

FREE AND FRAGMENTING FILLING LENGTH

M. R. BRIDSON AND T. R. RILEY

Abstract. The filling length of an edge-circuit η in the Cayley 2-complex of a finitely presented group is the least integer L such that there is a combinatorial null-homotopy of η down to a basepoint through loops of length at most L . We introduce similar notions in which the null-homotopy is not required to fix a basepoint, and in which the contracting loop is allowed to bifurcate. We exhibit groups in which the resulting filling invariants exhibit dramatically different behaviour to the standard notion of filling length. We also define the corresponding filling invariants for Riemannian manifolds and translate our results to this setting.

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

Key words and phrases: filling length, null-homotopy, van Kampen diagram

1. I

Consider a vertical cylinder $C \subseteq \mathbb{R}^3$ of height h whose base has diameter $d \ll h$. Let S be the surface formed by the curved portion of C and the disc capping off its top. Topologically, S is a closed 2-disc. The loop ∂S can be homotoped in S to a constant loop through loops of length at most πd by lifting it up the cylinder and then contracting it across the top of C . However, if we insist on keeping a basepoint on ∂S fixed in the course of the null-homotopy then we will encounter far longer loops, some of length at least $2h$.

In this article we will bring to light similar contrasts between basepoint-fixed and basepoint-free null-homotopies for loops in the Cayley 2-complex $\text{Cay}^2(\mathcal{P})$ of a finite presentation \mathcal{P} of a group Γ . Words w that represent 1 in Γ (*null-homotopic* words) correspond to edge-circuits η_w in $\text{Cay}^2(\mathcal{P})$. The filling length $\text{FL}(w)$ of w was defined by Gromov [13] and in a combinatorial context is the minimal length L such that there is a basepoint-preserving *combinatorial null-homotopy* of η_w through loops of length at most L . (A closely related notion called LNCH was considered by Gersten in [9].) We define $\text{FFL}(w)$, the *free filling length* of w , likewise but without holding a basepoint fixed, and $\text{FFFL}(w)$, the *fragmenting free filling length* of w , by also allowing the contracting loops to bifurcate. Detailed definitions are in Section 2.

We construct a finite presentation in which $\text{FL}(w)$ and $\text{FFL}(w)$ differ dramatically for an infinite sequence of null-homotopic words of increasing length. [Our conventions are $[a, b] = a^{-1}b^{-1}ab$, $a^b = b^{-1}ab$ and $a^{-b} = b^{-1}a^{-1}b$. For $f, g : \mathbb{N} \rightarrow \mathbb{N}$ or $f, g : \mathbb{R} \rightarrow \mathbb{R}$, we write $f \leq g$ when there exists $C > 0$ such that for all l we have $f(l) \leq Cg(Cl + C) + Cl + C$, which gives an equivalence relation expressing qualitative agreement of the growths of f and g : write $f \approx g$ if and only if $f \leq g$ and $g \leq f$. We compare functions $\mathbb{N} \rightarrow \mathbb{N}$ with functions $\mathbb{R} \rightarrow \mathbb{R}$ by extending the former to \mathbb{R} by considering them to be constant on the intervals $[n, n + 1)$ for all $n \in \mathbb{N}$.]

Theorem A. *Let Γ be the group given by the aspherical presentation*

$$Q := \langle a, b, r, s, t \mid a^b a^{-2}, [t, a], [r, at], [r, s], [s, t] \rangle.$$

Date: December 8, 2005; revised May 16, 2006. To appear in the Journal of Algebra.

For $n \in \mathbb{N}$ define $w_n := [s, a^{-b^n} r a^{b^n}]$. Then w_n is null-homotopic in \mathcal{Q} , has length $\ell(w_n) = 8n + 8$, and $\text{FFL}(w_n) \simeq n$, but $\text{FL}(w_n) \simeq 2^n$.

We assume the reader is familiar with *van Kampen diagrams* ([4] is a recent survey); they can be thought of as *combinatorial homotopy discs* for loops in the Cayley 2-complex of a presentation. We show that there are van Kampen diagrams $\hat{\Delta}_n$ for w_n that, owing to geometry like that in our cylinder example, have $\text{FFL}(\hat{\Delta}_n) \leq n$ and $\text{FL}(\hat{\Delta}_n) \simeq 2^n$. Indeed, we will show that every van Kampen diagram Δ for w_n has intrinsic diameter $\geq 2^n$; that is, there is a pair of vertices x, y such that every edge-path in $\Delta^{(1)}$ connecting x to y has length $\geq 2^n$. It will then follow that $\text{FL}(w_n) \geq 2^n$.

The area $\text{Area}(w)$ (resp. intrinsic diameter $\text{IDiam}(w)$) of w is the least integer K such that there is a van Kampen diagram for w with K 2-cells (resp. with intrinsic diameter K). For $M = \text{Area}, \text{IDiam}, \text{FL}, \text{FFL}$ or FFFL one defines *filling functions* $M : \mathbb{N} \rightarrow \mathbb{N}$ for finitely presented groups:

$$M(n) := \max \{ M(w) \mid w \text{ null-homotopic and } \ell(w) \leq n \}.$$

(The argument of M — diagram, word or integer — determines its meaning; the potential for ambiguity is tolerated as it spares us from a terminology over-load.) In the case $M = \text{Area}$, the function $M(n)$ is the *Dehn function*.

In spite of Theorem A, the filling functions FL and FFL for \mathcal{Q} are \simeq -equivalent: we will see that any van Kampen diagram for $w'_n := [s, a^{-b^n} r a^{b^n} r a^{b^n} r^{-1} a^{-b^n}]$ has two *peaks* and the savings that can be made by escaping the basepoint are no longer significant. On the other hand, if we allow our loops to bifurcate then they can pass over peaks independently, and so FFFL exhibits markedly different behaviour.

Theorem B. *The filling functions $\text{IDiam}, \text{FL}, \text{FFL}, \text{FFFL} : \mathbb{N} \rightarrow \mathbb{N}$ for \mathcal{Q} satisfy*

$$\begin{aligned} \text{IDiam}(n) &\simeq \text{FL}(n) \simeq \text{FFL}(n) \simeq \text{Area}(n) \simeq 2^n \\ \text{FFFL}(n) &\simeq n. \end{aligned}$$

The relations $\text{IDiam}(n) \simeq \text{Area}(n)$ and $\text{FL}(n) \simeq \text{Area}(n)$ are noteworthy in the non-hyperbolic setting. In contrast, for groups where $\text{Area}(n) \simeq n^\alpha$ for some $\alpha \geq 2$ one knows that $\text{IDiam}(n) \leq n^{\alpha-1}$ — see [11].

Theorems A and B are proved in Section 3, modulo a number of auxiliary propositions postponed to Section 4. The remaining sections are dedicated to establishing the credentials of FL, FFL and FFFL for inclusion in the pantheon of filling invariants: we relate them to other filling functions, we interpret them in terms of algorithmic complexity, and we relate them to functions concerning null-homotopies of loops in metric spaces (particularly, Riemannian manifolds).

FL can be thought of as a space-complexity measure in that $\text{FL}(w)$ is the minimal L such that w can be converted to the empty word through a sequence of words of length at most L , each obtained from the previous by free reduction, free expansion, or applying a relator (see [11]). We will show in Section 2 that FFL and FFFL can be interpreted in a similar way: by allowing an additional operation of conjugation we get $\text{FFL}(w)$, and further including the move that replaces a word $w = uv$ by a pair of words u, v we get $\text{FFFL}(w)$. This point of view is useful for calculations and allows us to prove that, as for FL , given a finite presentation, $\text{Area}(n)$ is at most an exponential of FFFL . This and other relationships are spelt out in the following theorem, which shows, in particular, that \mathcal{Q} provides an example of $\text{Area}(n)$ outgrowing $\text{FFFL}(n)$ as extremely as is possible.

Theorem C. *Let \mathcal{P} be a finite presentation. There is a constant C , depending only on \mathcal{P} , such that for all $n \in \mathbb{N}$ the Dehn, filling length, free filling length, and fragmenting free filling length functions $\text{Area}, \text{FL}, \text{FFL}, \text{FFFL} : \mathbb{N} \rightarrow \mathbb{N}$ of \mathcal{P} satisfy*

$$\begin{aligned} \text{Area}(n) &\leq C^{\text{FFFL}(n)} \\ \text{FFFL}(n) &\leq \text{FFL}(n) \leq \text{FL}(n) \leq C \text{Area}(n) + n. \end{aligned}$$

In Section 5 we prove that, up to \simeq -equivalence, $\text{FL}(n)$, $\text{FFL}(n)$ and $\text{FFFL}(n)$ are all quasi-isometry invariants among finitely presented groups and so, in particular, do not depend on a choice of finite presentation of a group – cf. [12, Theorem 8.1].

Theorem D. *If \mathcal{P} and \mathcal{P}' are finite presentations of quasi-isometric groups then*

$$\text{FL}_{\mathcal{P}} \simeq \text{FL}_{\mathcal{P}'}, \quad \text{FFL}_{\mathcal{P}} \simeq \text{FFL}_{\mathcal{P}'}, \quad \text{and} \quad \text{FFFL}_{\mathcal{P}} \simeq \text{FFFL}_{\mathcal{P}'}$$

In Section 6 we define analogous filling invariants $\text{FL}_X(n)$, $\text{FFL}_X(n)$ and $\text{FFFL}_X(n)$ concerning null-homotopies of rectifiable loops in arbitrary metric spaces X , and we prove:

Theorem E. *Suppose a group Γ with finite presentation $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ acts properly and cocompactly by isometries on a simply connected geodesic metric space X for which there exist $\mu, L > 0$ such that every loop of length less than μ admits a based null-homotopy of filling length less than L . Then $\text{FL}_{\mathcal{P}} \simeq \text{FL}_X$, $\text{FFL}_{\mathcal{P}} \simeq \text{FFL}_X$ and $\text{FFFL}_{\mathcal{P}} \simeq \text{FFFL}_X$.*

As the universal cover of any closed connected Riemannian manifold satisfies these conditions and (as is well known) every finitely presentable group is the fundamental group of such a manifold, we can use Theorem E and its analogues for Area and IDiam (proved in [4, 6]) to obtain from Theorems A and B:

Corollary F. *There exists a closed connected Riemannian manifold M such that*

$$\begin{aligned} \text{IDiam}_{\tilde{M}}(n) &\simeq \text{FL}_{\tilde{M}}(n) \simeq \text{FFL}_{\tilde{M}}(n) \simeq \text{Area}_{\tilde{M}}(n) \simeq 2^n \\ &\text{FFFL}_{\tilde{M}}(n) \simeq n. \end{aligned}$$

Moreover, there is an infinite sequence of loops c_n in \tilde{M} such that $\ell(c_n) \rightarrow \infty$, $\text{FFL}_{\tilde{M}}(c_n) \simeq n$, and $\text{FL}_{\tilde{M}}(c_n) \simeq 2^n$.

[Strictly speaking, the final part is not a direct corollary of the prior results, but it follows from the methods of Section 6: the w_n of Theorem A can be used to construct *word-like* loops c_n with $\ell(c_n) \simeq \ell(w_n)$; van Kampen diagrams witnessing to the upper bounds $\text{FFL}(w_n) \leq n$ and $\text{FL}(w_n) \leq 2^n$ can be translated into fillings of c_n showing $\text{FFL}_{\tilde{M}}(c_n) \leq n$ and $\text{FL}_{\tilde{M}}(c_n) \leq 2^n$; and every null-homotopy of c_n in \tilde{M} has filling length $\geq 2^n$ because the filling-disc corresponding to a null-homotopy facilitates a van Kampen diagram for w_n with \simeq similar filling length.]

Natural questions that remain open include the following.

Open problem 1.1. Does there exist a finite presentation for which $\text{FL}(n) \neq \text{FFL}(n)$?

Open problem 1.2. Does there exist a finite presentation for which $\text{IDiam}(n) \neq \text{FL}(n)$?

In reference to the first of these problems we note that for a finite presentation $\langle \mathcal{A} \mid \mathcal{R} \rangle$ of a group Γ , the presentation $\langle \mathcal{A} \cup \{t\} \mid \mathcal{R} \rangle$ for $\Gamma * \mathbb{Z}$ satisfies $\text{FL}(n) \simeq \text{FFL}(n)$ for the following reason. If w_n is a word over \mathcal{P} with $\text{FL}(w_n) = \text{FL}(n)$ then $\text{FFL}([w_n, t]) \geq 2 + \ell(w_n) + \text{FL}(w_n)$ since any diagram for $[w_n, t]$ is a diagram for w_n joined to a diagram for w_n^{-1} by a t -edge, and the most efficient way to shell such a diagram (from the point-of-view of FFL) is to shell the w_n diagram, then collapse the t -edge, and then shell the w_n^{-1} diagram.

The second problem was posed by Gromov in [13, §5C], where it is also attributed to Casson. A negative answer would imply that the *double exponential upper bound* [7, 8] on the Dehn (i.e. area) function in terms of $\text{IDiam}(n)$ could be improved to a single exponential using the single exponential upper bound [11, 13] for $\text{Area}(n)$ in terms of $\text{FL}(n)$.

Acknowledgement. The second author is grateful for support from NSF grant DMS–0404767 and for the hospitality of the Centre de Recerca Matemàtica in Barcelona during the writing of this article.

2. T

Let Δ be a *diagram*; that is, a finite, planar, contractible, combinatorial 2-complex; i.e. a van Kampen diagram bereft of all group theoretic decorations. Before discussing filling length, we recall a combinatorial notion of null-homotopy called a *shelling* from [11].

Definition 2.1. A *shelling* $\mathcal{S} = (\Delta_i)$ of Δ is a sequence of diagrams

$$\Delta = \Delta_0, \Delta_1, \dots, \Delta_m,$$


in which each Δ_{i+1} is obtained from Δ_i by one of the *shelling moves* defined below and depicted in Figure 1.

- *1-cell collapse.* Remove a pair (e^1, e^0) where e^1 is a 1-cell with $e^0 \in \partial e^1$ and e^1 is attached to the rest of Δ_i only by one of its end vertices $\neq e^0$. (We call such an e^1 a *spike*.)
- *1-cell expansion.* Cut along some 1-cell e^1 in Δ_i that has a vertex e^0 in $\partial\Delta_i$, in such a way that e^0 and e^1 are doubled.
- *2-cell collapse.* Remove a pair (e^2, e^1) where e^2 is a 2-cell which has some edge $e^1 \in (\partial e^2 \cap \partial\Delta_i)$. The effect on the boundary circuit is to replace e^1 with $\partial e^2 \setminus e^1$.

We say that the shelling \mathcal{S} is *full* when Δ_m is a single vertex. A *full shelling to a base vertex* $\star = \Delta_m$ on $\partial\Delta$ is a full shelling in which \star is preserved throughout the sequence (Δ_i) . In particular, in every *1-cell collapse* $e^0 \neq \star$, and in every *1-cell expansion* on Δ_i where $e_0 = \star$ a choice is made as to which of the two copies of e_0 is to be \star in Δ_{i+1} .

We define a *full fragmenting shelling* \mathcal{S} of Δ by adapting the definition of a full shelling by allowing each Δ_i to be a disjoint union of finitely many diagrams, insisting that Δ_m be a set of vertices, and allowing one extra type of move:

- *Fragmentation.* Δ_{i+1} is the disjoint union of Δ'_i and Δ''_i , where $\Delta_i = \Delta'_i \cup \Delta''_i$ and $\Delta'_i \cap \Delta''_i$ is a single vertex.



F 1. Shelling moves.

For a shelling \mathcal{S} , define $\ell(\mathcal{S}) := \max_i \ell(\partial\Delta_i)$, where $\ell(\partial\Delta_i)$ denotes the sum of the lengths of the boundary circuits of the components of Δ_i . Then the *filling length* $\text{FL}(\Delta, \star)$,

the *free filling length* $\text{FFL}(\Delta)$, and the *fragmented free filling length* $\text{FFFL}(\Delta)$ are the minima of $\ell(\mathcal{S})$ as \mathcal{S} ranges over all full shellings to \star , full shellings, and all full fragmenting free shellings of Δ , respectively. This notation emphasises the fact that $\text{FL}(\Delta, \star)$ is defined with respect to a base vertex $\star \in \partial\Delta$ but $\text{FFL}(\Delta)$ and $\text{FFFL}(\Delta)$ are not.

Let $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ be a finite presentation of a group Γ . Define the filling length $\text{FL}(w)$ of a null-homotopic word w in a presentation $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ by

$$\text{FL}(w) := \min \{ \text{FL}(\Delta, \star) \mid \Delta \text{ a van Kampen diagram for } w \},$$

and $\text{FFL}(w)$ and $\text{FFFL}(w)$ likewise.

A sequence $\mathcal{S} = w_0, \dots, w_m$ of null-homotopic words is a *null-sequence* if each w_{i+1} is obtained from w_i by one of the following moves:

- *Free reduction.* Remove a subword aa^{-1} or $a^{-1}a$ from w_i , where $a \in \mathcal{A}$.
- *Free expansion.* This is the inverse of a free reduction.
- *Application of a relator.* Replace a subword u of w_i by a word v such that a cyclic conjugate of uv^{-1} is in $\mathcal{R}^{\pm 1}$.

We define two more moves:

- *Cyclic conjugation.* Replace w_i by a cyclic permutation.
- *Fragmentation.* Replace a word $w = uv$ by a pair of words u, v . (In effect, insert a letter into w that represents a blank space.)

To employ the fragmentation move we must generalise our definition of a *null-sequence* so that each w_i is a finite sequence of words, and when we perform any of the operations listed above we execute it on one of the words in w_i .

The proof of the following reassuring lemma is straightforward. The “only if” part is well known (indeed, freely reducing, freely expanding, and applying relators suffices). The “if” part can be proved by an easy induction on the number of fragmentation moves used.

Lemma 2.2. *A word w over \mathcal{P} represents the identity if and only if it can be reduced to a sequence of empty words by free reductions, free expansions, applying relators, cyclically conjugating, and fragmenting.*

For a null-sequence \mathcal{S} , define $\ell(\mathcal{S}) := \max_i \ell(w_i)$ where, if fragmentation moves are employed, $\ell(w_i)$ is the sum of the lengths of words in the sequence w_i . Proposition 1 in [11] says that for all null-homotopic words w , we have $\text{FL}(w) = \min_{\mathcal{S}} \ell(\mathcal{S})$, quantifying over null-sequences \mathcal{S} for w that employ *free reductions, free expansions and applications of relators*. We add the following.

Proposition 2.3. *Quantifying over all null-sequences \mathcal{S} for a null-homotopic word w , where free reduction, free expansion, applications of relators, and cyclic conjugation are allowed, we have $\text{FFL}(w) = \min_{\mathcal{S}} \ell(\mathcal{S})$. If, additionally, we allow fragmentations we get $\text{FFFL}(w) = \min_{\mathcal{S}} \ell(\mathcal{S})$.*

Proof. The proof in [11] that $\min_{\mathcal{S}} \ell(\mathcal{S}) \leq \text{FL}(w)$ is straightforward — the words around the boundary of the van Kampen diagrams in the course of a full shelling form a null-sequence. Each word in the sequence of boundary words in the course of a *free* shelling of a van Kampen diagram is only defined up to cyclic conjugation, as a base vertex is not kept fixed during the shelling. It is then easy to see that $\min_{\mathcal{S}} \ell(\mathcal{S}) \leq \text{FFL}(w)$, quantifying over all all null-sequences \mathcal{S} for w that use *free reduction, free expansion, applications of relators, and cyclic conjugation*. Introducing *fragmentation* moves into the shelling, produces *fragmentations* in the corresponding null-sequence, and we see $\min_{\mathcal{S}} \ell(\mathcal{S}) \leq \text{FFFL}(w)$.

The reverse bounds require more care. Given a null-sequence \mathcal{S} for w involving *applications of relations* and *free expansions* and *reductions*, we seek to construct a van Kampen diagram with a shelling during which the lengths of the boundary circuit remain at most $\ell(\mathcal{S})$. This can be done by starting with an edge-loop labelled by w , and filling it in by attaching a 2-cell on every *application of a relator*, by attaching a 1-cell on every *free expansion*, by folding together two adjacent 1-cells on every *free reduction*. However it is possible that the resulting complex will not be planar: 2-spheres or other cycles may be pinched off (for example, when an inverse pair aa^{-1} is inserted and then removed). Removing these cycles gives a van Kampen diagram Δ with $\text{FL}(\Delta) \leq \ell(w)$. This is explained carefully in [11]. If \mathcal{S} also uses *cyclic conjugation* moves then no extra complications are added to the construction of Δ . If \mathcal{S} also uses *fragmentations* then the corresponding move in the course of the construction of Δ is to identify two vertices so that an inner boundary circuit is changed from a topological circle to a figure-eight. \square

Remark 2.4. (*Filling length and space complexity.*) Envisaging a null-sequence to be the course of a calculation on a Turing tape, we see that $\text{FL}(w)$ is the non-deterministic space complexity of the following approach to solving the word problem for \mathcal{P} : write w on the tape and then exhaustively apply relators and perform free reductions and free expansions. A sequence of moves that converts w into the empty word amounts to a proof that w represents 1 in \mathcal{P} and $\text{FL}(w)$ is the minimal upper bound on the number of places on the tape that have to be used in the calculation (see [10] for more details). If we allow cyclic conjugation then the non-deterministic space complexity is $\text{FFL}(w)$. If we also include fragmentation then the non-deterministic space complexity is $\text{FFFL}(w)$ plus the maximum number of blank spaces separating the words; that is, between $\text{FFFL}(w)$ and $2\text{FFFL}(w)$.

The following inequalities are easy

$$(1) \quad \text{FFFL}(w) \leq \text{FFL}(w) \leq \text{FL}(w) \leq K \text{Area}(w) + \ell(w),$$

$$(2) \quad \text{IDiam}(w) \leq \text{FL}(w),$$

where K is a constant depending only on \mathcal{P} . (See [11] or [13].) Similar inequalities, for example $\text{IDiam}(n) \leq \text{FL}(n)$ for all n , relate the corresponding filling functions (defined in Section 1).

It is well known [11, 13] that there is a constant C , that depends only on \mathcal{P} , such that the Dehn function $\text{Area} : \mathbb{N} \rightarrow \mathbb{N}$ of \mathcal{P} satisfies


$$\text{Area}(n) \leq C^{\text{FL}(n)}$$

for all n . This is essentially the Time-Space bound of algorithmic complexity: the number of different words of length $\text{FL}(n)$ on an alphabet of size C is $C^{\text{FL}(n)}$, and as w admits a null-sequence that does not include repeated words, $C^{\text{FL}(n)}$ is an upper bound on the number of times relations are applied in the null-sequence and hence on $\text{Area}(w)$. The same proof applies to FFFL and so with (1) gives Theorem C.

3. P T A B

Proof of Theorem A. We exploit a technique due to the first author in [3] to show that the intrinsic diameter $\text{IDiam}(w_n)$ of $w_n = [s, a^{-b^n} r a^{b^n}]$ is at least 2^n .

Suppose $\pi : \Delta \rightarrow \text{Cay}^2(Q)$ is a van Kampen diagram for w_n . Figure 2 is a schematic depiction of Δ and Figure 3 shows an explicit example when $n = 3$. We seek an edge-path μ in Δ , labelled by a word in which the exponent sum of the letters t is 2^n . A s -corridor


 F 2. A van Kampen diagram Δ for w_n with a subdiagram Σ .

through Δ connects the two letters s in w_n . Along each side of this corridor we read a word u in $\{r^{\pm 1}, t^{\pm 1}\}^*$. Let Σ be a subdiagram of Δ with boundary made up of one side of the s -corridor and a portion of $\partial\Delta$ labelled $a^{-b^n}ra^{b^n}$. An r -corridor in Σ joins the r in $a^{-b^n}ra^{b^n}$ to some edge-labelled r in u . Let u_0 be the prefix of u such that the letter immediately following u_0 is this r , and then let μ be the edge-path along the side of the r -corridor running from the vertex at the end of a^{-b^n} to the vertex at the end of u_0 . Let $v \in \{(at)^{\pm 1}\}^*$ be the word along μ . Then $u_0 = a^{b^n}v$ in \mathcal{Q} . Killing r , s and t , retracts \mathcal{Q} onto the subpresentation $\langle a, b \mid a^b = a^2 \rangle$. So $\bar{v} = a^{b^n}$ in $\langle a, b \mid a^b = a^2 \rangle$, where $\bar{v} \in \{a^{\pm 1}\}^*$ is v with all letters $t^{\pm 1}$ removed. It follows that the exponent sum of the letters in \bar{v} is 2^n and hence that μ has the desired property.


Killing all generators other than t defines a retraction ϕ of \mathcal{Q} onto $\langle t \rangle \cong \mathbb{Z}$ that is distance decreasing with respect to word metrics. But the image of $\phi \circ \pi : \Delta \rightarrow \mathbb{Z}$ has diameter at least 2^n on account of μ . So $\text{IDiam}(w_n) \geq 2^n$, as claimed.

It is easy to check that the van Kampen diagram $\hat{\Delta}_n$ for w_n constructed below admits a shelling down to their base vertex that realises the bound $\text{FL}(w_n) \leq 2^n$. So, as $\text{IDiam}(w_n) \leq \text{FL}(w_n)$, we deduce that $\text{IDiam}(w_n) \simeq \text{FL}(w_n) \simeq 2^n$.

The bound $\text{FFL}(w_n) \leq n$ follows from Proposition 4.4. Nonetheless we will sketch a proof since the salient ideas, developing the cylinder example from Section 1, appear here more transparently than in the more general contexts of Section 4.

The van Kampen diagram $\hat{\Delta}_3$ for w_3 is depicted in Figure 3. The analogous construction of the diagram $\hat{\Delta}_n$ for w_n should be clear. Within $\hat{\Delta}_n$ there are four triangular subdiagrams over the subpresentation $\langle a, t \mid [a, t] \rangle$, in which strings of a -edges run vertically (in the sense of Figure 3). Cut along each of these strings (except those of length 2^n at the left and right sides of the diagrams) and insert back-to-back copies of the $\langle a, b \mid a^b = a^2 \rangle$ -diagrams Ω_k that *shortcut* a^k to a word u_k of length $\sim \log_2 k \leq n$. These *shortcut diagrams* Ω_k are constructed in Proposition 4.1 and the way they are inserted is shown (in a more general context) in Figure 8. Call the resulting diagram Δ_n .

For $0 \leq k \leq 2^n$ let $\hat{\rho}_k$ be the edge-paths in $\hat{\Delta}_n$, forming concentric squares in Figure 3, labelled by $a^k s a^{-k} r^{-1} a^k s^{-1} a^{-k} r$. Next, for $0 \leq k \leq 2^n$ define ρ_k to be the edge-path in Δ_n that is obtained from $\hat{\rho}_k$ by replacing each subword $a^{\pm k}$ by its shortcut $u_k^{\pm 1}$. (In particular $\rho_{2^n} = \partial\Delta_n$.) Note that for all k , the length of ρ_k is $\leq n$.

F 3. The van Kampen diagram $\hat{\Delta}_3$ for w_3 .

We now briefly describe a full shelling of Δ_n that realises the bound $\text{FFL}(w_n) \leq n$. In the course of this shelling we encounter the subdiagrams of Δ_n that have ρ_k as their boundary loops. The following lemma concerning the diagrams Ω_k of Proposition 4.1 is the key to shelling the subdiagram bounded by ρ_{k+1} down to that bounded by ρ_k . During this shelling, the length of the boundary loop remains less than a linear function of n because $\log_2 k \leq n$.

Lemma 3.1. *Let Π be the van Kampen diagram comprising a copy of Ω_{k+1} and a copy Ω_k joined to each side of a t -corridor along sides labelled a^{k+1} . Let \star be a vertex on $\partial\Pi$ located at the start of either of the paths labelled a^k along the sides of the t -corridor. Then $\text{FL}(\Pi, \star) \leq \log_2 k$.*

The proof of this lemma becomes clear when one considers concurrently running the shellings of Ω_{k+1} and Ω_k of Proposition 4.1, and a shelling of the t -corridor.


We complete the proof of Theorem A by noting that the lower bound $\text{FFL}(w_n) \geq n$ is trivial as $\text{FFL}(w_n) \geq \ell(w_n) = 8n + 8$, and so $\text{FFL}(w_n) \simeq n$. \square

Proof of Theorem B. The lower bound of 2^n on $\text{IDiam}(w_n)$ established above proves that $2^n \leq \text{IDiam}(n)$. So by (1) in Section 2 we see that $2^n \leq \text{FL}(n) \leq \text{Area}(n)$ and $\text{FFL}(n) \leq \text{Area}(n)$ for all $n \in \mathbb{N}$.

To establish $2^n \leq \text{FFL}(n)$ we show that the words

$$w'_n := [s, a^{-b^n} r a^{b^n} r a^{b^n} r^{-1} a^{-b^n}]$$

satisfy $\text{FFL}(w'_n) \geq 2^n$. Figure 4 shows a van Kampen diagram for w'_n built from a copy of the diagram $\hat{\Delta}_n$ for w_n , an $[r, s]$ -2-cell, and a mirror image of $\hat{\Delta}_n$. Thus w'_n is null-homotopic.


 F 4. A van Kampen diagram for w'_n .

To show that $\text{FFL}(\Delta') \geq 2^n$ for *all* van Kampen diagrams Δ' for w'_n we develop the argument used to establish the lower bound on diameter in the proof of Theorem A. An s -corridor runs through Δ' , and as there are three occurrences of r in w'_n and three of r^{-1} , each r is joined to an r^{-1} by an r -corridor that crosses the s -corridor. This is illustrated in Figure 5; note that the behaviour of the corridors could be more complex because each r -corridor could cross the s -corridor multiple times. As shown in Figure 5, let u_1 , u_2 and u_3 be the words along the sides of the initial portions of these three r -corridors running from the first, second and third r in w'_n and ending where the corridor first meets the s -corridor. (So $u_1 = u_3 = (at)^{2^n}$ and u_2 is the empty word in the example of Figure 4.)


Retracting Q onto the subpresentation $\langle a, b \mid b^{-1}ab = a^2 \rangle$ by killing r , s and t we see that the exponent sum of the letter a , and hence also of t , in u_1 and in u_3 is 2^n , while the exponent sum in u_2 is 0. The word along the side of the s -corridor is of the form $v_1rv_2rv_3rv_4$ where $v_i \in \{r^{\pm 1}, t^{\pm 1}\}^*$ and v_i runs to the vertex at the end of u_i for $i = 1, 2, 3$ (see Figure 5). By considering the retraction of Q onto $\langle t \rangle$ in which a, b, r, s are killed, we see that the exponent sum of t in v_i is 2^n for $i = 1, 4$ and is -2^n for $i = 2, 3$.

Suppose $\mathcal{S} = (\Delta'_i)$ is a full shelling of Δ' , in the course of which the base vertex is not required to be kept fixed. The retraction ϕ in which all generators other than t are killed, defines a distance decreasing map ϕ of Q onto $\langle t \rangle \cong \mathbb{Z}$. And the edge-circuit $\phi \circ \pi|_{\partial\Delta'_i} : \partial\Delta'_i \rightarrow \mathbb{Z}$ has length at most $\ell(\partial\Delta'_i)$. There are natural combinatorial maps $\psi_i : \Delta'_i \rightarrow \Delta'$ (only prevented from being injective by 1 -cell expansion moves) under which $\psi_i(\partial\Delta'_i)$ forms a contracting sequence of edge-circuits. Let i be the least integer such that $\psi_i(\partial\Delta'_i)$ includes either the vertex at the end of v_1 or at the start of v_4 – these are ringed by small circles in Figure 5. We will explain why $\ell(\partial\Delta'_i)$ is at least 2^n when this vertex x is at the end of v_1 . A similar argument will show the same result to hold when x is the vertex at the start of v_4 . Some vertex y on v_4 must be included in $\psi_i(\partial\Delta'_i)$ because the contracting edge-circuit cannot have yet crossed the vertex at the start of v_4 . So

$$\text{FFL}(w'_n) \geq \ell(\partial\Delta'_i) \geq |\phi(x) - \phi(y)| \geq 2^n,$$

as required.

It is clear that $n \geq \text{FFFL}(n)$ because for all null-homotopic words w we have $\text{FFFL}(w) \geq \ell(w)$. Proposition 4.7, which is the culmination of a sequence of propositions in Section 4, will give $\text{FFFL}(n) \simeq n$. Theorem C will then give us $\text{Area}(n) \leq 2^n$ and the proof of the theorem will be complete. \square

F 5. A van Kampen diagram Δ' for w'_n .

4. A

In this section we provide a number of results which build up to Proposition 4.7 where we prove a linear upper bound on $\text{FFFL}(n)$ in the presentation \mathcal{Q} of Theorem A. We begin in Proposition 4.1 with a technical result giving carefully controlled shellings of diagrams Ω_k that *shortcut* a^k in $\langle a, b \mid a^b = a^2 \rangle$ to a word of length $\leq \log_2 k$. Next Proposition 4.2 claims the filling length of $\mathcal{Q}_0 := \langle a, b, t \mid a^b a^{-2}, [a, t] \rangle$ admits a linear upper bound.

Proposition 4.4 gives a linear upper bound on $\text{FFL}(w)$ for null-homotopic words in \mathcal{Q} that have exactly one pair of letters r, r^{-1} and one pair s, s^{-1} , with the occurrences of the $r^{\pm 1}$ alternating with those of the $s^{\pm 1}$. We show that such words have diagrams with one s - and one r -corridor with these corridors crossing only once. These diagrams can be thought of as *towers*, and are exemplified the diagrams Δ_n for w_n constructed in the proof of Theorem A. In the proof of Proposition 4.4 we refer back to Proposition 4.3 which provides controlled shellings for the four subdiagrams that *tower* diagram breaks up into if we remove the r - and s -corridors.

Next, in Proposition 4.5, we establish an upper bound on $\text{FFFL}(w)$ that is linear in $\ell(w)$, for null-homotopic words w in \mathcal{Q} that have exactly one pair of letters s, s^{-1} . The essential idea here is to find a diagram that can be fragmented into a number of subdiagrams, one for each r -corridor that crosses the s -corridor, and apply Proposition 4.4 to each of these subdiagrams. Proposition 4.6 takes care of the case where there is no letter s in w . Finally, we prove Proposition 4.7: given a word that is null-homotopic in \mathcal{Q} we construct a diagram that can be fragmented into subdiagrams each of which contain at most one s -corridor, and then we apply Propositions 4.5 and 4.6.


Proposition 4.1. *Let \mathcal{P} be the presentation $\langle a, b \mid a^b = a^2 \rangle$. Fix $k > 0$. There is a word u_k of length at most $12 + 4 \log_2 k$ such that $u_k = a^k$ in \mathcal{P} . Moreover, there is a \mathcal{P} -van Kampen diagram Ω_k with boundary word $a^k u_k^{-1}$ (as read from a vertex \star) satisfying the following condition: if μ is the subarc of $\partial\Omega_k$ along which one reads a^k , then there is a shelling $\mathcal{S} = \Omega_{k_0}, \dots, \Omega_{k_p}$ that collapses $\Omega_k = \Omega_{k_0}$ to $\star = \Omega_{k_p}$, such that each Ω_{k_i} is a subdiagram of Ω_k and, expressing the boundary circuit of Ω_{k_i} as μ_i followed by ν_i , where μ_i is the maximal length subarc of $\partial\Omega_{k_i}$ that starts at \star and follows μ , we have $\ell(\nu_i) \leq 12 + 4 \log_2 k$.*

Proof. Let m be the least integer such that $2^m \geq k$. So $m < 1 + \log_2 k$. Let Ξ be the standard van Kampen diagram that demonstrates the equality $a^{b^m} = a^{2^m}$ and is depicted in Figure 6 in the case $m = 4$. Let (Ξ_i) be the shelling of Ξ in which each Ξ_i is a subdiagram of Ξ and Ξ_{i+1} is obtained from Ξ_i as follows. A 1-cell collapse is performed on a *spike*

of Ξ_i if possible. Otherwise, of the rightmost 2-cells in Ξ_i , let e^2 be the lowest in the sense of Figure 6. Let e^1 be the right-most edge of the lower horizontal side of e^2 . Then $e_1 \in \partial\Xi_i$. Perform a 2-cell collapse on (e^2, e^1) . As an illustration, the numbers in the 2-cells in Figure 6 show the order in which the 2-cells are collapsed.

When there is no spike in Ξ_i , its anticlockwise boundary circuit, starting from \star , follows *horizontal* (in the sense of Figure 6) edges labelled by a , then travels *upwards* through the boundary of m 2-cells (visiting at most three 1-cells of each) and then *descends* back to \star along edges labelled b . Removing the initial horizontal path, the number of edges traversed is at most $4m$. And, as the total length of the 1-dimensional portions of $\partial\Xi_i$ is at most 8, and $m < 1 + \log_2 k$, we deduce that the length of $\partial\Phi_i$ minus the length of the initial horizontal arc, is at most $12 + 4 \log_2 k$.

In the case $k = 2^m$ we find that defining $\Omega_i := \Xi_i$ for all $i \geq 0$ gives the asserted result. For the case $k \neq 2^m$, let c be the maximum i such that a^k is a prefix of the word w one reads anticlockwise around $\partial\Xi_i$, starting from \star . Then defining $u_k := w_0^{-1}$, where $w = a^k w_0$ and $\Omega_{ki} := \Xi_{c+i}$ for all $i \geq 0$, we have our result. \square



F 6. The van Kampen diagram Ξ for $a^{2^4} a^{-b^4}$.

Proposition 4.2. *The filling length function $FL : \mathbb{N} \rightarrow \mathbb{N}$ of*

$$\mathcal{Q}_0 := \langle a, b, t \mid a^b a^{-2}, [a, t] \rangle$$

admits a linear upper bound.

Proof. Corollary E1 of [2] says that an HNN extension of a finitely generated free group with finitely many stable letters, in which the associated subgroups are all finitely generated, is asynchronously automatic. This applies to \mathcal{Q}_0 . Theorem 3.1 in [9] says (in different language) that if a group is asynchronously combable then its filling length function admits a linear upper bound. \square

Proposition 4.3. *Suppose $(at)^{-k} wt^j$ is a null-homotopic word in $\mathcal{Q}_0 = \langle a, b, t \mid a^b a^{-2}, [t, a] \rangle$. Let Λ be the 1-dimensional van Kampen diagram for*


$$t^j (at)^{-k} (at)^{k-1} at^{-(j-1)}$$

constructed by assembling 1-cells in \mathbb{R}^2 as depicted at the right in Figure 7. There is a \mathcal{Q}_0 -van Kampen diagram Δ for $(at)^{-k} wt^j$ with the following properties. There is a shelling of Δ through a sequence of diagrams $\Delta = \Delta_0, \Delta_1, \dots, \Delta_m = \Lambda$ with the portion $t^j (at)^{-k}$ of $\partial\Delta$ left undisturbed throughout (see Figure 7). Let v_i be the maximal length arc of the boundary circuit of Δ_i contained entirely in Λ . There exists $C > 0$, depending only on \mathcal{Q}_0 such that $L := \max_i (\ell(\partial\Delta_i) - \ell(v_i)) \leq C\ell(w)$.


Proof. We consider first the case $k \geq 0$. In \mathcal{Q}_0 we find that $w^{-1} = t^j(at)^{-k} = t^{j-k}a^{-k} = t^{j-k}u_k^{-1}$, where u_k is the word of Proposition 4.1 that has length at most $12 + 4 \log_2 k$. These equalities are displayed in the left-most diagram of Figure 7, which shows the framework of the van Kampen Δ : a union of a diagram Δ' for $t^{j-k}u_k^{-1}w$, a diagram Δ'' for $(at)^{-k}u_k t^k$ and a tripod; the lower triangular region in the figure folds up to give a tripod, the exact configuration of which depends on the relative signs of j, k and $j - k$.

Note that $j - k$ is the exponent sum of the $t^{\pm 1}$ in w as \mathcal{Q}_0 retracts onto $\langle t \rangle \cong \mathbb{Z}$, and so $|j - k| \leq \ell(w)$. And $k \leq 2^{\ell(w)}$ because killing t retracts \mathcal{Q}_0 onto \mathcal{P} . It follows that $\ell(u_k) \leq 12 + 4\ell(w)$ and $\ell(t^{j-k}u_k^{-1}w) \leq 2\ell(w) + \ell(u_k) \leq 12 + 6\ell(w)$. By Proposition 4.2 we can take Δ' to be a van Kampen diagram for $t^{j-k}u_k^{-1}w$ with filling length at most a constant times $12 + 6\ell(w)$. We can cut along the edge-path in Δ labelled by $u_k t^{-(j-k)}$, leaving Δ' attached to the rest of the diagram at only one vertex, and then shell Δ' , and in the process the length of the non- $t^j(at)^{-k}$ -portion of the boundary curve has length at most a constant times $\ell(w)$.

The word $t^k(at)^{-k}a^k$ admits an obvious diagram with vertical t -corridors (as shown in Figure 8) of height $k - 1, k - 2, \dots, 1$. We cut along the vertical paths labelled by powers of a and insert copies of the diagrams Ω_i of Proposition 4.1 and their mirror images, as shown in Figure 8 (illustrated in the case $k = 6$). The shellings of Lemma 3.1 can be composed to give a shelling down to Λ that realises the asserted bound on L . \square



F 7. Shelling Δ down to Λ .



F 8. Δ'' .

For a word w we use $\ell_i(w)$ to denote the number of letters $t^{\pm 1}$ in w .

Proposition 4.4. *Suppose w is a null-homotopic word over*

$$Q := \langle a, b, r, s, t \mid a^b a^{-2}, [t, a], [r, at], [r, s], [s, t] \rangle$$

and $\ell_r(w) = \ell_s(w) = 2$ with the occurrences of $r^{\pm 1}$ and $s^{\pm 1}$ alternating in w . Then $\text{FFL}(w) \leq C\ell(w)$ where $C > 0$ depends only on Q .

Proof. In any van Kampen diagram Δ for w there is one r -corridor and one s -corridor. The condition that the occurrences of $r^{\pm 1}$ and $s^{\pm 1}$ in w alternate is equivalent to saying that these corridors must cross at least once in Δ .

Let Δ be a minimal area diagram for w . Suppose, for a contradiction, that the s -corridor and r -corridor in Δ cross more than once. Then there is a subdiagram between the two corridors with boundary word $w_0 = uv$ where $u \in \{t^{\pm 1}\}^*$ and $v \in \{(at)^{\pm 1}\}^*$. Both u and v must be reduced words as otherwise Δ would not be a reduced diagram and hence not be of minimal area. Killing r , s and t retracts Q onto a group in which a has infinite order, and so v is the empty word. Killing all the generators other than t retracts Q onto $\langle t \rangle = \mathbb{Z}$ and so u is the empty word. It follows that there are cancelling $[r, s]$ -2-cells where the two corridors cross and so Δ is not of minimal area – a contradiction.

A similar argument shows that Δ contains no r - or s -annuli.

Conclude that Δ consists of an r -corridor, an s -corridor and four subdiagrams of the form where Proposition 4.3 applies. Produce a new van Kampen diagram Δ' for w by replacing the four subdiagrams that minimise the length L of Proposition 4.3. A shelling of Δ' realising the asserted bound is obtained by running shellings of the four subdiagrams and the two corridors concurrently in the obvious way so that the diagram is eventually shelled to the $[r, s]$ -2-cell, and then to a single vertex. \square

Proposition 4.5. *Suppose w is a null-homotopic word in Q (defined above) and $\ell_s(w) = 2$. Then $\text{FFL}(w) \leq C\ell(w)$ where $C > 0$ depends only on Q .*

Proof. Let Δ be a reduced van Kampen diagram for w . We will use the layout of the r -corridors and the one s -corridor in Δ as a template for the construction of another van Kampen diagram Δ_2 for w that will admit a shelling realising the asserted bound.

Suppose C is an r -corridor in Δ that does not cross the s -corridor. The word w_C along the sides of C is in $\{(at)\}^*$ and is reduced because Δ is a reduced diagram, and so must be $(at)^k$ for some $k \in \mathbb{Z}$. Killing all defining generators other than t retracts Q onto $\langle t \rangle$. So $k \leq \ell(\lambda)$, where λ is a portion of the boundary circuit of Δ connecting the end points of a side of C .

It follows that if we remove any number of r -corridors that do not cross the s -corridor from Δ , then the length of the boundary circuit of each connected component is at most $2\ell(w)$.

Suppose we remove *all* of the r -corridors that do not cross the s -corridor from Δ . Define Δ_0 to be the connected component that contains the s -corridor. All of the other connected components have boundary words that are null-homotopic in

$$Q_0 = \langle a, b, t \mid a^b a^{-2}, [t, a] \rangle,$$

which is both a retract and a subpresentation of Q . Obtain Δ_1 from Δ by replacing all these subdiagrams by Q_0 -diagrams of minimal FFL.


Repeating the following gives a shelling of Δ_1 down to Δ_0 in the course of which the boundary circuit has length at most $C_0\ell(w)$, where C_0 is a constant that depends only on Q . Choose an r -corridor C in Δ_1 such that, of the two components we get by removing C , that which does not contain the s -corridor contains no r -corridor. Cut along one side of C

using *1-cell expansion* moves, and one *fragmentation* move. Next use *1-cell collapse* and *2-cell collapse* moves to remove the 2-cells along C . By the remarks above, both connected components have boundary circuits of length at most $2\ell(w)$. Collapse the component that does not contain the s -corridor down to a single vertex using a minimal FFL shelling, in the course of which the boundary circuit has length at most a constant times $2\ell(w)$ by Proposition 4.2. (In fact, a shelling down to Δ_0 within the required bound, can be achieved without the fragmentation move if care is taken over basepoints.)

It remains to show that the word w_0 around $\partial\Delta_0$ admits a van Kampen diagram with a full, fragmenting, free shelling in which the sum of the lengths of the boundaries of the components are at most a constant times $\ell(w_0)$. For then we can take Δ_2 to be Δ_1 with Δ_0 replaced by this diagram.

First suppose $\ell_r(w_0) = 0$. Then w_0 is null-homotopic in the retract $\mathcal{Q}_1 := \langle a, b, s, t \mid a^b a^{-2}, [t, a], [s, t] \rangle$, and the length of the s -corridor in any reduced \mathcal{Q}_1 -diagram Δ'_0 for w_0 is at most $\ell(w_0)/2$ on account of the retraction onto $\langle t \rangle$. Assume that the two components of Δ'_0 we get on removing the s -corridor are \mathcal{Q}_0 -diagrams of minimal FFL. Then we can collapse Δ'_0 by shelling each of these components and the s -corridor in turn, and using Proposition 4.2 it is easy to check that the length of the boundary circuit remains at most a constant times $\ell(w_0)$.

Next suppose $\ell_r(w_0) = 2$. Then Proposition 4.4 applies and gives us the result we need.



F 9. Shelling away one r -corridor.

Finally, suppose $\ell_r(w_0) > 2$. Then there is a subword $r^\varepsilon w_1 r^{-\varepsilon}$ in w_0 , where $\varepsilon = \pm 1$, $\ell_r(w_1) = 0$ and $\ell_s(w_1) = 1$. As $r^\varepsilon w_1 r^{-\varepsilon} = w_1$ in \mathcal{Q} , there is a \mathcal{Q} -van Kampen diagram for w_0 that we can shell by cutting along an edge-path labelled by w_1 to cut the diagram into two, as shown in Figure 9, and the shelling the two components. One of these components is a diagram for $r^\varepsilon w_1 r^{-\varepsilon} w_1^{-1}$, and this we shell first as per Proposition 4.4. The remaining component has boundary length $\ell(w_0) - 2$ and includes two fewer letters $r^{\pm 1}$, and so by continuing inductively we can find a shelling for which FFFL is at most a constant times $\ell(w_0)$. \square

Proposition 4.6. *Suppose w is a null-homotopic word in*

$$\mathcal{Q}_2 := \langle a, b, r, t \mid a^b a^{-2}, [t, a], [r, at] \rangle.$$

Then $\text{FL}(w) \leq C\ell(w)$ where $C > 0$ depends only on \mathcal{Q}_2 .

Proof. The method used in the proof of Proposition 4.4, to reduce to the case where all the r -corridors cross the s -corridor, gives this result. \square

Proposition 4.7. *Suppose w is a null-homotopic word in \mathcal{Q} . Then $\text{FFFL}(w) \leq C\ell(w)$ where $C > 0$ depends only on \mathcal{Q} .*

Proof. The cases where $\ell_s(w) = 0$ and $\ell_s(w) = 2$ are dealt with by Propositions 4.6 and 4.4, respectively. For the case $\ell_s(w) > 2$ we take a similar approach to that used in the proof of Proposition 4.4 to control $\text{FFL}(w_0)$.

There is a subword $s^\varepsilon w_1 s^{-\varepsilon}$ in w , where $\varepsilon = \pm 1$ and $s^\varepsilon w_1 s^{-\varepsilon} = w_1$ in \mathcal{Q} . So we can find a \mathcal{Q} -van Kampen diagram for w that can be severed into two components, one of which is a diagram for $s^\varepsilon w_1 s^{-\varepsilon} w_1^{-1}$, and the other of which is a diagram for a word of length $\ell(w) - 2$ that has two fewer letters $s^{\pm 1}$. The former of these two components can be shelled as per Proposition 4.4. Continuing inductively we see that the other component can be taken to be a diagram that admits a shelling in which the boundary circuit has length at most a constant times $\ell(w)$. \square

5. Q -

In this section we prove Theorem D, the quasi-isometry invariance amongst finitely presented groups of the functions FL, FFL and FFFL up to \simeq -equivalence. Our approach is to monitor how filling length, in its three guises, behaves in the standard proof that finite presentability is a quasi-isometry invariant [5, page 143]; other quantified versions of this proof exist in the literature (the first being [1], where the focus is on Area), so our exposition here will be brief.

We have quasi-isometric groups Γ and Γ' with finite presentations $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ and $\mathcal{P}' = \langle \mathcal{A}' \mid \mathcal{R}' \rangle$, respectively. So there is a quasi-isometry $f : (\Gamma, d_{\mathcal{A}}) \rightarrow (\Gamma', d_{\mathcal{A}'})$ with quasi-inverse $g : (\Gamma', d_{\mathcal{A}'}) \rightarrow (\Gamma, d_{\mathcal{A}})$. We begin by showing $\text{FL}_{\mathcal{P}} \simeq \text{FL}_{\mathcal{P}'}$.

Suppose ρ' is an edge-circuit in the Cayley graph of \mathcal{P}' , visiting vertices v_0, v_1, \dots, v_n in order with $v_0 = v_n$. Consider a circuit ρ in the Cayley graph $\text{Cay}^1(\mathcal{P})$ obtained by choosing geodesics joining $g(v_i)$ to $g(v_{i+1})$ for $i = 0, \dots, n-1$. Note that $\ell(\rho)$ is at most a constant times $\ell(\rho')$. Fill ρ with a van Kampen diagram Δ over \mathcal{P} admitting a shelling \mathcal{S} of filling length $\text{FL}_{\mathcal{P}}(\ell(\rho))$ and use f to map $\Delta^{(0)}$ to Γ' . Then join $f(a)$ to $f(b)$ by a geodesic whenever a and b are the end points of an edge in Δ . The result is a combinatorial map $\pi'_0 : \Delta'_0 \rightarrow \text{Cay}^1(\mathcal{P}')$ filling a loop ρ'_0 , where Δ'_0 is obtained from Δ by subdividing each of its edges into edge-paths of length at most some constant.

Interpolate between ρ' and ρ'_0 by joining v_i to $f(g(v_i))$ for every i , to build a map $\pi'_1 : \Delta'_1 \rightarrow \text{Cay}^1(\mathcal{P}')$ where Δ'_1 is obtained from Δ'_0 by attaching an annulus A of n 2-cells around the boundary.

One obtains a shelling \mathcal{S}'_1 of Δ'_1 down to the base vertex with $\ell(\mathcal{S}'_1)$ at most a constant times $(1 + \text{FL}_{\mathcal{P}}(\ell(\rho)))$ by first shelling away A to leave just a stalk from v_0 to $f(g(v_0))$, and then running a shelling of Δ'_0 modelled on \mathcal{S} ; whenever \mathcal{S} demands the collapse of a 1-cell in Δ , one collapses all the 1-cells in the corresponding edge-path in Δ'_0 .

Although the words labelling the 2-cells of Δ'_1 are null-homotopic they may fail to be in \mathcal{R}' , so Δ'_1 is not yet a van Kampen diagram over \mathcal{P}' . To rectify this one should replace the 2-cells of Δ'_1 by van Kampen diagrams over \mathcal{P}' , each of area at most some uniform constant. But a problem arises in that van Kampen diagrams can be singular 2-discs, so gluing them in place of 2-cells may destroy planarity. One gets around this by replacing the 2-cells of Δ'_1 one at a time in the following manner. If the boundary circuit of the 2-cell e^2 that is to be replaced is not embedded, then we focus on the innermost embedded circuit σ in the 1-skeleton that encloses a disc containing e^2 (this has length less than the boundary circuit of e^2). We delete the entire subdiagram enclosed by σ and replace it with a van Kampen for the word labelling σ . The result is a van Kampen diagram Δ'_2 for ρ' over \mathcal{P}' . We obtain a shelling \mathcal{S}'_2 for Δ'_2 by altering \mathcal{S}'_1 : each time we discarded some connected component of the set of edges inside some σ we contract it (more strictly, its pre-image)

and all the 2-cells it encloses to a single vertex in every one of the diagrams comprising the shelling, and each time we fill some 2-cell e^2 with a van Kampen diagram D we shell out all of D when we had been due to perform a 2-cell collapse move on e^2 . [This could fail to give a shelling when a 2-cell collapse move in \mathcal{S}'_1 removes a pair (e^2, e^1) , where e^1 is one of the now contracted edges – but then a 1-cell expansion followed by a 2-cell collapse producing the same effect can be used instead.] The difference between $\ell(\mathcal{S}'_1)$ and $\ell(\mathcal{S}'_2)$ is then at most some additive constant. Deduce that $\text{FL}_{\mathcal{P}'} \leq \text{FL}_{\mathcal{P}}$. Interchanging \mathcal{P} and \mathcal{P}' we have $\text{FL}_{\mathcal{P}} \leq \text{FL}_{\mathcal{P}'}$ and so $\text{FL}_{\mathcal{P}} \simeq \text{FL}_{\mathcal{P}'}$.

The proof that $\text{FFL}_{\mathcal{P}} \simeq \text{FFL}_{\mathcal{P}'}$ is essentially the same, except we consider free shellings, we discard the stalk between v_0 and $f(g(v_0))$, and we replace $\text{FL}_{\mathcal{P}}(\ell(\rho))$ by $\text{FFL}_{\mathcal{P}}(\ell(\rho))$. To show $\text{FFFL}_{\mathcal{P}} \simeq \text{FFFL}_{\mathcal{P}'}$ we additionally allow free and fragmenting shellings and we use $\text{FFFL}_{\mathcal{P}}(\ell(\rho))$ in place of $\text{FL}_{\mathcal{P}}(\ell(\rho))$; no further technical concerns arise.

6. R

Suppose $c : [0, 1] \rightarrow X$ is a loop in a metric space X .

A *based null-homotopy* H of c is a continuous map $H : [0, 1]^2 \rightarrow X$ for which $H(0, t) = H(1, t) = c(0)$ for all t and, defining $H_t : [0, 1] \rightarrow X$ by $H_t(s) = H(s, t)$, we have $H_0 = c$ and $H_1(s) = c(0)$ for all s .

A *free null-homotopy* H of c is a continuous map $H : [0, 1]^2 \rightarrow X$ such that $H(0, t) = H(1, t)$ for all t , and H_0 and H_1 are c and a constant function, respectively.

Let \mathcal{S} be the set of subspaces S of $[0, 1]^2$ that have $[0, 1] \times \{0\} \subset S$ and are the union of a finite family of closed triangular discs $\Delta_i = [(a_i, 0), (b_i, 0), (c_i, d_i)]$ with $c_i \in [a_i, b_i]$ and $d_i \in (0, 1]$. The fibres S_t of the projection mapping points in S to their second co-ordinate are empty or are disjoint unions $\bigsqcup_{i=1}^{k_t} I_{t,i} \times \{t\}$ of closed intervals $I_{t,i}$. At finitely many critical t -values τ_t , some of the intervals comprising the fibre bifurcate or collapse and vanish.

A *free and fragmenting null-homotopy* H of c is a continuous map $H : S \rightarrow X$ for some $S \in \mathcal{S}$ where, defining H_t to be the restriction of H to S_t , we find $H_0 = c$ and $H_t(x) = H_t(y)$ for all t whenever x and y are the end points of some $I_{t,i}$. We define $\ell(H_t)$ to be the sum of the lengths of the k_t loops in X defined by H_t . Taking S to be a single triangle reduces to the case of a free null-homotopy.

In each of the three settings above define $\ell(H) := \sup_{t \in [0, 1]} \ell(H_t)$, and then

$$\begin{aligned} \text{FL}(c) &:= \inf \{ \ell(H) \mid \text{based null-homotopies } H \text{ of } c \} \\ \text{FFL}(c) &:= \inf \{ \ell(H) \mid \text{free null-homotopies } H \text{ of } c \} \\ \text{FFFL}(c) &:= \inf \{ \ell(H) \mid \text{free and fragmenting null-homotopies } H \text{ of } c \}. \end{aligned}$$

For $M = \text{FL}, \text{FFL}$ or FFFL , define $M_X : [0, \infty) \rightarrow [0, \infty)$ by

$$M_X(l) = \sup \{ M(c) \mid \text{null-homotopic loops } c \text{ with } \ell(c) \leq l \}.$$

The following lemma gives sufficient conditions for $\text{FL}_X, \text{FFL}_X$ and FFFL_X to be well-defined — conditions enjoyed by the universal cover of any closed connected Riemannian manifold, for example.

Lemma 6.1. *Suppose X is the universal cover of a compact geodesic space Y for which there exist $\mu, L > 0$ such that every loop of length less than μ admits a based null-homotopy of filling length less than L . Then $\text{FL}_X, \text{FFL}_X$ and FFFL_X are well-defined functions $[0, \infty) \rightarrow [0, \infty)$.*

Proof. The proof of Lemma 2.2 in [6] can readily be adapted to this context. In brief, we first show that every rectifiable loop c in X admits a based null-homotopy with finite

filling length – apply a compactness argument to an arbitrary based null-homotopy for c to partition c into finitely many loops of length at most μ ; by hypothesis each such loop has finite filling length and it follows that c has finite filling length.

Next suppose c has length l and assume (by reducing μ if necessary) that balls of radius μ in Y lift to X . Cover Y with a maximal collection of disjoint closed balls of radius $\mu/5 > 0$; let $\Lambda \subset X$ be the set of lifts of their centres. Subdivide c into $m \leq 1 + 10\ell/\mu$ arcs with end-points v_i , each of length at most $\mu/10$; each v_i lies within $\mu/5$ of some $u_i \in \Lambda$; form a piecewise geodesic loop c' approximating c by connecting-up these u_i . Loops made up of the portion of c from v_i to v_{i+1} and geodesics $[u_i, u_{i+1}]$, $[u_i, v_i]$ and $[u_{i+1}, v_{i+1}]$ have length at most μ . Homotopy discs for these loops together form a *collar* between c and c' . By passing across these discs one at a time, it is possible to homotop c across the collar to a loop made up of c' and a stalk of length $\mu/5$ from $c(0)$ to u_0 , encountering loops only of length at most a constant (depending on L and μ) times l en route. Modulo the action of $\pi_1 Y$, there are only finitely many such piecewise geodesic loops such as c' and, by our earlier argument, each one admits a filling of finite filling length. It follows that FL_X , and hence FFL_X and FFFL_X , are well-defined functions. \square

Proof of Theorem E. Fix a basepoint $p \in X$. Define a quasi-isometry Φ mapping the Cayley graph of $\mathcal{P} = \langle \mathcal{A} \mid \mathcal{R} \rangle$ to X by choosing a geodesic from p to its translate $a \cdot p$ for each $a \in \mathcal{A}$, and then extending equivariantly. Let Ψ be a quasi-isometry from X to Γ sending $x \in X$ to some γ such that $\gamma.p$ is a point of $\Gamma.p$ closest to x .

A path in X is called *word-like* (following [4]) if it is the image in X of an edge-path in the Cayley graph. For each $r \in \mathcal{R}$, let c_r denote the word-like loop in X , based at p that is the image of an edge-circuit in the Cayley graph labelled r . Map the Cayley 2-complex of \mathcal{P} to X by choosing a disc-filling arising from a based null-homotopy of finite filling length for each c_r .

We will show first that $\text{FL}_X \leq \text{FL}_{\mathcal{P}}$, $\text{FFL}_X \leq \text{FFL}_{\mathcal{P}}$ and $\text{FFFL}_X \leq \text{FFFL}_{\mathcal{P}}$. As in the proof of Lemma 6.1, a collar between an arbitrary rectifiable loop c in X and a word-like loop c' , can be used to show there is no change in the \simeq classes of FL_X , FFL_X or FFFL_X if one takes the suprema in their definitions to be over fillings of word-like loops only: for FL_X one notes that c can be homotoped across the collar to a loop based at $c(0)$ that is obtained from c' by attaching a stalk from $c(0)$ to $c'(0)$, and one need pass through loops of length no more than $C\ell(c) + C$ en route, where C is a constant independent of c ; for FFL_X and FFFL_X , the stalk is abandoned and the homotopy is between c and c' .

One gets an upper bound on the filling length of a word-like loop c in X by taking the image in X of a minimal filling length van Kampen diagram Δ . The progress of the boundary circuit in the course of a shelling of Δ dictates a sequence of stages in a null-homotopy of c . Using Lemma 6.1, we can interpolate between these stages in a way that increases the length of the curve by no more than an additive constant, and so we get $\text{FL}_X \leq \text{FL}_{\mathcal{P}}$. The proof that $\text{FFL}_X \leq \text{FFL}_{\mathcal{P}}$ and $\text{FFFL}_X \leq \text{FFFL}_{\mathcal{P}}$ can be completed likewise.

Now we address $\text{FL}_{\mathcal{P}} \leq \text{FL}_X$. Consider a word-like loop $c : [0, 1] \rightarrow X$ corresponding to a null-homotopic word w over \mathcal{P} of length n . Fix a constant $\lambda > \max_{a \in \mathcal{A}} d_X(p, a.p)$. Then $\ell(c) \leq \lambda n$. Let $H : [0, 1]^2 \rightarrow X$ be a based null-homotopy of c with filling length at most $1 + \text{FL}_X(\lambda n)$. By uniform continuity, there exists $\varepsilon > 0$ such that $\varepsilon^{-1} \in \mathbb{Z}$ and $d_X(H(a), H(b)) \leq 1$ for all $a, b \in [0, 1]^2$ with $d_{\mathbb{B}^2}(a, b) \leq \varepsilon$.

Subdivide $[0, 1]^2$ into ε^{-1} rectangles separated by the lines $t = t_j$ where $t_j = j\varepsilon$ and $j = 0, 1, \dots, \varepsilon^{-1}$. For all such $t = t_j$, take $0 = s_{t,0} < s_{t,1} < \dots < s_{t,k_t} = 1$ in such a way that

for all i , the restriction of H_t to $[s_{t,i}, s_{t,i+1}]$ is an arc of length at most λ and $k_t \leq 1 + \ell(H_t)/\lambda$. Mark the points $s_{t,0}, \dots, s_{t,k_t}$ on each of the lines $t = t_j$. Then, for all $j = 0, 1, \dots, \varepsilon^{-1} - 1$ and all $i = 1, 2, \dots, k_{t_j} - 1$, join $(s_{t,i}, t_j)$ to $(s_{t,i'}, t_{j+1})$ by a straight-line segment where $(s_{t,i'}, t_{j+1})$ is the first marked point reached from $(s_{t,i}, t_{j+1})$ by increasing the s -coordinate. Note that

$$(3) \quad d(H_{t_j}(s_{t,i}), H_{t_{j+1}}(s_{t,i'})) \leq 1 + \lambda.$$

In the same way, for all $j = 1, \dots, \varepsilon^{-1}$ and for all $i = 1, \dots, k_{t_j} - 1$ such that $(s_{t,i}, t_j)$ is not the terminal vertex of one of the edges we just connected, join $(s_{t,i}, t_j)$ to some $(x_{t,i'}, t_{j-1})$ for which

$$(4) \quad d(H_{t_j}(s_{t,i}), H_{t_{j-1}}(x_{t,i'})) \leq 1 + \lambda,$$

so as to produce a diagram Δ in which every 2-cell has boundary circuit of combinatorial length at most 4.

Orient every edge e of Δ arbitrarily and define

$$g_e := \Psi(H_{t'}(s'))\Psi(H_t(s))^{-1}$$

when the initial and terminal points of e are (s, t) and (s', t') , respectively. It follows from (3) and (4) that in the word metric associated to \mathcal{A} the distance from 1 to g_e , denoted $|g_e|$, is at most a constant $K = K(\mathcal{A})$. Subdivide e into a path of $|g_e|$ edges; give each of these new edges an orientation and a labelling by a letter in \mathcal{A} so that one reads a word representing g_e along the path. Make all the choices in the construction above in such a way that w labels the line $t = 0$ and all the other edges in $\partial\Delta$ are labelled by e .

The shelling (Δ_i) of Δ which strips away the rectangles from left to right, shelling each in turn from top to bottom, has

$$\max_i \ell(\Delta_i) \leq K(1 + \text{FL}_X(\lambda n)) + 4K.$$

Let w_i be the word one reads around the boundary circuit of Δ_i . Each 2-cell in each Δ_i has boundary circuit labelled by a null-homotopic word that may not be in \mathcal{R} , but has length at most $4K$. So it is possible to interpolate between the w_i to produce a null-sequence (as in Section 2) for $w = w_0$ with respect to \mathcal{P} ; in this sequence every word has length at most $K(1 + \text{FL}_X(\lambda n)) + 4K$ plus a universal constant. Thus $\text{FFL}_{\mathcal{P}} \leq \text{FL}_X$, as required.

That $\text{FFL}_{\mathcal{P}} \leq \text{FFL}_X$ can be proved in the same way. The argument needs to be developed further to show that $\text{FFFL}_{\mathcal{P}} \leq \text{FFFL}_X$. Given the word-like loop c , one takes a free and fragmenting null-homotopy $H : S \rightarrow X$ of c with $\ell(H)$ at most $\text{FFFL}_X(\ell(c)) + 1$ and with the property that whenever loops of length less than some prior fixed constant appear, those loops are contracted to points before any further bifurcations occur. This implies that for all t , the number of connected components k_t in the fibre $S_t = \bigsqcup_{i=1}^{k_t} I_{t,i} \times \{t\}$ is at most a constant times $(1 + \ell(H_t))$. For the construction of Δ we inscribe S with the arcs of its intersection with the lines $t = t_j$ and with the additional lines $t = \tau_l$, where τ_l are the critical t -values of H . Using H_t and Ψ as before, we subdivide the fibre S_t into edges and label each of its k_t connected components by a null-homotopic word – this works as before, except we additionally insist that the end points of the closed intervals $I_{t,i}$ comprising S_t be included amongst the s_i – this may add k_t to the total length the words along S_t , but the argument given above ensures that this additional cost is no more than a constant times $(1 + \ell(H_t))$. \square

R

- [1] J. M. Alonso. Inégalités isopérimétriques et quasi-isométries. *C.R. Acad. Sci. Paris Ser. I Math.*, 311:761–764, 1990.
- [2] G. Baumslag, S. M. Gersten, M. Shapiro, and H. Short. Automatic groups and amalgams. *J. Pure Appl. Algebra*, 76(3):229–316, 1991.
- [3] M. R. Bridson. Asymptotic cones and polynomial isoperimetric inequalities. *Topology*, 38(3):543–554, 1999.
- [4] 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.
- [5] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-positive Curvature*. Number 319 in Grundlehren der mathematischen Wissenschaften. Springer Verlag, 1999.
- [6] M. R. Bridson and T. R. Riley. Extrinsic versus intrinsic diameter for Riemannian filling-discs and van Kampen diagrams. Preprint, [arXiv:math.GR/0511004](https://arxiv.org/abs/math/0511004).
- [7] D. E. Cohen. Isoperimetric and isodiametric inequalities for group presentations. *Int. J. of Alg. and Comp.*, 1(3):315–320, 1991.
- [8] S. M. Gersten. The double exponential theorem for isoperimetric and isodiametric functions. *Int. J. of Alg. and Comp.*, 1(3):321–327, 1991.
- [9] S. M. Gersten. Finiteness properties of asynchronously automatic groups. In R. Charney, M. Davis, and M. Shapiro, editors, *Geometric Group Theory*, volume 3 of *Ohio State University, Mathematical Research Institute Publications*, pages 121–133. de Gruyter, 1995.
- [10] S. M. Gersten, D. F. Holt, and T. R. Riley. Isoperimetric functions for nilpotent groups. *GAF*, 13:795–814, 2003.
- [11] S. M. Gersten and T. R. Riley. Filling length in finitely presentable groups. *Geom. Dedicata*, 92:41–58, 2002.
- [12] S. M. Gersten and T. R. Riley. Some duality conjectures for finite graphs and their group theoretic consequences. *Proc. Edin. Math. Soc.*, 48(2):389–421, 2005.
- [13] 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.

M R. B

Mathematics, Huxley Building, Imperial College, London, SW7 2AZ, UK
m.bridson@imperial.ac.uk, <http://www.ma.ic.ac.uk/~mbrids/>

T R. R

Mathematics Department, 310 Malott Hall, Cornell University, Ithaca, NY 14853-4201, USA
tim.riley@math.cornell.edu, <http://www.math.cornell.edu/~riley/>