The Dissertation Committee for Theodore Joseph Weisman
certifies that this is the approved version of the following dissertation:

Higher-rank generalizations of convex cocompact and geometrically finite dynamics

Committee:

Jeffrey Danciger, Supervisor

Daniel Allcock

Samuel Ballas

Lewis Bowen

Higher-rank generalizations of convex cocompact and geometrically finite dynamics

by

Theodore Joseph Weisman

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2022

Acknowledgments

I am deeply indebted to my advisor, Jeff Danciger, who is responsible for much of my growth as a mathematician these last few years. In countless meetings where I waffled incoherently about things that I thought I understood (but didn't really), Jeff showed me again and again what good understanding and clear communication actually look like. This thesis owes a lot to Jeff, but I owe him even more: he poured a truly incredible amount of time into lifting me up and guiding me through the equally intimidating and inscrutable worlds of geometry and academia. I can't thank him enough. I also owe thanks to my committee, and to every one of the mathematical mentors I've had throughout my life—I could not have written this thesis without all of their encouragement.

This thesis specifically has also benefited greatly from discussions with many different people, including of course Jeff, but also Daniel Allcock, Fanny Kassel, Katie Mann, Jason Manning, Max Riestenberg, Feng Zhu, and Andy Zimmer. I am grateful to all of these people for sharing their insights with me.

I am beyond grateful for the community I found in the math department at the University of Texas at Austin. I consider myself unbelievably lucky to have spent my time in graduate school among so many people who are unwaveringly committed to supporting each other (mathematically and otherwise), who constantly make each others' lives richer, and who let me feel free to be the huge dork I've always been anyway. There are far too many wonderful people in this department for me to thank everyone individually, but I would like to shout out every one of the Jefflings: thanks to Martin Bobb, Casandra Monroe, Charlie Reid, Max Riestenberg, Florian Stecker, Luis Torres, and Neža Žager Korenjak for teaching me so much, and helping to build a delightful environment for research, learning, and silliness.

This thesis also owes its existence to the people who were a part of my life outside of mathematics. Special mention should go to my good friend Neža Žager Korenjak, who has been there for me so many times and in so many ways that any words feel inadequate. Thanks also to Tom Atchity, Mackenzie Bailey, Tom Gannon, Gill Grindstaff, Jonathan Johnson, Casandra Monroe, Charlie Reid, Max Riestenberg, Logan Stokols, Hannah Turner, Yixian Wu, and many, many others.

I'm especially grateful for all of the communities where I've found both purpose and friendship during my time in graduate school. I'd like to specifically acknowledge the support I received from my friends and teammates on Moontower—sharing the joy of ultimate (and many other shenanigans) with all of these fine folks has been a highlight of my time in Austin, and being part of such a tight-knit community has been vital to my sanity and happiness these last five years. Thanks also to every member of short story club and TV club, for your wisdom, humor, and insight, and thanks to everyone who ever threw a frisbee with me on SO(2). And thank you to the close families I've found myself part of in Austin: to my roommates these last few years, and to the Good Friends Dinner Rotation (a name I just made up, but you know what this is). All of these people made life feel meaningful through everything that was difficult (and wonderful) about my years in graduate school.

Finally, I thank my family: my sister Madeleine, and my parents. Thank you for being infinitely more patient with me than I have been with you, and for your endless support throughout this whole journey. I love you very much.

Higher-rank generalizations of convex cocompact and geometrically finite dynamics

Theodore Joseph Weisman, Ph.D.
The University of Texas at Austin, 2022

Supervisor: Jeffrey Danciger

We study several higher-rank generalizations of the dynamical behavior of convex cocompact groups in rank-one Lie groups, in the context of both convex projective geometry and relatively hyperbolic groups. Our results include a dynamical characterization of a notion of convex cocompact projective structure due to Danciger-Guéritaud-Kassel. This generalizes a dynamical characterization of Anosov representations of hyperbolic groups. Using topological dynamics, we also define a new notion of geometrical finiteness in higher rank which generalizes previous notions of relative Anosov representation due to Kapovich-Leeb and Zhu. We prove that these "extended geometrically finite" representations are stable under certain small relative deformations, and we provide various examples coming from the theory of convex projective structures.

Table of Contents

List of Figures

Chapter 1

Introduction

1.1 Convex cocompactness in rank one

The classical study of Fuchsian and Kleinian groups has yielded enormously important results in the field of geometric topology. In fact much of the modern understanding of manifolds in two and three dimensions rests on deep knowledge of the behavior of discrete subgroups of $\mathrm{SL}(2,\mathbb{R})$ and $\mathrm{SL}(2,\mathbb{C})$ and their connection to hyperbolic geometry. Among such subgroups, the most well-behaved examples are the *convex cocompact* groups.

Definition 1.1.1. Let Γ be a discrete subgroup of $\mathrm{SL}(2,\mathbb{R})$. The group Γ is *convex cocompact* if there is a nonempty Γ-invariant convex subset C of the hyperbolic plane $\mathbb{H}^2$ such that the quotient C/Γ is compact.

This definition generalizes without difficulty for discrete subgroups of $\mathrm{SL}(2,\mathbb{C})$ (replacing the action of Γ on $\mathbb{H}^2$ with an action on $\mathbb{H}^3$) or more generally discrete subgroups of $\mathrm{SO}(d,1)$ (which acts on d-dimensional hyperbolic space $\mathbb{H}^d$).

One way to study convex cocompact groups is to use the tools of *coarse metric geometry.* A convex subset C of d-dimensional hyperbolic space is by definition isometrically embedded. Since $\mathbb{H}^d$ has constant negative sectional curvature, this implies that when C is viewed as a metric space, it is *δ-hyperbolic* in the sense of Gromov [Gro87]. Then, an immediate consequence of the Milnor-Schwarz lemma is

Proposition 1.1.2. *Let $\Gamma \subset \mathrm{SO}(d,1)$ be a convex cocompact group. Then Γ is finitely generated, and for any point $x_0 \in \mathbb{H}^d$, the orbit map $\gamma \mapsto \gamma \cdot x_0$ is a quasi-isometric embedding. Thus, Γ is a* word-hyperbolic *group.*

Proposition 1.1.2 is one reason that convex cocompact groups are a satisfying object of study from the perspective of geometric group theory. Indeed the theory of word-hyperbolic groups has both led to greater understanding of convex cocompactness in general, and has also been enriched by the wide variety of interesting examples that appear in this context. Worth mentioning in particular are *Schottky groups* and the *quasifuchsian* subgroups of $\mathrm{SL}(2,\mathbb{C})$, i.e. deformations of convex cocompact subgroups of $\mathrm{SL}(2,\mathbb{R}) \subset \mathrm{SL}(2,\mathbb{C})$.

Beyond the realm of hyperbolic geometry, it is also possible to define a notion of a convex cocompact subgroup of any semisimple Lie group G with real rank equal to 1. Any discrete subgroup of G acts properly discontinuously on the *Riemannian symmetric space* $X = G/K$, where K is a maximal compact subgroup of G, which leads to the following

Definition 1.1.3. Let Γ be a discrete subgroup of a real semisimple Lie group G with rank one. The group Γ is *convex cocompact* if Γ acts cocompactly on a nonempty convex subset C of the Riemannian symmetric space $X = G/K$.

When $G = \mathrm{SO}(d,1)$ this definition agrees exactly with the definition of convex cocompactness given earlier. When G is a general rank-one Lie group, the Riemannian symmetric space G/K still has everywhere (pinched) negative sectional curvature, so a version of Proposition 1.1.2 also applies in this context.

1.1.1 Convex cocompactness in higher rank

Given that convex cocompact groups in rank-one exhibit a number of desirable geometric, algebraic, and dynamical properties (many of which will be discussed throughout this thesis), it is natural to ask if there is a good definition of a convex cocompact subgroup of a real Lie group G with real rank *greater* than one—possibly one that opens a doorway to understanding the (still very mysterious) world of discrete subgroups of higher-rank Lie groups. Definition 1.1.3 might appear to provide an immediate way forward: one should simply ask for a discrete subgroup $\Gamma \subset G$ to act cocompactly on an invariant convex subset of the Riemannian symmetric space G/K.

This is a perfectly sensible definition, but one problem is that it is too strong to give a theory with a rich variety of examples to work with. In particular, Kleiner-Leeb [KL06] and Quint [Qui05] independently showed the following

Theorem 1.1.4. *Let G be a semisimple Lie group with real rank ≥ 2. If $\Gamma \subset G$ is a discrete Zariski-dense subgroup which acts cocompactly on a nonempty convex subset $C \subset G/K$, then Γ is a uniform lattice in G.*

If we want to use our understanding of convex cocompact groups in rank one to study discrete subgroups of higher rank Lie groups, Theorem 1.1.4 tells us that we should not try generalizing Definition 1.1.3 directly: we will not find any examples beyond uniform lattices. To understand more, we should look for a broader class of discrete subgroup where our rank-one tools still apply, which leads us to the notion of an *Anosov representation*.

1.2 Anosov representations

The idea behind an Anosov representation is the fact that convex cocompact groups in rank one can be characterized by their *dynamical* properties. The original definition of an Anosov subgroup was due to Labourie [Lab06]. He proved that if $\pi_1 S$ is the fundamental group of a hyperbolic surface, and $\rho : \pi_1 S \to \mathrm{SL}(d, \mathbb{R})$ is a *Hitchin representation*, then the geodesic flow on the unit tangent bundle $T^1 S$ induces an *Anosov flow* on a certain bundle associated to the representation ρ. Labourie called representations with such a flow *Anosov representations*, and showed two important results. First, an Anosov representation always has discrete image and finite kernel. Second, Anosov representations are *stable*: they form an open subset of the representation variety $\mathrm{Hom}(\pi_1 S, \mathrm{SL}(d, \mathbb{R}))$.

Later, Guichard-Wienhard [GW12] extended Labourie's definition to cover representations of an arbitrary word-hyperbolic group Γ into a semisimple Lie group G. They further introduced a finer notion of *P-Anosov representation*, where P is a parabolic subgroup of G, and proved a general version of Labourie's stability result:

Theorem 1.2.1. *Let Γ be a word-hyperbolic group, let G be a semisimple Lie group, and let $P \subset G$ be a parabolic subgroup. The P-Anosov representations of Γ into G form an open subset of the representation variety $\mathrm{Hom}(\Gamma, G)$.*

When G has real rank one, it turns out that Anosov representations are basically equivalent to convex cocompact groups: in this case, a representation $\rho : \Gamma \to G$ is Anosov exactly when ρ has finite kernel and $\rho(\Gamma) \subset G$ is a convex cocompact subgroup. And in fact the stability theorems for Anosov representations recall classical work of Sullivan [Sul85] on the stability of convex cocompact groups.

The original proofs of Theorem 1.2.1 in [Lab06] and [GW12] differ significantly from Sullivan's proof of stability in rank-one, since verifying Labourie's original definition of Anosov representation involves working closely with the technicalities of a flow on a particular bundle associated to the representation.

However, there are now many equivalent definitions of an Anosov representation, and some (though not all) of these characterizations are amenable to proofs of Theorem 1.2.1 that use more of Sullivan's original techniques. Definitions have been given in terms of the asymptotics of *Cartan* and *Jordan* projections ([GGKW17], [Tso20]), the geometry of the Riemannian symmetric space G/K ([KLP14], [KLP17]), dynamics of hyperbolic group actions on their boundaries ([GGKW17], [KLP17]), convex projective geometry ([DGK17], [Zim21]), growth of singular value gaps ([BPS19]), and even the algorithmic properties of hyperbolic groups ([BPS19]). We will state the specific definitions that are most relevant to this thesis in Section 2.4.

We emphasize that in every one of the definitions listed above, an Anosov representation *must* be a representation of a word-hyperbolic group. One of the main aims of this thesis is to consider the following

Question 1.2.2. What generalizations of Anosov representations allow us to understand *non-word-hyperbolic* discrete subgroups of higher-rank Lie groups?

Throughout this thesis, we will address this question by considering properties of many of the definitions of Anosov representation listed above. Our focus, however, will be on the relationship between two of them in particular. The first is the characterization of Anosov representations in terms of certain *convex cocompact projective structures*, due to Danciger-Guéritaud-Kassel [DGK17] and Zimmer [Zim21]. The second is the perspective of Kapovich-Leeb-Porti [KLP17] and Guéritaud-Guichard-Kassel-Wienhard [GGKW17], which says that a representation of a hyperbolic group Γ is P-Anosov when it is equipped with an equivariant embedding of the *Gromov boundary* $\partial\Gamma$ of Γ into the flag variety G/P, satisfying certain *dynamical* conditions.

We will explain both of these viewpoints in much greater detail later on. For now, we comment that the relationship between them illustrates the deep connection between *geometric structures* and *hyperbolic boundary dynamics* discussed for instance in a recent survey of Kassel [Kas18]. A central theme of this thesis is that this connection should persist in a broader context. In particular, a version of this principle survives even when we leave the world of word-hyperbolic groups.

1.3 Expansion dynamics for projectively convex cocompact groups

In Chapter 3 of this thesis, we consider one situation in which "Anosov-like" dynamics emerge from a discrete representation of a non-word-hyperbolic group—specifically, a representation associated to a convex cocompact projective structure.

Many authors have previously explored the connection between Anosov representations and geometric structures. Notably, in [GW12], Guichard-Wienhard explained how to construct certain cocompact *domains of discontinuity* for an Anosov representation acting on a flag variety G/P, and in [KLP18], Kapovich-Leeb-Porti gave a systematic procedure for constructing similar domains of discontinuity and determining whether these domains are cocompact and nonempty.

These results show that many Anosov representations can be viewed as the holonomy representations of certain (G, X) structures on closed manifolds. However, the geometric structures in question can be difficult to describe, and indeed understanding the topology of these structures in general is currently an active area of research.

In [DGK17], Danciger-Guéritaud-Kassel introduced a notion of *convex cocompact projective structure* (see Definition 3.1.2 for the precise definition). They showed the following

Theorem 1.3.1 ([DGK17])**.** *Let* $\rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$ *be a* projectively convex cocompact *representation of a word-hyperbolic group* Γ. *Then* ρ *is a* P_1*-Anosov representation.*

In fact Danciger-Guéritaud-Kassel's result is stronger than what we have stated here, and they prove a partial converse as well. Independently, Zimmer [Zim21] also showed that a similar notion of convex cocompact projective structure can be associated to *any* Anosov representation $\rho : \Gamma \to G$ by composing ρ with an appropriate representation of G.

Via Theorem 1.3.1, Danciger-Guéritaud-Kassel aimed to facilitate the study of Anosov representations by associating them to concrete geometric structures whose features closely resemble those of convex cocompact manifolds in hyperbolic space. This echoes work of Choi-Goldman [Gol90] [CG93], which shows that Hitchin representations of surface groups into $\mathrm{SL}(3, \mathbb{R})$ are parameterized by convex real projective structures on surfaces. The approach is also closely related to results of Benoist, who in a series of papers (e.g. [Ben00] [Ben04] [Ben06]) extensively studied the properties of groups acting cocompactly on properly convex subsets of projective space; see [Ben08] for a survey of such results. Benoist's work (and the subsequent results of Danciger-Guéritaud-Kassel and Zimmer) demonstrates that many of the ideas used to understand closed hyperbolic manifolds can be adapted to the setting of projective geometry. This in turn allows for a study of Anosov representations which more closely aligns with the classical study of convex cocompactness in rank one.

However, convex projective structures also hint at a possible way to understand non-word-hyperbolic groups with "Anosov-like" properties. There are many fascinating examples of convex projective manifolds with non-hyperbolic fundamental group, which display an intriguing mix of "hyperbolic" and "nonpositively curved" behavior: see [Ben06], [BDL15], [CLM20], [DGK$^+$21] for some *compact* examples, and see [Bal21] [BM20], [Bob19] for examples with *cusps*. The compact examples cited above all also fall under Danciger-Guéritaud-Kassel's definition of projective convex cocompactness.

Our main result in Chapter 3 (Theorem 3.1.9) answers a question originally asked by Danciger-Guéritaud-Kassel, and says that these examples also have a dynamical characterization which is closely analogous to a dynamical characterization of Anosov representations described in [KLP17].

To be slightly more specific, [KLP17] states that a representation $\rho : \Gamma \to G$ is P-Anosov if and only if Γ has an *expanding action* about its *limit set* in the flag manifold G/P (and the limit set also obeys some additional technical requirements). Our result says that a representation $\rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$ is projectively convex cocompact if and only if Γ has an expanding action about the *faces* in its *full orbital limit set*. The upshot is that projective convex cocompactness is characterized by the expansion dynamics of Γ on a family of Grassmannians $\mathrm{Gr}(k, d)$, for *multiple values of* k. In other words, we are able to see "Anosov-like" dynamics in the action of an (a priori arbitrary) finitely generated group by considering expansion dynamics on several *different* flag manifolds simultaneously.

1.4 Convex cocompactness and relative hyperbolicity

Using the results of Chapter 3, we are able to draw a further connection between Anosov representations and projectively convex cocompact groups which are *relatively* hyperbolic, relative to a collection of convex cocompact subgroups. We explore this connection in detail in Chapter 4. (We provide some background on relatively hyperbolic groups in Section 2.2).

The key idea here is that an Anosov representation of a hyperbolic group Γ is characterized by the existence of a certain *equivariant embedding* of the Gromov boundary of Γ into a flag variety. When $\rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$ is a P_1-Anosov representation, this takes the form of a *pair* of equivariant embeddings of $\partial\Gamma$ into real projective space $\mathbb{RP}^{d-1}$ and the dual projective space $(\mathbb{RP}^{d-1})^*$.

Our main results in Chapter 4 (Theorem 4.1.2 and Theorem 4.1.8) can be summarized as follows. We let Γ be a relatively hyperbolic group, relative to a collection $\mathcal{H}$ of peripheral subgroups. We suppose that $\rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$ is a representation such that the restriction $\rho|_H : H \to \mathrm{PGL}(d, \mathbb{R})$ is projectively convex cocompact for each $H \in \mathcal{H}$. Together, Theorem 4.1.2 and Theorem 4.1.8 state that ρ is projectively convex cocompact if and only if there is an equivariant embedding of the *Bowditch boundary* $\partial(\Gamma, \mathcal{H})$ into a certain *quotient* of projective space. When the collection of peripheral subgroups $\mathcal{H}$ is empty, then Γ is a hyperbolic group and this embedding is precisely the equivariant embedding of the Gromov boundary $\partial\Gamma$ which characterizes an Anosov representation.

We note that our results are closely related to independent work of Islam-Zimmer [IZ20], who previously showed a version of Theorem 4.1.8 when the represented group Γ is isomorphic to the fundamental group of a 3-manifold. In addition, since our work first appeared, Islam-Zimmer have also adapted some of our techniques to prove stronger versions of some of these results (see [IZ22]).

1.4.1 Criterion for relative hyperbolicity

The results of Chapter 4 incidentally allow us to prove more than just a characterization of convex cocompactness for relatively hyperbolic groups. They also give us a way to verify that a group is relatively hyperbolic, relative to a given collection of subgroups.

This is motivated by an observation of Benoist, who showed that convex projective geometry can be used to give a criterion for *hyperbolicity* of a group with a representation $\rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$. Specifically, Benoist showed

Theorem 1.4.1 ([Ben04], Theorem 1.1)**.** *Let Ω be a properly convex subset of $\mathbb{RP}^{d-1}$, and suppose that $\Gamma \subset \mathrm{PGL}(d, \mathbb{R})$ is a discrete group acting cocompactly on Ω. Then Γ is word-hyperbolic if and only if Ω is* strictly convex*, i.e. there are no nontrivial projective segments in $\partial\Omega$.*

Danciger-Guéritaud-Kassel later proved a more general version of this result for a group Γ with a projectively convex cocompact action (see [DGK17], Theorem 1.4). We prove a result for relatively hyperbolic groups (also contained in the statement of Theorem 4.1.8) which has a similar statement in spirit, but a very different proof: whereas Benoist and Danciger-Guéritaud-Kassel directly use the *metric geometry* of convex domains in $\mathbb{RP}^{d-1}$ to prove that certain groups are word-hyperbolic, we instead exploit a result of Yaman [Yam04], which says that relatively hyperbolic groups can be characterized by the *topological dynamics* of their actions on their Bowditch boundaries.

1.5 Extended geometrically finite representations

The last two chapters of this thesis introduce a new generalization of Anosov representations for relatively hyperbolic groups. We refer to these as *extended geometrically finite* (EGF) representations.

1.5.1 Geometrically finite groups in rank one

As the name implies, EGF representations are also a generalization of *geometrically finite* subgroups of rank-one Lie groups, which in turn generalize convex cocompact groups. Compare the following definition to Definition 1.1.3:

Definition 1.5.1. Let $\Gamma \subset G$ be a *finitely generated* discrete subgroup of a rank-one semisimple Lie group G, and let X be the Riemannian symmetric space G/K. We say that Γ is *geometrically finite* if there is a Γ-invariant convex subset $C \subseteq X$ with nonempty interior such that the quotient C/Γ has finite volume.

Remark 1.5.2. In [Bow93] and [Bow95], Bowditch gave several definitions of a "geometrically finite" group in $\mathrm{SO}(d,1)$ and a rank-one Lie group G, respectively. He showed that all of these definitions are equivalent (although in general they *differ* from the historical definition of geometrical finiteness, which was stated in terms of finite-sided polyhedra).

The definition we give here is essentially the same as Bowditch's definition F5 in [Bow95]. Note than unlike the analogous definition for convex cocompactness (Definition 1.1.3), Definition 1.5.1 *assumes* that the discrete group Γ is finitely generated. This assumption is implied by Bowditch's other definitions of geometrical finiteness, but it is included in this one for technical reasons. Indeed Bowditch observed that the assumption follows automatically when Γ is torsion-free (or more generally, if there is a uniform bound on the orders of finite subgroups of Γ). Bowditch further conjectured that this assumption is unnecessary in general, although to the author's knowledge this is still unknown.

In general, a geometrically finite group Γ is always *relatively* hyperbolic, relative to a collection $\mathcal{H}$ of "cusp groups" (i.e. $(\Gamma, \mathcal{H})$ is a *relatively hyperbolic pair*). And in fact, geometrically finite Kleinian groups were one of the primary motivations behind Gromov's original definition [Gro87] of a relatively hyperbolic group. In that sense, geometrically finite groups are the *relatively hyperbolic* analogue of convex cocompact groups in rank one.

One can make this analogy more explicit by presenting two alternative definitions of convex cocompactness and geometrical finiteness, side-by-side; for simplicity we restrict to the context of discrete subgroups of $\mathrm{SO}(d,1)$.

We first need another definition:

Definition 1.5.3. Let Γ be a discrete subgroup of $\mathrm{SO}(d,1)$. The *limit set* Λ_Γ is the set of accumulation points in $\partial\mathbb{H}^d$ of some (any) Γ-orbit $\Gamma \cdot x$, for $x \in \mathbb{H}^d$.

The limit set is always a closed Γ-invariant subset of the ideal boundary $\partial\mathbb{H}^d$ of hyperbolic space. The result below is attributed to Coornaert [Coo90] and Bourdon [Bou95] (although it does not appear in either of those papers exactly as it is stated here).

Proposition 1.5.4. *Let Γ be a word-hyperbolic discrete subgroup of $\mathrm{SO}(d,1)$. Then Γ is convex cocompact if and only if there is an equivariant embedding $\partial\Gamma \to \partial\mathbb{H}^d$ with image Λ_Γ.*

The analogous statement for geometrically finite groups is essentially due to Bowditch. It follows from the first definition of relative hyperbolicity in [Bow12] and the equivalence of definitions GF2 and GF5 in [Bow93].

Proposition 1.5.5. *Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, with Γ a discrete subgroup of $\mathrm{SO}(d,1)$. Then Γ is geometrically finite if and only if there is an equivariant embedding $\partial(\Gamma, \mathcal{H}) \to \partial\mathbb{H}^d$ with image Λ_Γ.*

When $\mathcal{H}$ is empty, then Γ is a word-hyperbolic group, the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ is exactly the Gromov boundary $\partial\Gamma$, and Proposition 1.5.5 reduces exactly to Proposition 1.5.4.

1.5.2 Geometrical finiteness in higher rank: existing definitions

The similarity between Proposition 1.5.4 and Proposition 1.5.5 indicates one possible way to define a higher-rank version of geometrical finiteness—or put another way, a relatively hyperbolic version of an Anosov representation. Since (in light of the results of [GGKW17] and [KLP17]) a P-Anosov representation $\rho : \Gamma \to G$ is a representation equipped with a certain equivariant embedding $\partial\Gamma \to G/P$, a "relativized" Anosov representation of a relatively hyperbolic pair $(\Gamma, \mathcal{H})$ ought to be a representation of Γ equipped with a certain equivariant embedding of the *Bowditch* boundary $\partial(\Gamma, \mathcal{H})$ into G/P.

In [KL18], Kapovich-Leeb gave several possible definitions of an Anosov representation of a relatively hyperbolic group, including one (τ_{mod}-*asymptotic-embeddedness*) which follows this idea. And while not all of the definitions of relative Anosov representation given by Kapovich-Leeb are equivalent, every one of them is still at least implicitly equipped with an equivariant embedding of the Bowditch boundary into some flag manifold G/P.

In independent work, Zhu [Zhu19], [Zhu21] also defined a notion of relative Anosov representation, called a *relatively dominated* representation. Generalizing work of Bochi-Potrie-Sambarino [BPS19] in the Anosov setting, Zhu provided several equivalent characterizations of relatively dominated representations, and further showed that they are always relatively asymptotically embedded in the sense of Kapovich-Leeb. In particular, a relatively dominated representation is always equipped with an equivariant Bowditch boundary embedding.

The definitions of Kapovich-Leeb and Zhu cover several interesting examples of relatively hyperbolic discrete groups in higher rank. For instance, Zhu showed that a type of "geometrically finite" convex projective structure studied by Crampon-Marquis [CM14] gives rise to a relatively dominated representation. Further, in [CZZ21], Canary-Zhang-Zimmer showed that *cusped Hitchin* representations of geometrically finite Fuchsian groups are also relatively dominated. (Currently, there are no known examples of representations that are relatively asymptotically embedded but not relatively dominated.)

1.5.3 Representations without an equivariant boundary embedding

In rank one, the relatively asymptotic embedded representations of Kapovich-Leeb are exactly the same as geometrically finite representations. However, there are also many examples of discrete relatively hyperbolic groups in higher rank which are *not* relatively asymptotically embedded—but still exhibit a sort of "geometrically finite" behavior. Such examples often appear in the context of convex projective structures. For instance, given a geometrically finite group Γ in $\mathrm{SO}(d, 1)$, there are sometimes deformations of Γ inside of $\mathrm{SL}(d + 1, \mathbb{R})$ which do *not* give rise to an equivariant embedding of the Bowditch boundary of Γ—but *are* associated to some convex projective orbifold with a reasonably tame projective structure on its ends, called a *generalized cusp.* Cusps of this type were originally studied by Cooper-Long-

Tillmann [CLT15] [CLT18] and classified by Ballas-Cooper-Leitner [BCL20]; Ballas [Bal21] and Bobb [Bob19] have constructed many examples of projective manifolds with these cusps.

Further, there are also many examples of *projectively convex cocompact* representations of relatively hyperbolic groups which are not relatively asymptotically embedded in the sense of Kapovich-Leeb—and in fact these representations are *never* relatively asymptotically embedded unless the peripheral structure is trivial (so the group is word-hyperbolic and the representation is Anosov). As mentioned previously, explicit constructions of such groups have been given in [BDL15], [CLM20], [DGK+21].

The results of Chapter 4 of this thesis imply, however, that it is still possible to connect these examples to the dynamics of relatively hyperbolic group actions on their Bowditch boundaries. This motivates our definition of *extended geometrically finite representations*:

Definition 1.5.6. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and let P be a symmetric parabolic subgroup of a semisimple Lie group G. We say that a representation $\rho : \Gamma \to G$ is *extended geometrically finite* with respect to P if there exists a closed $\rho(\Gamma)$-invariant subset $\Lambda \subset G/P$ and a continuous ρ-equivariant surjective antipodal map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ which *extends the convergence dynamics* of Γ.

We will explain the technical conditions in Definition 1.5.6 fully in Chapter 5. For now, we note that the key difference between our definition and relative asymptotic embeddedness is that our associated Bowditch boundary map "goes the other way:"

instead of *embedding* $\partial(\Gamma, \mathcal{H})$ into a flag manifold, we instead look for an invariant subset of the flag manifold which *factors onto* the Bowditch boundary via the map ϕ. In the language of this thesis, we say the map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ is an *extension* of the topological dynamical system $(\Gamma, \partial(\Gamma, \mathcal{H}))$. It is the *boundary extension* associated to the EGF representation ρ. Ultimately, we will find (see Theorem 5.6.2) that EGF representations are equivalent to relatively asymptotically embedded representations precisely when the boundary extension is a homeomorphism (so the inverse map is the associated embedding).

The advantage of defining the map in this "backwards" way is that it is more flexible with respect to the action of the peripheral subgroups (the "cusp groups") of the relatively hyperbolic group Γ on the flag manifold G/P. In particular, unlike relatively asymptotic embedded representations, EGF representations allow for *non-unipotent* peripheral subgroups, and peripheral subgroups which are not P*-divergent* (see Section 2.3). This flexibility means that if M is a manifold with relatively hyperbolic fundamental group, convex cocompact projective structures on M, as well as convex projective structures with generalized cusps, can give rise to EGF representations. We prove these results (and provide further new examples of EGF representations) in Chapter 6.

1.5.4 Anosov relativization and stability

Chapter 5 of this thesis proves several properties of EGF representations. We first show that EGF representations interact well with the theory of Anosov representations by proving an *Anosov relativization theorem* (Theorem 5.1.10). This result says that if an EGF representation restricts to Anosov representations on its peripheral subgroups, then it is itself an Anosov representation.

Our main result about EGF representations (Theorem 5.6.2) is that they satisfy a *relative stability* property, in analogy to Theorem 1.2.1 for Anosov representations: we show that if $\rho : \Gamma \to G$ is an EGF representation, then we can define certain *subspaces* W of the representation variety $\mathrm{Hom}(\Gamma, G)$ such that EGF representations form an open subset of W about ρ. We call such subspaces *peripherally stable*, because their definition depends only on the restriction of a deformed representation to the peripheral subgroups of Γ. Peripheral stability should be thought of as a *local* property, giving a criterion which ensures that certain small deformations of ρ are still EGF representations. The condition is fairly flexible—in particular, in Section 6.2, we show that one can find peripherally stable deformation spaces which allow *unipotent* peripheral subgroups in $\mathrm{PGL}(d, \mathbb{R})$ to deform into *diagonalizable* subgroups. Explicit examples of deformations of this type have been found in [BDL15].

It seems desirable to understand the nature of peripherally stable subspaces in general. We know that deformations of EGF representations which restrict to *conjugacy* on the peripheral subgroups are always peripherally stable. And for general EGF representations, we do not expect $\mathrm{Hom}(\Gamma, G)$ itself to be peripherally stable subspace (or in other words, EGF representations are not *absolutely* stable in general). This fails already in rank one, since geometrical finiteness is itself an unstable property (see Remark 5.6.3).

Remark 1.5.7. As a corollary of Theorem 5.6.2 (Corollary 5.6.5), we are able to show that small deformations of relatively asymptotically embedded representations which restrict to conjugacy on the peripherals are still relatively asymptotically embedded.

In [KL18], Kapovich-Leeb asserted that a result along these lines holds, but did not make (or prove) any precise claim. In addition, in [CZZ21], Canary-Zhang-Zimmer study a notion of *P_k-Anosov representations* for geometrically finite Fuchsian groups, and prove that these representations are always *relatively dominated* in the sense of Zhu. Canary-Zhang-Zimmer then show that their P_k-Anosov representations are stable in the space of *type-preserving* deformations, which proves a special case of our relative stability result.

Chapter 2

Background and preliminaries

2.1 Convex projective geometry

Essentially all of the motivating examples in this thesis come from *convex projective geometry*, i.e. the geometry of *properly convex* subsets of real projective space and the projective transformations which preserve them. In this section we review the essential features of convex projective geometry, and establish notation which persists throughout the thesis. For a general reference on this material, see e.g. [Mar14].

2.1.1 Basic definitions

Notation 2.1.1. We will assume without comment that every real vector space V in this thesis is finite-dimensional.

- If V is a real vector space, then $\mathbb{P}(V)$ is projective space $(V - \{0\})/\sim$, where $u \sim v$ if $u = \lambda v$ for $\lambda \in \mathbb{R}$.

- For any $x \in V - \{0\}$, we let $[x]$ denote the image of x under the quotient map $V - \{0\} \to \mathbb{P}(V)$.

- If W is a subset of V, then we let $[W]$ denote the image of $W - \{0\}$ in $\mathbb{P}(V)$. If $W \subseteq V$ is a vector subspace, we will identify $\mathbb{P}(W)$ with $[W] \subset \mathbb{P}(V)$.

- When $B \subset \mathbb{P}(V)$, then the *span* of B, denoted $\mathrm{span}(B)$, is the subspace of V spanned by any lift of B in V. The *projective span* of B is the subset $pr\mathrm{span}(B) \subset \mathbb{P}(V)$.

- If w is any element of the *dual* projective space $\mathbb{P}(V^*)$, we let $\mathbb{P}(w)$ denote the image of $\ker \tilde{w}$ in $\mathbb{P}(V)$, where $\tilde{w}$ is any lift of w in V^*.

- If $w \in \mathbb{P}(V^*)$ and $v \in \mathbb{P}(V)$ satisfy $v \notin \mathbb{P}(w)$, then we say w and v are *transverse* and write $w \perp v$.

- We let $\mathcal{F}_\pm(V)$ denote the space of partial flags (V_-, V_+), where $V_- \subset V_+ \subset V$, and $\dim(V_-) = \mathrm{codim}(V_+) = 1$. If $\dim(V) = d+1$, then we say $\mathcal{F}_\pm(V)$ is the space of flags of *type* $(1,d)$. If V is understood from context we will sometimes just write $\mathcal{F}_\pm = \mathcal{F}_\pm(V)$.

- We let $\mathrm{Gr}(k,V)$ denote the *Grassmannian of k-dimensional subspaces* of V, and write $\mathrm{Gr}(k,d)$ for $\mathrm{Gr}(k,\mathbb{R}^d)$. If W is an element of $\mathrm{Gr}(k,V)$, we write $\mathbb{P}(W)$ for the image of $W - \{0\}$ in $\mathbb{P}(V)$.

2.1.1.1 Convex cones and convex domains

Definition 2.1.2. A *convex cone* in a real vector space V is a convex subset of $V - \{0\}$ which is closed under multiplication by positive scalars.

A convex cone is *sharp* if it does not contain any affine line.

The *boundary* of a convex cone C in a real vector space V is the boundary of C viewed as a cone in its linear span $\mathrm{span}(C) \subset V$; this boundary is homeomorphic to a cone over S^{k-2}, where $k = \dim \operatorname{span}(C)$. We will write ∂C for the boundary of C.

Definition 2.1.3. Let Ω be a subset of $\mathbb{P}(V)$ for a real vector space V. Ω is *convex* if it is the projectivization of some convex cone $\tilde{\Omega} \subset V$, and it is *properly convex* if $\tilde{\Omega}$ is sharp (equivalently, if $\overline{\Omega}$ is a convex subset of some affine chart in $\mathbb{P}(V)$). An open convex set is a *convex domain.*

The boundary $\partial\Omega$ is the projectivization of $\partial\tilde{\Omega} - \{0\}$. A convex set Ω is *strictly convex* if $\partial\Omega$ does not contain a nontrivial projective segment.

For most of this thesis, we will only consider convex domains $\Omega \subset \mathbb{P}(V)$ which are either *open*, or else open in their projective span $\mathbb{P}(\operatorname{span}(\Omega)) \subset \mathbb{P}(V)$. In both of these cases we can write $\overline{\Omega} = \partial\Omega \sqcup \Omega$.

Definition 2.1.4. Let Ω be a convex subset of $\mathbb{P}(V)$. A *supporting subspace* of Ω is a projective subspace which intersects $\partial\Omega$ but not Ω. In particular, a *supporting hyperplane* is a codimension-1 supporting subspace.

Proposition 2.1.5. *Let Ω be a convex subset of $\mathbb{P}(V)$. Every point $x \in \partial\Omega$ is contained in at least one supporting hyperplane.*

Proof. This follows immediately from the fact that every point $x \neq 0$ in the boundary of a convex cone C in V is contained in a codimension-1 subspace containing x but not intersecting C. □

We remark that a convex domain in $\mathbb{P}(V)$ has C^1 boundary precisely when every point in $\partial\Omega$ is contained in *exactly* one supporting hyperplane.

2.1.1.2 Projective line segments

When Ω is a properly convex set, and $x, y \in \overline{\Omega}$, we use $[x, y]$ to denote the unique (closed) projective line segment joining x and y which is contained in $\overline{\Omega}$. We similarly use (x, y), $[x, y)$, $(x, y]$ to denote open and half-open projective line segments.

2.1.1.3 Convex hull and ideal boundary

Definition 2.1.6. If Ω is a properly convex set in $\mathbb{P}(V)$ and $\Lambda \subset \partial\Omega$, then the *convex hull* of Λ is its convex hull in Ω in any affine chart containing Ω. We denote the convex hull of Λ by $\mathrm{Hull}_\Omega(\Lambda)$.

The *ideal boundary* of a set C in a properly convex set Ω is the set

$$\partial_i(C) := \overline{C} \cap \partial\Omega,$$

where the closure of C is taken in $\mathbb{P}(V)$.

2.1.2 Open faces in $\partial\Omega$

Definition 2.1.7. Let Ω be a properly convex domain. The *open face* of $\partial\Omega$ at a point x, which we denote $F_\Omega(x)$, is the set of points $y \in \partial\Omega$ such that x and y lie in an open segment $(a, b) \subset \partial\Omega$.

The *dimension* of a face F is the dimension of a minimal projective subspace containing F; such a minimal subspace is called the *support* of the face and is denoted $\mathrm{supp}(F)$.

A face is always a convex subset of projective space, open in its support. An open face is a closed subset of $\partial\Omega$ if and only if it is an extreme point of Ω.

Remark 2.1.8. Earlier versions of [DGK17] referred to what we call a *face* as a "stratum." Our definition of face agrees with the definition used by Islam and Zimmer. Notably, our faces are *not* the same as the *facettes* of Benoist and Benzecri.

In particular, our definition ensures that every point in the boundary of a properly convex domain Ω is contained in some face.

2.1.3 The Hilbert metric

Here we recall the definition of the Hilbert metric, a useful tool for understanding group actions on properly convex domains. See e.g. [Mar14] for more background.

Given four distinct points a, b, c, d in $\mathbb{R}\mathrm{P}^1$ (or four points in $\mathbb{R}\mathrm{P}^{d-1}$ lying on a single projective line), recall that the *cross-ratio* $[a, b; c, d]$ is given by

$$[a, b; c, d] := \frac{(d-a)\cdot(c-b)}{(c-a)\cdot(d-b)},$$

where the distances are measured in any identification of $\mathbb{R}\mathrm{P}^1$ with $\mathbb{R}\cup\{\infty\}$. Our convention is chosen so that $[0, \infty; 1, z] = z$.

The cross-ratio is a projective invariant on 4-tuples, and in fact it parameterizes the space of $\mathrm{PGL}(2, \mathbb{R})$-orbits of distinct 4-tuples in $\mathbb{R}\mathrm{P}^1$.

Definition 2.1.9. Let $\Omega \subset \mathbb{P}(V)$ be a properly convex domain. The *Hilbert metric*

$$d_\Omega(\cdot, \cdot) : \Omega^2 \to \mathbb{R}^{\geq 0}$$

is given by the formula

$$d_\Omega(x, y) = \frac{1}{2} \log[a, b; x, y],$$

where a, b are the two points in $\partial\Omega$ such that a, x, y, b lie on a projective line in that order.

When the domain Ω is an ellipsoid of dimension d, the Hilbert metric on Ω recovers the familiar Klein model for hyperbolic space $\mathbb{H}^d$. More generally we have the following:

Proposition 2.1.10. *Let Ω be a properly convex domain. Then:*

1. (Ω, d_Ω) *is a proper metric space.*

2. *If x and y are in Ω, then $[x, y]$ is the image of a geodesic (with respect to d_Ω) joining x and y.*

3. $\mathrm{Aut}(\Omega)$ *acts by isometries of d_Ω.*

This implies that $\mathrm{Aut}(\Omega)$ always acts *properly* on Ω. In particular, a subgroup of $\mathrm{Aut}(\Omega)$ is discrete in $\mathrm{PGL}(V)$ if and only if it acts properly discontinuously on Ω.

Part (2) of the above Proposition means that (Ω, d_Ω) is always a geodesic metric space. However, in general it need not be uniquely geodesic—this is one of many ways in which the geometry on a properly convex domain equipped with its Hilbert metric can differ from hyperbolic geometry.

The point of the Hilbert metric is that it allows us to understand many aspects of group actions on convex projective domains in terms of metric geometry; in particular, we may apply the Švarc-Milnor lemma when we have a convex cocompact action on a domain.

The Hilbert metric can also be used to characterize open faces in $\partial\Omega$:

Proposition 2.1.11. *Let Ω be a properly convex domain, let $x \in \partial\Omega$, and fix points $p_1, p_2 \in \Omega$. For any $y \in \partial\Omega$, we have $y \in F_\Omega(x)$ if and only if the Hausdorff distance (with respect to d_Ω) between $[p_1, x)$ and $[p_2, y)$ is finite.*

Since $[p_1, x)$ and $[p_2, y)$ are the images of geodesic rays in (Ω, d_Ω), the above is equivalent to the condition that, if c_x, c_y are unit-speed geodesic rays in (Ω, d_Ω) following projective line segments from p_1, p_2 to x, y, respectively, then

$$d_\Omega(c_x(t), c_y(t)) \leq k$$

for some fixed k independent of $t \in \mathbb{R}^{\geq 0}$.

Each face F of a properly convex projective domain Ω is itself a properly convex set. By viewing F as an open subset of its projective span, one can also define a Hilbert metric d_F on this face.

Proposition 2.1.12. *Let Ω be a properly convex domain, let F be a face of Ω, and let x_n be a sequence in Ω converging to some $x_\infty \in F$.*

For any $D > 0$, if $y_n \in \Omega$ is a sequence satisfying

$$d_\Omega(x_n, y_n) \leq D,$$

then any accumulation point y_∞ of y_n lies in F, and

$$d_F(x_\infty, y_\infty) \leq D.$$

2.1.4 Properly embedded simplices

Definition 2.1.13. A *projective k-simplex in* $\mathbb{RP}^{d-1}$ is the projectivization of the positive linear span of $k+1$ linearly independent vectors in $\mathbb{R}^d$.

A projective k-simplex Δ is an example of a properly convex set in $\mathbb{RP}^{d-1}$. If Δ is the span of standard basis vectors $e_1, \ldots, e_d$, the group $D^+ \subset \mathrm{PGL}(d, \mathbb{R})$ of projectivized diagonal matrices with positive entries (isomorphic to $\mathbb{R}^{d-1}$) acts simply transitively on Δ. Then, any discrete $\mathbb{Z}^{d-1}$ subgroup of D^+ acts properly discontinuously and cocompactly on Δ, so the Švarc-Milnor lemma implies that (Δ, d_Δ) is quasi-isometric to Euclidean space $\mathbb{E}^{d-1}$.

Definition 2.1.14. Let Ω be a properly convex domain. A convex projective simplex $\Delta \subset \Omega$ is *properly embedded* if $\partial\Delta$ is contained in $\partial\Omega$.

A properly embedded simplex in Ω gives an isometric embedding

$$(\Delta, d_\Delta) \to (\Omega, d_\Omega),$$

which in turn gives a quasi-isometric embedding

$$\mathbb{E}^k \to (\Omega, d_\Omega).$$

Maximal properly embedded simplices in Ω can be thought of as analogues of *maximal flats* in CAT(0) spaces; see e.g. [Ben06], [IZ19a], [IZ19b], [Bob20]. However, in general, the metric space (Ω, d_Ω) is not CAT(0); in fact this occurs if and only if Ω is an ellipsoid [KS58].

2.1.5 Duality for convex domains

Let V be a real vector space. Given a convex set $\Omega \subset \mathbb{P}(V)$, it is often useful to consider the *dual convex set* $\Omega^* \subset \mathbb{P}(V^*)$.

Definition 2.1.15. Let C be a convex cone in a real vector space V. The *dual convex cone* $C^* \subset V^* - \{0\}$ is

$$C^* = \{\alpha \in V^* : \alpha(x) > 0 \text{ for all } x \in \overline{C}\}.$$

The following is easily verified:

Proposition 2.1.16. *Let C be a convex cone in a real vector space V.*

1. *C^* is a convex cone in $V^* - \{0\}$.*

2. *$C^{**} = C$, under the canonical identification $V^{**} = V$.*

3. *C^* is sharp if and only if C has nonempty interior.*

If Ω is the projectivization of a convex cone in $\mathbb{P}(V)$, the *dual convex set* is the projectivization Ω^* of $\tilde{\Omega}^*$, where $\tilde{\Omega}$ is any cone over Ω. When Ω is a properly convex domain in $\mathbb{P}(V)$, Ω^* is a properly convex domain in $\mathbb{P}(V^*)$.

In general, Ω^* need not be projectively equivalent to Ω. However, the features of Ω affect the features of Ω^*. For instance, Ω is strictly convex if and only if the boundary of Ω^* is C^1 (and vice versa, since Ω^{**} is naturally identified with Ω). We also note that duality reverses inclusions of convex sets.

If Γ is a subgroup of $\mathrm{PGL}(V)$ preserving Ω, the dual action of Γ on $\mathbb{P}(V^*)$ preserves the dual domain Ω^*. So we can simultaneously view Γ as a subgroup of $\mathrm{Aut}(\Omega)$ and $\mathrm{Aut}(\Omega^*)$.

2.1.6 Dynamics of $\mathrm{Aut}(\Omega)$

When G is a semisimple Lie group, we say that a sequence $g_n \in G$ is *divergent* if it leaves every compact subset of G. A divergent sequence always has a P^+-*divergent* subsequence for some parabolic $P^+ \subset G$ (see Section 2.3). By taking a further subsequence, we can find $P^\pm$-*limit points* (Definition 2.3.5) for g_n and g_n^{-1} in the flag manifolds G/P^+ and G/P^-, for P^- opposite to P^+.

When $G = \mathrm{PGL}(V)$, we can interpret this in terms of *attracting and repelling subspaces*.

Proposition 2.1.17. *Let $g_n \in \mathrm{PGL}(V)$ be a divergent sequence. Up to subsequence, there exist subspaces E_+, E_- in V such that*

$$\dim(E_+) + \dim(E_-) = \dim(V),$$

and for any open set $U \subset \mathbb{P}(V)$ containing $\mathbb{P}(E_+)$ and any compact set $K \subset \mathbb{P}(V) - \mathbb{P}(E_-)$, for all sufficiently large n we have

$$g_n \cdot K \subset U.$$

Proof. For each k with $1 \leq k < d$, we let $P_k \subset \mathrm{PGL}(V)$ be the subgroup preserving a k-dimensional subspace of V. The maximal parabolic subgroups of $\mathrm{PGL}(V)$ are exactly the subgroups conjugate to some P_k. So any divergent sequence in $\mathrm{PGL}(V)$ is P_k-divergent for some k, and the result follows. □

The subspaces E_+, E_- are respectively referred to as *attracting* and *repelling* subspaces. They are uniquely determined once a subsequence of g_n is chosen and their dimensions are specified. They are *not* necessarily transverse.

The faces of Ω inform the limiting dynamical behavior of divergent sequences in $\mathrm{Aut}(\Omega)$. The following two results are likely well-known, but for convenience we provide proofs.

Proposition 2.1.18. *Let γ_n be a divergent sequence in $\mathrm{Aut}(\Omega)$ for a properly convex domain Ω, and suppose that for some $x \in \Omega$, the sequence $\gamma_n x$ accumulates on a face F_+ of Ω. Then, after extracting a subsequence, there is an attracting subspace E_+ of γ_n such that $\mathbb{P}(E_+) \subseteq \mathrm{supp}(F_+)$.*

Proof. Let $x_\infty \in F_+$ be an accumulation point of the point of the sequence $\gamma_n x$. Using a diagonal argument, we can replace γ_n with a subsequence so that for every y in a countable dense subset of $B_\Omega(x, 1)$, the sequence $\gamma_n \cdot y$ has a well-defined limit in the compact space $\mathbb{P}(V)$. Proposition 2.1.12 then implies that for every point $y \in B_\Omega(x, 1)$, the sequence $\gamma_n y$ converges to a unique point in F_+.

Let B_∞ be the set of accumulation points of $\gamma_n \cdot B_\Omega(x, 1)$, and let W_∞ be the subspace $\mathrm{span}(B_\infty) \subset V$.

Proposition 2.1.12 implies that B_∞ is a subset of the face $F = F_\Omega(x_\infty)$, so $\mathbb{P}(W_\infty)$ is a projective subspace of $\mathrm{supp}(F)$. Let $k = \dim(W_\infty)$. We claim that there is an open subset U of the Grassmannian $\mathrm{Gr}(k, V)$ so that

$$\gamma_n \cdot U \to \{W_\infty\}.$$

This implies the desired result by Proposition 2.3.7.

To see the claim, fix k points $z_1, \ldots, z_k \in B_\Omega(x, 1)$ so that the limits of the sequences $\gamma_n z_1, \ldots \gamma_n z_k$ span the projective subspace $\mathbb{P}(W_\infty)$. Proposition 2.1.12 implies that for some fixed $\varepsilon > 0$, if z_i' lies in $B_\Omega(z_i, \varepsilon)$, then the limits of the sequences

$$\gamma_n z_1', \ldots, \gamma_n z_k'$$

are in general position, and therefore also span $\mathbb{P}(W_\infty)$.

Then, if U is the open set

$$\{W \in \mathrm{Gr}(k, V) : W = u_1 \oplus \ldots \oplus u_k, u_i \in B_\Omega(z_i, \varepsilon)\},$$

we have that $\gamma_n U \to \{W_\infty\}$, as required. □

We also have a closely related lemma:

Lemma 2.1.19. *Let Ω be a properly convex domain in $\mathbb{P}(V)$, let Γ be a subgroup of $\mathrm{Aut}(\Omega)$, and let Λ be any closed Γ-invariant subset of $\partial\Omega$ with nonempty convex hull in Ω.*

If E_+ and E_- are attracting and repelling subspaces for some divergent sequence $\{\gamma_n\} \subset \Gamma$, then $\mathbb{P}(E_+)$ and $\mathbb{P}(E_-)$ are supporting subspaces of Ω that intersect Λ nontrivially.

Proof. It suffices to show the claim for E_+, because a repelling subspace for the sequence γ_n is an attracting subspace for the sequence γ_n^{-1}.

Since Ω is open, it is not contained in $\mathbb{P}(E_-)$. So, for some $x \in \Omega$, the limit of $\gamma_n \cdot x$ is contained in $\mathbb{P}(E_+)$. Since Ω is Γ-invariant, $\mathbb{P}(E_+)$ intersects $\overline{\Omega}$ nontrivially.

Let E_+^* be the subspace of V^* consisting of functionals which vanish on E_+. E_+^* is an attracting subspace for the sequence γ_n under the dual action of Γ on V^*. So, by the previous argument, $\mathbb{P}(E_+^*)$ intersects $\overline{\Omega^*}$ nontrivially, which means $\mathbb{P}(E_+)$ cannot intersect Ω.

This shows that $\mathbb{P}(E_+)$ is a supporting subspace of Ω (and therefore $\mathbb{P}(E_-)$ is as well). To see that $\mathbb{P}(E_+)$ intersects Λ nontrivially, note that since Λ has nonempty convex hull in Ω and $\mathbb{P}(E_-)$ is a supporting subspace of Ω, Λ is not a subset of $\mathbb{P}(E_-) \cap \partial\Omega$. So, for some $x \in \Lambda$, $\gamma_n \cdot x$ accumulates to a point y in $\mathbb{P}(E_+)$; since Λ is Γ-invariant and closed, y is in Λ also. $\square$

2.2 Relatively hyperbolic groups

In this section we discuss some of the basic theory of relatively hyperbolic groups, mostly to establish the notation and conventions we will use throughout the thesis. We refer to [BH99], [Bow12], [DS05] for background on hyperbolic groups and relatively hyperbolic groups. See also section 3 of [KL18] for an overview (which we follow in part here).

Notation 2.2.1. Throughout this thesis, if X is a metric space, A is a subset of X, and $r \geq 0$, we let $N_X(A, r)$ denote the open r-neighborhood in X about A. For a point $x \in X$, we let $B_X(x, r)$ denote the open r-ball about x.

When the metric space X is implied from context, we will often just write $N(A, r)$ or $B(x, r)$.

2.2.1 Geometrically finite actions

Recall that a finitely generated group Γ is *hyperbolic* (or *word-hyperbolic* or *δ-hyperbolic* or *Gromov-hyperbolic*) if and only if it acts properly discontinuously and cocompactly on a δ-hyperbolic proper geodesic metric space Y.

A *relatively hyperbolic group* is also a group with an action by isometries on a δ-hyperbolic proper geodesic metric space Y, but instead of asking for the action to cocompact, we ask for the action to be in some sense "geometrically finite."

To be precise, this means that Y has a Γ-invariant decomposition into a *thick part* Y_{th} and a countable collection $\mathcal{B}$ of *horoballs*, invariant under the action of Γ on Y. For a horoball B, we let $\mathrm{ctr}(B)$ denote the center of B in ∂Y, and we let Γ_p denote the stabilizer of any $p \in \partial Y$.

Definition 2.2.2. Let Γ be a group acting on a hyperbolic metric space Y, and let $\mathcal{B}$ be a countable collection of horoballs in Y. If:

1. The action of Γ on the closure of $Y_{\mathrm{th}} = Y - \bigcup_{B \in \mathcal{B}} B$ is cocompact, and

2. for each $B \in \mathcal{B}$, the stabilizer of $\mathrm{ctr}(B)$ in Γ is finitely generated and infinite,

then we say that Γ is a *relatively hyperbolic group*, relative to the collection $\mathcal{H} = \{\Gamma_p : p = \mathrm{ctr}(B) \text{ for } B \in \mathcal{B}\}$.

Definition 2.2.3. Let Γ be a relatively hyperbolic group, relative to a collection of subgroups $\mathcal{H}$.

- The centers of the horoballs in $\mathcal{B}$ are called *parabolic points* for the Γ-action on ∂Y. The set of parabolic points in ∂Y is denoted $\partial_{\mathrm{par}} Y$.

- The parabolic point stablizers $\mathcal{H} = \{\mathrm{Stab}_\Gamma(p) : p \in \partial_{\mathrm{par}} Y\}$ are called *peripheral subgroups*. We often write Γ_p for $\mathrm{Stab}_\Gamma(p)$.

A group Γ might be hyperbolic relative to different collections $\mathcal{H}$, $\mathcal{H}'$ of peripheral subgroups. The collection $\mathcal{H}$ of peripheral subgroups is sometimes called a *peripheral structure* for Γ.

Definition 2.2.4. Let Γ be a finitely generated group, and let $\mathcal{H}$ be a collection of subgroups. We say that $(\Gamma, \mathcal{H})$ is a *relatively hyperbolic pair* if Γ is hyperbolic relative to $\mathcal{H}$.

Remark 2.2.5. Whenever $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair, there are always only *finitely many* conjuacy classes of groups lying in the collection $\mathcal{H}$. Often, it will be convenient to consider a finite subset $\mathcal{P} \subset \mathcal{H}$ containing exactly one subgroup in each conjugacy class in $\mathcal{H}$.

In fact, many authors adopt the convention that a relatively hyperbolic group Γ is hyperbolic relative to a *finite* collection of subgroups $\mathcal{P}$, and work with the *conjugates* of $\mathcal{P}$ as needed. In this thesis, however, if $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair, then $\mathcal{H}$ is invariant under conjugacy. In particular, $\mathcal{H}$ is always either empty or infinite.

Our convention is somewhat more natural when considering the *action* of a relatively hyperbolic group Γ on various spaces. However, the tradeoff is that it is a little more unwieldy when we want to consider the metric geometry of either Γ itself, a Gromov model, or some other associated metric space (e.g. Definition 2.2.17 below).

2.2.2 The Bowditch boundary

Definition 2.2.6. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, so that $\mathcal{H}$ is the set of stabilizers of parabolic points for an action of Γ on a metric space Y as in Definition 2.2.2. We say that Y is a *Gromov model* for the pair $(\Gamma, \mathcal{H})$.

In general there is *not* a unique choice of Gromov model for a given relatively hyperbolic pair $(\Gamma, \mathcal{H})$, even up to quasi-isometry. There are various "canonical" constructions for a preferred quasi-isometry class of Gromov model, with certain desirable metric properties (see e.g. [Bow12], [GM08]).

Given *any* two Gromov models Y, Y' for $(\Gamma, \mathcal{H})$, there is always a Γ-equivariant homeomorphism $\partial Y \to \partial Y'$ [Bow12]. The Γ-space ∂Y is the *Bowditch boundary* of $(\Gamma, \mathcal{H})$. We will denote it by $\partial(\Gamma, \mathcal{H})$, or sometimes just $\partial\Gamma$ when the collection of peripheral subgroups is understood from context. Since a Gromov model Y is a proper hyperbolic metric space, $\partial(\Gamma, \mathcal{H})$ is always compact and metrizable.

Definition 2.2.7. We say a relatively hyperbolic pair $(\Gamma, \mathcal{H})$ is *elementary* if Γ is finite or virtually cyclic, or if $\mathcal{H} = \{\Gamma\}$.

Whenever $(\Gamma, \mathcal{H})$ is nonelementary, its Bowditch boundary contains at least three points. The convergence properties of the action of Γ on $\partial(\Gamma, \mathcal{H})$ (see below) imply that in this case, $\partial(\Gamma, \mathcal{H})$ is *perfect* (i.e. contains no isolated points).

2.2.2.1 Cocompactness on pairs

Let Y be a Gromov model for a relatively hyperbolic pair $(\Gamma, \mathcal{H})$. Since Y is hyperbolic, proper, and geodesic, for any compact subset $K \subset Y$, the space of bi-infinite geodesics passing through K is compact.

Given any distinct pair of points $u, v \in \partial Y$, there is a bi-infinite geodesic c in Y joining u to v. This geodesic must pass through the thick part Y_{th} of Y, so up to the action of Γ it passes through a fixed compact subset $K \subset Y_{\mathrm{th}}$.

This implies:

Proposition 2.2.8. *The action of* Γ *on the space of distinct pairs in* $\partial(\Gamma, \mathcal{H})$ *is cocompact.*

2.2.3 Convergence group actions

If a group Γ acts on a proper geodesic hyperbolic metric space Y, we can characterize the geometrical finiteness of the action entirely in terms of the topological dynamics of the action on ∂Y. In particular, we can understand geometrical finiteness by studying properties of *convergence group actions.* See [Tuk94], [Tuk98], [Bow99] for further detail on such actions, and justifications for the results stated in this section.

Definition 2.2.9. Let Γ act on a Hausdorff space Z. We say that Γ acts as a *convergence group* if for every sequence γ_n of pairwise distinct elements of Γ, there exists a subsequence γ_{n_k} and points $a, b \in Z$ such that the restriction of γ_{n_k} to $Z - \{a\}$ converges to the constant map b, uniformly on compacts in $Z - \{a\}$.

Definition 2.2.10. Let Γ act as a convergence group on a topological space Z.

1. A point $z \in Z$ is a *conical limit point* if there exists a sequence $\gamma_n \in \Gamma$ and distinct points $a, b \in Z$ such that $\gamma_n z \to a$ and $\gamma_n y \to b$ for any $y \neq z$.

2. An infinite subgroup H is a *parabolic subgroup* if it fixes a point $p \in Z$, and every infinite-order element of H fixes exactly one point in Z.

3. A point $p \in Z$ is a *parabolic point* if it is the fixed point of a parabolic subgroup.

4. A parabolic point p is *bounded* if its stabilizer Γ_p acts cocompactly on $Z - \{p\}$.

The name "conical limit point" makes more sense in the context of convergence group actions on boundaries of hyperbolic metric spaces.

Definition 2.2.11. Let Y be a hyperbolic metric space, and let $z \in \partial Y$. We say that a sequence $y_n \in Y$ *limits conically to* z if there is a geodesic ray $c : \mathbb{R}^+ \to Y$ limiting to z and a constant $D > 0$ such that

$$d_Y(y_n, c(t_n)) < D$$

for some sequence $t_n \to \infty$.

A bounded neighborhood of a geodesic in a hyperbolic metric space looks like a "cone," hence "conical limit."

Proposition 2.2.12. *Let Γ act properly discontinuously by isometries on a proper geodesic hyperbolic metric space Y, and fix a basepoint $y_0 \in Y$.*

Γ acts on ∂Y as a convergence group. Moreover, a point $z \in \partial Y$ is a conical limit point (in the dynamical sense of Definition 2.2.10) if and only if there is a sequence $\gamma_n \cdot y_0$ limiting conically to z (in the geometric sense of Definition 2.2.11).

In this case, there are distinct points $a, b \in \partial Y$ such that $\gamma_n^{-1} \cdot z \to a$ and $\gamma_n^{-1} z' \to b$ for any $z' \neq z$ in ∂Y.

If $\gamma_n \cdot y_0$ limits conically to a point $z \in \partial Y$ for some (hence any) basepoint $y_0 \in Y$, we just say that γ_n limits conically to z.

Theorem 2.2.13 ([Bow12])**.** *Let Γ be a group acting by isometries on a hyperbolic metric space Y. Then Γ is a relatively hyperbolic group, acting on Y as in Definition 2.2.2, if and only if:*

- *The induced action of Γ on ∂Y is a convergence group action, and*
- *Every point $z \in \partial Y$ is either a conical limit point or a bounded parabolic point.*

Whenever a group Γ acts as a convergence group on a perfect compact metrizable space Z, and every point in Z is either a conical limit point or a bounded parabolic point, we say the Γ-action on Z is *geometrically finite*. This is justified by the following theorem of Yaman:

Theorem 2.2.14 ([Yam04])**.** *Let Γ be a non-elementary group, and let $\mathcal{H}$ be the collection of all conjugates of a finite collection of finitely-generated proper subgroups of Γ.*

Then Γ is hyperbolic relative to $\mathcal{H}$ if and only if Γ acts on a compact, perfect, and metrizable space Z as a convergence group, every point in Z is either a conical limit point or a bounded parabolic point for the Γ-action, and the parabolic points in Z are exactly the fixed points of the groups in $\mathcal{H}$.

In this case, the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ is equivariantly homeomorphic to Z.

Definition 2.2.15. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair. We write

$$\partial(\Gamma, \mathcal{H}) = \partial_{\text{con}}(\Gamma, \mathcal{H}) \sqcup \partial_{\text{par}}(\Gamma, \mathcal{H}),$$

where $\partial_{\text{con}}(\Gamma, \mathcal{H})$ and $\partial_{\text{par}}(\Gamma, \mathcal{H})$ are respectively the conical limit points and parabolic points in $\partial(\Gamma, \mathcal{H})$.

2.2.3.1 Compactification of Γ and divergent sequences

When $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair, there is a natural topology on the set

$$\overline{\Gamma} = \Gamma \sqcup \partial(\Gamma, \mathcal{H})$$

making it into a *compactification* of Γ (i.e. $\overline{\Gamma}$ is compact, $\partial(\Gamma, \mathcal{H})$ and Γ are both embedded in $\overline{\Gamma}$, and Γ is an open dense subset of $\overline{\Gamma}$).

Definition 2.2.16. A sequence $\gamma_n \in \Gamma$ is *divergent* if it leaves every bounded subset of Γ (equivalently, if a subsequence of it consists of pairwise distinct elements).

Up to subsequence, a divergent sequence $\gamma_n \in \Gamma$ converges to a point $z \in \partial(\Gamma, \mathcal{H})$.

2.2.4 The coned-off Cayley graph

Fix a finite set $\mathcal{P}$ of conjugacy representatives for the groups in $\mathcal{H}$. The set $\mathcal{P}$ corresponds to a finite set $\Pi \subset \partial_{\text{par}}\Gamma$ of parabolic points, such that

$$\mathcal{P} = \{\Gamma_p : p \in \Pi\}.$$

Π contains exactly one point in each Γ-orbit in $\partial_{\text{par}}\Gamma$.

Definition 2.2.17. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and fix a finite generating set S for Γ and finite collection of conjugacy representatives $\mathcal{P}$ for $\mathcal{H}$.

The *coned-off Cayley graph* $\mathrm{Cay}(\Gamma, S, \mathcal{P})$ is a metric space obtained from the Cayley graph $\mathrm{Cay}(\Gamma, S)$ as follows: for each coset gP_i for $P_i \in \mathcal{P}$, we add a vertex $v(gP_i)$. Then, we add an edge of length 1 from each $h \in gP_i$ to $v(gP_i)$.

The quasi-isometry class of $\mathrm{Cay}(\Gamma, S, \mathcal{P})$ is independent of the choice of generating set S. When $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair, $\mathrm{Cay}(\Gamma, S, \mathcal{P})$ is a hyperbolic metric space. It is *not* a proper metric space if $\mathcal{H}$ is nonempty. The Gromov boundary of $\mathrm{Cay}(\Gamma, S, \mathcal{P})$ is equivariantly homeomorphic to the set $\partial_{\mathrm{con}}\Gamma$ of conical limit points in $\partial(\Gamma, \mathcal{H})$.

2.3 Semisimple Lie groups: structure and dynamics

The last two chapters of this thesis are more general than the rest: we will consider not just discrete subgroups of $\mathrm{PGL}(V)$ acting on real projective space $\mathbb{P}(V)$, but discrete subgroups Γ of semisimple Lie groups G acting on *flag manifolds* G/P, where P is a parabolic subgroup of G.

In this section, we give an overview of the definitions and notation we will use to describe the *dynamical behavior* such actions. We mostly follow the notation of [GGKW17], but we will also identify the connection to the language of [KLP17].

For the most part, we will avoid using the technical theory of semisimple Lie groups and their associated Riemannian symmetric spaces. In fact, in many cases, our approach will be to use a representation of G to reduce to the case $G = \mathrm{PGL}(n, \mathbb{R})$.

The most important part of this section is 2.3.6, which identifies the connection between P-*divergence* (or equivalently τ_{mod}-*regularity*) and *contracting dynamics* in G.

Standard references for the general theory are [Ebe96], [Hel01], and [Kna02]. See also section 3 of [Max21] for a careful discussion of the theory as it relates to Anosov representations and the work of Kapovich-Leeb-Porti.

2.3.1 Parabolic subgroups

Let K be a maximal compact subgroup of the semisimple Lie group G, and let X be the Riemannian symmetric space G/K. A subgroup $P \subset G$ is a *parabolic subgroup* if it is the stabilizer of a point in the visual boundary $\partial_\infty X$ of X. Two parabolic subgroups P, Q are *opposite* if there is a bi-infinite geodesic c in X so that P is the stabilizer of $c(\infty)$ and Q is the stabilizer of $c(-\infty)$.

The compact homogeneous G-space G/P is called a *flag manifold*. If P and Q are parabolic subgroups, then we say that two flags $\xi^+ \in G/P$ and $\xi^- \in G/Q$ are *opposite* if the stabilizers of ξ^+, ξ^- are opposite parabolic subgroups. (In particular a conjugate of Q must be opposite to P).

2.3.2 Root space decomposition

Let $\mathfrak{g}$ be the Lie algebra of G, and let $\mathfrak{k}$ be the Lie algebra of the maximal compact K. We can decompose $\mathfrak{g}$ as $\mathfrak{k} \oplus \mathfrak{p}$, and fix a maximal abelian subalgebra $\mathfrak{a} \subset \mathfrak{p}$ (the *Cartan subalgebra*). The restriction of the Killing form B to $\mathfrak{p}$ is positive definite, so any Cartan subalgebra $\mathfrak{a}$ is naturally endowed with a Euclidean structure.

Since $\mathfrak{a}$ is abelian, it acts semisimply on $\mathfrak{g}$. So we let $\Sigma \subset \mathfrak{a}^*$ denote the set of *roots* for this Cartan subalgebra. We have a *root space decomposition*

$$\mathfrak{g} = \mathfrak{g}_0 \oplus \bigoplus_{\alpha \in \Sigma} \mathfrak{g}_\alpha,$$

where $X \in \mathfrak{a}$ acts on $\mathfrak{g}_\alpha$ by multiplication by $\alpha(X)$.

We choose a set of *simple roots* $\Delta \subset \Sigma$ so that each $\alpha \in \Sigma$ can be uniquely written as a linear combination of elements of Δ with coefficients either all nonnegative or all nonpositive. We let Σ_+ denote the *positive roots*, i.e. roots which are nonnegative linear combinations of elements of Δ.

The simple roots Δ determine a *Euclidean Weyl chamber*

$$\mathfrak{a}^+ = \{x \in \mathfrak{a} : \alpha(x) \geq 0, \text{ for all } \alpha \in \Delta\}.$$

The kernels of the roots $\alpha \in \Delta$ are the *walls* of the Euclidean Weyl chamber.

Choosing a maximal compact K, a Cartan subalgebra $\mathfrak{a}$, and a Euclidean Weyl chamber $\mathfrak{a}^+$ determines a *Cartan projection*

$$\mu : G \to \mathfrak{a}^+,$$

uniquely determined by the equation $g = k \exp(\mu(g))k'$, where $k, k' \in K$ and $\mu(g) \in \mathfrak{a}^+$.

2.3.3 P-divergence

Fix a subset θ of the simple roots Δ. We define a *standard parabolic subgroup* P_θ^+ to be the normalizer of the Lie algebra

$$\bigoplus_{\alpha \in \Sigma_\theta^+} \mathfrak{g}_\alpha,$$

where Σ_θ^+ is the set of positive roots which are *not* in the span of $\Delta - \theta$. The *opposite subgroup* P^- is the normalizer of

$$\bigoplus_{\alpha \in \Sigma_\theta^+} \mathfrak{g}_{-\alpha}.$$

Every parabolic subgroup $P \subset G$ is conjugate to a unique standard parabolic subgroup P_θ^+, and every pair of opposite parabolics (P^+, P^-) is simultaneously conjugate to a unique pair (P_θ^+, P_θ^-).

For a fixed $\theta \subset \Delta$, the group P_θ^+ is the stabilizer of the endpoint of a geodesic ray $\exp(tZ) \cdot p$, where $p \in X$ is the image of the identity in G/K, and for any $\alpha \in \Delta$, the element $Z \in \mathfrak{a}^+$ satisfies

$$\alpha(Z) = 0 \iff \alpha \in \Delta - \theta.$$

Definition 2.3.1. Let g_n be a sequence in G. The sequence g_n is P_θ^+*-divergent* if for every $\alpha \in \theta$, we have

$$\alpha(\mu(g_n)) \to \infty.$$

That is, the Cartan projections of the sequence g_n *drift away from the walls* of $\mathfrak{a}$ determined by the subset $\theta \subset \Delta$.

For a general parabolic subgroup $P \subset G$, we say that g_n is *P-divergent* if g_n is P_θ^+-divergent for P_θ^+ conjugate to P.

2.3.4 τ_{mod}-regularity

P-divergent sequences are equivalent to the *τ_{mod}-regular* sequences discussed in the work of Kapovich-Leeb-Porti, where τ_{mod} is the unique face corresponding to P in a *spherical model Weyl chamber.* We explain the connection here.

For any point $p \in X$, we let $\mathfrak{p}$ be the uniquely determined subspace of $\mathfrak{g}$ such that $\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}$, where $\mathfrak{k}$ is the Lie algebra of the stabilizer of p in G.

Let $z \in \partial_\infty X$. There is a point $p \in X$, a maximal abelian subalgebra $\mathfrak{a} \subset \mathfrak{p}$, a Euclidean Weyl chamber $\mathfrak{a}^+ \subset \mathfrak{a}$, and a unit-length $Z \in \mathfrak{a}^+$ such that z is the endpoint of the geodesic ray $c(t) = \exp(tZ) \cdot p$.

Up to the action of the stabilizer of z, the point p, the Cartan subalgebra $\mathfrak{a}$, the Euclidean Weyl chamber $\mathfrak{a}^+$, and the unit vector $Z \in \mathfrak{a}^+$ are uniquely determined. In addition, the stabilizer in G of the triple $(p, \mathfrak{a}, \mathfrak{a}^+)$ acts trivially on $\mathfrak{a}^+$.

This means that we can identify the space $\partial_\infty X/G$ with the set of unit vectors in any Euclidean Weyl chamber $\mathfrak{a}^+$. This set has the structure of a *spherical simplex.* We let σ_{mod} denote the *model spherical Weyl chamber* $\partial_\infty X/G$.

We let $\pi : \partial_\infty X \to \sigma_{\mathrm{mod}}$ be the *type map* to the model spherical Weyl chamber. For fixed $z \in \partial_\infty X$, we let P_z denote the parabolic subgroup stabilizing z.

After choosing a maximal compact K, a Cartan subalgebra $\mathfrak{a}$, and a Euclidean Weyl chamber $\mathfrak{a}^+$, the data of a *face* τ_{mod} of the spherical simplex σ_{mod} is the same as the data of a *subset* of the simple roots of G: the set of roots identifies a collection of walls of the Euclidean Weyl chamber $\mathfrak{a}^+$. The intersection of those walls with the unit sphere in $\mathfrak{a}$ is uniquely identified with a face of σ_{mod}.

Definition 2.3.2. Let τ_{mod} be a face of the model spherical Weyl chamber σ_{mod}. We say that a sequence $g_n \in G$ is τ_{mod}-*regular* if g_n is P_z-divergent for some $z \in \partial_\infty X$ such that $\pi(z) \in \tau_{\mathrm{mod}}$.

For a fixed model face $\tau_{\mathrm{mod}} \subset \sigma_{\mathrm{mod}}$, we let $P_{\tau_{\mathrm{mod}}}$ denote any parabolic subgroup which is the stabilizer of a point $z \in \pi^{-1}(\tau_{\mathrm{mod}})$. All such parabolic subgroups are conjugate, so as a G-space the flag manifold $G/P_{\tau_{\mathrm{mod}}}$ depends only on the model face τ_{mod}.

2.3.5 Affine charts

Definition 2.3.3. Let P^+, P^- be opposite parabolic subgroups in G. Given a flag $\xi \in G/P^-$, we define

$$\mathrm{Opp}(\xi) = \{\eta \in G/P^+ : \xi \text{ is opposite to } \eta\}.$$

We call a set of the form $\mathrm{Opp}(\xi)$ for some $\xi \in G/P^-$ an *affine chart* in G/P^+.

An affine chart is the unique open dense orbit of $\mathrm{Stab}_G(\xi)$ in G/P^+. When $G = \mathrm{PGL}(d+1, \mathbb{R})$ and P^+ is the stabilizer of a line $\ell \subset \mathbb{R}^{d+1}$, G/P^+ is identified with $\mathbb{RP}^d$ and this notion of affine chart agrees with the usual one in $\mathbb{RP}^d$.

2.3.6 Dynamics in flag manifolds

There is a close connection between P-divergence in the group G and the *topological dynamics* of the action of G on the associated flag manifold G/P. Kapovich-Leeb-Porti frame this connection in terms of a *contraction property* for P-divergent sequences.

Definition 2.3.4 ([KLP17], Definition 4.1)**.** Let g_n be a sequence of group elements in G. We say that g_n is *P^+-contracting* if there exist $\xi \in G/P^+$, $\xi_- \in G/P^-$ such that g_n converges uniformly to ξ on compact subsets of $\mathrm{Opp}(\xi_-)$.

The flag ξ is the uniquely determined *limit* of the sequence g_n.

Definition 2.3.5. For an arbitrary sequence $g_n \in G$, a *P^+-limit point* of g_n in G/P^+ is the limit point of some P^+-contracting subsequence of g_n.

The *P^+-limit set* of a group $\Gamma \subset G$ is the set of P^+-limit points of sequences in Γ.

The importance of contracting sequences is captured by the following:

Proposition 2.3.6 ([KLP17], Proposition 4.15)**.** *A sequence $g_n \in G$ is P^+-divergent if and only if every subsequence of g_n has a P^+-contracting subsequence.*

Proposition 2.3.6 implies in particular that if $g_n \in G$ is P^+-divergent, then up to subsequence there is an open subset $U \subset G/P^+$ such that $g_n \cdot U$ converges to a singleton in G/P^+. It turns out that this "weak contraction property" is enough to characterize P^+-divergence.

Proposition 2.3.7. *Let g_n be a sequence in G, and suppose that for some nonempty open subset $U \subset G/P^+$, we have $g_n \cdot U \to \{\xi\}$ for $\xi \in G/P^+$. Then g_n is P^+-divergent, and has a unique P^+-limit point $\xi \in G/P^+$.*

We provide a proof of this fact in Section 6.3.

2.3.6.1 Dynamics of inverses of P^+-divergent sequences

When g_n is a P^+-divergent sequence, the inverse sequence is P^--divergent. Kapovich-Leeb-Porti show that this can be framed in terms of the dynamical behavior of the inverse sequence.

Lemma 2.3.8 ([KLP17], Lemma 4.19). *For $g_n \in G$ and flags $\xi_- \in G/P^-, \xi_+ \in G/P^+$, the following are equivalent:*

1. *g_n is P^+-contracting and $g_n|_{\mathrm{Opp}(\xi_-)} \to \xi_+$ uniformly on compacts.*
2. *g_n is P^+-divergent, g_n has unique P^+-limit point ξ_+, and g_n^{-1} has unique P^--limit point ξ_-.*

2.4 Anosov representations: working definitions

The notion of an *Anosov representation* is central to this thesis. In this section we briefly review several equivalent definitions of Anosov representations. Proofs of the various equivalences can be found in e.g. [GGKW17] and [KLP17]. For a comprehensive overview of the topic, see the lecture notes [Can21].

2.4.1 Definitions in terms of limit maps

Definition 2.4.1. Let G be a semisimple Lie group, let $P^{\pm}$ be a pair of opposite parabolic subgroups, and let Z be a topological space. We say that a pair of maps $\xi^{\pm} : Z \to G/P^{\pm}$ are *antipodal* if for any distinct $z_1, z_2 \in Z$, $\xi^+(z_1)$ is opposite to $\xi^-(z_2)$.

When P is a *symmetric* parabolic subgroup (i.e. P is conjugate to some opposite parabolic P^-), then we say a map $\xi : Z \to G/P$ is *antipodal* if $\xi(z_1)$ is opposite to $\xi(z_2)$ for distinct $z_1, z_2 \in Z$.

When G is a rank-one semisimple Lie group, then there is only one conjugacy class of parabolic subgroup, and a map $\xi : Z \to G/P$ is antipodal if and only if it is injective.

Definition 2.4.2. Let Γ be a word-hyperbolic group, let G be a semisimple Lie group, and let $P^{\pm} \subset G$ be a pair of opposite parabolic subgroups. We say that a representation $\rho : \Gamma \to G$ is *P-Anosov* if $\rho(\Gamma)$ is P^+-divergent, and there exists a pair of ρ-equivariant antipodal embeddings $\xi^{\pm} : \partial\Gamma \to G/P^{\pm}$ such that the image of ξ^+ is the P^+-limit set of $\rho(\Gamma)$ and the image of ξ^- is the P^--limit set of $\rho(\Gamma)$.

In [KLP17], Kapovich-Leeb-Porti refer to representations satisfying Definition 2.4.2 as *asymptotically embedded*, and prove that asymptotic embedded representations are equivalent to Anosov representations. We will take Definition 2.4.2 as our primary definition of Anosov representations.

It is immediate from the definition that a representation $\rho : \Gamma \to G$ is P^+-Anosov if and only if it is P^--Anosov, if and only if it is P-Anosov for the *symmetric* parabolic subgroup $P = P_{+\theta} \cap P_{-\theta}$, where $P_{\pm\theta}$ are the standard parabolic subgroups conjugate to P^+ and P^-.

The end of the previous section shows how asymptotic embeddedness can also be characterized via *topological dynamics*.

Proposition 2.4.3. *Let Γ be a word-hyperbolic group, let G be a semisimple Lie group, and let $P^\pm \subset G$ be a pair of opposite parabolics. A representation $\rho : \Gamma \to G$ is P^+-Anosov if and only if there are ρ-equivariant antipodal embeddings $\xi^\pm : \partial\Gamma \to G/P^\pm$, such that for any sequence $\gamma_n \in \Gamma$ such that $\gamma_n^{\pm 1} \to z^\pm \in \partial\Gamma$, we have*

$$\rho(\gamma_n)|_{\mathrm{Opp}(\xi^-(z^-))} \to \xi^+(z^+)$$

uniformly on compacts.

Proof. Suppose $\rho : \Gamma \to G$ is P^+-Anosov, and let $\gamma_n \in \Gamma$ be a sequence such that $\gamma_n^{\pm 1} \to z^\pm$. Then since $\rho(\Gamma)$ is P^+-divergent, for any subsequence of γ_n, a further subsequence satisfies $\rho(\gamma_n)|_{\mathrm{Opp}(F_-)} \to F_+$ for flags $F_\pm$ in $\xi^\pm(\partial\Gamma)$. And, since the embeddings $\xi^\pm$ are ρ-equivariant and antipodal, $\rho(\gamma_n)\xi^+(z')$ must converge to $\xi^+(z^+)$ for any $z' \neq z^-$. The only possibility is $\xi^+(z^+) = F_+$ and $\xi^-(z^-) = F_-$.

Conversely, suppose that ρ satisfies the hypotheses of the proposition. Any infinite sequence $\gamma_n \in \Gamma$ has a subsequence such that $\gamma_n^{\pm 1} \to z^\pm$. So, if F is a P^+-limit point of the sequence $\rho(\gamma_n)$, it is the image of $\xi^+(z^+)$. Thus, the image of ξ^+ is precisely the P^+ limit set of Γ (and similarly for ξ^-). $\square$

2.4.2 Expansion dynamics

Kapovich-Leeb-Porti have also shown that Anosov representations can be understood in terms of asymptotic *expansion* behavior.

Definition 2.4.4. Let Γ be a group acting on a metric space (Z, d_Z). Fix a constant $C > 1$. We say that Γ is *C-expanding* at a point $z \in Z$ if there exists an open subset $U \subset Z$ containing z and a group element $\gamma \in \Gamma$ such that

$$d_Z(\gamma x, \gamma y) \geq C \cdot d_Z(x, y)$$

for all $x, y \in U$. If Γ is C-expanding at z for some constant $C > 1$, we just say Γ is *expanding* at z.

If $\Lambda \subset Z$ is a closed subset, we say that Γ is *expanding* at Λ if Γ is expanding at every $z \in \Lambda$.

Note that if a group Γ is expanding on a *compact* set Λ, then an iterative argument shows that for *any* constant $C > 1$, Γ is C-expanding at every $z \in \Lambda$.

In particular, this means that if Γ is expanding at a closed subset Λ of a compact smooth manifold Z equipped with some smooth metric, then Γ is also expanding at Λ with respect to the distance induced by *any* smooth metric on Z (because all such metrics are bilipschitz equivalent).

Proposition 2.4.5 ([KLP17])**.** *Let Γ be discrete subgroup of a semisimple Lie group G, and let P be a symmetric parabolic subgroup. The inclusion $\Gamma \hookrightarrow G$ is P-Anosov if and only if Γ is P-divergent, any two distinct points of the P-limit set $\Lambda_P(\Gamma)$ are opposite, and Γ is expanding at $\Lambda_P(\Gamma)$ in G/P.*

Remark 2.4.6. We do *not* need to assume that Γ is an abstract word-hyperbolic group to apply Proposition 2.4.5: part of the statement of the proposition is that word-hyperbolicity of Γ follows from the other assumptions.

2.4.3 Reduction to the projective Anosov case

When $G = \mathrm{PGL}(d, \mathbb{R})$, and $P = P_1$ is the stabilizer in $\mathrm{PGL}(d, \mathbb{R})$ of the line spanned by the first standard basis vector, then G/P is identified (via the orbit-stabilizer theorem) with real projective space $\mathbb{RP}^{d-1}$. An opposite parabolic P_{d-1} is the stabilizer of a hyperplane in $\mathbb{R}^d$, so G/P_{d-1} is identified with the *dual* projective space $(\mathbb{RP}^{d-1})^*$, thought of as the space of hyperplanes in $\mathbb{R}^d$.

In this special case, we can combine Definition 2.4.2 and Proposition 2.4.3 to give the following definition:

Definition 2.4.7. Let $\rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$ be a representation of a word-hyperbolic group Γ. The representation ρ is P_1*-Anosov* (or *projective Anosov*) if there exist equivariant embeddings

$$\xi : \partial\Gamma \to \mathbb{RP}^{d-1}, \quad \xi^* : \partial\Gamma \to (\mathbb{RP}^{d-1})^*$$

such that for any sequence $\gamma_n \in \Gamma$ with $\gamma_n^{\pm} \to z^{\pm} \in \partial\Gamma$, we have

$$\rho(\gamma_n)|_{\mathbb{RP}^{d-1} - \xi^*(z_-)} \to \xi(z_+)$$

uniformly on compacts.

It turns out that understanding P_1-Anosov representations into $\mathrm{PGL}(d, \mathbb{R})$ is enough to understand *any* P-Anosov representation into any semisimple Lie group

G. Explicitly, we have the following theorem, due to Guéritaud-Guichard-Kassel-Wienhard:

Theorem 2.4.8 ([GGKW17], Proposition 3.5). *Let G be a semisimple Lie group, and let P be a parabolic subgroup of G. There exists $d \geq 1$ and a representation $\zeta : G \to \mathrm{PGL}(d, \mathbb{R})$ such that a representation $\rho : \Gamma \to G$ is P-Anosov if and only the composition $\zeta \circ \rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$ is P_1-Anosov.*

This theorem is part of the reason that much of this thesis focuses on generalizing properties of P_1-Anosov representations in particular.

Chapter 3

Expansion dynamics for projectively convex cocompact groups

3.1 Projective convex cocompactness

The goal of this chapter is to prove a dynamical characterization of *projective convex cocompactness*. Material from this chapter appeared previously in the arXiv preprint "Dynamical properties of convex cocompact groups in projective space" [Wei20].

Before stating the theorem, we need to give the precise definition of projective convex cocompactness. We let V be a real vector space.

Definition 3.1.1. Let Ω be a properly convex domain in $\mathbb{P}(V)$. Recall that the *automorphism group* $\mathrm{Aut}(\Omega)$ is the group $\{\gamma \in \mathrm{PGL}(V) : \gamma \cdot \Omega = \Omega\}$.

For a subgroup $\Gamma \subseteq \mathrm{Aut}(\Omega)$, the *full orbital limit set* $\Lambda_\Omega(\Gamma)$ is the set of accumulation points in $\partial\Omega$ of sequences of the form $\gamma_n \cdot x$, for $x \in \Omega$ and $\gamma_n \in \Gamma$.

The *hull* of Γ, denoted $\mathrm{Hull}_\Omega(\Gamma)$, is the convex hull in Ω of the full orbital limit set.

Using Proposition 2.1.12, one can check that when Ω is a *strictly* convex domain, then the full orbital limit set $\Lambda_\Omega(\Gamma)$ is the same as the accumulation set of *any* Γ-orbit in Ω. In general, however, this is *smaller* than $\Lambda_\Omega(\Gamma)$.

Definition 3.1.2 ([DGK17])**.** Let $\Omega \subset \mathbb{P}(V)$ be a properly convex domain, and let $\Gamma \subseteq \mathrm{Aut}(\Omega)$ be discrete. We say that Γ *acts convex cocompactly on* Ω if Γ acts cocompactly on $\mathrm{Hull}_\Omega(\Gamma)$, the convex hull of the full orbital limit set of Γ.

If Γ is a discrete subgroup of $\mathrm{PGL}(V)$ acting convex cocompactly on *some* properly convex domain $\Omega \subset \mathbb{P}(V)$, we say that Γ *acts convex cocompactly in* $\mathbb{P}(V)$ or is *projectively convex cocompact*.

If Γ is an abstract group, and $\rho : \Gamma \to \mathrm{PGL}(V)$ is a representation with finite kernel such that $\rho(\Gamma)$ is projectively convex cocompact, we say that ρ is a *projectively convex cocompact representation*.

In [DGK17], Danciger-Guéritaud-Kassel proved the following theorem, which shows that projectively convex cocompact representations are closely related to Anosov representations. Zimmer independently proved a related result in [Zim21].

Theorem 3.1.3 ([DGK17])**.** *Let $\rho : \Gamma \to \mathrm{PGL}(V)$ be a representation of a word-hyperbolic group, and suppose that $\rho(\Gamma)$ preserves some properly convex domain $\Omega \subset \mathbb{P}(V)$. Let $P_1 \subset \mathrm{PGL}(V)$ be the parabolic subgroup preserving a line in V.*

Then the representation ρ is P_1-Anosov if and only if ρ is projectively convex cocompact. In this case, the image of the Anosov boundary embedding $\xi : \partial\Gamma \to \mathbb{P}(V)$ is precisely the full orbital limit set $\Lambda_\Omega(\Gamma) \subset \partial\Omega$.

Remark 3.1.4. On its face, Theorem 3.1.3 does not give a general characterization of Anosov representations in terms of convex projective structures: it only applies to P_1-Anosov representations which preserve a properly convex domain in $\mathbb{RP}^{d-1}$. In light of Theorem 2.4.8, however, *any* P-Anosov representation $\rho : \Gamma \to G$ can be viewed as a P_1-Anosov representation in some $\mathrm{PGL}(V)$ after composing ρ with an appropriate representation $\zeta : G \to \mathrm{PGL}(V)$.

The composition $\zeta \circ \rho$ might still not preserve a properly convex set in $\mathbb{P}(V)$. However, one can further compose $\zeta \circ \rho$ with the symmetric square $\tau^2 : \mathrm{PGL}(V) \to \mathrm{PGL}(\mathrm{Sym}^2 V)$. The image of $\mathrm{PGL}(V)$ in $\mathrm{PGL}(\mathrm{Sym}^2 V)$ preserves the projectivization of the set of positive definite symmetric bilinear forms in $\mathrm{Sym}^2 V$, which is a properly convex open subset of $\mathbb{P}(\mathrm{Sym}^2 V)$. Altogether, this implies the following:

Corollary 3.1.5. *Let G be a semisimple Lie group, and let P be a parabolic subgroup. There exists $d \geq 1$ and a representation $\tau : G \to \mathrm{PGL}(d, \mathbb{R})$ such that a representation $\rho : \Gamma \to G$ is P-Anosov if and only if the composition $\tau \circ \rho : \Gamma \to \mathrm{PGL}(d, \mathbb{R})$ is projectively convex cocompact.*

The goal of this chapter of the thesis is to prove a version of Theorem 3.1.3 for *non-word-hyperbolic groups*, inspired by the characterization of Anosov representations in terms of *expansion dynamics* (see Proposition 2.4.5). To that end, we make the following

Definition 3.1.6. Let Γ be a discrete subgroup of $\mathrm{PGL}(V)$ preserving a properly convex domain Ω, and let Λ be a closed Γ-invariant subset of $\partial\Omega$. Fix a Riemannian metric d_k on each Grassmannian $\mathrm{Gr}(k, V)$. We say that Γ is *expanding at the faces* of Λ if for each $x \in \Lambda$, there is a constant $C > 1$, an open neighborhood U of $\mathrm{supp}(F_\Omega(x))$ in the Grassmannian $\mathrm{Gr}(k, V)$ and a group element $\gamma \in \Gamma$ such that for all $W_1, W_2 \in U$, we have

$$d_k(\gamma W_1, \gamma W_2) \geq C \cdot d_k(W_1, W_2).$$

If the constant C can be chosen uniformly for all $x \in \Lambda$, then we say that Γ is *C-expanding* or *uniformly expanding* at the faces of Λ.

Before we can state our main theorem, we need to define two more technical conditions on subsets of boundaries of properly convex domains.

Definition 3.1.7. Let Λ be a subset of a properly convex domain Ω. We say that Λ *contains all of its faces* if for each $x \in \Lambda$, we have $F_\Omega(x) \subset \Lambda$.

We say that Λ is *boundary-convex* if any supporting hyperplane of Ω intersects Λ in a convex set.

The following is a consequence of Lemma 4.1 (1) in [DGK17]:

Lemma 3.1.8. *Let Ω be a properly convex domain, and let Γ act convex cocompactly on Ω. The full orbital limit set $\Lambda_\Omega(\Gamma)$ is closed and boundary-convex, and contains all of its faces.*

Theorem 3.1.9. *Let Ω be a properly convex domain in $\mathbb{P}(V)$, and let Γ be a discrete subgroup of $\mathrm{Aut}(\Omega)$. The following are equivalent:*

1. *Γ acts convex cocompactly on Ω.*
2. *There is a closed, Γ-invariant, and boundary-convex subset $\Lambda \subset \partial\Omega$ with nonempty convex hull such that Λ contains all of its faces and Γ is uniformly expanding at the faces of Λ.*

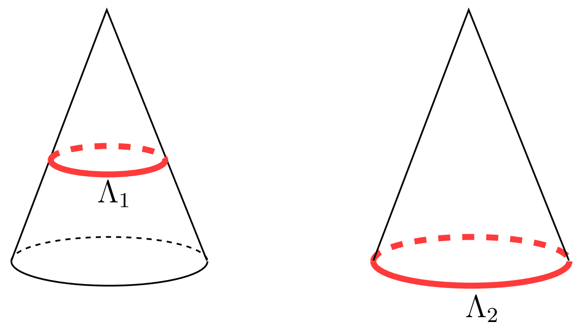

Figure 3.1: Left: $\Lambda_1 \subset \partial\Omega$ does not contain all of its faces. Right: $\Lambda_2 \subset \partial\Omega$ contains all of its faces, but is not boundary-convex: a line segment joining two points of Λ_2 intersects $\partial\Omega - \Lambda_2$.

In this case, the set Λ is the full orbital limit set $\Lambda_\Omega(\Gamma)$.

Remark 3.1.10. (a) When we prove the implication (2) $\implies$ (1) of Theorem 3.1.9, we will not actually need to assume that the expansion at the faces of Λ is uniform—only that the expansion occurs with respect to a particular choice of Riemannian metric on $\mathrm{Gr}(k, V)$. See Remark 3.2.2.

(b) The uniform expansion we get at the faces of $\Lambda_\Omega(\Gamma)$ will also allow us to give a description of the Cartan projection of certain sequences in Γ. See Proposition 3.4.14 for an exact statement.

A special case of convex cocompactness is when a discrete group $\Gamma \subset \mathrm{Aut}(\Omega)$ acts cocompactly on all of Ω. In this case, we say that Ω is *divisible*, and the group Γ *divides* the domain. As $\partial\Omega$ is always boundary convex and contains all of its faces, when $\Lambda = \partial\Omega$, Theorem 3.1.9 can be stated as the following:

Corollary 3.1.11. *Let Γ be a discrete subgroup of* $\mathrm{PGL}(d, \mathbb{R})$ *preserving a properly* *convex domain* Ω*. Then* Γ *divides* Ω *if and only if* Γ *is uniformly expanding at the faces of* $\partial\Omega$*.*

3.1.1 Outline of the chapter

Section 3.2 is devoted to the proof of the implication (2) $\implies$ (1) in Theorem 3.1.9, using a modified version of an analogous argument in [DGK17]. In Section 3.3, we recall the statement of the Benzécri cocompactness theorem for properly convex domains in $\mathbb{R}\mathrm{P}^{d-1}$, and prove a version of it relative to a direct sum decomposition of $\mathbb{R}^d$. Then in Section 3.4, we use the results of Section 3.3 to prove the implication (1) $\implies$ (2) of Theorem 3.1.9. We also prove some auxiliary technical results about certain divergent sequences in convex cocompact groups, which will be useful later in Chapter 4.

3.2 Expansion implies convex cocompactness

The goal of this section is to prove the implication (2) $\implies$ (1) of Theorem 3.1.9.

Lemma 3.2.1. *Fix a metric* $d_{\mathbb{P}}$ *on* $\mathbb{R}\mathrm{P}^{d-1}$*. Let* x *be a point in a convex domain* Ω*, and let* $W \in \mathrm{Gr}(k, d)$ *be a supporting subspace of* Ω*. For every* $y \in \mathbb{P}(W)$*, we have*

$$d_{\mathbb{P}}(x, y) \geq d_{\mathbb{P}}(x, \partial\Omega).$$

We emphasize that we do *not* need to assume that Ω is properly convex—only that it is convex in the sense of Definition 2.1.3.

Proof. Since $\mathbb{P}(W)$ is a supporting subspace, we can write

$$\mathbb{P}(W) = (\mathbb{P}(W) \cap \partial\Omega) \cup (\mathbb{P}(W) \cap (\mathbb{R}\mathrm{P}^{d-1} - \overline{\Omega})).$$

If $y \in \partial\Omega$, then the inequality is immediate. Otherwise, we can choose lifts $\tilde{\Omega}$, $\tilde{x}$, $\tilde{y}$ of Ω, x and y to the projective sphere S^{d-1} so that $d_S(\tilde{x}, \tilde{y}) = d_{\mathbb{P}}(x, y)$, where d_S is the metric on S^{d-1} induced by $d_{\mathbb{P}}$.

We can fix a closed hemisphere H of S^{d-1} containing both $\tilde{x}$ and $\tilde{y}$. The intersection $\tilde{\Omega} \cap H$ contains (at least) one lift of every point in Ω. Its closure is a closed ball in H with nonempty interior, whose boundary separates $\tilde{x}$ and $\tilde{y}$.

Thus, the distance from $\tilde{x}$ to some point $\tilde{z}$ lifting $z \in \partial\Omega$ is smaller than $d_S(\tilde{x}, \tilde{y}) = d_{\mathbb{P}}(x, y)$, implying $d_{\mathbb{P}}(x, z) < d_{\mathbb{P}}(x, y)$. □

Any (Riemannian) metric on $\mathbb{R}\mathrm{P}^{d-1}$ induces a (Riemannian) metric on each $\mathrm{Gr}(k, d)$, by viewing each $W \in \mathrm{Gr}(k, d)$ as closed subsets $\mathbb{P}(W)$ of $\mathbb{R}\mathrm{P}^{d-1}$, and taking Hausdorff distance. From this point forward, we will only work with the *angle metric* on projective space, which is induced by a choice of inner product on $\mathbb{R}^d$.

Remark 3.2.2. It is possible that a group action could be expanding at the faces of Λ with respect to some choice of Riemannian metric on $\mathrm{Gr}(k, d)$, but not with respect to another.

However, if Γ is C-expanding with respect to d for a uniform constant C, the choice of metric does not matter: since $\mathrm{Gr}(k, d)$ is compact, all Riemannian metrics on $\mathrm{Gr}(k, d)$ are bilipschitz-equivalent, and when Γ is C-expanding at the faces of Λ, one can apply expanding elements iteratively to see that Γ is also C'-expanding for an arbitrary constant C'.

When the set of supports of $(k-1)$-dimensional faces of Λ is compact in $\mathrm{Gr}(k,d)$ for each k, then a Γ-action is expanding at faces with respect to some choice of metric d if and only if it is uniformly expanding at faces with respect to that metric (and hence to every metric). For instance, this is the case when Λ is compact and does not contain any nontrivial segments (so the set of faces is the same as the set of points).

In our context, however, we will not be able to assume this kind of compactness. So, when we discuss expansion, we need to either specify the Riemannian metric or assume that the expansion is uniform.

Lemma 3.2.3. *Let $x \in \mathbb{R}\mathrm{P}^{d-1}$, and let $W \in \mathrm{Gr}(k,d)$. There exists $V \in \mathrm{Gr}(k,d)$ so that $x \in \mathbb{P}(V)$ and*

$$d_{\mathbb{P}}(x, \mathbb{P}(W)) = d_H(V, W),$$

where $d_{\mathbb{P}}$ is the angle metric on projective space, and d_H is the metric induced on $\mathrm{Gr}(k,d)$ by Hausdorff distance.

Proof. If $x \in \mathbb{P}(W)$, then we can just take $V = W$, so assume that $d_{\mathbb{P}}(x, \mathbb{P}(W)) > 0$. The definition of Hausdorff distance immediately implies that for any $\mathbb{P}(V)$ containing x, $d_H(V,W) \geq d_{\mathbb{P}}(x, \mathbb{P}(W))$, so we only need to find some V satisfying the other bound. The diameter of projective space in the angle metric is $\pi/2$, which gives an upper bound on the Hausdorff distance between any two closed subsets of $\mathbb{R}\mathrm{P}^{d-1}$. So we only need to consider the case where $d_{\mathbb{P}}(x, \mathbb{P}(W)) < \pi/2$.

In this case, we let $W' = x^{\perp} \cap W$, and then let $V = W' \oplus x$. Let $z \in \mathbb{RP}^{d-1}$ be the orthogonal projection of x onto W, so that $d_{\mathbb{P}}(x, z) = d_{\mathbb{P}}(x, \mathbb{P}(W))$. Let $\tilde{z}$ and $\tilde{x}$ be unit vector representatives of z and x, respectively, chosen so that if

$$\lambda = \langle \tilde{x}, \tilde{z} \rangle,$$

then

$$d_{\mathbb{P}}(x, z) = \cos^{-1}(\lambda).$$

Let $v \in V$. We want to show that $d_{\mathbb{P}}([v], \mathbb{P}(W)) \leq \cos^{-1}(\lambda)$, i.e. that for some $w \in W$,

$$\frac{\langle v, w \rangle}{||v|| \cdot ||w||} \geq \lambda.$$

If $v \in W$, then we can choose $w = v$. Otherwise, we can rescale v in order to write it as $w' + \tilde{x}$, for $w' \in W'$. Then let $w = w' + \tilde{z}$. Note that

$$||w|| = ||v|| = \sqrt{1 + ||w'||^2}.$$

Now we just compute:

$$\begin{aligned} \frac{\langle v, w \rangle}{||v|| \cdot ||w||} &= \frac{\langle w' + \tilde{x}, w' + \tilde{z} \rangle}{||v|| \cdot ||w||} = \frac{\langle \tilde{x}, \tilde{z} \rangle + \langle w', w' \rangle}{1 + ||w'||^2} \\ &\geq \frac{\langle \tilde{x}, \tilde{z} \rangle + \langle \tilde{x}, \tilde{z} \rangle ||w'||^2}{1 + ||w'||^2} = \langle \tilde{x}, \tilde{z} \rangle = \lambda. \end{aligned}$$

□

Remark 3.2.4. Lemma 3.2.3 still holds if we replace $d_{\mathbb{P}}$ with any metric on projective space which is an increasing function of the angle metric. In particular, the conclusion holds for any metric on projective space in which projective lines are geodesics and a maximal compact in $\mathrm{PGL}(d, \mathbb{R})$ acts by isometries.

Most of the work of proving the implication (2) $\implies$ (1) in Theorem 3.1.9 is contained in the following:

Proposition 3.2.5. *Let Ω be a convex domain preserved by a group $\Gamma \subset \mathrm{PGL}(d, \mathbb{R})$. Let C be a Γ-invariant subset of Ω, closed in Ω, with ideal boundary $\partial_i C$. Suppose that Γ is expanding at the faces of $\partial_i C$, with respect to the metrics on $\mathrm{Gr}(k, d)$ specified in Lemma 3.2.3.*

If either

1. *Γ is discrete and Ω is properly convex, or*

2. *Γ is uniformly expanding at the faces of $\partial_i C$,*

then Γ acts cocompactly on C.

Proof. Danciger-Guéritaud-Kassel [DGK17] give a proof of this fact in the case where $\partial_i C$ contains no segments, which is itself based on an argument of Kapovich, Leeb, and Porti in [KLP14] inspired by Sullivan [Sul85]. Our proof will be based on similar ideas.

We let $d_{\mathbb{P}}$ denote the angle metric on projective space, and we let d_H denote the metric on $\mathrm{Gr}(k, d)$ induced by Hausdorff distance.

For any $\varepsilon > 0$, the set

$$S_\varepsilon = \{x \in C : d_{\mathbb{P}}(x, \partial\Omega) \geq \varepsilon\}$$

is compact. So, supposing for a contradiction that the action of Γ on C is not cocompact, for a sequence $\varepsilon_n \to 0$, there exists x_n so that $\Gamma \cdot x_n$ lies in $C - S_{\varepsilon_n}$.

For any fixed constant $E > 1$, we can replace each x_n with an element in its orbit so that

$$d_{\mathbb{P}}(\gamma x_n, \partial\Omega) \leq E \cdot d_{\mathbb{P}}(x_n, \partial\Omega) \tag{3.1}$$

for all $\gamma \in \Gamma$. If Γ is discrete and Ω is properly convex, then $\Gamma \cdot x_n$ is a discrete subset of Ω and we can actually take $E = 1$. Otherwise, if the expansion at the faces of $\partial_i C$ is uniform, we pick E to be less than the uniform expansion constant.

Up to a subsequence, x_n converges in $\mathbb{RP}^{d-1}$ to some $x \in \partial_i C$. Let F be the open face of $\partial\Omega$ at x, and let $\mathbb{P}(V)$ be the support of F, for $V \in \mathrm{Gr}(k, d)$.

Let $U \subset \mathrm{Gr}(k, d)$ be an expanding neighborhood of V in $\mathrm{Gr}(k, d)$, with expanding element $\gamma \in \Gamma$ expanding by a constant $E(\gamma) > E$ on U.

Since $\partial_i C$ is compact and Γ-invariant, there is some $z_n \in \partial\Omega$ so that

$$d_{\mathbb{P}}(\gamma x_n, \gamma z_n) = d_{\mathbb{P}}(\gamma x_n, \partial\Omega).$$

Since $x_n \to x$, and the distance from γx_n to γz_n is at most ε_n, z_n converges to x as well.

Proposition 2.1.5 implies that there is *some* supporting hyperplane of Ω which intersects z_n. Any such sequence of supporting hyperplanes must subconverge to a supporting hyperplane of Ω at x. This supporting hyperplane contains $\mathbb{P}(V)$, so there is a sequence $V_n \in \mathrm{Gr}(k, d)$ with $\mathbb{P}(V_n)$ supporting Ω at z_n and V_n subconverging to V.

Since we know $\gamma \cdot z_n$ realizes the distance from $\gamma \cdot x_n$ to $\partial\Omega$, we must have

$$d_{\mathbb{P}}(\gamma x_n, \partial\Omega) \geq d_{\mathbb{P}}(\gamma x_n, \gamma\mathbb{P}(V_n)). \tag{3.2}$$

Then, Lemma 3.2.3 implies that we can choose subspaces $W_n \in \mathrm{Gr}(k, d)$, with $\mathbb{P}(W_n)$ containing x_n, so that

$$d_{\mathbb{P}}(\gamma x_n, \gamma \mathbb{P}(V_n)) = d_H(\gamma W_n, \gamma V_n). \tag{3.3}$$

Since $d_{\mathbb{P}}(\gamma x_n, \gamma \mathbb{P}(V_n))$ converges to 0, $d_H(\gamma W_n, \gamma V_n)$ does as well. Since γ is fixed, and V_n converges to V, W_n also converges to V. So eventually, both V_n and W_n lie in the $E(\gamma)$-expanding neighborhood U of V, meaning that we have

$$d_H(\gamma W_n, \gamma V_n) > E \cdot d_H(W_n, V_n). \tag{3.4}$$

The trivial bound on Hausdorff distance implies that

$$d_H(W_n, V_n) \geq d_{\mathbb{P}}(x_n, \mathbb{P}(V_n)), \tag{3.5}$$

and Lemma 3.2.1 implies that

$$d_{\mathbb{P}}(x_n, \mathbb{P}(V_n)) \geq d_{\mathbb{P}}(x_n, \partial\Omega). \tag{3.6}$$

Putting (3.2), (3.3), (3.4), (3.5), and (3.6) together, we see that

$$d_{\mathbb{P}}(\gamma x_n, \partial\Omega) > E \cdot d_{\mathbb{P}}(x_n, \partial\Omega),$$

which contradicts (3.1) above. □

We need one more lemma before we can show the main result of this section. The statement is closely related to [DGK17, Lemma 6.3], and gives a condition for when a Γ-invariant convex subset of a properly convex domain Ω contains $\mathrm{Hull}_\Omega(\Gamma)$. (The result in [DGK17] is stated for a cocompact action of a group Γ on a convex set C, but the proof only uses Γ-invariance.)

Lemma 3.2.6. *Let C be a nonempty convex set in Ω whose ideal boundary contains all of its faces, and suppose that $\Gamma \subseteq \mathrm{Aut}(\Omega)$ preserves C. Then $\partial_i C$ contains $\Lambda_\Omega(\Gamma)$, the full orbital limit set of Γ.*

In particular, if Γ is discrete, and the Γ action on C is cocompact, then the action of Γ on Ω is convex cocompact and $\partial_i C = \Lambda_\Omega(\Gamma)$.

Proof. We follow the proof of Lemma 6.3 in [DGK17].

Let $z_\infty \in \Lambda_\Omega(\Gamma)$, which is by definition the limit of a sequence $\gamma_n \cdot z$ for some $z \in \Omega$ and a sequence $\gamma_n \in \Gamma$. Fix $y \in C$, and consider the sequence $\gamma_n y$. Since $d(\gamma_n z, \gamma_n y) = d(z, y)$ for all n, Proposition 2.1.11 implies that up to a subsequence, $\gamma_n z$ and $\gamma_n y$ both converge to points in the same face of $\partial\Omega$. But any accumulation point of $\gamma_n y$ in $\partial\Omega$ lies in $\partial_i C$ and $\partial_i C$ contains its faces, so $z_\infty \in \partial_i C$.

Since $\partial_i C$ contains $\Lambda_\Omega(\Gamma)$, C must contain $\mathrm{Hull}_\Omega(\Gamma)$. [DGK17, Lemma 4.10 (3)] then implies that $\Lambda_\Omega(\Gamma) = \partial_i C$ is closed in C, which means that $\mathrm{Hull}_\Omega(\Gamma)$ is closed in C and the action on $\mathrm{Hull}_\Omega(\Gamma)$ is cocompact. $\square$

Proof of (2) $\implies$ (1) in Theorem 3.1.9. Let Ω be a properly convex domain, let Γ be a discrete subgroup of $\mathrm{Aut}(\Omega)$, and Λ be a Γ-invariant, closed and boundary-convex subset of $\partial\Omega$ with nonempty convex hull, such that Λ contains all of its faces and Γ is uniformly expanding at the faces of Λ.

Since Λ is boundary-convex and has nonempty convex hull, Λ is exactly the ideal boundary of $\mathrm{Hull}_\Omega(\Lambda)$. So, Proposition 3.2.5 implies that Γ acts cocompactly on

$\mathrm{Hull}_\Omega(\Lambda)$. Since Λ also contains its faces, applying Lemma 3.2.6 with $C = \mathrm{Hull}_\Omega(\Lambda)$ completes the proof. $\square$

3.3 A relativized Benzécri theorem

In this section we recall the statement of Benzécri's cocompactness theorem for convex projective domains, as well as prove a version of it (Proposition 3.3.4) that applies relative to a direct sum decomposition of $\mathbb{R}^d$.

3.3.1 The space of projective domains

Good references for this material include [Gol88] and [Mar14].

Let V be a real vector space. We denote the set of properly convex open subsets of $\mathbb{P}(V)$ by $\mathcal{C}(V)$. We topologize $\mathcal{C}(V)$ via the metric:

$$d(\Omega_1, \Omega_2) := d_{\mathrm{Haus}}(\overline{\Omega_1}, \overline{\Omega_2}),$$

where $d_{\mathrm{Haus}}(\cdot, \cdot)$ is the Hausdorff distance induced by any metric on $\mathbb{P}(V)$ (the choice of metric on $\mathbb{P}(V)$ does not affect the topology on $\mathcal{C}(V)$).

Definition 3.3.1. A *pointed properly convex domain* in $\mathbb{P}(V)$ is a pair (Ω, x), where $\Omega \in \mathcal{C}(V)$ and $x \in \Omega$. We denote the set of pointed properly convex domains in $\mathbb{P}(V)$ by $\mathcal{C}_*(V)$, and topologize $\mathcal{C}_*(V)$ by viewing it as a subspace of $\mathcal{C}(V) \times \mathbb{P}(V)$.

$\mathrm{PGL}(V)$ acts on both $\mathcal{C}(V)$ and $\mathcal{C}_*(V)$ by homeomorphisms. We have the following important result:

Theorem 3.3.2 (Benzécri, [Ben60])**.** *The action of* $\mathrm{PGL}(V)$ *on* $\mathcal{C}_*(V)$ *is proper and cocompact.*

3.3.2 Benzécri relative to a direct sum

We now let V_a, V_b be subspaces of V so that $V_a \oplus V_b = V$. The decomposition induces natural projection maps $\pi_{V_a} : V \to V_a$ and $\pi_{V_b} : V \to V_b$, as well as a decomposition of the dual V^* into $V_a^* \oplus V_b^*$. Here V_a^*, V_b^* are respectively identified with the linear functionals on V which vanish on V_b, V_a.

When Ω is a convex subset of $\mathbb{P}(V)$ which is disjoint from $\mathbb{P}(V_b)$, we let $\pi_{V_a}(\Omega)$ be the projectivization of $\pi_{V_a}(\tilde{\Omega})$, where $\tilde{\Omega}$ is a cone over Ω. A priori this is only a convex subset of $\mathbb{P}(V_a)$, although we will see (Proposition 3.3.6) that if Ω is properly convex and open, and $\overline{\Omega}$ is disjoint from $\mathbb{P}(V_b)$, then $\pi_{V_a}(\Omega)$ is properly convex and open in $\mathbb{P}(V_a)$.

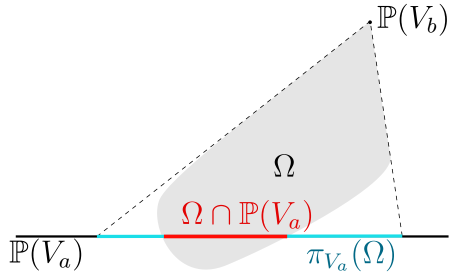

Figure 3.2: The domains $\Omega \cap \mathbb{P}(V_a)$ and $\pi_{V_a}(\Omega)$. In this case, $\pi_{V_a}(\Omega)$ is properly convex even though $\overline{\Omega}$ intersects $\mathbb{P}(V_b)$.

We remark that if $\Omega \cap \mathbb{P}(V_b)$ is nonempty, then $\pi_{V_a}(\Omega)$ is not even well-defined. On the other hand, if $\overline{\Omega} \cap \mathbb{P}(V_b)$ is nonempty, but $\Omega \cap \mathbb{P}(V_b)$ is empty, then $\pi_{V_a}(\Omega)$ does exist, and may or may not be a properly convex subset of $\mathbb{P}(V_a)$.

Definition 3.3.3. Let $V = V_a \oplus V_b$, and let $\mathcal{K}_a$ be a subset of $\mathcal{C}_*(V_a)$. We define the subset $\mathcal{C}_*(V_a, V_b, \mathcal{K}_a)$ by

$$\mathcal{C}_*(V_a, V_b, \mathcal{K}_a) := \left\{ (\Omega, x) \in \mathcal{C}_*(V) : \begin{array}{l} \mathbb{P}(V_b) \cap \Omega = \emptyset, \\ (\Omega \cap \mathbb{P}(V_a), x) \in \mathcal{K}_a, \\ (\pi_{V_a}(\Omega), x) \in \mathcal{K}_a \end{array} \right\}.$$

The groups $\mathrm{GL}(V_a)$ and $\mathrm{GL}(V_b)$ both have a well-defined action on $\mathcal{C}_*(V)$: we take $g \in \mathrm{GL}(V_a)$ and $h \in \mathrm{GL}(V_b)$ to act by the projectivizations of $g \oplus \mathrm{id}_{V_b}$, $\mathrm{id}_{V_a} \oplus h$ respectively, on $\mathbb{P}(V_a \oplus V_b)$.

Since the $\mathrm{GL}(V_b)$-action on $\mathbb{P}(V)$ fixes $\mathbb{P}(V_a)$ pointwise and commutes with projection to $\mathbb{P}(V_a)$, for any $\mathcal{K}_a \subset \mathcal{C}_*(V_a)$, $\mathrm{GL}(V_b)$ acts on the subset $\mathcal{C}_*(V_a, V_b, \mathcal{K}_a)$. The main result of this section is the following:

Proposition 3.3.4. *Let V_a, V_b be subspaces of a real vector space V such that $V_a \oplus V_b = V$. For any compact subset $\mathcal{K}_a \subset \mathcal{C}_*(V_a)$, the action of $\mathrm{GL}(V_b)$ on $\mathcal{C}_*(V_a, V_b, \mathcal{K}_a)$ is proper and cocompact.*

3.3.3 Convex cones in direct sums

Before proving Proposition 3.3.4, we explore some of the properties of convex cones in a vector space V which splits as a direct sum $V = V_a \oplus V_b$.

3.3.3.1 Duality

Suppose that C_a is a convex cone in $V - \{0\}$, for $V = V_a \oplus V_b$. The intersection $C_a^* \cap V_a^*$ consists of functionals in C_a^* which vanish on V_b. If we know that C_a lies inside of V_a, then any functional on V_a which does not vanish anywhere on $\overline{C_a} - \{0\}$ can be extended by zero on V_b to get an element of $C_a^* \cap V_a^*$. So in this case, $C_a^* \cap V_a^*$ is canonically indentified with the dual of the cone C_a *viewed as a cone in* V_a.

This view allows us to understand projection and intersection as dual operations on convex cones in $V = V_a \oplus V_b$, in the following sense:

Lemma 3.3.5. *Let V be a real vector space with $V = V_1 \oplus V_2$, and let C be a convex cone intersecting V_2 trivially. Then $C^* \cap V_1^* \subseteq \pi_{V_1}(C)^* \cap V_1^*$ and*

$$\overline{\pi_{V_1}(C)^*} \cap V_1^* \subseteq \overline{C^*} \cap V_1^*.$$

Moreover, if $\overline{C} - \{0\}$ intersects V_2 trivially, then in fact

$$\pi_{V_1}(C)^* \cap V_1^* = C^* \cap V_1^*.$$

Proof. First let $\alpha \in C^* \cap V_1^*$. Let v be any nonzero element of the closure of $\pi_{V_1}(C)$, so that $v + v_2 \in \overline{C}$ for some $v_2 \in V_2$. We know that $\alpha(v + v_2) \neq 0$ and $\alpha(v_2) = 0$, so $\alpha(v) \neq 0$. This shows that α is in $\pi_{V_1}(C)^*$.

Now let $\alpha \in \overline{\pi_{V_1}(C)^*} \cap V_1^* - \{0\}$, and let $v \in C$. We can write $v = v_1 + v_2$ for $v_1 \in V_1$, $v_2 \in V_2$; since we assume C does not intersect V_2, v_1 is nonzero. Then since $\alpha \in V_1^*$, $\alpha(v) = \alpha(v_1) \neq 0$. So, $\alpha \in \overline{C^*}$.

If we further assume that $\overline{C} \cap V_2 = \{0\}$, a similar argument shows that any $\alpha \in \pi_{V_1}(C)^* \cap V_1^*$ is nonzero on any $v \in \overline{C} - \{0\}$, implying $\alpha \in C$. $\square$

As a consequence of the above, we note:

Proposition 3.3.6. *Let C be a sharp (Definition 2.1.2) open convex cone in a vector space $V = V_1 \oplus V_2$. If $\overline{C} - \{0\}$ intersects V_2 trivially, then the projection $\pi_{V_1}(C)$ is sharp and open in V_1.*

Proof. Openness is immediate since projection is an open map. Since C is sharp, if $\overline{C}$ does not intersect V_2, then there is some $\alpha \in V^*$ whose kernel contains V_2 and does not intersect $\overline{C}$, i.e. $\alpha \in C^* \cap V_1^*$. Since non-intersection with $\overline{C}$ is an open condition, $C^* \cap V_1^*$ is a nonempty open subset of V_1^*. Then Lemma 3.3.5 implies that $\pi_{V_1}(C)^* \cap V_1^*$ is nonempty and open in V_1^*. So its dual in $V_1^{**} = V_1$ is sharp. □

3.3.3.2 Convex hulls

If Ω_1, Ω_2 are properly convex subsets of $\mathbb{P}(V)$, we cannot always find a minimal properly convex subset $\Omega \subset \mathbb{P}(V)$ which contains $\Omega_1 \cup \Omega_2$ (that is, convex hulls do not always exist). Here we describe some circumstances under which this is possible.

Definition 3.3.7. Let Ω_1, Ω_2 be properly convex sets in $\mathbb{P}(V)$. For each $W \in \Omega_1^* \cap \Omega_2^*$, we let $\mathrm{Hull}_W(\Omega_1, \Omega_2)$ denote the convex hull of Ω_1 and Ω_2 in the affine chart $\mathbb{P}(V) - \mathbb{P}(W)$.

The set $\mathrm{Hull}_W(\Omega_1, \Omega_2)$ is minimal among all convex subsets of $\mathbb{P}(V) - \mathbb{P}(W)$ containing $\Omega_1 \cup \Omega_2$. However, it is possible that for some other $W' \in \Omega_1^* \cap \Omega_2^*$, $\mathrm{Hull}_{W'}(\Omega_1, \Omega_2)$ is not contained in $\mathbb{P}(V) - \mathbb{P}(W)$. So, to guarantee minimality among all convex subsets of $\mathbb{P}(V)$, we need a little more:

Lemma 3.3.8. *If $\Omega_1 \cap \Omega_2$ is nonempty, then for any $W \in \Omega_1^* \cap \Omega_2^*$,* $\mathrm{Hull}_W(\Omega_1, \Omega_2)$ *is the unqiue minimal properly convex subset of $\mathbb{P}(V)$ containing $\Omega_1 \cup \Omega_2$.*

Proof. Let A be the affine chart $\mathbb{P}(V) - \mathbb{P}(W)$, and let H be any properly convex set containing $\Omega_1 \cup \Omega_2$. Since $\Omega_1 \cap \Omega_2$ is nonempty, $\Omega_1 \cup \Omega_2$ is a connected subset of A, so it is contained in a single connected component C of $H \cap A$. This component is a convex subset of A, so by definition C (hence H) contains $\mathrm{Hull}_W(\Omega_1, \Omega_2)$. $\square$

Lemma 3.3.8 allows us to define the convex hull of a pair of properly convex sets without reference to a particular affine chart.

Definition 3.3.9. When Ω_1, Ω_2 are properly convex sets such that $\Omega_1 \cap \Omega_2$ and $\Omega_1^* \cap \Omega_2^*$ are both nonempty, we let $\mathrm{Hull}(\Omega_1, \Omega_2)$ denote the minimal properly convex set containing $\Omega_1 \cup \Omega_2$.

3.3.4 Proving Benzécri for direct sums

We can now begin proving Proposition 3.3.4. As a first step, we consider the case where $\dim V_a = 1$, i.e. $\mathbb{P}(V_a)$ is identified with a single point in $\mathbb{P}(V)$.

Lemma 3.3.10. *Let $V = V_b \oplus x$ for a point $x \in \mathbb{P}(V)$. $\mathrm{GL}(V_b)$ acts cocompactly on the set of domains*

$$\mathcal{C}_*(x, V_b) := \{(\Omega, x) \in \mathcal{C}_*(V) : \mathbb{P}(V_b) \cap \Omega = \emptyset\}.$$

Proof. Let (Ω_n, x) be a sequence of domains in $\mathcal{C}_*(x, V_b)$. The Benzécri cocompactness theorem (Theorem 3.3.2) implies that we can find group elements $g_n \in \mathrm{PGL}(V)$ so that the sequence of pointed domains

$$(g_n\Omega_n, g_n x)$$

subconverges to a pointed domain (Ω, x'). We want to show that these group elements can be chosen to preserve the decomposition $V_b \oplus x$.

We view V_b as a point in $\overline{\Omega_n^*}$, so $g_n V_b$ lies in $g_n \overline{\Omega_n^*}$ for all n, and a subsequence of $g_n V_b$ converges to some $W \in \overline{\Omega^*}$. In particular, $\mathbb{P}(W)$ does not contain x'. This means that we can find a sequence of group elements g_n', lying in a fixed compact subset of $\mathrm{PGL}(V)$, so that

$$g_n' \cdot g_n V_b = V_b, \quad g_n' \cdot g_n x = x.$$

Since the g_n' lie in a compact subset of $\mathrm{PGL}(V_b)$, the domains

$$g_n' g_n \Omega_n$$

must also subconverge to some properly convex domain Ω', which contains x. So we can replace g_n with $g_n' g_n$ to get the desired sequence of group elements. $\square$

Lemma 3.3.10 gets us partway to proving Proposition 3.3.4. We see that if Ω is any domain in $\mathcal{C}_*(V_a, V_b, \mathcal{K}_a)$, we can always find some $h \in \mathrm{GL}(V_b)$ so that $h\Omega \cap \mathbb{P}(V_b \oplus x)$ lies in a fixed compact set of domains in $\mathcal{C}(V_b \oplus x)$. This is almost enough to ensure that $h\Omega$ itself lies in a fixed compact set of domains in $\mathcal{C}(V)$. The exact condition we'll need is the following:

Lemma 3.3.11. *Let V be a real vector space, and suppose $V = W_a \oplus V_b \oplus x$, for a point $x \in \mathbb{P}(V)$.*

Let Ω_a, Ω_a' be properly convex domains in $\mathbb{P}(W_a \oplus x)$, and let Ω_b, Ω_b' be properly convex domains in $\mathbb{P}(V_b \oplus x)$, such that

$$x \in \Omega_a \subset \Omega_a',$$

$$x \in \Omega_b \subset \Omega_b'.$$

There exist properly convex domains Ω_1, Ω_2 in $\mathbb{P}(V)$ such that any $\Omega \in \mathcal{C}(V)$ disjoint from $\mathbb{P}(W_a)$ and $\mathbb{P}(V_b)$ which satisfies:

1. $\Omega'_a \supset \pi_{W_a \oplus x}(\Omega)$,

2. $\Omega_a \subset \Omega \cap \mathbb{P}(W_a \oplus x)$,

3. $\Omega'_b \supset \pi_{V_b \oplus x}(\Omega)$

4. $\Omega_b \subset \Omega \cap \mathbb{P}(V_b \oplus x)$,

also satisfies $x \in \Omega_1$ and $\Omega_1 \subset \Omega \subset \Omega_2$.

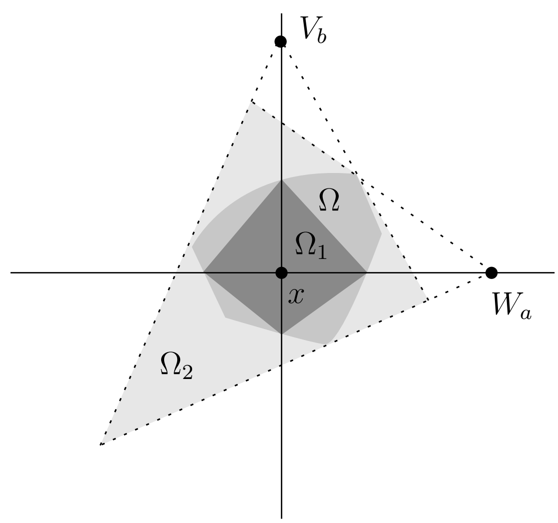

Figure 3.3: Ω fits between a pair of domains Ω_1 and Ω_2, which depend only on the intersections and projections between Ω and $\mathbb{P}(W_a \oplus x)$, $\mathbb{P}(V_b \oplus x)$.

Proof. We know $\Omega_a \cap \Omega_b = \{x\}$. If necessary, we can slightly shrink Ω_a and Ω_b so that $\overline{\Omega_a} \cap \mathbb{P}(V_a) = \emptyset$ and $\overline{\Omega_b} \cap \mathbb{P}(V_b) = \emptyset$, which means that $\mathbb{P}(W_a \oplus V_b)$ can be identified with an element of $\Omega_a^* \cap \Omega_b^*$. So, the convex hull (Definition 3.3.9) of Ω_a, Ω_b exists, and we can take

$$\Omega_1 = \mathrm{Hull}(\Omega_a, \Omega_b).$$

This is a well-defined convex domain containing x. Ω_1 is open because $W_a \oplus x$ and $V_b \oplus x$ span V.

To build Ω_2, we consider the *relative dual domains*

$$D(\Omega_a') = (\Omega_a')^* \cap \mathbb{P}((W_a \oplus x)^*), \quad D(\Omega_b') = (\Omega_b')^* \cap \mathbb{P}((V_b \oplus x)^*).$$

$D(\Omega_a')$, $D(\Omega_b')$ can also be obtained by taking the duals of Ω_a' and Ω_b' viewed as convex subsets of $\mathbb{P}(W_a \oplus x)$, $\mathbb{P}(V_b \oplus x)$. Since Ω_a' and Ω_b' are properly convex subsets of $\mathbb{P}(W_a \oplus x)$ and $\mathbb{P}(V_b \oplus x)$, $D(\Omega_a')$ and $D(\Omega_b')$ are open in $\mathbb{P}((W_a \oplus x)^*)$ and $\mathbb{P}((V_b \oplus x)^*)$.

We also know that x lies in $D(\Omega_a')^* \cap D(\Omega_b')^*$. So we can define the convex open set Ω_2^* to be the interior of

$$\mathrm{Hull}_x(D(\Omega_a'), D(\Omega_b')),$$

using Definition 3.3.7. The interior of this hull is nonempty because $(W_a \oplus x)^*$ and $(V_b \oplus x)^*$ span V^*.

Let Ω be any domain satisfying the hypotheses of the lemma. Since duality reverses inclusions, we know $D(\Omega_a') \subseteq \pi_{W_a \oplus x}(\Omega)^*$ and $D(\Omega_b') \subseteq \pi_{V_a \oplus x}(\Omega)^*$. Then, Lemma 3.3.5 implies

$$D(\Omega_a') \subseteq \overline{\pi_{W_a \oplus x}(\Omega)^*} \cap \mathbb{P}((W_a \oplus x)^*) \subseteq \overline{\Omega^*} \cap \mathbb{P}((W_a \oplus x)^*),$$
$$D(\Omega_b') \subseteq \overline{\pi_{V_a \oplus x}(\Omega)^*} \cap \mathbb{P}((V_b \oplus x)^*) \subseteq \overline{\Omega^*} \cap \mathbb{P}((V_b \oplus x)^*).$$

In particular, $D(\Omega'_a)$ and $D(\Omega'_b)$ are both contained in $\overline{\Omega^*}$. Since $\Omega^{**} = \Omega$ contains x, $\overline{\Omega^*}$ is contained in the affine chart $\mathbb{P}(V^*) - \mathbb{P}(x)$. So, $\overline{\Omega^*}$ contains the closure of

$$\mathrm{Hull}_x(D(\Omega'_a), D(\Omega'_b)),$$

meaning Ω^* contains Ω_2^* and Ω is contained in the properly convex set $\Omega_2 = \Omega_2^{**}$. $\square$

Remark 3.3.12. If $\overline{\Omega}_b$ does not intersect $\mathbb{P}(V_b)$ and $\overline{\Omega}_a$ does not intersect $\mathbb{P}(W_a)$, we can work in the affine chart $\mathbb{P}(V) - \mathbb{P}(W_a \oplus V_b)$, and Lemma 3.3.11 is equivalent to the fact that if a convex subset C of an affine space has open and bounded projections to and intersections with a pair of complementary affine subspaces, C is itself open and bounded in terms of the size of the projections and intersections.

We do not take this approach because we do *not* want to assume that $\overline{\Omega}_b$ and $\mathbb{P}(V_b)$ are disjoint.

Our next task is to show that we can sometimes replace assumption (3) in Lemma 3.3.11 with:

(3a) $\Omega'_b \supset \Omega \cap \mathbb{P}(V_b \oplus x)$.

This will be done in Proposition 3.3.14 below. We start with some Euclidean geometry.

We endow $\mathbb{R}^d$ with its standard inner product. For a subspace $W \subseteq \mathbb{R}^d$, we let $\pi_W : \mathbb{R}^d \to W$ denote the orthogonal projection, and for $R > 0$, let $B(R)$ denote the open ball around the origin of radius R.

Lemma 3.3.13. *Let Ω be a convex subset of $\mathbb{R}^d$ containing the origin, and let W be a subspace of $\mathbb{R}^d$.*

Suppose that there are $R_1, R_2 > 0$ so that:

- $B(R_1) \cap W^\perp \subset \Omega \cap W^\perp$,
- $\pi_{W^\perp}(\Omega) \subset B(R_2)$.

Then there exists a linear map $f : \mathbb{R}^d \to \mathbb{R}^d$, depending only on R_1 and R_2, so that $\pi_W(\Omega) \subset f(\Omega \cap W)$.

Proof. Let p be any point in $\pi_W(\Omega)$, and let z be some point in Ω so that $\pi_W(z) = p$. We can write $z = p + y$ for $y \in \pi_{W^\perp}(\Omega)$.

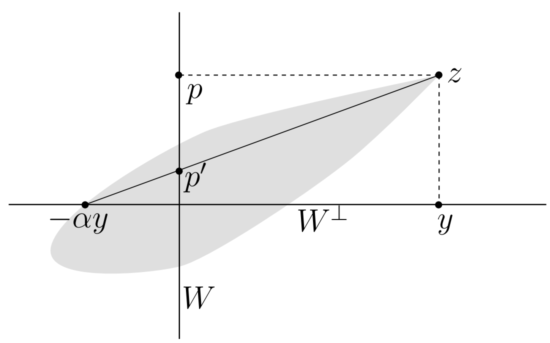

Figure 3.4: Illustration for the proof of Lemma 3.3.13. The ratio $||p||/||p'||$ is bounded in terms of α.

Let ℓ be the line through the origin passing through y. For some $\alpha > 0$, we know that ℓ intersects $\overline{(\Omega \cap W^\perp)} - B(R_1)$ at $y' = -\alpha y$. Note that

$$\alpha = \frac{||y'||}{||y||} > \frac{R_1}{R_2}.$$

Since Ω is convex and contains $\Omega \cap W^{\perp}$, it contains the open line segment

$$\{t(-\alpha y) + (1-t)(y+p) : t \in (0,1)\}.$$

This line segment passes through W when $t = \frac{1}{1+\alpha}$, meaning that Ω must contain the point

$$p' = \left(1 - \frac{1}{1+\alpha}\right) p.$$

Since Ω contains the origin, it also contains

$$\left(1 - \frac{1}{1 + R_1/R_2}\right) p = \frac{R_1}{R_1 + R_2} p.$$

This point lies in $\Omega \cap W$, meaning that p lies in $R_3 \cdot (\Omega \cap W)$ where

$$R_3 := \frac{R_1 + R_2}{R_1}.$$

So we can take our map f to be the linear rescaling about the origin by R_3. $\square$

Proposition 3.3.14. *Let $V = W_a \oplus V_b \oplus x$, for $x \in \mathbb{P}(V)$.*

Let Ω_a, Ω'_a be properly convex domains in $\mathbb{P}(W_a \oplus x)$, and let Ω''_b be a properly convex domain in $\mathbb{P}(V_b \oplus x)$ such that

$$x \in \Omega_a \subset \Omega'_a, \quad x \in \Omega''_b.$$

If $\overline{\Omega'_a}$ does not intersect $\mathbb{P}(W_a)$, then there exists a properly convex domain Ω'_b in $\mathbb{P}(V_b \oplus x)$ so that any $\Omega \in \mathcal{C}(V)$ which satisfies $\Omega \cap \mathbb{P}(V_b) = \emptyset$ and

1. *$\Omega'_a \supset \pi_{W_a \oplus x}(\Omega)$,*

2. *$\Omega_a \subset \Omega \cap \mathbb{P}(W_a \oplus x)$,*

(3a) $\Omega''_b \supset \Omega \cap \mathbb{P}(V_b \oplus x)$

also satisfies

(3) $\Omega'_b \supset \pi_{V_b \oplus x}(\Omega)$.

Proof. Let $H = W_a \oplus V_b$, and consider the affine chart $A = \mathbb{P}(V) - \mathbb{P}(H)$. We can choose coordinates and a Euclidean metric on this affine chart so that $W_a \oplus x$ and $V_b \oplus x$ map to complementary orthogonal subspaces W_a, V_b of A, meeting at the origin. In these coordinates, the projectivizations of the projection maps $\pi_{W_a \oplus x}$, $\pi_{V_b \oplus x}$ correspond to the orthogonal projections to W_a and V_b, respectively.

Since $\mathbb{P}(W_a)$ does not intersect $\overline{\Omega'_a}$, the images of Ω'_a and Ω_a in A are both bounded open convex subsets of W_a.

Let Ω be a properly convex domain not intersecting $\mathbb{P}(V_b)$ and satisfying assumptions (1), (2), (3a). Since $\pi_{W_a \oplus x}(\Omega)$ is contained in A, Ω cannot intersect $\mathbb{P}(W_a \oplus V_b)$, so Ω is contained in the affine chart A (although its closure need not be).

In particular, $\Omega \cap \mathbb{P}(V_b \oplus x)$ is contained in the unique connected component of $\Omega''_b \cap A$ which contains x. So, by replacing Ω''_b with this connected component, we may assume that the image of Ω''_b in A is a convex open subset of V_b.

Lemma 3.3.13 then implies that there is an affine map $f : A \to A$, depending only on Ω_a and Ω'_a, so that

$$\pi_{V_b \oplus x}(\Omega) \subseteq f(\Omega''_b).$$

So, we can take Ω'_b to be the properly convex domain $f(\Omega''_b)$. $\square$

We are now ready to prove Proposition 3.3.4.

Proof of Proposition 3.3.4. Properness follows immediately from the standard Benzécri theorem (Theorem 3.3.2), since the restriction of a proper action of a group G on X to a closed subgroup H and an H-invariant subset of X is always proper.

We let V_a, V_b, and $\mathcal{K}_a \subset \mathcal{C}_*(V_a)$ be as in the statement of the theorem. Let (Ω_n, x_n) be a sequence of properly convex domains in $\mathcal{C}_*(V_a, V_b, \mathcal{K}_a)$. We can choose a subsequence so that $x_n \to x$. Our goal is to find a pair of properly convex domains Ω_1, Ω_2 (with $x \in \Omega_1$) and $h_n \in \mathrm{GL}(V_b)$, so that up to a subsequence,

$$\Omega_1 \subset h_n \cdot \Omega_n \subset \Omega_2.$$

This will be sufficient, because $x_n \in \mathbb{P}(V_a)$, so $h_n x_n = x_n$ converges to x and $h_n \Omega_n$ subconverges to some properly convex domain Ω containing $\Omega_1 \ni x$.

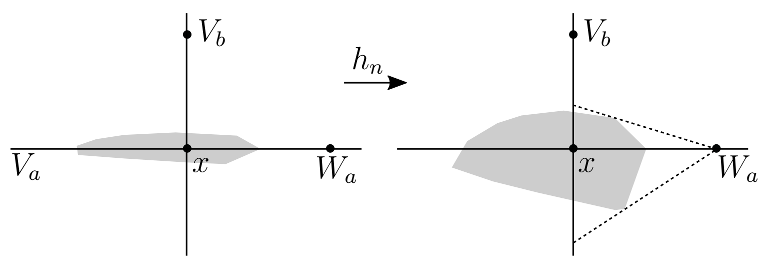

Figure 3.5: Applying an element $h_n \in \mathrm{GL}(V_b)$ "rescales" in $\mathbb{P}(V_b \oplus x)$ about x; if the size of the intersection $\Omega \cap \mathbb{P}(V_b \oplus x)$ is bounded, then the size of the projection to $V_b \oplus x$ (with respect to the decomposition $V = W_a \oplus V_b \oplus x$) is also bounded (Proposition 3.3.14).

Consider the sequence of domains $\Omega'_n = \Omega_n \cap \mathbb{P}(V_b \oplus x)$. We know $\mathbb{P}(V_b)$ is disjoint from Ω_n for all n. So, Lemma 3.3.10 implies that we can find $h_n \in \mathrm{GL}(V_b)$ so that the domains $h_n\Omega'_n$ subconverge in $\mathcal{C}(V_b \oplus x)$ to some domain Ω' in $\mathbb{P}(V_b \oplus x)$. In particular, up to a subsequence, we can find fixed domains $\Omega_b, \Omega''_b \subset \mathbb{P}(V_b \oplus x)$ such that for all n,

$$x \in \Omega_b \subset h_n\Omega'_n \subset \Omega''_b.$$

Since the intersections $\Omega_n \cap \mathbb{P}(V_a)$ and projections $\pi_{V_a}(\Omega_n)$ both lie in a fixed compact set in $\mathcal{C}(V_a)$, we can also assume that there are domains $\Omega_a, \Omega'_a \in \mathcal{C}(V_a)$ so that for all n,

$$\Omega_a \subset \Omega_n \cap \mathbb{P}(V_a), \quad \Omega'_a \supset \pi_{V_a}(\Omega_n).$$

Since the action of any $h_n \in \mathrm{GL}(V_b)$ fixes V_a pointwise and commutes with projection to V_a, this immediately implies that for all n,

$$\Omega_a \subset h_n\Omega_n \cap \mathbb{P}(V_a), \quad \Omega'_a \supset \pi_{V_a}(h_n\Omega_n).$$

Fix a subspace $W_a \subset V_a$ so that $V_a = W_a \oplus x$ and $\mathbb{P}(W_a)$ does not intersect the closure of Ω'_a. This allows us to define a projection map $\pi_{V_b \oplus x} : V \to V_b \oplus x$, whose kernel is W_a. Proposition 3.3.6 implies that $\pi_{V_b\oplus x}(h_n\Omega_n)$ is a properly convex open subset of $\mathbb{P}(V_b \oplus x)$, and Proposition 3.3.14 implies that for all n, $\pi_{V_b\oplus x}(h_n\Omega_n)$ is contained in a properly convex domain $\Omega'_b \subset \mathbb{P}(V_b \oplus x)$, depending only on Ω_a, Ω'_a, and Ω''_b. Then we can apply Lemma 3.3.11 to the domains $\Omega_a, \Omega'_a, \Omega_b$, Ω'_b to finish the proof. $\square$

3.4 Cocompactness implies expansion

The main goal of this section is to prove the implication (1) $\implies$ (2) of Theorem 3.1.9. In fact we will prove a slightly more general statement:

Proposition 3.4.1. *Let C be a convex subset of a properly convex domain Ω, and suppose that $\Gamma \subseteq \mathrm{Aut}(\Omega)$ acts cocompactly on C. Then Γ is uniformly expanding at the faces of the ideal boundary of C.*

Afterwards, we will use some of the ideas arising in the proof to show that a version of "north-south dynamics" holds for certain sequences of elements in a convex cocompact group (Proposition 3.4.13). We also describe the behavior of the *Cartan projection* of those sequences.

3.4.1 Pseudo-loxodromic elements

Our main inspiration comes from an observation in Sullivan's study [Sul79] of conformal densities on $\mathbb{H}^d$: if γ is any isometry of $\mathbb{H}^d$, and x is any point in $\mathbb{H}^d$, then γ expands a small ball in $\partial\mathbb{H}^d$ at the endpoint of the geodesic ray from x to $\gamma^{-1}x$, with expansion constant related to $d(x, \gamma^{-1}x)$.

This observation relies on the fact that, given distinct points $x, y \in \mathbb{H}^d$, there is a loxodromic isometry taking x to y whose axis is the geodesic joining x and y. The exact analogue of this fact for properly convex domains does not hold in general, since there is no reason to expect even the full automorphism group of a properly convex domain to act transitively on the domain. However, instead of looking for actual automorphisms of the domain, we can instead look for elements of $\mathrm{PGL}(d, \mathbb{R})$ that don't perturb the domain "too much." We make this precise below.

Definition 3.4.2. Let $\Omega \subset \mathbb{RP}^{d-1}$ be a properly convex domain, and let $\mathcal{K}$ be a compact subset of $\mathcal{C}(\mathbb{R}^d)$ containing Ω. An element $g \in \mathrm{PGL}(d, \mathbb{R})$ is a *$\mathcal{K}$-pseudo-automorphism* of Ω if $g\Omega \in \mathcal{K}$.

Definition 3.4.3. Let $\Omega \subset \mathbb{RP}^{d-1}$ be a properly convex domain. For a compact subset $\mathcal{K} \subset \mathcal{C}(\mathbb{R}^d)$ containing Ω, we say that a $\mathcal{K}$-pseudo-automorphism $g \in \mathrm{PGL}(d, \mathbb{R})$ is $\mathcal{K}$-*pseudo-loxodromic* if there is a g-invariant direct sum decomposition

$$\mathbb{R}^d = V_- \oplus V_0 \oplus V_+,$$

where V_-, V_+ are positive eigenspaces of g and supporting subspaces of Ω, and $\mathbb{P}(V_- \oplus V_+)$ intersects every $\Omega' \in \mathcal{K}$.

The subspaces V_- and V_+ are referred to as *endpoints* of g. The projective subspace $\mathbb{P}(V_- \oplus V_+)$ is the *axis* of the pseudo-loxodromic, and V_0 is the *neutral subspace*.

A pseudo-loxodromic element preserves its axis $\mathbb{P}(V_- \oplus V_+)$. When V_- and V_+ are points in $\mathbb{RP}^{d-1}$, this axis is an actual projective line.

We do *not* assume that an individual pseudo-loxodromic element attracts points on its axis towards either of its endpoints, since we are only interested in the dynamics of *sequences* of pseudo-loxodromics.

If g_n is a sequence of $\mathcal{K}$-pseudo-loxodromic elements with common endpoints, then, up to a subsequence, the domains $g_n\Omega$ converge to a domain Ω_∞ in $\mathcal{K}$ which intersects the common axis. In fact, we observe:

Proposition 3.4.4. *Let g_n be a sequence of $\mathcal{K}$-pseudo-loxodromic elements with common endpoints V_+, V_-. If $g_n\Omega$ converges to Ω_∞, then Ω_∞ contains the relative interior of the convex hull (in Ω) of $V_+ \cap \partial\Omega$ and $V_- \cap \partial\Omega$.*

Proof. This convex hull is invariant under g_n, so it is contained in $\overline{\Omega_\infty}$, and since Ω_∞ intersects the axis $\mathbb{P}(V_+ \oplus V_-)$ the conclusion follows. $\square$

Definition 3.4.5. Let Ω be a properly convex domain, and let g_n be a sequence of $\mathcal{K}$-pseudo-loxodromic elements with common endpoints V_+, V_- and common neutral subspace V_0. We say that V_- is a *repelling endpoint* of the sequence g_n if there is a sequence

$$x_n \in \Omega \cap \mathbb{P}(V_- \oplus V_+)$$

such that $g_n x_n = x$ for some $x \in \Omega$, and $x_n \to x_- \in \partial\Omega$ with

$$V_- = \operatorname{supp} F_\Omega(x_-).$$

3.4.2 Existence of repelling pseudo-loxodromics

We will use pseudo-loxodromics to state an analogue (Lemma 3.4.7) of the fact that any two points in $\mathbb{H}^d$ can be joined by the axis of a loxodromic isometry. First, we need a lemma:

Lemma 3.4.6. *Let x_+, x_- be a pair of points in the boundary of a properly convex domain $\Omega \subset \mathbb{P}(V)$ such that $(x_-, x_+) \subseteq \Omega$. Let $\mathbb{P}(H_+), \mathbb{P}(H_-)$ be supporting hyperplanes of Ω at x_+, x_-. Let $V_- = \operatorname{supp}(F_\Omega(x_-))$, and let $W = V_- \oplus x_+$.*

There exists a subspace $H_0 \subset H_+ \cap H_-$ such that

1. *$H_- = H_0 \oplus V_-$, and*
2. *$\pi_W(\Omega)$ is properly convex, where $\pi_W : V \to W$ is the projection with kernel H_0.*

Note that while $\mathbb{P}(H_0)$ does not intersect Ω, the intersection $\mathbb{P}(H_0) \cap \overline{\Omega}$ may be nonempty.

Proof. First suppose that $V_- = x_-$. In this case, we take $H_0 = H_+ \cap H_-$, and $\pi_W(\Omega)$ is exactly the line segment (x_-, x_+).

So now suppose that V_- is not a single point, and consider the properly convex set $\Omega_- = \partial\Omega \cap H_-$. $H_+ \cap H_-$ is a codimension-one projective subspace of H_-. Because $(x_-, x_+) \subseteq \Omega$, $H_+ \cap H_-$ does not contain V_-.

$H_+ \cap H_-$ intersects $\overline{\Omega_-}$ in a (possibly empty) properly convex set. So, since the projective subspace V_- has dimension $k \geq 1$, there exists a codimension-k projective subspace of $\mathbb{P}(H_+ \cap H_-)$ which does not intersect $\overline{\Omega_-}$ or V_-. Let $\mathbb{P}(H_0)$ be such a subspace; since $\mathbb{P}(H_0)$ is disjoint from $\overline{\Omega}$, we are done by Proposition 3.3.6. $\square$

The following lemma is the main technical result in this section. It implies in particular that every open face in the boundary of a properly convex domain is the repelling endpoint of *some* sequence of $\mathcal{K}$-pseudo-loxodromics.

Lemma 3.4.7. *Let Ω be a properly convex domain, let $x_- \in \partial\Omega$, and let L be a projective line intersecting Ω, joining x_- with some $x_+ \in \partial\Omega$, $x_+ \neq x_-$. Let $F_- = F_\Omega(x_-)$.*

For any sequence $\{x_n\} \subset L$, with $x_n \to x_-$, up to a subsequence, there exists a compact set $\mathcal{K} \subset \mathcal{C}(\mathbb{R}^d)$, a subspace $H_0 \subset \mathbb{R}^d$, and a sequence of $\mathcal{K}$-pseudo-loxodromic elements g_n in $\mathrm{PGL}(d, \mathbb{R})$*, with endpoints* $\mathrm{supp}(F_-)$ *and x_+ and neutral subspace H_0, such that $g_n x_n = x$ for a fixed $x \in L \cap \Omega$.*

Proof. Our strategy is to start with the case that F_- is codimension-one (so the neutral subspace H_0 is trivial), and then use Proposition 3.3.4 to extend to the general case.

F_- is codimension-one

Let V_- be the support of F_-. For each n, we let $s_n \in \mathrm{GL}(d, \mathbb{R})$ be the diagonal map

$$\lambda_n \mathrm{id}_{x_+} \oplus \mathrm{id}_{V_-} = \begin{bmatrix} \lambda_n & \\ & \mathrm{id}_{V_-} \end{bmatrix}$$

acting on $x_+ \oplus V_-$, where $\lambda_n \to \infty$ is chosen so that $s_n x_n = x$ for a fixed $x \in L \cap \Omega$.

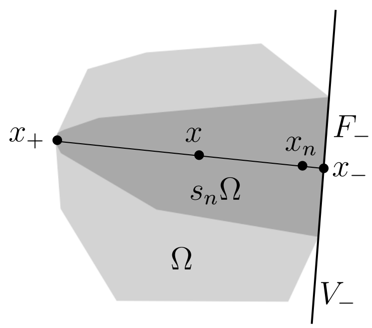

Figure 3.6: Since s_n attracts towards x_+ and repels from V_-, $s_n\Omega$ converges to the convex hull of F_- and x_+.

The sequence of domains $s_n \cdot \Omega_n$ converges to a cone over F_-, with a cone point at x_+ (see Figure 3.6). Since F_- is a codimension-one face of Ω, this cone is a properly convex domain containing x in its interior.

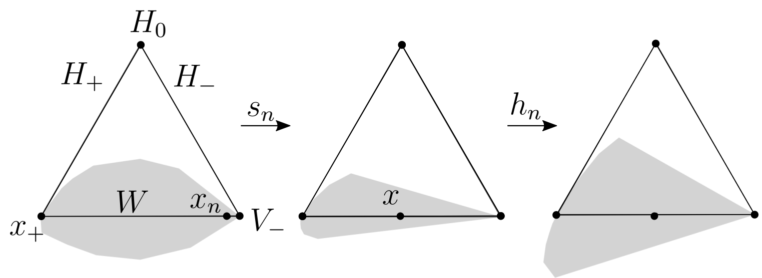

Figure 3.7: To build the sequence of pseudo-loxodromic elements g_n, we push x_n away from x_- with $s_n \in \mathrm{GL}(W)$, ensuring that $s_n\Omega \cap \mathbb{P}(W)$ and $\pi_W(s_n\Omega)$ converge, and then use a "correcting" element $h_n \in \mathrm{GL}(H_0)$ to keep the domain from degenerating. Both s_n and h_n preserve the decomposition $\mathbb{R}^d = x_+ \oplus H_0 \oplus V_-$.

The general case

Let V_- be the support of F_-, and let $\mathbb{P}(H_+)$, $\mathbb{P}(H_-)$ be supporting hyperplanes of Ω at x_+, F_-. Let $W = V_- \oplus x_+$. We choose a projective subspace $\mathbb{P}(H_0) \subset \mathbb{P}(H_+) \cap \mathbb{P}(H_-)$ as in Lemma 3.4.6 so that $H_- = V_- \oplus H_0$ and $\pi_W(\Omega)$ is properly convex, where $\pi_W : V \to W$ is the projection with kernel H_0.

The domains

$$\Omega \cap \mathbb{P}(W), \quad \pi_W(\Omega)$$

are both properly convex open subsets of $\mathbb{P}(W)$ containing F_- as a codimension-one face in their boundaries. Using the argument from the previous case, we can find group elements $s_n \in \mathrm{GL}(W)$ so that

$$s_n \cdot (\Omega \cap P(W)), \quad s_n \cdot \pi_W(\Omega)$$

both converge to properly convex domains in $\mathbb{P}(W)$ containing a fixed $x = s_n x_n$ in Ω.

We extend s_n linearly to the map $s_n \oplus \mathrm{id}_{H_0}$ on $W \oplus H_0$. Consider the sequence of properly convex domains

$$\Omega_n = (s_n \oplus \mathrm{id}_{H_0}) \cdot \Omega.$$

Since $s_n \oplus \mathrm{id}_{H_0}$ commutes with projection to W and intersection with W, the sequences of pointed properly convex domains

$$(\Omega_n \cap \mathbb{P}(W), x), \quad (\pi_W(\Omega_n), x)$$

both converge in $\mathcal{C}_*(W)$. In particular, both of these sequences are contained in a fixed compact $\mathcal{K}_W \subset \mathcal{C}_*(W)$, and the pointed domains (Ω_n, x) all lie in the subset

$$\mathcal{C}_*(W, H_0, \mathcal{K}_W)$$

from Definition 3.3.3.

Then, Proposition 3.3.4 (applied to the decomposition $\mathbb{R}^d = W \oplus H_0$) tells us that there is a sequence of group elements $h_n \in \mathrm{GL}(H_0)$ such that the pointed properly convex domains

$$(\mathrm{id}_W \oplus h_n) \cdot (\Omega_n, x)$$

lie in a fixed compact $\mathcal{K}$ in $\mathcal{C}_*(\mathbb{R}^d)$.

Then, we can take our sequence of $\mathcal{K}$-pseudo-loxodromic elements g_n to be the projectivizations of $(\mathrm{id}_W \oplus h_n) \cdot (s_n \oplus \mathrm{id}_{H_0}) = (s_n \oplus h_n)$. $\square$

Next we examine some of the dynamical behavior of pseudo-loxodromic sequences that have a repelling endpoint. Let V be a normed vector space. For any $g \in \mathrm{GL}(V)$, recall that the *norm* and *conorm* of g on V are defined by

$$||g|| = \sup_{v \in V - \{0\}} \frac{||gv||}{||v||}, \quad \mathbf{m}(g) = \inf_{v \in V - \{0\}} \frac{||gv||}{||v||}.$$

Proposition 3.4.8. *Let g_n be a sequence of $\mathcal{K}$-pseudo-loxodromic elements with common endpoints V_+, V_- and common neutral subspace V_0, and suppose that V_- is a repelling endpoint (Definition 3.4.5) of the sequence g_n. Let $E_+ = V_+ \oplus V_0$. The sequence g_n satisfies*

$$\frac{\mathbf{m}(g_n|_{E_+})}{||g_n|_{V_-}||} \to \infty. \tag{3.7}$$

The ratio (3.7) can be computed by fixing a norm on $\mathbb{R}^d$, and then choosing a lift of each g_n in $\mathrm{GL}(d, \mathbb{R})$. The value of (3.7) does not depend on the choice of lift, and the asymptotic behavior does not depend on the choice of norm.

Proof. We can fix lifts $\tilde{g}_n$ of g_n in $\mathrm{GL}(d, \mathbb{R})$ which restrict to the identity on V_-. Our goal is then to show that

$$\mathbf{m}(\tilde{g}_n|_{E_+}) \to \infty,$$

or equivalently, that

$$||\tilde{g}_n^{-1}|_{E_+}|| \to 0.$$

Suppose otherwise, so that for a sequence $v_n \in E_+$ with $||v_n|| = 1$, there is some $\varepsilon > 0$ so that

$$||\tilde{g}_n^{-1} \cdot v_n|| \geq \varepsilon.$$

Let $x_n \in \Omega \cap \mathbb{P}(V_+ \oplus V_-)$ be a sequence so that $g_n x_n = x$ for some $x \in \Omega$ and $x_n \to x_-$, where V_- is the support of $F_\Omega(x_-)$. We can choose a subsequence so that $g_n\Omega$ converges to some properly convex domain Ω_∞. Ω_∞ contains x by Proposition 3.4.4, so let U be an open neighborhood of x whose closure is contained in Ω_∞. We can find a lift $\tilde{x}$ of x in $\mathbb{R}^d$ so that the projectivizations of each vector

$$\tilde{x} \pm v_n$$

lie in U, and thus in Ω_n for all sufficiently large n. Since $\tilde{g}_n$ restricts to the identity on V_-, the sequence $\tilde{g}_n^{-1}\tilde{x}$ converges to a lift $\tilde{x}_-$ of x_-.

Then, up to a subsequence, the sequence of pairs of vectors

$$\tilde{g}_n^{-1} \cdot (\tilde{x} \pm v_n)$$

lies in a lift $\tilde{\Omega}$ of Ω, and converges in $\mathbb{R}^d$ to $\tilde{x}_- \pm v_\infty$, where $v_\infty \in E_+$ has norm at least ε. This pair of points spans a nontrivial projective line segment in $\overline{\Omega}$ whose interior intersects the face $F_\Omega(x_-)$ only at x_-, contradicting the definition of $F_\Omega(x_-)$. $\square$

Proposition 3.4.8 implies in particular that a sequence of $\mathcal{K}$-pseudo-loxodromic elements with repelling subspace V_- attracts generic points in $\mathbb{RP}^{d-1}$ to the projective subspace $\mathbb{P}(E_+)$, and repels points away from $\mathbb{P}(V_-)$. It also implies that the sequence g_n has *expansion* behavior on the Grassmannian in a neighborhood of V_-:

Corollary 3.4.9. *Let g_n be a sequence of $\mathcal{K}$-pseudo-loxodromic elements with common endpoints V_+, V_- and common neutral subspace V_0, and suppose that V_- is a repelling endpoint of the sequence, lying in* $\mathrm{Gr}(k,d)$.

Then for any Riemannian metric d_k on $\mathrm{Gr}(k,d)$*, and any $E > 1$, there exists $N \in \mathbb{N}$ such that if $n \geq N$, g_n is E-expanding (with respect to d_k) on some neighborhood of V_- in* $\mathrm{Gr}(k,d)$.

Proof. This follows from Proposition 3.4.8 via a computation in an appropriate metric on the Grassmannian, which is explicitly carried out in e.g. [BPS19, Lemma A.10]. □

3.4.3 Expansion

Before we proceed, we fix some additional terminology:

Definition 3.4.10. Given a properly convex domain Ω and a point $x \in \partial\Omega$, we say that a sequence $x_n \in \Omega$ *limits to x along a line L* if $x_n \to x$ in $\mathbb{RP}^{d-1}$, L is an open projective line segment $(x, x') \subseteq \Omega$, and there exists a constant $R > 0$ such that

$$d_\Omega(x_n, L) < R$$

for all n.

If the specific line L is implied (or not relevant), we will just say that x_n limits to x along a line.

If F is some open face of $\partial\Omega$, we say that x_n *limits to F along a line L* if every subsequence of x_n has a subsequence limiting to some $x \in F$ along L.

Remark 3.4.11. If Γ is a group acting on a properly convex domain Ω, and there are $\gamma_n \in \Gamma$ so that $\gamma_n x_0$ limits to x along a line for some $x_0 \in \Omega$, the point x is often referred to as a *conical limit point* for the action of Γ on $\partial\Omega$. We will avoid this terminology, since we will need to discuss conical limit points later in a way that is not exactly equivalent.

Proposition 3.4.12. *Let Ω be a properly convex domain and let $\Gamma \subseteq \mathrm{Aut}(\Omega)$. Let F_- be an open face of $\partial\Omega$, and let x_n be a sequence in Ω limiting to F_- along a line.*

If there exists $\gamma_n \in \Gamma$ so that $\gamma_n x_n$ is relatively compact in Ω, then:

(a) There exists a compact set $K \subseteq \mathrm{PGL}(d, \mathbb{R})$ such that $\gamma_n = k_n g_n$, where $k_n \in K$ and $g_n \in \mathrm{PGL}(d, \mathbb{R})$ is a sequence of $\mathcal{K}$-pseudo-loxodromics with repelling endpoint $\mathrm{supp}(F_-)$.

(b) For any Riemannian metric d_k on $\mathrm{Gr}(k, d)$, and any $E > 1$, for all sufficiently large n there is a neighborhood U of $\mathrm{supp}(F_-)$ in $\mathrm{Gr}(k, d)$ such that γ_n is E-expanding (with respect to d_k) on U.

Proof. Fix a compact $C \subset \Omega$ so that $\gamma_n x_n \in C$ for all n. We can move each x_n by a bounded Hilbert distance so that it lies on a fixed line segment L with an endpoint on F_-. So, by enlarging C if necessary, we can assume that the points x_n actually lie on the line L.

Let $\mathcal{K}' \subset \mathcal{C}_*(\mathbb{R}^d)$ be the compact set $\{\Omega\} \times C$. By assumption we know that for all n, we have

$$(\Omega, \gamma_n x_n) \in \mathcal{K}'.$$

Using Lemma 3.4.7, we can find a compact subset $\mathcal{K} \subset \mathcal{C}(\mathbb{R}^d)$ and a sequence g_n of $\mathcal{K}$-pseudo-loxodromic elements with repelling endpoint $\mathrm{supp}(F_-)$ taking x_n to x, for some $x \in \Omega \cap L$. The g_n can be chosen so that the axis contains L, implying that the set

$$\mathcal{K} \times \{x\} \subset \mathcal{C}_*(\mathbb{R}^d)$$

is compact.

Each group element $k_n = \gamma_n g_n^{-1}$ takes a pointed domain in the compact set $\mathcal{K} \times \{x\}$ to a pointed domain in the compact set $\mathcal{K}'$. But then, because $\mathrm{PGL}(d, \mathbb{R})$ acts properly on $\mathcal{C}_*(\mathbb{R}^d)$, the k_n lie in a fixed compact subset of $\mathrm{PGL}(d, \mathbb{R})$. This proves part (a).

Let V_- be the support of F_-, and let $k = \dim V_-$. The elements k_n can be viewed as lying in a compact subset of the diffeomorphisms of the compact manifold $\mathrm{Gr}(k, d)$. So, for any fixed Riemannian metric d on $\mathrm{Gr}(k, d)$, there is a constant $M > 0$ so that for all n and all $W_1, W_2 \in \mathrm{Gr}(k, d)$,

$$d_k(k_n W_1, k_n W_2) > M \cdot d_k(W_1, W_2).$$

Fix $E > 1$. Since g_n has repelling endpoint V_-, Corollary 3.4.9 implies that for some sufficiently large n, there is a neighborhood U of V_- in $\mathrm{Gr}(k, d)$ so that g_n satisfies

$$d_k(g_n W_1, g_n W_2) > \frac{E}{M} \cdot d_k(W_1, W_2)$$

for all $W_1, W_2 \in U$. But then we have

$$d_k(\gamma_n W_1, \gamma_n W_2) > E \cdot d_k(W_1, W_2)$$

giving us the required expansion. □

Proof of Proposition 3.4.1. Let Γ act cocompactly on some convex $C \subset \Omega$. Fix a Riemannian metric on $\mathrm{Gr}(k, d)$ and a constant $E > 1$.

For every face F of $\partial_i C$, there is a sequence x_n in C limiting to F along a line. Then part (b) of Proposition 3.4.12 implies that if $\gamma_n x_n$ is relatively compact in C for $\gamma_n \in \Gamma$, γ_n is E-expanding on a neighborhood of $\mathrm{supp}(F)$ for sufficiently large n. □

Proof of (1) $\implies$ (2) *in Theorem 3.1.9.* We apply Proposition 3.4.1 to $\mathrm{Hull}_\Omega(\Gamma)$, whose ideal boundary is the full orbital limit set $\Lambda_\Omega(\Gamma)$. Lemma 3.1.8 implies that $\Lambda_\Omega(\Gamma)$ contains all of its faces and is closed and boundary-convex, so it is the Γ-invariant subset required by the theorem. □

3.4.4 North-south dynamics

In Section 4.3, it will be useful to apply a consequence of part (a) of Proposition 3.4.12. The following can be thought of as a kind of weak version of north-south dynamics on the limit set of a group acting on a convex projective domain.

Proposition 3.4.13. *Let Ω be a properly convex domain, let $\Gamma \subset \mathrm{Aut}(\Omega)$, and let Λ be a closed Γ-invariant subset of $\partial\Omega$. Let F be an open face of Λ, and let x_n be a sequence limiting to F along a line.*

For any sequence γ_n such that $\gamma_n x_n$ is relatively compact in Ω, there exist subspaces E_+ and E_-, with $E_+ \oplus E_- = \mathbb{R}^d$, so that:

1. *$\mathbb{P}(E_+)$, $\mathbb{P}(E_-)$ are supporting subspaces of Ω, intersecting Λ,*
2. *for every compact $K \subset \partial\Omega - \overline{F}$, a subsequence of $\gamma_n K$ converges uniformly to a subset of $\mathbb{P}(E_+)$, and a subsequence of $\gamma_n F$ converges to a subset of $\mathbb{P}(E_-)$,*
3. *for every $x \in F$ and every $z \in \partial\Omega - \overline{F}$, the sequence of line segments*

$$\gamma_n \cdot [x, z]$$

subconverges to a line segment intersecting Ω.

Proof. Using Proposition 3.4.12, we decompose each γ_n as $k_n g_n$, for a sequence g_n of $\mathcal{K}$-pseudo-loxodromic elements with repelling endpoint $V_- = \mathrm{supp}(F)$, and k_n lying in a fixed compact in $\mathrm{PGL}(d, \mathbb{R})$. Taking a subsequence, we may assume that k_n converges to $k \in \mathrm{PGL}(d, \mathbb{R})$, so that

$$\gamma_n V_- = k_n g_n V_- = k_n V_- \to k V_-.$$

Let $E_- = kV_-$. We let V_+ be the other endpoint of the sequence g_n, let V_0 be the neutral subspace, and let $E_+ := k(V_+ \oplus V_0)$. Since Λ is closed and Γ-invariant, both $\mathbb{P}(E_+)$ and $\mathbb{P}(E_-)$ intersect Λ.

Fix a compact subset K in $\partial\Omega - F$. Proposition 3.4.8 implies that $g_n K$ converges uniformly to a subset of $\mathbb{P}(V_+ \oplus V_0)$. So, $k_n g_n K$ converges uniformly to a subset of $\mathbb{P}(E_+)$.

This shows parts (1) and (2). To see part (3), let L be the line segment $[x, z]$. By Proposition 2.1.11, we can find $R > 0$ and $x_n' \in L$ such that

$$d_\Omega(x_n, x_n') \leq R.$$

We know that $\gamma_n x_n$ lies in a fixed compact subset C of Ω. So, $\gamma_n x_n'$ lies in a closed and bounded Hilbert neighborhood of C. This is also a compact subset of Ω, so up to a subsequence, $\gamma_n x_n'$ converges to some $x_0' \in \Omega$.

The limit of the line segment $[\gamma_n x_-, \gamma_n x_n']$ is nontrivial, intersects Ω, and is a subsegment of the limit of $[\gamma_n x_-, \gamma_n z]$, so this implies the desired result. $\square$

3.4.5 Cartan projections of sequences in Γ

The decomposition given by part (a) of Proposition 3.4.12 also allows us to describe the behavior of the *Cartan projection* of certain sequences in a group Γ which acts convex cocompactly on a properly convex domain.

First we briefly review the definition of the Cartan projection. Recall that $\mathrm{GL}(d, \mathbb{R})$ has a *Cartan decomposition*

$$\mathrm{GL}(d, \mathbb{R}) = K \exp(\mathfrak{a}_+)K,$$

where $K = \mathrm{O}(d)$ and $\mathfrak{a}_+ \subset \mathfrak{gl}(d, \mathbb{R})$ is the set of diagonal matrices with nonincreasing entries. That is, each $g \in \mathrm{GL}(d, \mathbb{R})$ can be uniquely written

$$g = k \cdot \exp(\mathrm{diag}(\mu_1(g), \ldots, \mu_d(g))) \cdot k',$$

for $k, k' \in K$ and $\mu_1(g) \geq \mu_2(g) \geq \ldots \geq \mu_d(g)$.

$\mu_i(g)$ is the logarithm of the ith singular value of g. The map $\mathrm{GL}(d, \mathbb{R}) \to \mathbb{R}^d$ given by

$$g \mapsto \mu(g) = (\mu_1(g), \ldots, \mu_d(g))$$

is the Cartan projection.

While the map μ is not defined on $\mathrm{PGL}(d, \mathbb{R})$, the *gaps*

$$\mu_i(g) - \mu_{i+k}(g)$$

still make sense for any $g \in \mathrm{PGL}(d, \mathbb{R})$. A sequence $g_n \in \mathrm{PGL}(d, \mathbb{R})$ is divergent (i.e., leaves every compact set in $\mathrm{PGL}(d, \mathbb{R})$) if and only if the gaps

$$\mu_1(g_n) - \mu_d(g_n)$$

tend to infinity.

Proposition 3.4.14. *Let Ω be a properly convex domain, let Γ act convex cocompactly on Ω, and let F be an open face of $\Lambda_\Omega(\Gamma)$, $\dim(F) = k-1$.*

Fix a basepoint $x_0 \in \Omega$. For any sequence of group elements $\gamma_n \in \mathrm{Aut}(\Omega)$ such that $\gamma_n x_0$ limits to F along a line, the Cartan projections μ_i satisfy:

1. $$\mu_k(\gamma_n) - \mu_{k+1}(\gamma_n) \to \infty.$$

2. *For a constant D independent of n, we have*
$$\mu_1(\gamma_n) - \mu_k(\gamma_n) < D.$$

Proof. Using part (a) of Proposition 3.4.12, we can write

$$\gamma_n^{-1} = k_n g_n,$$

where g_n is a sequence of $\mathcal{K}$-pseudo-loxodromics with repelling endpoint $V_- = \mathrm{supp}(F)$.

The singular values of the sequence γ_n depend on a choice of inner product on $\mathbb{R}^d$, but changing the inner product only changes the singular values by a bounded amount. So, we may assume that the endpoints V_+, V_- and neutral subspace V_0 of the sequence g_n are orthogonal to each other.

Proposition 3.4.8 then implies that the smallest k singular values of g_n are identically the eigenvalue of g_n on V_-, and that g_n has an unbounded singular value gap at index $d-k$.

The Cartan projection $\mu : \mathrm{GL}(d, \mathbb{R}) \to \mathbb{R}^d$ satisfies the inequality

$$||\mu(gh) - \mu(g)|| \leq ||\mu(h)||$$

for all $g, h \in \mathrm{GL}(d, \mathbb{R})$. So, since the group elements k_n^{-1} lie in a fixed compact subset of $\mathrm{PGL}(d, \mathbb{R})$, and $\gamma_n = g_n^{-1} k_n^{-1}$, we can find lifts $\tilde{\gamma}_n$, $\tilde{g}_n$ in $\mathrm{GL}(d, \mathbb{R})$ so that the differences

$$||\mu(\tilde{\gamma}_n) - \mu(\tilde{g}_n^{-1})||$$

are bounded. So γ_n has an unbounded singular value gap at index k and bounded singular value gaps at indices $1, \ldots, k-1$. $\square$

Chapter 4

Projective convex cocompactness and relative hyperbolicity

4.1 Results in this chapter

In this chapter, we will use the dynamical properties of projective convex cocompactness studied in Chapter 3 to prove the projectively convex cocompact representations are characterized by the existence of a certain equivariant *Bowditch boundary embedding* into a *quotient* of the boundary of some projective domain. Material from this chapter previously appeared as part of the arXiv preprint "Dynamical properties of convex cocompact groups in projective space" [Wei20].

Definition 4.1.1. Let $\mathcal{H} = \{H_i\}$ be a collection of subgroups of $\mathrm{PGL}(d, \mathbb{R})$, each acting convex cocompactly on a fixed properly convex domain Ω with pairwise disjoint full orbital limit sets $\Lambda_\Omega(H_i)$.

We denote by $[\partial\Omega]_{\mathcal{H}}$ the space obtained from $\partial\Omega$ by collapsing all of the full orbital sets $\Lambda_\Omega(H_i)$ to points. Similarly, for $x \in \partial\Omega$, or a subset $\Lambda \subseteq \partial\Omega$, we use $[x]_{\mathcal{H}}$ and $[\Lambda]_{\mathcal{H}}$ to denote the images of x and Λ in $[\partial\Omega]_{\mathcal{H}}$.

When $\mathcal{H}$ is a conjugacy-invariant collection of subgroups of a group $\Gamma \subseteq \mathrm{Aut}(\Omega)$, the action of Γ on $\partial\Omega$ descends to an action on $[\partial\Omega]_{\mathcal{H}}$. More generally, if $\Lambda \subseteq \partial\Omega$ is Γ-invarant, Γ also acts on $[\Lambda]_{\mathcal{H}}$.

We show the following:

Theorem 4.1.2. *Let $\Gamma \subseteq \mathrm{PGL}(d, \mathbb{R})$ act on a properly convex domain Ω, and suppose that Γ is hyperbolic relative to a family of subgroups $\mathcal{H} = \{H_i\}$, such that the H_i each act convex cocompactly on Ω with pairwise disjoint full orbital limit sets.*

If there is a boundary-convex Γ-invariant subset $\Lambda \subseteq \partial\Omega$ containing all of its faces, and a Γ-equivariant embedding $\partial(\Gamma, \mathcal{H}) \to [\partial\Omega]_{\mathcal{H}}$ with image $[\Lambda]_{\mathcal{H}}$, then Γ acts convex cocompactly on Ω and Λ is the full orbital limit set $\Lambda_\Omega(\Gamma)$.

Remark 4.1.3. In Theorem 4.1.2, we do not need to assume that Γ is discrete in $\mathrm{PGL}(d, \mathbb{R})$: this will also follow from the existence of the equivariant boundary embedding.

There are two special cases of Theorem 4.1.2 worth considering, which we state separately as corollaries. The first is when the subset Λ is the entire boundary $\partial\Omega$.

Corollary 4.1.4. *Let Γ, Ω, and $\mathcal{H}$ be as in Theorem 4.1.2, and suppose that $\partial(\Gamma, \mathcal{H})$ is equivariantly homeomorphic to $[\partial\Omega]_{\mathcal{H}}$. Then Γ divides Ω.*

The second corollary is when the set of peripheral subgroups is empty, i.e. Γ is hyperbolic.

Corollary 4.1.5. *Let Γ be a word-hyperbolic group in $\mathrm{PGL}(d, \mathbb{R})$ acting on a properly convex domain Ω, and suppose that the Gromov boundary of Γ embeds equivariantly into $\partial\Omega$ with image Λ.*

If Λ is boundary-convex and contains all of its faces, then Γ acts convex cocompactly on Ω and $\Lambda = \Lambda_\Omega(\Gamma)$.

When a hyperbolic group acts convex cocompactly on a domain Ω, its full orbital limit set contains no segments. So in this case, Λ contains all of its faces whenever no point of Λ lies in the interior of any segment in $\partial\Omega$.

We also can phrase this corollary in terms of Anosov boundary maps. Recall that Theorem 3.1.3 (proved as [DGK17, Theorem 1.15]; see also [Zim21, Theorem 1.10]) states that if a word-hyperbolic group Γ acts convex cocompactly on some domain Ω, then the inclusion map $\Gamma \hookrightarrow \mathrm{PGL}(d, \mathbb{R})$ is a P_1-Anosov representation preserving Ω, and in this case the full orbital limit set is the image of the Anosov boundary map $\partial\Gamma \to \mathbb{RP}^{d-1}$. Thus Corollary 4.1.5 implies:

Corollary 4.1.6. *Let Γ be a word-hyperbolic subgroup of $\mathrm{PGL}(d, \mathbb{R})$ preserving a properly convex domain Ω, and suppose that there exists a Γ-equivariant embedding $\xi : \partial\Gamma \to \partial\Omega$ whose image is boundary-convex and contains all of its faces. Then the inclusion $\Gamma \hookrightarrow \mathrm{PGL}(d, \mathbb{R})$ is a P_1-Anosov representation with $\mathbb{RP}^{d-1}$ boundary map ξ.*

Note that it is not true in general that the $\mathbb{RP}^{d-1}$-boundary map ξ of a P_1-Anosov representation always embeds into the boundary of some properly convex domain $\Omega \subset \mathbb{RP}^{d-1}$. Moreover even if ξ does embed into $\partial\Omega$ for some Ω, it does not necessarily follow that the image of the embedding is boundary-convex. However, as mentioned in Remark 3.1.4, if $\rho : \Gamma \to G$ is *any* Anosov representation, there is always a representation $\tau : G \to \mathrm{PGL}(n, \mathbb{R})$ for some large n so that the composition $\tau \circ \rho$ is

P_1-Anosov, and the P_1 limit map embeds $\partial\Gamma$ into the boundary of *some* projective domain $\Omega \subseteq \mathbb{R}\mathrm{P}^{n-1}$ as a boundary-convex closed subset containing its faces.

During the proof of Theorem 4.1.2, we will see the following (see Proposition 4.2.1):

Proposition 4.1.7. *In the setting of Theorem 4.1.2, every nontrivial segment in the set Λ is contained in the full orbital limit set of some $H_i \in \mathcal{H}$.*

This leads us to a converse to Theorem 4.1.2.

Theorem 4.1.8. *Let Γ be a group acting convex cocompactly on a properly convex domain Ω, and suppose that Γ has a conjugacy-invariant collection of subgroups $\mathcal{H} = \{H_i\}$, such that the groups in $\mathcal{H}$ lie in finitely many conjugacy classes and each H_i acts convex cocompactly on Ω.*

Then Γ is hyperbolic relative to $\mathcal{H}$ if and only if

(i) the full orbital limit sets $\Lambda_\Omega(H_i)$, $\Lambda_\Omega(H_j)$ are disjoint for distinct $H_i, H_j \in \mathcal{H}$,

(ii) every nontrivial segment in $\Lambda_\Omega(\Gamma)$ is contained in $\Lambda_\Omega(H_i)$ for some $H_i \in \mathcal{H}$, and

(iii) each $H_i \in \mathcal{H}$ is its own normalizer in Γ.

Moreover, in this case, $\partial(\Gamma, \mathcal{H})$ equivariantly embeds into $[\partial\Omega]_{\mathcal{H}}$ with image $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$.

Remark 4.1.9. If conditions (i) and (ii) hold for a conjugacy-closed collection of subgroups $\mathcal{H}$ of Γ, then they also hold for the collection of normalizers, since $g \cdot \Lambda_\Omega(H_i) = \Lambda_\Omega(gH_ig^{-1})$ for any $H_i \in \mathcal{H}$.

Moreover, condition (iii) of Theorem 4.1.8 is always true for the peripheral subgroups of a relatively hyperbolic group, because then each $H_i \in \mathcal{H}$ can be exactly realized as the stabilizer of its unique fixed point in the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ (see Theorem 2.2.14).

Work of Islam-Zimmer [IZ19a] [IZ19b] implies that if Γ is hyperbolic relative to $\mathcal{H}$, all of the groups in $\mathcal{H}$ are virtually abelian, and ρ is a projectively convex cocompact representation, then in fact ρ *automatically* restricts to a projectively convex cocompact representation on each group $H \in \mathcal{H}$. In fact, since our results first appeared, Islam-Zimmer have also shown that this is true even without the assumption that each $H \in \mathcal{H}$ is virtually abelian. This implies the following:

Corollary 4.1.10 (See [IZ22], Theorem 1.6)**.** *Let Ω be a properly convex domain, and let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair. If Γ acts convex cocompactly on Ω, then there is an equivariant embedding from $\partial(\Gamma, \mathcal{H})$ to $[\partial\Omega]_{\mathcal{H}}$ whose image is $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$.*

We remark that Islam-Zimmer have also subsequently shown that versions of Theorem 4.1.8 and Corollary 4.1.10 hold in the more general context of *naive* convex cocompact groups.

4.1.1 Outline of the chapter

In Section 4.2, we use Yaman's dynamical characterization of relative hyperbolicity (Theorem 2.2.14), together with some of the results of the previous chapter, to prove Theorem 4.1.2.

In Section 4.3, we extend a result of Islam-Zimmer [IZ19b], showing that when Γ is a group acting on a properly convex domain Ω, and Γ is hyperbolic relative to a collection of subgroups acting convex cocompactly on Ω, conditions (i) and (ii) of Theorem 4.1.8 hold. Then, we prove the rest of Theorem 4.1.8, again using Yaman's dynamical definition of relative hyperbolicity.

4.2 Bowditch boundary embedding implies convex cocompactness

Our goal in this section is to prove Theorem 4.1.2. Our first step is the following:

Proposition 4.2.1. *Let Ω be a properly convex domain, and let $\Gamma \subset \mathrm{Aut}(\Omega)$ be hyperbolic relative to a collection of subgroups $\mathcal{H} = \{H_i\}$ each acting convex cocompactly on Ω with disjoint full orbital limit sets $\Lambda_\Omega(H_i)$.*

Suppose Λ is a Γ-invariant subset of $\partial\Omega$ containing all of its faces and containing $\Lambda_\Omega(H_i)$ for every H_i. If $[\Lambda]_{\mathcal{H}}$ is the image of a Γ-equivariant embedding $\phi : \partial(\Gamma, \mathcal{H}) \to [\partial\Omega]_{\mathcal{H}}$, then the set

$$\Lambda_c = \Lambda - \bigcup_{H_i \in \mathcal{H}} \Lambda_\Omega(H_i)$$

contains only extreme points in $\partial\Omega$.

Proof. The equivariant homeomorphism $\phi : \partial(\Gamma, \mathcal{H}) \to [\Lambda]_{\mathcal{H}}$ means that Γ acts on $[\Lambda]_{\mathcal{H}}$ as a convergence group as in Theorem 2.2.14. In particular, we can classify the points of $[\Lambda]_{\mathcal{H}}$ as either bounded parabolic points or conical limit points, where the parabolic points are exactly the points corresponding to $\Lambda_\Omega(H_i)$.

So, if x is a point in Λ_c, it represents a conical limit point in $[\Lambda]_{\mathcal{H}}$. Suppose for a contradiction that x is not an extreme point, i.e. x lies in the interior of a nontrivial segment $[a, b] \subset \partial\Omega$. Since Λ contains all of its faces, $(a, b) \subset \Lambda$, and we can find $w, z \in \Lambda$ such that w, x, z are pairwise distinct points lying on (a, b) in that order.

Lemma 3.1.8 tells us that each $\Lambda_\Omega(H_i)$ contains its faces, so we know that w and z cannot lie in any $\Lambda_\Omega(H_i)$. So w, x, and z represent three distinct points in $[\Lambda]_{\mathcal{H}}$.

This means that there exist group elements $\gamma_n \in \Gamma$ so that $\gamma_n[x]_{\mathcal{H}} \to a$, and $\gamma_n[z]_{\mathcal{H}}$, $\gamma_n[w]_{\mathcal{H}}$ both converge to some $b \in [\Lambda]_{\mathcal{H}}$, with a, b distinct.

This convergence is only in $[\Lambda]_{\mathcal{H}}$. However, since the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ is always compact, $[\Lambda]_{\mathcal{H}}$ is as well, and therefore its preimage Λ in the compact set $\partial\Omega$ is compact too. So, up to a subsequence, we can assume that $\gamma_n x \to u$, and $\gamma_n z \to v_1$, $\gamma_n w \to v_2$, with

$$[u]_{\mathcal{H}} = a, \; [v_1]_{\mathcal{H}} = [v_2]_{\mathcal{H}} = b.$$

The line segment $[w, z]$ must converge to the line segment $[v_1, v_2]$, which must contain u. If $v_1 = v_2$, this is clearly impossible without having $u = v_1 = v_2$. If $v_1 \neq v_2$, then v_1, v_2 both lie in $\Lambda_\Omega(H_i)$ for some H_i. Since each $\Lambda_\Omega(H_i)$ is boundary-convex (Lemma 3.1.8 again), u must lie in $\Lambda_\Omega(H_i)$ as well, a contradiction. □

The above is important partly because of the following proposition, which we will use repeatedly in the proof of both Theorem 4.1.2 and its converse.

Proposition 4.2.2. *Let Ω be a properly convex domain, and let Λ be a boundary-convex subset of $\partial\Omega$ containing all of its faces. Let $\mathcal{H}$ be a collection of subgroups of* $\mathrm{Aut}(\Omega)$ *acting convex cocompactly with disjoint full orbital limit sets in Ω.*

If every point in $\Lambda_c = \Lambda - \bigcup_{H_i \in \mathcal{H}} \Lambda_\Omega(H_i)$ is an extreme point, then for any $x, y \in \Lambda$ with $[x]_{\mathcal{H}} \neq [y]_{\mathcal{H}}$, the segment (x, y) lies in Ω.

Proof. Let x, y be distinct points in Λ. Boundary-convexity means that if the segment (x, y) is in $\partial\Omega$, it is also in Λ. Since we know Λ_c only contains extreme points, some $u \in (x, y)$ lies in $\Lambda_\Omega(H_i)$ for some $H_i \in \mathcal{H}$. Since H_i acts convex cocompactly on Ω, Lemma 3.1.8 implies that $[x, y]$ lies in $\Lambda_\Omega(H_i)$, which means that $[x]_{\mathcal{H}} = [y]_{\mathcal{H}}$. □

The following proposition explains why we do not need to assume that Γ is discrete in the statement of Theorem 4.1.2.

Proposition 4.2.3. *If Ω, Γ, Λ are as in Theorem 4.1.2, and Γ is non-elementary, then Γ is discrete.*

Proof. Γ acts as a convergence group on $[\Lambda]_{\mathcal{H}}$, so it acts properly discontinuously on the space of distinct triples in $[\Lambda]_{\mathcal{H}}$, which we denote $\mathcal{T}([\Lambda]_{\mathcal{H}})$.

The map

$$\Gamma \times \mathcal{T}([\Lambda]_{\mathcal{H}}) \to \mathcal{T}([\Lambda]_{\mathcal{H}})$$

given by the Γ-action is continuous, so Γ is discrete. □

We are now able to prove Theorem 4.1.2.

Proof of Theorem 4.1.2. Let Ω, Γ, Λ, $\mathcal{H}$ be as in the hypotheses for Theorem 4.1.2. We can assume that $\mathcal{H} \neq \{\Gamma\}$ and that Γ is infinite (if not then the theorem is trivial). This means that $\partial(\Gamma, \mathcal{H})$ contains at least two points, and Proposition 4.2.2 implies that $\mathrm{Hull}_\Omega(\Gamma)$ is nonempty.

If Γ is virtually infinite cyclic, the hypotheses of the theorem imply that the generator γ of a finite-index cyclic subgroup fixes a pair of points $\{x, y\}$ in $\partial\Omega$ with $(x, y) \subset \Omega$; γ acts as a translation in the Hilbert metric along the axis (x, y). This action is properly discontinuous (so Γ is discrete) and cocompact. Further, since x and y are extreme points, $\gamma^n z$ converges to either x or y as $n \to \pm\infty$ for all $z \in \Omega$, so $\Lambda_\Omega(\Gamma) = \{x, y\}$.

So we may assume Γ is non-elementary. Owing to Theorem 3.1.9, we only need to show that Γ is expanding at the faces of Λ; in fact we will show directly that the expansion is uniform.

Since each H_i acts convex cocompactly on Ω, Theorem 3.1.9 means that Γ is expanding in a neighborhood of the support of any face of $\Lambda_\Omega(H_i)$ for some H_i. In fact, we can assume that the expansion constants are uniform over all $H_i \in \mathcal{H}$ (see Remark 3.2.2), so we only need to consider the faces in

$$\Lambda_c = \Lambda - \bigcup_{H_i \in \mathcal{H}} \Lambda_\Omega(H_i).$$

Proposition 4.2.1 implies that each of these faces is actually just a point in $\partial\Omega$, whose support is equal to itself.

Let x be a point in Λ_c. We will build a sequence of points x_n in Ω limiting to x along a line (Definition 3.4.10), and show that the orbits $\Gamma \cdot x_n$ intersect a fixed compact set.

Since Γ is non-elementary, its Bowditch boundary contains at least three distinct conical limit points, so we can find $y, z \in \Lambda_c$ so that $[x]_{\mathcal{H}}, [y]_{\mathcal{H}}, [z]_{\mathcal{H}}$ are pairwise distinct.

Fix supporting hyperplanes W, V of Ω at x and z, respectively. Proposition 4.2.2 implies that $W \cap V$ does not contain x, y, or z, and that the line segment (x, z) is in Ω. The projective hyperplane spanned by $W \cap V$ and y intersects (x, z) at a point $w \in \Omega$.

Since $[x]_{\mathcal{H}}$ is a conical limit point, we can find a sequence $\gamma_n \in \Gamma$ so that

$$\gamma_n[x]_{\mathcal{H}} \to a$$

and

$$\gamma_n[z]_{\mathcal{H}},\ \gamma_n[y]_{\mathcal{H}} \to b$$

for a, b distinct. As in the proof of Proposition 4.2.1, we can pick subsequences so that $\gamma_n x$, $\gamma_n y$, and $\gamma_n z$ all converge to points $x_\infty, y_\infty, z_\infty$ in Λ, and $\gamma_n W$ and $\gamma_n V$ converge to supporting hyperplanes W_∞, V_∞ of Ω at x_∞ and z_∞.

Since x_∞ and z_∞ represent distinct points of $[\Lambda]_{\mathcal{H}}$, Proposition 4.2.2 implies that $W_\infty \cap V_\infty$ must not contain x_∞ or z_∞; for the same reason y_∞ is also not contained in $W_\infty \cap V_\infty$.

While $[z_\infty]_{\mathcal{H}} = [y_\infty]_{\mathcal{H}}$, it is not necessarily true that $y_\infty = z_\infty$. However, we do know that the segment (y_∞, x_∞) cannot lie in $\partial\Omega$. So, the sequence

$$\gamma_n \cdot (H \cap (x, z)) = \gamma_n w$$

cannot approach x_∞.

Propositon 4.2.3 means that we know Γ is discrete, and so its action on Ω is properly discontinuous. Thus $\gamma_n w$ must accumulate to an endpoint of $[x_\infty, z_\infty]$—and therefore to z_∞.

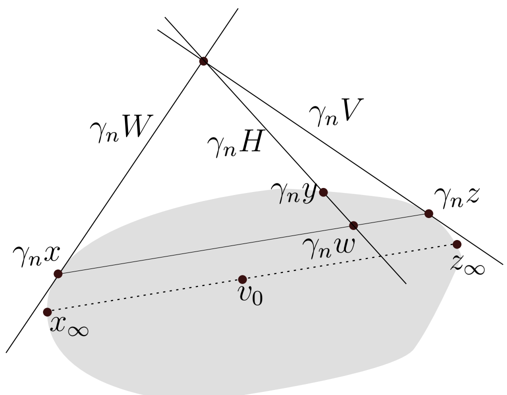

Figure 4.1: The sequence $\gamma_n w$ limits to z_∞, so the sequence $\gamma_n^{-1} v_0$ limits to x along a line.

Let ℓ be the line segment $[x, z]$. This segment has a well-defined total order, where $a < b$ if a is closer to x than b. If $\ell_n = [\gamma_n x, \gamma_n z]$, then γ_n is an order-preserving isometry from ℓ to ℓ_n, where the metric is the restricted Hilbert metric d_Ω.

Fix a basepoint v_0 on the line segment $\ell_\infty = [x_\infty, z_\infty]$, and choose $v_n \in \ell_n$ converging to v_0. Since $\gamma_n w$ converges to z_∞, we see that $v_n < \gamma_n w$ and

$$d_\Omega(v_n, \gamma_n w) \to \infty.$$

Thus we must have $\gamma_n^{-1} v_n \to x$.

But now we can apply part (b) of Proposition 3.4.12 to the sequence $\gamma_n^{-1} v_n \subset \ell$ to see that γ_n is eventually expanding in a neighborhood of x in $\mathbb{RP}^{d-1}$. $\square$

4.3 Convex cocompact groups which are relatively hyperbolic

The goal of this section is to prove Theorem 4.1.8.

4.3.1 Non-peripheral segments in the boundary

We start by showing that conditions (i) and (ii) of Theorem 4.1.8 are satisfied whenever Γ is a convex cocompact group hyperbolic relative to a collection of convex cocompact subgroups. That is, we will show:

Proposition 4.3.1. *Let Γ be a group hyperbolic relative to a collection $\mathcal{H}$ of subgroups, and suppose that Γ and each $H_i \in \mathcal{H}$ act on a properly convex domain Ω convex cocompactly.*

Then:

(i) The full orbital limit sets $\Lambda_\Omega(H_i)$ are disjoint for distinct $H_i, H_j \in \mathcal{H}$.

(ii) Every nontrivial segment in $\Lambda_\Omega(\Gamma)$ is contained in the full orbital limit set of some peripheral subgroup H_i.

We will closely follow the proof of a similar result of Islam and Zimmer [IZ19b, Theorem 1.8 (7)]. The main idea is that a nontrivial segment ℓ in the full orbital limit set $\Lambda_\Omega(\Gamma)$ of a convex cocompact group Γ is accumulated to by segments in the boundary of some maximal properly embedded simplices in $\mathrm{Hull}_\Omega(\Gamma)$. When Γ is hyperbolic relative to a collection $\mathcal{A}$ of virtually abelian subgroups of rank ≥ 2, Islam and Zimmer show that $\mathcal{A}$ is in one-to-one correspondence with the set of maximal properly embedded simplices in $\mathrm{Hull}_\Omega(\Gamma)$, and then use a coset separation property due to Druţu and Sapir [DS05] to see that these maximal properly embedded simplices are isolated. This ends up implying that ℓ lies in the boundary of one of the simplices that accumulate to it.

When we do *not* assume the peripheral subgroups are virtually abelian, we need to modify this approach slightly. First, we need to assume that the peripheral subgroups act convex cocompactly on Ω ([IZ19b] implies that this assumption is always satisfied in the virtually abelian case). Second, in our situation, the maximal properly embedded simplices in $\mathrm{Hull}_\Omega(\Gamma)$ do not need to be isolated. However, it is true that the convex cores $\mathrm{Hull}_\Omega(H_i)$ of the peripheral subgroups in $\mathcal{H}$ are isolated. So the desired result ends up following from the fact that every maximal properly embedded k-simplex ($k \geq 2$) in $\mathrm{Hull}_\Omega(\Gamma)$ lies in $\mathrm{Hull}_\Omega(H_i)$ for some $H_i \in \mathcal{H}$; this is Lemma 4.3.3 below.

4.3.1.1 Cosets and convex cores of peripheral subgroups

Let Γ be hyperbolic relative to a collection of subgroups $\mathcal{H}$, and suppose that Γ and each $H_i \in \mathcal{H}$ act convex cocompactly on a fixed properly convex domain Ω. We fix a basepoint $x \in \Omega$, and fix a finite set $\mathcal{P}$ of conjugacy representatives for $\mathcal{H}$.

The Švarc-Milnor lemma implies that Γ is finitely generated and that, under the word metric induced by any finite generating set, Γ is equivariantly quasi-isometric to the convex core $\mathrm{Hull}_\Omega(\Gamma)$ equipped with the restricted Hilbert metric d_Ω. The quasi-isometry can be taken to be the orbit map $\gamma \mapsto \gamma \cdot x$.

Since each P_i also acts convex cocompactly on Ω, each P_i is also finitely generated, and P_i is quasi-isometric to $\mathrm{Hull}_\Omega(P_i)$, which isometrically embeds into $\mathrm{Hull}_\Omega(\Gamma)$. We may assume that the quasi-isometry constants are uniform over all $P_i \in \mathcal{P}$, and fix a finite generating set for Γ containing generating sets for each P_i.

Since $g \cdot \mathrm{Hull}_\Omega(P_i) = \mathrm{Hull}_\Omega(gP_ig^{-1})$, if we fix a Γ-equivariant quasi-isometry

$$\phi : \mathrm{Hull}_\Omega(\Gamma) \to \Gamma,$$

we know ϕ restricts to a quasi-isometry

$$\mathrm{Hull}_\Omega(gP_ig^{-1}) \to gP_i,$$

with uniform quasi-isometry constants over all $g \in \Gamma$, $P_i \in \mathcal{P}$.

The cosets gP_i have a *separation property*: distinct cosets cannot stay "close" to each other over sets of large diameter. The precise statement is as follows. For any metric space X, and any $A \subseteq X$, recall that $N_X(A, r)$ denotes the open r-neighborhood of A in X with respect to the metric d_X, and $B_X(x, r)$ denotes the open r-ball about $x \in X$. When $X = \Omega$ for a properly convex domain Ω, we assume that Ω is equipped with the Hilbert metric.

Theorem 4.3.2 ([DS05, Theorem 4.1 (α_1)]). *Let Γ be hyperbolic relative to $\mathcal{H}$, and let $\mathcal{P}$ be a finite set of conjugacy representatives. For every $r > 0$, there exists $R > 0$ such that for every distinct pair of left cosets g_1P_1, g_2P_2, the diameter of the set*

$$N_\Gamma(g_1P_1, r) \cap N_\Gamma(g_2P_2, r)$$

is at most R.

In addition, Theorem 1.7 of [DS05] implies that if $k \geq 2$, any quasi-isometrically embedded k-flat in a relatively hyperbolic group Γ is contained in the D-neighborhood of a coset gP_i of some peripheral subgroup $P_i \in \mathcal{P}$. This allows us to see the following:

Lemma 4.3.3. *Suppose Γ acts convex cocompactly on Ω, and that Γ is hyperbolic relative to a collection of subgroups $\mathcal{H}$ also acting convex cocompactly on Ω. Every properly embedded k-simplex ($k \geq 2$) in Ω with boundary in $\Lambda_\Omega(\Gamma)$ is contained in* $\mathrm{Hull}_\Omega(H_i)$ *for some $H_i \in \mathcal{H}$.*

Proof. Each such embedded k-simplex Δ is a quasi-isometrically embedded k-flat in $\mathrm{Hull}_\Omega(\Gamma)$, so $\phi(\Delta)$ is a quasi-isometrically embedded k-flat in Γ. [DS05], Theorem 1.7 implies that $\phi(\Delta)$ is contained in a uniform neighborhood of gP for some $P \in \mathcal{P}$.

Applying a quasi-inverse of ϕ tells us that Δ is in a uniform Hilbert neighborhood of $\mathrm{Hull}_\Omega(gPg^{-1})$ in Ω. So the boundary of Δ is contained in $\partial_i \mathrm{Hull}_\Omega(gPg^{-1})$, and Δ itself lies in $\mathrm{Hull}_\Omega(gPg^{-1})$. $\square$

We now quote:

Lemma 4.3.4 ([IZ19b, Lemma 15.4])**.** *Let (u, v) be a nontrivial line segment in $\Lambda_\Omega(\Gamma)$, let $m \in (u, v)$ and $p \in \mathrm{Hull}_\Omega(\Gamma)$, and let V be the span of (u, v) and p. For any $r > 0$, $\varepsilon > 0$, there exists a neighborhood U of m in $\mathbb{P}(V)$ such that if $x \in U \cap \mathrm{Hull}_\Omega(\Gamma)$, then there is a properly embedded simplex $S_x \subset \mathrm{Hull}_\Omega(\Gamma)$ such that*

$$B_\Omega(x, r) \cap \mathbb{P}(V) \subset N_\Omega(S_x, \varepsilon).$$

Now we can prove Proposition 4.3.1. The proof of part (ii) is nearly identical to the proof of Lemma 15.5 in [IZ19b].

Proof of Proposition 4.3.1. (i). Let H_i, H_j be a pair of peripheral subgroups in $\mathcal{H}$, and suppose that $\Lambda_\Omega(H_i) \cap \Lambda_\Omega(H_j)$ contains a point $x \in \partial\Omega$. We can find a pair of projective-line geodesic rays in $\mathrm{Hull}_\Omega(H_i)$ and $\mathrm{Hull}_\Omega(H_j)$ with one endpoint at x. Proposition 2.1.11 implies that the images of these rays have finite Hausdorff distance.

Thus, in Γ, a uniform neighborhood of the coset g_iP_i corresponding to H_i contains an infinite-diameter subset of the coset g_jP_j corresponding to H_j. So Theorem 4.3.2 implies that $H_i = H_j$.

(ii). Consider any nontrivial segment $[u, v]$ in $\Lambda_\Omega(\Gamma)$, and fix $m \in (u, v)$ and $p \in \mathrm{Hull}_\Omega(\Gamma)$. Theorem 4.3.2 implies that for some $R > 0$, there exists $r > 0$ such that the diameter of

$$N_\Omega(\mathrm{Hull}_\Omega(H_i), r) \cap N_\Omega(\mathrm{Hull}_\Omega(H_j), r)$$

is less than R whenever H_i and H_j are distinct.

Let V be the span of u, v, and p. Lemma 4.3.4 implies that for some neighborhood U of m in $\mathbb{P}(V)$, for every $x \in U$, there is some properly embedded simplex S_x such that

$$B_\Omega(x, R) \cap \mathbb{P}(V) \subset N_\Omega(S_x, r).$$

Lemma 4.3.3 means that the simplex S_x is contained in the convex hull $\mathrm{Hull}_\Omega(H_x)$ of some peripheral subgroup H_x, and part (i) implies that this peripheral subgroup is unique.

We can shrink U so that it is convex, and claim that in this case $H_x = H_y$ for all $x, y \in U \cap \mathrm{Hull}_\Omega(\Gamma)$. By convexity, it suffices to show this when $d_\Omega(x, y) \leq R/2$.

Then

$$B_\Omega(x, R/2) \cap \mathbb{P}(V) \subset B_\Omega(y, R) \cap \mathbb{P}(V) \subset N_\Omega(S_y, r)$$

so the diameter of

$$N_\Omega(\mathrm{Hull}_\Omega(H_x), r) \cap N_\Omega(\mathrm{Hull}_\Omega(H_y), r)$$

is at least the diameter of $B_\Omega(x, R/2) = R$. Thus $H_x = H_y$.

Fix $H = H_x$ for some $x \in U \cap \mathrm{Hull}_\Omega(\Gamma)$. Then, if x_n is a sequence in $\mathrm{Hull}_\Omega(\Gamma)$ approaching m, there is a sequence $x_n' \in \mathrm{Hull}_\Omega(H)$ such that

$$d_\Omega(x_n, x_n') \leq k,$$

for k independent of n. Up to a subsequence, x_n' converges to some $x' \in \Lambda_\Omega(H)$. Proposition 2.1.11 implies that

$$F_\Omega(x') = F_\Omega(m) \supseteq (u, v).$$

$\Lambda_\Omega(H)$ contains x'. It is also closed and contains all of its faces (Lemma 3.1.8). So $[u, v] \subset \Lambda_\Omega(H)$.

□

4.3.2 Convex cocompact and no relative segment implies relatively hyperbolic

We now turn to the rest of Theorem 4.1.8. As in our proof of Theorem 4.1.2, the main tool will be Yaman's dynamical characterization of relative hyperbolicity (Theorem 2.2.14). If Γ is virtually cyclic, Yaman's theorem does not apply, but in this case Γ is hyperbolic and the result follows from [DGK17].

Throughout the rest of this section, we assume (as in the hypotheses to Theorem 4.1.8) that Ω is a properly convex domain preserved by a discrete non-elementary group Γ acting convex cocompactly with full orbital limit set $\Lambda_\Omega(\Gamma)$, and $\mathcal{H}$ is a conjugacy-invariant set of subgroups of Γ lying in finitely many conjugacy classes, with each $H_i \in \mathcal{H}$ acting convex cocompactly on Ω. We also assume $\mathcal{H} \neq \{\Gamma\}$, since the result is trivial in this case.

We will prove the following:

Proposition 4.3.5. *Suppose that conditions (i), (ii), and (iii) of Theorem 4.1.8 hold for the collection of subgroups $\mathcal{H}$. Then:*

1. *Γ acts as a convergence group on $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$,*
2. *$[\Lambda_\Omega(\Gamma)]_\mathcal{H}$ is compact, metrizable, and perfect,*
3. *the groups H_i are parabolic subgroups, and their fixed points are bounded parabolic,*
4. *every point in*

   $$[\Lambda_\Omega(\Gamma)]_\mathcal{H} - \{[\Lambda_\Omega(H_i)]_\mathcal{H} : H_i \in \mathcal{H}\}$$

   is a conical limit point for the Γ-action on $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$.

Since convex cocompact groups are always finitely generated, Theorem 4.1.8 is a direct consequence of Proposition 4.3.5, Proposition 4.3.1, and Theorem 2.2.14.

4.3.2.1 Dynamics of the Γ-action on $\Lambda_\Omega(\Gamma)$

We start by proving part (1) of Proposition 4.3.5, using the theory of *attracting* and *repelling* subspaces of sequences in Γ discussed at the end of Section 2.1.

Proposition 4.3.6. Γ *acts as a convergence group on* $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$.

Proof. Let γ_n be an infinite sequence in Γ. Since Γ is discrete, γ_n is divergent, and there is a subsequence of γ_n and projective subspaces E_+, E_- so that for any compact $K \subset \mathbb{RP}^{d-1} - E_-$, $\gamma_n \cdot K$ converges to a subset of E_+.

Lemma 2.1.19 implies that E_+ and E_- are both supporting subspaces of Ω and both intersect $\Lambda_\Omega(\Gamma)$ nontrivially. The intersections $E_+ \cap \Lambda_\Omega(\Gamma)$ and $E_- \cap \Lambda_\Omega(\Gamma)$ are respectively the closures of subsets of a pair of faces $F_+, F_- \subset \Lambda_\Omega(\Gamma)$. By assumption, every face in $\Lambda_\Omega(\Gamma)$ containing a nontrivial projective segment lies in some $\Lambda_\Omega(H_i)$, so each face in $\Lambda_\Omega(\Gamma)$ represents a single point of $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$. So we have

$$[E_- \cap \Lambda_\Omega(\Gamma)]_{\mathcal{H}} = a, \quad [E_+ \cap \Lambda_\Omega(\Gamma)]_{\mathcal{H}} = b$$

for (not necessarily distinct) points $a, b \in [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$.

Let $[K]_{\mathcal{H}}$ be a compact subset of $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{a\}$, where K is the preimage of $[K]_{\mathcal{H}}$ in $\Lambda_\Omega(\Gamma)$. K is compact since $\Lambda_\Omega(\Gamma)$ is compact. Moreover, K cannot intersect E_-. So, $\gamma_n \cdot K$ converges to a subset of $E_+ \cap \Lambda_\Omega(\Gamma)$, and $\gamma_n \cdot [K]_{\mathcal{H}}$ converges to b. $\square$

4.3.2.2 Topological properties of $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$

Next, we will check that $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ satisfies each of the properties in part (2) of Proposition 4.3.5. The first, compactness, is immediate from the compactness of $\Lambda_\Omega(\Gamma)$.

Showing that $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ is metrizable is equivalent to showing that it is Hausdorff, since it is a quotient of a compact metrizable space.

Let

$$\pi_{\mathcal{H}} : \Lambda_\Omega(\Gamma) \to [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$$

be the quotient map. We will show that if a is a point in $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$, then we can find arbitrarily small open neighborhoods of $\pi_{\mathcal{H}}^{-1}(a)$ in $\Lambda_\Omega(\Gamma)$ which are of the form $\pi_{\mathcal{H}}^{-1}(U)$ for $U \subset [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$.

Our first step is the following:

Lemma 4.3.7. *Fix any metric $d_{\mathbb{P}}$ on projective space. Let $a \in [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$.*

For any $\varepsilon > 0$, there exists a subset $W(a, \varepsilon) \subset \Lambda_\Omega(\Gamma)$ satisfying:

1. *$W(a, \varepsilon) = \pi_{\mathcal{H}}^{-1}(V)$ for some $V \subset [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$,*
2. *$W(a, \varepsilon)$ contains an open neighborhood of $\pi_{\mathcal{H}}^{-1}(a)$ in $\Lambda_\Omega(\Gamma)$, and*
3. *For every $z \in W(a, \varepsilon)$, we have*

$$d_{\mathbb{P}}(z, \pi_{\mathcal{H}}^{-1}(a)) < \varepsilon.$$

Proof. Let $X_a = \pi_{\mathcal{H}}^{-1}(a)$. For any open set U in $\Lambda_\Omega(\Gamma)$ containing X_a, we let $W(U)$ be the set

$$\pi_{\mathcal{H}}^{-1}([U]_{\mathcal{H}}) = U \cup \{x \in \Lambda_\Omega(H_i) : \Lambda_\Omega(H_i) \cap U \neq \emptyset\}.$$

$W(U)$ is a subset of $\Lambda_\Omega(\Gamma)$ satisfying conditions (1) and (2). We claim that for any given $\varepsilon > 0$, $W(U)$ also satisfies condition (3) as long as U is sufficiently small.

We proceed by contradiction. Suppose otherwise, so that there is some $\varepsilon > 0$ so that for a shrinking sequence of open neighborhoods U_n of X_a, there is some $H_n \in \mathcal{H}$ such that

$$\Lambda_\Omega(H_n) \cap U_n \neq \emptyset,$$

and $\Lambda_\Omega(H_n)$ contains a point z_n such that $d_\mathbb{P}(z_n, X_a) \geq \varepsilon$.

We write $\Lambda_n = \Lambda_\Omega(H_n)$. We can choose a subsequence so that in the topology on nonempty closed subsets of projective space, Λ_n converges to some closed subset of $\Lambda_\Omega(\Gamma)$, which we denote Λ_∞, and z_n converges to $z_\infty \in \Lambda_\infty$ such that $d_\mathbb{P}(z_\infty, X_a) \geq \varepsilon$.

Λ_∞ intersects every open subset of $\Lambda_\Omega(\Gamma)$ containing X_a, and since X_a is a closed subset of a metrizable space, this means Λ_∞ intersects X_a. We will get a contradiction by showing that in fact $z_\infty \in X_a$.

We consider two cases:

Case 1: $\mathrm{Hull}_\Omega(\Lambda_\infty)$ *is nonempty.* Since the groups in $\mathcal{H}$ lie in only finitely many conjugacy classes, up to a subsequence, the H_n are all conjugate to each other, and we may assume that $\Lambda_n = \gamma_n \Lambda_0$ for a sequence $\gamma_n \in \Gamma$.

We can find a sequence $x_n \in \mathrm{Hull}_\Omega(\Lambda_n)$ converging to some $x_\infty \in \mathrm{Hull}_\Omega(\Lambda_\infty)$. Since the action of H_0 on $\mathrm{Hull}_\Omega(\Lambda_0)$ is cocompact, there is some fixed $R > 0$ so that every H_0-orbit in $\mathrm{Hull}_\Omega(\Lambda_0)$ intersects the Hilbert ball of radius R about x_0. Since H_n is a conjugate of H_0 by an isometry of the Hilbert metric on Ω, the same is true (with the same R) for every x_n, H_n, and Λ_n.

So, we can find a sequence

$$\mu_n \in \gamma_n H_n \gamma_n^{-1}$$

so that $\mu_n\gamma_n \cdot x_0$ lies in the Hilbert ball of radius R about x_n. Since x_n converges to $x_\infty \in \Omega$, and Γ acts properly discontinuously on Ω, this means that a subsequence of $\mu_n\gamma_n$ is eventually constant. Because $\mu_n\gamma_n\Lambda_0 = \gamma_n\Lambda_0$, this means we can assume there is some fixed $\gamma \in \Gamma$ so that

$$\Lambda_\infty = \gamma\Lambda_0 = \Lambda_\Omega(\gamma H_0 \gamma^{-1}).$$

But then since the limit sets $\Lambda_\Omega(H_i)$ are disjoint, we must have $X_a = \Lambda_\infty$, which means $z_\infty \in X_a$.

Case 2: $\mathrm{Hull}_\Omega(\Lambda_\infty)$ *is empty.* In this case, Λ_∞ lies in the closure $\overline{F}$ of some face F of $\partial\Omega$; without loss of generality we can assume that F intersects Λ_∞. $\Lambda_\Omega(\Gamma)$ contains its faces, so it contains all of F.

If F is a single point, we must have $\Lambda_\infty = \{z_\infty\}$, so z_∞ lies in X_a. If F is *not* a single point, it contains a nontrivial segment. By assumption, this segment lies in $\Lambda_\Omega(H_i)$ for some H_i; since $\Lambda_\Omega(H_i)$ is closed and contains its faces, all of $\overline{F}$ lies in $\Lambda_\Omega(H_i)$ as well. But then $\Lambda_\Omega(H_i)$ intersects both X_a and z_∞, so $X_a = \Lambda_\Omega(H_i) = [z_\infty]_\mathcal{H}$ and $z_\infty \in X_a$ in this case as well.

$\square$

Proposition 4.3.8. $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$ *is Hausdorff.*

Proof. Let a, a' be distinct points in $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$, and let X_a, X_a' be the preimages of a and a' in $\Lambda_\Omega(\Gamma)$.

Since X_a and X'_a are closed disjoint subsets of the metrizable space $\Lambda_\Omega(\Gamma)$, there is some $\varepsilon > 0$ such that for any $x \in X_a$, $x' \in X'_a$,

$$d(x, x') > 2\varepsilon.$$

For each $n \in \mathbb{N}$, we define a sequence of sets U_n containing X_a as follows. We let $U_0 = X_a$. Then, for each $n > 0$, we take U_n to be the set

$$\bigcup_{b \in [U_{n-1}]_{\mathcal{H}}} W(b, \varepsilon/2^n),$$

where $W(b, \varepsilon/2^n)$ is the set given by Lemma 4.3.7. Note that each U_n is a set of the form $\pi_{\mathcal{H}}^{-1}(V)$ for some $V \subset [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$; moreover, if $z \in U_n$, then

$$d(z, U_{n-1}) < \varepsilon/2^n.$$

Consider the set $U = \bigcup_{n \in \mathbb{N}} U_n$. U is the preimage of some $V \subset [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$, and it must be contained in an ε-neighborhood of X_a. In addition, U is open in $\Lambda_\Omega(\Gamma)$: if z is in U_n, then U_{n+1} contains $W([z]_{\mathcal{H}}, \varepsilon/2^{n+1})$, which in turn contains an open neighborhood of z.

This means that $[U]_{\mathcal{H}}$ is an open set in $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ containing a. We can similarly construct an open set $[U']_{\mathcal{H}}$ containing a' such that U' is contained in an ε-neighborhood of X'_a. U and U'_∞ are disjoint, so $[U]_{\mathcal{H}}$ and $[U']_{\mathcal{H}}$ separate a and a'. $\square$

Next we show that the space $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ is perfect, i.e. $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ contains no isolated points.

Proposition 4.3.9. $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$ *is perfect.*

Proof. Fix $a \in [\Lambda_\Omega(\Gamma)]_\mathcal{H}$ and a representative x of a. Let $F = F_\Omega(x)$.

Let x_n be a sequence of points in $\mathrm{Hull}_\Omega(\Gamma)$ converging to x in $\mathbb{RP}^{d-1}$. Convex cocompactness means that for some $\gamma_n \in \Gamma$, $\gamma_n^{-1}x_n \in C$ for a fixed compact $C \subset \Omega$.

This means that (up to a subsequence) for fixed $x_0 \in \Omega$, $\gamma_n x_0$ converges to a point in F. And since γ_n acts by Hilbert isometries, Proposition 2.1.11 implies that if B is any open ball with finite Hilbert radius about x_0, $\gamma_n B$ converges uniformly to a subset of F.

γ_n is divergent in $\mathrm{PGL}(d, \mathbb{R})$, so let E_+ and E_- be a pair of attracting and repelling projective subspaces for the sequence γ_n. We know that E_+ and E_- are supporting subspaces of Ω, and that

$$[E_- \cap \Lambda_\Omega(\Gamma)]_\mathcal{H}, \quad [E_+ \cap \Lambda_\Omega(\Gamma)]_\mathcal{H}$$

are single points in $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$. Moreover, since an open subset of projective space converges under γ_n to F, E_+ intersects F, and $[E_+ \cap \Lambda_\Omega(\Gamma)]_\mathcal{H} = a$. Let $b = [E_- \cap \Lambda_\Omega(\Gamma)]_\mathcal{H}$.

Since we assume $\mathcal{H} \neq \{\Gamma\}$, $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$ cannot be a single point, and since Γ is non-elementary, $[\Lambda_\Omega(\Gamma)]_\mathcal{H}$ contains at least three points. So, we can find a pair of points $c_1, c_2 \in [\Lambda_\Omega(\Gamma)]_\mathcal{H}$ such that $\{b, c_1, c_2\}$ are pairwise distinct. Both c_1 and c_2 have a representative which does not lie in E_-, so both $\gamma_n \cdot c_1$ and $\gamma_n \cdot c_2$ converge to a; since $c_1 \neq c_2$, a cannot be isolated. $\square$

4.3.2.3 Parabolic points in $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$

Our next task is to verify part (3) of Proposition 4.3.5—that is, to show that points stabilized by our candidate peripheral subgroups are bounded parabolic points.

Proposition 4.3.10. *Each point $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$ in $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ is a parabolic point for the action of Γ, with stabilizer H_i.*

Proof. The fact that H_i is self-normalizing implies that H_i is exactly the stabilizer of $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$ in Γ: for general $g \in \mathrm{Aut}(\Omega)$,

$$g \cdot \Lambda_\Omega(H_i) = \Lambda_\Omega(gH_ig^{-1}),$$

and since we assume that the full orbital limit sets of distinct groups in $\mathcal{H}$ are disjoint, $g \in \Gamma$ preserves $\Lambda_\Omega(H_i)$ if and only if g normalizes H_i.

So we just need to check that the groups H_i are parabolic. Let $\gamma \in H_i$ be an infinite-order element, so that γ^n is a divergent sequence in $\mathrm{PGL}(d, \mathbb{R})$. We want to show that γ does not fix any point in $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ other than $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$.

Let E_+ and E_- be attracting and repelling subspaces for the sequence γ^n. Lemma 2.1.19 implies that both E_+ and E_- support Ω and intersect $\Lambda_\Omega(H_i)$ nontrivially.

Let $b \in [\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[\Lambda_\Omega(H_i)]_{\mathcal{H}}\}$, let $y \in \Lambda_\Omega(H_i) \cap E_-$, and let $x \in \Lambda_\Omega(\Gamma)$ be a representative of b. Proposition 4.2.2 implies that x cannot lie in E_-, so $\gamma^n x$ converges to a point in $\Lambda_\Omega(\Gamma) \cap E_+$. Then $\gamma^n b$ converges to $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$, and in particular γ does not fix b. □

We still need to show that the parabolic points $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$ are *bounded* parabolic points, i.e. that H_i acts cocompactly on

$$[\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[\Lambda_\Omega(H_i)]_{\mathcal{H}}\}.$$

Our strategy is to show that the set

$$\Lambda_i = \Lambda_\Omega(\Gamma) - \Lambda_\Omega(H_i)$$

is a closed subset of the interior of some convex open set Ω_{H_i}, such that the ideal boundary of Λ_i in Ω_{H_i} is exactly $\Lambda_\Omega(H_i)$. Then, we can use the fact that H_i is uniformly expanding in supports at $\Lambda_\Omega(H_i)$ to see that the action of H_i on Λ_i is cocompact.

If H_i is irreducible (or more generally, if we know that H_i contains a *proximal element*), then as a consequence of [Ben00, Proposition 3.1] (or [DGK17, Proposition 4.5]), we can simply take Ω_{H_i} to be the unique H_i-invariant *maximal* properly convex domain $\Omega_{\max}$ in $\mathbb{R}\mathrm{P}^{d-1}$. Since we do not know if H_i contains a proximal element in general, we do not know if such a maximal domain exists. So, we will construct Ω_{H_i} directly.

To do so, we consider the *dual* full orbital limit set $\Lambda_{\Omega^*}(\Gamma)$ of a group Γ acting on a properly convex domain Ω. i.e. the full orbital limit set in Ω^* of Γ viewed as a subgroup of $\mathrm{Aut}(\Omega^*)$. Each element $w \in \Lambda_{\Omega^*}(\Gamma)$ is an element of $\partial\Omega^*$, so $\mathbb{P}(w)$ is a supporting hyperplane of Ω.

Proposition 4.3.11. *Let* Γ *be any subgroup of* $\mathrm{Aut}(\Omega)$.

1. *For every* $x \in \Lambda_\Omega(\Gamma)$ *there exists* $w \in \Lambda_{\Omega^*}(\Gamma)$ *such that* $w(x) = 0$.

2. For every $w \in \Lambda_{\Omega^*}(\Gamma)$ there exists $x \in \Lambda_\Omega(\Gamma)$ such that $w(x) = 0$.

The statement follows from e.g. Proposition 5.6 in [IZ19a]; we provide an alternative proof for convenience.

Proof. The two statements are dual to each other, so we only need to prove (1).

Given a point $x \in \Omega$, and $W \in \Omega^*$, we consider a quantity

$$\delta_\Omega(x, W)$$

defined in [DGK17] as follows:

$$\delta_\Omega(x, W) = \inf_{z \in \mathbb{P}(W)} \{\min\{|[a_z, x; b_z, z]|, |[b_z, x; a_z, z]|\},$$

where a_z and b_z are the points in $\partial\Omega$ such that a_z, x, b_z, z lie on a projective line. $\delta_\Omega(x, W)$ can be thought of as an $\mathrm{Aut}(\Omega)$-invariant measure of how "close" x is to $\partial\Omega$, relative to the projective hyperplane W: it takes on nonzero values for $x \in \Omega$, $W \in \Omega^*$, and for fixed $W \in \Omega^*$ and x_n converging to $\partial\Omega$, $\delta_\Omega(x_n, W)$ converges to 0.

We now take $z \in \Lambda_\Omega(\Gamma)$, and choose $\gamma_n \in \Gamma$, $z_0 \in \Omega$ so that $\gamma_n \cdot z_0 \to z$. Fix some $W_0 \in \Omega^*$, and consider the sequence $\gamma_n \cdot W_0$. Up to a subsequence, this converges to some $W \in \Lambda^*_\Omega(\Gamma)$.

Since $\delta_\Omega(x, W)$ is Γ-invariant, for any sequence

$$y_n \in \gamma_n \cdot \mathbb{P}(W_0),$$

both of the cross-ratios

$$[a_{y_n}, \gamma_n \cdot z_0; b_{y_n}, y_n], [b_{y_n}, \gamma_n \cdot z_0; a_{y_n}, y_n]$$

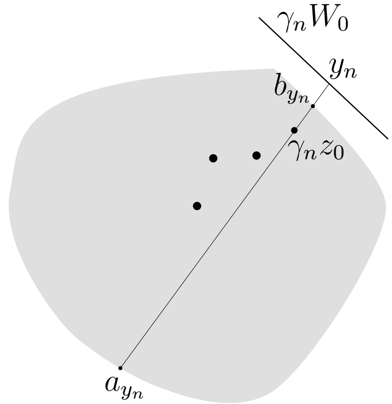

Figure 4.2: If $\gamma_n z_0$ approaches the boundary of Ω, and $\delta_\Omega(\gamma_n z_0, \gamma_n W_0)$ is bounded away from 0, $\gamma_n W_0$ must limit to a hyperplane containing the limit of $\gamma_n z_0$.

remain bounded away from 0 as $n \to \infty$. But since $\gamma_n \cdot z_0$ approaches $z \in \partial\Omega$, we can choose y_n so that exactly one of a_{y_n}, b_{y_n} also approaches z. Thus, y_n approaches z as well, and so $\mathbb{P}(W)$ contains z. $\square$

Next, we consider the *dual convex core* for the Γ-action on Ω.

Definition 4.3.12. Let $\Omega \subset \mathbb{R}\mathrm{P}^{d-1}$ be a properly convex domain, and let $\Gamma \subseteq \mathrm{Aut}(\Omega)$. The *dual convex hull* $\mathrm{Hull}^*_\Omega(\Gamma)$ is the convex set

$$[\mathrm{Hull}_{\Omega^*}(\Lambda_{\Omega^*}(\Gamma))]^*.$$

Equivalently, $\mathrm{Hull}^*_\Omega(\Gamma)$ is the unique connected component of

$$\mathbb{R}\mathrm{P}^{d-1} - \bigcup_{W \in \Lambda_{\Omega^*}(\Gamma)} W$$

which contains Ω.

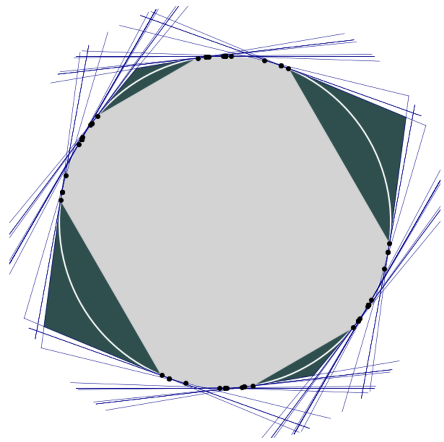

Figure 4.3: Part of the limit set and dual limit set for a group Γ acting convex compactly on the projective model for $\mathbb{H}^2$ (the interior of the white circle). $\mathrm{Hull}_\Omega(\Gamma)$ is the light region, and $\mathrm{Hull}^*_\Omega(\Gamma)$ is the dark region.

As long as $\Lambda_{\Omega^*}(\Gamma)$ contains at least two points, $\mathrm{Hull}^*_\Omega(\Gamma)$ does not contain all of $\mathbb{R}\mathrm{P}^{d-1}$. It can be viewed as an intersection of convex subspaces, so it is convex in the sense of Definition 2.1.3, but in general it is not properly convex.

We can use the dual convex core to finish proving part (3) of Proposition 4.3.5.

Proposition 4.3.13. *The stabilizer of* $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$ *acts cocompactly on*

$$\Lambda_i = [\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[\Lambda_\Omega(H_i)]_{\mathcal{H}}\}.$$

Proof. Let $\Omega_{H_i} = \mathrm{Hull}^*_\Omega(H_i)$ be the dual convex hull of H_i in Ω. Proposition 4.3.11 implies that $\Lambda_\Omega(H_i)$ lies in the boundary of Ω_{H_i}.

Moreover, the set $\Lambda_\Omega(\Gamma) - \Lambda_\Omega(H_i)$ lies in the interior of Ω_{H_i}—for, every point in the boundary of Ω_{H_i} is contained a projective hyperplane $\mathbb{P}(W)$ for $W \in \Lambda_{\Omega^*}(H_i)$, and each such hyperplane supports some $x \in \Lambda_\Omega(H_i)$. Since $\mathbb{P}(W)$ is also a supporting hyperplane of Ω, Proposition 4.2.2 implies that no $y \in \Lambda_\Omega(\Gamma) - \Lambda_\Omega(H_i)$ lies in $\mathbb{P}(W)$.

$\Lambda_\Omega(\Gamma)$ is thus a closed subset of Ω_{H_i} whose ideal boundary in Ω_{H_i} is contained in $\Lambda_\Omega(H_i)$. Since H_i acts convex cocompactly on Ω, it is uniformly expanding in supports at $\Lambda_\Omega(H_i)$ by Theorem 3.1.9. Then Proposition 3.2.5 (applied to the convex domain Ω_{H_i}) implies that the action of H_i on $\Lambda_\Omega(\Gamma) - \Lambda_\Omega(H_i)$ is cocompact—which means that the H_i-action on the quotient $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[\Lambda_\Omega(H_i)]_{\mathcal{H}}\}$ is cocompact as well. $\square$

4.3.2.4 Conical limit points in $[\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$

Finally we check part (4) of Proposition 4.3.5—that the remaining points in our candidate Bowditch boundary are indeed conical limit points. We will do this in two steps.

Lemma 4.3.14. *Let $H_i \in \mathcal{H}$, let*

$$x_n \in \Lambda_\Omega(\Gamma) - \Lambda_\Omega(H_i)$$

be a sequence approaching $x \in \Lambda_\Omega(H_i)$, and let $F = F_\Omega(x)$. If h_n is a sequence such that $h_n[x_n]_{\mathcal{H}}$ is relatively compact in

$$[\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[\Lambda_\Omega(H_i)]_{\mathcal{H}}\},$$

then for any compact

$$K \subset \partial\Omega - \overline{F},$$

h_n subconverges on $[K]_{\mathcal{H}}$ to the constant map $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$.

Proof. Any such sequence h_n must be divergent, so using Proposition 2.1.17, we can find nontrivial projective subspaces E_+, E_- so that if K is any compact subset of $\mathbb{RP}^{d-1} - E_-$, $h_n \cdot K$ subconverges uniformly to a subset of E_+. E_+ and E_- are supporting subspaces of Ω, and $E_+ \cap \partial\Omega$ and $E_- \cap \partial\Omega$ are both subsets of $\Lambda_\Omega(H_i)$.

E_- must contain x, since otherwise $h_n x_n$ would subconverge to a point in

$$E_+ \cap \partial\Omega \subseteq \Lambda_\Omega(H_i).$$

But then $E_- \cap \partial\Omega$ is a subset of $\overline{F}$ and the desired condition holds. □

Proposition 4.3.15. *Every element of the set*

$$[\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[\Lambda_\Omega(H_i)]_{\mathcal{H}} : H_i \in \mathcal{H}\}$$

is a conical limit point for the action of Γ *on* $\Lambda_\Omega(\Gamma)$.

Proof. By assumption, any point in this set has a unique representative $x \in \Lambda_\Omega(\Gamma)$ which is an extreme point in $\partial\Omega$. Fix a sequence $x_n \in \Omega$ limiting to x along a line, and let $\gamma_n \in \Gamma$ be group elements taking x_n back to some fixed compact in Ω.

Proposition 3.4.13 implies that there is a supporting subspace E_+ of Ω, intersecting $\Lambda_\Omega(\Gamma)$, so that $\gamma_n x$ limits to some $x' \in \Lambda_\Omega(\Gamma)$ not intersecting E_+, and if K is any compact subset of $\Lambda_\Omega(\Gamma) - x$, a subsequence of $\gamma_n K$ converges uniformly to a subset of $E_+ \cap \Lambda_\Omega(\Gamma)$. In particular, γ_n converges uniformly on compacts in

$$[\Lambda]_{\mathcal{H}} - \{[x]_{\mathcal{H}}\}$$

to the constant map $[E_+ \cap \Lambda_\Omega(\Gamma)]_{\mathcal{H}}$.

If $[x']_{\mathcal{H}} \neq [E_+ \cap \Lambda_\Omega(\Gamma)]_{\mathcal{H}}$, then we are done. However, it is also possible that x' and $E_+ \cap \Lambda_\Omega(\Gamma)$ both lie in the same full orbital limit set of some convex cocompact subgroup H_i.

In this case, we use the fact that $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$ is a bounded parabolic fixed point (Proposition 4.3.13) to find a sequence $h_n \in H_i$ such that $h_n \cdot [\gamma_n x]_{\mathcal{H}}$ lies in a fixed compact set C in

$$[\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[\Lambda_\Omega(H_i)]_{\mathcal{H}}\},$$

and consider the sequence of group elements $h_n\gamma_n$.

Fix a compact subset $[K]_{\mathcal{H}}$ of

$$[\Lambda_\Omega(\Gamma)]_{\mathcal{H}} - \{[x]_{\mathcal{H}}\},$$

where K is the (compact) preimage of $[K]_{\mathcal{H}}$ in $\Lambda_\Omega(\Gamma) - \{x\}$.

After taking a subsequence, $\gamma_n K$ must converge to a compact subset of $E_+ \cap \Lambda_\Omega(\Gamma)$, which does not intersect x'. In fact, part (3) of Proposition 3.4.13 implies that $\gamma_n K$ converges to a compact subset of $\Lambda_\Omega(\Gamma) - \overline{F'}$, where $F' = F_\Omega(x')$. So there is a fixed compact

$$K' \subset \Lambda_\Omega(\Gamma) - \overline{F'}$$

so that for sufficiently large n, $\gamma_n K \subset K'$. Then Lemma 4.3.14 implies that

$$h_n\gamma_n[K]_{\mathcal{H}} \subseteq h_n[K']_{\mathcal{H}}$$

subconverges to $[\Lambda_\Omega(H_i)]_{\mathcal{H}}$. But on the other hand,

$$[(h_n\gamma_n) \cdot x_n]_{\mathcal{H}} \in C$$

subconverges to some $b \neq [\Lambda_\Omega(H_i)]_{\mathcal{H}}$, so $h_n\gamma_n$ gives us the sequence of group elements we need. $\square$

Chapter 5

EGF representations

Material from this chapter and the next previously appeared in the arXiv preprint "An extended definition of Anosov representation for relatively hyperbolic groups" [Wei22].

In this chapter, we introduce another generalization of Anosov dynamics for non-hyperbolic groups: *extended geometrically finite* (EGF) representations. EGF representations generalize the *topological* dynamical behavior of Anosov representations, so their definition is based on the *convergence dynamics* definition of Anosov representations (Proposition 2.4.3).

Definition 5.0.1. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, with Γ acting on a connected compact metrizable space M by homeomorphisms. Let $\Lambda \subset M$ be a closed Γ-invariant set.

We say that an equivariant surjective map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ *extends the convergence dynamics of* Γ if for each $z \in \partial(\Gamma, \mathcal{H})$, there exists an open set $C_z \subset M$ containing $\Lambda - \phi^{-1}(z)$, satisfying the following:

If γ_n is a sequence in Γ with $\gamma_n^{\pm 1} \to z_\pm$ for $z_\pm \in \partial(\Gamma, \mathcal{H})$, then for any compact set $K \subset C_{z_-}$ and any open set U containing $\phi^{-1}(z_+)$, for sufficiently large n, $\gamma_n \cdot K$ lies in U.

EGF representations are defined with respect to a parabolic subgroup P of a semisimple Lie group G. For convenience, when working with EGF representations, we will always assume that P is a *symmetric* parabolic subgroup—i.e. P is conjugate to a subgroup P^- which is opposite to P.

When $P = P^+$ is symmetric, we can identify G/P^+ with G/P^-, so that it makes sense to say that two flags $\xi_1, \xi_2 \in G/P$ are *opposite*.

Definition 5.0.2. Let P be symmetric, and let A, B be two subsets of G/P. We say that A and B are *opposite* if every $\xi \in A$ is opposite to every $\nu \in B$.

Definition 5.0.3. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and let $\Lambda \subset G/P$. We say that a continuous surjective map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ is *antipodal* if for every pair of distinct points $z_1, z_2 \in \partial(\Gamma, \mathcal{H})$, $\phi^{-1}(z_1)$ is opposite to $\phi^{-1}(z_2)$.

The main definition of this chapter is:

Definition 1.5.6. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and let P be a symmetric parabolic subgroup of a semisimple Lie group G. We say that a representation $\rho : \Gamma \to G$ is *extended geometrically finite* with respect to P if there exists a closed $\rho(\Gamma)$-invariant subset $\Lambda \subset G/P$ and a continuous ρ-equivariant surjective antipodal map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ which *extends the convergence dynamics* of Γ.

Remark 5.0.4. Unfortunately, the boundary set $\Lambda \subset G/P$ is *not* necessarily uniquely determined by the representation ρ. In many contexts, we will be able to make a natural choice, but we do not give a procedure for doing so in general.

The main result of this chapter (Theorem 5.6.2) is that EGF representations satisfy a *relative stability* property, similar to the *stability* property for Anosov representations.

5.0.1 Outline of the chapter

In Section 5.1, we explain the precise relationship between EGF representations and the *relatively asymptotically embedded* representations defined by Kapovich-Leeb in [KL18]. In the same section, we also prove an *Anosov relativization theorem*, which says (roughly) that "EGF relative to Anosov is Anosov."

The results in Section 5.1 rely on some facts about EGF representations which are only proved later in the chapter. In Section 5.2 and Section 5.3, we introduce the main technical tools needed to prove these results: a *relative quasigeodesic automaton* associated to the action of any non-elementary relatively hyperbolic group on its Bowditch boundary. In Section 5.4, we relate relative quasigeodesic automata to the action of a relatively hyperbolic group on a flag manifold. We use these tools to develop an alternative characterization of EGF representations in Section 5.5, and finally prove our main theorem in Section 5.6.

5.0.2 Proof strategy

The general approach to our proof of Theorem 5.6.2 is loosely inspired by Sullivan's proof of stability for convex cocompact groups in rank-one [Sul85]. Sullivan shows that a discrete group $\Gamma \subset \mathrm{PO}(d,1)$ is convex cocompact if and only if the action of Γ on its limit set Λ in $\partial\mathbb{H}^d$ satisfies an *expansion* property. Then, he uses this expansion property to give a *symbolic coding* for infinite quasigeodesic rays in Γ. This coding gives a way to see that the correspondence between geodesic rays in Γ and points in Λ is stable under small perturbations of the representation. Notably, in [KKL19], Kapovich-Kim-Lee follow Sullivan's basic procedure to provide a detailed proof of his result in a much more general setting, in particular showing stability properties of uniform lattices in higher-rank Lie groups. The ideas in that paper guide some of our approach below.

In [BPS19], Bochi-Potrie-Sambarino use a related technique to prove that P_k-Anosov representations of hyperbolic groups into $\mathrm{PGL}(d+1,\mathbb{R})$ are stable. They consider the *geodesic automaton* for a hyperbolic group Γ. This is a finite directed graph $\mathcal{G}$ whose edges are labelled with generators of Γ. Bochi-Potrie-Sambarino show that a representation $\rho : \Gamma \to \mathrm{PGL}(d+1,\mathbb{R})$ is P_k-Anosov if and only if there is a way to assign an open subset U_v of $\mathbb{RP}^d$ to each vertex v of $\mathcal{G}$ so that each edge (a,b) of $\mathcal{G}$ corresponds to an inclusion

$$\rho(\gamma_{a,b}) \cdot \overline{U_b} \subset U_a. \tag{5.1}$$

This shows that it is possible to verify if a small deformation of an Anosov representation is Anosov by checking a finite set of open conditions.

In a sense, this approach can be thought of as a kind of *generalized ping-pong argument* for non-free groups. Indeed, when Γ is actually a free group, the inclusions in (5.1) are exactly the inclusions of sets needed to set up ping-pong.

We take a somewhat similar approach to prove our relative stability theorem. Given an extended geometrically finite representation ρ, we construct a *relative* *quasigeodesic automaton* $\mathcal{G}$, and open subsets $U_v \subset G/P$ for each vertex v of $\mathcal{G}$, satisfying certain inclusions corresponding to the edges of $\mathcal{G}$. We show that if such a system of open subsets also exists for a nearby representation ρ', then ρ' must also be extended geometrically finite.

It is helpful to keep an important special case in mind. If A, B are finitely generated groups, then $A * B$ is hyperbolic relative to the collection of conjugates of A and B. In this case, the inclusions of sets determined by a relative quasigeodesic automaton are nothing more than the inclusions of sets required to set up a ping-pong argument proving that a representation of $A * B$ is discrete and faithful.

For a general relatively hyperbolic group Γ, we essentially use the convergence dynamics of Γ to encode points in $\partial(\Gamma, \mathcal{H})$ using infinite paths in the graph $\mathcal{G}$. This is closely related to the idea of describing the *geodesic flow* of the group with symbolic dynamics; see for instance [Ser81].

Our approach only yields a *relative* stability result because, unlike in the non-relative case, some of the edges of the relative automaton $\mathcal{G}$ correspond to an infinite number of inclusions of open sets. The proof shows that if ρ is EGF, any deformation ρ' which respects the conditions imposed by these "parabolic edges" must also be EGF.

5.1 Basic properties

In this section we cover some basic properties of EGF representations. We then show that they generalize a notion of *relative Anosov representation* due to Kapovich-

Leeb (Theorem 5.1.7), and prove an *Anosov relativization theorem* (Theorem 5.1.10).

5.1.1 Discreteness and finite kernel

We first observe that EGF representations are always discrete with finite kernel. When $\rho : \Gamma \to G$ is an EGF representation, the action of $\rho(\Gamma)$ on the boundary set Λ is by definition an extension of the topological dynamical system $(\Gamma, \partial(\Gamma, \mathcal{H}))$. Convergence dynamics imply that the homomorphism $\Gamma \to \mathrm{Homeo}(\partial(\Gamma, \mathcal{H}))$ has finite kernel and discrete image. So the map $\Gamma \to \mathrm{Homeo}(\Lambda)$ must also have discrete image and finite kernel, and therefore so does the representation $\rho : \Gamma \to G$.

5.1.2 Shrinking the sets C_z

Let $\rho : \Gamma \to G$ be an EGF representation with boundary map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$. By assumption, we know there exists an open subset $C_z \subset G/P$ for each $z \in \partial(\Gamma, \mathcal{H})$, satisfying the *extended convergence dynamics* conditions (Definition 5.0.1). In general, there is not a canonical choice for the set C_z. We are able to make some assumptions about the properties of the C_z, however.

Proposition 5.1.1. *Let $\rho : \Gamma \to G$ be an EGF representation with boundary extension ϕ. For any $z \in \partial(\Gamma, \mathcal{H})$, we can choose the set C_z to be a subset of*

$$\mathrm{Opp}(\phi^{-1}(z)) := \{\xi \in G/P : \xi \textit{ is opposite to } \nu \textit{ for every } \nu \in \phi^{-1}(z)\}.$$

Proof. Since $\phi^{-1}(z)$ is closed, $\mathrm{Opp}(\phi^{-1}(z))$ is an open subset of G/P. And, transversality of ϕ implies that $\mathrm{Opp}(\phi^{-1}(z))$ contains $\Lambda - \phi^{-1}(z)$. So the intersection $C_z \cap \mathrm{Opp}(\phi^{-1}(z))$ is open and nonempty, meaning we can replace C_z with this intersection. $\square$

5.1.3 An equivalent characterization of EGF representations

Let $\rho : \Gamma \to G$ be a representation of a relatively hyperbolic group. It turns out than in order to check that ρ is an EGF representation, it suffices to look at the dynamics of the $\rho(\Gamma)$-action along *conical limit sequences* and sequences lying entirely in peripheral subgroups of Γ.

Proposition 5.1.2. *Let $\rho : \Gamma \to G$ be a representation of a relatively hyperbolic group, and let $\Lambda \subset G/P$ be a closed $\rho(\Gamma)$-invariant set, where $P \subset G$ is a symmetric parabolic subgroup. Suppose that $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ is a surjective ρ-equivariant antipodal map.*

Then ρ is an EGF representation if and only if for each $z \in \partial(\Gamma, \mathcal{H})$, there exists $C_z \subset G/P$, with $C_z \subset \mathrm{Opp}(\phi^{-1}(z))$ and $\Lambda - \phi^{-1}(z) \subset C_z$, such that:

1. *For any sequence $\gamma_n \in \Gamma$ limiting conically to z (with $\gamma_n^{-1} \to z_-$), any compact $K \subset C_{z_-}$, and any neighborhood U containing $\phi^{-1}(z)$, we have $\rho(\gamma_n) \cdot K \subset U$ for all sufficiently large n.*

2. *For any parabolic point p, any compact $K \subset C_p$, and any open set U containing $\phi^{-1}(p)$, for all but finitely many $\gamma \in \Gamma_p$, we have $\rho(\gamma) \cdot K \subset U$.*

In practice, Proposition 5.1.2 gives a criterion for extended geometrical finiteness which is much easier to check than the full definition, so we will use it throughout this chapter and the next. However, the proof requires the technical machinery of *relative quasigeodesic automata*, so we defer it to Section 5.5.

5.1.4 Properties of Λ

Proposition 5.1.3. *Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and let $\rho : \Gamma \to G$ be a representation which is EGF with respect to a symmetric parabolic P, with boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$. Then Λ contains the P-limit set of $\rho(\Gamma)$.*

Proof. Let $\xi \in G/P$ be a flag in the P-limit set of $\rho(\Gamma)$. Then there is a P-contracting sequence $\rho(\gamma_n)$ for $\gamma_n \in \Gamma$ and a flag $\xi_- \in G/P$ such that $\rho(\gamma_n)\eta$ converges to ξ for any η in $\mathrm{Opp}(\xi_-)$. Up to subsequence $\gamma_n^{\pm 1}$ converges to $z_\pm \in \partial(\Gamma, \mathcal{H})$, so for any flag $\eta \in C_{z_-}$, the sequence $\rho(\gamma_n)\eta$ subconverges to a point in $\phi^{-1}(z_+)$. But since $\mathrm{Opp}(\xi_-)$ is open and dense, for some $\eta \in C_{z_-}$ we have $\rho(\gamma_n)\eta \to \xi$ and hence $\xi \in \phi^{-1}(z_+)$. $\square$

In particular, Proposition 5.1.3 implies that the EGF boundary set $\Lambda \subset G/P$ of an EGF representation $\rho : \Gamma \to G$ must always contain the *P-proximal limit set* of $\rho(\Gamma)$. (Recall that $g \in G$ is *P-proximal* if it has a unique attracting fixed point in G/P; the *P-proximal limit set* of a subgroup of G is the closure of the set of attracting fixed points of P-proximal elements).

We will see that most of the power of EGF representations lies in the fact that their associated boundary extensions $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ do not have to be homeomorphisms (so the Bowditch boundary of Γ does not need to be equivariantly embedded in any flag manifold). However, it turns out that it is always possible to choose the boundary extension ϕ so that it has a well-defined inverse on *conical limit points* in $\partial(\Gamma, \mathcal{H})$. In fact, we can even get a somewhat precise description of all the fibers of ϕ. Concretely, we have the following:

Proposition 5.1.4. *Let $\rho : \Gamma \to G$ be an EGF representation, with boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$. There is a $\rho(\Gamma)$-invariant closed subset $\Lambda' \subset G/P$ and a ρ-equivariant map $\phi' : \Lambda' \to \partial(\Gamma, \mathcal{H})$ such that:*

1. *$\phi' : \Lambda' \to \partial(\Gamma, \mathcal{H})$ is also a boundary extension for ρ,*
2. *for every $z \in \partial_{\mathrm{con}}(\Gamma, \mathcal{H})$, $\phi'^{-1}(z)$ is a singleton, and*
3. *for every $p \in \partial_{\mathrm{par}}(\Gamma, \mathcal{H})$, $\phi'^{-1}(p)$ is the closure of the set of all accumulation points of orbits $\gamma_n \cdot x$ for γ_n a sequence of distinct elements in Γ_p and $x \in C_p$.*

We will prove Proposition 5.1.4 at the end of Section 5.6, where it will follow as a consequence of the proof of the relative stability theorem for EGF representations (Theorem 5.6.2)—see Remark 5.6.20.

We will rely on both Proposition 5.1.2 and Proposition 5.1.4 to prove the rest of the results in this section (which are not needed anywhere else in this chapter).

5.1.5 Relatively asymptotically embedded representations

EGF representations give a strict generalization of the *relatively asymptotically embedded* representations of Kapovich and Leeb. We recall the definition here.

Definition 5.1.5 ([KL18], Definition 7.1)**.** Let Γ be a subgroup of G and suppose $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair. Let $P_{\tau_{\mathrm{mod}}} \subset G$ be a symmetric parabolic subgroup.

The subgroup Γ is *relatively τ_{mod}-asymptotically embedded* if it is $P_{\tau_{\mathrm{mod}}}$-divergent, and there is a Γ-equivariant antipodal embedding $\partial(\Gamma, \mathcal{H}) \to \Lambda_{\tau_{\mathrm{mod}}} \subset G/P$, where $\Lambda_{\tau_{\mathrm{mod}}}$ is the $P_{\tau_{\mathrm{mod}}}$-limit set of Γ.

Here, we say an embedding $\psi : \partial(\Gamma, \mathcal{H}) \to G/P$ is *antipodal* if for every distinct ξ_1, ξ_2 in $\partial(\Gamma, \mathcal{H})$, $\psi(\xi_1)$ and $\psi(\xi_2)$ are opposite flags.

Remark 5.1.6. When Γ is a *hyperbolic* group (and the collection of peripheral subgroups $\mathcal{H}$ is empty), then the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ is identified with the Gromov boundary $\partial\Gamma$. In this case, a relatively asymptotically embedded representation is just called an *asymptotically embedded* representation. τ_{mod}-asymptotically embedded representations are exactly the same as $P_{\tau_{\mathrm{mod}}}$-Anosov representations (see [KLP17], Theorem 1.1).

In general, it is possible to define P-Anosov representations for a *non-symmetric* parabolic subgroup P. However, there is no loss of generality in assuming that P is symmetric: a representation $\rho : \Gamma \to G$ is P-Anosov if and only if it is P'-Anosov for a symmetric parabolic subgroup $P' \subset G$ depending only on P.

The relationship between EGF representations and asymptotically embedded representations is given by the following

Theorem 5.1.7. *Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair. A representation $\rho : \Gamma \to G$ is relatively τ_{mod}-asymptotically embedded if and only if ρ is EGF with respect to $P_{\tau_{\mathrm{mod}}}$ and has a boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ which is a homeomorphism.*

To prove this theorem, the main thing we need to show is:

Proposition 5.1.8. *Let $\rho : \Gamma \to G$ be an EGF representation with respect to $P_{\tau_{\mathrm{mod}}}$, and suppose that the boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ is a homeomorphism. Then:*

1. *$\rho(\Gamma)$ is $P_{\tau_{\mathrm{mod}}}$-divergent, and Λ is the $P_{\tau_{\mathrm{mod}}}$-limit set of $\rho(\Gamma)$.*

2. *The sets C_z for $z \in \partial(\Gamma, \mathcal{H})$ can be taken to be*

$$\mathrm{Opp}(\phi^{-1}(z)) = \{\nu \in G/P : \nu \textit{ is opposite to } \phi^{-1}(z)\}.$$

Proof. (1). Let γ_n be any infinite sequence of elements in Γ. After extracting a subsequence, we have $\gamma_n^{\pm 1} \to z_\pm$, and since ϕ is a homeomorphism, $\rho(\gamma_n)$ converges to the point $\phi^{-1}(z_+)$ uniformly on compacts in the open set C_{z_-}. Then Proposition 2.3.7 implies that $\rho(\gamma_n)$ is $P_{\tau_{\mathrm{mod}}}$-divergent, with unique $P_{\tau_{\mathrm{mod}}}$-limit point $\phi^{-1}(z_+) \in \Lambda$.

(2). The fact that ϕ is antipodal is exactly the statement that the sets $\mathrm{Opp}(\phi^{-1}(z))$ contain $\Lambda - \phi^{-1}(z)$ for every $z \in \partial(\Gamma, \mathcal{H})$, so we just need to see that the appropriate dynamics hold for these sets. Let γ_n be an infinite sequence in Γ with $\gamma_n^{\pm 1} \to z_\pm$ for $z_\pm \in \partial(\Gamma, \mathcal{H})$.

We know that for open subsets $U_\pm \subset G/P$, we have $\rho(\gamma_n) \cdot U_+ \to \phi^{-1}(z_+)$ and $\rho(\gamma_n^{-1})U_- \to \phi^{-1}(z_-)$, uniformly on compacts. Proposition 2.3.7 implies that $\rho(\gamma_n)$ and $\rho(\gamma_n^{-1})$ are both P-divergent with unique P-limit points $\phi^{-1}(z_+)$, $\phi^{-1}(z_-)$. So in fact by Lemma 2.3.8 $\rho(\gamma_n)$ converges to $\phi^{-1}(z_+)$ uniformly on compacts in $\mathrm{Opp}(\phi^{-1}(z_-))$. □

Proof of Theorem 5.1.7. Proposition 5.1.8 ensures that if ρ is an EGF representation, and the boundary extension ϕ is a homeomorphism, then ρ is $P_{\tau_{\mathrm{mod}}}$-divergent and ϕ^{-1} is an antipodal embedding whose image is the $P_{\tau_{\mathrm{mod}}}$-limit set.

On the other hand, if ρ is relatively τ_{mod}-asymptotically embedded, with boundary embedding $\psi : \partial(\Gamma, \mathcal{H}) \to \Lambda$, for each $z \in \partial(\Gamma, \mathcal{H})$, we can take

$$C_z = \mathrm{Opp}(\psi(z)).$$

Antipodality means that C_z contains $\Lambda - \psi(z)$, and $P_{\tau_{\mathrm{mod}}}$-divergence and Lemma 2.3.8 imply that $\rho(\Gamma)$ has the approriate convergence dynamics. $\square$

When Γ is a *hyperbolic* group (i.e. if Γ is relatively hyperbolic relative to an empty collection of peripheral subgroups), then the Bowditch boundary of Γ is identified with its Gromov boundary, and *relative* τ_{mod}-asymptotic embeddedness is the same as the notion of τ_{mod}-asymptotic embeddedness defined in [KLP17]. A representation $\rho : \Gamma \to G$ is τ_{mod}-asymptotically embedded if and only if it is $P_{\tau_{\mathrm{mod}}}$-Anosov (see [KLP17], Theorem 1.1).

In particular, Theorem 5.1.7 implies:

Corollary 5.1.9. *Let $\rho : \Gamma \to G$ be a representation of a hyperbolic group. Then, when Γ is equipped with the trivial peripheral structure $\mathcal{H} = \emptyset$, ρ is EGF with respect to P if and only if ρ is P-Anosov.*

5.1.6 Relativization

We now turn to the situation where we have an EGF representation of a *hyperbolic* group Γ with a *nonempty* collection of peripheral subgroups. That is, for some invariant set $\Lambda \subset G/P$, we have an EGF boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$, where $\partial(\Gamma, \mathcal{H})$ is the Bowditch boundary of Γ with peripheral structure $\mathcal{H}$.

We want to prove the following theorem:

Theorem 5.1.10. *Let Γ be a hyperbolic relative to a collection of subgroups $\mathcal{H}$, and suppose that each $H \in \mathcal{H}$ is hyperbolic.*

If $\rho : \Gamma \to G$ is an EGF representation with respect to P for the peripheral structure $\mathcal{H}$, and ρ restricts to a P-Anosov representation on each $H \in \mathcal{H}$, then ρ is a P-Anosov representation of Γ.

For the rest of this section, we assume that Γ is a hyperbolic group, and $\mathcal{H}$ is a collection of subgroups of Γ so that the pair $(\Gamma, \mathcal{H})$ is relatively hyperbolic. We let $\rho : \Gamma \to G$ be an EGF representation for the pair $(\Gamma, \mathcal{H})$ with respect to a symmetric parabolic subgroup $P \subset G$, and we assume that for each $H \in \mathcal{H}$, $\rho|_H : H \to G$ is P-Anosov, with Anosov limit map $\psi_H : \partial H \to G/P$.

The main step in the proof is to observe that it is always possible to choose the boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ so that Λ is equivariantly homeomorphic to the Gromov boundary of Γ (which we here denote $\partial\Gamma$).

Whenever Γ is a hyperbolic group and $\mathcal{H}$ is a collection of subgroups so that $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair, there is an explicit description of the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ in terms of the Gromov boundary $\partial\Gamma$ of Γ—see [Ger12], [GP13], or [Tra13]. Specifically, we can say:

Proposition 5.1.11. *There is an equivariant surjective map $\phi_\Gamma : \partial\Gamma \to \partial(\Gamma, \mathcal{H})$ such that for each conical limit point z in $\partial(\Gamma, \mathcal{H})$, $\phi_\Gamma^{-1}(z)$ is a singleton, and for each parabolic point $p \in \partial(\Gamma, \mathcal{H})$ with $H = \mathrm{Stab}_\Gamma(p)$, $\phi_\Gamma^{-1}(p)$ is an embedded copy of ∂H in $\partial\Gamma$.*

In our situation, we can see that the boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ satisfies similar properties.

Lemma 5.1.12. *There is a closed $\rho(\Gamma)$-invariant subset $\Lambda' \subset G/P$ and an EGF boundary extension $\phi' : \Lambda' \to \partial(\Gamma, \mathcal{H})$ such that:*

1. *For each conical limit point $z \in \partial(\Gamma, \mathcal{H})$, $\phi'^{-1}(z)$ is a singleton.*

2. *For each parabolic point $p \in \partial(\Gamma, \mathcal{H})$, with $H = \mathrm{Stab}_\Gamma(p)$, we have $\phi'^{-1}(p) = \psi_H(\partial H)$.*

Proof. We choose Λ' as in Proposition 5.1.4. The only thing we need to check is that for $H = \mathrm{Stab}_\Gamma(p)$, the set $\psi_H(\partial H)$ is exactly the closure of the set of accumulation points of $\rho(H)$-orbits in C_p. But since we may assume C_p is contained in $\mathrm{Opp}(\psi_H(\partial H))$, this follows immediately from the fact that $\rho(H)$ is P-divergent and the closed set $\psi_H(\partial H)$ is the P-limit set of $\rho(H)$.

□

Next we need a lemma which will allow us to characterize the Gromov boundary of Γ as an extension of the Bowditch boundary $\partial(\Gamma, \mathcal{H})$. First recall that if Γ acts as a convergence group on a space Z, the *limit set* of Γ is the set of points $z \in Z$ such that for some $y \in Z$ and some sequence $\gamma_n \in \Gamma$, we have

$$\gamma_n|_{Z - \{y\}} \to z$$

uniformly on compacts.

Lemma 5.1.13. *Let Γ act on compact metrizable spaces X and Y, and let $\phi_X : X \to \partial(\Gamma, \mathcal{H})$, $\phi_Y : Y \to \partial(\Gamma, \mathcal{H})$ be equivariant surjective maps such that for every conical limit point $z \in \partial(\Gamma, \mathcal{H})$, $\phi_X^{-1}(z)$ and $\phi_Y^{-1}(z)$ are both singletons, and for every parabolic point $p \in \partial(\Gamma, \mathcal{H})$, $H = \mathrm{Stab}_\Gamma(p)$ acts as a convergence group on X and Y, with limit sets $\phi_X^{-1}(p)$, $\phi_Y^{-1}(p)$ equivariantly homeomorphic to ∂H.*

Then for any sequences $z_n, z'_n \in \partial_{\mathrm{con}}(\Gamma, \mathcal{H})$, we have

$$\lim_{n\to\infty} \phi_X^{-1}(z_n) = \lim_{n\to\infty} \phi_X^{-1}(z'_n)$$

if and only if

$$\lim_{n\to\infty} \phi_Y^{-1}(z_n) = \lim_{n\to\infty} \phi_Y^{-1}(z'_n).$$

Proof. We proceed by contradiction, and suppose that for a pair of sequences $z_n, z'_n \in \partial_{\mathrm{con}}(\Gamma, \mathcal{H})$, we have

$$\lim_{n\to\infty} \phi_X^{-1}(z_n) = \lim_{n\to\infty} \phi_X^{-1}(z'_n) = x,$$

but

$$\lim_{n\to\infty} \phi_Y^{-1}(z_n) \neq \lim_{n\to\infty} \phi_Y^{-1}(z'_n).$$

After taking a subsequence we may assume z_n converges to $z \in \partial(\Gamma, \mathcal{H})$, and that $y_n = \phi_Y^{-1}(z_n)$ converges to y and $y'_n = \phi_Y^{-1}(z'_n)$ converges to y' for $y \neq y'$. By continuity, we have

$$\phi_Y(y) = \phi_Y(y') = \phi_X(x) = z.$$

Since ϕ_X and ϕ_Y are bijective on $\phi_X^{-1}(\partial_{\mathrm{con}}(\Gamma, \mathcal{H}))$ and $\phi_Y^{-1}(\partial_{\mathrm{con}}(\Gamma, \mathcal{H}))$ respectively, we must have $z = p$ for a parabolic point $p \in \partial_{\mathrm{par}}(\Gamma, \mathcal{H})$. Let $H = \mathrm{Stab}_\Gamma(p)$.

Since p is a bounded parabolic point, we can find sequences of group elements $h_n, h'_n \in H$ so that for a fixed compact subset $K \subset \partial(\Gamma, \mathcal{H}) - \{p\}$, we have

$$h_n z_n \in K, \quad h'_n z'_n \in K. \tag{5.2}$$

This implies that no subsequence of $h_n y_n$ or $h'_n y'_n$ converges to a point in $\phi_Y^{-1}(p)$.

Then, since H acts as a convergence group on Y with limit set $\phi_Y^{-1}(p)$, up to subsequence there are points $u, u' \in \phi_Y^{-1}(p)$ so that h_n converges to a point in $\phi_Y^{-1}(p)$ uniformly on compacts in $Y - \{u\}$, and h'_n converges to a point in $\phi_Y^{-1}(p)$ uniformly on compacts in $Y - \{u'\}$. So, we must have $u = y$ and $u' = y'$.

This means that the sequences h_n^{-1} and $h_n'^{-1}$ have distinct limits in the compactification $\overline{H} = H \sqcup \partial H$. So, there are distinct points $v, v' \in \phi_X^{-1}(p)$ so that (again up to subsequence) h_n converges to a point in $\phi_X^{-1}(p)$ uniformly on compacts in $X - \{v\}$, and h'_n converges to a point in $\phi_X^{-1}(p)$ uniformly on compacts in $X - \{v'\}$. Without loss of generality, we can assume $x \neq v$.

But then $\phi_X^{-1}(z_n)$ lies in a compact subset of $X - \{v\}$, so $h_n \phi_X^{-1}(z_n)$ converges to a point in $\phi_X^{-1}(p)$ and $h_n z_n$ converges to p. But this contradicts (5.2) above. $\square$

Proposition 5.1.14. *If the set Λ satisfies the conclusions of Lemma 5.1.12, then Λ is equivariantly homeomorphic to the Gromov boundary of Γ.*

Proof. Let $\phi_\Gamma : \partial\Gamma \to \partial(\Gamma, \mathcal{H})$ denote the quotient map identifying the limit set of each $H \in \mathcal{H}$ to the parabolic point p with $H = \mathrm{Stab}_\Gamma(p)$. For each conical limit point $z \in \partial(\Gamma, \mathcal{H})$, the fiber $\phi_\Gamma^{-1}(z)$ is a singleton. So, there is an equivariant bijection f from $\phi_\Gamma^{-1}(\partial_{\mathrm{con}}(\Gamma, \mathcal{H}))$ to $\phi^{-1}(\partial_{\mathrm{con}}(\Gamma, \mathcal{H}))$.

Moreover, since $\phi_\Gamma^{-1}(\partial_{\mathrm{con}}(\Gamma,\mathcal{H}))$ is Γ-invariant, and the action of Γ on its Gromov boundary $\partial\Gamma$ is minimal, $\phi_\Gamma^{-1}(\partial_{\mathrm{con}}(\Gamma,\mathcal{H}))$ is dense in $\partial\Gamma$. We claim that f extends to a continuous injective map $\partial\Gamma \to \Lambda$ by defining $f(x) = \lim f(x_n)$ for any sequence $x_n \to x$.

To see this, we can apply Lemma 5.1.13, taking $\partial\Gamma = X$ and $\Lambda = Y$. We know that Γ always acts on its own Gromov boundary as a convergence group (so in particular each $H \in \mathcal{H}$ acts on $\partial\Gamma$ as a convergence group with limit set ∂H). And, since ρ restricts to a P-Anosov representation on each $H \in \mathcal{H}$, for any infinite sequence $h_n \in H$, up to subsequence there are $u, u_- \in \psi_H(\partial H)$ so that $\rho(h_n)$ converges to u uniformly on compacts in $\mathrm{Opp}(u_-)$. Antipodality of ϕ implies that $\rho(h_n)$ converges to u uniformly on compacts in $\Lambda - \psi_H(\partial H)$. The other hypotheses of Lemma 5.1.13 follow from Proposition 5.1.11 and Lemma 5.1.12.

We still need to check that f is actually surjective. We know that f restricts to a bijection on $\phi_\Gamma^{-1}(\partial_{\mathrm{con}}(\Gamma,\mathcal{H}))$, and that f takes $\phi_\Gamma^{-1}(p)$ to $\phi^{-1}(p)$ for each parabolic point p in $\partial(\Gamma,\mathcal{H})$. So we just need to check that for every $H \in \mathcal{H}$, f restricts to a surjective map $\partial H \to \psi_H(\partial H)$. If H is non-elementary, this must be the case because the action of H on ∂H is minimal and f maps ∂H into $\psi_H(\partial H)$ as an invariant closed subset. Otherwise, H is virtually cyclic and ∂H, $\psi_H(\partial H)$ both contain exactly two points. Then injectivity of f implies surjectivity.

So we conclude that there is a continuous bijection $f : \partial\Gamma \to \Lambda$, and since $\partial\Gamma$ is compact and Λ is metrizable, f is a homeomorphism. $\square$

We let $f : \Lambda \to \partial\Gamma$ denote the equivariant homeomorphism from Proposition 5.1.14. The final step in the proof of Theorem 5.1.10 is the following:

Proposition 5.1.15. *The equivariant homeomorphism $f : \Lambda \to \partial\Gamma$ extends the convergence dynamics of Γ on its Gromov boundary $\partial\Gamma$.*

Proof. By Proposition 5.1.2, we just need to show that if $\gamma_n \in \Gamma$ is a conical limit sequence with $\gamma_n^{\pm 1} \to z_\pm$ for $z_\pm \in \partial\Gamma$, then $\rho(\gamma_n)$ converges to $\phi^{-1}(z_+)$ uniformly on compacts in $\mathrm{Opp}(\phi^{-1}(z_-))$.

We consider two cases:

Case 1: $\phi \circ f(z_+)$ is a parabolic point p in $\partial(\Gamma, \mathcal{H})$. In this case, γ_n lies along a quasigeodesic ray in Γ limiting to some $z_+ \in \partial H$. This means that for a bounded sequence $b_n \in \Gamma$, we have $\gamma_n b_n \in H$. Since ρ restricts to a P-Anosov representation on H, this means that $\rho(\gamma_n b_n)$ is P-divergent with unique P-limit point $\psi_H(z_+)$. But then $\rho(\gamma_n)$ is also P-divergent with unique P-limit point $\psi_H(z_+)$.

Since p is a parabolic point (in particular, *not* a conical limit point), and γ_n limits to p, the sequence γ_n^{-1} must also limit to p. So we apply the same reasoning to see that $\rho(\gamma_n^{-1})$ is also P-divergent with unique P-limit point $\psi_H(z_-)$. Then we are done by Lemma 2.3.8.

Case 2: $\phi \circ f(z_+)$ is a conical limit point in $\partial(\Gamma, \mathcal{H})$. In this case, subsequences of both γ_n and γ_n^{-1} are conical limit sequences for the action of Γ on $\partial(\Gamma, \mathcal{H})$. So, $\phi \circ f(z_-)$ is also a conical limit point in $\partial(\Gamma, \mathcal{H})$. So the desired result follows directly from the fact that ρ is an EGF representation.

$\square$

Proof of Theorem 5.1.10. Let Γ be hyperbolic, let $\mathcal{H}$ be a collection of subgroups such that $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair, and let $\rho : \Gamma \to G$ be an EGF representation with respect to P, for the peripheral structure $\mathcal{H}$.

Suppose that ρ restricts to a P-Anosov representation on each $H \in \mathcal{H}$. Proposition 5.1.15 implies that ρ is *also* an EGF representation of Γ for its *empty* peripheral structure, whose boundary extension can be chosen to be a homeomorphism. Then Theorem 5.1.7 says that ρ is relatively asymptotically embedded (again for the empty peripheral structure on Γ). Then we apply Corollary 5.1.9 to complete the proof. $\square$

5.2 Relative quasigeodesic automata

In the next three sections, we develop the technical tools needed to prove the main results of the chapter: namely, a *relative quasigeodesic automaton* for a relatively hyperbolic group Γ acting on a flag manifold G/P, and a *system of open sets* in G/P which is in some sense *compatible* with both the relative quasigeodesic automaton and the action of Γ on G/P.

The basic idea is motivated by the computational theory of hyperbolic groups. Given a hyperbolic group Γ with finite generating set S, it is always possible to find a finite directed graph $\mathcal{G}$, with edges labeled by elements of S, so that directed paths on $\mathcal{G}$ starting at a fixed vertex $v_{\mathrm{id}} \in \mathcal{G}$ are in one-to-one correspondence with geodesic words in Γ. The graph $\mathcal{G}$ is called a *geodesic automaton* for Γ.

Geodesic automata are really a manifestation of the local-to-global principle for geodesics in hyperbolic metric spaces: the fact that the automaton exists means that it is possible to recognize a geodesic path in a hyperbolic group just by looking at bounded-length subpaths.

In this section, we consider a *relative* version of a geodesic automaton. This is a finite directed graph $\mathcal{G}$ which encodes the behavior of quasigeodesics in the coned-off Cayley graph of a relatively hyperbolic group Γ. Eventually, our goal is to build such an automaton by looking at the dynamics of the action of Γ on its Bowditch boundary $\partial(\Gamma, \mathcal{H})$. The main result of this section is Proposition 5.2.13, which says that we can construct such a *relative quasigeodesic automaton* for a relatively hyperbolic pair $(\Gamma, \mathcal{H})$ using an open covering of the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ which satisfies certain technical conditions.

In this section and the next, we will work in the general context of a relatively hyperbolic group Γ acting by homeomorphisms on a connected compact metrizable space M, before returning to the case where M is a flag manifold G/P for the rest of the chapter.

Throughout the rest of this section, we fix a *non-elementary* relatively hyperbolic pair $(\Gamma, \mathcal{H})$, and let $\Pi \subset \partial_{\mathrm{par}}(\Gamma, \mathcal{H})$ be a finite set, containing exactly one point from each Γ-orbit in $\partial_{\mathrm{par}}(\Gamma, \mathcal{H})$. We also fix a finite generating set S for Γ, which allows us to refer to the coned-off Cayley graph $\mathrm{Cay}(\Gamma, S, \mathcal{P})$ (Definition 2.2.17).

Definition 5.2.1. A Γ-*graph* is a finite directed graph $\mathcal{G}$ where each vertex v is labelled with a subset $T_v \subset \Gamma$, which is either:

- A singleton $\{\gamma\}$, with $\gamma \neq \mathrm{id}$, or
- A cofinite subset of a coset $g\Gamma_p$ for some $p \in \Pi$, $g \in \Gamma$.

A sequence $\{\alpha_n\} \subset \Gamma$ is a *$\mathcal{G}$-path* if $\alpha_n \in T_{v_n}$ for a vertex path $\{v_n\}$ in $\mathcal{G}$.

Remark 5.2.2. We will often refer to "the" vertex path $\{v_n\}$ corresponding to a $\mathcal{G}$-path $\{\alpha_n\}$, although we will never actually verify that such a vertex path is uniquely determined by the sequence of group elements $\{\alpha_n\}$ in Γ.

A vertex of a Γ-graph which is labeled by a cofinite subset of a (necessarily unique) coset $g\Gamma_p$ is a *parabolic vertex*. If v is a parabolic vertex, we let $p_v = g \cdot p$ denote the corresponding parabolic point in $\partial_{\mathrm{par}}(\Gamma, \mathcal{H})$.

Remark 5.2.3. It will be convenient to allow parabolic vertices to be labeled by *cofinite* subsets of peripheral cosets (instead of just the entire coset) when we construct Γ-graphs using the convergence dynamics of the Γ-action on $\partial(\Gamma, \mathcal{H})$.

Definition 5.2.4. Let $z \in \partial(\Gamma, \mathcal{H})$. We say that a $\mathcal{G}$-path $\{\alpha_n\}$ *limits to* z if either:

- $z \in \partial_{\mathrm{con}}(\Gamma, \mathcal{H})$, $\{\alpha_n\}$ is infinite, and the sequence

$$\{\gamma_n = \alpha_1 \cdots \alpha_n\}_{n=1}^{\infty}$$

  limits to z in the compactification $\overline{\Gamma} = \Gamma \sqcup \partial(\Gamma, \mathcal{H})$, or

- $z \in \partial_{\mathrm{par}}(\Gamma, \mathcal{H})$, $\{\alpha_n\}$ is a finite $\mathcal{G}$-path whose corresponding vertex path $\{v_n\}$ ends at a parabolic vertex v_N, and

$$z = \alpha_1 \cdots \alpha_{N-1} p_{v_N}.$$

Definition 5.2.5. Let $\mathcal{G}$ be a Γ-graph. The *endpoint* of a finite $\mathcal{G}$-path $\{\alpha_n\}_{n=1}^{N}$ is

$$\alpha_1 \cdots \alpha_N.$$

Definition 5.2.6. A Γ-graph $\mathcal{G}$ is a *relative quasigeodesic automaton* if:

1. There is a constant $D > 0$ so that for any infinite $\mathcal{G}$-path α_n, the sequence

$$\{\gamma_n = \alpha_1 \cdots \alpha_n\} \subset \Gamma$$

lies Hausdorff distance at most D from a geodesic ray in $\mathrm{Cay}(\Gamma, S, \mathcal{P})$, based at the identity.

2. For every $z \in \partial(\Gamma, \mathcal{H})$, there exists a $\mathcal{G}$-path limiting to z.

One way to think of a relative quasigeodesic automaton is that it gives us a system for finding quasigeodesic representatives of every element in the group. More concretely, we have the following:

Lemma 5.2.7. *Let $\mathcal{G}$ be a relative quasigeodesic automaton. There is a constant $R > 0$ so that set of endpoints of $\mathcal{G}$-paths is R-dense in Γ.*

Proof. If Γ is hyperbolic and $\mathcal{H}$ is empty, then this is a consequence of the Morse lemma and the fact that the union of the images of all infinite geodesic rays based at the identity in Γ is coarsely dense in Γ (see [Bog97]).

If $\mathcal{H}$ is nonempty, there is some $R > 0$ so that the union of all of the cosets $g \cdot \Gamma_p$ for $p \in \Pi$ is R-dense in Γ. So it suffices to show that for each $p \in \Pi$, there is some $R > 0$ so that all but R elements in any coset $g \cdot \Gamma_p$ are the endpoints of a $\mathcal{G}$-path.

For any such coset $g \cdot \Gamma_p$, we can find a finite $\mathcal{G}$-path $\{\alpha_n\}_{n=1}^{N-1}$ limiting to the vertex $g \cdot p$. That is,

$$g \cdot p = \alpha_1 \cdots \alpha_{N-1} p_{v_N}.$$

By definition $p_{v_N} = g' \cdot p$ with T_{v_N} a cofinite subset of the coset $g'\Gamma_p$ That is,

$$g \cdot \Gamma_p = \alpha_1 \cdots \alpha_{N-1} g' \Gamma_p,$$

so for all but finitely many $\gamma \in g \cdot \Gamma_p$ (depending only on the size of the complement of T_{v_N} in $g' \cdot \Gamma_p$), we can find $\alpha_N \in g'\Gamma_p$ with

$$\alpha_1 \cdots \alpha_N = \gamma.$$

$\square$

Remark 5.2.8. In general, we do *not* require the set of elements in Γ labelling the vertices of a relative quasigeodesic automaton $\mathcal{G}$ to generate the group Γ (although the proposition above implies that they at least generate a finite-index subgroup).

5.2.1 Compatible systems of open sets

A relative quasigeodesic automaton always exists for any relatively hyperbolic group (although we will not prove this fact in full generality). We will give a way to construct a relative quasigeodesic automaton using the *convergence group dynamics* of a group acting on its Bowditch boundary.

Definition 5.2.9. Suppose that Γ acts on a metrizable space M by homeomorphisms, and let $\mathcal{G}$ be a Γ-graph. A *$\mathcal{G}$-compatible system of open sets* for the action of Γ on M

is an assignment of an open subset $U_v \subset M$ to each vertex v of $\mathcal{G}$ such that for each edge $e = (v, w)$ in $\mathcal{G}$, for some $\varepsilon > 0$, we have

$$\alpha \cdot N_M(U_w, \varepsilon) \subset U_v \tag{5.3}$$

for all $\alpha \in T_v$.

Remark 5.2.10. If $\mathcal{G}$ has no parabolic vertices (so each set T_v contains a single group element $\alpha_v \in \Gamma$), then (5.3) is equivalent to requiring $\alpha_v \cdot \overline{U_w} \subset U_v$ for every edge (v, w) in $\mathcal{G}$. When $\mathcal{G}$ has parabolic vertices (so T_v may be infinite), (5.3) is in general a stronger condition.

Proposition 5.2.11. *Let $\mathcal{G}$ be a Γ-graph, and let $\{U_v : v \text{ vertex of } \mathcal{G}\}$ be a $\mathcal{G}$-compatible system of subsets of $\partial(\Gamma, \mathcal{H})$ for the action of Γ on $\partial(\Gamma, \mathcal{H})$.*

There is a constant $D > 0$ satisfying the following: let $\{\alpha_n\}$ be an infinite $\mathcal{G}$-path, corresponding to a vertex path $\{v_n\}$, and suppose the sequence $\{\gamma_n = \alpha_1 \cdots \alpha_n\}$ is divergent in Γ. Then for any point z in the intersection

$$U_\infty = \bigcap_{n=1}^{\infty} \alpha_1 \cdots \alpha_n U_{v_{n+1}},$$

the sequence γ_n lies within Hausdorff distance D of a geodesic ray in $\mathrm{Cay}(\Gamma, S, \mathcal{P})$ *tending towards z.*

Proof. Fix a point $z \in U_\infty$, and write $z = z_+$ and $U_n = U_{v_n}$. We first claim that there is a uniform $\varepsilon > 0$ and a point $z_- \in \partial(\Gamma, \mathcal{H})$ such that

$$d(\gamma_n^{-1} z_+, \gamma_n^{-1} z_-) > \varepsilon \tag{5.4}$$

for all $n \geq 0$.

To prove the claim, choose a uniform $\varepsilon > 0$ so that for every vertex v in $\mathcal{G}$, we have $N(U_v, \varepsilon) \neq \partial(\Gamma, \mathcal{H})$, and for every edge (v, w) in $\mathcal{G}$ and every $\alpha \in T_v$, we have $\alpha \cdot N(U_w, \varepsilon) \subset U_v$. Then we choose some $z_- \in \partial(\Gamma, \mathcal{H}) - \overline{N}(U_1, \varepsilon)$.

By the $\mathcal{G}$-compatibility condition, we know that for any n, $\gamma_n U_{n+1} \subset \ldots \subset \gamma_1 U_2 \subset U_1$, so we know that $d(z_+, z_-) > \varepsilon$.

Then, for any $n \geq 1$, we have

$$\gamma_n^{-1} z_+ \in U_{n+1}.$$

Moreover since $\gamma_n N(U_{n+1}, \varepsilon) \subset U_1$, we also have

$$\gamma_n^{-1} z_- \in \partial(\Gamma, \mathcal{H}) - N(U_{n+1}, \varepsilon).$$

So for all n we have $d(\gamma_n^{-1} z_+, \gamma_n^{-1} z_-) > \varepsilon$, which establishes that (5.4) holds for all n.

Now, consider a bi-infinite geodesic c in a Gromov model Y for Γ joining z_+ and z_-. The sequence of geodesics $\gamma_n^{-1} \cdot c$ has endpoints in $\partial Y = \partial(\Gamma, \mathcal{H})$ lying distance at least ε apart, so each geodesic in the sequence passes within a uniformly bounded neighborhood of a fixed basepoint $y_0 \in Y$. Therefore $\gamma_n \cdot y_0$ lies in a uniformly bounded neighbood of the geodesic c.

Since γ_n is divergent, $\gamma_n y_0$ can only accumulate at either z_+ or z_-. But in fact $\gamma_n y_0$ can only accumulate at z_+—for in the construction of c above, we could have chosen any z_- in the nonempty open set $\partial(\Gamma, \mathcal{H}) - \overline{N}(U_1, \varepsilon)$, and since $\partial(\Gamma, \mathcal{H})$ is perfect there is at least one such $z'_- \neq z_-$.

This implies that γ_n is a conical limit sequence in Γ, limiting to z_+. Since the distance between γ_n and γ_{n+1} is bounded in $\mathrm{Cay}(\Gamma, S, \mathcal{P})$, the desired conclusion follows. $\square$

Definition 5.2.12. Let $\mathcal{G}$ be a Γ-graph. An infinite $\mathcal{G}$-path $\{\alpha_n\}$ is *divergent* if the sequence $\{\gamma_n = \alpha_1 \cdots \alpha_n\}$ leaves every bounded subset of Γ.

We say that a Γ-graph $\mathcal{G}$ is divergent if *every* infinite $\mathcal{G}$-path is divergent.

Whenever $\{U_v\}$ is a $\mathcal{G}$-compatible system of open sets for a Γ-graph $\mathcal{G}$, one can think of a $\mathcal{G}$-path $\{\alpha_n\}$ as giving a *symbolic coding* of a point in the intersection

$$\alpha_1 \cdots \alpha_n U_{n+1}.$$

The following proposition gives a way to *construct* such a coding for a given point $z \in \partial(\Gamma, \mathcal{H})$, given an appropriate pair of open coverings of the Bowditch boundary $\partial(\Gamma, \mathcal{H})$ compatible with a Γ-graph $\mathcal{G}$.

Proposition 5.2.13. *Let $\mathcal{G}$ be a* divergent *Γ-graph. Suppose that for each vertex $a \in \mathcal{G}$, there exist open subsets V_a, W_a of $\partial(\Gamma, \mathcal{H})$ such that the following conditions hold:*

1. *The sets $\{W_a\}$ give a $\mathcal{G}$-compatible system of sets for the action of Γ on $\partial(\Gamma, \mathcal{H})$.*
2. *For all vertices a, we have $V_a \subset W_a$ and $\overline{W_a} \neq \partial(\Gamma, \mathcal{H})$.*
3. *The sets V_a give an open covering of $\partial(\Gamma, \mathcal{H})$.*
4. *For every $z \in \partial(\Gamma, \mathcal{H})$ and every non-parabolic vertex a such that $z \in V_a$, there is an edge (a, b) in $\mathcal{G}$ such that $\alpha_a^{-1} \cdot z \in V_b$ for $\{\alpha_a\} = T_a$.*

5. *For every* $z \in \partial(\Gamma, \mathcal{H})$ *and every* parabolic *vertex* a *such that* $z \in V_a - \{p_a\}$, *there is an edge* (a, b) *in* $\mathcal{G}$ *and* $\alpha \in T_a$ *such that* $\alpha^{-1} \cdot z \in V_b$.

Then $\mathcal{G}$ *is a relative quasigeodesic automaton for* Γ.

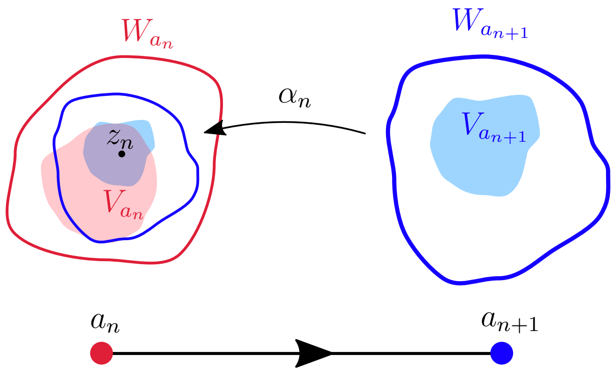

Figure 5.1: Illustration for the proof of Proposition 5.2.13. The group element α_n nests an ε-neighborhood of $W_{a_{n+1}}$ inside of W_{a_n} whenever $\alpha_n \cdot V_{a_{n+1}}$ intersects V_{a_n}.

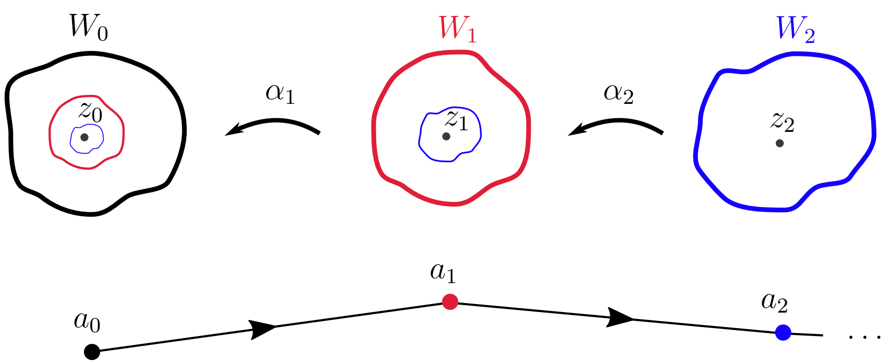

Figure 5.2: By iterating the nesting procedure backwards, we produce an infinite $\mathcal{G}$-path and a sequence of subsets intersecting in the initial point $z = z_0$.

Proof. Proposition 5.2.11 implies that any infinite $\mathcal{G}$-path lies finite Hausdorff distance from a geodesic ray in $\mathrm{Cay}(\Gamma, S, \mathcal{P})$. So, we just need to show that every $z \in \partial(\Gamma, \mathcal{H})$ is the limit of a $\mathcal{G}$-path.

The idea behind the proof is to use the fact that the sets V_a cover $\partial(\Gamma, \mathcal{H})$ to show that we can keep "expanding" a neighborhood of z in $\partial(\Gamma, \mathcal{H})$ to construct a path in $\mathcal{G}$ limiting to z. The $\{V_a\}$ covering tells us how to find the next edge in the path, and the $\{W_a\}$ cover gives us the $\mathcal{G}$-compatible system we need to show that the path is a geodesic.

We let A denote the vertex set of $\mathcal{G}$. When $a \in A$ is not a parabolic vertex, we write $T_a = \{\gamma_a\}$.

Case 1: z is a conical limit point. Fix $a \in A$ so that $z \in V_a$. We take $z_0 = z$, $a_0 = a$, and define sequences $\{z_n\}_{n=0}^{\infty} \subset \partial_{\mathrm{con}}\Gamma$, $\{a_n\}_{n=0}^{\infty} \subset A$, and $\{\alpha_n\}_{n=1}^{\infty} \subset \Gamma$ as follows:

- If a_n is not a parabolic vertex, then we choose $\alpha_{n+1} = \gamma_{a_n}$. Let $z_{n+1} = \alpha_{n+1}^{-1} \cdot z_n$. Since conical limit points are invariant under the action of Γ, z_{n+1} is a conical limit point. By condition 4, there is a vertex a_{n+1} satisfying $z_{n+1} \in V_{a_{n+1}}$ with (a_n, a_{n+1}) an edge in $\mathcal{G}$.

- If a_n is a parabolic vertex, then since z_n is a conical limit point, $z_n \neq p$ for $p = p_{a_n}$. Then condition 5 implies that there exists some $\alpha_{n+1} \in T_{a_n}$ so that $\alpha_{n+1}^{-1} \cdot z_n \in V_{a_{n+1}}$ for an edge (a_n, a_{n+1}) in $\mathcal{G}$. Again, $z_{n+1} = \alpha_{n+1}^{-1} \cdot z_n$ must be a conical limit point since $\partial_{\mathrm{con}}\Gamma$ is Γ-invariant.

The sequence $\{\alpha_n\}$ necessarily gives a $\mathcal{G}$-path. By assumption the sequence

$$\gamma_n = \alpha_1 \cdots \alpha_n$$

is divergent. And by construction $z = \gamma_n z_n$ lies in $\gamma_n W_{a_n}$ for all n. So, Proposition 5.2.11 implies that γ_n is a conical limit sequence, limiting to z. See Figure 5.2.

Case 2: z is a parabolic point. As before fix $a \in A$ so that $z \in V_a$, and take $z_0 = z$, $a_0 = a$. We inductively define sequences z_n, a_n, α_n as before, but we claim that for some finite N, a_N is a parabolic vertex with $z_N = p_{a_N}$. For if not, we can build an infinite $\mathcal{G}$-path (as in the previous case) limiting to z. But then, Proposition 5.2.11 would imply that z is actually a conical limit point. So, we must have

$$z = \gamma_N a_N = \alpha_1 \cdots \alpha_N a_N$$

as required.

□

Remark 5.2.14. In a typical application of Proposition 5.2.13, it will not be possible to construct the open coverings $\{V_a\}$ and $\{W_a\}$ so that $V_a = W_a$ for all vertices a. In particular we expect this to be impossible whenever $\partial(\Gamma, \mathcal{H})$ is connected.

To conclude this section, we make one more observation about systems of $\mathcal{G}$-compatible sets as in Proposition 5.2.13.

Lemma 5.2.15. *Let Γ be a relatively hyperbolic group, let $\mathcal{G}$ be a Γ-graph, and let $\{V_a\}$, $\{W_a\}$ be an assignment of open subsets of $\partial(\Gamma, \mathcal{H})$ to vertices of $\mathcal{G}$ satisfying the hypotheses of Proposition 5.2.13.*

Fix $z \in \partial_{\mathrm{con}}\Gamma$ and $N \in \mathbb{N}$. There exists $\delta > 0$ so that if $d(z, z') < \delta$, then there are $\mathcal{G}$-paths $\{\alpha_n\}, \{\beta_n\}$ limiting to z, z' respectively, with $\alpha_i = \beta_i$ for all $i < N$.

Proof. Let $\{\alpha_n\}$ be a $\mathcal{G}$-path limiting to z coming from the construction in Proposition 5.2.13, passing through vertices v_n. We choose $\delta > 0$ small enough so that if

$d(z, z') < \delta$, then z' lies in the set

$$\alpha_1 \cdots \alpha_N V_{v_{n+1}}.$$

Then for every $i < N$, we have

$$\alpha_i^{-1} \alpha_{i-1}^{-1} \cdots \alpha_1^{-1} z' \in V_{v_{i+1}}.$$

As in Proposition 5.2.13, we can then extend $\{\alpha_n\}_{n=1}^{N-1}$ to a $\mathcal{G}$-path limiting to z'. $\square$

5.3 Extended convergence dynamics

Let Γ be a relatively hyperbolic group acting on a connected compact metrizable space M. In this section, we will show that if the action of Γ on M extends the convergence dynamics of Γ (Definition 5.0.1), then we can construct a relative quasigeodesic automaton $\mathcal{G}$ and a $\mathcal{G}$-compatible system of open subsets of M which are in some sense reasonably well-behaved with respect to the group action.

To give the precise statement, we let $\Lambda \subset M$ be a closed Γ-invariant subset, and let $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ be an equivariant, surjective, and continuous map satisfying the following: for each $z \in \partial(\Gamma, \mathcal{H})$, there is an open set $C_z \subset M$ containing $\Lambda - \phi^{-1}(z)$ such that:

1. For any sequence $\gamma_n \in \Gamma$ limiting conically to z, with $\gamma_n^{-1} \to z_-$, any open set U containing $\phi^{-1}(z)$, and any compact $K \subset C_{z_-}$, we have $\gamma_n \cdot K \subset U$ for all sufficiently large n.

2. For any parabolic point p, any compact $K \subset C_p$, and any open set U containing $\phi^{-1}(p)$, for all but finitely many $\gamma \in \Gamma_p$, we have $\gamma \cdot K \subset U$.

Note that in particular, any map extending convergence dynamics satisfies these conditions. For the rest of this section, however, we *only* assume that (1) and (2) both hold for our map ϕ. In this context, we will show:

Proposition 5.3.1. *For any $\varepsilon > 0$, there is a relative quasigeodesic automaton $\mathcal{G}$ for Γ, a $\mathcal{G}$-compatible system of open sets $\{U_v\}$ for the action of Γ on M, and a $\mathcal{G}$-compatible system of open sets $\{W_v\}$ for the action of Γ on $\partial(\Gamma, \mathcal{H})$ such that:*

1. *For every vertex v, there is some $z \in W_v$ so that*

$$\phi^{-1}(W_v) \subset U_v \subset N_M(\phi^{-1}(z), \varepsilon).$$

2. *For every $p \in \Pi$, there is a parabolic vertex a with $p_a = p$. Moreover, for every parabolic vertex w with $p_w = g \cdot p$, (a, b) is an edge of $\mathcal{G}$ if and only if (w, b) is an edge of $\mathcal{G}$.*

3. *If $q = g \cdot p$ for $p \in \Pi$, a is a parabolic vertex with $p_a = q$, and (a, b) is an edge of $\mathcal{G}$, then $q \in W_a$, $U_b \subset C_p$ and $g\Gamma_p \cdot W_b$ contains $\partial(\Gamma, \mathcal{H}) - \{q\}$.*

Remark 5.3.2. By equivariance of ϕ, for each $p \in \partial_{\mathrm{par}}\Gamma$, we can replace C_p with $\Gamma_p \cdot C_p$ and assume that C_p is Γ_p-invariant (and that if $q = g \cdot p$, then $C_q = g \cdot C_p$).

The proof of Proposition 5.3.1 involves some technicalities, so we first outline the general approach:

1. For each $z \in \partial(\Gamma, \mathcal{H})$, we construct a pair V_z, W_z of small open neighborhoods of z and a subset $T_z \subset \Gamma$ so that for each $\alpha \in T_z$, α^{-1} is "expanding" about

some point in V_z. When z is a conical limit point, then we can choose a single element $\alpha_z \in \Gamma$ which expands about every point in V_z. When z is a parabolic point, we may use a different element of Γ to "expand" about each $u \in V_z - \{z\}$. We choose V_z, W_z, and T_z so that if α^{-1} is "expanding" about $u \in V_z$, and $\alpha^{-1}u \in V_y$, then $\alpha^{-1}W_z \supset W_y$. See Figure 5.3.

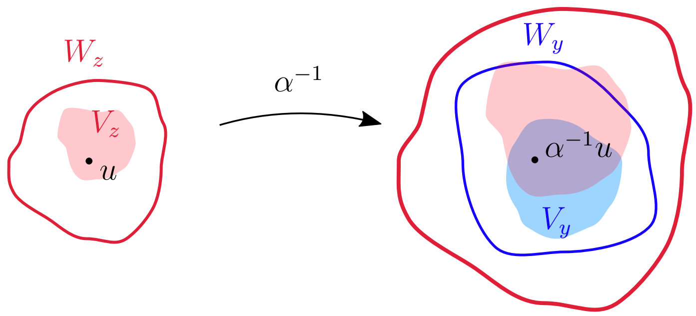

Figure 5.3: The group element α^{-1} is "expanding" about $u \in V_z$. We will construct V_z, W_z and V_y, W_y so that if $\alpha^{-1}u$ lies in V_y, then $\alpha^{-1}W_z$ contains W_y. Equivalently, we get the containment $\alpha W_y \subset W_z$ illustrated earlier in Figure 5.1.

2. Using compactness of $\partial(\Gamma, \mathcal{H})$, we pick a finite set of points $a \in \partial(\Gamma, \mathcal{H})$ so that the sets $\{V_a\}$ give an open covering of $\partial(\Gamma, \mathcal{H})$. These points in $\partial(\Gamma, \mathcal{H})$ are identified with the vertices of a Γ-graph $\mathcal{G}$. We define the edges of $\mathcal{G}$ in such a way so that if, for some $\alpha \in T_a$, α^{-1} expands about $u \in V_a$ and $\alpha^{-1}u \in V_b$, then there is an edge from a to b. This ensures that $\{W_a\}$ is a $\mathcal{G}$-compatible system of open subsets of $\partial(\Gamma, \mathcal{H})$.

3. Simultaneously, we construct a $\mathcal{G}$-compatible system $\{U_a\}$ of open sets in M by taking U_a to be a small neighborhood of $\phi^{-1}(a)$. The idea is to use the extended convergence dynamics to ensure that if, for some $\alpha \in T_z$, α^{-1} "expands"

about some $u \in V_z$ and the point $\alpha^{-1}u$ lies in V_y, then $\alpha^{-1}U_z$ contains U_y. See Figure 5.6 below.

4. Finally, we use Proposition 5.2.13 to prove that $\mathcal{G}$ is actually a relative quasi-geodesic automaton. The open sets V_a, W_a are constructed exactly to satisfy the conditions of the proposition, so the main thing to check in this step is that the graph $\mathcal{G}$ is actually divergent (using the action of Γ on M).

Throughout the rest of the section, we will work with fixed metrics on both $\partial(\Gamma, \mathcal{H})$ and M. Critically, none of our "expansion" arguments will depend sensitively on the precise choice of metric. That is, in the sketch above, when we say that some group element $\alpha \in \Gamma$ "expands" on a small open subset U of a metric space X, we just mean that αU is quantifiably "bigger" than U, and *not* that for any $x, y \in U$, we have $d(\alpha \cdot x, \alpha \cdot y) \geq C \cdot d(x, y)$ for some expansion constant C. Lemma 5.3.5 and Lemma 5.3.7 below describe precisely what we mean by "bigger." The general idea is captured by the following example.

Example 5.3.3. We consider the group $\mathrm{PGL}(2, \mathbb{Z})$. While $\mathrm{PGL}(2, \mathbb{Z})$ is virtually a free group (and therefore word-hyperbolic), it is also relatively hyperbolic, relative to the collection $\mathcal{H}$ of conjugates of the parabolic subgroup $\left\{ \begin{pmatrix} \pm 1 & t \\ 0 & 1 \end{pmatrix} : t \in \mathbb{Z} \right\}$.

Since $\mathrm{PGL}(2, \mathbb{Z})$ acts with finite covolume on the hyperbolic plane $\mathbb{H}^2$, the Bowditch boundary of the pair $(\mathrm{PGL}(2, \mathbb{Z}), \mathcal{H})$ is equivariantly identified with $\partial \mathbb{H}^2$, the visual boundary of $\mathbb{H}^2$. Given a non-parabolic point $w \in \partial \mathbb{H}^2$, we can find an element of $\mathrm{PGL}(2, \mathbb{Z})$ which "expands" a neighborhood of w. There are two distinct possibilities:

1. Suppose w is in a small neighborhood V_z of a conical limit point $z \in \partial\mathbb{H}^2$. Then choose some loxodromic element $\gamma \in \mathrm{PGL}(2, \mathbb{Z})$ whose attracting fixed point is close to z. Then, if W_z is a slightly larger neighborhood of z, $\gamma^{-1} \cdot W_z$ is large enough to contain a uniformly large neighborhood of $\gamma^{-1} \cdot w$. See Figure 5.4.

2. On the other hand, suppose w is in a small neighborhood V_q of a parabolic fixed point $q \in \partial\mathbb{H}^2$, but $w \neq q$. We can find some element $\gamma \in \Gamma_q = \mathrm{Stab}_\Gamma(q)$ so that γ^{-1} takes w into a fundamental domain for the action of Γ_q on $\partial\mathbb{H}^2 - \{q\}$. Then, if W_q is a slightly larger neighborhood of q, $\gamma^{-1} \cdot W_q$ is again large enough to contain a uniformly large neighborhood of $\gamma^{-1} \cdot w$. See Figure 5.5.

There is a slight issue with this approach: in the second case above (when w is close to a parabolic point q), it is actually not quite good enough to "expand" a neighborhood of w by using Γ_q to push w into a fundamental domain for Γ_q on $\partial\mathbb{H}^2 - \{q\}$. The reason is that there might be no such fundamental domain which is actually far away from $\partial\mathbb{H}^2 - \{q\}$. We resolve this issue by instead choosing γ to lie in a *coset* $g\Gamma_p$, where $q = gp$ for some $p \in \Pi$. Then $\gamma^{-1} \cdot w$ lies in a fundamental domain for Γ_p on $\partial\mathbb{H}^2 - \{p\}$, which allows us to get *uniform* control on the size of the expanded neighborhood $\gamma^{-1}W_q$.

The two technical lemmas below (Lemma 5.3.5 and Lemma 5.3.7) essentially say that one can set up this kind of expansion *simultaneously* on the Bowditch boundary of our relatively hyperbolic group Γ *and* in a neighborhood of the Γ-invariant set $\Lambda \subset M$. The precise formulation of the expansion condition found in these two lemmas is best motivated by the proof of Proposition 5.3.10 below, which

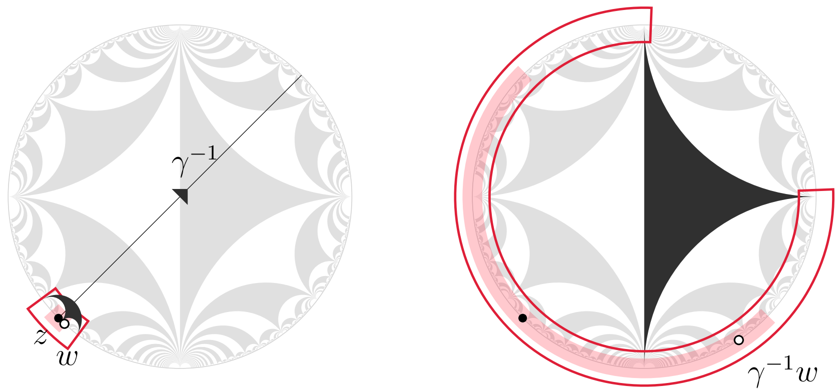

Figure 5.4: For any point w in a sufficiently small neighborhood V_z (pink) of z, the expanded neighborhood $\gamma^{-1}W_z$ (red) contains a uniform neighborhood of $\gamma^{-1}w$.

shows that the "expanding" open sets we construct give rise to a $\mathcal{G}$-compatible system of open sets on a Γ-graph $\mathcal{G}$.

Lemma 5.3.4. *There exists $\varepsilon > 0$ (depending on ϕ and D) so that for any $a, b \in \partial(\Gamma, \mathcal{H})$ with $d(a, b) > D$, the ε-neighborhood of $\phi^{-1}(a)$ in M is contained in C_b.*

Proof. Since $\phi^{-1}(z)$ is closed in M, such an $\varepsilon > 0$ exists for any fixed pair of distinct $(a, b) \in \partial(\Gamma, \mathcal{H})^2$. Then the result follows, since the space of pairs $(a, b) \in (\partial(\Gamma, \mathcal{H}))^2$ satisfying $d(a, b) > D$ is compact. $\square$

Lemma 5.3.5. *There exists $\varepsilon_{\text{con}} > 0, \delta_{\text{con}} > 0$ satisfying the following: for any $\varepsilon > 0$, $\delta > 0$ with $\varepsilon < \varepsilon_{\text{con}}$, $\delta < \delta_{\text{con}}$, and every conical limit point z, we can find:*

- *A group element $\gamma_z \in \Gamma$*
- *Open subsets $W_z, V_z \subset \partial(\Gamma, \mathcal{H})$ with $z \in V_z \subset W_z$*

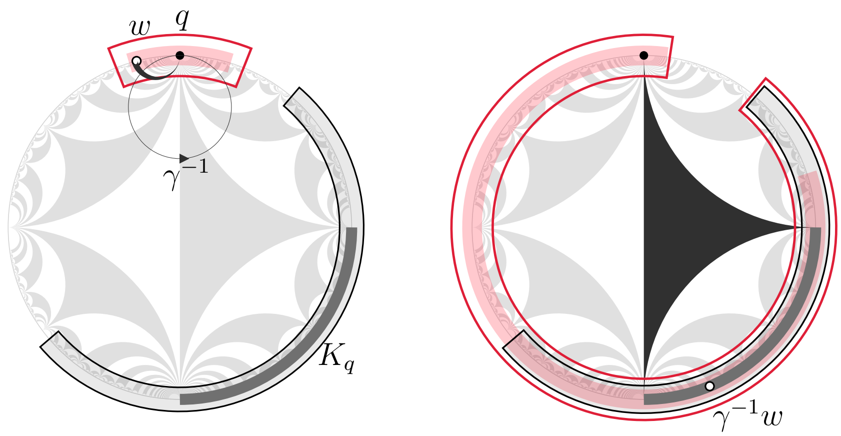

Figure 5.5: For any point $w \neq q$ in a neighborhood V_q (pink) of the parabolic point q, we find some $\gamma \in \Gamma_q$ so that $\gamma^{-1}w$ lies in K_q (dark gray), a fundamental domain for the action of Γ_q on $\partial\mathbb{H}^2 - \{q\}$. The expanded neighborhood $\gamma^{-1}W_q$ (red) contains a uniform neighborhood of K_q, so $\gamma^{-1}W_q$ contains a uniform neighborhood of $\gamma^{-1}w$.

such that:

1. $\operatorname{diam}(W_z) < \delta$,

2. *In* $\partial(\Gamma, \mathcal{H})$, *we have*

$$N_{\partial\Gamma}(\gamma_z^{-1}V_z, \delta) \subset \gamma_z^{-1}W_z.$$

3. *In* M *we have*

$$N_M(\gamma_z^{-1}\phi^{-1}(W_z), 2\varepsilon) \subset \gamma_z^{-1}N_M(\phi^{-1}(z), \varepsilon).$$

Remark 5.3.6. Conditions (1) and (2) together imply that for any $y, z \in \partial_{\text{con}}\Gamma$, if $\gamma_z^{-1}V_z$ intersects V_y, then $\gamma_z W_y \subset W_z$. Later, we will see that condition (3) implies that if $\gamma_z^{-1}V_z$ intersects V_y, then also $\gamma_z N_M(\phi^{-1}(y), 2\varepsilon) \subset N_M(\phi^{-1}(z), \varepsilon)$ (giving us the inclusion indicated by Figure 5.3).

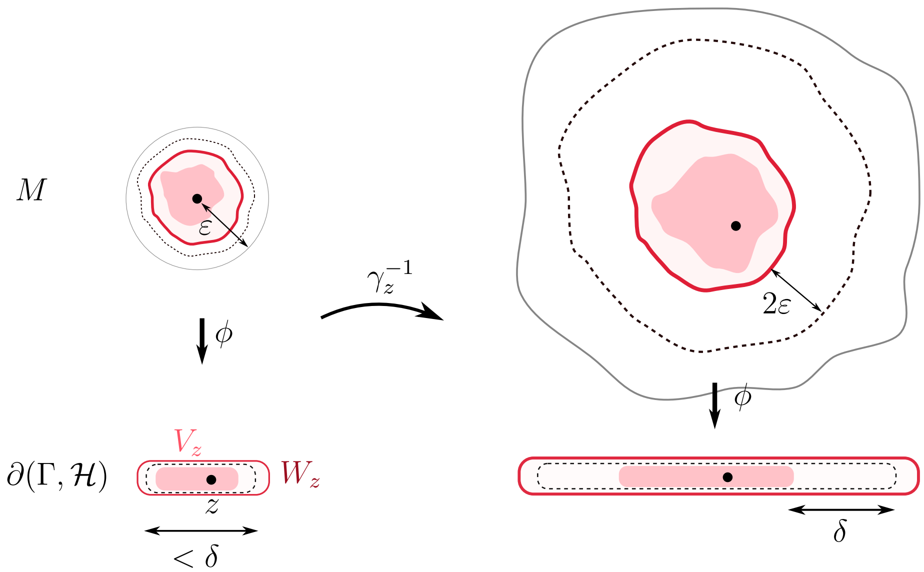

Figure 5.6: The group element γ_z^{-1} is "expanding" about $V_z \subset \partial(\Gamma, \mathcal{H})$: while W_z has diameter $< \delta$, $\gamma_z^{-1} W_z$ contains a δ-neighborhood of $\gamma_z^{-1} V_z$. At the same time, γ_z^{-1} enlarges an ε-neighborhood of $\phi^{-1}(z)$ in M, so that it contains a 2ε-neighborhood of $\gamma_z^{-1}\phi^{-1}(W_z)$.

Proof. For a conical limit point z, we choose a sequence γ_n so that for distinct $a, b \in \partial(\Gamma, \mathcal{H})$, we have $\gamma_n^{-1} z \to a$ and $\gamma_n^{-1} w \to b$ for any $w \neq z$. That is, γ_n limits conically to z in $\overline{\Gamma}$, and γ_n^{-1} limits conically to b. Since the Γ-action on distinct pairs in $\partial(\Gamma, \mathcal{H})$ is cocompact (Proposition 2.2.8), we may assume that $d(a, b) > D$ for a uniform constant $D > 0$.

We choose $\varepsilon_{\mathrm{con}} > 0$ from Lemma 5.3.4 so that if $a, b \in \partial(\Gamma, \mathcal{H})$ satisfy $d(a, b) > D/2$, then a $2\varepsilon_{\mathrm{con}}$-neighborhood of $\phi^{-1}(a)$ is contained in C_b. Let $\varepsilon > 0$ satisfy $\varepsilon < \varepsilon_{\mathrm{con}}$, and let δ satisfy $\delta < \delta_{\mathrm{con}} := D/4$.

By the triangle inequality, we have $d(c, b) > D/2$ for all $c \in B_{\partial\Gamma}(a, 2\delta)$, so the closed 2ε-neighborhood of $\phi^{-1}(B_{\partial\Gamma}(a, 2\delta))$ is contained in C_b. This means that we can choose n large enough so that

$$\gamma_n \cdot N_M(\phi^{-1}(B(a, 2\delta)), 2\varepsilon)$$

is contained in $N_M(\phi^{-1}(z), \varepsilon)$ and

$$\gamma_n \cdot B_{\partial\Gamma}(a, 2\delta)$$

is contained in $B_{\partial\Gamma}(z, \delta/2)$. We let $\gamma_z = \gamma_n$ for this large n, and take

$$W_z = \gamma_z \cdot B_{\partial\Gamma}(a, 2\delta)$$

and

$$V_z = \gamma_z \cdot B_{\partial\Gamma}(a, \delta).$$

□

The next lemma is a version of Lemma 5.3.5 for parabolic points. As before, we want to show that for a point q in the Bowditch boundary, we can find a neighborhood W_q of q in $\partial(\Gamma, \mathcal{H})$ with uniformly bounded diameter δ, and group elements $\gamma \in \Gamma$ so that γ^{-1} enlarges W_q enough to contain a 2δ-neighborhood of $\gamma^{-1}z$, for some z close to q. Simultaneously we want to choose γ so that γ^{-1} enlarges an ε-neighborhood of $\phi^{-1}(q)$ in a similar manner. This case is more complicated, because we need to allow γ to depend on the point $z \in W_q$: if q is a parabolic point in $\partial(\Gamma, \mathcal{H})$, then in general there is *not* a single group element in Γ which expands distances in a neighborhood of q.

Lemma 5.3.7. *For each point $p \in \Pi$, there exist constants $\varepsilon_p > 0$, $\delta_p > 0$ such that for any $q = g \cdot p \in \Gamma \cdot p$, any $\varepsilon < \varepsilon_p$, and any $\delta < \delta_p$, we can find:*

- *A cofinite subset* T_q *of the coset* $g\Gamma_p$*,*

- *Open neighborhoods* V_q, W_q *of* $\partial(\Gamma, \mathcal{H})$*, with* $q \in V_q \subset W_q$*,*

- *Open neighborhoods* $\hat{V}_q, \hat{W}_q$ *of* $\partial(\Gamma, \mathcal{H})$ *with* $\hat{V}_q \subset \hat{W}_q$

such that:

1. $\mathrm{diam}(W_q) < \delta$*, and* $\phi^{-1}(W_q) \subset N(\phi^{-1}(q), \varepsilon)$*.*

2. *in* $\partial(\Gamma, \mathcal{H})$*, we have*

$$N_{\partial\Gamma}(\hat{V}_q, \delta) \subset \hat{W}_q.$$

3. *For every* $z \in V_q - \{q\}$*, there exists* $\gamma \in T_q$ *with* $\gamma^{-1} \cdot z \in \hat{V}_q$*.*

4. *For every* $\gamma \in T_q$*, we have*

$$N_M(\phi^{-1}(\hat{W}_q), 2\varepsilon) \subset \gamma^{-1} N_M(\phi^{-1}(q), \varepsilon)$$

and

$$\hat{W}_q \subset \gamma^{-1} W_q.$$

5. $N_M(\phi^{-1}(\hat{W}_q), 2\varepsilon)$ *is contained in* C_p *and* $g\Gamma_p \cdot \hat{V}_q$ *contains* $\partial(\Gamma, \mathcal{H}) - \{q\}$*.*

Remark 5.3.8. If $z \in V_q - \{q\}$ and $\gamma^{-1} z \in \hat{V}_q$ for some $\gamma \in T_q$, we think of γ^{-1} as "expanding" about z. Conditions (1) and (2) imply that if $\gamma^{-1} z \in V_y$ for some $\gamma \in T_q$, then $\hat{W}_q$ contains W_y, and by condition (4), $\gamma^{-1} W_q$ contains W_y. Here V_y, W_y are the sets from either Lemma 5.3.5 or Lemma 5.3.7.

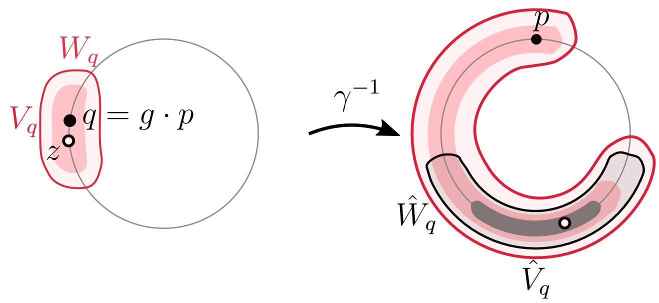

Figure 5.7: The behavior of sets in $\partial(\Gamma, \mathcal{H})$ described by Lemma 5.3.7. Given $z \in V_q$, we pick an element $\gamma \in g\Gamma_p$ so that a uniformly large neighborhood of $\gamma^{-1}z$ is contained in $\gamma^{-1}W_q$. We cannot pick γ^{-1} to expand the metric everywhere close to q—some points in V_q get contracted close to p.

Proof. Pick a compact set $K \subset \partial(\Gamma, \mathcal{H}) - \{p\}$ so that $\Gamma_p \cdot K$ covers $\partial(\Gamma, \mathcal{H}) - \{p\}$. Choose δ_p small enough so that the closure of $N_{\partial\Gamma}(K, 2\delta_p)$ does not contain p. Then, for any $\delta < \delta_p$, we can pick

$$\hat{V}_q = N_{\partial\Gamma}(K, \delta), \quad \hat{W}_q = N_{\partial\Gamma}(K, 2\delta).$$

We can choose ε_p sufficiently small so that a $2\varepsilon_p$-neighborhood of $\phi^{-1}(N_{\partial\Gamma}(K, 2\delta_p))$ is contained in C_p. Now, fix $\varepsilon < \varepsilon_p$. We claim that for a cofinite subset $T_q \subset g \cdot \Gamma_p$, for any $\gamma \in T_q$, we have

$$\gamma \cdot \hat{W}_q \subset \quad B_{\partial\Gamma}(q, \delta/2) \tag{5.5}$$

$$\gamma \cdot N_M(\phi^{-1}(\hat{W}_q), 2\varepsilon) \subset \quad N_M(\phi^{-1}(q), \varepsilon) \tag{5.6}$$

To see that this claim holds, it suffices to verify that for any infinite sequence γ_n of distinct group elements in $g\Gamma_p$, (5.5) and (5.6) both hold for all sufficiently large n.

We write $\gamma_n = g \cdot \gamma_n'$ for $\gamma_n' \in \Gamma_p$. Then γ_n' converges uniformly to p on compact subsets of $\partial(\Gamma, \mathcal{H}) - \{p\}$, so γ_n converges uniformly to q on compact subsets of $\partial(\Gamma, \mathcal{H}) - \{p\}$, implying that (5.5) eventually holds. And by our assumptions, we know that

$$\gamma_n' \cdot N_M(\phi^{-1}(\hat{W}_q), 2\varepsilon) \subset g^{-1} \cdot N_M(\phi^{-1}(q), \varepsilon)$$

for sufficiently large n, implying that (5.6) also eventually holds.

So we can take W_q to be the set

$$\{q\} \cup \bigcup_{\gamma \in T_q} \gamma \cdot \hat{W}_q,$$

and V_q to be the set

$$\{q\} \cup \bigcup_{\gamma \in T_q} \gamma \cdot \hat{V}_q.$$

To see that W_q and V_q are open we just need to verify that they each contain a neighborhood of q. Since $\hat{V}_q$ and $\hat{W}_q$ each contain K, and $\Gamma_p \cdot K$ covers $\partial(\Gamma, \mathcal{H}) - \{p\}$, V_q and W_q each contain the set

$$\partial(\Gamma, \mathcal{H}) - \bigcup_{\gamma \in g\Gamma_p - T_q} \gamma K.$$

Since T_q is cofinite in $g\Gamma_p$ this is an open set containing q. □

5.3.1 Construction of the relative automaton

We will construct the relative automaton $\mathcal{G}$ satisfying the conditions of Proposition 5.3.1 by choosing a suitable open covering of $\partial(\Gamma, \mathcal{H})$, and then using compactness to take a finite subcover. The subsets of this subcover will be the vertices of $\mathcal{G}$.

We choose constants $\varepsilon > 0$, $\delta > 0$ so that $\varepsilon < \varepsilon_{\mathrm{con}}$, $\delta < \delta_{\mathrm{con}}$ (where $\varepsilon_{\mathrm{con}}$, δ_{con} are the constants coming from Lemma 5.3.5) and $\varepsilon < \varepsilon_p$, $\delta < \delta_p$ for each $p \in \Pi$ (where ε_p, δ_p are the constants coming from Lemma 5.3.7).

Then:

- For each $z \in \partial_{\mathrm{con}}\Gamma$, we define W_z, V_z, γ_z as in Lemma 5.3.5, with parameters ε, δ.

- For each $q \in \partial_{\mathrm{par}}\Gamma$, we define $V_q, W_q, \hat{V}_q, \hat{W}_q$, and T_q as in Lemma 5.3.7, again with parameters ε, δ.

The collections of sets $\{V_z : z \in \partial_{\mathrm{con}}\Gamma\}$ and $\{V_q : q \in \partial_{\mathrm{par}}\Gamma\}$ together give an open covering of the Bowditch boundary $\partial(\Gamma, \mathcal{H})$. So we choose a finite subcover $\mathcal{V}$, which we can write as

$$\mathcal{V} = \{V_a : a \in A\}$$

where A is a finite subset of $\partial(\Gamma, \mathcal{H})$. We can in particular ensure that A contains the finite set Π.

We identify the vertices of our Γ-graph $\mathcal{G}$ with A. For each $a \in A$, the set T_a is either $\{\gamma_a\}$ (if a is a conical limit point) or T_q (if $a = q$ for a parabolic point q). Then, for each $a \in A$, we define the open sets U_a by

$$U_a = N_M(\phi^{-1}(a), \varepsilon).$$

The edges of the Γ-graph $\mathcal{G}$ are defined as follows:

- For $a, b \in A$ with $a \in \partial_{\mathrm{con}}\Gamma$, there is an edge from a to b if $(\gamma_a^{-1} \cdot V_a) \cap V_b$ is nonempty.

- If $a, b \in A$ with $a \in \partial_{\mathrm{par}}\Gamma$, there is an edge from a to b if $\hat{V}_a \cap V_b$ is nonempty.

Since $\mathcal{V}$ is an open covering of $\partial(\Gamma, \mathcal{H})$, and the sets $\hat{V}_a$ and $\gamma_a^{-1}V_a$ are nonempty, every vertex of $\mathcal{G}$ has at least one outgoing edge. Moreover, for any parabolic point a, the set $\hat{V}_a$ depends only on the orbit of a in $\partial(\Gamma, \mathcal{H})$, so $\mathcal{G}$ must satisfy condition (2) in Proposition 5.3.1.

Proposition 5.3.9. *For each $a \in A$, we have*

$$\phi^{-1}(W_a) \subset U_a.$$

Proof. When a is not a parabolic vertex, Part (3) of Lemma 5.3.5 implies:

$$\phi^{-1}(W_a) = \gamma_a\gamma_a^{-1}\phi^{-1}(W_a) \subset \gamma_a N(\gamma_a^{-1}\phi^{-1}(W_a), 2\varepsilon) \subset N_M(\phi^{-1}(a), \varepsilon) = U_a.$$

When a is a parabolic vertex, then the claim follows directly from Part (1) of Lemma 5.3.7. □

Next we verify:

Proposition 5.3.10. *The collection of sets $\{W_v\}$ and $\{U_v\}$ are both $\mathcal{G}$-compatible systems of open sets for the Γ-graph $\mathcal{G}$.*

Proof. First fix an edge (a, b) with $a \in \partial_{\mathrm{con}}\Gamma$. Since $(\gamma_a^{-1}V_a) \cap V_b$ is nonempty, part 2 of Lemma 5.3.5 implies that $\gamma_a^{-1} \cdot W_a$ contains the δ-neighborhood of some point $z \in V_b$. Since $\mathrm{diam}(W_b) < \delta$ and $V_b \subset W_b$, we can find a small $\varepsilon' > 0$ so that $\gamma_a N_{\partial\Gamma}(W_b, \varepsilon') \subset W_a$.

In particular, $\gamma_a^{-1} \cdot W_a$ contains b, which means that $N_M(\gamma_a^{-1}\phi^{-1}(W_a), 2\varepsilon)$ contains $N_M(\phi^{-1}(b), 2\varepsilon)$, which contains $N_M(U_b, \varepsilon)$. Then, part 3 of Lemma 5.3.5 implies that $\gamma_a \cdot N_M(U_b, \varepsilon)$ is contained in $N_M(\phi^{-1}(a), \varepsilon) = U_a$.

Next fix an edge (q, b) with $q \in \partial_{\mathrm{par}}\Gamma$. From part 2 of Lemma 5.3.7, we know that $\hat{W}_q$ contains the δ-neighborhood of a point $z \in \hat{V}_q \cap V_b$. Since $\mathrm{diam}(W_b) < \delta$ and $V_b \subset W_b$, this means that $\hat{W}_q$ contains an ε'-neighborhood of W_b for some small $\varepsilon' > 0$. So part 4 of Lemma 5.3.7 implies that for any $\gamma \in T_q$, we have $\gamma \cdot N(W_b, \varepsilon') \subset W_q$.

In particular $\hat{W}_q$ contains b, so $N_M(\phi^{-1}(\hat{W}_q), 2\varepsilon)$ contains $N_M(\phi^{-1}(b), 2\varepsilon)$, which contains $N_M(U_b, \varepsilon)$. Then, part 4 of Lemma 5.3.7 implies that

$$\gamma N_M(U_b, \varepsilon) \subset N_M(\phi^{-1}(q), \varepsilon) = U_q$$

for any $\gamma \in T_q$. □

We observe:

Proposition 5.3.11. *The $\mathcal{G}$-compatible systems of open subsets $\{U_v\}$ and $\{W_v\}$ satisfy (1) - (3) in Proposition 5.3.1.*

Proof. Part (1) follows directly from Proposition 5.3.9, and the fact that we defined each U_a to be the ε-neighborhood of $\phi^{-1}(a)$. Part (2) is true by the construction of the open covering $\mathcal{V}$ and the graph $\mathcal{G}$. Part (3) follows by construction and from part 5 of Lemma 5.3.7. □

To finish the proof of Proposition 5.3.1, we now just need to show:

Proposition 5.3.12. *$\mathcal{G}$ is a relative quasigeodesic automaton for the pair $(\Gamma, \mathcal{H})$.*

Proof. We apply Proposition 5.2.13, using the $\mathcal{G}$-compatible system $\{W_a\}$ and the sets $\{V_a\}$ we defined in the construction of $\mathcal{G}$.

The first three conditions of Proposition 5.2.13 are satisfied by construction. To see that conditions 4 and 5 hold, first observe that if $z \in V_a$ for a non-parabolic vertex a, then $\gamma_a^{-1} \cdot z$ lies in some V_b and (a, b) is an edge in $\mathcal{G}$. And if $z \in V_a - \{p_a\}$ for a parabolic vertex a, then part (3) of Lemma 5.3.7 says that there is some $\gamma \in T_a$ such that $\gamma^{-1} \cdot z \in \hat{V}_a$. If V_b contains $\gamma^{-1} \cdot z$, the edge (a, b) must be in $\mathcal{G}$.

It only remains to check that $\mathcal{G}$ is a divergent Γ-graph. Let $\{\alpha_n\}$ be an infinite $\mathcal{G}$-path, following a vertex path $\{v_n\}$. The $\mathcal{G}$-compatibility condition implies that $\gamma_n \overline{U}_{v_{n+1}}$ is a subset of $\gamma_{n-1} U_{v_n}$ for every n. Since M is connected and each U_v is a proper subset of M, this inclusion must be proper. This implies that in the sequence γ_n, no element can appear more than $\#A$ times and therefore γ_n is divergent. $\square$

Remark 5.3.13. This last step is the only part of the proof of Proposition 5.3.1 which uses the connectedness of M. This hypothesis is likely unnecessary, but omitting it would involve introducing additional technicalities in the construction of the sets V_a, W_a—and as stated, the proposition is strong enough for our purposes.

Note that with this hypothesis removed, Proposition 5.3.1 would imply that any non-elementary relatively hyperbolic group has a relative quasigeodesic automaton (by taking $M = \partial(\Gamma, \mathcal{H})$). As stated, the proposition only shows that such an automaton exists when $\partial(\Gamma, \mathcal{H})$ is connected.

We conclude this section by observing that one can slightly refine the construction in Proposition 5.3.1 as follows:

Proposition 5.3.14. *Fix a compact subset Z of the Bowditch boundary $\partial(\Gamma, \mathcal{H})$. Then, for any open set $U \subset M$ containing $\phi^{-1}(Z)$, there is a relative quasigeodesic* *automaton $\mathcal{G}$ and a pair of $\mathcal{G}$-compatible systems of open sets $\{U_a\}$, $\{W_a\}$ as in Proposition 5.3.1, additionally satisfying the following: any $z \in Z$ is the limit of a $\mathcal{G}$-path $\{\alpha_n\}$ (with corresponding vertex path $\{v_n\}$) such that $U_{v_1} \subset U$.*

Proof. We choose $\varepsilon > 0$ so that U contains $N_M(\phi^{-1}(Z), \varepsilon)$. We then construct our relative quasigeodesic automaton $\mathcal{G}$ as in the proof of Proposition 5.3.1, but we also choose a finite subset $A_Z \subset Z$ so that the sets V_a for $a \in A_Z$ give a finite open covering of Z. We can ensure that the vertex set A of $\mathcal{G}$ contains A_Z.

Then, for any $z \in Z$, by the construction in Proposition 5.2.13, we can find a $\mathcal{G}$-path limiting to z whose first vertex is some $a \in A_Z$. The corresponding open set for this vertex is $U_a = N_M(\phi^{-1}(a), \varepsilon) \subset U$. □

5.4 Contracting paths in flag manifolds

Let $\Gamma \subset G$ be a discrete relatively hyperbolic group, and let $\mathcal{G}$ be a Γ-graph. Fix a pair of opposite parabolic subgroups P^+, P^-. Our goal in this section is to show that under certain conditions, if $\{U_v\}$ is a $\mathcal{G}$-compatible system of open subsets of G/P^+ for the action of Γ on G/P^+, then the sequence of group elements lying along an infinite $\mathcal{G}$-path is P^+-divergent.

5.4.1 Contracting paths in Γ-graphs

Definition 5.4.1. Let Γ be a discrete subgroup of G, let $\mathcal{G}$ be a Γ-graph, and let $\{U_v\}_{v \in V(\mathcal{G})}$ be a $\mathcal{G}$-compatible system of open subsets of G/P^+. We say that a $\mathcal{G}$-path

$\{\alpha_n\}_{n\in\mathbb{N}}$ is *contracting* if the decreasing intersection

$$\bigcap_{n=1}^{\infty} \alpha_1 \cdots \alpha_n \cdot U_{v_{n+1}} \tag{5.7}$$

is a singleton in G/P^+.

Definition 5.4.2. We say that an open set $\Omega \subset G/P^+$ is a *proper domain* if the closure of Ω lies in an affine chart $\mathrm{Opp}(\xi) \subset G/P^+$ for some $\xi \in G/P^-$.

Here is the main result in this section:

Proposition 5.4.3. *Let $\mathcal{G}$ be a Γ-graph for $(\Gamma, \mathcal{H})$, and let $\{U_v\}_{v\in V(\mathcal{G})}$ be $\mathcal{G}$-compatible system of open subsets of G/P^+.*

If the set U_v is a proper domain for each vertex v of the automaton, then every infinite $\mathcal{G}$-path is contracting.

5.4.2 A metric property for bounded domains in flag manifolds

To prove Proposition 5.4.3, we consider a metric C_Ω defined by Zimmer [Zim18] on any proper domain $\Omega \subset G/P^+$. C_Ω is defined so that it is invariant under the action of G on G/P^+: for any x, y in some proper domain $\Omega \subset G/P^+$, and any $g \in G$, we have

$$C_\Omega(x, y) = C_{g\Omega}(gx, gy). \tag{5.8}$$

In general, C_Ω is not a complete metric. However, C_Ω induces the standard topology on Ω as an open subset of G/P. We will show that for a $\mathcal{G}$-path $\{\alpha_n\}$, the diameter of

$$\alpha_1 \cdots \alpha_n U_{v_{n+1}}$$

with respect to $C_{U_{v_1}}$ tends to zero as $n \to \infty$.

Zimmer's construction of C_Ω depends on an irreducible representation $\zeta : G \to \mathrm{PGL}(V)$ for some real vector space V. This is provided by a theorem of Guéritaud-Guichard-Kassel-Wienhard.

Theorem 5.4.4 ([GGKW17], see also [Zim18], Theorem 4.6)**.** *There exists a real vector space V, an irreducible representation $\zeta : G \to \mathrm{PGL}(V)$, a line $\ell \subset V$, and a hyperplane $H \subset V$ such that:*

1. *$\ell + H = V$.*

2. *The stabilizer of ℓ in G is P^+ and the stabilizer of H in G is P^-.*

3. *gP^+g^{-1} and hP^-h^{-1} are opposite if and only if $\zeta(g)\ell$ and $\zeta(h)H$ are transverse.*

The representation ζ determines a pair of embeddings $\iota : G/P^+ \to \mathbb{P}(V)$ and $\iota^* : G/P^- \to \mathbb{P}(V^*)$ by

$$\iota(gP^+) = \zeta(g)\ell, \qquad \iota^*(gP^-) = \zeta(g)H.$$

In this section, we will identify $\mathbb{P}(V^*)$ with the space of projective hyperplanes in $\mathbb{P}(V)$, by identifying the projectivization of a functional $w \in V^*$ with the projectivization of its kernel.

Definition 5.4.5. Let Ω be an open subset of G/P^+. The *dual domain* $\Omega^* \subset G/P^-$ is

$$\Omega^* = \{\nu \in G/P^- : \nu \text{ is opposite to } \xi \text{ for every } \xi \in \overline{\Omega}\}.$$

Note that Ω^* is open if and only if Ω is a proper domain.

Definition 5.4.6. Let $w_1, w_2 \in \mathbb{P}(V^*)$, and let $z_1, z_2 \in \mathbb{P}(V)$. The *cross-ratio* $[w_1, w_2; z_1, z_2]$ is defined by

$$\frac{\tilde{w}_1(\tilde{z}_2)\tilde{w}_2(\tilde{z}_1)}{\tilde{w}_1(\tilde{z}_1)\tilde{w}_2(\tilde{z}_2)},$$

where $\tilde{w}_i$, $\tilde{z}_i$ are respectively lifts of w_i and z_i in V^* and V.

Remark 5.4.7. When V is two-dimensional, we can identify the projective line $\mathbb{P}(V^*)$ with $\mathbb{P}(V)$ by identifying each $[w] \in \mathbb{P}(V^*)$ with $[\ker(w)] \in \mathbb{P}(V)$. In that case, the cross-ratio defined above agrees with the standard four-point cross-ratio on $\mathbb{RP}^1$, given by

$$[a, b; c, d] := \frac{(d-a)(c-b)}{(c-a)(d-b)}. \tag{5.9}$$

The differences in (5.9) can be measured in any affine chart in $\mathbb{RP}^1$ containing a, b, c, d. Our convention is chosen so that if we identify $\mathbb{RP}^1$ with $\mathbb{R} \cup \{\infty\}$, we have $[0, \infty; 1, z] = z$.

Definition 5.4.8. Let $\Omega \subset G/P^+$ be a proper domain. We define the function $C_\Omega : \Omega \times \Omega \to \mathbb{R}$ by

$$C_\Omega(x, y) = \sup_{\xi_1, \xi_2 \in \Omega^*} \log |[\iota^*(\xi_1), \iota^*(\xi_2); \iota(x), \iota(y)]|.$$

For any $g \in G$ and any proper domain $\Omega \subset G/P^+$, we have $(g\Omega)^* = (g\Omega^*)$. So C_Ω must satisfy the G-invariance condition (5.8).

If Ω is a properly convex subset of $\mathbb{P}(V)$, and ζ, ι, ι^* are the identity maps on $\mathrm{PGL}(V)$, $\mathbb{P}(V)$, and $\mathbb{P}(V^*)$ respectively, then C_Ω agrees with the *Hilbert metric* on Ω (see Definition 2.1.9). More generally we have:

Theorem 5.4.9 ([Zim18], Theorem 5.2)**.** *If Ω is open and bounded in an affine chart, then C_Ω is a metric on Ω which induces the standard topology on Ω as an open subset of G/P^+.*

Remark 5.4.10. This particular result in [Zim18] is stated only for noncompact simple Lie groups, but the proof only assumes that G is semisimple with no compact factor.

Since taking duals of proper domains reverses inclusions, it follows that if $\Omega_1 \subset \Omega_2$, then $C_{\Omega_1} \geq C_{\Omega_2}$. Our goal now is to sharpen this inequality, and show:

Proposition 5.4.11. *Let Ω_1, Ω_2 be proper domains in G/P^+, such that $\overline{\Omega_1} \subset \Omega_2$.*

There exists a constant $\lambda > 1$ (depending on Ω_1 and Ω_2) so that for all $x, y \in \Omega_1$,

$$C_{\Omega_1}(x,y) \geq \lambda \cdot C_{\Omega_2}(x,y).$$

A consequence is the following, which in particular implies Proposition 5.4.3.

Corollary 5.4.12. *Let $\mathcal{G}$ be a Γ-graph for a relatively hyperbolic group Γ, and let $\{U_v\}$ be a $\mathcal{G}$-compatible system of open subsets of G/P^+. If each U_v is a proper domain, then there are constants $\lambda_1, \lambda_2 > 0$ so that for any $\mathcal{G}$-path $\{\alpha_n\}$ in the Γ-graph $\mathcal{G}$, the diameter of*

$$\alpha_1 \cdots \alpha_n \cdot U_{v_{n+1}}$$

with respect to $C_{U_{v_1}}$ is at most

$$\lambda_1 \cdot \exp(-\lambda_2 \cdot n).$$

Proof. For any open set $U \subset G/P^+$ and $A \subset U$, we let $\mathrm{diam}_U(A)$ denote the diameter of A with respect to the metric C_U. We choose a uniform $\varepsilon > 0$ so that in some fixed metric on G/P^+, every edge (v, w) in $\mathcal{G}$, and every $\alpha \in T_v$, we have

$$\alpha N(U_w, \varepsilon) \subset U_v.$$

Then for each vertex set U_v, we write $U_v^\varepsilon = N(U_v, \varepsilon)$.

We take

$$\lambda_1 = \max\{\mathrm{diam}_{U_v^\varepsilon}(U_v)\}.$$

Proposition 5.4.11 implies that there exists $\lambda_v > 0$ such that for all $x, y \in U_v$, we have

$$C_{U_v}(x, y) \geq \exp(\lambda_v) \cdot C_{U_v^\varepsilon}(x, y).$$

Take $\lambda_2 = \min_v\{\lambda_v\}$. We claim that for all $n \geq 1$, we have

$$\mathrm{diam}_{U_1^\varepsilon}(\alpha_1 \cdots \alpha_n U_{v_n}) \leq \lambda_1 \exp(-\lambda_2 \cdot (n-1)).$$

We prove the claim via induction on the length of the $\mathcal{G}$-path $\{\alpha_n\}$. For $n = 1$, the claim is true because $\alpha_1 U_{v_2} \subset U_{v_1}$. For $n > 1$, we can assume

$$\lambda_1 \exp(-\lambda_2(n-2)) \geq \mathrm{diam}_{U_{v_2}^\varepsilon}(\alpha_2 \cdots \alpha_n \cdot U_{v_{n+1}}).$$

Then we have

$$\begin{aligned}
\mathrm{diam}_{U_{v_2}^\varepsilon}(\alpha_2 \cdots \alpha_n \cdot U_{v_{n+1}}) &= \mathrm{diam}_{\alpha_1 U_{v_2}^\varepsilon}(\alpha_1 \cdots \alpha_n \cdot U_{v_{n+1}}) \\
&\geq \mathrm{diam}_{U_{v_1}}(\alpha_1 \cdots \alpha_n \cdot U_{v_{n+1}}) \\
&\geq \exp(\lambda_2) \cdot \mathrm{diam}_{U_{v_1}^\varepsilon}(\alpha_1 \cdots \alpha_n \cdot U_{v_{n+1}}).
\end{aligned}$$

Finally, the claim implies the corollary because we know that

$$
\begin{aligned}
\mathrm{diam}_{U_1}(\alpha_1 \cdots \alpha_n U_{n+1}) &\leq \mathrm{diam}_{\alpha_1 U_2^{\varepsilon}}(\alpha_1 \cdots \alpha_n U_{n+1}) \\
&= \mathrm{diam}_{U_2^{\varepsilon}}(\alpha_2 \cdots \alpha_n U_{n+1}) \\
&\leq \lambda_1 \exp(\lambda_2(n-2)).
\end{aligned}
$$

So, we can replace λ_1 with $\lambda_1 \exp(-2\lambda_2)$ to get the desired result. □

We now proceed with the proof of Proposition 5.4.11. We first observe that in the special case where Ω_1, Ω_2 are properly convex subsets of real projective space, one can show the desired result essentially via the following:

Proposition 5.4.13. *Let a, b, c, d be points in $\mathbb{R}\mathrm{P}^1$, arranged so that $a < b < c < d \leq a$ in a cyclic ordering on $\mathbb{R}\mathrm{P}^1$. Then there exists a constant $\lambda > 1$, depending only on the cross-ratio $[a, b; c, d]$, so that for all distinct $x, y \in (b, c)$, we have*

$$
|\log[b, c; x, y]| \geq \lambda \cdot |\log[a, d; x, y]|.
$$

Proposition 5.4.13 is a standard fact in real projective geometry and can be verified by a computation. Note that we allow the degenerate case $a = d$: in this situation the right-hand side is identically zero for distinct $x, y \in (b, c)$. We allow no other equalities among a, b, c, d, so the cross-ratio $[a, b; c, d]$ lies in $\mathbb{R} - \{1\}$.

To apply Proposition 5.4.13 to our situation, we need to get some control on the behavior of the embeddings $\iota : G/P^+ \to \mathbb{P}(V)$ and $\iota^* : G/P^- \to \mathbb{P}(V^*)$. We do so in the next three lemmas below.

Lemma 5.4.14. *Let x, y be distinct points in G/P^+. There exists a one-parameter subgroup $g_t \subset G$ such that $\zeta(g_t)$ fixes $\iota(x)$ and $\iota(y)$, and acts nontrivially on the projective line L_{xy} spanned by $\iota(x)$ and $\iota(y)$.*

Proof. We can write $x = gP^+$ for some $g \in G$. Let $\mathfrak{a}$ denote a Cartan subalgebra of the Lie algebra $\mathfrak{g}$ of G. There is a conjugate $\mathfrak{a}'$ of $\mathfrak{a}$ such that the action of $\exp(\mathfrak{a}')$ on G/P^+ fixes both x and y (see [Ebe96], Proposition 2.21.14). So, up to the action of G on G/P^+, we can assume that x is fixed by a standard parabolic subgroup P_θ^+ conjugate to P^+, and that x, y are both fixed by the Cartan subgroup $\exp(\mathfrak{a})$.

We choose $Z \in \mathfrak{a}^+$ so that $\alpha(Z) \neq 0$ for all $\alpha \in \theta$. Then $g_t = \exp(tZ)$ is a 1-parameter subgroup of G fixing x. As $t \to +\infty$, g_t is P_θ^+-divergent, with unique attracting fixed point x.

Then [GGKW17], Lemma 3.7 implies that $\zeta(g_t)$ is P_1-divergent, where P_1 is the stabilizer of a line in V, and $\iota(x)$ is the unique one-dimensional eigenspace of $\zeta(g_t)$ whose eigenvalue has largest modulus. And, since $\zeta(g_t)$ fixes $\iota(x)$ and $\iota(y)$, $\zeta(g_t)$ preserves L_{xy}, and acts nontrivially since the eigenvalues of $\zeta(g_t)$ on $\iota(x)$ and $\iota(y)$ must be distinct. □

Lemma 5.4.15. *Let L be any projective line in $\mathbb{P}(V)$ tangent to the image of the embedding $\iota : G/P^+ \to \mathbb{P}(V)$ at a point $\iota(x)$ for $x \in G/P^+$. There exists a one-parameter subgroup g_t of G so that $\zeta(g_t)$ acts nontrivially on L with unique fixed point $\iota(x)$.*

Proof. Fix a sequence $y_n \in G/P^+$ such that $y_n \neq x$ and the projective line L_n

spanned by $\iota(x)$ and $\iota(y_n)$ converges to L. By Lemma 5.4.14, there exists $Z_n \in \mathfrak{g}$ so that $\zeta(\exp(tZ_n))$ acts nontrivially on L_n, with fixed points $\iota(x)$ and $\iota(y_n)$.

In the projectivization $\mathbb{P}(\mathfrak{g})$, $[Z_n]$ converges to some $[Z]$. Since $\zeta : G \to \mathrm{PGL}(V)$ has finite kernel, there is an induced map $\zeta : \mathbb{P}(\mathfrak{g}) \to \mathbb{P}(\mathfrak{sl}(V))$, which satisfies

$$\zeta([Z_n]) \to \zeta([Z]).$$

A continuity argument shows that the one-parameter subgroup $\zeta(\exp(tZ))$ acts nontrivially on the line L, and has unique fixed point at $\iota(x)$. $\square$

Lemma 5.4.16. *Let $\Omega \subset G/P^+$ be a proper domain, and let L be a projective line in $\mathbb{P}(V)$ which is either spanned by two points in $\iota(\Omega)$, or is tangent to $\iota(G/P^+)$ at a point $\iota(x)$ for $x \in \Omega$. Then*

$$W_L = \{v \in L : v = \iota^*(\xi) \cap L \textit{ for } \xi \in \Omega^*\}$$

is a nonempty open subset of L.

Proof. W_L is nonempty since Ω^* is nonempty. So let $v \in W_L$, and choose $\xi \in \Omega^*$ so that $\iota^*(\xi) \cap L = v$. We need to show that an open interval $I \subset L$ containing v is contained in W_L.

If L is spanned by $x, y \in \iota(\Omega)$, then Lemma 5.4.14 implies that we can find a one-parameter subgroup $g_t \in G$ such that $\zeta(g_t)$ fixes x and y, and acts nontrivially on L. Since Ω^* is open, we can find $\varepsilon > 0$ so that $g_t \cdot \xi \in \Omega^*$ for $t \in (-\varepsilon, \varepsilon)$. Since x and y are in $\iota(\Omega)$, $\iota^*(\xi)$ is transverse to both x and y, so we have $v \neq x$, $v \neq y$. Then as t varies from $-\varepsilon$ to ε,

$$\iota^*(g_t \cdot \xi) \cap L = \zeta(g_t) \cdot v$$

gives an open interval in W_L containing v.

A similar argument using Lemma 5.4.15 shows that the claim also holds if L is tangent to $\iota(\Omega)$. $\square$

We can now prove a slightly weaker version of Proposition 5.4.11, which we will then use to show the stronger version.

Lemma 5.4.17. *Let Ω_1, Ω_2 be proper domains in G/P^+, with $\overline{\Omega_1} \subset \Omega_2$, and let $K \subset \Omega_1$ be compact. There exists a constant $\lambda > 1$ such that for all $x, y \in K$,*

$$C_{\Omega_1}(x,y) \geq \lambda \cdot C_{\Omega_2}(x,y).$$

Proof. Since K is compact, it suffices to show that for fixed $x \in \Omega_1$, the ratio

$$\frac{C_{\Omega_1}(x,y)}{C_{\Omega_2}(x,y)}$$

is bounded below by some $\lambda > 1$ as y varies in $K - \{x\}$.

Let $y \in K - \{x\}$, and let L_{xy} denote the projective line spanned by $\iota(x)$ and $\iota(y)$. Choose $\xi, \eta \in \overline{\Omega_2^*}$ so that

$$C_{\Omega_2}(x,y) = \log |[\iota^*(\xi), \iota^*(\eta); \iota(x), \iota(y)]|.$$

That is, if $v = \iota^*(\xi) \cap L_{xy}$, $w = \iota^*(\eta) \cap L_{xy}$, we have

$$C_{\Omega_2}(x,y) = \log |[v, w; \iota(x), \iota(y)]| = \log \frac{|v - \iota(y)| \cdot |w - \iota(x)|}{|v - \iota(x)| \cdot |w - \iota(y)|},$$

where the distances are measured in any identification of L_{xy} with $\mathbb{RP}^1 = \mathbb{R} \cup \{\infty\}$.

We can choose an identification of L_{xy} with $\mathbb{R} \cup \{\infty\}$ so that either $v < \iota(x) < \iota(y) < w$ or $v < \iota(x) < w < \iota(y)$. In either case, for any $v' \in (v, \iota(x)) \subset L_{xy}$, we have

$$\log |[v', w; \iota(x), \iota(y)]| > \log |[v, w; \iota(x), \iota(y)]|.$$

We know that $\overline{\Omega_2^*} \subset \Omega_1^*$, so ξ, η lie in Ω_1^*. Then Lemma 5.4.16 implies that there exists $\xi' \in \Omega_1^*$ so that $v' = \iota^*(\xi') \cap L_{xy}$ lies in the interval $(v, \iota(x)) \subset L_{xy}$. See Figure 5.8.

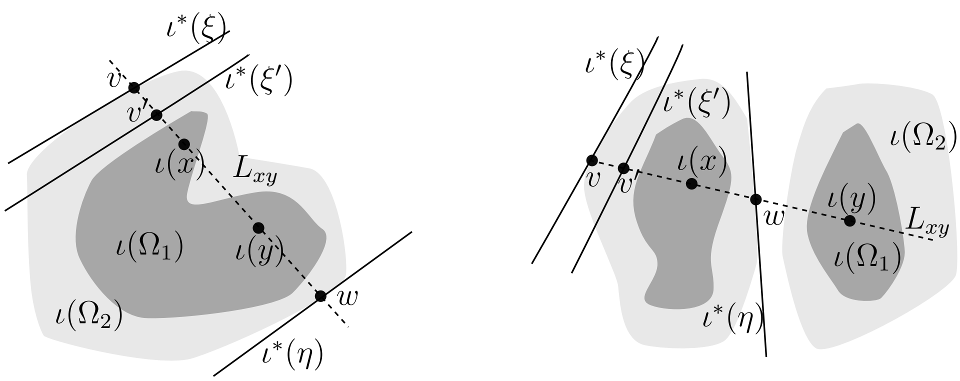

Figure 5.8: We can always find $\xi' \in \Omega_1^*$ close to ξ so that the absolute value of the cross-ratio $[\iota^*(\xi), \iota^*(\nu); \iota(x), \iota(y)]$ increases when we replace ξ with ξ'. In particular this is possible even when the sets $\iota(\Omega_1)$, $\iota(\Omega_2)$ fail to be convex (left) or even connected (right).

Then, we have

$$\begin{aligned} C_{\Omega_1}(x, y) &\geq \log |[\iota^*(\xi'), \iota^*(\eta); \iota(x), \iota(y)]| \\ &= \log |[v', w; \iota(x), \iota(y)] \\ &> \log |[v, w; \iota(x), \iota(y)] \\ &= C_{\Omega_2}(x, y). \end{aligned}$$

This shows that $\frac{C_{\Omega_1}(x,y)}{C_{\Omega_2}(x,y)} > 1$ for all $y \in K - \{x\}$. We still need to find some uniform $\lambda > 1$ so that $\frac{C_{\Omega_1}(x,y)}{C_{\Omega_2}(x,y)} \geq \lambda$ for all $y \in K - \{x\}$. To see this, suppose for the sake of a contradiction that for a sequence $y_n \in K - \{x\}$, we have

$$\frac{C_{\Omega_1}(x, y_n)}{C_{\Omega_2}(x, y_n)} \to 1. \tag{5.10}$$

Since K is compact, y_n must converge to x. Up to subsequence, the sequence of projective lines L_n spanned by $\iota(x)$ and $\iota(y_n)$ converges to a line L tangent to $\iota(G/P^+)$ at $\iota(x)$.

For each y_n, choose ξ_n, $\eta_n \in \overline{\Omega_2^*}$ so that

$$C_{\Omega_2}(x, y_n) = \log |[\iota^*(\xi_n), \iota^*(\eta_n); \iota(x), \iota(y_n)]|.$$

Let $v_n = \iota^*(\xi_n) \cap L_n$, $w_n = \iota^*(\eta_n) \cap L_n$. Then up to subsequence ξ_n converges to $\xi \in \overline{\Omega_2^*}$, η_n converges to $\eta \in \overline{\Omega_2^*}$, and v_n, w_n respectively converge to $v = \iota^*(\xi) \cap L$, $w = \iota^*(\eta) \cap L$.

Since x is in Ω_2, $\iota^*(\xi)$ and $\iota^*(\eta)$ are both transverse to $\iota(x)$—so in particular $x \neq w$ and $x \neq v$ (although a priori we could have $v = w$).

Since $\xi \in \overline{\Omega_2^*} \subset \Omega_1^*$, Lemma 5.4.16 implies that there exist $\xi', \eta' \in \Omega_1^*$ so that for some identification of L with $\mathbb{R} \cup \{\infty\}$, we have

$$v < \iota^*(\xi') \cap L < \iota(x) < \iota^*(\eta') \cap L < w.$$

Note that this is possible even if $v = w$, because then we can just identify both v and w with ∞. Let $v_n' = \iota^*(\xi') \cap L_n$, and let $w_n' = \iota^*(\eta') \cap L_n$. Respectively, v_n' and w_n' converge to $v' = \iota^*(\xi') \cap L$ and $w' = \iota^*(\eta') \cap L$.

This means that the cross-ratios $[v_n, v_n'; w_n', w_n]$ converge to $[v, v'; w', w] \in \mathbb{R} - \{1\}$, and in particular are bounded away from both ∞ and 1 for all n.

We choose identifications of L_n with $\mathbb{R} \cup \{\infty\}$ so that $v_n < v_n' < \iota(x) < w_n' < w_n$. Since y_n converges to x, for all sufficiently large n, we have $v_n' < \iota(y_n) < w_n'$. Then, Proposition 5.4.13 implies that for all n, we have

$$\log |[v_n', \iota(x), \iota(y_n), w_n']| \geq \lambda \cdot \log |[v_n, \iota(x), \iota(y_n), w_n]|$$

for some $\lambda > 1$ independent of n. But then since

$$C_{\Omega_1}(x, y_n) \geq \log |[\iota^*(\xi'), \iota^*(\eta'); \iota(x), \iota(y_n)]|,$$

we have $C_{\Omega_1}(x, y_n)/C_{\Omega_2}(x, y_n) \geq \lambda$ for all n, contradicting (5.10) above. $\square$

Proof of Proposition 5.4.11. We fix an open set $\Omega_{1.5}$ such that $\overline{\Omega_1} \subset \Omega_{1.5}$ and $\overline{\Omega_{1.5}} \subset \Omega_2$. Since $C_{\Omega_1}(x, y) \geq C_{\Omega_{1.5}}(x, y)$ for all $x, y \in \Omega_1$, we just need to see that there is some $\lambda > 1$ so that

$$\frac{C_{\Omega_{1.5}}(x, y)}{C_{\Omega_2}(x, y)} \geq \lambda$$

for all $x, y \in \overline{\Omega_1}$. This follows from Lemma 5.4.17. $\square$

5.4.3 Contracting paths are P^+-divergent

Proposition 5.4.18. *Let $\mathcal{G}$ be a Γ-graph for a group $\Gamma \subset G$, and let $\{U_v\}$ be a $\mathcal{G}$-compatible system of open sets of G/P^+ with each U_v a proper domain.*

If α_n is a contracting $\mathcal{G}$-path, then the sequence

$$\gamma_n = \alpha_1 \cdots \alpha_n$$

is P^+-divergent with unique limit point ξ, where $\{\xi\} = \bigcap_{n=1}^{\infty} \gamma_n U_{n+1}$.

Proof. Consider the sequence of open sets

$$\gamma_n \cdot U_{v_{n+1}}.$$

Up to subsequence, $U_{v_{n+1}}$ is a fixed open set $U \subset G/P^+$. By assumption $\{\alpha_n\}$ is a contracting path, so $\gamma_n \cdot U_{v_{n+1}}$ converges to a singleton $\{\xi\}$. So, we apply Proposition 2.3.7. □

5.5 A weaker criterion for EGF representations

We have now developed enough tools to be able to prove our weaker characterization of EGF representations, which we restate here:

Proposition 5.1.2. *Let $\rho : \Gamma \to G$ be a representation of a relatively hyperbolic group, and let $\Lambda \subset G/P$ be a closed $\rho(\Gamma)$-invariant set, where $P \subset G$ is a symmetric parabolic subgroup. Suppose that $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ is a surjective ρ-equivariant antipodal map.*

Then ρ is an EGF representation if and only if for each $z \in \partial(\Gamma, \mathcal{H})$, there exists $C_z \subset G/P$, with $C_z \subset \mathrm{Opp}(\phi^{-1}(z))$ and $\Lambda - \phi^{-1}(z) \subset C_z$, such that:

1. *For any sequence $\gamma_n \in \Gamma$ limiting conically to z (with $\gamma_n^{-1} \to z_-$), any compact $K \subset C_{z_-}$, and any neighborhood U containing $\phi^{-1}(z)$, we have $\rho(\gamma_n) \cdot K \subset U$ for all sufficiently large n.*

2. *For any parabolic point p, any compact $K \subset C_p$, and any open set U containing $\phi^{-1}(p)$, for all but finitely many $\gamma \in \Gamma_p$, we have $\rho(\gamma) \cdot K \subset U$.*

Proof. The "only if" part is immediate from the definition of EGF representations and Proposition 5.1.1, so assume that we have a map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ satisfying the conditions given.

Let γ_n be an infinite sequence in Γ, and suppose that in the compactification $\overline{\Gamma} = \Gamma \sqcup \partial(\Gamma, \mathcal{H})$, we have $\gamma_n \to z_+$ and $\gamma_n^{-1} \to z_-$ for $z_\pm \in \partial(\Gamma, \mathcal{H})$. It suffices to show that for any subsequence of γ_n, any compact $K \subset C_{z_-}$, and any open U containing $\phi^{-1}(z_+)$, we can find a further subsequence such that $\gamma_n K \subset U$. So, we can freely extract subsequences throughout this proof.

We consider two cases:

Case 1: γ_n is unbounded in the coned-off Cayley graph $\mathrm{Cay}(\Gamma, S, \mathcal{P})$. Our assumptions imply that we can construct a relative quasigeodesic automaton $\mathcal{G}$, together with a $\mathcal{G}$-compatible system of open sets $\{U_v\}$ for the action of Γ on G/P satisfying the conditions of Proposition 5.3.1. Since the map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ is antipodal, and antipodality is an open condition, the first condition of Proposition 5.3.1 means we can ensure that each U_v is a proper domain in G/P.

By Lemma 5.2.7, we may assume that for a bounded sequence $b_n \in \Gamma$, $\gamma_n b_n$ is the endpoint of a finite $\mathcal{G}$-path $\{\alpha_m^{(n)}\}_{m=1}^{M_n}$. Up to subsequence b_n is a constant b, independent of n.

Let $\{v_m^n\}$ be the vertex path associated to $\{\alpha_m^{(n)}\}$. Up to subsequence $v_{M_n+1}^n$ is a fixed vertex v, and v_1^n is a fixed vertex v'. Let $U_{v'}^\varepsilon$ be an ε-neighborhood of $U_{v'}$.

Since the sequence M_n is unbounded, Corollary 5.4.12 implies that the diameter of

$$\rho(\gamma_n b) \cdot U_v = \rho(\alpha_1^{(n)}) \cdots \rho(\alpha_{M_n}^{(n)}) U_v$$

with respect to the metric $C_{U^\varepsilon_{v'}}$ tends to zero, exponentially in n. Since this sequence of sets lies in the compact set $\overline{U_{v'}} \subset U^\varepsilon_{v'}$, up to subsequence it must converge to a singleton $\{\xi\}$ in G/P. In fact ξ must lie in Λ, because Λ is compact and ξ is the limit of a sequence of points in the sequence of nonempty closed sets $(\rho(\gamma_n b) \cdot \overline{U_v}) \cap \Lambda$. Then, since $\rho(\gamma_n) \cdot \rho(b) U_v$ converges to $\{\xi\}$, Proposition 2.3.7 implies that $\rho(\gamma_n)$ is P-divergent with unique P-limit ξ. The exact same agument applied to γ_n^{-1} implies that $\rho(\gamma_n^{-1})$ is also P-divergent, with unique P-limit $\xi_- \in \Lambda$.

Then, Lemma 2.3.8 implies that $\rho(\gamma_n)$ converges to ξ uniformly on compacts in $\mathrm{Opp}(\xi_-)$. But, since γ_n converges to z_+ uniformly on compacts in $\partial(\Gamma, \mathcal{H}) - \{z_-\}$, and ϕ is equivariant and surjective, the only possibility is $\xi \in \phi^{-1}(z_+)$ and $\xi_- \in \phi^{-1}(z_-)$. Then we are done, since by assumption C_{z_-} is contained in $\mathrm{Opp}(\xi_-)$.

Case 2: γ_n is bounded in $\mathrm{Cay}(\Gamma, S, \mathcal{P})$. We can write γ_n as an alternating product

$$\gamma_n = g_1^{(n)} h_1^{(n)} \cdots g_k^{(n)} h_k^{(n)} g_{k+1}^{(n)},$$

where $g_i^{(n)}$ is bounded in Γ, and $h_i^{(n)}$ lies in $\Gamma_{p_i^n}$ for a parabolic point $p_i^n \in \Pi$. Without loss of generality, the $h_i^{(n)}$ are unbounded in Γ as $n \to \infty$. Up to subsequence we can assume that $g_i^{(n)} = g_i$ and $p_i^n = p_i$ (independent of n). Since Π contains exactly one representative of each parabolic orbit, we can also assume that $g_{i+1} p_{i+1} \neq p_i$ for any i.

We claim that γ_n converges to $z_+ = g_1 p_1$, γ_n^{-1} converges to $z_- = g_{k+1}^{-1} p_k$, and for any compact $K \subset C_{z_-}$ and open U containing $\phi^{-1}(z_+)$, for large n, we have $\gamma_n \cdot K \subset U$.

Fix such a compact K and open U. We will prove the claim by inducting on k. When $k = 1$, then $p = p_1 = p_k$, and $\gamma_n = g_1 h_n g_2$ for $h_n \in \Gamma_p$ and $g_1, g_2 \in \Gamma$ fixed. The distance between $h_n g_2$ and h_n is bounded in any word metric on Γ, so $h_n g_2$ converges to p in $\overline{\Gamma}$ and $g_1 h_n g_2$ converges to $g_1 p = z_+$. We also know that $K \subset C_{z_-} = C_{g_2^{-1}p}$, so $h_n g_2 K$ eventually lies in a small neighborhood of $\phi^{-1}(p)$ by condition (2) of our hypotheses. Then $g_1 h_n g_2 K$ lies in any small neighborhood of $\phi^{-1}(g_1 p) = \phi^{-1}(z_+)$.

When $k > 1$, we consider the sequence

$$\gamma'_n = g_2 h_2^{(n)} \cdots g_k h_k^{(n)} g_{k+1}.$$

Inductively we can assume that for large n, $\gamma'_n \to g_2 p_2$ and $\rho(\gamma'_n) \cdot K$ is a subset of an arbitrarily small neighborhood of $\phi^{-1}(g_2 p_2)$. Then since $p_1 \neq g_2 p_2$, for large enough n, $\rho(\gamma'_n) \cdot K$ is a compact subset of C_{p_1}. So our hypotheses imply that for large n,

$$\rho(\gamma_n) \cdot K = \rho(g_1 h_1^{(n)}) \rho(\gamma'_n) \cdot K \subset U.$$

$\square$

5.6 Relative stability

In this section we prove the main *relative stability property* for EGF representations (Theorem 5.6.2).

5.6.1 Deformations of EGF representations

In general, the set of EGF representations is not an open subset of $\mathrm{Hom}(\Gamma, G)$. However, it is relatively open in a subspace of $\mathrm{Hom}(\Gamma, G)$ where we restrict the deformations of the peripheral subgroups appropriately. Roughly speaking, we want to consider subspaces $\mathcal{W} \subset \mathrm{Hom}(\Gamma, G)$ where the *dynamical* behavior of the peripheral subgroups changes continuously under deformation. That is, if ρ_t is a small deformation of a representation ρ_0, where $\rho_0(\Gamma_p)$ attracts points towards Λ_p at a particular "speed," then we want $\rho_t(\Gamma_p)$ to attract points towards a small deformation of Λ_p at a similar "speed."

The precise condition is the following:

Definition 5.6.1. Let $\rho_0 : \Gamma \to G$ be an EGF representation with boundary extension $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$, and let $\mathcal{W} \subset \mathrm{Hom}(\Gamma, G)$ contain ρ_0.

We say that $\mathcal{W}$ is *peripherally stable at* (ρ_0, ϕ) if for every $p \in \Pi$, every open set U containing $\phi^{-1}(p)$, every compact set $K \subset C_p$, and every cofinite set $T \subset \Gamma_p$ such that

$$\rho_0(T) \cdot K \subset U,$$

there is an open set $\mathcal{W}' \subset \mathcal{W}$ containing ρ_0, such that for every $\rho' \in \mathcal{W}'$, we have

$$\rho'(T) \cdot K \subset U.$$

The following is the main result of this section.

Theorem 5.6.2. *Let $\rho : \Gamma \to G$ be EGF with respect to P, let $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ be a boundary extension, and let $\mathcal{W} \subseteq \mathrm{Hom}(\Gamma, G)$ be peripherally stable at (ρ, ϕ). For any compact subset Z of $\partial(\Gamma, \mathcal{H})$ and any open set $V \subset G/P$ containing $\phi^{-1}(Z)$, there is an open subset $\mathcal{W}' \subset \mathcal{W}$ containing ρ such that each $\rho' \in \mathcal{W}'$ is EGF with respect to P, and has an EGF boundary extension ϕ' satisfying $\phi'^{-1}(Z) \subset V$.*

Remark 5.6.3. In [Bow98], Bowditch explored the deformation spaces of geometrically finite groups $\Gamma \subset \mathrm{PO}(d,1)$, and gave an explicit discription of semialgebraic subspaces of $\mathrm{Hom}(\Gamma, \mathrm{PO}(d,1))$ in which small deformations of Γ are still geometrically finite.

Bowditch's deformation spaces are peripherally stable, so it seems desirable to find a general algebraic description of peripherally stable subspaces.

Even in $\mathrm{PO}(d,1)$, the question is subtle, however. Bowditch also gives examples of geometrically finite representations $\rho : \Gamma \to \mathrm{PO}(d,1)$ (for $d \geq 4$) and deformations ρ_t of ρ in $\mathrm{Hom}(\Gamma, \mathrm{PO}(d,1))$ such that the restriction of ρ_t to each cusp group in Γ is discrete, faithful, and parabolic, but ρ_t is not even discrete; further examples exist where the deformation is discrete, but not geometrically finite.

The simplest examples of peripherally stable deformations are *cusp-preserving* deformations.

Definition 5.6.4. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and let $\rho : \Gamma \to G$ be a representation. The space of *cusp-preserving* representations

$$\mathrm{Hom}_{\mathrm{cp}}(\Gamma, G, \mathcal{H}, \rho)$$

is the set of representations $\rho' : \Gamma \to G$ such that for each peripheral subgroup $H \in \mathcal{H}$, we have

$$\rho'|_H = g \cdot \rho|_H \cdot g^{-1}$$

for some $g \in G$ (which may depend on H).

It is immediate from the definition that if $\rho = \rho_0$ is EGF, and $\rho_t : \Gamma \to G$ is a continuous family of representations such that $\rho_t(\Gamma_p) = g_t \rho_0(\Gamma_p) g_t^{-1}$ for a path $g_t \in G$ with $g_0 = \mathrm{id}$, then the set $\{\rho_t\}$ is peripherally stable at ρ_0. So one consequence of Theorem 5.6.2 is the following:

Corollary 5.6.5. *Let $\rho : \Gamma \to G$ be an EGF representation. Then there is a neighborhood of ρ in*

$$\mathrm{Hom}_{\mathrm{cp}}(\Gamma, G, \mathcal{H}, \rho)$$

consisting of EGF representations.

There are also often examples of peripherally stable deformations which are *not* cusp-preserving.

Example 5.6.6. Let $B \in \mathrm{SL}(d, \mathbb{R})$ be a d-dimensional Jordan block with eigenvalue 1 and eigenvector v, and let $A \in \mathrm{SL}(d+2, \mathbb{R})$ be the block matrix $\begin{pmatrix} B & & \\ & 1 & \\ & & 1 \end{pmatrix}$.

Although $[v]$ is not quite an attracting fixed point of A, it is still an "attracting subspace" in the sense that if K is any compact subset of $\mathbb{RP}^{d+1}$ which does not intersect a fixed hyperplane of A, then $A^n \cdot K$ converges to $\{[v]\}$ (see Proposition 2.1.17). Via a ping-pong argument, one can use this "attracting" behavior to show that for some $k \geq 1$ and some $M \in \mathrm{SL}(d+2, \mathbb{R})$, the group Γ generated by $\alpha = A^k$ and $\beta = MA^kM^{-1}$ is a discrete free group with free generators α, β. The group Γ is hyperbolic relative to the subgroups $\langle \alpha \rangle$, $\langle \beta \rangle$, and the inclusion $\Gamma \hookrightarrow \mathrm{SL}(d+2, \mathbb{R})$ is EGF with respect to $P_{1,d+1}$ (the stabilizer of a line in a hyperplane in $\mathbb{R}^{d+2}$).

Here, there are peripherally stable deformations of Γ which change the Jordan block decomposition of A. For instance, consider a continuous path $A_t : [0, 1] \to$

$\mathrm{SL}(d+2,\mathbb{R})$ given by $A_t = \begin{pmatrix} B_t & & \\ & 1 & \\ & & 1 \end{pmatrix}$, where $B_0 = B$ and B_t is a *diagonalizable* matrix in $\mathrm{SL}(d,\mathbb{R})$. For small values of t, the group Γ_t generated by $\alpha_t = A_t^k$ and β is still discrete and freely generated by α_t and β—since the "attracting" fixed points of A_t deform continuously with t, the same exact ping-pong setup works for all small $t \geq 0$. And indeed the path in $\mathrm{Hom}(\Gamma, \mathrm{SL}(d+2,\mathbb{R}))$ determined by the path A_t is a peripherally stable subspace.

On the other hand, consider the path $A_t' = \begin{pmatrix} B & & \\ & e^t & \\ & & e^{-t} \end{pmatrix}$, and let $\alpha_t' = A_t'^k$. In this case the corresponding subspace of $\mathrm{Hom}(\Gamma, \mathrm{SL}(d+2,\mathbb{R}))$ is *not* peripherally stable: while the group generated by α_t' is still discrete, the attracting fixed points of A_t' do *not* deform continuously in t. So, there is no way to use the ping-pong setup for Γ to ensure that $\Gamma_t' = \langle \alpha_t', \beta \rangle$ is a discrete group.

Example 5.6.7. Here is a somewhat more interesting example of a *non*-peripherally stable deformation. Let M be a finite-volume noncompact hyperbolic 3-manifold, with holonomy representation $\rho : \pi_1 M \to \mathrm{PSL}(2,\mathbb{C})$ (so there is an identification $M = \mathbb{H}^3/\rho(\pi_1 M)$). Then $\pi_1 M$ is hyperbolic relative to the collection $\mathcal{C}$ of conjugates of cusp groups (each of which is isomorphic to $\mathbb{Z}^2$), and the representation ρ is geometrically finite (in particular, EGF).

In this case, for *any* sufficiently small nontrivial deformation ρ' of ρ in the character variety $\mathrm{Hom}(\pi_1 M, \mathrm{PSL}(2,\mathbb{C}))/\mathrm{PSL}(2,\mathbb{C})$, the restriction of ρ' to some cusp group $C \in \mathcal{C}$ either fails to be discrete or has infinite kernel. So $\mathrm{Hom}(\pi_1 M, \mathrm{PSL}(2,\mathbb{C}))$ is *not* peripherally stable, because any sufficiently small deformation of ρ inside of a peripherally stable subspace must have discrete image and finite kernel on each $C \in \mathcal{C}$.

This is true despite the fact that arbitrarily small deformations of ρ are associated to complete hyperbolic structures on Dehn fillings of M.

The main ingredient in the proof of Theorem 5.6.2 is the relative quasigeodesic automaton $\mathcal{G}$ and the associated $\mathcal{G}$-compatible system of open sets $\{U_v\}$ we constructed in Proposition 5.3.1. The following proposition is immediate from the definition of peripheral stability:

Proposition 5.6.8. *Let $\rho : \Gamma \to G$ be an EGF representation with boundary extension ϕ, and let $\mathcal{W} \subset \mathrm{Hom}(\Gamma, G)$ be a subspace which is peripherally stable at (ρ, ϕ).*

If $\mathcal{G}$ is a relative quasigeodesic automaton for Γ, and $\{U_v\}$ is a $\mathcal{G}$-compatible system of open subsets of G/P for $\rho(\Gamma)$, then there is an open subset $\mathcal{W}' \subset \mathcal{W}$ containing ρ such that for every $\rho' \in \mathcal{W}'$, $\{U_v\}$ is also a $\mathcal{G}$-compatible system of open sets for $\rho'(\Gamma)$.

Theorem 5.6.2 then follows from a kind of converse to Proposition 5.3.1: we will show that we can reconstruct a map extending the convergence dynamics of Γ from the $\mathcal{G}$-compatible system $\{U_v\}$.

5.6.2 An equivariant map on conical limit points

For the rest of this section, we let $\rho : \Gamma \to G$ be a representation which is EGF with respect to a symmetric parabolic subgroup $P \subset G$. We let $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$ be a boundary extension for ρ, and assume that $\mathcal{W} \subset \mathrm{Hom}(\Gamma, G)$ is peripherally stable at (ρ, ϕ). We also let Z be a compact subset of $\partial(\Gamma, \mathcal{H})$, and let $V \subset G/P$ be an open subset containing $\phi^{-1}(Z)$. We again fix a finite subset $\Pi \subset \partial_{\mathrm{par}}(\Gamma, \mathcal{H})$, containing one point from every Γ-orbit in $\partial_{\mathrm{par}}(\Gamma, \mathcal{H})$.

Using Proposition 5.3.1, we can find a relative quasigeodesic automaton $\mathcal{G}$ and $\mathcal{G}$-compatible system $\{U_v\}$ of open subsets of G/P for $\rho(\Gamma)$. Using Proposition 5.3.14, we can ensure that for any $z \in Z$, there is a $\mathcal{G}$-path $\{\alpha_n\}$ limiting to z (with vertex path $\{v_n\}$) so that U_{v_1} is contained in V.

Antipodality of the map ϕ implies that for each $z \in \partial(\Gamma, \mathcal{H})$, each fiber $\phi^{-1}(z)$ is a closed subset of some affine chart in G/P. So, we can also assume that U_v is a proper domain for each vertex v of $\mathcal{G}$. In fact, by way of the following lemma, we can assume even more:

Lemma 5.6.9. *Let ρ be an EGF representation with boundary map $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$.*

For any $\delta > 0$, we can find a relative quasigeodesic automaton $\mathcal{G}$ with $\mathcal{G}$-compatible system $\{U_v\}$ of open sets in G/P as in Proposition 5.3.1, so that for any $x, y \in \partial(\Gamma, \mathcal{H})$ with $d(x, y) > \delta$, if $\phi^{-1}(x) \subset U_v$ and $\phi^{-1}(y) \subset U_w$, then $\overline{U_v}$ and $\overline{U_w}$ are opposite.

Proof. We choose $\varepsilon > 0$ so that if $d(v, w) > \delta/2$ for $v, w \in \partial(\Gamma, \mathcal{H})$, then the closed ε-neighborhoods of

$$\phi^{-1}(v), \qquad \phi^{-1}(w)$$

are opposite. This is possible for a fixed pair $v, w \in \partial(\Gamma, \mathcal{H})$ since antipodality is an open condition, and $\phi^{-1}(v)$, $\phi^{-1}(w)$ are opposite compact sets. Then we can pick a uniform ε for all pairs since the the subset $\{(u, v) \in (\partial(\Gamma, \mathcal{H}))^2 : d(u, v) > \delta/2\}$ is compact.

Consider $\mathcal{G}$-compatible systems of open subsets $\{U_v\}$ and $\{W_v\}$ for the action of Γ on G/P and $\partial(\Gamma, \mathcal{H})$, coming from Proposition 5.3.1. We can ensure that for each vertex a, the diameter of W_a is at most $\delta/4$, and $U_a \subset N(\phi^{-1}(w), \varepsilon)$ for some $w \in W_a$.

If $x, y \in \partial(\Gamma, \mathcal{H})$ satisfy $d(x, y) > \delta$, and $x \in W_a$, $y \in W_b$, we have

$$d(v, w) > \delta/2$$

for all $v \in W_a$, $w \in W_b$.

Then, if $\phi^{-1}(x) \subset U_a$ and $\phi^{-1}(y) \subset U_b$, we have

$$U_a \subset N(\phi^{-1}(v), \varepsilon), \qquad U_b \subset N(\phi^{-1}(w), \varepsilon)$$

for $v \in W_a$, $w \in W_b$ with $d(v, w) > \delta/2$. By our choice of ε, the closures of $N(\phi^{-1}(w), \varepsilon)$ and $N(\phi^{-1}(v), \varepsilon)$ are opposite. □

Using cocompactness of the action of Γ on the space of distinct pairs in $\partial(\Gamma, \mathcal{H})$, we know that there exists some fixed $\delta > 0$ such that for any distinct $z_1, z_2 \in \partial(\Gamma, \mathcal{H})$, we can find some $\gamma \in \Gamma$ such that $d(\gamma z_1, \gamma z_2) > \delta$. Then, in light of Lemma 5.6.9, we can make the following assumption:

Assumption 5.6.10. For any $z_1, z_2 \in \partial(\Gamma, \mathcal{H})$ satisfying $d(z_1, z_2) > \delta$, if $\phi^{-1}(z_1) \subset U_v$ and $\phi^{-1}(z_2) \subset U_w$ for v, w vertices of $\mathcal{G}$, then $\overline{U_v}$ and $\overline{U_w}$ are opposite.

With our relative quasigeodesic automaton $\mathcal{G}$ and compatible system of open sets $\{U_v\}$ fixed, we now choose an open subset $\mathcal{W}' \subset \mathcal{W}$ so that for any $\rho' \in \mathcal{W}'$,

$\{U_v\}$ is also a $\mathcal{G}$-compatible system for the action of $\rho'(\Gamma)$ on G/P. Our main goal for the rest of this section is to show that any $\rho' \in \mathcal{W}'$ is an EGF representation. So, we fix some $\rho' \in \mathcal{W}'$.

Let $\mathrm{Path}(\mathcal{G})$ denote the set of infinite $\mathcal{G}$-paths. Proposition 5.4.3 implies that every path in $\mathrm{Path}(\mathcal{G})$ is contracting for the ρ'-action, so we have a map

$$\psi_{\rho'} : \mathrm{Path}(\mathcal{G}) \to G/P,$$

where the path $\{\alpha_n\}$ maps to the unique element of

$$\bigcap_{n=1}^{\infty} \rho'(\alpha_1) \cdots \rho'(\alpha_n) U_{v_{n+1}}.$$

Lemma 5.6.11. *The map $\psi_{\rho'} : \mathrm{Path}(\mathcal{G}) \to G/P$ induces an equivariant map*

$$\psi_{\rho'} : \partial_{\mathrm{con}}\Gamma \to G/P.$$

Proof. We first need to see that $\psi_{\rho'}$ is well-defined, i.e. that if z is a conical limit point and $\{\alpha_n\}$, $\{\beta_n\}$ are $\mathcal{G}$-paths limiting to z, then $\psi_{\rho'}(\{\alpha_n\}) = \psi_{\rho'}(\{\beta_n\})$.

Let

$$\gamma_n = \alpha_1 \cdots \alpha_n, \qquad \eta_m = \beta_1 \cdots \beta_m.$$

We can use Proposition 5.2.11 to see that γ_n and η_m lie within bounded Hausdorff distance of a geodesic in $\mathrm{Cay}(\Gamma, S, \mathcal{P})$ limiting to z, so there is a fixed D so that for m, n tending to infinity,

$$d(\gamma_n, \eta_m) < D$$

in the Cayley graph of Γ. Proposition 5.4.18 implies that $\rho'(\gamma_n)$ and $\rho'(\eta_m)$ are both P-divergent sequences and each have a unique P-limit point in G/P, given by $\psi_{\rho'}(\{\alpha_n\})$, $\psi_{\rho'}(\{\beta_m\})$, respectively. Then, Lemma 4.23 in [KLP17] implies that because $\rho'(\gamma_n) = \rho'(\eta_m)g_n$ for a bounded sequence $g_n \in G$, the P-limit points of $\rho'(\gamma_n)$ and $\rho'(\eta_n)$ must agree and therefore $\psi_{\rho'}(\{\alpha_n\}) = \psi_{\rho'}(\{\beta_m\})$.

Next we observe that $\psi_{\rho'}$ is equivariant. Fix a finite generating set S for Γ. It suffices to show that $\psi_{\rho'}(s \cdot z) = \rho'(s) \cdot \psi_{\rho'}(z)$ for all $s \in S$.

Let $\{\alpha_n\}$ be a $\mathcal{G}$-path limiting to some $z \in \partial_{\text{con}}\Gamma$, and consider the sequence

$$\gamma'_n = s\alpha_1 \cdots \alpha_n.$$

Again, Proposition 5.2.11 implies that γ'_n lies bounded Hausdorff distance from a geodesic in $\text{Cay}(\Gamma, S, \mathcal{P})$, which must limit to $s \cdot z$. So if we fix a $\mathcal{G}$-path β_n limiting to $s \cdot z$, the same argument as above shows that $\psi_{\rho'}(\{\beta_n\}) = \rho'(s) \cdot \psi_{\rho'}(\{\alpha_n\})$. $\square$

It will turn out that $\psi_{\rho'}$ is also both continuous and injective. However, we do not prove this directly.

5.6.3 Extending $\psi_{\rho'}$ to parabolic points

We want to extend the map $\psi_{\rho'} : \partial_{\text{con}}\Gamma \to G/P$ to the entire Bowditch boundary $\partial(\Gamma, \mathcal{H})$. To do so, we need to view $\psi_{\rho'}$ as a map to the set of closed subsets of G/P.

The first step is to define $\psi_{\rho'}$ on the finite set $\Pi \subset \partial_{\text{par}}\Gamma$. For any vertex v in $\mathcal{G}$, we consider the set

$$B_v = \bigcup_{(v,y) \text{ edge in } \mathcal{G}} U_y.$$

Then, for each $p \in \Pi$, we pick a parabolic vertex w so that $p_w = p$. We define Λ'_p to be the closure of the set of accumulation points of sequences of the form $\rho'(\gamma_n) \cdot x$, for $x \in B_w$ and γ_n distinct elements of Γ_p. Part (3) of Proposition 5.3.1 guarantees that $B_w \subset C_p$, and $\mathcal{G}$-compatiblity of the system $\{U_v\}$ for the $\rho'(\Gamma)$-action on G/P implies that $\Lambda'_p \subset U_w$. By construction, Λ'_p is invariant under the action of $\rho(\Gamma_p)$.

Next, given a parabolic point $q \in \partial_{\text{par}}\Gamma$, we write $q = g \cdot p$ for $p \in \Pi$, and then define

$$\psi_{\rho'}(q) := \rho'(g)\Lambda'_p.$$

Since Λ'_p is Γ_p-invariant and Γ_p is exactly the stabilizer of p, this does not depend on the choice of coset representative in $g\Gamma_p$. Moreover $\psi_{\rho'}$ is still ρ'-equivariant.

In addition, if v is any parabolic vertex with parabolic point $p_v = g \cdot p$ for $p \in \Pi$, part (2) of Proposition 5.3.1 ensures that $B_v = B_w$ for any parabolic vertex w with $p_w = p$. So, $\rho'(g) \cdot \Lambda'_p$ is exactly the closure of the set of accumulation points of the form $\rho'(g\gamma_n) \cdot x$ for sequences $\gamma_n \in \Gamma_p$ and $x \in B_v$. Then $\mathcal{G}$-compatibility implies that $\psi_{\rho'}(p_v) = \rho(g)\Lambda'_p$ is a subset of U_v.

Remark 5.6.12. There is a natural topology on the space of closed subsets of G/P, induced by the Hausdorff distance arising from some (any) choice of metric on G/P. We emphasize that the map $\psi_{\rho'}$ is *not* necessarily continuous with respect to this topology.

Ultimately we want to use $\psi_{\rho'}$ to define a map extending the convergence dynamics of Γ, so we will need to also define the sets C'_z for each $z \in \partial(\Gamma, \mathcal{H})$. For now, we only define C'_p for $p \in \Pi$: this will be the set

$$\bigcup_{\gamma \in \Gamma_p} \rho'(\gamma) B_w.$$

We can immediately observe:

Proposition 5.6.13. *C'_p is $\rho'(\Gamma_p)$-invariant. Moreover, for any infinite sequence $\gamma_n \in \Gamma_p$, any compact $K \subset C'_p$, and any open $U \subset G/P$ containing Λ'_p, for sufficiently large n, $\rho'(\gamma_n) \cdot K$ lies in U.*

Proof. Γ_p-invariance follows directly from the definition.

Fix a compact $K \subset C'_p$ and an open $U \subset G/P$ containing Λ'_p. K is contained in finitely many sets $\rho'(\gamma)B_w$ for $\gamma \in \Gamma_p$, so any accumulation point of $\rho'(\gamma_n)x$ for $x \in K$ and $\gamma_n \in \Gamma_p$ lies in Λ'_p. In particular, for sufficiently large n, $\gamma_n x$ lies in U, and since K is compact we can pick n large enough so that $\gamma_n x \in U$ for all $x \in K$. $\square$

We next want to use $\psi_{\rho'}$ to define an antipodal extension from a subset of G/P to $\partial(\Gamma, \mathcal{H})$.

Lemma 5.6.14. *For any $z \in \partial(\Gamma, \mathcal{H})$, if $\{\alpha_n\}$ is a $\mathcal{G}$-path limiting to z with corresponding vertex path $\{v_n\}$, then $\phi^{-1}(z)$ and $\psi_{\rho'}(z)$ are both subsets of U_{v_1}.*

Proof. If z is a conical limit point, then this follows immediately from Proposition 5.2.11 and the definition of $\psi_{\rho'}$. On the other hand, if z is a parabolic point, then $z = \alpha_1 \cdots \alpha_N p_v$, where v is a parabolic vertex at the end of the vertex path

$\{v_n\}$. By part (3) of Proposition 5.3.1, we have $p_v \in W_v$ and thus $\phi^{-1}(p_v) \subset U_v$. By ρ-equivariance of ϕ we have

$$\phi^{-1}(z) = \rho(\alpha_1 \cdots \alpha_N)\phi^{-1}(p_v),$$

so by $\mathcal{G}$-compatibility we have $\phi^{-1}(z) \subset U_{v_1}$. On the other hand, we have constructed $\psi_{\rho'}$ so that $\psi_{\rho'}(p_v) \subset U_v$, so ρ'-equivariance of $\psi_{\rho'}$ and $\mathcal{G}$-compatibility also show that $\psi_{\rho'}(z) \subset U_{v_1}$. □

Lemma 5.6.15. *For any two distinct points z_1, z_2 in $\partial(\Gamma, \mathcal{H})$, the sets*

$$\psi_{\rho'}(z_1), \qquad \psi_{\rho'}(z_2)$$

are opposite (in particular disjoint).

Proof. We know that for any distinct $z_1, z_2 > 0$, we can find $\gamma \in \Gamma$ so that $d(\gamma z_1, \gamma z_2) > \delta$. So, since $\psi_{\rho'}$ is ρ'-equivariant, we just need to show that if $z_1, z_2 \in \partial(\Gamma, \mathcal{H})$ satisfy $d(z_1, z_2) > \delta$, then $\psi_{\rho'}(z_1)$ is opposite to $\psi_{\rho'}(z_2)$.

Let $\{\alpha_n\}$, $\{\beta_n\}$ be $\mathcal{G}$-paths respectively limiting to points $z_1, z_2 \in \partial(\Gamma, \mathcal{H})$ with $d(z_1, z_2) > \delta$, with corresponding vertex paths $\{v_n\}$ and $\{w_n\}$. By Lemma 5.6.14, we must have $\phi^{-1}(z_1) \subset U_{v_1}$ and $\phi^{-1}(z_2) \subset U_{w_2}$, so under Assumption 5.6.10, we know that U_{v_1} and U_{w_1} are opposite. But then we are done since Lemma 5.6.14 also implies that $\psi_{\rho'}(z_1) \subset U_{v_1}$ and $\psi_{\rho'}(z_2) \subset U_{w_1}$. □

5.6.4 The boundary set of the deformed representation

We define our candidate boundary set $\Lambda' \subset G/P$ by

$$\Lambda' = \bigcup_{z \in \partial(\Gamma, \mathcal{H})} \psi_{\rho'}(z).$$

We then have an equivariant map

$$\phi' : \Lambda' \to \partial(\Gamma, \mathcal{H}),$$

where $\phi'(x) = z$ if $x \in \psi_{\rho'}(z)$. Lemma 5.6.15 implies that ϕ' is well-defined and antipodal. It is necessarily both surjective and ρ'-equivariant, and its fibers are either singletons or translates of the sets Λ'_p for $p \in \Pi$.

Lemma 5.6.16. *For any $z \in Z$, we have $\phi'^{-1}(z) \subset V$.*

Proof. Recall that we used Proposition 5.3.14 to construct our automaton so that for any $z \in Z$, there is a $\mathcal{G}$-path limiting to z with vertex path $\{v_n\}$ such that $U_{v_1} \subset V$. Then Lemma 5.6.14 implies $\phi'^{-1}(z) \subset V$. □

Lemma 5.6.17. *Λ' is compact.*

Proof. Fix a sequence $y_n \in \Lambda'$, and let $x_n = \phi'(y_n)$. Since $\partial(\Gamma, \mathcal{H})$ is compact, up to subsequence $x_n \to x$. We want to see that a subsequence of y_n converges to some $y \in \Lambda'$. We consider two possibilities:

Case 1: x is a parabolic point. We can write $x = g \cdot p$, where $p \in \Pi$. Let w be a parabolic vertex with $p_w = p$.

Choose an edge (w, v) in $\mathcal{G}$. Part (3) of Proposition 5.3.1 implies that we can find a compact subset $K \subset W_v$ such that $\Gamma_p \cdot K = \partial(\Gamma, \mathcal{H}) - \{p\}$. So, we can find $\gamma_n \in \Gamma_p$ so that

$$z_n = \gamma_n^{-1} g^{-1} x_n \in K.$$

Then, since $\phi^{-1}(W_v) \subset U_v$, by definition C'_p contains $\phi'^{-1}(K)$, and in particular C'_p contains $\phi'^{-1}(z_n)$ for infinitely many z_n.

Then using Proposition 5.6.13, we know that up to subsequence,

$$\rho'(\gamma_n)\phi'^{-1}(z_n) = \rho'(\gamma_n)\phi'^{-1}(\gamma_n^{-1}g^{-1}x_n)$$

converges to a compact subset of Λ'_p, which means that

$$y_n \in \rho'(g)\rho'(\gamma_n)\phi'^{-1}(\gamma_n^{-1}g^{-1}x_n)$$

subconverges to a point in $\rho'(g)\Lambda'_p$.

Case 2: x is a conical limit point. We want to show that any sequence in $\phi'^{-1}(x_n)$ limits to $\phi'^{-1}(x)$, so fix any $\varepsilon > 0$. Using Corollary 5.4.12, we can choose N so that if $\{\alpha_m\}$ is any $\mathcal{G}$-path limiting to x, with corresponding vertex path $\{v_m\}$, then the diameter of

$$\rho'(\alpha_1 \cdots \alpha_N)U_{v_{N+1}}$$

is less than ε with respect to a metric on U_{v_1}. We fix such a $\mathcal{G}$-path $\{\alpha_m\}$. Then, we use Lemma 5.2.15 to see that for sufficiently large n, there is a $\mathcal{G}$-path $\{\beta_m^n\}$ limiting to x_n with $\beta_i = \alpha_i$ for $i \leq N$. Thus the Hausdorff distance (with respect to $C_{U_{v_1}}$) between $\phi'^{-1}(x_n)$ and $\phi'^{-1}(x)$ is at most ε. Since $\phi'^{-1}(x_n)$ and $\phi'^{-1}(x)$ both lie in the compact set $\rho'(\alpha_1)\overline{U_{v_2}} \subset U_{v_1}$, this proves the claim.

$\square$

Lemma 5.6.18. *ϕ' is continuous and proper.*

Proof. Since Λ' is compact, we just need to show continuity. Fix $y \in \Lambda'$ and a sequence $y_n \in \Lambda'$ approaching y. We want to show that $\phi'(y_n)$ approaches $\phi'(y) = x$.

Suppose otherwise. Since $\partial(\Gamma, \mathcal{H})$ is compact, up to a subsequence $z_n = \phi'(y_n)$ approaches $z \neq x$. Using the equivariance of ϕ', and cocompactness of the Γ-action on distinct pairs in $\partial(\Gamma, \mathcal{H})$, we may assume that $d(x, z) > \delta$. For sufficiently large n, we have $d(x, z_n) > \delta$ as well. Then, as in the proof of Lemma 5.6.15, by Assumption 5.6.10 we know that for any vertices v, w in $\mathcal{G}$ such that U_v contains $\psi_{\rho'}(x)$ and U_w contains $\psi_{\rho'}(z_n)$, the intersection $\overline{U_v} \cap \overline{U_w}$ is empty.

But by definition of ϕ', we have

$$y \in \psi_{\rho'}(x) \subset U_v, \qquad y_n \in \psi_{\rho'}(z_n) \subset U_w$$

for vertices v, w in $\mathcal{G}$. This contradicts the fact that $y_n \to y$. □

5.6.5 Dynamics on the deformation

To complete the proof of Theorem 5.6.2, we just need to show:

Proposition 5.6.19. *ϕ' extends the convergence dynamics of Γ.*

Proof. We will apply Proposition 5.1.2, and show separately that the required dynamics hold for conical limit sequences and sequences in peripheral subgroups.

For each conical limit point z, we define $C'_z = \mathrm{Opp}(\phi'^{-1}(z))$. Antipodality of ϕ' implies that C'_z contains $\Lambda' - \phi'^{-1}(z)$ for all $z \in \partial_{\mathrm{con}}\Gamma$.

First suppose that γ_n is a sequence limiting conically to a point $z \in \partial(\Gamma, \mathcal{H})$, and that γ_n^{-1} limits conically to some $z_- \in \partial(\Gamma, \mathcal{H})$.

Up to subsequence, γ_n lies a bounded distance away from the image of a $\mathcal{G}$-path $\{\alpha_n\}$ limiting to z. Proposition 5.4.18 implies that $\rho'(\gamma_n)$ is P-divergent, with unique P-limit $\phi'^{-1}(z)$. Similarly, $\rho'(\gamma_n^{-1})$ is P-divergent, with unique P-limit $\phi'^{-1}(z_-)$. Then Lemma 2.3.8 implies that $\rho'(\gamma_n)$ converges uniformly to $\phi'^{-1}(z)$ on compacts in C'_{z_-}.

For each parabolic point q, we write $q = g \cdot p$ for $p \in \Pi$, and then take $C'_q = \rho'(g) \cdot C'_p$. Proposition 5.6.13 says that for any $p \in \Pi$, any compact $K \subset C'_p$, and any open $U \subset G/P$ containing Λ'_p, if γ_n is an infinite sequence in Γ_p, then $\rho'(\gamma_n) \cdot K \subset U$ for sufficiently large n. Then, since $\Gamma_q = g\Gamma_p g^{-1}$, the same is true for any parabolic point q.

So, we just need to check that for each $p \in \Pi$, C'_p contains $\Lambda' - \Lambda'_p$. As in the proof of Lemma 5.6.17, we can find a compact subset $K \subset \partial(\Gamma, \mathcal{H}) - \{p\}$ so that $\Gamma_p K = \partial(\Gamma, \mathcal{H}) - \{p\}$ and C'_p contains $\phi'^{-1}(K)$. But then since C'_p is $\rho'(\Gamma_p)$-invariant (by Proposition 5.6.13), we have

$$\rho'(\Gamma_p) \cdot \phi'^{-1}(K) = \phi'^{-1}(\partial(\Gamma, \mathcal{H}) - \{p\}) \subset C'_p.$$

$\square$

Remark 5.6.20. The definition of the set Λ' and the map ϕ' immediately imply that the *fibers* of the deformed boundary extension $\phi' : \Lambda' \to \partial(\Gamma, \mathcal{H})$ satisfy the conclusions of Proposition 5.1.4: the fiber over each conical limit point is a singleton, and the fiber over each parabolic point p is the closure of the accumulation sets of Γ_p-orbits in C'_p. So, we obtain Proposition 5.1.4 by taking $\mathcal{W}$ to be the singleton $\{\rho\}$, and following the proof of Theorem 5.6.2 (using C_p for C'_p throughout).

Chapter 6

Examples of EGF representations

The last chapter of the thesis is devoted to examples of EGF representations. All of the examples we provide ultimately derive from *convex real projective structures* on manifolds.

We provide three families of examples in this chapter. First, in Section 6.1, we prove that *convex cocompact* projective structures give rise to EGF representations. To be precise, we prove:

Theorem 6.0.1. *Let* $(\Gamma, \mathcal{H})$ *be a relatively hyperbolic pair, and let* $\rho : \Gamma \to \mathrm{SL}(d+1, \mathbb{R})$ *be a group acting convex cocompactly on a properly convex domain* $\Omega \subset \mathbb{RP}^d$*. Suppose that each peripheral subgroup* $H \in \mathcal{H}$ *also acts convex cocompactly on* Ω*.*

Then ρ *is EGF with respect to* $P_{1,d}$*, where* $P_{1,d}$ *is the stabilizer of a flag of type* $(1, d)$ *in* $\mathbb{R}^{d+1}$*.*

In Section 6.2, we show that convex projective manifolds with *generalized cusps* also give rise to EGF representations.

Finally, in Section 6.3, we provide new examples of EGF representations by taking *symmetric powers* of certain projectively convex cocompact representations. We also show that for these examples, the entire space $\mathrm{Hom}(\Gamma, G)$ is peripherally stable. The main theorem of the previous chapter then implies that *any* sufficiently small deformation of one of these representations is still EGF.

6.1 Projective convex cocompact representations

Our first goal is to prove the following theorem, which shows that the relatively hyperbolic projectively convex cocompact groups explored in Chapter 4 fit into the framework of EGF representations.

Theorem 6.1.1. *Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and let $\rho : \Gamma \to \mathrm{SL}(d{+}1, \mathbb{R})$ be a group acting convex cocompactly on a properly convex domain $\Omega \subset \mathbb{RP}^d$. Then ρ is EGF with respect to $P_{1,d}$, where $P_{1,d}$ is the stabilizer of a flag of type $(1, d)$ in $\mathbb{R}^{d+1}$.*

The idea behind Theorem 6.1.1 is that the quotient map $\Lambda_\Omega(\Gamma) \to [\Lambda_\Omega(\Gamma)]_{\mathcal{H}} \simeq \partial(\Gamma, \mathcal{H})$ discussed in Section 4.3 should induce an EGF boundary extension. Since EGF representations are defined for *symmetric* parabolic subgroups, we need to augment this quotient map, and define it on an appropriate space of *flags*.

Definition 6.1.2. Let $\Gamma \subseteq \mathrm{Aut}(\Omega)$.

1. The *dual full orbital limit set* $\Lambda^*_\Omega(\Gamma)$ is the full orbital limit set of Γ^* in $\partial\Omega^*$.

2. The *flag-valued full orbital limit set* $\hat{\Lambda}_\Omega(\Gamma)$ is the set

$$\hat{\Lambda}_\Omega(\Gamma) := \{(x, w) \in \mathcal{F}_\pm(V) : x \in \Lambda_\Omega(\Gamma), w \in \Lambda^*_\Omega(\Gamma)\}.$$

3. The *maximal domain* $\Omega_{\max}(\Gamma)$ is the unique connected component of

$$\mathbb{P}(V) - \bigcup_{w \in \Lambda^*_\Omega(\Gamma)} \mathbb{P}(w)$$

containing Ω. Equivalently, $\Omega_{\max}(\Gamma)$ is the dual of the convex hull of $\Lambda^*_\Omega(\Gamma)$ in Ω^*.

We emphasize that $\Omega_{\max}(\Gamma)$ is *not* necessarily a properly convex set, which means that we cannot always define a Hilbert metric on it (so we do not have a guarantee that Γ acts properly discontinuously in general). However, when Γ acts convex cocompactly on Ω, we do get a properly discontinuous action, thanks to the following argument suggested by Jeff Danciger and Fanny Kassel:

Proposition 6.1.3. *Let Γ act convex cocompactly on a properly convex comain Ω. For any $x \in \Omega_{\max}(\Gamma)$, every accumulation point of $\Gamma \cdot x$ lies in $\Lambda_\Omega(\Gamma)$. In particular, Γ acts properly discontinuously on $\Omega_{\max}(\Gamma)$.*

Proof. When $\Omega_{\max}(\Gamma)$ is a *properly* convex domain, this follows immediately from Proposition 4.18 in [DGK17], which says that whenever Γ acts convex cocompactly on some domain Ω, and Ω' is any Γ-invariant properly convex domain containing Ω, then Γ acts convex cocompactly on Ω' and $\Lambda_\Omega(\Gamma) = \Lambda_{\Omega'}(\Gamma)$.

So, we consider the case where $\Omega_{\max}(\Gamma)$ is *not* properly convex. We may assume our domain Ω is chosen so that Γ acts convex cocompactly on both Ω and $\Omega^* \subset \mathbb{P}(V^*)$. Since $\Omega_{\max}(\Gamma)$ is not properly convex, its dual $\Omega_{\max}(\Gamma)^*$ (given by the convex hull of $\Lambda^*_\Omega(\Gamma)$ in Ω^*) has empty interior (i.e. it spans a proper projective subspace of $\mathbb{P}(V^*)$).

Given any $\varepsilon > 0$, we let Ω^*_ε be the uniform ε-neighborhood of $\Omega_{\max}(\Gamma)^*$, with respect to the Hilbert metric on Ω^*. We let $\Omega_\varepsilon \subset \mathbb{P}(V)$ denote the dual of Ω^*_ε. Note that Ω_ε is a Γ-invariant properly convex subset of $\Omega_{\max}(\Gamma)$, containing Ω.

Since duality reverses inclusions, and the intersection

$$\bigcap_{\varepsilon \to 0} \Omega_\varepsilon^*$$

is exactly the set $\Omega_{\max}(\Gamma)^*$, the union

$$\bigcup_{\varepsilon \to 0} \Omega_\varepsilon$$

is the set $\Omega_{\max}(\Gamma)$. So, if we fix $x \in \Omega_{\max}(\Gamma)$, for some $\varepsilon > 0$ we have $x \in \Omega_\varepsilon$. Then we apply Proposition 4.18 in [DGK17] to the *properly* convex domain Ω_ε to see that $\Gamma \cdot x$ must accumulate in the set $\Lambda_\Omega(\Gamma)$. $\square$

We are almost ready to prove Theorem 6.1.1. We first observe that, using a theorem of Danciger-Guéritaud-Kassel ([DGK17], Theorem 1.18), one can slightly strengthen Theorem 4.1.8 to obtain the following:

Corollary 6.1.4. *Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and suppose that $\rho : \Gamma \to \mathrm{PGL}(V)$ is a projectively convex cocompact representation, with $\rho|_H$ projectively convex cocompact for each $H \in \mathcal{H}$.*

Then there is a properly convex $\rho(\Gamma)$-invariant domain $\Omega \subset \mathbb{P}(V)$, with $\rho(\Gamma)$ acting convex cocompactly on Ω, such that every nontrivial projective segment in $\partial\Omega$ is contained in $\Lambda_\Omega(H)$ for some $H \in \mathcal{H}$.

We also observe:

Proposition 6.1.5. *Let $\psi : \partial(\Gamma, \mathcal{H}) \to [\Lambda_\Omega(\Gamma)]_{\mathcal{H}}$ be the equivariant homeomorphism coming from Corollary 4.1.10.*

If $z \in \partial(\Gamma, \mathcal{H})$ is a conical limit point, and γ_n is a sequence limiting to z in $\overline{\Gamma} = \Gamma \sqcup \partial(\Gamma, \mathcal{H})$, then $\psi(z) \in \mathbb{P}(V)$ is the unique attracting limit point for γ_n in $\mathbb{P}(V)$.

Proof. It suffices to show that any subsequence of γ_n has a further subsequence whose unique attracting limit point is $\psi(z)$. So, using the convergence group property, we can take a subsequence and assume that there is some point $y \in \partial(\Gamma, \mathcal{H})$ so that γ_n converges to z on every point in the set $\partial(\Gamma, \mathcal{H}) - \{y\}$. We can further assume that the pair $(\Gamma, \mathcal{H})$ is not elementary, so $\partial(\Gamma, \mathcal{H})$ contains infinitely many points.

So, by Corollary 6.1.4, we can find distinct points $u, v \in \Lambda_\Omega(\Gamma)$ so that the projective line segment (u, v) lies in Ω, and $\rho(\gamma_n)u$, $\rho(\gamma_n)v$ both converge to $\psi(z)$. Lemma 4.1 in [DGK17] implies that $F_\Omega(\psi(z)) \subset \Lambda_\Omega(\Gamma)$, and then Corollary 6.1.4 implies that $F_\Omega(\psi(z)) = \{\psi(z)\}$.

Then for any $x \in (u, v)$, $\rho(\gamma_n)x$ converges to $\psi(z)$, and we are done by Proposition 2.1.18. □

Proof of Theorem 6.1.1. Fix a $(d+1)$-dimensional real vector space V, let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair, and let $\rho : \Gamma \to \mathrm{PGL}(V)$ be a representation such that $\rho(\Gamma)$ acts convex cocompactly on a properly convex domain Ω. This implies (see [IZ22]) that each $H \in \mathcal{H}$ also acts convex cocompactly on Ω.

The first step in the proof is to define our boundary extension $\hat{\phi} : \hat{\Lambda} \to \partial(\Gamma, \mathcal{H})$, where $\hat{\Lambda}$ is the flag-valued full orbital limit set $\hat{\Lambda}_\Omega(\Gamma)$. We use the equivariant homeomorphism $\psi : \partial(\Gamma, \mathcal{H}) \to [\Lambda]_{\mathcal{H}}$ coming from Corollary 4.1.10 to define an equivariant surjection $\phi : \Lambda \to \partial(\Gamma, \mathcal{H})$, where the preimage of each parabolic point p in $\partial(\Gamma, \mathcal{H})$ is exactly $\Lambda_\Omega(H)$ for $H = \mathrm{Stab}_\Gamma(p)$.

Similarly, if Λ^* is the dual full orbital limit set, we can find an equivariant surjection $\phi^* : \Lambda^* \to \partial(\Gamma, \mathcal{H})$, where ${\phi^*}^{-1}(z)$ is a single hyperplane if z is a conical limit point, and the dual full orbital limit set of $\mathrm{Stab}_\Gamma(z)$ if z is a parabolic point.

Each element of Λ^* is a supporting hyperplane of the domain Ω. Corollary 6.1.4 implies that for every point (x, w) in the flag-valued full orbital limit set $\hat{\Lambda} = \hat{\Lambda}_\Omega(\Gamma)$, either:

1. $x = \phi^{-1}(z)$ and $w = {\phi^*}^{-1}(z)$ for a conical limit point $z \in \partial(\Gamma, \mathcal{H})$, or
2. $x \in \Lambda_H$ and $w \in \Lambda_H^*$ for a peripheral subgroup $H \in \mathcal{H}$.

This allows us to combine ϕ and ϕ^* to get a well-defined equivariant surjection $\hat{\phi} : \hat{\Lambda} \to \partial(\Gamma, \mathcal{H})$.

The next step is to define the open subsets $C_z \subset \mathcal{F}_\pm$ for each $z \in \partial(\Gamma, \mathcal{H})$. If $z \in \partial(\Gamma, \mathcal{H})$ is a conical limit point, we define the set C_z by

$$C_z = \{\nu \in \mathcal{F}_\pm : \nu \text{ is opposite to } \phi^{-1}(z)\}.$$

Otherwise, if z is a parabolic point, we consider the maximal domain $\Omega_{\max}(H) \subset \mathbb{P}(V)$ for $H = \mathrm{Stab}_\Gamma(z)$. Dually, we can define $\Omega^*_{\max}(H) \subset \mathbb{P}(V^*)$, viewing $\Lambda_\Omega(H)$ as the dual full orbital limit set of H^* acting on Ω^*.

Then, we define

$$C_z = \{(x, w) \in \mathcal{F}_\pm(V) : x \in \Omega_{\max}(H), w \in \Omega^*_{\max}(H)\}.$$

For every $z \in \partial(\Gamma, \mathcal{H})$, C_z is open, and Corollary 6.1.4 implies that C_z contains $\hat{\phi}^{-1}(z')$ for every $z' \neq z$ in $\partial(\Gamma, \mathcal{H})$.

The last step is to check that the map $\hat{\phi}$ actually extends convergence dynamics, using the sets C_z. Because of Proposition 5.1.2, it suffices to verify that $\hat{\phi}$ has the right dynamical behavior for conical limit sequences and sequences in peripheral subgroups.

Fix a conical limit point $z \in \partial(\Gamma, \mathcal{H})$, and let $\gamma_n \in \Gamma$ be a sequence limiting conically to z, with γ_n^{-1} limiting conically to some point z_-.

Proposition 6.1.5 implies that $\phi^{-1}(z)$ is the unique attracting subspace for $\rho(\gamma_n)$ in $\mathbb{P}(V)$, and $\phi^{-1}(z_-)$ is the unique attracting subspace for $\rho(\gamma_n^{-1})$ in $\mathbb{P}(V)$. Dually, $(\phi^*)^{-1}(z_-)$ is the unique attracting subspace for $\rho(\gamma_n^{-1})$ in $\mathbb{P}(V^*)$, and $(\phi^*)^{-1}(z)$ is the unique attracting subspace for $\rho(\gamma_n)$ in $\mathbb{P}(V^*)$. Then Lemma 2.3.8 ensures that $\rho(\gamma_n)$ converges to $\hat{\phi}^{-1}(z)$ uniformly on compacts in C_{z_-}.

On the other hand, if z is a parabolic point and γ_n is an infinite sequence in $\mathrm{Stab}_\Gamma(z)$, Proposition 6.1.3 implies that for any compact $K \subset C_z$, the set $\rho(\gamma_n) \cdot K$ eventually lies in any given neighborhood of $\phi^{-1}(z)$, as required. $\square$

6.2 Generalized cusps

In [CLT18], Cooper-Long-Tillmann studied relative stability properties for the holonomy of certain *noncompact* convex projective orbifolds—those with *generalized cusps*. They consider the situation of a convex projective manifold M (with strictly convex boundary) which is a union of a compact piece and finitely many ends homeomorphic to $N \times [0, \infty)$, where N is a compact manifold with virtually abelian fundamental group. The ends of such a manifold are called generalized cusps.

In [BCL20], Ballas-Cooper-Leitner classified generalized cusps in $\mathbb{RP}^d$ into $d+1$ different *cusp types*. A *type 0 cusp* is projectively equivalent to a cusp in $\mathbb{H}^d$ and has virtually unipotent holonomy, while a *type d cusp* has virtually diagonalizable holonomy. The remaining cusp types "interpolate" between these two.

If $M = \Omega/\Gamma$ is a finite-volume convex projective d-manifold, and $\Omega \subseteq \mathbb{RP}^d$ is a strictly convex domain, then work of Crampon-Marquis [CM14] (see also Cooper-Long-Tillmann [CLT15]) shows that M decomposes into a union of a compact piece and finitely many type 0 cusps, and $\pi_1 M$ is hyperbolic relative to the collection of cusp groups

$$\mathcal{C} = \{\pi_1 C_i : C_i \text{ a cusp in } M\}.$$

Corollary 6.2.5 implies that the holonomy $\rho : \pi_1 M \to \mathrm{SL}(d+1, \mathbb{R})$ of such a manifold is an EGF representation, and Theorem 5.6.2 says that peripherally stable deformations of ρ are also EGF.

In this situation, we can give an explicit example of a subspace of $\mathrm{Hom}(\pi_1 M, \mathrm{SL}(d+1))$ which is peripherally stable. Following Cooper-Long-Tillman, we let $\mathrm{VFG}(d+1)$ denote the set of *virtual flag groups*: discrete subgroups of $\mathrm{SL}(d+1, \mathbb{R})$ which have a finite-index subgroup which is conjugate to a group of upper-triangular matrices.

Definition 6.2.1. Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair. We let

$$\mathrm{Rep}_{\mathrm{VFG}}(\Gamma, \mathcal{H})$$

denote the space

$$\{\rho \in \mathrm{Hom}(\Gamma, \mathrm{SL}(d+1, \mathbb{R})) : \forall H_i \in \mathcal{H}, \rho(H_i) \in \mathrm{VFG}(d+1)\}.$$

Let $\rho : \pi_1 M \to \mathrm{SL}(d+1, \mathbb{R})$ be the holonomy of the manifold M. [CLT18, Theorem 0.1] says that the image of the holonomy of each cusp of M lies in $\mathrm{VFG}(d+1)$—i.e. $\rho \in \mathrm{Rep}_{\mathrm{VFG}}(\pi_1 M, \mathcal{C})$. And, if ρ' is any small deformation of ρ in $\mathrm{Rep}_{\mathrm{VFG}}(\pi_1 M, \mathcal{C})$, then ρ' is the holonomy of a nearby convex projective structure on M.

Our goal in this section is to prove that these nearby representations are also EGF:

Theorem 6.2.2. *Let $M = \Omega/\Gamma$ be a finite-volume convex projective manifold, and suppose that Ω is strictly convex (so the holonomy $\rho : \pi_1 M \to \mathrm{SL}(d+1, \mathbb{R})$ is EGF with boundary extension ϕ).*

Then the space $\mathrm{Rep}_{\mathrm{VFG}}(\pi_1 M, \mathcal{C})$ is peripherally stable at (ρ, ϕ). In particular, due to Theorem 5.6.2, an open subset of $\mathrm{Rep}_{\mathrm{VFG}}(\pi_1 M, \mathcal{C})$ containing ρ consists of EGF representations.

In [Bob19], Bobb produced examples of convex projective d-manifolds with cusps of any non-diagonalizable type for all dimensions d. These examples are all small deformations of finite-volume hyperbolic d-manifolds. So, a consequence of Theorem 6.2.2 is the following:

Corollary 6.2.3. *In every dimension d and for every $0 < t < d$, there exists a convex projective manifold with a type t generalized cusp and EGF holonomy.*

Remark 6.2.4. We expect that a version of Theorem 6.2.2 also holds with weaker assumptions on M, which are more in line with the original assumptions in the Cooper-Long-Tillmann stability result. For instance, we conjecture that whenever M

is a convex projective manifold with strictly convex boundary, each end of M is a generalized cusp, and $\pi_1 M$ is relatively hyperbolic (relative to a collection of subgroups $\mathcal{H}$), then the holonomy of M is an EGF representation, and $\mathrm{Rep}_{\mathrm{VFG}}(\pi_1 M, \mathcal{H})$ is a peripherally stable subspace.

6.2.0.1 Geometrically finite strictly convex projective manifolds

In [CM14], Crampon-Marquis define a notion of geometrical finiteness for groups acting on strictly convex domains in projective space. Zhu [Zhu19] showed that these groups are always relatively dominated, a stronger condition than relative asymptotic embeddedness. So an immediate corollary of Zhu's result and Theorem 5.1.7 is:

Corollary 6.2.5. *Let $\Omega \subseteq \mathbb{R}\mathrm{P}^d$ be a strictly convex domain, and suppose that $\Gamma \subseteq \mathrm{Aut}(\Omega)$ acts on Ω geometrically finitely in the sense of Crampon-Marquis.*

Then the inclusion $\Gamma \hookrightarrow \mathrm{SL}(d+1, \mathbb{R})$ is EGF.

It is worth noting that *unlike* generalized cusps, the cusps in the geometrically finite convex projective manifolds of Crampon-Marquis do *not* need to have compact cross-section. It seems possible that there is a general theory of cusped convex projective manifolds which incorporates both the examples of Crampon-Marquis and the generalized cusps classified by Ballas-Cooper-Leitner discussed below. It would be interesting to explore when such cusped convex projective manifolds have EGF holonomy representation.

6.2.1 Generalized cusps: definitions

Definition 6.2.6. Let $\Omega \subset \mathbb{RP}^d$ be a properly convex set with nonempty interior. In general, Ω might not be either open or closed.

- The *frontier* of Ω is $\mathrm{Fr}(\Omega) = \overline{\Omega} - \mathrm{int}(\Omega)$.
- The *nonideal boundary* of Ω is $\partial_n\Omega = \mathrm{Fr}(\Omega) \cap \Omega$.
- The *ideal boundary* of Ω is $\partial_i\Omega = \mathrm{Fr}(\Omega) - \partial_n\Omega$.

Some authors (e.g. [CLT18, BCL20]) just refer to the nonideal boundary of Ω as the *boundary* $\partial\Omega$. We avoid this since it conflicts with the notation used in the previous section (where we use $\partial\Omega$ to denote the *ideal* boundary of a properly convex domain Ω).

If $M = \Omega/\Gamma$ is a convex projective manifold with boundary, then ∂M is identified with $\partial_n\Omega/\Gamma$. When $\partial_n\Omega$ contains no nontrivial projective segments, then we say that the manifold M has *strictly convex boundary.*

Definition 6.2.7. Let $\Omega \subset \mathbb{RP}^d$ be a properly convex set with nonempty interior, and let $\Gamma \subseteq \mathrm{Aut}(\Omega)$ be discrete. A manifold $C = \Omega/\Gamma$ is a *generalized cusp* if C has compact and strictly convex boundary, $\Gamma \simeq \pi_1 C$ is virtually abelian, and C is homeomorphic to $\partial C \times [0, \infty)$.

Recall that $\mathrm{VFG}(d+1)$ is the set of *virtual flag groups* in $\mathrm{PGL}(d+1, \mathbb{R})$, i.e. the discrete groups which have a finite-index subgroup conjugate to a group of upper-triangular matrices.

Definition 6.2.8. Let C be a generalized cusp. We let

$$\mathrm{VFG}(C) = \{\rho : \pi_1 C \to \mathrm{PGL}(d+1, \mathbb{R}) : \rho(\pi_1 C) \in \mathrm{VFG}\}.$$

Theorem 6.2.9 ([CLT18], Theorem 6.25)**.** *Let C be a generalized cusp, and let $\mathcal{U}$ be the set of holonomies of convex projective structures on C with strictly convex boundary. Then $\mathcal{U}$ is an open subset of* $\mathrm{VFG}(C)$*.*

6.2.2 Generalized horospheres

In [CLT18], Cooper-Long-Tillmann show that if C is a generalized cusp with holonomy $\rho : \pi_1 C \to \mathrm{PGL}(d+1, \mathbb{R})$, there is a finite-index subgroup $\Gamma_1 \subseteq \pi_1 C$ (depending only on $\pi_1 C$ and d) so that $\rho(\Gamma_1)$ is a lattice in a *syndetic hull* of $\rho(\Gamma_1)$: a uniquely determined connected Lie group $T(\rho) \subset \mathrm{PGL}(d+1, \mathbb{R})$, conjugate into the group of upper triangular matrices. This group is called the *translation group* of the cusp.

We may assume that Γ_1 is free abelian, so it is a lattice in $\Gamma_1 \otimes_{\mathbb{Z}} \mathbb{R} \simeq \mathbb{R}^{d-1}$. The restriction

$$\rho|_{\Gamma_1} : \Gamma_1 \to \mathrm{PGL}(d+1, \mathbb{R})$$

extends uniquely to an embedding of Lie groups

$$\iota_\rho : \mathbb{R}^{d-1} \hookrightarrow \mathrm{PGL}(d+1, \mathbb{R})$$

with image $T(\rho)$.

We observe the following:

Proposition 6.2.10. *Let C be a generalized cusp. The embedding*

$$\iota_\rho : \mathbb{R}^{d-1} \to \mathrm{PGL}(d+1, \mathbb{R})$$

varies continuously with $\rho \in \mathrm{VFG}(C)$ in the compact-open topology on maps from $\mathbb{R}^{d-1}$ into $\mathrm{PGL}(d+1, \mathbb{R})$.

Proof. The Lie algebra of $\mathrm{PGL}(d+1, \mathbb{R})$ is identified with $\mathfrak{sl}(d+1, \mathbb{R})$. Fix a finite set S of generators for Γ_1. The subspace $\mathfrak{s}_\rho \subset \mathfrak{sl}(d+1, \mathbb{R})$ spanned by $\log(\rho(S))$ varies smoothly with ρ, and the induced map

$$\mathbb{R}^{d-1} \to \mathfrak{sl}(d+1, \mathbb{R})$$

with image $\mathfrak{s}_\rho$ varies continuously in the compact-open topology. The embedding ι_ρ is given by composition with the exponential map. $\square$

When $C = \Omega/\Gamma$ is projectively equivalent to a cusp in some hyperbolic manifold, we can assume that Ω is a closed horoball in $\mathbb{H}^d$. The nonideal boundary $\partial_n \Omega$ is a horosphere preserved by Γ, which carries a Γ-invariant Euclidean metric. In this case the translation group is the group of Euclidean translations on $\partial_n \Omega$.

When C is a generalized cusp, we can always find some orbit of the translation group $T(\rho)$ in $\mathbb{RP}^d$ which is a strictly convex hypersurface (see Proposition 6.22 in [CLT18]). This hypersurface is called a *generalized horosphere*. Its convex hull in $\mathbb{RP}^d$ is a *generalized horoball*. We can always choose this horoball so that its quotient by Γ is contained in the generalized cusp C.

Now consider a convex projective manifold $M = \Omega/\Gamma$ with strictly convex boundary which can be written as a union of a compact piece and finitely many generalized cusps. It is always possible to choose the cusps so that each cusp $C \subset M$ is a quotient $C = \Omega_C/\Gamma_C$, where $\Omega_C \subset \Omega$ is a convex subset whose nonideal boundary is a generalized horosphere, and $\Gamma_C = \pi_1 C$ is the cusp group. The boundary $\partial_n \Omega_C$ is *homogeneous*: the translation group $T(\Gamma_C)$ acts simply transitively on $\partial_n \Omega_C$.

6.2.3 The ideal boundary

The Ballas-Cooper-Leitner classification of generalized cusps allows us to get a more explicit description of the ideal boundary of Ω. Given a generalized cusp $C = \Omega/\Gamma$, we let Ω_C denote the "standard" Γ-invariant domain with homogeneous nonideal boundary, alluded to above.

Proposition 6.2.11 (Lemmas 1.24 and 1.25 in [BCL20])**.** *Let $C = \Omega_C/\Gamma$ be a generalized cusp. The ideal boundary of Ω_C is a projective k-simplex Δ_C. There is a unique supporting hyperplane H_C of Ω_C containing Δ_C, and the affine chart*

$$\mathbb{A}_C = \mathbb{R}\mathrm{P}^d - H_C$$

is the unique affine chart containing Ω_C as a closed subset.

The vertices of Δ_C must be preserved by Γ, and in fact they are all eigenvectors for the translation group $T(\Gamma)$.

Each generalized horosphere S_C for C is a strictly convex hypersurface contained in the affine chart $\mathbb{A}_C$. The closure of this hypersurface in $\mathbb{R}\mathrm{P}^d$ is either $S_C \cup \partial\Delta_C$ (if C is a "type d" cusp) or $S_C \cup \Delta_C$ (if C is any other type of cusp).

6.2.4 Deformations of convex hypersurfaces

The main ingredient in the proof of Theorem 6.2.2 is the following:

Lemma 6.2.12. *Let C be a hyperbolic cusp with holonomy ρ. Let p_C be the cusp point, and let H_C be the unique supporting hyperplane of Ω_C at p_C.*

Let $x \in \mathbb{A}_C = \mathbb{R}\mathrm{P}^d - H_C$, let $U \subset \mathbb{R}\mathrm{P}^d$ be an open subset containing p_C, and let $F \subset \pi_1 C$ be a cofinite subset such that $\rho(F) \cdot x \subset U$.

There exists a neighborhood $\mathcal{W}$ of ρ in $\mathrm{VFG}(C)$ so that for any $\rho' \in \mathcal{W}$, we have

$$\rho'(F) \cdot x \subset U.$$

Proof. [CLT18], Theorem 6.25 implies that we can choose a neighborhood $\mathcal{W}$ of ρ in $\mathrm{VFG}(C)$ consisting of holonomies of generalized cusps. For any $\rho' \in \mathcal{W}$, we let Ω' denote a "standard" properly convex set preserved by $\rho'(\pi_1 C)$, whose non-ideal boundary is a generalized horosphere.

Since p_C and H_C are respectively the unique eigenvector and fixed hyperplane of $\rho(\pi_1 C)$, we can choose our neighborhood $\mathcal{W}$ so that for any $\rho' \in \mathcal{W}$, any eigenvectors and fixed hyperplanes of $\rho'(\pi_1 C)$ are close to p_C, H_C.

In particular, we can choose $\mathcal{W}$ so that the ideal boundary of Ω' is a k-simplex Δ' contained in U. And, by applying a small conjugation in $\mathrm{PGL}(d+1, \mathbb{R})$, we can assume that H_C is the unique supporting hyperplane of Ω' containing Δ'.

Let $T(\rho)$ be the translation group of ρ. The orbit $T(\rho) \cdot x$ is a paraboloid in $\mathbb{A}_C$. We can write $\mathbb{A}_C = \mathbb{R}^{d-1} \times \mathbb{R}$, and then view this paraboloid as the graph of a function

$$f_\rho : \mathbb{R}^{d-1} \to \mathbb{R}.$$

The function f_ρ is determined by the condition

$$(u, f_\rho(u)) = \iota_\rho(u) \cdot x.$$

Here $\iota_\rho : \mathbb{R}^{d-1} \to \mathrm{PGL}(d+1, \mathbb{R})$ restricts to ρ on a finite-index subgroup $\Gamma_1 \subset \pi_1 C$, with Γ_1 identified with $\mathbb{Z}^{d-1} \subset \mathbb{R}^{d-1}$.

If $T(\rho')$ is the translation group of ρ', then Lemma 6.24 in [CLT18] implies that $T(\rho') \cdot x$ is a strictly convex hypersurface $S' \subset \mathbb{A}_C$. The hypersurface S' is the graph of a map $f_{\rho'} : \mathbb{R}^{d-1} \to \mathbb{R}$, satisfying

$$(u, f_{\rho'}(u)) = \iota_{\rho'}(u) \cdot x.$$

Proposition 6.2.10 implies that $f_{\rho'}$ varies continuously (in the compact-open topology on continuous maps $\mathbb{R}^{d-1} \to \mathbb{R}$) as ρ' varies in $\mathcal{W}$.

We fix a norm $||\cdot||$ on $\mathbb{R}^{d-1}$. There is a constant D so that for any $(u, v) \in \mathbb{A}_C$, if $||u|| > 1$ and $|v|/||u|| > D$, then (u, v) is contained in the neighborhood U of p_C.

$f_{\rho'}$ is a strictly convex function, which we can assume is nonnegative and uniquely minimized at the origin. So, if $f_{\rho'}(u)/||u|| > D$ on $\{u \in \mathbb{R}^{d-1} : ||u|| = M\}$ for some constant M, then $f_{\rho'}(u)/||u|| > D$ for all u with $||u|| \geq M$.

So, as long as W is sufficiently small, there is a fixed ball $B \subset \mathbb{R}^{d-1}$ so that if $u \in \mathbb{R}^{d-1} - B$, then

$$(u, f_{\rho'}(u)) \in U$$

for any $\rho' \in \mathcal{W}$.

The ball B contains at most finitely many elements of $\Gamma_1 - F$. So we can in fact choose $\mathcal{W}$ small enough so that for any $\rho' \in \mathcal{W}$, every element of $\rho'(\Gamma_1 - F) \cdot x$ lies in U. Then since $|\Gamma_1 : \pi_1 C| < \infty$ we can shrink $\mathcal{W}$ even further to guarantee that for any $\rho' \in \mathcal{W}$,

$$\rho'(\pi_1 C - F) \cdot x$$

lies in U as well. $\square$

6.2.5 Peripheral stability

Proof of Theorem 6.2.2. Let $\rho : \pi_1 M \to \mathrm{PGL}(d+1, \mathbb{R})$ be the holonomy of a finite-volume convex projective manifold M, and let Ω be a ρ-invariant strictly convex domain such that $M = \Omega/\rho(\pi_1 M)$. Write $\Gamma = \pi_1 M$, and let $\mathcal{H}$ be the collection of cusp groups, so $(\Gamma, \mathcal{H})$ is a relatively hyperbolic pair, and for each $H \in \mathcal{H}$, $\rho|_H$ is the holonomy of a hyperbolic cusp.

Since Ω is strictly convex, [CLT15], Theorem 11.6 also implies that Ω has C^1 boundary. So there is a ρ-equivariant homeomorphism $\partial\Omega \to \partial\Omega^*$ assigning the point $z \in \partial\Omega$ to the unique supporting hyperplane of Ω at z. We let $\partial\hat{\Omega}$ denote the space

$$\{(x, w) \in \mathcal{F}_{\pm} : x \in \partial\Omega, w \in \partial\Omega^*\}.$$

There is an equivariant homeomorphism $\psi : \partial(\Gamma, \mathcal{H}) \to \partial\hat{\Omega}$, with the parabolic points in $\partial(\Gamma, \mathcal{H})$ corresponding to the fixed flags of the cusp groups. The inverse map $\phi : \partial\hat{\Omega} \to \partial(\Gamma, \mathcal{H})$ extends the convergence dynamics of Γ. For each parabolic point $p \in \partial(\Gamma, \mathcal{H})$, the open set C_p is

$$\mathrm{Opp}(\psi(p)) = \{\xi \in \mathcal{F}_\pm : \xi \text{ is opposite to } \psi(p)\}.$$

Let $\pi : \mathcal{F}_\pm \to \mathbb{RP}^d$ be the canonical projection map. It suffices to show that for any compact set $K \subset \pi(C_p)$, any open neighborhood U of $\pi(\psi(p))$ in $\mathbb{RP}^d$, and any cofinite subset $F \subset \Gamma_p = \mathrm{Stab}_\Gamma(p)$ such that

$$\rho(F) \cdot K \subset U,$$

we can find a neighborhood $\mathcal{W} \subset \mathrm{VFG}(\Gamma_p)$ containing ρ such that

$$\rho'(F) \cdot K \subset U$$

for any $\rho' \in \mathcal{W}$. (The same argument applied dually will show that we can upgrade K to a compact subset of $C_p \subset \mathcal{F}_\pm$ and U to an open subset in $\mathcal{F}_\pm$).

For any x in such a compact set $K \subset \pi(C_p)$, Lemma 6.2.12 implies that there is an open neighborhood $\mathcal{W} \subset \mathrm{VFG}(\Gamma_p)$, containing the restriction $\rho|_{\Gamma_p}$, so that for any $\rho' \in \mathcal{W}$, we have

$$\rho'(F) \cdot x \subset U.$$

But then by compactness of K, we can find a *single* neighborhood $\mathcal{W} \subset \mathrm{VFG}(\Gamma_p)$ so that $\rho'(F) \cdot K \subset U$, as required. $\square$

6.3 Symmetric powers of convex cocompact groups

In this section, we construct new examples of extended geometrically finite representations by taking symmetric powers of convex cocompact representations of groups which are hyperbolic relative to virtually abelian subgroups. We also prove an absolute stability result for these representations.

6.3.1 Symmetric powers

Let V be a finite-dimensional real vector space. We let τ_m denote the symmetric representation

$$\mathrm{SL}(V) \to \mathrm{SL}(\mathrm{Sym}^m(V)).$$

Throughout this section, we will view $\mathrm{Sym}^m(V)$ as a quotient of the space of homogeneous degree-m polynomials in elements of V. We will always leave this quotient implicit. That is, if $v_1, \ldots, v_k \in V$, and $r_1, \ldots, r_k \in \mathbb{N} \cup \{0\}$ with $\sum r_i = m$, we will view the monomial $v_1^{r_1} \cdots v_k^{r_k}$ as an element of $\mathrm{Sym}^m(V)$.

There is a τ_m-equivariant embedding

$$\iota : \mathbb{P}(V) \to \mathbb{P}(\mathrm{Sym}^m(V))$$

given by $[v] \mapsto [v^m]$. There is also a corresponding dual embedding

$$\iota^* : \mathbb{P}(V^*) \to \mathbb{P}(\mathrm{Sym}^m(V)^*),$$

using the canonical identification $\mathrm{Sym}^m(V^*) \simeq \mathrm{Sym}^m(V)^*$. We observe that $v \in \mathbb{P}(V)$ and $w \in \mathbb{P}(V^*)$ are transverse if and only if their respective images under ι and ι^* are also transverse. This means that the maps ι and ι^* also give rise to a τ_m-equivariant map

$$\hat{\iota} : \mathcal{F}_{\pm}(V) \to \mathcal{F}_{1,d'}(\mathrm{Sym}^m(V))$$

given by $\hat{\iota}(v, w) = (\iota(v), \iota^*(w))$. Here, $d+1$ is the dimension of V and $d'+1$ is the dimension of $\mathrm{Sym}^m(V)$. In this section we will mostly use will use $\mathcal{F}(\mathrm{Sym}^m V)$ to denote $\mathcal{F}_{1,d'}(\mathrm{Sym}^m V)$.

6.3.1.1 Dynamics in symmetric powers

The dynamics of divergent sequences in $\mathrm{SL}(V)$ on $\mathbb{P}(V)$ and $\mathbb{P}(V^*)$ are respected in $\mathbb{P}(\mathrm{Sym}^m V)$, in the following sense:

Proposition 6.3.1. *Let $\{g_n\}$ be an infinite sequence in $\mathrm{SL}(V)$ such that for some $w \in \mathbb{P}(V^*)$, $x \in \mathbb{P}(V)$, we have $g_n|_{\mathrm{Opp}(w)} \to x$ uniformly on compacts.*

Then

$$\tau_m(g_n)|_{\mathrm{Opp}(\iota^*(w))} \to \iota(x)$$

uniformly on compacts.

6.3.2 Symmetric powers of relatively hyperbolic convex cocompact groups

Suppose that $\overline{\rho} : \Gamma \to \mathrm{PGL}(V)$ is a convex cocompact representation, so that there is a properly convex domain $\Omega \subset \mathbb{P}(V)$ with $\overline{\rho}(\Gamma)$ acting convex cocompactly on Ω. As mentioned in the introduction to this paper, we can always replace Γ with a finite-index subgroup so that $\overline{\rho}$ lifts to a representation $\rho : \Gamma \to \mathrm{SL}(V)$ with finite kernel and discrete image. In this situation, we also say that $\rho : \Gamma \to \mathrm{SL}(V)$ is *convex cocompact in* $\mathbb{P}(V)$.

Let $(\Gamma, \mathcal{H})$ be a relatively hyperbolic pair such that each group in $\mathcal{H}$ is virtually abelian, and let $\rho : \Gamma \to \mathrm{SL}(V)$ be a convex cocompact representation.

Representations of this form have been studied extensively by Islam-Zimmer [IZ19b], [IZ19a], who proved a number of strong structural results. In particular, Islam-Zimmer showed that in this situation, for each $H \in \mathcal{H}$, the image $\rho(H)$ acts cocompactly on a properly embedded k-simplex $\Delta_H \subset \Omega$, with $k = \mathrm{rank}(H)$. (A simplex $\Delta \subset \Omega$ is *properly embedded* if $\partial\Delta \subset \partial\Omega$.) Conversely, every properly embedded maximal k-simplex in the convex hull of $\Lambda_\Omega(\Gamma)$ always has a cocompact action by some $H \in \mathcal{H}$ with rank k.

We let

$$\rho^m : \Gamma \to \mathrm{SL}(\mathrm{Sym}^m V)$$

denote the composition $\tau_m \circ \rho$. We have two main goals in this section. The first is to prove the following:

Theorem 6.3.2. *The representation ρ^m is EGF, with respect to the parabolic subgroup $P \subset \mathrm{SL}(\mathrm{Sym}^m V)$ stabilizing a line in a hyperplane in $\mathrm{Sym}^m V$.*

Note that Theorem 6.3.2 does *not* follow directly from the fact that convex cocompact representations in $\mathbb{P}(V)$ are EGF (Theorem 6.1.1), because we do not know that the representations ρ^m are convex cocompact in $\mathbb{P}(\mathrm{Sym}^m V)$. In fact, Jeff Danciger and Fanny Kassel have indicated in personal communication to the author that ρ^m should *never* be convex cocompact in $\mathbb{P}(\mathrm{Sym}^m V)$ unless the collection $\mathcal{H}$ is empty: while $\rho^m(\Gamma)$ does preserve a properly convex domain in $\mathbb{P}(\mathrm{Sym}^m V)$, the convex hull of the full orbital limit set in any such domain seems "too big" for $\rho^m(\Gamma)$ to act cocompactly.

Our second goal is to show:

Theorem 6.3.3. *There is an open subset of ρ^m in*

$$\mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^m V))$$

consisting of EGF representations.

We will prove Theorem 6.3.3 by proving that the *entire* space $\mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^m V))$ is peripherally stable about ρ^m, and applying Theorem 5.6.2. In particular this shows that small perturbations of ρ^m still have finite kernel and discrete image, giving new examples of discrete subgroups of higher-rank Lie groups which are stable (as discrete groups).

6.3.2.1 Proof strategy

6.3.2.2 Proof strategy

To show that ρ^m is EGF, we will give an explicit description of the boundary set $\hat{\Lambda}_m \subset \mathcal{F}(\mathrm{Sym}^m V)$. The naive choice is to just take $\hat{\Lambda}_m$ to be $\hat{\iota}(\hat{\Lambda}_\Omega(\Gamma))$, where $\hat{\Lambda}_\Omega(\Gamma)$ is the flag-valued full orbital limit set giving the EGF boundary set for $\rho : \Gamma \to \mathrm{SL}(V)$ (see Section 6.1). While there is a equivariant surjective map from this set to $\partial(\Gamma, \mathcal{H})$, it turns out that we will have to *enlarge* it in order to ensure that the relevant dynamics hold.

The idea is the following: for each parabolic point $p \in \partial(\Gamma, \mathcal{H})$, with stabilizer H, we take the fiber over p in $\hat{\Lambda}_m$ to be the space of flags in the boundary of a simplex $S_H \subset \mathbb{P}(\mathrm{Sym}^m V)$, constructed using the simplex $\Delta_H \subset \Omega$ on which $H = \mathrm{Stab}_\Gamma(p)$ acts cocompactly. The simplex S_H is chosen so that if γ_n is a sequence in Γ converging to p, then a face of S_H spans a minimal attracting subspace of $\rho^m(\gamma_n)$. We also want to ensure that the simplex S_H is *stable*, i.e. if ρ_t^m is a small deformation of ρ^m in $\mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^m V))$, then $\rho_t^m(H)$ preserves a simplex S_H^t close to S_H. We verify that S_H has these properties by analyzing the relationship between the weights of ρ and ρ^m on the virtually abelian group H.

The other main steps in the proof involve checking that the boundary set $\hat{\Lambda}_m$ we construct is actually a compact space surjecting continuously onto $\partial(\Gamma, \mathcal{H})$, and constructing the open sets $\hat{C}_p$ also required by the definition of an EGF representation. For the latter, we make heavy use of the fact that the *dual* action of $\rho(\Gamma)$ on $\mathbb{P}(V^*)$ is also projectively convex cocompact, which allows us to construct a stable *dual* simplex S_H^* for each $H \in \mathcal{H}$. The vertices of S_H^* are thought of as hyperplanes in $\mathbb{P}(\mathrm{Sym}^m(V))$, cutting out a region C_p of $\mathbb{P}(\mathrm{Sym}^m(V))$ on which $\rho^m(H)$ attracts points towards S_H.

6.3.2.3 Example: symmetric squares of convex projective 3-manifold groups

We illustrate the general idea of our approach with a specific example. Let $\Omega \subset \mathbb{R}\mathrm{P}^3$ be a properly convex domain, and let $\Gamma \subseteq \mathrm{Aut}(\Omega)$ be a discrete group acting cocompactly on Ω. In [Ben06], Benoist produced examples of such groups which are hyperbolic relative to a nonempty collection $\mathcal{H}$ of virtually abelian subgroups of rank 2. Further examples were constructed by Ballas-Danciger-Lee in [BDL15]. Up to finite index, each $H \in \mathcal{H}$ acts diagonalizably on $\mathbb{R}\mathrm{P}^3$, preserving a projective tetrahedron $T_H \subset \mathbb{R}\mathrm{P}^3$ and acting cocompactly on a properly embedded triangle $\Delta_H \subset \Omega$. Each edge of Δ_H is contained in a unique supporting hyperplane of Ω. The common intersection of these hyperplanes is the fourth vertex of T_H.

More explicitly, up to finite index, each H acts diagonally on $\mathbb{R}^4$ in the basis $\{v_1, v_2, v_3, v_4\}$, where the v_i are the vertices of T_H, and v_1, v_2, v_3 are the vertices of Δ_H. We can consider the situation where (in this basis) H is the discrete group

$$\left\{ \begin{pmatrix} 2^a & & & \\ & 2^b & & \\ & & 2^c & \\ & & & 1 \end{pmatrix} : a, b, c \in \mathbb{Z},\ a + b + c = 0 \right\}.$$

The dual of H preserves the corresponding dual basis $\{v_1^*, v_2^*, v_3^*, v_4^*\}$, and acts cocompactly on a projective triangle $\Delta_H^* \subset (\mathbb{RP}^3)^*$ with vertices v_1^*, v_2^*, v_3^*. The kernels $\mathbb{P}(v_i^*)$ for $i = 1, 2, 3$ give three supporting hyperplanes of Ω which cut out a region R_H of projective space containing Ω. In fact, R_H also contains $\partial\Omega - \partial\Delta_H$.

Now let $\rho^2 : \Gamma \to \mathrm{SL}(\mathrm{Sym}^2(\mathbb{R}^4)) \simeq \mathrm{SL}(10, \mathbb{R})$ be the composition of the inclusion $\Gamma \hookrightarrow \mathrm{SL}(4, \mathbb{R})$ with the symmetric square $\tau_2 : \mathrm{SL}(4, \mathbb{R}) \to \mathrm{SL}(\mathrm{Sym}^2(\mathbb{R}^4))$. In this case, the induced map $\iota : \mathbb{RP}^3 \to \mathbb{RP}^9$ is the Veronese embedding.

For each $H \in \mathcal{H}$, $\rho^2(H)$ preserves a 9-simplex in $\mathbb{RP}^9$, with vertices

$$\{v_1^2, v_2^2, v_3^2, v_4^2, v_1v_2, v_1v_3, v_1v_4, v_2v_3, v_2v_4, v_3v_4\}.$$

In particular $\rho^2(H)$ also preserves the 5-simplex S_H with vertices

$$\{v_1^2, v_2^2, v_3^2, v_1v_2, v_2v_3, v_1v_3\}.$$

Any divergent sequence in $\rho^2(H)$ always has an attracting subspace spanned by a face of S_H, since the eigenvalues of elements of H on v_4 are always dominated by some eigenvalue of that element on either v_1, v_2, or v_3. For instance, if we consider the sequence

$$a_n = \begin{pmatrix} 2^{2n} & & & \\ & 2^{-n} & & \\ & & 2^{-n} & \\ & & & 1 \end{pmatrix},$$

then $\rho^2(a_n)$ attracts towards the subspace spanned by $\{v_1^2\}$. On the other hand, if a_n is the sequence

$$\begin{pmatrix} 2^{n} & & & \\ & 2^{n} & & \\ & & 2^{-2n} & \\ & & & 1 \end{pmatrix},$$

then $\rho^2(a_n)$ attracts towards the subspace spanned by $\{v_1^2, v_2^2, v_1 v_2\}$.

Moreover, the subspaces spanned by faces of S_H are transverse to $\iota(\partial\Omega - \partial\Delta_H)$, so large elements of H attract points in $\iota(\partial\Omega)$ that are "far from" ∂S_H towards ∂S_H. In fact, this dynamical behavior extends to an entire open subset of $\mathbb{RP}^9$: namely, a region cut out by the hyperplanes corresponding to the vertices of the dual 5-simplex $S_H^* \subset (\mathbb{RP}^9)^*$ with vertices

$$\{(v_1^*)^2, (v_2^*)^2, (v_3^*)^2, v_1^* v_2^*, v_2^* v_3^*, v_1^* v_3^*\}.$$

So, the simplex S_H serves as the "parabolic point" for the action of the peripheral subgroup H on $\mathbb{RP}^9$—and moreover, this behavior is stable under perturbations of $\rho^2(H)$ in $\mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^2 V))$. This means that we can construct our candidate boundary set for the representation ρ^m by taking

$$\hat{\iota}(\hat{\Lambda}_\Omega(\Gamma)) \cup \bigcup_{H \in \mathcal{H}} \partial \hat{S}_H,$$

where $\partial \hat{S}_H \subset \mathcal{F}(\mathrm{Sym}^2 \mathbb{R}^4)$ is a closed subset of the space of flags projecting to the boundary of the simplex S_H.

6.3.3 Generalized weight spaces

To carry out the general construction of the simplex S_H identified in the previous example, we need some description of *attracting subspaces* for the groups $\rho^m(H) \subset \mathrm{SL}(\mathrm{Sym}^m V)$. We obtain this description by recalling some of the properties of *weights* of representations of free abelian groups.

Definition 6.3.4. Let $\rho : H \to \mathrm{GL}(V)$ be a representation of a free abelian group H, and let $\rho_{\mathbb{C}} : H \to \mathrm{GL}(V \otimes \mathbb{C})$ be the complexification of ρ.

A *complex weight* of ρ is a homomorphism $\mu_{\mathbb{C}} : H \to \mathbb{C}$ such that the *weight space*

$$V_{\mu_{\mathbb{C}}} = \ker(\rho_{\mathbb{C}}(h) - \exp(\mu_{\mathbb{C}}(h))I)$$

is nontrivial for some (any) $h \in H$. A *generalized complex weight* is similarly a homomorphism $\mu_{\mathbb{C}} : H \to \mathbb{C}$ such that the *generalized weight space*

$$V_{\mu_{\mathbb{C}}} = \bigcup_{n=1}^{\dim V} \ker(\rho_{\mathbb{C}}(h) - \exp(\mu_{\mathbb{C}}(h))I)^n$$

is nontrivial.

For any generalized weight $\mu_{\mathbb{C}}$, the *nilpotence degree* of $\mu_{\mathbb{C}}$ is the minimal $\ell \in \mathbb{N}$ such that $V_{\mu_{\mathbb{C}}} = \ker(\rho_{\mathbb{C}}(h) - \exp(\mu_{\mathbb{C}}(h))I)^{\ell}$.

Given a representation $\rho : H \to \mathrm{GL}(V)$, the generalized complex weight spaces of $\rho_{\mathbb{C}}$ give a $\rho_{\mathbb{C}}$-invariant decomposition of $V \otimes \mathbb{C}$. This in turn gives a ρ-invariant decomposition of V, since when $\mu_{\mathbb{C}}$ is a weight which takes on complex values, the direct sum $V_{\mu_{\mathbb{C}}} \oplus V_{\overline{\mu_{\mathbb{C}}}}$ is a ρ-invariant real subspace of V. By a slight abuse of terminology we refer to this as the *generalized weight space decomposition* for the representation ρ. The associated *real weights* of the representation are homomorphisms $\mu : H \to \mathbb{R}$ of the form $\log|\exp \mu_{\mathbb{C}}|$, where $\mu_{\mathbb{C}}$ is a (generalized) complex weight. For the rest of the section, unless otherwise indicated, when we refer to (generalized) *weights* of a representation into $\mathrm{SL}(V)$, we will mean the (generalized) *real* weights, and similarly for weight spaces.

Generalized weight spaces of ρ are stable under deformations of ρ. To be precise, we observe the following:

Proposition 6.3.5. *Let $\rho : \Gamma \to \mathrm{GL}(V)$ be a representation of a free abelian group, and let*

$$V^0_{\mu_1} \oplus \ldots \oplus V^0_{\mu_s}$$

be the generalized weight space decomposition of V for ρ. Let ρ_t be a continuous family of representations in $\mathrm{Hom}(\Gamma, \mathrm{GL}(V))$, with $\rho = \rho_0$.

For all sufficiently small t, there is a ρ_t-invariant decomposition

$$V^t_1 \oplus \ldots \oplus V^t_s$$

such that V^t_i is a sum of generalized weight spaces for ρ_t, with V^t_i varying continuously with t, and each of the weights associated to V^t_i also varying continuously with t.

Proof. The weights vary continously as a set with multiplicity, because the roots of the characteristic polynomial of $\rho(\gamma)$ vary continuously in ρ for fixed $\gamma \in \Gamma$. And, if μ is a complex weight with multiplicity k, then for small t there are complex weights

$\mu_1^t, \ldots, \mu_k^t$ of ρ_t (possibly with repeats) close to μ such that the sum of the kernels $\ker(\rho_{\mathbb{C}}(\gamma) - \exp(\mu_i^t(\gamma))I)^{\dim V}$ is close to $\ker(\rho_{\mathbb{C}}(\gamma) - \exp(\mu(\gamma))I)^{\dim V}$. $\square$

Definition 6.3.6. Let $\rho : H \to \mathrm{SL}(V)$ be a representation of a free abelian group H, and let Φ be the set of generalized weights of ρ. For any subset $\theta \subseteq \Phi$, we let $V_\theta \subseteq V$ denote the span of the generalized weight spaces V_μ for $\mu \in \theta$, and we let $V_\theta^{\mathrm{opp}} \subseteq V$ denote the span of the generalized weight spaces $V_{\mu'}$ for $\mu' \in \Phi - \theta$.

6.3.3.1 Faces in the convex hull of the weights

Whenever $\rho : H \to \mathrm{SL}(V)$ is a representation of a free abelian group with rank k, we can extend any real (generalized) weight $\mu : H \to \mathbb{R}$ to a homomorphism $\mu : H \otimes \mathbb{R} \to \mathbb{R}$, and view it as an element of $(\mathbb{R}^k)^* \simeq \mathbb{R}^k$.

Definition 6.3.7. Let $\rho : H \to \mathrm{SL}(V)$ be a representation of a free abelian group, and let Φ be the set of generalized weights of ρ. We denote the closed convex hull of Φ in $(H \otimes \mathbb{R})^* \simeq (\mathbb{R}^k)^*$ by $\mathcal{C}(\rho)$; since Φ is a finite subset of $(\mathbb{R}^k)^*$, $\mathcal{C}(\rho)$ is a convex polytope in $(\mathbb{R}^k)^*$.

The convex polytope $\mathcal{C}(\rho)$ is important for our purposes because it tells us how to find attracting subspaces for $\rho(H)$. In particular, there is a correspondence between the *faces* of $\mathcal{C}(\rho)$ and attracting subspaces of $\rho(H)$.

Definition 6.3.8. Let $\rho : H \to \mathrm{SL}(V)$ be a represention of a free abelian group with generalized weight set Φ. Let F be a closed face of $\mathcal{C}(\rho)$. We let $\Phi(F)$ denote the set of generalized weights of ρ lying in F.

For a face F of $\mathcal{C}(\rho)$, we write V_F and V_F^{opp} for $V_{\Phi(F)}$ and $V_{\Phi(F)}^{\mathrm{opp}}$, respectively.

Proposition 6.3.9. *Let $\rho : H \to \mathrm{SL}(V)$ be a representation of a free abelian group. For any divergent sequence $h_n \in H$, there is a face F of $\mathcal{C}(\rho)$ such that V_F and V_F^{opp} are respectively attracting and repelling subspaces for $\rho(h_n)$.*

Conversely, for any face F of $\mathcal{C}(\rho)$, V_F is an attracting subspace for some sequence $\rho(h_n)$ with $h_n \in H$ divergent.

To prove Proposition 6.3.9, we first establish some estimates which will later help us show that the convergence to the spaces V_F is *uniform* and *stable*. To make our estimates explicit, we choose an inner product on V, which induces a norm $||\cdot||$ on V and a smooth metric $d_{\mathbb{P}}$ on $\mathbb{P}(V)$. Specifically, for any transverse subspaces $W, W' \subset V$, we define

$$\angle(W, W') = \inf_{\substack{w \in W - \{0\}, \\ w' \in W' - \{0\}}} \angle(w, w'),$$

and then take $d_{\mathbb{P}}([u],[v]) = \sin(\angle([u],[v]))$. We also choose a norm $|\cdot|$ on $H \otimes \mathbb{R} \simeq \mathbb{R}^k$.

Lemma 6.3.10. *Given a representation $\rho : H \to \mathrm{SL}(V)$ of a free abelian group H, there exists a constant $D = D(\rho) > 0$ (varying continuously with ρ) satisfying the following: for any generalized weight μ of ρ with nilpotence degree ℓ, any $h \in H \otimes \mathbb{R}$, and any v in V_μ, we have*

$$\frac{1}{D}(\exp(\mu(h))||v|| \leq ||\rho(h)v|| \leq (\exp(\mu(h)) \cdot D|h|^{\ell-1})||v||.$$

Proof. For any $h \in H \otimes \mathbb{R}$, we let $I + N_\rho(h)$ be the unipotent part of the (multiplicative) Jordan-Chevalley decomposition of $\rho_{\mathbb{C}}(h)$. That is, if we fix a generalized weight μ, and let $\mu_{\mathbb{C}}$ be an associated complex weight, we can write

$$\rho_{\mathbb{C}}(h)v = \exp(\mu_{\mathbb{C}}(h))(I + N_\rho(h))v$$

for any $v \in V_{\mu_{\mathbb{C}}}$. Moreover, for any $v \in V_{\mu_{\mathbb{C}}}$, we have $N_\rho(h)^\ell v = 0$, where ℓ is the nilpotence degree of $\mu_{\mathbb{C}}$.

By considering $\rho_{\mathbb{C}}(t \cdot h)$, we see that for any fixed $h \in H \otimes \mathbb{R}$, $N_\rho(t \cdot h)$ is a polynomial in t of degree at most $\dim V$, with coefficients varying continuously with ρ and h, and a root at $t = 0$. Moreover, the restriction of $N_\rho(t \cdot h)$ to $V_{\mu_{\mathbb{C}}}$ is a polynomial in t with degree $\ell - 1$.

Then, since $\mathbb{P}(H \otimes \mathbb{R})$ is compact, we can find $D = D(\rho)$ (varying continuously with ρ) so that for any $h \in H \otimes \mathbb{R}$ and any $v \in V_\mu \otimes \mathbb{C}$,

$$\frac{1}{D}||v|| \le ||(I + N_\rho(h))v|| \le D \cdot |h|^{\ell-1}||v||,$$

where $||\cdot||$ is the norm on $V \otimes \mathbb{C}$ induced by our norm on V.

Then, for any $v \in V_\mu$, we have

$$\begin{aligned}
||\rho(h)v|| = ||\rho_{\mathbb{C}}(h)v|| &= ||\exp(\mu_{\mathbb{C}}(h)(I + N_\rho(h))v|| \\
&\le |\exp(\mu_{\mathbb{C}}(h))| \cdot D|h|^{\ell-1}||v|| \\
&= \exp(\mu(h)) \cdot D|h|^{\ell-1}||v||.
\end{aligned}$$

The left-hand inequality follows similarly. $\square$

The lemma below gives us a uniform estimate for the amount a group element $\rho(h)$ "attracts" towards a direct sum of weight spaces of ρ, in terms of the weights of ρ.

Lemma 6.3.11. *Fix a constant $M \geq 0$, and let $\rho : H \to \mathrm{SL}(V)$ be a representation of a free abelian group H. Let $D = D(\rho)$ be the constant coming from Lemma 6.3.10, and for any subset θ of the generalized weights Φ of ρ, let $\ell = \ell(\theta)$ be the maximum nilpotence degree of any weight in $\Phi - \theta$. Then, if $h \in H \otimes \mathbb{R}$ satisfies*

$$\mu(h) - \mu^{\mathrm{opp}}(h) \geq M|h|$$

for any $\mu \in \theta$ and $\mu^{\mathrm{opp}} \in \Phi - \theta$, we have

$$\frac{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta))}{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta^{\mathrm{opp}}))} \leq \frac{D^2|h|^{\ell-1}}{\exp(M|h|)} \cdot \frac{d_{\mathbb{P}}(x, \mathbb{P}(V_\theta^{\mathrm{opp}}))^{-1}}{\sin^2 \angle(V_\theta, V_\theta^{\mathrm{opp}})}.$$

for any $x \in \mathbb{P}(V) - \mathbb{P}(V_\theta^{\mathrm{opp}})$.

Proof. Fix a representation $\rho : H \to \mathrm{SL}(V)$ and a subset θ of the generalized weights of ρ, giving us a pair of complementary subspaces $V_\theta, V_\theta^{\mathrm{opp}} \subset V$. Let $x = [v] \in \mathbb{P}(V) - \mathbb{P}(V_\theta^{\mathrm{opp}})$. We can uniquely write $v = v_\theta + v_\theta^{\mathrm{opp}}$ for $v_\theta \in V_\theta$ and $v_\theta^{\mathrm{opp}} \in V_\theta^{\mathrm{opp}}$, with $v_\theta \neq 0$. Then we have

$$\frac{||v_\theta^{\mathrm{opp}}||}{||v||} \sin \angle(V_\theta, V_\theta^{\mathrm{opp}}) \leq d_{\mathbb{P}}(x, \mathbb{P}(V_\theta)) \leq \frac{||v_\theta^{\mathrm{opp}}||}{||v||},$$

where $||\cdot||$ is the norm on V induced by our choice of inner product. Similarly we have

$$\frac{||v_\theta||}{||v||} \sin \angle(V_\theta, V_\theta^{\mathrm{opp}}) \leq d_{\mathbb{P}}(x, \mathbb{P}(V_\theta^{\mathrm{opp}})) \leq \frac{||v_\theta||}{||v||}.$$

So in particular, we have

$$\sin(\angle(V_\theta, V_\theta^{\mathrm{opp}}))\frac{||v_\theta^{\mathrm{opp}}||}{||v_\theta||} \leq \frac{d_{\mathbb{P}}(x, \mathbb{P}(V_\theta))}{d_{\mathbb{P}}(x, \mathbb{P}(V_\theta^{\mathrm{opp}}))} \leq \frac{1}{\sin \angle(V_\theta, V_\theta^{\mathrm{opp}})} \cdot \frac{||v_\theta^{\mathrm{opp}}||}{||v_\theta||}. \tag{6.1}$$

Now let $h \in H \otimes \mathbb{R}$ satisfy $\mu(h) - \mu^{\mathrm{opp}}(h) \geq M|h|$ for all $\mu \in \theta$ and $\mu^{\mathrm{opp}} \in \Phi - \theta$. Since V_θ and V_θ^{opp} are $\rho(H)$-invariant, we have

$$\frac{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta))}{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta^{\mathrm{opp}}))} \leq \frac{1}{\sin \angle(V_\theta, V_\theta^{\mathrm{opp}})} \cdot \frac{||\rho(h)v_\theta^{\mathrm{opp}}||}{||\rho(h)v_\theta||}. \tag{6.2}$$

Since v_θ and v_θ^{opp} are respectively linear combinations of elements of V_μ for $\mu \in \theta$ and $V_{\mu^{\mathrm{opp}}}$ for $\mu^{\mathrm{opp}} \in \Phi - \theta$, we can apply Lemma 6.3.10 to $||\rho(h)v_\theta||$ and $||\rho(h)v_\theta^{\mathrm{opp}}||$ to see that

$$\frac{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta))}{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta^{\mathrm{opp}}))} \leq \frac{D^2|h|^{\ell-1}}{\exp(M|h|)} \cdot \frac{1}{\sin \angle(V_\theta, V_\theta^{\mathrm{opp}})} \frac{||v_\theta^{\mathrm{opp}}||}{||v_\theta||}.$$

Then, applying the left-hand inequality in (6.1) and using the fact that $d_{\mathbb{P}}(x, \mathbb{P}(V_\theta)) \leq 1$, we see that

$$\frac{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta))}{d_{\mathbb{P}}(\rho(h)x, \mathbb{P}(V_\theta^{\mathrm{opp}}))} \leq \frac{D^2|h|^{\ell-1}}{\exp(M|h|)} \cdot \frac{d_{\mathbb{P}}(x, \mathbb{P}(V_\theta^{\mathrm{opp}}))^{-1}}{\sin^2 \angle(V_\theta, V_\theta^{\mathrm{opp}})}.$$

as required. □

Proof of Proposition 6.3.9. Let $\rho : H \to \mathrm{SL}(V)$ be a representation of a free abelian group, let Φ be the set of generalized weights, and let h_n be a divergent sequence in H. Up to subsequence, the sequence $h_n/|h_n|$ converges to some $h_\infty \in H \otimes \mathbb{R}$ with $|h_\infty| = 1$.

We can view h_∞ as a linear functional on the space $(H \otimes \mathbb{R})^*$. Since $\mathcal{C}(\rho) \subset (H \otimes \mathbb{R})^*$ is a convex polytope, this means there is a face F of $\mathcal{C}$ so that for any $\mu \in \Phi(F)$ and any $\mu^{\mathrm{opp}} \in \Phi - \Phi(F)$, we have $\mu(h_\infty) > \mu^{\mathrm{opp}}(h_\infty)$. Then for sufficiently large n we also have $\mu(h_n/|h_n|) > \mu^{\mathrm{opp}}(h_n/|h_n|)$. In fact, since Φ is finite, there is a constant $M > 0$ such that $\mu(h_n) - \mu^{\mathrm{opp}}(h_n) > M|h_n|$ for every $\mu \in \Phi(F)$ and every

$\mu^{\mathrm{opp}} \in \Phi - \Phi(F)$. Applying Lemma 6.3.11 we see that for any $x \in \mathbb{P}(V) - \mathbb{P}(V_F^{\mathrm{opp}})$, the distance

$$d_{\mathbb{P}}(\rho(h_n)x, \mathbb{P}(V_F)) \leq \frac{d_{\mathbb{P}}(\rho(h_n)x, \mathbb{P}(V_F))}{d_{\mathbb{P}}(\rho(h_n)x, \mathbb{P}(V_F^{\mathrm{opp}}))}$$

tends to 0 as $n \to \infty$, so V_F and V_F^{opp} must respectively be attracting and repelling subspaces for $\rho(h_n)$.

Conversely, if F is any face of $\mathcal{C}(\rho)$, we can choose $h \in H \otimes \mathbb{R}$ so that $\mu(h) > 0$ and $\mu(h) > \mu^{\mathrm{opp}}(h)$ for any $\mu \in \Phi(F)$ and $\mu^{\mathrm{opp}} \in \Phi - \Phi(F)$. Then if $h_n \in H$ is any divergent sequence with $[h_n] \to [h]$ in the projectivization $\mathbb{P}(H \otimes \mathbb{R})$, another application of Lemma 6.3.11 shows that V_F is an attracting subspace for h_n. $\square$

6.3.4 Weights of peripheral subgroups in convex cocompact groups

For the rest of this section, we fix a relatively hyperbolic pair $(\Gamma, \mathcal{H})$, where each $H \in \mathcal{H}$ is virtually abelian with rank at least 2. We also fix a representation $\rho : \Gamma \to \mathrm{SL}(V)$ which is convex cocompact in $\mathbb{P}(V)$, and let $\Omega \subset \mathbb{P}(V)$ be a properly convex domain where $\rho(\Gamma)$ acts convex cocompactly.

Our goal now is to describe the convex polytope in $(H \otimes \mathbb{R})^*$ associated to the restriction of ρ to each $H \in \mathcal{H}$, which we can use to understand the dynamics of both $\rho(H)$ and $\rho^m(H)$.

Definition 6.3.12. For each $H \in \mathcal{H}$, we let $\mathcal{V}_H \subset \mathbb{P}(V)$ denote the set of vertices of Δ_H.

Proposition 6.3.13. *Let $H \in \mathcal{H}$ be a peripheral subgroup of rank $k \geq 2$, and let $H_0 \subseteq H$ be a finite-index free abelian subgroup. Consider the restriction $\rho_0 = \rho|_{H_0}$.*

Then, the convex polytope $\mathcal{C}(\rho_0)$ is a k-simplex in $(H_0 \otimes \mathbb{R})^*$, and each vertex of $\mathcal{C}(\rho_0)$ is a weight μ whose associated weight space is a vertex of Δ_H.

Moreover, every weight of ρ_0 which is not a vertex of $\mathcal{C}(\rho_0)$ lies in the interior of $\mathcal{C}(\rho_0)$.

Proof. Each vertex $v \in \mathcal{V}_H$ lies in a weight space of ρ_0, with an associated weight μ_v. Let Φ denote the weights of ρ_0, and let $\Phi(\mathcal{V}_H) \subseteq \Phi$ be the set of weights of the form μ_v for $v \in \mathcal{V}_H$. We claim that for any $\mu \in \Phi - \Phi(\mathcal{V}_H)$ and any $h \in H_0 \otimes \mathbb{R}$, there is a vertex $v \in \mathcal{V}_H$ such that

$$\mu_v(h) > \mu(h).$$

Suppose for a contradiction that the claim does not hold, i.e. there exists $h \in H_0 \otimes \mathbb{R}$ and $\mu \in \Phi - \Phi(\mathcal{V}_H)$ such that $\mu(h) \geq \mu_v(h)$ for all $v \in \mathcal{V}_H$. Let $W_H = \operatorname{span}(\Delta_H)$, and let V' be the ρ_0-invariant subspace $V_\mu \oplus W_H$. Then ρ_0 induces a representation $\rho_0' : H_0 \to \mathrm{SL}(V')$.

Each vertex $v \in \mathcal{V}_H$ is an eigenspace for ρ_0, so the nilpotence degree of μ_v (viewed as a weight of ρ_0') is 1. Then, we can apply Lemma 6.3.11 to ρ_0' (with $\theta = \{\mu\}$) to see that for any divergent sequence $h_n \in H$ with $[h_n] \to [h]$ in $\mathbb{P}(H_0 \otimes \mathbb{R})$, and any fixed $x \in \mathbb{P}(V') - \mathbb{P}(W_H)$, the ratio

$$\frac{d_{\mathbb{P}}(\rho(h_n)x, \mathbb{P}(V_\mu))}{d_{\mathbb{P}}(\rho(h_n)x, \mathbb{P}(W_H))}$$

does not tend to infinity as $n \to \infty$. In particular this is true for some $x \in \Omega$, since $\mathbb{P}(V') \cap \Omega$ is relatively open and nonempty. But since $\Delta_H \subset \mathbb{P}(W_H)$, this contradicts the fact that $\partial\Delta_H$ is the full orbital limit set of $\rho(H_0)$ in Ω.

We have now proved our claim, which implies that any extreme point of the convex polytope $\mathcal{C}(\rho_0)$ is a weight μ_v for $v \in \mathcal{V}_H$. On the other hand, by Corollary 6.1.4, we may assume that each vertex $v \in \mathcal{V}_H$ is an extreme point in $\partial\Omega$, and by Proposition 2.1.18, this means that for each $v \in \mathcal{V}_H$, there is a sequence $h_n \in H_0$ such that v is an attracting subspace for $\rho_0(h_n)$. Since $\rho_0(H_0)$ acts diagonalizably on W_H, this implies that μ_v is an extreme point of $\mathcal{C}(\rho_0)$. The last assertion of the proposition follows directly from the claim. $\square$

Proposition 6.3.13 tells us that we can combinatorially identify the k-simplex Δ_H and the k-simplex $\mathcal{C}(\rho_0)$ for $\rho_0 = \rho|_{H_0}$. We write this identification explicitly:

Definition 6.3.14. Let $H \in \mathcal{H}$, and let H_0 be a finite-index free abelian subgroup. For each face F of Δ_H with vertices $\mathcal{V}(F)$, we let $\tilde{F}$ denote the face of $\mathcal{C}(\rho_0)$ whose vertices are the weights μ_v for $v \in \mathcal{V}(F)$.

6.3.5 Invariant simplices in the symmetric power

Our next step is to describe the simplices $S_H \subset \mathbb{P}(\mathrm{Sym}^m V)$ which give rise to the fibers in $\hat{\Lambda}_m$ over parabolic points, for our EGF boundary extension $\hat{\Lambda}_m \to \partial(\Gamma, \mathcal{H})$.

Let $H \in \mathcal{H}$ have rank k, and let $H_0 \subseteq H$ be a finite-index free abelian subgroup. We let ρ_0, ρ_0^m respectively denote the restrictions of ρ, ρ^m to H_0, and let Φ, Φ^m denote the sets of weights of ρ_0 and ρ_0^m. We observe the following:

Lemma 6.3.15. *The convex polytope $\mathcal{C}(\rho_0^m)$ is the k-simplex $m\mathcal{C}(\rho_0)$. Moreover, for every face $\tilde{F}$ of $\mathcal{C}(\rho_0)$, the weights in $\Phi^m \cap m\tilde{F}$ are exactly the vertices of the*

mth barycentric subdivision of $m\tilde{F}$, and each such weight has a one-dimensional generalized weight space.

Proof. The weights of ρ_0^m are exactly the set of homomorphisms of the form

$$\sum_{\mu \in \Phi} a_\mu \mu,$$

where $a_\mu \in \mathbb{N} \cup \{0\}$ and $\sum a_\mu = m$. In particular, the set of rescaled weights $\frac{1}{m}\Phi^m$ consists entirely of convex combinations of weights of ρ_0, and contains every weight in Φ. This (together with Proposition 6.3.13) implies that $\mathcal{C}(\rho_0^m)$ is a k-simplex.

Further, every (rescaled) weight in the boundary of the rescaled simplex $\frac{1}{m}\mathcal{C}(\rho_0^m)$ must be a convex combination of weights lying in a single face of the simplex $\mathcal{C}(\rho_0)$. But Proposition 6.3.13 says that every weight in $\Phi \cap \partial\mathcal{C}(\rho_0)$ is a vertex of $\mathcal{C}(\rho_0)$. So, if F is a face of the simplex Δ_H with vertices $\mathcal{V}(F)$, the weights in $\tilde{F} \cap \frac{1}{m}\Phi^m$ are exactly the convex combinations of the form

$$\frac{1}{m} \sum_{v \in \mathcal{V}(F)} a_v \mu_v, \tag{6.3}$$

where $a_v \in \mathbb{N} \cup \{0\}$ and $\sum a_v = m$. These are exactly the vertices in the mth barycentric subdivision of $\tilde{F}$, and in fact each such vertex has *unique* expression of the form (6.3). Since each weight μ_v for $v \in \mathcal{V}(F)$ has a one-dimensional generalized weight space, it follows that the weights in $\Phi^m \cap mF$ do as well. □

6.3.5.1 The simplices $S_H \subset \mathbb{P}(\operatorname{Sym}^m V)$

Using Lemma 6.3.15, we can define the vertices of the simplex S_H: they are exactly the weight spaces for the weights μ lying in $\Phi^m \cap \partial\mathcal{C}(\rho_0^m)$.

To define S_H as a subset of $\mathbb{P}(\operatorname{Sym}^m V)$, we choose lifts in $\operatorname{Sym}^m V$ of each vertex of S_H, and then take convex combinations. Our lifts are chosen as follows: we first pick a lift $\tilde{v} \in V$ of each vertex $v \in \mathcal{V}_H$, so that Δ_H is the projectivization of the convex hull in V of $\{\tilde{v} : v \in \mathcal{V}_H\}$. The weight space of μ for each $\mu \in \partial\mathcal{C}(\rho_0^m)$ is spanned by a unique vector in $\operatorname{Sym}^m V$ of the form

$$\tilde{v}_\mu = \prod_{v \in \mathcal{V}_H} \tilde{v}^{a_v},$$

where $a_v \in \mathbb{N} \cup \{0\}$ and $\sum a_v = m$. Then we can define S_H to be the projectivization of the convex hull in $\operatorname{Sym}^m V$ of the $\tilde{v}_\mu$'s.

6.3.5.2 Dynamics on the simplices S_H

By definition, the vertices of S_H are exactly the weight spaces for the weights in the boundary of the simplex $\mathcal{C}(\rho_0^m) \subset (H_0 \otimes \mathbb{R})^*$. So, Proposition 6.3.9 immediately implies the following:

Corollary 6.3.16. *Let $H \in \mathcal{H}$. For every divergent sequence $h_n \in H$, there is a face F of S_H which spans an attracting subspace for the sequence $\rho^m(h_n)$.*

6.3.6 Dual simplices in symmetric powers

As discussed in Section 6.1, [DGK17], Proposition 5.6 says that since $\Gamma \to \mathrm{SL}(V)$ is convex cocompact in $\mathbb{P}(V)$, the dual representation $\Gamma \to \mathrm{SL}(V^*)$ is convex cocompact in $\mathbb{P}(V^*)$, and in fact there is a domain $\Omega \subset \mathbb{P}(V)$ so that Γ acts convex cocompactly on both Ω and the dual domain Ω^*. By the work of Islam-Zimmer [IZ19b], each virtually abelian subgroup $H \in \mathcal{H}$ must act cocompactly on a properly embedded *dual* simplex $\Delta_H^* \subset \Omega^*$. And, for each vertex w of Δ_H^*, the projective hyperplane $\mathbb{P}(w)$ is a supporting hyperplane of Ω at $\partial\Delta_H$.

6.3.6.1 The simplices $S_H^* \subset \mathbb{P}(\mathrm{Sym}^m V^*)$

For each $H \in \mathcal{H}$, we can define an H-invariant *dual* simplex $S_H^* \subset \mathbb{P}(\mathrm{Sym}^m V^*)$, by carrying out the construction we used to find S_H (but this time for the dual representation $\rho^* : \Gamma \to \mathrm{SL}(V^*)$). We can describe the relationship between the simplices S_H and S_H^* a little more explicitly. For a finite-index free abelian subgroup $H_0 \subseteq H$, we let $\rho_0^* : H_0 \to \mathrm{SL}(V^*)$ be the dual of the restriction of ρ to H_0, and similarly define $(\rho_0^m)^* : H_0 \to \mathrm{SL}(V^*)$. Then the weights of ρ_0^* are the negative weights of ρ_0, and the weights of $(\rho_0^m)^*$ are the negative weights of ρ_0^m.

Suppose μ^m is a weight of ρ_0^m with a one-dimensional generalized weight space v^m. Then, the negative weight $-\mu^m$ also has a one-dimensional generalized weight space $w^m \in \mathbb{P}(\mathrm{Sym}^m V^*)$, and $\mathbb{P}(w^m)$ is the hyperplane spanned by the weight spaces of the weights in $\Phi^m - \{\mu^m\}$. In particular, we can consider the case where μ^m is a weight lying in the boundary of $\mathcal{C}(\rho_0^m)$. In this case, v^m is a vertex of S_H, w^m is a vertex of S_H^*, and $\mathbb{P}(w^m)$ is a hyperplane intersecting S_H in a codimension-1 face of S_H.

This allows us to define a simultaneous lift of the *boundaries* of the simplices S_H, S_H^* in the space of flags $\mathcal{F}(\mathrm{Sym}^m V)$.

Definition 6.3.17. For a peripheral subgroup $H \in \mathcal{H}$, we let $\partial\hat{S}_H$ denote the set

$$\partial\hat{S}_H = \{(v, w) \in \mathcal{F}(\mathrm{Sym}^m(V)) : v \in \partial S_H, w \in \partial S_H^*\}.$$

The discussion above shows that $\partial \hat{S}_H$ is a nonempty closed invariant subset of $\mathcal{F}(\mathrm{Sym}^m V)$, projecting to ∂S_H and ∂S_H^* under the canonical maps $\mathcal{F}(\mathrm{Sym}^m V) \to \mathbb{P}(\mathrm{Sym}^m V)$ and $\mathcal{F}(\mathrm{Sym}^m V^*) \to \mathbb{P}(V^*)$.

6.3.7 Defining the boundary set

Using the sets $\partial \hat{S}_H$, we can define our candidate for the EGF boundary set $\hat{\Lambda}_m \subset \mathcal{F}(\mathrm{Sym}^m V)$ as follows. We let $\hat{\phi} : \hat{\Lambda}_\Omega(\Gamma) \to \partial(\Gamma, \mathcal{H})$ denote the boundary extension for the EGF representation ρ coming from Theorem 6.1.1. For each $z \in \partial(\Gamma, \mathcal{H})$, we define the set $\hat{\psi}_m(z) \subset \mathcal{F}(\mathrm{Sym}^m(V))$ by:

$$\hat{\psi}_m(z) = \begin{cases} \hat{\iota}(\hat{\phi}^{-1}(z)), & z \in \partial_{\mathrm{con}}(\Gamma, \mathcal{H}) \\ \partial \hat{S}_H, & z \in \partial_{\mathrm{par}}(\Gamma, \mathcal{H}). \end{cases}$$

We define

$$\hat{\Lambda}_m = \bigcup_{z \in \partial(\Gamma, \mathcal{H})} \hat{\psi}_m(z),$$

and observe that $\hat{\iota}(\hat{\Lambda}_\Omega(\Gamma)) \subset \hat{\Lambda}_m$.

The set $\hat{\Lambda}_m$ is $\rho^m(\Gamma)$-invariant, since $\hat{\iota}$ is τ_m-equivariant and the construction of the set $\partial \hat{S}_H$ is invariant. Ultimately we want to see that $\hat{\Lambda}_m$ is compact, and that there is a well-defined antipodal map $\hat{\phi}_m : \hat{\Lambda}_m \to \partial(\Gamma, \mathcal{H})$ giving us our EGF boundary extension.

6.3.8 Defining the boundary extension

Our next immediate goal is to show:

Proposition 6.3.18. *For distinct $z_1, z_2 \in \partial(\Gamma, \mathcal{H})$, the sets*

$$\hat{\psi}_m(z_1), \quad \hat{\psi}_m(z_2)$$

are opposite (in particular, disjoint). Consequently, the map $\hat{\phi}_m : \hat{\Lambda}_m \to \partial(\Gamma, \mathcal{H})$ *given by*

$$\hat{\phi}_m(\xi) = z \iff \xi \in \hat{\psi}_m(z)$$

is well-defined, equivariant, surjective, and antipodal.

Lemma 6.3.19. *Let* X *be a closed subset of* $\Lambda_\Omega(\Gamma)$*, and let* $\mathcal{H}_X \subset \mathcal{H}$ *be the set* $\{H \in \mathcal{H} : \partial\Delta_H \cap X \neq \emptyset\}$*.*

Then the set

$$X_{\mathcal{H}} = X \cup \bigcup_{H \in \mathcal{H}_X} \Delta_H$$

is closed.

Proof. Let x_n be a sequence in $X_{\mathcal{H}}$. By compactness of $\Lambda_\Omega(\Gamma)$, we can choose a subsequence so that $x_n \to x \in \Lambda_\Omega(\Gamma)$. We wish to show that $x \in X_{\mathcal{H}}$. Since X is closed and $X \subset X_{\mathcal{H}}$, we may assume that for each n, we have $x_n \in \partial\Delta_{H_n}$ for some $H_n \in \mathcal{H}_X$.

Up to subsequence, the sets $\partial\Delta_{H_n}$ converge to a closed set $\partial\Delta_\infty$ which is a connected finite union of (possibly degenerate) projective simplices. We must have $x \in \partial\Delta_\infty \subset \Lambda_\Omega(\Gamma)$. By definition, $\partial\Delta_{H_n}$ intersects X nontrivially, and since X is closed we must also have $\partial\Delta_\infty \cap X \neq \emptyset$.

Suppose for a contradiction that $x \notin X_{\mathcal{H}}$. Then in particular $x \notin X$. Since $\partial\Delta_\infty$ intersects X, it must contain at least two points, which means that every point in $\partial\Delta_\infty$ lies in a nontrivial closed projective segment (since $\partial\Delta_\infty$ is a connected finite union of projective simplices). But then by Corollary 6.1.4, $\partial\Delta_\infty \subset \partial\Delta_H$ for some $H \in \mathcal{H}$, and since $\partial\Delta_\infty \cap X \neq \emptyset$ we have $H \in \mathcal{H}_X$ and therefore $x \in X_{\mathcal{H}}$, contradiction. $\square$

Proposition 6.3.20. *For each $H \in \mathcal{H}$, there is a connected component C_H of*

$$\mathrm{Opp}(\partial S_H^*) = \{x \in \mathbb{P}(\mathrm{Sym}^m V) : x \perp w \textit{ for every } w \in \partial S_H^*\}$$

such that for every closed subset $X \subset \Lambda_\Omega(\Gamma) - \Delta_H$, C_H contains the closure of

$$\iota(X) \cup \bigcup_{H' \in \mathcal{H}_X} S_{H'},$$

where $\mathcal{H}_X = \{H \in \mathcal{H} : \partial\Delta_H \cap X \neq \emptyset\}$.

Proof. Let $\mathcal{V}_H$, $\mathcal{V}_H^*$ denote the vertex sets of Δ_H and Δ_H^*, respectively. Using the convexity of Ω, we can find lifts $\tilde{w} \in V^*$ for each vertex $w \in \mathcal{V}_H^*$, a continuous lift $\tilde{\Lambda} \subset V$ of $\Lambda_\Omega(\Gamma)$, and a continuous lift $\tilde{\Delta}_H \subset V$ of Δ_H so that

$$\tilde{w}(\tilde{\Lambda} - \tilde{\Delta}_H) > 0 \tag{6.4}$$

for every $w \in \mathcal{V}_H^*$.

The lifts $\tilde{w}$ induce lifts $\tilde{w}^m \in (\mathrm{Sym}^m V)^*$ of each vertex w^m of S_H^*. We take the set C_H to be the projectivization of

$$\{v \in \mathrm{Sym}^m V : \tilde{w}^m(v) > 0 \text{ for all vertices } w^m \text{ of } S_H^*\}.$$

Every point in ∂S_H^* is the projectivization of a convex combination of the lifts $\tilde{w}^m$. So, C_H is a connected component of $\mathbb{P}(\mathrm{Sym}^m V) - \bigcup_{w \in \partial(S_H^*)} \mathbb{P}(w)$.

Now let $X \subset \Lambda_\Omega(\Gamma) - \Delta_H$ be closed and let Y be the set

$$\iota(X) \cup \bigcup_{H' \in \mathcal{H}_X} S_{H'}.$$

We wish to show that $\overline{Y} \subset C_H$. Let $X_{\mathcal{H}}$ be the set

$$X_{\mathcal{H}} = X \cup \bigcup_{H' \in \mathcal{H}_X} \Delta_{H'}.$$

By Lemma 6.3.19, we can find a *compact* lift $\tilde{X}_{\mathcal{H}}$ of $X_{\mathcal{H}}$ in V so that $\tilde{w}(\tilde{x}) > 0$ for every $\tilde{x} \in \tilde{X}_{\mathcal{H}}$ and every $w \in \mathcal{V}_H^*$. We consider the set

$$\operatorname{Sym}^m \tilde{X}_{\mathcal{H}} = \{\tilde{x}^m \in \operatorname{Sym}^m V : \tilde{x}^m = \prod_{i=1}^{m} \tilde{x}_i \text{ for } \tilde{x}_i \in \tilde{X}_{\mathcal{H}}\}.$$

This set is the image of the m-fold Cartesian product $(\tilde{X}_{\mathcal{H}})^m$ under the continuous map $V^m \to \operatorname{Sym}^m V$ given by $(v_1, \ldots, v_m) \mapsto v_1 \cdots v_m$, so it is compact. Moreover, since $\operatorname{Sym}^m \tilde{X}_{\mathcal{H}}$ contains a lift of every vertex of every $S_{H'}$ for $H' \in \mathcal{H}_X$, the projectivization of the convex hull of $\operatorname{Sym}^m \tilde{X}_{\mathcal{H}}$ contains Y, hence $\overline{Y}$.

But from (6.4), we know that $\tilde{w}^m(\tilde{x}^m) > 0$ for every vertex w^m of S_H^* and every $\tilde{x}^m \in \tilde{X}_{\mathcal{H}}$, so we see that C_H contains the projectivization of the convex hull of $\operatorname{Sym}^m \tilde{X}_{\mathcal{H}}$. $\square$

Proof of Proposition 6.3.18. Let $z_1, z_2 \in \partial(\Gamma, \mathcal{H})$ be distinct. If both z_1 and z_2 are conical limit points, the proposition follows from the antipodality of the EGF boundary extension $\hat{\phi} : \hat{\Lambda}_\Omega(\Gamma) \to \partial(\Gamma, \mathcal{H})$ and the fact that $\hat{\iota}$ preserves transversality. On the other hand, if z_1 is a parabolic point, this follows from Proposition 6.3.20 (and the equivalent dual statement). $\square$

6.3.9 Dynamics on S_H

We have now defined an equivariant antipodal surjective map $\hat{\phi}_m : \hat{\Lambda}_m \to \partial(\Gamma, \mathcal{H})$, but we do not yet know that the set $\hat{\Lambda}_m$ is compact, or even that $\hat{\phi}_m$ is continuous. However, it turns out that it is easier to verify these two facts after proving that $\hat{\phi}_m$ has certain dynamical properties.

Lemma 6.3.21. *For each $H \in \mathcal{H}$, there exists an open set $\hat{C}_H \subset \mathcal{F}(\operatorname{Sym}^m V)$ containing $\hat{\Lambda}_m - \partial\hat{S}_H$, such that for any infinite sequence $h_n \in H$ and $\xi \in \hat{C}_H$, we have*

$$\rho^m(h_n)\xi \to \partial\hat{S}_H.$$

Proof. For each $H \in \mathcal{H}$, we let $C_H \subset \mathbb{P}(\operatorname{Sym}^m V)$ be the set coming from Proposition 6.3.20. Let h_n be a divergent sequence in some $H \in \mathcal{H}$, and let H_0 be a finite-index free abelian subgroup. Corollary 6.3.16 says that some face F of S_H spans an attracting subspace for $\rho^m(h_n)$. The corresponding repelling subspace is a direct sum of weight spaces for the restriction $\rho^m|_{H_0}$, so it is contained in $\mathbb{P}(w^m)$ for a vertex w^m of the dual simplex S^*_H. So, for any $x \in C_H$, any subsequence of $\rho^m(h_n)x$ subconverges to a point in $[\operatorname{span}(F)]$. In fact, $\rho^m(h_n)x$ subconverges to a point in $\overline{F} \subset \partial S_H$, since C_H is $\rho^m(H)$-invariant and $C_H \cap \operatorname{supp}(F) = \overline{F}$.

Then, we can dually define a set $C^*_H \subset \mathbb{P}(\operatorname{Sym}^m V^*)$, and take

$$\hat{C}_H = \{(x, w) \in \mathcal{F}(\operatorname{Sym}^m V) : x \in C_H, w \in C^*_H\}.$$

$\square$

6.3.10 Continuity and compactness

Lemma 6.3.22. *The set $\hat{\Lambda}_m$ is closed.*

Proof. Let (x_n, w_n) be a sequence in $\hat{\Lambda}_m$, and let $z_n = \hat{\phi}_m(x_n, w_n)$. Up to subsequence, z_n converges to $z \in \partial(\Gamma, \mathcal{H})$.

If z is a conical limit point, let γ_n be a sequence limiting conically to z, chosen so that for any $z' \neq z$, we have $\gamma_n^{-1} z' \to b$ and $\lim \gamma_n^{-1} z_n = a \neq b$.

Then $\hat{\phi}^{-1}(\gamma_n^{-1} z_n)$ converges to $\hat{\phi}^{-1}(a)$, and thus $\hat{\phi}^{-1}(\gamma_n^{-1} z_n)$ lies in a fixed compact subset X of $\mathrm{Opp}(\hat{\phi}^{-1}(b)) \cap \hat{\Lambda}_\Omega(\Gamma)$. By definition, this means that for every n, $\hat{\phi}_m^{-1}(\gamma_n^{-1} z_n)$ lies in the set

$$\iota(X) \cup \bigcup_{H' \in \mathcal{H}_X} S_{H'},$$

Arguing as in Proposition 6.3.20, we see that this set is compact. So by antipodality of $\hat{\phi}_m$, the sets of flags $\hat{\phi}_m^{-1}(\gamma_n^{-1} z_n)$ lie in a fixed compact subset of $\mathrm{Opp}(\hat{\phi}_m^{-1}(b)) = \mathrm{Opp}(\hat{\iota}(\hat{\phi}^{-1}(b)))$ and by Proposition 6.3.1,

$$(x_n, w_n) \in \rho^m(\gamma_n) \hat{\phi}_m^{-1}(\gamma_n^{-1} z_n)$$

converges to $\hat{\iota}(\hat{\phi}^{-1}(z))$.

If z is a parabolic point, we let $H = \mathrm{Stab}_\Gamma(z)$, and choose $h_n \in H$ so that $h_n^{-1} z_n \in K$ for a fixed compact $K - \{z\}$. By Proposition 6.3.20, we know that for all n, $\hat{\phi}_m^{-1}(h_n^{-1} z_n)$ lies in a fixed compact subset of $\hat{C}_H$. Then, Lemma 6.3.21 implies that

$$(x_n, w_n) \in \rho^m(h_n) \hat{\phi}_m^{-1}(h_n^{-1} z_n)$$

subconverges to a point in $\partial \hat{S}_H$. $\square$

Proposition 6.3.23. *The map $\hat{\phi}_m$ is continuous.*

Proof. Let (x_n, w_n) be a sequence in $\hat{\Lambda}_m$, converging to (x, w) (which we know lies in $\hat{\Lambda}_m$ by the previous proposition). Let $z_n = \hat{\phi}_m(x_n, w_n)$, and suppose for a contradiction that up to subsequence $z_n \to z$ for $z \neq \hat{\phi}_m(x, w)$.

Proposition 6.3.20 then implies that z_n lies in a compact subset $K \subset \partial(\Gamma, \mathcal{H})$ so that the closure of $\hat{\phi}_m^{-1}(K)$ is opposite to (x, w). This contradicts the fact that (x_n, w_n) converges to (x, w). $\square$

At this point, we have shown that $\hat{\phi}_m : \hat{\Lambda}_m \to \partial(\Gamma, \mathcal{H})$ is a continuous equivariant surjective antipodal map, and that $\hat{\Lambda}_m$ is a compact subset of $\mathcal{F}(\mathrm{Sym}^m V)$. So, we can finish the proof of Theorem 6.3.2 by showing:

Proposition 6.3.24. *The map $\hat{\phi}_m : \hat{\Lambda}_m \to \partial(\Gamma, \mathcal{H})$ extends the convergence dynamics of Γ.*

Proof. We will use Proposition 5.1.2, and show that $\hat{\phi}_m$ has the required dynamics on conical limit sequences and sequences in peripheral subgroups.

For each conical limit point $z \in \partial(\Gamma, \mathcal{H})$, we define $\hat{C}_z \subset \mathcal{F}(\mathrm{Sym}^m V)$ to be the affine chart

$$\hat{C}_z = \{\xi \in \mathcal{F}(\mathrm{Sym}^m V) : \xi \text{ is opposite to } \hat{\phi}_m^{-1}(z)\}.$$

If γ_n is a sequence limiting conically to z, and γ_n^{-1} limits conically to z_-, then $\hat{\phi}^{-1}(z)$, $\hat{\phi}^{-1}(z_-)$ are respectively attracting flags for $\rho(\gamma_n)$, $\rho(\gamma_n^{-1})$. It follows that $\hat{\iota}(\hat{\phi}^{-1}(z))$,

$\hat{\iota}(\hat{\phi}^{-1}(z_-))$ are respectively attracting flags for $\rho^m(\gamma_n)$ and $\rho^m(\gamma_n^{-1})$, so the required dynamics for the sequence are a consequence of Lemma 2.3.8.

On the other hand, for each parabolic point p, we take $\hat{C}_p$ to be the open set $\hat{C}_H$ considered in Lemma 6.3.21, for $H = \mathrm{Stab}_\Gamma(p)$. Lemma 6.3.21 implies that $\hat{C}_p$ contains $\hat{\Lambda}_m - \hat{\phi}_m^{-1}(p)$ and that for any $(x, w) \in \hat{C}_p$ and any infinite sequence $h_n \in H$, $\rho^m(h_n)\xi$ subconverges to a point in $\partial\hat{S}_H$. □

6.3.11 Stability

We have now shown that the representations ρ^m are all extended geometrically finite. Theorem 5.6.2 says that small peripherally stable deformations of ρ^m are also extended geometrically finite. So we want to show the following:

Proposition 6.3.25. *The space* $\mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^m V))$ *is peripherally stable at* $(\rho^m, \hat{\phi}_m)$.

The main step in the proof is the following:

Lemma 6.3.26. *Let H_0 be a finite-index free abelian subgroup of some $H \in \mathcal{H}$, with $H = \mathrm{Stab}_\Gamma(p)$. For any open set $U \subset \mathbb{P}(\mathrm{Sym}^m V)$ containing S_H and any $x \in C_p$, there exists a cofinite subset $T \subset H_0$ and an open set $\mathcal{W} \subset \mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^m V))$ containing ρ^m such that for any $\sigma \in \mathcal{W}$, we have $\sigma(h) \cdot x \in U$ for any $h \in T$.*

Proof. We proceed by contradiction, and suppose that there exists a sequence of distinct group elements $h_n \in H_0$ and a sequence of representations $\sigma_n : \Gamma \to \mathrm{SL}(\mathrm{Sym}^m V)$ such that $\sigma_n \to \rho^m$ and $\sigma_n(h_n)x \notin U$. We let Φ^m denote the set of generalized weights of $\rho^m|_{H_0}$, and we let Φ_n^m denote the generalized weights of $\sigma_n|_{H_0}$.

We choose a norm $|\cdot|$ on $H_0 \otimes \mathbb{R}$. Then up to subsequence $h_n/|h_n|$ converges to $h_\infty \in H_0 \otimes \mathbb{R}$ with $|h_\infty| = 1$.

Since Φ^m is finite, there is a face $\tilde{F}$ of the k-simplex $\mathcal{C}(\rho^m|_{H_0})$ and a constant $M > 0$, such that for every weight $\mu \in \Phi^m(F) = \Phi^m \cap \tilde{F}$, and every weight $\mu^{\mathrm{opp}} \in \Phi^m - \Phi(F)$, we have

$$\mu(h_\infty) - \mu^{\mathrm{opp}}(h_\infty) > M.$$

We let $V_F^m \subset \operatorname{Sym}^m V$ denote the span of the weight spaces of the weights in $\Phi^m(F)$; by definition $\mathbb{P}(V_F^m)$ is the projective span of a face of S_H.

Proposition 6.3.5 implies that as a set with multiplicity, the weights Φ_n^m converge to the weights Φ^m. So, for each n, there is a subset $\theta_n \subset \Phi_n^m$ such that θ_n converges to $\Phi^m \cap \partial\mathcal{C}(\rho^m|_{H_0})$, and a subset $\theta_n(F) \subset \theta_n$ such that $\theta_n(F)$ converges to $\Phi^m(F)$. Proposition 6.3.5 also implies that for sufficiently large n, all of the weights in θ_n must have one-dimensional generalized weight spaces, converging to the vertices of S_H.

This means that for each n, there are simplices S_H^n and $(S_H^n)^*$, invariant under the action of $\sigma_n(H)$, such that $S_H^n \to S_H$ and $(S_H^n)^* \to S_H^*$. So, we can find a sequence of group elements $g_n \in \mathrm{SL}(V)$, with g_n converging to the identity, so that $\sigma_n' = g_n \sigma_n g_n^{-1}$ preserves the simplices S_H and S_H^*. Moreover, the vertices of S_H are the weight spaces V_μ of σ_n' for $\mu \in \theta_n$, and the space V_F^m is spanned by weight spaces $V_{\mu(F)}$ of σ_n' for $\mu(F) \in \theta_n(F)$.

Now, since $\theta_n(F)$ converges to $\Phi^m(F)$, for sufficiently large n we must have

$$\mu_n(h_\infty) - \mu_n^{\mathrm{opp}}(h_\infty) > M$$

for every $\mu_n \in \theta_n(F)$ and every $\mu_n^{\mathrm{opp}} \in \Phi_n^m - \theta_n(F)$. This also means that for sufficiently large n we have

$$\mu_n(h_n) - \mu_n^{\mathrm{opp}}(h_n) > M|h_n|.$$

Then Lemma 6.3.11 implies that for sufficiently large n, $\sigma_n'(h_n)x$ lies in a small neighborhood of $\mathbb{P}(V_F)$ (here we use the fact that the constant $D = D(\rho)$ in Lemma 6.3.11 varies continuously with ρ, and that the nilpotence degree ℓ of any weight of σ_n' is bounded by $\dim \operatorname{Sym}^m V$). Moreover, we know that the set C_p is $\sigma_n'(H)$-invariant, since the simplex S_H^* is $\sigma_n'(H)$-invariant. Since x lies in C_p, $\sigma_n'(h_n)x$ lies in an arbitrarily small neighborhood of $\mathbb{P}(V_F) \cap C_p$. By definition this intersection is a face of S_H, so for large enough n, $\sigma_n(h_n)x$ must lie in an arbitrarily small neighborhood of this face, giving a contradiction. $\square$

Proof of Proposition 6.3.25. We want to show that if $H = \operatorname{Stab}_\Gamma(p)$ for a parabolic point p, K is a compact subset of $\hat{C}_p$, U is a neighborhood of $\hat{S}_H$, and $T \subset H$ is a cofinite subset such that

$$\rho^m(T) \cdot K \subset U, \tag{6.5}$$

then we can find a neighborhood $\mathcal{W}$ of $\rho^m|_H$ in $\operatorname{Hom}(H, \operatorname{SL}(\operatorname{Sym}^m V))$ so that for any $\sigma \in \mathcal{W}$,

$$\sigma(T) \cdot K \subset U. \tag{6.6}$$

For simplicity, we will not work in the space of flags $\mathcal{F}(\operatorname{Sym}^m V)$. Instead we will just show that that if (6.5) holds for a compact $K \subset C_p \subset \mathbb{P}(\operatorname{Sym}^m V)$ and an open neighborhood U of S_H in $\mathbb{P}(\operatorname{Sym}^m V)$, then (6.6) holds also.

Fix a finite-index free abelian subgroup $H_0 \subseteq H$. It suffices to show that we can choose an open $\mathcal{W} \subset \mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^m V))$ so that

$$\sigma(T \cap H_0) \cdot K \subset U$$

for all $\sigma \in \mathcal{W}$. Using Lemma 6.3.26, we can find a cofinite set $T' \subset H_0$ and an open set $\mathcal{W} \subset \mathrm{Hom}(\Gamma, \mathrm{SL}(\mathrm{Sym}^m V))$ so that for all $\sigma \in \mathcal{W}'$, we have

$$\sigma(T') \cdot K \subset U.$$

But then since $(T \cap H_0) - T'$ is finite, we can just shrink $\mathcal{W}$ to get the desired result. $\square$

Appendix

Let V be a real vector space, and let A_n be a sequence of matrices in $\mathrm{PGL}(V)$. It is sometimes possible to determine the global dynamical behavior of A_n on $\mathbb{P}(V)$ by considering the action of A_n on a small open subset of $\mathbb{P}(V)$: if there is an open subset $U \subset \mathbb{P}(V)$ such that $A_n \cdot U$ converges to a point in $\mathbb{P}(V)$, then in fact there is a dense open subset $U_- \subset \mathbb{P}(V)$ (the complement of a hyperplane) on which A_n converges to the same point, uniformly on compacts.

A similar statement holds for the action of A_n on Grassmannians $\mathrm{Gr}(k, V)$. These claims can be proved by considering the behavior of the singular value gaps of A_n as $n \to \infty$.

In this appendix we give a general result along these lines, where we take sequences of group elements $g_n \in G$ for a semisimple Lie group G with no compact factor and trivial center, and consider the limiting behavior of g_n on open subsets of some flag manifold G/P^+, where P^+ is a parabolic subgroup.

Proposition 2.3.7. *Let g_n be a sequence in G, and suppose that for some nonempty open subset $U \subset G/P^+$, we have $g_n \cdot U \to \{\xi\}$ for $\xi \in G/P^+$. Then g_n is P^+-divergent, and has a unique P^+-limit point $\xi \in G/P^+$.*

We will prove Proposition 2.3.7 by reducing it to the case where $G = \mathrm{PGL}(d, \mathbb{R})$ and $P^+ = P_1$ is the stabilizer of $[e_1] \in \mathbb{RP}^{d-1} \simeq G/P_1$. In this situation, P^+-divergence can be understood in terms of the behavior of the *singular value gaps* of the sequence g_n.

Proposition 0.0.27. *Suppose that $G = \mathrm{PGL}(d, \mathbb{R})$, and let $P^+ = P_1 \subset G$ be the stabilizer of a line in $\mathbb{R}^d$. A sequence $g_n \in G$ is P_1-divergent if and only if*

$$\frac{\sigma_1(g_n)}{\sigma_2(g_n)} \to \infty,$$

where $\sigma_i(g_n)$ is the ith-largest singular value of g_n.

For convenience, we give a proof of Proposition 2.3.7 in this special case.

Lemma 0.0.28. *Let g_n be a sequence in $\mathrm{PGL}(d, \mathbb{R})$, and suppose that for a nonempty open subset $U \subset \mathbb{RP}^{d-1}$, g_nU converges to a point in $\mathbb{RP}^{d-1}$. Then, the singular value gap*

$$\frac{\sigma_1(g_n)}{\sigma_2(g_n)}$$

tends to ∞ as $n \to \infty$.

Proof. It suffices to show that any subsequence of g_n has a subsequence which satisfies the property. Using the Cartan decomposition of $\mathrm{PGL}(d, \mathbb{R})$, we can write

$$g_n = k_n a_n k_n',$$

for $k_n, k_n' \in K = \mathrm{PO}(d)$ and a_n a diagonal matrix whose diagonal entries are $\sigma_1, \ldots, \sigma_d$. Up to subsequence k_n and k_n' converge respectively to $k, k' \in K$. For sufficiently large

n, $k'_n U \cap k'U$ is nonempty, so by replacing U with $k'U$ we can assume that $k'_n = \mathrm{id}$ for all n. Furthermore, if $k_n a_n U$ converges to a point $z \in \mathbb{R}\mathrm{P}^{d-1}$, then $a_n U$ converges to $k^{-1}z$.

So, $a_n U$ converges to a point, and since a_n is a diagonal matrix, the gap between the moduli of its largest and second-largest eigenvalues must be unbounded. $\square$

To prove the general case of Proposition 2.3.7, we take an irreducible representation $\zeta : G \to \mathrm{PGL}(V)$ coming from Theorem 5.4.4, so that P^+ maps to the stabilizer of a point p in $\mathbb{P}(V)$, P^- maps to the stabilizer of a dual point $q \in \mathbb{P}(V^*)$, and gP^+g^{-1}, hP^-h^{-1} are opposite if and only if $\zeta(g)p$, $\zeta(h)q$ are transverse. As in section 5.4, this determines embeddings $\iota : G/P \to \mathbb{P}(V)$ and $\iota^* : G/P^- \to \mathbb{P}(V^*)$ by

$$\iota(gP^+) = \zeta(g)p, \qquad \iota^*(gP^-) = \zeta(g)q.$$

The representation ζ additionally has the property that for any sequence $g_n \in G$, the singular value gaps

$$\sigma_1(\zeta(g_n))/\sigma_2(\zeta(g_n))$$

are unbounded if and only if g_n is P^+-divergent (see [GGKW17], Lemma 3.7).

Proof of Proposition 2.3.7. By [Zim18], Lemma 4.7, there exist flags $\xi_1, \ldots, \xi_D \in U$ so that lifts of $\iota(\xi_i)$ give a basis of V. Since $g_n \cdot U$ converges to a point in G/P, the set

$$\{\zeta(g_n) \cdot \iota(\xi_i) : 1 \leq i \leq D\}$$

converges to a single point in $\mathbb{P}(V)$.

This means that we can fix lifts $\iota(\tilde{\xi}_i) \in V$ so that, up to a subsequence, $\zeta(g_n)$ takes the projective $(D-1)$-simplex

$$\left[\sum_{i=1}^{D} \lambda_i \iota(\tilde{\xi}_i) : \lambda_i > 0\right] \subset \mathbb{P}(V)$$

to a point. This simplex is an open subset of $\mathbb{P}(V)$. Now we can apply Lemma 0.0.28 to see that the sequence g_n is P^+-divergent.

We now just need to check that ξ is the unique P^+-limit point of g_n. Choose any subsequence of g_n. Then any P^+-contracting subsequence g_m of this subsequence satisfies

$$g_m|_{\mathrm{Opp}(\xi_-)} \to \xi'$$

uniformly on compacts for some $\xi_- \in G/P^-$ and $\xi' \in G/P^+$. But since $\mathrm{Opp}(\xi_-)$ is open and dense, it intersects U nontrivially and thus $\xi' = \xi$. $\square$

Bibliography

[Bal21] Samuel A. Ballas. Constructing convex projective 3-manifolds with generalized cusps. *J. Lond. Math. Soc. (2)*, 103(4):1276–1313, 2021.

[BCL20] Samuel A. Ballas, Daryl Cooper, and Arielle Leitner. Generalized cusps in real projective manifolds: classification. *J. Topol.*, 13(4):1455–1496, 2020.

[BDL15] Samuel Ballas, Jeffrey Danciger, and Gye-Seon Lee. Convex projective structures on non-hyperbolic three-manifolds. *Geometry & Topology*, 22, 08 2015.

[Ben60] Jean-Paul Benzécri. Sur les variétés localement affines et localement projectives. *Bull. Soc. Math. France*, 88:229–332, 1960.

[Ben00] Yves Benoist. Automorphismes des cônes convexes. *Invent. Math.*, 141(1):149–193, 2000.

[Ben04] Yves Benoist. Convexes divisibles. I. In *Algebraic groups and arithmetic*, pages 339–374. Tata Inst. Fund. Res., Mumbai, 2004.

[Ben06] Yves Benoist. Convexes divisibles. IV. Structure du bord en dimension 3. *Invent. Math.*, 164(2):249–278, 2006.

[Ben08] Yves Benoist. A survey on divisible convex sets. In *Geometry, analysis and topology of discrete groups*, volume 6 of *Adv. Lect. Math. (ALM)*, pages 1–18. Int. Press, Somerville, MA, 2008.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BM20] Samuel Ballas and Ludovic Marquis. Properly convex bending of hyperbolic manifolds. *Groups Geom. Dyn.*, 14(2):653–688, 2020.

[Bob19] Martin D. Bobb. Convex projective manifolds with a cusp of any non-diagonalizable type. *J. Lond. Math. Soc. (2)*, 100(1):183–202, 2019.

[Bob20] Martin D. Bobb. Codimension-1 Simplices in Divisible Convex Domains. *arXiv e-prints*, page arXiv:2001.11096, January 2020.

[Bog97] O. V. Bogopol'skiĭ. Infinite commensurable hyperbolic groups are bi-Lipschitz equivalent. *Algebra i Logika*, 36(3):259–272, 357, 1997.

[Bou95] Marc Bourdon. Structure conforme au bord et flot géodésique d'un CAT(−1)-espace. *Enseign. Math. (2)*, 41(1-2):63–102, 1995.

[Bow93] B. H. Bowditch. Geometrical finiteness for hyperbolic groups. *J. Funct. Anal.*, 113(2):245–317, 1993.

[Bow95] B. H. Bowditch. Geometrical finiteness with variable negative curvature. *Duke Math. J.*, 77(1):229–274, 1995.

[Bow98] Brian H. Bowditch. Spaces of geometrically finite representations. *Ann. Acad. Sci. Fenn. Math.*, 23(2):389–414, 1998.

[Bow99] B. H. Bowditch. Convergence groups and configuration spaces. In *Geometric group theory down under (Canberra, 1996)*, pages 23–54. de Gruyter, Berlin, 1999.

[Bow12] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.

[BPS19] Jairo Bochi, Rafael Potrie, and Andrés Sambarino. Anosov representations and dominated splittings. *J. Eur. Math. Soc. (JEMS)*, 21(11):3343–3414, 2019.

[Can21] Richard Canary. Informal lecture notes on Anosov representations. Available at http://www.math.lsa.umich.edu/ canary/lecnotespublic.pdf, 2021.

[CG93] Suhyoung Choi and William M. Goldman. Convex real projective structures on closed surfaces are closed. *Proc. Amer. Math. Soc.*, 118(2):657–661, 1993.

[CLM20] Suhyoung Choi, Gye-Seon Lee, and Ludovic Marquis. Convex projective generalized Dehn filling. *Ann. Sci. Éc. Norm. Supér. (4)*, 53(1):217–266, 2020.

[CLT15] D. Cooper, D. D. Long, and S. Tillmann. On convex projective manifolds and cusps. *Adv. Math.*, 277:181–251, 2015.

[CLT18] Daryl Cooper, Darren Long, and Stephan Tillmann. Deforming convex projective manifolds. *Geom. Topol.*, 22(3):1349–1404, 2018.

[CM14] Mickaël Crampon and Ludovic Marquis. Finitude géométrique en géométrie de Hilbert. *Ann. Inst. Fourier (Grenoble)*, 64(6):2299–2377, 2014.

[Coo90] Michel Coornaert. *Sur les groupes proprement discontinus d'isométries des espaces hyperboliques au sens de Gromov*, volume 444 of *Publication de l'Institut de Recherche Mathématique Avancée [Publication of the Institute of Advanced Mathematical Research]*. Université Louis Pasteur, Département de Mathématique, Institut de Recherche Mathématique Avancée, Strasbourg, 1990. Dissertation, Université Louis Pasteur, Strasbourg, 1990.

[CZZ21] Richard Canary, Tengren Zhang, and Andrew Zimmer. Cusped Hitchin representations and Anosov representations of geometrically finite Fuchsian groups. *arXiv e-prints*, page arXiv:2103.06588, March 2021.

[DGK17] Jeffrey Danciger, François Guéritaud, and Fanny Kassel. Convex cocompact actions in real projective geometry. *arXiv e-prints*, page arXiv:1704.08711, April 2017.

[DGK+21] Jeffrey Danciger, François Guéritaud, Fanny Kassel, Gye-Seon Lee, and Ludovic Marquis. Convex cocompactness for Coxeter groups. *arXiv e-prints*, page arXiv:2102.02757, February 2021.

[DS05] Cornelia Druţu and Mark Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by Denis Osin and Mark Sapir.

[Ebe96] Patrick B. Eberlein. *Geometry of nonpositively curved manifolds*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1996.

[Ger12] Victor Gerasimov. Floyd maps for relatively hyperbolic groups. *Geom. Funct. Anal.*, 22(5):1361–1399, 2012.

[GGKW17] François Guéritaud, Olivier Guichard, Fanny Kassel, and Anna Wienhard. Anosov representations and proper actions. *Geom. Topol.*, 21(1):485–584, 2017.

[GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.

[Gol88] William M. Goldman. Projective geometry on manifolds, 1988.

[Gol90] William M. Goldman. Convex real projective structures on compact surfaces. *J. Differential Geom.*, 31(3):791–845, 1990.

[GP13] Victor Gerasimov and Leonid Potyagailo. Quasi-isometric maps and Floyd boundaries of relatively hyperbolic groups. *J. Eur. Math. Soc. (JEMS)*, 15(6):2115–2137, 2013.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[GW12] Olivier Guichard and Anna Wienhard. Anosov representations: domains of discontinuity and applications. *Invent. Math.*, 190(2):357–438, 2012.

[Hel01] Sigurdur Helgason. *Differential geometry, Lie groups, and symmetric spaces*, volume 34 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001. Corrected reprint of the 1978 original.

[IZ19a] Mitul Islam and Andrew Zimmer. A flat torus theorem for convex co-compact actions of projective linear groups. *arXiv e-prints*, page arXiv:1907.03277, July 2019.

[IZ19b] Mitul Islam and Andrew Zimmer. Convex co-compact actions of relatively hyperbolic groups. *arXiv e-prints*, page arXiv:1910.08885, October 2019.

[IZ20] Mitul Islam and Andrew Zimmer. Convex co-compact representations of 3-manifold groups. *arXiv e-prints*, page arXiv:2009.05191, September 2020.

[IZ22] Mitul Islam and Andrew Zimmer. The structure of relatively hyperbolic groups in convex real projective geometry. *arXiv e-prints*, page arXiv:2203.16596, March 2022.

[Kas18] Fanny Kassel. Geometric structures and representations of discrete groups. *arXiv e-prints*, page arXiv:1802.07221, February 2018.

[KKL19] Michael Kapovich, Sungwoon Kim, and Jaejeong Lee. Structural stability of meandering-hyperbolic group actions. *arXiv e-prints*, page arXiv:1904.06921, April 2019.

[KL06] Bruce Kleiner and Bernhard Leeb. Rigidity of invariant convex sets in symmetric spaces. *Invent. Math.*, 163(3):657–676, 2006.

[KL18] Michael Kapovich and Bernhard Leeb. Relativizing characterizations of Anosov subgroups, I. *arXiv e-prints*, page arXiv:1807.00160, June 2018.

[KLP14] Michael Kapovich, Bernhard Leeb, and Joan Porti. Morse actions of discrete groups on symmetric space. *arXiv e-prints*, page arXiv:1403.7671, March 2014.

[KLP17] Michael Kapovich, Bernhard Leeb, and Joan Porti. Anosov subgroups: dynamical and geometric characterizations. *Eur. J. Math.*, 3(4):808–898, 2017.

[KLP18] Michael Kapovich, Bernhard Leeb, and Joan Porti. Dynamics on flag manifolds: domains of proper discontinuity and cocompactness. *Geom. Topol.*, 22(1):157–234, 2018.

[Kna02] Anthony W. Knapp. *Lie groups beyond an introduction*, volume 140 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, second edition, 2002.

[KS58] Paul Kelly and Ernst Straus. Curvature in Hilbert geometries. *Pacific J. Math.*, 8:119–125, 1958.

[Lab06] François Labourie. Anosov flows, surface groups and curves in projective space. *Invent. Math.*, 165(1):51–114, 2006.

[Mar14] Ludovic Marquis. Around groups in Hilbert geometry. In *Handbook of Hilbert geometry*, volume 22 of *IRMA Lect. Math. Theor. Phys.*, pages 207–261. Eur. Math. Soc., Zürich, 2014.

[Max21] J. Maxwell Riestenberg. A quantified local-to-global principle for Morse quasigeodesics. *arXiv e-prints*, page arXiv:2101.07162, January 2021.

[Qui05] J.-F. Quint. Groupes convexes cocompacts en rang supérieur. *Geom. Dedicata*, 113:1–19, 2005.

[Ser81] Caroline Series. Symbolic dynamics for geodesic flows. *Acta Math.*, 146(1-2):103–128, 1981.

[Sul79] Dennis Sullivan. The density at infinity of a discrete group of hyperbolic motions. *Inst. Hautes Études Sci. Publ. Math.*, (50):171–202, 1979.

[Sul85] Dennis Sullivan. Quasiconformal homeomorphisms and dynamics. II. Structural stability implies hyperbolicity for Kleinian groups. *Acta Math.*, 155(3-4):243–260, 1985.

[Tra13] Hung Cong Tran. Relations between various boundaries of relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 23(7):1551–1572, 2013.

[Tso20] Konstantinos Tsouvalas. Anosov representations, strongly convex cocompact groups and weak eigenvalue gaps. *arXiv e-prints*, page arXiv:2008.04462, August 2020.

[Tuk94] Pekka Tukia. Convergence groups and Gromov's metric hyperbolic spaces. *New Zealand J. Math.*, 23(2):157–187, 1994.

[Tuk98] Pekka Tukia. Conical limit points and uniform convergence groups. *J. Reine Angew. Math.*, 501:71–98, 1998.

[Wei20] Theodore Weisman. Dynamical properties of convex cocompact actions in projective space. *arXiv e-prints*, page arXiv:2009.10994, September 2020.

[Wei22] Theodore Weisman. An extended definition of Anosov representation for relatively hyperbolic groups. *arXiv e-prints*, page arXiv:2205.07183, May 2022.

[Yam04] Asli Yaman. A topological characterisation of relatively hyperbolic groups. *J. Reine Angew. Math.*, 566:41–89, 2004.

[Zhu19] Feng Zhu. Relatively dominated representations. *arXiv e-prints*, page arXiv:1912.13152, December 2019.

[Zhu21] Feng Zhu. Relatively dominated representations from eigenvalue gaps and limit maps. *arXiv e-prints*, page arXiv:2102.10611, February 2021.

[Zim18] Andrew M. Zimmer. Proper quasi-homogeneous domains in flag manifolds and geometric structures. *Ann. Inst. Fourier (Grenoble)*, 68(6):2635–2662, 2018.

[Zim21] Andrew Zimmer. Projective Anosov representations, convex cocompact actions, and rigidity. *J. Differential Geom.*, 119(3):513–586, 2021.

Index