$ such that every element $P \in \mathcal{P}$ is conjugate to one of P_i. We call the system of the left cosets $\{gP \ : \ P \in \mathcal{P}_0, \ g \in \Gamma\}$ the *system of horospheres* and each of its element is called a *horosphere* (see [GP3, Section 5] for more explanations on horospheres).

Refine the graph $\mathcal{G}$ by adding an edge of length 1 between each pair of vertices belonging to the same horosphere. The obtained graph is called the *relative Cayley* graph with respect to the pair $(S, \mathcal{P}_0)$ and we denote it by Δ. The choice of the subset $\mathcal{P}_0$ for Δ plays a similar role as the choice of a finite generator set S for the Cayley graph $\mathcal{G}$ and the graph Δ does not depend on $\mathcal{P}_0$ up to quasi-isometry.

Let d denote the word distance of $\mathcal{G}$ and $\overline{d}$ the path distance of Δ. To distinguish paths in $\mathcal{G}$ and Δ we call them $\mathcal{G}$*-paths* (or simply *paths*) and Δ*-paths* (or *relative paths*) respectively. Every Δ-path l *lifts* to a $\mathcal{G}$-path in the following way: its lift $\tilde{l}$ has the same non-horospherical edges as l and every horospherical edge of l is replaced by a geodesic interval in $\mathcal{G}$ with the same endpoints.

From now on, we assume that the group Γ is relatively hyperbolic with respect to a system of maximal parabolic subgroups $\mathcal{P}$. By [Tuk, Theorem 1B] there is a finite subset $\mathcal{P}_0 = \{P_1, \dots, P_k\}$ of $\mathcal{P}$ such that every $P \in \mathcal{P}$ is conjugate to one of P_i $(i \in \{1, \dots, k\})$. So $\{gP \ : \ P \in \mathcal{P}_0, \ g \in \Gamma\}$ is a system of horospheres. One of the main properties of the relative graph Δ in this case is that it is hyperbolic [Far], [Bow] (see also [GP3, Proposition 7.1] for a direct proof of this fact and Remark 7.2 concerning different definitions of the relative hyperbolicity).

Recall that a path $\gamma : J \to X$ in a metric space (X, d_X) is called quasi-geodesic if there is an affine function $\alpha(t) = At + B$ $(t \geq 0)$, called a distortion function, which satisfies

$$\mathrm{diam}(J) \leq \alpha(\mathrm{diam}(\gamma(\partial J)), \tag{11}$$

where $\mathsf{diam}(\cdot)$ denotes the diameter of a set. The constants A and B are called *parameters* of the quasi-geodesic. In our case the role of X is played by the graphs $\mathcal{G}$ or Δ, so $J \subset \mathbb{N}$ and we say respectively that γ is d*-quasi-geodesic* (or simply *quasi-geodesic*) or $\overline{d}$*-quasi-geodesic* (or *relative quasi-geodesic*). In particular if α is the identity function our curve γ is a geodesic. Since Δ is hyperbolic, every $\overline{d}$-quasi-geodesic with fixed parameters stays in a uniformly bounded distance from a $\overline{d}$-geodesic having the same endpoints [Gro, Proposition 7.2.A].

Even though the graph $\mathcal{G}$ is not hyperbolic in general, the lifts of $\overline{d}$-quasi-geodesics to $\mathcal{G}$ have properties close to those of d-quasi-geodesics in a hyperbolic space. The following lemma confirms this fact and will be used in the next subsection.

Lemma 7.1. *Let Γ be a hyperbolic group relative to a system $\mathcal{P}$ of parabolic subgroups. Consider two curves l and m in the Cayley graph $\mathcal{G}$ having the same endpoints o and x. Assume that l is a d-quasi-geodesic and m is a $\overline{d}$-quasi-geodesic both having the parameters A and B. Assume also that for every horosphere $S = gP$ $(g \in \Gamma, P \in \mathcal{P}_0)$ the curve m can have at most one edge in Δ with endpoints in S. Then there exists a constant C, only depending on A and B, such that for every vertex $v \in m$ we have $d(v, l) \le C$.*

When m is a $\overline{d}$-geodesic the lemma is proved by Hruska in [Hru, Lemma 8.8], whose proof uses a lot of preliminary results. We provide below a direct proof based on several tools already used in the paper.

We call the above condition of the intersections of relative quasi-geodesics with horospheres *one-edge horospherical intersection property.* An important subclass of relative quasi-geodesics form $\overline{d}$-geodesics which obviously satisfy this property.

Proof of the Lemma. Let m' denote a lift of m to $\mathcal{G}$. We first claim that m' is a d-quasi-geodesic whose parameters only depend on A and B and every vertex of $m' \cap m$ is R-transitional for a constant R which also only depends on A and B. Indeed if m was a $\overline{d}$-geodesic the claim would directly follow from Propositions 6.1 and 7.8 of [GP3]. Actually, the first part of the proof of Proposition 6.1 consists in proving that m' is an α-distorted path where the distortion function α (see formula (11)) is a quadratic polynomial. To check this statement in our situation, one needs to replace the $\overline{d}$-distance n between the endpoints of m by $An + B$ for the parameters A and B of the quasi-geodesic m. The distortion function α of the lift m' obviously remains a quadratic polynomial after such a modification. Then, using that m can contains at most one edge in each horosphere of Δ, the rest of the proof of Proposition 6.1 applies without any change. It follows that the whole curve m' is α-distorted for a quadratic polynomial α, and every vertex of $m \cap m'$ is R-transitional in $\mathcal{G}$. Then all assumptions of [GP3, Proposition 7.8] are satisfied and it implies that m' is a quasi-geodesic in $\mathcal{G}$ with uniformly bounded parameters (depending on those of m), confirming our claim.

By the claim every $v \in m$ is an R-transition point of the curve m' having the same endpoints o and x as the d-quasi-geodesic l. By Proposition 3.2, there is a constant $\delta > 0$ such that the Floyd distance $\delta_v^f(o,x)$ satisfies $\delta_v^f(o,x) \geq \delta$. Then Karlsson's lemma [Kar, Lemma 1] implies that $d(v,l) \leq C$ for a uniform constant C depending on δ and the parameters of the quasi-geodesic l.

We note that Karlsson proved this lemma in assumption that l is a d-geodesic but his proof works without any changes in the case of quasi-geodesics (see [GP1] where this and more general cases are discussed). This concludes the proof. □

To finish the discussion about the relative quasi-geodesics it is worth to mention that the proof of the lemma works for the first entry (or the last exit) horospherical point u of any relative quasi-geodesic m (without assuming that m has one-edge horospherical intersection property). Then the above argument also shows that the distance $d(u,l)$ is uniformly bounded if $u \in m$ is such a point and m is a relative quasi-geodesic with bounded parameters.

There is a more general assumption than our one-edge horospherical intersection property which is due to B. Farb [Far]. His condition that relative quasi-geodesics are *without backtracking* is equivalent to that each horospherical part of such a curve m in Δ is connected and has a uniformly bounded $\overline{d}$-length depending only on the parameters of m. Such curves were used in [Far] to define BCP-property which we do not need to use here. We simply mention that the hyperbolicity of the relative graph together with the BCP-property is equivalent to the fact that the group Γ is relatively hyperbolic. Furthermore, we note that this BCP-property follows from the above lemma, combined with the fact that the intersection of different horospheres has a uniformly bounded d-diameter, see [GP3, Corollary 5.7].

7.2. Dahmani's geometric boundary. In [Dah1], the author introduces a compactification of hyperbolic groups relative to a system of parabolic subgroups $\mathcal{P}$. Dahmani's construction has an inductive nature: once one knows a "good" compactification of parabolic subgroups then a "good" compactification is obtained for the whole group. To define Dahmani's compactification we need to introduce few more notions. Let Γ admits a minimal geometrically finite action on a compactum T. Denote by Par the fixed-point set of parabolic subgroups for this action and by Λ_c the set of conical points.

Assumptions. Every maximal parabolic subgroup $P \in \mathcal{P}$ admits a metrizable compactification $P \cup \partial P$ such that P is dense in it. Furthermore for every finite subset F of P and for every open cover $\mathcal{U}$ of $P \cup \partial P$, all translates of F by P but finitely many are contained in an element of $\mathcal{U}$. In this case we say that finite sets *fade* at infinity [Bes], [Dahl].

We also assume that the action of Γ on elements of $\mathcal{P}$ continuously extends to an action on their boundaries: for $x_n \in P$ and $g \in \Gamma$, $x_n \to \xi \in \partial P$ implies that $gx_n \to \xi' = g(\xi) \in g(\partial P)$. Furthermore we ask this extension to be equivariant, that is $g(\partial P) = \partial(gPg^{-1})$ and we also set $\partial(gP) = g(\partial P)$.
End of Assumptions.

The boundary $\partial\Gamma$ is defined in [Dahl] as follows:

$$\partial\Gamma := \Big(\bigsqcup_{P \in \mathcal{P}} \partial P\Big) \bigsqcup \Lambda_c = \Big(\bigsqcup_{\substack{P_i \in \mathcal{P}_0 \\ g \in \Gamma}} g(\partial P_i)\Big) \bigsqcup \Lambda_c, \tag{12}$$

where $\mathcal{P}_0$ is the maximal subset of $\mathcal{P}$ of non-conjugate subgroups.

The proof of Proposition 3.6 above, applying without changes in this case, shows that the topology on the space $\Gamma \cup \partial\Gamma$ is uniquely defined by the following definition of convergence of sequences of elements of Γ to the boundary points in $\partial\Gamma$.

Definition 7.2 (Dahmani's compactification [Dahl, Definition 3.3]). Let Γ be a hyperbolic group relatively to a system of parabolic subgroups $\mathcal{P}$. We fix a metrisable compactification for every $P \in \mathcal{P}$ satisfying the above Assumptions. Say that a sequence g_n in Γ converges to a point $\xi \in \partial\Gamma$ if one of the two following cases happens:

- either $\xi \in \Lambda_c$ is a conical point then g_n converges to ξ in the Bowditch compactification of Γ;
- or $\xi \in \partial(gP)$ where $P \in \mathcal{P}_0$ is a parabolic subgroup and $g \in \Gamma$. Then there exist a sequence $u_n = gh_n \in gP, h_n \in P$ such that h_n tends to $g^{-1} \cdot \xi$ and a relative geodesic l_n between g_n and u_n, for which u_n is the first entry point of l_n in the horosphere gP.

Remark. In [Dahl, Definition 3.3] the author uses curves l_n satisfying technical conditions. They are relative quasi-geodesics outside a compact set. They have uniformly bounded parameters. Furthermore, they are assumed to be *left reduced*, a condition which implies that the left endpoint of such a curve is not followed by a horospherical edge. By Lemma 7.1, it is easy to see that we can replace such a curve l_n by a relative geodesic having the same endpoints, as it is done in the definition above.

Theorem 7.3 ([Dahl, Theorem 3.1]). *Let Γ be a hyperbolic group relative to a system $\mathcal{P}$. Assume that each $P \in P$ admits a metrisable topology satisfying the assumptions above. Then there exists a topology on $\Gamma \cup \partial\Gamma$ satisfying Definition 7.2 such that the topological space is compact and metrisable.*

To compare Dahmani's compactification with that of PBU we will use the geometric compactification of parabolic subgroups introduced in Section 3.2. Recall that in our context, every parabolic subgroup $P \in \mathcal{P}$ is virtually abelian and convergence to ∂P is given as follows.

Identifying P with $\mathbb{Z}^k \times \{1, \dots, N\}$ and $p \in P$ with $(z, j) \in \mathbb{Z}^k \times \{1, \dots, N\}$, we say that a sequence (z_n, j_n) converges in ∂P if z_n tends to infinity and $\frac{z_n}{\|z_n\|}$ converges to a point θ in $\mathbb{S}^{k-1}$. Formally, one can choose sets of the form $U_{n,m}(\theta) \times \{1, \dots, N\}$ to form a countable system of neighborhoods of a point θ in the boundary, where

$$U_{n,m}(\theta) = V_n(\theta) \cup \left\{ z \in \mathbb{Z}^k, \|z\| \geq m, \frac{z}{\|z\|} \in V_n(\theta) \right\},$$

with $V_n(\theta)$ a neighborhood of θ in $\mathbb{S}^{k-1}$.

Defining in such a way the topology on maximal subset $\mathcal{P}_0 = \{\mathcal{P}_1, \dots, P_k\}$ of non-conjugate elements of $\mathcal{P}$ we then define the topology on every $P \in \mathcal{P}$ to verify $\partial P = g(\partial P_i)$ if $P = gP_ig^{-1}, g \in \Gamma, i \in \{1, \dots, k\}$. With the following lemma we obtain that the topology defined in this way on the set $\bigsqcup_{P \in \mathcal{P}}(P \cup \partial P)$ satisfies all the assumptions above.

Lemma 7.4. *Finite subsets of the space $P \cup \partial P$, equipped with this topology, fade at infinity.*

Proof. By Lemma 3.3, if z_n and z'_n are two sequences in $\mathbb{Z}^k$ such that $\|z_n - z'_n\|$ is bounded by a constant and z_n converges to a point θ in $\partial\mathbb{Z}^k$, then z'_n also converges to θ in $\partial\mathbb{Z}^k$. By induction, the same holds for a finite number of sequences. That is, if $z_n^{(1)}, \dots, z_n^{(j)}$ are sequences in $\mathbb{Z}^k$ such that $\|z_n^{(j_1)} - z_n^{(j_2)}\|$ is bounded for every j_1, j_2 and such that one of them converges in $\partial\mathbb{Z}^k$, then they all converge to the same point. This property is called *perspectivity* property in [Ger2].

To finish the proof, assume by contradiction that F is a finite subset of P and $\mathcal{U}$ is an open cover of $P \cup \partial P$ such that there are infinitely many translates of F that are not contained in one of the open sets in $\mathcal{U}$. Denote these translates by $p_n F$ $(p_n \in P)$. Let $f \in F$. The sequence $p_n \cdot f$ tends to infinity. Up to taking a sub-sequence, it converges to a point ξ in ∂P. By the above perspectivity property for every $f' \in F$, $p_n \cdot f'$ all eventually belong to an arbitrary small neighborhood U_ξ of ξ. Then choosing U_ξ inside of an element of $\mathcal{U}$ containing ξ we obtain a contradiction. □

The following proposition is the main result of this section.

Proposition 7.5. *Let* Γ *be a hyperbolic group relatively to a system* $\mathcal{P}$ *of virtually abelian subgroups. For every* $P \in \mathcal{P}$ *we fix a geometric topology on* $P \cup \partial P$ *as above. Then the topologies of Dahmani's compactification and of PBU-compactification of* Γ *coincide.*

Proof. The convergence to a conical point is defined in the same way, so the only case we need to consider is when a sequence $g_n \in \Gamma$ converges to a point in ∂P for some $P \in \mathcal{P}$. So we need to prove that the convergence in one of the following topologies implies the convergence in the other one.

(1) (Dahmani's topology) There is a sequence of relative geodesics l_n between g_n and their first entry points $u_n = gh_n$ to a fixed horosphere gP, where $P \in \mathcal{P}_0, g \in \Gamma$, such that h_n converges to a point $\xi \in \partial P$.

(2) (BPU topology) The projections $v_n = \pi_{gP}(g_n)$ of g_n on gP satisfy that $g^{-1}v_n$ converges to $\xi \in \partial P$.

The proof is a direct consequence of the following lemma due to A. Sisto whose original proof is based on the BCP-property and other results which we do not use. We obtain the Lemma as a simple consequence of Lemma 7.1.

Lemma 7.6 (Sisto's lemma [Sis, Lemma 1.15.2] [Bounded Geodesic Image]). *With the above notations, there exists a constant* $M > 0$ *such that* $d(u_n, v_n) \le M$ *for the word distance* d *and* $n \in \mathbb{N}$.

Proof of the Lemma. Since the points u_n and v_n belong to the same horosphere gP, there is an edge e_n between them in the relative graph Δ. Then the curve $\widetilde{l}_n = l_n \cup e_n$ is a relative quasi-geodesic with endpoints v_n and g_n and with uniform parameters $A = B = 1$ (see Section 7.1). Furthermore since l_n is a $\overline{d}$-geodesic, the curve $l'_n \subset \Delta$ satisfies the one-edge horospherical intersection assumption.

By Lemma 7.1, there exists a uniform constant C such that for $u_n \in l'_n$, one has $d(u_n, [g_n, v_n]) \le C$ where $[g_n, v_n]$ is a d-geodesic between g_n and v_n. Denote by $y_n \in [g_n, v_n]$ a point such that $d(y_n, u_n) = d(u_n, [g_n, v_n]) \le C$. Since the geodesic $[g_n, v_n]$ realizes the distance $d(g_n, gP)$ and $u_n \in gP$, we have $d(y_n, v_n) \le d(y_n, u_n)$. Hence, $d(u_n, v_n) \le d(y_n, v_n) + d(y_n, u_n) \le 2d(y_n, u_n) \le 2C = M$. The lemma is proved. □

Both sequences u_n and v_n belong to the same horosphere gP $(P \in \mathcal{P}, g \in \Gamma)$. By Lemma 7.6 we have $d(u_n, v_n) \leq M$, hence $d(g^{-1}u_n, g^{-1}v_n) \leq M$. Thus, by Lemma 7.4, $h_n = g^{-1}u_n$ converges to a point in ∂P if and only if $g^{-1}v_n$ converges to the same point. This concludes the proof. □

7.3. Some questions and remarks. Let us make here some comments and ask further questions related to the above results. Combining the results of Theorem 1.3 and [GGPY, Theorem 1.3] we have the following corollary, which seems to be interesting independently of the random walks context.

Corollary 7.7. *Let Γ be a finitely generated group, hyperbolic relative to a collection of infinite virtually abelian subgroups. Then, there exists an equivariant and continuous surjective map from the PBU-boundary to the Floyd boundary of Γ.*

Bestvina [Bes] introduced the notion of $\mathcal{Z}$-boundaries for groups. Whenever ∂P is a $\mathcal{Z}$-boundary for P, Dahmani [Dah1, Theorem 4.1] showed that the construction presented above yields a $\mathcal{Z}$-boundary $\partial\Gamma$ for Γ. When parabolic subgroups are virtually abelian, the geometric boundary ∂P coincides with the CAT(0) boundary of P as noted in Section 3.2, see also [BH, Remark 7.3 (2)]. Hence, it is a $\mathcal{Z}$-boundary, according to [AG, Lemma 8]. Combining now the results of Corollary 1.4, Corollary 7.7 and Proposition 7.5, we get the following.

Corollary 7.8. *Let Γ be a finitely generated group, hyperbolic relative to a collection of infinite virtually abelian subgroups. Then, there exists an equivariant and continuous surjective map from a $\mathcal{Z}$-boundary to the Floyd boundary. There also exists an equivariant and continuous surjective map from a $\mathcal{Z}$-boundary of Γ to the Bowditch boundary, which is 1-to-1 at conical points and the preimage of a parabolic point coincides with the $\mathcal{Z}$-boundary of its stabilizer.*

Since the Martin boundary does not depend on different peripheral structures, but only on the random walk, Theorem 1.3 also implies the following.

Corollary 7.9. *If Γ is hyperbolic relative to two different collections of infinite virtually abelian subgroups, then two corresponding PBU-boundaries, constructed for each relatively hyperbolic structure, are equivariantly homeomorphic.*

We conclude the discussion with few questions related to the above results. We only considered PBU-boundaries for relatively hyperbolic with respect to virtually abelian parabolic subgroups. Theorem 7.3 shows that one can extend the definition of a PBU-boundary to any relatively hyperbolic group. However its relation with the corresponding Martin boundary is not known when the parabolic subgroups are not virtually abelian.

1. *Are Corollary 7.7 and Corollary 7.8 true for more general relatively hyperbolic groups? We conjecture that they are still true for hyperbolic groups relatively to virtually nilpotent groups, but for more general relatively hyperbolic groups there might exist counter-examples.*

The following question, motivated by the proof of Proposition 5.5, also seems to be interesting.

2. *Assume that Γ is hyperbolic relative to a system $\mathcal{P}$. Describe the class $\mathcal{P}$ for which our key Proposition 5.5 is true. Namely, let (x_n) be a sequence of elements of Γ tending to infinity and y_n be a projection of x_n onto a horosphere corresponding to a parabolic point p in the Bowditch boundary. Then is it true that x_n converges to a preimage of p in the Martin boundary if and only if y_n converges to the same point?*

In a forthcoming preprint by Gerasimov, Potyagailo and de Souza, it is proved that this condition defines a unique PBU-compactification of a group hyperbolic relative to a system of parabolic subgroups with fixed boundaries.

Acknowledgements. We would like to thank Sebastien Gouëzel for helpful advice. We also thank Wenyuan Yang for helpful conversations.

The authors are thankful to the research grant LABEX CEMPI for the organization of a conference on the Geometric Group Theory at the University of Lille where we all met. The second and the third authors are grateful to the LABEX CEMPI grant for providing them research fellowships for their stay in Lille after the conference. The third and the fourth authors are thankful to the Mathematical Institute in Oberwolfach (MFO) and to the Max-Planck Institute of Mathematics in Bonn where a part of this work was done; they were also partly supported by MATH-AmSud grant (code 18-MATH-08). The second author was partially supported by NSF grant DMS-1401875.

We are very much thankful to the referee(s) for very useful remarks and corrections which certainly made our paper better.

References

[AG] F. Ancel and C. Guilbault, Z-compactifications of open manifolds. *Topology* **38** (1999), 1265–1280. Zbl 0951.57011 MR 1690157

[Anc] A. Ancona, Positive harmonic functions and hyperbolicity. In *Potential theory-surveys and problems*. Lecture notes in mathematics, Springer (1988). pp. 1–23. Zbl 0677.31006 MR 0973878

[BB] S. BLACHÈRE and S. BROFFERIO, Internal diffusion limited aggregation on discrete groups having exponential growth. *Probabability Theory and Related Fields* **137** (2007), 323–343. Zbl 1106.60078 MR 2278460

[Bes] M. BESTVINA, Local homology properties of boundaries of groups. *Michigan Mathematical Journal* **43** (1996), 123–129. Zbl 0872.57005 MR 1381603

[BH] M. BRIDSON and A. HÄFLIGER, *Metric Spaces of Non-Positive Curvature*. Springer (1999). Zbl 0988.53001 MR 1744486

[BL] W. BALLMANN and F. LEDRAPPIER, Discretization of positive harmonic functions on Riemannian manifolds and Martin boundary. In *Actes de la table ronde de géométrie différentielle*, vol. 1. Séminaires et Congrès, Société mathématique de France (1996). pp. 77–92. Zbl 0885.53037 MR 1427756

[Bow] B. BOWDITCH, Relatively hyperbolic group. *International Journal of Algebra and Computation* **22** (2012), 66 pp. Zbl 1259.20052 MR 2922380

[Dah1] F. DAHMANI, Classifying spaces and boundaries for relatively hyperbolic groups. *Proceedings of the London Mathematical Society* **86** (2003), 666–684. Zbl 1031.20039 MR 1974394

[Dah2] —— Combination of convergence groups. *Geometry and Topology* **7** (2003), 933–963. Zbl 1037.20042 MR 2026551

[Dus] M. DUSSAULE, The Martin boundary of a free product of abelian groups (2017). *Ann. Inst. Fourier* 70 (2020), 313–373. Zbl 1442.31007 MR 4105941

[Far] B. FARB, Relatively hyperbolic groups. *Geometric and Functional Analysis* (*GAFA*) **8** (1998), 810–840. Zbl 0985.20027 MR 1650094

[Ger1] V. GERASIMOV, Expansive convergence groups are relatively hyperbolic. *Geometric and Functional Analysis* **19** (2009), 137–169. Zbl 1226.20037 MR 2507221

[Ger2] —— Floyd maps for relatively hyperbolic groups. *Geometric and Functional Analysis* **22** (2012), 1361–1399. Zbl 1276.20050 MR 2989436

[GGPY] I. GEKHTMAN, V. GERASIMOV, L. POTYAGAILO and W. YANG, Martin boundary covers Floyd boundary. *Inventiones mathematicae* **223** (2021), 759–809.

[GP1] V. GERASIMOV and L. POTYAGAILO, Quasi-isometric maps and floyd boundaries of relatively hyperbolic groups. *Journal of the European Mathematical Society* **15** (2013), 2115–2137. Zbl 1292.20047 MR 3120738

[GP2] —— Non-finitely generated relatively hyperbolic groups and Floyd quasiconvexity. *Groups, Geometry, and Dynamics* **9** (2015), 369–434. Zbl 1364.20032 MR 3356972

[GP3] —— Quasiconvexity in the relatively hyperbolic groups. *Journal für die reine und angewandte Mathematik* **710** (2016), 95–135. Zbl 1361.20031 MR 3437561

[Gro] M. GROMOV, Hyperbolic groups. In S. GERSTEN (ed.), *Essays in Group Theory*. M.S.R.I. Publications **8**, Springer (1987). pp. 75–265. Zbl 0634.20015 MR 0919829

[Gui] V. GUIRARDEL, Limit groups and groups acting freely on $\mathbb{R}^n$-trees. *Geometry and Topology* **8** (2004), 1427–1470. Zbl 1114.20013 MR 2119301

[HK] G. Hruska and B. Kleiner, Hadamard spaces with isolated flats. *Geometry and Topology* **9** (2005), 1501–1538. Zbl 1087.20034 MR 2175151

[Hru] G. Hruska, Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebra, Geometry and Topology* **10** (2010), 1807–1856. Zbl 1202.20046 MR 2684983

[Kar] A. Karlsson, Free subgroups of groups with nontrivial Floyd boundary. *Communications in Algebra* **31** (2003), 5361–5376. Zbl 1036.20032 MR 2005231

[Kes] H. Kesten, Full Banach mean values on countable groups. *Mathematica Scandinavica* **7** (1959), 146–156. Zbl 0092.26704 MR 0112053

[KSK] J. Kemeny, L. Snell and R. Knapp, *Denumerable Markov Chains*. Springer (1976). Zbl 0348.60090 MR 0407981

[Mar] G. Margulis, Positive harmonic functions on nilpotent groups. *Soviet Mathematics Doklady* **7** (1966), 241–244. Zbl 0187.38902 MR 0222217

[NS] P. Ney and F. Spitzer, The Martin boundary for random walk. *Transactions of the American Mathematical Society* **11** (1966), 116–132. Zbl 0141.15601 MR 0195151

[Rua] K. Ruane, CAT(0) boundaries of truncated hyperbolic space. *Topology Proceedings* **29** (2005), 317–331. Zbl 1090.53040 MR 2182938

[Saw] S. Sawyer, Martin boundaries and random walks. *Contemporary mathematics* **206** (1997), 17–44. Zbl 0891.60073 MR 1463727

[Sen] E. Seneta, *Non-negative Matrices and Markov Chains*. Springer (1981). Zbl 0471.60001 MR 0719544

[Sis] A. Sisto, Projections and relative hyperbolicity. *L'Enseignement Mathématique* **59** (2013), 165–181. Zbl 1309.20036 MR 3113603

[Tuk] P. Tukia, Conical limit points and uniform convergence groups. *Crelle's Journal* **501** (1998), 71–98. Zbl 0909.30034 MR 1637829

[TW] B. Tshishiku and G. Walsh, On groups with S^2 Bowditch boundary (2017). arXiv:1710.09018v2 MR 4167021

[Woe1] W. Woess, A description of the Martin boundary for nearest neighbour random walks on free products. In *Probability Measures on Groups VIII*. Lecture notes in mathematics, Springer (1986). pp. 203–215. Zbl 0611.60006 MR 0879007

[Woe2] —— *Random Walks on Infinite Graphs and Groups*. Cambridge University Press (2000). Zbl 0951.60002 MR 1743100

[Yam] A. Yaman, A topological characterisation of relatively hyperbolic groups. *Journal for Pure and Applied Mathematics* **566** (2004), 41–89. Zbl 1043.20020 MR 2039323

(*Reçu le 18 août 2019*)

Matthieu Dussaule, UMR de Mathématiques,
Université de Nantes, 2 chemin de la Houssinière,
44300 Nantes, France

e-mail: matthieu.dussaule@univ-nantes.fr

Ilya GEKHTMAN, Department of Mathematics,
Technion – Israeli Institute of Technology,
32000 Haifa, Israel

e-mail: ilyagekh@gmail.com

Victor GERASIMOV, Departamento de Matemática,
Universidade Federal de Minas Gerais,
Av. Antônio Carlos 6627, Caixa Postal 702,
30161-970 Belo Horizonte, MG, Brasil

e-mail: victor.gerasimov@gmail.com

Leonid POTYAGAILO, UFR de Mathématiques,
Université de Lille,
59655 Villeneuve d'Ascq, France

e-mail: leonid.potyagailo@univ-lille.fr