Morse, contracting, and strongly contracting sets with applications to boundaries and growth of groups

Christopher H. Cashen, PhD

Faculty of Mathematics
University of Vienna
1090 Vienna, Austria
E-mail address: christopher.cashen@univie.ac.at
URL: http://www.mat.univie.ac.at/~cashen

Key words and phrases. Morse, contracting, strongly contracting, graphical small cancellation, group action, boundary, growth tightness, cogrowth

The author was supported by the Austrian Science Fund (FWF): P 30487-N35 during the writing of this thesis. Support for the individual papers can be found in the respective chapters.

Abstract. We investigate several quantitative generalizations of the notion of quasiconvex subsets of (Gromov) hyperbolic spaces to arbitrary geodesic metric spaces. Some of these, such as the Morse property, strong contraction, and superlinear divergence, had been studied before in more specialized contexts, and some, such as contraction, we introduce for the first time. In general, we prove that quasiconvexity is the weakest of the properties, strong contraction is the strongest, and all of the others are equivalent. However, in hyperbolic spaces all are equivalent, and we prove that in CAT(0) spaces all except quasiconvexity are equivalent.

Despite the fact that many of these properties are equivalent, they are useful for different purposes. For instance, it is easy to see that the Morse property is quasiisometry invariant, but the contraction property gives good control over the divagation behavior of geodesic rays with a common basepoint. We exploit this control to define a boundary for arbitrary finitely generated groups that shares some properties of the boundary of a hyperbolic group. Our boundary is a metrizable topological space that is invariant under quasiisometries of the group, and the group acts on it with simple dynamics.

We investigate the geometry of infinitely presented graphical small cancellation groups. Such groups include the so-called 'Gromov Monsters', which were introduced as a source of counter-examples to the Baum-Connes conjecture. We give a local-to-global characterization of contracting geodesics in these groups, which we think of as defining 'hyperbolic directions'. Our characterization depends on a beautiful interplay between combinatorial and geometric versions of negative curvature. The result shows that the geometry of these groups is reminiscent of the geometry of relatively hyperbolic groups in the sense that there are certain well-defined non-hyperbolic regions, and geodesics that avoid these regions behave like hyperbolic geodesics. However, the groups are in general not relatively hyperbolic.

Armed with our understanding of geodesics in graphical small cancellation groups, we construct examples of the wide varieties of contraction behaviors that occur: we show that every degree of contraction can be achieved by a periodic geodesic in some finitely generated group, that there are groups in which every element has a strongly contracting axis even though the group is not hyperbolic, and that there are examples of finitely generated groups in which the existence of a strongly contracting axis for a given element depends on the choice of generating set for the group.

Since the Morse property is invariant under quasiisometries, we can say that a subgroup of a finitely generated group is Morse (or equivalently, contracting, divergent,...) if it has this property as a subset in some/any Cayley graph of the group. However, we could also let the group act on some other geodesic metric space and ask which elements have a contracting/Morse/strongly contracting axis for that particular action. In particular, we explore actions that are not cocompact. To preserve a connection with the geometry of the group, we require the action to be metrically proper. We also introduce a condition known as 'complementary growth gap' that says that there is an orbit of the group in the space that, while metrically distorted, is not too badly distorted from a growth-theoretic point of view. Our condition generalizes the 'parabolic growth gap' condition for Kleinian groups, and includes additional examples such as the action of the mapping class group of a hyperbolic surface on its Teichmüller space. We prove growth and cogrowth results for the orbit pseudometric induced on the group by such an action under the hypothesis the group has one element that acts with a strongly contracting axis. Our results generalize results that were known for word metrics on hyperbolic groups to far more general situations. The generality of our results is even more striking considering that there are contemporaneous papers that achieve similar results only for actions on hyperbolic spaces.

Contents

Preface

I. Plan of the thesis

This is a cumulative Habilitation thesis on the topic of generalizations of the geometric notions of convexity and hyperbolicity and their applications to the study of infinite discrete groups. The thesis includes eight papers on this topic. This Preface contains a brief Synopsis of each paper including its abstract, publication details, and a list of works that cite it, as well as a CV. The following chapter is a Survey of the research topic that gives an exposition of our results, places them into context, and demonstrates the impact of our research on its field. Each of the remaining eight chapters contains one of the papers.

Each of the papers included in this thesis has undergone successful peer review and is either already published or accepted and in press. In all cases the version of the paper presented here is the final author version, post-refereeing, subjected to reformatting in a consistent style. The final version of record of each paper can be found at the listed address on the website of its publisher, and may have formatting and editorial changes.

II. Synopsis of the Publications

A. *Characterizations of Morse quasi-geodesics via superlinear divergence and sublinear contraction*,

Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume,

Documenta Mathematica **22** (2017), 1193–1224.

`https://www.math.uni-bielefeld.de/documenta/vol-22/36.html`

We introduce and begin a systematic study of sublinearly contracting projections. We give two characterizations of Morse quasi-geodesics in an arbitrary geodesic metric space. One is that they are sublinearly contracting; the other is that they have completely superlinear divergence.

We give a further characterization of sublinearly contracting projections in terms of projections of geodesic segments.

The authors were equal partners in the conception, execution, and writing of this paper.

Works citing Paper A.

[1] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham. *Largest acylindrical actions and stability in hierarchically hyperbolic groups*. preprint. 2017. arXiv: `1705.06219v2`.

[2] Tarik Aougab, Matthew Gentry Durham, and Samuel J. Taylor. *Pulling back stability with applications to Out(F_n) and relatively hyperbolic groups*. J. Lond. Math. Soc. **96**.3 (2017), pp. 565–583.

[3] Arthur Bartels and Mladen Bestvina. *The Farrell-Jones conjecture for mapping class groups*. Invent. Math. **215**.2 (2019), pp. 651–712.

[4] Jonas Beyrer and Elia Fioravanti. *Cross ratios and cubulations of hyperbolic groups*. preprint. arXiv: `1810.08087`.

[5] Christopher H. Cashen. *Morse subsets of CAT(0) spaces are strongly contracting*. Geom. Dedicata (in press).

[6] Matthew Cordes. *A survey on Morse boundaries and stability*. *Beyond Hyperbolicity*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2019, pp. 83–116.

[7] Matthew Cordes and David Hume. *Stability and the Morse boundary*. J. Lond. Math. Soc. **95**.3 (2017), pp. 963–988.
[8] Cornelia Druţu, Shahar Mozes, and Mark Sapir. *Corrigendum to "Divergence in lattices in semisimple Lie groups and graphs of groups"*. Trans. Amer. Math. Soc. **370**.1 (2018), pp. 749–754.
[9] Elisabeth Fink. *Morse geodesics in torsion groups*. preprint. 2017. arXiv: 1710.11191.
[10] Merlin Incerti-Medici. *Comparing topologies on the Morse boundary and quasi-isometry invariance*. preprint. 2019. arXiv: 1903.07048.
[11] Heejoung Kim. *Stable subgroups and Morse subgroups in mapping class groups*. preprint. 2017. arXiv: 1710.11617v2.
[12] Abhijit Pal and Rahul Pandey. *Acylindrical Hyperbolicity of Subgroups*. preprint. 2019. arXiv: 1903.00628v2.
[13] Abhijit Pal and Suman Paul. *A Note on Strongly Contracting Geodesics in tree of spaces*. preprint. 2019. arXiv: 1904.09906.
[14] Yulan Qing and Kasra Rafi. *Sub-linearly contracting boundary I: CAT(0) spaces*. preprint. 2019.
[15] Jacob Russell, Davide Spriano, and Hung Cong Tran. *Convexity in Hierarchically Hyperbolic Spaces*. preprint. 2018. arXiv: 1809.09303.
[16] Hung Tran. *On strongly quasiconvex subgroups*. Geom. Topol. **23**.3 (2019), pp. 1173–1235.

B. *Morse subsets of CAT(0) spaces are strongly contracting*,
Christopher H. Cashen,
Geometriae Dedicata (in press, accepted May 7, 2019).

We prove that Morse subsets of CAT(0) spaces are strongly contracting. This generalizes and simplifies a result of Sultan, who proved it for Morse quasi-geodesics. Our proof goes through the recurrence characterization of Morse subsets.

The applicant is sole author of this paper and is responsible for the conception, execution, and writing of it in its entirety.

Works citing Paper B.

[1] Jacob Russell, Davide Spriano, and Hung Cong Tran. *Convexity in Hierarchically Hyperbolic Spaces*. preprint. 2018. arXiv: 1809.09303.

C. *Negative curvature in graphical small cancellation groups*,
Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume,
Groups, Geometry and Dynamics **13** (2019), no. 2, 579–632.
DOI: 10.4171/GGD/498

We use the interplay between combinatorial and coarse geometric versions of negative curvature to investigate the geometry of infinitely presented graphical $Gr'(1/6)$ small cancellation groups. In particular, we characterize their 'contracting geodesics', which should be thought of as the geodesics that behave hyperbolically.
We show that every degree of contraction can be achieved by a geodesic in a finitely generated group. We construct the first example of a finitely generated group G containing an element g that is strongly contracting with respect to one finite generating set of G and not strongly contracting with respect to another. In the case of classical $C'(1/6)$ small cancellation groups we give complete characterizations of geodesics that are Morse and that are strongly contracting.
We show that many graphical $Gr'(1/6)$ small cancellation groups contain strongly contracting elements and, in particular, are growth tight. We construct uncountably many quasi-isometry classes of finitely generated, torsion-free groups in which every maximal cyclic subgroup is hyperbolically embedded. These are the first examples of this kind that are not subgroups of hyperbolic groups.
In the course of our analysis we show that if the defining graph of a graphical $Gr'(1/6)$ small cancellation group has finite components, then the elements of the group have translation lengths that are rational and bounded away from zero.

The authors were equal partners in the conception, execution, and writing of this paper

Works citing Paper C.

[1] Carolyn R Abbott and David Hume. *The geometry of generalized loxodromic elements*. Ann. Inst. Fourier (Grenoble) (in press). preprint. arXiv: 1802.03089v2.
[2] Carolyn Abbott and David Hume. *Actions of small cancellation groups on hyperbolic spaces*. preprint. 2018. arXiv: 1807.10524.
[3] Tarik Aougab, Matthew Gentry Durham, and Samuel J. Taylor. *Pulling back stability with applications to Out*(F_n) *and relatively hyperbolic groups*. J. Lond. Math. Soc. **96**.3 (2017), pp. 565–583.
[4] Goulnara N. Arzhantseva and Christopher H. Cashen. *Cogrowth for group actions with strongly contracting elements*. Ergodic Theory Dynam. Systems (in press).
[5] Christopher H. Cashen. *Morse subsets of CAT(0) spaces are strongly contracting*. Geom. Dedicata (in press).
[6] Christopher H. Cashen and John M. Mackay. *A Metrizable Topology on the Contracting Boundary of a Group*. Trans. Amer. Math. Soc. **372**.3 (2019), pp. 1555–1600.
[7] Matthew Cordes and David Hume. *Stability and the Morse boundary*. J. Lond. Math. Soc. **95**.3 (2017), pp. 963–988.
[8] Ilya Gekhtman and Wen-yuan Yang. *Counting conjugacy classes in groups with contracting elements*. preprint. 2018. arXiv: 1810.02969.
[9] Dominik Gruber and Alessandro Sisto. *Infinitely presented graphical small cancellation groups are acylindrically hyperbolic*. Ann. Inst. Fourier (Grenoble) **68**.6 (2018), pp. 2501–2552. arXiv: 1408.4488.
[10] Suzhen Han and Wen-yuan Yang. *Generic free subgroups and statistical hyperbolicity*. preprint. 2018. arXiv: 1812.06265.
[11] David Hume and Alessandro Sisto. *Groups with no coarse embeddings into hyperbolic groups*. New York J. Math **23** (2017), pp. 1657–1670.
[12] Merlin Incerti-Medici. *Comparing topologies on the Morse boundary and quasi-isometry invariance*. preprint. 2019. arXiv: 1903.07048.
[13] Wenyuan Yang. *Genericity of contracting elements in groups*. Math. Ann. (in press).
[14] Wen-yuan Yang. *Statistically convex-cocompact actions of groups with contracting elements*. Int. Math. Res. Not. (2018).

D. *Growth Tight Actions*,

Goulnara N. Arzhantseva, Christopher H. Cashen, and Jing Tao,

Pacific Journal of Mathematics **278** (2015), no. 1, 1–49.

DOI:10.2140/pjm.2015.278.1

We introduce and systematically study the concept of a growth tight action. This generalizes growth tightness for word metrics as initiated by Grigorchuk and de la Harpe. Given a finitely generated, non-elementary group G acting on a G–space $\mathcal{X}$, we prove that if G contains a strongly contracting element and if G is not too badly distorted in $\mathcal{X}$, then the action of G on $\mathcal{X}$ is a growth tight action. It follows that if $\mathcal{X}$ is a cocompact, relatively hyperbolic G–space, then the action of G on $\mathcal{X}$ is a growth tight action. This generalizes all previously known results for growth tightness of cocompact actions: every already known example of a group that admits a growth tight action and has some infinite, infinite index normal subgroups is relatively hyperbolic, and, conversely, relatively hyperbolic groups admit growth tight actions. This also allows us to prove that many CAT(0) groups, including flip-graph-manifold groups and many Right Angled Artin Groups, and snowflake groups admit cocompact, growth tight actions. These provide first examples of non relatively hyperbolic groups admitting interesting growth tight actions. Our main result applies as well to cusp uniform actions on hyperbolic spaces and to the action of the mapping class group on Teichmüller space with the Teichmüller metric. Towards the proof of our main result, we give equivalent characterizations of strongly contracting elements and produce new examples of group actions with strongly contracting elements.

The authors were equal partners in the conception, execution, and writing of this paper.

Works citing Paper D.

[1] Goulnara N. Arzhantseva and Christopher H. Cashen. *Cogrowth for group actions with strongly contracting elements*. Ergodic Theory Dynam. Systems (in press).
[2] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Negative Curvature in Graphical Small Cancellation Groups*. Groups Geom. Dyn. **13**.2 (2019), pp. 579–632.
[3] Anne Broise-Alamichel, Jouni Parkkonen, and Frédéric Paulin. *Equidistribution and counting under equilibrium states in negatively curved spaces and graphs of groups. Applications to non-Archimedean Diophantine approximation*. Vol. 329. Progress in Mathematics. preprint. Birkhäuser Basel, 2019. arXiv: 1612.06717.
[4] Christopher H. Cashen and Jing Tao. *Growth Tight Actions of Product Groups*. Groups Geom. Dyn. **10**.2 (2016), pp. 753–770.
[5] Rémi Coulon, Rhiannon Dougall, Barbara Schapira, and Samuel Tapie. *Twisted Patterson-Sullivan measures and applications to amenability and coverings*. preprint. 2018. arXiv: 1809.10881v2.
[6] François Dahmani, David Futer, and Daniel T Wise. *Growth of quasiconvex subgroups*. Math. Proc. Cambridge Philos. Soc. (2018), pp. 1–26.
[7] Shubhabrata Das and Mahan Mj. *Controlled Floyd separation and non relatively hyperbolic groups*. J. Ramanujan Math. Soc. **30**.3 (2015), pp. 267–294.
[8] Ilya Gekhtman, Samuel J Taylor, and Giulio Tiozzo. *Counting problems in graph products and relatively hyperbolic groups*. preprint. 2017. arXiv: 1711.04177.
[9] Ilya Gekhtman and Wen-yuan Yang. *Counting conjugacy classes in groups with contracting elements*. preprint. 2018. arXiv: 1810.02969.
[10] Suzhen Han and Wen-yuan Yang. *Generic free subgroups and statistical hyperbolicity*. preprint. 2018. arXiv: 1812.06265.
[11] Ilya Kapovich, Joseph Maher, Catherine Pfaff, and Samuel J Taylor. *Random outer automorphisms of free groups: Attracting trees and their singularity structures*. preprint. 2018. arXiv: 1805.12382.
[12] Katsuhiko Matsuzaki. *Growth and cogrowth tightness of Kleinian and hyperbolic groups*. *Geometry and Analysis of Discrete Groups and Hyperbolic Spaces*. Ed. by M. Fujii, N. Kawazumi, and K. Ohshika. Vol. B66. RIMS Kôkyûroku Bessatsu. RIMS, 2017, pp. 21–36.
[13] Mahan Mj and Parthanil Roy. *Stable Random Fields, Bowen-Margulis measures and Extremal Cocycle Growth*. preprint. 2018. arXiv: 1809.08295.
[14] Kasra Rafi and Yvon Verberne. *Geodesics in the mapping class group*. preprint. 2018. arXiv: 1810.12489.
[15] Wenyuan Yang. *Genericity of contracting elements in groups*. Math. Ann. (in press).
[16] Wen-yuan Yang. *Statistically convex-cocompact actions of groups with contracting elements*. Int. Math. Res. Not. (2018).

E. *Growth Tight Actions of Product Groups*,
Christopher H. Cashen and Jing Tao,
Groups, Geometry and Dynamics **10** (2016), no. 2, 753–770.
DOI:10.4171/GGD/364

A group action on a metric space is called growth tight if the exponential growth rate of the group with respect to the induced pseudo-metric is strictly greater than that of its quotients. A prototypical example is the action of a free group on its Cayley graph with respect to a free generating set. More generally, with Arzhantseva we have shown that group actions with strongly contracting elements are growth tight. Examples of non-growth tight actions are product groups acting on the L^1 products of Cayley graphs of the factors.

In this paper we consider actions of product groups on product spaces, where each factor group acts with a strongly contracting element on its respective factor space. We show that this action is growth tight with respect to the L^p metric on the product space, for all $1 < p \leq \infty$. In particular, the L^∞ metric on a product of Cayley graphs corresponds to a word metric on the product group. This gives the first examples of groups that are growth tight with respect to an action on one of their Cayley graphs and non-growth tight with respect to an action on another, answering a question of Grigorchuk and de la Harpe.

The authors were equal partners in the conception, execution, and writing of this paper.

Works citing Paper E.

[1] Goulnara N. Arzhantseva, Christopher H. Cashen, and Jing Tao. *Growth Tight Actions*. Pacific J. Math. **278**.1 (2015), pp. 1–49.

F. *Cogrowth for group actions with strongly contracting elements*,
Goulnara N. Arzhantseva and Christopher H. Cashen,
Ergodic Theory & Dynamical Systems (in press, accepted Oct 15, 2018).
DOI:10.1017/etds.2018.123

Let G be a group acting properly by isometries and with a strongly contracting element on a geodesic metric space. Let N be an infinite normal subgroup of G, and let δ_N and δ_G be the growth rates of N and G with respect to the pseudo-metric induced by the action. We prove that if G has purely exponential growth with respect to the pseudo-metric then $\delta_N/\delta_G > 1/2$. Our result applies to suitable actions of hyperbolic groups, right-angled Artin groups and other CAT(0) groups, mapping class groups, snowflake groups, small cancellation groups, etc. This extends Grigorchuk's original result on free groups with respect to a word metrics and a recent result of Jaerisch, Matsuzaki, and Yabuki on groups acting on hyperbolic spaces to a much wider class of groups acting on spaces that are not necessarily hyperbolic.

The authors were equal partners in the conception, execution, and writing of this paper.

Works citing Paper F.

[1] Ilya Gekhtman and Arie Levit. *Critical exponents of invariant random subgroups in negative curvature*. Geom. Funct. Anal. (in press). arXiv: 1804.02995.

G. *Quasi-isometries need not induce homeomorphisms of contracting boundaries with the Gromov product topology*,
Christopher H. Cashen,
Analysis and Geometry in Metric Spaces **4** (2016), no. 1, 278–281.
DOI:10.1515/agms-2016-0011

We consider a 'contracting boundary' of a proper geodesic metric space consisting of equivalence classes of geodesic rays that behave like geodesics in a hyperbolic space. We topologize this set via the Gromov product, in analogy to the topology of the boundary of a hyperbolic space. We show that when the space is not hyperbolic, quasi-isometries do not necessarily give homeomorphisms of this boundary. Continuity can fail even when the spaces are required to be CAT(0). We show this by constructing an explicit example.

The applicant is sole author of this paper and is responsible for the conception, execution, and writing of it in its entirety.

Works citing Paper G.

[1] Christopher H. Cashen and John M. Mackay. *A Metrizable Topology on the Contracting Boundary of a Group*. Trans. Amer. Math. Soc. **372**.3 (2019), pp. 1555–1600.
[2] Matthew Cordes. *A survey on Morse boundaries and stability*. *Beyond Hyperbolicity*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2019, pp. 83–116.
[3] Matthew Cordes and David Hume. *Stability and the Morse boundary*. J. Lond. Math. Soc. **95**.3 (2017), pp. 963–988.
[4] Merlin Incerti-Medici. *Comparing topologies on the Morse boundary and quasi-isometry invariance*. preprint. 2019. arXiv: 1903.07048.
[5] Yulan Qing and Kasra Rafi. *Sub-linearly contracting boundary I: CAT(0) spaces*. preprint. 2019.

H. *A metrizable topology on the contracting boundary of a group*,
Christopher H. Cashen and John M. Mackay,
Transactions of the American Mathematical Society, 372 (2019), no. 3, 1555–1600.
DOI:10.1090/tran/7544
The ‘contracting boundary’ of a proper geodesic metric space consists of equivalence classes of geodesic rays that behave like rays in a hyperbolic space. We introduce a geometrically relevant, quasi-isometry invariant topology on the contracting boundary. When the space is the Cayley graph of a finitely generated group we show that our new topology is metrizable.

The authors were equal partners in the conception, execution, and writing of this paper.

Works citing Paper H.

[1] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham. *Largest acylindrical actions and stability in hierarchically hyperbolic groups*. preprint. 2017. arXiv: 1705.06219v2.
[2] Jonas Beyrer and Elia Fioravanti. *Cross ratios and cubulations of hyperbolic groups*. preprint. arXiv: 1810.08087.
[3] Jonas Beyrer and Elia Fioravanti. *Cubulations and cross ratios on contracting boundaries*. preprint. 2018. arXiv: 1810.08087.
[4] Christopher H. Cashen. *Morse subsets of CAT(0) spaces are strongly contracting*. Geom. Dedicata (in press).
[5] Ruth Charney, Matthew Cordes, and Devin Murray. *Quasi-Mobius Homeomorphisms of Morse boundaries*. Bull. London Math. Soc. **51**.3 (2019), pp. 501–515.
[6] Ruth Charney and Devin Murray. *A rank-one CAT(0) group is determined by its Morse boundary*. preprint. 2017. arXiv: 1707.07028.
[7] Matthew Cordes. *A survey on Morse boundaries and stability*. *Beyond Hyperbolicity*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2019, pp. 83–116.
[8] Matthew Cordes and Matthew Gentry Durham. *Boundary convex cocompactness and stability of subgroups of finitely generated groups*. Int. Math. Res. Not. (2016).
[9] Merlin Incerti-Medici. *Comparing topologies on the Morse boundary and quasi-isometry invariance*. preprint. 2019. arXiv: 1903.07048.
[10] Abhijit Pal and Rahul Pandey. *Acylindrical Hyperbolicity of Subgroups*. preprint. 2019. arXiv: 1903.00628v2.
[11] Yulan Qing and Kasra Rafi. *Sub-linearly contracting boundary I: CAT(0) spaces*. preprint. 2019.

III. Curriculum Vitae

Christopher H. Cashen

Faculty of Mathematics – University of Vienna – Oskar-Morgenstern-Platz 1
1090 Vienna, Austria
☎ +43 1 4277 50656 • ✉ christopher.cashen@univie.ac.at
http://www.mat.univie.ac.at/~cashen

Academic Positions

University of Vienna — **Vienna, Austria**
Project Leader/Senior Postdoc — *2017–present*
Postdoctoral Fellow — *2016–2017*
Lise Meitner Fellow — *2014–2016*
Postdoctoral Fellow — *2012–2014*

Université de Caen Normandie — **Caen, France**
Postdoctoral Fellow — *2011–2012*

University of Utah — **Salt Lake City, USA**
Postdoctoral Fellow — *2008–2011*

Mathematical Sciences Research Institute — **Berkeley, USA**
Postdoctoral Fellow — *Fall 2007*

University of Illinois at Chicago — **Chicago, USA**
Coordinator, Mathematical Sciences Learning Center — *Spring and Fall 2006*

University of Illinois at Chicago — **Chicago, USA**
Graduate Student — *2000-2007*

Education

University of Illinois at Chicago — **Chicago, USA**
PhD in Mathematics — *2007*

University of Illinois at Chicago — **Chicago, USA**
MS in Mathematics — *2001*

Loyola University — **Chicago, USA**
BS in Mathematics — *2000*
Magna Cum Laude, Alpha Sigma Nu honor society

Grant Awards

Austrian Science Fund (FWF): P 30487-N35
Project Leader *Generalizations of Hyperbolic Boundaries* — *2017–2020*

Austrian Science Fund (FWF): M 1717-N25
Project Leader *Geometric and Analytic Aspects of Free Group Automorphisms* — *2014–2016*
Lise Meitner Fellowship

Publications

Peer-reviewed

1. *Quasi-isometries between tubular groups*, Groups, Geometry and Dynamics **4** (2010), no. 3, 473–516.
2. *Line patterns in free groups*, w/Macura, Geometry & Topology **15** (2011), no. 3, 1419–1475.
3. *Growth tight actions*, w/Arzhantseva and Tao, Pacific Journal of Mathematics **278** (2015), no. 1, 1–49.
4. *Virtual geometricity is rare*, w/Manning, LMS Journal of Computation and Mathematics **18** (2015), no. 1, 444–455.
5. *Quasi-isometries need not induce homeomorphisms of contracting boundaries with the Gromov product topology*, Analysis and Geometry in Metric Spaces **4** (2016), no. 1, 278–281.
6. *Splitting line patterns in free groups*, Algebraic & Geometric Topology **16** (2016), no. 2, 621–673.
7. *Mapping tori of free group automorphisms, and the Bieri-Neumann-Strebel invariant of graphs of groups*, w/Levitt, Journal of Group Theory **19** (2016), no. 2, 191–216.
8. *Growth tight actions of product groups*, w/Tao, Groups, Geometry and Dynamics **10** (2016), no. 2, 753–770.
9. *A geometric proof of the structure theorem for cyclic splittings of free groups*, Topology Proceedings **50** (2017), 335–349.
10. *Characterizations of Morse geodesics via superlinear divergence and sublinear contraction*, w/Arzhantseva, Gruber, and Hume, Documenta Mathematica **22** (2017), 1193–1224.
11. *Quasi-isometry classification for [right-angled Coxeter groups defined by suitable subdivisions of] complete graphs*, w/Dani and Thomas, Journal of Topology **10** (2017), no. 4, 1066–1106, appendix to *Bowditch's JSJ tree and the quasi-isometry classification of certain Coxeter groups* by Dani and Thomas.
12. *Quasi-isometries between groups with two-ended splittings*, w/Martin, Mathematical Proceedings of the Cambridge Philosopical Society **162** (2017), no. 2, 249–291.
13. *Negative curvature in graphical small cancellation groups*, w/Arzhantseva, Gruber, and Hume, Groups, Geometry and Dynamics **13** (2019), no. 2, 579–632.
14. *A metrizable topology on the contracting boundary of a group*, w/Mackay, Transactions of the American Mathematical Society **372** (2019), no. 3, 1555–1600.
15. *Morse subsets of CAT(0) spaces are strongly contracting*, Geometriae Dedicata (in press).
16. *Cogrowth for group actions with strongly contracting elements*, w/Arzhantseva, Ergodic Theory and Dynamical Systems (in press).

Software Development

17. *virtuallygeometric*, w/Manning, (2014), computer program, `https://bitbucket.org/christopher_cashen/virtuallygeometric`.

Invited Research Visits

Non-Positive Curvature, Group Actions, and Cohomology	Newton Institute, UK, 03.2017
Measured Group Theory	ESI, Austria, 02.2016
Low-dimensional Topology, Geometry, and Dynamics	ICERM, USA, 10.2013
The Geometry of Outer Space	Aix-Marseille U., France, 07.2013
Automorphisms of Free Groups	CRM, Spain, 11.2012
Geometric Group Theory	MSRI, USA, Fall 2007

Invited Conference Talks

Of coarse! Quasi-isometries and groups: rigidity and classification — Ventotene, Italy, 09.2019

Groups with hyperbolic features — ETH Zürich, Switzerland, 08.2019

Dubrovnik IX, Topology & Dynamical Systems — IUC Dubrovnik, Croatia, 06.2019
Group actions with a strongly contracting element

Non-positive Curvature in Action — Newton Institute, UK, 01.2017
The topology of the contracting boundary of a group

Dubrovnik VIII, Geometric Topology, Geometric Group Theory & Dynamical Systems — IUC Dubrovnik, Croatia, 06.2015
Contracting elements in infinitely presented small cancellation groups

Geometry of Computation in Groups — ESI, Austria, 03.2014
Growth tight actions

Groups and Geometry in the South East — U. Southampton, UK, 12.2012
Quasi-isometries of groups admitting certain cyclic JSJ decompositions

Automorphisms of Free Groups — CRM, Spain, 11.2012
Mapping Tori of Polynomially Growing Free Group Automorphisms

Journées sur $\mathrm{Out}(F_n)$ — U. Paris-Sud, France, 08.2011
Quasi-Isometries of Mapping Tori of Free Group Automorphisms

AMS Sectional Meeting — UNLV, USA, 04.2011
Virtually Geometric Multiwords

Spring Topology and Dynamics Conference — U. Texas at Tyler, USA, 03.2011
Virtually Geometric Multiwords

Wasatch Topology Conference — Park City, USA, 12.2010
Virtually Geometric Multiwords

Quasi-isometric Rigidity in Low Dimensional Topology — BIRS, Canada, 03.2010
Line Patterns in Free Groups

AMS Sectional Meeting — FAU, USA, 11.2009
Line Patterns in Free Groups and Quasi-isometries of Mapping Tori of Linearly Growing Free Group Automorphisms

AMS Sectional Meeting — LSU, USA, 03.2008
Quasi-isometries Between Tubular Groups

Student Supervision

Zachary Wilcox (Bachelor) "Introduction to topology via matrix groups" University of Utah, 2011.

Alexandra Edletzberger (Master) "Residual finiteness in hyperbolic groups" University of Vienna, in progress.
Charlotte Hoffmann (Master) University of Vienna, in progress.

Teaching Experience

University of Vienna..

5 years as co-organizer Algebra Project Seminar: "Geometry and Analysis on Groups".

1 semester external lecturer for Exercises in Probability Theory and Statistics.

University of Utah..

7 semesters as sole instructor for the following courses:

Semester	Course	Enrollment	ICS
2008 Spring	Calculus II	46	4.77
2008 Fall	Trigonometry	95	4.32
2009 Spring	Real Analysis II (undergrad)	26	5.03
2009 Fall	PDE's for Engineers	26	5.05
2010 Spring	Riemannian Geometry (grad)	8	n/a
2010 Fall	Linear Algebra	42	5.17
2011 Spring	Calculus I	120	5.15

ICS = Instructor Composite Score from student evaluations, 1-6 with 6 best.

University of Illinois at Chicago...

1 year as coordinator of Mathematical Sciences Learning Center.

1 semester as lecturer for Multivariable Calculus.

2 years as a teaching assistant for Calculus sequence and Ordinary Differential Equations.

Professional Activities

Referee for various journals

Reviewer for Mathematical Reviews

Member: American Math. Society, Austrian Math. Society, Association for Women in Math.

Organizer: 'Max Dehn Seminar', U. Utah, 2008-2011

Organizer: 'Geometry and Analysis on Groups Seminar', U. Vienna, 2014-present

Survey of the topic and the impacts of this thesis

1. Introduction

A subset $\mathcal{Z}$ of a geodesic metric space X is called *quasiconvex* if there exists $Q \geq 0$ such that every geodesic segment connecting two points in $\mathcal{Z}$ is contained in the Q–neighborhood of $\mathcal{Z}$. Quasiisometries between hyperbolic spaces preserve quasiconvex subsets, so it makes sense to say that a subgroup of a hyperbolic group G is, or is not, quasiconvex, according to whether it corresponds to a quasiconvex subset in some/every Cayley graph of G. Quasiconvex subgroups of hyperbolic groups have many nice properties. For example, let H and H' be infinite quasiconvex subgroups of a hyperbolic group G. The following are all true:

(1) H is finitely generated and undistorted.
(2) $H \cap H'$ is quasiconvex.
(3) H has finite height, finite width, and bounded packing.
(4) H has finite index in its commensurator.
(5) (H is hyperbolic and) The inclusion of H into G extends to an embedding $\partial H \hookrightarrow \partial G$ whose image coincides with the limit set $\Lambda(H)$.
(6) If H has infinite index in G then there exists $g \in G$ such that $\langle H, g\rangle \cong H * \langle g\rangle$ is quasiconvex.

When the ambient space is not hyperbolic, the quasiconvexity property loses much of its strength. In Section 3 we describe another property, the *Morse property*, that is equivalent to quasiconvexity in hyperbolic spaces, but remains interesting in nonhyperbolic spaces as well. It is also invariant under quasiisometries of the ambient space, so again we can apply it to subgroups, and many of the properties of quasiconvex subgroups of hyperbolic groups can be generalized to Morse subgroups of arbitrary finitely generated groups.

Theorem 1.1. *Suppose that H and H' are infinite Morse subgroups of a finitely generated group G.*

(1) H is finitely generated and undistorted.
(2) $H \cap H'$ is quasiconvex.
(3) H has finite height, finite width, and bounded packing.
(4) H has finite index in its commensurator.
(5) The inclusion of H into G extends to an embedding $\partial H \hookrightarrow \partial G$ whose image coincides with the limit set $\Lambda(H)$. Here ∂G and ∂H can refer to either the Morse or contracting boundaries of G and H.

(1), (2), and (4) are true by arguments similar to the quasiconvex case in hyperbolic groups. (3) is a result of Tran [77], and the interested reader is referred there for definitions.

For (5) to make sense we need a notion of a boundary of a finitely generated group that plays the role of the hyperbolic boundary. This is the topic of Section 5.

Item (6) from the hyperbolic case is not true in general for Morse subgroups. In the special case that H is cyclic, Item (6) allows us to build free subgroups of G. However, Osin, Ol′shanskii, and Sapir [59] constructed a so-called "Tarski monster" group, which is an infinite, finitely generated, non-virtually cyclic group with the properties that every cyclic subgroup is Morse and every proper subgroup is cyclic. The failure of Item (6) motivates a search for even stronger properties, some of which are described in Section 6.

In Section 7 we focus on one of these stronger properties. We study actions of groups on metric spaces where at least one element of the group acts as a *strongly contracting element* and prove some results about growth and cogrowth of group actions.

In Section 4 we apply the innovations of Section 3 to study the geometry of infinitely presented graphical small cancellation groups. These provide a useful source of examples for various metric properties in groups. In particular, we show that the property of being a strongly contracting element for an action on a Cayley graph can depend on the choice of generating set.

1.1. On the impacts of this thesis. In the coming pages we survey the constituent papers of this thesis and related literature, which are all tied together by the broad goal, exemplified by Theorem 1.1, of generalizing results from hyperbolic groups to arbitrary finitely generated groups that feature limited forms of hyperbolic behavior. Here we highlight our most important contributions.

In [Paper A]/Section 3 we introduce the contraction property and show it is equivalent to several previously known properties, including the Morse property. We initially, in [Paper A], proved the equivalence of contraction and the Morse property, while others extended our work to include recurrence [5] and divergence [77]. The equivalence of all of these properties was only known before in the case of quasigeodesics in CAT(0) spaces and quasiconvex sets in hyperbolic spaces, but using our characterization of contraction we can now prove them for arbitrary subsets of arbitrary geodesic metric spaces. These equivalences have been used to give characterizations of stable subgroups of mapping class groups, outer automorphism groups of free groups, some relatively hyperbolic groups, and right-angled Artin and Coxeter groups [5, 77], and to relate divergence to the topology of asymptotic cones [30].

The contraction property turned out to be versatile, but we original designed it specifically for the study of infinitely presented graphical small cancellation groups. In [Paper C]/Section 4 we prove foundational results about this very interesting class of groups. Some of these results are new even when restricted to the well-studied special case of classical small cancellation groups. Conversely, this class of groups is a well suited to constructing exotic examples to tease apart curvature properties in finitely generated groups.

The contraction property was also instrumental in the construction of a quasiisometry invariant metrizable topology on the contracting boundary of a group, described in [Paper H]/Section 5. Boundaries of hyperbolic groups have proven to be very useful tools, and we anticipate that our boundary will lead to many applications for nonhyperbolic groups. Qing and Rafi [64] have extended our construction to give a quasiisometry invariant topological model for the Poisson boundary of CAT(0) groups.

In [Paper D]/Section 7 we initiate the study of groups acting on metric spaces with a strongly contracting element. Such actions encompass geometric actions on hyperbolic and CAT(0) spaces, actions of relatively hyperbolic groups on hyperbolic spaces, and actions of mapping class groups on Teichmüller spaces. The latter two types of actions are not cocompact, and it turns out that even in the relatively hyperbolic case the kind of growth results we would like to prove break down if we allow arbitrary noncocompact proper actions. We introduce a property called *complementary growth gap* to allow noncocompact actions that still have some control on how badly the orbits of the group are distorted. This implies the ‘parabolic gap’ condition that had previously been utilized in the study of Kleinian groups, but can be defined for arbitrary proper actions, and, in particular, is satisfied by the action of a mapping class group on its Teichmüller space with the Teichmüller metric. With these two conditions, strongly contracting element + complementary growth gap, we are able to reprove and vastly generalize growth [Paper D] and [Paper E] and cogrowth [Paper F] theorems that were previously known only in the hyperbolic case, or sometimes only in the case of free groups! Yang [82] and Coulon et al. [25] have subsequently adopted our framework to generalize famous results such as pure exponential growth, genericity of hyperbolic elements, and the Grigorchuk amenability criterion.

2. Preliminaries

Throughout, $\mathcal{X}$ is a geodesic metric space, groups are finitely generated, and actions are isometric actions. Let $\mathbb{E}^n$ denote n–dimensional Euclidean space. We refer the reader to [16] for standard definitions such as hyperbolic and CAT(0) spaces, quasiisometries, and quasigeodesics.

An isometric action of a finitely generated group on a geodesic metric space is *geometric* if it is properly discontinuous and cocompact.

Let $I \subset \mathbb{R}$ denote some interval. A rectifiable path $\gamma\colon I \to \mathcal{X}$ with endpoints x and y is *E–efficient* if $\mathrm{len}(\gamma) \leq Ed(x, y)$. It is (E_1, E_2)–efficient if $\mathrm{len}(\gamma) \leq E_1 d(x, y) + E_2$. A path is *locally* (E_1, E_2)*–efficient* if every bounded subpath is (E_1, E_2)–efficient.

- $\mathcal{S}_r(\mathcal{Z}) := \{x \in \mathcal{X} \mid d(x, \mathcal{Z}) = r\}$
- $\mathcal{N}_r(\mathcal{Z}) := \{x \in \mathcal{X} \mid d(x, Z) < r\}$.
- $\mathcal{N}^c_r(\mathcal{Z}) := \mathcal{X} - \mathcal{N}_r(\mathcal{Z})$.
- $\bar{\mathcal{N}}_r(\mathcal{Z}) := \{x \in \mathcal{X} \mid d(x, Z) \leq r\}$.

If $\phi, \psi\colon \mathbb{R} \to \mathbb{R} \cup \{\pm\infty\}$ then we write $\phi \prec \psi$ if there exists constants $A, B, C, D > 0$ such that for all r sufficiently large we have $\phi(r) \leq A\psi(Br + C) + D$, and we write $\phi \asymp \psi$ if $\phi \prec \psi$ and $\psi \prec \phi$. When $\phi \asymp \psi$ we say the functions ϕ and ψ are *equivalent*.

We denote the identity function $\mathbb{R} \to \mathbb{R}$ by Id, so for all $r \in \mathbb{R}$ we have $\mathrm{Id}(r) = r$.

Lemma 2.1. *Let $\mathcal{X}$ be a geodesic metric space. Given $L \geq 1$ and $A \geq 0$ there exist $H \geq 0$, $E_1 \geq 1$, $E_2 \geq 0$, and $A' \geq 0$ such that every (L, A)–quasigeodesic is Hausdorff distance at most H from an (L, A')–quasigeodesic, locally (E_1, E_2)–efficient path with the same endpoints.*

Conversely, every locally (E_1, E_2)–efficient path, upon reparameterizing by arc length, is an (E_1, E_2)–quasigeodesic.

Proof. The first claim is [16, Lemma III.H.1.11]. The second is easy. □

3. The Morse property and related properties

In this section we fix a subset $\mathcal{Z}$ of a geodesic metric space $\mathcal{X}$ and define several analogues of quasiconvexity of $\mathcal{Z}$. Intuitively, these properties all say that the most efficient way to travel between points in $\mathcal{Z}$ is to travel in, or close to, $\mathcal{Z}$, and that it is increasingly inefficient to make progress in $\mathcal{Z}$ using paths that stray far from it. To each such property we associate a *gauge* function quantifying it. We are interested not only in deciding if two properties are equivalent, but also in establishing *effective* relations, in the sense that the gauge for one property can be bounded in terms of the gauge for another.

Some of the properties have a long history. Indeed, Morse [56] introduced the first version of such a property some 95 years ago. In this thesis, specifically, in [Paper A], we introduce a new property, *contraction*, that ties the others together and lets us prove that many of them are actually equivalent. A downside to this long history is that some of the properties have become more nuanced as their scope has expanded. Some results that were first proved in hyperbolic or CAT(0) settings, or were proved with the assumption that $\mathcal{Z}$ is a geodesic or a quasigeodesic, can be extended to more general situations. Sometimes this generalization is direct, but sometimes it comes at the cost of making the geometric hypothesis on $\mathcal{Z}$ stronger. However, the terminology used in the literature can blur the distinction between such enhanced properties. For instance, at least four of the properties we list can be found in the literature under the name ‘contraction’, yet are distinct in general. Conversely, sometimes different authors use different names for identical definitions. In this section we review the various properties from the literature, including from [Paper D] and [Paper A], and impose a terminology to distinguish them. We give a comprehensive statement of the relationships between these properties, see Theorem 3.2 and Figure 4. Following the theorem, in Definition 3.3 we define a set to be *Morse* if it satisfies a set of six equivalent properties.

We are not aware of any single source that has tackled all of these properties in this generality, but the theorem is not a new result. Rather, Theorem 3.2 is an attempt to ‘tidy-up’ the literature by casting all the necessary pieces in the same general setting. For completeness we have included a proof, which is broken into separate statements for each of the implications. These are either easy or are adaptations of results appearing in this thesis or elsewhere in the literature, and can be found in the appendix (Section 8). The appendix also contains computations of two families of examples that demonstrate nonimplications between some of the properties and show that some of the effective bounds we give are sharp.

The definitions and their history are given in Section 3.1. In Section 3.2 we comment on quasiisometry invariance and the natural extension of these metric properties to subgroups.

Throughout this section we make the assumption that the set-valued closest point projection map $\pi_{\mathcal{Z}}(x) := \{z \in \mathcal{Z} \mid d(x,z) \leq d(x,\mathcal{Z})\}$ does not have the empty set in its image. This is for convenience. If it is not the case, it can be arranged by choosing any $\epsilon > 0$ and using the map $\pi^{\epsilon}_{\mathcal{Z}}(x) := \{z \in \mathcal{Z} \mid d(x,z) \leq d(x,\mathcal{Z}) + \epsilon\}$ instead. Of course this will mean that the *quantitative* equivalences between our various properties will pick up some factors of ϵ in the computations. However, Lemma 8.5 shows that *qualitatively* the choice of ϵ does not matter.

3.1. Definitions. In the following definitions, the phrase "there is a function defined by…" is an assertion that the succeeding expression containing a supremum or infinum defines a (real-valued, unless specified) function, ie. that the extremum exists.

3.1.1. *Path-based definitions.* We first consider the behavior of various path systems with respect to $\mathcal{Z}$:

- $\mathcal{G}$ the set of all geodesic segments.
- $\mathcal{Q} := \cup_{L\geq 1, A\geq 0}\mathcal{Q}(L,A)$ where $\mathcal{Q}(L,A)$ is the set of all (L,A)–quasigeodesic segments.
- $\mathcal{L} := \cup_{E_1\geq 1, E_2\geq 0}\mathcal{L}(E_1,E_2)$ where $\mathcal{L}(E_1,E_2)$ is the set of all locally (E_1,E_2)–efficient segments.

We set up the definitions in a uniform way by defining $\mathcal{PS}(\underline{C})$ to be a collection of maps from intervals of $\mathbb{R}$ into $\mathcal{X}$ satisfying some condition depending on parameters $\underline{C} := (C_1, C_2, \dots)$, and let $\mathcal{PS}$ be the union of the $\mathcal{PS}(\underline{C})$ as $\underline{C}$ varies over some space of parameters, eg.: $\mathcal{PS} = \mathcal{Q}$, $\mathcal{G}$, or $\mathcal{L}$.

$\mathcal{PS}$–quasiconvex: There is a function χ defined by $\chi(\underline{C}) := \sup_\gamma \sup_{w\in\gamma} d(w,\mathcal{Z})$, where the first supremum is taken over those $\gamma \in \mathcal{PS}(\underline{C})$ with both endpoints on $\mathcal{Z}$.

$\mathcal{PS}$–trim: For each fixed $\underline{C}$ there is a sublinear function $r \mapsto \tau(r;\underline{C})$ defined by $\tau(r;\underline{C}) := \sup_\gamma \sup_z d(z,\gamma)$ where the first supremum is taken over $\gamma \in \mathcal{PS}(\underline{C})$ such that, if γ^+ and γ^- are the endpoints of γ, then $\gamma^- \in \mathcal{Z}$ and $d(\gamma^+,\mathcal{Z}) \leq r$. The second supremum is taken over points $z \in \pi_Z(\gamma^+)$.

If, in addition, for each $\underline{C}$ the function $\tau(r;\underline{C})$ is bounded, then we say $\mathcal{Z}$ is *strongly $\mathcal{PS}$*–trim.

The usual notion of 'quasiconvexity' is 'geodesic quasiconvexity' in this scheme. Since there is only one point in the parameter space for $\mathcal{G}$, the quasiconvexity gauge χ has just a single value $\chi(1) = C$. Thus, we say 'C–quasiconvex' to mean geodesically quasiconvex with quasiconvexity gauge $\chi(1) = C$.

Notice that all of these path-based properties are preserved upon passing to sub-path systems $\mathcal{PS}' \subset \mathcal{PS}$, and, more generally, upon passing to $\mathcal{PS}'$ such that for every $\underline{C}'$ there exists $\underline{C}$ and H such that for every $\gamma' \in \mathcal{PS}'(\underline{C}')$ there exists a $\gamma \in \mathcal{PS}(\underline{C})$ at Hausdorff distance at most H from γ'. An example of the former is that geodesics and, by Lemma 2.1, locally efficient paths, are quasigeodesic, so quasigeodesic quasiconvexity implies geodesic quasiconvex and locally efficient quasiconvexity. An example of the latter is that quasigeodesics are bounded Hausdorff distance from locally efficient paths, as in Lemma 2.1, so locally efficient quasiconvexity implies quasigeodesic quasiconvexity. Thus, quasigeodesic quasiconvexity and locally efficient quasiconvexity are equivalent, although the gauge χ of course depends on the specification.

Morse [56, Lemma 8] essentially shows that geodesics in the hyperbolic plane are locally efficiently quasiconvex. Gromov [43, Proposition 7.2.A] observed that this is true for geodesics in any hyperbolic space. He referred to this property as 'stability'. The term 'stable' continued to be used in the early geometric group theory literature to describe quasigeodesically quasiconvex quasigeodesics.

Druţu and Sapir [32] talk about geodesics in relatively hyperbolic groups satisfying the 'Morse property' to describe quasigeodesic quasiconvexity. Druţu, Mozes, and Sapir [31] introduce the terminology 'Morse quasigeodesic' for a quasigeodesically quasiconvex quasigeodesic.

Our paper [Paper A] characterizes the Morse property in the guise of quasigeodesic quasiconvexity for arbitrary subsets. Subsequently, both Genevois [36] and Tran [77] reintroduced this property. Genevois uses the terminology 'Morse subset'. Tran says 'Morse' for subsets, but 'strongly quasiconvex' for subgroups. Note that the terminology 'strongly quasiconvex' has also been used for different notions, for example, in [27].

We use the name 'trim' to be different from, but still call to mind, the terms 'slim' and 'thin' often used to describe hyperbolic triangles. A triangle whose base is contained in a geodesically trim set has a waist that is small relative to its height. 'Strongly geodesically trim', restricted to the case that $\mathcal{Z}$ is a quasigeodesic in a CAT(0) space, matches Sultan's [74] definition of 'slim'. Similarly, Bestvina and Fujiwara [13] prove a lemma that says a triangle in a CAT(0) space with one contracting side is 'thin'. In a CAT(0) space, a geodesic triangle with a strongly geodesically trim side is slim/thin in the sense as used in hyperbolic spaces [16, III.H.1], but this is not true in general metric spaces. Thus, the term 'trim' is still somewhat misleading; a triangle with a trim base can look like a fat triangle that has over tightened its belt, as in Figure 1.

Trimness is closely related to the 'divagation function' of Proposition 5.4.

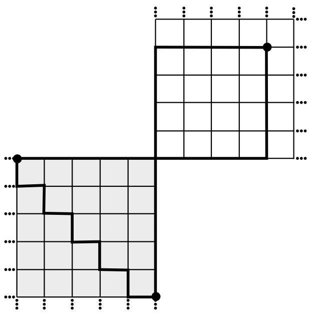

FIGURE 1. A geodesic triangle whose base is contained in a trim subset (shaded) has a small waist.

Durham and Taylor [34] (for subgroups, see also [5] for subsets) make a different choice of how to generalize quasigeodesic quasiconvexity to arbitrary subsets that they call 'stability'. We define stability as follows:

stable: There exists a family of uniformly quasigeodesically quasiconvex uniform quasigeodesics in $\mathcal{X}$ that are contained in $\mathcal{Z}$ and transitive on $\mathcal{Z}$.

The wording here is slightly different than the original definition, where is was additionally required that $\mathcal{Z}$ is an 'undistorted' subset of $\mathcal{X}$. This presupposes that $\mathcal{Z}$ comes equipped with its own metric $d_{\mathcal{Z}}$, and then requires that the inclusion of $\mathcal{Z}$ into $\mathcal{X}$ is a quasiisometric embedding with respect to $d_{\mathcal{Z}}$ and $d_{\mathcal{X}}$. The intuition is that geodesics in $\mathcal{Z}$ are sent to uniform quasigeodesics by a quasiisometric embedding, and we additionally require that the image quasigeodesics are uniformly quasigeodesically quasiconvex. The definition we have given makes sense for arbitrary subsets, without reference to an external metric on $\mathcal{Z}$, and is equivalent to Durham and Taylor's provided that $(\mathcal{Z}, d_{\mathcal{Z}})$ is a quasigeodesic space. Cordes and Hume [23] also modify the definition of stability to make sense for arbitrary subsets by saying that $\mathcal{Z}$ is quasiconvex and that there is a uniformly quasigeodesically quasiconvex family of geodesics in $\mathcal{X}$ that is transitive on $\mathcal{Z}$. Their definition is equivalent to ours.

The initial interest in stability was to characterize convex cocompact subgroups of mapping class groups [34]. It was subsequently shown to characterize convex cocompactness in groups with a nontrivial Morse boundary [22], and was also used to define a quasiisometry invariant notion of dimension [23].

3.1.2. *Efficiency based definitions.* One might have also considered the path system:

- $\mathcal{E}(E)$ the set of all E–efficient segments, with $\mathcal{E} := \cup_{E\geq 0}\mathcal{E}(E)$.

This does not lend itself to a useful notion of 'efficiently quasiconvex'. The problem is that subsegments of E–efficient segments need not be E–efficient. Even in a tree, for example, one can go off in some direction for distance D, then return to the starting point, then go off in some other direction for distance $2D$, to make an 2–efficient path that leaves the D–neighborhood of the geodesic between the endpoints. However, in the tree there is a kind of reversed version of convexity: paths do not need to stay close to the geodesic between the endpoints, but the geodesic does have to stay close to the path. Shchur [71, 39] calls the generalization of this property to hyperbolic spaces the 'anti-Morse property'. Druţu, Mozes, and Sapir [31], see also [5], generalize this idea by requiring that an efficient path with endpoints x and y on $\mathcal{Z}$ contains some point that is close to some point of $\mathcal{Z}$ away from x and y.

recurrent: For some $N > 2$ there exists a function ρ defined by $\rho(E; N) := \sup_\gamma \bar{d}(\gamma, \mathcal{Z}'_\gamma)$, where the supremum is taken over $\gamma \in \mathcal{E}(E)$ with endpoints on $\mathcal{Z}$, and if γ^+, γ^- denote the endpoints of γ then we define $\mathcal{Z}'_\gamma := \mathcal{Z} - \mathcal{N}_{d(\gamma^+,\gamma^-)/N}(\{\gamma^-, \gamma^+\})$. The modified distance $\bar{d}$ is given by:

$$\bar{d}(\gamma, \mathcal{Z}'_\gamma) = \begin{cases} d(\gamma^+, \gamma^-)/N & \text{if } \mathcal{Z}'_\gamma = \emptyset \\ \inf_{w\in\gamma} d(w, \mathcal{Z}'_\gamma) & \text{otherwise} \end{cases}$$

It is clear that for $2 < N < N'$ we have $\rho(E; N') \leq \rho(E; N)$. It is not clear from the definition, but follows from the proof of Theorem 3.2, that if $\rho(E; N')$ exists then so does $\rho(E; N)$, so the existence of a recurrence function does not depend on the choice of $N > 2$, although the equivalence class of the function ρ might. We will assume $N = 3$ and omit it from the notation unless otherwise specified.

The modified distance $\bar{d}$ in the definition allows us to deal with the case that some $\mathcal{Z}'_\gamma$ is empty. This degenerate case only occurs if $\mathcal{Z}$ is bounded and clustered about two points; for instance, if $\mathcal{Z}$ only has two points. If $\mathcal{Z}$ is unbounded or connected with more than one point then no $\mathcal{Z}'_\gamma$ is empty and $\bar{d}$ is the usual distance.

3.1.3. *Sphere-based definitions.* For the next two definitions, we make the convention that the infinum of the empty set is ∞, and let d_r be the induced length metric on $\mathcal{N}^c_r(\mathcal{Z}) \subset \mathcal{X}$, with the convention that d_r takes value ∞ on points from different components.

frilly:[1] There is a function $\mu\colon [0, \infty) \to [0, \infty]$ with $\lim_{r\to\infty} \mu(r) = \infty$ defined by $\mu(r) := \inf_{x,y} d_r(x, y)/d(x, y)$, where the infinum is taken over points $x, y \in \mathcal{S}_r(\mathcal{Z})$ such that $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) > 2r$.

Figure 2 illustrates frilliness.

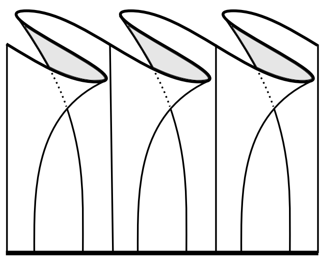

FIGURE 2. The bottom line is a frilly geodesic.

Note that if follows from the definition that if there exists r such that $\mu(r) = \infty$ then $\mu(s) = \infty$ for all $s > r$. In particular, if $\mathcal{Z}$ is bounded then $\mu(r) = \infty$ for all $r \geq (\operatorname{diam} \mathcal{Z})/2$.

More generally, for each $M \in (0, 1]$ and $N \geq 2$ define the (possibly ∞–valued) function $\mu(r; M, N) := \inf_{x,y} d_{Mr}(x, y)/d(x, y)$, where the infinum is taken over points $x, y \in \mathcal{S}_r(\mathcal{Z})$ such

[1]This definition needs to be modified in the case that the projection map is $\pi^\epsilon_{\mathcal{Z}}$ for $\epsilon > 0$. In this case, to guarantee x and y are distinct the condition should be: $\operatorname{diam} \pi^\epsilon_{\mathcal{Z}}(x) \cup \pi^\epsilon_{\mathcal{Z}}(y) > 2(r + \epsilon)$.

that $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) > Nr$. It turns out, see Proposition 8.3, that $\mu(r; M, N) \asymp \mu(r; M', N')$ for all M, N, M', and N', so up to equivalence it is enough to consider $\mu(r) = \mu(r; 1, 2)$.

Morse [56, Lemma 5] proved[2] that geodesics in the hyperbolic plane are frilly as a precursor to proving quasigeodesic quasiconvexity.

Frilliness is closely related to a particular generalized notion of 'divergence'. The concept of divergence has a long history, and many different definitions. The simplest is to define the divergence of a pair of geodesic rays with common basepoint to be the function that takes a positive number r to the infimal length of a path joining the points at distance r along the two rays, with the condition that the path stays outside the ball of radius r about the basepoint. This definition can easily be applied to geodesics by choosing a basepoint and considering the geodesic to be the union of two rays. However, in this case the choice of basepoint can affect the result if the geodesic is not periodic, so a more natural choice is to rechoose the basepoint for each given r so as to minimize the result. This is called 'lower divergence' by Charney and Sultan [19].

In hyperbolic spaces geodesics have at least exponential divergence, while in Euclidean spaces they have linear divergence. Gromov [42, §6.B2(h)] mused that in spaces of nonpositive curvature these could be the only possibilities. However, Gersten [38] gave a coarse geometric definition of divergence and constructed an example of a finite nonpositively curved 2–complex whose universal cover has rays diverging quadratically, but no faster. Gersten [37] and Kapovich and Leeb [49] used divergence to study nonpositive curvature in 3–manifold groups. Macura [52], and, independently, Behrstock and Druţu [10], constructed examples of CAT(0) spaces with arbitrary polynomial degrees of divergence.

In the realm of nonpositive curvature there is a gap between geodesics with linear divergence and those whose divergence grows faster than linear (see Corollary 3.6), even if the gap is not quite so dramatic as Gromov had guessed. It was well understood that the 'Euclidean like' geodesics were those with linear divergence, and so, naturally, the geodesics with nonlinear divergence are the 'hyperbolic like' ones. However, in the absence of a nonpositive curvature hypothesis, divergence functions can have exotic behavior[3], and it turns out that divergence 'not bounded by a linear function', which used to be a commonly used hypothesis, is not sufficient to imply quasigeodesic quasiconvexity. For example, Ol′shanskii, Osin, and Sapir [59] construct a group that contains quasigeodesics whose divergence is linear on an unbounded sequence but not bounded by any linear function. Moreover, the example has an asymptotic cone with no cut point, hence the group has no stable quasigeodesics.

In [Paper A] we clarified this situation by saying that ϕ is *completely superlinear* if for every choice of nonnegative constants C_1, C_2 the set $\{r \in [0,\infty) \mid \phi(r) \leq C_1 r + C_2\}$ is bounded. Equivalently, $\lim_{r\to\infty} \phi(r)/r = \infty$. We then formulate a version of divergence for quasigeodesics, and show that quasigeodesic quasiconvexity is equivalent to completely superlinear divergence.

Tran [79] invented a version of divergence, called 'relative lower divergence', that is equivalent to lower divergence for geodesics but also makes sense for arbitrary subsets of a geodesic metric space. He shows in [77] that the characterization of quasigeodesic quasiconvexity in terms of completely superlinear divergence goes through for this generalized version of divergence.

We propose the following simplified version of relative lower divergence:

divergent: There is a function $\delta\colon [0,\infty) \to [0,\infty]$ with $\lim_{r\to\infty} \delta(r)/r = \infty$ defined by $\delta(r) := \inf_{x,y} d_r(x, y)$, where the infinum is taken over points $x, y \in \mathcal{S}_r(\mathcal{Z})$ such that $d(x, y) \geq 3r$.

[2]What Morse actually proved is stronger. Specifically, he considered all points x and y in $\mathcal{S}_r(\mathcal{Z})$ whose projection diameter is at least 1. In general we cannot assume that closest point projection is so well behaved—for $x \in \mathcal{S}_r(\mathcal{Z})$ the set $\pi_{\mathcal{Z}}(x)$ can have diameter $2r$, so we need $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) > 2r$ even to distinguish x and y as distinct points.

[3]Gruber and Sisto [46] construct groups with even wilder behavior: for any countable collection of subexponential functions they construct a 2–generated group whose divergence exceed each of the given functions on some subsequence, but is at most quadratic on some subsequence.

Thus we say '$\mathcal{Z}$ is divergent' to mean that the lower divergence δ of $\mathcal{X}$ relative to $\mathcal{Z}$ is a completely superlinear function.

As in the frilly case, it is useful to define a family of divergence functions: for all $0 < M \leq 1$ and $N > 2$, define $\delta(r; M, N) := \inf_{x,y} d_{Mr}(x, y)$, where the infinum is taken over points $x, y \in \mathcal{S}_r(\mathcal{Z})$ such that $d(x, y) \geq Nr$.

Also as in the frilly case, all of these functions are equivalent, see Proposition 8.4, so up to equivalence it is enough[4] to consider $\delta(r) = \delta(r; 1, 3)$.

It turns out to be important in the proof of Proposition 8.4 that the ratios $d(x, y)/r$ are bounded away from 2, but here is a compelling geometric justification that the case of $N = 2$ is not suitable for our purposes: Let $\mathcal{X} := \mathbb{E}^2$, and let $\mathcal{Z}$ be a point in $\mathcal{X}$. Then $\mathcal{Z}$ is a frilly/quasigeodesically quasiconvex subset, but $\mathcal{S}_r(\mathcal{Z})$ is just the circle of radius r in the Euclidean plane centered at the point $\mathcal{Z}$, so $\delta(r; M, 2)$ is linear. On the other hand, for $N > 2$ there are no points in $\mathcal{S}_r(\mathcal{Z})$ at distance Nr from each other, so, by convention, $\delta(r; M, N) = \infty$. The $N = 2$ functions are not equivalent to the $N > 2$ functions, and the $N = 2$ functions cannot identify frilly subsets.

Worry not that this example is somehow cheating by taking $\mathcal{Z}$ to be a bounded set. A similar argument applies to $\mathcal{X}$ the universal cover of the wedge of two flat tori, with $\mathcal{Z}$ a connected component of the lift of one of the tori. In this case $\mathcal{Z}$ is again a frilly subset, but the $N = 2$ divergence functions are still linear.

3.1.4. *Contraction.*

contracting: There is a function $\kappa_1 \colon (0, \infty) \to (0, \infty)$ that is continuous, unbounded, nondecreasing, and bounded above by the identity, such that there is a function κ_2 satisfying $\lim_{r\to\infty} \frac{\kappa_2(r)}{\kappa_1(r)} = 0$ defined by $\kappa_2(r) := \sup_{x,y} \operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y)$ where the supremum is taken over points x and y such that $d(x, \mathcal{Z}) \leq r$ and $d(x, y) \leq \kappa_1(d(x, \mathcal{Z}))$.

Remark 3.1. Sometimes when we have κ_1 and κ_2 as in the definition it is convenient to replace κ_2 by $\kappa_2'(r) := \kappa_1(r) \cdot \sup_{s \geq r} \kappa_2(s)/\kappa_1(s)$. Then $\kappa_2'(r) \geq \kappa_2(r)$, κ_2' is nondecreasing, and κ_2'/κ_1 is continuous, nonincreasing, and $\lim_{r\to\infty} \frac{\kappa_2'(r)}{\kappa_1(r)} = 0$.

Special cases of note:

sublinearly contracting: Contracting with $\kappa_1(r) := r$, so that κ_2 is a sublinear function.
semistrongly contracting: Contracting with $\kappa_1(r) := r/2$ and κ_2 bounded.
strongly contracting: Contracting with $\kappa_1(r) := r$ and κ_2 bounded.

For brevity, we say *C–strongly contracting* or *C–semistrongly contracting* for strongly or semistrong contraction, respectively, with κ_2 bounded by C, and we say *κ–contracting* for sublinear contraction with $\kappa_2 = \kappa$. Figure 3, which is [Figure 2, Paper A], illustrates that geodesics in $\mathbb{H}^2$ are strongly contracting.

Minsky [55] proved that geodesics in the thick part of Teichmüller space satisfy a 'contraction property', which is what we have called 'strong contraction'. Masur and Minsky [53] defined a more general contraction property which corresponds in our terminology to (κ_1, κ_2)–contracting with κ_1 linear and κ_2 bounded. By enlarging κ_2 it is easy to see that this is equivalent to semistrong contraction. They show that the image of certain Teichmüller geodesics in the curve complex of a surface satisfy this weaker version of contraction, and use this fact to conclude the curve complex is hyperbolic. In the literature on mapping class groups and related areas [9, 17, 33, 2, 65] the term 'contraction' continues to be used for a version of the property we are calling semistrong contraction, where the projection map is not necessarily closest point projection.

Bestvina and Fujiwara [13] define a contraction property for geodesics in CAT(0) spaces. Their property is exactly strong contraction.

[4]In Tran's definition [77] the divergence is the whole family of functions $\{\delta(r; M, N)\}$ as the parameters $0 < M \leq 1$ and $2 \leq N$ vary, up to a notion of equivalence of families. Even his notion of equivalence of functions is different than ours. However, the two notions of equivalence of functions coincide on divergence functions. This uses the fact that $\delta(r) \geq 3r$ and an argument similar to the proof of Proposition 8.4. Then using Proposition 8.4, it follows that $\delta(r) \asymp \delta'(r)$ for divergence functions δ and δ' if and only if the families $\{\delta(r; M, N)\}$ and $\{\delta'(r; M, N)\}$ are equivalent families.

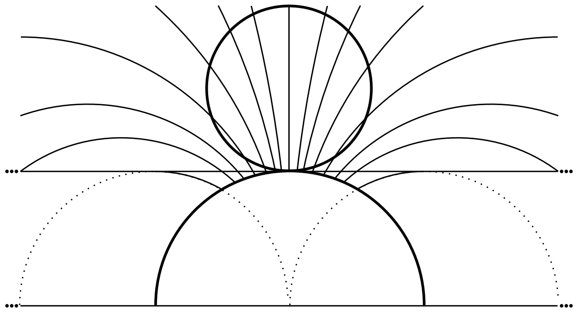

FIGURE 3. Geodesics in the hyperbolic plane are strongly contracting. Even the projection of a tangent horoball to a geodesic has bounded diameter.

Algom-Kfir [3] proves an Out(F_n) analogue of Minsky's result by showing that the axis of fully irreducible free group automorphisms are strongly contracting in outer space. This marks the first use of the term 'strong contraction' to distinguish between this property and the semistrong version. She also provides an explicit proof that strong contraction implies quasigeodesic quasiconvexity.

The most general notion of 'contraction' given above was introduced in our paper [Paper A].

3.1.5. *Relationships between the various definitions.* Figure 4 shows the relationships between the various notions defined in this section. Unlabelled arrows are implications that follow immediately from the definitions. Double arrows are proven in this thesis. Dashed arrows are implications that are true in CAT(0) spaces, but not in general. The boxed properties are equivalent.

THEOREM 3.2. *We have the following relationships, depicted in Figure 4, between properties of an arbitrary subset $\mathcal{Z}$ of a geodesic metric space $\mathcal{X}$. Furthermore, all the implications are effective. (Effective bounds for selected properties are stated in Corollary 3.5.) Finally, this list is complete, in the sense that implications that do not follow from those listed here are not true in general.*

① *Semistrongly contracting $\implies$ logarithmically contracting. (See remark following Lemma 2.4 of [Paper D].)*
② *If $\mathcal{X}$ is CAT(0) then geodesically trim $\implies$ contracting and strongly geodesically trim $\implies$ semistrongly contracting. (Proposition 8.12)*
③ *Quasigeodesically quasiconvex $\implies$ sublinearly contracting. ([Theorem 1.4, Paper A])*
④ *If $\mathcal{X}$ is CAT(0) then recurrent $\implies$ strongly contracting. ([Paper B])*
⑤ *Contracting $\implies$ quasigeodesically trim. (Proposition 8.6)*
⑥ *Sublinearly contracting $\implies$ recurrent for all $N > 2$. ([5, Theorem 3.2])*
⑦ *Recurrent for some $N > 2$ $\implies$ frilly. (Proposition 8.14)*
⑧ *Divergent $\implies$ quasigeodesically quasiconvex. (Proposition 8.15)*
⑨ *Quasigeodesically trim $\implies$ quasigeodesically quasiconvex, and geodesically trim $\implies$ geodesically quasiconvex. (Proposition 8.10)*
⑩ *Contracting $\implies$ frilly $\implies$ quasigeodesically quasiconvex and divergent. (Propositions 8.7 and 8.9)*
❷ *frilly $\implies$ divergent. (Proposition 8.8)*
❶ *If $\phi\colon \mathcal{Z} \to \mathcal{X}$ is a quasiisometric embedding between geodesic metric spaces then $\phi(Z)$ is stable if and only if $\phi(Z)$ is quasigeodesically quasiconvex and $\mathcal{Z}$ is hyperbolic. (Proposition 8.16)*

PROOF. The proofs of the various positive implications are referenced in the statement of the theorem.

It is well known that even in CAT(0) spaces geodesic quasiconvexity does not imply quasigeodesic quasiconvexity. To see that semistrong contraction does not imply strongly contracting

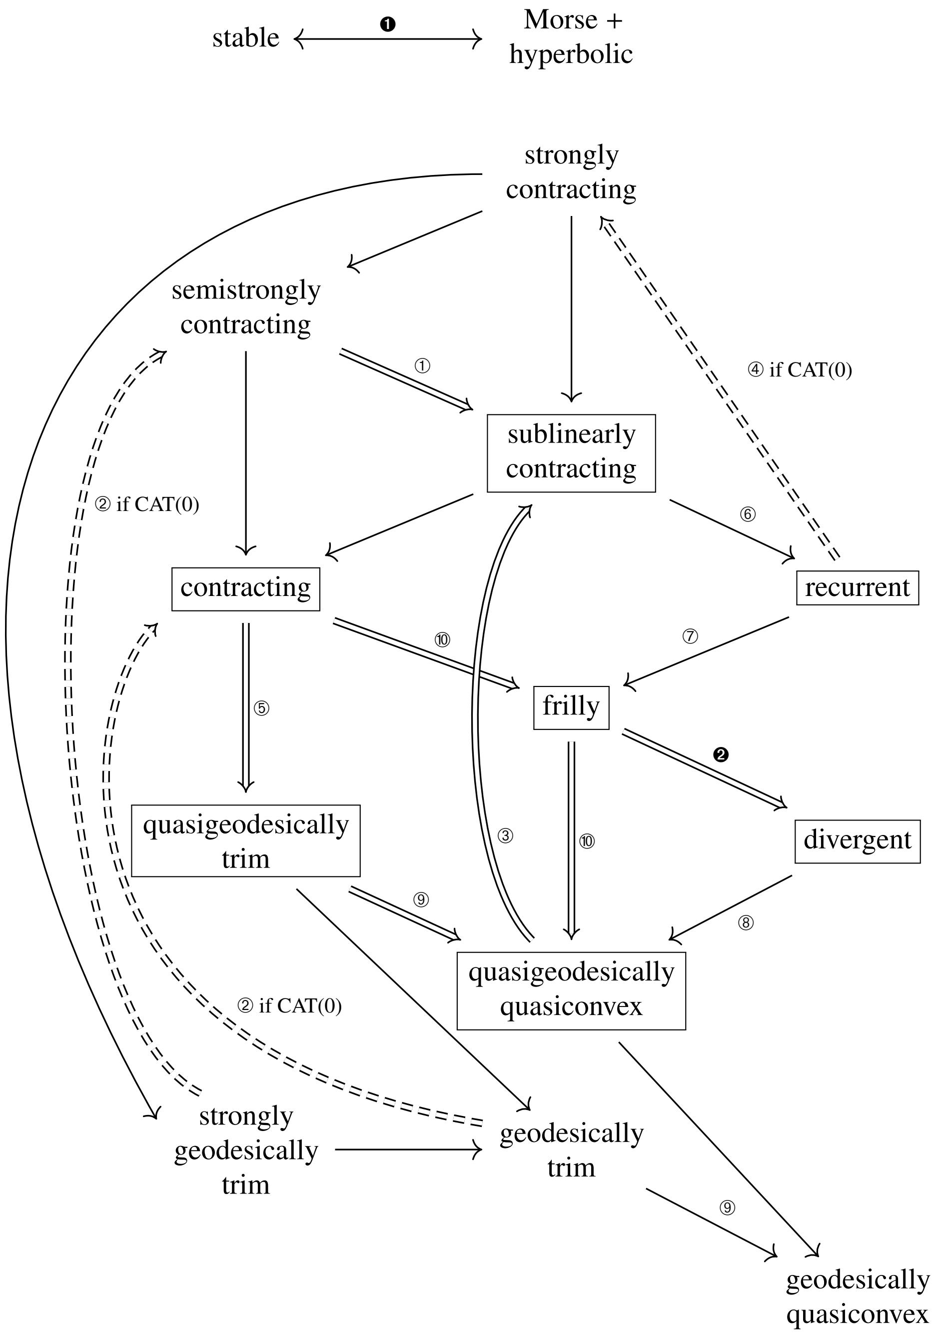

FIGURE 4. Relationships between the various definitions.

or strongly geodesically trim see Example 3.2 of [Paper A]. The remaining nonimplications are demonstrated in the examples of Section 8.1. □

DEFINITION 3.3. A subset of a geodesic metric space is *Morse* if it enjoys the following equivalent properties:

- contracting
- recurrent
- frilly
- divergent
- quasigeodesically trim
- quasigeodesically quasiconvex

COROLLARY 3.4. *If X is hyperbolic all of the properties in Figure 4 are equivalent. If X is CAT(0) then, of the properties in Figure 4, stability is the strongest, geodesic quasiconvexity is the weakest, and everything else is equivalent.*

Proof. The proof that quasiconvexity implies strong contraction in hyperbolic spaces is easy and is left to the reader. Moreover, a quasiconvex subset of a hyperbolic space is itself hyperbolic, so it is also stable.

In the CAT(0) case, we already know quasiconvex does not imply Morse. Also, if we consider $\mathcal{X} = \mathcal{Z} = \mathbb{E}^2$, then $\mathcal{Z}$ is Morse in $\mathcal{X}$ but not stable. □

Attributions/remarks on parts of Theorem 3.2:

②&④ Sultan [74] showed that in CAT(0) spaces quasigeodesically quasiconvex quasigeodesics are strongly contracting, but the proof does not give effective equivalence. In later unpublished work [75] he showed effective control and also that strongly geodesically trim quasigeodesics in CAT(0) spaces are semistrongly contracting. ② generalizes the latter to arbitrary subsets, and to work with a weaker hypothesis. The former we prove with the combination of ③ and ④. The advantage of this approach over Sultan's is that the part using the CAT(0) hypothesis, ④, is quite easy, while ③ works in general, so we get a clearer understanding of where the hard part of the argument is and where the part is that actually requires the CAT(0) hypothesis.

⑤ [Theorem 4.8, Paper H] is equivalent to 'sublinearly contracting implies quasigeodesically trim'. However, the proof given there is not effective.

⑥&⑦ Druţu, Mozes, and Sapir [31] claimed an equivalence between quasigeodesic quasiconvexity and recurrence for quasigeodesics. Their proof was not effective, and involved passage through some other intermediate characterizations. Aougab, Durham, and Taylor [5] found an error in one of these intermediate steps, and repaired it by giving an effective equivalence between quasigeodesic quasiconvexity and recurrence. Druţu, Mozes, and Sapir have also published their own correction [30]. Both [5] and [30] use our implication ③ in an essential way for the equivalence of quasigeodesic quasiconvexity and recurrence.

We prove ⑦ in Proposition 8.14 by adapting the argument of [5] to our definition of frilly and to arbitrary subsets.

[5, Theorem 3.2] proves ⑥ for quasigeodesics. They divide a quasigeodesic into 'left', 'middle', and 'right' pieces. For a general contracting subset $\mathcal{Z}$ and efficient path γ with endpoints $\gamma^+, \gamma^- \in \mathcal{Z}$ take the 'left' piece to be $\mathcal{Z}_l := \mathcal{Z} \cap \mathcal{N}_{d(\gamma^+,\gamma^-)/N}(\gamma^-)$, the 'right' piece to be $\mathcal{Z}_r := \mathcal{Z} \cap \mathcal{N}_{d(\gamma^+,\gamma^-)/N}(\gamma^+)$, and the 'middle' to be $\mathcal{Z}' := \mathcal{Z} - (\mathcal{Z}_l \cup \mathcal{Z}_r)$. With these definitions, the rest of their argument goes through without change.

⑧ The converse implication is proven for quasigeodesics in [Paper A] and generalized to arbitrary subsets in [77], but these proofs are not effective.

⑩ Algom-Kfir [3] proved strongly contracting geodesics are quasigeodesically quasiconvex. One can generalize her argument to show contraction implies quasigeodesic quasiconvexity. Our proof for ⑩ essentially factors the generalized argument into two steps to involve frilliness and also optimize the two steps so that the resulting effective bounds are sharp.

Corollary 3.5. *We have the following effective bounds among the Morse properties. For simplicity we only state the quasiconvexity gauge for* $(E, 0)$*–locally efficient paths. Shorthand versions of these inequalities are shown in Figure 5. In the figure, double arrows indicate that the given effective bounds are verified by example to be optimal.*

(1) κ–contracting implies μ–frilly for $\mu > r/\kappa'(r)$, where $\kappa'(r) := r \cdot \sup_{s\geq r} \kappa(s)/s$ as in Remark 3.1.
(2) κ–contracting implies ρ–recurrent for $\rho(E) < \sup\{r \mid \frac{\kappa'(r)}{r} \geq \frac{1}{8E}\}$ for κ' as in (1).
(3) ρ–recurrent implies μ–frilly for $\mu(r) > \inf\{s \mid r \leq 1 + \lim_{t\to s^+} \rho(1+2t)\}$.
(4) μ–frilly implies δ–divergent for $\delta(r) > r\mu(r)$.
(5) μ–frilly implies χ–$\mathcal{L}$–quasiconvex for $\chi(E) \leq (1+2E) \cdot \sup\{r \mid \mu(r) \leq E\}$.
(6) δ–divergent implies χ–$\mathcal{L}$–quasiconvex for $\chi(E) < E \cdot \sup\{s \mid \delta(s) \leq 2(E^2+3E+1)s\}$.
(7) χ–$\mathcal{L}$–quasiconvex implies κ–contracting for $\kappa(r) < \sup\{s \leq 4r \mid s \leq 18\chi(\frac{12r}{s})\}$.

(8) *τ–quasigeodesically trim implies χ–$\mathcal{L}$–quasiconvex for* $\chi(E) \leq \sup\{r \mid r \leq (E^2 + 1)(\tau(r; E, 0))\}$.

(9) *κ–contracting implies τ–quasigeodesically trim for* $\tau(r; L, A) \prec \left(\frac{\mathrm{Id}^2}{\kappa'}\right)^{-1}(r)$ *for κ' as in* (1).

Example 8.2 gives a family of examples for which the relations provided by items (1), (2), (3), and (4) are equivalences. Example 8.1 gives a family of examples for which the relations provided by items (3), and (4), (5) are equivalences.

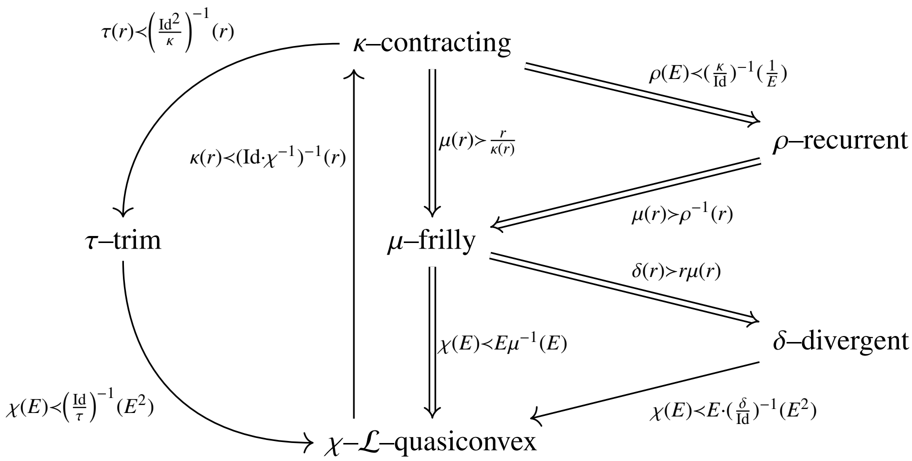

FIGURE 5. Effective relationships.

Since Morse subsets of CAT(0) and hyperbolic spaces are strongly contracting, Corollary 3.5 implies the following positive answer to [77, Question 1.8].

COROLLARY 3.6. *Strongly contracting sets, in particular, Morse subsets of hyperbolic and CAT(0) spaces, are at least linearly frilly and at least quadratically divergent.*

3.2. Quasiisometry invariance and Morse subgroups. It is well-known that the property of quasigeodesic quasiconvexity is invariant under quasiisometries, and it is easy to check that the quasigeodesic quasiconvexity gauge is invariant, up to equivalence. By Theorem 3.2, this implies that frilliness, divergence, contraction, etc., are also invariant. However, going back and forth with Theorem 3.2 leads to some loss in quantitative control between these properties under quasiisometry. In fact, it can be checked directly that the equivalence classes of the frilliness and divergence gauges are invariant under quasiisometry.

PROPOSITION 3.7. *Let $\phi\colon \mathcal{X} \to \mathcal{Y}$ be a quasiisometry between geodesic metric spaces and let $\mathcal{Z} \subset \mathcal{X}$. If $\mathcal{Z}$ is χ–quasigeodesically quasiconvex, μ–frilly, and δ–divergent then $\phi(\mathcal{Z})$ is χ'–quasigeodesically quasiconvex, μ'–frilly, and δ'–divergent, respectively, with $\chi \asymp \chi'$, $\mu \asymp \mu'$, and $\delta \asymp \delta'$.*

The trick to make this proposition work for μ and δ is to use Proposition 8.3 and Proposition 8.4 to consider $\mu(r; 1, N)$ and $\delta(r; 1, N)$, respectively, where N is large compared to the quasiisometry constants, and see that these functions push forward under the quasiisometry to give bounds on their $N = 2, 3$ counterparts in $\mathcal{Y}$. Further details are left to the reader.

COROLLARY 3.8. *Let $\mathcal{X}$ be a geodesic metric space and $\mathcal{Q}$ a quasiisometry type of Morse subspace of $\mathcal{X}$. For a Morse subset $\mathcal{Z} \subset \mathcal{X}$ in the quasiisometry class $\mathcal{Q}$, let $\chi_{\mathcal{Z}}$, $\mu_{\mathcal{Z}}$, and $\delta_{\mathcal{Z}}$ be quasigeodesic quasiconvexity, frilliness, and divergences gauges, respectively, of $\mathcal{Z}$. Then the following set of triples of equivalence classes of functions is a quasiisometry invariant of $\mathcal{X}$:*

$$\{([\chi_{\mathcal{Z}}], [\mu_{\mathcal{Z}}], [\delta_{\mathcal{Z}}]) \mid \mathcal{Z} \subset \mathcal{X} \textit{ Morse, with } \mathcal{Z} \in \mathcal{Q}\}$$

Note that although we can *bound* any one of these gauges in terms of one of the others, Corollary 3.5 does not determine its equivalence class, so the invariant described in Corollary 3.8 is, a priori, finer than if we replaced the triple by the equivalence class of just one of the three types of gauges. It would be interesting to know examples where the triples actually give a strictly finer invariant.

The case that Q consists of (quasi)geodesics and the triple is replaced by the equivalence class of the divergence gauge is the *divergence spectrum of $\mathcal{X}$* introduced in [78].

In contrast to Proposition 3.7, it is not so easy to argue directly that contraction is preserved by quasiisometries. In fact, the equivalence class of the contraction gauge is not preserved, in general: we exhibit in [Paper C] an example of a finitely generated group G and an infinite cyclic subgroup H such that there are finite generating sets S and S' of G such that H is a strongly contracting geodesic in $\mathrm{Cay}(G,S)$ and a geodesic that is logarithmically but not strongly contracting in $\mathrm{Cay}(G,S')$.

When the contraction gauge is sufficiently far from linear we can at least get a bound for the contraction gauge of the image set. The example of the previous paragraph shows this bound is sharp for strongly contracting sets.

Lemma 3.9. *Suppose $\mathcal{Z}\subset\mathcal{X}$ is κ–contracting with $\lim_{r\to\infty}\frac{\kappa(r)\cdot\log r}{r}=0$. If $\phi\colon\mathcal{X}\to\mathcal{Y}$ is a quasiisometry then $\phi(\mathcal{Z})$ is κ'–contracting for $\kappa'\prec\kappa\cdot\log$. In particular, if $\mathcal{Z}$ is strongly contracting then $\phi(\mathcal{Z})$ is at worst logarithmically contracting.*

Corollary 3.10. *Suppose G is a finitely generated group and $h\in G$ is an infinite order element. Suppose there exists a finite generating set for G in which $\langle h\rangle$ has contraction gauge κ for κ superlogarithmic. Then there does not exist a finite generating set for G such that $\langle h\rangle$ is strongly contracting.*

Proof of Lemma 3.9. Let $\bar\phi$ be a quasiinverse to ϕ. Suppose that ϕ and $\bar\phi$ are both (L,A)–quasiisometries. The map $\psi:=\phi\circ\pi_{\mathcal{Z}}\circ\bar\phi$ has the property that $d(y,y')\leq d(y,\phi(\mathcal{Z}))/L^2$ implies $\operatorname{diam}\psi(y)\cup\psi(y')\prec\kappa(d(y,\phi(\mathcal{Z})))$. Now if $d(y,y')\leq d(y,\phi(\mathcal{Z}))$ then a geodesic between y and y' can be covered by $\log d(y,\phi(\mathcal{Z}))$ many closed balls whose radii are at most $1/L^2$ times the distance from their respective centers to $\phi(\mathcal{Z})$. Since the geodesic from y to y' is contained in the $2d(y,\phi(\mathcal{Z}))$–neighborhood of $\phi(\mathcal{Z})$, this implies $\operatorname{diam}\psi(y)\cup\psi(y')\prec\kappa(d(y,\phi(\mathcal{Z}))\cdot\log d(y,\phi(\mathcal{Z}))$. On the other hand, if we apply this argument to a point $y'\in\pi_{\phi(\mathcal{Z})}(y)$ we get an equivalent bound for $\operatorname{diam}\pi_{\phi(\mathcal{Z})}(y)\cup\psi(y)$. Similarly for y' we get $\operatorname{diam}\pi_{\phi(\mathcal{Z})}(y)\cup\psi(y)\prec\kappa(d(y',\phi(\mathcal{Z}))\cdot\log d(y',\phi(\mathcal{Z}))$, but since $d(y',\phi(\mathcal{Z}))\leq 2d(y,\phi(\mathcal{Z}))$, this bound is equivalent to the previous two. Combining the three bounds gives us:

$$d(y,y')\leq d(y,\phi(\mathcal{Z}))\implies \operatorname{diam}\pi_{\phi(\mathcal{Z})}(y)\cup\pi_{\phi(\mathcal{Z})}(y')\prec\kappa(d(y,\phi(\mathcal{Z})))\cdot\log d(y,\phi(\mathcal{Z}))$$

Therefore, $\phi(\mathcal{Z})$ is κ'–contracting for some $\kappa'\prec\kappa\cdot\log$, which is sublinear by hypothesis. □

Proposition 3.7 immediately suggests the following definitions:

Definition 3.11. A subgroup H of a finitely generated group G is *Morse* if it is Morse as a subset of some, equivalently of every, Cayley graph of G with respect to a finite generating set.

Moreover, the equivalence classes of the frilliness, divergence, and quasigeodesic quasiconvexity gauges of H are well-defined properties of $H<G$, independent of the choice of generating set of G.

We call an element $g\in G$ a *Morse element* (resp. contracting element, divergent element) if $\langle g\rangle$ is an infinite Morse subgroup. By the example above, there is no well-defined property of being a *strongly contracting element* of G, only of being a strongly contracting element with respect to some Cayley graph. More generally, if G acts on a space $\mathcal{X}$ we say g is Morse, strongly contracting, etc with respect to $G\curvearrowright\mathcal{X}$ if $\langle g\rangle$ has an unbounded orbit with the corresponding property.

Let us recall some examples of nonelementary Morse subgroups. Quasiconvex subgroups of hyperbolic groups are Morse; indeed, in any Cayley graph they are strongly contracting. The factor groups of a free product, or, more generally, of a graph of groups with finite edge groups, are also strongly contracting in every Cayley graph. Druţu and Sapir [32] showed that peripheral subgroups of relatively hyperbolic groups are Morse. We will see in Section 6 that a more general class of subgroups called hyperbolically embedded subgroups are Morse. In large classes of right-angled Artin groups every infinite index Morse subgroups is free (follows from results of [51], see [77]). Free groups are hyperbolic, so in these RAAGs being infinite index and Morse is equivalent to being stable. The same is true for mapping class groups [50]. RAAGs and mapping class groups both belong to a class of groups called hierarchically hyperbolic groups. Russell, Spriano, and Tran [65] show that certain hierarchically hyperbolic groups have the property that Morse subgroups satisfy a variation of the contraction property (similar to the one that appears in the proof of Lemma 3.9) and pose the question of characterizing when a hierarchically hyperbolic group has the property that all of its infinite index Morse subgroups are stable. Tran [77] and Genevois [36] characterize which special subgroups of right-angled Coxeter groups are Morse.

The following extends [31, Lemma 3.25] from the case of cyclic subgroups:

Lemma 3.12 ([Lemma 6.6, Paper H]). *Suppose $\phi\colon \mathcal{X} \to \mathcal{Y}$ is a coarse Lipschitz uniformly proper*[5] *map between geodesic metric spaces and $\mathcal{Z} \subset \mathcal{X}$. If $\phi(\mathcal{X})$ is Morse and $\mathcal{Z}$ is Morse then $\phi(\mathcal{Z})$ is Morse. If $\phi(\mathcal{Z})$ is Morse then $\mathcal{Z}$ is Morse. Moreover, the quasigeodesic quasiconvexity gauge of $\mathcal{Z}$ bounds that of $\phi(\mathcal{Z})$, and vice versa, up to functions depending on ϕ.*

In particular, if G is a finitely generated group acting properly on a geodesic metric space $\mathcal{X}$ then the orbit map of G into $\mathcal{X}$ is coarse Lipschitz and uniformly proper with respect to the word metric on G coming from any finite generating set.

Corollary 3.13. *Suppose G is a finitely generated group acting properly on a based geodesic metric space $(\mathcal{X}, o)$ and $H < G$. If $G.o$ is Morse in $\mathcal{X}$ and H is Morse in G then $H.o$ is Morse in $\mathcal{X}$. If $H.o$ is Morse in $\mathcal{X}$ then H is Morse in G.*

Corollary 3.13 implies the main result of [5], which is the statement that if $H.o$ is stable in $\mathcal{X}$ then H is stable in G.

Corollary 3.14. *Let H be a subgroup of a finitely generated group G. If H is finitely generated then every subgroup of H that is Morse in G is also Morse in H. If H is Morse in G then every Morse subgroup of H is also Morse in G.*

4. Application: The geometry of graphical small cancellation groups

We designed the general version of the contraction condition introduced in the previous section as a tool for understanding the Morse directions in infinitely presented graphical small cancellation groups. These groups generalize classical (metric) small cancellation groups as follows: Fix a finite alphabet $\mathcal{A}$ and a potentially infinite sequence of directed connected graphs (Γ_i). Let $\Gamma := \coprod_i \Gamma_i$. Assign to each directed edge an element of $\mathcal{A}$. Define a finitely generated group $G = G(\Gamma)$ by taking the generating set to be $\mathcal{A}$, and by taking as relators the words in $\mathcal{A}$ that can be read along undirected cycles in Γ, where we consider that a directed edge labelled a traversed against its orientation contributes the letter a^{-1}. This gives a finitely generated group G, and $\mathcal{X} = \mathcal{X}(\Gamma) := Cay(G, \mathcal{A})$ is a Cayley graph of G. The 'classical' case is the case that all of the Γ_i are circles, so each component Γ_i gives a single relator, up to inversion and cyclic permutation.

The interest in such a definition is that by imposing a small cancellation condition on the labeling one guarantees that the graph Γ 'lives inside' $\mathcal{X}$, so we can engineer groups containing exotic geometries. The precise small cancellation condition and the quality of the corresponding embedding of Γ have evolved. The first version was used by Gromov [44] to construct finitely generated groups with a Cayley graph containing an expander graph. These groups, now known as 'Gromov monsters', were the first known counterexamples to the Baum-Connes conjecture with coefficients [47]. The theory was further developed by Ollivier [58], Arzhantseva and

[5] We say ϕ is *uniformly proper* if there is a nondecreasing unbounded function $\psi\colon [0,\infty) \to [0,\infty)$ such that $d(\phi(x), \phi(x')) \geq \psi(d(x, x'))$ for all $x, x' \in \mathcal{X}$.

Delzant [6], and Gruber [45]. There are two small cancellation conditions we are interested in here, the $Gr'(1/6)$ and $C'(1/6)$ conditions. Definitions can be found in [Section 2, Paper C]. The $C'(1/6)$ condition implies the $Gr'(1/6)$ condition. These conditions guarantee there are isometrically embedded copies of Γ in the Cayley graph of the group:

THEOREM 4.1 ([46, Lemma 2.15]). *Let $\Gamma := (\Gamma_i)$ be a $Gr'(1/6)$–labelled graph. For every i and every vertex $v \in \Gamma_i$ there is a unique label-preserving isometric embedding of Γ_i into $\mathcal{X}(\Gamma)$ sending v to the vertex corresponding to the identity element of $G(\Gamma)$.*

In the classical case, when Γ is finite the resulting group is hyperbolic. This remains true for graphical small cancellation groups as well, so the 'interesting' case is when Γ is infinite (more specifically, when Γ has infinitely generated first homology). In this case the group G is not finitely presentable, so cannot be hyperbolic. For intuition, let us specialize to the case that there are infinitely many components Γ_i and all are circles. The small cancellation condition forces $|\Gamma_i| \to \infty$ as $i \to \infty$, since the circles cannot have relatively long segments with common labels. The geometry is similar to the toy model from Example 8.1: if $\mathcal{Z}$ is a geodesic in the Cayley graph $\mathcal{X}$, at every vertex it encounters infinitely many loops coming from copies of the Γ_i. Fix such a loop γ. Suppose that $\alpha := \gamma \cap \mathcal{Z}$ contains an edge and let $\beta := \overline{\gamma - \alpha}$. Since $\mathcal{Z}$ is geodesic and γ is isometrically embedded, α and β are segments. The segment β provides a detour around the subsegment α of $\mathcal{Z}$. The efficiency of this detour is $|\beta|/|\alpha|$, so if we would like to show that $\mathcal{Z}$ is Morse, we must at least show that for every such γ the proportion of γ that intersects $\mathcal{Z}$ is small. It turns out that this kind of consideration gives a complete answer to the question of which geodesics are Morse. Let us say that $\mathcal{Z}$ is *locally κ–contracting* if for every embedded copy of every Γ_i its intersection with $\mathcal{Z}$ is κ–contracting as a subset of Γ_i.

THEOREM 4.2 (cf [Theorem 4.1, Paper C]). *Let $\Gamma = (\Gamma_i)$ be $Gr'(1/6)$ labelled. Let $\mathcal{Z}$ be a geodesic in $\mathcal{X}(\Gamma)$. Then there exists κ such that $\mathcal{Z}$ is locally κ–contracting if and only if there exists κ' such that $\mathcal{Z}$ is κ'–contracting, and when this is true we have $\kappa' \asymp \kappa$.*

The motivation for graphical small cancellation was to construct monsters, but a more positive description might be that it provides a toolset for constructing metric phenomena in Cayley graphs. Building the desired phenomenon into a Cayley graph may be difficult because of homogeneity. Instead, one can build the phenomenon into a graph, which should be much easier, and then invoke small cancellation to embed that graph into a Cayley graph. Theorem 4.2 tells us that this embedding preserves directional curvature.

Using this construction, we give the first example of a finitely generated group G and an infinite order element g such that g is strongly contracting for one choice of generating set and not strongly contracting for another. Recall the discussion following Definition 3.11. In general, it is difficult to tell if a Morse geodesic is strongly contracting or not, but we can build graphs specifically to verify or negate strong contraction. The trick then is to do this for graphs that are related in such a way that they yield isomorphic groups, see [Theorem 4.19, Paper C].

For another example of this workflow, we can verify that essentially every sublinear function arises as the contraction function of an infinite cyclic subgroup in a finitely generated group. Consult Example 8.1. This example can be discretized by restricting to $s \in 2\mathbb{Z}$, rounding the length of each γ_s to the nearest integer, and subdividing to make the result a graph Γ. In this case the graph Γ has a single, infinite component. Then $\mathcal{Z}$ sits as a geodesic in Γ, and we label all edges of $\mathcal{Z}$ by a letter a. Further, one can write down by hand labels for the γ_s in letters b and c such that the resulting labelling on the graph Γ is $Gr'(1/6)$. Since $\mathcal{Z}$ is geodesic in Γ and Γ isometrically embeds in $\mathcal{X}$, $\mathcal{Z}$ is a geodesic in $\mathcal{X}$, and it is rational, by construction, since the element a acts on it by translation. Then the computations from Example 8.1 combined with Theorem 4.2 show that $\mathcal{Z}$ is κ–contracting for some sublinear $\kappa \asymp \phi^{-1}$.

We will see more examples of this kind of construction in Section 6.

The previous example was convenient in two ways:

- We were able to easily find a rational geodesic.
- The graph was simple enough to find explicit $Gr'(1/6)$–labellings.

It turns out that we can replicate both of these facts for fairly general families of graphs.

LEMMA 4.3 ([Lemma 5.2, Paper C]). *Let* $\Gamma = (\Gamma_i)$ *be a disjoint union of finite connected graphs* Γ_i *with a* $Gr'(1/6)$*–labelling by a finite alphabet* $\mathcal{A}$*. Every infinite cyclic subgroup of* $G = G(\Gamma)$ *is at finite Hausdorff distance from a rational geodesic in* $Cay(G, \mathcal{A})$*.*

Along the way to proving Lemma 4.3 we also prove a result that is reminiscent of hyperbolic groups. Recall that the *translation length* of an element g of a group G with respect to a finite generating set $\mathcal{A}$ is $\tau_{\mathcal{A}}(g) := \lim_{n\to\infty} |g^n|/n$, where $|\cdot|$ is word length with respect to $\mathcal{A}$. Recall further that in hyperbolic groups infinite order elements have rational translation lengths with bounded denominators [43, 76, 29].

THEOREM 4.4 ([Theorem 5.4, Paper C]). *Let* $\Gamma = (\Gamma_i)$ *be a* $Gr'(1/6)$*–labelled graph, labelled by a finite alphabet* $\mathcal{A}$*, such that for each* i *the component* Γ_i *is finite. For every infinite order* $g \in G(\Gamma)$*,* $\tau_{\mathcal{A}}(g) \geq 1/3$ *and is rational.*

The existence of small cancellation labellings is furnished by the following result of Osajda [60]. Recall that the *girth* of a graph is the length of the shortest embedded cycle.

THEOREM 4.5. *Let* $\Gamma = (\Gamma_i)$ *be a sequence of bounded valence, connected, finite graphs, such that* $(\mathrm{girth}(\Gamma_i))$ *is an unbounded sequence and the ratios* $\frac{\mathrm{diam}(\Gamma_i)}{\mathrm{girth}(\Gamma_i)}$ *are bounded. Then there exists an infinite subsequence* $\Gamma' := (\Gamma_{i_j})$ *and a finite set* $\mathcal{A}$ *such that* Γ' *admits a* $C'(1/6)$*–labelling over* $\mathcal{A}$*.*

In [Section 6, Paper C] we give a simple method of combining two labellings, a so-called *push-out labelling*, on Γ that inherits properties from both, provided these properties pass to refined labellings. Here it is important that we use the $C'(1/6)$ condition, as it has this inheritance property, but the $Gr'(1/6)$ condition does not.

An example of another property of labellings is that of being *nonrepetitive*, which means that there does not exist an embedded path γ in Γ such that the label on γ can be written ww for some word w. Graph theory techniques [4] show that a bounded valence graph admits a nonrepetitive labelling by a finite alphabet, so by taking the push-out of such a labelling with the Osajda labelling we get a nonrepetitive $C'(1/6)$ labelling. For such a labelling every infinite order element g is strongly contracting, since by Lemma 4.3, up to finite Hausdorff distance we may assume that $\langle g \rangle$ is a geodesic, but since the labelling is nonrepetitive the longest subsegment that intersects an embedded component of Γ has length less than $2|g|$. Using this argument we can construct many groups that have the apparently nice property that every infinite order element is strongly contracting, but contain very strange embedded geometries. If, for example, Γ is an expander, then this construction shows:

COROLLARY 4.6 (cf [Corollary 6.10, Paper C]). *There exist Gromov monster groups such that every infinite order element is strongly contracting.*

5. Generalizing hyperbolic boundaries

Let X be proper hyperbolic space. One can define its Gromov boundary ∂X by picking a basepoint o and taking Hausdorff equivalence classes of geodesic rays based at o. One can then topologize this set of equivalence class by declaring that two points are close if representative geodesic rays fellow travel for a long time. This topology turns out to be independent of the choice of basepoint and invariant under quasiisometry. In particular, this means that a hyperbolic group has a well-defined homeomorphism type of boundary, independent of the choice of generating set of the group.

Quasiisometry invariance of the boundary fails for visual boundaries of CAT(0) groups [26]; that is, a group G can act geometrically on CAT(0) spaces X and X' whose visual boundaries are not homeomorphic to one another, so G does not have a well-defined CAT(0) boundary.

5.1. Morse boundaries. Charney and Sultan [19] wondered whether the flat subspaces in CAT(0) space are responsible for the failure of quasiisometry invariance, and whether perhaps there was a hyperbolic-like subset of the boundary that is preserved. Specifically, they define the *contracting boundary* $\partial_c X$ of a proper CAT(0) space to be the subset of the visual boundary consisting of points whose representative geodesic rays are Morse/(strongly) contracting sets. They ask if the subspace topology on the contracting boundary is preserved by quasiisometries. We gave a negative answer[6] to this question in [Paper G]. They go on to define a new topology on $\partial_c\mathcal{X}$ that they show is quasiisometry invariant. The idea of their construction is that in a hyperbolic space geodesics are uniformly contracting, that is, there is a single sublinear function bounding all of their contraction gauges. In CAT(0) spaces this is no longer true, but it is true within a Hausdorff equivalence class of ray. Therefore, given a contraction function κ, one can specify a subset $\partial_c^\kappa\mathcal{X}$ of the contracting boundary whose points are represented by rays that are κ–contracting. Such a subset can be topologized as in the hyperbolic case. Charney and Sultan define a topology on $\partial_c\mathcal{X}$ by realizing it as the direct limit of the system $\partial_c^\kappa\mathcal{X}$, related by inclusion, as κ varies over all possible contraction functions, topologizing the $\partial_c^\kappa\mathcal{X}$ as above, and then taking the direct limit topology, $\mathcal{DL}$. Their main result is that the resulting topology is invariant under quasiisometries.

When the space is hyperbolic there is a single κ such that $\partial_c^\kappa\mathcal{X}$ is the entire visual boundary, and in this case the contracting boundary is exactly the hyperbolic boundary. However, in nonhyperbolic spaces the direct limit topology is more mysterious. Murray [57] showed the contracting boundary is nonempty and compact if and only if the space is hyperbolic, which is true if and only if the contracting boundary is second countable. In fact, in the nonhyperbolic case it can even fail to be first countable.

Cordes [21] generalized Charney and Sultan's approach to work for Morse geodesic rays in arbitrary proper geodesic metric spaces, so this boundary has become known as the *Morse boundary*.

Theorem 5.1 ([19], [21]). *The Morse boundary of a proper geodesic metric space is independent of the choice of basepoint and invariant under quasiisometry. In particular, every finitely generated group G has a well-defined Morse boundary $\partial_M G$ obtained by taking $\partial_M G := \partial_c^{\mathcal{DL}}\mathcal{X}$, where $\mathcal{X}$ is a Cayley graph of G.*

In addition to being quasiisometry invariant, Morse boundaries enjoy some other properties reminiscent of hyperbolic boundaries. For instance [21, Proposition 4.2] implies:

Proposition 5.2. *Let H be a Morse subgroup of a finitely generated group G. The inclusion of H into G extends to an embedding $\partial_M H \hookrightarrow \partial_M G$.*

It also can be deduced from his methods that the image of $\partial_M H$ in $\partial_M G$ is precisely the limit set $\Lambda(H)$ of H acting on G (see also [77]). This proves Item (5) of Theorem 1.1 for the case of Morse boundaries.

As a topological space, the Morse boundary provides a quasiisometry invariant of a group: two groups with nonhomeomorphic Morse boundaries are not quasiisometric to each other. The converse is true after equipping the Morse boundary with some additional structure:

Theorem 5.3 ([18]). *Let $\phi\colon\ \partial_M\mathcal{X} \to \partial_M\mathcal{X}'$ be a homeomorphism between nontrivial Morse boundaries of proper geodesic metric spaces admitting geometric group actions such that ϕ is quasi-Möbius and 2–stable. Then ϕ is induced by a quasiisometry between $\mathcal{X}$ and $\mathcal{X}'$.*

A Morse boundary is *nontrivial* if it contains more than two points. A homeomorphism is quasi-Möbius if it coarsely preserves a certain version of a cross ratio defined in [18]. It is 2–stable if given χ there exists χ' such that for every pair of points (x, y) in $\partial_M\mathcal{X}$ that can be connected in $\mathcal{X}$ by a χ–quasigeodesically quasiconvex geodesic, the points $(\phi(x), \phi(y))$ in $\partial_M\mathcal{X}'$ can be connected by a χ'–quasigeodesically quasiconvex geodesic.

See Cordes's [20] survey article on Morse boundaries for further results and references.

[6]The question remains open for spaces $\mathcal{X}$ admitting a geometric group action, as our examples are not homogeneous.

5.2. Contracting boundaries revisited. Despite some success, the Morse boundary can be difficult to work with due to the definition of the topology as a direct limit. By Murray's [57] result, even for innocuous looking examples like $\mathbb{Z}^2 * \mathbb{Z}$ the direct limit topology fails to be first countable, in contrast to the boundary of a hyperbolic space, which is metrizable. This is despite the fact that $\mathbb{Z}^2 * \mathbb{Z}$ admits two other boundaries that would seem to capture its hyperbolic-like behavior: the Bowditch boundary for $\mathbb{Z}^2 * \mathbb{Z}$ viewed as a relatively hyperbolic group, and the boundary of the Bass-Serre tree for $\mathbb{Z}^2 * \mathbb{Z}$ viewed as an HNN extension of $\mathbb{Z}^2$ over the trivial group.

This led us to consider whether there could be a coarser quasiisometry invariant topology on the sat of Hausdorff equivalence classes of contracting geodesic rays.

The proof of quasiisometry invariance of hyperbolic boundaries uses hyperbolicity in two key places:

- A quasiisometry between hyperbolic spaces takes geodesics to within uniformly bounded distance of geodesics.
- There is a clear distinction between fellow-travelling and non-fellow-travelling.

The first one is not true for general spaces, and our topology will instead deal directly with quasigeodesics. The second, however, is true in general spaces when at least one of the rays is Morse:

Proposition 5.4 (cf [Paper H]). *Given a contraction gauge κ and constants $L \geq 1$, $A \geq 0$ there exists a threshold distance T and superlinear function λ such that if $\mathcal{Z}$ is κ–contracting and γ is a continuous (L, A)–quasigeodesic ray with basepoint in $\mathcal{Z}$ then either β is contained in the T–neighborhood of $\mathcal{Z}$ or there is a first time t_0 such that $d(\beta(t_0), \mathcal{Z}) = T$ and for $t \geq t_0$ we have $d(\beta(t), \mathcal{Z}) \geq \lambda(\operatorname{diam} \pi_Z(\beta(t)) \cup \pi_Z(\beta(t_0)))$.*

We call λ a *divagation function*. It quantifies the idea that a quasigeodesic ray cannot wander slowly away from a contracting set; it can either closely fellow travel or it can move away in an asymptotically 'orthogonal' direction. We can then say unambiguously how long β fellow-travels $\mathcal{Z}$: it does so up until the point where it hits the divagation threshold and starts moving away quickly.

If $\mathcal{Z}$ is a quasigeodesic ray and β is also contracting, then proving such an estimate is easy if we allow the divagation function λ to depend on the contracting gauges of both $\mathcal{Z}$ and β. The key point of Proposition 5.4 is that λ does not depend on a contraction function for β. This allows us to define a new topology on $\partial_c \mathcal{X}$ as follows: points of $\partial_c \mathcal{X}$ are Hausdorff equivalence classes of continuous contracting quasigeodesic rays based at a common basepoint. This set is in fact equivalent to the underlying set in Charney and Sultan's definition. Given $[\alpha] \in \partial_c \mathcal{X}$ represented by a κ–contracting geodesic ray α, we say that $[\beta]$ belongs to $U(r)$ if every quasigeodesic ray β in $[\beta]$ closely fellow travels α for distance at least r, where 'closely' means the divagation threshold of Proposition 5.4, which depends on κ and the quasigeodesic constants of β, but not on α or β.

The main result of [Paper H] is:

Theorem 5.5. *The sets $U(r)$ defined above form a neighborhood basis at $[\alpha]$ for a topology, the* topology of fellow traveling quasigeodesics*, denoted $\mathcal{FQ}$, on $\partial_c \mathcal{X}$. This topology is independent of the choice of basepoint and the choice of geodesic representative α of $[\alpha]$. It is also invariant under quasiisometries.*

Furthermore, if $\mathcal{X}$ admits a geometric group action then $\partial_c^{\mathcal{FQ}} \mathcal{X}$ is metrizable.

Henceforth we use 'contracting boundary' to mean $\partial_C \mathcal{X} := \partial_c^{\mathcal{FQ}} \mathcal{X}$ and 'Morse boundary' to mean $\partial_M \mathcal{X} := \partial_c^{\mathcal{DL}} \mathcal{X}$.

It is straightforward to prove that $\mathcal{DL}$ is a refinement of $\mathcal{FQ}$, so the identity map on the underlying set gives a continuous map $\partial_M \mathcal{X} \to \partial_C \mathcal{X}$, and that both coincide with the hyperbolic boundary when $\mathcal{X}$ is hyperbolic.

There are also other 'well-understood' cases in which our topology matches the expected topology:

Theorem 5.6 ([Theorem 7.6, Paper H]). *If G is hyperbolic relative to subgroups that are either hyperbolic or have empty contracting boundary, then the natural map from the contracting boundary of G into the Bowditch boundary is an embedding.*

Theorem 5.7. *If G is a finitely generated group acting cocompactly on a tree $\mathcal{T}$ with virtually Abelian vertex stabilizers and finite edge stabilizers then the natural map $\partial_C G \to \partial\mathcal{T}$ is an embedding.*

These two results show that for $G = \mathbb{Z} * \mathbb{Z}$ the topology $\mathcal{FQ}$ on $\partial_c G$ is the 'right' topology, in the sense that it agrees with the topology coming from the other two natural hyperbolic-like sources.

We prove the statement from Theorem 1.1 Item (5):

Proposition 5.8 (cf [Proposition 6.8, Paper H]). *If H is a Morse subgroup of a finitely generated group G then the inclusion map extends to an embedding $\partial_C H \hookrightarrow \partial_C G$. Moreover, the image coincides with $\Lambda(H)$.*

We also have analogues of Murray's [57] dynamical results from the CAT(0) case:

Theorem 5.9 ([Corollary 8.4 and Theorem 9.4, Paper H]). *If G is a finitely generated group and $\partial_C G$ has more than two points the action of G on $\partial_C G$ is minimal and satisfies a weak version of North-South dynamics.*

Pal and Pandey [62] further investigate the dynamics, and prove:

Theorem 5.10. *Let H be a finitely generated subgroup of a finitely generated group G such that $\Lambda(H) \subset \partial_C G$ is compact and has more than two points. Then $H \curvearrowright \Lambda(H)$ is a convergence action and H is acylindrically hyperbolic.*

As in the case of the Morse boundary, the contracting boundary of a group is nonempty and compact if and only if the group is hyperbolic. A component of the proof of this fact is the following statement, which is of independent interest:

Theorem 5.11 (cf [Theorem 10.1, Paper H]). *Let $\mathcal{X}$ be a geodesic metric space admitting a geometric action. The following are equivalent:*

(1) $\mathcal{X}$ is hyperbolic.
(2) Geodesic segments in $\mathcal{X}$ are uniformly quasigeodesically quasiconvex.
(3) Every geodesic ray in $\mathcal{X}$ is Morse.

The implications (1) $\implies$ (2) $\implies$ (3) is easy. We prove (3) $\implies$ (1). The implication (2) $\implies$ (1) was the main result of [15].

Note that hyperbolicity is not guaranteed if only the rational geodesics are Morse, nor if rational geodesics are strongly contracting, eg. Corollary 4.6.

6. Stronger properties

Ballmann [8, Theorem 3.5] proved that if a proper nonelementary CAT(0) space $\mathcal{X}$ contains a periodic geodesic that does not bound a half-flat, then every cocompact group of isometries of $\mathcal{X}$ contains a non-Abelian free subgroup. An element acting by translation on this periodic geodesic is strongly contracting for this action, and we might wonder if the existence of a strongly contracting element is enough to guarantee the existence of a free subgroup if we do not assume $\mathcal{X}$ is CAT(0). It turns out that this is true, by the quasitree construction of [12]. This fits in with several other related ideas, which we now recall.

Definition 6.1 ([27]). H is a *hyperbolically embedded subgroup* of G if there exists a subset S of $G - H$ such that $S \sqcup H$ generates G, the Cayley graph $\mathcal{X}$ of G with respect to the (infinite) generating set $S \sqcup H$ is hyperbolic, and $(H, \hat{d})$ is a proper metric space, where $\hat{d}(h, h')$ is the length of the shortest path in $\mathcal{X}$ between h and h' that does not use any edges of the induced (complete) subgraph of $\mathcal{X}$ with vertex set H.

Let G be a group acting on a space $\mathcal{X}$ and $\mathcal{Y}$ a subspace of $\mathcal{X}$. The action is *acylindrical along* $\mathcal{Y}$ if for every $r > 0$ there exist $R, N > 0$ such that $d(y, y') > R$ for $y, y' \in \mathcal{Y}$ implies $\#\{g \in G \mid d(y, gy) \leq r,\ d(y', gy') \leq r\} \leq N$. It is *acylindrical* if it is acylindrical along $\mathcal{X}$.

An action on a hyperbolic space is *elementary* if the axes of all loxodromic elements are coarsely equivalent; otherwise the action is nonelementary.

A group G is *acylindrically hyperbolic* if it admits a nonelementary acylindrical action on a hyperbolic space.

An element $h \in G$ is called *generalized loxodromic* if there exists a nonelementary acylindrical action G on a hyperbolic space $\mathcal{X}$ such that h acts loxodromically.

An element $h \in G$ is a *WPD element (satisfies the weak proper discontinuity condition)* for an action of G on a hyperbolic space $\mathcal{X}$ if for every $x \in \mathcal{X}$ and $r > 0$ there exists $n > 0$ such that $\#\{g \in G \mid d(x, gx) \leq r,\ d(h^n x, gh^n x) \leq r\} < \infty$. An action of G on a hyperbolic space $\mathcal{X}$ is called a *WPD* action if it is nonelementary and every loxodromic element is WPD.

The WPD property was introduced by Bestvina and Fujiwara [14]. Acylindrical hyperbolicity was systematized by Osin [61].

Theorem 6.2 ([73]). *Hyperbolically embedded subgroups of finitely generated groups are Morse.*

Theorem 6.3 ([27, Theorem 8.7]). *Let G be a nonelementary group acting on a hyperbolic space such that the element h is loxodromic and WPD. Then there exists positive n such that the normal closure of h^n is a non-Abelian free group.*

Definition 6.4. Given a nontrivial element $h \in G$, if h is contained in a unique maximal virtually cyclic subgroup of G then we call that subgroup the *elementary closure* of h and denote it $E(h)$.

Theorem 6.5. *Let h be an infinite order element of a nonelementary group G. Consider the following properties.*

(1) There exists an action of G on a space $\mathcal{X}$ such that orbits of $\langle h \rangle$ are unbounded and strongly contracting.
(2) There exists a Cayley graph $\mathcal{X}$ of G (with respect to a possibly infinite generating set) that is hyperbolic, $G \curvearrowright \mathcal{X}$ is acylindrical, and h acts loxodromically.
(3) h is generalized loxodromic.
(4) There is a nonelementary action of G on a hyperbolic space that is acylindrical along $\langle h \rangle$ orbits.
(5) There is an action of G on a hyperbolic space such that h is a WPD element.
(6) $E(h)$ exists and is hyperbolically embedded.
(7) $E(h)$ exists and is Morse.
(8) h is a Morse element.

Then conditions (1)–(6) *are equivalent,* (7) *and* (8) *are equivalent, and* (6) $\implies$ (7).

Proof. It is easy to see condition (i) $\implies$ (i+1). (2) $\implies$ (1) since quasigeodesics in hyperbolic spaces are strongly contracting. (1) $\implies$ (4) follows from [12]. (3) $\implies$ (6) is proved in [72]. Theorem 6.2 implies (6) $\implies$ (7). [61] proves the equivalence of (2), (3), (5), and (6). (7) and (8) are equivalent because $\langle h \rangle$ is a finite index subgroup of $E(h)$. □

By the previous two results and the existence of Tarski monsters whose elements are Morse, we see that a generalized loxodromic element is Morse, but the converse need not be true. By Theorem 3.2 if h is a Morse element of G then for any finite generating set of G there is a sublinear κ such that h is κ–contracting. If κ is bounded then h is strongly contracting, so it is generalized loxodromic. One might wonder if h must be generalized loxodromic if κ is unbounded but grows slowly enough. The answer, however, is 'no':

Theorem 6.6 ([Theorem 6.4, Paper C]). *For every unbounded sublinear κ there exists a 2–generator, infinitely presented graphical small cancellation group containing an infinite order element h that is κ'–contracting for some $\kappa' \asymp \kappa$ and no virtually cyclic subgroup containing h is hyperbolically embedded.*

To prove this theorem we construct examples with the property that all powers of h contain torsion in their normal closure, ruling out the conclusion of Theorem 6.3. Abbott and Hume [1] revisit this result and use similar methods to produce torsion-free examples.

Conversely, for a generalized loxodromic element h there exists a space (the one defining it as a generalized loxodromic) on which it acts as a strongly contracting element. For any finite generating set, h is Morse in the corresponding Cayley graph. One might wonder whether the generalized loxodromic hypothesis imposes some upper bound on such contraction functions that is better than just being sublinear. Say, logarithmic, as in Lemma 3.9? The answer to the latter question, at least, is 'no': Abbott and Hume construct graphical small cancellation groups containing generalized loxodromic elements h such that the contraction gauge κ for h satisfies $\kappa(r) \succ r/\log(r)$.

Let us now return to the quasitree construction mentioned at the beginning of this section. Let G be a group acting on a geodesic metric space $\mathcal{X}$ with basepoint o such that there exists an infinite order element $h \in G$ such that $\langle h\rangle.o$ is a strongly contracting quasigeodesic in $\mathcal{X}$. Then $E(h)$ exists and $\mathcal{H} := E(h).o$ is also a strongly contracting subset, as is any translate of $\mathcal{H}$. Let $\mathbb{Y}$ be the set of distinct G–translates of $\mathcal{H}$. Then for $\mathcal{Y}, \mathcal{Y}', \mathcal{Y}'' \in \mathbb{Y}$ we have $\pi_{\mathcal{Y}}(\mathcal{Y}')$ has bounded diameter when $\mathcal{Y} \neq \mathcal{Y}'$ and we can define projection distances for $\mathcal{Y}' \neq \mathcal{Y} \neq \mathcal{Y}''$:

$$d^{\pi}_{\mathcal{Y}}(\mathcal{Y}', \mathcal{Y}'') := \operatorname{diam} \pi_{\mathcal{Y}}(\mathcal{Y}') \cup \pi_{\mathcal{Y}}(\mathcal{Y}'')$$

Bestvina, Bromberg, and Fujiwara [12] axiomatize properties of these projections and projection distances and show that whenever the axioms are satisfied there exists a quasitree encoding the geometry of $\mathbb{Y}$. As a consequence, we can prove the following.

Definition 6.7. Let G be a group acting on a based geodesic metric space $(\mathcal{X}, o)$ such that there is an infinite order element h with $\mathcal{H} := E(h).o$ strongly contracting in $\mathcal{X}$. An (A, B)–*good path* is a finite concatenation of geodesics $\cdots + [a_i, b_i] + [b_i, a_{i+1}] + [a_{i+1}, b_{i+1}] + \cdots$ for $1 \leq i \leq n-1$ such that for each $1 \leq i \leq n$ there exists $g_i \in G$ such that:

- $a_i, b_i \in g_i\mathcal{H}$
- $g_i\mathcal{H} \neq g_{i+1}\mathcal{H}$
- $d(a_i, b_i) \geq A$
- $d^{\pi}_{g_i\mathcal{H}}(b_{i-1}, a_i) \leq B$
- $d^{\pi}_{g_i\mathcal{H}}(b_i, a_{i+1}) \leq B$

For appropriate choices of A and B, good paths are essentially the same as the *admissible paths* of [13] and [81].

Lemma 6.8. *With the setup as in Definition 6.7, if $\mathcal{H}$ is C–strongly contracting then there exists a C' such that for all sufficiently large B and sufficiently small A, if γ is an (A, B)–good path and α is a geodesic connecting $g_1\mathcal{H}$ to $g_n\mathcal{H}$, then α passes within distance C' of a_i and b_j for all $2 \leq i \leq n$ and $1 \leq j \leq n-1$.*

Corollary 6.9. *There exists C'' such that if x_i and y_i are points of α that are C'–close to a_i and b_i, respectively, then the subsegment of α between x_i and y_i is within Hausdorff distance C'' of a subsegment of $g_i\mathcal{H}$ between a_i and b_i.*

The interiors of segments $[b_i, a_{i+1}]$, however, do not necessarily stay close to α.

When $n \leq 3$ Lemma 6.8 can be proven directly from the definition of strong contraction. For large n some care is required because an inductive argument introduces a finite additive error at each step, and these errors erode the buffer provided by the constant B. This can be avoided by invoking the improved projection axioms of [11], as we do in [Paper F].

7. Application: Growth of groups

Let $(\mathcal{X}, o)$ be a based (pseudo)metric space, and let $\mathcal{Y}$ be a subset of $\mathcal{X}$ such that the intersection of $\mathcal{Y}$ with every metric ball is finite. Then we define the exponential growth rate of

$\mathcal{Y}$ to be:

$$\delta_{\mathcal{Y}} := \limsup_{r\to\infty} \frac{\log \#\mathcal{Y} \cap \bar{B}_r(o)}{r}$$

Let G be a finitely generated group acting metrically properly on a based geodesic metric space $(\mathcal{X}, o)$, in the sense that for every r the set $\{g \in G \mid d(o, g.o) \leq r\}$ is finite. We can define a pseudometric d on G by $d(g, h) = d_{\mathcal{X}}(g.o, h.o)$, and the growth rate of a subset $S \subset G$ with respect to this pseudometric is the same as the growth rate of $S.o$ as a subset of $\mathcal{X}$. If N is a normal subgroup of G then there is an induced pseudometric on G/N given by $\bar{d}(gN, hN) := \min_{n,n'\in N} d(gn, hn')$.

Definition 7.1. The *cogrowth of G/N with respect to $G \curvearrowright \mathcal{X}$* is the ratio δ_N/δ_G, for growth rates with respect to $G \curvearrowright \mathcal{X}$.

Definition 7.2. Let G act metrically properly on a based geodesic metric space $\mathcal{X}$. $G \curvearrowright \mathcal{X}$ is a *growth tight action* if for every infinite normal subgroup N of G, the growth rate of G with respect to d is strictly greater than the growth rate of G/N with respect to the induced pseudometric $\bar{d}$ on G/N.

When $\mathcal{X}$ is the Cayley graph of a finite rank free group G with respect to a basis, Grigorchuk [40] showed that the cogrowth of every proper quotient is strictly greater than 1/2, and Grigorchuk and de la Harpe [41] showed $G \curvearrowright \mathcal{X}$ is growth tight. They also showed that $G' \curvearrowright \mathcal{X}'$ is not growth tight if $G' = G \times G$ and $\mathcal{X}'$ is the Cayley graph of G' with respect to the generating set obtained by taking the union of bases for the factors. They ask:

(1) Is the action of a nonelementary hyperbolic group on (any) one of its Cayley graphs growth tight?
(2) Does growth tightness for actions on Cayley graphs depend on the choice of generating set?
(3) Is the cogrowth of a quotient of a nonelementary hyperbolic group by an infinite subgroup strictly greater than 1/2 with respect to an action on (any) one of its Cayley graphs?

Arzhantseva and Lysenok [7] answered the first question affirmatively. Sambusetti and collaborators [70, 69, 67, 28] proved that the action of a negatively curved Riemannian manifold on its Riemannian universal cover is a growth tight action. Sambusetti [68] also proved that any action on one of its Cayley graph of a nonelementary group that splits over a finite subgroup is growth tight. Sabourau [66] proved that a geometric action of a group on a hyperbolic space is growth tight. Note that Sabourau's result is not implied by Arzhantseva and Lysenok's, because there is no result saying growth tightness is preserved by equivariant quasiisometry.

The main result of [Paper D] generalizes all of these results. A simplified statement is:

Theorem 7.3. *If G is a finitely generated group acting geometrically on a geodesic metric space $\mathcal{X}$ with a strongly contracting element then $G \curvearrowright \mathcal{X}$ is growth tight.*

In [Paper E] we answer Grigorchuk and de la Harpe's second question:

Theorem 7.4. *Let $G_i \curvearrowright \mathcal{X}_i$ be finitely many geometric actions of finitely generated groups on geodesic metric spaces, each with a strongly contracting element. Let G be the product of the G_i. Let $\mathcal{X}$ be the product of the $\mathcal{X}_i$ with the L^p metric. Then $G \curvearrowright \mathcal{X}$ is not growth tight for $p = 1$ and growth tight for $p > 1$.*

In particular, if each $\mathcal{X}_i$ is a Cayley graph of G_i with respect to some finite generating set, then the L^p metric on $\mathcal{X}$ agrees with a word metric on G when $p = 1$ and when $p = \infty$, so, for example:

Corollary 7.5. *$F_2 \times F_2$ has a generating set for which it is growth tight and a generating set for which it is not growth tight.*

The answer to the third question is also 'yes'. There are two independent results that imply this as a special case, one of Matsuzaki, Yabuki, and Jaerisch [54], and one of ours from [Paper F], a simplified version of which is:

THEOREM 7.6. *If G is a finitely generated group acting geometrically on a geodesic metric space $\mathcal{X}$ with a strongly contracting element then for every infinite normal subgroup N of G the quotient G/N has cogrowth strictly greater than 1/2.*

We have claimed that Theorem 7.3 and Theorem 7.6 are 'simplified' theorems. The simplification comes by assuming that the action of G on $\mathcal{X}$ is geometric. Our techniques actually work in a more general setting than that of cocompact actions, but we need a technical condition to replace cocompactness. To see why some condition is necessary, consider, for example $\mathbb{Z}$ acting parabolically on the hyperbolic plane. In the induced metric, $\mathbb{Z}$ has exponential growth, but this is entirely due to the distortion of the embedding, not on any structural properties of $\mathbb{Z}$. Similarly, if G is a relatively hyperbolic group then it admits a cusp-uniform action on a hyperbolic space, but by modifying the cusp geometry we can arrange that the growth rate of a parabolic subgroup equals the growth rate of the entire group, and that growth tightness then fails. In [Paper D] we introduce the following growth-theoretic measure of the distortion of the group orbit: for constant $M \geq 0$ we define the *complementary growth* of $G \curvearrowright \mathcal{X}$, with respect to parameter M, to be the growth rate of the subset of the orbit $G.o$ consisting of points $g.o$ such that there exists a geodesic γ in $\mathcal{X}$ between o and $g.o$ that stays outside the M–neighborhood of $G.o$ except near the endpoints of γ, see [Definition 6.2, Paper D] for a precise definition.

DEFINITION 7.7. Let G be a finitely generated group acting metrically properly on a based geodesic metric space $(\mathcal{X}, o)$. We say $G \curvearrowright \mathcal{X}$ has *complementary growth gap* if there exists $M \geq 0$ such that δ_G is strictly greater than the complementary growth of $G \curvearrowright \mathcal{X}$ with parameter M.

Clearly if $G \curvearrowright \mathcal{X}$ is cobounded, or, more generally, has a quasiconvex orbit, then $G \curvearrowright \mathcal{X}$ has zero complementary growth, so $G \curvearrowright \mathcal{X}$ has complementary growth gap as long as $\delta_G > 0$. For cusp-uniform actions on hyperbolic spaces, complementary growth gap implies the *parabolic growth gap* condition of [28]. We also prove, using a theorem of Eskin, Mirzakhani, and Rafi [35], that the action of the mapping class group of a hyperbolic surface on its Teichmüller space with the Teichmüller metric has complementary growth gap.

THEOREM 7.8 (Main theorems of [Paper D] and [Paper F]). *Let G be a finitely generated group acting metrically properly on a based geodesic metric space (X, o) with complementary growth gap and with a strongly contracting element. Then $G \curvearrowright \mathcal{X}$ is growth tight and for every infinite normal subgroup the quotient has cogrowth strictly greater than 1/2.*

The idea of the proof of growth tightness has 3 steps:

(1) If $G \curvearrowright \mathcal{X}$ has a strongly contracting element then every infinite normal subgroup contains a strongly contracting element, so it suffices to assume N is the normal closure of a strongly contracting element. Moreover, by passing to a high power of the element we may assume that its translation length is much larger then the strong contraction constant of its elementary closure. Let $\mathcal{H}$ denote the orbit of the elementary closure $E(h)$ of the strongly contracting element $h \in N$.
(2) Take a minimal section $\mathcal{Y}$ of the quotient map $G \to G/N$. Minimality implies that $\mathcal{Y}$ has small projection to every translate of $\mathcal{H}$. Furthermore, the growth rate of $\mathcal{Y}$ in $\mathcal{X}$ is the same as the growth rate of G/N in the induced quotient pseudometric.
(3) Show that there is a tree's worth of copies of $\mathcal{Y}$ embedded in a metrically good way in $\mathcal{X}$. To be specific, we send tuples of $\mathcal{Y}$ into $\mathcal{X}$ by $(y_1, y_2, \dots, y_n) \mapsto y_1hy_2h\cdots y_nh.o \in \mathcal{X}$. To see that this is a 'good' metric embedding we apply the results on good paths from the previous section. Then we argue that the growth rate of a tree's worth of $\mathcal{Y}$ is strictly greater than the growth rate of $\mathcal{Y}$.

In the course of the argument, we also observe that whenever the translation length of h is long enough with respect its contraction constant, the growth rate of the conjugacy class $[h]$ is exactly $\delta_G/2$. This partially generalizes results of Parkkonen and Paulin [63] and Huber [48], and inspires the cogrowth part of the theorem.

To prove the cogrowth bound, we attempt to repeat the strategy for growth tightness using $[h]$ in place of $\mathcal{Y}$. However, it is definitely not true as in (2) above that the projection of $[h]$ to translates of $\mathcal{H}$ is small, and this breaks the embedding part of the argument in (3). We fix this issue by constructing a particular subset of $[h]$ that includes at least half of the points of $[h]$ in every ball about o and has the property that the desired map $(y_1, y_2, \dots, y_n) \mapsto y_1 h y_2 h \cdots y_n h.o \in \mathcal{X}$ is the good metric embedding we need. Since half of $[h]$ still has growth rate $\delta_G/2$ this completes the proof.

In comparison, Matsuzaki, Yabuki, and Jaerisch study actions of a group G on a hyperbolic space $\mathcal{X}$ satisfying a condition called *divergence type*, and one of their conclusions is the strict cogrowth bound for actions of divergence type on hyperbolic spaces. Cocompact actions have divergence type, so the special case of their theorem where $\mathcal{X}$ is the Cayley graph of a hyperbolic group answers Grigorchuk and de la Harpe's question. Our theorem has a stronger hypothesis *if we restrict $\mathcal{X}$ to be hyperbolic*, but our theorem also applies to nonhyperbolic spaces.

The complementary growth gap condition seems like a very natural condition to impose for growth-theoretic questions. Following our introduction of this property, Yang [82] embarked on a systematic study of such actions and proved, again when the action contains a strongly contracting element, properties similar to those for actions of hyperbolic groups such as pure exponential growth and genericity of contracting elements [80]. Coulon, Dougall, Schapira, and Tapie [25] also use this complementary growth gap condition to study the other end of the cogrowth spectrum; they show when $\mathcal{X}$ is hyperbolic and $G \curvearrowright \mathcal{X}$ has complementary growth gap then the cogrowth of G/N equals 1 if and only if G/N is amenable, generalizing a result of Grigorchuk [40] in the free group case and Coulon, Dalbo, and Sambusetti [24] in the case of hyperbolic groups.

8. Appendix

In the name of completeness, in this appendix we provide details of the proof of Theorem 3.2. In Section 8.1 we consider two families of examples. These demonstrate that the implications that do not follow from those depicted in Figure 4 are not true. They also verify that some of the effective bounds in Corollary 3.5 are sharp. In Section 8.2 we show that the choice of parameters in the various functions, eg. the frilliness and divergence gauges, does not matter, up to equivalence of functions. In Section 8.3 we prove Theorem 3.2, one implication at a time. With the exception of Proposition 8.6, each of the implications is either easy or is similar to arguments that can be found in the literature, see remarks following Theorem 3.2 for attributions. We present the details in order to make effective bounds explicit, and, when possible, optimal. The overall proof is also simplified and clarified by factoring some of the pieces through frilliness.

8.1. Examples.

Example 8.1 (Line with detours). Let ϕ be a continuous increasing positive function with $\phi(r) \geq r$ and with $\lim_{r\to\infty} \phi(r)/r = \infty$. Let $\mathcal{X}$ be the space built from $\mathcal{Z} := \mathbb{R}$ by attaching, for each $s \in [1, \infty)$, a new segment γ_s of length $\phi(s)$ connecting $-s/2$ and $s/2$.

Note that $\mathcal{Z}$ cannot be semistrongly contracting or strongly geodesically trim because both of these conditions imply that there is a uniform bound on the diameter of projections of singletons, but in this example for each s the midpoint of γ_s has projection diameter equal to s.

Contraction gauge κ. Given $r > 0$, we must consider points $x \in \bar{N}_r(\mathcal{Z})$) and the projection diameters of balls around such points. For s such that $\phi(s) > 4r$ and $x \in \gamma_s \cap \bar{N}_r(\mathcal{Z})$, the projection of $\bar{N}_r(x)$ to $\mathcal{Z}$ is a single point. Consider s such that $\phi(s) \leq 4r$. Let y be the point halfway around γ_s, and let x be a point one quarter of the way around γ_s. Then $d(x, y) = d(x, \mathcal{Z})$ and $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) = s$. This is the worst case; thus, $\kappa(r) = \phi^{-1}(4r) \asymp \phi^{-1}(r)$. Superlinearity of ϕ implies sublinearity of κ.

Recurrence gauge ρ. For every s we have $d(\gamma_s, \mathcal{Z} - (\mathcal{N}_{s/3}(-s/2) \cup \mathcal{N}_{s/3}(s/2))) = s/3$, realized by either of the two endpoints of γ_s. For $E \geq 1$, γ_s is E–efficient when $\phi(s)/s \leq E$. Superlinearity of ϕ implies there exists $S := \sup\{s \mid \phi(s)/s \leq E\}$. Therefore, $\rho(E) \leq S/3$. Continuity of ϕ implies this is the least upper bound for ρ, so $\rho(E) \asymp \sup\{s \mid \phi(s)/s \leq E\}$.

$\mathcal{L}$–quasiconvexity gauge χ. The path γ_s is locally $(\phi(s)/s)$–efficient and achieves distance $\phi(s)/2$ from $\mathcal{Z}$. Thus, an E–locally efficient path is confined to $\mathcal{Z} \cup \coprod_{\{s \mid \phi(s)/s \leq E\}} \gamma_s$, so $\chi(E) = \frac{1}{2}\phi(\sup\{s \mid \phi(s)/s \leq E\}) \asymp \phi(\sup\{s \mid \phi(s)/s \leq E\})$.

Frilliness gauge μ. Superlinearity of ϕ implies there exists S such that for all $s \geq S$ we have $\phi(s) \geq 3s$. Suppose $s > 2r \geq S$. Then there exist $x, y \in \gamma_s \cap \mathcal{S}_r(\mathcal{Z})$ such that $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) = s > 2r$, so $\mu(r) \leq d_r(x,y)/d(x,y)$. We know $d_r(x,y) = \phi(s) - 2r$. The definition of S implies $d(x,y) = s + 2r$. Thus $\mu(r) \leq (\phi(s) - 2r)/(s + 2r)$. We have this estimate for all $s > 2r$, and the infinum of these bounds occurs by substituting $2r$ for s, hence $\mu(r) \leq \frac{1}{2} \cdot \frac{\phi(2r)}{2r} - \frac{1}{2}$. Continuity of ϕ implies this is the least upper bound for μ, so $\mu(r) \asymp \frac{\phi(r)}{r}$. Superlinearity of ϕ implies $\lim_{r\to\infty} \mu(r) = \infty$.

Divergence gauge δ. By superlinearity of ϕ, for sufficiently large S and all $s \geq S$ we have $\phi(s) \geq 5s$. Suppose $r \geq S$, and suppose γ_s has points x and y in $\gamma_s \cap \mathcal{S}_r(\mathcal{Z})$ that participate in the definition of $\delta(r)$. This happens when $d(x,y) \geq 3r$, which occurs if and only if $s \geq r$. For such x and y, we have $d_r(x,y) = \phi(s) - 2r$. Thus, $\delta(r) = \inf_{s\geq r} \phi(s) - 2r = \phi(r) - 2r \asymp \phi(r)$.

Trimness gauge τ. Fix $E \geq 1$. Consider $r > 0$ and a point x such that $0 < d(x, \mathcal{Z}) \leq r$. The point x lies on some γ_s. Let β be a path from x to $\mathcal{Z}$. Either β goes through a point of $\pi_Z(x)$ or $d(x,\mathcal{Z}) < \phi(s)/2$ and β goes through the point $y \in \gamma_s \cap \mathcal{Z}$ at distance s from $z := \pi_Z(x)$. Thus, it suffices to consider $d(x,\mathcal{Z}) < \phi(s)/2$ and the path β that is the longer of the two subsegments of γ_s connecting x to $\mathcal{Z}$, which has length $\phi(s) - d(x,\mathcal{Z})$.

Let $R := \sup\{t \mid \phi(t) \leq (2E+1)t\}$, which exists since ϕ is superlinear. Notice that $d(\pi_Z(x), \beta) \leq s$, so without loss of generality we may assume $s > R$, since otherwise we already have a bound for $d(\pi_Z(x), \beta)$ that is independent of r.

For β to be E–efficient we need $d(x,\mathcal{Z}) \geq \frac{\phi(s) - Es}{E+1}$. Since $\phi(s)/s > 2E + 1$ this implies $d(x,z) \geq s$ so $d(z,\beta) = s = d(z,y)$. Therefore:

$$\tau(r;E) = \max\{R, \sup\{s \mid \phi(s) - Es \leq r(E+1)\}\} \asymp \sup\{s \mid \phi(s) - Es \leq r(E+1)\}$$

Superlinearity of ϕ implies $\tau(r;E)$ is sublinear.

If we assume $\phi(r)/r$ is invertible then the above computations give us:

- $\kappa(r) \asymp \phi^{-1}(r)$
- $\mu(r) \asymp \phi(r)/r$
- $\delta(r) \asymp \phi(r)$
- $\rho(E) \asymp (\frac{\phi}{\mathrm{Id}})^{-1}(E)$
- $\chi(E) \asymp \phi((\frac{\phi}{\mathrm{Id}})^{-1}(E)) = E \cdot (\frac{\phi}{\mathrm{Id}})^{-1}(E)$
- $\tau(r;E) \asymp \sup\{s \mid \phi(s) - Es \leq r(E+1)\}$

For these examples the bounds given by items (3), (4), and (5) of Corollary 3.5 are sharp.

Example 8.2 (Warped L^1 half-plane). Consider the upper half Cartesian plane. We define vertical paths to have their usual length, and rescale horizontal paths at height b by a factor of $\phi(b)$, where ϕ is a nondecreasing continuous function with $\phi(0) = 1$. Let $\mathcal{X}$ be this space, and define the distance between two points to be the infinum of lengths of piecewise vertical or horizontal paths connecting them. The conditions on ϕ imply this is a geodesic metric space. Moreover, the geodesic between two points consists of Euclidean line segments going first down, then horizontally, then up (where some of these segments may be trivial).

Let $\mathcal{Z} := \{(a,0) \mid a \in \mathbb{R}\}$. It is a convex geodesic, and $\pi_{\mathcal{Z}}((a,b)) = (a,0)$.

If ϕ is bounded then $\mathcal{X}$ is quasiisometric to a Euclidean half-plane, so $\mathcal{Z}$ is not Morse.

If $\phi(b) = 1$ for all $b \geq 0$ then $\mathcal{X}$ is just the upper half-plane with the L^1 metric, and it is easy to check that $\mathcal{Z}$ is not geodesically trim by looking at large squares.

If $\phi(0) < \phi(b)$ for all $b > 0$ then $\mathcal{Z}$ is strongly geodesically trim. In particular, if ϕ is also bounded then $\mathcal{Z}$ is strongly geodesically trim but not Morse.

Now assume that ϕ is unbounded. We estimate the various gauges, showing that $\mathcal{Z}$ is Morse. By symmetry it suffices to consider $x := (0, r)$ and $y := (a, r)$ as points on $\mathcal{S}_r(\mathcal{Z})$. Note $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) = a$.

Frilliness gauge μ. Suppose $a > 2r > 0$. We have $d_r(x, y) = a\phi(r)$, while $a \leq d(x, y) \leq 2r + a$, so:

$$\mu(r) = \inf_{a>2r} \frac{d_r((0, r), (a, r))}{d((0, r), (a, r))} \asymp \phi(r)$$

Since ϕ is nondecreasing and unbounded, it limits to ∞, so μ does too.

Divergence gauge δ. If $d(x, y) \geq 3r$ then $2r + a \geq d(x, y) \geq 3r$ implies $a \geq r$ and $d_r(x, y) = a\phi(r) \geq r\phi(r)$, so $\delta(r) \geq r\phi(r)$. Conversely, for $a = 5r$ we have $d(x, y) \geq 5r - 2r = 3r$ and $\delta(r) \leq d_r(x, y) = 5r\phi(r)$, so $\delta(r) \asymp r\phi(r)$.

Contraction gauge κ. Suppose $v \in \bar{\mathcal{N}}_{d(x,\mathcal{Z})}(x)$ maximizes $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(v)$. A geodesic from x to v cannot have an upward segment since in that case we could have instead extended the preceding horizontal segment to get a point still in the ball and having projection to $\mathcal{Z}$ farther from $\pi_{\mathcal{Z}}(x)$ than $\pi_{\mathcal{Z}}(v)$. Thus, if $v = (a, b)$ then $0 \leq b \leq r$, and $r = d(x, v) = r - b + a\phi(b)$. This implies $b/\phi(b) = a = \operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(v)$, so $\kappa(r) = \sup_{0\leq b\leq r} b/\phi(b)$. To see that κ is sublinear, suppose $\limsup_{r\to\infty} \kappa(r)/r > 0$ and produce a contradiction to the fact that ϕ is unbounded.

Recurrence gauge ρ. For every $E > 1$, every r such that $\phi(r) < E$, and every sufficiently long a, the piecewise Euclidean-geodesic path γ connecting $(0, 0)$, $(0, r)$, (a, r), and $(a, 0)$ is locally E–efficient and $d(\gamma, \mathcal{Z}'_\gamma) = r$, so $\rho(r) \geq \sup\{r \mid \phi(r) < E\}$.

For the converse, consider the points $z = (0, 0)$ and $z' = (a, 0)$ of $\mathcal{Z}$, with $r = a/3 = d(z, z')/3$ and $\mathcal{Z}' := \mathcal{Z} - (\mathcal{N}_r(z) \cup \mathcal{N}_r(z'))$. Consider the segment γ of $\mathcal{S}_r(Z')$ from z to z': it consists of a concatenation of Euclidean line segments connecting the points $(0, 0)$, (r, r), $(2r, r)$ and $(3r, 0)$. Every path from z to z' in $\mathcal{N}^c_r(\mathcal{Z}')$ projects to this segment, and the projection is length nonincreasing, so γ is the most efficient connection between z and z' avoiding $\mathcal{N}_r(\mathcal{Z}')$. Estimate:

$$E \geq \operatorname{len}(\gamma)/3r \geq (4r + r\phi(r))/3r = (4 + \phi(r))/3$$

Thus, $\phi(r) \leq 3E - 4$, which tells us:

$$\sup\{r \mid \phi(r) < E\} \leq \rho(E) \leq \sup\{r \mid \phi(r) \leq 3E - 4\}$$

Since ϕ is unbounded, this supremum exists.

If we assume ϕ is twice differentiable with $\phi''(r) \leq 0$ then the above computations give us:

- $\kappa(r) = r/\phi(r)$
- $\mu(r) \asymp \phi(r)$
- $\delta(r) = r\phi(r)$
- $\rho(E) \asymp \phi^{-1}(E)$

For these examples the bounds given by items (1), (2), (3), and (4) of Corollary 3.5 are sharp.

8.2. Robustness.

8.2.1. *Frilliness and divergence.*

Proposition 8.3. *Let $\mu(r; M, N)$ be the frilliness gauges of $\mathcal{Z} \subset \mathcal{X}$. For all $0 < M \leq M' \leq 1$ and $2 \leq N \leq N'$ we have $\mu(r; M, N) \asymp \mu(r; M', N')$.*

Proof. It is immediate from the definitions that $\mu(r; M, N) \leq \mu(r; M', N)$ and $\mu(r; M, N) \leq \mu(r; M, N')$.

Suppose $M < M'$, and, given $r > 0$, define $R := \frac{M'}{M} r$. Suppose that $\mu(R; M, N)$ is finite, so there exist $x, y \in \mathcal{S}_R(\mathcal{Z})$ with $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) > NR$ and $\frac{d_{MR}(x,y)}{d(x,y)} < \mu(R; M, N) + 1$. Let $x'' \in \pi_{\mathcal{Z}}(x)$ and $y'' \in \pi_{\mathcal{Z}}(y)$ be points such that $d(x'', y'') > NR$. Choose a geodesic from x to x'' and let x' be the unique point on it at distance r from $\mathcal{Z}$. Define y' similarly with respect to y. See Figure 6. Observe:

$$\frac{NM'}{M} r = NR < d(x'', y'') \leq \operatorname{diam} \pi_{\mathcal{Z}}(x') \cup \pi_{\mathcal{Z}}(y') \leq 2r + d(x', y') = \left(2 + \frac{d(x', y')}{r}\right) r$$

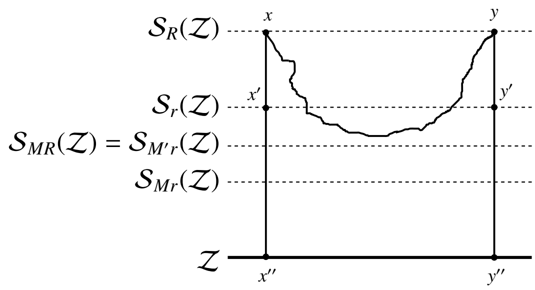

FIGURE 6. Setup for Proposition 8.3.

This implies $r/d(x',y') < M/(NM' - 2M)$, which further implies:

$$\frac{R-r}{d(x',y')} < \frac{M'-M}{NM'-2M} \leq \frac{1}{2} \tag{1}$$

Now, since $\operatorname{diam} \pi_{\mathcal{Z}}(x') \cup \pi_{\mathcal{Z}}(y') > NR > Nr$, we have that x' and y' participate in the infinum defining $\mu(r; M', N)$, so we can estimate:

$$\begin{aligned}
\mu(r; M', N) &\leq \frac{d_{M'r}(x',y')}{d(x',y')} \\
&\leq \frac{d(x',x) + d_{M'r}(x,y) + d(y,y')}{d(x',y')} \\
&= \frac{d(x',x) + d_{MR}(x,y) + d(y,y')}{d(x',y')} \\
&\leq \frac{2(R-r) + d(x,y)(1+\mu(R;M,N))}{d(x',y')} \\
&\leq \frac{2(R-r) + (d(x',y') + 2(R-r))\,(1+\mu(R;M,N))}{d(x',y')} \\
&= \left(1 + \frac{4(R-r)}{d(x',y')}\right) + \left(1 + \frac{2(R-r)}{d(x',y')}\right)\mu(R;M,N) \\
&\leq 3 + 2\mu(R;M,N)
\end{aligned}$$

This argument assumed $\mu(R; M, N)$ was finite, but the same inequality is true if $\mu(R; M, N) = \infty$, so:

$$\mu(r; M', N) \prec \mu(R; M, N) = \mu\left(\frac{M'}{M}r; M, N\right) \asymp \mu(r; M, N)$$

The proof for $N' > N$ uses a similar argument with $R := \frac{N'}{N}r$ to show $\mu(r; M, N') \prec \mu(R; M, N) \asymp \mu(r; M, N)$. This time, the analogue of (1) is:

$$\frac{R-r}{d(x',y')} < \frac{N'-N}{N(N'-2)} \leq \frac{1}{2} \tag{2}$$ □

PROPOSITION 8.4. *Let $\delta(r; M, N)$ be the lower divergence functions of $\mathcal{X}$ relative to $\mathcal{Z}$. For all $0 < M \leq M' \leq 1$ and $2 < N \leq N'$ we have $\delta(r; M, N) \asymp \delta(r; M', N')$.*

PROOF. It is immediate from the definitions that $\delta(r; M, N) \leq \delta(r; M', N')$.
Suppose $N' > N$, and given r let $R := \frac{N'-2}{N-2}r$. Arguing as in Proposition 8.3, we see:

$$\begin{aligned}
\delta(r; M, N') &\leq 2(R-r) + \delta(R; M, N) + 1 \\
&= 2\left(\frac{N'-N}{N-2}\right)r + \delta\left(\frac{N'-2}{N-2}r; M, N\right) + 1 \\
&\asymp \delta(r; M, N)
\end{aligned}$$

Note that the linear term disappears in the final equivalence because divergence functions grow at least linearly by construction.

Suppose $M' > M$. Let $N'' > N$ be the solution of $\frac{N''-2}{N-2} = \frac{M'}{M}$, and, given r, let $R := \frac{M'}{M}r$. Again, arguing as in Proposition 8.3 we get:

$$\delta(r; M', N) \leq \delta(r; M', N'') \leq 2(R - r) + \delta(R; M, N) + 1 \asymp \delta(r; M, N) \qquad \square$$

8.2.2. *The contraction properties.* We state some robustness results for contraction. One motivation for writing these results in terms of contraction instead of quasigeodesic quasiconvexity is that then they can also be used for strong contraction. These results show that using almost closest point projection instead of closest point projection does not change the equivalence class of the contraction gauge, nor does exchanging $\mathcal{Z}$ with a Hausdorff equivalent set.

Lemma 8.5 ([Section 6, Paper A]).

- *If $\pi_{\mathcal{Z}}$ is (κ_1, κ_2)–contracting for $\kappa_1(r) = \kappa_1'(r) - C$, with $\kappa_1'(r) \leqslant r$ and $C \geqslant 0$, then $\pi_{\mathcal{Z}}$ is (κ_1', κ_2')–contracting for some $\kappa_2' \asymp \kappa_2$.*
- *Suppose $\epsilon_0, \epsilon_1 \geq 0$ are constants such that the empty set is neither in the image of $\pi_{\mathcal{Z}}^{\epsilon_0}$ nor in the image of $\pi_{\mathcal{Z}}^{\epsilon_1}$. If $\pi_{\mathcal{Z}}^{\epsilon_0}$ is (κ_1, κ_2)–contracting then there exist κ_1' and κ_2' such that $\pi_{\mathcal{Z}}^{\epsilon_1}$ is (κ_1', κ_2')–contracting. If $\epsilon_1 \leqslant \epsilon_0$ or if $\kappa_1(r) := r$ then we can take $\kappa_1' = \kappa_1$ and $\kappa_2' \asymp \kappa_2$.*
- *Let $\mathcal{Y}$ and $\mathcal{Z}$ be subspaces of a geodesic metric space $\mathcal{X}$ at bounded Hausdorff distance from one another. Suppose that $\pi_{\mathcal{Y}}^{\epsilon}$ is (κ_1, κ_2)–contracting. Then $\pi_{\mathcal{Z}}^{\epsilon}$ is (r, κ_2')–contracting for some κ_2'. If $\kappa_1(r) = r$ then $\kappa_2' \asymp \kappa_2$.*

8.3. Proof of Theorem 3.2.

Proposition 8.6 (⑤). *Contracting implies quasigeodesically trim.*

Proof. Suppose $\mathcal{Z}$ is (κ_1, κ_2)–contracting. By Remark 3.1, we may assume that κ_2/κ_1 is a nonincreasing continuous function. There is also no loss in restricting our attention to continuous quasigeodesics, since any geodesic can be tamed to get a continuous quasigeodesic at bounded Hausdorff distance and with controlled quasigeodesic constants, as in Lemma 2.1. With these simplifications, if κ_2 is identically zero then a path with one endpoint x off $\mathcal{Z}$ must pass through the unique point $\pi_{\mathcal{Z}}(x)$, so we can take the trimness gauge to be 0. Therefore, we can assume κ_2 is not identically zero, which, since κ_2 is nondecreasing, means that for all sufficiently large r we have $\kappa_2(r) > 0$.

Fix $L \geq 1$ and $A \geq 0$. Throughout we assume r is large enough that $\kappa_1(r) \geq 2A$, which is true for all sufficiently large r. For $N > 0$ define:

$$C_N := \sup\{s \mid 2(N + L^2)\kappa_2(s) > \kappa_1(s)\}$$

This number exists since $\lim_{r\to\infty} \kappa_2(r)/\kappa_1(r) = 0$. Suppose that $\beta\colon [0, T] \to \mathcal{X}$ is a continuous (L, A)–quasigeodesic with $d(\beta, \mathcal{Z}) = d(\beta(T), \mathcal{Z}) = C_N$. The same argument as [Theorem 4.2, Paper H] yields:

$$\operatorname{diam} \pi_{\mathcal{Z}}(\beta(0)) \cup \pi_{\mathcal{Z}}(\beta(T)) \leq \frac{1}{N}\left(L^2 d(\beta(0), \mathcal{Z}) + (L^2 + 1)C_N + L^2 A\right) \tag{3}$$

Now suppose $N > 1$ and $\beta\colon [0, T] \to \mathcal{X}$ is a continuous (L, A)–quasigeodesic with $\beta(T) \in \mathcal{Z}$ and $d(\beta(0), \mathcal{Z}) > C_N$. Let s be the first time such that $d(\beta(s), \mathcal{Z}) = C_N$, and let t be the first time such that $d(\beta(t), \mathcal{Z}) = C_1$. Then by applying (3) twice, for N and 1, we have:

$$\begin{aligned}
d(\beta, \pi_{\mathcal{Z}}(\beta(0))) &\leq \operatorname{diam} \pi_{\mathcal{Z}}(\beta(0)) \cup \pi_{\mathcal{Z}}(\beta(s)) + \operatorname{diam} \pi_{\mathcal{Z}}(\beta(s) \cup \pi_{\mathcal{Z}}(\beta(t)) + C_1 \\
&\leq \frac{1}{N}\left(L^2 d(\beta(0), \mathcal{Z}) + (L^2 + 1)C_N + L^2 A\right) \\
&\qquad + \frac{1}{1}\left(L^2 C_N + (L^2 + 1)C_1 + L^2 A\right) + C_1 \\
&= \frac{L^2}{N} d(\beta(0), \mathcal{Z}) + \left[L^2 C_N + \frac{(L^2 + 1)C_N + L^2 A}{N} + (L^2 + 2)C_1 + L^2 A\right] \qquad (4)
\end{aligned}$$

Let $\phi(N)$ be the bracketed expression in (4).

Since κ_2/κ_1 is continuous, nonincreasing, and limits to zero, we can define:

$$\psi(r) := \left(\frac{\mathrm{Id}}{\kappa_2/\kappa_1}\right)^{-1}(r)$$

Since ψ is the inverse of a strictly increasing, continuous, superlinear function it is strictly increasing, continuous, unbounded, and sublinear.

Assume r is sufficiently large that $\kappa_1(\psi(r)) > 2L^2\kappa_2(\psi(r))$, which is possible since κ_1 grows faster than κ_2 and ψ is increasing and unbounded. Define:

$$M(r) := \frac{2L^2}{\frac{\kappa_1(\psi(r))}{2\kappa_2(\psi(r))} - L^2}$$

Observe:

$$\begin{aligned} rM(r) &= \frac{2L^2 r}{\frac{\kappa_1(\psi(r))}{2\kappa_2(\psi(r))} - L^2} \\ &= \frac{2L^2\frac{\mathrm{Id}}{\kappa_2/\kappa_1}(\psi(r))}{\frac{\kappa_1(\psi(r))}{2\kappa_2(\psi(r))} - L^2} \\ &= \frac{4L^2\psi(r)}{1 - 2L^2\frac{\kappa_2}{\kappa_1}(\psi(r))} \\ &\asymp \psi(r) \end{aligned}$$

Define $N(r) := 2L^2/M(r)$. Notice that $\phi(N(r)) \asymp C_{N(r)}$, and by backfilling the definitions of $N(r)$, $M(r)$, and C, we have that $C_{N(r)} = \sup\{s \mid \frac{\kappa_2(s)}{\kappa_1(s)} > \frac{\kappa_2(\psi(r))}{\kappa_1(\psi(r))}\}$, which, since κ_2/κ_1 is nonincreasing, is at most $\psi(r)$.

For r sufficiently large, as above, suppose that $\beta\colon [0,T] \to \mathcal{X}$ is a continuous (L,A)–quasigeodesic with $\beta(T) \in \mathcal{Z}$ and $d(\beta(0),\mathcal{Z}) \leq r$. If $d(\beta(0),\mathcal{Z}) \leq C_{N(r)}$, then:

$$d(\beta, \pi_{\mathcal{Z}}(\beta(0))) \leq d(\beta(0),\mathcal{Z}) \leq C_{N(r)} \leq \psi(r)$$

On the other hand, if $d(\beta(0),\mathcal{Z}) > C_{N(r)}$ then we can apply (4) to get:

$$\begin{aligned} d(\beta, \pi_{\mathcal{Z}}(\beta(0))) &\leq \frac{L^2}{N(r)} d(\beta(0),\mathcal{Z}) + \phi(N(r)) \\ &= \frac{M(r)}{2} d(\beta(0),\mathcal{Z}) + \phi(N(r)) \\ &\leq \frac{rM(r)}{2} + \phi(N(r)) \\ &\prec \psi(r) \end{aligned}$$

We have shown that there exists a sublinear function $\psi(r)$ such that for each fixed L and A we have $\tau(r; L, A) \prec \psi(r)$, so $\tau(r; L, A)$ is sublinear. This means $\mathcal{Z}$ is quasigeodesically trim. □

Proposition 8.7. *Contracting implies frilly.*

Proof. Suppose $\mathcal{Z}$ is (κ_1,κ_2)–contracting with κ_2/κ_1 nonincreasing as in Remark 3.1, and let $\mu(r)$ be the frilliness gauge. We must show $\lim_{r\to\infty}\mu(r) = \infty$. Let $r > 0$. If $\mu(r) = \infty$ we are done, so suppose not, which means there exist $x, y \in \mathcal{S}_r(\mathcal{Z})$ with $\operatorname{diam}\pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) > 2r$ and a path $\gamma \subset \mathcal{N}_r^c(\mathcal{Z})$ connecting x to y satisfying $\operatorname{len}(\gamma) < 1 + d_r(x,y)$. Let $x_0 := x$. Given $x_i \neq y$, define x_{i+1} to be the last point of γ at distance at most $\kappa_1(d(x_i,\mathcal{Z}))$ from x_i. Suppose k is such that $y = x_{k+1}$. Such a k exists because the distance between successive x_i's is at least $\kappa_1(r)$. Then $\operatorname{diam}\pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) \leq \sum_{i=0}^{k}\kappa_2(d(x_i,\mathcal{Z}))$. Note that if $\kappa_2(r) = 0$ then $\operatorname{diam}\pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) = 0$, which is a contradiction, so $\kappa_2(r) > 0$. Now estimate:

$$\begin{aligned}
d_r(x,y) + 1 > \operatorname{len}(\gamma) &\geq \sum_{i=0}^{k} d(x_i, x_{i+1}) \\
&= d(x_k, y) + \sum_{i=0}^{k-1} \kappa_1(d(x_i, \mathcal{Z})) \\
&= d(x_k, y) - d(x_k, \mathcal{Z}) + d(x_k, \mathcal{Z}) - \kappa_1(d(x_k, \mathcal{Z})) \\
&\qquad + \kappa_1(d(x_k, \mathcal{Z})) + \sum_{i=0}^{k-1} \kappa_1(d(x_i, \mathcal{Z})) \\
&> -r + \sum_{i=0}^{k} \frac{\kappa_1(d(x_i, \mathcal{Z}))}{\kappa_2(d(x_i, \mathcal{Z}))} \kappa_2(d(x_i, \mathcal{Z})) \\
&\geq -r + \sum_{i=0}^{k} \frac{\kappa_1(r)}{\kappa_2(r)} \kappa_2(d(x_i, \mathcal{Z})) \\
&\geq -r + \frac{\kappa_1(r)}{\kappa_2(r)} \operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y)
\end{aligned}$$

We then have:

$$\begin{aligned}
\frac{d_r(x,y)}{d(x,y)} &> \frac{\frac{\kappa_1(r)}{\kappa_2(r)} \operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) - r - 1}{2r + \operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y)} \\
&> \frac{\frac{\kappa_1(r)}{\kappa_2(r)}(2r) - r - 1}{4r} \\
&= \frac{1}{2} \cdot \frac{\kappa_1(r)}{\kappa_2(r)} - \frac{r+1}{4r}
\end{aligned}$$

Since x and y were arbitrary points participating in the infinum defining $\mu(r)$, this gives:

$$\mu(r) \geq \frac{1}{2} \cdot \frac{\kappa_1(r)}{\kappa_2(r)} - \frac{r+1}{4r} \asymp \frac{\kappa_1}{\kappa_2}(r)$$

Since κ_1 grows faster than κ_2 by hypothesis, $\mu(r)$ goes to ∞, as desired. □

Proposition 8.8. *Frilly implies divergent.*

Proof. Let μ be the frilliness gauge of $\mathcal{Z}$, and let δ be the divergence gauge. We must show $\lim_{r\to\infty} \delta(r)/r = \infty$, which will follow from showing $\delta(r) \succ r\mu(r)$. Let $r > 0$. Suppose there exist $x, y \in \mathcal{S}_r(\mathcal{Z})$ such that $d(x, y) \geq 5r$. Then $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) \geq 5r - 2r > 2r$. Since $\mathcal{Z}$ is frilly, $d_r(x, y) \geq \mu(r; 1, 2)d(x, y) \geq 5r\mu(r; 1, 2)$. Thus:

$$\delta(r; 1, 5) \geq 5r\mu(r; 1, 2) \tag{5}$$

If there are no such x and y then $\delta(r; 1, 5) = \infty$ and (5) still holds.

Combining (5) with Proposition 8.4 and Proposition 8.3, this gives:

$$\delta(r) = \delta(r; 1, 3) \asymp \delta(r; 1, 5) \geq 5r\mu(r; 1, 2) \asymp r\mu(r; 1, 2) = r\mu(r)$$ □

Proposition 8.9. *Frilly implies quasigeodesically quasiconvex.*

This is easy. For the purpose of exhibiting the effectiveness we give a proof for the locally efficient case and leave the general case to the reader.

Proof. Let μ be the frilliness gauge of $\mathcal{Z}$. Fix an $E \geq 1$. Since $\lim_{r\to\infty} \mu(r) = \infty$, there exists $R := \sup\{r \mid \mu(r) \leq E\}$ such that for all $r > R$ we have $\mu(r) > E$. Suppose γ is a locally E–efficient path with endpoints on $\mathcal{Z}$ such that $\max_{w\in\gamma} d(w, \mathcal{Z}) = D > R$. For any $R < r < D$, let α be a subsegment of γ such that $\alpha \subset \mathcal{N}^c_r(\mathcal{Z})$ and such that α has endpoints $x, y \in \mathcal{S}_r(\mathcal{Z})$.

If $\operatorname{diam}\pi_{\mathcal{Z}}(x)\cup\pi_{\mathcal{Z}}(y) > 2r$ then x and y participate in the infinum defining μ and we get a contradiction:

$$E < \mu(r) \le \frac{d_r(x,y)}{d(x,y)} \le \frac{\operatorname{len}(\alpha)}{d(x,y)} \le E$$

Thus, $\operatorname{diam}\pi_{\mathcal{Z}}(x)\cup\pi_{\mathcal{Z}}(y) \le 2r$, which implies $d(x,y) \le 4r$. Then:

$$\max_{w\in\alpha} d(w,\mathcal{Z}) \le r + \operatorname{len}(\alpha)/2 \le r + Ed(x,y)/2 \le (1+2E)r$$

Since this was true for every $r > R$ we have that $\mathcal{Z}$ is χ–locally efficiently quasiconvex for:

$$\chi(E) \le (1+2E)\cdot\sup\{r \mid \mu(r) \le E\}$$ □

Proposition 8.10. *Quasigeodesically trim implies quasigeodesically quasiconvex.*

Proof. Suppose $\mathcal{Z}$ is τ–quasigeodesically trim. Fix $L \ge 1$ and $A \ge 0$, and let $\tau(r)$ denote $\tau(r; L, A)$. By definition of trim, τ is sublinear, so $R := \sup\{r \mid \tau(r) \ge r\}$ exists.

Suppose $\gamma\colon [0,T] \to \mathcal{X}$ is an (L,A)–quasigeodesic segment with endpoints on $\mathcal{Z}$. If $\sup_{w\in\gamma} d(w,\mathcal{Z}) \le R$ then we are done, so suppose there exists $b \in (0,T)$ such that $d(\gamma(b),\mathcal{Z}) > \tau(d(\gamma(b),\mathcal{Z}))$. Choose a point $z \in \pi_{\mathcal{Z}}(\gamma(b))$. Since $\mathcal{Z}$ is trim, there exist $a \in [0,b)$ and $c \in (b,T]$ such that $d(\gamma(a),z) \le \tau(d(\gamma(b),\mathcal{Z}))$ and $d(z,\gamma(c)) \le \tau(d(\gamma(b),\mathcal{Z}))$. Thus, $c - a \le 2L\tau(d(\gamma(b),\mathcal{Z})) + LA$. On the other hand, we have:

$$d(\gamma(a),\mathcal{Z}) \le d(\gamma(a),z) \le \tau(d(\gamma(b),\mathcal{Z}))$$

This implies:

$$d(\gamma(a),\gamma(b)) \ge d(\gamma(b),\mathcal{Z}) - d(\gamma(a),\mathcal{Z}) \ge d(\gamma(b),\mathcal{Z}) - \tau(d(\gamma(b),\mathcal{Z}))$$

We have a similar bound for $d(\gamma(c),\gamma(b))$, and using the quasigeodesic property we get corresponding bounds on $c - b$ and $b - a$. Upon adding these we get:

$$c - a \ge \frac{2}{L}\left(d(\gamma(b),\mathcal{Z}) - \tau(d(\gamma(b),\mathcal{Z})) - A\right)$$

Combining this with the upper bound on $c - a$, we have:

$$R < d(\gamma(b),\mathcal{Z}) \le (L^2+1)(\tau(d(\gamma(b),\mathcal{Z})) + A)$$

Thus, $\chi(L,A) \le \sup\{r \mid r \le (L^2+1)(\tau(r;L,A)+A)\}$, which exists since τ is sublinear. □

Corollary 8.11. *Geodesically trim implies geodesically quasiconvex*

In general geodesically trim does not imply quasigeodesic quasiconvexity, nor does strongly geodesically trim (recall Example 8.2). However, this implication does hold in CAT(0) spaces:

Proposition 8.12 (②). *If $\mathcal{X}$ is CAT(0) and $\mathcal{Z} \subset \mathcal{X}$ is τ–geodesically trim then $\mathcal{Z}$ is $(r/2,\tau')$–contracting for a function $\tau' \prec \tau$.*

Proof. Suppose x and y are points with $2d(x,y) \le d(x,\mathcal{Z})$. Choose arbitrary $x' \in \pi_{\mathcal{Z}}(x)$ and $y' \in \pi_{\mathcal{Z}}(y)$. Since $2d(x,y) \le d(x,\mathcal{Z}) \le d(x,y')$, the angle at y' of the Euclidean comparison triangle to $xy'y$ is at most 30°.

Let $T := \tau(d(x,\mathcal{Z}))$. Since $\mathcal{Z}$ is τ–trim there exists a point z on the geodesic from x to y' such that $d(z,x') \le T$. For any $D \le \min\{2d(z,y')/\sqrt{3}, d(y,y')\}$ there exists a point u on the geodesic from x to y' such that $2d(u,y') = D\sqrt{3}$ and a point v on the geodesic from y to y' such that $d(v,y') = D$. The Euclidean distance between the comparison points for u and v in triangle $xy'y$ is at most $D/2$, since the angle at y' is at most 30°, so $d(u,v) \le D/2$.

Now let D achieve its upper bound, which happens when either $u = z$ or $v = y$.

First, consider the case that $D = d(y,y')$ and $v = y$. By Proposition 8.10, $\mathcal{Z}$ is Q–geodesically quasiconvex for $Q = \sup\{r \mid r \le 2\tau(r)\}$, which exists since τ is sublinear. Together with the CAT(0) hypothesis and the fact that $d(z,x') \le T$, this gives us $[z,y'] \subset \bar{\mathcal{N}}_T([x',y']) \subset \bar{\mathcal{N}}_{T+Q}(\mathcal{Z})$. Thus:

$$d(y,\mathcal{Z}) \le d(y,u) + d(u,\mathcal{Z}) \le D/2 + T + Q = d(y,\mathcal{Z})/2 + T + Q$$

The fact that $d(x,y) \le d(x,\mathcal{Z})/2$ then gives:

$$d(x,\mathcal{Z})/2 \le d(x,\mathcal{Z}) - d(x,y) \le d(y,\mathcal{Z}) \le 2T + 2Q$$

However, this means that this case can only occur if $d(x, \mathcal{Z}) \le R := \sup\{r \mid r \le 4\tau(r) + 4Q\}$. Since $d(x', y') \le d(x', x) + d(x, y) + d(y, y') \le 3d(x, \mathcal{Z})$, there is a uniform bound $d(x', y') \le 3R$ in this case.

Now consider the case that $D = 2d(z, y')/\sqrt{3}$ and $u = z$. We have:

$$d(v, y') = d(v, \mathcal{Z}) \le d(v, z) + d(z, \mathcal{Z}) \le d(v, z) + T \le D/2 + T \le d(v, y')/2 + T$$

This implies:

$$2(d(x', y') - T)/\sqrt{3} \le 2d(z, y')/\sqrt{3} = D = d(v, y') \le 2T$$

This gives $d(x', y') \le (1 + \sqrt{3})T = (1 + \sqrt{3})\tau(d(x, \mathcal{Z}))$. Combining the two cases, $\mathcal{Z}$ is $(r/2, \tau')$–contracting for $\tau'(r) \le (1 + \sqrt{3})\tau(r) + 3R \asymp \tau(r)$. □

In the special case that τ is bounded, we recover a result of Sultan, which in our terminology says:

Corollary 8.13 (cf. Sultan [75, Theorem 3.4]). *If X is CAT(0) and $\mathcal{Z}$ is strongly geodesically trim then $\mathcal{Z}$ is semistrongly contracting.*

Proposition 8.14. *Suppose there exists $N > 2$ such that $\mathcal{Z}$ is recurrent. Then $\mathcal{Z}$ is frilly.*

Proof. Let ρ be the recurrence gauge of $\mathcal{Z}$ with respect to parameter N, and let μ be the frilliness gauge. We must show $\lim_{r\to\infty} \mu(r) = \infty$. By Proposition 8.3, it suffices to show that $\lim_{r\to\infty} \mu(r; 1, 2N) = \infty$.

If $\mu(r; 1, 2N)$ attains the value ∞ we are done, so we may suppose that for each $r \ge 1$, there exist points $x, y \in \mathcal{S}_r(\mathcal{Z})$ satisfying $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) > 2Nr$ and $d_r(x, y) < \infty$, so there is a path $\gamma \subset \mathcal{N}^c_r(\mathcal{Z})$ from x to y with $\operatorname{len}(\gamma) < 1 + d_r(x, y)$. Pick $x' \in \pi_{\mathcal{Z}}(x)$ and $y' \in \pi_{\mathcal{Z}}(y)$ such that $d(x', y') > 2Nr$. Define α to be the concatenation of a geodesic from x' to x, then γ, then a geodesic from y to y'. We estimate the efficiency of α:

$$\begin{aligned}
\frac{\operatorname{len}(\alpha)}{d(x', y')} &\le \frac{d_r(x, y) + 1 + 2r}{d(x', y')} \\
&\le \frac{1 + 2r}{d(x', y')} + \frac{2r + d(x', y')}{d(x, y)} \cdot \frac{d_r(x, y)}{d(x', y')} \\
&\le \frac{1 + 2r}{d(x', y')} + \frac{2r + d(x', y')}{d(x', y')} \cdot \frac{d_r(x, y)}{d(x, y)} \\
&< 1 + (1 + 1) \cdot \frac{d_r(x, y)}{d(x, y)}
\end{aligned}$$

Let $\mathcal{Z}' := \mathcal{Z} - \mathcal{N}_{d(x',y')/N}\{x', y'\}$. Suppose $\mathcal{Z}' \ne \emptyset$. The choice of $2N$ guarantees that $d(u, \mathcal{Z}') \ge r$ for all u on a geodesic from x to x' or from y to y', and $d(v, \mathcal{Z}') \ge d(v, \mathcal{Z}) \ge r$ for all $v \in \gamma$, so $d(\alpha, \mathcal{Z}') \ge r$. We get an upper bound on $d(\alpha, \mathcal{Z}')$ from recurrence: there exists a point $w \in \alpha$ with $d(w, \mathcal{Z}') \le 1 + \rho\left(1 + 2\frac{d_r(x,y)}{d(x,y)}; N\right)$, so:

$$r \le d(\alpha, \mathcal{Z}') \le 1 + \rho\left(1 + 2\frac{d_r(x, y)}{d(x, y)}; N\right) \tag{6}$$

The bound (6) also holds in the case that $\mathcal{Z}' = \emptyset$, since then:

$$\rho\left(1 + 2\frac{d_r(x, y)}{d(x, y)}; N\right) \ge \bar{d}(\alpha, \mathcal{Z}') := d(x', y')/N > 2r > r$$

Define

$$\mu'(r) = \inf\{s \mid r \le 1 + \lim_{t\to s^+} \rho(1 + 2t; N)\}$$

Since ρ is nondecreasing, $\lim_{r\to\infty} \mu'(r) = \infty$. Since x and y were an arbitrary pair participating in the infinum defining $\mu(r; 1, 2N)$, (6) implies $\mu(r; 1, 2N) \ge \mu'(r)$. □

The next proposition is Tran's, with some trivial adjustments to make the result explicitly effective. For simplicity we will only consider the case of locally efficient quasiconvexity.

Proposition 8.15 ([77, Proposition 3.1]). *Divergent implies locally efficiently quasiconvex.*

Proof. Let $\delta(r)$ be the divergence gauge of $\mathcal{Z}$. Fix $E \geq 1$. Since δ is completely superlinear, there exists $R \geq 1$ such that $\delta(r) > (2E^2 + 6E + 1)r$ for all $r > R$.

Suppose there exists a locally E–efficient path γ with endpoints on $\mathcal{Z}$ and a point $w \in \gamma$ with $A := d(w, \mathcal{Z}) > 2(E + 1)R$. Let $r := \frac{A}{2(E+1)} > R$. Let α be a subsegment of γ containing w and contained in $\mathcal{N}_r^c(\mathcal{Z})$ with endpoints $x, y \in \mathcal{S}_r(\mathcal{Z})$.

First suppose $\mathrm{len}(\alpha) \leq 2E(E + 2)r$. Since α goes from $\mathcal{S}_r(\mathcal{Z})$ to $\mathcal{S}_A(\mathcal{Z})$ and back, $\mathrm{len}(\alpha) \geq 2(A - r) = 2r(2E + 1)$. Since α is E–efficient, $d(x, y) \geq \mathrm{len}(\alpha)/E > 4r$, which implies $\delta(r) \leq d_r(x, y) \leq \mathrm{len}(\alpha) \leq 2E(E + 2)r < (2E^2 + 6E + 1)r$. Since $r > R$, this is a contradiction, so we must have $\mathrm{len}(\alpha) > 2E(E + 2)r$.

In this case, let $u \in \alpha$ be the point such that the length of the segment of α between x and u is exactly $2E(E + 2)r$. Pick a geodesic from u to a point of $\pi_{\mathcal{Z}}(u)$ and let v be the point of intersection of the geodesic and $\mathcal{S}_r(\mathcal{Z})$. We have:

$$d(x, v) \geq d(x, u) - d(u, v) \geq \mathrm{len}(\alpha|_{x,u})/E - (A - r) = 2(E + 2)r - (2E + 1)r = 3r$$

This implies:

$$\delta(r) \leq d_r(x, v) \leq \mathrm{len}(\alpha|_{x,u}) + d(u, v) \leq 2E(E + 2)r + A - r = (2E^2 + 6E + 1)r$$

Again, since $r > R$ this contradicts the definition of R, so there is no such γ. Thus:

$$\chi(E) \leq 2(E + 1) \cdot \sup\{r \mid \delta(r) \leq (2E^2 + 6E + 1)r\}$$ □

The final piece of the proof of Theorem 3.2 is:

Proposition 8.16. *Let $\phi\colon \mathcal{X} \to \mathcal{Y}$ be a quasiisometric embedding between geodesic metric spaces. The set $\phi(\mathcal{X})$ is stable in $\mathcal{Y}$ if and only if $\phi(\mathcal{X})$ is Morse and $\mathcal{X}$ is hyperbolic.*

Proof. The 'if' direction is a consequence of Lemma 3.12. For the 'only if' direction, stable clearly implies quasigeodesically quasiconvex. The fact that it also implies hyperbolicity is sketched in [34, Lemma 3.3]. The idea is to pull back the transitive family on $\phi(\mathcal{X})$ given in the definition of stability to get such a family in $\mathcal{X}$, and argue that this implies the thin triangle definition of hyperbolicity. □

References

[1] Carolyn R Abbott and David Hume. *The geometry of generalized loxodromic elements*. Ann. Inst. Fourier (Grenoble) (in press). arXiv: 1802.03089v2.

[2] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham. *Largest acylindrical actions and stability in hierarchically hyperbolic groups*. preprint. 2017. arXiv: 1705.06219v2.

[3] Yael Algom-Kfir. *Strongly contracting geodesics in outer space*. Geom. Topol. **15**.4 (2011), pp. 2181–2233.

[4] Noga Alon, Jarosław Grytczuk, Mariusz Hałuszczak, and Oliver Riordan. *Nonrepetitive colorings of graphs*. Random Structures & Algorithms **21**.3-4 (2002), pp. 336–346.

[5] Tarik Aougab, Matthew Gentry Durham, and Samuel J. Taylor. *Pulling back stability with applications to Out(F_n) and relatively hyperbolic groups*. J. Lond. Math. Soc. **96**.3 (2017), pp. 565–583.

[6] Goulnara N. Arzhantseva and Thomas Delzant. *Examples of random groups*. preprint. 2008.

[7] Goulnara N. Arzhantseva and Igor G. Lysenok. *Growth tightness for word hyperbolic groups*. Math. Z. **241**.3 (2002), pp. 597–611.

[8] Werner Ballmann. *Lectures on spaces of nonpositive curvature*. Vol. 25. DMV Seminar. With an appendix by Misha Brin. Basel: Birkhäuser Verlag, 1995, pp. viii+112.

[9] Jason A. Behrstock. *Asymptotic geometry of the mapping class group and Teichmüller space*. Geom. Topol. **10** (2006), pp. 1523–1578.

[10] Jason Behrstock and Cornelia Druţu. *Divergence, thick groups, and short conjugators*. Illinois J. Math. **58**.4 (2014), pp. 939–980.

[11] Mladen Bestvina, Ken Bromberg, Koji Fujiwara, and Alessandro Sisto. *Acylindrical actions on projection complexes*. preprint. 2017. arXiv: 1711.08722v1.

[12] Mladen Bestvina, Kenneth Bromberg, and Koji Fujiwara. *Constructing group actions on quasi-trees and applications to mapping class groups*. Publ. Math. Inst. Hautes Études Sci. **122**.1 (2015), pp. 1–64.

[13] Mladen Bestvina and Koji Fujiwara. *A characterization of higher rank symmetric spaces via bounded cohomology*. Geom. Funct. Anal. **19**.1 (2009), pp. 11–40.

[14] Mladen Bestvina and Koji Fujiwara. *Bounded cohomology of subgroups of mapping class groups*. Geom. Topol. **6** (2002), pp. 69–89.

[15] Mario Bonk. *Quasi-geodesic segments and Gromov hyperbolic spaces*. Geom. Dedicata **62**.3 (1996), pp. 281–298.

[16] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Vol. 319. Grundlehren der mathematischen Wissenschaften. Berlin: Springer, 1999.

[17] Jeffrey Brock and Howard Masur. *Coarse and synthetic Weil-Petersson geometry: quasi-flats, geodesics and relative hyperbolicity*. Geom. Topol. **12**.4 (2008), pp. 2453–2495.

[18] Ruth Charney, Matthew Cordes, and Devin Murray. *Quasi-Mobius Homeomorphisms of Morse boundaries*. Bull. London Math. Soc. **51**.3 (2019), pp. 501–515.

[19] Ruth Charney and Harold Sultan. *Contracting boundaries of* CAT(0) *spaces*. J. Topol. **8**.1 (2015), pp. 93–117.

[20] Matthew Cordes. *A survey on Morse boundaries and stability*. *Beyond Hyperbolicity*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2019, pp. 83–116.

[21] Matthew Cordes. *Morse Boundaries of Proper Geodesic Metric Spaces*. Groups Geom. Dyn. **11**.4 (2017), pp. 1281–1306.

[22] Matthew Cordes and Matthew Gentry Durham. *Boundary convex cocompactness and stability of subgroups of finitely generated groups*. Int. Math. Res. Not. (2016).

[23] Matthew Cordes and David Hume. *Stability and the Morse boundary*. J. Lond. Math. Soc. **95**.3 (2017), pp. 963–988.

[24] Rémi Coulon, Françoise Dal'bo, and Andrea Sambusetti. *Growth gap in hyperbolic groups and amenability*. preprint. 2017. arXiv: 1709.07287v2.

[25] Rémi Coulon, Rhiannon Dougall, Barbara Schapira, and Samuel Tapie. *Twisted Patterson-Sullivan measures and applications to amenability and coverings*. preprint. 2018. arXiv: 1809.10881v2.

[26] Christopher B. Croke and Bruce Kleiner. *Spaces with nonpositive curvature and their ideal boundaries*. Topology **39** (2000), pp. 549–556.

[27] F. Dahmani, V. Guirardel, and D. Osin. *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*. Mem. Amer. Math. Soc. **245**.1156 (2017), pp. v+152.

[28] Françoise Dal'Bo, Marc Peigné, Jean-Claude Picaud, and Andrea Sambusetti. *On the growth of quotients of Kleinian groups*. Ergodic Theory Dynam. Systems **31**.3 (2011), pp. 835–851.

[29] Thomas Delzant. *Sous-groupes distingués et quotients des groupes hyperboliques*. Duke Math. J. **83**.3 (1996), pp. 661–682.

[30] Cornelia Druţu, Shahar Mozes, and Mark Sapir. *Corrigendum to "Divergence in lattices in semisimple Lie groups and graphs of groups"*. Trans. Amer. Math. Soc. **370**.1 (2018), pp. 749–754.

[31] Cornelia Druţu, Shahar Mozes, and Mark Sapir. *Divergence in lattices in semisimple Lie groups and graphs of groups*. Trans. Amer. Math. Soc. **362**.5 (2010), pp. 2451–2505.

[32] Cornelia Druţu and Mark Sapir. *Tree-graded spaces and asymptotic cones of groups*. Topology **44**.5 (2005). With an appendix by Denis Osin and Sapir, pp. 959–1058.

[33] Moon Duchin and Kasra Rafi. *Divergence of geodesics in Teichmüller space and the mapping class group*. Geom. Funct. Anal. **19**.3 (2009), pp. 722–742.

[34] Matthew Gentry Durham and Samuel J. Taylor. *Convex cocompactness and stability in mapping class groups*. Algebr. Geom. Topol. **15**.5 (2015), pp. 2839–2859.

[35] Alex Eskin, Maryam Mirzakhani, and Kasra Rafi. *Counting closed geodesics in strata*. Invent. Math. **215**.2 (2019), pp. 535–607.

[36] Anthony Genevois. *Hyperbolicities in CAT(0) cube complexes*. preprint. 2017. arXiv: 1709.08843v2.

[37] S. M. Gersten. *Divergence in* 3*-manifold groups*. Geom. Funct. Anal. **4**.6 (1994), pp. 633–647.

[38] S. M. Gersten. *Quadratic divergence of geodesics in* CAT(0) *spaces*. Geom. Funct. Anal. **4**.1 (1994), pp. 37–51.

[39] Sébastien Gouëzel and Vladimir Shchur. *Corrigendum: A corrected quantitative version of the Morse lemma*. J. Funct. Anal. **277**.4 (2019), pp. 1258–1268.

[40] R. I. Grigorchuk. *Symmetrical random walks on discrete groups*. *Multicomponent random systems*. Vol. 6. Adv. Probab. Related Topics. Dekker, New York, 1980, pp. 285–325.

[41] Rostislav Grigorchuk and Pierre de la Harpe. *On problems related to growth, entropy, and spectrum in group theory*. J. Dynam. Control Systems **3**.1 (1997), pp. 51–89.

[42] M. Gromov. *Asymptotic invariants of infinite groups*. *Geometric group theory, Vol. 2 (Sussex, 1991)*. Vol. 182. London Math. Soc. Lecture Note Ser. Cambridge: Cambridge Univ. Press, 1993, pp. 1–295.

[43] M. Gromov. *Hyperbolic groups*. *Essays in group theory*. Vol. 8. Math. Sci. Res. Inst. Publ. New York: Springer, 1987, pp. 75–263.

[44] Mikhael Gromov. *Random walk in random groups*. Geom. Funct. Anal. **13**.1 (2003), pp. 73–146.

[45] Dominik Gruber. *Groups with graphical $C(6)$ and $C(7)$ small cancellation presentations*. Trans. Amer. Math. Soc. **367**.3 (2015), pp. 2051–2078.

[46] Dominik Gruber and Alessandro Sisto. *Infinitely presented graphical small cancellation groups are acylindrically hyperbolic*. Ann. Inst. Fourier (Grenoble) **68**.6 (2018), pp. 2501–2552. arXiv: 1408.4488.

[47] N. Higson, V. Lafforgue, and G. Skandalis. *Counterexamples to the Baum-Connes conjecture*. Geom. Funct. Anal. **12**.2 (2002), pp. 330–354.

[48] Heinz Huber. *Über eine neue Klasse automorpher Funktionen und ein Gitterpunktproblem in der hyperbolischen Ebene. I*. Comment. Math. Helv. **30** (1956), pp. 20–62.

[49] Michael Kapovich and Bernhard Leeb. 3*-manifold groups and nonpositive curvature*. Geom. Funct. Anal. **8**.5 (1998), pp. 841–852.

[50] Heejoung Kim. *Stable subgroups and Morse subgroups in mapping class groups*. preprint. 2017. arXiv: 1710.11617v2.

[51] Thomas Koberda, Johanna Mangahas, and Samuel J. Taylor. *The geometry of purely loxodromic subgroups of right-angled Artin groups*. Trans. Amer. Math. Soc. **369**.11 (2017), pp. 8179–8208.

[52] Nataša Macura. *CAT(0) spaces with polynomial divergence of geodesics*. Geom. Dedicata **163** (2013), pp. 361–378.

[53] Howard A. Masur and Yair N. Minsky. *Geometry of the complex of curves. I. Hyperbolicity*. Invent. Math. **138**.1 (1999), pp. 103–149.

[54] Katsuhiko Matsuzaki, Yasuhiro Yabuki, and Johannes Jaerisch. *Normalizer, divergence type and Patterson measure for discrete groups of the Gromov hyperbolic space*. preprint. 2015. arXiv: 1511.02664v1.

[55] Yair N. Minsky. *Quasi-projections in Teichmüller space*. J. Reine Angew. Math. **473** (1996), pp. 121–136.

[56] Harold Marston Morse. *A fundamental class of geodesics on any closed surface of genus greater than one*. Trans. Amer. Math. Soc. **26**.1 (1924), pp. 25–60.

[57] Devin Murray. *Topology and dynamics of the contracting boundary of cocompact* CAT(0) *spaces*. Pacific J. Math. **299**.1 (2019), pp. 89–116.

[58] Yann Ollivier. *On a small cancellation theorem of Gromov*. Bull. Belg. Math. Soc. Simon Stevin **13**.1 (2006), pp. 75–89.

[59] Alexander Yu. Ol′shanskii, Denis V. Osin, and Mark V. Sapir. *Lacunary hyperbolic groups*. Geom. Topol. **13**.4 (2009). With an appendix by M. Kapovich and B. Kleiner, pp. 2051–2140.

[60] Damian Osajda. *Small cancellation labellings of some infinite graphs and applications*. preprint. 2014. arXiv: 1406.5015v2.

[61] D. Osin. *Acylindrically hyperbolic groups*. Trans. Amer. Math. Soc. **368**.2 (2016), pp. 851–888.

[62] Abhijit Pal and Rahul Pandey. *Acylindrical Hyperbolicity of Subgroups*. preprint. 2019. arXiv: 1903.00628v2.

[63] Jouni Parkkonen and Frédéric Paulin. *On the hyperbolic orbital counting problem in conjugacy classes*. Math. Z. **279**.3-4 (2015), pp. 1175–1196.

[64] Yulan Qing and Kasra Rafi. *Sub-linearly contracting boundary I: CAT(0) spaces*. preprint. 2019.

[65] Jacob Russell, Davide Spriano, and Hung Cong Tran. *Convexity in Hierarchically Hyperbolic Spaces*. preprint. 2018. arXiv: 1809.09303v2.

[66] Stéphane Sabourau. *Growth of quotients of groups acting by isometries on Gromov-hyperbolic spaces*. J. Mod. Dyn. **7**.2 (2013), pp. 269–290.

[67] Andrea Sambusetti. *Growth tightness in group theory and Riemannian geometry*. *Recent advances in geometry and topology*. Cluj Univ. Press, Cluj-Napoca, 2004, pp. 341–352.

[68] Andrea Sambusetti. *Growth tightness of free and amalgamated products*. Ann. Sci. École Norm. Sup. (4) **35**.4 (2002), pp. 477–488.

[69] Andrea Sambusetti. *Growth tightness of negatively curved manifolds*. C. R. Math. Acad. Sci. Paris **336**.6 (2003), pp. 487–491.

[70] Andrea Sambusetti. *Growth tightness of surface groups*. Expo. Math. **20**.4 (2002), pp. 345–363.

[71] Vladimir Shchur. *A quantitative version of the Morse lemma and quasi-isometries fixing the ideal boundary*. J. Funct. Anal. **264**.3 (2013), pp. 815–836.

[72] Alessandro Sisto. *Contracting elements and random walks*. J. Reine Angew. Math. **742** (2018), pp. 79–114.

[73] Alessandro Sisto. *Quasi-convexity of hyperbolically embedded subgroups*. Math. Z. (in press). arXiv: 1310.7753.

[74] Harold Sultan. *Hyperbolic quasi-geodesics in CAT(0) spaces*. Geom. Dedicata **169**.1 (2014), pp. 209–224.

[75] Harold Mark Sultan. *Explicit equivalences between CAT(0) hyperbolic type geodesics*. preprint. 2012. arXiv: 1211.6673.

[76] Eric L. Swenson. *Hyperbolic elements in negatively curved groups*. Geom. Dedicata **55**.2 (1995), pp. 199–210.

[77] Hung Tran. *On strongly quasiconvex subgroups*. Geom. Topol. **23**.3 (2019), pp. 1173–1235.

[78] Hung Cong Tran. *Divergence spectra and Morse boundaries of relatively hyperbolic groups*. Geometriae Dedicata (2017).

[79] Hung Cong Tran. *Relative divergence of finitely generated groups*. Algebr. Geom. Topol. **15**.3 (2015), pp. 1717–1769.

[80] Wenyuan Yang. *Genericity of contracting elements in groups*. Math. Ann. (in press).

[81] Wen-yuan Yang. *Growth tightness for groups with contracting elements*. Math. Proc. Cambridge Philos. Soc. **157**.2 (2014), pp. 297–319.

[82] Wen-yuan Yang. *Statistically convex-cocompact actions of groups with contracting elements*. Int. Math. Res. Not. (2018).

PAPER A

Characterizations of Morse quasi-geodesics via superlinear divergence and sublinear contraction

Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume
Documenta Mathematica **22** (2017), 1193–1224.
https://www.math.uni-bielefeld.de/documenta/vol-22/36.html

We introduce and begin a systematic study of sublinearly contracting projections.

We give two characterizations of Morse quasi-geodesics in an arbitrary geodesic metric space. One is that they are sublinearly contracting; the other is that they have completely superlinear divergence.

We give a further characterization of sublinearly contracting projections in terms of projections of geodesic segments.

1. Introduction

This paper initiates a systematic study of *contracting projections*. The aim is to clarify and quantify ways in which a subspace of a geodesic metric space can 'behave like' a convex subspace of a hyperbolic space.

The definition of hyperbolicity captures the notion that a space is uniformly negatively curved on all sufficiently large scales. Following Gromov's seminal paper [21], hyperbolic groups and spaces have been intensively studied and many generalizations of this notion have been considered.

One particular collection of ideas focus on finding 'hyperbolic directions', geodesics that have some of the features exhibited by geodesics in hyperbolic spaces, for instance, those that satisfy the *Morse lemma*, have *superlinear divergence* or satisfy some *contraction* hypothesis. These ideas find application to Mostow rigidity in rank 1 [29], the Rank Rigidity Conjecture for CAT(0) spaces [4, 8, 11], and hyperbolicity of the curve complex of a hyperbolic surface [24, 22]. Recently, the concept of *strongly contracting projection* has been a topic of intense interest in relation to mapping class groups and outer automorphisms of free groups [1, 7], acylindrically hyperbolic groups [16, 27], and contracting/Morse boundaries [30, 31, 13, 14, 25].

We introduce a more general version of contracting projection than has been previously studied. Our main result is that this new version of contraction is equivalent to the Morse property and to a certain superlinear divergence property. We give quantitative links between these various geometric properties. We also generalize several fundamental theorems about stronger versions of contraction to our new, more general, context.

In this paper we establish fundamental results in a very general setting, so that they will be broadly applicable. Indeed, the novel version of contracting projections we introduce here is essential in a subsequent paper [3], in which we explore the geometry of finitely generated graphical small cancellation groups, a class that includes the Gromov monster groups as

The authors thank the referee for the careful reading and helpful remarks, and the Isaac Newton Institute for Mathematical Sciences for support and hospitality during the program *Non-positive curvature: group actions and cohomology* where work on this paper was undertaken. This work is partially supported by EPSRC Grant Number EP/K032208/1.

Goulnara N. Arzhantseva was partially supported by the ERC grant "ANALYTIC" no. 259527 and a grant from the Simons Foundation. Christopher H. Cashen was supported by the Austrian Science Fund (FWF):M1717-N25. Dominik Gruber was supported by the Swiss National Science Foundation Professorship FN PP00P2-144681/1. David Hume was supported by the ERC grant no. 278469 and the grant ANR-14-CE25-0004 "GAMME".

notorious examples. In that paper we engineer finitely generated groups with Cayley graphs that mimic the surprising geometry of our examples from Section 3. In particular, the new spectrum of contracting behaviors in geodesic metric spaces that we discover here does appear in the setting of Cayley graphs of finitely generated groups. We also, in [3], use the equivalence between sublinear contraction and the Morse property established here in Theorem 1.4 to characterize Morse geodesics in certain families of graphical small cancellation groups.

Since the preprint version of this article appeared there have already been other applications of our results, including work of Cordes and Hume [15] and Cashen and Mackay [12] on Morse boundaries of finitely generated groups and work of Aougab, Durham, and Taylor [2] on cocompact subgroups of mapping class groups and $\mathrm{Out}(F_n)$.

We give detailed introductions to the three main geometric properties in Sections 1.1, 1.2, and 1.3 and make precise statements of our results in Sections 1.4 and 1.5.

1.1. Contracting projections. Let Y be a subspace of a geodesic metric space X, and let $\epsilon \geqslant 0$. The *ϵ–closest point projection of X to Y* is the map $\pi_Y^\epsilon\colon X \to 2^Y$ sending a point $x \in X$ to the set:

$$\pi_Y^\epsilon(x) := \{y \in Y \mid d(x,y) \leqslant d(x,Y) + \epsilon\} \subset Y$$

We do not assume the sets $\pi_Y^\epsilon(x)$ have uniformly bounded diameter. Note that given any $x \in X$, $\emptyset \neq Y \subset X$, and $\epsilon > 0$, the set $\pi_Y^\epsilon(x)$ is non-empty.

Definition 1.1. The ϵ–closest point projection $\pi_Y^\epsilon\colon X \to 2^Y$ is *(ρ_1,ρ_2)–contracting* if the following conditions are satisfied.

- The empty set is not in the image of π_Y^ϵ.
- The functions[1] ρ_1 and ρ_2 are non-decreasing and eventually non-negative.
- The function ρ_1 is unbounded and $\rho_1(r) \leqslant r$.
- For all $x,\ x' \in X$, if $d(x,x') \leqslant \rho_1(d(x,Y))$ then:
$$\operatorname{diam} \pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(x') \leqslant \rho_2(d(x,Y))$$
- $\lim_{r\to\infty} \frac{\rho_2(r)}{\rho_1(r)} = 0$.

We say that *Y is (ρ_1,ρ_2)–contracting* if there exists $\epsilon \geqslant 0$ such that π_Y^ϵ is (ρ_1,ρ_2)–contracting, see Definition 6.4. We say a collection of subspaces $\{Y_i\}_{i\in\mathcal{I}}$ is *uniformly contracting* if there exist ρ_1 and ρ_2 such that for all $i \in \mathcal{I}$, the subspace Y_i is (ρ_1,ρ_2)–contracting.

The rough idea is that, asymptotically as x gets far from Y, if B is a ball centered at x and disjoint from Y then the diameter of its projection is negligible compared to the diameter of B. More accurately, this is true at a specific scale — when the radius of B is $\rho_1(d(x,Y))$. We claim no finer control of the projection diameter when B has smaller radius.

For a simple, but conceptually useful, example, consider a circle X and an arc $Y \subset X$. Take $\rho_1(r) := r$, and let ρ_2 be the constant function whose value is the distance between the endpoints of Y. Then π_Y^0 is (ρ_1,ρ_2)–contracting. There is a unique point $x \in X$ farthest from Y. The ball B of radius $\rho_1(d(x,Y))$ about x is all of $X \setminus Y$, and $\pi_Y^0(B) = \pi_Y^0(x)$ has diameter $\rho_2(d(x,Y))$.

The simplest example that is not (ρ_1,ρ_2)–contracting for any choice of ρ_1 and ρ_2 is to take X to be the Euclidean plane and take Y to be a geodesic. Then the diameter of π_Y^0 of any ball is equal to the diameter of the ball, so we cannot satisfy $\lim_{r\to\infty} \frac{\rho_2(r)}{\rho_1(r)} = 0$.

The simplest contracting example with Y unbounded is to take X to be a tree and Y to be an unbounded convex subspace. Then $\operatorname{diam} \pi_Y^0(B_{d(x,Y)}(x)) = 0$ for every x, so π_Y^0 is $(r,0)$–contracting. In more general δ–hyperbolic spaces, ϵ–closest point projection to a geodesic is (r,D)–contracting for some D depending only on δ and ϵ. Such a case, when $\rho_1(r) := r$ and ρ_2 is bounded, is called *strongly contracting*.

Pseudo-Anosov axes in Teichmüller space are strongly contracting [24], as are iwip axes in the Outer Space of the outer automorphism group of a free group [1] and axes of rank 1 isometries of CAT(0) spaces [4, 8].

---

[1] The term 'function' always refers to a real valued function whose domain, unless otherwise noted, is the non-negative real numbers.

We say that π_Y^ϵ is *semi-strongly contracting* if it is (ρ_1, ρ_2)–contracting for $\rho_1(r) := r/2$ and ρ_2 bounded. Related notions have been considered in the context of the mapping class group of a hyperbolic surface [22, 5, 19].

We say that π_Y^ϵ is *sublinearly contracting* if it is (ρ_1, ρ_2)–contracting for $\rho_1(r) := r$. In this case the definition implies ρ_2 is a sublinear function, see Definition 2.1. Similarly, π_Y^ϵ is *logarithmically contracting* if it is (ρ_1, ρ_2)–contracting for $\rho_1(r) := r$ and ρ_2 logarithmic.

A schematic diagram of different contracting behaviors is given in Figure 1. A wide range of examples are presented in Section 3.

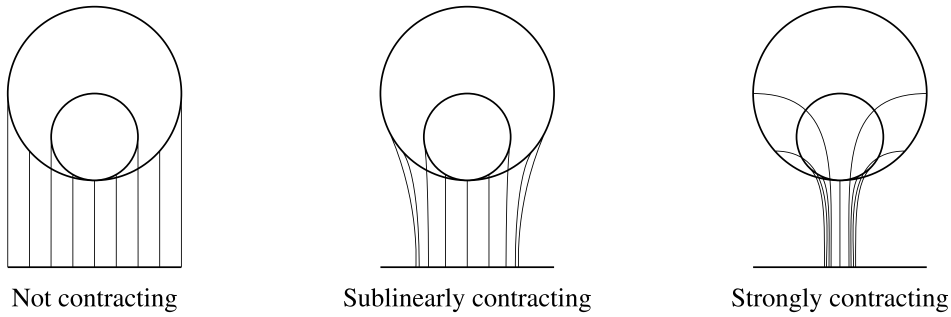

FIGURE 1. Types of contraction

1.2. The Morse property.

DEFINITION 1.2. A subspace Y of a geodesic metric space X is *μ–Morse* for a function $\mu\colon [1,\infty) \times [0,\infty) \to [0,\infty)$ if for every $L \geqslant 1$ and $A \geqslant 0$, every (L, A)–quasi-geodesic γ with endpoints on Y remains within distance $\mu(L, A)$ of Y.

The subspace Y is called *Morse*, or is said to *have the Morse property*, if it μ–Morse for some function μ. A collection of subspaces $\{Y_i\}_{i\in\mathcal{I}}$ is said to be *uniformly Morse* if there exists a function μ such that for all $i \in \mathcal{I}$ the subspace Y_i is μ–Morse.

Morse quasi-geodesics have been intensively studied: they play a key role in boundary theory for hyperbolic and relatively hyperbolic groups. Recently, the Charney school [30, 31, 13, 14, 25] has been generalizing such boundary theories to arbitrary proper geodesic metric spaces using the so called 'Morse boundary' consisting of asymptotic equivalence classes of Morse rays.

Morse quasi-geodesics have been characterized[2] in terms of cut-points in asymptotic cones [17]: a quasi-geodesic q in X is Morse if and only if every point x in the limit $\mathfrak{q}$ of q in any asymptotic cone C of X is a cut-point separating ends of $\mathfrak{q}$; that is, $C \setminus \{x\}$ has at least two connected components containing points of $\mathfrak{q}$. Cut-points in asymptotic cones are a key element of the proof of the quasi-isometry invariance of relatively hyperbolic (asymptotically tree-graded) spaces [18]. It remains a very important open question to determine whether a space in which every asymptotic cone admits a cut-point necessarily admits a Morse quasi-geodesic.

As a result, it is of great interest to find and classify Morse quasi-geodesics. If a solvable group admits a Morse quasi-geodesic then it is virtually cyclic, and the same holds for any other group satisfying a non-trivial law, for instance, a torsion group with bounded exponent [18]. At the other extreme, every quasi-geodesic in a hyperbolic space is Morse. There are non-trivial classifications of Morse quasi-geodesics for relatively hyperbolic groups [28] and CAT(0) spaces [4, 8, 31]. We use the tools of this paper to perform such a classification for graphical small cancellation groups in [3].

[2] See also the related "middle recurrence" characterization of the Morse property in [17].

1.3. Divergence. Closely related to the study of Morse quasi-geodesics is the notion of *divergence*. The definition is technical, so we postpone it until Definition 5.1. The idea is that the divergence of a quasi-geodesic γ in a space X is a function whose value at r is the minimal length of a path in X circumventing a ball of radius r centered on γ. In our version of divergence we allow the forbidden ball to be centered at different points of γ for different values of r. Some authors require the balls to have fixed center at $\gamma(0)$.

Morse geodesics were used to produce cut points in asymptotic cones. Divergence can be used to rule them out [17]: if G is a finitely generated group then no asymptotic cone of G admits a cut point if and only if there exists a constant K such that for any finite geodesic $[a, b]$ with midpoint c, there is a path from a to b avoiding the ball centered at c with radius $d(a, b)/4 - 2$ of length at most $Kd(a, b) + K$. The interplay between divergence and Morse quasi-geodesics is explored in [17] and [6].

Morally, for a quasi-geodesic γ the Morse property and linear divergence are opposites. The Morse property says good (quasi-geodesic) paths between points of γ stay close to γ, and linear divergence says it is easy for a path between points of γ to stray far from γ. However, there are some subtleties. There are groups that admit quasi-geodesics with superlinear divergence, yet have an asymptotic cone with no cut point, and therefore no Morse quasi-geodesics [26]. By construction, for each of these groups there is an unbounded sequence (r_n) such that the divergence is linear (it satisfies the above conditions for a fixed K) whenever $d(a, b) = r_n$ for some n. We say a geodesic metric space has *completely superlinear divergence* if no such unbounded sequence exists. We show in Theorem 1.5 that this is the precise divergence property that characterizes Morse quasi-geodesics.

1.4. Main theorems. Restricted to quasi-geodesics, our main results say:

Theorem 1.3. *Let X be a geodesic metric space. Let γ be a quasi-geodesic in X. The following are equivalent:*

(1) γ is sublinearly contracting.
(2) γ is Morse.
(3) γ has completely superlinear divergence.

Special cases of this theorem have appeared before. If X is hyperbolic then these conditions are well-known properties of arbitrary quasi-geodesics, and conditions (1) and (3) can be strengthened to 'strongly contracting' and 'at least exponential divergence', respectively. If X is CAT(0) and γ is a geodesic then this is a recent theorem of Charney and Sultan [13]. In that case, conditions (1) and (3) can be strengthened to 'strongly contracting' and 'at least quadratic divergence', respectively. Our theorem establishes these equivalences in full generality.

The Morse and contraction properties make sense for *subspaces* of X, not just quasi-geodesics. Our main theorem is:

Theorem 1.4. *Let Y be a subspace of a geodesic metric space X. Let $\epsilon \geqslant 0$ be such that π_Y^ϵ does not contain the empty set in its image. The following are equivalent:*

(1) There exists $\mu\colon [1, \infty) \times [0, \infty) \to [0, \infty)$ such that Y is μ–Morse.
(2) There exists $\mu'\colon [1, \infty) \to [0, \infty)$ such that every continuous $(L, 0)$–quasi-geodesic with endpoints on Y remains in the $\mu'(L)$-neighborhood of Y.
(3) There exists ρ such that π_Y^ϵ is (r, ρ)–contracting.
(4) There exist ρ_1 and ρ_2 such that π_Y^ϵ is (ρ_1, ρ_2)–contracting.

Moreover, in each implication we bound the parameters of the conclusion in terms of the parameters of the hypothesis, independent of Y.

Divergence, on the other hand, is specialized to quasi-geodesics.

Theorem 1.5. *Let γ be a quasi-geodesic in a geodesic metric space X. The following are equivalent:*

(1) γ is Morse.
(2) γ has completely superlinear divergence.

Moreover, the Morse function can be bounded in terms of the divergence function, independent of γ.

We mention a further characterization of Morse quasi-geodesics: It can be shown fairly easily that a quasi-geodesic $\gamma\colon I \to X$ is Morse if and only if the collection of its subsegments $\{\gamma_J \mid J$ is a subinterval of $I\}$ is uniformly Morse. Moreover, the Morse functions for γ and for the subsegments can be bounded in terms of one another and the quasi-geodesic constants of γ. The quantitative nature of the equivalences in Theorem 1.4 then implies that γ is Morse if and only if the collection of its subsegments is uniformly contracting.

1.5. Further applications. We consider several important theorems about strongly contracting projections that have appeared in the literature, and generalize them by proving sublinear analogues.

The first of these results is the 'Bounded Geodesic Image Property', cf [23, 8]. This says that if π_Y^ϵ is strongly contracting then there exist constants A and B such that if γ is a geodesic segment with $d(\gamma, Y) \geqslant A$, then $\operatorname{diam} \pi_Y^\epsilon(\gamma) \leqslant B$. In fact, this property is equivalent to strong contraction. We prove, in Theorem 7.1, that π_Y^ϵ is (r,ρ)–contracting if and only if there exist a constant A and a function $\rho' \asymp \rho$ such that if γ is a geodesic segment with $d(\gamma, Y) \geqslant A$ then

$$\operatorname{diam} \pi_Y^\epsilon(\gamma) \leqslant \rho'(\max\{d(x, Y), d(x', Y)\}),$$

where x and x' are the endpoints of γ.

The second strong contraction result is one of the 'Projection Axioms' of Bestvina, Bromberg, and Fujiwara [7]. It says that if π_Y^ϵ and $\pi_{Y'}^{\epsilon'}$ are both strongly contracting, and if Y and Y' are sufficiently far apart, then $\operatorname{diam} \pi_Y^\epsilon(Y')$ and $\operatorname{diam} \pi_{Y'}^{\epsilon'}(Y)$ are bounded in terms of the contraction constants. In Proposition 8.2 we prove that if 'strongly contracting' is weakened to '(r,ρ)–contracting' then $\operatorname{diam} \pi_Y^\epsilon(Y')$ and $\operatorname{diam} \pi_{Y'}^{\epsilon'}(Y)$ are bounded by an affine function of $\rho(d(Y, Y'))$. This is the best that can be expected, since even for a single point x we can only conclude $\operatorname{diam} \pi_Y^\epsilon(x) \leqslant \rho(d(x, Y))$.

Finally, a theorem of Masur and Minsky [22] says, approximately and in our language, that if for every pair of points in a geodesic metric space X there exists a path between them such that these paths all admit semi-strongly contracting projections, with contraction constants uniform over the family of paths, then the space X is hyperbolic. Our Corollary 8.4 says the conclusion still holds if 'semi-strongly contracting' is weakened to 'sublinearly contracting'.

1.6. Robustness. In Section 6 we investigate the following question: Let π_Y^ϵ be (ρ_1, ρ_2)–contracting. What effect does changing ρ_1, ϵ, or Y have on this property, in terms of ρ_2?

We obtain optimal answers when $\rho_1(r) = r$, see Lemma 6.2 and Lemma 6.3. It would be interesting to have good quantitative results in more general cases.

The Morse property is invariant under quasi-isometry, so, by Theorem 1.5, the property of being sublinearly contracting is also a quasi-isometry invariant. Very little is known, however, about how the contraction parameters vary under quasi-isometry. In a subsequent paper [3] we demonstrate that strong contraction is not preserved by quasi-isometries.

2. Preliminaries

Let $N_r(y) := \{x \in X \mid d(x, y) < r\}$ and $\overline{N}_r(y) := \{x \in X \mid d(x, y) \leqslant r\}$. If Y is a subspace of X, let $N_r(Y) := \cup_{y \in Y} N_r(y)$, and $\overline{N}_r(Y) := \cup_{y \in Y} \overline{N}_r(y)$.

Let $\operatorname{diam} Y := \sup\{d(y, y') \mid y,\ y' \in Y\}$.

A *geodesic* is an isometric embedding of an interval. A metric space X is *geodesic* if for every pair of points $x, x' \in X$ there exists a geodesic connecting them.

The *Hausdorff distance* between non-empty subspaces Y and Z of X is the infimal C such that $Y \subset \overline{N}_C(Z)$ and $Z \subset \overline{N}_C(Y)$. Two subspaces are *C–Hausdorff equivalent* if the Hausdorff distance between them is at most C.

Given $L \geqslant 1$ and $A \geqslant 0$, a map $\phi\colon X \to Y$ between metric spaces is an (L, A)*–quasi-isometric embedding* if $\frac{1}{L} d(x, x') - A \leqslant d(\phi(x), \phi(x')) \leqslant L d(x, x') + A$ for every $x,\ x' \in X$. It is an (L, A)*–quasi-isometry* if, in addition, $Y = \overline{N}_A(\phi(X))$.

An (L, A)–*quasi-geodesic* is an (L, A)–quasi-isometric embedding of an interval.

DEFINITION 2.1. A function f is *sublinear* if it is non-decreasing, eventually non-negative, and $\lim_{r\to\infty} \frac{f(r)}{r} = 0$.

We write $f \preceq g$ if there exist constants $C_1 > 0$, $C_2 > 0$, $C_3 \geqslant 0$, and $C_4 \geqslant 0$ such that $f(r) \leqslant C_1 g(C_2 r + C_3) + C_4$ for all r. This partial order gives an equivalence relation $f \asymp g$ if $f \preceq g$ and $f \succeq g$. If $f \asymp g$ we say f and g are *asymptotic*.

3. Examples of contraction

We begin with a classical example.

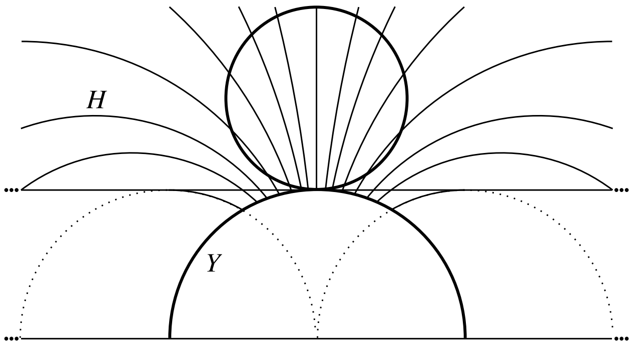

FIGURE 2. Contraction in $\mathbb{H}^2$.

EXAMPLE 3.1. Let X be the hyperbolic plane, with the upper half-space model, and let Y be the geodesic that is the upper half of the unit circle, see Figure 2. Pick any point $x \notin Y$. Up to isometry, we may assume x sits on the y–axis above Y. The ball of radius $d(x, Y)$ about x is contained in the horoball $H := \{(a, b) \in \mathbb{R}^2 \mid b \geqslant 1\}$. The closest point projection of H to Y has diameter $\ln(3 + 2\sqrt{2})$. Thus, π_Y^0 is $(r, \ln(3 + 2\sqrt{2}))$–contracting.

We now construct examples exhibiting a wider range of contracting behaviors than have appeared previously in the literature.

EXAMPLE 3.2. Let $\rho = \rho_1 \colon [0, \infty) \to [0, \infty)$ be an unbounded function such that $\rho(r) \leqslant r$, $\mathrm{Id} - \rho$ is unbounded, and there exists an $A \geqslant 0$ with $\rho(A) > 0$ such that $0 \leqslant \rho(a + b) - \rho(a) < b$ for all $a \geqslant A$ and $b > 0$. We construct a space X and $Y \subset X$ such that π_Y^0 is $(\rho, 2)$–contracting but not strongly contracting.

The map $\phi\colon [A, \infty) \to [A - \rho(A), \infty) : x \mapsto x - \rho(x)$ is a bijection by assumption. We set $\sigma(0) := \phi(A)$ and, for $i \in \mathbb{N}$, recursively define[3] $\sigma(i + 1) := \phi^{-1}(\sigma(i))$. This is well-defined since $[A, \infty) \subset [A - \rho(A), \infty)$. Rearranging this expression yields $\rho(\sigma(i + 1)) = \sigma(i + 1) - \sigma(i)$. Note that $\sigma(i + 1) - \sigma(i) \geqslant \rho(A) > 0$ for every $i \in \mathbb{N} \cup \{0\}$, whence, in particular, $\sigma(i) \to \infty$ as $i \to \infty$.

Let $Y := [0, \infty)$ be a ray. For $i \in \mathbb{N} \cup \{0\}$, let Z_i be a segment of length $\sigma(i)$ with endpoints labelled y_i and z_i. Identify y_i with the point i in Y. Let W_i be a segment of length $\sigma(i+1)-\sigma(i)+1$ connecting z_i to z_{i+1}. Let X be the resulting geodesic metric space. See Figure 3.

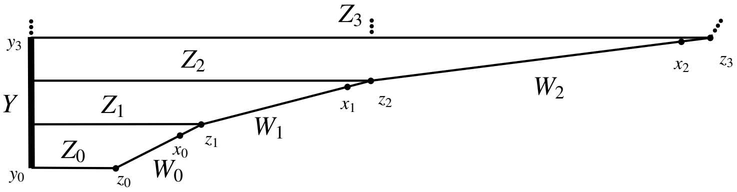

FIGURE 3. $(\rho_1, 2)$–contraction

[3] An *Abel function* for f is a function α such that $\alpha(f(x)) = \alpha(x) + 1$. The function σ is the inverse of an Abel function for ϕ^{-1}.

Let x_i be the point of W_i at distance $1/2$ from z_{i+1}. Clearly $\operatorname{diam}\pi^0_Y(x_i) = 1$. It is easy to see that each complementary component of $X \setminus (Y \cup \{x_i\}_{i\in\mathbb{N}\cup\{0\}})$ projects to a single point of Y. Now consider the ball of radius $\rho(d(x,Y))$ about some x. First assume $x \in W_i$ for some i. Our assumptions on ρ yield:

$$\overline{N}_{\rho(d(x,Y))}(x) \subseteq W_i \cup \overline{N}_{\rho(\sigma(i))+1/2}(z_i) \cup N_{\rho(\sigma(i+1))+1/2}(z_{i+1})$$

The latter may contain x_i and x_{i+1} but no other x_j. If, on the other hand, x is in some Z_i, then $\overline{N}_{\rho(d(x,Y))}(x)$ is contained in $Z_i \cup \overline{N}_{\rho(d(z_i,Y))}(z_i)$. Therefore, for any $x \in X$, we have that $\pi^0_Y(\overline{N}_{\rho(d(x,Y))}(x))$ has diameter at most 2.

Observe that $\overline{N}_{d(z_i,Y)}(z_i)$ contains $\{z_j, z_{j+1}, \dots z_i\}$ for $0 \leqslant i - j \leqslant \sigma(j)$. Since $\sigma(i) \to \infty$ as $i \to \infty$, this implies that Y is not strongly contracting.

Concrete examples include:

- $\rho(r) := 2\sqrt{r} - 1$ and $A = 1$ and $\sigma(r) := r^2$.
- $\rho(r) := r/2$ and $A = 2$ and $\sigma(r) := 2^r$. This is an example of semi-strong contraction.
- $\rho(r) := \min\{r, r - \log_2 r\}$ and $A = 2$ and $\sigma(r) := 2 \uparrow\uparrow r$.

In Knuth's 'up-arrow notation' $2 \uparrow\uparrow r$ denotes tetration, so that $2 \uparrow\uparrow r = \underbrace{2^{\cdot^{\cdot^{2}}}}_{r\text{ times}}$ when $r \in \mathbb{N} \cup \{0\}$.

The following proposition shows that it is sometimes possible to 'trade' between the input and output contraction functions, so we can use Example 3.2 to demonstrate further examples of (ρ_1, ρ_2)–contraction conditions.

Proposition 3.3. *Suppose that π^ϵ_Y is (ρ_1, B)–contracting, where B is a constant and $\rho = \rho_1$ is a non-decreasing, non-negative, unbounded function such that* $\mathrm{Id} - \rho$ *is unbounded and such that there exists a constant A such that $\rho(A) > 0$ and $0 \leqslant \rho(a + b) - \rho(a) < b$ for all $a \geqslant A$ and $b > 0$. Define $A' := A - \rho(A)$. For $x \in [A', \infty)$ define*[4] *$\alpha(x)$ to be the minimal non-negative integer such that $(\mathrm{Id} - \rho)^{\alpha(x)}(x) \in [A', A)$. Then π^ϵ_Y is $(r - A, \rho_2)$–contracting for some $\rho_2 \asymp \alpha$.*

Proof. Observe as in Example 3.2 that the map $\phi\colon\ x \mapsto x - \rho(x)$ is a bijection $[A, \infty) \to [A', \infty)$ and that, since ϕ is strictly increasing for $x \geqslant A$, the collection $\{[\phi^k(A'), \phi^{k-1}(A')) \mid k \leqslant 0\}$ is a partition of $[A', \infty)$.

We show that $\rho_2(r) := B\alpha(r)$ will suffice. It follows from unboundedness of ρ that ρ_2 is sublinear: we have $\rho_2 \asymp \alpha$. The map α is a step function with steps of height 1, so it is sufficient to show that the lengths of the steps go to infinity, ie $\phi^{-n-1}(A') - \phi^{-n}(A') \to \infty$ as $n \to \infty$. As computed in Example 3.2, we have $\rho(\phi^{-n-1}(A')) = \phi^{-n-1}(A') - \phi^{-n}(A')$. Since $\phi^{-n-1}(A') \to \infty$ as $n \to \infty$ as argued in Example 3.2 and since ρ goes to infinity, sublinearity follows.

Let x and y be points of X such that $d(x, y) \leqslant d(x, Y) - A$. Define $r_0 := d(x, Y)$ and while $r_0 - r_i \leqslant d(x, y)$, define $r_{i+1} := \phi(r_i)$. Note that this is well-defined, ie $r_i \geqslant A$, since $r_0 - r_i \leqslant d(x, y) \leqslant r_0 - A$. Let k be the largest index such that $r_0 - r_k \leqslant d(x, y)$. Then the fact that $\phi^{\alpha(r_0)}(r_0) < A$ and the observation we just made shows $k < \alpha(r_0)$.

Fix a geodesic from x to y and for $0 \leqslant i \leqslant k$ define x_i to be the point at distance $r_0 - r_i$ from x along this geodesic. Define $x_{k+1} := y$. For $0 \leqslant i \leqslant k$ we have $d(x_{i+1}, x_i) \leqslant \rho(d(x_i, Y))$ by construction, whence:

$$\operatorname{diam}\pi^\epsilon_Y(x) \cup \pi^\epsilon_Y(y) \leqslant \sum_{i=0}^{k} \operatorname{diam}\pi^\epsilon_Y(x_i) \cup \pi^\epsilon_Y(x_{i+1}) \leqslant B\alpha(r_0)$$

Thus, π^ϵ_Y is $(r - A, \rho_2)$–contracting. □

Applying Proposition 3.3 to the concrete examples in Example 3.2 we see:

- $(2\sqrt{r} - 1, 2)$–contracting implies $(r - 1, \rho_2)$–contracting for $\rho_2 \asymp \sqrt{\cdot}$.
- $(r/2, 2)$–contracting implies $(r - 2, \rho_2)$–contracting for $\rho_2 \asymp \log_2$.
- Finally, $(r - \log_2 r, 2)$–contracting implies $(r - 2, \rho_2)$–contracting for $\rho_2 \asymp \operatorname{superlog}_2$.

[4]The function $\alpha\colon\ [A', \infty) \to \mathbb{N} \cup \{0\}$ is an Abel function for $(\mathrm{Id} - \rho)^{-1}$. For instance, take α to be the inverse of $\sigma\colon\ \mathbb{N} \cup \{0\} \to \sigma(\mathbb{N} \cup \{0\})$ from Example 3.2 extended to all of $[A', \infty)$ by a rounding-off function.

That the converse to Proposition 3.3 can fail follows from the next example.

EXAMPLE 3.4. Let ρ_2 be a sublinear function such that $0 < \rho_2(r) < r$. Let Y be a line. Choose a collection of disjoint intervals $\{I_i\}_{i\in\mathbb{N}}$ of Y such that $|I_i| = \rho_2(i)$ and let y_i be the center of I_i. Connect the endpoints of I_i by attaching a segment J_i of length $4i$, and let x_i be the center of this segment. Let X be the resulting geodesic space, see Figure 4. We claim π_Y^0 is (r, ρ_2)–contracting.

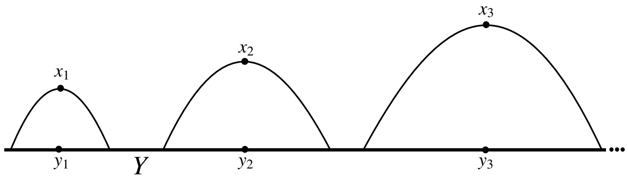

FIGURE 4. (r, ρ_2)–contracting

Suppose that $x \in J_i \subset X$ and $d(x, Y) < i$. Then $d(x, x_i) > d(x, Y)$, and $\operatorname{diam} \pi_Y^0(\overline{N}_{d(x,Y)}(x)) = 0$. For $x \in J_i \subset X$ with $d(x, Y) \geqslant i$ we have $d(x, x_i) \leqslant d(x, Y)$ and:

$$\operatorname{diam} \pi_Y^0(\overline{N}_{d(x,Y)}(x)) = \operatorname{diam} \pi_Y^0(x_i) = \rho_2(i) \leqslant \rho_2(d(x, Y))$$

This proves the claim. Furthermore, ρ_2 is optimal, in the following sense: Since $\operatorname{diam} \pi_Y^0(x_i) = \rho_2(d(x_i, Y)/2) = \rho_2(i)$, if ρ_1' and ρ_2' are some other functions such that π_Y^0 is (ρ_1', ρ_2')–contracting then $\rho_2(i) \leqslant \rho_2'(2i)$ for $i \in \mathbb{N}$.

4. The Morse property

The following two propositions establish our main result, Theorem 1.4.

PROPOSITION 4.1. *Let Y be a subspace of a geodesic metric space X. Suppose π_Y^ϵ is (ρ_1, ρ_2)–contracting. There exists a function μ, depending only on ϵ, ρ_1, and ρ_2, such that Y is μ–Morse.*

PROOF. Given L' and A' there exist L, A, and C such that every (L', A')–quasi-geodesic is C–Hausdorff equivalent to a continuous (L, A)–quasi-geodesic with the same endpoints [10, Lemma III.H.1.11]. Thus, it suffices to show there exists a bound B, depending only on ϵ, ρ_1 and ρ_2, such that every continuous (L, A)–quasi-geodesic connecting points on Y is contained in $N_B(Y)$. Then we set $\mu(L', A') := B + C$.

Let γ be a continuous (L, A)–quasi-geodesic with endpoints on Y. Take E to be sufficiently large so that $\rho_1(E) > 3A$ and for all $r \geqslant E$ we have $\frac{\rho_2(r)}{\rho_1(r)} < \frac{1}{3L^2}$.

Suppose $\gamma \not\subseteq N_E(Y)$, and let $[a, b]$ be a maximal subinterval of the domain of γ such that $\gamma|_{[a,b]} \subset X \setminus N_E(Y)$. We show there exists a T independent of γ and Y such that $b - a \leqslant T$. We conclude by setting $B := E + L \cdot \frac{T}{2} + A$.

Let $t_0 := a$. Supposing we have defined $t_0, \ldots, t_i$, if $d(\gamma(t_i), \gamma(b)) > \rho_1(d(\gamma(t_i), Y))$ define t_{i+1} to be the first time that $d(\gamma(t_i), \gamma(t_{i+1})) = \rho_1(d(\gamma(t_i), Y))$. Such a t_{i+1} exists because γ is continuous. Since $d(\gamma, Y) \geqslant E$ we have $d(\gamma(t_i), \gamma(t_{i+1})) = \rho_1(d(\gamma(t_i), Y)) \geqslant \rho_1(E) > 0$, so after finitely many steps we reach an index k such that $d(\gamma(t_k), \gamma(b)) \leqslant \rho_1(d(\gamma(t_k), Y))$. Applying the contraction condition to the points $\gamma(t_i)$, we see:

$$\operatorname{diam} \pi_Y^\epsilon(\gamma(a)) \cup \pi_Y^\epsilon(\gamma(b)) \leqslant \sum_{i=0}^{k} \rho_2(d(\gamma(t_i), Y))$$

This allows us to estimate:

$$\begin{aligned} d(\gamma(a), \gamma(b)) &\leqslant d(\gamma(a), \pi_Y^\epsilon(\gamma(a))) \\ &\qquad + \operatorname{diam} \pi_Y^\epsilon(\gamma(a)) \cup \pi_Y^\epsilon(\gamma(b)) + d(\gamma(b), \pi_Y^\epsilon(\gamma(b))) \\ &\leqslant 2(E + \epsilon) + \sum_{i=0}^{k} \rho_2(d(\gamma(t_i), Y)) \qquad (1) \end{aligned}$$

On the other hand, since γ is a (L, A)–quasi-geodesic, we have:

$$\begin{aligned}
Ld(\gamma(a), \gamma(b)) + LA \geqslant b - a &= b - t_k + \sum_{i=0}^{k-1}(t_{i+1} - t_i) \\
&\geqslant \frac{1}{L}(d(\gamma(b), \gamma(t_k)) - A) + \sum_{i=0}^{k-1}\frac{1}{L}(d(\gamma(t_{i+1}), \gamma(t_i)) - A) \\
&= \frac{1}{L}(d(\gamma(b), \gamma(t_k)) - \rho_1(d(\gamma(t_k), Y))) \\
&\quad + \sum_{i=0}^{k}\frac{1}{L}(\rho_1(d(\gamma(t_i), Y)) - A) \\
&\geqslant \frac{-d(\gamma(b), Y)}{L} + \sum_{i=0}^{k}\frac{1}{L}(\rho_1(d(\gamma(t_i), Y)) - A) \\
&= -\frac{E}{L} + \sum_{i=0}^{k}\frac{1}{L}(\rho_1(d(\gamma(t_i), Y)) - A)
\end{aligned}$$

Combining this with the previous inequality and rearranging terms, we have:

$$\sum_{i=0}^{k}\Big(\rho_1(d(\gamma(t_i), Y)) - L^2\rho_2(d(\gamma(t_i), Y)) - A\Big) \leqslant E + L^2A + 2L^2(E + \epsilon)$$

Now, left hand side is at least $L^2\sum_{i=0}^{k}\rho_2(d(\gamma(t_i), Y))$, by our choice of E; combined with (1), this gives us:

$$d(\gamma(a), \gamma(b)) \leqslant \frac{E}{L^2} + A + 4(E + \epsilon)$$

This estimate and the fact that γ is a quasi-geodesic give us a bound for $b - a$. □

Proposition 4.2. *Let Y be a subspace of a geodesic metric space X. Suppose there is a non-decreasing function μ such that every continuous $(L, 0)$-quasi-geodesic with endpoints on Y is contained in the closed $\mu(L)$–neighborhood of Y. Suppose the empty set is not in the image of π_Y^ϵ. Then there is a function ρ', depending only on μ and ϵ, such that π_Y^ϵ is (r, ρ')–contracting.*

We remark that since an $(L, 0)$–quasi-geodesic is also an $(L', 0)$–quasi-geodesic for any $L' > L$, there is no loss in requiring the Morse function to be non-decreasing.

Proof. Consider the optimal contraction function:

$$\rho(r) := \sup_{d(x,y)\leqslant d(x,Y)\leqslant r} \operatorname{diam}\pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(y) \leqslant 4r + 2\epsilon$$

Our goal is define a function ρ' depending on μ and ϵ that is non-negative, non-decreasing, and sublinear and such that ρ' is an upper bound for ρ.

Define $\rho'(r) := 0$ if $\epsilon = 0$ and $\mu \equiv 0$. In this case ρ' clearly has the first three properties. Otherwise, we first replace μ by $s \mapsto \inf_{t>s}\mu(s)$. The new μ still satisfies the hypotheses of the proposition and has that additional property that it is right continuous: $\lim_{t\to s^+}\mu(t) = \mu(s)$ for all $s \geqslant 1$. Define $\rho'(0) := 2\epsilon$ and for $r > 0$ define:

$$\rho'(r) := \sup\left\{s \leqslant 4r + 2\epsilon \mid s \leqslant 18\mu\left(\frac{3(4r + 2\epsilon)}{s}\right) + 12\epsilon\right\}$$

If $\mu \equiv 0$ then ρ' increases linearly from 2ϵ to 12ϵ and then remains constant, so it is non-negative, non-decreasing, and sublinear.

If $\mu \not\equiv 0$ then $\rho'(r) > 0$ when $r > 0$, and the conditions on μ ensure ρ' is actually a maximum. The fact that it is non-decreasing then follows by observing that $\rho'(r)$ participates in the supremum defining $\rho'(r')$ when $0 \leqslant r < r'$. To see ρ' is sublinear, we suppose that $\limsup_{r\to\infty}\rho'(r)/r > 0$ and derive a contradiction. Suppose that there exists some $\delta > 0$ and a sequence (r_i) of positive numbers increasing without bound such that $\rho'(r_i) > \delta r_i$ for all i. By definition of ρ', for each i there exists $\delta r_i < s_i \leqslant 4r_i + 2\epsilon$ such that:

$$s_i \leqslant 18\mu\left(\frac{3(4r_i + 2\epsilon)}{s_i}\right) + 12\epsilon \leqslant 18\mu\left(\frac{3(4r_i + 2\epsilon)}{\delta r_i}\right) + 12\epsilon$$

This is a contradiction, since the left-hand side grows without bound while the right-hand side is bounded above by $18\mu(\frac{12}{\delta} + 1) + 12\epsilon$ once i is sufficiently large.

Now we must show $\rho(r) \leqslant \rho'(r)$. It suffices to check this for those r such that $\rho(r) > 0$. The idea of the proof is to choose, for each such r, points x and y such that $d(x, y) \leqslant d(x, Y) \leqslant r$ whose projection diameters nearly realize $\rho(r)$. Take a path γ that is a concatenation of geodesics from a projection point of x to x, then from x to y, then from y to a projection point of y. For $L := \frac{3(4r+2\epsilon)}{\rho(r)} \geqslant 3$ we show that we can make γ into an $(L, 0)$–quasi-geodesic γ' by introducing at most two shortcuts in a particular way. The Morse hypothesis implies that γ' is contained in the $\mu(L)$–neighborhood of Y. We then argue that the condition $d(x, y) \leqslant d(x, Y)$ implies:

$$\rho(r) < 18\mu(L) + 12\epsilon \tag{2}$$

In the case that $\epsilon = 0$ and $\mu \equiv 0$, this gives a contradiction, which means that there is no r for which ρ takes a positive value, and we have $\rho(r) = \rho'(r) = 0$ for all r. Otherwise, plugging the value of L into (2), we conclude that $\rho(r)$ participates in the supremum defining $\rho'(r)$, whence $\rho(r) \leqslant \rho'(r)$.

First we show how to produce quasi-geodesics. Consider points x, y, $p_x \in \pi_Y^\epsilon(x)$, and $p_y \in \pi_Y^\epsilon(y)$. Let $\gamma := [p_x, x][x, y][y, p_y]$ be a concatenation of three geodesics. Let $[p, q]_\gamma$ denote the subsegment of γ from p to q, and let $|[p, q]_\gamma|$ denote its length. For this part of the argument we may use any $L \geqslant \frac{|\gamma|}{d(p_x,p_y)} \geqslant 1$. Consider the continuous function $D(p, q) := Ld(p, q) - |[p, q]_\gamma|$ defined on points $(p, q) \in \gamma \times \gamma$ such that p precedes q on γ. The restriction on L implies that $D(p_x, p_y) \geqslant 0$. We conclude that if $[p, q]_\gamma$ is a subsegment of γ that is maximal with respect inclusion among subsegments for which D takes non-positive values on the endpoints, then $Ld(p, q) = |[p, q]_\gamma|$. We consider several cases. Each carries the additional assumption that we are not in one of the previous cases.

Case 0: D is non-negative. Set $\gamma' := \gamma$, which is an $(L, 0)$–quasi-geodesics by definition of D.

Case 1: D takes a non-positive value on $[p_x, x]_\gamma \times [y, p_y]_\gamma$. In this case there exist points $x' \in [p_x, x]$ and $y' \in [p_y, y]$ such that the segment $[x', y']_\gamma$ is maximal with respect to inclusion among subsegments of γ with the property that D takes non-positive values on endpoints. Define γ' by replacing $[x', y']_\gamma$ by some geodesic segment with the same endpoints; $\gamma' := [p_x, x']_\gamma[x', y'][y', p_y]_\gamma$. We claim that γ' is an $(L, 0)$–quasi-geodesic. Since γ' is a concatenation of geodesic segments, it suffices to check that points on distinct segments are sufficiently far apart. We check distances between arbitrary points $x'' \in [p_x, x']_{\gamma'}$, $z \in [x', y']_{\gamma'}$, and $y'' \in [y', p_y]_{\gamma'}$.

Suppose, for contradiction, that $Ld(x'', y'') < |[x'', y'']_{\gamma'}|$. Since $[x', y']_\gamma$ has been replaced by a geodesic segment, $Ld(x'', y'') < |[x'', y'']_{\gamma'}| \leqslant |[x'', y'']_\gamma|$, so $D(x'', y'') < 0$. Since $D(x', y') = 0$ we have $x'' \in [p_x, x')_\gamma$ or $y'' \in (y', p_y]_\gamma$, but then $[x'', y'']_\gamma$ is a subsegment of γ strictly containing $[x', y']_\gamma$ such that D takes a non-positive value on its endpoints. This contradicts maximality of $[x', y']_\gamma$ among such subsegments, so $d(x'', y'') \geq \frac{|[x'',y'']_{\gamma'}|}{L}$.

Suppose, for contradiction, that $Ld(x'', z) < |[x'', z]_{\gamma'}|$. This implies $x'' \neq x'$, because x' and z lie on a geodesic subsegment of γ'. We estimate:

$$\begin{aligned} d(x'', y') &\leqslant d(x'', z) + d(z, y') \\ &< \frac{|[x'', z]_{\gamma'}|}{L} + d(z, y') \\ &= \frac{d(x'', x') + d(x', z)}{L} + d(x', y') - d(x', z) \\ &= \frac{|[x'', x']_{\gamma}|}{L} + \frac{|[x', y']_{\gamma}|}{L} - \left(\frac{L-1}{L}\right) d(x', z) \\ &\leqslant \frac{|[x'', y']_{\gamma}|}{L} \end{aligned}$$

Since $x'' \in [p_x, x')_\gamma$, we have exhibited a subsegment $[x'', y']_\gamma$ strictly containing $[x', y']_\gamma$ such that D takes a non-positive value on its endpoints. This contradicts maximality of $[x', y']_\gamma$ among such subsegments, so $d(x'', z) \geqslant \frac{|[x'', z]_{\gamma'}|}{L}$.

A symmetric argument shows $d(y'', z) \geqslant \frac{|[y'', z]_{\gamma'}|}{L}$, so γ' is an $(L, 0)$–quasi-geodesic.

Case 2: D takes a non-positive value on an element of $[p_x, x]_\gamma \times (x, y]_\gamma$. Let $[x', q_x]_\gamma$ be a subsegment of γ maximal with respect to inclusion among subsegments for which D takes non-positive values on endpoints, with $x' \in [p_x, x]_\gamma$. Since we are not in Case 1, $q_x \in (x, y)_\gamma$. Now consider whether or not $[q_x, p_y]_\gamma$ is an $(L, 0)$–quasi-geodesic. If so, define $\gamma' := [p_x, x']_\gamma [x', q_x][q_x, p_y]_\gamma$. Otherwise, D takes a negative value on an element of $[q_x, y)_\gamma \times (y, p_y]_\gamma$. Let $[q_y, y']_\gamma$ be a maximal subsegment of $[q_x, p_y]_\gamma$, with $q_y \in [q_x, y)_\gamma$ and $y' \in (y, p_y]$ on which D takes non-positive values on endpoints. We claim that $q_y \in (q_x, y)_\gamma$ and $D(q_y, y') = 0$, because if $D(q_y, y) < 0$ and $q_y \neq q_x$ then we can enlarge the subsegment, contradicting maximality, while if $q_y = q_x$ then $D(x', y') \leqslant 0$, contradicting the assumption that we are not in Case 1.

In either of these cases, we claim γ' is an $(L, 0)$–quasi-geodesic. This follows by verifying that the distance between points in distinct geodesic components of γ' have distance at least equal to the length of the subsegment of γ' they bound divided by L. The strategy is to suppose D attains a strictly negative value and then either derive a contradiction to maximality of $[x', q_x]_\gamma$ or $[q_y, y']_\gamma$ or to the assumption that we are not Case 1. The arguments are substantially similar to the computations in Case 1 and are left to the reader.

Case 3: D takes a non-positive value on an element of $[x, y)_\gamma \times [y, p_y]_\gamma$. The argument here is symmetric to the subcase of Case 2 in which only a corner at x is cut short.

We have shown how to produce an $(L, 0)$–quasi-geodesic γ' from γ. We now proceed to show $\rho(r) \leqslant \rho'(r)$ for any r such that $\rho(r) > 0$. Since $\rho(r) > 0$ there exist x and y such that $d(x, y) \leqslant d(x, Y) \leqslant r$ and $\operatorname{diam} \pi^\epsilon_Y(x) \cup \pi^\epsilon_Y(y) > \frac{2}{3}\rho(r)$. Choose $p_x \in \pi^\epsilon_Y(x)$, $p_y \in \pi^\epsilon_Y(y)$ such that $d(p_x, p_y) > \frac{2}{3}\rho(r)$.

Let $\gamma := [p_x, x][x, y][y, p_y]$. Let $L := \frac{12r+6\epsilon}{\rho(r)} \geqslant 2\frac{|\gamma|}{d(p_x, p_y)}$, and let γ' be the $(L, 0)$–quasi-geodesic produced from γ as above. By the Morse hypothesis, γ' is contained in the $\mu(L)$–neighborhood of Y.

Case a: γ' comes from Case 0 or Case 3. In this case $x \in \gamma'$, so $d(x, Y) \leqslant \mu(L)$, so $\rho(r) < \frac{3}{2} d(p_x, p_y) \leqslant \frac{3}{2}(4\mu(L) + 2\epsilon)$.

Case b: γ' comes from Case 1. In this case $p_x \in \pi^\epsilon_Y(x')$ and $p_y \in \pi^\epsilon_Y(y')$, so $d(x', p_x) \leqslant \mu(L) + \epsilon$ and $d(y', p_y) \leqslant \mu(L) + \epsilon$. Also, by definition of L we have:

$$d(x', y') = \frac{|[x', y']_\gamma|}{L} \leqslant \frac{|\gamma|}{L} \leqslant \frac{4r + 2\epsilon}{\frac{3(4r+2\epsilon)}{\rho(r)}} = \frac{\rho(r)}{3}$$

Since $d(p_x, p_y) > \frac{2}{3}\rho(r)$, we conclude $d(x', p_x) + d(y', p_y) > \frac{\rho(r)}{3}$, so that $\rho(r) < 6\mu(L) + 6\epsilon$.

Case c: γ' comes from Case 2. In this case γ' contains a geodesic segment from a point $x' \in [p_x, x]_\gamma$ to a point $q_x \in [x, y]_\gamma$. As in the previous case, $d(x', q_x) = \frac{|[x', q_x]_\gamma}{L} \leqslant \frac{|\gamma|}{L} \leqslant \frac{\rho(r)}{3}$. Consider a point $w \in \pi_Y^\epsilon(q_x)$. Since $d(x, y) \leqslant d(x, Y)$, we have $d(q_x, y) \leqslant d(q_x, Y) \leqslant \mu(L)$, which implies $d(w, p_y) \leqslant 4\mu(L) + 2\epsilon$. Thus $d(p_x, w) > \frac{2}{3}\rho(r) - (4\mu(L) + 2\epsilon)$. We also have $d(x', Y) \leqslant \mu(L)$ and $d(q_x, Y) \leqslant \mu(L)$, since both these points belong to γ', so:

$$\frac{\rho(r)}{3} \geqslant d(x', q_x) \geqslant d(p_x, w) - d(x', p_x) - d(q_x, w) > \frac{2}{3}\rho(r) - (6\mu(L) + 4\epsilon)$$

The resulting bound on $\rho(r)$ is the largest of the three cases, and establishes the bound of (2), completing the proof. □

5. Divergence

In this section we relate divergence to contraction and the Morse property, thereby proving Theorem 1.5.

There is a link between the Morse property and superlinear divergence via asymptotic cones [17]. Although this principle is well-known, there are competing definitions of 'superlinear' and 'divergence', so we give a detailed proof of Theorem 1.5 in terms of our definitions. Our analysis actually yields more. In the introduction we claimed that for a quasi-geodesic the Morse property, hence, sublinear contraction, is morally the opposite of high divergence. We prove a precise technical formulation of this claim in Proposition 5.5. Roughly speaking, the result we obtain is that if divergence of a quasi-geodesic γ is greater than a function f then almost closest point projection to γ is (r, f^{-1})–contracting.

Definition 5.1. Let X be a geodesic metric space and let $\gamma\colon \mathbb{R} \to X$ be an (L, A)–quasi-geodesic. Let $\lambda \in (0, 1]$, and let $\kappa \geqslant L + A$. Let $\Lambda_\gamma(r, s; L, A, \lambda, \kappa)$ be the infimal length of a path from $\gamma(s - r)$ to $\gamma(s + r)$ that is disjoint from the ball of radius $\lambda(L^{-1}r - A) - \kappa$ centered at $\gamma(s)$, or ∞ if no such path exists. The (L, A, λ, κ)*–divergence* of γ evaluated at r is $\Delta_\gamma(r; L, A, \lambda, \kappa) := \inf_s \Lambda_\gamma(r, s; L, A, \lambda, \kappa)$.

Notice that if γ is a geodesic, $\lambda := 1/2$, and $\kappa := 2$ we recover the definition of divergence we gave in the introduction.

We make the convention that $\infty \leqslant \infty$.

In light of the following lemma, γ has a well defined divergence, up to equivalence of functions, and we use $\Delta_\gamma(r)$ to denote the equivalence class of $\Delta_\gamma(r; L, A, \lambda, \kappa)$.

Lemma 5.2. *Let γ be an (L, A)–quasi-geodesic. Suppose γ is also an (L', A')–quasi-geodesic. Let $\lambda, \lambda' \in (0, 1]$, $\kappa \geqslant L + A$, and $\kappa' \geqslant L' + A'$. Then $\Delta_\gamma(r; L, A, \lambda, \kappa) \asymp \Delta_\gamma(r; L', A', \lambda', \kappa')$.*

Proof. Take $0 < M < 1$ small enough that $\frac{\lambda}{L} - \frac{\lambda'}{L'}M > 0$. Then for any sufficiently large $C \geqslant 0$ the affine function $\theta\colon\ r \mapsto Mr - C$ satisfies:

$$\lambda'((L')^{-1}\theta(r) - A') - 2\kappa' \leqslant \lambda(L^{-1}r - A) - \kappa$$

Fix $s \in \mathbb{R}$ and let P be any path from $\gamma(s - r)$ to $\gamma(s + r)$ that is disjoint from the ball of radius $\lambda(L^{-1}r - A) - \kappa$ centered at $\gamma(s)$. By the above inequality it is also disjoint from the ball of radius $\lambda'((L')^{-1}\theta(r) - A') - 2\kappa'$ about $\gamma(s)$.

Let $\{x_0, x_1, \ldots, x_l\}$ be the set $[s-r, s-\theta(r)] \cap (\mathbb{Z} \cup \{s-r, s-\theta(r)\})$ in descending order and let P_- be the path from $\gamma(s - \theta(r))$ to $\gamma(s - r)$ obtained by concatenating geodesics $[\gamma(x_i), \gamma(x_{i+1})]$. Define a path P_+ from $\gamma(s + r)$ to $\gamma(s + \theta(r))$ similarly. Since $\kappa' \geqslant L' + A'$, the paths P_- and P_+ are disjoint from the ball of radius $\lambda'((L')^{-1}\theta(r) - A') - \kappa'$ centered at $\gamma(s)$.

Define P' to be the path from $\gamma(s - \theta(r))$ to $\gamma(s + \theta(r))$ obtained by concatenating P_-, P, and P_+.

Now, for each r choose s and P so that $|P| \leqslant 1 + \Delta_\gamma(r; L, A, \lambda, \kappa)$. Then $\Delta_\gamma(\theta(r); L', A', \lambda', \kappa') \leqslant |P| + 2(L(r - \theta(r)) + A)$. Since γ is quasi-geodesic, $r \leqslant L|P| + LA$, so the right-hand side can be bounded by an affine function of $\Delta_\gamma(r; L, A, \lambda, \kappa)$. This proves one direction of the equivalence. The other follows immediately by reversing the roles in the above argument. □

We first give an example of the relationship between divergence and contraction.

Example 5.3. Let $f(r) \geqslant r$ be an increasing, invertible function. Consider the space X constructed in Example 3.4, but this time take $|I_i| := 2i$ and $|J_i| := f(i)$ for $i \in \mathbb{N}$. Let γ be a geodesic whose image is Y. Then $\Lambda_\gamma(i, \gamma^{-1}(y_i); 1, 0, 1, 1) = f(i)$, and this is optimal for radius i, so $\Delta_\gamma \asymp f$. On the other hand, the computation of Example 3.4 shows that $\operatorname{diam} \pi_Y^0(x_i) = 2f^{-1}(4r)$. Thus, π_Y^0 is sublinearly contracting if and only if f^{-1} is sublinear, and, in this case, it is (r, ρ)–contracting for $\rho \asymp f^{-1}$.

Our next proposition proves the implication (2) $\implies$ (1) of Theorem 1.5. It also gives a quantitave link between high divergence and contraction.

Definition 5.4. We say a function g is *completely super–f* if for every choice of $C_1 > 0$, $C_2 > 0$, $C_3 \geqslant 0$, and $C_4 \geqslant 0$ the collection of $r \in [0, \infty)$ such that $g(r) \leqslant C_1 f(C_2 r + C_3) + C_4$ is bounded.

Proposition 5.5. *Let γ be a quasi-geodesic in a geodesic metric space X. Suppose the empty set is not in the image of π_γ^ϵ. Let $f(r) \geqslant r$ be an increasing, invertible function. If γ has completely super–f divergence, then there exists a function ρ such that π_γ^ϵ is (r, ρ)–contracting and* $\lim_{r\to\infty} \frac{\rho(r)}{f^{-1}(r)} = 0$.

In particular, if γ has completely superlinear divergence then there exists a sublinear function ρ such that π_γ^ϵ is (r, ρ)–contracting.

Proof. Let γ be an (L, A)-quasi-geodesic. Define:

$$\rho(r) := \sup_{d(x,y)\leqslant d(x,\gamma)\leqslant r} \operatorname{diam} \pi_\gamma^\epsilon(x) \cup \pi_\gamma^\epsilon(y)$$

To see that π_γ^ϵ is (r, ρ)–contracting we must show that ρ is sublinear. Since $f(r) \geqslant r$, it suffices to prove the second claim:

$$\lim_{r\to\infty} \frac{\rho(r)}{f^{-1}(r)} = 0$$

Suppose for a contradiction that $\limsup_{r\to\infty} \frac{\rho(r)}{f^{-1}(r)} > 0$. Then there exist $c > 0$; sequences (x_n) and (y_n) with $x_n, y_n \in X$, $d(x_n, \gamma) \geqslant n$, and $d(x_n, y_n) \leqslant d(x_n, \gamma)$; and $x'_n \in \pi_\gamma^\epsilon(x_n)$ and $y'_n \in \pi_\gamma^\epsilon(y_n)$ such that:

$$cf^{-1}(d(x_n, \gamma)) \leqslant d(x'_n, y'_n) \tag{3}$$

Let a_n and b_n be such that $\gamma(a_n - b_n) = x'_n$ and $\gamma(a_n + b_n) = y'_n$. Define $m_n := \gamma(a_n)$ and $R_n := \frac{b_n}{L} - A$. Since γ is an (L, A)–quasi-geodesic, $d(m_n, \{x'_n, y'_n\}) \geqslant R_n$ and $b_n \geqslant \frac{d(x'_n,y'_n)-A}{2L}$. By (3) and the facts that f^{-1} is unbounded and increasing, $\lim_{n\to\infty} R_n = \infty$.

Choose $0 < \lambda < \frac{1}{4}$ and $\kappa := L + A$.

If there is a geodesic from x_n to y_n containing a point z such that $d(z, m_n) \leqslant \lambda R_n$, then:

$$\begin{aligned}
R_n &\leqslant d(y'_n, m_n) \\
&\leqslant d(y'_n, y_n) + d(y_n, z) + d(z, m_n) \\
&\leqslant d(y_n, \gamma) + \epsilon + d(y_n, z) + d(z, m_n) \\
&\leqslant \epsilon + 2(d(y_n, z) + d(z, m_n)) \\
&\leqslant \epsilon + 2\lambda R_n + 2d(y_n, z) \\
&= \epsilon + 2\lambda R_n + 2(d(x_n, y_n) - d(z, x_n)) \\
&\leqslant \epsilon + 2\lambda R_n + 2(d(x_n, \gamma) - (d(x_n, \gamma) - \lambda R_n)) \\
&= \epsilon + 4\lambda R_n
\end{aligned}$$

Thus, $R_n \leqslant \frac{\epsilon}{1-4\lambda}$.

If there is a geodesic from x_n to x'_n or from y_n to y'_n containing a point z such that $d(z, m_n) \leqslant \lambda R_n$, then a similar argument shows $R_n \leqslant \frac{\epsilon}{1-2\lambda}$.

Since $R_n \to \infty$, for all sufficiently large n and any choice of path p_n that is a concatenation of geodesics $[x'_n, x_n]$, $[x_n, y_n]$, $[y_n, y'_n]$, the path p_n remains outside the ball of radius λR_n about

m_n. This gives us a path of length at most $4d(x_n, \gamma) + 2\epsilon$ from $\gamma(a_n - b_n)$ to $\gamma(a_n + b_n)$ that remains outside the ball of radius $\lambda\left(\frac{b_n}{L} - A\right)$ about $\gamma(a_n)$.

On the other hand, (3) implies:

$$d(x_n, \gamma) \leqslant f\left(\frac{1}{c} d(x'_n, y'_n)\right) \leqslant f\left(\frac{2b_nL + A}{c}\right)$$

We conclude that for all sufficiently large n the (L, A, λ, κ)–divergence of γ evaluated at b_n is at most $2\epsilon + 4f\left(\frac{2b_nL+A}{c}\right)$, which contradicts the hypothesis that the divergence is completely super–f. □

The previous result can be strengthened to the statement:

Proposition 5.6. *Let f be an increasing, invertible, completely superlinear function satisfying the following additional condition:*

(∗) *For every C there exists some D such that for all $r > 1$ and $k > D$ we have $f(kr) > Cf(Cr + C) + C$.*

If the divergence of γ is at least f then γ is (r, ρ)–contracting for some function $\rho \preceq f^{-1}$.

Proof. For a contradiction we suppose that $\rho \npreceq f^{-1}$ and replace (3) with $d(x'_n, y'_n) \geqslant nf^{-1}(d(x_n, \gamma))$. Using the same method as in the proof of Proposition 5.5, we deduce that for all sufficiently large n the (L, A, λ, κ)–divergence of γ evaluated at b_n is at most $2\epsilon + 4f\left(\frac{2b_nL+A}{n}\right)$. Thus, $f(b_n) \leqslant 2\epsilon + 4f\left(\frac{2b_nL+A}{n}\right)$. Let $c_n := b_n/n$ and $M := \max\{2\epsilon, 4, 2L, A\}$. Then, since f is increasing:

$$f(nc_n) \leqslant Mf(Mc_n + M) + M \tag{4}$$

The left-hand side is unbounded as n grows, so we immediately obtain a contradiction if the sequence $(c_n)_{n\in\mathbb{N}}$ is bounded. If the sequence is unbounded then, by passing to a subsequence, we may assume $c_n > 1$ for all n. In this case the inequality (4) holds for all n, which contradicts condition (∗). □

Suitable functions f for Proposition 5.6 include $f(r) := r^d$, $r^d/\log(r)$, $r\log(r)$ and d^r for any $d > 1$. The function $f(r) := 2^{2^{2^{1+\lfloor \log_2 \log_2 r \rfloor}}}$ is completely superlinear, but does not satisfy (∗), since $f(n2^{2^{n-1}}) = f(2^{2^{n-1}})$ for all $n \in \mathbb{N}$.

Corollary 5.7. *If a quasi-geodesic γ has divergence at least r^k then γ is $(r, r^{1/k})$–contracting. If it has exponential divergence, then γ is logarithmically contracting. Finally, if it has infinite divergence, then it is strongly contracting.*

Here infinite divergence means $\Delta_\gamma(r) = \infty$ for all r large enough. Example 5.3 shows these conclusions are optimal.

We now address the implication (1) $\implies$ (2) of Theorem 1.5. In this direction we can show that the Morse property implies completely superlinear divergence, but we do not get explicit control of the divergence function in terms of the Morse function, see Proposition 5.10.

There is one special case in which we can say more. Charney and Sultan [13] recently gave a proof[5] that if α is a Morse geodesic in a CAT(0) space then α has at least quadratic divergence. Essentially the same argument gives a general result:

Proposition 5.8. *Let α be a geodesic in a geodesic metric space X. If α is (ρ_1, ρ_2)–contracting with ρ_2 bounded, then $\Delta_\alpha(r) \succeq r\rho_1(r)$.*

Lemma 5.9. *Let X be a geodesic metric space. Let $a, b, c, d \in X$ and $r > 0$ satisfy the following conditions:*

(1) $d(a, d) \geqslant r$

[5]The original proof of this fact is due to Behrstock and Druţu [6], by different methods.

(2) *There exists a path* γ *from* a *to* d *passing through* b *and* c *such that the length of* γ *is at most* Cr *and such that* $[a,b]_\gamma$, $[b,c]_\gamma$, *and* $[c,d]_\gamma$ *are continuous* $(L,0)$*–quasi-geodesics.*

(3) *The path* γ *does not contain a point within distance* λr *of* e*, where* e *is the midpoint of a geodesic from* a *to* d*.*

For any $L' > \max\{L, C, C/\lambda\} \geqslant 1$ *there exists a continuous* $(L',0)$*–quasi-geodesic* γ' *from* a *to* d *of length at most* $|\gamma|$ *such that* γ' *does not contain a point within distance* $\lambda r/2$ *of* e*.*

Proof. The construction of γ' is exactly as in Proposition 4.2 with L replaced by L'. This involves finding points p and q on γ such that $L'd(p,q) = |[p,q]_\gamma|$ and replacing $[p,q]_\gamma$ by a geodesic with the same endpoints. Now, $d(p,q) \leqslant |\gamma|/L' < \lambda r$, so for any point z on a newly introduced geodesic segment we have $d(z,e) \geqslant d(\gamma,e) - d(p,q)/2 > \lambda r/2$. □

Proposition 5.10. *Let* γ *be a Morse quasi-geodesic in a geodesic metric space* X*. Then the divergence of* γ *is completely superlinear.*

Proof. We prove the contrapositive. Let γ be an (L,A)–quasi-geodesic and suppose its divergence is not completely superlinear. Then there exists $C > 0$ for which there exists an unbounded sequence of numbers $r_n \geqslant 1$ and paths p_n such that:

(1) There exists a sequence of real numbers s_n such that the endpoints of p_n are $x_n = \gamma(s_n - r_n)$ and $y_n = \gamma(s_n + r_n)$.

(2) $|p_n| \leqslant Cr_n$.

(3) p_n does not intersect the $(\frac{r_n}{2L} - A)$–neighborhood of $\gamma(s_n)$.

We may assume all $r_n \geqslant 4AL$ so point (3) can be replaced by:

3′. p_n does not intersect the $(\frac{r_n}{4L})$–neighborhood of $m_n := \gamma(s_n)$.

Our goal is to construct uniform quasi-geodesics γ_n from x_n to y_n that avoid increasingly large balls around m_n.

Set $x_{n,0} := x_n$ and define $x_{n,1}$ to be the last point on p_n for which we have $d(x_{n,0}, x_{n,1}) = r_n/8L$.

Similarly define $x_{n,i}$ to be y_n if $d(x_{n,i-1}, y_n) < r_n/4L$ or to be the last point on p_n satisfying $d(x_{n,i-1}, x_{n,i}) = r_n/8L$ otherwise.

Note that $y_n = x_{n,k_n}$ for some $k_n \leqslant 8CL$. By construction, if $i \neq j$ then $d(x_{n,i}, x_{n,j}) \geqslant r_n/8L$.

Let γ_n^1 be a concatenation of geodesics $[x_{n,0}, x_{n,1}] \dots [x_{n,k_n-1}, y_n]$. We have that $|\gamma_n^1| \leqslant Cr_n$ and $d(\gamma_n^1, m_n) > r_n/8L$.

Applying Lemma 5.9 for each $1 \leqslant i \leqslant \lfloor k_n/3 \rfloor$ there are $(L_2,0)$–quasi-geodesics (where L_2 does not depend on n) from $x_{n,3(i-1)}$ to $x_{n,3i}$ such that the concatenation γ_n^2 of these with $[x_{n,3\lfloor k_n/3\rfloor}, y_n]_{\gamma_n^1}$ satisfies $d(\gamma_n^2, m_n) > r_n/16L$.

Repeating this procedure at most $d = \lceil \log_3 8CL \rceil$ times we obtain an $(L_d,0)$ quasi-geodesic γ_n^d from x_n to y_n satisfying $d(\gamma_n^d, m_n) > r_n/(2^{d+2}L)$. Again, L_d does not depend on n.

If γ is μ–Morse, then the γ_n^d are μ'–Morse for some μ' that does not depend on n. Then $d(\gamma_n^d, m_n) \leqslant \mu'(K,C)$, which is bounded, contradicting the lower bound above. □

A finitely generated group is called *constricted* if all of its asymptotic cones have cut points [18].

Corollary 5.11. *Suppose there exists a quasi-geodesic* γ *with completely superlinear divergence in a geodesic metric space* X*. In every asymptotic cone of* X *every point of the ultralimit of* γ *is a cut point.*

In particular, a finitely generated group is constricted if one of its Cayley graphs contains a quasi-geodesic with completely superlinear divergence.

Olshanskii, Osin, and Sapir [26, Corollary 6.4] build a group that has an asymptotic cone with no cut point such that the group has a Cayley graph with geodesics of superlinear divergence. These geodesics are therefore not Morse. They explicitly state that their construction yields geodesics that are not completely superlinear. Corollary 5.11 shows that this will be the case in any such construction.

6. Robustness

Suppose that π_Y^ϵ is (ρ_1, ρ_2)–contracting. In this section we investigate the extent to which ρ_2 is affected by changes to ρ_1, ϵ, or Y.

Clearly π_Y^ϵ is (ρ_1', ρ_2)–contracting for $\rho_1' \leqslant \rho_1$. From Theorem 1.4 we know that π_Y^ϵ is (r, ρ_2')–contracting for some ρ_2' depending on ρ_1 and ρ_2. For this ρ_2', it follows that π_Y^ϵ is (ρ_1', ρ_2')–contracting for every $\rho_1 \leqslant \rho_1' \leqslant r$.

In general ρ_2 and ρ_2' are not asymptotic. For example, if π_Y^ϵ is $(r/2, B_1)$–contracting it is (r, ρ_2)–contracting for $\rho_2 \asymp \log_2$, as in Proposition 3.3, but not necessarily (r, B_2)–contracting for some constant B_2, by Example 3.2. One well-known special case is that $(r/M, B_1)$–contracting for $M > 1$ and B_1 bounded implies $(r/2, B_2)$–contracting for some bounded B_2, see, eg, [30].

The output contraction functions are asymptotic when the input function is changed by an additive constant:

Lemma 6.1. *If π_Y^ϵ is (ρ_1, ρ_2)–contracting for $\rho_1(r) = \rho_1'(r) - C$, with $\rho_1'(r) \leqslant r$ and $C \geqslant 0$, then π_Y^ϵ is (ρ_1', ρ_2')–contracting for some $\rho_2' \asymp \rho_2$.*

Proof. Let $C' := \sup\{r \mid \rho_1(r) \leqslant C\}$. Suppose that x and y are points with $d(x,y) \leqslant \rho_1'(d(x,Y))$. If $d(x,y) \leqslant \rho_1(d(x,Y)) = \rho_1'(d(x,Y)) - C$ then we have $\operatorname{diam} \pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(y) \leqslant \rho_2(d(x,Y))$. Otherwise, let z be a point on a geodesic from x to y such that $d(x,z) = \rho_1(d(x,Y))$. This implies $d(y,z) \leqslant C$. Now:

$$\begin{aligned}\operatorname{diam} \pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(y) &\leqslant \operatorname{diam} \pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(z) + \operatorname{diam} \pi_Y^\epsilon(z) \cup \pi_Y^\epsilon(y)\\ &\leqslant \rho_2(d(x,Y)) + \operatorname{diam} \pi_Y^\epsilon(z) \cup \pi_Y^\epsilon(y)\end{aligned}$$

If $d(z,y) > \rho_1(d(z,Y))$ then $d(z,Y) \leqslant C'$, so $\operatorname{diam} \pi_Y^\epsilon(z) \cup \pi_Y^\epsilon(y) \leqslant 2(C + C' + \epsilon)$. If $d(z,y) \leqslant \rho_1(d(z,Y))$ then $\operatorname{diam} \pi_Y^\epsilon(z) \cup \pi_Y^\epsilon(y) \leqslant \rho_2(d(z,Y)) \leqslant \rho_2(2d(x,Y))$. Combining these cases, we see that $d(x,y) \leqslant \rho_1(d(x,Y))$ implies:

$$\operatorname{diam} \pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(y) \leqslant \rho_2(d(x,Y)) + \rho_2(2d(x,Y)) + 2(C + C' + \epsilon)$$

Thus, it suffices to take $\rho_2'(r) := 2\rho_2(2r) + 2(C + C' + \epsilon)$. □

Next, consider changes to the projection parameter.

Lemma 6.2. *Suppose ϵ_0 and ϵ_1 are constants such that the empty set is neither in the image of $\pi_Y^{\epsilon_0} \colon X \to 2^Y$ nor in the image of $\pi_Y^{\epsilon_1} \colon X \to 2^Y$. If $\pi_Y^{\epsilon_0}$ is (ρ_1, ρ_2)–contracting then there exist ρ_1' and ρ_2' such that $\pi_Y^{\epsilon_1}$ is (ρ_1', ρ_2')–contracting. If $\epsilon_1 \leqslant \epsilon_0$ or if $\rho_1(r) := r$ then we can take $\rho_1' = \rho_1$ and $\rho_2' \asymp \rho_2$.*

Proof. When $\epsilon_1 \leqslant \epsilon_0$ we have $\pi_Y^{\epsilon_1}(x) \subset \pi_Y^{\epsilon_0}(x)$, so the result is clear. In this case $\rho_1' = \rho_1$ and $\rho_2' = \rho_2$ will suffice.

The fact that $\pi_Y^{\epsilon_1}$ is sublinearly contracting follows from Theorem 1.4, since Y is Morse. It remains only to prove the asymptotic statement in the case that $\rho_1(r) := r$, so suppose $\pi_Y^{\epsilon_0}$ is (r, ρ_2)–contracting.

For any $x \in X \setminus Y$ and each $i \in \{0,1\}$, consider a point $x_i \in \pi_Y^{\epsilon_i}(x)$ and a point z_i on a geodesic from x to x_i with $d(x, z_i) = d(x,Y)$. Then:

$$\begin{aligned}d(x_0, x_1) &\leqslant d(x_0, z_0) + d(z_0, \pi_Y^{\epsilon_0}(z_0)) + \operatorname{diam} \pi_Y^{\epsilon_0}(z_0) \cup \pi_Y^{\epsilon_0}(x)\\ &\quad + \operatorname{diam} \pi_Y^{\epsilon_0}(x) \cup \pi_Y^{\epsilon_0}(z_1) + d(\pi_Y^{\epsilon_0}(z_1), z_1) + d(z_1, x_1)\\ &\leqslant \epsilon_0 + 2\epsilon_0 + \rho_2(d(x,Y)) + \rho_2(d(x,Y)) + \epsilon_0 + \epsilon_1 + \epsilon_1\\ &= 4\epsilon_0 + 2\epsilon_1 + 2\rho_2(d(x,Y))\end{aligned}$$

If $d(x,y) \leqslant d(x,Y)$ then:

$$\begin{aligned}\operatorname{diam} \pi_Y^{\epsilon_1}(x) \cup \pi_Y^{\epsilon_1}(y) &\leqslant \operatorname{diam} \pi_Y^{\epsilon_1}(x) \cup \pi_Y^{\epsilon_0}(x) + \operatorname{diam} \pi_Y^{\epsilon_0}(x) \cup \pi_Y^{\epsilon_0}(y)\\ &\quad + \operatorname{diam} \pi_Y^{\epsilon_0}(y) \cup \pi_Y^{\epsilon_1}(y)\\ &\leqslant 4\epsilon_0 + 2\epsilon_1 + 2\rho_2(d(x,Y)) + \rho_2(d(x,Y))\\ &\quad + 4\epsilon_0 + 2\epsilon_1 + 2\rho_2(d(y,Y))\end{aligned}$$

Since $d(y,Y) \leqslant 2d(x,Y)$, this means that $\pi_Y^{\epsilon_1}$ is (r,ρ_2')–contracting for:

$$\rho_2'(r) := 8\epsilon_0 + 4\epsilon_1 + 3\rho_2(r) + 2\rho_2(2r) \asymp \rho_2(r) \qquad \square$$

Finally, consider changes to the target of the projection map.

Lemma 6.3. *Let Y and Y' be subspaces of a geodesic metric space X at bounded Hausdorff distance from one another. Suppose that π_Y^ϵ is (ρ_1,ρ_2)–contracting. Then $\pi_{Y'}^\epsilon$ is (r,ρ_2')–contracting for some ρ_2'. If $\rho_1(r) = r$ then we can take $\rho_2' \asymp \rho_2$.*

Proof. Let C be the Hausdorff distance between Y and Y'.

For every $x \in X$ we have $\pi_{Y'}^\epsilon(x) \subset \overline{N}_C(\pi_Y^{\epsilon+2C}(x))$. The result now follows easily from Lemma 6.2. $\square$

In light of Lemma 6.2, we can speak of the set Y being a contracting set if some ϵ–closest point projection to Y is contracting.

Definition 6.4. We say *Y is (ρ_1,ρ_2)–contracting* if there exists an $\epsilon \geqslant 0$ such that the ϵ–closest point projection $\pi_Y^\epsilon\colon X \to 2^Y$ is (ρ_1,ρ_2)–contracting.

Equivalently, *Y is (ρ_1,ρ_2)–contracting* if for all sufficiently small $\epsilon \geqslant 0$, if π_Y^ϵ does not have the empty set in its image, then π_Y^ϵ is (ρ_1,ρ_2)–contracting.

7. Geodesic image theorem

In this section we give an additional characterization of sublinear contraction in terms of projections of geodesic segments.

Theorem 7.1. *Let Y be a subspace of a geodesic metric space X. Suppose the empty set is not in the image of π_Y^ϵ. The following are equivalent:*

(1) There exist a sublinear function ρ and a constant $C \geqslant 0$ such that for every geodesic segment $\gamma \subset X$, with endpoints denoted x and y, if $d(\gamma,Y) \geqslant C$ then $\operatorname{diam}\pi_Y^\epsilon(\gamma) \leqslant \rho(\max\{d(x,Y),d(y,Y)\})$.

(2) There exist a sublinear function ρ' and a constant $C' \geqslant 0$ such that for every geodesic segment $\gamma \subset X$, if $d(\gamma,Y) \geqslant C'$ then $\operatorname{diam}\pi_Y^\epsilon(\gamma) \leqslant \rho'(\max_{z\in\gamma} d(z,Y))$.

(3) There exists a sublinear function ρ'' such that π_Y^ϵ is (r,ρ'')–contracting.

Moreover, $\rho \asymp \rho' \asymp \rho''$.

See Figure 3, letting γ be a subsegment of $\cup_i W_i$.

The case that Y is strongly contracting, that is, ρ'' is bounded, recovers the well-known 'Bounded Geodesic Image Property', cf [23, 8].

Corollary 7.2. *If Y is strongly contracting, $R_2 \geqslant 1$ is a constant greater than twice the bound on the contraction function for Y, and γ is a geodesic segment that does not enter the R_2–neighborhood of Y then* $\operatorname{diam}\pi_Y^\epsilon(\gamma)$ *is bounded, with bound depending only on ϵ and ρ''.*

Alternatively, one could read Theorem 7.1 as saying that if π_Y^ϵ is sublinearly contracting and γ is a geodesic ray that is far from Y, but such that $\pi_Y^\epsilon(\gamma)$ is large, then $d(\gamma(t),Y)$ grows superlinearly with respect to $\operatorname{diam}\pi_Y^\epsilon(\gamma([0,t]))$.

Proof of Theorem 7.1.
(1) $\implies$ (3): Define $\rho_1(r) := r - C$ and $\rho_2(r) = \rho(2r - C)$. By Lemma 6.1, it suffices to show that π_Y^ϵ is (ρ_1,ρ_2)–contracting.

Suppose x and y are points of X with $d(x,y) \leqslant \rho_1(d(x,Y))$, and let γ be a geodesic from x to y. Then γ remains outside the C–neighborhood of Y, by the definition of ρ_1, so:

$$\begin{aligned}\operatorname{diam}\pi_Y^\epsilon(x)\cup\pi_Y^\epsilon(y) &\leqslant \operatorname{diam}\pi_Y^\epsilon(\gamma)\\ &\leqslant \rho(\max\{d(x,Y),d(y,Y)\})\\ &\leqslant \rho(2d(x,Y)-C) = \rho_2(d(x,Y))\end{aligned}$$

This proves (1) $\implies$ (3), and a similar argument proves (2) $\implies$ (3).

Now assume (3). If $d(x, y) \leqslant d(x, Y) + d(y, Y)$ then both (1) and (2) follow easily, so assume not. Let z_0 be the point of γ at distance $d(x, Y)$ from x. Our assumption says $d(z_0, y) > d(y, Y)$. Define points z_{i+1} inductively as follows: if $d(z_i, y) > d(y, Y) + d(z_i, Y)$ define z_{i+1} to be the point of γ between z_i and y at distance $d(z_i, Y)$ from z_i. Let k be the last index so defined. From these choices we estimate:

$$\begin{aligned} \operatorname{diam} \pi_Y^\epsilon(\gamma) &\leqslant \operatorname{diam} \pi_Y^\epsilon(\overline{N}_{d(x,Y)}(x)) + \sum_{i=0}^{k} \operatorname{diam} \pi_Y^\epsilon(\overline{N}_{d(z_i,Y)}(z_i)) \\ &\quad + \operatorname{diam} \pi_Y^\epsilon(\overline{N}_{d(y,Y)}(y)) \\ &\leqslant 2\left(\rho''(d(x,Y)) + \sum_{i=0}^{k} \rho''(d(z_i,Y)) + \rho''(d(y,Y))\right) \end{aligned} \tag{5}$$

Since γ is a geodesic:

$$\begin{aligned} d(x,y) &= d(x, z_0) + \sum_{i=0}^{k-1} d(z_i, z_{i+1}) + d(z_k, y) \\ &= d(x, Y) + \sum_{i=0}^{k-1} d(z_i, Y) + d(z_k, y) \end{aligned} \tag{6}$$

We can also bound $d(x, y)$ in terms of the projections to Y:

$$\begin{aligned} d(x,y) &\leqslant d(x, \pi_Y^\epsilon(x)) + \operatorname{diam} \pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(y) + d(\pi_Y^\epsilon(y), y) \\ &\leqslant d(x, \pi_Y^\epsilon(x)) + \operatorname{diam} \pi_Y^\epsilon(x) \cup \pi_Y^\epsilon(z_0) + \sum_{i=0}^{k-1} \operatorname{diam} \pi_Y^\epsilon(z_i) \cup \pi_Y^\epsilon(z_{i+1}) \\ &\quad + \operatorname{diam} \pi_Y^\epsilon(z_k) \cup \pi_Y^\epsilon(y) + d(\pi_Y^\epsilon(y), y) \\ &\leqslant d(x, Y) + \epsilon + \rho''(d(x,Y)) + \sum_{i=0}^{k-1} \rho''(d(z_i, Y)) \\ &\quad + \rho''(d(z_k, Y)) + \rho''(d(y,Y)) + d(y, Y) + \epsilon \end{aligned} \tag{7}$$

Combining (6) and (7) gives us the estimate:

$$\begin{aligned} &\sum_{i=0}^{k-1} d(z_i, Y) - \rho''(d(z_i, Y)) \leqslant \\ &\qquad 2\epsilon + \rho''(d(x,Y)) + \rho''(d(z_k,Y)) + \rho''(d(y,Y)) + d(y,Y) - d(z_k, y) \end{aligned} \tag{8}$$

Define $R_n \geqslant 0$ such that for all $r \geqslant R_n$ we have $0 \leqslant \rho''(r) \leqslant r/n$. Suppose that $d(\gamma, Y) \geqslant R_2$ so that $d(z_i, Y) - \rho''(d(z_i, Y)) \geqslant \rho''(d(z_i, Y))$ for all i. These bounds, along with (8), (5), and $E := d(z_k, y) - d(y, Y)$ give:

$$\operatorname{diam} \pi_Y^\epsilon(\gamma) \leqslant 2\left(2(\epsilon + \rho''(d(x,Y)) + \rho''(d(z_k,Y)) + \rho''(d(y,Y))) - E\right)$$

By construction, $E > 0$, so to prove (2) it suffices to take $C' := R_2$ and $\rho(r) := 4\epsilon + 12\rho''(r)$.

To prove (1) we suppose $d(\gamma, Y) \geqslant C := R_4 \geqslant R_2$ and bound $2\rho''(d(z_k, Y)) - E$ in terms of $\rho''(d(y, Y))$. There are two cases to consider. If $d(z_k, Y) \leqslant 4d(y, Y)$ then $2\rho''(d(z_k, Y)) - E \leqslant 2\rho''(4d(y, Y))$. Otherwise, $d(z_k, Y) > 4d(y, Y)$ implies $E > d(z_k, Y)/2$, so:

$$2\rho''(d(z_k, Y)) - E < 2\frac{d(z_k, Y)}{4} - \frac{d(z_k, Y)}{2} = 0$$

Thus, it suffices to take $\rho'(r) := 4\epsilon + 12\rho''(4r)$. □

8. Further applications

First, we prove a general result.

PROPOSITION 8.1. *Let X be a geodesic metric space. Suppose subspaces Y and Y' of X are μ–Morse. Let $\epsilon \geqslant 0$ be a constant such that there exist points $p \in Y$ and $p' \in Y'$ such that $d(p, p') \leqslant d(Y, Y') + \epsilon$. Then there exist a constant B and a sublinear function ρ, each depending only on μ and ϵ, satisfying the following conditions:*

- *If $d(Y, Y') \leqslant 2\mu(4, 0)$ then $Y \cup Y'$ is B–quasi-convex.*
- *If $d(Y, Y') > 2\mu(4, 0)$ then for every geodesic α from Y to Y' with $|\alpha| \leqslant d(Y, Y') + \epsilon$ and every geodesic γ from Y to Y' we have $d(\alpha, \gamma) < \rho(d(Y, Y'))$.*

PROOF. Take geodesics α and γ as hypothesized. Let β be a geodesic from α to γ with $|\beta| = d(\alpha, \gamma)$. See Figure 5. Let $\delta := [p, x]_\alpha \beta [y, q]_\gamma$ and $\delta' := [p', x]_\alpha \beta [y, q']_\gamma$. (Recall that $[p, x]_\alpha$

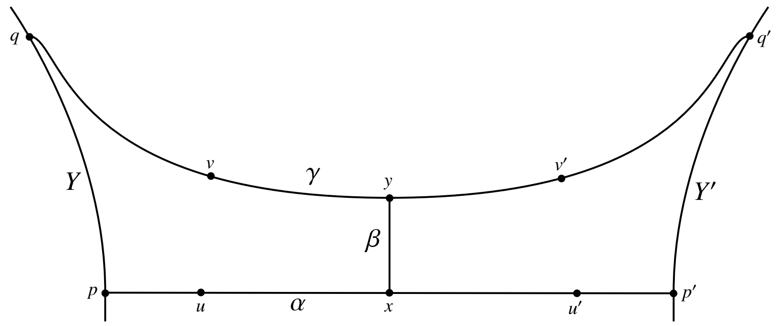

FIGURE 5. Setup for Proposition 8.1

denotes the subsegment of α from p to x.) Suppose that δ fails to be a $(k, 0)$–quasi-geodesic for some $k > 3$. Both $[p, x]_\alpha \beta$ and $\beta[y, q]_\gamma$ are $(3, 0)$–quasi-geodesics, by minimality of $d(x, y)$, so there exist points $u \in [p, x]_\alpha$ and $v \in [y, q]_\gamma$ such that $kd(u, v) < d(u, x) + d(x, y) + d(y, v)$. Now, $d(v, y) \leqslant d(v, u) + d(u, x) + d(x, y)$, so:

$$(k - 1)d(x, y) \leqslant (k - 1)d(u, v) < 2(d(u, x) + d(x, y))$$

Whence:

$$d(\alpha, \gamma) = d(x, y) < \frac{2d(u, x)}{k - 3} \leqslant \frac{2|\alpha|}{k - 3} \leqslant \frac{2(d(Y, Y') + \epsilon)}{k - 3} \tag{9}$$

If $d(Y, Y') \leqslant 2\mu(4, 0)$ and δ is not a $(4, 0)$–quasi-geodesic then $d(\alpha, \gamma) < 4\mu(4, 0) + 2\epsilon$, by (9). This means $[y, q]_\gamma$ is a geodesic with one endpoint on Y and one within distance $6\mu(4, 0) + 2\epsilon$ of Y. Since Y is μ–Morse there is a B_0 depending on μ such that such a geodesic segment is contained in the B_0–neighborhood of Y.

If δ is a $(4, 0)$–quasi-geodesic it is contained in the $\mu(4, 0)$–neighborhood of Y.

The same arguments apply for δ', and $\gamma \subset \delta \cup \delta'$, so if $d(Y, Y') \leqslant 2\mu(4, 0)$ then $Y \cup Y'$ is B–quasi-convex for $B := \max\{B_0, \mu(4, 0)\}$.

Now suppose $d(Y, Y') > 2\mu(4, 0)$. Then δ and δ' cannot both be $(4, 0)$–quasi-geodesics. By (9):

$$\begin{aligned} d(\alpha, \gamma) &< \frac{2(d(Y, Y') + \epsilon)}{\sup\{k \in \mathbb{R} \mid \delta \text{ or } \delta' \text{ is not a } (k, 0)\text{–quasi-geodesic}\} - 3} \\ &\leqslant \frac{2(d(Y, Y') + \epsilon)}{\sup\{k \in \mathbb{R} \mid d(Y, Y') > 2\mu(k, 0)\} - 3} \end{aligned}$$

Define:

$$\rho(r) := \frac{2(r + \epsilon)}{\sup\{k \in \mathbb{R} \mid r > 2\mu(k, 0)\} - 3}$$

We interpret $\rho(r)$ to be 0 if $\{2\mu(k, 0)\}_{k \in \mathbb{R}}$ is bounded above by r. For $r \geqslant \epsilon$ we have:

$$\frac{\rho(r)}{r} \leqslant \frac{4}{\sup\{k \in \mathbb{R} \mid r > 2\mu(k, 0)\} - 3}$$

The denominator is unbounded and non-decreasing as a function of r, so we have $\lim_{r\to\infty} \frac{\rho(r)}{r} = 0$. ◻

We first give an application of the second part of Proposition 8.1.

Proposition 8.2. *Let X be a geodesic metric space and let Y and Y' be μ–Morse subspaces of X. Let $\epsilon \geqslant 0$ be a constant such that the image of π^{ϵ}_{Y} does not contain the empty set, and such that there exist points $p \in Y$ and $p' \in Y'$ such that $d(p, p') \leqslant d(Y, Y') + \epsilon$.*

Suppose $d(Y, Y') > 2\mu(6, 0)$. Then there is a sublinear function ρ depending only on μ such that $\operatorname{diam} \pi^{\epsilon}_{Y}(Y') \leqslant \rho(d(Y, Y'))$.

Proof. Since Y is μ–Morse, there is a sublinear function ρ' depending only on μ such that Y is (r, ρ')–contracting, by Proposition 4.2.

Note that $p \in \pi^{\epsilon}_{Y}(p')$. Choose $q' \in Y'$ and $q \in \pi^{\epsilon}_{Y}(q')$. Let γ be a geodesic from q to q', let α be a geodesic from p to p', and let $x \in \alpha$ and $y \in \gamma$ be points such that $d(x, y) = d(\alpha, \gamma)$. The setup is the same as in Proposition 8.1, and we make the corresponding definitions of δ, δ', etc.

Suppose δ' is not a $(5, 0)$–quasi-geodesic. Define u' and v' as in Proposition 8.1, so that $d(u', x) + d(x, y) + d(y, v') > 5d(u', v')$. We have $p \in \pi^{\epsilon}_{Y}(u')$ and $q \in \pi^{\epsilon}_{Y}(v')$. By definition of x and y, we know $d(x, y) \leqslant d(u', v')$, so $d(u', x) + d(y, v') > 4d(u', v')$. In particular, we have $2d(u', v') < d(u', x)$ or $2d(u', v') < d(v', y)$. We suppose the former, the other case being similar.

First, suppose that $d(u', Y) < \epsilon$. Then:

$$\begin{aligned} d(p, q) &\leqslant d(p, v') + d(v', q) \\ &\leqslant 2d(p, v') + \epsilon \\ &\leqslant 2(d(p, u') + d(u', v')) + \epsilon \\ &\leqslant 2d(p, u') + d(u', x) + \epsilon \\ &\leqslant 3(d(u', Y) + \epsilon) + \epsilon < 7\epsilon \end{aligned}$$

Otherwise, if $d(u', Y) \geqslant \epsilon$, then we have:

$$d(u', v') < \frac{1}{2} d(u', x) \leqslant \frac{1}{2}(d(u', Y) + \epsilon) \leqslant d(u', Y)$$

By the contraction property:

$$d(p, q) \leqslant \operatorname{diam} \pi^{\epsilon}_{Y}(u') \cup \pi^{\epsilon}_{Y}(v') \leqslant \rho'(d(u', Y)) \leqslant \rho'(d(Y, Y') + \epsilon)$$

Suppose instead that δ' is a $(5, 0)$–quasi-geodesic. Then δ is not a $(6, 0)$–quasi-geodesic, since $d(Y, Y') > 2\mu(6, 0)$. By (9) we have:

$$d(x, y) < \frac{2}{3}(d(x, u)) \leqslant \frac{2}{3}(d(x, Y) + \epsilon)$$

If $d(x, Y) \leqslant 2\epsilon$ it follows that $d(x, y) \leqslant 2\epsilon$. Thus $d(y, Y) \leqslant d(y, x) + d(x, Y) \leqslant 4\epsilon$, and:

$$d(p, q) \leqslant d(q, y) + d(y, x) + d(x, p) \leqslant d(y, Y) + \epsilon + 2\epsilon + d(x, Y) + \epsilon \leqslant 10\epsilon$$

Otherwise $d(x, Y) > 2\epsilon$ and it follows that $d(x, y) \leqslant d(x, Y)$. We then use the contraction property to see:

$$d(p, q) \leqslant \operatorname{diam} \pi^{\epsilon}_{Y}(x) \cup \pi^{\epsilon}_{Y}(y) \leqslant \rho'(d(x, Y)) \leqslant \rho'(d(Y, Y') + \epsilon)$$

Since q' was an arbitrary point in Y' and q was an arbitrary point of $\pi^{\epsilon}_{Y}(q')$, we conclude $\operatorname{diam} \pi^{\epsilon}_{Y}(Y') \leqslant 2(\rho'(d(Y, Y') + \epsilon) + 10\epsilon)$. ◻

We also have the following applications of the first part of Proposition 8.1:

Corollary 8.3. *A geodesic triangle in which two of the sides are μ–Morse is δ–thin, with δ depending only on μ.*

Corollary 8.4. *Suppose X is a geodesic metric space and $\mathcal{P}$ is a collection of (ρ_1, ρ_2)–contracting paths such that for every pair of points x, $y \in X$ there exists a $\gamma \in \mathcal{P}$ with endpoints x and y. Then X is δ–hyperbolic, with δ depending only on ρ_1 and ρ_2.*

Corollary 8.4 is an analogue of [22, Theorem 2.3], which is roughly the same statement when the paths in $\mathcal{P}$ are all semi-strongly contracting with uniform contraction parameters.

Corollary 8.5. *Let G be a group generated by a finite set $\mathcal{S}$. Suppose there exist functions ρ_1 and ρ_2 and, for each $g \in G$, a path α_g from 1 to g in* $\mathrm{Cay}(G, \mathcal{S})$ *that is (ρ_1, ρ_2)–contracting. Then G is hyperbolic.*

We must assume uniform contraction in Corollary 8.5, even for finitely presented groups. Druţu, Mozes, and Sapir [17] show that if H is a finitely generated subgroup of a finitely generated group G and $h \in H$ is a Morse element in G, that is, $\langle h \rangle$ is Morse in some, hence, every, Cayley graph of G, then h is a Morse element in H. Thus, if H is a finitely generated subgroup of a torsion-free hyperbolic group then every element of H is Morse. However, Brady [9] constructed an example of a finitely presented subgroup H of a torsion-free hyperbolic group G such that H is not hyperbolic.

Fink [20] claims that if all geodesics in a homogeneous proper geodesic metric space are Morse, then the space is hyperbolic. First is an assertion, [20, Proposition 3.2], that if every geodesic is Morse then the collection of geodesics is uniformly Morse, ie, there exists a μ such that every geodesic is μ–Morse. Then an asymptotic cone argument is used to conclude the space is hyperbolic. This second step can now be accomplished via our Corollary 8.4 without resort to the asymptotic cone machinery.

References

[1] Yael Algom-Kfir. *Strongly contracting geodesics in outer space*. Geom. Topol. **15**.4 (2011), pp. 2181–2233.

[2] Tarik Aougab, Matthew Gentry Durham, and Samuel J. Taylor. *Middle recurrence and pulling back stability*. preprint. 2016. arXiv: 1609.06698.

[3] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Negative Curvature in Graphical Small Cancellation Groups*. preprint. 2016. arXiv: 1602.03767.

[4] Werner Ballmann and Michael Brin. *Orbihedra of Nonpositive Curvature*. Publ. Math. Inst. Hautes Études Sci. **82**.1 (1995), pp. 169–209.

[5] Jason A. Behrstock. *Asymptotic geometry of the mapping class group and Teichmüller space*. Geom. Topol. **10** (2006), pp. 1523–1578.

[6] Jason A. Behrstock and Cornelia Druţu. *Divergence, thick groups, and short conjugators*. Illinois J. Math. **58**.4 (2014), pp. 939–980.

[7] Mladen Bestvina, Kenneth Bromberg, and Koji Fujiwara. *Constructing group actions on quasi-trees and applications to mapping class groups*. Publ. Math. Inst. Hautes Études Sci. **122**.1 (2015), pp. 1–64.

[8] Mladen Bestvina and Koji Fujiwara. *A characterization of higher rank symmetric spaces via bounded cohomology*. Geom. Funct. Anal. **19**.1 (2009), pp. 11–40.

[9] Noel Brady. *Branched coverings of cubical complexes and subgroups of hyperbolic groups*. J. London Math. Soc. (2) **60**.2 (1999), pp. 461–480.

[10] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Vol. 319. Grundlehren der mathematischen Wissenschaften. Berlin: Springer, 1999.

[11] Pierre-Emmanuel Caprace and Koji Fujiwara. *Rank-one isometries of buildings and quasi-morphisms of Kac-Moody groups*. Geom. Funct. Anal. **19**.5 (2010), pp. 1296–1319.

[12] Christopher H. Cashen and John M. Mackay. *A Metrizable Topology on the Contracting Boundary of a Group*. preprint. 2017. arXiv: 1703.01482.

[13] Ruth Charney and Harold M. Sultan. *Contracting boundaries of* CAT(0) *spaces*. J. Topol. **8**.1 (2015), pp. 93–117.

[14] Matthew Cordes. *Morse Boundaries of Proper Geodesic Metric Spaces*. Groups Geom. Dyn. (in press). arXiv: 1502.04376v3.

[15] Matthew Cordes and David Hume. *Stability and the Morse boundary*. J. Lond. Math. Soc. (in press). arXiv: 1606.00129.

[16] F. Dahmani, V. Guirardel, and D. Osin. *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*. Mem. Amer. Math. Soc. **245**.1156 (2017), pp. v+152.

[17] Cornelia Druţu, Shahar Mozes, and Mark Sapir. *Divergence in lattices in semisimple Lie groups and graphs of groups*. Trans. Amer. Math. Soc. **362**.5 (2010), pp. 2451–2505.

[18] Cornelia Druţu and Mark Sapir. *Tree-graded spaces and asymptotic cones of groups*. Topology **44**.5 (2005). With an appendix by Denis Osin and Sapir, pp. 959–1058.

[19] Moon Duchin and Kasra Rafi. *Divergence of geodesics in Teichmüller space and the mapping class group*. Geom. Funct. Anal. **19**.3 (2009), pp. 722–742.

[20] Elisabeth Fink. *Hyperbolicity via Geodesic Stability*. preprint. 2015. arXiv: 1504.06863v1.

[21] Mikhael Gromov. *Hyperbolic groups*. *Essays in group theory*. Vol. 8. Math. Sci. Res. Inst. Publ. New York: Springer, 1987, pp. 75–263.

[22] Howard A. Masur and Yair N. Minsky. *Geometry of the complex of curves. I. Hyperbolicity*. Invent. Math. **138**.1 (1999), pp. 103–149.

[23] Howard A. Masur and Yair N. Minsky. *Geometry of the complex of curves. II. Hierarchical structure*. Geom. Funct. Anal. **10**.4 (2000), pp. 902–974.

[24] Yair N. Minsky. *Quasi-projections in Teichmüller space*. J. Reine Angew. Math. **473** (1996), pp. 121–136.

[25] Devin Murray. *Topology and Dynamics of the Contracting Boundary of Cocompact CAT(0) Spaces*. preprint. 2015. arXiv: 1509.09314v1.

[26] Alexander Yu. Olshanskii, Denis V. Osin, and Mark V. Sapir. *Lacunary hyperbolic groups*. Geom. Topol. **13**.4 (2009). With an appendix by Michael Kapovich and Bruce Kleiner, pp. 2051–2140.

[27] Denis V. Osin. *Acylindrically hyperbolic groups*. Trans. Amer. Math. Soc. **368**.2 (2016), pp. 851–888.

[28] Denis V. Osin. *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*. Mem. Amer. Math. Soc. **179**.843 (2006), pp. vi+100.

[29] Frédéric Paulin. *Un groupe hyperbolique est déterminé par son bord*. J. London Math. Soc. (2) **54**.1 (1996), pp. 50–74.

[30] Harold M. Sultan. *Explicit equivalences between CAT(0) hyperbolic type geodesics*. preprint. 2012. arXiv: 1211.6673v2.

[31] Harold M. Sultan. *Hyperbolic quasi-geodesics in CAT(0) spaces*. Geom. Dedicata **169**.1 (2014), pp. 209–224.

PAPER B

Morse subsets of CAT(0) spaces are strongly contracting

Christopher H. Cashen
Geometriae Dedicata (in press)
accepted May 7, 2019

We prove that Morse subsets of CAT(0) spaces are strongly contracting. This generalizes and simplifies a result of Sultan, who proved it for Morse quasi-geodesics. Our proof goes through the recurrence characterization of Morse subsets.

In this note we give a short proof of the following technical result:

Proposition 0.1. *If $\mathcal{Z}$ is a ρ–recurrent subset of a CAT(0) space X and the empty set is not in the image of the map $\pi_{\mathcal{Z}}(x) := \{z \in \mathcal{Z} \mid d(x,z) = d(x,\mathcal{Z})\}$ then $\mathcal{Z}$ is $12\rho(21)$–strongly contracting.*

This is the final piece of the following theorem, which says that a number of properties that are equivalent to quasi-convexity in hyperbolic spaces are also equivalent to one another in CAT(0) spaces:

Theorem 0.2. *Let X be a geodesic metric space. Let $\mathcal{Z}$ be an unbounded subset of X such that the empty set is not in the image of $\pi_{\mathcal{Z}}$. The following are equivalent:*

$\mathcal{Z}$ is Morse: *There is a function $\mu\colon [1,\infty)\times[0,\infty)\to[0,\infty)$ defined by:*

$$\mu(L,A) := \sup_{\gamma}\sup_{w\in\gamma} d(w,\mathcal{Z})$$

The first supremum is taken over (L,A)–quasi-geodesic segments γ with both endpoints on $\mathcal{Z}$.

$\mathcal{Z}$ is contracting: *There is a function $\sigma\colon [0,\infty)\to[0,\infty)$ with $\lim_{r\to\infty}\sigma(r)/r = 0$ defined by:*

$$\sigma(r) := \sup_{d(x,y)\leq d(x,\mathcal{Z})\leq r} \operatorname{diam} \pi_{\mathcal{Z}}(x)\cup\pi_{\mathcal{Z}}(y)$$

$\mathcal{Z}$ is recurrent: *There is a function $\rho\colon [1,\infty)\to[0,\infty)$ defined by:*

$$\rho(q) := \sup_{\Delta(\gamma)\leq q}\inf_{w\in\gamma} d(w,\mathcal{Z}')$$

The first supremum is taken over rectifiable segments γ with distinct endpoints on $\mathcal{Z}$ such that $\Delta(\gamma) := \frac{\mathrm{len}(\gamma)}{d(\gamma^+,\gamma^-)} \leq q$, where γ^+ and γ^- are the endpoints of γ and $\mathcal{Z}'$ is $\mathcal{Z}$ with the open balls of radius $d(\gamma^+,\gamma^-)/3$ about γ^+ and γ^- removed.

If X is hyperbolic or CAT(0) then these conditions are equivalent to:

$\mathcal{Z}$ is strongly contracting: *$\mathcal{Z}$ is contracting and the contraction gauge σ is a bounded function.*

Corollary 0.3. *Morse subsets of CAT(0) spaces are strongly contracting.*

We refer the reader to [3] for background on hyperbolic and CAT(0) spaces.

The Proposition and the Theorem can be extended to arbitrary non-empty subsets $\mathcal{Z}$ by suitable modification of the definitions. Specifically, if the empty set is in the image of π_Z then redefine $\pi_{\mathcal{Z}}(x) := \{z\in\mathcal{Z} \mid d(x,z)\leq d(x,\mathcal{Z})+1\}$. Extra bookkeeping is then required to compute an explicit contraction bound in the proof of the proposition. For bounded sets the four properties are trivially satisfied, with the possible exception that the given definition of

This research was supported by the Austrian Science Fund (FWF): P 30487-N35

recurrence does not make sense if some $\mathcal{Z}'$ is empty, which occurs, for instance, when $\mathcal{Z}$ is a two point set. We could redefine ρ to be the diameter of $\mathcal{Z}$ in this case.

The corollary confirms a conjecture of Russell, Spriano, and Tran [7] and generalizes a result of Sultan [8], who proved that Morse quasi-geodesics in CAT(0) spaces are strongly contracting.

Genevois [6] proved that Morse subsets of a finite dimensional CAT(0) cube complex $\mathcal{X}$ are strongly contracting *in the combinatorial metric*. While this is quasi-isometric to the CAT(0) metric, the property of being a strongly contracting subset is not, in general, preserved by quasi-isometries [2], so Genevois's result and our theorem are independent. However, since the Morse property *is* preserved by quasi-isometries, and since Morse equals strongly contracting in both metrics, $\mathcal{X}$ has the same strongly contracting subsets regardless of whether it is endowed with the CAT(0) or the combinatorial metric.

Proof of the theorem. The contraction condition was introduced in [1], where it was shown to be equivalent to the Morse condition. The recurrence condition was used to characterize Morse quasi-geodesics in [5], and this characterization can be extended to arbitrary subsets, as in [4, Theorem 2.2]. Strong contraction obviously implies contraction. It is easy to see that all of these properties are equivalent to quasi-convexity in hyperbolic spaces. The proposition supplies the remaining implication. □

There is extensive literature making use of the Morse property and equivalent characterizations in various settings, but a complete exposition would be longer than this paper, so we will not attempt it. Sultan's result uses a characterization of the images of Morse quasi-geodesics in asymptotic cones due to Druţu, Mozes, and Sapir [5]. Loosely speaking, this characterization depends on there being a sensible notion of one point being *between* two others, which we have for quasi-geodesics but not, at least in an obvious way, for arbitrary subsets. We avoid the use of asymptotic cones and instead use recurrence (which also comes from [5]). We construct curves in essentially the same way as Sultan, but our argument, in addition to applying to general subsets, is simpler and gives an explicit strong contraction bound.

Proof of the proposition. Define $D := \rho(21)$. Supposing the contraction gauge σ of $\mathcal{Z}$ is not bounded by $12D$, we derive a contradiction. Failure of the contraction bound means there exist points $x, y \in \mathcal{X}$ such that $d(x,y) \le d(x,\mathcal{Z})$ and such that $\operatorname{diam} \pi_{\mathcal{Z}}(x) \cup \pi_{\mathcal{Z}}(y) > 12D$. We may assume $d(x,\mathcal{Z}) \ge d(y,\mathcal{Z})$, because otherwise $d(x,y) \le d(y,\mathcal{Z})$ and we can swap the roles of x and y. Choose $x' \in \pi_{\mathcal{Z}}(x)$ and $y' \in \pi_{\mathcal{Z}}(y)$ such that $P := d(x',y') > 12D$. Let $\mathcal{Z}'$ denote the set $\mathcal{Z}$ with the open balls of radius $P/3$ about x' and y' removed.

For points $a, b \in \mathcal{X}$, let $[a,b]\colon [0,1] \to \mathcal{X}$ denote the geodesic segment from a to b, parameterized proportional to arc length. Concatenation is denoted '+'.

$$(*) \qquad \text{If } d(w,\mathcal{Z}') \le D \text{ for some } w \in \mathcal{X} \text{ then } w \notin [x',x]+[x,y]+[y,y'].$$

To see this, first suppose $w \in [x',x]$. Then $x' \in \pi_{\mathcal{Z}}(w)$, so $P/3 \le d(x',\mathcal{Z}') \le d(x',w) + d(w,\mathcal{Z}') = d(w,\mathcal{Z}) + d(w,\mathcal{Z}') \le 2d(w,\mathcal{Z}') \le 2D$, which is a contradiction, since $P > 12D$. Similarly, $w \notin [y',y]$. If $w \in [x,y]$ then:

$$d(x,w) + d(w,y) = d(x,y) \le d(x,\mathcal{Z}) \le d(x,w) + D$$

Thus, $d(w,y) \le D$, which implies:

$$P/3 \le d(y',\mathcal{Z}') \le d(y',y) + d(y,\mathcal{Z}') \le 2d(y,\mathcal{Z}') \le 2(d(y,w) + d(w,\mathcal{Z}')) \le 4D$$

Again, this contradicts the hypothesis that $P > 12D$, so $(*)$ is verified.

Now there are three cases to consider.

Case 1, $d(x,x') \le 6P$**:** Define $\gamma := [x',x]+[x,y]+[y,y']$. Then $\operatorname{len}(\gamma) \le 18P < 21P$, so recurrence says there is a point $w \in \gamma$ such that $d(w,\mathcal{Z}') \le D$. By $(*)$, this is impossible.

Case 2, $d(x,x') > 6P$ **and** $d(y,y') \le 4P$**:** Let $a := [x',x](\frac{6P}{d(x,x')})$ and $b := [y,x](\frac{6P}{d(x,x')})$, so that:

$$d(a,x') = \frac{6P}{d(x,x')} \cdot d(x,x') = 6P \quad \text{and} \quad d(b,y) = \frac{6P}{d(x,x')} \cdot d(x,y) \le 6P$$

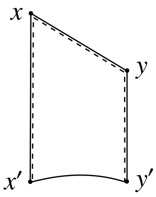

Case 1

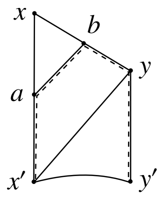

Case 2

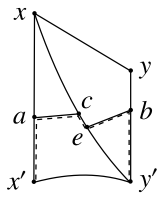

Case 3

Since $d(x', y) \leq 5P$, the CAT(0) condition implies $d(a, b) < 5P$. Define $\gamma := [x', a] + [a, b] + [b, y] + [y, y']$. Since $\mathrm{len}(\gamma) < 6P + 5P + 6P + 4P = 21P$, recurrence says there is a point $w \in \gamma$ with $d(w, \mathcal{Z}') \leq D$. By $(*)$, the only possibility is $w \in [a, b]$, but this is impossible because $d([a, b], \mathcal{Z}) \geq d(a, \mathcal{Z}) - d(a, b) > 6P - 5P = P > D$.

Case 3, $d(x, x') > 6P$ **and** $d(y, y') > 4P$**:** Let $a := [x', x](\frac{4P}{d(x,x')})$ and let $c := [y', x](\frac{4P}{d(x,x')})$. Then $d(x', a) = 4P$ and:

$$4P \leq d(y', c) = \frac{4P}{d(x, x')} \cdot d(y', x) \leq \frac{4P}{d(x, x')} \cdot (d(x, x') + P) < \frac{14}{3}P$$

Let b be the point of $[y', y]$ at distance $4P$ from y', and let e be the point of $[y', x]$ at distance $4P$ from y', so $d(c, e) < \frac{2}{3}P$. The CAT(0) condition implies that $d(a, c) < P$ and, since $d(x, y) \leq d(x, y')$, that $d(e, b) \leq 4\sqrt{2}P$.

Define $\gamma := [x', a] + [a, c] + [c, e] + [e, b] + [b, y']$. Then $\mathrm{len}(\gamma) < 4P + P + \frac{2}{3}P + 4\sqrt{2}P + 4P < 21P$, so recurrence demands a point $w \in \gamma$ with $d(w, \mathcal{Z}') \leq D$. By $(*)$, $w \notin [x', a]$, $[b, y']$. We cannot have $w \in [a, c] + [c, e]$ because $d([a, c] + [c, e], \mathcal{Z}) \geq d(a, \mathcal{Z}) - (d(a, c) + d(c, e)) > 4P - P - \frac{2}{3}P > D$. Thus, $w \in [e, b]$, so $d(e, b) = d(e, w) + d(w, b)$. However, $d(w, b) \geq d(b, \mathcal{Z}) - d(w, \mathcal{Z}) \geq 4P - D > \frac{47}{12}P$. By the same reasoning, $\frac{47}{12}P < d(a, w)$, but $d(a, w) < P + \frac{2}{3}P + d(e, w)$, so $d(e, w) > \frac{27}{12}P$. This gives us the desired contradiction:

$$6P < \frac{74}{12}P < d(e, w) + d(w, b) = d(e, b) \leq 4\sqrt{2}P < 6P$$

Since all three cases ended in contradiction, we conclude $12D$ bounds σ. □

References

[1] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Characterizations of Morse geodesics via superlinear divergence and sublinear contraction*. Doc. Math. **22** (2017), pp. 1193–1224.

[2] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Negative Curvature in Graphical Small Cancellation Groups*. Groups Geom. Dyn. **13**.2 (2019), pp. 579–632.

[3] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Vol. 319. Grundlehren der mathematischen Wissenschaften. Berlin: Springer, 1999.

[4] Christopher H. Cashen and John M. Mackay. *A Metrizable Topology on the Contracting Boundary of a Group*. Trans. Amer. Math. Soc. **372**.3 (2019), pp. 1555–1600.

[5] Cornelia Druţu, Shahar Mozes, and Mark Sapir. *Divergence in lattices in semisimple Lie groups and graphs of groups*. Trans. Amer. Math. Soc. **362**.5 (2010), pp. 2451–2505.

[6] Anthony Genevois. *Hyperbolicities in CAT(0) cube complexes*. preprint. 2017. arXiv: 1709.08843v2.

[7] Jacob Russell, Davide Spriano, and Hung Cong Tran. *Convexity in Hierarchically Hyperbolic Spaces*. preprint. 2018. arXiv: 1809.09303v2.

[8] Harold Sultan. *Hyperbolic quasi-geodesics in CAT(0) spaces*. Geom. Dedicata **169**.1 (2014), pp. 209–224.

PAPER C

Negative curvature in graphical small cancellation groups

Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume
Groups, Geometry and Dynamics **13** (2019), no. 2, 579–632
DOI: 10.4171/GGD/498

We use the interplay between combinatorial and coarse geometric versions of negative curvature to investigate the geometry of infinitely presented graphical $Gr'(1/6)$ small cancellation groups. In particular, we characterize their 'contracting geodesics', which should be thought of as the geodesics that behave hyperbolically.

We show that every degree of contraction can be achieved by a geodesic in a finitely generated group. We construct the first example of a finitely generated group G containing an element g that is strongly contracting with respect to one finite generating set of G and not strongly contracting with respect to another. In the case of classical $C'(1/6)$ small cancellation groups we give complete characterizations of geodesics that are Morse and that are strongly contracting.

We show that many graphical $Gr'(1/6)$ small cancellation groups contain strongly contracting elements and, in particular, are growth tight. We construct uncountably many quasi-isometry classes of finitely generated, torsion-free groups in which every maximal cyclic subgroup is hyperbolically embedded. These are the first examples of this kind that are not subgroups of hyperbolic groups.

In the course of our analysis we show that if the defining graph of a graphical $Gr'(1/6)$ small cancellation group has finite components, then the elements of the group have translation lengths that are rational and bounded away from zero.

1. Introduction

Graphical small cancellation theory was introduced by Gromov as a powerful tool for constructing finitely generated groups with desired geometric and analytic properties [25]. Its key feature is that it produces infinite groups with prescribed subgraphs in their Cayley graphs. The group properties are thus derived from the combinatorial or asymptotic properties of the embedded subgraphs. Over the last two decades, graphical small cancellation theory has become an increasingly prominent and versatile tool of geometric group theory with a wide range of striking examples and applications.

In this paper, we provide a thorough investigation of the hyperbolic-like geometry of graphical small cancellation constructions. Our theorems show that the constructed groups behave strongly like groups hyperbolic relative to their defining graphs. This contrasts the fact that in general they need not be Gromov hyperbolic or even relatively hyperbolic. With this geometric analogy in mind, we produce a variety of concrete examples and determine the spectrum of negative curvature possible in the realm of finitely generated groups.

Graphical small cancellation theory was first used by Gromov [25] in the description of the groups now known as 'Gromov monsters', which are finitely generated groups that contain sequences of expander graphs in their Cayley graphs. These monster groups do not coarsely embed into a Hilbert space, whence they are not coarsely amenable (i.e. do not have Yu's

This work is supported by the ERC grant of Goulnara Arzhantseva "ANALYTIC" no. 259527. The second author is supported by the Austrian Science Fund (FWF):M1717-N25. The third author is supported by the Swiss National Science Foundation Professorship FN PP00P2-144681/1 of Laura Ciobanu. The fourth author is supported by the ERC grant no. 278469 of Pierre–Emmanuel Caprace and the grant ANR-14-CE25-0004 "GAMME" of Damien Gaboriau and Romain Tessera.

property A), and they give counterexamples to the Baum-Connes conjecture with coefficients [31]. Graphical small cancellation theory is currently the only means of proving the existence of finitely generated groups with any of these three properties.

Since Gromov's initial impetus, the theory has gained in significance and variety of applications. Indeed, the construction is very versatile: every countable group embeds into a 2-generated graphical small cancellation group [28, Example 1.13], and graphical small cancellation theory has been used, for instance, to give many new groups without the unique product property [7, 46, 29], the first examples of non-coarsely amenable groups with the Haagerup property [6, 39], and new hyperbolic groups with Kazhdan's Property (T) [25, 43, 38], as well as to build a continuum of Gromov monsters [32], the first examples of finitely generated groups that do not coarsely embed into Hilbert space and yet do not contain a weakly embedded expander [8], and to analyze the Wirtinger presentations of prime alternating link groups [17]. Moreover, since the class of graphical small cancellation groups contains all classical small cancellation groups, it contains, for example, groups having no finite quotients [41], groups with prescribed asymptotic cones [51, 21], and groups with exceptional divergence functions [30].

Our paper has two purposes. The first purpose is to study geometric aspects of graphical small cancellation groups. An important property we focus on is the existence of subspaces that 'behave like' subspaces of a negatively curved space. We quantify this phenomenon by considering *contraction properties* of a subspace. Intuitively, this measures the asymptotic growth of closest point projections of metric balls to the subspace. In a recent work [3], we defined a general quantitative spectrum of contraction in arbitrary geodesic metric spaces and used it to produce new results on the interplay between contraction, divergence, and the property of being Morse. Our definitions generalize prior notions of contraction that have been instrumental in the study of numerous examples of finitely generated groups of current interest, such as mapping class groups [34, 9, 22], outer automorphism groups of free groups [36, 1], and, more generally, acylindrically hyperbolic groups [18]. In the present paper, we completely determine the contraction properties of geodesics in graphical small cancellation groups through their defining graphs. En route, we describe geodesic polygons and translation lengths in these groups.

The second purpose of this paper is to detect the range of possible contracting behaviors in finitely generated groups. To this end, we use graphical small cancellation theory to show that every degree of contraction can be achieved by a geodesic in a suitable group. Moreover, we give the first examples of strongly contracting geodesics that are not preserved under quasi-isometries of groups. These results further establish the graphical small cancellation technique as a fundamental source of novel examples of finitely generated groups.

Our main technical result is a *local-to-global* theorem for the contraction properties of geodesics in graphical small cancellation groups. It states that the contraction function of a geodesic is measured by its intersections with the defining graph. This confirms the analogy with relatively hyperbolic spaces and their peripheral subspaces. Beyond the applications alluded to above, the theorem also enables us to prove the general result that many infinitely presented graphical small cancellation groups contain strongly contracting elements and, in particular, are growth tight, and to provide a characterization of Morse geodesics in classical $C'(1/6)$ small cancellation groups. Furthermore, using the fact that strongly contracting elements give rise to hyperbolically embedded virtually cyclic subgroups, we produce the first examples of torsion-free groups in which every element is contained in a maximal virtually cyclic hyperbolically embedded subgroup but that are not subgroups of hyperbolic groups.

The proof of our local-to-global theorem rests on a meticulous analysis of the geometry of the Cayley graphs of graphical $Gr'(1/6)$ small cancellation groups. In particular, we provide a complete classification of the geodesic quadrangles in the Cayley graphs of these groups, which is of independent interest and is new even for classical $C'(1/6)$–groups.

The general tools that we establish have additional applications. For instance, we show that in many infinitely presented graphical small cancellation groups, the translation lengths of infinite order elements are rational and bounded away from zero. These tools will undoubtedly be useful towards further applications of this very interesting class of groups.

In the remainder of this introduction we explain the key concepts and main results of this paper in more detail, and give a brief overview of the proof of our local-to-global theorem.

Acknowledgements. A part of this work was developed during the program *Measured group theory* held at the Erwin Schrödinger Institute for Mathematics and Physics in Vienna in 2016. We thank its organizers and the institute for its hospitality. We also thank the Isaac Newton Institute for Mathematical Sciences for support and hospitality during the program *Non-positive curvature: group actions and cohomology*. This work is partially supported by EPSRC Grant Number EP/K032208/1.

Finally, we thank the referees for their careful reading of the paper and their helpful comments.

1.1. Contracting subspaces. In the following we shall assume that X is a geodesic metric space such that for every closed $Y \subset X$ and $x \in X$, the set $\pi(x) := \{y \in Y \mid d(x, y) = d(x, Y)\}$ is non-empty. This is true for a proper space X, but also for a connected graph (i.e. a connected 1-dimensional CW-complex) X. We call π *closest point projection to* Y. We do *not* assume the sets $\pi(x)$ have uniformly bounded diameter.

Definition 1.1 (Contracting). Let Y be a closed subspace of X, and denote by π the closest point projection to Y. Let ρ_1 and ρ_2 be non-decreasing, eventually non-negative functions, with $\rho_1(r) \leqslant r$ and ρ_1 unbounded. We say that Y is (ρ_1, ρ_2)*–contracting* if the following conditions are satisfied for all $x, x' \in X$:

- $d(x, x') \leqslant \rho_1(d(x, Y)) \implies \operatorname{diam} \pi(x) \cup \pi(x') \leqslant \rho_2(d(x, Y))$
- $\lim_{r\to\infty} \frac{\rho_2(r)}{\rho_1(r)} = 0$

If $\rho_1(r) = r$, then we say Y is *sublinearly contracting*, if $\rho_1(r) = r$ and $\rho_2(r) = C$ for some constant C we say it is *strongly contracting*, and if $\rho_1(r) = r/2$ and $\rho_2(r) = C$ for some constant C we say it is *semi-strongly contracting*

We say a function f is *sublinear* if it is non-decreasing, eventually non-negative, and $\lim_{r\to\infty} \frac{f(r)}{r} = 0$.

The most basic example of a strongly contracting subspace is a geodesic in a tree or, more generally, a geodesic in a δ–hyperbolic space. The opposite extreme occurs in a Euclidean space where there are no contracting geodesics (for any choice of ρ_1 and ρ_2). The contrast between hyperbolic and Euclidean type behavior is evident in the following well-known examples of contraction:

- A geodesic in a CAT(0) space is strongly contracting if and only if it is Morse [13, 48].
- A geodesic in a relatively hyperbolic space is strongly contracting if for every $C \geqslant 0$ there exists a $B \geqslant 0$ such that the geodesic spends at most time B in the C–neighborhood of a peripheral subset [44].

The common idea in these situations is that the given space has certain regions that are not hyperbolic, but geodesics that avoid these non-hyperbolic regions behave very much like geodesics in a hyperbolic space. Similar phenomena occur for pseudo-Anosov axes in the Teichmüller space of a hyperbolic surface and iwip axes in the Outer Space of the outer automorphism group of a free group. Such axes avoid the 'thin parts' of their respective spaces and therefore are strongly contracting [36, 1].

A version of semi-strong contraction, where the projection is not necessarily closest point projection, occurs for pseudo-Anosov axes in the mapping class group of a hyperbolic surface [34, 9, 22].

1.2. Local-to-global theorem. Given a directed graph Γ whose edges are labelled by the elements of a set $\mathcal{S}$, the group defined by Γ, denoted $G(\Gamma)$ is given by the presentation $\langle \mathcal{S} \mid \text{labels of embedded cycles in } \Gamma\rangle$. The graphical $Gr'(1/6)$ small cancellation condition, see Section 2.1, is a combinatorial requirement on the labelling of Γ, whose key consequence is that the connected components of Γ isometrically embed into $\mathrm{Cay}(G(\Gamma), \mathcal{S})$. In the case that Γ is a disjoint union of cycle graphs labelled by a set of words $\mathcal{R}$, the graphical $Gr'(1/6)$–condition for Γ corresponds to the classical $C'(1/6)$–condition for $\mathcal{R}$.

We show that, similar to the situations described above, geodesics in Cayley graphs of graphical $Gr'(1/6)$ small cancellation groups behave like hyperbolic geodesics as long as they avoid the embedded components of the defining graph. In fact, a geodesic is as hyperbolic as its intersections with the embedded components of Γ:

Theorem (Theorem 4.1). *Let Γ be a $Gr'(1/6)$–labelled graph. There exist ρ_1' and ρ_2' such that a geodesic α in $X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$ is (ρ_1', ρ_2')–contracting if and only if there exist ρ_1 and ρ_2 such that for every embedded component Γ_0 of Γ in X such that $\Gamma_0 \cap \alpha \neq \emptyset$, we have that $\Gamma_0 \cap \alpha$ is (ρ_1, ρ_2)–contracting as a subspace of Γ_0.*

Moreover, ρ_1' and ρ_2' can be bounded in terms of ρ_1 and ρ_2, and when $\rho_1(r) \geqslant r/2$ we can take $\rho_1' = \rho_1$ and $\rho_2' \asymp \rho_2$.

Here $\asymp$ denotes a standard notion of asymptotic equivalence, see Section 2. Our theorem gives the following explicit application to classical $C'(1/6)$–groups. We denote by $|\cdot|$ the number of edges of a path graph or cycle graph.

Theorem (Corollary 4.14). *Let Γ be a $Gr'(1/6)$–labelled graph whose components are cycle graphs. Let α be a geodesic in $X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$. Define $\rho(r) := \max_{|\Gamma_i| \leqslant r} |\Gamma_i \cap \alpha|$, where the Γ_i range over embedded components of Γ in X. Then α is sublinearly contracting if and only if ρ is sublinear, in which case α is (r, ρ)–contracting. In particular, α is strongly contracting if and only if ρ is bounded.*

1.3. Morse geodesics. A classically more well-studied notion of what it means to behave like a subspace of a hyperbolic space is the property of being *Morse*.

Definition 1.2 (Morse). A subspace Y of a geodesic metric space X is *μ–Morse* if every (L, A)–quasi-geodesic in X with endpoints on Y is contained in the $\mu(L, A)$–neighborhood of Y. A subspace is *Morse* if there exists some μ such that it is μ–Morse.

The property of being Morse is invariant under quasi-isometries, and the fact that quasi-geodesics in a Gromov hyperbolic space are Morse is known as the 'Morse Lemma'. These two results are main ingredients in the proof that hyperbolicity is preserved by quasi-isometries. Morse geodesics are of further interest due to their close connection with the geometry of asymptotic cones and relations with other important geometric concepts such as divergence, see for example [20, 10, 3] and references therein. In [3], we prove that being contracting is, in fact, equivalent to being Morse.

Theorem 1.3 ([3, Theorem 1.4]). *If Y is a subspace of a geodesic metric space such that the empty set is not in the image of closest point projection to Y, then Y is Morse if and only if Y is (ρ_1, ρ_2)-contracting for some ρ_1 and ρ_2.*

Thus, in the case of classical $C'(1/6)$–groups, Theorem 4.1 and Theorem 1.3 enable us to provide a complete characterization of Morse geodesics in the Cayley graph.

Theorem (Corollary 4.14). *Let Γ be a $Gr'(1/6)$–labelled graph whose components are cycle graphs. Let α be a geodesic in $\mathrm{Cay}(G(\Gamma), \mathcal{S})$. Define $\rho(r) := \max_{|\Gamma_i| \leqslant r} |\Gamma_i \cap \alpha|$, where the Γ_i range over embedded components of Γ. Then α is Morse if and only if ρ is sublinear.*

1.4. Range of contracting behaviors. As mentioned, in a CAT(0)-space, a geodesic is Morse if and only if it is strongly contracting. Thus, Theorem 1.3 says that in a CAT(0) space a geodesic is either strongly contracting or not contracting at all. We show that in finitely generated groups, the spectrum of contraction is, in fact, much richer: *every* degree of contraction can be attained.

Theorem (Theorem 4.15). *Let ρ be a sublinear function. There exists a group G with finite generating set $\mathcal{S}$ and a sublinear function $\rho' \asymp \rho$ such that there exists an (r, ρ')–contracting* *geodesic α in* $\mathrm{Cay}(G,\mathcal{S})$*, and ρ' is optimal, in the sense that if α is (r,ρ'')–contracting for some other function ρ'' then* $\limsup_{r\to\infty}\frac{\rho''(2r)}{\rho(r)} \geqslant 1$.

Furthermore α can be chosen to be within finite Hausdorff distance of a cyclic subgroup of G.

Our current examples are not finitely presentable. The contraction spectrum for finitely presented groups remains largely unexplored. Indeed, if one restricts to geodesics within finite Hausdorff distance of a cyclic subgroup, finitely presented groups can only display countably many degrees of contraction.

Question 1.4. For which functions ρ do there exist finitely presented groups G containing a geodesic in some Cayley graph that is (r,ρ)–contracting?

1.5. Non-stability of strong contraction. While the property of being Morse is stable under quasi-isometries, it has remained unknown whether the property of being strongly contracting is. We provide a negative answer by providing the first examples of spaces X and $\tilde{X}$ and geodesics γ and $\tilde{\gamma}$ such that there exists a quasi-isometry $X \to \tilde{X}$ mapping γ to $\tilde{\gamma}$ and such that γ is not strongly contracting, but $\tilde{\gamma}$ is strongly contracting.

Theorem (Theorem 4.19). *There exists a group G with finite generating sets $\mathcal{S} \subset \tilde{\mathcal{S}}$ and an infinite geodesic γ in $X :=$* $\mathrm{Cay}(G,\mathcal{S})$ *labelled by the powers of a generator such that γ is not strongly contracting, but its image $\tilde{\gamma}$ in $\tilde{X} :=$* $\mathrm{Cay}(G,\tilde{\mathcal{S}})$ *obtained from the inclusion $\mathcal{S} \subset \tilde{\mathcal{S}}$ is an infinite strongly contracting geodesic.*

Indeed, in many familiar settings, such as hyperbolic groups, CAT(0) groups, or toral relatively hyperbolic groups, such examples could not be obtained, since in those contexts, strong contraction is equivalent to the Morse property.

1.6. Strongly contracting elements and growth tightness. Another of our main results is the existence of strongly contracting elements in many graphical small cancellation groups:

Theorem (Theorem 5.1). *Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Assume that $G(\Gamma)$ is infinite. Then there exists an infinite order element $g \in G(\Gamma)$ such that $\langle g\rangle$ is strongly contracting in* $\mathrm{Cay}(G(\Gamma),\mathcal{S})$.

The element g is, in fact, the WPD element for the action on the hyperbolic coned-off space in Gruber and Sisto's proof of acylindrical hyperbolicity of these groups [30]. Theorem 5.1 has the following consequence (which does not follow from acylindrical hyperbolicity):

Arzhantseva, Cashen, and Tao [4] have shown that the action of a finitely generated group G on a Cayley graph X is *growth tight* if the action has a strongly contracting element, that is, an element g such that $\langle g\rangle$ is strongly contracting in X. Growth tightness means that the exponential growth rate of an orbit of G in X is strictly greater than the growth rate of an orbit of G/N in $N\backslash X$, for every infinite normal subgroup N. Theorem 5.1 therefore implies:

Theorem (Theorem 5.2). *Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Then the action of $G(\Gamma)$ on* $\mathrm{Cay}(G(\Gamma),\mathcal{S})$ *is growth tight.*

This has raised our interest in the following question, first asked by Grigorchuk and de la Harpe for hyperbolic fundamental groups of closed orientable surfaces [23]:

Question 1.5. Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Does $G(\Gamma)$ attain its infimal growth rate with respect to the generating set $\mathcal{S}$?

Together with Theorem 5.2, a positive answer, even for the subclass of classical $C'(1/6)$–groups, would establish small cancellation theory as an abundant source of *Hopfian* groups.

1.7. Contraction and hyperbolically embedded subgroups. An important application of the notion of strong contraction is the fact that an infinite order element whose orbit in the Cayley graph is strongly contracting is contained in a virtually cyclic hyperbolically embedded subgroup [18], see Definition 6.1. In particular, our proof of Theorem 5.1 gives a new argument that the WPD elements of [30] produce hyperbolically embedded subgroups. Furthermore, our methods detect strongly contracting elements that need not be hyperbolic elements for the action on the coned-off space of [30], see Example 4.18 and Remark 5.8. Admitting a proper infinite hyperbolically embedded subgroup is equivalent to being acylindrially hyperbolic [40], which implies a number of strong group theoretic properties.

Not all hyperbolically embedded virtually cyclic subgroups are strongly contracting; see [4] for an example. However, every hyperbolically embedded subgroup of a finitely generated group is Morse [45]. In light of Theorem 1.3, a natural question is whether there exists some critical rate of contraction that guarantees a subgroup is hyperbolically embedded. That is, does there exist an unbounded sublinear function ρ_2 such that every element g with a (r, ρ_2)–contracting orbit in some Cayley graph has a hyperbolically embedded virtually cyclic elementary closure? The elementary closure of g is the subgroup generated by all virtually cyclic subgroups containing g. We prove no such ρ_2 exists.

Theorem (Theorem 6.4). *Let ρ_2 be an unbounded sublinear function. There exists a $Gr'(1/6)$–labelled graph Γ with set of labels $\mathcal{S} := \{a, b\}$ whose components are all cycles such that $G(\Gamma)$ has the following properties: Any virtually cyclic subgroup E of $G(\Gamma)$ containing $\langle a\rangle$ is (r, ρ_2')–contracting in the Cayley graph* $\mathrm{Cay}(G(\Gamma), \mathcal{S})$ *for some $\rho_2' \asymp \rho_2$, but E is not hyperbolically embedded in $G(\Gamma)$.*

1.8. Hyperbolically embedded cycles. In a subgroup of a hyperbolic group, every infinite order element is contained in a maximal virtually cyclic, hyperbolically embedded subgroup, whence we define the following:

Definition 1.6 (HEC property). A group has the *hyperbolically embedded cycles property (HEC property)* if the elementary closure $E(g)$ of every infinite order element g is virtually cyclic and hyperbolically embedded.

It is natural to ask whether this property characterizes subgroups of hyperbolic groups. While torsion presents an obvious complication, see Section 6.2, we also present a negative answer to our question in the torsion-free case.

Theorem (Theorem 6.6). *There exist $2^{\aleph_0}$ pairwise non-quasi-isometric finitely generated torsion-free groups in which every non-trivial cyclic subgroup is strongly contracting and which, therefore, have the HEC property.*

These are the first examples of groups of this kind that do not arise as subgroups of hyperbolic groups. Our examples include exotic specimens such as Gromov monsters.

1.9. Translation lengths. Let $|\cdot|$ be the word length in $G(\Gamma)$ with respect to $\mathcal{S}$. The *translation length* of an element $g \in G(\Gamma)$ is:

$$\tau(g) := \lim_{n\to\infty} \frac{|g^n|}{n}$$

Conner [15] calls a group whose non-torsion elements have translation length bounded away from zero *translation discrete*. Hyperbolic groups [49], CAT(0) groups [16], and finitely presented groups satisfying various classical small cancellation conditions [33] are translation discrete.

We show that many (possibly infinitely presented) graphical small cancellation groups are also translation discrete:

Theorem (Theorem 5.4). *Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Then every infinite order element of $G(\Gamma)$ has rational translation length, and translation lengths are bounded away from zero.*

1.10. The idea of the proof of the local-to-global theorem. In a tree, geodesic quadrangles are degenerate, as seen in Figure 1. In a hyperbolic space, geodesic quadrangles can be approximated by geodesic quadrangles in a tree. If the base is a fixed geodesic α, the top is some given geodesic γ, and the sides are given by closest point projection from the endpoints

FIGURE 1. Tree quadrangles

of the top to the bottom, then the resulting geodesic quadrangle is either 'short' or 'thin', as in Figure 2.

FIGURE 2. Hyperbolic quadrangles from closest point projection

In Proposition 3.20, we show a combinatorial version of this dichotomy through an analysis of van Kampen diagrams in graphical $Gr'(1/6)$ small cancellation groups. Specifically, if $X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$ and $\alpha \subset X$ is a (ρ_1, ρ_2)–contracting geodesic, $\gamma \subset X$ is another geodesic, and each endpoint of γ is connected via a geodesic to a closest point of α, then the boundary word of the resulting geodesic quadrangle admits a van Kampen diagram that is either 'short' or 'thin' in terms of number of faces, as depicted in Figure 3.

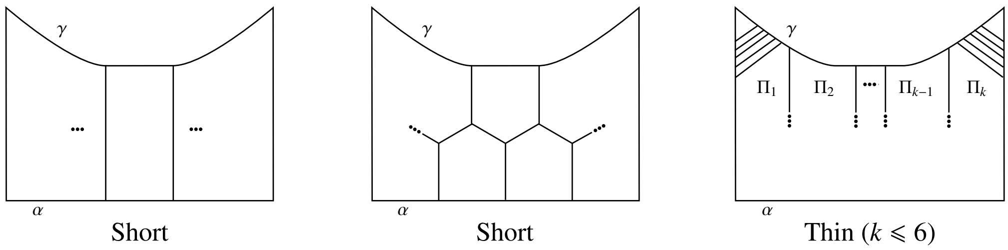

FIGURE 3. Combinatorially short and thin quadrangles

The main ingredient in establishing this dichotomy is a classification of 'special combinatorial geodesic quadrangles', see Theorem 3.18, that extends Strebel's classification of geodesic bigons and triangles in small cancellation groups, see Theorem 3.13. This classification is of independent interest, and is novel even within the class of classical small cancellation groups.

These combinatorial versions of short and thin quadrangles do not immediately imply their metric counterparts, because the faces in the van Kampen diagrams may have boundary words that are arbitrarily long relators. In Section 4 we use the fact that α is (ρ_1, ρ_2)–contracting to show that even if the faces have long boundaries, their projections to α are small with respect to their distance from α.

The essential trick that is used repeatedly is to play off the small cancellation condition against the contraction condition. Specifically, if Π is a face of the van Kampen diagram with few sides, one of which sits on α, we use the small cancellation condition to show that $|\alpha \cap \partial\Pi|$ is bounded below by a linear function of $|\partial\Pi|$. Then we use the contraction condition to say that $|\alpha \cap \partial\Pi|$ is bounded above by a sublinear function of $|\partial\Pi|$. Thus, we have a sublinear function of $|\partial\Pi|$ that gives an upper bound to a linear function of $|\partial\Pi|$. This is only possible if $|\partial\Pi|$ is smaller than some bound depending on the two functions.

2. Preliminaries

We set notation. Let $\mathcal{S}$ be a set.

- $\langle \mathcal{S} \rangle$ is the group generated by $\mathcal{S}$. If $\mathcal{S}$ is a subset of a group G then $\langle \mathcal{S} \rangle$ is the subgroup of G generated by $\mathcal{S}$. If $\mathcal{S}$ is a set of formal symbols then $\langle \mathcal{S} \rangle$ is the free group freely generated by $\mathcal{S}$.
- If $\mathcal{S}$ is a set of formal symbols then $\mathcal{S}^- := \{s^{-1} \mid s \in \mathcal{S}\}$ is the set of formal inverses, and $\mathcal{S}^\pm := \mathcal{S} \cup \mathcal{S}^-$.
- $\mathcal{S}^*$ is the free monoid over $\mathcal{S}$.
- $2^{\mathcal{S}}$ is the set of subsets of $\mathcal{S}$.
- $\mathcal{S}^{\mathbb{N}}$ is the set of infinite sequences with terms in $\mathcal{S}$.

We write $f \preceq g$ if there exist $C_1 > 0$, $C_2 > 0$, $C_3 \geqslant 0$, and $C_4 \geqslant 0$ such that $f(x) \leqslant C_1 g(C_2 x + C_3) + C_4$ for all x. If $f \preceq g$ and $g \preceq f$ then we write $f \asymp g$.

Note that if $f \preceq g$ and g is bounded then f is bounded, and if f is non-decreasing and eventually non-negative and $f \preceq g$ for a function g such that $\lim_{r\to\infty} \frac{g(r)}{r} = 0$, then f is sublinear.

The *girth* of a graph is the length of its shortest non-trivial cycle.

2.1. Graphical small cancellation. Graphical small cancellation theory is a generalization of classical small cancellation theory. The main application is an embedding of a desired sequence of graphs into the Cayley graph of a group. It was introduced by Gromov [25], and was later clarified and expanded by Ollivier [37], Arzhantseva and Delzant [5], and, in a systematic way, by Gruber [26].

2.1.1. *Basic facts.* Let Γ be a directed graph with edges labelled by a set $\mathcal{S}$. We allow paths to traverse edges against their given direction, with the convention that the label of an oppositely traversed edge is the formal inverse of the given label. Thus, given a finite path in Γ we can *read* a word in $(\mathcal{S}^\pm)^*$ by concatenating the labels of the edges along the path.

We require that the labelling is *reduced*, in the sense that no vertex has two incident outgoing edges with the same label, and no vertex has two incident incoming edges with the same label. This implies that the word read on an immersed path is freely reduced and, hence, an element of $\langle \mathcal{S} \rangle$. Also, the word read on an immersed cycle is cyclically reduced.

Let $\mathcal{R}$ be the set of words in $\langle \mathcal{S} \rangle$ read on embedded cycles in Γ. Note that this definition implies that elements of $\mathcal{R}$ are cyclically reduced and that $\mathcal{R}$ is closed under inversion and cyclic permutation of its elements.

Definition 2.1 (Group defined by a labelled graph). The group G defined by a reduced $\mathcal{S}$–labelled graph Γ is the group $G(\Gamma) := \langle \mathcal{S} \mid \mathcal{R} \rangle$.

The notion of a group defined by a labelled graph first appeared in Rips and Segev's construction of torsion-free groups without the unique-product property [42].

Definition 2.2 (Piece). A *piece* is a labelled path graph p that admits two distinct label-preserving maps $\phi_1, \phi_2\colon\ p \to \Gamma$ such that there is no label-preserving automorphism ψ of Γ with $\phi_2 = \psi \circ \phi_1$.

Definition 2.3 ($Gr'(\lambda)$ and $C'(\lambda)$ conditions). Let Γ be a reduced labelled graph, and let $\lambda > 0$.

Γ is *$Gr'(\lambda)$–labelled* if whenever p is a piece contained in a simple cycle c of Γ then $|p| < \lambda |c|$.

Γ is *$C'(\lambda)$–labelled* if it is $Gr'(\lambda)$–labelled and, in addition, every label-preserving automorphism of Γ restricts to the identity on every connected component with non-trivial fundamental group.

A presentation $\langle \mathcal{S} \mid \mathcal{R} \rangle$ satisfies the *classical $C'(\lambda)$-condition* if the disjoint union of cycle graphs labelled by the elements of $\mathcal{R}$ is a $Gr'(\lambda)$-labelled graph.

Actually, every group is defined by a $Gr'(1/6)$–labelled graph: simply take Γ to be its Cayley graph with respect to any generating set of the group [26, Example 2.2]. Therefore, general statements about groups defined by $Gr'(1/6)$–labelled graphs either require that some additional condition be imposed on Γ or are tautologically true when $\Gamma = \mathrm{Cay}(G(\Gamma), \mathcal{S})$.

A subspace Y of a geodesic metric space is *convex* if every geodesic segment between points of Y is contained in Y.

LEMMA 2.4 ([30, Lemma 2.15]). *Let Γ be a $Gr'(^1/_6)$–labelled graph. Let Γ_i be a component of Γ. For any choice of a vertex $x \in X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$ and any vertex $y \in \Gamma_i$ there is a unique label-preserving map $\Gamma_i \to X$ that takes y to x, and this map is an isometric embedding with convex image.*

DEFINITION 2.5 (Embedded component). An *embedded component* Γ_0 of Γ refers to the image of an isometric embedding of some Γ_i into $X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$ via a label-preserving map. Equivalently, it is a $G(\Gamma)$–translate in X of the image of Γ_i under the unique label-preserving map determined by an arbitrary choice of basepoints in X and Γ_i.

We consider a graph Γ as a sequence $(\Gamma_i)_i$ of its connected components.

2.2. Contraction terminology. Recall Definition 1.1. In [3] we considered, more generally, almost closest point projections $x \mapsto \{y \in Y \mid d(x, y) \leqslant d(x, Y) + \epsilon\}$ to ensure the empty set is not in the image of the projection. That is unnecessary in this paper as we are in the case that Y is a subgraph of a graph X, which guarantees $\emptyset \notin \mathrm{Im}\,\pi$. Here, and from now on, $\pi\colon X \to 2^Y$ denotes closest point projection to Y.

We say a geodesic α in $\mathrm{Cay}(G(\Gamma), \mathcal{S})$ is *locally (ρ_1, ρ_2)–contracting* if, for each embedded component Γ_0 of Γ such that $\Gamma_0 \cap \alpha$ is non-empty, closest point projection in Γ_0 of Γ_0 to $\Gamma_0 \cap \alpha$ is (ρ_1, ρ_2)–contracting.

We say a geodesic is *uniformly locally contracting* if there exist ρ_1 and ρ_2 such that it is locally (ρ_1, ρ_2)–contracting. We add 'uniform' here to stress that the intersection of the geodesic with each embedded component of Γ is contracting with respect to the same contraction functions. Similarly, a geodesic is *uniformly locally sublinearly contracting* if it is locally (r, ρ_2)–contracting, and is *uniformly locally strongly contracting* is it locally (r, ρ_2)–contracting for ρ_2 bounded.

3. Classification of quadrangles

In this section, we establish geometric results that will let us prove our theorems about contraction in graphical small cancellation groups. In particular, we provide a complete classification of the geodesic quadrangles in the Cayley graph of a group defined by a $Gr'(^1/_6)$–labelled graph. The main technical result to be used in our subsequent investigation will be recorded in Proposition 3.20.

3.1. Combinatorial geodesic polygons. One of the main tools of small cancellation theory are so-called 'van Kampen diagrams'.

DEFINITION 3.1 (Diagram). A (disc) *diagram* is a finite, simply-connected, 2–dimensional CW complex with an embedding into the plane, considered up to orientation-preserving homeomorphisms of the plane. It is $\mathcal{S}$*–labelled* if its directed edges are labelled by elements in $\mathcal{S}$. It is a diagram *over $\mathcal{R}$* if it is $\mathcal{S}$–labelled and the word read on the boundary of each 2–cell belongs to $\mathcal{R}$. A diagram is *simple* if it is homeomorphic to a disc.

If D is a diagram over $\mathcal{R}$, b is a basepoint in D, and g is an element of $G = \langle \mathcal{S} \mid \mathcal{R} \rangle$, then there exists a unique label-preserving map from the 1–skeleton of D into $\mathrm{Cay}(G, \mathcal{S})$ taking b to g. In general, this map need not be an immersion.

An *arc* in a diagram D is a maximal path of length at least 1 all of whose interior vertices have valence 2 in D. An *interior arc* is an arc whose interior is contained in the interior of D. An *exterior arc* is an arc contained in the boundary of D. A *face* is the image of a closed 2-cell of D. If Π is a face, its *interior degree* $i(\Pi)$ is the number of interior arcs in its boundary. Likewise, its *exterior degree* $e(\Pi)$ is the number of exterior arcs. An *interior face* is one with exterior degree 0; an *exterior face* is one with positive exterior degree.

If D is a finite, simply connected, planar, 2–dimensional CW-complex whose boundary is written as a concatenation of immersed subpaths $\gamma_1, \ldots, \gamma_k$, which we call *sides* of D, then there is a unique, up to orientation-preserving homeomorphism of $\mathbb{R}^2$, embedding $\phi\colon D \to \mathbb{R}^2$ such that the concatenation of the $\phi(\gamma_i)$ is the positively oriented boundary $\partial\phi(D)$. This claim follows easily from the Schoenflies Theorem. Thus, $(D, (\gamma_i)_i)$ uniquely determines a (not necessarily simple) diagram $\phi(D)$. We call ϕ the *canonical embedding* of $(D, (\gamma_i)_i)$. Having said this once, we omit ϕ from the notation and conflate D and the γ_i with their ϕ–images.

Definition 3.2 ($(3, 7)$–diagram). A $(3, 7)$–*diagram* is a diagram such that every interior vertex has valence at least three and every interior face has interior degree at least seven.

Definition 3.3 (Combinatorial geodesic polygon [30, Definition 2.11]). A *combinatorial geodesic n–gon* $(D, (\gamma_i)_i)$ is a $(3, 7)$–diagram D whose boundary is a concatenation of immersed subpaths $\gamma_0, \dots, \gamma_{n-1}$ such that each boundary face whose exterior part is a single arc that is contained in one of the sides γ_i has interior degree at least 4. A valence 2 vertex that belongs to more than one side is called a *distinguished vertex*. A face whose exterior part contains an arc not contained in one of the sides is a *distinguished face*.

The ordering of the sides of a combinatorial geodesic n–gon is considered up to cyclic permutation, with subscripts modulo n. We also refer to 'the combinatorial geodesic n–gon D' when the sides are clear from context or irrelevant. We can also say 'combinatorial geodesic polygon' when the number of sides is irrelevant. Following common usage, 2–gons, 3–gons, and 4–gons will respectively be denominated *bigons*, *triangles*, and *quadrangles*.

If D is a simple combinatorial geodesic n–gon then every distinguished face contains a distinguished vertex, so there are at most n distinguished faces.

We record the following crucial fact about diagrams over graphical small cancellation presentations. In the following, $\mathcal{R}$ is the set of labels of simple cycles on a $Gr'(1/6)$–graph Γ labelled by the set $\mathcal{S}$, and $X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$.

Lemma 3.4 ([26, Lemma 2.13]). *Let Γ be a $Gr'(1/6)$–labelled graph, and let $w \in \langle \mathcal{S} \rangle$ represent the identity in $G(\Gamma)$. Then, there exists an $\mathcal{S}$-labelled diagram over $\mathcal{R}$ with boundary word w in which every interior arc is a piece.*

The sides of a combinatorial geodesic polygon are *not* assumed to be geodesic. The definition and choice of terminology are motivated by the following proposition. An n–gon P in X is a closed edge path that decomposes into immersed simplicial subpaths $\gamma'_0, \dots, \gamma'_{n-1}$, which are called *sides of P*.

Proposition 3.5. *If P is an n–gon in X with sides $\gamma'_0, \dots, \gamma'_{n-1}$ that are geodesics then there is an $\mathcal{S}$–labelled diagram D over $\mathcal{R}$ with sides $\gamma_0, \dots, \gamma_{n-1}$ such that for each $0 \leqslant i < n$ the word of $\langle \mathcal{S} \rangle$ read on γ_i is the same as the word read on γ'_i. Furthermore, we can choose D in such a way that after forgetting interior vertices of valence 2, we obtain a combinatorial geodesic n–gon $(D, (\gamma_i)_i)$*

Here, forgetting one interior vertex of valence 2 means replacing its two incident edges by a single one. Note that, when performing this operation, we consider D merely as (unlabelled) diagram, i.e. we ignore the orientations and labels of edges. Forgetting interior vertices of valence 2 means iterating this operation, such that we end up with a diagram without interior vertices of valence 2.

Proof. The existence of an $\mathcal{S}$–labelled diagram over $\mathcal{R}$ whose boundary label matches the label of P is the well-known van Kampen Lemma. Lemma 3.4 guarantees that the diagram can be chosen such that all interior arcs are pieces. The small cancellation condition then implies interior faces have interior degree at least 7. If there is a face Π with $e(\Pi) = 1$ whose exterior part is contained in a single γ_i, then the exterior part is a geodesic. Thus, the length of the interior part is at least half of the length of $\partial\Pi$. Since interior arcs are pieces, the small cancellation condition implies there must be at least four of them to account for half the length of $\partial\Pi$. Now if we forget interior vertices of valence 2, we have the desired combinatorial geodesic polygon. □

Remark 3.6. The word $w \in \langle \mathcal{S} \rangle$ read on a cycle in X represents the trivial element in $G(\Gamma)$. The combinatorial geodesic n–gon of Proposition 3.5 is a special type of van Kampen diagram witnessing the triviality of the word w labelling an n–gon in X whose sides are geodesics. In the remainder of Section 3 we make combinatorial arguments about arbitrary $(3, 7)$–diagrams, not necessarily $\mathcal{S}$–labelled diagrams over $\mathcal{R}$. In Section 4 we use Proposition 3.5 to apply results of this section to graphical small cancellation groups.

We record an equivalent formulation of the Euler characteristic formula for certain diagrams. Recall that for a face Π of a diagram, $e(\Pi)$ is the exterior degree of Π, which is the number of exterior arcs in its boundary. Similarly, $i(\Pi)$ is the interior degree of Π.

LEMMA 3.7 (Strebel's curvature formula, [47, p.253]). *Let D be a simple diagram without vertices of degree 2. Then:*

$$6 = 2\sum_{v}(3 - d(v)) + \sum_{e(\Pi)=0}(6 - i(\Pi)) + \sum_{e(\Pi)=1}(4 - i(\Pi)) + \sum_{e(\Pi)\geqslant 2}(6 - 2e(\Pi) - i(\Pi)).$$

Here $d(v)$ denotes the degree of a vertex v.

It readily follows from Lemma 3.7 that any $(3, 7)$–diagram with more than one face has at least 2 faces with exterior degree 1 and interior degree at most 3. (This is usually known as Greendlinger's lemma.) Therefore:

LEMMA 3.8. *The sides of a combinatorial geodesic polygon are embedded, and every combinatorial geodesic polygon has at least two sides.*

The same argument gives the following well-known fact, which greatly simplifies many considerations:

LEMMA 3.9. *Let D be a $(3, 7)$–diagram. Then any face is simply connected.*

We also state an immediate consequence of [27, Lemma 4.14]:

LEMMA 3.10. *If Π is a face of a combinatorial geodesic polygon D and α is a side of D then $\Pi \cap \alpha$ is empty or connected. If $\Pi_1, \dots, \Pi_k$ is a sequence of faces of a combinatorial geodesic polygon D such that $\Pi_i \cap \Pi_{i+1} \neq \emptyset$ for all $1 \leqslant i < k$ and α is a side of D such that $\Pi_i \cap \alpha \neq \emptyset$ for all i then $\cup_{1\leqslant i\leqslant k}\Pi_i \cap \alpha$ is connected.*

DEFINITION 3.11 (Degenerate). A combinatorial geodesic n–gon $(D, (\gamma_i)_i)$ is *degenerate* if there exists an i such that $D, \gamma_0, \dots, \gamma_i\gamma_{i+1}, \dots, \gamma_{n-1}$ is a combinatorial geodesic $(n - 1)$–gon. In this case the terminal vertex of γ_i is called a *degenerate vertex*.

It will be useful to minimize the number of sides of a diagram D by replacing a degenerate combinatorial geodesic n–gon $(D, (\gamma_i)_i)$ with a non-degenerate combinatorial geodesic k–gon $(D, (\gamma'_i)_i)$ for some $k < n$.

3.2. Reducibility. In this section we define operations for combining and reducing combinatorial geodesic n–gons. We stress that the setting is only combinatorial — these operations need not preserve the property of being a diagram over $\mathcal{R}$.

First note that if $(D, (\gamma_i)_i)$ is a combinatorial geodesic n–gon, and if D' is obtained from D by subdividing an edge, then $(D', (\gamma_i)_i)$ is still a combinatorial geodesic n–gon. The new vertex produced by subdivision is a non-distinguished vertex of valence 2. Conversely, if v is a non-distinguished vertex of valence 2 then we can 'forget' it by replacing the two incident edges with a single edge.

If D is simple and non-degenerate then by forgetting all non-distinguished vertices of valence 2 we can arrange that the distinguished vertices are exactly the vertices of valence 2 and all other vertices have valence at least 3.

DEFINITION 3.12 (Reducible). A combinatorial geodesic l–gon P is *reducible* if it admits a vertex, edge, or face reduction, as defined below, see Figure 4. It is *irreducible* otherwise.

In all of the following cases, let $(D, (\gamma_i)_i)$ be a combinatorial geodesic n–gon, and let $(D', (\gamma'_i)_i)$ be a combinatorial geodesic n'–gon.

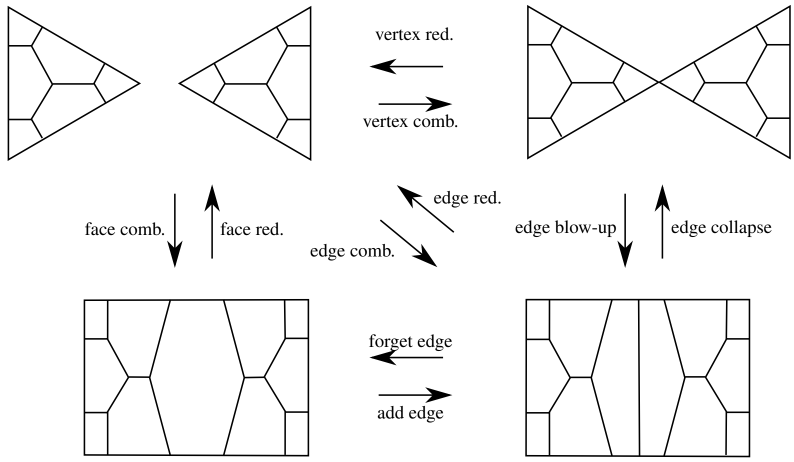

FIGURE 4. Combination and reduction

3.2.1. *Vertex reduction.* Suppose $v \in D$ is a separating vertex that is in the boundary of exactly two faces and that these two faces are the only maximal cells containing v. Suppose that there are exactly two sides, γ_i and γ_j, containing v, which is necessarily true if $(D, (\gamma_i)_i)$ is non-degenerate. Let γ_i^{-v} denote the initial path of γ_i ending at v, and let γ_i^{v+} denote the terminal path of γ_i beginning at v, and similarly for γ_j. Define the *vertex reduction of* $(D, (\gamma_i)_i)$ *at* v to be the two combinatorial geodesic polygons whose underlying CW-complexes are each D minus one of the complementary components of $D \setminus v$, respectively, and whose sides are, respectively, $\gamma_i^{v+}, \gamma_{i+1}, \dots, \gamma_{j-1}, \gamma_j^{-v}$ and $\gamma_j^{v+}, \gamma_{j+1}, \dots \gamma_{i-1}, \gamma_i^{-v}$.

Note that each of the resulting combinatorial polygons contains a distinguished vertex corresponding to v.

The inverse operation of vertex reduction we denominate a *vertex combination*. Suppose that $v \in D$ is a distinguished vertex that is contained in the boundary of a single face and is not contained in any other maximal cell. Then there is some i such that $v = \gamma_i \cap \gamma_{i+1}$. Make corresponding assumptions for $v' = \gamma'_{i'} \cap \gamma'_{i'+1} \subset D'$. The *vertex combination of* $(D, (\gamma_i)_i)$ *and* $(D', (\gamma'_i)_i)$ *at* v *and* v' is the combinatorial geodesic $(n + n' - 2)$–gon whose underlying CW-complex is the wedge sum of D and D' at v and v', and whose sides are:

$$\gamma_0, \dots, \gamma_{i-1}, \gamma_i\gamma'_{i'+1}, \gamma'_{i'+2}, \dots, \gamma'_{i'-1}, \gamma'_{i'}\gamma_{i+1}, \gamma_{i+1}, \dots, \gamma_{n-1}$$

This is the unique way to define the sides so that the canonical embeddings of $(D, (\gamma_i)_i)$ and $(D', (\gamma'_i)_i)$ factor through inclusion into the wedge sum and the canonical embedding of the resulting combinatorial geodesic $(n + n' - 2)$–gon.

3.2.2. *Edge reduction.* Suppose $e \subset D$ is an interior edge of D such that the boundary of e is contained in the boundary of D. Suppose further that the only maximal cells that intersect e are the two faces that intersect its interior. The hypotheses imply that e separates D into two components, and the $(3, 7)$–condition implies that e has two distinct boundary vertices. Suppose that each of these boundary vertices belongs to exactly one side, which is necessarily true if $(D, (\gamma_i)_i)$ is non-degenerate. Define the *edge reduction of* $(D, (\gamma_i)_i)$ *at* e to be the two combinatorial geodesic polygons obtained by collapsing e to a vertex and then performing vertex reduction at the resulting vertex.

The inverse operation to edge reduction we denominate *edge combination*. Suppose each of $v \in D$ and $v' \in D'$ is a distinguished vertex that is contained in a single face and in no other maximal cell. First perform a vertex combination at v and v' and then blow up the wedge point to an interior edge, while keeping the same sides. As before, we require that v and v' each belong to a single face and no other maximal cell, which implies that the resulting combinatorial geodesic polygon is uniquely determined.

3.2.3. *Face reduction.* Suppose $\Pi \subset D$ is a face with $e(\Pi) \geqslant 2$. Suppose that there are boundary edges e and e' of Π that are boundary edges of D such that removing the union of the interiors of Π, e, and e' separates D into two components, D_1 and D_2. Suppose that e and e' each intersect only one side of D, which is necessarily true if $(D,(\gamma_i)_i)$ is non-degenerate. Finally, suppose that D_1 and D_2 each contain a distinguished vertex. Define the *face reduction of* $(D,(\gamma_i)_i)$ *at* (Π, e, e'), or at Π, when e and e' are clear, to be the two combinatorial geodesic polygons obtained by subdividing e and e', subdividing Π by adding a new edge connecting the subdivision points of e and e', and then performing an edge reduction on this new edge.

The inverse operation to face reduction, which we denominate *face combination*, is to first perform edge combination, which results in a new interior edge in the boundary of exactly two faces, and then forget this new edge, replacing the two incident faces by a single face, and replacing the four resulting edges by two.

3.3. Combinatorial geodesic bigons and triangles. We state Strebel's classification of combinatorial geodesic bigons and triangles. Let us stress again that we are working in the combinatorial setting, cf Remark 3.6. Strebel's original statement includes that the diagram D comes from a small cancellation presentation, but what is actually used in the proof are the properties of the diagram that we have encapsulated in the definition of 'combinatorial geodesic polygon', Definition 3.3. This observation was first made in [30, Section 2.5 and Remark 4.7] and [27, Lemma 4.7].

THEOREM 3.13 (Strebel's classification[1], [47, Theorem 43]). *Let D be a simple diagram that is not a single face.*

- *If D is a combinatorial geodesic bigon, then D has shape* I_1 *in Figure 5.*
- *If D is a combinatorial geodesic triangle, then D has one of the shapes* I_2, I_3, II, III_1, IV, *or* V *in Figure 5.*

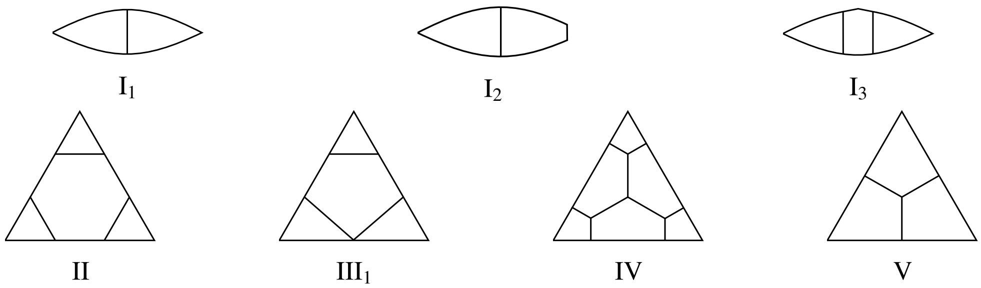

FIGURE 5. Strebel's classification of combinatorial geodesic bigons and triangles.

Note that shapes I_2 and I_3 degenerate to combinatorial geodesic bigons.

Each of these shapes represents an infinite family of combinatorial geodesic bigons or triangles obtained by performing face combination at a non-degenerate, distinguished vertex with a shape I_1 bigon arbitrarily many times. Figure 6 shows alternate examples of each shape.

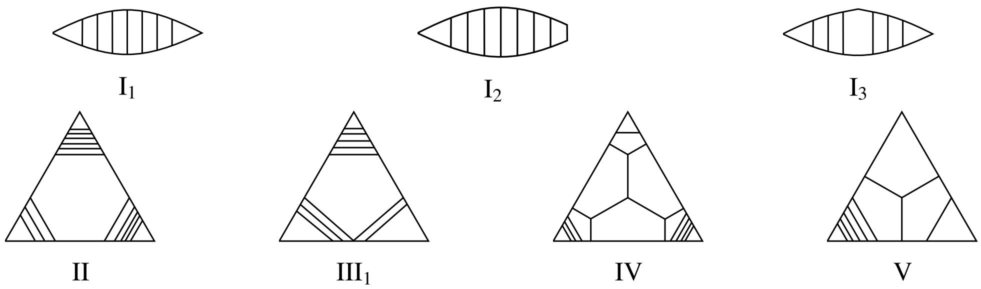

FIGURE 6. Alternate examples of each shape.

[1]Strebel also considers a second definition of combinatorial geodesic polygon that yields one additional shape III_2. This is not relevant for us, but we retain the subscript for shape III_1 for consistency with Strebel's notation.

3.4. Special combinatorial geodesic quadrangles. In this section diagram faces are labelled with their contribution to the curvature sum (Lemma 3.7) if this contribution is non-zero.

Definition 3.14 (Special). A combinatorial geodesic n–gon, for $n > 2$, is *special* if it is simple, non-degenerate, irreducible, and every non-distinguished vertex has valence 3.

The only special combinatorial geodesic triangles are the representatives of shapes IV and V pictured in Figure 5. In this section we classify special combinatorial geodesic quadrangles.

Let D be a special combinatorial geodesic polygon. Simplicity and trivalence imply that $\Pi \cap \partial D$ is a disjoint union of arcs, for each face Π. Irreducibility implies $\Pi \cap \partial D$ consists of at most one arc. Non-degeneracy implies that every distinguished face has exactly one distinguished vertex and either 2 or 3 interior arcs.

The curvature formula of Lemma 3.7 can be simplified as follows.

Lemma 3.15 (Special curvature formula). *Let D be a special combinatorial geodesic polygon that is not a single face. Then:*

$$6 = \sum_{e(\Pi)=0} (6 - i(\Pi)) + \sum_{e(\Pi)=1} (4 - i(\Pi))$$

Proof. By definition, every non-distinguished vertex has valence 3. To apply Lemma 3.7, we must address the possible existence of degree 2 vertices in the boundary: we may iteratively remove such vertices, always replacing the two adjacent edges by a single edge. Since D is not a single face, this makes sense for every degree 2 vertex, and we thus remove all degree 2 vertices. Since the degree of a face counts *arcs*, not edges, the operation does not alter the sum. The formula follows by applying Lemma 3.7, and noting, as a consequence of irreducibility, that every face has exterior degree at most 1. □

For the remainder of this section, let D be a special combinatorial geodesic quadrangle. Then D has exactly four distinguished faces, each of which contributes either 1 or 2 to the curvature sum, and every other face makes a non-positive contribution. Let D_k refer to the set of distinguished faces of D that contribute k to the curvature sum, ie, with $e(\Pi) = 1$ and $i(\Pi) = 4 - k$.

An *ordinary* face will refer to a non-distinguished face Π with $e(\Pi) = 1$ and $i(\Pi) = 4$, which contributes 0 to the curvature sum. An *extraordinary* face will refer to a non-distinguished face Π with $e(\Pi) = 0$ or with $e(\Pi) = 1$ and $i(\Pi) > 4$. Note that if $i(\Pi) > 6$, we must have $e(\Pi) = 0$.

3.4.1. *Zippers.* Ordinary faces can fit together to make arbitrarily long sequences of subsequent faces we call *zippers*, as in Figure 7.

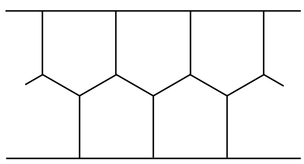

Figure 7. A zipper.

The ordinary faces in a zipper are called *teeth*. We define a *zipper Z with zero teeth* to be three consecutive interior edges that separate the diagram into two parts, each of which contains two distinguished faces. For example, the bold edges of Figure 8 form a zipper with zero teeth. Since D is special, the two interior edges incident to the interior vertices of Z and not belonging to Z must be contained in opposite complementary components of Z, otherwise D would admit a face reduction.

Using a symmetry argument, we show that the portions of the diagram on opposite sides of a zipper each contribute 3 to the curvature sum: first, consider a zipper Z of length 0. Then the two interior edges incident at Z are not on the same side of Z, for otherwise we would have a face with exterior degree at least 2 (and hence face-reducibility) or only one distinguished face on that side. Now assume that one of the two sides S contributes $k \neq 3$ to the curvature sum.

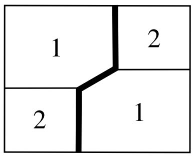

FIGURE 8. Zipper with zero teeth (in bold).

Then we may rotate a copy of S by 180 degrees and attach it to S by identifying the respective copies of Z, thus obtaining a special combinatorial quadrangle for which the curvature formula amounts to $2k \neq 6$; a contradiction. The case of an arbitrary zipper Z now follows similarly by attaching a rotated (or in the case that Z has an odd number of faces reflected) copy of S to $S \cup Z$.

We need to see how to terminate a zipper. Let Π be the face with two edges on the zipper. If Π is ordinary then the zipper just gets longer, so assume not. One possibility is that Π is distinguished, in which case it is a D_1 and there is only one other face, which is a D_2. Otherwise, since the end of the zipper containing Π must contribute 3 to the curvature sum, we must have $i(\Pi) = 5$ and $e(\Pi) = 1$, both distinguished faces must be D_2's, and the other two faces sharing edges with Π must be ordinary.

The two possibilities are shown in Figure 9. In conclusion:

LEMMA 3.16. *There are six infinite families of configurations of special combinatorial geodesic quadrangles containing zippers, determined by the choice of two zipper ends from Figure 9 and the parity of the number of teeth.*

(1) Zipper end 1. (2) Zipper end 2.

FIGURE 9. Terminating a zipper

3.4.2. *Extraordinary configurations.*

LEMMA 3.17. *The six configurations shown in Figure 10 are the only special combinatorial geodesic quadrangles containing an extraordinary face and no zipper.*

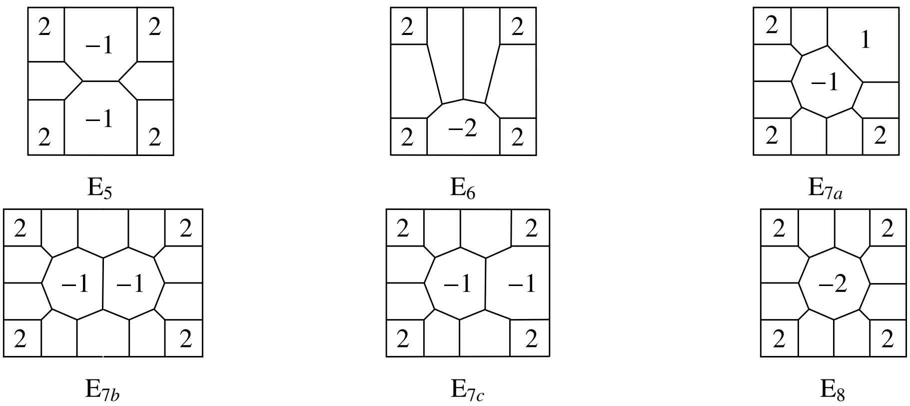

FIGURE 10. Extraordinary special combinatorial geodesic quadrangles.

PROOF. Let D be a special combinatorial geodesic quadrangle without zippers. First, suppose that D contains no interior faces.

Case E_5: D contains an extraordinary face Π with $i(\Pi) = 5$. In this case Π contributes -1 to the curvature sum. Consider the third interior edge e of Π. Let Π' be the face on the opposite side of e. If e is the second interior edge of Π' then we get a zipper, so it must be at least the third. (It cannot be the first, for this would give a vertex of degree at least 4.)

By symmetry, we see that $i(\Pi') \geqslant 5$. As Π' contributes at least -1 to the curvature sum and is not interior, we deduce $i(\Pi') = 5$ and $e(\Pi') = 1$, with e the third interior edge of Π'. This then implies that the distinguished faces are all D_2's, and every other face is ordinary. There are two ordinary faces bordering both Π and Π'. Since no face has more than one exterior edge, we must then fill in the four D_2's on the corners.

Case E_6: D contains an extraordinary face Π with $i(\Pi) = 6$. In this case Π contributes -2 to the curvature sum, so the distinguished faces are all D_2's, and every other face is ordinary. Consider the third interior vertex of Π. Let e be the edge incident to this vertex that does not belong to Π. If e does not have a vertex on the boundary then there is a zipper contained in the boundary of $\Pi \cup \Pi'$, where Π' is either one of the ordinary faces with side e. Since we have assumed no zippers, e does have a boundary vertex, and there is a unique way to fill in the rest of D with ordinary faces and D_2's.

Now we move on to the interior face cases. According to the curvature formula of Lemma 3.7, interior faces have either 7 or 8 sides.

Case E_8: D contains an extraordinary face Π with $i(\Pi) = 8$. In this case, $e(\Pi) = 0$, and Π contributes -2 to the curvature sum, so all four distinguished faces are D_2 and every other face is ordinary. Since the interior sides of a D_2 have a vertex on the boundary, they cannot share an edge with Π, so every face sharing an edge with Π is ordinary. There is only one way to pack 8 ordinary faces around Π, up to symmetry, and this determines the placement of the four D_2's.

Case E_7: D contains an extraordinary face Π with $i(\Pi) = 7$. In this case, $e(\Pi) = 0$, and Π contributes -1 to the curvature sum. The distinguished faces are therefore either three D_2's and one D_1 or four D_2's. If there are three D_2's and one D_1 then every other face is ordinary. A D_2 cannot share an edge with an interior face, so Π has at least 6 edges that are shared by ordinary faces.

If there are four D_2's then none of them share a face with Π, and there is exactly one other face that is not ordinary.

In either case, Π has at least 6 edges that are shared by ordinary faces. Let $\Pi_1, \dots, \Pi_7$ be the consecutive faces sharing an edge with Π, and assume all except possibly Π_7 are ordinary. Let e be the edge shared by Π_1 and Π_7. If e has a vertex on the boundary then there is only one way to fit 6 ordinary faces around Π. In this case, Π_7 is a D_1, and the configuration is shown in Figure 11(1).

If e does not have a vertex on the boundary then, again, there is only one way to fit six ordinary faces around Π, shown in Figure 11(2). We see that $i(\Pi_7) \geqslant 5$. Since Π_7 contributes at least -1 to the curvature sum, the two possibilities are $i(\Pi_7) = 7$ and $e(\Pi_7) = 0$ or $i(\Pi_7) = 5$ and $e(\Pi_7) = 1$. In both cases, the remaining distinguished faces are D_2's, all other faces are ordinary, and there is a unique way to complete the 4–gon. These are types E_{7b} and E_{7c}, respectively, of Figure 10. □

(1) Case e does have boundary vertex. (2) Case e does not have boundary vertex.

FIGURE 11. Interior 7–gon.

3.4.3. *Classification of special combinatorial geodesic quadrangles.*

THEOREM 3.18. *Every special combinatorial geodesic quadrangle is either one of the six extraordinary configurations of Lemma 3.17 or belongs to one of the six zippered families of Lemma 3.16.*

The theorem is proven by Lemma 3.17, Lemma 3.16, and the following:

LEMMA 3.19. *Every special combinatorial geodesic quadrangle contains an extraordinary face or a zipper.*

PROOF. Suppose D is a special combinatorial geodesic quadrangle that does not contain an extraordinary face. Then no face makes a negative contribution to the curvature sum, so D is composed of two D_2's, two D_1's, and some number of ordinary faces.

Pick a side of the quadrangle. Let S be the union of faces along the side. Let A be the union of interior edges separating S from $D \setminus S$. Since every vertex has valence 3, both S and $D \setminus S$ are connected, and each contains two distinguished faces.

Consider the edges incident to interior vertices of A. Each one is contained either in S or in $D \setminus S$. At least one is contained in S and one in $D \setminus S$, since each side contains two distinguished vertices and no face separates D. Two consecutive edges cannot point into $D \setminus S$, because the face $\Pi \subset S$ containing the edge between them would either be extraordinary, contradicting the hypothesis, or a distinguished face with interior degree at least four, contradicting non-degeneracy.

Two consecutive edges cannot point into S, because the face $\Pi \subset S$ between them would be non-distinguished with $i(\Pi) = 3$. Therefore the edges along A alternate, and A has length at least 3.

If A consists of 3 edges then it is a zipper, and we are done, so suppose it consists of at least 4 edges.

If the first face Π adjacent to A on the $D \setminus S$ side is distinguished then we get a zipper, so suppose it is not. Let Π' be the next face along A on the $D \setminus S$ side. The two possibilities are shown in Figure 12(1) and Figure 12(2).

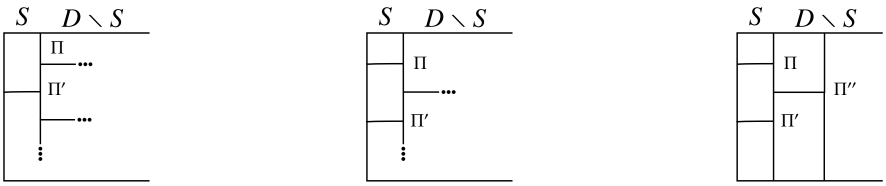

(1) Case 1 (2) Case 2 (3) Case 2 - next level

FIGURE 12. Quadrangle with no zipper or extraordinary face.

Since Π is not distinguished, the edge shared by Π and Π' does not contain a boundary vertex. In the first case $i(\Pi') \geqslant 5$, using that A has length at least 4, contrary to hypothesis. In the second case, $i(\Pi) \geqslant 4$ and $i(\Pi') \geqslant 4$, so both are ordinary, and we have the situation in Figure 12(3). Let Π'' be the face of $D \setminus S$ adjacent to Π and Π'. Either Π'' contains more than one boundary arc, or Π'' is distinguished and contains two distinguished vertices, but both of these are contrary to hypothesis. □

3.5. Quadrangle dichotomy. The following proposition says that a combinatorial geodesic quadrangle must be either *short*, conditions (1) or (2), or *thin*, (3).

PROPOSITION 3.20. *Let D be a simple combinatorial geodesic quadrangle with boundary path $\gamma_1\delta_1\gamma_2^{-1}\delta_2^{-1}$. Then one of the following holds:*

(1) *There exists a face Π that intersects both γ_1 and γ_2 in edges with $2 = e(\Pi) = i(\Pi)$.*
(2) *There exist faces Π intersecting γ_1 in edges, and Π' and Π'', both intersecting γ_2 in edges, such that $1 = e(\Pi) = e(\Pi') = e(\Pi'')$, $4 = i(\Pi) = i(\Pi') = i(\Pi'')$, and any two of Π, Π' and Π'' pairwise intersect in edges. Moreover, $\Pi' \cap \Pi''$ is an arc connecting γ_2 to Π, and $\Pi \cap (\Pi' \cup \Pi'')$ is connected.*

(3) There exist $k \leqslant 6$ and faces $\Pi_1, \Pi_2, \ldots, \Pi_k$, each intersecting γ_1 in edges, such that $\Pi_1 \cap \delta_1 \neq \emptyset$ and $\Pi_k \cap \delta_2 \neq \emptyset$ and, for each $1 \leqslant i < k$, we have $\Pi_i \cap \Pi_{i+1} \neq \emptyset$.

The three cases are pictured in Figures 13(1), 13(2), and 13(3), respectively.

Lemma 3.10 implies that in case (3) of Proposition 3.20 the set $\bigcup_{i=1}^{k} \Pi_i \cap \gamma_1$ is a path subgraph of γ_1.

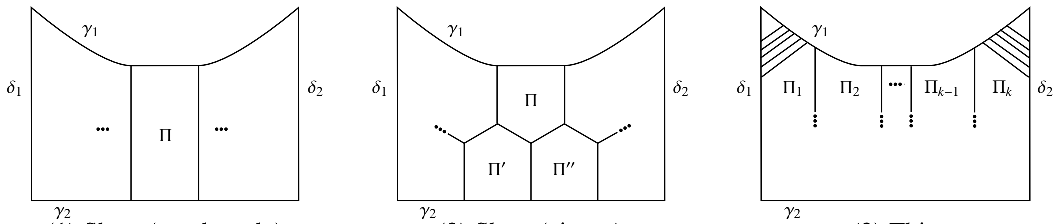

(1) Short (quadrangle) (2) Short (zipper) (3) Thin

FIGURE 13. Quadrangle dichotomy

PROOF OF PROPOSITION 3.20. If D is a single face we are done, so suppose not.

Suppose that the first two conditions do not hold. Let $k > 0$ be minimal such that $\Pi_1, \ldots, \Pi_k$ is a sequence of faces satisfying condition (3). Minimality is equivalent to requiring $\Pi_i \cap \delta_1 = \emptyset$ and $\Pi_i \cap \delta_2 = \emptyset$ for all $1 < i < k$.

We may also assume D is non-degenerate, since for triangles we have $k \leqslant 3$, by Theorem 3.13.

Since D is a simple $(3,7)$–diagram it has no valence 1 vertices. If it has non-distinguished valence 2 vertices we can forget them. We then arrange that non-distinguished vertices of D have valence 3 as follows.

Suppose v is a vertex of valence greater than 3. Since D is simple, for any two faces Π and Π' that contain v in their boundaries and do not share an edge incident to v, there is a unique way to blow up v to an interior edge e_v in such a way that the images of Π and Π' are the faces with e_v in their boundaries. Let β denote the blow-up map. Since Π and Π' do not share an edge incident to v, the two vertices of e_v each have valence at least 3 and strictly less than that of v.

The only effect of this blow-up on faces is to increase the interior degrees of Π and Π' by 1 each, so $\beta(D)$ is still a simple combinatorial geodesic quadrangle. The faces $\beta(\Pi_1), \ldots, \beta(\Pi_k)$ still satisfy condition (3), since β only introduces an interior edge. We now check that β does not produce a diagram satisfying conditions (1) or (2) and not condition (3).

We argue for the face Π. The same arguments apply for Π'. Since D is simple Π cannot have $e(\Pi) = 2$ and $i(\Pi) = 1$, so $\beta(\Pi)$ does not satisfy condition (1).

If $e(\Pi) > 1$ then $e(\beta(\Pi)) = e(\Pi) > 1$, so $\beta(\Pi)$ cannot be one of the faces satisfying condition (2). Suppose now that $e(\Pi) = 1$. The blow-up does not change the boundary of D, so Π is distinguished if and only if $\beta(\Pi)$ is distinguished. If Π is distinguished with $i(\Pi) \geqslant 3$ then $\beta(\Pi)$ is distinguished with $i(\beta(\Pi)) \geqslant 4$, so $\beta(D)$ degenerates to a triangle, which implies $k \leqslant 3$, and we are done. Otherwise, if Π is distinguished then $i(\beta(\Pi)) < 4$, so $\beta(\Pi)$ cannot be one of the faces satisfying condition (2). Finally, if Π is non-distinguished then the combinatorial geodesic polygon condition requires $i(\Pi) \geqslant 4$, so $i(\beta(\Pi)) \geqslant 5$, and $\beta(\Pi)$ cannot be one of the faces satisfying condition (2).

We conclude that $k \leqslant 3$, in which case we are done, or $\beta(D)$ is a simple non-degenerate combinatorial geodesic quadrangle not satisfying conditions (1) or (2) and containing a sequence $\beta(\Pi_1), \ldots, \beta(\Pi_k)$ of faces satisfying the requirements of condition (3). Moreover, since β only introduced an interior edge, $\beta(\Pi_i) \cap \beta(\delta_j) = \emptyset$ for all $1 < i < k$ and $j \in \{1, 2\}$, so $\beta(\Pi_1), \ldots, \beta(\Pi_k)$ is a minimal length sequence satisfying condition (3) in $\beta(D)$.

We now repeat blowing up higher valence vertices until either the quadrangle becomes degenerate, and we are done, or there are no higher valence vertices left. Thus, we may assume non-distinguished vertices of D have valence 3.

If D is irreducible then it is special. If D is special and has no zipper then $k \leqslant 5$, by considering the possibilities given by Theorem 3.18.

If D is special and has a zipper then, by considering the possible zippered configurations of Lemma 3.16, k is at most four plus the number of teeth of the zipper adjacent to γ_1. Since D does not satisfy condition (2), there can be at most two teeth of the zipper adjacent to γ_1, so $k \leqslant 6$. In fact, $k = 6$ can occur, so 6 is the best possible bound for k.

The remaining possibility is that D is reducible. It is not vertex reducible, since it is simple. If an edge reduction is possible then so is a face reduction, so suppose Π is a face with edges e, $e' \subset \partial\Pi$ such that D admits a face reduction at (Π, e, e'). Since D is non-degenerate there are four distinguished faces.

Suppose the face reduction separates one distinguished face from the other three, which occurs when one of e or e' is an edge of one of the γ_i's and the other is an edge of one of the δ_i's. Then D is a union of a bigon and a quadrangle with fewer faces than D. By minimality of k, either $\Pi = \Pi_1$ or $\Pi = \Pi_k$ or $\Pi \neq \Pi_i$. Therefore $\Pi_1, \ldots, \Pi_k$ corresponds to a sequence of faces in the new quadrangle still satisfying condition 3.

Repeat the argument for the new quadrangle. Since the number of faces in the quadrangle decreases, this process stops after finitely many steps, so we may assume that D is not reducible into a bigon and a quadrangle. Since D is non-degenerate, this implies that every distinguished face has one exterior arc and either 2 or 3 interior arcs, and the distinguished faces are the only ones that intersect both one of the δ_i's and one of the γ_i's.

Suppose there is a face Π containing edges $e \subset \partial\Pi \cap \delta_1$ and $e' \subset \partial\Pi \cap \delta_2$. Face reduction at (Π, e, e') sends the Π_i's to a minimal length sequence of faces along one side of a combinatorial geodesic triangle connecting the other two sides. Thus, $k \leqslant 3$.

Next, suppose there is an edge e that meets both γ_1 and γ_2. Then e separates δ_1 from δ_2, so there is some $i < k$ such that $\Pi_1, \ldots, \Pi_i$ are on one side of e and $\Pi_{i+1}, \ldots, \Pi_k$ are on the other. Edge reduction of D at e results in two combinatorial geodesic triangles. In one of these triangles there is a sequence of faces corresponding to $\Pi_1, \ldots, \Pi_i$ that run along the side corresponding to γ_1 and connect the side corresponding to δ_1 to the opposite distinguished vertex. Since we were not in case (1), $i(\Pi_i) > 2$, so the face corresponding to Π_i in the triangle has more than one interior arc. It follows that $i \leqslant 3$. The same argument applied to Π_{i+1} shows $k - i \leqslant 3$, so $k \leqslant 6$. (The lower right diagram in Figure 4 shows $k = 6$ can be achieved in this case.)

Finally, suppose D has a separating face Π_i and no separating edge. Face reduction of D at Π_i yields two combinatorial geodesic triangles. Since D has no separating edge, each of these triangles has a distinguished face corresponding to Π_i with interior degree greater than one. The same argument as in the previous case then tells us $i \leqslant 3$ and $k - i + 1 \leqslant 3$, so $k \leqslant 5$. □

4. Contraction in $Gr'(1/6)$–groups

In this section, we study contraction in groups defined by $Gr'(1/6)$–labelled graphs. We give a characterization of strongly contracting geodesics in the Cayley graph in terms of their intersections with embedded components of the defining graph. Throughout this section, Γ is an arbitrary $Gr'(1/6)$–labelled graph with a set of labels $\mathcal{S}$, and $X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$.

Recall that a geodesic α in X is *locally (ρ_1, ρ_2)–contracting* if for every embedded component Γ_0 of Γ such that $\Gamma_0 \cap \alpha \neq \emptyset$, closest point projection in Γ_0 of Γ_0 to $\Gamma_0 \cap \alpha$ is (ρ_1, ρ_2)–contracting.

We will prove the following theorem:

THEOREM 4.1. *A geodesic α in X is locally (ρ_1, ρ_2)–contracting if and only if there exist ρ_1' and ρ_2' such that α is (ρ_1', ρ_2')–contracting.*

Moreover, ρ_1' and ρ_2' can be bounded in terms of ρ_1 and ρ_2, and when $\rho_1(r) \geqslant r/2$ we can take $\rho_1' = \rho_1$ and $\rho_2' \asymp \rho_2$.

Consequently, α is strongly contracting if and only if it is uniformly locally strongly contracting, α is semi-strongly contracting if and only if it is locally uniformly semi-strongly contracting, and α is sublinearly contracting if and only if it is uniformly locally sublinearly contracting.

REMARK 4.2. Theorem 4.1 provides an analogy with the geometry present in many of the preeminent examples of spaces with a mixture of hyperbolic and non-hyperbolic behavior, such as Teichmüller space, the Culler-Vogtmann Outer Space, or relatively hyperbolic spaces. The simplest situation is that geodesics that avoid spending a long time in a non-hyperbolic region behave like hyperbolic geodesics. The analogy here is strongest with relatively hyperbolic spaces, with the embedded components of the defining graph corresponding to the peripheral regions of a relatively hyperbolic space. The peripheral regions are not necessarily non-hyperbolic, so hyperbolic geodesics do not necessarily have to avoid them completely. Rather, a geodesic is roughly as hyperbolic as its intersections with peripheral regions.

Our original motivation for the present paper was to make this analogy precise, and, in particular, to determine whether graphical small cancellation groups contained strongly contracting elements, and therefore fit into the scheme of groups with growth tight actions introduced by Arzhantseva, Cashen, and Tao [4], see Section 5.

In concrete terms, this analogy was first suggested by results of Gruber and Sisto, who proved that the Cayley graph of a graphical small cancellation group is weakly hyperbolic relative to the embedded components of the defining graph [30] in the sense of [18], see also Section 5.2.

Note, however, that in general the groups we consider in the present paper need not be non-trivially relatively hyperbolic [11, 30].

4.1. Contraction and Morse quasi-geodesics. We quote some of our technical results from [3] that let us simplify and reformulate the contraction conditions.

LEMMA 4.3 ([3, Lemma 6.3]). *Let Y and Y' be closed subspaces of X at bounded Hausdorff distance from one another. Suppose Y is (r,ρ_2)–contracting. Then Y' is (r,ρ_2')–contracting for some $\rho_2' \asymp \rho_2$. In particular, if Y is strongly contracting then so is Y'.*

THEOREM 4.4 ([3, Theorem 1.4, Theorem 7.1]). *Let Y be a closed subspace of X. Then the following are equivalent:*

(1) Y is Morse.
(2) Y is sublinearly contracting.
(3) There exist ρ_1 and ρ_2 such that Y is (ρ_1,ρ_2)–contracting.
(4) There exist a constant $C \geqslant 0$ and a sublinear function ρ such that if γ is a geodesic segment then $d(\gamma, Y) \geqslant C$ implies $\operatorname{diam}\pi(\gamma) \leqslant \rho(\max_{z\in\gamma} d(z, Y))$.

Moreover, for each implication the function of the conclusion depends only on the function of the hypothesis, not on Y.

We remark that in the case $\rho_1(r) := r$ the proof of '(3) implies (4)' yields $\rho \asymp \rho_2$. When, in addition, ρ_2 is bounded, which is the strongly contracting case, this recovers the well-known 'Bounded Geodesic Image Property', cf [35, 13].

COROLLARY 4.5. *If Y is strongly contracting then every geodesic segment that stays sufficiently far from Y has uniformly bounded projection diameter.*

4.2. Proof of Theorem 4.1.

LEMMA 4.6. *Let α be a geodesic in X, and let x be a vertex not in α. Let γ be a path from x to a vertex of α such that $|\gamma| = d(x,\alpha)$. Then, if p is a path from x to a vertex of α such that p is a piece, then $p = \gamma$. In particular, the closest point projection of p to α is $p \cap \alpha$.*

PROOF. If p has the same terminal vertex as γ and if $p \neq \gamma$, then there exist subpaths γ' of γ and p' of p, each of length at least 1, such that $c := \gamma' p'^{-1}$ is a simple cycle. Since γ' is a geodesic, we have $|\gamma'| \leqslant |c|/2$. Since p' is a piece we have $|p'| < |c|/6$. This is a contradiction.

If p has a different terminal vertex, then there exist terminal subpaths p' of p and γ' of γ, respectively, each having length at least 1, and a path α' contained in α such that $c := \gamma'\alpha' p'^{-1}$ is a simple cycle. By assumption on γ, we have $|\gamma'| \leqslant |p'|$ and $|\alpha'| \leqslant |c|/2$, whence we conclude $|p'| \geqslant |c|/4$, contradicting the fact that p' is a piece.

The final claim follows from the fact that a subpath of a piece is a piece. □

LEMMA 4.7. *Let Γ_0 be an embedded component of Γ in X, and let α be a geodesic in X such that Γ_0 intersects α. Closest point projection of Γ_0 to α in X agrees with closest point projection to $\Gamma_0 \cap \alpha$ in Γ_0.*

PROOF. Consider a vertex $x \in \Gamma_0 \setminus \alpha$ such that there exists a point $v \in \alpha$ with $d(x, \alpha) = d(x, v)$ and $v \notin \Gamma_0$. Let γ be a geodesic from x to v, let p be a path in Γ_0 from x to a vertex v' of α. Let D be a diagram over $\mathcal{R}$ as in Lemma 3.4 filling the triangle $\gamma[v, v']p^{-1}$, where $[v, v']$ is the reduced path in α from v to v'.

Among all possible choices of x, v, v', γ, p, and D as above, make the choice for which D has the minimal possible number of edges. Note that, by minimality, D is a simple disc diagram.

Let Π be a face of D that intersects the side of D corresponding to p in an arc a. Then a has a lift to Γ via being a subpath of (a copy of) p, and one via being a subpath of $\partial\Pi$. If the lifts coincide (up to a label-preserving automorphism of Γ), then we can remove the edges of a from D, thus obtaining a path p' in Γ_0 as above, contradicting minimality. Hence, a is a piece. Therefore, Lemma 4.6 implies that the side of D corresponding to p is not contained in a single face of D.

We make a diagram D' by attaching a new face Π' to D by identifying a proper subpath of the boundary of Π' with the side of D corresponding to p. This operation is purely combinatorial, the boundary of Π' is not labelled. Note that if Π is a face of D with $e(\Pi) = 1$ whose exterior arc is contained in p then, by the previous paragraph, that exterior arc is a piece. Since interior arcs of D are pieces, $i(\Pi) \geqslant 6$ in D. Thus, Π becomes an interior face of D' with $i(\Pi) \geqslant 7$. It follows that D' is a combinatorial geodesic bigon.

Apply Theorem 3.13 to D': it has at most two distinguished faces, one is Π' and the other, if it exists, is at the vertex corresponding to v. Any other face of D' came from D, and, in particular, the side of D corresponding to p is contained in a single face of D. This is a contradiction. □

Lemma 4.7 immediately implies the global-to-local direction of Theorem 4.1.

LEMMA 4.8. *Let α be a geodesic in X. Let $p_1 p_2$ be a simple path starting at a vertex y in α and terminating at a vertex $x \in X$ such that $p_1 p_2$ is contained in some embedded component Γ_0 and such that each p_i is a piece that is not a single vertex. Then $d(x, \alpha) > \max\{|p_1|, |p_2|\}$.*

PROOF. Let q be a path starting at x and terminating at a vertex y' in α such that $|q| = d(x, \alpha)$. If $y' = y$, then the claim follows from the convexity of Γ_0, noting that any simple path that is a concatenation of two pieces must be a geodesic. Similarly, if $q \cap p_1 \neq \emptyset$ then Lemma 4.6 implies $q = p_1 p_2$. Hence, assume that $q \cap p_1 = \emptyset$.

Without loss of generality, we may assume $q \cap p_2 = x$. Consider a diagram D over $\mathcal{R}$ as in Lemma 3.4 filling the embedded geodesic quadrangle $p_1 p_2 q[y', y]$. We may stick onto the 2-complex D two 2-cells Π_1 and Π_2 by identifying proper subpaths of their boundaries with p_1 and p_2, respectively. By construction, this yields a combinatorial geodesic triangle, and by Theorem 3.13, it has shape III_1, where Π_1 and Π_2 are the distinguished faces intersecting each other only in a valence 4 vertex. Using Lemma 4.6, there are no faces with exterior degree 2, so D is a single face.

This shows that there exists a simple cycle of the form $c := p_1 p_2 q[y', y]$ in some embedded component Γ_1. Since $|p_1|$, $|p_2| < |c|/6$ and $|[y', y]| \leqslant |c|/2$, we have $|q| > |c|/6 > \max\{|p_1|, |p_2|\}$. □

LEMMA 4.9. *Let α be a geodesic in X. Let Γ_0 be an embedded component of Γ intersecting α. Suppose $\alpha \cap \Gamma_0$ is (ρ_1, ρ_2)–contracting in Γ_0. Let c be a simple cycle in Γ_0 such that $c = p_1 p_2 q a$, where each p_i is a piece, q realizes the distance of the terminal vertex x of p_2 to α, and a is a subpath of α. Then $|c|$ is bounded, with bound depending only on ρ_1 and ρ_2.*

PROOF. Let $x := p_2 \cap q$, and let $y := p_2 \cap p_1$. Since any subpath of a piece is a piece itself, we may assume without loss of generality that x is the point of p_2 maximizing $\operatorname{diam} \pi(x) \cup \pi(y)$, where π denotes closest point projection to α in X. By Lemma 4.6, we have $|a| \leqslant \operatorname{diam} \pi(x) \cup \pi(y)$.

Since $p_1 p_2$ is a path from α to x and q minimizes distance, $|q| \leqslant |p_1| + |p_2| < |c|/3$, which implies $|a| = |c| - |p_1| - |p_2| - |q| > |c|/3$.

Let $C \geqslant 0$ be as Theorem 4.4 (4). If $|p_1| < C$, then $|q| \leqslant |p_1| + |p_2| < C + {}^{|c|}/_{6}$, which implies $|c| < 12C$.

Otherwise, Lemma 4.8 implies that for every point $z \in p_2$ we have $d(z, \alpha) \geqslant C$, so p_2 is a geodesic that stays outside the C–neighborhood of α. Theorem 4.4 (4) says $\operatorname{diam} \pi(p_2)$ is bounded by a sublinear function of $\max_{z \in p_2} d(z, \alpha) < {}^{|c|}/_{3}$. Thus, ${}^{|c|}/_{3} < |a| \leqslant \operatorname{diam} \pi(p_2)$ is bounded above by a sublinear function of $|c|$, which implies $|c|$ is bounded. □

Lemma 4.10. *Let α be a geodesic in X. Let Γ_0 be an embedded component of Γ intersecting α. Suppose $\alpha \cap \Gamma_0$ is (ρ_1, ρ_2)–contracting in Γ_0. Let c be a simple cycle in Γ_0 of the form $c = p_1 p_2 p_3 p_4 a$, where each p_i is a piece and a is a subpath of α. Then $|c|$ is bounded, with bound depending only on ρ_1 and ρ_2.*

Proof. Let $x := p_2 \cap p_3$. Let $y_1 := p_1 \cap \alpha$. Let $y_2 := p_4 \cap \alpha$. By Lemma 4.7, there is a path $q \subset \Gamma_0$ such that $|q| = d(x, \alpha)$. Let $y' := q \cap \alpha$.

By symmetry, we may suppose $q \cap p_3 = x$. Apply Lemma 4.9 to see that $|q| + |p_3| + |p_4| + d(y_2, y')$ is uniformly bounded. If q coincides with $p_1 p_2$ we are done. Otherwise it must be that $q \cap p_1 = \emptyset$, for otherwise we would have a simple cycle composed of a geodesic and one or two pieces, which is impossible. Let $p_2' := \overline{p_2 \setminus q}$, and let $q' := \overline{q \setminus p_2}$. Apply Lemma 4.9 to see $|q'| + |p_2'| + |p_1| + d(y_1, y')$ is uniformly bounded. Thus, $|c|$ is uniformly bounded. □

Lemma 4.11. *Let α be a geodesic in X. Let Γ_0 be an embedded component of Γ intersecting α. Suppose $\alpha \cap \Gamma_0$ is (ρ_1, ρ_2)–contracting in Γ_0. Let c be a simple cycle in Γ_0 such that $c = q_1 p q_3 a$, where p is a piece, q_1 and q_3 are geodesics realizing the closest point projections of the endpoints of p to α, and a is a subpath of α. Then there is a sublinear function ρ_2' depending only on ρ_1 and ρ_2 such that* $\operatorname{diam} \pi(c) \leqslant \rho_2'(|c|)$. *If $\rho_1(r) \geqslant {}^{r}/_{2}$ then we can take $\rho_2' := 2\rho_2$.*

Proof. Among all path subgraphs of p, consider those with maximal projection diameter. Among those, choose one, p', with minimal length, and let x' and y' be its endpoints. Let $x'' \in \pi(x')$ and $y'' \in \pi(y')$ be vertices such that $d(x'', y'') = \operatorname{diam} \pi(x') \cup \pi(y')$. By Lemma 4.7 there are geodesics $q_1' \subset \Gamma_0$ connecting x' to x'' and $q_3' \subset \Gamma_0$ connecting y' to y''. Let a' be the subpath of α from x'' to y''. Let $c' := q_1' p' q_3' a'$. The maximality hypothesis on p' implies $\pi(p') \subset a'$ and $\operatorname{diam} \pi(p') = \operatorname{diam} \pi(p)$. It is immediate from the definitions that $\operatorname{diam} \pi(c) = \operatorname{diam} \pi(p)$, so it suffices to bound $d(x'', y'')$.

Let x and y be the endpoints of p, and note that $|q_1'| + |q_3'| \leqslant |q_1| + d(x, x') + d(y, y') + |q_3| \leqslant |c|$. If there exists a vertex $z \in q_1' \cap q_3'$ then both x'' and y'' are in $\pi(z)$, so $d(x'', y'') \leqslant \rho_2(d(z, \alpha)) \leqslant \rho_2(|q_1'|) \leqslant \rho_2(|c|)$, and we are done. Otherwise, minimality of $|p'|$ implies c' is a simple cycle.

Since p' is a piece and a' is geodesic, we have $|q_1'| + |q_3'| > {}^{|c'|}/_{3} > 2|p'|$, so there exists a point $z' \in p'$ such that $d(x', z') \leqslant {}^{|q_1'|}/_{2}$ and $d(z', y') \leqslant {}^{|q_3'|}/_{2}$.

If $\rho_1(r) \geqslant {}^{r}/_{2}$, then take $\rho_1' := \rho_1$ and $\rho_2'' := \rho_2$. Otherwise, by Theorem 4.4 (2) there exists a sublinear function ρ_2'' such that $\alpha \cap \Gamma_0$ is (ρ_1', ρ_2'')–contracting in Γ_0 for $\rho_1'(r) := r$. Let $\rho_2' := 2\rho_2''$. Since $\pi(p') \subset a'$, we conclude:

$$\begin{aligned} d(x'', y'') &\leqslant \operatorname{diam} \pi(x') \cup \pi(z') + \operatorname{diam} \pi(y') \cup \pi(z') \\ &\leqslant \rho_2''(d(x', \alpha)) + \rho_2''(d(y', \alpha)) \\ &= \rho_2''(|q_1'|) + \rho_2''(|q_3'|) \leqslant 2\rho_2''(|c|) = \rho_2'(|c|) \end{aligned}$$ □

Lemma 4.12. *Let α be a locally (ρ_1, ρ_2)–contracting geodesic in X. Let $Y \subset X$ be either an embedded component of Γ, a piece, or a single vertex. Then there is a sublinear function ρ_2' depending only on ρ_1 and ρ_2 such that if Y is disjoint from α then* $\operatorname{diam} \pi(Y) \leqslant \rho_2'(d(Y, \alpha))$.

If $\rho_1(r) \geqslant {}^{r}/_{2}$ we can take $\rho_2' \asymp \rho_2$.

Proof. Suppose Y is disjoint from α and choose a vertex $y \in Y$ such that $d(y, \alpha) = d(Y, \alpha)$. Let y' be a point in $\pi(y)$. It suffices to show that there exists a sublinear function ρ_2'' such that for every $x' \in \pi(Y)$ we have $d(x', y') \leqslant \rho_2''(d(Y, \alpha))$. Given such a ρ_2'', set $\rho_2' := 2\rho_2''$, and the lemma follows from the triangle inequality.

If $\pi(Y) = \{y'\}$ we are done. Otherwise, let x' be an arbitrary point in $\pi(Y) \setminus \{y'\}$. Let α' be the subpath of α from x' to y'.

Choose $x \in Y$ such that $x' \in \pi(x)$. Choose a path p from x to y in Y, and geodesics β_1 and β_2 connecting x to x' and y to y', respectively. Choose a diagram D over $\mathcal{R}$ as in Lemma 3.4 filling $\alpha'\beta_2^{-1}p^{-1}\beta_1$. Assume that we have chosen x, p, β_1, β_2, and D so that D has the minimal number of edges among all possible choices.

In the case that β_1 and β_2 intersect, let D' be the disc component of D intersecting the side corresponding to α. Then D' is a combinatorial geodesic triangle. Apply Theorem 3.13 to D'. Since β_1 and β_2 are geodesics realizing closest point projection, the only possibilities are that D' is:

(1) a single face,
(2) shape I_2, where α' is the side joining two vertices in the same distinguished face,
(3) shape IV with exactly three faces incident to α', two corners with interior degree 2 each and one ordinary face, or
(4) shape V with exactly two faces incident to α', the two corners with interior degree 2 each.

In all cases, $d(x', y')$ is bounded by a sublinear function ρ_2'' of $d(Y, \alpha)$ that depends only on ρ_1 and ρ_2: For case (1) this follows from the fact that α is uniformly locally contracting. For case (3) this follows from Lemma 4.9 and Lemma 4.10, and for case (4) this follows from Lemma 4.9. In these two cases, the bounds are in fact constants depending only on ρ_1 and ρ_2. Now consider case (2). Let Π be the face of D' containing α'. Let c be the embedded quadrangle in X whose sides are α', a subpath q_1 of β_1, a piece p', and a subpath q_2 of β_2. Apply Lemma 4.11 to c, and observe that $||q_1| - |q_2|| \leqslant |p'| < |c|/6$, and $|q_1| + |q_2| > |c|/3$, whence $d(Y, \alpha) = |\beta_2| \geqslant |q_2| > |c|/6$.

Now suppose that β_1 and β_2 do not intersect. In this case minimality of D and the fact that y minimizes the distance from Y to α imply that D is simple. If Y is an embedded component of Γ, it follows as in the proof of Lemma 4.7 that any arc in the side of D corresponding to p is a piece. The same is true if Y is a piece since subpaths of pieces are pieces. Thus, we can stick a new face onto p to obtain a combinatorial geodesic triangle D', and we make the same argument as above, noting this time that case (1) cannot hold, for it would imply that Y intersects α. □

Proof of Theorem 4.1. Recall that the global-to-local direction of Theorem 4.1 follows from Lemma 4.7.

Suppose that α is locally (ρ_1, ρ_2)–contracting.

Let x and y be points of X such that $d(x, y) \leqslant \rho_1(d(x, \alpha))$. Let γ be a geodesic from x to y. Let $x' \in \pi(x)$ and $y' \in \pi(y)$ be points realizing $\operatorname{diam} \pi(x) \cup \pi(y)$. Let δ_1 be a geodesic from x' to x, and let δ_2 be a geodesic from y' to y. Let α' be the path subgraph of α from x' to y'.

First, assume that α' does not enter the C-neighborhood of γ, where C is the constant from Theorem 4.4 (4) associated to (ρ_1, ρ_2).

If δ_1 and δ_2 intersect, Lemma 4.12 yields the claim, whence we will assume that they do not. Moreover, by removing initial and terminal subpaths of γ that do not increase the size of the closest point projection to α, we may assume that γ, δ_1, δ_2 and α' can be concatenated to a simple closed path c_0. Let D be a diagram as in Lemma 3.4 for the label of c_0, and, by identifying ∂D with c_0, we can consider γ, δ_1, δ_2 and α' as subpaths of ∂D. As the interior arcs of D are pieces and its four sides are geodesics, we can apply Proposition 3.20.

The first possibility is that there is a face Π with $e(\Pi) = 2$ and $i(\Pi) = 2$. Its boundary is a cycle $c = p_1 q p_3 a$ in some embedded component such that the p_i are pieces, a is a path subgraph of α, and q is a path subgraph of γ. We have $|c|/6 < |a| \leqslant |c|/2$. Therefore, $\max_{z \in c} d(z, \alpha) \leqslant 5|c|/12$.

We now choose $R_2 \geqslant 0$ so that for all $r \geqslant R_2$ we have $2\rho_2(r) < \rho_1(r)$. Suppose γ does not enter the R_2–neighborhood of α.

By Theorem 4.4 (4), $|c|/6 < |a|$ is bounded by a sublinear function of $\max_{z \in c} d(z, \alpha) \leqslant 5|c|/12$. This implies $|c|$ is uniformly bounded, as are $|p_1|$, $|p_3| < |c|/6$. Therefore, γ enters a uniformly bounded neighborhood of α.

The second possibility is that D contains a zipper with two teeth on α. By Lemma 4.10, the boundary lengths of the two teeth on α are uniformly bounded. It follows that the boundary length of the upper tooth is also uniformly bounded, since its intersection with the bottom teeth accounts for more than $1/6$–th of its length. Therefore, γ enters a uniformly bounded neighborhood of α.

Let C' be larger than C, R_2, and the bounds from the first two cases. Let z be the first point of γ such that $d(z, \alpha) \leqslant C'$, if such a point exists. Otherwise let $z := y$. The projection diameter of γ is at most the projection diameter of the path subgraph of γ from x to z plus the projection diameter of the path subgraph from z to y. The latter is at most $4C'$, since every point of this path is within $2C'$ of α, so it suffices to bound the former. Thus, we may assume that the geodesic from x to y does not enter the C' neighborhood of α.

By our choice of C', we conclude that D must fall into the third case of Proposition 3.20. Therefore, there exist $k \leqslant 6$ and a path graph $p := p_0p_1p_2 \dots p_kp_{k+1}$ (recall Lemma 3.10) from δ_1 to δ_2 such that:

- If $1 \leqslant i \leqslant k$, then p_i is a path subgraph of $\gamma \cap \Pi_i$, and
- p_0 is empty or a piece in Π_1, and p_{k+1} is empty or a piece in Π_k.

The second claim follows since, in D, the corner that is separated from the rest of D by removing Π_1 is either empty, a face, or has shape I_1. The same observation holds for the corner at Π_k.

Notice that every point of γ is within $d(x, \alpha) + \rho_1(d(x, \alpha)) \leqslant 2d(x, \alpha)$ of α.

For $1 \leqslant i \leqslant k$, the path graph p_i is a geodesic subsegment of γ that is outside the R_2–neighborhood of α. If p_i is contained in an embedded component Γ_0 of Γ disjoint from α, then $\operatorname{diam} \pi(p_i)$ is bounded by a sublinear function of $d(\Gamma_0, \alpha) \leqslant d(p_i, \alpha) < 2d(x, \alpha)$, by Lemma 4.12. If p_i is not contained in an embedded component disjoint from α, then Theorem 4.4 (4) says $\operatorname{diam} \pi(p_i)$ is bounded by a sublinear function of $\max_{z \in p_i} d(z, \alpha) \leqslant 2d(x, \alpha)$. When $i \in \{0, k+1\}$, the diameter of $\pi(p_i)$ is bounded by a sublinear function of $d(p_i, \alpha)$ by Lemma 4.12.

We have that $\operatorname{diam} \pi(x) \cup \pi(y) \leqslant \sum_{i=0}^{k+1} \operatorname{diam} \pi(p_i)$, and each of the at most 8 terms is bounded by a sublinear function of $d(x, \alpha)$, so $\operatorname{diam} \pi(x) \cup \pi(y)$ is bounded by a sublinear function ρ_2' of $d(x, \alpha)$. Thus, α is (r, ρ_2')–contracting. Moreover, each of the constituent sublinear functions of ρ_2' is determined by ρ_1 and ρ_2 and, when $\rho_1(r) \geqslant r/2$, is either bounded or asymptotic to ρ_2, so ρ_2' depends only on ρ_1 and ρ_2, and we can take $\rho_2' \asymp \rho_2$ when $\rho_1(r) \geqslant r/2$. □

4.3. First applications of Theorem 4.1.

Theorem 4.1 and Theorem 4.4 show:

Theorem 4.13. *A geodesic in X is Morse if and only if it is uniformly locally contracting.*

In the classical small cancellation case we have more explicit criteria:

Corollary 4.14. *Let Γ be a $Gr'(1/6)$–labelled graph whose components are cycle graphs. Let α be a geodesic in X. Define $\rho(r) := \max_{|\Gamma_i| \leqslant r} |\Gamma_i \cap \alpha|$, where the Γ_i range over embedded components of Γ. Then α is Morse if and only if ρ is sublinear, and α is strongly contracting if and only if ρ is bounded.*

In hyperbolic spaces and in CAT(0) spaces Morse geodesics are known to be strongly contracting. In graphical small cancellation groups we build the first examples with a wide range of degrees of contraction:

Theorem 4.15. *Let ρ be a sublinear function. There exists a group G with finite generating set $\mathcal{S}$ and a function $\rho' \asymp \rho$ such that there exists an (r, ρ')–contracting geodesic α in the Cayley graph X of G with respect to $\mathcal{S}$.*

Furthermore, ρ' is optimal, in the following sense: If α is (r, ρ'')–contracting for some ρ'' then $\limsup_{r \to \infty} \frac{\rho''(2r)}{\rho(r)} \geqslant 1$.

If $\lambda > 0$, a *$C'(\lambda)$-collection of words W* is a subset of $\langle \mathcal{S} \rangle$ such that the disjoint union of cycle graphs labelled by the elements of W satisfies the graphical $C'(\lambda)$-condition.

Proof. We may assume ρ is unbounded and integer valued. Since ρ is sublinear, there exists an R such that $5\rho(r) \leqslant 2r$ for all $r \geqslant R$. Let $\mathcal{S} := \{a, b, c\}$. Let $I \subset \{z \in \mathbb{Z} \mid z \geqslant R\}$ be an infinite set such that there exists a $C'(1/12)$-collection $\{w_i\}_{i \in I}$ of words $w_i \in \langle b, c \rangle$ with $|w_i| = 4i$. For $i \in I$, define $R_i := a^{\rho(i)} w_i$.

Let $\Gamma := (\Gamma_i)_{i\in I}$ be a disjoint union of $\mathcal{S}$–labelled cycle graphs, with Γ_i labelled by R_i. Let $G := G(\Gamma)$ and $X := \mathrm{Cay}(G, \mathcal{S})$. There are no non-trivial label-preserving automorphisms of any component Γ_i because of the unique a–labelled path subgraph. There are no non-trivial label-preserving automorphisms of Γ that exchange components since $\{w_i\}_{i\in I}$ have distinct lengths.

If p is a piece contained in Γ_i and labelled by l then l can be written $l = l' + l'' + l'''$ where l' is a suffix of w_i, l'' is a subword of $a^{\rho(i)}$, and l''' is a prefix of w_i. If l' or l''' is non-empty then $l'l'''$ is a piece for $\{w_i\}_{i\in I}$. Therefore $|p| < \rho(i) + |w_i|/12$. Since $5\rho(i) \leqslant 2i$ this implies $|p| < |R_i|/6$, so Γ is $C'(1/6)$–labelled.

Let α be the geodesic with all edge labels a. By construction, α is locally (r, ρ)–contracting, so Theorem 4.1 says there exists $\rho' \asymp \rho$ such that α is (r, ρ')–contracting.

Conversely, if α is (r, ρ'')–contracting then it is locally (r, ρ'')–contracting. By construction, for $i \in I$ there exists a point x_i such that $d(x_i, \alpha) = 2i$ and $\mathrm{diam}\,\pi(x_i) = \rho(i)$, so we must have $\rho(i) \leqslant \rho''(2i)$. □

Theorem 4.16 shows that in classical small cancellation groups the geometry of cycle graphs dictates that only the output contraction function ρ_2 plays a role. In Theorem 4.19 we construct a graphical example for which the input contraction function ρ_1 also carries non-trivial information.

Theorem 4.16. *Let Γ be a $Gr'(1/6)$–labelled graph whose components are cycle graphs. A geodesic α in X that is (ρ_1, ρ_2)–contracting is (r, ρ_2')–contracting for $\rho_2' \asymp \rho_2$.*

Proof. Let Γ_0 be an embedded component of Γ. Since Γ_0 is a cycle graph and $\Gamma_0 \cap \alpha$ is connected, there is a unique point y for which the closest point projection $\pi(y)$ in Γ_0 has positive diameter, and $\mathrm{diam}\,\pi(y) = \mathrm{diam}\,\pi(\Gamma_0 \cap \alpha) \leqslant \rho_2(d(y, \alpha))$.

Let x be a point of Γ_0 equidistant from y and α. Then $d(x, y) = d(x, \alpha)$ and:

$$\mathrm{diam}\,\pi(x) \cup \pi(y) = \mathrm{diam}\,\pi(y) \leqslant \rho_2(d(y, \alpha)) = \rho_2(2d(x, \alpha)) = \rho_2'(d(x, \alpha))$$

This is the worst case, since for x' closer to α the ball of radius $d(x', \alpha)$ about x' does not include y. We conclude that α is locally $(r, \rho_2(2r))$–contracting. Now apply Theorem 4.1 to see that α is (r, ρ_2')–contracting for some $\rho_2'(r) \asymp \rho_2(2r) \asymp \rho_2(r)$. □

Corollary 4.17. *Let Γ be a $Gr'(1/6)$–labelled graph whose components are cycle graphs. A geodesic α in X that is (ρ_1, ρ_2)–contracting with ρ_2 bounded is strongly contracting.*

Having uniformly bounded intersection with every embedded component of Γ is a sufficient condition for a geodesic α in the Cayley graph of a $Gr'(1/6)$–group to be strongly contracting. It is not a necessary condition. Consider the following example:

Example 4.18. We start with a simple example of an infinite classical $C'(1/6)$–small cancellation presentation: $\langle a, b \,|R_n : n \in \mathbb{N}\rangle$ where, for each $n \geqslant 0$ we define:

$$R_n = ab^{20n+1}ab^{20n+2} \cdots ab^{20n+19}ab^{-(20n+20)}$$

The graphs Γ_i for $i \geqslant 1$ are defined by taking the disjoint union of oriented cycles R_n for $(i-1)(i) \leqslant 2n \leqslant -2 + i(i+1)$ and identifying the unique subpath with label b^{20k} in R_{k-1} to the unique subpath in R_k with label b^{20k} which is preceded by a and succeeded by ba. This means we identify the paths labelled by the bold words $ab^{20(k-1)+1}ab^{20(k-1)+2} \ldots a\mathbf{b^{-20k}}$ and $a\mathbf{b^{20k}}ba \ldots ab^{-(20k+20)}$.

The $\{a, b\}$–labelled graph $\sqcup_{i\in\mathbb{N}}\Gamma_i$ satisfies the $Gr'(1/6)$–condition and gives rise to the same Cayley graph as $\langle a, b \,|\, R_n : n \in \mathbb{N}\rangle$. Consider a path α in the Cayley graph labelled by the powers of a. For every i, there exists an embedded copy Γ_i' of Γ_i with $|\alpha \cap \Gamma_i'| = i$. Since paths in Γ_i labelled by powers of a are geodesic, this implies that α is a geodesic. It has intersection of length at most 1 with relators R_n, so it is strongly contracting, but we have $|\alpha \cap \Gamma_i'| = i$.

We also provide the first examples spaces X and $\tilde{X}$ and geodesics γ and $\tilde{\gamma}$ such that there exists a quasi-isometry $X \to \tilde{X}$ mapping γ to $\tilde{\gamma}$ and such that γ is not strongly contracting, but $\tilde{\gamma}$ is strongly contracting.

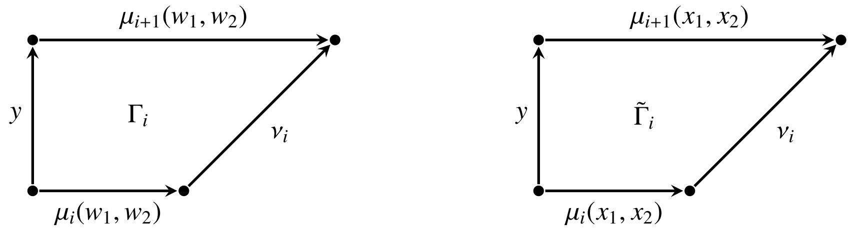

FIGURE 14. The graphs Γ_i and $\tilde{\Gamma}_i$.

THEOREM 4.19 (Non-stability of strong contraction). *There exists a group G with finite generating sets $\mathcal{S} \subset \tilde{\mathcal{S}}$ and an infinite geodesic γ in $X := \mathrm{Cay}(G, \mathcal{S})$ labelled by the powers of a generator such that γ is not strongly contracting, but its image $\tilde{\gamma}$ in $\tilde{X} := \mathrm{Cay}(G, \tilde{\mathcal{S}})$ obtained from the inclusion $\mathcal{S} \subset \tilde{\mathcal{S}}$ is an infinite strongly contracting geodesic.*

The idea is to turn [3, Example 3.2] into a $Gr'(1/6)$–labelled graph Γ. By construction, Γ will contain a non-strongly contracting geodesic γ that will be labelled by the powers of a generator y. Theorem 4.1 then ensures that the image of γ in the Cayley graph is not strongly contracting. By adding additional edges, corresponding to new generators, to Γ and cutting the resulting graph apart into cycle graphs, we obtain a classical $C'(1/6)$–presentation of the same group in which no relator contains more than one occurrence of the letter y. Thus, a geodesic labelled by the powers of y will be strongly contracting in the Cayley graph with respect to the new generating set.

PROOF OF THEOREM 4.19. Assume the sets $\{y\}$, $\{a, b\}$, $\{x_1, x_2\}$, $\mathcal{S}_1$, $\mathcal{S}_2$ are pairwise disjoint sets, and $|\mathcal{S}_1| \geqslant 2$ and $|\mathcal{S}_2| \geqslant 2$. A *classical piece* with respect to a set of words is the label of a piece in the disjoint union of cycle graphs labelled by the words. It is an exercise to explicitly construct words with the following properties:

Let $\omega := \{w_1, w_2\} \subset \langle \mathcal{S}_1 \rangle$ be a $C'(1/6)$-collection of words (as defined for the proof of Theorem 4.15) such that $|w_1| = |w_2| \geqslant 24$. Moreover, assume that the words $w_1 w_2$, $w_1 w_2^{-1}$, and $w_1^{-1} w_2$ are freely reduced.

Let $\mu(a, b) := \{\mu_i(a, b) \mid i \in \mathbb{N}\} \subset \langle \{a, b\} \rangle$ be a $C'(1/6)$-collection of words such that there exists $C \in \mathbb{N}$, $C \geqslant 6$ with $|\mu_i(a, b)| = (C/|w_1|) \cdot 2^i$. Note that by our assumptions on ω, the set $\mu(w_1, w_2)$ also satisfies the classical $C'(1/6)$–condition: Since no two $w_1^{\pm 1}$ and $w_2^{\pm 2}$ start with the same letter, any piece with respect to $\mu(w_1, w_2)$ comes from a piece with respect to $\mu(a, b)$.

Let $v := \{v_i : i \in \mathbb{N}\} \subset \langle \mathcal{S}_2 \rangle$ be a $C'(1/12)$-collection of words such that $|v_i| = C \cdot 2^i + 1$.

The graph Γ labelled over $\mathcal{S} := \{y\} \cup \mathcal{S}_1 \cup \mathcal{S}_2$ is obtained by taking the disjoint union of the Γ_i as in Figure 14 and, for each i, identifying the 'top' of Γ_i with the 'bottom' of Γ_{i+1}. These both have the same label. Note that any simple closed path γ in Γ is a path going around a finite union $\Gamma_i \cup \Gamma_{i+1} \cup \cdots \cup \Gamma_j$ for $i \leqslant j$, i.e., the label of γ is, up to inversion and cyclic shift, of the form:

$$\mu_i(w_1, w_2) v_i v_{i+1} \ldots v_j \mu_{j+1}(w_1, w_2)^{-1} y^{i-j-1}$$

A piece that is a simple subpath of a cyclic shift of γ has a label that is a subword of one of the following words.

- A piece in some $\mu_k(w_1, w_2)$ or a product of two pieces in some v_k or $v_k v_{k+1}$ (since, by gluing together words v_k when constructing Γ, new pieces may have arisen, each labelled by a product of two pieces in v), or
- y^{i-j-1}, or
- a product of two words from the first bullet, or a product $p y^{i-j-1} q$, where p and q are pieces in $\mu_{j+1}(w_1, w_2)$ and $\mu_i(w_1, w_2)$ respectively.

Here, a piece in $\mu_i(w_1, w_2)$ means a classical piece with respect to the collection of words $\mu(w_1, w_2)$, a piece in v_i means a piece with respect to v, and so on. A piece may also be the empty word.

Note that any piece p in the first bullet has length less than $1/6$ of its ambient word $\mu_k(x_1, x_2)$ or v_k or $v_k v_{k+1}$, and that $6|i - j - 1| \leqslant C2^j < |v_j|$. Therefore, Γ satisfies the $Gr'(1/6)$-condition.

Let α be the geodesic ray in Γ labelled by positive powers of y. Let β be the geodesic ray in Γ labelled by $v_1 v_2 \cdots$. The ray β leaves every bounded neighborhood of α, but has unbounded image under the closest point projection to α in Γ. By Corollary 4.5, α is not strongly contracting in Γ. Thus, by Theorem 4.1, the geodesic ray labelled by positive powers of y is not strongly contracting in $\mathrm{Cay}(G(\Gamma), \mathcal{S})$.

We now define a graph $\tilde{\Gamma}$ labelled over $\tilde{\mathcal{S}} := \{y\} \cup \{x_1, x_2\} \cup \mathcal{S}_1 \cup \mathcal{S}_2$ as follows: Let c_1 be a cycle graph labelled by $x_1 w_1^{-1}$ and c_2 a cycle graph labelled by $x_2 w_2^{-1}$. Set:

$$\tilde{\Gamma} := c_1 \sqcup c_2 \sqcup \bigsqcup_{i \in \mathbb{N}} \tilde{\Gamma}_i$$

where the $\tilde{\Gamma}_i$ are given in Figure 14. Note that while, in this new graph, the paths labelled by $\mu_i(x_1, x_2)$ are pieces, no path labelled by $\mu_{i+1}^{-1}(x_1, x_2)y^{-1}$ or $y^{-1}\mu_i(x_1, x_2)$ is a piece.

By construction, $|\mu_i(x_1, x_2)|, |\mu_{i+1}(x_1, x_2)| < |v_i|/12$, and any piece in v_i has length less than $|v_i|/12$. This, together with observations as above and the fact that any simple piece in c_1 or c_2 has length less than $|c_i|/6$ by construction, shows that $\tilde{\Gamma}$ satisfies the $Gr'(1/6)$–condition. Since its components are cycle graphs, the geodesic ray labelled by the positive powers of y is strongly contracting in $\mathrm{Cay}(G(\tilde{\Gamma}), \tilde{\mathcal{S}})$, by Theorem 4.1.

The presentation for $G(\tilde{\Gamma})$ coming from $\tilde{\Gamma}$ includes the generators and relations of the presentation of $G(\Gamma)$ coming from $G(\Gamma)$, as well as generators x_1 and x_2 and relations $x_1 w_1^{-1}$ and $x_2 w_2^{-1}$. Rewriting the presentation by Tietze transformations, we see $G(\Gamma) \cong G(\tilde{\Gamma})$. In particular, the inclusion $\mathrm{Cay}(G(\Gamma), \mathcal{S}) \hookrightarrow \mathrm{Cay}(G(\tilde{\Gamma}), \tilde{\mathcal{S}})$ is a quasi-isometry. □

We record the following consequence of Theorem 4.1 and Lemma 4.3 for reference:

Corollary 4.20. *Let α and α' be infinite geodesic rays in* $\mathrm{Cay}(G(\Gamma), \mathcal{S})$*, where Γ is a $Gr'(1/6)$–labelled graph labelled by $\mathcal{S}$, such that $d_{\mathrm{Hausdorff}}(\alpha, \alpha') < \infty$. Then α is uniformly locally strongly contracting if and only if α' is.*

5. Strongly contracting elements

In this section, we show the existence of strongly contracting elements in graphical small cancellation groups.

Theorem 5.1. *Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Assume that $G(\Gamma)$ is infinite. Then there exists an infinite order element $g \in G(\Gamma)$ such that $\langle g \rangle$ is strongly contracting in* $\mathrm{Cay}(G(\Gamma), \mathcal{S})$.

The element g is the WPD element for the action of $G(\Gamma)$ on the hyperbolic coned-off space of Gruber and Sisto [30] (see also Section 5.2).

A recent theorem of Arzhantseva, Cashen, and Tao [4] says that if G is a group acting cocompactly on a proper metric space X, and if $g \in G$ is an infinite order element such that closest point projection to an orbit of $\langle g \rangle$ is strongly contracting, then the action of G on X is *growth tight*. This means that the rate of exponential growth of G with respect to the pseudo-metric induced by the metric of X is strictly greater than the growth rate of a quotient of G by any infinite normal subgroup, with respect to the induced pseudo-metric on the quotient group. Thus, a corollary of Theorem 5.1, is:

Theorem 5.2. *Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Then the action of $G(\Gamma)$ on* $\mathrm{Cay}(G(\Gamma), \mathcal{S})$ *is growth tight.*

5.1. Infinite cyclic subgroups are close to periodic geodesics.

In the previous section we deduced contraction results for geodesics in a group defined by a $Gr'(1/6)$-labelled graph. In order to show that cyclic subgroups are strongly contracting, we show that they are actually close to bi-infinite geodesics. As a by-product, we also obtain a result about translation lengths in graphical small cancellation groups.

We glean from [27] the following:

Lemma 5.3. *Let Γ be a $Gr'(1/6)$-labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Every infinite cyclic subgroup of $G := G(\Gamma)$ is at bounded Hausdorff distance from a periodic bi-infinite geodesic in* $\mathrm{Cay}(G, \mathcal{S})$.

Recall that a bi-infinite path graph α in $\mathrm{Cay}(G, \mathcal{S})$ is *periodic* if there is a cyclic subgroup of G that stabilizes α and acts cocompactly on it. In the proof we have surjections $\langle \mathcal{S} \rangle \twoheadrightarrow H \twoheadrightarrow G$. Let $|\cdot|_H$ denote the word length in H with respect to the image of $\mathcal{S}$. Similarly, let $|\cdot|_G$ denote the word length in G with respect to the image of $\mathcal{S}$.

Recall that the translation length of an element $g \in G$ is defined by:

$$\tau_G(g) := \lim_{n\to\infty} \frac{|g^n|_G}{n}$$

This limit always exists since the map $n \mapsto |g^n|_G$ is subadditive.

Proof of Lemma 5.3. Let x be an infinite order element of G.

Claim. There exist $w \in \langle \mathcal{S} \rangle$ and $N \in \mathbb{N}$ such that w represents an element $g \in G$ such that x is conjugate to g^N and the following property is satisfied: There exist a hyperbolic group H, also a quotient of $\langle \mathcal{S} \rangle$, and an epimorphism $\phi\colon H \twoheadrightarrow G$ induced by the identity on $\mathcal{S}$ such that, if h denotes the element of H represented by w, then ϕ restricts to an isometry $\langle h \rangle \to \langle g \rangle$ with respect to the subspace metrics in H and G, respectively.

We show how to deduce the statement of the lemma from the claim: By a theorem of Swenson [49, Theorem 8], since H is hyperbolic, there exist $h_0 \in H$ and $M \in \mathbb{N}$ such that h^M is conjugate to h_0 and such that $|h_0^n|_H = n|h_0|_H$ for all $n > 0$. Therefore, $\tau_H(h^M) = |h_0|_H$, and the bi-infinite path graph in H labelled by the powers of a shortest element $w_0 \in \langle \mathcal{S} \rangle$ representing h_0 is geodesic. Consider $g_0 := \phi(h_0)$, the element of G represented by w_0. Then g^M is conjugate to g_0 by construction, and we have $\tau_G(g_0) = \tau_G(g^M) = \tau_H(h^M) = |w_0|$ since $\langle h \rangle \to \langle g \rangle$ is an isometry. This implies $|g^{nM}|_G = n|w_0|$ for every n, i.e., the bi-infinite path graph in $\mathrm{Cay}(G, \mathcal{S})$ starting at 1 and labelled by the powers of w_0 is geodesic. This proves the lemma, assuming our claim.

It remains to show the claim: By [27, Section 4], there exists $g \in G$ such that x is conjugate to g^N for some N and such that, if w is a shortest word representing g, we have the following possibilities: Any bi-infinite periodic path graph γ labelled by the powers of w is a convex geodesic (Case 1 in [27, Section 4]), or there exists some C_0 such that for any copy c in $\mathrm{Cay}(G, \mathcal{S})$ of a simple cycle in Γ, the length of the intersection of c with γ is at most C_0 (Cases 2a and 2b in [27, Section 4]). If the first possibility is true, then the claim holds for $H = \langle \mathcal{S} \rangle$.

Assume there exists a C_0 as in the second possibility. Then we have the following property by [27, Proof of Theorem 4.2 in Case 2a] and [27, Lemma 4.17]: Whenever v is a geodesic word representing g^n for some n, then the equation $v = w^n$ already holds in $G(\Gamma^{<6C_0})$, where $\Gamma^{<6C_0}$ denotes the subgraph of Γ that is the union of all components with girth less than $6C_0$. Therefore, the epimorphism $G(\Gamma^{<6C_0}) \twoheadrightarrow G$ induced by the identity on $\mathcal{S}$ restricts to an isometry on the cyclic subgroup generated by the element of $G(\Gamma^{<6C_0})$ represented by w. The graph $\Gamma^{<6C_0}$ is (up to identifying isomorphic components) a finite $Gr'(1/6)$–labelled graph, so $G(\Gamma^{<6C_0})$ is a hyperbolic group. □

A result on translation lengths. We record another application of our investigations. We show:

Theorem 5.4. *Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Then every infinite order element of* $G(\Gamma)$ *has rational translation length, and translation lengths are bounded away from zero.*

Recall that similar theorems are true for hyperbolic groups [24, 49, 19] and for the action of the mapping class group of a surface on its curve complex [14].

The rationality statement is a direct consequence of Lemma 5.3. We prove the remaining statement:

Proposition 5.5. *Let Γ be a $Gr'(1/6)$–labelled graph whose components are finite, labelled by a finite set $\mathcal{S}$. Every infinite order element x of $G := G(\Gamma)$ has $\tau_G(x) \geqslant 1/3$.*

Proof. Let x be an infinite order element of G. By [27, Section 4], there exists $g \in G$ such that x is conjugate to g^N for some N and such that, if w is a shortest word representing g, we have the following possibilities: Any bi-infinite periodic path graph γ labelled by the powers of w is a convex geodesic (Case 1 in [27, Section 4]), or, for each n and any shortest word g_n representing g^n, there exists a diagram B_n over Γ whose boundary word is $g_n w^{-n}$, such that B_n is a combinatorial geodesic bigon with respect to the obvious decomposition of ∂B_n (Cases 2a, 2b in [27, Section 4]). In particular, every disk component of B_n is a single face, or has shape I_1.

If the first possibility is true, then $\tau_G(g) = |w| \geqslant 1$. Now consider the second possibility. In Case 2a, it follows from [27, Proof of Theorem 4.2 in Case 2a] that any face Π of B_n, the length of the intersection of Π with the side of B_n corresponding to g_n (the bottom) is more than $1/3$ times the length of its intersection with the side corresponding to w^n (the top). Therefore, in this case, $|g_n| > n|w|/3 \geqslant n/3$. In Case 2b, if a face Π intersects the top in at most $|\partial\Pi|/2$, then it intersects the bottom in more than $|\partial\Pi|/6$ by the small cancellation condition. If Π intersects the top in more than $|\partial\Pi|/2$, then [27, Lemma 4.8] implies that the intersection of Π with the top has length less than $2|w|$. The intersection with the bottom has length at least 1, since every disk component of B_n has shape I_1. Therefore, we have $|g_n| > n|w| \cdot \min\{\frac{1}{3}, \frac{1}{2|w|}\} \geqslant \frac{n}{3}$. We conclude $\tau_G(g) \geqslant 1/3$. Hence, $\tau_G(x) = \tau_G(g^N) = |N|\tau_G(g) \geqslant N/3 \geqslant 1/3$. □

5.2. The coned-off space. In [30], Gruber and Sisto prove that non-elementary groups defined by $Gr(7)$-labelled graphs, which, in particular, includes those defined by $Gr'(1/6)$-labelled graphs, are acylindrically hyperbolic. They prove this result by studying the action of $G(\Gamma)$ on what we call the *coned-off space* Y defined as follows: given a graph Γ labelled by $\mathcal{S}$, let $\mathcal{W}$ denote the set of all elements of $G(\Gamma)$ represented by words read on (not necessarily closed!) paths in Γ. We set $Y := \mathrm{Cay}(G(\Gamma), \mathcal{S} \cup \mathcal{W})$. Thus, we obtain Y from $X := \mathrm{Cay}(G(\Gamma), \mathcal{S})$ by attaching to every embedded component of Γ in X a complete graph.

The proof in [30] shows hyperbolicity of the space Y and existence of an element of $G(\Gamma)$ whose action on Y is hyperbolic and weakly properly discontinuous (WPD). By a theorem of Osin [40], this yields acylindrical hyperbolicity.

Lemma 5.6. *Let Γ be a $Gr'(1/6)$–labelled graph. Let g be an infinite-order element of $G := G(\Gamma)$. Let $X := \mathrm{Cay}(G, \mathcal{S})$, and let Y be the coned-off space. Let γ be a bi-infinite geodesic in X that is at finite Hausdorff distance (in X) from $\langle g\rangle$. Then g is hyperbolic for the action of G on Y if and only if there exists $C > 0$ such that for every embedded component Γ_0 of Γ in X we have* $\mathrm{diam}_X(\gamma \cap \Gamma_0) < C$.

Combining this with Lemma 4.3 and Theorem 4.1, we obtain the following.

Corollary 5.7. *If $\langle g\rangle$ is within bounded Hausdorff distance of a bi-infinite geodesic in X and acts hyperbolically on Y, then $\langle g\rangle$ is strongly contracting in X.*

Remark 5.8. The converse of Corollary 5.7 fails. Consider the group defined by $\sqcup_{i\in\mathbb{N}}\Gamma_i$ constructed in Example 4.18. The subgroup $\langle a\rangle$ is strongly contracting, but in the coned-off space Y its orbit has diameter at most 2, so a does not act hyperbolically on Y. Thus, the methods of [30] do not detect that $\langle a\rangle$ is strongly contracting.

Proof of Lemma 5.6 . Let γ be a bi-infinite geodesic in X whose Hausdorff distance from $\langle g\rangle$ is equal to $\epsilon < \infty$.

Suppose for every $n \in \mathbb{N}$ there exists an embedded component Γ_n of Γ in X such that $\mathrm{diam}_X(\Gamma_n \cap \gamma) \geqslant n$. Denote by γ_n the path graph $\Gamma_n \cap \gamma$. (This is a connected set since each Γ_n is convex by [30, Lemma 2.15].) Let $\iota\gamma_n$ and $\tau\gamma_n$ denote, respectively, the initial and terminal vertices of γ_n. Then, by assumption, for each n, there exist m_n and l_n such that $d_X(\iota\gamma_n, g^{m_n}) < \epsilon$ and $d_X(\tau\gamma_n, g^{l_n}) < \epsilon$ and, since $|\gamma_n| \to \infty$, we have $|m_n - l_n| \to \infty$. We have $d_Y(g^{m_n}, g^{l_n}) < 2\epsilon + 1$, since the vertex set of Γ_n has diameter at most 1 in the metric of Y. Therefore, the map $\mathbb{Z} \to Y\colon z \mapsto g^z$ is not a quasi-isometric embedding.

On the other hand, suppose there exists C such that, for every embedded component Γ_0 of Γ, we have $\operatorname{diam}_X(\gamma \cap \Gamma_0) < C$. Then, by [30, Proposition 3.6], we have, for any k and l, that $d_Y(g^k, g^l) > \frac{1}{C}(d_X(g^k, g^l) - 2\epsilon) - 2\epsilon$. Since $\langle g \rangle$ is undistorted in G by [27, Theorem 4.2], this gives the lower quasi-isometry bound for the map $\mathbb{Z} \to Y\colon\ z \mapsto g^z$. Since this map is obviously Lipschitz, it is, in fact, a quasi-isometric embedding, whence g is hyperbolic. □

We will use the following result of [30]:

Proposition 5.9 ([30, Section 4]). *Let Γ be a $Gr'(1/6)$-labelled graph. Suppose that Γ has at least two non-isomorphic components that each contain an embedded cycle of length at least 2. Then $G(\Gamma)$ contains a hyperbolic element g for the action of $G(\Gamma)$ on the coned-off space Y.*

5.3. Proof of Theorem 5.1. If Γ has only finitely many pairwise non-isomorphic components with non-trivial fundamental groups, then $G(\Gamma)$ is Gromov hyperbolic [37, 26] and, hence, the result holds.

If Γ has infinitely many pairwise non-isomorphic components with non-trivial fundamental groups, then, since $\mathcal{S}$ is finite, there exist at least two such components each containing an embedded cycle of length at least 2. Therefore, by Proposition 5.9, there exists a hyperbolic element g for the action of $G(\Gamma)$ on Y. Since the components of Γ are finite, Lemma 5.3 yields that $\langle g \rangle$ is within bounded Hausdorff distance of a bi-infinite geodesic in X. Therefore, Corollary 5.7 implies the result. □

6. Hyperbolically embedded subgroups

6.1. Contracting subgroups and hyperbolically embedded subgroups. We recall a definition of hyperbolically embedded subgroups from [18].

Definition 6.1 (Hyperbolically embedded subgroup). Let G be a group and H a subgroup. Then H is *hyperbolically embedded* in G if there exists a subset $\mathcal{S}$ of G with the properties below. We denote by X^H the Cayley graph of H with respect to H, considered as subgraph of $X := \operatorname{Cay}(G, H \sqcup \mathcal{S})$, and by $\hat{d}$ the metric on H obtained as follows: we define $\hat{d}(x, y)$ to be the length of a shortest path from x to y in X that does not use any edges of X^H. If no such path exists, set $\hat{d}(x, y) = \infty$. The required properties are:

- $H \sqcup \mathcal{S}$ generates G (i.e. X is connected).
- X is Gromov hyperbolic.
- The metric $\hat{d}$ is proper on H, i.e. $\hat{d}$–balls are finite.

Note that if $s \in H \cap \mathcal{S}$, then it is our convention that $\operatorname{Cay}(G, H \sqcup \mathcal{S})$ will have a double-edge corresponding to s (once considered as element of H and once considered as element of $\mathcal{S}$); hence the symbol $\sqcup$ for disjoint union.

One approach to finding hyperbolically embedded subgroups is provided by [12, Theorem H]. A special case of this theorem states that if G is a finitely generated group and a hyperbolic element $g \in G$ has a strongly contracting orbit in a Cayley graph of G, then the elementary closure $E(g)$ of g is an infinite, virtually cyclic, hyperbolically embedded subgroup of G.

Combining this with Lemma 4.3 and Theorem 4.1, we obtain the following.

Theorem 6.2. *Let $G := G(\Gamma)$ be the group defined by a $Gr'(1/6)$–labelled graph Γ. Let $g \in G$ be of infinite order, and let γ be a bi-infinite geodesic at finite Hausdorff distance from $\langle g \rangle$. If γ is uniformly locally strongly contracting, then $\langle g \rangle$ is strongly contracting and, in particular, the elementary closure $E(g)$ of g is a virtually cyclic hyperbolically embedded subgroup.*

Remark 6.3. By [12, Theorem H], we can also view Theorem 5.1 as an alternative proof of acylindrical hyperbolicity for the groups considered in that theorem. The initial proof of acylindrical hyperbolicity of these groups in [30] relies on the hyperbolicity of the coned-off space and on showing that the element g satisfies a certain weak proper discontinuity condition (and, in fact, applies to a larger class of groups). Our Theorem 6.2 gives an alternative proof that the element g from [30] gives rise to a hyperbolically embedded virtually cyclic subgroup.

Every hyperbolically embedded subgroup of a finitely generated group is Morse [45], and every element with a strongly contracting orbit has a virtually cyclic hyperbolically embedded elementary closure. In graphical small cancellation groups, Morse and strongly contracting elements have been classified in terms of the defining graph, by our results in Section 4. It is natural to ask whether the collection of cyclic hyperbolically embedded subgroups can be classified in a similar way. We give one negative result in this direction.

Theorem 6.4. *Let ρ_2 be an unbounded sublinear function. There exists a $Gr'(1/6)$–labelled graph Γ with set of labels $\mathcal{S} := \{a, b\}$ whose components are all cycles such that the group $G := G(\Gamma)$ has the following properties: Any virtually cyclic subgroup E of G containing $\langle a\rangle$ is (r, ρ_2')–contracting in the Cayley graph $X := \mathrm{Cay}(G, \mathcal{S})$ for some $\rho_2' \asymp \rho_2$, but E is not hyperbolically embedded in G.*

Proof. For every $r > 0$, choose $N_r \geqslant 6$ such that $\rho_2(\frac{1}{2}(1 + r)N_r) \geqslant r$. This is possible because ρ_2 is unbounded. For every $r > 0$, let $R_r := (a^r ba^r b^{-1})^{N_r}$, and consider the graph Γ that is the disjoint union of cycles γ_r labelled by the R_r. Any reduced path that is a piece in γ_r has length at most $2r+1$, which is less than $|\gamma_r|/6$. Therefore, Γ satisfies the $Gr'(1/6)$-condition. Denote by α a bi-infinite path in X labelled by the powers of a. For every copy γ of a γ_r embedded in X, we have $\operatorname{diam} \alpha \cap \gamma < |\gamma|/6$. Therefore, it is readily seen from considering diagrams of shape I_1 that α is a bi-infinite geodesic. Also note that γ intersects α in a path of length at most r.

Suppose γ is a relator intersecting α, and x is a point in γ with $\delta := d(x, \alpha)$. If $|\gamma| > 4\delta$, then, for any point y in γ with $d(x, y) \leqslant \delta$, we have $\pi(y) = \pi(x)$, and $\pi(x)$ is a singleton, where π denotes closest point projection to α. Now assume $|\gamma| \leqslant 4\delta$, and $y \in \gamma$. Then $\operatorname{diam} \pi(x) \cup \pi(y) \leqslant \operatorname{diam} \gamma \cap \alpha \leqslant \rho_2(\delta)$, by construction, since ρ_2 is non-decreasing. Therefore, α is locally (r, ρ_2)-contracting and, hence, by Theorem 4.1 and Lemma 4.3 there exists ρ_2' with $\rho_2' \asymp \rho_2$ such that any virtually cyclic subgroup E of G containing $\langle a\rangle$ is (r, ρ_2')-contracting.

If E is a hyperbolically embedded subgroup then, by [18, Theorem 5.3], there exists $r > 0$ such that the normal closure N of a^r is a free group. By construction, the element of G represented by $a^r ba^r b^{-1}$ is non-trivial and has finite order, whence N is not torsion-free. Therefore, E is not hyperbolically embedded. □

Another natural question is the following.

Question 6.5. Let G be a group generated by a finite set $\mathcal{S}$, and suppose $g \in G$ has a (ρ_1, ρ_2)–contracting orbit in $X := \mathrm{Cay}(G, \mathcal{S})$ where ρ_2 is bounded. Is the elementary closure $E(g)$ a hyperbolically embedded virtually cyclic subgroup of G?

Is the statement true if $G := G(\Gamma)$ for a $Gr'(1/6)$–labelled graph Γ whose components are finite, with finite set of labels $\mathcal{S}$?

In the case of classical $C'(1/6)$-groups, an affirmative answer follows from Theorem 4.16.

6.2. Hyperbolically embedded cycles. Recall that a group has the *hyperbolically embedded cycles property (HEC property)* if the elementary closure $E(g)$ of every infinite order element g is virtually cyclic and hyperbolically embedded. Hyperbolic groups have this property: $E(g)$ is the stabilizer of a g-axis. The HEC property in fact passes to any subgroup of a hyperbolic group G using the action of the subgroup on a Cayley graph of G.

It is therefore very natural to ask whether this classifies subgroups of hyperbolic groups. Torsion presents one complication: A free product of infinite torsion groups, or, more generally, a group hyperbolic relative to infinite torsion subgroups, has the HEC property but cannot be a subgroup of a hyperbolic group. We show that even among torsion-free groups there are many examples of groups with the HEC property that are not subgroups of any hyperbolic group.

Theorem 6.6. *There exist $2^{\aleph_0}$ pairwise non-quasi-isometric finitely generated torsion-free groups in which every non-trivial cyclic subgroup is strongly contracting and which, therefore, have the HEC property.*

Since there are only countably many finitely generated subgroups of finitely presented groups, most of the groups of Theorem 6.6 do not occur as subgroups of hyperbolic groups.

We show in Corollary 6.10 that there are even exotic examples of groups with the HEC property such as the Gromov monster groups.

The theorem is proven by building small cancellation groups in which no power of any element has long intersection with an embedded component of the defining graph. In Theorem 6.7 we give a condition on the labelling that guarantees this property. In Theorem 6.8 we show that this condition can be satisfied. Then we construct specific examples satisfying our condition and apply a version of a construction of Thomas and Velickovic [51], proven in Proposition 6.13, to show that we get uncountably many quasi-isometry classes of groups.

A labelling of the edges of an undirected graph is said to be *non-repetitive* if there does not exist a non-trivial embedded path graph that is labelled by a word of the form ww. (Here, the label of a path is just the concatenation of the labels of the edges in the free monoid on the labelling set.) The *Thue number* of a graph is the minimal cardinality of a labelling set for which the graph admits a non-repetitive labelling.

We define a labelling of a directed graph to be *non-repetitive* if there does not exist a non-trivial embedded path graph $\gamma = e_1, \dots, e_{2n}$ such that for all $1 \leqslant i \leqslant n$ the label of the directed edge e_i is equal to the label of the directed edge e_{i+n}. Note that, given an undirected graph with a non-repetitive labelling, any choice of orientation gives rise to a non-repetitive labelling of the resulting oriented graph.

Theorem 6.7. *Let $\mathcal{S}$ be a finite set and let Γ be a graph with finite components and a labelling by $\mathcal{S}$ that is both $Gr'(1/6)$ and non-repetitive. Let $G := G(\Gamma)$. Every infinite cyclic subgroup H of G is strongly contracting in $X := \mathrm{Cay}(G, \mathcal{S})$. Thus, G has the HEC property.*

Proof. By Lemma 5.3, every infinite cyclic subgroup is bounded Hausdorff distance from a periodic geodesic α. By Lemma 4.3, it is enough to show α is strongly contracting, so we may assume that $H = \langle h \rangle$ acts cocompactly on α, and that h is represented by a cyclically reduced word v whose powers label α.

Suppose α intersects some embedded component Γ_0 in more than a single vertex. Since α is geodesic, $\Gamma_0 \cap \alpha$ is an embedded path graph. Since the labelling of Γ is non-repetitive, the label of this path does not contain a subword of the form ww. However, the labelling of α is repetitive — it is $\cdots vv \cdots$. Thus, $|\Gamma_0 \cap \alpha| < 2|v|$. We conclude α is strongly contracting by Theorem 4.1.

By [12, Theorem H], the elementary closure of H is a virtually cyclic hyperbolically embedded subgroup. □

Suppose that Γ is a directed graph with two labellings $L_1 \colon \mathcal{E}\Gamma \to \mathcal{S}_1$ and $L_2 \colon \mathcal{E}\Gamma \to \mathcal{S}_2$. We define the *push-out labelling* $L \colon \mathcal{E}\Gamma \to \mathcal{S}$, where $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2$, by $L(e) := (L_1(e), L_2(e))$. We will write $\mathcal{S}_1$ alphabetically and $\mathcal{S}_2$ as numerical index, i.e. $\mathcal{S}^{\pm} = \{a_n^{\epsilon} \mid \epsilon \in \pm 1,\ a \in \mathcal{S}_1,\ n \in \mathcal{S}_2\}$.

If $L_2 \colon \mathcal{E}\Gamma \to \mathcal{S}_2$ is a non-repetitive labelling of Γ, then for any labelling $L_1 \colon \mathcal{E}\Gamma \to \mathcal{S}_1$ the push-out labelling is a non-repetitive labelling of Γ. Similarly, if $L_1 \colon \mathcal{E}\Gamma \to \mathcal{S}_1$ satisfies the $C'(1/6)$-condition, then so does the push-out labelling.

Theorem 6.8. *Let $\Gamma = (\Gamma_i)_{i \in \mathbb{N}}$ be a sequence of finite, connected graphs satisfying the following conditions:*

- *Γ has bounded valence.*
- *$(\mathrm{girth}(\Gamma_i))_{i \in \mathbb{N}}$ is an unbounded sequence.*
- *The ratios $\frac{\mathrm{girth}(\Gamma_i)}{\mathrm{diam}(\Gamma_i)}$ are bounded, uniformly over i, away from 0.*

Then there exist an infinite subsequence $(\Gamma_{i_j})_{j \in \mathbb{N}}$ of graphs and a finite set $\mathcal{S}$ such that $(\Gamma_{i_j})_{j \in \mathbb{N}}$ admits an labelling by $\mathcal{S}$ that is both $C'(1/6)$ and non-repetitive.

Proof. A theorem of Osajda [39] says that, given the hypotheses on Γ, there exists an infinite subsequence $(\Gamma_{i_j})_{j \in \mathbb{N}}$ that admits a choice of orientation and a labelling by a finite set $\mathcal{S}_1$ satisfying the $C'(1/6)$-condition. Alon, et al. [2] show that the Thue number of a bounded valence undirected graph is bounded by a polynomial function of the valence bound. Since Γ has bounded valence, there exists a finite set of labels $\mathcal{S}_2$ such that Γ admits a non-repetitive $\mathcal{S}_2$–labelling (as undirected graph and, hence, also as directed graph). The push-out of these two labellings satisfies the theorem. □

Combining this with Theorem 6.7, we have:

Corollary 6.9. *If* Γ *is as in Theorem 6.8 then* $G(\Gamma)$ *has the HEC property.*

Corollary 6.10. *There exist Gromov monster groups with the HEC property.*

More generally, if Γ is a bounded valence graph with a labelling satisfying some property $\mathcal{P}$ such that $\mathcal{P}$ is preserved upon passing to a refinement of the labelling, then Γ admits a labelling that is both $\mathcal{P}$ and non-repetitive. For instance, Arzhantseva and Osajda [6] introduced a 'lacunary walling condition' to produce first examples of non-coarsely amenable groups[2] with the Haagerup property. This condition is preserved upon passing to refinements of the labelling, so the same argument as in Theorem 6.8 yields:

Corollary 6.11. *There exist finitely generated, non-coarsely amenable groups with the Haagerup property and the HEC property.*

Not every interesting property of labellings is preserved upon passing to refinements. For example, the $Gr'(\lambda)$ condition is not preserved, because in the original labelling there may be a long labelled path p with distinct label-preserving maps ϕ_1, $\phi_2\colon\ p \to \Gamma$ and a label-preserving automorphism ψ of Γ such that $\phi_2 = \psi \circ \phi_1$. If ψ fails to be a label-preserving automorphism of Γ with the refined labelling then p may be too long a piece.

For another example, passing to a refinement can yield a group with non-trivial free factors. In particular, no property of a labelling that implies the resulting group has Kazhdan's Property (T) is preserved by passing to refinements.

Both of the factor labellings in the proof of Theorem 6.8 are produced probabilistically using the Lovász Local Lemma. For the purpose of proving Theorem 6.6 we construct *explicit* examples of labelled graphs satisfying Theorem 6.7, so we get concrete examples of groups satisfying Theorem 6.6. To construct such examples we use the fact, first observed by Alon, et al. [2], that a cycle graph has Thue number at most 4, as follows:

Definition 6.12 (Thue-Morse sequence). Define $\sigma\colon \{0,1\}^* \to \{0,1\}^*$ by $\sigma(0) := 01$, $\sigma(1) := 10$. The sequence $(x_i)_{i\in\mathbb{N}} := \sigma^\infty(0)$ is called the *Thue-Morse sequence*[3].

The Thue-Morse sequence famously does not contain any subword of the form www. The 'first difference' sequence[4] $(y_i)_{i\in\mathbb{N}} \in \{-1,0,1\}^{\mathbb{N}}$ defined by $y_i := x_{i+1} - x_i$, where $(x_i)_{i\in\mathbb{N}}$ is the Thue-Morse sequence, does not contain any subword of the form ww, which means that it gives a non-repetitive labelling of the ray graph. Both of these facts are due to Thue [52].

A cycle graph of length n with edges $e_0, e_1, \dots, e_{n-1}$ admits a non-repetitive labelling by $\{-1,0,1,\infty\}$ by labelling edge e_0 with ∞ and labelling edge e_i for $i > 0$ with term y_i of the first difference of the Thue-Morse sequence.

Define $\Gamma := (\Gamma_n)_{n\in\mathbb{N}}$ to be a disjoint union of cycle graphs such that $|\Gamma_n| = 11(44n-19)$. Give each of these cycles the non-repetitive $\{-1,0,1,\infty\}$–labelling defined above. Also give Γ_n the $\{a,b\}$–labelling $R_n := \prod_{i=22n-21}^{22n} ab^i$. This is a $C'(1/6)$–labelling. There are no label-preserving automorphisms of Γ since the components are cycles of different lengths and are labelled by positive words that are not proper powers. If p is a piece contained in Γ_n then, since the gaps between a's in the R_i are all different, p contains at most one edge labelled a. The longest subword of R_n containing at most one a is $b^{22n-1}ab^{22n}$ of length $44n$. For $n \in \mathbb{N}$, the ratio $\frac{44n}{11(44n-19)}$ takes maximum value $4/25 < 1/6$ at $n = 1$.

The push-out of these two labellings is an $\mathcal{S}$–labelling of Γ that is both non-repetitive and $C'(1/6)$, with $\mathcal{S} := \{a_{-1}, a_0, a_1, a_\infty, b_{-1}, b_0, b_1, b_\infty\}$.

For each subset $I \subset \mathbb{N}$ define $\Gamma_I := (\Gamma_n)_{n\in I}$. All of these graphs have non-repetitive, $C'(1/6)$–labellings inherited from Γ. By Theorem 6.7, every non-trivial cyclic subgroup of the group G_I defined by Γ_I is strongly contracting. The proof of Theorem 6.6 is completed by the following proposition and the fact that groups defined by $C'(1/6)$-labelled graphs are torsion-free [26].

[2] These are groups G whose reduced C^*–algebra $C^*_{red}(G)$ is not exact.

[3] Sequence A010060 of [50].

[4] Sequence A029883 of [50].

Proposition 6.13 ([51]). *There is a subset $\mathcal{I} \subseteq 2^{\mathbb{N}}$ of cardinality $2^{\aleph_0}$ such that given $I, J \in \mathcal{I}$, the groups G_I and G_J are quasi-isometric if and only if $I = J$.*

Proof. Choose $\mathcal{I}$ to be a collection of infinite subsets of $\{2^{2^n} \mid n \in \mathbb{N}\}$ with infinite pairwise symmetric difference.

Let $I, J \in \mathcal{I}$ be distinct. Without loss of generality, assume $I \setminus J$ is infinite, and let μ be a non-atomic ultrafilter on $\mathbb{N}$ with $\mu(I) = 1$ and $\mu(J) = 0$.

The asymptotic cone of G_J over μ with scaling sequence $(|R_n|)_{n \in I}$ is an $\mathbb{R}$-tree, while the asymptotic cone of G_I over μ with the same scaling contains a loop of length 1.

If G_I and G_J are quasi-isometric then their asymptotic cones are bi-Lipschitz equivalent. Thus these groups are not quasi-isometric. □

References

[1] Yael Algom-Kfir. *Strongly contracting geodesics in outer space*. Geom. Topol. **15**.4 (2011), pp. 2181–2233.

[2] Noga Alon, Jarosław Grytczuk, Mariusz Hałuszczak, and Oliver Riordan. *Nonrepetitive colorings of graphs*. Random Structures & Algorithms **21**.3-4 (2002), pp. 336–346.

[3] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Characterizations of Morse geodesics via superlinear divergence and sublinear contraction*. Doc. Math. **22** (2017), pp. 1193–1224.

[4] Goulnara N. Arzhantseva, Christopher H. Cashen, and Jing Tao. *Growth Tight Actions*. Pacific J. Math. **278**.1 (2015), pp. 1–49.

[5] Goulnara N. Arzhantseva and Thomas Delzant. *Examples of random groups*. preprint, available on the authors' webpages. 2008.

[6] Goulnara Arzhantseva and Damian Osajda. *Graphical small cancellation groups with the Haagerup property*. preprint. 2014. arXiv: 1404.6807v2.

[7] Goulnara Arzhantseva and Markus Steenbock. *Rips construction without unique product*. preprint. 2014. arXiv: 1407.2441.

[8] Goulnara Arzhantseva and Romain Tessera. *Admitting a coarse embedding is not preserved under group extensions*. Int. Math. Res. Not. IMRN (in press). arXiv: 1605.01192.

[9] Jason A. Behrstock. *Asymptotic geometry of the mapping class group and Teichmüller space*. Geom. Topol. **10** (2006), pp. 1523–1578.

[10] Jason Behrstock and Cornelia Druţu. *Divergence, thick groups, and short conjugators*. Illinois J. Math. **58**.4 (2014), pp. 939–980.

[11] Jason Behrstock, Cornelia Druţu, and Lee Mosher. *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*. Math. Ann. **344**.3 (2009), pp. 543–595.

[12] Mladen Bestvina, Kenneth Bromberg, and Koji Fujiwara. *Constructing group actions on quasi-trees and applications to mapping class groups*. Publ. Math. Inst. Hautes Études Sci. **122**.1 (2015), pp. 1–64.

[13] Mladen Bestvina and Koji Fujiwara. *A characterization of higher rank symmetric spaces via bounded cohomology*. Geom. Funct. Anal. **19**.1 (2009), pp. 11–40.

[14] Brian H. Bowditch. *Tight geodesics in the curve complex*. Invent. Math. **171**.2 (2008), pp. 281–300.

[15] Gregory R. Conner. *A class of finitely generated groups with irrational translation numbers*. Arch. Math. (Basel) **69**.4 (1997), pp. 265–274.

[16] Gregory R. Conner. *Translation numbers of groups acting on quasiconvex spaces*. *Computational and geometric aspects of modern algebra (Edinburgh, 1998)*. Vol. 275. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 2000, pp. 28–38.

[17] Rémi Cunéo and Hamish Short. *Graph small cancellation theory applied to alternating link groups*. Journal of Knot Theory and Its Ramifications **21**.11 (2012), 1250113, 15 pages.

[18] François Dahmani, Vincent Guirardel, and Denis Osin. *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*. Mem. Amer. Math. Soc. **245**.1156 (2017), pp. v+152.

[19] Thomas Delzant. *Sous-groupes distingués et quotients des groupes hyperboliques*. Duke Math. J. **83**.3 (1996), pp. 661–682.

[20] Cornelia Druţu, Shahar Mozes, and Mark Sapir. *Divergence in lattices in semisimple Lie groups and graphs of groups*. Trans. Amer. Math. Soc. **362**.5 (2010), pp. 2451–2505.

[21] Cornelia Druţu and Mark Sapir. *Tree-graded spaces and asymptotic cones of groups*. Topology **44**.5 (2005). With an appendix by D. Osin and Sapir, pp. 959–1058.

[22] Moon Duchin and Kasra Rafi. *Divergence of geodesics in Teichmüller space and the mapping class group*. Geom. Funct. Anal. **19**.3 (2009), pp. 722–742.

[23] Rostislav Grigorchuk and Pierre de la Harpe. *On problems related to growth, entropy, and spectrum in group theory*. J. Dynam. Control Systems **3**.1 (1997), pp. 51–89.

[24] Mikhael Gromov. *Hyperbolic groups*. *Essays in group theory*. Vol. 8. Math. Sci. Res. Inst. Publ. New York: Springer, 1987, pp. 75–263.

[25] Mikhael Gromov. *Random walk in random groups*. Geom. Funct. Anal. **13**.1 (2003), pp. 73–146.

[26] Dominik Gruber. *Groups with graphical $C(6)$ and $C(7)$ small cancellation presentations*. Trans. Amer. Math. Soc. **367**.3 (2015), pp. 2051–2078.

[27] Dominik Gruber. *Infinitely presented C(6)-groups are SQ-universal*. J. Lond. Math. Soc. **92**.1 (2015), pp. 178–201.

[28] Dominik Gruber. *Infinitely presented graphical small cancellation groups: Coarse embeddings, acylindrical hyperbolicity, and subgroup constructions*. PhD thesis. University of Vienna, 2015.

[29] Dominik Gruber, Alexandre Martin, and Markus Steenbock. *Finite index subgroups without unique product in graphical small cancellation groups*. Bull. London Math. Soc. **47**.4 (2015), pp. 631–638.

[30] Dominik Gruber and Alessandro Sisto. *Infinitely presented graphical small cancellation groups are acylindrically hyperbolic*. Ann. Inst. Fourier (in press). arXiv: `1408.4488v3`.

[31] Nigel Higson, Vincent Lafforgue, and George Skandalis. *Counterexamples to the Baum-Connes conjecture*. Geom. Funct. Anal. **12**.2 (2002), pp. 330–354.

[32] David Hume. *A continuum of expanders*. Fund. Math. **238**.2 (2017), pp. 143–152.

[33] Ilya Kapovich. *Small cancellation groups and translation numbers*. Trans. Amer. Math. Soc. **349**.5 (1997), pp. 1851–1875.

[34] Howard A. Masur and Yair N. Minsky. *Geometry of the complex of curves. I. Hyperbolicity*. Invent. Math. **138**.1 (1999), pp. 103–149.

[35] Howard A. Masur and Yair N. Minsky. *Geometry of the complex of curves. II. Hierarchical structure*. Geom. Funct. Anal. **10**.4 (2000), pp. 902–974.

[36] Yair N. Minsky. *Quasi-projections in Teichmüller space*. J. Reine Angew. Math. **473** (1996), pp. 121–136.

[37] Yann Ollivier. *On a small cancellation theorem of Gromov*. Bull. Belg. Math. Soc. Simon Stevin **13**.1 (2006), pp. 75–89.

[38] Yann Ollivier and Daniel T. Wise. *Kazhdan groups with infinite outer automorphism group*. Trans. Amer. Math. Soc. **359**.5 (2007), pp. 1959–1976.

[39] Damian Osajda. *Small cancellation labellings of some infinite graphs and applications*. preprint. 2014. arXiv: `1406.5015v1`.

[40] Denis V. Osin. *Acylindrically hyperbolic groups*. Trans. Amer. Math. Soc. **368**.2 (2016), pp. 851–888.

[41] Stephen Pride. *Some problems in combinatorial group theory*. *Groups—Korea 1988 (Pusan, 1988)*. Vol. 1398. Lecture Notes in Math. Springer, Berlin, 1989, pp. 146–155.

[42] Eliyahu Rips and Yoav Segev. *Torsion-free group without unique product property*. J. Algebra **108**.1 (1987), pp. 116–126.

[43] Lior Silberman. *Addendum to: "Random walk in random groups" [Geom. Funct. Anal.* **13** *(2003), no. 1, 73–146] by M. Gromov*. Geom. Funct. Anal. **13**.1 (2003), pp. 147–177.

[44] Alessandro Sisto. *Projections and relative hyperbolicity*. Enseign. Math. (2) **59**.1-2 (2013), pp. 165–181.

[45] Alessandro Sisto. *Quasi-convexity of hyperbolically embedded subgroups*. Math. Z. **283**.3-4 (2016), pp. 649–658.

[46] Markus Steenbock. *Rips-Segev torsion-free groups without the unique product property*. J. Algebra **438** (2015), pp. 337–378.

[47] Ralph Strebel. *Appendix. Small cancellation groups*. *Sur les groupes hyperboliques d'après Mikhael Gromov (Bern, 1988)*. Vol. 83. Progr. Math. Birkhäuser Boston, Boston, MA, 1990, pp. 227–273.

[48] Harold Sultan. *Hyperbolic quasi-geodesics in CAT(0) spaces*. Geom. Dedicata **169**.1 (2014), pp. 209–224.

[49] Eric L. Swenson. *Hyperbolic elements in negatively curved groups*. Geom. Dedicata **55**.2 (1995), pp. 199–210.

[50] *The on-line encyclopedia of integer sequences*. URL: `http://oeis.org` (visited on 2015).

[51] Simon Thomas and Boban Velickovic. *Asymptotic cones of finitely generated groups*. Bull. London Math. Soc. **32**.2 (2000), pp. 203–208.

[52] Axel Thue. *Über die gegenseitige Lage gleicher Teile gewisser Zeichenreihen.* German. Kristiania: J. Dybwad. 67 S. Lex. 8°. 1912.

PAPER D

Growth Tight Actions

Goulnara N. Arzhantseva, Christopher H. Cashen, and Jing Tao
Pacific Journal of Mathematics **278** (2015), no. 1, 1–49.
DOI:10.2140/pjm.2015.278.1

We introduce and systematically study the concept of a growth tight action. This generalizes growth tightness for word metrics as initiated by Grigorchuk and de la Harpe. Given a finitely generated, non-elementary group G acting on a G–space $\mathcal{X}$, we prove that if G contains a strongly contracting element and if G is not too badly distorted in $\mathcal{X}$, then the action of G on $\mathcal{X}$ is a growth tight action. It follows that if $\mathcal{X}$ is a cocompact, relatively hyperbolic G–space, then the action of G on $\mathcal{X}$ is a growth tight action. This generalizes all previously known results for growth tightness of cocompact actions: every already known example of a group that admits a growth tight action and has some infinite, infinite index normal subgroups is relatively hyperbolic, and, conversely, relatively hyperbolic groups admit growth tight actions. This also allows us to prove that many CAT(0) groups, including flip-graph-manifold groups and many Right Angled Artin Groups, and snowflake groups admit cocompact, growth tight actions. These provide first examples of non relatively hyperbolic groups admitting interesting growth tight actions. Our main result applies as well to cusp uniform actions on hyperbolic spaces and to the action of the mapping class group on Teichmüller space with the Teichmüller metric. Towards the proof of our main result, we give equivalent characterizations of strongly contracting elements and produce new examples of group actions with strongly contracting elements.

0. Introduction

The growth exponent of a set $\mathcal{A} \subset \mathcal{X}$ with respect to a pseudo-metric d is

$$\delta_{\mathcal{A},d} := \limsup_{r\to\infty} \frac{\log \#\{a \in \mathcal{A} \mid d(o,a) \leqslant r\}}{r}$$

where # denotes cardinality and $o \in \mathcal{X}$ is some basepoint. The limit is independent of the choice of basepoint.

Let G be a finitely generated group. A left invariant pseudo-metric d on G induces a left invariant pseudo-metric $\bar{d}$ on any quotient G/Γ of G by $\bar{d}(g\Gamma, g'\Gamma) := d(g\Gamma, g'\Gamma)$.

Definition 0.1. G is *growth tight* with respect to d if $\delta_{G,d} > \delta_{G/\Gamma,\bar{d}}$ for every infinite normal subgroup $\Gamma \trianglelefteqslant G$.

One natural way to put a left invariant metric on a finitely generated group is to choose a finite generating set and consider the word metric. More generally, pseudo-metrics on a group are provided by actions of the group on metric spaces. Let $\mathcal{X}$ be a G–space, that is, a proper, geodesic metric space with a properly discontinuous, isometric G–action $G \curvearrowright \mathcal{X}$. The choice of a basepoint $o \in \mathcal{X}$ induces a left invariant pseudo-metric on G by $d_G(g, g') := d_{\mathcal{X}}(g.o, g'.o)$.

The first two authors supported by the European Research Council (ERC) grant of Goulnara ARZHANTSEVA, grant agreement #259527 and the Erwin Schrödinger Institute workshop "Geometry of Computation in Groups". The second author partially supported by the Austrian Science Fund (FWF):M1717-N25. The third author supported by NSF grant DMS-1311834.

Define the growth exponent δ_G of G with respect to $\mathcal{X}$ to be the growth exponent of G with respect to an induced pseudo-metric d_G. This depends only on the G–space $\mathcal{X}$, since a different choice of basepoint in $\mathcal{X}$ defines a pseudo-metric that differs from d_G by an additive constant. Likewise, let $\delta_{G/\Gamma}$ denote the growth exponent of G/Γ with respect to a pseudo-metric on G/Γ induced by $d_{\mathcal{X}}$.

Definition 0.2. $G \curvearrowright \mathcal{X}$ is a *growth tight action* if $\delta_G > \delta_{G/\Gamma}$ for every infinite normal subgroup $\Gamma \trianglelefteq G$.

Some groups admit growth tight actions for the simple reason that they lack any infinite, infinite index normal subgroups. For such a group G, every action on a G–space with positive growth exponent will be growth tight. Exponentially growing simple groups are examples, as, by the Margulis Normal Subgroup Theorem [41], are irreducible lattices in higher rank semi-simple Lie groups.

Growth tightness[1] for word metrics was introduced and studied by Grigorchuk and de la Harpe [32], who showed, for example, that a finite rank free group equipped with the word metric from a free generating set is growth tight. On the other hand, they showed that the product of a free group with itself, generated by free generating sets of the factors, is not growth tight. Together with the Normal Subgroup Theorem, these results suggest that for interesting examples of growth tightness we should examine 'rank 1' type behavior. Further evidence for this idea comes from the work of Sambusetti and collaborators, who in a series of papers [52, 51, 49, 26] prove growth tightness for the action of the fundamental group of a negatively curved Riemannian manifold on its Riemannian universal cover.

In the study of non-positively curved, or CAT(0), spaces there is a well established idea that a space may be non-positively curved but have some specific directions that look negatively curved. More precisely:

Definition 0.3 ([7]). A hyperbolic isometry of a proper CAT(0) space is *rank 1* if it has an axis that does not bound a half-flat.

In Definition 2.17, we introduce the notion for an element of G to be *strongly contracting* with respect to $G \curvearrowright \mathcal{X}$. In the case that $\mathcal{X}$ is a CAT(0) G–space, the strongly contracting elements of G are precisely those that act as rank 1 isometries of $\mathcal{X}$ (see Theorem 9.1).

In addition to having a strongly contracting element, we will assume that the orbit of G in $\mathcal{X}$ is not too badly distorted. There are two different ways to make this precise.

We say a G–space is *C–quasi-convex* if there exists a C–quasi-convex G–orbit (see Definition 1.3 and Definition 1.4). This means that it is possible to travel along geodesics joining points in the orbit of G without leaving a neighborhood of the orbit.

Theorem ((Theorem 6.4)). *Let G be a finitely generated, non-elementary group. Let $\mathcal{X}$ be a quasi-convex G–space. If G contains a strongly contracting element then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

Alternatively, we can assume that the growth rate of the number of orbit points that can be reached by geodesics lying entirely, except near the endpoints, outside a neighborhood of the orbit is strictly smaller than the growth rate of the group:

Theorem ((Theorem 6.3)). *Let G be a finitely generated, non-elementary group. Let $\mathcal{X}$ be a G–space. If G contains a strongly contracting element and there exists a $C \geqslant 0$ such that the C–complementary growth exponent of G is strictly less than the growth exponent of G, then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

See Definition 6.2 for a precise definition of the C–complementary growth exponent.

The proof of Theorem 6.4 is a special case of the proof of Theorem 6.3.

Using Theorem 6.4, we prove:

[1]Grigorchuk and de la Harpe define growth tightness in terms of 'growth rate', which is just the exponentiation of our growth exponent. The growth exponent definition is analogous to the notion of 'volume entropy' familiar in Riemannian geometry, and is more compatible with the Poincaré series in Section 1.2.

THEOREM ((Theorem 8.6)). *If $\mathcal{X}$ is a quasi-convex, relatively hyperbolic G–space and G does not coarsely fix a peripheral subspace then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

This generalizes all previously known results for growth tightness of cocompact actions: every already known example of a group that admits a growth tight action and has some infinite, infinite index normal subgroups is relatively hyperbolic, and, conversely, relatively hyperbolic groups admit growth tight actions [4, 50, 61, 51, 48, 26].

We also use Theorem 6.4 to prove growth tightness for actions on non relatively hyperbolic spaces. For instance, we prove that a group action on a proper CAT(0) space with a rank 1 isometry is growth tight:

THEOREM ((Theorem 9.2)). *If G is a finitely generated, non-elementary group and $\mathcal{X}$ is a quasi-convex, CAT(0) G–space such that G contains an element that acts as a rank 1 isometry on $\mathcal{X}$, then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

Two interesting classes of non relatively hyperbolic groups to which Theorem 9.2 applies are non-elementary Right Angled Artin Groups, which are non relatively hyperbolic when the defining graph is connected, and flip-graph-manifolds. These are the first examples of non relatively hyperbolic groups that admit non-trivial growth tight actions.

THEOREM ((Theorem 9.3)). *Let Θ be a finite graph that is not a join and has more than one vertex. The action of the Right Angled Artin Group G defined by Θ on the universal cover $\mathcal{X}$ of the Salvetti complex associated to Θ is a growth tight action.*

THEOREM ((Theorem 9.4)). *Let M be a flip-graph-manifold. Let G and $\mathcal{X}$ be the fundamental group and universal cover, respectively, of M. Then the action of G on $\mathcal{X}$ by deck transformations is a growth tight action.*

We even exhibit an infinite family of non relatively hyperbolic, non-CAT(0) groups that admit cocompact, growth tight actions:

THEOREM ((Theorem 11.1)). *The Brady-Bridson snowflake groups $BB(1, r)$ for $r \geqslant 3$ admit cocompact, growth tight actions.*

We prove growth tightness for interesting non-quasi-convex actions using Theorem 6.3. We generalize a theorem of Dal'bo, Peigné, Picaud, and Sambusetti [26] for Kleinian groups satisfying an additional Parabolic Gap Condition, see Definition 8.8, to cusp-uniform actions on arbitrary hyperbolic spaces satisfying the Parabolic Gap Condition:

THEOREM ((Theorem 8.9)). *Let G be a finitely generated, non-elementary group. Let $G \curvearrowright \mathcal{X}$ be a cusp uniform action on a hyperbolic space. Suppose that G satisfies the Parabolic Gap Condition. Then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

Once again, our theorems extend beyond actions on relatively hyperbolic spaces, as we use Theorem 6.3 to prove:

THEOREM ((Theorem 10.2)). *The action of the mapping class group of a hyperbolic surface on its Teichmüller space with the Teichmüller metric is a growth tight action.*

Mapping class groups, barring exceptional low complexity cases, are neither relatively hyperbolic nor CAT(0).

In the first seven sections of this paper we prove our main results, Theorem 6.3 and Theorem 6.4. We show in Proposition 3.1 that if there exists a strongly contracting element for $G \curvearrowright \mathcal{X}$ then every infinite normal subgroup Γ contains a strongly contracting element h. We prove growth tightness by bounding the growth exponent of a subset that is orthogonal, in a coarse sense, to every translate of an axis for h.

A dual problem, which is of independent interest, is to find the growth exponent of the conjugacy class of h. In Section 7 we show that the growth exponent of the conjugacy class of a strongly contracting element is exactly half the growth exponent of the group, provided the strongly contracting element moves the base point far enough.

In Sections 8–11 we produce new examples of group actions with strongly contracting elements. These include groups acting on relatively hyperbolic metric spaces (Section 8), certain CAT(0) groups (Section 9), mapping class groups (Section 10), and snowflake groups (Section 11). Our main theorems imply that all these groups admit growth tight actions. These are first examples of growth tight actions and groups which do not come from and are not relatively hyperbolic groups.

0.1. Invariance. Growth tightness is a delicate condition. A construction of Dal'bo, Otal, and Peigné [25], see Observation 8.1, shows that there exist groups G and non-cocompact, hyperbolic, equivariantly quasi-isometric G–spaces $\mathcal{X}$ and $\mathcal{X}'$ such that $G \curvearrowright \mathcal{X}$ is growth tight and $G \curvearrowright \mathcal{X}'$ is not.

In subsequent work [21], we extend the techniques of this paper to produce the first examples of groups that admit a growth tight action on one of their Cayley graphs and a non-growth tight action on another. This answers in the affirmative the following question of Grigorchuk and de la Harpe [32]:

Question 0.4. Does there exist a word metric for which $F_2 \times F_2$ is growth tight?

Recall that $F_2 \times F_2$ is not growth tight with respect to a generating set that is a union of free generating sets of the two factors.

More generally, a product of infinite groups acting on the l^1 product of their Cayley graphs is not growth tight. Such l^1 products and the Dal'bo, Otal, Peigné examples are the only known general constructions of non-growth tight examples. It would be interesting to have a condition to exclude growth tightness. One can not hope to bound the growth exponents of quotients away from that of the group, as Shukhov [55] and Coulon [23] have given examples of hyperbolic groups and sequences of quotients whose growth exponents limit to that of the group. At present, growth tightness can only be excluded for a particular action by exhibiting a quotient of the group by an infinite normal subgroup whose growth exponent is equal to that of the group.

0.2. The Hopf Property. A group G is *Hopfian* if there is no proper quotient of G isomorphic to G.

Let $\mathfrak{D}$ be a set of pseudo-metrics on G that is *quotient-closed*, in the sense that if Γ is a normal subgroup of G such that there exists an isomorphism $\phi\colon\ G \to G/\Gamma$, then for every $d \in \mathfrak{D}$, the pseudo-metric on G obtained by pulling back via ϕ the pseudo-metric on G/Γ induced by d is also in $\mathfrak{D}$. For example, the set of word metrics on G coming from finite generating sets is quotient-closed.

Suppose further that $\mathfrak{D}$ contains a minimal growth pseudo-metric d_0, i.e., $\delta_{G,d_0} = \inf_{d\in\mathfrak{D}} \delta_{G,d}$, and that G is growth tight with respect to d_0.

Proposition 0.5. *Let G be a finitely generated group with a bound on the cardinalities of its finite normal subgroups. Suppose that there exists a quotient-closed set $\mathfrak{D}$ of pseudo-metrics on G that contains a growth tight, minimal growth element d_0 as above. Then G is Hopfian.*

The hypothesis on bounded cardinalities of finite normal subgroups holds for all groups of interest in this paper, see Theorem 1.12.

Proof. Suppose that Γ is a normal subgroup of G such that $G \cong G/\Gamma$. Let d be the pseudo-metric on G obtained from pulling back the pseudo-metric on G/Γ induced by d_0. Since $\mathfrak{D}$ is quotient-closed, $d \in \mathfrak{D}$. By minimality, $\delta_{G,d_0} \leqslant \delta_{G,d}$, but by growth tightness, $\delta_{G,d} \leqslant \delta_{G,d_0}$, with equality only if Γ is finite. Thus, the only normal subgroups Γ for which we could have $G \cong G/\Gamma$ are finite. However, if $G \cong G/\Gamma$ for some finite Γ then G has arbitrarily large finite normal subgroups, contrary to hypothesis. □

Grigorchuk and de la Harpe [32] suggested this as a possible approach to the question of whether a non-elementary Gromov hyperbolic group is Hopfian, in the particular case that $\mathfrak{D}$ is the set of word metrics on G. Arzhantseva and Lysenok [4] proved that every word metric on a non-elementary hyperbolic group is growth tight. They conjectured that the growth exponent of such a group achieves its infinum on some finite generating set and proved a step towards this conjecture [3]. Sambusetti [50] gave an examples of a (non-hyperbolic) group for which the set of word metrics does not realize its infimal growth exponent. In general it is difficult to determine whether a given group has a generating set that realizes the infimal growth exponent among word metrics. Part of our motivation for studying growth tight actions is to open new possibilities for the set $\mathfrak{D}$ of pseudo-metrics considered above.

Torsion free hyperbolic groups are Hopfian by a theorem of Sela [54]. Reinfeldt and Weidmann [47] have announced a generalization of Sela's techniques to hyperbolic groups with torsion, and concluded that all hyperbolic groups are Hopfian.

0.3. The Rank Rigidity Conjecture. The Rank Rigidity Conjecture [20, 8] asserts that if $\mathcal{X}$ is a locally compact, irreducible, geodesically complete CAT(0) space, and G is an infinite discrete group acting properly and cocompactly on $\mathcal{X}$, then one of the following holds:

(1) $\mathcal{X}$ is a higher rank symmetric space.
(2) $\mathcal{X}$ is a Euclidean building of dimension at least 2.
(3) G contains a rank 1 isometry.

In case (1), the Margulis Normal Subgroup Theorem implies that G is trivially growth tight, since it has no infinite, infinite index normal subgroups. Conjecturally, the Margulis Normal Subgroup Theorem also holds in case (2). Our Theorem 9.2 says that if $\mathcal{X}$ is proper then $G \curvearrowright \mathcal{X}$ is a growth tight action in case (3). Thus, a non-growth tight action of a non-elementary group on a proper, irreducible CAT(0) space as above would provide a counterexample either to the Rank Rigidity Conjecture or to the conjecture that the Margulis Normal Subgroup Theorem applies to Euclidean buildings.

The Rank Rigidity Conjecture is known to be true for many interesting classes of spaces, such as Hadamard manifolds [6], 2–dimensional, piecewise-Euclidean cell complexs [7], Davis complexes of Coxeter groups [19], universal covers of Salvetti complexes of Right Angled Artin Groups [9], and finite dimensional CAT(0) cube complexes [20], so Theorem 9.2 provides many new examples of growth tight actions.

It is unclear when growth tightness holds if $\mathcal{X}$ is reducible. A direct product of infinite groups acting via a product action on a product space with the l^1 metric fails to be growth tight. However, there are also examples [18] of infinite simple groups acting cocompactly on products of trees. In [21] we find partial results in the case that the group action is a product action.

0.4. Outline of the Proof of the Main Theorems. Sambusetti [50] proved that a non-elementary free product of non-trivial groups has a greater growth exponent than that of either factor. Thus, a strategy to prove growth tightness is to find a subset of G that looks like a free product, with one factor that grows like the quotient group we are interested in. Specifically:

(1) Find a subset $A \subset G \subset \mathcal{X}$ such that $\delta_A = \delta_{G/\Gamma}$. We will obtain A as a coarsely dense subset of a minimal section of the quotient map $G \to G/\Gamma$, see Definition 4.4.
(2) Construct an embedding of a free product set $A * \mathbb{Z}_2$ into $\mathcal{X}$. The existence of a strongly contracting element $h \in \Gamma$ is used in the construction of this embedding, see Proposition 5.1.
(3) Show that $\delta_{G/\Gamma} = \delta_{A,d_{\mathcal{X}}} < \delta_{A*\mathbb{Z}/2\mathbb{Z},d_{\mathcal{X}}} \leqslant \delta_G$. In this step it is crucial that A is divergent, see Definition 1.7 and Lemma 6.1. We use quasi-convexity/complementary growth exponent to establish divergence.

This outline, due to Sambusetti, is nowadays standard. Typically step (2) is accomplished by a Ping-Pong argument, making use of fine control on the geometry of the space $\mathcal{X}$. Our methods are coarser than such a standard approach, and therefore can be applied to a wider variety of spaces. We use, in particular, a technique of Bestvina, Bromberg, and Fujiwara [11] to construct an action of G on a quasi-tree. Verifying that the map from the free product set into $\mathcal{X}$ is an embedding amounts to showing that elements in A do not cross certain coarse edges of the quasi-tree.

1. Preliminaries

Fix a G–space $\mathcal{X}$. From now on, d is used to denote the metric on $\mathcal{X}$ as well as the induced pseudo-metric on G and G/Γ. Since there will be no possibility of confusion, we suppress d from the growth exponent notation.

We denote by $\mathcal{B}_r(x)$ the open ball of radius r about the point x and by $\mathcal{B}_r(\mathcal{A}) := \cup_{x\in\mathcal{A}}\mathcal{B}_r(x)$ the open r–neighborhood about the set $\mathcal{A}$. The closed r–ball and closed r–neighborhood are denoted $\overline{\mathcal{B}}_r(x)$ and $\overline{\mathcal{B}}_r(\mathcal{A})$, respectively.

1.1. Coarse Language. All of the following definitions may be written without specifying C to indicate that some such $C \geqslant 0$ exists: Two subsets $\mathcal{A}$ and $\mathcal{A}'$ of $\mathcal{X}$ are *C–coarsely equivalent* if $\mathcal{A} \subset \overline{\mathcal{B}}_C(\mathcal{A}')$ and $\mathcal{A}' \subset \overline{\mathcal{B}}_C(\mathcal{A})$. A subset $\mathcal{A}$ of $\mathcal{X}$ is *C–coarsely dense* if it is C–coarsely equivalent to $\mathcal{X}$. A subset $\mathcal{A}$ of $\mathcal{X}$ is *C–coarsely connected* if for every a and a' in $\mathcal{A}$ there exists a chain $a = a_0, a_1, \ldots, a_n = a'$ of points in $\mathcal{A}$ with $d(a_i, a_{i+1}) \leqslant C$.

A *pseudo-map* $\phi\colon \mathcal{X} \to \mathcal{Y}$ assigns to each point in $\mathcal{X}$ a subset $\phi(x)$ of $\mathcal{Y}$. A pseudo-map is *C–coarsely well defined* if for every $x \in \mathcal{X}$ the set $\phi(x)$ of $\mathcal{Y}$ has diameter at most C. Pseudo-maps ϕ and ϕ' with the same domain and codomain are *C–coarsely equivalent* or *C–coarsely agree* if $\phi(x)$ is C–coarsely equivalent to $\phi'(x)$ for every x in the domain. A C–coarsely well defined pseudo-map is called a *C–coarse map*. From a C–coarse map we can obtain a C–coarsely equivalent map by selecting one point from every image set. Conversely:

Lemma 1.1. *If $\phi\colon \mathcal{X} \to \mathcal{Y}$ is coarsely G–equivariant then there is an equivariant coarse map coarsely equivalent to ϕ.*

Proof. Suppose there is a C such that $d(g.\phi(x), \phi(g.x)) \leqslant C$ for all $x \in \mathcal{X}$ and $g \in G$. Define $\phi'(x) := \bigcup_{g\in G} g^{-1}.\phi(g.x)$. Then ϕ' is G–equivariant and C–coarsely equivalent to ϕ. □

Definition 1.2. If $\phi\colon \mathcal{X} \to \mathcal{Y}$ is a pseudo-map and $\mathcal{A}$ and $\mathcal{A}'$ are subsets of $\mathcal{X}$, let $d^{\phi}(\mathcal{A}, \mathcal{A}')$ denote the diameter of $\phi(\mathcal{A}) \cup \phi(\mathcal{A}')$.

Definition 1.3. A subset $\mathcal{A} \subset \mathcal{X}$ is *C–quasi-convex* if for every $a_0, a_1 \in \mathcal{A}$ there exists a geodesic γ between a_0 and a_1 such that $\gamma \subset \overline{\mathcal{B}}_C(\mathcal{A})$. It is *$C$–strongly quasi-convex* if every geodesic with endpoints in $\mathcal{A}$ stays in $\overline{\mathcal{B}}_C(\mathcal{A})$.

Definition 1.4. A G–space $\mathcal{X}$ is *C–quasi-convex* if it contains a C–quasi-convex G–orbit.

For convenience, if $\mathcal{X}$ is a quasi-convex G–space we assume we have chosen a basepoint $o \in \mathcal{X}$ such that $G.o$ is quasi-convex.

A group is *elementary* if it has a finite index cyclic subgroup.

Definition 1.5. Let $g \in G$. The *elementary closure* of g, denoted by $E(g)$, is the largest virtually cyclic subgroup containing g, if such a subgroup exists.

A map $\phi\colon \mathcal{X} \to \mathcal{Y}$ is an *(M, C)–quasi-isometric embedding*, for some $M \geqslant 1$ and $C \geqslant 0$, if, for all $x_0, x_1 \in \mathcal{X}$:

$$\frac{1}{M}d(x_0, x_1) - C \leqslant d(\phi(x_0), (x_1)) \leqslant Md(x_0, x_1) + C$$

A map ϕ is *C–coarsely M–Lipschitz* if the second inequality holds, and is a *quasi-isometry* if it is a quasi-isometric embedding whose image is C–coarsely dense.

An *(M, C)–quasi-geodesic* is an (M, C)–quasi-isometric embedding of a coarsely connected subset of $\mathbb{R}$. If $\gamma\colon I \to \mathcal{X}$ is a quasi-geodesic we let γ_t denote the point $\gamma(t)$, and let γ denote the image of γ in $\mathcal{X}$.

Definition 1.6. A quasi-geodesic Q is *Morse* if for every $M \geqslant 1$ there exists a $K \geqslant 0$ such that every (M, M)–quasi-geodesic with endpoints on Q is contained in the K–neighborhood of Q.

We will use notation to simplify some calculations. Let C be a ‘universal constant’. For us this will usually mean a constant that depends on $G \curvearrowright \mathcal{X}$ and a choice of $o \in \mathcal{X}$, but not on the point in $\mathcal{X}$ at which quantities a and b are calculated.

- For $a \leqslant Cb$ we write $a \stackrel{*}{\prec} b$.
- For $\frac{1}{C}b \leqslant a \leqslant Cb$ we write $a \stackrel{*}{\asymp} b$.
- For $a \leqslant b + C$ we write $a \stackrel{+}{\prec} b$.
- For $b - C \leqslant a \leqslant b + C$ we write $a \stackrel{+}{\asymp} b$.
- For $a \leqslant Cb + C$ we write $a \prec b$.
- For $\frac{1}{C}b - C \leqslant a \leqslant Cb + C$ we write $a \asymp b$.

1.2. Poincaré Series and Growth. Let $(\mathcal{X}, o, d)$ be a pseudo-metric space with choice of basepoint. Let $|x| := d(o, x)$ be the induced semi-norm. Define the *Poincaré series* of $\mathcal{A} \subset \mathcal{X}$ to be

$$\Theta_{\mathcal{A}}(s) := \sum_{a \in \mathcal{A}} \exp(-s|a|)$$

Another related series is:

$$\Theta'_{\mathcal{A}}(s) := \sum_{n=0}^{\infty} \#(\overline{\mathcal{B}}_n(o) \cap \mathcal{A}) \cdot \exp(-sn)$$

The series $\Theta_{\mathcal{A}}$ and $\Theta'_{\mathcal{A}}$ have the same convergence behavior, since $\Theta_{\mathcal{A}}(s) = \Theta'_{\mathcal{A}}(s) \cdot (1 - \exp(-s))$. It follows that the growth exponent of $\mathcal{A}$ is a *critical exponent* for $\Theta'_{\mathcal{A}}$ and $\Theta_{\mathcal{A}}$: the series converge for s greater than the critical exponent and diverge for s less than the critical exponent.

Definition 1.7. $\mathcal{A} \subset \mathcal{X}$ is *divergent* if $\Theta_{\mathcal{A}}$ diverges at its critical exponent.

Since point stabilizers are finite, if $A < G$ and we set $\mathcal{A} := A.o$ then $\Theta_A \stackrel{*}{\asymp} \Theta_{\mathcal{A}}$ and $\Theta'_A \stackrel{*}{\asymp} \Theta'_{\mathcal{A}}$. This implies $\delta_A = \delta_{\mathcal{A}}$, so we can compute the growth exponent of A with respect to the pseudo-metric on A induced by $G \curvearrowright \mathcal{X}$ by computing the growth exponent of the A–orbit as a subset of $\mathcal{X}$.

1.3. The Quasi-tree Construction. We recall the method of Bestvina, Bromberg, and Fujiwara [11] for producing group actions on quasi-trees. A *quasi-tree* is a geodesic metric space that is quasi-isometric to a simplicial tree. Manning [40] gave a characterization of quasi-trees as spaces satisfying a 'bottleneck' property. We use an equivalent formulation:

Definition 1.8 ((Bottleneck Property)). A geodesic metric space satisfies the *bottleneck property* if there exists a number Δ such that for all x and y in $\mathcal{X}$, and for any point m on a geodesic segment from x to y, every path from x to y passes through $\overline{\mathcal{B}}_{\Delta}(m)$.

Theorem 1.9 ([40, Theorem 4.6]). *A geodesic metric space is a quasi-tree if and only if it satisfies the bottleneck property.*

Let $\mathbb{Y}$ be a collection of geodesic metric spaces, and suppose for each $X, Y \in \mathbb{Y}$ we have a subset $\pi_Y(X) \subset Y$, which is referred to as the *projection of X to Y*. Let $d^{\pi}_Y(X, Z) := \operatorname{diam} \pi_Y(X) \cup \pi_Y(Z)$.

Definition 1.10 ((Projection Axioms)). A set $\mathbb{Y}$ with projections as above satisfies the *projection axioms* if there exist $\xi \geqslant 0$ such that for all distinct $X, Y, Z \in \mathbb{Y}$:

(P0) $\operatorname{diam} \pi_Y(X) \leqslant \xi$
(P1) At most one of $d^{\pi}_X(Y, Z)$, $d^{\pi}_Y(X, Z)$, or $d^{\pi}_Z(X, Y)$ is strictly greater than ξ.
(P2) $|\{V \in \mathbb{Y} \mid d^{\pi}_V(X, Y) > \xi\}| < \infty$

For a motivating example, let G be the fundamental group of a closed hyperbolic surface, and let $\mathcal{H}$ be the axis in $\mathbb{H}^2$ of $h \in G$. Let $\mathbb{Y}$ be the distinct G–translates of $\mathcal{H}$, and for each $Y \in \mathbb{Y}$ let π_Y be closest point projection to Y. In this example, projection distances arise as closest point projection in an ambient space containing $\mathbb{Y}$. Bestvina, Bromberg, and Fujiwara consider abstractly the collection $\mathbb{Y}$ and projections satisfying the projection axioms, and build an ambient space containing a copy of $\mathbb{Y}$ such that closest point projection agrees with the given projections, up to bounded error:

THEOREM 1.11 ([11, Theorem A and Theorem B]). *Consider a set $\mathbb{Y}$ of geodesic metric spaces and projections satisfying the projection axioms. There exists a geodesic metric space $\mathcal{Y}$ containing disjoint, isometrically embedded, totally geodesic copies of each $Y \in \mathbb{Y}$, such that for X, $Y \in \mathbb{Y}$, closest point projection of X to Y in $\mathcal{Y}$ is uniformly coarsely equivalent to $\pi_Y(X)$.*

The construction is equivariant with respect to any group action that preserves the projections. Also, if each $Y \in \mathbb{Y}$ is a quasi-tree, with uniform bottleneck constants, then $\mathcal{Y}$ is a quasi-tree.

The basic idea is that Z is 'between' X and Y in $\mathbb{Y}$ if $d^\pi_Z(X, Y)$ is large, and X and Y are 'close' if there is no Z between them. Essentially, $\mathcal{Y}$ is constructed by choosing parameters C and K and connecting every point of $\pi_Y(X)$ to every point of $\pi_X(Y)$ by an edge of length K if there does not exist $Z \in \mathbb{Y}$ with $d^\pi_Z(X, Y) > C$. For technical reasons one actually must perturb the projection distances a bounded amount first. Then, if C is chosen sufficiently large and K is chosen sufficiently large with respect to C, the resulting space is the $\mathcal{Y}$ of Theorem 1.11.

1.4. Hyperbolically Embedded Subgroups. Dahmani, Guirardel, and Osin [24] define the concept of a *hyperbolically embedded subgroup*. This is a generalization of a peripheral subgroup of a relatively hyperbolic group. We will not state the definition, as it is technical and we will not work with this property directly, but it follows from [24, Theorem 4.42] that $E(h)$ is hyperbolically embedded in G for any strongly contracting element h. The proof of this theorem proceeds by considering the action of $E(h)$ on a quasi-tree constructed via the method of Bestvina, Bromberg, and Fujiwara.

We state some results on hyperbolically embedded subgroups that are related to the work in this paper. These are not used in the proofs of the main theorems.

THEOREM 1.12 ([24, Theorem 2.23]). *If G has a hyperbolically embedded subgroup then G has a maximal finite normal subgroup.*

Recall that this theorem guarantees one of the hypotheses of Proposition 0.5.

THEOREM 1.13. *If G contains an infinite order element h such that $E(h)$ is hyperbolically embedded then G has an infinite, infinite index normal subgroup.*

PROOF. By [24, Theorem 5.15], for a sufficiently large n, the normal closure of $\langle h^n \rangle$ in G is the free product of the conjugates of $\langle h^n \rangle$. □

This theorem says that our main results are true for interesting reasons, not simply for lack of normal subgroups.

Minasyan and Osin [42] produce hyperbolically embedded subgroups in certain graphs of groups. We use these to produce growth tight examples in Theorem 9.5.

THEOREM 1.14 ([42, Theorem 4.17]). *Let G be a finitely generated, non-elementary group that splits non-trivially as a graph of groups and is not an ascending HNN-extension. If there exist two edges of the corresponding Bass-Serre tree whose stabilizers have finite intersection then G contains an infinite order element h such that $E(h)$ is hyperbolically embedded in G.*

2. Contraction and Constriction

In this section we introduce properties called 'contracting' and 'constricting' that generalize properties of closest point projection to a geodesic in hyperbolic space, and verify that the 'strong' versions of these properties are sufficient to satisfy the projection axioms of Definition 1.10. These facts are well known to the experts[2], but as there is currently no published general treatment of this material, we provide a detailed account.

[2]For example, [56] shows the projection axioms are satisfied for constricting elements, without assuming that $\mathcal{X}$ is proper.

2.1. Contracting and Constricting. In this section we define contracting and constricting maps and show that the strong versions of these properties are equivalent.

Definition 2.1. A C–coarse map $\pi\colon \mathcal{X} \to \mathcal{A}$ is C*–coarsely a closest point projection* if for all x there exists an $a \in \mathcal{A}$ with $d(x, \mathcal{A}) = d(x, a)$ such that $\operatorname{diam}\{a\} \cup \pi(x) \leqslant C$.

Recall $d^\pi(x_0, x_1) := \operatorname{diam} \pi(x_0) \cup \pi(x_1)$.

Definition 2.2. $\pi\colon \mathcal{X} \to \mathcal{A}$ is (M, C)*–contracting* for $C \geqslant 0$ and $M \geqslant 1$ if

(1) π and $\mathrm{Id}_{\mathcal{A}}$ are C–coarsely equivalent on $\mathcal{A}$, and
(2) $d(x_0, x_1) < \frac{1}{M} d(x_0, \mathcal{A}) - C$ implies $d^\pi(x_0, x_1) \leqslant C$ for all $x_0, x_1 \in \mathcal{X}$.

We say π is *strongly contracting* if it is $(1, C)$–contracting and $d(x, \pi(x)) - d(x, \mathcal{A}) \leqslant C$ for all $x \in \mathcal{X}$.

Another formulation of strong contraction says that geodesics far from $\mathcal{A}$ have bounded projections to $\mathcal{A}$:

Definition 2.3. A coarse map $\pi\colon \mathcal{X} \to \mathcal{A}$ has the *Bounded Geodesic Image Property* if there is a constant C such that for every geodesic $\mathcal{L}$, if $\mathcal{L} \cap \mathcal{B}_C(\mathcal{A}) = \emptyset$ then $\operatorname{diam}(\pi(\mathcal{L})) \leqslant C$.

Lemma 2.4. *If $d(x, \pi(x)) - d(x, \mathcal{A})$ is uniformly bounded then π has the Bounded Geodesic Image Property if and only if it is strongly contracting.*

Proof. First, assume that π has the Bounded Geodesic Image Property, for some constant C. Let x be any point in $\mathcal{X} \setminus \mathcal{B}_C(\mathcal{A})$. For any y such that $d(x, y) < d(x, \mathcal{A}) - C$, every geodesic from x to y remains outside $\mathcal{B}_C(\mathcal{A})$, so its projection has diameter at most C.

For the converse, suppose $\pi\colon \mathcal{X} \to \mathcal{A}$ is a C–coarse map that is $(1, C)$–contracting and $d(x, \pi(x)) - d(x, \mathcal{A}) \leqslant C$ for all $x \in \mathcal{X}$. If $C = 0$ then balls outside of $\mathcal{B}_C(\mathcal{A})$ project to a single point, and we are done, so assume $C > 0$. Let $\mathcal{L}\colon [0, T] \to \mathcal{X}$ be a geodesic that stays outside $\mathcal{B}_{3C}(\mathcal{A})$. Let $t_0 := d(\mathcal{L}_0, \mathcal{A}) - C$, and let $s := T - d(\mathcal{L}_T, \mathcal{A}) + C$. If $s \leqslant t_0$ then $d^\pi(\mathcal{L}_0, \mathcal{L}_T) \leqslant 2C$. Otherwise, define $t_{i+1} := t_i + d(\mathcal{L}_{t_i}, \mathcal{A}) - C$, provided $t_{i+1} < s$. Each $t_{i+1} - t_i \geqslant 2C$, so we have a partition of $[0, T]$ into subintervals $[0, t_0], [t_0, t_1], \ldots, [t_{k-1}, t_k], [t_k, s], [s, T]$ with $k < \frac{s - t_0}{2C}$, and if $[a, b]$ is one of these intervals then $d^\pi(\mathcal{L}_a, \mathcal{L}_b) \leqslant C$, by strong contraction.

Now,

$$\begin{aligned} d(\mathcal{L}_0, \mathcal{L}_T) &\leqslant d(\mathcal{L}_0, \pi(\mathcal{L}_0)) + d(\pi(\mathcal{L}_0), \pi(\mathcal{L}_{t_0})) + d(\pi(\mathcal{L}_{t_0}), \pi(\mathcal{L}_s)) \\ &\qquad + d(\pi(\mathcal{L}_s), \pi(\mathcal{L}_T)) + d(\pi(\mathcal{L}_T), \mathcal{L}_T) \\ &\leqslant d(\mathcal{L}_0, \pi(\mathcal{L}_0)) + d(\pi(\mathcal{L}_T), \mathcal{L}_T) + C(3 + \frac{s - t_0}{2C}), \end{aligned}$$

and

$$\begin{aligned} d(\mathcal{L}_0, \mathcal{L}_T) &= d(\mathcal{L}_0, \mathcal{L}_{t_0}) + d(\mathcal{L}_{t_0}, \mathcal{L}_s) + d(\mathcal{L}_s, \mathcal{L}_T) \\ &= d(\mathcal{L}_0, \mathcal{A}) - C + s - t_0 + d(\mathcal{L}_T, \mathcal{A}) - C, \end{aligned}$$

so

$$s - t_0 \leqslant 2(5C + d(\mathcal{L}_0, \pi(\mathcal{L}_0)) - d(\mathcal{L}_0, \mathcal{A}) + d(\mathcal{L}_T, \pi(\mathcal{L}_T)) - d(\mathcal{L}_T, \mathcal{A})) \leqslant 14C.$$

This means $k < 7$, so $d^\pi(\mathcal{L}_0, \mathcal{L}_T) \leqslant C(3 + k) < 10C$. □

If π is only (M, C)–contracting then a similar argument shows that $d^\pi(\mathcal{L}_0, \mathcal{L}_T)$ is bounded in terms of C and $\log_{\frac{M+1}{M-1}}(d(\mathcal{L}_0, \mathcal{A}) d(\mathcal{L}_T, \mathcal{A}))$.

We now introduce the notion of a constricting map. Using constricting maps will simplify some of our proofs, but it turns out that the strong versions of constricting and contracting are equivalent.

Definition 2.5. A *path system* is a transitive collection of quasi-geodesics with uniform constants that is closed under taking subpaths.

A path system is *minimizing* if, for some $C \geqslant 0$, it contains a path system consisting of $(1, C)$–quasi-geodesics.

Definition 2.6. Let $\mathcal{PS}$ be a path system. A coarse map $\pi\colon \mathcal{X} \to \mathcal{A}$ is (M, C)–$\mathcal{PS}$–*constricting*[3] for $M \geqslant 1$ and $C \geqslant 0$ if:

(1) $\mathcal{PS}$ contains a path system consisting of (M, C)–quasi-geodesics,
(2) π and $\mathrm{Id}_{\mathcal{A}}$ are C–coarsely equivalent on $\mathcal{A}$, and
(3) for every $\mathcal{P} \in \mathcal{PS}$ with endpoints x_0 and x_1, if $d^\pi(x_0, x_1) > C$ then $d(\pi(x_i), \mathcal{P}) \leqslant C$ for both $i \in \{0, 1\}$.

A coarse map is *constricting* if it is (M, C)–$\mathcal{PS}$–constricting for some path system $\mathcal{PS}$ and *strongly constricting* if it is $(1, C)$–constricting for the path system consisting of all geodesics.

Lemma 2.7. *If* $\pi\colon \mathcal{X} \to \mathcal{A}$ *is constricting then it is contracting.*

Proof. Suppose π is (M, C)–$\mathcal{PS}$-constricting C–coarse map for a path system $\mathcal{PS}$ consisting of (M, C)–quasi-geodesics. Suppose $\mathcal{P}\colon [0, T] \to \mathcal{X}$ is a path in $\mathcal{PS}$ with $\mathcal{P}_0 = x$ and $\mathcal{P}_T = y$, and suppose $z = \mathcal{P}_s \in \overline{\mathcal{B}}_C(\mathcal{A})$. Using the fact that $\mathcal{P}$ is an (M, C)–quasi-geodesic on the intervals $[0, T]$, $[0, s]$, and $[s, T]$, one sees that $d(x, y) \geqslant \frac{1}{M^2}(d(x, \mathcal{A}) + d(y, \mathcal{A}) - 4C)$. Therefore, if $d(x, y) < \frac{1}{M^2} d(x, \mathcal{A}) - \frac{4C}{M^2}$ then $\mathcal{P}$ can not enter $\overline{\mathcal{B}}_C(\mathcal{A})$. This would contradict the constricting property, unless $d^\pi(x, y) \leqslant C$. Therefore π is $(M^2, \max\{C, \frac{4C}{M^2}\})$–contracting. □

Lemma 2.8. *Let* $\pi\colon \mathcal{X} \to \mathcal{A}$ *be an* C*–coarse map that is* $(1, C)$*–*$\mathcal{PS}$*–constricting. For all* $x \in \mathcal{X}$ *and all* $r \geqslant 0$ *we have* $\{a \in \mathcal{A} \mid d(x, a) \leqslant d(x, \mathcal{A}) + r\} \subset \{a \in \mathcal{A} \mid d(a, \pi(x)) \leqslant r + 5C\}$.

In particular, setting $r = 0$ *shows that closest point projection to* $\mathcal{A}$ *is coarsely well defined and coarsely equivalent to* π.

Proof. For $x \in \mathcal{X}$ and $r \geqslant 0$, let $a \in \mathcal{A}$ be a point such that $d(x, a) \leqslant d(x, \mathcal{A}) + r$. Let $\mathcal{P}$ be a $(1, C)$–quasi-geodesic from x to a in $\mathcal{PS}$. If $d(a, \pi(x)) > 2C$ then $d^\pi(a, x) > C$, so there is a point $z \in \mathcal{P} \cap \overline{\mathcal{B}}_C(\pi(x))$. Now $d(x, z) + C \geqslant d(x, \pi(x)) \geqslant d(x, \mathcal{A}) \geqslant d(x, a) - r$. Since $\mathcal{P}$ is a $(1, C)$–quasi-geodesic, $d(x, a) \geqslant d(x, z) + d(z, a) - 3C$, so $d(z, a) \leqslant r + 4C$, and $d(a, \pi(x)) \leqslant r + 5C$. □

Proposition 2.9. *Let* $\pi\colon \mathcal{X} \to \mathcal{A}$*. The following are equivalent:*

(1) π *is strongly constricting.*
(2) π *is constricting for some minimizing path system.*
(3) π *is strongly contracting.*
(4) π *has the Bounded Geodesic Image Property and* $d(x, \pi(x)) - d(x, \mathcal{A})$ *is uniformly bounded.*

Proof. (1) implies (2) is immediate.

Suppose π is $(1, C)$–$\mathcal{PS}$–constricting for a minimizing path system $\mathcal{PS}$ consisting of $(1, C)$–quasi-geodesics. Lemma 2.7 shows π is $(1, C')$–contracting. By Lemma 2.8, π is coarsely a closest point projection, so $d(x, \pi(x)) - d(x, \mathcal{A})$ is uniformly bounded. Thus, (2) implies (3).

Now suppose π is $(1, C)$–contracting and $d(x, \pi(x)) - d(x, \mathcal{A}) \leqslant C$ for all $x \in \mathcal{X}$. Take any geodesic $\mathcal{L}\colon [0, T] \to \mathcal{X}$. If $d^\pi(\mathcal{L}_0, \mathcal{L}_T) > 10C$ then $\mathcal{L} \cap \mathcal{B}_{3C}(\mathcal{A}) \neq \emptyset$, as in Lemma 2.4. Let $t = t_0$, t_1 be the first and last times, respectively, such that $d(\mathcal{L}_t, \mathcal{A}) \leqslant 3C$. By Lemma 2.4, $d^\pi(\mathcal{L}_0, \mathcal{L}_{t_0}) \leqslant 10C$. Thus, $d(\pi(\mathcal{L}_0), \mathcal{L}_{t_0}) \leqslant d^\pi(\mathcal{L}_0, \mathcal{L}_{t_0}) + d(\pi(\mathcal{L}_{t_0}), \mathcal{L}_{t_0}) \leqslant 14C$. The same argument shows $d(\pi(\mathcal{L}_T), \mathcal{L}_{t_1}) \leqslant 14C$, so π is $(1, 14C)$–constricting for the path system of all geodesics. Thus, (3) implies (1).

(3) is equivalent to (4) by Lemma 2.4. □

2.2. Additional Properties of Contracting and Constricting Maps. We establish some properties of contracting and constricting maps that will be useful in the sequel.

Lemma 2.10. *If* π *is a* $(1, C)$*–strongly constricting* C*–coarse map and* $d^\pi(x, y) > C$ *then* $d(x, y) \geqslant d(x, \pi(x)) + d^\pi(x, y) + d(\pi(y), y) - 6C$.

[3] Sisto [56] calls this property '$\mathcal{PS}$–contracting'. We change the name to avoid conflict with the better established 'contracting' terminology of Definition 2.2.

Proof. Let $\mathcal{L}$ be a geodesic from x to y. By strong constriction, there exist s and t such that $d(\mathcal{L}_s, \pi(x)) \leqslant C$ and $d(\mathcal{L}_t, \pi(y)) \leqslant C$. The lemma follows from the triangle inequality and the fact that $\pi(x)$ and $\pi(y)$ have diameter at most C. □

Lemma 2.11. *If $\pi\colon \mathcal{X} \to \mathcal{A}$ is strongly constricting then it is coarsely 1–Lipschitz.*

Proof. Let π be an C–coarse map that is $(1, C)$–constricting on the path system of geodesics. Let x_0 and x_1 be arbitrary points, and let $\mathcal{L}$ be a geodesic from x_0 to x_1. If $d^\pi(x_0, x_1) > 4C$ then $\mathcal{L} \cap \mathcal{B}_C(x_i) \neq \emptyset$ for each i, which implies $d(x_0, x_1) \geqslant d(x_0, \pi(x_0)) + d^\pi(x_0, x_1) + d(\pi(x_1), x_1) - 8C$. Thus, for all x_0 and x_1, we have $d^\pi(x_0, x_1) \leqslant d(x_0, x_1) + 8C$. □

Lemma 2.12. *Let $\pi\colon \mathcal{X} \to \mathcal{A}$ be an (M, C)–contracting C–coarse map such that $d(x, \pi(x)) - d(x, \mathcal{A}) \leqslant C$ for all $x \in \mathcal{X}$. Fix $K \geqslant 1$. For all sufficiently large D there exists a T_{max} such that if $Q\colon [0, T] \to \mathcal{X}$ is a (K, K)–quasi-geodesic with $d(Q_0, A) = D = d(Q_T, \mathcal{A})$ and $Q \cap \mathcal{B}_D(\mathcal{A}) = \emptyset$ then $T \leqslant T_{max}$.*

Proof. Let $D > M(K^2C + C + K)$. Let $t_0 := 0$ and let t_{i+1} be the last time that $d(Q_{t_i}, Q_{t_{i+1}}) = \frac{1}{M} d(Q_{t_i}, \mathcal{A}) - C$, provided $t_{i+1} < T$. This subdivides $[0, T]$ into at most $1 + \frac{TK}{\frac{D}{M} - C - K}$ intervals $[t_0, t_1], \ldots, [t_k, T]$, each of which has endpoints whose π–images are distance at most C apart.

Since Q is a quasi-geodesic, $T \leqslant K d(Q_0, Q_T) + K^2$. On the other hand:

$$d(Q_0, Q_T) \leqslant 2D + 2C + d^\pi(Q_0, Q_T) \leqslant 2D + 2C + C\left(1 + \frac{TK}{\frac{D}{M} - C - K}\right)$$

Combined with the condition on D, this yields an upper bound on T. □

Corollary 2.13. *If $\pi\colon \mathcal{X} \to \mathcal{A}$ is contracting and $d(x, \pi(x)) - d(x, \mathcal{A})$ is uniformly bounded, then for all $M \geqslant 1$ and $D \geqslant 0$ there exists a K such that every (M, M)–quasi-geodesic with endpoints at distance at most D from $\mathcal{A}$ is contained in $\overline{\mathcal{B}}_K(\mathcal{A})$.*

In particular, if $\mathcal{A}$ is a quasi-geodesic then it is Morse.

Lemma 2.14. *Let $Q\colon \mathbb{R} \to \mathcal{X}$ be a quasi-geodesic, and let $\pi\colon \mathcal{X} \to Q$ be a strongly contracting projection. For all $D \geqslant 0$ there exists a K such that if $\mathcal{P}\colon [0, T] \to \mathcal{X}$ is a geodesic and t_0 and t_1 are such that $d(\mathcal{P}_0, Q_{t_0}) \leqslant D$ and $d(\mathcal{P}_T, Q_{t_1}) \leqslant D$ then $Q_{[t_0, t_1]} \subset \overline{\mathcal{B}}_K(\mathcal{P})$.*

Proof. By Proposition 2.9, π is strongly constricting, so $\mathcal{P}$ passes close to every point in $\pi(\mathcal{P})$. Let i and j be numbers in the domain of $\mathcal{P}$, with $0 < j - i \leqslant 1$. Let s_i and s_j be such that $Q_{s_i} \in \pi(\mathcal{P}_i)$ and $Q_{s_j} \in \pi(\mathcal{P}_j)$. Then s_i and s_j are boundedly far apart, since π is coarsely 1-Lipschitz, by Lemma 2.11, and Q is a quasi-geodesic. Therefore, the diameter of $Q_{[s_i, s_j]}$ is bounded, and we have already noted that $Q(s_i)$ and $Q(s_j)$ are close to $\mathcal{P}$, since they are in the image of π. □

Lemma 2.15. *Let $\mathcal{A}$ and $\mathcal{A}'$ be coarsely equivalent subsets of $\mathcal{X}$. Let $\sigma\colon \mathcal{A} \to \mathcal{A}'$ and $\bar{\sigma}\colon \mathcal{A}' \to \mathcal{A}$ be C–coarse maps such that $d(a, \sigma(a)) \leqslant C$ for all $a \in \mathcal{A}$ and $d(a', \bar{\sigma}(a')) \leqslant C$ for all $a' \in \mathcal{A}'$. Then $\pi_{\mathcal{A}}\colon \mathcal{X} \to \mathcal{A}$ is strongly contracting if and only if $\pi_{\mathcal{A}'} := \sigma \circ \pi_{\mathcal{A}}\colon \mathcal{X} \to \mathcal{A}'$ is strongly contracting.*

Proof. Suppose $\pi_{\mathcal{A}}$ is $(1, C)$–contracting and $d(x, \pi(x)) - d(x, \mathcal{A}) \leqslant C$ for all $x \in \mathcal{X}$. If $d(x, y) \leqslant d(x, \mathcal{A}') - 2C \leqslant d(x, \mathcal{A}) - C$ then $d^\pi_{\mathcal{A}'}(x, y) \leqslant d^\pi_{\mathcal{A}}(x, y) + 2C \leqslant 3C$, so $\pi_{\mathcal{A}'}$ is $(1, 3C)$–contracting.

Take a point x and let $a' \in \mathcal{A}'$ such that $d(x, \mathcal{A}') = d(x, a')$. Then $d(x, \bar{\sigma}(a')) - C \leqslant d(x, a') \leqslant d(x, \pi_{\mathcal{A}'}(x)) \leqslant d(x, \pi_{\mathcal{A}}(x)) + 2C$, so $d(x, \bar{\sigma}(a')) \leqslant d(x, \mathcal{A}) + 3C$. By Proposition 2.9, $\pi_{\mathcal{A}}$ is strongly constricting, so by Lemma 2.8, there is a constant D such that $d(\pi_{\mathcal{A}}(x), \bar{\sigma}(a')) \leqslant 3C + D$. Thus, $\pi_{\mathcal{A}'}$ is $(5C + D)$–coarsely a closest point projection, hence, strongly contracting. □

Lemma 2.16. *Let $\pi\colon \mathcal{X} \to \mathcal{A}$ be strongly constricting. There exists a number K such that if $d(\mathcal{A}, g\mathcal{A}) > K$ then* diam $\pi(g\mathcal{A})$ *is bounded, independent of g.*

Proof. Let π be $(1, C)$–strongly constricting. By Proposition 2.9, π is strongly contracting, so by Corollary 2.13 there is a constant K such that a geodesic with endpoints in $\mathcal{A}$ stays in the

$(K-C)$–neighborhood of $\mathcal{A}$. Therefore, a geodesic with endpoints in $g\mathcal{A}$ stays in $\overline{\mathcal{B}}_{K-C}(g\mathcal{A})$. Choose $x \in g\mathcal{A}$ such that $d(x,\mathcal{A}) = d(g\mathcal{A},\mathcal{A})$. For all $y \in g\mathcal{A}$, if $d^\pi(x,y) > C$ then a geodesic from x to y passes within distance C of $\pi(x)$ and $\pi(y)$. This means $\overline{\mathcal{B}}_C(\mathcal{A}) \cap \overline{\mathcal{B}}_{K-C}(g\mathcal{A}) \neq \emptyset$, so $d(\mathcal{A}, g\mathcal{A}) \leqslant K$. Thus, if $d(\mathcal{A}, g\mathcal{A}) > K$, then $d^\pi(x,y) \leqslant C$, so $\operatorname{diam}\pi(g\mathcal{A}) \leqslant 2C$. □

2.3. Strongly Contracting Elements. We have defined contraction and constriction for maps. We now give definitions for group elements:

Definition 2.17. An element $h \in G$ is called *contracting*, with respect to $G \curvearrowright \mathcal{X}$, if $i \mapsto h^i.o$ is a quasi-geodesic and if there exists a subset $\mathcal{A} \subset \mathcal{X}$ on which $\langle h\rangle$ acts cocompactly and a map $\pi\colon \mathcal{X} \to \mathcal{A}$ that is contracting.

An element $h \in G$ is called *constricting*, with respect to $G \curvearrowright \mathcal{X}$, if $i \mapsto h^i.o$ is a quasi-geodesic and if there exists a subset $\mathcal{A} \subset \mathcal{X}$ on which $\langle h\rangle$ acts cocompactly, a G–invariant path system $\mathcal{PS}$, and a map $\pi\colon \mathcal{X} \to \mathcal{A}$ that is $\mathcal{PS}$–constricting.

An element is *strongly contracting* or *strongly constricting* if the projection π is, respectively, strongly contracting or strongly constricting.

For π and $\mathcal{A}$ as in the definition, Proposition 2.9 says π is strongly contracting if and only if it is strongly constricting. Thus, Lemma 2.8 says closest point projection to $\mathcal{A}$ is coarsely well defined and coarsely equivalent to π. Lemma 2.15 says that the choice of the set $\mathcal{A}$ only affects the constants of strong contraction. It follows that an element h is strongly contracting if and only if $i \mapsto h^i.o$ is a quasi-geodesic and closest point projection to $\langle h\rangle.o$ is strongly contracting. In the remainder of this section we produce more finely tailored choices for $\mathcal{A}$ and π. In particular, we would like π to be compatible with the group action, see Remark 2.22.

Proposition 2.18 ((cf. [24, Lemma 6.5])). *Let G be a finitely generated group, and let $\mathcal{X}$ be a G–space. Let $h \in G$ be an infinite order element. If there exists a strongly constricting $\pi\colon \mathcal{X} \to \langle h\rangle.o$ then:*

$$E(h) = H := \{g \in G \mid g\langle h\rangle.o \text{ is coarsely equivalent to } \langle h\rangle.o\}$$

Proof. H is a group containing every finite index supergroup of $\langle h\rangle$. Let D be the constant of Lemma 2.16, and let $S := \{g \in G \mid d(g\langle h\rangle.o, \langle h\rangle.o) \leqslant D\}$. Then Lemma 2.16 implies $H \subset S$. Since $G \curvearrowright \mathcal{X}$ is properly discontinuous, S is contained in finitely many h–orbits, so $\langle h\rangle < H$ has finite index. Therefore, $E(h)$ exists and is equal to H. □

Definition 2.19. If h is a strongly contracting element, define the *(quasi)-axis* of h, with respect to the basepoint o, to be $\mathcal{H} := E(h).o$.

Lemma 2.20. *If h is a strongly contracting element then there exists an $E(h)$–equivariant, strongly contracting coarse map $\pi_{\mathcal{H}}\colon \mathcal{X} \to \mathcal{H}$.*

Proof. By Proposition 2.9, Lemma 2.8, and Lemma 2.15 any choice of closest point projection map to $\mathcal{H}$ is strongly contracting and coarsely $E(h)$–equivariant, so, by Lemma 1.1, we can replace it by a coarsely equivalent, $E(h)$–equivariant coarse map, which will still be strongly contracting, by Lemma 2.15. □

Definition 2.21. From the projection $\pi_{\mathcal{H}}$ of Lemma 2.20 define strongly contracting projections onto each translate of $\mathcal{H}$ by $\pi_{g\mathcal{H}}\colon \mathcal{X} \to g\mathcal{H} : x \mapsto g.\pi_{\mathcal{H}}(g^{-1}.x)$.

If $g'\mathcal{H} = g\mathcal{H}$ then $g^{-1}g' \in E(h)$ so Lemma 2.20 implies $\pi_{g'\mathcal{H}}(x) = \pi_{g\mathcal{H}}(x)$ for all $x \in \mathcal{X}$.

Remark 2.22. The projections of Definition 2.21 satisfy $g.\pi_{\mathcal{H}}(x) = \pi_{g\mathcal{H}}(g.x)$ for all $x \in \mathcal{X}$ and $g \in G$.

2.4. Strongly Contracting Elements and the Projection Axioms. Let $h \in G$ be a strongly contracting element with respect to $G \curvearrowright \mathcal{X}$. Let $\mathcal{H}$ be a quasi-axis of h defined in Definition 2.19. We wish to apply Theorem 1.11 to the collection of G–translates of $\mathcal{H}$ with the projections of Definition 2.21. To see that the hypotheses of the theorem are satisfied, we first embed $\mathcal{H}$ into a geodesic metric space and then verify the projection axioms of Definition 1.10.

Choose representatives $1 = g_0, \dots, g_{n-1}$ for $\langle h\rangle\backslash E(h)$, so that for each i we have $d(g_i.o, o) = \min_{g\in\langle h\rangle g_i} d(g.o, o)$. Let $g_n := h$. Let $\hat{\mathcal{H}}$ be the Cayley graph of $E(h)$ with respect to the generating set $\{g_1, \dots, g_n\}$. The graph $\hat{\mathcal{H}}$ becomes a geodesic metric space by assigning each edge length one, and it is a quasi-tree since $E(h)$ is virtually cyclic.

Choose representatives $1 = f_0, f_1, \dots$ for $G/E(h)$. Let $\mathbb{Y}$ be a disjoint union of copies of $\hat{\mathcal{H}}$, one for each $f_iE(h) \in G/E(h)$, denoted $f_i\hat{\mathcal{H}}$. The orbit map $f_i\hat{\mathcal{H}} \to f_i\mathcal{H} := f_ie \mapsto f_ie.o$ is a quasi-isometric embedding, so its inverse $\phi_{f_i\mathcal{H}}\colon f_i\mathcal{H} \to f_i\hat{\mathcal{H}}$ is a coarse map that is a quasi-isometry. Define $\pi_{f_i\hat{\mathcal{H}}}(f_j\hat{\mathcal{H}}) := \phi_{f_i}(\pi_{f_i\mathcal{H}}(f_j\mathcal{H}))$. Since ϕ_{f_i} is a quasi-isometry it suffices to check the projection axioms on translates of $\mathcal{H}$ in $\mathcal{X}$.

Lemma 2.23 (Axiom (P0)). *There is a uniform bound on the diameter of $\pi_{\mathcal{H}}(g\mathcal{H})$ for $g \notin E(h)$.*

Proof. Let $\pi_{\mathcal{H}}\colon \mathcal{X} \to \mathcal{H}$ be $(1, C')$–strongly constricting. Let $\mathcal{Q}\colon \mathbb{R} \twoheadrightarrow \mathcal{H}$ be an (M, C'')–quasi-geodesic parameterization that agrees with $i \mapsto h^i.o$ on the integers. Replace C' and C'' by $C := \max\{C', C''\}$.

Let $D := \operatorname{diam}\langle h\rangle\backslash\mathcal{H}$. Let K be large enough so that if $\mathcal{P}$ is a geodesic with $d(\mathcal{P}_{s_0}, \mathcal{Q}_{t_0}) \leqslant C$ and $d(\mathcal{P}_{s_1}, \mathcal{Q}_{t_1}) \leqslant C$ then $\mathcal{P}_{[s_0,s_1]} \subset \overline{\mathcal{B}}_K(\mathcal{Q}_{[t_0,t_1]})$ and $\mathcal{Q}_{[s_0,s_1]} \subset \overline{\mathcal{B}}_K(\mathcal{P}_{[t_0,t_1]})$, as in Corollary 2.13 and Lemma 2.14.

Suppose $g \notin E(h)$. For a pair of points x_0, $x_1 \in g\mathcal{H}$, take t_0 and t_1 such that $\mathcal{Q}_{t_i} \in \pi_{\mathcal{H}}(x_i)$ for each i. Let $\mathcal{P}$ be a geodesic connecting x_0 to x_1. If $d^{\pi}_{\mathcal{H}}(x_0, x_1) > C$ then for each i there exists s_i such that $d(\mathcal{P}_{s_i}, \mathcal{Q}_{t_i}) \leqslant C$.

Now $\mathcal{Q}_{[t_0,t_1]}$ is K–close to $\mathcal{P}_{[s_0,s_1]}$, which in turn is K–close to a subinterval of $g\mathcal{H}$. Therefore, for each integer $i \in [t_0, t_1]$ there is an integer α_i such that $d(h^i.o, gh^{\alpha_i}g^{-1}.o) \leqslant 2K + D$.

If for some $i \neq j$ we have $h^{-i}gh^{\alpha_i}g^{-1}.o = h^{-j}gh^{\alpha_j}g^{-1}.o$, then $h^{j-i} = gh^{\alpha_j-\alpha_i}g^{-1}$, which implies $\langle h\rangle$ and $\langle ghg^{-1}\rangle$ are commensurable. However, this would imply $g \in E(h)$, contrary to hypothesis. Therefore, for each integer i in $[t_0, t_1]$ we get a distinct point $h^{-i}gh^{\alpha_i}g^{-1}.o \in \overline{\mathcal{B}}_{2K+D}(o)$. Since the action of G is properly discontinuous, the number of orbit points in $\overline{\mathcal{B}}_{2K+D}(o)$ is finite, so $\operatorname{diam} \pi_{\mathcal{H}}(g\mathcal{H})$ is bounded, independent of g. □

Lemma 2.24 (Axiom (P1)). *For all sufficiently large ξ and for any X, Y, $Z \in \mathbb{Y}$, at most one of $d^{\pi}_X(Y, Z)$, $d^{\pi}_Y(X, Z)$, and $d^{\pi}_Z(X, Y)$ is greater than ξ.*

Proof. Suppose π_Y is $(1, C)$–strongly constricting. Let ξ' be the constant from Lemma 2.23. Let $\xi \geqslant 2\xi' + 14C$. Suppose that $d^{\pi}_X(Y, Z) > \xi$. We show $d^{\pi}_X(Y, Z) \leqslant \xi$; the inequality $d^{\pi}_Z(X, Y) \leqslant \xi$ follows by a similar argument.

Take any point $z \in Z$, and let $y \in Y$ be a point such that $d(z, y) = d(z, Y)$. Let $\mathcal{L}\colon [0, T] \to \mathcal{X}$ be a geodesic from z to y. For every point of $\mathcal{L}$, y is the closest point of Y. By Lemma 2.8, $\pi_Y(\mathcal{L}) \subset \overline{\mathcal{B}}_{5C}(y)$. Now, $d^{\pi}_X(Y, Z) > \xi$ implies $d^{\pi}_X(\mathcal{L}_0, \mathcal{L}_T) > C$, so there is a $z' \in \mathcal{L}$ and $x \in X$ with $d(x, z') \leqslant D$. By Lemma 2.11, π_Y is $8C$–coarsely 1–Lipschitz, which means $d^{\pi}_Y(x, z') \leqslant 9C$. Thus, $d^{\pi}_Y(X, Z) \leqslant 2\xi' + d^{\pi}_Y(x, z) \leqslant 2\xi' + 5C + d^{\pi}_Y(x, z') \leqslant 2\xi' + 14C \leqslant \xi$. □

Lemma 2.25 (Axiom (P2)). *For all sufficiently large ξ and for all X, $Y \in \mathbb{Y}$, the set $\{V \in \mathbb{Y} \mid d^{\pi}_V(X, Y) > \xi\}$ is finite.*

Proof. Let ξ' be the constant of Lemma 2.23. Suppose $\pi_{\mathcal{H}}$ is $(1, C)$–strongly constricting. Let $\xi > C + 2\xi'$. Take arbitrary X, $Y \in \mathbb{Y}$, and let $\mathcal{L}$ be a geodesic from some point in $\pi_X(Y)$ to some point in $\pi_Y(X)$. If $d^{\pi}_V(X, Y) > \xi$ then $d^{\pi}_V(\mathcal{L}_0, \mathcal{L}_T) > C$, so $\mathcal{L}$ comes within distance C of V. By proper discontinuity of the action, there are only finitely many elements of $\mathbb{Y}$ that come within distance C of the finite geodesic $\mathcal{L}$. □

Definition 2.26. Let $\mathcal{Y}$ be the quasi-tree produced by Theorem 1.11 from $\mathbb{Y}$. Let $\star \in \mathcal{Y}$ be the vertex corresponding to $o \in \mathcal{X}$. Let $\hat{\pi}_{g\hat{\mathcal{H}}}\colon \mathcal{Y} \to g\hat{\mathcal{H}}$ be closest point projection to the isometrically embedded copy of $g\hat{\mathcal{H}}$ in $\mathcal{Y}$, which the theorem says coarsely agrees with $\pi_{g\hat{\mathcal{H}}}$.

Definition 2.27. Define uniform quasi-isometric embeddings $\phi_{g\mathcal{H}}\colon g\mathcal{H} \to \mathcal{Y}$ for each translate $g\mathcal{H}$ of $\mathcal{H}$ by sending $g\mathcal{H}$ to $f_i\hat{\mathcal{H}}$ via ϕ_{f_i}, where $g \in f_iE(h)$, and postcomposing by the isometric embedding of $f_i\hat{\mathcal{H}}$ into $\mathcal{Y}$ provided by Theorem 1.11.

Proposition 2.28. *If there is a strongly contracting element for $G \curvearrowright \mathcal{X}$ then G has non-zero growth exponent.*

Proof. [11, Proposition 3.23] says G contains a free subgroup, so it has exponential growth. □

3. Abundance of Strongly Contracting Elements

In this section we show that strongly contracting elements are abundant:

Proposition 3.1. *If G contains a strongly contracting element for $G \curvearrowright \mathcal{X}$ then so does every infinite normal subgroup.*

In effect, the proposition reduces the problem of growth tightness for arbitrary quotients of G to quotients by the normal closure of a strongly contracting element.

Given a strongly contracting element $h \in \mathcal{G}$ and an infinite normal subgroup Γ of G we find an element $g \in \Gamma$ such that $f := gh^ng^{-1}h^{-n} \in \Gamma$ is strongly contracting for all sufficiently large n. To prove f is strongly contracting we follow a standard strategy by showing that an axis for f has 'long' ($\asymp n$) segments in contracting sets, separated by 'short' ($= d(o, g.o)$) hops between such segments. For each $x \in \mathcal{X}$ there is, coarsely, a unique one of these segments such that the projection of x transitions from landing at the end of the segment to landing at the beginning of the segment. We use this transition point to define the projection to the f–axis, and verify that this projection is strongly contracting.

We first prove some preliminary lemmas.

Lemma 3.2. *Let $h \in G$ be an infinite order element and $\pi\colon \mathcal{X} \to \langle h\rangle.o$ a contracting coarse map such that $d(x, \pi(x)) - d(x, \mathcal{A})$ is uniformly bounded. Then $i \mapsto h^i.o$ is a quasi-geodesic.*

Proof. Take any $\alpha < \beta$ in $\mathbb{Z}$. By the triangle inequality, $d(h^\alpha.o, h^\beta.o) \overset{*}{\prec} (\beta - \alpha)$. We now prove the opposite inequality. Let $\mathcal{L}\colon [0, T] \to \mathcal{X}$ be a geodesic from $h^\alpha.o$ to $h^\beta.o$. By Corollary 2.13, there exists a D such that for every $i \in [0, T] \cap \mathbb{Z}$ there exists an $\alpha \leqslant \alpha_i \leqslant \beta$ such that $d(\mathcal{L}_i, h^{\alpha_i}.o) \leqslant D$. Since the action of G on $\mathcal{X}$ is properly discontinuous, there exists a maximum γ such that $d(o, h^\gamma.o) \leqslant 2D + 1$, so $\alpha_{i+1} - \alpha_i \leqslant \gamma$ for all i. Setting $\alpha_0 := \alpha$ and $\alpha_{\lceil T\rceil} := \beta$, we have $\beta - \alpha = \sum_{i=0}^{\lceil T\rceil-1} \alpha_{i+1} - \alpha_i \leqslant \gamma\lceil T\rceil \leqslant \gamma(d(h^\alpha.o, h^\beta.o) + 1)$. □

Fix a strongly contracting element h, and let $\mathcal{Y}$ be the quasi-tree of Definition 2.26, with bottleneck constant Δ.

Lemma 3.3. *There exists $K \geqslant 0$ such that $d^\pi_{\mathcal{H}}(o, g_1.o) - d^\pi_{\mathcal{H}}(g_1.o, g_0.o) \geqslant K$ implies $g_0.\star$ and $g_1.\star$ are contained in the same component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(\star)$.*

Proof. Let $D := \operatorname{diam}\langle h\rangle\backslash\hat{\mathcal{H}}$ in $\mathcal{Y}$. For each $i \in \{0, 1\}$, choose an m_i such that we have $d(h^{m_i}.\star, \hat{\pi}_{\hat{\mathcal{H}}}(g_i.\star)) \leqslant D$. Choose a geodesic $\mathcal{L}$ from $\star$ to $h.\star$. Take $M > 0$ such that $h^m.\mathcal{L} \cap \overline{\mathcal{B}}_\Delta(\star) = \emptyset$ when $|m| \geq M$.

For each i, $|m_i| \asymp d(h^{m_i}.\star, \star) \succ K$, so for sufficiently large K we have $d(h^{m_i}.\star, \star) > 2\Delta + D$ and $|m_i| > M$. Furthermore, m_0 and m_1 must have the same sign if K is large enough: by Lemma 2.14, the interval of $\mathcal{H}$ between $h^{m_0}.o$ and $h^{m_1}.o$ stays close to a geodesic between $h^{m_0}.o$ and $h^{m_1}.o$, so if m_0 and m_1 have different signs:

$$d^\pi_{\mathcal{H}}(g_0.o, g_1.o) \overset{+}{\asymp} d(h^{m_0}.o, h^{m_1}.o) \overset{+}{\asymp} d(o, h^{m_0}.o) + d(o, h^{m_1}.o) \overset{+}{\asymp} d^\pi_{\mathcal{H}}(o, g_0.o) + d^\pi_{\mathcal{H}}(o, g_1.o)$$

However, $d^\pi_{\mathcal{H}}(g_0.o, g_1.o) \leqslant d^\pi_{\mathcal{H}}(o, g_1.o) - K$, so this would imply

$$K \overset{+}{\prec} d^\pi_{\mathcal{H}}(o, g_0.o) \overset{+}{\prec} -K,$$

which is false for sufficiently large K.

No geodesic between $g_i.\star$ and $h^{m_i}.\star$ enters $\overline{\mathcal{B}}_\Delta(\star)$, since this would imply:

$$d(h^{m_1}.\star, \star) \leqslant 2\Delta + D$$

For $\min\{m_0, m_1\} \leqslant m \leqslant \max\{m_0, m_1\} - 1$ the geodesic $h^m.\mathcal{L}$ stays outside $\overline{\mathcal{B}}_\Delta(\star)$ since m_0 and m_1 have the same sign and magnitude at least M, which implies $|m| \geqslant M$.

By concatenating such geodesics, we construct a path from $g_0.\star$ to $g_1.\star$ in $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(\star)$. ◻

Corollary 3.4. *There exists an $N > 0$ such that for all $n \geqslant N$ the points $h^n.\star$ and $h^N.\star$ are in the same component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(\star)$.*

Proof. Take N large enough so that $d^\pi_{\mathcal{H}}(o, h^n.o) \geqslant K + d(o, h.o) + 2C$ for all $n \geqslant N$. Then $d^\pi_{\mathcal{H}}(o, h^{n+1}.o) - d^\pi_{\mathcal{H}}(h^n.o, h^{n+1}.o) \geqslant K$. Apply Lemma 3.3. ◻

Definition 3.5. Call the component of $\mathcal{Y} \setminus \mathcal{B}_\Delta(g.\star)$ containing $gh^n.\star$ for all sufficiently large n *the gh^∞ component* and the component containing $gh^{-n}.\star$ for all sufficiently large n *the $gh^{-\infty}$ component*.

Lemma 3.6. *For some $K \geqslant 0$ suppose g_0 and g_1 are elements of G such that $g_0\mathcal{H} \neq g_1\mathcal{H}$ and $d^\pi_{g_0\mathcal{H}}(g_0.o, g_1.o) \leqslant K$ and $d^\pi_{g_1\mathcal{H}}(g_0.o, g_1.o) \leqslant K$. Then there exists an $N > 0$ such that for all $n \geqslant N$, ϵ_0, $\epsilon_1 \in \{\pm 1\}$, and $f_0, f_1 \in \{g_0, g_1\}$*

- *the balls $\overline{\mathcal{B}}_\Delta(f_0 h^{\epsilon_0 n/2}.\star)$ and $\overline{\mathcal{B}}_\Delta(f_1 h^{\epsilon_1 n/2}.\star)$ in $\mathcal{Y}$ are disjoint unless $f_0 = f_1$ and $\epsilon_0 = \epsilon_1$,*
- *$f_0.\star$ and $f_1.\star$ are in the $f_0 h^{-\epsilon_0\infty}$ component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(f_0 h^{\epsilon_0 n/2}.\star)$, and*
- *$f_0 h^{\epsilon_0 n}.\star$ and $f_0 h^{\epsilon_0 n} f_1.\star$ are in the $f_0 h^{\epsilon_0\infty}$ component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(f_0 h^{\epsilon_0 n/2}.\star)$.*

Proof. $\overline{\mathcal{B}}_\Delta(f_0 h^{n/2}.\star)$ and $\overline{\mathcal{B}}_\Delta(f_0 h^{-n/2}.\star)$ are disjoint for all sufficiently large n since $i \mapsto h^i.\star$ is a quasi-geodesic. In the other cases, $f_0\mathcal{H}$ and $f_1\mathcal{H}$ are distinct axes, so $f_0\hat{\mathcal{H}}$ and $f_1\hat{\mathcal{H}}$ are disjoint. For each $i \in \{0, 1\}$, the bounds $d^\pi_{f_i\mathcal{H}}(f_i.o, f_{1-i}.o) \leqslant K$ imply that the closest point projection $\hat{\pi}_{f_i\hat{\mathcal{H}}}(f_{1-i}\hat{\mathcal{H}})$ of $f_{1-i}\hat{\mathcal{H}}$ to $f_i\hat{\mathcal{H}}$ is contained in a bounded neighborhood of $f_i.\star$. For any point $y.\star \in \overline{\mathcal{B}}_\Delta(f_1 h^{\epsilon_1 n/2}.\star) \setminus f_1\hat{\mathcal{H}}$, we have that $\hat{\pi}_{f_1\hat{\mathcal{H}}}(y\hat{\mathcal{H}})$ is 2Δ–close to $f_1 h^{\epsilon_1 n/2}.\star$. Therefore

$$d^\pi_{f_1\hat{\mathcal{H}}}(f_0\hat{\mathcal{H}}, y\hat{\mathcal{H}}) \stackrel{+}{\asymp} d(f_1.\star, f_1 h^{\epsilon_1 n/2}.\star) \asymp n,$$

so for n sufficiently large we can make $d^\pi_{f_1\hat{\mathcal{H}}}(f_0\hat{\mathcal{H}}, y\hat{\mathcal{H}})$ larger then the constant ξ of projection axiom (P1), which implies $d^\pi_{f_0\hat{\mathcal{H}}}(f_1\hat{\mathcal{H}}, y\hat{\mathcal{H}}) \leqslant \xi$. On the other hand, $\overline{\mathcal{B}}_\Delta(f_0 h^{\epsilon_0 n/2}.\star)$ projects close to $f_0 h^{\epsilon_0 n/2}.\star$ in $f_0\hat{\mathcal{H}}$, so for large enough n the balls have disjoint projections, which means the balls are disjoint.

For the second statement, suppose N is large enough so that for all $n \geqslant N$ we have $d(o, h^{n/2}.o) \geqslant K' + K + 2C$, where K' is the constant of Lemma 3.3. Then

$$d^\pi_{f_0\mathcal{H}}(f_0 h^{\epsilon_0 n/2}.o, f_0.o) - d^\pi_{f_0\mathcal{H}}(f_0.o, f_1.o) \geqslant K',$$

so Lemma 3.3 implies $f_0.\star$ and $f_1.\star$ are in the same component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(f_0 h^{\epsilon_0 n/2}.\star)$. If, in addition, N is at least twice the constant of Corollary 3.4, then this is the $f_0 h^{-\epsilon_0\infty}$ component.

The proof of the third statement is similar. ◻

Proof of Proposition 3.1. Strongly constricting is the same as strongly contracting, by Proposition 2.9, so suppose h is a $(1, C)$–strongly constricting element. By Lemma 2.8, there exists a D such that $\pi_{\mathcal{H}}$ is D–coarsely equivalent to closest point projection. Recall that $D > C$. By Lemma 2.11, there exists a D' such that $\pi_{\mathcal{H}}$ is D'–coarsely 1-Lipschitz.

Let Γ be an infinite normal subgroup of G. Every infinite order element of $E(h)$ is strongly contracting, so if Γ contains such an element we are done. Otherwise, $\Gamma \cap E(h)$ is finite. Since Γ is infinite, there exists an element $g \in \Gamma$ such that $g \notin E(h)$. We claim that for sufficiently large n the element $f := gh^n g^{-1} h^{-n} \in \Gamma$ is strongly constricting.

For brevity, let $f^{i+1/2}$ denote $f^i gh^n$. Let $\hat{\mathcal{A}}_i := f^{i/2}\hat{\mathcal{H}}$ and $\mathcal{A}_i := f^{i/2}\mathcal{H}$. Define $\mathcal{B}_0 := \overline{\mathcal{B}}_\Delta(h^{n/2}.\star)$, $\mathcal{B}_1 := \overline{\mathcal{B}}_\Delta(f^{1/2} h^{-n/2}.\star)$, and $\mathcal{B}_{2k+i} := f^k \mathcal{B}_i$ for $k \in \mathbb{Z}$. Let $\hat{\mathcal{Z}}_i := f^{i/2} h^{(-1)^i n}.\star \in \mathcal{Y}$ and $\mathcal{Z}_i := f^{i/2} h^{(-1)^i n}.o \in \mathcal{X}$. Let $\hat{\mathcal{V}}_i := f^{i/2}.\star \in \mathcal{Y}$ and $\mathcal{V}_i := f^{i/2}.o \in \mathcal{X}$. See Figure 1.

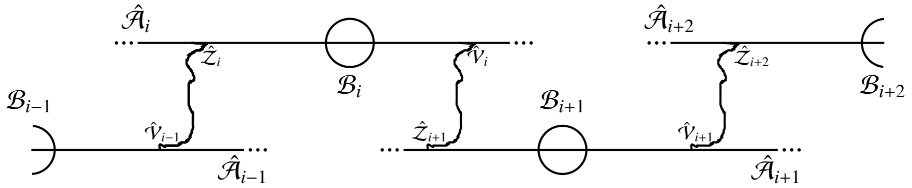

FIGURE 1. Disjoint balls in $\mathcal{Y}$.

By repeated applications of Lemma 3.6, for large enough n the balls $\mathcal{B}_i$ are pairwise disjoint. There are two orbits of these balls under the f–action, so f is an infinite order element. Furthermore, the balls are linearly ordered by separation, consistent with the subscripts, since for all i we have that $\mathcal{B}_j$ is contained in the $f^{i/2}h^{(-1)^{i+1}\infty}$ component of $\mathcal{Y} \setminus \mathcal{B}_i$ for all $j > i$, and in the $f^{i/2}h^{(-1)^i\infty}$ component for all $j < i$.

For any i and any $j < i-1$ the ball $\mathcal{B}_{i-1}$ separates $\hat{\mathcal{A}}_j$ from $\hat{\mathcal{A}}_i$ in $\mathcal{Y}$, so $\hat{\pi}_{\hat{\mathcal{A}}_i}(\hat{\mathcal{A}}_j)$ is contained in a bounded neighborhood of $\hat{\pi}_{\hat{\mathcal{A}}_i}(\hat{\mathcal{A}}_{i-1})$, which in turn we know is contained in a bounded neighborhood of $\hat{\mathcal{Z}}_i$. Conversely, $\hat{\pi}_{\hat{\mathcal{A}}_i}(\hat{\mathcal{A}}_j)$ is contained in a bounded neighborhood of $\hat{\mathcal{V}}_i$ for $j > i$. Since $\hat{\pi}_{\hat{\mathcal{A}}_i}$ agrees with $\pi_{\mathcal{A}_i}$ up to bounded error, the same statements are true for the axes in $\mathcal{X}$. That is, there exists a K, independent of n, such that for all i we have

- $d^{\pi}_{\mathcal{A}_i}(\mathcal{Z}_i, \mathcal{A}_j) \leqslant K$ if $j < i$, and
- $d^{\pi}_{\mathcal{A}_i}(\mathcal{V}_i, \mathcal{A}_j) \leqslant K$ if $j > i$.

Define $K' := 2K + C + 2D + D'$.

Suppose that for some $x \in \mathcal{X}$ there exists an i such that $d^{\pi}_{\mathcal{A}_i}(x, \mathcal{V}_i) > K'$. Then for any $j > i$ we have $d(\pi_{\mathcal{A}_i}(x), \pi_{\mathcal{A}_i}(\mathcal{A}_j)) > D > C$. Let y be a point of $\mathcal{A}_j$ closest to x. On any given geodesic from x to y there is a point $z \in \overline{\mathcal{B}}_{C+K}(\mathcal{V}_i)$, since $d^{\pi}_{\mathcal{A}_i}(x, y) > C$. Now $\pi_{\mathcal{A}_j}$ is D–coarsely equivalent to closest point projection, and y is closest to both x and z, so $d^{\pi}_{\mathcal{A}_j}(x, z) \leqslant 2D$. However, z is $(C + K)$–close to $\mathcal{V}_i$, and $d^{\pi}_{\mathcal{A}_j}(\mathcal{V}_i, \mathcal{Z}_j) \leqslant K$, so $d^{\pi}_{\mathcal{A}_j}(x, \mathcal{Z}_j) \leqslant 2D + C + K + D' + K = K'$.

We have shown that $d^{\pi}_{\mathcal{A}_i}(x, \mathcal{V}_i) > K'$ implies $d^{\pi}_{\mathcal{A}_j}(x, \mathcal{Z}_j) \leqslant K'$ for all $j > i$. A similar argument shows that $d^{\pi}_{\mathcal{A}_i}(x, \mathcal{Z}_i) > K'$ implies $d^{\pi}_{\mathcal{A}_j}(x, \mathcal{V}_j) \leqslant K'$ for all $j < i$.

Assume that n is large enough so that $d^{\pi}_{\mathcal{A}_0}(\mathcal{Z}_0, \mathcal{V}_0) = d^{\pi}_{\mathcal{H}}(h^n.o, o) > 2K' + 2C + 2D + d(o, g.o)$. Define $\mathcal{F} := \cup_{i\in\mathbb{Z}}\{\mathcal{V}_i\}$. We wish to define $\pi_{\mathcal{F}} \colon \mathcal{X} \to \mathcal{F}$ by sending a point x to the point $\mathcal{V}_\alpha$ where α is the greatest integer such that $d^{\pi}_{\mathcal{A}_\alpha}(x, \mathcal{V}_\alpha) \leqslant K'$, but we must verify that such an α exists. Fix an $x \in \mathcal{X}$, and suppose that $\iota \in \mathbb{Z}$ is such that $d(x, \mathcal{A}_\iota) = \min_{j\in\mathbb{Z}} d(x, \mathcal{A}_j)$. Such an ι exists since the action is properly discontinuous. Suppose that $d^{\pi}_{\mathcal{A}_\iota}(x, \mathcal{V}_\iota) \leqslant K'$. By the assumption on n, $d^{\pi}_{\mathcal{A}_\iota}(x, \mathcal{Z}_\iota) > K'$, so $d^{\pi}_{\mathcal{A}_j}(x, \mathcal{V}_\iota) \leqslant K'$ for all $j < \iota$. A brief computation shows that $d^{\pi}_{\mathcal{A}_{\iota+1}}(x, \mathcal{Z}_{\iota+1}) \leqslant d(x, \mathcal{A}_{\iota+1}) + d(o, g.o) + K' + 2C + D$. By Lemma 2.8, $d(\mathcal{Z}_{\iota+1}, \pi_{\mathcal{A}_{\iota+1}}(x)) \leqslant d(o, g.o) + K' + 2C + 2D$, which, again by our assumption on n, implies $d^{\pi}_{\mathcal{A}_{\iota+1}}(x, \mathcal{V}_{\iota+1}) > K'$. We conclude that $\alpha \leqslant \iota$. The previous paragraph then tells us that $d^{\pi}_{\mathcal{A}_j}(x, \mathcal{Z}_j) \leqslant K'$ for all $j > \alpha + 1$.

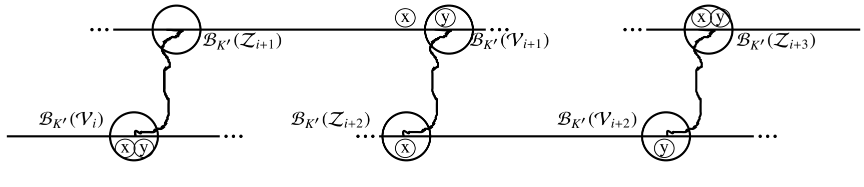

FIGURE 2. Projections ⓧ of x and ⓨ of y to each axis.

Now suppose x and y are points with $\pi_{\mathcal{F}}(x) = \mathcal{V}_i$ and $\pi_{\mathcal{F}}(y) = \mathcal{V}_j$ for $j > i + 1$. Then for each $i + 2 \leqslant k \leqslant j$ we have $d^{\pi}_{\mathcal{A}_k}(x, y) \geqslant d^{\pi}_{\mathcal{A}_k}(\mathcal{Z}_k, \mathcal{V}_k) - 2K' > C$. Figure 2 depicts a situation with $j = i + 2$ that shows $j > i + 1$ is necessary, since the projections to $\mathcal{A}_{i+1}$ may be close. By the strong constriction property for each $\mathcal{A}_k$, every geodesic from x to y passes $(C + K')$–close to $\mathcal{Z}_k$ and $\mathcal{V}_k$. So every geodesic passes within $C + K'$ of $\pi_{\mathcal{F}}(y) = \mathcal{V}_j$ and within $C + K'$ of $\mathcal{Z}_{i+2}$, which boundedly close to $\pi_{\mathcal{F}}(x) = \mathcal{V}_i$.

Therefore, $\pi_{\mathcal{F}}$ is $(1, \max\{d(\mathcal{V}_0, \mathcal{V}_2), C+K'+d(\mathcal{V}_0, \mathcal{Z}_2)\})$–strongly constricting. Lemma 3.2 says $i \mapsto f^i.o$ is a quasi-geodesic, so $f \in \Gamma$ is a strongly contracting element. □

4. A Minimal Section

Let $\mathcal{X}$ be a G–space with basepoint o. Suppose that there exists a strongly contracting element for $G \curvearrowright \mathcal{X}$. Let Γ be an infinite normal subgroup of G. By Proposition 3.1, there exists a strongly contracting element $h \in \Gamma$. Let $\mathcal{H} = E(h).o$ be an axis for h, and define equivariant projections to translates of $\mathcal{H}$ as in Definition 2.21. Suppose $\pi_{\mathcal{H}}$ is a $(1, C)$–strongly constricting C–coarse map.

Definition 4.1. For each element $g\Gamma \in G/\Gamma$ choose an element $\bar{g} \in g\Gamma$ such that $d(o, \bar{g}.o) = d(o, g\Gamma.o) = d(\Gamma.o, g\Gamma.o)$. Let $\bar{G} := \{\bar{g} \mid g\Gamma \in G/\Gamma\}$. We call $\bar{G}$ a *minimal section*, and let $\bar{\mathcal{G}}$ denote $\bar{G}.o$.

Observe that $\Theta'_{G/\Gamma}(s) = \Theta'_{\bar{\mathcal{G}}}(s)$, so $\delta_{G/\Gamma} = \delta_{\bar{G}}$.

The next lemma says, coarsely, that the minimal section is orthogonal to translates of $\mathcal{H}$.

Lemma 4.2. *For every $\bar{g} \in \bar{G}$ and for every $f \in G$ we have $d^{\pi}_{f\mathcal{H}}(o, \bar{g}.o) \leqslant 8C + D$, where* $D := \operatorname{diam}\langle h\rangle\backslash\mathcal{H}$.

Proof. Suppose not. Then there exists an $n \neq 0$ such that:

$$\begin{aligned} D &\geqslant d(\pi_{f\mathcal{H}}(o), fh^n f^{-1}.\pi_{f\mathcal{H}}(\bar{g}.o)) \\ &\geqslant d^{\pi}_{f\mathcal{H}}(o, fh^n f^{-1}\bar{g}.o) - 2C \end{aligned}$$

Thus, $d^{\pi}_{f\mathcal{H}}(o, \bar{g}.o) - d^{\pi}_{f\mathcal{H}}(o, fh^n f^{-1}\bar{g}.o) > 6C$. However:

$$\begin{aligned} &d(o, fh^n f^{-1}\bar{g}.o) \\ &\quad\leqslant d(o, \pi_{f\mathcal{H}}(o)) + d^{\pi}_{f\mathcal{H}}(o, fh^n f^{-1}\bar{g}.o) + d(\pi_{f\mathcal{H}}(fh^n f^{-1}\bar{g}.o), fh^n f^{-1}\bar{g}.o) \\ &\quad< d(o, \pi_{f\mathcal{H}}(o)) + d^{\pi}_{f\mathcal{H}}(o, \bar{g}.o) + d(\pi_{f\mathcal{H}}(fh^n f^{-1}\bar{g}.o), fh^n f^{-1}\bar{g}.o) - 6C \\ &\quad= d(o, \pi_{f\mathcal{H}}(o)) + d^{\pi}_{f\mathcal{H}}(o, \bar{g}.o) + d(\pi_{f\mathcal{H}}(\bar{g}.o), \bar{g}.o) - 6C \\ &\quad\leqslant d(o, \bar{g}.o) \quad \text{(by Lemma 2.10)} \end{aligned}$$

This contradicts minimality of $\bar{G}$, since $fh^n f^{-1}\bar{g} = \bar{g}\bar{g}^{-1}fh^n f^{-1}\bar{g} \in \bar{g}\Gamma$. □

Corollary 4.3. *If $d(\bar{g}.o, \bar{g}'.o) \geqslant 18C + 2D$ for $\bar{g}$, $\bar{g}' \in \bar{G}$ then there is no $f \in G$ such that $\bar{g}.o \in f\mathcal{H}$ and $\bar{g}'.o \in f\mathcal{H}$.*

Proof. If there were such an f, we would have $d^{\pi}_{f\mathcal{H}}(\bar{g}.o, \bar{g}'.o) \geqslant 2(8C + D)$, which means either $\bar{g}$ or $\bar{g}'$ would contradict Lemma 4.2. □

In light of Corollary 4.3, it will be convenient to pass to a coarsely dense subset of $\bar{\mathcal{G}}$ whose elements yield distinct translates of $\mathcal{H}$:

Definition 4.4. Let $K \geqslant 18C + 2D$, and let A be a maximal subset of $\bar{G}$ such that $1 \in A$ and $d(\bar{g}.o, \bar{g}'.o) \geqslant K$ for all distinct $\bar{g}$, $\bar{g}' \in A$. Let $\mathcal{A} := A.o$.

By maximality, for every $\bar{g} \in \bar{G}$ there is some $a \in A$ such that $d(a.o, \bar{g}.o) \leqslant K$. There are boundedly many points of $\bar{\mathcal{G}}$ in a ball of radius K, so $\Theta_{\bar{\mathcal{G}}}(s)$ is bounded below by $\Theta_{\mathcal{A}}(s)$ and above by a constant multiple of $\Theta_{\mathcal{A}}(s)$. In particular, $\Theta_{\mathcal{A}}(s)$ has the same convergence behavior as $\Theta_{\bar{\mathcal{G}}}(s)$, so $\delta_A = \delta_{\bar{G}} = \delta_{G/\Gamma}$.

Corollary 4.3 implies $a\mathcal{H} \neq a'\mathcal{H}$ for distinct a, $a' \in A$.

5. Embedding a Free Product Set

Let $A \subset \bar{G}$ as in Definition 4.4, and let $A^* := A \setminus \{1\}$. Consider the free product set $A^* * \mathbb{Z}_2 := \bigcup_{k=1}^{\infty}\{(a_1, \ldots, a_k) \mid a_i \in A^*\}$. For any $n > 0$ we can map the free product set into G by $\phi_n\colon (a_1, \ldots, a_k) \mapsto a_1h^na_2h^n \cdots a_kh^n$. Our goal is to show $\delta_{\phi_n(A^* * \mathbb{Z}_2)} > \delta_A$. We establish the inequality in the next section. In this section we show ϕ_n is an injection for all sufficiently large n. In fact, we prove something stronger:

Proposition 5.1. *The map $A^* * \mathbb{Z}_2 \to G : (a_1, \ldots, a_k) \mapsto a_1h^n \cdots a_nh^n.o$ is an injection for all sufficiently large n.*

The map is an injection because we have an action of G on the quasi-tree $\mathcal{Y}$, and for large enough n we have "quasi-edges" of the form $[y, yh^n]$. We have set things up so that the a's do not backtrack across such edges. See Figure 3. We make this precise:

Proof. Let $\underline{a} = (a_1, \ldots, a_k) \in A^* * \mathbb{Z}_2$.

By Lemma 4.2, there is a K such that $d^{\pi}_{f\mathcal{H}}(o, \bar{g}.o) \leqslant K$ for every $f \in G$ and every $\bar{g} \in \bar{G}$. The choice of $A \subset \bar{G}$ in Definition 4.4 guarantees that the axes $a\mathcal{H}$ for $a \in A$ are distinct. Let N be the constant of Lemma 3.6 for this K, and choose $n \geqslant N$.

Note that the proof of Lemma 3.6 includes the fact that $d(o, h^{n/2}.o) \geqslant K' + K + 2C$, where K' is the constant of Lemma 3.3. Therefore, if $\phi_n(\underline{a}).o = \phi_n(\underline{a}').o$ then

$$d^{\pi}_{\phi_n(\underline{a})\mathcal{H}}(\phi_n(\underline{a}).o, \phi_n(\underline{a})h^{-n/2}.o) - d^{\pi}_{\phi_n(\underline{a})\mathcal{H}}(\phi_n(\underline{a}).o, \phi_n(\underline{a}').o) \geq K' + C > K',$$

so Lemma 3.3 implies $\phi_n(\underline{a}).\star$ and $\phi_n(\underline{a}').\star$, though they might not be equal, are at least contained in the same component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(\phi_n(\underline{a})h^{-n/2}.\star)$.

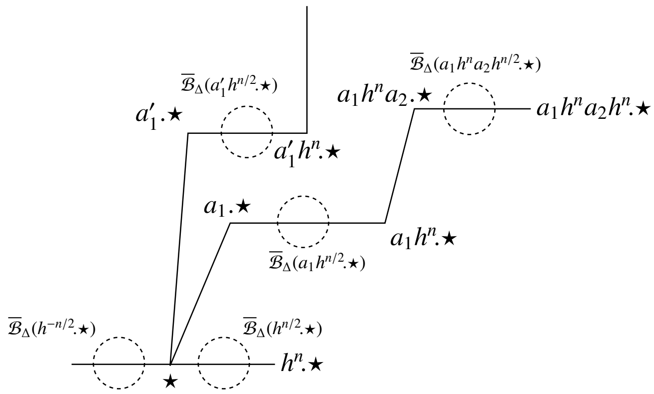

Figure 3. A does not cross h^n quasi-edges

Define $\mathcal{V}_i(\underline{a})$ to be the $a_1h^n \cdots a_ih^\infty$ component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(a_1h^n \cdots a_ih^{n/2}.\star)$ for $i \leqslant k$ (recall Definition 3.5). Lemma 3.6 implies that $\mathcal{V}_i(\underline{a}) \supset \mathcal{V}_{i+1}(\underline{a})$ and $\phi_n(\underline{a}).\star \in \mathcal{V}_k(\underline{a})$. Moreover, for $i \leq \min\{k, k'\}$, $\mathcal{V}_i(\underline{a})$ and $\mathcal{V}_i(\underline{a}')$ are disjoint unless $a_j = a_j'$ for all $j \leqslant i$.

If $\phi_n(\underline{a}).o = o$ then Lemma 3.3 implies $\star \in \mathcal{V}_k(\underline{a}) \subset \mathcal{V}_1(\underline{a})$. This contradicts the fact that $\star$ is contained in the $a_1h^{-\infty}$ component of $\mathcal{Y} \setminus \overline{\mathcal{B}}_\Delta(a_1h^{n/2}.\star)$. The same argument shows that if $\underline{a}$ is a proper prefix of $\underline{a}'$, that is, if $\underline{a} = (a_1, \ldots, a_k)$ and $\underline{a}' = (a_1, \ldots, a_k, a_{k+1}', \ldots, a_{k'}')$ with $k' > k$, then $\phi_n(\underline{a}).o \neq \phi_n(\underline{a}').o$.

Suppose $\phi_n(\underline{a}).o = \phi_n(\underline{a}').o$ with $k \leqslant k'$. Lemma 3.3 implies $\phi_n(\underline{a}).\star \in \mathcal{V}_{k'}(\underline{a}')$, so $a_i = a_i'$ for all $i \leqslant k$. Since $\underline{a}$ cannot be a proper prefix of $\underline{a}'$, $k = k'$. Hence, $\phi_n(\underline{a}).o = \phi_n(\underline{a}').o$ implies $\underline{a} = \underline{a}'$ for all sufficiently large n. □

6. Growth Gap

A free product of groups has greater growth exponent than the factor groups, with respect to a word metric, so we expect that $\phi_n(A^* * \mathbb{Z}_2)$ should have a larger growth exponent than A. To verify this intuition, one must show that the Poincaré series for $\phi_n(A^* * \mathbb{Z}_2)$ diverges at $\delta_A + \epsilon$ for some $\epsilon > 0$. A clever manipulation of Poincaré series yields the following criterion:

LEMMA 6.1 ([26, Criterion 2.4],[50, Proposition 2.3]). *If the map*

$$\phi_n\colon\ A^* * \mathbb{Z}_2 \to G : (a_1, \dots, a_k) \mapsto a_1 h^n \cdots a_k h^n$$

is an injection, and if $\exp(|h^n| \cdot \delta_A) < \Theta_A(\delta_A)$*, then* $\delta_{\phi_n(A^* * \mathbb{Z}_2)} > \delta_A$.

Because our methods are coarse we have passed to a high power h^n of h and therefore do not have control over $|h^n|$. However, the criterion is satisfied automatically if A, or, equivalently, $\bar{G}$, is divergent, which, recalling Definition 1.7, means Θ_A diverges at δ_A. The following definition will be used in a condition to guarantee divergence of $\bar{G}$.

DEFINITION 6.2. Let $Comp^G_{Q,r} \subset G.o$ be the set of points $g.o$ such that there exists a geodesic $[x, y]$ of length r with $x \in \overline{\mathcal{B}}_Q(o)$ and $y \in \overline{\mathcal{B}}_Q(g.o)$ whose interior is contained in $\mathcal{X} \setminus \overline{\mathcal{B}}_Q(G.o)$.

Define the *Q–complementary growth exponent* of G to be:

$$\delta^c_G := \limsup_{r\to\infty} \frac{\log \#Comp^G_{Q,r}}{r}$$

THEOREM 6.3. *Let G be a finitely generated, non-elementary group. Let $\mathcal{X}$ be a G–space. If G contains a strongly contracting element and there exists a $Q \geqslant 0$ such that the Q–complementary growth exponent of G is strictly less than the growth exponent of G, then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

The proof of Theorem 6.3 follows in part the proof of [26, Theorem 1.4] for geometrically finite Kleinian groups. For the divergence part of the proof, the Kleinian group ingredients of [26, Theorem 1.4] are inessential, and our changes are mostly cosmetic. The real generalization is in the use of Proposition 5.1 instead of a Ping-Pong argument.

PROOF. Let Γ be an infinite, infinite index normal subgroup of G. By Proposition 3.1, there is a strongly contracting element in Γ. Let $\bar{G}$ be a minimal section of G/Γ. If $\delta_{\bar{G}} \leqslant \delta^c_G$ then we are done, since $\delta^c_G < \delta_G$, so suppose $\delta_{\bar{G}} > \delta^c_G$.

CLAIM 6.3.1. $\bar{G}$ is divergent.

Assume the claim, and let A be a maximal separated set in $\bar{G}$ as in Definition 4.4. Then A and $\bar{G}$ have the same critical exponent, and are both divergent. By Proposition 5.1, for sufficiently large n the map $\phi_n\colon\ A^* * \mathbb{Z}_2 \to G$ is an injection. By Lemma 6.1, $\delta_A < \delta_{\phi_n(A^* * \mathbb{Z})}$. Thus, $\delta_{G/\Gamma} = \delta_{\mathcal{A}} < \delta_{\phi_n(A^* * \mathbb{Z}_2)} \leqslant \delta_G$.

It remains to prove the claim.

Let $r > 0$, and suppose $d(o, \bar{g}.o) = r$. Let $0 \leqslant M_0 \leqslant r$ and $M_1 = r - M_0$. Choose a geodesic $[o, \bar{g}.o]$ from o to $\bar{g}.o$, and let $[o, \bar{g}.o](M_0)$ denote the point of $[o, \bar{g}.o]$ at distance M_0 from o.

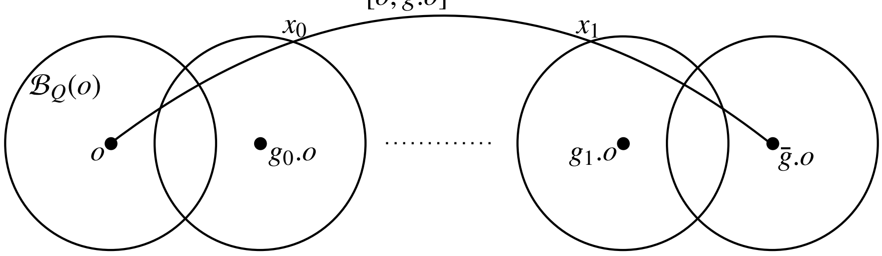

FIGURE 4. Splitting a geodesic into three subsegments

First, we suppose that $[o, \bar{g}.o](M_0) \in \mathcal{X} \setminus \overline{\mathcal{B}}_Q(G.o)$. Let $[x_0, x_1] \subset [o, \bar{g}.o]$ be the largest subsegment containing $[o, \bar{g}.o](M_0)$ such that $(x_0, x_1) \subset \mathcal{X} \setminus \overline{\mathcal{B}}_Q(G.o)$. Let $m_0 = d(o, x_0)$, and let $m_1 = d(x_1, \bar{g}.o)$. There exist group elements $g_i \in G$ such that $d(g_i.o, x_i) \leqslant Q$. See Figure 4. We have $\bar{g}.o = \overline{g_0} \cdot g_0^{-1} g_1 \cdot \overline{g_1^{-1} \bar{g}}.o$. Now $m_0 - Q \leqslant d(o, \bar{g}_0.o) \leqslant d(o, g_0.o) \leqslant m_0 + Q$, and $m_1 - Q \leqslant d(o, \overline{g_1^{-1}\bar{g}}.o) \leqslant d(o, g_1^{-1}\bar{g}.o) \leqslant m_1 + Q$. Furthermore, $g_0^{-1} g_1 \in Comp^G_{Q,r-(m_0+m_1)}$. Thus, the point $\bar{g}.o$ can be expressed as a product of an element of $\bar{G}$ of length $m_0 \pm Q$, an element of $\bar{G}$ of length $m_1 \pm Q$, and the quotient of an element of $Comp^G_{Q,r-(m_0+m_1)}$.

(†) The same is also true if $[o,\bar{g}.o](M_0)\in\overline{\mathcal{B}}_Q(G.o)$, in which case we can take $m_0=M_0$ and $m_1=r-m_0$. Then choose $g_0=g_1$ so that the contribution from $Comp^G_{Q,r-(m_0+m_1)}$ is trivial.

Let $V_{r,Q}:=\#\left(\bar{G}.o\cap\overline{\mathcal{B}}_{r+Q}(o)\setminus\mathcal{B}_{r-Q}(o)\right)$. For every $M_0+M_1=r$ we have:

$$V_{r,Q}\overset{*}{\prec}\sum_{m_0=0}^{M_0}\sum_{m_1=0}^{M_1}V_{m_0,Q}\cdot V_{m_1,Q}\cdot\#Comp^G_{Q,r-(m_0+m_1)}$$

Choose $\xi>0$ such that $\delta_{\bar{G}}\geqslant 2\xi+\delta^c_G$. Since

$$\#Comp^G_{Q,r-(m_0+m_1)}\overset{*}{\prec}\exp((r-(m_0+m_1))(\delta_{\bar{G}}-\xi))$$

whenever $r-(m_0+m_1)$ is sufficiently large, it follows that:

(1) $$V_{r,Q}\cdot\exp(-r(\delta_{\bar{G}}-\xi))\overset{*}{\prec}\left(\sum_{m_0=0}^{M_0}V_{m_0,Q}\cdot\exp(-m_0(\delta_{\bar{G}}-\xi))\right)\cdot\left(\sum_{m_1=0}^{M_1}V_{m_1,Q}\cdot\exp(-m_1(\delta_{\bar{G}}-\xi))\right)$$

Set $w_i:=V_{i,Q}\cdot\exp(-i(\delta_{\bar{G}}-\xi))$ and $W_i:=\sum_{j=1}^{i}w_i$. Then (1) and [26, Lemma 4.3] imply that $\sum_i w_i\cdot\exp(-is)$ diverges at its critical exponent, which is:

$$\limsup_i\frac{\log w_i}{i}=\left(\limsup_i\frac{\log V_{i,Q}}{i}\right)-(\delta_{\bar{G}}-\xi)=\xi$$

So $\infty=\sum_i w_i\cdot\exp(i\xi)=\sum_i V_{i,Q}\cdot\exp(-i\delta_{\bar{G}})\overset{*}{\asymp}\Theta_{\bar{G}}(\delta_{\bar{G}})$. □

THEOREM 6.4. *Let G be a finitely generated, non-elementary group. Let X be a quasi-convex G–space. If G contains a strongly contracting element then $G\curvearrowright X$ is a growth tight action.*

PROOF. The proof is an easier special case of the proof of Theorem 6.3. If X is Q–quasi-convex then we can always choose to be in case (†) of the proof. □

7. Growth of Conjugacy Classes

Parkkonen and Paulin [46] ask: given a finitely generated group G with a word metric and an element $h\in G$, what is growth rate of the conjugacy class $[h]$ of h? In a hyperbolic group G there is a finite subgroup, the *virtual center*, consisting of elements whose centralizer is finite index in G. The growth exponent of a conjugacy class in the virtual center is clearly zero. Parkkonen and Paulin show that for every element h not in the virtual center, $\delta_{[h]}=\frac{1}{2}\delta_G$. This generalized an old result of Huber [38] for the case of G acting cocompactly on the hyperbolic plane and h loxodromic.

Since strongly contracting elements behave much like infinite order elements in hyperbolic groups, it is natural to ask whether the growth exponent of the conjugacy class of a strongly contracting element h also satisfies $\delta_{[h]}=\frac{1}{2}\delta_G$.

We show that the lower bound holds, and the upper bound holds if h moves the basepoint sufficiently far with respect to the contraction constant for the axis.

THEOREM 7.1. *Let G be a non-elementary, finitely generated group, and let X be a G–space. Let h be a strongly contracting element for $G\curvearrowright X$. Then $\delta_{[h]}\geqslant\frac{1}{2}\delta_G$.*

Let $D:=\operatorname{diam}Z(h)\backslash\mathcal{H}$, where $Z(h)$ is the centralizer of h in G. Suppose $\pi_{\mathcal{H}}$ is a $(1,C)$–strongly constricting, C–coarse map. If $d(o,h.o)>15C+2D$ then $\delta_{[h]}=\frac{1}{2}\delta_G$.

COROLLARY 7.2. *For h strongly contracting, $\delta_{[h^n]}=\frac{1}{2}\delta_G$ for all sufficiently large n.*

PROOF. For n nonzero, $E(h^n)=E(h)$ and $Z(h^n)\supset Z(h)$, so the same C and D work for h^n as work for h. On the other hand, $\langle h\rangle$ is quasi-isometrically embedded, so $d(o,h^n.o)\asymp n$. Thus, $d(o,h^n.o)>15C+2D$ for large enough n. □

It would be interesting to know whether the restriction on $d(o,h.o)$ is really necessary:

Question 7.3. Does there exist an action $G \curvearrowright \mathcal{X}$ such that h is a strongly contracting element with $\delta_{[h]} > \frac{1}{2}\delta_G$?

Proof of Theorem 7.1. Define $K := 6C + D$ and $F := \{g \in G \mid d^{\pi}_{g\mathcal{H}}(o, g.o) \leqslant K\}$.

First, we will show $\delta_F = \delta_G$. Then, we will relate $\delta_{[h]}$ to δ_F.

For any $r \geqslant 0$ consider $\phi\colon \{f \in F \setminus E(h) \mid d(o, f.o) \leqslant r\} \to \{g\mathcal{H} \mid g \in G \setminus E(h) \text{ and } g\mathcal{H} \cap \overline{\mathcal{B}}_r(o) \neq \emptyset\}$ defined by $\phi(f) := f\mathcal{H}$. For each axis $g\mathcal{H}$ meeting $\overline{\mathcal{B}}_r(o)$ there exists a $g' \in gE(h)$ such that $d(o, g'.o) = d(o, g\mathcal{H}) \leqslant r$. Since $\pi_{g\mathcal{H}}$ is within $5C$ of closest point projection, by Lemma 2.8, we have $d^{\pi}_{g'\mathcal{H}}(o, g'.o) \leqslant 6C \leqslant K$. Therefore, $g' \in F$ with $\phi(g') = g\mathcal{H}$, so ϕ is surjective.

We estimate:

$$\#\text{axes meeting } \overline{\mathcal{B}}_r(o) \geqslant \frac{|G.o \cap \overline{\mathcal{B}}_r(o)| \times \#\text{axes per orbit point}}{\text{maximum number of orbit points per axis}}$$

The basepoint belongs to $[\mathrm{Stab}_G(o) : E(h) \cap \mathrm{Stab}_G(o)]$ distinct translates of $\mathcal{H}$, so the number of axes per orbit point is constant. The maximum number of orbit points in $\overline{\mathcal{B}}_r(o)$ contained in a single axis is proportional to r, since each axis is a quasi-isometrically embedded image of a virtually cyclic group. Combined with surjectivity of ϕ this gives:

$$|F.o \cap \overline{\mathcal{B}}_r(o)| \overset{*}{\succ} \frac{|G.o \cap \overline{\mathcal{B}}_r(o)|}{r}$$

Thus:

$$\begin{aligned}
\delta_F &= \limsup_{r\to\infty} \frac{1}{r} \log |F.o \cap \overline{\mathcal{B}}_r(o)| \\
&\geqslant \limsup_{r\to\infty} \frac{1}{r} \log \frac{|G.o \cap \overline{\mathcal{B}}_r(o)|}{r} \\
&= \limsup_{r\to\infty} \frac{1}{r} \log |G.o \cap \overline{\mathcal{B}}_r(o)| = \delta_G
\end{aligned}$$

The reverse inequality is trivial, since $F \subset G$, so $\delta_F = \delta_G$.

Now consider the map $\psi\colon F \setminus E(h) \to [h] \setminus E(h)$ defined by $\psi(f) := fhf^{-1}$. Choose minimal length representatives $e_1, \ldots, e_m$ of $Z(h)\backslash E(h)$. For each $g \in G \setminus E(h)$ there exists a $g' \in gE(h)$ such that $d(o, g\mathcal{H}) = d(o, g'.o)$. There exist $z \in Z(h)$ and i such that $g' = gze_i$. Let $f := g'e_i^{-1}$, so that $fhf^{-1} = gze_ie_i{}^{-1}he_ie_i{}^{-1}z^{-1}g^{-1} = ghg^{-1}$. Since e_i has length at most D and $\pi_{g\mathcal{H}}$ is $5C$–close to closest point projection, it follows that $f \in F$, so ψ is surjective. Furthermore, $d(o, fhf^{-1}.o) \leqslant 2d(o, f.o) + d(o, h.o)$, by the triangle inequality.

On the other hand, ψ is boundedly many-to-one, since if $fhf^{-1} = f'hf'^{-1}$ then $f' \in fE(h)$, so $f\mathcal{H} = f'\mathcal{H}$. By definition of F, we then have $d^{\pi}_{f\mathcal{H}}(o, f.o) \leqslant K$ and $d^{\pi}_{f\mathcal{H}}(o, f'.o) \leqslant K$, so $d(f.o, f'.o) \leqslant 2(C + K)$. There are uniformly boundedly many such f' for each f.

Hence, ψ is a surjective, boundedly-many-to-one map such that $d(o, \psi(f).o) \overset{+}{\prec} 2d(o, f.o)$ for all f. We excluded $E(h)$ from the domain and range, but its growth exponent is zero, since it embeds quasi-isometrically into $\mathcal{X}$, so $\delta_{[h]} = \delta_{[h]\setminus E(h)} \geqslant \frac{1}{2}\delta_{F\setminus E(h)} = \frac{1}{2}\delta_F = \frac{1}{2}\delta_G$.

Now, $d^{\pi}_{f\mathcal{H}}(fh.o, fhf^{-1}.o) = d^{\pi}_{f\mathcal{H}}(o, f.o) \leqslant K$ for $f \in F$, so $d^{\pi}_{f\mathcal{H}}(o, fhf^{-1}.o) > d(f.o, fh.o) - 2(C+K)$. If $d(o, h.o) > 15C + 2D = C + 2(C+K)$ then we have $d^{\pi}_{f\mathcal{H}}(o, fhf^{-1}.o) > C$, so by strong constriction, $d(o, fhf^{-1}.o) \geqslant 2d(o, f.o) + d(o, h.o) - 4(C + K)$. Thus, $d(o, \psi(f).o) \overset{+}{\asymp} 2d(o, f.o)$ and $\delta_{[h]} = \frac{1}{2}\delta_G$. □

8. Actions on Relatively Hyperbolic Spaces

Yang [61] proved that the action of a finitely generated group G with a non-trivial Floyd boundary on any of its Cayley graphs is growth tight. Relatively hyperbolic groups have non-trivial Floyd boundaries by a theorem of Gerasimov [31], so the action of a relatively hyperbolic group on any of its Cayley graphs is growth tight. It is an open question whether there exists a group with a non-trivial Floyd boundary that is not relatively hyperbolic.

There is also a notion of relative hyperbolicity of metric spaces, which we will review in Section 8.1. One motivating example of a relatively hyperbolic metric space is a Cayley graph of a relatively hyperbolic group. Another is the universal cover $\tilde{M}$ of a complete, finite volume hyperbolic manifold M. The fundamental group $\pi_1(M)$ of such a manifold is a relatively hyperbolic group, so the action of $\pi_1(M)$ on any of its Cayley graphs is growth tight by Yang's theorem. However, this does not tell us whether the action of $\pi_1(M)$ on $\tilde{M}$ is growth tight. This question was addressed for a more general class of manifolds by Dal'bo, Peigné, Picaud, and Sambusetti [26], who proved growth tightness results for geometrically finite Kleinian groups. Using our main theorems, Theorem 6.3 and Theorem 6.4, we generalize their results to all groups acting on relatively hyperbolic metric spaces.

8.1. Relatively Hyperbolic Metric Spaces.

Definition 8.1 (cf. [27, 57]). Let $\mathcal{X}$ be a geodesic metric space and let $\underline{\mathcal{P}}$ be a collection of uniformly coarsely connected subsets of $\mathcal{X}$. We say $\mathcal{X}$ is *hyperbolic relative to the peripheral sets* $\underline{\mathcal{P}}$ if the following conditions are satisfied:

1. For each A there exists a B such that $\operatorname{diam}(\overline{\mathcal{B}}_A(\mathcal{P}_0)\cap\overline{\mathcal{B}}_A(\mathcal{P}_1)) \leqslant B$ for distinct $\mathcal{P}_0,\ \mathcal{P}_1 \in \underline{\mathcal{P}}$.
2. There exists an $\epsilon \in (0, \frac{1}{2})$ and $M \geqslant 0$ such that if $x_0,\ x_1 \in \mathcal{X}$ are points such that for some $\mathcal{P} \in \underline{\mathcal{P}}$ we have $d(x_i, \mathcal{P}) \leqslant \epsilon \cdot d(x_0, x_1)$ for each i, then every geodesic from x_0 to x_1 intersects $\overline{\mathcal{B}}_M(\mathcal{P})$.
3. There exist σ and δ so that for every geodesic triangle either:
   (a) there exists a ball of radius σ intersecting all three sides, or
   (b) there exists a $\mathcal{P} \in \underline{\mathcal{P}}$ such that $\overline{\mathcal{B}}_\sigma(\mathcal{P})$ intersects all three sides and for each corner of the triangle, the points of the outgoing geodesics from that corner which first enter $\overline{\mathcal{B}}_\sigma(\mathcal{P})$ are distance at most δ apart.

We say $\mathcal{X}$ is *hyperbolic* if it hyperbolic relative to $\underline{\mathcal{P}} = \emptyset$.

If $\mathcal{X}$ is hyperbolic in the sense of Definition 8.1 then the only non-trivial condition is 3a, which is equivalent to the usual definition of hyperbolic metric space.

Definition 8.2. A group G is *hyperbolic relative to a collection of finitely generated peripheral subgroups* if a Cayley graph of G is hyperbolic relative to the cosets of the peripheral subgroups.

Sisto [57] shows Definition 8.2 is equivalent to Bowditch's [13] definition of relatively hyperbolic groups.

Definition 8.3 (cf. [35]). Let $\mathcal{X}$ be a connected graph with edges of length bounded below. A *combinatorial horoball* based on X with parameter $a > 0$ is a graph whose vertex set is $\mathrm{Vert}\mathcal{X} \times (\{0\} \cup \mathbb{N})$, contains an edge of length 1 between (v, n) and $(v, n + 1)$ for all $v \in \mathrm{Vert}\mathcal{X}$ and all $n \in \{0\} \cup \mathbb{N}$, and for each edge $[v, w] \in \mathcal{X}$ contains an edge $[(v, n), (w, n)]$ of length $e^{-an} \cdot \mathrm{length}([v, w])$.

Let $\mathcal{X}$ be hyperbolic relative to $\underline{\mathcal{P}}$. An *augmented space* is a space obtained from $\mathcal{X}$ as follows. By definition, there exists a constant C such that each $\mathcal{P} \in \underline{\mathcal{P}}$ is C–coarsely connected. For each $\mathcal{P} \in \underline{\mathcal{P}}$ choose a maximal subset of points that pairwise have distance at least C from one another. Let these points be the vertex set of a graph. For edges, choose a geodesic connecting each pair of vertices at distance at most $2C$ from each other. Use this graph as the base of a combinatorial horoball with parameter $a_{\mathcal{P}} > 0$. The augmented space is the space obtained from the union of $\mathcal{X}$ with horoballs $\mathcal{X}_{\mathcal{P}}$ for each $\mathcal{P} \in \underline{\mathcal{P}}$ by identifying the base of $\mathcal{X}_{\mathcal{P}}$ with the graph constructed in $\mathcal{P}$.

Definition 8.4. Let $\mathcal{X}$ be a hyperbolic G–space, and let $\underline{P}$ be the collection of maximal parabolic subgroups of G. Suppose there exists a G–invariant collection of disjoint open horoballs centered at the points fixed by the parabolic subgroups. The *truncated space* is $\mathcal{X}$ minus the union of these open horoballs. We say $G \curvearrowright \mathcal{X}$ is *cusp uniform* if G acts cocompactly on the truncated space.

If G acts cocompactly on a G–space $\mathcal{X}'$ that is hyperbolic relative to a G–invariant peripheral system $\underline{\mathcal{P}}$, then an augmented space $\mathcal{X}$ can be constructed G–equivariantly, and $G \curvearrowright \mathcal{X}$ will be a cusp uniform action.

Several different versions of the following theorem occur in the literature on relatively hyperbolic groups:

Theorem 8.5 ([13, 34, 57]). *If $\mathcal{X}$ is hyperbolic relative to $\underline{\mathcal{P}}$ then any augmented space with horoball parameters bounded below is hyperbolic.*

If $G \curvearrowright \mathcal{X}$ is a cusp uniform action then G is hyperbolic relative to the maximal parabolic subgroups and the truncated space is hyperbolic relative to boundaries of the deleted horoballs.

8.2. Quasi-convex Actions.

Theorem 8.6. *If $\mathcal{X}$ is a quasi-convex, relatively hyperbolic G–space and G does not coarsely fix a peripheral subspace then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

Proof. It follows from [57, Lemma 5.4] that every infinite order element of G that does not coarsely fix a peripheral subspace is strongly constricting. We conclude by Theorem 6.4. □

Theorem 8.6 unifies the existing proofs of growth tightness for cocompact actions on hyperbolic spaces [48] and for actions of a relatively hyperbolic group on its Cayley graphs [61], and extends to actions on a more general class of spaces.

Corollary 8.7. *The action of a finitely generated group G with infinitely many ends on any one of its Cayley graphs is growth tight.*

Proof. Stallings' Theorem [59] says that G splits non-trivially over a finite subgroup. G is hyperbolic relative to the factor groups of this splitting. Since the splitting is non-trivial, G does not fix any factor group, so the result follows from Theorem 8.6. □

Corollary 8.7 generalizes a result of Sambusetti [50, Theorem 1.4], who proved it with additional constraints on the factor groups.

8.3. Cusp Uniform Actions. Theorem 8.6 and Theorem 8.5 show that if $G \curvearrowright \mathcal{X}$ is a cusp uniform action on a hyperbolic space then the action of G on the truncated space is a growth tight action. A natural question is whether $G \curvearrowright \mathcal{X}$ is a growth tight action. This action is not quasi-convex if the parabolic subgroups are infinite, as geodesics in $\mathcal{X}$ will travel deeply into horoballs, and, indeed, an example of Dal'bo, Otal, and Peigné [25] shows $G \curvearrowright \mathcal{X}$ need not be growth tight.

To see how growth tightness can fail, consider the combinatorial horoball from Definition 8.3. If $\mathcal{X}$ is, say, the Cayley graph of some group and we build the combinatorial horoball with parameter $a > 0$ based on $\mathcal{X}$, then the r–ball about a basepoint $o \in \mathcal{X}$ in the horoball metric intersected with $\mathcal{X} \times \{0\}$ contains the ball of radius $C \cdot \exp(\frac{ar}{2})$ in the $\mathcal{X}$–metric, for a constant C depending only on a. Thus, if the number of vertices of balls in $\mathcal{X}$ grows faster than polynomially in the radius, then the growth exponent with respect to the horoball metric will be infinite. Furthermore, even if growth in $\mathcal{X}$ is polynomial we can make the growth exponent in the horoball be as large as we like by taking a to be sufficiently large. Dal'bo, Otal, and Peigné construct non-growth tight examples of relatively hyperbolic groups with two cusps by taking one of the horoball parameters to be large enough so that the corresponding parabolic subgroup dominates the growth of the group; that is, the growth exponent of the parabolic subgroup is equal to the growth exponent of the whole group. Quotienting by the second parabolic subgroup then does not decrease the growth exponent, so this action is not growth tight.

Not only does this provide an example of a non-growth tight action on a hyperbolic space, but since augmented spaces with different horoball parameters are equivariantly quasi-isometric to each other, we have:

Observation 8.1. Growth tightness is not invariant among equivariantly quasi-isometric G–spaces.

Dal'bo, Peigné, Picaud, and Sambusetti [26, Theorem 1.4] show that this is essentially the only way that growth tightness can fail for cusp uniform actions. Their proof is for geometrically finite Kleinian groups, but our Theorem 6.3 generalizes this result.

Definition 8.8. Let $G \curvearrowright \mathcal{X}$ be a cusp uniform action on a hyperbolic space. Let $\underline{P}$ be the collection of maximal parabolic subgroups of G. Then $G \curvearrowright \mathcal{X}$ satisfies the *Parabolic Gap Condition* if $\delta_P < \delta_G$ for all $P \in \underline{P}$.

Theorem 8.9. *Let G be a finitely generated, non-elementary group. Let $G \curvearrowright \mathcal{X}$ be a cusp uniform action on a hyperbolic space. Suppose that $G \curvearrowright \mathcal{X}$ satisfies the Parabolic Gap Condition. Then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

Proof. Let Q be the diameter of the quotient of the truncated space. The Q–complementary growth exponent is the maximum of the parabolic growth exponents, which, by the Parabolic Gap Condition, is strictly less than the growth exponent of G. Apply Theorem 6.3. □

Theorem 8.10. *Let G be a finitely generated group hyperbolic relative to a collection $\underline{P}$ of virtually nilpotent subgroups. Then there exists a hyperbolic G–space $\mathcal{X}$ such that $G \curvearrowright \mathcal{X}$ is cusp uniform and growth tight.*

Proof. Construct $\mathcal{X}$ as an augmented space by taking a Cayley graph for G and attaching combinatorial horoballs to the cosets of the peripheral subgroups. Since the peripheral groups are virtually nilpotent, they have polynomial growth in any word metric [33]. It follows that the growth exponent of each parabolic group with respect to the horoball metric is bounded by a multiple of the horoball parameter. By choosing the horoball parameters small enough, we can ensure $G \curvearrowright \mathcal{X}$ satisfies the Parabolic Gap Condition, and apply Theorem 8.9. □

8.4. Non Relative Hyperbolicity. In subsequent sections we provide further examples of growth tight actions. To show these are not redundant we will verify that the groups are not relatively hyperbolic.

In this section we recall a technique for showing that a group is not relatively hyperbolic, due to Anderson, Aramayona, and Shackleton [2]. Another approach to non relative hyperbolicity, contemporaneous to and more general than [2], and also implying Theorem 8.12, was developed by Behrstock, Druţu, and Mosher [10].

Theorem 8.11 ([2, Theorem 2]). *Let G be a finitely generated, non-elementary group, and let S be a (possibly infinite) generating set consisting of infinite order elements. Consider the 'commutativity graph' with one vertex for each element of S and an edge between vertices s and s' if some non-trivial powers of s and s' commute. If this graph is connected and there is at least one pair s, $s' \in S$ such that $\langle s, s'\rangle$ contains a rank 2 free abelian subgroup, then G is not hyperbolic relative to any finite collection of proper finitely generated subgroups.*

To prove this theorem, one shows that the subgroup generated by S is contained in one of the peripheral subgroups. Since S generates G this gives a contradiction, because the peripheral subgroups are proper subgroups of G.

We will actually use a mild generalization of Theorem 8.11 to the case when S generates a proper subgroup of G:

Theorem 8.12. *Let G be a finitely generated, non-elementary group. Let S be a set of infinite order elements whose commutativity graph is connected and such that there is a pair $s, s' \in S$ such that $\langle s, s'\rangle$ contains a rank 2 free abelian subgroup. Consider the 'coset graph' whose vertices are cosets of $\langle S\rangle$, with an edge connecting $g\langle S\rangle$ and $h\langle S\rangle$ if $g\langle S\rangle g^{-1} \cap h\langle S\rangle h^{-1}$ is infinite. If this graph is connected, then G is not hyperbolic relative to any finite collection of proper finitely generated subgroups.*

Proof. Suppose G is hyperbolic relative to $\{P_1, \ldots, P_k\}$. As in the proof of Theorem 8.11, $\langle S\rangle$ is contained in a conjugate of some P_i. We assume, without loss of generality, that $\langle S\rangle \subset P_1$. Condition (1) of Definition 8.1 implies $P_i \cap gP_ig^{-1}$ is finite for $g \notin P_i$. Thus, for $g\langle S\rangle$ adjacent to $\langle S\rangle$ in the coset graph, $g \in P_1$ and $g\langle S\rangle g^{-1} \subset P_1$. Connectivity of the coset graph implies that every element of G is contained in P_1, contradicting the hypothesis that P_1 is a proper subgroup. □

We also note that Theorem 8.11 and Theorem 8.12 imply the, a priori, stronger result that G has trivial Floyd boundary.

9. Rank 1 Actions on CAT(0) Spaces

A metric space is *CAT(0)* if every geodesic triangle is at least as thin as a triangle in Euclidean space with the same side lengths. An isometry ϕ of a CAT(0) space $\mathcal{X}$ is *hyperbolic* if $\inf_{x\in\mathcal{X}} d(x,\phi(x))$ is positive and is attained. See, for example, [17] for more background.

Let $\mathcal{X}$ be a CAT(0) G–space. Recall that our definition of 'G–space' includes the hypothesis that $\mathcal{X}$ is proper, so an element is strongly contracting if and only if it acts as a rank 1 isometry:

THEOREM 9.1 ([12, Theorem 5.4]). *Let h be a hyperbolic isometry of a proper CAT(0) space $\mathcal{X}$ with axis $\mathcal{A}$. Closest point projection to $\mathcal{A}$ is strongly contracting if and only if $\mathcal{A}$ does not bound an isometrically embedded half-flat in $\mathcal{X}$.*

Theorem 9.1 and Theorem 6.4 show:

THEOREM 9.2. *If G is a non-elementary, finitely generated group and $\mathcal{X}$ is a quasi-convex, CAT(0) G–space such that G contains an element that acts as a rank 1 isometry on $\mathcal{X}$, then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

Recall from Section 0.3 that there are many interesting classes of CAT(0) spaces that admit rank 1 isometries. In the remainder of this section we highlight a few examples.

Let Θ be a simple graph. The *Right Angled Artin Group $G(\Theta)$ defined by* Θ is the group defined by the presentation $\langle g_v \text{ for } v \in \mathrm{Vert}(\Theta) \mid g_v g_w g_v^{-1} g_w^{-1} = 1 \text{ for } [v,w] \in \mathrm{Edge}(\Theta)\rangle$. The graph Θ also determines a cube complex constructed by taking a rose with one loop for each vertex of Θ, and then gluing in a k–cube to form a k–torus for each complete k–vertex subgraph of Θ. The resulting complex is called the *Salvetti complex*, and its fundamental group is $G(\Theta)$. The universal cover of the Salvetti complex turns out to be a CAT(0) cube complex. See [22] for more background on Right Angled Artin Groups.

If Θ is a single vertex then $G(\Theta) \cong \mathbb{Z}$ is elementary. If Θ is a join, that is, if it is a complete bi-partite graph, then $G(\Theta)$ is a direct product of Right Angled Artin Groups defined by the two parts. In all other cases, we find a growth tight action:

THEOREM 9.3. *Let Θ be a finite simple graph that is not a join and has more than one vertex. The action of the Right Angled Artin Group $G(\Theta)$ defined by Θ on the universal cover $\mathcal{X}$ of the Salvetti complex associated to Θ is a growth tight action.*

PROOF. The universal cover $\mathcal{X}$ of the Salvetti complex of Θ is a cocompact, CAT(0) $G(\Theta)$–space. If Θ is not connected then $\mathcal{X}$ is hyperbolic relative to subcomplexes defined by the components of Θ, so $G(\Theta) \curvearrowright \mathcal{X}$ is growth tight by Theorem 8.6. If Θ is connected then $G(\Theta)$ contains a rank 1 isometry by a theorem of Behrstock and Charney [9]. The result follows from Theorem 9.2. □

The defining graph of a Right Angled Artin Group is a commutativity graph. If this graph is connected then the group is not relatively hyperbolic by Theorem 8.11.

A *flip-graph-manifold* is a compact three dimensional manifold M with boundary obtained from a finite collection of Seifert fibered pieces that are each a product of a circle with a compact oriented hyperbolic surface with boundary. These are glued together along boundary tori by a map exchanging the fiber and base directions. Such manifolds were studied by Kapovich and Leeb [39], who show that the universal cover of M admits a CAT(0) metric, and that an element of $\pi_1(M)$ that acts hyperbolically is rank 1 if and only if it is not represented by a loop contained in a single Seifert fibered piece. Thus, Theorem 9.2 implies:

THEOREM 9.4. *The action of the fundamental group of a flip-graph-manifold by deck transformations on its universal cover with its natural CAT(0) metric is a growth tight action.*

To see that the fundamental group of a flip-graph-manifold is not-relatively hyperbolic, apply Theorem 8.12 where S is the set of elliptic elements for the action of G on the Bass-Serre tree of the defining graph of groups decomposition.

Theorem 9.3 and Theorem 9.4 give the first non-trivial examples of growth tight actions on spaces that are not relatively hyperbolic.

The idea of the proof for flip-graph-manifolds generalizes to other CAT(0) graphs of groups via Theorem 1.14:

THEOREM 9.5. *Let G be a non-elementary, finitely generated group that splits non-trivially as a graph of groups and is not an ascending HNN-extension. Suppose that the corresponding action of G on the Bass-Serre tree of the splitting has two edges whose stabilizers have finite intersection. Suppose there exists a cocompact, CAT(0) G–space $\mathcal{X}$. Then $G \curvearrowright \mathcal{X}$ is a growth tight action.*

PROOF. By Theorem 1.14, G contains an infinite order element h such that $E(h)$ is hyperbolically embedded. A theorem of Sisto [58] implies that any axis of h is a Morse quasi-geodesic. An element with an axis that bounds a half-flat is not Morse, so h is rank 1, and the result follows by Theorem 9.2. □

10. Mapping Class Groups

Let $\mathcal{S} = \mathcal{S}_{g,p}$ be a connected and oriented surface of genus g with p punctures. We require $\mathcal{S}$ to have negative Euler characteristic.

Given two orientation-preserving homeomorphisms $\phi, \psi \colon \mathcal{S} \to \mathcal{S}$, we will consider ϕ and ψ to be equivalent if $\phi \circ \psi^{-1}$ is isotopic to the identity map on $\mathcal{S}$. Each equivalence class is called a *mapping class* of $\mathcal{S}$, and the set $\mathrm{Mod}(\mathcal{S})$ of all equivalence classes naturally forms a group called the *mapping class group* of $\mathcal{S}$.

A mapping class $f \in \mathrm{Mod}(\mathcal{S})$ is called *reducible* is there exists an f–invariant curve system on $\mathcal{S}$ and *irreducible* otherwise. By the Nielsen-Thurston classification of elements of $\mathrm{Mod}(\mathcal{S})$, a mapping class is irreducible and infinite order if and only if it is pseudo-Anosov [60].

Let $\mathcal{X}$ be the Teichmüller space of marked hyperbolic structures on $\mathcal{S}$, equipped with the Teichmüller metric. See [37] and [45] for more information.

THEOREM 10.1 ([43]). *Every pseudo-Anosov element is strongly contracting for* $\mathrm{Mod}(\mathcal{S}) \curvearrowright \mathcal{X}$.

For each $\epsilon > 0$ there is a decomposition of $\mathcal{X}$ into a 'thick part' $\mathcal{X}^{\geqslant\epsilon}$ and a 'thin part' $\mathcal{X}^{<\epsilon}$ according to whether the hyperbolic structure on $\mathcal{S}$ corresponding to the point $x \in \mathcal{X}$ has any closed curves of length $< \epsilon$. This decomposition is $\mathrm{Mod}(\mathcal{S})$–invariant, and $\mathrm{Mod}(\mathcal{S}) \curvearrowright \mathcal{X}^{\geqslant\epsilon}$ is cocompact, see [44] and [30]. Geodesics between points in the thick part can travel deeply into the thin part, so the action of $\mathrm{Mod}(\mathcal{S})$ on Teichmüller space is not quasi-convex. To prove growth tightness need a bound on the complementary growth exponent. Such a bound is provided by a recent theorem of Eskin, Mirzakhani, and Rafi [29, Theorem 1.7].

THEOREM 10.2. *The action of the mapping class group* $\mathrm{Mod}(\mathcal{S})$ *of* $\mathcal{S} = \mathcal{S}_{g,p}$ *on its Teichmüller space $\mathcal{X}$ with the Teichmüller metric is a growth tight action.*

PROOF. Let $\zeta = 6g - 6 + 2p \geqslant 2$. The growth exponent of $\mathrm{Mod}(\mathcal{S})$ with respect to its action on $\mathcal{X}$ is ζ [5]. (We remark that the result of [5] is stated for closed surfaces, but their proof works in general. For our interest, it is enough that the growth exponent of $\mathrm{Mod}(\mathcal{S})$ is bounded below by ζ. This can be obtained from [36] and [29].)

Choose an r_0 and a maximal r_0–separated set in moduli space $\mathrm{Mod}(\mathcal{S})\backslash\mathcal{X}$, and let $\mathcal{A}$ be its full lift to $\mathcal{X}$. Given r_0 as above and $\delta = \frac{1}{2}$, let ϵ be sufficiently small as in [29, Theorem 1.7]. Let Q be the smallest number such that the ϵ–thick part of $\mathcal{X}$ is contained in $\overline{\mathcal{B}}_Q(\mathrm{Mod}(\mathcal{S}).o)$. Choose a finite subset $\{a_1, \ldots, a_n\} \subset \mathcal{A}$ such that:

$$\overline{\mathcal{B}}_Q(o) \setminus \mathcal{B}_Q(\mathrm{Mod}(\mathcal{S}).o) \subset \bigcup_{i=1}^{n} \mathcal{B}_{r_0}(a_i)$$

Suppose that $g \in \mathrm{Mod}(\mathcal{S})$ is such that there exists a geodesic $[x,y]$ between $\overline{\mathcal{B}}_Q(o)$ and $\overline{\mathcal{B}}_Q(g.o)$ whose interior stays in $\mathcal{X} \setminus \overline{\mathcal{B}}_Q(\mathrm{Mod}(\mathcal{S}).o)$. Then there are indices i and j such that $x \in \mathcal{B}_{r_0}(a_i)$ and $y \in \mathcal{B}_{r_0}(g.a_j)$. This means that every element contributing to $Comp_{Q,r}^{\mathrm{Mod}(\mathcal{S})}$ of Definition 6.2 also contributes to some $N_1(Q_{1,\epsilon}, a_i, a_j, r)$ of [29, Theorem 1.7]. The conclusion of [29, Theorem 1.7] is that $N_1(Q_{1,\epsilon}, a_i, a_j, r) \leqslant G(a_i)G(a_j)\exp(r \cdot (\zeta - \frac{1}{2}))$ for all sufficiently large r, where G is a particular function on $\mathcal{X}$. There are finitely many such sets, and the function G is bounded on $\{a_1, \ldots, a_n\}$, so there is a constant C such that $Comp_{Q,r}^{\mathrm{Mod}(\mathcal{S})} \leqslant C \cdot \exp(r \cdot (\zeta - \frac{1}{2}))$ for all sufficiently large r. Thus, the Q–complementary growth exponent is at most $\zeta - \frac{1}{2} < \zeta$. The theorem now follows from Theorem 10.1 and Theorem 6.3. □

When the genus of $\mathcal{S}$ is at least 3 then there does not exist a cocompact, CAT(0) $\mathrm{Mod}(\mathcal{S})$–space [16]. The fact that such an $\mathrm{Mod}(\mathcal{S})$ is not relatively hyperbolic (in fact, has trivial Floyd boundary) is an application of Theorem 8.11 appearing in [2]. Therefore, Theorem 10.2 does not follow from the results of the previous sections.

A natural question is whether the action of a mapping class group on its Cayley graphs is growth tight. There is also a combinatorial model for the mapping class group known as the *marking complex*. Finally, a mapping class group acts cocompactly on a thick part of Teichmüller space. All of these spaces are quasi-isometric, and Duchin and Rafi [28] show that pseudo-Anosov elements are contracting for the action of a mapping class group on any one of its Cayley graphs, but we do not know whether one of these actions admits a strongly contacting element.

Question 10.3. Is the action of a mapping class group of a hyperbolic surface on one of its Cayley graphs/marking complex/thick part of Teichmüller space growth tight?

The outer automorphism group of a finite rank non-abelian free group, $\mathrm{Out}(F_n)$ is often studied in analogy with $\mathrm{Mod}(\mathcal{S})$. Algom-Kfir [1] has proven an analogue of Minsky's theorem that says that a *fully irreducible* outer automorphism class is strongly contracting for the action of $\mathrm{Out}(F_n)$ on its Outer Space, which is the analogue of the Teichmüller space. However, we lack the analogue of the Eskin-Mirzakhani-Rafi theorem that was used to control the complementary growth exponent in the mapping class group case.

There is also an analogue of the thick part of Teichmüller space called the *spine* of the Outer Space, on which $\mathrm{Out}(F_n)$ acts cocompactly.

Question 10.4. Is the action of $\mathrm{Out}(F_n)$ on one of its Cayley graphs/Outer Space/spine of Outer Space growth tight?

11. Snowflake Groups

Let $G := BB(1,r) = \langle a,b,s,t \mid aba^{-1}b^{-1} = 1,\ s^{-1}as = a^rb,\ t^{-1}at = a^rb^{-1}\rangle$ be a Brady-Bridson snowflake group with $r \geqslant 3$. Let $L := 2r$. These groups have an interesting mixture of positive and negative curvature properties. G splits as an amalgam of $\mathbb{Z}^2 = \langle a,b\rangle$ by two cyclic groups $\langle a^rb\rangle$ and $\langle a^rb^{-1}\rangle$, and the action of G on the Bass-Serre tree $\mathcal{T}$ of this splitting satisfies Theorem 1.14, so G has hyperbolically embedded subgroups. However, we can not automatically conclude that such a hyperbolically embedded subgroup gives rise to a strongly contracting element, as there does not exist a cocompact, CAT(0) G–space. If such a space existed, then the Dehn function of G would be at most quadratic, but Brady and Bridson [14] have shown that the Dehn function of $BB(1,r)$ is $n^{2\log_2 L} > n^2$.

We will fix a G–space $\mathcal{X}$ and demonstrate two different elements of G that act hyperbolically on $\mathcal{T}$ such that the pointwise stabilizer of any length 3 segment of their axes is finite. One of these elements will be strongly contracting for the action on $\mathcal{X}$, and the other will not. Hence:

Theorem 11.1. *G admits a cocompact growth tight action.*

Observe that Theorem 8.12 with $S := \{a,b\}$ shows that G is not relatively hyperbolic.

11.1. The Model Space $\mathcal{X}$. Let $\mathcal{X}$ be the Cayley graph for G with respect to the generating set $\{a, a^rb, a^rb^{-1}, s, t\}$, where the edges corresponding to a^rb and a^rb^{-1} have been rescaled to have length $L := 2r$. The point of scaling these edges is that a^rb, a^rb^{-1}, and a^{2r} form an equilateral triangle of side length L, which will facilitate finding geodesics in this particular model.

It is also useful to consider G as the fundamental group of the topological space obtained from a torus by gluing on two annuli. Choose a basepoint for the torus and for each boundary component of the annuli. For one annulus, the s–annulus, glue the two boundary curves to the curves a and a^rb in the torus, gluing basepoints to the basepoint of the torus. For the other annulus, the t–annulus, glue the two boundary curves to the curves a and a^rb^{-1} of the torus. The resulting space is a graph of spaces [53] associated to the given graph of groups decomposition of G. The fundamental group of this space is G, which acts freely by deck transformations on the universal cover $\mathcal{X}'$. Choose the basepoint o of $\mathcal{X}'$ to be a lift of the basepoint of the torus. The correspondence between a vertex $g \in \mathcal{X}$ and the point $g.o \in \mathcal{X}'$ inspires the following terminology: A *plane* is a coset $g\langle a, b\rangle \in G/\langle a, b\rangle$, which corresponds to a lift of the torus at the point $g.o \in \mathcal{X}'$. An *s-wall* is the set of outgoing s–edges incident to a coset $g\langle a\rangle \in G/\langle a\rangle$. This corresponds to a lift of the s–annulus at the point $g.o \in \mathcal{X}'$. A *t-wall* is the set of outgoing t–edges incident to a coset $g\langle a\rangle \in G/\langle a\rangle$. This corresponds to a lift of the t–annulus at the point $g.o \in \mathcal{X}'$. Each wall separates $\mathcal{X}$ (and $\mathcal{X}'$) into two complementary components. Notice that the origins of consecutive edges in an s–wall are connected by a single a–edge of length 1, while the termini of those edges are connected by a single a^rb–edge of length L. We say that crossing an s–wall in the positive direction scales distance by a factor of L. The same is true for the t–walls.

11.2. Geodesics Between Points in a Plane. We will define a family of $\mathcal{X}$–geodesics joining 1 to every point of $\langle a, b\rangle$. This is similar to the argument of [14].

From the fact that $\langle a, b\rangle$ is abelian, for every point a^xb^y there is a geodesic from 1 to a^xb^y of the form:

$$[1, (a^rb)^m] + (a^rb)^m[1, (a^rb^{-1})^n] + (a^rb)^m(a^rb^{-1})^n[1, a^p]$$

where $[g, h]$ indicates a geodesic from g to h.

For a point of the form $(a^rb)^m$ there is an a^rb–edge path from 1 to $(a^rb)^m$ of length mL. This path is clearly inefficient, as it lies along the boundary of an s^{-1}–wall that scales distance by $1/L$, so we can push the original edge path across the wall to a path $s^{-1}a^ms$ of length $2 + m$. We claim there is a geodesic from 1 to $(a^rb)^m$ of the form $[1, s^{-1}] + s^{-1}[1, a^m] + s^{-1}a^m[s^{-1}, 1]$. We have already exhibited a wall-crossing path of length $2 + m$, which is shorter than any path from 1 to $(a^rb)^m$ that stays in the plane $\langle a, b\rangle$. Thus, a geodesic must cross some walls. Every path from 1 to $(a^rb)^m$ can, by rearranging subsegments and eliminating backtracking, be replaced by a path of at most the same length and having the form $\gamma_s + \gamma_t + \gamma'$ where:

- $\gamma_s = [1, s^{-1}] + s^{-1}[1, a^n] + s^{-1}a^n[s^{-1}, 1]$ if non-trivial.
- $\gamma_t = s^{-1}a^ns[1, t^{-1}] + s^{-1}a^nst^{-1}[1, a^p] + s^{-1}a^nst^{-1}a^p[t^{-1}, 1]$ if non-trivial.
- $\gamma' = s^{-1}a^nst^{-1}a^pt[1, a^q]$ if non-trivial.

The path $\gamma = \gamma_s + \gamma_t + \gamma'$ is a path from 1 to $s^{-1}a^nst^{-1}a^pta^q = (a^rb)^n(a^rb^{-1})^pa^q = a^{r(n+p)+q}b^{n-p} = a^{rm}b^m$, so $p = n - m$ and $q = -Lp$. Since p and q are proportional, γ_t and γ' are either both trivial or both non-trivial. Suppose they are non-trivial. There is a symmetry that exchanges γ_t with a path $\gamma_t' = s^{-1}a^ns[1, s^{-1}] + s^{-1}a^nss^{-1}[1, a^{-p}] + s^{-1}a^nss^{-1}a^{-p}[s^{-1}, 1]$ of the same length. However, γ_t' and $\gamma_t + \gamma'$ have the same endpoints, and γ_t' is shorter, so γ could not have been geodesic if γ_t and γ' are non-trivial. Thus, if γ is geodesic then $\gamma = \gamma_s$. This reduces the problem of finding a geodesic from 1 to $(a^rb)^m$ to finding a geodesic from 1 to a^n.

A similar argument holds for geodesics from 1 to $(a^rb^{-1})^m$, so we can find geodesics from 1 to any point in $\langle a, b\rangle$ if we know geodesics from 1 to powers of a.

For powers of a the idea is that a^{mL}, $(a^rb)^m$, and $(a^rb^{-1})^m$ form an equilateral triangle in the plane, but the latter two can be shortened by a factor of L by pushing across a wall. Since $L \geq 6$, the savings of a factor of $L/2$ in length outweighs the added overhead from crossing walls.

For small powers of a we can find geodesics by inspection of the Cayley graph. For $0 \leqslant |p| \leqslant {}^{L}\!/_{2} + 3$, the edge path a^p from 1 to a^p is a geodesic of length $|p|$. For ${}^{L}\!/_{2} + 3 \leqslant p \leqslant L$ the edge path $s^{-1}ast^{-1}ata^{p-L}$ is a geodesic from 1 to a^p of length $6 + L - p$. We conclude that for $m > 0$ and $-{}^{L}\!/_{2} + 3 \leq p \leq {}^{L}\!/_{2} + 3$ there is a geodesic from 1 to a^{mL+p} of the form:

$$\begin{aligned}&[1, s^{-1}] + s^{-1}[1, a^m] + s^{-1}a^m[s^{-1}, 1]\\ &\quad + s^{-1}a^m s[1, t^{-1}] + s^{-1}a^m st^{-1}[1, a^m] + s^{-1}a^m st^{-1}a^m[t^{-1}, 1]\\ &\quad + s^{-1}a^m st^{-1}a^m t[1, a^p]\end{aligned}$$

We can now find geodesics from 1 to powers of a by induction, and from these we know a geodesic from 1 to any $a^x b^y$. We see an example in Figure 5, where trapezoids are walls and triangles are contained in planes. The top half boundary and bottom half boundary of the figure each give geodesics of length $5 \cdot 2^5 - 4$ between 1 and a^{L^5}. (This form of geodesic loop bears witness to the Dehn function [14], and inspired the name 'snowflake group' [15].)

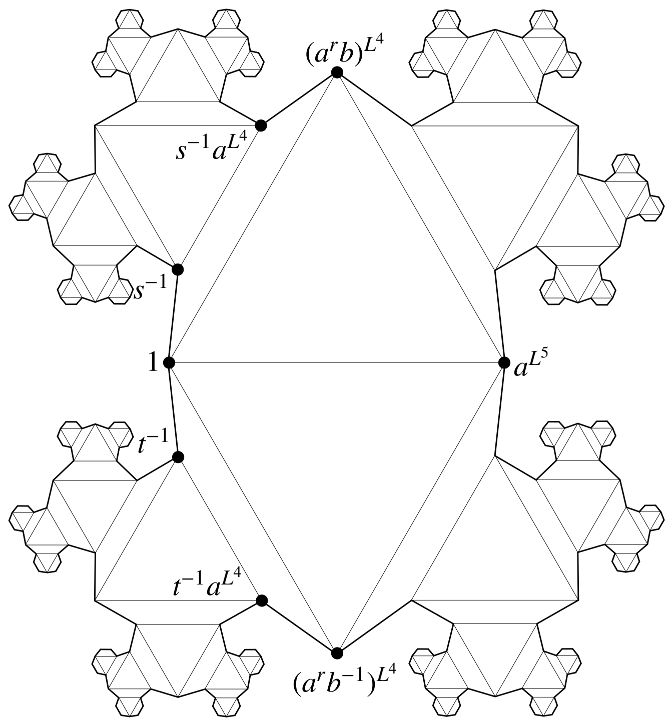

FIGURE 5. Snowflake - The boundary is a geodesic loop of length $2(5 \cdot 2^5 - 4)$

11.3. Projections to Geodesics in $\mathcal{X}$. In this section we consider two different geodesics: $\alpha(2n) = (s^{-1}t)^n$ and $\beta(n) = s^{-n}$. These are geodesics since for each of these paths, every edge crosses a distinct wall. Let $\mathcal{T}$ be the Bass-Serre tree of G, and let $o \in \mathcal{T}$ be the vertex fixed by the subgroup $\langle a, b\rangle$. The orbit map $g \mapsto g.o$ sends each of α and β isometrically to a geodesic in $\mathcal{T}$. We will use π_α to denote closest point projection to α, both in $\mathcal{X}$ and in $\mathcal{T}$, and similarly for β.

Both of these geodesics have the property that for any vertices at distance at least three in the corresponding geodesic of the Bass-Serre tree, the pointwise stabilizers of the pair of vertices is trivial. We might hope, in analogy to Theorem 9.5, that these would be strongly contracting geodesics. As in Theorem 9.5, $\langle s^{-1}t\rangle$ and $\langle s\rangle$ are hyperbolically embedded subgroups in G, but, of the two, we will see only $s^{-1}t$ is strongly contracting.

11.3.1. α. We claim that closest point projection $\pi_\alpha \colon \mathcal{X} \to \alpha$ is coarsely well defined and strongly contracting. First, consider π_α on $\langle a, b\rangle$. The geodesic α enters $\langle a, b\rangle$ through the incoming t–wall V at 1, and exits through the outgoing s^{-1}–wall W at 1.

Lemma 11.2. *For every $v \in V$ and every $w \in W$ there exists a geodesic from v to w that includes the vertex 1.*

Proof. The lemma follows from the discussion of geodesics in Section 11.2. □

Lemma 11.3. *The orbit map $\mathcal{X} \to \mathcal{T} : g \mapsto g.o$ coarsely commutes with closest point projection to α. In particular, closest point projection to α in $\mathcal{X}$ is coarsely well defined.*

Proof. Suppose $z \in \mathcal{X}$ is some vertex that is separated from 1 by V, and suppose there is an $n \geqslant 0$ such that $\alpha(n) \in \pi_\alpha(z)$. Let σ be a geodesic from z to $\alpha(n)$. Write $\sigma = \sigma_1 + \sigma_2 + \sigma_3$, where σ_2 is the subsegment of σ from the first time σ crosses V until the first time σ reaches W. By Lemma 11.2, we can replace σ_2 by a geodesic segment $\sigma_2' + \sigma_2''$ where the concatenation point is 1. This means that z is connected to $1 = \alpha(0)$ by a path $\sigma_1 + \sigma_2'$. By hypothesis, the length of this path is at least the length of σ, so σ_2'' and σ_3 are trivial and $n = 0$. It follows immediately that the orbit map $\mathcal{X} \to \mathcal{T}$ commutes with π_α up to an error of 4. (In fact, a little more work will show the error is at most 2.) □

Lemma 11.4 (Bounded Geodesic Image Property for π_α). *For any geodesic σ in $\mathcal{X}$, if the diameter of $\pi_\alpha(\sigma.o)$ is at least 5, then $\sigma \cap \alpha \neq \emptyset$.*

Proof. Suppose $\alpha([-1,3]).o \subset \pi_\alpha(\sigma.o)$. Then σ crosses the walls V, W, $s^{-1}tV$ and $s^{-1}tW$. Write σ as a concatenation of geodesic subsegments $\sigma_1 + \sigma_2 + \sigma_3 + \sigma_4 + \sigma_5$, where σ_1 is all of σ prior to the first V crossing, σ_2 is the part of σ between the first V crossing and the last W crossing, σ_3 is the part between the last W crossing and the first $s^{-1}tV$ crossing, which included edges labeled s^{-1} and t, σ_4 is the part from the first $s^{-1}tV$ crossing until the last $s^{-1}tW$ crossing, and σ_5 is the remainder of σ. We can apply Section 11.2 to replace σ_2 by a geodesic $\sigma_2' + \sigma_2''$ with the same endpoints and concatenated at 1. Similarly, we can replace σ_4 by a geodesic $\sigma_4' + \sigma_4''$ with the same endpoints and concatenated at $s^{-1}t$. But then we can replace the subsegment $\sigma_2 + \sigma_3 + \sigma_4$ of σ by the path $\sigma_2'' + [1, s^{-1}t] + \sigma_4''$ with the same endpoints. This path is strictly shorter unless σ_2'' and σ_4'' are trivial. This means that $[1, s^{-1}t] \subset \sigma \cap \alpha$. □

By Proposition 2.9, this means:

Corollary 11.5. *The element $s^{-1}t$ is strongly contracting for $G \curvearrowright \mathcal{X}$.*

Together with Theorem 6.4, this proves Theorem 11.1.

11.3.2. *β.* Using out knowledge of geodesics from Section 11.2, we see that the closest point of the s^{-1}–wall at 1 to the point a^{L^k} is $(a^r b)^{L^{k-1}}$, which is the midpoint of a geodesic from 1 to a^{L^k}. This geodesic coincides with β on the interval from 1 to s^{-k}. It follows that $\pi_\beta(a^{L^j}) = \beta(j)$ for all $j \geqslant 0$.

For $0 < j < k$ there is a geodesic $\sigma_{j,k}$ from a^{L^j} to a^{L^k} such that $d(\sigma_{j,k}, \beta) = d(a^{L^j}, \beta)$. See Figure 6. Letting j and $k - j$ grow large, the geodesics $\sigma_{j,k}$ for stay outside large neighborhoods of β but have large projections to β. Therefore, π_β is not strongly contracting, since it does not enjoy the Bounded Geodesic Image Property.

References

[1] Yael Algom-Kfir. *Strongly contracting geodesics in outer space*. Geom. Topol. **15**.4 (2011), pp. 2181–2233.

[2] James W. Anderson, Javier Aramayona, and Kenneth J. Shackleton. *An obstruction to the strong relative hyperbolicity of a group*. J. Group Theory **10**.6 (2007), pp. 749–756.

[3] Goulnara N. Arzhantseva and Igor G. Lysenok. *A lower bound on the growth of word hyperbolic groups*. J. London Math. Soc. (2) **73**.1 (2006), pp. 109–125.

[4] Goulnara N. Arzhantseva and Igor G. Lysenok. *Growth tightness for word hyperbolic groups*. Math. Z. **241**.3 (2002), pp. 597–611.

[5] Jayadev Athreya, Alexander Bufetov, Alex Eskin, and Maryam Mirzakhani. *Lattice point asymptotics and volume growth on Teichmüller space*. Duke Math. J. **161**.6 (2012), pp. 1055–1111.

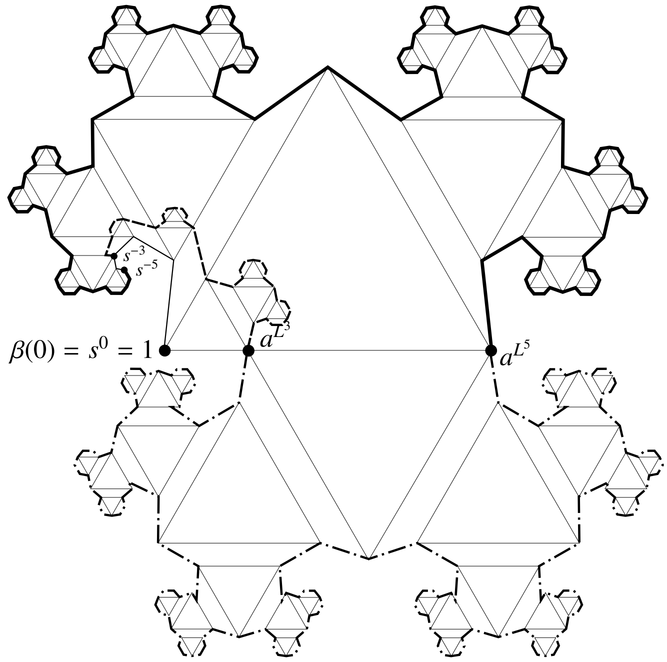

FIGURE 6. Geodesics $[a^{L^3}, \pi_\beta(a^{L^3})]$ (dashed), $[a^{L^5}, \pi_\beta(a^{L^5})]$ (solid), and $\sigma_{3,5} = [a^{L^3}, a^{L^5}]$ (dash-dot)

[6] Werner Ballmann. *Lectures on spaces of nonpositive curvature*. Vol. 25. DMV Seminar. With an appendix by Misha Brin. Basel: Birkhäuser Verlag, 1995, pp. viii+112.

[7] Werner Ballmann and Michael Brin. *Orbihedra of Nonpositive Curvature*. Publ. Math. Inst. Hautes Études Sci. **82**.1 (1995), pp. 169–209.

[8] Werner Ballmann and Sergei Buyalo. *Periodic rank one geodesics in Hadamard spaces*. *Geometric and probabilistic structures in dynamics*. Vol. 469. Contemp. Math. Providence, RI: Amer. Math. Soc., 2008, pp. 19–27.

[9] Jason Behrstock and Ruth Charney. *Divergence and quasimorphisms of right-angled Artin groups*. Math. Ann. **352**.2 (2012), pp. 339–356.

[10] Jason Behrstock, Cornelia Druţu, and Lee Mosher. *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*. Math. Ann. **344**.3 (2009), pp. 543–595.

[11] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. *Constructing group actions on quasi-trees and applications to mapping class groups*. Publ. Math. Inst. Hautes Études Sci. (2014), pp. 1–64.

[12] Mladen Bestvina and Koji Fujiwara. *A characterization of higher rank symmetric spaces via bounded cohomology*. Geom. Funct. Anal. **19**.1 (2009), pp. 11–40.

[13] Brian H. Bowditch. *Relatively hyperbolic groups*. Internat. J. Algebra Comput. **22**.3 (2012), pp. 1250016, 66.

[14] Noel Brady and Martin R. Bridson. *There is only one gap in the isoperimetric spectrum*. Geom. Funct. Anal. **10**.5 (2000), pp. 1053–1070.

[15] Noel Brady, Martin R. Bridson, Max Forester, and Krishnan Shankar. *Snowflake groups, Perron-Frobenius eigenvalues and isoperimetric spectra*. Geom. Topol. **13**.1 (2009), pp. 141–187.

[16] Martin R. Bridson. *Semisimple actions of mapping class groups on* CAT(0) *spaces*. *Geometry of Riemann surfaces*. Vol. 368. London Math. Soc. Lecture Note Ser. Cambridge: Cambridge Univ. Press, 2010, pp. 1–14.

[17] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Vol. 319. Grundlehren der mathematischen Wissenschaften. Berlin: Springer, 1999.

[18] Marc Burger and Shahar Mozes. *Finitely presented simple groups and products of trees*. C. R. Acad. Sci. Paris Sér. I Math. **324**.7 (1997), pp. 747–752.

[19] Pierre-Emmanuel Caprace and Koji Fujiwara. *Rank-one isometries of buildings and quasi-morphisms of Kac-Moody groups*. Geom. Funct. Anal. **19**.5 (2010), pp. 1296–1319.

[20] Pierre-Emmanuel Caprace and Michah Sageev. *Rank rigidity for CAT(0) cube complexes*. Geom. Funct. Anal. **21**.4 (2011), pp. 851–891.

[21] Christopher H. Cashen and Jing Tao. *Growth Tight Actions of Product Groups*. Groups Geom. Dyn. (in press). arXiv: `1404.4335v2`.

[22] Ruth Charney. *An introduction to right-angled Artin groups*. Geom. Dedicata **125** (2007), pp. 141–158.

[23] Rémi Coulon. *Growth of periodic quotients of hyperbolic groups*. Algebr. Geom. Topol. **13**.6 (2013), pp. 3111–3133.

[24] François Dahmani, Vincent Guirardel, and Denis Osin. *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*. preprint. 2011. arXiv: `1111.7048v5`.

[25] Françoise Dal'bo, Jean-Pierre Otal, and Marc Peigné. *Séries de Poincaré des groupes géométriquement finis*. Israel J. Math. **118** (2000), pp. 109–124.

[26] Françoise Dal'Bo, Marc Peigné, Jean-Claude Picaud, and Andrea Sambusetti. *On the growth of quotients of Kleinian groups*. Ergodic Theory Dynam. Systems **31**.3 (2011), pp. 835–851.

[27] Cornelia Druţu. *Relatively hyperbolic groups: geometry and quasi-isometric invariance*. Comment. Math. Helv. **84**.3 (2009), pp. 503–546.

[28] Moon Duchin and Kasra Rafi. *Divergence of geodesics in Teichmüller space and the mapping class group*. Geom. Funct. Anal. **19**.3 (2009), pp. 722–742.

[29] Alex Eskin, Maryam Mirzakhani, and Kasra Rafi. *Counting closed geodesics in strata*. preprint. 2012. arXiv: `1206.5574v1`.

[30] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Vol. 49. Princeton Mathematical Series. Princeton, NJ: Princeton University Press, 2012, pp. xiv+472.

[31] Victor Gerasimov. *Floyd maps for relatively hyperbolic groups*. Geom. Funct. Anal. **22**.5 (2012), pp. 1361–1399.

[32] Rostislav Grigorchuk and Pierre de la Harpe. *On problems related to growth, entropy, and spectrum in group theory*. J. Dynam. Control Systems **3**.1 (1997), pp. 51–89.

[33] Michael Gromov. *Groups of Polynomial Growth and Expanding Maps*. Publ. Math. Inst. Hautes Études Sci. **53** (1981). with an appendix by Jacques Tits, pp. 53–78.

[34] Daniel Groves and Jason Fox Manning. *Dehn filling in relatively hyperbolic groups*. Israel Journal of Mathematics **168**.1 (2008), pp. 317–429.

[35] Daniel Groves and Jason Fox Manning. *Fillings, finite generation and direct limits of relatively hyperbolic groups*. Groups Geom. Dyn. **1**.3 (2007), pp. 329–342.

[36] Ursula Hamenstädt. *Bowen's construction for the Teichmüller flow*. J. Mod. Dyn. **7**.4 (2013), pp. 489–526.

[37] John Hamal Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 1*. Matrix Editions, Ithaca, NY, 2006, pp. xx+459.

[38] Heinz Huber. *Über eine neue Klasse automorpher Funktionen und ein Gitterpunktproblem in der hyperbolischen Ebene. I*. Comment. Math. Helv. **30** (1956), 20–62 (1955).

[39] Michael Kapovich and Bernhard Leeb. 3-*manifold groups and nonpositive curvature*. Geom. Funct. Anal. **8**.5 (1998), pp. 841–852.

[40] Jason Fox Manning. *Geometry of pseudocharacters*. Geom. Topol. **9** (2005), pp. 1147–1185.

[41] Gregori A. Margulis. *Discrete subgroups of semisimple Lie groups*. Vol. 17. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Berlin: Springer-Verlag, 1991, pp. x+388.

[42] Ashot Minasyan and Denis Osin. *Acylindrical hyperbolicity of groups acting on trees*. Math. Ann. (2014), pp. 1–51.

[43] Yair N. Minsky. *Quasi-projections in Teichmüller space*. J. Reine Angew. Math. **473** (1996), pp. 121–136.

[44] David Mumford. *A remark on Mahler's compactness theorem*. Proc. Amer. Math. Soc. **28** (1971), pp. 289–294.

[45] Athanase Papadopoulos. *Introduction to Teichmüller theory, old and new*. *Handbook of Teichmüller theory. Vol. I*. Vol. 11. IRMA Lect. Math. Theor. Phys. Eur. Math. Soc., Zürich, 2007, pp. 1–30.

[46] Jouni Parkkonen and Frédéric Paulin. *On the hyperbolic orbital counting problem in conjugacy classes*. Math. Z. **279**.3-4 (2015), pp. 1175–1196.

[47] Cornelius Reinfeldt and Richard Weidmann. *Makanin-Razborov diagrams for hyperbolic groups*. preprint. 2010. URL: `http://www.math.uni-kiel.de/algebra/weidmann/mr.pdf`.

[48] Stéphane Sabourau. *Growth of quotients of groups acting by isometries on Gromov-hyperbolic spaces*. J. Mod. Dyn. **7**.2 (2013), pp. 269–290.

[49] Andrea Sambusetti. *Growth tightness in group theory and Riemannian geometry*. *Recent advances in geometry and topology*. Cluj Univ. Press, Cluj-Napoca, 2004, pp. 341–352.

[50] Andrea Sambusetti. *Growth tightness of free and amalgamated products*. Ann. Sci. École Norm. Sup. (4) **35**.4 (2002), pp. 477–488.

[51] Andrea Sambusetti. *Growth tightness of negatively curved manifolds*. C. R. Math. Acad. Sci. Paris **336**.6 (2003), pp. 487–491.

[52] Andrea Sambusetti. *Growth tightness of surface groups*. Expo. Math. **20**.4 (2002), pp. 345–363.

[53] Peter Scott and Terry Wall. *Topological methods in group theory*. *Homological group theory (Proc. Sympos., Durham, 1977)*. Vol. 36. London Math. Soc. Lecture Note Ser. Cambridge: Cambridge Univ. Press, 1979, pp. 137–203.

[54] Zlil Sela. *Endomorphisms of hyperbolic groups. I. The Hopf property*. Topology **38**.2 (1999), pp. 301–321.

[55] Alexei G. Shukhov. *On the dependence of the growth exponent on the length of the defining relation*. Mat. Zametki **65**.4 (1999), pp. 612–618.

[56] Alessandro Sisto. *Contracting elements and random walks*. preprint. 2011. arXiv: `1112.2666v2`.

[57] Alessandro Sisto. *On metric relative hyperbolicity*. preprint. 2012. arXiv: `1210.8081v1`.

[58] Alessandro Sisto. *Quasi-convexity of hyperbolically embedded subgroups*. Math. Z. (in press). arXiv: `1310.7753`.

[59] John R. Stallings. *Group theory and three-dimensional manifolds*. A James K. Whittemore Lecture in Mathematics given at Yale University, 1969, Yale Mathematical Monographs, 4. New Haven, Conn.: Yale University Press, 1971, pp. v+65.

[60] William P. Thurston. *Hyperbolic Structures on 3–Manifolds, II: Surface Groups and 3–manifolds which Fiber Over the Circle*. preprint. 1986. arXiv: `math/9801045v1`.

[61] Wenyuan Yang. *Growth tightness of groups with nontrivial Floyd boundary*. preprint. 2013. arXiv: `1301.5623v1`.

PAPER E

Growth Tight Actions of Product Groups

Christopher H. Cashen and Jing Tao
Groups, Geometry and Dynamics **10** (2016), no. 2, 753–770.
DOI:10.4171/GGD/364

A group action on a metric space is called growth tight if the exponential growth rate of the group with respect to the induced pseudo-metric is strictly greater than that of its quotients. A prototypical example is the action of a free group on its Cayley graph with respect to a free generating set. More generally, with Arzhantseva we have shown that group actions with strongly contracting elements are growth tight.

Examples of non-growth tight actions are product groups acting on the L^1 products of Cayley graphs of the factors.

In this paper we consider actions of product groups on product spaces, where each factor group acts with a strongly contracting element on its respective factor space. We show that this action is growth tight with respect to the L^p metric on the product space, for all $1 < p \leq \infty$. In particular, the L^∞ metric on a product of Cayley graphs corresponds to a word metric on the product group. This gives the first examples of groups that are growth tight with respect to an action on one of their Cayley graphs and non-growth tight with respect to an action on another, answering a question of Grigorchuk and de la Harpe.

1. Introduction

The *growth exponent* of a set A with respect to a pseudo-metric d is

$$\delta_{A,d} = \limsup_{r\to\infty} \frac{1}{r} \cdot \log \#\{a \in A \mid d(o,a) \leq r\}$$

where # denotes cardinality and $o \in A$ is some basepoint. The limit is independent of the choice of basepoint.

Let G be a finitely generated group, and let $(\mathcal{X}, d, o)$ be a proper, based, geodesic metric space on which G acts properly discontinuously and cocompactly by isometries.

The metric d induces a left invariant pseudo-metric $\bar{d}$ on any quotient G/N of G by $\bar{d}(gN, g'N) = \min_{n,n' \in N} d(gn.o, g'n'.o)$. When $(\mathcal{X}, d, o)$ is clear we let $\delta_{G/N}$ denote $\delta_{G/N,\bar{d}}$ and let δ_G denote $\delta_{G/\{1\},\bar{d}}$.

Definition 1.1 ([1]). $G \curvearrowright \mathcal{X}$ is a *growth tight action* if $\delta_G > \delta_{G/N}$ for every infinite normal subgroup $N \trianglelefteq G$.

If S is a finite generating set of G, we say G is *growth tight with respect to S* if the action of G via left multiplication on the Cayley graph of G with respect to S is growth tight.

The first examples of such actions were given by Grigorchuk and de la Harpe [9], who showed that a finite rank, non-abelian free group $\mathbb{F}$ is growth tight with respect to any free generating set S. In the same paper, they observe that the product $\mathbb{F} \times \mathbb{F}$ is not growth tight with respect to the generating set $S \times \{1\} \cup \{1\} \times S$, and ask whether there exists a finite generating set with respect to which $\mathbb{F} \times \mathbb{F}$ is growth tight.

The first author is supported by the European Research Council (ERC) grant of Goulnara ARZHANTSEVA, grant agreement #259527 and the Erwin Schrödinger Institute workshop “Geometry of Computation in Groups”. The second author is partially supported by NSF grant DMS-1311834.

We answer this question affirmatively. This is the first example of a group that is growth tight with respect to one generating set and not growth tight with respect to another.

Our main result is for growth tightness of product groups $G_1 \times \cdots \times G_n$. We require that each factor G_i acts cocompactly with a strongly contracting element on a space $\mathcal{X}_i$, see Definition 2.2. Examples include actions of hyperbolic or relatively hyperbolic groups by left multiplication on any of their Cayley graphs, and groups acting cocompactly on proper CAT(0) spaces with rank 1 isometries. With Arzhantseva [1], we have shown that such actions are growth tight.

THEOREM 1.2. *For $1 \leq i \leq n$, let G_i be a non-elementary, finitely generated group acting properly discontinuously and cocompactly by isometries on a proper, based, geodesic metric space $(\mathcal{X}_i, d_i, o_i)$ with a strongly contracting element. Let $G = G_1 \times \cdots \times G_n$. Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, with $o = (o_1, \ldots, o_n)$ and let d be the L^p metric on $\mathcal{X}$ for some $1 \leq p \leq \infty$. Let $G \curvearrowright \mathcal{X}$ be the coordinate-wise action. Then $G \curvearrowright \mathcal{X}$ is growth tight unless $p = 1$ and $n > 1$.*

REMARK 1.3. Cocompactness of the factor actions is not strictly necessary. We use it to prove a subadditivity result, Lemma 4.4. There are weaker conditions than cocompactness of the action that can be used to prove such a result. These are discussed in [1, Section 6]. For simplicity, we will stick to cocompact actions in this paper, since this suffices for our main applications.

In the case that $\mathcal{X}_i$ is the Cayley graph of G_i with respect to a finite, symmetric generating set S_i, there is a natural bijection between vertices of $\mathcal{X}$ and elements of G. This bijection is an isometry between vertices of $\mathcal{X}$ with the L^1 metric and elements of G with the word metric corresponding to the generating set:

$$S^1 = \bigcup_{1 \leq i \leq n} \{(s_1, \ldots, s_n) \mid s_j = 1 \text{ for } j \neq i \text{ and } s_i \in S_i\}$$

The same bijection is also an isometry between vertices of $\mathcal{X}$ with the L^∞ metric and elements of G with the word metric corresponding to the generating set:

$$S^\infty = \{(s_1, \ldots, s_n) \mid s_i \in S_i \cup \{1\}\}$$

COROLLARY 1.4. *For $1 \leq i \leq n$, let G_i be a non-elementary group with a finite, symmetric generating set S_i. Let $\mathcal{X}_i$ be the Cayley graph of G_i with respect to S_i, and suppose that the action of G_i on $\mathcal{X}_i$ by left multiplication has a strongly contracting element. When $n \geq 2$, the product $G = G_1 \times \cdots \times G_n$ admits a finite generating set S^1 for which the action on the corresponding Cayley graph is not growth tight and another finite generating set S^∞ for which the action on the corresponding Cayley graph is growth tight.*

Non-elementary, finitely generated, relatively hyperbolic groups, and finite rank free groups in particular, act with a strongly contracting element on any one of their Cayley graphs, so:

COROLLARY 1.5. *If $\mathbb{F}$ is a finite rank free group and S is a finite, symmetric free generating set of $\mathbb{F}$ then $\mathbb{F} \times \mathbb{F}$ is growth tight with respect to the generating set $(S \cup \{1\}) \times (S \cup \{1\})$.*

Another common way to think of $\mathbb{F} \times \mathbb{F}$ is as the Right Angled Artin Group with defining graph the join of two sets of vertices of cardinality equal to the rank of $\mathbb{F}$. The universal cover of the corresponding Salvetti complex is the product of Cayley graphs of $\mathbb{F}$ with respect to free generating sets. There are two natural metrics to consider on the vertex set of the universal cover of the Salvetti complex: the induced length metric from the piecewise Euclidean structure, which is the restriction of the L^2 metric on the product, and the induced length metric in the 1–skeleton, which is the restriction of the L^1 metric on the product.

COROLLARY 1.6. *The action of $\mathbb{F} \times \mathbb{F}$ on the universal cover of its Salvetti complex is growth tight with respect to the piecewise Euclidean metric but not growth tight with respect to the 1–skeleton metric.*

We sketch a direct proof of Corollary 1.5. The proof of Theorem 1.2 follows the same outline.

Sketch proof of Corollary 1.5. Let $\mathcal{X}$ be the Cayley graph of $\mathbb{F}$ with respect to S. Let $G = \mathbb{F} \times \mathbb{F}$ be generated by $(S \cup \{1\}) \times (S \cup \{1\})$, which induces the L^∞ metric on $\mathcal{X} \times \mathcal{X}$. We have $\delta_G = 2\delta_{\mathbb{F}} > 0$.

Let N be a non-trivial normal subgroup of G. If N has trivial projection to, say, the first factor, then $G/N = \mathbb{F} \times (\mathbb{F}/\pi_2(N))$. Since $\mathbb{F}$ is growth tight with respect to every word metric, $\delta_{\mathbb{F}/\pi_2(N)} < \delta_{\mathbb{F}}$, so $\delta_{G/N} = \delta_{\mathbb{F}} + \delta_{\mathbb{F}/\pi_2(N)} < 2\delta_{\mathbb{F}} = \delta_G$.

If N has non-trivial projection to both factors, then there is an element $(h_1, h_2) \in N$ with both coordinates non-trivial. For each $(a_1, a_2)N \in (\mathbb{F} \times \mathbb{F})/N = G/N$, choose an element $(a_1', a_2') \in (a_1, a_2)N$ such that

$$d((a_1', a_2'), (1, 1)) = d((a_1, a_2)N, (1, 1)).$$

Let $A = \{(a_1', a_2') \mid (a_1, a_2)N \in G/N\}$. We call A a *minimal section of the quotient map*. We have $\delta_{A,d} = \delta_{G/N,\bar{d}}$.

Given a non-trivial, reduced word f, let $W(f)$ be the subset of elements of $\mathbb{F}$ whose expression as a reduced word in S contains f as a subword. Denote by $\bar{a}$ the inverse of a word a in $\mathbb{F}$. If $(a_1', a_2') \in W(h_1) \times W(h_2)$ then there exist b_i and c_i such that $a_i' = b_i h_i c_i$ for $i = 1$, 2, and

$$(a_1', a_2') = (b_1 h_1 c_1, b_2 h_2 c_2) = (b_1 c_1, b_2 c_2) \cdot (\overline{c_1} h_1 c_1, \overline{c_2} h_2 c_2)$$

So $(b_1 c_1, b_2 c_2)N = (a_1', a_2')N$, but this contradicts the fact that $(a_1', a_2') \in A$, since $|(b_1 c_1, b_2 c_2)|_\infty < |(a_1', a_2')|_\infty$. Therefore, $A \subset (\mathbb{F} - W(h_1)) \times \mathbb{F} \cup \mathbb{F} \times (\mathbb{F} - W(h_2))$. However, for any non-trivial f the growth exponent of $\mathbb{F} - W(f)$ is strictly less than that of $\mathbb{F}$, so the growth exponent of A is strictly less than that of $\mathbb{F} \times \mathbb{F}$. □

The fact that the growth exponent of $F - W(f)$ is strictly less than that of F has analogues in formal language theory. A language $\mathcal{L}$ over a finite alphabet is known as 'growth-sensitive' or 'entropy-sensitive' if for every finite set of words in $\mathcal{L}$, called the *forbidden words*, the sub-language of words that do not contain one of the forbidden words as a subword has strictly smaller growth exponent than $\mathcal{L}$. It has been a topic of recent interest to decide what kinds of languages are growth-sensitive [5, 6, 10].

Our approach to growth tightness is to prove a coarse-geometric version of growth sensitivity, where the forbidden word is a power of a strongly contracting element.

The first coarse-geometric version of growth sensitivity was used by Arzhantseva and Lysenok [2] to prove growth tightness for hyperbolic groups. With Arzhantseva, [1] we gave a more general construction that applied to group actions with strongly contracting elements. The idea is that the action of a strongly contracting element closely resembles the action of an infinite order element of a hyperbolic group on a Cayley graph.

In [1] we proved a coarse-geometric version of the statement that the growth exponent of the set of reduced words in $\mathbb{F}$ that do not contain f or $\overline{f}$ as subwords is strictly less than the growth exponent of $\mathbb{F}$. For products this is not enough, since, for example, if $(f, f) \in N \trianglelefteq \mathbb{F} \times \mathbb{F}$ we cannot make the element $(f, \overline{f})$ shorter by applying powers of (f, f). We really want to forbid only positive occurrences of f in each coordinate, so we need to strengthen our coarse-geometric statement to take orientation into account.

After preliminaries in Section 2, we show in Section 3 that an infinite normal subgroup of G that has infinite projection to each factor contains an element $\underline{h}$ for which each coordinate is strongly contracting for the action of the factor group on the factor space.

In Section 4 we prove the main technical lemma, Lemma 4.7, which is our oriented growth sensitivity result.

In Section 5 we complete the proof of Theorem 1.2.

2. Preliminaries

For any group G, we use $\bar{g}$ to denote the multiplicative inverse of $g \in G$.

A group is *elementary* if it is finite or has an infinite cyclic subgroup of finite index.

A *quasi-map* $\pi\colon \mathcal{X} \to \mathcal{Y}$ between metric spaces assigns to each point $x \in \mathcal{X}$ a subset $\pi(x) \subset \mathcal{Y}$ of uniformly bounded diameter.

2.1. Strongly Contracting Elements. We define strongly contracting elements following Sisto[1] [11]. See also [1] for additional reference.

Definition 2.1. Let $(\mathcal{X}, d)$ be a proper geodesic metric space, and let $\mathcal{A} \subset \mathcal{X}$ be a subset. Given a constant $C > 0$, a map $\pi_{\mathcal{A}} \colon \mathcal{X} \to \mathcal{A}$ is called a *C–strongly contracting projection* if $\pi_{\mathcal{A}}$ satisfies the following properties:

- For every $a \in \mathcal{A}$, $d(a, \pi_{\mathcal{A}}(a)) \leq C$.
- For every $x, y \in \mathcal{X}$, if $d(\pi_{\mathcal{A}}(x), \pi_{\mathcal{A}}(y)) > C$, then for every geodesic segment $\mathcal{P}$ with endpoints x and y, we have $d(\pi_{\mathcal{A}}(x), \mathcal{P}) \leq C$ and $d(\pi_{\mathcal{A}}(y), \mathcal{P}) \leq C$.

We say the map $\pi_{\mathcal{A}}$ is a *strongly contracting projection* if it is C–strongly contracting for some $C > 0$.

Fix a base point $o \in \mathcal{X}$. Let G be a finitely-generated group that admits a proper, cocompact, and isometric action on $\mathcal{X}$.

Definition 2.2. An element $h \in G$ is *C–strongly contracting* if $i \mapsto h^i.o$ is a quasi-geodesic and if there exists $C > 0$ such that, for every geodesic segment $\mathcal{P}$ with endpoints on $\langle h \rangle.o$, there exists a C–strongly contracting projection $\pi_{\mathcal{P}} \colon \mathcal{X} \to \mathcal{P}$.

An element $h \in G$ *strongly contracting* if there exists a $C > 0$ such that h is C–strongly contracting.

The property of strongly contracting is independent of the base point o. Since the action is by isometries, a conjugate of a strongly contracting element is strongly contracting.

Let $h \in G$ be a strongly contracting element. Let $E(h) < G$ be the subgroup such that $g \in E(h)$ if and only if the Hausdorff distance between $\langle h \rangle.o$ and $g\langle h \rangle.o$ is bounded. Then $E(h)$ is *hyperbolically embedded* in the sense of Dahmani-Guirardel-Osin [7], and $E(h)$ is the unique maximal virtually cyclic subgroup containing h [7, Lemma 6.5]. Thus, $E(h)$ is the subgroup that often called the *elementarizer* or *elementary closure* of $\langle h \rangle$.

Definition 2.3. Given a strongly contracting element $h \in G$ and a point $o \in \mathcal{X}$, the set $\mathcal{H} = E(h).o$ is called a *quasi-axis* in $\mathcal{X}$ for h.

Lemma 2.4 ([1, Lemma 2.20]). *If $h \in G$ is strongly contracting, then there exists a strongly contracting projection quasi-map $\pi_{\mathcal{H}} \colon G \to \mathcal{H}$ such that $\pi_{\mathcal{H}}$ is $E(h)$–equivariant.*

Definition 2.5. If $h \in G$ is strongly contracting and $g \notin E(h)$ define $\pi_{g\mathcal{H}} \colon \mathcal{X} \to g\mathcal{H}$ by $\pi_{g\mathcal{H}}(x) = g.\pi_{\mathcal{H}}(\bar{g}.x)$.

Combining Lemma 2.4 and Definition 2.5, we may assume that the strongly contracting projection quasi-maps $\pi_{g\mathcal{H}}$ to translates of $\mathcal{H}$ are G–equivariant.

Lemma 2.6. *If $h \in G$ is C–strongly contracting there exist non-negative constants λ, ϵ, and μ such that $i \to h^i.o$ is a (λ, ϵ)–quasi-geodesic and for $0 \leq \alpha \leq \beta$ every geodesic from o to $h^\beta.o$ passes within distance μ of $h^\alpha.o$.*

Proof. There exist λ and ϵ such that $i \to h^i.o$ is a (λ, ϵ)–quasi-geodesic by definition of contracting element. Let γ be a geodesic segment from o to $h^\beta.o$. Then γ is Morse, by [11, Lemma 2.9]. Thus, there is a μ depending only on C, λ, and ϵ such that every (λ, ϵ)–quasi-geodesic segment with endpoints on γ is contained in the μ–neighborhood of γ. But $i \mapsto h^i.o$ for $i \in [0, \beta]$ is such a (λ, ϵ)–quasi-geodesic, so there is a point of γ at distance at most μ from $h^\alpha.o$. □

2.2. Actions on Quasi-trees. Let h be a contracting element for $G \curvearrowright \mathcal{X}$ as in the previous section, and let $\mathcal{H}$ be the quasi-axis of h.

In Lemma 4.7 we will consider a *free product subset*

$$Z^* * h^m = \bigcup_{i=1}^{\infty} \{z_1 h^m \cdots z_i h^m \mid z_j \in Z - \{1\}\}$$

[1] Sisto considers '$\mathcal{PS}$–contracting projections'. We use 'strong' to indicate the special case that $\mathcal{PS}$ is the collection of all geodesic segments in $\mathcal{X}$.

for a certain subset $Z \subset G$ and a sufficiently large m. We wish to know that the orbit map from G into $\mathcal{X}$ is an embedding on this free product set.

This statement recalls the following well known result:

Proposition 2.7 (Baumslag's Lemma [3]). *If $z_1, \dots, z_k$ and h are elements of a free group such that h does not commute with any of the z_i, then $z_1 h^{m_1} \cdots z_k h^{m_k} \neq 1$ if all the $|m_i|$ are sufficiently large.*

A convenient way to prove such an embedding result is to work in a tree, so that the global result, that $z_1 h^{m_1} \cdots z_k h^{m_k} \neq 1$, can be certified by a local 'no-backtracking' condition. In our situation, we do not have an action on a tree to work with, but a construction of Bestvina, Bromberg, and Fujiwara [4] produces an action of G on a quasi-tree, a space quasi-isometric to a simplicial tree, from the action of G on the G–translates of $\mathcal{H}$. In [1] we use this quasi-tree construction and a no-backtracking argument to prove that the orbit map is an embedding of a certain free product subset. The proof of Lemma 4.7 consists of choosing an appropriate free product set to which we can apply the argument from [1]. The details of the construction of the quasi-tree and the proof of the free product subset embedding are somewhat technical, so we will not repeat them here (see [4, Section 3] and [1, Section 2.4] for more details). However, we will make use of some of Bestvina, Bromberg, and Fujiwara's 'projection axioms', which hold for quasi-axes of contracting elements by work of Sisto [11], as recounted below.

Let $\mathbb{Y}$ be the collection of all distinct G–translates of $\mathcal{H}$. For each $Y \in \mathbb{Y}$, let π_Y be the projection map from the above. Set

$$d^{\pi}_Y(x, y) = \operatorname{diam}\{\pi_Y(x), \pi_Y(y)\}.$$

Lemma 2.8 ([1, Section 2.4], cf [11, Theorem 5.6]). *There exists $\xi \geqslant 0$ such that for all distinct $X, Y, Z \in \mathbb{Y}$:*

(P0) $\operatorname{diam} \pi_Y(X) \leqslant \xi$

(P1) At most one of $d^{\pi}_X(Y, Z)$, $d^{\pi}_Y(X, Z)$ and $d^{\pi}_Z(X, Y)$ is strictly greater than ξ.

(P2) $|\{V \in \mathbb{Y} \mid d^{\pi}_V(X, Y) > \xi\}| < \infty$

3. Elements that are Strongly Contracting in each Coordinate

Let G be a finitely generated, non-elementary group acting properly discontinuously and cocompactly by isometries on a based proper geodesic metric space $(\mathcal{X}, d, o)$ such that there exists an element $h \in G$ that is strongly contracting for $G \curvearrowright \mathcal{X}$. Let $\mathcal{H} = E(h).o$. Let C be the contraction constant for $\pi_{\mathcal{H}}$ from Lemma 2.4, and let ξ be the constant of Lemma 2.8. For any $x \in \mathcal{X}$ and any $r > 0$), denote by $B_r(x)$ the open ball of radius r about x.

Lemma 3.1. *Let p be a point of $\mathcal{H}$. Let g be an element of G. There exists a constant D such that either some non-trivial power of g is contained in $\langle h \rangle$ or for all $n > 0$ we have $d^{\pi}_{\mathcal{H}}(g^n.p, p) \leq 2d(p, g.p) + D$.*

Proof. Since $\langle h \rangle$ is a finite index subgroup of $E(h)$, if some non-trivial power of g is contained in $E(h)$ then some non-trivial power of g is contained in $\langle h \rangle$, and we are done. Thus, we may assume that no non-trivial power of g is contained in $E(h)$. This implies that if $m \neq n$ then $g^m\mathcal{H} \neq g^n\mathcal{H}$.

Let z be a point on a geodesic from p to $g.p$ in $B_C(\pi_{\mathcal{H}}(g.p))$. Let ξ be the constant of Lemma 2.8. Axiom (P0) of Lemma 2.8 says $\operatorname{diam} \pi_{\mathcal{H}}(g\mathcal{H}) \leq \xi$.

$$\begin{aligned} d(p, g.p) &= d(p, z) + d(z, g.p) \\ &\geq d(p, \pi_{\mathcal{H}}(g.p)) - C + d(z, g.p) \\ &\geq d^{\pi}_{\mathcal{H}}(p, g\mathcal{H}) - C - \xi + d(z, g.p) \\ &\geq d^{\pi}_{\mathcal{H}}(p, g\mathcal{H}) - C - \xi \end{aligned}$$

By a similar argument, for every $k \neq 0, \pm 1$,

$$d(p, g.p) \geq d^{\pi}_{g^k\mathcal{H}}(\mathcal{H}, g^{\pm 1}\mathcal{H}) - 2C - 2\xi$$

Using the above we obtain that, for any $n > 1$,

$$\begin{aligned} d^{\pi}_{\bar{g}\mathcal{H}}(\mathcal{H}, g^{n-1}\mathcal{H}) = d^{\pi}_{\mathcal{H}}(g\mathcal{H}, g^{n}\mathcal{H}) &\geq d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, p) - d^{\pi}_{\mathcal{H}}(g\mathcal{H}, p) \\ &\geq d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, p) - d(g.p, p) - C - \xi \end{aligned}$$

Suppose that $d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, p) - d(g.p, p) - C - \xi > \xi$. The previous inequality says $d^{\pi}_{\bar{g}\mathcal{H}}(\mathcal{H}, g^{n-1}\mathcal{H}) > \xi$, so (P1) of Lemma 2.8 implies $\xi \geq d^{\pi}_{\mathcal{H}}(\bar{g}\mathcal{H}, g^{n-1}\mathcal{H})$, hence:

$$\begin{aligned} \xi \geq d^{\pi}_{\mathcal{H}}(\bar{g}\mathcal{H}, g^{n-1}\mathcal{H}) &\geq d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, p) - d^{\pi}_{\mathcal{H}}(p, \bar{g}\mathcal{H}) - d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, g^{n-1}\mathcal{H}) \\ &\geq d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, p) - d^{\pi}_{\mathcal{H}}(p, \bar{g}\mathcal{H}) - d^{\pi}_{\bar{g}^{n}\mathcal{H}}(\mathcal{H}, \bar{g}\mathcal{H}) \\ &\geq d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, p) - 2d(g.p, p) - 3C - 3\xi \end{aligned}$$

Thus, $d^{\pi}_{\mathcal{H}}(g^{n}\mathcal{H}, p) \leq 2d(g.p, p) + D$ for $D = 3C + 4\xi$. □

LEMMA 3.2. *For every $g \in G$ there exists an $l > 0$ and an $n' \geq 0$ such that for all $m > 0$ and all $n \geq n'$, except possibly one, the elements $g^{lm}h^{n}$ and $h^{n}g^{lm}$ are strongly contracting.*

PROOF. Suppose there exists a minimal $a > 0$ and b such that $g^{a} = h^{b}$. If $b > 0$ let $l = a$ and let $n' = 0$, so that $g^{lm}h^{n} = h^{bm+n}$ is a positive power of h. If $b = 0$ let $l = a$ and $n' = 1$ so that $g^{lm}h^{n} = h^{n}$ is a positive power of h. If $b < 0$ let $l = a$, $n' = 0$, and $n \geq n'$ such that $n \neq -mb$. Then $g^{lm}h^{n}$ is a non-zero power of h.

If no non-trivial power of g is contained in $\langle h\rangle$, let $l = 1$. By Lemma 3.1, there exists a D' such that for every $p \in \mathcal{H}$ and every $m > 0$ we have $d^{\pi}_{\mathcal{H}}(g^{m}.p, p) \leq 2d(p, g.p) + D'$. Let D be the maximum of D' and the constant D from [11, Lemma 5.2]. Let $p \in \mathcal{H}$ be a point such that $d(p, g.p)$ is minimal. Let n' be large enough so that $d(p, h^{n'}.p) \geq 4d(p, g.p) + 3D$. Then for $n \geq n'$ we have $d(p, h^{n}.p) \geq d(\pi_{\mathcal{H}}(g^{m}.p), p) + d(p, \pi_{\mathcal{H}}(\bar{g}^{m}.p)) + D$. This implies $g^{lm}h^{n}$ is strongly contracting by [11, Lemma 5.2]. $h^{n}g^{lm}$ is also strongly contracting as it is conjugate to $g^{lm}h^{n}$. □

For $i = 1, \dots, n$, let G_i be a non-elementary group acting properly discontinuously and cocompactly by isometries on a proper, based, geodesic metric space $(\mathcal{X}_i, d_i, o_i)$. Assume, for each i, that $G_i \curvearrowright \mathcal{X}_i$ has a strongly contracting element. Let $G = G_1 \times \cdots \times G_n$. Let $\chi_i \colon G \to G_i$ be projection to the i–th coordinate.

LEMMA 3.3. *Let N be an infinite normal subgroup of G such that $\chi_i(N)$ is infinite for all i. There exists an element $\underline{h} = (h_1, \dots, h_n) \in N$ such that h_i is a strongly contracting element for $G_i \curvearrowright \mathcal{X}_i$.*

PROOF. $\chi_i(N)$ is an infinite normal subgroup of G_i, so it contains a strongly contracting element by [1, Proposition 3.1]. For each i, let $\underline{g}_i = (a_{i,1}, \dots, a_{i,n}) \in N$ such that $a_{i,i}$ is a strongly contracting element for $G_i \curvearrowright \mathcal{X}_i$.

We will show by induction that there is a product of the $\underline{g}_i$ that gives the desired element $\underline{h}$. The element $\underline{g}_1$ has a strongly contracting element in its first coordinate. Suppose that there is a product $\underline{f} = (f_1, \dots, f_n)$ of $\underline{g}_1, \dots, \underline{g}_i$ such that the first i coordinates are strongly contracting elements in their coordinate spaces.

For $1 \leq j \leq i$ there exists an l_j and an n'_j as in Lemma 3.2 such that for all m and all $n \geq n'_j$, except possibly one, we have $a^{l_j m}_{i+1,j} f^{n}_{j}$ is strongly contracting. Similarly, there are l_{i+1} and n'_{i+1} such that $a^{n}_{i+1,i+1} f^{l_{i+1}m}_{i+1}$ is strongly contracting for all $m > 0$ and $n \geq n'_{i+1}$.

Let l be the least common multiple of $l_1, \dots, l_i$. Let m be large enough so that $ml \geq n'_{i+1}$. Let $\lambda_k = l_{i+1}(k + \max_{j=1,\dots,i} n'_j)$, where $k \geq 0$ varies.

Consider $\underline{g}^{ml}_{i+1}\underline{f}^{\lambda_k}$. For $1 \leq j \leq i$, the j-th coordinate is strongly contracting for all except possibly one value of k, since ml is a multiple of l_j and $\lambda_k \geq n'_j$. Similarly, the $(i+1)$–st coordinate is strongly contracting for all except possibly one value of k since λ_k is a multiple of l_{i+1} and $ml \geq n'_{i+1}$. By choosing a k that is not among the at most $i+1$ forbidden values, we have that the first $i+1$ coordinates of $\underline{g}^{ml}_{i+1}\underline{f}^{\lambda_k}$ are strongly contracting in their coordinate space. □

We will say an element $g \in G_i$ has a K–long h_i–projection if there exists an $f \in G_i$ such that $d^\pi_{f\mathcal{H}_i}(o_i, g.o_i) \geq K$.

Lemma 3.4. *Given $\underline{h}$ as in Lemma 3.3, there exists an element $\underline{h'} = (h'_1, \ldots, h'_n) \in N$ such that h'_i is strongly contracting for each $G_i \curvearrowright X_i$ and there exists a K such that powers of h'_i have no K–long h_i–projections and powers of h_i have no K–long h'_i–projections.*

Proof. For each i, the group G_i is non-elementary, so there exists a $g_i \in G_i - E(h_i)$. Let $\underline{g} = (g_1, \ldots, g_n)$. [1, Proposition 3.1] shows that $\underline{h'} = \underline{g}\underline{h}^m\underline{\bar{g}}\underline{\bar{h}}^m \in N$ is strongly contracting in each coordinate for any sufficiently large m, so K can be taken to be $\max_i d^\pi_{g_i\mathcal{H}_i}(\mathcal{H}_i, g_i h_i^m \bar{g}_i \mathcal{H}_i) + 2\xi_i$, where ξ_i is chosen by Lemma 2.8. □

4. Elements without Long, Positive Projections

In the following, let G be any finitely generated, non-elementary group (not necessarily a product) acting properly discontinuously and cocompactly by isometries on a based proper geodesic metric space (X, d, o). Suppose there exists a strongly contracting element $h \in G$ for $G \curvearrowright X$. Let $\mathcal{H} = E(h).o$ and let C be the contraction constant for $\pi_\mathcal{H}$.

Let $D = \operatorname{diam}(G\backslash X)$ and let $D' = \operatorname{diam}(\langle h \rangle \backslash \mathcal{H})$.

Definition 4.1. For x_0 and x_1 in X, the ordered pair (x_0, x_1) has a *K–long, positive h–projection* if there exists a $k \in G$ such that $d^\pi_{k\mathcal{H}}(x_0, x_1) \geq K$ and $d(k.o, \pi_{k\mathcal{H}}(x_0)) \leq D'$ and there exists $\alpha > 0$ such that $d(kh^\alpha.o, \pi_{k\mathcal{H}}(x_1)) \leq D'$.

It is immediate that the property of having a K–long, positive h–projection is invariant under the G–action. We also remark that the 'positive' restriction is vacuous if $K > 2D'$ and there exists an element of G that flips the ends of $\mathcal{H}$.

Definition 4.2. Let $\hat{G}(K)$ be the elements $g \in G$ such that there exist points $x \in B_D(o)$ and $y \in B_D(g.o)$ and a geodesic γ from x to y such that no subsegment of γ has a K–long, positive h–projection.

For any $g \in G$, set $|g| = d(o, g.o)$.

Lemma 4.3. *For all sufficiently large K and for every $g \in G - \hat{G}(K)$ there exists a $k \in G$ and an interval $[\alpha', \alpha''] \subset \mathbb{Z}^+$ such that $|kh^{-\alpha}\bar{k}g| < |g|$ for all $\alpha' \leq \alpha \leq \alpha''$. The lower bound α' depends only on h and the upper bound α'' depends linearly on K.*

Proof. Let γ be a geodesic from $\gamma(0) = o$ to $\gamma(T) = g.o$. Since $g \notin \hat{G}(K)$, there are times t_0 and t_1 in $[0, T]$ such that $(\gamma(t_0), \gamma(t_1))$ has a K–long, positive h–projection. Let $k \in G$ such that $d^\pi_{k\mathcal{H}}(\gamma(t_0), \gamma(t_1)) \geq K$ and $d(\pi_{k\mathcal{H}}(\gamma(t_0)), k.o)$, and let $\beta > 0$ be such that $d(kh^\beta, \pi_{k\mathcal{H}}(\gamma(t_1))) \leq D'$.

Let λ, ϵ, and μ be the constants of Lemma 2.6 for h. Let ξ be the constant of Lemma 2.8 Since $i \mapsto h^i.o$ is (λ, ϵ)–quasi-geodesic and $d(1, h^\beta.o) \geq K - 2D' - 2\xi$ we have $\beta \geq \frac{1}{\lambda}(K - 2D' - 2\xi) - \epsilon$.

Set $\alpha'' = \beta$ and $\alpha' = \lambda(4(C + D' + \xi) + \epsilon + 2\mu + 1))$. We assume that K is large enough so that $\alpha'' \geq \alpha'$. For all $\alpha' \leq \alpha \leq \alpha''$ we have:

$$d(k.o, kh^\beta.o) \geq d(k.o, kh^\alpha.o) + d(kh^\alpha.o, kh^\beta.o) - 2\mu$$

Rearranging, and using the quasi-geodesic condition for $k\mathcal{H}$:

$$\begin{aligned} d(kh^\alpha.o, kh^\beta.o) &\leq d(k.o, kh^\beta.o) - d(k.o, kh^\alpha.o) + 2\mu \\ &\leq d(k.o, kh^\beta.o) - (\alpha/\lambda - \epsilon) + 2\mu \\ &< d(k.o, kh^\beta.o) - 4(C + D' + \xi) \end{aligned}$$

Now we use the fact that γ passes $C + D' + \xi$ close to $k.o$ and $kh^\beta.o$:

$$\begin{aligned} |g| &= d(\gamma(0), \gamma(t_0)) + d(\gamma(t_0), \gamma(t_1)) + d(\gamma(t_1), \gamma(T)) \\ &\geq d(\gamma(0), \gamma(t_0)) + d(\gamma(t_0), k.o) + d(k.o, kh^\beta.o) \\ &\quad + d(kh^\beta.o, \gamma(t_1)) + d(\gamma(t_1), \gamma(T)) - 4(C + D' + \xi) \end{aligned}$$

So:

$$\begin{aligned}
|kh^{-\alpha}\bar{k}g| &\leq d(\gamma(0), \gamma(t_0)) + d(\gamma(t_0), k.o) + d(k.o, kh^{-\alpha}\bar{k}kh^{\beta}.o) \\
&\quad + d(kh^{-\alpha}\bar{k}kh^{\beta}.o, kh^{-\alpha}\bar{k}\gamma(t_1)) + d(kh^{-\alpha}\bar{k}\gamma(t_1), kh^{-\alpha}\bar{k}\gamma(T)) \\
&= d(\gamma(0), \gamma(t_0)) + d(\gamma(t_0), k.o) + d(kh^{\alpha}.o, kh^{\beta}.o) \\
&\quad + d(kh^{\beta}.o, \gamma(t_1)) + d(\gamma(t_1), \gamma(T)) \\
&\leq |g| + 4(C + D' + \xi) - d(k.o, kh^{\beta}.o) + d(kh^{\alpha}.o, kh^{\beta}.o) \\
&< |g|
\end{aligned}$$

□

LEMMA 4.4. *Fix K and let $P(r) = \#(B_r(o) \cap \hat{G}(K).o)$. The function* $\log P(r)$ *is subadditive in r, up to bounded error.*

PROOF. Let $g.o \in B_r(o) \cap \hat{G}(K).o$. Let x, y, and γ be as in Definition 4.2. Let $m + n = r$. If $d(x, y) > m$ let z be the point on γ at distance m from x. Otherwise, let $z = y$. There exists an $f \in G$ such that $d(z, f.o) \leq D$.

We claim that f contributes to $P(m + 2D)$ and $\bar{f}g$ contributes to $P(n + 2D)$. This is because $d(o, f.o) \leq m + 2D$, and the subsegment of γ from x to z is a geodesic for f satisfying Definition 4.2. Similarly, $d(o, \bar{f}g.o) = d(f.o, g.o) \leq n + 2D$, and the subsegment of $\bar{f}.\gamma$ from $\bar{f}.z$ to $\bar{f}.y$ is a geodesic for $\bar{f}g$ satisfying Definition 4.2.

This shows that for any $m + n = r$ we have $P(r) \leq P(m + 2D) \cdot P(n + 2D)$. Applying this relation for $(m - 2D) + 4D = m + 2D$ and $(n - 2D) + 4D = n + 2D$ yields:

$$P(r) \leq (P(6D))^2 \cdot P(m) \cdot P(n)$$

Thus:

$$\log P(m + n) \leq \log P(m) + \log P(n) + 2 \log P(6D).$$

□

There is a result known as Fekete's Lemma that says if (a_i) is a subadditive sequence then $\lim_{i\to\infty} \frac{a_i}{i}$ exists and is equal to $\inf_i \frac{a_i}{i}$. We will need the following generalization for almost subadditive sequences:

LEMMA 4.5. *Let (a_i) be an unbounded, increasing sequence of positive numbers. Suppose there exists b such that $a_{m+n} \leq a_m + a_n + b$ for all m and n. Then $L = \lim_{i\to\infty} \frac{a_i}{i}$ exists and $a_i \geq Li - b$ for all i.*

PROOF. Let $L^+ = \limsup_i \frac{a_i}{i}$. Let $L^- = \liminf_i \frac{a_i}{i}$. Suppose that $L^+ > L^-$. Let $\epsilon = \frac{L^+ - L^-}{3}$. Since the sequence is increasing and unbounded, there exists an I such that for all $i > I$ we have $\frac{a_i + b}{a_i} < \sqrt{\frac{L^+ - \epsilon}{L^- + \epsilon}}$. Fix an $i > I$ such that $\frac{a_i}{i} < L^- + \epsilon$. Choose a j such that $\frac{a_j}{j} > L^+ - \epsilon$ and $\frac{q+1}{q} < \sqrt{\frac{L^+ - \epsilon}{L^- + \epsilon}}$, where $qi \leq j < (q + 1)i$.

$$L^+ - \epsilon < \frac{a_j}{j} \leq \frac{a_j}{qi} < \frac{(q + 1)(a_i + b)}{qi} < \frac{L^+ - \epsilon}{L^- + \epsilon} \cdot \frac{a_i}{i} \leq \frac{L^+ - \epsilon}{L^- + \epsilon}(L^- + \epsilon) = L^+ - \epsilon$$

This is a contradiction, so $L = L^+ = L^-$.

If for some i we have $a_i < Li - b$ then

$$L = \lim_{j\to\infty} \frac{a_{ij}}{ij} \leq \lim_{j\to\infty} \frac{j(a_i + b)}{ij} = \frac{a_i + b}{i} < L,$$

which is a contradiction. □

4.1. Divergence. For any subset $A \subset G$, define:

$$\Theta_A(s) = \sum_{r=0}^{\infty} \#(B_r(o) \cap A.o)e^{-rs}$$

The growth exponent δ_A is the *critical exponent* of Θ_A, that is, Θ_A diverges for all $s < \delta_A$ and converges for all $s > \delta_A$. We say A is *divergent* if Θ_A diverges at δ_A.

Lemma 4.6. *$\hat{G}(K)$ is divergent.*

Proof. Let $P(r) = \#(B_r(o) \cap \hat{G}(K).o)$. By Lemma 4.4 and Lemma 4.5, $\log P(r) \geq r\delta_{\hat{G}(K)} - 2\log P(6D)$ for all r. Thus:

$$\Theta_{\hat{G}(K)}(\delta_{\hat{G}(K)}) = \sum_{r=0}^{\infty} P(r)\exp(-r\delta_{\hat{G}(K)}) \geq \sum_{r=0}^{\infty} \frac{1}{P(6D)^2} = \infty. \qquad \square$$

Lemma 4.7. *For sufficiently large K, the growth exponent of $\hat{G}(K)$ is strictly smaller than the growth exponent of G.*

Proof. Let $h' \in G$ and D be the element and constant, respectively, of Lemma 3.4 (in this case the product has only one factor). Let $K > D$.

Define a map ϕ on $\hat{G}(K)$ as follows.

$$\phi(g) = \begin{cases} h'g\bar{h}' & \text{if } d^{\pi}_{\mathcal{H}}(o, g.o) \geq K \text{ and } d^{\pi}_{g\mathcal{H}}(o, g.o) \geq K \\ h'g & \text{if } d^{\pi}_{\mathcal{H}}(o, g.o) \geq K \\ g\bar{h}' & \text{if } d^{\pi}_{g\mathcal{H}}(o, g.o) \geq K \\ g & \text{otherwise} \end{cases}$$

Then for all $g \in \hat{G}(K)$ we have $d^{\pi}_{\mathcal{H}}(o, \phi(g).o) < K$ and $d^{\pi}_{\phi(g)\mathcal{H}}(o, \phi(g).o) < K$.

Let $\hat{G}'(K)$ be the image of ϕ. Then ϕ is a bijection between $\hat{G}(K)$ and $\hat{G}'(K)$, and for all $g \in \hat{G}(K)$ we have $|g| = |\phi(g)| \pm 2|h'|$. It follows that $\delta_{\hat{G}(K)} = \delta_{\hat{G}'(K)}$ and $\hat{G}'(K)$ is divergent.

Let Z be a maximal $2K$–separated subset of $\hat{G}'(K)$. Then $\delta_Z = \delta_{\hat{G}'(K)}$ and Z is divergent. For z and z' in Z, if $z\mathcal{H} = z'\mathcal{H}$ then since $d^{\pi}_{z\mathcal{H}}(o, z.o) < K$ and $d^{\pi}_{z'\mathcal{H}}(o, z'.o) < K$ we have $d(z.o, z'.o) < 2K$, so $z = z'$.

Consider the free product set

$$Z^* * h^m = \bigcup_{i=1}^{\infty} \{z_1 h^m \cdots z_i h^m \mid z_j \in Z - \{1\}\}.$$

By the same arguments as [1, Proposition 4.1], for all sufficiently large m, the orbit map is an injection of $Z^* * h^m$ into $\mathcal{X}$. This fact, together with divergence of Z, implies that $\delta_Z < \delta_G$, by [8, Criterion 2.4]. $\square$

5. Proof of the Main Theorem

Let $(\mathcal{X}_1, d_1, o_1), \ldots, (\mathcal{X}_n, d_n, o_n)$ a finite collection of proper geodesic metric spaces. Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, and let $o = (o_1, \ldots, o_n)$. Let $\underline{x} = (x_1, \ldots, x_n)$ and $\underline{y} = (y_1, \ldots, y_n)$ be any points in $\mathcal{X}$. For any $1 \leq p < \infty$, the L^p metric on $\mathcal{X}$ is defined to be:

$$d^p(\underline{x}, \underline{y}) = \left(\sum_{i=1}^{n} (d_i(x_i, y_i))^p \right)^{1/p}$$

The L^∞ metric on $\mathcal{X}$ is:

$$d^\infty(\underline{x}, \underline{y}) = \max_i d_i(x_i, y_i)$$

Proposition 5.1. *For $i = 1, \ldots, n$, let G_i be a non-elementary, finitely generated group acting properly discontinuously and cocompactly by isometries on a proper geodesic metric space $\mathcal{X}_i$. Let $G = G_1 \times \cdots \times G_n$. For each i, let A_i be a subset of G_i such that $\log P_i(r)$ is subadditive in r, up to bounded error, for $P_i(r) = \#(B_r(o_i) \cap A_i.o_i)$. Let $\delta_i = \delta_{A_i.o_i}$ be the growth exponent of A_i. For $1 \leq p \leq \infty$, the growth exponent δ_A of $A = \prod_i^n A_i$ with respect to the L^p metric on $\mathcal{X}$ is the L^q–norm of $(\delta_1, \ldots, \delta_n)$, where ${}^1/_p + {}^1/_q = 1$, and $\frac{1}{\infty}$ is understood to be 0.*

Proof. For each $g \in G_i$ let $|g|_i = d_i(o_i, g.o_i)$. For $\underline{g} = (g_1, \ldots, g_n) \in G$, let $|\underline{g}|_p = d^p(o, \underline{g}.o)$. Let B^p_r be the closed r–ball with respect to the L^p metric.

Let $P(r) = \#B^p_r(o) \cap A.o$.

Let $\mathbb{R}^n$ be equipped with the L^p norm $\|\cdot\|_p$, and let S^p_r be the vectors of norm r. Let $\phi\colon \mathbb{R}^n \to \mathbb{R}$ be the linear function $\phi(x_1,\ldots,x_n) = \sum_{i=1}^n \delta_i x_i$. For every $r > 0$ the duality of L^q and L^p implies:

$$\|(\delta_1,\ldots,\delta_n)\|_q = \|\phi\|_p = \sup_{(x_1,\ldots,x_n)\in S^p_r} \frac{|\phi(x_1,\ldots,x_n)|}{r}$$

Since $\delta_i \geq 0$ for all i, the supremum can be restricted to the positive sector of S^p_r. Furthermore, letting

$$Z^p_r = \left\{(r_1,\ldots,r_n) \mid \|(r_1,\ldots,r_n)\|_p \leq r,\ r_i \in \mathbb{N}\right\},$$

we have:

$$\|\phi\|_p = \lim_{r\to\infty} \max_{(r_1,\ldots,r_n)\in Z^p_r} \frac{\phi(r_1,\ldots,r_n)}{r}$$

Given two positive valued functions $f(r)$ and $g(r)$, we write $f(r) \sim g(r)$ if $\lim_{r\to\infty} \frac{\log f(r)}{\log g(r)} = 1$. Lemma 4.4 and Lemma 4.5 imply $P_i(r) \sim e^{\delta_i r}$ for each $i = 1,\ldots,n$, so:

$$\begin{aligned}\|\phi\|_p &= \lim_{r\to\infty} \max_{(r_1,\ldots,r_n)\in Z^p_r} \frac{\phi(r_1,\ldots,r_n)}{r}\\ &= \lim_{r\to\infty} \max_{(r_1,\ldots,r_n)\in Z^p_r} \frac{\log \prod_{i=1}^n P_i(r_i)}{r}\end{aligned}$$

For any fixed r there is $(z_{r,1},\ldots,z_{r,n}) \in Z^p_r$ such that:

$$\prod_{i=1}^n P_i(z_{r,i}) = \max_{(r_1,\ldots,r_n)\in Z^p_r} \prod_{i=1}^n P_i(r_i)$$

We also note that:

$$\prod_{i=1}^n P_i(z_{r,i}) \leq P(r) \leq \sum_{(r_1,\ldots,r_n)\in Z^p_r} \prod_{i=1}^n P_i(r_i) \leq \#Z^p_r \cdot \prod_{i=1}^n P_i(z_{r,i})$$

Since $\#Z^p_r \leq r^n$, this means $P(r) \sim \prod_i^n P_i(z_{r,i})$.
Therefore:

$$\begin{aligned}\delta_A = \limsup_{r\to\infty} \frac{\log P(r)}{r} &= \lim_{r\to\infty} \frac{\log \prod_{i=1}^n P_i(z_{r,i})}{r}\\ &= \lim_{r\to\infty} \max_{(r_1,\ldots,r_n)\in Z^p_r} \frac{\log \prod_{i=1}^n P_i(r_i)}{r}\\ &= \|\phi\|_p = \|(\delta_1,\ldots,\delta_n)\|_q\end{aligned}$$

□

Proof of Theorem 1.2. The existence of a strongly contracting element implies that each factor group has strictly positive growth exponent, and the main theorem of [1] says that $G_i \curvearrowright \mathcal{X}_i$ is growth tight, so we are done if $n = 1$.

Assume $n > 1$ and let $1 \leq q \leq \infty$ be such that $1/p + 1/q = 1$. If $p = 1$, then by Proposition 5.1 the growth exponent of G is the maximum of the growth exponents of the G_i. Thus, we may kill the slowest growing factor without changing the growth exponent, and the action of G on $\mathcal{X}$ with the L^1 metric is not growth tight.

Now assume $p > 1$. Let $\chi_i\colon G \to G_i$ be projection to the i–th coordinate. Let N be an infinite normal subgroup of G.

First we assume that $\chi_i(N)$ is infinite for all i.

By Lemma 3.3, there exists an element $\underline{h} = (h_1,\ldots,h_n) \in N$ such that h_i is a strongly contracting element for $G_i \curvearrowright X_i$ for each i.

Let A be a *minimal section* of the quotient map $G \to G/N$. That is, A consists of a representative for each coset gN and $d(o,\underline{a}.o) = d(N.o,\underline{a}N.o)$ for all $\underline{a} \in A$, where d is the L^p metric on $\mathcal{X}$.

Proposition 5.2. *For all sufficiently large K and for all $\underline{a} = (a_1,\ldots,a_n) \in A$ there exists an index $1 \leq i \leq n$ such that $a_i \in \hat{G}_i(K)$.*

Proof. For each i, let $\hat{G}_i(K)$ be as in Definition 4.2 for each G_i. Assume K is greater than the constants K from Lemma 4.7 and Lemma 4.3 applied to each G_i.

Suppose $\underline{a}$ is such that for all i we have $a_i \in G_i - \hat{G}_i(K)$. For each i, let $k_i \in G_i$ and $[\alpha_i', \alpha_i'']$ be the k and interval, respectively, from Lemma 4.3 applied to a_i. The α_i' depend only on their respective h_i, while the α_i'' depend linearly on K. By choosing K large enough, we may choose α such that $\max_i \alpha_i' \leq \alpha \leq \min_i \alpha_i''$, so that $\alpha \in [\alpha_i', \alpha_i'']$ for all i. Let $\underline{k} = (k_1, \dots, k_n)$. The i–th coordinate of $\underline{k}\underline{\bar{h}}^{\alpha}\underline{\bar{k}}\underline{a}$ is $k_i\bar{h}_i^{\alpha}\bar{k}_i a_i$, which is shorter than a_i by Lemma 4.3. But this means that $\underline{k}\underline{\bar{h}}^{\alpha}\underline{\bar{k}}\underline{a}$ is shorter than $\underline{a}$. This contradicts the fact that $\underline{a}$ belongs to a minimal section, since $\underline{k}\underline{\bar{h}}^{\alpha}\underline{\bar{k}}\underline{a} = \underline{a}(\underline{\bar{a}}\underline{k}\underline{\bar{h}}^{\alpha}\underline{\bar{k}}\underline{a}) \in \underline{a}N$. □

Continuing the proof of Theorem 1.2, by Proposition 5.2,

$$A \subset \bigcup_{i=1}^{n} G_1 \times \cdots \times \hat{G}_i \times \cdots \times G_n,$$

where $\hat{G}_i = \hat{G}_i(K)$ for some sufficiently large K. By Proposition 5.1, the growth exponent of $G_1 \times \cdots \times \hat{G}_i \times \cdots \times G_n$ is $\|(\delta_1, \dots, \hat{\delta}_i, \dots, \delta_n)\|_q$, where δ_i is the growth exponent of G_i and $\hat{\delta}_i$ is the growth exponent of $\hat{G}_i$. Thus, the growth exponent of A is $\max_i \|(\delta_1, \dots, \hat{\delta}_i, \dots, \delta_n)\|_q$. By Lemma 4.7, $\hat{\delta}_i < \delta_i$ for each i, so, since $q < \infty$:

$$\delta_{G/N} = \delta_A = \max_i \|(\delta_1, \dots, \hat{\delta}_i, \dots, \delta_n)\|_q < \|(\delta_1, \dots, \delta_n)\|_q = \delta_G$$

It remains to consider the case that some $\chi_i(N)$ is finite. By reordering, if necessary, we may assume $\chi_i(N)$ is finite for $i \leq m$ and infinite for $i > m$. Since N is infinite, $m < n$. Let $G^1 = G_1 \times \cdots \times G_m$ with $\chi^1 = \chi_1 \times \cdots \times \chi_m \colon G \to G^1$. Let $G^\infty = G_{m+1} \times \cdots \times G_n$ with $\chi^\infty = \chi_{m+1} \times \cdots \times \chi_n \colon G \to G^\infty$.

Now $\ker(\chi^1) \cap N$ is a finite index subgroup of N that is normal in G, so G/N is a quotient of $G/(\ker(\chi^1) \cap N)$ by a finite group, and they have the same growth rates. Replacing N with $\ker(\chi^1) \cap N$, we can assume that $\chi_i(N)$ is trivial for $1 \leq i \leq m$ and infinite for $m < i \leq n$. The theorem applied to G^∞ shows that $\delta_{G^\infty/\chi^\infty(N)} < \delta_{G^\infty}$, so, since $q < \infty$:

$$\delta_{G/N} = \|(\delta_{G^1}, \delta_{G^\infty/\chi^\infty(N)})\|_q < \|(\delta_{G^1}, \delta_{G^\infty})\|_q = \delta_G$$ □

In the case that the normal subgroup has infinite projection to each factor, our proof uses the existence of a contracting element in each factor in an essential way. One wonders if the theorem is still true without this hypothesis:

Question. If, for $1 \leq i \leq n$, G_i is a non-elementary, finitely generated group acting properly discontinuously and cocompactly by isometries on a proper geodesic metric space $\mathcal{X}_i$, and if, for all i, $G_i \curvearrowright \mathcal{X}_i$ is growth tight, is it still true that the product group is growth tight with respect to the action on the product space with the L^p metric for some/all $p > 1$?

References

[1] Goulnara N. Arzhantseva, Christopher H. Cashen, and Jing Tao. *Growth Tight Actions*. Pacific J. Math. (in press). arXiv: 1401.0499v3.

[2] Goulnara N. Arzhantseva and Igor G. Lysenok. *Growth tightness for word hyperbolic groups*. Math. Z. **241**.3 (2002), pp. 597–611.

[3] Gilbert Baumslag. *On generalised free products*. Math. Z. **78** (1962), pp. 423–438.

[4] Mladen Bestvina, Kenneth Bromberg, and Koji Fujiwara. *Constructing group actions on quasi-trees and applications to mapping class groups*. Publ. Math. Inst. Hautes Études Sci., in press. 2010. arXiv: 1006.1939v5.

[5] Tullio Ceccherini-Silberstein and Wolfgang Woess. *Growth and ergodicity of context-free languages*. Trans. Amer. Math. Soc. **354**.11 (2002), pp. 4597–4625.

[6] Tullio Ceccherini-Silberstein and Wolfgang Woess. *Growth-sensitivity of context-free languages*. Theoret. Comput. Sci. **307**.1 (2003). Words, pp. 103–116.

[7] François Dahmani, Vincent Guirardel, and Denis Osin. *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*. preprint. 2011. arXiv: 1111.7048v5.

[8] Françoise Dal'Bo, Marc Peigné, Jean-Claude Picaud, and Andrea Sambusetti. *On the growth of quotients of Kleinian groups*. Ergodic Theory Dynam. Systems **31**.3 (2011), pp. 835–851.

[9] Rostislav Grigorchuk and Pierre de la Harpe. *On problems related to growth, entropy, and spectrum in group theory*. J. Dynam. Control Systems **3**.1 (1997), pp. 51–89.

[10] Wilfried Huss, Ecaterina Sava, and Wolfgang Woess. *Entropy sensitivity of languages defined by infinite automata, via Markov chains with forbidden transitions*. Theoret. Comput. Sci. **411**.44-46 (2010), pp. 3917–3922.

[11] Alessandro Sisto. *Contracting elements and random walks*. preprint. 2011. arXiv: 1112.2666v2.

PAPER F

Cogrowth for group actions with strongly contracting elements

Goulnara N. Arzhantseva and Christopher H. Cashen
Ergodic Theory & Dynamical Systems (in press)
accepted Oct 15, 2018
DOI:10.1017/etds.2018.123

Let G be a group acting properly by isometries and with a strongly contracting element on a geodesic metric space. Let N be an infinite normal subgroup of G, and let δ_N and δ_G be the growth rates of N and G with respect to the pseudo-metric induced by the action. We prove that if G has purely exponential growth with respect to the pseudo-metric then $\delta_N/\delta_G > 1/2$. Our result applies to suitable actions of hyperbolic groups, right-angled Artin groups and other CAT(0) groups, mapping class groups, snowflake groups, small cancellation groups, etc. This extends Grigorchuk's original result on free groups with respect to a word metrics and a recent result of Jaerisch, Matsuzaki, and Yabuki on groups acting on hyperbolic spaces to a much wider class of groups acting on spaces that are not necessarily hyperbolic.

1. Introduction

We consider the exponential *growth rate* δ_G of the orbit of a group G acting properly on a geodesic metric space X. In various notable contexts this asymptotic invariant is related to the Hausdorff dimension of the limit set of G in ∂X and to analytical and dynamical properties of $G\backslash X$ such as the spectrum of the Laplacian, divergence rates of random walks, volume entropy, and ergodicity of the geodesic flow.

In some cases of special interest, the value of half the growth rate of the ambient space X is distinguished. For example, when $X = \mathbb{H}^n$ and H is a torsion free discrete group of isometries of X, the Elstrodt-Patterson-Sullivan formula [24] for the bottom of the spectrum of the Laplacian of $H\backslash X$ has a phase change when the ratio of δ_H to the volume entropy of X is $1/2$. Similarly, if X is a Cayley tree of a finite rank free group F_n and H is a subgroup, then the Grigorchuk cogrowth formula [13] for the spectral radius of $H\backslash X$ has a phase change at $\delta_H/\delta_{F_n} = 1/2$. Our main result says that, in great generality, normal subgroups land decisively on one side of this distinguished value:

Theorem 1.1. *Suppose G is a group acting properly by isometries on a geodesic metric space X with a strongly contracting element and with purely exponential growth. If N is an infinite normal subgroup of G then $\delta_N/\delta_G > 1/2$, where the growth rates δ_G and δ_N are computed with respect to $G \curvearrowright X$.*

The ratio δ_N/δ_G is known as the *cogrowth* of $Q := G/N$. The hypotheses will be explained in detail in the next section. Briefly, the existence of a strongly contracting element means that some element of G acts hyperbolically on X, though X itself need not be hyperbolic, and pure exponential growth is guaranteed if the action has a strongly contracting element and an orbit of G in X is not too badly distorted.

This research was partially supported by the European Research Council (ERC) grant of Goulnara Arzhantseva, “ANALYTIC” grant agreement no. 259527. The second author is supported by the Austrian Science Fund (FWF): P 30487-N35.

In negative curvature, the strict lower bound on cogrowth has been shown in various special cases [23, 21, 5, 16]. For $X = G = F_n$, the strict lower bound on cogrowth is due to Grigorchuk [13].

Grigorchuk and de la Harpe [14, page 69] (see also [15, Problem 36]) asked whether the strict lower cogrowth bound also holds when F_n is replaced by a non-elementary Gromov hyperbolic group, and X is one of its Cayley graphs. This long-open problem was recently answered affirmatively by Jaerisch, Matsuzaki, and Yabuki [19] (see also a survey by Matsuzaki [18]). Their result applies more generally to groups of divergence type acting on hyperbolic spaces. Theorem 1.1 gives an alternative proof of the positive answer to Grigorchuk and de la Harpe's question, and goes much beyond. In comparison, Jaerisch, Matsuzaki, and Yabuki's result applies to more general actions if one restricts to actions on *hyperbolic spaces*, while Theorem 1.1 applies to many renowned non-hyperbolic examples.

Corollary 1.2. *For the following $G \curvearrowright X$, for every infinite normal subgroup N of G we have $\delta_N/\delta_G > 1/2$.*

(1) G is a non-elementary hyperbolic group acting cocompactly on a hyperbolic space X.
(2) G is a relatively hyperbolic group, and X is hyperbolic such that $G \curvearrowright X$ is cusp uniform and satisfies the parabolic gap condition.
(3) G is a right-angled Artin group defined by a finite simple graph that is neither a single vertex nor a join, and X is the universal cover of its Salvetti complex.
(4) X is a CAT(0) space, and G acts cocompactly with a rank 1 isometry on X.
(5) G is the mapping class group of a surface of genus g and p punctures, with $6g-6+2p \geqslant 2$, and X is the Teichmüller space of the surface with the Teichmüller metric.

Results (3)-(5) are new, only known as consequences of Theorem 1.1. Further new examples include wide classes of snowflake groups [2] and of infinitely presented graphical and classical small cancellation groups [1], hence, many so-called infinite 'monster' groups.

The generality of Theorem 1.1 is striking. Previous successes in showing the strict lower bound on cogrowth have relied on fairly sophisticated results concerning Patterson-Sullivan measures on the boundary of a hyperbolic space or ergodicity of the geodesic flow on $G\backslash X$. These tools are not available in our general setting. Instead, we use the geometry of the group action directly to estimate orbit growth. The idea of our argument is as follows.

(1) If G contains a strongly contracting element for $G \curvearrowright X$ then so does every infinite normal subgroup N of G. Let $c \in N$ be such an element.
(2) By passing to a high power of c, if necessary, we may assume that its translation length is much larger than the constants describing its strong contraction properties. In this case the growth $\delta_{[c]}$ of the set $[c]$ of conjugates of c is exactly $\delta_G/2$.
(3) A 'tree's worth' of copies of $[c]$ injects into the normal closure $\langle\langle c\rangle\rangle$ of c, which is a subgroup of N. It follows that the growth rate of $\langle\langle c\rangle\rangle$, hence of N, is strictly greater than $\delta_{[c]} = \delta_G/2$. In this step we use the 'hyperbolicity' of the action of c, as quantified by strong contraction, to provide geometric separation between copies of $[c]$.

We used this strategy in our paper with Tao [2] (see also references therein) to prove *growth tightness* of $G \curvearrowright X$ for actions having a strongly contracting element. The key point was to estimate the growth rate of the quotient of G by the normal closure of c. We chose a section A of the quotient map and built a tree's worth of copies of it by translating by a high power of c. By construction, the set A did not contain words containing high powers of c as subwords, so translates of A by powers of c were geometrically separated. There is a serious difficulty in applying step (3) for cogrowth, because $[c]$ *does* contain words with arbitrarily large powers of c as subwords. Indeed, any word of G can occur as a subword of an element of $[c]$, so we do not get the same nice geometric separation as hoped for in step (3), and consequently our abstract tree's worth of copies of $[c]$ does not inject into G. We overcome this difficulty by quantifying how this mapping fails to be an injection. We show there is asymptotically at least half of $[c]$ for which the map is an injection, and we use this half of $[c]$ to complete step (3).

For an example where the conclusion of the theorem does not hold, consider the group $G = F_2 \times F_2$ acting on its Cayley graph X with respect to the generating set $(S \cup 1) \times (S \cup 1)$, where S is a free generating set of F_2. The F_2 factors are normal and have growth rate exactly half the growth rate of G. The action $G \curvearrowright X$ does not have a strongly contracting element.

We thank the referee for the careful reading and helpful comments.

2. Preliminaries

We write $x \overset{*}{\prec} y$, $x \overset{+}{\prec} y$, or $x \prec y$ if there is a universal constant $C > 0$ such that $x < Cy$, $x < y + C$, or $x < Cy + C$, respectively. We define $\overset{*}{\succ}$, $\overset{+}{\succ}$, $\succ$, $\overset{*}{\asymp}$, $\overset{+}{\asymp}$, and $\asymp$ similarly.

Throughout, we let (X, d, o) be a based geodesic metric space and let G be a group acting isometrically on X. For $Y \subset X$ and $r \geqslant 0$, let $B_r(Y) := \{x \in X \mid \exists y \in Y, d(x, y) < r\}$ and $\bar{B}_r(Y) := \{x \in X \mid \exists y \in Y, d(x, y) \leqslant r\}$. Let $B_r := B_r(o)$, and let $S_r^\Delta := B_{r+\Delta} - B_r$.

There are induced pseudo-metric and semi-norm on G given by $d(g, h) := d(g.o, h.o)$ and $|g| := d(o, g.o)$.

2.1. Growth. The *(exponential) growth rate* of a subset $Y \subset X$ is:

$$\delta_Y := \limsup_{r\to\infty} \frac{\log \#Y \cap \bar{B}_r}{r}$$

The *Poincaré series* of a countable subset Y of X is:

$$\Theta_Y(s) := \sum_{y \in Y} \exp(-sd(y, o))$$

For any $\Delta > 0$ we also consider the series:

$$\Theta_Y^{S,\Delta}(s) := \sum_{i=0}^{\infty} (\#Y \cap S_{\Delta i}^{\Delta(i+1)}) \exp(-s\Delta i)$$

$$\Theta_Y^{B,\Delta}(s) := \sum_{i=0}^{\infty} (\#Y \cap \bar{B}_{\Delta i}) \exp(-s\Delta i)$$

The series $\Theta_Y^{B,\Delta}(s)$ and $\Theta_Y^{S,\Delta}(s)$ agree with $\Theta_Y(s)$ up to multiplicative error depending on Δ and s, so they all converge and diverge together. Now, $\Theta_Y(s)$ converges for $s > \delta_Y$ and diverges for $s < \delta_Y$. The set Y is said to be *divergent*, or *of divergent type*, if $\Theta_Y(s)$ diverges at $s = \delta_Y$.

We say that $Y \subset X$ has *purely exponential growth* if there exist $\delta > 0$ and $\Delta > 0$ such that $\#Y \cap S_r^\Delta \overset{*}{\asymp} \exp(\delta r)$. Recall this means there is a constant $C > 0$, independent of r, such that $\exp(\delta r)/C \leqslant \#Y \cap S_r^\Delta \leqslant C \exp(\delta r)$.

An action $G \curvearrowright X$ is *(metrically) proper* if for all $x \in X$ and $r \geqslant 0$ the set $\{g \in G \mid d(x, g.o) \leqslant r\}$ is finite. When $G \curvearrowright X$ is proper we extend all the preceding definitions to subsets H of G by taking $Y = H.o$, e.g.:

$$\delta_H := \limsup_{r\to\infty} \frac{\log \#H.o \cap \bar{B}_r}{r} = \limsup_{r\to\infty} \frac{\log \#\{h \in H \mid |h| \leqslant r\}}{r}$$

When $G \curvearrowright X$ is cocompact, or, more generally, has a quasi-convex orbit, the growth of $\#S_r^\Delta \cap G.o$ is coarsely sub-multiplicative, which, when $\delta_G > 0$, implies an exponential lower bound on $\#S_r^\Delta \cap G.o$. Conversely, if $G \curvearrowright X$ contains a strongly contracting element then the growth of $\#S_r^\Delta \cap G.o$ is coarsely super-multiplicative, which implies the corresponding exponential upper bound. For instance, Coornaert [9] proved that a quasi-convex-cocompact, exponentially growing subgroup of a hyperbolic group has purely exponential growth. More generally, in [2] we introduced the following condition that implies the pseudo-metric induced by a group action behaves like a word metric for growth purposes: the *complementary growth* of $G \curvearrowright X$ is the growth rate of the set of points of $G.o$ that can be reached from o by a geodesic segment in X that stays completely outside of a neighborhood of $G.o$, except near its endpoints. We say that $G \curvearrowright X$ has *complementary growth gap* if the complementary growth is strictly less than δ_G. Yang [25] proved that if G acts properly with a strongly contracting element and $0 < \delta_G < \infty$ then complementary growth gap implies purely exponential growth.

For relatively hyperbolic groups the complementary growth gap specializes to the *parabolic growth gap* of [11], which requires that the growth of parabolic subgroups of a relatively hyperbolic group is strictly less than the growth rate of the whole group. For another non-cocompact example, we showed in [2] that the action of the mapping class group of a hyperbolic surface on its Teichmüller space has complementary growth gap.

For a non-example, consider the integers $\mathbb{Z}$ acting parabolically on the hyperbolic plane. Hyperbolic geodesics connecting o to $n.o$ for large n travel deeply into a horoball at the fixed point of $\mathbb{Z}$ on $\partial\mathbb{H}^2$, far from the orbit of $\mathbb{Z}$. Although $\mathbb{Z}$ has 0 exponential growth in any word metric, in terms of this action on $\mathbb{H}^2$ it has exponential growth due entirely to the distortion of the orbit.

2.2. Contraction. A subset Y of X is *C–strongly contracting*, for a 'contraction constant' $C \geqslant 0$, if for all $x, x' \in X$, if $d(x, x') \leqslant d(x, Y)$ then the diameter of $\pi_Y(x) \cup \pi_Y(x')$ is at most C, where $\pi_Y(x) := \{y \in Y \mid d(x, y) = d(x, Y)\}$. A set is called *strongly contracting* if there exists a $C \geqslant 0$ such that it is C–strongly contracting. The *projection distance in Y* is $d^\pi_Y(x, x') := \operatorname{diam} \pi_Y(x) \cup \pi_Y(x')$. We extend these definitions to sets $Z \subset X$ by $\pi_Y(Z) := \cup_{z \in Z} \pi_Y(z)$ and $d^\pi_Y(Z, Z') := \operatorname{diam} \pi_Y(Z) \cup \pi_Y(Z')$.

Strong contraction of Y is equivalent [2, Lemma 2.4] to the *bounded geodesic image property*: For all $C \geqslant 0$ there exists $C' \geqslant C$ such that if Y is C–strongly contracting then for every geodesic γ in X, if $\gamma \cap B_{C'}(Y) = \emptyset$ then $\operatorname{diam} \pi_Y(\gamma) \leqslant C'$.

Corollary 2.1. *Suppose Y is C–strongly contracting and C' is as above. Suppose γ is a geodesic defined on an interval $[a, b]$, possibly infinite. Let $t_0 := \inf\{t \mid d(\gamma(t), Y) < C'\}$, and let $t_1 := \sup\{t \mid d(\gamma(t), Y) < C'\}$. Then $\operatorname{diam} \pi_Y(\gamma([a, t_0])) \leqslant C'$ and $\operatorname{diam} \pi_Y(\gamma([t_1, b])) \leqslant C'$, while $\gamma([t_0, t_1]) \subset \bar{B}_{3C'}(Y)$. If a and b are finite and $\operatorname{diam} \pi_Y(\gamma(a)) \cup \pi_Y(\gamma(b)) > C'$ then $\pi_Y(\gamma(a)) \subset \bar{B}_{2C'}(\gamma(t_0))$ and $\pi_Y(\gamma(b)) \subset \bar{B}_{2C'}(\gamma(t_1))$.*

An infinite order element $c \in G$ is said to be a *strongly contracting element* for $G \curvearrowright X$ if the set $\langle c\rangle.o$ is strongly contracting. In this case $\mathbb{Z} \to X : i \mapsto c^i.o$ is a quasi-isometric embedding and c is contained in a maximal virtually cyclic subgroup $E(c)$. This subgroup, which is alternately known as the *elementarizer* or *elementary closure* of c, can also be characterized as the maximal subgroup consisting of elements $g \in G$ such that $g^{-1}\langle c\rangle g$ is at bounded Hausdorff distance from $\langle c\rangle$. Since $E(c).o$ is coarsely equivalent to $\langle c\rangle.o$, the set $E(c).o$ is also strongly contracting. Note that $E(c) = E(c^n)$ for every $n \neq 0$. Thus, when considering $E(c).o$, we can pass to powers of c freely without changing the set $E(c^n).o$, and in particular without changing its contraction constant.

For a strongly contracting element c, let $\mathcal{E} := E(c).o$, and let $\mathbf{Y}$ be the collection of distinct G–translates of $\mathcal{E}$. Bestvina, Bromberg, and Fujiwara [4] axiomatized the geometry of projection distances in $\mathbf{Y}$. With Sisto [3] they showed that by a small change in the projections and projection distances, a cleaner set of axioms is satisfied—these will allow us to make an inductive argument in the next section. The following is [3, Theorem 4.1] applied to $\mathbf{Y}$. We list here only those axioms that we will make use of and that are not immediate from our particular definitions of $\mathbf{Y}$, $\pi_{\mathcal{Y}}$, and $d^\pi_{\mathcal{Y}}$. A detailed verification that $\mathbf{Y}$ satisfies the hypotheses of [3, Theorem 4.1] can be found in [2].

Theorem 2.2. *There exists $\theta \geqslant 0$ such that for each $\mathcal{Y} \in \mathbf{Y}$ there is a projection $\pi'_{\mathcal{Y}}$ taking elements of $\mathbf{Y}$ to subsets of $\mathcal{Y}$ such that for all $\mathcal{X} \in \mathbf{Y}$ and $g \in G$ we have $\pi'_{\mathcal{Y}}(\mathcal{X}) \subset B_\theta(\pi_{\mathcal{Y}}(\mathcal{X}))$ and $\pi'_{g\mathcal{Y}}(g\mathcal{X}) = g\pi'_{\mathcal{Y}}(\mathcal{X})$. Furthermore, there are distance maps $d_{\mathcal{Y}}(\mathcal{X}, \mathcal{Z}) = \operatorname{diam} \pi'_{\mathcal{Y}}(\mathcal{X}) \cup \pi'_{\mathcal{Y}}(\mathcal{Z})$ with $|d_{\mathcal{Y}} - d^\pi_{\mathcal{Y}}| \leqslant 2\theta$ such that, for $\theta' := 11\theta$, the following axioms are satisfied for all $\mathcal{X}$, $\mathcal{Y}$, $\mathcal{Z}$, $\mathcal{W} \in \mathbf{Y}$:*

(P 0): *$d^\pi_{\mathcal{Y}}(\mathcal{X}, \mathcal{X}) \leqslant \theta$ when $\mathcal{X} \neq \mathcal{Y}$.*

(P 1): *If $d^\pi_{\mathcal{Y}}(\mathcal{X}, \mathcal{Z}) > \theta$ then $d^\pi_{\mathcal{X}}(\mathcal{Y}, \mathcal{Z}) \leqslant \theta$ for all distinct $\mathcal{X}$, $\mathcal{Y}$, $\mathcal{Z}$.*

(SP 3): *If $d_{\mathcal{Y}}(\mathcal{X}, \mathcal{Z}) > \theta'$ then $d_{\mathcal{Z}}(\mathcal{X}, \mathcal{W}) = d_{\mathcal{Z}}(\mathcal{Y}, \mathcal{W})$ for all $\mathcal{W} \in \mathbf{Y} - \{\mathcal{Z}\}$.*

(SP 4): $d_{\mathcal{Y}}(\mathcal{X}, \mathcal{X}) \leqslant \theta'$ *when* $\mathcal{X} \neq \mathcal{Y}$.

For more details on strongly contracting elements and many examples, see [2].

Proposition 2.3 ([3, Lemma 2.2 and Proposition 2.3]). *With* θ' *as in Theorem 2.2, for each* $\mathcal{X}$ *and* $\mathcal{Z}$ *in* $\mathbf{Y}$ *define* $\mathbf{Y}(\mathcal{X}, \mathcal{Z}) := \{\mathcal{Y} \in \mathbf{Y} - \{\mathcal{X}, \mathcal{Z}\} \mid d_{\mathcal{Y}}(\mathcal{X}, \mathcal{Z}) > 2\theta'\}$ *and* $\mathbf{Y}[\mathcal{X}, \mathcal{Z}] := \mathbf{Y}(\mathcal{X}, \mathcal{Z}) \cup \{\mathcal{X}, \mathcal{Z}\}$. *There is a total order* $\sqsubset$ *on* $\mathbf{Y}[\mathcal{X}, \mathcal{Z}]$ *such if* $\mathcal{Y}_0 \sqsubset \mathcal{Y}_1 \sqsubset \mathcal{Y}_2$ *then* $d_{\mathcal{Y}_1}(\mathcal{Y}_0, \mathcal{Y}_2) = d_{\mathcal{Y}_1}(\mathcal{X}, \mathcal{Z})$. *The relation* $\mathcal{Y}_0 \sqsubset \mathcal{Y}_1$ *is defined by each of the following equivalent conditions:*

- $d_{\mathcal{Y}_0}(\mathcal{X}, \mathcal{Y}_1) > \theta'$
- $d_{\mathcal{Y}_1}(\mathcal{X}, \mathcal{Y}_0) \leqslant \theta'$
- $d_{\mathcal{Y}_1}(\mathcal{Y}_0, \mathcal{Z}) > \theta'$
- $d_{\mathcal{Y}_0}(\mathcal{Y}_1, \mathcal{Z}) \leqslant \theta'$

3. Embedding a tree's worth of copies of [c]

For a subset $H \subset G$, let $H^* := H - \{1\}$, and consider $\hat{H} := \bigcup_{k=1}^{\infty} (H^*)^k$. We consider $\hat{H}$ to be a *'tree's worth of copies of H'* in allusion to the case of the free product $H * \mathbb{Z}/2\mathbb{Z}$ when H is a group. The group $H * \mathbb{Z}/2\mathbb{Z}$ acts on a tree with vertex stabilizers conjugate to H, and every element that is not equal to 1 or the generator z of $\mathbb{Z}/2\mathbb{Z}$ has a unique expression as $z^{\alpha} h_1 z h_2 z \cdots h_k z^{\beta}$ for some $k \in \mathbb{N}$, $\alpha, \beta \in \{0, 1\}$, and $h_i \in H^*$.

The naïve map $\hat{H} \to X : (h_1, \ldots, h_k) \mapsto h_1 c \cdots h_k c.o$, where c is a strongly contracting element, is clearly not an injection for $H = [c]$, as it gives collisions $(h^{-1}, h) \mapsto h^{-1}chc.o \leftarrow\!\shortmid (h^{-1}ch)$. To avoid collisions we remove a fraction of [c] in four steps, and use a slightly different map. The main technical result is:

Proposition 3.1. *Under the hypothesis of Theorem 1.1, let c be a strongly contracting element. After possibly passing to a power of c, there is a subset* $G_4 \subset [c]$ *that is divergent, has* $\delta_{G_4} = \delta_G/2$, *and for which the map* $\hat{G}_4 \to X : (g_1, \ldots, g_k) \mapsto (\prod_{i=1}^{k} g_i c^2).o$ *is an injection.*

The main theorem follows by an argument analogous to the one we used in [2], which we reproduce for the reader's convenience.

Proof of Theorem 1.1. Let $c' \in G$ be a strongly contracting element for $G \curvearrowright X$. Suppose that $N < E(c')$. Since N is infinite, it has a finite index subgroup in common with $\langle c' \rangle$. But conjugation by an element of G fixes N, so it moves $\langle c' \rangle$ by a bounded Hausdorff distance, which means $G = E(c')$ is virtually cyclic and N is a finite index subgroup of G. However, $\langle c' \rangle$ has an undistorted orbit in X. Since this is a finite index subgroup of G, the growth of G is only linear, contradicting the exponential growth hypothesis. Thus, we may assume that G is not virtually cyclic and that N contains an element g that is not in $E(c')$. We showed in [2, Proposition 3.1] that for sufficiently large n the element $c := g^{-1}(c')^{-n} g (c')^n$ is a strongly contracting element of N.

Consider G_4 as provided by Proposition 3.1 with respect to c. Then $\hat{G}_4$ injects into X, and, moreover, the image is contained in $\langle\langle c \rangle\rangle.o \subset N.o$. Therefore, the growth rate of N is at least as large as the growth rate of the image of $\hat{G}_4$, which we estimate using its Poincaré series:

$$\begin{aligned}\Theta_{\hat{G}_4}(s) &= \sum_{k=1}^{\infty} \sum_{(g_1,\dots,g_k)\in(G_4^*)^k} \exp(-s|g_1c^2\cdots g_kc^2|) \\ &\geqslant \sum_{k=1}^{\infty} \sum_{(g_1,\dots,g_k)\in(G_4^*)^k} \exp\left(-sk|c^2| - s\sum_{i=1}^{k}|g_i|\right) \\ &= \sum_{k=1}^{\infty} \exp(-sk|c^2|) \sum_{(g_1,\dots,g_k)\in(G_4^*)^k} \prod_{i=1}^{k} \exp(-s|g_i|) \\ &= \sum_{k=1}^{\infty} \exp(-sk|c^2|) \left(\sum_{g\in G_4^*} \exp(-s|g|)\right)^k \\ &= \sum_{k=1}^{\infty} \left(\exp(-s|c^2|)\Theta_{G_4^*}(s)\right)^k\end{aligned}$$

Since G_4 is divergent, for sufficiently small positive ϵ we have $\Theta_{G_4^*}(\delta_{G_4}+\epsilon) \geqslant \exp((\delta_{G_4}+\epsilon)|c^2|)$, so $\Theta_{\hat{G}_4}(\delta_{G_4}+\epsilon)$ diverges, which implies $\delta_{\hat{G}_4} \geqslant \delta_{G_4}+\epsilon$. Thus, $\delta_N \geqslant \delta_{\hat{G}_4} \geqslant \delta_{G_4}+\epsilon > \delta_{G_4} = \delta_G/2$. □

The remainder of this section is devoted to the construction of the set G_4 satisfying the conclusion of Proposition 3.1. Here is a brief overview. We need a subset of $[c]$ such that the given map is an injection. It would be preferable if we could take conjugates of c by elements g that have no long projection to any element of $\mathbf{Y}$. It is easy to build an injection based on such elements, but, unfortunately, there are too few of them in our setting—the growth rate of the set of such elements is strictly smaller than δ_G, so the growth rate of c–conjugates by such elements is strictly smaller than $\delta_G/2$. Instead, we consider elements g that do not have long projections to $\mathcal{E}$ and $g\mathcal{E}$; in a sense, these are elements 'orthogonal to $\mathbf{Y}$ at their endpoints', rather than 'orthogonal to $\mathbf{Y}$' throughout. The desired condition can be achieved with a small modification near the ends of g, so this does not change the growth rate. We call this set of elements G_1 and the conjugates of (a power of) c by these elements G_2. We define G_3 by passing to a maximal subset of G_2 such that elements are sufficiently far apart. This does not change the set much; in particular, the growth rate is unchanged. However, it will be an important point for the injection argument, because we show in Lemma 3.5 that if g and h are in G_3 then $g\mathcal{E} = h\mathcal{E}$ implies $g = h$. The final refinement is to pass to the subset G_4 of G_3 of elements that are not 'in the shadow' of some other element of G_3, that is to say, elements g such that there does not exist h such that a geodesic from o to $g.o$ passes close to $h.o$. The crux of the argument, Lemma 3.6, is to show that at least half of G_3 is unshadowed, so G_4 is divergent with growth rate $\delta_G/2$. Finally, in Lemma 3.7, we check that G_4 gives the desired injection.

Fix an element $f_0 \in G$ such that $f_0\mathcal{E}$ is disjoint from $\mathcal{E}$, $o \in \pi_{\mathcal{E}}(f_0.o)$, and $f_0.o \in \pi_{f_0\mathcal{E}}(o)$. To see that such an element exists, first note that there exists $g \in G - E(c)$, for instance, as in the first paragraph of the proof of Theorem 1.1. If $\mathcal{E}$ and $g\mathcal{E}$ are disjoint, let f_1 and f_2 be elements of G such that $f_1.o \in \mathcal{E}$ and $f_2.o \in g\mathcal{E}$ realize the minimum distance between $\mathcal{E}$ and $g\mathcal{E}$. Then the element $f_0 := f_1^{-1}f_2$ satisfies our requirements. If $g\mathcal{E}$ and $\mathcal{E}$ are not disjoint consider $g\mathcal{E}$ and $c^ng\mathcal{E}$, for some n. If they intersect then, by (P 0):

$$2\theta \geqslant d^\pi_{\mathcal{E}}(g\mathcal{E}, g\mathcal{E}) + d^\pi_{\mathcal{E}}(c^ng\mathcal{E}, c^ng\mathcal{E}) \geqslant d^\pi_{\mathcal{E}}(g\mathcal{E}, c^ng\mathcal{E}) \geqslant |c^n|$$

This is impossible once n is sufficiently large as c is strongly contracting. So, $g\mathcal{E}$ and $c^ng\mathcal{E}$ are disjoint for such n, and we get f_0 by the previous argument after replacing g with $g^{-1}c^ng$.

Since $\mathcal{E}$ and $f_0\mathcal{E}$ are disjoint and o and $f_0.o$ are contained in one another's projections, strong contraction of c, and hence of $\mathcal{E}$, gives a constant $C \geqslant 0$ such that:

$$d^\pi_{f_0\mathcal{E}}(o, f_0.o) = \operatorname{diam}\pi_{f_0\mathcal{E}}(o) \leqslant C \quad \text{and} \quad d^\pi_{\mathcal{E}}(o, f_0.o) = \operatorname{diam}\pi_{\mathcal{E}}(f_0.o) \leqslant C \tag{1}$$

In the sequel, we use the following notation: $|f_0|$ is the length of the element f_0 just defined; Δ is as in the definition of purely exponential growth of G; C is a contraction constant for $\mathcal{E}$; C' is the corresponding constant from Corollary 2.1; θ and θ' are as in Theorem 2.2; K is a fixed constant strictly greater than $\max\{C, \theta + \theta'/2\}$. We call these, collectively, 'the constants'. The terms 'small' and 'close' mean bounded by some combination of the constants. When possible we decline to compute these explicitly since only finitely many such combinations appear in the proof, except where noted. Furthermore, Δ depends only on G, and the others depend only on $\mathcal{E} = E(c).o$. Since $E(c) = E(c^p)$ for all $p \neq 0$, we can, and will, pass to high powers of c to make $|c^p|$ much larger than all of the constants and combinations of them that we encounter.

Set $G_1 := \{g \in G \mid d^\pi_{\mathcal{E}}(o, g.o) \leqslant 2K \text{ and } d^\pi_{g\mathcal{E}}(o, g.o) \leqslant 2K \text{ and } g\mathcal{E} \neq \mathcal{E}\}$. This is a subset of G that is closed under taking inverses.

Lemma 3.2. *For every $g \in G$ at least one of the elements g, $f_0 g$, $g f_0$, or $f_0 g f_0$ belongs to G_1.*

Proof. First, consider $g \notin E(c)$ with $|g| \leqslant K$. Recall $g \in E(c)$ if and only if $g\mathcal{E} = \mathcal{E}$. By definition, $\pi_{\mathcal{E}}(g.o)$ is the set of points of $\mathcal{E}$ minimizing the distance to $g.o$. By hypothesis, o is a point of $\mathcal{E}$ at distance at most K from $g.o$ so $d(g.o, \pi_{\mathcal{E}}(g.o)) \leqslant K$, and $d^\pi_{\mathcal{E}}(o, g.o) = \operatorname{diam}\{o\} \cup \pi_{\mathcal{E}}(g.o) \leqslant 2K$. The same argument for o projecting to $g\mathcal{E}$ gives $d^\pi_{g\mathcal{E}}(o, g.o) \leqslant 2K$. Thus, elements g of this form already belong to G_1.

Next, consider an element $g \in E(c)$ such that $|g| \leqslant K$. Since $g \in E(c)$, we have $f_0 g\mathcal{E} = f_0\mathcal{E} \neq \mathcal{E}$ and $\pi_{\mathcal{E}}(g.o) = g.o$, so $d^\pi_{\mathcal{E}}(o, g.o) = d(o, g.o) \leqslant K$. Using this estimate and (1), we see:

$$d^\pi_{f_0 g\mathcal{E}}(o, f_0 g.o) \leqslant d^\pi_{f_0 g\mathcal{E}}(o, f_0.o) + d^\pi_{f_0 g\mathcal{E}}(f_0.o, f_0 g.o) = d^\pi_{f_0\mathcal{E}}(o, f_0.o) + d^\pi_{\mathcal{E}}(o, g.o) \leqslant C + K < 2K$$

In the other direction, using the fact that $o \in \pi_{\mathcal{E}}(f_0.o) \subset \pi_{\mathcal{E}}(f_0\mathcal{E})$, along with (P 0):

$$d^\pi_{\mathcal{E}}(o, f_0 g.o) \leqslant d^\pi_{\mathcal{E}}(o, f_0\mathcal{E}) \leqslant d^\pi_{\mathcal{E}}(f_0\mathcal{E}, f_0\mathcal{E}) \leqslant \theta < K$$

Note that we did not use $d^\pi_{\mathcal{E}}(o, g.o) \leqslant K$ for this direction—the inequality is valid for any $g \in E(c)$.

Suppose $g \notin E(c)$ and $d^\pi_{\mathcal{E}}(o, g.o) > K$ then:

$$\theta < K < d^\pi_{\mathcal{E}}(o, g.o) = d^\pi_{f_0\mathcal{E}}(f_0.o, f_0 g.o) \leqslant d^\pi_{f_0\mathcal{E}}(\mathcal{E}, f_0 g\mathcal{E})$$

This contradicts (P 0) if $\mathcal{E} = f_0 g\mathcal{E}$, since, by hypothesis, $f_0\mathcal{E} \neq \mathcal{E}$ and $f_0 g\mathcal{E} \neq f_0\mathcal{E}$. Thus, $\mathcal{E}$, $f_0\mathcal{E}$, and $f_0 g\mathcal{E}$ are distinct, and we can apply (P 1) to get:

$$d^\pi_{\mathcal{E}}(o, f_0 g.o) \leqslant d^\pi_{\mathcal{E}}(f_0\mathcal{E}, f_0 g\mathcal{E}) \leqslant \theta < K$$

For $|g| \leqslant K$ we are done, either g or $f_0 g$ is in G_1, and for $|g| > K$ we have shown that there is at least one choice of $g' \in \{g, f_0 g\}$ such that $g'\mathcal{E} \neq \mathcal{E}$ and $d^\pi_{\mathcal{E}}(o, g'.o) \leqslant K$. If $d^\pi_{g'\mathcal{E}}(o, g'.o) \leqslant K$ then we are done, so suppose not. Consider the possibility that $g' f_0\mathcal{E} = \mathcal{E}$. Then $g' f_0.o \in \mathcal{E}$, so $o \in \pi_{\mathcal{E}}(f_0.o)$ implies $g'.o \in \pi_{g'\mathcal{E}}(g' f_0.o) \subset \pi_{g'\mathcal{E}}(\mathcal{E})$. Since $g'\mathcal{E} \neq \mathcal{E}$, (P 0) says $d^\pi_{g'\mathcal{E}}(\mathcal{E}, \mathcal{E}) \leqslant \theta$, so:

$$K < d^\pi_{g'\mathcal{E}}(o, g'.o) \leqslant d^\pi_{g'\mathcal{E}}(\mathcal{E}, \mathcal{E}) \leqslant \theta < K$$

This is a contradiction, so $\mathcal{E}$, $g'\mathcal{E}$, and $g' f_0\mathcal{E}$ are distinct. Observe, since $g'.o \in \pi_{g'\mathcal{E}}(g' f_0.o)$:

$$d^\pi_{g'\mathcal{E}}(\mathcal{E}, g' f_0\mathcal{E}) \geqslant d^\pi_{g'\mathcal{E}}(o, g' f_0.o) \geqslant d^\pi_{g'\mathcal{E}}(o, g'.o) > K > \theta$$

Thus, by (P 1) and the fact that $g' f_0.o \in \pi_{g' f_0\mathcal{E}}(g'.o)$, we have $d^\pi_{g' f_0\mathcal{E}}(o, g' f_0.o) \leqslant d^\pi_{g' f_0\mathcal{E}}(\mathcal{E}, g'\mathcal{E}) \leqslant \theta < K$.

To check that the first inequality has not been spoiled, use the fact that $d^\pi_{g'\mathcal{E}}(\mathcal{E}, g' f_0\mathcal{E}) > \theta$, so (P 1) implies $d^\pi_{\mathcal{E}}(g'\mathcal{E}, g' f_0\mathcal{E}) \leqslant \theta$, which gives:

$$d^\pi_{\mathcal{E}}(o, g' f_0.o) \leqslant d^\pi_{\mathcal{E}}(o, g'.o) + d^\pi_{\mathcal{E}}(g'.o, g' f_0.o) \leqslant K + d^\pi_{\mathcal{E}}(g'\mathcal{E}, g' f_0\mathcal{E}) < K + \theta < 2K \qquad \square$$

Define $\phi_0\colon G \to G_1$ by fixing G_1 and sending an element $g \in G - G_1$ to an arbitrary element of the nonempty set $\{f_0 g,\ g f_0,\ f_0 g f_0\} \cap G_1$. The map ϕ_0 is surjective, at most 4-to-1, and changes norm by at most $2|f_0|$.

For each $p \in \mathbb{N}$, define $G_{2,p} := \{g^{-1} c^p g \mid g \in G_1\}$ and $\phi_{1,p}\colon G_1 \to G_{2,p} : g \mapsto g^{-1} c^p g$.

Lemma 3.3. *If p is sufficiently large then for every $g \in G_1$ we have:*

$$2|g| + |c^p| - 8C' - 8K \leqslant |\phi_{1,p}(g)| \leqslant 2|g| + |c^p|$$

Proof. The upper bound is clear. We derive a lower bound from strong contraction. From the definition of G_1 it follows that $\pi_{g^{-1}\mathcal{E}}(o) \subset \bar{B}_{2K}(g^{-1}.o)$ and $\pi_{g^{-1}\mathcal{E}}(g^{-1}c^p g.o) \subset \bar{B}_{2K}(g^{-1}c^p.o)$, so:

$$|c^p| - 4K \leqslant d^{\pi}_{g^{-1}\mathcal{E}}(o, g^{-1}c^p g.o) \leqslant |c^p| + 4K \tag{2}$$

Let γ be a geodesic from o to $g^{-1}c^p g.o$. Its endpoints have projection to $g^{-1}\mathcal{E}$ at distance at least $|c^p| - 4K \gg C'$ from one another, for p sufficiently large, as c is strongly contracting. Thus, for t_0 and t_1 as in Corollary 2.1, we have $d(\gamma(t_0), \pi_{g^{-1}\mathcal{E}}(o)) \leqslant 2C'$, so $d(\gamma(t_0), g^{-1}.o) \leqslant 2C' + 2K$, and, similarly, $d(\gamma(t_1), g^{-1}c^p.o) \leqslant 2C' + 2K$.

$$\begin{aligned}
|\phi_{1,p}(g)| = |\gamma| &= d(o, \gamma(t_0)) + d(\gamma(t_0), \gamma(t_1)) + d(\gamma(t_1), g^{-1}c^p g.o) \\
&\geqslant \Big(d(o, g^{-1}.o) - (2C' + 2K)\Big) + \Big(d(g^{-1}.o, g^{-1}c^p.o) - 2(2C' + 2K)\Big) \\
&\qquad + \Big(d(g^{-1}c^p.o, g^{-1}c^p g.o) - (2C' + 2K)\Big) \\
&= 2|g| + |c^p| - 8C' - 8K
\end{aligned}$$

□

The following lemma also follows from (2).

Lemma 3.4. *Let $g^{-1}c^p g = \phi_{1,p}(g) \in G_{2,p}$. If p is sufficiently large then $g^{-1}\mathcal{E} \in \mathbf{Y}(\mathcal{E}, g^{-1}c^p g\mathcal{E})$.*

We also claim $\phi_{1,p}$ is bounded-to-one, independent of p. To see this, fix $g \in G_1$ and consider $h \in G_1$ such that $\phi_{1,p}(g) = \phi_{1,p}(h)$. Then gh^{-1} commutes with c^p, so $gh^{-1} \in E(c^p) = E(c)$. Thus:

$$\begin{aligned}
|gh^{-1}| &= d^{\pi}_{\mathcal{E}}(o, gh^{-1}.o) \\
&\leqslant d^{\pi}_{\mathcal{E}}(o, g.o) + d^{\pi}_{\mathcal{E}}(g.o, gh^{-1}.o) \\
&= d^{\pi}_{\mathcal{E}}(o, g.o) + d^{\pi}_{hg^{-1}\mathcal{E}}(h.o, o) \\
&= d^{\pi}_{\mathcal{E}}(o, g.o) + d^{\pi}_{\mathcal{E}}(h.o, o) \\
&\leqslant 4K
\end{aligned}$$

So, h satisfies $h^{-1}.o \in \bar{B}_{4K}(g^{-1}.o)$. By properness of $G \curvearrowright X$, $\#G.o \cap \bar{B}_{4K}(g^{-1}.o) = \#G.o \cap \bar{B}_{4K}(o)$ is finite.

Let $G_{3,p}$ be a maximal $(6K + 1)$–separated subset of $G_{2,p}$, that is, a subset that is maximal for inclusion among those with the property that $d(g.o, h.o) \geqslant 6K + 1$ for distinct elements g and h. Let $\phi_{2,p}\colon G_{2,p} \to G_{3,p}$ be a choice of closest point. This map is surjective. By maximality, $\phi_{2,p}$ moves points a distance less than $6K + 1$. Thus, by properness of $G \curvearrowright X$, the map $\phi_{2,p}$ is bounded-to-one, independent of p.

Lemma 3.5. *If p is sufficiently large then $g^{-1}c^p g\mathcal{E} = h^{-1}c^p h\mathcal{E}$ for $g^{-1}c^p g$ and $h^{-1}c^p h$ in $G_{3,p}$ implies $g^{-1}c^p g = h^{-1}c^p h$.*

Proof. Since $g \in G_1$, $d^{\pi}_{g\mathcal{E}}(o, g.o) \leqslant 2K$, and:

$$\begin{aligned}
d^{\pi}_{g^{-1}c^p g\mathcal{E}}(o, g^{-1}c^p g.o) &\leqslant d^{\pi}_{g^{-1}c^p g\mathcal{E}}(o, g^{-1}c^p.o) + d^{\pi}_{g^{-1}c^p g\mathcal{E}}(g^{-1}c^p.o, g^{-1}c^p g.o) \\
&\leqslant d^{\pi}_{g^{-1}c^p g\mathcal{E}}(\mathcal{E}, g^{-1}\mathcal{E}) + 2K
\end{aligned}$$

Furthermore, $g \in G_1$ implies $\mathcal{E} \neq g^{-1}\mathcal{E} \neq g^{-1}c^p g\mathcal{E}$. By (2), $d^{\pi}_{g^{-1}\mathcal{E}}(\mathcal{E}, g^{-1}c^p g\mathcal{E}) \geqslant |c^p| - 4K \gg \theta$, so by (P 0), $\mathcal{E} \neq g^{-1}c^p g\mathcal{E}$. Thus $\mathcal{E}$, $g^{-1}\mathcal{E}$, and $g^{-1}c^p g\mathcal{E}$ are distinct and we can apply (P 1) to see $d^{\pi}_{g^{-1}c^p g\mathcal{E}}(\mathcal{E}, g^{-1}\mathcal{E}) \leqslant \theta < K$. Plugging this into previous inequality gives:

$$d^{\pi}_{g^{-1}c^p g\mathcal{E}}(o, g^{-1}c^p g.o) < 3K \tag{3}$$

The same computation applies for h, so $\pi_{g^{-1}c^pg\mathcal{E}}(o) \subset \bar{B}_{3K}(g^{-1}c^pg.o) \cap \bar{B}_{3K}(h^{-1}c^ph.o)$. Thus, $g^{-1}c^pg$ and $h^{-1}c^ph$ are elements at distance at most $6K$ in a $(6K+1)$–separated set; hence, they are equal. □

For each $D \geqslant 0$, consider the set $G'_{4,p,D}$ consisting of elements $g^{-1}c^pg \in G_{3,p}$ such that there exists a different element $h^{-1}c^ph \in G_{3,p}$ such that $h^{-1}c^phc^{2p}.o$ is within distance D of a geodesic γ from o to $g^{-1}c^pg.o$. Define $G_{4,p,D} := G_{3,p} - G'_{4,p,D}$.

Lemma 3.6. *For all $D \geqslant 0$, for p sufficiently large, $G_{4,p,D}$ is divergent and $\delta_{G_{4,p,D}} = \delta_G/2$.*

Proof. The maps $\phi_{2,p}$, $\phi_{1,p}$, and ϕ_0 are surjective and bounded-to-one, with bound independent of p, so their composition is as well. Furthermore, we know how they change norm: ϕ_0 moves points at most $2|f_0|$, $\phi_{2,p}$ moves less than $6K+1$, and $|\phi_{1,p}(g)|$ is estimated in Lemma 3.3. Putting these together, for any $r \geqslant 0$ and $g \in G \cap S^\Delta_r$ we have:

$$(4)\quad 2r + |c^p| - 4|f_0| - 8C' - 14K - 1 \leqslant |\phi_{2,p} \circ \phi_{1,p} \circ \phi_0(g)| < 2r + |c^p| + 2\Delta + 4|f_0| + 6K + 1$$

Let $t := 2r + |c^p| - 4|f_0| - 8C' - 14K - 1$, $E := 4|f_0| + 4C' + 10K + 1$, and $\Delta' := 2(\Delta + E)$, so that (4) shows:

$$\phi_{2,p} \circ \phi_{1,p} \circ \phi_0(G \cap S^\Delta_r) \subset G_{3,p} \cap S^{\Delta'}_t \subset \phi_{2,p} \circ \phi_{1,p} \circ \phi_0(G \cap S^{\Delta+2E}_{r-E})$$

This lets us compare the size of spherical shells in $G_{3,p}$ and G:

$$(5)\qquad \#G \cap S^{\Delta+2E}_{r-E} \geqslant \#G_{3,p} \cap S^{\Delta'}_t \overset{*}{\succ} \#G \cap S^\Delta_r$$

Pure exponential growth of G says that $\#G \cap S^\Delta_r \overset{*}{\asymp} \exp(r\delta_G)$. Combining this with (5), we have:

$$(6)\qquad \#G_{3,p} \cap S^{\Delta'}_t \overset{*}{\asymp} \exp(\delta_G r) \overset{*}{\asymp} \exp(-\delta_G|c^p|/2)\exp(t\delta_G/2)$$

This tells us that $\delta_{G_{3,p}} = \delta_G/2$ and $G_{3,p}$ is divergent.

Now we will estimate an upper bound for $\#G'_{4,p,D} \cap S^{\Delta'}_r$ and see that for large p and r it is less than half of $\#G_{3,p} \cap S^{\Delta'}_r$. Thus, to get $G_{4,p,D}$ we threw away less than half of $G_{3,p}$, at least outside a sufficiently large radius. We conclude that $\delta_{G_{4,p,D}} = \delta_G/2$ and $G_{4,p,D}$ is divergent.

Consider $g^{-1}c^pg \in G'_{4,p,D} \cap S^{\Delta'}_r$ for any $r > 7|c^p|$. By definition of $G'_{4,p,D}$, there exists $h^{-1}c^ph \in G_{3,p}$ such that $h^{-1}c^ph \neq g^{-1}c^pg$ and $h^{-1}c^phc^{2p}.o$ is close to a geodesic γ from o to $g^{-1}c^pg.o$.

Let $\sqsubset$ be the order of Proposition 2.3 on $\mathbf{Y}[\mathcal{E}, g^{-1}c^pg\mathcal{E}]$. The first step of the proof is to show that $\mathcal{E}$, $g^{-1}\mathcal{E}$, $g^{-1}c^pg\mathcal{E}$, $h^{-1}\mathcal{E}$, and $h^{-1}c^ph\mathcal{E}$ are distinct elements of $\mathbf{Y}[\mathcal{E}, g^{-1}c^pg\mathcal{E}]$, and that the ordering is one of the two possibilities shown in Figure 1 and Figure 2.

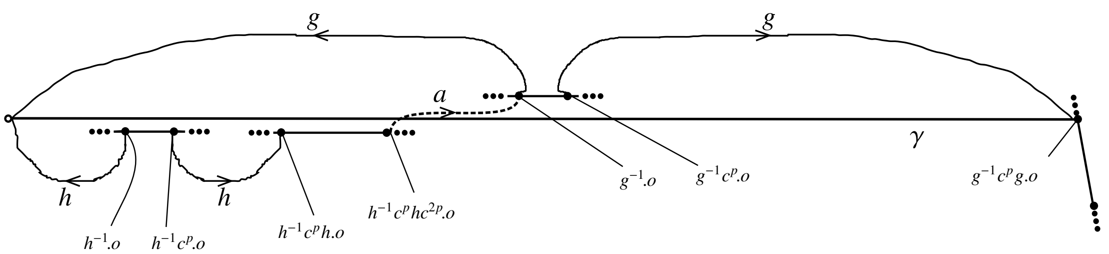

Figure 1. $h^{-1}c^ph\mathcal{E}$ before $g^{-1}\mathcal{E}$, that is, $h^{-1}c^ph\mathcal{E} \sqsubset g^{-1}\mathcal{E}$

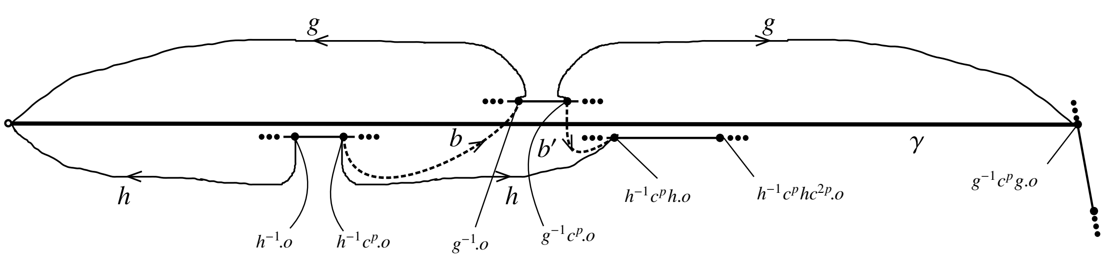

Figure 2. $h^{-1}c^ph\mathcal{E}$ after $g^{-1}\mathcal{E}$, that is, $g^{-1}\mathcal{E} \sqsubset h^{-1}c^ph\mathcal{E}$

By Lemma 3.4, $\mathcal{E} \sqsubset g^{-1}\mathcal{E} \sqsubset g^{-1}c^p g\mathcal{E}$, so these three are distinct. Similarly, $\mathcal{E}$, $h^{-1}\mathcal{E}$, and $h^{-1}c^p h\mathcal{E}$ are distinct. Lemma 3.5 implies $g^{-1}c^p g\mathcal{E} \neq h^{-1}c^p h\mathcal{E}$.

We have $|c^p| + 2|g| \geqslant |g^{-1}c^p g| \overset{+}{\succ} |h^{-1}c^p hc^{2p}|$ since $h^{-1}c^p hc^{2p}.o$ is close to a geodesic from o to $g^{-1}c^p g.o$. On the other hand, any geodesic from o to $h^{-1}c^p hc^{2p}.o$ has projection to $h^{-1}c^p h\mathcal{E}$ of diameter greater than $|c^{2p}| - 3K$ by (3). This is much larger than C' when p is large, so $|h^{-1}c^p hc^{2p}| \overset{+}{\asymp} |h^{-1}c^p h| + |c^{2p}| \overset{+}{\succ} 3|c^p| + 2|h|$ by Corollary 2.1 and Lemma 3.3. Thus:

$$|g| \overset{+}{\succ} |h| + |c^p| \tag{7}$$

However, by definition of G_1, if $h^{-1}\mathcal{E} = g^{-1}\mathcal{E}$, then:

$$4K \geqslant d^{\pi}_{g^{-1}\mathcal{E}}(o, g^{-1}.o) + d^{\pi}_{h^{-1}\mathcal{E}}(o, h^{-1}.o) \geqslant d(g^{-1}.o, h^{-1}.o) \geqslant |g| - |h| \overset{+}{\succ} |c^p|$$

This is a contradiction for sufficiently large p. Similar considerations show $h^{-1}\mathcal{E} \neq g^{-1}c^p g\mathcal{E}$, since o projects close to $h^{-1}.o$ in $h^{-1}\mathcal{E}$, by definition of G_1, and close to $g^{-1}c^p g.o$ in $g^{-1}c^p g\mathcal{E}$, by (3), but $|h| \ll |g^{-1}c^p g|$, by Lemma 3.3 and (7).

Next we show that $h^{-1}\mathcal{E}$ and $h^{-1}c^p h\mathcal{E}$ belong to $\mathbf{Y}[\mathcal{E}, g^{-1}c^p g\mathcal{E}]$, and in the course of the proof we will observe $g^{-1}\mathcal{E} \neq h^{-1}c^p h\mathcal{E}$. By hypothesis, there exists t such that $d(\gamma(t), h^{-1}c^p hc^{2p}.o) \leqslant D$. This implies $d^{\pi}_{h^{-1}c^p h\mathcal{E}}(\gamma(t), h^{-1}c^p hc^{2p}.o) \leqslant 2D$. Since $d^{\pi}_{h^{-1}c^p h\mathcal{E}}(o, h^{-1}c^p h.o) < 3K$, by (3), we have $d^{\pi}_{h^{-1}c^p h\mathcal{E}}(o, \gamma(t)) \geqslant |c^{2p}| - 2D - 3K$, which is large for p sufficiently large. Let t_0 and t_1 be the first and last times γ is distance C' from $h^{-1}c^p h\mathcal{E}$, as in Corollary 2.1 with respect to $h^{-1}c^p h\mathcal{E}$. We cannot have $t \leqslant t_0$, since then $d^{\pi}_{h^{-1}c^p h\mathcal{E}}(o, \gamma(t)) \leqslant C'$, which is a contradiction for large p.

If $t \geqslant t_1$ then $d^{\pi}_{h^{-1}c^p h\mathcal{E}}(\gamma(t), g^{-1}c^p g.o) \leqslant C'$, so:

$$\begin{aligned} d^{\pi}_{h^{-1}c^p h\mathcal{E}}(o, g^{-1}c^p g.o) &\geqslant d^{\pi}_{h^{-1}c^p h\mathcal{E}}(o, \gamma(t)) - d^{\pi}_{h^{-1}c^p h\mathcal{E}}(\gamma(t), g^{-1}c^p g.o) \\ &\geqslant |c^{2p}| - 3K - 2D - C' \end{aligned}$$

If $t_0 < t < t_1$ then we use Corollary 2.1 to say $d^{\pi}_{h^{-1}c^p h\mathcal{E}}(o, g^{-1}c^p g.o) \geqslant |\gamma(t_0, t_1)| - 4C'$, and then estimate:

$$\begin{aligned} |\gamma(t_0, t_1)| &\geqslant d(\gamma(t_0), \gamma(t)) \\ &\geqslant d\left(\pi_{h^{-1}c^p h\mathcal{E}}(\gamma(t_0)), \pi_{h^{-1}c^p h\mathcal{E}}(\gamma(t))\right) - C' - D \\ &\geqslant d^{\pi}_{h^{-1}c^p h\mathcal{E}}(\gamma(t_0), \gamma(t)) - \operatorname{diam} \pi_{h^{-1}c^p h\mathcal{E}}(\gamma(t_0)) - \operatorname{diam} \pi_{h^{-1}c^p h\mathcal{E}}(\gamma(t)) - C' - D \\ &\geqslant d^{\pi}_{h^{-1}c^p h\mathcal{E}}(\gamma(t_0), \gamma(t)) - 2C - C' - D \\ &\geqslant d^{\pi}_{h^{-1}c^p h\mathcal{E}}(o, \gamma(t)) - d^{\pi}_{h^{-1}c^p h\mathcal{E}}(\gamma(t_0), o) - 2C - C' - D \\ &\geqslant |c^{2p}| - 2D - 3K - C' - 2C - C' - D \end{aligned}$$

Thus, $h^{-1}c^p h\mathcal{E} \in \mathbf{Y}[\mathcal{E}, g^{-1}c^p g\mathcal{E}]$ once p is sufficiently large. Additionally, this shows $g^{-1}\mathcal{E} \neq h^{-1}c^p h\mathcal{E}$ because, by (2) and (P 0), we have $d^{\pi}_{g^{-1}\mathcal{E}}(\mathcal{E}, g^{-1}c^p g\mathcal{E}) \overset{+}{\asymp} |c^p|$, while from the estimates above we have $d^{\pi}_{h^{-1}c^p h\mathcal{E}}(\mathcal{E}, g^{-1}c^p g\mathcal{E}) \overset{+}{\succ} |c^{2p}|$, and these are incompatible for sufficiently large p. Thus, the five axes are distinct.

From Corollary 2.1 we deduce that:

$$d(h^{-1}c^p hc^{2p}.o, h^{-1}\mathcal{E}) \overset{+}{\succ} d^{\pi}_{h^{-1}c^p h\mathcal{E}}(h^{-1}c^p hc^{2p}.o, h^{-1}\mathcal{E}) \overset{+}{\asymp} |c^{2p}|$$

Thus, for large enough p we have $d(h^{-1}c^p hc^{2p}.o, h^{-1}\mathcal{E}) \geqslant D \geqslant d(\gamma(t), h^{-1}c^p hc^{2p}.o)$, so strong contraction of $h^{-1}\mathcal{E}$ implies $d^{\pi}_{h^{-1}\mathcal{E}}(\gamma(t), h^{-1}c^p hc^{2p}.o) \leqslant C$. Since o projects close to $h^{-1}.o$ in $h^{-1}\mathcal{E}$ and $h^{-1}c^p hc^{2p}.o \in h^{-1}c^p h\mathcal{E}$ projects close to $h^{-1}c^p.o$, Corollary 2.1 says γ must pass close to $h^{-1}c^p.o$. Now we can run the same argument as for $h^{-1}c^p h\mathcal{E}$ to see $h^{-1}\mathcal{E} \in \mathbf{Y}[\mathcal{E}, g^{-1}c^p g\mathcal{E}]$ once p is sufficiently large.

The first step of the proof is completed by observing that $g^{-1}\mathcal{E} \sqsubset h^{-1}\mathcal{E}$ implies $|h| \overset{+}{\asymp} |g| + |c^p|$, which cannot be true when p is sufficiently large, by (7). Thus, $h^{-1}\mathcal{E}$ comes before $g^{-1}\mathcal{E}$ and $h^{-1}c^p h\mathcal{E}$ under $\sqsubset$, and we are left with the possibilities that $h^{-1}c^p h\mathcal{E} \sqsubset g^{-1}\mathcal{E}$, as in Figure 1, or the converse, as in Figure 2.

In the case of Figure 1, we have $h^{-1}c^p h\mathcal{E} \sqsubset g^{-1}\mathcal{E}$, so the projection of $h^{-1}c^p hc^{2p}.o$ to $g^{-1}\mathcal{E}$ is close to the projection of o, which we know to be close to $g^{-1}.o$. Write $g^{-1}.o = h^{-1}c^p hc^{2p}a.o$ as in Figure 1 with $|g| \stackrel{+}{\asymp} 2|h| + 3|c^p| + |a|$.

In the case of Figure 2, we have $h^{-1}\mathcal{E} \sqsubset g^{-1}\mathcal{E}$ and $g^{-1}\mathcal{E} \sqsubset h^{-1}c^p h\mathcal{E}$. The former implies the projection of $h^{-1}c^p.o$ to $g^{-1}\mathcal{E}$ is close to the projection of o, which we know to be close to $g^{-1}.o$, while the latter implies the projection of $h^{-1}c^p h.o$ to $g^{-1}\mathcal{E}$ is close to the projection of $g^{-1}c^p g.o$, which we know to be close to $g^{-1}c^p.o$. Write $g^{-1}.o = h^{-1}c^p b.o$ with $|g| \stackrel{+}{\asymp} |h| + |c^p| + |b|$ and write $h.o = bc^p b'.o$ as in Figure 2 with $|h| \stackrel{+}{\asymp} |b| + |c^p| + |b'|$; together these give $|g| \stackrel{+}{\asymp} 2|b| + 2|c^p| + |b'|$.

Suppose we are in the case of Figure 2, so there are elements b and b' such that $(r - |c^p|)/2 \stackrel{+}{\asymp} |g| \stackrel{+}{\asymp} 2|b| + 2|c^p| + |b'|$. Since G has purely exponential growth, if $i \leqslant |b| < i + 1$ there are, up to a bounded multiplicative error independent of p, r, and i, at most $\exp(\delta_G i)$ possible choices for b and at most $\exp(\delta_G(\frac{r-5|c^p|}{2} - 2i))$ choices of b', so there is an upper bound for the number of possible elements g by a multiple of:

$$\sum_{i=0}^{\frac{r-5|c^p|}{4}} \exp(\delta_G i)\exp\left(\delta_G\left(\frac{r-5|c^p|}{2} - 2i\right)\right) < \frac{\exp(r\delta_G/2)}{\exp(5\delta_G|c^p|/2)(1-\exp(-\delta_G))} \tag{8}$$

The case of Figure 1 is similar, but gives an even smaller upper bound[1]. Thus, for all sufficiently large p and r:

$$\#G'_{4,p,D} \cap S_r^{\Delta'} \stackrel{*}{\prec} \exp(-5\delta_G|c^p|/2)\exp(r\delta_G/2) \tag{9}$$

Combining (6) and (9) gives:

$$\#G'_{4,p,D} \cap S_r^{\Delta'} \stackrel{*}{\prec} \exp(-2|c^p|\delta_G) \cdot \#G_{3,p} \cap S_r^{\Delta'} \tag{10}$$

Crucially, the multiplicative constant in this asymptotic inequality does not depend on p, so for p sufficiently large, $\exp(2|c^p|\delta_G)$ is more than twice the multiplicative constant, and (10) becomes a true inequality $\#G'_{4,p,D} \cap S_r^{\Delta'} < \frac{1}{2}\#G_{3,p} \cap S_r^{\Delta'}$. We conclude that to get $G_{4,p,D}$ from $G_{3,p}$ we threw away fewer than half of the points of $G_{3,p}$ in each spherical shell $S_r^{\Delta'}$ such that $r > 7|c^p|$. □

LEMMA 3.7. *For all sufficiently large D, for all sufficiently large p, the map* $\hat{G}_{4,p,D} \to X : (g_1, \ldots, g_k) \mapsto (\prod_{i=1}^k g_i c^{2p}).o$ *is an injection.*

PROOF. Consider a point $(\prod_{i=1}^k g_i c^{2p}).o$ in the image. Set $g_0 := c^{-2p}$. Suppose that for each i we have $g_i = e_i^{-1}c^p e_i$ for $e_i \in G_1$. For $0 \leqslant i \leqslant k$ set $z'_{2i} := (\prod_{j=0}^i g_j c^{2p}).o$, $z_{2i} := (\prod_{j=0}^i g_j c^{2p})c^{-2p}.o$, and $\mathcal{Z}_{2i} := (\prod_{j=0}^i g_j c^{2p})\mathcal{E}$. For $0 < i \leqslant k$ set $z_{2i-1} := (\prod_{j=0}^{i-1} g_j c^{2p})e_{2i-1}^{-1}.o$, $z'_{2i-1} := (\prod_{j=0}^{i-1} g_j c^{2p})e_{2i-1}^{-1}c^p.o$, and $\mathcal{Z}_{2i-1} := (\prod_{j=0}^{i-1} g_j c^{2p})e_{2i-1}^{-1}\mathcal{E}$. See Figure 3.

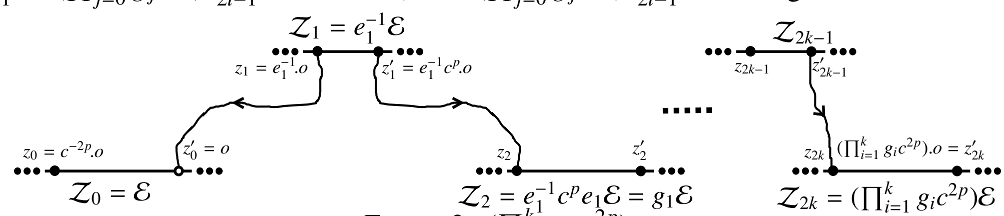

FIGURE 3. $(\prod_{i=1}^k g_i c^{2p}).o$

Let us complete the proof assuming the following claim, to which we shall return:

$$\forall 0 \leqslant i < j \leqslant 2k, \quad d^\pi_{\mathcal{Z}_i}(z'_i, \mathcal{Z}_j) < 5K \quad \text{and} \quad d^\pi_{\mathcal{Z}_j}(z_j, \mathcal{Z}_i) < 5K \tag{11}$$

When p is sufficiently large, $d(z_i, z'_i) \gg 10K$ for all i, so (11) implies that $\mathcal{Z}_i \sqsubset \mathcal{Z}_j$ for all $0 \leqslant i < j \leqslant 2k$, where $\sqsubset$ is the order of Proposition 2.3 on $\mathbf{Y}[\mathcal{Z}_0, \mathcal{Z}_{2k}]$.

[1]Replace each '5' in (8) with a '7'. This accounts for the restriction that $r - 7|c^p| > 0$.

Suppose that the map $\hat{G}_{4,p,D} \to X$ is not an injection; there exist distinct elements $(g_1, \ldots, g_m)$ and $(h_1, \ldots, h_n)$ of $\hat{G}_{4,p,D}$ with the same image $z \in X$. Suppose $m + n$ is minimal among such tuples. If $h_1\mathcal{E} = g_1\mathcal{E}$ then $h_1 = g_1$ by Lemma 3.5. This contradicts minimality of $m + n$, so we must have $h_1\mathcal{E} \neq g_1\mathcal{E}$. Let $\mathcal{Z}_0, \ldots, \mathcal{Z}_{2m}$ be as in Figure 3 for $(g_1, \ldots, g_m)$. By definition, $o \in \mathcal{Z}_0$ and $z \in \mathcal{Z}_{2m}$. By (11), $\pi_{\mathcal{Z}_{2m}}(o)$ is close to z_{2m}. By Corollary 2.1, any geodesic from o to z ends with a segment that stays close to the subsegment of $\mathcal{Z}_{2m}$ between z_{2m} and $z = z'_{2m}$. However, if $\mathcal{Z}'_0, \ldots, \mathcal{Z}'_{2n}$ are as in Figure 3 for $(h_1, \ldots, h_n)$, then the same is true for $\mathcal{Z}'_{2n}$, which implies $d^\pi_{\mathcal{Z}_{2m}}(\mathcal{Z}'_{2n}, \mathcal{Z}'_{2n}) \overset{+}{\succ} d(z_{2m}, z'_{2m}) = |c^{2p}|$. Once p is sufficiently large, (P 0) requires $\mathcal{Z}_{2m} = \mathcal{Z}'_{2n}$. Thus, $\mathbf{Y}[\mathcal{Z}_0, \mathcal{Z}_{2m}] = \mathbf{Y}[\mathcal{Z}'_0, \mathcal{Z}'_{2n}]$, and all of the $\mathcal{Z}_i$ and $\mathcal{Z}'_j$ are comparable in the order $\sqsubset$ on $\mathbf{Y}[\mathcal{Z}_0, \mathcal{Z}_{2m}]$. In particular, $\mathcal{Z}'_2 = h_1\mathcal{E} \neq g_1\mathcal{E} = \mathcal{Z}_2$, so one of them comes before the other. Suppose, without loss of generality, that $h_1\mathcal{E} \sqsubset g_1\mathcal{E}$. Then $d_{h_1\mathcal{E}}(g_1\mathcal{E}, \mathcal{Z}_{2m}) \leqslant \theta'$, by Proposition 2.3, and $d^\pi_{h_1\mathcal{E}}(\mathcal{Z}_{2m}, h_1c^{2p}.o) < 5K$ by (11), so:

$$\begin{aligned} d^\pi_{h_1\mathcal{E}}(g_1.o, h_1c^{2p}.o) &\leqslant d^\pi_{h_1\mathcal{E}}(g_1\mathcal{E}, \mathcal{Z}_{2m}) + d^\pi_{h_1\mathcal{E}}(\mathcal{Z}_{2m}, h_1c^{2p}.o) \\ &< \theta' + 2\theta + 5K < 7K \end{aligned}$$

On the other hand, $d^\pi_{h_1\mathcal{E}}(o, h_1.o) < 3K$, by (3), so $d^\pi_{h_1\mathcal{E}}(o, g_1.o) \geqslant |c^{2p}| - 10K \gg C'$. By Corollary 2.1, any geodesic from o to $g_1.o$ passes within distance $2C'$ of $\pi_{h_1\mathcal{E}}(g_1.o)$, which is less than $7K$ from $h_1c^{2p}.o$. This means $g_1 \in G'_{4,p,(7K+2C')}$, which is a contradiction if $D \geqslant 7K + 2C'$. Thus, if $D \geqslant 7K + 2C'$ then for sufficiently large p the map is injective.

We prove (11) by induction on $m = j - i$. For each $0 \leqslant i < 2k$ we have that z'_i and z_{i+1} differ by an element of G_1, so $\mathcal{Z}_i \neq \mathcal{Z}_{i+1}$ and $d^\pi_{\mathcal{Z}_{i+1}}(z_{i+1}, z'_i) \leqslant 2K$. Furthermore, by (P 0), $d^\pi_{\mathcal{Z}_{i+1}}(\mathcal{Z}_i, \mathcal{Z}_i) \leqslant \theta$. Thus:

$$d^\pi_{\mathcal{Z}_{i+1}}(z_{i+1}, \mathcal{Z}_i) \leqslant d^\pi_{\mathcal{Z}_{i+1}}(z_{i+1}, z'_i) + d^\pi_{\mathcal{Z}_{i+1}}(z'_i, \mathcal{Z}_i) \leqslant d^\pi_{\mathcal{Z}_{i+1}}(z_{i+1}, z'_i) + d^\pi_{\mathcal{Z}_{i+1}}(\mathcal{Z}_i, \mathcal{Z}_i) \leqslant 2K + \theta < 3K$$

Similarly, $d^\pi_{\mathcal{Z}_i}(z'_i, \mathcal{Z}_{i+1}) < 3K$.

Now extend m to $m + 1$: Suppose that for some $m \geqslant 1$ and all $0 < j - i \leqslant m$ we have $d^\pi_{\mathcal{Z}_j}(z_j, \mathcal{Z}_i) < 5K$ and $d^\pi_{\mathcal{Z}_i}(z'_i, \mathcal{Z}_j) < 5K$. (Note that this implies $\mathcal{Z}_i \neq \mathcal{Z}_j$.) Then for all $0 \leqslant i \leqslant 2k - m - 1$:

$$d_{\mathcal{Z}_{i+1}}(\mathcal{Z}_{i+m+1}, \mathcal{Z}_i) \geqslant d^\pi_{\mathcal{Z}_{i+1}}(\mathcal{Z}_{i+m+1}, \mathcal{Z}_i) - 2\theta > d(z_{i+1}, z'_{i+1}) - 10K - 2\theta \gg \theta'$$

The final inequality is true for sufficiently large p, because the distance between z_{i+1} and z'_{i+1} is either $|c^p|$ or $|c^{2p}| \overset{+}{\asymp} 2|c^p|$, according to whether i is even or odd. Thus, by (SP 3) and (SP 4):

$$d^\pi_{\mathcal{Z}_i}(\mathcal{Z}_{i+m+1}, \mathcal{Z}_{i+1}) \leqslant d_{\mathcal{Z}_i}(\mathcal{Z}_{i+m+1}, \mathcal{Z}_{i+1}) + 2\theta = d_{\mathcal{Z}_i}(\mathcal{Z}_{i+1}, \mathcal{Z}_{i+1}) + 2\theta \leqslant \theta' + 2\theta < 2K$$

which implies:

$$d^\pi_{\mathcal{Z}_i}(z'_i, \mathcal{Z}_{i+m+1}) \leqslant d^\pi_{\mathcal{Z}_i}(z'_i, \mathcal{Z}_{i+1}) + d^\pi_{\mathcal{Z}_i}(\mathcal{Z}_{i+1}, \mathcal{Z}_{i+m+1}) < 3K + 2K = 5K$$

A similar argument gives $d^\pi_{\mathcal{Z}_{i+m+1}}(z_{i+m+1}, \mathcal{Z}_i) < 5K$. This completes the induction. □

Proof of Proposition 3.1. Take D and p as in Lemma 3.7. For this D, enlarge p if necessary to satisfy the hypotheses of Lemma 3.6. Set $G_4 := G_{4,p,D}$. □

4. Questions

Question 4.1. Can we replace purely exponential growth of G by divergence of G in Theorem 1.1?

By [19], the answer is 'yes' when X is hyperbolic.

Recall in (5) we showed $\Theta_G(s)$ is comparable to $\Theta_{G_{3,p}}(s/2)$, while it is clear that $\Theta_{G_{3,p}}(s/2) \leqslant \Theta_N(s/2)$. If G is divergent then $\Theta_G(s)$ diverges at $s = \delta_G$, which means $\Theta_N(t)$ diverges at $t = \delta_G/2$. There are two possible circumstances in which $\Theta_N(t)$ diverges at $t = \delta_G/2$:

(12) Either $\delta_N > \delta_G/2$, or $\delta_N = \delta_G/2$ and N is divergent.

We proved the first case of (12) directly, with the additional assumption of purely exponential growth of G. The approach of [19] is to prove, if X is hyperbolic, that $\delta_N = \delta_G$ when N is divergent, so, since $\delta_G > \delta_G/2$, the second case of (12) is impossible. Thus, a positive answer to Question 4.1 would be implied by a positive answer to the following question, which is also interesting in its own right.

Question 4.2. If G is a group acting properly by isometries with a strongly contracting element on a geodesic metric space X and $G \curvearrowright X$ is divergent, is it true that for every divergent normal subgroup N of G we have $\delta_N = \delta_G$?

Jaerisch and Matsuzaki [17] show that if F is a finite rank free group and N is a non-trivial normal subgroup of F then, with respect to a word metric defined by a free generating set of F, there is a inequality $\delta_N + \frac{1}{2}\delta_{F/N} \geqslant \delta_F$. Notice, $\delta_N > \delta_F/2$ by the lower cogrowth bound, and $\delta_{F/N} < \delta_F$ by growth tightness of F.

Question 4.3. Is there an analogue of Jaerisch and Matsuzaki's inequality for G acting with a strongly contracting element and complementary growth gap? Note that we know both growth tightness, by [2], and lower cogrowth bound, by Theorem 1.1, for such actions.

For $G = X = F_n$ [13, 20, 7] and $X = \mathbb{H}^2$ and G a closed surface group [5], there exists a sequence $(N_i)_{i\in\mathbb{N}}$ of normal subgroups of G such that δ_{N_i}/δ_G limits to $1/2$, so the lower cogrowth bound is optimal.

Question 4.4. Is the lower cogrowth bound optimal in Theorem 1.1?

We must mention that the *upper* cogrowth bound is also very interesting. Grigorchuk [13] and Cohen [8] showed that when F is a finite rank free group, with respect to a word metric defined by a free generating set the upper cogrowth bound $\delta_N/\delta_F = 1$ is achieved for $N \lhd F$ if and only if F/N is amenable. There have been several generalizations [6, 21, 22, 12, 10] to growth rates defined with respect to an action $G \curvearrowright X$, but the most general to date [10] still requires G to be hyperbolic, the action to be cocompact, and X to be either a Cayley graph of G or a CAT(-1) space. In the vein of our theorem, it would be very interesting to generalize such a result to a non-hyperbolic setting.

References

[1] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Negative Curvature in Graphical Small Cancellation Groups*. Groups Geom. Dyn. (in press). arXiv: 1602.03767.

[2] Goulnara N. Arzhantseva, Christopher H. Cashen, and Jing Tao. *Growth Tight Actions*. Pacific J. Math. **278**.1 (2015), pp. 1–49.

[3] Mladen Bestvina, Ken Bromberg, Koji Fujiwara, and Alessandro Sisto. *Acylindrical actions on projection complexes*. preprint. 2017. arXiv: 1711.08722v1.

[4] Mladen Bestvina, Kenneth Bromberg, and Koji Fujiwara. *Constructing group actions on quasi-trees and applications to mapping class groups*. Publ. Math. Inst. Hautes Études Sci. **122**.1 (2015), pp. 1–64.

[5] Petra Bonfert-Taylor, Katsuhiko Matsuzaki, and Edward C. Taylor. *Large and small covers of a hyperbolic manifold*. J. Geom. Anal. **22**.2 (2012), pp. 455–470.

[6] Robert Brooks. *The bottom of the spectrum of a Riemannian covering*. J. Reine Angew. Math. **357** (1985), pp. 101–114.

[7] Christophe Champetier. *Cocroissance des groupes à petite simplification*. Bull. London Math. Soc. **25**.5 (1993), pp. 438–444.

[8] Joel M. Cohen. *Cogrowth and amenability of discrete groups*. J. Funct. Anal. **48**.3 (1982), pp. 301–309.

[9] Michel Coornaert. *Mesures de Patterson-Sullivan sur le bord d'un espace hyperbolique au sens de Gromov*. Pacific J. Math. **159**.2 (1993), pp. 241–270.

[10] Rémi Coulon, Françoise Dal'Bo, and Andrea Sambusetti. *Growth gap in hyperbolic groups and amenability*. preprint. 2017. arXiv: 1709.07287v1.

[11] Françoise Dal'Bo, Marc Peigné, Jean-Claude Picaud, and Andrea Sambusetti. *On the growth of quotients of Kleinian groups*. Ergodic Theory Dynam. Systems **31**.3 (2011), pp. 835–851.

[12] Rhiannon Dougall and Richard Sharp. *Amenability, critical exponents of subgroups and growth of closed geodesics*. Math. Ann. **365**.3-4 (2016), pp. 1359–1377.

[13] R. I. Grigorchuk. *Symmetrical random walks on discrete groups*. *Multicomponent random systems*. Vol. 6. Adv. Probab. Related Topics. Dekker, New York, 1980, pp. 285–325.

[14] Rostislav Grigorchuk and Pierre de la Harpe. *On problems related to growth, entropy, and spectrum in group theory*. J. Dynam. Control Systems **3**.1 (1997), pp. 51–89.

[15] Pierre de la Harpe. *Topics in geometric group theory*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 2000, pp. vi+310.

[16] Johannes Jaerisch. *A lower bound for the exponent of convergence of normal subgroups of Kleinian groups*. J. Geom. Anal. **25**.1 (2015), pp. 298–305.

[17] Johannes Jaerisch and Katsuhiko Matsuzaki. *Growth and cogrowth of normal subgroups of a free group*. Proc. Amer. Math. Soc. **145**.10 (2017), pp. 4141–4149.

[18] Katsuhiko Matsuzaki. *Growth and cogrowth tightness of Kleinian and hyperbolic groups*. *Geometry and Analysis of Discrete Groups and Hyperbolic Spaces*. Ed. by M. Fujii, N. Kawazumi, and K. Ohshika. Vol. B66. RIMS Kôkyûroku Bessatsu. 2017, pp. 21–36.

[19] Katsuhiko Matsuzaki, Yasuhiro Yabuki, and Johannes Jaerisch. *Normalizer, divergence type and Patterson measure for discrete groups of the Gromov hyperbolic space*. preprint. 2015. arXiv: 1511.02664v1.

[20] Yann Ollivier. *Cogrowth and spectral gap of generic groups*. Ann. Inst. Fourier (Grenoble) **55**.1 (2005), pp. 289–317.

[21] Thomas Roblin. *Un théorème de Fatou pour les densités conformes avec applications aux revêtements Galoisiens en courbure négative*. Israel J. Math. **147** (2005), pp. 333–357.

[22] Thomas Roblin and Samuel Tapie. *Exposants critiques et moyennabilité*. *Géométrie ergodique*. Vol. 43. Monogr. Enseign. Math. Enseignement Math., Geneva, 2013, pp. 61–92.

[23] Yehuda Shalom. *Rigidity, unitary representations of semisimple groups, and fundamental groups of manifolds with rank one transformation group*. Ann. of Math. (2) **152**.1 (2000), pp. 113–182.

[24] Dennis Sullivan. *Related aspects of positivity in Riemannian geometry*. J. Differential Geom. **25**.3 (1987), pp. 327–351.

[25] Wenyuan Yang. *Statistically convex-cocompact actions of groups with contracting elements*. preprint. 2016. arXiv: 1612.03648v4.

PAPER G

Quasi-isometries need not induce homeomorphisms of contracting boundaries with the Gromov product topology

Christopher H. Cashen

Analysis and Geometry in Metric Spaces **4** (2016), no. 1, 278–281.
DOI:10.1515/agms-2016-0011

We consider a 'contracting boundary' of a proper geodesic metric space consisting of equivalence classes of geodesic rays that behave like geodesics in a hyperbolic space. We topologize this set via the Gromov product, in analogy to the topology of the boundary of a hyperbolic space. We show that when the space is not hyperbolic, quasi-isometries do not necessarily give homeomorphisms of this boundary. Continuity can fail even when the spaces are required to be CAT(0). We show this by constructing an explicit example.

1. Introduction

In an extremely influential paper, Gromov [7] introduced hyperbolic spaces and their boundaries. Among myriad applications, the topological type of the boundary provides a quasi-isometry invariant of the space, since quasi-isometries of hyperbolic spaces extend to homeomorphisms of their boundaries.

Recently Charney and Sultan [4] introduced a quasi-isometry invariant 'contracting boundary' for CAT(0) spaces, consisting of those equivalence classes of geodesic rays that are 'contracting', which is to say that they behave like geodesic rays in a hyperbolic space in a certain quantifiable way. As a set, the contracting boundary of a CAT(0) space can be naturally viewed as a subset of the visual boundary of the space. A quasi-isometry does induce a bijection of this contracting subset, even though it does not necessarily induce a homeomorphism of the entire visual boundary. Charney and Sultan were unable to determine if this bijection is a homeomorphism with respect to the subspace topology. Instead, they define a finer topology that they show to be quasi-isometrically invariant. We answer their question in the negative: quasi-isometries of CAT(0) spaces do not, in general, induce homeomorphisms of the contracting boundary with the subspace topology. We do so by constructing an explicit example.

2. The contracting boundary and Gromov product topology

Let X be a proper geodesic metric space. Let γ be a geodesic ray in X, and define the closest point projection map $\pi_\gamma\colon X \to 2^X$ by $\pi_\gamma(x) := \{y \in \gamma \mid d(x,y) = d(x,\gamma)\}$. Properness of X guarantees that the empty set is not in the image of π_γ.

A geodesic ray γ in X is *contracting* if there exists a non-decreasing, eventually non-negative function ρ such that $\lim_{r\to\infty} \rho(r)/r = 0$ and such that for all x and y in X, if $d(x,y) \leq d(x,\gamma)$ then $\operatorname{diam} \pi_\gamma(x) \cup \pi_\gamma(y) \leq \rho(d(x,\gamma))$. The ray is *strongly contracting* if the function ρ can be chosen to be bounded.

A geodesic ray γ in X is *Morse* if there exists a function μ such that if α is a (λ,ϵ)–quasi-geodesic with endpoints on γ, then α is contained in the $\mu(\lambda,\epsilon)$–neighborhood of γ.

It is not hard to show that a contracting ray is Morse. Cordes [5] generalizes the Charney-Sultan construction by building a 'Morse boundary' consisting of equivalence classes of Morse geodesic rays in an arbitrary geodesic metric space. In fact, the Morse and contracting properties

Supported by the Austrian Science Fund (FWF):M1717-N25.

are equivalent in geodesic metric spaces [1], so we can just as well call Cordes's construction the contracting boundary, where we allow rays satisfying the more general version of contraction defined above.

Let us describe the points of the contracting boundary. For points x, y, $z \in X$, the *Gromov product* of x and y with respect to z is defined by:

$$(x \cdot y)_z := \frac{1}{2}(d(x,z) + d(y,z) - d(x,y))$$

Fix a basepoint $o \in X$ and consider contracting geodesic rays based at o. Define an equivalence relation by $\alpha \sim \beta$ if $\lim_{i,j\to\infty}(\alpha(i) \cdot \beta(j))_o = \infty$. This relation is transitive on contracting geodesic rays because contracting rays are Morse and Morse rays are related if and only if they are at bounded Hausdorff distance from one another. Define the contracting boundary $\partial_c X$ to be the set of equivalence classes. It is easy to see that a quasi-isometry ϕ of X induces a bijection $\partial_c\phi$ of $\partial_c X$. It remains to define a topology on $\partial_c X$ and check continuity of $\partial_c\phi$.

The topology is defined by restricting to $\partial_c X$ the usual construction of the 'ideal' or 'Gromov' boundary (cf. [7, 2, 3]). Extend the Gromov product to $\partial_c X$ by:

$$(\eta \cdot \zeta)_o := \sup_{\alpha\in\eta,\beta\in\zeta} \liminf_{i,j\to\infty}(\alpha(i) \cdot \beta(j))_o$$

Given $\eta \in \partial_c X$ and $r > 0$, define $U(\eta, r) := \{\zeta \in \partial_c X \mid (\eta \cdot \zeta)_o \geq r\}$. Define the *Gromov product topology on* $\partial_c X$ to be the topology such that a set $U \subset \partial_c X$ is open if for every $\eta \in U$ there exists an $r > 0$ such that $U(\eta, r) \subset U$. Denote the contracting boundary with this topology $\partial_c^{Gp} X$.

When X is hyperbolic $\partial_c^{Gp} X$ is the usual Gromov boundary. When X is CAT(0) $\partial_c^{Gp} X$ is homeomorphic to the contracting subset of the visual boundary with the subspace topology.

Note that $\mathcal{U}_\eta := \{U(\eta, r) \mid r > 0\}$ is not necessarily a neighborhood basis at η in this topology. We do not need this fact for the conclusions of Section 3, but, as it may be of separate interest, we give a sufficient condition. Some spaces satisfy a *contraction alternative* in the sense that every geodesic ray is either strongly contracting or not contracting. We will say that such a space *is CA*. By [1], CA is equivalent to "every Morse geodesic ray is strongly contracting." Examples of CA spaces include hyperbolic spaces, in which geodesic rays are uniformly strongly contracting, and CAT(0) spaces [8].

PROPOSITION 2.1. *If X is a proper geodesic CA metric space then for all $\eta \in \partial_c X$ the set $\mathcal{U}_\eta$ is a neighborhood basis at η in $\partial_c^{Gp} X$.*

PROOF. A standard topological argument shows that $\mathcal{U}_\eta$ is a neighborhood basis at η if and only if:

(◎) $\quad \forall r > 0,\ \exists R_\eta > r,\ \forall \zeta \in U(\eta, R_\eta),\ \exists R_\zeta > 0$ such that $U(\zeta, R_\zeta) \subset U(\eta, r)$

Suppose α is a contracting geodesic ray based at o. By the contraction alternative it is strongly contracting, so there exists a $C \geq 0$ bounding its contraction function. For brevity, let us say that α is 'C–strongly contracting'. The Geodesic Image Theorem (GIT), [1, cf. Theorem 7.1], implies that if β is a geodesic segment that stays at least distance $2C$ from α then the diameter of $\pi_\alpha(\beta)$ is at most $4C$. It follows easily that if β is a geodesic ray based at o then α and β are asymptotic if and only if β is contained in the closed $6C$–neighborhood of α. In fact, this can be improved to $5C$ by a further application of the definition of strong contraction.

If A is a contracting set and B is bounded Hausdorff distance from A then B is also contracting, with contraction function determined by that of A and the Hausdorff distance [1, Lemma 6.3]. In particular, if α is C–strongly contracting then there exists a C' depending only on C such that every geodesic ray α' based at o and asymptotic to α is C'–strongly contracting. Thus, for a given $\eta \in \partial_c X$ there exists a C_η such that every geodesic ray $\alpha \in \eta$ is C_η–strongly contracting.

Claim: Given $\eta \in \partial_c X$ there exists $K_\eta \geq 0$ such that for all $\zeta \in \partial_c X \setminus \{\eta\}$ and all $\alpha \in \eta$, $\beta \in \zeta$, if $T(\alpha,\beta) := \max\{t \mid d(\beta(t), \alpha) = 2C_\eta\}$ then $|T(\alpha,\beta) - (\eta \cdot \zeta)_o| \leq K_\eta$.

Assuming the Claim, we show that condition (◎) is satisfied. Let $\eta \in \partial_c X$ and $r > 0$. Set $R_\eta := r + 2K_\eta + 13C_\eta$. For $\zeta \in U(\eta, R_\eta)$, set $R_\zeta := (\zeta \cdot \eta)_o + K_\eta + K_\zeta + 6C_\eta + 4C_\zeta$. Suppose that $\xi \in U(\zeta, R_\zeta)$. Choose $\alpha \in \eta$, $\beta \in \zeta$, and $\gamma \in \xi$. Let $x := \gamma(T(\beta,\gamma))$ and let y be a point of β at distance $2C_\eta$ from x. Let $z := \beta(T(\alpha,\beta))$. Let $w := \gamma(T(\alpha,\gamma))$.

$$\begin{aligned}
d(y,\alpha) &\geq d(y,z) - 6C_\eta && \text{by the GIT}\\
&= d(o,y) - d(o,z) - 6C_\eta\\
&\geq d(o,x) - d(o,z) - 6C_\eta - 2C_\zeta\\
&\geq (\xi\cdot\zeta)_o - (\zeta\cdot\eta)_o - K_\eta - K_\zeta - 6C_\eta - 2C_\zeta && \text{by the Claim, twice}\\
&\geq R_\zeta - (\zeta\cdot\eta)_o - K_\eta - K_\zeta - 6C_\eta - 2C_\zeta && \text{since } \xi \in U(\zeta, R_\zeta)\\
&= 2C_\zeta = d(x,y)
\end{aligned}$$

Since $d(x,y) \leq d(y,\alpha)$, the contraction property for α says the diameter of $\pi_\alpha(x) \cup \pi_\alpha(y)$ is at most C_η. With the GIT, this tells us the diameter of $\pi_\alpha(\beta([T(\alpha,\beta),\infty))) \cup \pi_\alpha(\gamma([T(\alpha,\gamma),\infty)))$ is at most $9C_\eta$. Thus, $d(w,z) \leq 13C_\eta$. The Claim gives us $d(o,z) \geq R_\eta - K_\eta$, so $d(o,w) \geq d(o,z) - 13C_\eta \geq R_\eta - K_\eta - 13C_\eta > r + K_\eta$, which, by the Claim again, yields $(\xi\cdot\eta)_o \geq r$. Hence, $U(\zeta, R_\zeta) \subset U(\eta, r)$.

It remains to prove the claim. Let α, $\alpha' \in \eta$ and β, $\beta' \in \zeta$ be arbitrary.

Consider s, $t \gg T(\alpha,\beta)$. Let γ be a geodesic from $\alpha(s)$ to $\beta(t)$. Let z be the last point on γ at distance $2C_\eta$ from α. Let $y \in \pi_\alpha(\beta(t))$. Let $x := \beta(T(\alpha,\beta))$. The GIT says the projection of the subsegment of β between x and $\beta(t)$ has diameter at most $4C_\eta$, as does the projection of the subsegment of γ from z to $\beta(t)$. Thus $d(x,y) \leq 6C_\eta$ and $d(y,z) \leq 6C_\eta$. It follows that $|(\alpha(s)\cdot\beta(t))_o - d(o,y)| \leq 6C_\eta$, so:

$$|(\alpha(s)\cdot\beta(t))_o - T(\alpha,\beta)| \leq 12C_\eta \tag{1}$$

Consider the effect of replacing α with α'. For every t we have $d(\beta(t),\alpha) \geq t - T(\alpha,\beta) - 6C_\eta$, and α and α' have Hausdorff distance at most $5C_\eta$, so $d(\beta(t),\alpha') \geq t - T(\alpha,\beta) - 11C_\eta$. Since $d(\beta(T(\alpha',\beta)),\alpha') = 2C_\eta$ we have $T(\alpha',\beta) \leq T(\alpha,\beta) + 13C_\eta$. The argument is symmetric in α and α', so we conclude:

$$|T(\alpha,\beta) - T(\alpha',\beta)| \leq 13C_\eta \tag{2}$$

Now consider the effect of replacing β with β'. The Hausdorff distance between them is at most $5C_\zeta$. This does not admit any a priori bound in terms of C_η. However, eventually points of β are closer to β' than they are to α, so we can invoke strong contraction of α and the GIT, twice, to say:

$$\operatorname{diam} \pi_\alpha(\beta([T(\alpha,\beta),\infty))) \cup \pi_\alpha(\beta'([T(\alpha,\beta'),\infty))) \leq 9C_\eta$$

Which tells us:

$$|T(\alpha,\beta) - T(\alpha,\beta')| \leq 13C_\eta \tag{3}$$

Combining equations (1), (2), and (3), we have, for all s, s', t, t' sufficiently large, that $|(\alpha(s)\cdot\beta(t))_o - (\alpha'(s')\cdot\beta'(t'))_o| \leq 50C_\eta$. Thus, for any $\alpha \in \eta$ and $\beta \in \zeta$ and for all sufficiently large s and t we have $|(\alpha(s)\cdot\beta(t))_o - (\eta\cdot\zeta)_o| \leq 50C_\eta$. A further application of equation (1) completes the proof of the Claim with $K_\eta := 62C_\eta$. □

3. Pathological Examples

Construct a proper geodesic metric space X from rays α, β, and γ_i for $i \in \mathbb{N}$ as follows. Identify $\alpha(0)$ and $\beta(0)$, and take this to be the basepoint o. For each i connect $\gamma_i(0)$ to $\alpha(i)$ and $\beta(i)$ by segments of length 2^i. Then the γ_i are strongly contracting, and α and β are contracting rays whose contracting function ρ can be taken to be logarithmic. The essential point is that the projection of $\gamma_i(0)$ to $\alpha \cup \beta$ has diameter $2i$, while the distance from $\gamma_i(0)$ to $\alpha \cup \beta$ is 2^i.

The contracting boundary of X consists of one point for each of the rays α, β, and γ_i, which we denote $\alpha(\infty)$, $\beta(\infty)$, and $\gamma_i(\infty)$, respectively. Compute the Gromov products of boundary points: $(\alpha(\infty)\cdot\gamma_i(\infty))_o = i = (\beta(\infty)\cdot\gamma_i(\infty))_o$, while $(\alpha(\infty)\cdot\beta(\infty))_o = 0$. The sequence $(\gamma_i(\infty))_i$ converges to both $\alpha(\infty)$ and $\beta(\infty)$ in $\partial_c^{Gp}X$. In this example $\partial_c^{Gp}X$ is compact but not Hausdorff.

Now consider the space Y obtained from X by redefining, for each i, the length of the segment connecting $\gamma_i(0)$ to β to be $2^i - 2i$. The identity map is a quasi-isometry, but in the new metric $(\alpha(\infty)\cdot\gamma_i(\infty))_o = 0$. The sequence $(\gamma_i(\infty))_i$ does not converge to $\alpha(\infty)$ in $\partial_c^{Gp}Y$. Thus, $\partial_c\mathrm{Id}\colon \partial_c^{Gp}X \to \partial_c^{Gp}Y$ is not continuous.

Next, we construct a CAT(0) example. Let X' be the universal cover of the Euclidean plane minus a ball of radius one. Parameterize X' by polar coordinates $\mathbb{R}\times[1,\infty)$. Let $\alpha\colon [0,\infty)\to X' : t\mapsto (t,1)$ and $\beta\colon [0,\infty)\to X' : t\mapsto(-t,1)$. Each of these geodesic rays is π–strongly contracting.

Let X be the proper CAT(0) space obtained from X' by attaching, for each $i\in\mathbb{N}$, a geodesic ray γ_i with $\gamma_i(0) = (i, 2^i)\in X'$. These rays are also strongly contracting.

The contracting boundary of X consists of points corresponding to the $\gamma_i(\infty)$ and the two points $\alpha(\infty)$ and $\beta(\infty)$. Since α is strongly contracting, it follows that $(\alpha(\infty)\cdot\gamma_i(\infty))_o = i$ up to bounded error. Thus, the sequence $(\gamma_i(\infty))_i$ converges to $\alpha(\infty)$ in $\partial_c^{Gp}X$.

Let Y be the proper CAT(0) space obtained from X' by attaching, for each $i\in\mathbb{N}$, a geodesic ray γ_i' with $\gamma_i'(0) = (0, 2^i)\in X'$. Define ϕ to be the map $(t,r)\mapsto(t-\log_2(r), r)$ on X', so that $\phi(\gamma_i(0)) = \gamma_i'(0)$. This is a variation of the well-known logarithmic spiral quasi-isometry of the Euclidean plane. Extend ϕ to all of X by isometries $\gamma_i\to\gamma_i'$ for each i. This gives a quasi-isometry $\phi\colon X\to Y$, but points in $\partial_c^{Gp}Y$ are isolated, so $\partial_c\phi\colon \partial_c^{Gp}X\to\partial_c^{Gp}Y$ is not continuous.

Interesting open questions remain: If $\phi\colon X\to Y$ is a quasi-isometry between proper geodesic spaces that have cocompact isometry groups and such that X and Y are CAT(0) (or, more generally, CA), must $\partial_c\phi\colon \partial_c^{Gp}X\to\partial_c^{Gp}Y$ be a homeomorphism? Must $\partial_c^{Gp}X$ and $\partial_c^{Gp}Y$ be homeomorphic? Note that for visual boundaries of CAT(0) spaces the second question is much harder than the first [6].

References

[1] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Characterizations of Morse geodesics via superlinear divergence and sublinear contraction*. preprint. 2016. arXiv: 1601.01897.

[2] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Vol. 319. Grundlehren der mathematischen Wissenschaften. Berlin: Springer, 1999.

[3] Stephen M. Buckley and Simon L. Kokkendorff. *Comparing the Floyd and ideal boundaries of a metric space*. Trans. Amer. Math. Soc. **361**.2 (2009), pp. 715–734.

[4] Ruth Charney and Harold Sultan. *Contracting boundaries of* CAT(0) *spaces*. J. Topol. **8**.1 (2015), pp. 93–117.

[5] Matthew Cordes. *Morse Boundaries of Proper Geodesic Metric Spaces*. preprint. 2015. arXiv: 1502.04376.

[6] Christopher B. Croke and Bruce Kleiner. *Spaces with nonpositive curvature and their ideal boundaries*. Topology **39** (2000), pp. 549–556.

[7] M. Gromov. *Hyperbolic groups. Essays in group theory*. Vol. 8. Math. Sci. Res. Inst. Publ. New York: Springer, 1987, pp. 75–263.

[8] Harold Sultan. *Hyperbolic quasi-geodesics in CAT(0) spaces*. Geom. Dedicata **169**.1 (2014), pp. 209–224.

PAPER H

A metrizable topology on the contracting boundary of a group

Christopher H. Cashen and John M. Mackay
Transactions of the American Mathematical Society, 372 (2019), no. 3, 1555–1600.
DOI:10.1090/tran/7544

The 'contracting boundary' of a proper geodesic metric space consists of equivalence classes of geodesic rays that behave like rays in a hyperbolic space. We introduce a geometrically relevant, quasi-isometry invariant topology on the contracting boundary. When the space is the Cayley graph of a finitely generated group we show that our new topology is metrizable.

1. Introduction

There is a long history in geometry of attaching a 'boundary at infinity' or 'ideal boundary' to a space. When a group acts geometrically on a space we might wonder to what extent the group and the boundary of the space are related. In the setting of (Gromov) hyperbolic groups this relationship is very strong: the boundary is determined by the group, up to homeomorphism. In particular, the boundaries of all Cayley graphs of a hyperbolic group are homeomorphic, so it makes sense to call any one of these boundaries the boundary *of the group*. This is not true, for example, in the case of a group acting geometrically on a non-positively curved space: Croke and Kleiner [19] gave an example of a group acting geometrically on two different CAT(0) spaces with non-homeomorphic visual boundaries, so there is not a well-defined visual boundary associated to the group.

Charney and Sultan [14] sought to rectify this problem by defining a 'contracting boundary' for CAT(0) spaces. Hyperbolic boundaries and visual boundaries of CAT(0) spaces can be constructed as equivalence classes of geodesic rays emanating from a fixed basepoint. These represent the metrically distinct ways of 'going to infinity'. Charney and Sultan's idea was to restrict attention to ways of going to infinity in hyperbolic directions: They consider equivalence classes of geodesic rays that are 'contracting', which is a way of quantifying how hyperbolic such rays are. They topologize the resulting set using a direct limit construction, and show that this topology is preserved by quasi-isometries. However, their construction has drawbacks: basically, it has too many open sets. In general it is not first countable.

In this paper we define a bordification of a proper geodesic metric space by adding a contracting boundary with a quasi-isometry invariant topology. When the space is a Cayley graph of a finitely generated group, we prove that the topology on the boundary is metrizable, which is a significant improvement over the direct limit topology. (See Example 1.1 for a motivating example.) Furthermore, our topology more closely resembles the topology of the boundary of a hyperbolic space, which we hope will make it easier to work with.

Our contracting boundary consists of equivalence classes of 'contracting quasi-geodesics'. The definition of contraction we use follows that of Arzhantseva, Cashen, Gruber, and Hume [4];

The first author thanks the Isaac Newton Institute for Mathematical Sciences for support and hospitality during the program *Non-positive curvature: group actions and cohomology* where work on this paper was undertaken. This work is partially supported by EPSRC Grant Number EP/K032208/1 and by the Austrian Science Fund (FWF): P30487-N35.
The second author was supported in part by the National Science Foundation under Grant DMS-1440140 while visiting the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester, and in part by EPSRC grant EP/P010245/1.

this is weaker than that of Charney and Sultan, so our construction applies to more general spaces. For example, we get contracting quasi-geodesics from cyclic subgroups generated by non-peripheral elements of relatively hyperbolic groups [23], pseudo-Anosov elements of mapping class groups [32, 24], fully irreducible free group automorphisms [2], and generalized loxodromic elements of acylindrically hyperbolic groups [35, 20, 8, 39]. On CAT(0) spaces the two definitions agree, so our boundary is the same as theirs *as a set*, but our topology is coarser.

Cordes [16] has defined a 'Morse boundary' for proper geodesic metric spaces by applying Charney and Sultan's direct limit construction to the set of equivalence classes of Morse geodesic rays. This boundary has been further studied by Cordes and Hume [18], who relate it to the notion of 'stable subgroups' introduced by Durham and Taylor [25]; for a recent survey of these developments[1], see Cordes [15]. It turns out that our notion of contracting geodesic is equivalent to the Morse condition, and our contracting boundary agrees with the Morse boundary as a set, but, again, our topology is coarser.

If the underlying space is hyperbolic then all of these boundaries are homeomorphic to the Gromov boundary. At the other extreme, all of these boundaries are empty in spaces with no hyperbolic directions. In particular, it follows from work of Drutu and Sapir [23] that groups that are *wide*, that is, no asymptotic cone contains a cut point, will have empty contracting boundary. This includes groups satisfying a law: for instance, solvable groups or bounded torsion groups.

The boundary of a proper hyperbolic space can be topologized as follows. If ζ is a point in the boundary, an equivalence class of geodesic rays issuing from the chosen basepoint, we declare a small neighborhood of ζ to consist of boundary points η such that if $\alpha \in \zeta$ and $\beta \in \eta$ are representative geodesic rays then β closely fellow-travels α for a long time. In proving that this topology is invariant under quasi-isometries, hyperbolicity is used at two key points. The first is that quasi-isometries take geodesic rays uniformly close to geodesic rays. In general a quasi-isometry only takes a geodesic ray to a quasi-geodesic ray, but hyperbolicity implies that this is within bounded distance of a geodesic ray, with bound depending only on the quasi-isometry and hyperbolicity constants. The second use of hyperbolicity is to draw a clear distinction between fellow-travelling and not, which is used to show that the time for which two geodesics fellow-travel is roughly preserved by quasi-isometries. If α and β are non-asymptotic geodesic rays issuing from a common basepoint in a hyperbolic space, then closest point projection sends β to a bounded subset $\alpha([0, T_0])$ of α, and there is a transition in the behavior of β at time T_0. For $t < T_0$ the distance from $\beta(t)$ to α is bounded and the diameter of the projection of $\beta([0, t])$ to α grows like t. After this time β escapes *quickly* from α, that is, $d(\beta(t), \alpha)$ grows like $t - T_0$, and the diameter of the projection of $\beta([T_0, t])$ is bounded.

We recover the second point for non-hyperbolic spaces using the contraction property. Our definition of a *contracting* set Z, see Definition 3.2, is that the diameter of the projection of a ball tangent to Z is bounded by a function of the radius of the ball whose growth rate is less than linear. Essentially this means that sets far from Z have large diameter compared to the diameter of their projection. In contrast to the hyperbolic case, it is not true, in general, that if α is a contracting geodesic ray and β is a geodesic ray not asymptotic to α then β has bounded projection to α. However, we *can* still characterize the escape of β from α by the relation between the growth of the projection of $\beta([0, t])$ to α and the distance from $\beta(t)$ to α. The main technical tool we introduce is a divagation estimate that says if α is contracting and β is a quasi-geodesic then β cannot wander slowly away from α; if it is to escape, it must do so quickly. More precisely, once β exceeds a threshold distance from α, depending on the quasi-geodesic constants of β and the contraction function for α, then the distance from $\beta(t)$ to α grows superlinearly compared to the growth of the projection of $\beta([0, t])$ to α. In fact, for the purpose of proving that fellow-travelling time is roughly preserved by quasi-isometries it will be enough to know that the this relationship is at least a fixed linear function.

[1]In even more recent developments, Behrstock [6] produces interesting examples of right-angled Coxeter groups whose Morse boundaries contain a circle, and Charney and Murray [13] give conditions that guarantee that a homeomorphism between Morse boundaries of CAT(0) spaces is induced by a quasi-isometry.

The first point cannot be recovered, and, in fact, the topology as described above, using only geodesic rays, is not quasi-isometry invariant for non-hyperbolic spaces [12]. Instead, we introduce a finer topology that we call the *topology of fellow-travelling quasi-geodesics*. The idea is that η is close to ζ if all *quasi*-geodesics tending to η closely fellow-travel quasi-geodesics tending to ζ for a long time. See Definition 5.3 for a precise definition. Using our divagation estimates we show that this topology is quasi-isometry invariant.

The use of quasi-geodesic rays in our definition is quite natural in the setting of coarse geometry, since then the rays under consideration do not depend on the choice of a particular metric within a fixed quasi-isometry class. Geodesics, on the other hand, are highly sensitive to the choice of metric, and it is only the presence of a very strong hypothesis like global hyperbolicity that allows us to define a quasi-isometry invariant boundary topology using geodesics alone.

Example 1.1. Consider $H := \langle a, b \mid [a, b] = 1\rangle * \langle c\rangle$, which can be thought of as the fundamental group of a flat, square torus wedged with a circle. Let X be the universal cover, with basepoint o above the wedge point.

Connected components of the preimage of the torus are Euclidean planes isometrically embedded in X. Geodesic segments contained in such a plane behave more like Euclidean geodesics than hyperbolic geodesics. In fact, a geodesic ray α based at o is contracting if and only if there exists a bound B_α such that α spends time at most time B_α in any one of the planes. Let $\alpha(\infty)$ denote the equivalence class of this ray as a point in the contracting boundary.

In Charney and Sultan's topology, if $(\alpha^n)_{n\in\mathbb{N}}$ is a sequence of contracting geodesic rays with the B_{α^n} unbounded, then $(\alpha^n(\infty))$ is not a convergent sequence in the contracting boundary. Murray [33] uses this fact to show that the contracting boundary is not first countable.

In the topology of fellow-travelling quasi-geodesics it will turn out that $(\alpha^n(\infty))$ converges if and only if there exists a contracting geodesic α in X such that the projections of the α^n to geodesics in the Bass-Serre tree of H (with respect to the given free product splitting of H) converge to the projection of α.

From another point of view, H is hyperbolic relative to the Abelian subgroup $A := \langle a, b\rangle$. We show in Theorem 7.6 that this implies that there is a natural map from the contracting boundary of H to the Bowditch boundary of the pair (H, A), and, with the topology of fellow-travelling quasi-geodesics, that this map is a topological embedding. The embedding statement cannot be true for Charney and Sultan's topology, since it is not first countable.

After some preliminaries in Section 2, we define the contraction property and recall/prove some basic technical results in Section 3 concerning the behavior of geodesics relative to contracting sets. In Section 4 we extend these results to quasi-geodesics, and derive the key divagation estimates, see Corollary 4.4 and Lemma 4.7.

In Section 5 introduce the topology of fellow-travelling quasi-geodesics and show that it is first countable, Hausdorff, and regular. In Section 6 we prove that it is also quasi-isometry invariant.

We compare other possible topologies in Section 7.

In Section 8 we consider the case of a finitely generated group. In this case we prove that the contracting boundary is second countable, hence metrizable.

We also prove a weak version of North-South dynamics for the action of a group on its contracting boundary in Section 9, in the spirit of Murray's work [33].

Finally, in Section 10 we show that the contracting boundary of an infinite, finitely generated group is non-empty and compact if and only if the group is hyperbolic.

We thank the referee for a careful reading of our paper.

2. Preliminaries

Let X be a metric space with metric d. For $Z \subset X$, define:

- $N_r Z := \{x \in X \mid \exists z \in Z,\ d(z, x) < r\}$

- $N_r^c Z := \{x \in X \mid \forall z \in Z,\ d(z,x) \geq r\}$
- $\bar{N}_r Z := \{x \in X \mid \exists z \in Z,\ d(z,x) \leq r\}$
- $\bar{N}_r^c Z := \{x \in X \mid \forall z \in Z,\ d(z,x) > r\}$

For $L \geq 1$ and $A \geq 0$, a map $\phi\colon\ (X, d_X) \to (X', d_{X'})$ is an (L, A)–*quasi-isometric embedding* if for all $x,\ y \in X$:

$$\frac{1}{L} d_X(x,y) - A \leq d_{X'}(\phi(x), \phi(y)) \leq L d_X(x,y) + A$$

If, in addition, $\bar{N}_A \phi(X) = X'$ then ϕ is an (L, A)–*quasi-isometry*. A *quasi-isometry inverse* $\bar{\phi}$ of a quasi-isometry $\phi\colon\ X \to X'$ is a quasi-isometry $\bar{\phi}\colon\ X' \to X$ such that the compositions $\phi \circ \bar{\phi}$ and $\bar{\phi} \circ \phi$ are both bounded distance from the identity map on the respective space.

A *geodesic* is an isometric embedding of an interval. A *quasi-geodesic* is a quasi-isometric embedding of an interval. If $\alpha\colon\ I \to X$ is a quasi-geodesic, we often use α_t to denote $\alpha(t)$, and conflate α with its image in X. When I is of the form $[a, b]$ or $[a, \infty)$ we will assume, by precomposing α with a translation of the domain, that $a = 0$. We use $\alpha + \beta$ and $\bar{\alpha}$ to denote concatenation and reversal, respectively.

A metric space is *geodesic* if every pair of points can be connected by a geodesic.

A metric space if *proper* if closed balls are compact.

It is often convenient to improve quasi-geodesics to be continuous, which can be accomplished by the following lemma.

Lemma 2.1 (Taming quasi-geodesics [11, Lemma III.H.1.11]). *If X is a geodesic metric space and $\gamma\colon\ [a, b] \to X$ is an (L, A)–quasi-geodesic then there exists a continuous $(L, 2(L+A))$–quasi-geodesic γ' such that $\gamma_a = \gamma'_a$, $\gamma_b = \gamma'_b$ and the Hausdorff distance between γ and γ' is at most $L + A$.*

Proof. Define γ' to agree with γ at the endpoints and at integer points of $[a, b]$, and then connect the dots by geodesic interpolation. □

A subspace Z of a geodesic metric space X is A–*quasi-convex* for some $A \geq 0$ if every geodesic connecting points in Z is contained in $\bar{N}_A Z$.

If f and g are functions then we say $f \preceq g$ if there exists a constant $C > 0$ such that $f(x) \leq Cg(Cx + c) + C$ for all x. If $f \preceq g$ and $g \preceq f$ then we write $f \asymp g$.

We will give a detailed account of the contracting property in the next section, but let us first take a moment to recall alternate characterizations, which will prove useful later in the paper.

A subspace Z of a metric space X is μ–*Morse* for some $\mu : [1, \infty) \times [0, \infty) \to \mathbb{R}$ if for every $L \geq 1$ and every $A \geq 0$, every (L, A)–quasi-geodesic with endpoints in Z is contained in $\bar{N}_{\mu(L,A)} Z$. We say Z is *Morse* if there exists μ such that it is μ–Morse. It is easy to see that the property of being Morse is invariant under quasi-isometries. In particular, a subset of a finitely generated group G is Morse in one Cayley graph of G if and only if it is Morse in every Cayley graph of G. Thus, we can speak of a *Morse subset* of G without specifying a finite generating set.

A set Z is called *t–recurrent*[2], for $t \in (0, 1/2)$, if for every $C \geq 1$ there exists $D \geq 0$ such that if p is a path with endpoints x and y on Z such that the ratio of the length of p to the distance between its endpoints is at most C, then there exists a point $z \in Z$ such that $d(p, z) \leq D$ and $\min\{d(z, x),\ d(z, y)\} \geq td(x, y)$. The set Z is called *recurrent* if it is t–recurrent for every $t \in (0, 1/2)$.

Theorem 2.2. *Let Z be a subset of a geodesic metric space X. The following are equivalent:*

(1) Z is Morse.
(2) Z is contracting.
(3) Z is recurrent.
(4) There exists $t \in (0, 1/2)$ such that Z is t-recurrent.

[2]This characterization was introduced in [22] with $t = 1/3$ for Z a quasi-geodesic. The idea is that a short curve must pass near the 'middle third' of the subsegment of Z connecting its endpoints. The property, again only for quasi-geodesics, but for variable t, is called 'middle recurrence' in [3].

Moreover, each of the equivalences are 'effective', in the sense that the defining function of one property determines the defining functions of each of the others.

Proof. The equivalence of (1) and (2) is proved in [4]. That (3) implies (4) is obvious. The implications '(2) implies (3)' and '(4) implies (1)' are proved in [3] for the case that Z is a quasi-geodesic, but their proofs go through with minimal change for arbitrary subsets Z. □

3. Contraction

Definition 3.1. We call a function ρ *sublinear* if it is non-decreasing, eventually non-negative, and $\lim_{r\to\infty}\rho(r)/r = 0$.

Definition 3.2. Let X be a proper geodesic metric space. Let Z be a closed subset of X, and let $\pi_Z\colon X \to 2^Z : x \mapsto \{z \in Z \mid d(x,z) = d(x,Z)\}$ be closest point projection to Z. Then, for a sublinear function ρ, we say that Z is *ρ–contracting* if for all x and y in X:

$$d(x,y) \leq d(x,Z) \implies \operatorname{diam} \pi_Z(x) \cup \pi_Z(y) \leq \rho(d(x,Z))$$

We say Z is *contracting* if there exists a sublinear function ρ such that Z is ρ–contracting. We say a collection of subsets $\{Z_i\}_{i\in\mathcal{I}}$ is *uniformly contracting* if there exists a sublinear function ρ such that for every $i \in \mathcal{I}$ the set Z_i is ρ–contracting.

We shorten π_Z to π when Z is clear from context.

Let us stress that the closest point projection map is set-valued, and there is no bound on the diameter of image sets other than that implied by the definition.

In a tree every convex subset is ρ–contracting where ρ is identically 0. More generally, in a hyperbolic space a set is contracting if and only if it is quasi-convex. In fact, in this case more is true: the contraction function is bounded in terms of the hyperbolicity and quasi-convexity constants. We call a set *strongly contracting* if it is contracting with bounded contraction function.

The more general Definition 3.2 was introduced by Arzhantseva, Cashen, Gruber, and Hume to characterize Morse geodesics in small cancellation groups [5].

The concept of strong contraction (sometimes simply called 'contraction' in the literature) has been studied before, notably by Minsky [32] to describe axes of pseudo-Anosov mapping classes in Teichmüller space, by Bestvina and Fujiwara [9] to describe axes of rank-one isometries of CAT(0) spaces (see also Sultan [40]), and by Algom-Kfir [2] to describe axes of fully irreducible free group automorphisms acting on Outer Space.

Masur and Minsky [31] introduced a different notion of contraction that requires the existence of constants A and B such that:

$$d(x,y) \leq d(x,Z)/A \implies \operatorname{diam} \pi_Z(x) \cup \pi_Z(y) \leq B$$

This is satisfied, for example, by axes of pseudo-Anosov elements in the mapping class group (as opposed to Teichmüller space). Some authors refer to this property as 'contraction', eg [7, 24, 1]. It is not hard to show that this version implies the version in Definition 3.2 with the contraction function ρ being logarithmic.

We now recall some further results about contracting sets in a geodesic metric space X.

Lemma 3.3 ([4, Lemma 6.3]). *Given a sublinear function ρ and a constant $C \geq 0$ there exists a sublinear function $\rho' \asymp \rho$ such that if $Z \subset X$ and $Z' \subset X$ have Hausdorff distance at most C and Z is ρ–contracting then Z' is ρ'–contracting.*

Theorem 3.4 (Geodesic Image Theorem [4, Theorem 7.1]). *For $Z \subset X$, there exists a sublinear function ρ so that Z is ρ–contracting if and only if there exists a sublinear function ρ' and a constant κ_ρ so that for every geodesic segment γ, with endpoints denoted x and y, if $d(\gamma, Z) \geq \kappa_\rho$ then* $\operatorname{diam} \pi(\gamma) \leq \rho'(\max\{d(x,Z),\ d(y,Z)\})$. *Moreover ρ' and κ_ρ depend only on ρ and vice-versa, with $\rho' \asymp \rho$.*

An easy consequence is that there exists a κ'_ρ such that if γ is a geodesic segment with endpoints at distance at most κ_ρ from a ρ–contracting set Z then $\gamma \subset \bar{N}_{\kappa'_\rho}(Z)$.

The following is a special case of [4, Proposition 8.1].

Lemma 3.5. *Given a sublinear function ρ and a constant $C \geq 0$ there exists a constant B such that if α and β are ρ–contracting geodesics such that their initial points α_0 and β_0 satisfy $d(\alpha_0, \beta_0) = d(\alpha, \beta) \leq C$ then $\alpha \cup \beta$ is B–quasi-convex.*

The next two lemmas are easy-to-state generalizations of results that are known for strong contraction. The proofs are rather tedious, due to the weak hypotheses, so we postpone them until after Lemma 3.8.

Lemma 3.6. *Given a sublinear function ρ there is a sublinear function $\rho' \asymp \rho$ such that every subsegment of a ρ–contracting geodesic is ρ'–contracting.*

Lemma 3.7. *Given a sublinear function ρ there is a sublinear function $\rho' \asymp \rho$ such that if α and β are ρ–contracting geodesic rays or segments such that $\gamma := \bar{\alpha} + \beta$ is geodesic, then γ is ρ'–contracting.*

Given $C \geq 0$ a *geodesic C–almost triangle* is a trio of geodesics $\alpha^i\colon [a_i, b_i] \to X$, for $i \in \{0, 1, 2\}$ and $a_i \leq 0 \leq b_i \in \mathbb{R} \cup \{-\infty, \infty\}$, such that for each $i \in \{0, 1, 2\}$, with scripts taken modulo 3, we have:

- $b_i < \infty$ if and only if $a_{i+1} > -\infty$.
- If b_i and a_{i+1} are finite then $d(\alpha^i_{b_i}, \alpha^{i+1}_{a_{i+1}}) \leq C$.
- If b_i and a_{i+1} are not finite then $\alpha^i_{[0,\infty)}$ and $\bar{\alpha}^{i+1}_{[0,\infty)} = \alpha^{i+1}_{(-\infty,0]}$ are asymptotic.

Lemma 3.8. *Given a sublinear function ρ and constant $C \geq 0$ there is a sublinear function $\rho' \asymp \rho$ such that if α, β, and γ are a geodesic C–almost triangle and α and β are ρ–contracting then γ is ρ'–contracting.*

Proof. First suppose α, β, and γ are segments. By Lemma 3.5, there exists a B depending only on ρ and C such that $\alpha \cup \beta$ is B–quasi–convex. Thus, we can replace $\alpha \cup \beta$ by a single geodesic segment δ whose endpoints are C–close to the endpoints of γ. Furthermore, δ is a union of two subsegments, one of which has endpoints within distance B of α, and the other of which has endpoints within distance B of β. Consequently, by Theorem 3.4 there exists B' so that these two subsegments are B'–Hausdorff equivalent to subsegments of α and of β, respectively. Applying Lemma 3.6, Lemma 3.3, and Lemma 3.7, there is a $\rho'' \asymp \rho$ depending on ρ and B' such that δ is ρ''–contracting. Theorem 3.4 implies that since γ and δ are close at their endpoints, they stay close along their entire lengths, so their Hausdorff distance is determined by ρ'' and C, hence by ρ and C. Applying Lemma 3.3 again, we conclude γ is ρ'–contracting with $\rho' \asymp \rho'' \asymp \rho$ depending only on ρ and C.

In the case of an ideal triangle, where not all three sides are segments, replace C by $\max\{C, \kappa_\rho\}$. Theorem 3.4 implies that if, say, γ and $\bar{\alpha}$ have asymptotic tails then the set of points γ that come κ_ρ–close to α is unbounded. Truncate the triangle at such a C–close pair of points. Doing the same for other ideal vertices, we get a C–almost triangle to which we can apply the previous argument and conclude that a subsegment of γ is ρ'–contracting. Since γ comes κ_ρ–close to α on an unbounded set, we can repeat the argument for larger and larger almost triangles approximating α, β, γ, and find that every subsegment of γ is contained in a ρ'–contracting subsegment, which implies that γ itself is ρ'–contracting. □

Definition 3.9. If Z is a subset of $\mathbb{R}$ define the *interval of Z*, $\operatorname{invl}(Z)$, to be the smallest closed interval containing Z. If $\gamma\colon I \to X$ is a geodesic and Z is a subset of γ let $\operatorname{invl}(Z) := \gamma(\operatorname{invl}(\gamma^{-1}(Z)))$.

Proof of Lemma 3.6. Let $\gamma\colon I \to X$ be a ρ–contracting geodesic. Let $J := [j_0, j_1]$ be a subinterval of I. Let $\rho'' \asymp \rho$ be the function given by Theorem 3.4, and let κ'_ρ be the constant defined there. We claim it suffices to take $\rho'(r) := 2(2\kappa'_\rho + \rho''(2r) + \rho(2r))$.

First we show that if $\pi_{\gamma_I}(x)$ misses γ_J then $\pi_{\gamma_I}(x)$ is relatively close to one of the endpoints of γ_J. This is automatic if $\operatorname{diam} \gamma_J \leq \rho(d(x, \gamma_J))$, so assume not. With this assumption, $\pi_{\gamma_I}(x)$ cannot contain points on both sides of γ_J, that is, if $\gamma^{-1}(\pi_{\gamma_I}(x))$ contains a point less than j_0 then it does not also contain one greater than j_1, and vice versa. Suppose that $\gamma^{-1}(\pi_{\gamma_I}(x))$ is contained in $(-\infty, j_0)$. Let β be a geodesic from x to a point y in $\pi_{\gamma_J}(x)$. There exists a first time s such that $d(\beta_s, \gamma_I) = \kappa_\rho$. By Theorem 3.4, $\operatorname{diam} \pi_{\gamma_I}(\beta|_{[0,s]}) \leq \rho''(d(x, \gamma_I))$. Suppose that $\gamma_{j_0} \in \operatorname{invl}(\pi_{\gamma_I}(\beta|_{[0,s]}))$. Then there is a first time $s' \in [0, s]$, such that $\pi_I(\beta_{s'})$ contains a point in $\gamma_{[j_0,\infty)}$. By the assumption on the diameter of γ_J, we actually have $\pi_{\gamma_I}(\beta_{s'}) \cap \gamma_J \neq \emptyset$, so $y \in \pi_{\gamma_J}(\beta_{s'}) \subset \pi_{\gamma_I}(\beta_{s'}) \subset \pi_{\gamma_I}(\beta|_{[0,s]})$ and $\operatorname{diam} \gamma_{j_0} \cup \pi_{\gamma_J}(x) \leq \rho''(d(x, \gamma_I))$. Otherwise, if $\gamma_{j_0} \notin \operatorname{invl}(\pi_{\gamma_I}(\beta|_{[0,s]}))$, then let $t > s$ be the first time such that $\gamma_{j_0} \in \operatorname{invl} \pi_{\gamma_I}(\beta|_{[s,t]})$. Again, $y \in \pi_{\gamma_I}(\beta_t)$. Since the points of β after β_s, are contained in $\bar{N}_{\kappa'_\rho}\gamma$, for all small $E > 0$ we have $\operatorname{diam} \pi_{\gamma_I}\beta_{t-E} \cup \pi_{\gamma_I}\beta_t \leq E + 2\kappa'_\rho$. Therefore, $d(\gamma_{j_0}, y) \leq d(y, \pi_{\gamma_I}(\beta_{t-E})) \leq E + 2\kappa'_\rho$, for all sufficiently small E. We conclude:

$$\operatorname{diam} \gamma_{j_0} \cup \pi_{\gamma_J}(x) \leq \max\{2\kappa'_\rho, \rho''(d(x, \gamma_I))\} \tag{1}$$

Now suppose x and y are points such that $d(x, y) \leq d(x, \gamma_J)$. Note that $d(y, \gamma_J) \leq 2d(x, \gamma_J)$. We must show $\operatorname{diam} \pi_{\gamma_J}(x) \cup \pi_{\gamma_J}(y)$ is bounded by a sublinear function of $d(x, \gamma_J)$. There are several cases, depending on whether $\pi_{\gamma_I}(x)$ and $\pi_{\gamma_I}(y)$ hit γ_J.

Case 1: $\gamma_J \cap \pi_{\gamma_I}(x) \neq \emptyset$ and $\gamma_J \cap \pi_{\gamma_I}(y) \neq \emptyset$. In this case $\pi_{\gamma_J}(x) \subset \pi_{\gamma_I}(x)$, and likewise for y, so:

$$\operatorname{diam} \pi_{\gamma_J}(x) \cup \pi_{\gamma_J}(y) \leq \operatorname{diam} \pi_{\gamma_I}(x) \cup \pi_{\gamma_I}(y) \leq \rho(d(x, \gamma_I)) = \rho(d(x, \gamma_J))$$

Case 2: $\gamma^{-1}(\pi_{\gamma_I}(x)) < j_0$ and $\gamma^{-1}(\pi_{\gamma_I}(y)) < j_0$. By (1) twice:

$$\begin{aligned}\operatorname{diam} \pi_{\gamma_J}(x) \cup \pi_{\gamma_J}(y) &\leq 2 \max\{2\kappa'_\rho, \rho''(d(x, \gamma_I)), \rho''(d(y, \gamma_I))\} \\ &\leq 2 \max\{2\kappa'_\rho, \rho''(2d(x, \gamma_J))\}\end{aligned}$$

Case 3: $\gamma^{-1}(\pi_{\gamma_I}(y)) < j_0$ and $\gamma_J \cap \pi_{\gamma_I}(x) \neq \emptyset$. In this case $\pi_{\gamma_J}(x) \subset \pi_{\gamma_I}(x)$ and $d(x, y) \leq d(x, \gamma_J) = d(x, \gamma_I)$, so $\operatorname{diam} \pi_{\gamma_I}(y) \cup \pi_{\gamma_J}(x) \leq \rho(d(x, \gamma_J))$. By hypothesis, $\gamma_{j_0} \in \operatorname{invl}(\pi_{\gamma_I}(y) \cup \pi_{\gamma_J}(x))$, and by (1): $\operatorname{diam} \gamma_{j_0} \cup \pi_{\gamma_J}(y) \leq \max\{2\kappa'_\rho, \rho''(2d(x, \gamma_J))\}$. Thus:

$$\operatorname{diam} \pi_{\gamma_J}(x) \cup \pi_{\gamma_J}(y) \leq \max\{\rho(d(x, \gamma_J)), 2\kappa'_\rho, \rho''(2d(x, \gamma_J))\}$$

Case 4: $\gamma^{-1}(\pi_{\gamma_I}(x)) < j_0$ and $\pi_{\gamma_I}(y) \cap \gamma_{[j_0,\infty)} \neq \emptyset$. If $j_1 - j_0 \leq 2\rho(2d(x, \gamma_J))$ then there is nothing more to prove, so assume not. Let β be a geodesic from x to y. For all $z \in \beta$:

$$d(x, z) + d(z, y) = d(x, y) \leq d(x, \gamma_J) \leq d(x, z) + d(z, \gamma_J)$$

This implies $d(z, y) \leq d(z, \gamma_J)$. Let z be the first point on β such that $\gamma^{-1}(\pi_{\gamma_I}(z))$ contains a point greater than or equal to j_0. By the hypothesis on $|J|$, $\gamma^{-1}(\pi_{\gamma_I}(z)) < j_1$. This means $\operatorname{diam} \pi_{\gamma_J}(z) \cup \pi_{\gamma_J}(y)$ is controlled by one of the previous cases, and it suffices to control $\operatorname{diam} \pi_{\gamma_J}(x) \cup \pi_{\gamma_J}(z)$.

We know from (1) that $\pi_{\gamma_J}(x)$ is $\max\{2\kappa'_\rho, \rho''(d(x, \gamma_J))\}$-close to γ_{j_0}, so it suffices to control $\operatorname{diam} \gamma_{j_0} \cup \pi_{\gamma_J}(z)$. Take a point $w \neq z$ on β before z such that $d(z, w) \leq d(z, \gamma_I)$. By hypothesis, $\gamma_{j_0} \in \operatorname{invl} \pi_{\gamma_I}(w) \cup \pi_{\gamma_I}(z)$, but $\operatorname{diam} \pi_{\gamma_I}(w) \cup \pi_{\gamma_I}(z) \leq \rho(d(z, \gamma_I)) = \rho(d(z, \gamma_J)) \leq \rho(2d(x, \gamma_J))$.

Up to symmetric arguments, this exhausts all the cases. □

Proof of Lemma 3.7. Let α and β be ρ–contracting geodesic segments or rays with $\alpha_0 = \beta_0$ such that $\gamma := \bar{\alpha} + \beta$ is geodesic.

First suppose that x is a point such that $\pi_\gamma(x) \cap \alpha \neq \emptyset$ and $\pi_\gamma(x) \cap \beta \neq \emptyset$. Let δ be a geodesic from x to $\alpha_0 = \beta_0$. Recall from Theorem 3.4 that once δ enters the κ_ρ–neighborhood of either α or β then it cannot leave the κ'_ρ–neighborhood. Thus, δ intersects at most one of $\bar{N}_{\kappa_\rho}\alpha \setminus N_{2\kappa'_\rho}\alpha_0$ or $\bar{N}_{\kappa_\rho}\beta \setminus N_{2\kappa'_\rho}\beta_0$. Without loss of generality, suppose δ does not intersect $\bar{N}_{\kappa_\rho}\beta \setminus N_{2\kappa'_\rho}\beta_0$. Let t be the first time such that $d(\delta_t, \beta) = \kappa_\rho$. Then $d(\delta_t, \beta_0) \leq 2\kappa'_\rho$ and, by Theorem 3.4, there is a sublinear $\rho'' \asymp \rho$ such that $\operatorname{diam} \pi_\beta(\delta|_{[0,t]}) \leq \rho''(d(x, \beta)) = \rho''(d(x, \gamma))$. In particular, this means $\operatorname{diam} \pi_\beta(x) \cup \beta_0 \leq \operatorname{diam} \pi_\beta(\delta) \leq 4\kappa'_\rho + \rho''(d(x, \gamma))$. Now let δ' be a geodesic from x to a point $x' \in \pi_\beta(x)$, and project δ' to α. Since $\bar{\alpha} + \beta$ is geodesic, $\operatorname{diam} \pi_\alpha(x) \cup \alpha_0 \leq \operatorname{diam} \pi_\alpha\delta' \leq \rho''(\max\{d(x, \alpha), d(x', \alpha)\})$ by Theorem 3.4. We have already established that $d(x', \alpha) \leq 4\kappa'_\rho + \rho''(d(x, \gamma))$. Since ρ'' grows sublinearly, $d(x, \alpha) > 4\kappa'_\rho + \rho''(d(x, \gamma))$ except for $d(x, \alpha)$ less than some bound depending only on ρ and ρ''. We conclude that there is a sublinear function $\rho''' \asymp \rho$ depending only on ρ such that $\operatorname{diam} \pi_\alpha(x) \cup \alpha_0 \leq \rho'''(d(x, \gamma))$ and $\operatorname{diam} \pi_\beta(x) \cup \beta_0 \leq \rho'''(d(x, \gamma))$, hence $\operatorname{diam} \pi_\gamma(x) \leq 2\rho'''(d(x, \gamma))$.

Now suppose $x, y \in X$ are points such that $d(x,y) \leq d(x,\gamma)$. There are several cases according to where $\pi_\gamma(x)$ and $\pi_\gamma(y)$ lie.

Case 1: $\pi_\gamma(x) \cap \alpha \neq \emptyset \neq \pi_\gamma(y) \cap \alpha$. Then $d(x,y) \leq d(x,\gamma) = d(x,\alpha)$, so contraction for α implies $\operatorname{diam} \pi_\alpha(x) \cup \pi_\alpha(y) \leq \rho(d(x,\alpha)) = \rho(d(x,\gamma))$. There are four sub-cases to check, according to whether $\pi_\gamma(x)$ and $\pi_\gamma(y)$ hit β. These are easy to check, with the worst bound being $\operatorname{diam} \pi_\gamma(x) \cup \pi_\gamma(y) \leq \rho(d(x,\gamma)) + 2\rho'''(2d(x,\gamma))$.

Case 2: $\pi_\gamma(x) \cap \beta = \emptyset = \pi_\gamma(y) \cap \alpha$. Let δ be a geodesic from x to y. Let w be the first point on δ such that $\pi_\gamma(w) \cap \beta \neq \emptyset$. Then $d(w,\alpha) = d(w,\beta) = d(w,\gamma) \leq 2d(x,\gamma)$ and $d(y,\beta) = d(y,\gamma) \leq 2d(x,\gamma)$. We can apply that α is ρ-contracting to the pair x, w since $d(x,w) \leq d(x,y) \leq d(x,\gamma) = d(x,\alpha)$. Likewise, we can apply that β is ρ-contracting to w, y since $d(x,w) + d(w,y) = d(x,y) \leq d(x,\gamma) \leq d(x,w) + d(w,\gamma)$ so $d(w,y) \leq d(w,\gamma)$. We conclude:

$$\begin{aligned}\operatorname{diam} \pi_\gamma(x) \cup \pi_\gamma(y) &\leq \operatorname{diam} \pi_\alpha(x) \cup \pi_\alpha(w) + \operatorname{diam} \pi_\gamma(w) \\ &\qquad + \operatorname{diam} \pi_\beta(w) \cup \pi_\beta(y) \\ &\leq \rho(d(x,\alpha)) + 2\rho'''(d(w,\gamma)) + \rho(d(w,\beta)) \\ &\leq \rho(d(x,\gamma)) + 2\rho'''(2d(x,\gamma)) + \rho(2d(x,\gamma))\end{aligned}$$

By symmetry these two cases cover all possibilities, so it suffices to define $\rho'(r) := 2\rho(2r) + 2\rho'''(2r)$. □

4. Contraction and Quasi-geodesics

In this section we explore the behavior of a quasi-geodesic ray based at a point in a contracting set Z. The main conclusion is that such a ray can stay close to Z for an arbitrarily long time, but once it exceeds a certain threshold distance depending on the quasi-geodesic constants and the contraction function then the ray must escape Z at a definite linear rate.

Definition 4.1. Given a sublinear function ρ and constants $L \geq 1$ and $A \geq 0$, define:

$$\kappa(\rho, L, A) := \max\{3A, 3L^2, 1 + \inf\{R > 0 \mid \forall r \geq R,\ 3L^2\rho(r) \leq r\}\}$$

Define:

$$\kappa'(\rho, L, A) := (L^2 + 2)(2\kappa(\rho, L, A) + A)$$

Remark 4.2. For the rest of the paper κ and κ' always refer to the functions defined in Definition 4.1. We use them frequently and without further reference.

This definition implies that for $r \geq \kappa(\rho, L, A)$ we have:

$$r - L^2\rho(r) - A \geq \frac{1}{3}r \geq L^2\rho(r) \tag{2}$$

An inspection of the proof of [4, Theorem 7.1] gives that $\kappa(\rho, 1, 0) \geq \kappa_\rho$ and $\kappa'(\rho, 1, 0) \geq \kappa'_\rho$, so the results of the previous section still hold using $\kappa(\rho, 1, 0)$ and $\kappa'(\rho, 1, 0)$. Enlarging the constants lets us give unified proofs for geodesics and quasi-geodesics.

Theorem 4.3 (Quasi-geodesic Image Theorem). *Let $Z \subset X$ be ρ–contracting. Let $\beta\colon [0,T] \to X$ be a continuous (L,A)–quasi-geodesic segment. If $d(\beta, Z) \geq \kappa(\rho, L, A)$ then:*

$$\operatorname{diam} \pi(\beta_0) \cup \pi(\beta_T) \leq \frac{L^2+1}{L^2}(A + d(\beta_T, Z)) + \frac{L^2-1}{L^2} d(\beta_0, Z) + 2\rho(d(\beta_0, Z))$$

The proof generalizes the proof of the Geodesic Image Theorem to work for quasi-geodesics. We typically apply the result when $d(\beta_T, Z) = \kappa(\rho, L, A)$, in which case the theorem says that for fixed ρ, L, and A the projection diameter of β is bounded in terms of $d(\beta_0, Z)$. In particular, when β is geodesic, or, more generally, when $L = 1$, the bound is sublinear in $d(\beta_0, Z)$, and we recover a version of the Geodesic Image Theorem. With a little more work we can prove this stronger statement for quasi-geodesics as well. Although we do not need it in this paper, the stronger version may be of independent interest, so we include a proof at the end of this section (see Theorem 4.9).

Proof of Theorem 4.3. Let $t_0 := 0$. For each $i \in \mathbb{N}$ in turn, let t_{i+1} be the first time such that $d(\beta_{t_i}, \beta_{t_{i+1}}) = d(\beta_{t_i}, Z)$, or set $t_{i+1} = T$ if no such time exists. Let j be the first index such that $d(\beta_{t_j}, \beta_T) \leq d(\beta_{t_j}, Z)$.

$$\begin{aligned} T &= T - t_j + \sum_{i=0}^{j-1}(t_{i+1} - t_i) \\ &\geq \frac{1}{L}\left(d(\beta_{t_j}, \beta_T) - d(\beta_{t_j}, Z) + \sum_{i=0}^{j}(d(\beta_{t_i}, Z) - A)\right) \\ &\geq \frac{1}{L}\left(-d(\beta_T, Z) + \sum_{i=0}^{j}(d(\beta_{t_i}, Z) - A)\right) \end{aligned}$$

On the other hand:

$$\begin{aligned} \frac{T}{L} - A &\leq d(\beta_0, \beta_T) \\ &\leq d(\beta_0, Z) + \operatorname{diam} \pi(\beta_0) \cup \pi(\beta_T) + d(Z, \beta_T) \\ &\leq d(\beta_0, Z) + d(\beta_T, Z) + \sum_{i=0}^{j} \rho(d(\beta_{t_i}, Z)) \end{aligned}$$

Combining these gives:

$$\begin{aligned} &\sum_{i=1}^{j}\left(d(\beta_{t_i}, Z) - L^2\rho(d(\beta_{t_i}, Z)) - A\right) \\ &\qquad \leq d(\beta_T, Z) + L^2\left(A + d(\beta_0, Z) + d(\beta_T, Z)\right) \\ &\qquad\qquad - \left(d(\beta_0, Z) - L^2\rho(d(\beta_0, Z)) - A\right) \end{aligned}$$

By (2), the left-hand side is at least $L^2 \sum_{i=1}^{j} \rho(d(\beta_{t_i}, Z))$, so:

$$\begin{aligned} \operatorname{diam} \pi(\beta_0) \cup \pi(\beta_T) &\leq \sum_{i=0}^{j} \rho(d(\beta_{t_i}, Z)) \\ &\leq \frac{L^2+1}{L^2}\left(A + d(\beta_T, Z)\right) + \frac{L^2-1}{L^2} d(\beta_0, Z) + 2\rho(d(\beta_0, Z)) \end{aligned}$$ □

Corollary 4.4. *Let Z be ρ–contracting and let β be a continuous (L, A)–quasi-geodesic ray with $d(\beta_0, Z) \leq \kappa(\rho, L, A)$. There are two possibilities:*

(1) *The set $\{t \mid d(\beta_t, Z) \leq \kappa(\rho, L, A)\}$ is unbounded and β is contained in the $\kappa'(\rho, L, A)$–neighborhood of Z.*
(2) *There exists a last time T_0 such that $d(\beta_{T_0}, Z) = \kappa(\rho, L, A)$ and:*

$$(\star) \qquad \forall t, \quad d(\beta_t, Z) \geq \frac{1}{2L}(t - T_0) - 2(A + \kappa(\rho, L, A))$$

Proof. Let $\kappa := \kappa(\rho, L, A)$. Let $[a, b]$ be a maximal interval such that $d(\beta_t, Z) \geq \kappa$ for $t \in [a, b]$ and $d(\beta_a, Z) = d(\beta_b, Z) = \kappa$.

For $t \in [a, b]$ we have $d(\beta_t, Z) \leq \kappa + L \cdot (b - a)/2 + A$. Since β is quasi-geodesic:

$$(b - a) \leq L(A + d(\beta_a, \beta_b)) \leq L(A + 2\kappa + \operatorname{diam} \pi(\beta_a) \cup \pi(\beta_b))$$

Theorem 4.3 implies:

$$\begin{aligned} \operatorname{diam} \pi(\beta_a) \cup \pi(\beta_b) &\leq \frac{L^2+1}{L^2}(A + \kappa) + \frac{L^2-1}{L^2}\kappa + \frac{2\kappa}{3L^2} \\ &= \frac{L^2+1}{L^2}A + \frac{6L^2+2}{3L^2}\kappa \end{aligned}$$

Putting these estimates together yields:

$$d(\beta_t, Z) < (L^2+2)(2\kappa + A) = \kappa'(\rho, L, A)$$

Thus, once β leaves the $\kappa'(\rho, L, A)$–neighborhood of Z it can never return to the $\kappa(\rho, L, A)$–neighborhood of Z. If $\{t \mid d(\beta_t, Z) \leq \kappa\}$ is unbounded then β never leaves the $\kappa'(\rho, L, A)$–neighborhood of Z.

Suppose now that there does exist some last time T_0 such that $d(\beta_{T_0}, Z) = \kappa$. Any segment $\beta_{[T_0,t]}$ stays outside $N_\kappa Z$, so apply Theorem 4.3 to see:

$$\begin{aligned}\frac{t-T_0}{L} - A &\leq d(\beta_t, \beta_{T_0})\\ &\leq d(\beta_t, Z) + \operatorname{diam}\pi(\beta_t)\cup\pi(\beta_{T_0}) + \kappa\\ &\leq \frac{6L^2-1}{3L^2} d(\beta_t, Z) + \frac{L^2+1}{L^2}(A+\kappa) + \kappa\end{aligned}$$

Thus:

$$d(\beta_t, Z) \geq \frac{3L}{6L^2-1}(t-T_0) - \frac{6L^2+3}{6L^2-1}(A+\kappa)$$ □

Lemma 4.5. *Suppose α is a continuous, ρ–contracting (L, A)–quasi-geodesic and β is a continuous (L, A)–quasi-geodesic ray such that $d(\alpha_0, \beta_0) \leq \kappa(\rho, L, A)$. If there are r, $s \in [0, \infty)$ such that $d(\alpha_r, \beta_s) \leq \kappa(\rho, L, A)$ then $d_{Haus}(\alpha_{[0,r]}, \beta_{[0,s]}) \leq \kappa'(\rho, L, A)$. If $\alpha_{[0,\infty)}$ and $\beta_{[0,\infty)}$ are asymptotic then their Hausdorff distance is at most $\kappa'(\rho, L, A)$.*

Proof. Corollary 4.4 (1) reduces the asymptotic case to the bounded case and shows that β is contained in $\bar{N}_{\kappa'(\rho,L,A)}\alpha$.

For the other direction, suppose that (a, b) is a maximal open subinterval of the domain of α such that $\alpha_{(a,b)} \cap \pi_\alpha(\beta \cap \bar{N}_{\kappa(\rho,L,A)}\alpha) = \emptyset$. For subsegments of β contained in $\bar{N}_{\kappa(\rho,L,A)}\alpha$ the projection to α has jumps of size at most $2\kappa(\rho, L, A)$. For subsegments of β outside $\bar{N}_{\kappa(\rho,L,A)}\alpha$ the largest possible gap in the projection is bounded by Theorem 4.3 by:

$$\begin{aligned}d(\alpha_a, \alpha_b) &\leq \frac{L^2+1}{L^2}(A + \kappa(\rho, L, A)) + \frac{L^2-1}{L^2}\kappa(\rho, L, A) + 2\rho(\kappa(\rho, L, A))\\ &\leq \frac{L^2+1}{L^2}(A + \kappa(\rho, L, A)) + \frac{L^2-1}{L^2}\kappa(\rho, L, A) + \frac{2\kappa(\rho, L, A)}{3L^2}\\ &= \frac{L^2+1}{L^2}A + \frac{6L^2+2}{6L^2}\cdot 2\kappa(\rho, L, A)\end{aligned} \tag{3}$$

This is greater than $2\kappa(\rho, L, A)$.

In either case, for $c \in (a, b)$ we have:

$$\begin{aligned}d(\alpha_c, \beta) &\leq \kappa(\rho, L, A) + \min\{d(\alpha_a, \alpha_c),\ d(\alpha_b, \alpha_c)\}\\ &\leq \kappa(\rho, L, A) + A + L\frac{b-a}{2}\\ &\leq \kappa(\rho, L, A) + A + L\frac{LA + Ld(\alpha_a, \alpha_b)}{2}\end{aligned} \tag{4}$$

Substitute (3) into (4) and observe that the resulting bound is less than $\kappa'(\rho, L, A)$, which was defined to be $(L^2+2)(A + 2\kappa(\rho, L, A))$. □

Lemma 4.6. *If α is a ρ–contracting geodesic ray and β is a continuous (L, A)–quasi-geodesic ray asymptotic to α with $\alpha_0 = \beta_0$ then β is ρ'–contracting where $\rho' \asymp \rho$ depends only on ρ, L, and A.*

Proof. Lemma 4.5 says the Hausdorff distance between α and β is bounded in terms of ρ, L, and A, so the claim follows from Lemma 3.3. □

The next lemma gives the key divagation estimate, which gives us lower bounds on fellow-travelling distance.

LEMMA 4.7. *Let α be a ρ-contracting geodesic ray, and let β be a continuous (L, A)–quasi-geodesic ray with $\alpha_0 = \beta_0 = o$. Given some R and J, suppose there exists a point $x \in \alpha$ with $d(x, o) \geq R$ and $d(x, \beta) \leq J$. Let y be the last point on the subsegment of α between o and x such that $d(y, \beta) = \kappa(\rho, L, A)$. There is a constant $M \leq 2$ and a function $\lambda(\phi, p, q)$ defined for sublinear ϕ, $p \geq 1$, and $q \geq 0$ such that λ is monotonically increasing in p and q and:*

$$d(x, y) \leq MJ + \lambda(\rho, L, A)$$

Thus:

$$d(o, y) \geq R - MJ - \lambda(\rho, L, A)$$

PROOF. If $d(x, \beta) \leq \kappa(\rho, L, A)$ then $y = x$ and we are done. Otherwise, let a be the last time such that β_a is $\kappa(\rho, L, A)$–close to α between o and x, and let $y' \in \alpha$ be the last point of α with $d(\beta_a, y') = \kappa(\rho, L, A)$. Note $d(y, x) \leq d(y', x)$.

Now let b be the first time such that $d(\beta_b, x) = J$. The subsegment $\beta_{[a,b]}$ stays outside $N_{\kappa(\rho,L,A)}\alpha$. Pick a geodesic from β_b to x and let w be the first point such that $d(w, \alpha) = \kappa(\rho, 1, 0)$. Pick $z \in \pi(\beta_b)$ and $v \in \pi(w)$, and let $W := d(\beta_b, w)$, $Y := d(y', z)$, $Z := d(z, v)$, and $X := d(v, x)$, see Figure 1.

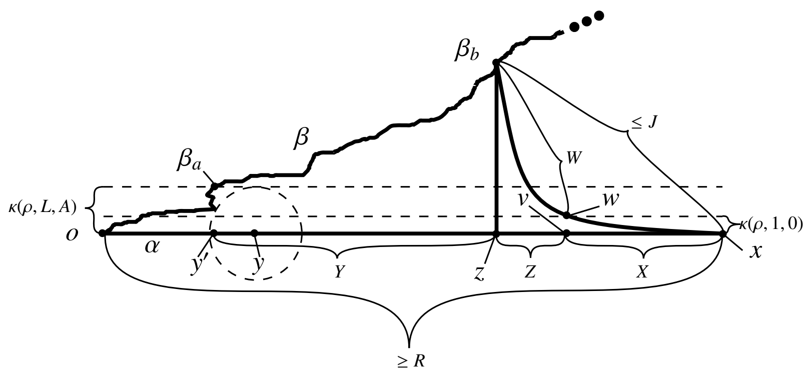

FIGURE 1. Setup for Lemma 4.7

We have $W \geq d(\beta_b, \alpha) - Z - \kappa(\rho, 1, 0)$, and $X \leq J - W + \kappa(\rho, 1, 0)$, so:

$$\begin{aligned} d(y', x) &\leq X + Y + Z \\ &\leq Y + Z + J - d(\beta_b, \alpha) + Z + 2\kappa(\rho, 1, 0) \end{aligned}$$

Apply Theorem 4.3 to the subsegment of γ between γ_a to γ_b to bound Y. Apply Theorem 4.3 to the subsegment of the chosen geodesic from β_b to x between β_b and w to bound Z. (Note in the latter case that we are applying Theorem 4.3 to a geodesic, so use $L = 1$ and $A = 0$ for this case.) Combining these bounds for Y and Z with the bound on $d(x, y')$ above yields:

$$d(y', x) \leq J + 6\kappa(\rho, 1, 0) + \frac{L^2 + 1}{L^2}(A + \kappa(\rho, L, A)) - \frac{1}{L^2} d(\beta_b, \alpha) + 6\rho(d(\beta_b, \alpha))$$

Now use the facts that $\rho(d(\beta_b, \alpha)) \leq \frac{d(\beta_b,\alpha)}{3L^2}$ and $d(\beta_b, \alpha) \leq J$ to achieve:

$$\begin{aligned} d(y, x) \leq d(y', x) &\leq \frac{L^2 + 1}{L^2} J + 6\kappa(\rho, 1, 0) + \frac{L^2 + 1}{L^2}(A + \kappa(\rho, L, A)) \\ &\leq 2J + 6\kappa(\rho, 1, 0) + 2(A + \kappa(\rho, L, A)) \end{aligned}$$

Set $M := 2$ and $\lambda(\phi, p, q) := 6\kappa(\phi, 1, 0) + 2(q + \kappa(\phi, p, q))$. □

Here is an application of Lemma 4.7 that we will use in Section 9.

LEMMA 4.8. *Given a sublinear function ρ and constants $L \geq 1$, $A \geq 0$ there exist constants $L' \geq 1$ and $A' \geq 0$ such that if α is a ρ-contracting geodesic ray or segment and β is a continuous (L, A)–quasi-geodesic ray not asymptotic to α with $\alpha_0 = \beta_0 = o$, then we obtain a continuous (L', A')–quasi-geodesic by following α backward until α_{s_0}, then following a geodesic from α_{s_0} to β_{t_0}, then following β, where β_{t_0} is the last point of β at distance $\kappa(\rho, L, A)$ from α, and where α_{s_0} is the last point of α at distance $\kappa(\rho, L, A)$ from β_{t_0}.*

Proof. Define $\kappa := \kappa(\rho, L, A)$ and M and $\lambda := \lambda(\rho, L, A)$ from Lemma 4.7. Recall $M \leq 2 \leq 2L$. It suffices to take $A' := \left(\frac{(4L+1)\kappa+\lambda}{4L} + A\right)$ and $L' := 4L$. Since we have constructed a concatenation of three quasi-geodesic segments, it suffices to check that points on different segments are not too close together. Since $A' > A + \kappa$ we may ignore the short middle segment. Thus, we need to check for $s \geq s_0$ and $t \geq t_0$ that $d(\alpha_s, \beta_t) \geq \frac{s-s_0+t-t_0+\kappa}{L'} - A'$.

For such s and t, let $x := \alpha_s$, $y := \alpha_{s_0}$, and $z := \beta_t$. By Lemma 4.7, $s - s_0 = d(x, y) \leq Md(x, z) + \lambda < 2Ld(x, z) + \lambda$. Choose some point $z' \in \pi_\alpha(z)$. By Corollary 4.4 (★) we have $d(z, x) \geq d(z, z') \geq \frac{t-t_0}{2L} - 2(A + \kappa)$. Now average these two lower bounds for $d(x, z)$:

$$\begin{aligned} d(\alpha_s, \beta_t) = d(x, z) &\geq \frac{1}{2}\left(\frac{s - s_0}{2L} - \frac{\lambda}{2L} + \frac{t - t_0}{2L} - 2(A + \kappa)\right) \\ &\geq \frac{s - s_0 + t - t_0 + \kappa}{4L} - \left(\frac{\lambda}{4L} + \frac{4L + 1}{4L}\kappa + A\right) \end{aligned}$$ □

To close this section we give the stronger formulation of the Quasi-geodesic Image Theorem:

Theorem 4.9. *Given a sublinear function ρ and constants $L \geq 1$ and $A \geq 0$ there is a sublinear function ρ' such that if Z is ρ–contracting and $\beta\colon [0, T] \to X$ is a continuous (L, A)–quasi-geodesic segment with $d(\beta, Z) = d(\beta_T, Z) = \kappa(\rho, L, A)$ then* $\operatorname{diam} \pi(\beta_0) \cup \pi(\beta_T) \leq \rho'(d(\beta_0, Z))$.

Proof. Define $\rho'(r) := \sup_\beta \operatorname{diam} \pi(\beta_0) \cup \pi(\beta_T)$ where the supremum is taken over all continuous (L, A)–quasi-geodesic segments β such that $d(\beta, Z) = \kappa(\rho, L, A)$ is realized at one endpoint of β and the other endpoint is at distance at most r from Z. Suppose that ρ' is not sublinear, so suppose $\limsup_{r\to\infty} \rho'(r)/r = 2\epsilon > 0$. Then there exists a sequence $(r_i) \to \infty$ such that for each i there exists a continuous (L, A)–quasi-geodesic segment $\beta^{(i)}\colon [0, T_i] \to X$ with $d(\beta^{(i)}_{T_i}, Z) = \kappa(\rho, L, A)$ and $d(\beta^{(i)}_0, Z) \leq r_i$ and $\operatorname{diam} \pi(\beta^{(i)}_0) \cup \pi(\beta^{(i)}_{T_i}) \geq \epsilon r_i$, so that $\operatorname{diam} \pi(\beta^{(i)}_0) \cup \pi(\beta^{(i)}_{T_i}) \geq \epsilon d(\beta^{(i)}_0, Z)$.

For $n \in \mathbb{N}$ define κ_n large enough so that for all $r \geq \kappa_n$ we have $r - L^2\rho(r) - A \geq \frac{1}{3}r \geq nL^2\rho(r)$ (recalling (2), $\kappa_1 = \kappa(\rho, L, A)$). The proof of Theorem 4.3 shows that if a continuous (L, A)–quasi-geodesic segment stays outside the κ_n–neighborhood of Z then:

$$\begin{aligned} \operatorname{diam} \pi(\beta_0) \cup \pi(\beta_T) &\leq \frac{L^2 + 1}{nL^2}(A + d(\beta_T, Z)) + \frac{L^2 - 1}{nL^2}d(\beta_0, Z) + \frac{n + 1}{n}\rho(d(\beta_0, Z)) \\ &\leq \frac{1}{n}(2A + 2d(\beta_T, Z) + d(\beta_0, Z)) \end{aligned} \tag{5}$$

For $\epsilon > 0$ as above, choose $n \in \mathbb{N}$ large enough that $n\epsilon > 2$. For all sufficiently large i we have that $\operatorname{diam} \pi(\beta^{(i)}_0) \cup \pi(\beta^{(i)}_{T_i}) > 2A + 2\kappa_1 + \kappa_n$. By (5) for $n = 1$, we have $d(\beta^{(i)}_0, Z) > \kappa_n$. Let $s_i > 0$ be the first time such that $d(\beta^{(i)}_{s_i}, Z) = \kappa_n$.

$$\begin{aligned} \epsilon d(\beta^{(i)}_0, Z) &\leq \operatorname{diam} \pi(\beta^{(i)}_{T_i}) \cup \pi(\beta^{(i)}_0) \\ &\leq \operatorname{diam} \pi(\beta^{(i)}_{T_i}) \cup \pi(\beta^{(i)}_{s_i}) + \operatorname{diam} \pi(\beta^{(i)}_{s_i}) \cup \pi(\beta^{(i)}_0) \\ &\leq (2A + 2\kappa_1 + \kappa_n) + \frac{1}{n}\left(2A + 2\kappa_n + d(\beta^{(i)}_0, Z)\right) \\ &\leq (2A + 2\kappa_1 + \kappa_n) + \frac{\epsilon}{2}\left(2A + 2\kappa_n + d(\beta^{(i)}_0, Z)\right) \end{aligned}$$

Solving for $d(\beta^{(i)}_0, Z)$, we find that it is bounded, independent of i. By (5) for $n = 1$, this would bound $\operatorname{diam} \pi(\beta^{(i)}_0) \cup \pi(\beta^{(i)}_{T_i})$, independent of i, whereas we have assumed $\operatorname{diam} \pi(\beta^{(i)}_0) \cup \pi(\beta^{(i)}_{T_i}) \geq \epsilon r_i \to \infty$. This is a contradiction, so we conclude $\lim_{r\to\infty} \rho'(r)/r = 0$. □

5. The contracting boundary and the topology of fellow-travelling quasi-geodesics

Definition 5.1. Let X be a proper geodesic metric space with basepoint o. Define $\partial_c X$ to be the set of contracting quasi-geodesic rays based at o modulo Hausdorff equivalence.

Lemma 5.2. *For each $\zeta \in \partial_c X$:*

- *The set of contracting geodesic rays in ζ is non-empty.*
- *There is a sublinear function:*

$$\rho_\zeta(r) := \sup_{\alpha,x,y} \operatorname{diam} \pi_\alpha(x) \cup \pi_\alpha(y)$$

Here the supremum is taken over geodesics $\alpha \in \zeta$ and points x and y such that $d(x, y) \leq d(x, \alpha) \leq r$.

- *Every geodesic in ζ is ρ_ζ–contracting.*

Proof. By definition, ζ is an equivalence class of contracting quasi-geodesic rays, so there exists some ρ'–contracting (L, A)–quasi-geodesic ray $\beta \in \zeta$ based at o. Since X is proper, a sequence of geodesic segments connecting o to β_i for $i \in \mathbb{N}$ has a subsequence that converges to a geodesic α'. By Theorem 3.4, all of these geodesic segments, hence α' as well, are contained in a bounded neighborhood of β, with bound depending only on ρ', so there do exist geodesics asymptotic to β. Furthermore, Corollary 4.4 implies that geodesic rays asymptotic to β have uniformly bounded Hausdorff distance from β, with bound depending on ρ', L, and A. By Lemma 3.3, all such geodesics are ρ''–contracting for some $\rho'' \asymp \rho'$ depending on ρ', L, and A.

The function ρ_ζ is non-decreasing and bounds projection diameters by definition. The fact that there exists a sublinear function ρ'' such that all geodesics in ζ are ρ''–contracting implies $\rho_\zeta \leq \rho''$, so ρ_ζ is also sublinear. □

Definition 5.3. Let X be a proper geodesic metric space. Take $\zeta \in \partial_c X$. Fix a geodesic ray $\alpha^\zeta \in \zeta$. For each $r \geq 1$ define $U(\zeta, r)$ to be the set of points $\eta \in \partial_c X$ such that for all $L \geq 1$ and $A \geq 0$ and every continuous (L, A)–quasi-geodesic ray $\beta \in \eta$ we have $d(\beta, \alpha^\zeta \cap N_r^c o) \leq \kappa(\rho_\zeta, L, A)$.

Informally, $\eta \in U(\zeta, r)$ means that inside the ball of radius r about the basepoint quasi-geodesics in η fellow-travel α^ζ just as closely as quasi-geodesics in ζ do. Alternatively, quasi-geodesics in η do not escape from α^ζ until after they leave the ball of radius r about the basepoint.

Definition 5.4. Define the *topology of fellow-travelling quasi-geodesics* on $\partial_c X$ by:

$$\mathcal{FQ} := \{U \subset \partial_c X \mid \forall \zeta \in U,\ \exists r \geq 1,\ U(\zeta, r) \subset U\}$$

The contracting boundary equipped with this topology is denoted $\partial_c^{\mathcal{FQ}} X$.

We do not assume that the sets $U(\zeta, r)$ are open in the topology $\mathcal{FQ}$. Indeed, from the definition it is not even clear that $U(\zeta, r)$ is a neighborhood of ζ, but we will show that this is the case.

Proposition 5.5. *$\mathcal{FQ}$ is a topology on $\partial_c X$, and for each $\zeta \in \partial_c X$ the collection $\{U(\zeta, n) \mid n \in \mathbb{N}\}$ is a neighborhood basis at ζ.*

Corollary 5.6. *$\partial_c^{\mathcal{FQ}} X$ is first countable.*

Observation 5.1. Suppose $\eta \notin U(\zeta, r)$. By definition, for some L and A there exists a continuous (L, A)–quasi-geodesic $\beta \in \eta$ such that $d(\beta, \alpha^\zeta \cap N_r^c o) > \kappa(\rho_\zeta, L, A)$. Since $o \in \beta$, this is not possible if $\kappa(\rho_\zeta, L, A) \geq r$. Thus, in light of Definition 4.1, the quasi-geodesic β witnessing $\eta \notin U(\zeta, r)$ must be an (L, A)–quasi-geodesic with $L^2 < r/3$ and $A < r/3$.

The proof of Proposition 5.5 depends on two lemmas. The first is a recombination result for quasi-geodesics. Its key feature is that the quasi-geodesic constants of the result depend only on the quasi-geodesic constants of the input, not on the contraction function.

Lemma 5.7 (Tail wagging). *Let ρ be a sublinear function. Let $L \geq 1$ and $A \geq 0$. Let $T \geq 11\kappa'(\rho, L, A)$ and $S \geq T + 6\kappa'(\rho, L, A) + 6\kappa'(\rho, 1, 0)$. Suppose α is a ρ–contracting geodesic ray based at o, γ is a continuous (L, A)–quasi-geodesic ray based at o such that $d(\gamma, \alpha_{[T,\infty)}) \leq \kappa(\rho, L, A)$, and β is a geodesic ray based at o such that $d(\beta, \alpha_{[S,\infty)}) \leq \kappa(\rho, 1, 0)$. Then there are continuous $(2L + 1, A)$–quasi-geodesic rays that agree with γ until a point within distance $11\kappa'(\rho, L, A)$ of α_T and share tails with α and β, respectively.*

Proof. We construct the quasi-geodesic ray that shares a tail with β. The construction for the α–tail is similar, but with easier estimates.

Let $T'' := T - 3\kappa'(\rho, 1, 0) - \kappa'(\rho, L, A)$. Let T' be the first time at which γ comes within distance $\kappa'(\rho, L, A)$ of $\alpha_{[T'',\infty)}$. Let S' be such that $d(\beta_{S'}, \alpha_{[S,\infty)}) \leq \kappa(\rho, 1, 0)$. Let $t_0 \leq T'$ and $r_0 \geq S'$ be times such that $d(\gamma_{t_0}, \beta_{r_0}) = d(\gamma_{[0,T']}, \beta_{[S',\infty)})$, and let δ be a geodesic from γ_{t_0} to β_{r_0}. There are times b, c, b', and c' such that $d(\gamma_{t_0}, \alpha_b)$, $d(\gamma_{T'}, \alpha_c) \leq \kappa'(\rho, L, A)$, $d(\beta_{b'}, \alpha_b)$, $d(\beta_{c'}, \alpha_c) \leq \kappa'(\rho, 1, 0)$. For any $t \leq t_0$ there exist a and a' such that $d(\gamma_t, \alpha_a) \leq \kappa'(\rho, L, A)$ and $d(\alpha_a, \beta_{a'}) \leq \kappa(\rho, 1, 0)$. See Figure 2.

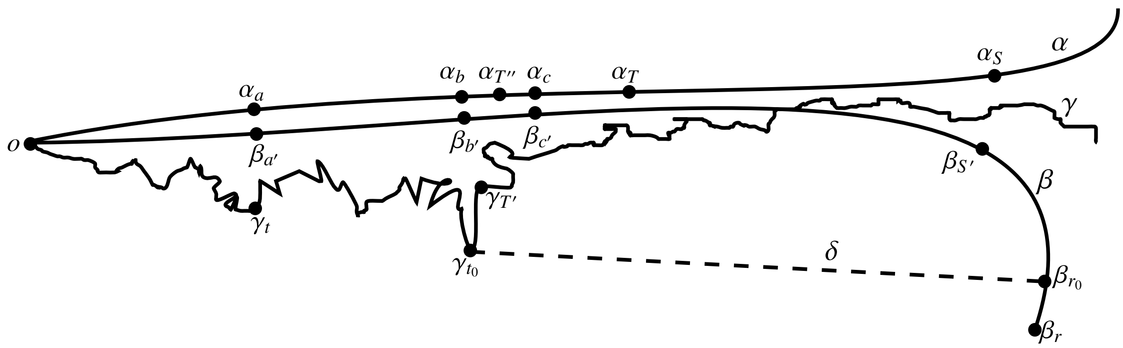

Figure 2. Wagging the tail of γ.

The desired quasi-geodesic ray is $\gamma_{[0,t_0]} + \delta + \beta_{[r_0,\infty)}$.

First, we verify $d(\gamma_{t_0}, \alpha_T) \leq 11\kappa'(\rho, L, A)$. The definitions of t_0 and r_0 demand $d(\gamma_{t_0}, \beta_{r_0}) \leq d(\gamma_{T'}, \beta_{S'})$. The left-hand side is at least $S' - b' - (\kappa'(\rho, L, A) + \kappa'(\rho, 1, 0))$, while the right-hand side is no more than $S' - c' + (\kappa'(\rho, L, A) + \kappa'(\rho, 1, 0))$, so $c' - b' \leq 2(\kappa'(\rho, L, A) + \kappa'(\rho, 1, 0))$. Since $T'' \leq c \leq T'' + 2\kappa'(\rho, L, A)$ we have $d(\alpha_c, \alpha_T) \leq \kappa'(\rho, L, A) + 3\kappa'(\rho, 1, 0)$. Together, these allow us to estimate:

$$\begin{aligned} d(\gamma_{t_0}, \alpha_T) &\leq d(\gamma_{t_0}, \beta_{b'}) + d(\beta_{b'}, \beta_{c'}) + d(\beta_{c'}, \alpha_c) + d(\alpha_c, \alpha_T) \\ &\leq (\kappa'(\rho, L, A) + \kappa'(\rho, 1, 0)) + c' - b' + \kappa'(\rho, 1, 0) \\ &\qquad + (\kappa'(\rho, L, A) + 3\kappa'(\rho, 1, 0)) \\ &\leq 7\kappa'(\rho, 1, 0) + 4\kappa'(\rho, L, A) \leq 11\kappa'(\rho, L, A) \end{aligned}$$

Next we verify the quasi-geodesic constants. Since we have a concatenation of quasi-geodesics, we only need to check that points on different pieces are not closer than they ought to be with respect to the parameterization.

First we claim $\gamma_{[0,t_0]} + \delta$ is an (L', A)–quasi-geodesic for $L' := 2L + 1$. This is true for $\gamma_{[0,t_0]}$ and δ individually. Suppose there are $0 \leq t < t_0$ and $0 < u \leq |\delta|$ such that $d(\gamma_t, \delta_u) < \frac{t_0 - t + u}{L'} - A$. Now, $d(\delta_u, \gamma_t) \geq d(\delta_u, \gamma_{t_0}) = u$, which implies $u < \frac{L'}{L'-1}(\frac{t_0 - t}{L'} - A)$. But then:

$$\begin{aligned} \frac{t_0 - t}{L} - A &\leq d(\gamma_t, \gamma_{t_0}) \leq d(\gamma_t, \delta_u) + d(\delta_u, \gamma_{t_0}) \\ &\leq \left(\frac{t_0 - t + u}{L'} - A\right) + u \end{aligned}$$

Plugging in the value for L' and the bound for u yields a contradiction.

The same argument shows $\delta + \beta_{[r_0,\infty)}$ is a $(3, 0)$–quasi-geodesic.

Now consider points γ_t and β_r for $t \leq t_0$ and $r \geq r_0$.

$$\begin{aligned}
d(\gamma_t,\beta_r) &\geq r - a' - (\kappa'(\rho,L,A) + \kappa'(\rho,1,0))\\
&= r - r_0 + r_0 - b' + b' - a' - (\kappa'(\rho,L,A) + \kappa'(\rho,1,0))\\
&\geq r - r_0 + d(\gamma_{t_0},\beta_{r_0}) + d(\gamma_t,\gamma_{t_0}) - 4(\kappa'(\rho,L,A) + \kappa'(\rho,1,0))\\
&\geq \frac{t_0 - t + r - r_0 + |\delta|}{2L+1} - A + |\delta|\frac{2L}{2L+1} - 4(\kappa'(\rho,L,A) + \kappa'(\rho,1,0))
\end{aligned}$$

Thus, the ray we have constructed is a $(2L + 1, A)$–quasi-geodesic, since $|\delta| \geq 6(\kappa'(\rho,L,A) + \kappa'(\rho,1,0))$, as we now verify:

$$\begin{aligned}
|\delta| &\geq r_0 - b' - (\kappa'(\rho,L,A) + \kappa'(\rho,1,0))\\
&\geq r_0 - b - (\kappa'(\rho,L,A) + 2\kappa'(\rho,1,0))\\
&\geq S' - T'' - (\kappa'(\rho,L,A) + 2\kappa'(\rho,1,0))\\
&\geq S - T'' - (\kappa'(\rho,L,A) + 3\kappa'(\rho,1,0))\\
&\geq 6(\kappa'(\rho,L,A) + \kappa'(\rho,1,0))
\end{aligned}$$

□

Lemma 5.8. *For every sublinear function ρ and $r \geq 1$ there exists a number $\psi(\rho,r) > r$ such that for every $R \geq \psi(\rho,r)$ and every $\zeta \in \partial_c X$ such that $\rho_\zeta \leq \rho$ we have that for every $\eta \in U(\zeta,R)$ there exists an R' such that $U(\eta,R') \subset U(\zeta,r)$.*

Proof. It suffices to take $\psi(\rho,r) := r + M\kappa(\rho, 2\sqrt{r/3} + 1, r/3) + \lambda(\rho, \sqrt{r/3}, r/3)$, where M and λ are as in Lemma 4.7.

Suppose $R \geq \psi(\rho,r)$ and ζ is a point in $\partial_c X$ such that $\rho_\zeta \leq \rho$. Suppose that $\eta \in U(\zeta,R)$ with $\eta \neq \zeta$. Let T_0 be the last time such that $d(\alpha^\eta_{T_0}, \alpha^\zeta) = \kappa(\rho_\zeta, 1, 0)$. Set:

$$R' := T_0 + 2\kappa'(\rho_\zeta, 2\sqrt{r/3} + 1, r/3) + 4\kappa(\rho_\zeta,1,0) + 28\kappa'(\rho_\eta, \sqrt{r/3}, r/3) + 6\kappa'(\rho_\eta,1,0)$$

Suppose that there exists a point $\xi \in U(\eta,R')$ such that $\xi \notin U(\zeta,r)$. The latter implies there exists an $L \geq 1$ and $A \geq 0$ and a continuous (L,A)–quasi-geodesic $\gamma \in \xi$ such that $d(\gamma, N^c_r o \cap \alpha^\zeta) > \kappa(\rho_\zeta, L, A)$. By Observation 5.1, we have L^2, $A < r/3$. Set $\alpha := \alpha^\eta, \beta := \alpha^\xi$, $T := T_0 + 4\kappa(\rho_\zeta,1,0) + 22\kappa'(\rho_\eta, \sqrt{r/3}, r/3) + 2\kappa'(\rho_\zeta, 2\sqrt{r/3}+1, r/3)$, and $S := R' \geq T + 6\kappa'(\rho_\eta, L, A) + 6\kappa'(\rho_\eta,1,0)$. Apply Lemma 5.7 to α, β, γ, T, and S to produce a continuous $(2L+1, A)$–quasi-geodesic $\delta \in \eta$ that agrees with γ at least until a point z in the ball of radius $11\kappa'(\rho_\eta, L, A)$ about α^η_T.

By Corollary 4.4 (★) we have $d(\alpha_T, \alpha^\zeta) \geq (T - T_0)/2 - 2\kappa(\rho_\zeta,1,0)$, which implies $d(z,\alpha^\zeta) \geq \kappa'(\rho_\zeta, 2L+1, A)$, so by point z the ray δ has already escaped α^ζ and can never return to its $\kappa(\rho_\zeta, 2L+1, A)$–neighborhood. Therefore, the only points of δ in the $\kappa(\rho_\zeta, 2L+1, A)$–neighborhood of α^ζ are those that were contributed by γ. By construction, γ does not come $\kappa(\rho_\zeta, L, A)$–close to α^ζ outside the ball of radius r. By applying Lemma 4.7, we see that δ is a witness to $\eta \notin U(\zeta,R)$. This is a contradiction, so $U(\eta,R') \subset U(\zeta,r)$. □

Proof of Proposition 5.5. For every $\zeta \in \partial_c X$ and $1 \leq r < r'$ we have $\zeta \in U(\zeta,r') \subset U(\zeta,r)$. The nesting is immediate from Definition 5.3, and $\zeta \in U(\zeta,r)$ by Corollary 4.4. Now it is easy to see that $\mathcal{FQ}$ is a topology. That a set of the form $U(\zeta,r)$ is a neighborhood of ζ in this topology follows from showing that the set

$$U := \{\eta \in U(\zeta,r) \mid \exists R_\eta,\ U(\eta,R_\eta) \subset U(\zeta,r)\}$$

is open, since then $\zeta \in U \subset U(\zeta,r)$. Now if $\eta \in U$ then there exists R_η so that $U(\eta,R_\eta) \subset U(\zeta,r)$. Lemma 5.8 says that for all $\xi \in U(\eta, \psi(\rho_\eta, R_\eta))$ there exists R' with $U(\xi,R') \subset U(\eta,R_\eta) \subset U(\zeta,r)$. Therefore $U(\eta,\psi(\rho_\eta,R_\eta)) \subset U$ and so U is open. □

From this proof we observe the following consequence.

Corollary 5.9. *For every $\zeta \in \partial_c X$ and $r \geq 1$ there exists an open set U such that $U(\zeta, \psi(\rho_\zeta, r)) \subset U \subset U(\zeta,r)$.*

Proposition 5.10. *The topology $\mathcal{FQ}$ does not depend on the choice of basepoint or on the choices of the representative geodesic rays for each point in $\partial_c X$.*

Proof. Let C be the set of contracting quasi-geodesic rays based at o and let C' be the set of contracting quasi-geodesic rays based at o'. There is a map $\phi\colon C \to C'$ by prefixing $\gamma \in C$ with a chosen geodesic segment from o' to o. The map ϕ clearly induces a bijection $\partial_c\phi$ between contracting boundaries of X with respect to different basepoints, and the inverse map can be achieved by simply prefixing quasi-geodesic rays by a geodesic from o to o'. We check that $\partial_c\phi$ is an open map. For $\zeta \in \partial_c^{\mathcal{FQ}}X$ and $r \geq 1$ we show for sufficiently large R that $U'(\partial_c\phi(\zeta), R) \subset \partial_c\phi(U(\zeta, r))$, where $U'(\partial_c\phi(\zeta), R)$ denotes the appropriate neighborhood of $\partial_c\phi(\zeta)$ defined with o' as basepoint.

Let $\alpha := \alpha^\zeta$ be the reference geodesic for ζ based at o, and let α' be the reference geodesic for $\partial_c\phi(\zeta)$ based at o'. Then α' is bounded Hausdorff distance from α. Suppose α is ρ–contracting and α' is ρ'–contracting. Theorem 3.4 implies that α' eventually comes within distance $\kappa(\rho, 1, 0)$ of α, and Theorem 4.3 implies that this first happens at some time no later than $d(o', \alpha) + 3\kappa(\rho, 1, 0) + 2\rho(d(o', \alpha))$. After that time α' remains in the $\kappa'(\rho, 1, 0)$–neighborhood of α. Assume $R > d(o', \alpha) + 3\kappa(\rho, 1, 0) + 2\rho(d(o', \alpha))$.

Assume further that $R > r + 2d(o, o')$ and suppose $\eta \in U'(\partial_c\phi(\zeta), R)$. Let $\gamma \in \partial_c\phi^{-1}(\eta)$ be an arbitrary continuous (L, A)–quasi-geodesic. Our goal is to show that if R is chosen sufficiently large with respect to ρ, ρ', and r, then such a γ must come within distance $\kappa(\rho, L, A)$ of α outside $N_r o$. We then conclude $\partial_c\phi^{-1}(U'(\partial_c\phi(\zeta), R)) \subset U(\zeta, r)$. By Observation 5.1, it suffices to consider the case L^2, $A < r/3$.

Now, $\gamma' := \phi(\gamma) \in \eta$ is a continuous $(L, A + 2d(o, o'))$–quasi-geodesic. Since $\eta \in U'(\partial_c\phi(\zeta), R)$ there exists a point $x' \in \alpha'$ such that $d(\gamma', x') \leq \kappa(\rho', L, A + 2d(o, o'))$ and $d(x', o') \geq R$. The first restriction on R implies there is a point $x \in \alpha$ such that $d(x, x') \leq \kappa'(\rho, 1, 0)$, so $d(\gamma', x) \leq \kappa'(\rho, 1, 0) + \kappa(\rho', L, A + 2d(o, o'))$. We also have $d(x, o) \geq R - \kappa'(\rho, 1, 0) - d(o, o')$. Assuming further that $R > 2d(o, o') + 2\kappa'(\rho, 1, 0) + \kappa(\rho', L, A + 2d(o, o'))$, we have that the point of γ' close to x is actually a point of γ. Let y be the last point of α at distance $\kappa(\rho, L, A)$ from γ (see Figure 3), and apply Lemma 4.7 to find:

$$\begin{aligned} d(o, y) \geq R - \kappa'(\rho, 1, 0) - d(o, o') \\ - M(\kappa'(\rho, 1, 0) + \kappa(\rho', \sqrt{r/3}, r/3 + 2d(o, o'))) - \lambda(\rho, \sqrt{r/3}, r/3) \end{aligned}$$

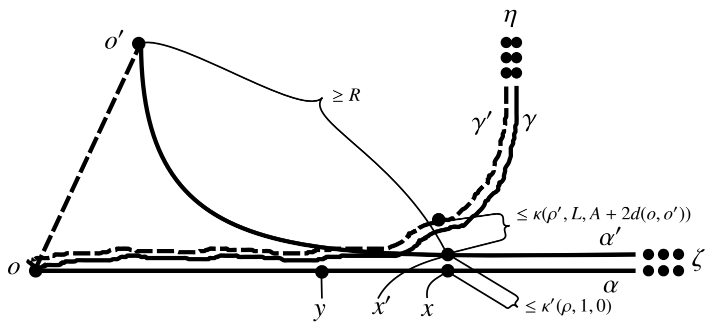

Figure 3. Change of basepoint

Assuming that R was chosen large enough to guarantee the right-hand side is at least r, we have that γ comes within distance $\kappa(\rho, L, A)$ of α outside $N_r o$. □

Proposition 5.11. *$\partial_c^{\mathcal{FQ}}X$ is Hausdorff.*

Proof. Let ζ and η be distinct points in $\partial_c X$. Let $\alpha := \alpha^\zeta$ and $\beta := \alpha^\eta$ be representative geodesic rays. Let R be large enough that the $\kappa'(\rho_\zeta, 1, 0)$–neighborhood of $\alpha_{[R,\infty)}$ is disjoint from the $\kappa'(\rho_\eta, 1, 0)$–neighborhood of $\beta_{[R,\infty)}$. Such an R exists by Corollary 4.4.

Choose $\xi \in U(\zeta, R)$. Let $\gamma \in \xi$ be a geodesic ray. Since $\xi \in U(\zeta, R)$ there exists a point $x \in \alpha$ and $y \in \gamma$ with $d(x, o) \geq R$ and $d(x, y) \leq \kappa(\rho_\zeta, 1, 0)$. By construction $d(y, \beta) > \kappa'(\rho_\eta, 1, 0)$, so, by Corollary 4.4, the final visit of γ to the $\kappa(\rho_\eta, 1, 0)$–neighborhood of β must have occurred inside the ball of radius R about o. Thus, $\xi \notin U(\eta, R)$. □

Proposition 5.12. *$\partial_c^{\mathcal{FQ}}X$ is regular.*

Proof. Suppose $C \subset \partial_c^{\mathcal{FQ}}X$ is closed and $\zeta \in C^c$. Then C^c is a neighborhood of ζ, so there exists r' such that for all $r \geq r'$ we have $U(\zeta, r) \subset C^c$. Suppose:

$$\forall \zeta \in \partial_c^{\mathcal{FQ}}X,\ \exists r' \geq 1,\ \forall r \geq r',\ \exists R > r,\ \overline{U(\zeta, R)} \subset U(\zeta, r) \tag{6}$$

Then there exists an $R > r$ such that $\overline{U(\zeta, R)} \subset U(\zeta, r) \subset C^c$, so C is contained in an open set $\overline{U(\zeta, R)}^c$ that is disjoint from $U(\zeta, R)$. By Proposition 5.5, $U(\zeta, R)$ is a neighborhood of ζ, so it contains an open set U that contains ζ. The disjoint open sets U and $\overline{U(\zeta, R)}^c$ separate ζ and C, so (6) implies regularity.

The proof of (6) is similar to the proof of Lemma 5.8: suppose given r and ζ there is no R satisfying the claim. Then there exists a point $\eta \in \overline{U(\zeta, R)} \cap U(\zeta, r)^c$. Now $\eta \in \overline{U(\zeta, R)}$ implies that for all $n \in \mathbb{N}$ there exists $\xi_n \in U(\zeta, R) \cap U(\eta, n)$, while $\eta \notin U(\zeta, r)$ implies there exist L^2, $A < r/3$ and a continuous (L, A)–quasi-geodesic $\gamma \in \eta$ such that $d(\gamma, N_r^c o \cap \alpha^\zeta) > \kappa(\rho_\zeta, L, A)$. For sufficiently large n we wag the tail of γ by Lemma 5.7 to produce a continuous $(2L+1, A)$–quasi-geodesic $\delta \in \xi_n$ that agrees with γ on a long initial segment. If R is large enough this sets up a contradiction between the fact that $\xi_n \in U(\zeta, R)$ and the fact that γ witnesses $\eta \notin U(\zeta, r)$, so for large enough R we have $\overline{U(\zeta, R)} \subset U(\zeta, r)$, as desired. □

Generally in this paper we will work directly with the topology on the contracting boundary. However, it is worth mentioning that this object that we have called a 'boundary' really is a topological boundary.

Definition 5.13. A *bordification* of a Hausdorff topological space X is a Hausdorff space $\hat{X}$ containing X as an open, dense subset.

The contracting boundary of a proper geodesic metric space provides a bordification of X by $\hat{X} := X \cup \partial_c X$ as follows. For $x \in X$ take a neighborhood basis for x to be metric balls about x. For $\zeta \in \partial_c X$ take a neighborhood basis for ζ to be sets $\hat{U}(\zeta, r)$ consisting of $U(\zeta, r)$ and points $x \in X$ such that we have $d(\gamma, N_r^c o \cap \alpha^\zeta) \leq \kappa(\rho_\zeta, L, A)$ for every $L \geq 1$, $A \geq 0$, and continuous (L, A)–quasi-geodesic segment γ with endpoints o and x.

Proposition 5.14. *$\hat{X} := X \cup \partial_c X$ topologized as above defines a first countable bordification of X such that the induced topology on $\partial_c X$ is the topology of fellow-travelling quasi-geodesics.*

Proof. A similar argument to that of Proposition 5.5 shows we have defined a neighborhood basis in a topology for each point in $\hat{X}$, and the topology agrees with the metric topology on X and topology $\mathcal{FQ}$ on $\partial_c X$ by construction. That $\hat{X}$ is Hausdorff follows from Proposition 5.11. X is clearly open in $\hat{X}$. To see that X is dense, consider $\zeta \in \partial_c X$, which, by definition, is an equivalence class of contracting quasi-geodesic rays. For any quasi-geodesic ray $\gamma \in \zeta$ we have that $(\gamma_n)_{n \in \mathbb{N}}$ is a sequence in X converging to γ, because the subsegments $\gamma_{[0,n]}$ are uniformly contracting. □

Definition 5.15. If G is a finitely generated group acting properly discontinuously on a proper geodesic metric space X with basepoint o we define the *limit set* $\Lambda(G) := \overline{Go} \setminus Go$ of G to be the topological frontier in $\hat{X}$ of the orbit Go of the basepoint.

6. Quasi-isometry invariance

In this section we prove quasi-isometry invariance of the topology of fellow-travelling quasi-geodesics.

Theorem 6.1. *Suppose $\phi\colon X \to X'$ is a quasi-isometric embedding between proper geodesic metric spaces. If ϕ takes contracting quasi-geodesics to contracting quasi-geodesics then it induces an injection $\partial_c\phi\colon \partial_c^{\mathcal{FQ}}X \to \partial_c^{\mathcal{FQ}}X'$ that is an open mapping onto its image with the subspace topology. If $\phi(X)$ is a contracting subset of X' then $\partial_c\phi$ is continuous.*

We will see in Lemma 6.6 that if $\phi(X)$ is contracting then ϕ does indeed take contracting quasi-geodesics to contracting quasi-geodesics, so we get the following corollary of Theorem 6.1.

Corollary 6.2. *If $\phi\colon\ X \to X'$ is a quasi-isometric embedding between proper geodesic metric spaces and $\phi(X)$ is contracting in X' then $\partial_c\phi$ is an embedding. In particular, if ϕ is a quasi-isometry then $\partial_c\phi$ is a homeomorphism.*

Remark 6.3. Cordes [16] proves a version of Theorem 6.1 and Corollary 6.2 for the Morse boundary. The construction of the injective map is exactly the same. For continuity, he defines a map between contracting boundaries to be *Morse-preserving* if for each μ there is a μ' such that the map takes boundary points with a μ–Morse representative to boundary points with a μ'–Morse representative, and shows that if ϕ is a quasi-isometric embedding that induces a Morse-preserving map $\partial_c\phi$ on the contracting boundary then $\partial_c\phi$ is continuous in the direct limit topology.

Similarly, let us say that ϕ is *Morse-controlled* if for each μ there exists μ' such that ϕ takes μ–Morse geodesics to μ'–Morse geodesics. A Morse-controlled quasi-isometric embedding induces a Morse-preserving boundary map. We will see in Lemma 6.6 that the hypothesis that $\phi(X)$ is a contracting set implies that ϕ is Morse-controlled.

Cayley graphs of a fixed group with respect to different finite generating sets are quasi-isometric, so Corollary 6.2 allows us to define the contracting boundary of a finitely generated group, independent of a choice of generating set.

Definition 6.4. If G is a finitely generated group define $\partial_c^{\mathcal{FQ}}G$ to be $\partial_c^{\mathcal{FQ}}X$ where X is any Cayley graph of G with respect to a finite generating set.

The hypothesis in Theorem 6.1 that $\phi(X)$ is contracting already implies that it is undistorted, so in fact we do not need to explicitly require ϕ to be a quasi-isometric embedding. We can relax the hypotheses by only requiring ϕ to be coarse Lipschitz and uniformly proper. This is illustrated by the following easy lemma.

A map $\phi\colon\ X \to X'$ between metric spaces is *coarse Lipschitz* if there are constants $L \geq 1$ and $A \geq 0$ such that $d(\phi(x), \phi(x')) \leq Ld(x, x') + A$ for all $x, x' \in X$. It is *uniformly proper* if there exists a non-decreasing function $\chi\colon\ [0,\infty) \to [0,\infty)$ such that $d(x, x') \leq \chi(d(\phi(x), \phi(x')))$ for all $x, x' \in X$. Note that if X is geodesic and ϕ is coarse Lipschitz and uniformly proper then $\chi(r) > 0$ once $r > A$.

Lemma 6.5. *If $\phi\colon\ X \to X'$ is a coarse Lipschitz, uniformly proper map between geodesic metric spaces and $Z \subset X$ has quasi-convex image in X' then $\phi|_Z\colon\ Z \to X'$ is a quasi-isometric embedding.*

We will prove a stronger statement than this in Lemma 6.7.

Lemma 6.6. *Suppose $\phi\colon\ X \to X'$ is a coarse Lipschitz, uniformly proper map between geodesic metric spaces and $Z \subset X$. If $\phi(X)$ is Morse and Z is Morse then $\phi(Z)$ is Morse. If $\phi(Z)$ is Morse then Z is Morse. Moreover, the Morse function of Z determines the Morse function of $\phi(Z)$, and vice versa, up to functions depending on ϕ.*

Before proving Lemma 6.6 let us consider some examples to motivate the hypotheses. If X is a Euclidean plane, X' is a line, Z is a geodesic in X, and ϕ is the composition of projection of X onto Z and an isometry between Z and X' then ϕ is Lipschitz and $\phi(Z)$ is Morse, but ϕ is not proper and Z is not Morse. If X is a line, X' is a plane, $Z = X$, and ϕ is an isometric embedding then ϕ is Lipschitz and uniformly proper and Z is Morse, but $\phi(X) = \phi(Z)$ is not Morse.

In this paper we will only use the lemma in the case that $\phi(X)$ is Morse, and in this case it is easy to prove that Z is Morse when $\phi(Z)$ is. However, the more general statement might be of independent interest, and requires only mild generalizations of known results. The proof of the first claim uses essentially the same argument as the well-known result that quasi-convex subspaces are quasi-isometrically embedded. The key technical point for this direction is made in Lemma 6.7 (in a more general form than needed for Lemma 6.6, for later use). The second claim is proved using the same strategy as used by Drutu, Mozes, and Sapir [22, Lemma 3.25], who proved it in the case that X' is a finitely generated group, $\phi\colon X \to X'$ is inclusion of a finitely generated subgroup, and Z is an infinite cyclic group.

LEMMA 6.7. *If $\phi\colon X \to X'$ is a coarse Lipschitz, uniformly proper map between geodesic metric spaces and $Z \subset X$ has Morse image in X' then for every $L \geq 1$ and $A \geq 0$ there exist $L' \geq 1$, $A' \geq 0$, $D' \geq 0$, and $D \geq 0$ such that for every (L, A)–quasi-geodesic γ in X' with endpoints on $\phi(Z)$ there is an (L', A')–quasi-geodesic δ in X with endpoints in Z such that:*

- $\delta \subset N_{D'}(Z)$
- $\gamma \subset N_D(\phi(\delta))$
- *γ and $\phi(\delta)$ have the same endpoints.*

The proof, briefly, is to project γ to $\phi(Z)$ and then pull the image back to X.

PROOF. Suppose ϕ is χ–uniformly proper, (L_ϕ, A_ϕ)–coarse-Lipschitz, and $\phi(Z)$ is μ–Morse. Suppose the domain of γ is $[0, T]$. For $z \in \{0, T\}$ choose $\delta_z \in Z$ such that $\phi(\delta_z) = \gamma_z$. For $z \in \mathbb{Z} \cap (0, T)$ choose $\delta_z \in Z$ such that $d(\phi(\delta_z), \gamma_z) \leq \mu(L, A)$. Complete δ to a map on $[0, T]$ by connecting the dots by geodesic interpolation in X. For $D := L/2 + A + \mu(L, A)$ we have $\gamma \subset \bar{N}_D(\phi(\delta))$. Since the reparameterized geodesic segments used to build δ have endpoints on Z and length at most $\chi(L + A + 2\mu(L, A))$, by choosing $D' := \chi(L + A + 2\mu(L, A))/2$ we have $\delta \subset \bar{N}_{D'}(Z)$, and, furthermore, δ is $\chi(L + A + 2\mu(L, A))$–Lipschitz. For any $a \in [0, T]$ we have $d(\phi(\delta_a), \gamma_a) \leq L_\phi D' + A_\phi + D$. Finally, for $a, b \in [0, T]$:

$$\begin{aligned} L_\phi d(\delta_a, \delta_b) + A_\phi &\geq d(\phi(\delta_a), \phi(\delta_b)) \\ &\geq d(\gamma_a, \gamma_b) - 2(L_\phi D' + A_\phi + D) \\ &\geq \frac{|b - a|}{L} - A - 2(L_\phi D' + A_\phi + D) \end{aligned}$$

Thus, δ is an (L', A')–quasi-geodesic for $L' := \max\{L_\phi L, \chi(L + A + 2\mu(L, A))\}$ and $A' := (3A_\phi + A + 2L_\phi D' + 2D)/L_\phi$. □

Lemma 6.5 follows by the same argument applied to a geodesic.

PROOF OF LEMMA 6.6. Suppose $\phi(X)$ and Z are Morse. A quasi-geodesic γ in X' with endpoints on $\phi(Z)$ has endpoints on $\phi(X)$, which is Morse. Apply Lemma 6.7 to get a quasi-geodesic δ in X such that $\phi(\delta)$ is coarsely equivalent to γ. We can, and do, choose δ so that it has endpoints on Z. Since Z is Morse, δ stays close to Z, so $\phi(\delta)$ stays close to $\phi(Z)$, so γ is close to $\phi(Z)$. Thus, $\phi(Z)$ is Morse.

Now suppose $\phi(Z)$ is Morse. By Lemma 6.5, ϕ restricted to Z is a quasi-isometric embedding. Suppose that it is an (L, A)–quasi-isometric embedding. (These constants are at least the coarse Lipschitz constants, so we will also assume ϕ is (L, A)–coarse Lipschitz on all of X.) Suppose $\phi(Z)$ is μ–Morse. The Morse property implies that $\phi(Z)$ is $(2\mu(1, 0) + 1)$–coarsely connected, so Z is E–coarsely connected for $E := L(2\mu(1, 0) + 1 + A)$. Let $t := \frac{1}{6L^2}$. If Z has diameter at most $\frac{1+2t}{1-2t}E$ then it is μ'–Morse for μ' the function $\frac{1+2t}{1-2t}E$, which depends only on L, A, and μ. In this case we are done. Otherwise we prove that Z is t–recurrent and apply Theorem 2.2.

We fix $C \geq 1$ and produce the corresponding D from the definition of recurrence.

Since the diameter of Z is bigger than $\frac{1+2t}{1-2t}E$, the fact that Z is E–coarsely connected implies that for every a, $b \in Z$ there exists a point $c \in Z$ such that $td(a, b) + E \geq \min\{d(a, c),\ d(b, c)\} \geq td(a, b)$: if $d(a, b) \geq \frac{1}{1-2t}E$ then c may be found on a coarse path from a to b, otherwise c may be found on a coarse path joining a to one of two points separated by more than $\frac{1+2t}{1-2t}E$. Such a point c is within distance $td(a, b) + E$ of every path with endpoints a and b, so for any fixed $K \geq 0$ we may restrict our attention to the case $d(a, b) > K$ by assuming D is at least $tK + E$.

Suppose p is path in X with endpoints a and b on Z such that p has length $|p|$ at most $Cd(a, b)$ and $d(a, b) > 8L(A + 1)$. Subdivide p into $\lceil |p| \rceil$ many subsegments, all but possibly the last of which has length 1. Denote the endpoints of these subsegments $a = x_0, x_1, \ldots, x_{\lceil |p| \rceil} = b$. Let q be a path in X' obtained by connecting each $\phi(x_i)$ to $\phi(x_{i+1})$ by a geodesic. Then q is a path of length at most $(L + A)\lceil |p| \rceil \leq (L + A)\frac{9}{8}|p|$ that coincides with $\phi(p)$ on $\phi(\{x_0, \ldots, x_{\lceil |p| \rceil}\})$. Since

ϕ is an (L, A)–quasi-geodesic embedding of Z we have that the distance between the endpoints of q is at least $d(a, b)/L - A \geq \frac{7}{8L}d(a, b)$, so that $|q| < C'd(\phi(a), \phi(b))$ for $C' = 9CL(L + A)/7$. Since $\phi(Z)$ is Morse it is recurrent, so given $t' := 1/3$ and C' as above there is a $D' \geq 0$ and $z \in Z$ such that $\min\{d(\phi(z), \phi(a)), d(\phi(z), \phi(b))\} \geq d(\phi(a), \phi(b))/3$ and $d(\phi(z), q) \leq D'$. Thus, there is some i such that $d(\phi(x_i), \phi(z)) \leq D' + (L + A)/2$. If ϕ is χ–uniformly proper then $d(x_i, z) \leq D := \max\{\chi(D' + (L + A)/2),\ 8tL(A + 1) + E\}$. It remains to check that z is sufficiently far from the endpoints of p. This follows easily from our choice of t, the distance bound between $\phi(z)$ and the endpoints of q, and the assumption $d(a, b) > 8LA$, by using the fact that $\phi|_Z$ is an (L, A)–quasi-isometric embedding. □

Corollary 6.8. *If G is a finitely generated group and Z is a subset of a finitely generated subgroup H of G such that Z is Morse in G then Z is Morse in H. If G is a finitely generated group and H is a Morse subgroup of G then every Morse subset Z of H is also Morse in G.*

Proof of Theorem 6.1. Since the topology is basepoint invariant we choose $o \in X$ and let $o' := \phi(o) \in X'$.

Suppose ϕ is an (L, A)–quasi-isometric embedding, and suppose $\bar{\phi}\colon \phi(X) \to X$ is an (L, A)–quasi-isometry inverse to ϕ. We assume $\sup_{x \in X'} d(\phi \circ \bar{\phi}(x), x) \leq A$.

The quasi-isometric embedding ϕ induces an injective map between equivalence classes of quasi-geodesic rays based at o and equivalence classes of quasi-geodesic rays based at o'. The hypothesis that ϕ sends contracting quasi-geodesics to contracting quasi-geodesics implies that it takes equivalence classes of contracting quasi-geodesic rays to equivalence classes of contracting quasi-geodesic rays, so ϕ induces an injection $\partial_c\phi\colon \partial_c X \to \partial_c X'$.

Continuity: Assume $\phi(X)$ is ρ–contracting. By Lemma 6.6, ϕ sends contracting quasi-geodesic rays to contracting quasi-geodesic rays, so we have an injective map $\partial_c\phi$ as above. We claim that:

$$\forall \zeta \in \partial_c\phi(\partial_c X),\ \forall r > 1,\ \exists R' > 1,\ \forall R \geq R',\ U((\partial_c\phi)^{-1}(\zeta), R) \subset (\partial_c\phi)^{-1}(U(\zeta, r)) \qquad (7)$$

Given the claim, let U' be an open set in $\partial_c^{\mathcal{FQ}} X'$. For each $\zeta \in U' \cap \partial_c\phi(\partial_c X)$ there exists an r_ζ such that $U(\zeta, r_\zeta) \subset U'$. Apply (7) to get an R_ζ, and choose an open neighborhood of $(\partial_c\phi)^{-1}(\zeta)$ contained in $U((\partial_c\phi)^{-1}(\zeta), R_\zeta)$. Let U be the union of these open sets for all $\zeta \in U' \cap \partial_c\phi(\partial_c X)$. Then U is an open set and (7) implies $U = (\partial_c\phi)^{-1}(U')$.

To prove the claim we play our usual game of supposing the converse, deriving a bound on R, and then choosing R to be larger than that bound. The key point is that all of the constants involved are bounded in terms of ζ, $(\partial_c\phi)^{-1}(\zeta)$, r, and ρ.

Suppose for given $\zeta \in \partial_c\phi(\partial_c X)$ and $r > 1$ there exists an $R > 1$ and a point $\eta \in U((\partial_c\phi)^{-1}(\zeta), R)$ such that $\eta \notin (\partial_c\phi)^{-1}(U(\zeta, r))$. The latter implies there exists a continuous (L', A')–quasi-geodesic $\gamma \in \partial_c\phi(\eta)$ witnessing $\partial_c\phi(\eta) \notin U(\zeta, r)$. By Observation 5.1, the quasi-geodesic constants of γ are bounded in terms of r. We must adjust γ to get it into the domain of $\bar{\phi}$. Since $\partial_c\phi(\eta)$ is in the image of $\partial_c\phi$, the quasi-geodesic γ is asymptotic to a quasi-geodesic contained in $\phi(X)$, so γ is contained in a bounded neighborhood of $\phi(X)$. Since $\phi(X)$ as a whole is ρ–contracting, we can replace γ by a projection γ' of γ to $\phi(X)$ as Lemma 6.7. The Hausdorff distance between γ and γ' is bounded[3] in terms of ρ and the quasi-geodesic constants of γ, hence by r, and the additive quasi-geodesic constant of γ' increases by at most twice the Hausdorff distance.

Tame $\bar{\phi}(\gamma')$ to get a continuous quasi-geodesic $\hat{\gamma} \in \eta$. The Hausdorff distance between them and the quasi-geodesic constants (L'', A'') of $\hat{\gamma}$ are bounded in terms of the quasi-isometry constants of γ' and ϕ.

Let $\alpha := \alpha^{(\partial_c\phi)^{-1}(\zeta)}$. Since $\hat{\gamma} \in \eta \in U((\partial_c\phi)^{-1}(\zeta), R)$, there exists $x \in \alpha$ such that $d(o, x) \geq R$ and $d(x, \hat{\gamma}) \leq \kappa(\rho_{(\partial_c\phi)^{-1}(\zeta)}, L'', A'')$. By the argument of the previous paragraph, $\kappa(\rho_{(\partial_c\phi)^{-1}(\zeta)}, L'', A'')$ can be bounded in terms of L, A, r, ρ, ρ_ζ, and $\rho_{(\partial_c\phi)^{-1}(\zeta)}$. This bound

[3] If we had only assumed ϕ to be Morse-controlled this bound would depend on the Morse/contraction function of η, which can be arbitrarily bad, even for η in a small neighborhood of ζ.

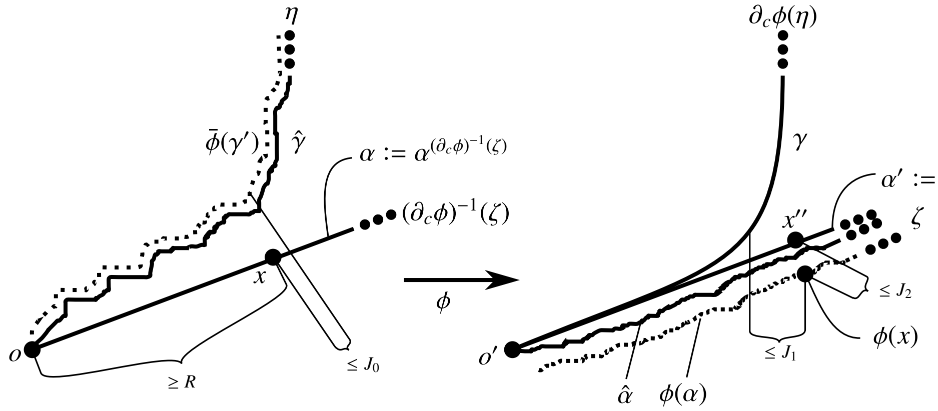

FIGURE 4. Setup for Theorem 6.1

plus the Hausdorff distance to $\bar\phi(\gamma')$ give a bound $d(\bar\phi(\gamma'), x) \leq J_0$. Push forward by ϕ to get $d(\gamma, \phi(x)) \leq J_1 := LJ_0 + 2A + d_{\mathrm{Haus}}(\gamma, \gamma')$. We also know $d(\phi(x), o') \geq R/L - A$.

The quasi-isometric embedding ϕ sends the geodesic α to an (L, A)–quasi-geodesic $\phi(\alpha)$ asymptotic to $\alpha' := \alpha^\zeta$ with $\phi(\alpha)_0 = \phi(o) = o'$. Tame $\phi(\alpha)$ to produce a continuous $(L, 2L+2A)$–quasi-geodesic $\hat\alpha$ at Hausdorff distance at most $L + A$ from $\phi(\alpha)$. Since $\hat\alpha \in \zeta$ we have that $\hat\alpha$ is contained in the $\kappa'(\rho_\zeta, L, 2L + 2A)$–neighborhood of α', so $\phi(\alpha)$ is contained in the J_2–neighborhood of α' for $J_2 := \kappa'(\rho_\zeta, L, 2L + 2A) + L + A$. In particular, $d(\phi(x), \alpha') \leq J_2$. Let x'' be the closest point of α' to $\phi(x)$, so that $d(\gamma, x'') \leq J_1 + J_2$ and $d(o', x'') \geq R/L - A - J_2$. By Lemma 4.7, since γ is an (L', A')–quasi-geodesic, if y is the last point of α' such that $d(\gamma, y) = \kappa(\rho_\zeta, L', A')$ then:

$$d(o', y) \geq R/L - A - J_2 - M(J_1 + J_2) - \lambda(\rho_\zeta, L', A') \tag{8}$$

Everything except R in (8) can be bounded in terms of L, A, r, ρ, ρ_ζ, and $\rho_{(\partial_c\phi)^{-1}(\zeta)}$, so, given r and ζ we can choose R large enough to guarantee $d(o', y) > r$. For such an R, we have $\partial_c\phi(\eta) \in U(\zeta, r)$ for every $\eta \in U((\partial_c\phi)^{-1}(\zeta), R)$. This finishes the proof of claim (7), so we conclude $\partial_c\phi$ is continuous if $\phi(X)$ is contracting.

Open mapping: The image of $\bar\phi$ is coarsely dense in X, so it is contracting. Thus, we can apply the argument of the proof of continuity above to $\bar\phi$ to get the following analogue of (7), noting that $\partial_c\bar\phi = (\partial_c\phi)^{-1}$:

$$\forall \zeta \in \partial_c X,\ \forall r > 1,\ \exists R' > 1,\ \forall R \geq R',\ \partial_c\phi(U(\zeta, r)) \supset U(\partial_c\phi(\zeta), R) \cap \partial_c\phi(\partial_c X) \tag{9}$$

Let U be an open set in $\partial_c^{\mathcal{FQ}}X$. For every $\zeta \in U$ there exists r_ζ such that $U(\zeta, r_\zeta) \subset U$. Apply (9) to get R_ζ, and let U_ζ be an open neighborhood of $\partial_c\phi(\zeta)$ contained in $U(\partial_c\phi(\zeta), R_\zeta)$. Then $U' := \bigcup_{\zeta \in U} U_\zeta$ is an open set in $\partial_c^{\mathcal{FQ}}X'$ containing $\partial_c\phi(U)$. The choices of the R_ζ, by (9), imply that $U' \cap \partial_c\phi(\partial_c^{\mathcal{FQ}}X) = \partial_c\phi(U)$. □

One reason it may be convenient to weaken the stated quasi-isometric embedding hypothesis is that the orbit map of a properly discontinuous group action of a finitely generated group on a proper geodesic metric space is always coarse Lipschitz and uniformly proper, so we get the following consequences of Corollary 6.2.

PROPOSITION 6.9. *Suppose G acts properly discontinuously on a proper geodesic metric space X. Suppose the orbit map $\phi\colon\ g \mapsto go$ takes contracting quasi-geodesics to contracting quasi-geodesics and has quasi-convex image. Then:*

- *G is finitely generated.*
- *The orbit map $\phi\colon\ g \mapsto go$ is a quasi-isometric embedding.*
- *The orbit map induces an injection $\partial_c\phi\colon\ \partial_c^{\mathcal{FQ}}G \to \partial_c^{\mathcal{FQ}}X$ that is an open mapping onto its image, which is $\Lambda(G)$ (recall Definition 5.15). In particular, if $\Lambda(G)$ is compact then so is $\partial_c^{\mathcal{FQ}}G$.*

If $\phi(G)$ is contracting in X then the above are true and $\partial_c\phi$ is an embedding, so $\partial_c^{\mathcal{FQ}}G$ is homeomorphic to $\Lambda(G)$.

Proof. Suppose Go is Q–quasi-convex. A standard argument shows that G is finitely generated and ϕ is a quasi-isometric embedding.

Since ϕ takes contracting quasi-geodesics to contracting quasi-geodesics it induces an open injection $\partial_c\phi$ onto its image by Theorem 6.1.

A point in $\zeta \in \partial_c^{\mathcal{FQ}}G$ is sent to the equivalence class of the contracting quasi-geodesic ray $\phi(\alpha^\zeta)$. The sequence $(\phi(\alpha^\zeta_n))_{n\in\mathbb{N}}$ converges to $\partial_c\phi(\zeta)$ in $\hat{X}$, so the image of $\partial_c\phi$ is contained in $\Lambda(G)$.

Conversely, suppose $\zeta \in \Lambda(G)$. Then there is a sequence $(g_no)_{n\in\mathbb{N}}$ converging in $\hat{X}$ to $\zeta \in \partial_c^{\mathcal{FQ}}X$. By passing to a subsequence, we may assume $g_no \in \hat{U}(\zeta, n)$ for all n. The definition of $\hat{U}(\zeta, n)$ implies that for any chosen geodesic γ^n from o to g_no there are points x_n on α^ζ and y_n on γ^n such that $d(x_n, y_n) \leq \kappa := \kappa(\rho_\zeta, 1, 0)$ and $d(o, x_n) \geq n$. Since Go is quasi-convex, there exists $g'_n \in G$ such that $d(y_n, g'_no) \leq Q$. Thus, the set $Go \cap N_{\kappa+Q}\alpha^\zeta$ is unbounded. If $go \in N_{\kappa+Q}\alpha^\zeta$ then an application of Lemma 4.7 implies that a geodesic from o to go has an initial segment that is $\kappa'(\rho_\zeta, 1, 0)$–Hausdorff equivalent to an initial segment of α^ζ, and the length of these initial segments is $d(o, go)$ minus a constant depending on ζ and Q, but not g. Since we can take $d(o, go)$ arbitrarily large, and since every geodesic from o to go is contained in the Q–neighborhood of Go, we conclude that α^ζ is contained in a bounded neighborhood of Go. Now project α^ζ to Go and pull back to G to get a contracting quasi-geodesic ray whose ϕ–image is asymptotic to α^ζ, which shows $\zeta \in \partial_c\phi(\partial_c^{\mathcal{FQ}}G)$.

If $\phi(G)$ is contracting then ϕ does indeed take contracting quasi-geodesics to contracting quasi-geodesics, by Lemma 6.6, and have quasi-convex image, so the previous claims are true and Corollary 6.2 says $\partial_c\phi$ is an embedding. □

Corollary 6.10. *If H is a subgroup of a finitely generated group G and H is contracting in G then H is finitely generated and the inclusion $\iota\colon H \to G$ induces an embedding $\partial_c\iota\colon \partial_c^{\mathcal{FQ}}H \to \partial_c^{\mathcal{FQ}}G$.*

Properly speaking, we ought to require that H is a contracting subset of the Cayley graph of G with respect to some specified generating set, but it follows from Theorem 2.2 that the property of being a contracting subset does not depend on the choice of metric within a quasi-isometry class.

Corollary 6.11. *If H is a hyperbolically embedded subgroup (in the sense of [20]) in a finitely generated group G then $\partial_c\iota\colon \partial_c^{\mathcal{FQ}}H \to \partial_c^{\mathcal{FQ}}G$ is an embedding. A special case is that of a peripheral subgroup of a relatively hyperbolic group.*

Proof. Sisto [39] shows hyperbolically embedded subgroups are Morse, hence contracting. Peripheral subgroups of relatively hyperbolic groups are a motivating example for the definition of hyperbolically embedded subgroups in [20], but in this special case the fact they are Morse was already shown by Drutu and Sapir [23]. □

Together with Corollary 6.2, Corollary 6.8 implies:

Corollary 6.12. *If G is a finitely generated group and Z is a Morse subset of G then $\partial_c^{\mathcal{FQ}}Z$ embeds into $\partial_c^{\mathcal{FQ}}H$ for every finitely generated subgroup H of G containing Z. In particular, if $\partial_c^{\mathcal{FQ}}Z$ is non-empty then so is $\partial_c^{\mathcal{FQ}}H$, and if $\partial_c^{\mathcal{FQ}}Z$ contains a non-trivial connected component then so does $\partial_c^{\mathcal{FQ}}H$.*

7. Comparison to other topologies

Definition 7.1. Let X be a proper geodesic metric space. Take $\zeta \in \partial_cX$. Fix a geodesic ray $\alpha \in \zeta$. For each $r \geq 1$ define $V(\zeta, r)$ to be the set of points $\eta \in \partial_cX$ such that for every geodesic ray $\beta \in \eta$ we have $d(\beta, \alpha \cap N_r^co) \leq \kappa(\rho_\zeta, 1, 0)$.

The same argument as Proposition 5.5 shows that $\{V(\zeta, n) \mid n \in \mathbb{N}\}$ gives a neighborhood basis at ζ for a topology $\mathcal{FG}$ on ∂_cX. We call $\mathcal{FG}$ the *topology of fellow-travelling geodesics*. It is immediate from the definitions that $\mathcal{FQ}$ is a refinement of $\mathcal{FG}$. The topology $\mathcal{FG}$ need not be preserved by quasi-isometries of X [12]. It is an open question whether $\mathcal{FG}$ is preserved by quasi-isometries when X is the Cayley graph of a finitely generated group.

One might also try to take $V'(\zeta, r)$ to be the set of points $\eta \in \partial_c X$ such that for *some* geodesic ray $\beta \in \eta$ we have $d(\beta, \alpha \cap N_r^c o) \leq \kappa(\rho_\zeta, 1, 0)$. Let $\mathcal{FG}'$ denote the resulting topology. Beware that in general $\{V'(\zeta, r) \mid r \geq 1\}$ is only a filter base converging to ζ, not necessarily a neighborhood base of ζ in $\mathcal{FG}'$; the sets $V'(\zeta, r)$ might not be neighborhoods of ζ.

Lemma 7.2. *Let X be a proper geodesic metric space. Let $\partial_\rho X = \{\zeta \in \partial_c X \mid \rho_\zeta \leq \rho\}$, i.e. $\zeta \in \partial_\rho X$ if all geodesics $\alpha \in \zeta$ are ρ–contracting. The topologies on $\partial_\rho X$ generated by taking, for each $\zeta \in \partial_\rho X$ and $r \geq 1$, the sets $U(\zeta, r) \cap \partial_\rho X$, $V(\zeta, r) \cap \partial_\rho X$, or $V'(\zeta, r) \cap \partial_\rho X$, are equivalent.*

Proof. For each ζ and r we have $U(\zeta, r) \cap \partial_\rho X \subset V(\zeta, r) \cap \partial_\rho X \subset V'(\zeta, r) \cap \partial_\rho X$ by definition.

Given that points in $V'(\zeta, r) \cap \partial_\rho X$ and $U(\zeta, r) \cap \partial_\rho X$ are uniformly contracting, a straightforward application of Lemma 4.7 shows that for all ζ and r, for all sufficiently large R we have $V'(\zeta, R) \cap \partial_\rho X \subset U(\zeta, r) \cap \partial_\rho X$. Also since points of $V'(\zeta, R) \cap \partial_\rho X$ are uniformly contracting, these do, in fact, give a neighborhood basis at ζ for the induced topology, as in Proposition 5.5. □

Proposition 7.3. *Let X be a proper geodesic metric space. If X is hyperbolic then $\partial_c^{\mathcal{FQ}} X \cong \partial_c^{\mathcal{FG}} X \cong \partial_c^{\mathcal{FG}'} X$, and these are homeomorphic to the Gromov boundary. If X is CAT(0) then $\partial_c^{\mathcal{FG}} X \cong \partial_c^{\mathcal{FG}'} X$, and these are homeomorphic to the subset of the visual boundary of X consisting of endpoints of contracting geodesic rays, topologized as a subspace of the visual boundary.*

Proof. For a description of a neighborhood basis for points in the Gromov or visual boundary see [11, III.H.3.6] and [11, p. II.8.6], respectively. Note that these are equivalent to the neighborhood bases for $\mathcal{FG}'$.

The claim for hyperbolic spaces follows from Lemma 7.2, because geodesics in a hyperbolic space are uniformly contracting.

If X is CAT(0) then $\partial_c^{\mathcal{FG}} X \cong \partial_c^{\mathcal{FG}'} X$ because there is a unique geodesic ray in each asymptotic equivalence class. □

More generally, $\partial_c^{\mathcal{FG}} X \cong \partial_c^{\mathcal{FG}'} X$ if X is a proper geodesic metric space with the property that every geodesic ray in X is either not contracting or has contraction function bounded by a constant. This follows by the same argument as in [12].

Next, we recall the *direct limit topology*, $\mathcal{DL}$, on $\partial_c X$ of Charney and Sultan [14] and Cordes [16].

For a given contraction function ρ consider the set $\partial_\rho X$ of points ζ in $\partial_c X$ such that one can take $\rho_\zeta \leq \rho$, as in Lemma 7.2. The topologies $\mathcal{FQ}$, $\mathcal{FG}$, and $\mathcal{FG}'$ on $\partial_\rho X$ coincide by Lemma 7.2. For $\rho \leq \rho'$ the inclusion $\partial_\rho X \hookrightarrow \partial_{\rho'} X$ is continuous, and $\partial_c X$, as a set, is the direct limit of this system of inclusions over all contraction functions.

Let $\mathcal{DL}$ be the direct limit topology on $\partial_c X$, that is, the finest topology on $\partial_c X$ such that all of the inclusion maps $\partial_\rho X \hookrightarrow \partial_c X$ are continuous.

Proposition 7.4. *$\mathcal{DL}$ is a refinement of $\mathcal{FQ}$.*

Proof. The universal property of the direct limit topology says that a map from the direct limit is continuous if and only if the precomposition with each inclusion map is continuous. Thus, it suffices to show the inclusion $\partial_\rho X \hookrightarrow \partial_c^{\mathcal{FQ}} X$ is continuous. This is clear from Lemma 7.2, since we can take the topology on $\partial_\rho X$ to be the subspace topology induced from $\partial_\rho X \hookrightarrow \partial_c^{\mathcal{FQ}} X$. □

Lemma 7.5. *$\partial_c^{\mathcal{DL}} X$ is homeomorphic to Cordes's Morse boundary.*

Proof. Cordes considers Morse geodesic rays, and defines the Morse boundary to be the set of asymptotic equivalence classes of Morse geodesic rays based at o, topologized by taking the direct limit topology of the system of uniformly Morse subsets. By Theorem 2.2, a collection of uniformly Morse rays is contained in a collection of uniformly contracting rays, and vice versa.

It follows as in [16, Remark 3.4] that the direct limit topology over uniformly Morse points and the direct limit topology over uniformly contracting points agree on $\partial_c X$. □

As with the other topologies, if X is hyperbolic then $\partial_c^{\mathcal{DL}} X$ is homeomorphic to the Gromov boundary. Thus, if X is a proper geodesic hyperbolic metric space then all of the above topologies yield a compact contracting boundary. Conversely, Murray [33] showed if X is a complete CAT(0) space admitting a properly discontinuous, cocompact, isometric group action, and if $\partial_c^{\mathcal{DL}} X$ is compact and non-empty, then X is hyperbolic. Work of Cordes and Durham [17] shows that if the contracting boundary, with topology $\mathcal{DL}$, of a finitely generated group is non-empty and compact then the group is hyperbolic. We will prove this for $\mathcal{FQ}$ in Section 10.

We have shown that all of the topologies we consider agree for hyperbolic groups. More generally, we could ask about relatively hyperbolic groups. There are many ways to define relatively hyperbolic groups [26, 10, 36, 27, 21, 23, 29, 38], all of which are equivalent in our setting. Let G be a finitely generated group that is hyperbolic relative to a collection of *peripheral* subgroups $\mathcal{P}$. Fix a finite generating set for G. We again use G to denote the Cayley graph of G with respect to this generating set. Let $\tilde{G}$ be the *cusped space* obtained by gluing a combinatorial horoball onto each left coset of a peripheral subgroup, as in [27]. The cusped space is hyperbolic, and its boundary $\partial\tilde{G}$ is the Bowditch boundary of $(G, \mathcal{P})$. Points in the Bowditch boundary that are fixed by a conjugate of a peripheral subgroup are known as *parabolic points*, and the remaining points are known as *conical points*. As described, G sits as a subgraph in $\tilde{G}$.

Theorem 7.6. *If a finitely generated group G is hyperbolic relative to $\mathcal{P}$, then the inclusion $\iota\colon G \hookrightarrow \tilde{G}$ induces a continuous, G–equivariant map $\iota_*\colon \partial_c^{\mathcal{FG}} G \to \partial\tilde{G}$ that is injective at conical points.*

For $\iota_*\colon \partial_c^{\mathcal{FQ}} G \to \partial\tilde{G}$, the preimage of a parabolic point is the contracting boundary of its stabilizer subgroup embedded in $\partial_c^{\mathcal{FQ}} G$ as in Corollary 6.11.

Let $q\colon \partial_c^{\mathcal{FQ}} G \to \partial_c^{\mathcal{FQ}} G/\iota_*$ be the quotient map from $\partial_c^{\mathcal{FQ}} G$ to its ι_*–decomposition space, that is, the quotient space of $\partial_c^{\mathcal{FQ}} G$ obtained by collapsing to a point the preimage of each point in $\iota_*(\partial_c^{\mathcal{FQ}} G)$. If each peripheral subgroup is hyperbolic or has empty contracting boundary then $\iota_* \circ q^{-1}$ is an embedding.

Theorem 7.6 for $\mathcal{DL}$, without the embedding result, was first observed by Tran [41]. Recall from the introduction that the embedding statement is not true for $\mathcal{DL}$ (cf [41, Remark 8.13]).

Corollary 7.7. *If G is a finitely generated group that is hyperbolic relative to subgroups with empty contracting boundaries then $\partial_c^{\mathcal{FQ}} G = \partial_c^{\mathcal{FG}} G$.*

Since the contracting boundary of a hyperbolic group is the same as the Gromov boundary, we also recover the following well-known result (see [30] and references therein).

Corollary 7.8. *If G is hyperbolic and hyperbolic relative to $\mathcal{P}$ then the Bowditch boundary of $(G, \mathcal{P})$ can be obtained from the Gromov boundary of G by collapsing to a point each embedded Gromov boundary of a peripheral subgroup.*

The following example shows that the embedding statement of Theorem 7.6 can fail when a peripheral subgroup is non-hyperbolic with non-trivial contracting boundary.

Example 7.9. Let $A := \langle a, b \mid [a, b] = 1\rangle$, $H := A * \langle c\rangle$, and $G := H * \langle d\rangle$. Since G is a free product of H and a hyperbolic group, G is hyperbolic relative to H.

A geodesic α in G (or H) is contracting if and only if there is a bound B such that α spends at most time B in any given coset of A.

Consider the sequence $(a^n d^\infty)_{n\in\mathbb{N}}$ in $\partial_c^{\mathcal{FQ}} G$. We have $(\iota_*(a^n d^\infty)) \to \iota_*(\partial_c^{\mathcal{FQ}} H)$, which is a parabolic point in $\partial\tilde{G}$. However, $(q(a^n d^\infty))$ does not converge in $\partial_c^{\mathcal{FQ}} G/\iota_*$. To see this, note that every edge e in the Cayley graph of G with one incident vertex in A determines a clopen subset U_e of $\partial_c^{\mathcal{FQ}} G$ consisting of all $\zeta \in \partial_c^{\mathcal{FQ}} G$ such that α^ζ crosses e. Let U be the union of the U_e for every edge e incident to A and labelled by c or c^{-1}. This is an open set containing $\partial_c^{\mathcal{FQ}} H$ such that $q^{-1}(q(U)) = U$ and $a^n d^\infty \notin U$ for all $n \in \mathbb{N}$. Therefore, $q(U)$ is an open set in $\partial_c^{\mathcal{F}Q} G/\iota_*$ containing the point $q(\partial_c^{\mathcal{F}Q} H)$ but not containing any $q(a^n d^\infty)$.

Before proving the theorem let us recall some of the necessary machinery for relatively hyperbolic groups. Any bounded set meets finitely many cosets of the peripherals, and projections of peripheral sets to one another are uniformly bounded.

Given an (L, A)–quasi-geodesic γ, Drutu and Sapir [23] define the *saturation* $\mathrm{Sat}(\gamma)$ of γ to be the union of γ and all cosets gP of peripheral subgroups $P \in \mathcal{P}$ such that γ comes within distance M of gP, where M is a number depending on L and A. [23, Lemma 4.25] says there exists μ independent of γ such that the saturation of γ is μ–Morse. It follows that the analogous saturation of γ in $\tilde{G}$, that is, the union of γ and all horoballs sufficiently close to γ, is also Morse.

Sisto [38] extends these results, showing, in particular, that peripheral subgroups are strongly contracting.[4]

The other key definition is that of a *transition point* of γ, as defined by Hruska [29]. The idea is that a point of γ is *deep* if it is contained in a long subsegment of γ that is contained in a neighborhood of some gP, and a point is a transition point if it is not deep. A quasi-geodesic in γ is bounded Hausdorff distance from a path of the form $\beta_0 + \alpha_1 + \beta_1 + \alpha_2 + \cdots$ where the β_i are shortest paths connecting some $g_i P_i$ to some $g_{i+1} P_{i+1}$ and the α_i are paths in $g_i P_i$. The transition points are the points close to the β–segments. In $\tilde{G}$ there is an obvious way to shorten such a path by letting the α–segments relax into the corresponding horoballs. If the endpoints of α_i are x and y, this replaces α_i with a segment of length roughly $2\log_2 d_G(x, y)$ in $\tilde{G}$. This is essentially all that happens: if γ is a quasi-geodesic in G then take a geodesic $\hat{\gamma}$ with the same endpoints as γ in the coned-off space $\hat{G}$ obtained by collapsing each coset of a peripheral subgroup. Lift $\hat{g}$ to a nice α-β path in G as above. The β–segments are coarsely well-defined, because the cosets of peripheral subgroups are strongly contracting, and the union of the β segments is Hausdorff equivalent to the set of transition points of γ. Only the endpoints of the α–segments are coarsely well defined, but relaxing the α-segments to geodesics in the corresponding horoball yields a uniform quasi-geodesic in $\tilde{G}$ (see [10, Section 7]). Since $\tilde{G}$ is hyperbolic, this is within bounded Hausdorff distance of any $\tilde{G}$ geodesic $\tilde{\gamma}$ with the same endpoints as γ. In particular, $\tilde{\gamma}$ comes boundedly close to the transition points of γ.

Proof of Theorem 7.6. We omit ι from the notation and think of G sitting as a subgraph of $\tilde{G}$. First we show that for $\zeta \in \partial_c G$ the sequence $(\alpha_n^\zeta)_{n\in\mathbb{N}}$ converges to a point of $\partial\tilde{G}$. Distances in G give an upper bound for distances in $\tilde{G}$, so all quasi-geodesics in G asymptotic to α^ζ also converge to this point in $\partial\tilde{G}$, which we define to be $\iota_*(\zeta)$. Let $(x\cdot y) := \frac{1}{2}(d(\mathbf{1}, \mathbf{x}) + \mathbf{d}(\mathbf{1}, \mathbf{y}) - \mathbf{d}(\mathbf{x}, \mathbf{y}))$ denote the Gromov product of $x, y \in \tilde{G}$ with respect to the basepoint $\mathbf{1}$ corresponding to the identity element of G. (See [11, Section III.H.3] for background on boundaries of hyperbolic spaces.)

To see that the sequence $(\alpha_n^\zeta)_{n\in\mathbb{N}}$ does indeed converge, there are two cases. If α^ζ has unbounded projection to gP for some $g \in G$ and $P \in \mathcal{P}$, then a tail of α^ζ is contained in a bounded neighborhood of gP, but leaves every bounded subset of gP. It follows that (α_n^ζ) converges to the parabolic point in $\partial\tilde{G}$ fixed by gPg^{-1} corresponding to the horoball attached to gP. Furthermore, the projection of the tail of α^ζ to gP is a contracting quasi-geodesic ray in gP (by Corollary 6.8), so P has non-trivial contracting boundary.

The other case is that α^ζ has bounded (not necessarily uniformly!) projection to every gP. Now, given any r there are only finitely many horoballs in $\tilde{G}$ that meet the r–neighborhood of $\mathbf{1}$. Since α^ζ has bounded projection to each of these, for sufficiently large s none of these are in $\mathrm{Sat}(\alpha^\zeta_{[s,\infty)})$. Since $\mathrm{Sat}(\alpha^\zeta_{[s,\infty)})$ is μ–Morse in $\tilde{G}$ for some μ independent of α^ζ, for any $m,\ n \geq s$, geodesics connecting α_m^ζ and α_n^ζ in $\tilde{G}$ stay outside the $(r - \mu(1, 0))$–ball about $\mathbf{1}$. We conclude $\lim_{m,n\to\infty}(\alpha_m^\zeta \cdot \alpha_n^\zeta)_{\tilde{G}} = \infty$, so $(\alpha_n^\zeta)_{n\in\mathbb{N}}$ converges to a point in $\partial\tilde{G}$, which, in this case, is a conical point.

[4] Sisto does not use the term 'strongly contracting', but observe it is equivalent to the first two conditions of [38, Definition 2.1].

If α^ζ and α^η tend to the same conical point in $\partial\tilde{G}$ then the sets of transition points of α^ζ and α^η are unbounded and at bounded Hausdorff distance from one another in G. Since they are contracting geodesics in G they can only come close on unbounded sets if they are in fact asymptotic, so ι_* is injective at conical points.

Continuity: To show $\iota_*\colon \partial_c^{\mathcal{FG}}G \to \partial\tilde{G}$ is continuous we show that for all $\zeta \in \partial_c G$ and all r there exists an R such that for all $\eta \in V(\zeta, R))$ we have $(\iota_*(\zeta) \cdot \iota_*(\eta))_{\tilde{G}} > r$.

Recall that there is a bound B such that a $\tilde{G}$ geodesic comes B–close to the transition points of a G geodesic with the same endpoints. There exists B' so that $\operatorname{diam}\pi_{gP}(x) \leq B'$ for each $x \in G, g \in G, P \in \mathcal{P}$, and so that for any deep point x of a geodesic along gP we have $\operatorname{diam}\{x\} \cup \pi_{gP}(x) \leq B'$. Finally, there exists a constant B'' depending on B so that if $x, y \in G$ satisfy $d(\pi_{gP}(x), \pi_{gP}(y)) \geq B''$ for some $g \in G, P \in \mathcal{P}$, then any geodesic from x to y has a deep component along gP whose transition points at the ends are within B'' of $\pi_{gP}(x)$ and $\pi_{gP}(y)$, respectively.

Suppose ζ and r are given. If $\iota_*(\zeta)$ is conical then given any $r' \geq 0$ there is an R' such that for all $n \geq R'$ we have $d_{\tilde{G}}(\mathbf{1}, \alpha^\zeta_\mathbf{n}) > \mathbf{r}'$. Choose $R \geq R'$ such that α^ζ_R is a transition point, and moreover that any deep component along α^ζ within $\kappa'(\rho_\zeta, 1, 0) + B' + B''$ of α^ζ_R has distance at least R' from $\mathbf{1}$. If $\eta \in V(\zeta, R)$ then α^η comes $\kappa(\rho_\zeta, 1, 0)$–close to $\alpha^\zeta_{[R,\infty)}$ in G, so there is a point α^η_t that is $\kappa'(\rho_\zeta, 1, 0)$–close to α^ζ_R. If α^η_t is a deep point of α^η, let $g'P'$ be the corresponding coset. If $d(\pi_{g'P'}(\mathbf{1}), \alpha^\eta_\mathbf{t}) > \mathbf{J} := \mathbf{B}'' + \mathbf{2B}' + \mathbf{2}\kappa'(\rho, \mathbf{1}, \mathbf{0}))$ then the geodesic α^ζ must also have a deep component along $g'P'$ with one endpoint $(\kappa'(\rho_\zeta, 1, 0) + B' + B'')$–close to α^ζ_R and the other $z := \alpha^\zeta_s \in \bar{N}_{B''}\pi_{g'P'}(\mathbf{1})$; by assumption, $s \geq R'$. If α^η_t is a transition point of α^η, or if $d(\pi_{g'P'}(\mathbf{1}), \alpha^\eta_\mathbf{t}) \leq \mathbf{J}$, then $z := \alpha^\zeta_R$ is $(J + B'')$–close to a transition point of α^η. In either case then, z is a transition point of α^ζ which is $(J + B'')$–close to a transition point of α^η, and has $d(z, \mathbf{1}) \geq \mathbf{R}'$. Thus, by the choice of B, there are points $x \in [\mathbf{1}, \iota_*(\zeta)]$ and $y \in [\mathbf{1}, \iota_*(\eta)]$ with $d_{\tilde{G}}(\mathbf{1}, \mathbf{x}) \geq \mathbf{r}' - \mathbf{B}$ and $d(x, y) \leq 2B + J + B''$. This allows us to bound $(\iota_*(\zeta) \cdot \iota_*(\eta))$ below in terms of these constants and the hyperbolicity constant for $\tilde{G}$, and by choosing r' large enough we guarantee $(\iota_*(\zeta) \cdot \iota_*(\eta)) > r$.

Now suppose $\iota_*(\zeta)$ is parabolic. Then there is some $R_0 \geq 0$, $M \geq 0$, $g \in G$, and $P \in \mathcal{P}$ such that $\alpha^\zeta_{[R_0,\infty)} \subset N_M gP$. For $R \gg R_0$, if $\eta \in V(\zeta, R)$ then α^η comes within distance $\kappa(\rho_\zeta, 1, 0)$ of $\alpha^\zeta_{[R,\infty)}$. If $\iota_*(\zeta) \neq \iota_*(\eta)$ then eventually α^η escapes from gP, so it has a transition point at G–distance greater than $R - R_0 - \kappa(\rho_\zeta, 1, 0)$ from $\alpha^\zeta_{R_0}$. This implies $\operatorname{diam}\pi_{gP}([\mathbf{1}, \iota_*(\eta)]) > \mathbf{R} - \mathbf{R_0} - \mathbf{C}$, where C depends on M, B, the contraction function of gP, and $\kappa(\rho_\zeta, 1, 0)$. It follows from the geometry of the horoballs that for $\iota_*(\zeta)$ is the parabolic boundary point corresponding to gP and $y \in \partial\tilde{G}$ we have $(\iota_*(\zeta) \cdot y)$ is roughly $d_{\tilde{G}}(\mathbf{1}, \mathbf{gP}) + \log_\mathbf{2} \operatorname{diam}\pi_\mathbf{gP}(\mathbf{1}) \cup \pi_\mathbf{gP}(\mathbf{y})$, so by choosing $R > 2^r + R_0 + C$ we guarantee $(\iota_*(\zeta) \cdot \iota_*(\eta)) > r + d_{\tilde{G}}(\mathbf{1}, \mathbf{gP}) \geq \mathbf{r}$.

Embedding: Suppose that $U' \subset \partial_c^{\mathcal{FQ}}G/\iota_*$ is open. Define $U := q^{-1}(U')$, which is open in $\partial_c^{\mathcal{FQ}}G$. We claim that for each $p \in \iota_*(U)$ there exists an $R_p > 0$ such that for $p' \in \iota_*(U)$, if $(p \cdot p') > R_p$ then $\iota_*^{-1}(p') \subset U$. Given the claim, the proof concludes by choosing, for each $p \in \iota_*(U)$, an open neighborhood V_p of p such that $V_p \subset \{p' \in \partial\tilde{G} \mid (p \cdot p') > R_p\}$, and setting $V = \bigcup_{p \in \iota_*(U)} V_p$. Then V is open and $\iota_*(\partial_c^{\mathcal{FQ}}G) \cap V = \iota_*(U)$, so that $\iota_* \circ q^{-1}(\partial_c^{\mathcal{FQ}}G/\iota_*) \cap V = \iota_* \circ q^{-1}(U')$.

First we prove the claim when p is conical. In this case there is a unique point $\zeta \in \iota_*^{-1}(p)$, and since U is open there exists $r_\zeta > 1$ such that $U(\zeta, r_\zeta) \subset U$. Let x be a transition point of α^ζ and choose R such that $d_{\tilde{G}}(\mathbf{1}, \mathbf{x}) \leq \mathbf{R}$. If the claim is false then there exists an $\eta \in \partial_c^{\mathcal{FQ}}G$ such that $(\iota_*(\zeta) \cdot \iota_*(\eta)) > R$ and $\eta \notin U(\zeta, r_\zeta)$. Since $\eta \notin U(\zeta, r_\zeta)$, there exists an L and A and a continuous (L, A)–quasi-geodesic $\gamma \in \eta$ such that the last point $y \in \alpha^\zeta$ such that $d_G(y, \gamma) = \kappa(\rho_\zeta, L, A)$ satisfies $d(\mathbf{1}, \mathbf{y}) < \mathbf{r}_\zeta$. By Observation 5.1, we can take $L < \sqrt{r_\zeta/3}$ and $A < r_\zeta/3$.

By hyperbolicity, geodesics in $\tilde{G}$ tending to $\iota_*(\zeta)$ and $\iota_*(\eta)$ remain boundedly close together for distance approximately $(\iota_*(\zeta) \cdot \iota_*(\eta)) > R$. Since x is a transition point of α^ζ there is a B such that any geodesic $[\mathbf{1}, \iota_*(\zeta)]$ comes B–close to x, so some point z' in a geodesic $[\mathbf{1}, \iota_*(\eta)]$

also comes boundedly close to x. If the point z' lies in a horoball along which γ has a deep component, whose transition points at both ends are close to $[\mathbf{1}, \iota_*(\zeta)]$, then this deep component must be of bounded size else $x \in \alpha^\zeta$ would not be a transition point. It follows that γ must contain a point z at bounded distance from x. Since x and z are transition points, we also get a bound on $d_G(x, z)$. Then, by applying Lemma 4.7, we get an upper bound on $d_G(\mathbf{1}, \mathbf{x})$ depending on r_ζ and ρ_ζ, but independent of γ and η. However, if the set of transition points of α^ζ is bounded in G then it is bounded in $\tilde{G}$, which would imply $\iota_*(\zeta) = p$ is parabolic, contrary to hypothesis.

Now suppose p is parabolic. By hypothesis, its stabilizer G_p is a hyperbolic group conjugate into $\mathcal{P}$. Since the maps are G–equivariant we may assume $G_p \in \mathcal{P}$. We may assume that we have chosen a generating set for G extending one for G_p. Since G_p is quasi-isometrically embedded, by Lemma 6.5, there exist $L_p \geq 1$, $A_p \geq 0$ with $\frac{1}{L_p} d_{G_p}(x, y) - A_p \leq d(x, y) \leq d_{G_p}(x, y)$ for all x, $y \in G_p$. The contracting boundary $\partial_c^{\mathcal{FQ}} G_p$ embeds into $\partial_c^{\mathcal{FQ}} G$ and is compact—it is homeomorphic to the Gromov boundary of G_p. Geodesic rays in G_p are uniformly contracting, by hyperbolicity, so there exists a contracting function ρ such that for all $\zeta \in \partial_c^{\mathcal{FQ}} G_p \subset \partial_c^{\mathcal{FQ}} G$ we have that ζ is ρ–contracting.

We will verify the following fact at the end of this proof:

$$\exists R'_p > 1 \, \forall \xi \in \partial_c^{\mathcal{FQ}} G_p, \; U(\xi, R'_p) \subset U \tag{10}$$

Assuming (10), let $\eta \in \iota_*^{-1}(p')$ and let $\gamma \in \eta$ be a continuous (L, A)–quasi-geodesic for some $L < \sqrt{R'_p/3}$, $A < R'_p/3$. Since G_p is strongly contracting, there exist C and C' such that the diameter of $\pi_{G_p}(\alpha^\eta \cap N_C^c G_p)$ is at most C'. We define $\pi_{G_p}(\eta)$ to be this finite diameter set. Since γ stays $\kappa'(\rho_\eta, L, A)$–close to α^η, strong contraction implies $\pi_{G_p}(\gamma \cap N^c_{\kappa'(\rho_\eta, L, A)} G_p) \subset N_{C'} \pi_{G_p}(\eta)$. We apply [38, Lemma 1.15] to any sufficiently long[5] initial subsegment of γ to conclude there is a function K, a point $z \in \gamma$, and a point $x \in \pi_{G_p}(\eta)$ such that $d(x, z) \leq K(L, A)$.

Since G_p is a hyperbolic group there exists a constant D such that every point is within D of a geodesic ray based at $\mathbf{1}$. Let $\xi \in \partial_c^{\mathcal{FQ}} G_p$ be a point such that there is a G_p–geodesic $[\mathbf{1}, \xi]$ containing a point w with $d(w, x) \leq D$. Since this G_p–geodesic is a (L_p, A_p)–quasi-geodesic in G, there exists $y' \in \alpha^\xi$ such that $d(y', w) \leq \kappa'(\rho, L_p, A_p)$.

We have $d(z, y') \leq K(L, A) + D + \kappa'(\rho, L_p, A_p)$ and:

$$\begin{aligned} d(y', \mathbf{1}) &\geq d(x, \mathbf{1}) - \mathbf{D} - \kappa'(\rho, \mathbf{L_p}, \mathbf{A_p}) \\ &\geq \frac{d_{G_p}(\mathbf{1}, \pi_{\mathbf{G_p}}(\eta))}{L_p} - A_p - D - \kappa'(\rho, L_p, A_p) \end{aligned}$$

Lemma 4.7 implies that, for M and λ as in the lemma, γ comes within distance $\kappa(\rho, L, A)$ of α^ξ outside the ball of radius:

$$\frac{d_{G_p}(\mathbf{1}, \pi_{\mathbf{G_p}}(\eta))}{L_p} - A_p - D - \kappa'(\rho, L_p, A_p) - M(K(L, A) + D + \kappa'(\rho, L_p, A_p)) - \lambda(\rho, L, A) \tag{11}$$

Now, $d_{G_p}(\mathbf{1}, \pi_{\mathbf{G_p}}(\eta)) \asymp \mathbf{2^{(p \cdot p')}} > \mathbf{2^{R_p}}$. Since $L < \sqrt{R'_p/3}$, $A < R'_p/3$, all the negative terms are bounded in terms of R'_p, so we can guarantee (11) is greater than R'_p by taking R_p sufficiently large. This means the quasi-geodesic γ does not witness $\eta \notin U(\xi, R'_p)$. Since γ was arbitrary, $\eta \in U(\xi, R'_p)$, which, by (10), is contained in U. Thus, $\iota_*^{-1}(p') \subset U$ when $(p \cdot p') > R_p$ for R_p sufficiently large with respect to R'_p.

It remains to determine R'_p and verify (10). Define:

$$\theta(s) := s + M(\kappa'(\rho, \sqrt{s/3}, s/3) + \kappa(\rho, 1, 0)) + \lambda(\rho, \sqrt{s/3}, s/3)$$

Since U is open, for every $\zeta \in \partial_c^{\mathcal{FQ}} G_p$ there exists r_ζ such that $U(\zeta, r_\zeta) \subset U$. For each $\zeta \in \partial_c^{\mathcal{FQ}} G_p$, let U_ζ be an open neighborhood of ζ such that $U_\zeta \subset U(\zeta, \theta(r_\zeta))$. Then $\{U_\zeta\}_{\zeta \in \partial_c^{\mathcal{FQ}} G_p}$

[5] Long enough to leave the $\max\{D_0(L, A), \kappa'(\rho_\eta, L, A)\}$–neighborhood of G_p where D_0 is as in [38, Lemma 1.15].

is an open cover of $\partial_c^{\mathcal{FQ}}G_p$, which is compact, so there exists a finite set $F \subset \partial_c^{\mathcal{FQ}}G_p$ such that $\partial_c^{\mathcal{FQ}}G_p \subset \bigcup_{\zeta\in F} U_\zeta \subset \bigcup_{\zeta\in F} U(\zeta, \theta(r_\zeta))$. Define $r := \max_{\zeta\in F} r_\zeta$ and:

$$R'_p := r + \kappa'(\rho, 1, 0) + M(\kappa(\rho, \sqrt{r/3}, r/3) + \kappa'(\rho, 1, 0)) + \lambda(\rho, \sqrt{r/3}, r/3)$$

Suppose that $\xi \in \partial_c^{\mathcal{FQ}}G_p$ and $\eta \in U(\xi, R'_p)$. There exists $\zeta \in F$ such that $\xi \in U_\zeta \subset U(\zeta, \theta(r_\zeta))$. Let $\gamma \in \eta$ be a continuous (L, A)–quasi-geodesic for some $L < \sqrt{r_\zeta/3}$, $A < r_\zeta/3$. Since $\eta \in U(\xi, R'_p)$, there exist $z \in \gamma$ and $x \in \alpha^\xi$ such that $d(x, z) \leq \kappa(\rho, L, A)$ and $d(x, \mathbf{1}) \geq \mathbf{R'_p}$.

There are now two cases to consider. First, suppose that there exists $y' \in \alpha^\zeta$ with $d(x, y) \leq \kappa'(\rho, 1, 0)$. Then $d(y', \mathbf{1}) \geq \mathbf{R'_p} - \kappa'(\rho, \mathbf{1}, \mathbf{0})$. By Lemma 4.7, γ comes $\kappa(\rho, L, A)$ close to α^ζ outside the ball of radius:

$$R'_p - \kappa'(\rho, 1, 0) - M(\kappa'(\rho, 1, 0) + \kappa(\rho, L, A)) - \lambda(\rho, L, A)$$

By definition of R'_p and the conditions $L < \sqrt{r_\zeta/3}$, $A < r_\zeta/3$, this radius is at least r, which is at least r_ζ, so γ does not witness $\eta \notin U(\zeta, r_\zeta)$.

The second case, where the above y' does not exist, is the case that x occurs after α^ξ has already escaped α^ζ. In this case there exists x' between $\mathbf{1}$ and x on α^ξ and $y' \in \alpha^\zeta$ such that $d(x', y') \leq \kappa(\rho, 1, 0)$ and $d(\mathbf{1}, \mathbf{y}') \geq \theta(\mathbf{r}_\zeta)$. Moreover, by Lemma 4.5, there exists $z' \in \gamma$ such that $d(z', x') \leq \kappa'(\rho, L, A)$. By Lemma 4.7, γ comes within distance $\kappa(\rho, L, A)$ of α^ζ outside the ball of radius:

$$\theta(r_\zeta) - M(\kappa(\rho, 1, 0) + \kappa'(\rho, L, A)) - \lambda(\rho, L, A)$$

By definition of θ and the conditions $L < \sqrt{r_\zeta/3}$, $A < r_\zeta/3$, this radius is at least r_ζ, so γ does not witness $\eta \notin U(\zeta, r_\zeta)$. This verifies (10). □

8. Metrizability for group boundaries

In this section let G be a finitely generated group with nonempty contracting boundary. Consider the Cayley graph of G with respect to some fixed finite generating set, which is a proper geodesic metric space we again denote G, and take the basepoint to be the vertex $\mathbf{1}$ corresponding to the identity element of the group.

There is a natural action of G on $\partial_c^{\mathcal{FQ}}G$ by homeomorphisms defined by sending $g \in G$ to the map that takes $\zeta \in \partial_c G$ to the equivalence class of the quasi-geodesic that is the concatenation of a geodesic from $\mathbf{1}$ to g and the geodesic $g\alpha^\zeta$.

The following two results generalize results of Murray [33] for the case of $\partial_c^{\mathcal{DL}}X$ when X is CAT(0). See also [28].

Remark 8.1. If $\beta\colon I \to G$ is a geodesic and β_m is a vertex for some $m \in \mathbb{Z} \cap I$ then β_n is a vertex for every $n \in \mathbb{Z} \cap I$. Vertices in the Cayley graph are in one-to-one correspondence with group elements. If Z is a subset of the Cayley graph we use $\beta_n Z$ to denote the image of Z under the action by the group element corresponding to the vertex β_n.

We will always parameterize bi-infinite geodesics in G so that integers go to vertices.

Proposition 8.2. *G is virtually (infinite) cyclic if and only if $G \curvearrowright \partial_c^{\mathcal{FQ}}G$ has a finite orbit.*

Proof. If G is virtually cyclic then $|\partial_c G| = 2$ and every orbit is finite.

Conversely, if G has a finite orbit then it has a finite index subgroup that fixes a point in $\partial_c G$. The inclusion of a finite index subgroup is a quasi-isometry, so we may assume that G fixes a point $\zeta \in \partial_c G$.

Let $\alpha \in \zeta$ be geodesic and ρ–contracting. Let β be an arbitrary geodesic ray or segment with $\beta_0 = \mathbf{1}$. Since $G\zeta = \zeta$, for all $n \in \mathbb{N}$ the geodesic rays α and $\beta_n\alpha$ are asymptotic. By Theorem 3.4, α and $\beta_n\alpha$ eventually stay within distance κ'_ρ of one another. Truncate α and $\beta_n\alpha$ when their distance is κ'_ρ. By Lemma 3.6, these segments are contracting, and they form a geodesic almost triangle with $\beta_{[0,n]}$, so, by Lemma 3.8, $\beta_{[0,n]}$ is ρ'–contracting for some $\rho' \asymp \rho$ depending only on ρ. Since this is true uniformly for all n, β is ρ'–contracting. Since β was arbitrary and G is homogeneous, every geodesic in G is uniformly contracting, which means

G is hyperbolic and $\partial_c^{\mathcal{F}Q}G$ is the Gromov boundary. If G is hyperbolic and not virtually cyclic then its boundary is uncountable and every orbit is dense, hence infinite. □

PROPOSITION 8.3. *Suppose $|\partial_c^{\mathcal{F}Q}G| > 2$, and fix a point $\eta \in \partial_c G$. For every $\zeta \in \partial_c G$ and every $r \geq 1$ there exists an $R' \geq 1$ such that for all $R_2 \geq R_1 \geq R'$ there exist $g \in G$ such that $\zeta \in U(g\eta, R_2) \subset U(g\eta, R_1) \subset U(\zeta, r)$.*

COROLLARY 8.4. *$\partial_c^{\mathcal{F}Q}G$ is separable.*

COROLLARY 8.5. *If G is not virtually cyclic then $G \curvearrowright \partial_c^{\mathcal{F}Q}G$ is minimal, that is, every orbit is dense.*

REMARK 8.6. For the corollaries we just need to know that we can push η into $U(\zeta, r)$ via the group action. The stronger statement of Proposition 8.3 is used in Proposition 8.7 to upgrade first countable and separable to second countable. The reason for having two parameters R_1 and R_2 is to be able to apply Corollary 5.9 in case $U(g\eta, R_1)$ is not an open set.

PROOF OF PROPOSITION 8.3. By Proposition 8.2, $G \curvearrowright \partial_c^{\mathcal{F}Q}G$ does not have a finite orbit, so there exists a $g' \in G$ with $\eta' := g'\eta \neq \eta$. Let β be a geodesic joining η' and η. It suffices to assume $\beta_0 = \mathbf{1}$; otherwise, we could consider $\beta' := \beta_0^{-1}\beta$, which is a geodesic with $\beta'_0 = \mathbf{1}$ and endpoints in $G\eta$.

Let $\alpha := \alpha^{\zeta}$ be the geodesic representative of ζ. Choose ρ so that α, $\beta_{[0,\infty)}$, and $\bar{\beta}_{[0,-\infty)}$ are all ρ–contracting.

For each integral $t \gg 0$, at most one of $\alpha_t\beta_{[0,\infty)}$ and $\alpha_t\bar{\beta}_{[0,-\infty)}$ remains in the closed $\kappa'(\rho, 1, 0)$–neighborhood of $\alpha_{[0,t]}$ for distance greater than $2\kappa'(\rho, 1, 0)$, otherwise we contradict the fact that $\alpha_t\beta$ is a geodesic. Define $g_t := \alpha_t$ if $\alpha_t\beta_{[0,\infty)}$ does not remain in the closed $\kappa'(\rho, 1, 0)$–neighborhood of $\alpha_{[0,t]}$ for distance greater than $2\kappa'(\rho, 1, 0)$. Otherwise, define $g_t := \alpha_t g'$. For each $s \in \mathbb{N}$ consider a geodesic triangle with sides $\alpha_{[0,t]}$, $g_t\beta_{[0,s]}$, and a geodesic $\delta^{s,t}$ joining α_0 to $g_t\beta_s$. By Lemma 3.6, the first two sides are uniformly contracting, so $\delta^{s,t}$ is as well, by Lemma 3.8. Since G is proper, for each fixed t a subsequence of the $\delta^{s,t}$ converges to a contracting geodesic ray $\delta^t \in g_t\eta$. See Figure 5. Moreover, since the $\delta^{s,t}$ are uniformly

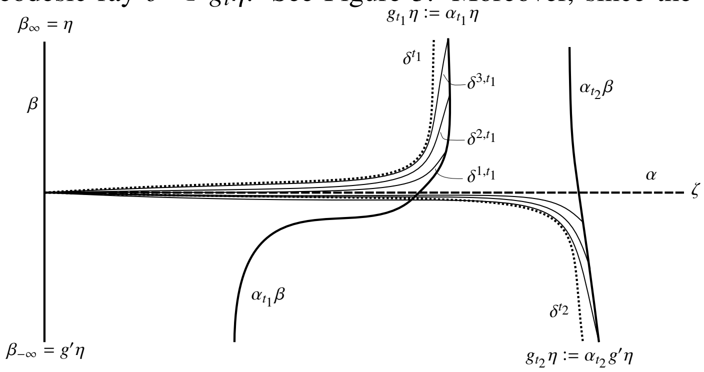

FIGURE 5

contracting, the contraction function for δ^t does not depend on t. Now, for any given t it is possible that δ^t does not coincide with the chosen representative $\alpha^{g_t\eta}$ of $g_t\eta$, but they are asymptotic, so Lemma 4.6 tells us that uniform contraction for the δ^t implies uniform contraction for the $\alpha^{g_t\eta}$. Thus, there is a ρ' independent of t such that $\alpha^{g_t\eta}$ is ρ'–contracting. Furthermore, the defining condition for g_t guarantees that there is a C independent of t such that the geodesic representative $\alpha^{g_t\eta}$ comes within distance $\kappa(\rho, 1, 0)$ of α outside of $N_{t-C}\mathbf{1}$, which implies that $\alpha^{g_t\eta}_{[0,t-C]} \subset \bar{N}_{2\kappa'(\rho,1,0)}\alpha_{[0,t-C]}$.

First we give a condition that implies $\zeta \in U(g_t\eta, R)$. Suppose:

$$t \geq R + C + 2M(\kappa'(\rho, \sqrt{R/3}, R/3) + \kappa'(\rho, 1, 0)) + \lambda(\rho', \sqrt{R/3}, R/3) \tag{12}$$

Suppose that $\gamma \in \zeta$ is a continuous (L, A)–quasi-geodesic. By Observation 5.1 is suffices to consider L^2, $A < R/3$. By Corollary 4.4, $\gamma \subset \bar{N}_{\kappa'(\rho,L,A)}\alpha$, so there is a point γ_a that is $(2\kappa'(\rho, L, A) + 2\kappa'(\rho, 1, 0))$–close to $\alpha^{g_t\eta}_{t-C}$. By Lemma 4.7, γ comes $\kappa(\rho', L, A)$–close to $\alpha^{g_t\eta}$

outside the ball around $\mathbf{1}$ of radius $t - C - M(2\kappa'(\rho, L, A) + 2\kappa'(\rho, 1, 0)) - \lambda(\rho', L, A)$. By (12), this is at least R. Since $\gamma \in \zeta$ was arbitrary, $\zeta \in U(g_t\eta, R)$.

Next, we give a condition that implies $U(g_t\eta, R) \subset U(\zeta, r)$. Suppose:

$$t - C \geq R \geq r + M(2\kappa'(\rho, 1, 0) + 2\kappa'(\rho', \sqrt{r/3}, r/3)) + \lambda(\rho, \sqrt{r/3}, r/3) \tag{13}$$

Suppose that γ is a continuous (L, A)–quasi-geodesic such that $[\gamma] \in U(g_t\eta, R)$. By Observation 5.1, it suffices to consider L^2, $A < r/3$. By definition, γ comes $\kappa(\rho', L, A)$ close to $\alpha^{g_t\eta}$ outside $N_R\mathbf{1}$, so some point γ_b is $2\kappa'(\rho', L, A)$–close to $\alpha_R^{g_t\eta}$, which implies that $d(\gamma_b, \alpha_R) \leq (2\kappa'(\rho', L, A) + 2\kappa'(\rho, 1, 0))$. Now apply Lemma 4.7 to see that γ comes $\kappa(\rho, L, A)$–close to α outside the ball around $\mathbf{1}$ of radius $R - M(2\kappa'(\rho, 1, 0) + 2\kappa'(\rho', L, A)) - \lambda(\rho, L, A)$, which is at least r, by (13). Thus, $U(g_t\eta, R) \subset U(\zeta, r)$.

Equipped with these two conditions, we finish the proof. The contraction functions ρ and ρ' are determined by ζ and η. Given these and any $r \geq 1$, define R' to be the right hand side of (13). Given any $R_2 \geq R_1 \geq R'$, it suffices to define $g := g_t$ for any t large enough to satisfy both condition (12) for $R = R_2$ and condition (13) for $R = R_1$. □

Proposition 8.7. *$\partial_c^{\mathcal{FQ}}G$ is second countable.*

Proof. If G is virtually cyclic then $\partial_c^{\mathcal{FQ}}G$ is the discrete space with two points, and we are done. Otherwise, fix any $\eta \in \partial_c G$. For each $g \in G$ and $n \in \mathbb{N}$ choose an open set $U_{g,n}$ such that $U(g\eta, \psi(\rho_{g\eta}, n)) \subset U_{g,n} \subset U(g\eta, n)$ as in Corollary 5.9.

Let U be a non-empty open set and let ζ be a point in U. By definition of $\mathcal{FQ}$, there exists an $r \geq 1$ such that $U(\zeta, r) \subset U$. Let R' be the constant of Proposition 8.3 for ζ and r, and let $R' \leq R_1 \in \mathbb{N}$. As noted there, there exists a sublinear ρ' such that the points $g_t\eta$ in the proof of Proposition 8.3 are all ρ'–contracting. Define $R_2 := \psi(\rho', R_1) \geq \psi(\rho_{g_t\eta}, R_1)$. Combining Proposition 8.3 and the definition of the sets $U_{g,n}$, there exists $g \in G$ such that:

$$\zeta \in U(g\eta, R_2) \subset U_{g,R_1} \subset U(g\eta, R_1) \subset U(\zeta, r) \subset U$$

Thus, $\mathcal{U} := \{U_{g,n} \mid g \in G,\ n \in \mathbb{N}\}$ is a countable basis for $\partial_c^{\mathcal{FQ}}G$. □

Corollary 8.8. *$\partial_c^{\mathcal{FQ}}G$ is metrizable.*

Proof. $\partial_c^{\mathcal{FQ}}G$ is second countable by Proposition 8.7, regular by Proposition 5.12, and Hausdorff by Proposition 5.11. The Urysohn Metrization Theorem says every second countable, regular, Hausdorff space is metrizable. □

It is an interesting open problem to describe, in terms of the geometry of G, a metric on $\partial_c^{\mathcal{FQ}}G$ that is compatible with $\mathcal{FQ}$.

9. Dynamics

Definition 9.1. An element $g \in G$ is *contracting* if $\mathbb{Z} \to G : n \mapsto g^n$ is a quasi-isometric embedding whose image is a contracting set.

We use g^∞ and $g^{-\infty}$ to denote the equivalence classes of the contracting quasi-geodesic rays based at $\mathbf{1}$ corresponding to the non-negative powers of g and non-positive powers, respectively. These are distinct points in $\partial_c G$.

Lemma 9.2. *Given a contracting element $g \in G$, an $r \geq 1$, and a point $\zeta \in \partial_c^{\mathcal{FQ}}G \setminus \{g^\infty,\ g^{-\infty}\}$ there exists an $R' \geq 1$ such that for every $R \geq R'$ and every $n \in \mathbb{N}$ we have $U(\zeta, R) \subset g^{-n}U(g^n\zeta, r)$.*

Proof. Since g is contracting there is a sublinear function ρ such that all geodesic segments joining powers of g as well as geodesic rays based at $\mathbf{1}$ going to g^∞, $g^{-\infty}$, or $g^m\zeta$ for any $m \in \mathbb{Z}$ are all ρ–contracting.

Consider a geodesic triangle with sides $g^{-m}\alpha^{g^m\zeta}$, α^ζ, and a geodesic from $\mathbf{1}$ to g^{-m} for arbitrary $m \in \mathbb{Z}$. All sides are ρ–contracting, and such a triangle is B–thin for some B independent of m. Thus, for sufficiently large s', independent of m, the point $\alpha_{s'}^\zeta$ is more than B–far from $\langle g\rangle$, hence B–close to $g^{-m}\alpha^{g^m\zeta}$. Since α^ζ and $g^{-m}\alpha^{g^m\zeta}$ are asymptotic, they eventually come

$\kappa(\rho,1,0)$–close and then stay $\kappa'(\rho,1,0)$–close thereafter. Theorem 3.4 says the first time they come $\kappa(\rho,1,0)$ close occurs no later than $s'+\rho'(B)$. Take $R'\geq s'+\rho'(B)$, which guarantees $d(\alpha^\zeta_s, g^{-m}\alpha^{g^m\zeta})\leq\kappa'(\rho,1,0)$ for all $m\in\mathbb{Z}$ and all $s\geq R'$.

Let L' and A' be the constants of Lemma 4.8 for ρ, $L=\sqrt{r/3}$ and $A=r/3$, and let M and λ be as in Lemma 4.7. Take $T:=1+r+M(\kappa(\rho,L',A')+\kappa'(\rho,1,0))+\lambda(\rho,1,0)$. We require further that R' is larger than L' and A' and large enough so that for all $s\geq R'$ we have $d(\alpha^\zeta_s,\langle g\rangle)>T+\kappa'(\rho,1,0)$.

Suppose, for a contradiction, that there exist some $n\in\mathbb{N}$ and $R\geq R'$ such that there exists a point $\eta\in U(\zeta,R)\setminus g^{-n}U(g^n\zeta,r)$. Since $\eta\notin g^{-n}U(g^n\zeta,r)$, for some L, A there exists a continuous (L,A)–quasi-geodesic $\gamma\in g^n\eta$ that does not come $\kappa(\rho,L,A)$–close to $\alpha^{g^n\zeta}$ outside $N_r\mathbf{1}$. By Observation 5.1, it suffices to consider the case that $L<\sqrt{r/3}$ and $A<r/3$.

As in Lemma 4.8, construct a continuous (L',A')–quasi-geodesic δ that first follows a geodesic from $\mathbf{1}$ towards g^{-n}, then a geodesic segment of length $\kappa(\rho,L,A)$, and then follows a tail of $g^{-n}\gamma$. Since it shares a tail with $g^{-n}\gamma$, we have $\delta\in\eta$. Since $\eta\in U(\zeta,R)$, there is some $s\geq R$ such that δ comes within distance $\kappa(\rho,L',A')$ of α^ζ_s. Our choice of R' guarantees that $d(\alpha^\zeta_s,g^{-n}\alpha^{g^n\zeta})\leq\kappa'(\rho,1,0)$ and $d(\alpha^\zeta_s,\langle g\rangle)>T+\kappa'(\rho,1,0)$. The latter implies that the point of δ close to α^ζ_s is a point of $g^{-n}\gamma$, so there is a point of $g^{-n}\gamma$ that comes within distance $\kappa(\rho,L',A')+\kappa'(\rho,1,0)$ of a point x of $g^{-n}\alpha^{g^n\zeta}$ such that $d(x,\langle g\rangle)\geq T$. Lemma 4.7 says that there is a point y on $g^{-n}\alpha^{g^n\zeta}$ such that $d(y,g^{-n}\gamma)=\kappa(\rho,L,A)$ and $d(x,y)\leq M(\kappa(\rho,L',A')+\kappa'(\rho,1,0))+\lambda(\rho,L,A)$. The definition of T implies $d(y,g^{-n})\geq d(y,\langle g\rangle)\geq d(x,\langle g\rangle)-d(x,y)>r$. But then g^ny is a point of $\alpha^{g^n\zeta}$ with $d(g^ny,\mathbf{1})>\mathbf{r}$ and $d(g^ny,\gamma)=\kappa(\rho,L,A)$, contradicting the definition of γ. □

Lemma 9.3. *Given a contracting element $g\in G$, an $r\geq 1$, and a point $\zeta\in\partial^{\mathcal{FQ}}_cG\setminus\{g^{-\infty}\}$ there exist constants $R'\geq 1$ and N such that for all $R\geq R'$ and $n\geq N$ we have $g^nU(\zeta,R)\subset U(g^\infty,r)$.*

Proof. The lemma is easy if $\zeta=g^\infty$, so assume not. Since g is contracting the geodesics $\alpha^{g^n\zeta}$ are uniformly contracting. Let ρ be a sublinear function such that α^{g^∞}, $\alpha^{g^{-\infty}}$, and all $\alpha^{g^n\zeta}$ are ρ–contracting. Since these geodesics are uniformly contracting, ideal geodesic triangles with vertices g^∞, $g^{-\infty}$, and $g^n\zeta$ are uniformly thin, independent of n. Thus, for N sufficiently large and for all $n\geq N$ we have that $\alpha^{g^n\zeta}$ stays $\kappa'(\rho,1,0)$ close to α^{g^∞} for distance greater than $S':=1+r+\lambda(\rho,\sqrt{r/3},r/3)+M(\kappa'(\rho,\sqrt{r/3},r/3)+\kappa'(\rho,1,0))$, where M and λ are as in Lemma 4.7.

Suppose that $\eta\in U(g^n\zeta,S)$ for some $n\geq N$ and $S\geq S'$. Let $\gamma\in\eta$ be a continuous (L,A)–quasi-geodesic for some L^2, $A\leq r/3$. By hypothesis, γ comes $\kappa(\rho,L,A)$–close to $\alpha^{g^n\zeta}$ outside $N_S\mathbf{1}$. Therefore, γ stays $\kappa'(\rho,L,A)$–close to $\alpha^{g^n\zeta}$ in $N_S\mathbf{1}$. By our choice of N, this implies γ stays $(\kappa'(\rho,L,A)+\kappa'(\rho,1,0))$–close to α^{g^∞} in $N_S\mathbf{1}$. By our choice of S and Lemma 4.7, γ comes $\kappa(\rho,L,A)$–close to α^{g^∞} outside the neighborhood of $\mathbf{1}$ of radius:

$$\begin{aligned}S&-M(\kappa'(\rho,L,A)+\kappa'(\rho,1,0))-\lambda(\rho,L,A)\\&\geq S'-M(\kappa'(\rho,\sqrt{r/3},r/3)+\kappa'(\rho,1,0))-\lambda(\rho,\sqrt{r/3},r/3)>r\end{aligned}$$

Since γ was arbitrary, $\eta\in U(g^\infty,r)$, thus $U(g^n\zeta,S)\subset U(g^\infty,r)$.

By Lemma 9.2, given g, S', and ζ there exists an R' such that for all $R\geq R'$ and every $n\in\mathbb{N}$ we have $U(\zeta,R)\subset g^{-n}U(g^n\zeta,S')$. Thus, for this R' and N as above we have, for all $n\geq N$ and $R\geq R'$, that $g^nU(\zeta,R)\subset U(g^n\zeta,S')\subset U(g^\infty,r)$. □

Theorem 9.4 (Weak North-South dynamics for contracting elements). *Let $g\in G$ be a contracting element. For every open set V containing g^∞ and every compact set $C\subset\partial^{\mathcal{FQ}}_cG\setminus\{g^{-\infty}\}$ there exists an N such that for all $n\geq N$ we have $g^nC\subset V$.*

We remark that if Theorem 9.4 were true for *closed* sets and G contained contracting elements without common powers then we could play ping-pong to produce a free subgroup of G. Such a result cannot be true in this generality because there are Tarski Monsters, non-cyclic groups such that every proper subgroup is cyclic, such that every non-trivial element is Morse, hence, contracting [34, Theorem 1.12].

Proof of Theorem 9.4. Since V is an open set containing g^∞ there exists some $r > 0$ such that $U(g^\infty, r) \subset V$. For this r and for each $\zeta \in C$ there exist R_ζ and N_ζ as in Lemma 9.3 such that for all $n \geq N_\zeta$ we have $g^n U(\zeta, R_\zeta) \subset U(g^\infty, r)$. By Proposition 5.5, $U(\zeta, R_\zeta)$ is a neighborhood of ζ, so there exists an open set U'_ζ such that $\zeta \in U'_\zeta \subset U(\zeta, R_\zeta)$. The collection $\{U'_\zeta \mid \zeta \in C\}$ is an open cover of C, which is compact, so there exists a finite subset C' of C such that $\{U'_\zeta \mid \zeta \in C'\}$ covers C. Define $N := \max_{\zeta \in C'} N_\zeta$. For every $n \geq N$ we then have:

$$\begin{aligned} g^n C \subset g^n(\bigcup_{\zeta \in C'} U'_\zeta) &\subset g^n(\bigcup_{\zeta \in C'} U(\zeta, R_\zeta)) \\ = \bigcup_{\zeta \in C'} g^n U(\zeta, R_\zeta) &\subset U(g^\infty, r) \subset V \end{aligned}$$

□

10. Compactness

In this section we characterize when the contracting boundary of a group is compact, Theorem 10.1, and give a partial characterization of when the limit set of a group is compact, Proposition 10.4.

Theorem 10.1. *Let G be an infinite, finitely generated group. Consider the Cayley graph of G, which we again denote G, with respect to a fixed finite generating set. The following are equivalent:*

(1) Geodesic rays in G are uniformly contracting.
(2) Geodesic segments in G are uniformly contracting.
(3) G is hyperbolic.
(4) $\partial_c^{\mathcal{DL}} G$ is non-empty and compact.
(5) $\partial_c^{\mathcal{FQ}} G$ is non-empty and compact.
(6) Every geodesic ray in G is contracting.

Remark 10.2. Work of Cordes and Durham [17] implies '(4) implies (3)'. Roughly the same argument we use for '(1) implies (2)' is contained in the proof of [17, Proposition 4.2]. More interestingly, they prove [17, Lemma 4.1] that compact subsets of the Morse boundary (of a space) are uniformly Morse, which is a more general version of '(4) implies (1)'. We specifically designed the topology of fellow-travelling quasi-geodesics to allow sequences with decaying contraction/Morse functions to converge when geometrically appropriate, so the corresponding statement cannot be true in our setting. In particular, the equivalence of (5) and (6) with (1)-(4) does not follow from their result.

Remark 10.3. Previous attempts have been made to prove results similar to '(6) implies (1)'. We point out a difficulty in the obvious approach. Suppose that $(\gamma^n)_{n\in\mathbb{N}}$ is a sequence of geodesics with decaying Morse functions. Let δ^n be paths witnessing the decaying Morse functions, by which we mean that there exist L and A such that for each n the path δ^n is an (L, A)–quasi-geodesic with endpoints on γ^n, and that δ^n is not contained in $N_n\gamma^n$. Let β^n denote the subsegment of γ^n between the endpoints of δ^n. We may assume by translation that for all n the basepoint **1** is approximately the midpoint of β^n. By properness, there is a subsequence of (γ^n) that converges to a geodesic γ through **1**. One would guess that γ is not Morse, but this is not true in general; explicit counterexamples can be constructed. The problem is that the convergence to γ can be much slower than the growth of $|\beta^n|$ so that the subsegment of γ^n that agrees with γ can be a vanishingly small fraction of β^n. In this case the segments δ^n may not have endpoints close to γ, so no conclusion can be drawn.

It seems difficult to fix this argument. Instead, our strategy, roughly, will be to construct for each i a translate $g_i\delta^i$ of δ^i and for each n a geodesic ray that passes suitably close to both endpoints of $g_i\delta^i$ for all $i \leq n$. We argue that a subsequence of these rays converges to a non-Morse geodesic ray.

Proof. Assume (1). Since G is infinite and finitely generated, there exists a geodesic ray α based at **1**. Recall that for $n \in \mathbb{N}$ the point α_n is a vertex of the Cayley graph, so it corresponds to a unique element of the group G. Thus, $\alpha_n^{-1}\alpha$ is simply the translate of α by the isometry of G defined by left multiplication by the element α_n^{-1}. Since $\mathbf{1} = \alpha_\mathbf{n}^{-\mathbf{1}}\alpha_\mathbf{n} \in \alpha_\mathbf{n}^{-\mathbf{1}}\alpha$, the sequence $(\alpha_n^{-1}\alpha)_{n\in\mathbb{N}}$ has a subsequence that converges to a bi-infinite geodesic β containing $\mathbf{1}$. By construction, subsegments of β are close to subsegments of translates of α, so by Lemma 3.6 and Lemma 3.7, β is contracting, with contraction function determined by the uniform bound for rays. Let β^+ and β^- denote the two rays based at $\mathbf{1}$ such that $\beta = \beta^+ \cup \beta^-$.

Let g be an arbitrary non-trivial element of G. Consider the ideal geodesic triangle with one side $g\beta$ and whose other two sides are geodesic rays based at $\mathbf{1}$ with endpoints $g\beta^+_\infty$ and $g\beta^-_\infty$. The sides of this triangle are uniformly contracting, so it is uniformly thin, so there is some constant C such that every point on $g\beta$ is C–close to one of the other two sides. In particular, g is C–close to one of the other sides. Since the constant C is independent of g, we have that that for every $g \in G$ there exists a geodesic ray γ^g based at $\mathbf{1}$ and passing within distance C of g.

Let δ be a geodesic segment with endpoints h and hg for some $g, h \in G$. The contraction function of the geodesic $h^{-1}\delta$ from $\mathbf{1}$ to g can be bounded in terms of C and the contraction function of γ^g, but since rays have uniform contraction this gives us a bound for the contraction function for $h^{-1}\delta$, hence for δ. Since every geodesic segment is at Hausdorff distance at most 1/2 from a geodesic segment with endpoints at vertices, Lemma 3.3 tells us that geodesic segments are uniformly contracting. Thus, (1) implies (2).

If geodesic segments in G are uniformly contracting then geodesic bigons are uniformly thin, so G is hyperbolic by a theorem of Papasoglu [37]. Thus, (2) implies (3).

By [16, Theorem 3.10], if G is hyperbolic then $\partial_c^{\mathcal{DL}}G$ agrees with the Gromov boundary, which is compact, so (3) implies (4).

$\mathcal{DL}$ is a refinement of $\mathcal{FQ}$ by Proposition 7.4, so (4) implies (5).

If G is virtually cyclic then (6) is true. If G is not virtually cyclic then, by Proposition 8.2, if $\partial_c^{\mathcal{FQ}}G$ is non-empty then it is infinite. In particular, there are distinct points in $\partial_c^{\mathcal{FQ}}G$. Choose two of them and connect them by a geodesic β, which is necessarily contracting. By translating β we may assume that $\beta_0 = \mathbf{1}$. Suppose that α is an arbitrary geodesic ray based at $\mathbf{1}$. As in the proof of Proposition 8.3, after possibly exchanging β with $\bar{\beta}$ there is increasing $\sigma'\colon \mathbb{N} \to \mathbb{N}$ such that for all $n \in \mathbb{N}$ we have that $\alpha_{\sigma'(n)}\beta_{[0,\infty)}$ does not backtrack far along $\alpha_{[0,\sigma'(n)]}$. This means there are L and A, independent of n, such that the concatenation of $\alpha_{[0,\sigma'(n)]}$ and $\alpha_{\sigma'(n)}\beta_{[0,\infty)}$ is a continuous (L, A)–quasi-geodesic.

If $\partial_c^{\mathcal{FQ}}G$ is compact then the sequence $(\alpha_{\sigma'(n)}\beta_\infty)_{n\in\mathbb{N}}$ has a convergent subsequence, so there is an increasing $\sigma''\colon \mathbb{N} \to \mathbb{N}$ such that for $\sigma := \sigma' \circ \sigma''$ we have $(\alpha_{\sigma(n)}\beta_\infty)_{n\in\mathbb{N}}$ converges to a point $\zeta \in \partial_c^{\mathcal{FQ}}G$. Then for every $r > 1$ there exists an N such that for all $n \geq N$ we have $\alpha_{\sigma(n)}\beta_\infty \in U(\zeta, r)$. For $r > 3L^2$, $3A$ we then have that the continuous (L, A)–quasi-geodesic $\alpha_{[0,\sigma(n)]} + \alpha_{\sigma(n)}\beta_{[0,\infty)} \in \alpha_{\sigma(n)}\beta_\infty$ comes $\kappa(\rho_\zeta, L, A)$–close to α^ζ outside the ball of radius r about $\mathbf{1}$, for all sufficiently large n, and therefore has initial segment of length at least r contained in the $\kappa'(\rho_\zeta, L, A)$–neighborhood of α^ζ. Since this is true for all sufficiently large n and since longer and longer initial segments of the $\alpha_{[0,\sigma(n)]} + \alpha_{\sigma(n)}\beta_{[0,\infty)}$ are initial segments of α, we conclude that α is asymptotic to α^ζ, which implies that α is contracting. Thus, (5) implies (6).

Finally, we prove (6) implies (1). We do so by assuming (6) is true and (1) is false, and deriving a contradiction. The strategy is as follows. The fact that $\partial_c^{\mathcal{FQ}}G$ is not uniformly contracting implies that no non-empty open subset of $\partial_c^{\mathcal{FQ}}G$ is uniformly contracting. We construct a nested decreasing sequence of neighborhoods focused on points with successively worse contraction behavior. We use properness of G to conclude that a subsequence of representative geodesics of these focal points converges to a geodesic ray. The assumption (6) implies the limiting ray is contracting, so it represents a point in $\partial_c^{\mathcal{FQ}}G$, and we claim that this point is in the intersection of the nested sequence of neighborhoods. Furthermore, the details of the construction ensure that the limiting ray actually experiences the successively worse contraction behavior of the construction, with the conclusion that it is not a contracting ray, contradicting (6).

If rays in G fail to be uniformly contracting then so do bi-infinite geodesics in G. To see this, fix a ray α. Any other ray has a translate β with the same basepoint as α. Since rays are contracting there is a contracting bi-infinite geodesic γ with endpoints α_∞ and β_∞. Now α, β, and γ make a geodesic triangle, so the contraction function of β can be bounded in terms of those of α and γ. If bi-infinite geodesics are all ρ–contracting then this would mean the contraction function for β can be bounded in terms of only ρ and ρ_α, so rays would be uniformly contracting. Combining this with Theorem 2.2, we have that $\neg(1)$ implies bi-infinite geodesics in G are not uniformly Morse. They are all Morse, as each bi-infinite geodesic can be written as a union of two rays, which are contracting, by (6). For each $L \geq 1, A \geq 0$, and bi-infinite geodesic γ in G, define $D(\gamma, L, A)$ to be the supremum of the set:

$$\{d(z,\gamma) \mid z \text{ is a point on a continuous } (L,A)\text{–quasi-geodesic with endpoints on } \gamma\}$$

Since γ is Morse, $D(\gamma, L, A)$ exists for each L and A. If $\sup_\gamma D(\gamma, L, A)$ exists for every L and A then we can define $\mu(L, A) := \sup_\gamma D(\gamma, L, A)$ and we have that all bi-infinite geodesics are μ–Morse, contrary to hypothesis, so there exist some $L \geq 1$ and $A \geq 0$ such that for all $n \in \mathbb{N}$ there exists a bi-infinite geodesic γ^n and a continuous (L, A)–quasi-geodesic δ^n with endpoints on γ^n such that δ^n is not contained in $N_n\gamma^n$. By translating and shifting the parameterization of γ^n we may assume that $\gamma^n_0 = \mathbf{1}$ and that the distances from $\mathbf{1}$ to the two endpoints of δ^n differ by at most 1.

Now we make a claim and use it to finish the proof:

(14) Let γ be a bi-infinite ρ_γ–contracting geodesic. Given $\zeta \in \partial_c^{\mathcal{FQ}}G$, $R > 1$, $r \geq 0$ there exists $\eta \in U(\zeta, R)$ and $g \in G$ such that α^η passes within distance $\kappa(\rho_\gamma, 1, 0)$ of both endpoints of a segment of $g\gamma$ containing $g\gamma_{[-r,r]}$.

See Figure 6 and Figure 7 for illustrations of (14). Assuming (14), we construct a decreasing nested sequence of neighborhoods in $\partial_c^{\mathcal{FQ}}G$ focusing on points with successively worse contraction behavior. The key trick is to build extra padding into our constants to give the contraction function of the eventual limiting geodesic time to dominate. Let M and λ be as in Lemma 4.7; recall that $\lambda(\phi, 1, 0) = 8\kappa(\phi, 1, 0)$. Let ψ be as in Corollary 5.9. Let $\zeta^0 \in \partial_c^{\mathcal{FQ}}G$ and $R_0 > 1$ be arbitrary. Now, supposing ζ^i and R_i have been defined, consider γ^{i+1}. Let r_{i+1} be the least integer greater than half the distance between endpoints of δ^i plus the quantity $(M + 1)\kappa(\rho_{\gamma^i}, 1, 0) + (6M + 15)(i + 1)$. Apply (14) to γ^{i+1}, ζ^i, $\psi(\zeta^i, R_i)$, r_{i+1} and get output $\zeta^{i+1} \in U(\zeta^i, \psi(\zeta^i, R_i)) \subset U(\zeta^i, R_i)$ and $g_{i+1} \in G$. Let R''_{i+1} be $\kappa(\rho_{\gamma^i}, 1, 0)$ plus the larger of the distances to $\mathbf{1}$ of the endpoints of the subsegment of $g_{i+1}\gamma^{i+1}$ given by (14). Define $R'_{i+1} := R''_{i+1} + (M + 1)\kappa(\rho_{\zeta^{i+1}}, 1, 0) + 9(i + 1)$. By Corollary 5.9, there is an open set U_i such that $U(\zeta^i, \psi(\zeta^i, R_i)) \subset U_i \subset U(\zeta^i, R_i)$, so since $\zeta^{i+1} \in U(\zeta^i, \psi(\zeta^i, R_i))$ we can choose $R_{i+1} \geq R'_{i+1}$ large enough to guarantee $U(\zeta^{i+1}, R_{i+1}) \subset U \subset U(\zeta^i, R_i)$.

Consider the sequence of geodesic rays $(\alpha^{\zeta^n})_{n\in\mathbb{N}}$. Some subsequence converges to a geodesic ray α. By hypothesis, all rays are contracting, so there exists some ρ_α such that α is ρ_α–contracting.

Pick any $i \geq \kappa'(\rho_\alpha, 1, 0) > \kappa(\rho_\alpha, 1, 0)$. There is some $n \gg i$ such that α agrees with α^{ζ^n} for distance $R_i + \kappa(\rho_{\zeta^i}, 1, 0)$. Since the neighborhoods are nested, by construction, $\zeta^n \in U(\zeta^i, R_i)$, which implies that α comes $\kappa(\rho_{\zeta^i}, 1, 0)$–close to α^{ζ^i} outside the ball of radius $R_i \geq R'_i$ about $\mathbf{1}$. The definition of R'_i and the fact that $i > \kappa(\rho_\alpha, 1, 0)$ gives us, by Lemma 4.7, that α^{ζ^i} comes $\kappa(\rho_\alpha, 1, 0)$–close to α outside the ball of radius R''_i about $\mathbf{1}$. In particular, by Lemma 4.5, α passes $(\kappa'(\rho_\alpha, 1, 0) + \kappa(\rho_{\zeta^i}, 1, 0))$-close to both endpoints of a subsegment of $g_i\gamma^i$ containing $g_i\gamma^i_{[-r_i,r_i]}$. The definition of r_i and the fact that $i \geq \kappa(\rho_\alpha, 1, 0)$ give us, by a second application of Lemma 4.7 and Lemma 4.5, that α comes within distance $\kappa'(\rho_\alpha, 1, 0)$ of both endpoints of $g_i\delta^i$. Connect the endpoints of $g_i\delta^i$ to α by shortest geodesic segments. For $A' := A + 2\kappa'(\rho_\alpha, 1, 0)$, the resulting path δ'_i is a continuous (L, A')–quasi-geodesic that is contained in $\bar{N}_{\kappa'(\rho_\alpha,L,A')}(\alpha)$ but leaves the $(i - \kappa'(\rho_\alpha, 1, 0))$–neighborhood of the subsegment of α between its endpoints. For sufficiently large i this contradicts the fact that δ'_i is (L, A')–quasi-geodesic.

We now prove (14). The idea is to take an element g that pushes γ far out along α^ζ and take η to be one of the endpoints of $g\gamma$. Then α^η forms a geodesic triangle with a subsegment of α^ζ and a subray of $g\gamma$. Additionally, we arrange for $g\gamma_{[-r,r]}$ to be suitably far from the quasi-center of this triangle so that it is in one of the thin legs of the triangle, parallel to a segment of α^η.

Let $\alpha := \alpha^\zeta$. Let $r' > r$ represent a number to be determined, and choose any $s > r' + 2\kappa'(\rho_\gamma, 1, 0)$. First, suppose that for arbitrarily large t there exists $g \in N_{\kappa'(\rho_\gamma,1,0)}\alpha_t$ such that $d(g\gamma_s^{-1}\gamma_{r'}, \alpha_{[0,t]}) \geq \kappa'(\rho_\gamma, 1, 0)$. We claim that for any sufficiently large t we can take such a g and $\eta := g\gamma_s^{-1}\gamma_{-\infty}$ as the output of (14). To see this, define a continuous quasi-geodesic by following α until we reach the first of either α_t or a point of $\pi_\alpha(g\gamma_s^{-1}\gamma_{r'})$, then follow a geodesic to $g\gamma_s^{-1}\gamma_{r'}$, then follow $g\gamma_s^{-1}\bar{\gamma}$ towards η. By an argument similar to Lemma 4.8, β is an (L, A)–quasi-geodesic, for some L and A not depending on g, r, or s. Now, $\alpha_{[0,t]}$, $g\gamma_s^{-1}\gamma_{(-\infty,s]}$, and α^η form a $\kappa'(\rho_\gamma, 1, 0)$–almost geodesic triangle, so the contraction function of α^η is bounded in terms of the contraction functions of α and γ. Thus, there is some E such that β and α^η are bounded Hausdorff distance E from one another, independent of our choices. By the hypothesis on g and Corollary 4.4 the two sides $\alpha_{[0,t]}$ and $g\gamma_s^{-1}\gamma_{(-\infty,s]}$ of the almost geodesic triangle are diverging at a linear rate, and so $g\gamma_s^{-1}\gamma_{r'}$ is H–close to some point of α^η for some H, again independent of our choices. Assume that we chose $r' \geq r + (M + 1)H + 9\kappa(\rho_\gamma, 1, 0)$. Then, by Lemma 4.7, we have that α^η passes within distance $\kappa(\rho_\gamma, 1, 0)$ of some point of $g\gamma_s^{-1}\gamma_{[r,r']}$, and also of some point in $g\gamma_s^{-1}\gamma_{[-r',-r]}$.

We also need to show $\eta \in U(\zeta, R)$. Any continuous (L', A')–quasi-geodesic in η stays bounded Hausdorff distance H' from β, with bound depending on L' and A', but not g, s, or t. We only need to consider $L' < \sqrt{R/3}$ and $A' < R/3$, so we can bound H' in terms of R (and ρ_α and ρ_γ). We therefore have that such a quasi-geodesic passes $(H' + H)$–close to $g\gamma_s^{-1}\gamma_{r'}$, which is $(s - r' + \kappa'(\rho_\gamma, 1, 0))$–close to α_t. Applying Lemma 4.7 we see that such a geodesic passes $\kappa(\rho_\alpha, L', A')$–close to α outside the ball of radius R provided that t is chosen sufficiently large with respect to R, s, and the contraction functions for α and γ. By hypothesis, we can choose t as large as we like, so in this case we are done.

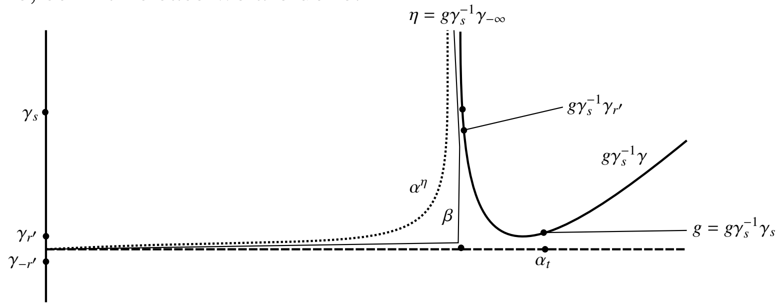

FIGURE 6. First case for (14).

The other case is that there exists T such that for every $g \in N_{\kappa'(\rho_\gamma,1,0)}\alpha_{[T,\infty]}$ we have $d(g\gamma_s^{-1}\gamma_{r'}, \alpha_{t'}) < \kappa'(\rho_\gamma, 1, 0)$ for some t' with $t > T$ and $t - t' = s - r' \pm 2\kappa'(\rho_\gamma, 1, 0) > 0$. Let w be the word in the generators for G read along the path $\gamma_{[r',s]}$. Let $t_0 > T$ be arbitrary, and let $g_0 := \alpha_{t_0}$. Let $g_1 := g_0\gamma_s^{-1}\gamma_{r'}$. By hypothesis there is a $t_1 < t_0$ such that $d(g_1, \alpha_{t_1}) \leq \kappa'(\rho_\gamma, 1, 0)$. The segment $g_0\gamma_s^{-1}\gamma_{[r',s]}$ has edge label w and, by Theorem 3.4, is contained in $\bar{N}_K(\alpha)$ for some K depending only on $\kappa'(\rho_\gamma, 1, 0)$ and ρ_γ. If $t_1 > T$ we can repeat, setting $g_2 := g_1\gamma_s^{-1}\gamma_{r'}$, so that the initial vertex g_1 of $g_1\gamma_s^{-1}\gamma_{[r',s]}$ agrees with the terminal vertex of $g_0\gamma_s^{-1}\gamma_{[r',s]}$. Repeating this construction until $t_i \leq T$, we construct a path from α_t to $\bar{N}_{\kappa'(\rho_\gamma,1,0)}\alpha_{[0,T]}$ that is contained in the K–neighborhood of α and whose edge label is a power of w^{-1}. Since this is true for arbitrarily large t, we conclude that w is a contracting element in G and $\zeta = hw^\infty$ for some $h \in G$ that is $(s - r' + \kappa'(\rho_\gamma, 1, 0))$–close to α_T. Furthermore, we can also take $s' > s$ arbitrarily large and run the same argument to conclude that either we find the g and η we are looking for from the first case, or else arbitrarily long segments $\gamma_{[r',s']}$ can be sent into $\bar{N}_K\alpha_{[T,\infty)}$. We already know this tube contains an infinite path labelled by powers of w. Therefore, there is f which is $(s - r' + 2\kappa'(\rho_\gamma, 1, 0))$–close to γ_s such that $\gamma_\infty = fw^\infty$.

If $fw^{-\infty} = \gamma_{-\infty}$ then we can take $\eta := \zeta$ and $g := hw^a f^{-1}$ for a sufficiently large. Otherwise, for any sufficiently large t and a we can take $g := \alpha_t\gamma_s^{-1} fw^{-a}$ and $\eta := g\gamma_{-\infty}$, see Figure 7. □

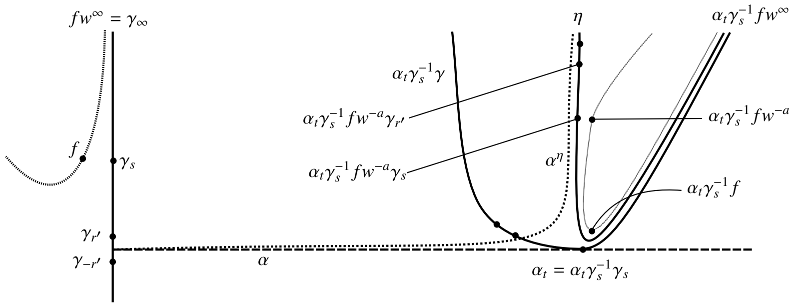

Figure 7. Second case for (14).

Proposition 10.4. *Let G be a group acting properly discontinuously on a proper geodesic metric space X. Suppose that the orbit map $\phi\colon\ g \mapsto go$ takes contracting quasi-geodesics to contracting quasi-geodesics and has quasi-convex image. If Go has compact closure in $\hat{X}$ then G is hyperbolic. If G is infinite and hyperbolic and ϕ is Morse-controlled then $\Lambda(G)$ is non-empty and compact and Go has compact closure in $\hat{X}$.*

Proof. If G is finite then Go is compact and G is hyperbolic, so assume G is infinite. Since it has a quasi-convex orbit, G is finitely generated and ϕ is a quasi-isometric embedding. Since G is infinite, Go is unbounded, so if it has compact closure then $\Lambda(G)$ is non-empty and compact. By Proposition 6.9, $\partial_c^{\mathcal{FQ}}G$ is compact. By Theorem 10.1, G is hyperbolic.

Conversely, if G is hyperbolic then, with respect to any finite generating set, geodesics in G are uniformly Morse. Therefore, there is a μ such that for any two points g, $h \in G$ there is a μ–Morse geodesic $\gamma^{g,h}$ in G from g to h. If ϕ is Morse-controlled, there exists a μ' depending only on μ such that $\phi(\gamma^{g,h})$ is μ'–Morse. Any (L, A)–quasi-geodesic with endpoints go and ho therefore stays $\mu'(L, A)$–close to $\phi(\gamma^{g,h})$, but $\phi(\gamma^{g,h})$ is a quasi-geodesic with integral points on Go, so it remains close to Go. Therefore, Go is a Morse subset of X. By Proposition 6.9, $\partial_c\phi\colon\ \partial_c^{\mathcal{FQ}}G \to \Lambda(G)$ is a homeomorphism. Since G is an infinite hyperbolic group, $\partial_c^{\mathcal{FQ}}G$ is non-empty and compact, so $\Lambda(G)$ is as well.

It remains to show $\overline{Go}$ is compact. Suppose $\mathcal{U}$ is an open cover of $\overline{Go}$. Only finitely many elements of $\mathcal{U}$ are required to cover $\Lambda(G)$. We claim that the part of Go not covered by these finitely many sets is bounded, hence, finite, so only finitely many more elements of $\mathcal{U}$ are required to cover all of $\overline{Go}$. To see this, suppose $(g_n o)_{n\in\mathbb{N}}$ is an unbounded sequence in Go that does not enter the chosen finite cover of $\Lambda(G)$. By passing to a subsequence, we may assume $d(o, g_n o) \geq n$, in which case $(g_n o)_{n\in\mathbb{N}}$ is a sequence with no convergent subsequence in $\hat{X}$. For each n pick a geodesic γ^n from o to $g_n o$. A subsequence $(\gamma^{\sigma(n)})_{n\in\mathbb{N}}$ converges to a geodesic ray γ in X based at o, but since the geodesics γ^n were uniformly contracting, γ is contracting. Moreover, by uniform contraction the endpoints $g_{\sigma(n)}o$ converge to γ_∞ in $\hat{X}$, which is a contradiction. □

References

[1] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham. *Largest acylindrical actions and stability in hierarchically hyperbolic groups*. preprint. 2017. arXiv: 1705.06219.

[2] Yael Algom-Kfir. *Strongly contracting geodesics in outer space*. Geom. Topol. **15**.4 (2011), pp. 2181–2233.

[3] Tarik Aougab, Matthew Gentry Durham, and Samuel J. Taylor. *Pulling back stability with applications to Out(F_n) and relatively hyperbolic groups*. J. Lond. Math. Soc. **96**.3 (2017), pp. 565–583.

[4] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Characterizations of Morse geodesics via superlinear divergence and sublinear contraction*. Doc. Math. **22** (2017), pp. 1193–1224. arXiv: 1601.01897.

[5] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. *Contracting Geodesics in Infinitely Presented Graphical Small Cancellation Groups*. preprint. 2016. arXiv: 1602.03767v2.

[6] Jason Behrstock. *A counterexample to questions about boundaries, stability, and commensurability*. preprint. 2017. arXiv: 1705.03984v2.

[7] Jason A. Behrstock. *Asymptotic geometry of the mapping class group and Teichmüller space*. Geom. Topol. **10** (2006), pp. 1523–1578.

[8] Mladen Bestvina, Kenneth Bromberg, and Koji Fujiwara. *Constructing group actions on quasi-trees and applications to mapping class groups*. Publ. Math. Inst. Hautes Études Sci. **122**.1 (2015), pp. 1–64.

[9] Mladen Bestvina and Koji Fujiwara. *A characterization of higher rank symmetric spaces via bounded cohomology*. Geom. Funct. Anal. **19**.1 (2009), pp. 11–40.

[10] Brian H. Bowditch. *Relatively hyperbolic groups*. Internat. J. Algebra Comput. **22**.3 (2012), pp. 1250016, 66.

[11] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Vol. 319. Grundlehren der mathematischen Wissenschaften. Berlin: Springer, 1999.

[12] Christopher H. Cashen. *Quasi-isometries Need Not Induce Homeomorphisms of Contracting Boundaries with the Gromov Product Topology*. Anal. Geom. Metr. Spaces **4**.1 (2016), pp. 278–281.

[13] Ruth Charney and Devin Murray. *A rank-one CAT(0) group is determined by its Morse boundary*. preprint. 2017. arXiv: 1707.07028v1.

[14] Ruth Charney and Harold Sultan. *Contracting boundaries of* CAT(0) *spaces*. J. Topol. **8**.1 (2015), pp. 93–117.

[15] Matthew Cordes. *A survey on Morse boundaries & stability*. preprint. 2017. arXiv: 1704.07598v1.

[16] Matthew Cordes. *Morse Boundaries of Proper Geodesic Metric Spaces*. Groups Geom. Dyn. **11**.4 (2017), pp. 1281–1306.

[17] Matthew Cordes and Matthew Gentry Durham. *Boundary convex cocompactness and stability of subgroups of finitely generated groups*. preprint. 2016. arXiv: 1607.08899v1.

[18] Matthew Cordes and David Hume. *Stability and the Morse boundary*. J. Lond. Math. Soc. (in press).

[19] Christopher B. Croke and Bruce Kleiner. *Spaces with nonpositive curvature and their ideal boundaries*. Topology **39** (2000), pp. 549–556.

[20] François Dahmani, Vincent Guirardel, and Denis V. Osin. *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*. Mem. Amer. Math. Soc. **245**.1156 (2017), pp. v+152.

[21] Cornelia Druţu. *Relatively hyperbolic groups: geometry and quasi-isometric invariance*. Comment. Math. Helv. **84**.3 (2009), pp. 503–546.

[22] Cornelia Druţu, Shahar Mozes, and Mark Sapir. *Divergence in lattices in semisimple Lie groups and graphs of groups*. Trans. Amer. Math. Soc. **362**.5 (2010), pp. 2451–2505.

[23] Cornelia Druţu and Mark Sapir. *Tree-graded spaces and asymptotic cones of groups*. Topology **44**.5 (2005). With an appendix by Denis Osin and Sapir, pp. 959–1058.

[24] Moon Duchin and Kasra Rafi. *Divergence of geodesics in Teichmüller space and the mapping class group*. Geom. Funct. Anal. **19**.3 (2009), pp. 722–742.

[25] Matthew Gentry Durham and Samuel J. Taylor. *Convex cocompactness and stability in mapping class groups*. Algebr. Geom. Topol. **15**.5 (2015), pp. 2839–2859.

[26] Benson Farb. *Relatively Hyperbolic Groups*. Geom. Funct. Anal. **8**.5 (1998), pp. 810–840.

[27] Daniel Groves and Jason Fox Manning. *Dehn filling in relatively hyperbolic groups*. Israel J. Math. **168**.1 (2008), pp. 317–429.

[28] Ursula Hamenstädt. *Rank-one isometries of proper* CAT(0)*-spaces*. *Discrete groups and geometric structures*. Vol. 501. Contemp. Math. Amer. Math. Soc., Providence, RI, 2009, pp. 43–59.

[29] G. Christopher Hruska. *Relative hyperbolicity and relative quasiconvexity for countable groups*. Algebr. Geom. Topol. **10**.3 (2010), pp. 1807–1856.

[30] Jason Fox Manning. *The Bowditch boundary of* $(G,\mathcal{H})$ *when G is hyperbolic*. preprint. 2015. arXiv: `1504.03630v2`.

[31] Howard A. Masur and Yair N. Minsky. *Geometry of the complex of curves. I. Hyperbolicity*. Invent. Math. **138**.1 (1999), pp. 103–149.

[32] Yair N. Minsky. *Quasi-projections in Teichmüller space*. J. Reine Angew. Math. **473** (1996), pp. 121–136.

[33] Devin Murray. *Topology and Dynamics of the Contracting Boundary of Cocompact CAT(0) Spaces*. preprint. 2015. arXiv: `1509.09314v2`.

[34] Alexander Yu. Ol′shanskii, Denis V. Osin, and Mark V. Sapir. *Lacunary hyperbolic groups*. Geom. Topol. **13**.4 (2009). With an appendix by Michael Kapovich and Bruce Kleiner, pp. 2051–2140.

[35] Denis V. Osin. *Acylindrically hyperbolic groups*. Trans. Amer. Math. Soc. **368**.2 (2016), pp. 851–888.

[36] Denis V. Osin. *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*. Mem. Amer. Math. Soc. **179**.843 (2006), pp. vi+100.

[37] Panos Papasoglu. *Strongly geodesically automatic groups are hyperbolic*. Invent. Math. **121**.2 (1995), pp. 323–334.

[38] Alessandro Sisto. *Projections and relative hyperbolicity*. Enseign. Math. (2) **59**.1-2 (2013), pp. 165–181.

[39] Alessandro Sisto. *Quasi-convexity of hyperbolically embedded subgroups*. Math. Z. **283**.3-4 (2016), pp. 649–658.

[40] Harold Sultan. *Hyperbolic quasi-geodesics in CAT(0) spaces*. Geom. Dedicata **169**.1 (2014), pp. 209–224.

[41] Hung Cong Tran. *Divergence spectra and Morse boundaries of relatively hyperbolic groups*. Geometriae Dedicata (2017).