

FREE-GROUP AUTOMORPHISMS, TRAIN TRACKS, AND THE BEADED DECOMPOSITION

MARTIN R. BRIDSON AND DANIEL GROVES

ABSTRACT. We study the automorphisms ϕ of a finitely generated free group F . Building on the train-track technology of Bestvina, Feighn and Handel, we provide a topological representative $f : G \rightarrow G$ of a power of ϕ that behaves very much like the realization on the rose of a positive automorphism. This resemblance is encapsulated in the *Beaded Decomposition Theorem* which describes the structure of paths in G obtained by repeatedly passing to f -images of an edge and taking subpaths. This decomposition is the key to adapting our proof of the quadratic isoperimetric inequality for $F \rtimes_{\phi} \mathbb{Z}$, with ϕ positive, to the general case.

The study of automorphisms of free groups is informed greatly by the analogies with automorphisms of free-abelian groups and surface groups, but one often has to work considerably harder in the free group case in order to obtain the appropriate analogues of familiar results from these other contexts. Nowhere is this more true than in the quest for suitable normal forms and geometric representatives. One can gain insight into the nature of individual elements of $\mathrm{GL}(n, \mathbb{Z})$ by realizing them as diffeomorphisms of the n -torus. Likewise, one analyzes individual elements of the mapping class group by realizing them as diffeomorphisms of a surface. The situation for $\mathrm{Aut}(F_n)$ and $\mathrm{Out}(F_n)$ is more complicated: the natural choices of classifying space $K(F_n, 1)$ are finite graphs of genus n , and no element of infinite order in $\mathrm{Out}(F_n)$ is induced by the action on $\pi_1(Y)$ of a homeomorphism of Y . Thus the best that one can hope for in this situation is to identify a graph Y_{ϕ} that admits a homotopy equivalence inducing ϕ and that has additional structure well-adapted to ϕ . This is the context of the *train track technology* of Bestvina, Feighn and Handel [4, 1, 3].

Their work results in a decomposition theory for elements of $\mathrm{Out}(F_n)$ that is closely analogous to (but more complicated than) the Nielsen-Thurston theory for surface automorphisms. The finer features of the topological normal forms

Date: 27 July 2006.

2000 Mathematics Subject Classification. 20F65, (20F06, 20F28, 57M07).

Key words and phrases. automorphisms of free groups, train tracks, free-by-cyclic groups.

The first author's work was supported in part by Research Fellowships from the EPSRC and a Royal Society Wolfson Research Merit Award. The second author was supported in part by a Junior Research Fellowship at Merton College, Oxford, and by NSF Grant DMS-0504251. We thank these organisations for their support.

that they obtain are adapted to the problems that they wished to solve in each of their papers: the Scott conjecture in [4] and the Tits alternative in the series of papers [1, 3, 2]. The problem that we wish to solve in this series of papers (of which [5] is the first, this is the second and [6] is the third), is that of determining the Dehn functions of all free-by-cyclic groups. This requires a further refinement of the train-track technology. Specifically, we must adapt our topological representatives so as to make tractable the problem of determining the isoperimetric properties of the mapping torus of the homotopy equivalence $f : Y_\phi \rightarrow Y_\phi$ realizing an iterate of ϕ .

The prototype for a train-track representative is the obvious realization of a positive automorphism on the rose. This motivates the following strategy for the Dehn-function problem: first we proved the theorem in the case of positive automorphisms [5], where one already encounters most of the web of large-scale cancellation phenomena that explain why the general theorem is true; then, in the general case, we follow the architecture of the proof from [5], using a suitably refined train-track description of the automorphism in place of the positivity assumption. We shall see in [6] that this approach works remarkably well. However, in order to bring it to fruition, one must deal with myriad additional complexities arising from the intricacies of cancellation that do not arise in the positive case.

The properties of the topological representative $f : G \rightarrow G$ constructed in [1] allow one to control the manner in which a path σ evolves as one looks at its iterated images under f , and one might naively suppose that this is the key issue that one must overcome in translating the arguments from the positive case [5] to the general case [6]. However, upon close inspection one discovers this is actually only a fraction of the story, the point being that when a corridor evolves in the time flow on a van Kampen diagram, the interaction of the forward iterates of the individual edges is such that the basic *splitting* of paths established in [1] gets broken. It is to overcome this difficulty that we need the notion of *hard splitting* introduced in Section 2; such splittings are denoted $\sigma_1 \odot \sigma_2$.

In the analysis of van Kampen diagrams, the class of “broken” paths that one must understand are the residues of the images of a single edge that survive repeated cancellation during the corridor flow. In the language of the topological representative $f : G \rightarrow G$, this amounts to understanding *monochromatic paths*, as defined below. Every edge-path ρ in G admits a unique maximal splitting into edge paths; our purpose in this article is to understand the nature of factors in the case where ρ is monochromatic (grouping certain of the factors into larger units).

To this end, we identify a small number of basic units into which the iterated images of monochromatic paths split; the key feature of this splitting is that it is robust enough to withstand the difficulties caused by cancellation in van

Kampen diagrams. The basic units are defined so as to ensure that they enjoy those features of individual edges that proved important in the positive case [5]. We call the units *beads*. The vocabulary of beads is as follows.

Let $f : G \rightarrow G$ be a topological representative and let $f_{\#}(\sigma)$ denote the tightening rel endpoints of the image of an edge-path σ . Following [4], if $f_{\#}(\tau) = \tau$ we call τ a *Nielsen path*. A path ρ in G is called a *growing exceptional path* (GEP) if either ρ or $\bar{\rho}$ is of the form $E_i \bar{\tau}^k \bar{E}_j$ where τ is a Nielsen path, $k \geq 1$, E_i and E_j are parabolic edges, $f(E_i) = E_i \odot \tau^m$, $f(E_j) = E_j \odot \tau^n$, and $n > m > 0$. If it is ρ (resp. $\bar{\rho}$) that is of this form, then proper initial (resp. terminal) sub edge-paths of ρ are called Ψ EPs (*pseudo-exceptional paths*).

Let $f : G \rightarrow G$ be an improved relative train track map and $r, J \geq 1$ integers. Then *r-monochromatic* paths in G are defined by a simple recursion: edges in G are *r-monochromatic* and if ρ is an *r-monochromatic* path then every sub edge-path of $f_{\#}^r(\rho)$ is *r-monochromatic*. A (J, f) -*atom* is an *f-monochromatic* edge-path of length at most J that admits no non-vacuous hard splitting into edge-paths.

An edge-path ρ is (J, f) -*beaded* if it admits a hard splitting $\rho = \rho_1 \odot \cdots \odot \rho_k$ where each ρ_i is a GEP, a Ψ EP, a (J, f) -atom, or an indivisible Nielsen path of length at most J (where GEPs, Ψ EPs and Nielsen paths are defined with respect to the map f).

The following is the main result of this paper.

Beaded Decomposition Theorem: *For every $\phi \in \text{Out}(F_n)$, there exist positive integers k, r and J such that ϕ^k has an improved relative train-track representative $f : G \rightarrow G$ with the property that every *r-monochromatic* path in G is (J, f) -beaded.*

(See Subsection 1.2 below for a precise description of what we mean by the map $f_{\#}^r$.)

As is clear from the preceding discussion, our main motivation for developing the Beaded Decomposition is its application in [6]. The import of the current paper in [6] has been deliberately distilled into this single statement, and the technical addenda in Section 8, so that a reader who is willing to accept these as articles of faith can proceed directly from [5] to [6].

We expect that our particular refinement of the train-track technology may prove useful in other contexts. This expectation stems from the general point that the development of refined topological representatives leads to insights into purely algebraic questions about free-group automorphisms. See [7] for a concrete illustration of this.¹

¹[7] contains some results about the growth of words under iterated automorphisms. A previous version of this paper contained an incorrect version of these results. We thank Gilbert Levitt for pointing out our error.

CONTENTS

1. Improved relative train track maps	4
2. Hard splittings	10
3. A small reduction	13
4. Nibbled futures	14
5. Passing to an iterate of f	23
6. The nibbled futures of GEPs	25
7. Proof of the Beaded Decomposition Theorem	31
8. Refinements of the Main Theorem	31
References	35

1. IMPROVED RELATIVE TRAIN TRACK MAPS

In this section we collect and refine those elements of the train-track technology that we shall need. Most of the material here is drawn directly from [4] and [1].

The philosophy behind train tracks is to find an *efficient* topological representative for an outer automorphism of F . Precisely what it means for a graph map to be *efficient* is spelled out in this section.

1.1. Paths, Splittings, Turns and Strata. Let G be a graph. Following [1], we try to reserve the term *path* for a map $\sigma : [0, 1] \rightarrow G$ that is either constant or an immersion (i.e. *tight*). The reverse path $t \mapsto \sigma(1 - t)$ will be denoted $\bar{\sigma}$. We conflate the map σ with its monotone reparameterisations (and even its image, when this does not cause confusion). Given an arbitrary continuous map $\rho : [0, 1] \rightarrow G$, we denote by $[\rho]$ the unique (tight) *path* homotopic rel endpoints to ρ . In keeping with the notation of the previous section, given $f : G \rightarrow G$ and a path σ in G , we write $f_{\#}(\sigma)$ to denote $[f(\sigma)]$. We are primarily concerned with *edge paths*, i.e. those paths σ for which $\sigma(0)$ and $\sigma(1)$ are vertices.

We consider only maps $f : G \rightarrow G$ that send vertices to vertices and edges to edge-paths (not necessarily to single edges). If there is an isomorphism $F \cong \pi_1 G$ such that f induces $\mathcal{O} \in \text{Out}(F)$, then one says that f *represents* \mathcal{O} . **Notation:** Given a map f of graphs and a path ρ in the domain, we'll follow the standard practice of denoting by $f_{\#}(\rho)$ the unique locally-injective edge path that is homotopic rel endpoints to $f(\rho)$.

1.2. Replacing f by an Iterate. In order to obtain good topological representatives of outer automorphisms, one has to replace the given map by a large iterate. It is important to be clear what one means by *iterate* in this context, since we wish to consider only topological representatives whose restriction to

each edge is an immersion and this property is not inherited by (naive) powers of the map.

Thus we deem the phrase² *replacing f by an iterate*, to mean that for fixed $k \in \mathbb{N}$, we pass from consideration of $f : G \rightarrow G$ to consideration of the map $f_{\#}^k : G \rightarrow G$ that sends each edge E in G to the tight edge-path $f_{\#}^k(E)$ that is homotopic rel endpoints to $f^k(E)$.

1.3. (Improved) Relative train tracks. We now describe the properties of Improved Relative Train Track maps, as constructed in [4] and [1].

Splitting: Suppose that $\sigma = \sigma_1\sigma_2$ is a decomposition of a path into nontrivial subpaths (we do not assume that σ_1 and σ_2 are edge-paths, even if σ is). We say that $\sigma = \sigma_1\sigma_2$ is a k -*splitting* if

$$f_{\#}^k(\sigma) = f_{\#}^k(\sigma_1)f_{\#}^k(\sigma_2)$$

is a decomposition into sub-paths (i.e. for *some* choice of tightening, there is no folding between the f -images of σ_1 and σ_2 for at least k iterates). If $\sigma = \sigma_1\sigma_2$ is a k -splitting for all $k > 0$ then it is called a *splitting*³ and we write $\sigma = \sigma_1 \cdot \sigma_2$. If one of σ_1 or σ_2 is the empty path, the splitting is said to be *vacuous*.

A *turn* in G is an unordered pair of half-edges originating at a common vertex. A turn is *non-degenerate* if it is defined by distinct half-edges, and is *degenerate* otherwise. The map $f : G \rightarrow G$ induces a self-map Df on the set of oriented edges of G by sending an oriented edge to the first oriented edge in its f -image. Df induces a map Tf on the set of turns in G .

A turn is *illegal* with respect to $f : G \rightarrow G$ if its image under some iterate of Tf is degenerate; a turn is *legal* if it is not illegal.

Associated to f is a *filtration* of G ,

$$\emptyset = G_0 \subset G_1 \subset \cdots \subset G_{\omega} = G,$$

consisting of f -invariant subgraphs of G . We call the sets $H_r := \overline{G_r \setminus G_{r-1}}$ *strata*. To each stratum H_r is associated M_r , the *transition matrix for H_r* ; the (i, j) th entry of M_r is the number of times the f -image of the j th edge crosses the i th edge in either direction. By choosing a filtration carefully one may ensure that for each r the matrix M_r is either the zero matrix or is irreducible. If M_r is the zero matrix, then we say that H_r is a *zero stratum*. Otherwise, M_r has an associated Perron-Frobenius eigenvalue $\lambda_r \geq 1$, see [8]. If $\lambda_r > 1$ then we say that H_r is an *exponential stratum*; if $\lambda_r = 1$ then we say that

²and obvious variations on it

³In the next section, we introduce a stronger notion of *hard* splittings.

H_r is a *parabolic stratum*⁴. The edges in strata inherit these adjectives, e.g. “exponential edge”.

A turn is defined to be *in* H_r if both half-edges lie in the stratum H_r . A turn is a *mixed turn in* (G_r, G_{r-1}) if one edge is in H_r and the other is in G_{r-1} . A path with no illegal turns in H_r is said to be *r-legal*. We may emphasize that certain turns are in H_r by calling them *r-(il)legal turns*.

Definition 1.1. [4, Section 5, p.38] We say that $f : G \rightarrow G$ is a *relative train track map* if the following conditions hold for every exponential stratum H_r :

- (RTT-i) Df maps the set of oriented edges in H_r to itself; in particular all mixed turns in (G_r, G_{r-1}) are legal.
- (RTT-ii) If α is a nontrivial path in G_{r-1} with endpoints in $H_r \cap G_{r-1}$, then $f_{\#}(\alpha)$ is a nontrivial path with endpoints in $H_r \cap G_{r-1}$.
- (RTT-iii) For each legal path β in H_r , $f(\beta)$ is a path that does not contain any illegal turns in H_r .

The following lemma is “the most important consequence of being a relative train track map” [1, p.530]; it follows immediately from Definition 1.1.

Lemma 1.2. [4, Lemma 5.8, p.39] *Suppose that $f : G \rightarrow G$ is a relative train track map, that H_r is an exponential stratum and that $\sigma = a_1 b_1 a_2 \dots b_l$ is the decomposition of an r -legal path σ into subpaths a_j in H_r and b_j in G_{r-1} . (Allow for the possibility that a_1 or b_l is trivial, but assume the other subpaths are nontrivial.) Then $f_{\#}(\sigma) = f(a_1) f_{\#}(b_1) f(a_2) \dots f_{\#}(b_l)$ and is r -legal.*

Definition 1.3. Suppose that $f : G \rightarrow G$ is a topological representative, that the parabolic stratum H_i consists of a single edge E_i and that $f(E_i) = E_i u_i$ for some path u_i in G_{i-1} . We say that the paths of the form $E_i \gamma \bar{E}_i$, $E_i \gamma$ and $\gamma \bar{E}_i$, where γ is in G_{i-1} , are *basic paths of height i* .

Lemma 1.4. [1, Lemma 4.1.4, p.555] *Suppose that $f : G \rightarrow G$ and E_i are as in Definition 1.3. Suppose further that σ is a path or circuit in G_i that intersects H_i nontrivially and that the endpoints of σ are not contained in the interior of E_i . Then σ has a splitting each of whose pieces is either a basic path of height i or is contained in G_{i-1} .*

Definition 1.5. A *Nielsen path* is a nontrivial path σ such that $f_{\#}^k(\sigma) = \sigma$ for some $k \geq 1$.

Nielsen paths are called *periodic Nielsen paths* in [1], but Theorem 1.8 below allows us to choose an f so that any periodic Nielsen path has period 1 (which is to say that $f_{\#}(\sigma) = \sigma$), and we shall assume that f satisfies the properties outlined in Theorem 1.8. Thus we can assume that $k = 1$ in the

⁴Bestvina *et al.* use the terminology *exponentially-growing* and *non-exponentially-growing* for our exponential and parabolic. This difference in terminology explains the names of the items in Theorem 1.8 below.

above definition. A Nielsen path is called *indivisible* if it cannot be split as a concatenation of two non-trivial Nielsen paths.

Definition 1.6 (cf. 5.1.3, p.531 [1]). Suppose that H_i is a single edge E_i and that $f(E_i) = E_i\tau^l$ for some closed Nielsen path τ in G_{i-1} and some $l > 0$. The *exceptional paths of height i* are those paths of the form $E_i\tau^k\bar{E}_j$ or $E_i\bar{\tau}^k\bar{E}_j$ where $k \geq 0$, $j \leq i$, H_j is a single edge E_j and $f(E_j) = E_j\tau^m$ for some $m > 0$.

Remark 1.7. In [1] the authors require that the path τ is an indivisible Nielsen path. However, exceptional paths are defined so that condition ne-(iii) of Theorem 1.8 holds, and an analysis of the proof of this theorem in [1] shows that the restriction to indivisible Nielsen paths in exceptional paths is invalid.

In Definition 1.6, the paths do not have a preferred orientation. Thus it is important to note that the paths of the form $E_j\tau^k\bar{E}_i$ and $E_j\bar{\tau}^k\bar{E}_i$ with E_i, E_j and τ as above are also exceptional paths of height i .

1.4. The Theorem of Bestvina, Feighn and Handel. A matrix is *aperiodic* if it has a power in which every entry is positive. The map f is *eg-aperiodic* if every exponential stratum has an aperiodic transition matrix.

Theorem 5.1.5 in [1] is the main structural theorem for improved relative train track maps. We shall use it continually in what follows, often without explicit mention. We therefore record those parts of it which we need. A map f which satisfies the statements of Theorem 1.8 is called an *improved relative train track map*.

Theorem 1.8. (cf. Theorem 5.1.5, p.562, [1]) *For every outer automorphism $\mathcal{O} \in \text{Out}(F)$ there is an eg-aperiodic relative train track map $f : G \rightarrow G$ with filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_\omega = G$ such that f represents an iterate of \mathcal{O} , and f has the following properties.*

- *Every periodic Nielsen path has period one.*
- *For every vertex $v \in G$, $f(v)$ is a fixed point. If v is an endpoint of an edge in a parabolic stratum then v is a fixed point. If v is the endpoint of an edge in an exponential stratum H_i and if v is also contained in a noncontractible component of G_{i-1} , then v is a fixed point.*
- *H_i is a zero stratum if and only if it is the union of the contractible components of G_i .*
- *If H_i is a zero stratum, then*
 - z-(i) *H_{i+1} is an exponential stratum.*
 - z-(ii) *$f|H_i$ is an immersion.*
- *If H_i is a parabolic stratum, then*
 - ne-(i) *H_i is a single edge E_i .*
 - ne-(ii) *$f(E_i)$ splits as $E_i \cdot u_i$ for some closed path u_i in G_{i-1} whose base-point is fixed by f .*

- ne-(iii) *If σ is a basic path of height i that does not split as a concatenation of two basic paths of height i or as a concatenation of a basic path of height i with a path contained in G_{i-1} , then either: (i) for some k , the path $f_{\#}^k(\sigma)$ splits into pieces, one of which equals E_i or \bar{E}_i ; or (ii) u_i is a Nielsen path and, for some k , the path $f_{\#}^k(\sigma)$ is an exceptional path of height i .*
- *If H_i is an exponential stratum then*
 - eg-(i) *There is at most one indivisible Nielsen path ρ_i in G_i that intersects H_i nontrivially. The initial edges of ρ_i and $\bar{\rho}_i$ are distinct (possibly partial) edges in H_i .*

Suppose that $f : G \rightarrow G$ is an improved relative train track map representing some iterate ϕ^k of $\phi \in \text{Out}(F_n)$, and that ρ is a Nielsen path in G_r that intersects H_r nontrivially, and suppose that ρ is not an edge-path. Then subdividing the edges containing the endpoints of ρ at the endpoints, gives a new graph G' , and the map $f' : G' \rightarrow G'$ induced by f is an improved relative train track map representing ϕ^k . To ease notation, it is convenient to assume that this subdivision has been performed. Under this assumption, all Nielsen paths will be edge-paths, and all of the paths which we consider in the remainder of this paper will also be edge-paths.

Convention 1.9. Since all Nielsen paths in the remainder of this paper will be edge paths, we will use the phrase ‘*indivisible Nielsen path*’ to mean a Nielsen edge-path which cannot be decomposed nontrivially as a concatenation of two non-trivial Nielsen *edge*-paths. In particular, a single edge fixed pointwise by f will be considered to be an indivisible Nielsen path.

For the remainder of this article, we will concentrate on an improved relative train track map $f : G \rightarrow G$ and repeatedly pass to iterates $f_{\#}^k$ in order to better control its cancellation properties.

Recall the following from [1, Section 4.2, pp.558-559].

Definition 1.10. If $f : G \rightarrow G$ is a relative train track map and H_r is an exponential stratum, then define P_r to be the set of paths ρ in G_r that are such that:

- (i) For each $k \geq 1$ the path $f_{\#}^k(\rho)$ contains exactly one illegal turn in H_r .
- (ii) For each $k \geq 1$ the initial and terminal (possibly partial) edges of $f_{\#}^k(\rho)$ are contained in H_r .
- (iii) The number of H_r -edges in $f_{\#}^k(\rho)$ is bounded independently of k .

Lemma 1.11. [1, Lemma 4.2.5, p.558] *P_r is a finite $f_{\#}$ -invariant set.*

Lemma 1.12. [1, Lemma 4.2.6, p.559] *Suppose that $f : G \rightarrow G$ is a relative train track map, that H_r is an exponential stratum, that σ is a path or circuit in G_r and that, for each $k \geq 0$, the path $f_{\#}^k(\sigma)$ has the same finite number of*

illegal turns in H_r . Then σ can be split into subpaths that are either r -legal or elements of P_r .

Definition 1.13. If ρ is a path and r is the least integer such that ρ is in G_r then we say that ρ has weight r .

If ρ has weight r and H_r is exponential, we will say that ρ is an *exponential path*. We define *parabolic paths* similarly.

Lemma 1.14. Suppose that σ is an edge-path and that, for some $k \geq 1$, $f_{\#}^k(\sigma)$ is a Nielsen path. Then $f_{\#}(\sigma)$ is a Nielsen path.

Proof. Suppose that the endpoints of σ are u_1 and v_1 and that the endpoints of $f_{\#}^k(\sigma)$ are u_2 and v_2 . For each vertex $v \in G$, $f(v)$ is fixed by f , so $f(u_1) = u_2$ and $f(v_1) = v_2$. If $f_{\#}(\sigma) \neq f_{\#}^k(\sigma)$ then we have two edge-paths with the same endpoints which eventually get mapped to the same path. Thus there is some nontrivial circuit which is killed by f , contradicting the fact that f is a homotopy equivalence. Therefore $f_{\#}(\sigma) = f_{\#}^k(\sigma)$ and so is a Nielsen path. \square

Always, L will denote the maximum of $|f(E)|$, for E an edge in G .

Later, we will pass to further iterates of f in order to find a particularly nice form.

An analysis of the results in this section allows us to see that there are three kinds of irreducible Nielsen paths. The first are those which are single edges; the second are certain exceptional paths; and the third lie in the set P_r . We will use this trichotomy frequently without mention. The first two cases are where the path is parabolic-weight, the third where it is exponential-weight. It is not possible for Nielsen path to have weight r where H_r is a zero stratum.

Observation 1.15. Let ρ be an indivisible Nielsen path of weight r . Then the first and last edges in ρ are contained in H_r .

Because periodic Nielsen paths have period 1, the set of Nielsen paths does not change when f is replaced by a further iterate of itself. We will use this fact often.

Lemma 1.16. Suppose E is an edge such that $|f_{\#}^j(E)|$ grows linearly with j . Then $f(E) = E \cdot \tau^k$, where τ is a Nielsen path that is not a proper power. The edge-path τ decomposes into indivisible Nielsen paths each of which is itself an edge-path.

Proof. The fact that $f(E) = E \cdot \tau^k$, where τ is a Nielsen path follows from conditions ne-(ii) and ne-(iii) of Theorem 1.8. ⁵

The final sentence follows because τ is an edge-path. Any Nielsen path admits a splitting into indivisible Nielsen paths. If there were an indivisible

⁵If Theorem 1.8, ne-(iii) held with the Nielsen path τ in the definition of exceptional paths being indivisible, we could also insist that τ be indivisible here.

Nielsen path in this decomposition which was not an edge-path then it would have to be of exponential weight, and there is at most one indivisible Nielsen path of each exponential weight. Therefore, the final Nielsen path of this weight would end at a half-edge, and the remainder of this edge could not be contained in an indivisible Nielsen path, contradicting the decomposition of τ into indivisible Nielsen paths. \square

Lemma 1.17. *Let τ be a Nielsen path and τ_0 a proper initial (or terminal) sub edge-path of τ . No image $f_{\#}^k(\tau_0)$ contains τ as a sub-path.*

Proof. It is sufficient to prove the lemma for indivisible Nielsen paths, as the result for arbitrary Nielsen paths then follows immediately.

If τ is an indivisible Nielsen path and τ_0 is a proper non-trivial subpath of τ then τ cannot be a single edge. Therefore, either τ is either an indivisible Nielsen path of exponential weight, or an exceptional path.

In case τ is an indivisible Nielsen path of exponential weight, suppose the weight is r . Then, by Lemma 1.12 τ contains a single illegal turn in H_r . If τ_0 does not contain this illegal turn then τ_0 is r -legal, and so no iterate of τ_0 contains an illegal turn in H_r , and therefore no iterate of τ_0 can contain τ as a subpath. If τ_0 does contain the r -illegal turn in τ then, being a proper subpath of τ , the path on one side of the illegal turn in τ_0 and its iterates is strictly smaller than the corresponding path in τ , and again τ cannot be contained as a subpath of any iterate of τ_0 .

If τ is an exceptional path, then $\tau = E_i \rho^k \bar{E}_j$ where ρ is a Nielsen path and E_i and E_j are of weight greater than ρ . Therefore, any sub edge-path τ_0 of τ contains at most one edge of weight greater than ρ , and the same is true for any iterate of τ_0 , and once again no iterate of τ_0 contains τ as a sub-path. \square

2. HARD SPLITTINGS

In this section we introduce a new concept for improved relative train tracks: *hard splittings*. This plays an important role in the subsequent sections of this paper, and also in [6].

Recall that a decomposition of a path $\sigma = \sigma_1 \sigma_2$ is a k -splitting if $f_{\#}^k(\sigma) = f_{\#}^k(\sigma_1) f_{\#}^k(\sigma_2)$; which means that, for *some* choice of tightening, the images of σ_1 and σ_2 do not interact with each other. This leads to the concept of *splittings*. We need a more restrictive notion, where the decomposition is preserved for *every* choice of tightening. For this purpose, we make the following

Definition 2.1 (Hard splittings). We say that a k -splitting $\rho = \rho_1 \rho_2$ is a *hard k -splitting* if for *any* choice of tightening of $f^k(\rho) = f^k(\rho_1) f^k(\rho_2)$ there is no cancellation between the image of ρ_1 and the image of ρ_2 .

A decomposition which is a hard k -splitting for all $k \geq 1$ is called a *hard splitting*. If $\rho_1 \cdot \rho_2$ is a hard splitting, we write $\rho_1 \odot \rho_2$.

An edge path is *hard-indivisible* (or *h-indivisible*) if it admits no non-vacuous hard splitting into edge paths.

Example 2.2. Suppose that G is the graph with a single vertex and edges E_1, E_2 and E_3 . Suppose that $f(E_1) = E_1$, $f(E_2) = E_2E_1$ and $f(E_3) = E_3\bar{E}_1\bar{E}_2$. Then f is an improved relative train track. Then $E_3E_2 \cdot \bar{E}_1$ is a 1-splitting, since

$$f(E_3E_2\bar{E}_1) = E_3\bar{E}_1\bar{E}_2E_2E_1\bar{E}_1,$$

which tightens to $E_3\bar{E}_1 = f_{\#}(E_3E_2)f_{\#}(\bar{E}_1)$. In fact this is a splitting. However, there is a choice of tightening which first cancels the final $E_1\bar{E}_1$ and then the subpath \bar{E}_2E_2 . Therefore the splitting $E_3E_2 \cdot \bar{E}_1$ is not a hard 1-splitting.

The following lemma shows the main utility of hard splittings, and the example above shows that it is not true in general for splittings.

Lemma 2.3. *Suppose that $\sigma_1 \odot \sigma_2$ is a hard splitting, and that ρ is an initial subpath of σ_2 . Then $\sigma_1 \odot \rho$ is a hard splitting.*

Proof. If there were any cancellation between images of σ_1 and ρ then there would be a possible tightening between the images of σ_1 and σ_2 . \square

The following two lemmas will also be crucial for our applications of hard splittings in [6].

Lemma 2.4. *Every edge path admits a unique maximal hard splitting into edge paths.*

Proof. This follows by an obvious induction on length from the observation that if $\rho = \rho_1\rho_2\rho_3$, where the ρ_i are edge paths, and if $\rho = \rho_1 \odot \rho_2\rho_3$ and $\rho = \rho_1\rho_2 \odot \rho_3$ then $\rho = \rho_1 \odot \rho_2 \odot \rho_3$. \square

Lemma 2.5. *If $\rho = \rho_1 \odot \rho_2$ and σ_1 and σ_2 are, respectively, terminal and initial subpaths of $f_{\#}^k(\rho_1)$ and $f_{\#}^k(\rho_2)$ for some $k \geq 0$ then $\sigma_1\sigma_2 = \sigma_1 \odot \sigma_2$.*

Proof. For all $i \geq 0$, the untightened path $f^i(\sigma_1)$ is a terminal subpath of the untightened path $f^{k+i}(\rho_1)$, while $f^i(\sigma_2)$ is an initial subpath of $f^{k+i}(\rho_2)$.

The hardness of the splitting $\rho = \rho_1 \odot \rho_2$ ensures that no matter how one tightens $f^{k+i}(\rho_1)f^{k+i}(\rho_2)$ there will be no cancellation between $f^{k+i}(\rho_1)$ and $f^{k+i}(\rho_2)$. In particular, one is free to tighten $f^i(\sigma_1)f^i(\sigma_2)$ first, and there can be no cancellation between them. (It may happen that when one goes to tighten $f^{k+i}(\rho_1)$ completely, the whole of $f^i(\sigma_1)$ is cancelled, but this does not affect the assertion of the lemma.) \square

The purpose of the remainder of this section is to sharpen results from the previous section to cover hard splittings. ⁶

⁶Bestvina *et al.* make no explicit mention of the distinction between splittings and hard splittings, however condition (3) of Proposition 5.4.3 on p.581 (see Lemma 2.8 below) indicates that they are aware of the distinction and that the term ‘splitting’ has the same meaning for them as it does here.

Lemma 2.6 (cf. Lemma 4.1.1, p.554 [1]). *If $\sigma = \sigma_1 \odot \sigma_2$ is a hard splitting, and $\sigma_1 = \sigma'_1 \odot \sigma'_2$ is a hard splitting then $\sigma = \sigma'_1 \odot \sigma'_2 \odot \sigma_2$ is a hard splitting. The analogous result with the roles of σ_1 and σ_2 reversed also holds.*

Remark 2.7. The possible existence of an edge-path σ_2 so that $f_{\#}(\sigma_2)$ is a single vertex means that $\sigma_1\sigma_2 = \sigma_1 \odot \sigma_2$ and $\sigma_2\sigma_3 = \sigma_2 \odot \sigma_3$ need *not* imply that $\sigma_1\sigma_2\sigma_3 = \sigma_1 \odot \sigma_2 \odot \sigma_3$.

Indeed if σ_2 is an edge-path so that $f_{\#}(\sigma_2)$ is a vertex then $f_{\#}(\sigma_1)$ and $f_{\#}(\sigma_3)$ come together in a tightening of $f(\sigma_1\sigma_2\sigma_3)$, possibly cancelling.

In contrast, if $f_{\#}(\sigma_2)$ (and hence each $f_{\#}^k(\sigma_2)$) contains an edge, then the hardness of the two splittings ensures that in any tightening $f_{\#}(\sigma_1\sigma_2\sigma_3) = f_{\#}(\sigma_1)f_{\#}(\sigma_2)f_{\#}(\sigma_3)$, that is $\sigma_1\sigma_2\sigma_3 = \sigma_1 \odot \sigma_2 \odot \sigma_3$.

The following strengthening of Theorem 1.8 ne-(ii) is a restatement of (a weak form of) [1, Proposition 5.4.3.(3), p.581].

Lemma 2.8. *Suppose f is an improved relative train track map and E is a parabolic edge with $f(E) = Eu$. For any initial subpath w of u , $E \cdot w$ is a splitting.*

Corollary 2.9. *Suppose f is an improved relative train track map, E is a parabolic edge and $f(E) = Eu$. Then $f(E) = E \odot u$.*

Lemma 2.10. *Suppose H_i is a parabolic stratum and σ is a path in G_i that intersects H_i nontrivially, and that the endpoints of σ are not contained in the interior of E_i . Then σ admits a hard splitting, each of whose pieces is either a basic path of height i or is contained in G_{i-1} .*

Lemma 2.11. *If σ is a basic path of height i that does not admit a hard splitting as a concatenation of two basic paths of height i or as a concatenation of a basic path of height i with a path of weight less than i , then either; (i) for some k , the path $f_{\#}^k(\sigma)$ admits a hard splitting into pieces, one of which is E_i or \bar{E}_i ; or (ii) $f(E_i) = E_i \odot u_i$, where u_i is a Nielsen path and, for some k , the path $f_{\#}^k(\sigma)$ is an exceptional path of height i .*

Proof. Follows from the proof of [1, Lemma 5.5.1, pp.585–590]. \square

Lemma 2.12 (cf. Lemma 1.12 above). *Suppose that $f : G \rightarrow G$ is a relative train track map, that H_r is an exponentially-growing stratum, that σ is a path or circuit in G_r , and that each $f_{\#}^k(\sigma)$ has the same finite number of illegal turns in H_r . Then σ can be decomposed as $\sigma = \rho_1 \odot \dots \odot \rho_k$, where each ρ_i is either (i) an element of P_r ; (ii) an r -legal path which starts and ends with edges in H_r ; or (iii) of weight at most $r - 1$.*

Proof. Consider the splitting of σ given by Lemma 1.12. The pieces of this splitting are either (i) elements of P_r , or (ii) r -legal paths. By Definition 1.1 RTT-(i), any r -legal path admits a hard splitting into r -legal paths which start

and end with edges in H_r , and paths of weight at most $r - 1$. The turn at the end of a Nielsen path in the splitting of σ is either a mixed turn (with the edge from H_r coming from the Nielsen path and the other edge being of weight at most $r - 1$) or a legal turn in H_r . In either case, σ admits a hard splitting at the vertex of this turn. \square

The next result follows from a consideration of the form of indivisible Nielsen paths, noting Definition 1.1 and Lemma 2.12.

Lemma 2.13. *Any Nielsen path admits a hard splitting into indivisible Nielsen paths.*

Remark 2.14. If $\rho = \rho_1 \odot \rho_2$ is a hard splitting for the map f then it is a hard splitting for $f_{\#}^k$ for any $k \geq 1$.

We record a piece of terminology for later use.

Definition 2.15. A sub edge-path ρ of a path χ is *displayed* if there is a hard splitting of χ immediately on either side of ρ .

3. A SMALL REDUCTION

In this section we clarify a couple of issues about monochromatic paths, and state Theorem 3.2, which immediately implies the Beaded Decomposition Theorem.

Our strategy for proving the Beaded Decomposition Theorem is as follows: given an automorphism $\phi \in \text{Aut}(F_n)$, we start with an improved relative train track representative $f : G \rightarrow G$ for some iterate ϕ^k of ϕ , as obtained from the conclusion of Theorem 1.8. We analyse the evolution of monochromatic paths, and eventually pass to an iterate of f in which we can prove the Beaded Decomposition Theorem. However, it is crucial to note that monochromatic paths for f are not necessarily monochromatic paths for $f_{\#}^k$ when $k > 1$. See Section 5 for further discussion about some of these issues.

These concerns lead to the following definition, where we are concentrating on a fixed IRTT $f : G \rightarrow G$, and so omit mention of f from our notation.

Definition 3.1. For a positive integer r , we define *r -monochromatic paths* by recursion: edges in G are r -monochromatic and if ρ is an r -monochromatic path then every sub edge-path of $f_{\#}^r(\rho)$ is r -monochromatic.

Note that if r' is a multiple of r then every r' -monochromatic path is r -monochromatic but not *vice versa*. Thus if we replace f by an iterate then, for fixed n , the set of n -monochromatic paths may get smaller. The content of the Beaded Decomposition Theorem is that one need only pass to a bounded iterate in order to ensure that all monochromatic paths admit a beaded decomposition. In particular, the Beaded Decomposition Theorem is an immediate consequence of the following theorem.

Theorem 3.2 (Monochromatic paths are beaded). *Let $f : G \rightarrow G$ be an improved relative train track map. There exist constants r and J , depending only on f , so that every r -monochromatic path in G is (J, f) -beaded.*

4. NIBBLED FUTURES

Monochromatic paths arise as *nibbled futures* in the sense defined below. Thus in order to prove Theorem 3.2 we must understand how *nibbled futures* evolve. The results in this section reduce this challenge to the task of understanding the nibbled futures of GEPs.

Definition 4.1 (Nibbled Futures). Let ρ be a (tight) edge path. The *0-step nibbled future of ρ* is ρ .

For $k \geq 1$, a *k -step nibbled future of ρ* is a sub edge-path of $f_{\#}(\sigma)$, where σ is a $(k-1)$ -step nibbled future of ρ . A *nibbled future of ρ* is a k -step nibbled future for some $k \geq 0$.

For $k \geq 0$, the *k -step entire future of ρ* is $f_{\#}^k(\rho)$.

Remark 4.2. The 1-monochromatic paths are precisely the nibbled futures of single edges.

Theorem 4.3 (First Decomposition Theorem). *For any $n \geq 1$ there exists an integer $V = V(n, f)$ such that if ρ is an edge path of length at most n then any nibbled future of ρ admits a hard splitting into edge paths, each of which is either the nibbled future of a GEP or else has length at most V .*

The remainder of this section is dedicated to proving Theorem 4.3. We begin by examining the entire future of a path of fixed length (Lemma 4.5) and then refine the argument to deal with nibbling. In the proof of the first of these lemmas we require the following observation.

Remark 4.4. Suppose that ρ is a tight path of weight r . The immediate past of an r -illegal turn in $f_{\#}(\rho)$ (under any choice of tightening) is an r -illegal turn, and two r -illegal turns cannot have the same r -illegal turn as their past. In particular, the number of r -illegal turns in $f_{\#}^l(\rho)$ is a non-increasing function of l , bounded below by 0.

Lemma 4.5. *There is a function $D : \mathbb{N} \rightarrow \mathbb{N}$, depending only on f , such that, for any $r \in \{1, \dots, \omega\}$, if ρ is a path of weight r , and $|\rho| \leq n$, then for any $i \geq D(n)$ the edge path $f_{\#}^i(\rho)$ admits a hard splitting into edge paths, each of which is either*

- (1) a single edge of weight r ;
- (2) an indivisible Nielsen path of weight r ;
- (3) a GEP of weight r ; or
- (4) a path of weight at most $r-1$.

Proof. If H_r is a zero stratum, then $f_{\#}(\rho)$ has weight at most $r - 1$, and $D(n) = 1$ will suffice for any n .

If H_r is a parabolic stratum, then ρ admits a hard splitting into pieces which are either basic of height r or of weight at most $r - 1$ (Lemma 2.10). Thus it is sufficient to consider the case where ρ is a basic path of weight r and $|\rho| \leq n$. By at most 2 applications of Lemma 2.11, we see that there exists a k such that $f_{\#}^k(\rho)$ admits a hard splitting into pieces which are either (i) single edges of weight r , (ii) exceptional paths of height r , or (iii) of weight at most $r - 1$. By taking the maximum of such k over all basic paths of height r which are of length at most n , we find an integer k_0 so that we have the desired hard splitting of $f_{\#}^{k_0}(\rho)$ for all basic paths of height r of length at most n . Any of the exceptional paths in these splittings which are not GEPs have bounded length and are either indivisible Nielsen paths or are decreasing in length. A crude bound on the length of the exceptional paths which are not GEPs is $L^{k_0}n$ where L is the maximum length of $f(E)$ over all edges $E \in G$. Thus, those exceptional paths which are decreasing in length will become GEPs within less than $L^{k_0}n$ iterations. Therefore, replacing k_0 by $k_0 + L^{k_0}n$, we may assume all exceptional paths in the hard splitting are GEPs.

Finally, suppose that H_r is an exponential stratum. As noted in Remark 4.4, the number of r -illegal turns in $f_{\#}^l(\rho)$ is a non-increasing function of l bounded below by 0. Therefore, there is some j so that the number of r -illegal turns in $f_{\#}^{j'}(\rho)$ is the same for all $j' \geq j$. By Lemma 2.12, $f_{\#}^j(\rho)$ admits a hard splitting into pieces which are either (i) elements of P_r , (ii) single edges in H_r , or (iii) paths of weight at most $r - 1$. To finish the proof of the lemma it remains to note that if $\sigma \in P_r$ then $f_{\#}(\sigma)$ is a Nielsen path by Lemma 1.14.

Therefore, the required constant for H_r may be taken to be the maximum of $j + 1$ over all the paths of weight r of length at most n .

To find $D(n)$ we need merely take the maximum of the constants found above over all of the strata H_r of G . \square

In the extension of the above proof to cover nibbled futures, we shall need the following straightforward adaptation of Lemma 1.17.

Lemma 4.6. *Let τ be a Nielsen path and τ_0 a proper initial (or terminal) sub-path of τ . No nibbled future of τ_0 contains τ as a sub-path.*

Proposition 4.7. *There exists a function $D' : \mathbb{N} \rightarrow \mathbb{N}$, depending only on f , so that for any $r \in \{1, \dots, \omega\}$, if ρ is a path of weight r and $|\rho| \leq n$, then for any $i \geq D'(n)$ any i -step nibbled future of ρ admits a hard splitting into edge paths, each of which is either*

- (1) a single edge of weight r ;
- (2) a nibbled future of a weight r indivisible Nielsen path;
- (3) a nibbled future of a weight r GEP; or
- (4) a path of weight at most $r - 1$.

Remark 4.8. Each of the conditions (1) – (4) stated above is stable in the following sense: once an edge in a k -step nibbled future lies in a path satisfying one of these conditions, then any future of this edge in any further nibbled future will also lie in such a path (possibly the future will go from case (1) to case (4), but otherwise which case it falls into is also stable). Thus we can split the proof of Proposition 4.7 into a number of cases, deal with the cases separately by finding some constant which suffices, and finally take a maximum to find $D'(n)$. An entirely similar remark applies to a number of subsequent proofs, in particular Theorem 8.1.

Proof (Proposition 4.7). Let $\rho_0 = \rho$ and for $j > 0$ let ρ_j be a sub edge-path of $f_{\#}(\rho_{j-1})$.

If H_r is a zero stratum, then $f_{\#}(\rho)$ has weight at most $r - 1$ and it suffices to take $D'(n) = 1$.

Suppose that H_r is an exponential stratum. By Lemma 4.5, the $D(n)$ -step entire future of ρ admits a hard splitting of the desired form. We consider how nibbling can affect this splitting. As we move forwards through the nibbled future of ρ , cancellation of H_r -edges can occur only at r -illegal turns and at the ends, where the nibbling occurs.

Remark 4.4 implies that we can trace the r -illegal turns forwards through the successive nibbled futures of ρ (whilst the r -illegal continues to exist). We compare the r -illegal turns in ρ_k to those in $f_{\#}^k(\rho)$, the entire future of ρ . We say that the nibbling *first cancels an r -illegal turn at time k* if the collection of r -illegal turns in ρ_{k-1} is the same as the collection in $f_{\#}^{k-1}(\rho)$, but the collection in ρ_k is *not* the same as that of $f_{\#}^k(\rho)$. The first observation we make is that if, at time k , the nibbling has not yet cancelled any r -illegal turn then the sequence of H_r -edges ρ_k is a subsequence of the H_r -edges in $f_{\#}^k(\rho)$. Therefore, any splitting of the desired type for $f_{\#}^k(\rho)$ is inherited by ρ_k .

Since there is a splitting of the $D(n)$ -step entire future of ρ of the desired form, either there is a splitting of $\rho_{D(n)}$, or else $\rho_{D(n)}$ has fewer r -illegal turns than $f_{\#}^{D(n)}(\rho)$, and hence than ρ . However, $|\rho_{D(n)}| \leq n \cdot L^{D(n)}$. We apply the above argument to $\rho_{D(n)}$, going forwards a further $D(nL^{D(n)})$ steps into the future. Since the number of illegal turns in H_r in ρ was at most $n - 1$, we will eventually find a splitting of the required form within an amount of time bounded by a function of n (this function depends only on f , as required). Denoting this function by D_0 , we have that any $D_0(n)$ -step nibbled future of any path of exponential weight whose length is at most n admits a hard splitting of the desired form.

Now suppose that H_r is a parabolic stratum. By Lemma 2.10, ρ admits a hard splitting into basic edge paths. Therefore we may assume (by reversing the orientation of ρ if necessary) that $\rho = E_r\sigma$ or $\rho = E_r\sigma\overline{E_r}$ where E_r is the unique edge in H_r and σ is in G_{r-1} . For the nibbled future of ρ to have weight

r , the nibbling must occur only on one side (since the only edges of weight r in any future of ρ occur on the ends). We assume that all nibbling occurs from the right. Once again, the $D(n)$ -step entire future of ρ admits a hard splitting of the desired form. If $\rho = E_r \sigma \overline{E_r}$ then the $D(n)$ -step nibbled future of ρ either admits a hard splitting of the required form, or is of the form $E_r \sigma_1$, where σ_1 is in G_{r-1} . Hence we may assume that $\rho = E_r \sigma$. Suppose that $f(E_r) = E_r u_r$, and that u_r has weight $s < r$.

We first consider the possibility that $f_{\#}(\sigma)$ has weight $q > s$ (but less than r by hypothesis). There are two cases to consider here. The first is that H_q is an exponential stratum. The future of E_r cannot cancel any edges of weight q or higher in the future of σ , so the edges of weight q in the nibbled future of ρ are exactly the same as the edges of weight q in the corresponding nibbled future of σ (recall we are assuming that nibbling only occurs from the right). This $D_0(|\sigma|)$ -step nibbled future of σ admits a hard splitting into edge paths which are either⁷ single edges of weight q , the nibbled future of an indivisible Nielsen path of weight q , or of weight at most $q - 1$. Let σ_2 be the path from the right endpoint of $\overline{E_r}$ up to but not including the first edge of weight q . Then, since mixed turns are legal, the $D_0(n)$ -step nibbled future of ρ admits a hard splitting into edge paths, the leftmost of which is $E_r \sigma_2$.

Suppose now that H_q is a parabolic stratum. Then arguing as in Lemma 2.10, we see that ρ admits a hard splitting into edge paths, the leftmost of which is either $E_r \sigma_2$ or $E_r \sigma_2 \overline{E_q}$, where σ_2 has weight at most $q - 1$. Thus we may suppose that ρ itself has this form. Again, either the $D(n)$ -step nibbled future of ρ admits a hard splitting of the required form, or the $D(n)$ -step nibbled future of ρ has the form $E_r \sigma_3$, where σ_3 has weight at most $q - 1$. The above considerations cover the possibility that a GEP of weight r occurs as a factor of the hard splitting. Thus we may assume that in some nibbled future of ρ there will necessarily be a hard splitting on each side of the edge of weight r .

In this fashion, going forwards into the nibbled future an amount of time bounded by a function of n , we may assume that ρ has the form $E_r \sigma_4$, where σ_4 has weight exactly s (if σ_4 has weight less than s then $f_{\#}(E_r \sigma_4) = E_r \odot \sigma_5$ where σ_5 has weight less than r , and this is a splitting of the required form which is inherited by the nibbled future).

We now consider what kind of stratum H_s is. Suppose that H_s is parabolic. There are only two ways in which cancellation between weight s edges in the nibbled future of ρ can occur (see [5, Lemma 5.5]): they might be cancelled by edges whose immediate past is the edge of weight r on the left end of the previous nibbled future; alternatively, they can be nibbled from the right. The $D(n)$ -step entire future of ρ admits a hard splitting as $E_r \odot \sigma_6$, where σ_6 has weight at most $r - 1$. There is no way that nibbling can affect this splitting.

⁷GEPs have parabolic weight

Finally, suppose that H_s is an exponential stratum. We follow a similar argument to the case when H_r was an exponential stratum. Either the $D(n)$ -step nibbled future of ρ admits a hard splitting of the desired kind (which means $\rho_{D(n)} = E_r \odot \sigma_7$ where σ_7 has weight at most $r - 1$), or there are fewer s -illegal turns in the future of σ_4 in $\rho_{D(n)}$ than there are s -illegal turns in σ_4 . We then apply the same argument to the nibbled future of $\rho_{D(n)}$ until eventually we achieve a hard splitting of the required form. This completes the proof of Proposition 4.7. \square

We are now in a position to prove Theorem 4.3. For this we require the following definition.

Definition 4.9. Suppose that H_r is a stratum, and $E \in H_r$. An r -seed is a non-empty subpath ρ of $f(E)$ which is maximal subject to lying in G_{r-1} .

If the stratum H_r is not relevant, we just refer to *seeds*.

Note that seeds are edge-paths and that the set of all seeds is finite.

The following is an immediate consequence of Lemma 2.12 and RTT-(i) of Definition 1.1.

Lemma 4.10. *If $E \in H_r$ is an exponential edge and ρ is an r -seed in $f(E)$ then $f(E) = \sigma_1 \odot \rho \odot \sigma_2$ where σ_1 and σ_2 are r -legal paths which start and finish with edges in H_r .*

Proof (Theorem 4.3). Suppose that ρ is a path of length n and that ρ_k is a k -step nibbled future of ρ . Denote by $\rho_0 = \rho, \rho_1, \dots, \rho_{k-1}$ the intermediate nibbled futures of ρ used in order to define ρ_k .

We begin by constructing a van Kampen diagram⁸ Δ_k which encodes the ρ_i , proceeding by induction on k . For $k = 1$ the diagram Δ_1 has a single (folded) corridor with the bottom labelled by ρ and the path ρ_1 a subpath of the top of this corridor. Suppose that we have associated a van Kampen diagram Δ_{k-1} to ρ_{k-1} , with a unique corridor at each time $t = 0, \dots, k - 2$, such that ρ_{k-1} is a subpath of the top of the latest (folded) corridor. Then we attach a new folded corridor to Δ_{k-1} whose bottom is labelled by ρ_k . The path ρ_k is, by definition, a subpath of the top of this new latest corridor. By convention, we consider ρ_i to occur at time i .

Choose an arbitrary edge ε in ρ_k on the (folded) top of the latest corridor in Δ_k . We will prove that there is a path σ containing ε in ρ_k so that ρ_k admits a hard splitting immediately on either side of σ and so that σ is either suitably *short* or a nibbled future of a GEP. The purpose of this proof is to find a suitable notion of *short*.

Consider the embedded ‘family forest’ \mathcal{F} for Δ_k , tracing the histories of edges lying on the folded tops of corridors (see [5, 3.2]). Let p be the path in

⁸in fact, just a stack of corridors

\mathcal{F} which follows the history of ε . We denote by $p(i)$ the edge which intersects p and lies on the bottom of the corridor at time i . The edges $p(i)$ form the *past* of ε . We will sometimes denote the edge ε by $p(k)$. It will be an analysis of the times at which the weight of $p(i)$ decreases that forms the core of the proof of the theorem.

The weights of the edges $p(0), p(1), \dots, p(k)$ form a non-increasing sequence. Suppose this sequence is $W = \{w_0, \dots, w_k\}$. A *drop* in W is a time t such that $w_{t-1} > w_t$. At such times, the edge $p(t)$ is contained in a (folded) seed in the bottom of a corridor of Δ_k .

We will show that either successive drops occur rapidly, or else we reach a situation wherein each time a drop occurs we lose no essential information by restricting our attention to a small subpath of ρ_i .

To make this localisation argument precise, we define *incidents*, which fall into two types.

An *incident of Type A* is a time t which (i) is a drop; and (ii) is such that there is a hard splitting of ρ_t immediately on either side of the folded seed containing $p(t)$.

An *incident of Type B* is a time t such that $p(t-1)$ lies in an indivisible Nielsen path with a hard splitting of ρ_{t-1} immediately on either side, but $p(t)$ does not; except that we do not consider this to be an incident if some ρ_i , for $i \leq t-1$ admits a hard splitting $\rho_i = \sigma_1 \odot \sigma_2 \odot \sigma_3$ with $p(i) \subseteq \sigma_2$ and σ_2 a GEP. In case of an incident of Type B, necessarily $p(t)$ lies in the nibbled future of a Nielsen path on one end of ρ_t with a hard splitting of ρ_t immediately on the other side.

Define the time t_1 to be the last time at which there is an incident (of Type A or Type B). If there are no incidents, let $t_1 = 0$. If this incident is of Type A, the edge $p(t_1)$ lies in a folded seed, call it π , and there is a hard splitting of ρ_{t_1} immediately on either side of π . If the incident is of Type B, the edge $p(t_1)$ lies in the 1-step nibbled future of a Nielsen path, call this nibbled future π also. In case $t_1 = 0$, let $\pi = \rho$. We will see that there is a bound, α say, on the length of π which depends only on f and n , and not on the choice of π , or the choice of nibbled future. We postpone the proof of the existence of the bound α while we examine the consequences of its existence.

The purpose of isolating the path π is that it is a path of controlled length and the hard splitting⁹ of ρ_{t_1} immediately on either side of π means that we need only consider the nibbled future or π . Suppose that π has weight r .

Claim 1: There exists a constant $\beta = \beta(n, \alpha, f)$ so that one of the following must occur:

- (i) for some $t_1 < i < k$, the edge $p(i)$ lies in a GEP in $f_{\#}(\rho_{i-1})$ with a hard splitting immediately on either side;

⁹this splitting is vacuous in case $t_1 = 0$ and at various other points during this proof which we do not explicitly mention

- (ii) at some time $i \leq t_1 + \beta$, the edge $p(i)$ lies in an indivisible Nielsen path τ in $f_{\#}(\rho_{i-1})$ with a hard splitting immediately on either side;
- (iii) $k - t_1 \leq \beta$; or
- (iv) there is a hard splitting of ρ_k immediately on either side of ε .

This claim implies the theorem, modulo the bound on α , as we shall now explain. In case (i), for all $j \geq i$, the edge $p(j)$ lies in the nibbled future of a GEP, so in particular this is true for $\varepsilon = p(k)$. If case (ii) arises then the definition of t_1 implies that for $j \geq i$, the edge $p(j)$ always lies in a path labelled τ with a hard splitting immediately on either side, for otherwise there would be a subsequent incident. Also, the length of this Nielsen path is at most αL^β . If case (iii) arises, then the nibbled future of π at time k has length at most αL^β .

To prove the claim, we define two sequences of numbers $V_\omega, V_{\omega-1}, \dots, V_1$ and $V'_\omega, V'_{\omega-1}, \dots, V'_1$, depending on n and f , as follows (where $D'(n)$ is the function from Proposition 4.7):

$$\begin{aligned} V_\omega &:= D'(\alpha), \\ V'_\omega &:= V_\omega + \alpha L^{V_\omega}. \end{aligned}$$

For $\omega > i \geq 1$, supposing V'_{i+1} to be defined,

$$V_i := V'_{i+1} + D'(\alpha L^{V'_{i+1}}).$$

Also, supposing V_i to be defined, we define

$$V'_i := V_i + \alpha L^{V_i}.$$

We consider the situation at time $t_1 + V_r$ (recall that r is the weight of π). Possibly $k \leq t_1 + V_r$, which is covered by case (iii) of our claim. Therefore, suppose that $k > t_1 + V_r$.

According to Proposition 4.7, and the definition of t_1 , at time $t_1 + V_r$ the V_r -step nibbled future of π which exists in $\rho_{t_1+V_r}$ admits a hard splitting into edge paths, each of which is either:

- (1) a single edge of weight r ;
- (2) a nibbled future of a weight r indivisible Nielsen path;
- (3) a nibbled future of a weight r GEP; or
- (4) a path of weight at most $r - 1$.

We need to augment possibility (3) by noting that the proof of Proposition 4.7 shows that the GEP referred to lies in the j -step nibbled future of π for some $j \leq V_r$.

We analyse what happens when the edge $p(t_1 + V_r)$ lies in each of these four types of path.

Case (1): In the first case, by the definition of t_1 , there will be a hard splitting of ρ_k immediately on either side of ε , since in this case if there is

a drop in W after $t_1 + V_r$ then there is an incident of Type A, contrary to hypothesis.

Case (3): If $p(t_1 + V_r)$ lies in a path of the third type then we are in case (i) of our claim, and hence content.

The fourth type of path will lead us to an inductive argument on the weight of the path under consideration. But first we consider the nibbled futures of Nielsen paths.

Case (2): Suppose that in $\rho_{t_1+V_r}$ the edge $p(t_1+V_r)$ lies in the nibbled future of a Nielsen path of weight r , with a hard splitting of $\rho_{t_1+V_r}$ immediately on either side. Suppose that this nibbled future is π_r . If π_r is actually a Nielsen path then we lie in case (ii) of our claim. Thus suppose that π_r is not a Nielsen path. It has length at most αL^{V_r} , and within time αL^{V_r} any nibbled future of π_r admits a hard splitting into edge paths of types (1), (3) and (4) from the above list.

To see this, consider the three types of indivisible Nielsen paths. If τ is a Nielsen path which is a single edge fixed pointwise by f , then any nibbled future of τ is either a single edge or empty. Suppose that τ is an indivisible Nielsen path of weight r and H_r is exponential, and suppose that τ' is a proper subpath of τ . Then there is some iterated image $f_{\#}^l(\tau')$ of τ' which is r -legal with $l < \alpha L^{V_r}$. Finally suppose that $E_i \tau^k \overline{E_j}$ is an indivisible Nielsen path of parabolic weight. Thus τ is a Nielsen path of weight less than r , and E_i, E_j are edges such that $f(E_i) = E_i \odot \tau^m$, $f(E_j) = E_j \odot \tau^m$. A 1-step nibbled future of $E_i \tau^k \overline{E_j}$ has one of three forms: (I) $E_i \tau^{k_1} \tau'$, where τ' is a proper sub edge-path of τ ; (II) $\tau' \tau^{k_2} \tau''$ where τ' and τ'' are proper sub edge-paths of τ ; or (III) $\tau' \tau^{k_3} \overline{E_j}$, where τ' is a proper sub edge-path of τ .

Case 2(I): In this case, $E_i \tau^{k_1} \tau'$ admits a hard splitting into E_i and $\tau^{k_1} \tau'$, which is of the required sort.

Case 2(II): In this case the path already had weight less than r .

Case 2(III): Suppose we are in case (III), and that μ , the αL^{V_r} -step nibbled future of $\tau' \tau^{k_3} \overline{E_j}$ has a copy of $\overline{E_j}$. Lemma 4.6 assures us that no nibbled future of τ' can contain τ as a subpath, and therefore there is a splitting of μ immediately on the right of $\overline{E_r}$, and we are done. If there is no copy of $\overline{E_j}$ in μ , we are also done, since this nibbled future must have weight less than r .

Case 4: Having dealt with cases (1) and (3), we may now suppose that at time $t_1 + V_r + \alpha L^{V_r} = t_1 + V'_r$ the edge $p(t_1 + V'_r)$ lies in an edge-path of weight at most $r - 1$ with a hard splitting of $\rho_{t_1+V'_r}$ immediately on either side.¹⁰ Denote this path by π'_r , chosen to be in the future of π . Note that π'_r has length at most $\alpha L^{V'_r}$.

¹⁰Note that again it is possible that $k < t_1 + V'_r$, in which case we are in case (iii) of our claim. We suppose therefore that this is not the case.

By Proposition 4.7 again, either $k < t_1 + V_{r-1}$ or at time $t_1 + V_{r-1}$ the nibbled future of π'_r admits a hard splitting into edge-paths each of which is either:

- (1) a single edge of weight $r - 1$;
- (2) a nibbled future of a weight $r - 1$ indivisible Nielsen path;
- (3) a nibbled future of a weight $r - 1$ GEP; or
- (4) a path of weight at most $r - 2$.

We continue in this manner. We may conceivably fall into case (4) each time until $t_1 + V_1$ when it is not possible to fall into a path of weight at most $1 - 1$! Thus at some stage we must fall into one of the first three cases. This completes the proof of Claim 1.

The existence of α . We must find a bound, in terms of n and f , on the length of indivisible Nielsen paths that arise in the nibbled future of ρ with a hard splitting immediately on either side.¹¹ To this end, suppose that ε' is an edge which lies in an indivisible Nielsen path τ in a k' -step nibbled future of ρ , and that there is a hard splitting immediately on either side of τ . We again denote the i -step nibbled future of ρ by ρ_i for $0 \leq i \leq k'$.

As above, we associate a diagram $\Delta_{k'}$ to $\rho_{k'}$.¹² Denote by q the path in the family forest of $\Delta_{k'}$ which follows the past of ε' . Let $q(i)$ be the edge in ρ_i which intersects q . Let the sequence of weights of the edges $q(i)$ be $W' = \{w'_0, \dots, w'_{k'}\}$.

Define incidents of Type A and B for W' in exactly the same way as for W , and let t_2 be the time of the last incident of Type A for W' . If there is no incident of Type A for W' let $t_2 = 0$. Let κ be the folded seed containing $q(t_2)$; in case $t_2 = 0$ let $\kappa = \rho$. Define $\theta = \max\{n, L\}$ and note that $|\kappa| \leq \theta$. The path τ must lie in the nibbled future of κ , so it suffices to consider the nibbled future of κ . Suppose that κ has weight r' .

We deal with the nibbled future of κ in the same way as we dealt with that of π . Let $\kappa_0 = \kappa, \kappa_1, \dots$ be the nibbled futures of κ .

Claim 2: There exists a constant $\beta' = \beta'(n, f)$ so that one of the following must occur:

- (i) for some $t_2 < i < k'$, the edge $q(i)$ lies in a GEP in $f_{\#}(\kappa_{i-1})$ that has a hard splitting immediately on either side;
- (ii) not in case (i), *and* at some time $i \leq k'$ the edge $q(i)$ lies in an indivisible Nielsen path τ_0 in $f_{\#}(\kappa_{i-1})$ so that $|\tau_0| \leq \theta L^{\beta'}$ and immediately on either side of τ_0 there is a hard splitting, *and* there are no incidents of Type B after time i ;
- (iii) $k' - t_2 \leq \beta'$; or

¹¹Recall that the definition of Type B incidents excluded the case of Nielsen paths which lie in the nibbled future of a GEP with a hard splitting immediately on either side.

¹²If we are considering Nielsen paths arising in the past of ε above, then we can assume $k' \leq k$ and that $\Delta_{k'}$ is a subdiagram of Δ_k in the obvious way.

(iv) there is a hard splitting of $\kappa_{k'}$ immediately on either side of ε' .

Let us prove that this claim implies the existence of α and hence completes the proof of the theorem. By definition, α is required to be an upper bound on the length of an arbitrary Nielsen path τ involved in a Type B incident. We assume this incident occurs at time k' and use Claim 2 to analyse what happens.

Case (i) of Claim 2 is irrelevant in this regard. If case (ii) occurs, the futures of τ_0 are unchanging up to time k' , so $\tau = \tau_0$ and we have our required bound. In case (iii) the length of τ is at most $\theta L^{\beta'}$, and in case (iv) τ is a single edge. It suffices to let $\alpha = \theta L^{\beta'}$.

It remains to prove Claim 2. The proof of Claim 2 follows that of Claim 1 almost verbatim, with θ in place of α and κ in place of ρ , etc., *except* that the third sentence in Case (2) of the proof becomes invalid because Type B incidents after time $t_2 + V_r$ may occur.

In this setting, suppose π_r (which occurs at time $t_2 + V_r$) is a Nielsen path, but that we are not in case (ii) of Claim 2, and there is a subsequent Type B incident at time j , say. The length of π_r is at most θL^{V_r} . The Nielsen path at time $j - 1$ has the same length as the one at time $t_2 + V_r$. We go forward to time j , where the future of π_r is no longer a Nielsen path, and continue the proof of Case (2) from the fourth sentence of the proof.

Otherwise, the proof of Claim 2 is the same as that of Claim 1 (the above modification is required at each weight, but at most *once* for each weight). The only way in which the length bounds change is in the replacement of θ by α (including in the definitions of V_i and V'_i). This finally completes the proof of Theorem 4.3. \square

5. PASSING TO AN ITERATE OF f

In this section we describe what happens to various definitions when we replace f by an iterate. Suppose that $k \geq 1$, and consider the relationship between f and $f_0 = f_{\#}^k$.

First, for any integer $j \geq 1$, the set of kj -monochromatic paths for f is the same as the set of j -monochromatic paths for f_0 . Therefore, once Theorem 3.2 is proved, we will pass to an iterate so that r -monochromatic becomes 1-monochromatic. However, the story is not quite as simple as that.

It is not hard to see that if $\sigma \odot \nu$ is a hard splitting for f , then it is also a hard splitting for f_0 .

When f is replaced by f_0 , the set of GEPs is unchanged, as are the sets of Ψ EPs and indivisible Nielsen paths. Also, the set of indivisible Nielsen paths which occur as sub-paths of $f(E)$ for some linear edge E remains unchanged.

With the definition as given, the set of (J, f_0) -atoms may be smaller than the set of (J, f) -atoms. This is because an atom is required to be 1-monochromatic.

However, we continue to consider the set of (J, f) -atoms even when we pass to f_0 , and we also consider paths to be *beaded* if they are (J, f) -beaded.

Since we are quantifying over a smaller set of paths the constant $V(n, f_0)$ in Theorem 4.3 is assumed, without loss of generality, to be $V(n, f)$. This is an important point, because the constant V is used to find the appropriate J when proving Theorem 3.2. When passing from f to f_0 , we need this J to remain unchanged, for the appropriate iterate k which we eventually choose depends crucially upon J . See Remark 7.1 below.

It is also clear that if $m \leq n$ then without loss of generality we may assume that $V(m, f) \leq V(n, f)$. Once again, this is because we are considering a smaller set of paths when defining $V(m, f)$.

We now want to replace f by a fixed iterate in order to control some of the cancellation within monochromatic paths. The following lemma is particularly useful in the proof of Proposition 6.9 below, and also for Theorem 8.1.

Lemma 5.1. *There exists $k_1 \geq 1$ so that $f_1 = f_{\#}^{k_1}$ satisfies the following. Suppose that E is an exponential edge of weight r and that σ is an indivisible Nielsen path of weight r (if it exists, σ is unique up to a change of orientation). Then*

- (1) $|f_1(E)| > |\sigma|$;
- (2) Moreover, if σ is an indivisible Nielsen path of exponential weight r and σ_0 is a proper subedge-path of σ , then $(f_1)_{\#}(\sigma_0)$ is r -legal;
- (3) If σ_0 is a proper initial sub edge-path of σ then $(f_1)_{\#}(\sigma_0)$ admits a hard splitting, $f(E) \odot \xi$, where E is the edge on the left end of σ ;
- (4) Finally, if σ_1 is a proper terminal sub edge-path of σ then $(f_1)_{\#}(\sigma_1) = \xi' \odot f(E')$ where E' is the edge on the right end of σ .

Now suppose that σ is an indivisible Nielsen path of parabolic weight r and that σ is a sub edge-path of $f(E_1)$ for some linear edge E_1 . The path σ is either of the form $E\eta^{m_\sigma}\overline{E'}$ or of the form $E\overline{\eta}^{m_\sigma}E'$, for some linear edges E and E' . Then

- (1) If σ_0 is a proper initial sub edge-path of σ then

$$(f_1)_{\#}(\sigma_0) = E \odot \eta \odot \cdots \odot \eta \odot \xi'',$$

where there are more than m_σ copies of η visible in this splitting.

- (2) If σ_1 is a proper terminal sub edge-path of σ then

$$(f_1)_{\#}(\sigma_1) = \xi' \odot \overline{\eta} \odot \cdots \odot \overline{\eta} \odot \overline{E'},$$

where there are more than m_σ copies of $\overline{\eta}$ visible in this splitting;

Proof. First suppose that H_r is an exponential stratum, that σ is an indivisible Nielsen path of weight r , and that E is an edge of weight r . Since $|f_{\#}^j(E)|$ grows exponentially with j , and $|f_{\#}^j(\sigma)|$ is constant, there is certainly some d_0 so that $|f_{\#}^d(E)| > |\sigma|$ for all $d \geq d_0$.

There is a single r -illegal turn in σ , and if σ_0 is a proper sub edge-path of σ . By Lemma 1.17, no future of σ_0 can contain σ as a subpath. The number of r -illegal turns in iterates of σ_0 must stabilise, so by Lemma 1.12 there is an iterate of σ_0 which is r -legal. Since there are only finitely many paths σ_0 , we can choose an iterate of f which works for all such σ_0 .

Suppose now that σ_0 is a proper initial sub edge-path of σ , and that E is the edge on the left end of σ . It is not hard to see that every (entire) future of σ_0 has E on its left end. We have found an iterate of f so that $f_{\#}^{d'}(\sigma_0)$ is r -legal. It now follows immediately that

$$f_{\#}^{d'+1}(\sigma_0) = f(E) \odot \xi,$$

for some path ξ . The case when σ_1 is a proper terminal sub edge-path of σ is identical.

Now suppose that H_r is a parabolic stratum and that σ is an indivisible Nielsen path of weight r of the form in the statement of the lemma. The claims about sub-paths of σ follow from the hard splittings $f(E) = E \odot u_E$ and $f(E') = E' \odot u_{E'}$, and from the fact that m_σ is bounded because σ is a subpath of some $f(E_1)$.

As in Remark 4.8, we can treat each of the cases separately, and finally take a maximum. \square

6. THE NIBBLED FUTURES OF GEPS

The entire future of a GEP is a GEP but a nibbled future need not be and Theorem 4.3 tells us that we need to analyse these nibbled futures. This analysis will lead us to define *proto- Ψ EPs*. In Proposition 6.9, we establish a normal form for proto- Ψ EPs which proves that proto- Ψ EPs are in fact the Ψ EPs which appear in the Beaded Decomposition Theorem.

To this end, suppose that

$$\zeta = E_i \bar{\tau}^n \bar{E}_j$$

is a GEP, where τ is a Nielsen path, $f(E_i) = E_i \odot \tau^{m_i}$ and $f(E_j) = E_j \odot \tau^{m_j}$. As in Definition 1.6, we consider $E_i \bar{\tau}^n \bar{E}_j$ to be unoriented, but here we do not suppose that $j \leq i$. However, we suppose $n > 0$ and thus, since $E_i \bar{\tau}^n \bar{E}_j$ is a GEP, $m_j > m_i > 0$.

The analysis of GEPS of the form $E_j \tau^n \bar{E}_i$ is entirely similar to that of GEPS of the form $E_i \bar{\tau}^n \bar{E}_j$ except that one must reverse all left-right orientations. Therefore, we ignore this case until Definition 6.2 below (and often afterwards also!).

We fix a sequence of nibbled futures $\zeta = \rho_{-l}, \dots, \rho_0, \rho_1, \dots, \rho_k, \dots$ of ζ , where ρ_0 is the first nibbled future which is not the entire future. Since the entire future of a GEP is a GEP, we restrict our attention to the nibbled futures of ρ_0 .

There are three cases to consider, depending on the type of sub-path on either end of ρ_0 .

- (1) $\rho_0 = \bar{\tau}_0 \bar{\tau}^m \bar{E}_j$;
- (2) $\rho_0 = \bar{\tau}_0 \bar{\tau}^m \bar{\tau}_1$.
- (3) $\rho_0 = E_i \bar{\tau}^m \bar{\tau}_1$;

where τ_0 is a (possibly empty) initial sub edge-path of τ , and τ_1 is a (possibly empty) terminal sub-edge path of τ .

In case (1) ρ_0 admits a hard splitting

$$\rho_0 = \bar{\tau}_0 \odot \bar{\tau} \odot \cdots \odot \bar{\tau} \odot \bar{E}_j.$$

Since τ_0 is a sub edge-path of $f(E_i)$, it has length less than L and its nibbled futures admit hard splittings as in Theorem 4.3 into nibbled futures of GEPs and paths of length at most $V(L, f)$. These GEPs will necessarily be of strictly lower weight than ρ_0 , since $\bar{\tau}_0$ is. Thus, case (1) is easily dealt with by an induction on weight, supposing that we have a nice splitting of the nibbled futures of lower weight GEPs; this is made precise in Proposition 6.10. Case (2) is entirely similar.

Case (3) is by far the most troublesome of the three, and it is this case which leads to the definition of *proto- Ψ EPs* in Definition 6.2 below. Henceforth assume $\rho_0 = E_i \bar{\tau}^m \bar{\tau}_1$.

Each of the nibbled futures of ρ_0 (up to the moment of death, Subsection 6.1) has a nibbled future of $\bar{\tau}_1$ on the right. If the latter becomes empty at some point, the nibbled future of ρ_0 at this time has the form $E_i \bar{\tau}^{n'} \bar{\tau}_2$, where τ_2 is a proper sub edge-path of τ . We restart our analysis at this moment. Hence we make the following

Working Assumption 6.1. *We make the following two assumptions on the nibbled futures considered:*

- (1) *All nibbling of ρ_i occurs on the right; and*
- (2) *the i -step nibbled future $\bar{\tau}_{1,i}$ of $\bar{\tau}_1$ inherited from ρ_i is non-empty.*

We will deal with the case $m - km_i < 0$ later, in particular with the value of k for which $m - (k - 1)m_i \geq 0$ but $m - km_i < 0$. For now suppose that $m - km_i \geq 0$.

In this case, the path ρ_k has the form

$$\rho_k = E_i \bar{\tau}^{m - km_i} \bar{\tau}_{1,k}.$$

There are (possibly empty) Nielsen edge paths ι and ν , and an indivisible Nielsen edge path σ so that

$$(6.1) \quad \tau = \iota \odot \sigma \odot \nu \text{ and } \tau_1 = \sigma_1 \odot \nu,$$

where σ_1 is a proper terminal sub edge-path of σ . Now, as in Working Assumption 6.1, there is no loss of generality in supposing that

$$\rho_k = E_i \bar{\tau}^{m - km_i} \bar{\nu} \bar{\sigma}_{1,k},$$

where $\bar{\sigma}_{1,k}$ is the nibbled future of $\bar{\sigma}_1$ inherited from ρ_k , and that $\bar{\sigma}_{1,k}$ is non-empty.

Since $|\sigma_1| < L$, by Theorem 4.3 the path $\sigma_{1,k}$ admits a hard splitting into edge-paths each of which is either the nibbled future of a GEP, or of length at most $V(L, f)$; we take the (unique) maximal hard splitting of $\sigma_{1,k}$ into edge paths.

Let $s = \lfloor m/m_i \rfloor + 1$. In ρ_s (but not before) there may be some interaction between the future of E_i and $\bar{\sigma}_{1,s}$. We denote by $\gamma_{\sigma_1}^{k,m}$ the concatenation of those factors in the hard splitting of $\bar{\sigma}_{1,k}$ which contain edges any part of whose future is eventually cancelled by some edge in the future of E_i under any choice of nibbled futures of ρ_k (not just the ρ_{k+t} chosen earlier) and any choice of tightening. Below we will analyse more carefully the structure of the paths $\bar{\sigma}_{1,k}$ and $\gamma_{\sigma_1}^{k,m}$.

We now have $\bar{\sigma}_{1,k} = \gamma_{\sigma_1}^{k,m} \odot \sigma_{1,k}^\bullet$. From (6.1), we also have

$$(6.2) \quad \rho_k = E_i \bar{\tau}^{m-km_i} \bar{\nu} \gamma_{\sigma_1}^{k,m} \odot \sigma_{1,k}^\bullet.$$

Definition 6.2 (Proto- Ψ EPs). Suppose that τ is a Nielsen edge path, E_i a linear edge such that $f(E_i) = E_i \odot \tau^{m_i}$ and τ_1 a proper terminal sub edge-path of τ such that $\tau_1 = \sigma_1 \odot \nu$ as in (6.1). Let $k, m \geq 0$ be such that $m - km_i \geq 0$ and let $\gamma_{\sigma_1}^{k,m}$ be as in (6.2). A path π is called a *proto- Ψ EP* if either π or $\bar{\pi}$ is of the form

$$E_i \bar{\tau}^{m-km_i} \bar{\nu} \gamma_{\sigma_1}^{k,m}.$$

Remarks 6.3.

- (1) The definition of proto- Ψ EPs is intended to capture those paths which remain when a GEP is partially cancelled, leaving a path which may shrink in size of its own accord.
- (2) By definition, a proto- Ψ EP admits no non-vacuous hard splitting into edge paths.

We now introduce two distinguished kinds of proto- Ψ EPs.

Definition 6.4. Suppose that

$$\pi = E_i \bar{\tau}^{m-km_i} \bar{\nu} \gamma_{\sigma_1}^{k,m},$$

is a proto- Ψ EP as in Definition 6.2.

The path π is a *transient* proto- Ψ EP if $k = 0$.

The path π is a *stable* proto- Ψ EP if $\gamma_{\sigma_1}^{k,m}$ is a single edge.

Lemma 6.5. *A transient proto- Ψ EP is a Ψ EP.*

Proof. With the notation of Definition 6.2, in this case $\gamma_{\sigma_1}^{0,m}$ is visibly a sub-path of $\bar{\tau}$, and the proto- Ψ EP is visibly a sub-path of a GEP. \square

Lemma 6.6. *A stable proto- Ψ EP is a Ψ EP.*

Proof. Since $\bar{\sigma}$ is a Nielsen path, if α is a nibbled future of $\bar{\sigma}$ where all the nibbling has occurred on the right, then the first edge in α is the same as the first edge in $\bar{\sigma}$.

However, $\gamma_{\sigma_1}^{k,m}$ is a nibbled future of $\bar{\sigma}$ where all the nibbling has occurred on the right. Therefore, if $\gamma_{\sigma_1}^{k,m}$ is a single edge then it must be a sub-path of $\bar{\sigma}$. It is now easy to see that any stable proto- Ψ EP must be a Ψ EP. \square

Remark 6.7. We will prove in Proposition 6.9 that after replacing f by a suitable iterate all proto- Ψ EPs are either transient or stable, and hence are Ψ EPs.

6.1. The Death of a proto- Ψ EP. Suppose that $\pi = E_i \bar{\tau}^{m-km_i} \bar{\nu} \gamma_{\sigma_1}^{k,m}$ is a proto- Ψ EP which satisfies Assumption 6.1. Let $q = \lfloor \frac{m-km_i}{m_i} \rfloor + 1$, and consider, π_{q-1} , a $(q-1)$ -step nibbled future of π . As before, we assume that the $(q-1)$ -step nibbled future of $\gamma_{\sigma_0}^{k,m}$ inherited from a π_{q-1} is not empty and that the edge labelled E_i on the very left is not nibbled.

In π_{q-1} , the edge E_j has consumed all of the copies of $\bar{\tau}$ and begins to interact with the future of $\bar{\nu} \gamma_{\sigma_1}^{k,m}$. Also, the future of π at time q need not contain a Ψ EP. Hence we refer to the time q as the *death of the Ψ EP*. Recall that $\tau = \iota \odot \sigma \odot \nu$ and that $\gamma_{\sigma_1}^{k,m}$ is a k -step nibbled future of $\bar{\sigma}_1$, where σ_1 is a proper subbath of σ . Let $p = m - (k + q - 1)m_i$, so that $0 \leq p < m_i$.

The path π_{q-1} has the form

$$\pi_{q-1} = E_i \bar{\tau}^p \bar{\nu} \gamma_{\sigma_1}^{k+q-1,m}.$$

Suppose that π_q is a 1-step nibbled future of π_{q-1} . In other words, π_q is a subpath of $f_{\#}(\pi_{q-1})$. Consider what happens when $f(\pi_{q-1})$ is tightened to form $f_{\#}(\pi_{q-1})$ (with any choice of tightening). The p copies of $\bar{\tau}$ (possibly in various stages of tightening) will be consumed by E_i , leaving $\bar{\nu} \odot f(\gamma_{\sigma_1}^{k+q-1,m})$ to interact with at least one remaining copy of $\tau = \iota \odot \sigma \odot \nu$. The paths ν and $\bar{\nu}$ will cancel with each other¹³.

Lemma 4.6 states that $\gamma_{\sigma_1}^{k,m}$ cannot contain σ as a subpath. Therefore, once ν and $\bar{\nu}$ have cancelled, not all of $\bar{\sigma}$ will cancel with $f(\gamma_{\sigma_1}^{k+q-1,m})$. A consequence of this discussion (and the fact that $f(E_i) = E_i \odot \tau^{m_i}$) is the following

Lemma 6.8. *Suppose that $\pi = E_i \bar{\tau}^{m-km_i} \bar{\nu} \gamma_{\sigma_0}^{k,m}$ is a proto- Ψ EP, and let $q = \lfloor \frac{m-km_i}{m_i} \rfloor + 1$. Suppose that π_{q-1} is a $(q-1)$ -step nibbled future of π satisfying Assumption 6.1. If π_q is an immediate nibbled future of π_{q-1} and π_q contains E_i then π_q admits a hard splitting*

$$\pi_q = E_i \odot \lambda.$$

¹³The hard splittings imply that this cancellation must occur under *any* choice of tightening.

We now analyse the interaction between $f(\gamma_{\sigma_1}^{k+q-1,m})$ and σ more closely. As usual, there are two cases to consider, depending on whether σ has exponential or parabolic weight¹⁴.

In the following proposition, f_1 is the iterate of f from Lemma 5.1 and we are using the definitions as explained in Section 5. Also, we assume that proto- Ψ EPs are defined using f_1 , not f .

Proposition 6.9. *Every proto- Ψ EP for f_1 is either transient or stable. In particular, every proto- Ψ EP for f_1 is a Ψ EP.*

Proof. Let $\pi = E_i \bar{\tau}^{m-km_i} \bar{\nu} \gamma_{\sigma_1}^{k,m}$ be a proto- Ψ EP for f_1 .

Lemma 6.5 implies that if $k = 0$ then π is a Ψ EP. Consider Working Assumption 6.1. If Assumption 6.1.(2) fails to hold at any point, then we can restart our analysis, and in particular we have a transient proto- Ψ EP at this moment. Thus we may suppose that π is an initial sub-path of a k -step nibbled future of a GEP, where $k \geq 1$ and we may further suppose that π satisfies Assumption 6.1.(2). We prove that in this case π is a stable proto- Ψ EP.

First suppose that σ has exponential weight, r say. If $\bar{\sigma}_0$ is a proper initial sub edge-path of $\bar{\sigma}$ then Lemma 5.1 asserts that

$$(f_1)_{\#}(\sigma_0) = f(E) \odot \xi,$$

and $|f(E)| > |\sigma|$. Note also that $f(E) = E \odot \xi''$ for some path ξ'' .

Now, at the death of the proto- Ψ EP, the nibbled future of $\gamma_{\sigma_0}^{k,m}$ interacts with a copy of E_i , and in particular with a copy of $f(\sigma)$ (in some stage of tightening). Now the above hard splitting, and the fact that σ is not r -legal whilst $f(E)$ is, shows that $\gamma_{\sigma_1}^{k,m}$ must be a single edge (namely E).

Suppose now that σ has parabolic weight. Since σ has proper sub edge-paths, it is not a single edge and so σ or $\bar{\sigma}$ has the form $E\eta^{m\sigma}\bar{E}'$. The hard splittings guaranteed by Lemma 5.1 now imply that $\gamma_{\sigma_1}^{k,m}$ is a single edge in this case also.

Therefore, every proto- Ψ EP for f_1 is transient or stable, proving the first assertion of the proposition. The second assertion follows from the first assertion, and Lemmas 6.5 and 6.6. \square

Finally, we can prove the main result of this section. In the following, L_1 is the maximum length of $f_1(E)$ over all edges E of G .

The following statement assumes the conventions of Section 5.

Proposition 6.10. *Under iteration of the map f_1 constructed in Lemma 5.1, any nibbled future of a GEP admits a hard splitting into edge paths, each of which is either a GEP, a Ψ EP, or of length at most $V(2L_1, f)$.*

¹⁴Recall that there are three kinds of indivisible Nielsen paths: constant edges, parabolic weight and exponential weight. If σ has nontrivial proper sub edge-paths, then it is certainly not a single edge, constant or not.

Proof. Suppose that $E_i\bar{\tau}^n\bar{E}_j$ is a GEP of weight r . We may suppose by induction that any nibbled future of any GEP of weight less than r admits a hard splitting of the required form (the base case $r = 1$ is vacuous, since there cannot be a GEP of weight 1).

Suppose that ρ is a nibbled future of $E_i\bar{\tau}^n\bar{E}_j$. If ρ is the entire future, it is a GEP and there is nothing to prove. Otherwise, as in the analysis at the beginning of this section, we consider the first time when a nibbled future is not the entire future. Let the nibbled future be ρ_0 . In cases (1) and (2) from that analysis, ρ_0 admits a hard splitting into edge-paths, each of which is either (i) \bar{E}_i ; (ii) $\bar{\tau}$; or (iii) a proper sub edge-path of τ . In each of these cases, Theorem 4.3 asserts that there is a hard splitting of ρ into edge-paths, each of which is either of length at most $V(L, f)$ or is the nibbled future of a GEP. Any nibbled future of a GEP which occurs in this splitting is necessarily of weight strictly less than r , and so admits a hard splitting of the required form by induction.

Suppose then that ρ_0 satisfies Case (3), the third of the cases articulated at the beginning of this section. In this case, ρ_0 is a transient proto- Ψ EP. Also, any time that Assumption 6.1.(2) is not satisfied, the nibbled future of ρ_0 is a transient proto- Ψ EP. Thus, we may assume that Assumption 6.1 is satisfied. If $m - km_i \geq 0$ then we have

$$\rho = E_i\bar{\tau}^{m-km_i}\bar{\nu}\gamma_{\sigma_1}^{k,m} \odot \sigma_{1,k}^\bullet.$$

The first path in this splitting is a stable Ψ EP by Proposition 6.9. Once again, Theorem 4.3 and the inductive hypothesis yield a hard splitting of $\sigma_{1,k}^\bullet$ of the required form.

Finally, suppose that Case (3) pertains and $m - km_i < 0$. Let $q = \lfloor \frac{m - km_i}{m_i} \rfloor + 1$ (the significance of this moment – “the death of the Ψ EP” – was explained at the beginning of this subsection). By the definition of a Ψ EP (Definition 6.2), the q -step nibbled future of ρ_0 admits a hard splitting as

$$E_i\bar{\tau}^{m-qm_i}\bar{\nu}\gamma_{\sigma_1}^{q,m} \odot \sigma_{1,q}^\bullet.$$

By Lemma 6.8, the immediate future of $E_i\bar{\tau}^{m-qm_i}\bar{\nu}\gamma_{\sigma_1}^{q,m}$ admits a hard splitting as $E_i \odot \xi$. Since $\gamma_{\sigma_1}^{q,m}$ is a single edge, we have a bound of $2L_1$ on the length of ξ . Any nibbled future of $E_i \odot \xi$ now admits a hard splitting into edge paths, each of which is either a GEP, a Ψ EP or of length at most $V(2L_1, f)$, by induction on weight and Theorem 4.3. \square

We highlight one consequence of Proposition 6.10:

Corollary 6.11. *Suppose that $\rho = E_i\bar{\tau}^{m-km_i}\bar{\nu}\gamma$ is a Ψ EP. Any immediate nibbled future of ρ (with all nibbling on the right) has one of the following two forms:*

- (1) $\rho' \odot \sigma$, where ρ' is a Ψ EP and σ admits a hard splitting into atoms; or
- (2) $E_i \odot \sigma$, where σ admits a hard splitting into atoms.

In particular, this is true of $f_{\#}(\rho)$.

There are entirely analogous statements in case ρ is a Ψ EP where $\bar{\rho}$ has the above form and all nibbling occurs on the left.

7. PROOF OF THE BEADED DECOMPOSITION THEOREM

In this section, we finally prove Theorem 3.2. As noted in Section 3, this immediately implies the Beaded Decomposition Theorem.

Proof (Theorem 3.2). Take $r = k_1$, the constant from Lemma 5.1, and $J = V(2L_1, f)$, where V is the constant from Theorem 4.3, and L_1 is the maximum length of $f_{\#}^{k_1}(E)$ for any edge $E \in G$.

Suppose that ρ is an r -monochromatic path. Then ρ is a 1-monochromatic path for $f_1 = f_{\#}^{k_1}$. By Proposition 6.9, every proto- Ψ EP for f_1 is a Ψ EP.

By Theorem 4.3, ρ admits a hard splitting into edge paths, each of which is either the nibbled future of a GEP or else has length at most $V(1, f)$. By Proposition 6.10, if we replace f by f_1 then any nibbled future of a GEP admits a hard splitting into edge paths, each of which is either a GEP, a Ψ EP or else has length at most $V(2L_1, f)$. By Lemma 2.6, the splitting of the nibbled future of a GEP is inherited by ρ .

We have shown that ρ is (J, f) -beaded, as required. \square

Remark 7.1. We have already remarked that, for a fixed m , the constant $V(m, f)$ from Theorem 4.3 remains unchanged when f is replaced by an iterate.

As in Section 5, we retain the notion of (J, f) -beaded with the original f when passing to an iterate of f .

Therefore, when f is replaced by an iterate, Theorem 3.2 remains true with the same constant J . This remark is important in our applications, for we pass to iterates of f , and the iterate chosen will depend on J .

8. REFINEMENTS OF THE MAIN THEOREM

The Beaded Decomposition Theorem is the main result of this paper. In this section, we provide a few further refinements that will be required for future applications.

Throughout this section we suppose that f has been replaced with f_1 from Lemma 5.1, whilst maintaining the conventions for definitions from Section 5. When we refer to f we mean this iterate f_1 . With this in mind, a *monochromatic path* is a 1-monochromatic path for f . Similarly, armed with Theorem 3.2, we refer to (J, f) -beads, simply as *beads*, and a path which is (J, f) -beaded will be referred to simply as *beaded*. The constant L now refers to the maximum length $|f(E)|$ for edges $E \in G$ with the new f .

In the following theorem, the *past* of an edge is defined with respect to an arbitrary choice of tightening.

Theorem 8.1. *There exists a constant D_1 , depending only on f , with the following properties. Suppose $i \geq D_1$, that χ is a monochromatic path and that ε is an edge in $f_{\#}^i(\chi)$ of weight r whose past in χ is also of weight r . Then ε is contained in an edge-path ρ so that $f_{\#}^i(\chi)$ has a hard splitting immediately on either side of ρ and ρ is one of the following:*

- (1) a Nielsen path;
- (2) a GEP;
- (3) a Ψ EP; or
- (4) a single edge.

Proof. Let χ be a monochromatic path. For any $k \geq 0$, denote $f_{\#}^k(\chi)$ by χ_k . In a sense, we prove the theorem ‘backwards’, by fixing an edge ε_0 of weight r in $\chi_0 = \chi$ and considering its futures in the paths χ_k , $k \geq 1$. The purpose of this proof is to find a constant D_1 so that if ε is any edge of weight r in χ_i with past ε_0 , and if $i \geq D_1$ then we can find a path ρ around ε satisfying one of the conditions of the statement of the theorem.

Fix $\varepsilon_0 \in \chi_0$. By Theorem 3.2, there is an edge path π containing ε_0 so that χ admits a hard splitting immediately on either side of π and π either (I) is a GEP; (II) has length at most J ; or (III) is a Ψ EP. In the light of Remark 4.8, it suffices to establish the existence of a suitable D_1 in each case. To consider the futures of ε_0 in the futures $f_{\#}^k(\chi)$ of χ , it suffices to consider the futures of ε_0 within the (entire) futures of π . Therefore, for $k \geq 0$, let $\pi_k = f_{\#}^k(\pi)$. Suppose that we have chosen, for each k , an edge ε_k in π_k such that: (i) ε_k lies in the future of ε_0 ; (ii) ε_k has the same weight as ε_0 ; and (iii) ε_k is in the future of ε_{k-1} for all $k \geq 1$.

Case (I): π is a GEP. In this case, the path π_k is a GEP for all k , any future of ε_0 lies in π_k , and there is a hard splitting of χ_k immediately on either side of π_k . Therefore, the conclusion of the theorem holds in this case with $D_1 = 1$.

Case (II): $|\pi| \leq J$. Denote the weight of π by s . Necessarily $s \geq r$. By Lemma 4.5 the path $\pi_{D(J)}$ admits a hard splitting into edge paths, each of which is either

- (1) a single edge of weight s ;
- (2) an indivisible Nielsen path of weight s ;
- (3) a GEP of weight s ; or
- (4) a path of weight at most $s - 1$.

We consider which of these types of edge paths our chosen edge $\varepsilon_{D(J)}$ lies in. In case (1) there is a hard splitting of $\pi_{D(J)}$ immediately on either side of the edge $\varepsilon_{D(J)}$, so for all $i \geq D(J)$ there is a hard splitting of π_i immediately on either side of ε_i , since ε_i and $\varepsilon_{D(J)}$ both have the same weight as ε_0 . For cases (2) and (3), $\varepsilon_{D(J)}$ lies in an indivisible Nielsen path or GEP with a hard splitting of $\pi_{D(J)}$ immediately on either side, so for all $i \geq D(J)$ any future of ε_0 in π_i ,

and in particular ε_i , lies in an indivisible Nielsen path of GEP immediately on either side of which there is a hard splitting of π_i .

Finally, suppose $\varepsilon_{D(J)}$ lies in an edge path $\tilde{\rho}$ with a hard splitting of $\pi_{D(J)}$ immediately on either side, and that $\tilde{\rho}$ is not a single edge, an indivisible Nielsen path, or a GEP¹⁵. We need only consider the future of $\tilde{\rho}$. For $k \geq 0$, let $\rho_{D(J)+k} = f_{\#}^k(\tilde{\rho})$ be the future of $\tilde{\rho}$ in $\pi_{D(J)+k}$. Now, $|\tilde{\rho}| \leq JL^{D(J)}$ so by Lemma 4.5 the edge path $\rho_{D(J)+D(JL^{D(J)})}$ admits a hard splitting into edges paths, each of which is either

- (1) a single edge of weight $s - 1$;
- (2) an indivisible Nielsen path of weight $s - 1$;
- (3) a GEP of weight $s - 1$; or
- (4) a path of weight at most $s - 2$.

We proceed in this manner. If we ever fall into one of the first three cases, we are done. Otherwise, after $s - r + 1$ iterations of this argument, the fourth case describes a path of weight strictly less than r . Since the weight of each ε_i is r , it cannot lie in such a path, and one of the first three cases must hold. Thus we have found the required bound D_1 in the case that $|\pi| \leq J$.

Case (III): π is a Ψ EP.

Let $\pi = E_i \bar{\tau}^{m-km_i} \bar{\nu} \gamma_{\sigma_1}^{k,m}$ as in Definition 6.2. We consider where in the path π the edge ε_0 lies. First of all, suppose that ε_0 is the unique copy of E_i . Since ε_0 is parabolic, it has a unique weight s future at each moment in time. Let $q = \lfloor \frac{m-km_i}{m_i} \rfloor + 1$, the moment of death. For $1 \leq p \leq q - 1$, the edge ε_p is the leftmost edge in a Ψ EP and there is a hard splitting of π_p immediately on either side of this Ψ EP. For $p \geq q$, Lemma 6.8 ensures that there is a hard splitting of π_p immediately on either side of ε_p . Therefore in this case the conclusion of the theorem holds with $D_1 = 1$.

Now suppose that the edge ε_0 lies in one of the copies of $\bar{\tau}$ in π , or in the visible copy of $\bar{\nu}$. Then any future of ε_0 lies in a copy of τ or ν respectively, which lies in a Ψ EP with a hard splitting immediately on either side, until this copy of $\bar{\tau}$ or $\bar{\nu}$ is consumed by E_i . Again, the conclusion of the theorem holds with $D_1 = 1$.

Finally, suppose that ε_0 lies in $\gamma_{\sigma_1}^{k,m}$. For ease of notation, for the remainder of the proof γ will denote $\gamma_{\sigma_1}^{k,m}$. By Proposition 6.9 γ is a single edge. Until the q -step nibbled future of π , any future of γ of the same weight is either γ or will have a splitting of π immediately on either side.

Since σ is an indivisible Nielsen path, and γ is a single edge, γ is the leftmost edge of $\bar{\sigma}$. Therefore $[\sigma\gamma]$ is a proper sub edge-path of σ .

Suppose that σ has exponential weight (this weight is r). By Lemma 5.1 and the above remark, $f_{\#}(\sigma\gamma)$ is r -legal. Therefore, any future of γ which

¹⁵In this case necessarily $s \leq r - 1$

has weight r will have, at time q and every time afterwards, a hard splitting immediately on either side.

Suppose now that σ has parabolic weight r . Since $[\sigma E]$ is a proper sub edge-path of σ , and since there is a single edge of weight r in $f(E)$ and this is cancelled, it is impossible for γ to have a future of weight r after time q . \square

Recall that the number of strata for the map $f : G \rightarrow G$ is ω . Recall also the definition of *displayed* from Definition 2.15

Lemma 8.2. *Let χ be a monochromatic path. Then the number of displayed Ψ EPs in χ of length more than J is less than 2ω .*

Proof. Suppose that χ is a monochromatic path, and that ρ is a subpath of χ , with a hard splitting immediately on either side, such that ρ is a Ψ EP, and $|\rho| > J$. Then, tracing through the past of χ , the past of ρ must have come into existence because of nibbling on one end of the past of χ . Suppose this nibbling was from the left. Then all edges to the left of ρ in χ have weight strictly less than that of ρ , since it must have come from a proper subpath of an indivisible Nielsen path in the nibbled future of the GEP which became ρ . Also, any Ψ EP to the left of ρ must have arisen due to nibbling from the left. Therefore, there are at most ω Ψ EPs of length more than J which came about due to nibbling from the left. The same is true for Ψ EPs which arose through nibbling from the right. \square

Lemma 8.3. *Let D_1 be the constant from Theorem 8.1, and let $f_2 = (f_1)_{\#}^{D_1}$. If ρ is an atom, then either $(f_2)_{\#}^{\omega}(\rho)$ is a beaded path all of whose beads are Nielsen paths and GEPs, or else there is some displayed edge $\varepsilon \subseteq (f_2)_{\#}^{\omega}(\rho)$ so that all edges in $(f_2)_{\#}^{\omega}(\rho)$ whose weight is greater than that of ε lie in Nielsen paths and GEPs.*

Proof. Suppose that ρ is an atom of weight r . If H_r is a zero stratum and $(f_2)_{\#}(\rho)$ has weight s then H_s is not a zero stratum. Thus, by going forwards one step in time if necessary, we suppose that H_r is not a zero stratum, so $(f_2)_{\#}(\rho)$ has weight r .

By Theorem 8.1, all edge of weight r in $(f_2)_{\#}(\rho)$ are either displayed or lie in Nielsen paths or GEPs (since we are considering the entire future of an atom, Ψ EPs do not arise here). If all edge of weight r in $(f_2)_{\#}(\rho)$ lie in Nielsen paths or GEPs then we consider the atoms in $(f_2)_{\#}(\rho)$ of weight less than r (this hard splitting exists since ρ and hence $(f_2)_{\#}(\rho)$ are monochromatic paths). We now consider the immediate future of these atoms in $(f_2)_{\#}^2(\rho)$, etc. It is now clear that the statement of the lemma is true. \square

Finally, we record an immediate consequence of the Beaded Decomposition Theorem and Proposition 6.10:

Theorem 8.4. *Suppose that σ is a beaded path. Any nibbled future of σ is also beaded.*

REFERENCES

- [1] M. Bestvina, M. Feighn and M. Handel, The Tits alternative for $\text{Out}(F_n)$ I: Dynamics of exponentially growing automorphisms, *Ann. of Math. (2)*, **151** (2000), 517–623.
- [2] M. Bestvina, M. Feighn and M. Handel, Solvable subgroups of $\text{Out}(F_n)$ are virtually abelian, *Geom. Ded.* **104** (2004), 71–96.
- [3] M. Bestvina, M. Feighn and M. Handel, The Tits alternative for $\text{Out}(F_n)$, II: A Kolchin type theorem, *Ann. of Math. (2)*, **161** (2005), 1–59.
- [4] M. Bestvina and M. Handel, Train tracks and automorphisms of free groups, *Ann. of Math. (2)*, **135**, 1–51.
- [5] M.R. Bridson and D. Groves, The quadratic isoperimetric inequality for mapping tori of free group automorphisms 1: Positive automorphisms, preprint at <http://arxiv.org/math.GR/0211459>.
- [6] M.R. Bridson and D. Groves, The quadratic isoperimetric inequality for mapping tori of free group automorphisms, II: The general case, in preparation.
- [7] M.R. Bridson and D. Groves, The growth of conjugacy classes under free-group automorphisms, in preparation.
- [8] E. Seneta, Non-negative matrices and Markov chains, Springer-Verlag, New York, 1981.

MARTIN R. BRIDSON, MATHEMATICS DEPARTMENT, 180 QUEEN'S GATE, LONDON, SW7 2BZ, U.K.

E-mail address: `m.bridson@ic.ac.uk`

DANIEL GROVES, DEPARTMENT OF MATHEMATICS, CALIFORNIA INSTITUTE OF TECHNOLOGY, PASADENA, CA 9110, USA

E-mail address: `groves@caltech.edu`