UNIVERSITÄT DES SAARLANDES

Isometric maps between Outer Spaces

Dissertation
zur Erlangung des Grades
des Doktors der Naturwissenschaften
der Fakultät für Mathematik und Informatik
der Universität des Saarlandes

von

Christian Steinhart

Saarbrücken
2021

Tag des Kolloquiums:	01.02.2022
Dekan:	Prof. Dr. Thomas Schuster
Prüfungsausschuss:	Vorsitzender
	Prof. Dr. Mark Groves
	Berichterstatter
	Prof. Dr. Gabriela Weitze-Schmithüsen
	Prof. Dr. Hannah Markwig
	Prof. Dr. Armando Martino
	Akademische Mitarbeiterin
	Dr. Isabel Stenger

Acknowledgements

First and most of all I want to thank *Professor Gabriela Weitze-Schmithüsen* for supporting me throughout the whole thesis and giving me a lot of freedom and time while being always available for questions and discussions. Furthermore, I want to thank her for all the proofreading and helpful remarks.

Many thanks go to *Professor Armando Martino* for a very warm welcome in Southampton, for many helpful discussions and for accepting to become a referee of this theses. Also I want to thank the working group in Southampton for their hospitality and the discussions about Outer Space during and after my stays.

Also many thanks go to *Professor Hannah Markwig* for hosting me in Tübingen, her encouragement and support for my first paper and agreeing to be a referee for this thesis.

I would also like to thank all the other persons who have contributed to this thesis, in particular *Samuel Lelièvre* for helping me with my Sage problems and helping me to clean up my code and *Sven Caspart* for his thorough proofreading of the whole thesis and his LaTeX-support. I also want to thank *Benedikt Hewer* for his help with several IT problems and *Professor Karen Vogtmann* for allowing me to use her wonderful picture of CV_2.

Furthermore, I want to thank *Christine Wilk-Pitz* for all the organisatorial help and the rest of our working group as well as the working group *Groves* for a very open and welcoming atmosphere and the many discussions during coffee breaks and on the floor.

Finally, I want to thank my family for their support during the whole time and just for being there, in particular my nephew *Leopold Steinhart* for taking care that I do not oversleep whenever I visited.

This work was supported by the SFB-TRR 195 “Symbolic Tools in Mathematics and their Application” of the German Research Foundation (DFG).

Abstract

In this thesis we study the geometry of Culler–Vogtmann Outer Space CV_n with regard to the Lipschitz metric. We prove that the theorem of Francaviglia and Martino in [FM12b] about the isometry group of Outer Space also holds for reduced Outer Space, that is any isometry of CV_n^{red} comes from the $\mathrm{Out}(F_n)$-action on Outer Space. For this purpose we introduce *envelopes* in CV_n and use them to show how to construct a piecewise rigid geodesic between two points and all rigid geodesics emanating from a given point. Another application of these envelopes is the construction of a local geodesic which passes through a given sequence of points.

Furthermore, we introduce two families of isometric embeddings between Outer Spaces of different rank, which we call *naive embeddings* and *natural embeddings*. We show that natural embeddings from CV_n to CV_k exhibit some kind of rigidity for $n > 2$ while natural embeddings from CV_2 to CV_k can be deformed into other isometric embeddings.

Zusammenfassung

In dieser Dissertation untersuchen wir die Geometrie von Culler–Vogtmann Outer Space CV_n bezüglich der Lipschitz-Metrik. Francaviglia und Martino zeigten in [FM12b], dass jede Isometrie von CV_n bereits von der $\mathrm{Out}(F_n)$-Aktion kommt. Wir zeigen, dass dies auch für den reduzierten Outer Space gilt. Hierfür führen wir *Einhüllende* in CV_n ein und zeigen wie wir mit diesen eine stückweise starre Geodätische zwischen zwei Punkten konstruieren und alle starren Geodätischen die von einem Punkt ausgehen berechnen können. Als weitere Anwendung der Einhüllenden geben wir eine Konstruktion für eine lokal Geodätische an, die durch eine beliebige zuvor gegebene Folge von Punkten geht.

Weiterhin führen wir zwei Familien von isometrischen Einbettung zwischen Outer Spaces von verschiedenen Rängen ein, welche wir *naive Einbettungen* und *natürliche Einbettungen* nennen. Anschließend zeigen wir, dass natürliche Einbettungen von CV_n nach CV_k für $n > 2$ eine gewisse Art von Starrheit innehaben wohingegen natürliche Einbettungen von CV_2 nach CV_k zu weiteren isometrischen Einbettungen deformiert werden können.

Contents

Preface

Inspired by the study of the mapping class group via its action on Teichmüller Space (see [FM12a, Chapter 10]), Marc Culler and Karen Vogtmann introduced in 1986 in their paper [CV86] a topological space in order to study the outer automorphism group of the free group F_n. This space is nowadays known as Culler–Vogtmann Outer Space and denoted by CV_n. Outer Space can be described as an analogon of Teichmüller Space for finite, metric graphs. More precisely, a point in Outer Space can be written as a finite, metric graph without leaves and with a marking up to some equivalence. In this context a marking is a homotopy equivalence from the standard rose graph R_n, which is the graph with one vertex and n edges, to the finite, metric graph. In particular, we identify in this way the free group F_n with the fundamental group of the finite, metric graph. Two points in CV_n are identified, if there exists a homothety, that is an isometry up to scaling, between the two underlying marked, metric graphs which carries over the markings up to free homotopy.

Outer Space comes with a natural simplicial structure. Namely an (open) simplex of Outer Space consists of all marked, metric graphs which differ only by their edge lengths. For this description we use that we can normalise points in CV_n by scaling such that the edge lengths have sum one. By setting some of the edge lengths to zero, that is collapsing the corresponding edges, we obtain a face of the simplex. Observe that there are some faces missing as we are not allowed to collapse a loop (see Figure 4 for the simplicial structure of CV_2). A short calculation using the Euler characteristic shows that the maximal dimensional simplex of CV_n has dimension $3n-4$.

In analogy to the action of the mapping class group on Teichmüller Space the outer automorphism group $\mathrm{Out}(F_n)$ acts on CV_n by change of marking. Culler and Vogtmann showed that the $\mathrm{Out}(F_n)$-action on CV_n is properly discontinuous and $\mathrm{Out}(F_n)$ acts cocompactly on a contractible deformation retract of CV_n. Since then $\mathrm{Out}(F_n)$ and its action on Outer Space have been actively studied. For example Mladen Bestvina and Michael Handel introduced train tracks in [BH92] and proved Scott's conjecture, namely that every fixed subgroup of an automorphism of F_n has at most rank n. Bestvina, Handel and Mark Feighn showed in [BFH00, BFH05] that the Tits alternative holds for $\mathrm{Out}(F_n)$. The bordification of Outer Space was introduced by Bestvina and Feighn in [BF00], to mention only some of the work. A good place to start reading are the surveys [Vog02] from Vogtmann and [Bes02] from Bestvina.

Similarly to the asymmetric Thurston metric in Teichmüller Space introduced by William Thurston in [Thu98], Stefano Francaviglia and Armando Martino introduced in [FM11] an *asymmetric Lipschitz metric* for Outer Space. Both metrics can be described in terms of the minimal Lipschitz constant of change of marking maps. Equivalently they can be described by the supremal stretching of loops (see [Thu98, Theorem 8.5] for the Thurston metric and [FM11, Proposition 3.15] for CV_n).

While for the Thurston metric the supremum of the stretching of simple loops might not be attained, Francaviglia and Martino showed that in Outer Space the maximal stretching is always attained by an element of a finite set of loops called the *candidates*, which depend only on one of the marked graphs. Consequently, given two points in Outer Space we can algorithmically calculate their distance as described in Section 1.6. We have implemented the described algorithm in Sage [Sag] and the implementation can be found in [Ste18] and in the Appendix.

It follows from the definition of the Lipschitz metric in terms of maximal stretchings of loops that the action of $\mathrm{Out}(F_n)$ on CV_n is by isometries. Francaviglia and Martino showed in [FM12b] that all isometries of CV_n with regard to the asymmetric Lipschitz metric d_R and its symmetrised version $d(A, B) := d_R(A, B) + d_R(B, A)$ come from the action of an element in $\mathrm{Out}(F_n)$. The analogue statment for Teichmüller space was proven up to a few sporadic cases in [Wal14] and for the once punctured torus in [DLRT20].

To name two applications for the Lipschitz metric, observe that we have for an automorphism $\phi \in \mathrm{Out}(F_n)$ the displacement function $A \mapsto d_R(A, A \cdot \phi)$. This displacement function was used by Bestvina [Bes11] to give an alternative proof that every irreducible automorphism can be represented as train track map and by Francaviglia and Martino in [FM21] to reprove an algorithm to solve the conjugacy problem for irreducible automorphisms. For a survey about the Lipschitz metric of Outer Space we refer the reader to [Vog15].

The goal of this thesis is to continue the study of the geometry of Outer Space with regard to the Lipschitz metric. Our emphasis lies on isometric maps between Outer Spaces. We will first study the isometry group of reduced Outer Space. Afterwards we turn to the question, how isometric embeddings between Outer Spaces of different rank look like and whether they inherit similar rigidity properties as isometries. Furthermore, we obtain results about rigid geodesics and local geodesics.

Reduced Outer Space is a deformation retract of CV_n which consists of the marked, metric graphs in CV_n without separating edges. We will show in Theorem 3.9 that the isometry group of reduced Outer Space equals the isometry group of Outer Space. In order to prove Theorem 3.9 we want to apply the proof of Francaviglia and Martino for non-reduced Outer Space described in [FM12b]. The missing step is to show that the isometries of reduced Outer Space are simplicial. We will prove this in Theorem 3.8. The proof of this theorem and the needed tools are published as preprint in [Ste19].

The crucial tool in the proof of Theorem 3.8 are so called *envelopes* introduced in Section 2, which are the sets of all points lying on geodesics between two given points. We will see that an envelope is always a compact subset of CV_n, whose intersection with a closed simplex of CV_n is a polytope. In particular, the intersection of an envelope with a simplex has a well-defined dimension. However the dimension can differ from simplex to simplex (see Figure 13).

We will introduce the notion of *general position* of two points (see Definition 3.4) and show in Lemma 3.7 that two points in the same simplex and in general position to each other have in their simplex a full dimensional envelope. In contrast, for any open subset U intersecting at least two simplices there exist two open subsets U_A and U_B of U such that for any pair $A \in U_A$ and $B \in U_B$ their envelope has lower dimension near B. Since envelopes are preserved under isometries we can conclude that the $(3n-5)$-skeleton of CV_n is preserved under isometries. To see that an isometry preserves also the skeleton of lower dimension, observe that any simplex of dimension lower than $3n-5$ is the facet of more than two simplices. Hence, we get by induction that an isometry preserves all skeletons and is in particular simplicial.

Instead of an envelope between two points, we can also consider envelopes with a starting point and a *coarse direction*, which yields the notion of *in-envelopes* and *out-envelopes* (see Definition 2.11). An in-envelope of a point $B \in CV_n$ along a coarse direction $S \subset F_n$ consists of all points $A \in CV_n$ such that all loops in S are maximally stretched from A to B. Accordingly an out-envelope of a point $A \in CV_n$ along $S \subset F_n$ consists of all points $B \in CV_n$ such that all loops in S are maximally stretched from A to B. That is A lies in the in-envelope of B along S if and only if B lies in the out-envelope of A along S. As in the case of envelopes their intersection with simplices are again polytopes and can be explicitly computed.

As it turns out envelopes are also a useful tool to understand rigid geodesics. Here we call a geodesic *rigid* if each subpath is the unique geodesic joining its two endpoints. Rigid geodesics are quite rare in Outer Space, namely for any given point there exists only finitely many rigid geodesics in a simplex passing through this point. While a geodesic between two points in Outer Space is typically not unique, we will show in Theorem 2.14 how to construct for any two points a piecewise rigid geodesic between them by going along the edges of envelopes.

In Theorem 2.21 we fully classify rigid geodesics. Namely, we show that all rigid geodesics can be written as concatenation of edges of in- and out-envelopes. Hence, we can compute all rigid geodesic emanating from a given point.

In Section 4 we will briefly discuss local geodesics in Outer Space. As another application of envelopes we will see in Theorem 4.4 that given a sequence of points in CV_n we can construct a local geodesic passing through them. Furthermore, we will show in Proposition 4.6 that we can approximate a given path in Outer Space arbitrarily well by local geodesics with regard to the symmetric and asymmetric metric.

We will discuss isometric maps between Outer Spaces of different rank in Section 5 and show that there exist continuous families of isometric embeddings of CV_n into CV_k for arbitrary $n < k$. Since we have seen in Section 3 that the simplicial structure is determined by the Lipschitz metric, it is natural to ask, if all isometric embeddings are simplicial. We will give examples which show that this is not always the case.

The isometric embeddings we will construct can be divided into the following two types:

- The *naive embeddings*: They correspond to identifications of F_n with a free factor of F_k and come from attaching a graph with fundamental group F_{k-n} to the points in CV_n.
- The *natural embeddings*: They come from finite coverings and correspond to identifying F_k with a finite index subgroup of F_n. By the Nielsen-Schreier formula we only have natural embeddings for $k = 1 + d(n-1)$ with $d \in \mathbb{N}$.

While we inductively construct a naive isometric embedding for $n = 2$ in Section 5.1, we need for higher rank a coherent choice of a base point where we attach the graph.

We will discuss two ways how to get such a choice of a base point and the resulting embeddings in Section 5.2. For the first we make use of translation axes. The second is due to Skora [Sko90] and generalises the length function to a set of loops. It should be noted here that only the first choice of base points leads to an isometric embedding, while for the second we obtain isometric embeddings from finite subcomplexes of CV_n into CV_k.

We will see in Section 5.2 that from a naive embedding we get a continuous family of isometric embeddings by slightly deforming the attached graph. We will discuss in Section 6 whether we can also deform natural embeddings in a similar manner to obtain new isometric embeddings. As it turns out, all the natural embeddings from CV_2 to a CV_k with $k > 2$ can be deformed. On the other hand the natural embeddings from CV_n have for $n \geq 3$ some sort of rigidity: We will see in Theorem 6.8 that if an isometric embedding differs only in a bounded subset from such a natural embedding, then they already have to be equal.

1 Preliminaries

1.1 What is Outer Space?

This section will give a quick introduction into the basic definitions and properties of Culler–Vogtmann Outer Space introduced in [CV86]. The section mainly follows the surveys [Vog15] and [Vog02]. Typically a point in Outer Space CV_n consists of three data, a finite graph Γ without leaves, a marking on Γ and the lengths of its edges.

Definition 1.1 (i) For $n \in \mathbb{N}$ the *rose* R_n is the graph with one vertex and n edges, also called petals. We identify the fundamental group $\pi_1(R_n, \star)$ with the free group F_n generated by $x_1, \ldots, x_n$ by assigning each (oriented) petal an element x_i.

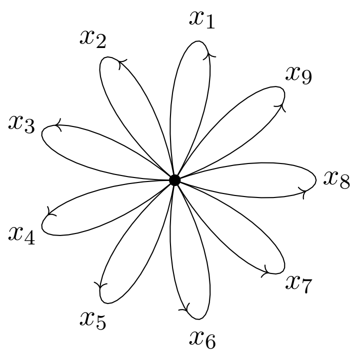

Figure 1: The rose graph R_9 with labelled petals

(ii) Let Γ be a finite graph where each vertex has at least valency three. A *marking on* Γ is a homotopy equivalence $m\colon R_n \to \Gamma$. We use the induced isomorphism $m_*\colon F_n = \pi_1(R_n, \star) \to \pi_1(\Gamma, m(\star))$ to identify F_n with the fundamental group $\pi_1(\Gamma, m(\star))$. As we will be interested only in the homotopy class of the marking, we will from now on omit the base points in the fundamental groups and use this identification to read conjugacy classes of elements of F_n as reduced cycles in Γ (see Notation 1.2 (ii)).

(iii) Let Γ be a graph and let $E(\Gamma)$ be the edges of Γ. A function

$$l : E(\Gamma) \to \mathbb{R}_{>0}$$

which assigns to each edge $e \in E(\Gamma)$ a positive real length $l(e)$ is called a *length function* and the pair (Γ, l) is called a *metric graph.* We can think of (Γ, l) as the metric simplicial 1-complex obtained by glueing intervals $[0, l(e)]$ glued together at their endpoints. The *volume* of a metric graph is defined as the sum over all edge lengths, that is

$$\mathrm{vol}(\Gamma, l) := \sum_{e \in E(\Gamma)} l(e).$$

We call a metric graph *normalised* if it has volume 1.

(iv) The *(projectivised) Outer Space of rank n* is defined as the set of marked, metric graphs (Γ, l, m) together with some equivalence

$$\begin{aligned} CV_n := \{(\Gamma, l, m) \mid \Gamma & \text{ is a finite graph with all vertices having at least valency 3} \\ & \text{together with a length function } l : E(\Gamma) \to \mathbb{R}_{>0} \\ & \text{and a marking } m : R_n \to \Gamma\}/\sim \end{aligned}$$

where two points $(\Gamma, l, m), (\Gamma', l', m')$ are equivalent $(\Gamma, l, m) \sim (\Gamma', l', m')$, if there exists a homothety $h\colon (\Gamma, l) \to (\Gamma', l')$ such that the induced marking $h \circ m$ is homotopic to m', that is the following diagram commutes up to homotopy:

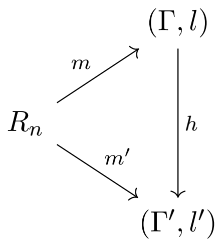

(v) The *reduced Outer Space* CV_n^{red} *of rank n* is the subset of CV_n where the graphs have no separating edges, i.e. a point in CV_n^{red} is a 2-edge-connected metric graph with a marking.

To rephrase the above definition, a point in CV_n can be represented as a finite metric graph Γ without leaves together with a fixed identification of the free group F_n with its fundamental group $\pi_1(\Gamma)$.

Notation 1.2 (i) Conventionally, the marking is written down by a homotopy inverse $\tilde{m}$ in the following way: Fix a spanning tree of Γ which will be collapsed by $\tilde{m}$ to the vertex of R_n. The rest of the edges are labelled with an orientation and a basis of F_n depending on the marking. Each edge will then be sent by $\tilde{m}$ to the sequence of petals in R_n corresponding to its label (see Figure 2).

Keep in mind that such a homotopy inverse only defines the marking up to homotopy. Since two homotopic markings are identified under the equivalence in CV_n, the homotopy inverse defines the marking of a point in CV_n.

(ii) For a marked graph (Γ, m) we can identify each element of F_n with its image under the marking in $\pi_1(\Gamma)$ and vice versa. We call a cycle $\alpha : S^1 \to \Gamma$ *cyclically reduced* if the map α is an immersion. Observe that each cycle in Γ is homotop to a cyclically reduced loop by iteratively retracting the edges where α is not locally injective. This cyclically reduced loop is unique up to change of the start- and endpoint and

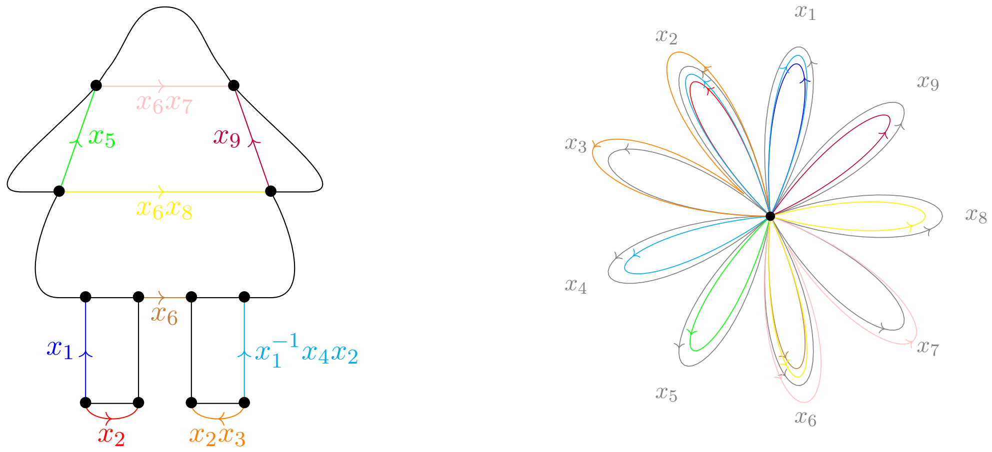

Figure 2: A typical element in Outer Space and its homotopy inverse image in R_9

parametrisation. As we are only interested in the marking up to homotopy we will typically identify each (conjugacy class of an) element $\alpha \in F_n$ with a cyclically reduced representation of $m(\alpha)$.

(iii) Let $A := (\Gamma, l, m) \in CV_n$ be a point in Outer Space and $\alpha \in F_n$. As in (ii) the element α has a unique cyclically reduced realization. We will denote its length by $l_A(\alpha)$.

There is an alternative definition of Outer Space as actions of F_n on metric trees, which we will use in Section 5.2 and Section 5.3. This definition comes from the action of the fundamental group of Γ on its universal covering. To distinguish the two definitions for Outer Space here we will denote Outer Space in the latter viewpoint as X_n.

Definition 1.3
A point in Outer Space X_n is a metric, simplicial tree (T, l) together with a free, minimal action $m : F_n \to \mathrm{Isom}(T, l)$ up to some equivalence. Here free means that the action has only trivial stabilisers and minimal means that there exist no non-trivial m-invariant subtree.

We identify two points $(T, l, m), (T', l', m') \in X_n$, if there exists an F_n-equivariant homothety $h\colon T \to T'$, i.e. for all $\alpha \in F_n$ we have $m'(\alpha) \circ h = h \circ m(\alpha)$.

It is easy to check that the map from CV_n to X_n which sends a metric graph Γ to its universal cover T and its fundamental group to its deck transformation group is a bijection. A proof for this statement can for example be found in [AK19, Section 2.3]. The length of an element $\alpha \in F_n$ for some $(T, l, m) \in X_n$ is its minimal displacement of points in T by m. By [CM87] the length of some non-trivial element $\alpha \in F_n \setminus \{\mathrm{id}\}$ is exactly its translation distance along its translation axis:

Lemma 1.4 ([CM87, 1.3])
Let $(T, l, m) \in X_n$ be a point in Outer Space and let $\alpha \in F_n \setminus \{\mathrm{id}\}$ be a non-trivial element. We denote with

$$l_{(T,l,m)}(\alpha) := \inf_{p \in T} d_T(p, \alpha \cdot p)$$

the *length of* α in (T, l, m) and by

$$T_\alpha := \{p \in T \mid d_T(p, \alpha \cdot p) = l_{(T,l,m)}(\alpha)\}$$

its *characteristic set*. We have then that T_α is an embedding of the real line and α acts on T_α by translation. We call the set T_α the *translation axis* of α.

Furthermore we have for any $p \in T$ the equality $d_T(p, \alpha \cdot p) = l_{(T,l,m)}(\alpha) + 2d_T(p, T_\alpha)$. Observe that the translation axis does not depend on the lengths of the edges but only on the tree and the action m.

1.2 The topology of Outer Space

There is a natural topology on Outer Space defined as follows.

The length function gives an embedding into the projective space of the product of $\mathbb{R}$ over F_n:

$$l : CV_n \to \mathbf{P}(\mathbb{R}^{F_n}) \quad , \quad A = (\Gamma, l, m) \mapsto \{l_A(\alpha)\}_{\alpha \in F_n}$$

where the injectivity follows from the fact that minimal, free isometric actions of F_n on trees are determined by their length function ([AB16] Proposition 7.13 (ii)). Alternatively, the injectivity also follows from the Lipschitz metric see Definition 1.10. Hence, we can give CV_n the subspace topology of $\mathbf{P}(\mathbb{R}^{F_n})$, where we endow $\mathbb{R}^{F_n}$ with the product topology.

Alternatively we can describe the topology from the fact that we can build CV_n as simplicial complex with some missing faces. Each simplex will correspond to a marked graph where we vary edge lengths. As each simplex is a face of at most finitely many simplices, it follows from Remark 1.6 that these two descriptions are indeed equivalent.

Definition 1.5
The data (Γ, m) of an element $A = (\Gamma, l, m) \in CV_n$ is called the *topological type* of A. We will denote this from now on by $\Delta(A)$.

After normalising the volume we will see in Remark 1.6 that points of the same topological type correspond to an open simplex in CV_n. Hence, we will also denote by $\Delta(A)$ the corresponding open simplex in CV_n.

Remark 1.6
Let (Γ, m) be a marked graph with edges $E(\Gamma) = \{e_1, \ldots, e_k\}$. For the open simplex $\Delta = \{(l_1, \ldots, l_k) \mid l_1, \ldots, l_k \in \mathbb{R}_{>0}, \sum_{i=1}^k l_i = 1\}$ we consider the map

$$\begin{aligned} \iota : \Delta &\to CV_n, \\ (l_1, \ldots, l_k) &\mapsto (\Gamma, m, l) \quad \text{with } l(e_i) = l_i \text{ for all } i \in \{1, \ldots k\}. \end{aligned}$$

Then ι is a homeomorphism onto its image.

Proof. It is clear that ι is continuous as the lengths of cycles continuously depend on the lengths of the edges. From Lemma 5.7 it will follow directly that ι is also injective and open. □

For each topological type we then get by Remark 1.6 an embedding from an open simplex to CV_n. Their images cover the whole Outer Space. The dimension of a simplex corresponding to a marked graph is one less than the number of edges in the marked graph. That means that maximal dimensional simplices correspond to the 3-regular, marked graphs and by the Euler characteristic such a simplex has dimension $3n - 4$. Accordingly if a simplex has lower dimension, then its corresponding marked graph has a vertex of valency at least four.

The simplices are glued together in the following natural way. If we can pass from one marked graph (Γ, m) to another marked graph (Γ', m') by collapsing a forest (see Figure 3), then we identify the open simplex of (Γ', m') with the missing face of the simplex of (Γ, m) were the edges of the forest have length 0.

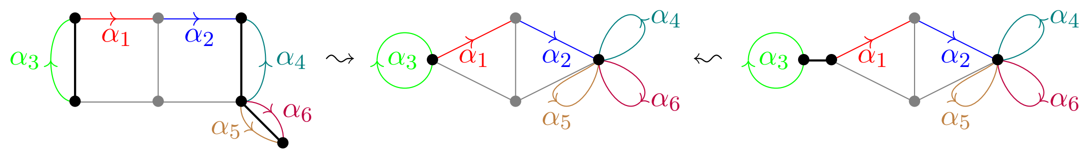

Figure 3: Collapsing the black forest

Hence, we can represent CV_n as a simplicial complex (see Figure 4) with some missing faces. This will induce the same topology on CV_n as the length function. By contracting separating edges in each graph we get that CV_n^{red} is a deformation retract of CV_n. In Figure 4 this corresponds to retracting the yellow fins into the blue plane.

To denote how far or close we are to such a missing face we introduce the notion of thinness and thickness. Here a graph is thin if it has a loop of a very short length compared to its volume. Accordingly if all loops have at least some certain length we call it thick. By saying we go to the thin part of CV_n we mean that we gradually decrease the length of a loop to be arbitrarily short.

Definition 1.7
Let $\varepsilon > 0$ and $A := (\Gamma, l, m) \in CV_n$ be a point in Outer Space. We say A is ε-*thin* if there exists an $\alpha \in F_n \setminus \{\text{id}\}$ such that $l_A(\alpha)/\operatorname{vol}(A) < \varepsilon$, otherwise we call A ε-*thick*. That means a normalised graph is always as thick as its shortest loop.

The following two important facts about Outer Space have been proven by Culler and Vogtmann:

Theorem 1.8 (Culler–Vogtmann [CV86])
CV_n is contractible.

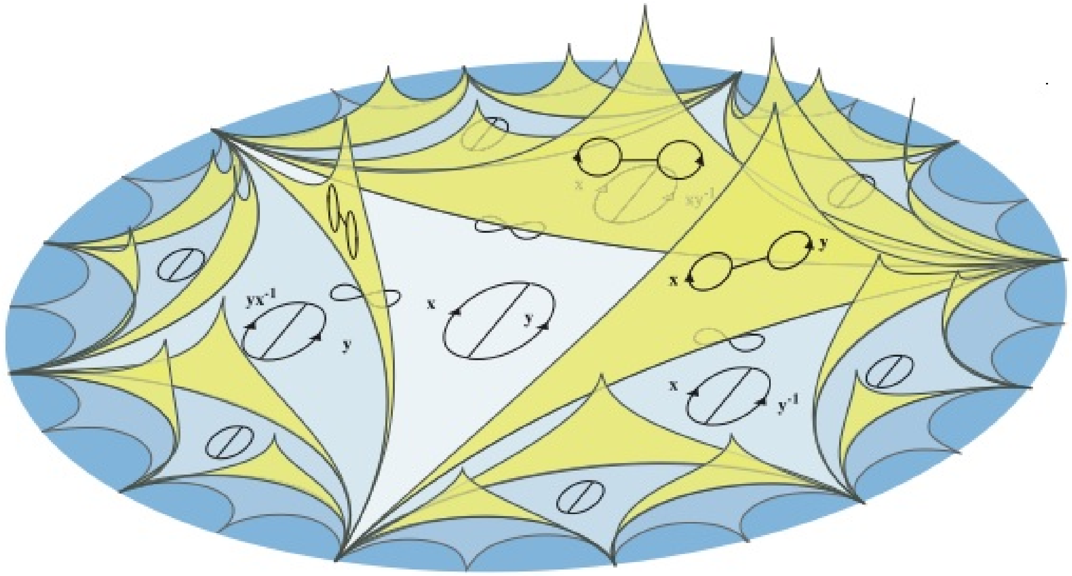

Figure 4: A picture of CV_2 by Karen Vogtmann [Vog15]

On CV_n we have a natural right action of $\mathrm{Aut}(F_n)$ by changing the marking. Each element $\phi \in \mathrm{Aut}(F_n)$ can be realised as a homotopy equivalence $\phi' : R_n \to R_n$ and precomposing to the marking yields the action on CV_n: $(\Gamma, l, m) \bullet \phi := (\Gamma, l, m \circ \phi')$. Furthermore, inner automorphisms act trivially on CV_n since the marking is only defined up to homotopy, hence we actually have an action of $\mathrm{Out}(F_n) := \mathrm{Aut}(F_n)/\mathrm{Inn}(F_n)$.

Theorem 1.9 (Culler–Vogtmann [CV86])
The $\mathrm{Out}(F_n)$-action on CV_n is fix-point free and each point has a finite stabiliser.

Although this action is not cocompact on CV_n, Culler and Vogtmann showed that it acts cocompactly on the spine K_n of Outer Space. The spine is a simply connected deformation retract and is the barycentric decomposition of CV_n without the simplices that would hit the missing faces of CV_n.

1.3 The metric on Outer Space

Similarly to the Thurston metric on Teichmüller space as introduced in [Thu98], Francaviglia and Martino introduced an asymmetric metric on CV_n in [FM11]. The content of this section is from their paper [FM11].

Definition 1.10
Let $A := (\Gamma, l, m)$ and $B := (\Gamma', l', m') \in CV_n$ be two points in CV_n. We call a continuous map $h\colon (\Gamma, l) \to (\Gamma', l')$ a *change of marking map* from A to B if we have $h \circ m \cong m'$,

that is the following diagram commutes up to homotopy:

$$\begin{array}{ccc} R_n & \xrightarrow{m} & \Gamma \\ & \searrow^{m'} & \downarrow h \\ & & \Gamma' \end{array}$$

Consider the set S of all change of marking maps h from A to B. Since finite metric graphs are compact, each $h \in S$ is Lipschitz continuous with Lipschitz constant $L(h)$. We call then

$$\Lambda_R(A, B) := \inf_{h \in S} L(h)$$

the *stretching factor* from A to B and define the *Lipschitz distance* from A to B as

$$d_R(A, B) := \log \left(\Lambda_R(A, B) \cdot \frac{\mathrm{vol}(A)}{\mathrm{vol}(B)} \right)$$

The symmetric Lipschitz distance is defined as the sum

$$d(A, B) := d_R(A, B) + d_R(B, A).$$

The Arzela-Ascoli theorem yields that the infimum $\Lambda_R(A, B)$ is actually attained by a map $h \in S$. Thus, it is clear that d_R satisfies the triangle inequality. As any surjective 1-Lipschitz map between two volume 1 graphs has to be already an isometry we have that $d_R(A, B) = 0$ implies $A = B$. Example 1.13 shows that d_R is not symmetric. Hence, we have that d_R is an asymmetric metric on CV_n and d is a symmetric metric on CV_n.

Similar to the Thurston metric we can calculate Λ_R as the supremal stretching of paths in the graph. By Proposition 1.12 from [FM11] it turns out that in Outer Space this supremum is always attained and can explicitly be calculated by the stretching of a finite set of paths, namely the candidates of a graph which are defined as follows:

Definition 1.11
For a given point $A \in CV_n$ a *candidate of A* is a cycle in A whose image is a topological embedding of a simple loop, a figure of eight or a barbell (see Figure 5). We will denote the set of candidates of A with $\mathrm{cand}(A)$ and identify it via the marking as the corresponding subset (of conjugacy classes) in F_n. Clearly the candidates of A only depend on the topological type $\Delta(A)$. As we have seen in Remark 1.6 a topological type corresponds to a simplex $\Delta \subset CV_n$, hence we will write accordingly $\mathrm{cand}(\Delta) := \mathrm{cand}(A)$ for any $A \in \Delta$.

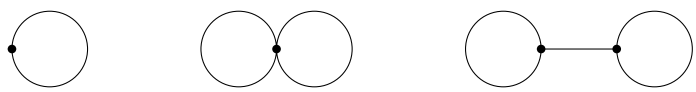

Figure 5: A simple loop, a figure of eight and a barbell

Proposition 1.12 (Francaviglia–Martino [FM11, Proposition 3.15])
For $A, B \in CV_n$ we have:

$$\Lambda_R(A,B) = \sup_{\alpha \in F_n} \frac{l_B(\alpha)}{l_A(\alpha)} = \max_{\alpha \in \mathrm{cand}(A)} \frac{l_B(\alpha)}{l_A(\alpha)}$$

Since there are only finitely many candidates in a finite graph, we can compute $\Lambda_R(A, B)$. The algorithm for doing this is described in Section 1.6. An implementation for Sage [Sag] can be found under [Ste18]. Furthermore this identification also shows that the topology from d_R is the topology on CV_n described in Section 1.2.

Example 1.13 (i) As an example that d_R is not symmetric, consider the two metric graphs

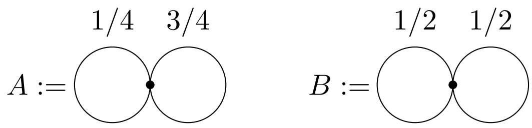

with the indicated edge lengths and the same marking. Observe that the candidates in A are up to orientation the two simple loops and the two figures of eight, namely the figure of eight with both loops in clockwise direction and the figure of eight with one loop in clockwise and the other in anticlockwise direction. As the figures of eight all have length 1, we have as stretching factors

$$\Lambda_R(A,B) = \max\left\{\frac{1/4}{1/2}, \frac{3/4}{1/2}, 1\right\} = 3/2$$
$$\Lambda_R(B,A) = \max\left\{\frac{1/2}{1/4}, \frac{1/2}{3/4}, 1\right\} = 2.$$

(ii) As a second example involving the thin part of Outer Space let $A, B \in CV_n$ be of the same topological type and let A be ε-thin and B be δ-thick for some $\delta \gg \varepsilon > 0$. After normalising we have $\Lambda_R(A, B) \geq \frac{\delta}{\varepsilon}$ and $\Lambda_R(B, A) \leq \frac{1}{\delta}$. That means "It's easy to get to the boundary of Outer Space, but really hard to come back".

As we will frequently compare the stretching of different candidates we will regularly use the following easy but very useful inequality:

Lemma 1.14
Let $a, b, c, d \in \mathbb{R}$ and $c, d > 0$, then we have

$$\min\left\{\frac{a}{c}, \frac{b}{d}\right\} \leq \frac{a+b}{c+d} \leq \max\left\{\frac{a}{c}, \frac{b}{d}\right\}$$

Furthermore we have $\frac{a+b}{c+d} = \max\{\frac{a}{c}, \frac{b}{d}\}$ if and only if $\frac{a}{c} = \frac{b}{d}$.

Proof. By symmetry we assume $\frac{a}{c} \geq \frac{b}{d}$. We have

$$\begin{aligned}\frac{a}{c} &= \frac{a\cdot(c+d)}{c\cdot(c+d)} = \frac{ac+ad}{c\cdot(c+d)} = \frac{a}{c+d} + \frac{a}{c}\cdot\frac{d}{c+d}\\ &\geq \frac{a}{c+d} + \frac{b}{d}\cdot\frac{d}{c+d} = \frac{a}{c+d} + \frac{b}{c+d} = \frac{a+b}{c+d}\\ &= \frac{a}{c}\cdot\frac{c}{c+d} + \frac{b}{c+d} \geq \frac{b}{d}\cdot\frac{c}{c+d} + \frac{b}{c+d} = \frac{bc+bd}{d\cdot(c+d)} = \frac{b}{d},\end{aligned}$$

where the inequalities are equalities if and only if $\frac{a}{c} = \frac{b}{d}$ holds. □

As a direct application of this inequality we get the following:

Corollary 1.15
Let $A, B \in CV_n$ have the same topological type. Let furthermore $\alpha\beta$ be a figure of eight in A with loops α and β. Then $\alpha\beta$ is exactly then maximally stretched from A to B if also α and β are maximally stretched.

Therefore it is only interesting to consider the stretching of a figure of eights if the two points lie in different simplices of CV_n.

Definition 1.16
We say $\alpha \in F_n$ is a *witness* from A to B for the asymmetric metric, if α is maximally stretched from A to B, i.e. if $\Lambda_R(A,B) = \frac{l_B(\alpha)}{l_A(\alpha)}$ holds. We denote the set of witnesses from A to B as $W_R(A,B)$.

Most of the time the witnesses we consider will be candidates of A. Such witnesses will be called *candidate witnesses*. The set of candidate witnesses is denoted by $CW_R(A,B) := \text{cand}(A) \cap W_R(A,B)$.

A pair $(\alpha,\beta) \in F_n \times F_n$ is called a *(symmetric) witness* for A and B for the symmetric metric, if α is an asymmetric witness from A to B and β is an asymmetric witness from B to A. We denote the set of symmetric witnesses for A and B as $W(A,B)$ and the set of symmetric candidate witnesses as $CW(A,B) := W(A,B) \cap (\text{cand}(A) \times \text{cand}(B))$.

As all conjugates of an element $\alpha \in F_n$ have the same length and for powers of α we have $l_A(\alpha^k) = |k| \cdot l_A(\alpha)$ we will from now on assume that a given witness is a simple element in F_n and in particular it has no proper roots. Accordingly we will ignore conjugated and inverted candidates, that is for $(\alpha\beta) \in \text{cand}(A)$ we will consider $\beta\alpha$ and $\beta^{-1}\alpha^{-1}$ to be the same candidate as $\alpha\beta$.

By Proposition 1.12 the set of candidate witnesses and in particular the set of witnesses is never empty. We will see that in most of the cases there is only one candidate witness.

As we will see in the next section, witnesses play a crucial role in the study of geodesics and can be considered as coarse directions of geodesics.

1.4 Geodesics in Outer Space

To understand a metric space it is quite often important to understand its geodesics. For example in [FM12b] Francaviglia and Martino used rigid geodesics in CV_n to determine the isometry group of a simplex corresponding to a multitheta-graph and thereby the isometry group of CV_n. We will also see in Section 3 that the simplicial structure of CV_n^{red} is encoded in its geodesics. This section contains the basic definitions and properties of geodesics, which we will use in the later sections. Up to Lemma 1.25 and its corollary the following statements are well known and can for example be found in [FM11].

As in spaces with a symmetric metric one defines length and geodesic paths in regard to an asymmetric metric:

Definition 1.17
Let (X, d) be a space with an asymmetric metric d, $I \subseteq \mathbb{R}$ an interval and $\gamma\colon I \to X$ a path.

(i) Let $J := [s, t] \subseteq I$ be a closed subinterval. Then the length of the arc $\gamma|_J$ is defined as the supremum

$$l(\gamma|_J) := \sup \left\{ \sum_{i=1}^{N} d(\gamma(t_{i-1}), \gamma(t_i)) \mid N \in \mathbb{N}, s =: t_0 \leq t_1 \leq \cdots \leq t_N := t \right\}.$$

We call γ *rectifiable* if and only if the arc length is finite for every arc of γ.

(ii) A rectifiable curve γ is called a *geodesic* if and only if for every arc $J := [s, t] \subseteq I$ we have $l(\gamma|_J) = d(\gamma(s), \gamma(t))$.

Be aware that - as d is not symmetric - these definitions highly depend on the orientation of γ. For example we can have $d(\gamma(s), \gamma(t)) \neq d(\gamma(t), \gamma(s))$. In particular, if γ is a d_R-geodesic in CV_n, we can still have that $\overline{\gamma}(t) := \gamma(-t)$ is not a d_R-geodesic. This is for example the case in Lemma 1.25. We can alternatively characterise geodesics by the triangle equality:

Lemma 1.18 ([FM11, Lemma 5.1])
Let (X, d) be a space with an asymmetric metric d, $I \subseteq \mathbb{R}$ an interval and $\gamma\colon I \to X$ a path. Then we have that γ is a geodesic if and only if it realises the triangle equality

$$d(\gamma(s), \gamma(t)) = d(\gamma(s), \gamma(r)) + d(\gamma(r), \gamma(t))$$

for all $s, t, r \in I$ with $s \leq r \leq t$.

Proof. Let $s \leq r \leq t \in I$.
"$\Rightarrow$": By the triangle inequality of the metric we have

$$\begin{aligned} d(\gamma(s), \gamma(t)) &\leq d(\gamma(s), \gamma(r)) + d(\gamma(r), \gamma(t)) \\ &\leq l(\gamma|_{[s,t]}) = d(\gamma(s), \gamma(t)) \end{aligned}$$

where the last inequality is the definition of length and the last equality comes from the definition of geodesic. Hence, equality holds.

"$\Leftarrow$": Let $s =: t_0 \leq \cdots \leq t_N := t$ be any subdivision of $[s,t]$. Then iteratively applying the triangle equality yields

$$\sum_{i=1}^{N} d(\gamma(t_{i-1}), \gamma(t_i)) = d(\gamma(t_0), \gamma(t_2)) + \sum_{i=3}^{N} d(\gamma(t_{i-1}), \gamma(t_i)) = \cdots = d(\gamma(t_0), \gamma(t_N))$$

and hence $l(\gamma|_{[s,t]}) = d(\gamma(s), \gamma(t))$ □

Remark 1.19
The definition of a geodesic differs slightly in different contexts, for example we do not require here a geodesic to be parametrised by length. It is also quite common to require a geodesic to only locally minimise the distance and not globally as we do.

The reason we consider only globally minimising geodesics is that local geodesics might "change directions" in CV_n and behave not the way you normally expect from geodesics. For example any path in CV_n can be approximated arbitrarily close with a locally minimising geodesic. We will discuss this shortly in Section 4.

To emphasise with respect to which metric of CV_n a geodesic is considered we will stick to the following notation:

Notation 1.20
From now on a geodesic in CV_n means a geodesic with respect to the asymmetric metric d_R (see Definition 1.10) and a *symmetric geodesic* is a geodesic with respect to the symmetric metric d.

Observe that symmetric geodesics are exactly the paths which are geodesics independent of the orientation. Namely by Lemma 1.18 and Definition 1.10 of the symmetric metric we have the following corollary:

Corollary 1.21
Let $\gamma\colon I \to CV_n$ be a path and let $\overline{\gamma}(t) := \gamma(-t)$ be the path γ with reversed orientation. Then γ is a symmetric geodesic if and only if γ and $\overline{\gamma}$ are geodesics.

As mentioned before witnesses play the role of coarse directions of geodesics in the sense that witnesses are preserved along geodesics and encode how to continue a geodesic.

Lemma 1.22
Let $A, B \in CV_n$, $\gamma\colon I \to CV_n$ be a geodesic from A to B and $t \in I$. Then we have $W_R(A,B) = W_R(A,\gamma(t)) \cap W_R(\gamma(t), B)$.

Proof. Without loss of generality we assume that all representants are normalised to have volume 1. By Lemma 1.18 γ satisfies the triangle equality and we have for all $t \in I$:

$$\begin{aligned}
\alpha \in W_R(A,B) &\iff \log\left(\frac{l_B(\alpha)}{l_A(\alpha)}\right) = d_R(A,B) \\
&\iff \log\left(\frac{l_{\gamma(t)}(\alpha)}{l_A(\alpha)}\right) + \log\left(\frac{l_B(\alpha)}{l_{\gamma(t)}(\alpha)}\right) = d_R(A,\gamma(t)) + d_R(\gamma(t),B) \\
&\overset{(3)}{\iff} \log\left(\frac{l_{\gamma(t)}(\alpha)}{l_A(\alpha)}\right) = d_R(A,\gamma(t)) \text{ and } \log\left(\frac{l_B(\alpha)}{l_{\gamma(t)}(\alpha)}\right) = d_R(\gamma(t),B) \\
&\iff \alpha \in W_R(A,\gamma(t)) \cap W_R(\gamma(t),B)
\end{aligned}$$

Where we used in the implication (3) that the inequalities

$$\log\left(\frac{l_{\gamma(t)}(\beta)}{l_A(\beta)}\right) \leq d_R(A,\gamma(t)) \text{ and } \log\left(\frac{l_B(\beta)}{l_{\gamma(t)}(\beta)}\right) \leq d_R(\gamma(t),B)$$

hold for any $\beta \in F_n$. □

The converse is also true, namely any path preserving a witness is already a geodesic:

Lemma 1.23
Let $A, B \in CV_n$ and $\gamma\colon I \to CV_n$ be a path from A to B. Then γ is a geodesic if and only if there exists an $\alpha \in W_R(A,B)$ such that for all $s,t \in I$ with $s \leq t$ we have $\alpha \in W_R(\gamma(s),\gamma(t))$.

Proof. If γ is a geodesic then Lemma 1.22 implies $W_R(\gamma(s),\gamma(t)) \supseteq W_R(A,\gamma(t)) \supseteq W_R(A,B)$. By Proposition 1.12 $W_R(A,B)$ has at least one element $\alpha \in W_R(A,B)$.
On the other hand if $\alpha \in W_R(\gamma(s),\gamma(t))$ for all $s \leq t \in I$, then we have after normalising for all $s \leq r \leq t \in I$

$$\begin{aligned}
d_R(\gamma(s),\gamma(r)) + d_R(\gamma(r),\gamma(t)) &= \log\left(\frac{l_{\gamma(r)}(\alpha)}{l_{\gamma(s)}(\alpha)}\right) + \log\left(\frac{l_{\gamma(t)}(\alpha)}{l_{\gamma(r)}(\alpha)}\right) \\
&= \log\left(\frac{l_{\gamma(t)}(\alpha)}{l_{\gamma(s)}(\alpha)}\right) = d_R(\gamma(s),\gamma(t)).
\end{aligned}$$

Hence, γ is a geodesic by Lemma 1.18. □

In particular we have that geodesics are exactly the paths which preserve a witness and so by Lemma 1.22 all witnesses between its endpoints. As there always exists a candidate witness we can also say that geodesics from A to B are the paths which maximally stretch a candidate from A along its way.

Using the Lemmas 1.22 and 1.23 one can see that two geodesics concatenate to a geodesic if and only if they share a witness. By Lemma 1.18 this is exactly the case when their endpoints satisfy the triangle equality.

Corollary 1.24
Let $\gamma\colon [r,s] \to CV_n$ and $\sigma\colon [s,t] \to CV_n$ be two geodesics with $\gamma(s) = \sigma(s)$. Then the following are equivalent:

(i) Their concatenation $\gamma * \sigma \colon [r,t] \to CV_n$ is a geodesic.

(ii) $W_R(\gamma(r), \gamma(s)) \cap W_R(\sigma(s), \sigma(t)) \neq \emptyset$.

(iii) $d_R(\gamma(r), \sigma(t)) = d_R(\gamma(r), \gamma(s)) + d_R(\sigma(s), \sigma(t))$.

By [FM11, Theorem 5.5] we have that (CV_n, d_R) is a geodesic space. The more concrete statement together with a sketch of the proof can be found later as Theorem 1.35.

While CV_n with the asymmetric metric is a geodesic space, this does not hold for the symmetric metric. As an example we have that for each neighbourhood U of a figure of eight graph in CV_2 we can find two points $A, B \in U$ without a symmetric geodesic joining them as shown in the following lemma.

Lemma 1.25
Let $U \subset CV_2^{\text{red}}$ be an open set which intersects at least two simplices. Then U contains points $A, B \in U$ such that there exists no symmetric geodesic from A to B.

Proof. The idea is to choose A and B in two different adjacent simplices and consider the possible intersection points C and C' of geodesics from A to B respectively from B to A with the common face of the simplices containing A and B (see Figure 6). If there exists a symmetric geodesic from A to B, then there exists a common intersection point $C = C'$. We will show that choosing A and B appropriately such a common intersection point does not exist.

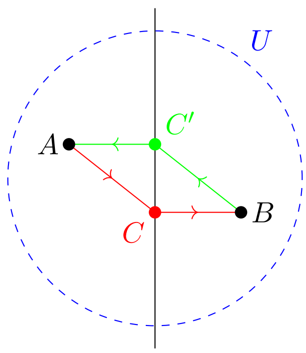

Figure 6: A, B and U in CV_2^{red} with asymmetric geodesics

As U intersects two simplices in CV_2 it contains a figure of eight graph $X \in U$. We denote by α and β its two loops and its edge lengths by a and $1-a$ as in Figure 7. For small enough $\delta, \varepsilon > 0$ and $a^{\pm} := a \pm \delta$ let A, B be as in Figure 7 the two differently marked theta-graphs near X with the edge lengths $a^{\pm}, \varepsilon$ and $1 - a^{\pm} - \varepsilon$ such that we have $A, B \in U$. That is A and B are attained from X by first slightly stretching and shrinking the edge a respectively and afterwards relaxing the 4-valent vertex to an edge of length ε such that they have different marking.

Observe that reduced Outer Space without the figure of eight-simplex containing X has two components, one containing A and the other containing B. Hence, each geodesic

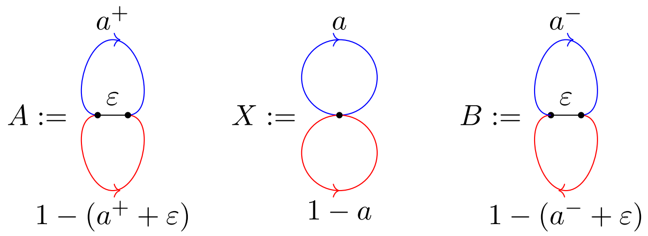

Figure 7: The points A, X and B in U

from A to B must pass through a point $C \in \Delta(X)$. Let C be such a graph in the same simplex as X with edge lengths b and $1-b$.

By Lemma 1.22 there exists a candidate of A which has to be maximally stretched from A to C and from C to B. Similarly there exists a candidate of B which is maximally shrunk along the geodesic. Hence, we consider the stretching factors of the candidates $\alpha, \beta, \alpha\beta$ and $\alpha\beta^-$. The length ratios of the candidates in A and B are as shown in Table 1.

	α	β	$\alpha\beta$	$\alpha\beta^{-1}$
$\frac{l_B(\cdot)}{l_A(\cdot)}$	$\frac{a^-+\varepsilon}{a^++\varepsilon}$	$\frac{1-a^-}{1-a^+}$	$\frac{1-\varepsilon}{1+\varepsilon}$	$\frac{1+\varepsilon}{1-\varepsilon}$
$\frac{l_C(\cdot)}{l_A(\cdot)}$	$\frac{b}{a^++\varepsilon}$	$\frac{1-b}{1-a^+}$	$\frac{1}{1+\varepsilon}$	$\frac{1}{1-\varepsilon}$
$\frac{l_B(\cdot)}{l_C(\cdot)}$	$\frac{a^-+\varepsilon}{b}$	$\frac{1-a^-}{1-b}$	$\frac{1-\varepsilon}{1}$	$\frac{1+\varepsilon}{1}$

Table 1: The stretching factors of candidates in Lemma 1.25

Observe that we have as stretching factors $\frac{l_B(\alpha)}{l_A(\alpha)} < 1 < \frac{l_B(\beta)}{l_A(\beta)}$. As A and B are normalised this means that α is not stretched from A to B and hence is not a witness from A to B. Similarly β is not a witness from B to A.

As $\alpha\beta$ is not a candidate in A we only need to compare the length ratios $\frac{l_B(\beta)}{l_A(\beta)}$ and $\frac{l_B(\alpha\beta^{-1})}{l_A(\alpha\beta^{-1})}$ to get the candidate witnesses from A to B and compare $\frac{l_B(\alpha)}{l_A(\alpha)}$ with $\frac{l_B(\alpha\beta)}{l_A(\alpha\beta)}$ to get the candidate witnesses from B to A.

- $\alpha\beta^{-1}$ is the candidate witness from A to B if $\frac{l_B(\alpha\beta^{-1})}{l_A(\alpha\beta^{-1})} \geq \frac{l_B(\beta)}{l_A(\beta)}$. We have:

$$\begin{aligned}\frac{l_B(\alpha\beta^{-1})}{l_A(\alpha\beta^{-1})} \geq \frac{l_B(\beta)}{l_A(\beta)} &\iff \frac{1+\varepsilon}{1-\varepsilon} \geq \frac{1-a^-}{1-a^+}\\ &\iff (1+\varepsilon)(1-a^+) \geq (1-\varepsilon)(1-a^-)\\ &\iff (1+\varepsilon)(1-a-\delta) \geq (1-\varepsilon)(1-a+\delta)\\ &\iff 2\delta - 2\varepsilon + 2\varepsilon a \leq 0\\ &\iff \delta \leq (1-a)\varepsilon\end{aligned}$$

  So we have that $\alpha\beta^{-1}$ is the only candidate witness from A to B if $\delta < (1-a)\varepsilon$.

- $\alpha\beta$ is the candidate witness from B to A if

$$\begin{aligned}\frac{l_A(\alpha\beta)}{l_B(\alpha\beta)} \geq \frac{l_A(\alpha)}{l_B(\alpha)} &\iff \frac{1+\varepsilon}{1-\varepsilon} \geq \frac{a^+ + \varepsilon}{a^- + \varepsilon}\\ &\iff (1+\varepsilon)(a^- + \varepsilon) \geq (1-\varepsilon)(a^+ + \varepsilon)\\ &\iff 2(-\delta + \varepsilon(a+\varepsilon)) \geq 0\\ &\iff \delta \leq \varepsilon a + \varepsilon^2.\end{aligned}$$

  Hence, $\alpha\beta$ is the only candidate witness from B to A if $\delta < a\varepsilon$ holds.

If we choose $\delta < \min\{(1-a)\varepsilon, a\varepsilon\}$ we have that $\alpha\beta$ and $\alpha\beta^{-1}$ are witnesses from B to A and A to B respectively. By Lemma 1.22 this means that for any intermediate point C on a symmetric geodesic from A to B they are also maximally stretched and shrunk towards A respectively. Hence, we have the following conditions on the lengths of the edges of C:

- $\alpha\beta$ is maximally stretched from C to A:

$$\begin{aligned}\frac{l_A(\alpha\beta)}{l_C(\alpha\beta)} \geq \frac{l_A(\alpha)}{l_C(\alpha)} &\iff 1+\varepsilon \geq \frac{a^+ + \varepsilon}{b}\\ &\iff b \geq \frac{a^+ + \varepsilon}{1+\varepsilon}\end{aligned}$$

- $\alpha\beta^{-1}$ is maximally stretched from C to B:

$$\begin{aligned}\frac{l_B(\alpha\beta^{-1})}{l_C(\alpha\beta^{-1})} \geq \frac{l_B(\beta)}{l_C(\beta)} &\iff 1+\varepsilon \geq \frac{1-a^-}{1-b}\\ &\iff 1-b \geq \frac{1-a^-}{1+\varepsilon}\\ &\iff b \leq \frac{1+\varepsilon-(1-a^-)}{1+\varepsilon} = \frac{a^- + \varepsilon}{1+\varepsilon}\end{aligned}$$

As we have $a^- < a^+$ we get the desired contradiction $b \leq \frac{a^-+\varepsilon}{1+\varepsilon} < \frac{a^++\varepsilon}{1+\varepsilon} \leq b$. $\square$

Thinking of a simplex $\Delta \subset CV_n$ with coordinates in $\mathbb{R}^k$, we can take straight lines in a simplex. Recall that for $A, B \in \Delta$ they have the same topological type $\Delta(A)$ and differ only in their edge lengths. Then the straight line between them is the path $\gamma\colon [0,1] \to \Delta$ such that $\Delta(\gamma(t)) = \Delta(A)$ and for an edge $e \in E(\Delta(A))$ the length $l_{\gamma(t)}(e) = tl_A(e) + (1-t)l_B(e)$ varies linearly along γ. We can check that $\gamma(t)$ satisfies Lemma 1.23 and hence is a geodesic from A to B. In particular we have:

Lemma 1.26 ([FM11, Proposition 5.9])
Let $\gamma\colon [0,1] \to \Delta \subset CV_n$ be a straight line in a simplex, then γ is a geodesic between its two endpoints and hence by Corollary 1.21 a symmetric geodesic.

Since geodesics are preserved under isometries we have by Lemma 1.26 and Lemma 1.25 that an isometry of CV_2^{red} preserves maximal simplices. This already implies the special case of Theorem 3.9 for the case $n = 2$.

Corollary 1.27
Isometries of CV_2^{red} are simplicial.

Remark 1.28
While Outer Space with the asymmetric metric is geodesic, these geodesics are typically far from being unique. That is for two points $A, B \in CV_n$ there are in general several geodesics from A to B. For example let $A, B, C \in CV_2$ be theta-graphs with the same marking α and β and the following edge lengths:

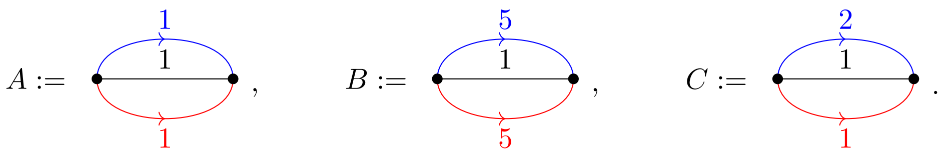

By Lemma 1.26 the straight line between A and B that is stretching the blue and red edges simultaneously, is a geodesic. On the other hand consider the concatenation of the two straight lines from A to C and from C to B. The candidates of A, B and C are the paths α, β and $\alpha\beta^{-1}$. Comparing their lengths we have:

$$\frac{l_C(\alpha\beta^{-1})}{l_A(\alpha\beta^{-1})} = \frac{3}{2}, \qquad \frac{l_C(\alpha)}{l_A(\alpha)} = \frac{3}{2}, \qquad \frac{l_C(\beta)}{l_A(\beta)} = \frac{2}{2} = 1$$
$$\frac{l_B(\alpha\beta^{-1})}{l_C(\alpha\beta^{-1})} = \frac{10}{3}, \qquad \frac{l_B(\alpha)}{l_C(\alpha)} = \frac{6}{3} = 2, \qquad \frac{l_B(\beta)}{l_C(\beta)} = \frac{6}{2} = 3.$$

By Proposition 1.12 one of these candidates is a maximally stretched loop and hence $\alpha\beta^{-1}$ is a (candidate) witness from A to C and from C to B. By Corollary 1.24 the concatenation of the straight lines from A to C and from C to B is also a geodesic.

On the other hand there exists (up to reparametrisation) only one geodesic from B to A: Similarly to before we have that α and β are both witnesses from B to A. Let now C' be a theta-graph with the same marking as A and B lying on a geodesic from B to

A. By Lemma 1.22 the elements α and β are also witnesses from B to C'. So we have $\frac{l_{C'}(\alpha)}{l_B(\alpha)} = \frac{l_{C'}(\beta)}{l_B(\beta)}$ and in particular the blue and red edge in C' have the same length. This implies that C' already has to lie on the straight line from B to A and hence there can not exist another geodesic from B to A.

Definition 1.29
Let $I \subset \mathbb{R}$ be an interval and $\gamma: I \to CV_n$ be a geodesic. We say γ is *rigid* if it is up to reparametrisation the unique geodesic from $\gamma(s)$ to $\gamma(t)$ for any $s \leq t \in I$.
If $I = [s,t]$ is a closed interval, then by Corollary 1.24 $\gamma: I \to CV_n$ is rigid if and only if it is the unique geodesic from $\gamma(s)$ to $\gamma(t)$.

Rigid geodesics are typically quite rare, e.g. will see in Section 3 or more exactly in Lemma 3.7 that given two arbitrary points they are typically not joined by a rigid geodesic. On the other hand by Proposition 2.13 there exists for a given point $A \in CV_n$ a finite number of rigid geodesics in $\Delta(A)$ starting at A. Moreover by Theorem 2.14 any two points in CV_n can be joined by a geodesic which is piecewise rigid. By Theorem 2.21 we have a full classification of rigid geodesic in terms of envelopes, which gives us a method to calculate all rigid geodesics starting or ending at a point.

1.5 Train tracks and folding paths

The important tools used in [FM11] to prove Proposition 1.12 and that (CV_n, d_R) is geodesic are so-called train tracks and folding paths. They were first introduced on graphs in the work of Bestvina and Handel [BH92] inspired by Thurston's train tracks on surfaces.

We will roughly sketch the work in [FM11] and use these tools to prove Lemma 1.36, which states that each geodesic between two points in CV_n can be extended to a geodesic ray.

Definition 1.30
Let Γ be a graph and $v \in V(\Gamma)$ a vertex. A *direction* at v is the germ of a geodesic ray starting at v, that is each time v is an endpoint of an edge e we consider a small part of the edge as a direction, in particular the valency of v is the cardinality of $T(v)$. We denote the set of directions at v by $T(v)$.

For each vertex v we want to endow $T(v)$ with an equivalence relation. For such equivalence relations we call the equivalence classes *gates*. If two edges are in the same gate at v one typically draws them tangentially close to v as in figure 8.

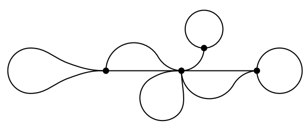

Figure 8: A visualisation of a train track structure with three gates at each vertex.

A *train track structure* on Γ is a collection of equivalence relations on $T(v)$, one for each vertex v of Γ. We say a train track structure *admissible* if it has at each vertex at least two different gates.

A *turn* is a pair of distinct directions at v. We call a turn *legal* if the two directions are in different gates, else we call it *illegal.* Likewise, we say a path α in a graph with a train track structure is called *legal* if at each vertex the entering and leaving directions of α are in different gates, else we say it is *illegal.* Observe that for an admissible train track on a finite graph, there always exist legal cycles, as we can extend a path with legal turns until we found a cycle.

Train tracks naturally arise from maps between graphs in the following way. Let Γ and Γ' be two graphs and $h\colon \Gamma \to \Gamma'$ be a continuous map, which is locally injective on the interior of each edge of Γ. For a vertex $v \in V(\Gamma)$ a gate consists then of the preimage of a direction of $h(v)$ (we might introduce here $h(v)$ as an artificial vertex of Γ'). As h is locally injective on each edge, this is well-defined. We call this the *train track structure induced by h.*

Before we apply induced train tracks to CV_n we need to introduce the notion of tension graph and optimal change of marking map.

Definition 1.31 (see [FM11, Definition 3.8])
Let $A, B \in CV_n$ and let S be the set of change of marking maps as in Definition 1.10.

(i) Let $h \in S$ an edge-wise linear change of marking map. That means along each edge $e \in E(A)$ the map $h|_e$ has a constant stretching factor $L(h|_e)$. We call the subgraph $\Delta_{\max}(h) := \{e \in E(A) \mid L(h|_e) = L(h)\}$ consisting of the edges of A which are maximally stretched under h the *tension graph* of h.

(ii) We call a change of marking map $h \in S$ *optimal*, if it satisfies the following conditions:
   - h has minimal Lipschitz constant $L(h) = \Lambda_R(A, B)$.
   - h is edge-wise linear.
   - The train track structure on $\Delta_{\max}(h)$ induced by h is admissible.

As can be seen in [FM11], an optimal change of marking map always exists: Clearly "tightening" along edges to make a map edge-wise linear at most decreases the Lipschitz constant of a map. If $\Delta_{\max}(h)$ has only one gate at a vertex v, we can slightly relax h at v to a new change of marking map h' with $L(h') \leq L(h)$ and $v \notin \Delta_{\max}(h') \subset \Delta_{\max}(h)$.

We can use the induced train track structure from an optimal change of marking map to distinguish witnesses:

Lemma 1.32
Let $A, B \in CV_n$ and $h\colon A \to B$ an optimal change of marking map and $\alpha \in F_n \setminus \{\mathrm{id}\}$. Then α is a witness from A to B if and only if α is a legal, closed path in $\Delta_{\max}(h)$ with respect to the induced train track by h.

Proof. Let α be a closed, legal path in the train track induced by h. Asα is legal $h|_\alpha$ immerses α at each vertex. Since h is also already locally injective on each edge, we have that $h(\alpha)$ is cyclically reduced. As $\alpha \subset \Delta_{\max}(h)$ lies in the tension graph we have then $l_B(\alpha) = l_B(h(\alpha)) = L(h) \cdot l_A(\alpha)$ which concludes that α is a witness.

Let now α be a witness from A to B, then we have that α is maximally stretched that means $l_B(\alpha) = \Lambda_R(A, B) l_A(\alpha) = L(h) l_A(\alpha)$. This can only happen if α already lies in the tension graph and there is no cancellation happening at each vertex, i.e. α makes no illegal turns. □

As the induced train track of an optimal change of marking map is admissible, we now have that there always exists a witness. For example we can construct an infinite long, legal path ρ in $\Delta_{\max}(h)$ and take as loop any finite subpath which starts with the same gate as it would end in ρ.

Similarly to Stallings' folding in [Sta83] we can continuously identify or "fold" edges close to a vertex, namely some part of edges close to a vertex v are identified if they are in the same gate. Gradually folding then yields a continuous path in CV_n (see Figure 9).

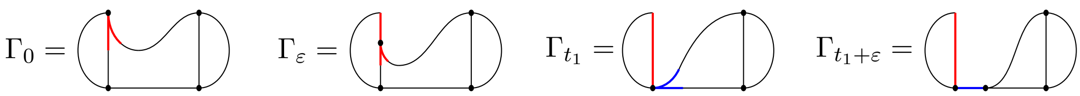

Figure 9: Folding first the red and then the blue edges of a graph

Definition 1.33
Let Γ be a metric graph with an admissible train track structure. For every vertex $v \in V(\Gamma)$ and small enough $t > 0$ we define on points $p, q \in \Gamma$ the equivalence relation $p \sim_{v,t} q$ if

- $d(v, p) = d(v, q) \leq t$
- and the geodesics from v to p and q respectively are in the same gate.

Increasing t means we slowly fold all edges at v which are in the same gate.

Let $\sim_t$ be the union of these equivalence relations over $v \in V(\Gamma)$ and set $\Gamma_t := \Gamma / \sim_t$. Keep in mind that (after possibly introducing some new vertices) Γ_t is itself again a metric graph with smaller volume than Γ.

Let t_1 be the first time when an edge is completely folded and $s \leq t_1$. We call a path of the form $\Gamma_* : [0, s] \to \{\text{metric graphs}\}, t \mapsto \Gamma_t$ a *simple folding path* (of metric graphs). A *folding path* is then a path $\Gamma_* : I \subseteq \mathbb{R}_{\geq 0} \to \{\text{metric graphs}\}$ such that on each closed interval it is the concatenation of finitely many simple folding paths.

Observe that folding gives a quotient map $\sigma_t \colon \Gamma \to \Gamma_t$. As we only identify edges close to a common vertex, we can lift any closed loop and hence its induced map $\sigma_{t,*} \colon \pi_1(\Gamma) \to \pi_1(\Gamma_t)$ is surjective. As long as we do not completely fold as in Figure 10 then $\sigma_{t,*}$ is also injective and by the Whitehead theorem we have that σ_t is a homotopy equivalence. Hence, any marking on Γ carries over to a marking on Γ_t via the projection.

Moreover, as we continuously change the lengths, folding a point $A = (\Gamma, l, m) \in CV_n$ along an admissible train track without completely folding loops will yield a continuous path $A_t := (\Gamma_t, l_t, m_t)$ in CV_n with the marking $m_t := \sigma_t \circ m$. Here admissible ensures that the resulting graph Γ_t does not have leaves.

Figure 10: Foldings which are not homotopy equivalences

As in Lemma 1.32 we can now determine when a folding path is a geodesic in terms of legal paths.

Lemma 1.34 (i) Let $A_0 = (\Gamma, l, m) \in CV_n$ be a point with a train track structure on Γ. Let $\gamma\colon [0, t_1] \to CV_n, t \mapsto A_t$ be a simple fold and $\alpha \in F_n \setminus \{\mathrm{id}\}$. Then α is a legal path in the train track if and only if α is a witness from A to A_{t_1}.

(ii) Let $\gamma\colon I \subseteq \mathbb{R} \to CV_n, t \mapsto A_t$ be a folding path in CV_n. Then γ is a geodesic if and only if there exists an $\alpha \in F_n \setminus \{\mathrm{id}\}$ which is a legal path in each train track of the simple folds.

Proof. (i) Let α be a cyclically reduced path in A_0. As we only identify points close to a vertex, the length of α will at most decrease at these foldings. Such a decrease only happens at the vertices where α has an illegal turn since we have then some backtracking. That means all legal cycles have the same length in A_0 and in A_t while all illegal cycles have smaller length in A_t than in A_0. In particular the legal cycles are exactly the witnesses from A_0 to A_t as the maximal stretching is 1. The Lipschitz distance from A_0 to A_t comes here solely from the decrease of the volume.

(ii) Follows directly from (i) and Lemma 1.22 and Corollary 1.24 respectively.

□

We can use these folding paths to construct geodesics in Outer Space. As we want to extend such a geodesic in Lemma 1.36 in a similar manner we also roughly sketch the proof the following theorem.

Theorem 1.35 (Francaviglia–Martino [FM11, Theorem 5.5])
Let $A, B \in CV_n$. Then there exists a geodesic from A to B which is a concatenation of a path which linearly shrinks edges in A and a folding path.

Proof. Let $h\colon A \to B$ be an optimal change of marking map. After shrinking all the edge lengths of A which do not lie in $\Delta_{\max}(h)$ and eventually relaxing h at some vertices we get an optimal change of marking map $h_0 : A_0 \to B$ with $\Delta_{\max}(h_0) = A_0$ where A_0 is obtained from A by shrinking the edge lengths.

We now fold A_0 along the train track induced by h_0, that means we identify $p \sim_{v,t} q$ if $h_0(p) = h_0(q)$ and p and q are less than t away from a vertex v. Let t_1 be the first time we have completely folded an edge. As we only have identified points with the same image, this implies that h_0 factors through the folded graph $A_t := A_0/\sim_t$ for any $0 \leq t \leq t_1$. In particular we get a change of marking map $h_t : A_t \to B, [p]_t \mapsto h_0(p)$. Since we have folded along an admissible train track and h_t is a homotopy equivalence we get that A_t is an element in CV_n. Observe that h_t has on each edge the same Lipschitz constant as h_0. Furthermore, its induced train track has at each vertex at least two gates and thus is admissible. So we have that h_t is again an optimal change of marking map.

We can then either continue folding according to h_t or we have already that h_t is an immersion. If h_t is an immersion, then it stretches each closed path by $L(h_t)$, so we have already that h_t is a homothety and we have $A_t = B$ in CV_n. If we can continue to fold, observe that by each fold from A_s to A_{s+t} we decrease the volume by at least t. Given that h_0 and h_t have the same Lipschitz constant, we have $\operatorname{vol}(B) = \operatorname{vol}(h_t(A_t)) \leq L(h_t)\operatorname{vol}(A_t) \leq L(h_0)(\operatorname{vol}(A_0) - t)$ and thus the folding has to stop at some point.

Now let $\alpha \in F_n$ be a legal path in A with respect to the train track structure induced by h, then it is clearly a witness from A to A_0 as we only shrink edges outside of $\Delta_{\max}(h)$. As it is a legal path in A, it is still a legal path in A_0 and by Lemma 1.34 a witness from A_0 to A_t for some small enough $t > 0$.

Being a legal path means that the restriction of the optimal change of marking map to α is an immersion. As h_0 factors through any h_t, we also have that the restriction of h_t to α is still an immersion. So α is still a legal path in each A_t and again a witness for any A_t to A_{s+t}. By Corollary 1.24 we have that the constructed path in CV_n is a geodesic. □

We will use this to extend a given geodesic to a geodesic ray of arbitrarily long diameter, that is it will either go to the thin part or has infinite length.

Lemma 1.36
Let $A, B \in CV_n$ and $\gamma\colon [0,1] \to CV_n$ be a geodesic from A to B. Then there exists a geodesic $\sigma\colon \mathbb{R}_{\geq 0} \to CV_n$ such that $\sigma|_{[0,1]} = \gamma$ and for any $L \geq 0$ there exists a $t > 0$ with $d(\sigma(0), \sigma(t)) = d_R(\sigma(0), \sigma(t)) + d_R(\sigma(t), \sigma(0)) > L$.

Furthermore, for every $\alpha \in W_R(\gamma(0), \gamma(1))$ and $t > 0$ we also have $\alpha \in W_R(\sigma(0), \sigma(t))$.

Proof. By Corollary 1.24 and Theorem 1.35 we can assume that γ is either a folding path from A to B or just shrinking edges outside of the tension graph.

If γ only shrinks edges, let $E \subset E(A)$ be the edges shrunk by γ. Thus, an element $\alpha \in F_n$ is a witness from A to B if and only if it is disjoint from E. Hence, we can continue shrinking the edges in E without changing the witnesses to get a geodesic σ.

If E contains a loop, then we can continue shrinking the edges until the loop is arbitrarily short, that is σ goes into the thin part. Hence, σ already has infinite diameter, i.e. for any L we have $d(\sigma(0), \sigma(t)) > d_R(\sigma(t), \sigma(0)) > L$ for large enough t as we have seen in Example 1.13 (ii).

Else E is a forest, hence we can continue to shrink E until we collapse all edges in E at $t = 2$ to attain $B' := \sigma(2)$. Afterwards we extend σ by folding two arbitrary germs of edges in B' which were separated by E as depicted in Figure 11:

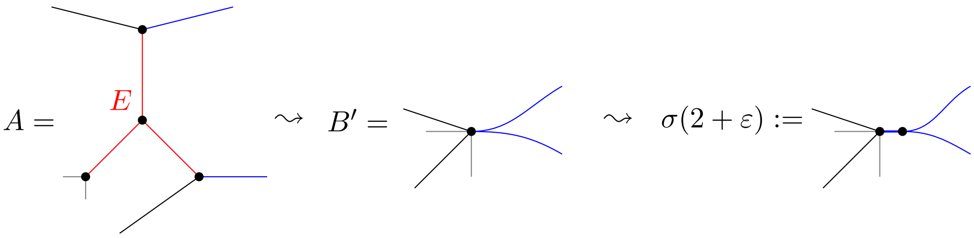

Figure 11: Folding after collapsing a forest

Such two germs always exist, as E is a finite forest and A and B have no leaves. Clearly, any witness from A to B' has to be a legal path along this folding, otherwise it would have crossed the edges in E connecting the two germs. Hence, we can assume that γ is a folding path by Corollary 1.24.

By Lemma 1.34 $\alpha \in F_n \setminus \{\mathrm{id}\}$ is a witness from A to B if and only if it is a legal path in the train track for each fold. We will show that we can continue a folding path in such a manner that each legal path in the previous folds is also a legal path in the next fold. In particular, extending γ to σ any witness $\alpha \in W_R(A, B)$ will then also be a witness from A to $\sigma(t)$ for $t > 0$ and by Lemma 1.34 σ will be a geodesic.

Let $\sigma\colon [0, t] \to CV_n$ be a folding path. As the order of folds does not matter we will assume that σ only folds at one vertex at a time. We will parametrise the folding path by the folded length of the edges, that is we consider $\sigma(s)$ as a graph of volume e^{-s}. Observe that as any legal path has the same length in $\sigma(s)$ for all s we have parametrised σ by its asymmetric length. Let $t_k < t$ be the last time when we started folding at a vertex v, namely we obtain $\sigma(t)$ by folding the same gate in $\sigma(t_k)$ by some amount ε:

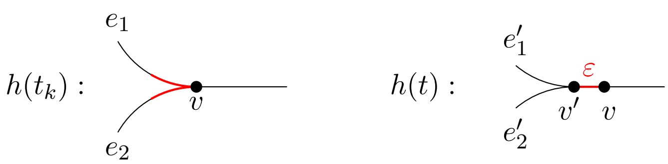

If here or at some point of the continuation of σ we might fold a loop as in Figure 10, we continue folding the loop until it becomes arbitrarily small and as before end up in the thin part and hence are done.

Otherwise we will continue folding until one edge e_1 is completely folded at some $t_{k+1} \geq t$. At this point we can continue folding any of the adjacent edges l_1 of e_1 with e_2 (or to an adjacent edge of e_2 if e_1 and e_2 had the same length), that is l_1 plays now the role of e_1 as depicted in Figure 12.

An important fact here is that any legal path $\alpha \in F_n \setminus \{\mathrm{id}\}$ from $\sigma(t_k)$ to $\sigma(t)$ is clearly legal from $\sigma(t)$ to $\sigma(t_{k+1})$ and still legal from $\sigma(t_{k+1})$ to $\sigma(t_{k+1} + \varepsilon)$. Namely, if α took in $\sigma(t_{k+1})$ the only illegal turn from l_1 to e_2', then it must have already taken the edge-sequence l_1, e_1, e_2 in $\sigma(t)$, which is also illegal. In particular, by Lemma 1.34 if $\sigma|_{[0,t]}$ is a geodesic, then this continuation is still a geodesic.

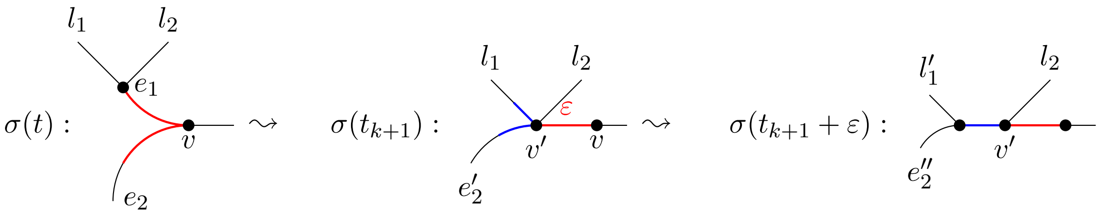

Figure 12: Extension of a folding path

The described folding decreases the length of the longer edge by the length of the smaller one and leaves the rest of the edge lengths the same. So after finitely many steps this longer edge is also completely folded. This implies that unless we run into a fold as in Figure 10 we can, by choosing the folded adjacent edges accordingly, fold until each edge is arbitrarily small. That is we can continue σ to the half-open interval $[0, \infty)$. Two such choices of folded edges would be for example to always choose the longest adjacent edge or to choose the edges which give the shortest path to the longest edge in the current graph $\sigma(t_i)$.

We now either end up in the thin part of Outer Space or we get for any $0 < t < \infty$ that $\text{vol}(\sigma(t)) = e^{-t}$. Moreover any original witness α is a legal path along all folds and hence the length of α in $\sigma(t)$ for $t \in [0, \infty)$ is constant. Hence, we have $d_R(\sigma(0), \sigma(t)) \geq \log(\frac{l_{\sigma(t)}(\alpha)}{l_{\sigma(0)}(\alpha)} \cdot \frac{\text{vol}(\sigma(0))}{\text{vol}(\sigma(t))}) = \log(1 \cdot \frac{1}{e^{-t}}) = t \to \infty$ as t goes to ∞. □

1.6 Algorithm to determine the distance of two points

Given two points $A, B \in CV_n$, we will give a description of an algorithm to determine their Lipschitz distance $d_R(A, B)$ based on Proposition 1.12. The algorithm takes as input two marked, metric graphs given as in Figure 2 and returns the candidate witnesses from A to B, their stretching and the Lipschitz distance from A to B. An implementation of this algorithm in Sage [Sag] can be found in [Ste18] and is added in the Appendix. We will split the description of the algorithm in the following steps:

A. Describe elements of CV_n for a computer.

B. Find the candidates in a marked graph.

C. Translate words of F_n into immersed loops in a graph.

D. Compare lengths of paths in graphs.

A. An element of CV_n

During this subsection a marked, metric graph $(\Gamma, l, m) \in CV_n$ consists of the following data:

- A set of vertices $V(\Gamma)$, which we typically enumerate from 1 up to $2n-2$.
- A list of (oriented) edges $E(\Gamma)$, where each edge $e \in E(\Gamma)$ is a tuple (v_1, v_2, α) of two vertices $v_1, v_2 \in V(\Gamma)$ and a label $\alpha \in F_n$.
- A length function $l : E(\Gamma) \to \mathbb{R}_{\geq 0}$ which assigns to each edge a length.

We require additionally that the edges $(v_1, v_2, \mathrm{id}) \in E(\Gamma)$ with the neutral element $\mathrm{id} \in F_n$ as label form a spanning tree T of Γ and the labels of the other edges form a basis of F_n.

For an edge (v_1, v_2, α) we will write $\overline{(v_1, v_2, \alpha)} = (v_2, v_1, \overline{\alpha}) \in E(\Gamma)$ for the (same) edge with flipped orientation. Here $\overline{\alpha} := \alpha^{-1} \in F_n$ denotes the inverse of α in F_n.

B. Computing the candidates of a marked graph

The easiest way to compute all candidates is to first compute all simple loops, i.e. loops which cross each vertex and edge at most once. Afterwards we can use these simple loops to compute all barbells and figures of eight. As there is already a rich literature about effectively calculating all simple loops and paths between two points in a graph, we will omit this part here. Instead we use that we already have a function or list *simpleLoops* of all simple loops in a given graph. Also we assume that we have a function *allPaths* which yields all simple edge-paths between two given points, these are paths which do not cross the same vertex twice. For example Sage [Sag] already implements these functions.

Notation 1.37
For a simple loop α and a vertex $v \in \alpha$ we will denote with α_v the cyclic permutation of α starting at v and with $\overline{\alpha}$ the loop with inverse orientation. For example for the loop $\alpha = ((v_1, v_2, \alpha_{1,2}), (v_2, v_3, \alpha_{2,3}), \ldots, (v_k, v_1, \alpha_{k,1}))$ we have

$$\begin{aligned} \alpha_{v_i} &= ((v_i, v_{i+1}, \alpha_{i,i+1}), (v_{i+1}, v_{i+2}, \alpha_{i+1,i+2}), \ldots, (v_k, v_1, \alpha_{k,1}), \ldots, (v_{i-1}, v_i, \alpha_{i-1,i})) \\ \text{and } \overline{\alpha}_{v_i} &= ((v_i, v_{i-1}, \overline{\alpha_{i-1,i}}), (v_{i-1}, v_{i-2}, \overline{\alpha_{i-2,i-1}}), \ldots, (v_1, v_k, \overline{\alpha_{k,1}}), \ldots, (v_{i+1}, v_i, \overline{\alpha_{i+1,i}})). \end{aligned}$$

For two edge-paths $\alpha = (e_1, \ldots, e_k)$ and $\beta = (f_1, \ldots, f_m)$ we will denote with $\alpha \star \beta = (e_1, \ldots, e_k, f_1, \ldots, f_m)$ their concatenation.

To compute all figures of eight and barbells it is enough to go through all pairs α, β of simple loops and check if they intersect in a vertex. Depending on their intersection we get the following candidates:

- If they intersect in exactly one vertex $\{v\} = \alpha \cap \beta$, then we get two figures of eight, namely $\alpha_v + \beta_v$ and $\alpha_v + \overline{\beta}_v$.
- If they do not intersect we get the barbells where α and β are the simple loops. That is for each simple edge-path $\chi_{v,w}$ between a vertex $v \in \alpha$ and a vertex $w \in \beta$ we get as barbells $\alpha_v + \chi_{v,w} + \beta_w + \overline{\chi_{v,w}}$ and $\alpha_v + \chi_{v,w} + \overline{\beta}_w + \overline{\chi_{v,w}}$
- If they intersect in more than one vertex we do not get a candidate.

This yields Algorithm 1 to compute all candidates.

Algorithm 1 compute figures of eight and barbells as list of edges

Input: A graph Γ and a list *simpleLoops* of all simple loops in Γ
Output: A list of all candidates in Γ

Let *eights* and *barbells* be two empty lists.
for $0 \leq i < length(simpleLoops)$ **do**
  for $i < j < length(simpleLoops)$ **do**
    Set $\alpha := simpleLoops[i]$ and $\beta := simpleLoops[j]$.
    Set $common := vertexIntersection(\alpha, \beta)$.
    if $length(common) \geq 2$ **then**
      Skip
    else if $length(common) == 1$ **then**
      Set $v := common[0]$
      Add $\alpha_v + \beta_v$ to *eights*
      Add $\alpha_v + \overline{\beta}_v$ to *eights*
    else
      for $v \in \alpha, w \in \beta$ **do**
        Set Γ' to be a copy of Γ without the vertices of $\alpha \setminus \{v\}$ and $\beta \setminus \{w\}$.
        for $\chi \in allPaths$ from v to w in Γ' **do**
          Add $\alpha_v + \chi \setminus \{w\} + \beta_w + \overline{\chi} \setminus \{v\}$ to *barbells*.
          Add $\alpha_v + \chi \setminus \{w\} + \overline{\beta}_w + \overline{\chi} \setminus \{v\}$ to *barbells*.
        end for
      end for
    end if
  end for
end for
return *simpleLoops*, *eights* and *barbells*

Remark 1.38
It should be noted here that there are more efficient algorithms to determine candidates in a graph. For example we can use that we already have a spanning tree T. Let $\alpha_1, \ldots, \alpha_n \in F_n$ be the labels of the edges outside of the spanning tree. Then each simple loop corresponds to a word $\alpha = \alpha_{i_1}^{\varepsilon_1} \ldots \alpha_{i_k}^{\varepsilon_k}$ with pairwise different α_i and $\varepsilon \in \{\pm 1\}$ such that each vertex v occurs at most once in all paths in T joining the terminal vertex of $\alpha_i^{\varepsilon_i}$ with the initial vertex $\alpha_{i+1}^{\varepsilon_{i+1}}$ for $1 \leq i < k$ and joining the terminal vertex of $\alpha_k^{\varepsilon_k}$ with the initial vertex of $\alpha_1^{\varepsilon_1}$. Similarly, we get that figures of eight are words α as above with exactly one vertex occurring twice.

For barbells we consider all words of the form $\alpha\chi\beta\overline{\chi}$ where α, χ and β are disjoint words as above such that each vertex occurs at most once in one of the following five sets:

- The three (non-closed) paths of each word α, χ and β.
- The union of the paths from the start and ending vertices of α to the first vertex of χ (this may look like a tripod).

- The union of the paths from the start and ending vertices of β to the last vertex of χ.

To convert a closed edge-path in our graph Γ to an element in F_n we multiply all the edge labels of the edges along the path. Hence, we have now an algorithm to determine all candidates of a marked graph as elements of F_n.

C. Translating words of F_n into immersed loops in a graph

To compare the length of a candidate in two different graphs we now need a way to write a given word as an edge-path in a marked graph. As the labels of the edges outside of our spanning tree form a basis of F_n this practically boils down to rewriting a word in a different basis and afterwards connecting the letters via the spanning tree. Again there already exist efficient algorithms to rewrite a given word as reduced word in another basis for example in GAP [GAP21]. We will assume that we have such an algorithm at our disposal and just cyclically reduce the resulting word.

To translate this cyclically reduced word into a loop, we connect the corresponding sequence of labelled edges, which corresponds to the letters of the word, along our spanning tree T. Here we use that any two vertices u, v are joined by a unique path in a our spanning tree T. We will denote this path with $\chi_{u,v}$. This gives the following algorithm:

Algorithm 2 Immerse a word as an edge-path

Input: A marked graph Γ and a word $\alpha \in F_n$.
Output: An immersed loop representing the conjugacy class of α as list of edges.

```
Let (α_1, ..., α_n) be the non-trivial labels of edges of Γ.
compute α as reduced word ω = (α_{j_1}^±, ..., α_{j_k}^±) of α_1, ..., α_n.
while \overline{ω_1}=ω_{length(ω)} do
    Remove ω_1 and ω_{length(ω)} from ω.
end while
Set α to be an empty list of (oriented) edges.
Add ω_1 at the end of ω.
for 1 ≤ i < length(ω): do
    Add the edge corresponding to ω_i to α.
    Let u be the terminal vertex of the edge corresponding to ω_i
    Let v be the initial vertex of the edge corresponding to ω_{i+1}
    Add the edge-path χ_{u,v} to α
end for
return α
```

D. Calculating the distance

To calculate the distance between two points in Outer Space we only need to put the previous algorithms together. We use as an abbreviation $l_A(\alpha) = \sum_{e \in \alpha} l_A(e)$ for the length of some edge-path α.

Algorithm 3 Calculate the distance of two points in Outer Space

Input: Two points $A, B \in CV_n$

Output: Their distance $d_R(A, B)$.

Use Algorithm 1 to get a list *cand* of candidates of A.
Set $fractions$ to be an empty list.
for $\alpha \in cand$: **do**
    Use Algorithm 2 to write α as an edge-path α' in B.
    Add $\frac{l_B(\alpha')}{l_A(\alpha)}$ to $fractions$.
end for
Set $\Lambda_R := \max(fractions)$.
Set $vol_A := \sum_{e \in E(A)} l_A(e)$ and $vol_B := \sum_{e \in E(B)} l_B(e)$.
Set $d_R := \log(\Lambda_R \cdot \frac{vol_A}{vol_B})$
return d_R

2 Envelopes in Outer Space

As we have seen in Remark 1.28 geodesics in Outer Space are not necessarily unique. To give a measure of how much two points fail to have a unique geodesic, it seems reasonable to look at the all geodesics at the same time. We borrow the notion of an *envelope* from [DLRT20].

Definition 2.1
Let (X, d) be a metric space and $A, B \in X$. Then we define the *envelope from A to B* as the set

$$\operatorname{Env}_d(A, B) := \{C \in X \mid C \text{ lies on a geodesic from } A \text{ to } B\}.$$

To distinguish the envelopes in CV_n coming from the two different metrics, we will write $\operatorname{Env}_R := \operatorname{Env}_{d_R}$ for envelopes with respect to the asymmetric metric and $\operatorname{Env} := \operatorname{Env}_d$ for the symmetric metric.

Envelopes in CV_n have the following important properties which we will use later.

Remark 2.2 (i) It is clear that isometries preserve envelopes as they send geodesics to geodesics.

(ii) The diameter of an envelope is bounded. More explicitly:
- If the metric is symmetric, the diameter of an envelope is the distance between the two endpoints by the triangle inequality.
- If d is an asymmetric metric we get for all $C, C' \in \operatorname{Env}_d(A, B)$

$$d(C, C') \leq d(C, B) + d(B, A) + d(A, C') \leq 2d(A, B) + d(B, A).$$

(iii) By Lemma 1.18 we have that intermediate envelopes are subsets, namely for $A, B \in X$ and $C \in \operatorname{Env}(A, B)$ we have

$$\operatorname{Env}(A, C) \cup \operatorname{Env}(C, B) \subseteq \operatorname{Env}(A, B).$$

(iv) Because (CV_n, d_R) is a geodesic space and by the equivalence of (i) and (iii) in Corollary 1.24, we can write envelopes in CV_n as

$$\operatorname{Env}_R(A, B) = \{C \in CV_n \mid d_R(A, B) = d_R(A, C) + d_R(C, B)\}.$$

The following example shows that the upper bound in Remark 2.2 (ii) is the best we can obtain for CV_n.

Example 2.3
Let $0 < \varepsilon < 1$ be a small number and $A \in CV_3$ be the rose R_3 with loop lengths 1, $\frac{1}{\varepsilon}$ and ε. Similarly let $B \in CV_3$ be the rose obtained from A by shrinking the second and third petal by the factor ε, that is we have edge lengths $(1, 1, \varepsilon^2)$ in B. Let furthermore

C and C' be the two rose graphs obtained from A by shrinking only one of the petals as in B, that is we have the edge lengths $(1, \frac{1}{\varepsilon}, \varepsilon^2)$ and $(1, 1, \varepsilon)$, respectively. Clearly C and C' lie in the envelope $\mathrm{Env}_R(A, B)$ as the first petal is maximally stretched from A to C and C' respectively and from C and C' to B:

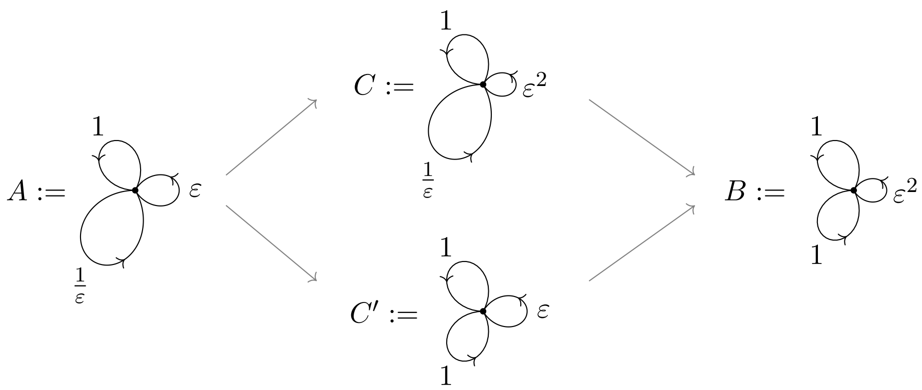

That means while the distance from A to B comes purely from the fraction of their volumes, the distance from C to C' includes both, the stretching from B to A as well as the fraction of the volumes. More exactly we have:

$$d_R(A, B) = \log\left(1 \cdot \frac{\mathrm{vol}(A)}{\mathrm{vol}(B)}\right) = \log\left(\frac{1 + \frac{1}{\varepsilon} + \varepsilon}{2 + \varepsilon^2}\right)$$

$$d_R(B, A) = \log\left(\frac{\varepsilon}{\varepsilon^2} \cdot \frac{\mathrm{vol}(B)}{\mathrm{vol}(A)}\right) = \log\left(\frac{1}{\varepsilon}\right) - \log\left(\frac{1 + \frac{1}{\varepsilon} + \varepsilon}{2 + \varepsilon^2}\right)$$

$$\begin{aligned} d_R(C, C') &= \log\left(\frac{\varepsilon}{\varepsilon^2} \cdot \frac{\mathrm{vol}(C)}{\mathrm{vol}(C')}\right) \\ &= \log\left(\frac{1}{\varepsilon}\right) + \log\left(\frac{1 + \frac{1}{\varepsilon} + \varepsilon^2}{2 + \varepsilon}\right) \end{aligned}$$

Letting ε go to zero we have that $\log\left(\frac{1+\frac{1}{\varepsilon}+\varepsilon^2}{2+\varepsilon}\right)$ approaches $\log\left(\frac{1+\frac{1}{\varepsilon}+\varepsilon}{2+\varepsilon^2}\right)$, hence have that the diameter of $\mathrm{Env}_R(A, B)$ may approach $2d_R(A, B) + d_R(B, A)$ arbitrarily close.

Before we look into envelopes in CV_n we will introduce two notations:

Notation 2.4
Let $A \in CV_n$, $e \in E(A)$ be an edge in the underlying graph of A and $\alpha \in F_n$. Then we denote by $\#(e, \alpha)$ the number of times the cyclically reduced path corresponding to α passes through e (without considering the orientation). This number depends only on the topological type $\Delta(A)$ of A and not its lengths.

Definition 2.5
Let $A, B \in CV_n$ be two points. We define the *supporting simplices* of the envelope $\mathrm{Env}_R(A, B)$ as the set of all simplices $\Delta \subset CV_n$ which intersect the envelope:

$$\mathrm{supp}(\mathrm{Env}_R(A, B)) := \{\Delta \subset CV_n \mid \Delta \text{ is a simplex and } \Delta \cap \mathrm{Env}_R(A, B) \neq \emptyset\}.$$

As the diameter of an envelope is bounded, we have by the following Lemma that the support is always finite.

Lemma 2.6
For any $B \in CV_n$ and $r > 0$ the *ingoing ball* $B_r^{\mathrm{in}}(B) := \{A \in CV_n \mid d_R(A, B) < r\}$ intersects only finitely many simplices of CV_n.

Proof. Without loss of generality we can assume that B is the standard marked rose R_n with edge lengths all $1/n$ since we have $B_r^{\mathrm{in}}(B) \subseteq B_{r+d_R(B,R_n)}^{\mathrm{in}}(R_n)$. We will show that for a fixed graph Γ there are only finitely many markings $m\colon R_n \to \Gamma$ such that there exists a length function l the point (Γ, l, m) lies in $B_r^{\mathrm{in}}(R_n)$. Since the unmarked graphs $CV_n/\operatorname{Out}(F_n)$ are finite, the claim then follows.

Let Γ be a fixed graph and T a spanning tree of Γ. Recall that a marking corresponds to labelling the edges outside of T with a basis $\omega_1, \ldots, \omega_n \in F_n$. As conjugating every word ω_i corresponds to a free homotopy we can always choose $\omega_1, \ldots, \omega_n$ in such a way that conjugating does not decrease the sum of the word-lengths.

Let $A := (\Gamma, l, m) \in CV_n$ be a normalised representant. Since each ω_i corresponds to a simple loop it has length at most 1 in A. Moreover, for each word $\omega \in F_n$ its length in B is $\frac{1}{n}$ of its cyclically reduced word-length. For $A \in B_r^{\mathrm{in}}(B)$ we have $\log(\frac{l_B(\omega_i)}{l_A(\omega_i)}) \leq r$, hence cyclically reduced edge labels of A have at most word-length ne^r.

On the other hand assume ω_i is not cyclically reduced that is we can write $\omega_i = \alpha\tilde{\omega_i}\alpha^{-1}$ for some letter α and a shorter word $\tilde{\omega_i}$. Then there exists an $\omega_j \neq \omega_i$ such that the word concatenation $\omega_j\omega_i$ is a cyclically reduced word: Assume $\omega_j\omega_i$ is not a cyclically reduced word for all ω_j, then we either have $\omega_j = \tilde{\omega_j}\alpha^{-1}$ or $\omega_j = \alpha\tilde{\omega_j}$. But that means we can simultaneously conjugate all labels of A by α to decrease the word length of ω_i without increasing the word length of any other ω_j which contradicts our choice of $\omega_1, \ldots, \omega_n$.

This means up to free homotopy we have that each word ω_i is either cyclically reduced or there exists an $\omega_j \neq \omega_i$ such that $\omega_j\omega_i$ is cyclically reduced. Again in the latter case the length of $\omega_i\omega_j$ in A is at most 2, thus ω_i has word-length of at most $2ne^r$. As there are only finitely many words of a given word-length we have only finitely many possible markings for Γ such that $A := (\Gamma, l, m) \in B_r^{\mathrm{in}}(B)$. $\square$

Remark 2.7
The statement of Lemma 2.6 does not hold for *outgoing balls*, i.e. sets of the form $B_r^{\mathrm{out}}(A) := \{B \in CV_n \mid d_R(A, B) < r\}$. For example let A be the rose R_2 with edge labels α and β and edge lengths $\frac{1}{2}$. For $k \in \mathbb{N}_{\geq 2}$ let B_k be the rose with edge labels $\alpha\beta^k$ and β and with edge lengths $1 - \frac{1}{k}$ and $\frac{1}{k}$. Then a short calculation shows

$$d_R(A, B_k) = \log\left(\frac{l_{B_k}(\alpha)}{l_A(\alpha)}\right) = \log\left(\frac{2 - \frac{1}{k}}{\frac{1}{2}}\right) < \log(4)$$

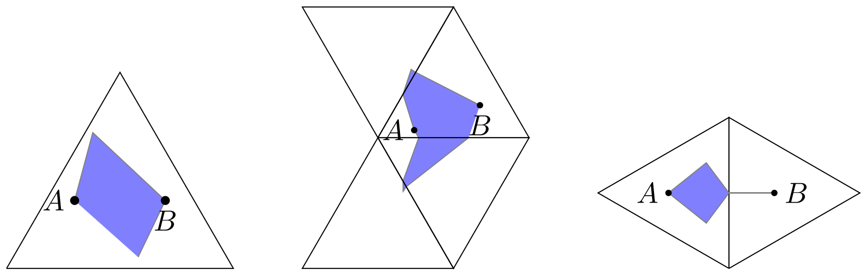

Figure 13: Some envelopes $\mathrm{Env}_R(A, B)$ in CV_2

but all the B_k have a different topological type.

Envelopes in CV_n behave nicely with regard to the simplicial structure of CV_n. In particular, their intersection with any simplex is a polytope. We will use this in the proof of Theorem 2.21 to construct locally rigid geodesics. In the next chapter we will then use envelopes to determine the simplicial structure of Outer Space.

Lemma 2.8
Let $A, B \in CV_n$. Then their envelope $\mathrm{Env}_R(A, B)$ is a polytope in the sense that it can be written as a finite union of polytopes in simplices.

Proof. Let α be a witness from A to B and $\Delta \in \mathrm{supp}(\mathrm{Env}_R(A, B))$ be a supporting simplex. For a point $C \in \Delta$ we have by Corollary 1.24 and Lemma 1.22 $C \in \mathrm{Env}_R(A, B)$ if and only if α is a witness from A to C and from C to B. So each $\beta \in \mathrm{cand}(A)$ and $\omega \in \mathrm{cand}(\Delta)$ yields a linear inequality in the simplex Δ: The sets of candidates are finite,

$$\frac{l_C(\alpha)}{l_A(\alpha)} \geq \frac{l_C(\beta)}{l_A(\beta)} \iff \sum_{e_i \in E(C)} l_C(e_i) \cdot \big(l_A(\beta) \cdot \#(e_i, \alpha) - l_A(\alpha) \cdot \#(e_i, \beta)\big) \geq 0 \qquad (\star)$$

$$\frac{l_B(\alpha)}{l_C(\alpha)} \geq \frac{l_B(\omega)}{l_C(\omega)} \iff \sum_{e_i \in E(C)} l_C(e_i) \cdot \big(l_B(\alpha) \cdot \#(e_i, \omega) - l_B(\omega) \cdot \#(e_i, \alpha)\big) \geq 0 \qquad (\star\star)$$

Since the terms $\big(l_A(\beta) \cdot \#(e_i, \alpha) - l_A(\alpha) \cdot \#(e_i, \beta)\big)$ and $\big(l_B(\alpha) \cdot \#(e_i, \omega) - l_B(\omega) \cdot \#(e_i, \alpha)\big)$ do not depend on the edge lengths we get indeed a linear inequality. As there are only finitely many candidates, the points $C \in \Delta$ satisfying the above inequalities form a polytope in Δ.

By Proposition 1.12 there always exists a maximally stretched candidate β or ω. This means α is maximally stretched and thus a witness from A to C and from C to B if and only if the above inequalities are satisfied for all $\beta \in \mathrm{cand}(A)$ and $\omega \in \mathrm{cand}(\Delta)$, that is if C lies in the above polytope. As the support of $\mathrm{Env}_R(A, B)$ is finite this concludes the proof. □

Remark 2.9
In the following, when we talk about the faces of an envelope we only mean faces arising

from equalities in the inequalities ($\star$) and ($\star\star$) above. In particular that means we will exclude the faces arising solely from intersections of the envelope with a simplex.

Corollary 2.10
$\mathrm{Env}_R(A,B)$ is compact.

Proof. Since the diameter of an envelope is bounded, it stays away from missing faces in CV_n and has non-empty intersection with at most finitely many simplices. Hence, the intersection of the envelope with a simplex is closed in the simplicial closure and therefore compact. Thus the envelope is compact as finite union of compact sets. $\square$

Interpreting the proof of Lemma 2.8 we can view each envelope as the intersection of two cones coming from the two end-points, namely the intersection of half-spaces belonging to the inequalities of type ($\star$) and the inequalities of type ($\star\star$). We can see these cones as set of points a geodesic ray can reach if we fix a maximally stretched loop as a coarse direction.

Definition 2.11
Let $S \subseteq F_n$ be a subset. We consider S as a coarse direction or as a set of wanted witnesses. We call the set

$$\mathrm{Env}_R^{\mathrm{out}}(A,S) := \{C \in CV_n \mid S \subseteq W_R(A,C)\}$$

the *out-envelope of A in the direction of S* and

$$\mathrm{Env}_R^{\mathrm{in}}(B,S) := \{C \in CV_n \mid S \subseteq W_R(C,B)\}$$

the *in-envelope of B in the direction of S*. If $S = \{\alpha\}$ is a singleton, we will just write $\mathrm{Env}_R^{\mathrm{out}}(A,\alpha)$ and $\mathrm{Env}_R^{\mathrm{in}}(B,\alpha)$ respectively.

Remark 2.12 (i) By Corollary 1.24 the in- and out-envelopes tell, how we can extend geodesics in either direction and furthermore we have

$$\mathrm{Env}_R(A,B) = \mathrm{Env}_R^{\mathrm{out}}(A,S) \cap \mathrm{Env}_R^{\mathrm{in}}(B,S)$$

for all non-empty $S \subseteq W_R(A,B)$.

(ii) We will see in Proposition 2.16 (iv) and Proposition 2.18 (iv) that in- and out-envelopes are polytopes in each simplex, namely the out-envelopes are parametrised by the ($\star$)-inequalities and the in-envelopes by the ($\star\star$)-inequalities of Lemma 2.8 for some subset of the corresponding candidates.

(iii) By definition the intersection of out-envelopes is the out-envelope of the union of their directions, that is we have $\mathrm{Env}_R^{\mathrm{out}}(A,S) = \bigcap_{\alpha\in S}\mathrm{Env}_R^{\mathrm{out}}(A,\alpha)$ or more generally $\mathrm{Env}_R^{\mathrm{out}}(A,S) = \bigcap_i \mathrm{Env}_R^{\mathrm{out}}(A,S_i)$ for some cover $S = \bigcup_i S_i$. The same also holds for in-envelopes.

By Remark 1.28 we know that two points are almost never joined by unique geodesics. But it turns out that edges of out-envelopes are rigid geodesics. This can be shown by considering an equality instead of an inequality ($\star$) in Lemma 2.8. Any such equality yields a hyperplane in the corresponding simplex. We will see that a geodesic which ends in such a hyperplane already has to stay in the hyperplane the whole time.

Lemma 2.13
Let $A, B \in CV_n$ be two points in Outer Space, $\Delta \in \operatorname{supp}(\operatorname{Env}_R(A, B))$ a supporting simplex of their envelope and $\overline{\Delta}$ its closure in CV_n.

(i) Let $H := H(\beta) \subset \overline{\Delta}$ be a hyperplane which comes from an equality of the form ($\star$) for some $\beta \in \operatorname{cand}(A)$ and $\alpha \in W_R(A, B)$ in Lemma 2.8. Then we have for any $C \in H$ that the envelope $\operatorname{Env}_R(A, C)$ restricted to $\overline{\Delta}$ stays in the hyperplane H, namely we have $\operatorname{Env}_R(A, C) \cap \overline{\Delta} \subseteq H$.

Similarly we have $\operatorname{Env}_R(C', B) \cap \overline{\Delta} \subset H'$ for hyperplanes $H' := H'(\omega)$ coming from equalities of the form ($\star\star$) and $C' \in H'$.

(ii) Let $\gamma\colon [0, 1] \to \overline{\Delta}$ be an edge of the envelope $\operatorname{Env}_R(A, B)$, that is a 1-dimensional face of the polytope $\operatorname{Env}_R(A, B) \cap \overline{\Delta}$. If γ has a witness from A to B as a coarse direction, that is we have $W_R(\gamma(0), \gamma(1)) \cap W_R(A, B) \neq \emptyset$, then γ is a rigid geodesic. In particular all emanating edges from A and incoming edges to B are rigid geodesics.

(iii) Let $S \subset F_n$. Then consecutive edges in $\operatorname{Env}_R^{\text{out}}(A, S)$ form rigid geodesics. Similarly consecutive edges in $\operatorname{Env}_R^{\text{in}}(B, S)$ form rigid geodesics.

Proof. (i) Let $\beta \in \operatorname{cand}(A)$ be the candidate corresponding to H and $C \in H$. By Lemma 1.22 each geodesic γ from A to C must lie completely in H since each point on γ also has β as a witness from A to C. In other words a geodesic from A to B never re-enters a hyperplane of the form ($\star$) and never leaves a hyperplane of the form ($\star\star$).

(ii) Let A' and B' be the endpoints of γ and $\sigma\colon [0, 1] \to \overline{\Delta}$ be any geodesic from A' to B'. By Corollary 1.24 we can extend σ to a geodesic $\tilde{\sigma}\colon [-1, 2] \to CV_n$ from A to B.

Since γ is an edge, we can write it as an intersection of hyperplanes $H(\beta_i)$ coming from equalities in ($\star$) for some $\beta_i \in \operatorname{cand}(A)$ and hyperplanes $H'(\omega_j)$ coming from equalities in ($\star\star$) for some $\omega_j \in \operatorname{cand}(\Delta)$. By (i) $\tilde{\sigma}|_{[-1,1]} \cap \overline{\Delta}$ lies in all hyperplanes $H(\beta_i)$ and $\tilde{\sigma}|_{[0,2]} \cap \overline{\Delta}$ lies in all hyperplanes $H'(\omega_j)$. In particular σ lies on their intersection, i.e. $\sigma = \gamma$.

(iii) Let $\gamma_1, \ldots, \gamma_k$ be consecutive edges in $\operatorname{Env}_R^{\text{out}}(A, S)$, i.e. each γ_i is an edge of the cone $\operatorname{Env}_R^{\text{out}}(A, S) \cap \overline{\Delta_i}$ in some closed simplex $\overline{\Delta_i}$ and the endpoint A_i of γ_i is the starting point of γ_{i+1}. Let γ be a geodesic from A to A_k. By (i) the restriction of γ to the simplex $\overline{\Delta_k}$ must be contained in all hyperplanes containing A_k. Therefore γ must contain γ_k and hence A_{k-1} lies on γ. Inductively, γ has to go through all $A_1, \ldots, A_k$ and all γ_i, so we already have $\gamma = \gamma_1 * \cdots * \gamma_k$ which means that $\gamma_1 * \cdots * \gamma_k$ is the unique geodesic from A to A_k.

□

We can now use this to construct piecewise rigid geodesics between any two points.

Theorem 2.14
For any $A, B \in CV_n$ there exist rigid geodesics $\gamma_1, \ldots, \gamma_k$ such that their concatenation $\gamma := \gamma_1 * \gamma_2 * \cdots * \gamma_k$ is a geodesic from A to B.

Proof. Let $A, B \in CV_n$ be any two points. We will construct γ_i for i inductively starting with $A_1 := A$:

Starting at A_i choose any consecutive edges $\sigma_1, \ldots, \sigma_{k_i}$ in $\mathrm{Env}_R(A_i, B)$ until they hit the first time a hyperplane $H'(\omega)$ coming from an equality of type $(\star\star)$ and denote this point with A_{i+1}. This means that $\sigma_1, \ldots, \sigma_{k_i}$ are actually edges of $\mathrm{Env}_R^{\mathrm{out}}(A_i, W_R(A_i, B))$ (σ_{k_i} might be only a part of an edge). By Lemma 2.13 (iii) their concatenation $\gamma_i := \sigma_1 * \cdots * \sigma_{k_i}$ is a rigid geodesic from A_i to A_{i+1}. Since there are only finitely many edges in $\mathrm{Env}_R(A_i, B)$ such a sequence of edges is always finite.

As each A_{i+1} lies in the envelope of $\mathrm{Env}_R(A_i, B)$ we have by Corollary 1.24 that any concatenation $\gamma_1 * \cdots * \gamma_k$ of these rigid geodesics is again a geodesic.

To see that this induction stops after finitely many steps consider the set $S := \bigcup_{C \in \mathrm{Env}_R(A,B)} \mathrm{cand}(C)$, i.e. S is the set of all possible candidates in the support of $\mathrm{Env}_R(A, B)$. By construction there exists an $\omega \in S \cap W_R(A_{i+1}, B)$ with $\omega \notin W_R(A_i, B)$ so we get by Lemma 1.22 the strictly increasing sequence:

$$S \cap W_R(A_0, B) \subsetneq S \cap W_R(A_1, B) \subsetneq \cdots \subseteq S$$

The set S is finite, as there are only finitely many candidates per simplex and the support of $\mathrm{Env}_R(A, B)$ is finite by Lemma 2.6. In particular this sequence and hence the algorithm has to stop at some point. □

Apart from being useful for the construction of piecewise rigid geodesics the out-envelopes have other interesting properties. For example we can use them to show that we can slightly vary two points such that we at most decrease the number of their candidate witnesses. The reason for this is that by continuously varying the base point A in a simplex we also continuously vary the out- and in-going envelopes. This is due to the fact that they are parametrised by the inequalities $(\star)$ and $(\star\star)$ in Lemma 2.8.

Lemma 2.15
Let $A, B \in CV_n$ and $S := \{\alpha \in \mathrm{cand}(\Delta) \mid \Delta \text{ is a simplex in } CV_n \text{ and } A \in \overline{\Delta}\}$ be the candidates close to A. Then there exist neighbourhoods $U_A \ni A$ and $U_B \ni B$ such that for all $A' \in U_A$ and $B' \in U_B$ we have $W_R(A', B') \cap S \subseteq W_R(A, B) \cap S$.

In particular, if A lies in a maximal simplex we can always choose $U_A \subset \Delta(A)$ and hence at most decrease the candidate witnesses, i.e. we have $CW_R(A', B') \subseteq CW_R(A, B)$ for all $A' \in U_A$ and $B' \in U_B$.

Proof. Let $A', B' \in CV_n$ be two points and $\alpha \in W_R(A, B)$. We will use that a $\beta \in S$ can only be a witness from A' to B' if we have $\frac{l_{B'}(\alpha)}{l_{A'}(\alpha)} \leq \frac{l_{B'}(\beta)}{l_{A'}(\beta)}$, that is in the $(\star)$-inequality in Lemma 2.8 the sum is less or equal zero.

We can choose U_A small enough such that $\text{cand}(A') \subseteq S$ for all $A' \in U_A$, that is U_A only intersects simplices whose closure contain A. We consider the finite set S instead of $\text{cand}(A')$ in the $(\star)$-inequalities. Varying A continuously varies the coefficients of $l_B(e_i)$ in the $(\star)$-inequality continuously. Since S is finite we can choose U_A and U_B in such a matter that if for any $\beta \in S$ the sum $l_A(\beta)l_B(\alpha) - l_A(\alpha)l_B(\beta)$ is strictly greater than zero, it stays so for all $A' \in U_A$ and $B' \in U_B$. That means that any $\beta \in S$ which was not already a witness from A to B is less stretched from A' and B' than α and hence does not become a witness. □

Previous Lemma 2.15 means that out-envelopes depend continuously on the base point. As another useful property we will now see that the out-envelopes of a given point yield a partition of Outer Space into polytopes. That is two out-envelopes of the same point only intersect at their faces and faces are again out-envelopes. Furthermore, it is enough to restrict ourselves to out-envelopes in coarse direction of candidates. We will see that this is also locally true for in-envelopes.

Proposition 2.16
Let $A \in CV_n$, then we have:

(i) $CV_n = \bigcup_{\alpha \in \text{cand}(A)} \text{Env}_R^{\text{out}}(A, \alpha)$

(ii) $\{A\} = \bigcap_{\alpha \in \text{cand}(A)} \text{Env}_R^{\text{out}}(A, \alpha) = \text{Env}_R^{\text{out}}(A, \text{cand}(A))$

(iii) For all $\alpha \in \text{cand}(A)$ the interior $\text{Env}_R^{\text{out}}(A, \alpha)^{\text{int}}$ is non-empty.

(iv) For all subsets $M \subseteq F_n$ and all simplices $\Delta \in \text{supp}(\text{Env}_R^{\text{out}}(A, M))$, there exists a subset of candidates $S \subseteq \text{cand}(A)$ such that the out-envelopes of M and S are the same in Δ, i.e. $\Delta \cap \text{Env}_R^{\text{out}}(A, M) = \Delta \cap \text{Env}_R^{\text{out}}(A, S)$.

(v) Let $S_1, S_2 \subseteq F_n$ and let Δ be a simplex in CV_n, then $\text{Env}_R^{\text{out}}(A, S_1) \cap \text{Env}_R^{\text{out}}(A, S_2) \cap \Delta$ is a face of $\text{Env}_R^{\text{out}}(A, S_1) \cap \Delta$.

Proof. (i) Follows directly from the fact that for each $B \in CV_n$ there exists a candidate witness $\alpha \in CW_R(A, B)$.

(ii) Let $B \in \text{Env}_R^{\text{out}}(A, \text{cand}(A))$ be a point in the out-envelope and $\beta \in W_R(B, A)$ be a witness from B to A. By (iv) we have $\text{Env}_R^{\text{out}}(A, \beta) \cap \Delta(B) = \text{Env}_R^{\text{out}}(A, S) \cap \Delta(B)$ for some $S \subseteq \text{cand}(A)$. By Remark 2.12 we have $B \in \text{Env}_R^{\text{out}}(A, \text{cand}(A)) \subseteq \text{Env}_R^{\text{out}}(A, S)$. So we have that β is maximally stretched from A to B and from B to A. In particular we have $\Lambda_R(A, B) = \Lambda_R(B, A)^{-1}$ and hence B and A is the same point in CV_n.

(iii) We will construct an open set contained in $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha)$.
If A is in a maximal simplex, let $\varepsilon > 0$ be smaller than each edge length of A and $\alpha \in \mathrm{cand}(A)$ be a candidate of A. Recall that as A is in a maximal simplex it does not contain any figures of eight. Let B be an element in CV_n obtained from A by changing the edge lengths in the following way:

- Each edge not contained in α is shrunk by more than ε, i.e. $l_B(e) < l_A(e) - \varepsilon$.
- If α is a simple loop: Let n_α be the number of edges in α. Then each edge contained in α is stretched but at most by $\varepsilon/2n_\alpha$, i.e. $l_A(e) < l_B(e) < l_A(e) + \varepsilon/2n_\alpha$.
- If α is a barbell: We denote by α_1, α_2 the two circles of α and by ρ the barbell handle as edge-paths, i.e. looking at α as a sequence of edges we have $\alpha = \alpha_1 * \rho * \alpha_2 * \rho^{-1}$. We shrink each edge contained in one of the two circles and the circles are shrunk in total by at most $\varepsilon/4$, i.e. $l_A(\alpha_i) - \varepsilon/4 < l_B(\alpha_i) < l_A(\alpha_i)$. Let n_ρ be the number of edges in ρ, then stretch each edge contained in ρ by less than $\varepsilon/2n_\rho$ such that ρ is stretched more than $\varepsilon/4$, that is we have then $l_B(\alpha) > l_A(\alpha)$.

We have now that each candidate $\beta \in \mathrm{cand}(A)$ which is not completely contained in α crosses an edge which is shrunk in B. As we stretch the other edges in sum less than $\varepsilon/2$ we have that such a β is shrunk from A to B. Thus α, and in the case of the barbell its counterpart with a "flipped" circle, is the only and hence maximally stretched path from A to B, i.e. $B \in \mathrm{Env}_R^{\mathrm{out}}(A, \alpha)$ holds. As we had strict inequalities for the edge lengths, the set of such constructed B is open and so the claim follows. In particular we have shown that $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta(A)$ has full dimension when A lies in a maximal simplex.

Let now A be not in a maximal simplex. We can construct similar to B from above an $A' \in \mathrm{Env}_R^{\mathrm{out}}(A, \alpha)$ in a maximal simplex with $\alpha \in \mathrm{cand}(A')$. Namely we shrink all edges of A which are not covered by α by some ε and relax all vertices of valency larger than three with edges of some length $\varepsilon/2k$ along α, where k is the number of newly introduced edges and along α means that α passes over the new edges whenever possible. We only need to take care in case α is a figure of eight to relax it into a barbell so that we still have $\alpha \in \mathrm{cand}(A')$. Since we have then $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \subseteq \mathrm{Env}_R^{\mathrm{out}}(A', \alpha)$ the claim follows.

(iv) By Remark 2.12 (iii) we can assume without loss of generality that $M = \{\alpha\}$ is a singleton. Let $B \in \mathrm{relint}(\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta)$ be a point in the relative interior of the polytope $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta$. For $S := CW_R(A, B)$ we claim that $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta = \mathrm{Env}_R^{\mathrm{out}}(A, S) \cap \Delta$. To prove the claim we show that α has in the corresponding out-envelopes the same stretching as any $\beta \in S$.

Let $B + V \subset \Delta$ be the affine subset in Δ spanned by $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta$ and let $v \in V$ be small enough such that $B \pm v \in \mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta$ holds. Here we use that B is in the relative interior of the out-envelope. For $\beta \in S$ the inequality in $(\star)$ is an equality for B, i.e. we have $\frac{l_B(\alpha)}{l_A(\alpha)} = \frac{l_B(\beta)}{l_A(\beta)}$.

Recall that the length functions $l_C(\alpha)$ linearly depend on the edge lengths of C. Since $B+v$ and $B-v$ are in $\mathrm{Env}^{\mathrm{out}}_R(A,\alpha)$ we also have the equalities $\frac{l_{B\pm v}(\alpha)}{l_A(\alpha)} = \frac{l_{B\pm v}(\beta)}{l_A(\beta)}$. Elsewise we would have $\frac{l_{B+v}(\alpha)}{l_A(\alpha)} > \frac{l_{B+v}(\beta)}{l_A(\beta)}$ or $\frac{l_{B-v}(\alpha)}{l_A(\alpha)} > \frac{l_{B-v}(\beta)}{l_A(\beta)}$ which contradicts $B \pm v \in \mathrm{Env}^{\mathrm{out}}_R(A,\alpha)$.

By the linearity of the length function we obtain equality in ($\star$) for α and $\beta \in S$ for all multiples of v. In particular this equality holds for all $B+V$, i.e. for all $\beta \in S$ and $C \in B+V$ we have $\frac{l_C(\alpha)}{l_A(\alpha)} = \frac{l_C(\beta)}{l_A(\beta)}$. As $B+V$ was the affine subspace spanned by $\mathrm{Env}^{\mathrm{out}}_R(A,\alpha)$ we have $\mathrm{Env}^{\mathrm{out}}_R(A,\alpha) \subseteq \mathrm{Env}^{\mathrm{out}}_R(A,\alpha)$ and thus $\mathrm{Env}^{\mathrm{out}}_R(A,\alpha) \cap \Delta \subseteq \mathrm{Env}^{\mathrm{out}}_R(A,S)$.

Similarly we prove the other inclusion. Let $B+W$ be the affine subset spanned by $\mathrm{Env}^{\mathrm{out}}_R(A,S) \cap \Delta$ in Δ, and let $w \in W$ be small enough. As $B+W$ is spanned by $\mathrm{Env}^{\mathrm{out}}_R(A,S)$, we assume without loss of generality that $B+w \in \mathrm{Env}^{\mathrm{out}}_R(A,S)$ holds. If we have $B+w \in \mathrm{Env}^{\mathrm{out}}_R(A,\alpha)$ for all small enough $w \in W$, we conclude as before equality in ($\star$) for α and all $\beta \in S$ for all $C \in \mathrm{Env}^{\mathrm{out}}_R(A,S) \subseteq B+W$ and thus $\mathrm{Env}^{\mathrm{out}}_R(A,S) \cap \Delta \subseteq \mathrm{Env}^{\mathrm{out}}_R(A,\alpha)$.

Assume $B+w \notin \mathrm{Env}^{\mathrm{out}}_R(A,\alpha)$, then we have $\frac{l_{B+w}(\alpha)}{l_A(\alpha)} < \frac{l_{B+w}(\beta)}{l_A(\beta)}$ for one and thus for all $\beta \in S$. By the linearity of the lengths we have hence $\frac{l_{B-w}(\alpha)}{l_A(\alpha)} > \frac{l_{B-w}(\beta)}{l_A(\beta)}$ and in particular $S \cap CW_R(A, B-w) = \emptyset$. But by choosing $|w|$ small enough we have by Lemma 2.15 $CW_R(A, B-w) \subset S$, which contradicts $CW_R(A, B-w) \neq \emptyset$.

(v) Follows directly from (iv) and the inequalities ($\star$) in Lemma 2.8, since the defining half spaces of the polytope come from the candidates.

□

Remark 2.17
Statement (iii) does not hold in reduced Outer Space, since we might not be able to relax a figure of eight into a barbell. As example in CV_2^{red} let A be the rose with petals α and β. Then the out-envelope $\mathrm{Env}^{\mathrm{out}}_R(A,\alpha\beta)$ of its figure of eight has dimension one: For any $B \in CV_2^{\mathrm{red}}$ the two loops α and β have to intersect in at least one vertex or edge, hence we have $l_B(\alpha\beta) \leq l_B(\alpha) + l_B(\beta)$. We use here that α and β are generators of F_2, hence we can assume by Stallings folding that they can be simultaneously cyclically reduced in B by conjugation. By Lemma 1.14 we have then that $\alpha\beta$ is exactly then a witness if the equality $l_B(\alpha\beta) = l_B(\alpha) + l_B(\beta)$ holds and α and β are also witnesses. As α and β form a basis of F_2, their edge counts can not be multiples in any simplex Δ. Hence, their equality in the ($\star$)-inequalities of Lemma 2.8 yield a proper restriction, that is we have $\dim(\mathrm{Env}^{\mathrm{out}}_R(A,\alpha\beta)) \leq \dim(\Delta) - 1 \leq 1$. Similarly, let $\omega \in F_2$ be an element that is not a power of α or β. Then the out-envelope $\mathrm{Env}^{\mathrm{out}}_R(A,\omega)$ has at most dimension one.

For the in-envelopes we get similar results as in Proposition 2.16.

Proposition 2.18
Let $B \in CV_n$ and $S := \{\alpha \in F_n \mid \alpha \text{ can be extended to a free generating set of } F_n\} = \bigcup_{A \in CV_n} \mathrm{cand}(A)$ the set of possible candidates, then we have:

(i) $CV_n = \bigcup_{\alpha \in S} \mathrm{Env}^{\mathrm{in}}_R(B, \alpha)$

(ii) $\{B\} = \bigcap_{\alpha \in \mathrm{cand}(B)} \mathrm{Env}^{\mathrm{in}}_R(B, \alpha) = \mathrm{Env}^{\mathrm{in}}_R(B, \mathrm{cand}(B))$

(iii) For all $\alpha \in S$ the interior $\mathrm{Env}^{\mathrm{in}}_R(B, \alpha)^{\mathrm{int}}$ is non-empty.

(iv) For all sets of witnesses $M \subseteq F_n$ and all simplices $\Delta \in \mathrm{supp}(\mathrm{Env}^{\mathrm{out}}_R(B, M))$, there exists a subset of candidates $S' \subseteq \mathrm{cand}(\Delta)$ such that $\Delta \cap \mathrm{Env}^{\mathrm{in}}_R(B, M) = \Delta \cap \mathrm{Env}^{\mathrm{in}}_R(B, S')$.

(v) Let $S_1, S_2 \subseteq F_n$ and let Δ be a simplex in CV_n, then $\mathrm{Env}^{\mathrm{in}}_R(B, S_1) \cap \mathrm{Env}^{\mathrm{in}}_R(B, S_2) \cap \Delta$ is a face of $\mathrm{Env}^{\mathrm{in}}_R(B, S_1) \cap \Delta$.

Proof. All statements except (iii) are proven as in Proposition 2.16.

For (iii) let $\alpha \in S$. As α is a candidate it can be extended to a basis $(\alpha, \alpha_2, \ldots, \alpha_n)$ of F_n. Consider the marked rainbow graph A as in Figure 14, where the two edges covered by α have length less than some small $\varepsilon > 0$ and all the other edges have length greater than 1.

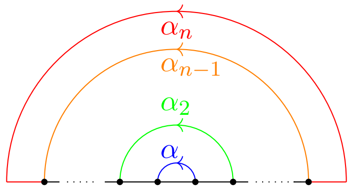

Figure 14: marked rainbow graph

Hence, all candidates of A except α have length greater than 2 and α has length less than 2ε. Since there are only finitely many candidates, we can choose ε small enough such that $\frac{l_B(\alpha)}{2\varepsilon} > \frac{l_B(\beta)}{2}$ holds for every other candidate $\beta \in \mathrm{cand}(A)$. As we have $\frac{l_B(\alpha)}{l_A(\alpha)} > \frac{l_B(\alpha)}{2\varepsilon} > \frac{l_B(\beta)}{2} > \frac{l_B(\beta)}{l_A(\beta)}$ and there always exists a candidate witness, we have that all the rainbow graphs with the above edge lengths lie in $\mathrm{Env}^{\mathrm{in}}_R(A, \alpha)$. As the conditions on the edge lengths are strict inequalities these rainbow graphs form an open set. □

As a direct application of Proposition 2.18 (iii) and its proof we get a slightly weaker counterpart to Lemma 1.36, namely we can extend a given geodesic backwards.

Corollary 2.19
Let $B, C \in CV_n$ be two points and $\gamma\colon [0,1] \to CV_n$ be a geodesic from B to C. Then there exists a geodesic $\sigma\colon \mathbb{R}_{\leq 1}$ such that $\sigma|_{[0,1]} = \gamma$ and for any $L \geq 0$ there exists a $t < 0$ with $d_R(\sigma(t), B) > L$.

While Lemma 2.13 (iii) tells us that consecutive edges of out- and in-envelope are rigid geodesics, we will see that in fact all rigid geodesics are of that form. To prove this, we will need the following lemma.

Lemma 2.20
Let $\overline{\Delta} \subset CV_n$ be a closed simplex, $A \in \overline{\Delta}$ and $S \subseteq F_n$ such that $\mathcal{E} := \mathrm{Env}_R^{\mathrm{out}}(A, S) \cap \overline{\Delta}$ has dimension k. Then for each relative interior point $B \in \mathrm{relint}(\mathrm{Env}_R^{\mathrm{out}}(A, S)) \cap \overline{\Delta}$ we have $\dim(\mathrm{Env}_R(A, B) \cap \overline{\Delta}) = k$. In particular, if we have $k \geq 2$, then there exists no rigid geodesic from A to B.

The analogue statement holds for in-envelopes.

Proof. Let $B \in \mathrm{relint}(\mathrm{Env}_R^{\mathrm{out}}(A, S)) \cap \overline{\Delta}$, $S' := W_R(A, B)$ and $\mathcal{E}' := \mathrm{Env}_R^{\mathrm{out}}(A, S') \cap \overline{\Delta}$. By $B \in \mathcal{E}$ we have $S \subseteq S'$ and thus $\mathcal{E}' \subseteq \mathcal{E}$. Applying Proposition 2.16 (v) to $\mathcal{E}$ and $\mathcal{E}'$ we get that $\mathcal{E}'$ is a face of $\mathcal{E}$. But $\mathcal{E}'$ contains an interior point B of $\mathcal{E}$ and thus we have $\mathcal{E}' = \mathcal{E}$.

Now let U_A be as in Lemma 2.15 and $A' \in U_A \cap \mathcal{E}$. By $A' \in \mathcal{E} = \mathcal{E}'$ we get $S' \subseteq W_R(A, A')$ and by $A' \in U_A$ and Lemma 2.15 we have $CW_R(A', B) \subseteq S'$. But then Corollary 1.24 implies that A' lies on a geodesic from A to B, i.e. $A' \in \mathrm{Env}_R(A, B) \cap \overline{Delta}$.

This means we have $U_A \cap \mathcal{E} \subset \mathrm{Env}_R(A, B) \cap \overline{\Delta} \subset \mathcal{E}$ and thus $k = \dim(\mathcal{E}) = \dim(\mathrm{Env}_R(A, B) \cap \overline{\Delta})$. In particular, for $k \geq 2$ there exist several geodesic from A to B as otherwise $\mathrm{Env}_R(A, B)$ would be the unique geodesic and hence would have dimension 1.

Similar we get the statement for in-envelopes. □

We have now the tools to classify rigid geodesics in terms of envelopes.

Theorem 2.21
Let $I \subset \mathbb{R}$ be any interval, $t \in I$ and $\gamma \colon I \to CV_n$ a geodesic. Then the following is equivalent:

(i) γ is a rigid geodesic.

(ii) For all $t \in I$, $\gamma|_{\geq t}$ is the concatenation of consecutive edges of an out-envelope of $\gamma(t)$ and $\gamma|_{\leq t}$ is the concatenation of consecutive edges of in-envelopes of $\gamma(t)$.

Proof. Recall Definition 1.29 that γ is rigid, if each arc $\gamma|_{[s,t]}$ is rigid for any $s \leq t \in I$. By Lemma 2.13 (iii) we have then (ii) $\Rightarrow$ (i).

For the converse assume that for some $t \in I$ $\gamma|_{\geq t}$ does not lie on the edges of an out-envelope. Let $s \geq t$ be minimal such that $\gamma|_{>s}$ does not lie on edges of an out-envelope of $\gamma(t)$. Let $\varepsilon > 0$ be small enough such that for $s' := s + \varepsilon$ the points $\gamma(s)$ and $\gamma(s')$ lie in the same closed simplex $\overline{\Delta}$ and that $\gamma(s')$ does not lie on an edge of an out-envelope. Let $S := W_R(\gamma(t), \gamma(s'))$ be the set of witnesses from $\gamma(t)$ to $\gamma(s')$. As $\gamma(s')$ does not lie on an edge we have that the envelope $\mathcal{E} := \mathrm{Env}_R^{\mathrm{out}}(\gamma(t), S)$ has at least dimension two in $\overline{\Delta}$ and $\gamma(s')$ lies in the relative interior of $\mathcal{E}$. By Lemma 2.20 there exist several geodesics from $\gamma(s)$ to $\gamma(s')$, which contradicts that γ is rigid.

Similarly, we get that γ is not rigid, if $\gamma|_{\leq t}$ is not the concatenation of consecutive edges of in-envelopes. □

Remark 2.22
In Theorem 2.21 observe that while each bounded segment $\gamma|_{[s,t]}$ lies in a single in-envelope of $\gamma(t)$, we will see that the whole negative part $\gamma|_{\leq t}$ might not lie in a single in-envelope of

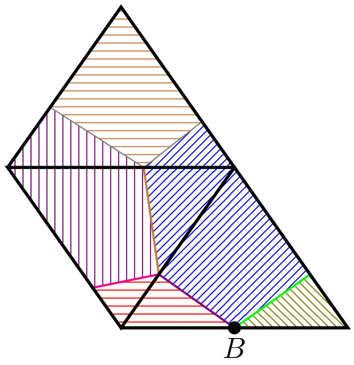

Figure 15: In-envelopes for a figure of eight B in CV_2^{red}.

$\gamma(t)$. That means we have to consider in Theorem 2.21 edges of possibly several different in-envelopes but only need the edges of a single out-envelope. The reason one out-envelope is enough is that for a given geodesic $\gamma\colon I \to CV_n$ and $t \in I$ we have by Lemma 1.22 $CW_R(\gamma(t), \gamma(s_2)) \subset CW_R(\gamma(t), \gamma(s_1)) \neq \emptyset$ for all $t \leq s_1 \leq s_2$. Since $CW_R(\gamma(t), \gamma(s))$ is finite for all $s > t$ this implies $CW_R(\gamma|_{>t}) := \bigcap_{s>t} CW_R(\gamma(t), \gamma(s)) \neq \emptyset$ and $CW_R(\gamma|_{>t}) \subseteq CW_R(\gamma(t), \gamma(s)) \subseteq$ for all $s > t$, thus the out-envelope $\mathrm{Env}_R(\gamma(t), CW_R(\gamma|_{>t}))$ suffices.

This argument does not hold for the negative direction $\gamma|_{<t}$, since the candidates here depend on $\gamma(s)$ and not $\gamma(t)$. We might even get that $\bigcap_{s<t} W_R(\gamma(s), \gamma(t))$ is an empty set. We can construct such a counterexample in CV_2^{red}. The idea for the construction is that in each theta-simplex we have three in-envelopes coming from the three candidates. That means each time we cross a face in CV_2^{red} along the edge of two in-envelopes we have the two in-envelopes "fanning out" to three in-envelopes. So at each face we can choose two ways how to continue our incoming geodesic. Choosing appropriately we can then eliminate all elements as witnesses:

Enumerate all possible witness as $(\alpha_i)_{i\in\mathbb{N}}$ and start with any maximal simplex $\Delta_0 \subset CV_2^{\text{red}}$. Let $A_0 \in \overline{\Delta_0}$ be a figure of eight graph. A short calculation shows that we have three in-envelopes of A_0 in Δ_0 where the middle in-envelope comes from a figure of eight in A_0. Furthermore, the two rigid geodesics in Δ_0 ending in A_0 will pass the two other faces of $\overline{\Delta_0}$. That means we can choose inductively A_i to be the figure of eight in $\overline{\Delta_0}$ such that the geodesic γ_i from A_i to A_{i-1} is rigid and we have $\alpha_i \notin W_R(A_i, A_{i-1})$ unless α_i is a figure of eight in A_i. In this case observe that α_i is not a figure of eight in A_{i+1}, so we can simply rearrange α_i with α_{i+1}. We define then Δ_i to be the other maximal simplex with $A_i \in \overline{\Delta_i}$.

We now need to show that that the concatenation of the γ_i yields a geodesic. Let $\alpha \in CW_R(A_i, A_0)$ be a candidate witness. By choosing i large enough we can assume that $\alpha \notin \mathrm{cand}(\Delta_0)$. Similar to Remark 2.17 we have $\dim(\mathrm{Env}_R^{\text{in}}(A_0, \alpha) \cap \Delta_0) < 2$. As CV_2^{red} is simply connected, any geodesic from A_i to A_0 has to pass through Δ_0 and $\Delta(A_1)$. By Proposition 2.18 (v) the envelope then has to be the rigid geodesics from A_1 to A_0 in Δ_0. Inductively, we get that the concatenation $\gamma_{i-1} * \cdots * \gamma_0$ is the unique geodesic from A_i to A_0. In particular all the γ_i concatenate to a rigid geodesic γ which has by construction $\bigcap_{s\leq t} W_R(\gamma(s), \gamma(t)) = \emptyset$.

3 Simplicial structure of CV_n^{red}

In this section we will see how to distinguish faces in CV_n^{red} by the use of envelopes. The important observation we use is that envelopes may have different dimensions in different simplices. You can see an example of this behaviour in the third image in Figure 13. In reduced Outer Space we can construct such envelopes near any face as follows:

Lemma 3.1
Let $C \in CV_n^{\text{red}}$ be in a face, i.e. the underlying graph of C has a vertex of degree at least four. Furthermore, let $U \ni C$ be a neighbourhood of C in $CV_n{}^{\text{red}}$. Then there exist open sets $U_A, U_B \subset U$ such that for any $A \in U_A, B \in U_B$ the envelope $\text{Env}_R(A, B)$ has near A full dimension while it has lower dimension near B, i.e. we have $\dim(\text{Env}_R(A, B) \cap U_A) = 3n - 4 > \dim(\text{Env}_R(A, B) \cap U_B)$.

Proof. Since we can slightly relax all vertices of valency four in any open set, we can without loss of generality assume that C has exactly one vertex v of valency four. Let us look at the star around v:

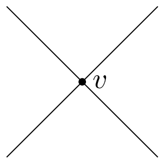

Since $C \in CV_n^{\text{red}}$, there exists no separating edge and hence we find embedded circles α and β passing through v:

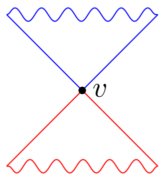

Furthermore, we can assume that α and β are disjoint apart from v by cutting out common edges and glueing the paths back together:

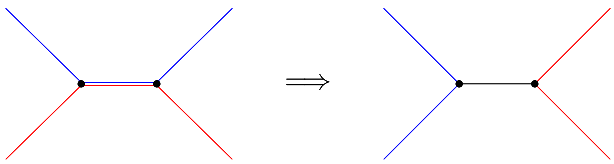

Giving α and β an orientation, we can see them as elements in F_n.
For small enough $\varepsilon > 0$ we construct the following graphs in U:

- A is obtained by relaxing v to an edge (v_1, v_2) of length ε as in figure 16 and keeping the rest as in C.
- B is obtained by relaxing v to an edge (v_1, v_2) of length ε in a different manner than in A by swapping the edges belonging β. Thus we effectively reverse the

orientation at which β passes through (v_1, v_2) when compared to A (see Figure 16). Additionally we shrink each edge which does not lie in α or β by 2ε, i.e. we have $l_B(e) = l_C(e) - 2\varepsilon$ for such an edge e.

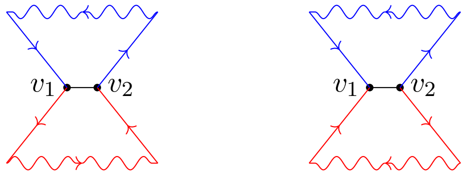

Figure 16: v relaxed to get A and B

Observe that A and B now lie in maximal simplices. Since topologically A and B are the same outside of v and most of the edges of B are shorter than in A, each candidate of A can only be stretched from A to B if it crosses v_1 or v_2 and does not cross edges outside of α and β. Hence, $\alpha\beta$ is the only maximally stretched candidate from A to B. Considering the inequalities for the envelope around A, we have that in a small neighbourhood around A the inequalities of type $(\star\star)$ are always satisfied and as in Proposition 2.16 (iii) it has full dimension $3n - 4$.

On the other hand let $D \in \mathrm{Env}_R(A, B)$ be close to B, i.e. with the same topological type as B. By Lemma 1.22 $\alpha\beta$ has to be a witness from D to B. By Lemma 1.14 α and β are also witnesses from D to B. But this yields a non trivial equality condition in $(\star\star)$ for the edge lengths of D, namely $\frac{l_B(\alpha)}{l_D(\alpha)} = \frac{l_B(\beta)}{l_D(\beta)}$, thus $\mathrm{Env}_R(A, B)$ has close to B at most dimension $3n - 4 - 1$.

By Lemma 2.15 we can choose small enough neighbourhoods $U_A \ni A$ and $U_B \ni B$ such that we do not change the candidate witness, that is we have $CW_R(A', B') = \{\alpha\beta\}$ and hence as before $\dim(\mathrm{Env}_R(A', B') \cap U_A) > \dim(\mathrm{Env}_R(A', B') \cap U_B)$ for all $A' \in U_A$ and $B' \in U_B$. $\square$

We saw an example for such a dimension change in Figure 13. In that example A and B are theta-graphs with marking as depicted in Figure 16 and $\alpha\beta$ is the candidate witness from A to B.

Remark 3.2
The argumentation of Lemma 3.1 does not necessarily hold in non-reduced Outer Space, for example consider C as a doubled barbell graph in Figure 17.

Regardless on how we resolve the 4-valent vertex, we will always get the same sets of candidates. Fixing the edge lengths of C we can now easily find a neighbourhood of C which contains only the topological types gained by slightly stretching the 4-valent vertex.

Recall that envelopes are polytopes in a simplex and thus have a fixed dimension in each simplex. With previous Lemma 3.1 and the fact that envelopes are preserved under isometries we get that isometries of CV_n^{red} send maximal simplices to maximal simplices. The lower dimensional skeleton of CV_n^{red} is preserved for topological reasons and we get:

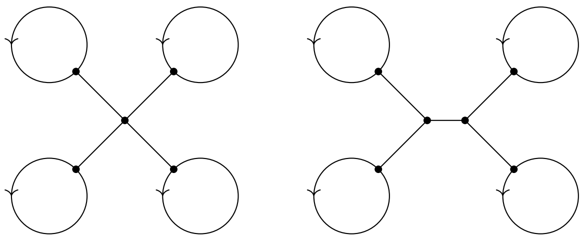

Figure 17: doubled barbell graph and its adjacent topological type

Theorem 3.3
An isometry in regard to the asymmetric metric of CV_n^{red} is simplicial.

For the symmetric version of this theorem, which also implies the last theorem, we will need a little bit more work. We will first show that almost all points $A, B \in \Delta$ in the same maximal simplex Δ have a symmetric envelope $\text{Env}(A, B)$ with full dimension $3n - 4$. On the other hand by Lemma 3.1 there are open sets at faces, where this is not satisfied. To quantify what "almost all points" means we will introduce a notion of general position.

Definition 3.4
Let $A, B \in CV_n$ be two points in maximal simplices. We say B is in *general position* to A if there exist neighbourhoods $U_A \ni A$ and $U_B \ni B$ such that for all $A' \in U_A$ and $B' \in U_B$ the sets of candidate witnesses $CW_R(A, B) = CW_R(A', B')$ are the same. Otherwise we say B is in *special position* to A.

It should be remarked that our notion of general position is not symmetric, i.e. B can be in general position to A while A is in special position to B. Before we go into more detail how to use this notion, we want to point out that, as the name suggests, almost every pair is in general position. More precisely we have:

Proposition 3.5
Let $A \in CV_n$ be in a maximal simplex then the set of points B in general position to A is dense and open. The same statement holds if we fix B and vary A.

Proof. The open property follows directly from the definition. The dense property follows from the following Lemma 3.6 and Proposition 2.16. □

We can find points in general position with the help of in- or out-envelopes, namely a point is in general position to another if the two points lie in the interior of the others in-envelope and out-envelopes, respectively:

Lemma 3.6
Let $A, B \in CV_n$ be points in maximal simplices, then the following are equivalent:

(i) B is in general position to A.

(ii) $B \in \mathrm{Env}_R^{\mathrm{out}}(A, \alpha)^{\mathrm{int}}$ for an $\alpha \in F_n$.

(iii) $A \in \mathrm{Env}_R^{\mathrm{in}}(B, \alpha)^{\mathrm{int}}$ for an $\alpha \in F_n$.

Proof. "(i)⇒(ii) and (iii)" follows directly from the definition of general position for every $\alpha \in CW_R(A, B)$.

"(iii)⇒(i)": By Proposition 2.18 (iv) we can assume that α is a candidate of A. If there exists no other candidate witness $\beta \in CW_R(A, B)$, we are done by Lemma 2.15. Hence, assume $\beta \in CW_R(A, B) \setminus \{\alpha\}$ is another candidate witness. Since A is in the interior of $\mathrm{Env}_R^{\mathrm{in}}(B, \alpha)$, we have by Proposition 2.18 (v) $\mathrm{Env}_R^{\mathrm{in}}(B, \beta) \cap \Delta(B) = \mathrm{Env}_R^{\mathrm{in}}(B, \alpha) \cap \Delta(B)$. Consider the edge-counts of α and β in A in the ($\star\star$)-equality. As we have that the equality holds for any C in a small neighbourhood of A, the coefficient in the ($\star\star$)-equality for every edge length have to be zero, that is we gave $\big(l_B(\alpha) \cdot \#_A(e_i, \omega) - l_B(\omega) \cdot \#_A(e_i, \alpha)\big) = 0$. In particular the edge-counts of α and β in A are multiples of each other and thus α and β are in A either the same candidate or a barbell and its counterpart, that is we reverse the orientation of a petal e.g. $\omega_1\omega_2$ and $\omega_1\omega_2^{-1}$) – as A is in the a maximal simplex there exist no 4-valent vertices and hence no figures of eight. Furthermore we have shown that $CW_R(A, B) = \{\alpha, \beta\}$ consists only of the barbell and its counterpart β.

We will now show that α and β also have the the same edge-counts in B. Let $h\colon A \to B$ be an optimal change of marking map (see Definition 1.31). Since α and β are witnesses, their images in B under h are reduced paths. As they have the same edge counts in A, this implies they have also the same edge counts in B. In particular they also have the same lengths in the simplices of A and B, i.e. $l_{A'}(\beta) = l_{A'}(\alpha)$ and $l_{B'}(\beta) = l_{B'}(\alpha)$ for all $A' \in \Delta(A), B' \in \Delta(B)$.

By Lemma 2.15 we can choose neighbourhoods U_A and U_B small enough such that $CW_R(A', B') \subseteq CW_R(A, B) = \{\alpha, \beta\}$. As α and β are equally stretched between two points in the corresponding simplices we conclude $CW_R(A', B') = CW_R(A, B)$.

"(ii)⇒(i)": As before assume $\beta \neq \alpha$ with $\beta \in CW_R(A, B)$ exists. We show that α and β are are again a barbell and its counterpart in A. As in (iii) we have that the edge-counts of α and β in B are multiples of each other. In particular the letter counts of the word $\alpha * \beta \in F_n$ are multiples of some number $k \geq 2$. This implies that $\alpha * \beta$ and hence the pair α and β can not be extended to a basis of F_n: Any basis of F_n is send by abelianisation to a basis of $\mathbb{Z}^n$, but $\alpha * \beta$ would be a multiple of k in $\mathbb{Z}^n$ and hence can not be extended to a basis of $\mathbb{Z}^n$.

Assume that α and β are two different candidates and the pair (α, β) is not a barbell or a figure of eight and its counterpart. Then the set $\{\alpha, \beta\}$ can be extended to a basis of F_n, as the following sketches:

- Let $\alpha, \beta \in F_n$ be such a pair and fix a spanning tree which contains as much as possible of α that means all up to one or two edges.
- Label the edges according to the marking with elements $\omega_j \in F_n$. The labels form a basis of the fundamental group. If α is a simple closed loop, it is one of the labels $\alpha = \omega_{i_1}$, otherwise α is a barbell or figure of eight and we have $\alpha = \omega_{i_1}\omega_{i_2}$ for two edge labels ω_{i_1} and ω_{i_2}.

- If β is a barbell we write it as a word

  $$\beta = (\omega_1 \dots \omega_k)(\omega_{k+1} \dots \omega_l)(\omega_{l+1} \dots \omega_m)(\omega_{k+1} \dots \omega_l)^{-1}$$

  in this basis. As the loops and the handle of β are disjoint all the ω_i are different edge labels and only the handle $(\omega_{k+1} \dots \omega_l)$ might be empty.

  As β is not the counterpart of α, there exists an $i \in \{1, \dots, m\}$ such that the letter ω_i is neither ω_{i_1} nor ω_{i_2}. Hence, we can exchange ω_i with β and ω_{i_1} with α and get again a basis of F_n.

- If β is not a barbell its word has either again a letter not contained in α or it is one of the two letters of α so we can again extend α and β to a basis.

Therefore, α and β must be as in the case (iii) and the claim follows. □

Looking at the proof of Proposition 2.16 (iii) we get that $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta(A)$ has already full dimension for all $\alpha \in \mathrm{cand}(A)$ near a point $A \in CV_n$. Comparing this to the proof of Lemma 3.6, we will see the following Lemma 3.7 (i), namely we can distinguish for two points A, B if B is in general position to A by the dimension of their envelope near A. As a more direct application of Lemma 3.6 we will also see that the property of general position of two points in a common maximal simplex is preserved along a straight line joining them.

Lemma 3.7
Let $A, B \in CV_n$ be in maximal simplices.

(i) B is in general position to A if and only if $\mathrm{Env}_R(A, B)$ has full dimension in $\Delta(A)$.

(ii) Let A, B be in the same maximal simplex Δ and in general position to each other. Furthermore, let $\gamma \colon [0, 1] \to \Delta$ be the straight line in this simplex joining A and B. Then for any $0 \leq s \leq t \leq 1$ we have that $\gamma(s)$ and $\gamma(t)$ are also in general position to each other.

(iii) Let A, B be as in (ii), then their envelope in regard to the symmetric metric has full dimension, i.e. we have $\dim(\mathrm{Env}(A, B)) \cap \Delta = 3n - 4$.

Proof. (i) Let B be in general position to A and $\alpha \in CW_R(A, B)$. By definition of general position there exists a neighbourhood U_A with $\alpha \in CW_R(A', B)$ for all $A' \in U_A$. In Proposition 2.16 (iii) we have seen that $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \cap \Delta(A)$ has full dimension and as $\mathrm{Env}_R^{\mathrm{out}}(A, \alpha)$ is a cone at A we have $U_A \cap \mathrm{Env}_R^{\mathrm{out}}(A, \alpha)$ has full dimension. By Corollary 1.24 we have $U_A \cap \mathrm{Env}_R^{\mathrm{out}}(A, \alpha) \subset \mathrm{Env}_R(A, B)$ which implies the claim.

For the converse statement observe that A is always a point in the relative interior of $\mathrm{Env}_R^{\mathrm{in}}(A, CW_R(A, B))$. Since we have $\mathrm{Env}_R(A, B) \subset \mathrm{Env}_R^{\mathrm{in}}(A, CW_R(A, B))$, $\mathrm{Env}_R^{\mathrm{in}}(A, CW_R(A, B))$ has full dimension in $\Delta(A)$ and thus A lies in $\mathrm{Env}_R^{\mathrm{in}}(A, CW_R(A, B))^{\mathrm{int}}$. In particular, we have $A \in \mathrm{Env}_R^{\mathrm{in}}(A, \alpha)^{\mathrm{int}}$ for any $\alpha \in CW_R(A, B)$, thus B is in general position to A by Lemma 3.6.

(ii) By Lemma 3.6 we have that A lies in the interior of in- and out-envelopes coming from B. These envelopes are cones in Δ and γ is a straight line in Δ, hence γ lies in the interior of the same in- and out-envelopes and in particular $\gamma(s)$ and B are in general position to each other. Similarly, we have that $\gamma(t)$ and $\gamma(s)$ lie in general position to each other.

(iii) Let γ be the straight segment as in (ii) and $C := \gamma(t)$ for any $0 < t < 1$. By (ii) C lies in general position to A and B and vice versa. Hence, there exists a neighbourhood U_C of C such that all $C' \in U_C$ have the same candidate witnesses as C from and to A and B, respectively. By Corollary 1.24 that means $U_C \subset \text{Env}(A, B)$, which concludes the proof.

□

We now use Lemma 3.7 and Lemma 3.1 to distinguish faces with envelopes in the symmetric metric. Since envelopes are preserved under isometries we get the following theorem:

Theorem 3.8
An isometry in regard to the symmetric Lipschitz metric of CV_n^{red} is simplicial.

Proof. Let $\varphi \in \text{Isom}(CV_n^{\text{red}})$ be an isometry with respect to the symmetric metric. We will first show that φ preserves maximal simplices. Therefore let $C \in CV_n^{\text{red}}$ be a point in a maximal simplex. Assume that $\varphi(C)$ lies on a face and let $U \subset \Delta(C)$ be a neighbourhood of C contained in the maximal simplex.

Let $\varphi(A), \varphi(B) \in \varphi(U)$ be as in Lemma 3.1. By the dense property of general position we assume that their preimages A and B are in general position to each other. Since φ is an isometry, it restricts to an isometry of $\text{Env}(A, B)$ to $\text{Env}(\varphi(A), \varphi(B))$. In particular, it preserves the dimension of envelopes close to its endpoints. But by Lemma 3.7 (iii) $\text{Env}(A, B)$ has near B full dimension while by Lemma 3.1 $\text{Env}(\varphi(A), \varphi(B))$ has smaller dimension near $\varphi(B)$.

We have thus proven that φ sends maximal simplices to maximal simplices and thus preserves the $(3n-5)$-skeleton of CV_n^{red}. If C belongs to the $(3n-6)$-skeleton of CV_n^{red}, then C has either two 4-valent vertices or one 5-valent vertex, which can be resolved to points in at least 4 different $(3n-5)$-dimensional simplices of CV_n^{red}. This means C belongs to a face of at least 4 simplices of dimension $3n-5$. So by topological reasons $\varphi(C)$ must belong to the $(3n-6)$-skeleton of CV_n^{red}. Inductively, φ sends each simplex to a simplex of the same dimension. □

We can now apply the proof of Stefano Francaviglias and Armando Martinos in [FM12b] for non-reduced Outer Space and get the same result for reduced Outer Space, namely:

Theorem 3.9
The isometry groups of CV_n^{red} with regard to the symmetric and the asymmetric Lipschitz metric are the same as in non-reduced case:

$$\text{Isom}(CV_n^{\text{red}}) = \text{Isom}(CV_n) = \begin{cases} \text{Out}(F_n), & \text{if } n \geq 3, \\ \text{PGL}(2, \mathbb{Z}), & \text{if } n = 2. \end{cases}$$

Remark 3.10
Another way to distinguish faces is directly by the property of general position. By Lemma 3.7 the property of general position of two points is preserved under isometries if both points are sent into maximal simplices. If $\gamma\colon [0,1] \to CV_n$ is a straight line in a maximal simplex, then by Lemma 3.7 (ii) $\gamma(1)$ is in general position to $\gamma(0)$ if and only if $\gamma(t)$ is in general position to $\gamma(s)$ for all $0 \le s < t \le 1$. On the other hand let A and B be as in Lemma 3.1 and $\gamma\colon [0,1] \to CV_n$ any geodesic from a A to B. Then there exists $0 < s < t < 1$ such that $\gamma(s)$ and $\gamma(t)$ are not in general position since as soon as γ passes the face, the envelope $\mathrm{Env}_R(\gamma(s), \gamma(t)) \cap \Delta(\gamma(s))$ lies in $\mathrm{Env}_R(A, B) \cap \Delta(B)$ which has not full dimension. This behaviour just relies on the fact that the coarse direction of γ is not a candidate of (the simplex of) B.

Having this in mind we can actually deduce similar behaviour for all geodesic rays, namely that if a geodesic runs long enough, it looses one dimension of freedom or in other words gains a small amount of rigidity:

Proposition 3.11
Let $\gamma\colon \mathbb{R}_{\ge 0} \to CV_n$ be an asymmetric geodesic ray parametrised by length, then there exists a $T > 0$ such that for all $T \le s < t$ we have $\dim\big(\mathrm{Env}_R(\gamma(s), \gamma(t))\big) \le \dim(CV_n) - 1$. In particular $\gamma(t)$ is not in general position to $\gamma(s)$.

Proof. Let $\alpha \in \bigcap_{t \in \mathbb{R}} CW_R(\gamma(0), \gamma(t))$ be a coarse direction of γ. Such an α exists since γ is a geodesic ray. By Lemma 1.22 we have $\alpha \in W(\gamma(s), \gamma(t))$ for all $0 \le s < t$. Since α is a witness for all elements of γ, its length is stretched exponentially to the length of γ. Hence, there exists a $T > 0$ with $l_{\gamma(s)}(\alpha) > 2\,\mathrm{vol}(\gamma(s))$ for all $s \ge T$. In particular, α is not a candidate for all $\gamma(s)$ with $s \ge T$ as at least one edge is covered by α more than three times.

Let now $\beta \in CW_R(\gamma(s), \gamma(t))$ for some $T \le s < t$. We will show that the equalities in $(\star)$ for β and α yield a non-trivial restriction for $\gamma(t)$. Assuming the contrary we would have $(l_{\gamma(s)}(\beta) \cdot \#(e_i, \alpha) - l_{\gamma(s)}(\alpha) \cdot \#(e_i, \beta)) = 0$ for all edges e_i. In other words in the abelianisation of F_n the element α would be the $(l_{\gamma(s)}(\alpha)/l_{\gamma(s)}(\beta))$-multiple of β. By construction we have$l_{\gamma(s)}(\alpha) > 2\,\mathrm{vol}(\gamma(s)) > l_{\gamma(s)}(\beta)$, hence α would be a proper multiple in the abelianisation. But α can be extended to a basis of F_n and hence its image in the abelianisation can not be a proper multiple of another element. Since the geodesic ray has to satisfy this non-trivial equality for $t > T$, we have $\dim(\mathrm{Env}_R^{\mathrm{out}}(\gamma(s))) < 3n - 4$ and by Lemma 3.7 $\gamma(t)$ is not in general position to $\gamma(s)$. □

Since a geodesic is rigid if and only if the envelope has everywhere dimension 1, we get the following corollary:

Corollary 3.12
Let $\gamma\colon \mathbb{R}_{\ge 0} \to CV_2$ be an asymmetric geodesic ray parametrised by length, then there exists a $T > 0$ such that $\gamma|_{\ge T}$ is a rigid geodesic ray.

Remark 3.13
Observe that Proposition 3.11 is sharp for every $n \ge 2$, namely there exist long geodesic rays $\gamma\colon \mathbb{R}_{\ge 0} \to CV_n$ with $\dim\big(\mathrm{Env}_R(\gamma(s), \gamma(t))\big) = \dim(CV_n) - 1$ for all $0 < s < t$ as we can see in the following example.

Example 3.14
Let $A_0 \in CV_2$ be the figure of eight with marking α and β and edge lengths a and $1-a$ for $a = \frac{\sqrt{5}-1}{2}$. If we choose γ to be the geodesic ray starting at A_0 with direction $\{\alpha\beta\}$,

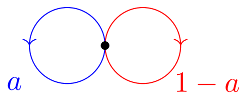

Figure 18: marked figure of eight

which is the same direction as $\{\alpha, \beta\}$, a short calculation shows that γ is the unique geodesic from A_0 to A_1, where A_1 is the figure of eight with marking $\alpha\beta^{-1}, \beta$ and again with edge lengths $1-a$ and a, respectively. In terms of envelopes the closed simplex $\overline{\Delta_0}$ containing A_0 and A_1 is covered by the envelopes of $\text{Env}_R^{\text{out}}(A_0, \alpha)$ and $\text{Env}_R^{\text{out}}(A_0, \beta)$ and γ is exactly the intersections of those two envelopes and hence rigid.

Furthermore, for the adjacent theta-simplex $\overline{\Delta_1}$ containing A_1 the out-envelopes $\text{Env}_R^{\text{out}}(A_1, \alpha)$ and $\text{Env}_R^{\text{out}}(A_1, \beta)$ lie in the interior of Δ_1, hence the envelope $\text{Env}_R^{\text{out}}(A_1, \{\alpha, \beta\}) \cap \Delta_1$ is one dimensional and by Proposition 2.16 (iv) it is the intersection of out-envelopes of candidates of A_1. But these look exactly like for A_0 and Δ_0, i.e. there are only two envelopes belonging to $\alpha\beta^{-1}$ and β and hence $\text{Env}_R^{\text{out}}(A_1, \{\alpha, \beta\}) \cap \Delta_1$ intersects the next face at A_2 with again edge lengths a and $1-a$.

Since we can continue γ along the out-envelopes with coarse direction $\{\alpha, \beta\}$ we inductively get that γ is a rigid geodesic ray with infinite length. Figure 19 is a picture of γ where the letters in the brackets denote the marking of the figures of eight.

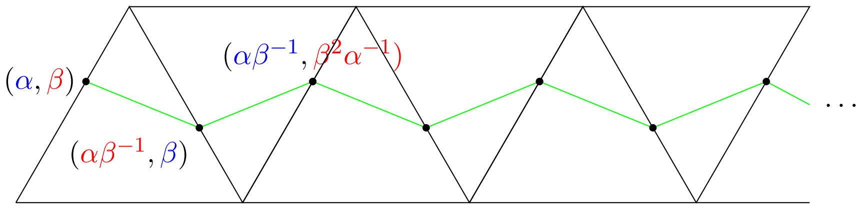

Figure 19: An infinite long rigid geodesic ray in CV_2

We consider the graphs $A_i' \in CV_n$ where we add to each A_i the same additional marked subgraph, for example a bouquet of roses with petal length 1 and the same marking:

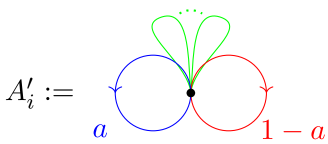

It is easy to check that the added green graph does not distribute anything to the distance between the A_j'. Hence, we can slightly vary the corresponding green edges, that is their lengths and their endpoints, without changing the witnesses. In particular we get that the corresponding ray γ' is an infinite geodesic ray with $\text{Env}_R(\gamma'(s), \gamma'(t)) = \dim(CV_n) - 1$ for all $0 \leq s < t$.

4 Local geodesics

In this section we will discuss how symmetric, locally minimising geodesics alias local geodesics look like in Outer Space. An *asymmetric, local geodesic* is here a path $\gamma\colon I \to CV_n$ such that for each $t \in I$ there exists an open subinterval $t \in J \subset I$ such that $\gamma|_J$ is a geodesic and a *local geodesic* is a path which is an asymmetric, local geodesic independent of orientation. In particular, we have that a path is an asymmetric, local geodesic if and only if it can be written as a concatenation of countably many geodesics $(\gamma_i)_{i\in\mathbb{Z}} : [0,1] \to CV_n$ such that for all $i \in \mathbb{Z}$ we have $\gamma_i(1) = \gamma_{i+1}(0)$ and there exists an $\varepsilon_i > 0$ with $W_R(\gamma_i(1-\varepsilon_i), \gamma_i(1)) \cap W_R(\gamma_{i+1}(0), \gamma_{i+1}(\varepsilon_i)) \neq \emptyset$ (see Corollary 1.24). With this in mind we can construct a locally geodesic loop:

Example 4.1
Consider the marked theta-graph $\Gamma :=$ 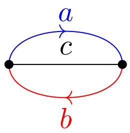

with marking α and β and the three points $A, B, C \in CV_2$ corresponding to Γ with edge lengths $(a, c, b) = (1,1,1), (2,1,1)$ and $(1, \frac{1}{3}, 1)$, respectively.

Observe that the candidates of Γ are α, β and $\alpha\beta$, hence for each ordered pair of A, B, C at least one of them is a witness. A short calculation shows $\{\alpha, \alpha\beta\} \subset W_R(A,B), \{\beta, \alpha\beta\} \subset W_R(B,C)$ and $\{\alpha, \beta\} \subset W_R(C,A)$. By Lemma 1.23 the concatenation of the straight lines $\overrightarrow{AB}, \overrightarrow{BC}$ and $\overrightarrow{CA}$ yields a closed, asymmetric, local geodesic as any two consecutive lines share a witness.

Similarly we can construct a local geodesic for the symmetric metric in the shape of a hexagon. For example the hexagon in $\Delta(\Gamma)$ where the vertices have edge lengths $(1,2,2), (1,1,2), (2,1,2), (2,1,1), (2,2,1)$ and $(1,2,1)$ is a local geodesic (see Figure 20) as each two consecutive lines (in either direction) share at least one candidate witness.

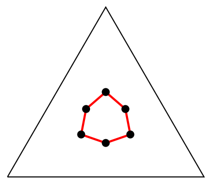

Figure 20: A local geodesic

We can make such changes of directions of a local geodesic arbitrarily small. That means for any two points we can give a local geodesic between them which starts and ends with a given direction. We will first give the statement only for points in the same simplex and in general position and construct from this the general case.

Lemma 4.2
Let $\Delta \subset CV_n$ be a maximal simplex, $A, B \in \Delta$ two points in general position to each other and γ_A, γ_B be two local geodesic respectively ending at A and starting at B, i.e.

we have for $a, b \in \mathbb{R}_{>0} \cup \{\infty\}$ that $\gamma_A \colon [-a, 0] \to CV_n$, $\gamma_B : [1, 1+b] \to CV_n$ are local geodesics with $\gamma_A(0) = A$ and $\gamma_B(1) = B$.

Then there exists a path $\sigma \colon [0,1] \to CV_n$ such that the concatenation $\gamma_A * \sigma * \gamma_B$ is a local geodesic. Moreover, σ may approximate the distance between A and B arbitrarily well, that is for every ε we can find such a σ that $l(\sigma) < d(A, B) + \varepsilon$.

Proof. For a point C in Δ consider the fan $F(C)$ in Δ originating at C and generated by the in- and out-envelopes of C, that is each cone in $F(C)$ can be written as an intersection of an in-envelope and an out-envelope of C. As in Proposition 2.16 and Lemma 3.7 each full dimensional cone in $F(C)$ corresponds to a pair of candidates of C and we write $C(\alpha, \beta) := \mathrm{Env}_R^{\mathrm{out}}(C, \alpha) \cap \mathrm{Env}_R^{\mathrm{in}}(C, \beta)$ for such a cone. Since the in- and out-envelopes in Δ depend continuously on the point C there exists a neighbourhood $U_A \subseteq \Delta$ such that for all points $A' \in U_A$ their fans $F(A)$ and $F(A')$ have the same structure. This means two $(3n-4)$ dimensional cones in $F(A)$ have a $(3n-5)$-dimensional intersection if and only if their corresponding cones in $F(A')$ have a $(3n-5)$-dimensional intersection.

As A and B are in general position, we can choose U_A small enough that A' and B are in general position for all $A' \in U_A$. Let now $\varepsilon > 0$ be small enough such that $\gamma_A(-\varepsilon) \in \Delta$ and $\gamma_A|_{[-\varepsilon, 0]}$ is a geodesic. Let $(\beta, \alpha) \in \mathrm{cand}(A)^2$ correspond to a full dimensional cone of $F(A)$ containing $\gamma_A(-\varepsilon)$. A and B are in general position, thus we have by Lemma 3.6 that B lies in the interior of a full dimensional cone $A(\omega, \mu)$ for $(\omega, \mu) \in CW(A, B)$. In $F(A)$ we fix a sequence of full dimensional cones $A(\alpha_0, \beta_0), \dots, A(\alpha_N, \beta_N)$, which represents a path from the cone $A(\alpha, \beta)$ to the cone $A(\omega, \mu)$, that is for all i the cones $A(\alpha_i, \beta_i)$ and $A(\alpha_{i+1}, \beta_{i+1})$ share a facet and $(\alpha_0, \beta_0) = (\alpha, \beta), (\alpha_N, \beta_N) = (\omega, \mu)$. As we have chosen U_A small enough this path of cones is the same for any $A' \in U_A$. Without loss of generality let $U_A = B_\delta(A)$ be the δ-ball around A with respect to the symmetric distance for some $\delta > 0$. We define inductively $A_0 := A$ and $A_i \in A_{i-1}(\alpha_{i-1}, \beta_{i-1}) \cap A_{i-1}(\alpha_i, \beta_i) \cap B_{\delta/N}(A_{i-1})$ for $1 \le i \le N$. In particular we have $A_i \in U_A$ for all i and thus the intersection of the cones is not empty. By construction we have furthermore $(\alpha_i, \beta_i) \in CW(A_{i-1}, A_i) \cap CW(A_i, A_{i+1})$.

Let $\sigma_{A,i}$ be a geodesic from A_{i-1} to A_i, for example the straight line between them. Since they share a witness, their concatenation $\sigma_{A,1} * \cdots * \sigma_{A,N}$ yields a local geodesic by Corollary 1.24. As $(\alpha, \beta) \in CW(A, A_1) \cap CW(\gamma_A(-\varepsilon), A)$, the concatenation $\gamma_A * \sigma_{A,1} * \cdots * \sigma_{A,N}$ is still a local geodesic. Furthermore, A_N and B are still in general position and we have $(\omega, \mu) \in CW(A_N, B) = CW(A, B)$ and by construction $(\omega, \mu) \in CW(A_{N-1}, A_N)$.

Similarly, we construct points $B_1, \dots, B_M \in \Delta$ and obtain a local geodesic $\sigma_{B,M} * \cdots * \sigma_{B,1} * \gamma_B$ such that $(\mu, \omega) \in CW(B_{M-1}, B_M)$. Let $\sigma_{A,B}$ be a symmetric geodesic from A_N to B_M, e.g. a straight line in Δ. As A and B are in general position and we chose the neighbourhoods small enough, we still have $(\omega, \mu) \in CW(A_N, B_M)$. By construction we also have $(\omega, \mu) \in CW(A_{N-1}, A_N) \cap CW(B_M, B_{M-1})$ and hence the concatenation $\gamma_A * \sigma_{A,1} * \cdots * \sigma_{A,N} * \sigma_{A,B} * \sigma_{B,M} * \cdots * \sigma_{B,1} * \gamma_B$ is a local geodesic. In particular we have that $\sigma \colon = \sigma_{A,1} * \cdots * \sigma_{A,N} * \sigma_{A,B} * \sigma_{B,M} * \cdots * \sigma_{B,1}$ is the desired path.

Observe that the path σ has length $l(\sigma) = l(\sigma_{A,1}) + \cdots + l(\sigma_{A,B}) + \cdots + l(\sigma_{B,1})$. Each of the segments $\sigma_{A,i}$ has at most length δ/N and A_N, B_M lie in the δ-neighbourhoods of A and B, that is we have by the triangle inequality $l(\sigma_{A,B}) = d(A_N, B_M) \le 2\delta + d(A, B)$.

In particular we have then $l(\sigma) = \delta + 2\delta + d(A,B) + \delta = 4\delta + d(A,B)$, that is by choosing δ small enough σ approximates the distance arbitrarily close. □

While we have seen in Lemma 1.25 that not all points in different simplices are joined by a symmetric geodesic, we will now see that we can join any two adjacent simplices of CV_n by a symmetric geodesic:

Lemma 4.3
Let $\Delta_1, \Delta_2 \subset CV_n$ be two adjacent simplices and $C \in F := \overline{\Delta_1} \cap \overline{\Delta_2}$ be a point on their common face F. Then there exist points $A \in \Delta_1$, $B \in \Delta_2$ such that C lies on a symmetric geodesic from A to B.

Proof. The idea of the proof is to go “skew” over the face F that means we have two witnesses α and β which are so much stretched that the collapsed edges of Δ_1 and Δ_2 do not matter.

Similar to the proof of Lemma 3.1 there exist two edge-disjoint simple loops $\alpha, \beta \in F_n$ in C. Assume that C is normalised an let $l := \min_{e \in E(C)} l_C(e)$ be the minimal length of an edge in C.

Recall that the underlying marked graph of F is obtained by collapsing a forest in the marked graph of Δ_1. Hence, we can define the point $A \in \Delta_1$ with the following edge lengths

$$l_A(e) = \begin{cases} 2 \cdot l_C(e), & \text{if } e \text{ is contained in } \alpha, \\ \frac{1}{2} \cdot l_C(e), & \text{if } e \text{ is contained in } \beta, \\ \delta, & \text{if } e \text{ is collapsed from } \Delta_1 \text{ to } F, \\ l_C(e), & \text{otherwise} \end{cases},$$

for some small enough $0 < \delta < \frac{l}{2(2n-3)}, \frac{l_C(\beta)2l}{4(2n-3)(2+l)}$. We will show that α is maximally stretched from C to A and β is maximally stretched from A to C. Observe that the collapsed forest contains at most $2n-3$ edges, hence we have for the stretching of α and β in A the following inequalities:

$$\frac{l_A(\alpha)}{l_C(\alpha)} \geq \frac{2l_C(\alpha)}{l_C(\alpha)} = 2$$

$$\begin{aligned} \frac{l_C(\beta)}{l_A(\beta)} &\geq \frac{l_C(\beta)}{\frac{1}{2}l_C(\beta) + (2n-3)\delta} = 2 - \frac{4(2n-3)\delta}{l_C(\beta) + 2(2n-3)\delta} \\ &> 2 - \frac{4(2n-3)\delta}{l_C(\beta)} \qquad > 2 - \frac{2l}{2+l}. \end{aligned}$$

Here we used in the last inequality $\delta < \frac{l_C(\beta)2l}{4(2n-3)(2+l)}$.

Let $\omega \in (\mathrm{cand}(A) \cup \mathrm{cand}(C)) \setminus \{\alpha, \beta\}$ be another candidate. Since α and β are simple loops in C, there exist edges $e_\alpha, e_\beta \in E(C)$ which are covered by ω but e_α is not covered by α and e_β is not covered by β. So we have the inequalities:

$$
\begin{aligned}
& l_A(\omega) \leq 2l_C(\omega) - l_C(e_\alpha) + 2 \cdot (2n-3)\delta < 2l_C(\omega) - l + l = 2l_C(\omega) \\
\Rightarrow \quad & \frac{l_A(\omega)}{l_C(\omega)} < 2 \leq \frac{l_A(\alpha)}{l_C(\alpha)} \\
\\
& l_A(\omega) \geq \frac{1}{2}l_C(\omega) + \frac{1}{2}l_C(e_\beta) \geq \frac{1}{2}l_C(\omega) + \frac{1}{2}l \\
\Rightarrow \quad & \frac{l_C(\omega)}{l_A(\omega)} \leq \frac{l_C(\omega)}{\frac{1}{2}l_C(\omega) + \frac{1}{2}l} \leq 2 - \frac{2l}{l_C(\omega) + l} \leq 2 - \frac{2l}{2+l} < \frac{l_C(\beta)}{l_A(\beta)}
\end{aligned}
$$

We used here $\delta < \frac{l}{2(2n-3)}$ in the first line.

We construct the point $B \in \Delta_2$ in the same manner with the roles of α and β exchanged and get by Corollary 1.24 that C lies on a symmetric geodesic from A to B. □

Using these two lemmas, we can construct a local geodesic through any given sequence of points in CV_n.

Theorem 4.4
Let $(A_i)_{i \in \mathbb{Z}} \subset CV_n$ be a sequence of points in CV_n. Then there exists a local geodesic $\gamma\colon \mathbb{R} \to CV_n$ passing through these points, i.e. we have $\gamma(i) = A_i$ for all $i \in \mathbb{Z}$.

Proof. By Proposition 3.5 being in general position is an open and dense property. Thus, after possibly introducing some additional intermediate points, we can assume that the two consecutive points A_i and A_{i+1} are either as in Lemma 4.2, that is in the same simplex and in general position to each other, or belong to a triplet as in Lemma 4.3. Hence, we can construct inductively a local geodesic passing through these points. □

As each simplex contains a countable dense set and there are countably many simplices in CV_n we get that:

Corollary 4.5
There exists a local geodesic which is dense in CV_n.

Similarly to above we can approximate any given path by a local geodesic:

Proposition 4.6
Let $\gamma\colon I \to CV_n$ be a continuous path and $\varepsilon > 0$. Then there exists a local geodesic $\sigma\colon J \to CV_n$ such that γ lies in the ε-neighbourhood of σ in regard to the symmetric metric and vice versa.

Proof. As points in general position are dense, we can find a sequence $(A_i)_{i\in\mathbb{Z}}$ in the ε-neighbourhood of γ such that $d(A_i, A_{i+1}) < \varepsilon$ and the points A_i lie either as in Lemma 4.2 or Lemma 4.3 and we have $\gamma \subset \bigcup_{i\in\mathbb{Z}} B_\varepsilon(A_i)$. By Lemma 4.2 we can assume that the constructed local geodesic $\sigma|_{A_i,A_{i+1}}$ has at most length $l(\sigma|_{A_i,A_{i+1}}) \leq d(A_i, A_{i+1}) + \varepsilon$ and hence is contained in a 2ε-neighbourhood of A_i. In particular we have that the constructed local geodesic is contained in the 3ε-neighbourhood of γ and γ is contained in the ε-neighbourhood of σ. □

5 Isometric embeddings

While we know by [FM12b] and by Section 3 that the isometry group of Outer Space and reduced Outer Space is (virtually) $\mathrm{Out}(F_n)$, there is so far nothing known about isometric embeddings from CV_n to CV_k for $k > n$. We will discuss and explicitly construct two different types of such embeddings in this section. Each of these types of embeddings correspond to a different way of identifying one free group as a subgroup of another. The first type introduced in Sections 5.1 and 5.2 is the "naive way" to increase the rank of the fundamental group, which is attaching a rose to a marked graph and corresponds to F_n being a free factor of F_k. The second type is the more natural embedding from Section 5.3, which derives from finite coverings and corresponds to identifying F_k with a finite index subgroup of F_n. While the naive embedding leaves a lot of freedom where and which graph we attach, the natural embedding is restricted by the Nielsen-Schreier formula, that is we find such an embedding only for $k = 1 + d(n-1)$ for some $d \in \mathbb{N}$.

One of the questions we will answer is: Are embeddings from CV_n to CV_k in some sense discrete? For comparison the isometry group acts properly discontinuously on CV_n as each simplex has a finite stabiliser. Furthermore, the automorphism group of F_n is finitely presented by [Nie24] and hence countable. In contrast the naive embeddings can be continuously deformed as we will see in Example 5.6 and we will see in Section 6 that we can also continuously and locally deform a natural embedding from CV_2 to CV_k. Hence, we have families of arbitrarily close isometric embeddings. However, we will also see that one can not locally deform a natural embedding from CV_n to CV_k for $n \geq 3$, so in this sense these natural embeddings are more rigid.

Before we start with isometric embeddings from CV_n to CV_k recall that we have seen in Section 3 that the simplicial structure of CV_n is not only determined topologically but in reduced Outer Space also by the Lipschitz metric, as we can locally distinguish faces with envelopes (see Theorem 3.8). In contrast we will see that there is no local reason that an isometric embedding should be simplicial. For instance we can isometrically embed three adjacent simplices of CV_2 into a single rose-simplex of CV_k for any large enough k, as we will see in Example 5.1.

The idea of this example is that the distance of two points is determined by the stretching of their candidates. Since the distance inside a rose-simplex is determined by the stretching of its petals (see Corollary 1.15), we will give each petal the length of a candidate. To avoid different volumes of the representants in the image we will add another petal which normalises the volume. Here we have to be careful that this "normalising petal" is never a witness in the image. One way to avoid that the "normalising petal", is a witness is to add a large constant to the length, i.e. give it length $K - \sum_\alpha l(\alpha)$ for some large enough $K \gg 0$, where α ranges over the set of all candidates in the preimage.

Example 5.1
Let $\Delta_1, \Delta_2, \Delta_3, \Delta_4 \subset CV_2$ be the simplices corresponding to following the topological types with marking α and β (see Figure 21).

Let $U = \bigcup_{1 \leq i \leq 4} \Delta_i \subset CV_2$ be their union. Observe that Δ_4 is the shared face of

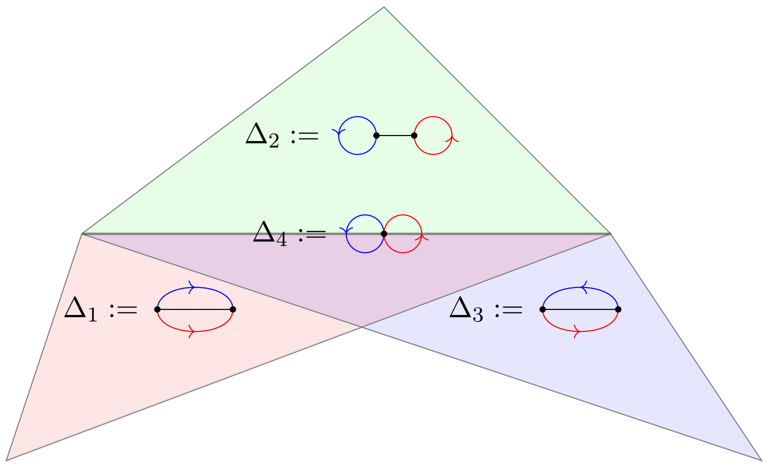

Figure 21: Three simplices in CV_2 sharing a common face.

Δ_1, Δ_2 and Δ_3. We denote for any $A \in U$ with $R_5(A) := R_5(l_1(A), \dots, l_5(A))$ the rose with petal-lengths

$$\begin{aligned} l_1(A) &:= l_A(\alpha), & l_2(A) &:= l_A(\beta), \\ l_3(A) &:= l_A(\alpha\beta), & l_4(A) &:= l_A(\alpha\beta^{-1}), \\ l_5(A) &:= 4\,\mathrm{vol}(A) - \big(l_1(A) + l_2(A) + l_3(A) + l_4(A)\big) \end{aligned}$$

and standard marking. We will show that the map

$$\iota : U \to \Delta(R_5) \subset CV_5 \quad , \quad A \mapsto R_5(A)$$

is isometric in regard of the Lipschitz metric. In particular when we pass the common face of Δ_1, Δ_2 and Δ_3 in U, we do not pass a face in its image.

Proof. For any element $A \in U$ we have for the lengths of the candidates $l_i(A) \leq 2\,\mathrm{vol}(A)$ for $1 \leq i \leq 4$. Hence, $R_5(A)$ has strict positive lengths, which scale linearly with $\mathrm{vol}(A)$. Thus, ι is a well defined map. Since $\alpha, \beta, \alpha\beta, \alpha\beta^{-1}$ are exactly all the candidates occurring in U, we now want to show that their stretching factors determine the distance in the image of ι.

By Corollary 1.15 we have that the maximal stretching of two roses of the same topological type is the maximal stretching of their petals. That means we have

$$\Lambda_R(R_5(l_1, \dots, l_5), R_5(l'_1, \dots, l'_5)) = \max\left\{\frac{l'_1}{l_1}, \dots, \frac{l'_5}{l_5}\right\}.$$

These fractions are exactly the stretching factors of the candidates in U and of l_5. So we only need to show that l_5 is never maximally stretched in the image of ι. Now let $A, B \in U$

be two normalised representants with edge lengths a, b, c and a', b', c', respectively. Then we have as the edge lengths in $R(A)$

$$\begin{aligned} \text{for } A \in \Delta_1 \cup \Delta_3: \quad l_5(A) &= 4\,\mathrm{vol}(A) - (\ l_A(\alpha) + l_A(\beta) + l_A(\alpha\beta) + l_A(\alpha\beta^{-1})\ ) \\ &= 4 - (\ (a+c) + (b+c) + (a+2c+b) + (a+b)\ ) \\ &= 4 - (3(a+b+c)+c) = 1 - c = a+b \\ \text{for } A \in \Delta_2: \quad l_5(A) &= 4 - (a+b+2(1+c)) = 1-c = a+b \end{aligned}$$

We used here the normalisation $a+b+c=1$ and consider Δ_4 as a special case of any of the above with $c=0$.

For any topological type of B we also have $l_B(\alpha\beta) \geq a'+b'$, hence we have for $A \in \Delta_1$ the inequality $\frac{l_4(B)}{l_4(A)} \geq \frac{a'+b'}{a+b} = \frac{l_5(B)}{l_5(A)}$ and for $A \in \Delta_3$ the inequality $\frac{l_3(B)}{l_3(A)} \geq \frac{a'+b'}{a+b} = \frac{l_5(B)}{l_5(A)}$. For $A \in \Delta_2$ we have by Lemma 1.14 the inequality $\frac{l_5(B)}{l_5(A)} = \frac{a'+b'}{a+b} \leq \max\{\frac{a'}{a}, \frac{b'}{b}\} \leq \max\{\frac{l_1(B)}{l_1(A)}, \frac{l_2(B)}{l_2(A)}\}$.

As the maximal stretchings of the corresponding candidates in A and $R_5(A)$ are the same, we have that $\Lambda_R(R_5(A), R_5(B)) = \Lambda_R(A, B)$. Since $\mathrm{vol}(R_5(A)) = 4\,\mathrm{vol}(A)$, this already implies $d_R(R_5(A), R_5(B)) = d_R(A, B)$, i.e. ι is an isometric embedding.

□

While we can construct such examples for any given finite number of simplices, we can not construct a global isometric embedding in this manner. The reason for this is that we have infinitely many candidates in CV_n and by Proposition 2.16 (iii) each of them appears as the only witness between two points. Furthermore, we can never embed CV_n isometrically into a simplex for the following reason: If a point $A \in CV_k$ is ε-thick (see Definition 1.7), then any point $B \in CV_k$ with distance $d_R(A, B) > \log(\frac{2}{\varepsilon})$ has a different topological type than A as the candidate witness from A to B has length bigger than $2\,\mathrm{vol}(B)$ and thus can not be a witness in B.

Nevertheless the above example implies that there is no local reason to expect isometric embeddings to be simplicial. We will also see in Corollary 5.15 how to construct global examples where the image of a simplex intersects multiple simplices.

5.1 Naive embedding for CV_2

The first naive idea to embed CV_n into CV_k is to attach a bouquet of $n-k$ petals for a given representative in CV_n as in Figure 22:

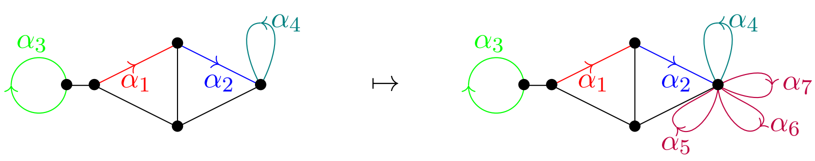

Figure 22: Attaching a bouquet of petals to a marked graph

This naive approach has the following obstructions, which we need to take care of:

- The map depends on the choice of the representant of a point in CV_n. Since the attached rose prevents free homotopy, two different representants of the same point in CV_n may have different images in CV_k if we attach the same bouquet of petals.
- For each representant there is a choice of the point where we attach the rose.
- It introduces new candidates which might be more stretched, so this might not be an isometric embedding.

The first problem comes from the fact that attaching a rose prevents some of the free homotopy which would occur otherwise, namely each free homotopy would move around the attaching point. For example conjugating with an element in F_n would also conjugate the attached petal.

The last problem can be partially solved by choosing the attached petals large enough similar to Example 5.1.

We will visualise these problems in the following examples. In Example 5.2 (i) we will see an isometric, local embedding, which we will use later for a global isometric embedding. In comparison we will choose in Example 5.2 (ii) a different representant and will see that it can not be extended to an isometric embedding. Example 5.2 (iii) shows that we might increase the distance by introducing a too small petal.

Example 5.2 (i) Let $\Delta_1, \Delta_2 \subset CV_2$ be the two adjacent (open) simplices corresponding to a marked theta-graph. That is as before there exist $\alpha, \beta \in F_2$ such that Δ_1 and Δ_2 are all the marked graphs with a representant of the form

$\Delta_1 :=$ 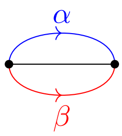

and $\Delta_2 :=$ 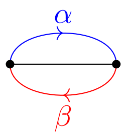

.

Consider the map $\phi_1 : \Delta_1 \cup \Delta_2 \to CV_3$, which attaches to each normalised representant a loop with marking $\omega \in F_3$ of fixed length $l(\omega) > 0$, that is:

$\phi_1 : \Delta_1 \cup \Delta_2 \to CV_3,$
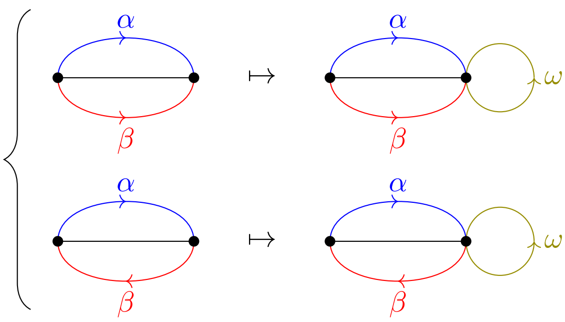

We will show that ϕ_1 preserves the distance. As ϕ_1 adds a loop of constant length to normalised representants, it is enough to consider the stretching factor Λ_R for the distance. Observe that each candidate in the image $\phi_1(\Delta_1 \cup \Delta_2)$ is of the form ω, μ or $\mu\omega^{\pm 1}$, where $\mu \in \{\alpha, \beta, \beta\alpha, \beta^{-1}\alpha\}$ is a loop, i.e. one of the candidates in the preimage. For the candidates of the form $\mu\omega^{\pm 1}$ we have $l(\mu\omega) = l(\mu) + l(\omega)$. Since $l(\omega)$ is fixed in each image, we have by Lemma 1.14 that for two points $\phi_1(A), \phi_1(B) \in \phi_1(\Delta_1 \cup \Delta_2)$ all the $\mu\omega$ are either less stretched than $l(\mu)$ or not stretched at all from $\phi_1(A)$ to $\phi_1(B)$. Thus, the stretching factor from $\phi_1(A)$ to $\phi_1(B)$ is the same as the stretching for A to B. By allowing edges with length zero we can extend ϕ_1 to an isometric embedding of the closure of the two simplices and in particular to their shared face

$\Delta_3 := \overline{\Delta_1} \cap \overline{\Delta_2} =$ 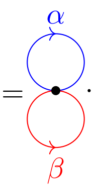

.

(ii) Let $\Delta_1, \Delta_2 \subset CV_2$ be the same simplices as in Example (i). Observe that we can choose a different representative for the graphs in Δ_2, namely we have

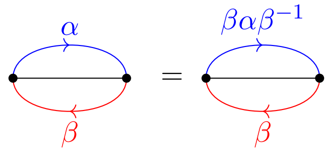

$=: \Delta_2.$

As before we define a map:

$\phi_2 : \Delta_1 \cup \Delta_2 \to CV_3,$ 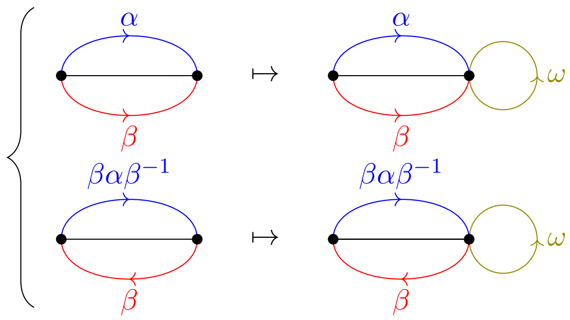

In contrast to before, ϕ_2 can not be extended continuously to $\overline{\Delta_1} \cup \overline{\Delta_2}$. Any such extension would have to send the common face Δ_3 of Example (i) to two different faces: to the rose with marked petals α, β, ω and the rose with marked petals $\beta\alpha\beta^{-1}, \beta, \omega$. The problem is that the free homotopy which sends the marking α to $\beta\alpha\beta^{-1}$ would also conjugate ω and so these faces are indeed different. In particular we have that ϕ_2 is not an isometric map, since points $A \in \Delta_1$ and $B \in \Delta_2$ may be arbitrarily close near Δ_3 but their images are distant to each other.

(iii) Let $\Delta \subset CV_2$ be the simplex in CV_2 corresponding to a barbell graph and let (a, c, b) be the edge lengths for some normalised point $A \in \Delta$. For some $\lambda \in [0, 1]$ let $c_1 := \lambda c$ and $c_2 := (1 - \lambda)c$. We define the map

$\psi_\lambda : \Delta \to CV_3$,

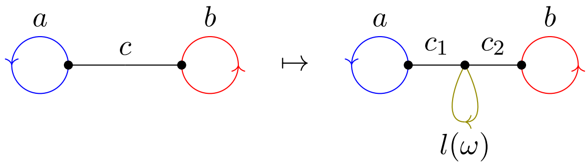

that attaches to each (normalised) point in Δ a petal of some fixed length $l(\omega)$ at the separating edge. We will show that ψ_λ is an isometric embedding if and only if $l(\omega) \geq \max\{2\lambda, 2(1 - \lambda)\}$ holds. In particular, if $l(\omega) \geq 2$, then ψ_λ is an isometric embedding for every λ. As $l(\omega)$ is constant, it is again enough to consider the stretching factors.

For the proof let α, β and ω denote the marking of the petals as depicted in the above figure and assume by symmetry that $\lambda \geq \frac{1}{2}$.

First we show that for an $l(\omega) \geq 2\lambda$ the map ψ_λ is an isometric embedding. Let $\psi_\lambda(A), \psi_\lambda(B =\in \psi_\lambda(\Delta)$ be two points in the image. We show that from A to B there is at least one of the candidates $\alpha, \beta, \alpha\beta$ as much stretched as one of the new candidates $\omega, \alpha\omega$ and $\beta\omega$ in $\psi_\lambda(A)$. As α, β and $\alpha\beta$ are also candidates in the preimage with the same stretching factor we can then again conclude that ψ_λ is an isometric embedding. Let $A, B \in \Delta$ be the points with lengths a, c, b, respectively a', c', b'. As ω has constant length $l(\omega)$, it can never be a witness. Assume $\alpha\omega$ is a witness from $\psi_\lambda(A)$ to $\psi_\lambda(B)$ and α is not maximally stretched. Then we have the following inequalities:

$$\begin{aligned}
& \frac{l_{\psi_\lambda(B)}(\alpha)}{l_{\psi_\lambda(A)}(\alpha)} < \frac{l_{\psi_\lambda(B)}(\alpha\omega)}{l_{\psi_\lambda(A)}(\alpha\omega)} \\
\iff \quad & \frac{a'}{a} < \frac{a' + 2c_1' + l(\omega)}{a + 2c_1 + l(\omega)} \\
& \overset{1.14}{\leq} \max\left\{\frac{a'}{a}, \frac{2\lambda c' + l(\omega)}{2\lambda c + l(\omega)}\right\} = \frac{2\lambda c' + l(\omega)}{2\lambda c + l(\omega)} \\
& = \frac{2\lambda c' + 2\lambda + (l(\omega) - 2\lambda)}{2\lambda c + 2\lambda + (l(\omega) - 2\lambda)} \\
& \overset{1.14}{\leq} \frac{2\lambda(c' + 1)}{2\lambda(c + 1)} = \frac{l_{\psi_\lambda(B)}(\alpha\beta)}{l_{\psi_\lambda(A)}(\alpha\beta)}
\end{aligned}$$

where we used in the last inequality that the maximal stretching and thus the fraction $\frac{l_{\psi_\lambda(B)}(\alpha\omega)}{l_{\psi_\lambda(A)}(\alpha\omega)}$ has to be at least 1 and $l(\omega) - 2\lambda \geq 0$. Hence, $\alpha\beta$ is also a witness from $\psi_\lambda(A)$ to $\psi_\lambda(B)$ and thus we have $\Lambda_R(\psi_\lambda(A), \psi_\lambda(B)) = \Lambda_R(A, B)$. A similar calculation holds if $\beta\omega$ is a witness and β is not a witness. Since the stretching factors are the same as in the preimage, we have that ψ_λ is indeed an isometry.

We now give an example that if we have $l(\omega) < 2\lambda$, then ψ_λ is not an isometry. For some small $\varepsilon > 0$ let $A, B \in \Delta$ be the two (normalised) points with lengths $a = c = b = \frac{1}{3}$ and $a' = \frac{1}{3} + \varepsilon, c' = \frac{1}{3} + 4\varepsilon, b' = \frac{1}{3} - 5\varepsilon$. Then we have as stretching factors for α, β, $\alpha\beta$ and $\alpha\omega$:

$$\begin{aligned}
\frac{l_B(\alpha)}{l_A(\alpha)} &= \frac{1/3 + \varepsilon}{1/3} = 1 + 3\varepsilon \\
\frac{l_B(\beta)}{l_A(\beta)} &= \frac{1/3 - 5\varepsilon}{1/3} < 1 \\
\frac{l_B(\alpha\beta)}{l_A(\alpha\beta)} &= \frac{1 + \frac{1}{3} + 4\varepsilon}{1 + \frac{1}{3}} = 1 + 3\varepsilon \\
\frac{l_{\psi_\lambda(B)}(\alpha\omega)}{l_{\psi_\lambda(A)}(\alpha\omega)} &= \frac{\frac{1}{3} + \varepsilon + \frac{2}{3}\lambda + 8\varepsilon\lambda + l(\omega)}{\frac{1}{3} + \frac{2}{3}\lambda + l(\omega)} \\
&= 1 + \frac{\varepsilon + 8\varepsilon\lambda}{\frac{1}{3} + \frac{2}{3}\lambda + l(\omega)} \\
&> 1 + \frac{\varepsilon + 8\varepsilon\lambda}{\frac{1}{3} + \frac{2}{3}\lambda + 2\lambda} \\
&= 1 + \frac{3\varepsilon(1 + 8\lambda)}{1 + 8\lambda} = 1 + 3\varepsilon
\end{aligned}$$

In particular we have that $\Lambda_R(\psi_\lambda(A), \psi_\lambda(B) > \Lambda_R(A, B)$, which shows that ψ_λ is not an isometric embedding for $l(\omega) < 2\lambda$.

To avoid the first two obstructions of the naive embedding mentioned at the beginning of this section we need to make a coherent choice of representants and attaching points. For the case $n = 2$ this can be done inductively as we will see in the following section. To make such a choice for any $n \geq 2$ we will consider Outer Space as in Definition 1.3 as action on trees and choose the base point on the tree instead. We will discuss two such constructions of a base point in the tree in Section 5.2.

Recall that by the third obstruction we may increase the distance of two points by attaching a petal. Observe that we solved this in Example 5.1 and Example 5.2 (iii) by attaching large enough petals. This will also be the solution for the upcoming embeddings, namely we need to take care that the attached graph is thick enough (see Lemma 5.8). It will turn out that we can even slightly deform the attached graph along the preimage without changing the distance in the image of an embedding.

Example 5.3
Starting with the isometric embeddings ϕ_1 and ψ_λ from Example 5.2 we will construct inductively a global embedding from CV_2 into CV_3 by extending the constructed embedding step by step to adjacent simplices.

The reason we can do this construction is the following: Given a non-empty connected simplicial subcomplex $\Sigma \subseteq CV_2$ and $\Delta' \subset CV_2$ a simplex adjacent to Σ, then Σ and Δ' share exactly one face of dimension 1 since otherwise we would have a closed loop around the missing vertex of two common faces of Σ and Δ', which contradicts that CV_2 is contractible. So we extend a given embedding of a subcomplex one simplex after another without worrying about the order of the extended simplices. We start to construct such an embedding from CV_2^{red} to CV_3 and extend it afterwards to CV_2.

Let $\Sigma \subseteq CV_2^{\text{red}}$ be a connected simplicial subcomplex and $\phi\colon \Sigma \to CV_3$ a locally isometric embedding which looks on each simplex like ϕ_1, that is for a maximal simplex $\Delta \subset \Sigma$ we have $\alpha, \beta \in F_2$ such that for the given representants of points in Δ we have

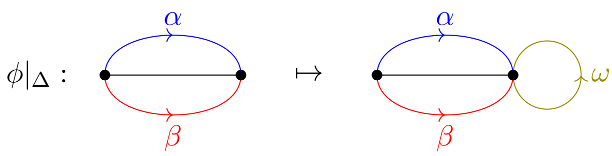

Any adjacent simplex $\Delta' \subseteq CV_2^{\text{red}}$ shares a face with a unique simplex $\Delta \subset \Sigma$. Recall that changing the spanning tree just corresponds to renaming α and β:

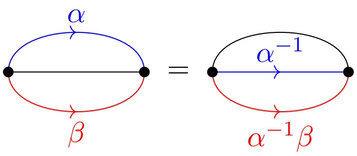

and

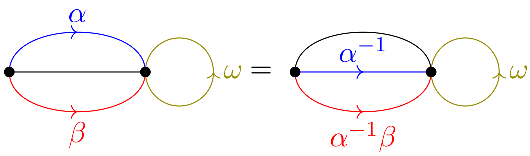

Hence, we can up to relabelling assume that Δ and Δ' look like in Example 5.2. That means we can extend ϕ like ϕ_1 to a locally isometric embedding $\phi\colon \Sigma \cup \overline{\Delta'} \to CV_3$ such that $\phi|_{\overline{\Delta \cup \Delta'}}$ looks like ϕ_1. Inductively we get a global map $\phi\colon CV_2^{\text{red}} \to CV_3$, which is locally isometric.

To get an embedding for non-reduced Outer Space CV_2 observe that ϕ_1 and ψ_λ are equal on their common face i.e. the simplex corresponding the rose R_2. Hence, we can extend ϕ to the whole Outer Space CV_2 via ψ_λ for some $\lambda \in [0,1]$ on the barbell graphs.

To justify that ϕ from Example 5.3 is indeed an isometric embedding, observe that ϕ at most increases the distance. As we just attach a loop to a given graph $A \in CV_2$ we can identify any original loop $\alpha \in F_2$ of A with a loop $\tilde{\alpha}$ in $\phi(A)$. In particular we have for a loop $\alpha \in F_2$ and a normalised representant $A \in CV_2$ that the length $l_A(\alpha) = l_{\phi(A)}(\alpha)$ is preserved. By the following Lemma we have then that ϕ is indeed an isometric embedding.

Lemma 5.4
Let $\phi\colon CV_n \to CV_k$ be a map not decreasing the distance, that means for all $A, B \in CV_k$ we have $d_R(\phi(A), \phi(B)) \geq d_R(A, B)$. Then ϕ is an isometric embedding if and only if it is restricted to each closed simplex $\overline{\Delta} \subset CV_k$ an isometric embedding.

Proof. Assume $\phi\colon CV_n \to CV_k$ does not decrease the distance and is restricted to each simplex an isometric embedding. Let $A, B \in CV_n$. Since (CV_n, d_R) is a geodesic space there exists a geodesic γ from A to B. Let $A_1, \ldots, A_N$ be the points when γ enters or leaves a simplex and set $A_0 := A, A_{N+1} := B$. Since A_i and A_{i+1} lie in a common closed simplex we have $d_R(A_i, A_{i+1}) = d_R(\phi(A_i), \phi(A_{i+1}))$ for all i. By the triangle inequality we have

$$d_R(\phi(A), \phi(B)) \leq \sum_i d_R(\phi(A_i), \phi(A_{i+1})) = \sum_i d_R(A_i, A_{i+1}) = d_R(A, B).$$

On the other hand we have by assumption $d_R(\phi(A), \phi(B)) \geq d_R(A, B)$. Thus equality holds and ϕ is an isometric embedding. □

Proposition 5.5
Let $\phi\colon CV_2 \to CV_3$ be constructed as in Example 5.3 with $l(\omega) \geq 2$. Then ϕ is an isometry.

Proof. The proof follows directly form Example 5.2 and Lemma 5.4. □

To get an embedding from CV_2 to CV_k for any $k > 2$ we can instead of a single petal attach any marked, metric graph C to gain an isometric embedding, as long as the smallest loop of C has at least length 2. Note here that changing the volume of C also changes the corresponding map and in particular C is not normalised. Furthermore, observe that C does not have to be fixed along CV_2, namely we can slightly vary the graph we attach as long as the changing of C is reasonably small compared to the actual distance in CV_2 as can be seen in the following example.

Example 5.6
Let $\phi\colon CV_2 \to CV_4$ be constructed similar to Proposition 5.5 but instead of a loop we attach a (fixed) marked theta-graph with varying edge lengths to a normalised representant. That is for a full dimensional simplex $\Delta \subset CV_2$ we have either

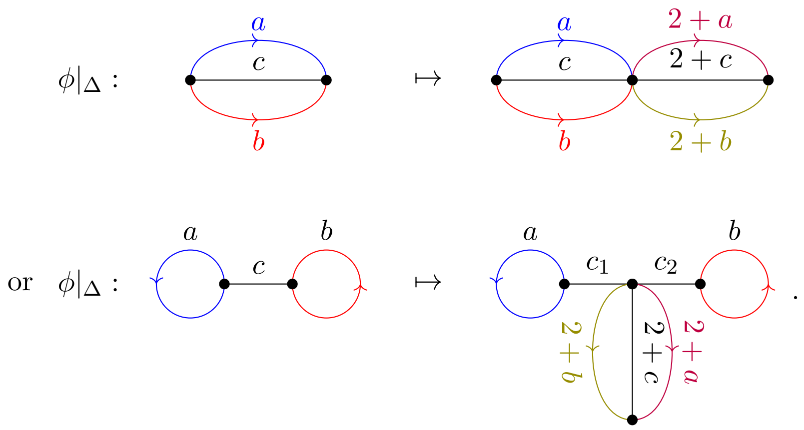

Here we need to take care that we name the edge lengths a, b, c such that they coincide when we pass a face in CV_2. As before we can check that the new candidates are less stretched than the old candidates. Let a, b, c and a', b', c' be edge lengths of two normalised points $A, B \in \Delta$. By Lemma 1.14 we have

$$\begin{aligned} \frac{1+a'}{1+a} &= \frac{a'+b'+c'+a'}{a+b+c+a} \\ &\leq \max\left\{\frac{a'+b'}{a+b}, \frac{a'+c'}{a+c}\right\}. \end{aligned}$$

This means we have $\frac{1+a'}{1+a}, \frac{1+b'}{1+b}, \frac{1+c'}{1+c} \leq \Lambda_R(A, B)$ if Δ is a theta-graph. By Lemma 1.14 we thus have that $\phi|_\Delta$ is an isometric embedding if Δ corresponds to a theta-graph.

Let now Δ correspond to a barbell. By rewriting $1 + a = a + b + c + a = 3/2a + 1/2b + 1/2(a + 2c + b)$ we get by Lemma 1.14 $\frac{1+a'}{1+a}, \frac{1+b'}{1+b}, \frac{1+c'}{1+c} \leq \Lambda_R(A, B)$. Let α be a simple loop in A and ω be a loop in the attached graph. Then we have by Lemma 1.14:

$$\begin{aligned} \frac{l_{\phi(B)}(\alpha\omega)}{l_{\phi(B)}(\alpha\omega)} &= \frac{l_B(\alpha) + 2c'_1 + 2 + l_{\phi(B)}(\omega) - 2}{l_A(\alpha) + 2c_1 + 2 + l_{\phi(A)}(\omega) - 2} \\ &\leq \max\left\{\frac{l_B(\alpha) + 2\lambda c' + 2}{l_A(\alpha) + 2\lambda c + 2}, \frac{1+a'}{1+a}, \frac{1+b'}{1+b}, \frac{1+c'}{1+c}\right\} \\ &\leq \Lambda_R(A, B) \end{aligned}$$

where we use in the last inequality that all of the fractions in the maximum are bounded by $\Lambda_R(A, B)$ (see Example 5.2 (iii) for the first term), hence ϕ is again an isometric embedding.

In a similar manner we can also construct an example where the attached graph crosses a face inside a simplex $\Delta \subset CV_2$, e.g. by attaching a figure of eight which slightly deforms to a theta-graph near the center of Δ.

Similarly we can attach a graph or petal at any given vertex of a marked graph to get an embedding from a simplex of $\Delta \subset CV_n$ to some CV_k for $k > n$. We will see in Lemma 5.8 that attaching a large enough petal yields again an isometric embedding. To give a rough bound what large enough means, we will use the following lemma:

Lemma 5.7
Let $\overline{\Delta} = (\Gamma, m) \subset CV_n$ be a closed simplex, $A, B \in \overline{\Delta}$ be two normalised representants and $e \in E(\Gamma)$ an edge of Γ. Then we can bound the shrinking and stretching of e in terms of the distance from A to B:

(i) If e is shrunken by $s := l_A(e) - l_B(e) > 0$, then the stretching from A to B is at least $\Lambda_R(A, B) \geq 1 + s$.

(ii) If e is stretched by $s := l_B(e) - l_A(e) > 0$, then the stretching from A to B is at least $\Lambda_R(A, B) \geq 1 + \frac{s}{3n-4}$.

Proof. (i) Let A' be a connected component of $A \setminus \{e\}$ and B' be the corresponding connected component of $B \setminus \{e\}$. We can view A' and B' as elements of the same

CV_k for some $k \le n-1$. Hence, we get by the definition of the metric that there exists a loop α in A' with

$$\frac{l_{B'}(\alpha)}{l_{A'}(\alpha)} = \Lambda_R(A', B') \ge \frac{\mathrm{vol}(B')}{\mathrm{vol}(A')}$$
$$\Rightarrow \Lambda_R(A, B) \ge \frac{l_B(\alpha)}{l_A(\alpha)} = \frac{l_{B'}(\alpha)}{l_{A'}(\alpha)} \ge \frac{\mathrm{vol}(B')}{\mathrm{vol}(A')}.$$

If $A \setminus \{e\}$ has two components, we choose A' to be the component such that the fraction $\frac{\mathrm{vol}(B')}{\mathrm{vol}(A')}$ is maximal. We have then by Lemma 1.14

$$\frac{\mathrm{vol}(B')}{\mathrm{vol}(A')} \overset{1.14}{\ge} \frac{1 - l_B(e)}{1 - l_A(e)} = \frac{1 - l_A(e) + s}{1 - l_A(e)} \overset{1.14}{\ge} \frac{1 - l_A(e) + s + l_A(e)}{1 - l_A(e) + l_A(e)} = 1 + s,$$

which concludes the proof.

(ii) There are at most $3n-3$ edges in Γ. Since A and B are normalised, we have $1 = \sum_{e \in E(\Gamma)} l_A(e) = \sum_{e \in E(\Gamma)} l_B(e)$, hence there exists at least one edge $e' \in E(\Gamma)$ with $l_A(e') - l_B(e') \ge \frac{s}{3n-4}$. The inequality follows then by part (i).

□

As before we get now an isometric embedding for any simplex:

Lemma 5.8
Let $\overline{\Delta} \subset CV_n$ be a closed simplex with topological representant (Γ, m) and $K \ge 2(2n-3)(3n-4)$. Then for each vertex $p \in V(\Gamma)$ the map

$$\psi_{p,K} : \overline{\Delta} \to CV_{n+1} \quad , \quad (\Gamma, l, m) \mapsto \psi_p(\Gamma, l, m)$$

that attaches to a normalised representant (Γ, l, m) a loop of length K with marking ω is an isometric embedding.

Proof. As in Example 5.2 we only need to check that none of the new candidates are witnesses. The attached loop has constant length and by Corollary 1.15 a figure of eight containing the attached loop will never be a witness. Hence, we only need to consider barbells containing the attached loop. Let $A, B \in \Delta$ be two normalised points, $\omega\beta$ be such a barbell and ρ be the path between its two loops ω and α. As a spanning tree of Γ has at most $2n-3$ edges also the path ρ contains at most $2n-3$ edges. By Lemma 5.7 we have then $l_B(\rho) - l_A(\rho) \le (2n-3)(3n-4)(\Lambda_R(A, B) - 1)$. Assuming $\omega\beta$ is a witness from A to B, then it is at least as stretched as α and we have the following inequalities:

$$\begin{aligned}\frac{l_B(\omega\beta)}{l_A(\omega\beta)} &= \frac{l_B(\omega) + 2l_B(\rho) + l_B(\alpha)}{l_A(\omega) + 2l_A(\rho) + l_A(\alpha)} \\ &= \frac{K + 2(l_B(\rho) - l_A(\rho)) + 2l_A(\rho) + l_B(\alpha)}{K + 2l_A(\rho) + l_A(\alpha)} \\ &\overset{1.14}{\leq} \frac{K + 2(l_B(\rho) - l_A(\rho)) + 2l_A(\rho)}{K + 2l_A(\rho)} \\ &\overset{5.7}{\leq} \frac{K + 2(2n-3)(3n-4)(\Lambda_R(A,B) - 1) + 2l_A(\rho)}{K + 2l_A(\rho)} \\ &= 1 + (\Lambda_R(A,B) - 1) \cdot \frac{2(2n-3)(3n-4)}{K + 2l_A(\rho)} \\ &< 1 + (\Lambda_R(A,B) - 1) = \Lambda_R(A,B),\end{aligned}$$

which contradicts that $\omega\alpha$ is maximally stretched. Hence, any barbell containing ω is never a witness between two points in Δ and so $\psi_{p,K}$ is an isometric embedding. □

We will see in the next section how to extend this embedding to a global isometric embedding.

5.2 Naive embedding via trees

To get a similar embedding as in Proposition 5.5 for all CV_n we need to construct a coherent choice of representant and attaching point for each element in CV_n. Recall that points in Outer Space can also be written as metric trees with free, minimal, isometric actions of F_n (see Definition 1.3) by passing to their universal cover. So instead of choosing a representant and a point in the finite graph, we will choose a point in the covering tree. The glueing of two finite graphs corresponds then to the following "glueing" or interweaving of two trees:

Definition 5.9
Let $(T^{(1)}, p_1)$ and $(T^{(2)}, p_2)$ be two pointed, metric trees with a minimal, free action of $G_1 := F_{n_1}$ and $G_2 := F_{n_2}$ respectively. For $n = n_1 + n_2$ we identify $F_n \cong G := G_1 * G_2$ and define their *interweaving at* p_1 *and* p_2 as

$$\begin{aligned}T^{(1)}\#_{(p_1,p_2)}T^{(2)} &:= G \times (T^{(1)} \sqcup T^{(2)})/\sim \\ &= \{(\alpha, p) \mid p \in T^{(1)} \sqcup T^{(2)}, \alpha \in G\}/\sim,\end{aligned}$$

where $\sim$ is the equivalence relation generated by the two relations

$$\begin{aligned}(\alpha,p)R_1(\beta,q) &:\iff \exists i \in \{1,2\} : \beta^{-1}\alpha \in G_i, p, q \in T^{(i)} \text{ and } \beta^{-1}\alpha \cdot p = q \\ \text{and } (\alpha,p)R_2(\beta,q) &:\iff p = p_1, q = p_2 \text{ and } \alpha = \beta.\end{aligned}$$

On $T^{(1)}\#_{(p_1,p_2)}T^{(2)}$ we define an G-action by left multiplication $\alpha \cdot (\beta, p) := (\alpha\beta, p)$. This is well defined as the relations R_1 and R_2 are invariant under left-multiplication.

By the following lemma $T^{(1)}\#_{(p_1,p_2)}T^{(2)}$ is the universal cover of the glueing of the two quotient graphs of $T^{(1)}/G_1$ and $T^{(2)}/G_2$. In particular $T^{(1)}\#_{(p_1,p_2)}T^{(2)}$ is again a tree and the G-action on it is minimal. Recall that given a point in CV_n as a marked, metric graph, its corresponding point in terms of trees with a minimal, free F_n action is its universal cover together with the deck transformation group. That means after identifying once and for all F_n with G we have that $T^{(1)}\#_{(p_1,p_2)}T^{(2)}$ together with the F_n-action is a point in CV_n and the interweaving yields a similar map as the glueing of finite, marked graphs in Section 5.1.

Lemma 5.10
Let $(T^{(1)}, p_1), (T^{(2)}, p_2)$ and G_1, G_2 be as in Definition 5.9 and $\Gamma_1 := T^{(1)}/G_1, \Gamma_2 := T^{(2)}/G_2$ be their quotient graphs. Then the interweaved tree $T := T^{(1)}\#_{(p_1,p_2)}T^{(2)}$ is the universal cover of the wedge sum $\Gamma_1 \vee \Gamma_2$, that is Γ_1 and Γ_2 glued at the points $[p_1]$ and $[p_2]$, and the F_n action defined in 5.9 is the corresponding deck transformation group. In particular we have that T, together with the F_n action and the metric induced by $T^{(1)}$ and $T^{(2)}$, is an element of Outer Space CV_n.

Proof. Observe that the relation R_1 only identifies points in $(F_n, T^{(i)})$ by the G_i-action on $T^{(i)}$, that is for $\alpha \in F_n, \beta \in G_i, p \in T^{(i)}$ we have $(\alpha\beta, p) = (\alpha, \beta \cdot p)$. Furthermore, R_2 only identifies pairs of glued points (α, p_1) and (α, p_2). Hence, we get for each $\alpha \in F_n$ an embedding $\phi_\alpha : T^{(1)} \to T^{(1)}\#_{(p_1,p_2)}T^{(2)}, p \mapsto (\alpha, p)$. Two such embeddings ϕ_α, ϕ_β have the same image if and only if α and β are in the same right G_1-coset. Indeed, if we have $\alpha = \beta\omega_1$ for some $\omega_1 \in G_1$, then we have $\phi_\alpha(p) = (\alpha, p) = (\alpha, \omega_1\omega_1^{-1}p) = (\alpha\omega_1, \omega_1^{-1}p) = \phi_\beta(\omega_1^{-1}p)$ and on the other hand if $\phi_\alpha(p) = \phi_\beta(q)$ for some point $p \in T^{(1)}$ then (α, p) and (β, q) are at most identified by relation R_1 and we have $\beta^{-1}\alpha =: \omega_1 \in G_1$. In particular all these embeddings are either disjoint or have the same image.

Similar we have embeddings $\psi_\alpha : T^{(2)} \to T^{(1)}\#_{(p_1,p_2)}T^{(2)}$ for $\alpha \in F_n$. By construction it is clear that the images of the ϕ_α and ψ_α for all $\alpha \in F_n$ cover the whole space $T^{(1)}\#_{(p_1,p_2)}T^{(2)}$. In particular we have that at each vertex T looks locally like a part of the wedge sum of $T^{(1)}$ and $T^{(2)}$ at p_1 and p_2.

We will write for a point $p \in T^{(1)} \sqcup T^{(2)}$ its image in the glued quotient graph as $[p] \in \Gamma := \Gamma_1 \vee \Gamma_2$ and define the map $\pi : T \to \Gamma, (\alpha, p) \mapsto [p]$. As R_1 and R_2 only identify points which are also identified under the quotient maps and wedge sum, this is well defined. Since T looks locally like the wedge sum of $T^{(1)}$ and $T^{(2)}$, we have that π is a covering map. As we have embedded $T^{(1)}$ and $T^{(2)}$ into T we can see the monodromy actions as right multiplication in these embeddings, i.e. for a path $\omega \in \pi_1(\Gamma_1, [p_1])$ we have by the anti-isometry between $\pi_1(\Gamma_1, [p_1])$ and the deck transformation group $\mathrm{deck}(T_1/\Gamma_1) = G_1$ the equalities $(\alpha, p_1) \bullet \omega = \phi_\alpha(p_1) \bullet \omega = \phi_\alpha(\omega \cdot p_1) = (\alpha\omega, p_1)$ and similar for a path in $\pi_1(\Gamma_2, [p_2]) = G_2$. As we can write a path $\omega \in \pi_1(\Gamma, [p_1])$ uniquely as a word of paths $\omega = \omega_1 \dots \omega_k \in \pi_1(\Gamma_1, [p_1]) * \pi_1(\Gamma_2, [p_2])$, we get inductively $(\mathrm{id}, p_1) \bullet \omega = (\omega_1, p_1) \bullet \omega_2 \dots \omega_k = \dots = (\omega_1 \dots \omega_k, p_1) = (\omega, p_1)$. Since each element in the fibre $\pi^{-1}([p_1])$ can be uniquely written as an element (ω, p_1) for some $\omega \in F_n$, the monodromy action gives a bijection between $\pi_1(\Gamma, [p_1])$ and the fibre $\pi^{-1}([p_1])$, which means that T is the universal cover of Γ. As the F_n action on (id, p_1) is the dual of the monodromy action of $\pi_1(\Gamma, [p_1])$, we have that F_n is the deck transformation group of π. □

Observe here that interweaving $(T^{(1)}, p_1)$ and $(T^{(2)}, p_2)$ yields a different element in CV_n than interweaving $(T^{(1)}, \alpha p_1)$ and $(T^{(2)}, p_2)$ for some $\alpha \in G_1$, namely in terms of finite graphs we conjugated Γ_1 by α while keeping Γ_2 the same. This means when we lift the wedge sum of two finite graphs to an interweaving of their universal covers the fibres of the attaching points play an important role.

A direct corollary of Lemma 5.10 is that interweaving trees does not change length:

Corollary 5.11
Let $T^{(1)} \in CV_{n_1}$ and $T^{(2)} \in CV_{n_2}$ be two elements in Outer Space, $p_1 \in T^{(1)}$ and $p_2 \in T^{(2)}$ two points and let $T = T^{(1)} \#_{p_1,p_2} T^{(2)}$ be their interweaving. For $i \in \{1, 2\}$ and an element $\alpha \in G_i := F_{n_i}$ we identify it with its corresponding element $\alpha \in G_1 * G_2$. Then α has the same length in T and T_i, i.e. we have $l_T(\alpha) = l_{T_i}(\alpha)$.

Proof. There are now two easy ways to see this: For example we can calculate the length of α as the length of its cyclically reduced representation in the quotient graph. By Lemma 5.10 these representations are equal and hence have the same length.

Alternatively let without loss of generality $\alpha \in G_1$ and take the length of α as the translation length along its axis $T^{(1)}_\alpha \subset T^{(1)}$ or $T_\alpha \subset T$, respectively. Recall that $T_\alpha \subset T$ is the unique embedding of $\mathbb{R}$ in T on which α acts as translation. Since we already have the embedding $\phi_{\mathrm{id}}(T^{(1)}_\alpha) \subset T$ on which α acts via translation, we have $T_\alpha = \phi_{\mathrm{id}}(T^{(1)}_\alpha)$ and so for any $p \in T_\alpha$:

$$l_T(\alpha) = d((\mathrm{id}, p), \alpha(\mathrm{id}, p)) = d(p, \alpha p) = l_{T^{(1)}}(\alpha).$$

□

In our setting $T^{(1)}$ will be a (normalised) element of CV_n with a chosen base point $p_1 \in T^{(1)}$ and $T^{(2)}$ will be a fixed marked graph with a base point p_2, for example the standard Cayley-graph of F_k with the origin as base point. We will present two slightly different ways to coherently construct such base points p_1 for each element of CV_n. For the first construction we use translation axes to determine a point in the tree. The second constructions is due to Skora in [Sko90, Chapter 5]. It should be noted here that the second construction only yields an isometric embedding for bounded subsets of CV_n and not an isometric embedding for the whole Outer Space.

We will use the translation axes from Lemma 1.4 to define a base point for each element of CV_n. Observe that two elements of F_n have the same translation axes in a marked graph $(T, l, m) \in CV_n$ if and only if they have a non-trivial common power in F_n, otherwise we would have a non-trivial stabiliser of a vertex. Since T is a tree and the translation axes are embeddings of $\mathbb{R}$, we have in particular that two different translation axes intersect at most on finitely many edges.

Definition 5.12
Let $\alpha, \beta \in F_n$ be two non-trivial elements with no common power. Then we define for each $(T, l, m) \in CV_n$ the base point $p \in T$ as follows:

(i) If the translation axes T_α and T_β are disjoint, then $p \in T_\alpha$ is the point in T_α which is the closest to T_β.

(ii) If the translation axes T_α and T_β intersect, then $p \in T_\alpha \cap T_\beta$ is the first point of $T_\alpha \cap T_\beta$ in the direction of the α-translation.

As T_α and T_β are different translation axes, their intersection is always at most a finite segment and hence the base point is a well-defined vertex. Furthermore, an optimal change of marking map between two representants of the same point in CV_n preserve the translation axes and hence also the base point.

Observe that this base point does not depend on the lengths of the edges and collapsing edges sends base points to base points. In particular we have a coherent choice of a base point: Let $\Delta, \Delta' \subset CV_n$ be two simplices sharing a face $F \subset CV_n$. Then the base point for any $T \in F$ identifies with the base point from $T' \in \Delta$ or Δ' after collapsing the corresponding edges. Similar to Proposition 5.5 we get the following isometric embedding.

Theorem 5.13
Let $K \geq 2(2n-3)(3n-4)$ and $\alpha, \beta \in F_n \setminus \{\text{id}\}$ be two elements without a non-trivial common power. For $(T, l, m) \in CV_n$ let $p \in T$ be the base point defined as in Definition 5.12. Let R be the Cayley-graph of $\mathbb{Z}$ with generator 1 endowed with the metric such that each edge has length K. As in Definition 5.9 we now interweave the marked tree $(T, l, m) \in CV_n$ at p with R at the point 0, that is in terms of finite graphs we attach a single loop to $\Gamma = T/F_n$. Then the resulting map

$$\psi_K : CV_n \to CV_{n+1} \quad , \quad T \mapsto T\#_{(p,0)}R.$$

is an isometric embedding.

Proof. The map $\psi_{p,K}$ is by Lemma 5.8 restricted to each closed simplex $\overline{\Delta} \subset CV_n$ an isometric embedding. As the translation lengths for elements $\alpha \in F_n$ are equal for $(T, l, m) \in CV_n$ and $\psi_{p,K}(T, l, m)$ we have that ψ_K at most increases the distance. Hence, by Lemma 5.4 it follows that $\psi_{p,K}$ is an isometric embedding. $\square$

It is clear that instead of the Cayley-graph of $\mathbb{Z}$ we can attach any thick enough marked tree and slightly vary the attached tree as we did in Example 5.6.

Instead of attaching one marked tree to an element in CV_n, recall that Definition 5.9 gives an embedding $CV_n \times CV_k \to CV_{n+k}$ by interweaving two points together. To get an isometric embedding we will, instead of glueing two graphs directly on each other, introduce a long enough edge in between the two elements. That is we first add a leaf to an element of CV_n and afterwards attach an element of CV_k to the other vertex of that leaf. As before this yields the following proposition:

Proposition 5.14
Let $K \geq 2(2n-3)(3n-4) + 2(2k-3)(3k-4)$ and for $(T, l, m) \in CV_n, (T', l', m') \in CV_k$

let $p \in T$ and $q \in T'$ be the base points as defined in Definition 5.12 for two fixed tuples $\alpha, \beta \in F_n, \alpha', \beta' \in F_k$. We define the map

$$\psi \colon CV_n \times CV_k \to CV_{n+k} \quad , \quad (T, T') \mapsto T\#_{(p,0)}[0, K]\#_{K,q}T',$$

which interweaves two normalised T and T' and adds an edge of length K between each glued pair of base points. Then ψ is an isometric embedding in regard to the product metric on $CV_n \times CV_k$, that is for $A, B \in CV_n, A', B' \in CV_k$ we have $d_R(\phi(A, A'), \phi(B, B')) = \max\{d_R(A, B), d_R(A', B')\}$.

We can use this to construct isometric embeddings where we vary the attached graph. Namely we have:

Corollary 5.15
Let ψ be as in Proposition 5.14.

(i) Let $\phi \colon CV_n \to CV_k$ be a map which at most decreases the distance, i.e. $d_R(A, B) \geq d_R(\phi(A), \phi(B))$. Then the map

$$\psi \circ (\mathrm{id}, \phi) \colon CV_n \to CV_{n+k} \quad , \quad A \mapsto \psi(A, \phi(A))$$

is an isometric embedding.

(ii) Let $\gamma \colon \mathbb{R}_{\geq 0} \to CV_k$ be a rectifiable path parametrised by its symmetric length and $C \in CV_n$. Then the map

$$\phi \colon CV_n \to CV_{n+k} \quad , \quad A \mapsto \psi\big(A, \gamma(d(A, C))\big)$$

is an isometric embedding.

Proof. (i) follows directly from Proposition 5.14. For (ii) we have for $A, B \in CV_n$ by the triangle inequality and since γ is parametrised by length

$$d(\gamma(d(A, C)), \gamma(d(B, C))) \leq |d(A, C) - d(B, C)| \leq d(A, B) \leq d_R(A, B),$$

which by (i) concludes the proof. □

As a direct application of Corollary 5.15(ii) we can now construct isometric embeddings from CV_n to CV_{n+k} such that the image of a simplex in CV_n intersects arbitrarily many simplices in CV_{n+k} by choosing γ accordingly.

An alternative way to construct a base point is due to Skora. The following definitions and lemmas up to and including Definition 5.20 are besides slight variations from [Sko90] Chapter 5. We start with a generalisation of length functions and the characteristic set from Lemma 1.4.

Definition 5.16
Let $S \subset F_n \setminus \{\mathrm{id}\}$ be a non-empty, finite subset. For $A = (T, l, m) \in CV_n$ we define the *length function* as

$$l_A(S) := \min_{p \in T} \max_{\alpha \in S} d_A(p, \alpha \cdot p).$$

The *characteristic set* of S is then the set of minimally displaced points

$$A_S := \{p \in T \mid \max_{\alpha \in S} d_A(p, \alpha \cdot p) = l_A(S)\}.$$

We will see in Lemma 5.17 that the characteristic set is non-empty and hence the minimum in the length function is indeed attained. While the translation axes T_α depend only on the topological type (T, l, m), the characteristic set also depends on the edge lengths, which is reflected by the notion A_S instead of possibly T_S.

Observe that if $S = \{\alpha\}$ is a singleton, then its length is just the length of α and its characteristic set is the translation axis T_α of α. It is a useful fact that α acts on its translation axis via a translation and on T as an isometry, hence for any $p \in T$ we have by [CM87, 1.3] for its displacement length $d_A(p, \alpha \cdot p) = 2d_A(p, T_\alpha) + l_A(\alpha)$:

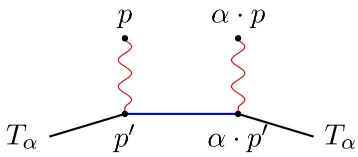

Figure 23: Displacement of a point relative to the translation axis of α.

We will now see that the characteristic set A_S is either a translation axis, a bounded subset of a translation axis or a point.

Lemma 5.17
Let $S \subset F_n \setminus \{\mathrm{id}\}$ be a non-empty, finite subset and $A = (T, l, m) \in CV_n$. Then the characteristic set is non-empty and of one of the following types:

(i) If for all $\alpha, \beta \in S$ the translation axes $T_\alpha = T_\beta$ are the same, then the characteristic set $A_S = T_\alpha$ is this translation axis and we have $l_A(S) = \max_{\alpha \in S} l_A(\alpha)$.

(ii) If there exist $\alpha, \beta \in S$ with different translation axes $T_\alpha \neq T_\beta$ and $l_A(S) = l_A(\alpha)$, then the characteristic set is a closed, bounded segment $A_S \subset T_\alpha$ of the translation axis T_α.

(iii) If $l_A(S) > \max_{\alpha \in S} l_A(\alpha)$ holds, then $A_S = \{p_S\}$ is a single point and there exist two elements $\alpha, \beta \in S$ with disjoint translation axes $T_\alpha \cap T_\beta = \emptyset$ which realise the length of S, i.e. we have $d_A(p_S, \alpha \cdot p_S) = l_A(S) = d_A(p_S, \beta \cdot p_S)$. In particular p_S lies on the shortest path connecting T_α and T_β.

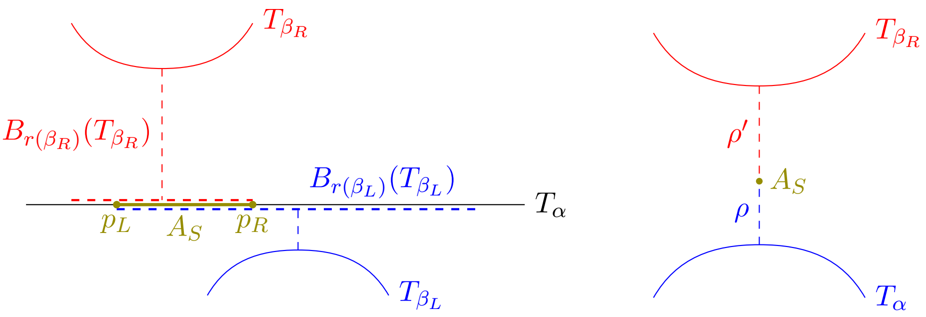

Figure 24: Characteristic sets of type (ii) and (iii) from Lemma 5.17

Proof. (i) If all elements of S have the same translation axis T_α, then clearly the minimum of $d_A(p, \alpha \cdot p)$ is attained only by points in T_α and (i) follows.

Let now $\alpha, \beta \in S$ be two elements with different translation axes T_α and T_β. For some positive number $R > 0$ we denote as usual by

$$\overline{B_R(T_\alpha)} := \{p \in T \mid \exists q \in T_\alpha \text{ with } d_A(p, q) \leq R\}$$

the closed ball of radius R around T_α. As T_α and T_β are two different isometric embeddings of $\mathbb{R}$, we have that T_α and T_β diverge, i.e. for every $R > 0$ the intersection of their closed R-neighbourhoods $\overline{B_R(T_\alpha)} \cap \overline{B_R(T_\beta)}$ is a finite subtree of T.

We will first show that A_S is not empty. Let $p \in T$ be some point on the tree and let $R = \max_{\omega \in S} d_A(p, \omega \cdot p)/2$ be half of its maximal displacement by elements of S. Let furthermore $\chi := \bigcap_{\omega \in S} \overline{B_R(T_\omega)}$ be the intersection of the closed R-neighbourhoods of the axes of the elements in S.

Let $q \in T \setminus \chi$ be a point outside of χ, then we have $d_A(q, T_\omega) > R$ for some $\omega \in S$ and hence

$$d_A(q, \omega \cdot q) = 2d_A(q, T_\omega) + l_A(\omega) > 2R = \max_{\omega \in S} d_A(p, \omega \cdot p) \geq l_A(S).$$

This means we only need to consider points in χ to calculate $l_A(S)$. As χ is a finite subtree, the continuous function $\max_{\omega \in S} d_A(\bullet, \omega \cdot \bullet)$ attains its minimum in χ and hence A_S is non-empty.

(ii): Consider again the equality $d_A(p, \alpha \cdot p) = 2d_A(p, T_\alpha) + l_A(\alpha)$. By the definition of characteristic set we have $A_S = \bigcap_{\omega \in S} \overline{B_{r(\omega)}(T_\omega)}$ with $r(\omega) := (l_A(S) - l_A(\omega))/2$. If we have now $l_A(\alpha) = l_A(S)$, then we have clearly $A_S \subset T_\alpha$. Let now $\beta \in S$ be an element with a different translation axis than α. We have then $A_S \subset T_\alpha \cap \overline{B_{r(\beta)}(T_\beta)}$ and hence A_S is a bounded subsegment of T_α. As each $\overline{B_{r(\beta)}(T_\beta)} \cap T_\alpha$ is a closed subsegment the characteristic set A_S is also a closed subsegment of T_α.

(iii): Assume now the conditions of statement (iii). Let $p \in A_S$ and $\alpha \in S$ be such that $d_A(p, \alpha \cdot p) = l_A(S)$ holds. Since we have $l_A(S) > l_A(\alpha)$ we have $p \notin T_\alpha$. Let ρ be the shortest path from p to T_α. There exists at least one $\beta \in S \setminus \{\alpha\}$ with $d_A(p, \beta \cdot p) = l_A(S)$ such that the shortest path ρ' from p to T_β is disjoint from ρ, else we could slightly move p along ρ and ρ' to decrease the maximal displacement length. In particular we have that T_α and T_β are disjoint and their shortest connection is along the paths ρ and ρ'. For the uniqueness of p let $q \in T \setminus \{p\}$. As T is a tree q lies in the component of $T \setminus \{p\}$ which does not contain T_α or T_β. Hence, we have without loss of generality $d_A(q, T_\alpha) = d(q, p) + d(p, T_\alpha)$ and in particular $d_A(q, \alpha \cdot q) = 2d_A(q, T_\alpha) + l_A(\alpha) > 2d_A(p, T_\alpha) + l_A(\alpha) = d_A(p, \alpha \cdot p) = l_A(S)$, thus $q \notin A_S$. □

As the displacement length $d_A(p, \alpha \cdot p)$ depends continuously on $p \in T$ and the lengths of the edges of A, we can convince ourselves that the length $l_A(S)$ and the characteristic set depend continuously on $A \in CV_n$. For completeness we will give a short proof for this as the following lemma:

Lemma 5.18
Let $S \subset F_n \setminus \{\mathrm{id}\}$ be a non-empty finite subset. Then the function $l_\bullet(S)\colon CV_n \to \mathbb{R}$, $A \mapsto l_A(S)$ is continuous.

Proof. If S is a singleton $\{\alpha\}$ or all elements of S have a common root α, then the statement follows as $l_A(S) = l_A(\alpha)^k$ is continuous. Hence, we can assume that S has at least two elements with different translation axes.

As CV_n is a locally finite simplicial complex (with some faces missing), it is enough to show that $l_\bullet(S)|_{\overline{\Delta}}$ is continuous for closed simplices. Let $\overline{\Delta} = (T, \cdot, m) \subset CV_n$ be a closed simplex in CV_n. For each $\alpha \in S$ choose a vertex $p_\alpha \in T_\alpha$ on the translation axis and for distinct $\alpha, \beta \in S$ let $\rho_{\alpha,\beta}$ be the shortest path between p_α and p_β. Each of these $\rho_{\alpha,\beta}$ is a finite subtree of T and any tuple of them can be connected by a third one at their endpoints. Hence, their union $\chi = \bigcup_{\alpha,\beta \in S} \rho_{\alpha,\beta}$ is also a finite subtree of T. By Lemma 5.17 the characteristic set lies either on a translation axis T_α or on a path $\rho_{\alpha,\beta}$.

Let $A \in \overline{\Delta}$ be some point. We will now show that in the former case A_S intersects with a $\rho_{\alpha,\beta}$. To this end assume $p_\alpha \notin A_S$ and let $q \in A_S$ be the closest point in A_S to p_α. As we have seen in the proof of Lemma 5.17 (ii) there exists a $\beta \in S$ with $d(q, T_\beta) = (l_A(S) - l_A(\beta))/2$ and $d(q', T_\beta) > d(q, T_\beta)$ for any point $q' \in T_\alpha$ between p_α and q. Thus, any path from p_α to T_β and in particular $\rho_{\alpha,\beta}$ has to pass through $q \in A_S$.

This means we have $\chi \cap A_S \neq \emptyset$ and hence $l_A(S) = \min_{p \in \chi} \max_{\omega \in S} d_A(p, \omega \cdot p)$. Since χ is a finite subtree and $d_A(p, \alpha \cdot p)$ depends continuously on $p \in \chi$ and the edge lengths, the length function $l_\bullet(S)$ is continuous. □

Example 5.19
Let $S = \{\alpha, \beta\alpha\beta^{-1}\}$ for some $\alpha, \beta \in F_n$. Then the characteristic set can be described nicely as we will see in the following. Recall that conjugation does not change the length, i.e. α and $\beta\alpha\beta^{-1}$ have the same lengths $l_A(\alpha) = l_A(\beta\alpha\beta^{-1})$ for all $A \in CV_n$. Hence, we have the following two cases.

If the intersection $T_\alpha \cap T_{\beta\alpha\beta^{-1}} \neq \emptyset$ is non-empty, then we have by Lemma 5.17 (ii) $A_S = T_\alpha \cap T_{\beta\alpha\beta^{-1}}$ (see Figure 25).

Otherwise the two translation axes T_α and $T_{\beta\alpha\beta^{-1}} = \beta \cdot T_\alpha$ are disjoint, hence there exists a unique shortest path from T_α to $\beta \cdot T_\alpha$ with midpoint p. Each other point

$p' \in T \setminus \{p_S\}$ is further away from T_α or $\beta \cdot T_\alpha$ as p, thus the characteristic set is $A_S = \{p\}$.

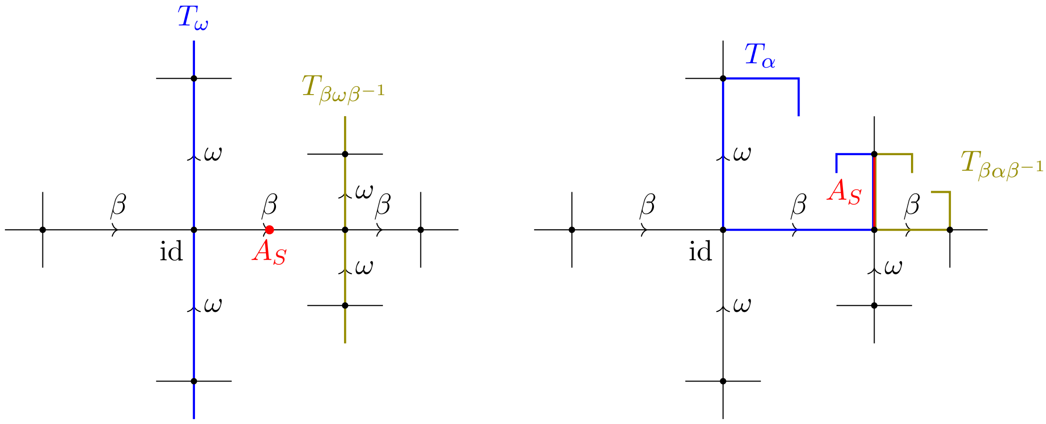

Figure 25: The Cayley-graph of F_2 with generators ω and β and its characteristic set with respect to $S = \{\omega, \beta\omega\beta^{-1}\}$ and $S' = \{\alpha, \beta\alpha\beta^{-1}\}$ for $\alpha := \beta\omega\beta^{-1}\omega^{-1}$.

From now on let $S \subset F_n$ be a fixed finite subset with at least two elements $\alpha, \beta \in S$ with different translation axes $T_\alpha \neq T_\beta$. For example we can take a basis $\{\alpha_1, \dots, \alpha_n\} \subset F_n$ or $\{\alpha, \beta\alpha\beta^{-1}\} \subset F_n$ an element and its conjugate where α and β generate a subgroup of rank two.

Definition 5.20
By Lemma 5.17 A_S is either a point or a bounded segment of a translation axis. We call the middle point of A_S the *base point* $p_S \in A_S$. Observe that for two different representants $(T, l, m) = (T', l', m') \in CV_n$ the base point p_S is preserved under the corresponding homothety between T and T' as the translation axes are also preserved, that is the homothety has to send the base point p_S of T to the base point of T'.

We will now apply the base point constructed from Skora to get again an embedding from CV_n to CV_{n+k}.

Definition 5.21
Let $K > 0$ and R_k be the Cayley-graph of the free group F_k with k generators where each edge has length K and $p_{\text{id}} \in R_k$ is the vertex corresponding to the identity. In terms of marked graphs this means R_k corresponds to the rose with edge length K. We define the map

$$\phi_{S,K} : CV_n \to CV_{n+k} \quad , \quad T \mapsto T\#_{(p_S, p_{\text{id}})} R_k,$$

where T is any normalised representant and $p_S \in T$ be as in Definition 5.20.

Observe that the basepoint $p_S \in T$ moves continuously with respect to on $A \in CV_n$ so the map $\phi_{S,K}$ is continuous. As the base point moves around, we can not use Lemma 5.8 to get an isometric embedding. Nevertheless, if we restrict to a simplex Δ, we still have that $\phi_{S,K}|_\Delta$ is an isometric embedding for large enough K:

Proposition 5.22
Let $\overline{\Delta} \subset CV_n$ be a closed simplex. Then there exists a $\kappa > 0$ such that for all $K \geq 2\kappa$ the map $\phi_{S,K}|_{\overline{\Delta}}$ from Definition 5.21 is an isometric embedding.

We will later see that this κ is determined by the simplex and by Proposition 5.26 we can not find a fixed κ for all simplices, in particular this construction does not yield an isometric embedding for the whole CV_n as in Theorem 5.13. Keep in mind that by Corollary 5.11 the map $\phi_{S,K}$ at most increases the distance of two points in CV_n, as the lengths of the original paths are the same after interweaving. The main idea of the proof of Proposition 5.22 is that in some sense the base point moves Lipschitz continuously along the underlying tree corresponding to Δ. This movement of the base point can be absorbed by stretching the edges of R_k along the movement (see Lemma 5.23).

To make this argument precise, we will need the following rather technical lemmas to bound the movement of the base point. For two points $A, B \in CV_n$ we get an upper bound for the distance of their images by an optimal change of marking map:

Lemma 5.23
Let $\phi_{S,K} : CV_n \to CV_{n+k}$ be as in Definition 5.21, $A, B \in CV_n$ and let $h\colon A \to B$ be a change of marking map with Lipschitz constant Λ. Furthermore, let $p_{A,S} \in A$ and $p_{B,S} \in B$ be the base points of A and B, respectively, and $r := d_B(p_{B,S}, h(p_{A,S}))$ the distance of the base point of B to the image of the base point of A.

Then there exists a change of marking map $\tilde{h} : \phi_{S,K}(A) \to \phi_{S,K}(B)$ with Lipschitz constant $\max\{\Lambda, \frac{K+2r}{K}\}$. In particular if $r \leq (\Lambda - 1)\frac{K}{2}$ holds, then h and $\tilde{h}$ have the same Lipschitz constant.

Proof. Recall that the interweaving $\phi_{S,K}(A)$ consists of copies of the trees A and R_k. We choose $\tilde{h}$ to map each copy of A like h, i.e. $\tilde{h}((\alpha, p)) = (\alpha, h(p))$ for all $p \in A$. This is well-defined as h is equivariant. Afterwards we extend $\tilde{h}$ linearly to edges of R_k such that it connects their endpoints accordingly. That is $\tilde{h}$ maps each edge of R_k linearly to the path connecting the image of its two endpoints, i.e. for a generator $\beta \in F_k$ $\tilde{h}$ sends the edge $e := [(\alpha, p_{A,S}), (\alpha\beta, p_{A,S})]$ linearly to the path $[(\alpha, h(p_{A,S})), (\alpha\beta, h(p_{A,S}))]$. Since for a point $p \in B$ any path from (α, p) to $(\alpha\beta, p)$ passes through the points $(\alpha, p_{B,S})$ and $(\alpha\beta, p_{B,S})$, we get for the length of the image of e:

$$\begin{aligned} l_{\phi_{S,K}(B)}(\tilde{h}(e)) &= d_{\phi_{S,K}(B)}\big((\alpha\beta, h(p_{A,S})), (\alpha, h(p_{A,S}))\big) \\ &= d_{\phi_{S,K}(B)}\big((\alpha\beta, h(p_{A,S})), (\alpha\beta, p_{B,S})\big) + d_{\phi_{S,K}(B)}\big((\alpha\beta, p_{B,S}), (\alpha, p_{B,S})\big) \\ &\quad + d_{\phi_{S,K}(B)}\big((\alpha, p_{B,S}), (\alpha, h(p_{A,S}))\big) \\ &= d_B(h(p_{A,S}), p_{B,S}) + l(\beta) + d_B(p_{B,S}, h(p_{A,S})) \\ &= 2r + K \end{aligned}$$

By construction $\tilde{h}$ is continuous and commutes with the F_{n+k}-action. As it stretches the copies of A like h and the edges of R_k by the factor $\frac{2r+K}{K}$ it satisfies the claim.

As the following inequalities are equivalent

$$\frac{K+2r}{K} \leq \Lambda \iff K + 2r \leq \Lambda K \iff 2r \leq (\Lambda - 1)K,$$

h and $\tilde{h}$ have the same Lipschitz constant if and only if $r \leq (\Lambda - 1)\frac{K}{2}$ holds. □

Lemma 5.23 in particular implies that if the base points are mapped under an optimal change of marking map close to each other, i.e. if their distance r can be bounded by a linear function $r \leq (\Lambda_R - 1)\kappa$, then for any $K \geq 2\kappa$ the embedding $\phi_{S,K}$ is isometric.

Lemma 5.24
Let $A := (T, l, m_A)$, $B := (T, l, m_B) \in CV_n$ be two normalised points in CV_n with the same topological type and $h\colon A \to B$ be an optimal change of marking map. Furthermore, let $\alpha \in F_n$, $p \in A$ and $r := d_A(p, T_\alpha)$ be the distance of p to the translation axis T_α of α. Then we have

$$d_B(h(p), T_\alpha) \leq \Lambda_R(A, B)r + (\Lambda_R(A, B) - 1) \cdot \#_T(\alpha) =: r',$$

where $\#_T(\alpha)$ is the number of edges of α in T, i.e. for any vertex $q \in T_\alpha$ the number of edges of the shortest path from q to $\alpha \cdot q$.

In other words h sends r-neighbourhoods of T_α in A into r'-neighbourhoods of T_α in B.

Proof. Let $\Lambda_R := \Lambda_R(A, B)$. Since h is F_n-equivariant and Λ_R-Lipschitz we have

$$d_B(h(p), \alpha \cdot h(p)) = d_B(h(p), h(\alpha \cdot p)) \leq \Lambda_R \cdot d_A(p, \alpha \cdot p).$$

Applying this to the distance formulas

$$\begin{aligned} d_A(p, \alpha \cdot p) &= 2 \cdot d_A(p, T_\alpha) + l_A(\alpha) = 2r + l_A(\alpha), \\ d_B(h(p), \alpha \cdot h(p)) &= 2 \cdot d_B(h(p), T_\alpha) + l_B(\alpha), \end{aligned}$$

we get the following inequality

$$\begin{aligned} 2 \cdot d_B(h(p), T_\alpha) &= d_B(h(p), \alpha \cdot h(p)) - l_B(\alpha) \\ &\leq \Lambda_R d_A(p, \alpha \cdot p) - l_B(\alpha) \\ &= \Lambda_R 2r + \Lambda_R l_A(\alpha) - l_B(\alpha) \\ &= \Lambda_R 2r + (\Lambda_R - 1) l_A(\alpha) + (l_A(\alpha) - l_B(\alpha)) \\ &\leq \Lambda_R 2r + 2(\Lambda_R - 1)\#_T(\alpha), \end{aligned}$$

where we used $l_A(\alpha) \leq \#_T(\alpha)$ in the last inequality and that by Lemma 5.7(i) each edge of α is at most shrunken by $\Lambda_R - 1$, hence $l_A(\alpha) - l_B(\alpha) \leq \#_T(\alpha)(\Lambda_R - 1)$. □

We now have the tools we need in order to prove Proposition 5.22:

Proof of Proposition 5.22. Let $A, B \in \Delta$ and $\gamma\colon [0, 1] \to \Delta$ be the straight line from A to B. As the length functions along γ depend linearly on t, we can split $[0, 1]$ into subintervals $0 = t_0 < t_1 < \cdots < t_m = 1$ such that for each $i \in \{1, \ldots m\}$ we have

(i) either $l_{\gamma(t)}(S) > \max_{\alpha \in S} l_{\gamma(t)}(\alpha)$ for all $t_{i-1} < t < t_i$

(ii) or there is an $\alpha \in S$ such that $l_{\gamma(t)}(S) = l_{\gamma(t)}(\alpha)$ for all $t_{i-1} \leq t \leq t_i$.

We will show that for all these subsegments the distance $d_R(\gamma(t_{i-1}), \gamma(t_i))$ is preserved under $\phi_{S,K}$ for some large enough K. Recall that γ is a geodesic, thus $d_R(A, B)$ is realised by the length of γ. Then the triangle inequality implies $d_R(A, B) \geq d_R(\phi_{S,K}(A), \phi_{S,K}(B))$ and since $\phi_{S,K}$ at most increases the distance, equality follows. Let $C := \gamma(t_{i-1})$ and $D := \gamma(t_i)$ be the endpoints of a segment and $\Lambda_R := \Lambda_R(C, D)$ be their stretching factor.

To prove the claim we want to apply Lemma 5.23, i.e. we show that there exists a constant $\kappa = \kappa(\Delta)$ depending only on Δ such that for any optimal change of marking map $h\colon C \to D$ with Lipschitz constant Λ_R the inequality $d_D(p_{D,S}, h(p_{C,S})) \leq (\Lambda_R - 1) \cdot \kappa$ holds. For the following rather technical calculation of κ observe that the following constants only depend on the topological type of Δ.

Let $(T, \cdot, m)$ be a marked tree corresponding to Δ and let throughout the proof $N := \max_{\alpha \in S} \#_T(\alpha)$ be the maximal number of edges of the elements $\alpha \in S$ in T, L be the maximal number of edges in the intersections $T_\alpha \cap T_\beta$ for $\alpha, \beta \in S$ and M be the maximal number of edges between two translation axes T_α and T_β for $\alpha, \beta \in S$, that is the distance of T_α and T_β in T by edges. As usual we consider normalised representants, that means in particular each edge has at most length 1.

The main idea in the two cases is that we approximate $d_D(p_{D,S}, h(p_{C,S}))$ by the distances of $p_{C,S}$ and $p_{D,S}$ to two translation axes whose shortest connection always contain the base point.

For **case (i)** recall that by Lemma 5.17 the base point p_S lies on the shortest path between the two translation axes of two elements $\alpha, \beta \in S$. In particular we have:

$$\begin{aligned} & l_{\gamma(t)}(S) = l_{\gamma(t)}(\alpha) + 2d_{\gamma(t)}(p_S, T_\alpha) = l_{\gamma(t)}(\beta) + 2d_{\gamma(t)}(p_S, T_\beta) \\ \Rightarrow \quad & 2l_{\gamma(t)}(S) = l_{\gamma(t)}(\alpha) + l_{\gamma(t)}(\beta) + 2d_{\gamma(t)}(T_\alpha, T_\beta) \end{aligned}$$

Again as there are only finitely many pairs $\alpha, \beta \in S$ and all these distances and lengths are continuous we can, up to further subdividing of the interval, assume that in the interval $[t_{i-1}, t_i]$ the above equalities are always satisfied for some fixed pair $\alpha, \beta \in S$. In particular the base point p_S of $\gamma(t)$ always lies on the shortest path between the translation axes T_α and T_β. We denote the endpoints of this shortest path as p_β and p_α, respectively (see Figure 26).

Recall that we want to estimate some $\kappa > 0$ such that we have an upper bound for the distance $d_D(h(p_{C,S}), p_{D,S}) \leq (\Lambda_R - 1)\kappa$. For this consider the triangle in D spanned by the points $h(p_{C,S})$, p_α and p_β. We denote the distances of the base points to p_α and p_β as in Figure 26, that is we have

$$\begin{aligned} d_\alpha &:= d_C(p_\alpha, p_{C,S}), & d_\beta &:= d_C(p_\beta, p_{C,S}) \\ d'_\alpha &:= d_D(p_\alpha, p_{D,S}), & d'_\beta &:= d_D(p_\beta, p_{D,S}) \\ r_\alpha &:= d_D(p_\alpha, h(p_{C,S})), & r_\beta &:= d_D(p_\beta, h(p_{C,S})). \end{aligned}$$

As $p_{D,S}$ lies on the path from p_α to p_β and D is a tree, we have that the point $p_{D,S}$ also lies on the path from p_α to $h(p_{C,S})$ or on the path from p_β to $h(p_{C,S})$. We will assume without loss of generality the former and thus get $d_D(p_{D,S}, h(p_{C,S})) = r_\alpha - d'_\alpha$.

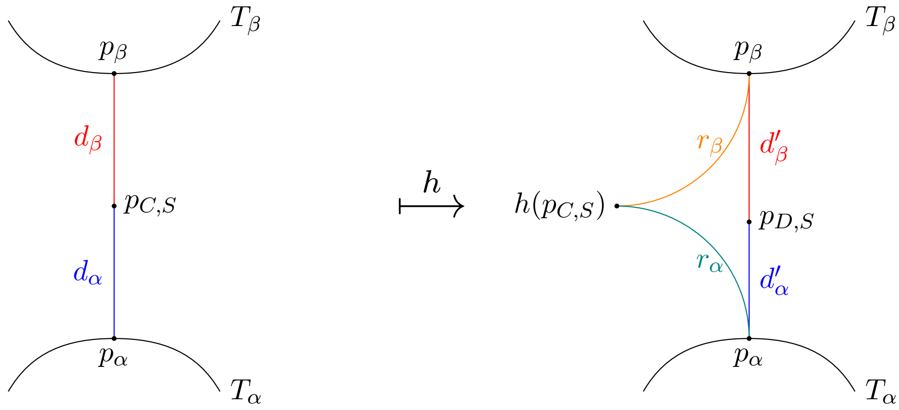

Figure 26: Image of base point under change of marking.

By Lemma 5.24 we have the inequality

$$r_\alpha \leq \Lambda_R d_\alpha + (\Lambda_R - 1) \cdot \#_T(\alpha) \leq \Lambda_R d_\alpha + (\Lambda_R - 1)N.$$

Recall the equalities $l_C(S) = l_C(\alpha) + 2d_\alpha = l_C(\beta) + 2d_\beta$ and similarly for D. Applying these to Lemma 5.7 (i) we get the following inequality on the length of S

$$\begin{aligned}
2(l_C(S) - l_D(S)) &= \big(l_C(\alpha) - l_D(\alpha) + 2(d_\alpha - d'_\alpha)\big) + \big(l_C(\beta) - l_D(\beta) + 2(d_\beta - d'_\beta)\big) \\
&= (l_C(\alpha) - l_D(\alpha)) + (l_C(\beta) - l_D(\beta)) + 2\big(d_\alpha + d_\beta - (d'_\alpha + d'_\beta)\big) \\
&= (l_C(\alpha) - l_D(\alpha)) + (l_C(\beta) - l_D(\beta)) + 2\big(d_C(T_\alpha, T_\beta) - d_D(T_\alpha, T_\beta)\big) \\
&\leq (\Lambda_R - 1)N + (\Lambda_R - 1)N + 2(\Lambda_R - 1)M \\
&= 2(\Lambda_R - 1)(N + M).
\end{aligned}$$

Again using the equalities $l_C(S) = l_C(\alpha) + 2d_\alpha$ and $l_D(S) = l_D(\alpha) + 2d'_\alpha$, we can now bound the difference of d_α and d'_α by the above inequality and Lemma 5.7 (ii):

$$\begin{aligned}
2(d_\alpha - d'_\alpha) &= (l_C(S) - l_D(S)) - (l_C(\alpha) - l_D(\alpha)) \\
&\leq (\Lambda_R - 1)(N + M) + (\Lambda_R - 1)(3n - 4)N \\
&= (\Lambda_R - 1)((3n - 3)N + M)
\end{aligned}$$

Using the two above inequalities for r_α and $d_\alpha - d'_\alpha$ we get:

$$\begin{aligned}
d_D(p_{D,S}, h(p_{C,S})) &= r_\alpha - d'_\alpha \\
&\leq \Lambda_R d_\alpha + (\Lambda_R - 1)N - d_\alpha + (d_\alpha - d'_\alpha) \\
&\leq (\Lambda_R - 1)\left(d_\alpha + N + \frac{3n-3}{2}N + \frac{1}{2}M\right) \\
&\leq (\Lambda_R - 1)\left(\frac{3}{2}M + \frac{3n-1}{2}N\right)
\end{aligned}$$

We used here in the last inequality that each edge has at most length 1 and hence we have $d_\alpha \leq d_C(T_\alpha, T_\beta) \leq M$.

In **case (ii)** let $\alpha \in S$ be such that $l_{\gamma(t)}(S) = l_{\gamma(t)}(\alpha)$ for all $t \in [t_{i-1}, t_i]$. As in Figure 24 let $p_{\gamma(t),L} \in \gamma(t)_S$ be the "leftmost" point of the characteristic set, that is we have a $\beta_L \in S$ with $d_{\gamma(t)}(p_{\gamma(t),L}, \beta_L \cdot p_{\gamma(t),L}) = l_{\gamma(t)}(S)$ and $d_{\gamma(t)}(p, \beta_L \cdot p) > l_{\gamma(t)}(S)$ for any point $p \in T_\alpha$ further left along T_α. Similarly we denote $p_{\gamma(t),R}$ and β_R and again up to further subdividing the interval we assume that β_L and β_R are the same for all $t \in [t_{i-1}, t_i]$.

Let $u := d_C(p_{C,L}, p_{C,R})$ and $u' := d_D(p_{D,L}, p_{D,R})$ be the diameter of the characteristic sets C_S and D_S, respectively. We now want to bound the difference $u - u'$ by some inequality of the form $u - u' \leq (\Lambda_R - 1)\kappa_u$. If T_α and T_{β_L} do not intersect, let $q_L \in T_\alpha$ be the vertex in T_α closest to T_β, otherwise let $q_L \in T_\alpha \cap T_{\beta_L}$ be the leftmost vertex on the intersection, i.e. we have $d_C(p_{C,L}, T_{\beta_L}) = d_C(p_{C,L}, q_L) + d_C(q_L, T_{\beta_L})$. Similarly we have a vertex q_R. Since $l_C(S) = d_C(p_{C,L}, \beta_L \cdot p_{C,L}) \geq d_C(p, \beta_L \cdot p)$ for any point $p \in C_S$ and T_{β_L} is an embedded line in a tree, we have that q_L lies to the right of $p_{C,L}$ and similarly to the right of $p_{D,L}$. Similarly q_R lies to the left of $p_{C,R}$ and $p_{D,R}$. As all four points lie on the embedded line T_α we have the following two possible equalities for u: If q_L lies left of q_R we get the equality $u = d_C(p_{C,L}, q_L) + d_C(q_L, q_R) + d_C(q_R, p_{C,R})$ and if q_L lies right of q_R we get the equality $u = d_C(p_{C,L}, q_L) + d_C(q_R, p_{C,R}) - d_C(q_L, q_R)$. Hence we conclude:

$$\begin{aligned} u - u' = \ & \big(d_C(p_{C,L}, q_L) - d_D(p_{D,L}, q_L)\big) \\ & \pm\big(d_C(q_L, q_R) - d_D(q_L, q_R)\big) \\ & +\big(d_C(p_{C,R}, q_R) - d_D(p_{D,R}, q_R)\big) \end{aligned}$$

By definition of β_L we have

$$\begin{aligned} l_C(S) &= 2d_C(p_{C,L}, T_{\beta_L}) + l_C(\beta_L) \\ &= 2(d_C(p_{C,L}, q_L) + d_C(T_\alpha, T_{\beta,L}) + l_C(\beta_L) \\ \Rightarrow\ 2d_C(p_{C,L}, q_L) &= l_C(S) - 2d_C(T_\alpha, T_{\beta_L}) - l_C(\beta_L) \end{aligned}$$

and thus by Lemma 5.7:

$$\begin{aligned} 2\big(d_C(p_{C,L}, q_L) - d_D(p_{D,L}, q_L)\big) &= \big(l_C(S) - l_D(S)\big) + 2\big(d_D(T_\alpha, T_{\beta_L}) - d_C(T_\alpha, T_{\beta_L})\big) \\ &\quad + (l_D(\beta_L) - l_C(\beta_L)) \\ &\leq (\Lambda_R - 1)\big(\#_T(\alpha) + (3n-4)M + (3n-4)\#_T(\beta_L)\big) \\ &\leq (\Lambda_R - 1)\big(N + (3n-4)(M+N)\big) \end{aligned}$$

Recall that for any $\beta \in S$ and $p \in T$ we have the equality $d_{\gamma(t)}(p, \beta \cdot p) = 2d_{\gamma(t)}(p, T_\beta) + l_{\gamma(t)}(\beta)$. That means we have $\gamma(t)_S \subseteq \overline{B_r(T_\beta)} \cap T_\alpha$ for any $r \geq l_{\gamma(t)}(\alpha)$. In particular we have for each $\gamma(t)$ that its characteristic set is contained in the same subsegment ρ of T_α with at most $L + 2N$ edges as depicted in Figure 27, namely at most L edges from $T_\alpha \cap T_\beta$ and at most $N \geq \#_T(\alpha)$ edges on each side from the $l(\alpha)$-neighbourhood in T_α.

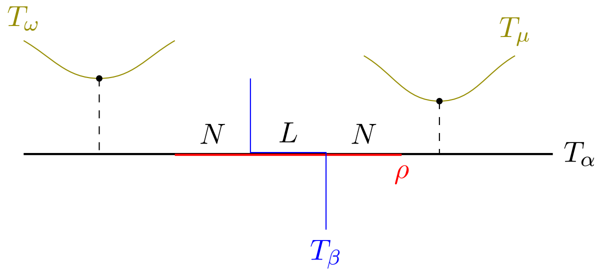

Figure 27: A subsegment ρ of T_α containing the characteristic set.

That is we have by Lemma 5.7 $|d_C(q_L,q_R) - d_D(q_L,q_R)| \le (\Lambda_R - 1)(3n-4)(2N+L)$ and hence we get as a bound for $u - u'$:

$$\begin{aligned} u - u' &\le (\Lambda_R - 1)\big(N + (3n-4)(M+N)\big) + (\Lambda_R - 1)(3n-4)(2N+L) \\ &= (\Lambda_R - 1)\underbrace{\big(N + (3n-4)(M+L+3N)\big)}_{=:\kappa_u} = (\Lambda_R - 1)\kappa_u \end{aligned}$$

To bound how far a point in ρ moves we construct two elements $\omega, \mu \in F_n$ such that – similar to the situation we already had in case (i) – their two translation axes T_ω and T_μ lie in the two complements of ρ as depicted in Figure 27.

Let e be the leftmost edge of ρ and choose an infinitely long path ζ in $T \setminus (F_n \cdot e)$ starting at the leftmost vertex of ρ which does not cross any edge in the orbit of e and is locally injective. Let $v_1, v_2, \ldots \in V(T)$ be the vertices of ζ and define $\omega \in F_n$ by the first time when two such vertices are in the same F_n orbit, i.e. we have $\omega \cdot v_j = v_{j+l}$. As T has at most $2n-2$ different vertex orbits, we have $j + l \le 2n-1$. Recall that we have $d_T(v_i, \omega v_i) = l_T(\omega) + 2d_T(v_i, T_\omega)$ and hence as in Figure 23 we have that the edge count of ω is at most $\#_T(\omega) \le 2n-1$. Furthermore, we have because of $j + l \le 2n-1$ that the edge distance from T_ω to ρ is at most $2n-2$. Moreover, each edge of T_ω lies in the orbit of an edge in ζ. As ζ does not contain an orbit of the first edge of ρ, we have that e is not an edge in T_ω and thus T_ω intersects ρ at most at the starting vertex of ζ. Similarly we construct $\mu \in F_n$ for the right-hand side of ρ.

As the base point $p_{C,S}$ lies in the middle of the characteristic set C_S, we have for the distance $d_C(T_\omega, p_{C,S}) = d_C(T_\omega, p_{C_L}) + \frac{1}{2}u \le 2n - 2 + N + L$ and accordingly for D. We will use this to get a $\kappa > 0$ as in case (i).

As T_ω lies left of the characteristic set, each path from T_ω to T_{β_L} has to pass through $p_{C,L}$. Hence we have the equality $d_C(T_\omega, T_{\beta_L}) = d_C(T_\omega, p_{C,L}) + d_C(p_{C,L}, T_{\beta_L})$ and similarly for D.

Observe that the shortest path from T_{β_L} to T_ω contains at most $2n - 2 + N + M$ edges and we have $2d_C(p_{C,L}, T_{\beta_L}) = l_C(S) - l_C(\beta_L)$. Hence, we get by Lemma 5.7 the following inequalities:

$$\begin{aligned} d_C(T_\omega, p_{C,L}) - d_D(T_\omega, p_{D,L}) &= \big(d_C(T_\omega, T_{\beta_L}) - d_C(p_{C,L}, T_{\beta_L})\big) \\ &\quad - \big(d_D(T_\omega, T_{\beta_L}) - d_D(p_{D,L}, T_{\beta_L})\big) \\ &= \big(d_C(T_\omega, T_{\beta_L}) - d_D(T_\omega, T_{\beta_L})\big) \\ &\quad + \frac{1}{2}\big((l_C(S) - l_D(S)) + (l_D(\beta_L) - l_C(\beta_L))\big) \\ &\leq (\Lambda_R - 1)(2n - 2 + N + M) + (\Lambda_R - 1)\frac{3n-3}{2}N \\ &= (\Lambda_R - 1)\kappa_L, \end{aligned}$$

with $\kappa_L := 2n - 2 + N + M + \frac{3n-3}{2}N$. In particular we have the inequality

$$\begin{aligned} d_C(T_\omega, p_{C,S}) - d_D(T_\omega, p_{D,S}) &= d_C(T_\omega, p_{C,L}) + \frac{1}{2}u - \big(d_D(T_\omega, p_{D,L}) + \frac{1}{2}u'\big) \\ &\leq (\Lambda_R - 1)\big(\kappa_L + \frac{1}{2}\kappa_u\big). \end{aligned}$$

Accordingly we get $d_C(T_\mu, p_{C,S}) - d_D(T_\mu, p_{D,S}) \leq (\Lambda_R - 1)\big(\kappa_L + \frac{1}{2}\kappa_u\big)$. We are now in a similar situation as in case (i), where T_ω and T_μ play the role of T_α and T_β, namely we fixate $p_{C,S}$ and $p_{D,S}$ by their distance to T_ω and T_μ. By Lemma 5.24 we have the inequality

$$d_D(T_\omega, h(p_{C,S})) \leq \Lambda_R d_C(T_\omega, p_{C,S}) + (\Lambda_R - 1)(2n - 1).$$

Up to renaming T_ω and T_μ we can again assume that $p_{D,S}$ lies on the path from T_ω to $h(p_{C,S})$, hence we have

$$\begin{aligned} d_D(p_{D,S}, h(p_{C,S})) &= d_D(T_\omega, h(p_{C,S})) - d_D(T_\omega, p_{D,S}) \\ &\leq \Lambda_R d_C(T_\omega, p_{C,S}) + (\Lambda_R - 1)(2n - 1) - d_D(T_\omega, p_{D,S}) \\ &\leq (\Lambda_R - 1)\big(2n - 1 + \kappa_L + \frac{1}{2}\kappa_u + d_C(T_\omega, p_{C,S})\big) \\ &\leq (\Lambda_R - 1)\big(2n - 1 + \kappa_L + \frac{1}{2}\kappa_u + 2n - 2 + N + L\big). \end{aligned}$$

Now we can combine the results for the cases (i) and (ii) to get an upper bound

$$d_D(p_{D,S}, h(p_{C,S})) \leq (\Lambda_R - 1) \cdot \underbrace{\kappa(n, N, L, M)}_{=:\kappa(\Delta)}$$

where the numbers N, L, M only depend on Δ. By Lemma 5.23 we then have that $d_R(\phi_{S,K}(C), \phi_{S,K}(D)) = d_R(C, D)$ for all $K \geq \frac{\kappa(\Delta)}{2}$. □

We also get a version of Proposition 5.22 for bounded subsets by taking the maximum the corresponding $\kappa(\Delta)$.

Corollary 5.25
Let $U \subset CV_n$ be a bounded subset, then there exists a $K > 0$ such that $\phi_{S,K}|_U$ is an isometry.

Proof. Let $U \subset CV_n$ be a (non-empty) bounded subset, $r := \operatorname{diam}(U)$ its diameter and $D \in U$ any point. Let $A, B \in U$ be two points, then the envelope $\operatorname{Env}_R(A, B)$ is contained in the $2r$-ingoing neighbourhood $B_{2r}^{\text{in}}(D)$ of D, since we have for every $C \in \operatorname{Env}_R(A, B)$ by the triangle inequality:

$$d_R(C, D) \leq d_R(C, B) + d_R(B, D) \leq d_R(A, B) + d_R(B, D) \leq 2 \operatorname{diam}(U) = 2r.$$

By Lemma 2.6 the ingoing ball $B_{2r}^{\text{in}}(D)$ intersects finitely many simplices $\Delta \subset CV_n$, hence we can choose K to be the maximum of the $\kappa(\Delta)$ from Proposition 5.22 over all these Δ. As any geodesic from A to B is now piecewise isometrically send to a (piecewise) geodesic, we have that $\phi_{S,K}$ at most decreases the distance from A to B. But as $\phi_{S,K}$ also preserves the stretching factor of candidates in A, it also at most increases the distance from A to B. Hence, $\phi_{S,K}|_U$ is an isometric embedding. □

While for any given simplex or bounded subset $U \subset CV_n$ we can find a large enough K such that $\phi_{S,K}|_U$ is an isometric embedding, it is in general not a global isometric embedding from CV_n to CV_{n+k} as we can see in the following example.

Proposition 5.26
Let $n \geq 3$, $S = \{s_1, \dots, s_n\} \subset F_n$ be a generating set and $K > 0$. Then the map $\phi_{S,K} : CV_n \to CV_{n+1}$ is not an isometric embedding.

Proof. To construct a counterexample, consider for some $k \in \mathbb{N}$ the set of generators $(\alpha_1, \dots, \alpha_n)$ with $\alpha_i := s_i$ for $i > 2$, $\alpha_2 := s_3^{-k} s_2 s_3^k$ and $\alpha_1 := \alpha_2^{-1} s_1 \alpha_2$, that is we have $s_2 = \alpha_3^k \alpha_2 \alpha_3^{-k}$, $s_1 = \alpha_2 \alpha_1 \alpha_2^{-1}$ and $s_i = \alpha_i$ for $i > 2$. Let R be the Cayley-graph of F_n with respect to the basis $(\alpha_1, \dots, \alpha_n)$:

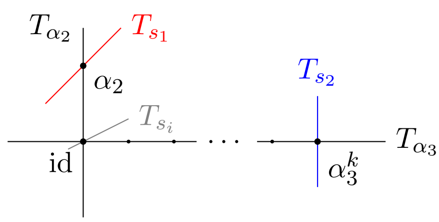

Figure 28: A Cayley-graph of F_n with two shifted translation axes.

For k large enough and $0 < \varepsilon < \frac{1}{k}$ let $A = (R, l_A, m), B = (R, l_B, m) \in CV_n$ be the two points where m denotes the standard left-multiplication on the Cayley-graph together with the following length functions:

$$l_A(\alpha_i) = \begin{cases} \frac{1}{k}, & \text{if } i = 3, \\ 1, & \text{if } i \neq 3, \end{cases}$$

$$l_B(\alpha_i) = \begin{cases} 1 + \varepsilon, & \text{if } i = 1, \\ \frac{1}{k} - \varepsilon, & \text{if } i = 3, \\ 1, & \text{if } i \neq 3. \end{cases}$$

We have $\mathrm{vol}(A) = \mathrm{vol}(B) = n - \frac{k-1}{k}$. Note that R can also be seen as the rose with $\alpha_1, \dots, \alpha_n$ as marking on the petals and α_1 is the only stretched loop from A to B. Hence, we have by Corollary 1.15

$$d_R(A, B) = \log(\Lambda_R(A, B)) = \log\left(\frac{l_B(\alpha_1)}{l_A(\alpha_1)}\right) = \log(1 + \varepsilon).$$

We will show that for large enough k we have $\Lambda_R(\phi_{S,K}(A), \phi_{S,K}(B)) \geq \Lambda_R(A, B)$.

To do so we first calculate their characteristic sets. We assume that $p_{A,S}$ and $p_{B,S}$ lie on the path between the translation axes T_{s_1} and T_{s_2} and show in afterwards that it is indeed the base point. Let $S' := \{s_1, s_2\}$. As T_{s_1} and T_{s_2} are shifts of T_{α_1} and T_{α_2} by α_2 and α_3^k, respectively, we get the following identities:

$$\begin{aligned} & l_A(S') = l_A(s_1) + 2d_A(T_{s_1}, p_{A,S'}) = l_A(s_2) + 2d_A(T_{s_2}, p_{A,S'}) \\ \Rightarrow \quad & 2l_A(S') = l_A(s_1) + l_A(s_2) + 2d_A(T_{s_1}, T_{s_2}) = 2 + 2(1 + k \cdot \frac{1}{k}) = 6 \\ & d_A(T_{s_1}, p_{A,S'}) = \frac{l_A(S') - l_A(s_1)}{2} = \frac{3-1}{2} = 1 \\ \Rightarrow \quad & p_{A,S'} = \mathrm{id} \end{aligned}$$

where id denotes the vertex corresponding to the neutral element of F_n. As for each $s_i \in S$ we have $d_A(p_{A,S'}, s_i \cdot p_{A,S'}) \leq 3$, we get $l_A(S) \leq 3$. On the other hand for any point $p \in R \setminus \{p_{A,S'}\}$ we have $d_A(p, s_1 \cdot p) > 3$ or $d_A(p, s_2 \cdot p) > 3$ and hence we have that $p_{A,S} = p_{A,S'} = \mathrm{id}$ is indeed our base point. Similarly we get for B:

$$\begin{aligned} 2l_B(S') &= l_B(s_1) + l_B(s_2) + 2d_B(T_{s_1}, T_{s_2}) \\ &= 1 + \varepsilon + 1 + 2(1 + k(1/k - \varepsilon)) = 6 - (2k-1)\varepsilon \end{aligned}$$

$$\begin{aligned} d_B(T_{s_1}, p_{B,S'}) &= \frac{l_B(S') - l_B(s_1)}{2} \\ &= \frac{3 - (k - \frac{1}{2})\varepsilon - (1 + \varepsilon)}{2} \\ &= 1 - \frac{2k+1}{4}\varepsilon. \end{aligned}$$

Since we have $d_B(T_{s_1}, p_{B,S'}) < d_B(T_{s_i}, \mathrm{id})$, the point $p_{B,S'}$ lies on T_{α_2} and we have

$$d_B(\mathrm{id}, p_{B,S'}) = d_B(T_{s_1}, \mathrm{id}) - d_B(T_{s_1}, p_{B,S'}) = \frac{2k+1}{4}\varepsilon$$

If we choose ε small enough such that $\frac{2k+1}{4}\varepsilon < \frac{1}{2}$ holds, we have as before that $p_{B,S'}$ is less displaced by any other s_i. In particular, we have $p_{B,S} = p_{B,S'}$ is indeed the base point for B.

We apply now the map $\phi_{S,K}$ and calculate the stretching factor of the newly introduced candidate $\alpha_3\alpha_{n+1}$:

$$\begin{aligned} l_{\phi_{S,K}(B)}(\alpha_3\alpha_{n+1}) &= l_B(\alpha_3) + 2d_B(\mathrm{id}, p_{B,S}) + K \\ &= \frac{1}{k} - \varepsilon + \frac{2k+1}{2}\varepsilon + K \end{aligned}$$

$$\begin{aligned} \Rightarrow \quad \Lambda_R(\phi_{S,K}(A), \phi_{S,K}(B)) &\geq \frac{l_{\phi_{S,K}(B)}(\alpha_3\alpha_{n+1})}{l_{\phi_{S,K}(A)}(\alpha_3\alpha_{n+1})} \\ &= \frac{\frac{1}{k} - \varepsilon + \frac{2k+1}{2}\varepsilon + K}{\frac{1}{k} + K} \\ &= 1 + \frac{2k-1}{2(\frac{1}{k} + K)} \cdot \varepsilon \\ &> 1 + \varepsilon = \Lambda_R(A, B) \end{aligned}$$

where we get the last inequality by choosing k large enough such that $2k - 1 > 2(\frac{1}{k} + K)$ holds. Since $\phi_{S,K}(A)$ and $\phi_{S,K}(B)$ have again the same volume, we have $d_R(A, B) < d_R(\phi_{S,K}(A), \phi_{S,K}(B))$, thus $\phi_{S,K}$ is not an isometric embedding. □

5.3 Embeddings via coverings

Another more natural embedding from CV_n to CV_k comes from finite coverings. Recall that by universal covering theory the finite covers of a finite, connected graph Γ correspond to finite index subgroups of the fundamental group $\pi_1(\Gamma)$. That is after fixing a finite index subgroup $F_k \cong G$ in F_n we get for every $A \in CV_n$ its corresponding finite cover in CV_k. In terms of marked, metric trees as in Definition 1.3 this means that a free, minimal F_n-action on a tree T is also a free, minimal G-action on T.

Lemma 5.27
Let $G \leq F_n$ be a finite index subgroup, T a simplicial, metric tree and $m\colon F_n \to \mathrm{Isom}(T)$ a free, minimal action by isometries. Then its restriction $m|_G\colon G \to \mathrm{Isom}(T)$ is a free, minimal action.

Proof. That $m|_G$ acts freely is clear. Since F_n acts minimally on T, we have that T is covered by the translation axes of F_n, that is

$$T = \bigcup_{\alpha \in F_n \setminus \{\mathrm{id}\}} T_\alpha,$$

as we can extend any finite subsegment of T to be a part of a translation axis. Recall that taking powers of α does not change the translation axis, i.e. we have $T_\alpha = T_{\alpha^r}$ for any $r \in \mathbb{N}$. Since G has finite index in F_n, there exists for each $\alpha \in F_n$ a number $r \in \mathbb{N}$ with $\alpha^r \in G$ and thus we also have

$$T = \bigcup_{\alpha \in G \setminus \{\mathrm{id}\}} T_\alpha.$$

An α-invariant subtree has to contain all translation axes T_α, thus T has no non-trivial G-invariant subtrees, that is the action $m|_G$ is minimal. □

This means, given a finite index subgroup $G \leq F_n$ and identifying it with F_k, gives us an embedding from CV_n to CV_k. Since the distance can be calculated as the maximal quotient of the translation lengths, we get the following natural isometric embedding.

Proposition 5.28
Let $\phi\colon F_k \to F_n$ be an injective group homomorphism such that $\phi(F_k)$ has finite index in F_n. Then the map

$$\phi^*\colon CV_n \to CV_k \quad , \quad (T, l, m) \mapsto (T, l, m \circ \phi)$$

is an isometric embedding.

Proof. By Lemma 5.27 we have that $m \circ \phi$ is a free, minimal action of F_k on T so ϕ^* is well defined. Let $A, B \in CV_n$, then $\Lambda_R(A, B) = \sup_{\alpha \in F_n} \frac{l_B(\alpha)}{l_A(\alpha)}$. Since $l_A(\alpha)$ is the translation length of α along its translation axis T_α, we have $l_A(\alpha^m) = m \cdot l_A(\alpha)$ for all $m \in \mathbb{N}$. Again $\phi(F_k)$ is a finite index subgroup of F_n and so there exists for each $\alpha \in F_n$ an $m \in \mathbb{N}$ with $\alpha^m \in \phi(F_k)$. Hence, we have:

$$\begin{aligned}\Lambda_R(A, B) &= \sup_{\alpha \in F_n} \frac{l_B(\alpha)}{l_A(\alpha)} = \sup_{\alpha \in F_n} \frac{m l_B(\alpha)}{m l_A(\alpha)} \\ &= \sup_{\alpha \in F_n} \frac{l_B(\alpha^m)}{l_A(\alpha^m)} = \sup_{\beta \in \phi(F_k)} \frac{l_B(\beta)}{l_A(\beta)} \\ &= \Lambda_R(\phi^*(A), \phi^*(B)).\end{aligned}$$

As $\phi^*(A)$ corresponds to a finite cover of A, we have $\operatorname{vol}(\phi^*(A)) = [F_n : \phi(F_k)] \operatorname{vol}(A)$ and similarly $\operatorname{vol}(\phi^*(B)) = [F_n : \phi(F_k)] \operatorname{vol}(B)$ and thus $d_R(A, B) = d_R(\phi^*(A), \phi^*(B))$. □

Sometimes it is more convenient to write down the isometric embedding from Proposition 5.28 in terms of finite coverings $T/(m \circ \phi) \twoheadrightarrow T/m$ of finite marked graphs.

Remark 5.29
Let $\phi\colon F_k \to F_n$ and $\phi^*\colon CV_n \to CV_k$ be as in Proposition 5.28 and $A = (\Gamma, l, m) \in CV_n$. Consider a representant of A as in Notation 1.2(i), that is a finite, metric graph Γ with labelled (oriented) edges as in Figure 2, where we label the edges of the spanning tree with $\mathrm{id} \in F_n$. As slight abuse of notation we will also denote the corresponding labelling of an edge $e \in E(\Gamma)$ as $m(e) \in F_n$.

Recall from the theory of coverings that we can describe the covering Υ of Γ which corresponds to $\phi^*(A)$ the following way: Let $R := \phi(F_k)\backslash F_n$ be the right cosets of $\phi(F_k)$. Then vertices and edges of Υ are R-multiples of vertices and edges of Γ, respectively:

$$\begin{aligned}V(\Upsilon) &:= \{(r, v) \mid r \in R, v \in V(\Gamma)\} \\ E(\Upsilon) &:= \{(r, e) \text{ is an edge from } (r, v_1) \text{ to } (r \cdot m(e), v_2) \\ &\qquad\qquad \mid r \in R, e \in E(\Gamma) \text{ is an edge from } v_1 \text{ to } v_2\}\end{aligned}$$

We set the same labels and lengths on edges in Υ as for Γ, i.e. $m_\Upsilon((r,e)) := m(e)$ and $l_\Upsilon((r,e)) := l(e)$.

Observe that in Remark 5.29 the edges of Υ with trivial labels under m_Υ do not form a spanning tree. In order to write down the marking as in Notation 1.2(i) we need to do the same as if we change the spanning tree: Fix a vertex $(1,v) \in V(\Upsilon)$ and a spanning tree $T \subset \Upsilon$ and label all edges in T by the neutral element of F_k. Let $(r,e) \in E(\Upsilon \setminus T)$ be an oriented edge outside of T, then there exists a unique reduced, closed path $\rho = (r_1,e_1)\dots(r_s,e_s) \subset T \cup \{e\}$ starting and ending at $(1,v)$ and crossing e exactly once. We label then (r,e) as the product $\phi^{-1}(m_\Upsilon(r_1,e_1)\cdot\ldots\cdot m_\Upsilon(r_s,e_s))$ of the original edge labels along ρ. Note here that the product of the labels is indeed in $\phi(F_k)$ since we have chosen a loop starting and ending at $(1,v)$.

Example 5.30
Let F_2 be the free group generated by α and β and $U := \langle \alpha, \beta\alpha\beta, \beta^2 \rangle$ one of its finite index subgroups isomorphic to F_3. Let $\phi^* : CV_2 \to CV_3$ be the corresponding isometric embedding as in Proposition 5.28. Then we have:

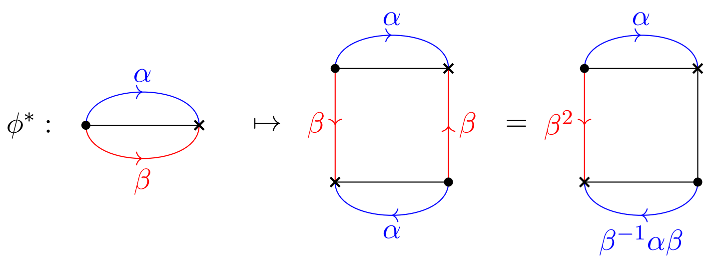

To distinguish for which $k \in \mathbb{N}$ we have such an embedding, recall that the free rank of a finite index subgroup of a free group is given by the Nielsen-Schreier formula (see [Ser80, Section 3.4]).

Theorem 5.31 (Nielsen–Schreier)
Let F_n be a free group of rank n and $U \leq F_n$ a subgroup with finite index $[F_n : U] = d < \infty$. Then U is a free group of rank $1 + d(n-1)$.

As there exists for any $d \in \mathbb{N}$ a degree d covering of a finite graph, there exists such an isometric embedding $\phi^* \colon CV_n \to CV_k$, if and only if $k = 1 + d(n-1)$ holds for some $d \in \mathbb{N}$. In particular we can embed CV_2 into any CV_k for $k \geq 2$.

6 Deformations of embeddings

As we have seen in Example 5.6 and Corollary 5.15 we can slightly deform naive embeddings from Section 5.2 to get continuous families of isometric embeddings. We will see in Theorem 6.7 that we can also locally deform a natural embedding from Section 5.3 from CV_2 to CV_k to get a continuous family of isometric embeddings. In contrast we will see in Theorem 6.8 that the natural embeddings from CV_n to CV_k for $n \geq 3$ have some sort of rigidity, that is they can not be locally deformed into another isometric embedding.

Definition 6.1
Let $k, n \in \mathbb{N}$ be two natural numbers and $\phi, \psi \colon CV_n \to CV_k$ be isometric embeddings. We say ϕ and ψ are *deformations* of each other, if there exists a homotopy of isometries $H \colon CV_n \times [0,1] \to CV_k$ between them. That means H is continuous, the map $H(\,\cdot\,, t)$ is for all $0 \leq t \leq 1$ an isometry and we have $H(\,\cdot\,, 0) = \phi$ and $H(\,\cdot\,, 1) = \psi$.

We say ϕ and ψ are *bounded deformations* of each other, if the homotopy additionally satisfies that the difference set $\{A \in CV_n \mid H(A, \cdot\,) \text{ is not constant}\}$ is a bounded subset of CV_n.

Observe slight varying of the attached graph in Section 5.2 yields deformations of isometric embeddings in the above sense.

The following Lemma 6.2 can be used to construct explicit bounded deformations of isometric embeddings. Namely, if we deform an isometric embedding in the interior of a bounded area such that its restriction to the interior is isometric, then the deformed embedding is globally still isometric. The idea for this is to skewer interior points with a geodesic ray which enters and leaves the deformed area.

Lemma 6.2
Let $U \subseteq CV_n$ be a subset of finite diameter and $U^c := CV_n \setminus U$ be its complement. Furthermore, let $\psi \colon CV_n \to CV_k$ be a map such that its restrictions to the closures $\psi|_{\overline{U^c}}$ and $\psi|_{\overline{U}}$ are isometric embeddings, then ψ is an isometric embedding.

Proof. We will show that for any $A \in U^c$ and $B \in U$ we have $d_R(A, B) = d_R(\psi(A), \psi(B))$. Let $\gamma \colon [0,1] \to CV_n$ be a geodesic from A to B. By Lemma 1.36 we can continue γ to a geodesic ray $\gamma \colon \mathbb{R}_{\geq 0} \to CV_n$ leaving U. Let $s < 1 < t$ be two times when γ passes the boundary of U.

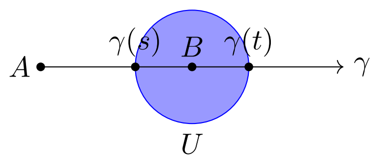

In particular, we have $\gamma(s), \gamma(t) \in \overline{U^c} \cap \overline{U}$ and hence ψ preservers their distances to A

and B. By the triangle inequality and Lemma 1.18 we have then the following inequalities:

$$\begin{aligned} d_R(\psi(A),\psi(B)) &\leq d_R(\psi(A),\psi(\gamma(s))) + d_R(\psi(\gamma(s)),\psi(B)) \\ &= d_R(A,\gamma(s)) + d_R(\gamma(s),B) = d_R(A,B), \\ d_R(\psi(A),\psi(B)) &\geq d_R(\psi(A),\psi(\gamma(t))) - d_R(\psi(B),\psi(\gamma(t))) \\ &= d_R(A,\gamma(t)) - d_R(B,\gamma(t)) = d_R(A,B) \end{aligned}$$

Hence, equality holds. Analogously, we have $d_R(\psi(B),\psi(A)) = d_R(B,A)$ and thus ψ is an isometric embedding. □

Lemma 6.2 also holds for more general cases. Namely, when X is a geodesic space and $U \subset X$ a subset such that each geodesic starting outside and ending inside U can be continued to a geodesic ending outside U.

Remark 6.3 (i) The condition in Lemma 6.2 that the area U is bounded, is essential, otherwise the geodesic may not leave U and we could, for example, flip U onto its complement. As concrete example of such a flip let $\Delta \in CV_2$ be a simplex corresponding to the standard figure of eight and $U \subset CV_2$ be a component of $CV_2 \setminus \Delta$. Then the automorphism $\phi\colon F_2 := \langle \alpha, \beta \rangle \to F_2$ with $\phi(\alpha) = \alpha^{-1}$ and $\phi(\beta) = \beta$ yields an isometric embedding $\phi^*\colon U \to CV_2$ with $\phi^*|_\Delta = \mathrm{id}_\Delta$. But if U contains a theta-graph the map

$$\psi\colon CV_2 \to CV_2 \quad , \quad A \mapsto \begin{cases} A, & A \notin U \\ \phi^*(A), & A \in U \end{cases}$$

is not injective as we have $\phi^*(U) \subset CV_2 \setminus U$.

(ii) To deform a bounded set by Lemma 6.2 it is actually enough that for all $A \in \partial U, B \in \overline{U}$ we have $d_R(A,B) = d_R(\psi(A),\psi(B))$ and $d_R(B,A) = d_R(\psi(B),\psi(A))$. From this follows similarly to the proof of Lemma 6.2 that $\psi|_{\overline{U}}$ is an isometric embedding. We use here that we can extend a geodesic in both directions by Corollary 2.19 and Lemma 1.36.

We will see in the following example, that we can also locally deform natural embeddings from CV_2 to CV_3.

Example 6.4
Let ϕ^* be the isometric embedding from Example 5.30 and $0 < \varepsilon, \delta < \frac{1}{3}$ be some small numbers. Let $U \subset CV_2$ be the theta-graphs with some fixed marking such that after normalising each edge has at least length δ. For $A \in U$ we set

$$\varepsilon(A) := \varepsilon \cdot (\min\{l_A(e) \mid e \in E(A)\} - \delta)$$

to be the desired deformation depending on the minimal edge length of A.

Then we define $\psi\colon CV_2 \to CV_3$ by $\psi(A) := \phi^*(A)$ for $A \notin U$ and on U as follows:

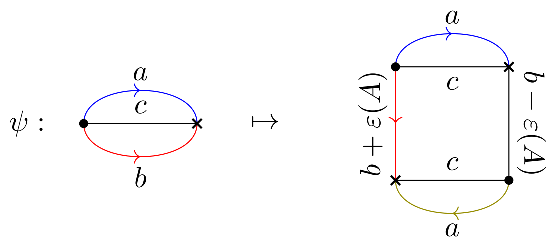

with marking α, β and ω. For $A \in \partial U$ we have $\varepsilon(A) = 0$, hence ψ is also equal to ϕ^* on the boundary of U and in particular restricted to the closure of the complement of U an isometric embedding. Since U is contained in the 2δ-thick part of a simplex it has finite diameter.

We now want to apply Lemma 6.2. To do this, it is left to show that $\psi|_U$ is an isometric embedding. There are up to orientation only two candidates of $\psi(A)$ which have a different length than in $\phi^*(A)$, namely the barbells $\mu^+ := \alpha\beta\omega\beta^{-1}$ and $\mu^- := \alpha\omega$. We will show that for $A, B \in U$ they are never maximally stretched from $\psi(A)$ to $\psi(B)$, that means we have $d_R(\psi(A), \psi(B)) = d_R(A, B)$. Let A, B be normalised and (a, c, b) and (a', c', b') be their edge lengths, respectively. We then have:

$$\frac{l_{\psi(B)}(\mu^\pm)}{l_{\psi(A)}(\mu^\pm)} = \frac{2 \cdot (a' + b' + c' \pm \varepsilon(B))}{2 \cdot (a + b + c \pm \varepsilon(A))} = \frac{1 \pm \varepsilon(B)}{1 \pm \varepsilon(A)}$$

We now want to compare these quotients with the stretching factors of $\alpha, \beta, \alpha\beta^{-1}$ from A to B. For the case μ^+ we assume $a = \min\{a, b, c\}$, that is we have $\frac{1+\varepsilon(B)}{1+\varepsilon(A)} \leq \frac{1+\varepsilon(a'-\delta)}{1+\varepsilon(a-\delta)}$. Likewise, for the case μ^- we assume $a' = \min\{a', b', c'\}$ and get $\frac{1-\varepsilon(B)}{1-\varepsilon(A)} \leq \frac{1-\varepsilon(a'-\delta)}{1-\varepsilon(a-\delta)}$. It follows:

$$\begin{aligned}
\frac{1 \pm \varepsilon(B)}{1 \pm \varepsilon(A)} &\leq \frac{1 \pm \varepsilon(a' - \delta)}{1 \pm \varepsilon(a - \delta)} = \frac{(1 \mp \delta\varepsilon - \varepsilon) + \varepsilon(1 \pm a')}{(1 \mp \delta\varepsilon - \varepsilon) + \varepsilon(1 \pm a)} \overset{1.14}{\leq} \max\left\{1, \frac{1 \pm a'}{1 \pm a}\right\} \\
&= \max\left\{1, \frac{a' + b' + c' \pm a'}{a + b + c \pm a}\right\} \overset{1.14}{\leq} \max\left\{\frac{a' + b'}{a + b}, \frac{a' + c'}{a + c}, \frac{b' + c'}{b + c}\right\} \\
&= \Lambda_R(A, B)
\end{aligned}$$

where we used in the last inequality that at least one candidate from B is maximally stretched, that is at least one of the stretching factors of the candidates in the last term is greater or equal to 1.

So we have that $\psi|_U$ is an isometric embedding and by Lemma 6.2 ψ is an isometric embedding.

We can even set $\delta = 0$ and deform the whole (unbounded) simplex in the same way. As in the example we then have that $\psi|_\Delta$ is an isometric embedding. Furthermore, for a theta-graph $A \in \Delta$ there exists a small $\delta' > 0$ and an $\varepsilon' > 0$ such that the image $\psi(A)$ is equal to $\psi'(A)$ with ψ' as in example 6.4 with ε' and δ'. Since we have $d_R(A, B) = d_R(\psi'(A), \psi'(B)) = d_R(\psi(A), \psi(B))$ and $d_R(B, A) = d_R(\psi(B), \psi(A))$ for $A \in \Delta, B \in CV_2 \setminus \Delta$ we have that ψ is also globally an isometric embedding.

By continuously varying either ε or δ in Example 6.4 it is clear that the example gives a (bounded) deformation of ϕ^*. Alternatively we can use the following lemma for bounded deformations.

Lemma 6.5
Let $\phi, \psi\colon CV_n \to CV_k$ be two isometric embeddings such that their difference set $U := \{A \in CV_n \mid \phi(A) \neq \psi(A)\}$ is bounded and for every $A \in U$ its images $\phi(A), \psi(A) \in CV_k$ lie in the same simplex $\Delta \subset CV_k$. Then they are bounded deformations of each other.

Proof. For $A \in CV_n$ let $\gamma_A : [0,1] \to CV_k$ be the straight line between $\phi(A)$ and $\psi(A)$ and set $H(A,t) := \gamma_A(t)$. This is well defined as $\phi(A)$ and $\psi(A)$ lie in the same simplex. Clearly we have $H(\,\cdot\,,0) = \phi$ and $H(\,\cdot\,,1) = \psi$.

We want to apply Remark 6.3 (ii) to show that $H(\,\cdot\,,t)$ is an isometric embedding. Let therefore be $B \in \partial U$ and $A \in U$. Keep in mind that for $B \in \partial U$ we have $\phi(B) = \psi(B)$ and hence we have $H(B,t) = \phi(B) = \psi(B)$ for all $t \in [0,1]$. By Lemma 1.36 there exists a geodesic passing through B and A and a point $C \in \partial U$. Let $\alpha \in W_R(\psi(B), \psi(C)) = W_R(\phi(B), \phi(C))$. By Lemma 1.22 we have $\alpha \in W_R(\psi(B), \psi(A)) \cap W_R(\phi(B), \phi(A))$. In particular, we have $\gamma_A(0), \gamma_A(1) \in \mathrm{Env}_R^{\mathrm{out}}(\psi(B), \alpha)$ and since envelopes are polytopes and γ_A is a straight line in Δ, we have $\alpha \in W_R(\psi(B), \gamma_A(t)) = W_R(H(B,t), H(A,t))$ for all $t \in [0,1]$.

As α is a witness from $\phi(B) = \psi(B)$ to $\phi(A)$ and $\psi(A)$ and ϕ and ψ are isometric embeddings, we get $l_{\phi(A)}(\alpha) = \Lambda_R(B,A) l_{\phi(B)} = l_{\psi(A)}(\alpha)$. In particular as the length of α is linear along the straight segment γ_A, we have that $l_{H(A,t)}(\alpha) = l_{\phi(A)}(\alpha)$ is constant along γ_A. Since α is a witness from $H(B,t)$ to $H(A,t)$ for all $t \in [0,1]$, we have

$$\begin{aligned} d_R(H(B,t), H(A,t)) &= \log(\Lambda_R(H(B,t), H(A,t)) \\ &= \log\left(\frac{l_{H(A,t)}(\alpha)}{l_{H(B,t)}(\alpha)}\right) \\ &= d_R(\phi(B), \phi(A)) = d_R(B,A). \end{aligned}$$

Similarly we get $d_R(H(A,t), H(B,t)) = d_R(A,B)$. □

We can see Example 6.4 as an elementary case for a family of bounded deformations. Keep in mind that as described in Remark 5.29 a natural embedding $\phi^*\colon CV_n \to CV_k$ gives a covering $\pi : \phi^*(A) \to A$.

This means that for a basepoint $p \in A$, we have the monodromy action of the fundamental group $\pi_1(A,p)$ on the fibres $L := \pi^{-1}(p) \subset \phi^*(A)$, i.e. a loop α moves a point $q \in L$ along the lift of α starting at p. We say that the monodromy action of α is transitive, if the only α-invariant subsets of L are trivial. This means there exists only one closed loop $\tilde{\alpha}$ in $\phi^*(A)$ which projects to a power of α.

Lemma 6.6
Let $\phi^*\colon CV_2 \to CV_k$ be an isometric embedding as in Proposition 5.28 and $\Delta \subset CV_2$ a simplex corresponding to a marked theta-graph (Γ, m). If Γ has a simple loop $\alpha \in F_2$ with non-transitive monodromy action, then ϕ^* has a bounded deformation ψ such that their difference set is a subset of Δ.

Proof. Let $U \subseteq \Delta$ and $0 < \varepsilon, \delta < \frac{1}{3}$ be as in Example 6.4 and $A \in U$ normalised. Let furthermore α be as in the statement.

We colour the edges of $\phi^*(A)$ depending on which edge of A they cover as usually with red, blue and black. Similarly we get a two-colouring of the vertices and since Γ has no loops we have that $\phi^*(A)$ is bipartite.

Let $\tilde{\alpha}$ be a simple loop in $\phi^*(A)$ which covers a power of α, namely it is a closed two coloured path. We may assume that $\tilde{\alpha}$ consists of alternating red and blue edges. By assumption the red-blue subgraph of $\phi^*(A)$ has at least two such simple loops and hence at least two components. In particular $\tilde{\alpha}$ does not span the whole graph.

Similar to Example 6.4 we construct a deformation ψ of ϕ^* on U. Let $\varepsilon(A) := \varepsilon \cdot (\min\{l_A(e) \mid e \in E(A)\} - \delta)$. We deform $\phi^*(A)$ to $\psi(A)$ by alternatingly adding $\pm\frac{1}{2}\varepsilon(A)$ to the lengths of the black edges adjacent to $\tilde{\alpha}$:

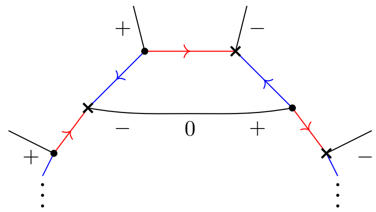

This is well defined since $\tilde{\alpha}$ has an even number of vertices as $\phi^*(A)$ is bipartite and $\tilde{\alpha}$ is a closed loop. If an edge connects two vertices of $\tilde{\alpha}$, these two vertices cover different vertices in A and we do not change the length of the edge. As $\tilde{\alpha}$ does not span the whole graph there is at least one edge non-trivially deformed.

Observe that any two-coloured loop does not change its length in $\psi(A)$ since any two-coloured path entering and leaving the subgraph spanned by $\tilde{\alpha}$ does so at two differently coloured vertices, that is it has to cross two edges which are oppositional deformed.

In the following we will subdivide any other loop in $\phi^*(A)$ into segments similar to those in Example 6.4, namely into undeformed lifts of simple loops, which are always two-coloured, and possible deformed segments which contain all three edge colours. As in Example 6.4 each of these segments is less or equally stretched between $\phi^*(A)$ and $\phi^*(B)$ than a simple loop between A and B. By Lemma 1.14 we will then have that any loop is less or equally stretched from $\psi(A)$ to $\psi(B)$ than from A to B.

It is clear that any loop contained in the subgraph spanned by either $\tilde{\alpha}$ or the vertices not covered by $\tilde{\alpha}$ is not deformed at all, hence without loss of generality let β be a cyclically reduced loop in $\phi^*(A)$ entering and leaving the subgraph spanned by $\tilde{\alpha}$.

As we might slightly rotate β without changing its length, we assume that the first edge of β is a black edge entering $\tilde{\alpha}$. We cut β at the beginning of each black edge which leaves or enters the subgraph spanned by $\tilde{\alpha}$. Namely, we have the subdivision $\beta = \beta_1 \dots \beta_{2m}$ as edge-paths such that for $1 \leq l \leq m$ all but the first vertices of β_{2l-1} lie in $\tilde{\alpha}$ and all but the first vertices of β_{2l} lie outside $\tilde{\alpha}$. We denote by v_i the deformation of each β_i:

$$v_i := l_{\psi(A)}(\beta_i) - l_{\phi^*(A)}(\beta_i) = \pm\frac{1}{2}\varepsilon(A).$$

Keep in mind that two adjacent edges have different colours. As the first edge of each β_i is black, each β_i contains at least two edges and any two adjacent edges are a lift of a simple loop in A. If $\beta_i = e_1 \dots e_{2r+1}$ contains an odd number of edges, then let $e_{2r'+1}$ be the first odd and non-black edge. Such an edge exists as β_i can not end with a black edge. By definition of r' the edge $e_{2r'-1}$ is black and has a different colour than $e_{2r'+1}$. Since adjacent edges have different colours, the three edges $e_{2r'-1}, e_{2r'}$ and $e_{2r'+1}$ have all different colours. This means we can subdivide every β_i into lifts of simple loops in A and possibly a segment of three differently coloured edges if β_i contains an odd number of edges.

Observe that we have for the deformations of the subpaths $v_{2i-1} + v_{2i} \neq 0$ if and only if β_{2i-1} contains an odd number of edges. This implies that if the deformations of β_{2i-1} and β_{2i} do not cancel each other out, then we have three differently coloured subsequent edges in β_{2i-1}.

Let now $A, B \in U$ be normalised and β be a loop segmented into $\beta_1 \dots \beta_{2m}$ as before. We get then for the stretching factor of β by the previous paragraphs and by Lemma 1.14

$$\begin{aligned}\frac{l_{\psi(B)}(\beta)}{l_{\psi(A)}(\beta)} &= \frac{\sum_{i=1}^{m}\big(l_{\psi(B)}(\beta_{2i-1}) + l_{\psi(B)}(\beta_{2i})\big)}{\sum_{i=1}^{m}\big(l_{\psi(A)}(\beta_{2i-1}) + l_{\psi(A)}(\beta_{2i})\big)} \\ &\leq \max_{1\leq i\leq m}\left\{\frac{l_{\psi(B)}(\beta_{2i-1}) + l_{\psi(B)}(\beta_{2i})}{l_{\psi(A)}(\beta_{2i-1}) + l_{\psi(A)}(\beta_{2i})}\right\} \\ &\leq \max_{1\leq i\leq m}\left\{\Lambda_R(A,B), 1, \frac{1\pm\varepsilon(B)}{1\pm\varepsilon(A)}\right\} = \Lambda_R(A,B).\end{aligned}$$

The last inequality follows from the decomposition of the β_{2i-1} and β_{2i} into lifts of simple loops and a possible segment with three colours and some deformation $v_{2i-1}+v_{2i} = \pm\varepsilon(\cdot)$. Recall that three differently coloured edges have up to some deformation length 1. The last equality comes from Example 6.4. Since any two-coloured simple loop in $\psi(A)$ has the same length as in $\phi^*(A)$, we have $\Lambda_R(\psi(A), \psi(B)) = \Lambda_R(A,B)$. By Lemma 6.2 we have then that ψ is an isometric embedding and by Lemma 6.5 that ψ is indeed a bounded deformation of ϕ^*. □

Using Lemma 6.6 we can now construct bounded deformations of any natural embedding $\phi^*\colon CV_2 \to CV_k$.

Theorem 6.7
Let $\phi^*\colon CV_2 \to CV_k$ be an isometric embedding as in Proposition 5.28 for some $k \geq 3$. Then ϕ^* can be locally deformed.

Proof. Let $\Delta \subset CV_2$ be the simplex corresponding to the theta-graph (Γ, m) with the standard marking, i.e. we have the two simple loops α and β. We now want to apply Lemma 6.6, that is we need a non-transitive monodromy action of a simple loop. Assume the monodromy action of α and β are both transitive, then there exists a $r \in \mathbb{N}$ such that $\beta\alpha^r$ fixes a lift of the basepoint. In particular the simplex $\Delta' \subset CV_2$ corresponding to the theta-graph (Γ, m') with the edge labels $(\beta\alpha^r, \alpha)$ satisfies the conditions of Lemma 6.6. □

While we can deform any natural embedding from CV_2 to CV_k, the same does not hold if we naturally embed CV_n to CV_k for $n \geq 3$. As it turns out these embeddings are rigid in the sense that they have no bounded deformation. The informal reason for this is that the three-coloured path in the proof of Lemma 6.6 can now be a witness and deforming it would properly change the distance in CV_k.

Theorem 6.8
Let $\phi\colon F_k \to F_n$ and $\phi^*\colon CV_n \to CV_k$ be as in Proposition 5.28 and $n \geq 3$. Then ϕ^* has no non-trivial bounded deformations.

We will use the remainder of this section to prove Theorem 6.8. Let from now on $n \geq 3$, $\phi^*\colon CV_n \to CV_k$ be as in Proposition 5.28 a natural embedding and $\psi\colon CV_n \to CV_k$ another isometric embedding such that their difference set $U := \{A \in CV_n \mid \phi^*(A) \neq \psi(A)\}$ is bounded. Assume $\phi^* \neq \psi$, that is we suppose $U \neq \emptyset$. We will split the proof into the following smaller steps:

(i) Show that there exists an $A \in U$ such that $\phi^*(A)$ and $\psi(A)$ differ only in their edge lengths.

(ii) Create families of closed paths in $\phi^*(A)$ which have the same lengths in $\phi^*(A)$ and $\psi(A)$.

(iii) Use these families to show that each edge in $\phi^*(A)$ and $\psi(A)$ already has the same length, which contradicts $\phi^*(A) \neq \psi(A)$.

Lemma 6.9
There exists an $A \in U$ such that $\phi^*(A)$ and $\psi(A)$ have the same topological type.

Proof. Since isometries are continuous, U is open and hence has non-empty intersection with an open maximal simplex Δ. As U has finite diameter there exists a small enough ε such that the ε-thin graphs are disjoint from U. In particular we have $\partial U \cap \Delta \neq \emptyset$. ϕ^* sends points of maximal simplices to maximal simplices Since ϕ^* and ψ are continuous, we have for any $A \in U \cap \Delta$ close enough to ∂U that the images $\phi^*(A)$ and $\psi(A)$ have the same topological type. □

We will from now on fix such an A in a maximal simplex and write $\phi^*(A) = (\Gamma, l_1, m)$ and $\psi(A) = (\Gamma, l_2, m)$ such that the (Γ, l_i) are normalised. For any not necessarily closed edge-path $\tilde{\alpha}$ in Γ we will denote the difference of their length as $v(\tilde{\alpha}) := l_2(\tilde{\alpha}) - l_1(\tilde{\alpha})$.

Recall that by construction Γ is a finite cover of A and the fundamental group of A acts via deck transformation. Hence, we consider as in Remark 5.29 the vertices and edges of Γ as $[F_n : \phi(F_k)] = \frac{k-1}{n-1} =: d$ copies of the vertices respectively edges of A.

Similarly we will denote the lifts of a path α in A by $\alpha_1, \dots \alpha_d$. For a path in Γ we will write $\tilde{\alpha}$ and denote its projection in A again as α. Keep in mind that for closed paths this projection is up to conjugation exactly the identification of elements in F_k as elements of F_n via ϕ. The projection of a cyclically reduced path α will be again cyclically reduced since the cover is unramified. For the purpose of this section identifying paths via the projection will be clearer and avoids ambiguities in regard to the base point.

Since the projection of (Γ, l_1, m) onto A is locally isometric, we have $l_1(\alpha_i) = l_A(\alpha)$, hence witnesses are preserved under this projection. In particular we have for a closed loop $\tilde{\alpha}$ in Γ that $\tilde{\alpha}$ is a witness from $\phi^*(A)$ to $\phi^*(B)$ if and only if α is a witness from A to B. We will see that this implies that witnesses can not be deformed.

Lemma 6.10
Let $\tilde{\alpha}$ be a cyclically reduced, closed path in Γ and α its projection in A. If there exist $B_L, B_R \in CV_n \setminus U$ such that α is a witness from B_L to A and from A to B_R, i.e. $\alpha \in W_R(B_L, A) \cap W_R(A, B_R)$, then we have $v(\tilde{\alpha}) = 0$.

Proof. By Corollary 1.24 there exists a geodesic γ from B_L to B_R which passes through A. Since ψ is isometric, it also sends γ to a geodesic. Furthermore, $\tilde{\alpha}$ is a witness from $\phi^*(B_L) = \psi(B_L)$ to $\phi^*(B_R) = \psi(B_R)$ and thus by Lemma 1.22 also a witness from $\psi(B_L)$ to $\psi(A)$.

As ϕ^* and ψ are isometric, we have (after normalising) that the maximal stretching of paths are equal:

$$\Lambda_R(\phi^*(B_L), \phi^*(A)) = \Lambda_R(B_L, A) = \Lambda_R(\psi(B_L), \psi(A))$$

Since $\tilde{\alpha}$ is a witness, we have

$$\frac{l_1(\tilde{\alpha})}{l_{\phi^*(B_L)}(\tilde{\alpha})} = \Lambda_R(\phi^*(B_L), \phi^*(A)) = \Lambda_R(\psi(B_L), \psi(A)) = \frac{l_2(\tilde{\alpha})}{l_{\psi(B_L)}(\tilde{\alpha})}$$

and by using $\phi^*(B_L) = \psi(B_L)$ we have the desired equality $l_1(\tilde{\alpha}) = l_2(\tilde{\alpha})$. □

We can now use Lemma 6.10 to show that any path $\tilde{\alpha}$ whose projection does not cover the whole graph A has no deformation:

Lemma 6.11
Let $\tilde{\alpha}$ be a cyclically reduced, closed path in Γ and α its projection in A. If there exists an edge $e \in E(A)$ which is not covered by α, then we have $v(\tilde{\alpha}) = 0$.

Proof. We will construct two points B_L, B_R as in Lemma 6.10.

To obtain B_L we shrink each edge of A covered by $\tilde{\alpha}$ by a factor $\varepsilon > 0$ and leave the rest with the same length, that is we have $B_L \in \Delta(A)$ with length function

$$l_{B_L}(e) = \begin{cases} \varepsilon \cdot l_A(e), & \text{if } e \text{ is an edge contained in } \alpha \\ l_A(e), & \text{if } e \text{ is an edge not contained in } \alpha \end{cases}$$

Since $l_{B_L}(\alpha) = \varepsilon l_A(\alpha)$, we have $\Lambda_R(B_L, A) \geq \frac{1}{\varepsilon}$. On the other hand there exists an unshrunk edge e and hence we have

$$d_R(B_L, A) = \log\left(\frac{\mathrm{vol}(B_L)}{\mathrm{vol}(A)} \cdot \Lambda_R(B_L, A)\right) \geq \log\left(\frac{l_A(e)}{1 \cdot \varepsilon}\right).$$

As U is bounded, we have that $B_L \notin U$ for small enough ε. Furthermore, each path $\beta \in F_n$ is at most stretched by $1/\varepsilon$ from B_L to A and hence we have $\alpha \in W_R(B_L, A)$.

We construct B_R similarly, that is we shrink the edges not covered by $\tilde{\alpha}$ by a factor ε to obtain B'_R. Note that B'_R might be still in U, as the shrunk edges might be a forest and thus even setting $\varepsilon = 0$ can yield a point in U. Nevertheless we have $\alpha \in W_R(A, B'_R)$. By Lemma 1.36 there exists a geodesic ray $\gamma \colon \mathbb{R}_{\geq 0} \to CV_n$ continuing the geodesic from A to B'_R and keeping α as a witness. Hence, for large enough t we have $B_R := \gamma(t) \notin U$ and by construction $\alpha \in W_R(A, B_R)$. The claim follows now from Lemma 6.10. □

Observe that for $n \geq 3$ and a candidate α in $A \in CV_n$ there exists an edge $e \in E(A)$ which is not covered by α. For example if α is a barbell, there exists a spanning tree $T \subset A$ which contains all but two edges of α. As $\pi_1(A)$ has at least rank three, there exist at least three edges in $e \setminus T$ and hence α does not cover all edges of A. This means that by Lemma 6.11 closed loops which project to a power of a candidate are never deformed. Nevertheless, we still have that a lift of a candidate α may be deformed if it is not closed in Γ. For instance we have seen this in Example 6.4 where we have $v((\beta^2)_i) = 0$ but $v(\beta_i) = \pm\varepsilon(A)$. To get that non-closed paths are not deformed, we consider them as handles of barbells:

Lemma 6.12
Let α and β be two not necessarily distinct, cyclically reduced loops in A and $\rho \subset A$ a path connecting them such that $\alpha\rho\beta\overline{\rho}$ is cyclically reduced and not every edge of A is covered by α, β and ρ. Furthermore, let $\tilde{\rho}$ be a lift of ρ in Γ. Then $\tilde{\rho}$ is not deformed, that is we have $v(\tilde{\rho}) = 0$.

Proof. Let $\tilde{\alpha}^d, \tilde{\beta}^d$ be the lifts of α^d and β^d starting at the endpoints of $\tilde{\rho}$. As d is the index of the subgroup, both $\tilde{\alpha}^d$ and $\tilde{\beta}^d$ are loops in Γ. By Lemma 6.11 we have $v(\tilde{\alpha}^d) = 0 = v(\tilde{\beta}^d)$ and as edge-path $v(\tilde{\alpha}^d\tilde{\rho}\tilde{\beta}^d\overline{\tilde{\rho}}) = 0$, where $\overline{\tilde{\rho}}$ denotes as usual the path $\tilde{\rho}$ with reversed orientation. On the other hand the variation is additive, hence we have

$$0 = v(\tilde{\alpha}^d\tilde{\rho}\tilde{\beta}^d\overline{\tilde{\rho}}) = v(\tilde{\alpha}^d) + v(\tilde{\rho}) + v(\tilde{\beta}^d) + v(\overline{\tilde{\rho}}) = 2v(\tilde{\rho}),$$

which concludes the proof. □

Lemma 6.12 in particular implies that the handle of a barbell is never deformed. We will now finish the proof of Theorem 6.8.

Proof of Theorem 6.8. Let $e \in E(A)$ be an edge and $e_1 \in E(\Gamma)$ one of its lifts. We will now show, that $v(e_1) = 0$, that is there exists no deformed edge. We have the following three cases for e:

e **is a loop in** A**:** As A is in a maximal simplex and e is a loop, e can be separated from the rest of A by its adjacent edge. That means we can see it as a petal of a barbell, i.e. we have a disjoint simple loop α and a path ρ between e and α such that the edge-path $e\rho\alpha$ is a barbell in A. Keep in mind that for each lift α_i of a closed loop α there exists a lift $\alpha_i^d := (\alpha^d)_i$ which is a closed loop starting with α_i.

We will enumerate the lifts in Γ in such a manner that $e_1 \dots e_m$ is a simple loop in Γ and the ρ_i share their starting vertex with e_i and their ending vertex with the loops α_i^d for $i = 1, \dots m$. Namely we have the following image in Γ:

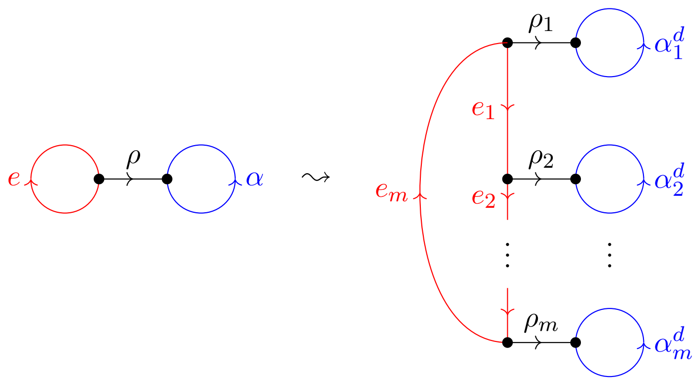

As $e\rho\alpha$ is a barbell in A and $n \geq 3$, it does not cover the whole graph A. Hence, by Lemma 6.12 we have $v(\rho_1) = 0 = v(\rho_2)$ and $v(\overline{\rho_1} e_1 \rho_2) = 0$. By additivity of v we have then $0 = v(\overline{\rho_1} e_1 \rho_2) = v(\rho_1) + v(e_1) + v(\rho_2) = v(e_1)$.

e **is a separating edge:** If the edges adjacent to e are not separating, then there exists in each component of $A \setminus e$ a simple loop sharing a vertex with e. Hence, by Lemma 6.12 e is not deformed.

If some edges adjacent to e are separating, then there exist two disjoint paths ρ_α, ρ_β connecting e to two simple loops α, β in one component of $A \setminus e$. Let ω be a simple loop in the other component connected to e by a path ρ_ω:

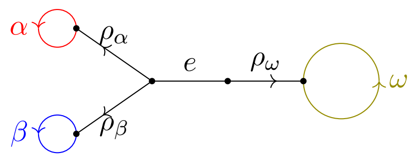

By Lemma 6.12 each lift of the paths $\overline{\rho_\alpha}\rho_\beta, \overline{\rho_\alpha} e \rho_\omega, \overline{\rho_\beta} e \rho_\omega$ has trivial deformation. Let $\rho_{\alpha,1}, \rho_{\beta,1}$ and $\rho_{\omega,1}$ be the lifts of these paths adjacent to e_1. By additivity of the deformation we then have

$$
\begin{aligned}
& v(\overline{\rho_{\alpha,1}} \rho_{\beta,1}) = 0 \\
\Rightarrow \quad & v(\rho_{\alpha,1}) = -v(\rho_{\beta,1}) \\
& v(\overline{\rho_{\alpha,1}} e_1 \rho_{\omega,1}) = 0 = v(\overline{\rho_{\beta,1}} e_1 \rho_{\omega,1}) \\
\Rightarrow \quad & v(\rho_{\alpha,1}) = v(\rho_{\beta,1})
\end{aligned}
$$

and hence $v(\rho_{\alpha,1}) = v(\rho_{\beta,1}) = 0$.

Similar we get $v(\rho_{\omega,1}) = 0$, either by choosing ρ_ω to be trivial if the adjacent edge of e is no separating or as in the case of α and β if the adjacent edge of e is separating. By $v((\overline{\rho_\alpha} e_1 \rho_\omega)_1) = 0$ we have then $v(e_1) = 0$.

e is a non-separating edge: We will construct two disjoint simple paths α and β starting and ending at the endpoints of e. That is we have the two cases:

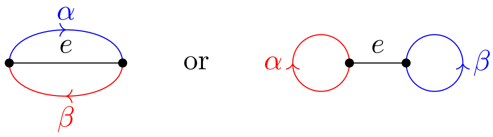

In the first case $\alpha\beta$ is a simple loop, hence there exists an edge not covered by α, β and e. So in both cases we can apply Lemma 6.12 to get $v(e_1) = 0$.

The existence of one such simple path α is given since e is non-separating. For example take a shortest path between the endpoints of e in $A \setminus \{e\}$. Furthermore as each vertex has valency 3, we have two additional edges f and g adjacent to e and disjoint from α:

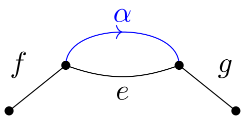

We assume $f \neq g$, else we set $\beta = f$ and are done. If f and g are separating edges, we extend them to be maximal paths consisting of separating edges and let β, ω be two simple loops adjacent to the endpoints of f and g. By applying Lemma 6.12 to the three barbells with simple loops $\beta, \alpha e, \omega$ we have for their lifts $v(f_1) = 0 = v(g_1)$ and $v(\overline{f_1} e_1 g_1) = 0$, from which follows directly $v(e_1) = 0$. Similarly if f is a separating edge and g is non-separating, then there exists a simple loop ω in $A \setminus \{e, f\}$ containing g and we get from Lemma 6.12 $v(\overline{f_1} e_1) = 0$ for the barbell with loops β and ω and as before $v(f_1) = 0$, thus $v(e_1) = 0$:

Assume now that f and g are non-separating edges, then we can extend them to a path β connecting the endpoints of e:

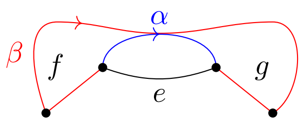

The path β may intersect α on some edges. As in the proof of Lemma 3.1 we can construct from α and β two disjoint paths α' and β' by iteratively cutting out common edges. Since α and β are disjoint in the neighbourhood of e, the disjoint paths created α' and β' still join the vertices of e. Hence, we have also for the final case $v(e_1) = 0$. $\square$

Appendix: An implementation in Sage of the algorithm described in Section 1.6

The following code represents a point in Outer Space as an object in Sage [Sag] and includes the methods to compute the candidates of a point and the Lipschitz distance between two points as described in Section 1.6. The code together with some examples can be found in [Ste18].

```
class MarkedGraph(object):
    """
    Class for marked graphs, e.g. elements of Outer Space aka Culler-Vogtmann space.

    This class provides some additional functions, e.g. to find candidates for
    maximal stretching ((topologically) embedded simple curves, figure of eights
    and barbells) and conversion of cycles to elements of fundamental group and back.
    The calculation of the Lipschitz distance comes from the paper
    https://arxiv.org/pdf/0803.0640.pdf

    CONDITIONS:

    The labels of the edges must be elements of a group (to make sense the
    fundamental group, but you can abuse this) or None.

    The graph is directed in the following sense: an edge is oriented from
    the smaller index to the bigger one (i.e. (2, 3, g) means if you go from
    vertex 2 to 3 corresponds g and from 3 to 2 corresponds to g^-).

    None will be converted to the neutral element and at least one label
    must be not None.

    The giving weights function is a map from edges to your weights, where
    an edge is given by the triple ('smaller vertex', 'bigger vertex', 'label').
    Only tested, when vertices of graph are (real) numbers (they should be
    at least comparable/totally ordered).
    """
    # Variables:
    # * _candidates = [[], [], []]
    # * _weights = {}
    # * _graph = Graph()

    # INIT
    def __init__(self, graph=Graph(), weights=None):
        """
        Initialize

        If no weights are given, all edges are given weight 'wi_j_l'
        (edge = (i, j, l) and change the label accordingly) before making it nice
        """
        # setting the standard variables to trivial
```

```
        # first entry are simple curves, second are eights, third are barbells
        self._candidates = [[], [], []]
        if weights:
            # since we actually change the graph for computation, use a copy
            self._weights = deepcopy(weights)
        else:
            # self._weights = {edge: 1 for edge in graph.edges()}
            self._weights = {edge: var('w%s_%s_%s' % (edge[0], edge[1],
                        str(edge[2]).replace('*', '').replace('^-1','_inv'))
                        ) for edge in graph.edges()}
        self.setGraph(graph)

    # HELPER CLASSES
    def _getOne(self):
        """
        Return the 1 of the corresponding group for the labels or 't' if there are no edges.

        If first not None edges label is not in a group, an error is raised.
        """
        edges = self._graph.edges()
        for edge in edges:
            if edge[2]:
                return edge[2].parent().one()
        print('WARNING: _getOne() was called without any edges/not None labels. '
              'Returned None as value.')
        return None

    # GRAPH AND SPANNING TREE
    # The methods to define/get the graph, spanning tree and to make graph nice

    def setGraph(self, graph=Graph()):
        """
        Set the graph to given variable and change None-labels to the neutral element.
        """
        self._graph = graph
        self._convertNone()

    def _convertNone(self):
        """
        Convert all None edges to neutral element.
        """
        one = self._getOne()
        for edge in self._graph.edges():
            if edge[2] == None:
                self._graph.delete_edge((edge[0], edge[1], None))
                self._graph.add_edge(edge[0], edge[1], one)
                self._weights.update(
                    {(edge[0], edge[1], one):
                     self._weights.pop((edge[0], edge[1], None), None)})
```

```
def getGraph(self):
    """
    Return (non-copy) graph on which the element is based
    """
    return self._graph


def getSpanningTree(self):
    """
    Return subgraph of the edges with neutral element label
    """
    one = self._getOne()
    self._convertNone()
    spanEdges = [e for e in self._graph.edges() if e[2] == one]
    tree = self._graph.subgraph(edges=spanEdges)
    if tree.has_loops() or tree.has_multiple_edges():
        print('WARNING: Spanning tree has loops or multiple edges.'
              'Check the labeling. This might cause errors later on.')
    return(tree)

def makeGraphNice(self, graph=None, stretchVariable=None, considerWeights=False):
    """
    Split edges so graph has no loops or multiple edges.

    This is done by introducing the minimal number of vertices on edges midpoints.

    Respects marking (and weights) (new introduced edges have label
    'stretchVariable' and (weight 0) and 'old' edge keeps marking(&weight)).

    WARNING: This method might throw an error, if weights and marking are
    not correctly set, e.g. if edges are indistuingishable. To keep weights,
    set it on True, else it will break the weights!
    """
    if graph is None:
        graph = self._graph
    vertices = graph.vertices()
    if not vertices:
        print 'Warning: Given graph has has no vertices!'
        return
    newVertexNumber = max(vertices) + 1

    for edge in graph.loops():
        start = min([edge[0], edge[1]])
        end = max([edge[0], edge[1]])
        graph.add_vertices([newVertexNumber, newVertexNumber + 1])

        # introduce edges of three and assign and markings
        graph.add_edge(start, newVertexNumber, edge[2])
```

```
            graph.add_edge(newVertexNumber, newVertexNumber + 1, stretchVariable)
            graph.add_edge(end,newVertexNumber+1, stretchVariable)

            # weightspart
            if considerWeights:
                self._weights.update({(start, newVertexNumber, edge[2]): self._weights.pop(
                    edge)})
                self._weights.update({(newVertexNumber, newVertexNumber + 1,
                    stretchVariable): 0})
                self._weights.update({(end, newVertexNumber + 1, stretchVariable): 0})

            # cleanup
            graph.delete_edge(edge)
            newVertexNumber = newVertexNumber + 2

        for edge in graph.multiple_edges():
            start=min([edge[0], edge[1]])
            end=max([edge[0], edge[1]])
            graph.add_vertex(newVertexNumber)

            # introduce edges and assign marking
            graph.add_edge(start, newVertexNumber, edge[2])
            graph.add_edge(newVertexNumber, end, stretchVariable)

            # weightspart
            if considerWeights:
                self._weights.update({(start, newVertexNumber, edge[2]): self._weights.get(edge
                    )})
                self._weights.update({(end, newVertexNumber, stretchVariable): 0})
                self._weights.pop(edge)

            #cleanup
            graph.delete_edge(edge)
            newVertexNumber = newVertexNumber + 1

    # MARKINGS AND WEIGHTS
    # Methods to set and get markings in different ways
    # Every method that changes marking also resets fundamental group.

    def setWeights(self, weights):
        self._weights = weights

    def updateWeights(self, weights):
        self._weights.update(weights)

    def getWeights(self):
        return deepcopy(self._weights)

    def setCentralWeights(self):
```

```
        """
        Set all weights to 1 and return copy of the weighting function
        """
        weights = {}
        G = Graph(self._graph)
        for edge in G.edges():
            weights.update({edge: 1})
        self._weights = weights
        return deepcopy(weights)

    # CANDIDATES
    # This part involves how to get the candidates and simple cycles
    def simple_cycles(self, Gamma=None, non_duplicates=True):
        """
        Return a copylist of all simple cycles in the given graph Gamma as lists vertices

        OPTIONS:

        - non_duplicates=True enforces the list to have unique elements (only important for
            multiedges).
        """
        if Gamma is None:
            Gamma = self._graph
        G = Graph(Gamma) # make copy of given graph
        cycleList = []

        multi = []

        if non_duplicates:
            for edge in G.multiple_edges() + G.loops():
                if not [edge[0], edge[1]] in multi:
                    multi.append([edge[0], edge[1]])
        else:
            multi = G.multiple_edges() + G.loops()
        for mult in multi:
            # case of loops
            if mult[0] == mult[1]:
                cycleList.append([mult[0], mult[0]])
            # case of multiple edge
            else:
                cycleList.append([mult[0], mult[1], mult[0]])
        G.allow_loops(False)
        G.allow_multiple_edges(False)
        edges = G.edges()
        for edge in edges:
            G.delete_edge(edge) # delete starting edge to avoid duplicates
            # add all cycles containing this edge
            cycleList.extend([[edge[0]] + path for path in
                              G.all_paths(edge[1], edge[0])])
```

```
        self._candidates[0] = deepcopy(cycleList)
        return cycleList


    def candidates(self, Gamma=None, reevaluate=True):
        """
        Return the candidates as group elements.
        """
        if not reevaluate and self._candidates[0]:
            return self._candidates
        if Gamma is None:
            Gamma = self._graph
        simples = self.simple_cycles(Gamma=Gamma)
        candidates = []
        for cycle in simples:
            candidates.extend(self.path_to_elementList(cycle))
        candidates = [candidates] # the simple loops
        for path in simples: # delete last vertex in loops
            path.pop()

        eights = [] # list of eights
        barbells = [] # list of barbells

        loops = Gamma.loops()
        for i in range(len(loops)):
            for j in range(i + 1, len(loops)):
                if loops[i][0] == loops[j][0]:
                    eights.append(loops[i][2]*loops[j][2])
                    eights.append(loops[i][2]/loops[j][2])

        number_simples = len(simples)
        while simples:
            cyc = simples.pop()
            for other_cyc in simples:
                intersec = list(set(cyc) & set(other_cyc))
                if len(intersec) == 1: # single point intersection gives eight
                    firstCyc = cyc[cyc.index(intersec[0]):]
                    firstCyc.extend(cyc[:cyc.index(intersec[0]) + 1])
                    firstCyc = self.path_to_elementList(firstCyc)
                    i = other_cyc.index(intersec[0])
                    secondCyc = (other_cyc[i:])
                    secondCyc.extend(other_cyc[:i+1])
                    secondCyc = self.path_to_elementList(secondCyc)
                    for first in firstCyc:
                        for second in secondCyc:
                            eights.append(first*second)
                            eights.append(first/second)
                if len(intersec) == 0: # no intersection might gives barbells
                    length = len(other_cyc)
```

```
                barb = []
                for vertcyc in cyc:
                    firstCyc = cyc[cyc.index(vertcyc):]
                    firstCyc.extend(cyc[:cyc.index(vertcyc)+1])
                    firstCyc = self.path_to_elementList(firstCyc)
                    for i in range(length):
                        G = Graph(Gamma)
                        G.delete_vertices([vert for vert in cyc if vert != vertcyc])
                        G.delete_vertices([vert for vert in other_cyc if vert != other_cyc[i
                            ]])
                        secondCyc = (other_cyc[i:])
                        secondCyc.extend(other_cyc[:i+1])
                        secondCyc = self.path_to_elementList(secondCyc)
                        for path in G.all_paths(vertcyc, other_cyc[i]):
                            middlePart = self.path_to_elementList(path)
                            for first in firstCyc:
                                for second in secondCyc:
                                    for mid in middlePart:
                                        barbells.append(first*mid*second/mid)
                                        barbells.append(first*mid/second/mid)
        candidates.append(eights)
        candidates.append(barbells)
        self._candidates = deepcopy(candidates)
        return candidates


    def path_to_elementList(self, path):
        """
        Take a path as list of vertices and return a list of words each representing the path.
        """
        self._convertNone() # if there were changes in between
        one = self._getOne()
        letters = []
        for i in range(len(path)-1):
            if path[i] < path[i+1]:
                letters.append(self._graph.edge_label(path[i], path[i+1]))
            else:
                letters.append([letter.inverse() for letter in
                                self._graph.edge_label(path[i], path[i+1])])
        words = []
        word = one

        treshHolds = [len(l) for l in letters]
        treshHolds.append(0)
        i = [0] * (len(letters) + 1)
        j = 0
        nextIterate = True
        while nextIterate:
            if (j > 0 and path[j-1] == path[j+1] and
```

```
                letters[j][i[j]]*letters[j-1][i[j-1]] == one):
                # don't walk back and forth again
                i[j] += 1
            else:
                word *= letters[j][i[j]]
                j += 1
                if j == len(letters):
                    words.append(word)

            #increase list indices and shorten word back again:
            while i[j] == treshHolds[j]:
                i[j] = 0
                j -= 1
                if j >= 0:
                    word /= letters[j][i[j]]
                    i[j] += 1
                else:
                    nextIterate = False
                    break
        # if the path is a loop or 2 multiedges we didn't sort out the inverse yet.
        oldwords = words
        words = []
        for word in oldwords:
            if not (word.inverse() in words or word in words):
                words.append(word)
        return words

    # FUNDAMENTAL GROUP
    # Functions to generate the fundamental group and to translate it to closed paths in the
        graph and back

    def generateFundGroup(self):
        """
        Return the subgroup generated by the labels.

        If there are no edges, return None.
        """
        generators = self._graph.edge_labels()
        if generators:
            return generators[0].parent().subgroup(generators)
        else:
            return None

    def groupelementAsPath(self, word):
        """
        Take an element of the fundamental group and return a cyclically reduced path (list of
            edges) realizing the word.
        """
        tree = self.getSpanningTree() # also converts None to 1
```

```
        tree.allow_multiple_edges(False) # force it to have no multiple edges (implies unique
            labels)

        edges = self._graph.edges()

        converter = {} # dictionary of generator to edge
        generators = [] # the edge-labels generate the group, these are svaed here
        for edge in self._graph.edges():
            if edge[2] == word.parent().one():
                continue
            generators.append(edge[2])
            converter.update({len(generators):edge})
        # write a word as tietze with the new generators
        tietze = self.convertGroupelementToHereTietze(word, gens=generators)
        tietze = self.cyclycallyReduced(tietze) # cyclically reduce word

        number = tietze[-1] # throw an error if the word was trivial
        # check orientation of last edge to get first vertex of cycle
        if number > 0:
            lastVertex = converter.get(abs(number))[1]
        else:
            lastVertex = converter.get(abs(number))[0]

        path = []
        for number in tietze:
            edge = converter.get(abs(number))
            if number > 0:
                start = edge[0]
                end = edge[1]
            else:
                start = edge[1]
                end = edge[0]
            # add path in tree
            treeWalk = tree.shortest_path(lastVertex, start)
            for i in range(len(treeWalk) - 1):
                if treeWalk[i] < treeWalk[i+1]:
                    path.append((treeWalk[i], treeWalk[i+1],
                                 tree.edge_label(treeWalk[i], treeWalk[i+1])))
                else:
                    path.append((treeWalk[i+1], treeWalk[i],
                                 tree.edge_label(treeWalk[i+1], treeWalk[i])))
            path.append(edge)
            # save lastVertex of path
            lastVertex = end
        return path

    def convertGroupelementToHereTietze(self, word, gens):
        """
        Take a word and write it in terms of the given generators.
```

```
    """
    # make it compatible for translation into GAP
    rank = len(gens)
    xTietze = list(word.Tietze())
    tgens = [list(gen.Tietze()) for gen in gens]

    # translate word and generators to GAP
    gap.eval('tWord := %s;; tgens := %s;; rank := %s;;'
             % (xTietze, tgens, rank))
    gap.eval('G := FreeGroup(rank);; '
             'Ggens := GeneratorsOfGroup(G);;')
    gap.eval('word := One(G);; for i in tWord do '
             'nextLetter := Ggens[AbsInt(i)]; '
             'if i>0 then word := word*nextLetter; '
             'else word := word/nextLetter; fi;  od;')
    gap.eval('newGens := [];; '
             'for tWord in tgens do newGen := One(G);; '
             'for i in tWord do nextLetter := Ggens[AbsInt(i)]; '
             'if i>0 then newGen := newGen*nextLetter; '
             'else newGen := newGen/nextLetter; fi;  '
             'od; Add(newGens, newGen); od;')
    gap.eval('H := Subgroup(G, newGens);;')

    # change the base in GAP
    gap.eval('hom := EpimorphismFromFreeGroup(H);;')
    gap.eval('newWord := PreImagesRepresentative(hom, word);;')

    # and translate new Tietze-word back to Sage
    newTietze = gap('TietzeWordAbstractWord(newWord);').sage()
    return newTietze


def cyclycallyReduced(self, tietzeWord):
    """
    Take a Tietze represenation of a word and return a cyclycally reduced copy of it.

    Does only conjugate and not check interior of word.
    """
    tietzeList = list(tietzeWord)
    while tietzeList and tietzeList[0] == -tietzeList[-1]:
        tietzeList.pop()
        tietzeList.pop(0)
    newWord = tuple(tietzeList)
    return newWord

# DISTANCES
# Functions to determine the Lipschitz distance and lengths of paths.

def lengthOfElement(self, groupElement):
```

```
        """
        Return length of the cyclycally reduced path corresponding to the element of the
            fundamental group.
        """
        return(self.lengthOfPath(self.groupelementAsPath(groupElement)))

    def lengthOfPath(self, path, weights=None, graph=None):
        """
        Take a list of edges (=path) and return the length, i.e. sum of edge weights.

        If path is empty or point, return 0.
        """
        if not weights:
            weights = self._weights
        if not graph:
            graph = self._graph
        if not path:
            return 0
        length = 0
        for edge in path:
            weight = weights.get(edge)
            length = length + weight
        return length

    def distanceTo(self, markedGraph, lazy=False):
        """
        Compute asymmetric Lipschitz distance to given markedGraph

        If lazy, it takes the saved candidates if possible.
        """
        distance = log(max(self.candFractions(markedGraph, lazy).values()
                            )*self.volume()/markedGraph.volume())
        return distance

    def volume(self):
        """
        Return the volume, i.e. the sum of edge lengths of this graph
        """
        edges = self._graph.edges()
        sum = self._weights.get(edges[0])
        for edge in edges:
            sum += self._weights.get(edge)
        sum -= self._weights.get(edges[0])
        return sum


    def candFractions(self, markedGraph, lazy=False):
        """
        Return a dictionary of all ratios of lengths
```

```
These are ratios of lenghts of the candidates here and their image in the given marked
    graph

Keys are elements in fundamental group
If lazy, it will not reevaluate the candidates (if there are at least some simple cycles).
"""
candidates = self.candidates(reevaluate= (not lazy))
candidates = candidates[0] + candidates[1] + candidates[2]
fractionCand = {}
for candidate in candidates:
    denominator = self.lengthOfElement(candidate)
    numerator = markedGraph.lengthOfElement(candidate)
    fractionCand.update({candidate: numerator/denominator})
return fractionCand
```

References

[AB16] Roger Alperin and Hyman Bass. *Length functions of group actions on Λ-trees*, pages 265–378. Princeton University Press, 2016.

[AK19] Yael Algom-Kfir. The metric completion of Outer Space. *Geometriae Dedicata*, 204:191–230, 2019.

[Bes02] Mladen Bestvina. The topology of $\mathrm{Out}(F_n)$. In *Proceedings of the International Congress of Mathematicians, Vol. II (Beijing, 2002)*, pages 373–384. Higher Ed. Press, Beijing, 2002.

[Bes11] Mladen Bestvina. A Bers-like proof of the existence of train tracks for free group automorphisms. *Fundamenta Mathematicae*, 214(1):1–12, 2011.

[BF00] Mladen Bestvina and Mark Feighn. The topology at infinity of $\mathrm{O}ut(F_n)$. *Inventiones Mathematicae*, 140(3):651–692, 2000.

[BFH00] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $\mathrm{O}ut(F_n)$. I. Dynamics of exponentially-growing automorphisms. *Annals of Mathematics. Second Series*, 151(2):517–623, 2000.

[BFH05] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $\mathrm{O}ut(F_n)$. II. A Kolchin type theorem. *Annals of Mathematics. Second Series*, 161(1):1–59, 2005.

[BH92] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Annals of Mathematics*, 135(1):1–51, 1992.

[CM87] Marc Culler and John Morgan. Group actions on $\mathbb{R}$-trees. *Proceedings of the London Mathematical Society. Third Series*, 55:571–604, 1987.

[CV86] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones Mathematicae*, 84(1):91–119, 1986.

[DLRT20] David Dumas, Anna Lenzhen, Kasra Rafi, and Jing Tao. Coarse and fine geometry of the Thurston metric. *Forum Math. Sigma*, 8:Paper No. e28, 58, 2020.

[FM11] Stefano Francaviglia and Armando Martino. Metric properties of Outer Space. *Publicacions Matemàtiques*, 55(2):433–473, 2011.

[FM12a] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.

[FM12b] Stefano Francaviglia and Armando Martino. The isometry group of Outer Space. *Advances in Mathematics*, 231(3-4):1940–1973, 2012.

[FM21] Stefano Francaviglia and Armando Martino. Displacements of automorphisms of free groups II: Connectivity of level sets and decision problems. *ArXiv Mathematics e-prints*, 2021. arXiv:1807.02782.

[GAP21] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.11.1*, 2021. https://www.gap-system.org.

[Nie24] Jakob Nielsen. Die Isomorphismengruppe der freien Gruppen. *Mathematische Annalen*, 91:169–209, 1924.

[Sag] Sage Developers. *SageMath, the Sage Mathematics Software System (Version 8.0).* https://www.sagemath.org.

[Ser80] Jean-Pierre Serre. *Trees*. Springer-Verlag, Berlin, 1980. Translated from the French by John Stillwell.

[Sko90] Richard K. Skora. Deformations of length functions in groups. 1990. preprint, https://levitt.users.lmno.cnrs.fr/SkoraDeformation.pdf.

[Sta83] John R. Stallings. Topology of finite graphs. *Inventiones Mathematicae*, 71(3):551–565, 1983.

[Ste18] Christian Steinhart. CV-Lipschitz-Calculator, 2018. https://gitlab.com/ctst/CV-Lipschitz-Calculator.

[Ste19] Christian Steinhart. Envelopes in Outer Space. *ArXiv Mathematics e-prints*, 2019. arXiv:1907.06402.

[Thu98] William P. Thurston. Minimal stretch maps between hyperbolic surfaces. *ArXiv Mathematics e-prints*, 1998. arXiv:math/9801039.

[Vog02] Karen Vogtmann. Automorphisms of free groups and Outer space. In *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000)*, volume 94, pages 1–31, 2002.

[Vog15] Karen Vogtmann. On the geometry of outer space. *American Mathematical Society. Bulletin. New Series*, 52(1):27–46, 2015.

[Wal14] Cormac Walsh. The horoboundary and isometry group of Thurston's Lipschitz metric. In *Handbook of Teichmüller theory. Vol. IV*, volume 19 of *IRMA Lect. Math. Theor. Phys.*, pages 327–353. Eur. Math. Soc., Zürich, 2014.