Cannon–Thurston maps, subgroup distortion, and hyperbolic hydra

Owen Baker[*]

*Department of Mathematics, 900 University Avenue,
University of California, Riverside, CA 92521, USA*

Timothy Riley[†]

*Department of Mathematics, 310 Malott Hall,
Cornell University, Ithaca, NY 14853, USA*

May 9, 2018

Abstract

There is a family of hyperbolic groups known as *hyperbolic hydra* which contain heavily distorted free subgroups. We prove the existence of Cannon–Thurston maps (that is, maps of the boundaries induced by subgroup inclusion) for these free subgroups. It is known that Cannon–Thurston maps between hyperbolic *space* boundaries can exist even in the presence of arbitrarily heavy (even non-recursive) distortion. The hyperbolic hydra examples show that Cannon–Thurston maps can exist even between hyperbolic *group* boundaries in the presence of arbitrarily heavy primitive recursive distortion.

1 Introduction

An isometry of hyperbolic n-space $\mathbb{H}^n$ induces a homeomorphism on $\partial \mathbb{H}^n = S^{n-1}$, the sphere at infinity. More generally, an isometric embedding $\mathbb{H}^m \hookrightarrow \mathbb{H}^n$ induces an embedding $S^{m-1} \hookrightarrow S^{n-1}$; a quasi-isometric embedding $X \to Y$ of (Gromov-)hyperbolic spaces induces an embedding $\partial X \hookrightarrow \partial Y$ of the Gromov boundaries. Remarkably, many natural embeddings which are far from being isometric nonetheless induce maps on the boundaries. These are known as *Cannon–Thurston maps* in honor of the researchers who gave the first exotic example: let M be a hyperbolic 3-manifold fibering over the circle with hyperbolic surface fiber S and pseudo-Anosov monodromy. The inclusion $S \hookrightarrow M$ of a fiber induces an embedding $f : \mathbb{H}^2 = \widetilde{S} \hookrightarrow \widetilde{M} = \mathbb{H}^3$ of the universal covers, which they showed induces a (surjective!) map $\hat{f} : S^1 \to S^2$ [12].

[*]otb2@cornell.edu

Cannon & Thurston's example can be viewed as a map $\partial \pi_1 S = S^1 \to S^2 = \partial \mathbb{H}^3$ induced by an orbit map $\pi_1 S \to \mathbb{H}^3$ of the surface Kleinian group $\pi_1 S$. Thus one way of generalizing the Cannon–Thurston example involves replacing $\pi_1 S$ with an arbitrary finitely generated Kleinian group. In a series of papers culminating in [27], Mahan Mj (formerly Mitra) showed that for any finitely generated Kleinian group G, the orbit map $G \to \mathbb{H}^3$ always extends continuously to the the boundary, inducing a *Cannon–Thurston map* $\partial G \to S^2$.

Another interpretation of Cannon & Thurston's example is as a map $\partial \pi_1 S = S^1 \to S^2 = \partial \pi_1 M$ induced by the inclusion of (Gromov)-hyperbolic groups $\pi_1 S \hookrightarrow \pi_1 M$. (Section 2 contains background on hyperbolic groups and their boundaries.) This leads to another direction of generalization: Let $\Lambda \le \Gamma$ be hyperbolic groups. If the inclusion map $f : \Lambda \hookrightarrow \Gamma$ extends to a continuous map $\overline{f} : \Lambda \cup \partial \Lambda \to \Gamma \cup \partial \Gamma$ of the Gromov compactifications, then $\overline{f}$ (or its restriction $\hat{f} : \partial \Lambda \to \partial \Gamma$) is called a *Cannon–Thurston map*. When a Cannon–Thurston map exists, it is unique.

Mj showed this map exists when Λ is an infinite hyperbolic normal subgroup of a hyperbolic group Γ [23]. He also showed it exists when Γ is a hyperbolic group which is a finite graph of hyperbolic groups, with Λ one of the (infinite) vertex- or edge-groups, under the assumption that the edge inclusions are quasi-isometric embeddings [24]. We gave the first example of a hyperbolic group with hyperbolic subgroup, for which there is no Cannon–Thurston map [1]. Matsuda & Oguni showed our example leads to examples where the subgroup in question can be any non-elementary hyperbolic group or, even, relatively hyperbolic group [20].

The fact that the original Cannon–Thurston map $\hat{f} : S^1 \to S^2$ of [12] is surjective (space-filling) stems from the difference between the intrinsic metric of the hyperbolic surface fiber S and the ambient metric from the hyperbolic 3-manifold M. That is, *from the fact that $\widetilde{S} \hookrightarrow \widetilde{M}$ (or $\pi_1 S \hookrightarrow \pi_1 M$) is distorted*.

For a finitely generated subgroup Λ of a finitely generated group Γ, define the *distortion function*

$$\text{Dist}_\Lambda^\Gamma(n) := \max\{d_\Lambda(e,h) \mid h \in \Lambda, d_\Gamma(e,h) \le n\},$$

where d_Γ and d_Λ are word metrics with respect to some finite generating sets. We say that $f \preceq g$ for $f, g : \mathbb{N} \to \mathbb{N}$ when there exists $C > 0$ such that $f(n) \le Cg(Cn+C)+Cn+C$ for all $n \ge 0$. We say $f \simeq g$ when $f \preceq g$ and $g \preceq f$. Up to $\simeq$, $\text{Dist}_\Lambda^\Gamma(n)$ does not depend on the choices of finite generating sets. A similar definition applies in the Kleinian groups context:

$$\text{Dist}_G^{\mathbb{H}^3}(n) := \max\left\{ d_G(e,g) \mid g \in G, d_{\mathbb{H}^3}(x_0, g \cdot x_0) \le n \right\}.$$

If Λ is an undistorted subgroup in a hyperbolic group Γ (that is, $\text{Dist}_\Lambda^\Gamma(n) \preceq n$), then Λ is also hyperbolic (e.g. [8, page 461]) and the Cannon–Thurston map $\partial \Lambda \to \partial \Gamma$ is readily seen to exist and be injective.

It is natural to ask whether extreme distortion is an obstacle to the existence of a Cannon–Thurston map. In the Kleinian group setting, it is not. Cannon–Thurston maps always exist for surface Kleinian groups: McMullen [21] proved this in the punctured torus case and Mj [27] in the general case. And Mj proved that distortion functions can be arbitrarily wild in this setting [24, p.160–161], even non-recursive, by a construction based on ideas of Minsky from [22].

For hyperbolic subgroups of hyperbolic groups, the relation (or lack thereof) between distortion and the existence of Cannon–Thurston maps is less clear. As we mentioned earlier, Mj proved they exist for infinite normal hyperbolic subgroups of hyperbolic groups. These are never more than exponentially distorted. Further, Mj established in [24, p.159–160] that they exist for all $k \ge 1$ for certain examples (based on a construction of Bestvina, Feighn, & Handel [4]) which display k-fold iterated exponential distortion. There are examples of heavier distortion where the existence of the Cannon–Thurston map remains unknown, specifically Mj's example where the distortion exceeds a k-fold iterated exponential for all k (but by Corollary 3 of [3] (unpublished) is dominated by the function A_4 discussed below), and the closely related CAT(-1) examples of Barnard, Brady & Dani [2].

In this article we show that arbitrarily fast-growing primitive recursive distortion is no barrier to the existence of Cannon–Thurston maps for hyperbolic subgroups of hyperbolic groups.

The *hyperbolic hydra* Γ_k ($k = 1, 2, \ldots$) of [7] are a family of hyperbolic groups with finite-rank free subgroups Λ_k exhibiting the fastest-growing distortion functions known for hyperbolic subgroups of hyperbolic groups: $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k}$ grows at least like A_k, the k-th of Ackermann's family of increasingly fast-growing functions that begins with $A_1(n) = 2n$, $A_2(n) = 2^n$, and $A_3(n)$ equalling the height-n tower of powers of 2. Any primitive recursive function is dominated by some A_k (see for example pages 11–21 of [11]). We will prove in Section 4:

Theorem 1.1. *Hyperbolic hydra have Cannon–Thurston maps $\partial \Lambda_k \to \partial \Gamma_k$ for all k.*

While heavy distortion may fail to obstruct the existence of Cannon–Thurston maps, it is natural to expect it to manifest as some sort of wildness in the map: witness, for example, how Cannon & Thurston's original example is a space-filling curve $S^1 \twoheadrightarrow S^2$ [12]. More generally, wildness manifests in Cannon–Thurston maps in the relationship between "ε" and "δ" in their continuity (having given the boundaries natural metrics). This was made precise by Miyachi [25] in the Kleinian group setting when the group is a finitely generated Fuchsian group of the first kind with bounded geometry and no parabolic elements: Miyachi [25] gives an upper bound on the modulus of continuity for the Cannon–Thurston map, and shows it is not Hölder continuous when the group is geometrically infinite.

In this article we will establish the corresponding result for hyperbolic subgroups of hyperbolic groups. Here are more details. The *modulus of continuity* $\varepsilon : [0, \infty) \to [0, \infty]$ of a function $f : U \to V$ between metric spaces is

$$\varepsilon(\delta) := \sup\{d_V(f(a), f(b)) \mid a, b \in U \text{ with } d_U(a, b) \le \delta\}.$$

This notion goes back at least to Lebesgue in 1909 [19]. An upper bound on $\varepsilon(\delta)$ expresses a degree of good behaviour: f is uniformly continuous if $\varepsilon(\delta) \to 0$ as $\delta \to 0$; is Lipschitz if $\varepsilon(\delta) \le C\delta$ for a constant $C > 0$; and is α-Hölder if $\varepsilon(\delta) \le C\delta^\alpha$ for a constant $C > 0$. Wildness manifests in lower bounds on $\varepsilon(\delta)$, expressing that $\varepsilon(\delta)$ is extravagantly larger than δ when δ is close to 0. We will prove:

Theorem 1.2. *Suppose $\Lambda \le \Gamma$ are hyperbolic, Λ is non-elementary, the Cannon–Thurston map $\partial \Lambda \to \partial \Gamma$ exists, and $r, s > 1$ are any visual parameters for visual metrics on $\partial \Lambda$ and $\partial \Gamma$, respectively. Then there exist $\alpha, \beta > 0$ so that for all $n \ge 0$ the modulus of continuity satisfies*

$$\varepsilon\left(\frac{\beta}{r^{\mathrm{Dist}_\Lambda^\Gamma(n)}}\right) \ge \frac{\alpha}{s^n}.$$

While we have not been able to find this theorem elsewhere in the literature, we understand that it is known to some experts. Indeed, the second author first heard of a suggestion of a relationship between the modulus of continuity of Cannon–Thurston map and subgroup distortion from Mahan Mj. The moduli of continuity of Cannon–Thurston maps have received attention before in the Kleinian group setting. For a faithful discrete representation to $\mathrm{PSL}_2(\mathbb{C})$ of a finitely generated Fuchsian group of the first type with bounded geometry and no parabolic elements, Miyachi [25] gives an upper bound on the modulus of continuity of the Cannon–Thurston map.

The reason we insist that Λ be non-elementary (i.e. contains an F_2 subgroup; see e.g. [16, Theorem 2.28]) in this theorem is that elementary Λ are not interesting in this context: if Λ is finite, then $\partial \Lambda$ is empty; if Λ is virtually $\mathbb{Z}$, then Λ is quasiconvex in G, its boundary $\partial \Lambda$ is two points, and the Cannon–Thurston map exists and is an embedding.

Theorem 1.1 of [7] says that $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k} \succeq A_k$. It combines with Theorems 1.1 and 1.2 to give that the Cannon–Thurston map $\partial \Lambda_k \to \partial \Gamma_k$ has the property that for δ incredibly small, $\varepsilon(\delta)$ is, by comparison, huge. More precisely, we will prove in Section 5:

Corollary 1.3. *Fix $k \ge 3$ and any visual metrics on $\partial \Lambda_k$ and $\partial \Gamma_k$. The modulus of continuity for the Cannon–Thurston map $\partial \Lambda_k \to \partial \Gamma_k$ for hyperbolic hydra satisfies*

$$\varepsilon\left(\frac{1}{A_{k-1}(n)}\right) \ge \frac{1}{n}$$

for all sufficiently large n. In particular, for any primitive recursive function f there exists a Cannon–Thurston map between hyperbolic groups satisfying

$$\varepsilon\left(\frac{1}{f(n)}\right) \geq \frac{1}{n}$$

for all sufficiently large n.

(The final part of the corollary follows immediately, since every primitive recursive function is dominated by some A_k, as we remarked earlier, citing [11].)

A detailed understanding of the Cannon–Thurston Map $\partial\Lambda_k \to \partial\Gamma_k$ appears hard to obtain. Whilst $\partial\Lambda_k$ is a Cantor set (as Λ_k is free), $\partial\Gamma_k$ is not so readily identified. (I. Kapovich & M. Lustig [17] recently made advances in the understanding of Cannon–Thurston maps for certain free-by-cyclic groups, but $\Lambda_k \leq \Gamma_k$ do not fall within the scope of their work.) Here is what we can say about $\partial\Gamma_k$.

Splittings of hyperbolic free-by-cyclic groups $F \rtimes_\varphi \mathbb{Z}$ are studied in [18] and [9], the former dealing with the case where φ is an irreducible hyperbolic free group automorphism, and the latter with φ a general hyperbolic free group automorphism. The argument preceding Corollary 15 in [18] shows that any hyperbolic free-by-cyclic group has a one-dimensional boundary: the cohomological dimension of any (finitely generated free)-by-cyclic group is 2 (see e.g. [10, pp.185–7]), so [5, Corollary 1.4(d)] implies $\partial\Gamma_k$ has dimension $2 - 1 = 1$.

The argument of [18, Corollary 15] shows that any hyperbolic free-by-cyclic group has connected, locally connected boundary without global cut points. To see this, it suffices by [16, Theorems 7.1 and 7.2] to check that $F \rtimes_\varphi \mathbb{Z}$ is freely indecomposable, which is true for *any* free group automorphism φ. Indeed, the Bass–Serre tree T for any graph of groups decomposition of $F \rtimes_\varphi \mathbb{Z}$ admits a minimal action by the normal subgroup F with quotient a finite graph by Grushko's Theorem. This shows the edge stabilizers for the action of $F \rtimes_\varphi \mathbb{Z}$ on T are non-trivial, so the decomposition cannot be free.

On the other hand, Γ_k splits as an HNN-extension over $\mathbb{Z}$ for every k, so [16, Theorem 7.2] implies $\partial\Gamma_k$ has local cut points. Indeed, Γ_1 splits over $\mathbb{Z}$ as $\langle B, a_1 \rangle *_{\langle a_1^{-1} t^2 u a_1 = t^2 v^{-1} \rangle}$, where B denotes the subgroup generated by all the defining generators other than a_0 and a_1 (then a_0 appears as $t^{-1} a_1 t$), and for $k \geq 2$, Γ_k splits as an HNN-extension over $\mathbb{Z}$ with the stable letter a_k conjugating t to $t a_{k-1}^{-1}$.

For additional background on Cannon–Thurston maps we recommend Mj's recent survey [26].

The organization of this article. In Section 2 we give background on hyperbolic groups and their boundaries. In Section 3 we define Cannon–Thurston maps and prove an embellished version of a lemma of Mitra giving necessary and sufficient conditions for their existence. In Section 4 we review the construction of hyperbolic hydra groups and prove Theorem 1.1. In Section 5 we prove Theorem 1.2 and Corollary 1.3.

2   Hyperbolic groups and their boundaries

This section contains a brief account of some pertinent background. More general treatments can be found in, for example, [8, 14, 16, 28] and Gromov's foundational article [15].

For a metric space X, the *Gromov product* $(a \cdot b)_e$ (or $(a \cdot b)_e^X$ if there is ambiguity) of $a, b \in X$ with respect to $e \in X$ is

$$(a \cdot b)_e = \frac{1}{2}(d(a,e) + d(b,e) - d(a,b)).$$

One says X is (δ)-*hyperbolic* when

$$(a \cdot b)_e \geq \min\{(a \cdot c)_e, (b \cdot c)_e\} - \delta$$

for all $e, a, b, c \in X$, and X is hyperbolic when it is (δ)-hyperbolic for some $\delta \geq 0$. When X is a geodesic space this is equivalent to other standard definitions of hyperbolicity (such as δ-thin or δ-slim triangles), although the δ involved may not agree.

When X is (0)-hyperbolic and geodesic—that is, an $\mathbb{R}$-tree—$(a \cdot b)_e$ is the distance from e to the geodesic between a and b. Correspondingly, in a (δ)-hyperbolic geodesic space every pair of geodesics from e to a and to b, both parametrized by arc-length, 6δ-fellow-travel for approximately $(a \cdot b)_e$ and then diverge (by the *insize* characterization of hyperbolicity of [8, page 408]). Indeed:

Lemma 2.1. *In a (δ)-hyperbolic geodesic metric space, for every geodesic $[a, b]$ connecting a and b*

$$|d(e, [a, b]) - (a \cdot b)_e| \leq 6\delta.$$

Proof. See [8]: the proof of Proposition 1.22 on page 411 shows that *insizes* of geodesic triangles are at most 6δ, and the proof of Proposition 1.17(3) $\implies$ (2) on page 409 shows that all geodesic triangles are 6δ-*thin*, and the claimed inequality follows. $\square$

The *(Gromov-) boundary* ∂X of a hyperbolic metric space X is defined with reference to, but is in fact independent of, a point $e \in X$. It is the set of equivalence classes of sequences (a_n) in X such that $(a_m \cdot a_n)_e \to \infty$ as $m, n \to \infty$, where two such sequences (a_n) and (b_n) are equivalent when $(a_m \cdot b_n)_e \to \infty$ as $m, n \to \infty$. Indeed, they are equivalent when $(a_n \cdot b_n)_e \to \infty$ as $n \to \infty$ since

$$(a_n \cdot b_m)_e \geq \min\{(a_n \cdot b_n)_e, (b_m \cdot b_n)_e\} - \delta$$

by (δ)-hyperbolicity. Denote the equivalence class of (a_n) by $\lim a_n$.

When X is a *geodesic* hyperbolic metric space, there are equivalent definitions of ∂X, such as ∂X is the set of equivalence classes of geodesic rays emanating from x, where two such rays are equivalent when they stay uniformly close. The condition $(a_m \cdot a_n)_e \to \infty$ is what makes a sequence (a_n) ray-like, and the condition $(a_m \cdot b_n)_e \to \infty$ corresponds to uniform closeness.

Extend the Gromov product to $\overline{X} := X \cup \partial X$ by

$$(a \cdot b)_e = \sup \liminf_{m,n \to \infty} (a_m \cdot b_n)_e$$

where the sup is over all sequences (a_m) and (b_n) in X representing (when in ∂X) or tending to (when in X) a and b, respectively. (The "$\sup \liminf$" is necessary—see [8, page 432].)

We note, for (3) in the following lemma, that in a proper geodesic hyperbolic metric space X, each pair of distinct points $a, b \in \partial X$ is joined by a bi-infinite geodesic line $[a, b]$ (Lemma 3.2 on page 428 of [8]).

Lemma 2.2. *Suppose X is a proper geodesic (δ)-hyperbolic metric space.*

(0). *If $x, y \in \overline{X}$ and $e \in X$, then there exist sequences (x_n) and (y_n) in X with $x = \lim x_n$, $y = \lim y_n$, and $(x \cdot y)_e = \lim_n (x_n \cdot y_n)_e$.*

(1). *If $a, b, c \in \overline{X}$ and $e \in X$, then $(a \cdot b)_e \geq \min\{(a \cdot c)_e, (c \cdot b)_e\} - 2\delta$.*

(2). *If $a, b \in X$ and $c \in \partial X$, then $|d(a, b) - (a \cdot c)_b - (b \cdot c)_a| \leq \delta$.*

(3). *If $e \in X$ and $[a, b]$ is any geodesic joining any $a, b \in \partial X$, then $|d(e, [a, b]) - (a \cdot b)_e| \leq 8\delta$.*

Proof. (0) and (1) are parts 3 and 4 of [8, page 433, Remark 3.17]. (Alternatively, see parts 3 and 5 of [28, Lemma 4.6] but note that there $\inf\liminf$ is used in place of $\sup\liminf$ and so the constants differ.)

For (2), using (0) take sequences c_n, c_n' both approaching c such that $(a \cdot c)_b = \lim(a \cdot c_n)_b$ and $(b \cdot c)_a = \lim(b \cdot c_n')_a$. Now, (δ)-hyperbolicity yields $(a \cdot c_n)_b \geq \min\{(a \cdot c_n')_b, (c_n \cdot c_n')_b\} - \delta$. As $n \to \infty$, we have $(c_n \cdot c_n')_b \to \infty$, but $(a \cdot c_n')_b$ is bounded above by $d(a, b)$. So $(a \cdot c_n')_b \leq (a \cdot c_n)_b + \delta$ for all sufficiently large n. Interchanging the roles of c_n and c_n' we find $|(a \cdot c_n)_b - (a \cdot c_n')_b| \leq \delta$ for all sufficiently large n. Hence:

$$|d(a, b) - (a \cdot c_n)_b - (b \cdot c_n')_a| \leq \delta + |d(a, b) - (a \cdot c_n')_b - (b \cdot c_n')_a| = \delta + |0| = \delta$$

for all sufficiently large n. Taking the limit as $n \to \infty$ gives the result.

For (3) (cf. Exercise 3.18(3) [8, page 433]), choose sequences $a_n \to a$ and $b_n \to b$ along $[a, b]$. Also choose $a_n', b_n' \in X$ as in (0) so that $(a \cdot b)_e = \lim(a_n' \cdot b_n')_e$. For large enough n, the closest point of $[a, b]$ to e lies on $[a_n, b_n]$, so

$$|d(e, [a, b]) - (a_n \cdot b_n)_e| = |d(e, [a_n, b_n]) - (a_n \cdot b_n)_e| \leq 6\delta, \tag{1}$$

with the inequality being by Lemma 2.1. By the (δ)-hyperbolicity condition,

$$(a_n \cdot b_n)_e \geq \min\{(a_n \cdot a_n')_e, (a_n' \cdot b_n')_e, (b_n \cdot b_n')_e\} - 2\delta, \text{ and}$$
$$(a_n' \cdot b_n')_e \geq \min\{(a_n \cdot a_n')_e, (a_n \cdot b_n)_e, (b_n \cdot b_n')_e\} - 2\delta.$$

As $n \to \infty$ both $(a_n \cdot a_n')_e \to \infty$ and $(b_n \cdot b_n')_e \to \infty$, but $\limsup(a_n \cdot b_n)_e$ and $\limsup(a_n' \cdot b_n')_e$ are bounded above by $(a \cdot b)_e + 1$ (else, passing to subsequences, we can assume $(a_n \cdot b_n)_e > (a \cdot b)_e + 1/2$ for all n, and so $\liminf(a_n \cdot b_n)_e \geq (a \cdot b)_e + 1/2$ contrary to the definition of $(a \cdot b)_e$). So these two inequalities together give $|(a_n \cdot b_n)_e - (a_n' \cdot b_n')_e| \leq 2\delta$ for all sufficiently large n. Combining this with (1) gives the result. □

Visual metrics are natural metrics on the boundary ∂X of a (δ)-hyperbolic space X. Their essence is that $a, b \in \partial X$ are close when geodesics from a basepoint $e \in X$ to a and from e to b fellow travel for a long distance. One might hope that if $r > 1$, then $d(a, b) = r^{-(a \cdot b)_e}$ would define such a metric, but unfortunately, as such, transitivity can fail. Instead, say that a metric d on ∂X is a *visual metric* with *visual parameter* $r > 1$ when there exist $k_1, k_2 > 0$ such that for all $a, b \in \partial X$,

$$k_1 r^{-(a \cdot b)_e} \leq d(a, b) \leq k_2 r^{-(a \cdot b)_e}. \tag{2}$$

Lemma 2.3. *Suppose X is a (δ)-hyperbolic space, $r > 1$, and $e \in X$ is the base point with respect to which ∂X is defined. Then there is a visual metric on ∂X with parameter r. Moreover, any two visual metrics d and d' on ∂X (perhaps with different r and e) are Hölder-equivalent in that there exists $\alpha > 0$ such that the identity map $(\partial X, d) \to (\partial X, d')$ is α-Hölder and its inverse is $(1/\alpha)$-Hölder. In particular, the visual topology on ∂X is independent of these choices.*

The existence claim is [8, Proposition 3.21, page 435]. Hölder-equivalence is known (see [16, Theorem 2.18]), and follows, in this generality, immediately from the definition of visual metric and the fact that $(a \cdot b)_e \geq (a \cdot b)_{e'} - d(e, e')$.

We will need that $\overline{X} := X \cup \partial X$ is a compactification of X:

Lemma 2.4. *If X is a proper (δ)-hyperbolic geodesic metric space, then there is a unique compact metrizable topology on $\overline{X} := X \cup \partial X$ such that: the inclusions of X and of ∂X are homeomorphic onto their images, ∂X is closed, and a sequence x_n in X converges to $x \in \partial X$ if and only if (x_n) is in the equivalence class x.*

Proof. Uniqueness follows from the fact that for metrizable spaces, the topology is determined by knowledge of which sequences converge and to which points they converge. For a sequence x_n in $\overline{X}$, $x_n \to x \in X$ in the topology of $\overline{X}$ precisely when all but finitely many $x_n \in X$ and $x_n \to x$ in the topology of X. Also, $x_n \to x \in \partial X$ in the topology of $\overline{X}$ precisely when: (i) the subsequence of x_n consisting of points in ∂X is

finite or converges to x in the topology of ∂X, AND (ii) the subsequence of x_n consisting of points in X is finite or represents the equivalence class x.

The topology on $\overline{X}$ is constructed and sequential compactness is proved in [8, page 430, III.H.3.7] and metrizability in [8, page 433 III.H.3.18(4)]. For metric spaces, sequential compactness is equivalent to compactness. The agreement with the topology on ∂X coming from the visual metric is [8, page 435 III.H.3.21]. The characterization of convergent sequences follows from [8, page 431, III.H.3.13]. □

Given a finitely generated group G with finite generating set A, one forms the Cayley graph $C_A(G)$ with vertex set G and edge set $\{\{v, w\} \mid v, w \in G, \, v\,w^{-1} \in A \cup A^{-1}\}$. The graph metric on $C_A(G)$ induces the *word metric* on the vertex set G and G is called a (δ)-*hyperbolic group* if the metric space G is (δ)-hyperbolic.

Lemma 2.5. *If G is a hyperbolic group, then there is a unique compact metrizable topology on $\overline{G} := G \cup \partial G$ such that: the inclusions of G and of ∂G are homeomorphic onto their images, ∂G is closed, and a sequence g_n in G converges to $g \in \partial G$ if and only if (g_n) is in the equivalence class g.*

Proof. In this case, $C_A(G)$ is a proper hyperbolic geodesic metric space, to which Lemma 2.4 applies. One may identify ∂G with $\partial C_A(G)$ as visual metric spaces and thus identify $\overline{G} := G \cup \partial G$ with a subspace of $\overline{C_A(G)} = C_A(G) \cup \partial G$, and so endow $\overline{G}$ with the subspace topology. The stated properties of $\overline{G}$ are now a consequence of Lemma 2.4, and uniqueness follows as it does in the proof of Lemma 2.4. □

Changing the finite generating set A induces a quasi-isometry and so does not affect whether G is hyperbolic (Theorem 1.9 of [8, page 402, III.H.1]), nor does it affect the topology on ∂G (Theorem 3.9 of [8, page 430, III.H.3.7]) or $\overline{G}$ (Lemma 3.1(b) ⟺ (c) below).

3 Cannon–Thurston maps and Mitra's Lemma

Given two hyperbolic groups Λ and Γ and an injective homomorphism $\iota : \Lambda \to \Gamma$, one may ask if it extends to a continuous map $\overline{\iota} : \overline{\Lambda} \to \overline{\Gamma}$. Equivalently (Lemma 3.1), one may ask whether ι induces a well-defined map $\hat{\iota} : \partial \Lambda \to \partial \Gamma$ sending $[(g_n)]$ to $[(\iota g_n)]$. When one (and hence both) exist, $\hat{\iota}$ is the restriction of $\overline{\iota}$ to $\partial \Lambda$ and is called the *Cannon–Thurston map*. The next section will prove the existence of Cannon–Thurston maps for the heavily distorted free subgroups of hyperbolic hydras. This section deals with general tools for showing the existence of Cannon–Thurston maps.

For now, let us deal with a more general setting. Consider a map $f : Y \to X$ where (X, d_X) is (δ_X)–hyperbolic and (Y, d_Y) is (δ_Y)–hyperbolic. We will assume X and Y are proper and geodesic, or else Lemma 2.4 will not provide us a topology on $\overline{X}$ and $\overline{Y}$ (though one could still ask whether $(x_n) \to (f(x_n))$ gives a well-defined continuous map $\partial X \to \partial Y$). The following lemma is an embellished version of Mitra's criterion for the Cannon–Thurston map to exist ([23, Lemma 2.1] and [24, Lemma 2.1]).

Our notation is that $B_X(e, R) = \{x \in X \mid d_X(e, x) < r\}$ and $\overline{B}_X(e, R) = \{x \in X \mid d_X(e, x) \leq r\}$, and we write $\gamma = [x, y]_X$ to mean γ *is a geodesic in X from x to y*. The metrics on ∂X and ∂Y implicit in this lemma are any visual metrics $d_{\partial X}$ and $d_{\partial Y}$.

The additional hypothesis for (f) can be removed if δ_Y, the constant of hyperbolicity for Y, is zero; for $\delta_Y > 0$, we do not know whether (f) is equivalent to or strictly weaker than (a)–(e) in its absence. An inclusion map of a subgroup into an ambient group is Lipschitz when both have word metrics coming from some finite generating sets, and so it is satisfied in that setting.

We remark that the equivalence (a)–(d) also follows for X and Y hyperbolic groups and $f : X \to Y$ an injective homomorphism, with the same proof (replacing Lemma 2.4 with Lemma 2.5).

Lemma 3.1. *Suppose (X, d_X) and (Y, d_Y) are infinite proper geodesic hyperbolic metric spaces, and $f : Y \to X$ is a proper map. Fix a basepoint $e \in Y$. Define $M, M', M'' : [0, \infty) \to [0, \infty)$ by*

$$M(N) := \inf \left\{ (f(x) \cdot f(y))^X_{f(e)} \mid x, y \in Y \text{ and } (x \cdot y)^Y_e \geq N \right\},$$

$$M'(N) := \inf \left\{ d_X(f(e), \gamma) \mid \gamma = [f(x), f(y)]_X \text{ for some } [x, y]_Y \text{ in } Y \smallsetminus B_Y(e, N) \right\},$$

$$M''(N) := \inf \left\{ d_X(f(e), \gamma) \mid \gamma = [f(z), f(y)]_X \text{ for some } z \text{ on some } [e, y]_Y \text{ with } d_Y(e, z) \geq N \right\}.$$

The following are equivalent:

(a). *$(a_n) \mapsto (f(b_n))$ induces a well-defined function $\hat{f} : \partial Y \to \partial X$.*

(b). *$(a_n) \mapsto (f(b_n))$ induces a well-defined, continuous map $\hat{f} : \partial Y \to \partial X$.*

(c). *There exists a continuous extension $\overline{f} : \overline{Y} \to \overline{X}$ of f.*

(d). *$M(N) \to \infty$ as $N \to \infty$.*

(e). *$M'(N) \to \infty$ as $N \to \infty$.*

Moreover, if $\sup\{ d_X(f(x), f(y)) \mid d_Y(x, y) \leq r \} < \infty$ for all $r \geq 0$, then these are also equivalent to

(f). *$M''(N) \to \infty$ as $N \to \infty$.*

Proof. Each of $M(N)$, $M'(N)$, and $M''(N)$ is a non-decreasing function, so is either bounded or tends to ∞ as $N \to \infty$.

That (b) $\implies$ (a) is immediate.

Here is why (a) $\implies$ (d). Suppose $M(N) \leq C$ for all N. So there are sequences (p_n) and (q_n) in Y with $(p_n \cdot q_n)^Y_e \to \infty$ but $(f(p_n) \cdot f(q_n))^X_{f(e)} \leq C$ for all n. As $Y \cup \partial Y$ is sequentially compact by Lemma 2.4, both (p_n) and (q_n) have subsequences which converge in $Y \cup \partial Y$. But the condition $(p_n \cdot q_n)^Y_e \to \infty$ precludes any such subsequence from converging in Y, so those subsequences converge to points in ∂Y, indeed to the same point.

Next we prove (d) $\implies$ (b). Suppose sequences (p_n) and (q_n) in Y both represent the same point in ∂Y. Then $(p_n \cdot q_n)^Y_e \to \infty$ as $n \to \infty$, and so $(f(p_n) \cdot f(q_n))^X_{f(e)} \to \infty$, since $M(N) \to \infty$ as $N \to \infty$. Thus if $(f(p_n))$ and $(f(q_n))$ represent points in ∂X, then those points are the same.

So, to prove $\hat{f}$ is well-defined, it suffices to show that if a sequence (a_n) in Y represents a point in ∂Y (and so $d_Y(e, a_n) \to \infty$, since $d_Y(e, a_n) \geq (a_n \cdot a_m)^Y_e$), then $(f(a_n))$ represents a point in ∂X. Indeed, it suffices to show that a *subsequence* of $(f(a_n))$ represents a point in ∂X. By sequential compactness of $X \cup \partial X$ (Lemma 2.4), a subsequence of $(f(a_n))$ converges. If it converges to a point in X, then a subsequence of $(f(a_n))$ is in some compact (by properness of X) ball $\overline{B}_X(e, R)$. But then, by properness of f, a subsequence of (a_n) would be contained in some ball $B_Y(e, R')$, which would contradict $d_Y(e, a_n) \to \infty$. So some subsequence of $(f(a_n))$ converges to (that is, *represents*—see Lemma 2.4) a point in ∂X.

To establish continuity, suppose $p, q \in \partial Y$. By definition of the visual metrics $d_{\partial X}$ and $d_{\partial Y}$, there exist constants $r, s > 1$ and $k, l > 0$ (independent of p, q) such that

$$d_{\partial X}(\hat{f}(p), \hat{f}(q)) \leq k r^{-(\hat{f}(p) \cdot \hat{f}(q))^X_{f(e)}} \tag{3}$$

and

$$d_{\partial Y}(p, q) \geq l s^{-(p \cdot q)^Y_e}. \tag{4}$$

Since $(p \cdot q)_e^Y = \sup \liminf_{m,n \to \infty} (p_m \cdot q_n)_e^Y$, there exist sequences (p_m) and (q_n) in Y representing p and q, respectively, with $\liminf_{m,n \to \infty} (p_m \cdot q_n)_e^Y \geq (p \cdot q)_e^Y - 1$. So $(p_m \cdot q_n)_e^Y \geq (p,q)_e^Y - 2$ for all sufficiently large m, n. By definition of M we have $(f(p_m) \cdot f(q_n))_{f(e)}^X \geq M((p \cdot q)_e^Y - 2)$ for such m, n, and hence

$$(\hat{f}(p) \cdot \hat{f}(q))_{f(e)}^X \geq M((p \cdot q)_e^Y - 2). \tag{5}$$

Combining (3) and (5), we have

$$d_{\partial X}(\hat{f}(p), \hat{f}(q)) \leq k r^{-M((p \cdot q)_e^Y - 2)}. \tag{6}$$

So, by (4), if we make $d_{\partial Y}(p, q)$ sufficiently small, we can make $(p \cdot q)_e^Y$ arbitrarily large, so by hypothesis make $M((p \cdot q)_e^Y - 2)$ arbitrarily large, and so by (6) make $d_{\partial X}(\hat{f}(p), \hat{f}(q))$ arbitrarily small. Thus $\hat{f}$ is continuous.

That (c) $\Longrightarrow$ (b) is an immediate consequence of Lemma 2.4, for we just take $\hat{f}$ to be the restriction of $\overline{f}$. Properness of f guarantees that $\overline{f}(\partial Y) \subseteq \partial X$.

To see that (b) $\Longrightarrow$ (c) we show that the function $\overline{f} := f \cup \hat{f}$ is continuous. Since $\overline{X}$ and $\overline{Y}$ are metric spaces, it suffices to show that $p_n \to p$ implies $\overline{f}(p_n) \to \overline{f}(p)$ whenever $p_n, p \in \overline{Y}$. Since Y is an open subset of $\overline{Y}$ on which $\overline{f}$ restricts to the continuous function f, we may assume $p \in \partial Y$. Since $\hat{f}$ is continuous, we may assume each $p_n \in Y$. But then Lemma 2.4 says (p_n) *represents* p, so by (b), $(f(p_n))$ represents $\hat{f}(p)$. Using Lemma 2.4 again, we see $f(p_n) \to \hat{f}(p)$. That is, $\overline{f}(p_n) \to \overline{f}(p)$.

The equivalence of (d) and (e) comes from Lemma 2.1, which implies that there exists $C > 0$ such that $M'(N) \leq M(N + C) + C$ and $M(N) \leq M'(N + C) + C$ for all N.

That (e) $\Longrightarrow$ (f) is immediate as $[z, y]_Y$ is a geodesic segment in Y lying outside $B_Y(e, N)$.

Here is a proof that (f) $\Longrightarrow$ (e) under the assumption that $\sup\{d_X(f(x), f(y)) \mid d_Y(x, y) \leq r\} < \infty$ for all $r \geq 0$. Suppose $\lambda = [h_1, h_2]_Y$. As $t := (h_1 \cdot h_2)_e^Y$ approximates $d_Y(\lambda, e)$ with error at most a constant (Lemma 2.1), it is enough to show $d_X([f(h_1), f(h_2)]_X, e) \to \infty$ as $t \to \infty$. Let $\alpha_i = [e, h_i]_Y$ for $i = 1, 2$. By the slim-triangles condition, $[f(h_1), f(h_2)]_X$ lies in a C-neighborhood of a piecewise-geodesic path $[f(h_1), f(\alpha_1(t))]_X \cup [f(\alpha_1(t)), f(\alpha_2(t))]_X \cup [f(\alpha_2(t)), f(h_2)]_X$ for some constant C. So it is enough to show that the distance of each of these three segments from $f(e)$ in X tends to ∞ as $t \to \infty$. This is so for $[f(h_1), f(\alpha_1(t))]_X$ and $[f(\alpha_2(t)), f(h_2)]_X$ by (f). By the thin-triangles condition (see e.g. [8, pages 408–409]), $d_Y(\alpha_1(l), \alpha_2(l))$ is at most a constant for all $0 \leq l \leq t$, and so in particular $d_Y(\alpha_1(t), \alpha_2(t))$ is at most a constant. The assumption that $\sup\{d_X(f(x), f(y)) \mid d_Y(x, y) \leq r\} < \infty$ for all $r \geq 0$ gives an upper bound, independent of t, on the length of $[f(\alpha_1(t)), f(\alpha_2(t))]_X$. Since the distances of the endpoints of this segment from $f(e)$ in X tend to ∞ as $t \to \infty$, so does the distance of the whole segment. $\qquad\square$

The first part of the following lemma shows that Theorem 1.1 is not a quirk of the choice of generating sets. The second establishes the sense in which the function $\varepsilon(\delta)$ of Section 1 is an invariant for Cannon–Thurston maps. The third will allow us to reinterpret Lemma 3.1 (as Corollary 3.3) in a manner well suited to analyzing hyperbolic hydra.

Lemma 3.2. *Suppose Λ is a hyperbolic subgroup of a hyperbolic group Γ.*

(i). *Whether the Cannon–Thurston map $\partial \Lambda \to \partial \Gamma$ exists does not depend on the choice of finite generating sets giving the word metrics.*

(ii). *If the Cannon–Thurston map $\iota : \partial \Lambda \to \partial \Gamma$ exists for a hyperbolic subgroup Λ of a hyperbolic group Γ, the modulus of continuity for ι does not depend on the finite generating sets and the choices of visual metrics up to the following Hölder-type equivalence. If $\varepsilon(\delta)$ and $\varepsilon'(\delta)$ are the moduli of continuity of ι defined with respect to different such choices, then there are functions $f_1, f_2 : (0, \infty) \to (0, \infty)$, each of the form $x \mapsto C_i x^{\alpha_i}$ for some $C_i, \alpha_i > 0$, such that $\varepsilon'(\delta) \leq (f_1 \circ \varepsilon \circ f_2)(\delta)$ for all $\delta > 0$.*

(iii). *Suppose A and B are finite generating sets for Γ and Λ, respectively. Suppose $f : C_B(\Lambda) \to C_A(\Gamma)$ is any map between the respective Cayley graphs which restricts to the inclusion $\Lambda \hookrightarrow \Gamma$ on the vertices of $C_B(\Lambda)$ and sends edges to geodesics in $C_A(\Gamma)$. The Cannon–Thurston map $\partial \Lambda \to \partial \Gamma$, defined in terms of finite generating sets A for Γ and B for Λ, exists if and only if the Cannon–Thurston map $\partial C_B(\Lambda) \to \partial C_A(\Gamma)$ does.*

Proof. Suppose $\partial_1\Gamma$, $\partial_2\Gamma$, $\partial_1\Lambda$, and $\partial_2\Lambda$ are boundaries of Γ and Λ defined with respect to different finite generating sets and the Cannon–Thurston map $\partial_1\Lambda \to \partial_1\Gamma$ exists. The identity maps on Λ and Γ changing the word metrics are Lipschitz, so induce maps $\partial_2\Lambda \to \partial_1\Lambda$ and $\partial_1\Gamma \to \partial_2\Gamma$ (Theorem 3.9 of [8, page 430, III.H.3.7]). The composite map $\partial_2\Lambda \to \partial_1\Lambda \to \partial_1\Gamma \to \partial_2\Gamma$ satisfies Lemma 3.1(a), so the Cannon–Thurston map $\partial_2\Lambda \to \partial_2\Gamma$ exists.

If X, Y, Z are metric spaces and $g : Y \to Z$ and $h : X \to Y$ are maps with moduli of continuity ε_g and ε_h, respectively, then it follows from the definition that $\varepsilon_{g \circ h}(\delta) \leq (\varepsilon_g \circ \varepsilon_h)(\delta)$ for all $\delta > 0$. Specializing to the case of $\partial_2\Lambda \to \partial_1\Lambda \to \partial_1\Gamma \to \partial_2\Gamma$ in the previous paragraph, (ii) follows from the fact (Proposition 5.5 and Theorem 6.5 of [6]) that $\partial_2\Lambda \to \partial_1\Lambda$ and $\partial_1\Gamma \to \partial_2\Gamma$ are Hölder.

For (iii), a Cannon–Thurston map $\overline{C_B(\Lambda)} \to \overline{C_A(\Gamma)}$ restricts to a Cannon–Thurston map $\overline{\Lambda} \to \overline{\Gamma}$. Conversely, the quasi-isometries $\Gamma \hookrightarrow C_A(\Gamma)$ and $\Lambda \hookrightarrow C_B(\Lambda)$ induce isometries $\partial\Gamma \to \partial C_A(\Gamma)$ and $\partial C_B(\Lambda) \to \partial\Lambda$. So a Cannon–Thurston map $\partial\Lambda \to \partial\Gamma$ induces a composite map $\partial C_B(\Lambda) \to \partial\Lambda \to \partial\Gamma \to \partial C_A(\Gamma)$ satisfying Lemma 3.1(a). □

Corollary 3.3. *Suppose Λ is a finite-rank free subgroup of a hyperbolic group Γ. Suppose A is a finite generating set for Γ and B is a free basis for Λ. The Cannon–Thurston map $\partial\Lambda \to \partial\Gamma$ exists if and only if for all $M'' > 0$, there exists N such that whenever $\alpha\beta$ is a reduced word on B with $|\alpha| \geq N$, every geodesic in the Cayley graph $C_A(\Gamma)$ joining α to $\alpha\beta$ lies outside the ball of radius M'' about e.*

Proof. By Lemma 3.2(iii), the Cannon–Thurston map $\partial\Lambda \to \partial\Gamma$ exists if and only if the Cannon–Thurston map $\partial C_B(\Lambda) \to \partial C_A(\Gamma)$ for f (as defined in that lemma) does. Now applying condition (f) of Lemma 3.1 to f gives the result, since reduced words correspond to geodesics in $C_B(\Lambda)$. □

4 Cannon–Thurston maps for hyperbolic hydra groups

The *hyperbolic hydra* Γ_k are a family of hyperbolic groups with distorted (when $k > 1$) free subgroups Λ_k. In this section we will review some pertinent details from [7] and [13] of the construction and properties of the hyperbolic hydra groups Γ_k. We then show the existence of Cannon–Thurston maps $\partial\Lambda_k \to \partial\Gamma_k$ for hyperbolic hydra. Throughout, we fix an integer $k \geq 1$.

The hyperbolic hydra Γ_k of [7] is an elaboration of the *hydra group* G_k of [13]:

$$G_k = F(a_1, \ldots, a_k) \rtimes_\varphi \mathbb{Z}$$

where φ is the automorphism

$$\varphi(a_i) = \begin{cases} a_1 & i = 1, \\ a_i a_{i-1} & 1 < i \leq k \end{cases}$$

of the free group $F_k = F(a_1, \ldots, a_k)$. Let t denote a generator of the $\mathbb{Z}$-factor, so $\varphi(a_i) = t^{-1} a_i t$ in G_k for all i. In [13] it is proved that G_k is CAT(0) and has a rank-k free subgroup $H_k = \langle a_1 t, \ldots, a_k t \rangle$, distorted so that $\mathrm{Dist}_{H_k}^{G_k} \simeq A_k$.

Since the restriction of φ to $\langle a_1, \ldots, a_{i-1} \rangle = F_{i-1}$ is an automorphism for each i and $\varphi(a_i) \in a_i \langle a_1, \ldots, a_{i-1} \rangle$, we have

Lemma 4.1. *For any integer j (positive or negative),*

$$\varphi^j(a_i) \in a_i \langle a_1, \ldots, a_{i-1} \rangle \qquad and \qquad \varphi^j(a_i^{-1}) \in \langle a_1, \ldots, a_{i-1} \rangle a_i^{-1}.$$

For example, $\varphi^{-1}(a_7) = a_7 a_5 a_3 a_1 a_2^{-1} a_4^{-1} a_6^{-1}$.

The *normal form* of g in G_k is the unique $\widehat{w} t^m$ such that $\widehat{w}$ is a reduced word on $a_1, \ldots, a_k$ and $g = \widehat{w} t^m$ in G_k. For any $r \le k$, an H_r-*word* is a reduced word on $a_1 t, \ldots, a_r t$. For example, $(a_3 t)(a_3 t)(a_2 t)(a_3 t)^{-1}$ is an H_3-word and its normal form is $a_3 a_3 (a_1^{-1} a_2 a_1^2 a_2 a_1^{-1}) a_3^{-1} t^2$, since in G_3

$$
\begin{aligned}
(a_3 t)(a_3 t)(a_2 t)(a_3 t)^{-1} &= a_3(t a_3 t^{-1})(t^2 a_2 t^{-2})(t^2 a_3^{-1} t^{-2}) t^2 \\
&= a_3 \varphi^{-1}(a_3) \varphi^{-2}(a_2) \varphi^{-2}(a_3^{-1}) t^2 \\
&= a_3(a_3 a_1 a_2^{-1})(a_2 a_1^{-2})(a_2 a_1^2 a_2 a_1^{-1} a_3^{-1}) t^2 \\
&= a_3 a_3 (a_1^{-1} a_2 a_1^2 a_2 a_1^{-1}) a_3^{-1} t^2.
\end{aligned}
$$

Lemma 6.1 of [13], which says that $H_k \cap \langle t \rangle = \{1\}$ in G_k, implies the following two lemmas:

Lemma 4.2. *Given $g \in G_k$, if there exists j such that $g t^j \in H_k$, then that j is unique.*

Lemma 4.3. *If $g_1, g_2 \in H_k$ have normal forms $\widehat{w} t^{n_1}$ and $\widehat{w} t^{n_2}$, respectively, then $g_1 = g_2$.*

Observe that the relative locations of the $(a_3 t)^{\pm 1}$ in $(a_3 t)(a_3 t)(a_2 t)(a_3 t)^{-1}$ are the same as the relative locations of the $a_3^{\pm 1}$ in $a_3 a_3 (a_1^{-1} a_2 a_1^2 a_2 a_1^{-1}) a_3^{-1}$. That is, ignoring all other symbols, the former word has two $(a_3 t)$ symbols followed by $(a_3 t)^{-1}$, while the latter has two a_3 symbols followed by a_3^{-1}. This is an instance of the following lemma.

Lemma 4.4. *Consider an H_r-word w of the form*

$$u_0 (a_r t)^{\epsilon_1} u_1 (a_r t)^{\epsilon_2} u_2 \cdots (a_r t)^{\epsilon_n} u_n$$

where $u_0, \ldots, u_n \in H_{r-1}$ and $\epsilon_1, \ldots, \epsilon_n \in \{\pm 1\}$, and $u_i \ne 1$ whenever $\epsilon_i = -\epsilon_{i+1}$. Then, for all $s \in \mathbb{Z}$, the normal form $\widehat{w} t^m$ of w satisfies

$$\varphi^s(\widehat{w}) = v_0 a_r^{\epsilon_1} v_1 a_r^{\epsilon_2} v_2 \cdots a_r^{\epsilon_n} v_n$$

for some $v_0, \ldots, v_n \in \langle a_1, \ldots, a_{r-1} \rangle$ where $v_i \ne 1$ whenever $\epsilon_i = -\epsilon_{i+1}$.

The case $s = 0$ is Lemma 6.2 of [13]. We will not need to use Lemma 4.4; we state it because it sets the scene for the analogous Lemma 4.8 below. A proof of 4.4 can be extracted from the proof we will give for 4.8 by replacing $\Gamma_k, \Lambda_r, \Lambda_{r-1}, \varphi$ with $G_k, H_r, H_{r-1}, \theta$, respectively, and invoking Lemmas 4.1 and 4.3 instead of 4.5 and 4.7, respectively.

The construction of G_k above is elaborated in [7] to give the hyperbolic hydra Γ_k. It involves additional variables $a_0, b_1, \ldots, b_l$ and has the form

$$\Gamma_k = F \rtimes_\theta \mathbb{Z}$$

where F is the free group $F(a_0, \ldots, a_k, b_1, \ldots, b_l)$, and θ is an automorphism of F whose restriction to $F(b_1, \ldots, b_l)$ is an automorphism and

$$\theta(a_i) = \begin{cases} U a_1 V & i = 0, \\ a_0 & i = 1, \\ a_i a_{i-1} & 1 < i \le k, \end{cases}$$

where U and V are words on $b_1, \ldots, b_l$. We will prove here that Cannon–Thurston maps exist for all hyperbolic Γ_k of this form. In [7], U, V, l and $\theta|_{F(b_1, \ldots, b_l)}$ are chosen carefully to ensure Γ_k is hyperbolic. (In fact, in [7], $l = 17$, and U and V depend on k, but $\theta|_{F(b_1, \ldots, b_l)}$ does not.)

Let t denote a generator of the $\mathbb{Z}$-factor in $\Gamma_k = F \rtimes_\theta \mathbb{Z}$, so $t^{-1} a_i t = \theta(a_i)$ and $t^{-1} b_j t = \theta(b_j)$ for all i and j. For $1 \le r \le k$, let Λ_r be the subgroup $\langle a_0 t, \ldots, a_r t, b_1, \ldots, b_l \rangle$ of Γ_k. It is proved in [7] that Λ_k is free of rank $k + l + 1$ and is distorted so that $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k} \succeq A_k$. Let $\Lambda_0 = \langle b_1, \ldots, b_l \rangle$. (Note that Λ_r actually depends on both k and r, but k is fixed throughout this section.)

Understand $\theta^n(\widehat{w})$ to mean the reduced word on $a_0, \ldots, a_k, b_1, \ldots, b_l$ that represents $\theta^n(\widehat{w})$ in F.

Mapping $a_i \mapsto a_{\max\{1,i\}}$, $b_j \mapsto 1$ and $t \mapsto t$ for all i, j defines a surjection $\Gamma_k \twoheadrightarrow G_k$ such that $\Phi(\Lambda_r) = H_r$.

Corresponding to Lemmas 4.1–4.4 for G_k, we have the following Lemmas 4.5–4.8 for Γ_k.

The definition of θ immediately gives:

Lemma 4.5. *For any integer $j \in \mathbb{Z}$ (positive or negative), and any $i > 1$*

$$\theta^j(a_i) \in a_i \langle a_0, a_1, \ldots, a_{i-1}, b_1, \ldots, b_l \rangle \qquad and \qquad \theta^j(a_i^{-1}) \in \langle a_0, a_1, \ldots, a_{i-1}, b_1, \ldots, b_l \rangle a_i^{-1}.$$

The *normal form* of $g \in \Gamma_k$ is the unique $\widehat{g} t^m$ such that $\widehat{g}$ is a reduced word on $a_0, a_1, \ldots, a_k, b_1, \ldots, b_l$ and $g = \widehat{g} t^m$ in Γ_k. For any $1 \le r \le k$, a Λ_r-*word* is a reduced word on $a_0 t, a_1 t, \ldots, a_r t, b_1, \ldots, b_l$. Likewise, a Λ_0-word is a reduced word on $b_1, \ldots, b_l$.

Proposition 4.8 of [7] says that $\Lambda_k \cap \langle t \rangle = \{1\}$. So we immediately have the following analogues of Lemmas 4.2 and 4.3:

Lemma 4.6. *Given $g \in \Gamma_k$, if there exists j such that $g t^j \in \Lambda_k$, then that j is unique.*

Lemma 4.7. *If $g_1, g_2 \in \Lambda_k$ have normal forms $\widehat{w} t^{n_1}$ and $\widehat{w} t^{n_2}$, respectively, then $g_1 = g_2$.*

Finally, we have analogues of Lemma 4.4. We treat the $r > 1$ and $r = 1$ cases separately. First–

Lemma 4.8. *Let $r > 1$. Consider a Λ_r-word w of the form*

$$u_0 (a_r t)^{\epsilon_1} u_1 (a_r t)^{\epsilon_2} u_2 \cdots (a_r t)^{\epsilon_n} u_n$$

where $u_0, \ldots, u_n \in \Lambda_{r-1}$ and $\epsilon_1, \ldots, \epsilon_n \in \{\pm 1\}$, and $u_i \ne 1$ whenever $\epsilon_i = -\epsilon_{i+1}$. Then for any $s \in \mathbb{Z}$,

$$\theta^s(\widehat{w}) = v_0 a_r^{\epsilon_1} v_1 a_r^{\epsilon_2} v_2 \cdots a_r^{\epsilon_n} v_n$$

for some $v_0, \ldots, v_n \in \langle a_0, a_1, \ldots, a_{r-1}, b_1, \ldots, b_l \rangle$ where $v_i \ne 1$ whenever $\epsilon_i = -\epsilon_{i+1}$.

In short, Lemma 4.8 says that for a reduced Λ_r word w and its normal form $\widehat{w} t^n$, the occurrences of $a_r t$ in w correspond to occurrences of a_r in $\widehat{w}$ (and indeed in $\theta^s(\widehat{w})$ for any s) and the occurrences of $(a_r t)^{-1}$ in w correspond to occurrences of a_r^{-1} in $\widehat{w}$ (and indeed in $\theta^s(\widehat{w})$ for any s): the count of each and the order in which they occur in their respective words is preserved.

Proof of Lemma 4.8. Write each u_j in normal form: $u_j = \widehat{u_j} t^{m_j}$ with $\widehat{u_j} \in \langle a_0, a_1, \ldots, a_{r-1}, b_1, \ldots, b_l \rangle$. Then

$$\theta^s(\widehat{w}) = \theta^s(\widehat{u_0}) \theta^{p_1}(a_r^{\epsilon_1}) \theta^{q_1}(\widehat{u_1}) \theta^{p_2}(a_r^{\epsilon_2}) \theta^{q_2}(\widehat{u_2}) \cdots \theta^{p_n}(a_r^{\epsilon_n}) \theta^{q_n}(\widehat{u_n}) \tag{7}$$

where $q_i = s - (m_0 + m_1 + \cdots + m_{i-1}) - (\epsilon_1 + \cdots + \epsilon_i)$ and where

$$p_i = \begin{cases} s - (m_0 + m_1 + \cdots + m_{i-1}) - (\epsilon_1 + \cdots + \epsilon_{i-1}) & \text{if } \epsilon_i = 1 \\ s - (m_0 + m_1 + \cdots + m_{i-1}) - (\epsilon_1 + \cdots + \epsilon_{i-1}) + 1 & \text{if } \epsilon_i = -1. \end{cases}$$

By Lemma 4.5, the result follows unless $a_r^{\pm 1} a_r^{\mp 1}$ is cancelled during the reduction of the right side of equation (7). For an $a_r^{-1} a_r$ cancellation to occur in (7), it would have to occur within $\theta^{p_i}(a_r^{-1}) \theta^{q_i}(\widehat{u_i}) \theta^{p_{i+1}}(a_r)$

for some i. But then Lemma 4.5 would yield $\theta^{q_i}(\widehat{u_i}) = 1$, so $\widehat{u_i} = 1$. By Lemma 4.7, we would have $u_i = 1$, a contradiction. For an $a_r a_r^{-1}$ cancellation to occur in (7), we would have $\theta^{p_i}(a_r)\theta^{q_i}(\widehat{u_i})\theta^{p_{i+1}}(a_r^{-1}) = 1$ for some i such that $\epsilon_i = 1$ and $\epsilon_{i+1} = -1$. But then $q_i - p_i = -1$ and $p_{i+1} - p_i = -m_i$. So

$$
\begin{aligned}
1 &= \theta^{-p_i}(\theta^{p_i}(a_r)\theta^{q_i}(\widehat{u_i})\theta^{p_{i+1}}(a_r^{-1})) \\
&= a_r \theta^{q_i - p_i}(\widehat{u_i})\theta^{p_{i+1}-p_i}(a_r^{-1}) \\
&= a_r \theta^{-1}(\widehat{u_i})\theta^{-m_i}(a_r^{-1}) \\
&= a_r t\,\widehat{u_i}\,t^{-1} t^{m_i} a_r^{-1} t^{-m_i} \\
&= a_r t\, u_i (a_r t)^{-1} t^{-m_i}.
\end{aligned}
$$

So $t^{m_i} = (a_r t)u_i(a_r t)^{-1} \in \Lambda_k$. But then $t^{m_i} = 1$ by Lemma 4.7. This would contradict the fact that $u_i \in \Lambda_{r-1} \setminus \{1\}$. Thus no $a_r^{\pm 1} a_r^{\mp 1}$ cancellation occurs in the reduction of the right side of equation (7) and the lemma is proved. $\qquad\square$

Next we give the $r = 1$ analogue to Lemma 4.8. Recall that $\theta(a_0) = U a_1 V$ and $\theta(a_1) = a_0$, so Lemma 4.5 does not apply in the $r = 1$ case. Roughly speaking, Lemma 4.8 could be expanded to accommodate the $r = 1$ case by allowing occurrences of a_0 to sometimes swap with occurrences of a_1. More precisely:

Lemma 4.9. *Consider a Λ_1–word w of the form*

$$
w = u_0(a_{\mu_1} t)^{\epsilon_1} u_1 (a_{\mu_2} t)^{\epsilon_2} u_2 \cdots (a_{\mu_n} t)^{\epsilon_n} u_n
$$

where $u_0, \ldots, u_n \in \Lambda_0 = \langle b_1, \ldots, b_l \rangle$ with $\mu_1, \ldots, \mu_n \in \{0, 1\}$ and $\epsilon_1, \ldots, \epsilon_n \in \{\pm 1\}$, and with $u_i \neq 1$ whenever $\mu_i = \mu_{i+1}$ and $\epsilon_i = -\epsilon_{i+1}$. Then for all $s \in \mathbb{Z}$,

$$
\theta^s(\widehat{w}) = v_0 a_{\xi_1}^{\epsilon_1} v_1 a_{\xi_2}^{\epsilon_2} v_2 \cdots a_{\xi_n}^{\epsilon_n} v_n
$$

for some $v_0, \ldots, v_n \in \langle b_1, \ldots, b_l \rangle$ and some $\xi_1, \ldots, \xi_n \in \{0, 1\}$ with $v_i \neq 1$ whenever $\xi_i = \xi_{i+1}$ and $\epsilon_i = -\epsilon_{i+1}$.

Proof. We have
$$
\theta^s(\widehat{w}) = \theta^s(u_0)\theta^{p_1}(a_{\mu_1}^{\epsilon_1})\theta^{q_1}(u_1)\theta^{p_2}(a_{\mu_2}^{\epsilon_2})\theta^{q_2}(u_2)\cdots\theta^{p_n}(a_{\mu_n}^{\epsilon_n})\theta^{q_n}(u_n) \tag{8}
$$
where $q_i = s - (\epsilon_1 + \cdots + \epsilon_i)$ and where

$$
p_i = \begin{cases} s - (\epsilon_1 + \cdots + \epsilon_{i-1}) & \text{if } \epsilon_i = +1 \\ s - (\epsilon_1 + \cdots + \epsilon_{i-1}) + 1 & \text{if } \epsilon_i = -1. \end{cases}
$$

Each $\theta^{q_i}(u_i) \in \langle b_1, \ldots, b_l \rangle$ and $\theta^{p_i}(a_{\mu_i}^{\epsilon_i}) \in \langle b_1, \ldots, b_l \rangle a_{\xi_i}^{\epsilon_i} \langle b_1, \ldots, b_l \rangle$ where $\xi_i = \mu_i + p_i \pmod 2$. The result follows unless some $a_0^{\pm 1} a_0^{\mp 1}$ or $a_1^{\pm 1} a_1^{\mp 1}$ cancels in the reduction of the right side of equation (8). Such a cancellation would have to occur within $\theta^{p_i}(a_{\mu_i}^{\epsilon_i})\theta^{q_i}(u_i)\theta^{p_{i+1}}(a_{\mu_{i+1}}^{\epsilon_{i+1}})$ for some i. But then

$$
\theta^{p_i}(a_{\mu_i}^{\epsilon_i})\theta^{q_i}(u_i)\theta^{p_{i+1}}(a_{\mu_{i+1}}^{\epsilon_{i+1}}) \in \langle b_1, \ldots, b_l \rangle.
$$

Applying θ^{-p_i}, we get

$$
a_{\mu_i}^{\epsilon_i}\theta^{q_i - p_i}(u_i)\theta^{p_{i+1}-p_i}(a_{\mu_{i+1}}^{\epsilon_{i+1}}) \in \langle b_1, \ldots, b_l \rangle.
$$

For $a_0^{\pm 1} a_0^{\mp 1}$ or $a_1^{\pm 1} a_1^{\mp 1}$ to cancel, we would have to have $\epsilon_{i+1} = -\epsilon_i$. But then $p_{i+1} = p_i$ (check the two cases $\epsilon_i = \pm 1$), and so

$$
a_{\mu_i}^{\epsilon_i}\theta^{q_i - p_i}(u_i)a_{\mu_{i+1}}^{-\epsilon_i} \in \langle b_1, \ldots, b_l \rangle.
$$

But then $\theta^{q_i - p_i}(u_i) = 1$ and $\mu_i = \mu_{i+1}$, because $\langle a_0, a_1, b_1, \ldots, b_l \rangle$ is free on the given generators. This would contradict the fact that $u_i \neq 1$. Thus no such cancellation occurs, and the lemma is proved. $\qquad\square$

Next we give a technical lemma comparing the location of the final $(a_r t)^{\pm 1}$ in a Λ_r-word w to the location of the corresponding $a_r^{\pm 1}$ in its normal form $\widehat{w} t^m$, and moreover in $\theta^n(\widehat{w})$. A point of terminology: the 'R' in the word 'WORD' *occurs 3 letters in*.

We use $|u|_F$ to denote the length of a word u on $a_0, \ldots, a_r, b_1, \ldots, b_l$. And we use $|u|_{\Lambda_r}$ to denote the length of a Λ_r-word—that is, length as a word on $(a_0 t), \ldots, (a_r t), b_1, \ldots, b_l$, not as a word on $a_0, \ldots, a_r, b_1, \ldots, b_l$.

Lemma 4.10. *For all integers $A, B \geq 0$ and r with $k \geq r > 1$, there exists N such that if $|n| \leq B$ and $w = u(a_r t)^{\pm 1}$ is a reduced Λ_r-word with $|u|_{\Lambda_r} \geq N$, then the final $a_r^{\pm 1}$ in $\theta^n(\widehat{w})$ occurs at least A symbols in.*

Similarly, such an N exists for the final $a_0^{\pm 1}$ or $a_1^{\pm 1}$ in $\theta^n(\widehat{w})$ when $r = 1$ and w is $u(a_0 t)^{\pm 1}$ or $u(a_1 t)^{\pm 1}$.

Proof. We may suppose n is a fixed integer such that $|n| \leq B$.

Assume first that $r > 1$. In the manner explained in the comment following Lemma 4.8, the $(a_r t)^{\pm 1}$ in w correspond to the $a_r^{\pm 1}$ in $\theta^n(\widehat{w})$.

First we will address the case where $w = u(a_r t)^{-1}$. Since the final letter of $\theta^{\pm 1}(a_r^{-1})$ is a_r^{-1}, the final a_r^{-1} in $\theta^n(\widehat{w})$ is in fact the final letter of $\theta^n(\widehat{w})$. So we are seeking to prove that $|\theta^n(\widehat{w})|_F \geq A$ when N is sufficiently large. Since θ is an automorphism, only finitely many $v \in F$ satisfy $|\theta^n(v)|_F < A$. By Lemma 4.7, each such v equals $\widehat{w}$ for at most one Λ_r-word w, so the lemma is proved by taking N sufficiently large to avoid these finitely many w.

Next we address the case where $w = u(a_r t)$. The normal forms of u and w are $\widehat{w} t^{j+1}$ and $\widehat{u} t^j$ for some j. They are related in that

$$\widehat{w} t^{j+1} \;=\; w \;=\; u(a_r t) \;=\; \widehat{u} t^j a_r t \;=\; \widehat{u}\, \theta^{-j}(a_r)\, t^{j+1}.$$

Thus $\widehat{w} = \widehat{u}\, \theta^{-j}(a_r)$.

Since the first letter of $\theta^{n-j}(a_r)$ is a_r, Lemma 4.8 says there is no cancellation between $\theta^n(\widehat{u})$ and $\theta^{n-j}(a_r)$ in $\theta^n(\widehat{w}) = \theta^n(\widehat{u}) \cdot \theta^{n-j}(a_r)$. (Otherwise there would be too few instances of a_r in $\theta^n(\widehat{w})$.) Thus the final a_r in $\theta^n(\widehat{w})$ occurs $|\theta^n(\widehat{u})|_F + 1$ letters in. So we are now seeking to prove that $|\theta^n(\widehat{u})|_F + 1 \geq A$ when N is sufficiently large, and this can be handled as in the previous case. This completes the proof when $r > 1$.

Next, we do the case $r = 1$. Suppose $w = u(a_\mu t)^{\pm 1}$ is reduced with $u = \widehat{u} t^j$ and $\mu \in \{0, 1\}$. This time we have $\widehat{w} = \widehat{u}\, \theta^{-J}(a_\mu^{\pm 1})$ where $|J - j| \leq 1$. This time there may be some cancellation between $\theta^n(\widehat{u})$ and $\theta^{n-J}(a_\mu^{\pm 1})$ in $\theta^n(\widehat{w}) = \theta^n(\widehat{u}) \cdot \theta^{n-J}(a_\mu^{\pm 1})$, but the cancellation is restricted (by Lemma 4.9) to only symbols from $\{b_1^{\pm 1}, \ldots, b_l^{\pm 1}\}$. We have two cases: $|J| \geq A$ or $|J| < A$. If $|J| \geq A$, then $|j| \geq A - 1$, so Lemma 4.9 tells us that $\theta^n(\widehat{u})$ contains at least $A - 1$ symbols from among $\{a_0^{\pm 1}, a_1^{\pm 1}\}$. Thus the final $a_0^{\pm 1}$ or $a_1^{\pm 1}$ of $\theta^n(\widehat{w})$ occurs at least A symbols in, as desired. On the other hand, if $|J| < A$, then $|n - J| < n + A$ by the triangle inequality. This gives a bound on $|\theta^{n-J}(a_\mu^{\pm 1})|_F$ and hence on the amount of cancellation that can occur in $\theta^n(\widehat{w}) = \theta^n(\widehat{u}) \cdot \theta^{n-J}(a_\mu^{\pm 1})$. The lemma now follows by the same argument as in the $r > 1$ case. $\square$

Let $C(F)$ and $C(\Gamma_k)$ denote the Cayley graphs of F and Γ_k with respect to $a_0, \ldots, a_k, b_1, \ldots, b_l$ and $a_0, \ldots, a_k, b_1, \ldots, b_l, t$, respectively. Let $B_F(e, R)$ denote the open ball of radius R about e in $C(F)$. Write $[x, y]_F$ or $[x, y]_{\Gamma_k}$ for a geodesic between x and y in $C(F)$ or $C(\Gamma_k)$, respectively. In the case of $C(F)$, which is a tree, geodesics between any given pair of points are unique. Let d_F and d_{Γ_k} be the associated metrics.

The *shadow* of the suffix β of a reduced Λ_k-word $\alpha\beta$ is the set of all geodesic segments $[\widehat{\alpha \cdot \beta(i)}, \widehat{\alpha \cdot \beta(i+1)}]_F$ where $\beta(i)$ denotes the length-i prefix of β and $0 \leq i < |\beta|_{\Lambda_k}$.

Lemma 4.11. *For all $K > 0$, there exist integers $C, R > 0$ such that if $\alpha\beta$ is a reduced Λ_k-word and $|\alpha|_{\Lambda_k} \geq C$ and the shadow of β is outside $B_F(e, R)$, then every $[\alpha, \alpha\beta]_{\Gamma_k}$ satisfies $d_{\Gamma_k}([\alpha, \alpha\beta]_{\Gamma_k}, e) \geq K$.*

Proof. As Γ_k is hyperbolic, there is some $\delta > 0$ such that every geodesic triangle in $C(\Gamma_k)$ is δ-slim.

Any geodesic segment $[\alpha, \alpha\beta]_{\Gamma_k}$ is in a 2δ-neighborhood of any piecewise-geodesic path $[\alpha, \widehat{\alpha}]_{\Gamma_k} \cup [\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k} \cup [\widehat{\alpha\beta}, \alpha\beta]_{\Gamma_k}$ in $C(\Gamma_k)$. So it suffices to have $[\alpha, \widehat{\alpha}]_{\Gamma_k}$, $[\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k}$, and $[\widehat{\alpha\beta}, \alpha\beta]_{\Gamma_k}$ stay at least $K + 2\delta$ away from the identity element e.

To ensure $d(e, [\alpha, \widehat{\alpha}]_{\Gamma_k}) \geq K + 2\delta$, we need $[\alpha, \widehat{\alpha}]_{\Gamma_k}$ to avoid finitely many elements of Γ_k, say $g_1, \ldots, g_m$. Since $\alpha = \widehat{\alpha} t^n$ for some n, there is a unique geodesic $[\alpha, \widehat{\alpha}]$ joining α to $\widehat{\alpha}$ in $C(\Gamma_k)$ and it is a succession of edges all labelled t. So if g_i is on $[\alpha, \widehat{\alpha}]$ then $g_i t^{j_i} = \alpha \in \Lambda_k$ for some j_i. But then by Lemma 4.6, it suffices for α not to be one of at most m elements of Λ_k. So it suffices to ensure $|\alpha|_{\Lambda_k}$ is sufficiently long.

Since $|\alpha\beta|_{\Lambda_k} \geq |\alpha|_{\Lambda_k}$, we find $d([\widehat{\alpha\beta}, \alpha\beta]_{\Gamma_k}, e) \geq K + 2\delta$ also.

Finally, we consider $[\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k}$. The strategy is to use the existence of the Cannon–Thurston map $\partial F \to \partial\Gamma_k$ (not $\partial\Lambda_k \to \partial\Gamma_k$!) to ensure this geodesic stays far (at least $K + 2\delta$ from the identity. As previously mentioned, the main theorem in [23] is that Cannon–Thurston maps always exist for infinite hyperbolic normal subgroups of hyperbolic groups. Thus condition (e) of Lemma 3.1 must hold where $X = C(\Gamma_k)$ and $Y = C(F)$ and f is the inclusion map. Since $C(F)$ is a tree, the geodesic segment $[\widehat{\alpha}, \widehat{\alpha\beta}]_F$ is a subset of the union of the geodesic segments comprising the shadow of β, which do not intersect $B_F(e, R)$ by assumption. Thus $[\widehat{\alpha}, \widehat{\alpha\beta}]_F$ is disjoint from $B_F(e, R)$. So condition (e) says that choosing R large enough makes $[\widehat{\alpha}, \widehat{\alpha\beta}]_{\Gamma_k}$ arbitrarily far (so at least $K + 2\delta$) from e, as desired. $\qquad\square$

We are now ready to use Corollary 3.3 to show the Cannon–Thurston map $\partial\Lambda_k \to \partial\Gamma_k$ exists.

Proof of Theorem 1.1. Fix some $\delta > 0$ so that all geodesic triangles in $C(\Gamma_k)$ are δ-slim.

For integers $A, B \geq 0$ and r with $k \geq r \geq 1$, let $N(r, A, B)$ be the least integer N as per Lemma 4.10. Given an integer $R > 0$, recursively define a sequence $N_k(R), N_{k-1}(R), \ldots, N_1(R)$ of positive integers by:

$$N_k(R) := N(k, R, 0),$$

and for $r = k-1, k-2, \ldots, 1$

$$N_r(R) := N\left(r, R + \max\left\{|\widehat{w}|_F \,\middle|\, \Lambda_k\text{-words } w \text{ with } |w|_{\Lambda_k} = \sum_{j=r+1}^{k} N_j(R)\right\}, \sum_{j=r+1}^{k} N_j(R)\right).$$

Suppose $M'' > 0$ is given. Let $K = M'' + (2\delta + 1)k$. Let $C, R > 0$ be obtained from K as per Lemma 4.11. Recall the map Φ from the hyperbolic hydra group Γ_k to the hydra group G_k defined by $a_i \mapsto a_{\max\{1, i\}}$, $b_j \mapsto 1$ and $t \mapsto t$ for all i, j. Let L be the maximum of $|\theta^{-n}(s)|_F$ ranging over all $s \in \{b_1, \ldots, b_l\}$ and all $n \in \mathbb{Z}$ for which there exists $\widehat{u} \in F$ such that $\widehat{u} t^n \in \Lambda_k$ and the reduced word representing $\Phi(\widehat{u})$ in $F(a_1, \ldots, a_k)$ has length less than R. (There are only finitely many such n since $\widehat{u} t^n \in \Lambda_k$ implies $\Phi(\widehat{u} t^n) = \Phi(\widehat{u}) t^n \in H_k$, and for any $x \in F(a_1, \ldots, a_k)$, there is at most one $m \in \mathbb{Z}$ such that $x t^m \in H_k$ by Lemma 4.2.) Choose $C' > C$ so that every Λ_k-word u of length $|u|_{\Lambda_k} \geq C'$ has free-by-cyclic normal form $\widehat{u} t^n$ with $|\widehat{u}|_F \geq R + (L/2)$.

Suppose $\alpha\beta$ is a reduced Λ_k-word such that $|\alpha|_{\Lambda_k} \geq N := C' + \sum_{r=1}^{k} N_r(R)$. Express α as $w_k \cdots w_0$ where $|w_r|_{\Lambda_k} = N_r(R)$ for $1 \leq r \leq k$. In particular, $|\alpha|_{\Lambda_k} \geq |w_0|_{\Lambda_k} \geq C' > C$. Let β_r denote the longest prefix of β in Λ_r and let γ_r denote (any) geodesic $[\alpha, \alpha\beta_r]_{\Gamma_k}$. In particular, γ_k is an arbitrary geodesic $[\alpha, \alpha\beta]_{\Gamma_k}$. We will show that γ_k lies at least a distance M'' from e in $C(\Gamma_k)$. Corollary 3.3 will then complete the proof.

Suppose, for a contradiction, that $d_{\Gamma_k}(\gamma_k, e) < M''$. Let $\widehat{\alpha} t^n$ be the free-by-cyclic normal form of α in Γ_k.

We claim that the shadow of the suffix β_0 of $\alpha\beta_0$ does not intersect $B_F(e, R)$. The endpoints of the geodesic segments in F comprising this shadow are all of the form $\widehat{\alpha}\theta^{-n}(x)$ for various $x \in F(b_1, \ldots, b_l)$. There are two cases to consider: the length of the reduced word in $F(a_1, \ldots, a_k)$ representing $\Phi(\widehat{\alpha})$ is at least R and is less than R. In the former case, because $\widehat{\alpha}$ contains at least R letters $a_i^{\pm 1}$ $(0 \leq i \leq k)$, the closest approach of any such geodesic to e (i.e. the Gromov product of its endpoints) is at least R. In the latter case, L is an

upper bound for the length of the constituent geodesics in the shadow of β_0, and so, by definition of C', the shadow of β_0 does not intersect $B_F(e, R)$. In either case, the shadow stays outside of $B_F(e, R)$.

On the other hand, the following claim, in the case $r = 0$, shows that

$$d_{\Gamma_k}([\alpha, \alpha\beta_0]_{\Gamma_k}, e) < M'' + (2\delta + 1)k = K,$$

so Lemma 4.11 implies the shadow of β_0 intersects $B_F(e, R)$. This contradiction will prove the theorem.

Claim. For $r = k, k-1, \ldots, 1, 0$,

  (i_r). $d_{\Gamma_k}(\gamma_r, e) < M'' + (2\delta + 1)(k - r)$, and

(ii_r). $w_r w_{r-1} \cdots w_0 \in \Lambda_r$.

We prove this claim using downward induction on r.

The base case $r = k$ is straightforward: $\gamma_k = [\alpha, \alpha\beta]_{\Gamma_k}$ and $w_k \cdots w_0 = \alpha \in \Lambda_k$ by definition, and $d_{\Gamma_k}(\gamma_k, e) < M''$ by hypothesis.

Now we prove that (i_{r+1}) and (ii_{r+1}) implies (i_r) and (ii_r) for $r = k-1, \ldots, 1, 0$. (In the case $r = 0$, we must interpret each instance of "$(a_{r+1}t)^{\pm 1}$" in the following to mean "$(a_0 t)^{\pm 1}$ or $(a_1 t)^{\pm 1}$" and each occurrence of "$a_{r+1}^{\pm 1}$" to mean "$a_0^{\pm 1}$ or $a_1^{\pm 1}$".)

We will make repeated use of the following lemma.

Lemma 4.12. *Suppose $w_k, \ldots, w_{r+1}$ are as defined earlier, and $w_k \cdots w_{r+1} x(a_{r+1}t)^{\pm 1} y z$ is a reduced Λ_k-word in which the subwords x, y and z are Λ_{r+1}-words and $|w_k \cdots w_{r+1} x(a_{r+1}t)^{\pm 1} y|_{\Lambda_k} \geq C$. If γ is any geodesic in $C(\Gamma_k)$ from $w_k \cdots w_{r+1} x(a_{r+1}t)^{\pm 1} y$ to $w_k \cdots w_{r+1} x(a_{r+1}t)^{\pm 1} y z$, then $d_{\Gamma_k}(\gamma, e) \geq K$.*

To prove this lemma we consider the free-by-cyclic normal form $\widehat{w_k \cdots w_{r+2}} t^{n_r}$ for $w_k \cdots w_{r+2}$ in Γ_k. Now, $|w_k \cdots w_{r+2}|_{\Lambda_k} = \sum_{j=r+2}^{k} N_j(R)$ by construction, so $|n_r| \leq \sum_{j=r+2}^{k} N_j(R)$. So, by definition,

$$N_{r+1}(R) \geq N(r+1, R + |\widehat{w_k \cdots w_{r+2}}|_F, |n_r|).$$

By hypothesis (ii_{r+1}), we see $w_{r+1} x(a_{r+1}t)^{\pm 1} \in \Lambda_{r+1}$, and $|w_{r+1}| = N_{r+1}(R)$ by construction, so the final $a_{r+1}^{\pm 1}$ in $\theta^{-n_r}(\widehat{w_{r+1} x(a_{r+1}t)}^{\pm 1})$ is at least $R + |\widehat{w_k \cdots w_{r+2}}|_F$ letters in. Therefore, the final $a_{r+1}^{\pm 1}$ in $\widehat{w_k \cdots w_{r+2}} \, \theta^{-n_r}(\widehat{w_{r+1} x(a_{r+1}t)}^{\pm 1})$ is at least $(R + |\widehat{w_k \cdots w_{r+2}}|_F) - |\widehat{w_k \cdots w_{r+2}}|_F = R$ letters in. Since y and z are Λ_{r+1}-words and $w_k \cdots w_{r+1} x(a_{r+1}t)^{\pm 1} y z$ is reduced, this implies the shadow of z cannot intersect $B_F(e, R)$. So, by Lemma 4.11, $d_{\Gamma_k}(\gamma, e) \geq K$, completing the proof of the lemma.

Returning to the proof of the claim, we will consider two cases: $\beta_r = \beta_{r+1}$ and $\beta_r \neq \beta_{r+1}$. In the former case, we may assume $\gamma_r = \gamma_{r+1}$. So (i_r) follows immediately from (i_{r+1}). Since (i_r) implies $d_{\Gamma_k}(\gamma_r, e) < K$, (ii_{r+1}) and Lemma 4.12 with $z = \beta_r$ shows that $w_r \cdots w_0$ cannot be expressed as $x(a_{r+1}t)^{\pm 1} y$. So $w_r \cdots w_0 \in \Lambda_r$ and we have (ii_r).

Next, assume $\beta_r \neq \beta_{r+1}$. The (reduced) word β_{r+1} can be expressed as $\beta_r(a_{r+1}t)^{\pm 1} \beta_r'$ for some $\beta_r' \in \Lambda_{r+1}$. Let ρ_r be the geodesic segment in $C(\Gamma_k)$ labelled $(a_{r+1}t)^{\pm 1}$ connecting $\alpha\beta_r$ to $\alpha\beta_r(a_{r+1}t)^{\pm 1}$. Let $\gamma_r' = [\alpha\beta_r(a_{r+1}t)^{\pm 1}, \alpha\beta_{r+1}]_{\Gamma_k}$. Then, by Lemma 4.12 with $x = w_r \cdots w_0 \beta_r$, y the empty word, and $z = \beta_r'$,

$$d_{\Gamma_k}(\gamma_r', e) \geq K. \tag{9}$$

If $(a_{r+1}t)^{\pm 1}$ occurs in $w_r \cdots w_0$, then Lemma 4.12 with $x(a_{r+1}t)^{\pm 1} y = w_r \cdots w_0$ and $z = \beta_r$ shows that

$$d_{\Gamma_k}(\gamma_r, e) \geq K. \tag{10}$$

By the slim-triangles condition for $C(\Gamma_k)$, γ_{r+1} is contained in the 2δ-neighborhood of $\gamma_r \cup \rho_r \cup \gamma'_r$ and hence in the $(2\delta+1)$-neighborhood of $\gamma_r \cup \gamma'_r$. So

$$
\begin{aligned}
\min\{d(e,\gamma_r),d(e,\gamma'_r)\} &\leq d(e,\gamma_{r+1})+(2\delta+1) \\
&< M'' + (2\delta+1)(k-(r+1))+(2\delta+1) \\
&= M'' + (2\delta+1)(k-r),
\end{aligned}
$$

the second inequality coming from (i_{r+1}).

But by (9), $d(e,\gamma'_r) \geq K \geq M'' + (2\delta+1)(k-r)$. So $\min\{d(e,\gamma_r),d(e,\gamma'_r)\} = d(e,\gamma_r)$ and (i_r) follows.

Moreover, (10) cannot be true since it contradicts (i_r), so $(a_{r+1}t)^{\pm 1}$ does not occur in $w_r \cdots w_0$ and (ii_r) follows. This completes the induction step of the claim, and thus proves the theorem by contradiction. □

5   Wildness of Cannon–Thurston maps

Proof of Theorem 1.2. We have that Γ and Λ are (δ_Γ)- and (δ_Λ)-hyperbolic, respectively, for some $\delta_\Gamma, \delta_\Lambda > 0$. Let $\iota : \Lambda \to \Gamma$ denote the inclusion map and $\hat{\iota} : \partial\Lambda \to \partial\Gamma$ denote the Cannon–Thurston map.

Since Λ is non-elementary, $|\hat{\iota}(\partial\Lambda)| = \infty$ by [16, Thm. 12.2(1)]. We may thus choose $p_1, p_2, p_3 \in \partial\Lambda$ with $\hat{\iota}p_1, \hat{\iota}p_2, \hat{\iota}p_3 \in \partial\Gamma$ distinct. Let $C = 2\delta_\Lambda + \max\{(p_i \cdot p_j)_e^\Lambda \mid 1 \leq i < j \leq 3\}$.

By definition of the distortion function (see Section 1), we can take a sequence $h_n \in \Lambda$ with $d_\Gamma(e,h_n) \leq n$ and $d_\Lambda(e,h_n) = \mathrm{Dist}_\Lambda^\Gamma(n)$. By Lemma 2.2(1)

$$
(p_i \cdot p_j)_e^\Lambda = (h_n p_i \cdot h_n p_j)_{h_n}^\Lambda \geq \min\{(e \cdot h_n p_i)_{h_n}^\Lambda, (e \cdot h_n p_j)_{h_n}^\Lambda\} - 2\delta_\Lambda,
$$

so we can choose $i = i(n)$ and $j = j(n) \in \{1,2,3\}$ with $i \neq j$ such that

$$
(e \cdot h_n p_i)_{h_n}^\Lambda, (e \cdot h_n p_j)_{h_n}^\Lambda \leq C.
$$

Combined with Lemma 2.2(2) this gives that for $k = i, j$,

$$
(h_n \cdot h_n p_k)_e^\Lambda \geq d_\Lambda(e,h_n) - (e \cdot h_n p_k)_{h_n}^\Lambda - \delta_\Lambda \geq \mathrm{Dist}_\Lambda^\Gamma(n) - C - \delta_\Lambda.
$$

So, by Lemma 2.2(1),

$$
(h_n p_i \cdot h_n p_j)_e^\Lambda \geq \min\{(h_n \cdot h_n p_i)_e^\Lambda, (h_n \cdot h_n p_j)_e^\Lambda\} - 2\delta_\Lambda \geq \mathrm{Dist}_\Lambda^\Gamma(n) - C - 3\delta_\Lambda.
$$

Writing $\beta := k_2 r^{C+3\delta_\Lambda}$, where k_2 is as per (2) in Section 2 applied to $\partial\Lambda$, we get

$$
d_{\partial\Lambda}(h_n p_i, h_n p_j) \leq k_2 r^{-(h_n p_i \cdot h_n p_j)_e^\Lambda} \leq \frac{\beta}{r^{\mathrm{Dist}_\Lambda^\Gamma(n)}}. \tag{11}
$$

On the other hand, using Lemma 2.2(3) for the first inequality,

$$
\begin{aligned}
(\hat{\iota}(h_n p_i) \cdot \hat{\iota}(h_n p_j))_e^\Gamma &= (h_n \hat{\iota} p_i \cdot h_n \hat{\iota} p_j)_e^\Gamma \\
&\leq d_\Gamma(e, [h_n \hat{\iota} p_i, h_n \hat{\iota} p_j]_\Gamma) + 8\delta_\Gamma \\
&\leq d_\Gamma(e, h_n) + d_\Gamma(h_n, [h_n \hat{\iota} p_i, h_n \hat{\iota} p_j]_\Gamma) + 8\delta_\Gamma \\
&\leq n + d_\Gamma(e, [\hat{\iota} p_i, \hat{\iota} p_j]_\Gamma) + 8\delta_\Gamma.
\end{aligned}
$$

Writing $\alpha := k_1 / s^{8\delta_\Gamma + \max\{d_\Gamma(e,[\hat{\iota}p_i,\hat{\iota}p_j]_\Gamma) \mid 1 \leq i < j \leq 3\}}$, where k_1 is as per (2) in Section 2 applied to $\partial\Gamma$, we then get

$$
d_{\partial\Gamma}(\hat{\iota}(h_n p_i), \hat{\iota}(h_n p_j)) \geq k_1 s^{-(\hat{\iota}(h_n p_i), \hat{\iota}(h_n p_j))_e^\Gamma} \geq \frac{\alpha}{s^n}. \tag{12}
$$

Combining (11) and (12) yields the inequality claimed in Theorem 1.2. □

Corollary 1.3 will follow from the following proposition.

Proposition 5.1. *For all $k \geq 2$, the modulus of continuity $\varepsilon(\delta)$ of the Cannon–Thurston map $\partial \Lambda_k \to \partial \Gamma_k$ for hyperbolic hydra has the property that there exist $C_0, C_1 > 0$ and $C_2 > 1$ such that for all $\eta \in (0, C_0)$,*

$$\varepsilon \left(\frac{1}{C_2^{A_k\left(\lfloor C_1 \log(1/\eta) \rfloor\right)}} \right) \geq \eta. \tag{13}$$

Proof. For convenience, extend the domains of the functions $A_k : \mathbb{N} \to \mathbb{N}$ and $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k} : \mathbb{N} \to \mathbb{N}$ to $[1, \infty)$ by declaring the functions to be constant on the half-open intervals $[n, n+1)$.

From Theorem 1.2 we have

$$\varepsilon \left(\frac{\beta}{r^{\mathrm{Dist}_{\Lambda_k}^{\Gamma_k}(n)}} \right) \geq \frac{\alpha}{s^n}$$

for all real $n \geq 0$ and some constants $\alpha, \beta > 0$ and $r, s > 1$. Thus for all $\eta \leq \alpha$,

$$\varepsilon \left(\frac{1}{\exp \left(\log(r) \mathrm{Dist}_{\Lambda_k}^{\Gamma_k}(\log_s(\alpha/\eta)) - \log(\beta) \right)} \right) \geq \eta. \tag{14}$$

As $\mathrm{Dist}_{\Lambda_k}^{\Gamma_k} \succeq A_k$ by [7], there exists $C > 0$ such that $A_k(n) \leq C \mathrm{Dist}_{\Lambda_k}^{\Gamma_k}(Cn + C) + Cn + C$ for all real n, and therefore

$$\mathrm{Dist}_{\Lambda_k}^{\Gamma_k}(N) \geq \frac{1}{C} A_k \left(\frac{N - C}{C} \right) - \frac{N}{C}$$

for $N \geq 2C$. So

$$\log(r) \mathrm{Dist}_{\Lambda_k}^{\Gamma_k}(\log_s(\alpha/\eta)) - \log(\beta) \geq K_1 A_k(K_2 \log(1/\eta) - K_3) - K_4 \log(1/\eta) - K_5$$

for all $\eta \in (0, K_0)$, for suitable constants $K_0, \ldots, K_5 > 0$. By shrinking K_0 if necessary, we can make $\log(1/\eta)$ arbitrarily large, so that we may absorb the constant K_3 into K_2. Moreover, A_k grows faster than a linear function as $k \geq 2$, so (by further shrinking K_0) the constants K_4 and K_5 can be absorbed into K_1. Thus

$$\log(r) \mathrm{Dist}_{\Lambda_k}^{\Gamma_k}(\log_s(\alpha/\eta)) - \log(\beta) \geq K_1 A_k(K_2 \log(1/\eta)). \tag{15}$$

So combining (14) and (15) and setting $C_0 = K_0$, $C_1 = K_2$, and $C_2 = e^{K_1}$, we have the result claimed. $\square$

Plugging $\eta = 1/n$ into inequality (13) yields

$$\varepsilon \left(\frac{1}{C_2^{A_k\left(\lfloor C_1 \log(n) \rfloor\right)}} \right) \geq \frac{1}{n}$$

for all sufficiently large n. So Corollary 1.3, which asserts

$$\varepsilon \left(\frac{1}{A_{k-1}(n)} \right) \geq \frac{1}{n}$$

for all sufficiently large n, follows from the fact that

$$C_2^{A_k\left(\lfloor C_1 \log(n) \rfloor\right)} \gg A_k \left(\lfloor C_1 \log(n) \rfloor \right) = A_{k-1} \left(A_k \left(\lfloor C_1 \log(n) \rfloor - 1 \right) \right) \geq A_{k-1} \left(A_3 \left(\lfloor C_1 \log(n) \rfloor - 1 \right) \right) \gg A_{k-1}(n).$$

For the last $\gg$ inequality, recall that $A_3(m) = 2^{2^{\cdot^{\cdot^{\cdot^2}}}} \Big\} m$.

References

[1] O. Baker and T. R. Riley. Cannon-Thurston maps do not always exist. *Forum Math. Sigma*, 1:e3, 11, 2013.

[2] J. Barnard, N. Brady, and P. Dani. Super-exponential distortion of subgroups of CAT(-1) groups. *Algebr. Geom. Topol.*, 7:301–308, 2007.

[3] A. A. Bernasconi. *On HNN-extensions and the complexity of the word problem for one-relator groups*. PhD thesis, University of Utah, 1994. http://www.math.utah.edu/~sg/Papers/bernasconi-thesis.pdf.

[4] M. Bestvina, M. Feighn, and M. Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geom. Funct. Anal.*, 7(2):215–244, 1997.

[5] M. Bestvina and G. Mess. The boundary of negatively curved groups. *J. Amer. Math. Soc.*, 4(3):469–481, 1991.

[6] M. Bonk and O. Schramm. Embeddings of Gromov hyperbolic spaces. *Geom. Funct. Anal.*, 10(2):266–306, 2000.

[7] N. Brady, W. Dison, and T. Riley. Hyperbolic hydra. *Groups Geom. Dyn.*, 7(4):961–976, 2013.

[8] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[9] P. Brinkmann. Splittings of mapping tori of free group automorphisms. *Geom. Dedicata*, 93:191–203, 2002.

[10] K. S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer–Verlag, New York, 1994. Corrected reprint of the 1982 original.

[11] C. Calude. *Theories of computational complexity*, volume 35 of *Annals of Discrete Mathematics*. North-Holland Publishing Co., Amsterdam, 1988.

[12] J. W. Cannon and W. P. Thurston. Group invariant Peano curves. *Geom. Topol.*, 11:1315–1355, 2007.

[13] W. Dison and T. R. Riley. Hydra groups. *Comment. Math. Helv.*, 88(3):507–540, 2013.

[14] E. Ghys and P. de la Harpe, editors. *Sur les Groups Hyperbolic d'après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser, 1990.

[15] M. Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in group theory*, volume 8 of *MSRI publications*, pages 75–263. Springer–Verlag, 1987.

[16] I. Kapovich and N. Benakli. Boundaries of hyperbolic groups. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, volume 296 of *Contemp. Math.*, pages 39–93. Amer. Math. Soc., Providence, RI, 2002.

[17] I. Kapovich and M. Lustig. Cannon–Thurston fibers for iwip automorphisms of F_N. *J. Lond. Math. Soc. (2)*, 91(1):203–224, 2015.

[18] M. Kapovich and B. Kleiner. Hyperbolic groups with low-dimensional boundary. *Ann. Scient. Éc. Norm. Sup.*, pages 647–669, 2000.

[19] H. Lebesgue. Sur les intégrales singulières. *Ann. Fac. Sci. Toulouse Sci. Math. Sci. Phys. (3)*, 1:25–117, 1909.

[20] Y. Matsuda and S. Oguni. On Cannon-Thurston maps for relatively hyperbolic groups. *J. Group Theory*, 17(1):41–47, 2014.

[21] T. C. McMullen. Local connectivity, Kleinian groups and geodesics on the blowup of the torus. *Invent. Math.*, 146(1):35–91, 2001.

[22] Y. N. Minsky. The classification of punctured-torus groups. *Ann. of Math. (2)*, 149(2):559–626, 1999.

[23] M. Mitra. Cannon–Thurston maps for hyperbolic group extensions. *Topology*, 37(3):527–538, 1998.

[24] M. Mitra. Cannon–Thurston maps for trees of hyperbolic metric spaces. *J. Differential Geometry*, 48:135–164, 1998.

[25] H. Miyachi. Moduli of continuity of Cannon–Thurston maps. In *Spaces of Kleinian groups*, volume 329 of *London Math. Soc. Lecture Note Ser.*, pages 121–149. Cambridge Univ. Press, Cambridge, 2006.

[26] M. Mj. Cannon–Thurston maps. Survey article for the proceedings of the 2018 ICM, arXiv:1712.00760.

[27] M. Mj. Cannon–Thurston maps for Kleinian groups. *Forum Math. Pi*, 5:e1, 49, 2017.

[28] H. Short *et al.*. Notes on word hyperbolic groups. In E. Ghys, A. Haefliger, and A. Verjovsky, editors, *Group Theory from a Geometrical Viewpoint (Trieste, 1990)*, volume 8 of *MSRI publications*, pages 3–63. World Scientific Publishing, River Edge, N.J., 1991. www.cmi.univ-mrs.fr/~hamish/Papers/MSRInotes2004.pdf.