5-2023

Intrinsic Tame Filling Functions and Other Refinements of Diameter Functions

Andrew Quaisley
University of Nebraska-Lincoln, andrew.quaisley@gmail.com

INTRINSIC TAME FILLING FUNCTIONS AND OTHER REFINEMENTS OF DIAMETER FUNCTIONS

by

Andrew Quaisley

A DISSERTATION

Presented to the Faculty of

The Graduate College at the University of Nebraska

In Partial Fulfilment of Requirements

For the Degree of Doctor of Philosophy

Major: Mathematics

Under the Supervision of Professors Susan Hermiller and Mark Brittenham

Lincoln, Nebraska

May, 2023

INTRINSIC TAME FILLING FUNCTIONS AND OTHER REFINEMENTS OF DIAMETER FUNCTIONS

Andrew Quaisley, Ph.D.

University of Nebraska, 2023

Adviser: Susan Hermiller and Mark Brittenham

Tame filling functions are quasi-isometry invariants that are refinements of the diameter function of a group. Although tame filling functions were defined in part to provide a proper refinement of the diameter function, we show that every finite presentation of a group has an intrinsic tame filling function that is equivalent to its intrinsic diameter function. We then introduce some alternative filling functions—based on concepts similar to those used to define intrinsic tame filling functions—that are potential proper refinements of the intrinsic diameter function.

ACKNOWLEDGMENTS

I'd like to begin by thanking my advisors, Susan and Mark, for all of their work, guidance, and support over the past five years. Their knowledge and expertice were obviously invaluable to my research, but I am even more grateful for their patience and kindness while I learned to navigate the challenges of grad school, math research, and preparation for life after graduation.

I'd also like to thank all of my friends from the math department for all the relief they've provided from the stress of grad school in the form of game and movie nights, interesting conversations, and hugs. Special thanks go to George, Emily, Stephen, and John, for helping me take care of myself through my struggles with mental health. Allan Donsig and Nathan Wakefield also deserve big thank-yous for their patience, understanding, and support. I am also eternally grateful to everyone who ever covered one of my classes on basically zero notice.

Finally, I'd like to thank my family for all of their love and support. I never questioned my ability to continue my pursuit of math as far as I wanted to take it, in large part due to my parents' continuous support and encouragement. I am also deeply grateful to my sister for convincing me early on that "weird" is a compliment and that nerds are awesome. She was my first and most important role-model in the fields of math, learning, and curiosity, among many others. She has also been a life-long source of love, acceptance, and support.

Of course, I can't leave out one of the newest members of my family, my partner Kelsey. Her daily injections of joy, fun, and silliness into my life have been absolutely necessary for my mental stability during grad school. And her presence, attention, and emotional support have gotten me through many difficult days. She has somehow single-handedly made grad school the happiest time of my life.

Table of Contents

List of Figures

Chapter 1

Introduction

1.1 Background and Motivation

Given a finite presentation for a group and a word w in the generators equal to the identity, a van Kampen diagram for w is a planar 2-complex that demonstrates how the relators of the presentation can reduce w to the empty word. As such, van Kampen diagrams serve as one connection between combinatorial and geometric group theory; geometric properties of van Kampen diagrams of a presentation can have implications for combinatorial properties of the presentation, and vice-versa. See Section 2 of this chapter for the formal definition.

One way to study properties of groups via their van Kampen diagrams is through filling functions. These are functions that measure the growth of a given property of van Kampen diagrams with respect to the length of the word that the diagram represents. More formally, suppose M is a function from van Kampen diagrams to $\mathbb{N}$, called a diagram measurement. Then we can define a filling function $\mathcal{M} : \mathbb{N} \to \mathbb{N}$ in the following way. For each word $w =_G 1$, let Δ_w be a van Kampen diagram for w such that $M(\Delta_w)$ is minimal over all van Kampen diagrams for w. Then for all $n \in \mathbb{N}$, define $\mathcal{M}(n)$ to be the maximum $M(\Delta_w)$ over all words $w =_G 1$ with length at most n. Intuitively, $\mathcal{M}$ measures how quickly van Kampen diagrams grow with respect to the measurement M.

The definition of a filling function implicitly depends on the presentation of the group, since different presentations yield different sets of words and van Kampen diagrams. However, under an appropriate equivalence relation that distinguishes between functions that grow qualitatively differently (see Definition 1.2.2), the equivalence class of a filling function may be a group invariant and tell us something about the group as a whole. Many well-studied filling functions are also invariant over groups that have "similar geometry", formalized in the notion of *quasi-isometric* groups. These invariants often tell us something about the complexity of the word problem for the group. For an introduction to filling functions, see [3, Chapter II] by Riley.

One well-studied filling function is the intrinsic diameter (or isodiametric) function, which measures the diameter of van Kampen diagrams. The intrinsic diameter function is one of several filling functions that is recursive if and only if the word problem of the group is solvable (see [16, Th. 2.1]). As a result, there are groups whose intrinsic diameter functions grow faster than any recursive function. It is also known that the collection of intrinsic diameter functions includes a variety of different kinds of growth. For example, [26, Cor. 1.4] shows that this collection includes a wide variety of power functions.

The intrinsic diameter function is intrinsic in the sense that the diameter is measured using the path metric on the 1-skeleton of the van Kampen diagram itself. Replacing this metric with the path metric of the Cayley graph gives an extrinsic diameter function, which is bounded above by the intrinsic diameter function. Extrinsic diameter functions were introduced by Bridson and Riley in [5], where they also showed that there can exist arbitrarily large polynomial gaps between intrinsic and extrinsic diameter functions: for any $\alpha > 0$, there is a finite presentation with intrinsic diameter function bounded below by a polynomial p and extrinsic diameter function bounded above by a polynomial q, with $\deg(p) - \deg(q) = \alpha$.

However, many groups are known to have at most linear intrinsic—and therefore also at most linear extrinsic—diameter functions, meaning that the equivalence class of the diameter functions cannot distinguish between these groups. Among them are all combable groups and almost convex groups [16, Prop. 3.4, Cor. 4.3], and therefore all automatic groups [14, Lemma 2.3.2]. These classes include word hyperbolic groups [14, Th. 3.4.5], groups acting geometrically on CAT(0) cube complexes [24], and Euclidean groups [14, Cor. 4.1.6], [9, Th. 2.5]; Coxeter groups [7, 11]; Artin groups of finite type [10] and sufficiently large type [18]; and mapping class groups of surfaces of finite type [23], among others. The fundamental groups of all closed 3-manifolds and 3-manifolds with toroidal boundary also have at most linear diameter functions [16, Th. 3.5]. The class of groups with at most linear diameter functions is also closed under graph products [20].

In part because such a large class of groups have linear diameter functions, Brittenham and Hermiller defined new quasi-isometry invariants in [8] called intrinsic and extrinsic tame filling functions. The following is an informal description of tame filling functions; see Section 2 for formal definitions. Tame filling functions are defined using a *filling* of a presentation—a choice of a van Kampen diagram for each word over the generators that is equal to the identity in the group—and a *1-combing* for each of these van Kampen diagrams, which amounts to a choice of continuously-varying paths from the basepoint of a van Kampen diagram to its boundary. Given a point on one of these paths a distance n from the basepoint of the diagram, a tame filling function f mandates that no prior point on the path can be a distance more than $f(n)$ from the basepoint; a 1-combing satisfying this property is called *f-tame*. Informally, a van Kampen diagram can be considered a terrain where elevation is represented by distance from the basepoint. A group with a slow-growing tame filling function is intended to have van Kampen diagrams that are relatively smooth inclines, extending steadily outward from the basepoint to the boundary. A group with a fast-growing tame filling function, on the other hand, may require van Kampen diagrams with tall mountains and deep valleys.

Brittenham and Hermiller showed that any intrinsic or extrinsic tame filling function for a presentation grows at least as fast as the respective diameter function, leaving open the question of whether or not tame filling functions ever grow strictly faster than their respective diameter functions, especially in the case of groups with linear diameter functions. In fact, it was not even known if every group admits a finite-valued intrinsic or extrinsic tame filling function, whereas every group admits a finite-valued intrinsic and extrinsic diameter function by definition. The main purpose of this dissertation is to prove the following theorem in Chapter 2, and thereby resolve these questions for intrinsic tame filling functions:

Theorem 2.4.4. *Given a finite presentation $\mathcal{P} = \langle A|R\rangle$ such that for all $a \in A$, a is not equal to the identity, there is an intrinsic tame filling function for $\mathcal{P}$ that is equivalent to the intrinsic diameter function of $\mathcal{P}$.*

As a result, all groups admit a finite-valued intrinsic tame filling function. This was proven concurrently by Nu'Man and Riley; they showed that, given a finite presentation, if $D : \mathbb{N} \to \mathbb{N}$ is the corresponding Dehn function, then $f(n) = D(e^n)$ is an intrinsic tame filling function for the presentation [25]. However, since the Dehn function of a presentation grows at least as fast as, and often strictly faster than, the intrinsic diameter function (up to the equivalence relation defined in Section 2), Theorem 2.4.4 provides a tighter bound on the growth of minimal intrinsic tame filling functions.

Beyond that, Theorem 2.4.4 strengthens the notion of intrinsic diameter. If f is the intrinsic diameter function of a presentation, then by definition, for any word w of length n, there exists a van Kampen diagram for w with intrinsic diameter at most $f(n)$. However, Theorem 2.4.4 implies the existence of an alternative van Kampen diagram for w that is g-tame, for some function g equivalent to f. This provides an additional option with more structure than the diameter restriction, and that may prove more useful for some purposes.

Theorem 2.4.4 also shows that intrinsic tame filling functions do not distinguish between groups with equivalent intrinsic diameter functions, as they were defined, in part, to do. This motivates us to define several new filling functions in Chapters 3 and 4 that may fulfill the desired purpose.

In Chapter 3, we focus on the metaphor of van Kampen diagrams as terrains with elevation represented by distance from the basepoint, and attempt to define filling functions that are more sensitive to the "hilliness" of the terrain than intrinsic tame filling functions are. We define a notion of "contour lines" on a van Kampen diagram, viewing the diagram as a kind of topographic map, and use them to define the aggregate variation of a van Kampen diagram—a diagram measurement that takes into account both the height and number of hills in the diagram. We name the corresponding filling functions aggregate variation functions. We then use the solvable Baumslag-Solitar groups as examples to test out intrinsic aggregate variation functions and consider whether or not they are equivalent to intrinsic diameter functions. Motivated in part by these examples, we are able to prove the following theorem.

Theorem 3.2.3. *If G is a group with finite presentation $\mathcal{P}$, then there is a presentation $\mathcal{P}'$ that is $\mathcal{P}$ with finitely many relators added such that the intrinsic diameter function of $\mathcal{P}$ is an intrinsic aggregate variation function for $\mathcal{P}'$.*

However, we note that the proof of Theorem 3.2.3 requires us to use unreduced van Kampen diagrams (diagrams with adjacent 2-cells that are mirror images of each other, and that can therefore be removed to simplify the diagram), as well as use relators that may not be cyclically reduced. When we consider the standard, reduced van Kampen diagrams for the standard presentations of the solvable Baumslag-Solitar groups, we are able to prove Proposition 3.3.1: that these standard diagrams have intrinsic aggregate variation that grows faster than the intrinsic diameter function of the group. This suggests that we may be on the right track towards finding a proper refinement of intrinsic diameter functions if we require reduced diagrams and relators.

In Chapter 4, we define another filling function, called a subdiagram diameter function, which can be viewed as a variation on tame filling functions that breaks the proof of Theorem 2.4.4. If f is a subdiagram diameter function for a presentation, then each word representing the identity has a van Kampen diagram such that every subdiagram has diameter bounded by f. By definition, this filling function is bounded below by the corresponding diameter function, and we also show in Proposition 4.2.1 that it is a quasi-isometry invariant. We leave it to future work to determine whether or not subdiagram diameter functions can distinguish between groups with the same corresponding diameter functions.

The questions resolved for intrinsic tame filling functions in this dissertation are still open in the case of extrinsic tame filling functions; it is quite possible that there exists a group whose extrinsic tame filling functions all grow strictly faster than its extrinsic diameter function. This possibility is supported by a conjecture of Tschantz—that there exists a finitely presented group that is not tame combable [8, p. 3]. Brittenham and Hermiller showed that such a group would not even admit a finite-valued extrinsic tame filling function, and therefore it would not have an extrinsic tame filling function that is equivalent to its extrinsic diameter function.

An argument similar to the one used to prove Theorem 2.4.4 would not be sufficient to answer these questions. See the remark at the end of Chapter 2 for a brief explanation.

1.2 Preliminary Definitions and Notation

Most of the definitions from this section are taken from [3, Chapter II] and [8].

Throughout this dissertation, let $\mathcal{P} = \langle A|R\rangle$ be a finite presentation of a group G, with no generator in A representing the identity of G. Given a word $w \in (A \cup A^{-1})^*$, let $\ell(w)$ denote the length of w. Given $v \in (A \cup A^{-1})^*$, we will write $w =_G v$ if w and v represent the same element of G, and similarly for $g \in G$ we will write $w =_G g$ if w represents g in G. Given an edge path γ in a 2-complex, let $|\gamma|$ denote the length of γ, that is, the number of edges in γ. Given $g_1, g_2 \in G$, let $d_G(g_1, g_2)$ be the length of the shortest edge path in the Cayley graph of $\mathcal{P}$ between the 0-cells corresponding to g_1 and g_2; equivalently, $d_G(g_1, g_2)$ is the length of the shortest word in $(A \cup A^{-1})^*$ representing $g_2 g_1^{-1}$. Given a path $\gamma : [0, 1] \to X$ into any space X, let $\overline{\gamma}$ denote the path traveled in the opposite direction, i.e. $\overline{\gamma}(t) = \gamma(t - 1)$ for all $t \in [0, 1]$.

Definition 1.2.1. Let G and H be groups. A function $\phi : G \to H$ is a *quasi-isometry* if there exists a constant $k > 0$ such that

- for all $g_1, g_2 \in G$, $\frac{1}{k} d_G(g_1, g_2) - k \leq d_H(\phi(g_1), \phi(g_2)) \leq k d_G(g_1, g_2) + k$, and
- for all $h \in H$ there exists $g_3 \in G$ such that $d_H(h, \phi(g_3)) \leq k$.

Note that, if there is an quasi-isometry from G to H, then there is also a quasi-isometry from H to G given by $h \mapsto g_3$, where g_3 is any element of G with $d_H(h, \phi(g_3)) \leq k$, as guaranteed to exist by the definition above. So G and H are called *quasi-isometric* if a quasi-isometry between them exists. A group invariant is a *quasi-isometry invariant* if it has the same value for any two groups that are quasi-isometric.

Definition 1.2.2. Let $S, T \subseteq \mathbb{N}[\frac{1}{4}]$ and let $f : S \to [0, \infty)$ and $g : T \to [0, \infty)$ be functions. Write $f \preceq g$ if there exist constants $A, B, C, D, E \geq 0$ such that for all $s \in S$, $f(s) \leq Ag(Bt + C) + Dt + E$, where $t = \max([0, s] \cap T)$. If $f \preceq g$ and $g \preceq f$, write $f \simeq g$ and say that f and g are *equivalent*. If P is a property of functions and g has P, say that f *is at most* P if $f \preceq g$. For example, f is at most linear if $f \preceq g$ and g is linear.

Note that this notion of equivalence can distinguish between, for example, linear, quadratic, polynomial, and exponential growth, among others. Using this equivalence relation, the equivalence class of many well-studied filling functions is a quasi-isometry invariant (up to adding some additional relators to the presentation in some cases); see [3, Chapter II] for details. Although all the filling functions listed in [3, Chapter II] have a domain of $\mathbb{N}$, the above definition includes functions with domains in $\mathbb{N}[\frac{1}{4}]$ specifically in order to compare diameter functions with tame filling functions, defined below.

The following definition of van Kampen diagram is taken from [6].

Given CW-complexes C_1 and C_2, a continuous map $f : C_1 \to C_2$ is *combinatorial* if its restriction to each open cell of C_1 is a homeomorphism onto an open cell of C_2.

A *combinatorial complex* is a CW-complex C with the following restriction on its attaching maps. For each n-complex σ of C, there is an $(n-1)$-dimensional combinatorial complex S_σ, a homeomorphism $h : \partial D^n \to S_\sigma$, and a combinatorial map $f : S_\sigma \to C^{(n-1)}$ such that $\psi_\sigma = f \circ h$ is σ's attaching map.

A *singular disk diagram* D is a compact, contractible, combinatorial 2-complex embedded in $\mathbb{R}^2$. Each 1-cell has associated to it two directed edges. A *boundary circuit* of D is a directed edge circuit γ starting and ending at any 0-cell $*$ on the boundary of D, constructed by traveling around the boundary counterclockwise until every directed edge associated to a 1-cell in the boundary has been used. Note that γ may not be unique, since if $*$ is a cut vertex there will be multiple directed edges that will serve to begin the circuit.

Given a word $w \in (A \cup A^{-1})^*$, a *van Kampen diagram* Δ for w with respect to the presentation $\langle A|R\rangle$ is a singular disk diagram that has a specified basepoint $* \in \Delta^{(0)} \cap \partial\Delta$, and that has every directed edge labeled by an element of $A \cup A^{-1}$ in the following way. For each 1-cell of Δ, one associated directed edge is labeled by an element $a \in A$ and the other is labeled by a^{-1}, such that a boundary circuit of each 2-cell is labeled by an element of R or R^{-1}, and such that a boundary circuit of Δ starting at $*$ is labeled by w. We will also require that, for each 2-cell σ of Δ, the combinatorial map $f : S_\sigma \to \Delta^{(n-1)}$ from the definition of combinatorial complex is a homeomorphism, implying that σ's attaching map is a homeomorphism.

Let X be the Cayley 2-complex of the presentation $\mathcal{P}$. Then for every van Kampen diagram Δ with respect to $\mathcal{P}$ there is a combinatorial map $\pi_\Delta : \Delta \to X$ with $\pi_\Delta(*) = 1_G$ that preserves the directions and labels of these edges.

Given a van Kampen diagram Δ, define the *coarse distance function* $d_\Delta : \Delta^{(0)} \times \Delta \to \mathbb{N}[\frac{1}{4}]$ in the following way. Let $v \in \Delta^{(0)}$ and $x \in \Delta$. If $x \in \Delta^{(0)}$, let $d_\Delta(v, x)$ be the length of the shortest edge path between v and x. If $x \in \Delta^{(1)} \setminus \Delta^{(0)}$, then let p and q be the endpoints of the edge on which x lies and let

$$d_\Delta(v, x) = \min(d_\Delta(v, p), d_\Delta(v, q)) + \frac{1}{2}.$$

If $x \in \Delta \setminus \Delta^{(1)}$, then let $e_1, \ldots, e_n$ be the 1-cells on the boundary of the 2-cell in which x lies and for each $i \in [n]$ choose some $y_i \in e_i$. Then let

$$d_\Delta(v, x) = \max\{d_\Delta(v, y_i) | i \in [n]\} - \frac{1}{4}.$$

If C is a 2-complex and $x, y \in C^{(0)}$, then define a *C-geodesic* from x to y to be an edge path from x to y in $C^{(1)}$ of length $d_C(x, y)$. If the complex being referred to is clear, we will just call such a path a geodesic. If T is a tree in C, and $* \in C^{(0)} \cap T$, we will say that T is a *tree of C-geodesics out of* $*$ if, for every $x \in T^{(0)}$, the unique simple edge path from $*$ to x in T is a C-geodesic. If $D \subseteq C$, we will write $E(D)$ for the set of 1-cells of C contained in D.

Given a van Kampen diagram Δ with basepoint $*$, the *intrinsic diameter* of Δ is

$$\mathrm{IDiam}(\Delta) = \max\{d_\Delta(*, x) | x \in \Delta^{(0)}\},$$

and the *extrinsic diameter* of Δ is

$$\mathrm{EDiam}(\Delta) = \max\{d_X(1_G, \pi_\Delta(x)) | x \in \Delta^{(0)}\}.$$

Given $w \in (A \cup A^{-1})^*$ with $w =_G 1$, the *intrinsic diameter* of w is

$$\mathrm{IDiam}(w) = \min\{\mathrm{IDiam}(\Delta) | \Delta \text{ is a van Kampen diagram for } w \text{ with basepoint } *\},$$

and the *extrinsic diameter* of w is

$$\mathrm{EDiam}(w) = \min\{\mathrm{EDiam}(\Delta) | \Delta \text{ is a van Kampen diagram for } w \text{ with basepoint } *\}.$$

Definition 1.2.3. The *intrinsic diameter function* of $\langle A | R \rangle$ is the function IDiam : $\mathbb{N} \to \mathbb{N}$ given by

$$\mathrm{IDiam}(n) = \max\{\mathrm{IDiam}(w) | w \in (A \cup A^{-1})^* \text{ with } \ell(w) \leq n \text{ and } w =_G 1\},$$

and the *extrinsic diameter function* of $\langle A|R\rangle$ is the function EDiam : $\mathbb{N} \to \mathbb{N}$ given by

$$\mathrm{EDiam}(n) = \max\{\mathrm{EDiam}(w) | w \in (A \cup A^{-1})^* \text{ with } \ell(w) \leq n \text{ and } w =_G 1\}.$$

Note that in [3, Chapter II], these diameter functions are defined using distance between arbitrary vertices, rather than fixing one of them as the basepoint, and the function defined above is referred to as the *based* intrinsic diameter. However, these two definitions give equivalent functions, and the above definition is more convenient for the purposes of this work.

The following definitions of 1-combing and f-tame were introduced by Mihalik and Tschantz in [21]:

Given a 2-complex C with basepoint $*$ and a subcomplex $D \subseteq C^{(1)}$ containing $*$, a *1-combing* of the pair (C, D) based at $*$ is a continuous function $\Psi : D \times I \to C$ such that

- for all $t \in I$, $\Psi(*, t) = *$,
- for all $x \in D$, $\Psi(x, 0) = *$ and $\Psi(x, 1) = x$, and
- for all $p \in D^{(0)}$ and all $t \in I$, $\Psi(p, t) \in C^{(1)}$.

Let $f : \mathbb{N}[\frac{1}{4}] \to \mathbb{N}[\frac{1}{4}]$ be a function. A 1-combing Ψ of the pair (C, D) is *f-tame* if for all $n \in \mathbb{N}[\frac{1}{4}]$ and for all $x \in D$ and $s, t \in [0, 1]$ such that $s \leq t$, if $d_C(*, \Psi(x, t)) \leq n$, then $d_C(*, \Psi(x, s)) \leq f(n)$.

Finally, the definition of intrinsic tame filling function was introduced by Brittenham and Hermiller in [8]:

Definition 1.2.4. Let $f : \mathbb{N}[\frac{1}{4}] \to \mathbb{N}[\frac{1}{4}]$ be non-decreasing. Then f is an *intrinsic tame filling function* for $\langle A|R\rangle$ if, for all $w \in (A \cup A^{-1})^*$ with $w =_G 1$, there is a van Kampen diagram Δ_w for w with basepoint $*$ and a 1-combing Ψ_w of $(\Delta_w, \partial\Delta_w)$ based at $*$ such that Ψ_w is f-tame. Similarly, f is an *extrinsic tame filling function* for $\langle A|R\rangle$ if, for all $w \in (A \cup A^{-1})^*$ with $w =_G 1$, there is a van Kampen diagram Δ_w for w with basepoint $*$ and a 1-combing Ψ_w of $(X, \pi_\Delta(\partial\Delta_w))$ based at 1_G such that the image of Ψ_w is contained in the image of π_Δ and Ψ_w is f-tame.

Chapter 2

Intrinsic Tame Filling Functions are Equivalent to Intrinsic Diameter Functions

The purpose of this chapter is to prove Theorem 2.4.4. To do so, we will first build up a number of definitions and properties related to 1-combings and certain nice subcomplexes of van Kampen diagrams. We will then bring them all together in the last section to prove the theorem.

2.1 Icicles

Let Δ be a van Kampen diagram and T a spanning tree of Δ. Given a 1-cell $e \in E(\Delta \setminus T)$ with a corresponding directed edge $\overrightarrow{e}$ directed from a vertex x to a vertex y, let γ_x and γ_y be the unique simple paths in T from $*$ to x and y, respectively. Let α be the longest initial segment on which γ_x and γ_y agree, and define β_x and β_y by $\gamma_x = \alpha \cdot \beta_x$ and $\gamma_y = \alpha \cdot \beta_y$. Then $\eta = \beta_x \cdot \overrightarrow{e} \cdot \overline{\beta_y}$ is a simple circuit. Since Δ is planar, by the Jordan Curve Theorem, we know that η splits the plane into two components, the inside and outside of η. See Figure 2.1.

Definition 2.1.1. Given the above notation, the *T-icicle at e* is the union of α, η, and the inside of η. It will just be called the icicle at e and be denoted I_e if the tree

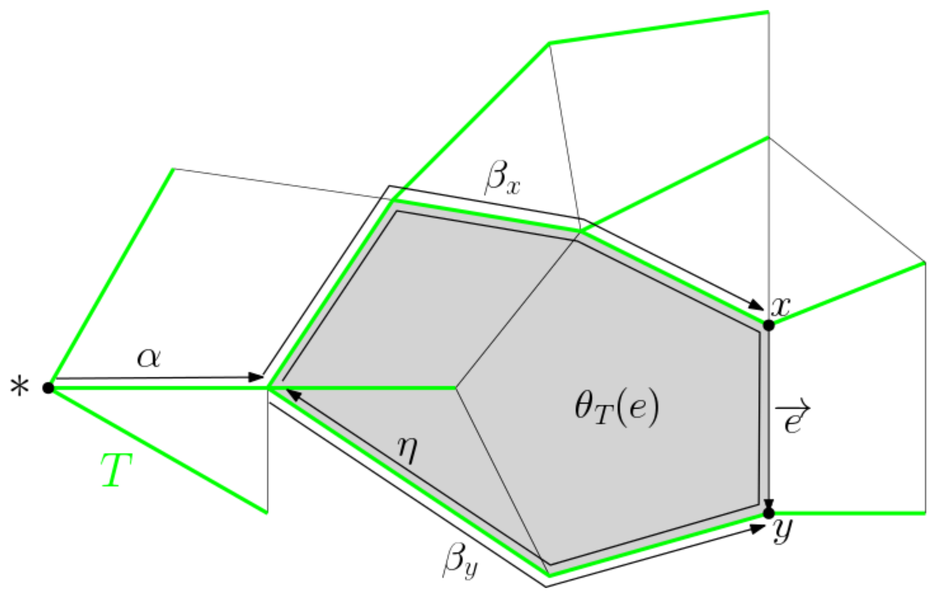

Figure 2.1: The T-icicle at e. Edges in the tree T are thickened. The body of the icicle is shaded and α is the tail.

being used is clear. Define the *tail* of the icicle to be α and the *body* of the icicle to be the union of η and the inside of η.

Definition 2.1.2. Let $\theta_T(e)$ be the 2-cell in the body of I_e with $e \subset \partial\theta_T(e)$. Then θ_T is a function from $E(\Delta \setminus T)$ to the 2-cells of Δ, which we will call the *icicle flow function* for T.

Icicles are the building blocks we will use to construct van Kampen diagrams with tame 1-combings, and are also the basis for constructing said 1-combings. Icicles are particularly useful tools for these purposes in part because of the way they intersect each other: for every pair of icicles, either one is entirely contained in the other or they do not intersect, except possibly at their boundary.

Lemma 2.1.3. *Let Δ be a van Kampen diagram and T a spanning tree of Δ. Let $e, e' \in E(\Delta \backslash T)$. If $e' \subseteq I_e$, then $I_{e'} \subseteq I_e$. If $e' \not\subseteq I_e$ and $e \not\subseteq I_{e'}$, then* $\text{int}(I_e) \cap \text{int}(I_{e'}) = \emptyset$.

Proof. Let e have endpoints x and y and e' have endpoints x' and y'. Let γ_x, γ_y, $\gamma_{x'}$, and $\gamma_{y'}$ be the unique simple paths in T from $*$ to x, y, x', and y', respectively. Suppose $e' \subseteq I_e$. We will first show that $\gamma_{x'}$ and $\gamma_{y'}$ are contained in I_e. Suppose by way of contradiction that there is a point $p \in \gamma_{x'} \setminus I_e$. Since $x' \in I_e$, there is also some vertex $q \in \gamma_{x'} \cap \partial I_e \cap \Delta^{(0)}$ such that p is between $*$ and q in $\gamma_{x'}$. Since $\partial I_e \cap \Delta^{(0)} \subset \gamma_x \cup \gamma_y$, without loss of generality, let $q \in \gamma_x$. Then the unique simple path in T from $*$ to q must be contained in γ_x. But since p is between $*$ and q in $\gamma_{x'}$, this path must also contain p. This contradicts the fact that $p \notin I_e$. So $\gamma_{x'} \subset I_e$. By the same argument, $\gamma_{y'} \subset I_e$. So $\partial I_{e'} = \gamma_{x'} \cup \gamma_{y'} \cup e' \subset I_e$. Since I_e is simply connected, this implies that $I_{e'} \subseteq I_e$.

Now suppose $e' \nsubseteq I_e$ and $e \nsubseteq I_{e'}$. Let B_e and $B_{e'}$ be the bodies of I_e and $I_{e'}$, respectively. Note that $\text{int}(I_e) = \text{int}(B_e)$ and $\text{int}(I_{e'}) = \text{int}(B_{e'})$. I will first show that $\gamma_{x'}$ does not intersect $\text{int}(B_e)$. For suppose there were some $p \in \gamma_{x'} \cap \text{int}(B_e)$. Now since $e' \nsubseteq I_e$, $e' \cap \text{int}(B_e) = \emptyset$. In particular, $x' \notin \text{int}(B_e)$. So there is some vertex $q \in \partial B_e \cap \Delta^{(0)}$ such that p is between $*$ and q in $\gamma_{x'}$. Since $\partial B_e \cap \Delta^{(0)} \subset \gamma_x \cup \gamma_y$, without loss of generality let $q \in \gamma_x$. Then the unique simple path in T from $*$ to q must be contained in γ_x. But since p is between $*$ and q in $\gamma_{x'}$, this path must also contain p. This contradicts the fact that $p \in \text{int}(B_e)$, since $\text{int}(B_e)$ does not intersect γ_x. So $\gamma_{x'} \cap \text{int}(B_e) = \emptyset$. By the same argument, $\gamma_{y'} \cap \text{int}(B_e) = \gamma_x \cap \text{int}(B_{e'}) = \gamma_y \cap \text{int}(B_{e'}) = \emptyset$. Then since $\partial B_e \subseteq e \cup \gamma_x \cup \gamma_y$, we have that $\partial B_e \cap \text{int}(B_{e'}) = \emptyset$. In the same way, $\partial B_{e'} \cap \text{int}(B_e) = \emptyset$.

Now consider $C = \text{int}(B_e) \cap \text{int}(B_{e'})$, and suppose by way of contradiction that $C \neq \emptyset$. Note that in this case, $\bar{C} = B_e \cap B_{e'}$ and that $\partial C \subseteq \partial B_e \cup \partial B_{e'}$. Suppose there is some $z \in \partial C \setminus \partial B_e$. Then $z \in \partial B_{e'}$ and since $z \in \bar{C} \setminus \partial B_e$, we must have that $z \in \text{int}(B_e)$. This contradicts that $\partial B_{e'} \cap \text{int}(B_e) = \emptyset$, so it must be that case that $\partial C \setminus \partial B_e = \emptyset$. By the same argument, $\partial C \setminus \partial B_{e'} = \emptyset$. Therefore, $\partial C = \emptyset$. This implies that $\bar{C} = \text{int}(C)$, which makes $C = \emptyset$.

□

2.2 1-combings Respecting a Spanning Tree

In this section, we will construct a special type of 1-combing that will end up being particularly tame. This is because the combing paths will respect a spanning tree, growing outwards from the basepoint of the diagram along the icicles of the tree and never crossing a branch of the tree. If the spanning tree chosen is a tree of geodesics out of the basepoint, then such a 1-combing has combing paths that are in some way close to being geodesics themselves, resulting in a 1-combing that is about as tame as possible.

Given a 2-complex C, a subcomplex $D \subseteq C^{(1)}$, a point $x \in C$, and a 1-combing Ψ of the pair (C, D), define

$$P(\Psi, x) = \{y \in C | \text{there is } p \in D \text{ and } 0 \leq s \leq t \leq 1 \text{ with } \Psi(p,t) = x \text{ and } \Psi(p,s) = y\},$$

the set of points prior to x in Ψ. This is the set of points y such that there is some combing path containing x where y appears on that path before x.

Definition 2.2.1. Given a van Kampen diagram Δ with basepoint $*$ and a spanning tree T of Δ, a *1-combing respecting* T is a 1-combing Ψ of the pair $(\Delta, \partial\Delta)$ based at $*$ satisfying two properties:

1. For all $e \in E(\partial\Delta \setminus T)$, $\bigcup_{x \in e} P(\Psi, x) = I_e$.

2. For $p \in \partial\Delta$ and $t \in [0,1]$, if $\Psi(p,t) \in T$, then $\Psi(p,[0,t])$ is the unique simple path in T from $*$ to $\Psi(p,t)$.

The main purpose of this section is to show that such 1-combings exist:

Proposition 2.2.2. *Let Δ be a van Kampen diagram and let T be a spanning tree of Δ. Then there exists a 1-combing Ψ_T that respects T.*

Proof. We will first demonstrate how to construct such a 1-combing, and then prove that it is a 1-combing that respects the given tree. Figure 2.2 shows an example of what the combing paths of our constructed 1-combings could look like.

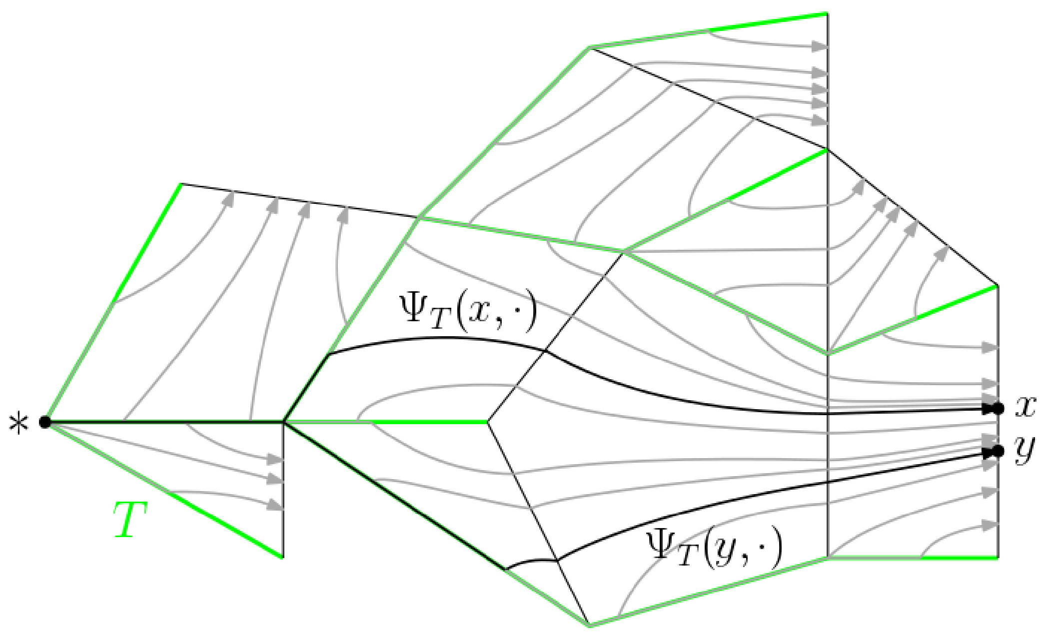

Figure 2.2: Representation of Ψ_T

We will define our combing paths in segments, one for each 2-cell, starting from the boundary of the diagram and working inwards towards the basepoint, following the icicles. At each step, we identify a 1-cell $e \in E(\Delta \setminus T)$ that is on the boundary of the subcomplex on which combing paths have not yet been defined, as well as its associated 2-cell $\sigma = \theta_T(e)$, the one that is inside the icicle at e with $e \subseteq \partial\sigma$. Then we determine the segments of the combing paths that intersect $\mathrm{int}(\sigma)$ with a homotopy H from e to $\partial\sigma \setminus \mathrm{int}(e)$. Note that the combing paths will run in the opposite direction to the paths $H(x, \cdot)$ for each $x \in e$. Once all the segments have been defined, we stitch them together into a continuous 1-combing Ψ_T.

Let Δ be a van Kampen diagram and T a spanning tree of Δ. Let $\Delta_0 = \Delta$, and define $\psi_0 : \partial\Delta \times [0,1] \to \Delta$ by $\psi_0(x,t) = x$. Now suppose we have defined a function $\psi_i : \partial\Delta \times [0,1] \to \Delta$ and a subcomplex Δ_i of Δ with the following four properties:

- $T \subseteq \Delta_i$.
- $\mathbb{R}^2 \setminus \Delta_i$ is connected.
- ψ_i is continuous.
- For all $x \in \partial\Delta$, $\psi_i(x,0) \in \Delta_i^{(1)}$.

Note that these properties all hold for $i = 0$. Now if all 1-cells of $\partial\Delta_i$ are in T, the process ends. Otherwise, we will define ψ_{i+1} and Δ_{i+1} such that these same properties hold, as in Figure 2.3. Let e_{i+1} be a 1-cell of $\partial\Delta_i$ that is not in T, and let I_{i+1} be the T-icicle at e_{i+1}. Let $\sigma_{i+1} = \theta_T(e_{i+1})$, the 2-cell in I_{i+1} whose boundary contains e_{i+1}. Note that $\partial I_{i+1} \subseteq \Delta_i$, so since $\mathbb{R}^2 \setminus \Delta_i$ is connected, we must have that $I_{i+1} \subseteq \Delta_i$. Therefore, $\sigma_{i+1} \subseteq \Delta_i$. Let $\Delta_{i+1} = \Delta_i \setminus (\mathrm{int}(\sigma_{i+1}) \cup \mathrm{int}(e_{i+1}))$. Since Δ_{i+1} is a subcomplex of Δ_i and we have not removed any vertices or any 1-cells in T from Δ_i, Δ_{i+1} also contains T. Additionally, since e_{i+1} is on the boundary of Δ_i,

$$\mathbb{R}^2 \setminus \Delta_{i+1} = (\mathbb{R}^2 \setminus \Delta_i) \cup \mathrm{int}(\sigma_{i+1}) \cup \mathrm{int}(e_{i+1})$$

is connected.

We will define a continuous $H_{i+1} : e_{i+1} \times [0,1] \to \sigma_{i+1}$ with the following properties:

i) For all $x \in e_{i+1}$, $H_{i+1}(x,0) = x$ and $H_{i+1}(x,1) \in \partial\sigma_{i+1} \setminus \mathrm{int}(e_{i+1})$.

ii) If x is an endpoint of e_{i+1}, then $H_{i+1}(x,t) = x$ for all $t \in [0,1]$.

iii) For all $x \in \text{int}(e_{i+1})$ and all $t \in (0,1)$, $H_{i+1}(x,t) \in \text{int}(\sigma_{i+1})$.

Since σ_{i+1} is a 2-cell, we know that it is homeomorphic to the unit disk

$$D^2 = \{(a,b) \in \mathbb{R}^2 | a^2 + b^2 \leq 1\}.$$

Define

$$S^1_+ = \{(a,b) \in \mathbb{R}^2 | a^2+b^2 = 1 \text{ and } b \geq 0\} \text{ and } S^1_- = \{(a,b) \in \mathbb{R}^2 | a^2+b^2 = 1 \text{ and } b \leq 0\}.$$

Then let $\phi_{\sigma_{i+1}} : \sigma_{i+1} \to D^2$ be a homeomorphism such that $\phi_{\sigma_{i+1}}(e_{i+1}) = S^1_+$. Note that this implies that $\phi_{\sigma_{i+1}}(\partial\sigma_{i+1} \setminus \text{int}(e_{i+1})) = S^1_-$. Then define $h : S^1_+ \times [0,1] \to D^2$ by

$$h((a,b),t) = (a,(1-2t)b).$$

Note that h is a straight-line homotopy rel boundary from S^1_+ to S^1_-. Finally, let

$$H_{i+1} = \phi^{-1}_{\sigma_{i+1}} \circ h \circ (\phi_{\sigma_{i+1}}|_{e_{i+1}} \times \text{id}_{[0,1]}).$$

This definition gives the desired properties, for:

i) Let $x \in e_{i+1}$. Then

$$H_{i+1}(x,0) = \phi^{-1}_{\sigma_{i+1}}(h(\phi_{\sigma_{i+1}}(x),0)) = \phi^{-1}_{\sigma_{i+1}}(\phi_{\sigma_{i+1}}(x)) = x.$$

Also, since $\phi_{\sigma_{i+1}}(x) \in S^1_+$, this implies that $h(\phi_{\sigma_{i+1}}(x),1) \in S^1_-$, so

$$H_{i+1}(x,1) = \phi^{-1}_{\sigma_{i+1}}(h(\phi_{\sigma_{i+1}}(x),1)) \in \phi^{-1}_{\sigma_{i+1}}(S^1_-) = \partial\sigma_{i+1} \setminus \text{int}(e_{i+1}).$$

ii) Let x be an endpoint of e_{i+1} and let $t \in [0,1]$. Then $\phi_{\sigma_{i+1}}(x) \in \{(1,0),(-1,0)\}$. Since the vertical component of $\phi_{\sigma_{i+1}}(x)$ is 0, $h(\phi_{\sigma_{i+1}}(x),t) = \phi_{\sigma_{i+1}}(x)$. Thus,

$$H_{i+1}(x,t) = \phi^{-1}_{\sigma_{i+1}}(h(\phi_{\sigma_{i+1}}(x),t)) = \phi^{-1}_{\sigma_{i+1}}(\phi_{\sigma_{i+1}}(x)) = x.$$

iii) Let $x \in \operatorname{int}(e_{i+1})$ and $t \in (0,1)$. Let $\phi_{\sigma_{i+1}}(x) = (a,b)$. Since $x \in \operatorname{int}(e_{i+1})$, we know that $b > 0$. Since $t \in (0,1)$, this implies that $|(1-2t)b| < |b|$. Therefore, $a^2 + ((1-2t)b)^2 < 1$, which means that $h(\phi_{\sigma_{i+1}}(x),t) \in \operatorname{int}(D^2)$. Hence,

$$H_{i+1}(x,t) = \phi^{-1}_{\sigma_{i+1}}(h(\phi_{\sigma_{i+1}}(x),t)) \in \operatorname{int}(\sigma_{i+1}).$$

Now define $\psi_{i+1} : \partial\Delta \times [0,1] \to \Delta$ by

$$\psi_{i+1}(x,t) = \begin{cases} H_{i+1}(\psi_i(x,0), 1-t), & \psi_i(x,0) \in e_{i+1} \\ \psi_i(x,0), & \psi_i(x,0) \in \Delta^{(1)} \setminus \operatorname{int}(e_{i+1}). \end{cases}$$

Note that the two pieces of ψ_{i+1} agree on their intersection because H_{i+1} fixes the endpoints of e_{i+1}, so ψ_{i+1} is well-defined and, by the pasting lemma, continuous.

Finally, let $x \in \partial\Delta$. We know that $\psi_i(x,0) \in \Delta_i^{(1)}$, so if $\psi_i(x,0) \notin e_{i+1}$, then $\psi_{i+1}(x,0) = \psi_i(x,0) \in \Delta_{i+1}^{(1)}$. But if $\psi_i(x,0) \in e_{i+1}$, then

$$\psi_{i+1}(x,0) = H_{i+1}\left(\psi_i(x,0), 1\right) \in \partial\sigma_{i+1} \setminus \operatorname{int}(e_{i+1}) \subseteq \Delta_{i+1}^{(1)}.$$

In either case, $\psi_{i+1}(x,0) \in \Delta_{i+1}^{(1)}$. So Δ_{i+1} and ψ_{i+1} satisfy the desired properties.

Now, since there are only finitely many 1-cells of Δ outside T, this construction must end with some Δ_n such that $\partial\Delta_n \subseteq T$. I claim that, in fact, $\Delta_n = T$. We know that $T \subseteq \Delta_n$. Suppose by way of contradiction that there exists $x \in \Delta_n \setminus T$. Then

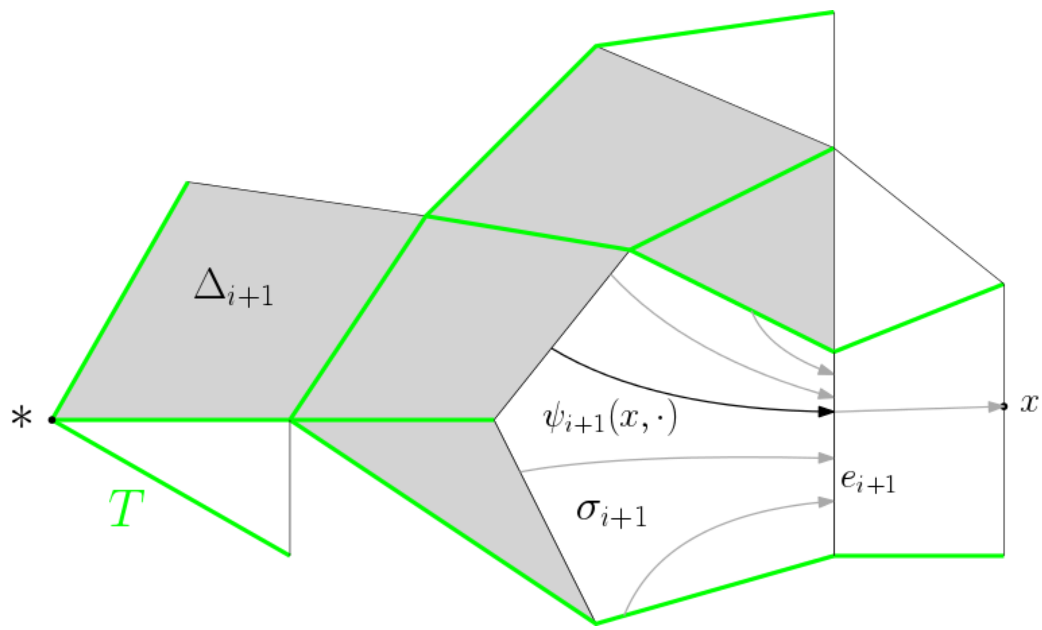

Figure 2.3: Defining Δ_{i+1} and ψ_{i+1}. The entire diagram is Δ. Δ_{i+1} is the shaded region along with T, and $\Delta_i = \Delta_{i+1} \cup \sigma_{i+1}$.

in particular $x \in \text{int}(\Delta_n)$, which means that $\mathbb{R}^2 \setminus \partial\Delta_n$ is disconnected, and therefore $\mathbb{R}^2 \setminus T$ is disconnected. This is a contradiction, since a finite, planar tree cannot disconnect the plane, so $\Delta_n = T$.

Now let $H_{n+1} : T \times [0,1] \to T$ be the homotopy from T to $*$ such that, for each $x \in T$, $H_{n+1}(x,\cdot)$ is the constant speed parametrization of the unique simple path from x to $*$ in T. Note that for all $x \in \partial\Delta$, $\psi_n(x,0) \in \Delta_n = T$, so we can define $\psi_{n+1} : \partial\Delta \times [0,1] \to \Delta$ by $\psi_{n+1}(x,t) = H_{n+1}(\psi_n(x,0), 1-t)$.

Finally, define $\Psi_T : \partial\Delta \times [0,1] \to \Delta$ by $\Psi_T(x,t) = \psi_i(x,(n+1)t + i - n - 1)$ for $t \in \left[1 - \frac{i}{n+1}, 1 - \frac{i-1}{n+1}\right]$. Note that when $t = 1 - \frac{i}{n+1}$, we have $(n+1)t + i - n - 1 = 0$ and when $t = 1 - \frac{i-1}{n+1}$, we have $(n+1)t + i - n - 1 = 1$. So to show that Ψ_T is well-defined, we must show that for all $i \in [n]$ and $x \in \partial\Delta$, $\psi_{i+1}(x,1) = \psi_i(x,0)$. This follows immediately from the definitions of ψ_{i+1} and the fact that $H_{i+1}(y,0) = y$ for all $y \in e_{i+1}$. Since each ψ_i is continuous, Ψ_T is continuous by the pasting lemma.

We will now establish that Ψ_T is a 1-combing that respects T. It will be convenient to first establish the second property, that for all $p \in \partial\Delta$ and $t \in [0,1]$, if $\Psi(p,t) \in T$, then $\Psi(p,[0,t])$ is the unique simple path in T from $*$ to $\Psi(p,t)$.

Let $p \in \partial\Delta$ and $t \in [0,1]$, and suppose that $x = \Psi(p,t) \in T$. There is some $i \in [n+1]$ such that $t \in \left[1 - \frac{i}{n+1}, 1 - \frac{i-1}{n+1}\right]$, so $x = \psi_i(p,s)$ where $s = (n+1)(t-1)+i$. If $t \leq \frac{1}{n+1}$, we may choose $i = n+1$ so that $x = \psi_{n+1}(p,s)$. Then $\Psi(p,[0,t]) = \psi_{n+1}(p,[0,s])$, which is by definition the unique simple path in T from $*$ to x. Otherwise, we may choose $i \leq n$ such that $t \neq 1 - \frac{i}{n+1}$, and therefore $s \neq 0$. In that case, if $s = 1$, then $x = \psi_{i-1}(p,0)$ from the definition of ψ_i. If $0 < s < 1$, we cannot have that $x = H_i(\psi_{i-1}(x,0), 1-s)$, since $x \in T \subseteq \Delta^{(1)}$ and we know that $H_i(\psi_{i-1}(x,0), 1-s) \notin \Delta^{(1)}$ because $1-s \in (0,1)$. So from the definition of ψ_i we must have that $x = \psi_{i-1}(p,0) \in T$. Therefore, $\psi_{i-1}(p,0) \notin \mathrm{int}(e_i)$, so $\psi_i(p,r) = \psi_{i-1}(p,0) = x$ for all $r \in [0,1]$. Now for all $j \in [n]$ with $j > i$, suppose by induction that $\psi_{j-1}(p,0) = x$. Then by the same logic, $\psi_{j-1}(p,0) \notin \mathrm{int}(e_j)$, so $\psi_j(p,r) = x$ for all $r \in [0,1]$. As a result, $\Psi_T\left(p, \left[\frac{1}{n+1}, t\right]\right) = \{x\}$. So $\Psi_T(p,[0,t]) = \Psi_T\left(p, \left[0, \frac{1}{n+1}\right]\right)$, and we already know from the case with $t \leq \frac{1}{n+1}$ that $\Psi_T\left(p, \left[0, \frac{1}{n+1}\right]\right)$ is the unique simple path from $*$ to x. This establishes the desired property.

Now we will show that Ψ_T is in fact a 1-combing. We already know that Ψ_T is continuous. For any $p \in \partial\Delta$, $\Psi_T(p,0) = \psi_{n+1}(p,0) = *$ and $\Psi_T(p,1) = \psi_1(p,1) = \psi_0(p,0) = p$. Also, $p \in T$, so $\Psi(p,1) \in T$. Then by the second property of a 1-combing respecting T, we know that $\Psi(p,[0,1]) \subseteq T \subseteq \Delta^{(1)}$. In particular, if $p = *$, the unique simple path from $*$ to $*$ in T is the constant path, so $\Psi(*,[0,1]) = \{*\}$. Hence, Ψ_T is a 1-combing.

Finally, we will show that Ψ_T satisfies the first property of a 1-combing respecting T, that for all $e \in E(\partial\Delta \setminus T)$, $\bigcup_{x \in e} P(\Psi, x) = I_e$.

First we will show that $I_e \supseteq \bigcup_{x \in e} P(\Psi_T, x)$. Let $p \in \partial\Delta$ and $t \in [0,1]$ such that $\Psi_T(p,t) \in e$. Suppose by way of contradiction that there is some $s \leq t$ so that

$\Psi_T(p,s) \notin I_e$. Then there is some smallest $s' \in [s,t]$ such that $\Psi_T(p,s') \in I_e$. Note that this implies that $\Psi_T(p,s')$ is on the boundary of the icicle. There are two possible cases for where $\Psi_T(p,s')$ lands, as shown in Figure 2.4.

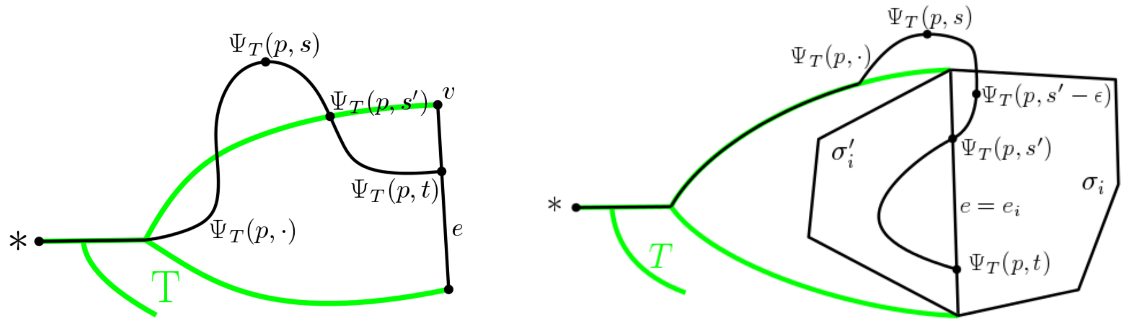

(a) Case 1 (b) Case 2

Figure 2.4: If a prior point of e in Ψ_T were outside the icicle at e

Case 1: $\Psi_T(p,s') \in T$. By the definition of I_e, $\Psi_T(p,s')$ is on the simple path in T from $*$ to some endpoint v of e. Therefore, the simple path in T from $*$ to $\Psi_T(p,s')$ is in the icicle at e. So $\Psi_T(p,[0,s']) \subseteq I_e$. So $\Psi_T(p,s) \in I_e$, which is a contradiction.

Case 2: $\Psi_T(p,s') \notin T$. Since $\Psi_T(p,s')$ is on the boundary of the icicle at e, this implies that $\Psi_T(p,s') \in \text{int}(e)$. Let $\epsilon = \frac{1}{2(n+1)}$ and let $e = e_i$ from the construction of Ψ_T. Then since $\Psi_T(p,s') \in \Delta^{(1)}$, we must have that $\Psi_T(p,s') = \psi_j(p,0)$ for some j. In fact, $j < i$, since we know that $\psi_j(p,0) \in \Delta_j$ but $\text{int}(e_i) \notin \Delta_k$ for $k \geq i$. Since for all k with $j \leq k < i$, if $\psi_k(p,0) \in e_i \neq e_k$ then $\psi_{k+1}(p,0) = \psi_k(p,0)$, we know that $\psi_i(p,1) = \psi_{i-1}(p,0) = \psi_j(p,0) = \Psi_T(p,s')$. Then $\Psi_T(p,s'-\epsilon) = \psi_i(p,\frac{1}{2}) \in \sigma_i$ by the definition of H_i. Now since s' was the smallest element of $[s,t]$ in the icicle at e_i, either $s \leq s'-\epsilon < s'$, in which case $\Psi_T(p,s'-\epsilon)$ is not in the icicle at e_i, or $s'-\epsilon < s < s'$, in which case $\Psi_T(p,s) = \psi_i(p,r)$ for some $r \in (\frac{1}{2},1)$, and thus $\Psi_T(p,s) \in \sigma_i$. Either case implies that σ_i is not in the icicle at e_i, which contradicts the definition of σ_i. This final contradiction completes case 2, implying that no such s exists.

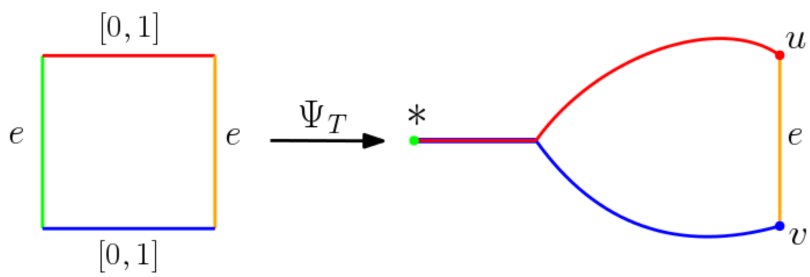

Figure 2.5: Ψ_T mapping the boundary of $e \times [0, 1]$ to the boundary of I_e when $e \subseteq \partial\Delta$.

Now we will show the opposite containment. Let $y \in I_e$.

Case 1: $e \subseteq \partial\Delta$. See Figure 2.5. Let u and v be the endpoints of e. Now $\Psi_T|_{e\times[0,1]}$ is a continuous function from the ball $e \times [0, 1]$ into Δ where $\Psi_T(e \times \{0\}) = *$, $\Psi_T(\{u\} \times [0, 1])$ is the simple path from $*$ to u in T, $\Psi_T(e \times \{1\}) = e$, and $\Psi_T(\{v\} \times [0, 1])$ is the simple path from $*$ to v in T. So the image of the path around the boundary of $e \times [0, 1]$ is the path around the boundary of the icicle at e. Therefore, since $\Psi_T|_{e\times[0,1]}$ is continuous, every point in the icicle at e is in $\Psi_T(e \times [0, 1])$. Hence, there is some $x \in e$ and $s \in [0, 1]$ such that $y = \Psi_T(x, s)$. Of course, $x = \Psi_T(x, 1)$, so $y \in P(\Psi_T, x)$.

Case 2: $e \not\subseteq \partial\Delta$. First, note that

$$\Delta = \bigcup_{f \in E(\partial\Delta \setminus T)} I_f,$$

so for each $x \in \Delta$, there exists a 1-cell $f \subseteq \partial\Delta$ with $x \in I_f$. Then as a result of Case 1, there must be some $p \in \partial\Delta$ and $s \in [0, 1]$ such that $y = \Psi_T(p, s)$. Since $y \in I_e$, there is some $t \geq s$ such that $\Psi_T(p, t) \in \partial I_e$. If $\Psi_T(p, t) \in e$, this shows that $y \in P(\Psi_T, x)$, where $x = \Psi_T(p, t) \in e$. Otherwise, we must have that $\Psi_T(p, t) \in T$. As a result, $\Psi_T(p, [0, t]) \subseteq T$, and in particular $y \in T$. Since $y \in I_e$, y must be on the unique simple path in T from $*$ to an endpoint of e; call this endpoint x. By the same reasoning as for y, there must be some $p' \in \partial\Delta$ and $t' \in [0, 1]$ such that

$x = \Psi_T(p', t')$. Then since $x \in T$ and y is on the simple path from $*$ to x in T, there must be some $s' < t'$ such that $y = \Psi_T(p', s')$. Hence, $y \in P(\Psi_T, x)$. □

Note that, in the construction of Ψ_T, although it was convenient to deal with the 1-cells of $\Delta \setminus T$ in a somewhat arbitrary order to construct a continuous Ψ_T, fewer choices are required to determine the images of individual combing paths. Since for each $i \in [n]$, $\sigma_i \in \Delta_{i-1} \setminus \Delta_i$, we have that the σ_i's are all distinct; in other words, the icicle flow function θ_T is injective. Since $\Delta_n = T$ and does not contain any 2-cells, we have that θ_T is a bijection between the 1-cells of $\Delta \setminus T$ and the 2-cells of Δ. This bijection, along with a choice of homeomorphism $\phi_\sigma : \sigma \to D^2$ for each 2-cell σ, determines the images of the combing paths of Ψ_T.

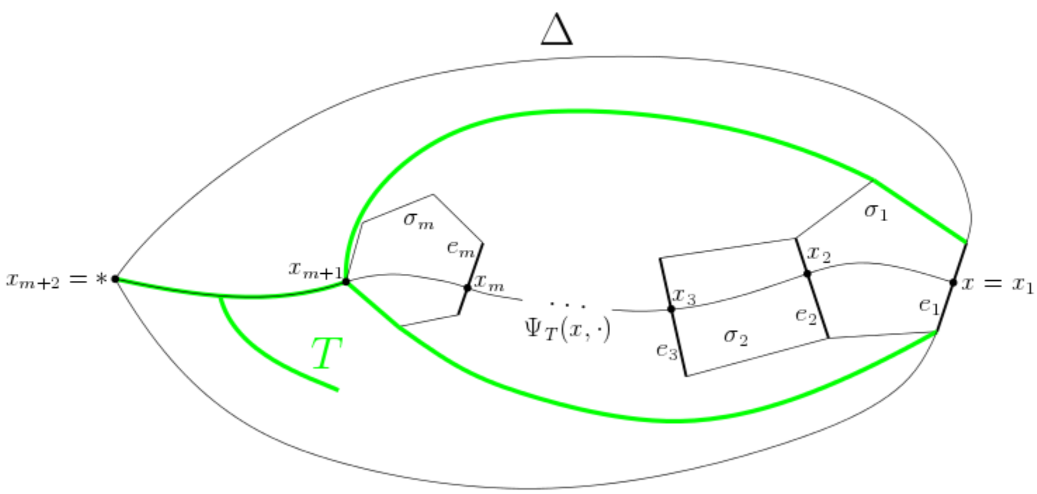

Figure 2.6: The decomposition of $\Psi_T(x, \cdot)$ given by Lemma 2.2.3.

In particular, the following lemma will be useful when dealing with individual combing paths. It decomposes a combing path into an initial segment in the tree T and subsequent segments in each 2-cell that it crosses, as shown in Figure 2.6.

Lemma 2.2.3. *Let Δ be a van Kampen diagram, T a spanning tree for Δ, and Ψ_T a 1-combing constructed as in the proof of Proposition 2.2.2. Given $x \in \partial\Delta$, there exists a sequence $1 \geq t_1 > \cdots > t_{m+1} > t_{m+2} = 0$ with the following properties. Let*

$x_j = \Psi_T(x, t_j)$ for all $j \in [m+2]$. Then for each $j \in [m]$ there is a 1-cell $e_j \in E(\Delta \setminus T)$ with $x_j \in \operatorname{int}(e_j)$. Furthermore, $\Psi_T(x, [0, t_{m+1}])$ is the simple path from $*$ to x_{m+1} in T, $\Psi_T(x, [t_1, 1]) = \{x\}$, and for all $j \in [m]$, $t_j = min\{t \in [0,1] | \Psi_T(x,t) = x_j\}$, and $\Psi_T(x, [t_{j+1}, t_j]) = \phi_{\sigma_j}^{-1}(h(\phi_{\sigma_j}(x_j), [0,1]))$, where $\sigma_j := \theta_T(e_j)$.

Proof. Let $x \in \partial\Delta$. For each $i \in \{0, 1, \ldots, n+1\}$, let $s_i = 1 - \frac{i}{n+1}$ and let $y_i = \Psi_T(x, s_i)$. Also, for $i \in \{0, 1, \ldots, n-1\}$, either $y_i \in \operatorname{int}(e_{i+1})$, in which case $\Psi_T(x, [s_{i+1}, s_i]) = H_{i+1}(y_i, [0,1])$, or $y_i \notin \operatorname{int}(e_{i+1})$, in which case $\Psi_T(x, [s_{i+1}, s_i]) = \{y_i\} = \{y_{i+1}\}$. Let $i_1 < i_2 < \cdots < i_m$ be the collection of indices such that $y_{i_j} \in \operatorname{int}\big(e_{i_j+1}\big)$, and let $i_{m+1} = n$. Then for all $j \in [m+1]$, let $t_j = s_{i_j}$ and $x_j = y_{i_j}$ Then we know that $\Psi_T(x, [0, t_{m+1}]) =: \gamma_x$, the unique simple path from $*$ to x_{m+1} in T. Then for all $i \in \{0, \ldots, i_1 - 1\}$, since $\Psi_T(x, [s_{i+1}, s_i]) = \{y_i\} = \{y_{i+1}\}$, we have that $x = y_i = x_1$. Hence, $\Psi_T(x, [t_1, 1]) = \{x\}$. Similarly, for all $j \in [m]$ and $i \in \{i_j + 1, \ldots, i_{j+1} - 1\}$, $y_{i_j+1} = y_i = x_{j+1}$, so $\Psi_T(x, [t_{j+1}, s_{i_j+1}]) = \{x_{j+1}\}$. Now for $j \in [m]$,

$$H_{i_j+1}(x_j, 1) = \Psi_T(x, s_{i_j+1}) = y_{i_j+1} = x_{j+1},$$

and therefore

$$\begin{aligned}
\Psi_T(x, [t_{j+1}, t_j]) &= \Psi_T(x, [t_{j+1}, s_{i_j+1}]) \cup \Psi_T(x, [s_{i_j+1}, t_j]) \\
&= \{x_{j+1}\} \cup H_{i_j+1}(x_j, [0,1]) \\
&= H_{i_j+1}(x_j, [0,1]) \\
&= \phi_{\sigma_{i_j+1}}^{-1}(h(\phi_{\sigma_{i_j+1}}(x_j), [0,1])).
\end{aligned}$$

Note that by definition $\sigma_{i_j+1} = \theta_T(e_{i_j+1})$, and e_{i_j+1} is the 1-cell containing x_j, so this matches the statement of the lemma.

We have only left to show that $t_j = \min\{t \in [0,1] | \Psi_T(x,t) = x_j\}$. Note that the above implies that

$$\Psi_T(x,[0,1]) = \gamma_x \cup \bigcup_{j\in[m]} H_{i_j+1}(x_j,[0,1]).$$

Note also that for each $j,k \in [m]$ with $k > j$, $x_j \neq x_k$, since $x_j \in \mathrm{int}\big(e_{i_j+1}\big)$ and $\mathrm{int}\big(e_{i_j+1}\big)$ is not in Δ_{i_j+1}, whereas $x_k \in \mathrm{int}(e_{i_k+1})$, and $e_{i_k+1} \subseteq \Delta_{i_k} \subseteq \Delta_{i_j+1}$ since $i_k \geq i_j + 1$. Furthermore, for all $j,k \in [m]$, $x_j \notin \Psi_T(x,(s_{i_k+1},t_k)) = H_{i_k+1}(x_k,(0,1))$, since $H_{i_k+1}(x_k,(0,1)) \cap \Delta^{(1)} = \emptyset$. Also, for all $j \in [m]$, $x_j \notin \gamma_x$, since $\gamma_x \subseteq T$ but $x_j \in \mathrm{int}\big(e_{i_j+1}\big)$, and $e_{i_j+1} \in E(\Delta \setminus T)$. This shows that $x_j \notin \Psi_T(x,[0,t_j))$. In other words, $t_j = \min\{t \in [0,1] | \Psi_T(x,t) = x_j\}$. □

2.3 Constructing Simply, Geodesically Bounded Van Kampen Diagrams of Small Diameter

Our main strategy for proving Theorem 2.4.4 will be to take a van Kampen diagram with a geodesic spanning tree and start replacing the bodies of icicles whose diameter is too large with diagrams that have smaller intrinsic diameter. This will cause 1-combings respecting the tree to become more and more tame. In order for this replacement to happen while preserving the necessary structure, the diagrams with which we replace the bodies of icicles must have two important properties.

Definition 2.3.1. A van Kampen diagram is *simply bounded* if it is bounded by a simple circuit.

Note that the bodies of icicles are always simply bounded by definition, so the replacement diagrams must also be simply bounded in order for them to be bodies of icicles in the same tree.

Definition 2.3.2. Let Δ be a van Kampen diagram with basepoint $*$. Then we say that Δ is *geodesically bounded* if $d_\Delta(*, x) = d_{\partial\Delta}(*, x)$ for all $x \in \Delta^{(0)} \cap \partial\Delta$.

Note that if we consider the body of an icicle of a tree of geodesics to be a van Kampen diagram with basepoint at the intersection of the tail and body of the icicle, it is geodesically bounded. The replacement diagrams must share this property in order to preserve distances in the resulting diagram.

The purpose of this section is to show that, given a van Kampen diagram that could be the body of an icicle—in other words, a diagram that is simply bounded—there is always another van Kampen diagram for the same word that is both simply bounded and geodesically bounded and has relatively small intrinsic diameter. In order to say what "relatively small" means, we need a few more definitions.

Recall that $\mathcal{P} = \langle A|R\rangle$ is a finite presentation. Given $w \in (A \cup A^{-1})^*$, let SB_w be the set of all simply bounded van Kampen diagrams for w. Define the simply bounded diameter of w by

$$\mathrm{IDiam_{sb}}(w) = \inf\{\mathrm{IDiam}(D)|D \in \mathrm{SB}_w\}$$

and choose D_w to be a van Kampen diagram for w that attains this infimum if such a diagram exists. Note that the infimum is attained if and only if $\mathrm{SB}_w \neq \emptyset$. Let

$$M_{\mathcal{P}} = \max(\{0\} \cup \{\mathrm{IDiam_{sb}}(w)|w \in (A \cup A^{-1})^* \text{ such that } \ell(w) \leq 4 \text{ and } \mathrm{SB}_w \neq \emptyset\}).$$

Proposition 2.3.3. *Let $\mathcal{P} = \langle A|R\rangle$ be a finite presentation for a group G such that no generator is equal to the identity. Let $w \in (A \cup A^{-1})^*$ with $w =_G 1$. Suppose that there exists a simply bounded van Kampen diagram for w. Then there exists a simply* *and geodesically bounded van Kampen diagram* Δ *for* w *such that*

$$\mathrm{IDiam}(\Delta) \leq max\left(\mathrm{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right).$$

Proof. Case 1: Suppose that $\ell(w) \leq 2$. Since there exists a simply bounded van Kampen diagram for w, $\mathrm{SB}_w \neq \emptyset$. So let $\Delta = D_w$. Then $\mathrm{IDiam}(\Delta) \leq M_{\mathcal{P}} \leq \max(\mathrm{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}})$. Also, if there exists a vertex $x \neq *$ on the boundary, then $d_\Delta(*, x) = 1 = d_{\partial\Delta}(*, x)$, so Δ is geodesically bounded. Thus, Δ satisfies the desired conditions.

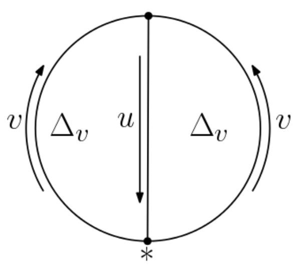

Figure 2.7: Showing $\mathrm{SB}_{vv^{-1}} \neq \emptyset$ in Case 2

Case 2: Suppose that $\ell(w) > 2$. We will first show that for any subword v of w with $\ell(v) = 2$, $\mathrm{SB}_{vv^{-1}} \neq \emptyset$. We will therefore be free to use the diagrams $D_{vv^{-1}}$ in order to construct Δ from the statement of the proposition. Construct a diagram Δ_v by taking a simply bounded van Kampen diagram for w and moving the basepoint so that Δ_v is a van Kampen diagram for vu, for some nonempty word u. Then since vu labels a simple circuit, u labels a simple path on the boundary. Therefore, taking two copies of Δ_v and gluing them along their respective paths labeled by u results in a van Kampen diagram for vv^{-1}, shown in Figure 2.7. Note that the boundary of this van Kampen diagram is a simple circuit since u is nonempty. Therefore, $\mathrm{SB}_{vv^{-1}}$ contains this diagram.

Let Δ_0 be a van Kampen diagram for w of minimum intrinsic diameter. Let

$* = p_0, p_1, \ldots, p_n = *$ be the sequence of vertices traversed in order by the path labeled by w along the boundary of Δ_0. Let $I_0 = \{i_1, \ldots, i_m\} \subseteq [n-1]$ be the set of indices such that the vertices they index appeared earlier on the boundary; that is,

$$I_0 = \{i \in [n-1] | \text{there exists } k < i \text{ such that } p_i = p_k\}.$$

The vertices at these indices are cut vertices of Δ_0, and therefore obstructions to making it simply bounded. If $I_0 = \emptyset$, then Δ_0 is simply bounded and $m = 0$. Otherwise, we will construct a sequence of van Kampen diagrams for w $\Delta_0, \Delta_1, \ldots, \Delta_m$ such that Δ_m is simply bounded and $\mathrm{IDiam}(\Delta_m) \leq \max\left(\mathrm{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right)$. Our strategy will be to add a bit of "padding" to the diagram at a cut vertex at each step, reducing the number of indices corresponding to a cut vertex by 1.

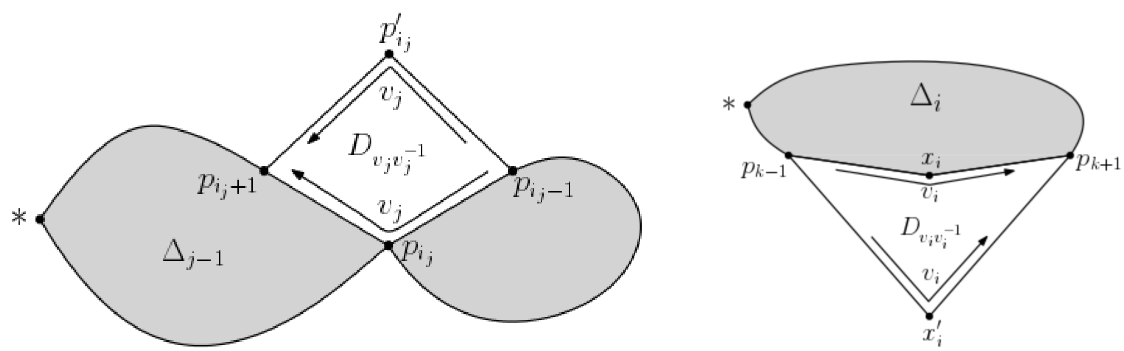

(a) Defining Δ_j for $0 \leq j < m$ (b) Defining Δ_{i+1} for $m \leq i < l$

Figure 2.8: Defining $\Delta_1, \ldots \Delta_l$ in Case 2 of Proposition 2.3.3

For $j \in [m]$, assume by induction that we have constructed a van Kampen diagram Δ_{j-1} for w such that $* = p_0, p_1, \ldots, p_n = *$ is the sequence of vertices traversed in order by the path labeled by w along the boundary of Δ_{j-1}, that $p_0, p_1, \ldots, p_{i_j-1}$ are all distinct, and that $\mathrm{IDiam}(\Delta_{j-1}) \leq \max\left(\mathrm{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right)$. Let

$$I_{j-1} = \{i \in [n-1] | \text{there exists } k < i \text{ such that } p_i = p_k\}$$

and assume by induction that $I_{j-1} = \{i_j, \ldots, i_m\}$. Then construct Δ_j from Δ_{j-1} in the following way. Let γ_j be the path along the boundary of Δ_{j-1} from p_{i_j-1} to p_{i_j+1} and let v_j be the subword of w that labels γ_j. Note that $\ell(v_j) = 2$, so we have shown above that $\mathrm{SB}_{v_j v_j^{-1}} \neq \emptyset$ and therefore $D_{v_j v_j^{-1}}$ is defined. Let η_j be a path in $D_{v_j v_j^{-1}}$ along the boundary starting at the basepoint labeled by v_j. Since $D_{v_j v_j^{-1}}$ is simply bounded, η_j is a simple path, and I claim that γ_j is as well. Since no generator is equal to the identity, we know that $p_{i_j-1} \neq p_{i_j} \neq p_{i_j+1}$, so we only need to show that $p_{i_j-1} \neq p_{i_j+1}$. Since $i_j \in I_{j-1}$, there is some $k < i_j$ such that $p_k = p_{i_j}$. Using Lemma 2.3.4, proven below, this implies that $p_{i_j-1} \neq p_{i_j+1}$, making γ_j a simple path.

Now construct Δ_j from Δ_{j-1} and $D_{v_j v_j^{-1}}$ by gluing η_j along γ_j, as in Figure 2.8A, and let $q_j : \Delta_{j-1} \sqcup D_{v_j v_j^{-1}} \to \Delta_j$ be the corresponding quotient map. Because γ_j and η_j are both simple paths along the boundaries of their respective diagrams, Δ_j is planar and simply connected. Let $p'_k = q_j(p_k)$ for $k \neq i_j$, and let p'_{i_j} be the vertex on the boundary of $q_j(D_{v_j v_j^{-1}})$ that is not in $q_j(\eta_j)$. Then the path along the boundary of Δ_j from p'_{i_j-1} to p'_{i_j+1} through p'_{i_j} is the image under q_j of a path of length 2 along the boundary of $D_{v_j v_j^{-1}}$ starting at the basepoint, and is therefore labeled by v_j. So Δ_j is a van Kampen diagram for w, and $* = p'_0, p'_1, \ldots, p'_n = *$ is the sequence of vertices traversed in order by the path labeled by w along the boundary of Δ_j. Additionally, $p'_{i_j} \neq p'_k$ for all $k < i_j$, so $p'_0, p'_1, \ldots, p'_{i_j}$ are all distinct. Furthermore, since $d_{\Delta_{j-1}}(*, p_{i_j-1}) \leq \left\lfloor \frac{\ell(w)}{2} \right\rfloor$ and q_j identifies the basepoint of $D_{v_j v_j^{-1}}$ with p_{i_j-1}, we know that for any $y \in q_j(D_{v_j v_j^{-1}})^{(0)}$,

$$d_{\Delta_j}(*, y) \leq d_{\Delta_{j-1}}(*, p_{i_j-1}) + \mathrm{IDiam}(D_{v_j v_j^{-1}}) \leq \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}.$$

Hence,

$$\begin{aligned}\mathrm{IDiam}(\Delta_j) &\leq \max\left(\mathrm{IDiam}(\Delta_{j-1}), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right) \\ &\leq \max\left(\mathrm{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right).\end{aligned}$$

Let

$$I_j = \{i \in [n-1] | \text{there exists } k < i \text{ such that } p'_i = p'_k\}.$$

Since $q_j|_{\Delta_{j-1}}$ is a bijection, for $0 \leq k < i < n$ with $i, k \neq i_j$, we have that $p'_i = p'_k$ if and only if $p_i = p_k$. Also, p'_{i_j} is only the image under q_j of the one boundary vertex of $D_{v_j v_j^{-1}}$ that is not on η_j, and not the image of any p_i. Therefore, $p'_{i_j} \neq p'_i$ for any $i \neq i_j$. As a result, $I_j = I_{j-1} \setminus \{i_j\} = \{i_{j+1}, \ldots, i_m\}$. This completes the induction.

So Δ_m is a van Kampen diagram for w, and $I_m = \emptyset$. Since $\ell(w) > 2$, this implies that the boundary of Δ_m is a simple circuit. Furthermore,

$$\mathrm{IDiam}(\Delta_m) \leq \max\left(\mathrm{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right).$$

However, Δ_m may not be geodesically bounded; we may need to continue with a similar process in order to acquire a geodesically bounded diagram with these properties.

Given any van Kampen diagram Δ, define the *cumulative boundary diameter* of Δ to be

$$\mathrm{bdiam}(\Delta) := \sum_{x \in \Delta^{(0)} \cap \partial\Delta} d_\Delta(*, x).$$

Note that the cumulative boundary diameter is bounded above with respect to the length of the boundary circuit of Δ. If Δ is a van Kampen diagram for a word w, then for each $x \in \Delta^{(0)} \cap \partial\Delta$ we have $d_\Delta(*, x) \leq d_{\partial\Delta}(*, x) \leq \frac{1}{2}\ell(w)$, and therefore $\mathrm{bdiam}(\Delta) \leq \frac{1}{2}\ell(w)^2$.

We will construct a sequence of van Kampen diagrams $\Delta_m, \Delta_{m+1}, \ldots, \Delta_l$ for w such that if $m \leq i \leq l$ then $\text{IDiam}(\Delta_i) \leq \max\left(\text{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right)$, Δ_i is simply bounded, and, for $i > m$, $\text{bdiam}(\Delta_i) > \text{bdiam}(\Delta_{i-1})$. The fact that the cumulative boundary diameter keeps increasing implies that at some point, each point on the boundary of the diagram will reach its maximum possible distance from the basepoint in a van Kampen diagram for w, resulting in a geodesically bounded diagram. Assuming that these properties hold for $i \geq m$, we will construct Δ_{i+1} as follows. Let

$$X_i = \{x \in \Delta_i^{(0)} \cap \partial\Delta_i | d_{\Delta_i}(*, x) < d_{\partial\Delta_i}(*, x)\},$$

the set of vertices x along the boundary of Δ_i such that there is no Δ_i-geodesic from $*$ to x in $\partial\Delta_i$. If $X_i = \emptyset$, then Δ_i is geodesically bounded, so we will let $l := i$ and the construction of the sequence ends. Otherwise, let $m_i = \min(\{d_{\Delta_i}(*, x) | x \in X_i\})$ and let $x_i \in X_i$ be a vertex that attains this minimum. Let $* = p_0, p_1, \ldots, p_{n-1}, p_n = *$ be the sequence of vertices traversed in order by the path labeled by w along the boundary of Δ_i, and let $k \in [n]$ such that $p_k = x_i$. Note that by the definition of X_i, $* \notin X_i$, so $0 < k < n$. Let γ_i be the path along the boundary from p_{k-1} to x_i to p_{k+1}, and let v_i be the subword of w of length 2 that labels γ_i. Let η_i be a path in $D_{v_i v_i^{-1}}$ along the boundary starting at the basepoint labeled by v_i. Note that γ_i and η_i are both simple paths, since they are paths of length 2 along the boundaries of simply bounded diagrams of a word w with $\ell(w) > 2$.

Construct Δ_{i+1} from Δ_i and $D_{v_i v_i^{-1}}$ by gluing η_i along γ_i, as in Figure 2.8B, and let $q_i : \Delta_i \sqcup D_{v_i v_i^{-1}} \to \Delta_{i+1}$ be the corresponding quotient map. Because γ_i and η_i are both simple paths along the boundaries of their respective diagrams, Δ_{i+1} is planar and simply connected. Let $p'_j = q_i(p_j)$ for $j = 0, \ldots, n$ and let x'_i be the vertex on the boundary of $q_i(D_{v_i v_i^{-1}})$ that is not in $q_i(\eta_i)$. Note that, by the exact same reasoning as in the construction of Δ_1 through Δ_m, Δ_{i+1} is a van Kampen diagram for w, Δ_{i+1} is simply bounded, and $\mathrm{IDiam}(\Delta_{i+1}) \leq \max\left(\mathrm{IDiam}(\ell(w)), \left\lfloor \frac{\ell(w)}{2} \right\rfloor + M_{\mathcal{P}}\right)$.

We have only to show that $\mathrm{bdiam}(\Delta_{i+1}) > \mathrm{bdiam}(\Delta_i)$. First, I claim that for any $y \in \Delta_i^{(0)}$, we have that $d_{\Delta_{i+1}}(*, q_i(y)) = d_{\Delta_i}(*, y)$; gluing $D_{v_i v_i^{-1}}$ to Δ_i has not changed the distance from $*$ to the image of y. It is sufficient to show that $d_{\Delta_{i+1}}(*, q_i(y)) = d_{\Delta_i}(*, y)$ for $y \in \{x_i, p_{k-1}, p_{k+1}\}$. Now since $d_{\Delta_i}(*, x_i) = m_i$, $d_{\Delta_i}(*, p_{k-1}) \in \{m_i - 1, m_i, m_i + 1\}$. Suppose that $d_{\Delta_i}(*, p_{k-1}) = m_i - 1$. Then by the definition of m_i, $d_{\Delta_i}(*, p_{k-1}) = d_{\partial\Delta_i}(*, p_{k-1})$. So there is a Δ_i-geodesic along the boundary of Δ_i from $*$ to p_{k-1}. Since x_i is adjacent to p_{k-1} via a 1-cell on the boundary of Δ_i and $d_{\Delta_i}(*, p_{k-1}) + 1 = d_{\Delta_i}(*, x_i)$, this geodesic extends to a geodesic from $*$ to x_i along the boundary of Δ_i. This contradicts the fact that $x_i \in X_i$. Thus, $d_{\Delta_i}(*, p_{k-1}) \in \{m_i, m_i + 1\}$. The same argument shows that $d_{\Delta_i}(*, p_{k+1}) \in \{m_i, m_i + 1\}$. So the distances from $*$ to x_i, p_{k-1}, and p_{k+1} in Δ_i all differ from each other by at most 1. Now let $y \in \{x_i, p_{k-1}, p_{k+1}\}$ and let γ be a Δ_{i+1}-geodesic from $*$ to $q_i(y)$. Let $z \in \{x_i, p_{k-1}, p_{k+1}\}$ such that $q_i(z)$ is the first point at which γ enters $q_i(D_{v_i v_i^{-1}})$. If $z = y$, γ is contained in $q_i(\Delta_i)$, so $d_{\Delta_{i+1}}(*, q_i(y)) = d_{\Delta_i}(*, y)$. Otherwise, $z \neq y$. Then since γ is a Δ_{i+1}-geodesic which, up until $q_i(z)$, is contained in $q_i(\Delta_i)$, $d_{\Delta_{i+1}}(*, q_i(z)) = d_{\Delta_i}(*, z) \geq m_i$. Since distance increases along geodesics, this implies that $d_{\Delta_{i+1}}(*, q_i(y)) \geq m_i + 1 \geq d_{\Delta_i}(*, y)$. So $d_{\Delta_{i+1}}(*, q_i(y)) = d_{\Delta_i}(*, y)$.

Since $\partial\Delta_{i+1} \setminus q_i(\partial\Delta_i) = \{x_i'\}$ and $q_i(\partial\Delta_i) \setminus \partial\Delta_{i+1} = \{q_i(x_i)\}$, and $d_{\Delta_{i+1}}(*, q_i(y)) = d_{\Delta_i}(*, y)$ for all $y \in \partial\Delta_i$, we have that

$$\mathrm{bdiam}(\Delta_{i+1}) - \mathrm{bdiam}(\Delta_i) = d_{\Delta_{i+1}}(*, x_i') - d_{\Delta_i}(*, x_i).$$

Now any path from $*$ to x_i' in Δ_{i+1} passes through one of $q_i(x_i)$, $q_i(p_{k-1})$, or $q_i(p_{k+1})$. Hence, $d_{\Delta_{i+1}}(*, x_i') > m_i = d_{\Delta_i}(*, x_i)$. Thus, $\mathrm{bdiam}(\Delta_{i+1}) - \mathrm{bdiam}(\Delta_i) > 0$. Also, since Δ_i is a van Kampen diagram for w for all i, we showed above that $\text{bdiam}(\Delta_i) \leq \frac{1}{2}\ell(w)^2$.

Since $\text{bdiam}(\Delta_i)$ is increasing and bounded above, the construction of this sequence must end with some Δ_l. Then $X_l = \emptyset$, which implies that, for all $x \in \Delta_l^{(0)} \cap \partial\Delta_l$, $d_{\Delta_l}(*, x) = d_{\partial\Delta_l}(*, x)$. So Δ_l is geodesically bounded.

□

We used the following technical lemma to show that γ_j was a simple path when constructing Δ_{j+1} for $j+1 \in [m]$. We will now tie up this loose end.

Lemma 2.3.4. *Let Δ be a van Kampen diagram for a word w in a finite presentation in which no generator is equal to the identity. Let $p_1, p_2, \ldots, p_n = p_1$ be the sequence of vertices traversed in order by the path labeled by w along the boundary of Δ. Suppose there exist $i, j \in [n]$ with $i < j$ such that $p_i = p_j$ and $p_k \neq p_l$ whenever $i \leq k < l < j$. Then $p_{j-1} \neq p_{j+1}$.*

Proof. Note first the following general facts about the boundary circuit of a van Kampen diagram. Recall that each 1-cell in a van Kampen diagram is associated to two directed edges, one going each direction. The boundary of each 2-cell is a circuit which can be directed either clockwise or counterclockwise. It is possible to choose a directed boundary circuit for each 2-cell (going either clockwise or counterclockwise) such that each directed edge in the diagram appears exactly once in either the directed boundary circuit of one of the 2-cells or the boundary circuit for the diagram, but not both; see, for example, [19, p. 236]. So each directed edge corresponding to a 1-cell on the boundary of a van Kampen diagram appears at most once in the boundary circuit of the diagram. If both directed edges appear in the boundary circuit, the corresponding 1-cell must be incident to the complement of the diagram on both sides, since otherwise one of the directed edges would appear in the directed boundary circuit of a 2-cell. If only one of the two directed edges appears in the boundary circuit, then the other appears in the directed boundary circuit of a 2-cell, so the corresponding 1-cell is incident to a 2-cell of the diagram on one side and incident to the complement of the diagram on the other.

If e is a directed edge, let $\bar{e}$ denote the other edge corresponding to the same 1-cell, going in the opposite direction.

Let $i, j \in [n]$ with $i < j$ such that $p_i = p_j$ and $p_k \neq p_l$ whenever $i \leq k < l < j$. Suppose by way of contradiction that $p_{j-1} = p_{j+1}$. Let e be the directed edge in the boundary circuit from p_{j-1} to p_j, and let f be the directed edge in the boundary circuit from p_j to p_{j+1}. Let $\gamma = ef$, the directed path in the boundary circuit from p_{j-1} to p_{j+1}. Since no generator is equal to the identity, we know that $p_j \neq p_{j-1} = p_{j+1}$. So either γ is a simple circuit or $e = \bar{f}$. Let η be the directed path in the boundary circuit from the ith vertex p_i to the jth vertex p_j. Note that e is the last edge of η and η is a circuit that does not repeat any vertices, by the assumption that $p_k \neq p_l$ whenever $i \leq k < l < j$. Therefore, either η is a simple circuit or $\eta = \bar{e}e$.

Suppose first that $e = \bar{f}$. In this case, both e and $\bar{e}$ appear in γ, and therefore both appear in the boundary circuit, so e is incident to $\mathbb{R}^2 \setminus \Delta$ on both sides. If η is a simple circuit and e appears in η, e is incident to the interior of η. This would imply that the interior of η, a simple circuit in Δ, contains a point in $\mathbb{R}^2 \setminus \Delta$, contradicting the fact that Δ is simply-connected. If $\eta = \bar{e}e$, then $\bar{e}$ would appear twice in the boundary circuit: once on η and once on γ. This is a contradiction, so $e \neq \bar{f}$.

So suppose instead that γ is a simple circuit. In this case, e and f are both incident to a 2-cell on one side, and thus neither $\bar{e}$ nor $\bar{f}$ appear in the boundary circuit. Therefore, $\eta \neq \bar{e}e$. If instead η is a simple circuit, we are in the case shown in Figure 2.9. Then $\text{int}(f)$ must either be in the interior or exterior of η; f cannot be a part of η, since it appears only once in the boundary circuit in γ, and we know that

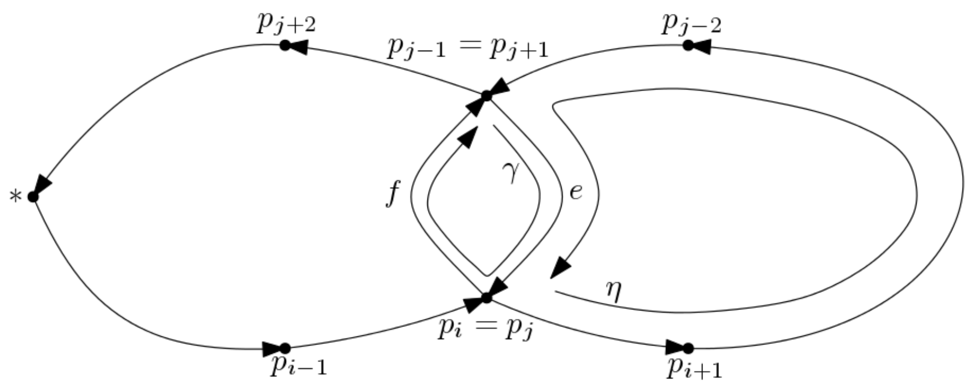

Figure 2.9: One possible drawing of Δ in the case where γ and η are both simple circuits. The arrows indicate the direction of the boundary circuit as well as the orientations of e and f.

$\bar{f}$ does not appear on the boundary circuit. f cannot be in the interior of η, since it is incident to the complement of Δ, so $\text{int}(f)$ is in the exterior of η. In this case, either $\eta \setminus e$ is in the interior of γ or neither contains the other. Now each edge of η is incident to the complement of Δ, so $\eta \setminus e$ cannot be in the interior of γ; then its interior would contain a point in the complement of Δ. So the interiors of γ and η do not intersect. But since e is on both γ and η, it is incident to the interiors of both. Since these interiors do not intersect, e must be incident to the interior of γ on one side and the interior of η on the other. This contradicts the fact that e is incident to the complement of Δ. This final contradiction implies that $p_{j-1} \neq p_{j+1}$.

□

2.4 Main Theorem

In this section, we will finish the proof of Theorem 2.4.4. In order to simplify the proof, we will first define a variant of intrinsic tame filling functions that only relies on distance to the 1-skeleton of a van Kampen diagram and prove that such functions are equivalent to tame filling functions. This will allow us to ignore distance to 2-cells in the remainder of the proof.

Definition 2.4.1. Let $f : \mathbb{N}[\frac{1}{4}] \to \mathbb{N}[\frac{1}{4}]$ be a function, C a 2-complex with basepoint $*$ and D a subcomplex of $C^{(1)}$. A 1-combing Ψ of the pair (C, D) is *graph f-tame* if for all $n \in \mathbb{N}[\frac{1}{4}]$ and for all $x \in D$ and $s, t \in [0, 1]$ such that $s \leq t$ and $\Psi(x, s), \Psi(x, t) \in C^{(1)}$, if $d_C(*, \Psi(x, t)) \leq n$, then $d_C(*, \Psi(x, s)) \leq f(n)$.

Definition 2.4.2. Let $f : \mathbb{N}[\frac{1}{4}] \to \mathbb{N}[\frac{1}{4}]$ be non-decreasing. f is an *intrinsic graph tame filling function* for $\langle A|R\rangle$ if, for all $w \in (A \cup A^{-1})^*$ with $w =_G 1$, there is a van Kampen diagram Δ_w for w with basepoint $*$ and a 1-combing Ψ_w of $(\Delta_w, \partial\Delta_w)$ based at $*$ such that Ψ_w is graph f-tame.

Lemma 2.4.3. *Every intrinsic graph tame filling function for $\langle A|R\rangle$ is equivalent to an intrinsic tame filling function for $\langle A|R\rangle$. Conversely, every intrinsic tame filling function is itself an intrinsic graph tame filling function.*

Proof. We will first show that all intrinsic tame filling functions are intrinsic graph tame filling functions. Let h be an intrinsic tame filling function and let $w \in (A \cup A^{-1})^*$ with $w =_G 1$. Let Δ_w be a van Kampen diagram for w and Ψ_w an h-tame 1-combing of Δ_w. Since Ψ_w is h-tame, we have that for all $n \in \mathbb{N}[\frac{1}{4}]$, $p \in \partial\Delta_w$, and $s, t \in [0, 1]$ such that $s \leq t$, if $d_{\Delta_w}(*, \Psi_w(p, t)) \leq n$, then $d_{\Delta_w}(*, \Psi_w(p, s)) \leq h(n)$. Since this is true for all such p, s, and t, it is also true whenever $\Psi(p, s), \Psi(p, t) \in \Delta_w^{(1)}$. Therefore, Ψ_w is graph h-tame. So h is also an intrinsic graph tame filling function.

For the other direction, let f be an intrinsic graph tame filling function. Let $\rho = \max\{\ell(r)|r \in R\}$. Define $g : \mathbb{N}[\frac{1}{4}] \to \mathbb{N}[\frac{1}{4}]$ by $g(n) = f\left(n + \frac{3}{4}\right) + \frac{\rho}{2} - \frac{1}{4}$. Note that g is equivalent to f. We will show that g is an intrinsic tame filling function for $\langle A|R\rangle$.

Let $w \in (A \cup A^{-1})^*$ with $w =_G 1$. Let Δ_w be a van Kampen diagram for w and Ψ_w a graph f-tame 1-combing of Δ_w. Let $n \in \mathbb{N}[\frac{1}{4}]$, $p \in \partial\Delta_w$, and $s, t \in [0, 1]$ such that $s \leq t$ and $d_{\Delta_w}(*, \Psi_w(p,t)) \leq n$. We will show that Ψ_w is g-tame by showing that $d_{\Delta_w}(*, \Psi_w(p,s)) \leq g(n)$. See Figure 2.10.

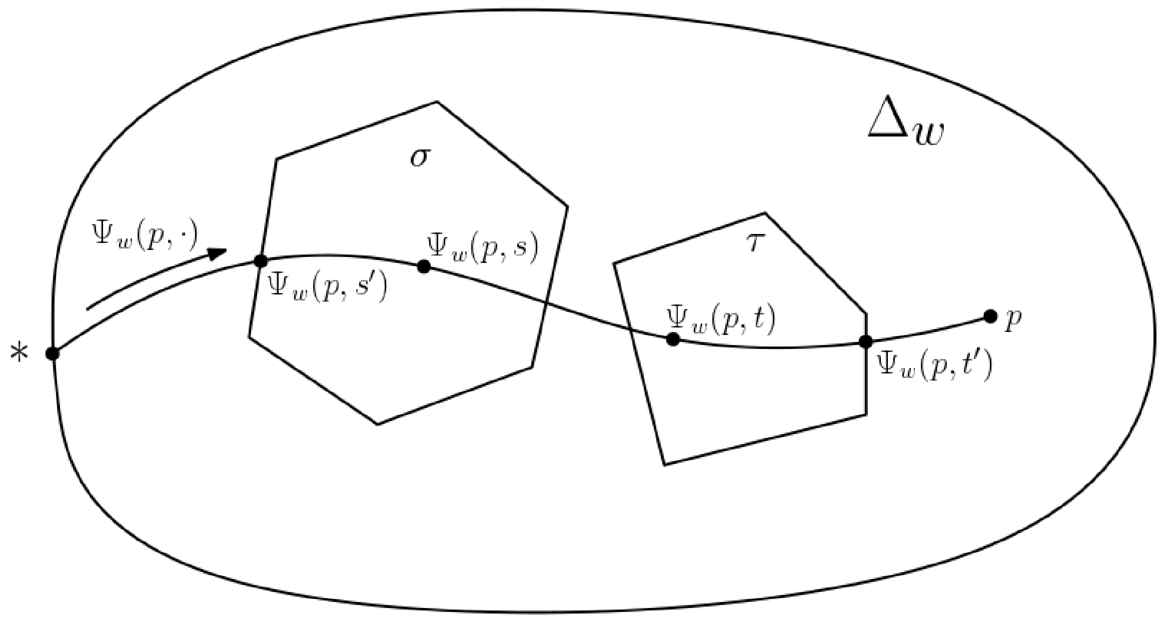

Figure 2.10: Showing Ψ_w is g-tame.

Choose $t' \in [0,1]$ such that $t' \geq t$ in the following way. If $\Psi_w(p,t) \in \Delta_w^{(1)}$, let $t' = t$. Otherwise, let τ be the 2-cell with $\Psi_w(p,t)$ in its interior. Since $p \in \Delta_w^{(1)}$, and $\Psi_w(p,\cdot)$ is continuous, there is some $t' > t$ such that $\Psi_w(p,t') \in \partial\tau$. Then by the definition of the coarse distance function, $d_{\Delta_w}(*, \Psi_w(p,t')) \leq d_{\Delta_w}(*, \Psi_w(p,t)) + \frac{3}{4} \leq n + \frac{3}{4}$.

Now choose $s' \in [0,1]$ such that $s' \leq s$ in the following way. If $\Psi_w(p,s) \in \Delta_w^{(1)}$, let $s' = s$. Otherwise, let σ be the 2-cell with $\Psi_w(p,s)$ in its interior. Since $* \in \Delta_w^{(1)}$, and $\Psi_w(p,\cdot)$ is continuous, there is some $s' < s$ such that $\Psi_w(p,s') \in \partial\sigma$. Now there are at most ρ 1-cells and ρ vertices in $\partial\sigma$. Suppose q is a point in $\partial\sigma$ with maximum coarse distance from $*$. Then $d_{\Delta_w}(*, q) \geq d_{\Delta_w}(*, \Psi_w(p,s)) + \frac{1}{4}$. Now, there is a path in $\partial\sigma$ from q to $\Psi_w(p,s')$ that passes through at most ρ 1-cells not containing q and ρ vertices not equal to q. At each step from a vertex to the interior of a 1-cell or the interior of a 1-cell to a vertex, the coarse distance changes by $\frac{1}{2}$. Since there are at most ρ such steps from q to $\Psi_w(p, s')$,

$$d_{\Delta_w}(*, \Psi_w(p, s')) \geq d_{\Delta_w}(*, q) - \frac{\rho}{2} \geq d_{\Delta_w}(*, \Psi_w(p, s)) + \frac{1}{4} - \frac{\rho}{2}.$$

But since Ψ_w is graph f-tame, $s' \leq t'$, and $d_{\Delta_w}(*, \Psi_w(p, t')) \leq n + \frac{3}{4}$, we know that $d_{\Delta_w}(*, \Psi_w(p, s')) \leq f\left(n + \frac{3}{4}\right)$. This and the above inequality imply that

$$d_{\Delta_w}(*, \Psi_w(p, s)) \leq f\left(n + \frac{3}{4}\right) + \frac{\rho}{2} - \frac{1}{4} = g(n).$$

So Ψ_w is g-tame for all such w. Therefore, g is an intrinsic tame filling function for $\langle A|R\rangle$.

□

Theorem 2.4.4. *Given a finite presentation $\mathcal{P} = \langle A|R\rangle$ such that for all $a \in A$, a is not equal to the identity, there is an intrinsic tame filling function for $\mathcal{P}$ that is equivalent to the intrinsic diameter function of $\mathcal{P}$.*

Proof. Define $f : \mathbb{N}[\frac{1}{4}] \to \mathbb{N}[\frac{1}{4}]$ by

$$f(n) = \text{IDiam}(\lceil 2n + 1\rceil) + n + M_{\mathcal{P}} + 1.$$

Note that f is strictly increasing, since IDiam is an increasing function. Note also that f is equivalent to IDiam. We will show that f is an intrinsic graph tame filling function, and then apply Lemma 2.4.3.

Let $w =_G 1$ and let Δ_0 be any van Kampen diagram for w. Let T_0 be a tree of Δ_0-geodesics out of $*$ and let Ψ_{T_0} be a 1-combing of $(\Delta_0, \partial\Delta_0)$ based at $*$ respecting T_0. We will construct a sequence $(\Delta_0, T_0, \Psi_{T_0}), (\Delta_1, T_1, \Psi_{T_1}), \ldots, (\Delta_n, T_n, \Psi_{T_n})$ where each Δ_i is a van Kampen diagram for w, each T_i is a spanning tree of Δ_i-geodesics out of $*$, and each Ψ_{T_i} is a 1-combing of $(\Delta_i, \partial\Delta_i)$ respecting T_i, and then show that Ψ_{T_n} is graph f-tame.

Given that we have constructed Δ_i, T_i, and Ψ_{T_i}, if Ψ_{T_i} is graph f-tame, then $n = i$ and we are done. Otherwise, construct Δ_{i+1}, T_{i+1}, and $\Psi_{T_{i+1}}$ as follows. We will start with the construction of Δ_{i+1} by finding 1-cells in Δ_i whose T_i-icicles contain points too far away from $*$ for Ψ_{T_i} to be f-tame. We will then replace the bodies of these icicles in Δ_i with diagrams that have smaller diameter using Proposition 2.3.3; the resulting diagram will be Δ_{i+1}.

Recall the definition of the points prior to x in a 1-combing Ψ of the pair (C, D):

$$P(\Psi, x) = \{y \in C | \text{there is } p \in D \text{ and } 0 \le s \le t \le 1 \text{ with } \Psi(p,t) = x \text{ and } \Psi(p,s) = y\}.$$

For $x \in \Delta_i^{(1)}$, let

$$M_i(x) = \max\{d_{\Delta_i}(*, y) | y \in \Delta_i^{(1)} \cap P(\Psi_{T_i}, x)\},$$

the largest distance from $*$ that occurs in the 1-skeleton prior to x in Ψ_{T_i}. Note that, since Ψ_{T_i} is not graph f-tame, there exists $x \in \Delta_i^{(1)}$ such that $M_i(x) > f(d_{\Delta_i}(*, x))$. Then let

$$N_i = \max\{d_{\Delta_i}(*, x) | x \in \Delta_i^{(1)} \text{ and } M_i(x) > f(d_{\Delta_i}(*, x))\}.$$

Note that this maximum exists, since Δ_i is a finite complex, and therefore distances in the complex are bounded. Also note that if $x \in T_i$ and $\Psi_{T_i}(p,t) = x$, since Ψ_{T_i} respects T_i, we have that for all $s \le t$, $\Psi_{T_i}(p,s)$ is on the simple path in T_i from $*$ to x. Since this path is a Δ_i-geodesic, every point prior to x is at least as close to $*$ as x, so $M_i(x) = d_{\Delta_i}(*, x) < f(d_{\Delta_i}(*, x))$. Hence, if $d_{\Delta_i}(*, x) = N_i$ and $M_i(x) > f(N_i)$, then $x \notin T_i$.

Therefore, let

$$E_i = \{e \in E(\Delta_i \setminus T_i) | d_{\Delta_i}(*, e) = N_i \text{ and there exists } x \in e \text{ such that } M_i(x) > f(N_i)\}.$$

Define a partial order $\leq_i$ on E_i by $e \leq_i e'$ if and only if e is contained in the T_i-icicle at e' (and therefore the T_i-icicle at e is contained in the T_i-icicle at e' by Lemma 2.1.3). Let $F_i \subseteq E_i$ be the set of maximal elements of E_i with respect to $\leq_i$. Let $F_i = \{e_1, \ldots, e_{m_i}\}$. It is the bodies of the T_i-icicles at these 1-cells that we will replace to construct Δ_{i+1}.

For each $j = 1, \ldots, m_i$, let I_j be the T_i-icicle at e_j and let D_j be the body of I_j. Let α_j be the tail of I_j, and let $*_j$ be the vertex at which the tail and body of I_j meet. Let x_j and y_j be the endpoints of e_j and let $\vec{e_j}$ be the directed edge from x_j to y_j corresponding to e_j. Let β_{x_j} and β_{y_j} be the unique simple paths in T from $*_j$ to x_j and y_j, respectively. Let $\gamma_j = \beta_{x_j} \cdot \vec{e_j} \cdot \overline{\beta_{y_j}}$. Note that γ_j is a simple circuit that bounds D_j. Let w_j be the word labeling γ_j. Then we may consider D_j to be a van Kampen diagram for w_j with basepoint $*_j$.

Since we know that D_j is simply bounded, let D'_j be a simply and geodesically bounded van Kampen diagram for w_j such that

$$\mathrm{IDiam}(D'_j) \leq \max\left(\mathrm{IDiam}(\ell(w_j)), \left\lfloor \frac{\ell(w_j)}{2} \right\rfloor + M_{\mathcal{P}}\right),$$

as is guaranteed to exist by Proposition 2.3.3.

Let

$$\widehat{\Delta_i} = \Delta_i \setminus \bigcup_{j \in [m_i]} \mathrm{int}(D_j).$$

Note that since each $e \in F_i$ is maximal in $\leq_i$, for distinct 1-cells $e, e' \in F_i$, the

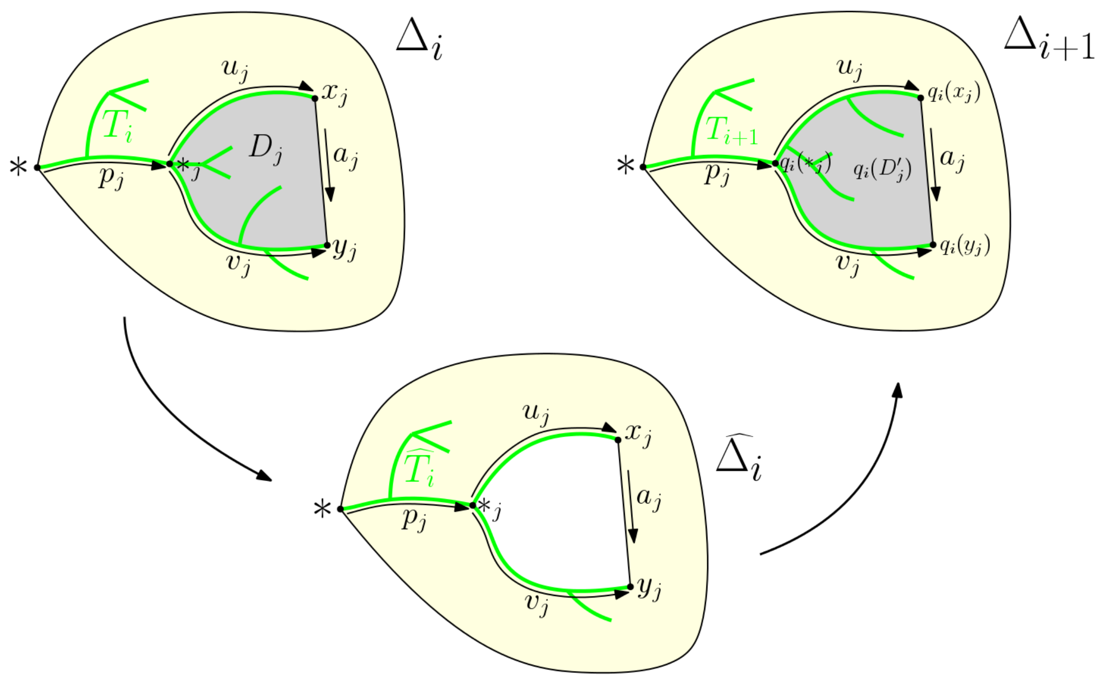

Figure 2.11: Constructing Δ_{i+1} and T_{i+1}

interiors of the T_i-icicles at e and e' do not intersect by Lemma 2.1.3. So if $k \neq j$, $D_k \cap \operatorname{int}(D_j) = \emptyset$ in Δ_i. Therefore, for each $j \in [m_i]$, $\partial D_j \subseteq \widehat{\Delta_i}$.

Construct Δ_{i+1} from $\widehat{\Delta_i}$ by gluing the basepoint of D'_j to $*_j$ and then gluing the boundary of D'_j to γ_j, gluing vertices to vertices and 1-cells to 1-cells with the same labels, for each $j \in [m_i]$, as depicted in Figure 2.11. Let

$$q_i : \widehat{\Delta_i} \sqcup \bigsqcup_{j \in [m_i]} D'_j \to \Delta_{i+1}$$

be the corresponding quotient map. Note that, since each D_j and D'_j are simply bounded, q_i does not identify 1-cells or vertices that were distinct in Δ_i, or in any D'_j. As a result, Δ_{i+1} is planar and simply connected. Since we have only replaced a subset of the interior of Δ_i, Δ_{i+1} is still a van Kampen diagram for w.

We now move on to the construction of T_{i+1}. Let $\widehat{T_i} = T_i \cap \widehat{\Delta_i}$. T_{i+1} will be an extension of $q_i(\widehat{T_i})$. I claim that $\widehat{T_i}$ is a spanning tree of $\widehat{\Delta_i}$. Note that since T_i is a tree and $\widehat{T_i} \subseteq T_i$, there are no simple circuits in $\widehat{T_i}$. Also, since T_i is a spanning tree of Δ_i, $\widehat{T_i}$ contains every vertex in $\widehat{\Delta_i}$. So we only need to show that $\widehat{T_i}$ is connected. It suffices to show that, for every $p \in \widehat{\Delta_i}^{(0)}$, there is a path from $*$ to p in $\widehat{T_i}$. Let η be the unique simple path from $*$ to p in T_i. Suppose that for some $j \in [m_i]$, η intersects D_j. Then let q be the last vertex of η that is in D_j. Since $p \notin \mathrm{int}(D_j)$, we know that $q \in \partial D_j$. Without loss of generality, assume that q lies on the simple path from $*$ to x_j in T_i. Therefore, the segment of η from $*$ to q is contained in the path from $*$ to x_j in T_i, which does not intersect $\mathrm{int}(D_j)$. Since q is the last vertex of η in D_j, this implies that η does not intersect $\mathrm{int}(D_j)$. Therefore, $\eta \subseteq \widehat{T_i}$.

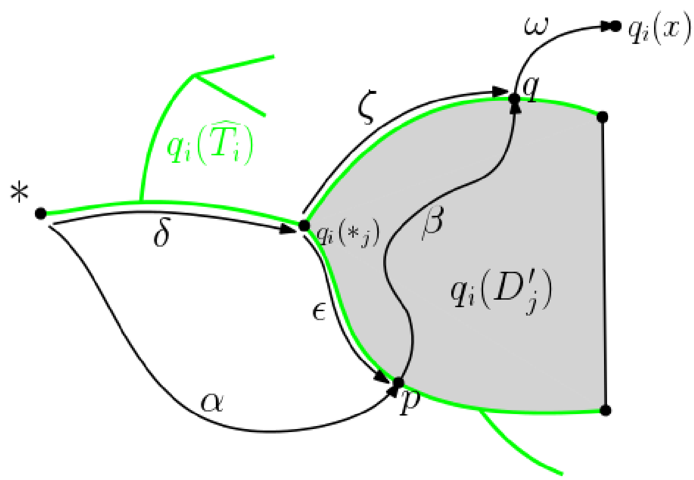

Figure 2.12: Constructing a Δ_{i+1}-geodesic η to stay inside $\widehat{\Delta_i}$ longer than γ.

I also claim that $q_i(\widehat{T_i})$ is a tree of Δ_{i+1}-geodesics out of $*$. Since $\widehat{T_i}$ is a tree of $\widehat{\Delta_i}$-geodesics out of $*$, it is sufficient to show that for every $x \in \widehat{\Delta_i}^{(0)}$, we have that $d_{\widehat{\Delta_i}}(*, x) \leq d_{\Delta_{i+1}}(*, q_i(x))$. Suppose by way of contradiction that there is some $x \in \widehat{\Delta_i}^{(0)}$ such that no Δ_{i+1}-geodesic from $*$ to $q_i(x)$ is contained in $q_i(\widehat{\Delta_i})$. Since $* \in q_i(\widehat{\Delta_i})$, each Δ_{i+1}-geodesic from $*$ to $q_i(x)$ has some initial segment contained in $q_i(\widehat{\Delta_i})$. Let γ be a Δ_{i+1}-geodesic from $*$ to $q_i(x)$ having the longest such initial segment. Let α be this initial segment of γ, and let p be its endpoint. Let q be the next vertex in $q_i(\widehat{\Delta_i})$ on γ after p, let β be the segment of γ from p to q, and let ω be the final segment of γ from q to $q_i(x)$. Note that by assumption $p \neq q$, so $|\beta| \geq 1$ and all but the endpoints of β are not contained in $q_i(\widehat{\Delta_i})$. Hence, $\beta \subseteq q_i(D'_j)$ for some $j \in [m]$ and $p, q \in q_i(\partial D'_j)$. Let δ be the simple path from $*$ to $q_i(*_j)$ in $q_i(\widehat{T_i})$ and let ϵ be the simple path from $q_i(*_j)$ to p in $q_i(\widehat{T_i})$. Finally, since D'_j is geodesically bounded, there is a $q_i(D'_j)$-geodesic ζ from $q_i(*_j)$ to q in $q_i(\partial D'_j)$. See Figure 2.12. Now let $\eta = \delta \cdot \zeta \cdot \omega$. Note that $|\alpha| = |\delta \cdot \epsilon|$, since $\delta \cdot \epsilon$ is the simple path from $*$ to p in $q_i(\widehat{T_i})$, and therefore a $q_i(\widehat{\Delta_i})$-geodesic from $*$ to p. Also, $|\epsilon \cdot \beta| \geq |\zeta|$, since $\epsilon \cdot \beta$ is a path from $*_j$ to q in $q_i(D'_j)$. Therefore,

$$|\gamma| = |\alpha \cdot \beta \cdot \omega| = |\delta \cdot \epsilon \cdot \beta \cdot \omega| \geq |\delta \cdot \zeta \cdot \omega| = |\eta|.$$

Hence, η is also a Δ_{i+1}-geodesic from $*$ to x. Now, the initial segment of η contained in $q_i(\widehat{\Delta_i})$ contains $\delta \cdot \zeta$, a Δ_{i+1}-geodesic from $*$ to q. Since $\alpha \cdot \beta$ is also a Δ_{i+1}-geodesic from $*$ to q and $|\beta| \geq 1$, we have that $\delta \cdot \zeta$ is strictly longer than α. So η's initial segment inside $q_i(\widehat{\Delta_i})$ is longer than that of γ. This contradicts the way that γ was chosen. Therefore, for all $x \in \widehat{\Delta_i}^{(0)}$, there is a Δ_{i+1}-geodesic from $*$ to $q_i(x)$ contained in $q_i(\widehat{\Delta_i})$. So $d_{\widehat{\Delta_i}}(*, x) \leq d_{\Delta_{i+1}}(*, q_i(x))$ as desired.

Finally, let T_{i+1} be an extension of $q_i(\widehat{T_i})$ to a geodesic spanning tree of Δ_{i+1}. Now we will choose $\Psi_{T_{i+1}}$ to be a particular 1-combing of $(\Delta_{i+1}, \partial\Delta_{i+1})$ based at $*$ respecting T_{i+1}, as guaranteed to exist by Proposition 2.2.2. I claim that we may construct $\Psi_{T_{i+1}}$ such that, for all $x \in \partial\Delta_i$ and $t \in [0,1]$ with $\Psi_{T_{i+1}}(q_i(x), [t, 1]) \subset q_i(\widehat{\Delta_i})$, there exists $t' \in [0,1]$ such that $\Psi_{T_{i+1}}(q_i(x), [t, 1]) = q_i(\Psi_{T_i}(x, [t', 1]))$. In other words, the images of the combing paths of $q_i \circ \Psi_{T_i}$ and $\Psi_{T_{i+1}}$ agree on end portions that stay in $q_i(\widehat{\Delta_i})$. This will be useful for showing that $\Psi_{T_{i+1}}$ retains the progress that Ψ_{T_i} has made towards achieving f-tameness.

To choose $\Psi_{T_{i+1}}$ in this way, recall that the images of individual combing paths of the 1-combings constructed in Section 2.2 are determined by the bijection θ_T between 1-cells outside of the spanning tree T and 2-cells, as well as a choice of homeomorphism from each 2-cell σ to D^2 that maps the edge $\theta_T^{-1}(\sigma)$ to S^1_+. Since we chose T_{i+1} to agree with $q_i(T_i)$ on $\widehat{\Delta_i}$, for any 1-cell $e \in E(\widehat{\Delta_i} \setminus T_i)$, we have that $q_i(e) \notin T_{i+1}$ and $q_i(\partial I_e) = \partial I_{q_i(e)}$. Therefore, for any 2-cell $\sigma \subseteq \widehat{\Delta_i}$, if $\sigma \subseteq I_e$, then $q_i(\sigma) \subseteq I_{q_i(e)}$. Hence,

$$q_i \circ \theta_{T_i}|_{E(\widehat{\Delta_i} \setminus T_i) \setminus F_i} = \theta_{T_{i+1}} \circ q_i|_{E(\widehat{\Delta_i} \setminus T_i) \setminus F_i}.$$

As a result, for each 2-cell $\sigma \subseteq \widehat{\Delta_i}$, if $\phi_\sigma : \sigma \to D^2$ is the homeomorphism chosen for σ in the construction of Ψ_{T_i}, we may choose $\phi_{q_i(\sigma)} = \phi_\sigma \circ (q_i|_\sigma)^{-1}$ as the homeomorphism from $q_i(\sigma)$ to D^2 in the construction of $\Psi_{T_{i+1}}$. Figure 2.13 represents this choice of $\phi_{q_i(\sigma)}$ pictorially. Having chosen these homeomorphisms, the images of the combing paths of $\Psi_{T_{i+1}}$ on $q_i(\widehat{\Delta_i})$ are determined, and we are free to make any other choices necessary to finish the construction of $\Psi_{T_{i+1}}$ as in Section 2.2.

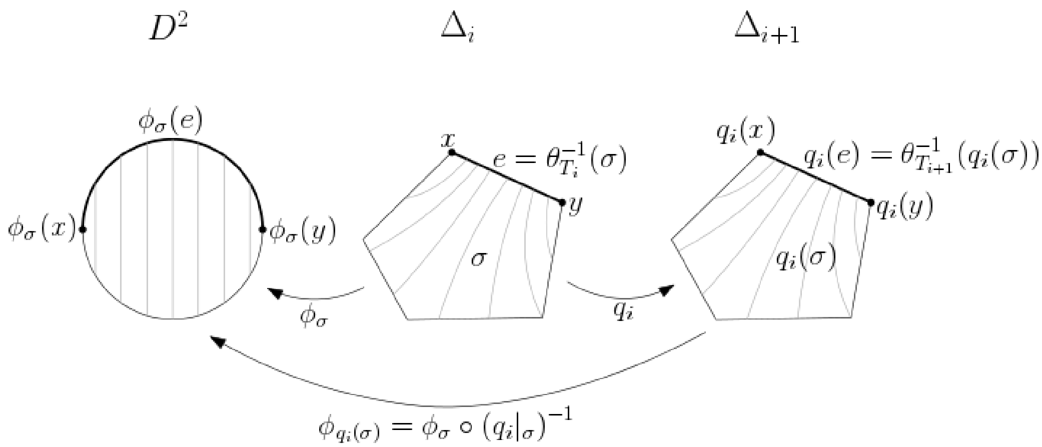

Figure 2.13: Choosing $\phi_{q_i(\sigma)}$ for a 2-cell σ in $\widehat{\Delta_i}$. The lighter curves depict some combing paths of Ψ_{T_i} and their images. As a result of this choice, the combing paths of $\Psi_{T_{i+1}}$ on $q_i(\sigma)$ will agree with those of $q_i \circ \Psi_{T_i}$.

Having chosen $\Psi_{T_{i+1}}$, let $x \in \partial\Delta_i$ and $t \in [0,1]$ with $\Psi_{T_{i+1}}(q_i(x), [t,1]) \subset q_i(\widehat{\Delta_i})$.

We need to show that there exists $t' \in [0,1]$ such that $\Psi_{T_{i+1}}(q_i(x),[t,1]) = q_i(\Psi_{T_i}(x,[t',1]))$.

By Lemma 2.2.3, we know there exists a sequence $1 \geq t_1 > \cdots > t_{m+1} > t_{m+2} = 0$ and a sequence of 1-cells $e_1, e_2, \ldots, e_m$ in $\Delta_{i+1} \setminus T_{i+1}$ such that, if we let $x_j = \Psi_{T_{i+1}}(q_i(x), t_j)$ and $\sigma_j = \theta_{T_{i+1}}(e_j)$, we have that

- $\Psi_{T_{i+1}}(q_i(x),[0,t_{m+1}])$ is the simple path from $*$ to x_m in T_{i+1},
- $\Psi_{T_{i+1}}(q_i(x),[t_1,1]) = \{q_i(x)\}$, and
- for all $j \in [m]$,
  - $x_j \in \operatorname{int}(e_j)$,
  - $t_j = \min\{s \in [0,1] | \Psi_{T_{i+1}}(q_i(x),s) = x_j\}$, and
  - $\Psi_{T_{i+1}}(q_i(x),[t_{j+1},t_j]) = \phi^{-1}_{\sigma_j}(h(\phi_{\sigma_j}(x_j),[0,1]))$.

Similarly, there is a sequence $1 \geq t'_1 > \cdots > t'_{l+1} > t'_{l+2} = 0$ and a sequence of 1-cells $e'_1, e'_2, \ldots, e'_l$ in $\Delta_i \setminus T_i$ such that, if we let $x'_j = \Psi_{T_i}(x, t'_j)$ and $\sigma'_j = \theta_{T_i}(e'_j)$, we have that

- $\Psi_{T_i}(x,[0,t'_{m+1}])$ is the simple path from $*$ to x'_m in T_i,
- $\Psi_{T_i}(x,[t'_1,1]) = \{x\}$, and
- for all $j \in [l]$,
  - $x'_j \in \operatorname{int}\big(e'_j\big)$,
  - $t'_j = \min\{s \in [0,1] | \Psi_{T_i}(x,s) = x'_j\}$, and
  - $\Psi_{T_i}(x,[t'_{j+1},t'_j]) = \phi^{-1}_{\sigma'_j}(h(\phi_{\sigma'_j}(x'_j),[0,1]))$.

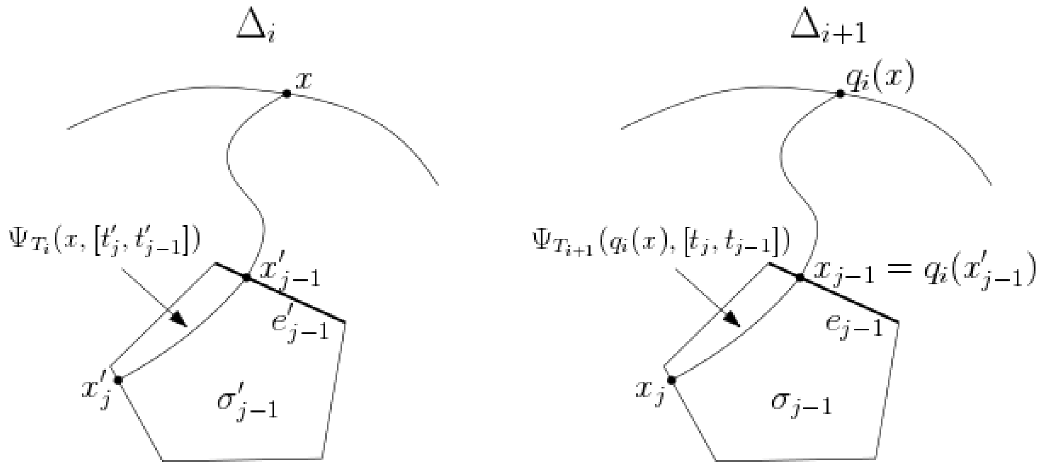

Figure 2.14: The situation at step j of the induction, demonstrating that $\Psi_{T_{i+1}}(q_i(x),[t_j,1]) = q_i(\Psi_{T_i}(x,[t'_j,1]))$.

If $t_1 \leq t \leq 1$, then we may simply let $t' = 1$. Then

$$\Psi_{T_{i+1}}(q_i(x),[t,1]) = \{q_i(x)\} = q_i(\Psi_{T_i}(x,[t',1])).$$

Otherwise, there is some $k \in [m+1]$ such that $t_{k+1} \leq t < t_k$. For $j \in [k]$, suppose by induction that $\Psi_{T_{i+1}}(q_i(x),[t_{j-1},1]) = q_i(\Psi_{T_i}(x,[t'_{j-1},1]))$ and that $x_{j-1} = q_i(x'_{j-1})$. This setup is shown in Figure 2.14. Then $e_{j-1} = q_i(e'_{j-1})$. Since $\theta_{T_{i+1}}(e_{j-1}) = \sigma_{j-1}$ and t_{j-1} is the smallest time with $\Psi_{T_{i+1}}(q_i(x),t_{j-1}) = x_{j-1}$, we must have that $\mathrm{int}(\sigma_{j-1}) \cap \Psi_{T_{i+1}}(q_i(x),[t_j,t_{j-1}]) \neq \emptyset$. We know that $\Psi_{T_{i+1}}(q_i(x),[t_j,t_{j-1}]) \subseteq q_i(\widehat{\Delta_i})$, so we must have $\sigma_{j-1} \subseteq q_i(\widehat{\Delta_i})$. Hence,

$$\sigma_{j-1} = \theta_{T_{i+1}}(e_{j-1}) = \theta_{T_{i+1}}(q_i(e'_{j-1})) = q_i(\theta_{T_i}(e'_{j-1})) = q_i(\sigma'_{j-1}).$$

Then since we chose $\phi_{\sigma_{j-1}} = \phi_{\sigma'_{j-1}} \circ (q_i|_{\sigma'_{j-1}})^{-1}$ in the construction of $\Psi_{T_{i+1}}$,

$$\begin{aligned}\Psi_{T_{i+1}}(q_i(x),[t_j,t_{j-1}]) &= \phi^{-1}_{\sigma_{j-1}}(h(\phi_{\sigma_{j-1}}(x_{j-1}),[0,1]))\\ &= q_i(\phi^{-1}_{\sigma'_{j-1}}(h(\phi_{\sigma'_{j-1}}(x'_{j-1}),[0,1])))\\ &= q_i(\Psi_{T_i}(x,[t'_j,t'_{j-1}])).\end{aligned}$$

Combining this result with the induction hypothesis gives that

$$\Psi_{T_{i+1}}(q_i(x),[t_j,1]) = q_i(\Psi_{T_i}(x,[t'_j,1])).$$

Also,

$$\begin{aligned}x_j = \Psi_{T_{i+1}}(q_i(x),t_j) &= \phi^{-1}_{\sigma_{j-1}}(h(\phi_{\sigma_{j-1}}(x_{j-1}),1))\\ &= q_i(\phi^{-1}_{\sigma'_{j-1}}(h(\phi_{\sigma'_{j-1}}(x'_{j-1}),1)))\\ &= q_i(\Psi_{T_i}(x,t'_j)) = q_i(x'_j).\end{aligned}$$

This completes the induction, establishing that $\Psi_{T_{i+1}}(q_i(x),[t_k,1]) = q_i(\Psi_{T_i}(x,[t'_k,1]))$ and that $x_k = q_i(x'_k)$.

Suppose $k \leq m$. Since $x_k = q_i(x'_k)$, we again know that $e_k = q_i(e'_k)$. Now since $t_k = \min\{s \in [0,1] | \Psi_{T_{i+1}}(q_i(x),s) = x_k\}$ and $t < t_k$, we know that $\Psi_{T_{i+1}}(q_i(x),t) \neq x_k$. Hence, $\Psi_{T_{i+1}}(q_i(x),[t,t_k])$ contains a point in $\text{int}(\sigma_k)$. Just as above, this implies that $\sigma_k = q_i(\sigma'_k)$ and therefore $\phi_{\sigma_k} = \phi_{\sigma'_k} \circ (q_i|_{\sigma'_k})^{-1}$. Hence,

$$\Psi_{T_{i+1}}(q_i(x),[t,t_k]) = \phi^{-1}_{\sigma_k}(h(\phi_{\sigma_k}(x_{k-1}),[0,s])), \text{ for some } s \in (0,1],$$
$$= q_i(\phi^{-1}_{\sigma'_k}(h(\phi_{\sigma'_k}(x'_{k-1}),[0,s])))$$
$$= q_i(\Psi_{T_i}(x,[t',t'_k])),$$

where $t' \in [t'_{k+1}, t'_k)$ such that $\phi_{\sigma'_k}(\Psi_{T_i}(x,t')) = h(\phi_{\sigma'_k}(x'_{k-1}), s) = \phi_{\sigma_k}(\Psi_{T_{i+1}}(q_i(x),t))$.

If instead $k = m+1$, then $\Psi_{T_{i+1}}(q_i(x),[0,t_k])$ is the unique simple path from $x_{m+2} = *$ to x_{m+1} in T_{i+1}, and $y := \Psi_{T_{i+1}}(q_i(x),t)$ is a point along that path. Since $x_{m+1} \in q_i(\widehat{\Delta_i})$ and we chose $T_{i+1} \subseteq q_i(\widehat{T_i})$, we have that $\Psi_{T_{i+1}}(q_i(x),[0,t_k]) \subseteq q_i(\widehat{\Delta_i})$. Let $y' = (q_i|_{\widehat{\Delta_i}})^{-1}(y)$. Again since $T_{i+1} \subseteq q_i(\widehat{T_i})$, we have that y' is on the unique simple path from $*$ to x'_{m+1}, so $y' \in \Psi_{T_i}(x,[0,t'_k])$. So let $t' \in [0,t'_k]$ such that $\Psi_{T_i}(x,t') = y'$. Then $\Psi_{T_{i+1}}(q_i(x),[t,t_k]) = q_i(\Psi_{T_{i+1}}(x,[t',t'_k]))$.

So in either case, there is some $t' \in [t'_{k+1}, t'_k]$ such that

$$\Psi_{T_{i+1}}(q_i(x),[t,t_k]) = q_i(\Psi_{T_{i+1}}(x,[t',t'_k])).$$

Therefore,

$$\Psi_{T_{i+1}}(q_i(x),[t,1]) = \Psi_{T_{i+1}}(q_i(x),[t,t_k]) \cup \Psi_{T_{i+1}}(q_i(x),[t_k,1])$$
$$= q_i(\Psi_{T_i}(x,[t',t'_k])) \cup q_i(\Psi_{T_i}(x,[t'_k,1]))$$
$$= q_i(\Psi_{T_i}(x,[t',1])).$$

So $\Psi_{T_{i+1}}$ has the desired property.

Our aim will now be to show that, if $\Psi_{T_{i+1}}$ is not graph f-tame, then $N_{i+1} < N_i$.

Recall that N_i is the greatest distance from the basepoint of any $x \in \Delta_i^{(1)}$ such that x demonstrates that Ψ_{T_i} is not f-tame, i.e., such that there is a point y prior to x in Ψ_{T_i} with $d_{\Delta_i}(*, y) > f(d_{\Delta_i}(*, x))$. If the sequence N is strictly decreasing, then at some point there will be no such x, and we will be left with a graph f-tame 1-combing.

Recall that each e_j is a 1-cell in Δ_i such that points in the interior of e_j are a distance N_i away from $*$ and there is a point prior to e_j—and therefore, in the icicle at e_j—of distance greater than $f(N_i)$ away from $*$. I claim that these obstructions to f-tameness have been removed in Δ_{i+1}: that for each $j \in [m_i]$, every $y \in \Delta_{i+1}^{(1)}$ in the T_{i+1}-icicle at $q_i(e_j)$ has $d_{\Delta_{i+1}}(*, y) \leq f(N_i)$.

For $j \in [m_i]$, we know that

$$\text{IDiam}(D'_j) \leq \max\left(\text{IDiam}(\ell(w_j)), \left\lfloor \frac{\ell(w_j)}{2} \right\rfloor + M_{\mathcal{P}}\right).$$

Now recall that α_j is the tail of the icicle at e_j and that $\gamma_j = \beta_{x_j} \cdot \vec{e_j} \cdot \overline{\beta_{y_j}}$ bounds the body of that icicle. Also recall that $N_i \in \mathbb{N}\left[\frac{1}{2}\right]$. Since $d_{\Delta_{i+1}}(*, q_i(e_j)) = N_i$ and $q_i(\alpha_j \cdot \beta_{x_j})$ and $q_i(\alpha_j \cdot \beta_{y_j})$ are Δ_{i+1}-geodesics, $|\alpha_j \cdot \gamma_j \cdot \overline{\alpha_j}| = |\alpha_j \cdot \beta_{x_j} \cdot \vec{e_j} \cdot \overline{\beta_{y_j}} \cdot \overline{\alpha_j}| \leq 2N_i + 1$. So $\ell(w_j) = |\gamma_j| \leq 2N_i + 1 - 2|\alpha_j|$. Therefore,

$$\begin{aligned}
\max&\left(\text{IDiam}(\ell(w_j)), \left\lfloor \frac{\ell(w_j)}{2} \right\rfloor + M_{\mathcal{P}}\right) \\
&\leq \text{IDiam}(\ell(w_j)) + \left\lfloor \frac{\ell(w_j)}{2} \right\rfloor + M_{\mathcal{P}} \\
&\leq \text{IDiam}(2N_i + 1 - 2|\alpha_j|) + \left\lfloor \frac{2N_i + 1 - 2|\alpha_j|}{2} \right\rfloor + M_{\mathcal{P}} \\
&\leq \text{IDiam}(2N_i + 1) + N_i + \frac{1}{2} - |\alpha_j| + M_{\mathcal{P}}.
\end{aligned}$$

Now let y' be a vertex in the T_{i+1}-icicle at $q_i(e_j)$ in Δ_{i+1}. If y' is on the simple path from $*$ to $q_i(*_j)$ in T_{i+1}, then $d_{\Delta_{i+1}}(*, y') \leq N_i$. Otherwise, since $q_i(*_j)$ is on the simple path from $*$ to y' in T_{i+1},

$$\begin{aligned} d_{\Delta_{i+1}}(*, y') &= d_{\Delta_{i+1}}(*, q_i(*_j)) + d_{\Delta_{i+1}}(q_i(*_j), y') \\ &\leq |\alpha_j| + \mathrm{IDiam}(D'_j) \\ &\leq |\alpha_j| + \mathrm{IDiam}(2N_i + 1) + N_i + \frac{1}{2} - |\alpha_j| + M_{\mathcal{P}} \\ &= \mathrm{IDiam}(2N_i + 1) + N_i + \frac{1}{2} + M_{\mathcal{P}} = f(N_i) - \frac{1}{2}. \end{aligned}$$

Every point on a 1-cell in the icicle has coarse distance differing by $\frac{1}{2}$ from some vertex in the icicle, so for all $y \in \Delta_{i+1}^{(1)}$ in the T_{i+1}-icicle at $q_i(e_j)$,

$$d_{\Delta_{i+1}}(*, y) \leq \max\left(N_i + \frac{1}{2}, f(N_i)\right) = f(N_i) \tag{2.4.1}$$

Now if $\Psi_{T_{i+1}}$ is not graph f-tame, then there exists some $x \in \Delta_{i+1}^{(1)}$ such that $M_{i+1}(x) > f(d_{\Delta_{i+1}}(*, x))$. Then let $p \in \partial\Delta_{i+1}$ and $s, t \in [0, 1]$ such that $s < t$, $\Psi_{T_{i+1}}(p, t) = x$, $y := \Psi_{T_{i+1}}(p, s) \in \Delta_{i+1}^{(1)}$, and $d_{\Delta_{i+1}}(*, y) = M_{i+1}(x)$. I claim that $d_{\Delta_{i+1}}(*, x) < N_i$, which will be sufficient to show that $N_{i+1} < N_i$. We will prove this by considering two cases.

Case 1: y is in the T_{i+1}-icicle at $q_i(e_j)$ for some $j \in [m_i]$. Then by inequality 2.4.1, $d_{\Delta_{i+1}}(*, y) \leq f(N_i)$. Since we chose x and y such that $f(d_{\Delta_{i+1}}(*, x)) < M_{i+1}(x) = d_{\Delta_{i+1}}(*, y)$, this shows that $f(d_{\Delta_{i+1}}(*, x)) < f(N_i)$. Since f is a strictly increasing function, this implies that $d_{\Delta_{i+1}}(*, x) < N_i$.

Case 2: y is not in a T_{i+1}-icicle at $q_i(e_j)$ for any $j \in [m_i]$. Then since $\Psi_{T_{i+1}}$ respects T_{i+1}, we know that for all $r \in [s, 1]$, $\Psi_{T_{i+1}}(p, r)$ is also not in the T_{i+1}-icicle at $q_i(e_j)$ for any $j \in [m_i]$. So $\Psi_{T_{i+1}}(p, [s, 1]) \subseteq q_i(\widehat{\Delta_i})$. Based on the way that we chose $\Psi_{T_{i+1}}$ such that its combing paths agree with those of $q_i \circ \Psi_{T_i}$ on $q_i(\widehat{\Delta_i})$, we know that there exist $p' \in \partial\Delta_i$ and $s' \in [0, 1]$ such that $q_i(\Psi_{T_i}(p', [s', 1])) = \Psi_{T_{i+1}}(p, [s, 1])$. In particular,

$q_i(p') = p$, $q_i(\Psi_{T_i}(p', s')) = y$, and there exists $t' \in [s', 1]$ such that $q_i(\Psi_{T_i}(p', t')) = x$. For ease of notation, let $y' = \Psi_{T_i}(p', s')$ and let $x' = \Psi_{T_i}(p', t')$. Now we know that $d_{\Delta_i}(*, x') = d_{\Delta_{i+1}}(*, x)$ and $d_{\Delta_i}(*, y') = d_{\Delta_{i+1}}(*, y)$. Since $y' \in P(\Psi_{T_i}, x')$, we have that

$$M_i(x') \geq d_{\Delta_i}(*, y') = d_{\Delta_{i+1}}(*, y) = M_{i+1}(x) > f(d_{\Delta_{i+1}}(*, x)) = f(d_{\Delta_i}(*, x')).$$

This implies that $d_{\Delta_i}(*, x') \leq N_i$, by the definition of N_i. But if $d_{\Delta_i}(*, x') = N_i$, then x' would be on some 1-cell in E_i, and therefore in the T_i-icicle at e_j for some $j \in [m_i]$. Since $T_{i+1} \cap q_i(\widehat{\Delta_i}) = q_i(\widehat{T_i})$, this would imply that x is in the T_{i+1}-icicle at $q_i(e_j)$, which is a contradiction. Therefore, $d_{\Delta_{i+1}}(*, x) = d_{\Delta_i}(*, x') < N_i$. This concludes case 2 and proves that $d_{\Delta_{i+1}}(*, x) < N_i$.

Since N_{i+1} is the maximum of $d_{\Delta_{i+1}}(*, x)$ for all such x, this implies that $N_{i+1} < N_i$. But for all i such that N_i is defined, $N_i > 0$. Therefore, this sequence $(\Delta_i, T_i, \Psi_{T_i})$ must end. Based on the way the sequence was constructed, this implies that for some $n \in \mathbb{N}$, Ψ_{T_n} is graph f-tame. Since w was an arbitrary word with $w =_G 1$ and there exists a van Kampen diagram Δ_n for w with a graph f-tame 1-combing, f is an intrinsic graph tame filling function for $\langle A|R \rangle$. Then by Lemma 2.4.3, f is equivalent to an intrinsic tame filling function for $\mathcal{P}$. Since this equivalence of functions is transitive, there is an intrinsic tame filling function for $\mathcal{P}$ that is equivalent to IDiam.

□

Remark. Unfortunately, a similar proof strategy for the extrinsic version of this theorem would fall apart. It is necessary to choose the spanning trees T_i to be trees of Δ_i-geodesics out of $*$. The fact that the paths used are Δ_i-geodesics is crucial to attain inequality 2.4.1. Reformulating this inequality for the extrinsic case would require paths that are geodesics in the Cayley graph, since extrinsic diameter is measured using distance in the Cayley graph. Since it is known that the intrinsic diameter functions of some groups grow strictly faster then their extrinsic diameter functions, it is not possible in general to find van Kampen diagrams with spanning trees of Cayley graph geodesics. In fact, in the extrinsic case, it is not even clear if it is possible in general to find a filling with an f-tame path to each vertex of each diagram, no matter how fast the function f grows; in the intrinsic case, optimally tame paths to each vertex are simply handed to us in the form of geodesics in the diagram.

Chapter 3

Aggregate Variation Functions

Given that intrinsic tame filling functions do not provide a proper refinement of the invariant given by intrinsic diameter functions, we are left with the question of whether or not there is any way to refine this invariant by measuring something along the same lines as intrinsic tame filling functions. This statement is intentionally vague, given that an intuitive sense of “what intrinsic tame filling functions measure” is rather subjective. In this chapter, we focus on the notion of a van Kampen diagram as a terrain with elevation represented by distance from the basepoint. We will view intrinsic tame filling functions as measuring the “hilliness” of the terrain, and consider ways of getting more detailed information about the geography of the terrain to see if we can distinguish between more groups.

3.1 Motivation and Definitions

Although a 1-combing of a van Kampen diagram does have a combing path that travels up and down every hill in the diagram, tame filling functions only tell us how “bad” individual hills can be. One way to try to get more information about the hilliness of a van Kampen diagram is to consider all the hills together and sum the change in elevation that would be required to travel up all of them. We will call this sum the aggregate variation of a van Kampen diagram (once it has been formally defined).

In order to formalize this notion into an actual definition, we need some way of identifying hills within a van Kampen diagram. To do so, we will first define a way to draw "contour lines" on a van Kampen diagram to turn it into a topographic map. Since these "contour lines" can split into multiple branches, we will refer to them as contour graphs.

Let Δ be a van Kampen diagram and let $h : \Delta^{(0)} \to \mathbb{Z}$ be a function such that, if $x, y \in \Delta^{(0)}$ are adjacent, then $|h(x) - h(y)| \leq 1$. (For example: given a basepoint $*$, we could define $h(x)$ to be the length of the shortest edge path from $*$ to x.) We will call such a function h a *height function* on Δ. Extend h to a function $\widetilde{h} : \Delta^{(1)} \to \mathbb{Z}[\frac{1}{2}]$ in the following way. For $x \in \Delta^{(1)} \setminus \Delta^{(0)}$, let u and v be the endpoints of the 1-cell containing x. Then define $\widetilde{h}(x) = \min(h(u), h(v)) + \frac{1}{2}$.

For each 2-cell σ of Δ, let $\phi_\sigma : D^2 \to \Delta$ be the characteristic map of σ—the map that restricts to the attaching map of σ on ∂D^2. From the definition of van Kampen diagram, we are assuming that ϕ_σ is a homeomorphism. Given $x, y \in \phi_\sigma^{-1}(\partial\sigma^{(0)})$ with $\widetilde{h}(\phi_\sigma(x)) = \widetilde{h}(\phi_\sigma(y)) =: n$, we will say that x and y are *contour partners* if there is a path γ from x to y in ∂D^2 such that for all $z \neq x, y$ on γ, we have $\widetilde{h}(\phi_\sigma(z)) < n$. If x and y are contour partners, then let s be the line segment in D^2 connecting x and y. Then we will call s a *contour segment* between x and y, and we will call $\phi_\sigma(s)$ a *contour path* between $\phi_\sigma(x)$ and $\phi_\sigma(y)$ in σ. Note that x and y are not necessarily distinct, and as a result every 0-cell of Δ that is contained in some 2-cell is itself a contour path. In fact, we will consider every 0-cell to be a contour path, even if it is not contained in any 2-cell.

Intuitively, two vertices of a 2-cell will be connected by a contour path if they share the same height and one side of the 2-cell is "downhill" from both of them. For example, Figures 3.1a and 3.1b both show vertices p and q of a 2-cell that have the same height of 10. In Figure 3.1a, every vertex between p and q on the left side of the 2-cell is below their shared height. We interpret this to mean that this left side of the 2-cell is downhill from some path connecting p and q. The contour path connecting p and q represents the path along this hill that runs perpendicular to the hill's slope and stays at a constant height of 10. In Figure 3.1b, however, neither side of the 2-cell stays below a height of 10; p and q are cut off from each other on each side by a vertex of height 12. We interpret this to mean that there is a ridge of height 12 that runs through the 2-cell, separating p and q. Therefore, we would not expect p and q to be connected by a contour line that goes through this 2-cell, so we do not draw a contour path between them.

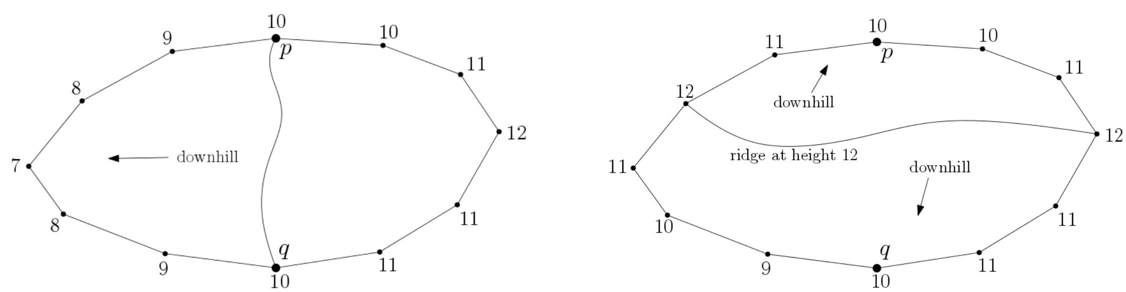

(a) p and q have the same height and are connected by a contour path.

(b) p and q have the same height but are separated by a ridge.

Figure 3.1: Contour path interpretations.

Let P be the union of all the contour paths in Δ. Each path component of P is called a *contour graph* of the pair (Δ, h). Let $\widehat{\Delta} = \Delta \setminus P$. Each path component of $\widehat{\Delta}$ is called a *contour component* of the pair (Δ, h). Note that since each 0-cell of Δ is itself a contour path, every 0-cell is part of some contour graph, even if that contour graph is composed only of that one 0-cell.

Finally, extend $\widetilde{h}$ to a function $\widehat{h} : \Delta \to \mathbb{N}[\frac{1}{2}]$ in the following way. Let $z' \in \text{int}(\sigma)$ for some 2-cell σ. Since ϕ_σ is a homeomorphism onto σ, let $z = \phi_\sigma^{-1}(z')$. If z is on a contour segment connecting contour partners x and y, define $\widehat{h}(z') = h(\phi_\sigma(x))$. Otherwise, let S_σ be the union of all contour segments between contour partners in $\phi_\sigma^{-1}(\sigma)$, let $\widehat{D^2_\sigma} = D^2 \setminus S_\sigma$, and let C_z be the path component of $\widehat{D^2_\sigma}$ containing z. If there is some $x \in C_z \cap \partial D^2$, then define $\widehat{h}(z') = \widetilde{h}(\phi_\sigma(x))$. Otherwise, ∂C_z is made up entirely of contour segments. So let x and y be any pair of contour partners such that the contour segment joining them is in ∂C_z, and define $\widehat{h}(z') = h(\phi_\sigma(x)) + \frac{1}{2}$.

Proposition 3.1.1. *$\widehat{h}$ is well-defined.*

Proof. We will first show that contour segments can only intersect each other at their endpoints. Suppose by way of contradiction that two contour segments of σ, s_1 and s_2, do cross somewhere between their endpoints. Then let $x_1, y_1, x_2, y_2 \in \partial D^2$ with x_1 and y_1 the endpoints of s_1 and x_2 and y_2 the endpoints of s_2. The fact that s_1 and s_2 cross implies that x_1 and y_1 are on opposite sides of s_2 and that x_2 and y_2 are on opposite sides of s_1.

Now for $i \in \{1, 2\}$, since s_i is a contour segment, we must have that $h(\phi_\sigma(x_i)) = h(\phi_\sigma(y_i)) =: n_i$, and there must exist a path γ_i in ∂D^2 from x_i to y_i such that, for all $z \neq x_i, y_i$ on γ_i, we have $\widetilde{h}(\phi_\sigma(z)) < n_i$. Suppose without loss of generality that $n_1 \leq n_2$. Now since x_1 and y_1 are on opposite sides of s_2, we must have that one of x_2 and y_2 is on γ_1. Since $h(\phi_\sigma(x_2)) = h(\phi_\sigma(y_2)) = n_2 \geq n_1$, this contradicts the property of γ_1 stated above. So contour segments only meet at their endpoints.

This shows that $\widehat{h}$ is well-defined on contour graphs, since no $z \in \text{int}(D^2)$ can be on two different contour segments of σ.

This also implies that $\widehat{h}(z')$ is well-defined in the case where there is no path in $\widehat{D^2_\sigma}$ from z to ∂D^2. For if ∂C_z is a union of contour segments, and contour segments only meet at their endpoints, then ∂C_z is a polygon such that each vertex maps to a 0-cell of σ. Since contour segments only connect endpoints with the same height in

σ, every vertex of ∂C_z maps to a 0-cell with the same height.

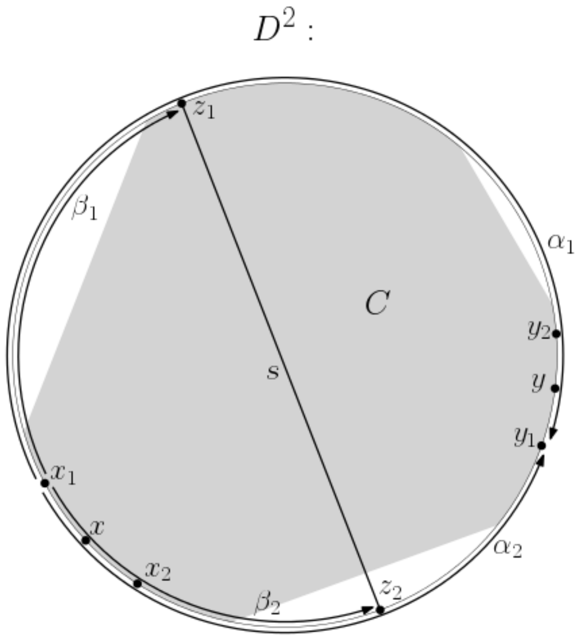

Figure 3.2: An example of the contradiction if $\widetilde{h}(\phi_\sigma(x)) \neq \widetilde{h}(\phi_\sigma(y))$. The path component C_z of $\widehat{D^2_\sigma}$ is shaded.

Finally, suppose $z \in \widehat{D^2_\sigma}$ and the path component C_z of $\widehat{D^2_\sigma}$ containing z intersects ∂D^2. See Figure 3.2 for an example diagram of this case. Suppose by way of contradiction that there exist $x, y \in C_z \cap \partial D^2$ with $\widetilde{h}(\phi_\sigma(x)) \neq \widetilde{h}(\phi_\sigma(y))$. Assume without loss of generality that $\widetilde{h}(\phi_\sigma(x)) < \widetilde{h}(\phi_\sigma(y))$. Note that $\phi_\sigma(x), \phi_\sigma(y) \in \Delta^{(1)} \cap \phi_\sigma(\widehat{D^2_\sigma})$, and $\phi_\sigma(\widehat{D^2_\sigma})$ does not contain any 0-cells because every 0-cell is contained in a contour path. So $\phi_\sigma(x)$ and $\phi_\sigma(y)$ are each contained in the interior of a 1-cell. Therefore, there is some $n \in \mathbb{N}$ with $\widetilde{h}(\phi_\sigma(x)) < n < \widetilde{h}(\phi_\sigma(y))$. Let x_1' and x_2' be the endpoints of the 1-cell containing $\phi_\sigma(x)$ such that $h(x_1') \leq h(x_2')$ and let y_1' and y_2' be the endpoints of the 1-cell containing $\phi_\sigma(y)$ such that $h(y_1') \leq h(y_2')$. Note that this implies that $h(x_1') = \widetilde{h}(\phi_\sigma(x)) - \frac{1}{2} < n$, and that $h(y_1') = \widetilde{h}(\phi_\sigma(y)) - \frac{1}{2} \geq n$. Since ϕ_σ is a homeomorphism, let $x_1 = \phi_\sigma^{-1}(x_1')$, $x_2 = \phi_\sigma^{-1}(x_2')$, $y_1 = \phi_\sigma^{-1}(y_1')$, and $y_2 = \phi_\sigma^{-1}(y_2')$. Let α_1 be the path in ∂D^2 from x_1 to y_1 not containing x_2 and let α_2 be the path in

∂D^2 from x_1 to y_1 containing x_2. Let z_1 be the first point on α_1 with $h(\phi_\sigma(z_1)) = n$ and let z_2 be the first point on α_2 with $h(\phi_\sigma(z_2)) = n$; these points must exist because the value of h can only differ by at most 1 between adjacent 0-cells and because $h(x_1') < n \leq h(y_1')$. Let β_1 be the initial segment of α_1 from x_1 to z_1 and let β_2 be the initial segment of α_2 from x_1 to z_2. Then $\gamma = \overline{\beta_1} \cdot \beta_2$ is a path in ∂D^2 from z_1 to z_2 such that for all $z_3 \neq z_1, z_2$ on γ, we have $\widetilde{h}(\phi_\sigma(z_3)) < n$. Therefore, z_1 and z_2 are contour partners and there is a contour segment s that connects them. Now note that y cannot be on β_1 or β_2 because $n \leq h(y_1'), h(y_2')$, and x must be on β_2 because $h(x_1') < n$. So if s is the line segment connecting z_1 and z_2, then x and y are on opposite sides of s. This contradicts the fact that x and y are in the same path component of $\widehat{D^2_\sigma}$. Therefore, for all $x, y \in C_z \cap \partial D^2$, we have $\widetilde{h}(\phi_\sigma(x)) = \widetilde{h}(\phi_\sigma(y))$.

□

Remark 3.1.2. Note that the fact that the contour segments of a given 2-cell only meet at their endpoints implies that contour paths can only meet at their endpoints. Therefore, every point in a given contour graph has the same height. Similarly, contour components are made up of subsets of 2-cells that can only meet each other at the interior of 1-cells, implying that every point in a given contour component has the same height.

It is convenient at this point to prove the following lemma, which will be useful later on. It confirms that contour graphs work similarly to the way we expect contour lines on a topographic map to behave, in that they always separate points of different elevations.

Lemma 3.1.3. *Let Δ be a van Kampen diagram. If $x, y \in \Delta$ and $n \in \mathbb{N}_0$ with $\widehat{h}(x) \leq n \leq \widehat{h}(y)$, then every path from x to y contains a point of height n.*

Proof. Let $x, y \in \Delta$ and $n \in \mathbb{Z}$ with $\widehat{h}(x) \leq n \leq \widehat{h}(y)$. Let $\gamma : [0,1] \to \Delta$ be a path from x to y. Note that every contour path of Δ is either a single point (which is closed in Δ), or a continous image in Δ of a line segement from D^2. Since line segments in D^2 are closed and bounded, and therefore compact, so are the continuous images of a line segment. In particular, every contour path is closed in Δ, and since van Kampen diagrams have finitely many 0-cells, they have finitely many contour paths. So every contour graph is closed in Δ, as is P, since each of these is a finite union of closed sets.

Let $A = \{t \in [0,1] | \gamma(t) \in P \cup \Delta^{(1)} \text{ and } \widehat{h}(\gamma(t)) \leq n\}$. Note that A is a compact subset of $[0,1]$, since every contour graph has constant, integer height and is closed in Δ, as is $\Delta^{(1)}$. So let $a = \max(\{0\} \cup A)$. Similarly, $B = \{t \in [a,1] | \gamma(t) \in P \cup \Delta^{(1)} \text{ and } \widehat{h}(\gamma(t)) \geq n\}$ is compact in $[0,1]$, so let $b = \min(\{1\} \cup B)$. Note that $a \leq b$, and that $a = b$ if and only if $\widehat{h}(\gamma(a)) = n$.

Suppose by way of contradiction that $a < b$. Then $\widehat{h}(\gamma(a)) < n < \widehat{h}(\gamma(b))$, and by the definitions of a and b we have $\gamma((a,b)) \cap (P \cup \Delta^{(1)}) = \emptyset$. So there is some path component D of $\Delta \setminus (P \cup \Delta^{(1)})$ such that $\gamma((a,b)) \subseteq D$ and $\gamma(a)$ and $\gamma(b)$ are in ∂D. Note that for some 2-cell σ and some contour component C, we have that $D \subseteq \text{int}(\sigma) \cap C$.

Since contour paths only intersect each other and 1-cells at 0-cells, we can view D as a face of a plane graph whose vertices are 0-cells of Δ and whose edges are 1-cells and contour paths of Δ. As such, ∂D is path connected. Now consider $\partial D \setminus C$. This removes the interior of every 1-cell in ∂D, leaving only contour paths. So each path component of $\partial D \setminus C$ is a subset of a contour graph, and therefore has constant height by Remark 3.1.2. Since ∂D is path connected and there are no 0-cells in C (since all 0-cells of Δ are in some contour graph), if $\partial D \setminus C$ has multiple path components, then each one is connected to another by some 1-cell whose interior is contained in

C.

Now we know that every point in C has the same height by Remark 3.1.2, and in particular this implies that the interior of every 1-cell in C has a height of $m + \frac{1}{2}$ for some $m \in \mathbb{Z}$. So if the interior of some 1-cell is in C, then the 0-cells at its endpoints have heights in $\{m, m+1\}$. We've shown above that every path component of $\partial D \setminus C$ contains such a 0-cell, implying that every path component of $\partial D \setminus C$ has height in $\{m, m+1\}$. Therefore, $\widehat{h}(\partial D) \subseteq \{m, m+\frac{1}{2}, m+1\}$. Since $\gamma(a), \gamma(b) \in \partial D$, this implies that $\widehat{h}(\gamma(b)) - \widehat{h}(\gamma(a)) \leq 1$. Given that $\widehat{h}(\gamma(a)) < n < \widehat{h}(\gamma(b))$, we must have $\widehat{h}(\gamma(a)) = n - \frac{1}{2}$ and $\widehat{h}(\gamma(b)) = n + \frac{1}{2}$. But this is not possible since m and n are both integers.

This is a contradiction, implying that $a = b$, and therefore that $\widehat{h}(\gamma(a)) = n$. So γ contains a point of height n.

□

Now that we can draw a topographic map on a van Kampen diagram, we are ready to return to our initial motivation. We wanted to add together the change in height that would be required to travel up each hill in the diagram. If we cut the diagram along a contour graph, we will view each path component (aside from the component containing the basepoint) as a hill. A tall hill, then, has many hills nested within it. To define the aggregate variation, we will simply count the total number of hills.

This definition will mostly match our initial motivation. For a path to reach the highest point of a hill, it must cross into a number of nested hills at least equal to the difference in elevation between the base of the hill and the highest point. Therefore, the height of each hill will be accounted for in the aggregate variation.

We will also need to account for one more detail in the definition of aggregate variation. It is possible to produce many short hills in a van Kampen diagram, depending on what relators are used in the presentation. If we leave things as described above, these short hills will all be counted by the aggregate variation. These would amount to noise that can drown out how truly flat or hilly the van Kampen diagrams of a group are, and prevent the desired filling function from being a group invariant. Therefore, we will add to the definition a notion of *sensitivity* that will allow us to ignore sufficiently short hills.

Here, then, is the final definition:

Definition 3.1.4. Let Δ be a van Kampen diagram with a basepoint $*$, and let h be a height function on Δ. Given $n \in \mathbb{Z}$, let P_n be the union of the contour graphs of (Δ, h) that have height n. Given $s \in \mathbb{N}[\frac{1}{2}]$ and $n \in \mathbb{Z}$, let C_n^s be the set of path components of $\Delta \setminus P_n$ that do not contain $*$ and do contain a point x with $|\widehat{h}(x) - n| \geq s$. Then define the *aggregate variation of* (Δ, h) *with sensitivity* s by

$$\mathrm{AV}_s(\Delta, h) = \sum_{n \in \mathbb{Z}} |C_n^s|.$$

Note that this sum is always finite, since Δ is finite.

We will now consider the specific case where the height function is given by distance to the basepoint. Although our motivation has so far been based on intrinsic distance (that is, distance measured within the van Kampen diagram), extrinsic distance (which is measured in the Cayley complex) also produces a height function, so we can define intrinsic and extrinsic versions of the new filling function.

Recall that $\mathcal{P} = \langle A | R \rangle$ is a finite presentation of a group G and X is the Cayley 2-complex of $\mathcal{P}$. Let ε be the identity vertex of X.

Let Δ be a van Kampen diagram with respect to $\mathcal{P}$ with basepoint $*$. Let π_Δ :

$\Delta \to X$ be the unique cellular map that preserves the directions and labels of the directed edges of Δ and sends $*$ to ϵ.

Define intrinsic and extrinsic distance functions $I, E : \Delta^{(0)} \to \mathbb{N}$ by $I(x) = d_\Delta(*, x)$ and $E(x) = d_X(\varepsilon, \pi_\Delta(x))$ (See Section 1.2 for the definition of this notation). Note that I and E are both height functions on Δ.

For $s \in \mathbb{N}[\frac{1}{2}]$, let $\mathrm{IAV}_s(\Delta) = \mathrm{AV}_s(\Delta, I)$, and let $\mathrm{EAV}_s(\Delta) = \mathrm{AV}_s(\Delta, E)$. These are diagram measurements, so we may define filling functions from them in the following way.

Let $w \in (A \cup A^{-1})^*$ with $w =_G 1$. For $s \in \mathbb{N}[\frac{1}{2}]$, define

$$\mathrm{IAV}_s(w) = \min\{\mathrm{IAV}_s(\Delta) | \Delta \text{ is a van Kampen diagram for } w\}$$

and define

$$\mathrm{EAV}_s(w) = \min\{\mathrm{EAV}_s(\Delta) | \Delta \text{ is a van Kampen diagram for } w\}.$$

For $s \in \mathbb{N}[\frac{1}{2}]$ and $n \in \mathbb{N}$, define

$$\mathrm{IAV}_s(n) = \max\{\mathrm{IAV}_s(w) | w \in (A \cup A^{-1})^* \text{ with } \ell(w) \leq n \text{ and } w =_G 1\},$$

and define

$$\mathrm{EAV}_s(n) = \max\{\mathrm{EAV}_s(w) | w \in (A \cup A^{-1})^* \text{ with } \ell(w) \leq n \text{ and } w =_G 1\}.$$

Definition 3.1.5. A non-decreasing function $f : \mathbb{N} \to \mathbb{N}$ is an *intrinsic aggregate variation function* (IAV function) for $\mathcal{P}$ if there is some sensitivity $s \in \mathbb{N}[\frac{1}{2}]$ such that $f(n) \geq \mathrm{IAV}_s(n)$ for all $n \in \mathbb{N}$. Similarly, a non-decreasing function $f : \mathbb{N} \to \mathbb{N}$ is an *extrinsic aggregate variation function* (EAV function) for $\mathcal{P}$ if there is some sensitivity $s \in \mathbb{N}[\frac{1}{2}]$ such that $f(n) \geq \mathrm{EAV}_s(n)$ for all $n \in \mathbb{N}$.

As an example to illustrate what IAV functions measure, consider the group $\mathrm{BS}(1,2)$ with standard presentation $\langle a, t | tat^{-1} = a^2 \rangle$. A van Kampen diagram D_5 for the word $w_5 = at^5at^{-5}a^{-1}t^5a^{-1}t^{-5}$ with respect to this presentation is shown in Figure 3.3, with contour paths with respect to the height function I drawn in color. We will compute $\mathrm{IAV}_1(D_5)$. Each component that contributes to $\mathrm{IAV}_1(D_5)$ is indicated with a number labeling the contour graph that separates the corresponding component from the basepoint $*$. There are also four components shaded in gray that would otherwise contribute to the aggregate variation, but are excluded by the sensitivity condition with a sensitivity of 1: all points in each of these components have height within one unit of the height of the contour graph that separates the given component from the basepoint. Since there are 11 components that contribute to $\mathrm{IAV}_1(D_5)$, we have that $\mathrm{IAV}_1(D_5) = 11$.

3.2 Comparisons to other filling functions

Having defined IAV and EAV functions, we will now consider how they compare to the Dehn function and diameter functions.

Given a van Kampen diagram Δ, let $\mathrm{Area}(\Delta)$ be the number of 2-cells of Δ. Given a finite presentation $\mathcal{P} = \langle A | R \rangle$ and a word $w \in (A \cup A^{-1})^*$ with $w =_G 1$, define

$$\mathrm{Area}_{\mathcal{P}}(w) = \min\{\mathrm{Area}(\Delta) | \Delta \text{ is a van Kampen diagram for } w\}.$$

Then the *Dehn function* (also known as the area function) of $\mathcal{P}$ is the function $\mathrm{Area}_{\mathcal{P}}$:

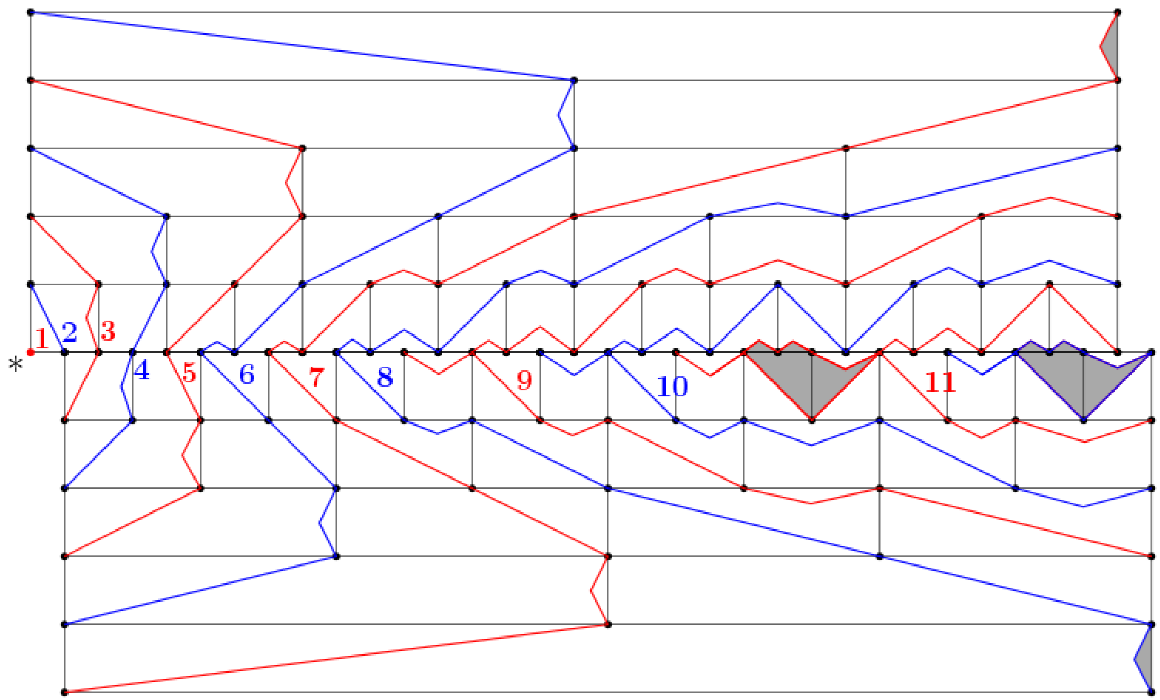

Figure 3.3: A van Kampen diagram D_5 for the word $w_5 = t^5at^{-5}at^5a^{-1}t^{-5}a^{-1}$. The 0-cells and 1-cells are drawn in black, with contour paths drawn thicker in blue and red. The edges labeled by an a are drawn horizontally directed to the right, and the edges labeled by a t are drawn vertically, directed upwards in the top half of the diagram and downwards in the bottom half.

$\mathbb{N} \to \mathbb{N}$ given by

$$\text{Area}_{\mathcal{P}}(n) = \max\{\text{Area}(w) | w \in (A \cup A^{-1})^* \text{ with } \ell(w) \leq n \text{ and } w =_G 1\}$$

.

Proposition 3.2.1. *Given a finite presentation $\mathcal{P} = \langle A|R\rangle$, there exists an IAV function f and an EAV function g for $\mathcal{P}$ such that $f, g \preceq \text{Area}_{\mathcal{P}}$.*

Proof. We will show that $\text{IAV}_1 \preceq \text{Area}_{\mathcal{P}}$. The proof that $\text{EAV}_1 \preceq \text{Area}_{\mathcal{P}}$ is identical.

Let $\rho = \max\{\ell(r)|r \in R\}$. Let Δ be a van Kampen diagram for a word w with basepoint $*$. Each 1-cell of Δ appears on the boundary of a 2-cell or on the boundary of the diagram (or both), so there are at most $\rho\,\mathrm{Area}(\Delta) + \ell(w)$ 1-cells in Δ. Each contour component of (Δ, I) either contains an open 1-cell or is entirely contained in the interior of a single 2-cell. Since each 1-cell is in a single contour component, the number of contour components that contain a 1-cell is at most $\rho\,\mathrm{Area}(\Delta) + \ell(w)$.

Recall from Definition 3.1.4 that, for each $n \in \mathbb{N}$, P_n is the union of the contour graphs of (Δ, I) of height n, and C_n^1 is the set of components of $\Delta \setminus P_n$ that do not contain $*$ and do contain a point x with $|\widehat{I}(x) - n| \geq 1$. Then

$$\mathrm{IAV}_1(\Delta) = \sum_{n=0}^{\mathrm{IDiam}(\Delta)} |C_n^1|.$$

Now for each $n \in \{0, 1, \ldots, \mathrm{IDiam}(\Delta)\}$, each component $C \in C_n^1$ contains at least one contour component D such that ∂D intersects a contour graph of height n. If D were contained in the interior of a 2-cell, then we would have $C = D$ and every $x \in C$ would have $\widehat{I}(x) = n + \frac{1}{2}$. But this contradicts the definition of C_n^1, so we must have that D contains an open 1-cell. Therefore, $|C_n^1| \leq |\mathcal{D}_n|$, where $\mathcal{D}_n$ is the set of pairs (D, n) where D is a contour component of (Δ, I) such that ∂D intersects a contour graph in P_n and D contains an open 1-cell. Let $\mathcal{D} = \cup_{n=0}^{\mathrm{IDiam}(\Delta)} D_n$. Then from the previous inequality, we have that $\mathrm{IAV}_1(\Delta) \leq |\mathcal{D}|$. Now since each contour component of (Δ, I) is adjacent to contour graphs of at most two different heights, each contour component containing an open 1-cell contributes at most two elements to $\mathcal{D}$. Therefore, $|\mathcal{D}| \leq 2(\rho\,\mathrm{Area}(\Delta) + \ell(w))$, which implies that $\mathrm{IAV}_1(\Delta) \leq 2\rho\,\mathrm{Area}(\Delta) + 2\ell(w)$. Minimizing over all van Kampen diagrams of w gives that $\mathrm{IAV}_1(w) \leq 2\rho\,\mathrm{Area}(w) + 2\ell(w)$. Then given $n \in \mathbb{N}$, maximizing over all $w \in (A \cup A^{-1})^*$ with $w =_G 1$ and $\ell(w) \leq n$ gives that $\mathrm{IAV}_1(n) \leq 2\rho\,\mathrm{Area}(n) + 2n$. Hence, $\mathrm{IAV}_1 \preceq \mathrm{Area}$.

□

Proposition 3.2.2. *Given a finite presentation $\mathcal{P}$, let f be an IAV function for $\mathcal{P}$*

and let g be an EAV function for $\mathcal{P}$. Then $\mathrm{IDiam}_{\mathcal{P}} \preceq f$ *and* $\mathrm{EDiam}_{\mathcal{P}} \preceq g$.

Proof. We will show that $\mathrm{IDiam}_{\mathcal{P}} \preceq f$. The proof that $\mathrm{EDiam}_{\mathcal{P}} \preceq g$ is identical.

Let $s \geq 0$ such that $f \geq IAV_s$. Let Δ be a van Kampen diagram with basepoint $*$. For each $n \in \mathbb{N}$, let P_n be the union of the contour graphs of (Δ, I) of height n. Let $x \in \Delta^{(0)}$ with $I(x) = \mathrm{IDiam}(\Delta)$. Now $I(*) = 0$, so $\Delta \setminus P_0$ has a component C_0 containing x and not $*$. Furthermore, for $n \in \{1, \ldots, I(x) - 1\}$, we have $I(*) < n < I(x)$, so by Lemma 3.1.3, the component C_n of $\Delta \setminus P_n$ containing x does not contain $*$. Also, for $n \in \{0, \ldots, I(x) - \lceil s \rceil\}$, we have that C_n contains the point x with $|I(x) - n| \geq s$, so C_n counts towards the sum in the definition of $\mathrm{AV}_s(\Delta, I)$. Therefore, $\mathrm{IAV}_s(\Delta) \geq \mathrm{IDiam}(\Delta) - \lceil s \rceil$. As a result, $f(n) \geq \mathrm{IAV}_s(n) \geq \mathrm{IDiam}(n) - \lceil s \rceil$, so $f \preceq \mathrm{IDiam}$.

□

The following theorem demonstrates that IAV functions are equivalent to intrinsic diameter functions, and therefore fail to be a proper refinement.

Theorem 3.2.3. *If G is a group with finite presentation $\mathcal{P}$, then there is a presentation $\mathcal{P}'$ that is $\mathcal{P}$ with finitely many relators added such that* $\mathrm{IDiam}_{\mathcal{P}}$ *is an IAV function for $\mathcal{P}'$.*

The following definition and proposition will be useful for proving Theorem 3.2.3.

Definition 3.2.4. Let Δ be a van Kampen diagram with height function h. For $n \in \mathbb{N}[\frac{1}{2}]$, define $\Gamma_{\geq n}(\Delta)$ to be the following graph. Let σ be a closed cell of Δ that is not contained in any higher-dimensional cell. Then σ is a vertex of $\Gamma_{\geq n}(\Delta)$ if there is some $x \in \sigma$ with $\widehat{h}(x) \geq n$. Two vertices σ and τ are adjacent if and only if there is some $x \in \sigma \cap \tau$ with $\widehat{h}(x) \geq n$.

Proposition 3.2.5. *If Δ is a van Kampen diagram, then for every $n \in \mathbb{N}[\frac{1}{2}]$, there is a bijection between the components of $\Gamma_{\geq n}(\Delta)$ and the path components of $\Delta_{\geq n} = \{x \in \Delta | \widehat{h}(x) \geq n\}$.*

Proof. We will first show that for every $n \in \mathbb{N}[\frac{1}{2}]$ and every vertex σ of $\Gamma_{\geq n}(\Delta)$, the set $\sigma_{\geq n} = \{x \in \sigma | \widehat{h}(x) \geq n\}$ is path connected.

So let $n \in \mathbb{N}[\frac{1}{2}]$ and let σ be a vertex of $\Gamma_{\geq n}(\Delta)$.

If σ is a 0-cell, then $\sigma_{\geq n} = \sigma$, which is path connected. Suppose σ is a 1-cell. Let y and z be the 0-cells in σ. Then because $\widehat{h}$ is constant on $\text{int}(\sigma)$, we have that $\sigma_{\geq n}$ is the union of some nonempty subset of $\{\text{int}(\sigma), \{y\}, \{z\}\}$. Every such union other than $\{y, z\}$ is path connected. However, if $\widehat{h}(y), \widehat{h}(z) \geq n$, then $\widehat{h}(\text{int}(\sigma)) \geq n$, so we cannot have $\sigma_{\geq n} = \{y, z\}$. Therefore, $\sigma_{\geq n}$ is path connected.

Finally, suppose σ is a 2-cell. Let $\phi_\sigma : D^2 \to \Delta$ be σ's characteristic map—the homeomorphism that restricts to σ's attaching map on ∂D^2. Let $D^2_{\geq n} = \phi_\sigma^{-1}(\sigma_{\geq n})$.

Note first that, by the definition of $\widehat{h}$, every $x \in D^2$ has a path γ from x to ∂D^2 such that $\widehat{h}(\phi_\sigma(\gamma)) = \{\widehat{h}(\phi_\sigma(x))\}$, with one exception: if x is not in a contour segment and the path component C of $\widehat{D^2_\sigma}$ containing x does not intersect ∂D^2. In this case, C is the interior of a polygon bounded by contour segments of height $\widehat{h}(\phi_\sigma(x)) - \frac{1}{2}$. So in this case there is a path γ from x to ∂D^2 with $\widehat{h}(\phi_\sigma(\gamma)) = \{\widehat{h}(\phi_\sigma(x)) - \frac{1}{2}, \widehat{h}(\phi_\sigma(x))\}$. If such an x exists and $\widehat{h}(\phi_\sigma(x)) = n$, then $\partial D^2 \cap D^2_{\geq n} = \emptyset$, and therefore $C = D^2_{\geq n}$, making $\sigma_{\geq n} = \phi_\sigma(C)$ path connected. Otherwise, for every other $x \in D^2_{\geq n}$, we have that $\widehat{h}(\phi_\sigma(x)) \geq n$, meaning that every element of $\widehat{h}(\gamma)$ has height greater than or equal to n. Therefore, for each $x \in D^2_{\geq n}$, there is a path γ from x to $\partial D^2 \cap D^2_{\geq n}$ contained in $D^2_{\geq n}$. So we need only show that $\partial D^2 \cap D^2_{\geq n}$ is contained in a single path component of $D^2_{\geq n}$.

Let $X = \{x_0 = x_m, x_1, \ldots, x_{m-1}\} = \{x \in \partial D^2 | \widehat{h}(\phi_\sigma(x)) = \lfloor n \rfloor\}$, indexed such that the x_i's appear in order counter-clockwise around ∂D^2. Note that because $\lfloor n \rfloor \in \mathbb{N}$, we have that $\phi_\sigma(X)$ is a subset of the 0-cells of σ. Note that if $X = \emptyset$, then $\sigma_{\geq n} = \sigma$, since we must have that all 0-cells of σ have height at least $\lfloor n \rfloor + 1 > n$. As a 2-cell, σ is path connected. Otherwise, for $i \in [m]$, let α_i be the simple path (or loop, if $m = 1$) in ∂D^2 from x_{i-1} to x_i not containing any other elements of X. Let $I_<$ be the set of indices $i \in [m]$ such that $\max(\widehat{h}(\phi_\sigma(\alpha_i \setminus \{x_{i-1}, x_i\}))) < \lfloor n \rfloor$, and let $I_>$ be the set of indices $i \in [m]$ such that $\min(\widehat{h}(\phi_\sigma(\alpha_i \setminus \{x_{i-1}, x_i\}))) > \lfloor n \rfloor$. Note that $I_< \cup I_> = [m]$. See Figure 3.4 for an example of what the x_i's and α_i's can look like in D^2.

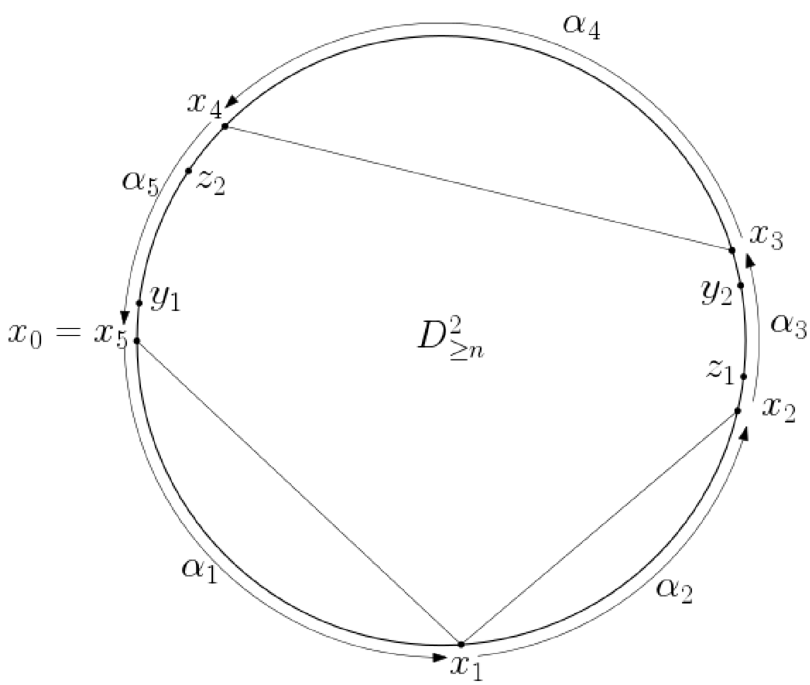

Figure 3.4: An example of what $D^2_{\geq n}$ can look like. The line segments shown are the contour segments of height n. Note that in this case, $I_< = \{1, 2, 4\}$ and $I_> = \{i_0 = i_2 = 3, i_1 = 5\}$.

Case 1: $n \in \mathbb{N}$, so that $\lfloor n \rfloor = n$. Let $i \in [m]$. If $i \in I_<$, then α_i demonstrates that x_{i-1} and x_i are contour partners, and the line segment s_i between them is a contour segment. So $s_i \subseteq D^2_{\geq n}$, putting x_{i-1} and x_i in the same path component of $D^2_{\geq n}$.

If instead $i \in I_>$, then $\alpha_i \subseteq D^2_{\geq n}$, putting every point of α_i (including x_{i-1} and x_i) in the same path component of $D^2_{\geq n}$. Then every element of X is in a single path component of $D^2_{\geq n}$, and since $\partial D^2 \cap D^2_{\geq n} = \bigcup_{i \in I_>} \alpha_i$, this implies that $\partial D^2 \cap D^2_{\geq n}$ is contained in a single path component of $D^2_{\geq n}$, as desired.

Case 2: $n \notin \mathbb{N}$, so that $\lfloor n \rfloor = n - \frac{1}{2}$. Let $I_> = \{i_0 = i_k, i_1, \ldots, i_{k-1}\}$, where $i_{j-1} < i_j$ for $j \in [k-1]$. Let $j \in [k]$. We want to show that $\alpha_{i_{j-1}}$ and α_{i_j} are contained in the same path component of $D^2_{\geq n}$. Let $y_j \in \alpha_{i_{j-1}}$ be a point such that $\phi_\sigma(y_j)$ is in the interior of a 1-cell containing $\phi_\sigma(x_{i_{j-1}})$. Since $\widehat{h}(\phi_\sigma(x_{i_{j-1}})) = \lfloor n \rfloor$, this implies that $\widehat{h}(\phi_\sigma(y_j)) = n$. Similarly choose $z_j \in \alpha_{i_j}$ to be a point such that $\phi_\sigma(z_j)$ is in the interior of a 1-cell containing $\phi_\sigma(x_{i_{j-1}-1})$, and note that $\widehat{h}(\phi_\sigma(z_j)) = n$ for the same reason as for y_j. Now consider the line segment s_j connecting y_j and z_j, directed from y_j to z_j. Note that every point p on the right side of s_j such that $\phi_\sigma(p)$ is a 0-cell is in some α_i where $i_{j-1} < i < i_j$. This implies that $i \in I_<$, meaning that $\widehat{h}(\phi_\sigma(p)) \leq \lfloor n \rfloor < n$.

Now suppose by way of contradiction that there is some point $q \in s_j$ such that q is on a contour segment s. This contour segment must have one endpoint p on the right side of s_j. Then $\widehat{h}(\phi_\sigma(p)) < n$, so s is a contour segment with height less than n. This implies that one of the two components of $\partial D^2 \setminus s$ contains only points of height less than n. But because s and s_j intersect in $\text{int}(D^2)$, we must have that y_j and z_j are on opposite sides of s. Therefore, one component of $\partial D^2 \setminus s$ contains y_j, and the other contains z_j. Since $\widehat{h}(\phi_\sigma(y_j)) = \widehat{h}(\phi_\sigma(y_j)) = n$, this is a contradiction.

Therefore, $s_j \subset \widehat{D^2_\sigma}$ (definition on p. 58), putting every point of s_j in the same path component of $\widehat{D^2_\sigma}$. By the definition of $\widehat{h}$, this implies that $\widehat{h}(\phi_\sigma(s_j)) = \{n\}$. Hence, $\alpha_{i_{j-1}}$ and α_{i_j} (minus their endpoints) are in the same path component of $\widehat{D^2_\sigma}$. So every $\alpha_{i_j} \setminus \{x_{i_j-1}, x_{i_j}\}$ for $j \in [k]$ is contained in a single path component of $\widehat{D^2_\sigma}$.

Since

$$\partial D^2 \cap D^2_{\geq n} = \bigcup_{i \in I_>} \alpha_i \setminus \{x_{i-1}, x_i\},$$

this implies that $\partial D^2 \cap D^2_{\geq n}$ is contained in a single path component of $D^2_{\geq n}$, as desired.

So we've shown that for every $n \in \mathbb{N}[\frac{1}{2}]$ and every vertex σ of $\Gamma_{\geq n}(\Delta)$, we have that $\sigma_{\geq n}$ is path connected.

Now let $\mathcal{C}_\Delta$ be the set of path components of $\Delta_{\geq n}$ and $\mathcal{C}_\Gamma$ the set of path components of $\Gamma_{\geq n}(\Delta)$. Define a function $f : \mathcal{C}_\Delta \to \mathcal{C}_\Gamma$ in the following way. Given a $C \in \mathcal{C}_\Delta$, let $x \in C$. Then there is some vertex σ of $\Gamma_{\geq n}(\Delta)$ containing x. So let $f(C)$ be the component of $\Gamma_{\geq n}(\Delta)$ containing σ. We have that f is onto because every component of $\Gamma_{\geq n}(\Delta)$ contains a vertex of $\Gamma_{\geq n}(\Delta)$, which contains a point in $\Delta_{\geq n}$. The fact that f is well-defined and 1-1 is equivalent to the following claim.

Let σ and τ be vertices of $\Gamma_{\geq n}(\Delta)$, and let $x \in \sigma$ and $y \in \tau$ with $\widehat{h}(x), \widehat{h}(y) \geq n$. I claim that there is a path from σ to τ in $\Gamma_{\geq n}(\Delta)$ if and only if there is a path from x to y in $\Delta_{\geq n}$.

Suppose there is an edge path δ from σ to τ in $\Gamma_{\geq n}(\Delta)$. Then construct a path β in $\Delta_{\geq n}$ from x to y as follows. Let $\sigma = \sigma_1, \ldots, \sigma_m = \tau$ be the sequence of vertices in order along δ. Then for $i \in [m-1]$, since σ_i and σ_{i+1} are adjacent in $\Gamma_{\geq n}(\Delta)$, there is some $x_i \in \sigma_i \cap \sigma_{i+1}$ with $\widehat{h}(x_i) \geq n$. Also let $x_0 = x$ and $x_m = y$. Then for $i \in [m]$, we have $x_{i-1}, x_i \in \sigma_{i \geq n}$. Since we've shown that $\sigma_{i \geq n}$ is path connected, there is a path β_i from x_{i-1} to x_i in $\Delta_{\geq n}$. Therefore, $\beta = \beta_0 \cdot \beta_1 \cdot \cdots \cdot \beta_m$ is a path from x to y in $\Delta_{\geq n}$.

Now suppose there is a path $\alpha : [0,1] \to \Delta_{\geq n}$ from x to y. Then construct a path in $\Gamma_{\geq n}(\Delta)$ from σ to τ as follows. We will define a sequence $\sigma = \sigma_1, \ldots, \sigma_m = \tau$ of vertices of $\Gamma_{\geq n}(\Delta)$ and a sequence $0 = t_0, \ldots, t_m \in [0,1]$ such that for $i \in [m]$,

1. $\beta(t_{i-1}), \beta(t_i) \in \sigma_i$ and

2. $\beta((t_i, 1]) \cap \sigma_i = \emptyset$.

Having defined t_{i-1} and σ_{i-1}, if $\sigma_{i-1} = \tau$, then $i - 1 = m$ and we are finished. Otherwise, define t_i and σ_i as follows. Since $\beta((t_{i-1}, 1]) \cap \sigma_i = \emptyset$, we have $\beta(t_{i-1}) \in \partial\sigma$ and there exist other vertices of $\Gamma_{\geq n}(\Delta)$ containing $\beta(t_{i-1})$. Let Σ_i be the set of vertices of $\Gamma_{\geq n}(\Delta)$ other than σ_{i-1} containing $\beta(t_{i-1})$. Define

$$t_i = \max(\{t \in [t_{i-1}, 1] | \beta(t) \in \sigma' \text{ for some } \sigma' \in \Sigma_i\})$$

and let $\sigma_i \in \Sigma_i$ such that $\beta(t_i) \in \sigma_i$. So by definition, $\beta((t_i, 1]) \cap \sigma_i = \emptyset$. Note that $\beta((t_{i-1}, 1])$ must contain some point in an element of Σ_i, implying that $t_{i-1} < t_i$. Therefore, for $j \in [i-1]$, since $\beta((t_j, 1]) \cap \sigma_j = \emptyset$ and $t_j < t_i$, we must have that $\sigma_i \neq \sigma_j$.

Because Δ has finitely many 2-cells, this construction must end with a sequence $\sigma = \sigma_1, \ldots, \sigma_m = \tau$. Note that for $i \in [m-1]$, we have $\beta(t_i) \in \sigma_i \cap \sigma_{i+1}$ and $\widehat{h}(\beta(t_i)) \geq n$. Therefore, σ_i and σ_{i+1} are adjacent in $\Gamma_{\geq n}(\Delta)$. Therefore, this sequence gives a path from σ to τ in $\Gamma_{\geq n}(\Delta)$.

□

The following lemma about connected graphs will be useful for arguments about $\Gamma_{\geq n}$. Given a graph Γ, let $V(\Gamma)$ denote the vertex set of Γ and $E(\Gamma)$ denote the edge set of Γ.

Lemma 3.2.6. *Let Γ and C be connected graphs. Let $V \subseteq V(\Gamma)$, and let D be the induced subgraph on the vertices in $V(\Gamma) \setminus V$. Let Γ' be a graph with $V(\Gamma') = V(C) \sqcup V(D)$, $E(\Gamma') \supseteq E(C) \sqcup E(D)$, and such that there exists an edge of Γ' connecting a*

$d \in V(D)$ to some $c \in V(C)$ whenever there is an edge in Γ connecting d to some $v \in V$. Then Γ' is connected.

Proof. We will show that every vertex $d \in V(D)$ of Γ' is in the same component of Γ' as C. Let $d \in V(D)$. Let α be a shortest path in Γ from d to any vertex $v \in V$. Then let d' be the second-to-last vertex on α and let α' be the initial segment of α from d to d'. We must have that $\alpha' \subseteq D$, since otherwise there is a shorter path from d to V. Then we know that there is a vertex $c \in V(C)$ such that d' and c are adjacent in Γ'. So d, d', and c are all in the same component of Γ'. Since C is connected, this means that d is in the same component of Γ' as C. Therefore, every vertex of Γ' is in the same component. $\square$

Our strategy for the proof of Theorem 3.2.3 will be, for an arbitrary word $w =_G 1$, to start with a van Kampen diagram Δ_0 for w of minimal intrinsic diameter and make adjustments to it until $\Gamma_{\geq n}(\Delta_0)$ is connected for all $n \in \mathbb{N}[\frac{1}{2}]$ less than $\mathrm{IDiam}(\Delta_0)$. The resulting van Kampen diagram will have only one component that contributes to $\mathrm{IAV}(\Delta_0)$ for each $n \in \mathbb{N}$ less than $\mathrm{IDiam}(\Delta_0)$. We will make these adjustments inductively, by taking van Kampen diagrams that have connected $\Gamma_{\geq n-\frac{1}{2}}$ and removing any obstructions to $\Gamma_{\geq n}$ also being connected. Below we define the relevant obstructions to the connectedness of $\Gamma_{\geq n}$, which depend on whether $n \in \mathbb{N}$ or $n \in \mathbb{N} + \frac{1}{2}$.

The obstruction for $n \in \mathbb{N}$: Given a van Kampen diagram Δ and $n \in \mathbb{N}$, let $E_{n-1}(\Delta)$ be the subset of 1-cells of Δ defined as follows. Let e be a 1-cell in Δ connecting two 0-cells of height $n-1$. Then $e \in E_{n-1}(\Delta)$ if and only if e is in the boundary of two 2-cells, σ_1 and σ_2, and for $k \in \{1,2\}$, we have that $\sup(\widehat{I}(\partial\sigma_k \setminus e)) > n-1$.

Suppose Δ is a van Kampen diagram and $n \in \mathbb{N}$ such that $\Gamma_{\geq n-\frac{1}{2}}(\Delta)$ is connected but $\Gamma_{\geq n}(\Delta)$ is not. Figure 3.5 demonstrates the relevance of the 1-cells in $E_{n-1}(\Delta)$ to the connectedness of $\Gamma_{\geq n}(\Delta)$. Each such 1-cell potentially represents an edge between σ_1 and σ_2 (viewed as vertices of $\Gamma_{\geq n-\frac{1}{2}}(\Delta)$) that does not exist in $\Gamma_{\geq n}(\Delta_{n-\frac{1}{2}})$. To construct a van Kampen diagram Δ' from Δ such that $\Gamma_{\geq n}(\Delta')$ is connected, our strategy will be to take Δ and replace each 1-cell of $E_{n-1}(\Delta)$ with an edge path containing two 0-cells of height n, which will ensure that σ_1 and σ_2 are adjacent in $\Gamma_{\geq n}(\Delta')$.

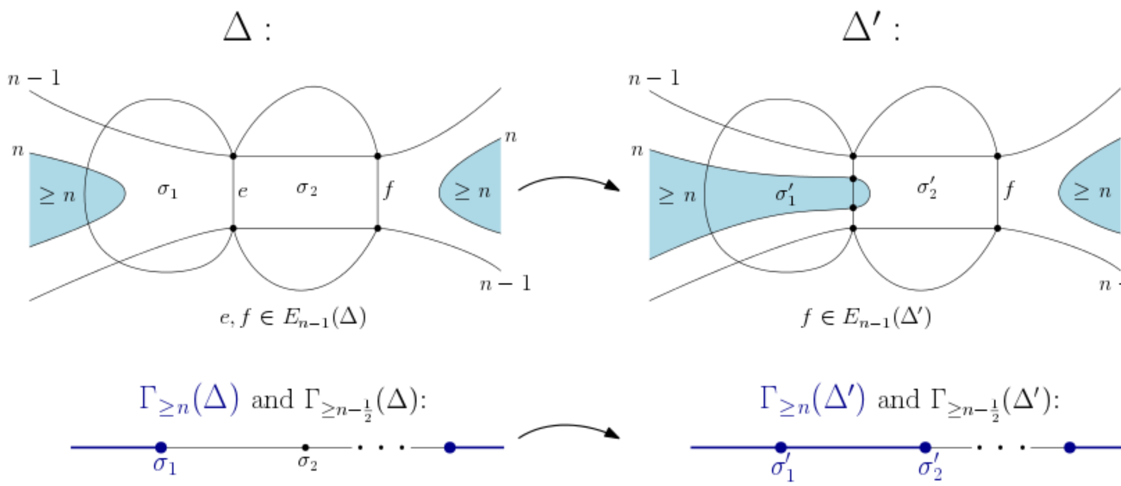

Figure 3.5: An example of the changes to $\Gamma_{\geq n}$, where Δ is a diagram with some 1-cell $e \in E_{n-1}(\Delta)$, and Δ' is the diagram after e has been replaced with an edge path. We highlight $\Gamma_{\geq n}$ as a subgraph of $\Gamma_{\geq n-\frac{1}{2}}$ by making it thicker and blue.

The obstruction for $n \in \mathbb{N}+\frac{1}{2}$: Given a van Kampen diagram Δ and $n \in \mathbb{N}+\frac{1}{2}$, let $X_{\lfloor n \rfloor}(\Delta)$ be the subset of 0-cells of Δ defined as follows. A 0-cell x with $I(x) = \lfloor n \rfloor$ is in $X_{\lfloor n \rfloor}(\Delta)$ if and only if there exist distinct vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Delta)$, σ and τ, both containing x such that σ is a vertex of $\Gamma_{\geq \lfloor n \rfloor+\frac{1}{2}}(\Delta)$ and one of the following properties is true:

(1) τ is not a vertex of $\Gamma_{\geq \lfloor n \rfloor+\frac{1}{2}}(\Delta)$ and there is a 0-cell $y \in \tau$ with $I(y) = \lfloor n \rfloor$ and $y \neq x$, or

(2) σ and τ are in different components of $\Gamma_{\geq \lfloor n \rfloor+\frac{1}{2}}(\Delta)$.

Suppose Δ is a van Kampen diagram and $n \in \mathbb{N} + \frac{1}{2}$ such that $\Gamma_{\geq \lfloor n \rfloor}(\Delta)$ is connected but $\Gamma_{\geq n}(\Delta)$ is not. Figure 3.6 demonstrates the relevance of the 0-cells in $X_{\lfloor n \rfloor}(\Delta)$ to the connectedness of $\Gamma_{\geq n}(\Delta)$. Our strategy for constructing a van Kampen diagram Δ' from Δ such that $\Gamma_{\geq n}(\Delta')$ is also connected will be, for each $x \in X_{\lfloor n \rfloor}(\Delta)$, to replace 1-cells of height less than n that contain x with edge paths containing a point of height at least n. This ensures that that every 2-cell containing x is in the same component of $\Gamma_{\geq n}(\Delta')$.

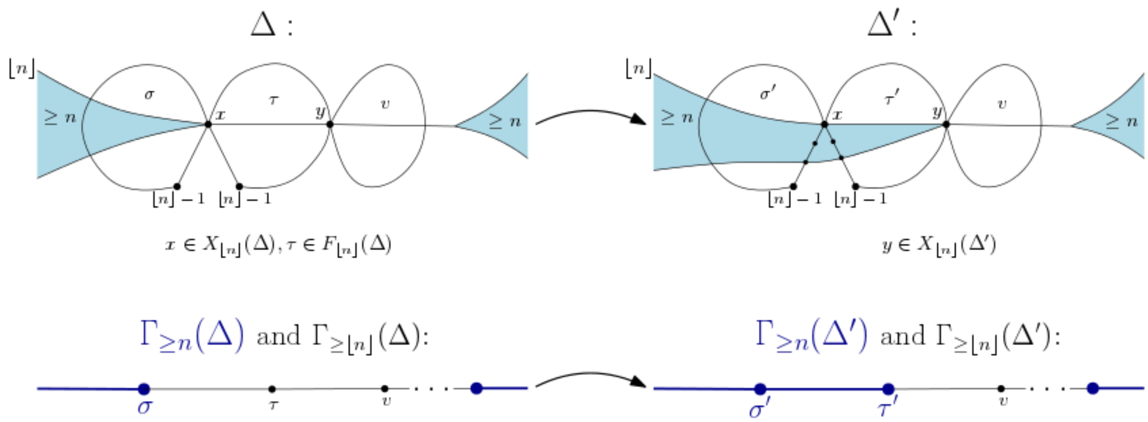

Figure 3.6: An example of the changes to $\Gamma_{\geq n}$, where Δ is a diagram with some 0-cell $x \in X_{\lfloor n \rfloor}(\Delta)$, and Δ' is the diagram after some 1-cells around x have been replaced with edge paths. In this case, $y \in X_{\lfloor n \rfloor}(\Delta) \setminus X_{\lfloor n \rfloor}(\Delta')$, but this only happens because $\tau \in F_{\lfloor n \rfloor}(\Delta) \setminus F_{\lfloor n \rfloor}(\Delta')$. We highlight $\Gamma_{\geq n}$ as a subgraph of $\Gamma_{\geq \lfloor n \rfloor}$ by making it thicker and blue.

Now, this replacement process can create additional elements of $X_{\lfloor n \rfloor}$, as shown in Figure 3.6. However, this can only happen finitely many times, once for each element of a specific subset of the vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Delta)$, which we define below.

Given a van Kampen diagram Δ and $n \in \mathbb{N} + \frac{1}{2}$, define $F_{\lfloor n \rfloor}(\Delta)$ to be the set of vertices τ of $\Gamma_{\geq \lfloor n \rfloor}(\Delta)$ not in $\Gamma_{\geq n}(\Delta)$ such that τ contains at least two 0-cells x_1 and x_2 with $I(x_1), I(x_2) = \lfloor n \rfloor$ and there are vertices $\upsilon_1, \upsilon_2 \neq \tau$ of $\Gamma_{\geq \lfloor n \rfloor}(\Delta)$ with $x_1 \in \upsilon_1$ and $x_2 \in \upsilon_2$.

The vertices in $F_{\lfloor n \rfloor}(\Delta)$ are, intuitively, vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Delta)$ that *could* be cut vertices. As such, they may separate different components of $\Gamma_{\geq n}(\Delta)$ in $\Gamma_{\geq \lfloor n \rfloor}(\Delta)$, and may need to become part of $\Gamma_{\geq n}(\Delta')$ in order to make it connected. Whenever removing one element of $X_{\lfloor n \rfloor}$ creates another, we have also added a vertex of $F_{\lfloor n \rfloor}$ to $\Gamma_{\geq n}$, thus still making progress towards a diagram with connected $\Gamma_{\geq n}$.

To formalize this progress, we will use the lexicographic order on $\mathbb{N}^2$: for $(n_1, n_2), (m_1, m_2) \in \mathbb{N}^2$, let $(n_1, n_2) < (m_1, m_2)$ if and only if $n_1 < m_1$ or $n_1 = m_1$ and $n_2 < m_2$. Then $(|F_{\lfloor n \rfloor}|, |X_{\lfloor n \rfloor}|)$ will decrease in this order for each element of $X_{\lfloor n \rfloor}$ that we remove.

We are now prepared to prove Theorem 3.2.3.

Proof of Theorem 3.2.3. Let G be a group with finite presentation $\mathcal{P} = \langle A|R \rangle$. Let $\rho = \max\{\ell(r) : r \in R\}$, and let $\bar{\rho} = 2\lceil \frac{\rho}{2} \rceil$, the smallest even number greater than or equal to ρ. Let R' be the (finite) set of words $r \in (A \cup A^{-1})^*$ of length at most $\max(\overline{\rho}, 4)$ such that $r =_G 1$. Note that $R \subseteq R'$. Then $\mathcal{P}' = \langle A|R' \rangle$ is a finite presentation for G.

Let $w \in (A \cup A^{-1})^*$ with $w =_G 1$. Then there is a van Kampen diagram Δ_0 for w with respect to $\mathcal{P}$ such that $\text{IDiam}(\Delta_0) \leq \text{IDiam}_{\mathcal{P}}(\ell(w))$.

Step A: We will construct a sequence of van Kampen diagrams $\Delta_0, \Delta_{0.5}, \Delta_1, \ldots, \Delta_{\text{IDiam}(\Delta_0)-\frac{1}{2}}$ for w with respect to $\mathcal{P}'$ by induction such that for each $n \in \{0, 0.5, \ldots, \text{IDiam}(\Delta_0) - \frac{1}{2}\}$, Δ_n has the following properties:

(3) For all $i \in \{0, 0.5, \ldots, n\}$, $\Gamma_{\geq i}(\Delta_n)$ is connected.

(4) $\text{IDiam}(\Delta_n) \leq \max(\text{IDiam}(\Delta_0), n + 1 + \lfloor \frac{1}{4} \max(\overline{\rho}, 4) \rfloor)$.

After constructing these van Kampen diagrams, Step B of the proof will be to show that $\text{AV}_s(\Delta_{\text{IDiam}(\Delta_0)-\frac{1}{2}}) \leq \text{IDiam}_{\mathcal{P}}(w)$.

For the base case, note that $\{x \in \Delta_0 | \widehat{h}(x) \geq 0\} = \Delta_0$ is path connected, so $\Gamma_{\geq 0}(\Delta_0)$ is connected by Proposition 3.2.5. Also, $\mathrm{IDiam}(\Delta_0) \leq \mathrm{IDiam}(\Delta_0)$, so Δ_0 satisfies both properties.

Now let $n \in \{0.5, \ldots, \mathrm{IDiam}(\Delta_0) - \frac{1}{2}\}$ and suppose that $\Delta_{n-\frac{1}{2}}$ is a van Kampen diagram for w with the desired properties. Then we will construct a van Kampen diagram Δ_n for w from $\Delta_{n-\frac{1}{2}}$ with the same properties and such that $\Gamma_{\geq n}(\Delta_n)$ is connected. How to do so depends on whether $n \in \mathbb{N}$ or not.

Case 1: Suppose $n \in \mathbb{N}$.

Notational Aside: Note that we use $\Delta^{(j)}$ for the j-skeleton of a van Kampen diagram Δ. Below we define some van Kampen diagrams $\Delta^j_{n-\frac{1}{2}}$ for some natural numbers j, and $\Delta^j_{n-\frac{1}{2}}$ is not to be confused with the j-skeleton of $\Delta_{n-\frac{1}{2}}$.

Let $\Delta^0_{n-\frac{1}{2}} = \Delta_{n-\frac{1}{2}}$. We will construct a sequence of van Kampen diagrams $\Delta^0_{n-\frac{1}{2}}, \ldots, \Delta^m_{n-\frac{1}{2}} = \Delta_n$ for w with respect to $\mathcal{P}'$ by induction such that for each $j \in \{0, \ldots, m\}$, $\Delta^j_{n-\frac{1}{2}}$ has the following properties:

(5) For all $i \in \{0, 0.5, \ldots, n - \frac{1}{2}\}$, $\Gamma_{\geq i}(\Delta^j_{n-\frac{1}{2}})$ is connected.

(6) $\mathrm{IDiam}(\Delta^j_{n-\frac{1}{2}}) \leq \max(\mathrm{IDiam}(\Delta_0), n + 1 + \lfloor \frac{1}{4} \max(\overline{\rho}, 4) \rfloor)$.

(7) If $j \geq 1$, then $E_{n-1}(\Delta^j_{n-\frac{1}{2}}) \subsetneq E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}})$.

Note that $\Delta^0_{n-\frac{1}{2}}$ already satisfies these properties by the inductive hypothesis of the induction over n.

Case 1 Step 1: Construct $\Delta^j_{n-\frac{1}{2}}$ from $\Delta^{j-1}_{n-\frac{1}{2}}$.

Let $j \geq 1$ and suppose we have constructed $\Delta^{j-1}_{n-\frac{1}{2}}$ with the above properties. If $E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}}) = \emptyset$, then $m = j - 1$ and we are done. Otherwise, construct $\Delta^j_{n-\frac{1}{2}}$ from $\Delta^{j-1}_{n-\frac{1}{2}}$ in the following way. (See Figure 3.7 for an example.) Let $e \in E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}})$. Then e is in the boundary of two 2-cells, σ_1 and σ_2. Let x and y be the endpoints of e, and let $a \in A \cup A^{-1}$ be the letter labeling the directed edge path $\overrightarrow{e}$ from x to y along e. For $k \in \{1, 2\}$, let γ_k be the directed edge path such that $\overrightarrow{e} \cdot \gamma_k$ is the boundary circuit of σ_k.

Now replace e with a directed edge path $\overrightarrow{e}'$ from x to y labeled by the word $aa^{-1}a$. Let x' be the 0-cell adjacent to y on $\overrightarrow{e}'$, and let y' be the 0-cell adjacent to x on $\overrightarrow{e}'$. For $k \in \{1, 2\}$, this replaces σ_k with a new 2-cell σ'_k with boundary circuit $\overrightarrow{e}' \cdot \gamma_k$. Let $w_k \in (A \cup A^{-1})^*$ be the word that labels γ_k. Then the boundary of σ'_k is labeled by $aa^{-1}aw_k =_G aw_k =_G 1$, since aw_k labels the boundary circuit of σ_k.

Now if $\ell(aa^{-1}aw_k) \leq \max(\overline{\rho}, 4)$, then $aa^{-1}aw_k \in R'$. Otherwise, $|\gamma_k| > 1$, so there are vertices on the boundary of σ_k other than x and y. Then by the fact that $e \in E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}})$, we know that $\sup(\widehat{I}(\partial\sigma_k \setminus e)) \geq n - 1$. So there is some 0-cell $z_k \neq x, y$ on γ_k with $I(z_k) \geq n - 1$. In fact, because height differs by at most 1 between adjacent 0-cells, we can choose z_k such that $I(z_k) \in \{n - 1, n\}$. Let α_k be the path from y to z_k and β_k the path from x to z_k such that $\alpha_k \cdot \overline{\beta_k} = \gamma_k$. If $|\alpha_k| \leq |\beta_k|$, then let u_k be the word labeling α_k and let $s_k = y'$. Otherwise, let u_k be the word labeling β_k and let $s_k = x'$. Then add a directed edge path α'_k through the interior of σ'_k from s_k to z_k labeled by u_k. This splits σ'_k into two 2-cells, one with boundary circuit labeled by the word $au_ku_k^{-1}a^{-1}$ and the other labeled by aw_k. Let σ^x_k be the one containing x, and σ^y_k the one containing y.

We already know that $aw_k \in R'$ and $au_ku_k^{-1}a^{-1} =_G 1$. Since $\ell(u_k) = \min(|\alpha_k|, |\beta_k|)$,

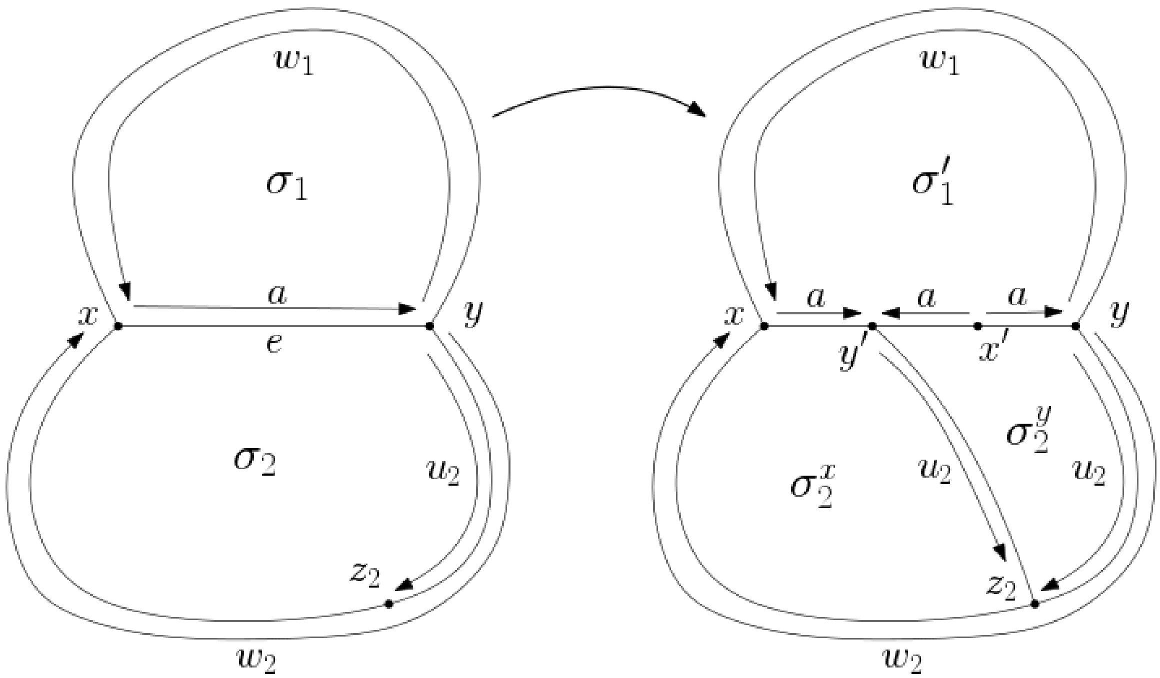

Figure 3.7: An example of the changes made to $\Delta^{j-1}_{n-\frac{1}{2}}$ to produce $\Delta^{j}_{n-\frac{1}{2}}$

we have that

$$\begin{aligned}\ell(au_ku_k^{-1}a^{-1}) &= 2\ell(u_k)+2\\ &\leq |\alpha_k|+|\beta_k|+2\\ &= \ell(aw_k)+1\\ &\leq \max(\overline{\rho},4)+1\end{aligned}$$

Now, since $\overline{\rho}$ is even, $\max(\overline{\rho},4)+1$ is odd. Note that $\ell(au_ku_k^{-1}a^{-1})$ is even, so we can therefore strengthen the inequality to $\ell(au_ku_k^{-1}a^{-1}) \leq \max(\overline{\rho},4)$. Therefore, $au_ku_k^{-1}a^{-1} \in R'$.

Having made these changes to $\Delta^{j-1}_{n-\frac{1}{2}}$, we will call the resulting diagram $\Delta^{j}_{n-\frac{1}{2}}$. We have shown that the boundary circuit of each new 2-cell of $\Delta^{j}_{n-\frac{1}{2}}$ is labeled by a word in R', so $\Delta^{j}_{n-\frac{1}{2}}$ is a van Kampen diagram with respect to $\mathcal{P}'$. We have also not added or removed any 1-cells on the boundary, so $\Delta^{j}_{n-\frac{1}{2}}$ is a van Kampen diagram for the word w, the same as $\Delta^{j-1}_{n-\frac{1}{2}}$.

Case 1 Step 2: Show that $\Delta^{j}_{n-\frac{1}{2}}$ satifies properties (6) and (7) from the induction on j: that $\mathrm{IDiam}(\Delta^{j}_{n-\frac{1}{2}}) \leq \max(\mathrm{IDiam}(\Delta_0), n+1+\lfloor \frac{1}{4}\max(\overline{\rho}, 4)\rfloor)$, and that $E_{n-1}(\Delta^{j}_{n-\frac{1}{2}}) \subsetneq E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}})$.

We will first show that every $p \in {\Delta^{j-1}_{n-\frac{1}{2}}}^{(0)}$ has the same height in $\Delta^{j}_{n-\frac{1}{2}}$ as in $\Delta^{j-1}_{n-\frac{1}{2}}$. See Figure 3.8 for an example of the paths used in this argument. Let γ be a geodesic edge path from the basepoint $*$ to p in ${\Delta^{j}_{n-\frac{1}{2}}}^{(1)}$. Suppose that γ is not contained in ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$. Then γ must at some point enter, and at some point exit, ${\Delta^{j}_{n-\frac{1}{2}}}^{(1)} \setminus {\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$ through one of the 0-cells x, y, z_1, or z_2, since these are the only 0-cells of $\Delta^{j-1}_{n-\frac{1}{2}}$ on 1-cells not in $\Delta^{j-1}_{n-\frac{1}{2}}$. So let q_{F} be the first point on γ in $\{x, y, z_1, z_2\}$, and let q_{L} be the last point on γ in $\{x, y, z_1, z_2\}$. Recall that these points have height either $n-1$ or n in $\Delta^{j-1}_{n-\frac{1}{2}}$. Let α be the path from $*$ to q_{F}, β the path from q_{F} to q_{L}, and ω the path from q_{L} to p, such that $\alpha \cdot \beta \cdot \omega = \gamma$. Note that, since γ contains a 1-cell not in ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$, we must have that $|\beta| \geq 1$. So $|\alpha \cdot \beta| \geq |\alpha| + 1$. But $|\alpha| \geq n-1$, since α is a path from $*$ to q_{F} in ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$. So $|\alpha \cdot \beta| \geq n$. Therefore, there is some path δ from $*$ to q_L in ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$ with $|\delta| \leq |\alpha \cdot \beta|$. Then $\delta \cdot \omega$ is a path from $*$ to p in ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$ with $|\delta \cdot \omega| \leq |\gamma|$. So the height of p has not decreased. Furthermore, the only 1-cell of ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$ not in ${\Delta^{j}_{n-\frac{1}{2}}}^{(1)}$ is e, but since e connects two 0-cells of height $n-1$, it cannot be in any geodesic in ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$. Therefore, every geodesic in ${\Delta^{j-1}_{n-\frac{1}{2}}}^{(1)}$ is in ${\Delta^{j}_{n-\frac{1}{2}}}^{(1)}$, meaning that the height of p has not increased, either.

In particular, we still have that $I(x), I(y), I(z_1), I(z_2) \in \{n-1, n\}$. Now let $p \in {\Delta^{j}_{n-\frac{1}{2}}}^{(0)} \setminus {\Delta^{j-1}_{n-\frac{1}{2}}}^{(0)}$. Since every path from $*$ to p contains one of x, y, z_1 and z_2, we have that $I(p) \geq n$. If p is in $\overrightarrow{e}'$, then it is adjacent to one of x and y, so $I(p) = n$. Otherwise, p is on an edge path of length at most $\frac{1}{2}\max(\overline{\rho}, 4)$ connecting two 0-cells of height at most n. Therefore, $I(p) \leq n + \lfloor \frac{1}{4}\max(\overline{\rho}, 4)\rfloor$.

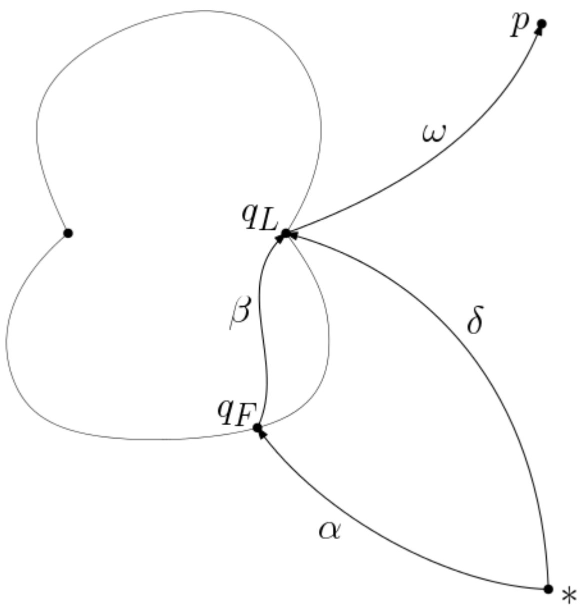

Figure 3.8: Showing that every 0-cell p of $\Delta^{j-1}_{n-\frac{1}{2}}$ has the same height in $\Delta^{j}_{n-\frac{1}{2}}$.

Since we know that $\text{IDiam}(\Delta^{j-1}_{n-\frac{1}{2}}) \leq \max(\text{IDiam}(\Delta_0), n+1+\lfloor\frac{1}{4}\max(\overline{\rho},4)\rfloor)$, this implies that $\text{IDiam}(\Delta^{j}_{n-\frac{1}{2}}) \leq \max(\text{IDiam}(\Delta_0), n+1+\lfloor\frac{1}{4}\max(\overline{\rho},4)\rfloor)$. Furthermore, each 1-cell in $\Delta^{j}_{n-\frac{1}{2}}{}^{(1)} \setminus \Delta^{j-1}_{n-\frac{1}{2}}{}^{(1)}$ has at least one endpoint in $\Delta^{j}_{n-\frac{1}{2}}{}^{(0)} \setminus \Delta^{j-1}_{n-\frac{1}{2}}{}^{(0)}$. Since each such endpoint has height at least n, we have not added any 1-cells that connect two 0-cells of height $n-1$. However, we have removed e, a 1-cell in $E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}})$. Therefore, $E_{n-1}(\Delta^{j}_{n-\frac{1}{2}}) \subsetneq E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}})$, as desired.

Case 1 Step 3: Show that $\Delta^{j}_{n-\frac{1}{2}}$ satisfies property (5) from the induction on j: that for all $i \in \{0, 0.5, \ldots, n-\frac{1}{2}\}$, $\Gamma_{\geq i}(\Delta^{j}_{n-\frac{1}{2}})$ is connected.

Let $i \in \{0, 0.5, \ldots, n-\frac{1}{2}\}$. By the induction hypothesis, we know that $\Gamma_{\geq i}(\Delta^{j-1}_{n-\frac{1}{2}})$ is connected. Let $V = \{\sigma_1, \sigma_2\}$ and let $D = \Gamma_{\geq i}(\Delta^{j-1}_{n-\frac{1}{2}}) \setminus V$. Let C be the largest subgraph of $\Gamma_{\geq i}(\Delta^{j}_{n-\frac{1}{2}})$ that does not contain any vertices of $\Gamma_{\geq i}(\Delta^{j-1}_{n-\frac{1}{2}})$. First note that C is connected. The vertices of C are exactly the 2-cells of $\Delta^{j}_{n-\frac{1}{2}}$ that are not in $\Delta^{j-1}_{n-\frac{1}{2}}$, since each of these 2-cells contains one of x' and y', and $I(x') = I(y') = n > i$. If σ_1' is a 2-cell of $\Delta^{j}_{n-\frac{1}{2}}$, then it has both x' and y' on its boundary, meaning that it is adjacent to either σ_2' or both of σ_2^x and σ_2^y in $\Gamma_{\geq i}(\Delta^j_{n-\frac{1}{2}})$, putting these all in the same component of C. Otherwise, σ_1^x and σ_1^y are both 2-cells of $\Delta^j_{n-\frac{1}{2}}$. Since they both have s_1 on their boundaries, they are adjacent to each other in $\Gamma_{\geq i}(\Delta^j_{n-\frac{1}{2}})$. Furthermore, one of these has both x' and y' on its boundary. So by the same reasoning as above, all of these vertices are in the same component of C.

Let $d \in D$ be adjacent to σ_k in $\Gamma_{\geq i}(\Delta^j_{n-\frac{1}{2}})$ for some $k \in \{1, 2\}$. Then there is some point $q \in d \cap \sigma_k$ with $I(q) \geq i$. Note that q is not in $\text{int}(e)$, since $d \notin V$. Since $q \in \partial\sigma_k \setminus \text{int}(e)$, we have $q \in \partial c$ for some vertex c in C. So d is adjacent to c in $\Gamma_{\geq i}(\Delta^j_{n-\frac{1}{2}})$. Therefore, Proposition 3.2.6 implies that $\Gamma_{\geq i}(\Delta^j_{n-\frac{1}{2}})$ is connected.

This concludes the induction on j.

Case 1 Step 4: Define Δ_n and show that Δ_n satisfies properties (3) and (4) of the induction on n.

Since $E_{n-1}(\Delta^j_{n-\frac{1}{2}}) \subsetneq E_{n-1}(\Delta^{j-1}_{n-\frac{1}{2}})$ for all $j \geq 1$, and $\Delta_{n-\frac{1}{2}}$ has finitely many 1-cells, the process from Case 1 Steps 1-3 must terminate in a van Kampen diagram Δ_n for w with respect to $\mathcal{P}'$ with the following properties:

- For all $i \in \{0, 0.5, \ldots, n - \frac{1}{2}\}$, $\Gamma_{\geq i}(\Delta_n)$ is connected (property (3) from the induction on n, but excluding $i = n$).
- $\text{IDiam}(\Delta_n) \leq \max(\text{IDiam}(\Delta_0), n + 1 + \left\lfloor \frac{1}{4} \max(\overline{\rho}, 4) \right\rfloor)$ (property (4) from the induction on n).
- $E_{n-1}(\Delta_n) = \emptyset$.

So to finish this step, we need only show that $\Gamma_{\geq n}(\Delta_n)$ is connected. Note that $\Gamma_{\geq n}(\Delta_n)$ is a subgraph of $\Gamma_{\geq n-\frac{1}{2}}(\Delta_n)$. Let σ and τ be vertices of $\Gamma_{\geq n}(\Delta_n)$. Then since $\Gamma_{\geq n-\frac{1}{2}}(\Delta_n)$ is connected, there is a simple path α from σ to τ in $\Gamma_{\geq n-\frac{1}{2}}(\Delta_n)$.

Suppose by way of contradiction that α is not contained in $\Gamma_{\geq n}(\Delta_n)$. Then let α' be the longest initial segment of α contained in $\Gamma_{\geq n}(\Delta_n)$. Let σ' be the last vertex of α' and let τ' be the next vertex after σ' along α. Then because σ' and τ' are adjacent in $\Gamma_{\geq n-\frac{1}{2}}(\Delta_n)$, there is some point $x \in \sigma' \cap \tau'$ with $\widehat{I}(x) \geq n - \frac{1}{2}$. But because either τ' is not in $\Gamma_{\geq n}(\Delta_n)$ or it is but σ' and τ' are not adjacent in $\Gamma_{\geq n}(\Delta_n)$, we must have that $\widehat{I}(x) < n$. So $\widehat{I}(x) = n - \frac{1}{2}$, meaning that $x \in \text{int}(e)$ for some 1-cell e of Δ_n with e in $\partial\sigma'$ and $\partial\tau'$. Therefore, one endpoint of e has height $n-1$, and the other has height either $n-1$ or n. Note also that this implies that σ' and τ' are 2-cells.

Suppose by way of contradiction that both endpoints of e have height $n-1$. Since σ' is a vertex of $\Gamma_{\geq n}(\Delta_n)$, we know it contains a point of height at least n. So $\sup(\widehat{I}(\partial\sigma' \setminus e)) > n-1$. If $\tau' = \tau$, then by the same reasoning, τ' contains a point of height at least n, making $\sup(\widehat{I}(\partial\tau' \setminus e)) > n-1$.

If $\tau' \neq \tau$, then let υ be the next vertex after τ' along α. So there is some point $y \in \partial\tau' \cap \partial\upsilon'$ with $\widehat{I}(y) \geq n - \frac{1}{2}$. Suppose that $y \in e$. Then since both endpoints of e have height $n-1$, we must have that $y \in \text{int}(e)$. This implies that e is in $\partial\upsilon$. But since e is also in $\partial\sigma'$ and $\partial\tau'$, and e can only be on the boundary of two distinct 2-cells of a planar 2-complex, we must have that two of σ', τ', and υ are equal. This contradicts the fact that α is a simple path, so we must have $y \notin e$. Therefore, $\sup(\widehat{I}(\partial\tau' \setminus e)) > n-1$.

Since both endpoints of e have height $n-1$, $\sup(\widehat{I}(\partial\sigma' \setminus e)) > n-1$, and $\sup(\widehat{I}(\partial\tau' \setminus e)) > n-1$, we must have that $e \in E_{n-1}(\Delta_n)$. This contradicts the fact that $E_{n-1}(\Delta_n) = \emptyset$. Therefore, one endpoint of e has height n. Call this endpoint z. Now $z \in \tau'$, which implies that τ' is a vertex of $\Gamma_{\geq n}(\Delta_n)$. Furthermore, $z \in \sigma' \cap \tau'$, meaning that σ' and τ' are adjacent in $\Gamma_{\geq n}(\Delta_n)$. This contradicts that α' is the longest initial segment of α contained in $\Gamma_{\geq n}(\Delta_n)$. Therefore, α is contained in $\Gamma_{\geq n}(\Delta_n)$. This shows that $\Gamma_{\geq n}(\Delta_n)$ is connected, as desired, completing Case 1.

Case 2: Suppose instead that $n \in \mathbb{N}_0 + \frac{1}{2}$.

We will construct a sequence of van Kampen diagrams $\Delta^0_{\lfloor n \rfloor}, \ldots, \Delta^m_{\lfloor n \rfloor}$ for w with respect to $\mathcal{P}'$ by induction such that the following properties hold for $j = 0, \ldots, l$:

(8) For all $i \in \{0, 0.5, \ldots, \lfloor n \rfloor\}$, $\Gamma_{\geq i}(\Delta^j_{\lfloor n \rfloor})$ is connected.

(9) $\mathrm{IDiam}(\Delta^j_{\lfloor n \rfloor}) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rfloor + 1 + \lfloor \frac{1}{4} \max(\overline{\rho}, 4) \rfloor)$.

(10) For $j \geq 1$, we have $(|F_{\lfloor n \rfloor}(\Delta^j_{\lfloor n \rfloor})|, |X_{\lfloor n \rfloor}(\Delta^j_{\lfloor n \rfloor})|) < (|F_{\lfloor n \rfloor}(\Delta^{j-1}_{\lfloor n \rfloor})|, |X_{\lfloor n \rfloor}(\Delta^{j-1}_{\lfloor n \rfloor})|)$.

For the base case, note that we have already shown that $\Delta^0_{\lfloor n \rfloor} = \Delta_{\lfloor n \rfloor}$ satisfies these properties.

Case 2 Step 1: For $j \geq 1$, we will construct $\Delta^j_{\lfloor n \rfloor}$ from $\Delta^{j-1}_{\lfloor n \rfloor}$ by induction.

Let $x \in X_{\lfloor n \rfloor}(\Delta^{j-1}_{\lfloor n \rfloor})$. Let N_x be the number of 1-cells containing x. Let e_1 be one such 1-cell with other endpoint y_1 such that $I(y_1) = \lfloor n \rfloor - 1$; such a 1-cell must exist because any geodesic from the basepoint to x has such a 1-cell as its last edge. Let $e_1, \ldots, e_{N_x}$ be the sequence of 1-cells containing x in counter-clockwise order around x starting at e_1. For $k \in [N_x]$, let y_k be the endpoint of e_k that is not x, and let $a_k \in A \cup A^{-1}$ be the letter labeling the directed edge $\overrightarrow{e_k}$ from x to y_k along e_k.

We will construct another sequence of van Kampen diagrams $\Delta^{j-1}_{\lfloor n \rfloor} = \Lambda^0_x, \ldots, \Lambda^{N_x}_x = \Delta^j_{\lfloor n \rfloor}$ for w with respect to $\mathcal{P}'$ by induction such that the following properties hold for $k \in [N_x]$:

(11) For all $i \in \{0, 0.5, \ldots, \lfloor n \rfloor\}$, $\Gamma_{\geq i}(\Lambda^k_x)$ is connected.

(12) $\mathrm{IDiam}(\Lambda^k_x) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rfloor + 1 + \lfloor \frac{1}{4} \max(\overline{\rho}, 4) \rfloor)$.

(13) For $k > 0$, we have $F_{\lfloor n \rfloor}(\Lambda^k_x) \subseteq F_{\lfloor n \rfloor}(\Lambda^{k-1}_x)$.

(14) For $k > 0$, if $X_{\lfloor n \rfloor}(\Lambda^k_x) \not\subseteq X_{\lfloor n \rfloor}(\Lambda^{k-1}_x)$, then $F_{\lfloor n \rfloor}(\Lambda^k_x) \subsetneq F_{\lfloor n \rfloor}(\Lambda^{k-1}_x)$.

(15) For $0 < k < N_x$, Λ_x^k has a 2-cell $\sigma_{k+1}^{\mathrm{R}}$ adjacent to $\overrightarrow{e_{k+1}}$ on its right side.

For the base case, note that we have already shown that $\Lambda_x^0 = \Delta_{\lfloor n \rfloor}^{j-1}$ satisfies these properties.

Case 2 Step 1: Construct Λ_x^1 from Λ_x^0 such that it satisfies properties (11) - (15).

Case 2 Step 1(a): Construct Λ_x^1 from Λ_x^0 such that it satisfies property (15) above, and show that it also satisfies property (12).

If $\overrightarrow{e_2}$ is adjacent to a 2-cell on its right, name this 2-cell σ_2^{R} and let $\Lambda_x^1 = \Lambda_x^0$. Otherwise, $\overrightarrow{e_2}$ is adjacent to $\mathbb{R}^2 \setminus \Lambda_x^0$ on its right and $\overrightarrow{e_1}$ is adjacent to $\mathbb{R}^2 \setminus \Lambda_x^0$ on its left. So let σ_2^{R} be a 2-cell with boundary circuit labeled by $a_2^{-1}a_1a_1^{-1}a_2$, and glue σ_2^{R} to Λ_x^0 by gluing the initial path along its boundary circuit labeled by $a_2^{-1}a_1$ along the path $\overleftarrow{\overrightarrow{e_2}} \cdot \overrightarrow{e_1}$. Let Λ_x^1 be the resulting diagram. Let $y_{1.5}$ be the 0-cell in σ_2^{R} replacing x on the boundary of the Λ_x^1, and let α be the directed edge path from y_1 to $y_{1.5}$ to y_2 along the boundary of Λ_x^1 labeled by $a_1^{-1}a_2$. See Figure 3.9.

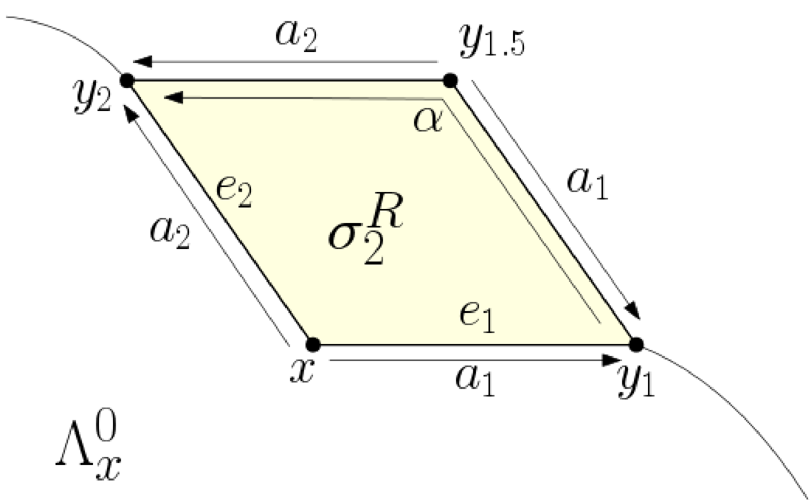

Figure 3.9: Adding σ_2^{R} to Λ_x^0 if e_2 is not adjacent to a 2-cell on its right.

First note that we have replaced $\overleftarrow{\overrightarrow{e_1}} \cdot \overrightarrow{e_2}$, part of the boundary circuit of Λ_x^0, with α. Since these two paths are both labeled by $a_1^{-1}a_2$, the boundary circuit of Λ_x^1 is still w. Also, the boundary circuit of σ_2^{R} is labeled by the word $a_1^{-1}a_2a_2^{-1}a_1$, which is in R' since it freely reduces to the identity and has length 4. So Λ_x^1 is a van Kampen diagram for w with respect to $\mathcal{P}'$.

Also note that every 0-cell of Λ_x^0 still has the same height in Λ_x^1. For suppose γ is a geodesic in Λ_x^1 from the basepoint $*$ to any 0-cell $p \in {\Lambda_x^0}^{(0)}$. If γ is not already contained in Λ_x^0, then it must include $y_{1.5}$. Since geodesics are simple paths, this implies that either α or $\overline{\alpha}$ is the only segment of γ outside of Λ_x^0. So replacing α or $\overline{\alpha}$ with $\overline{\overrightarrow{e_1}} \cdot \overrightarrow{e_2}$ or $\overline{\overrightarrow{e_2}} \cdot \overrightarrow{e_1}$, respectively, in γ results in a path of the same length from $*$ to p contained in Λ_x^0. So the distance from $*$ to p has not decreased. Since $\Lambda_x^0 \subset \Lambda_x^1$, this distance cannot increase, either, so it remains the same.

Now since y_1 and y_2 each share a 1-cell with x and $I(x) = \lfloor n \rfloor$, we have that $I(y_1), I(y_2) \in \{\lfloor n \rfloor - 1, \lfloor n \rfloor, \lfloor n \rfloor + 1\}$. Since every path from $*$ to $y_{1.5}$ contains either y_1 or y_2, this implies that $I(y_{1.5}) \in \{\lfloor n \rfloor, \lfloor n \rfloor + 1, \lfloor n \rfloor + 2\}$. So

$$\begin{aligned} \mathrm{IDiam}(\Lambda_x^1) &\leq \max(\mathrm{IDiam}(\Lambda_x^0), \lfloor n \rfloor + 2) \\ &\leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rfloor + 1 + \left\lfloor \frac{1}{4} \max(\overline{\rho}, 4) \right\rfloor). \end{aligned}$$

Case 2 Step 1(b): Show that Λ_x^1 satisfies property (11) from the induction on k: that for all $i \in \{0, 0.5, \ldots, \lfloor n \rfloor\}$, we have that $\Gamma_{\geq i}(\Lambda_x^1)$ is connected.

Let $i \in \{0, 0.5, \ldots, \lfloor n \rfloor\}$. We know by the inductive hypothesis that $\Gamma_{\geq i}(\Lambda_x^0)$ is connected.

Case 2(a): Suppose that neither e_1 nor e_2 is contained in any 2-cell of Λ_x^0. Since e_1 and e_2 both contain x and $I(x) = \lfloor n \rfloor$, this implies that e_1 and e_2 are both vertices of $\Gamma_{\geq i}(\Lambda_x^0)$. Since they share x, they are also adjacent in $\Gamma_{\geq i}(\Lambda_x^0)$. Now e_1 and e_2 are both contained in σ_2^{R}, so they are not vertices of $\Gamma_{\geq i}(\Lambda_x^1)$, but are replaced by σ_2^{R}, which contains x. However, since σ_2^{R} contains e_1 and e_2, it is adjacent in $\Gamma_{\geq i}(\Lambda_x^1)$ to every vertex that is adjacent to either e_1 or e_2 in $\Gamma_{\geq i}(\Lambda_x^0)$. So by Lemma 3.2.6,

$\Gamma_{\geq i}(\Lambda_x^1)$ is connected.

Case 2(b): Suppose that only one of e_1 and e_2 is contained in a 2-cell of Λ_x^0. Suppose without loss of generality that e_1 is contained in a 2-cell of Λ_x^0 and e_2 is not. Then e_2 is a vertex of $\Gamma_{\geq i}(\Lambda_x^0)$. Since e_2 is contained in σ_2^{R}, e_2 is not a vertex of $\Gamma_{\geq i}(\Lambda_x^1)$, but is replaced by σ_2^{R}. Furthermore, since σ_2^{R} contains e_2, it is adjacent in $\Gamma_{\geq i}(\Lambda_x^1)$ to every vertex that is adjacent to e_2 in $\Gamma_{\geq i}(\Lambda_x^0)$. Therefore, $\Gamma_{\geq i}(\Delta_{\lfloor n \rfloor}^{j-1})$ is isomorphic to a subgraph of $\Gamma_{\geq i}(\Lambda_x^1)$, where the isomorphism is given by the identity on every vertex other than e_2, which gets mapped to σ_2^{R}. This subgraph is connected, and since it contains every vertex of $\Gamma_{\geq i}(\Lambda_x^1)$, so is $\Gamma_{\geq i}(\Lambda_x^1)$.

Case 2(c): Suppose that both e_1 and e_2 are contained in 2-cells of Λ_x^0. Then $\Gamma_{\geq i}(\Lambda_x^0)$ is a subgraph of $\Gamma_{\geq i}(\Lambda_x^1)$, and σ_2^{R} is the only vertex of $\Gamma_{\geq i}(\Lambda_x^1)$ not in $\Gamma_{\geq i}(\Lambda_x^0)$. Since σ_2^{R} contains x, as does the 2-cell containing e_1, σ_2^{R} is adjacent to this 2-cell in $\Gamma_{\geq i}(\Lambda_x^1)$. Since $\Gamma_{\geq i}(\Lambda_x^0)$ is connected, this implies that $\Gamma_{\geq i}(\Lambda_x^1)$ is as well.

Case 2 Step 1(c): Show that $X_{\lfloor n \rfloor}(\Lambda_x^1) \subseteq X_{\lfloor n \rfloor}(\Lambda_x^0)$, implying that Λ_x^1 satisfies property (14) from the induction on k

Since the only change from $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$ to $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^1)$ is the addition of the vertex σ_2^{R} and possible removal of vertices e_1 and/or e_2 (if either of them were not already contained in a 2-cell), the only 0-cells that could be in $X_{\lfloor n \rfloor}(\Lambda_x^1) \triangle X_{\lfloor n \rfloor}(\Lambda_x^0)$ are those contained in σ_2^{R}, e_1, and e_2, that is, x, $y_{1.5}$, y_1, and y_2.

Firstly, we know that $x \in X_{\lfloor n \rfloor}(\Lambda_x^0)$, so $x \notin X_{\lfloor n \rfloor}(\Lambda_x^1) \setminus X_{\lfloor n \rfloor}(\Lambda_x^0)$. Also, the only vertex of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^1)$ containing $y_{1.5}$ is σ_2^{R}, so $y_{1.5} \notin X_{\lfloor n \rfloor}(\Lambda_x^1)$. Furthermore, $I(y_1) = \lfloor n \rfloor - 1$, so $y_1 \notin X_{\lfloor n \rfloor}(\Lambda_x^1)$. Finally, suppose by way of contradiction that $y_2 \in X_{\lfloor n \rfloor}(\Lambda_x^1) \setminus X_{\lfloor n \rfloor}(\Lambda_x^0)$. Then by the definition of $X_{\lfloor n \rfloor}(\Lambda_x^1)$, we have that $I(y_2) = \lfloor n \rfloor$ and there exist distinct vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^1)$, σ and τ, both containing y_2, such that $\sigma \in \Gamma_{\geq n}(\Lambda_x^1)$ and one of properties (1) or (2) from the definition of $X_{\lfloor n \rfloor}$ is true:

(1) τ is not a vertex of $\Gamma_{\geq n}(\Lambda_x^1)$ and there is a 0-cell $z \in \tau$ with $I(z) = \lfloor n \rfloor$ and $z \neq y_2$, or

(2) σ and τ are in different components of $\Gamma_{\geq n}(\Lambda_x^1)$.

But since $y_2 \notin X_{\lfloor n \rfloor}(\Lambda_x^0)$, σ and τ do not satisfy these properties in Λ_x^0. There are two cases for how this could happen.

Case 2(d): Suppose that one of σ or τ is not a vertex of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$. Since heights have been preserved from Λ_x^0 to Λ_x^1, the only way this can happen is if one of σ or τ is not in Λ_x^0, i.e., if one of them is σ_2^{R}. Note that the other must be in Λ_0, since σ and τ are distinct. Since $I(y_2) = \lfloor n \rfloor$ and $I(x) = \lfloor n \rfloor$, any point $z \in \mathrm{int}(e_2)$ has $\widehat{I}(z) = n$. So since $e_2 \subset \sigma_2^{\mathrm{R}}$, σ_2^{R} is a vertex of $\Gamma_{\geq n}(\Lambda_x^1)$. Therefore, either $\sigma = \sigma_2^{\mathrm{R}}$ and property (1) above holds, or $\tau = \sigma_2^{\mathrm{R}}$ and property (2) above holds. Let υ_1 be the vertex of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$ containing e_1, and let υ_2 be the vertex of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$ containing e_2.

Subcase 2(d)(i): $\sigma = \sigma_2^{\mathrm{R}}$ and property (1) above holds. Then υ_2 and τ are both vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$ containing y_2, and τ still satisfies property (1) in Λ_0. This implies that $y_2 \in X_{\lfloor n \rfloor}(\Lambda_x^0)$, a contradiction.

Subcase 2(d)(ii): $\tau = \sigma_2^{\mathrm{R}}$ and property (2) above holds. Then σ is a vertex of $\Gamma_{\geq n}(\Lambda_x^0)$, and since $e_2 \subset \upsilon_2$, so is υ_2. We know that σ and σ_2^{R} are in different components of $\Gamma_{\geq n}(\Lambda_x^1)$. But since $y_2 \notin X_{\lfloor n \rfloor}(\Lambda_x^0)$, we must have that σ and υ_2 are in the same component of $\Gamma_{\geq}(\Lambda_x^0)$. So there is a simple path α from σ to υ_2 in $\Gamma_{\geq n}(\Lambda_x^0)$. Let β be the longest initial segment of α contained in $\Gamma_{\geq n}(\Lambda_x^1)$. If $\beta = \alpha$, then υ_2 is a vertex of $\Gamma_{\geq n}(\Lambda_x^1)$, and since $\upsilon_2 \cap \sigma_2^{\mathrm{R}} \supseteq e_2$, υ_2 and σ_2^{R} are adjacent in $\Gamma_{\geq n}(\Lambda_x^1)$, putting them in the same component. But this implies that σ and σ_2^{R} are in the same component of $\Gamma_{\geq n}(\Lambda_x^1)$, a contradiction. So we must have $\beta \neq \alpha$. Let χ be the endpoint of β, so that χ and σ are in the same component of $\Gamma_{\geq n}(\Lambda_x^1)$. Since every pair of vertices of $\Gamma_{\geq n}(\Lambda_x^0)$ that are also both vertices of $\Gamma_{\geq n}(\Lambda_x^1)$ are either adjacent or not in both graphs, we must have that α contains a vertex of $\Gamma_{\geq n}(\Lambda_x^0)$ that is not a vertex of $\Gamma_{\geq n}(\Lambda_x^1)$. We know that υ_1 and υ_2 are the only possible vertices of $\Gamma_{\geq n}(\Lambda_x^0)$, that may not be in $\Gamma_{\geq n}(\Lambda_x^1)$, which can only happen if $\upsilon_1 = e_1$ or $\upsilon_2 = e_2$. So let $\varepsilon \in \{1, 2\}$ such that $\upsilon_\varepsilon = e_\varepsilon$ and e_ε is the vertex on α that comes after χ. Then because e_ε and χ are adjacent in $\Gamma_{\geq n}(\Lambda_x^0)$ and $e_\varepsilon \subset \sigma_2^{\mathrm{R}}$, we must have that χ and σ_2^{R} are adjacent in $\Gamma_{\geq n}(\Lambda_x^1)$. Therefore, σ_2^{R}, χ, and σ are all in the same component of $\Gamma_{\geq n}(\Lambda_x^1)$, which is a contradiction.

Case 2(e): Suppose that both σ and τ are vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$. Then σ is also a vertex of $\Gamma_{\geq n}(\Lambda_x^0)$, since it is a vertex of $\Gamma_{\geq n}(\Lambda_x^1)$. Now if τ satisfies property (1) above, then it satisfies that property in Λ_x^0 as well. This would imply that $y_2 \in X_{\lfloor n \rfloor}(\Lambda_x^0)$, a contradiction. So we must have that σ and τ are in different components of $\Gamma_{\geq n}(\Lambda_x^1)$. But since $y_2 \notin X_{\lfloor n \rfloor}(\Lambda_x^0)$, σ and τ must be in the same component of $\Gamma_{\geq n}(\Lambda_x^0)$. So let α be a simple path from σ to τ in $\Gamma_{\geq n}(\Lambda_x^0)$. We must have that α is not contained in $\Gamma_{\geq n}(\Lambda_x^1)$. Again let υ_1 be the vertex of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$ containing e_1, and let υ_2 be the vertex of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$ containing e_2. By the same reasoning as in Case 2(d)(ii), the only edges of α not contained in $\Gamma_{\geq n}(\Lambda_x^1)$ are those with one endpoint either e_1 or e_2, and where $\upsilon_1 = e_1$ or $\upsilon_2 = e_2$. But since every vertex of $\Gamma_{\geq n}(\Lambda_x^0)$ that is adjacent to υ_1 or υ_2 is also adjacent to σ_2^{R} in $\Gamma_{\geq n}(\Lambda_x^0)$, we may replace every such edge with one that has an endpoint of σ_2^{R}, resulting in a path from σ to τ in $\Gamma_{\geq n}(\Lambda_x^1)$. This is a contradiction.

Our assumption has led to a contradiction in every case, so we must have that $y_2 \notin X_{\lfloor n \rfloor}(\Lambda_x^1) \setminus X_{\lfloor n \rfloor}(\Lambda_x^0)$. Therefore, $X_{\lfloor n \rfloor}(\Lambda_x^1) \subseteq X_{\lfloor n \rfloor}(\Lambda_x^0)$.

Case 2 Step 1(d): Show that Λ_x^1 satisfies property (13) of the induction on k: that $F_{\lfloor n \rfloor}(\Lambda_x^1) \subseteq F_{\lfloor n \rfloor}(\Lambda_x^0)$.

Let $\tau \in F_{\lfloor n \rfloor}(\Lambda_x^1)$. Suppose by way of contradiction that $\tau = \sigma_2^{\mathrm{R}}$. In this case, we must have $I(y_1) = I(y_2) = \lfloor n \rfloor - 1$ so that e_1 and e_2 do not contain points of height greater than $\lfloor n \rfloor$. Then $I(y_{1.5}) = \lfloor n \rfloor$. So σ_2^{R} has exactly two 0-cells of height $\lfloor n \rfloor$, x and $y_{1.5}$. But $y_{1.5}$ is not on any vertex of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^1)$, contradicting the fact that $\tau \in F_{\lfloor n \rfloor}(\Lambda_x^1)$. So $\tau \neq \sigma_2^{\mathrm{R}}$.

As a result, $\tau \in \Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$. Now let $z_1, z_2 \in \tau^{(0)}$ and let υ_1 and υ_2 be vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^1)$ such that $I(z_1), I(z_2) = \lfloor n \rfloor$, $z_1 \in \upsilon_1$, and $z_2 \in \upsilon_2$. If υ_1 and υ_2 are both vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^0)$, then $\tau \in F_{\lfloor n \rfloor}(\Lambda_x^0)$. Otherwise, assume without loss of generality that $\upsilon_1 = \sigma_2^{\mathrm{R}}$. Then $z_1 \in \{x, y_1, y_2\}$. Since $I(z_1) = \lfloor n \rfloor$ and $I(y_1) = \lfloor n \rfloor - 1$, we know that $z_1 \neq y_1$. If $z_1 = x$, then since $x \in X_{\lfloor n \rfloor}(\Lambda_x^0)$, there is some vertex υ_1' of $\Gamma_{\geq n}(\Lambda_x^0)$ containing x. If instead $z_1 = y_2$, then $I(y_2) = \lfloor n \rfloor$, so $\max(I(e_2)) = \lfloor n \rfloor$. So let υ_1' be the vertex of $\Gamma_{\geq n}(\Lambda_x^0)$ containing e_2. In both cases, υ_1' contains z_1, and since $\upsilon_1' \in \Gamma_{\geq n}(\Lambda_x^0)$, we have $\tau \neq \upsilon_1'$. If we also have $\upsilon_2 = \sigma_2^{\mathrm{R}}$, we can choose υ_2', a vertex of $\Gamma_{\geq n}(\Lambda_x^0)$, in the same way such that $z_2 \in \upsilon_2'$ and $\upsilon_2' \neq \tau$. Then υ_1' and υ_2 or υ_2' demonstrate that $\tau \in F_{\lfloor n \rfloor}(\Lambda_x^0)$. Therefore, $F_{\lfloor n \rfloor}(\Lambda_x^1) \subseteq F_{\lfloor n \rfloor}(\Lambda_x^0)$.

Case 2 Step 2: For $k \geq 2$, contruct Λ_x^k from Λ_x^{k-1} such that it satisfies properties (11) - (15).

Case 2 Step 2(a): For $k \geq 2$, construct Λ_x^k from Λ_x^{k-1} such that it satisfies property (15) of the induction on k: that Λ_x^k has a 2-cell $\sigma_{k+1}^{\mathrm{R}}$ adjacent to $\overrightarrow{e_{k+1}}$ on its right side.

Let $k \geq 2$ and assume we have constructed Λ_x^{k-1}. We know that $\overrightarrow{e_k}$ is adjacent to a 2-cell σ_k^{R} on its right. If $\overrightarrow{e_k}$ is adjacent to a 2-cell on its left, name this 2-cell σ_k^{L}. Otherwise, $\overrightarrow{e_k}$ is adjacent to $\mathbb{R}^2 \setminus \Lambda_x^{k-1}$ on its left and $\overrightarrow{e_{k+1}}$ is adjacent to $\mathbb{R}^2 \setminus \Lambda_x^{k-1}$ on its right. So let σ_k^{L} be a 2-cell with boundary circuit labeled by $a_{k+1}^{-1} a_k a_k^{-1} a_{k+1}$, and glue σ_k^{L} to Λ_x^{k-1} by gluing the initial path along its boundary circuit labeled by $a_{k+1}^{-1} a_k$ along the path $\overleftarrow{e_{k+1}} \cdot \overrightarrow{e_k}$. Let $y_{k+\frac{1}{2}}$ be the 0-cell in σ_2^{L} replacing x on the boundary of the diagram, and let $\widetilde{\Lambda_x^{k-1}}$ be the resulting diagram.

We may use essentially the same argument used above in the construction of Λ_x^1

to show that $\widetilde{\Lambda_x^{k-1}}$ is a van Kampen diagram for w with respect to $\mathcal{P}'$ satisfying the following properties:

- For all $i \in \{0, 0.5, \ldots, \lfloor n \rceil\}$, $\Gamma_{\geq i}(\widetilde{\Lambda_x^{k-1}})$ is connected.
- $\mathrm{IDiam}(\widetilde{\Lambda_x^{k-1}}) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rceil + 1 + \left\lfloor \frac{1}{4} \max(\overline{\rho}, 4) \right\rfloor)$.
- $F_{\lfloor n \rceil}(\widetilde{\Lambda_x^{k-1}}) \subseteq F_{\lfloor n \rceil}(\Lambda_x^{k-1})$ and $X_{\lfloor n \rceil}(\widetilde{\Lambda_x^{k-1}}) \subseteq X_{\lfloor n \rceil}(\Lambda_x^{k-1})$.
- σ_k^{L} is adjacent to $\overrightarrow{e_k}$ on its left side and $\overrightarrow{e_{k+1}}$ (or $\overrightarrow{e_1}$, if $k = N_x$) on its right side.

The only difference between these two arguments is that, for Λ_x^1, we know that $I(y_1) = \lfloor n \rceil - 1$. In Case 2 Step 1(c), when determining which 0-cells can be in $X_{\lfloor n \rceil}(\Lambda_x^1)$, this fact allows us to immediately conclude that y_1 is not in $X_{\lfloor n \rceil}(\Lambda_x^1)$ without going through all the same work required to show that y_2 is not in $X_{\lfloor n \rceil}(\Lambda_x^1)$. However, in the $k \geq 2$ case, we do not know the height of either y_k or y_{k+1}, meaning that we must use the same, longer argument for y_k and y_{k+1} as we did for y_2. Otherwise, the arguments are identical.

Now construct Λ_x^k from $\widetilde{\Lambda_x^{k-1}}$ in the following way. (See Figure 3.10 for an example.) If $I(y_k) \geq \lfloor n \rceil$ or if $x \notin X_{\lfloor n \rceil}(\widetilde{\Lambda_x^{k-1}})$, then let $\sigma_{k+1}^{\mathrm{R}} = \sigma_k^{\mathrm{L}}$ and let $\Lambda_x^k = \widetilde{\Lambda_x^{k-1}}$. Otherwise, $I(y_k) = \lfloor n \rceil - 1$ and $x \in X_{\lfloor n \rceil}(\widetilde{\Lambda_x^{k-1}})$. For $\mathrm{S} \in \{\mathrm{R}, \mathrm{L}\}$, let γ_{S} be the directed edge path such that $\overrightarrow{e_k} \cdot \gamma_{\mathrm{S}}$ is the boundary circuit of σ_k^{S}. Then replace $\overrightarrow{e_k}$ with a directed edge path $\overrightarrow{e_k}'$ labeled by the word $a_k a_k^{-1} a_k$. Let x_k be the 0-cell adjacent to y_k on $\overrightarrow{e_k}'$, and let y_k' be the 0-cell adjacent to x on $\overrightarrow{e_k}'$. For $\mathrm{S} \in \{\mathrm{R}, \mathrm{L}\}$, this replaces σ_k^{S} with a new 2-cell $\sigma_k^{\mathrm{S}\prime}$ with boundary circuit $\overrightarrow{e_k}' \cdot \gamma_{\mathrm{S}}$. Let $w_{\mathrm{S}} \in (A \cup A^{-1})^*$ be the word that labels γ_{S}. Then the boundary of $\sigma_k^{\mathrm{S}\prime}$ is labeled by $a_k a_k^{-1} a_k w_{\mathrm{S}} =_G a_k w_{\mathrm{S}} =_G 1$, since $a_k w_{\mathrm{S}}$ labels the boundary circuit of σ_k^{S}.

Now if $\ell(a_k a_k^{-1} a_k w_S) \leq \max(\overline{\rho}, 4)$, then $a_k a_k^{-1} a_k w_S \in R'$. Otherwise, we have that $|\gamma_S| = \ell(w_S) \geq 2$, and we will split σ_k^S into two 2-cells whose boundary circuits are in R', as follows.

If there is some 0-cell other than x in γ_S with height at least $\lfloor n \rfloor$, then because height differs by at most 1 between adjacent 0-cells, we can choose a 0-cell $z_S \neq x, y_k$ on γ_S such that $I(z_S) = \lfloor n \rfloor$. Otherwise, every 0-cell on γ_S other than x has height less than or equal to $\lfloor n \rfloor - 1$. In particular, since $I(x) = \lfloor n \rfloor$, we know that the the next 0-cell along γ_S, which is adjacent to x, has height $\lfloor n \rfloor - 1$. Therefore, we can choose a 0-cell $z_S \neq x, y_k$ on γ_S such that $I(z_S) = \lfloor n \rfloor - 1$. Let α_S be the path from y_k to z_S and β_S the path from x to z_S such that $\alpha_S \cdot \overline{\beta_S} = \gamma_S$. If $|\alpha_S| \leq |\beta_S|$, then let u_S be the word labeling α_S and let $s_S = x_k$. Otherwise, let u_S be the word labeling β_S and let $s_S = y_k'$. Then add a directed edge path δ_S through the interior of $\sigma_k^{S\prime}$ from s_S to z_S labeled by u_S. This splits $\sigma_k^{S\prime}$ into two 2-cells, one with boundary circuit labeled by the word $a_k u_S u_S^{-1} a_k^{-1}$ and the other labeled by $a_k w_S$. Let $\sigma_k^S(x)$ be the one containing x, and $\sigma_k^S(y_k)$ the one containing y_k.

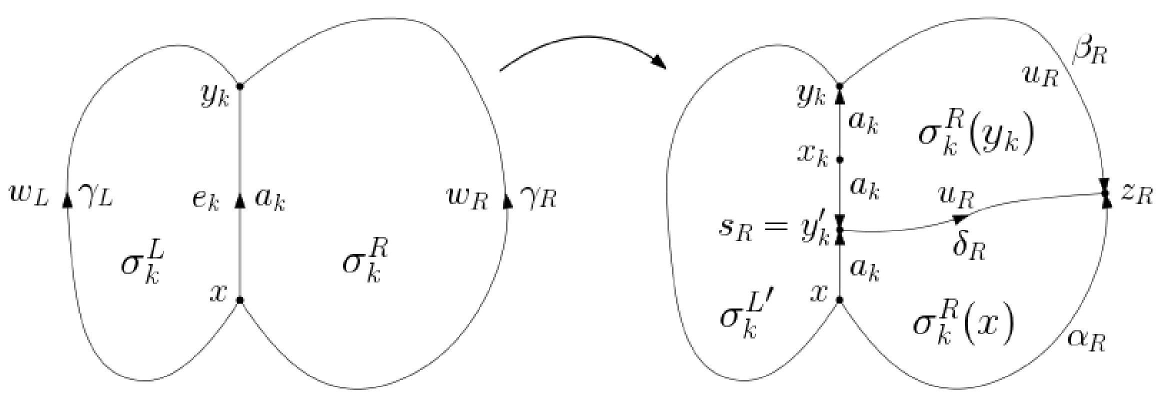

Figure 3.10: An example of the changes made to $\widetilde{\Lambda_x^{k-1}}$ to construct Λ_x^k.

We already know that $a_k w_S \in R'$. Of course, $a_k u_S u_S^{-1} a^{-1} =_G 1$. Since $\ell(u_S) =$

$\min(|\alpha_{\mathrm{S}}|, |\beta_{\mathrm{S}}|)$, we have that

$$\begin{aligned}
\ell(a_k u_{\mathrm{S}} u_{\mathrm{S}}^{-1} a_k^{-1}) &= 2\ell(u_{\mathrm{S}}) + 2 \\
&\leq |\alpha_{\mathrm{S}}| + |\beta_{\mathrm{S}}| + 2 \\
&= \ell(a_k w_{\mathrm{S}}) + 1 \\
&\leq \max(\overline{\rho}, 4) + 1
\end{aligned}$$

Now, since $\overline{\rho}$ is even, $\max(\overline{\rho}, 4) + 1$ is odd. Note that $\ell(a_k u_{\mathrm{S}} u_{\mathrm{S}}^{-1} a_k^{-1})$ is even, so we can therefore strengthen the inequality to $\ell(a_k u_{\mathrm{S}} u_{\mathrm{S}}^{-1} a_k^{-1}) \leq \max(\overline{\rho}, 4)$. Therefore, $a_k u_{\mathrm{S}} u_{\mathrm{S}}^{-1} a_k^{-1} \in R'$.

Having made these changes to $\widetilde{\Lambda_x^{k-1}}$, we will call the resulting diagram Λ_x^k. We have shown that the boundary circuit of each new 2-cell of Λ_x^k is labeled by a word in R', so Λ_x^k is a van Kampen diagram with respect to $\mathcal{P}'$. We have also not added or removed any 1-cells on the boundary, so Λ_x^k is a van Kampen diagram for the word w, the same as $\widetilde{\Lambda_x^{k-1}}$.

If $k < N_x$, let $\sigma_{k+1}^{\mathrm{R}}$ be the one of $\sigma_k^{\mathrm{L}\prime}$ and $\sigma_k^{\mathrm{L}}(x)$ that is a 2-cell of Λ_x^k. Then note that $\overrightarrow{e_{k+1}}$ is adjacent to $\sigma_{k+1}^{\mathrm{R}}$ on its right.

Case 2 Step 2(b): Show that Λ_x^k satisfies property (12) from the induction on k: that $\mathrm{IDiam}(\Lambda_x^k) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rfloor + 1 + \left\lfloor \frac{1}{4}\max(\overline{\rho}, 4) \right\rfloor)$.

We will first show that every $p \in \widetilde{\Lambda_x^{k-1}}^{(0)}$ has the same height in Λ_x^k as in $\widetilde{\Lambda_x^{k-1}}$. Let γ be an edge path from the basepoint $*$ to p in $\Lambda_x^{k(1)}$. Suppose that γ is not contained in $\widetilde{\Lambda_x^{k-1}}^{(1)}$. Then γ must at some point enter, and at some point exit, $\Lambda_x^{k(1)} \setminus \widetilde{\Lambda_x^{k-1}}^{(1)}$ through one of x, y_k, z_{L}, and z_{R}, since these are the only 0-cells of $\widetilde{\Lambda_x^{k-1}}$ on 1-cells not in $\widetilde{\Lambda_x^{k-1}}$. So let q_{F} be the first point on γ in $\{x, y_k, z_{\mathrm{L}}, z_{\mathrm{R}}\}$, and let q_{L} be the last point on γ in $\{x, y_k, z_{\mathrm{L}}, z_{\mathrm{R}}\}$. Recall that these points have height either $\lfloor n \rfloor - 1$ or $\lfloor n \rfloor$ in $\widetilde{\Lambda_x^{k-1}}$. Let α be the path from $*$ to q_{F}, β the path from q_{F} to q_{L}, and ω the path from q_{L} to p, such that $\alpha \cdot \beta \cdot \omega = \gamma$. Note that, since γ contains a 1-cell not in $\widetilde{\Lambda_x^{k-1}}^{(1)}$, we must have that $|\beta| \geq 1$. So $|\alpha \cdot \beta| \geq |\alpha| + 1$. But $|\alpha| \geq \lfloor n \rfloor - 1$, since α is a path from $*$ to q_{F} in $\widetilde{\Lambda_x^{k-1}}^{(1)}$. So $|\alpha \cdot \beta| \geq \lfloor n \rfloor$. Therefore, there is some path δ from $*$ to q_L in $\widetilde{\Lambda_x^{k-1}}^{(1)}$ with $|\delta| \leq |\alpha \cdot \beta|$. Then $\delta \cdot \omega$ is a path from $*$ to p in $\widetilde{\Lambda_x^{k-1}}^{(1)}$ with $|\delta \cdot \omega| \leq |\gamma|$. So the height of p has not decreased.

Now let γ be a geodesic from the basepoint $*$ to p in $\widetilde{\Lambda_x^{k-1}}$. Suppose that γ is not contained in Λ_x^k. The only 1-cell of $\widetilde{\Lambda_x^{k-1}}$ not in Λ_x^k is e_k, so e_k is the only edge of γ not in Λ_x^k. In particular, since γ is a geodesic, $\overrightarrow{e_k}$ appears in γ exactly once (since this is the direction of increasing height). So there are paths α and β in Λ_x^k such that $\gamma = \alpha \cdot \overrightarrow{e_k} \cdot \beta$ and $|\alpha| = I(y_k) = \lfloor n \rfloor - 1$. Now since $I(y_1) = \lfloor n \rfloor - 1$ in $\widetilde{\Lambda_x^{k-1}}$, there is a directed edge path δ from $*$ to y_1 in $\widetilde{\Lambda_x^{k-1}}$ with $|\delta| = \lfloor n \rfloor - 1$. Note also that δ cannot contain e_k, since $\max(I(e_k)) = \lfloor n \rfloor > \lfloor n \rfloor - 1$. So δ is also a path in Λ_x^k. Then $\delta \cdot \overrightarrow{e_1} \cdot \beta$ is a path from $*$ to p in Λ_x^k with the same length as γ. So the height of p in Λ_x^k has not increased, either.

In particular, we still have that $I(x), I(y_k), I(z_{\mathrm{R}}), I(z_{\mathrm{L}}) \in \{\lfloor n \rfloor - 1, \lfloor n \rfloor\}$. Now let $p \in \Lambda_x^{k(0)} \setminus \widetilde{\Lambda_x^{k-1}}^{(0)}$. Since every path from $*$ to p contains one of x, y_k, z_{R} and z_{L}, we have that $I(p) \geq \lfloor n \rfloor$. If p is in $\overrightarrow{e_k}'$, then it is adjacent to one of x and y_k, so $I(p) \in \{\lfloor n \rfloor, \lfloor n \rfloor + 1\}$. Otherwise, p is on an edge path of length at most $\frac{1}{2}\max(\overline{\rho}, 4)$ connecting two 0-cells of height at most $\lfloor n \rfloor$. Therefore, $I(p) \leq \lfloor n \rfloor + 1 + \left\lfloor \frac{1}{4}\max(\overline{\rho}, 4) \right\rfloor$. Since we know that $\mathrm{IDiam}(\widetilde{\Lambda_x^{k-1}}) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rfloor + 1 + \left\lfloor \frac{1}{4}\max(\overline{\rho}, 4) \right\rfloor)$, this implies that $\mathrm{IDiam}(\Lambda_x^k) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rfloor + 1 + \left\lfloor \frac{1}{4}\max(\overline{\rho}, 4) \right\rfloor)$.

Case 2 Step 2(c): Show that Λ_x^k satisfies property (11) from the induction on k: that for all $i \in \{0, 0.5, \ldots, \lfloor n \rfloor\}$, $\Gamma_{\geq i}(\Lambda_x^k)$ is connected.

Let $i \in \{0, 0.5, \ldots, \lfloor n \rfloor\}$. We know that $\Gamma_{\geq i}(\widetilde{\Lambda_x^{k-1}})$ is connected.

Let C be the largest subgraph of $\Gamma_{\geq i}(\Lambda_x^k)$ containing all the vertices of $\Gamma_{\geq i}(\Lambda_x^k)$ that are not in $\Gamma_{\geq i}(\widetilde{\Lambda_x^{k-1}})$. So the vertices of C are all of the following 2-cells that have been defined and exist in Λ_x^k: $\sigma_k^{\mathrm{R}\prime}$, $\sigma_k^{\mathrm{R}}(x)$, $\sigma_k^{\mathrm{R}}(y_k)$, $\sigma_k^{\mathrm{L}\prime}$, $\sigma_k^{\mathrm{L}}(x)$, and $\sigma_k^{\mathrm{L}}(y_k)$. First note that C is connected. If $\sigma_k^{\mathrm{R}\prime}$ is a 2-cell of Λ_x^k, then it has both x_k and y_k' on its boundary. Since $I(x), I(x_k') \geq i$, $\sigma_k^{\mathrm{R}\prime}$ is adjacent to either $\sigma_k^{\mathrm{L}\prime}$ or both of $\sigma_k^{\mathrm{L}}(x)$ and $\sigma_k^{\mathrm{L}}(y_k)$ in $\Gamma_{\geq i}(\Lambda_x^k)$, making C connected. Otherwise, $\sigma_k^{\mathrm{R}}(x)$ and $\sigma_k^{\mathrm{R}}(y_k)$ are both 2-cells of Λ_x^k. Since they both have s_{R} on their boundaries, they are adjacent to each other in $\Gamma_{\geq i}(\Lambda_x^k)$. Furthermore, one of these has both x_k and y_k' on its boundary. So by the same reasoning as above, C is connected.

Now let $V = \{\sigma_k^{\mathrm{R}}, \sigma_k^{\mathrm{L}}\}$, a subset of the vertices of $\Gamma_{\geq i}(\widetilde{\Lambda_x^{k-1}})$, and let $D = \Gamma_{\geq i}(\widetilde{\Lambda_x^{k-1}}) \setminus V$. Note that D is also $\Gamma_{\geq i}(\Lambda_x^k) \setminus C$. Let d be a vertex of D adjacent to some $v \in V$ in $\Gamma_{\geq i}(\widetilde{\Lambda_x^{k-1}})$. Then there is some point $q \in \partial d \cap \partial v$ with $I(q) \geq i$. Note that $q \notin \mathrm{int}(e)$, since $d \notin v$. Since $q \in \partial v \setminus \mathrm{int}(e)$, we have that $q \in \partial c$ for some $c \in C$. So d is adjacent to c in $\Gamma_{\geq i}(\Lambda_x^k)$. Then by Lemma 3.2.6, there is a subgraph of $\Gamma_{\geq i}(\Lambda_x^k)$ that is connected and contains all the vertices in C and D, i.e., all vertices of $\Gamma_{\geq i}(\Lambda_x^k)$. Therefore, $\Gamma_{\geq i}(\Lambda_x^k)$ is connected.

Case 2 Step 2(d): Show that Λ_x^k satisfies property (13) from the induction on k: that $F_{\lfloor n \rfloor}(\Lambda_x^k) \subseteq F_{\lfloor n \rfloor}(\Lambda_x^{k-1})$.

Since we have already shown that $F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}}) \subseteq F_{\lfloor n \rfloor}(\Lambda_x^{k-1})$, we need only show that $F_{\lfloor n \rfloor}(\Lambda_x^k) \subseteq F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. Let $\tau \in F_{\lfloor n \rfloor}(\Lambda_x^k)$. Suppose by way of contradiction that $\tau \notin \Gamma_{\geq \lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. Then $\tau \in \{\sigma_k^{\mathrm{R}\prime}, \sigma_k^{\mathrm{R}}(x), \sigma_k^{\mathrm{R}}(y_k), \sigma_k^{\mathrm{L}\prime}, \sigma_k^{\mathrm{L}}(x), \sigma_k^{\mathrm{L}}(y_k)\}$. For $\mathrm{S} \in \{\mathrm{R}, \mathrm{L}\}$, we cannot have $\tau = \sigma_k^{\mathrm{S}\prime}$ or $\tau = \sigma_k^{\mathrm{S}}(x)$, since both contain y_k', which has height $\lfloor n \rfloor + 1$. So for some $\mathrm{S} \in \{\mathrm{R}, \mathrm{L}\}$, $\tau = \sigma_k^{\mathrm{S}}(y_k)$. We must also have that $s_{\mathrm{S}} = x_k$, since otherwise $y_k' \in \sigma_k^{\mathrm{S}}(y_k)$. Now $I(x_k) = \lfloor n \rfloor$, but there must be some other 0-cell $z \in \sigma_k^{\mathrm{S}}(y_k)$ with $I(z) = \lfloor n \rfloor$. If $z \notin \delta_{\mathrm{S}}$, then $z \in \sigma_k^{\mathrm{S}} \setminus \{x\}$. Then by the way z_{S} was defined, we must have $I(z_{\mathrm{S}}) = \lfloor n \rfloor$. So there is some 0-cell z' on δ_{S} other than x_k with $I(z') = \lfloor n \rfloor$. Then every point between z' and x_k on δ_{S} has height at least n, which implies that $\sigma_k^{\mathrm{S}}(y_k) \in \Gamma_{\geq n}(\Lambda_x^k)$. This contradicts that $\sigma_k^{\mathrm{S}}(y_k) \in F_{\lfloor n \rfloor}(\Lambda_x^k)$.

So $\tau \in \Gamma_{\geq \lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. Now there are 0-cells $z_1, z_2 \in \tau$ with $I(z_1), I(z_2) = \lfloor n \rfloor$ and there are vertices $\upsilon_1, \upsilon_2 \neq \tau$ of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^k)$ with $z_1 \in \upsilon_1$ and $z_2 \in \upsilon_2$. If $\upsilon_1, \upsilon_2 \in \Gamma_{\geq \lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$, then $\tau \in F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. Otherwise, we may assume without loss of generality that $\upsilon_1 \in \{\sigma_k^{\mathrm{S}\prime}, \sigma_k^{\mathrm{S}}(x), \sigma_k^{\mathrm{S}}(y_k)\}$ for some $\mathrm{S} \in \{\mathrm{R}, \mathrm{L}\}$. Since $\tau \in \Gamma_{\geq \lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$, this implies that $\partial\tau \cap \partial\upsilon_1 \subseteq \sigma_k^{\mathrm{S}} \setminus \mathrm{int}(e_k)$. Therefore, $z_2 \in \sigma_k^{\mathrm{S}}$. So let $\upsilon_1' = \sigma_k^{\mathrm{S}}$. If we also have $\upsilon_2 \notin \Gamma_{\geq \lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$, we may similarly choose a vertex υ_2' of $\Gamma_{\geq \lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$ with $z_2 \in \upsilon_2'$. Therefore, υ_1' and υ_2 or υ_2' demonstrate that $\tau \in F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$, implying that $F_{\lfloor n \rfloor}(\Lambda_x^k) \subseteq F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$.

Case 2 Step 2(e): Show that Λ_x^k satisfies property (14) from the induction on k: that if $X_{\lfloor n \rfloor}(\Lambda_x^k) \not\subseteq X_{\lfloor n \rfloor}(\Lambda_x^{k-1})$, then $F_{\lfloor n \rfloor}(\Lambda_x^k) \subsetneq F_{\lfloor n \rfloor}(\Lambda_x^{k-1})$.

We have already shown in Case 2 Step 2(a) that $F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}}) \subseteq F_{\lfloor n \rfloor}(\Lambda_x^{k-1})$ and $X_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}}) \subseteq X_{\lfloor n \rfloor}(\Lambda_x^{k-1})$. So it is sufficient to show that if $X_{\lfloor n \rfloor}(\Lambda_x^k) \not\subseteq X_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$, then $F_{\lfloor n \rfloor}(\Lambda_x^k) \subsetneq F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$.

If $X_{\lfloor n \rfloor}(\Lambda_x^k) \not\subseteq X_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$, then there is some $p \in X_{\lfloor n \rfloor}(\Lambda_x^k) \setminus X_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. So by the definition of $X_{\lfloor n \rfloor}$, there are distinct vertices of $\Gamma_{\geq \lfloor n \rfloor}(\Lambda_x^k)$, σ and τ, both containing p such that σ is a vertex of $\Gamma_{\geq n}(\Lambda_x^k)$ and one of properties (1) or (2) from the definition of $X_{\lfloor n \rfloor}$ is true:

(1) τ is not a vertex of $\Gamma_{\geq n}(\Lambda_x^k)$ and there is a 0-cell $q \in \tau$ with $I(q) = \lfloor n \rfloor$ and $q \neq p$, or

(2) σ and τ are in different components of $\Gamma_{\geq n}(\Lambda_x^k)$.

Then there are several possibilities for how p could fail to be a member of $X_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$.

Suppose by way of contradiction that neither σ nor τ are in $\widetilde{\Lambda_x^{k-1}}$. In the proof that $F_{\lfloor n \rfloor}(\Lambda_x^k) \subseteq F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$, we showed that any element of $F_{\lfloor n \rfloor}(\Lambda_x^k)$ is in $\widetilde{\Lambda_x^{k-1}}$, but the same reasoning implies that τ cannot satisfy property (1) above. So we must have that σ and τ are in different components of $\Gamma_{\geq n}(\Lambda_x^k)$. However, I claim that this is impossible. Every 2-cell containing y'_k is in the same component of $\Gamma_{\geq n}(\Lambda_x^k)$, so the only new 2-cells that might be in a different component are $\sigma_k^{\mathrm{S}}(y_k)$ for $\mathrm{S} \in \{\mathrm{R}, \mathrm{L}\}$. But we have shown above that either $\sigma_k^{\mathrm{S}}(y_k)$ does not contain a point of height at least n, or it contains such a point on δ_{S}, making $\sigma_k^{\mathrm{S}}(y_k)$ and $\sigma_k^{\mathrm{S}}(x)$ adjacent in $\Gamma_{\geq n}(\Lambda_x^k)$. So all the vertices of $\Gamma_{\geq n}(\Lambda_x^k)$ that are not in $\widetilde{\Lambda_x^{k-1}}$ are in the same component, a contradiction.

Now suppose by way of contradiction that both σ and τ are in $\widetilde{\Lambda_x^{k-1}}$. Now if τ satisfies property (1) above, then it would also satisfy property (1) in $\widetilde{\Lambda_x^{k-1}}$, implying that $p \in X_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. So we must instead have that σ and τ are in different components of $\Gamma_{\geq n}(\Lambda_x^k)$, but the same component of $\Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$. Let $V = \{\sigma_k^{\mathrm{R}}, \sigma_k^{\mathrm{R}}\}$, let D be the largest subgraph containing all other vertices in the component of $\Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$ containing σ and τ, and let $C = \Gamma_{\geq n}(\Lambda_x^k) \setminus \Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$. We have shown above that C is connected. For any vertices $d \in D$ and $v \in V$ with d and v adjacent in $\Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$, there is a point $q \in \partial d \cap \partial v$. Since $d \notin V$, $q \notin \mathrm{int}(e_k)$. Therefore, there is some vertex $c \in C$ with $q \in C$, making d and c adjacent in $\Gamma_{\geq n}(\Lambda_x^k)$. So by Lemma 3.2.6, we have that σ and τ are in the same component of $\Gamma_{\geq n}(\Lambda_x^k)$, a contradiction.

Therefore, we must have that exactly one of σ and τ is in $\widetilde{\Lambda_x^{k-1}}$. Since we know that every vertex of $\Gamma_{\geq n}(\Lambda_x^k)$ that is not in $\widetilde{\Lambda_x^{k-1}}$ cannot be a vertex satisfying property (1) above, and since property (2) is symmetric between σ and τ, we may assume without loss of generality that τ is in $\widetilde{\Lambda_x^{k-1}}$ and σ is not. Then since $p \in \sigma \cap \tau$, we must have that $p \in \partial\sigma_k^{\mathrm{S}}$ for some $\mathrm{S} \in \{\mathrm{R}, \mathrm{L}\}$. We will show that $\sigma_k^{\mathrm{S}} \in F_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$.

By the same argument as above, we cannot have that σ and τ are in different components of $\Gamma_{\geq n}(\Lambda_x^k)$, but σ_k^{S} and τ are in the same component of $\Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$. We also know that $\tau \in \Gamma_{\geq \lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. So if $\sigma_k^{\mathrm{S}} \in \Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$, then σ_k^{S} and τ would satisfy the properties needed to put $p \in X_{\lfloor n \rfloor}(\widetilde{\Lambda_x^{k-1}})$. So we must have that $\sigma_k^{\mathrm{S}} \notin \Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$, i.e., that $\max(I(\sigma_k^{\mathrm{S}})) = \lfloor n \rfloor$.

Now one of σ_2^{R}, $\sigma_2^{\mathrm{R}\prime}$, or $\sigma_2^{\mathrm{R}}(x)$ is a vertex of $\Gamma_{\geq\lfloor n\rfloor}(\widetilde{\Lambda_x^{k-1}})$ containing x. If this vertex contains a point of height at least n, name it υ_1. Otherwise, we must have that $I(y_2) = \lfloor n\rfloor - 1$ and $k = 2$. Then the fact that $x \in X_{\lfloor n\rfloor}(\Lambda_x^0)$ implies that there is some vertex υ_1' of $\Gamma_{\geq n}(\Lambda_x^0)$ containing x, which must also be a vertex of $\Gamma_{\geq n}(\Lambda_x^1)$. If υ_1' is not a vertex of $\Gamma_{\geq n}(\Lambda_x^{1\prime})$, then it must be contained in a vertex υ_1 of $\Gamma_{\geq n}(\Lambda_x^{1\prime})$. Otherwise, let $\upsilon_1 = \upsilon_1'$. So υ_1 is a vertex of $\Gamma_{\geq n}(\widetilde{\Lambda_x^{k-1}})$ containing x. Since υ_1 contains a point of height at least n, $\upsilon_1 \neq \sigma_k^{\mathrm{S}}$.

Now we cannot have $p = x$, since $x \in X_{\lfloor n\rfloor}(\widetilde{\Lambda_x^{k-1}})$. So x and p are distinct 0-cells of σ_k^{S} with $I(x) = I(p) = \lfloor n\rfloor$ and there are vertices $\upsilon_1, \tau \neq \sigma_k^{\mathrm{S}}$ of $\Gamma_{\geq\lfloor n\rfloor}(\widetilde{\Lambda_x^{k-1}})$ with $x \in \upsilon_1$ and $p \in \tau$. Therefore, $\sigma_k^{\mathrm{S}} \in F_{\lfloor n\rfloor}(\widetilde{\Lambda_x^{k-1}})$. Since we know that $F_{\lfloor n\rfloor}(\Lambda_x^k) \subseteq F_{\lfloor n\rfloor}(\widetilde{\Lambda_x^{k-1}})$, and σ_k^{S} is not in Λ_x^k, this implies that $F_{\lfloor n\rfloor}(\Lambda_x^k) \subsetneq F_{\lfloor n\rfloor}(\widetilde{\Lambda_x^{k-1}})$. This completes the induction on k.

The result is a van Kampen diagram $\Lambda_x^{N_x}$ for w with respect to $\mathcal{P}'$ with the following properties:

- For all $i \in \{0, 0.5, \ldots, \lfloor n\rfloor\}$, $\Gamma_{\geq i}(\Lambda_x^{N_x})$ is connected (property (8) from the induction on j).
- $\mathrm{IDiam}(\Lambda_x^{N_x}) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n\rfloor + 1 + \lfloor \frac{1}{4}\max(\overline{\rho}, 4)\rfloor)$ (property (9) from the induction on j).
- For all $k \in \{0, \ldots, N_x - 1\}$, we have $F_{\lfloor n\rfloor}(\Lambda_x^{N_x}) \subseteq F_{\lfloor n\rfloor}(\Lambda_x^k)$ and if $X_{\lfloor n\rfloor}(\Lambda_x^{N_x}) \nsubseteq X_{\lfloor n\rfloor}(\Lambda_x^k)$, then $F_{\lfloor n\rfloor}(\Lambda_x^{N_x}) \subsetneq F_{\lfloor n\rfloor}(\Lambda_x^k)$
- Every 1-cell containing x has a 2-cell adjacent to it on both sides.

So let $\Delta_{\lfloor n\rfloor}^j = \Lambda_x^{N_x}$.

Case 2 Step 3: Show that $\Delta^j_{\lfloor n \rfloor}$ satisfies property (10) from the induction on j: that

$$(|F_{\lfloor n \rfloor}(\Lambda_x^{N_x})|, |X_{\lfloor n \rfloor}(\Lambda_x^{N_x})|) < (|F_{\lfloor n \rfloor}(\Lambda_x^0)|, |X_{\lfloor n \rfloor}(\Lambda_x^0)|).$$

Case 2(f): Suppose that in $\Lambda_x^{N_x}$, every 1-cell e containing x other than e_1 has $\max(I(e)) \geq n$. Then every 2-cell σ containing x is a vertex of $\Gamma_{\geq n}(\Lambda_x^{N_x})$. Since every such 1-cell e has a 2-cell adjacent to it on both sides, all the 2-cells containing x are in the same component of $\Gamma_{\geq n}(\Lambda_x^{N_x})$. Therefore, $x \notin X_{\lfloor n \rfloor}(\Lambda_x^{N_x})$.

Case 2(g): Suppose that there is some $k \in [N_x] \setminus \{1\}$ such that e_k is in $\Lambda_x^{N_x}$ and $I(y_k) = \lfloor n \rfloor - 1$. So we must have that $x \notin X_{\lfloor n \rfloor}(\Lambda_x^k)$. Then for each $i > k$, assume by induction that $x \notin X_{\lfloor n \rfloor}(\Lambda_x^{i-1})$. We have shown that $X_{\lfloor n \rfloor}(\Lambda_x^{i-1\prime}) \subseteq X_{\lfloor n \rfloor}(\Lambda_x^{i-1})$, so $x \notin X_{\lfloor n \rfloor}(\Lambda_x^{i-1\prime})$. But then we have $X_{\lfloor n \rfloor}(\Lambda_x^i) = X_{\lfloor n \rfloor}(\Lambda_x^{i-1\prime})$, so $x \notin X_{\lfloor n \rfloor}(\Lambda_x^i)$. Therefore, $x \notin X_{\lfloor n \rfloor}(\Lambda_x^{N_x})$.

So in either case $x \notin X_{\lfloor n \rfloor}(\Lambda_x^{N_x})$, but we know that $x \in X_{\lfloor n \rfloor}(\Lambda_x^0)$. Hence, either $X_{\lfloor n \rfloor}(\Lambda_x^{N_x}) \not\subset X_{\lfloor n \rfloor}(\Lambda_x^0)$, or else $X_{\lfloor n \rfloor}(\Lambda_x^{N_x}) \subsetneq X_{\lfloor n \rfloor}(\Lambda_x^0)$, which implies that $F_{\lfloor n \rfloor}(\Lambda_x^{N_x}) \not\subset F_{\lfloor n \rfloor}(\Lambda_x^0)$. So we have that either $|X_{\lfloor n \rfloor}(\Lambda_x^{N_x})| < |X_{\lfloor n \rfloor}(\Lambda_x^0)|$ or $|F_{\lfloor n \rfloor}(\Lambda_x^{N_x})| < |F_{\lfloor n \rfloor}(\Lambda_x^0)|$. Since we know that $|F_{\lfloor n \rfloor}(\Lambda_x^{N_x})| \leq |F_{\lfloor n \rfloor}(\Lambda_x^0)|$, this implies that

$$(|F_{\lfloor n \rfloor}(\Lambda_x^{N_x})|, |X_{\lfloor n \rfloor}(\Lambda_x^{N_x})|) < (|F_{\lfloor n \rfloor}(\Lambda_x^0)|, |X_{\lfloor n \rfloor}(\Lambda_x^0)|).$$

This completes the induction on j.

Case 2 Step 4: Define Δ_n and show that it satisfies properties (3) and (4) of the induction on n.

So we have a sequence $\Delta^0_{\lfloor n \rfloor}, \Delta^1_{\lfloor n \rfloor}, \ldots$ such that for $j \geq 1$, we have

$$(|F_{\lfloor n \rfloor}(\Delta^j_{\lfloor n \rfloor})|, |X_{\lfloor n \rfloor}(\Delta^j_{\lfloor n \rfloor})|) < (|F_{\lfloor n \rfloor}(\Delta^{j-1}_{\lfloor n \rfloor})|, |X_{\lfloor n \rfloor}(\Delta^{j-1}_{\lfloor n \rfloor})|).$$

Since the lexicographic order on $\mathbb{N}^2$ is a well-order, this sequence must end with a van Kampen diagram $\Delta_n := \Delta^m_{\lfloor n \rfloor}$ satisfying the following properties:

- For all $i \in \{0, 0.5, \ldots, \lfloor n \rfloor\}$, $\Gamma_{\geq i}(\Delta^m_{\lfloor n \rfloor})$ is connected (property (3) from the induction on n, but excluding $i = n$).
- $\mathrm{IDiam}(\Delta^m_{\lfloor n \rfloor}) \leq \max(\mathrm{IDiam}(\Delta_0), \lfloor n \rfloor + 1 + \left\lfloor \frac{1}{4} \max(\overline{\rho}, 4) \right\rfloor)$ (property (4) from the induction on n).
- $X_{\lfloor n \rfloor}(\Delta^m_{\lfloor n \rfloor}) = \emptyset$.

So we need only show that $\Gamma_{\geq n}(\Delta_n)$ is connected. Suppose by way of contradiction that σ and τ are vertices of $\Gamma_{\geq n}(\Delta_n)$ that are in different components. Let α be a path in $\Gamma_{\geq \lfloor n \rfloor}(\Delta_n)$ from σ to τ containing the fewest edges not in $\Gamma_{\geq n}(\Delta_n)$. Let β be the longest initial segment of α contained in $\Gamma_{\geq n}(\Delta_n)$, let σ' be the last vertex on β, and let τ' be the vertex after σ' on α. Then there is a point $x \in \partial\sigma' \cap \partial\tau'$ with $\widehat{I}(x) \geq \lfloor n \rfloor$ but $I(x) < n$. So x is a 0-cell with $I(x) = \lfloor n \rfloor$.

If $\tau' \notin \Gamma_{\geq n}(\Delta_n)$, then $\tau' \neq \tau$. So let υ be the vertex after τ' on α. If $x \in \upsilon$, we could replace the edges from σ' to τ' and τ' to υ with a single edge from σ' to υ. This would reduce the number of edges not in $\Gamma_{\geq n}(\Delta_n)$, contradicting the way α was chosen. So we must have $x \notin \upsilon$. Since τ' and υ are adjacent in $\Gamma_{\geq \lfloor n \rfloor}(\Delta_n)$, there must be some $y \neq x$ with $y \in \partial\tau' \cap \partial\upsilon$ and $I(y) \geq \lfloor n \rfloor$. Since $\tau' \notin \Gamma_{\geq n}(\Delta_n)$, y must be a 0-cell with $I(y) = \lfloor n \rfloor$. So τ' and σ' satisfy the properties required to make $x \in X_{\lfloor n \rfloor}(\Delta_n)$. This contradicts that $X_{\lfloor n \rfloor}(\Delta_n) = \emptyset$.

So we must have that $\tau' \in \Gamma_{\geq n}(\Delta_n)$. If τ' and σ' are in different components of $\Gamma_{\geq n}(\Delta_n)$, then we would again have $x \in X_{\lfloor n \rfloor}(\Delta_n)$. So there must be some path δ from σ' to τ' in $\Gamma_{\geq n}(\Delta_n)$. But then we can replace the edge on α from σ' to τ' with δ to acquire a path from σ to τ with fewer edges not in $\Gamma_{\geq n}(\Delta_n)$. This contradicts the way that α was chosen. Therefore, no such vertices σ and τ exist, implying that $\Gamma_{\geq n}(\Delta_n)$ is connected as desired. We have now finished Case 2.

This concludes the induction on n.

Step B: Define a van Kampen diagram Δ_w for the word w and show that there is an $s \in \mathbb{N}[\frac{1}{2}]$ such that $\mathrm{AV}_s(\Delta_w, I) \leq \mathrm{IDiam}(\Delta_0)$.

With the induction in Step A, we have constructed a van Kampen diagram $\Delta_w := \Delta_{\mathrm{IDiam}(\Delta_0)-\frac{1}{2}}$ for w with respect to $\mathcal{P}'$ with the following properties:

- For all $i \in \{0, 0.5, \ldots, \mathrm{IDiam}(\Delta_0) - \frac{1}{2}\}$, $\Gamma_{\geq i}(\Delta_w)$ is connected. By Proposition 3.2.5, this implies that the set $(\Delta_w)_{\geq i} = \{x \in \Delta_w | \widehat{I}(x) \geq i\}$ is path connected.

- $\mathrm{IDiam}(\Delta_w) \leq \mathrm{IDiam}(\Delta_0) + \left\lfloor \frac{1}{4} \max(\overline{\rho}, 4) \right\rfloor$.

Now let $s = 1 + \left\lfloor \frac{1}{4} \max(\overline{\rho}, 4) \right\rfloor$. Let $n \in \mathbb{N}_0$. Recall from the definition of AV_s that P_n is the union of the contour graphs of (Δ_w, I) that have height n and C_n^s is the set of path components of $\Delta_w \setminus P_n$ that do not contain $*$ and contain a point x with $|\widehat{I}(x) - n| \geq s$. Define $(\Delta_w)_{<n} := \{x \in \Delta_w | \widehat{I}(x) < n\}$ and note that $\Delta_w \setminus P_n = (\Delta_w)_{<n} \cup (\Delta_w)_{\geq n+\frac{1}{2}}$ (recall that the notation $\Delta_{\geq k}$ was defined in Proposition 3.2.5. We will first show that $(\Delta_w)_{<n}$ is path connected.

Let $x \in \Delta_w$ with $\widehat{I}(x) < n$. I claim that there is a path γ from x to the basepoint $*$ with $\max(\widehat{I}(\gamma([0,1]))) < n$. By the definition of $\widehat{I}$, if x is in the interior of a 2-cell σ, there is a point $y \in \partial\sigma$ with $\widehat{I}(y) \in \{\widehat{I}(x), \widehat{I}(x) - \frac{1}{2}\}$ and a path α from x to y with $\widehat{I}(\alpha) \subseteq \{\widehat{I}(x), \widehat{I}(x) - \frac{1}{2}\}$. Then $\max(\widehat{I}(\alpha)) < n$.

So we may reduce to the case where $x \in \Delta_w^{(1)}$. If x is in the interior of a 1-cell e, then $\widehat{I}(\mathrm{int}(e)) = \{\widehat{I}(x)\}$, and there is some 0-cell $y \in e$ with $I(y) \leq \widehat{I}(x)$. So the portion of e between x and y forms a path β from x to y with $\max(\widehat{I}(\beta)) = \widehat{I}(x) < n$.

Then we may reduce to the case where $x \in \Delta_w^{(1)}$. Then any shortest edge path γ from x to $*$ has $\max(\widehat{I}(\gamma)) = \widehat{I}(x) < n$. Therefore, x is in the same path component of $(\Delta_w)_{<n}$ as $*$. Since x was arbitrary, this means that there is only one path component of $(\Delta_w)_{<n}$.

Now suppose $n < \text{IDiam}(\Delta_0)$. Then we have shown—by Proposition 3.2.5 and property (1) of Step A—that the only path component of $\Delta_w \setminus P_n$ not containing $*$ is $(\Delta_w)_{\geq n+\frac{1}{2}}$. Therefore, $|C_n^s| = 1$ (where C_n^s is defined as in Definition 3.1.4).

Now suppose $n \geq \text{IDiam}(\Delta_0)$. We know that $(\Delta_w)_{<n}$ is the path component of $\Delta_w \setminus P_n$ containing $*$, so every other path component is a subset of $(\Delta_w)_{\geq n+\frac{1}{2}}$. Because $\text{IDiam}(\Delta_w) \leq \text{IDiam}(\Delta_0) + s - 1$, for any $x \in \Delta_w$, we have $\widehat{I}(x) \leq \text{IDiam}(\Delta_0) + s - \frac{1}{2}$. So for any point x in a path component of $\Delta_w \setminus P_n$ not containing $*$, we have $n < \widehat{I}(x) \leq \text{IDiam}(\Delta_0) + s - \frac{1}{2}$. Therefore, $|\widehat{I}(x) - n| < s$. This implies that $C_n^s = \emptyset$.

Therefore, $\text{AV}_s(\Delta_w, I) = \sum_{n=0}^{\text{IDiam}(\Delta_0)-1} |C_n^s| = \text{IDiam}(\Delta_0) \leq \text{IDiam}_{\mathcal{P}}(w)$. Since w was an arbitrary word representing the identity of G, this implies that for all $n \in N$, $\text{IAV}_s(n) \leq \text{IDiam}_{\mathcal{P}}(n)$. So $\text{IDiam}_{\mathcal{P}}$ is an IAV function for $\mathcal{P}'$.

$\square$

Remark 3.2.7. It is noteworthy that, by making $\Gamma_{\geq i}$ connected for every i less than the intrinsic diameter of the diagram, this proof essentially constructs a van Kampen diagram with only a single hill, which has height approximately equal to the intrinsic diameter of the diagram. In some sense, this is the simplest possible terrain we could ask for; knowing the intrinsic diameter of the diagram tells us almost everything there is to know about the terrain. This gives us additional structure to work with whenever we know the intrinsic diameter function of a presentation. This perspective also makes it seem very unlikely that there is a proper refinement of intrinsic diameter based on "hilliness", since we can always find van Kampen diagrams with only one hill.

3.3 A Notable Example: The Solvable Baumslag-Solitar Groups

Note that Theorem 3.2.3 is possible to prove only due to the power of adding relators to the presentation that are not cyclically reduced and using unreduced van Kampen diagrams, i.e., van Kampen diagrams containing a pair of 2-cells that share part of their boundaries and are mirror images of each other. If we instead restrict ourselves to using relators that are cyclically reduced and reduced van Kampen diagrams, it is plausible that we may obtain a proper refinement of IDiam. This conjecture is motivated by considering some standard van Kampen diagrams of the standard presentations for the solvable Baumslag-Solitar groups.

The Baumslag-Solitar group $\mathrm{BS}(m,n)$, where $m,n \in \mathbb{Z} \setminus \{0\}$, is defined by the presentation

$$\langle a,t | ta^m t^{-1} = a^n \rangle.$$

In 1962, Baumslag and Solitar showed that this family of groups contains groups that are non-Hopfian, disproving a conjecture that no such 1-relator groups existed [2]. Since then, these groups have proven to have a number properties of interest to combinatorial and geometric group theorists, often serving as interesting examples and counterexamples. For example, these groups are all asynchrously automatic [1], but not automatic when $|m| \neq |n|$, since they do not have a quadratic (or even polynomial) Dehn function [15]. The groups $\mathrm{BS}(1,n)$—the solvable Baumslag-Solitar groups—have rational growth like automatic groups, which Brazil proved in [4] using a regular language of geodesic normal forms for a subset of the elements. However, if $n > 1$, there is no regular language of geodesic normal forms for these groups [17]. These groups are also not almost convex [22], but $\mathrm{BS}(1,2)$ is minimally almost convex

[13].

Given the role Baumslag-Solitar groups have played as interesting examples and counterexamples, it is natural to consider what their IAV functions look like. In this section, I focus on the solvable Baumslag-Solitar groups—BS$(1, n)$ for $n \geq 2$.

These groups have a certain notable set of van Kampen diagrams, described as follows. Given $n \in \mathbb{N}$ with $n \geq 2$, let $\mathcal{P}_n$ be the standard presentation for BS$(1, n)$, i.e., $\mathcal{P}_n = \langle a, t | tat^{-1} = a^n \rangle$. For $m \in \mathbb{N}$, let $w_m = at^m at^{-m} a^{-1} t^m a^{-1} t^{-m}$. Then there is a standard van Kampen diagram D_m for w_m with respect to $\mathcal{P}_n$ defined inductively as follows. Let D_1 consist of two 2-cells, $\sigma_{\rm T}$ and $\sigma_{\rm B}$, that are glued together as shown in Figure 3.11. The basepoint $*$ of D_1 is the initial 0-cell on the boundary circuit of $\sigma_{\rm T}$. Having defined D_{m-1}, let D_m consist of n copies of D_{m-1} along with two 2-cells, $\sigma_{\rm T}$ and $\sigma_{\rm B}$, glued together as shown in Figure 3.11. The basepoint $*$ of D_m is the basepoint of the leftmost copy of D_{m-1}.

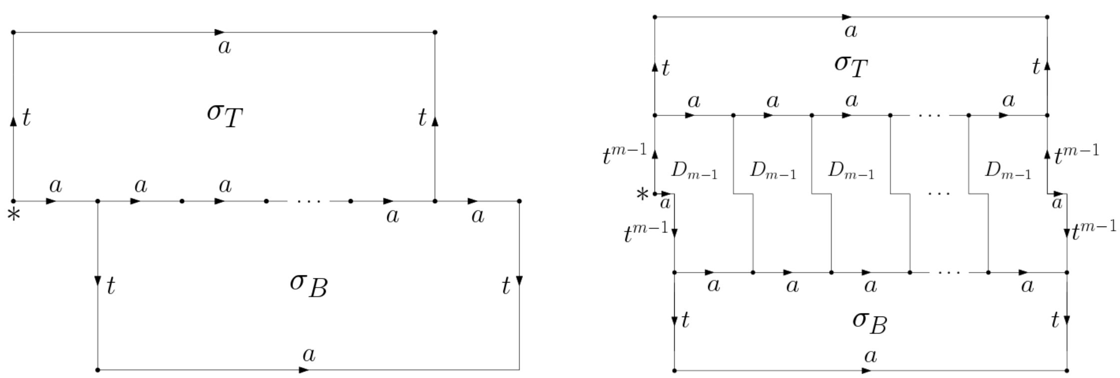

(a) D_1 (b) Inductive step

Figure 3.11: The inductive definition of D_m for $m \in \mathbb{N}$.

In these van Kampen diagrams, the number of 2-cells in D_m is exponential in m, while the length of w_m is linear in m. However, IDiam(BS$(1, n)$) is linear, and indeed the diameter of D_m is linear in m. This large difference between the area and diameter of these diagrams is notable, since, intuitively, trying to "squeeze" exponentially-increasing area into a ball in the Cayley complex with linearly-increasing radius might result in mountainous terrain in a van Kampen diagrams. This intuition motivates me to consider if there might be some refinement of IDiam that is not linear for $\mathrm{BS}(1,n)$, and helps to motivate the definition of IAV functions.

In order to help study IAV functions of $\mathrm{BS}(1,n)$, I wrote a Python program to draw D_m, find the height of each 0-cell, and draw contour graphs on top of the diagram. The code can be found at `https://github.com/andrew-quaisley/Contour-Drawing`.

Looking closely at the contour graphs of D_m in $\mathrm{BS}(1,2)$ for $m = 6, \ldots, 10$, patterns emerged to suggest that for any given sensitivity s, $IAV_s(D_m)$ grows exponentially with respect to m. For example, consider the contour graphs of D_8, shown in Figure 3.12. A "mountain range" runs through the middle of these diagrams, with many different peaks. Note that there is a repeating pattern to these peaks, with the same shapes appearing again and again at regular intervals along the diagram from left to right.

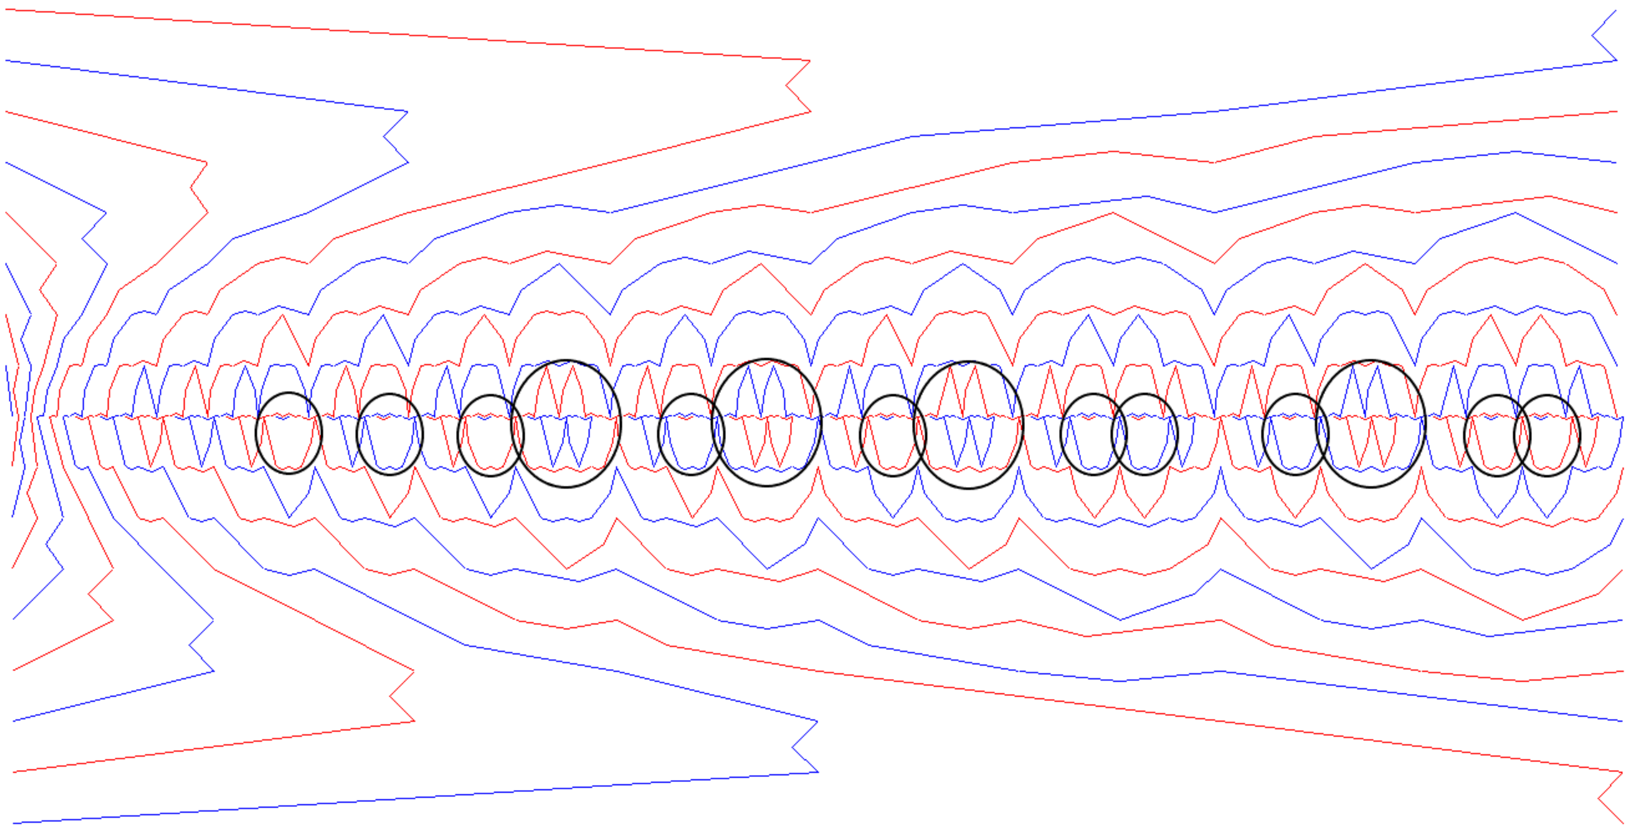

Figure 3.12: Just the contour graphs for D_8 in BS(1, 2), with ovals drawn around some of the components that contribute to $\mathrm{IAV}_1(D_8)$.

These patterns led to a proof of the following proposition:

Proposition 3.3.1. *For $n \in \mathbb{N}$ with $n \geq 2$ and $m \in \mathbb{N}$, let D_m be the van Kampen diagram with respect to the standard presentation for BS$(1, n)$ defined above. Then for all $s \in \mathbb{N}[\frac{1}{2}]$, $\mathrm{IAV}_s(D_m)$ grows at least exponentially as a function of m. In particular, for $m \in \mathbb{N}$, $\mathrm{IAV}_s(D_m) \geq n^{m-(2s+8)} - 2$.*

Before proving it, we will need a few more definitions, and some results about distance in $\mathrm{BS}(1, n)$. First, note that for each $i \in [m-1]$, D_m contains n^{m-i} copies of D_i glued end-to-end, starting at the basepoint. It will be useful to let D_i^j denote the jth such copy of D_i, counting out from the basepoint, as shown in Figure 3.13.

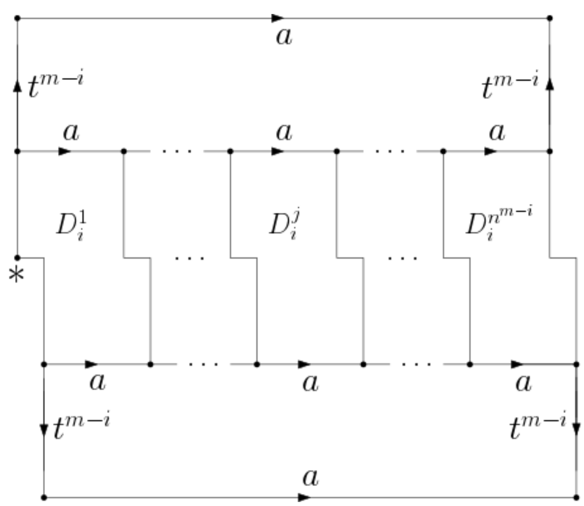

Figure 3.13: D_m, with $D_i^1, \ldots, D_i^{n^{m-i}}$ labeled.

Also note that any element $g \in \mathrm{BS}(1, n)$ can be written uniquely in the form $g = t^{-u}a^v t^w$, for some $u, v, w \in \mathbb{Z}$ where $u, w \geq 0$ and $n|v$ implies that $uw = 0$. We will call vn^{-u} the *a-length* of g and $w - u$ the *t-height* of g.

In [27], Taback and Walker build on [12] and [13] to describe a set of geodesics of the standard presentation of $\mathrm{BS}(1, n)$ based on the normal forms above. In fact, in [12], Elder descibed an algorithm to find the same geodesics. However, [12] focuses on finding a fast algorithm, so the definitions and lemmas of [27] are simpler to use for our purposes.

Taback and Walker first note that integer sequences $\mathbf{x} = (x_0, x_1, \ldots, x_{k_\mathbf{x}})$ such that $\Sigma(\mathbf{x}) := \sum_{j=0}^{k_\mathbf{x}} x_j n^j = v$ correspond to words for g in the following way. They define $\mathcal{L}_v$ to be the set of integer sequences $\mathbf{x}$ with $\Sigma(\mathbf{x}) = v$ and $x_{k_x} \neq 0$, and then define a function $\eta_{u,v,w} : \mathcal{L}_v \to \{a^{\pm 1}, t^{\pm 1}\}^*$ as follows:

$$\eta_{u,v,w}(\mathbf{x}) = \begin{cases} t^{-u} a^{x_0} t a^{x_1} \ldots t a^{x_{k_\mathbf{x}}} t^{w-k_\mathbf{x}} & \text{if } k_\mathbf{x} \leq w \text{ or } u \leq w < k_\mathbf{x} \\ t^{k_\mathbf{x}-u} a^{x_{k_\mathbf{x}}} t^{-1} a^{x_{k_\mathbf{x}-1}} t^{-1} \ldots t^{-1} a^{x_0} t^w & \text{if } w < k_\mathbf{x} \leq u \text{ or } w < u < k_\mathbf{x} \end{cases} \tag{3.3.1}$$

Then $\eta_{u,v,w}(\mathbf{x}) =_{\mathrm{BS}(1,n)} t^{-u} a^v t^w$.

We will be using a result from [13], rewritten in the notation above in [27, Lemma 3.9] as follows: given $\mathbf{x} \in \mathcal{L}_v$, the word $\eta_{u,v,w}(\mathbf{x})$ is a geodesic for the element g if $\ell(\eta_{u,v,w}(\mathbf{x})) \leq \ell(\eta_{u,v,w}(\mathbf{y}))$ for all $\mathbf{y} \in \mathcal{L}_v$. Note that $(v) \in \mathcal{L}_v$, which means $\mathcal{L}_v \neq \emptyset$, and therefore the lemma implies that g has a geodesic in $\eta_{u,v,w}(\mathcal{L}_v)$.

It will also be useful to note [27, Lemma 3.7]: given $\mathbf{x} \in \mathcal{L}_v$, the length of $\eta_{u,v,w}(\mathbf{x})$ is given by the formula

$$\ell(\eta_{u,v,w}(\mathbf{x})) = \begin{cases} ||\mathbf{x}||_1 + u + w & \text{if } k_\mathbf{x} \leq \max(u, w) \\ ||\mathbf{x}||_1 + 2k_\mathbf{x} - |u - w| & \text{otherwise} \end{cases}, \tag{3.3.2}$$

where $||\mathbf{x}||_1 = \sum_{j=0}^{k_\mathbf{x}} |x_j|$.

Note that the two cases give the same value when $k_\mathbf{x} = \max(u, w)$.

We are now ready to prove the proposition.

Proof of Proposition 3.3.1. Let $m \in \mathbb{N}$ and $s \in \mathbb{N}[\frac{1}{2}]$. If $m \leq 2s + 8$, then $n^{m-(2s+8)} - 2 < 0 \leq \mathrm{IAV}_s(D_m)$. Otherwise, for $i \in \{2, 3, \ldots, n^{m-(2s+8)}\}$, let $u_i^{\mathrm{T}} = t^{2s+8} a^i$ and

$u_i^{\mathrm{B}} = au_i^{\mathrm{T}}$. Note that both u_i^{T} and u_i^{B} label paths from $*$ to 0-cells p_i^{T} and p_i^{B} in D_m, because $m > 2s+8$, meaning that D_m contains $n^{m-(2s+8)}$ non-overlapping copies of D_{2s+8} glued together, side-by-side. Note also that $u_i^{\mathrm{B}} =_{\mathrm{BS}(1,n)} a^{in^{2s+8}+1} t^{2s+8}$.

Now let $\mathbf{x} \in \mathcal{L}_{in^{2s+8}+1}$ such that for all $\mathbf{y} \in \mathcal{L}_{in^{2s+8}+1}$, we have that $||\mathbf{x}||_1 + 2k_{\mathbf{x}} \leq ||\mathbf{y}||_1 + 2k_{\mathbf{y}}$ and, if $||\mathbf{x}||_1 + 2k_{\mathbf{x}} = ||\mathbf{y}||_1 + 2k_{\mathbf{y}}$, then $k_{\mathbf{x}} \geq k_{\mathbf{y}}$. I claim that $k_{\mathbf{x}} \geq 2s+8$. Suppose by way of contradiction that $k_{\mathbf{x}} < 2s+8$. We know that $\Sigma(\mathbf{x}) = in^{2s+8}+1 > 2n^{k_{\mathbf{x}}+1}$, so let $l \in \{0, \ldots, k_{\mathbf{x}}\}$ be the largest index with $\Sigma_l := \sum_{j=l}^{k_{\mathbf{x}}} x_j n^j \geq 2n^{k_{\mathbf{x}}+1}$. Then we have $\Sigma_{l+1} := \sum_{j=l+1}^{k_{\mathbf{x}}} x_j n^j < 2n^{k_{\mathbf{x}}+1}$. Since $\Sigma_l > \Sigma_{l+1}$, we must have that $x_l > 0$. Furthermore, since $l \leq k_{\mathbf{x}}$, $2n^{k_{\mathbf{x}}+1}$ and Σ_{l+1} are both multiples of n^l, and therefore so is $2n^{k_{\mathbf{x}}+1} - \Sigma_{l+1}$. Now $2n^{k_{\mathbf{x}}+1} - \Sigma_{l+1} \leq \Sigma_l - \Sigma_{l+1} = x_l n^l$. Therefore, there is some $x_l' \in \mathbb{N}$ with $x_l' < x_l$ such that $(x_l - x_l')n^l + \Sigma_{l+1} = 2n^{k_{\mathbf{x}}+1}$.

So let $\mathbf{x}' = (x_0, x_1, \ldots, x_{l-1}, x_l', 0, \ldots, 0, 2)$ (where $k_{\mathbf{x}'} = k_{\mathbf{x}} + 1$). Then $\mathbf{x}' \in \mathcal{L}_{in^{2s+8}+1}$ and $k_{\mathbf{x}} < k_{\mathbf{x}'}$. I claim that $||\mathbf{x}'||_1 + 2k_{\mathbf{x}'} \leq ||\mathbf{x}||_1 + 2k_{\mathbf{x}}$. Note that $2k_{\mathbf{x}'} - 2k_{\mathbf{x}} = 2$, so we need only show that $2 \leq ||\mathbf{x}||_1 - ||\mathbf{x}'||_1$. Now,

$$\begin{aligned}
2n^{k_{\mathbf{x}}+1} &= (x_l - x_l')n^l + \Sigma_{l+1} \\
&\leq (x_l - x_l')n^l + \sum_{j=l+1}^{k_{\mathbf{x}}} |x_j| n^j \\
&\leq (x_l - x_l')n^{k_{\mathbf{x}}} + \sum_{j=l+1}^{k_{\mathbf{x}}} |x_j| n^{k_{\mathbf{x}}}.
\end{aligned}$$

Then dividing both sides of this inequality by $n^{k_{\mathbf{x}}}$ gives us

$$\begin{aligned}2n &\leq x_l - x'_l + \sum_{j=l+1}^{k_{\mathbf{x}}} |x_j| \\ 2n-2 &\leq x_l - x'_l + \sum_{j=l+1}^{k_{\mathbf{x}}} |x_j| - 2 \\ &= ||\mathbf{x}||_1 - ||\mathbf{x}'||_1.\end{aligned}$$

Since $n \geq 2$, we have $2 \leq 2n-2$, and thus $2 \leq ||\mathbf{x}||_1 - ||\mathbf{x}'||_1$ as desired. This proves that $||\mathbf{x}'||_1 + 2k_{\mathbf{x}'} \leq ||\mathbf{x}||_1 + 2k_{\mathbf{x}}$, which contradicts the way that $\mathbf{x}$ was chosen since $k_{\mathbf{x}'} > k_{\mathbf{x}}$. Therefore, we must have that $k_{\mathbf{x}} \geq 2s+8$.

I also claim that $x_0 = 1$ and $x_1, \ldots, x_{2s+7} = 0$. First, note that since $\Sigma(\mathbf{x}) = in^{2s+8} + 1 \equiv 1 \bmod n$, we must have that $x_0 \equiv 1 \bmod n$. So there is some $c \in \mathbb{Z}$ with $x_0 = cn+1$. Let $\Sigma_< = \sum_{j=0}^{2s+7} x_j n^j$ and $\Sigma_> = \sum_{j=2s+8}^{k_{\mathbf{x}}} x_j n^j$. Now $\Sigma_< + \Sigma_> = \Sigma(\mathbf{x}) = in^{2s+8} + 1$, $\Sigma_>$ is divisible by n^{2s+8}, and $in^{2s+8}+1$ is congruent to 1 mod n^{2s+8}. Therefore, we must have that $\Sigma_< = dn^{2s+8} + 1$ for some $d \in \mathbb{Z}$. Suppose by way of contradiction that $d \neq 0$. Let $\mathbf{x}' = (1, 0, \ldots, 0, x_{2s+8} + d, x_{2s+9}, \ldots, x_{k_{\mathbf{x}}})$. Then $\mathbf{x}' \in \mathcal{L}_{in^{2s+8}+1}$. Note that $k_{\mathbf{x}} = k_{\mathbf{x}'}$. Then

$$||\mathbf{x}'||_1 - ||\mathbf{x}||_1 = 1 + |x_{2s+8} + d| - \sum_{j=0}^{2s+8} |x_j| \leq 1 + |d| - \sum_{j=0}^{2s+7} |x_j|.$$

Now

$$\begin{aligned} dn^{2s+8}+1 &= \sum_{j=0}^{2s+7} x_j n^j \\ dn^{2s+8} &= cn + \sum_{j=1}^{2s+7} x_j n^j \\ |d|n^{2s+8} &\le |c|n + \sum_{j=1}^{2s+7} |x_j| n^j \\ &\le |c|n + \sum_{j=1}^{2s+7} |x_j| n^{2s+7}. \end{aligned}$$

Then dividing both sides of the inequality by n^{2s+7} gives us

$$\begin{aligned} |d|n &\le |c|n^{-(2s+6)} + \sum_{j=1}^{2s+7} |x_j| \\ |d|n + |x_0| - |c|n^{-(2s+6)} &\le \sum_{j=0}^{2s+7} |x_j|. \end{aligned}$$

Since $n \geq 2$, $d \neq 0$, and $x_0 \neq 0$, this implies that

$$||\mathbf{x}'||_1 - ||\mathbf{x}||_1 \le 1 + |d| - |d|n - |x_0| + |c|n^{-(2s+6)} < 1 + |d|(1-n) - |x_0| + 1 \le 0.$$

Therefore, $||\mathbf{x}'||_1 < ||\mathbf{x}||_1$, which contradicts the way that $\mathbf{x}$ was chosen. Hence, we must have $d = 0$, meaning that $\Sigma_< = 1$ and that $\mathbf{x}'$ matches $\mathbf{x}$ from the $2s+8$th entry on. So if $||\mathbf{x}||_1 \le ||\mathbf{x}'||_1$, we must have that $\sum_{j=0}^{2s+7} |x_j| \le 1$. Given that $|x_0| \ge 1$ and $\Sigma_< = 1$, we must have $x_0 = 1$ and $x_1, \ldots, x_{2s+7} = 0$.

Let $\overline{u_i}^{\mathrm{T}} = t^{k_{\mathbf{x}}} a^{x_{k_{\mathbf{x}}}} t^{-1} a^{x_{k_{\mathbf{x}}-1}} \ldots t^{-1} a^{x_{2s+8}}$, and let $\overline{u_i}^{\mathrm{B}} = a\overline{u_i}^{\mathrm{T}}$. Note that because $x_0 = 1$ and $x_1, \ldots, x_{2s+7} = 0$, we have that $\overline{u_i}^{\mathrm{B}} =_{\mathrm{BS}(1,n)} u_i^{\mathrm{B}}$, and, as a result, $\overline{u_i}^{\mathrm{T}} =_{\mathrm{BS}(1,n)} u_i^{\mathrm{T}}$. Also note that $\ell(\overline{u_i}^{\mathrm{B}}) = ||\mathbf{x}||_1 + 2k_{\mathbf{x}} - (2s+8)$. I claim that

$\overline{u_i}^{\mathrm{T}} t^{-(2s+8)}$ and $\overline{u_i}^{\mathrm{B}} t^{-(2s+8)}$ label geodesics in D_m. First I will show that $\eta_{0,in^{2s+8}+1,0}(\mathbf{x})$ is a geodesic in $\mathcal{C}_n$, the Cayley 2-complex of $\mathcal{P}_n$.

Let $\mathbf{y} \in \mathcal{L}_{in^{2s+8}+1}$ such that $\eta_{0,in^{2s+8}+1,0}(\mathbf{y})$ is a geodesic for $a^{in^{2s+8}+1}$ in $\mathcal{C}_n$. Since $k_{\mathbf{y}} \geq 0$, we have by [27, Lemma 3.7] that $\ell(\eta_{0,in^{2s+8}+1,0}(\mathbf{y})) = ||\mathbf{y}||_1 + 2k_{\mathbf{y}}$. By the way that $\mathbf{x}$ was chosen, we have that $||\mathbf{x}||_1 + 2k_{\mathbf{x}} \leq ||\mathbf{y}||_1 + 2k_{\mathbf{y}}$. Then since $k_{\mathbf{x}} \geq 0$, again by [27, Lemma 3.7] we have that $\ell(\eta_{0,in^{2s+8}+1,0}(\mathbf{x})) = ||\mathbf{x}||_1 + 2k_{\mathbf{x}} \leq \ell(\eta_{0,in^{2s+8}+1,0}(\mathbf{y}))$. Hence, $\eta_{0,in^{2s+8}+1,0}(\mathbf{x})$ is a geodesic in $\mathcal{C}_n$.

As we know, $\ell(\eta_{0,in^{2s+8}+1,0}(\mathbf{x})) = ||\mathbf{x}||_1 + 2k_{\mathbf{x}} = \ell(\overline{u_i}^{\mathrm{B}} t^{-(2s+8)})$. So $\overline{u_i}^{\mathrm{B}} t^{-(2s+8)}$ is also a geodesic in $\mathcal{C}_n$. Since $\overline{u_i}^{\mathrm{T}} t^{-(2s+8)}$ is a subword of $\overline{u_i}^{\mathrm{B}} t^{-(2s+8)}$, it is also a geodesic in $\mathcal{C}_n$. Therefore, if both of these words label paths in D_m, then they are both geodesics in D_m.

Recall that $\pi_{D_m} : D_m \to \mathcal{C}_n$ is the unique cellular map preserving labels that takes the basepoint of D_m to the identity in $\mathcal{C}_n$. Then $\overline{u_i}^{\mathrm{T}} t^{-(2s+8)}$ labels a path in D_m if the path γ that it labels in $\mathcal{C}_n$ is a subset of the image of π_{D_m}.

Suppose by way of contradiction that γ is not contained in the image of π_{D_m}. We know that γ starts and ends in the image of π_{D_m}. Let α be the longest initial segment of γ contained in the image of π_{D_m}, and let p be the endpoint of α. Let β be the longest subpath of γ starting at p such that every point on β other than p and the endpoint of β are outside of the image of π_{D_m}, and let q be the endpoint of β. We will consider the possible positions of p and q.

Let M be the path in D_m starting at the basepoint and labeled by a^{n^m+1}. Then $D_m \setminus M$ has two path components, D_m^{T} and D_m^{B} (the top and bottom halves of D_m), where all the 0-cells in $\pi_{D_m}(D_m^{\mathrm{T}})$ have a-length divisible by n and all the 0-cells in $\pi_{D_m}(D_m^{\mathrm{B}})$ have a-length congruent to 1 mod n. We will let $\overline{D_m^{\mathrm{T}}} = D_m^{\mathrm{T}} \cup M$ and $\overline{D_m^{\mathrm{B}}} = D_m^{\mathrm{B}} \cup M$. Note that every 0-cell on γ has an a-length that is divisible by n^{2s+8}, since $x_0, \ldots, x_{2s+7} = 0$. Therefore, p and q are not in $\pi_{D_m}(D_m^{\mathrm{B}})$.

Also note that γ has an initial segment labeled by t^{2s+8} and a final segment labeled by $t^{-(2s+8)}$. These segments are both contained in the image of π_{D_m} because they are also initial and final segments of the path labeled by $t^{2s+8}a^it^{-(2s+8)}$, which is contained in the image of π_{D_m} because $m > 2s+8$. Note that every 0-cell on γ outside of these initial and final segments, including p and q, has a t-height of at least $2s+8$. So $p, q \in \pi_{D_m}(D_m^{\mathrm{T}})$, and the first letter along β cannot be t^{-1}, since every 0-cell in $\pi_{D_m}(D_m^{\mathrm{T}})$ has a 1-cell of $\pi_{D_m}(D_m)$ labeled by t^{-1} out of it.

Case 1: Suppose the first letter along β is a t. See Figure 3.14 for an example of this case. Since the letter t only appears in $\overline{u_i}^{\mathrm{T}}$ in the prefix $t^{k_{\mathbf{x}}}$, and the path starting at the basepoint labeled by t^m is contained in D_m, this implies that $k_{\mathbf{x}} > m$ and $p =_{\mathrm{BS}(1,n)} t^m$. Then since every 0-cell past p along the prefix $t^{k_{\mathbf{x}}}$ has t-height greater than m, and is therefore not in the image of π_{D_m}, the last letter along β cannot be t. If the last letter along β is t^{-1}, then we must have that $q \in \{t^m, t^m a\}$, since these are the only 0-cells in $\pi_{D_m}(D_m^{\mathrm{T}})$ that do not have a 1-cell of $\pi_{D_m}(D_m)$ labeled by t^{-1} coming into them. If the last letter of β is a or a^{-1}, then we must have $q \in \{t^j\}_{j=1}^m \cup \{a^{n^m}t^j\}_{j=1}^m$, since these are the only 0-cells of $\pi_{D_m}(D_m^{\mathrm{T}})$ that do not have a 1-cell of $\pi_{D_m}(D_m)$ labeled by a or a^{-1} coming into them. Finally, we also cannot have that $q \in \{t^j\}_{j=1}^m$, since γ is a geodesic and therefore cannot intersect itself. Hence, there is some $j \in [m]$ such that $q =_{\mathrm{BS}(1,n)} a^{n^m}t^j$. Then at^{j-m} labels a path from p to q in $\pi_{D_m}(D_m)$. Now any path from p to q, including β, must be labeled by a word containing an a (since p and q do not have the same a-length). Since the first letter along β is a t, β reaches a t-height of at least $m+1$, and therefore must contain at least $m+1-j$ edges labeled by t^{-1} to get back down to j, the t-height of q. Therefore, $|\beta| \geq m-j+3 > m-j+1 = \ell(at^{j-m})$. This contradicts the fact that γ is a geodesic.

Case 2: Suppose the first letter along β is a or a^{-1}. See Figure 3.15 for an

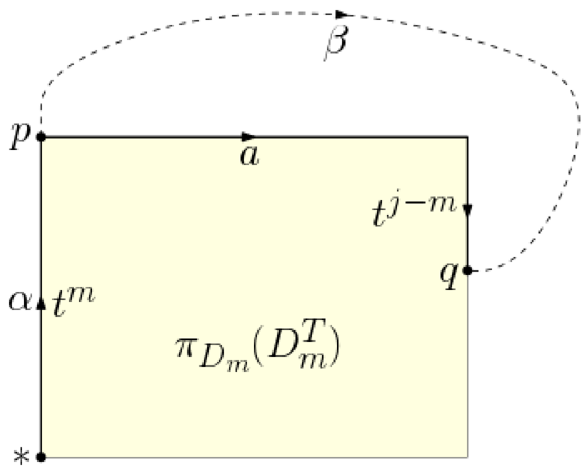

Figure 3.14: Showing γ is contained in $\pi_{D_m}(D_m^T)$ if the first letter along β is a t.

example of this case. By the same reasoning as in Case 1 for q, this implies that $p \in \{t^j\}_{j=1}^{k_{\mathbf{x}}} \cup \{a^{n^m} t^j\}_{j=1}^{k_{\mathbf{x}}}$. This also implies that $k_{\mathbf{x}} \leq m$, and neither p nor q can appear until the end of the initial segment of γ labeled by $t^{k_{\mathbf{x}}}$. Therefore, the last letter along β again cannot be t, and by the same reasoning as in Case 1, this implies that $q \in \{a^{n^m} t^j\}_{j=1}^{k_{\mathbf{x}}}$. So there is some $j_p \in [k_{\mathbf{x}}]$ such that $p \in \{t^{j_p}, a^{n^m} t^{j_p}\}$ and some $j_q \in [k_{\mathbf{x}}]$ such that $q =_{\mathrm{BS}(1,n)} a^{n^m} t^{j_q}$. Note that $j_p \geq j_q$, since p and q both appear after the initial segment of γ labeled by $t^{k_{\mathbf{x}}}$, with p appearing before q.

Suppose the first letter along β is a. Then $p =_{\mathrm{BS}(1,n)} a^{n^m} t^{j_p}$. So $t^{j_q - j_p}$ labels a path from p to q. Since p has t-height j_p and q has t-height j_q, any path from p to q, including β, must contain $j_p - j_q$ edges labeled by t^{-1}. But we know that β also contains an a, so $|\beta| \geq j_p - j_q + 1 > \ell(t^{j_q - j_p})$. This contradicts the fact that γ is a geodesic.

Suppose instead that the first letter along β is a^{-1}. Then $p =_{\mathrm{BS}(1,n)} t^{j_p}$. So $a^{n^{m-j_p}} t^{j_q - j_p}$ labels a path from p to q. As above, β must contain $j_p - j_q$ edges labeled by t^{-1}. Furthermore, since β comes after the initial segment of γ labeled by $t^{k_{\mathbf{x}}}$, this implies that every 0-cell on β has t-height at most j_p. Also, since the a-length of p is 0 and the a-length of q is m, the a-length of the word labeling β must be m. Therefore,

β must contain at least n^{m-j_p} edges labeled by an a. Given that the word labeling β also contains an a^{-1}, this implies that $|\beta| \geq j_p - j_q + n^{m-j_p} + 1 > \ell(a^{n^{m-j_p}} t^{j_q - j_p})$. This contradicts the fact that γ is a geodesic.

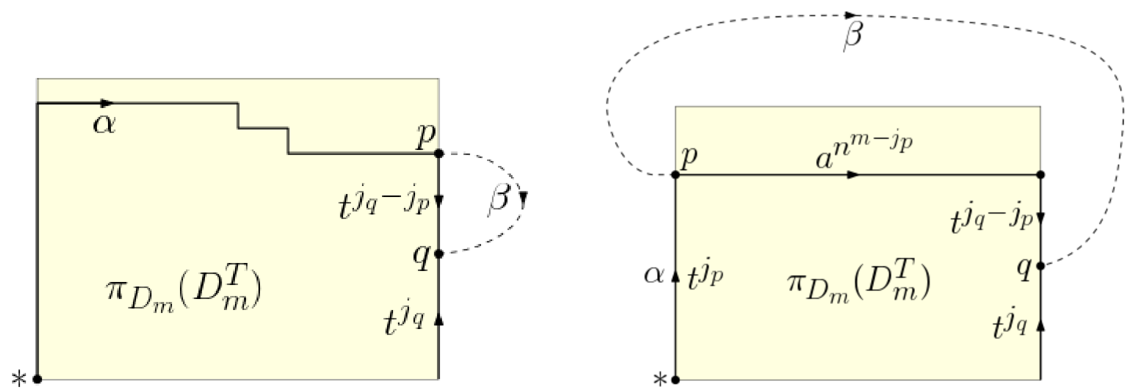

(a) The first letter along γ is an a (b) The first letter along γ is an a^{-1}

Figure 3.15: Showing γ is contained in $\pi_{D_m}(D_m^T)$ if the first letter along β is a or a^{-1}.

Therefore, γ is contained in the image of π_{D_m}, and $\overline{u_i}^{\mathrm{T}} t^{-(2s+8)}$ labels a path in D_m. Hence, we have that $\overline{u_i}^{\mathrm{T}} t^{-(2s+8)}$ labels a geodesic in D_m. Note that since this path is contained in D_m^{T}, and by the symmetry of D_m, we have that the path labeled by $\overline{u_i}^{\mathrm{B}} t^{-(2s+8)}$ is contained in D_m^{B}. So $\overline{u_i}^{\mathrm{B}} t^{-(2s+8)}$ also labels a geodesic in D_m. For $\mathrm{S} \in \{\mathrm{T}, \mathrm{B}\}$, let α_i^{S} be the D_m-geodesic from $*$ to p_i^{S} labeled by $\overline{u_i}^{\mathrm{S}}$, and let β_i^{S} be the D_m-geodesic starting at p_i^{S} labeled by $t^{-(2s+8)}$.

Let $\mathbf{y} = (1, 0, 1, 0, \ldots, 1, 0, 1)$, where $k_{\mathbf{y}} = 2s+6$. Then $\eta_{2s+8,\Sigma(\mathbf{y}),0}(\mathbf{y}) = (t^{-2}a)^{s+4}$. Note that for all $i \in \{3, \ldots, n^{m-(2s+8)}\}$, this word labels a path γ_i^{T} in $\overline{D_{2s+8}^{i}}^{\mathrm{T}}$ from p_{i-1}^{T} to a 0-cell q_i (since $\eta_{2s+8,\Sigma(\mathbf{y}),0}(\mathbf{y})$ has t-height between 0 and $-(2s+8)$ and a-length between 0 and $2s+8$). Also let γ_i^{B} be the path starting at p_{i-1}^{B} labeled by $(t^{-2}a)^{s+3}t^{-2}$. All of these paths are depicted in Figure 3.16.

Note that for $j \in \{0, \ldots, k_{\mathbf{y}}\}$, we have $|y_j| \leq \lfloor \frac{n}{2} \rfloor$, and that $k_{\mathbf{y}} < \max(2s+8, 0)$. Then if n is odd, by [27, Lemmas 3.10, 3.13], $\eta_{2s+8,\Sigma(\mathbf{y}),0}(\mathbf{y})$ is a geodesic in $\mathcal{C}_n$. If instead n is even, then also note that there are no adjacent non-zero entries in $\mathbf{y}$. Then

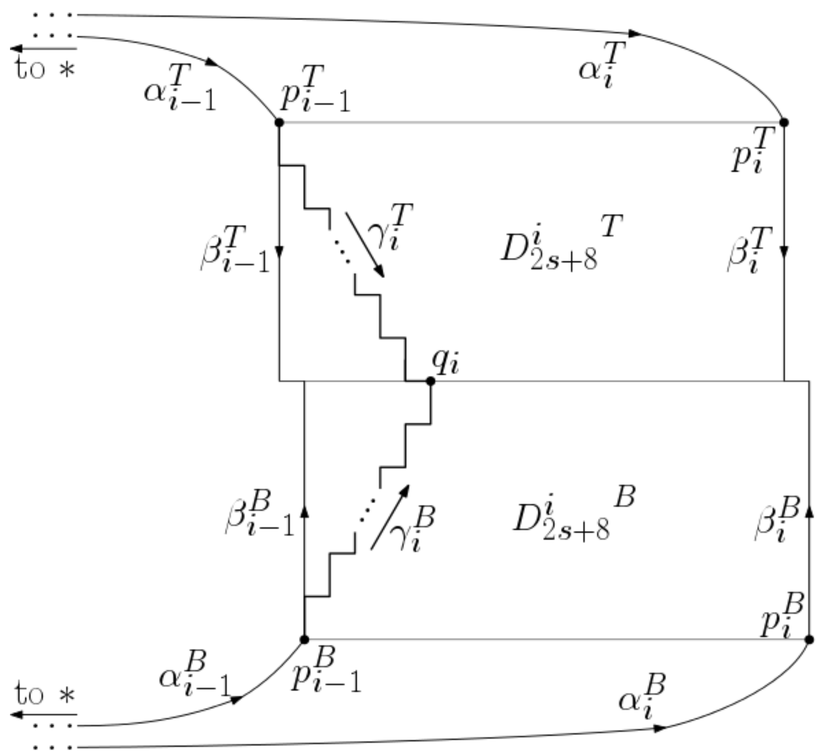

Figure 3.16: D_{2s+8}^i in D_m.

by [27, Props 3.22, 3.29], $\eta_{2s+8,\Sigma(\mathbf{y}),0}(\mathbf{y})$ is a geodesic in $\mathcal{C}_n$ in this case as well. Since $\eta_{2s+8,\Sigma(\mathbf{y}),0}(\mathbf{y})$ is a geodesic in $\mathcal{C}_n$, γ_i^{T} is a geodesic in $\overline{D_{2s+8}^i}^{\mathrm{T}}$. Now since $(t^{-2}a)^{s+3}t^{-2}$ is a subword of $\eta_{2s+8,\Sigma(\mathbf{y}),0}(\mathbf{y})$, it also labels a geodesic in $\mathcal{C}_n$, and since γ_i^{T} is contained in $\overline{D_{2s+8}^i}^{\mathrm{T}}$, we have that γ_i^{B} is contained in $\overline{D_{2s+8}^i}^{\mathrm{B}}$. Note that $\alpha_{i-1}^{\mathrm{B}} \cdot \gamma_i^{\mathrm{B}}$ is labeled by the word

$$\begin{aligned} \overline{u_{i-1}}^{\mathrm{B}}(t^{-2}a)^{s+3}t^{-2} &=_{\mathrm{BS}(1,n)} aa^{(i-1)n^{2s+8}}a^{\Sigma(\mathbf{y})-1} \\ &=_{\mathrm{BS}(1,n)} a^{(i-1)n^{2s+8}}a^{\Sigma(\mathbf{y})} \\ &=_{\mathrm{BS}(1,n)} \overline{u_{i-1}}^{\mathrm{T}}(t^{-2}a)^{s+4}, \end{aligned}$$

which labels the path $\alpha_{i-1}^T \cdot \gamma_i^T$, ending at q_i. Therefore, γ_i^{B} also ends at q_i.

Now let $i \in \{3, \ldots, n^{m-(2s+8)}\}$, and let $d_i = \min(I(p_{i-1}^{\mathrm{T}}), I(p_i^{\mathrm{T}}))$. I claim that

$I(q_i) \geq d_i + 3s + 10$. Let ζ be a geodesic from $*$ to q_i. Note that β_{i-1}^{T}, β_{i-1}^{B}, β_i^{T}, and β_i^{B} together contain all the 0-cells in the boundary of D_{2s+8}^i. Since $q_i \in D_{2s+8}^i$, we must have that ζ intersects β_j^{S} at a 0-cell z for some $j \in \{i-1, i\}$ and $\mathrm{S} \in \{\mathrm{T}, \mathrm{B}\}$. So let δ be an initial segment of ζ ending at z and ϵ a final segment of ζ starting at z with $\delta \cdot \epsilon = \zeta$. For some $k \in \{0, \ldots, 2s+8\}$, we have that z is the endpoint of the path labeled by $\overline{u_j}^{\mathrm{S}} t^{-k}$, so let $\beta(k)$ be the path from p_j^{S} to z labeled by t^{-k}. Then $\alpha_j^{\mathrm{S}} \cdot \beta(k)$ is a geodesic from $*$ to z, meaning that $|\alpha_j^{\mathrm{S}} \cdot \beta(k)| = |\delta|$. We also know that γ_i^{S} is a geodesic from p_{i-1}^{S} to q_i. If $j = i-1$ then let e be the constant path at p_{i-1}^{S}. If instead $j = i$, then let e be the path from p_{i-1}^{S} to p_j^{S} labeled by a. So $\overline{e} \cdot \gamma_i^{\mathrm{S}}$ is a path from p_j^{S} to q_i, and $e \cdot \beta(k) \cdot \epsilon$ is a path from p_{i-1}^{S} to q_i. Then $|\overline{e} \cdot \gamma_i^{\mathrm{S}}| \leq 1 + d_{D_m}(p_{i-1}^{\mathrm{S}}, q_i) \leq 1 + |e \cdot \beta(k) \cdot \epsilon| \leq 2 + |\beta(k) \cdot \epsilon|$. Also note that $|\gamma_i^{\mathrm{S}}| = \ell((t^{-2}a)^{s+3}t^{-2}) = 3(s+3) + 2 = 3s + 11$. Then we have that

$$\begin{aligned}
I(q_i) &= |\delta \cdot \epsilon| \\
&\geq |\alpha_j^{\mathrm{S}} \cdot \beta(k) \cdot \epsilon| \\
&\geq |\alpha_j^{\mathrm{S}}| + |\overline{e} \cdot \gamma_i^{\mathrm{S}}| - 2 \\
&\geq I(p_j^{\mathrm{S}}) + 1 + 3s + 11 - 2 \\
&\geq I(p_j^{\mathrm{T}}) + 3s + 10 \\
&\geq d_i + 3s + 10.
\end{aligned}$$

Now note that

$$\begin{aligned}
\max(I(\partial {D_{2s+8}^i}^{(0)})) &= \max_{j \in \{i-1,i\}, \mathrm{S} \in \{\mathrm{T},\mathrm{B}\}}(I(p_j^{\mathrm{S}})) + 2s + 8 \\
&= \max_{j \in \{i-1,i\}}(I(p_j^{\mathrm{T}})) + 2s + 9 \\
&\leq d_i + 2s + 10
\end{aligned}$$

Also note that this maximum can only be acheived at up to two 0-cells, the endpoints of β^{B}_{i-1} and β^{B}_{i}, which are not connected by a 1-cell. Therefore, $\max(\widehat{I}(\partial D^i_{2s+8})) \leq d_i + 2s + 10$.

Therefore, by Lemma 3.1.3, any path from q_i to ∂D^i_{2s+8} contains a point of height $d_i + 2s + 10$. So $\operatorname{int}\big(D^i_{2s+8}\big)$ contains a path component C_i of $D_m \setminus P_{d_i+2s+10}$. Also note that $*$ is not in C_i, and we have $q_i \in C_i$ with $|\widehat{I}(q_i) - (d_i + 2s + 10)| \geq s$. Therefore, $C_i \in C^s_{d_i+2s+10}$ from the definition of $IAV_s(D_m)$. Since $C_3, C_4, \ldots, C_{n^{m-(2s+8)}}$ are all distinct, they each contribute 1 to $IAV_s(D_m)$. Hence, $IAV_s(D_m) \geq n^{m-(2s+8)} - 2$.

□

This proposition suggests that we may be on the right track towards a proper refinement of IDiam, if we only allow reduced diagrams and relators. We will refer to this potential refinement as rIAV, even though I propose no formal definition and it is unclear if rIAV functions would give a quasi-isometry invariant or even a group invariant.

3.4 Further Questions

I think aggregate variation functions are fascinating, and I still have many unanswered questions about them. I will close out this section with some questions and conjectures about them:

- Can we define rIAV—a version of IAV using reduced diagrams and/or presentations with reduced relators—such that it is exponential for the standard presentation of $\mathrm{BS}(1, n)$ but still a quasi-isometry and/or group invariant?

- Do EAV functions give a quasi-isometry invariant?

- Is there a group whose rIAV or EAV function lands strictly between its corresponding diameter and area functions, but is not equivalent to either?
- How are (r)IAV and EAV related?
- Do all combable and almost convex groups have linear rIAV and EAV functions, as they have linear diameter functions?
- Do (r)IAV and EAV functions have interesting closure properties?
- Do (r)IAV and EAV functions have any connections to algorithms for the word problem?

Conjecture 3.4.1. All finitely generated free and abelian groups have linear rIAV and EAV functions.

Conjecture 3.4.2. All combable groups have linear EAV functions.

Chapter 4

Subdiagram Diameter Functions

4.1 Motivation and Definitions

Another possible refinement of intrinsic diameter comes from a completely different perspective on what intrinsic tame filling functions measure. The proof of Theorem 2.4.4 essentially comes down to two facts:

1. For any geodesic spanning tree of a van Kampen diagram, there is a 1-combing that is f-tame where $f(n) = \max\{\mathrm{IDiam}(I_e) | d_\Delta(*, e) \leq n\}$. In a vague sense, the 1-combing is as tame as the diameter of the icicles of the tree.

2. It is possible to replace the icicles with versions of themselves that have almost minimal diameter.

From a certain perspective, these facts imply that the 1-combing portion of the definition of intrinsic tame filling function is superfluous; there is a particular type of 1-combing that we can restrict ourselves to and acquire intrinsic tame filling functions that grow as slowly as possible, and their tameness can be measured by the diameter of the icicles of the tree without talking about 1-combings at all. In other words, defining intrinsic tame filling functions based only on diameters of icicles would result in exactly the same group invariant.

From this perspective, the proof of (2) works because icicles of the same tree intersect each other so nicely. For any two icicles, either one is inside the other or they do not intersect (besides at their boundary). This makes it easier to find a good order in which to replace icicles without having to worry about intersections messing up an area that has already been fixed; just replace bigger icicles first. This way, there is no possibility of ruining the nice diameter of a previously-replaced icicle by replacing the next icicle; if any part of a previously-replaced icicle is later replaced again, the latter replacement must be replacing some icicle entirely contained within the previously-replaced icicle, and this can never increase the diameter of the previously-replaced icicle.

This gives us another way to try to strengthen the notion of intrinsic tame filling functions. We could define new types of filling functions based on the diameters of a set of subcomplexes that need not intersect each other nicely, with the intention of breaking the proof of (2). We define such a type of filling function below, simply using the set of all subcomplexes that are themselves van Kampen diagrams.

Since subcomplexes do not come equipped with a basepoint, it will first be convenient to define a type of diameter that does not take a basepoint into account. Given a van Kampen diagram Δ, the *unbased intrinsic diameter* of Δ is

$$\overline{\mathrm{IDiam}}(\Delta) = \max\{d_\Delta(x, y) | x, y \in \Delta^{(0)}\}.$$

Definition 4.1.1. A *subdiagram* of Δ is a simply-connected subcomplex of Δ. Note that any subdiagram D of a van Kampen diagram can itself be thought of as a van Kampen diagram without a chosen basepoint. It will be useful to let $|\partial D|$ refer to the length of a boundary circuit of D without specifying a basepoint.

Definition 4.1.2. Given a function $f : \mathbb{N} \to \mathbb{N}$ and a van Kampen diagram Δ, f is an *intrinsic subdiagram diameter (ISD) function* for Δ if, for every subdiagram D of Δ, $\overline{\text{IDiam}}(D) \leq f(|\partial D|)$.

Definition 4.1.3. A non-decreasing function $f : \mathbb{N} \to \mathbb{N}$ is an *intrinsic subdiagram diameter (ISD) function* for $\langle A|R\rangle$ if, for all $w \in (A \cup A^{-1})^*$ with $w =_G 1$, there is a van Kampen diagram Δ_w for w such that f is an ISD function for Δ_w.

Note that any ISD function for a presentation is bounded below by the intrinsic diameter function of the presentation, since every van Kampen diagram is a subdiagram of itself. So there is potential for ISD functions to give a proper refinement of the intrinsic diameter function. Also note that, as with other distance-based filling functions, we could also define an extrinsic version of subdiagram diameter functions, measuring distance in the Cayley complex.

4.2 Quasi-isometry invariance

We will now confirm that ISD functions give a quasi-isometry invariant.

Proposition 4.2.1. *Let G and H be quasi-isometric groups with finite presentations $\mathcal{P}$ and $\mathcal{Q}$. If f is an ISD function for $\mathcal{P}$, then there is a presentation $\mathcal{Q}''$ for H with an ISD function equivalent to f, where $\mathcal{Q}''$ is $\mathcal{Q}$ with finitely many relators added.*

The proof follows a similar structure to the proof that IDiam is a quasi-isometry invariant. We take an arbitrary word w from $\mathcal{Q}$, use the quasi-isometry between G and H to convert it into a word w' in $\mathcal{P}$ that has a "nice" van Kampen diagram $\Delta_{w'}$, and then use the quasi-isometry to convert $\Delta_{w'}$ into a van Kampen diagram $\Delta_{w''}$ for a word w'' that is "close" to w in some way. We then glue a collar around $\Delta_{w''}$ to convert it into a van Kampen diagram Δ_w for the original word w, and show that Δ_w is also "nice" as a result of the "niceness" of $\Delta_{w'}$. The main difference is that, in the case of ISD functions, it takes considerably more work to show that Δ_w is "nice" compared to the corresponding proof for IDiam.

Proof. Let $\mathcal{P} = \langle A|R\rangle$ and $\mathcal{Q} = \langle B|S\rangle$, and let $\rho = \max\{\ell(r) : r \in R\}$. Let f be an ISD function for $\mathcal{P}$. Let $w = b_1 \dots b_n \in (B \cup B^{-1})^*$ with $w =_H 1$. We want to show that there is a van Kampen diagram Δ_w for w with respect to a presentation $\mathcal{Q}''$ for H such that the diameter of every subdiagram of Δ_w is bounded by a function equivalent to f.

Step 1: Using the fact that G and H are quasi-isometric, "convert" w into a word $w' \in (A \cup A^{-1})^*$ and choose a van Kampen diagram $\Delta_{w'}$ for w' with respect to $\mathcal{P}$ that is guaranteed to exist by the fact that f is an ISD function for $\mathcal{P}$.

Since G and H are quasi-isometric, there exist functions $\phi : G \to H$ and $\theta : H \to G$ and a constant $K > 0$ such that for all $g_1, g_2 \in G$ and $h_1, h_2 \in H$,

1. $\frac{1}{K} d_G(g_1, g_2) - K \le d_H(\phi(g_1), \phi(g_2)) \le K d_G(g_1, g_2) + K$,
2. $\frac{1}{K} d_H(h_1, h_2) - K \le d_G(\theta(h_1), \theta(h_2)) \le K d_H(h_1, h_2) + K$,
3. $d_G(g_1, \theta(\phi(g_1))) \le K$, and
4. $d_H(h_1, \phi(\theta(h_1))) \le K$.

Then define the functions $\widetilde{\phi} : G \times (A \cup A^{-1})^* \to (B \cup B^{-1})^*$ and $\widetilde{\theta} : H \times (B \cup B^{-1})^* \to (A \cup A^{-1})^*$ in the following way. For each $a \in A$ and $g \in G$, there is some $v \in (B \cup B^{-1})^*$ with $\phi(g)v =_H \phi(ga)$ and $\ell(v) \le 2K$, by property (1). So let $\widetilde{\phi}(g, a) = v$ and $\widetilde{\phi}(ga, a^{-1}) = v^{-1}$. Similarly, for each $b \in B$ and $h \in H$, there is some $u \in (A \cup A^{-1})^*$ with $\theta(h)u =_G \theta(hb)$ and $\ell(u) \le 2K$, by property (2). So let $\widetilde{\theta}(h, b) = u$ and $\widetilde{\theta}(hb, b^{-1}) = u^{-1}$.

Now given $u = a_1 \dots a_m \in (A \cup A^{-1})^*$ and $v = b_1 \dots b_n \in (B \cup B^{-1})^*$, for $i \in \{0, \dots, m\}$ define $u_i = a_1 \dots a_i$ and for $i \in \{0, \dots, n\}$ define $v_i = b_1 \dots b_i$. Then define

$$\widetilde{\phi}(g, u) = \widetilde{\phi}(gu_0, a_1) \dots \widetilde{\phi}(gu_{i-1}, a_i) \dots \widetilde{\phi}(gu_{m-1}, a_m).$$
$$\widetilde{\theta}(h, v) = \widetilde{\theta}(hv_0, b_1) \dots \widetilde{\theta}(hv_{i-1}, b_i) \dots \widetilde{\theta}(hv_{n-1}, b_n).$$

Note that as a result of these definitions, $\phi(g)\widetilde{\phi}(g, u) =_H \phi(gu)$ and $\theta(h)\widetilde{\theta}(h, v) =_G \theta(hv)$. In particular, if $u =_G 1$ and $v =_H 1$, then $\phi(g)\widetilde{\phi}(g, u) =_H \phi(g)$, making $\widetilde{\phi}(g, u) =_H 1$, and similarly $\widetilde{\theta}(h, v) =_G 1$.

Let $w' = \widetilde{\theta}(1, w)$. By the above note, $w' =_G 1$. So there exists a van Kampen diagram $\Delta_{w'}$ for w' with respect to the presentation $\mathcal{P}$ such that, for every subdiagram D of $\Delta_{w'}$, we have that $\overline{\text{IDiam}}(D) \leq f(|\partial D|)$.

Step 2: "Convert" $\Delta_{w'}$ into a van Kampen diagram $\Delta_{w''}$ for a word $w'' \in (B \cup B^{-1})^*$ with respect to a finite presentation for H.

Let $w'' = \widetilde{\phi}(1, w')$. Construct a van Kampen diagram $\Delta_{w''}$ for w'' from $\Delta_{w'}$ in the following way. First, for notation's sake, replace each vertex x of $\Delta_{w'}$ with a vertex x'. Then for each edge e in $\Delta_{w'}$ directed from a vertex x to a vertex y and labeled by a letter $a \in A$, replace e with a directed edge path from x' to y' labeled by $\widetilde{\phi}(\pi_{\Delta_{w'}}(x), a)$. These replacements result in the complex $\Delta_{w''}$. Since $\Delta_{w'}$ can be thought of as a subset of $\mathbb{R}^2$, so can $\Delta_{w''}$, and let $h : \Delta_{w'} \to \Delta_{w''}$ be the homeomorphism that extends to the identity on $\mathbb{R}^2$. Note that h replaces the boundary circuit of $\Delta_{w'}$, which is labeled by w', with a circuit labeled by $\widetilde{\phi}(1, w') = w''$. It also replaces the boundary circuit of each 2-cell, labeled by a relator $r \in R$, with a circuit labeled by $\widetilde{\phi}(g, r)$ for some $g \in G$. Since $r =_G 1$, we have $\widetilde{\phi}(g, r) =_H 1$. As a result, $\Delta_{w''}$ is a van

Kampen diagram for w'' with respect to the presentation $\mathcal{Q}' := \langle B | S \cup R' \rangle$, where $R' = \{\widetilde{\phi}(g,r) : g \in G, r \in R\}$. Note that R' is finite, since $\ell(\widetilde{\phi}(g,r)) \leq 2K\rho$ for all $g \in G$ and $r \in R$. Also note that $\mathcal{Q}'$ is a presentation for H, since each element of R' is equal to the identity in H.

Step 3: Construct a van Kampen diagram Δ_w for the original word w from $\Delta_{w''}$ by gluing a collar around $\Delta_{w''}$.

Recall that $w = b_1 \ldots b_n$, and for $i \in \{0, \ldots, n\}$ we've defined $w_i = b_1 \ldots b_i$. For $i \in \{0, \ldots, n\}$, let x'_i be the vertex on the boundary of $\Delta_{w'}$ such that $w'_i = \widetilde{\theta}(1, w_i)$ labels the initial segment of the boundary circuit of $\Delta_{w'}$ starting at the basepoint and ending at x'_i (so $x'_0 = x'_n$ is the basepoint of $\Delta_{w'}$). Then for $i \in [n]$, we have that x'_{i-1} is connected to x'_i by a directed edge path on the boundary labelled by the word $b'_i = \widetilde{\theta}(w_{i-1}, b_i)$. Note that $\ell(b'_i) \leq 2K$. Since x''_i is the vertex in $\Delta_{w''}$ corresponding to x'_i, we have that x''_{i-1} is connected to x''_i by a directed edge path α_i on the boundary labelled by $b''_i = \widetilde{\phi}(w'_{i-1}, b'_i)$. Note that $\ell(b''_i) \leq 2K\ell(b'_i) \leq 4K^2$.

Now by property (4) of the quasi-isometries ϕ and θ above, for each $i \in \{0, \ldots, n\}$ we have that $d_H(w_i, w''_i) \leq K$. So let $v_i \in (B \cup B^{-1})^*$ be a word of length at most K such that $w''_i v_i =_H w_i$, and attach to $\Delta_{w''}$ a directed edge path β_i labelled by v_i starting at x''_i. Call the endpoint of this path x_i. In particular, since $w_0 =_H 1$ and $w''_0 =_H 1$, we may let $v_0 = v_n$ be the empty word so that $x_0 = x''_0$. Then for each $i \in [n]$, we have that

$$\begin{aligned} w''_{i-1} v_{i-1} b_i &=_H w_{i-1} b_i \\ &=_H w_i \\ &=_H w''_i v_i \end{aligned}$$

Since $w''_i =_H w''_{i-1}b''_i$, this implies that $v_{i-1}b_i =_H b''_i v_i$. So $v_{i-1}b_i(b''_i v_i)^{-1} =_H 1$

So for each $i \in [n]$, glue a 2-cell σ_i with boundary circuit labeled by $v_{i-1}b_i(b''_i v_i)^{-1}$ along the directed edge path on the boundary of the diagram from x_i to x_{i+1} given by $\overline{\alpha_{i-1}} \cdot \beta_i \cdot \alpha_i$. Note that the length of the boundary circuit of each of these 2-cells is at most $4K^2 + 2K + 1$. Call the resulting singular disk diagram Δ_w.

Then the boundary circuit of Δ_w is labelled by $b_1 b_2 \dots b_n = w$. Let T be the (finite) set of words $v \in (B \cup B^{-1})^*$ such that $v =_H 1$ and $\ell(v) \leq 4K^2 + 2K + 1$. Then the boundary circuit of each 2-cell added to $\Delta_{w''}$ to make Δ_w is contained in T, so Δ_w is a van Kampen diagram for w with respect to the presentation $\mathcal{Q}'' := \langle B | S \cup R' \cup T \rangle$. Again, $\mathcal{Q}''$ is a presentation for H because all the relators in T are equal to the identity in H.

Step 4: Show that there is a function g equivalent to f (that does not depent on w) such that any subdiagram D of Δ_w has $\overline{\text{IDiam}}(D) \leq g(|D|)$.

Let D be a subdiagram of Δ_w. Let $D'' = D \cap \Delta_{w''}$. Note that D'' may not be a subdiagram of $\Delta_{w''}$, since D'' may not be path connected.

Step 4(a): Show that every path component of D'' is a subdiagram of $\Delta_{w''}$. In other words, every path component of D'' has trivial fundamental group.

Step 4(a)(i): Show that removing from D each 1-cell in $D \cap \partial\Delta_w$ that is not on the boundary of any 2-cell in D splits D into simply-connected components.

For $i \in [n]$ let e_i be the 1-cell connecting x_{i-1} and x_i in the boundary of Δ_w. Then, let $B = \{e_{i_1}, \dots, e_{i_m}\}$ be the set of 1-cells in $\partial\Delta_w \cap D$ that are not on the boundary of any 2-cell in D, in no particular order. For $j \in \{0, \dots, m\}$, let $B_j = \bigcup_{k \leq j} \text{int}(e_{i_k})$. We will show by induction that, for each $j \in \{0, \dots, m\}$, each path component of $D \setminus B_j$ has trivial fundamental group. So in particular this is true for $j = m$.

For the base case, we know that $D \setminus B_0 = D$ is simply connected by the definition of subdiagram, so its single path component has trivial fundamental group. Now let

$j \in [m]$ and suppose that each path component of $D \setminus B_{j-1}$ has trivial fundamental group. Let C_j be the path component of $D \setminus B_{j-1}$ containing e_{i_j} and choose any $y_j \in \operatorname{int}(e_{i_j})$. I claim that $C_j \setminus y_j$ has two path components, one containing x_{i_j-1} and one containing x_{i_j}.

First note that for all $j \in [m]$, $D \setminus B_j$ is a subcomplex of D, since we are only removing the interiors of 1-cells that are not in the boundary of any 2-cell. In particular, then, C_j is a subcomplex of D, as is $C_j \setminus \operatorname{int}(e_{i_j})$. Suppose by way of contradiction that x_{i_j-1} and x_{i_j} are in the same path component of $C_j \setminus y_j$. Note that $C_j \setminus y_j$ is homotopy equivalent to $C_j \setminus \operatorname{int}(e_{i_j})$, since the segments of e_{i_j} on either side of y_j can deformation retract onto x_{i_j-1} and x_{i_j}. So x_{i_j-1} and x_{i_j} are in the same component of $C_j \setminus \operatorname{int}(e_{i_j})$. Since this path component of $C_j \setminus \operatorname{int}(e_{i_j})$ is a connected 2-complex, its 1-skeleton is a connected graph. So there is a simple edge path γ from x_{i_j-1} to x_{i_j} in $C_j \setminus \operatorname{int}(e_{i_j})$. Now let $\overrightarrow{e_{i_j}}$ denote the directed edge path from x_{i_j} to x_{i_j-1} in C_j along e_j. Then $\gamma \cdot \overrightarrow{e_{i_j}}$ is a simple edge circuit in C_j containing e_{i_j}. By the Jordan Curve Theorem, $\gamma \cdot \overrightarrow{e_{i_j}}$ considered as a simple closed curve in $\mathbb{R}^2$ has an inside and an outside.

Note that since C_j is simply connected, we must have that every point on the inside of $\gamma \cdot \overrightarrow{e_{i_j}}$ is in C_j. For suppose instead that there is some point p in the inside of $\gamma \cdot \overrightarrow{e_{i_j}}$ with $p \notin C_j$. Then $\gamma \cdot \overrightarrow{e_{i_j}}$ represents a (nontrivial) generator of $\pi_1(\mathbb{R}^2 \setminus p)$. So the inclusion-induced map $\widetilde{i} : \pi_1(C_j) \to \pi_1(\mathbb{R}^2 \setminus p)$ sends the element of $\pi_1(C_j)$ represented by $\gamma \cdot \overrightarrow{e_{i_j}}$ to a nontrivial element, and therefore $\gamma \cdot \overrightarrow{e_{i_j}}$ represents a nontrivial element of $\pi_1(C_j)$. This contradicts the fact that C_j is simply connected.

Now since $\gamma \cdot \overrightarrow{e_{i_j}}$ is a simple edge circuit containing e_{i_j}, the region on one side of e_{i_j} is in the inside of $\gamma \cdot \overrightarrow{e_{i_j}}$, and the region on the other side of e_{i_j} is in the outside of $\gamma \cdot \overrightarrow{e_{i_j}}$. In particular, since e_{i_j} is on the boundary of Δ_w, it is adjacent to $\mathbb{R}^2 \setminus \Delta_w$ on one side, which must be contained in the outside of $\gamma \cdot \overrightarrow{e_{i_j}}$. So σ_{i_j} is adjacent to e_{i_j}

on the inside of $\gamma \cdot \overrightarrow{e_{i_j}}$. However, we chose e_{i_j} such that σ_{i_j} is not in D, and therefore not in C_j. This contradicts the fact above that every point on the inside of $\gamma \cdot \overrightarrow{e_{i_j}}$ is in C_j. Therefore, by this final contradiction, we must have that x_{i_j-1} and x_{i_j} are in different path components of $C_j \setminus y_j$.

To show that these are the only path components of $C_j \setminus y_j$, let q be a 0-cell that is not in the path component of $C_j \setminus y_j$ that contains x_{i_j-1}. Since C_j is a path connected subcomplex, there is a simple directed edge path η from q to x_{i_j-1} in C_j. Since q and x_{i_j-1} are in different path components of $C_j \setminus y_j$, we must have that η contains y_j, and therefore that e_{i_j} is a 1-cell on η. Since η is a simple path that ends at x_{i_j-1}—one of the 0-cells in e_{i_j}—this implies that $\overrightarrow{e_{i_j}}$ is the last directed edge of η. Therefore, removing $\overrightarrow{e_{i_j}}$ from η gives a path from q to x_{i_j}. Since η does not repeat edges, this path does not contain e_{i_j}, and therefore does not contain y_j, meaning that q and x_{i_j} are in the same path component of $C_j \setminus y_j$. Since every 0-cell of $C_j \setminus y_j$ is either in the path component of x_{i_j-1} or x_{i_j}, these are the only two path components of $C_j \setminus y_j$.

Let X_j be the path component of $C_j \setminus y_j$ containing x_{i_j-1} and let Y_j be the path component of $C_j \setminus y_j$ containing x_{i_j}. Since C_j is a 2-complex, it is locally path connected. This implies that every point in $C_j \setminus y_j$—an open subset of C_j—has a path connected neighborhood contained in $C_j \setminus y_j$. The fact that these neighborhoods are path connected means that each one is contained entirely in either X_j or Y_j, and therefore X_j is the union of these neighborhoods that it contains, as is Y_j. Hence, X_j and Y_j are open subsets of $C_j \setminus y_j$. Furthermore, $\operatorname{int}\big(e_{i_j}\big)$ is an open subset of C_j, since e_{i_j} is not on the boundary of any 2-cell in C_j.

So let $X_j' = X_j \cup \operatorname{int}\big(e_{i_j}\big)$ and $Y_j' = Y_j \cup \operatorname{int}\big(e_{i_j}\big)$. Then X_j' and Y_j' form an open cover of C_j with $X_j' \cap Y_j' = \operatorname{int}\big(e_{i_j}\big)$. Then by the Seifert-Van Kampen Theorem,

$$\pi_1(C_j) = \pi_1(X_j') *_{\pi_1\left(\operatorname{int}\left(e_{i_j}\right)\right)} \pi_1(Y_j').$$

We know that $\pi_1(C_j) = 1$ and $\pi_1\big(\mathrm{int}(e_{i_j})\big) = 1$, so this implies that $\pi_1(X'_j) * \pi_1(Y'_j) = 1$. Therefore, both X'_j and Y'_j have trivial fundamental group. Since X'_j and Y'_j deformation retract onto $X_j \setminus \mathrm{int}(e_{i_j})$ and $Y_j \setminus \mathrm{int}(e_{i_j})$, respectively, we must have that $X_j \setminus \mathrm{int}(e_{i_j})$ and $Y_j \setminus \mathrm{int}(e_{i_j})$ are the path components of $C_j \setminus \mathrm{int}(e_{i_j})$, and they each have trivial fundamental group.

Since every path component of $D \setminus B_{j-1}$ other than C_j is a path component of $D \setminus B_j$, and all of them have trivial fundamental group by the inductive hypothesis, this implies that every path component of $D \setminus B_j$ has trivial fundamental group. This completes the induction, giving us that every path component of $D \setminus B_m$ has trivial fundamental group.

Step 4(a)(ii): Show that each path component of $D \setminus B_m$ (i.e. each of the remaining components after splitting D up in the previous step) either deformation retracts into $D \cap \Delta_{w''} = D''$, or has no intersection with $\Delta_{w''}$, implying that every path component of D'' has trivial fundamental group.

First we will define a deformation retraction $R : \Delta_w \times I \to \Delta_w$ onto $\Delta_{w''}$ as follows.

Let

$$S^1_+ = \{(a,b) \in \mathbb{R}^2 | a^2 + b^2 = 1 \text{ and } b \geq 0\} \text{ and } S^1_- = \{(a,b) \in \mathbb{R}^2 | a^2 + b^2 = 1 \text{ and } b \leq 0\},$$

and define $r_1 : S^1_+ \times I \to D^2$ by

$$r_1((a,b),t) = (a, (1-2t)b).$$

Now given $i \in [n]$, let $\psi_i : \sigma_i \to D^2$ be a homeomorphism such that e_i is mapped to S^1_+. Then $R^i_1 = \psi_i^{-1} \circ r_1 \circ (\psi_i \times \mathrm{id}_I) : \sigma_i \times I \to \sigma_i$ is a deformation retraction of

σ_i onto $\partial\sigma_i \setminus \text{int}(e_i)$. Now $\partial\sigma \setminus \Delta''_w = e_i \cup \text{int}(\alpha_{i-1}) \cup \text{int}(\alpha_i)$, so we now need only retract the $\alpha_i's$ into Δ''_w. For $i \in [n]$, choose a homeomorphism $A_i : \alpha_i \to I$ with $A_i(x''_i) = 0$ and $A_i(x_i) = 1$, and let $r_2 : I \times I \to I$ be given by $\varphi(s,t) = s(1-t)$. Then $R_2^i = A_i^{-1} \circ r_2 \circ (A_i \times \text{id}_I) : \alpha_i \times I \to \alpha_i$ is a deformation retraction of α_i onto x''_i.

Now to define R, let $t \in I$ and $x \in \Delta_w$. If $x \in \Delta_{w''}$, then let $R(t,x) = x$. Otherwise, for any $x \in \Delta_w \setminus \Delta_{w''}$, we must have that x is in a 2-cell σ_i for some $i \in [n]$. If $t \leq \frac{1}{2}$, let $R(x,t) = R_1^i(x, 2t)$. Now as noted above, $R_1^i(x,1)$ is either in $\Delta_{w''}$, α_{i-1}, or α_i. For $t \geq \frac{1}{2}$, let

$$R(x,t) = \begin{cases} R_1^i(x,1), & \text{if } R_1^i(x,1) \in \Delta_{w''} \\ R_2^{i-1}(R_1^i(x,1), 2t-1), & \text{if } R_1^i(x,1) \in \alpha_{i-1} \\ R_2^i(R_1^i(x,1), 2t-1), & \text{if } R_1^i(x,1) \in \alpha_i \end{cases}$$

Note that the definitions of R on the different pieces of its domain agree wherever they intersect, making R well-defined and continuous.

Now let C be any path component of $D \setminus B_m$ that intersects $\Delta_{w''}$ and let $x \in C \setminus \Delta_{w''}$. If x is in a 2-cell σ contained in C, we know from the definition of R that $R(x,t) \in \sigma$ for all $t \in I$, so $R(x,t) \in C$. Otherwise, since $x \notin B_m$, we have that x is on α_i for some $i \in [n]$. Since x is not in any 2-cell contained in C, both 2-cells containing x, σ_i and σ_{i+1}, are not in D. Then e_i and e_{i+1} are therefore both in B, and thus not contained in C. Now since C intersects $\Delta_{w''}$, there must be a path in $C^{(1)}$ from x to some point in $\Delta_{w''}$. By the construction of Δ_w, this path must contain $A_i([0,t_x])$ where $A_i(t_x) = x$. But by the definition of R, $A_i([0,t_x])$ is exactly the image of $R(x,\cdot)$. So $R(x,t) \in C$ for all $t \in I$. This implies that $R|_{C\times I}$ is a deformation retract from C to $C \cap \Delta_{w''}$, making $C \cap \Delta_{w''}$ path connected, and therefore a path component of D''. Hence, every path component of D'' has trivial fundamental group, and is a subdiagram of $\Delta_{w''}$.

Step 4(b): Show that each path component C of D'' viewed as a subdiagram of $\Delta_{w''}$ has diameter bounded by a function equivalent to f, by considering a subdiagram that corresponds to C in $\Delta_{w'}$.

Let C be a path component (and therefore a subdiagram) of D'' and let x be a 0-cell of C. We know that C is a subdiagram of D''. Recall that $h : \Delta_{w'} \to \Delta_{w''}$ is a homeomorphism, and that it maps 2-cells to 2-cells, 0-cells to 0-cells, and 1-cells to edge paths. Let $C' = h^{-1}(C)$. Since h is a homeomorphism, C' is simply connected. However, C' may not be a subcomplex of $\Delta_{w'}$, because C may not contain every 1-cell of some edge path that replaced a 1-cell of $\Delta_{w'}$. So suppose that e is a 1-cell of $\Delta_{w'}$ such that $\emptyset \subset \text{int}(e) \cap C' \subset \text{int}(e)$. Note that any 2-cells in $\Delta_{w'}$ containing e must not be in C', since otherwise $e \subseteq C'$. Then since h maps 2-cells to 2-cells, this implies that no 2-cell of $\Delta_{w''}$ containing a 1-cell in $h(e)$ can be in C. Therefore, $h(e) \cap C \subseteq \partial C$. So if $x \in h(e)$, for some such 1-cell e, then $d_C(x, \partial C) = 0$.

Otherwise, let p and q be the endpoints of e. Since $h^{-1}(x) \notin e$ but C' is path connected, this implies that for every $y \in e \cap C'$ there is a path in $e \cap C'$ from y to one of p or q. So $e \cap C'$ is composed of two paths (each of which may contain only a single point), one of which has p as an endpoint, and one of which has q as an endpoint. Since these paths may be deformation retracted down to p and q, respectively, we have that $C' \setminus \text{int}(e)$ is still simply connected. Removing the interior of each such 1-cell of C' leaves us with a simply connected subcomplex $\widetilde{C}$ of $\Delta_{w'}$, i.e., a subdiagram of $\Delta_{w'}$ that still contains $h^{-1}(x)$.

Now recall that we chose $\Delta_{w'}$ such that $\overline{\text{IDiam}}(\widetilde{C}) \leq f(|\partial\widetilde{C}|)$. Therefore, $d_{\widetilde{C}}(h^{-1}(x), \partial\widetilde{C}) \leq f(|\partial\widetilde{C}|)$. Note that $\partial\widetilde{C} \subseteq \partial C'$, $\widetilde{C} \subseteq C'$, and $|\partial\widetilde{C}| \leq |\partial C'|$, so this implies that $d_{C'}(h^{-1}(x), \partial C') \leq f(|\partial C'|)$. Now since each edge in $\Delta_{w'}$ corresponds to an edge path of length at most $2K$ in $\Delta_{w''}$, this implies that $d_C(x, \partial C) \leq 2Kf(|\partial C'|) \leq 2Kf(|\partial C|)$.

So for any $x \in C$, we have that $d_C(x, \partial C) \leq 2Kf(|\partial C|)$.

Step 4(c): Show that the previous step implies that the diameter of D is bounded above by a function equivalent to f.

Let x be a 0-cell in D. If $x \notin D''$, then by the construction of Δ_w, x is in α_i for some $i \in [n]$. So $d_{\Delta_w}(x, x_i) \leq |\alpha_i| \leq K$. Since x_i is on the boundary of Δ_w, this implies that $d_D(x, \partial D) \leq K$.

Otherwise, $x \in D''$, and let C be the path component of D'' containing x. Since we know that $d_C(x, \partial C) \leq 2Kf(|\partial C|)$, let $y \in \partial C$ with $d_C(x, y) \leq 2Kf(|\partial C|)$. If $y \in \partial D$, then $d_D(x, \partial D) \leq 2Kf(|\partial C|)$. Otherwise, by the construction of Δ_w, we must have that y is on the boundary of a 2-cell in the collar. Since the boundary circuits of these 2-cells are of length at most $4K^2 + 2K + 1$, there is an edge path from y to $\partial\Delta_w$ of length at most $2K^2 + K$. So $d_D(y, \partial D) \leq 2K^2 + K$. This implies that

$$d_D(x, \partial D) \leq 2Kf(|\partial C|) + 2K^2 + K.$$

Now to compare $|\partial C|$ to $|\partial D|$, note that a directed edge e in the boundary circuit of C is not in the boundary circuit of D only if its reverse $\bar{e}$ is part of the boundary circuit of a 2-cell σ in D that is in the collar of Δ_w. In this case, the one directed edge in the boundary circuit of σ that is also in the boundary circuit of Δ_w will appear in the boundary circuit of ∂D instead. Each 2-cell in the collar of Δ_w contains at most $4K^2$ 1-cells from ∂C. So this implies that $|\partial C| \leq 4K^2|\partial D|$.

Putting everything together, we get that for all $x \in D$,

$$\begin{aligned} d_D(x, \partial D) &\leq \max(K, 2Kf(|\partial C|) + 2K^2 + K) \\ &\leq 2Kf(4K^2|\partial D|) + 2K^2 + K \end{aligned}$$

Now let $x, y \in D$. Then there are vertices $\bar{x}, \bar{y} \in \partial D$ such that both $d_D(x, \bar{x})$ and $d_D(y, \bar{y})$ are less than or equal to $2Kf(4K^2|\partial D|) + 2K^2 + K$. Since $\bar{x}, \bar{y} \in \partial D$, there is an edge path in $|\partial D|$ from $\bar{x}$ to $\bar{y}$ of length at most $\frac{1}{2}|\partial D|$. Therefore,

$$d_D(x, y) \leq 4Kf(4K^2|\partial D|) + 4K^2 + 2K + \frac{1}{2}|\partial D|.$$

This implies that

$$\overline{\text{IDiam}(D)} \leq 4Kf(4K^2|\partial D|) + 4K^2 + 2K + \frac{1}{2}|\partial D|.$$

So define $g : \mathbb{N} \to \mathbb{N}$ by $g(n) = 4Kf(4K^2 n) + 4K^2 + 2K + \lceil \frac{1}{2} n \rceil$. Note that g is equivalent to f. Then for every subdiagram D of Δ_w, we have that $\overline{\text{IDiam}}(D) \leq g(|\partial D|)$. So g is an ISD function for Δ_w. Since we've shown that we can construct such a van Kampen diagram for every word $w =_H 1$ with respect to the presentation $\mathcal{Q}''$, g is an ISD function for $\mathcal{Q}''$.

□

4.3 Further Questions

It is left for future work to determine if there is a group whose ISD functions all grow faster than its intrinsic diameter. However, I suspect that the type of van Kampen diagram constructed in the proof of Theorem 3.2.3—a van Kampen diagram that has essentially only one hill the height of its diameter—could be used to prove that ISD functions are equivalent to intrinsic diameter functions.

The way we have defined ISD functions suggests a way to define a subdiagram version of any filling function defined from a diagram measurement, as follows. Given a diagram measurement M, a function $f : \mathbb{N} \to \mathbb{N}$ is a subdiagram-M function for a van Kampen diagram Δ if, for every subdiagram D of Δ, $M(D) \leq f(|\partial D|)$. Then f is a subdiagram-M function for a presentation $\langle A|R \rangle$ if, for all $w \in (A \cup A^{-1})^*$ with $w =_G 1$, there is a van Kampen diagram Δ_w for w such that f is a subdiagram-M function for Δ_w. In the same way as for ISD functions, any subdiagram function would be bounded below by the original filling function, since every van Kampen diagram is a subdiagram of itself. We leave it to future work to determine if other such subdiagram filling functions give quasi-isometry invariants, and if they give proper refinements of the original filling function used to define them.

Bibliography

[1] G. Baumslag, S. M. Gersten, M. Shapiro, and H. Short. Automatic groups and amalgams. *J. Pure Appl. Algebra*, 76(3):229–316, 1991.

[2] G. Baumslag and D. Solitar. Some two-generator one-relator non-Hopfian groups. *Bull. Amer. Math. Soc.*, 68:199–201, 1962.

[3] N. Brady, T. Riley, and H. Short. *The geometry of the word problem for finitely generated groups*. Advanced Courses in Mathematics. CRM Barcelona. Birkhäuser Verlag, Basel, 2007. Papers from the Advanced Course held in Barcelona, July 5–15, 2005.

[4] M. Brazil. Growth functions for some nonautomatic Baumslag-Solitar groups. *Trans. Amer. Math. Soc.*, 342(1):137–154, 1994.

[5] M. Bridson and T. Riley. Extrinsic versus intrinsic diameter for Riemannian filling-discs and van Kampen diagrams. *J. Differential Geom.*, 82(1):115–154, 2009.

[6] M. R. Bridson. The geometry of the word problem. In *Invitations to geometry and topology*, volume 7 of *Oxf. Grad. Texts Math.*, pages 29–91. Oxford Univ. Press, Oxford, 2002.

[7] B. Brink and R. B. Howlett. A finiteness property and an automatic structure for Coxeter groups. *Math. Ann.*, 296(1):179–190, 1993.

[8] M. Brittenham and S. Hermiller. Tame filling invariants for groups. *Internat. J. Algebra Comput.*, 25(5):813–854, 2015.

[9] J. W. Cannon. Almost convex groups. *Geom. Dedicata*, 22(2):197–210, 1987.

[10] R. Charney. Artin groups of finite type are biautomatic. *Math. Ann.*, 292(4):671–683, 1992.

[11] M. W. Davis and M. Shapiro. Coxeter groups are almost convex. *Geom. Dedicata*, 39(1):55–57, 1991.

[12] M. Elder. A linear-time algorithm to compute geodesics in solvable Baumslag-Solitar groups. *Illinois J. Math.*, 54(1):109–128, 2010.

[13] M. Elder and S. Hermiller. Minimal almost convexity. *J. Group Theory*, 8(2):239–266, 2005.

[14] D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston. *Word processing in groups.* Jones and Bartlett Publishers, Boston, MA, 1992.

[15] S. M. Gersten. Dehn functions and l_1-norms of finite presentations. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 195–224. Springer, New York, 1992.

[16] S. M. Gersten. Isoperimetric and isodiametric functions of finite presentations. In *Geometric group theory, Vol. 1 (Sussex, 1991)*, volume 181 of *London Math. Soc. Lecture Note Ser.*, pages 79–96. Cambridge Univ. Press, Cambridge, 1993.

[17] J. R. J. Groves. Minimal length normal forms for some soluble groups. *J. Pure Appl. Algebra*, 114(1):51–58, 1996.

[18] D. F. Holt and S. Rees. Shortlex automaticity and geodesic regularity in Artin groups. *Groups Complex. Cryptol.*, 5(1):1–23, 2013.

[19] R. C. Lyndon and P. E. Schupp. *Combinatorial group theory.* Classics in Mathematics. Springer-Verlag, Berlin, 2001. Reprint of the 1977 edition.

[20] J. Meier. The topology of graph products of groups. *Proc. Edinburgh Math. Soc. (2)*, 37(3):539–544, 1994.

[21] M. L. Mihalik and S. T. Tschantz. Tame combings of groups. *Trans. Amer. Math. Soc.*, 349(10):4251–4264, 1997.

[22] C. F. Miller, III and M. Shapiro. Solvable Baumslag-Solitar groups are not almost convex. *Geom. Dedicata*, 72(2):123–127, 1998.

[23] L. Mosher. Mapping class groups are automatic. *Ann. of Math. (2)*, 142(2):303–384, 1995.

[24] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.

[25] A. Nu'Man and T. Riley. Private communication, 2019.

[26] M. V. Sapir, J.-C. Birget, and E. Rips. Isoperimetric and isodiametric functions of groups. *Ann. of Math. (2)*, 156(2):345–466, 2002.

[27] J. Taback and A. Walker. Conjugation curvature in solvable Baumslag-Solitar groups. *J. Topol. Anal.*