

Extrinsic versus intrinsic diameter for Riemannian filling–discs and van Kampen diagrams

M. R. Bridson and T. R. Riley*

Abstract

The diameter of a disc filling a loop in the universal covering of a Riemannian manifold M may be measured *extrinsically* using the distance function on the ambient space or *intrinsically* using the induced length metric on the disc. Correspondingly, the diameter of a van Kampen diagram Δ filling a word that represents the identity in a finitely presented group Γ can either be measured *intrinsically* in the 1–skeleton of Δ or *extrinsically* in the Cayley graph of Γ . We construct the first examples of closed manifolds M and finitely presented groups $\Gamma = \pi_1 M$ for which this choice — intrinsic versus extrinsic — gives rise to qualitatively different min–diameter filling functions.

2000 Mathematics Subject Classification: 53C23, 20F65, 20F10

Key words and phrases: diameter, isodiametric inequality, van Kampen diagram

1 Introduction and preliminaries

A continuous map $F : \mathbb{D}^2 \rightarrow \widetilde{M}$ from the unit disc $\mathbb{D}^2 \subset \mathbb{R}^2$ to the universal cover of a Riemannian manifold M is called a *disc–filling* for a loop c in \widetilde{M} when the restriction of F to $\partial\mathbb{D}^2$ is a monotone reparameterisation of c . According to context, one might measure the diameter of such a disc–filling *intrinsically* using the pseudo–distance $D(a, b)$ on \mathbb{D}^2 defined as the infimum of the lengths of curves $F \circ p$ such that $p : [0, 1] \rightarrow \mathbb{D}^2$ is a path from a to b , or *extrinsically* using the distance function d on the ambient manifold:

$$\begin{aligned} \text{IDiam}(F) &:= \sup \{ D(a, b) \mid a, b \in \mathbb{D}^2 \} \quad \text{and} \\ \text{EDiam}(F) &:= \sup \{ d(F(a), F(b)) \mid a, b \in \mathbb{D}^2 \}. \end{aligned}$$

*The first author was supported in part by fellowships from the EPSRC of Great Britain and the second author in part by NSF grant 0404767.

These two options give different functionals on the space of rectifiable loops in \widetilde{M} : the *intrinsic* and *extrinsic* diameter functionals $\text{IDiam}_{\widetilde{M}}, \text{EDiam}_{\widetilde{M}} : [0, \infty) \rightarrow [0, \infty)$, defined by

$$\begin{aligned}\text{IDiam}_{\widetilde{M}}(l) &:= \sup_c \inf_F \text{IDiam}(F) \quad \text{and} \\ \text{EDiam}_{\widetilde{M}}(l) &:= \sup_c \inf_F \text{EDiam}(F)\end{aligned}$$

where, in each case, the supremum ranges over loops c in \widetilde{M} of length at most l and the infimum over disc-fillings of c .

There is an obvious inequality $\text{IDiam}(F) \geq \text{EDiam}(F)$, which passes to the functionals: $\text{IDiam}_{\widetilde{M}}(l) \geq \text{EDiam}_{\widetilde{M}}(l)$ for all l . One anticipates that for certain compact manifolds M , families of minimal intrinsic diameter filling-discs might *fold-back* on themselves so as to have smaller extrinsic than intrinsic diameter, and the two functionals might then differ asymptotically. However, it has proved hard to animate this intuition with examples. In this paper we overcome this difficulty.

Theorem 1.1 *For every $\alpha > 0$, there is a closed connected Riemannian manifold M and some $\beta > 0$ such that $\text{EDiam}_{\widetilde{M}}(\ell) \preceq \ell^\beta$ and $\text{IDiam}_{\widetilde{M}}(\ell) \succeq \ell^{\alpha+\beta}$.*

[For functions $f, g : [0, \infty) \rightarrow [0, \infty)$ we write $f \preceq g$ when there exists $C > 0$ such that $f(l) \leq Cg(Cl + C) + Cl + C$ for all $l \in [0, \infty)$, and we say $f \simeq g$ when $f \preceq g$ and $g \preceq f$. We extend these relations to accommodate functions with domain \mathbb{N} by declaring such functions to be constant on the half-open intervals $[n, n + 1)$.]

Our main arguments will not be cast in the language of Riemannian manifolds but in terms of combinatorial filling problems for finitely presented groups. There is a close correspondence, given full voice by M. Gromov [7, 8], between the geometry of filling-discs in compact Riemannian manifolds M and the complexity of the word problem in $\pi_1 M$. Most famously, the 2-dimensional, genus-0 isoperimetric function of M has the same asymptotic behaviour as (i.e. is \simeq equivalent to) the Dehn function of any finite presentation of $\pi_1 M$. (See [3] for details.) But the correspondence is wider: most natural ways of bounding the geometry of filling-discs in a manifold correspond to measures of the complexity of the word problem in $\pi_1 M$. The following theorem, proved in Section 2, says that it applies to $\text{IDiam}_{\widetilde{M}}$ and $\text{EDiam}_{\widetilde{M}}$, and their group theoretic analogues $\text{IDiam}_{\mathcal{P}}$ and $\text{EDiam}_{\mathcal{P}}$ (as defined below).

Theorem 1.2 *If \mathcal{P} is a finite presentation of the fundamental group of a closed, connected, Riemannian manifold M then*

$$\text{IDiam}_{\mathcal{P}} \simeq \text{IDiam}_{\widetilde{M}} \quad \text{and} \quad \text{EDiam}_{\mathcal{P}} \simeq \text{EDiam}_{\widetilde{M}}.$$

Let $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ be a finite presentation of a group Γ . The length of a word w in the free monoid $(\mathcal{A}^{\pm 1})^*$ is denoted $\ell(w)$. Let d denote the word metric on Γ associated to \mathcal{A} . If $w = 1$ in Γ (in which case w is called *null-homotopic*) then there exist equalities in the free group $F(\mathcal{A})$ of the form

$$w = \prod_{i=1}^N u_i^{-1} r_i u_i,$$

where $r_i \in \mathcal{R}^{\pm 1}$ and $u_i \in (\mathcal{A}^{\pm 1})^*$. Ranging over all such products w_0 freely equal to w , define $\text{IDiam}_{\mathcal{P}}(w)$ to be the minimum of $\max_{1 \leq i \leq N} \ell(u_i)$, and define $\text{EDiam}_{\mathcal{P}}(w)$ to be the minimum of

$$\max \{ d(1, \bar{p}) \mid p \text{ a prefix of } w_0 \}$$

where \bar{p} denotes the element of Γ represented by p . Define $\text{IDiam}_{\mathcal{P}}, \text{EDiam}_{\mathcal{P}} : \mathbb{N} \rightarrow \mathbb{N}$ by letting $\text{IDiam}_{\mathcal{P}}(n)$ and $\text{EDiam}_{\mathcal{P}}(n)$ be the maxima of $\text{IDiam}(w)$ and $\text{EDiam}(w)$, respectively, over all null-homotopic words w with $\ell(w) \leq n$.

As we will explain in Section 2, one can reinterpret $\text{IDiam}_{\mathcal{P}}(w)$ and $\text{EDiam}_{\mathcal{P}}(w)$ using “combinatorial filling-discs” (*van Kampen diagrams*) for loops in the Cayley 2-complex of \mathcal{P} . Then relating the Cayley 2-complex of \mathcal{P} to \widetilde{M} will lead to a proof of Theorem 1.2.

In Section 3 we describe an infinite family of finite presentations whose diameter functionals exhibit the range of behaviour described in the following theorem.

Theorem 1.3 *For every $\alpha > 0$, there is a group with finite presentation \mathcal{P} and some $\beta > 0$ such that $\text{EDiam}_{\mathcal{P}}(n) \preceq n^{\beta}$ and $\text{IDiam}_{\mathcal{P}}(n) \succeq n^{\alpha+\beta}$.*

Theorem 1.1 is an immediate consequence of Theorems 1.2 and 1.3 because, as is well known, every finitely presentable group is the fundamental group of a closed connected Riemannian manifold.

Our results prompt the following questions. What is the optimal upper bound for $\text{IDiam}_{\mathcal{P}}$ in terms of $\text{EDiam}_{\mathcal{P}}$ for general finite presentations \mathcal{P} ? How might one construct a finite presentation \mathcal{P} for which $\text{EDiam}_{\mathcal{P}}(n)$ is bounded above by a polynomial function and $\text{IDiam}_{\mathcal{P}}(n)$ is bounded below by an exponential function?

1.1 A sketch of the proof of Theorem 1.3

We prove Theorem 1.3 by constructing a novel family of groups $\Psi_{k,m}$. They are obtained by amalgamating auxiliary groups Γ_m and Φ_k along an infinite cyclic subgroup generated by a letter t .

For each positive integer n , one finds an edge-loop in the Cayley 2-complex of Γ_m that has length roughly n and intrinsic diameter roughly n^m . And the reason for this large intrinsic diameter is that in any filling (that is, any *van Kampen diagram*), there is a family of concentric rings of 2-cells (specifically *t-rings*, as defined in Section 4) that nest to a depth of approximately n^m .

Let us turn our attention to the role of Φ_k . For each integer n there are “shortcut words” of length roughly $n^{m/k}$ that equal t^{n^m} in Φ_k . In the Cayley 2-complex of $\Psi_{k,m}$ these mitigate the effect of the nested *t-rings* and cause our large intrinsic-diameter diagrams to have smaller extrinsic diameter. This is illustrated in the leftmost diagram of Figure 1.

However there remains a significant problem. In the Cayley 2-complex of Φ_k , the path corresponding to t^{n^m} and the path of its shortcut word, together form a loop. This loop can be filled. We call any filling-disc for this loop a *shortcut diagram* — Δ of the middle diagram of Figure 1 is an illustration. One could insert two back-to-back copies of the shortcut diagram to get a new filling-disc for our original loop, as illustrated in the rightmost diagram of the figure. The danger is that this lowers intrinsic diameter and thereby stymies our efforts to separate the two diameter functionals of $\Psi_{k,m}$. But Φ_k is built in such a way that the shortcut diagrams are *fat* — that is, they themselves have large intrinsic diameter (see Section 5.2). So inserting shortcut diagrams decreases *intrinsic* diameter far less than the presence of shortcut words decreases extrinsic diameter.

We would like to say that the upshot is that for $\Psi_{k,m}$, the intrinsic diameter functional is at least n^m and the extrinsic diameter functional at most $n^{m/k}$. But, in truth, technicalities make our bounds, detailed in Theorem 7.3, more complicated.

A noteworthy aspect of the construction of Φ_k is the use of an *asymmetric* HNN extension J (cf. Section 3.1). We anticipate such extensions will prove useful in other contexts. We build $\Phi_k \cong \Theta_k *_{\mathbb{Z}} J *_{\mathbb{Z}} \hat{\Theta}_k$ by amalgamating J with groups $\Theta_k \cong \hat{\Theta}_k$ of the form $\mathbb{Z}^k \rtimes F_2$. We build Γ_m by starting with a finitely generated free group and taking a number of HNN extensions and amalgamated free products. All of these groups have compact classifying spaces; $\Psi_{k,m}$ has geometric dimension $k + 1$. The Dehn function of $\Psi_{k,m}$ is at most $n \mapsto C^{n^k}$ for some constant $C > 0$ — see Remark 7.5.




Figure 1: Inserting a *shortcut diagram* Δ .

1.2 The organisation of this article

In Section 2 we prove Theorem 1.2, establishing the qualitative agreement of the group theoretic and Riemannian definitions of diameter. In Section 3 we present the groups $\Psi_{k,m} = \Phi_k *_{\mathbb{Z}} \Gamma_m$ that will be used to prove Theorem 1.3. Section 4 is a brief discussion of *rings* and *corridors* — key tools for analysing van Kampen diagrams. In Sections 5 and 6 we explain the salient isodiametric and distortion properties of Φ_k and Γ_m . In Section 7 we prove the main theorem, modulo generalities about the diagrammatic behaviour of retracts and amalgams, which we postpone to Section 8. Section 9 is dedicated to a proof that, up to \simeq equivalence, extrinsic and intrinsic diameter are quasi-isometry invariants amongst finitely presentable groups.

2 Riemannian versus combinatorial diameter

In this section we reformulate the combinatorial intrinsic and extrinsic diameter filling functions for finitely presented groups Γ in terms of van Kampen diagrams, and then relate them to their Riemannian analogues.

2.1 Diameter of van Kampen diagrams

Let $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ be a finite presentation of a group Γ . Let Δ be a van Kampen diagram over \mathcal{P} with base vertex $\star \in \partial\Delta$. (We assume that the reader is familiar with basic definitions and properties of van Kampen diagrams and Cayley 2-complexes — [3] is a recent survey.) Let ρ be the path metric on $\Delta^{(1)}$ in which every edge has length one and let d denote the word metric on Γ associated to \mathcal{A} . Let \mathcal{C} be the Cayley graph of Γ with respect to \mathcal{A} and let $\pi : \Delta^{(1)} \rightarrow \mathcal{C}$ be the label-preserving graph-morphism with $\pi(\star) = 1$.

Define the *intrinsic* and *extrinsic diameter* of Δ by

$$\begin{aligned}\text{IDiam}(\Delta) &:= \max \left\{ \rho(\star, v) \mid v \in \Delta^{(0)} \right\}, \\ \text{EDiam}_{\mathcal{P}}(\Delta) &:= \max \left\{ d(1, \pi(v)) \mid v \in \Delta^{(0)} \right\}.\end{aligned}$$

We shall show that the algebraic definitions of diameter for null-homotopic words given in Section 1 are equivalent to

$$\begin{aligned}\text{IDiam}_{\mathcal{P}}(w) &:= \min \{ \text{IDiam}(\Delta) \mid \Delta \text{ a van Kampen diagram for } w \}, \text{ and} \\ \text{EDiam}_{\mathcal{P}}(w) &:= \min \{ \text{EDiam}_{\mathcal{P}}(\Delta) \mid \Delta \text{ a van Kampen diagram for } w \}.\end{aligned}$$

Given a van Kampen diagram Δ for w , one can cut it open along a maximal tree \mathcal{T} in $\Delta^{(1)}$ to produce a “lollipop” diagram Δ_0 (with a map $\pi_0 : \Delta_0 \rightarrow \mathcal{C}$) whose boundary circuit is labelled by a word $w_0 = \prod_{i=1}^N u_i^{-1} r_i u_i$ that equals w in $F(\mathcal{A})$ and has $r_i \in \mathcal{R}^{\pm 1}$ for all i . There is a natural combinatorial folding map $\varphi : \Delta_0 \twoheadrightarrow \Delta$ such that $\pi \circ \varphi = \pi_0$.

It follows that $\max \{ d(1, \bar{p}) \mid p \text{ a prefix of } w_0 \}$, where \bar{p} denotes the element of Γ represented by p , is at most $\text{EDiam}_{\mathcal{P}}(\Delta)$, because p corresponds to a vertex on $\partial\Delta_0$. Thus the algebraic version of $\text{EDiam}_{\mathcal{P}}(w)$ is bounded above by the geometric version. For the opposite inequality, suppose that x is a word of the form $\prod_{i=1}^N u_i^{-1} r_i u_i$ where $r_i \in \mathcal{R}^{\pm 1}$ for all i , and $x = w$ in $F(\mathcal{A})$. Define $m := \max \{ d(1, \bar{p}) \mid p \text{ a prefix of } x \}$. Suppose x minimises m and is of minimal length amongst all words that minimise m . There is an obvious lollipop diagram with boundary word x . A van Kampen diagram Δ for w with $\text{EDiam}_{\mathcal{P}}(\Delta) \leq m$ can be obtained from this by successively folding together pairs of edges with common initial point and identical labels. (A concern is that in the course of this folding, 2-spheres might be pinched off, but minimising $\ell(x)$ avoids this.)

$\text{IDiam}_{\mathcal{P}}(w)$, as defined above, differs from $\text{IDiam}_{\mathcal{P}}(w)$ of Section 1 by at most $\frac{1}{2} \max \{ \ell(r) \mid r \in \mathcal{R} \} + \ell(w)$. This can be shown by an argument similar to the one above, except that for the first inequality the tree \mathcal{T} should be chosen to be a maximal geodesic tree based at \star , and when proving the reverse inequality one must specify first that x minimises $\max_i \ell(u_i)$ and then that it has minimal length subject to this constraint. This discrepancy is of no great consequence for us since it has no effect on the \simeq -class of the resulting functional.

We will use the van Kampen-diagram interpretations of $\text{IDiam}_{\mathcal{P}}(w)$ and $\text{EDiam}_{\mathcal{P}}(w)$ in the remainder of this article.

2.2 Generality

With an eye to future applications and to highlight the essential ingredients of our arguments, in this subsection and the next we will work in a more general setting than that of Riemannian manifolds.

Definition 2.1 *Define the intrinsic and extrinsic diameters of a disc-filling $F : \mathbb{D}^2 \rightarrow X$ in a metric space (X, d_X) by*

$$\begin{aligned} \text{IDiam}(F) &:= \sup \{ d_F(a, b) \mid a, b \in \mathbb{D}^2 \}, \text{ and} \\ \text{EDiam}(F) &:= \sup \{ d_X(a, b) \mid a, b \in \mathbb{D}^2 \} \end{aligned}$$

where d_F is the pseudometric

$$d_F(a, b) := \inf \{ \ell(F \circ p) \mid p \text{ a path in } \mathbb{D}^2 \text{ from } a \text{ to } b \}$$

on \mathbb{D}^2 . Define $\text{IDiam}_X, \text{EDiam}_X : [0, \infty) \rightarrow [0, \infty)$ by

$$\begin{aligned} \text{IDiam}_X(l) &:= \sup_c \inf_F \{ \text{IDiam}(F) \mid \text{disc-fillings } F \text{ of loops } c \text{ with } \ell(c) \leq l \}, \\ \text{EDiam}_X(l) &:= \sup_c \inf_F \{ \text{EDiam}(F) \mid \text{disc-fillings } F \text{ of loops } c \text{ with } \ell(c) \leq l \}. \end{aligned}$$

[In closer analogy with the definitions in the previous subsection, one could deal instead with based loops and discs and define diameter functions in terms of distance from the basepoint. This makes little difference: the resulting functions $[0, \infty) \rightarrow [0, \infty)$ are \simeq .]

Lemma 2.2 *Let X be the universal cover of a compact geodesic space Y for which there exist $\mu, A > 0$ such that every loop of length less than μ admits a disc-filling of intrinsic diameter less than A . Equip X with the induced length metric. Then IDiam_X and EDiam_X are well-defined functions $[0, \infty) \rightarrow [0, \infty)$.*

Proof. First observe that every rectifiable loop in X admits a disc-filling of finite intrinsic diameter: given a disc-filling $F : \mathbb{D}^2 \rightarrow X$ of a rectifiable loop one can triangulate the disc so that the image of each edge has diameter at most $\mu/3$, and then one can modify F away from $\partial\mathbb{D}^2$ and the vertex set so that its restriction to each internal edge is a geodesic and every triangle has intrinsic diameter at most A .

Next, suppose c is a loop in X of length l . Shrinking μ if necessary, we may assume that balls of radius μ in Y lift to X . Cover Y with a maximal collection of disjoint balls of radius $\mu/10 > 0$; let $\Lambda \subset X$ be the set of lifts

of their centres. Partition c into $m \leq 1 + 10\ell/\mu$ arcs, each of length at most $\mu/10$ and with end points v_0, v_1, \dots, v_m . Each v_i is a distance at most $\mu/10$ from a point $u_i \in \Lambda$, and the distance from u_i to u_{i+1} is less than $\mu/2$ (indices modulo $m+1$). The loops made by joining v_i to u_i , then u_i to u_{i+1} , then u_{i+1} to v_{i+1} (each by a geodesic), and then v_{i+1} to v_i by an arc of c , each have total length at most μ . So, by hypothesis, they admit disc-fillings with intrinsic diameter at most A . Such disc-fillings together form a collar between c and a piecewise geodesic loop formed by concatenating at most $1 + 10\ell/\mu$ geodesic segments, each of length at most $\mu/2$ with endpoints in Λ . Modulo the action of $\pi_1 Y$, there are only finitely many such piecewise geodesic loops, and by the argument in the previous paragraph each one admits a filling of finite intrinsic diameter. It follows that IDiam_X and (hence) EDiam_X are well-defined functions. \blacksquare

Remark 2.3 *If $X = \widetilde{M}$, the universal cover of a closed connected Riemannian manifold M , then it satisfies the conditions of Lemma 2.2. Indeed, the required μ and A exist for any cocompact space X that is locally uniquely geodesic, for example a space with upper curvature bound in the sense of A.D. Alexandrov — i.e. a $\text{CAT}(\kappa)$ space [5]: by cocompactness, there exists $\eta > 0$ such that geodesics are unique in balls of radius η , and any loop in such a ball can be filled by coning it off to the centre of the ball using geodesics.*

2.3 The Translation Theorem

Theorem 1.2 is a special case of the following result.

Theorem 2.4 *Suppose a group Γ with finite presentation \mathcal{P} acts properly and cocompactly by isometries on a simply connected geodesic metric space X for which there exist $\mu, A > 0$ such that every loop of length less than μ admits a disc-filling of intrinsic diameter less than A . Then $\text{IDiam}_{\mathcal{P}} \simeq \text{IDiam}_X$ and $\text{EDiam}_{\mathcal{P}} \simeq \text{EDiam}_X$.*

We shall first establish the assertion concerning extrinsic diameter and the relation $\text{IDiam}_X \preceq \text{IDiam}_{\mathcal{P}}$.

Map the Cayley graph of Γ to X as follows: fix a basepoint $p \in X$, choose a geodesic from p to its translate $a \cdot p$ for each generator a of \mathcal{P} , and then extend equivariantly. Following [3], a path in X is called *word-like* if it is the image in X of an edge-path in the Cayley graph.

Each 2-cell in the Cayley 2-complex $\text{Cay}^2(\mathcal{P})$ is attached to the Cayley graph by an edge-loop labelled by one of the defining relations r of \mathcal{P} .

For each r we choose a filling disc F_r of finite intrinsic diameter for the corresponding word-like loop in X based at p . We then map $\text{Cay}^2(\mathcal{P})$ to X by the Γ -equivariant map that sends each 2-cell with boundary label r to a translate of F_r .

Using a collar between an arbitrary rectifiable loop in X and a word-like loop, as in the proof of Lemma 2.2, one can show that there is no change in the \simeq classes of either IDiam_X or EDiam_X if one takes the infima in their definitions to be over fillings of word-like loops only. Having made this reduction, the relation $\text{IDiam}_X \preceq \text{IDiam}_{\mathcal{P}}$ becomes obvious, since one gets an upper bound on the intrinsic diameter of discs filling a word-like loop simply by taking the image in X of a minimal intrinsic diameter van Kampen diagram Δ for the appropriate word. A technical concern here is that Δ need not be a topological 2-disc, but rather may be a finite planar tree-like arrangement of topological 2-discs and 1-dimensional arcs. This can be overcome by extending $\Delta \rightarrow X$ to a small regular neighbourhood V of $\Delta \subset \mathbb{R}^2$ by a map that is constant on the slices $\{x\} \times [0, 1)$ of the annulus $V \setminus \Delta \approx \partial V \times [0, 1)$.

The same argument yields $\text{EDiam}_X \preceq \text{EDiam}_{\mathcal{P}}$, and the converse is an easy approximation argument (the details of which are included in Lemma 2.5). Briefly, noting that the map of the Cayley graph of Γ to X is a quasi-isometry, we need only show that a disc-filling $F : \mathbb{D}^2 \rightarrow X$ of a word-like loop gives rise to a van Kampen diagram for the corresponding word that is C_0 -close to F . Such an approximating disc is obtained by simply taking a fine triangulation of \mathbb{D}^2 , and then labelling it as in the proof of Lemma 2.5.

The remainder of this section is devoted to the most difficult relation in Theorem 2.4, namely $\text{IDiam}_{\mathcal{P}} \preceq \text{IDiam}_X$. Like the argument sketched above, the proof of this assertion involves constructing a suitable combinatorial approximation to a Riemannian filling of a word-like loop. The control required in this approximation is spelt out in the following lemma.

Lemma 2.5 *To prove $\text{IDiam}_{\mathcal{P}} \preceq \text{IDiam}_X$, it suffices to exhibit constants L_0, L_1, λ such that, given a filling $F : \mathbb{D}^2 \rightarrow X$ of intrinsic diameter δ for a word-like loop c , one can construct a combinatorial cellulation Δ of \mathbb{D}^2 and a map $f : \Delta^{(0)} \rightarrow X$ with the following properties:*

- (1) *the attaching map of each 2-cell of Δ has combinatorial length at most L_0 ,*
- (2) *adjacent vertices in Δ are mapped by f to points that are a distance at most L_1 apart in X ,*

- (3) from each vertex of Δ , one can reach $\partial\Delta$ by traversing a path consisting of at most $\lambda(\delta + 1)$ edges, and
- (4) the vertices $\partial\Delta^{(0)}$ on $\partial\Delta$ are mapped by f to points on c , and their cyclic ordering is preserved; moreover, every vertex of c is in $f(\partial\Delta^{(0)})$.

Proof. In light of Theorem 9.1, we may take \mathcal{P} to be a finite presentation well adapted to the geometry of X . Thus we fix $\xi > 0$ so that X is the ξ -neighbourhood of $\Gamma \cdot p$, and as generators of Γ we take the set \mathcal{A} of those γ such that $d(p, \gamma(p)) \leq L_1 + 2\xi$. As relations \mathcal{R} we take all words of length at most L_0 in the letters \mathcal{A} that equal the identity in Γ . The remainder of our proof shows that these relations suffice; cf. [5, page 135].

Given a word w in the generators \mathcal{A} that equals 1 in Γ , we consider the corresponding word-like loop c in X . By construction, there are constants α, β such that the length of c is at most $\alpha\ell(w) + \beta$. So we will be done if, given a disc-filling of c with intrinsic diameter δ , we can exhibit a van Kampen diagram for w with intrinsic diameter bounded above by a linear function of δ .

Let Δ be as described in the statement of the lemma and label each vertex $v \in \Delta^{(0)} \setminus \partial\Delta$ by a group element γ_v such that $\gamma_v(p) \in \Gamma \cdot p$ is within ξ of $f(v)$. Label an edge from a vertex v to a vertex v' by $\gamma_v^{-1}\gamma_{v'}$, which is in \mathcal{A} . For vertices v of $\partial\Delta$, we choose γ_v to be one of the endpoints of the edge of the word-like loop c in which (the image of) v lies.

By hypothesis, the boundary of each 2-cell in the cellulation is then labelled by a word that belongs to \mathcal{R} . This process does not quite yield a van Kampen diagram for the original word w , but rather for a word w' that is equal to w modulo the cancellation of edge-labels 1. (Our convention on the choosing of γ_v for vertices on $\partial\mathbb{D}^2$ ensures this equality.) Collapsing such edges yields the desired diagram. ■

Continuing the proof that $\text{IDiam}_{\mathcal{P}} \preceq \text{IDiam}_X$, we concentrate on a disc-filling $F : \mathbb{D}^2 \rightarrow X$ of intrinsic diameter δ for a word-like loop c with basepoint $F(\star) = p$. We will construct f and Δ satisfying the conditions of Lemma 2.5 with $L_0 = 6$, $L_1 = 5/2$ and $\lambda = 1$ in three steps.

(i) *Constructing the tree of a fine triangulation.*

Since F is continuous and \mathbb{D}^2 is compact, \mathbb{D}^2 admits a finite triangulation $\hat{\Delta}$ such that F maps each triangle to a subset of X of diameter at most $1/2$. Fix $\varepsilon > 0$. Fix an ordering v_0, v_1, v_2, \dots on the vertices V of $\hat{\Delta} \setminus \partial\hat{\Delta}$ and, proceeding along the list, choose an embedded arc σ_v in \mathbb{D}^2 joining v to \star as follows so that $T := \bigcup_{v \in V} \sigma_v$ is a tree. Having defined σ_{v_i} for

each $i < m$, to define σ_{v_m} first choose an embedded arc from v_m to \star of length at most $\varepsilon + d_F(\star, v_m)$, and arrange that it pass through no vertices in $V \setminus \{v_m, \star\}$ (more on this in a moment) and finally replace its terminal segment beginning at the first point x it meets in $\bigcup_{i < m} \sigma_{v_i}$, with the arc in $\bigcup_{i < m} \sigma_{v_i}$ between x and \star . By construction V is the set of leaves of T . Were it not for the ε -error terms, the diameter of T would be at most 2δ . We choose ε small enough so that the diameter is no more than $2\delta + 1$.

[Technicalities: We introduced ε to avoid a discussion of the existence of geodesics in the pseudometric space (\mathbb{D}^2, d_F) . We can arrange that σ_{v_m} not pass through $V \setminus \{v_m, \star\}$ by replacing F with a map F' that is constant on a small disc about each $v \in V$ and which, on the complement E of the union of these disjoint discs, is $F \circ h$ where $h : E \rightarrow \mathbb{D}^2$ is a homeomorphism stretching an annular neighbourhood of each deleted disc to a punctured disc with the same centre. The distances between distinct vertices in (\mathbb{D}^2, d_F) and $(\mathbb{D}^2, d_{F'})$ are the same and, in the latter, a path can take a detour near any vertex without adding any length.]

We intend to combine T and $\hat{\Delta}$ to give a cellulation Δ with the properties required for Lemma 2.5. However, as things stand, the intersection of T with the cells of $\hat{\Delta}$ could be extremely complicated. To cope with this we introduce a hyperbolic device to extract the essential features of the pattern of intersections.

(ii) *Taming the intersection of T with $\hat{\Delta}$.*

As above, V is the set of vertices of $\hat{\Delta}$ that lie in the interior of \mathbb{D}^2 . We impose a complete hyperbolic metric on $\mathbb{D}^2 \setminus V$ in such a way that $\partial\mathbb{D}^2$ is a regular geodesic polygon and the 1-cells of $\hat{\Delta}$ are geodesics in the hyperbolic metric. For each $v \in V$ let τ_v be the unique geodesic in the same homotopy class of ideal paths from v to \star in $\mathbb{D}^2 \setminus V$ as σ_v . To aid visualisation, we are going to talk of these τ_v arcs as being *red*.

The point of what we have just done is that it vastly simplifies the pattern of intersections of the σ_v with $\hat{\Delta}^{(1)}$: the pattern of intersections of the τ_v with a triangle is the simple pattern shown in Figure 2 (except if the triangle has \star as one of its vertices then some of the red edges can meet \star , indeed the sides of the triangle may be red). Moreover, the larger-scale geometry has been retained in a way that allows us to construct a cellulation with combinatorial diameter at most a linear function of δ , as we will explain.

(iii) *Constructing the good cellulation Δ and map f .*

We begin by taking the cellulation Δ_0 consisting of $\hat{\Delta}$ overlaid with $\hat{T} := \bigcup_{v \in V} \tau_v$. That is, the vertices are the vertices of $\hat{\Delta}$ and the points of




Figure 2: The intersections of $\bigcup_{v \in V} \tau_v$ with a triangle of $\hat{\Delta}$.

intersection of \hat{T} with edges of $\hat{\Delta}$, and the edges come in two types: *black* edges which are edges of $\hat{\Delta}$ (which may be partitioned into multiple edges by vertices on \hat{T}), and *red* edges, which are arcs in \hat{T} between vertices. Define $f_0 : \Delta_0^{(0)} \rightarrow X$ to be the restriction of F to $\Delta_0^{(0)}$.

We will alter Δ_0 (and f_0) and the paths τ_v until we can control their combinatorial lengths in terms of the lengths of the σ_v and thereby get an upper bound on the intrinsic diameter of Δ_0 .

It will be convenient if whenever $v, v' \in V$ are distinct, τ_v and $\tau_{v'}$ meet only at \star . Achieve this by doubling $(\tau_v \cap \tau_{v'}) \setminus \{\star\}$, and joining doubled vertices by a new black edge — see Figure 3. Call the new combinatorial disc Δ_1 and define $f_1 : \Delta_1^{(0)} \rightarrow \hat{M}$ by letting $f_1(u)$ be $f_0(\bar{u})$ where \bar{u} is the image of u under the obvious combinatorial retraction of Δ_1 onto Δ_0 .




Figure 3: Doubling along the intersection of two paths τ_v and $\tau_{v'}$.

Repeat the following for each $v \in V$ such that τ_v has combinatorial length at least 2. As τ_v is in the same homotopy class of ideal paths from v to \star of $\hat{\Delta} \setminus V$ as σ_v , the sequence of edges τ_v crosses is also crossed by σ_v in the same order, except σ_v may make additional intersections en route. That is, if $\{x_i\}_{i=0}^k$ are the points of intersection of τ_v with $\hat{\Delta}$ numbered in order along τ_v with $x_0 = \star$, $x_k = v$ and all other x_i on so a black edge e_i , then there are points $\{y_i\}_{i=0}^k$ occurring in order along σ_v with $y_0 = x_0$, $y_k = x_k$ and y_i on e_i for all other i . For all i let n_i be the integer part of the length of the segment of σ_v from y_i to \star . Note that $i \mapsto n_i$ increases monotonically from 0 to $\lfloor \ell(\sigma_v) \rfloor$, and takes all possible integer values in that range because

of the condition that the diameters of the images of the triangles in $\hat{\Delta}$ are at most $1/2$. For each $n \in \mathbb{N}$, collapse to a single vertex the minimal arc of τ_v that includes all red vertices x_i such that $n_i = n$. Note that as each arc is collapsed the complex remains topologically a disc since the arc does not run between two boundary vertices.

The result is the desired cellulation Δ of \mathbb{D}^2 and we define $f : \Delta^{(0)} \rightarrow \widetilde{M}$ by setting $f(\bar{u})$ to be $f(u)$ for any choice of preimage u of \bar{u} under the quotient map $\Delta_1 \twoheadrightarrow \Delta$. All that remains is to explain why Δ and f satisfy the conditions of Lemma 2.5 with $L_0 = 6$, $L_1 = 5/2$, and $\lambda = 1$.

The faces of Δ each have at most six sides since this was the case for Δ_1 (six being realised when red edges cut across each of the corners of a triangle of $\hat{\Delta}$) and the subsequent collapses of edges can only decrease the number of sides, which proves (1).

For (2), first suppose x and x' are vertices in Δ joined by a red edge. Then they are images of vertices x_i and $x_{i'}$ on the same path τ_v in Δ_1 such that $f(x) = f_1(x_i)$, $f(x') = f_1(x_{i'})$ and $|n_i - n_{i'}| = 1$. But then, as $|n_i - n_{i'}| = 1$,

$$d(f(x), f(x')) = d(f_1(x_i), f_1(x_{i'})) \leq 2.$$

Next suppose x and x' are vertices in Δ joined by a black edge. Assume they are images of vertices x_i and $x_{i'}$ on paths τ_v and $\tau_{v'}$ in Δ_1 such that x_i and $x_{i'}$ are joined by a black edge. Then $d(f_1(x_i), f_1(x_{i'})) \leq 1/2$ and there are x_j and $x_{j'}$ on τ_v and $\tau_{v'}$ such that $f(x) = f_1(x_j)$, $f(x') = f_1(x_{j'})$, $n_i = n_j$, and $n_{i'} = n_{j'}$. So

$$\begin{aligned} d(f(x), f(x')) &= d(f_1(x_j), f_1(x_{j'})) \\ &= d(f_1(x_j), f_1(x_i)) + d(f_1(x_i), f_1(x_{i'})) + d(f_1(x_{i'}), f_1(x_{j'})) \\ &\leq 1 + \frac{1}{2} + 1 = \frac{5}{2}. \end{aligned}$$

The remaining possibility is that x and x' are joined by a black edge but one or both is in $\partial\Delta \setminus \{\star\}$ and so is not on any τ_v in Δ_1 . Similar considerations to the previous case yield $d(f(x), f(x')) \leq 1 + 1/2$.

By construction, for all $v \in V$, the combinatorial length of the images in Δ of τ_v is at most the length of σ_v , and so at most $\delta + 1$ and we have (3). And finally, (4) is immediate from the way we constructed f and Δ .

3 The groups for Theorem 1.3

The groups that we use form a 2-parameter family

$$\Psi_{k,m} = \Phi_k *_\mathbb{Z} \Gamma_m,$$

where $m > k > 1$ are integers.

3.1 The “Fat-Shortcuts” Groups Φ_k

We construct Φ_k by amalgamating J , a double HNN-extension of \mathbb{Z} , with two isomorphic copies of $\Theta_k = \mathbb{Z}^k \rtimes F_2$, where the action of F_2 is chosen so that a certain cyclic subgroup $\langle s_k \rangle$ of \mathbb{Z}^k is distorted in a precisely controlled manner as will be required in Lemma 5.2. The two copies of Θ_k are amalgamated with J along the cyclic subgroups $\langle s_k \rangle$. [The distortion of a finitely generated subgroup S of a finitely generated group G measures the difference between the word metric on S and the restriction to S of the word metric of G — see [5], page 506.]

We first take a presentation \mathcal{O}_k of $\Theta_k = \mathbb{Z}^k \rtimes F_2$ by choosing a basis s_1, \dots, s_k for \mathbb{Z}^k and specifying the action of $F_2 = \langle f, g \rangle$:

$$\begin{array}{ll} \text{generators} & s_1, \dots, s_k, f, g \\ \\ \text{relations} & f^{-1}s_k f = s_k, \text{ and } \forall i < k, f^{-1}s_i f = s_i s_{i+1} \\ & g^{-1}s_k g = s_k, g^{-1}s_{k-1}g = s_{k-1}, \text{ and } \forall i < k-1, g^{-1}s_i g = s_i s_{i+1} \\ & \forall i \neq j, [s_i, s_j] = 1. \end{array}$$

Let $\widehat{\Theta}_k$ be a second copy of Θ_k with corresponding presentation $\widehat{\mathcal{O}}_k$ in which the generators are $\hat{s}_1, \dots, \hat{s}_k, \hat{f}, \hat{g}$.

We define J to be the amalgam along $\langle b \rangle \cong \mathbb{Z}$ of the *asymmetric HNN extension* $\langle b, t, s \mid t^{-1}bs = b^3, s^{-1}bt = b^3 \rangle$ with the standard HNN extension $\langle b, \hat{s} \mid \hat{s}^{-1}b\hat{s} = b^3 \rangle$. [This is obtained from $\langle b, t \mid t^{-1}b^2t = b^6 \rangle$ by the Tietze move introducing $s = b^{-1}tb^3$.] Thus J has presentation

$$\langle b, t, s, \hat{s} \mid t^{-1}bs = b^3, s^{-1}bt = b^3, \hat{s}^{-1}b\hat{s} = b^3 \rangle.$$

Finally, we define

$$\Phi_k := \Theta_k *_\mathbb{Z} J *_\mathbb{Z} \widehat{\Theta}_k,$$

where the amalgamations identify $s, \hat{s} \in J$ with $s_k \in \Theta_k$ and $\hat{s}_k \in \widehat{\Theta}_k$, respectively. Thus, writing $\mathcal{O}_k = \langle \Sigma_k \mid \Upsilon_k \rangle$ and $\widehat{\mathcal{O}}_k = \langle \widehat{\Sigma}_k \mid \widehat{\Upsilon}_k \rangle$, a presentation for Φ_k is

$$\mathcal{P}_k := \langle b, t, \Sigma_k, \widehat{\Sigma}_k \mid \Upsilon_k, \widehat{\Upsilon}_k, t^{-1}bs_k = b^3, s_k^{-1}bt = b^3, \hat{s}_k^{-1}b\hat{s}_k = b^3 \rangle.$$

The asymmetric nature of the quasi-HNN presentation of $\langle b, s, t \rangle \subset J$ will be important in Section 5.1 when we wish to force shortcuts to be fat — see Figure 5 for an example and Section 5.4 for the general argument. As we indicated in the introduction, this fatness is the key to forcing the behaviour of the intrinsic and extrinsic diameter functionals to diverge.

3.2 The t -rings groups Γ_m

The groups Γ_m , whose presentations \mathcal{Q}_m we give below, are 2-fold HNN extensions of the free-by-free groups $B_m = F_m \rtimes F_2$. By deleting from \mathcal{Q}_m the generator T and the defining relations in which it appears, one recovers the groups constructed by the first author in [2]; these have isodiametric properties that prevent their asymptotic cones from being simply connected.

We define $B_m := F_m \rtimes F_2$, where the first generator σ of $F_2 = \langle \sigma, t \rangle$ acts on F_m as an automorphism with polynomial growth of degree $m - 1$, and t acts trivially. Specifically, B_m has presentation

$$\mathcal{B}_m := \langle \mathcal{A}_m, \sigma, t \mid \sigma^{-1} a_m \sigma = a_m; \forall i < m, \sigma^{-1} a_i \sigma = a_i a_{i+1}; \forall j, [t, a_j] = 1 \rangle,$$

where $\mathcal{A}_m := \{a_1, \dots, a_m\}$. We then obtain Γ_m by taking two successive HNN extensions of B_m : in the first, the stable letter τ commutes with a skewed copy of F_m and, in the second, the stable letter T commutes with τ and t . Thus, abbreviating $\mathcal{B}_m = \langle \mathcal{A}_m, \sigma, t \mid \mathcal{R} \rangle$, we define Γ_m to be the group with (aspherical) presentation

$$\mathcal{Q}_m := \langle \mathcal{A}_m, \sigma, t, \tau, T \mid \mathcal{R}, [t, T] = 1, [\tau, T] = 1, [\tau, a_m t] = 1; \forall i < m, [\tau, a_i] = 1 \rangle.$$

3.3 Preliminary lemmas

The following two lemmas are self-evident.

Lemma 3.1 Γ_m retracts onto $B_m = F_m \rtimes \langle \sigma, t \rangle$ and hence onto $\langle t \rangle$.

Lemma 3.2 The groups Θ_k and $\hat{\Theta}_k$ are retracts of Φ_k via

$$\begin{array}{ccccc} \Theta_k & \xhookrightarrow{\iota} & \Phi_k & \xrightarrow{\psi} & \Theta_k \\ \hat{\Theta}_k & \xhookrightarrow{\hat{\iota}} & \Phi_k & \xrightarrow{\hat{\psi}} & \hat{\Theta}_k, \end{array}$$

where ι and $\hat{\iota}$ are the obvious inclusions, ψ kills $b, \hat{f}, \hat{g}, \hat{s}_1, \dots, \hat{s}_k$, maps t to s_k , and f, g, s_1, \dots, s_k to themselves, and $\hat{\psi}$ kills $b, t, f, g, s_1, \dots, s_k$ and maps $\hat{f}, \hat{g}, \hat{s}_1, \dots, \hat{s}_k$ to themselves.

Since \mathcal{O}_k is the obvious presentation of the semidirect product $\Theta_k = \mathbb{Z}^k \rtimes F_2$ we have:

Lemma 3.3 *If $w \in \{s_1^{\pm 1}, \dots, s_k^{\pm 1}\}^*$ is null-homotopic in \mathcal{O}_k , then it is also null-homotopic in*

$$\langle s_1, \dots, s_k \mid \forall i \neq j, [s_i, s_j] \rangle \cong \mathbb{Z}^k.$$

Since $\langle b \rangle$ and $\langle s_1, \dots, s_{k-1} \rangle \cong \mathbb{Z}^{k-1}$ do not intersect the amalgamated subgroups of $\Phi_{k,m} = \Theta_k *_{\mathbb{Z}} J *_{\mathbb{Z}} \hat{\Theta}_k$, they generate their free product. In more detail:

Lemma 3.4 *If $w \in \{b^{\pm 1}, s_1^{\pm 1}, \dots, s_{k-1}^{\pm 1}\}^*$ is a null-homotopic word in \mathcal{P}_k , then it is also null-homotopic in the subpresentation*

$$\langle b, s_1, \dots, s_{k-1} \mid \forall i \neq j, [s_i, s_j] \rangle \cong \mathbb{Z} * \mathbb{Z}^{k-1}.$$

Lemma 3.5 *For all $n \in \mathbb{N}$, there exists a (positive) word $w_n \in \{s_1, \dots, s_k\}^*$ such that $w_n = f^{-n} s_1^n f^n$ in*

$$F_k \rtimes \mathbb{Z} \cong \langle s_1, \dots, s_k, f \mid f^{-1} s_k f = s_k; \forall i < k, f^{-1} s_i f = s_i s_{i+1} \rangle$$

and

$$w_n = \prod_{i=1}^k s_i^{n \binom{n}{i-1}} \quad \text{in} \quad \mathbb{Z}^k = \langle s_1, \dots, s_k \mid \forall i \neq j, [s_i, s_j] \rangle.$$

Proof. Let w_n be the unique positive word in the s_i such that $w_n = f^{-n} s_1^n f^n$ in $F_k \rtimes \mathbb{Z}$. The stable letter f acts on the abelianisation \mathbb{Z}^k of F_k as left multiplication by the $k \times k$ matrix with ones on the diagonal and subdiagonal and zeros elsewhere. So the result follows from the calculation

$$\begin{pmatrix} 1 & & & & \\ 1 & 1 & & & \\ & 1 & \ddots & & \\ & & \ddots & 1 & \\ & & & 1 & 1 \end{pmatrix}^n = \begin{pmatrix} 1 & & & & \\ \binom{n}{1} & 1 & & & \\ \binom{n}{2} & \binom{n}{1} & \ddots & & \\ \vdots & \vdots & \ddots & 1 & \\ \binom{n}{k-1} & \binom{n}{k-2} & \cdots & \binom{n}{1} & 1 \end{pmatrix}.$$

■

4 t -corridors and t -rings in van Kampen diagrams

The use of t -corridors and t -rings for analysing van Kampen diagrams is well established. Suppose $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ is a presentation, Δ is a van Kampen diagram over \mathcal{P} and $t \in \mathcal{A}$. An edge in Δ is called a t -edge if it is labelled by t . Let Δ^* denote the dual graph to the 1-skeleton of Δ (including a vertex v_∞ dual to the region exterior to Δ).

Suppose σ is either a simple edge-loop in Δ^* and that the edges of σ are all dual to t -edges in Δ . Then the subdiagram C of Δ consisting of all the (closed) 1-cells and 2-cells of Δ dual to vertices and edges of $\sigma \setminus \{v_\infty\}$, is called a t -ring if σ does not include v_∞ and a t -corridor if it does. Let e_1, \dots, e_r be the duals of the edges of σ , in the order they are crossed by σ , and such that in the ring case $e_1 = e_r$ and in the corridor case e_1 and e_r are in ∂e_∞ . The *length* of the t -ring or t -corridor is $r - 1$ (which is zero for t -corridors arising from t -edges not lying in the boundary of any 2-cell). In the corridor case, e_1 and e_r are called the *ends* of the corridor.

Let c_1, \dots, c_{r-1} be the 2-cells of C numbered so that e_i and e_{i+1} are part of ∂c_i for all i . Refer to one vertex of e_i as the left and the other as the right, depending on where it lies as we travel along σ . The *left (right) side* of C is the edge-path in Δ that follows the boundary-cycle ∂c_1 from the left (right) vertex of e_1 to the left (right) vertex of e_2 , and then ∂c_2 from the left (right) vertex of e_2 to the left (right) vertex of e_3 , and so on, terminating at the left (right) vertex of e_r . Note that the sides of C need not be embedded paths in $\Delta^{(1)}$. In the case of the ring, orienting σ anticlockwise in Δ , we call the left side is the *inside* and the right side is the *outside*.

The following lemma contains the basic, crucial observations about t -rings and t -corridors.

Lemma 4.1 *Suppose each $r \in \mathcal{R}$ contains precisely zero or two letters $t^{\pm 1}$. Then every t -edge of Δ lies in either a t -corridor or a t -ring, and the interiors of distinct t -rings/corridors are disjoint.*

In the case of the presentation \mathcal{P}_k of Section 3.1 it will be profitable to have more general definitions: we define an $s_k t$ -corridor or $s_k t$ -ring in a \mathcal{P}_k -van Kampen diagram Δ by allowing the duals of edges in σ to be s_k - or t -edges and allowing σ to be a simple edge-path whose initial and terminal vertices are either v_∞ or dual to 2-cells with boundary label $f^{-1}s_{k-1}fs_k^{-1}s_{k-1}^{-1}$. So an end of an $s_k t$ -corridor is either in $\partial\Delta$ or in the boundary of a 2-cell labelled $f^{-1}s_{k-1}fs_k^{-1}s_{k-1}^{-1}$.

5 Salient features of Φ_k

5.1 How to shortcut powers of t in Φ_k

We fix a positive integer k . The following proposition concerns the existence of *shortcut diagrams* that distort $\langle t \rangle$ in \mathcal{P}_k .

Proposition 5.1 *There is a constant C , depending only on k , such that for all $m > 0$ there is a word u_m of length $\ell(u_m) \leq Cm^{1/k} + C$ and a van Kampen diagram Δ_m for $t^m u_m^{-1}$ over \mathcal{P}_k with $\text{IDiam}(\Delta_m) \leq Cm + C$.*

Before proving this proposition we establish a similar result about the distortion of $\langle s_k \rangle$ in \mathcal{O}_k .

Lemma 5.2 *There is a constant C , depending only on k , such that for all $m > 0$ there is a word v_m of length $\ell(v_m) \leq Cm^{1/k} + C$ and a van Kampen diagram Σ_m for $s_k^m v_m^{-1}$ over \mathcal{O}_k with every vertex of Σ_m a distance (in the 1-skeleton of Σ_m) at most $Cm + C$ from the portion of $\partial\Sigma_m$ labelled v_m .*

Proof. Lemma 3.5 tells us that over the sub-presentation

$$\langle s_1, \dots, s_k, f \mid f^{-1}s_k f = s_k; \forall i < k, f^{-1}s_i f = s_i s_{i+1} \rangle$$

of \mathcal{O}_k , the word $f^{-n}s_1^n f^n$ equals a (positive) word $w_n \in \{s_1, \dots, s_k\}^*$ such that $w_n = \prod_{i=1}^k s_i^{n \binom{n}{i-1}}$ in $\langle s_1, \dots, s_k \rangle \cong \mathbb{Z}^k$. Fix $m > 0$ and let n be the least integer such that $n \binom{n}{k-1} \geq m$. Then $(n - k + 1)^k < (k - 1)!m$ because

$$(n - 1) \binom{n - 1}{k - 1} = \frac{1}{(k - 1)!} (n - 1)^2 (n - 2) \dots (n - k + 1) < m.$$

And so

$$n < (k - 1)!^{1/k} m^{1/k} + k - 1 < k m^{1/k} + k. \quad (1)$$

Let D_n be the van Kampen diagram for $f^{-n}s_1^n f^n w_n^{-1}$ obtained by stacking f -corridors n high, embedded in the plane as illustrated in Figure 4. Let w_{nm} be the shortest prefix of w_n in which the letter s_k occurs m times. Let ρ be the edge-path from the vertex of ∂D_n at which w_{nm} ends, to the portion of ∂D_n labelled by $f^{-n}s_1^n f^n$, that proceeds by travelling up where possible and left otherwise. As ρ never travels left twice consecutively, its length is at most $2n$. Cutting D_n along ρ we obtain two diagrams; one, D_{nm} (which




Figure 4: The \mathcal{O}_k -van Kampen diagrams D_n and Σ_m .

appears shaded in Figure 4), shows that w_{nm} is equal in Θ_k to a word of length at most $4n$.

Next we use the relations $[s_i, s_k]$ to gather all the letters s_k in w_{nm} to the left and produce a word $s_k^m \hat{w}_{nm}$. Diagrammatically, this is achieved by attaching s_k -corridors to D_{nm} . (These fill the triangular region in Figure 4.) Then we use the second stable letter g of Θ_k , which acts on $\langle s_1, \dots, s_k \rangle$ thus:

$$g^{-1}s_{k-1}g = s_{k-1}, \quad g^{-1}s_kg = s_k \quad \text{and} \quad g^{-1}s_i g = s_i s_{i+1}, \forall i < k-1.$$

The word $g^{-n}s_1^n g^n$ and the word obtained from w_n by removing all letters s_k represent the same element in $F_{k-1} \rtimes \langle g \rangle$. This equality is exhibited by a van Kampen diagram \hat{D}_n obtained by retracting D_n in the obvious manner, and a subdiagram \hat{D}_{nm} (a retraction of D_{nm}) of \hat{D}_n portrays an equality in $F_{k-1} \rtimes \langle g \rangle$ of \hat{w}_{nm} with a word of length at most $4n$. Attach \hat{D}_{nm} along \hat{w}_{nm} to get a diagram Σ_m short-cutting s_k^m to a word v_m of length at most $8n$ as shown in Figure 4. Because D_{nm} and \hat{D}_{nm} are stacks of $O(m^{1/k})$ corridors, it is possible to reach the portion of $\partial\Sigma_m$ labelled by v_m from any vertex of D_{nm} or \hat{D}_{nm} by traversing $O(m^{1/k})$ edges of $\Sigma^{(1)}$. The triangular region is a stack of $O(m)$ corridors which can be crossed to reach ∂D_{nm} . The assertion of the lemma then follows from the hypothesis that $\ell(v_m) \leq Cm^{1/k} + C$ and the inequality (1). \blacksquare

Proof of Proposition 5.1. Using stacks of s_k - t -corridors and \hat{s}_k -corridors in the obvious way, one can construct van Kampen diagrams demonstrating that $t^{-m}bs_k^m = b^{3^m}$ and $\hat{s}_k^{-m}b\hat{s}_k^m = b^{3^m}$ over the subpresentations $\langle b, s_k, t \mid t^{-1}bs_k = b^3, s_k^{-1}bt = b^3 \rangle$ and $\langle b, \hat{s}_k \mid \hat{s}_k^{-1}b\hat{s}_k = b^3 \rangle$ of \mathcal{P}_k , respectively. Join these diagrams to give a diagram that demonstrates that $\hat{s}_k^{-m}b\hat{s}_k^m = t^{-m}bs_k^m$ in Φ_k and that has intrinsic diameter at most a constant times m . Then obtain Δ_m as shown in Figure 5: attach the \mathcal{O}_k -van Kampen diagram Σ_m of Lemma 5.2 along s_k^m , and attach a copy of the corresponding $\hat{\mathcal{O}}_k$ -van Kampen diagram $\hat{\Sigma}_m$ and its mirror-image along \hat{s}_k^m and \hat{s}_k^{-m} .

The asserted bound on $\text{IDiam}_{\mathcal{P}_k}(\Delta_m)$ holds for the following reasons. Every vertex in the $t^{-m}bs_k^m = b^{3^m}$ and $\hat{s}_k^{-m}b\hat{s}_k^m = b^{3^m}$ subdiagrams is a distance $O(m)$ from Σ_m or from one of the copies of $\hat{\Sigma}_m$, or from the portion of Δ_m labelled t^m . The distance from any vertex of a subdiagram Σ_m or $\hat{\Sigma}_m$ to $\partial\Delta_m$ is $O(m)$ by Lemma 5.2. Thus one can reach $\partial\Delta_m$ from any of its vertices within the claimed bound. As $\ell(\partial\Delta_m) = O(m)$, one can then follow the boundary circuit to the base point. ■




Figure 5: The \mathcal{P}_k -van Kampen diagram Δ_m shortcutting t^m .

Remark 5.3 *The area of the diagram Δ_m in the above proof is exponential in m .*

5.2 Diagrams short-cutting powers t in Φ_k are *fat*

A key feature of the diagram Δ_m constructed in the previous section is that it contains a path labelled by a large power b^{3^m} of b . In the following proposition we will show that such behaviour is common to *all* \mathcal{P}_k -diagrams Δ that greatly shortcut t^m and we will deduce that all such Δ have large extrinsic (and hence intrinsic) diameter — that is, are *fat*. This means that such *shortcut diagrams* cannot be inserted to significantly reduce the intrinsic diameter of a van Kampen diagram.

For a word w and letter t , we denote the exponent sum of letters t in w by $h_t(w)$.

Proposition 5.4 *There exists a constant C , depending only on k , with the following property: if u and w are words such that $u = w$ in \mathcal{P}_k and $w \in \{t, t^{-1}\}^*$, and Δ is a \mathcal{P}_k -van Kampen diagram for wu^{-1} , homeomorphic to a 2-disc, then*

$$\text{IDiam}(\Delta) \geq \text{EDiam}_{\mathcal{P}_k}(\Delta) \geq \frac{C|h_t(w)|}{1 + \ell(u)}.$$

Proof. We may assume, without loss of generality, that $h_t(w) \geq 0$. Let γ_w and γ_u be the edge-paths in $\partial\Delta$ labelled w and u , respectively, such that the anticlockwise boundary circuit $\partial\Delta$ is γ_w followed by γ_u^{-1} . Let \star and \star' be the initial and terminal vertices of γ_w , respectively — see the left diagram of Figure 6.

The essential idea in this proof is straightforward: the $s_k t$ -corridors (defined in Section 4) emanating from γ_w stack up with the length of the corridors growing to roughly $3^{h_t(w)}$ at the top of the stack; taking the logarithm of this length gives a lower bound on extrinsic diameter of Δ . However, fleshing this idea out into a rigorous proof requires considerable care.

The first complication we face is that $s_k t$ -corridors need not run right across Δ , but can terminate at an edge in the side of an f - or g -corridor or ring. As every letter in w is $t^{\pm 1}$, the edges of $\partial\Delta$ connected in pairs by f - and g -corridors all lie in γ_u . Thus γ_w lies in a single connected component Δ_0 of the planar complex obtained by deleting from Δ the interiors and ends of all the f - and g -corridors. (In the left diagram of Figure 6, Δ_0 is shaded.) Note that $\text{EDiam}_{\mathcal{P}_k}(\Delta) \geq \text{EDiam}_{\mathcal{P}_k}(\Delta_0)$. Thus it suffices to prove that $\text{EDiam}_{\mathcal{P}_k}(\Delta) \geq C|h_t(w)|/(1 + \ell(u))$. We will do so by considering a further \mathcal{P}_k -van Kampen diagram Δ_1 obtained from Δ_0 by removing the interiors of a collection of 2-disc subdiagrams D_i and gluing in replacement diagrams E_i .




Figure 6: Corridors in the van Kampen diagrams Δ and Δ_1 .

Let γ_{u_0} be the edge-path in $\partial\Delta_0$ from \star to \star' such that γ_w followed by $\gamma_{u_0}^{-1}$ is the anticlockwise boundary circuit of Δ_0 .

Suppose A_i is an f - or g -ring in Δ_0 that is not enclosed by another f - or g -ring. The outer boundary circuit γ_i of A_i is labelled by a word w_i in $\{s_1^{\pm 1}, \dots, s_k^{\pm 1}\}^*$. The retraction $\psi : \Phi_k \rightarrow \Theta_k$ of Lemma 3.2 maps w_i to itself and satisfies the hypotheses of Lemma 8.2. So ψ induces a distance decreasing singular combinatorial map from D_i , the diagram enclosed by γ_i , onto an \mathcal{O}_k -van Kampen diagram. Therefore, without increasing extrinsic diameter, we may assume D_i to be an \mathcal{O}_k -diagram. By Lemma 3.3, w_i is null-homotopic in the subpresentation $\langle s_1, \dots, s_k \mid \forall i \neq l, [s_i, s_l] \rangle$ of \mathcal{P}_k .

Let E_i be a topological 2-disc van Kampen diagram for w_i in $\langle s_1, \dots, s_k \mid \forall i \neq j, [s_i, s_j] \rangle$.

We will show that there is an $s_k t$ -corridor in Δ_1 with one side labelled by a word w_T such that $|h_b(w_T)|$ is large. Recall that each $s_k t$ -corridor in a \mathcal{P}_k -van Kampen diagram connects two edges that are either in the boundary of the diagram or the boundary of a 2-cell with label $f^{-1}s_{k-1}fs_k^{-1}s_{k-1}^{-1}$.

But every such 2-cell is part of an f -corridor, and as Δ_1 contains no f - or g -corridor, each of its $s_k t$ -corridors C connects two edges in $\partial\Delta_1$. Each t -edge in γ_w is part of an $s_k t$ -corridor that must either return to γ_w or end on γ_{u_0} . An $s_k t$ -corridor of the former type connects two oppositely oriented t -edges in γ_w . Consider Δ_1 with γ_{u_0} and γ_w running upwards with γ_{u_0} on the left and γ_w on the right (as in Figure 6). So $s_k t$ -corridors of the latter type are *horizontal*, stacked one above another.

Define a horizontal corridor to be an *up*- or *down*-corridor according to whether the edge it meets on γ_w is oriented upwards or downwards. And call an up-corridor a *last-up*-corridor when it meets an edge on γ_w that is labelled by the final letter of a prefix w_1 of w with the property that $h_t(w_2) \geq h_t(w_1)$ for all prefixes w_2 of w with $\ell(w_2) \geq \ell(w_1)$. Note that there are exactly $h_t(w)$ last-up-corridors in Δ_1 . In the right diagram of Figure 6 the up- and down-corridors are shaded and the last-up-corridors are darker. That figure also depicts the scenario addressed in the following lemma.

Lemma 5.5 *Suppose that C_1 and C_2 are two horizontal corridors running from γ_w to γ_{u_0} in Δ_1 , that C_1 is below C_2 , and that there are no horizontal corridors between C_1 and C_2 . Assume that the subarc γ_{u_1} of γ_{u_0} connecting, but not including, the two edges where C_1 and C_2 meet γ_{u_0} is labelled by a word $u_1 \in \{s_1^{\pm 1}, \dots, s_k^{\pm 1}\}^*$. Let w_1 and w_2 be the words read right to left along the top of C_1 and along the bottom of C_2 , respectively. Then we have equality of the exponential sums $h_b(w_1) = h_b(w_2)$.*

Proof of Lemma 5.5. Let γ_{w_0} be the subarc of γ_w connecting, but not including, the two edges where C_1 and C_2 meet γ_w . Let w_0 be the subword of w that we read along γ_{w_0} .

Edges labelled by $s_k^{\pm 1}$ in γ_{u_1} and by $t^{\pm 1}$ in γ_{w_0} are the start of $s_k t$ -corridors that must return to γ_{u_1} or γ_{w_0} , respectively, because there are no horizontal corridors between C_1 and C_2 . So $h_{s_k}(u_1) = h_t(w_0) = 0$ and w_0 freely reduces to the empty word. Moreover in \mathcal{P}_k we find that $u_1 = u_2$ where u_2 is obtained from u_1 by removing all occurrences $s_k^{\pm 1}$. Also, $w_1 u_2 w_2^{-1}$ is null-homotopic.

Now $w_1, w_2 \in \{b^{\pm 1}, s_1^{\pm 1}, \dots, s_{k-1}^{\pm 1}\}^*$ as they are the labels of the sides of $s_k t$ -corridors and $u_2 \in \{s_1^{\pm 1}, \dots, s_{k-1}^{\pm 1}\}^*$. So by Lemma 3.4, $w_1 u_2 w_2^{-1}$ is null-homotopic in

$$\langle b, s_1, \dots, s_{k-1} \mid \forall i \neq j, [s_i, s_j] \rangle \cong \mathbb{Z} * \mathbb{Z}^{k-1}.$$

As $b^{\pm 1}$ does not occur in u_2 , we deduce that $h_b(w_1) = h_b(w_2)$. ■

Returning to our proof of Proposition 5.4, we note that edges that are part of γ_{u_0} but not of γ_u are labelled by letters in $\{s_1^{\pm 1}, \dots, s_k^{\pm 1}\}$ because they are in the sides of f - or g -corridors. So Lemma 5.5 applies to adjacent horizontal corridors that meet γ_{u_0} at edges connected by an arc of γ_{u_0} that does not include edges from γ_u . By the pigeonhole principle, there must be a stack S of $H := h_t(w)/(1 + \ell(u))$ last-up-corridors that all meet a fixed subarc of γ_{u_0} that includes no edges from γ_u . (The stack may also include horizontal corridors that are not last-up-corridors.)

If C is a horizontal s_k t -corridor and c_1 and c_2 are the words along the top and bottom sides of C , then $|h_b(c_1)| = 3|h_b(c_2)|$ when C is an up-corridor, and $3|h_b(c_1)| = |h_b(c_2)|$ when C is a down-corridor. Let $K \neq 0$ be the exponent sum of the letters b in the word along the bottom edge of the lowest last-up-corridor in S . Note that K is non-zero because the s_k t -corridor's two ends are labelled differently. The number of up-corridors minus the number of down corridors in S is H and so by Lemma 5.5, defining w_T to be the word we read along the edge-path γ_{w_T} along the top of the highest corridor in S , we find $h_b(w_T) = K 3^H$.

Suppose w' and w'' are prefixes of w_T . If $w' = w''$ in \mathcal{P}_k then $w' = w''$ in $\langle b, s_1, \dots, s_{k-1} \mid \forall i \neq j, [s_i, s_j] \rangle$ by Lemma 3.4, and so $h_b(w') = h_b(w'')$. So if $h_b(w') \neq h_b(w'')$, then the vertices v', v'' of γ_{w_T} reached after reading w' and w'' map to different vertices in $\text{Cay}^1(\mathcal{P}_k)$. So, because $h_a(w_T) = K 3^H$ and $|K| \geq 1$, the number of vertices in the image of Δ_1 in $\text{Cay}^1(\mathcal{P}_k)$ is at least 3^H . But we must say more: there are 3^H prefixes w' of w_T that all end $b^{\pm 1}$ and for which $h_b(w')$ are all different. The vertices v' at the end of the arc labelled by these w' are not in the interior of any of the diagrams E_i defined prior to Lemma 5.5 (as E_i contain no b -edges) — so the number of vertices in the image of $\Delta_0 \setminus (\bigcup \text{Int} D_i) = \Delta_1 \setminus (\bigcup \text{Int} E_i)$ in $\text{Cay}^1(\mathcal{P}_k)$ is at least 3^H .

The number of vertices in a closed ball $B(r)$ of radius r in $\text{Cay}^1(\mathcal{S}_{k,m})$ is at most c^r for some constant c depending only on the valence of the vertices in $\text{Cay}^1(\mathcal{S}_{k,m})$, and hence on the number of defining generators in $\mathcal{S}_{k,m}$. Thus

$$\text{EDiam}_{\mathcal{P}_k}(\Delta) \geq \text{EDiam}_{\mathcal{P}_k}(\Delta_0) \geq \log_c 3^H = \frac{h_t(w)}{1 + \ell(u)} \log_c 3.$$

■

5.3 An upper bound for the distortion of $\langle t \rangle$ in Φ_k

One conclusion of Proposition 5.1 was that it is possible to express $t^m \in \Phi_k$ as a word of length $\preceq m^{1/k}$. The main result of this section, Proposition 5.7, is that no greater distortion of $\langle t \rangle$ is possible; this will be crucial in Section 7 when we come to analyse $\mathcal{S}_{k,m}$ -van Kampen diagrams by breaking them down into \mathcal{P}_k - and \mathcal{Q}_m -subdiagrams meeting along arcs labelled by powers of t .

Lemma 5.6 *If u_0 is a word in the generators of Θ_k that equals s_k^m in the group, then*

$$|m| \leq \sum_{i=1}^k \binom{c+k-1-i}{c-1} \ell_{s_i}(u_0),$$

where $c := (\ell_f(u_0) + \ell_g(u_0))/2$.

Proof. If we represent elements of $\langle s_1, \dots, s_k \mid \forall i, j, [s_i, s_j] \rangle \cong \mathbb{Z}^k$ as column-vectors, then the actions of f and f^{-1} by conjugation are given by left multiplication by the $k \times k$ matrices

$$\begin{pmatrix} 1 & & & & \\ 1 & 1 & & & \\ & 1 & \ddots & & \\ & & \ddots & 1 & \\ & & & 1 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & & & & \\ -1 & 1 & & & \\ & 1 & -1 & \ddots & \\ -1 & 1 & \ddots & 1 & \\ \vdots & \vdots & \ddots & -1 & 1 \end{pmatrix},$$

respectively. Similarly, we can give matrices for the actions of g and g^{-1} .

We will inductively obtain words u_1, \dots, u_c , all of which equal s_k^m in \mathcal{O}_k , as follows. In every van Kampen diagram for $u_i s_k^{-m}$, the letters $f^{\pm 1}, g^{\pm 1}$ in u_i occur in pairs f, f^{-1} or g, g^{-1} connected by f -corridors and g -corridors. The geometry of these corridors necessitates that u_i have a subword of the form $f^{\mp 1} v f^{\pm 1}$ or $g^{\mp 1} v g^{\pm 1}$, where $v \in \{s_1^{\pm 1}, \dots, s_k^{\pm 1}\}^*$. Obtain u_{i+1} from u_i by replacing every letter s_j ($j = 1, 2, \dots, k$) in v by a word of minimal length in \mathbb{Z}^k that equals $f^{\mp 1} s_j f^{\pm 1}$ or $g^{\mp 1} s_j g^{\pm 1}$, as appropriate. Let A be the $k \times k$ matrix with ones in every entry on and below the diagonal and zeros elsewhere. Comparing A with the four matrices discussed above we see that $\ell_{s_i}(u_{i+1})$ is at most the i -th entry in the column vector $A(\ell_{s_1}(u_0), \dots, \ell_{s_k}(u_0))^{\text{tr}}$.

Using the identity

$$\sum_{n=0}^r \binom{n}{j} = \binom{r+1}{j+1},$$

we calculate that

$$A^c = \begin{pmatrix} \begin{pmatrix} c-1 \\ c-1 \end{pmatrix} & & & & \\ \begin{pmatrix} c \\ c-1 \end{pmatrix} & \begin{pmatrix} c-1 \\ c-1 \end{pmatrix} & & & \\ \vdots & \begin{pmatrix} c \\ c-1 \end{pmatrix} & \ddots & & \\ \begin{pmatrix} c+k-3 \\ c-1 \end{pmatrix} & \vdots & \ddots & \begin{pmatrix} c-1 \\ c-1 \end{pmatrix} & \\ \begin{pmatrix} c+k-2 \\ c-1 \end{pmatrix} & \begin{pmatrix} c+k-3 \\ c-1 \end{pmatrix} & \cdots & \begin{pmatrix} c \\ c-1 \end{pmatrix} & \begin{pmatrix} c-1 \\ c-1 \end{pmatrix} \end{pmatrix}.$$

Now $|m| = \ell_{s_k}(u_c)$ and so is at most the k -th entry in $A^c(\ell_{s_1}(u_0), \dots, \ell_{s_k}(u_0))^{\text{tr}}$. The asserted bound follows. \blacksquare

Recall that $h_t(w)$ denotes the exponent sum of letters $t^{\pm 1}$ in w .

Proposition 5.7 *Suppose w is a word equalling t^m in \mathcal{P}_k . Let $n := \ell(w) - \ell_t(w)$. Then*

$$|m| \leq K n^k + h_t(w),$$

where K is a constant depending only on k .

Proof. The retraction $\psi : \Phi_k \twoheadrightarrow \Theta_k$ of Lemma 3.2 maps w , letter-by-letter, to a word u equalling s_k^m in Θ_k . Each t is mapped by ψ to s_k which is central in Θ_k . So if we define w_0 to be the word obtained from w by removing all letters $t^{\pm 1}$ and u_0 to be the word obtained by applying ψ , letter-by-letter, to w_0 then $u_0 = s_k^{m-h_t(w)}$ in \mathcal{O}_k .

By Lemma 5.6

$$|m - r| \leq \sum_{i=1}^k \binom{n+k-1-i}{n-1} \ell_{s_i}(u_0) \leq \binom{n+k-2}{n-1} n.$$

The asserted bound then follows because, for a suitable constant K depending only on k , one has $\binom{n+k-2}{n-1} n \leq K n^k$ for all $n \in \mathbb{N}$. \blacksquare

5.4 The intrinsic diameter of diagrams for Φ_k

The results in this section culminate with an upper bound on the intrinsic diameter functional of Φ_k . This will be used in Section 7.1, when we establish the upper bound on the extrinsic diameter of $\Psi_{k,m}$.

Lemma 5.8 *If w is a length n word in \mathcal{P}_k representing an element of the subgroup $\mathbb{Z}^k \cong \langle s_1, \dots, s_k \rangle$ or $\mathbb{Z}^k \cong \langle \hat{s}_1, \dots, \hat{s}_k \rangle$ of Φ_k , then there is a word u in $\{s_1^{\pm 1}, \dots, s_k^{\pm 1}\}^*$ or in $\{\hat{s}_1^{\pm 1}, \dots, \hat{s}_k^{\pm 1}\}^*$, respectively, such that $w = u$ in Φ_k and $\ell(s) \leq K n^k$, where K depends only on k .*

Proof. The retracts of Lemma 3.2 mapping Φ_k onto Θ_k and $\hat{\Theta}_k$ mean that it is enough to prove this lemma for Θ_k instead of Φ_k . The result for Θ_k can be proved in the same manner as Lemma 5.6. \blacksquare

Define

$$\begin{aligned} \hat{\mathcal{P}}_k := \langle b, t, s_1, \dots, s_k, \hat{s}_1, \dots, \hat{s}_k \mid & t^{-1}bs_k = b^3, s_k^{-1}bt = b^3, \\ & \hat{s}_k^{-1}b\hat{s}_k = b^3; \ \forall i \neq j, [s_i, s_j] = 1, [\hat{s}_i, \hat{s}_j] = 1 \rangle, \end{aligned}$$

a subpresentation of \mathcal{P}_k . Note that the 2-dimensional portions of every $\hat{\mathcal{P}}_k$ -van Kampen are comprised of intersecting $s_k t$ -, s_i -, and \hat{s}_j -rings and corridors (where $1 \leq i < k$ and $1 \leq j \leq k$).

Lemma 5.9 *If Δ is a minimal area $\hat{\mathcal{P}}_k$ -van Kampen diagram for a word w then*

- (i) *amongst the corridors and rings in Δ , no two cross twice,*
- (ii) *Δ contains no s_i - or \hat{s}_i -rings ($1 \leq i < k$),*
- (iii) *Δ contains no \hat{s}_k -rings,*
- (iv) *Δ contains no $s_k t$ -rings, and*
- (v) *the length of each s_i - and \hat{s}_i -corridor ($1 \leq i < k$) in Δ is less than $\ell(w)/2$.*

Proof. For (i), first suppose that for some $i \neq j$ with $1 \leq i, j < k$, there is an s_i -corridor or s_i -ring C_i that crosses an s_j -corridor or s_j -ring C_j twice, intersecting at two 2-cells e and e' . Let \tilde{C}_i and \tilde{C}_j be portions of C_i and C_j between (but not including) e and e' , as shown in Figure 7. By an *innermost argument* we may assume that no s_j -corridor or s_j -ring intersects \tilde{C}_i twice and that no s_i -corridor or s_i -ring intersects \tilde{C}_j twice. Removing e and e' , relabelling all the s_i -edges in \tilde{C}_i by s_j and all the s_j -edges in \tilde{C}_j by s_i , and then gluing up as shown in Figure 7, would produce a van Kampen diagram for w of lesser area than Δ . This would contradict the minimality of the area of Δ . (Topologically, the effect of the surgery on Δ is to collapse to points arcs running through e and e' between opposite vertices in ∂e and $\partial e'$. This does not spoil planarity because no opposite pair of vertices were identified in Δ .)

Next suppose that for some $1 \leq i < k$, an $s_k t$ -corridor or ring C crosses an s_i -corridor or ring C_i twice at 2-cells e and e' . Let \tilde{C} and \tilde{C}_i be portions of C and C_i between (but not including) e and e' . We may assume that




Figure 7: Removing twice-crossing corridors.

no $s_k t$ -corridor or $s_k t$ -ring intersects \tilde{C}_i twice and that no s_j -corridor or s_j -ring ($1 \leq j < k$) intersects \tilde{C} twice. Let R be the subdiagram between C and C_i . For reasons we are about to explain we may assume \tilde{C} includes no 2-cells labelled by $t^{-1}bs_kb^{-3}$, (equivalently, ∂R contains no b -edge). This leads to a contradiction, as above.

No 2-cell of R labelled $t^{-1}bs_kb^{-3}$ lies in an s_k, t -corridor, as such a corridor would intersect \tilde{C}_i twice. Also R includes no 2-cell that is part of an s_k, t -ring — such a ring could intersect \tilde{C}_i and so would have to be entirely in R ; but then, by another *innermost argument*, this s_k, t -ring would only contain 2-cells labelled by $t^{-1}bs_kb^{-3}$ as otherwise it would twice intersect an s_j -corridor for some $1 \leq j < k$. So R includes no 2-cells labelled by $t^{-1}bs_kb^{-3}$, and for similar reasons, no 2-cells labelled by $\hat{s}_k^{-1}b\hat{s}_kb^{-3}$. So all b -edges in ∂R are identified in pairs and are amongst the 1-dimensional portions of R . It follows that there are no b -edges in R , or there is a b -edge in R that is connected to the rest of R at only one vertex, or there is a 2-disc component of $R \setminus \partial R$ that does not adjoin C_i .

The second case is impossible because it would imply that Δ was not reduced. The third case cannot happen because for some $1 \leq j < k$, an s_j -corridor would have to cross C , travel through this 2-disc portion and then cross C again. Thus R has no b -edges.

The same arguments tell us that for all $1 \leq i, j \leq k$, no \hat{s}_i - and \hat{s}_j -corridors or rings in Δ can cross twice. No other combination of rings and corridors can cross even once.

Now (ii) follows immediately from (i), as do (iii) and (iv) in the cases where the rings include 2-cells other than those labelled by $t^{-1}bs_kb^{-3}$ or $\hat{s}_k^{-1}b\hat{s}_kb^{-3}$. In the remaining cases the outer boundary of the ring is labelled by a freely reducible word in $\{b^{\pm 1}\}^*$, which contradicts the minimality of the area of Δ .

For (v), suppose C is an s_i - or \hat{s}_i -corridor. The boundary circuit of Δ is comprised of the two ends of C and two edge-paths, one of which, call it α , must have length less than $\ell(w)/2$. By (i)–(iv) a different corridor connects each on a side of C to α . So C has length less than $\ell(w)/2$. ■

Lemma 5.10 *There exists a constant K such that $\text{IDiam}_{\hat{\mathcal{P}}_k}(n) \leq Kn$ for all $n \in \mathbb{N}$.*

Proof. Suppose Δ is a minimal area $\hat{\mathcal{P}}_k$ -van Kampen diagram for a word w . All $s_k t$ -corridors in Δ are embedded: if a portion (of non-zero length) of the path along the side of an $s_k t$ -corridor formed an edge-loop then that would have to enclose a zero area subdiagram (by the results of Lemma 5.9) and therefore be labelled by a non-reduced path — but then Δ would not be a minimal area diagram as there would be an inverse pair of 2-cells on the corridor. So from any given vertex in Δ , travelling across $s_k t$ -corridors at most $n/2$ times, we meet either an s_i -corridor ($1 \leq i < k$), or an \hat{s}_i -corridor ($1 \leq i < k$), or $\partial\Delta$. In the former two cases one can follow a path of length at most $n/4$ along a side of the s_i - or \hat{s}_i -corridor to $\partial\Delta$. From any point on $\partial\Delta$ one can reach the base vertex by following the boundary circuit. ■

Proposition 5.11 $\text{IDiam}_{\mathcal{P}_k}(n) = O(n^k)$.

Proof. Suppose w is a null-homotopic word in \mathcal{P}_k and Δ is a van Kampen diagram for w . We will construct a new van Kampen diagram Δ' for w that satisfies the claimed bound on intrinsic diameter. We begin the construction of Δ' by taking an edge-circuit ξ of length $\ell(w)$ in the plane to serve as $\partial\Delta'$. We direct and label the edges of ξ so that one reads w around it.

The occurrences of f and f^{-1} in w are paired so that the corresponding edges of $\partial\Delta$ are joined by corridors in Δ . For each such pair, there must be a subword $f^{-1}w_0f$ in some cyclic conjugate of w or w^{-1} such that both w_0 and $f^{-1}w_0f$ represent elements of the subgroup $\mathbb{Z}^k = \langle s_1, \dots, s_k \rangle$. Join the pairs of f - and f^{-1} -edges in w by f -corridors running through the interior of ξ . It follows from Lemma 5.8 that the lengths of both sides of these f -corridors are at most a constant times $(2 + \ell(w_0))^k$.

Likewise, insert corridors into the interior of ξ joining pairs of g - and g^{-1} -edges, of \hat{f} - and \hat{f}^{-1} -edges, and of \hat{g} and \hat{g}^{-1} -edges. There is no obstruction to planarity in the 2-complex because we are mimicking the layout of corridors in Δ .

To complete the construction of Δ' , we shall fill the 2-disc holes inside ξ . The boundary circuit η of each hole is made up of the sides of corridors

C_1, \dots, C_r and a number of disjoint subarcs of ξ . The length of η is at most n^k , up to a multiplicative constant, because these disjoint subarcs contribute at most n and the lengths of the sides of C_1, \dots, C_r each contribute at most a constant times n_i^k , where $\sum_{i=1}^r n_i \leq n$. By Lemma 5.10, these circuits can be filled by $\hat{\mathcal{P}}_k$ -van Kampen diagrams with intrinsic diameter at most a constant times n^k . And, as the corridors have length at most a constant times n^k , we deduce that $\text{IDiam}_{\mathcal{P}_k}(\Delta') = O(n^k)$, as required. \blacksquare

Remark 5.12 *Let us consider why the Dehn function of \mathcal{P}_k is at most $n \mapsto K^{n^k}$ for some constant $K > 0$.*

Lemma 5.9 implies that the Dehn function of $\hat{\mathcal{P}}_k$ is at most K_1^n for some $K_1 > 0$: the total contribution of the 2-cells labelled by $[s_i, s_j]$ or $[\hat{s}_i, \hat{s}_j]$ for some $i \neq j$ is at most n^2 ; removing all s_i - and \hat{s}_i -corridors (for all $1 \leq i < k$) leaves components with linear length boundary circuits filled by minimal area van Kampen diagrams over

$$\langle b, t, s_k, \hat{s}_k \mid t^{-1}bs_k = b^3, s_k^{-1}bt = b^3, \hat{s}_k^{-1}b\hat{s}_k = b^3 \rangle,$$

and a standard corridors argument shows the Dehn function of this subpresentation is bounded above by an exponential function. The construction of diagrams in Proposition 5.11 then establishes that the Dehn function of \mathcal{P}_k is at most $n \mapsto K^{n^k}$ for some $K > 0$.

6 Diameter in Γ_m

In this section we establish an upper bound on the intrinsic, and hence extrinsic, diameter of null-homotopic words w in the presentation \mathcal{Q}_m for Γ_m . Also we show how the *shortcuts* of Proposition 5.1 lead to an improved bound on extrinsic diameter when we regard w as a word in the presentation $\mathcal{S}_{k,m}$ for $\Psi_{k,m}$.

Proposition 6.1 *Fix integers $n > 0$ and $1 < k < m$. Suppose w is a null-homotopic word in the presentation \mathcal{Q}_m for Γ_m , that $\ell(w) - \ell_t(w) \leq n$, and that $\ell_t(w) = O(n^k)$. Then there is a \mathcal{Q}_m -van Kampen diagram Δ for w with $\text{IDiam}_{\mathcal{Q}_m}(\Delta) = O(n^{m+1})$. Moreover, as an $\mathcal{S}_{k,m}$ -van Kampen diagram,*

$$\text{EDiam}_{\mathcal{S}_{k,m}}(\Delta) = O\left(n^{\max\{1+\frac{m}{k}, k\}}\right).$$

To prove this result we will need some purchase on the geometry of τ - and σ -corridors. This is provided by the following two lemmas.

Lemma 6.2 *Suppose Δ is a minimal area \mathcal{Q}_m -van Kampen diagram Δ for a null-homotopic word w and that \mathcal{C} is a τ -corridor in Δ . Let w_0 be the word along the sides of \mathcal{C} . If w_1 is a subword of w_0 and $w_1 \neq t^{\pm 1}$ then w_1 cannot represent the same group element as a non-zero power of t . In particular, \mathcal{C} is embedded; that is, the sides of \mathcal{C} are simple paths in Δ .*

Proof. As Δ is of minimal area, w_0 is freely reduced as a word in

$$\{a_1^{\pm 1}, \dots, a_{m-1}^{\pm 1}, (a_m t)^{\pm 1}, T^{\pm 1}\}^*.$$

Killing t and τ retracts \mathcal{Q}_m onto

$$\mathcal{U}_m := \langle a_1, \dots, a_m, \sigma, T \mid \sigma^{-1} a_m \sigma = a_m, \forall i < m, \sigma^{-1} a_i \sigma = a_i a_{i+1} \rangle,$$

in which a_1, \dots, a_m, T generate a free subgroup. This retraction sends w_1 to a word in $\{a_1^{\pm 1}, \dots, a_{m-1}^{\pm 1}, a_m^{\pm 1}, T^{\pm 1}\}^*$ that is freely reduced and therefore non-trivial. The result follows. \blacksquare

The analogue of this result for σ -corridors is more complex.

Lemma 6.3 *Suppose Δ is a minimal area \mathcal{Q}_m -van Kampen diagram Δ for a null-homotopic word w and that \mathcal{C} is a σ -corridor in Δ . Then \mathcal{C} is a topological 2-disc subdiagram of Δ with boundary label $\sigma^{-1} w_0 \sigma w_1^{-1}$, where w_0 and w_1 are freely reduced words in $\{a_1^{\pm 1}, \dots, a_m^{\pm 1}\}^*$. Moreover, defining ρ_0 and ρ_1 to be the arcs of $\partial \mathcal{C}$ along which one reads w_0 and w_1 , there exists $K > 0$, depending only on m , such that from any point on ρ_i ($i = 0, 1$) there is a path in $\mathcal{C}^{(1)}$ of length at most K to $\rho_{|1-i|}$. Also, no subword of w_0 or w_1 represents the same group element as a non-zero power of t .*

Proof. That w_0 is reduced, that ρ_0 is a simple path in $\Delta^{(1)}$, and that no subword of w_0 or w_1 represents the same group element as a non-zero power of t , are all proved as in Lemma 6.2. And \mathcal{C} is a topological 2-disc by a similar argument involving the retraction \mathcal{U}_m .

So w_1 is obtained from w_0 by replacing each $a_i^{\pm 1}$ by $(a_i a_{i+1})^{\pm 1}$ for all $i < m$, and then freely reducing. The constant K exists by the *Bounded Cancellation Lemma* of [6]. \blacksquare

Proof of Proposition 6.1. Let Δ be a \mathcal{Q}_m -van Kampen diagram for w that is of minimal area. Note that Δ contains no t -, τ - or σ -rings. Removing the τ - and σ -corridors from Δ leaves a disjoint union of subdiagrams Δ_i over

$$\langle a_1, \dots, a_m, t, T \mid \forall j, [t, a_j] = 1; [t, T] = 1 \rangle.$$

Since there are no t -rings in Δ_i and the words along the sides of the t -corridors are reduced, from any vertex in Δ_i one can reach $\partial\Delta_i$ by following at most $\ell(\partial\Delta_i)/2$ t -edges across successive t -corridors.

We claim that if \mathcal{C} is a τ - or σ -corridor in Δ and u is the word along a side of \mathcal{C} then $\ell(u) \leq O(\min\{n_1^m, n_2^m\})$, where n_1 and n_2 are the lengths of the two arcs that together comprise $\partial\Delta$ and have the same end points as one side of \mathcal{C} . This is because killing t and τ retracts \mathcal{Q}_m onto \mathcal{U}_m , and the calculation used in the proof of Proposition 5.6 shows that the conjugation action of the stable letter σ is an automorphism of polynomial growth of degree $m-1$. And on page 452 of [4] the first author shows that the growth of the inverse automorphism is also polynomial of degree $m-1$.

It follows that $\ell(\partial\Delta_i) \leq O(n^m)$ because the boundary of Δ_i consists of portions of $\partial\Delta$, and the sides of σ - and τ -corridors.

We are now ready to estimate $\text{IDiam}_{\mathcal{Q}_m}(\Delta)$. Suppose v is a vertex of Δ . In the light of Lemmas 6.2 and 6.3, there is no loss of generality in assuming v is not in the interior of a τ - or σ -corridor. Move from v to $\partial\Delta$ by successively crossing σ - and τ -corridors as follows. When located on a σ - or τ -corridor that has not just been crossed, follow at most K (the constant of Lemma 6.3) edges to cross to the other side; otherwise follow a maximal length embedded path of t -edges across some Δ_i . Finally follow $\partial\Delta$ to the base vertex \star . Let ρ denote the resulting edge-path from v to \star .

It must be verified that we can indeed reach $\partial\Delta$ by moving in the manner described and that $\ell(\rho) = O(n^m)$. First note that there are no embedded edge-loops in Δ labelled by words in $\{t^{\pm 1}\}^*$ because the interior of such a loop could be removed and the hole glued up (as \mathcal{Q}_m retracts onto $\langle t \rangle$), reducing the area of Δ . Next observe that ρ does not cross the same τ - or σ -corridor twice, for otherwise there would have to be an *innermost* τ - or σ -corridor that ρ crosses twice, contradicting either Lemma 6.2 or Lemma 6.3.

So crossing the σ - and τ -corridors, of which there are at most $n/2$, contributes at most $Kn/2$ to the length of ρ . We have already argued that each section of ρ between an adjacent pair of σ - or τ -corridors has length $O(n^m)$. So these sections together contribute at most $(n/2)O(n^m) = O(n^{m+1})$ to the total length. The final section of ρ is part of $\partial\Delta$ and so has length at most $O(n^k) < O(n^m)$. The total is $O(n^{m+1})$ as required.

For the bound on $\text{EDiam}_{\mathcal{S}_{k,m}}(\Delta)$, we note that in the word metric (i.e. measured in the Cayley graph) the distance from the initial to the terminal vertex of each arc of ρ whose edges are all labelled by $t^{\pm 1}$, is $O(n^{\frac{m}{k}})$ by Proposition 5.1. Therefore $\text{EDiam}_{\mathcal{S}_{k,m}}(\Delta) = O(n)O(n^{\frac{m}{k}}) = O(n^{1+\frac{m}{k}})$, as required. ■

Remark 6.4 *The proof of Proposition 6.1 also shows that the Dehn function of Γ_m is $O(n^{2m+1})$ (cf. Theorem 3.4 of [2]).*

7 Proof of the main theorem

In this section we prove Theorem 7.3 (modulo some technicalities postponed to Section 8). Theorem 1.3 then follows because k and m can be chosen so that the ratio of the intrinsic and extrinsic diameter filling functionals of $\Psi_{k,m}$ grows faster than any prescribed polynomial.

Before stating the theorem we recall the definition of an *alternating product expression* in an amalgam and a well known lemma — see Lemma 6.4 in Section III.F of [5] or Section 5.2 of [9] .

Definition 7.1 *Let $\Gamma = A *_C B$ be an amalgam of groups A and B along a common subgroup C . Suppose \mathcal{A} and \mathcal{B} are generating sets for A and B , respectively. An alternating product expression for $w \in (\mathcal{A}^{\pm 1} \cup \mathcal{B}^{\pm 1})^*$ is a cyclic conjugate $u_1 v_1 \dots u_p v_p$ of w in which for all i we have $u_i \in (\mathcal{A}^{\pm 1})^*$, $v_i \in (\mathcal{B}^{\pm 1})^*$, and if $p \neq 1$ then neither u_i nor v_i is the empty word.*

Lemma 7.2 *In the notation of Definition 7.1, if $w \in (\mathcal{A}^{\pm 1} \cup \mathcal{B}^{\pm 1})^*$ represents 1 in Γ then in any alternating product expression $u_1 v_1 \dots u_p v_p$ for w , some u_i or v_i represents an element of C .*

Theorem 7.3 *For all integers $m > k > 1$, the extrinsic and intrinsic diameter filling functions of the group $\Psi_{k,m}$ presented by $\mathcal{S}_{k,m}$, satisfy*

$$\begin{aligned} \text{EDiam}_{\mathcal{S}_{k,m}}(n) &= O\left(n^{\max\{1+\frac{m}{k}, k\}}\right), \\ n^{m/3} &= O(\text{IDiam}_{\mathcal{S}_{k,m}}(n)). \end{aligned}$$

7.1 Proof of the upper bound on $\text{EDiam}_{\mathcal{S}_{k,m}}(n)$

Suppose w is a null-homotopic word in $\mathcal{S}_{k,m}$. Let $u_1 v_1 \dots u_p v_p$ be an *alternating product decomposition* for w where each u_i and v_i is a word on the generators of \mathcal{P}_k and \mathcal{Q}_m respectively. Take a planar edge-circuit η around which, after directing and labelling the edges, one reads w . Decompose η into arcs along which we read the u_i and v_i , and call the vertices at which these meet *alternation vertices*. Say that an edge-path is a t -arc when it is made up of t -edges orientated the same way.

Repeated appeals to Lemma 7.2, the first of which tells us that some u_i or v_i represents a word in $\{t^{\pm 1}\}^*$, allow us to deduce that t -arcs with the

following properties can be inserted into the interior of η : each t -arc connects two distinct alternation vertices of η ; any two t -arcs are disjoint; and the t -arcs partition the interior of η into topological 2-disc regions, whose boundary loops are labelled by words that are null-homotopic either in \mathcal{P}_k or in \mathcal{Q}_m . Call these bounding loops \mathcal{P}_k -loops and \mathcal{Q}_m -loops, respectively.

Lemma 7.4 *For each t -arc α , let n_α be the length of the shorter of the two subarcs of η that share their two end vertices with α . Then the length $\ell(\alpha)$ of α satisfies*

$$\ell(\alpha) \leq K n_\alpha^k$$

where $K \geq 1$ is the constant of Proposition 5.7, which depends only on k .




Figure 8: Partitioning η with t -arcs.

Proof of Lemma 7.4. We induct on n_α . The base case, $n_\alpha = 0$ holds vacuously. For the inductive step, of the two regions adjoining α , choose \mathcal{R} to be that which is in the interior of the disc bounded by α and the length n_α subarc of η — see Figure 8. The boundary of \mathcal{R} is comprised of α , disjoint subarcs of η , and some t -arcs that we call $\alpha_1, \dots, \alpha_r$. Define n_0 to be the total length of these subarcs of η . Then

$$n_\alpha = n_0 + \sum_{i=1}^r n_{\alpha_i} \quad (2)$$

By induction hypothesis, $\ell(\alpha_i) \leq K n_{\alpha_i}^k$ for all i . We address two cases.

Case: $\partial\mathcal{R}$ is a \mathcal{Q}_m -loop. The exponent sum of the occurrences of t in any null-homotopic word in \mathcal{Q}_m is zero because \mathcal{Q}_m retracts onto $\langle t \rangle \cong \mathbb{Z}$ via the map that kills all generators other than t . So

$$\ell(\alpha) \leq n_0 + \sum_{i=1}^r K n_{\alpha_i}^k \leq K n_{\alpha}^k,$$

with the second inequality following from (2).

Case: $\partial\mathcal{R}$ is a \mathcal{P}_k -loop. Proposition 5.7 and (2) give

$$\ell(\alpha) \leq K n_0^k + \sum_{i=1}^r K n_{\alpha_i}^k \leq K n_{\alpha}^k.$$

This completes the proof of the lemma. ■

Next we fill the \mathcal{P}_k - and \mathcal{Q}_m -loops inside η to produce a van Kampen diagram that satisfies the asserted bound on extrinsic diameter. Let $n := \ell(w)$. By Lemma 7.4, the length of each of their boundary circuits is at most $O(n^k)$, with the contributions from portions not on η coming entirely from t -arcs. First we fill the \mathcal{Q}_m -loops as per Proposition 6.1, with diagrams each of which meets η and has extrinsic diameter $O\left(n^{\max\{1+\frac{m}{k}, k\}}\right)$ as an $\mathcal{S}_{k,m}$ -van Kampen diagram.

Next we glue a *shortcut diagram* along each of the t -arcs as per Proposition 5.1. These diagrams have intrinsic diameter $O(n^k)$, measured from base vertices on η . Finally, fill the remaining \mathcal{P}_k -loops — all have length $O(n)$ and meet η , and by Proposition 5.11 they can be filled by van Kampen diagrams of intrinsic, and hence extrinsic, diameter $O(n^k)$.

The result is an $\mathcal{S}_{k,m}$ -van Kampen diagram that admits the asserted bound on extrinsic diameter. ■

Remark 7.5 *This proof, together with Remarks 5.3, 5.12 and 6.4 establish that the Dehn function of $\Psi_{k,m}$ is at most $n \mapsto C n^k$ for some constant C .*

7.2 Proof of the lower bound on $\text{IDiam}_{\mathcal{S}_{k,m}}(n)$.

Define

$$w_n := [\tau, (\sigma^{-n} a_1^n \sigma^n) T(\sigma^{-n} a_1^{-n} \sigma^n)].$$

Let $u = u(a_1, \dots, a_m)$ be the (positive) word such that $\sigma^{-n} a_1^n \sigma^n = u$ in \mathcal{Q}_m . Define $u_t := u(a_1, \dots, a_{m-1}, a_m t)$. For some $q \in \mathbb{Z}$ we have $u t^q = u_t$. This plays a key role in the construction of the \mathcal{Q}_m -van Kampen for w_n , an outline for which is shown in Figure 9.




Figure 9: A \mathcal{Q}_m -van Kampen for w_n .

Lemma 7.6 *The word w_n has extrinsic diameter at least $Cn^m - C$ in \mathcal{Q}_m , where C is a constant depending only on m .*

Proof of Lemma 7.6. Our approach builds on the proof by the first author of Theorem 3.4 in [2]. Suppose $\pi : \Delta \rightarrow \text{Cay}^2(\mathcal{Q}_m)$ is a \mathcal{Q}_m -van Kampen diagram for w_n . First we find an edge-path ρ in Δ , along which one reads a word in which the exponent sum of the letters t is at least $Cn^m - C$ for some constant $C = C(m)$. A T -corridor connects the two letters T in w_n , and along each side of this corridor we read a word u in $\{t^{\pm 1}, \tau^{\pm 1}\}^*$ that equals $(\sigma^{-n}a_1^{-n}\sigma^n)\tau^{-1}(\sigma^{-n}a_1^n\sigma^n)$ in \mathcal{Q}_m . A τ -corridor joins the τ^{-1} in this latter word to some τ^{-1} in u . Let u_0 be the prefix of u such that the letter immediately following u_0 is this τ^{-1} , and then let ρ be the edge-path along the side of the τ -corridor running from the vertex at the end of $\sigma^{-n}a_1^{-n}\sigma^n$ to the vertex at the end of u_0 . Let $v = v(a_1, \dots, a_{m-1}, a_mt)$ be the word one reads along ρ . Then $u_0 = (\sigma^{-n}a_1^{-n}\sigma^n)v$ in \mathcal{Q}_m . Killing T , t and τ , retracts \mathcal{Q}_m onto

$$\langle a_1, \dots, a_m, s \mid \forall i \neq j, [a_i, a_j] = 1; \sigma^{-1}a_m\sigma = a_m, \forall i < m, \sigma^{-1}a_i\sigma = a_ia_{i+1} \rangle,$$

in which $\bar{v} = \sigma^{-n}a_1^n\sigma^n$, where $\bar{v} := v(a_1, \dots, a_{m-1}, a_m)$. By Lemma 3.5, the exponent sum of a_m in \bar{v} , and hence of t in v , is $n\binom{n}{m-1}$, which is at least $Cn^m - C$ for some constant C depending only on m .

Killing all generators other than t defines a retraction ϕ of Γ_m onto $\langle t \rangle \cong \mathbb{Z}$. The 0-skeleton of $\text{Cay}^2(\mathcal{Q}_m)$ is Γ_m and the existence of ρ shows that image of $\phi \circ \pi : \Delta \rightarrow \mathbb{Z}$ has diameter at least Cn^m , since the retraction

ϕ from the 0-skeleton of $\text{Cay}^2(\mathcal{Q}_m)$ to \mathbb{Z} decreases distance. This completes the proof of the lemma. \blacksquare

Let Δ be a minimal intrinsic diameter $\mathcal{S}_{k,m}$ -van Kampen diagram for w_n . On account of the retraction ϕ , we may assume Δ to have the properties described in Proposition 8.4 (2). That is, the \mathcal{P}_k -cells within Δ comprise a subcomplex whose connected components are all simply connected unions of topological disc subcomplexes any two of which meet at no more than one vertex. Refer to these topological disc subcomplexes as \mathcal{P}_k -islands. Around the boundary of each \mathcal{P}_k -island we read a word in $\{t^{\pm 1}\}^*$ that freely reduces to the empty word because t has infinite order in $\Phi_{k,m}$.

We obtain a \mathcal{Q}_m -van Kampen diagram $\overline{\Delta}$ for w_n from Δ by cutting out the \mathcal{P}_k -islands and then gluing up the attaching cycles by identifying adjacent, oppositely-oriented edges (i.e. successively cancelling pairs tt^{-1} or $t^{-1}t$ in the attaching words). The removal of the \mathcal{P}_k -islands and subsequent gluing is described by a collapsing map $\theta : \Delta \twoheadrightarrow \overline{\Delta}$ that is injective and combinatorial except on the \mathcal{P}_k -islands. This is depicted in Figure 10. Note that, no matter what choice of θ we make, $\overline{\Delta}$ will be planar by Lemma 8.1, since after cutting out and gluing up a number of the \mathcal{P}_k -islands, the boundary circuit of every remaining \mathcal{P}_k -island is a simple loop.

Let T be a maximal geodesic tree in the 1-skeleton of Δ , based at \star . Suppose γ is a geodesic in T from a vertex v not in the interior of a \mathcal{P}_k -island to \star . Define an edge-path $\overline{\gamma}$ in the 1-skeleton of $\overline{\Delta}$ from $\theta(v)$ to $\theta(\star)$ to follow the arcs of $\theta \circ \gamma$ outside the interior of \mathcal{P}_k -islands, and to follow the geodesic path in the tree $\theta(\partial I)$ whenever γ enters a \mathcal{P}_k -island I .

It will be important (in *Case 2* below) that the arcs of $\overline{\gamma}$ defined as geodesics in the images of the boundaries of \mathcal{P}_k -islands follow edge-paths labelled by reduced words in $\{t, t^{-1}\}^*$. That is, they must not traverse a pair of edges labelled by tt^{-1} or $t^{-1}t$. The gluing θ involved choices that, according to the following lemma (illustrated in Figure 10), we can exploit to ensure that the paths $\overline{\gamma}$ satisfy the conditions we require.

Lemma 7.7 *The gluing map $\theta : \Delta \twoheadrightarrow \overline{\Delta}$ can be chosen so as to satisfy the following. Suppose $\gamma_{u,v}$ is an edge-path in T from a vertex u to a vertex v and is an initial segment of some geodesic in T from u to \star . Assume, further, that $\gamma_{u,v}$ lies in some \mathcal{P}_k -island I and meets ∂I at u and v and nowhere else. Define $\delta_{u,v}$ to be the geodesic in the tree $\theta(\partial I)$ from $\theta(u)$ to $\theta(v)$. Then along $\delta_{u,v}$ we read a reduced word in $\{t, t^{-1}\}^*$.*

Proof. Start with any choice of gluing map θ . Suppose that I is a \mathcal{P}_k -island in Δ and that $u, v \in \partial I$ are as per the lemma, but that $\delta_{u,v}$ follows a word in




Figure 10: Collapsing the \mathcal{P}_k -islands.

$\{t, t^{-1}\}^*$ that includes an inverse pair, tt^{-1} or $t^{-1}t$. Then perform a diamond move as illustrated in Figure 11 to remove the inverse pair. This, in effect, amounts to changing the choice of θ .




Figure 11: A diamond move.

Suppose that $\hat{u}, \hat{v} \in \partial I$ are another pair as per the lemma. We claim that when we do the diamond move to remove a pair of edges from $\delta_{u,v}$ the effect, if any, on the word one reads along $\delta_{\hat{u}, \hat{v}}$ is also the removal of an inverse pair. Consider the ways in which $\delta_{\hat{u}, \hat{v}}$ could meet the pair of edges on which the diamond move is to be performed. The danger is that an inverse pair might be inserted into the word along $\delta_{\hat{u}, \hat{v}}$. But this could only happen when $\delta_{\hat{u}, \hat{v}}$ crosses $\delta_{u,v}$ at the vertex between the two edges, and such crossing is impossible because $\gamma_{u,v}$ and $\gamma_{\hat{u}, \hat{v}}$ are both part of geodesic arcs based at \star in the tree T .

There are only finitely many pairs u, v on ∂I . So after a finite number of diamond moves, θ is transformed to a map that satisfies the requirements

of the lemma. ■

Returning once more to the proof of the theorem, let \bar{v} be a vertex in $\bar{\Delta}$ for which $\rho(\bar{v}, \bar{\star}) = \text{IDiam}(\bar{\Delta})$ and let v be a vertex of Δ such that $\theta(v) = \bar{v}$. Let μ be the geodesic in T from v to \star . And let $\bar{\mu}$ be the edge-path from \bar{v} to $\bar{\star}$, obtained by connecting up the images under θ of the portions of μ not in the interior of \mathcal{P}_k -islands I with geodesic edge-paths in the trees $\theta(I)$. So $\bar{\mu}$ is a concatenation of two types of geodesic arcs: those that run through the image under θ of the boundary of some \mathcal{P}_k -island (call these *island arcs*), and those arcs from μ .

We fix $\alpha \geq \beta > 0$ and examine the following two cases.

Case 1) The island-arcs all have length at most n^α .

In place of each island-arc in $\bar{\mu}$, we find an arc in μ of length at least one (in fact, this is a crude lower bound) because the word along each such arc in μ equals some non-zero power of t in $\Psi_{k,m}$. So $\ell(\mu) \geq \ell(\bar{\mu})/n^\alpha \geq C n^{m-\alpha}$ and therefore $\text{IDiam}(\Delta) \succeq n^{m-\alpha}$, because μ is a geodesic in Δ , based at \star .

Case 2) Some island-arc through $\theta(\partial I)$, where I is some \mathcal{P}_k -island in Δ , has length more than n^α .

Let γ be the corresponding subarc of μ through I . If $\ell(\gamma) \geq n^\beta$ then it follows immediately that $\text{IDiam}(\Delta) \geq \ell(\mu) \geq n^\beta$.

Assume that $\ell(\gamma) < n^\beta$. Then γ divides I into two subdiagrams each of which has boundary circuit made up of γ together with a portion of $\partial\Delta$ around which we read a word in $\{t, t^{-1}\}^*$ with exponent sum more than n^α . Applying Proposition 5.4 to either of these two subdiagrams we learn that the intrinsic diameter of I is at least a constant times $n^\alpha/(1 + n^\beta)$.

So

$$\text{IDiam}(\Delta) \succeq \min \left\{ n^{m-\alpha}, n^\beta, \frac{n^\alpha}{1 + n^\beta} \right\},$$

and taking $\alpha = 2m/3$ and $\beta = m/3$ we finally have our result. ■

8 Amalgams and retractions

The main results in this section concern amalgams (that is, free products with amalgamation). But first we give some technical results on cutting,

gluing and collapsing operations on 2-complexes. We perform such operations on van Kampen diagrams and our concern is that planarity should not be lost.

The proof of the following lemma is straight-forward and we omit it.

Lemma 8.1 *Suppose that Δ is a finite, combinatorial complex embedded in the plane \mathbb{E}^2 and that ρ is the (not necessarily embedded) edge-circuit in $\partial\Delta$ around the boundary of the closure \overline{C} of a component C of $\mathbb{E}^2 - \Delta$ for which \overline{C} is compact.*

1. *If Υ is a topological 2-disc combinatorial complex with $\ell(\partial\Upsilon) = \ell(\rho)$, then gluing Υ to Δ by identifying the boundary circuit of Υ with ρ produces a planar 2-complex.*
2. *If ρ is simple and Υ is a planar contractible combinatorial 2-complex with $\ell(\partial\Upsilon) = \ell(\rho)$, then gluing Υ to Δ by identifying the boundary circuit of Υ with ρ gives a planar 2-complex.*
3. *Identifying two adjacent edges e_1 and e_2 in ρ , as illustrated in Figure 12, produces a planar 2-complex unless e_1 and e_2 together comprise the boundary of a subdiagram of Δ (as in the rightmost diagram of the figure).*




Figure 12: Identification of adjacent edges.

A *singular combinatorial map* Θ from one complex to another is a continuous map in which every closed n -cell e^n is either mapped homeomorphically onto an n -cell or is mapped onto $\Theta(\partial e^n)$.

We leave the proof of the following technical lemma to the reader.

Lemma 8.2 *Let $\mathcal{Q} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ and $\mathcal{Q}_0 = \langle \mathcal{A}_0 \mid \mathcal{R}_0 \rangle$ be finite presentations. Suppose θ is a map $\mathcal{A} \rightarrow \mathcal{A}_0^{\pm 1} \cup \{1\}$ and let $\bar{\theta} : (\mathcal{A}^{\pm 1})^* \rightarrow (\mathcal{A}_0^{\pm 1})^*$ be the extension of θ defined by*

$$\bar{\theta}(a_1^{\varepsilon_1} a_2^{\varepsilon_2} \dots a_r^{\varepsilon_r}) = \theta(a_1)^{\varepsilon_1} \theta(a_2)^{\varepsilon_2} \dots \theta(a_r)^{\varepsilon_r}.$$

Suppose that for all $r \in \mathcal{R}$, the word $\bar{\theta}(r)$ is either freely reducible or has a cyclic conjugate in $\mathcal{R}_0^{\pm 1}$.

If Δ is a \mathcal{Q} -van Kampen diagram for a word w , then there is a singular combinatorial map $\Theta : \Delta \rightarrow \bar{\Delta}$ to a \mathcal{Q}_0 -van Kampen diagram for $\bar{\theta}(w)$ that is distance decreasing with respect to the path metrics on $\Delta^{(1)}$ and $\bar{\Delta}^{(1)}$. Moreover, suppose e is an edge of Δ labelled by a and $\Theta(e)$ is not a single vertex; if $\theta(a) \in \mathcal{A}$ then Θ preserves the orientation of e and $\Theta(e)$ is labelled by $\theta(a)$, and if $\theta(a) \in \mathcal{A}^{-1}$ then Θ reverses the orientation of e and $\Theta(e)$ is labelled by $\theta(a)^{-1}$.

Corollary 8.3 *In addition to the hypotheses of Lemma 8.2 assume that \mathcal{Q}_0 is a subpresentation of \mathcal{Q} . If w is a null-homotopic word in \mathcal{Q}_0 , and if $\bar{\theta}(w) = w$, then there is a \mathcal{Q}_0 -van Kampen diagram for w that is of minimal intrinsic diameter (or radius) amongst all \mathcal{Q} -van Kampen diagrams for w .*

In the following proposition, the hypothesis that no cyclic conjugate of a word in $\mathcal{R}_1^{\pm 1}$ is in $\mathcal{R}_2^{\pm 1}$ is mild. It could only fail for freely reducible words in $\{t^{\pm 1}\}^*$. It allows the 2-cells of a \mathcal{P} -van Kampen Δ to be partitioned into \mathcal{R}_1 -cells and \mathcal{R}_2 -cells; that is, 2-cells that have boundary words in \mathcal{R}_1 or in \mathcal{R}_2 , respectively. Note that whenever an \mathcal{R}_1 -cell shares an edge with an \mathcal{R}_2 -cell, that edge is labelled by t .

Proposition 8.4 *Let $\langle \mathcal{A}_1 \mid \mathcal{R}_1 \rangle$ and $\langle \mathcal{A}_2 \mid \mathcal{R}_2 \rangle$ be presentations of groups A_1 and A_2 , such that $\mathcal{A}_1 \cap \mathcal{A}_2 = \{t\}$, where t has infinite order in both A_1 and A_2 . The amalgam $A_1 *_{\langle t \rangle} A_2$ has presentation $\mathcal{P} := \langle \mathcal{A}_1 \cup \mathcal{A}_2 \mid \mathcal{R}_1 \cup \mathcal{R}_2 \rangle$. Assume that no cyclic conjugate of a word in $\mathcal{R}_1^{\pm 1}$ is in $\mathcal{R}_2^{\pm 1}$.*

Suppose that A_2 retracts to $\langle t \rangle$ via a homomorphism ϕ that maps t to t and maps all other $a \in \mathcal{A}_2$ to 1 or $t^{\pm 1}$.

1. *Suppose w_1 is a null-homotopic word in $(\mathcal{A}_1^{\pm 1})^*$. Then w_1 has an $\langle \mathcal{A}_1 \mid \mathcal{R}_1 \rangle$ -van Kampen diagram that is of minimal intrinsic diameter (or radius) amongst all \mathcal{P} -van Kampen diagrams for w_1 .*
2. *Suppose w_2 is a null-homotopic word in $(\mathcal{A}_2^{\pm 1})^*$. Then w_2 has a minimal intrinsic diameter \mathcal{P} -van Kampen diagram $\hat{\Delta}$ such that whenever γ is a simple edge-circuit in $\hat{\Delta}$ around which we read a word in $(\mathcal{A}_1^{\pm 1})^*$, the subdiagram it bounds is an $\langle \mathcal{A}_1 \mid \mathcal{R}_1 \rangle$ -van Kampen diagram.*

Proof. The first part is a consequence of Corollary 8.3. For the second part we take a \mathcal{P} -van Kampen diagram Δ for w_2 of minimal intrinsic diameter

and obtain a \mathcal{P} -van Kampen diagram $\hat{\Delta}$ for w_2 with the required properties by repeating the following procedure.

Suppose ρ is a simple edge-circuit in Δ around which we read a word in $(\mathcal{A}_1^{\pm 1})^*$. Let w_ρ be the word we read around ρ , and let Δ_ρ be the \mathcal{P} -van Kampen subdiagram bounded by ρ . Glue the $\langle \mathcal{A}_1 \mid \mathcal{R}_1 \rangle$ -van Kampen diagram $\overline{\Delta}_\rho$ for w_ρ supplied by Lemma 8.2, in place of Δ_ρ . This does not increase intrinsic diameter because the map $\Theta : \Delta_\rho \rightarrow \overline{\Delta}_\rho$ of Lemma 8.2 is distance decreasing. Also the gluing cannot destroy planarity because ρ is simple — see Lemma 8.1 (2). ■

9 Quasi-isometry invariance

A special case of the following theorem is that for any two finite presentations $\mathcal{P}_1, \mathcal{P}_2$ of the same group, $\text{IDiam}_{\mathcal{P}_1} \simeq \text{IDiam}_{\mathcal{P}_2}$ and $\text{EDiam}_{\mathcal{P}_1} \simeq \text{EDiam}_{\mathcal{P}_2}$.

Theorem 9.1 *If \mathcal{P}_1 and \mathcal{P}_2 be finite presentations for quasi-isometric groups Γ_1 and Γ_2 , respectively, then $\text{IDiam}_{\mathcal{P}_1} \simeq \text{IDiam}_{\mathcal{P}_2}$ and $\text{EDiam}_{\mathcal{P}_1} \simeq \text{EDiam}_{\mathcal{P}_2}$.*

Proof. The theorem is proved by keeping track of diameters as one follows the standard proof that finite presentability is a quasi-isometry invariant [5, page 143]. The first quantified version of this proof (in the context of Dehn functions) appeared in [1].

Fix word metrics d_1 and d_2 for Γ_1 and Γ_2 , respectively, and quasi-isometries $f : (\Gamma_1, d_1) \rightarrow (\Gamma_2, d_2)$ and $g : (\Gamma_2, d_2) \rightarrow (\Gamma_1, d_1)$ with constants $\lambda \geq 1$ and $\mu \geq 0$, such that for all $u_1, v_1 \in \Gamma_1$ and all $u_2, v_2 \in \Gamma_2$,

$$\frac{1}{\lambda}d_1(u_1, v_1) - \mu \leq d_2(f(u_1), f(v_1)) \leq \lambda d_1(u_1, v_1) + \mu, \quad (3)$$

$$\frac{1}{\lambda}d_2(u_2, v_2) - \mu \leq d_1(g(u_2), g(v_2)) \leq \lambda d_2(u_2, v_2) + \mu, \quad (4)$$

$$d(u_1, g \circ f(u_1)), d(u_2, f \circ g(u_2)) \leq \mu. \quad (5)$$

Suppose ρ_2 is an edge-circuit in the Cayley graph $\text{Cay}^1(\Gamma_2, \mathcal{A}_2)$, visiting vertices $v_0, v_1, \dots, v_n = v_0$ in order. Consider the circuit ρ_1 in $\text{Cay}^1(\Gamma_1, \mathcal{A}_1)$ obtained by joining $g(v_0), g(v_1), \dots, g(v_n)$ by geodesics of length at most $\lambda + \mu$, using (4). Fill ρ_1 with a minimal-intrinsic-diameter van Kampen diagram $\pi_1 : \Delta_1 \rightarrow \text{Cay}^2(\mathcal{P}_1)$. Extend $f \circ \left(\pi_1 \Big|_{\Delta_1^{(0)}} \right)$ by joining the images of adjacent vertices by geodesics in $\text{Cay}^1(\Gamma_2, \mathcal{A}_2)$, each of length at most $\lambda + \mu$ by (3), to give a combinatorial map $\pi : \Delta^{(1)} \rightarrow \text{Cay}^1(\Gamma_2, \mathcal{A}_2)$ from the 1-skeleton of a diagram Δ obtained by subdividing each of the edges of Δ_1 .

Each 2-cell in Δ has boundary length at most $L_1 := (\lambda + \mu) \max_{r \in \mathcal{R}_1} \ell(r)$. Extend π to a map $\pi'_2 : \Delta_2^{(1)} \rightarrow \text{Cay}^1(\Gamma_2, \mathcal{A}_2)$ filling ρ_2 by joining each vertex u_2 on ρ_2 to $f \circ g(u_2)$ on $\pi(\partial\Delta)$ by a geodesic, which has length at most μ by (5). So Δ_2 is obtained from Δ by attaching a collar of $\ell(\rho_2)$ 2-cells around its boundary. Adjacent vertices in $\text{Cay}^1(\Gamma_2, \mathcal{A}_2)$ are mapped by g to vertices at most $\lambda + \mu$ apart by (4) and then by f to vertices at most $\lambda(\lambda + \mu) + \mu$ by (3). So the lengths of the boundaries of the 2-cells in the collar are each at most $L_2 := \lambda^2 + \lambda\mu + 3\mu + 1$.

It follows that if we define \mathcal{R}_2 to be the set of all null-homotopic words in $(\mathcal{A}_2^{\pm 1})^*$ of length at most $L := \max\{L_1, L_2\}$ then π_2 extends to a van Kampen diagram $\pi_2 : \Delta_2 \rightarrow \text{Cay}^2(\mathcal{P}_2)$ filling ρ_2 , where $\mathcal{P}_2 = \langle \mathcal{A}_2 \mid \mathcal{R}_2 \rangle$. So \mathcal{P}_2 is a finite presentation for Γ_2 . Now, $\ell(\partial\Delta_1) \leq (\lambda + \mu)\ell(\rho_2)$ and so $\text{IDiam}(\Delta_1)$ is at most $\text{IDiam}_{\mathcal{P}_1}((\lambda + \mu)\ell(\rho_2))$. By (3) we can multiply this by $(\lambda + \mu)$ to get an upper bound on the intrinsic diameter of Δ . Adding a further 2μ for the collar, we get

$$\text{IDiam}(\Delta_2) \leq (\lambda + \mu) \text{IDiam}_{\mathcal{P}_1}((\lambda + \mu)\ell(\rho_2)) + 2\mu, \quad (6)$$

which establishes $\text{IDiam}_{\mathcal{P}_2} \preceq \text{IDiam}_{\mathcal{P}_1}$ for this particular \mathcal{P}_2 .

However, the theorem concerns *arbitrary* \mathcal{R}_2 for which $\mathcal{P}_2 = \langle \mathcal{A}_2 \mid \mathcal{R}_2 \rangle$ is a finite presentation for Γ_2 . The boundary of each 2-cell of Δ_2 is mapped by π_2 to an edge-circuit in $\text{Cay}^1(\Gamma_2, \mathcal{A}_2)$ of length at most L . So, to make Δ_2 into a van Kampen $\hat{\Delta}_2$ diagram over \mathcal{P}_2 , we fill each of its 2-cells with a (possibly singular) van Kampen diagram over \mathcal{P}_2 . But, a technical concern here is that gluing a *singular* 2-disc diagram along a non-embedded edge-circuit of Δ_2 could destroy planarity. The way we deal with this is to fill the 2-cells of Δ_2 one at a time. And, on each occasion, if the 2-cell C to be filled has non-embedded boundary circuit then we discard all the edges inside the simple edge-circuit σ in ∂C such that no edge of ∂C is outside σ , and then we fill σ .

Discarding the edges inside all such σ does not stop the estimate (6) holding. Adding $\text{IDiam}_{\mathcal{P}_2}(L)$ to account for each of the fillings gives an upper bound on $\text{IDiam}_{\mathcal{P}_2}(\hat{\Delta}_2)$ and so $\text{IDiam}_{\mathcal{P}_2} \preceq \text{IDiam}_{\mathcal{P}_1}$. Interchanging the roles of \mathcal{P}_1 and \mathcal{P}_2 , we immediately deduce that $\text{IDiam}_{\mathcal{P}_1} \preceq \text{IDiam}_{\mathcal{P}_2}$ and so we have $\text{IDiam}_{\mathcal{P}_1} \simeq \text{IDiam}_{\mathcal{P}_2}$, as required.

That $\text{EDiam}_{\mathcal{P}_1} \simeq \text{EDiam}_{\mathcal{P}_2}$ can be proved the same way, except we take Δ_1 to be a minimal-extrinsic-diameter filling of ρ_1 , and then by (3)

$$\text{EDiam}(\Delta_2) \leq \lambda \text{EDiam}_{\mathcal{P}_1}((\lambda + \mu)\ell(\rho_2)) + \mu + 2L,$$

and adding a further constant $\text{EDiam}_{\mathcal{P}_2}(L)$ gives an upper bound on $\text{EDiam}_{\mathcal{P}_2}(\hat{\Delta}_2)$. ■

References

- [1] J. M. Alonso. Inégalités isopérimétriques et quasi-isométries. *C.R. Acad. Sci. Paris Ser. 1 Math.*, 311:761–764, 1990, MR1082628, Zbl 0726.57002
- [2] M. R. Bridson. Asymptotic cones and polynomial isoperimetric inequalities. *Topology*, 38(3):543–554, 1999, MR1670404, Zbl 0929.20032
- [3] M. R. Bridson. The geometry of the word problem. In M. R. Bridson and S. M. Salamon, editors, *Invitations to Geometry and Topology*, pages 33–94. O.U.P., 2002, MR1967746, Zbl 0996.54507
- [4] M. R. Bridson. Polynomial Dehn functions and the length of asynchronously automatic structures. *Proc. London Math. Soc.*, 85(2):441–465, 2002, MR1912057, Zbl 1046.20027
- [5] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-positive Curvature*. Number 319 in Grundlehren der mathematischen Wissenschaften. Springer Verlag, 1999, MR1744486, Zbl 0988.53001
- [6] D. Cooper. Automorphisms of free groups have finitely generated fixed point sets. *J. Algebra*, 111:453–456, 1987, MR0916179, Zbl 0628.20029
- [7] M. Gromov. Asymptotic invariants of infinite groups. In G. Niblo and M. Roller, editors, *Geometric group theory II*, number 182 in LMS lecture notes. Camb. Univ. Press, 1993, MR1253544, Zbl 0841.20039
- [8] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*, volume 152 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1999. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by S. M. Bates, MR1699320, Zbl 1113.53001
- [9] J.-P. Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, 2003. Translated from the French original by J. Stillwell, Corrected 2nd printing of the 1980 English translation, MR1954121, Zbl 1013.20001

MARTIN R. BRIDSON

Mathematical Institute, 24–29 St Giles’, Oxford, OX1 3LB, UK

bridson@maths.ac.uk, <http://people.maths.ox.ac.uk/~bridson/>

TIMOTHY R. RILEY

Department of Mathematics, University Walk, Bristol, BS8 1TW, UK

tim.riley@bristol.ac.uk, <http://www.maths.bris.ac.uk/~matrr/>