Forum of Mathematics, Sigma (2025), **Vol. 13**:e150 1–50
doi:10.1017/fms.2025.10094

CAMBRIDGE UNIVERSITY PRESS

RESEARCH ARTICLE

Rigidity phenomena and the statistical properties of group actions on CAT(0) cube complexes

Stephen Cantrell[1] and Eduardo Reyes[2]

[1]Department of Mathematics, University of Warwick, Coventry, CV4 7AL, UK;
E-mail: stephen.cantrell@warwick.co.uk.
[2]Facultad de Matemáticas, Pontificia Universidad Católica de Chile (PUC), Avenida Vicuña Mackenna 4860, Santiago, Chile;
E-mail: eduardoreyes@uc.cl (corresponding author).

Received: 17 October 2023; **Revised:** 23 May 2025; **Accepted:** 23 May 2025

2020 Mathematical Subject Classification: *Primary* – 20F65; *Secondary* – 20F67, 37D35

Abstract
We compare the marked length spectra of some pairs of proper and cocompact cubical actions of a nonvirtually cyclic group on CAT(0) cube complexes. The cubulations are required to be virtually co-special, have the same sets of convex-cocompact subgroups, and admit a contracting element. There are many groups for which these conditions are always fulfilled for any pair of cubulations, including nonelementary cubulable hyperbolic groups, many cubulable relatively hyperbolic groups, and many right-angled Artin and Coxeter groups.

For these pairs of cubulations, we study the Manhattan curve associated to their combinatorial metrics. We prove that this curve is analytic and convex, and a straight line if and only if the marked length spectra are homothetic. The same result holds if we consider invariant combinatorial metrics in which the lengths of the edges are not necessarily one. In addition, for their standard combinatorial metrics, we prove a large deviations theorem with shrinking intervals for their marked length spectra. We deduce the same result for pairs of word metrics on hyperbolic groups.

The main tool is the construction of a finite-state automaton that simultaneously encodes the marked length spectra of both cubulations in a coherent way, in analogy with results about (bi)combable functions on hyperbolic groups by Calegari-Fujiwara [14]. The existence of this automaton allows us to apply the machinery of thermodynamic formalism for suspension flows over subshifts of finite type, from which we deduce our results.

1. Introduction

In this work we study rigidity phenomena and the statistical properties of group actions on CAT(0) cube complexes and the methods we use exploit the interplay between geometric group theory and dynamics. Group actions on CAT(0) cube complexes are nowadays a central object of study. Since the influential work of Sageev [75], we have known that many groups admit proper and cubical actions on CAT(0) cube complexes (if in addition the actions are cocompact, in the sequel they are referred to as *cubulations*). The list includes small cancellation groups [60, 85], many 3-manifold groups [3, 43, 69, 70, 84], Coxeter groups [62], many Artin groups [21, 40], random groups at low density [63, 64], 1-relator groups with torsion [57, 83, 86], hyperbolic free-by-cyclic groups [44, 45], and so on. In particular, the fundamental groups of (compact) *special cube complexes* introduced by Haglund and Wise [47] form a very rich class of convex-cocompact subgroups of right-angled Artin groups, and they played a key role in the resolution of the Virtual Haken and Virtual Fibering Conjectures [1, 86].

In general, when nonempty, the space of geometric actions of a given group on CAT(0) cube complexes is quite large. For example, each filling multicurve on a closed hyperbolic surface is dual to a cubulation of its fundamental group [75]. Similarly, cubulations for fundamental groups of cusped hyperbolic 3-manifolds can be obtained from their vast sets of (relatively) quasiconvex surface subgroups [3, 24, 52]. Moreover, cubulations can be used to define deformation spaces, such as the classical Culler-Vogtmann outer space [25] that encodes geometric actions of free groups on trees. This perspective has been extended to right-angled Artin groups, for which outer spaces have been constructed using cubulations with some particular special cube complexes as quotients [8, 22].

Under some reasonable irreducibility assumptions, actions on CAT(0) cube complexes are *marked length-spectrum rigid* [4, 5]. More precisely, let $\mathcal{X}$ be a cubulation of a group Γ and let $\mathbf{conj} = \mathbf{conj}(\Gamma)$ denote the set of conjugacy classes of Γ. The *(stable) translation length* of this action is the function $\ell_{\mathcal{X}} : \mathbf{conj} \to \mathbb{R}$ given by

$$\ell_{\mathcal{X}}[g] = \lim_{n\to\infty} \frac{d_{\mathcal{X}}(g^n x, x)}{n},$$

where $d_{\mathcal{X}}$ denotes the combinatorial metric on the 1-skeleton of $\mathcal{X}$ and the limit above is independent of the representative $g \in [g]$ and the vertex $x \in \mathcal{X}$.

For two cubulations $\mathcal{X}, \mathcal{X}_*$ of Γ, marked length-spectrum rigidity states that the equality of translation length functions $\ell_{\mathcal{X}} = \ell_{\mathcal{X}_*}$ implies the existence of a Γ-equivariant cubical isometry from $\mathcal{X}$ onto $\mathcal{X}_*$. Since in general the translation length functions $\ell_{\mathcal{X}}$ and $\ell_{\mathcal{X}_*}$ will not coincide, it is natural to ask about the behavior of $\ell_{\mathcal{X}_*}[g]$ when $\ell_{\mathcal{X}}[g]$ is large. The goal of this paper is to address this question for "compatible" pairs of virtually co-special cubulations, that is, those having quotients with a special cube complex as a finite cover. Such compatibility is described in Definition 5.1, and is guaranteed for any group in the following class.

Definition 1.1. Let $\mathfrak{G}$ be the class of nonvirtually cyclic groups Γ satisfying the following:

(1) Γ admits a proper, cocompact and virtually co-special action on a CAT(0) cube complex.
(2) The class of convex-cocompact subgroups of Γ is the same with respect to any proper and cocompact action on a CAT(0) cube complex. That is, given any two proper, cocompact actions of Γ on CAT(0) cube complexes $\mathcal{X}, \mathcal{X}_*$ then the restricted action of a subgroup $H < \Gamma$ on $\mathcal{X}$ is convex-cocompact if and only if the restricted action of H on $\mathcal{X}_*$ is convex-cocompact.
(3) Some (equivalently, any) proper and cocompact action of Γ on a CAT(0) cube complex has a contracting element.

By [39, Lemma 4.6], contracting elements for proper and cocompact actions on CAT(0) cube complexes are those having invariant geodesics that satisfy the conclusion of the Morse lemma. In particular, the notion of being a contracting element is independent of the cubulation of Γ.

By Agol's theorem [1, Theorem 1.1] and the characterization of convex-cocompact subgroups in terms of quasiconvexity [47, Proposition 7.2], we see that every cubulable nonelementary hyperbolic group belongs to $\mathfrak{G}$. Moreover, the class $\mathfrak{G}$ is closed under relative hyperbolicity, in the sense that a cubulable relatively hyperbolic group belongs to $\mathfrak{G}$ as long as its peripheral subgroups belong to $\mathfrak{G}$. In particular, any $C'(1/6)$ small cancellation quotient of the free product of finitely many groups in $\mathfrak{G}$ belongs to $\mathfrak{G}$ [60]. However, this class is much larger since it contains some infinite families of right-angled Artin and Coxeter groups, most of them not being relatively hyperbolic with respect to any collection of proper subgroups. For instance, any right-angled Artin group with finite outer automorphism group belongs to $\mathfrak{G}$. See Proposition 5.2 for the precise statement.

1.1. Manhattan curves

Let $\mathcal{X}, \mathcal{X}_*$ be two cubulations of the group Γ. We endow these cubulations with Γ-invariant orthotope structures $\mathfrak{w}, \mathfrak{w}_*$ respectively, consisting of (non-necessarily integer) positive lengths assigned to the hyperplanes which are invariant under the action of Γ. This induces isometric actions of Γ on the cuboid complexes $\mathcal{X}^{\mathfrak{w}} = (\mathcal{X}, \mathfrak{w})$ and $\mathcal{X}_*^{\mathfrak{w}_*} = (\mathcal{X}_*, \mathfrak{w}_*)$, see Subsection 2.2 for further details. The *Manhattan curve* for the pair $(\mathcal{X}^{\mathfrak{w}}, \mathcal{X}_*^{\mathfrak{w}_*})$ is the boundary of the convex set

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

$$\mathcal{C}_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}} := \left\{(a,b)\in\mathbb{R}^2 : \sum_{[g]\in\mathbf{conj}(\Gamma)} e^{-a\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]-b\ell_{\mathcal{X}}^{\mathfrak{w}}[g]} < \infty\right\},$$

where $\ell_{\mathcal{X}}^{\mathfrak{w}}$ and $\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}$ are the respective translation length functions of the actions of Γ on the 1-skeleta of $\mathcal{X}^{\mathfrak{w}}$ and $\mathcal{X}_*^{\mathfrak{w}_*}$. We can parameterize this curve as $s\mapsto\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(s)$, where for $s\in\mathbb{R}$, $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(s)$ is the abscissa of convergence of the series

$$t\mapsto\sum_{[g]\in\mathbf{conj}(\Gamma)} e^{-t\ell_{\mathcal{X}}^{\mathfrak{w}}[g]-s\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}.$$

By abuse of notation, we also call the parametrization $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}$ the Manhattan curve of $(\mathcal{X}^{\mathfrak{w}},\mathcal{X}_*^{\mathfrak{w}_*})$.

Manhattan curves are useful tools for studying pairs of actions and are related to rigidity results, as well as recovering asymptotic invariants. They were introduced by Burger [12] for pairs of convex-cocompact representations of a group on rank 1 symmetric spaces, and later Sharp [81] proved that they are analytic for pairs of cocompact Fuchsian representations. Sharp also extended these results for pairs of points in the outer space of a free group [79, 80]. Recently, Manhattan curves have been studied for pairs of cusped Fuchsian representations [54, 55], pairs of cusped quasi-Fuchsian representations [6], comparing quasi-Fuchsian representations with negatively curved metrics on surfaces [53], pairs of cusped Hitchin representations [7], and pairs of geometric actions on hyperbolic groups [18, 19].

In some sense, the Manhattan curve can be seen as a function that ‘interpolates’ between its pair of defining isometric actions. In particular, the regularity (i.e., differentiability) properties of the Manhattan curve somehow measure the “compatibility” of such actions. Moreover, when Manhattan curves are known to be analytic, then (as they are convex) they are either straight lines or strictly convex everywhere. This convexity characterization leads to length spectrum rigidity and other rigidity results for pairs of actions. See [19, Theorem 1] for some examples of such results. In addition, when Manhattan curves are known to be analytic, we immediately obtain precise large deviations principles comparing isometric actions. We consider such results in this work. Lastly, the C^2-regularity of Manhattan curves can be use to construct *pressure metrics*, which recover the Weil-Petersson metric on Teichmüller space [61] and generalize it to other geometric settings [2, 6, 9, 55, 67].

Our first main theorem fits into the aforementioned results, and to the authors’ knowledge, is the first to address the analyticity of Manhattan curves outside the scope of relatively hyperbolic groups or representation theory.

Theorem 1.2. *Let Γ be a group in the class $\mathfrak{G}$ and let it act properly and cocompactly on the cuboid complexes $\mathcal{X}^{\mathfrak{w}}=(\mathcal{X},\mathfrak{w})$ and $\mathcal{X}_*^{\mathfrak{w}_*}=(\mathcal{X}_*,\mathfrak{w}_*)$. Then the Manhattan curve $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}:\mathbb{R}\to\mathbb{R}$ is convex, decreasing, and analytic. In addition, the following limit exists and equals $-\theta'_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(0)$:*

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) := \lim_{T\to\infty}\frac{1}{\#\{[g]\in\mathbf{conj}\colon \ell_{\mathcal{X}}^{\mathfrak{w}}[g]<T\}}\sum_{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]<T}\frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{T}.$$

Moreover, we always have

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}})\geq v_{\mathcal{X}^{\mathfrak{w}}}/v_{\mathcal{X}_*^{\mathfrak{w}_*}},$$

for $v_{\mathcal{X}^{\mathfrak{w}}}, v_{\mathcal{X}_*^{\mathfrak{w}_*}}$ the corresponding exponential growth rates, and the following are equivalent:

(1) *$\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}$ is a straight line;*
(2) *there exists $\Lambda>0$ such that $\ell_{\mathcal{X}}^{\mathfrak{w}}[g]=\Lambda\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]$ for all $[g]\in\mathbf{conj}(\Gamma)$; and*
(3) *$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}})=v_{\mathcal{X}^{\mathfrak{w}}}/v_{\mathcal{X}_*^{\mathfrak{w}_*}}$.*

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Remark 1.3. In the result above, the group Γ does not have to belong to $\mathfrak{G}$ as long as the triplet $(\Gamma, \mathcal{X}, \mathcal{X}_*)$ belongs to the class $\mathfrak{X}$ in Definition 5.1. In particular, the action on $\mathcal{X}_*$ does not have to be proper. See Theorem 6.1 for the more general statement.

1.2. An automaton for pairs of cubulations

The main tool in the proof of Theorems 1.2 and 6.1 is the construction of a *finite-state automaton* that simultaneously encodes translation lengths for the actions on both $\mathcal{X}$ and $\mathcal{X}_*$. Roughly speaking, an automaton is a finite directed graph $\mathcal{G}$ that encodes a group Γ equipped with a finite generating set S. The edges of $\mathcal{G}$ are labeled by elements of S, so that finite paths in $\mathcal{G}$ correspond to group elements whose word length (with respect to S) equals the length of the corresponding path. Well-known examples include the Bowen-Series coding for Fuchsian groups [13], Cannon's automatic structure for hyperbolic groups with arbitrary generating set [15] (see Example 2.4) and Hermiller-Meier's automatic structures for right-angled Artin and Coxeter groups with the standard generating sets [49]. Recently, the combinatorial structure of automata has been key to deduce strong counting results for some groups acting isometrically on δ-hyperbolic spaces. Examples include genericity of loxodromic elements [37, 38] and central limit theorems [36].

In principle, an automatic structure encodes a single length function associated to a group. However, for a hyperbolic group Γ with two input generating sets S, S_*, it is possible to enhance the automaton associated to S and equipping it with an extra (non-negative) integer-valued edge labeling. With respect to the new label, paths in the refined automaton record the word length with respect to S_* of the group element associated to the path. This was achieved by Calegari and Fujiwara in [14], and having access to an automatic structure and labeling such as this means that one can apply powerful tools and techniques from thermodynamic formalism and symbolic dynamics to study pairs word metrics on hyperbolic groups. For example, this construction was used by Cantrell and Tanaka [18] to show that Manhattan curves for pairs of word metrics on hyperbolic groups are analytic. Calegari-Fujiwara's construction was the main inspiration for the construction of the automaton for pairs of actions on CAT(0) cube complexes, which we now proceed to describe.

In our setting, we start with a group Γ in the class $\mathfrak{G}$ acting properly and cocompactly on the CAT(0) cube complexes $\mathcal{X}$ and $\mathcal{X}_*$. The main result is Theorem 5.11, which is stated using the formalism of automatic structures (see Subsection 2.3). As the statement of this theorem is rather technical, we provide a simplified version that also incorporates Lemma 6.7, Lemma 6.11, and Lemma 6.12.

Theorem 1.4. *Let Γ be a group in the class $\mathfrak{G}$ and let it act properly and cocompactly on the* CAT(0) *cube complexes $\mathcal{X}$ and $\mathcal{X}_*$. To the triplet $(\Gamma, \mathcal{X}, \mathcal{X}_*)$ we associate the following data:*

i) *A finite index subgroup $\overline{\Gamma} < \Gamma$ such that the quotient $\overline{\mathcal{X}} = \overline{\Gamma} \backslash \mathcal{X}$ is a special cube complex.*
ii) *A finite directed graph $\mathcal{G} = \mathcal{G}(\Gamma, \mathcal{X}, \mathcal{X}_*)$ equipped with a labeling map π that assigns to each edge of $\mathcal{G}$ an oriented hyperplane of $\overline{\mathcal{X}}$.*
iii) *An integer-valued functional ψ on the edges of $\mathcal{G}$.*

From this data, any closed loop ω in $\mathcal{G}$ is assigned to a closed (combinatorial) geodesic $\overline{\gamma}_\omega$ in $\overline{\mathcal{X}}$, and hence to a conjugacy class $[g_\omega] \in$ **conj**$(\overline{\Gamma})$*, in such a way that:*

(1) *If ω is determined by the sequence of edges $e_1, \ldots, e_n$ in $\mathcal{G}$, then the loop $\overline{\gamma}_\omega$ consists of edges that are dual to the hyperplanes $\pi(e_1), \ldots, \pi(e_n)$. In particular, the length of ω equals $\ell_{\mathcal{X}}[g_\omega]$.*
(2) *If ω is as in (1), then $\psi(e_1) + \cdots + \psi(e_n) = \ell_{\mathcal{X}_*}[g_\omega]$.*
(3) *The assignment $\omega \mapsto [g_\omega]$ from the set of closed loops of $\mathcal{G}$ into* **conj**$(\overline{\Gamma})$ *is*
   - *polynomial (in length)-to-one; and*
   - *has image with positive lower density with respect to the action of $\overline{\Gamma}$ on $\mathcal{X}$.*

We briefly sketch how Theorem 1.4 implies Theorem 1.2 in the case that $\mathcal{X}^{\mathfrak{w}} = \mathcal{X}$ and $\mathcal{X}_*^{\mathfrak{w}_*} = \mathcal{X}_*$. From the adjacency matrix of $\mathcal{G}$ we define a subshift of finite-type (Σ, σ), whose periodic orbits correspond to loops in $\mathcal{G}$ and hence induce conjugacy classes in **conj**$(\overline{\Gamma})$. The periods of these periodic

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

orbits correspond to $\ell_{\mathcal{X}}$-translation lengths. The functional ψ induces a potential $\Phi : \Sigma \to \mathbb{Z}$ that is constant on 2-cylinders, and whose Birkhoff sums correspond to $\ell_{\mathcal{X}_*}$-translation lengths. Item (3) then allows us to describe the Manhattan curve $\theta_{\mathcal{X}_*/\mathcal{X}}$ in terms of pressure functions associated with Φ (see Proposition 6.8), and the theorems then follow by standard results in thermodynamic formalism.

1.3. Large deviations

For a pair $\mathcal{X}, \mathcal{X}_*$ of cubulations of a group $\Gamma \in \mathfrak{G}$, we also study large deviations for their translation lengths.

That is, we estimate the number of conjugacy classes $[g]$ for which

$$\left|\frac{\ell_{\mathcal{X}_*}[g]}{\ell_{\mathcal{X}}[g]} - \eta\right| < \epsilon \text{ for some } \eta \in \mathbb{R} \text{ and a small } \epsilon > 0. \tag{1.1}$$

We also study the set of conjugacy classes $[g]$ such that

$$|\ell_{\mathcal{X}_*}[g] - \eta\ell_{\mathcal{X}}[g]| < \epsilon \text{ for some } \eta \in \mathbb{R} \text{ and a small } \epsilon > 0. \tag{1.2}$$

This latter comparison is more delicate than the corresponding quotient comparison in (1.1) above, as (when $\ell_{\mathcal{X}}[g]$ is bounded away from 0) (1.2) implies (1.1) but not vice versa.

Both of these equations, (1.1) and (1.2), represent the conjugacy classes $[g]$ for which $\ell_{\mathcal{X}_*}[g]$ is approximately $\eta\ell_{\mathcal{X}}[g]$. Therefore, understanding the growth rate/number of conjugacy classes satisfying these inequalities allows us to form a natural comparison between the actions of Γ on $\mathcal{X}$ and $\mathcal{X}_*$. These growth rates are also related to rigidity phenomena. See, for example, [17, Theorem 1.4] that classifies word metrics on hyperbolic groups in terms of growth rates coming from such large deviations estimates. In fact, this current work was motivated by [17, Theorem 4.1] and [17, Remark 4.3] as we now explain. The result [17, Theorem 4.1] shows that there is a precise large deviations theorem that compares certain 'compatible' word metrics on hyperbolic groups. Here 'compatible' is a condition regarding the exponential growth rates of the word metrics. In [17, Remark 4.3] the authors asked if this compatibility condition was necessary. In this work we show that it is not and that we can deduce precise large deviations results for all pairs of word metrics on hyperbolic groups. We present and discuss these results below in Subsection 1.4.

Despite the fact that estimating the number of conjugacy classes satisfying (1.2) is significantly harder than studying the analogous question for (1.1), there are previous works that tackle this problem in other settings. For example, let Σ be a closed surface with negative Euler characteristic and fundamental group Γ, and suppose that $\mathfrak{g}$ and $\mathfrak{g}_*$ are two hyperbolic metrics on Σ. These metrics induce isometric actions of Γ on $\widetilde{\Sigma}$ with translation length functions $\ell_{\mathfrak{g}}$ and $\ell_{\mathfrak{g}_*}$. A result of Schwartz and Sharp [78] states that there is an interval $(\alpha, \beta) \subset \mathbb{R}$ and constants $C, \lambda > 0$ such that any $\eta \in (\alpha, \beta)$ satisfies

$$\#\big\{[g] \in \mathbf{conj}(\Gamma) : \ell_{\mathfrak{g}}[g] < T : |\ell_{\mathfrak{g}_*}[g] - \eta\ell_{\mathfrak{g}}[g]| < \epsilon\big\} \sim \frac{Ce^{\lambda T}}{T^{3/2}}$$

as $T \to \infty$ for any fixed $\epsilon > 0$ (here '$\sim$' represents that the quotients of the two quantities converge to 1 as $T \to \infty$). Similar results are known to hold for surfaces of variable negative curvature by Dal'bo [28], for Hitchin representations by Dai and Martone [27], for Green metrics by Cantrell [16] and for some pairs of points in outer space by Sharp [79]. These asymptotics are often referred to as correlation results, and to prove them thermodynamic formalism is usually employed. To apply thermodynamic formalism one needs to know that the length spectra of the two considered metrics are not *rationally related*. That is, if ℓ_1, ℓ_2 are the length spectra that we want to compare then we would need to know that there do not exist nonzero $a, b \in \mathbb{R}$ with $a\ell_1[g] + b\ell_2[g] \in \mathbb{Z}$ for all $[g] \in \mathbf{conj}(\Gamma)$. This property is vital as it implies bounds on the operator norm of families of transfer operators, which are then used in the proof of the correlation asymptotic.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

On the other hand, the length spectra of a pair of cubical actions on CAT(0) cube complexes are always rationally related. Indeed, after possibly performing one cubical barycentric subdivision, every cubical isometry of a CAT(0) cube complex either fixes a vertex or preserves a bi-infinite geodesic on which the isometry acts by nontrivial translations [46]. In particular, the translation length function associated to any of these actions has image belonging to $\frac{1}{2}\mathbb{Z}$.

However, by means of the automaton from Theorem 5.11 (Theorem 1.4), we are still able to estimate the number of conjugacy classes satisfying (1.2). As a consequence of Theorem 3.2 in the setting of subshifts of finite type, we can prove the following result that can be seen as large deviations with shrinking intervals.

Theorem 1.5. *Let* Γ *be a group in the class* $\mathfrak{G}$ *and let it act properly and cocompactly on the* CAT(0) *cube complexes* $\mathcal{X}$ *and* $\mathcal{X}_*$. *Let* $\mathbf{conj}' \subset \mathbf{conj}$ *be the set of nontorsion conjugacy classes and consider the dilations*

$$\mathrm{Dil}(\mathcal{X}_*, \mathcal{X}) = \sup_{[g]\in\mathbf{conj}'} \frac{\ell_{\mathcal{X}_*}[g]}{\ell_{\mathcal{X}}[g]} \text{ and } \mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1} = \inf_{[g]\in\mathbf{conj}'} \frac{\ell_{\mathcal{X}_*}[g]}{\ell_{\mathcal{X}}[g]}.$$

Then there exists an analytic function

$$\mathcal{I} : [\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X})] \to \mathbb{R}$$

and $C > 0$ *such that for any* $\eta \in (\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X}))$ *we have*

$$0 < \limsup_{T\to\infty} \frac{1}{T} \log\left(\#\left\{[g] \in \mathbf{conj} : \ell_{\mathcal{X}}[g] < T, |\ell_{\mathcal{X}_*}[g] - \eta\ell_{\mathcal{X}}[g]| < \frac{C}{T}\right\}\right) = \mathcal{I}(\eta) \le v_{\mathcal{X}}. \tag{1.3}$$

Furthermore, we have equality in the above inequality if and only if $\eta = \tau(\mathcal{X}_*/\mathcal{X})$.

We have referred to this result as large deviations with shrinking intervals because it estimates the growth of the number of group elements satisfying (1.2) opposed to (1.1). It is desirable to prove this refinement, as it provides a much more precise comparison between the actions of Γ on $\mathcal{X}$ and $\mathcal{X}_*$.

Intuitively, the function $\mathcal{I}$ can be seen as measuring how similar the geometries actions on $\mathcal{X}$ and $\mathcal{X}_*$ are. That is, the closer $\mathcal{I}$ is to the constant function with value $v_{\mathcal{X}}$, the more similar the length functions $v_{\mathcal{X}}\ell_{\mathcal{X}}$ and $v_{\mathcal{X}_*}\ell_{\mathcal{X}_*}$ are. We also note that the real analyticity of $\mathcal{I}$ is a useful property. Indeed, as shown in [17], when $\mathcal{I}$ is analytic, it is possible to obtain rigidity results that compare metrics through the values taken by $\mathcal{I}$.

Remark 1.6. As in Theorem 1.2, the conclusion above still holds for triplets $(\Gamma, \mathcal{X}, \mathcal{X}_*)$ in the class $\mathfrak{X}$; see Theorem 6.2. However, for our arguments (particularly Theorem 3.2) it is crucial that the translation length functions belong to a lattice in $\mathbb{R}$. We still expect Theorem 6.2 to hold for arbitrary cuboid complexes $\mathcal{X}^{\mathfrak{w}}$ and $\mathcal{X}_*^{\mathfrak{w}_*}$, but we will not pursue this in this work.

As an application of Theorem 1.5 we deduce large deviations with shrinking intervals for the intersection of curves on hyperbolic surfaces. Let Σ be a closed orientable surface of negative Euler characteristic and fundamental group Γ. If α, β are immersed closed oriented curves in Σ, then the *(geometric) intersection number* is the minimal number $i_\Sigma(\alpha, \beta)$ of intersections of closed curves in the free homotopy classes of α and β. The function i_Σ can be extended by bilinearity to *weighted multicurves*, which are finite sums of the form $\sum_j \lambda_j\alpha_j$ with α_j immersed oriented closed curves in Σ and a set $(\lambda_j \ge 0)_j$ of *weights*. Any nontrivial element in $\mathbf{conj}(\Gamma)$ is represented by a unique free homotopy class of immersed oriented closed curves, so we can talk of the intersection number between a weighted multicurve in Σ and a conjugacy class in $\mathbf{conj}(\Gamma)$. For more details about the intersection number, see [33].

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

A generating set S for Γ is *simple* if there exists a point $p \in \Sigma$ such that elements of $S \subset \Gamma = \pi_1(\Sigma, p)$ can be represented by simple loops that are pairwise nonhomotopic and disjoint except at the base point p. For example, the generating set for the standard presentation

$$\Gamma = \left\langle a_1, b_1, \ldots, a_g, b_g \colon [a_1, b_1] \cdots [a_g, b_g] \right\rangle$$

is simple. In [31, Theorem 1.2], Erlandsson proved that the translation length function ℓ_S of the word metric of a simple generating set S can be recovered by pairing in the intersection number against a carefully chosen weighted multicurve α_S with weights in $\frac{1}{2}\mathbb{Z}$. Up to scaling, this translation length function can also be recovered by looking at the CAT(0) cube complex dual to the multicurve α_S. Therefore, Theorem 6.2 applies and we obtain the following.

Corollary 1.7. *Let Γ be the fundamental group of the closed orientable hyperbolic surface Σ and consider a simple generating set S of exponential growth rate v_S. Let α be a nontrivial weighted multicurve on Σ with integer weights, and define*

$$a_{\inf} := \inf_{[g] \in \mathbf{conj}'} \frac{i_\Sigma(\alpha, [g])}{\ell_S[g]} \quad \textit{and} \quad a_{\sup} := \sup_{[g] \in \mathbf{conj}'} \frac{i_\Sigma(\alpha, [g])}{\ell_S[g]}.$$

Then there exists an analytic convex function $\mathcal{I} : [a_{\inf}, a_{\sup}] \to \mathbb{R}$ and $C > 0$ such that for any $\eta \in (a_{\inf}, a_{\sup})$ we have

$$0 < \limsup_{T \to \infty} \frac{1}{T} \log\left(\#\left\{[g] \in \mathbf{conj} : \ell_S[g] < T, |i_\Sigma(\alpha, [g]) - \eta \ell_S[g]| < \frac{C}{T}\right\}\right) = \mathcal{I}(\eta) \leq v_S.$$

As discussed above, we can see this corollary as providing a precise comparison between the intersection number (with α) and word length of closed geodesics. In some sense, $\mathcal{I}$ is encoding how well the intersection number can be approximated by the simple word metric for S.

1.4. Word metrics on hyperbolic groups

Since the proof of Theorem 1.5 relies on Theorem 3.2 (a purely dynamical statement), the existence of an automaton encoding both actions on $\mathcal{X}$ and $\mathcal{X}_*$, and the arithmeticity of the translation length functions, we can deduce large deviations with shrinking intervals for any pair of group actions fulfilling similar conditions. That is the case of word metrics on hyperbolic groups, and in fact, this was the author's main motivation at the beginning of this project.

Let Γ be a nonelementary hyperbolic group and let $S, S_* \subset \Gamma$ be finite generating sets with corresponding word metrics d_S, d_{S_*}. By Cannon's theorem [15], for a total order on S, the language of lexicographically first geodesics in Γ is regular, so it is parametrized by a finite-state automaton. As a consequence of [14, Lemma 3.8], Calegari and Fujiwara are able to modify this automaton (without modifying the parameterized language), and find an integer functional on the edges of the graph of the automaton so that its sum over paths recovers the S_*-word length for the corresponding element in Γ (this was our main motivation to construct the automaton in Theorem 5.11).

By studying the subshift of finite type associated to this automaton, Cantrell and Tanaka [19, 18] deduced analyticity of the Manhattan curve for S, S_* as well as a large deviations principle. More precisely, let ℓ_S, ℓ_{S_*} be the corresponding translation length functions and consider the dilations

$$\mathrm{Dil}(S_*, S) = \sup_{[g] \in \mathbf{conj}'} \frac{\ell_{S_*}[g]}{\ell_S[g]} \quad \text{and} \quad \mathrm{Dil}(S, S_*)^{-1} = \inf_{[g] \in \mathbf{conj}'} \frac{\ell_{S_*}[g]}{\ell_S[g]}.$$

Then there exists a real analytic, concave function $\mathcal{I} : [\mathrm{Dil}(S, S_*)^{-1}, \mathrm{Dil}(S_*, S)] \to \mathbb{R}_{>0}$ such that for $\eta \in (\mathrm{Dil}(S, S_*)^{-1}, \mathrm{Dil}(S_*, S))$ we have

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

$$\lim_{\epsilon\to 0^+}\limsup_{n\to\infty}\frac{1}{T}\log\left(\#\left\{[g]\in \mathbf{conj}(\Gamma)\colon \ell_S[g]<T, \left|\frac{\ell_{S_*}[g]}{\ell_S[g]}-\eta\right|<\epsilon\right\}\right)=\mathcal{I}(\eta). \tag{1.4}$$

The rate function $\mathcal{I}$ is a Legendre transform constructed from the Manhattan curve for S, S_*, see [19, Theorem 4.23]. By applying Theorem 3.2 to this subshift, we can improve this result and obtain a large deviations theorem with shrinking intervals.

Theorem 1.8. *Let Γ be a nonelementary hyperbolic group and consider two finite generating sets S, S_* for Γ with exponential growth rates v_S, v_{S_*}, and $\mathcal{I}$ as above. Then there exists $C>0$ such that for any $\eta\in(\mathrm{Dil}(S,S_*)^{-1},\mathrm{Dil}(S_*,S))$ we have*

$$0<\limsup_{T\to\infty}\frac{1}{T}\log\left(\#\left\{[g]\in\mathbf{conj}: \ell_S[g]<T, |\ell_{S_*}[g]-\eta\ell_S[g]|<\frac{C}{T}\right\}\right)=\mathcal{I}(\eta)\le v_S.$$

Furthermore, we have equality in the above inequality if and only if

$$\eta=\tau(S_*/S):=\lim_{T\to\infty}\frac{1}{\#\{[g]\in\mathbf{conj}\colon \ell_S[g]<T\}}\sum_{\ell_S[g]<T}\frac{\ell_{S_*}[g]}{T}.$$

This result implies that, after scaling a pair of word metrics by their exponential growth rates, there is always an exponentially growing set for which their translation lengths are close, that is, for any $\epsilon>0$

$$0<\limsup_{T\to\infty}\frac{1}{T}\log\left(\#\{[g]\in\mathbf{conj}: \ell_S[g]<T, |v_S\ell_S[g]-v_{S_*}\ell_{S_*}[g]|<\epsilon\}\right)\le v_S. \tag{1.5}$$

This extends the recent work [17, Theorem 4.1] where the authors proved a correlation result for pairs of word metrics under an additional rationality assumption on the exponential growth rates. This result also answers a question raised by the authors in [17, Remark 4.3]. We also deduce the following corollary.

Corollary 1.9. *Let Γ be a nonelementary hyperbolic group, and let S and S_* be two finite generating sets on Γ. Then there exists $C>0$ such that for any $\eta\in[\mathrm{Dil}(S,S_*)^{-1},\mathrm{Dil}(S_*,S)]$ we can find an infinite sequence $(g_n)_{n\ge1}\subset\Gamma$ such that*

$$\left|\frac{\ell_{S_*}[g_n]}{\ell_S[g_n]}-\eta\right|\le\frac{C}{|g_n|_S^2}. \tag{1.6}$$

If $\eta\in[\mathrm{Dil}(S,S_*)^{-1},\mathrm{Dil}(S_*,S)]$ is rational then there exists $g\in\Gamma$ such that

$$\frac{\ell_{S_*}[g]}{\ell_S[g]}=\eta.$$

This result is concerned with understanding how well the values of the quotient of ℓ_{S_*} with ℓ_S can approximate a given $\eta\in[\mathrm{Dil}(S,S_*)^{-1},\mathrm{Dil}(S_*,S)]$. It shows that the approximation rate is optimal. Indeed, it is well-known that the translation length function for word metrics takes values in a lattice $\frac{1}{N}\mathbb{Z}$ for some $N\in\mathbb{N}$ and therefore by Hurwitz's Theorem [51] we cannot find a sequence g_n for which the convergence rate in (1.6) is faster.

In Subsection 4.2 we present an example for a pair of word metrics on a free group. In particular, we compute the limit supremum in (1.5) for this pair of word metrics.

Organization

The organization of the paper is as follows. In Section 2 we cover preliminary material about Manhattan curves, CAT(0) cube complexes and cubulable groups, finite-state automata, symbolic dynamics and suspension flows.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

In Section 3 we prove Theorem 3.2, a large deviations result with shrinking intervals for lattice potentials on mixing subshifts of finite type that are constant on 2-cylinders. We apply this theorem to pairs of word metrics on hyperbolic groups in Section 4 and prove Theorem 1.8.

In Section 5 we prove Proposition 5.2 that describes large classes of groups included in $\mathfrak{G}$. There we also prove Theorem 5.11, in which we construct a finite-state automaton for pairs of compatible actions on CAT(0) cube complexes. We use this automaton in Section 6 to prove Theorem 6.1 and Theorem 6.2, from which we deduce Theorem 1.2 and Theorem 1.5. In this section we also prove Theorem 1.4.

Finally, in the appendix we prove a Proposition A.1, a criterion for convex-cocompactness of subgroups of cubulable relatively hyperbolic groups, which may be of independent interest.

2. Preliminaries

2.1. Isometric group actions

Let Γ be a finitely generated group acting by isometries on the metric space (X, d_X) and let $x \in X$ be an arbitrary base point. The *(stable) translation length* of this action is the function $\ell_X : \mathbf{conj} \to \mathbb{R}$ given by

$$\ell_X[g] = \lim_{n\to\infty} \frac{d_X(g^n x, x)}{n} \quad \text{for } g \in [g] \text{ in } \mathbf{conj}.$$

The *exponential growth rate* of this action is the quantity

$$v_X := \limsup_{T\to\infty} \frac{\log(\#\{g \in \Gamma\colon d_X(gx, x) < T\})}{T} \in [0, +\infty].$$

As for the translation lengths, the exponential growth rate is independent of the chosen base point x. If X is geodesic (or more generally roughly geodesic) and the action of Γ on X is proper and cocompact, then v_X is finite. In some cases we can recover the exponential growth rate as the limit

$$\limsup_{T\to\infty} \frac{1}{T} \log \#\mathfrak{C}_X(T),$$

where each $T > 0$ we denote

$$\mathfrak{C}_X(T) = \{[g] \in \mathbf{conj}\colon \ell_X[g] < T\}.$$

This happens for example when Γ is hyperbolic and the action on X is proper and cocompact.

Given two isometric actions of Γ on the metric spaces X and X_*, the *Manhattan curve* for the pair (X, X_*) is the boundary of the convex set

$$\mathcal{C}_{X_*/X} := \left\{(a, b) \in \mathbb{R}^2 : \sum_{[g]\in\mathbf{conj}} e^{-a\ell_{X_*}[g] - b\ell_X[g]} < \infty\right\},$$

assuming it is nonempty. Equivalently, $\mathcal{C}_{X_*/X}$ is the set of points $(s, \theta_{X_*/X}(s))$ where $\theta_{X_*/X}(s)$ is the abscissa of convergence of the series

$$t \mapsto \sum_{[g]\in\mathbf{conj}} e^{-t\ell_X[g] - s\ell_{X_*}[g]}.$$

By abuse of notation, $\theta_{X_*/X}$ is also called the Manhattan curve for (X, X_*).

2.2. CAT(0) *cube complexes*

For bibliography about CAT(0) cube complexes and groups acting on them, we refer the reader to [10, 74]. A *nonpositively curved (NPC) cube complex* is a metric polyhedral complex in which all

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

polyhedra are unit-length Euclidean cubes, and satisfies Gromov's link condition: the link of each vertex is a flag complex. If this complex is simply connected we say that it is a CAT(0) *cube complex*.

Let $\mathcal{X}$ be an NPC cube complex. Consider the minimal equivalence relations on the set of edges (resp. oriented edges) of $\mathcal{X}$ such that the edges $e = \{v, w\}$ and $e = \{v', w'\}$ (resp. oriented edges $v \xrightarrow{e} w$ and $v' \xrightarrow{e'} w'$) are in the same equivalence class if v, w, v', w' span a square in $\mathcal{X}$ (resp. $v, w, v'w'$ span a square in $\mathcal{X}$ with v, v' adjacent and w, w' adjacent). Equivalence classes of these equivalence relations are called *hyperplanes* (resp. *oriented hyperplanes*), and we let $\mathbb{H}(\mathcal{X})$ denote the set of hyperplanes of $\mathcal{X}$. If the equivalence class of an (oriented) edge e is the (oriented) hyperplane $\mathfrak{h}$, then we say that e *dual* to $\mathfrak{h}$. A hyperplane is *2-sided* if it corresponds to exactly two oriented hyperplanes; otherwise it is *1-sided*.

Suppose now that $\mathcal{X}$ is a CAT(0) cube complex, in which case all hyperplanes are 2-sided. A *combinatorial path* in $\mathcal{X}$ is a sequence $\gamma = (\gamma_0, \dots, \gamma_n)$ of vertices in $\mathcal{X}$ such that γ_i is adjacent to γ_{i+1} for $i = 0, \dots, n-1$. In that case, we say that γ_0 is the initial vertex of γ and γ_n is its final vertex. The length of $\gamma = (\gamma_0, \dots, \gamma_n)$ is defined as n. This path is often seen as a continuous path by also considering the edges e_{i+1} joining each γ_i with γ_{i+1}. Such a path is *geodesic* if no two distinct edges e_i are dual to the same hyperplane. The combinatorial metric on $\mathcal{X}$ is the graph metric $d_{\mathcal{X}}$ on its 1-skeleton $\mathcal{X}^1$ so that each edge has length 1. It follows that a combinatorial path is geodesic if and only if it is geodesic for the metric $d_{\mathcal{X}}$.

A hyperplane $\mathfrak{h}$ in the CAT(0) cube complex $\mathcal{X}$ *separates* two vertices of $\mathcal{X}$ if some (any) combinatorial path connecting these vertices has an edge dual to $\mathfrak{h}$. It follows that the combinatorial distance of any two vertices in $\mathcal{X}$ equals the number of hyperplanes separating them. Also, a hyperplane $\mathfrak{h}$ determines the equivalence relation of "not being separated by $\mathfrak{h}$" on the set of vertices of $\mathcal{X}$. This equivalence relation has exactly 2 equivalence classes $\{\mathfrak{h}^-, \mathfrak{h}^+\}$, which are the *halfspaces* determined by $\mathfrak{h}$. A subcomplex of $\mathcal{X}$ is *convex* if its vertex set is the intersection of halfspaces. Equivalently, $Z \subset \mathcal{X}$ is convex if any combinatorial geodesic joining points in Z^0 is contained in Z^0.

An *orthotope structure* on the CAT(0) cube complex $\mathcal{X}$ is a function

$$\mathfrak{w} : \mathbb{H}(\mathcal{X}) \to \mathbb{R}_{>0},$$

and the pair $\mathcal{X}^{\mathfrak{w}} = (\mathcal{X}, \mathfrak{w})$ is called a *cuboid complex*. An orthotope structure induces a metric $d^{\mathfrak{w}}_{\mathcal{X}}$ on $\mathcal{X}^1$ by declaring each edge e to have length $\mathfrak{w}(\mathfrak{h})$ for $\mathfrak{h}$ the hyperplane dual to e. In this way, for any two vertices $x, y \in \mathcal{X}^0$ we have

$$d^{\mathfrak{w}}_{\mathcal{X}}(x, y) = \sum_{\mathfrak{h} \in \mathbb{H}(x|y)} \mathfrak{w}(\mathfrak{h}),$$

where $\mathbb{H}(x|y) \subset \mathbb{H}(\mathcal{X})$ is the collection of hyperplanes separating x and y. Note that if $\mathfrak{w}$ is the constant function equal to 1, then $d^{\mathfrak{w}}_{\mathcal{X}}$ is just the standard combinatorial metric $d_{\mathcal{X}}$.

Remark 2.1. It is clear that a geodesic in $\mathcal{X}$ with respect to $d_{\mathcal{X}}$ is also geodesic with respect to $d^{\mathfrak{w}}_{\mathcal{X}}$ for any orthotope structure $\mathfrak{w}$.

Now let Γ be a group acting on the CAT(0) cube complex $\mathcal{X}$. We always assume that the action is cubical, meaning that it preserves the cube complex structure. Under this assumption, the action is isometric on $\mathcal{X}^1$ with the combinatorial metric $d_{\mathcal{X}}$. Similarly, $\ell_{\mathcal{X}}$ always denotes the stable translation length of Γ with respect to the action on $(\mathcal{X}^1, d_{\mathcal{X}})$. If the action of Γ is proper and cocompact, we say that $\mathcal{X}$ is a *cubulation* of Γ.

The action of Γ on $\mathcal{X}$ induces a natural action on $\mathbb{H}(\mathcal{X})$. If $\mathfrak{w}$ is a Γ-invariant orthotope structure on $\mathcal{X}$ in the sense that $\mathfrak{w}(\mathfrak{h}) = \mathfrak{w}(g\mathfrak{h})$ for all $g \in \Gamma$ and $\mathfrak{h} \in \mathbb{H}(\mathcal{X})$, then we say that Γ *acts* on the cuboid complex $(\mathcal{X}, \mathfrak{w})$. In that case the action of Γ on $(\mathcal{X}^1, d^{\mathfrak{w}}_{\mathcal{X}})$ is also by isometries. We let $\ell^{\mathfrak{w}}_{\mathcal{X}}$ denote the stable translation length of Γ for its action on $(\mathcal{X}^1, d^{\mathfrak{w}}_{\mathcal{X}})$.

By a *hyperplane stabilizer* we mean a subgroup of Γ consisting of group elements g such that $g\mathfrak{h} = \mathfrak{h}$ for some fixed hyperplane $\mathfrak{h} \in \mathbb{H}(\mathcal{X})$. We denote the hyperplane stabilizer of $\mathfrak{h}$ by $\Gamma_{\mathfrak{h}}$. A hyperplane $\mathfrak{h}$ is

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

essential for the action of Γ if for any vertex x of $\mathcal{X}$, the halfspaces $\mathfrak{h}^{\pm}$ contain elements in the Γ-orbit of x arbitrarily far from $\mathfrak{h}^{\mp}$. The action of Γ on $\mathcal{X}$ is *essential* if every hyperplane is essential.

If $\mathcal{X}$ is a cubulation of Γ, a subgroup $H < \Gamma$ is called *convex-cocompact* with respect to $\mathcal{X}$ if there exists a convex subcomplex $Z \subset \mathcal{X}$ that is H-invariant and so that the action of H on Z is cocompact. Such a subcomplex Z is called a *convex core* for H. Note that the hyperplane stabilizer of any hyperplane $\mathfrak{h}$ is convex-cocompact since it acts cocompactly on the (convex subcomplex spanned by the) set of vertices in edges dual to $\mathfrak{h}$ [47, Lemma 13.4].

If $\mathcal{X}$ is a CAT(0) cube complex and $\mathbb{W} \subset \mathbb{H}(\mathcal{X})$ is any collection of hyperplanes, in [20, Section 2.3] Caprace and Sageev introduced the *restriction quotient*, which is a CAT(0) cube complex $\mathcal{X}(\mathbb{W})$ equipped with a surjective cellular map $\phi : \mathcal{X} \longrightarrow \mathcal{X}(\mathbb{W})$ satisfying the following: an edge in $\mathcal{X}$ is collapsed to a single vertex under ϕ if and only if it is dual to a hyperplane not in $\mathbb{W}$. The projection ϕ induces a natural bijection between $\mathbb{W}$ and $\mathbb{H}(\mathcal{X}(\mathbb{W}))$, and hence any orthotope structure $\mathfrak{w}$ on $\mathcal{X}$ induces an orthotope structure $\phi_*(\mathfrak{w})$ on $\mathcal{X}(\mathbb{W})$.

Remark 2.2. Note that (pre)images of convex subcomplexes under restriction quotients remain convex. In particular, images of geodesic paths remain geodesic, although some subpaths are allowed to collapse to points. Also note that if Γ acts on $\mathcal{X}$ and $\mathbb{W}$ is Γ-invariant, then there is a natural action of Γ on $\mathcal{X}(\mathbb{W})$.

A cubulation $\mathcal{X}$ of Γ is *co-special* if the quotient $\overline{\mathcal{X}} = \Gamma\backslash\mathcal{X}$ is a *special cube complex* in the sense of Haglund-Wise [47]. Equivalently, $\mathcal{X}$ is co-special if Γ injects into a right-angled Artin group A_G inducing a Γ-equivariant isometric embedding of $\mathcal{X}$ into R_G as a convex subcomplex, where R_G is the universal cover of the *Salvetti complex* $\overline{R}_G$ associated to the graph G. Among other properties of special cube complexes, hyperplanes are 2-sided and embedded, and they do not self-osculate [47, Def. 3.2]. In particular, different oriented edges with the same initial vertex are dual to different oriented hyperplanes. The cubulation $\mathcal{X}$ of Γ is *virtually co-special* if there exists a finite-index subgroup $\overline{\Gamma} < \Gamma$ such that the action of $\overline{\Gamma}$ on $\mathcal{X}$ is co-special.

Fundamental groups of compact special cube complexes are residually finite, and more generally, their convex-cocompact subgroups are separable [47, Corollary 7.9]. For our purposes, co-special cubulations will be used to construct finite-state automata parameterizing combinatorial geodesics, as we explain in Example 2.5 in the next subsection.

2.3. *Finite-state automata*

For references on automatic structures and some of their connections with group theory, see [32]. For the relation between automatic structures and special cube complexes, we refer the reader to [58, Section 5]. Let S be a finite set and let S^* denote the set of finite words over the alphabet S. If $w = h_1 \cdots h_n$ and $w' = h_1' \cdots h_m'$ are words in S^*, then its *concatenation* is the word $ww' := h_1 \cdots h_n h_1' \cdots h_m'$. The *length* of a word is the number of letters in S composing it. We let the empty set correspond to the unique word of length 0 in S^*. A *language* over S is any subset of words in S^*.

A *(finite-state)* automaton over S is a tuple

$$\mathcal{A} = (\mathcal{G}, \pi, I, F),$$

where $\mathcal{G} = (V, E)$ is a finite directed graph, $\pi : E \longrightarrow S$ is a *labeling function* and $I, F \subset V$ are nonempty sets of *initial* and *final* states.

Remark 2.3. This convention differs from the standard definition of automaton (e.g., [58, Definition 5.2]), where it is required for I to consist of a single vertex. This difference in convention does not significantly change the discussion, but it will be useful in Theorem 5.11 when we construct an automaton for which we have no control on the number of initial states.

By a *path* in $\mathcal{G}$ we mean a sequence ω of (always directed) edges $e_1, \ldots, e_n$ in E such that the final vertex v_i of e_i is the initial vertex of e_{i+1} for $i = 1, \ldots, n-1$. If v_0 is the initial vertex of e_1, we denote this path ω either by $\omega = (\xrightarrow{e_1} \cdots \xrightarrow{e_n})$ or $\omega = (v_0 \xrightarrow{e_1} \cdots \xrightarrow{e_n} v_n)$ depending on the emphasis we

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

want to give to the vertices. If there is no ambiguity on the edges, we can also denote this path by $(v_0 \to \cdots \to v_n)$. The *length* of a path is the number of edges that determine it. Note that paths of length 1 correspond to the edges in E. We also allow the degenerate case of paths of length 0, which are the vertices in V.

If ω, ω' are paths in $\mathcal{G}$ such that the final vertex of ω is the initial vertex of ω', then the concatenation $\omega\omega'$ is the path in $\mathcal{G}$ defined in the expected way. Similarly we define the concatenation of any finite number of paths.

A word w in S^* is *represented* by the path $\omega = (v_0 \xrightarrow{e_1} v_1 \cdots \xrightarrow{e_n} v_n)$ in $\mathcal{G}$ if $w = \pi(\omega) := \pi(e_1)\cdots\pi(e_n)$. If in addition $v_0 \in I$ and $v_n \in F$, we say that w *accepted* by $\mathcal{A}$ and that ω is *admissible*. It is clear that the word represented by a concatenation $\omega\omega'$ is the concatenation of the words represented by ω and ω'. Let $L = L_{\mathcal{A}}$ be the language consisting of the words accepted by $\mathcal{A}$. In this case we say that L is *parametrized* by $\mathcal{A}$.

The automaton $\mathcal{A}$ is *deterministic* if any two distinct edges in $\mathcal{G}$ with the same initial vertex have different labels. In that case, for any $w \in L_{\mathcal{A}}$ and any initial state $v \in I$ there exists at most one path in $\mathcal{G}$ representing w and starting at v. The automaton is *pruned* if any vertex in $\mathcal{G}$ is the final vertex of a path starting at an initial state.

Example 2.4 (Automatic structures on hyperbolic groups). Let Γ be a hyperbolic group and consider a finite set $S \subset \Gamma$ generating Γ as a semi-group and with word length $|\cdot|_S$. We consider S as our alphabet, so that there is a natural evaluation map $\mathrm{ev} : S^* \to \Gamma$. For a fixed total order on S we obtain the lexicographic order $\prec$ on S^*. Let $L = L_S$ be language of *lexicographically first geodesics*. That is, for each $g \in \Gamma$, a word $w \in S^*$ with $\mathrm{ev}(w) = g$ is in L if and only if w has length $|g|_S$ and $w \prec w'$ for all other $w' \in S^*$ with $\mathrm{ev}(w') = g$ and of length $|g|_S$. Cannon [15] showed that the language L defined above is *regular*, in the sense that $L = L_{\mathcal{A}}$ for $\mathcal{A} = (\mathcal{G} = (V, E), \pi, \{*\}, V)$ a deterministic finite-state automaton over S. In particular, the evaluation map gives us a length-preserving bijection from L onto Γ.

Consider now another finite generating subset $S_* \subset \Gamma$. In [14, Lemma 3.8], Calegari and Fujiwara constructed a new deterministic automaton $\mathcal{A}' = (\mathcal{G}' = (V', E'), \pi', \{*'\}, V')$ over S parameterizing L_S and an integer-valued function $\phi : V' \to \mathbb{Z}$ such that for any word $w \in L_S$ represented by the path $\omega = (*' \xrightarrow{e_1} \cdots \xrightarrow{e_n} v_n)$ in $\mathcal{G}'$ we have

$$|\mathrm{ev}(w)|_{S_*} = \sum_{i=1}^{n} \phi(v_i).$$

We note that it is equivalent to define the labeling ϕ on the directed edge set E' instead.

We will see in the following subsection that the automatic structure discussed above gives rise to a dynamical system (Σ, σ) called a subshift of finite type. The labeling ϕ above corresponds to a real-valued function on Σ which satisfies the Markovian property of being 'constant on 2-cylinders': see Subsection 2.4.

Example 2.5 (Automatic structures on special cube complexes). For this example we follow Sections 5.2 and 5.3 in [58]. Let $\overline{\mathcal{Z}}$ be a compact special cube complex with universal cover $\mathcal{Z}$ and fix a base vertex $o \in \mathcal{Z}$. Let $S_{\overline{\mathcal{Z}}}$ be the set of the oriented hyperplanes of $\overline{\mathcal{Z}}$. In [58], Li and Wise constructed a deterministic pruned finite-state automaton

$$\mathcal{A}_{\overline{\mathcal{Z}}} = (\mathcal{G}_{\overline{\mathcal{Z}}} = (V, E), \pi, \{*\}, V)$$

over $S_{\overline{\mathcal{Z}}}$ and parameterizing a language $L_{\overline{\mathcal{Z}}}$. This language describes geodesics in $\mathcal{Z}$ based at o in the following sense. Given any word $w \in L_{\overline{\mathcal{Z}}}$ represented by the (necessarily unique) admissible path ω there exists a geodesic path γ_w in $\mathcal{Z}$ starting at o and ending at the vertex $\tau_{\mathcal{Z}}(w)$ such that:

- if $w = \mathfrak{h}_1 \cdots \mathfrak{h}_n \in L_{\overline{\mathcal{Z}}}$ and $\omega = (*' \xrightarrow{e_1} \cdots \xrightarrow{e_n} v_n)$, then $\gamma_w = (o = x_0, \ldots, x_n)$ is a geodesic of length n;

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

- for each $1 \leq i \leq n$ we have $\pi(e_i) = \mathfrak{h}_i$ and the oriented hyperplane in $\mathcal{Z}$ dual to the oriented edge from x_{i-1} to x_i maps to $\mathfrak{h}_i$ under the quotient $\mathcal{Z} \to \overline{\mathcal{Z}}$; and,
- the map $\tau_{\mathcal{Z}} : L_{\overline{\mathcal{Z}}} \to \mathcal{Z}^0$ is a bijection.

The automaton $\mathcal{A}_{\overline{\mathcal{Z}}}$ being deterministic implies that the geodesic γ_w is uniquely determined by w.

Remark 2.6. The language constructed in [58] actually depends on an injection of $\Gamma = \pi_1(\overline{\mathcal{Z}})$ into a right-angled Artin group A_G inducing a Γ-equivariant isometric embedding of $\mathcal{Z}$ into R_G as a convex subcomplex, where R_G is the universal cover of the Salvetti complex $\overline{R}_G$ associated to G. In that case, the language obtained is over the alphabet of oriented hyperplanes of $\overline{R}_G$. The language $L_{\overline{\mathcal{Z}}}$ described above is a particular case of this construction, when we consider the local isometric immersion $\overline{\mathcal{Z}} \to \overline{R}_{G_{\overline{\mathcal{Z}}}}$ for $G_{\overline{\mathcal{Z}}}$ being the *crossing graph* of $\overline{\mathcal{Z}}$. For this immersion there is a natural bijection between $S_{\overline{\mathcal{Z}}}$ and the set of oriented hyperplanes in $\overline{R}_{G_{\overline{\mathcal{Z}}}}$, see for instance [47, Lemma 4.1].

2.4. Symbolic dynamics

In this subsection we introduce the preliminary material we need from symbolic dynamics. See Chapter 1 of [66] for more information regarding the basic definitions we now present. Let A be a $k \times k$ matrix with entries 0 or 1. This matrix is said to be *aperiodic* if there exists $N \geq 1$ such that all of the entries of A^N are strictly positive. We say that A is *irreducible* if for any $i, j \in \{1, \dots, k\}$ there exists $n \geq 1$ such that $(A^n)_{i,j}$ (i.e., the (i, j)th entry of A^n) is strictly positive.

The (one-sided) *subshift of finite type* Σ_A associated to A is the set of infinite sequences

$$\Sigma_A = \left\{(x_n)_{n=0}^{\infty} : x_n \in \{1, \dots, k\} \text{ and } A_{x_n, x_{n+1}} = 1 \text{ for all } n \geq 0\right\}.$$

These infinite sequences can be seen as infinite paths in a directed graph $\mathcal{G}_A$ with vertices labeled $1, \dots, k$ and a directed edge from vertex i to j if and only if $A_{i,j} = 1$. We will therefore refer to the numbers $1, \dots, k$ as the *states* of Σ_A. We equip Σ_A with the *shift map* $\sigma : \Sigma_A \to \Sigma_A$ defined by

$$\sigma((x_n)_{n=0}^{\infty}) = (x_{n+1})_{n=0}^{\infty}$$

to obtain a dynamical system (Σ_A, σ).

Consider a finite ordered string $x_0, \dots, x_{m-1} \in \{1, \dots, k\}$ where $A_{x_j, x_{j+1}} = 1$ for each $j = 0, \dots, m-2$. The *cylinder set* associated to this string is the subset of Σ_A given by

$$[x_0, \dots, x_{m-1}] := \left\{(y_n)_{n=0}^{\infty} \in \Sigma_A : y_j = x_j \text{ for } j = 0, \dots, m-1\right\}.$$

We endow Σ_A with a topology by declaring the set of all cylinder sets to be an open basis.

The system (Σ_A, σ) is said to be *mixing* if for any two open sets $U, V \subset \Sigma_A$ there is $N \geq 1$ such that $\sigma^n(U) \cap V \neq \emptyset$ for all $n \geq N$. We say that (Σ_A, σ) is *transitive* if for any two open sets $U, V \subset \Sigma_A$ there exists $n \geq 1$ such that $\sigma^n(U) \cap V \neq \emptyset$. We have that (Σ_A, σ) is mixing if and only if A aperiodic and (Σ_A, σ) is transitive if and only if A is irreducible. We will often suppress the dependence of A in the notation for a subshift and will write (Σ, σ).

Example 2.7. Let Γ be a hyperbolic group equipped with finite generating set S. Consider a corresponding automatic structure $\mathcal{A} = (\mathcal{G} = (V, E), \pi, \{*\}, V)$ as discussed in Example 2.4. Suppose we have labeled the vertices in V by 1 to k where k is the cardinality of V. Then the graph $\mathcal{G}$ is encoded by a $k \times k$ transition matrix A where the (i, j)th entry of A is 1 if there is a directed edge joining vertex i to j and is 0 otherwise. This matrix gives a subshift (Σ_A, σ) that encodes (Γ, S). A subshift obtained in this way is never transitive (as $*$ only has outgoing edges) and it is not known whether, after removing $*$, it is always possible to find a connected graph $\mathcal{G}$ representing a given pair (Γ, S). In general it is possible to decompose the graph $\mathcal{G}$ into connected components (i.e., maximal connected subgraphs). If the the transition matrices for these subgraphs are $\mathcal{C}_1, \dots, \mathcal{C}_m$ then the subshifts $\Sigma_{\mathcal{C}_j}$ are each transitive.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

We call a connected component $\mathcal{C}$ *maximal* if the number of paths in $\mathcal{C}$ consisting of n edges grows like λ^n, where λ^n is the growth rate of the n spheres in the Cayley graph $\mathrm{Cay}(\Gamma, S)$, that is, the growth rate of the number of paths of length n in $\mathcal{C}$ is as large as possible.

Throughout the rest of the section (Σ, σ) will be a mixing subshift of finite type, and consider a function (which we will, at some points, refer to as a *potential*) $\psi : \Sigma \to \mathbb{R}$. We say that ψ is *constant on 2-cylinders* if ψ is constant on each set of the form $[x_0, x_1]$ where $x_0, x_1 \in \{1, \dots, k\}$ and $A_{x_0,x_1} = 1$. The assumption that ψ is constant on 2-cylinders guarantees that ψ has Markovian behaviour: the value that ψ takes at $x \in \Sigma$ depends only on the initial cylinder that x belongs to. That is, $\psi(x)$ does not depend on the future cylinders that x visits under the iterates of σ. Although this is a restrictive condition for a general function on Σ, it is not restrictive for our purposes. This is because we are interested in understanding the growth rate properties of functions that have lattice image: their image lies in $\alpha\mathbb{Z}$ for some $\alpha \in \mathbb{R}$ (as discussed in Subsection 1.3). Hölder continuous functions on Σ that have a lattice image have the property that their image only depends on finitely many symbols of the input. Then after relabeling the subshift Σ (i.e., moving to a topologically conjugate subshift) we can assume that the function is in fact constant on 2-cylinders. See for example the proof of Proposition 5.1 in [66] for an example of this argument. To summarize, functions that are constant on 2-cylinders naturally arise when studying discrete geometries. Lastly, it is worth mentioning that such functions have particularly nice properties: see for example Lemma 3.4 below.

For each $n \geq 1$, the nth *Birkhoff sum* of ψ is the function

$$\psi^n : \Sigma \to \mathbb{R} \ \text{ such that } \ \psi^n(x) := \psi(x) + \psi(\sigma(x)) + \cdots + \psi(\sigma^{n-1}(x)).$$

A point $x \in \Sigma$ is said to be *periodic* if $\sigma^n(x) = x$ for some $n \geq 1$. Such an n is called a *period* of x. Note that a periodic point has infinitely many periods. Given a periodic point we will assume that it comes with with a choice of period (which may not be its least period) which we will label $|x|$ (so that $\sigma^{|x|}(x) = x$).

Example 2.8. Consider the automaton $\mathcal{A}' = (\mathcal{G}' = (V', E'), \pi', \{*'\}, V')$ and the integer-valued function $\phi : V' \to \mathbb{Z}$ introduced in Example 2.4. Then, as discussed in Example 2.7, $\mathcal{A}'$ gives rise to a subshift of finite type (Σ, σ). Furthermore, the labeling ϕ defines a function $f : \Sigma \to \mathbb{Z}$ by

$$f(x) = \phi(v_{x_0}),$$

where $x = (x_n)_{n=0}^{\infty}$ and $v_{x_0} \in V$ is the vertex corresponding to the symbol x_0. The function f is constant on 2-cylinders and furthermore if $x = (x_n)_{n=0}^{\infty} \in \Sigma$ then

$$f^n(x) = |g_{x,n}|_{S_*}$$

where $g_{x,n} \in \Gamma$ is the group element obtained from multiplying the first n labelings in the infinite path corresponding to x, that is, if the first n edges in the path corresponding to x are $e_1, \dots, e_n$ then $f^n(x) = |\pi(e_1) \cdots \pi(e_n)|_{S_*}$. We note that, when $x \in \Sigma$ satisfies that $\sigma^n(x) = x$ then

$$mf^n(x) = f^{mn}(x) = |g_{x,mn}|_{S_*} = |g_{x,n}^m|_{S_*}$$

and so

$$f^n(x) = \lim_{m \to \infty} |g_{x,n}^m|_{S_*}/m = \ell_{S_*}[g_{x,n}].$$

Two functions $\psi, \varphi : \Sigma \to \mathbb{R}$, which we are assuming to be constant on 2-cylinders, are said to be *cohomologous* if there exists a continuous function $u : \Sigma \to \mathbb{R}$ such that $\psi(x) = \varphi(x) + u(\sigma(x)) - u(x)$ for all $x \in \Sigma$. By Livsic's Theorem [66, Proposition 3.7], ψ and φ are cohomologous if and only if $\psi^n(x) = \varphi^n(x)$ whenever $\sigma^n(x) = x$.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

The variational principle states that there is a unique σ-invariant Borel probability measure on (Σ, σ) that achieves the supremum

$$\mathrm{P}(\psi) := \sup_{\mu\in\mathcal{M}_\sigma}\left\{h_\mu(\sigma) + \int \psi\, d\mu\right\},$$

where $\mathcal{M}_\sigma$ is the collection of all σ-invariant Borel probability measures on Σ and $h_\mu(\sigma)$ denotes the *(metric) entropy* of σ with respect to the measure μ [66, Theorem 3.5]. The quantity $\mathrm{P}(\psi)$ is referred to as the *pressure* of ψ and the measure attaining the supremum is called the *equilibrium state* of ψ. When ψ is a constant function, the measure achieving the supremum for the pressure of ψ is the *measure of maximal entropy*. Furthermore the *topological entropy* $h = h(\sigma)$ of (Σ, σ) is given by $h = \mathrm{P}(0)$.

Consider the quantities

$$\alpha_{\min} := \inf_{\mu\in\mathcal{M}_\sigma}\int_\Sigma \psi\, d\mu \quad\text{and}\quad \alpha_{\max} := \sup_{\mu\in\mathcal{M}_\sigma}\int_\Sigma \psi\, d\mu.$$

The *large deviations principle* implies (since functions that are constant on 2-cylinders are Hölder) that there exists a real analytic, concave function $\mathcal{L}(\psi,\cdot) : \mathbb{R} \to \mathbb{R}_{>0}\cup\{\infty\}$ such that, for any nonempty sets $U \subset V \subset \mathbb{R}$ with U open and V closed we have

$$\begin{aligned} -\inf_{s\in U}\mathcal{L}(\psi,s) &\le \liminf_{n\to\infty}\frac{1}{n}\log\mu\left(x\in\Sigma : \frac{\psi^n(x)}{n}\in U\right)\\ &\le \limsup_{n\to\infty}\frac{1}{n}\log\mu\left(x\in\Sigma : \frac{\psi^n(x)}{n}\in V\right) \le -\inf_{s\in V}\mathcal{L}(\psi,s). \end{aligned} \tag{2.1}$$

See [56, Theorem in Section 2.1] for this result. The same result holds, with the same rate function $\mathcal{L}(\psi,\cdot)$, when we replace the sets

$$\mu\left(x\in\Sigma : \left|\frac{\psi^n(x)}{n} - \eta\right| < \epsilon\right)$$

with the sequence of (normalized) cardinalities

$$\frac{1}{\#\{x\in\Sigma : \sigma^n(x) = x\}}\#\left\{x\in\Sigma : \sigma^n(x) = x \text{ and } \left|\frac{\psi^n(x)}{n} - \eta\right| < \epsilon\right\}.$$

This is a well-known result that follows from the same proof as that of [56, Theorem in Section 2.1]. The function $\mathcal{L}(\psi,\cdot)$ is the Legendre transform of $t \mapsto \mathrm{P}(t\psi) - h$. That is,

$$-\mathcal{L}(\psi,s) = \inf_{t\in\mathbb{R}}(\mathrm{P}(t\psi) - h - ts). \tag{2.2}$$

Furthermore, $\mathcal{L}(\psi,\cdot)$ is finite on $[\alpha_{\min}, \alpha_{\max}]$ and is infinite otherwise. An alternative characterization for $\mathcal{L}$ is the following:

$$-\mathcal{L}(\psi,\eta) = \sup\left\{h_\mu(\sigma) : \mu\in\mathcal{M}_\sigma \text{ and } \int\psi\, d\mu = \eta\right\} - h.$$

A function $\psi : \Sigma \to \mathbb{R}$ is *lattice* if there are $a, b \in \mathbb{R}$ satisfying

$$\{\psi^n(x) + an : x\in\Sigma \text{ and } \sigma^n(x) = x \text{ for some } n \ge 1\} \subset b\mathbb{Z}.$$

If this is not the case then we say that ψ is *nonlattice*.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Remark 2.9. Suppose that ψ is lattice. Then ψ is cohomologous to a function of the form $a + b\varphi$ where $a, b \in \mathbb{R}$ and $\varphi : \Sigma \to \mathbb{Z}$ [66, Proposition 5.2]. When this is the case, the large deviations behaviour of ψ and φ over periodic orbits is the same, since $\psi^n(x) = an + b\varphi^n(x)$ when $\sigma^n(x) = x$.

2.5. Suspension flows

In this subsection we define suspension flows of subshifts of finite type. See Chapters 1 to 6 of [66] for more details on the results stated in this subsection. Let Σ_A be a transitive subshift of finite type and $r : \Sigma_A \to \mathbb{R}_{>0}$ a function that is constant on 2-cylinders. We note that in [66, Chapter 6] suspension flows are considered over mixing subshifts, however the same proofs (with some minor modifications) work when the subshift is transitive. We define the *suspension flow* of Σ_A^r to be the space

$$\Sigma_A^r = \{(x,t) \in \Sigma_A \times \mathbb{R}_{\geq 0} : 0 \leq t \leq r(x)\}/\sim$$

where $(x,t) \sim (r(x), 0)$, equipped with the flow $\sigma^r = (\sigma_t^r)_{t>0}$ so that σ_t^r sends (x,s) to $(x, s+t)$ for $s \in \mathbb{R}$. There is a natural metric on Σ_A^r which can be constructed as in [66]. We will not present the construction of this metric here as it is a little technical.

For a Hölder continuous function $\Phi : \Sigma_A^r \to \mathbb{R}$ we can define its *pressure* as

$$\mathrm{P}_{\sigma^r}(\Phi) = \sup_{m \in \mathcal{M}_{\sigma^r}} \left\{ h_m(\sigma^r) + \int_{\Sigma_A^r} \Phi \, dm \right\},$$

where $\mathcal{M}_{\sigma^r}$ is the space of σ_t^r- invariant Borel probability measures on Σ_A^r and $h_m(\sigma^r)$ is the entropy of the time-one map σ_1^r for the measure m [66, Section 6].

Remark 2.10. For such Φ, the pressure function $s \mapsto \mathrm{P}_{\sigma^r}(s\Phi)$ is real analytic. This is a well-known result that follows from Proposition 6.1 in [66] and the implicit function theorem.

Let $\delta_r > 0$ be the unique number such that $\mathrm{P}(-\delta_r r) = 0$ and write $\mu_{-\delta_r r}$ for the *equilibrium state* of $-\delta_r r$ on Σ_A. The *measure of maximal entropy* for Σ_A^r is (locally) given by

$$\frac{\mu_{-\delta_r r} \times \mathrm{Leb}}{\int r \, d\mu_{-\delta_r r}},$$

where Leb represents the Lebesgue measure along $\mathbb{R}_{\geq 0}$ ([66, Proposition 6.1]). That is, up to normalization, the measure of maximal entropy is the measure that acts as Lebesgue along the fibers of the suspension and as $\mu_{-\delta_r r}$ on the base. If we write m for the measure of maximal entropy then we have that

$$\left.\frac{d}{ds}\right|_{s=0} \mathrm{P}_{\sigma^r}(s\Phi) = \int \Phi \, dm$$

(see [66, Proposition 4.10]).

For $T > 0$ we will write $P(\Sigma_A^r, T)$ for the collection of periodic orbits of σ^r of length less than T. Given $R > 0$ we will write $P(\Sigma_A^r, R, T)$ for the collection of periodic orbits of length between $T - R$ and $T + R$.

Given a Hölder continuous function $\Phi : \Sigma_A^r \to \mathbb{R}$ it is a standard result that for any $R > 0$ sufficiently large

$$\lim_{T \to \infty} \frac{1}{T} \log \left(\sum_{\tau \in P(\Sigma_A^r, R, T)} e^{-s \int_\tau \Phi} \right) = \mathrm{P}_{\sigma^r}(-s\Phi)$$

for any $s \in \mathbb{R}$ (see for example [66, Proposition 5.10]).

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Lastly we recall that two functions $\Phi, \Psi : \Sigma_A^r \to \mathbb{R}$ are *cohomologous* if $\Phi - \Psi = u'$ where $u : \Sigma_A^r \to \mathbb{R}$ is continuously differentiable (along flow lines) and

$$u'(x) = \lim_{t \to 0} \frac{u(\sigma_t^r(x)) - u(x)}{t}.$$

Further the function $s \mapsto \mathrm{P}_{\sigma^r}(s\Phi)$ is a straight line if and only if Φ is cohomologous to a constant function.

3. Large deviations

In this section we discuss large deviations with shrinking intervals for potentials on mixing subshifts of finite type. The main result of the section is Theorem 3.2, and it will be used in the proof of Theorems 1.8 and 6.2 in subsequent sections.

Suppose that (Σ, σ) is a mixing subshift of finite type with $k \times k$ transition matrix A, and let μ denote its measure of maximal entropy. Also, let M be the least number such that A^M has strictly positive entries. The large deviations principle (2.1) from Subsection 2.4 implies that there is a real analytic, concave function $\mathcal{L}(\psi, \cdot) : [\alpha_{\min}, \alpha_{\max}] \to \mathbb{R}_{>0}$ such that the following holds: for any $\eta \in (\alpha_{\min}, \alpha_{\max})$

$$\lim_{\epsilon \to 0^-} \limsup_{n \to \infty} \frac{1}{n} \log \mu\left(x \in \Sigma : \left|\frac{\psi^n(x)}{n} - \eta\right| < \epsilon\right) = -\mathcal{L}(\psi, \eta). \tag{3.1}$$

Instead of taking two limits as above, first with respect to n and then with respect to ϵ, it is natural to ask the following.

Question 3.1. How quickly can a sequence δ_n decay to 0 as $n \to \infty$ so that we have

$$\lim_{n \to \infty} \frac{1}{n} \log \mu\left(x \in \Sigma : \left|\frac{\psi^n(x)}{n} - \eta\right| < \delta_n\right) = -\mathcal{L}(\psi, \eta) \tag{3.2}$$

for each $\eta \in (\alpha_{\min}, \alpha_{\max})$?

We refer to this problem as *large deviations with shrinking intervals*. We can ask the same question when the limit in (3.2) is replaced with the limit supremum.

Large deviations with shrinking intervals are best understood for functions $\psi : \Sigma \to \mathbb{R}$ that are nonlattice. For example, when ψ is nonlattice the local central limit theorem [42] implies that (3.2) holds when $\delta_n^{-1} = O(n)$. In [68] Pollicott and Sharp improved this result under an additional assumption. They showed that if ψ satisfies a non-Diophantine condition then there exist $\kappa > 0$ such that (3.2) holds when $\delta_n^{-1} = O(n^{1+\kappa})$.

When ψ is lattice, large deviations with shrinking intervals are not as well-understood. The aim of this section is to study (3.2) for functions ψ that are constant on 2-cylinders and are lattice.

We now state our large deviations theorem with shrinking intervals. Write $\mathcal{L}(\psi, \cdot) : [\alpha_{\min}, \alpha_{\max}] \to \mathbb{R}_{>0}$ be the function introduced above.

Theorem 3.2. *Suppose that (Σ, σ) is a mixing subshift of finite type and that $\psi : \Sigma \to \mathbb{R}$ is a function that is constant on 2-cylinders. Then there exists $C > 0$ such that for any $\eta \in (\alpha_{\min}, \alpha_{\max})$*

$$\limsup_{n \to \infty} \frac{1}{n} \log\left(\#\left\{x \in \Sigma : \sigma^n(x) = x \text{ and } \left|\frac{\psi^n(x)}{n} - \eta\right| < \frac{C}{n^2}\right\}\right) = h - \mathcal{L}(\psi, \eta) \tag{3.3}$$

where h is the topological entropy of (Σ, σ). Furthermore we can take

$$C = \frac{4M^2(1 + k^2)^2(\alpha_{\max} - \alpha_{\min})}{\sqrt{5}}.$$

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

In the case that ψ takes values in $\mathbb{Z}$ then there exists $\eta \in (\alpha_{\min}, \alpha_{\max})$ and $\epsilon > 0$ such that

$$\limsup_{n\to\infty} \frac{1}{n} \log\left(\#\left\{x \in \Sigma : \sigma^n(x) = x \text{ and } \left|\frac{\psi^n(x)}{n} - \eta\right| < \frac{\epsilon}{n^2}\right\}\right) = 0.$$

Remark 3.3. i) This theorem also holds if we only assume that (Σ, σ) is transitive (opposed to mixing). Indeed, if (Σ, σ) is transitive we can find an integer $p \geq 1$ such that (Σ, σ^p) decomposes into p disjoint σ^p-invariant sets. These σ^p-invariant sets are mixing subshifts of finite type when equipped with σ^p. We can then apply the mixing version of our theorem above to these subshifts to deduce the transitive version. ii) Our result significantly improves the decay rate implied by the local limit theorem under the nonlattice assumption [73, Theoreme 5]. Furthermore, the case that ψ takes values in a lattice shows that the decay rates obtained in the first part of Theorem 3.2 are optimal.

For the rest of the section we note the correspondence between periodic orbits of (Σ, σ) and cycles (i.e., closed paths) in the adjacency graph $\mathcal{G}_A$. We say that a cycle is *simple* if it does not visit any state more than once.

To prove the above result we start with the following observation.

Lemma 3.4. *Suppose that (Σ, σ) is a mixing subshift of finite type and that $\psi : \Sigma \to \mathbb{R}$ is a function that is constant on 2-cylinders. Then*

$$\alpha_{\min} = \inf_{\sigma^n(x)=x} \frac{\psi^n(x)}{n} \quad \text{and} \quad \alpha_{\max} = \sup_{\sigma^n(x)=x} \frac{\psi^n(x)}{n},$$

and furthermore there exist periodic orbits $\overline{x}, \overline{y}$ that achieve these values, that is, $\psi^{|\overline{x}|}(\overline{x}) = |\overline{x}|\alpha_{\min}, \psi^{|\overline{y}|}(\overline{y}) = |\overline{y}|\alpha_{\max}$. Here the infimum and supremum are over all the periodic orbits.

Proof. The first statement follows from a result of Sigmund [82, Theorem 1] which states that the set of probability measures supported on periodic orbits is dense in $\mathcal{M}_\sigma$ (equipped with the weak-$*$ topology). For the furthermore statement note that each periodic orbit can be written as a disjoint union of simple cycles. It follows easily that the above infimum and supremum are attained by simple cycles (and powers of them). □

We now want to use the orbits $\overline{x}, \overline{y}$ from the lemma above to construct periodic orbits along which the average value of ψ approximates a given $\eta \in (\alpha_{\min}, \alpha_{\max})$ (see Proposition 3.8). We begin with the following observation.

Lemma 3.5. *Take an interval $(s, t) \subset \mathbb{R}$ and a number $\eta \in (s, t)$. Then there are infinitely many $n \geq 1$ for which there exist integers $0 \leq a, b \leq n$ with $a + b = n$ and such that*

$$\left|\frac{as + bt}{n} - \eta\right| \leq \frac{t - s}{\sqrt{5}\, n^2}.$$

Proof. Note that it suffices to prove this result when $s = 0$, $t = 1$. The general result then follows by shifting and rescaling the interval $(0, 1)$ into (s, t). When $s = 0$, $t = 1$ the result follows from the well-known Hurwitz's Theorem [51] from Diophantine approximation. □

We also need the following lemma. Suppose that the simple periodic orbit $\overline{x}$ realizing $\alpha_{\max}$ has initial state i and that the initial state for $\overline{y}$ realizing $\alpha_{\min}$ is j (from Lemma 3.4). Further assume that we repeat $\overline{x}$ and $\overline{y}$ by each other's periods so that they both have period l satisfying $1 < l \leq k^2$.

Lemma 3.6. *Suppose that (Σ, σ) is a mixing subshift of finite type and that $\psi : \Sigma \to \mathbb{R}$ is a function that is constant on 2-cylinders such that $\alpha_{\min} < \alpha_{\max}$. Then we can find periodic orbits x, y both*

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

with period at most $2M(1+k^2)$ *such that the initial state for* x *is* i*, the initial state for* y *is* j *and we have that*

$$\alpha_{\min} < \frac{\psi^{|x|}(x)}{|x|} < \frac{\psi^{|y|}(y)}{|y|} < \alpha_{\max}.$$

Proof. Consider the start state i and find a path p from i to j of length M and a path q from j to i of length M. By composing the path p with r (to be chosen later) repeats of a single cycle of $\overline{y}$ and q we obtain a periodic orbit x of period $2M + lr$. Further we have that

$$\frac{\psi^{|x|}(x)}{|x|} \leq \frac{2M\alpha_{\max} + lr\alpha_{\min}}{2M + lr}$$

and x starts at i. Similarly we find y starting at j with $|y| = 2M + lr$ and

$$\frac{\psi^{|y|}(y)}{|y|} \geq \frac{2M\alpha_{\min} + lr\alpha_{\max}}{2M + lr}.$$

Now, as long as r is chosen so that

$$2M\alpha_{\min} + lr\alpha_{\max} > 2M\alpha_{\max} + lr\alpha_{\min}$$

x, y will satisfy the final inequality in the lemma. Note that this inequality is satisfied for $r = 2M$ in which case x and y have periods $2M + 2Ml \leq 2M(1+k^2)$ as required. □

We also require the following.

Lemma 3.7. *Suppose that* (Σ, σ) *is a mixing subshift of finite type and that* $\psi : \Sigma \to \mathbb{R}$ *is a function that is constant on* 2*-cylinders. Suppose that there exist periodic orbits* x, y *both with period* l *and same initial state such that* $lA = \psi^l(x) < \psi^l(y) = lB$*. Then there exists* $\overline{C} > 0$ *such that for any* $\eta \in (A, B)$ *we can find an infinite sequence of periodic orbits* $x_n \in \Sigma$ *such that*

$$\left|\frac{\psi^{|x_n|}(x_n)}{|x_n|} - \eta\right| \leq \frac{\overline{C}}{|x_n|^2}.$$

Furthermore we can take

$$\overline{C} = \frac{(\alpha_{\max} - \alpha_{\min})l^2}{\sqrt{5}}.$$

Proof. By Lemma 3.5 there exist infinitely many $n \geq 1$ such that the following holds. There are integers $n_1, n_2 \geq 0$ with $n_1 + n_2 = n$ and such that

$$\left|\frac{n_1A + n_2B}{n_1 + n_2} - \eta\right| \leq \frac{B - A}{\sqrt{5}n^2} \leq \frac{(\alpha_{\max} - \alpha_{\min})l^2}{\sqrt{5}(ln)^2}.$$

Since x and y have the same initial vertex, we can form a new periodic orbit by composing n_1 copies of x followed by n_2 copies of y. This creates a periodic orbit z of orbit length nl with the property that

$$\frac{\psi^{nl}(z)}{nl} = \frac{n_1A + n_2B}{n_1 + n_2} \quad \text{and so} \quad \left|\frac{\psi^{nl}(z)}{nl} - \eta\right| \leq \frac{(\alpha_{\max} - \alpha_{\min})l^2}{\sqrt{5}|z|^2}.$$

Since we can run this construction for infinitely many n, the result follows. □

To obtain uniformity over η, that is, to show the existence of C in Theorem 3.2, we need to upgrade Lemma 3.7 using Lemma 3.6.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Proposition 3.8. *Suppose that* (Σ, σ) *is a mixing subshift of finite type and that* $\psi : \Sigma \to \mathbb{R}$ *is a function that is constant on 2-cylinders. Then there exists* $C > 0$ *such that for any* $\eta \in (\alpha_{\min}, \alpha_{\max})$ *there exists an infinite sequence of periodic orbits* $x_n \in \Sigma$ *such that*

$$\left| \frac{\psi^{|x_n|}(x_n)}{|x_n|} - \eta \right| \le \frac{C}{|x_n|^2}.$$

Furthermore we can take

$$C = \frac{4M^2(1+k^2)^2(\alpha_{\max} - \alpha_{\min})}{\sqrt{5}}.$$

Proof. We can assume that ψ is not cohomologous to a constant function (otherwise the conclusion is clear), so that $\alpha_{\min} < \alpha_{\max}$.

We split the interval $(\alpha_{\min}, \alpha_{\max}) = I_1 \cup I_2$ into the two (nondisjoint) intervals

$$I_1 = \left(\alpha_{\min}, \frac{\psi^{|y|}(y)}{|y|}\right) \text{ and } I_2 = \left(\frac{\psi^{|x|}(x)}{|x|}, \alpha_{\max}\right)$$

where x, y are the orbits constructed in Lemma 3.6. In particular, both $|x|$ and $|y|$ are bounded above by $2M(1+k^2)$. We can now apply Lemma 3.7 to both of the intervals I_1 and I_2 to deduce the result. □

Definition 3.9. Let ψ, Σ be as above. Suppose $w \in (\alpha_{\min}, \alpha_{\max})$ is chosen so that there exists x with $\sigma^n(x) = x$ and $\psi^n(x) = nw$. We define $d(\psi, w)$ to be the greatest common divisor of all numbers $n \ge 1$ such that

$$\#\{x \in \Sigma : \sigma^n(x) = x, \psi^n(x) = wn\} > 0.$$

If $\#\{x \in \Sigma : \sigma^n(x) = x, \psi^n(x) = wn\} = 0$ for all n then we set $d(\psi, w) = 0$.

Note that $d(\psi, w) = 0$ for all but countably many values of w. This is because the values for which $d(\psi, w) > 0$ are contained in the rational span of the values that the averaged Birkhoff sum of ψ attains on simple cycles. We now state the key result of Marcus and Tuncel that we need to prove large deviations with shrinking intervals. Recall that $\mathcal{L}$ represents the large deviations rate function introduced in Subsection 2.4 (see (2.1)).

Theorem 3.10 (Theorem 14 [59]). *For any* $\xi > 0$ *and any closed subset* $W \subset (\alpha_{\min}, \alpha_{\max})$ *there exist* $r, N \in \mathbb{Z}_{\ge 0}$ *and* $\delta > 0$ *such that, for any* $n \ge N$ *with* $d(\psi, w)|n$,

$$\#\{x \in \Sigma : \sigma^n(x) = x, \psi^n(x) = wn\} \ge \delta n^{-r} e^{n(h-(\mathcal{L}(\psi,w))-\xi)} \tag{3.4}$$

for each $w \in W$ *satisfying* $d(\psi, w) > 0$.

To make use of this result we need to control the values of $d(\psi, w)$ as w takes values in a shrinking interval.

Lemma 3.11. *There exists* $C > 0$ *such that for any* $\eta \in (\alpha_{\min}, \alpha_{\max})$ *we can find a sequence* $w_n \in (\alpha_{\min}, \alpha_{\max})$ *and a sequence* $x_n \in \Sigma$ *of periodic orbits such that*

$$w_n \in \left[\eta - \frac{C}{|x_n|^2}, \eta + \frac{C}{|x_n|^2}\right] \text{ with } \frac{\psi^{|x_n|}(x_n)}{|x_n|} = w_n$$

and $d(\psi, w_n) \,\big|\, |x_n|$.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Proof. By Proposition 3.8 there exist $C > 0, N \geq 1$ depending only on ψ, Σ such that for any $\eta \in (\alpha_{\min}, \alpha_{\max})$ and $n \geq N$ there is a sequence of periodic orbits $x_n \in \Sigma$ with periods $|x_n|$ (i.e., $\sigma^{|x_n|}(x_n) = x_n$) such that

$$\left|\frac{\psi^{|x_n|}(x_n)}{|x_n|} - \eta\right| \leq \frac{C}{|x_n|^2}.$$

We define $w_n = \frac{\psi^n(x_n)}{n}$ and note that $d(\psi, w_n) \,\big|\, |x_n|$ if and only if

$$\#\left\{x \in \Sigma : \sigma^{|x_n|}(x) = x, \frac{\psi^{|x_n|}(x_n)}{|x_n|} = w_n\right\} > 0.$$

Hence we are done. □

For a sequence $(\delta_n)_{n\geq 1}$ of positive numbers and $n \geq 1$ we define

$$F_n(\eta, \delta_n) = \left\{x \in \Sigma : \sigma^n(x) = x, \left|\frac{\psi^n(x)}{n} - \eta\right| < \delta_n\right\}$$

for $\eta \in (\alpha_{\min}, \alpha_{\max})$. Intuitively, this set contains the periodic orbits of period n along which the average value of ψ approximates η up to an error less than δ_n. Theorem 3.2 is concerned with the exponential growth rate of the cardinality of $F_n(\eta, Cn^{-2})$ for some constant $C > 0$ as $n \to \infty$. We therefore want to obtain bounds on $\#F_n(\eta, Cn^{-2})$ which we achieve in the following proposition.

Proposition 3.12. *Suppose that (Σ, σ) is a mixing subshift of finite type and that $\psi : \Sigma \to \mathbb{R}$ is a function that is constant on 2-cylinders. Then there exists $C > 0$ such that for any $\eta \in (\alpha_{\min}, \alpha_{\max})$ and $\xi > 0$ there exist $\delta, r, M > 0$ and a sequence of integers n_l with $n_l \to \infty$ as $l \to \infty$ such that*

$$\#F_{n_l}(\eta, Cn_l^{-2}) \geq \delta n_l^{-r} e^{n_l((h-\mathcal{L}(\psi,\eta))-\xi)}$$

for all $l \geq 1$.

Remark 3.13. For η with $d(\psi, \eta) > 0$ this result follows immediately from estimates due to Marcus and Tuncel [59]. The main strength of the above estimates are that they hold for all values of $\eta \in (\alpha_{\min}, \alpha_{\max})$.

Proof of Proposition 3.12. We use Proposition 3.8 to find a sequence of periodic orbits $(x_n)_n$ such that

$$|w_n - \eta| \leq \frac{C}{|x_n|^2} \quad \text{where} \quad w_n = \frac{\psi^{|x_n|}(x_n)}{|x_n|}$$

for each $n \geq 1$. By (3.4), for any $\xi > 0$ and for all n sufficiently large we have that

$$\#\left\{x \in \Sigma : \sigma^{|x_n|}(x) = x \text{ and } \left|\frac{\psi^{|x_n|}(x_n)}{|x_n|} - \eta\right| < \frac{C}{|x_n|^2}\right\} \geq \delta |x_n|^{-r} e^{|x_n|((h-\mathcal{L}(\psi,w_n))-\xi)} \tag{3.5}$$

for some $\delta, r > 0$. Now, by analyticity of $\mathcal{L}$, there is $\widetilde{C} > 0$ such that

$$|\mathcal{L}(\psi, \eta) - \mathcal{L}(\psi, w_n)| \leq \widetilde{C}|\eta - w_n| = O(|x_n|^{-2})$$

(where the implied error constant is independent of n). Substituting this into the right hand side of (3.5) provides the required bound concluding the proof. □

Proof of Theorem 3.2. Note that the large deviations principle where we count over periodic orbits (see the discusion after (2.1)) implies that

$$\limsup_{n\to\infty} \frac{1}{n} \log\left(\#\left\{x \in \Sigma : \sigma^n(x) = x, \left|\frac{\psi^n(x)}{n} - \eta\right| < \frac{C'}{n^2}\right\}\right) \leq h - \mathcal{L}(\psi, \eta)$$

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

for any $C' > 0$ and $\eta \in (\alpha_{\min}, \alpha_{\max})$. Proposition 3.12 also implies that for $C > 0$ (as in the proposition)

$$\limsup_{n\to\infty} \frac{1}{n} \log\left(\#\left\{x \in \Sigma_C : \sigma^n(x) = x, \left|\frac{\psi^n(x)}{n} - \eta\right| < \frac{C}{n^2}\right\}\right) \geq h - \mathcal{L}(\psi, \eta) \tag{3.6}$$

for any $\eta \in (\alpha_{\min}, \alpha_{\max})$. This proves the identity (3.3), and the estimate for C follows from Proposition 3.8.

We now finish by proving the final statement of the theorem which assumes that ψ takes values in $\mathbb{Z}$. When this is the case, the set $\left\{\frac{\psi^n(x)}{n} : \sigma^n(x) = x\right\}$ contains rational numbers all with denominator at most n. By [51, Satz II] there is $\eta \in (\alpha_{\min}, \alpha_{\max})$ such that if $\epsilon > 0$ is sufficiently small then for all but finitely many values of n,

$$\left|\frac{\psi^n(x)}{n} - \eta\right| > \frac{\epsilon}{n^2} \text{ for all } x \in \Sigma \text{ with } \sigma^n(x) = x.$$

Hence $F_n(\eta, \epsilon n^{-2})$ is empty for all n sufficiently large. □

4. Large deviations for pairs of word metrics

In this section we prove Theorem 1.8.

Throughout the proof, we will follow the same terminology and notation that we established above. Before presenting the proof, we restate the theorem for the reader's convenience.

Theorem 4.1. *Let Γ be a nonelementary hyperbolic group and consider two finite generating sets S, S_* for Γ with exponential growth rates v_S, v_{S_*}. Then there exist $C > 0$ and a real analytic, concave function $\mathcal{I} : [\mathrm{Dil}(S, S_*)^{-1}, \mathrm{Dil}(S_*, S)] \to \mathbb{R}_{>0}$ such that for any $\eta \in (\mathrm{Dil}(S, S_*)^{-1}, \mathrm{Dil}(S_*, S))$ we have*

$$0 < \limsup_{T\to\infty} \frac{1}{T} \log\left(\#\left\{[g] \in \mathbf{conj} : \ell_S[g] < T, |\ell_{S_*}[g] - \eta\ell_S[g]| < \frac{C}{T}\right\}\right) = \mathcal{I}(\eta) \leq v_S.$$

Furthermore, we have equality in the above inequality if and only if

$$\eta = \tau(S_*/S) := \lim_{T\to\infty} \frac{1}{\#\{[g] \in \mathbf{conj} \colon \ell_S[g] < T\}} \sum_{\ell_S[g]<T} \frac{\ell_{S_*}[g]}{T}.$$

Proof. Without loss of generality assume that d_S, d_{S_*} are not roughly similar (i.e., there does not exist $\tau, C > 0$ such that $|d_S(g,h) - \tau d_{S_*}(g,h)| < C$ for all $g, h \in \Gamma$). As discussed in Example 2.8, by Lemma 3.8 and Example 3.9 in [14] we can find a Cannon coding for (Γ, S) with corresponding shift space (Σ, σ) and a constant on 2-cylinders function $\psi : \Sigma \to \mathbb{Z}$ satisfying the following: if $z = (z_0, z_1, \ldots, z_{n-1}, z_0, \ldots) \in \Sigma$ is a periodic orbit of period n, then the nth Birkhoff sum of ψ on z outputs the value $\ell_{S_*}[g]$, where $g \in \Gamma$ is the group element obtained by multiplying the labelings in the finite path $z_0, z_1, \ldots, z_{n-1}$.

Fix a maximal component Σ_C in Σ (as in Example 2.7). We know that the function ψ satisfies a large deviations principle (as discussed at (3.1)) over the periodic orbits on Σ_C, as discussed after (2.1). We set $\mathcal{I} = h - \mathcal{L}$, where h is the topological entropy of (Σ, σ) and $\mathcal{L} = \mathcal{L}(\psi, \cdot)$ is the Legendre transform (as defined in (2.2) above). By [16, Lemma 3.3] we know that $\mathcal{I}$ is also the Legendre transform of the Manhattan curve for d_S, d_{S_*}. In particular,

$$\alpha_{\min} = \inf_{\sigma^n(x)=x} \frac{\psi^n(x)}{n} = \mathrm{Dil}(S, S_*)^{-1} \text{ and } \alpha_{\max} = \sup_{\sigma^n(x)=x} \frac{\psi^n(x)}{n} = \mathrm{Dil}(S_*, S)$$

where the infimum/supremum is taken over all periodic orbits in $\Sigma_{\mathcal{C}}$. Now, by Theorem 3.2 we obtain the same limits with shrinking intervals, that is, $\eta \in (\alpha_{\min}, \alpha_{\max})$

$$\limsup_{n\to\infty} \frac{1}{n} \log\left(\#\left\{x \in \Sigma : \sigma^n(x) = x \text{ and } \left|\frac{\psi^n(x)}{n} - \eta\right| \le \frac{C}{n^2}\right\}\right) = \mathcal{I}(\eta)$$

for some C independent of η. It is possible that the system $(\Sigma_{\mathcal{C}}, \sigma)$ is not mixing but is instead transitive. This is no issue as explained in Remark 3.3. Now note that each periodic orbit has a corresponding conjugacy class as described above: if $z = (z_0, z_1, \ldots, z_{n-1}, z_0, \ldots) \in \Sigma$ is a periodic orbit of period n then we associate to it the conjugacy class $[g]$ with $\ell_S[g] = n$ where g is the group element obtained by multiplying the labelings along the path $z_0, z_1, \ldots, z_{n-1}$. Furthermore [17, Lemma 4.2] says that the periodic orbits of period n overcount the number of conjugacy classes by at most a linear factor in n. Hence we must have

$$\limsup_{n\to\infty} \frac{1}{n} \log(\#\{[g] \in \mathbf{conj} : \ell_S[g] = n, |\ell_{S_*}[g] - \eta\ell_S[g]| < \epsilon\}) = \mathcal{I}(\eta) \le v_S$$

for any $\epsilon > 0$ and $\eta \in (\alpha_{\min}, \alpha_{\max})$ (note that $h = v_S$). By [19, Theorem 1.1] we have that $\mathcal{I}$ is maximized (and equals v_S) when

$$\eta = \lim_{T\to\infty} \frac{1}{\#\{[g] \in \mathbf{conj} : \ell_S[g] < T\}} \sum_{\ell_S[g]<T} \frac{\ell_{S_*}[g]}{T}.$$

This concludes the proof. □

As a consequence we deduce Corollary 1.9.

Proof of Corollary 1.9. The first part of the theorem, when η is irrational, follows as a direct corollary of Theorem 1.8. To deduce the final statement when η is rational it suffices to show that, when (Σ, σ) is a mixing subshift of finite type and $\psi : \Sigma \to \mathbb{Z}$ is a function that is constant of 2-cylinders, then the following holds: if $p/q \in (\alpha_{\min}, \alpha_{\max}) \cap \mathbb{Q}$ then there exists a periodic orbit $x \in \Sigma$ such that $\psi^{|x|}(x) = |x|p/q$. It is an easy exercise to verify this and so we leave it to the reader. □

Example 4.2. In this example we show how to apply Theorem 1.8 to a pair of word metrics on a free group. We calculate the exact value of $\mathcal{I}$ (evaluated at a natural value) and provide explicit constants which determine how similar the length spectra of the two word metrics are.

Let $F_2 = \langle a, b\rangle$ be the free group on two generators and consider the generating sets

$$S = \{a, b, a^{-1}, b^{-1}\} \text{ and } S_* = \{a, b, ab, a^{-1}, b^{-1}, (ab)^{-1}\}.$$

The corresponding word metrics have exponential growth rates $v_{S_*} = \log(4)$ and $v_S = \log(3)$. The Manhattan curve θ_{S/S_*} was computed in [19] and is given by

$$\theta_{S/S_*}(t) = \log\left(\frac{1}{2}e^{-t}\left(e^{-t} + \sqrt{e^{-t}(e^{-t}+8)} + 4\right)\right).$$

From the definition we see that $\mathcal{I}(v_{S_*}/v_S) = \log(4) - \Lambda/\log(3)$ where Λ is the constant

$$\Lambda = \sup_{t\in\mathbb{R}}\left\{\log(4) - \frac{\log(4)}{\log(3)}t - \theta_{S/S_*}(t)\right\}.$$

This can be computed (using, say Wolfram Alpha) to be

$$\log(16)\log\left(\log\left(\frac{4}{3}\right)\right) + \log(9)\log\left(\frac{\log\left(\frac{3}{2}\right)}{\log\left(\frac{4}{3}\right)}\right) + \log(4)\left(\log(2) - \log\left(\log\left(\frac{3}{2}\right)\log(2)\right)\right).$$

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Theorem 1.8 then implies that there is $C > 0$ such that

$$\limsup_{T\to\infty} \frac{1}{T}\log\left(\#\left\{[g]\in\mathbf{conj} : \ell_{S_*}[g] < T, |v_S\ell_S[g] - v_{S_*}\ell_{S_*}[g]| \leq \frac{C}{T}\right\}\right) = \log(4) - \frac{\Lambda}{\log(3)}$$
$$\approx 1.3679878759$$

and similarly we also have

$$\limsup_{T\to\infty} \frac{1}{T}\log\left(\#\left\{[g]\in\mathbf{conj} : \ell_{S}[g] < T, |v_S\ell_S[g] - v_{S_*}\ell_{S_*}[g]| \leq \frac{C}{T}\right\}\right) = \log(3) - \frac{\Lambda}{\log(4)}$$
$$\approx 1.0841047424.$$

In this case we can set $k = 6, M = 2, \alpha_{\max} = 2, \alpha_{\min} = 1$ by [19]. Also, it was computed in [19] that $-\theta'_{S/S_*}(0) = 4/3$. We therefore have that

$$\frac{4M^2(1+k^2)^2(\alpha_{\max} - \alpha_{\min})}{\sqrt{5}} = \frac{4(2^2)(1+6^2)(2-1)}{\sqrt{5}} = \frac{592}{\sqrt{5}} \approx 264.7 \leq 300$$

and

$$\limsup_{T\to\infty} \frac{1}{T}\log\left(\#\left\{[g]\in\mathbf{conj} : \ell_{S_*}[g] < T, |3\ell_S[g] - 4\ell_{S_*}[g]| \leq \frac{300}{T}\right\}\right) = v_{S_*}.$$

Therefore the length spectra of $3d_S$ and $4d_{S_*}$ are within distance 300 on a set of full exponential growth rate for d_{S_*}. In particular, for any $\eta \in (1, 2)$ we can find an infinite sequence of conjugacy classes $[g_n] \in \mathbf{conj}(\Gamma)$ such that

$$\left|\frac{\ell_S[g_n]}{\ell_{S_*}[g_n]} - \eta\right| \leq \frac{300}{|g_n|^2_{S_*}}.$$

5. Encoding cubulations via finite-state automata

In this section we study further the class $\mathfrak{G}$ defined in the introduction. Recall that a group belongs to $\mathfrak{G}$ if it is not virtually cyclic, it admits a virtually co-special cubulation with a contracting element, and its set of convex-cocompact subgroups is independent of the cubulation. See Subsection 2.2 for the notions of virtual co-specialness and convex-cocompact subgroups. We prove Proposition 5.2, which provides plenty of examples of groups in this class. Then we construct a finite-state automaton that encodes a pair of cubulations of a group in $\mathfrak{G}$, our main result being Theorem 5.11, that implies most of the claims in Theorem 1.4 from the introduction. This will be done in the greater generality of the class $\mathfrak{X}$ of pairs of compatible cubulations. Theorem 5.11 is also key to prove our main results in Section 6.

5.1. The classes $\mathfrak{G}$ and $\mathfrak{X}$

Throughout this and the next section we will work with the following notion of compatibility of pairs of group actions on CAT(0) cube complexes.

Definition 5.1. Let $\mathfrak{X}$ be the class of triplets $(\Gamma, \mathcal{X}, \mathcal{X}_*)$, where Γ is a nonvirtually cyclic group acting cocompactly on the CAT(0) cube complexes $\mathcal{X}, \mathcal{X}_*$ and satisfying:

(1) the action on $\mathcal{X}$ is proper and virtually co-special;
(2) every hyperplane stabilizer for the action on $\mathcal{X}_*$ is convex-cocompact for the action on $\mathcal{X}$; and,
(3) the action of Γ on $\mathcal{X}$ has a contracting element.

Note that in the definition above we do not require the action on $\mathcal{X}_*$ to be proper.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Since hyperplane stabilizers are always convex-cocompact and being a contracting element does not depend on the cubulation [39, Lemma 4.6], it follows that $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ whenever $\Gamma \in \mathfrak{G}$ and $\mathcal{X}, \mathcal{X}_*$ are cubulations of Γ.

If $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$, we should think of $\mathcal{X}$ and $\mathcal{X}_*$ as "compatible" cubulations of Γ with well-behaved counting properties. Compatibility is guaranteed by (2), by means of Proposition 5.7 that allows us to combine $\mathcal{X}$ and $\mathcal{X}_*$ into a new cubulation of Γ. Then (1) implies that the new cubulation is also virtually co-special, allowing us to use the automaton from Example 2.5. This automaton is upgraded in Theorem 5.11 to an automaton that "remembers" both $\mathcal{X}$ and $\mathcal{X}_*$. Then (3) is used to relate counting computations over this automaton to counting properties of the length functions $\ell_{\mathcal{X}}, \ell_{\mathcal{X}_*}$ over conjugacy classes of Γ. This counting is done in Section 6.

The main result of this subsection is the following proposition, which tells us that the class $\mathfrak{G}$ is much bigger than the class of cubulable hyperbolic groups.

Proposition 5.2. *The following classes of groups are contained in* $\mathfrak{G}$.

i) *Hyperbolic cubulable groups that are nonelementary.*
ii) *Right-angled Artin groups of the form* $\Gamma = A_G$*, where* G *is a finite graph with more than one vertex, that is not a join and such that* $\mathrm{st}(v)$ *is not contained in* $\mathrm{st}(w)$ *for every pair of distinct vertices* v, w *of* G.
iii) *Right-angled Coxeter groups that are not virtually cyclic or direct products and are of the form* $\Gamma = W_G$*, where* G *is finite and does not have any loose squares.*

Moreover, if Γ *is cubulable and hyperbolic relative to groups belonging to* $\mathfrak{G}$*, then* Γ *belongs to* $\mathfrak{G}$.

For $\Gamma = W_G$ a right-angled Coxeter group, a *loose square* is a full subgraph $\Delta \subset G$ that is a square, and such that for every maximal subgraph $\Lambda \subset G$ with W_Λ virtually abelian, either $\Delta \subset \Lambda$ or $\Delta \cap \Lambda$ generates a finite subgroup of Γ.

Many nonrelatively hyperbolic right-angled Artin and Coxeter groups satisfy assumptions ii) and iii) of Proposition 5.2. Indeed, a RAAG is nontrivially relatively hyperbolic if and only if it is a nontrivial free product, and hence RAAGs with underlying graphs n-agons for $n \geq 5$ are not relatively hyperbolic and belong to $\mathfrak{G}$.

For the proof of Proposition 5.2 we require two results, the first one being an observation about virtual specialness.

Lemma 5.3. *Let* $\mathcal{X}, \mathcal{X}_*$ *be cubulations of the group* Γ *and assume that the action on* $\mathcal{X}$ *is virtually co-special. If every hyperplane stabilizer for the action on* $\mathcal{X}_*$ *is convex-cocompact with respect to* $\mathcal{X}$*, then the action of* Γ *on* $\mathcal{X}_*$ *is virtually co-special.*

Proof. If $\mathcal{X}$ and $\mathcal{X}_*$ satisfy the assumptions of the lemma, then all the double cosets of the hyperplane stabilizers for the action of Γ on $\mathcal{X}_*$ are separable by [71, Theorem A.1]. Then the action of on $\mathcal{X}_*$ is virtually co-special by the double-cosets criterion [47, Theorem 9.19]. □

The second result we need is a criterion of convex-cocompactness for subgroups of cubulable relatively hyperbolic groups, which may be of independent interest and whose proof is postponed to the appendix.

Proposition 5.4. *Let* Γ *be a relatively hyperbolic group acting properly and cocompactly on the* $\mathrm{CAT}(0)$ *cube complex* $\mathcal{X}$*. Then the following are equivalent for a subgroup* $H < \Gamma$.

(1) H *is convex-cocompact for the action of* Γ *on* $\mathcal{X}$.
(2) H *is relatively quasiconvex and* $H \cap P$ *is convex-cocompact for the action of* Γ *on* $\mathcal{X}$ *for any maximal parabolic subgroup* $P < \Gamma$.

Proof of Proposition 5.2. By Agol's Theorem [1, Theorem 1.1] every cubulation of a hyperbolic group is virtually co-special. Moreover, the class of convex-cocompact subgroups for any such cubulation coincides with the class of quasiconvex subgroups [47, Proposition 7.2]. Since any loxodromic element in a hyperbolic group is contracting, this solves the proposition for groups in class i).

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Now, let Γ be a group belonging to the class ii) (respectively iii)). In [35, Theorem A] (resp. [35, Corollary C]) it was proven that Γ has a unique *cubical coarse median structure*. By [35, Theorem 2.15] this uniqueness result is equivalent to the class of convex-cocompact subgroups being independent of the cubulation. Since finitely generated right-angled Artin (resp. Coxeter) groups are virtually special, Γ satisfies Items (1) and (2) of Definition 1.1. The existence of a contracting element for Γ for a geometric group action on a CAT(0) cube complex follows from [88, Section 5.2] and the references therein.

To prove the moreover statement, let Γ be a group that is hyperbolic relative to groups belonging to $\mathfrak{G}$ and let $\mathcal{X}, \mathcal{X}_*$ be two cubulations of Γ. If $P < \Gamma$ is a maximal parabolic subgroup, then by [76, Theorem 1.1] it has convex cores $Z_P \subset \mathcal{X}$ and $(Z_P)_* \subset \mathcal{X}_*$. Also, since P belongs to $\mathfrak{G}$, by Lemma 5.3 the action of P on Z_P is virtually co-special. As this holds for every maximal parabolic subgroup, the action of Γ on $\mathcal{X}$ is virtually co-special either by [41, Theorem A] or [71, Theorem 1.2].

Consider a group $H < \Gamma$ that is convex-cocompact for the action of Γ on $\mathcal{X}$. By Proposition 5.4, H is relatively quasiconvex and $H \cap P$ is a convex-cocompact subgroup for any maximal parabolic subgroup $P < \Gamma$. This implies that $H \cap P < P$ is convex-cocompact for the action of P on Z_P, see, for example, Lemmas 2.14 and 2.15 in [71]. But each such P belongs to $\mathfrak{G}$, so $H \cap P$ is also convex-cocompact for the action of P on $(Z_P)_*$, implying that $H \cap P$ is convex-cocompact for the action of Γ on $\mathcal{X}_*$. By Proposition 5.4 we deduce that H is convex-cocompact for the action of Γ on $\mathcal{X}_*$, so that the actions on $\mathcal{X}$ and $\mathcal{X}_*$ have the same sets of convex-cocompact subgroups. Since any contracting element in a maximal parabolic subgroup is contracting for Γ, we have proven that Γ belongs to $\mathfrak{G}$. □

Example 5.5. Let M be a cusped hyperbolic 3-manifold with cusps $V_1, \dots, V_r \subset M$. We affirm that $\Gamma = \pi_1(M)$ does not belong to $\mathfrak{G}$. For each $i = 1, \dots, r$ choose distinct slopes α_i, β_i for the cusp V_i. Equivalently, each pair α_i, β_i represents a pair of cyclic subgroups (up to commensurability) such that their union generates a finite-index subgroup of $\pi_1(V_i)$. By [24, Corollary 1.3], there exists a cubulation $\mathcal{X}$ of Γ such that each α_i and β_i represents a convex-cocompact subgroup. Since any cubulation of $\mathbb{Z}^2$ has a subgroup that is not convex-cocompact (this follows for instance from [87, Theorem 3.6]), for Γ as above we can produce two cubulations $\mathcal{X}, \mathcal{X}_*$ not satisfying (2) in Definition 1.1.

However, the data of the slopes that are convex-cocompact in cubulations of Γ are the only obstruction for them determining triplets in $\mathfrak{X}$. That is, if $\mathcal{X}, \mathcal{X}_*$ are two cubulations of Γ, then $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ if and only if they have the same pairs of convex-cocompact slopes for each cusp subgroup. The proof of this follows the same lines as the proof of Proposition 5.2.

Example 5.6. There exist groups Γ not necessarily belonging to $\mathfrak{G}$ for which we still can find essentially distinct cubulations $\mathcal{X}, \mathcal{X}_*$ such that $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$.

As an example, let Γ be a finite graph product of finite groups and let $\mathcal{X}$ be the graph-product complex with the standard action by Γ. If $\phi \in \mathrm{Aut}(\Gamma)$ is any automorphism, then we let $\mathcal{X}_*$ be the cubulation obtained by precomposing by ϕ the action of Γ on $\mathcal{X}$. Then, as long as Γ has a contracting element, $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ by combining Theorem D (1) and Theorem 2.15 in [35]. We note that there are many right-angled Coxeter groups with *large* outer automorphism groups [77].

If instead we apply [35, Theorem D (2)], the same conclusion holds for Γ any Coxeter group with a contracting element, with $\mathcal{X}$ being its Niblo-Reeves cubulation, and $\mathcal{X}_*$ obtained from $\mathcal{X}$ after twisting by an automorphism of Γ.

5.2. Constructing the appropriate automaton

In this subsection we construct a finite-state automaton for a triplet $(\Gamma, \mathcal{X}, \mathcal{X}_*)$ in $\mathfrak{X}$, our main result being Theorem 5.11. This is the automaton we will use to prove our main results for the length functions $\ell^{\mathfrak{w}}_{\mathcal{X}}, \ell^{\mathfrak{w}_*}_{\mathcal{X}_*}$ in Section 6, and it implies most of Theorem 1.4. This automaton can be thought of as a cubical analog of the automaton for pairs of word metrics on hyperbolic groups constructed by Calegari-Fujiwara in [14, Lemma 3.8] and explained in Example 2.4.

Our starting point in the construction is the automaton for special cube complexes highlighted in Example 2.5. To use this automaton, our first step is the construction of a (virtually co-special)

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

cubulation for Γ that simultaneously encodes the actions on $\mathcal{X}$ and $\mathcal{X}_*$. This is the content of the next proposition.

Proposition 5.7. *Given* $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ *there exists a proper, cocompact, and essential action of* Γ *on a* CAT(0) *cube complex* $\mathcal{Z}$*,* Γ*-invariant subsets* $\mathbb{W}, \mathbb{W}_* \subset \mathbb{H}(\mathcal{Z})$ *such that* $\mathbb{W} \cup \mathbb{W}_* = \mathbb{H}(\mathcal{Z})$*, and a finite index subgroup* $\overline{\Gamma} < \Gamma$ *satisfying the following.*

- *Let* $\hat{\mathcal{X}} = \mathcal{Z}(\mathbb{W})$ *and* $\hat{\mathcal{X}}_* = \mathcal{Z}(\mathbb{W}_*)$*. Then* $\hat{\mathcal{X}}$ *and* $\hat{\mathcal{X}}_*$ *embed* Γ*-equivariantly as convex subcomplexes of* $\mathcal{X}$ *and* $\mathcal{X}_*$ *respectively. In particular, the triplet* $(\Gamma, \hat{\mathcal{X}}, \hat{\mathcal{X}}_*)$ *belongs to* $\mathfrak{X}$ *and we have the equalities* $\ell_{\mathcal{X}} = \ell_{\hat{\mathcal{X}}}$ *and* $\ell_{\mathcal{X}_*} = \ell_{\hat{\mathcal{X}}_*}$.
- *The action of* $\overline{\Gamma}$ *on both* $\mathcal{Z}$ *and* $\hat{\mathcal{X}}$ *is co-special.*

Remark 5.8. As we will see in the proof, the first conclusion of the proposition above only uses Item (2) in Definition 5.1. Item (1) is only used to find the finite index subgroup $\overline{\Gamma}$ satisfying the second conclusion.

The proof of Proposition 5.7 uses the formalism of *median algebras*, from which refer the reader to [34, Subsection 2.1.5]. We require the following lemma, which is a slight generalization of the implication (3) $\Rightarrow$ (1) in [34, Proposition 7.9].

Lemma 5.9. *Let* Γ *act on the* CAT(0) *cube complexes* $\mathcal{X}, \mathcal{X}_*$ *so that action on* $\mathcal{X}$ *is proper and cocompact, and the action on* $\mathcal{X}_*$ *is essential and has only finitely many orbits of hyperplanes. If every hyperplane stabilizer for the action of* Γ *on* $\mathcal{X}_*$ *is convex-cocompact for the action on* $\mathcal{X}$*, then for any finite subset* $F \subset \mathcal{X}^0 \times \mathcal{X}_*^0$ *the median algebra generated by the* Γ*-translates of* F *is* Γ*-cofinite.*

If Γ acts on a CAT(0) cube complex and $\mathfrak{h}$ is a hyperplane, recall that $\Gamma_{\mathfrak{h}}$ denotes the hyperplane stabilizer of $\mathfrak{h}$.

Proof. Without loss of generality, suppose that $F = P \times P_*$ for P, P_* some finite sets of vertices. The main idea in the proof of (3) $\Rightarrow$ (1) in [34, Proposition 7.9] is the construction, for each hyperplane $\mathfrak{h} \in \mathbb{H}(\mathcal{X}_*)$ with halfspaces $\mathfrak{h}^+$ and $\mathfrak{h}^-$, of a partition $C(\mathfrak{h}^-) \sqcup C(\mathfrak{h}) \sqcup C(\mathfrak{h}^+)$ of $\mathcal{X}$ such that:

- $C(\mathfrak{h})$ is a $\Gamma_{\mathfrak{h}}$-invariant convex subcomplex of $\mathcal{X}$ and the action of $\Gamma_{\mathfrak{h}}$ on $C(\mathfrak{h})$ is cocompact;
- $C(\mathfrak{h}^+)$ and $C(\mathfrak{h}^-)$ are $\Gamma_{\mathfrak{h}}$-invariant unions of connected components of $\mathcal{X} \backslash C(\mathfrak{h})$;
- if $g \in \Gamma$, $y \in P_*$ and $gy \in \mathfrak{h}^+$, then $gx \in C(\mathfrak{h}^+) \cup C(\mathfrak{h})$ for all $x \in P$;
- if $g \in \Gamma$, $y \in P_*$ and $gy \in \mathfrak{h}^-$, then $gx \in C(\mathfrak{h}^-) \cup C(\mathfrak{h})$ for all $x \in P$; and,
- for any $g \in \Gamma$ we have $C(g\mathfrak{h}) = gC(\mathfrak{h})$ and $C(g\mathfrak{h}^+) = gC(\mathfrak{h}^+)$.

Let $M \subset \mathcal{X}^0 \times \mathcal{X}_*^0$ be the median algebra generated by F, so that any wall of M is induced by a hyperplane in $\mathcal{X}$ or $\mathcal{X}_*$. Consider two transverse walls $\mathfrak{v}, \mathfrak{w}$ of M with $\mathfrak{v}$ induced by $\mathfrak{h} \in \mathbb{H}(\mathcal{X}_*)$ and $\mathfrak{w}$ induced by $\mathfrak{l} \in \mathbb{H}(\mathcal{X}) \sqcup \mathbb{H}(\mathcal{X}_*)$. As in the proof of (3) $\Rightarrow$ (1) in [34, Proposition 7.9], we can verify that $\mathfrak{l} \in \mathbb{H}(\mathcal{X})$ implies $C(\mathfrak{h}) \cap \mathfrak{l} \neq \emptyset$, and $\mathfrak{l} \in \mathbb{H}(\mathcal{X}_*)$ implies $C(\mathfrak{h}) \cap C(\mathfrak{l}) \neq \emptyset$. Then we can argue as in the proof of [34, Proposition 7.9] to conclude that M is the 0-skeleton of a CAT(0) cube complex on which the induced action of Γ is cocompact, implying that M is Γ-cofinite.

On the other hand, the construction of the partitions $C(\mathfrak{h}^-) \sqcup C(\mathfrak{h}) \sqcup C(\mathfrak{h}^+)$ is possible by the following slight generalization of [34, Lemma 7.11]: if $H < \Gamma$ is convex-cocompact for the action on $\mathcal{X}$ and $A \subset \Gamma$ is an *H-almost invariant set* (see [34, Lemma 7.11]), then there exists a partition $\mathcal{X} = C_- \sqcup C_0 \sqcup C_+$ such that:

- C_0 is an H-invariant convex subcomplex and the action of H on C_0 is cocompact;
- C_- and C_+ are H-invariant unions of connected components of $X \backslash C_0$; and,
- $A \cdot P \subset C_0 \cup C_+$ and $(\Gamma \backslash A) \cdot P \subset C_0 \cup C_-$.

The proof of this generalization follows from the expected modifications of the proof of [34, Lemma 7.11] and is left to the reader. From this, for a halfspace $\mathfrak{h}^+$ of $\mathcal{X}_*$ with bounding hyperplane $\mathfrak{h}$ in a complete

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

set of representatives of Γ-orbits of hyperplanes, we consider $(C(\mathfrak{h}^+), C(\mathfrak{h}), C(\mathfrak{h}^-)) = (C_+, C_0, C_-)$ for $A = \{g \in \Gamma\colon gy \in \mathfrak{h}^+ \text{ for some } y \in P_*\}$. We note that A is a $\Gamma_{\mathfrak{h}}$-almost invariant set by [34, Remark 7.12]. The proof of the lemma concludes after extending this construction Γ-equivariantly. □

Proof of Proposition 5.7. By [20, Proposition 3.5], let $\hat{\mathcal{X}}$ and $\hat{\mathcal{X}}_*$ be the Γ-essential cores of $\mathcal{X}$ and $\mathcal{X}_*$ respectively, which Γ-equivariantly embed in $\mathcal{X}$ and $\mathcal{X}_*$ as convex subcomplexes.

We claim that there exists a cubulation $\mathcal{Z}'$ of Γ and Γ-invariant subsets $\mathbb{W}, \mathbb{W}_* \subset \mathbb{H}(\mathcal{Z}')$ so that $\hat{\mathcal{X}}$ and $\hat{\mathcal{X}}_*$ are Γ-equivariantly isometric to $\mathcal{Z}'(\mathbb{W})$ and $\mathcal{Z}'(\mathbb{W}_*)$. Under the additional assumption that the action on $\mathcal{X}_*$ is proper, this is the content of the implication (3) $\Rightarrow$ (4) in [35, Theorem 2.17], so we now explain how to use Lemma 5.9 the prove the general case.

Let $P \subset \hat{\mathcal{X}}^0$ and $P_* \subset \hat{\mathcal{X}}_*^0$ be the vertex sets of compact connected subcomplexes K, K_* such that $\Gamma \cdot K = \hat{\mathcal{X}}^0$ and $\Gamma \cdot K_* = \hat{\mathcal{X}}_*^0$. By Lemma 5.9, the median algebra $M \subset \hat{\mathcal{X}}^0 \times \hat{\mathcal{X}}_*^0$ generated by the Γ-translates of $P \times P_*$ is Γ-cofinite. By Chepoi–Roller duality [23, 72], M is the 0-skeleton of a CAT(0) cube complex $\mathcal{Z}'$ equipped with a proper and cocompact cubical action of Γ. The restriction to M of the natural projection $\hat{\mathcal{X}}^0 \times \hat{\mathcal{X}}_*^0 \to \hat{\mathcal{X}}^0$ then induces a Γ-equivariant cubical map $\mathcal{Z}' \to \hat{\mathcal{X}}$ that is surjective because M contains $K \times K_*$. This map is also a *median morphism* on the 0-skeleton, hence a restriction quotient by the discussion preceding the proof of Theorem 2.17 in [35]. The same argument gives us a restriction quotient $\mathcal{Z}' \to \hat{\mathcal{X}}_*$.

To finish the proof of the first assertion, note that the action of Γ on $\mathcal{Z}'$ may not be essential, so instead we consider the projection quotient $\mathcal{Z} = \mathcal{Z}'(\mathbb{W} \cup \mathbb{W}_*)$, which is essential and cocompact since the action of Γ on both $\hat{\mathcal{X}}$ and $\hat{\mathcal{X}}_*$ is essential and the action on $\mathcal{Z}'$ is cocompact. The complex $\mathcal{Z}$ still Γ-equivariantly projects onto $\hat{\mathcal{X}}$ and $\hat{\mathcal{X}}_*$, so the action of Γ on $\mathcal{Z}$ is proper because the action of Γ on $\hat{\mathcal{X}}$ is proper.

To prove the second assertion, note that by [35, Theorem 2.17] the actions of Γ on $\hat{\mathcal{X}}$ and $\mathcal{Z}$ have the same sets of convex-cocompact subgroups, so any hyperplane stabilizer for the action of Γ on $\mathcal{Z}$ will be convex-cocompact with respect to $\hat{\mathcal{X}}$. But the action of Γ on $\hat{\mathcal{X}}$ is virtually co-special (it is a convex core for Γ acting on $\mathcal{X}$), so $\mathcal{Z}$ is virtually co-special by Lemma 5.3. Finally, co-specialness is preserved under taking finite-index subgroups, so we can choose $\overline{\Gamma}$ so that both quotients $\overline{\Gamma}\backslash\mathcal{Z}$ and $\overline{\Gamma}\backslash\hat{\mathcal{X}}$ are special. □

By virtue of the proposition above, throughout the rest of the section we will work under the following convention.

Convention 5.10. Γ *is a nonvirtually cyclic group acting properly, cocompactly and essentially on the* CAT(0) *cube complex* $\mathcal{Z}$*, and* $\mathbb{W} \subset \mathbb{H}(\mathcal{Z})$ *is a* Γ*-invariant subset such that the action of* Γ *on* $\mathcal{X} = \mathcal{Z}(\mathbb{W})$ *is proper, cocompact and essential. Let* $\phi : \mathcal{Z} \to \mathcal{X}$ *be the restriction quotient.*

Also, let $\overline{\Gamma} < \Gamma$ *be a finite index subgroup such that the quotients* $\overline{\mathcal{Z}} = \overline{\Gamma}\backslash\mathcal{Z}$ *and* $\overline{\mathcal{X}} = \overline{\Gamma}\backslash\mathcal{X}$ *are special cube complexes. We fix a base vertex* $\widetilde{o} \in \mathcal{Z}$ *and set* $o = \phi(\widetilde{o}) \in \mathcal{X}$.

Let $S_{\overline{\mathcal{Z}}}$ and $S_{\overline{\mathcal{X}}}$ be the set of oriented hyperplanes in $\overline{\mathcal{Z}}$ and $\overline{\mathcal{X}}$ respectively. By specialness all the hyperplanes in $\overline{\mathcal{Z}}$ and $\overline{\mathcal{X}}$ are 2-sided, so there exist two orientations for each hyperplane. Since each hyperplane in $S_{\overline{\mathcal{X}}}$ corresponds to a $\overline{\Gamma}$-orbit of oriented hyperplanes in $\mathbb{W} \subset \mathbb{H}(\mathcal{Z})$, there is a natural injection of $S_{\overline{\mathcal{X}}}$ into $S_{\overline{\mathcal{Z}}}$, so often we will consider $S_{\overline{\mathcal{X}}}$ as a subset of $S_{\overline{\mathcal{Z}}}$. The *label* of an oriented hyperplane in $\mathcal{Z}$ (resp. $\mathcal{X}$) is its projection in $\overline{\mathcal{Z}}$ (resp. $\overline{\mathcal{X}}$).

We say that a word $w = \mathfrak{h}_1 \cdots \mathfrak{h}_n$ in $(S_{\overline{\mathcal{Z}}})^*$ is *represented* by a (combinatorial) path $\gamma = (\gamma_0, \dots, \gamma_n)$ in $\mathcal{Z}$ if for each i the oriented hyperplane $\mathfrak{h}_i$ is the image in $\overline{\mathcal{Z}}$ of the oriented hyperplane dual to the edge from γ_{i-1} to γ_i. Similarly, we define when a word in $(S_{\overline{\mathcal{X}}})^*$ is represented by a path in $\mathcal{X}$. A consequence of specialness is that if a word is represented by two paths with the same initial vertex, then the paths must coincide.

The next theorem gives us a finite-state automaton over $S_{\overline{\mathcal{X}}}$ that keeps track of the action of $\overline{\Gamma}$ on both $\mathcal{X}$ and $\mathcal{Z}$.

Theorem 5.11. *Let* $\Gamma, \mathcal{Z}$ *and* $\mathcal{X}$ *satisfy Convention* 5.10*. Then there exists a language* $L = L_{\overline{\Gamma},\phi} \subset (S_{\overline{\mathcal{X}}})^*$ *parametrized by the pruned finite-state automaton*

$$\mathcal{A}_{\overline{\Gamma},\phi} = (\mathcal{G}_\phi = (V_\phi, E_\phi), \pi_\phi, I_\phi, V_\phi)$$

satisfying the following.

(1) *There exists* $C \geq 1$ *depending only on* $\overline{\Gamma}$ *and* $\phi : \mathcal{Z} \to \mathcal{X}$ *such that any* $w \in L_{\overline{\Gamma},\phi}$ *is represented by at most* C *paths in* $\mathcal{G}_\phi$ *starting at an initial state.*
(2) *Every* $w \in L_{\overline{\Gamma},\phi}$ *is represented by a unique combinatorial geodesic* $\gamma_w \subset \mathcal{X}$ *starting at the vertex* o*. We let* $\tau_{\mathcal{X}}(w)$ *denote the final vertex of* γ_w*.*
(3) *The map* $\tau_{\mathcal{X}} : L_{\overline{\Gamma},\phi} \to \mathcal{X}^0$ *is a bijection.*

Moreover, there exist maps $\Psi : V_\phi \to (S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}})^*$ *and* $\Xi : V_\phi \to \overline{\mathcal{X}}^0$ *satisfying the following.*

(4) *If* $w \in L_{\overline{\Gamma},\phi}$ *is represented by the path* $\omega = (v_0 \xrightarrow{e_1} v_1 \cdots \xrightarrow{e_n} v_n)$ *in* $\mathcal{G}_\phi$ *starting at an initial state, then the concatenation*

$$\alpha(\omega) := \Psi(v_0)\pi_\phi(e_1)\Psi(v_1)\cdots\pi_\phi(e_n) \in (S_{\overline{\mathcal{Z}}})^* \tag{5.1}$$

can be represented by a unique geodesic path $\widetilde{\gamma}_{\alpha(\omega)}$ *in* $\mathcal{Z}$ *with initial vertex* $\widetilde{o}$ *and final vertex* $\tau_{\mathcal{Z}}(\alpha(\omega))$*, so that* $\phi(\tau_{\mathcal{Z}}(\alpha(\omega))) = \tau_{\mathcal{X}}(w)$*.*
(5) *If* $w \in L_{\overline{\Gamma},\phi}$ *is represented by the path* $\omega = (v_0 \to \cdots \to v_n)$ *in* $\mathcal{G}_\phi$ *starting at an initial state, then the path* $\gamma_w = (\gamma_0, \ldots, \gamma_n)$ *in* $\mathcal{X}$ *projects to* $(\Xi(v_0), \ldots, \Xi(v_n))$ *in* $\overline{\mathcal{X}}$*.*

Recall from Subsection 2.3 that an automaton with underlying graph $\mathcal{G}$ is deterministic if any two edges of $\mathcal{G}$ with the same initial vertex have different labels, and that the automaton is pruned if any vertex in $\mathcal{G}$ is the final vertex of an admissible path.

Remark 5.12. i) The notation $L_{\overline{\Gamma},\phi}$ and $\mathcal{A}_{\Gamma,\phi}$ in the theorem above is chosen to emphasize the dependence on $\overline{\Gamma}, \mathcal{X}$ and $\mathcal{Z}$, but also on the restriction quotient map $\phi : \mathcal{Z} \to \mathcal{X}$. We have suppressed some notation for simplicity, but this is also the case for the data involved in the definition of $\mathcal{A}_{\overline{\Gamma},\phi}$, as well as for $\tau_{\mathcal{X}}, \tau_{\mathcal{Z}}, \Psi, \Xi$ and α. For simplicity, often we will use the simplified notation $L_\phi = L_{\overline{\Gamma},\phi}$ and $\mathcal{A}_\phi = \mathcal{A}_{\overline{\Gamma},\phi}$. ii) The specialness of $\overline{\mathcal{X}}$ and $\overline{\mathcal{Z}}$ is used at several steps in the construction of the automaton $\mathcal{A}_{\overline{\Gamma},\phi}$. Crucially, specialness of $\overline{\mathcal{Z}}$ allows us to use the automaton constructed by Li and Wise in [58] (see Example 2.5). This automaton does not remember the quotient $\Gamma\backslash\mathcal{Z}$ covered by $\overline{\mathcal{Z}}$, and in particular, the automaton $\mathcal{A}_{\overline{\Gamma},\phi}$ does not remember Γ. iii). We note that Convention 5.10, and hence the construction of $\mathcal{A}_{\overline{\Gamma},\phi}$ does not assume that Γ has a contracting element. The full strength of Item (3) in Definition 5.1 is used in Section 6 when we prove our counting theorems (compare with Convention 6.3). For these counting results, passing to the finite index subgroup $\overline{\Gamma}$ is not an issue, as it suffices for the automaton to see a subset of Γ having positive *lower density*, see Lemma 6.11.

We start the construction of $\mathcal{A}_{\overline{\Gamma},\phi}$ by considering a regular language $L_{\overline{\mathcal{Z}}} \subset (S_{\overline{\mathcal{Z}}})^*$ parameterized by the (pruned and deterministic) automaton

$$\mathcal{A}_{\overline{\mathcal{Z}}} = (\mathcal{G}_{\overline{\mathcal{Z}}} = (V_{\overline{\mathcal{Z}}}, E_{\overline{\mathcal{Z}}}), \pi_{\overline{\mathcal{Z}}}, \{*_{\overline{\mathcal{Z}}}\}, V_{\overline{\mathcal{Z}}})$$

over $S_{\overline{\mathcal{Z}}}$, constructed by Li and Wise in [58] and discussed in Example 2.5. This automaton satisfies:

- every $\widetilde{w} \in L_{\overline{\mathcal{Z}}}$ is represented by a unique geodesic path $\widetilde{\gamma}_{\widetilde{w}}$ in $\mathcal{Z}$ starting at $\widetilde{o}$ and ending at the vertex $\tau_{\mathcal{Z}}(\widetilde{w})$; and,
- the map $\tau_{\overline{\mathcal{Z}}} : L_{\overline{\mathcal{Z}}} \to \mathcal{Z}^0$ is a bijection.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

We use the automaton $\mathcal{A}_{\overline{\mathcal{Z}}}$ to produce a language $L_\phi = L_{\overline{\Gamma},\phi}$ over the alphabet $S_{\overline{\mathcal{X}}}$. We do this by first constructing an automaton

$$\hat{\mathcal{A}}_\phi = (\hat{\mathcal{G}}_\phi = (\hat{V}_\phi, \hat{E}_\phi), \hat{\pi}_\phi, \hat{I}_\phi, \hat{V}_\phi)$$

as follows.

Definition 5.13. Let $\hat{V}_\phi$ be the set of finite directed paths (possibly of length 0) in $\mathcal{G}_{\overline{\mathcal{Z}}}$ of the form $\omega = (v_0 \xrightarrow{e_1} \cdots \xrightarrow{e_n} v_n)$ and satisfying:

- $\pi_{\overline{\mathcal{Z}}}(e_i) \in S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}}$ for every $1 \leq i \leq n$;
- either $v_0 = *_{\overline{\mathcal{Z}}}$ or there exists an edge in $E_{\overline{\mathcal{Z}}}$ with label in $S_{\overline{\mathcal{X}}}$ and final vertex v_0; and,
- either there exists an edge in $E_{\overline{\mathcal{Z}}}$ with label in $S_{\overline{\mathcal{X}}}$ and initial vertex v_n or there are no edges in $E_{\overline{\mathcal{Z}}}$ with initial vertex v_n.

We consider an edge $\hat{e}$ from the vertex ω to the vertex ω' in $\hat{V}_\phi$ if there exists an edge $e \in E_{\overline{\mathcal{Z}}}$ with $\pi_{\overline{\mathcal{Z}}}(e) \in S_{\overline{\mathcal{X}}}$ and such that the concatenation $\omega e \omega'$ is a path in $\mathcal{G}_{\overline{\mathcal{Z}}}$. We define $\hat{\pi}_\phi(\hat{e}) := \pi_{\overline{\mathcal{Z}}}(e)$ and let $\hat{E}_\phi$ be the set of all of the edges defined in this way. Finally, a vertex of $\hat{V}_\phi$ is an initial state if its initial vertex (as a path in $\mathcal{G}_{\overline{\mathcal{Z}}}$) is $*_{\overline{\mathcal{Z}}}$. We let $\hat{I}_\phi$ be the set of all the initial states.

Lemma 5.14. *The set $\hat{V}_\phi$ is finite and nonempty. Therefore, $\hat{\mathcal{A}}_\phi$ defines a pruned finite-state automaton over $S_{\overline{\mathcal{X}}}$.*

Proof. The set $S_{\overline{\mathcal{X}}}$ is nonempty because Γ is nonelementary and $\mathcal{X}$ is essential, and since $\mathcal{A}_{\overline{\mathcal{Z}}}$ is pruned we can find a vertex in $\hat{I}_\phi$, so that $\hat{V}_\phi$ is nonempty.

To show finiteness, let M be the maximum cardinality of a preimage $\phi^{-1}(x) \cap \mathcal{Z}^0$ among $x \in \mathcal{X}^0$, which is finite since ϕ is Γ-invariant and the action of Γ on $\mathcal{Z}$ is cocompact. If $\omega = (v_0 \xrightarrow{e_1} \cdots \xrightarrow{e_n} v_n)$ is a vertex in $\hat{V}_\phi$, the fact that $\mathcal{A}_{\overline{\mathcal{Z}}}$ is pruned implies the existence of a geodesic path $\widetilde{\gamma}$ in $\mathcal{Z}$ representing the word $\pi_{\overline{\mathcal{Z}}}(e_1) \cdots \pi_{\overline{\mathcal{Z}}}(e_n) \in (S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}})^*$. Since ϕ collapses the hyperplanes not belonging to $\mathbb{W}$, the image $\phi(\widetilde{\gamma})$ consists of a single point, implying that $n + 1 \leq M$. We conclude that every vertex in $\hat{V}_\phi$ has uniformly bounded length as a path in $\mathcal{G}_{\overline{\mathcal{Z}}}$, so $\hat{V}_\phi$ is finite because $\mathcal{G}_{\overline{\mathcal{Z}}}$ is.

Finally, $\mathcal{A}_{\overline{\mathcal{Z}}}$ being pruned implies that $\hat{\mathcal{A}}_\phi$ is pruned, and by construction this automaton is over the alphabet $S_{\overline{\mathcal{X}}}$. □

Definition 5.15. We let $L_\phi = L_{\overline{\Gamma},\phi} \subset (S_{\overline{\mathcal{X}}})^*$ be the language parametrized by $\hat{\mathcal{A}}_\phi$.

If $\hat{\omega} = (\omega_0 \xrightarrow{\hat{e}_1} \cdots \xrightarrow{\hat{e}_n} \omega_n)$ is a path in $\hat{\mathcal{G}}_\phi$, then the concatenation

$$c(\hat{\omega}) := \omega_0 e_1 \cdots \omega_{n-1} e_n \tag{5.2}$$

is a path in $\mathcal{G}_{\overline{\mathcal{Z}}}$. Let $\hat{\alpha}(\hat{\omega}) \in (S_{\overline{\mathcal{Z}}})^*$ be the word represented by $c(\hat{\omega})$.

Lemma 5.16.

(1) *If $\hat{\omega}, \hat{\omega}'$ are paths in $\hat{\mathcal{G}}_\phi$ starting at an initial state and representing the same word $w \in L_\phi$, then $\phi(\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}))) = \phi(\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}'))) \in \mathcal{X}^0$. We denote this vertex by $\tau_{\mathcal{X}}(w)$.*
(2) *Any w in L_ϕ is represented by a unique geodesic path γ_w in $\mathcal{X}$ starting at o and ending at $\tau_{\mathcal{X}}(w)$.*
(3) *The map $\tau_{\mathcal{X}} : L_\phi \to \mathcal{X}^0$ is a bijection.*

Proof. By induction on the length of w we will prove simultaneously assertion (1) and that $d_{\mathcal{X}}(o, \tau_{\mathcal{X}}(w))$ equals the length of w. Suppose that w has length n and is represented by the paths

$$\hat{\omega} = (\omega_0 \xrightarrow{\hat{e}_1} \cdots \xrightarrow{\hat{e}_n} \omega_n) \text{ and } \hat{\omega}' = (\omega_0' \xrightarrow{\hat{e}_1'} \cdots \xrightarrow{\hat{e}_n'} \omega_n')$$

in $\hat{\mathcal{G}}_\phi$. If $n = 0$ then $\hat{\alpha}(\hat{\omega}) = \omega_0$ has no letters in $S_{\overline{\mathcal{X}}}$, and hence the projection of $\widetilde{\gamma}_{\hat{\alpha}(\hat{\omega})}$ under ϕ consists of a single vertex, which must be o. As the same happens for ω', this solves the base case.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

If $n \geq 1$ we consider $\hat{\xi} = (\omega_0 \xrightarrow{\hat{e}_1} \cdots \xrightarrow{\hat{e}_{n-1}} \omega_{n-1})$ and $\hat{\xi}' = (\omega_0' \xrightarrow{\hat{e}_1'} \cdots \xrightarrow{\hat{e}_{n-1}'} \omega_{n-1}')$. Our inductive assumption implies that $\phi(\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\xi}))) = \phi(\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\xi}'))) =: x$, so that no hyperplane with label in $S_{\overline{\mathcal{X}}}$ separates $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\xi}))$ and $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\xi}'))$. Since $\overline{\mathcal{X}}$ is special there exists at most one edge in $\mathcal{X}$ with initial vertex x and dual to a hyperplane with label $\mathfrak{h} = \hat{\pi}_\phi(\hat{e}_n) = \hat{\pi}_\phi(\hat{e}_n') \in S_{\overline{\mathcal{X}}}$. But ϕ is injective on $\mathbb{W}$, so that the hyperplane labeled $\mathfrak{h}$ that separates $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\xi}))$ and $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}))$ in $\mathcal{Z}$ is the same as the hyperplane labeled $\mathfrak{h}$ that separates $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\xi}'))$ and $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}'))$. Since this is the only hyperplane with a label in $S_{\overline{\mathcal{X}}}$ that separates these pairs, we conclude that every hyperplane separating $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}))$ and $\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}'))$ has a label in $S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}}$, which gives us $\phi(\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}))) = \phi(\tau_{\mathcal{Z}}(\hat{\alpha}(\hat{\omega}')))$. Moreover, if $w' \in L_\phi$ is the word represented by $\hat{\xi}$, then by induction we have $d_{\mathcal{X}}(o, \tau_{\mathcal{X}}(w')) = n-1$ and $d_{\mathcal{X}}(\tau_{\mathcal{X}}(w'), \tau_{\mathcal{X}}(\xi)) = 1$. But all these points belong to the projection under ϕ of the geodesic $\widetilde{\gamma}_{\hat{\alpha}(\hat{\omega})}$, so Remark 2.2 implies that $d_{\mathcal{X}}(o, \tau_{\mathcal{X}}(w)) = n$ and concludes the proof by induction, proving (1).

It is not hard to see that if $w = \mathfrak{h}_1 \cdots \mathfrak{h}_n \in L_\phi$ then

$$\gamma_w := (o, \tau_{\mathcal{X}}(\mathfrak{h}_1), \tau_{\mathcal{X}}(\mathfrak{h}_1\mathfrak{h}_2), \ldots, \tau_{\mathcal{X}}(\mathfrak{h}_1 \cdots \mathfrak{h}_n))$$

is the unique geodesic representing w in $\mathcal{X}$ and starting at o, which settles (2).

To prove (3), injectivity can be deduced by induction on the length of words in L_ϕ combined with the fact that no two distinct edges with the same initial vertex in $\mathcal{X}$ are dual to hyperplanes with the same label in $S_{\overline{\mathcal{X}}}$. This last statement is true by specialness of $\overline{\mathcal{X}}$.

To prove surjectivity, let $x \in \mathcal{X}^0$ and consider $\widetilde{w} \in L_{\overline{\mathcal{Z}}}$ such that $\phi(\tau_{\mathcal{Z}}(\widetilde{w})) = x$. Such an $\widetilde{w}$ exists because both $\tau_{\mathcal{Z}}$ and ϕ are surjective. We write $\widetilde{w} = w_0 e_1 \cdots e_n w_n$, where each e_i is a letter in $S_{\overline{\mathcal{X}}}$ and each w_i is a (possibly empty) word in $(S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}})^*$. Then $\widetilde{w}' := w_0 e_1 \cdots w_{n-1} e_n$ equals $\hat{\alpha}(\hat{\omega})$ for some path $\hat{\omega}$ in $\hat{\mathcal{G}}_\phi$ representing the word $w \in L_\phi$, for which $\tau_{\mathcal{X}}(w) = \phi(\tau_{\mathcal{Z}}(\widetilde{w}')) = \phi(\tau_{\mathcal{Z}}(\widetilde{w})) = x$. □

Lemma 5.17. *There exists $C \geq 1$ such that every $w \in L_\phi$ is represented by at most C paths in $\hat{\mathcal{G}}_\phi$ starting at an initial vertex.*

Proof. Since $\tau_{\mathcal{Z}}$ is injective and ϕ is uniformly finite-to-one when restricted to vertices, it is enough to prove that the assignment $\hat{\omega} \mapsto \hat{\alpha}(\hat{\omega})$ from the paths in $\hat{\mathcal{G}}_\phi$ starting at an initial vertex into $L_{\overline{\mathcal{Z}}}$ is uniformly finite-to one. To show this, note that such $\hat{\omega}$ is completely determined by its concatenation $c(\hat{\omega})$ in $\mathcal{G}_{\overline{\mathcal{Z}}}$ (defined in (5.2)) and its final vertex in $\hat{V}_\phi$, and that $c(\hat{\omega})$ is completely determined by $\hat{\alpha}(\hat{\omega})$ since $\mathcal{A}_{\overline{\mathcal{Z}}}$ is deterministic. The lemma then follows from Lemma 5.14. □

Proof of Theorem 5.11. First we note that $\overline{\mathcal{X}}^1$ can be seen as the finite-state automaton

$$\overline{\mathcal{X}}^1 = ((\overline{\mathcal{X}}^0, E(\overline{\mathcal{X}})), \mathrm{pr}_{\overline{\mathcal{X}}}, \{\overline{o}\}, \overline{\mathcal{X}}^0)$$

over $S_{\overline{\mathcal{X}}}$, where $E(\overline{\mathcal{X}})$ is the set of oriented edges of $\overline{\mathcal{X}}$, $\mathrm{pr}_{\overline{\mathcal{X}}}$ labels each directed edge of $\overline{\mathcal{X}}$ with its corresponding dual oriented hyperplane, and $\overline{o}$ is the image of o under the quotient $\mathcal{X} \to \overline{\mathcal{X}}$. This automaton is deterministic since $\overline{\mathcal{X}}$ is special.

Let $\mathcal{A}_\phi = \mathcal{A}_{\overline{\Gamma}, \phi} = (\mathcal{G}_\phi = (V_\phi, E_\phi), \pi_\phi, I_\phi, V_\phi)$ be the fiber product of $\hat{\mathcal{A}}_\phi$ and $\overline{\mathcal{X}}^1$. That is, in $\hat{V}_\phi \times \overline{\mathcal{X}}^0$ consider a directed edge from $(\hat{\omega}, \overline{x})$ to $(\hat{\omega}', \overline{x}')$ if there exists an edge $\hat{e}$ from ω to ω' in $\hat{\mathcal{G}}_\phi$ such that $\hat{\pi}_\phi(\hat{e})$ is the oriented hyperplane dual to an edge from $\overline{x}$ to $\overline{x}'$ in $\overline{\mathcal{X}}$ (so that $\overline{x}, \overline{x}'$ must be adjacent). By abuse of notation we will call this edge e and define $\pi_\phi(e) := \hat{\pi}_\phi(\hat{e})$. Let $I_\phi = \hat{I}_\phi \times \{\overline{o}\}$ be the set of initial states of $\mathcal{G}_\phi$ and let $V_\phi \subset \hat{V}_\phi \times \overline{\mathcal{X}}^0$ be the set of all the vertices in some directed path in $\hat{V}_\phi \times \overline{\mathcal{X}}^0$ starting at an initial state. Let E_ϕ be the set of all the directed edges between vertices in V_ϕ as defined above. Clearly $\mathcal{A}_\phi$ is pruned and finite.

There exists a label-preserving map $\mathfrak{p}$ from the set of paths in $\mathcal{G}_\phi$ starting at an initial state into the set of paths in $\hat{\mathcal{G}}_\phi$ starting at an initial state, which sends the path $((\omega_0, \overline{x}_0) \xrightarrow{e_1} \cdots \xrightarrow{e_n} (\omega_n, \overline{x}_n))$ to the

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

path $\hat{\omega} = (\omega_0 \xrightarrow{\hat{e}_1} \cdots \xrightarrow{\hat{e}_n} \omega_n)$. The map $\mathfrak{p}$ is a bijection since the sequence $\overline{x}_0, \ldots, \overline{x}_n$ is the image of γ_w under $\mathcal{X} \to \overline{\mathcal{X}}$, where w is the word represented by $\hat{\omega}$. In particular, the language parametrized by $\mathcal{A}_\phi$ is precisely L_ϕ, so Lemma 5.17 implies Item (1). Also, Items (2) and (3) follow from Lemma 5.16.

For Item (4) we consider the map $\Psi : V_\phi \to (S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}})^*$ that sends the vertex $(\omega, \overline{x})$ to the word represented by ω, seen as a path in $\mathcal{G}_{\overline{\mathcal{Z}}}$. From the definition of $\hat{\mathcal{G}}_\phi$ it is clear that for the path $\omega = (v_0 \xrightarrow{e_1} \cdots \xrightarrow{e_n} v_n)$ in $\mathcal{G}_\phi$ starting at an initial state, the concatenation

$$\alpha(\omega) := \hat{\alpha}(\mathfrak{p}(\omega)) = \Psi(v_0)\pi_\phi(e_1)\Psi(v_1)\cdots\pi_\phi(e_n)$$

belongs to $L_{\overline{\mathcal{Z}}}$, so it is represented by the geodesic path $\widetilde{\gamma}_{\alpha(\omega)}$ in $\mathcal{Z}$ starting at $\widetilde{o}$. The word w represented by ω is also represented by $\mathfrak{p}(\omega)$, so Lemma 5.16 implies that $\tau_{\mathcal{X}}(w) = \phi(\tau_{\mathcal{Z}}(\alpha(\omega)))$. $\mathcal{A}_{\overline{\mathcal{Z}}}$ being deterministic implies that $\widetilde{\gamma}_{\alpha(\omega)}$ is the unique geodesic in $\mathcal{Z}$ starting at $\widetilde{o}$ and representing w.

Finally, we define $\Xi : V_\phi \to \overline{\mathcal{X}}^0$ as the coordinate projection $\Xi(\omega, \overline{x}) = \overline{x}$, and the same argument for the proof that $\mathfrak{p}$ is a bijection implies Item (5). □

6. Proof of the main theorems

In this section we prove our main results about pairs of actions on CAT(0) cube complexes from the introduction. Theorem 1.2 and Theorem 1.5 are consequences of more general statements, given by Theorems 6.1 and 6.2 respectively. The strategy is to use the automaton from Theorem 5.11 to define an appropriate suspension flow, and then use Proposition 6.8 to relate the Manhattan curves with pressure functions for potentials on this suspension (at this point we have enough formalism to deduce Theorem 1.4). This relation will allow us to use the tools from symbolic dynamics and thermodynamic formalism discussed in Section 3 to deduce our main results.

The following are the main theorems of the section, and they are proven in Subsection 6.2. For their statements, we interpret the quantities $\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}$ and $v_{\mathcal{X}^{\mathfrak{w}}}/v_{\mathcal{X}_*^{\mathfrak{w}_*}}$ as zero if the action of Γ on $\mathcal{X}_*$ is not proper.

Theorem 6.1. *Let $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ and let $\mathfrak{w}, \mathfrak{w}_*$ be Γ-invariant orthotope structures on $\mathcal{X}$ and $\mathcal{X}_*$ respectively. Then the Manhattan curve $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}} : \mathbb{R} \to \mathbb{R}$ is convex, decreasing, and analytic. In addition, the following limit exists and equals $-\theta'_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(0)$:*

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) := \lim_{T\to\infty} \frac{1}{\#\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)} \sum_{[g]\in\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)} \frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]}.$$

Moreover, we always have

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) \geq v_{\mathcal{X}^{\mathfrak{w}}}/v_{\mathcal{X}_*^{\mathfrak{w}_*}}.$$

If the action of Γ on $\mathcal{X}_*$ is proper then the following are equivalent:

(1) *$\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}$ is a straight line;*
(2) *there exists $\Lambda > 0$ such that $\ell_{\mathcal{X}}^{\mathfrak{w}}[g] = \Lambda\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]$ for all $[g] \in$ **conj**(Γ); and,*
(3) *$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) = v_{\mathcal{X}^{\mathfrak{w}}}/v_{\mathcal{X}_*^{\mathfrak{w}_*}}$.*

Theorem 6.2. *Let $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$. Then there exists an analytic function*

$$\mathcal{I} : [\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X})] \to \mathbb{R}$$

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

and $C > 0$ *such that for any* $\eta \in (\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X}))$ *we have*

$$0 < \limsup_{T\to\infty} \frac{1}{T} \log\left(\#\left\{[g] \in \mathfrak{C}_{\mathcal{X}}(T) : |\ell_{\mathcal{X}_*}[g] - \eta\ell_{\mathcal{X}}[g]| < \frac{C}{T}\right\}\right) = \mathcal{I}(\eta) \le v_{\mathcal{X}}. \tag{6.1}$$

Furthermore, we have equality in the above inequality if and only if $\eta = \tau(\mathcal{X}_*/\mathcal{X})$.

As we noted in the previous section, $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ whenever $\Gamma \in \mathfrak{G}$ and $\mathcal{X}, \mathcal{X}_*$ are cubulations of Γ. From this we see that Theorems 6.1 and 6.2 imply Theorems 1.2 and 1.5 from the introduction.

The outline to prove these theorems is as follows. Given $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ and corresponding orthotope structures $\mathfrak{w}, \mathfrak{w}_*$, our goal is to relate the Manhattan curve of $(\mathcal{X}^{\mathfrak{w}}, \mathcal{X}_*^{\mathfrak{w}_*})$ to a pressure function for a potential on a suspension flow (Proposition 6.8), which is done in Subsection 6.1. The suspension flow and the potential are constructed using the automaton $\mathcal{A}_{\overline{\Gamma},\phi}$ from Theorem 5.11, and the proof of Proposition 6.8 relies on showing that any closed orbit in the suspension flow has associated a conjugacy class in Γ in an "almost bijective" way (Lemma 6.12 and Lemma 6.13). These two lemmas are the last pieces we need to prove Theorem 1.4. Then in Subsection 6.2, and with Proposition 6.8 at our disposal, we can deduce a standard large deviations principle for the pair $(\mathcal{X}^{\mathfrak{w}}, \mathcal{X}_*^{\mathfrak{w}_*})$ (Corollary 6.16). Combining this with tools from thermodynamic formalism, we prove Theorem 6.1. Finally, we apply Theorem 3.2 to deduce a large deviations theorem with shrinking intervals for $(\mathcal{X}, \mathcal{X}_*)$, which is Theorem 6.2.

6.1. Manhattan curves for pairs of cubulations

In this subsection we use the finite-state automaton given by Theorem 5.11 to describe the Manhattan geodesics for a pair of cubulations in terms of pressure functions. The main result is Proposition 6.8, which will allow us to use thermodynamic formalism to prove our main results in the next subsection. For this section we keep the notation from the previous section and assume the following.

Convention 6.3. *Let* $(\Gamma, \mathcal{X}, \mathcal{Z})$ *be a triplet satisfying Convention 5.10. Consider a nonempty* Γ*-invariant subset* $\mathbb{W}_* \subset \mathbb{H}(\mathcal{Z})$ *such that* $\mathbb{W} \cup \mathbb{W}_* = \mathbb{H}(\mathcal{Z})$ *and set* $\mathcal{X}_* = \mathcal{Z}(\mathbb{W}_*)$*. Then the action of* Γ *on* $\mathcal{X}_*$ *is cocompact, but not necessarily proper, and we further assume that this action is essential. Let* $\mathfrak{w}$ *and* $\mathfrak{w}_*$ *be* Γ*-invariant orthotope structures on* $\mathcal{X}$ *and* $\mathcal{X}_*$ *respectively. We also require the action of* Γ *on* $\mathcal{X}$ *to have a contracting element, so that* $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$.

Let $S_{\overline{\mathcal{X}}_*} \subset S_{\overline{\mathcal{Z}}}$ *be the set of all the oriented hyperplanes in* $\overline{\mathcal{Z}}$ *whose lifts to* $\mathcal{Z}$ *correspond to hyperplanes in* $\mathbb{W}_*$*, and let* $\phi_* : \mathcal{Z} \to \mathcal{X}_*$ *be the projection quotient with* $o_* = \phi_*(\tilde{o})$.

Since the structures $\mathfrak{w}, \mathfrak{w}_*$ are Γ-invariant, there exist natural weighting maps $\overline{\mathfrak{w}}, \overline{\mathfrak{w}}_* : S_{\overline{\mathcal{Z}}} \to \mathbb{R}$ defined according to:

$$\overline{\mathfrak{w}}(\mathfrak{h}) := \begin{cases} \mathfrak{w}(\tilde{\mathfrak{h}}) & \text{if } \mathfrak{h} \in S_{\overline{\mathcal{X}}} \text{ and } \tilde{\mathfrak{h}} \in \mathbb{H}(\mathcal{X}) \text{ is oriented and projects to } \mathfrak{h} \text{ under } \mathcal{X} \to \overline{\mathcal{X}}, \\ 0 & \text{if } \mathfrak{h} \in S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}}, \end{cases}$$

and

$$\overline{\mathfrak{w}}_*(\mathfrak{h}) := \begin{cases} \mathfrak{w}_*(\tilde{\mathfrak{h}}) & \text{if } \mathfrak{h} \in S_{\overline{\mathcal{X}}_*} \text{ and } \tilde{\mathfrak{h}} \in \mathbb{H}(\mathcal{X}) \text{ is oriented and projects to } \mathfrak{h} \text{ under } \mathcal{Z} \to \overline{\mathcal{Z}}, \\ 0 & \text{if } \mathfrak{h} \in S_{\overline{\mathcal{Z}}} \backslash S_{\overline{\mathcal{X}}_*}. \end{cases}$$

By abuse of notation we extend these weightings to $\overline{\mathfrak{w}}, \overline{\mathfrak{w}}_* : (S_{\overline{\mathcal{Z}}})^* \to \mathbb{R}$ by declaring the empty word to have weights 0 and assigning

$$\overline{\mathfrak{w}}(\mathfrak{h}_1 \cdots \mathfrak{h}_n) = \overline{\mathfrak{w}}(\mathfrak{h}_1) + \cdots + \overline{\mathfrak{w}}(\mathfrak{h}_n) \ \text{ and } \ \overline{\mathfrak{w}}_*(\mathfrak{h}_1 \cdots \mathfrak{h}_n) = \overline{\mathfrak{w}}_*(\mathfrak{h}_1) + \cdots + \overline{\mathfrak{w}}_*(\mathfrak{h}_n)$$

for a word $\mathfrak{h}_1 \cdots \mathfrak{h}_n \in (S_{\overline{\mathcal{Z}}})^*$ of positive length.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

We consider the automaton $\mathcal{A}_{\overline{\Gamma},\phi}$ given by Theorem 5.11, and we let $\Sigma^\times$ be the set of all the finite directed paths in the underlying graph $\mathcal{G}_\phi$ of $\mathcal{A}_{\overline{\Gamma},\phi}$. We can see $\Sigma^\times$ as a set of finite sequences of edges in E_ϕ. Similarly, let $\Sigma \subset (E_\phi)^{\mathbb{N}}$ be the set of infinite sequences $(e_i)_{i\geq 1}$ such that $(e_i)_{1\leq i\leq k} \in \Sigma^\times$ for all k. We also let $\sigma : \Sigma \to \Sigma$ denote the shift map $\sigma((e_i)_{i\geq 1}) = (e_{i+1})_{i\geq 1}$. For each n we let $P_n(\Sigma^\times) \subset \Sigma^\times$ be the subset of all the closed paths of length n, and set $P_{\leq n}(\Sigma^\times) = \bigcup_{j\leq n} P_j(\Sigma^\times)$ and $P(\Sigma^\times) = \bigcup_j P_j(\Sigma^\times)$. We will often identify the set $\mathrm{Fix}_n(\Sigma)$ of sequences $\omega \in \Sigma$ satisfying $\sigma^n(\omega) = \omega$ with $P_n(\Sigma^\times)$ via the truncation $\omega = (e_1, \dots, e_n, e_1, \dots) \mapsto t_n(\omega) := (e_1, \dots, e_n)$.

The next definition will be useful for the rest of the section.

Definition 6.4. A combinatorial path γ in $\mathcal{X}$ is a *good representative* of $\omega \in \Sigma^\times$ if there exists a path $\omega_0 \in \Sigma^\times$ starting at an initial state and ending at the initial vertex of ω and satisfying the following. If $w_0, w_0 w \in L_\phi$ are the words corresponding to ω_0 and $\omega_0\omega$ respectively, then γ is the portion of the path $\gamma_{w_0 w}$ representing $w_0 w$ from $\gamma^- = \tau_{\mathcal{X}}(w_0)$ to $\gamma^+ = \tau_{\mathcal{X}}(w_0 w)$.

Note that good representatives are geodesic. Also, by Theorem 5.11 (5), any two good representatives of the same path in $\Sigma^\times$ differ by a translation by an element in $\overline{\Gamma}$. In consequence, if $\omega \in P(\Sigma^\times)$ then there exists a well-defined conjugacy class $\beta(\omega) \in \mathbf{conj}(\overline{\Gamma})$ represented by any $g \in \overline{\Gamma}$ such that $\gamma^+ = g\gamma^-$ for γ a good representative of ω. Clearly $\omega \in P_n(\Sigma^\times)$ implies $\ell_{\mathcal{X}}[\beta(\omega)] = n$.

We also consider lifts of paths in $\Sigma^\times$ to $\mathcal{Z}$. First, we extend the equation (5.1) to define a map $\alpha : \Sigma^\times \to (S_{\overline{\mathcal{Z}}})^*$. If γ is a good representative of $\omega \in \Sigma^\times$ defined using the path ω_0 as above, then $\widetilde{\gamma}$ is the portion of $\widetilde{\gamma}_{\alpha(\omega_0\omega)}$ starting at $\widetilde{\gamma}^+_{\alpha(\omega_0)}$ and ending at $\widetilde{\gamma}^+_{\alpha(\omega_0\omega)}$, where $\widetilde{\gamma}_{\alpha(\omega_0)}$ and $\widetilde{\gamma}_{\alpha(\omega_0\omega)}$ are given by Theorem 5.11 (4). In this way the path $\widetilde{\gamma}$ represents the word $\alpha(\omega)$. Different choices of ω_0, ω_0' may give different lifts $\widetilde{\gamma}, \widetilde{\gamma}'$ even if $\tau_{\mathcal{X}}(w_0) = \tau_{\mathcal{X}}(w_0')$, but under this assumption we have $\phi(\widetilde{\gamma}) = \phi(\widetilde{\gamma}') = \gamma$.

A key feature of the automaton $\mathcal{A}_{\overline{\Gamma},\phi}$ is that it keeps track of translation lengths for the actions of Γ on the cuboid complexes $\mathcal{X}^{\mathfrak{w}}$ and $\mathcal{X}_*^{\mathfrak{w}_*}$, via the potential on (Σ, σ) defined below.

Definition 6.5. Let $r = r_{\mathcal{X}}^{\mathfrak{w}}, \psi = \psi_{\mathcal{X}_*}^{\mathfrak{w}_*} : E_\phi \to \mathbb{R}$ be the functions such that

$$r(e) = \overline{\mathfrak{w}}(\pi_\phi(e)) \quad \text{and} \quad \psi(e) = \psi(v_0 \xrightarrow{e} v_1) = \overline{\mathfrak{w}}_*(\Psi(v_0)) + \overline{\mathfrak{w}}_*(\pi_\phi(e)), \tag{6.2}$$

where Ψ is the function from Theorem 5.11 (4).

Remark 6.6. In the definition of ψ above, we note that $\overline{\mathfrak{w}}_*(\pi_\phi(e))$ is not necessarily zero since $S_{\overline{\mathcal{X}}}$ and $S_{\overline{\mathcal{X}}_*}$ are not necessarily disjoint. Also, our assumption that $\mathbb{W} \cup \mathbb{W}_* = \mathbb{H}(\mathcal{Z})$ implies that $S_{\overline{\mathcal{X}}} \cup S_{\overline{\mathcal{X}}_*} = S_{\overline{\mathcal{Z}}}$, so that $\Psi(v_0)$ always belongs to $(S_{\overline{\mathcal{X}}_*})^*$.

By abuse of notation we extend these functions to potentials $r, \psi : \Sigma \to \mathbb{R}$ via

$$r(e_1, e_2, \dots) = r(e_1) \ \text{ and } \ \psi(e_1, e_2, \dots) = \psi(e_1).$$

Clearly r and ψ are constant on 2-cylinders, and r is positive. The next lemma can be seen as a weak analog of [14, Lemma 3.8] for pairs of word metrics on hyperbolic groups.

Lemma 6.7. *Let $r, \psi : \Sigma \to \mathbb{R}$ be the potentials defined above. If $n \geq 1$ and $\omega \in \mathrm{Fix}_n(\Sigma)$ has truncation $t_n(\omega) \in P_n(\Sigma^\times)$, then $\ell_{\mathcal{X}}[\beta(t_n(\omega))] = n$ and the nth Birkhoff sums at ω satisfy*

$$r^n(\omega) = r(\omega) + r(\sigma(\omega)) + \cdots + r(\sigma^{n-1}(\omega)) = \ell_{\mathcal{X}}^{\mathfrak{w}}[\beta(t_n(\omega))] \tag{6.3}$$

and

$$\psi^n(\omega) = \psi(\omega) + \psi(\sigma(\omega)) + \cdots + \psi(\sigma^{n-1}(\omega)) = \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[\beta(t_n(\omega))]. \tag{6.4}$$

Proof of Lemma 6.7. Let $\omega \in \mathrm{Fix}_n(\Sigma)$ be as in the statement of the lemma and let $\omega_n = t_n(\omega) = (v_0 \xrightarrow{e_1} \cdots \xrightarrow{e_n} v_n) \in P_n(\Sigma^\times)$. As we noted previously we have $\ell_{\mathcal{X}}[\beta(\omega_n)] = n$, so the main content of the lemma are the identities (6.3) and (6.4), which we now prove.

For any $k \geq 1$, let $\omega_n^{(k)} \in P(\Sigma^{\times})$ be the concatenation of k copies of ω_n and let $(\gamma^{(k)})_k = (\gamma_{(\omega_n{}^{(k)})})_k \subset \mathcal{X}$ be a sequence of good representatives of $\omega_n^{(k)}$ with a common starting vertex $\gamma^- = (\gamma^{(k)})^-$. Let $q \in \Gamma$ satisfy $(\gamma^{(1)})^+ = q\gamma^-$, so that $[q] = \beta(\omega_n)$ and $(\gamma^{(k)})^+ = q^k\gamma^-$ for each k. Also, since ω is a closed path, we have that $\gamma^{(1)}$ is a fundamental domain for a q-invariant geodesic in $\mathcal{X}^{\mathfrak{w}}$, and hence

$$\begin{aligned}r^n(\omega) &= \overline{\mathfrak{w}}(\pi_\phi(e_1)) + \cdots + \overline{\mathfrak{w}}(\pi_\phi(e_n))\\ &= d_{\mathcal{X}}^{\mathfrak{w}}((\gamma^{(1)})^-, (\gamma^{(1)})^+) = d_{\mathcal{X}}^{\mathfrak{w}}(\gamma^-, q\gamma^-) = \ell_{\mathcal{X}}^{\mathfrak{w}}[q] = \ell_{\mathcal{X}}^{\mathfrak{w}}[\beta(t_n(\omega))].\end{aligned}$$

This proves (6.3).

To prove (6.4), let L be such that $\phi_* : \mathcal{Z} \to \mathcal{X}_*^{\mathfrak{w}_*}$ is L-Lipschitz and consider the lifts $\widetilde{\gamma}^{(k)} \subset \mathcal{Z}$ of the geodesic $\gamma^{(k)}$ with a common starting vertex $\widetilde{\gamma}^- = (\widetilde{\gamma}^{(k)})^-$. We project these paths to $\mathcal{X}_*$ by defining $\gamma_*^- := \phi_*(\widetilde{\gamma}^-)$ and $(\gamma_*^{(k)})^+ := \phi_*((\widetilde{\gamma}^{(k)})^+)$, and for all k we get

$$|d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^-, (\gamma_*^{(k)})^+) - d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^-, q^k\gamma_*^-)| \leq d_{\mathcal{X}_*}^{\mathfrak{w}_*}((\gamma_*^{(k)})^+, q^k\gamma_*^-) \leq L d_{\mathcal{Z}}((\widetilde{\gamma}^{(k)})^+, q^k\widetilde{\gamma}^-). \tag{6.5}$$

The last term in the inequality above is bounded by a number independent of k. Indeed, since $\phi((\widetilde{\gamma}^{(k)})^+) = \phi(q^k\widetilde{\gamma}^-) = q^k\gamma^-$, we have that both $(\widetilde{\gamma}^{(k)})^+$ and $q^k\widetilde{\gamma}^-$ belong to the preimage of $q^k\gamma^-$ under ϕ, so their distance is bounded by a number independent of k because $\phi : \mathcal{Z} \to \mathcal{X}$ is a quasi-isometry.

Also, we note that $d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^-, (\gamma_*^{(k)})^+) = \overline{\mathfrak{w}}_*(\alpha(\omega_n^{(k)}))$, which equals $k \cdot \overline{\mathfrak{w}}_*(\alpha(\omega_n))$ since $\alpha(\omega_n^{(k)})$ is the concatenation of k copies of $\alpha(\omega_n)$. By Theorem 5.11 (4), it follows that

$$\begin{aligned}d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^-, (\gamma_*^{(k)})^+) &= k \cdot \overline{\mathfrak{w}}_*(\alpha(\omega_n))\\ &= k \cdot (\overline{\mathfrak{w}}_*(\Psi(v_0)\pi_\phi(e_1)) + \overline{\mathfrak{w}}_*(\Psi(v_1)\pi_\phi(e_2)) + \cdots + \overline{\mathfrak{w}}_*(\Psi(v_{n-1})\pi_\phi(e_n)))\\ &= k \cdot (\psi(v_0 \xrightarrow{e_1} v_1) + \psi(v_1 \xrightarrow{e_2} v_2) + \cdots + \psi(v_{n-1} \xrightarrow{e_n} v_n))\\ &= k \cdot (\psi(\omega) + \psi(\sigma(\omega)) + \cdots + \psi(\sigma^{n-1}(\omega))) = k\psi^n(\omega).\end{aligned}$$

Therefore, combining this with (6.5) and after dividing by k and letting k tend to infinity we obtain $\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[\beta(t_n(\omega))] = \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[q] = \psi^n(\omega)$, as desired. □

To apply the results from Section 3 we require a mixing (or at least transitive) dynamical system. To obtain such a system we consider a maximal recurrent component $\mathcal{C}$ of the graph $\mathcal{G}_\phi$. As before we let $\Sigma_{\mathcal{C}}^{\times} \subset \Sigma^{\times}$ and $\Sigma_{\mathcal{C}} \subset \Sigma$ be the subsets corresponding to paths in $\mathcal{C}$, and note that $\Sigma_{\mathcal{C}}$ is σ-invariant. Similarly we define $P(\Sigma_{\mathcal{C}}^{\times})$, $P_n(\Sigma_{\mathcal{C}}^{\times})$ and $P_{\leq n}(\Sigma_{\mathcal{C}}^{\times})$, and we identify $P_n(\Sigma_{\mathcal{C}}^{\times})$ with $\mathrm{Fix}_n(\Sigma_{\mathcal{C}})$.

Let $r_{\mathcal{X}}^{\mathfrak{w}} : \Sigma_{\mathcal{C}} \to \mathbb{R}_{>0}$ be the (constant on 2-cylinders) restriction to $\Sigma_{\mathcal{C}}$ of the potential introduced in Definition 6.5 and consider the suspension flow:

$$\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}} := \{(\omega, t) \in \Sigma_{\mathcal{C}} \times \mathbb{R} : 0 \leq t \leq r_{\mathcal{X}}^{\mathfrak{w}}(\omega)\}/\sim,$$

where each $(\omega, r_{\mathcal{X}}^{\mathfrak{w}}(\omega))$ is identified with $(\sigma(\omega), 0)$ and the flow $\sigma^{r_{\mathcal{X}}^{\mathfrak{w}}} = (\sigma_s^{r_{\mathcal{X}}^{\mathfrak{w}}})_{s \in \mathbb{R}_{>0}}$ acts as $\sigma_s^{r_{\mathcal{X}}^{\mathfrak{w}}}(\omega, t) = (\omega, t + s)$.

Note that any closed $\sigma^{r_{\mathcal{X}}^{\mathfrak{w}}}$-orbit τ in $\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}}$ corresponds to a closed σ-orbit in $\Sigma_{\mathcal{C}}$. More precisely, such an orbit τ must be of the form $\tau = \{(\omega, t) : 0 \leq t \leq (r_{\mathcal{X}}^{\mathfrak{w}})^n(\omega)\}$ for some $\omega \in \Sigma_{\mathcal{C}}$ such that $\sigma^n(\omega) = \omega$. In this case the period of τ equals $l_\tau = (r_{\mathcal{X}}^{\mathfrak{w}})^n(\omega)$.

We fix a smooth function $\Delta : [0,1] \to \mathbb{R}_{\geq 0}$ such that $\Delta(0) = \Delta(1) = 1$ and $\int_0^1 \Delta(t)\, dt = 1$ and define $\Phi : \Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}} \to \mathbb{R}_{\geq 0}$ according to

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

$$\Phi(\omega,t)=\Delta\left(\frac{t}{r^{\mathfrak{w}}_{\mathcal{X}}(\omega)}\right)\frac{\psi^{\mathfrak{w}_*}_{\mathcal{X}_*}(\omega)}{r^{\mathfrak{w}}_{\mathcal{X}}(\omega)} \text{ for each } (\omega,t)\in\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}} \text{ with } 0\le t\le r^{\mathfrak{w}}_{\mathcal{X}}(\omega). \tag{6.6}$$

This function has the property that, for any closed $\sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}$-orbit τ in $\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}}$ with period l_τ and corresponding periodic σ-orbit $\omega,\sigma(\omega),\dots,\sigma^n(\omega)=\omega$ in $\Sigma_{\mathcal{C}}$ we have

$$\int_\tau\Phi:=\int_0^{l_\tau}\Phi(\sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_t(\omega,0))\,dt=\ell^{\mathfrak{w}_*}_{\mathcal{X}_*}[\beta(t_n(\omega))],$$

where we have used Lemma 6.7. For τ,ω, and n as above we adopt the notation

$$\beta(\tau)=\beta(t_n(\omega)),$$

which defines a map $\beta:P(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}})\to\mathbf{conj}(\overline{\Gamma})$ from the set $P(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}})$ of periodic orbits of $\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}}$ into $\mathbf{conj}(\overline{\Gamma})$. By Lemma 6.7, the period of any $\tau\in P(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}})$ equals $\ell^{\mathfrak{w}}_{\mathcal{X}}[\beta(\tau)]$.

For such a suspension flow $(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}},\sigma^{r^{\mathfrak{w}}_{\mathcal{X}}})$, the Manhattan curve for $(\mathcal{X}^{\mathfrak{w}},\mathcal{X}^{\mathfrak{w}_*}_*)$ can be described in terms of the pressures related to Φ, as stated in the next proposition.

Proposition 6.8. *Let $\Gamma,\mathcal{X},\mathcal{X}_*,\mathcal{Z},\mathfrak{w},\mathfrak{w}_*$ satisfy Convention 6.3 and let $\mathcal{A}_{\overline{\Gamma},\phi}$ and $r^{\mathfrak{w}}_{\mathcal{X}},\psi^{\mathfrak{w}_*}_{\mathcal{X}_*}:\Sigma\to\mathbb{R}$ be given by Theorem 5.11 and Definition 6.5 respectively. If $\mathcal{C}$ is a maximal recurrent component of $\mathcal{G}_\phi$ and $\Phi:\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}}\to\mathbb{R}$ is given by (6.6), then for any $s\in\mathbb{R}$ we have*

$$\theta_{\mathcal{X}^{\mathfrak{w}_*}_*/\mathcal{X}^{\mathfrak{w}}}(s)=\mathrm{P}_{\mathcal{C}}(-s\Phi),$$

where $\mathrm{P}_{\mathcal{C}}(-s\Phi)$ is the pressure of the potential $-s\Phi$ on the suspension $(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}},\sigma^{r^{\mathfrak{w}}_{\mathcal{X}}})$.

In particular, this result implies that the pressure $s\mapsto\mathrm{P}_{\mathcal{C}}(-s\Phi)$ is independent of the choice of maximal recurrent component $\mathcal{C}$.

Remark 6.9. As it will be clear from the proof of Proposition 6.8, if $\mathcal{C}$ is any maximal recurrent component of $\mathcal{G}_\phi$ then for any $s\in\mathbb{R}$ we also have

$$\theta_{\mathcal{X}_*/\mathcal{X}}(s)=\mathrm{P}_{\mathcal{C}}(-s\psi),$$

where $\psi=\psi^{\mathfrak{w}_*}_{\mathcal{X}_*}$ for $\mathfrak{w}_*\equiv 1$ the constant orthotope structure and $\mathrm{P}_{\mathcal{C}}(-s\psi)$ is the pressure of the potential $-s\psi$ on $(\Sigma_{\mathcal{C}},\sigma)$.

The rest of this subsection is devoted to proving this proposition, and using the formalism necessary in its proof to deduce Theorem 1.4. For the sequel we fix a compact set $K\subset\mathcal{X}$ such that $\mathcal{X}=\Gamma K$ and assume that $o\in K$. Since $\mathcal{G}_\phi$ is finite and pruned we also fix $N>0$ such that any $\omega\in\Sigma^\times$ has a good representative γ satisfying $d_{\mathcal{X}}(\gamma^-,o)\le N$. Otherwise explicit, for any $\omega\in\Sigma^\times$ we fix a good representative $\gamma=\gamma_\omega$ that minimizes $d_{\mathcal{X}}(\gamma^-,o)$ and define

$$\Gamma_{\mathcal{C}}:=\{g\in\Gamma:\ \text{there exists } \omega\in\Sigma^\times_{\mathcal{C}} \text{ such that } \gamma^+_\omega\in gK\}.$$

We also write $B_n=\{g\in\Gamma:d_{\mathcal{X}}(go,o)\le n\}$ for each $n\ge 0$.

Lemma 6.10. *Let C be the constant from Theorem 5.11 (1). Then*

$$\sup_{x\in\mathcal{X}^0}\#\{\omega\in\Sigma^\times:\ \gamma^+_\omega=x\}\le C(N+1).$$

Proof. Let $x\in\mathcal{X}^0$ and $\omega\in\Sigma^\times$ be such that $\gamma^+_\omega=x$, so that γ_ω is constructed from a path $\omega_0\in\Sigma^\times$ such that $\omega'=\omega_0\omega$ also belongs to $\Sigma^\times$. Then ω_0 is a prefix of ω' of length at most N and ω' represents the unique word $w \in L_\phi$ satisfying $\tau_{\mathcal{X}}(w) = x$. From this we deduce

$$\begin{aligned}\#\{\omega \in \Sigma^\times : \gamma_\omega^+ = x\} &\leq (N+1) \cdot \#\{\omega' \in \Sigma^\times : \omega' \text{ starts at an initial state and } \gamma_{\omega'}^+ = x\} \\ &\leq (N+1) \cdot \#\{\omega' \in \Sigma^\times : \omega' \text{ represents } w \text{ and starts at an initial state}\},\end{aligned}$$

and the lemma follows from Theorem 5.11 (1). □

Lemma 6.11. *The set $\Gamma_{\mathcal{C}}$ has positive lower density for the action on $\mathcal{X}$. That is*

$$\liminf_{n\to\infty} \frac{\#(\Gamma_{\mathcal{C}} \cap B_n)}{\#B_n} > 0.$$

Proof. For each n we let $(\Sigma_{\mathcal{C}}^\times)_{\leq n}$ denote the set of paths in $\mathcal{C}$ of length at most n. First we claim that there exist $B > 1$ and $0 \leq \lambda < e^{v_{\mathcal{X}}}$ (depending only on the adjacency matrix of $\mathcal{C}$) such that

$$\#(\Sigma_{\mathcal{C}}^\times)_{\leq n} \geq B^{-1} e^{n v_{\mathcal{X}}} - B\lambda^n \tag{6.7}$$

for all n large enough. To show this, let A be the adjacency matrix of $\mathcal{C}$, which is irreducible since $\mathcal{C}$ is recurrent. Moreover, $\mathcal{C}$ being maximal implies that the spectral radius of A equals $e^{v_{\mathcal{X}}}$. Suppose that A has $p \geq 1$ eigenvalues of absolute value $e^{v_{\mathcal{X}}}$ and let $0 \leq \lambda < e^{v_{\mathcal{X}}}$ be any number greater than the absolute value of all the other eigenvalues of A. By [26, Theorem 3.1], for each $k \geq 1$ the matrix A^{kp} has $e^{kpv_{\mathcal{X}}}$ as eigenvalue of multiplicity p and all its other eigenvalues have absolute value less than λ^{kp}. In particular, its trace satisfies

$$\operatorname{tr}(A^{kp}) \geq p e^{kpv_{\mathcal{X}}} - (\dim A - p)\lambda^{kp}.$$

But $\operatorname{tr}(A^{kp})$ equals the number of closed paths of length kp in $\mathcal{C}$, so if $n = kp + r$ with $0 \leq r < p$ an integer and $k \geq 1$, then

$$\begin{aligned}\#(\Sigma_{\mathcal{C}}^\times)_{\leq n} \geq \#(\Sigma_{\mathcal{C}}^\times)_{\leq kp} \geq \operatorname{tr}(A^{kp}) &\geq p e^{kpv_{\mathcal{X}}} - (\dim A - p)\lambda^{kp} \\ &\geq (pe^{-pv_{\mathcal{X}}}) e^{nv_{\mathcal{X}}} - (\dim A - p)\lambda^n.\end{aligned}$$

This concludes the proof of the claim. Now consider n large enough and $\omega \in (\Sigma_{\mathcal{C}}^\times)_{\leq n}$. Since the Γ-translates of K cover $\mathcal{X}$ we have $\gamma_\omega^+ \in gK$ for some $g \in \Gamma$, so that $d_{\mathcal{X}}(\gamma_\omega^+, go) \leq D$ with D being the diameter of K. In addition, by definition we have $d_{\mathcal{X}}(\gamma_\omega^-, o) \leq N$ and hence $d_{\mathcal{X}}(o, go) \leq n + D + N$. By Lemma 6.10 this implies that

$$\#(\Sigma_{\mathcal{C}}^\times)_{\leq n} \leq \#(\Gamma_{\mathcal{C}} \cap B_{n+D+N}) \cdot \sup_{g\in\Gamma} \#\{\omega \in \Sigma_{\mathcal{C}}^\times : \gamma_\omega^+ \in gK\} \leq (N+1)C \cdot \#K \cdot \#(\Gamma_{\mathcal{C}} \cap B_{n+D+N}),$$

where C is the constant from Theorem 5.11 (1).

Combining this with (6.7) we get

$$\begin{aligned}\#(\Gamma_{\mathcal{C}} \cap B_n) &\geq ((N+1)C\#K)^{-1} B^{-1} e^{(n-D-N)v_{\mathcal{X}}} - ((N+1)C\#K)^{-1} B\lambda^{n-D-N} \\ &= [((N+1)(C\#K)B)^{-1} e^{-(D+N)v_{\mathcal{X}}}] e^{nv_{\mathcal{X}}} - [((N+1)C\#K)^{-1} B\lambda^{-D-N}]\lambda^n.\end{aligned}$$

Finally, the since the action on $\mathcal{X}$ has a contracting element, [88, Theorem 1.8 (2)] implies that there exists $C' > 0$ such that $\#B_n \leq C' e^{nv_{\mathcal{X}}}$ for all n large enough, and the conclusion follows. □

The next two results are used to give a uniform comparison of the number of conjugacy classes in Γ with a bound on their translation lengths (with respect to $\mathcal{X}^{\mathfrak{w}}$) and the number of periodic orbits in the suspension $(\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}}, \sigma^{r_{\mathcal{X}}^{\mathfrak{w}}})$ with bounded period. The assumption of having a contracting element is

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

essential in the proof of the Lemma 6.13. For $T \geq R \geq 0$, recall that $P(\Sigma_C^{r_{\mathcal{X}}^{\mathfrak{w}}}, R, T)$ denotes the set of periodic orbits in $(\Sigma_C^{r_{\mathcal{X}}^{\mathfrak{w}}}, \sigma^{r_{\mathcal{X}}^{\mathfrak{w}}})$ with period on the interval $[T-R, T+R]$.

Lemma 6.12. *For any $R > 0$ there exists a polynomial Q (depending on R) that is nondecreasing on $\mathbb{R}_{>0}$ and such that for any $[g] \in \mathbf{conj}(\overline{\Gamma})$ we have*

$$\#\{\tau \in P(\Sigma_C^{r_{\mathcal{X}}^{\mathfrak{w}}}, R, T) : \beta(\tau) = [g]\} \leq Q(T). \tag{6.8}$$

Proof. To solve the lemma it is enough to show that there exists a polynomial $\widetilde{Q}$ that is increasing on $\mathbb{R}_{>0}$ and such that for any $[g] \in \mathbf{conj}(\overline{\Gamma})$ we have

$$\#\{\omega \in P(\Sigma^{\times}) : \beta(\omega) = [g]\} \leq \widetilde{Q}(\ell_{\mathcal{X}}[g]). \tag{6.9}$$

Indeed, since $\mathfrak{w}$ is nonvanishing, the identity map $\mathrm{Id} : \mathcal{X} \to \mathcal{X}^{\mathfrak{w}}$ is a quasi-isometry and $\ell_{\mathcal{X}}[g] \leq L\ell_{\mathcal{X}}^{\mathfrak{w}}[g]$ for any $[g] \in \mathbf{conj}(\Gamma)$, for $L = (\min\{\mathfrak{w}(\mathfrak{h}) : \mathfrak{h} \in \mathbb{H}(\mathcal{X})\})^{-1}$. Also, if $\tau \in P(\Sigma_C^{r_{\mathcal{X}}^{\mathfrak{w}}}, R, T)$ satisfies $\beta(\tau) = [g]$, then $l_\tau = \ell_{\mathcal{X}}^{\mathfrak{w}}[g] \in [T-R, T+R]$. Since any periodic orbit in $P(\Sigma_C^{r_{\mathcal{X}}^{\mathfrak{w}}})$ corresponds to an (orbit determined by an) element in $P(\Sigma_C^{\times}) \subset P(\Sigma^{\times})$, for any $[g]$ such that the left-hand side in (6.8) is nonzero we have

$$\begin{aligned}\#\{\tau \in P(\Sigma_C^{r_{\mathcal{X}}^{\mathfrak{w}}}, R, T) : \beta(\tau) = [g]\} &\leq \#\{\omega \in P(\Sigma^{\times}) : \beta(\omega) = [g]\} \\ &\leq \widetilde{Q}(\ell_{\mathcal{X}}[g]) \leq \widetilde{Q}(LT + LR) =: Q(T).\end{aligned}$$

To prove (6.9) we claim that for any group $\overline{\Gamma}$ acting properly, cocompactly and co-specially on a CAT(0) cube complex $\mathcal{X}$ and for any $x \in \mathcal{X}^0$ and $R \geq 0$ there exists a polynomial $\hat{Q}$ that is increasing on $\mathbb{R}_{>0}$ and such that

$$\#\{g \in [g] : d_{\mathcal{X}}(gx, x) \leq \ell_{\mathcal{X}}[g] + R\} \leq \hat{Q}(\ell_{\mathcal{X}}[g]) \tag{6.10}$$

for any $[g] \in \mathbf{conj}(\overline{\Gamma})$.

To see how this claim proves the lemma, fix $[g] \in \mathbf{conj}(\overline{\Gamma})$ and let $\omega \in P(\Sigma^{\times})$ be such that $\beta(\omega) = [g]$. If $\gamma = \gamma_\omega$ is a good representative such that $d_{\mathcal{X}}(\gamma^-, o) \leq N$, we let $q \in [g]$ be such that $\gamma^+ = q\gamma^-$. Then $d_{\mathcal{X}}(qo, \gamma^+) = d_{\mathcal{X}}(o, \gamma^-) \leq N$ and we have

$$|d_{\mathcal{X}}(o, qo) - \ell_{\mathcal{X}}[g]| = |d_{\mathcal{X}}(o, qo) - d_{\mathcal{X}}(\gamma^-, \gamma^+)| \leq 2N.$$

In addition, there exists a constant $\hat{C} > 0$ such that for any $[g] \in \mathbf{conj}(\overline{\Gamma})$ and any $q \in [g]$ satisfying $|d_{\mathcal{X}}(o, qo) - \ell_{\mathcal{X}}[g]| \leq 2N$, the set

$$\{\omega \in P(\Sigma^{\times}) : \beta(\omega) = [g] \text{ and } d_{\mathcal{X}}(\gamma_\omega^+, qo) \leq N\}$$

has cardinality at most $\hat{C}$. Indeed, if B is the set of vertices at distance at most N from o, then this cardinality is bounded above by

$$\sum_{x \in B} \#\{\omega \in P(\Sigma^{\times}) : \gamma_\omega^+ = qx\} \leq \#B \cdot \sup_{x \in \mathcal{X}^0} \#\{\omega \in P(\Sigma^{\times}) : \gamma_\omega^+ = x\} \leq \#B \cdot C(N+1) =: \hat{C},$$

where for the last inequality we used Lemma 6.10.

Applying this to our case of interest, we deduce

$$\#\{\omega \in P(\Sigma^{\times}) : \beta(\omega) = [g]\} \leq \hat{C} \cdot \#\{q \in [g] : |d_{\mathcal{X}}(qo, o) - \ell_{\mathcal{X}}[g]| \leq 2N\} \leq \hat{C} \cdot \hat{Q}(\ell_{\mathcal{X}}[g]),$$

where $\hat{Q}$ is the polynomial given by the claim for $x = o$ and $R = 2N$.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

To prove the claim (6.10), by equivariantly embedding $\mathcal{X}$ as a convex subcomplex of the universal cover of a Salvetti complex we can assume that $\overline{\Gamma}$ is a right-angled Artin group with standard (symmetric) generating set S and $\mathcal{X}$ is the universal cover of its Salvetti complex. Then $\mathcal{X}^1$ is the Cayley graph for $\overline{\Gamma}$ with respect to S, and since the expected conclusion is independent of the base point we can assume that $x = o$ is the identity element of $\overline{\Gamma}$, so that $d_{\mathcal{X}}(gx, x) = |g|_S$ is the word length of g for any $g \in \overline{\Gamma}$.

Now we fix $[g] \in \mathbf{conj}(\overline{\Gamma})$, set $\ell = \ell_{\mathcal{X}}[g] = \ell_S[g]$ and consider the sets $E_n[g] = \{g \in [g] : |g|_S \leq \ell + n\}$. Note that $\#E_0[g] \leq \ell$ since any two conjugate elements that minimize the word length are actually cyclically conjugated with respect to some minimal word representations. Also, if $g \in E_n[g]$, and $n > 0$, then indeed $n \geq 2$ and g is represented by a word of the form $x_1 a^{\pm} x_2 a^{\mp} x_3$, where $a \in S$ is a standard generator and all the letters in the words x_1 and x_3 commute with a. Then the element g' represented by the word $x_1 x_2 x_3$ belongs to $E_{n-2}[g]$, and there are at most $\#S \cdot (\ell + n)(\ell + n - 1)/2$ ways to reconstruct g from g'. Therefore, we have

$$\#E_n[g] \leq \#S \cdot (\ell + n)(\ell + n - 1)/2 \cdot \#E_{n-2}[g]$$

for each n, and hence $\#\{g \in [g] : |g|_S \leq \ell + R\} \leq \#E_{2R}[g] \leq \hat{Q}(\ell)$ for

$$\hat{Q}(t) = (\#S)^R \cdot (t + 2R)(t + 2R - 1) \cdots (t + 1)t/2^R.$$

This concludes the proof of the claim and the lemma. □

Lemma 6.13. *There exists $C' > 0$ such that for any nontorsion conjugacy class $[g] \in \mathbf{conj}(\Gamma)$ we can find a representative $\hat{g} \in [g]$ and a closed path $\omega_{[g]} \in P(\Sigma_{\mathcal{C}}^{\times})$ satisfying*

$$\min\{d_{\mathcal{X}}(\hat{g}o, \gamma^+_{\omega_{[g]}}), d_{\mathcal{X}}(\hat{g}^{-1}o, \gamma^+_{\omega_{[g]}})\} \leq C',$$

and additionally

$$\max\{|\ell^{\mathfrak{w}}_{\mathcal{X}}[g] - \ell^{\mathfrak{w}}_{\mathcal{X}}[\beta(\omega_{[g]})]|, |\ell^{\mathfrak{w}_*}_{\mathcal{X}_*}[g] - \ell^{\mathfrak{w}_*}_{\mathcal{X}_*}[\beta(\omega_{[g]})]|\} \leq C'.$$

Remark 6.14. As we will see in the proof, the constant C' above depends on the initial data from Convention 6.3 and the maximal component $\mathcal{C}$. The constant also depends on the existence of a contracting element in Γ, as we rely on the work of Yang [88, Theorem C].

We will need the next lemma, which follows immediately from Remarks 2.1 and 2.2.

Lemma 6.15. *Let γ be a g-invariant geodesic in $\mathcal{Z}$ for some $g \in \Gamma$. Then the image of γ under ϕ (resp. ϕ_*) in $\mathcal{X}$ (resp. $\mathcal{X}_*$) is a (possibly nonparametrized) g-invariant geodesic. For $\mathcal{X}_*$ we allow the degenerate case that $\phi_*(\gamma)$ is a point (which happens if and only if $\ell_{\mathcal{X}_*}[g] = 0$).*

Proof of Lemma 6.13. Before starting with the proof we provide a brief sketch. Given the conjugacy class $[g]$, we first find an appropriate representative $g \in [g]$ such that $d_{\mathcal{Z}}(\widetilde{o}, g\widetilde{o})$ is uniformly comparable to $\ell_{\mathcal{Z}}[g]$, with similar versions for $\ell^{\mathfrak{w}}_{\mathcal{X}}$ and $\ell^{\mathfrak{w}_*}_{\mathcal{X}_*}$. Then, we use Lemma 6.11 and the existence of a contracting element to find an element $s \in \overline{\Gamma}$ and a path $\omega' \in \Sigma^*_{\mathcal{C}}$ such that both so and sgo are within uniformly bounded distance from the good representative path $\gamma_{\omega'}$. Using the recurrence of $\mathcal{C}$, we construct the close path $\omega_{[g]}$ as the concatenation of an appropriate subpath of ω' and a path in $\Sigma^*_{\mathcal{C}}$ of uniformly bounded length. For the path representative $\gamma = \gamma_{\omega_{[g]}}$ we then find $s' \in \overline{\Gamma}$ such that $d_{\mathcal{X}}(s'so, \gamma^-)$ and $d_{\mathcal{X}}(s'sgo, \gamma^+)$ are uniformly bounded. The rest of the proof consists of verifying that $\hat{g} = s'sg(s's)^{-1}$ satisfies all the desired inequalities for the appropriate constant C'.

Now we start with the proof of the lemma, for which we consider the following constants. Let M_1 be such that any two vertices in $\mathcal{C}$ can be joined by a path in $\mathcal{C}$ of length at most M_1 (in both directions). This number exists since $\mathcal{C}$ is recurrent. Also, the projection $\phi : \mathcal{Z} \to \mathcal{X}$ is a quasi-isometry since it is Γ-equivariant and the action of Γ on both $\mathcal{Z}$ and $\mathcal{X}$ is proper and cocompact, so let $M_2 > 0$ be such that

$$d_{\mathcal{Z}}(x, y) \leq M_2 d_{\mathcal{X}}(\phi(x), \phi(y)) + M_2$$

for all $x, y \in \mathcal{Z}$. In addition, let M_3 be the diameter of $\phi^{-1}(K) \subset \mathcal{Z}$ and fix a constant M_4 larger than N and the diameter of K. Finally, let L be the maximum of all the weights $\mathfrak{w}(\mathfrak{h})$ or $\mathfrak{w}_*(\mathfrak{h}_*)$ among hyperplanes $\mathfrak{h} \in \mathbb{H}(\mathcal{X})$ and $\mathfrak{h}_* \in \mathbb{H}(\mathcal{X}_*)$, so that the four functions

$$\phi : \mathcal{Z} \to \mathcal{X}^{\mathfrak{w}}, \quad \phi_* : \mathcal{Z} \to \mathcal{X}_*^{\mathfrak{w}_*}, \quad \mathrm{Id} : \mathcal{X} \to \mathcal{X}^{\mathfrak{w}}, \quad \mathrm{Id} : \mathcal{X}_* \to \mathcal{X}_*^{\mathfrak{w}_*}$$

are L-Lipschitz.

We now start the proof, so we let $g \in \Gamma$ represent the nontorsion conjugacy class $[g] \in \mathbf{conj}(\Gamma)$. Then g fixes a bi-infinite combinatorial axis $\widetilde{\lambda}$ in the cubical barycentric subdivision $\dot{\mathcal{Z}}$ [46, Theorem 1.4]. After conjugating by an element of Γ we can assume that $d_{\mathcal{Z}}(\widetilde{o}, \widetilde{\lambda}) \leq M_3$, so in particular we have $|d_{\mathcal{Z}}(\widetilde{o}, g\widetilde{o}) - \ell_{\mathcal{Z}}[g]| \leq 2M_3$.

By Lemma 6.15 and the fact that ϕ, ϕ_* are Lipschitz, the images $\lambda = \phi(\widetilde{\lambda})$ and $\lambda_* = \phi_*(\widetilde{\lambda})$ are also g-invariant (unparametrized) geodesics satisfying $d^{\mathfrak{w}}_{\mathcal{X}}(o, \lambda) \leq LM_3$ and $d^{\mathfrak{w}_*}_{\mathcal{X}_*}(o_*, \lambda_*) \leq LM_3$, which gives us

$$|d^{\mathfrak{w}}_{\mathcal{X}}(o, go) - \ell^{\mathfrak{w}}_{\mathcal{X}}[g]| \leq 2LM_3 \quad \text{and} \quad |d^{\mathfrak{w}_*}_{\mathcal{X}_*}(o_*, go_*) - \ell^{\mathfrak{w}_*}_{\mathcal{X}_*}[g]| \leq 2LM_3. \tag{6.11}$$

The action of Γ on $\mathcal{X}$ is proper, cocompact and has a contracting element, and hence by [88, Theorem C] there exists a constant $\epsilon > 0$ satisfying the following for any $h \in \Gamma$. Let $\mathcal{V}_h$ denote the set of all the group elements $k \in \Gamma$ such that if $\gamma \subset \mathcal{X}$ is a combinatorial geodesic path with endpoints $\gamma^{\pm}$ verifying $d_{\mathcal{X}}(\gamma^-, o) \leq M_4$ and $d_{\mathcal{X}}(\gamma^+, ko) \leq M_4$, then there exists no $s \in \Gamma$ such that $d_{\mathcal{X}}(so, \gamma) \leq \epsilon$ and $d_{\mathcal{X}}(sho, \gamma) \leq \epsilon$. Then

$$\lim_{n \to \infty} \frac{\#(\mathcal{V}_h \cap B_n)}{\# B_n} = 0$$

(the freedom in our choice for M_4 comes from the Remark after Theorem C in [88]). In virtue of Lemma 6.11 we conclude that the set $\Gamma_{\mathcal{C}} \backslash \mathcal{V}_h$ is nonempty for every $h \in \Gamma$. Applying this to $h = g$ we deduce the existence of a path $\omega' \in \Sigma^{\times}_{\mathcal{C}}$ and $s \in \Gamma$ such that $d_{\mathcal{X}}(so, \gamma_{\omega'}) \leq \epsilon$ and $d_{\mathcal{X}}(sgo, \gamma_{\omega'}) \leq \epsilon$.

Let $u, v \in \gamma_{\omega'}$ be such that $d_{\mathcal{X}}(so, u) \leq \epsilon$ and $d_{\mathcal{X}}(sgo, v) \leq \epsilon$, and without loss of generality assume that u belongs to the portion of $\gamma_{\omega'}$ from $\gamma^-_{\omega'}$ to v. Let $\omega = \omega_{[g]} \in \Sigma^{\times}_{\mathcal{C}}$ be a closed path composed by the concatenation of the subpath $\overline{\omega}'$ of ω' that determines the portion of $\gamma_{\omega'}$ from u to v and a path in $\Sigma^{\times}_{\mathcal{C}}$ of length at most M_1 from the final vertex of $\overline{\omega}'$ to its initial vertex. Let $\gamma = \gamma_\omega \subset \mathcal{X}$ be the good representative of ω with

$$L^{-1} d^{\mathfrak{w}}_{\mathcal{X}}(\gamma^-, o) \leq d_{\mathcal{X}}(\gamma^-, o) \leq N, \tag{6.12}$$

and let $s' \in \overline{\Gamma}$ be such that $\gamma^- = s'u$. This implies

$$L^{-1} d^{\mathfrak{w}}_{\mathcal{X}}(s'so, \gamma^-) \leq d_{\mathcal{X}}(s'so, \gamma^-) \leq \epsilon \quad \text{and} \quad L^{-1} d^{\mathfrak{w}}_{\mathcal{X}}(s'sgo, \gamma^+) \leq d_{\mathcal{X}}(s'sgo, \gamma^+) \leq \epsilon + M_1. \tag{6.13}$$

Since ω is a loop we have $\gamma^+ = q\gamma^-$ for $[q] = \beta(\omega) \in \mathbf{conj}(\overline{\Gamma})$, and in particular from (6.11) we get

$$|\ell^{\mathfrak{w}}_{\mathcal{X}}[g] - \ell^{\mathfrak{w}}_{\mathcal{X}}[\beta(\omega)]| \leq L(2\epsilon + M_1 + 2M_3).$$

Also, for $k \geq 1$ let $\omega^{(k)} \in \Sigma^{\times}_{\mathcal{C}}$ be the concatenation of k copies of ω, and let $\gamma^{(k)} \subset \mathcal{X}$ be a good representative of $\omega^{(k)}$ so that $(\gamma^{(k)})^- = \gamma^-$ and $(\gamma^{(k)})^+ = q^k\gamma^-$. Note that $\gamma^{(k)}$ is always a subpath of $\gamma^{(k+1)}$.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Now we lift each $\gamma^{(k)}$ to $\mathcal{Z}$ to get a sequence $\widetilde{\gamma}^{(k)}$ of geodesic paths in $\mathcal{Z}$. Then $\phi(\widetilde{\gamma}^{(k)}) = \gamma^{(k)}$ (up to parametrization), so that $(\gamma^{(k)})^{\pm} = \phi((\widetilde{\gamma}^{(k)})^{\pm})$ for all k. We denote $\widetilde{\gamma}^{\pm} = (\widetilde{\gamma}^{(1)})^{\pm}$ and we assume $(\widetilde{\gamma}^{(k)})^{-} = \widetilde{\gamma}^{-}$ for all k. By Γ-equivariance of ϕ we have $\phi((\widetilde{\gamma}^{(k)})^{+}) = (\gamma^{(k)})^{+} = q^k\gamma^{-} = \phi(q^k\widetilde{\gamma}^{-})$, and hence

$$d_{\mathcal{Z}}((\widetilde{\gamma}^{(k)})^{+}, q^k\widetilde{\gamma}^{-}) \leq M_3 \tag{6.14}$$

for all k. Also, the inequalities (6.13) imply

$$d_{\mathcal{Z}}(s'sð\widetilde{o}, \widetilde{\gamma}^{-}) \leq M_2\epsilon + M_2 \quad \text{and} \quad d_{\mathcal{Z}}(s'sg\widetilde{o}, \widetilde{\gamma}^{+}) \leq M_2(\epsilon + M_1) + M_2. \tag{6.15}$$

We project the geodesics $\widetilde{\gamma}^{(k)}$ to $\mathcal{X}_*^{\mathfrak{w}_*}$ via ϕ_*, so we consider $\gamma_*^{(k)} := \phi_*(\widetilde{\gamma}^{(k)})$, which are (unparametrized) geodesics by Lemma 6.15, and as before we denote $\gamma_*^{\pm} = \phi_*(\gamma_*^{\pm})$.

The length of $\gamma_*^{(k)}$ in $\mathcal{X}_*^{\mathfrak{w}_*}$ equals $\overline{\mathfrak{w}}_*(\alpha(\omega^{(k)}))$, and since the word $\alpha(\omega^{(k)})$ is the concatenation of k copies of $\alpha(\omega)$ we have

$$d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, (\gamma_*^{(k)})^{+}) = k d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, \gamma_*^{+}) \tag{6.16}$$

for all k. In addition, by (6.14) and (6.15) we obtain

$$d_{\mathcal{X}_*}^{\mathfrak{w}_*}((\gamma_*^{(k)})^{+}, q^k\gamma_*^{-}) \leq LM_3, \tag{6.17}$$

and

$$d_{\mathcal{X}_*}^{\mathfrak{w}_*}(s'so_*, \gamma_*^{-}) \leq LM_2(\epsilon + 1) \quad \text{and} \quad d_{\mathcal{X}_*}^{\mathfrak{w}_*}(s'sgo_*, \gamma_*^{+}) \leq LM_2(\epsilon + M_1 + 1). \tag{6.18}$$

From these inequalities and (6.11) we get

$$\begin{aligned} \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[q] \leq d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, q\gamma_*^{-}) &\leq d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, \gamma_*^{+}) + LM_3 \\ &\leq d_{\mathcal{X}_*}^{\mathfrak{w}_*}(o_*, go_*) + L(M_3 + M_2(2\epsilon + M_1 + 2)) \\ &\leq \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g] + L(3M_3 + M_2(2\epsilon + M_1 + 2)). \end{aligned}$$

On the other hand, (6.16) and (6.17) imply

$$k d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, \gamma_*^{+}) = d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, (\gamma_*^{(k)})^{+}) \leq d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, q^k\gamma_*^{-}) + LM_3,$$

and after dividing by k and letting k tend to infinity we get

$$d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, \gamma_*^{+}) \leq \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[q].$$

Combining this inequality with (6.11) and (6.18) gives us

$$\begin{aligned} \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g] &\leq d_{\mathcal{X}_*}^{\mathfrak{w}_*}(s'so_*, s'sgo_*) + 2LM_3 \\ &\leq d_{\mathcal{X}_*}^{\mathfrak{w}_*}(\gamma_*^{-}, \gamma_*^{+}) + L(2M_3 + M_2(2\epsilon + M_1 + 2)) \\ &\leq \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[q] + L(2M_3 + M_2(2\epsilon + M_1 + 2)), \end{aligned}$$

and we deduce

$$|\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g] - \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[\beta(\omega)]| = |\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g] - \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[q]| \leq L(3M_3 + M_2(2\epsilon + M_1 + 2)).$$

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Finally, if we define $\hat{g} = s'sg(s's)^{-1} \in [g]$, then by (6.12) and (6.13) we get

$$d_{\mathcal{X}}(\hat{g}o, \gamma^+) \leq d_{\mathcal{X}}(\hat{g}o, s'sgo) + \epsilon + M_1 \leq d_{\mathcal{X}}((s's)^{-1}o, o) + \epsilon + M_1 \leq 2\epsilon + M_1 + N.$$

In conclusion, the lemma follows with $C' = (L+1)(N + 3M_3 + M_2(2\epsilon + M_1 + 2))$. □

Now we prove Theorem 1.4.

Proof of Theorem 1.4. Given $\Gamma \in \mathfrak{G}$ and cubulations $\mathcal{X}$ and $\mathcal{X}_*$ of Γ, we know that $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$. Applying Proposition 5.7 we obtain a finite index subgroup $\overline{\Gamma} < \Gamma$ and a cubulation $\mathcal{Z}$ of Γ, so that the Γ-essential cores $\hat{\mathcal{X}}, \hat{\mathcal{X}}_*$ of $\mathcal{X}, \mathcal{X}_*$, respectively, are restriction quotients of $\mathcal{Z}$. Taking a further finite index subgroup if necessary, and by applying Lemma 5.3, we can assume that $\overline{\mathcal{X}} = \overline{\Gamma}\backslash\mathcal{X}$ is special, and we choose this group $\overline{\Gamma}$ as data for part i).

The graph $\mathcal{G}$ for part ii) is the underlying graph for the automaton $\mathcal{A}_{\overline{\Gamma},\phi}$ from Theorem 5.11, applied to the triplet $\Gamma, \hat{\mathcal{X}}, \mathcal{Z}$ which by construction satisfies Convention 5.10. We note that the labeling $\pi = \pi_\phi$ still has image in $S_{\overline{\mathcal{X}}}$ since $\overline{\Gamma}\backslash\hat{\mathcal{X}} \to \overline{\mathcal{X}}$ is a convex isometric embedding (so the set of oriented hyperplanes in $\overline{\Gamma}\backslash\hat{\mathcal{X}}$ injects into $S_{\overline{\mathcal{X}}}$).

The triplet $(\Gamma, \hat{\mathcal{X}}, \mathcal{Z})$ also satisfies Convention 6.3 with $\hat{\mathcal{X}}_* = \mathcal{Z}(\mathbb{W}_*)$ (here $\mathfrak{w} \equiv \mathfrak{w}_* \equiv 1$ are the trivial orthotope structures), and the function $\psi = \psi^{\mathfrak{w}_*}_{\hat{\mathcal{X}}_*}$ is given as in Definition 6.5. This completes the data in part iii), and for a path ω in $\mathcal{G}$, the associated loop $\overline{\gamma}_\omega$ in $\overline{\mathcal{X}}$ is the image under $\mathcal{X} \to \overline{\Gamma}\backslash\hat{\mathcal{X}} \subset \overline{\mathcal{X}}$ of any good representative $\gamma = \gamma_\omega$ as in Definition 6.4.

We are left to prove the claims (1)-(3) from the statement. Item (1) follows from Theorem 5.11 (2), and Item (2) follows from Lemma 6.7. Finally, Item (3) follows from Lemma 6.12 and Lemma 6.11. □

To end the subsection we prove Proposition 6.8.

Proof of Proposition 6.8. For each $s \in \mathbb{R}$ and $R > 0$ we consider the sums

$$\mathcal{P}(R,T,s) = \sum_{|\ell^{\mathfrak{w}}_{\mathcal{X}}[g] - T| \leq R} e^{-s\ell^{\mathfrak{w}_*}_{\mathcal{X}_*}[g]} \quad \text{and} \quad \mathcal{P}_{\mathcal{C}}(R,T,s) = \sum_{\tau \in P(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}}, R, T)} e^{-s\int_\tau \Phi}.$$

Since $l_\tau = \ell^{\mathfrak{w}}_{\mathcal{X}}[\beta(\tau)]$ for any closed orbit τ, by Lemmas 6.7 and 6.12 there exists a polynomial Q depending only on R such that $\mathcal{P}_{\mathcal{C}}(R,T,s) \leq Q(T)\mathcal{P}(R,T,s)$ for each $s \in \mathbb{R}$ and $T > 0$.

For an inequality in the other direction, for any $[g]$ we use Lemma 6.13 to find a path $\omega_{[g]} \in P(\Sigma^{\times}_{\mathcal{C}})$ and a representative $\hat{g}$ of $[g]$ satisfying

$$\min\{d_{\mathcal{X}}(\hat{g}o, \gamma^+_{\omega_{[g]}}), d_{\mathcal{X}}(\hat{g}^{-1}o, \gamma^+_{\omega_{[g]}})\} \leq C' \tag{6.19}$$

and

$$\max\{|\ell^{\mathfrak{w}}_{\mathcal{X}}[g] - \ell^{\mathfrak{w}}_{\mathcal{X}}[\beta(\omega_{[g]})]|, |\ell^{\mathfrak{w}_*}_{\mathcal{X}_*}[g] - \ell^{\mathfrak{w}_*}_{\mathcal{X}_*}[\beta(\omega_{[g]})]|\} \leq C'. \tag{6.20}$$

From (6.19) we get that the association $[g] \mapsto \omega_{[g]}$ is uniformly finite-to-1. We extend this association to $[g] \mapsto \omega_{[g]} \mapsto \tau_{[g]}$, where $\tau_{[g]} \in P(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}})$ is the periodic orbit corresponding to the path $\omega_{[g]}$. Since changing the initial vertex of a closed path in $P(\Sigma^{\times}_{\mathcal{C}})$ does not change the periodic orbit in $P(\Sigma^{r^{\mathfrak{w}}_{\mathcal{X}}}_{\mathcal{C}})$, the association $\omega_{[g]} \mapsto \tau_{[g]}$ is at most (linear in $\ell_{\mathcal{X}}[g]$)-to-1. But $\ell_{\mathcal{X}}[g]$ is comparable to $\ell^{\mathfrak{w}}_{\mathcal{X}}[g] = \ell^{\mathfrak{w}}_{\mathcal{X}}[\beta(\tau_{[g]})] = l_{\tau_{[g]}}$ (recall that $\mathcal{X}$ and $\mathcal{X}^{\mathfrak{w}}$ are quasi-isometric), and so from (6.20) we deduce that for each $s \in \mathbb{R}$ there is $C_s > 0$ such that $\mathcal{P}(R,T,s) \leq C'_s F(T)\mathcal{P}_{\mathcal{C}}(R + C', T, s)$ for each $T > 0$, where F is a degree 1 polynomial depending only on R.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

It follows that for any fixed R sufficiently large and for any $s \in \mathbb{R}$

$$\lim_{T\to\infty} \frac{1}{T} \log \mathcal{P}(R,T,s) = \lim_{T\to\infty} \frac{1}{T} \log\left(\sum_{\tau \in P(\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}}, R, T)} e^{-s\int_\tau \Phi} \right) = \mathrm{P}_{\mathcal{C}}(-s\Phi), \tag{6.21}$$

where $\mathrm{P}_{\mathcal{C}}(-s\Phi)$ is the pressure of the potential $-s\Phi$ on the suspension $(\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}}, \sigma^{r_{\mathcal{X}}^{\mathfrak{w}}})$.

Also, note that

$$\sum_{[g]\in \mathbf{conj}(\Gamma)} e^{-t\ell_{\mathcal{X}}^{\mathfrak{w}}[g] - s\ell_{\mathcal{X}_*^{\mathfrak{w}_*}}[g]} \le e^{R|t|} \sum_{T=1}^{\infty} \mathcal{P}(R,T,s) e^{-tT}$$

assuming the right-hand side of the above converges. Similarly we have

$$\sum_{T=1}^{\infty} \mathcal{P}(R,T,s) e^{-tT} \le 2Re^{R|t|} \sum_{[g]\in \mathbf{conj}(\Gamma)} e^{-t\ell_{\mathcal{X}}^{\mathfrak{w}}[g] - s\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}$$

when the right-hand side converges. We deduce that for each $s \in \mathbb{R}$ the series

$$\sum_{T=1}^{\infty} \mathcal{P}(R,T,s) e^{-tT} \quad \text{and} \quad \sum_{[g]\in \mathbf{conj}(\Gamma)} e^{-t\ell_{\mathcal{X}}^{\mathfrak{w}}[g] - s\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}$$

have the same abscissa of convergence as t varies. Hence by (6.21) we deduce

$$\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(s) = \mathrm{P}_{\mathcal{C}}(-s\Phi),$$

as desired. □

6.2. Analyticity and Large deviations for pairs of cubulations

In this subsection we prove Theorems 6.1 and 6.2. For a triplet $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ we always assume that it satisfies Convention 6.3, which is possible by Proposition 5.7. In particular, all the results and notations from this and the previous section are valid for this triplet. We first prove a large deviations principle that follows from Proposition 6.8.

Corollary 6.16. *Let $(\Gamma, \mathcal{X}, \mathcal{X}_*) \in \mathfrak{X}$ and $\mathfrak{w}, \mathfrak{w}_*$ be Γ-invariant orthotope structures on $\mathcal{X}, \mathcal{X}_*$, and let $\mathcal{L} : [\mathrm{Dil}(\mathcal{X}^{\mathfrak{w}}, \mathcal{X}_*^{\mathfrak{w}_*})^{-1}, \mathrm{Dil}(\mathcal{X}_*^{\mathfrak{w}_*}, \mathcal{X}^{\mathfrak{w}})] \to \mathbb{R}$ be the Legendre transform of $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}$. Then for any nonempty open set $U \subset \mathbb{R}$ and closed set $V \subset \mathbb{R}$ with $U \subset V$ we have that*

$$\begin{aligned} -\inf_{s\in U} \mathcal{L}(s) &\le \liminf_{T\to\infty} \frac{1}{T} \log\left(\frac{1}{\#\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)} \# \left\{ [g] \in \mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T) : \frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]} \in U \right\} \right) \\ &\le \limsup_{T\to\infty} \frac{1}{T} \log\left(\frac{1}{\#\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)} \# \left\{ [g] \in \mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T) : \frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]} \in V \right\} \right) \le -\inf_{s\in V} \mathcal{L}(s). \end{aligned}$$

In consequence, the limit

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) := \lim_{T\to\infty} \frac{1}{\#\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)} \sum_{[g]\in \mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)} \frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]}$$

exists and equals $-\theta'_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(0)$.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

Proof. Recall that for $T > R > 0$ and $s \in \mathbb{R}$ we defined

$$\mathcal{P}(R,T,s) = \sum_{|\ell_{\mathcal{X}}^{\mathfrak{w}}[g]-T|\leq R} e^{-s\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}$$

in the proof of Proposition 6.8 above. We saw during that proof that

$$\lim_{T\to\infty}\frac{1}{T}\log\mathcal{P}(R,T,s) = \theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(s).$$

It follows from the Gärtner-Ellis Theorem [29, Theorem 2.3.6] that the large deviations principle stated in this corollary holds but with $\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)$ replaced by

$$\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T,R) = \{[g] \in \mathbf{conj}\colon |\ell_{\mathcal{X}}^{\mathfrak{w}}[g] - T| < R\}$$

for any fixed $R > 0$ sufficiently large. It is then easy to check that this large deviations principle implies the one stated in the corollary.

By this large deviations principle we know that for any $\epsilon > 0$ the cardinality of the set

$$E_\epsilon(T) := \left\{[g] \in \mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T) : \left|\frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]} + \theta'_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(0)\right| > \epsilon\right\}$$

grows strictly exponentially slower than $\#\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)$ as $T \to \infty$, that is, the quotient $\#E_\epsilon(T)/\#\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)$ decays to 0 exponentially as $T \to \infty$. It is then standard to deduce that

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) := \lim_{T\to\infty}\frac{1}{\#\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)}\sum_{[g]\in\mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T)}\frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]}$$

exists and is equal to $-\theta'_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(0)$ as required. □

Proof of Theorem 6.1. We showed in Proposition 6.8 that $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(s)$ is equal to the pressure $\mathrm{P}_{\mathcal{C}}(-s\Phi)$ for any s. It follows that $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}$ is analytic, convex and decreasing (*see also* Rmark 2.10). Also, by Corollary 6.16 we know that the limit labeled $\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}})$ in the theorem exists. Further by comparing the exponential growth rates of both sides of the inequality

$$\#\left\{[g] \in \mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T) : \left|\frac{\ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g]}{\ell_{\mathcal{X}}^{\mathfrak{w}}[g]} - \tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}})\right| \leq \epsilon\right\} \leq \#\left\{[g] \in \mathfrak{C}_{\mathcal{X}^{\mathfrak{w}}}(T) : \ell_{\mathcal{X}_*}^{\mathfrak{w}_*}[g] \leq (\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) + \epsilon)T\right\}$$

we see that

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) \geq \frac{v_{\mathcal{X}^{\mathfrak{w}}}}{v_{\mathcal{X}_*^{\mathfrak{w}_*}}}.$$

Therefore to conclude the proof we need to check the equivalence of the statements (1), (2) and (3) when the action of Γ on $\mathcal{X}_*$ (and hence on $\mathcal{X}_*^{\mathfrak{w}_*}$) is proper. When this is the case we have $0 < v_{\mathcal{X}_*^{\mathfrak{w}_*}} < \infty$ and the Manhattan curve $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(s)$ is 0 at $s = v_{\mathcal{X}_*^{\mathfrak{w}_*}}$.

We will prove the implications (1) $\Rightarrow$ (3) $\Rightarrow$ (2) $\Rightarrow$ (1). Note that the implication (1) $\Rightarrow$ (3) follows easily from the facts that $\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) = -\theta'_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(0), \theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(0) = v_{\mathcal{X}^{\mathfrak{w}}}, \theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}(v_{\mathcal{X}_*^{\mathfrak{w}_*}}) = 0$ and $\theta_{\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}}$ is convex so has nonincreasing derivative. Also, the implication (2) $\Rightarrow$ (1) follows from the definition of the Manhattan curve. Hence we just need to prove the implication (3) $\Rightarrow$ (2).

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

To do so we note that

$$\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) = \int_{\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}}} \Phi \, dm$$

where m is the measure of maximal entropy for $(\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}}, \sigma^{r_{\mathcal{X}}^{\mathfrak{w}}})$. However we saw in Subsection 2.5 that

$$\int_{\Sigma_{\mathcal{C}}^{r_{\mathcal{X}}^{\mathfrak{w}}}} \Phi \, dm = \frac{\int_{\Sigma_{\mathcal{C}}} \psi_{\mathcal{X}_*}^{\mathfrak{w}_*} \, d\mu_1}{\int_{\Sigma_{\mathcal{C}}} r_{\mathcal{X}}^{\mathfrak{w}} \, d\mu_1}$$

where μ_1 is the equilibrium state of $-\delta_{r_{\mathcal{X}}^{\mathfrak{w}}} r_{\mathcal{X}}^{\mathfrak{w}}$ on $\Sigma_{\mathcal{C}}$. To simplify notation going forward we will also write $r = r_{\mathcal{X}}^{\mathfrak{w}}$, $\psi = \psi_{\mathcal{X}_*}^{\mathfrak{w}_*}$ and μ_2 for the equilibrium state of $-\delta_{\psi_{\mathcal{X}_*}^{\mathfrak{w}_*}} \psi_{\mathcal{X}_*}^{\mathfrak{w}_*}$ on $\Sigma_{\mathcal{C}}$. We now note that by Proposition 6.8 we have that

$$\mathrm{P}_{\mathcal{C}}(-v_{\mathcal{X}^{\mathfrak{w}}} r_{\mathcal{X}}^{\mathfrak{w}}) = \mathrm{P}_{\mathcal{C}}(-v_{\mathcal{X}_*^{\mathfrak{w}_*}} \psi_{\mathcal{X}_*}^{\mathfrak{w}_*}) = 0.$$

Here the pressures are the pressures of the potentials over the subshift (not suspension). Hence the inequality $\tau(\mathcal{X}_*^{\mathfrak{w}_*}/\mathcal{X}^{\mathfrak{w}}) \geq v_{\mathcal{X}^{\mathfrak{w}}}/v_{\mathcal{X}_*^{\mathfrak{w}_*}}$ can be rewritten as

$$\frac{h_{\mu_2}(\sigma)}{\int_{\Sigma_{\mathcal{C}}} \psi \, d\mu_2} \geq \frac{h_{\mu_1}(\sigma)}{\int_{\Sigma_{\mathcal{C}}} \psi \, d\mu_1}$$

where $h_{\mu_1}(\sigma), h_{\mu_2}(\sigma)$ are the entropies of μ_1, μ_2 over the component $\mathcal{C}$. This inequality is true by the variational principle. Furthermore this inequality is a strict equality unless r and ψ are cohomologous. This implies by Lemmas 6.7 and 6.13 that there exist $\Lambda, C > 0$ such that

$$|\ell_{\mathcal{X}^{\mathfrak{w}}}^{\mathfrak{w}}[g] - \Lambda \ell_{\mathcal{X}_*^{\mathfrak{w}_*}}^{\mathfrak{w}_*}[g]| < C$$

for all $[g] \in \mathbf{conj}(\Gamma)$. This can only happen if (2) holds. □

Proof of Theorem 6.2. Let $\psi = \psi_{\mathcal{X}_*}^{\mathfrak{w}_*} : \Sigma \to \mathbb{Z}$ be the potential associated to the constant orthotope structure $\mathfrak{w}_* \equiv 1$. Let $\mathcal{L} : [\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X})] \to \mathbb{R}$ be the Legendre transform of $\theta_{\mathcal{X}_*/\mathcal{X}}$, which by Remark 6.9 equals the Legendre transform of $s \mapsto \mathrm{P}_{\mathcal{C}}(-s\psi)$ for $\mathcal{C}$ any maximal recurrent component of $\mathcal{G}_\phi$. Hence $\mathcal{L}$ is analytic.

From our large devation principle in Corollary 6.16 we have that

$$\limsup_{T\to\infty} \frac{1}{T} \log\left(\#\left\{[g] \in \mathbf{conj} : \ell_{\mathcal{X}}[g] < T, |\ell_{\mathcal{X}_*}[g] - \eta \ell_{\mathcal{X}}[g]| < \frac{C}{T}\right\}\right) \leq v_{\mathcal{X}} - \mathcal{L}(\eta)$$

for all $\eta \in (\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X}))$.

We now prove the lower bound. Fix a maximal component $\mathcal{C}$. By Lemma 6.7 and Lemma 6.12 there exists a polynomial Q such that for any $C > 0$ and $\eta \in (\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X}))$

$$\#\left\{\omega \in P_n(\Sigma_{\mathcal{C}}^{\times}) : \left|\frac{\psi^n(\omega)}{n} - \eta\right| < \frac{C}{n}\right\} \leq Q(n) \cdot \#\left\{[g] \in \mathbf{conj} : \ell_{\mathcal{X}}[g] \leq n, \left|\frac{\ell_{\mathcal{X}_*}[g]}{\ell_{\mathcal{X}}[g]} - \eta\right| < \frac{C}{n}\right\}$$

where ψ is the potential from Definition 6.5. However, by Theorem 3.2 (and Remark 3.3 as $\mathcal{C}$ may only be transitive) we have that

$$\limsup_{n\to\infty} \frac{1}{n} \log\left(\#\left\{\omega \in P_n(\Sigma_{\mathcal{C}}^{\times}) : \left|\frac{\psi^n(\omega)}{n} - \eta\right| < \frac{C}{n}\right\}\right) = h - \mathcal{I}(\eta)$$

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

where $\mathcal{I}$ is the Legendre transform of the map $s \mapsto \mathrm{P}_{\mathcal{C}}(-s\psi)$ and h is the topological entropy of the subshift $(\Sigma_{\mathcal{C}}, \sigma)$. However, as we saw above, $\mathcal{I}$ is precisely $\mathcal{L}$ and further by Lemma 6.11 we have $h = v_{\mathcal{X}}$. Hence we deduce that

$$\limsup_{T\to\infty} \frac{1}{T} \log \#\left\{[g] \in \mathbf{conj} : \ell_{\mathcal{X}}[g] < T, |\ell_{\mathcal{X}_*}[g] - \eta\ell_{\mathcal{X}}[g]| < \frac{C}{T}\right\} \geq v_{\mathcal{X}} - \mathcal{L}(\eta)$$

for each $\eta \in (\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X}))$. We have shown that the limit supremum in the statement of the theorem is equal to $v_{\mathcal{X}} - \mathcal{L}$ as required.

To conclude the proof we need to explain the additional conditions mentioned in the theorem. In particular we need to show that

$$0 < v_{\mathcal{X}} - \mathcal{L}(\eta) \leq v_{\mathcal{X}} \text{ for all } \eta \in (\mathrm{Dil}(\mathcal{X}, \mathcal{X}_*)^{-1}, \mathrm{Dil}(\mathcal{X}_*, \mathcal{X}))$$

and that the upper bound inequality is an equality if and only if $\eta = \tau(\mathcal{X}_*/\mathcal{X})$. All of these properties follow from the definition of $\mathcal{L}$ and the fact that $s \mapsto \mathrm{P}_{\mathcal{C}}(-s\psi)$ is strictly convex. □

Appendix A. Convex-cocompact subgroups of cubulable relatively hyperbolic groups

In this appendix we prove Proposition 5.4. First, we recall the statement.

Proposition A.1. *Let* Γ *be a relatively hyperbolic group acting properly and cocompactly on the* CAT(0) *cube complex* $\mathcal{X}$. *Then the following are equivalent for a subgroup* $H < \Gamma$.

(1) H *is convex-cocompact for the action on* $\mathcal{X}$.
(2) H *is relatively quasiconvex and* $H \cap P$ *is convex-cocompact for the action of* Γ *on* $\mathcal{X}$ *for any maximal parabolic subgroup* $P < \Gamma$.

Proof. Under these assumptions Γ is finitely generated, so fix $S \subset \Gamma$ a finite symmetric generating set and a Γ-equivariant quasi-isometry $\phi : \mathcal{X} \to \mathrm{Cay}(\Gamma, S)$. We also fix a vertex $x_0 \in \mathcal{X}$ such that $\phi(x_0) = o$ is the identity element in Γ. Let $\mathbb{P}$ be a complete collection of representatives of conjugacy classes of maximal parabolic subgroups in Γ, and let $\mathcal{P} = (\bigcup \mathbb{P})\backslash\{o\}$. We let d_S denote the (graph) word metric on $\mathrm{Cay}(\Gamma, S)$ and let $H < \Gamma$ be any subgroup.

If H is convex-cocompact, then it is undistorted, hence finitely generated and relatively quasiconvex by [50, Theorem 1.5]. Also, any maximal parabolic subgroup $P < \Gamma$ is convex-cocompact by [76, Theorem 1.1], and hence $H \cap P$ is also convex-cocompact by [71, Lemma 2.14 & Lemma 2.15]. This proves the implication (1) $\Rightarrow$ (2).

The implication (2) $\Rightarrow$ (1) is more involved, and for its proof we adopt the following convention. If γ' is a parameterized curve and $x = \gamma'_{t^-}, y = \gamma'_{t^+}$ belong to γ' with $t^- \leq t^+$, then $\gamma'|_{[x,y]} = \gamma'|_{[y,x]}$ is a set of points of form γ'_t, with $t^- \leq t \leq t^+$ (if there is more than one option for $t^\pm$, we consider any of them).

By [39, Lemma 4.3] it is enough to prove the following: there exists $K > 0$ such that if $\overline{\gamma} \subset \mathcal{X}$ is a (continuous) combinatorial geodesic with endpoints in Hx_0, then $\overline{\gamma} \subset N_K(Hx_0)$.

To find such K, consider constants L, C such that the image under ϕ of any combinatorial geodesic $\overline{\gamma}$ in $\mathcal{X}$ is at Hausdorff distance at most L from an L-Lipschitz (L, C)-quasigeodesic $\gamma : [0, \ell] \to \mathrm{Cay}(\Gamma, S)$ with same endpoints as $\phi(\overline{\gamma})$ (see e.g., [11, Proposition 8.3.4]).

Let $\overline{\gamma} \subset \mathcal{X}$ be a geodesic with endpoints in Hx_0 and let $\gamma = \gamma([0, \ell]) \subset \mathrm{Cay}(\Gamma, S)$ be as above, so that the endpoints of γ belong to H. Also, let $\hat{c}$ be a geodesic in $\mathrm{Cay}(\Gamma, S \cup \mathcal{P})$ with same endpoints as γ and let $\hat{c}_0, \ldots, \hat{c}_n$ be the (ordered) vertex set of $\hat{c}$. We define $I = \{j_0 < j_i < \cdots < j_k\}$ to be the set of all $0 \leq i \leq n-1$ such that $\hat{c}_{i+1}^{-1}\hat{c}_i \in \mathcal{P}$.

By quasiconvexity of H, there exists κ (independent of γ) and $h_i \in H$ such that $d_S(\hat{c}_i, h_i) \leq \kappa$ for all $0 \leq i \leq n$, see e.g., [50, Def. 6.10]. Also, by [50, Lemma 8.8] there exists A_0 depending only on L, C such that for any $0 \leq i \leq n$ there exists $c_i = \gamma_{t_i} \in \gamma$ satisfying $d_S(c_i, \hat{c}_i) \leq A_0$, for which we assume

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

$c_0 = \hat{c}_0$ and $c_n = \hat{c}_n$. Up to increasing A_0 (only in terms of L, C), we can always assume that c_i is a group element.

Since Γ is finitely generated, by [50, Proposition 9.4] there exists B_0 depending only on A_0 and κ (so only on L, C) such that if $g_1, g_2 \in \Gamma$ satisfy $|g_1|_S, |g_2|_S \leq A_0 + \kappa$, then

$$N_{A_0+\kappa}(g_1 H) \cap N_{A_0+\kappa}(g_2 P) \subset N_{B_0}(g_1 H g_1^{-1} \cap g_2 P g_2^{-1}) \tag{A.1}$$

for any $P \in \mathbb{P}$, where the neighborhoods are considered in $\mathrm{Cay}(\Gamma, S)$.

Let $\widetilde{p}$ be a geodesic lift of $\hat{c}$ to $\mathrm{Cay}(\Gamma, S)$. That is, $\widetilde{p}$ is obtained from $\hat{c}$ by replacing each edge corresponding to an element of $\mathcal{P}$ by a geodesic in $\mathrm{Cay}(\Gamma, S)$ with the same endpoints. For a point $x \in \gamma$ we distinguish two cases.

Case 1: $x \in \gamma|_{[c_{j_i+1}, c_{j_{i+1}}]}$ for some $j_i \in I$ (with the convention that $j_{-1} = -1$ and $j_{k+1} = n$). Consider geodesic paths $[c_{j_i+1}, \hat{c}_{j_i+1}]$ and $[c_{j_{i+1}}, \hat{c}_{j_{i+1}}]$ in $\mathrm{Cay}(\Gamma, S)$, and the quasigeodesic triangle with sides

$$\ell_1 = [\hat{c}_{j_i+1}, c_{j_i+1}] \cup \gamma|_{[c_{j_i+1}, x]}, \quad \ell_2 = \gamma|_{[x, c_{j_{i+1}}]} \cup [c_{j_{i+1}}, \hat{c}_{j_{i+1}}], \quad \ell_3 = \widetilde{p}|_{[c_{j_{i+1}}, c_{j_i+1}]}.$$

We also set

$$\ell_1^- = \hat{c}_{j_i+1}, \ell_1^+ = x, \quad \ell_2^- = \hat{x}, \ell_2^+ = \hat{c}_{j_{i+1}}, \quad \text{and} \quad \ell_3^- = \hat{c}_{j_{i+1}}, \ell_3^+ = \hat{c}_{j_i+1}.$$

Note that ℓ_1, ℓ_2, ℓ_3 are Lipschitz quasigeodesics with constants depending only on L, C and A_0 (hence only on L, C). Then by [30, Lemma 8.19] there exists R depending on L, C such that either:

- there exists $z \in \mathrm{Cay}(\Gamma, S)$ with $d_S(z, \ell_i) \leq R$ for $i = 1, 2, 3$; or,
- there exist $g \in \Gamma$ and $P \in \mathbb{P}$ such that $d_S(gP, \ell_i) \leq R$ for $i = 1, 2, 3$.

In the first subcase, let $u_i \in \ell_i$ be such that $d_S(z, u_i) \leq R$. Then $d_S(u_a, u_b) \leq 2R$ for all $1 \leq a, b \leq 3$, and since γ is (L, C)-quasigeodesic and $x \in \gamma|_{[u_1, u_2]}$, we have that $d_S(x, \ell_3) \leq d_S(x, u_3) \leq d_S(x, u_1) + d_S(u_1, u_3)$ is bounded above in terms of L, C and R (thus only in terms of L, C).

In the second subcase, by [30, Lemma 8.15] we can find M and $\mathfrak{d}$ depending only on L, C and R (so only on L, C) and points $u_i^-, u_i^+ \in \ell_i$ for $i = 1, 2, 3$ that satisfy:

- $d_S(u_i^\pm, gP) \leq M$; and,
- $\mathrm{diam}(\ell_i|_{[\ell_i^\pm, u_i^\pm]} \cap \overline{N}_M(gP)) \leq \mathfrak{d}$.

Take $v_i^\pm \in gP$ such that $d_S(u_i^\pm, v_i^\pm) \leq M$. Then by the definition of I and after considering vertices in $\hat{c}$ that are closest to $u_3^\pm$ in $\mathrm{Cay}(\Gamma, S)$ we get

$$d_S(u_3^+, u_3^-) \leq d_{S \cup \mathcal{P}}(v_3^-, v_3^+) + 2(1 + M) + 2 \leq 5 + 2M.$$

Also, [30, Lemma 8.14] implies the existence of D_1 depending only on L, C, M and $\mathfrak{d}$ (so only on L, C) with $d_S(u_i^+, u_{i+1}^-) \leq D_1$ (mod 3) for all i. In particular,

$$d_S(u_1^-, u_2^+) \leq d_S(u_1^-, u_3^+) + d_S(u_3^+, u_3^-) + d_S(u_3^-, u_2^+),$$

and as in the first subcase we conclude that x belongs to a neighborhood of ℓ_3 depending only on L and C.

In both subcases, we deduce that $d_S(x, \ell_3)$ is bounded in terms of L, C, and since ℓ_3 is contained in a neighborhood of H depending only on κ, we have that

$$d_S(x, H) \leq K_0 \tag{A.2}$$

for some K_0 depending only on L, C and κ (hence only in terms of L, C).

Case 2: $x \in \gamma|_{[c_j, c_{j+1}]}$ for some $j \in I$. Suppose $\hat{c}_j^{-1} \hat{c}_{j+1} = p \in P$ for $P \in \mathbb{P}$. Then

$$d_S(c_j^{-1} c_{j+1}, (c_j^{-1} h_j) H) \leq d_S(c_j^{-1} c_{j+1}, c_j^{-1} h_{j+1}) = d_S(c_{j+1}, h_{j+1}) \leq A_0 + \kappa,$$

and

$$d_S(c_j^{-1}c_{j+1},(c_j^{-1}\hat{c}_j)P) \leq d_S(c_j^{-1}c_{j+1},c_j^{-1}\hat{c}_jp) = d_S(c_{j+1},\hat{c}_{j+1}) \leq A_0.$$

Since $\max\{|c_j^{-1}h_j|_S,|c_j^{-1}\hat{c}_j^{-1}|_S\} \leq A_0+\kappa$, by (A.1) we conclude

$$d_S(c_j^{-1}c_{j+1},(c_j^{-1}h_j)H(c_j^{-1}h_j)^{-1}\cap(c_j^{-1}\hat{c}_j)P(c_j^{-1}\hat{c}_j)^{-1}) \leq B_0. \tag{A.3}$$

Note that any point $x \in \gamma$ satisfies the assumptions of one of the two cases above. Indeed, for $x=\gamma_t \in \gamma$, let I_- be the set of all the $j \in I$ such that c_j is not of the form $\gamma_{t'}$ with $t'>t$. Suppose first that I_- is nonempty and let j be its maximal element. If x does not satisfy Case 2, then c_{j+1} does not belong to $\gamma|_{[x,\gamma_\ell]}$. But if $j=j_i<j_k$, then $j_{i+1}\notin I_-$, so that $x\in\gamma|_{[c_{j_i+1},c_{j_{i+1}}]}$ and x satisfies Case 1.

Also, if $j=j_k$, then $x\in\gamma|_{[c_{j_i+1},c_n]}$ and x also satisfies Case 1. Therefore, we can assume that I_- is empty. But if I is nonempty then $x\in\gamma|_{[c_0,c_{j_1}]}$ and x satisfies Case 1, and if I is empty then $x\in\gamma|_{[c_0,c_n]}$ and x also satisfies Case 1.

Now, take $\overline{x}\in\overline{\gamma}$ and let $x\in\gamma$ within r from $\phi(\overline{x})$ in $\mathrm{Cay}(\Gamma,S)$, where r is independent of $\overline{x}$ and $\overline{\gamma}$. If x satisfies Case 1, by (A.2) we conclude that $d_{\mathcal{X}}(\overline{x},Hx_0)\leq K_1$ for K_1 a constant independent of $\overline{x}$ and $\overline{\gamma}$.

If x satisfies Case 2, suppose that $x\in\gamma|_{[c_j,c_{j+1}]}$ for $j\in I$. Then by (A.3) there exist vertices $\overline{x}^-,\overline{x}^+\subset\overline{\gamma}$ satisfying $d_{\mathcal{X}}(c_jx_0,\overline{x}^-)\leq\hat{r}$ and $d_{\mathcal{X}}(c_{j+1}x_0,\overline{x}^+)\leq\hat{r}$, where $\hat{r}$ depends only on ϕ and L,C.

Let F be the set of pairs $\alpha,\beta\in\Gamma$ satisfying $|\alpha|_S,|\beta|_S\leq A_0+\kappa$. By our assumption and [76, Theorem 1.1] we can find a convex core $Z_{\alpha,\beta}\subset\mathcal{X}$ for the group $\alpha H\alpha^{-1}\cap\beta P\beta^{-1}$ that contains the $\hat{r}$-neighborhood of x_0. By cocompactness, we can find $K_2>0$ such that

$$Z_{\alpha,\beta}\subset N_{K_2}((\alpha H\alpha^{-1}\cap\beta P\beta^{-1})x_0)\subset N_{K_2}((\alpha H\alpha^{-1})x_0)$$

for all $(\alpha,\beta)\in F$. Note that K_2 is independent of $\overline{\gamma}$. In particular we have

$$\overline{x}\in c_jZ_{c_j^{-1}h_j,c_j^{-1}\hat{c}_j}\subset c_jN_{K_3}(c_j^{-1}H(h_j^{-1}c_j)x_0)\subset N_{K_2}(Hx_0),$$

where $K_3:=K_2+\max\{d_{\mathcal{X}}(\alpha x_0,x_0)\colon|\alpha|_S\leq A_0+\kappa\}$ is independent of $\overline{x}$ and $\overline{\gamma}$. In conclusion, $\overline{\gamma}\subset N_K(Hx_0)$ for $K:=\max\{K_1,K_3\}$, and the implication (2) $\Rightarrow$ (1) follows. □

Acknowledgments. Both authors are grateful to Richard Sharp for his comments and suggestions on a preliminary version of this paper. The second author would like to thank the Max Planck Institut für Mathematik for its hospitality and financial support.

Competing interest. The authors have no competing interests to declare.

References

[1] I. Agol, 'The virtual Haken conjecture'. With an appendix by I. Agol, D. Groves and J. F. Manning. *Doc. Math.* 18 (2013), 1045–1087.

[2] T. Aougab, M. Clay and Y. Rieck, 'Thermodynamic metrics on outer space', *Ergodic Theory Dynam. Systems* **43**(3) (2023), 729–793.

[3] N. Bergeron and D. T. Wise, 'A boundary criterion for cubulation', *Amer. J. Math.* **134**(3) (2012), 843–859.

[4] J. Beyrer and E. Fioravanti, 'Cross ratios and cubulations of hyperbolic groups', *Math. Ann.* **384**(3–4) (2022), 1547–1592.

[5] J. Beyrer and E. Fioravanti, 'Cross-ratios on CAT(0) cube complexes and marked length-spectrum rigidity', *J. Lond. Math. Soc. (2)* **104**(5) (2021), 1973–2015.

[6] H. Bray, R. Canary and L.-Y. Kao, 'Pressure metrics for deformation spaces of quasifuchsian groups with parabolics', *Algebr. Geom. Topol.* **23**(8) (2023), 3615–3653.

[7] H. Bray, R. Canary, L.-Y. Kao and G. Martone, 'Counting, equidistribution and entropy gaps at infinity with applications to cusped Hitchin representations', *J. Reine Angew. Math.* **791** (2022), 1–51.

[8] C. Bregman, R. Charney and K. Vogtmann, 'Outer space for RAAGs', *Duke Math. J.* **172**(6) (2023), 1033–1108.

[9] M. Bridgeman, R. Canary, F. Labourie and A. Sambarino, 'The pressure metric for Anosov representations', *Geom. Funct. Anal.* **25**(4) (2015), 1089–1179.

[10] M. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren der Mathematischen Wissenschaften, vol. 319 (Springer-Verlag, Berlin, 1999).

[11] D. Burago, Y. Burago and S. Ivanov, *A Course in Metric Geometry*. Graduate Studies in Mathematics, vol. 33 (American Mathematical Society, Providence, RI, 2001).

[12] M. Burger, 'Intersection, the Manhattan curve, and Patterson-Sullivan theory in rank 2', *Int. Math. Res. Not.* 1993(7), 217–225.

[13] R. Bowen and C. Series, 'Markov maps associated with Fuchsian groups', *Inst. Hautes Études Sci. Publ. Math.* **50** (1979), 153–170.

[14] D. Calegari and K. Fujiwara, 'Combable functions, quasimorphisms, and the central limit theorem', *Ergodic Theory Dynam. Systems* **30**(5) (2010), 1343–1369.

[15] J. W. Cannon, 'The combinatorial structure of cocompact discrete hyperbolic groups', *Geom. Dedicata* **16**(2) (1984), 123–148.

[16] S. Cantrell, 'Mixing of the Mineyev flow, orbital counting and Poincaré series for strongly hyperbolic metrics', *Math. Ann.* **392** (2025), 1253–1288.

[17] S. Cantrell and E. Reyes, Marked length spectrum rigidity from rigidity on subsets. arXiv:2304.13209, 2023.

[18] S. Cantrell andR, 'Tanaka, Invariant measures of the topological flow and measures at infinity on hyperbolic groups', *J. Mod. Dyn.* **20** (2024), 215–274.

[19] S. Cantrell and R. Tanaka, 'The Manhattan curve, ergodic theory of topological flows and rigidity', To appear in *Geom. Topol.*, arXiv:2104.13451, 2021.

[20] P.-E. Caprace and M. Sageev, 'Rank rigidity for CAT(0) cube complexes', *Geom. Funct. Anal.* **21** (2011), 851–891.

[21] R. Charney and M. W. Davis, 'The K(π,1)-problem for hyperplane complements associated to infinite reflection groups', *J. Amer. Math. Soc.* **8**(3) (1995), 597–627.

[22] R. Charney, N. Stambaugh and K. Vogtmann, 'Outer space for untwisted automorphisms of right-angled Artin groups', *Geom. Topol.* **21** (2017), 1131–1178.

[23] V. Chepoi, 'Graphs of some CAT(0) complexes', *Adv. Appl. Math.* **24**(2) (2000), 125–179.

[24] D. Cooper and D. Futer, 'Ubiquitous quasi-Fuchsian surfaces in cusped hyperbolic 3–manifolds', *Geom. Topol.* **23**(1) (2019), 241–298.

[25] M. Culler and K. Vogtmann, 'Moduli of graphs and automorphisms of free groups', *Invent. Math.* **84**(1) (1986), 91–119.

[26] F. Dahmani, D. Futer and D. T. Wise, 'Growth of quasiconvex subgroups', *Math. Proc. Cambridge Philos. Soc.* **167**(3) (2019), 505–530.

[27] X. Dai and G. Martone, 'Correlation of the renormalized Hilbert length for convex projective surfaces', *Ergodic Theory Dynam. Systems* **43** (2023), 2938–2973.

[28] F. Dal'bo, 'Remarques sur le spectre des longueurs d'une surface et comptages', *Bol. Soc. Brasil. Mat. (N.S.)* **30**(2) (1999), 199–221.

[29] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. Applications of Mathematics (New York), vol. 38, 2nd ed. (Springer-Verlag, New York, 1998).

[30] C. Druţu and M. Sapir, 'Tree-graded spaces and asymptotic cones of groups'. With an appendix by D. V. Osin and M. Sapir. *Topology* 44 (2005), 959–1058.

[31] V. Erlandsson, 'A remark on the word length in surface groups'. *Trans. Amer. Math. Soc.* **327** (2019), 441–455.

[32] D. B. A. Epstein, J. W. Cannon, D. Holt, S. Levy, M. Paterson and W. Thurston, *Word Processing in Groups* (Jones and Bartlett Publishers, Boston, MA, 1992).

[33] B. Farb and D. Margalit, *A Primer on Mapping Class Groups*. Princeton Mathematical Series, vol. 49 (Princeton University Press, 2012).

[34] E. Fioravanti, 'On automorphisms and splittings of special groups', *Compos. Math.* **159**(2) (2023), 232–305.

[35] E. Fioravanti, I. Levcovitz and M. Sageev, 'Coarse cubical rigidity', *J. Topol.* **17**(3) (2024), Paper No. e12353, 50.

[36] I. Gekhtman, S. J. Taylor and G. Tiozzo, 'Central limit theorems for counting measures in coarse negative curvature', *Compos. Math.* **158**(10) (2022), 1980–2013.

[37] I. Gekhtman, S. J. Taylor and G. Tiozzo, 'Counting loxodromics for hyperbolic actions', *J. Topol.* **11**(2) (2018), 379–419.

[38] I. Gekhtman, S. J. Taylor and G. Tiozzo, 'Counting problems in graph products and relatively hyperbolic groups', *Israel J. Math.* **237** (2020), 311–371.

[39] A. Genevois, 'Hyperbolicities in CAT(0) cube complexes', *Enseign. Math.* **65**(1/2) (2019), 33–100.

[40] E. Godelle and L. Paris, 'K(π,1) and word problems for infinite type Artin-Tits groups, and applications to virtual braid groups', *Math. Z.* **272** (2012), 1339–1364.

[41] D. Groves and J. F. Manning, 'Specializing cubulated relatively hyperbolic groups', *J. Topol.* **15** (2022), 398–442.

[42] Y. Guivarc'h and J. Hardy, 'Théorèmes limites pour une classe de chaines de Markov et applications aux difféomorphismes d'Anosov', *Ann. Inst. Henri Poincaré Probab. Stat.* **24**(1) (1988), 73–98.

[43] M. F. Hagen and P. Przytycki, 'Cocompactly cubulated graph manifolds', *Israel J. Math.* **207**(1) (2015), 377–394.

[44] M. F. Hagen and D. T. Wise, 'Cubulating hyperbolic free-by-cyclic groups: the general case', *Geom. Funct. Anal.* **25** (2015), 134–179.

[45] M. F. Hagen and D. T. Wise, 'Cubulating hyperbolic free-by-cyclic groups: the irreducible case', *Duke Math. J.* **165**(9) (2016), 1753–1813.

[46] F. Haglund, 'Isometries of CAT(0) cube complexes are semi-simple', *Ann. Math. Québec* (2021). https://doi.org/10.1007/s40316-021-00186-2.

https://doi.org/10.1017/fms.2025.10094 Published online by Cambridge University Press

[47] F. Haglund and D. T. Wise, 'Special cube complexes', *Geom. Funct. Anal.* **17**(5) (2008), 1551–1620.
[48] Y. M. He, H. Lee and I. Park, Pressure metrics in geometry and dynamics. arXiv:2407.18441, 2024.
[49] S. Hermiller and J. Meier, 'Algorithms and geometry for graph products of groups', *J. Algebra* **171**(1) (1995), 230–257.
[50] G. C. Hruska, 'Relative hyperbolicity and relative quasiconvexity for countable groups', *Algebr. Geom. Topol.* **10**(3) (2010), 1807–1856.
[51] A. Hurwitz, 'Über die angenäherte Darstellung der Irrationalzahlen durch rationale Brüche', *Math. Ann.* **39**(2) (1891), 279–284.
[52] J. Kahn and V. Markovic, 'Immersing almost geodesic surfaces in a closed hyperbolic three manifold', *Ann. Math. (2)* **175**(3) (2012), 1127–1190.
[53] L.-Y. Kao, 'Entropy rigidity, pressure metric, and immersed surfaces in hyperbolic 3-Manifolds'. In *Thermodynamic Formalism: CIRM Jean-Morlet Chair, Fall 2019* (Springer, International Publishing, 2021), pp. 351–393.
[54] L.-Y. Kao, 'Manhattan curves for hyperbolic surfaces with cusps', *Ergodic Theory Dynam. Systems* **40**(7) (2020), 1843–1874.
[55] L.-Y. Kao, 'Pressure metrics and Manhattan curves for Teichmüller spaces of punctured surfaces', *Israel J. Math.* **240**(2) (2020), 567–602.
[56] Y. Kifer, 'Large deviations, averaging and periodic orbits of dynamical systems', *Comm. Math. Phys.* **162** (1994), 33–46.
[57] J. Lauer and D. T. Wise, 'Cubulating one-relator groups with torsion', *Math. Proc. Cambridge Philos. Soc.* **155**(3) (2013), 411–429.
[58] J. Li and D. T. Wise, 'No growth-gaps for special cube complexes', *Groups Geom. Dyn.* **14**(1) (2020), 117–135.
[59] B. Marcus and S. Tuncel, 'Entropy at a weight-per-symbol and embeddings of Markov chains', *Invent. Math.* **102** (1990), 235–266.
[60] A. Martin and M. Steenbock, 'A combination theorem for cubulation in small cancellation theory over free products', *Ann. Inst. Fourier (Grenoble)* **67** (2017), 1613–1670.
[61] C. McMullen, 'Thermodynamics, dimension and the Weil–Petersson metric', *Invent. Math.* **173** (2008), 365–425.
[62] G. A. Niblo and L. D. Reeves, 'Coxeter groups act on CAT(0) cube complexes', *J. Group Theory* **6**(3) (2003), 399–413.
[63] T. Odrzygóźdź, 'Cubulating random groups in the square model', *Israel J. Math.* **227**(2) (2018), 623–661.
[64] Y. Ollivier and D. T. Wise, 'Cubulating random groups at density less than 1/6', *Trans. Amer. Math. Soc.* **363**(9) (2011), 4701–4733.
[65] W. Parry, 'Intrinsic Markov chains, *Trans. Amer. Math. Soc.* **112** (1964), 55–66.
[66] W. Parry and M. Pollicott, 'Zeta functions and the periodic orbit structure of hyperbolic dynamics', *Astérisque* **187–188** (1990), 1–268.
[67] M. Pollicott and R. Sharp, 'A Weil–Petersson metric on spaces of metric graphs', *Geom. Dedicata* **172** (2014), 229–244.
[68] M. Pollicott and R. Sharp, 'Large deviations, fluctuations and shrinking intervals', *Comm. Math. Phys.* **290** (2009), 321–334.
[69] P. Przytycki and D. T. Wise, 'Graph manifolds with boundary are virtually special', *J. Topol.* **7**(2) (2014), 419–435.
[70] P. Przytycki and D. T. Wise, 'Mixed 3-manifolds are virtually special', *J. Amer. Math. Soc.* **31**(2) (2018), 319–347.
[71] E. Reyes, 'On cubulated relatively hyperbolic groups', *Geom. Topol.* **27** (2023), 575–640.
[72] M. A. Roller, *Poc sets, median algebras and group actions. An extended study of Dunwoody's construction and Sageev's theorem.* arXiv:1607.07747; University of Southampton preprint, 1998.
[73] J. Rousseau-Egele, 'Un théoreme de la limite locale pour une classe de transformations dilatantes et monotones par morceaux', *Ann. Probab.* **11** (1983), 772–788.
[74] M. Sageev, 'CAT(0) cube complexes and groups'. In *Geometric Group Theory*, IAS/Park City Math. Ser., vol. 21 (American Mathematical Society, Providence, RI, 2014), pp. 7–54.
[75] M. Sageev, 'Ends of group pairs and non-positively curved cube complexes', *Proc. Lond. Math. Soc.* **71**(3) (1995), 585–617.
[76] M. Sageev and D. T. Wise, 'Cores for quasiconvex actions', *Proc. Amer. Math. Soc.* **143** (2015), 2731–2741.
[77] A. Sale and T. Susse, 'Outer automorphism groups of right-angled Coxeter groups are either large or virtually abelian', *Trans. Amer. Math. Soc.* **372**(11) (2019), 7785–7803.
[78] R. Schwartz and R. Sharp, 'The correlation of length spectra of two hyperbolic surfaces', *Comm. Math. Phys.* **154** (1993), 423–430.
[79] R. Sharp, 'Comparing length functions on free groups'. In *Spectrum and Dynamics*, CRM Conf. Proc. Lecture Notes, vol. 52 (Amer. Math. Soc., Providence, RI, 2010), pp. 185–207.
[80] R. Sharp, 'Distortion and entropy for automorphisms of free groups', *Discrete Contin. Dyn. Syst.* **26** (2010), 347–363.
[81] R. Sharp, 'The Manhattan curve and the correlation of length spectra on hyperbolic surfaces', *Math. Z.* **228** (1998), 745–750.
[82] K. Sigmund, 'Generic properties of invariant measures for Axiom A diffeomorphisms', *Invent. Math.* **11** (1970), 99–109.
[83] B. Stucky, 'Cubulating one-relator products with torsion', *Groups Geom. Dyn.* **15** (2021), 691–754.
[84] J. Tidmore, 'Cocompact cubulations of mixed 3-manifolds', *Groups Geom. Dyn.* **12**(4) (2018), 1429–1460.
[85] D. T. Wise, 'Cubulating small cancellation groups', *Geom. Funct. Anal.* **14**(1) (2004), 150–214.
[86] D. T. Wise, 'The structure of groups with a quasiconvex hierarchy', *Ann. of Math. Stud.* **209** (2021), 1–376.
[87] D. T. Wise and D. J. Woodhouse, 'A cubical flat torus theorem and the bounded packing property', *Israel J. Math.* **217** (2017), 263–281.
[88] W. Yang, 'Statistically convex-cocompact actions of groups with contracting elements', *Int. Math. Res. Not. IMRN* 2019(23), 7259–7323.